

DISSERTATION

SUBNETWORK ENSEMBLES

Submitted by

Timothy J. Whitaker

Department of Computer Science

In partial fulfillment of the requirements

For the Degree of Doctor of Philosophy

Colorado State University

Fort Collins, Colorado

Fall 2023

Doctoral Committee:

Advisor: Darrell Whitley

Charles Anderson

Nikhil Krishnaswamy

Michael Kirby

Copyright by Timothy J. Whitaker 2023

All Rights Reserved

# ABSTRACT

## SUBNETWORK ENSEMBLES

Neural network ensembles have been effectively used to improve generalization by combining the predictions of multiple independently trained models. However, the growing scale and complexity of deep neural networks have led to these methods becoming prohibitively expensive and time consuming to implement. Low-cost ensemble methods have become increasingly important as they can alleviate the need to train multiple models from scratch while retaining the generalization benefits that traditional ensemble learning methods afford. This dissertation introduces and formalizes a low-cost framework for constructing Subnetwork Ensembles, where a collection of child networks are formed by sampling, perturbing, and optimizing subnetworks from a trained parent model. We explore several distinct methodologies for generating child networks and we evaluate their efficacy through a variety of ablation studies and established benchmarks. Our findings reveal that this approach can greatly improve training efficiency, parametric utilization, and generalization performance while minimizing computational cost. Subnetwork Ensembles offer a compelling framework for exploring how we can build better systems by leveraging the unrealized potential of deep neural networks.

## ACKNOWLEDGEMENTS

I would like to start by thanking my advisor Dr. Darrell Whitley for his mentorship and guidance. It's been an honor and a privilege to work together for the last several years. Thank you for taking a chance on me and for pushing me to become a better research scientist. Thank you to the members of my dissertation committee: Dr. Chuck Anderson, Dr. Nikhil Krishnaswamy, and Dr. Michael Kirby. I am so grateful to have been able to learn from each of you. Your classes have been an integral part of this research and I appreciate the feedback, advice, and guidance that you've offered along the way. I am appreciative of financial support from The Computer Science Department at Colorado State University (Artificial Intelligence and Evolutionary Computation Fellowship) and The National Science Foundation (Grant IIS-1908866).

I am thankful for my family and friends who have stood by me throughout this journey. A special thanks to my mom and dad, Sandy and Brian, for their unwavering support and unconditional love. I am eternally grateful for the sacrifices you have made to give me the opportunity to be where I am today. You've set an example for what one can accomplish with hard work, determination and the courage to pursue your dreams. Shout out to my brother, Dylan, and my best friend, Kyler, for always having my back. Thank you for giving me perspective and focus when I needed it most. Thank you to my extended family and in-laws for their encouragement and continued interest in my work. It has meant a lot to know I have so many loved ones rooting for me.

Last but not least, my deepest gratitude is reserved for my wife, Alyssa. I feel so blessed to have you in my life. You were the inspiration for taking on this endeavor and you were the foundation upon which it was all built. You've been there to celebrate my triumphs and you've been there to help me overcome insurmountable challenges. You've been my constant source of strength and my guiding light when all seemed dark. Thank you for always believing in me. I could not have done this without you.

## DEDICATION

*For Alyssa*

*In this grand cosmic ensemble, it is my deepest joy to share this time and space with you.*

# TABLE OF CONTENTS

ABSTRACT . . . . .	ii
ACKNOWLEDGEMENTS . . . . .	iii
DEDICATION . . . . .	iv
LIST OF TABLES . . . . .	vii
LIST OF FIGURES . . . . .	ix
Chapter 1    Introduction . . . . .	1
1.1        Overview . . . . .	1
1.2        Challenges . . . . .	2
1.3        Contributions . . . . .	4
Chapter 2    Background . . . . .	7
2.1        Neuroscience . . . . .	7
2.2        Artificial Neural Networks . . . . .	8
2.2.1    Optimization . . . . .	10
2.2.2    Architecture . . . . .	12
2.2.3    Subnetworks . . . . .	14
2.3        Ensemble Learning . . . . .	16
2.3.1    Low-Cost Ensemble Methods . . . . .	18
Chapter 3    Theoretical Framework . . . . .	26
3.1        Introduction . . . . .	26
3.2        Formalization . . . . .	28
3.2.1    Subnetwork Sampling . . . . .	29
3.2.2    Neural Partitioning . . . . .	30
3.2.3    Perturbative Processes . . . . .	32
3.2.4    Network Optimization . . . . .	34
3.2.5    Ensemble Predictions . . . . .	34
3.2.6    Evaluation Metrics . . . . .	35
Chapter 4    Noisy Subnetwork Ensembles . . . . .	37
4.1        Introduction . . . . .	37
4.2        Implementation . . . . .	38
4.2.1    Trust-Region Search . . . . .	39
4.2.2    Mirrored and Partitioned Noise . . . . .	40
4.3        Experiments . . . . .	42
4.3.1    Decision Boundaries . . . . .	42
4.3.2    Ablations . . . . .	44
4.3.3    Benchmarks . . . . .	46
4.4        Discussion . . . . .	50
Chapter 5    Sparse Subnetwork Ensembles . . . . .	52

5.1	Introduction . . . . .	52
5.2	Implementation . . . . .	53
5.3	Experiments . . . . .	58
5.3.1	Ablations . . . . .	58
5.3.2	Landscapes . . . . .	63
5.3.3	Benchmarks . . . . .	64
5.4	Discussion . . . . .	68
Chapter 6	Stochastic Subnetwork Ensembles . . . . .	70
6.1	Introduction . . . . .	70
6.2	Implementation . . . . .	71
6.2.1	Random Annealing . . . . .	72
6.2.2	Temperature Annealing . . . . .	74
6.3	Experiments . . . . .	76
6.3.1	Ablations . . . . .	76
6.3.2	Benchmarks . . . . .	80
6.4	Discussion . . . . .	82
Chapter 7	Diversity Analysis . . . . .	83
7.1	Introduction . . . . .	83
7.2	Output Diversity . . . . .	84
7.2.1	Prediction Clusters . . . . .	85
7.3	Interpretable Diversity . . . . .	87
7.3.1	Perceptual Hashing . . . . .	88
7.4	Experiments . . . . .	89
7.5	Discussion . . . . .	92
Chapter 8	Conclusions . . . . .	94
8.1	Summary . . . . .	94
8.2	Discussion . . . . .	95
Bibliography	. . . . .	98

## LIST OF TABLES

2.1	A collection of popular activation functions used in neural networks. Logistic Sigmoids and Hyperbolic Tangents have long been the standard for shallow multilayer perceptrons thanks to their smooth differentiable qualities. Deep neural networks have primarily moved towards rectifiers as their unbounded nature helps preserve gradient flow over many layers. . . . .	9
2.2	A generalized overview of the benefits and drawbacks of various low-cost ensemble learning strategies. Specific methods within each category may differ significantly in regards to their strengths or weaknesses. All categories are relatively and subjectively graded according to the primary idea behind their cost reduction, where a dash represents baseline performance, a single checkmark corresponds to relatively good performance, and two checkmarks corresponds to best performance. 19	19
4.1	Results for sparse mutation ensembles on ImageNet with a wide variety of models. Mutation sparsity and strength are determined from a small hyperparameter grid search. Accuracy (Acc), Negative Log Likelihood (NLL), and expected calibration error (ECE) are reported for the parent and the ensemble. Metrics prepended with $e$ refer to ensemble results. $\Delta\text{Acc}$ is the change in accuracy between the parent and the ensemble, $\sigma$ is the mutation strength, and $\overline{KL}$ is the average output divergence between the mutated models and the parent models. 16 models are generated and the 8 most accurate candidates on a validation set are used together as an ensemble. . . . .	47
4.2	Proximal Policy Optimization hyperparameters and evaluation results over 10 episodes. . . . .	49
5.1	Accuracy of various methods for the small training budget benchmark experiment. All ensembles contain 8 models and are trained for a total of 16 epochs. Dropout rate is 50%, Bootstrap Aggregation is trained on 90% subsets. . . . .	65
5.2	Results for ensembles of WideResNet-28-10 models on both CIFAR-10 and CIFAR-100. Methods with * denote results obtained from [61, 106]. Best low-cost ensemble results are <b>bold</b> . $c\text{Acc}$ , $c\text{NLL}$ , and $c\text{ECE}$ correspond to corrupted test sets. . . . .	66
5.3	Results for Sparse Subnetwork Ensembles on CIFAR-10 with different child network partition sizes. The parent network is trained for 140 epochs. 6 children are then created through neural partitioning and each is trained for an additional 10 epochs. Total Params counts all parameters in the ensemble. . . . .	68
6.1	Comparison between various baseline approaches when used with random/magnitude pruning on CIFAR-100. We report the best accuracy for each method at various levels of target sparsity when tuned with both a constant learning rate and a one-cycle learning rate schedule. Our annealing strategies consistently outperform established one-shot and iterative pruning methods. . . . .	77

6.2	Results for ensembles of WideResNet-28-10 models on both CIFAR-10 and CIFAR-100. Methods with * denote results obtained from [61, 106]. Best low-cost ensemble results are <b>bold</b> . cAcc, cNLL, and cECE correspond to corrupted test sets. . . . .	81
7.1	Average prediction disagreement ratio and KL divergence between ensemble members for several benchmark methods on CIFAR-10 with WideResNet-28x10. Results reported from [61, 106, 161] . . . . .	84

## LIST OF FIGURES

2.1	A diagram of the McCulloch and Pitts computational model of a neuron used in the Perceptron. The neuron computes a weighted sum of the inputs and then passes the output through a binary step function for classification. . . . .	9
2.2	Visualizations of how convolutional layers operate on input images. Each convolutional filter in a layer scans over the input image according to a stride length. The middle image displays how a 7x7 filter with a stride of 7 applies to the input image. The right most image displays a 7x7 filter with overlap using a stride of 2.	13
3.1	A visualization of child networks formed through three distinct methodologies that we investigate in Chapters 4, 5, and 6. Noisy Subnetwork Ensembles perturb subnetworks with noise. Sparse Subnetwork Ensembles prune subnetworks from the parent. Stochastic Subnetwork Ensembles use probability scores to determine active parameters on every pass through the network. . . . .	27
4.1	Neuroevolution excels at fine tuning fully-trained networks that get stuck in flat loss basins. The left graph displays a typical SGD training trajectory where models tend to converge to the edges of flat optima [77]. The middle graph displays white dots as child networks that are generated from sparse mutations. The right graph shows the final ensemble consisting of top candidates selected from evaluation on a separate validation set. A key insight to the success of gradient-free methods for fine tuning is that the loss landscapes of test distributions rarely match the training distribution exactly. Ensembling over wide areas of good validation performance is key to improving generalization. . . . .	38
4.2	The images above detail decision boundaries for a trained three layer multilayer perceptron after being perturbed with various mutations. With dense perturbations, there is a comparatively small window between a functional decision boundary and complete performance collapse. As mutations become more sparse, the model is better able to retain its behavior while the strength of parameter mutation increases. . . . .	43
4.3	Results of mutation ablations for a trained WideResNet-28x10 model on CIFAR-10. Lines of different colors correspond to the percentage of sparsity for the mutations. The top left graph reports the average KL divergence as a result of mutation strength. The top right graph explores the relationship between KL divergence and accuracy. The bottom left reports the accuracy of an ensemble evaluated with majority vote predictions. The bottom right reports accuracy of an ensemble evaluated with averaged predictions. . . . .	45
4.4	Training graphs on all 16 environments in the ProcGen benchmark. Orange lines correspond to single model evaluations and blue lines correspond to noisy subnetwork ensembles. . . . .	49

5.1	A diagram explaining network partitions in Sparse Subnetwork Ensembles. The parent network is optimized and settles into some optima. Child networks are created via neural partitioning. That is, two sparse children each inherit a different set of parameters from the parent. The optimization landscapes for each child are diverse, yet they converge quickly due to inheritance from a fully trained parent. . . . .	53
5.2	A convolutional parent network filter split into two partitioned child filters. The child networks learn diverse feature representations due to the geometric constraints of their unique topologies. . . . .	55
5.3	An example of the one-cycle learning rate and momentum schedule we use to fine-tune the child networks in Sparse Subnetwork Ensembles. . . . .	58
5.4	Ablations over pruning sparsity and structure for various model sizes. Each graph reports performance for an ensemble containing 8 models. The left figure reports the effect that pruning sparsity has on models of different sizes. The right figure displays the efficacy of neuron and connection pruning for LeNet-L. . . . .	60
5.5	Ablations exploring the effect of ensemble size and various ensemble pruning/tuning strategies. The left graph displays the results for ever increasing ensemble sizes spawned from a single parent network. The right figure contains results for various configurations of random pruning and neural partitioning with constant rate fine tuning and one-cycle fine tuning. . . . .	61
5.6	Exploring the early training efficiency of Sparse Subnetwork Ensembles with Mothernets and Dropout. The left figure compares Mothernets on the MNIST dataset with a multilayer perceptron. The right figure compares Dropout on CIFAR-10 with LeNet-L. . . . .	62
5.7	Visualizations of the loss landscapes of a ResNet-56 model as it undergoes pruning and tuning. The first column is the landscape of the fully trained parent model. The middle figure is the landscape of the model after removing 50% of the parameters. The final figure is the landscape of the pruned model after three additional epochs of fine tuning. . . . .	64
6.1	Ablations exploring how the initial temperature and the number of annealing epochs affect convergence behavior. . . . .	77
6.2	Ablations over the number of annealing epochs for random annealing vs temperature annealing and temperature annealing vs iterative pruning. . . . .	78
6.3	The rightmost graph displays the best performing annealing configuration for each sparsity and plots it against the one-shot pruning baseline. Stochastic annealing outperformed the other methods for each sparsity, with more drastic improvements appearing in extremely sparse networks. . . . .	79
7.1	Pairwise correlation heatmaps for ResNet-18 on CIFAR-10. (R/NP) denotes random pruning/neural partitioning. (BAG) denotes bootstrap aggregation. . .	85
7.2	T-SNE cluster representations of CIFAR-10 and WideResNet-28x10. The first row is the representation of classes in CIFAR-10 that each of our models produce. The second row displays the results of our model predictions and the final row displays the incorrect predictions. . . . .	86

7.3	Feature visualizations of a random selection of neurons in child networks from a Snapshot Ensemble and a Sparse Subnetwork Ensemble. Values reported are the mean hash distances between the two child networks over all 1024 visualizations in the final convolutional layer. Higher values represent greater distance and more diversity between representations. The sparsity from Subnetwork Ensembles results in more diverse feature representations than two subsequent checkpoints in Snapshot Ensembles. . . . .	90
7.4	Mean distances between Snapshot Ensemble and Sparse Subnetwork Ensemble saliency maps for all 50,000 samples in the ILSVRC2012 ImageNet validation set. The visualization includes SmoothGrad saliency maps for a shared parent network, two snapshot ensemble checkpoints, and two sparse subnetwork children. The visualization is an example of how saliency maps can be useful with ambiguous input as both snapshot checkpoints focus on the dog while the subnetworks split their attention between the dog and the cat. The diversity encouraged by sparse network structures help to reduce bias present in the parent network. . .	91

# Chapter 1

## Introduction

### 1.1 Overview

Ensemble Learning has long been an active area of research in machine learning as it is one of the most reliable, powerful, and flexible techniques for improving the predictive power of machine learning systems. [57, 88]. Rather than training a single model to make predictions, ensemble methods instead train several models and use them in combination with one another in order to make predictions. Combining the outputs of several models allows for ensembles to improve generalization by leveraging the complementary strengths and weaknesses of its members [28]. The principles behind ensemble learning algorithms are general and have been used with many types of parametric models (e.g. decision trees, support vector machines, neural networks) [10, 11, 32]. Recent research has shown that ensembles of deep neural networks in particular are especially effective, often being implemented as a part of winning solutions in high profile machine learning competitions [39, 76, 117, 121, 127, 133, 174].

Improvements in network architecture and training methodologies have led to a large increase in both model size and predictive power. Networks with billions of parameters are becoming increasingly commonplace as overparameterization is shown to be effective for improved training dynamics. However, the resource intensive nature of these large models brings about important challenges related to computational requirements, memory constraints, and energy efficiency. As deep neural networks continue to grow in both scale and complexity, the cost associated with traditional ensemble learning techniques quickly becomes untenable. The development of low-cost ensemble methods have become increasingly important as they offer the opportunity to significantly reduce computational costs while retaining the generalization benefits that full size ensembles afford.

Meanwhile, pruning has emerged as a critical avenue of research for reducing the size of deep neural networks. Large numbers of parameters can be removed from previously trained models with little discernible loss in predictive power after a few epochs of continued training [6, 43]. Saliency metrics are often used to select the most unnecessary parameters to remove in order to produce the most accurate subnetwork [43, 98, 149]. However, there is no single magic subnetwork that works best for every situation [20, 166]. A sufficiently overparameterized network contains vast numbers of equally capable subnetworks and even random sampling has been shown to result in highly accurate subnetworks at moderate levels of sparsity [29, 107].

In light of these observations, we introduce a novel low-cost ensemble learning framework that involves the dynamic generation of a collection of child networks spawned through the random sampling of subnetworks in a trained parent model. This research introduces and formalizes this process of constructing *Subnetwork Ensembles* with the goal of exploring how we can improve the training efficiency, parametric utilization, and generalization of deep ensembles while minimizing computational cost.

## 1.2 Challenges

The benefits of improved generalization have been well established in ensemble learning literature as performance has been shown to increase with larger numbers of both accurate and diverse members [28]. Diversity is an important concept that pertains to the correlation of predictions made by individual models in an ensemble [13, 91, 154]. If each model in the ensemble makes the same predictions or learns the same feature representations, then there is little benefit to be gained by combining their outputs.

Traditional ensemble methods encouraged diversity by training large collections of fast and simple classifiers on different subsets of the feature space [10, 32, 66]. However, these dataset splitting techniques do not work well with deep neural networks due to their data hungry nature. Deep neural network ensembles are more commonly trained on the entire

dataset, instead relying on random weight initialization, heterogeneous network architectures and stochastic ordering of training samples to encourage diversity.

Diversity is an especially important consideration for low-cost ensemble learning methods as they typically alleviate the need to train multiple deep neural networks from scratch by sharing network structure, gradient information, or learned parameters among ensemble members. This sharing of information invariably increases correlation and reduces diversity, which can have a significant negative impact on generalization [13, 91, 154].

Low-Cost ensemble methods often struggle to balance this trade-off between member diversity, member accuracy and ensemble size. Prioritizing one aspect often comes at the expense of another when working with a fixed compute budget. Low-cost methods can also lack algorithmic flexibility as they typically require particular training procedures or specialized network architectures to construct ensemble members [50, 70, 100, 159]. This can make it challenging to adapt these ensembles to accommodate evolving problem spaces.

Subnetwork Ensembles alleviate many of the challenges that typically plague competing methods. Our approach builds on the idea that we can train a single parent network and spawn diverse child networks by sampling, perturbing, and optimizing unique subnetworks from the parent. The decoupled nature between the parent and child networks allow for highly flexible implementations. The parent network can use any standard training procedure or network architecture without modification. This offers an opportunity to leverage pre-trained models which is becoming more valuable as networks continue to scale in size and cost. This also enables the ability to generate new child networks dynamically, which can lead to larger and more robust ensembles than are typically seen with deep neural networks. Child networks can be generated and optimized in parallel which allows for wide scaling as independent members don't rely on the weights of other child networks in the ensemble. Using sparse subnetworks allows for the potential to greatly reduce computational complexity while maintaining accuracy. The investigation of Subnetwork Ensembles could prove to be an important avenue of research for not only improving the efficiency and generalization of

low-cost ensemble methods, but also for advancing the theoretical underpinnings of neural computation.

## 1.3 Contributions

This work introduces, generalizes and formalizes the process of constructing, training, and evaluating Subnetwork Ensembles. We explore several approaches for generating child networks through noise perturbations, network pruning, and stochastic masking. We conduct a thorough and in depth evaluation of these methods across a number of benchmark tasks and with a wide variety of deep neural network architectures.

Chapter 2 introduces relevant background information and related work grounded in the foundations of modern deep learning. We explore how this work relates to the natural structures observed in biological neural networks. We discuss the state of deep neural network architectures and optimization methods. We also discuss traditional ensemble learning techniques and conduct a literature review of recent low-cost deep neural network ensemble methods.

Chapter 3 introduces the theoretical framework and formalizes Subnetwork Ensembling at a high level. We provide a mathematical overview of this approach and discuss various hyperparameters that govern how Subnetwork Ensembles could be formulated through different sampling, perturbation, and optimization methodologies. We introduce Neural Partitioning as a novel technique for encouraging diversity by dividing the parent network into independent partitions such that sets of child networks are formed with no inherited parameter overlap.

Chapter 4 introduces Noisy Subnetwork Ensembles, where child networks are generated by applying noise perturbations to subnetworks. We connect the ideas behind this approach to popular neuroevolutionary methods that use dense mutations and model selection strategies to generate populations of child networks. We conduct several ablation studies exploring the impact of noisy subnetwork mutations on network behavior. We introduce a grid search

algorithm for finding good perturbation hyperparameters by using trust region constraints, which limits divergence in the output space as a result of perturbation [130]. We evaluate Noisy Subnetwork Ensembles with many different deep neural network architectures on the difficult ImageNet dataset and we explore how this technique can be used on difficult reinforcement learning tasks with the ProcGen benchmark [22, 27].

Chapter 5 introduces Sparse Subnetwork Ensembles, where we generate child networks by pruning randomly sampled subnetworks from the trained parent network. We explore the efficacy of Neural Partitioning and demonstrate how it can reduce variance and improve performance by creating sets of child networks with opposing topologies. We illustrate how sparse child networks can be tuned to encourage diversity by using a phasic learning rate schedule that uses large learning rates to move child networks further apart in weight space before converging to local optima with small learning rates [69, 138]. We conduct several other ablations exploring the impact of pruning structure, pruning severity, ensemble size, and tuning procedure. We benchmark Sparse Subnetwork Ensembles against several state of the art low-cost and traditional ensemble methods on both in-distribution and out-of-distribution datasets, where we find that our approach proves to be more accurate and more robust than traditional ensemble methods that require up to 4x more training compute, all while using fewer parameters.

Chapter 6 introduces Stochastic Subnetwork Ensembles, which extends work in the previous chapter with a novel approach to representing subnetworks with probabilistic masks. Rather than pruning all the parameters in one-shot, instead each parameter is assigned a probability score that determines whether the connection will be retained on any forward pass. These probabilities are then annealed over several epochs such that the target subnetwork slowly reveals itself over time. Diversity is encouraged by assigning different probability distributions to different child networks. The inclusion of extra dimensionality in the early phases of tuning can help to provide smooth optimization of the target subnetwork. This

proves to be highly effective technique for extremely sparse (98%) and moderately sparse (50%) subnetworks, outperforming established one-shot and iterative pruning strategies.

Chapter 7 includes an in depth analysis of diversity in Sparse Subnetwork Ensembles. We include a standard evaluation of pairwise output diversity metrics and compare those to other benchmark ensemble algorithms. We additionally introduce a new approach to analyzing diversity by leveraging interpretability and visualization methods. We discuss how feature visualization and saliency maps can be used to qualitatively compare the differences between knowledge representations among sibling child networks in an ensemble. This allows us to meaningfully explore how and why diversity develops and contributes to robust performance. We introduce several perceptual hashing algorithms used to measure the distances between feature visualizations, improving on the typically noisy results seen with common image difference metrics like the squared error or cosine similarity. We conduct a large scale evaluation of this interpretable approach with both Sparse Subnetwork Ensembles and a benchmark temporal ensemble method called Snapshot Ensembles. We find that Sparse Subnetwork Ensembles produce significantly more diversity in both the feature space and output space among child networks while maintaining better accuracy.

Chapter 8 wraps up with a summary of the ideas, experiments, and observations found through this research. We discuss of how the Subnetwork Ensemble framework fits into the current landscape of ensemble methods for deep neural networks. We include our thoughts on the broader impacts and potential limitations that this work may have in other domains, and we explore several promising avenues and ideas for furthering this research in the future.

# Chapter 2

## Background

### 2.1 Neuroscience

Biological neural networks are far larger and more complex than the artificial neural networks used in machine learning research. Even the largest artificial neural networks pale in comparison to the human brain, which is a vast and complex web of interconnected subsystems containing over 86 billion neurons and 100 trillion synaptic connections [31]. Individual neurons within this framework tend to be very noisy, which makes it difficult to understand the relation between neuronal activity and behavior. A new conceptual paradigm is reframing how we view neuronal activity by instead exploring how groups of structurally connected neurons form the functional building blocks of the brain [15]. The organizational structure of these interconnected subsystems appears to play an important role for both energy efficiency and the emergence of complex, resilient, and adaptable learning behavior.

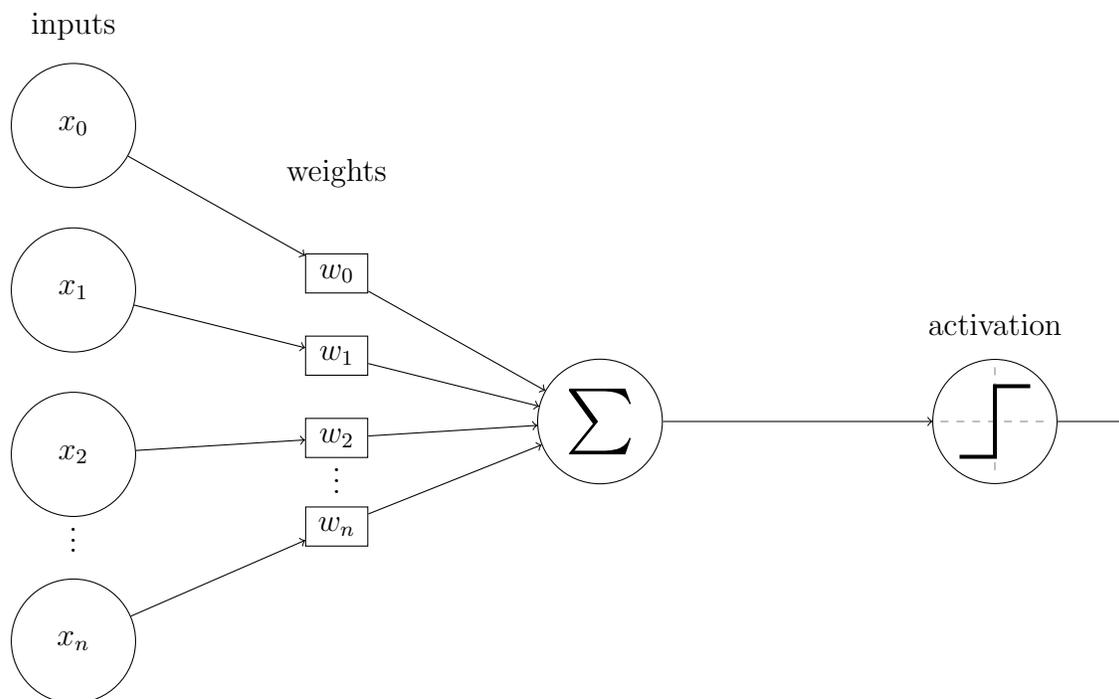
These structures emerge from an early process in brain development that involves a massive proliferation of neurons and synaptic connections followed by a widespread pruning of billions and trillions of those neurons and connections respectively. Neuronal proliferation creates an environment in which neurons and connections compete for resources. Neuronal apoptosis results in widespread neural cell death, resulting in a highly specialized structure consisting of many sparsely connected neuronal groups. These neuronal groups form the primary repertoire from which the rest of brain development occurs [35]. This overgrowth and subsequent pruning maximizes memory performance under metabolic energy constraints [18]. Additionally, this connective structure allows for information to be encoded in a sparse and distributed manner which optimizes a tradeoff between energy expenditure and expressivity [1, 53].

Several other important processes rely on the flexibility and resilience of these sparse subnetwork structures to encourage healthy brain function [139]. Neuroplasticity serves as a fundamental mechanism for learning and memory that relies on reorganizing synaptic connections in response to stimulus [25]. Long-term depression and long-term potentiation facilitate this process by strengthening or weakening the synaptic connections between neurons [75, 152]. Sparse subnetwork structures also provide a protective framework for mitigating the consequences of cellular damage or neurodegeneration. Synaptic Stripping is one such process where immunocompetent cells called microglia constantly scan the brain and selectively remove dysfunctional synapses from injured neurons [81, 136]. This is critical for preventing cascading failures that could result from synaptic dysfunction.

## 2.2 Artificial Neural Networks

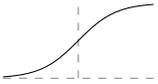
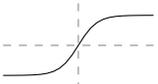
The origins of artificial neural network research can be traced back to the work of Warren McCulloch and Walter Pitts in 1943, who proposed a mathematical model of a biological neuron based on the all-or-nothing character of nervous activity [116]. This simplified computational model casts the neuron as a binary threshold unit that computes a weighted sum of its inputs to produce an output. Frank Rosenblatt later expanded on this work by developing the first single layer artificial neural network, consisting of McCulloch and Pitts neurons, called the Perceptron [125]. While the Perceptron was limited to modeling simple linearly separable functions, it was groundbreaking work that introduced adaptability to neural models through learning algorithms [114].

The transition to multilayer perceptrons greatly expanded the learning capabilities of neural networks [67]. Networks would now be organized into several successive layers, each consisting of multiple neurons with weighted connections to neurons in adjacent layers. The output of a neuron is computed by taking the weighted sum of all the inputs into the neuron and passing it through some non-linear activation function. Table 2.1 includes a small sample of some of the most popular activation functions used in neural networks over time.



**Figure 2.1:** A diagram of the McCulloch and Pitts computational model of a neuron used in the Perceptron. The neuron computes a weighted sum of the inputs and then passes the output through a binary step function for classification.

**Table 2.1:** A collection of popular activation functions used in neural networks. Logistic Sigmoids and Hyperbolic Tangents have long been the standard for shallow multilayer perceptrons thanks to their smooth differentiable qualities. Deep neural networks have primarily moved towards rectifiers as their unbounded nature helps preserve gradient flow over many layers.

Name	Function ( $\sigma(x)$ )	Derivative ( $\sigma'(x)$ )	Graph
Logistic	$\frac{1}{1 + e^{-x}}$	$\sigma(x)(1 - \sigma(x))$	
Hyperbolic Tangent	$\frac{e^x - e^{-x}}{e^x + e^{-x}}$	$1 - \sigma(x)^2$	
Binary Step	$\begin{cases} 0 & \text{if } x < 0 \\ 1 & \text{if } x \geq 0 \end{cases}$	0	
Rectified Linear Unit	$\begin{cases} 0 & \text{if } x < 0 \\ x & \text{if } x \geq 0 \end{cases}$	$\begin{cases} 0 & \text{if } x < 0 \\ 1 & \text{if } x \geq 0 \end{cases}$	

### 2.2.1 Optimization

Learning happens by adjusting the weights of the neural network to minimize error over a set of training samples. The error is described with a loss function that serves to measure the differences between model predictions and the true target values. The loss function can vary according to the task and training procedure. The mean squared error (MSE) and the mean absolute error (MAE) are both popular measures for regression tasks, where the goal is to predict a single scalar value.

$$MSE = \frac{1}{N} \sum_{i=1}^N (Y - \hat{Y})^2 \quad (2.1)$$

$$MAE = \frac{1}{N} \sum_{i=1}^N |Y - \hat{Y}| \quad (2.2)$$

Different loss functions may result in different optimization behavior according to their mathematical properties. For example, the mean squared error will tend to magnify the penalty for outliers compared to the mean absolute error. Regularization and penalty terms, such as weight decay, are often included in these loss functions to encourage optimization with respect to additional constraints [180].

The cross entropy loss (CE) is commonly used for classification tasks. Consider a dataset with samples that belong to one of  $C$  potential classes. The target label for each sample is encoded into a  $C$ -dimensional one-hot vector where each element in the vector has a value of 0 except for the correct class which has a value of 1. Neural networks are configured to output a  $C$ -dimensional vector where each dimension corresponds to the likelihood of that particular class being predicted. The outputs of the network are normalized and converted into a probability distribution by using the softmax function. The error is then determined by measuring the log likelihood between the predicted probability distribution  $\hat{Y}$  and the true distribution  $Y$ .

$$CE = - \sum_{i=1}^C Y_i \log(\hat{Y}_i) \quad (2.3)$$

Optimization is guided by algorithms that iteratively adjust the weights throughout training. Supervised deep learning tasks overwhelmingly use variations of gradient descent [126]. Data is fed through the network, the error is calculated according to the loss function, and updates are made by adjusting the weights according to a learning rate  $\eta$  with respect to the gradient of the loss  $\nabla F(x)$ .

$$w_{t+1} = w_t - \eta \nabla F(x) \tag{2.4}$$

The advent of the backpropagation algorithm provided an efficient method for calculating the gradient of neural networks by employing chain rule calculus layer by layer from the output to the input [160]. Consider the simple network in Figure 2.1 that contains a single neuron and a set of weights  $W$ . The output of the neuron  $O$ , feeds into an activation function with an output  $A$ . The derivative of the error  $E$  with respect to particular weight  $W_i$ , can be described by a chain of partial derivatives backpropagating through the network.

$$\frac{\partial E}{\partial W_i} = \frac{\partial E}{\partial A} \frac{\partial A}{\partial O} \frac{\partial O}{\partial W_i} \tag{2.5}$$

Classical gradient descent computes the gradient with respect to all samples in the training set for each weight update. While this approach leads directly to the optima of the loss function, it is computationally expensive for large datasets. Stochastic Gradient Descent instead updates the weights for every training sample individually. This approach is more efficient but adds variance as the gradient for each individual sample will not necessarily be in the direction of the true gradient of the training set. Minibatch Gradient Descent is most often used in deep learning, where the gradient is calculated over small subsets of the training data. This approach is more computationally efficient than full gradient descent while being less noisy than stochastic gradient descent [126].

Second order and derivative free methods have also been explored for optimizing neural networks. Second order approaches like Newton’s method and the conjugate gradient use

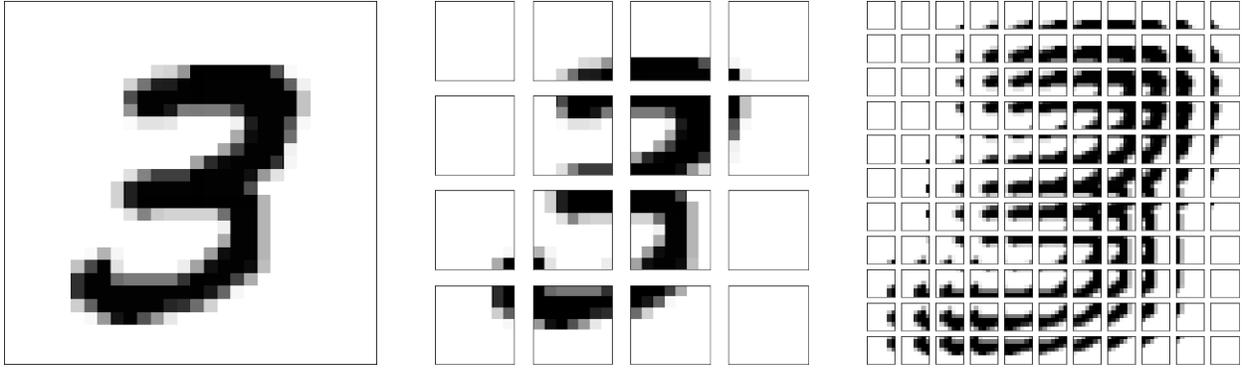
curvature information through second order derivatives to achieve faster convergence [4]. However, the computational cost of these methods with deep neural networks can make them impractical for large-scale problems [9]. Derivative free methods like evolutionary strategies and particle swarm optimization eschew gradient information and require only some objective fitness measure. These approaches excel on tasks where the gradient is non-differentiable, unavailable, or unreliable. However, these methods tend to be sample inefficient on supervised learning tasks compared to gradient methods [168].

Adaptive optimizers like Adagrad, Adadelata, and Adam use running statistics of previous gradients to scale the learning rates for each parameter [33,84,178]. This results in very fast convergence, however the solutions found by these methods tend to generalize worse than those trained with stochastic gradient descent [165]. While these alternatives are effective in certain contexts, stochastic gradient descent is the most popular optimization algorithm used in deep learning for its reliability, simplicity, and generality.

## 2.2.2 Architecture

Recent advancements in neural network architectures have led to breakthrough results across a wide variety of machine learning tasks. Many of these have been driven by the multilayer perceptron’s notoriously poor performance with high dimensional image data. These networks do not inherently account for spatial hierarchies as the input is flattened into a large one dimensional vector where each input neuron corresponds to a single feature. As a consequence, these networks are incredibly prone to overfitting noisy pixel level details and even small images lead to an explosion in the number of trainable weights.

Convolutional layers were introduced to mitigate these problems by drawing inspiration from the biological network structures found in animal visual cortexes [45]. These layers scan the input through a convolution operation, using a set of filters to produce feature maps that capture spatial relationships between neighboring pixels. Each filter uses the same shared window of weights for the entire input which drastically reduces the number of



**Figure 2.2:** Visualizations of how convolutional layers operate on input images. Each convolutional filter in a layer scans over the input image according to a stride length. The middle image displays how a 7x7 filter with a stride of 7 applies to the input image. The right most image displays a 7x7 filter with overlap using a stride of 2.

parameters required. These layers enabled much deeper networks that are able to capture complex hierarchical features without overfitting.

Convolutional networks have continued to establish state of the art results on many benchmarks, spurred by the introduction of MNIST and the five layer LeNet in 1989 [97]. Improvements have been progressive in the following years and limited mostly by the computational limitations of hardware at the time. Notable architectures include the eight layer AlexNet, which kickstarted the deep learning revolution by winning the ImageNet competition in 2012, and the sixteen layer VGG model that won the ImageNet localization challenge in 2014 [85, 134]. However, the use of deeper neural networks presented optimization challenges in the form of vanishing gradients [120]. This happens because backpropagation involves the repeated multiplication of partial derivatives through each layer from the output back to the input. If these derivative are small, then the repeated multiplication can cause the gradient to diminish quickly such that weight updates in the early layers approach zero. This is exacerbated by sigmoidal activation functions that saturate as their inputs become larger.

Residual Networks have since mitigated this problem by introducing skip connections between blocks of convolutional layers [62]. These connections bypass one or more layers by performing identity mapping where the input to a block is summed with the output of the

layers being bypassed. This reframes the underlying mapping  $H(x)$  for a group of layers to a residual mapping  $F(x) = H(x) - x$ , creating shortcuts that allow gradients to propagate more effectively throughout the network. Residual connections have enabled the effective training of substantially deeper networks while also exhibiting better convergence behavior during training, usually requiring fewer epochs to reach a stable and more accurate solution.

Vision Transformers have recently emerged as a powerful alternative to convolutional networks [30]. Unlike convolutional layers that use small kernels to capture local features, self-attention mechanisms in Vision Transformers allow for global interactions where pixels or image patches are weighted in relation to every other pixel or image patch. This approach offers the potential for better capturing spatial hierarchies, however it can be computationally expensive and require much more training data to avoid overfitting. Despite the promising results for Vision Transformers on massive datasets, deep residual convolutional networks are the de facto standard for computer vision tasks as they continue to achieve state-of-the-art results on many benchmarks.

### 2.2.3 Subnetworks

There has been a long history investigating the role of sparse subnetwork structures in deep learning research. Drawing inspiration from the high degree of sparsity observed in biological neural networks, these structures offer significant potential for improving parametric efficiency, inference speed, memory requirements, and generalization.

A large body of research has emerged to explore how network pruning can find accurate subnetworks within deep neural networks by selectively removing unimportant weights [6]. These methods offer significant potential for reducing computational complexity as a large number of parameters can be removed from a trained network while maintaining high performance after a small amount of continued training [6]. Many techniques have been introduced to select unnecessary weights by using importance measures like weight magnitude, gradient saliency, connection sensitivity, and synaptic flow [98,99,149]. While these methods can find

extremely sparse subnetworks that achieve high accuracy, even random pruning has been shown to result in accurate subnetworks up to moderate levels of sparsity [103, 107].

Sparse subnetwork structures are also implicitly leveraged in several important algorithmic advancements in deep learning. The rectified linear unit (ReLU), defined as  $f(x) = \max(0, x)$ , has gained widespread adoption as an activation function in deep neural networks. The activation patterns of ReLU neurons exhibit high levels of sparsity which are thought to improve the separability of high dimensional data [53]. Certain regularization methods leverage sparsity in parameter space in order to reduce overfitting and prevent feature reuse. Dropout, DropConnect, and Stochastic Depth Networks are the most popular examples where random neurons, connections, or layers are masked on each forward pass through the network [72, 140, 156]. Other effective regularization techniques include L1 and L2 weight decay penalties that aim to penalize large weights and push unimportant ones towards zero during training [87]. Weight decay can help to prevent overfitting, identify important features, and compress model sizes by learning sparse structures [111].

Other promising areas of neural network research aim to leverage the unique qualities of sparse networks to explore new models of neural computation. Spiking neural networks use biologically inspired action potentials over time where neurons only activate if a threshold is crossed [51, 150]. These models hold potential for significantly faster inference and reduced energy consumption for resource constrained environments. Echo State Networks and Liquid State Machines use recurrent reservoirs of sparsely connected neurons with fixed weights that have been proven to be particularly effective for dynamic time series modeling [78, 112]. Neuroevolutionary algorithms often evolve sparse network architecture along with the weights, excelling on tasks where gradient information is noisy or unavailable [40, 47, 141]. Super Networks are a theorized concept about how giant networks of the future will be too big to be trained by any one system, and instead portions of the network will need to be trained in a distributed manner by isolating sparse pathways [38].

## 2.3 Ensemble Learning

Many supervised learning algorithms can be described as search algorithms in a hypothesis space, where the goal is to train a set of model parameters to some candidate solution that best fits the target distribution [7]. These hypothesis spaces are very large and complex in deep learning which makes finding the global optimum very difficult. In practice, there are often many local optima that are good enough, however each will invariably have its own set of inherent biases [21]. Ensemble learning tries to remedy this problem by training several models that each converge to different local optima that are then combined together to form a better overall hypothesis.

This approach to improving generalization error is often explained in ensemble literature with an example of the bias-variance decomposition of the mean squared error (MSE) [13,122,154]. Given a model prediction  $\hat{y}$  and a true target value from an unknown test distribution  $y$ , the mean squared error is defined to be the expectation of the squared distance between the models predictions and the true target distribution. Bias is the difference between the expectation of the model and the true targets and variance is the squared difference between the models predictions and its mean.

$$bias = E[\hat{y}] - y \tag{2.6}$$

$$var = E[(\hat{y} - E[\hat{y}])^2] \tag{2.7}$$

$$MSE = E[(\hat{y} - y)^2] = (E[\hat{y}] - y)^2 + E[(\hat{y} - E[\hat{y}])^2] \tag{2.8}$$

$$MSE = bias^2 + var \tag{2.9}$$

Given an ensemble of  $M$  equally weighted estimators, the decomposition can be further extended to produce the bias-variance-covariance decomposition [13].

$$\overline{bias} = \frac{1}{M} \sum_i (E[\hat{y}_i] - y) \quad (2.10)$$

$$\overline{var} = \frac{1}{M} \sum_i E[(\hat{y}_i - E[\hat{y}_i])^2] \quad (2.11)$$

$$\overline{covar} = \frac{1}{M(M-1)} \sum_i \sum_{j \neq i} E[(\hat{y}_i - E[\hat{y}_i])(\hat{y}_j - E[\hat{y}_j])] \quad (2.12)$$

$$MSE = \overline{bias}^2 + \frac{1}{M} \overline{var} + (1 - \frac{1}{M}) \overline{covar} \quad (2.13)$$

Generalization error for single models relies upon the optimization of both bias and variance, where the tuning of a model towards high bias can cause it to miss important features and the tuning of a model toward high variance can cause it to be highly sensitive to noise. When the decomposition is extended to an ensemble, the generalization performance additionally depends on the covariance between models. Ideally, ensemble methods that promote diversity will be able to reduce covariance without increasing the bias or variance of individual models [13, 154].

Diversity is primarily encouraged through variation of the training data, network architecture, or training procedure for each ensemble member. Several of the most popular traditional ensemble methods predate the deep learning revolution and instead use simple and fast classifiers (e.g. decision trees, support vector machines, logistic regression, etc.) to produce large numbers of models optimized on different variations of the training data [129]. Bootstrap Aggregation does this by training each model on a random subset of the training data with replacement [10]. The Random Subspace Method instead trains each model on a random subset of features for each training sample [66]. Boosting strategically samples subsets of data to train models in sequence, where each subsequent model is trained to correct the mistakes of its predecessors [19, 32, 44, 80].

Traditional ensemble methods are well founded mathematically by leveraging large numbers of diverse models to overcome weak base classifiers. However, this approach is ill-suited to the computationally intensive nature of deep neural networks, which tend to be data hungry and require large numbers of training samples. These methods instead focus on training fewer models for longer periods, aiming to achieve high individual accuracy. Data variation techniques that reduce the size of the training set for each member tend to result in lower individual accuracy and worse ensemble performance overall [94]. Deep ensembles instead rely mostly on variations to network architecture (e.g. weight initializations, layer configurations, etc.) and training procedure (e.g. optimizer hyperparameters, loss functions, stochastic order of training samples, etc.) to induce diversity for each member [41,57,94,108].

### 2.3.1 Low-Cost Ensemble Methods

Efficient ensemble algorithms have become increasingly important as they hold the potential to significantly reduce the computational cost associated with traditional ensemble learning, while retaining many generalization benefits that ensemble learning affords. Several methods have been introduced that we broadly categorize as pseudo-ensembles, temporal ensembles, or evolutionary ensembles. Pseudo-ensembles train several models under the guise of a single monolithic architecture. Temporal ensembles encompass methods that save snapshots or checkpoints of a single model throughout training. Evolutionary ensembles leverage biologically inspired concepts like mutation, recombination, and selection to spawn new child networks. While these methods vary in their approach to reducing cost, the primary idea behind each is to share either network structure, gradient information, or learned parameters among ensemble members. Table 2.2 provides a broad overview on the relative benefits and drawbacks of each type of method across a number of categories: training cost, inference cost, storage cost, member diversity, and ensemble size.

**Table 2.2:** A generalized overview of the benefits and drawbacks of various low-cost ensemble learning strategies. Specific methods within each category may differ significantly in regards to their strengths or weaknesses. All categories are relatively and subjectively graded according to the primary idea behind their cost reduction, where a dash represents baseline performance, a single checkmark corresponds to relatively good performance, and two checkmarks corresponds to best performance.

Method	Training	Inference	Storage	Diversity	Ensemble Size
Pseudo	✓	✓✓	✓✓	-	-
Temporal	✓✓	-	-	✓	✓
Evolutionary	✓	-	-	✓✓	✓✓

### Pseudo-Ensembles

Pseudo-ensembles encompass a variety of methods that implicitly train ensembles as if they were a part of a single architecture [2]. These methods typically exploit variations of computational paths within the model to simulate the existence of multiple submodels. Pseudo-Ensembles also tend to share parameters among members, distinguishing these methods from other ensembles where independent members are created and trained separately.

Dropout is the canonical example of a pseudo-ensemble where neurons are randomly masked during training on each forward pass [140]. This technique effectively trains a unique subnetwork, or ensemble member, for each batch of training data. Dropout is deactivated during inference so that the entire network is utilized, which results in a computation where each of the subnetworks are implicitly ensembled together. Several variants of Dropout have since been introduced that mask different network structures. DropConnect masks individual connections rather than neurons [156]. Spatial Dropout masks feature maps in convolutional layers rather than neurons [153]. Stochastic Depth Networks mask entire layers of deep residual networks [73]. Since these techniques mask large parts of the network during training, significantly more forward passes are needed to fully train a network. Given enough training time, the resulting networks tend to converge to more robust optima than standard deep neural networks.

Some pseudo-ensemble methods explicitly encode relationships into the network architecture design. TreeNets are built on this idea where the network architecture splits into several independent branches that each have their own output head [100]. The early layers in the network serve as generalized feature extractors, while the independent branches learn specialized representations. During training, the shared layers accumulate the gradients from each independent branch and during testing, each path from input to one of the outputs can be treated as an independent ensemble member. Multiple-Input Multiple-Output architectures extend the ideas introduced with TreeNets to also incorporate independent input branches. Each input head is fed different samples and each of the output heads is trained to predict the corresponding input [61]. During inference, the same input is fed to all of the branches and the resulting outputs are ensembled together.

BatchEnsembles utilize parameter sharing by decomposing weight matrices into an element wise product between a single shared set of base weights and member-specific rank-1 matrices [159]. During training, the shared weights capture features that are common across all ensemble members while the member-specific weights can specialize. The rank-1 matrices for each ensemble member is represented with two single column vectors, which significantly reduces the storage cost compared to storing full rank weight matrices. Additionally, the element wise product is much cheaper computationally than matrix multiplication which results in little computational overhead compared to inference with a single network.

Knowledge Distillation was introduced to alleviate the inference and storage costs of ensembles and large networks. This approach is based on the idea that you can transfer the knowledge learned from a large network to a small network by training the small network using the predictions of a trained large network, instead of the raw training labels. This enables the small network to learn features that would be hard to disentangle if trained from scratch. This can be done in the context of ensemble learning by training a single network on the predictions made by a fully trained ensemble. The single network then approximates the behavior of the ensemble without needing to store and test all of the models separately [65].

Stochastic Weight Averaging is a recent optimization technique where the weights of a model are averaged over several states taken from the final training epochs [77]. This approach effectively integrates several local optima into a composite model with implicit pseudo-ensemble like inference. Model Soups operate on a similar principle and have recently shown state-of-the-art performance on large image classification benchmarks. Instead of taking the final epoch checkpoints of a single model, Model Soups average the weights of several independent models that are derived from the same pre-trained network but are each fine tuned with different hyperparameters [167].

Pseudo-ensembles are much more memory efficient than other low-cost ensemble methods as members are embedded within a single network structure. These techniques are generally very fast at performing inference, as forward passes do not need to be independently computed for each ensemble member. This makes these techniques highly valuable for environments where memory and inference cost is prioritized, like in embedded systems or edge computing. However, pseudo-ensembles do tend to need more training epochs to converge and the resulting models are less diverse than other ensemble methods, as large parts of the network structure are shared.

## **Temporal Ensembles**

Temporal ensembles are built on the idea that ensemble members can be created by selecting and saving checkpoints over a single network’s training trajectory [145]. These approaches capitalize on the variability of model weights as they converge to several distinct optima throughout training. Temporal methods primarily differ in how they determine which checkpoints to save and how to encourage the model to converge to diverse and accurate optima.

Fast Committee Learning pioneered this concept by selecting time slices at even intervals throughout training [145]. Horizontal Voting Ensembles were later introduced that saved a contiguous number of states from every epoch taken late in the training phase [170].

This sacrifices diversity for member accuracy as contiguous checkpoints tend to be highly correlated, while checkpoints taken early in training will be much less accurate.

Snapshot Ensembles introduced a highly effective technique for manipulating the learning rate schedule by alternating between large and small values over predetermined cycles [70]. The large rates encourage the model to move further in parameter space before converging to local optima with small learning rates. The model is saved at the end of each cycle before repeating, striking a nice balance between member diversity and accuracy. Fast Geometric Ensembles extend the algorithm introduced in Snapshot Ensembles by analyzing the loss surfaces found between the local optima saved at the end of each Snapshot checkpoint [50]. The landscape analysis revealed that these local optima tend to be connected by high accuracy pathways. Fast Geometric Ensembles leverage these high accuracy pathways by saving model states along them. Interpolating the weights between these optima allows for much larger ensemble sizes as they are not constrained by the cycle length used in Snapshot Ensembles.

FreeTickets is a recent method that combines ideas from sparse neural network training and temporal ensemble learning [106]. This approach uses a dynamic sparse training algorithm that isolates a unique sparse subnetwork at each cycle which is then optimized with a cyclic learning rate schedule. At the end of every optimization cycle, the sparse model state is saved before updating connectivity for the next cycle. The unique topological structure of each member encourages greater diversity between members, however training sparse subnetworks from scratch tends to result in slower convergence.

Temporal Ensembles tend to be very computationally efficient for training as these algorithms save model states over a typical network training trajectory. The performance of these temporal ensembles often compete with traditional dense ensembles at a significant fraction of the compute cost. Inference and storage is more expensive than pseudo-ensembles since each ensemble member needs to be evaluated independently. Temporal ensembles also tend to have smaller ensemble sizes than evolutionary ensembles as member generation is de-

pendent on training budget and the predetermined approach for selecting checkpoints. The members of certain temporal ensembles can struggle with diversity or accuracy as model checkpoints taken later in the training process tend to be highly correlated while earlier checkpoints tend to be less accurate.

## Evolutionary Ensembles

Evolutionary methods utilize concepts inspired by biological evolution like natural selection, mutation, and recombination to optimize populations over time [162]. These approaches have traditionally been concerned with single hypothesis testing, where optimization is geared towards selecting a single network with the best generalization performance from a population. However, several researchers have noted the natural connection between evolutionary populations and ensemble learning [46, 109, 172].

Evolution Strategies represent a popular subclass of evolutionary methods that describe a family of algorithms that generate child networks by selecting elite candidate models according to their fitness rankings and mutating them with noise perturbations [5]. These approaches are known to excel with wide scaling as candidate members can be cheaply generated and quickly evaluated in parallel [128].

Several variations of evolution strategies have been introduced that alter the ways that child networks are created. Natural Evolution Strategies generates new populations according to a weighted average over the previous population distribution [164]. This approach approximates a natural gradient as models are weighted according to their fitness such that mutations encourage models to move in the direction of previously successful models. Covariance Matrix Adaptation Evolution Strategies improves on evolution strategies by tracking dependencies between member parameters with a pairwise covariance matrix. The standard deviation can then control for the overall scale of the distribution, greatly improving exploration [58, 59]. However, despite the exploratory power of this approach, constructing the pairwise covariance matrix is computationally expensive and impractical for large networks.

Evolutionary algorithms are also commonly employed to evolve the topology of the network along with the weights. Neuroevolution of Augmenting Topologies is perhaps the most well known approach where networks are incrementally evolved from simple minimalist structures to more complex ones [141]. This approach also incorporates concepts like speciation to maintain a diverse population of network architectures with the goal of preventing premature convergence to suboptimal solutions. Weight Agnostic Neural Networks employ a similar approach, however the weights between neurons remain fixed. [48]. The focus then becomes entirely on evolving the architecture, where this approach interestingly found that network topology alone can encode solutions to complex reinforcement learning tasks.

Neuroevolutionary methods have primarily seen success with tasks that can be solved by relatively small networks, as these methods tend to be sample inefficient when scaling to larger models [49]. As a result, hybrid gradient/evolutionary methods have grown in popularity for deep neural networks as gradient optimization can be used to rapidly train the model while evolutionary processes can be used to improve exploration [26]. Evolutionary Ensembles with Negative Correlation Learning is a recent example of this hybrid methodology where gradient based optimization is applied to each member at each evolutionary generation. Multiple independent networks are trained simultaneously with the inclusion of a correlation penalty error term to encourage diversity. At each generation, the fittest individuals are selected as parents for the next generation, and child networks are spawned with normally distributed weight perturbations [109]. MotherNets is another method where a small base network is trained and used as a source to hatch larger child networks, where additional neurons and layers are added around the core network. The weights of the child networks are adjusted using function preserving network morphisms, in order to ensure that the added layers and neurons do not catastrophically affect the functionality of the core network. Diversity is encouraged through heterogenous network architectures where each child is optimized with different topological configurations [158].

These evolutionary approaches to optimization have several advantages over gradient based methods that are typically used in deep learning. They are inherently scalable as population members can be evaluated independently and they excel at exploring landscapes where gradient information is noisy, flat, or unavailable [128]. These low-cost evolutionary methods offer a unique advantage over other approaches in that it is generally cheap to generate new ensemble members. Once a parent network is trained, child networks tend to converge quickly while maintaining diversity in the parameter space. The decoupled nature of child generation and parent training enables more flexibility as large ensembles can be generated dynamically and pre-trained models could be used to save significant amounts of compute.

## Discussion

Ensemble learning has long been known to be an effective and powerful way to improve performance and these low-cost methods have introduced ways to share information between members that can significantly reduce computational cost. While there is no established taxonomy for low-cost methods in the literature, we note several commonalities where they can be organized under the categories of pseudo, temporal, and evolutionary ensembles. With fixed training budgets and competing research objectives, these approaches often struggle to balance several tradeoffs relating to member accuracy, member diversity, training cost, inference cost, and storage cost. These methods can suffer from restrictive algorithms that require special network structures or training procedures, which can make it difficult to utilize off the shelf models. It can also be costly to generate and train new ensemble members which results in smaller ensemble sizes that can limit generalization potential. These are increasingly important properties for ensemble methods as deep neural networks continue to grow in size. To address these multifaceted challenges, we aim to introduce a framework that allows for flexibility in its implementation while also optimizing these tradeoffs to produce accurate, robust, and efficient ensembles.

# Chapter 3

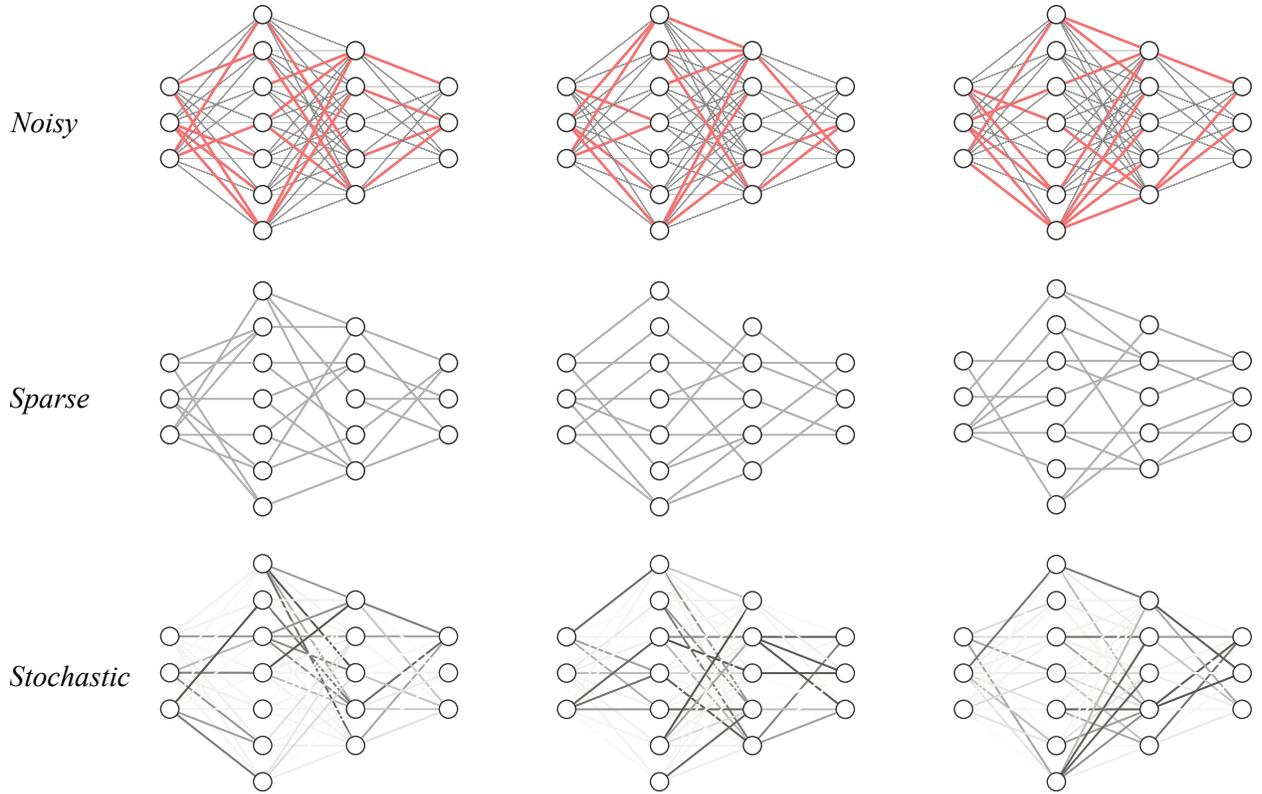
## Theoretical Framework

### 3.1 Introduction

We introduce Subnetwork Ensembles as an efficient low-cost approach to creating ensembles of deep neural networks. The goal with this research is to develop a generalized framework for describing a process where a large parent network is trained to convergence and then used to generate child networks through a perturbative process applied to randomly sampled subnetworks in the parent model. The spawned child networks can then be further optimized or tuned in order to recover the lost accuracy as a result of the perturbations.

The perturbative process, subnetwork sampling method, and tuning procedure used to create child networks can take many forms. We explore several approaches in Chapters 4, 5, and 6 called Noisy Subnetwork Ensembles, Sparse Subnetwork Ensembles, and Stochastic Subnetwork Ensembles respectively. Noisy Subnetwork Ensembles apply noise mutations to modify the weights of sampled subnetworks. Sparse Subnetwork Ensembles prune the weights of sampled subnetworks to create sparse child networks. Stochastic Subnetwork Ensembles uses probability distributions and Bernoulli realizations to create children with stochastic properties. Figure 3.1 provides a visualization of child networks formed through these three methodologies.

Subnetwork Ensembles are motivated by several principles observed in deep learning optimization research. Using a large parent network enables improved training dynamics and better convergence efficiency. Smaller networks trained from scratch often fail to achieve the same performance [64, 105, 148]. Large parent networks are extremely overparameterized and there's a notion that trained models contain large amounts of redundancy [6, 107]. Creating children from the subspaces of parent networks leverages unrealized capacity to improve parametric utilization. Spawned child networks inherit optimized parameters that



**Figure 3.1:** A visualization of child networks formed through three distinct methodologies that we investigate in Chapters 4, 5, and 6. Noisy Subnetwork Ensembles perturb subnetworks with noise. Sparse Subnetwork Ensembles prune subnetworks from the parent. Stochastic Subnetwork Ensembles use probability scores to determine active parameters on every pass through the network.

enable convergence to new local optima quickly. This property has long been observed in both pruning and transfer learning literature [6, 175, 183]. Perturbing different subspaces of the parent network encourages diverse representations among child networks. This is especially true for sparse subnetworks where altered network topologies impact the underlying optimization landscape [142].

This approach also offers several unique advantages over competing low-cost ensemble methods. Since the parent and child network procedures are completely decoupled, any parent network architecture or optimization algorithm can be used. Modern deep learning best practices (dropout, normalization layers, weight decay, adaptive optimizers, data augmentations, learning rate schedules, etc.) can be used to maximize performance of the parent [89]. This also opens the door to leveraging open source pre-trained models, which

is an increasingly valuable area of research as models and datasets continue to grow in size. It is also relatively fast and cheap to generate and train new child networks. This allows for the ability to dynamically grow ensemble sizes as needed, which can be very useful for some types of tasks where the problem space grows [16, 157]. This also allows for wide scaling as child networks can be generated and trained in parallel, avoiding the limitations of low-cost ensemble methods that are restricted to taking models along a sequential training trajectory [50, 70, 93]. Sparse Subnetwork Ensembles that leverage pruning mechanisms result in much smaller ensembles that significantly reduces memory requirements.

## 3.2 Formalization

Consider a parent neural network  $F$  parameterized with a set of weights  $\theta$  optimized via some standard training procedure. We create an ensemble  $E$  consisting of  $k$  child networks, where each child network  $f_i$  has weights  $\hat{\theta}_i$  obtained through a perturbative process  $\xi$  applied to a subset of weights from the parent network determined using a stochastic sampling process  $\phi$ .

$$E = \{f_1, f_2, \dots, f_k\} \tag{3.1}$$

$$f_i(x; \hat{\theta}_i) = F(x; \theta; \xi, \phi) \tag{3.2}$$

The generality of this approach offers a lot of freedom for deciding how Subnetwork Ensembles are formed. The perturbative and subnetwork sampling processes used to create child networks can take arbitrary forms which can be useful for tailoring this algorithm to a wide variety of tasks. Any network architecture can be used with this approach, however it is important to note that the parent network does serve a foundational role in ensemble performance, such that the parent network should be both accurate and have sufficient capacity to enable the creation of high performing children.

### 3.2.1 Subnetwork Sampling

Subnetworks are represented with binary bit mask matrices with the same dimensions of the weight matrices for each layer in the parent network. Consider a weight matrix  $W \in \mathbb{R}^{m \times n}$  representing the weights of a particular layer in a neural network. We generate a matrix  $M \in \{0, 1\}^{m \times n}$ , using a stochastic process  $\phi$ , with the same dimensions as  $W$ . The elements of  $M$  are binary values, where  $M_{ij} = 1$  if the corresponding weight  $W_{ij}$  is retained, and  $M_{ij} = 0$  if the weight is masked. The subnetwork weights  $\hat{W}$  can then be computed by taking the element-wise (Hadamard) product of  $W$  and  $M$ .

$$M \sim \phi^{m \times n} \tag{3.3}$$

$$\hat{W} = W \circ M \tag{3.4}$$

The topology of the subnetwork can be further described by the granularity in which parameters are masked. This granularity can be unstructured or structured, where the former refers to weight-level or connection-level pruning and the latter refers to neuron-level, channel-level, or layer-level pruning. Unstructured pruning tends to result in networks with better predictive capacity while structured pruning results in better computational efficiency. Removing entire rows, columns, or blocks from a network is more easily leveraged by hardware optimizations [6, 101]. However, emerging hardware technology offers a promising future for exploiting unstructured sparse computation [17, 36, 55, 119].

$$(\hat{W}_{uns}, \hat{W}_{str}) = \left( \begin{bmatrix} w_{11} & w_{12} & 0 & w_{14} \\ 0 & w_{22} & 0 & 0 \\ w_{31} & 0 & 0 & w_{34} \\ 0 & 0 & w_{43} & w_{44} \end{bmatrix}, \begin{bmatrix} w_{11} & w_{12} & w_{13} & w_{14} \\ 0 & 0 & 0 & 0 \\ w_{31} & w_{32} & w_{33} & w_{34} \\ 0 & 0 & 0 & 0 \end{bmatrix} \right) \tag{3.5}$$

It's important to also consider the distribution of masked weights throughout non-homogenous networks. It's common for layer configurations in deep neural networks to vary

significantly in size. Severe pruning of small layers may result in bottlenecks that restrict gradient flow. The distribution of masked weights can be controlled with global and local masking methods. Global methods are applied uniformly across the entire network while local methods are applied independently within each layer or sub-region. Global methods tend to result in higher compression rates as more parameters from the larger layers are removed whereas local methods offer more fine-grained control and reduced variance.

### 3.2.2 Neural Partitioning

The transformation of unique subnetworks among child networks is critical at encouraging diverse behavior. However, the random nature of subnetwork sampling can be a large source of variance when constructing child networks. It is prudent to reduce the amount of overlap between the subnetworks inherited by the parent in order produce less correlated models. We introduce Neural Partitioning as an effective tool for achieving variance reduction by dividing neural networks into sets of opposed child networks.

These ideas are inspired by antirandom testing and antithetic sampling [56, 113, 169]. This is also known as mirrored sampling in evolutionary communities, where whenever a random vector  $v = \{x_1, \dots, x_n\}$  is generated, another vector is generated with the opposite signs  $v' = \{-x_1, \dots, -x_n\}$  [12]. Antithetic sampling takes advantage of the fact that the variance of dependent random variables depends on the covariance between them.

$$\text{var}(X) = E[(X - E[X])^2] \tag{3.6}$$

$$\text{var}(X + Y) = E[((X + Y) - E[(X + Y)])^2] \tag{3.7}$$

$$= E[(X - E[X])^2] + E[(Y - E[Y])^2] + E[2(X - E[X])(Y - E[Y])] \tag{3.8}$$

$$= \text{var}(X) + \text{var}(Y) + 2\text{cov}(X, Y) \tag{3.9}$$

If  $X$  and  $Y$  are independent, then the variance of the sum of them is equal to the sum of each. However, if we sample  $Y$  such that it's dependent on  $X$ , then the variance is reduced if the covariance between them is negative.

With Neural Partitioning, we can leverage this idea by generating groups of subnetworks such that the inherited parameters are antithetical and form a disjoint union over the parameters of the parent network. Consider a weight matrix  $W \in \mathbb{R}^{m \times n}$  for a particular layer in a network. Consider a mask  $M \in \{0, 1\}^{m \times n}$  that has been generated via some subnetwork sampling process. The corresponding subnetwork mask with no parameter overlap, also called the anti-mask  $M'$ , is one in which the polarity of all of the mask bits are flipped  $M' = 1 - M$ .

$$(M, M') = \left( \begin{array}{c} \begin{bmatrix} 1 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 1 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 1 & 0 \\ 1 & 0 & 1 & 1 \\ 0 & 1 & 1 & 0 \\ 1 & 1 & 0 & 0 \end{bmatrix} \end{array} \right) \quad (3.10)$$

For large networks, this idea can be extended to partition the parameter space into  $k$  sibling subnetworks with a sparsity level of  $1 - \frac{1}{k}$ . However, creating too many partitions can lead to increasingly sparse subnetworks that exhibit bottleneck effects. This can reduce information flow throughout the network and adversely affect both model performance and training efficiency.

$$(M_1, M_2, M_3) = \left( \begin{array}{c} \begin{bmatrix} 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 \end{bmatrix} \end{array} \right) \quad (3.11)$$

$$(M_1, M_2, M_3) = \left( \begin{bmatrix} 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 \end{bmatrix} \right) \quad (3.12)$$

The goal of Neural Partitioning is to generate subnetwork masks that maximize the Cartesian distance between them. For an ensemble of  $k$  subnetwork masks  $E = \{M_1, M_2, \dots, M_k\}$ , we wish to maximize the total pairwise distance between all ensemble subnetworks. Consider the masks for weight matrices flattened into vectors  $A = \{a_1, \dots, a_N\}$  and  $B = \{b_1, \dots, b_N\}$ . The Cartesian distance ( $CD$ ) and total Cartesian distance ( $TD$ ) for an ensemble can be described with:

$$CD(A, B) = \sqrt{|a_1 - b_1| + \dots + |a_N - b_N|} \quad (3.13)$$

$$TD = \sum_{i=1}^k \sum_{j=1}^k CD(M_i, M_j) \quad (3.14)$$

### 3.2.3 Perturbative Processes

There are a wide variety of ways for defining and applying a perturbative process to a parent network in order to generate child networks. The choice of perturbation mechanism can influence distinct aspects of ensemble performance. We explore three methodologies in Chapters 4, 5, and 6 that have various levels of regularization effects and computational requirements.

In chapter 4, we explore how child networks can be created simply by perturbing subnetwork weights with random noise. Consider a noise matrix  $N \in \mathbb{R}^{m \times n}$ , with the same dimensions of the corresponding layer’s weight matrix, generated with random values drawn from a Gaussian distribution  $\mathcal{N}(\mu, \sigma^2)$ . Assume that a bit mask  $M \in \{0, 1\}^{m \times n}$  was generated with a sampling process  $\phi$  to determine the subnetwork that the perturbations should apply to. The perturbed weights  $\hat{W}$  are then computed by adding the masked noise values

to the original weights  $W$ .

$$M \sim \phi^{m \times n} \tag{3.15}$$

$$N \sim \mathcal{N}(\mu, \sigma^2)^{m \times n} \tag{3.16}$$

$$\hat{W} = W + (N \circ M) \tag{3.17}$$

In chapter 5, we prune subnetwork weights to create sparse topological structures. The perturbation can be described as setting certain parameters indicated by the subnetwork mask to zero and freezing them so that no further weight updates affect them during optimization. This form of perturbation offers a stronger regularization effect than those created with noise perturbations. However, the child networks need to be optimized further to recover the accuracy lost via pruning.

$$\hat{W} = \begin{bmatrix} w_{11} & w_{12} & 0 & w_{14} \\ 0 & w_{22} & 0 & 0 \\ w_{31} & 0 & 0 & w_{34} \\ 0 & 0 & w_{43} & w_{44} \end{bmatrix} \tag{3.18}$$

In chapter 6, we explore how stochasticity can be injected into the perturbation process by leveraging probability matrices that determine how likely it is that a given parameter is retained. Consider a subnetwork mask  $M$  generated from a random sampling process  $\phi$ . The mask is then used as a foundation for a function  $\psi(M)$  to construct a probability matrix  $P \in [0, 1]^{m \times n}$  where each value corresponds to the probability that a parameter is retained or masked. The subnetwork mask is determined on every forward pass through a Bernoulli realization of the probability matrix. We further experiment with annealing the probability matrix over time such that the subnetwork mask evolves to become deterministic throughout

training.

$$M \sim \phi^{m \times n} \tag{3.19}$$

$$P \sim \psi(M)^{m \times n} \tag{3.20}$$

$$\hat{W} = W \circ \text{Bernoulli}(P) \tag{3.21}$$

### 3.2.4 Network Optimization

Once we have a collection of child networks, they can be further optimized to recover lost accuracy as a result of the perturbative processes. The fine tuning procedure generally consists of continued training of each child network independently. Since the child networks inherit parameters from a fully trained parent, they tend to converge in relatively few epochs. Moreover, unique topological transformations as a result of pruning or stochastic masking can promote convergence to diverse local optima. Neuroevolutionary methods that leverage fitness evaluations and model selection strategies can also be used as a form of ensemble optimization. Ultimately, the goal is to construct an ensemble that minimizes the expected loss  $\mathcal{L}(\hat{y}, y)$  of the averaged predictions  $\bar{f}(x)$  for each  $(x, y)$  pair over the dataset  $d_{xy}$ .

$$\text{minimize } \mathbb{E}_{(x,y) \sim d_{xy}} \mathcal{L}(\bar{f}(x), y) \tag{3.22}$$

### 3.2.5 Ensemble Predictions

There are several approaches for combining model predictions in the ensemble. The two most popular methods include majority vote and weighted model averaging, where voting methods operate on the predicted class labels and averaging methods take the weighted

mean of the output distributions of each member [34, 42, 92].

$$\bar{f}_{\text{vote}}(x) = \text{mode}\{\text{argmax}(f_i(x)) : i \in 1, 2, \dots, k\} \quad (3.23)$$

$$\bar{f}_{\text{avg}}(x) = \text{argmax} \left( \frac{1}{k} \sum_{i=1}^k w_i f_i(x) \right) \quad (3.24)$$

While majority vote has a long history with traditional ensemble methods, averaged predictions have become the standardized approach for ensembles of deep neural networks. Throughout this research, we consider ensemble predictions given by the mean of equally weighted ensemble members. Perturbations can result in networks that produce outputs of very different magnitudes. In order to prevent the outsized influence by specific models, the raw output logits of each model are normalized by a softmax function before averaging. This is typically done in classification tasks to convert outputs into probability distributions.

$$\text{softmax}(x_i) = \frac{e^{x_i}}{\sum_j e^{x_j}} \quad (3.25)$$

### 3.2.6 Evaluation Metrics

The bulk of our benchmark experiments in the following chapters focus on typical image classification tasks. In these contexts we report standardized metrics including the accuracy, negative log likelihood, and expected calibration error of ensemble predictions on holdout test sets [90]. Accuracy reports the fraction of correct predictions (Eqn. 3.15). Negative Log Likelihood is a common loss function used for classification tasks in deep learning that measures the difference between predicted probability distributions and true output distributions (Eqn. 3.16). Expected Calibration Error is used for uncertainty estimation to measure how well a model’s estimated probabilities match the true observed probabilities of its predictions (Eqn. 3.17). This involves splitting the data into several bins and taking the average

difference between the accuracy and confidence of the predictions within each bin.

$$Acc = \frac{1}{N} \sum_{i=1}^N 1(\operatorname{argmax} \bar{f}(x_i) \equiv \operatorname{argmax} y_i) \quad (3.26)$$

$$NLL = -\frac{1}{N} \sum_{i=1}^N \sum_{c=1}^C y_{i,c} \log(\bar{f}(x_i)_c) \quad (3.27)$$

$$ECE = \sum_{i=1}^B \frac{|B_i|}{N} |(\operatorname{acc}(B_i) - \operatorname{conf}(B_i))| \quad (3.28)$$

where  $\bar{f}(x)$  represents the prediction made by the ensemble,  $y$  represents the correct output for input  $x$ ,  $N$  is the total number of test samples,  $|B_i|$  is the number of predictions in the bin,  $\operatorname{acc}(B_i)$  is the accuracy of the predictions made in the bin and  $\operatorname{conf}(B_i)$  is the average confidence of model predictions in the bin.

# Chapter 4

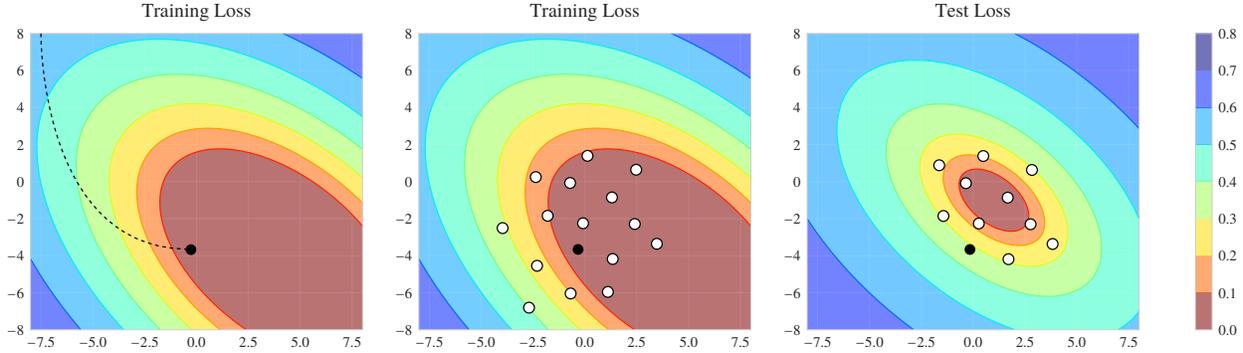
## Noisy Subnetwork Ensembles

### 4.1 Introduction

Neuroevolution is an area of machine learning research that aims to combine evolutionary algorithms and neural networks by using biologically inspired concepts like natural selection, mutation, and reproduction [3, 162, 163]. A core component of evolution strategies involves the generation of child networks by applying noise perturbations to one or more parent networks [5]. However, as neural networks continue to grow in size, the effect of noise perturbations on network behavior becomes increasingly pronounced.

There is often a critical region of mutation where weak perturbations fail to provide any diversity while strong perturbations lead to complete performance collapse. The optimal mutation window can be vanishingly small with modern deep neural networks that contain millions or billions of parameters. Due to the difficulty in mutating these large models, neuroevolutionary methods have primarily seen success with tasks that can be solved by relatively small networks [49]. One approach to alleviating this challenge with large models is to implement sparse mutations, where only a small subset of parameters are modified. This can widen the critical mutation window which allows models to handle stronger perturbations before succumbing to performance collapse. While sparse mutations have long been used in other areas of evolutionary and genetic programming, they have rarely been explored with evolution strategies using deep neural networks [5, 58, 82].

These observations of neuroevolutionary methods connect naturally to our work with Subnetwork Ensembles. In this chapter, we introduce a formulation that we call *Noisy Subnetwork Ensembles* where child networks are generated through noise perturbations applied to random subnetworks in a large, pre-trained parent network. Generating child networks in this manner can alleviate the problems with sample inefficiency that typically plague neu-



**Figure 4.1:** Neuroevolution excels at fine tuning fully-trained networks that get stuck in flat loss basins. The left graph displays a typical SGD training trajectory where models tend to converge to the edges of flat optima [77]. The middle graph displays white dots as child networks that are generated from sparse mutations. The right graph shows the final ensemble consisting of top candidates selected from evaluation on a separate validation set. A key insight to the success of gradient-free methods for fine tuning is that the loss landscapes of test distributions rarely match the training distribution exactly. Ensembling over wide areas of good validation performance is key to improving generalization.

roevolutionary methods. Since child networks retain their network structure and are only perturbed by noise, we eschew continued training of the children and instead implement a standard evolutionary model selection strategy, where we evaluate a large population of models and form the ensemble with those that perform best on a small validation set.

We conduct several ablation studies and visualizations designed to gain insight into the effect that mutation strength and sparsity have on network behavior. We evaluate Noisy Subnetwork Ensembles on the benchmark reinforcement learning environment ProcGen as well as the large scale ImageNet classification task [22,85]. We observe reliable and consistent generalization improvements with over a dozen different deep neural network architectures on several difficult tasks. This approach is flexible and strikes a nice balance with computational cost as these ensembles can improve performance over fully trained parent networks with no additional training required.

## 4.2 Implementation

Noisy Subnetwork Ensembles begin with the standard training of a large parent network  $F$ , parameterized with weights  $\theta$ . We then spawn  $k$  child networks by copying and perturbing

a subset of the weights. The binary subnetwork mask  $M \in \{0, 1\}$  is generated for each layer with the same dimensions as the layer weights  $W$ . The mask can be determined according to any random process  $\phi$  for a given sparsity target  $\rho$ . In our implementation, we use a layerwise unstructured approach where we randomly assign the value for each parameter in the mask by sampling a Bernoulli distribution with probability  $\rho$ , where a successful trial denotes that the parameter will be retained. We also generate a noise matrix  $N$  with the same dimensions as the weight matrix according to some given random distribution. We use a Gaussian distribution with tunable parameters for the mean  $\mu$  and variance  $\sigma^2$ . The sparse mutation for each layer is created by taking the Hadamard product  $\circ$  of the subnetwork mask  $M$  with the noise matrix  $N$ .

$$M \sim \text{Bernoulli}(\rho) \tag{4.1}$$

$$N \sim \mathcal{N}(\mu, \sigma^2) \tag{4.2}$$

$$\hat{W} = W + (N \circ M) \tag{4.3}$$

### 4.2.1 Trust-Region Search

It is important to determine appropriate values for both the mean  $\mu$  and variance  $\sigma^2$  of the noise distribution as well as the amount of sparsity  $\rho$  in the subnetwork mask. This can be challenging to determine before hand due to differences between the size and complexity of datasets, model architectures, optimization hyperparameters, etc.

We implement a grid search over a small validation set, where we measure the average differences between the outputs of a network before and after mutations are applied. Ensuring that the outputs of the network diverge within some bound is referred to as trust region optimization in machine learning contexts [24]. This is a popular technique in reinforcement learning algorithms as they promote stable policy updates in environments where too large of a change may lead to performance collapse such that the model will struggle to recover [130, 131].

A standard approach for establishing the trust region would be to measure the squared difference between output predictions before and after mutation. Given a network  $F$  that is parameterized by  $\theta$  with  $C$  outputs and a network that is perturbed with a noise mutation  $\gamma$ , the mean squared difference over a set of  $N$  total input samples  $X$  can be described as:

$$MSE = \frac{1}{N} \sum_{n=1}^N \sum_{c=1}^C (F(X_n; \theta)_c - F(X_n; \theta + \gamma)_c)^2 \quad (4.4)$$

However, since classification networks are trained with cross entropy loss, where outputs are treated as probability distributions, the mean squared error may not be the best fit for approximating the effect of that perturbation as it can be prone to increased variance. Instead, it is more common to describe the difference between output distributions with the Kullback-Leibler Divergence, also known as relative entropy. This allows for a stable measure of comparison that works regardless of network architecture or noise distribution.

$$KL = \frac{1}{N} \sum_{n=1}^N \sum_{c=1}^C F(X_n; \theta)_c \log \left( \frac{F(X_n; \theta)_c}{F(X_n; \theta + \gamma)_c} \right) \quad (4.5)$$

A grid search algorithm can then be implemented to try out several values for the mutation strength and sparsity parameters in order to target a divergence score that represents the boundary of a trust region. One nice property revealed in our experiments is that there is a roughly linear relationship between KL divergence and accuracy degradation in trained models. This insight can be used to tweak mutation parameters according to problem complexity, ensemble size, or parameter sensitivity where a larger KL target can allow for stronger mutations and better exploration at the cost of potentially less accurate candidate models.

### 4.2.2 Mirrored and Partitioned Noise

Using the above techniques, we can then generate a collection of child networks with subnetworks perturbed according to parameters found through this trust region grid search.

---

**Algorithm 1:** Noisy Subnetwork Ensembles

---

**Input:**  $N$ , the number of candidate models  
**Input:**  $K$ , the number of ensemble members  
**Input:**  $(X_{train}, X_{val}, X_{test})$ , the training, validation, and test data  
**Input:**  $\eta$ , the optimizer hyperparameters  
**Input:**  $\Delta$ , trust region boundary target  
**Input:**  $\phi(\rho)$ , the subnetwork sampling process with target sparsity  $\rho$

$O = \text{sgd}(\eta)$

$F = \text{initialize}()$

$F.w = \text{train}(F, O, X_{train})$

$\rho, \mu, \sigma^2 = \text{trustRegionSearch}(F, X_{val}, \Delta)$

$\text{candidates} = []$

**for**  $i$  in 1 to  $N$  **do**

$f_i = \text{initialize}()$

$f_i.w = F.w$

**for**  $j$  in  $f_i.layers$  **do**

$M \sim \phi(\rho)^{f_{i,j}.w}$

$N \sim \mathcal{N}(\mu, \sigma^2)^{f_{i,j}.w}$

$f_{i,j}.w = F_{i,j}.w + (N \circ M)$

**end**

$\text{fitness} = \text{evaluate}(f_i, X_{val})$

$\text{candidates} \leftarrow (f_i, \text{fitness})$

**end**

$\text{ensemble} = \text{select}(\text{candidates}, K)$

**for**  $(x, y)$  in  $X_{test}$  **do**

$\text{outputs} = [F(x)]$

**for**  $c$  in  $\text{ensemble}$  **do**

$\text{outputs} \leftarrow c(x)$

**end**

$\text{predictions} = \text{softmax}(\text{mean}(\text{outputs}))$

$\text{loss} = \text{nll}(\text{predictions}, y)$

$\text{accuracy} = \text{acc}(\text{predictions}, y)$

**end**

---

Neural Partitioning and Antithetical Sampling can also be leveraged when generating child networks [12]. A single generated subnetwork mask  $M$  and noise matrix  $N$  can be used to generate four child networks  $\hat{W}_i$  perturbed in opposite directions and subspaces, resulting in reduced variance and a balanced exploration of the parameter space.

$$M \sim \text{Bernoulli}(\rho) \tag{4.6}$$

$$N \sim \mathcal{N}(\mu, \sigma^2) \tag{4.7}$$

$$\hat{W}_1 = W + (N \circ M) \tag{4.8}$$

$$\hat{W}_2 = W + (N \circ (1 - M)) \tag{4.9}$$

$$\hat{W}_3 = W - (N \circ M) \tag{4.10}$$

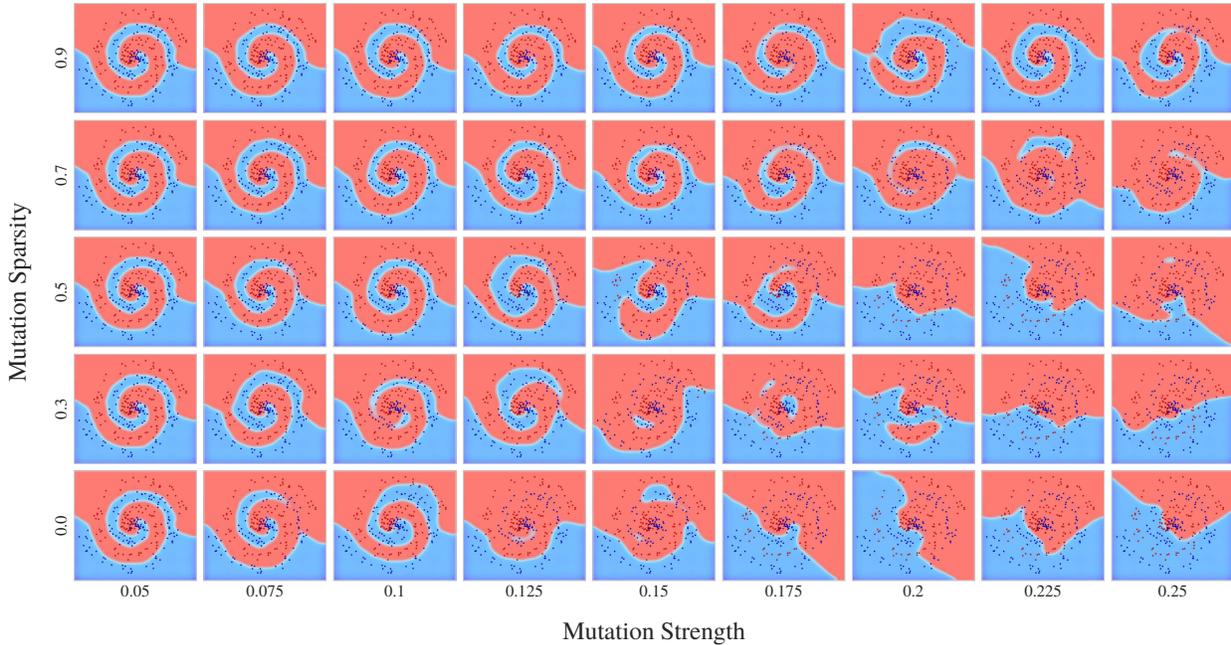
$$\hat{W}_4 = W - (N \circ (1 - M)) \tag{4.11}$$

We can gain further improvements (without conducting any additional optimization) by implementing a model selection process. For each candidate child network in the population, we evaluate its performance on a holdout validation set where the accuracy is recorded as a measure of its fitness. The top  $k$  models with the best performance will then be selected to become members of the final ensemble.

## 4.3 Experiments

### 4.3.1 Decision Boundaries

We begin by exploring the interplay between mutation sparsity and mutation strength by visualizing how changing these values affects network predictions. We use a simple three layer fully connected multilayer perceptron that contains 64 neurons in each layer. We train this model on a binary interleaved spiral dataset that contains 2500 sample (x, y) points. We train for 10 epochs using the Adam optimizer [84] with a learning rate 0.001 and use this trained model as the starting point for all mutations.



**Figure 4.2:** The images above detail decision boundaries for a trained three layer multilayer perceptron after being perturbed with various mutations. With dense perturbations, there is a comparatively small window between a functional decision boundary and complete performance collapse. As mutations become more sparse, the model is better able to retain its behavior while the strength of parameter mutation increases.

We then perturb the model with mutations sampled from a Gaussian Distribution  $N \sim \mathcal{N}(0, \sigma^2)$  applied to random subnetworks with varying levels of sparsity. The subnetwork masks are generated in a layerwise unstructured fashion by sampling  $M \sim \text{Bernoulli}(\rho)$  where  $\rho$  corresponds to the sparse probability. We evaluate several values for the variance of the noise perturbation  $\sigma \in [0.05, 0.25]$  and for subnetwork sparsities  $\rho \in [0.0, 0.9]$ .

Using the perturbed models, we make predictions on a holdout test set containing 250 samples. In figure 4.2, we display a grid of the decision boundaries of these perturbed models on the test set as we ablate between mutation strength and sparsity. When mutations are dense, we see a very quick collapse of performance for small perturbations. The window for optimal dense mutation is quite small even for this toy network and simple dataset. When mutations become more sparse, the model is able to maintain good classification boundaries while displaying small variations of prediction diversity. This kind of behavior is desirable

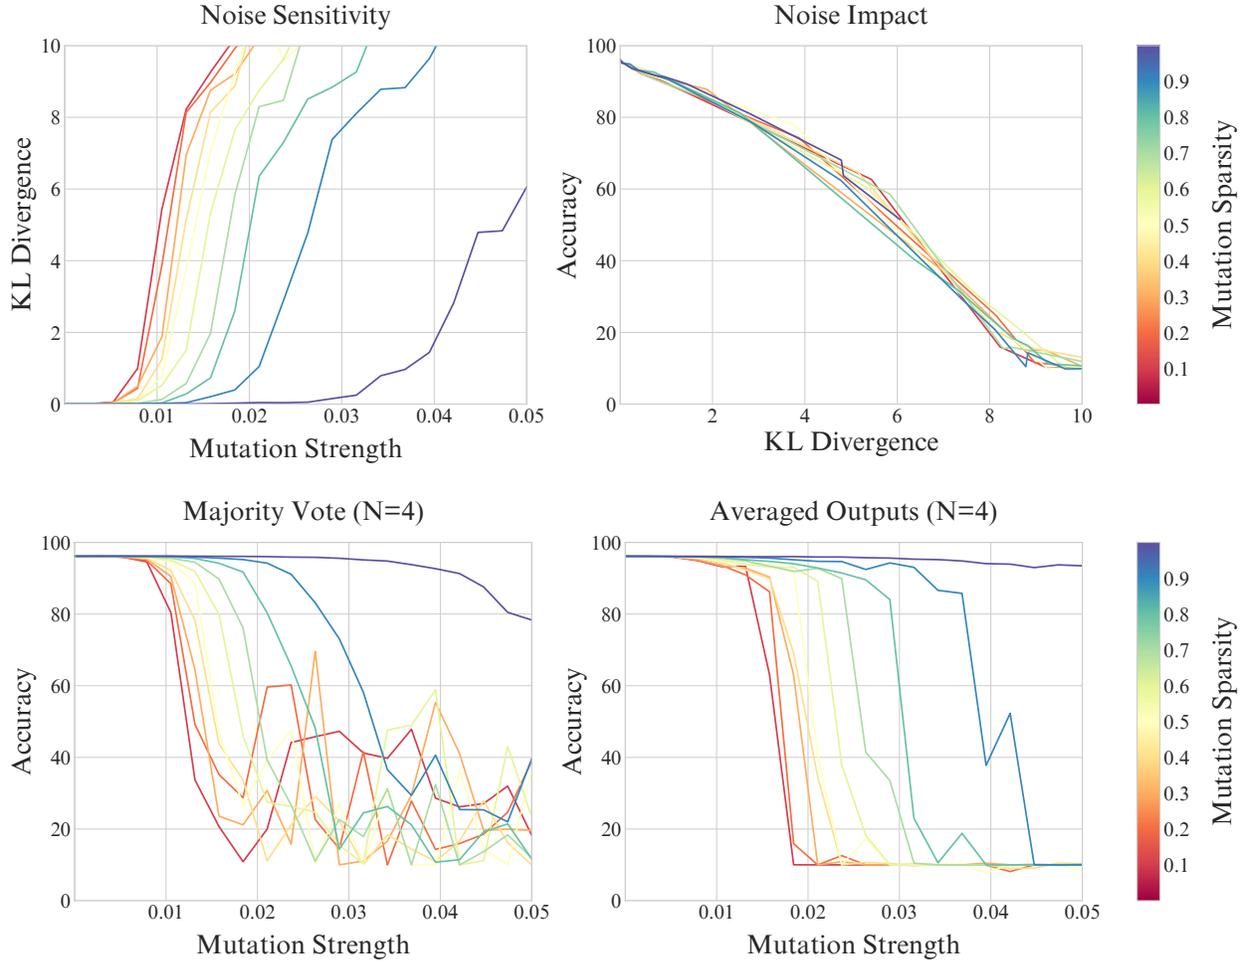
for neuroevolutionary populations as ensemble learning research has shown that ensembles perform better with large numbers of both accurate and diverse members [8].

### 4.3.2 Ablations

Next, we aim to explore whether the intuitions about how mutations affect predictions translate to a much larger convolutional network on benchmark computer vision datasets. For this experiment, we conduct ablations on the CIFAR-10 and CIFAR-100 datasets [85]. These are popular benchmark datasets and their use is widespread in computer vision research. They each contain 50,000 training and 10,000 test samples of colored 32x32 pixel images. CIFAR-10 contains samples belonging to one of 10 classes while CIFAR-100 contains samples belonging to one of 100 classes. We use a WideResNet-28x10 model for our parent network, which is a highly accurate network architecture that contains  $\sim 36M$  parameters. This network is a variant on the popular ResNet that decreases the depth and increases the width of each convolutional layer [176].

We implement a standard training algorithm for this type of model where we train for 100 epochs using Stochastic Gradient Descent with momentum [144]. A stepwise learning rate decay is used where an initial value of 0.1 decays to 0.01 after 50% of training and decays again to 0.001 for the final 10% of training. We use standard data augmentation schemes for CIFAR that includes a random crop and random horizontal flip along with mean standard normalization. We split the test set in half in order to construct a validation set used for model selection.

The parent networks achieve an accuracy of approximately 96% on CIFAR-10 and approximately 80% on CIFAR-100. Using these saved models as parent networks, we then perturb the models with varying amounts of mutation strengths,  $\sigma \in (0.0, 0.05]$ , and sparsities,  $\rho \in [0.01, 0.99]$ , and we evaluate their performance on the test sets. Figure 4.3 displays the results of these experiments where we start by measuring the effect that mutations have on KL Divergence. Predictably, we see that KL divergence quickly increases as the density



**Figure 4.3:** Results of mutation ablations for a trained WideResNet-28x10 model on CIFAR-10. Lines of different colors correspond to the percentage of sparsity for the mutations. The top left graph reports the average KL divergence as a result of mutation strength. The top right graph explores the relationship between KL divergence and accuracy. The bottom left reports the accuracy of an ensemble evaluated with majority vote predictions. The bottom right reports accuracy of an ensemble evaluated with averaged predictions.

and strength of mutations increase. The rapid increase in KL divergence for dense perturbations illustrates how quickly performance collapses between very small changes in mutation strength. We also observe an interesting linear relationship between KL divergence and model accuracy.

We then implement Noisy Subnetwork Ensembles where a population of 16 models are created by perturbing the parent model according to a given mutation strength and sparsity.

The top 4 models with the best accuracy on the validation set are then selected. We report the accuracy on the test set with this ensemble when predictions are decided via majority vote and averaged output distributions. The averaged predictions approach is much more consistent than majority vote, which display large amounts of variance as mutations become stronger. The accuracy of both combination methods are comparable with no noticeable difference at low levels of mutation. In settings with strong mutations, the averaged predictions model exhibits more robust behavior and better accuracy than majority vote. The point where the mutation started to negatively impact member accuracy corresponds to an observed KL divergence approximately between 0.02 to 0.05.

### 4.3.3 Benchmarks

#### ImageNet

We then explore our approach on the difficult benchmark computer vision dataset, ImageNet [27,86]. ImageNet is a large scale collection of images that have been hand labelled for use in machine learning tasks and is organized according to the wordnet dataset hierarchy. Over 14 million images and 20,000 labels have been collected in total. We use the 2012 ImageNet collection which consists of a training set of 1.2 million images and a validation set of 50,000 images, each belonging to one of 1000 categories. Images have varying sizes and aspect ratios and consist of both colored and grayscale photos. Images are normalized with mean standard normalization and we implement standard data augmentations which consists of resizing to 256x256 pixels and center cropping to 224x224 pixels.

We evaluate our approach with ten popular deep neural network architectures of varying sizes and generalization capacities in order to demonstrate the generality of this approach across different contexts [62, 68, 71, 74, 85, 134, 146, 171, 176, 181]. All networks are pretrained and available from the torchvision repository [123].

We break the ImageNet competition test collection (ILSVRC 2012) of 50,000 images into a 80/20 split between test and validation sets that now each contain 40,000 and 10,000

**Table 4.1:** Results for sparse mutation ensembles on ImageNet with a wide variety of models. Mutation sparsity and strength are determined from a small hyperparameter grid search. Accuracy (Acc), Negative Log Likelihood (NLL), and expected calibration error (ECE) are reported for the parent and the ensemble. Metrics prepended with  $e$  refer to ensemble results.  $\Delta\text{Acc}$  is the change in accuracy between the parent and the ensemble,  $\sigma$  is the mutation strength, and  $\overline{KL}$  is the average output divergence between the mutated models and the parent models. 16 models are generated and the 8 most accurate candidates on a validation set are used together as an ensemble.

Model	Acc $\uparrow$	NLL $\downarrow$	ECE $\downarrow$	eAcc $\uparrow$	eNLL $\downarrow$	eECE $\downarrow$	$\Delta\text{Acc}$ $\uparrow$	Parameters	$\sigma$	$\overline{KL}$
AlexNet	56.46	1.904	0.021	56.52	1.903	0.019	0.06	61.1M	0.0025	0.046
DenseNet-121	74.43	1.014	0.024	74.68	1.004	0.021	0.25	8.0M	0.0012	0.040
Inception-V3	69.57	1.819	0.184	69.98	1.681	0.169	0.41	27.2M	0.0007	0.049
MobileNet-V2	72.12	1.136	0.072	72.14	1.132	0.024	0.02	3.5M	0.0010	0.045
ResNet-18	69.76	1.247	0.026	69.93	1.238	0.022	0.17	11.7M	0.0020	0.046
ResNext-50	77.64	0.945	0.065	77.72	0.929	0.059	0.08	25.0M	0.0012	0.051
ShuffleNet-V2	69.51	1.354	0.072	69.57	1.351	0.071	0.06	2.3M	0.0020	0.052
SqueezeNet	58.10	1.852	0.017	58.14	1.852	0.017	0.04	1.2M	0.0020	0.026
VGG-16	71.62	1.140	0.027	71.64	1.138	0.028	0.02	138.4M	0.0025	0.048
WideResNet-50	78.47	0.879	0.054	78.60	0.852	0.038	0.13	68.9M	0.0015	0.049

samples respectively. All fitness evaluations and model selections use the validation set and all reported results are evaluated on the holdout test set. We report accuracy, negative log likelihood, and expected calibration error for the parent network and the ensemble [90]. We ran each model twice and report the best results.

We conduct a hyperparameter grid search in order to find appropriate mutation strengths for each model. Using a separate holdout dataset of 1000 samples we measure the average KL Divergence and accuracy while we ablate mutation strength  $\sigma \in [0.001, 0.015]$  with a mutation sparsity of 0.5. We then choose both a mutation strength that maximizes accuracy while targeting a KL Divergence of approximately 0.02 which was found empirically to be a safe range for mutation where the the accuracy of candidate models is balanced with exploration. We then construct a population of 16 models with neural partitioning and mirrored sampling. We first generate a random noise mutation sampled from a Gaussian distribution modified with mutation strength. We construct a mirrored version of this random vector by flipping the signs of each value. We then use neural partitioning to generate two opposed unstructured subnetwork masks each with 50% sparsity Each child network is then evaluated on a validation set and the top 8 candidates with the best accuracy are selected. We evaluate

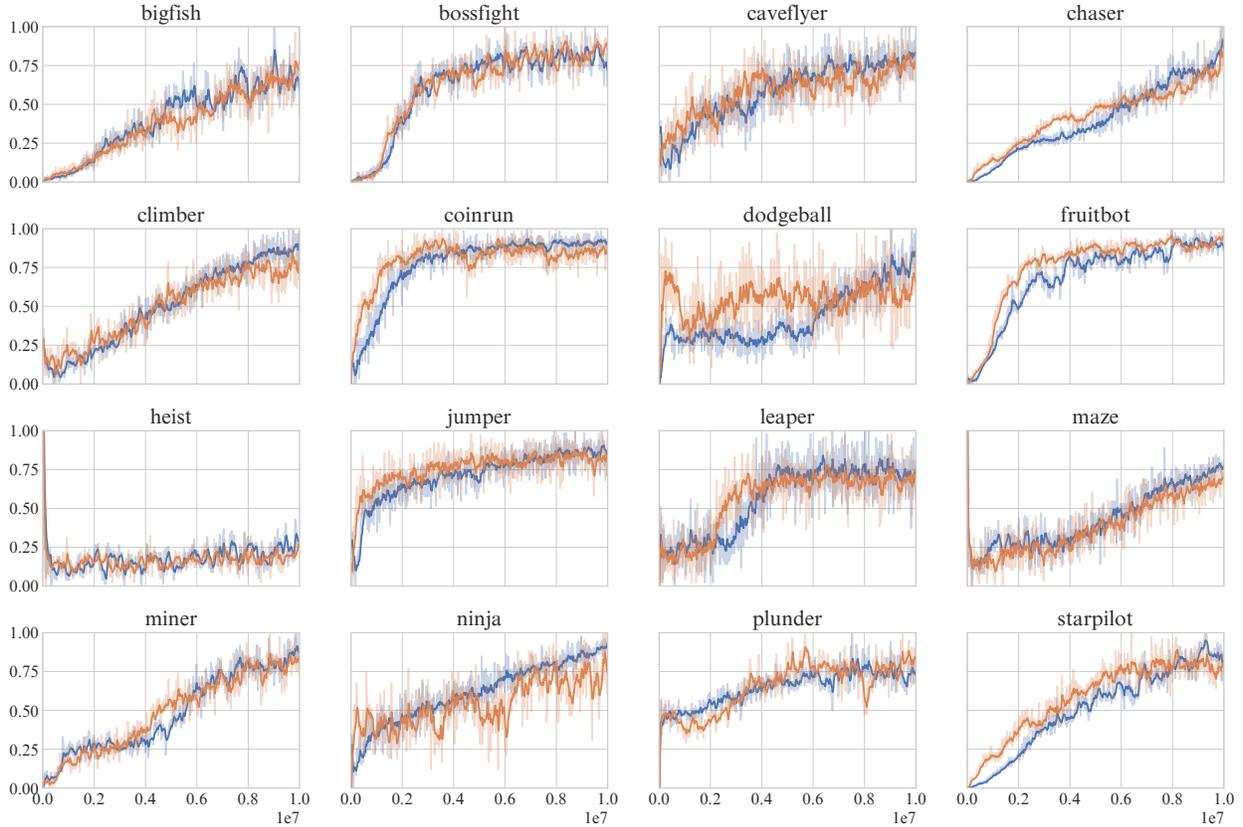
the top 8 candidates and combine their output predictions together to report the ensemble accuracy, negative log likelihood and expected calibration error.

Table 4.1 contains both the parent results and the ensemble results. We see consistent improvement in all metrics for each model with Noisy Subnetwork Ensembles. While the difference between generalization performance is very small in many cases, it is notable that convergence is stable and appears to be monotonic with optimal mutation hyperparameters. However, there’s a question as to how much potential is able to be squeezed out of these fully trained networks without any further optimization to the child networks. There is no accepted method for determining the theoretical ceiling of generalization capacity for complex network architectures. These small improvements may be considered significant in some contexts where an improvement of 0.5% accuracy on a dataset of 50,000 images corresponds to 250 more correct predictions.

## **Progen**

We additionally explore the efficacy of Noisy Subnetwork Ensembles on the difficult Progen reinforcement learning benchmark [22]. Progen consists of 16 unique game environments that use procedural generation to produce distinct and diverse levels for every run. The observation space for all environments are 64x64x3 RGB images and the action space consists of 15 discrete options that correspond to behavior within the game (e.g. left, right, jump, shoot). Progen can generate endless numbers of unique levels which allows for a more robust evaluation of generalization, which has previously been challenging to evaluate in other reinforcement learning benchmarks [179].

We use the Proximal Policy Optimization algorithm to train our models [131]. This is one of the most popular and reliable baseline algorithms in deep reinforcement learning. Proximal Policy Optimization updates policies by taking the largest step possible to improve performance while satisfying a constraint on how close new and old policies can be. The motivation is to avoid a common failure point in reinforcement learning where too significant of a change in policy can result in catastrophic performance collapse.



**Figure 4.4:** Training graphs on all 16 environments in the ProcGen benchmark. Orange lines correspond to single model evaluations and blue lines correspond to noisy subnetwork ensembles.

**Table 4.2:** Proximal Policy Optimization hyperparameters and evaluation results over 10 episodes.

Parameter	Value	Environment	Baseline	Noisy Subnetworks	$\sigma$
$\gamma$	0.999	Bigfish	$13.7 \pm 5.47$	$25.9 \pm 5.59$	0.06
$\lambda$	0.95	Bossfight	$9.1 \pm 1.93$	$12.4 \pm 0.32$	0.05
Epochs/rollout	3	Caveflyer	$3.9 \pm 1.53$	$7.3 \pm 1.32$	0.07
Timesteps/rollout	256	Chaser	$2.3 \pm 0.22$	$3.7 \pm 1.0$	0.01
Minibatches	8	Climber	$5.1 \pm 1.61$	$7.0 \pm 1.75$	0.01
Batch size	2048	Coinrun	$8.0 \pm 1.26$	$7.0 \pm 1.44$	0.06
Entropy Bonus	0.01	Dodgeball	$0.8 \pm 0.51$	$2.0 \pm 0.56$	0.05
Target KL	0.01	Fruitbot	$30.2 \pm 3.16$	$32.4 \pm 0.69$	0.04
Clip Range	0.2	Heist	$3.0 \pm 1.54$	$4.0 \pm 1.55$	0.10
Framestack	2	Jumper	$6.0 \pm 1.55$	$7.0 \pm 1.45$	0.02
Normalization	Reward	Leaper	$6.0 \pm 1.55$	$9.0 \pm 0.94$	0.01
Learning Rate	5e-4	Maze	$6.0 \pm 1.54$	$5.0 \pm 1.58$	0.04
Workers	1	Miner	$7.4 \pm 1.78$	$4.7 \pm 1.74$	0.09
Envs/worker	64	Ninja	$5.0 \pm 1.58$	$7.0 \pm 1.45$	0.01
Total Timesteps	10m	Plunder	$10.6 \pm 3.02$	$13.6 \pm 2.97$	0.06
Network	Impala	Starpilot	$20.1 \pm 5.42$	$39.1 \pm 6.85$	0.01
Block Config	[32, 64, 128]	Mean	$8.58 \pm 1.86$	$11.69 \pm 2.74$	$0.04 \pm 0.01$

We train an Impala Convolutional Neural Network for 10 million time steps on each environment [37]. The Impala architecture is a benchmark residual convolutional network for reinforcement learning tasks that contains 15 layers and approximately 1.6 million parameters. On evaluation rollouts, we create a Noisy Subnetwork Ensemble consisting of 8 models where a random selection of 20% of the weights are perturbed by sampling from a Gaussian distribution with mutation strength scaled by  $\sigma \in [0.01, 0.10]$ . Perturbations are applied in a layerwise unstructured fashion. Actions are chosen by averaging the softmax outputs of the ensemble members.

Figure 4.4 and Table 4.2 includes evaluation graphs and final results on each environment for both Noisy Subnetwork Ensembles as well as the single parent model. All results are reported as the mean and standard deviation over 10 runs. Noisy Subnetwork Ensembles improved results in most environments with a few exceptions, potentially due to the relation between ensembles and probabilistic actions. Proximal Policy Optimization is a stochastic algorithm that chooses actions according to a probability distribution. When combining the predictions of child networks, the output distribution tends to magnify the most likely action. This can be detrimental in environments where stochasticity is helpful for encouraging exploration as the deterministic choice may lead to situations where an agent gets stuck in a state where it is unable to recover.

## 4.4 Discussion

We introduce Noisy Subnetwork Ensembles as a formulation of Subnetwork Ensembles where child networks are generated by perturbing sampled subspaces of a parent network with random noise mutations. This proves to be an effective approach to alleviating the challenges of mutating deep neural networks in evolution strategies. Subnetwork perturbations widen the critical mutation window which allows for stronger mutations and more diversity before succumbing to performance collapse. We explore how these subnetwork mutations reduce

variance and we show how a model selection strategy can further improve performance in these ensembles.

We conduct several ablation studies in order to explore the interplay between sparsity and mutation strength on network behavior. We conduct a decision boundary experiment where visualizations are created that show the predictions of a model on a binary classification dataset. After being perturbed with dense mutations, the model sees a rapid decline in performance while sparsity significantly helps in maintaining accurate predictions. We then explore how these insights translate to a wide residual network on CIFAR where we explore parameter sensitivity and differences between averaged weights and averaged predictions for various levels of mutation strength and sparsity. Our findings reaffirm the idea that subnetwork mutations produce accurate models more reliably than dense mutations.

We then introduce two large scale benchmark evaluations of Noisy Subnetwork Ensembles on ImageNet and Procgen, with a wide variety of deep neural network architectures. We use a trained parent network to generate child networks with noise perturbation hyperparameters found through a short grid search algorithm. We observe a small yet reliable improvement in generalization for every model on ImageNet and for most environments on Procgen. This formulation of Noisy Subnetwork Ensembles offers an intriguing computational tradeoff as child networks can be quickly generated to obtain small generalization improvements, without requiring any additional training.

# Chapter 5

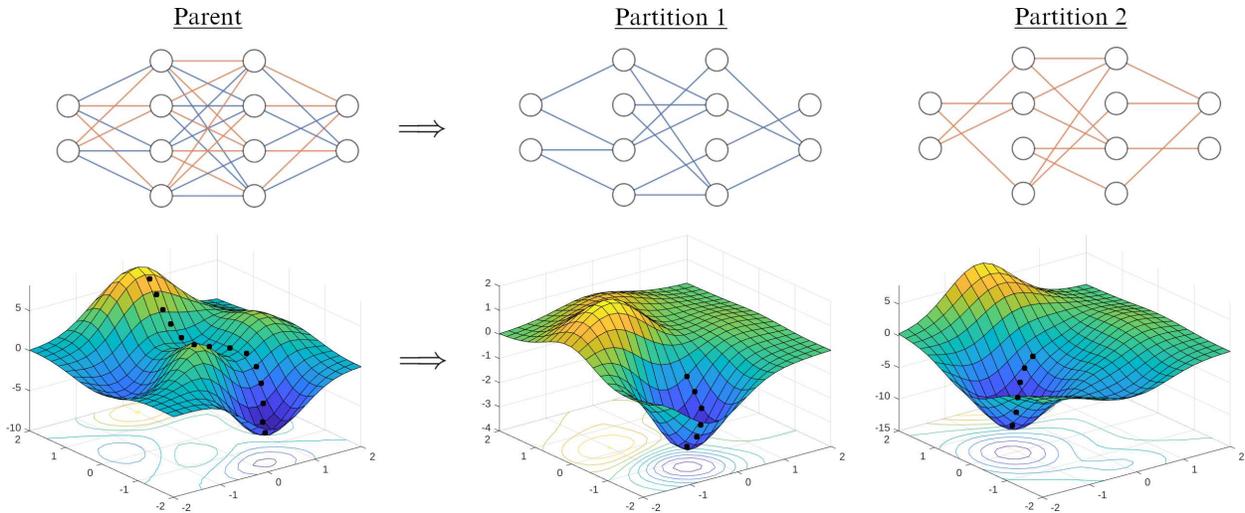
## Sparse Subnetwork Ensembles

### 5.1 Introduction

Noisy Subnetwork Ensembles were introduced as a no-cost formulation of Subnetwork Ensembles where child networks are generated through subnetwork perturbations and ensemble together without applying any additional training. The combination of these perturbed models was shown to reliably improve generalization performance on holdout data, however the differences were relatively small. Without further optimization of the child networks, there is a limited potential for improving the predictive capacity of each child network. Additionally, continued training of the child networks in Noisy Subnetwork Ensembles would naturally lead to convergence towards highly correlated areas of the optimization landscape as they share the vast majority of their parameters and network structure. This lack of diversity can limit performance and lead to less robust ensembles.

We extend this work by introducing *Sparse Subnetwork Ensembles*, a formulation of Subnetwork Ensembles where child networks are created through the pruning of subnetworks in a trained parent network. The sparse child networks are then subsequently trained for a small number of epochs in order to recover the lost accuracy from pruning. The unique network topology of each child encourages convergence to unique and diverse minima, resulting in more robust ensembles. We also explore several methods for improving diversity among child networks, including Neural Partitioning for generating children with mirrored network structures, and a cyclic learning rate schedule that alternates between large and small values to promote more exploration in weight space before converging [138].

This approach is highly flexible and offers several unique benefits over other low-cost ensemble methods. Like Noisy Subnetwork Ensembles, additional child networks can be dynamically generated with little additional computation. Generating ensemble members



**Figure 5.1:** A diagram explaining network partitions in Sparse Subnetwork Ensembles. The parent network is optimized and settles into some optima. Child networks are created via neural partitioning. That is, two sparse children each inherit a different set of parameters from the parent. The optimization landscapes for each child are diverse, yet they converge quickly due to inheritance from a fully trained parent.

through pruning results in much smaller networks that significantly reduce memory usage and computational cost during forward passes. This approach attempts to strike the optimal balance between leveraging the improved training dynamics associated with using a large parent network, while also creating accurate and diverse child networks that converge quickly to good optima.

We compare Sparse Subnetwork Ensembles to several popular state-of-the-art ensemble learning algorithms with benchmark deep neural network architectures and datasets. We conduct hyperparameter ablation studies with child sparsity, pruning structures, tuning schedules, and ensemble sizes of up to 128 members. We demonstrate that Sparse Subnetwork Ensembles are both highly efficient, accurate, and robust with both small and large training budgets.

## 5.2 Implementation

Sparse Subnetwork Ensembles begin with the standard training of a large parent network  $F$ , parameterized with weights  $\theta$ . Given a fixed training budget, it can be difficult to deter-

mine the appropriate amount of time to allot for training the parent as model architecture, dataset complexity, total training budget, ensemble size, and inference compute limitations can all have a strong impact on performance for different ensemble configurations. As a general guideline, our large scale image classification experiments with Wide Resnets suggest a 70/30 split between parent/children training time for a budget of 200 epochs.

We then spawn  $k$  child networks by copying the parent and pruning a subset of the weights. The binary subnetwork mask  $M \in \{0, 1\}$  is generated for each layer with the same dimensions as the layer weights  $W$ . The values for each parameter in the mask is determined by sampling a Bernoulli distribution with probability  $\rho$ , where a successful trial denotes that the parameter will be retained. We also explore an implementation where the mask is determined by sampling such that neurons are masked rather than individual connections. The sparse subnetwork is created by taking the Hadamard product  $\circ$  of the mask  $M$  with the weights  $W$  at each layer.

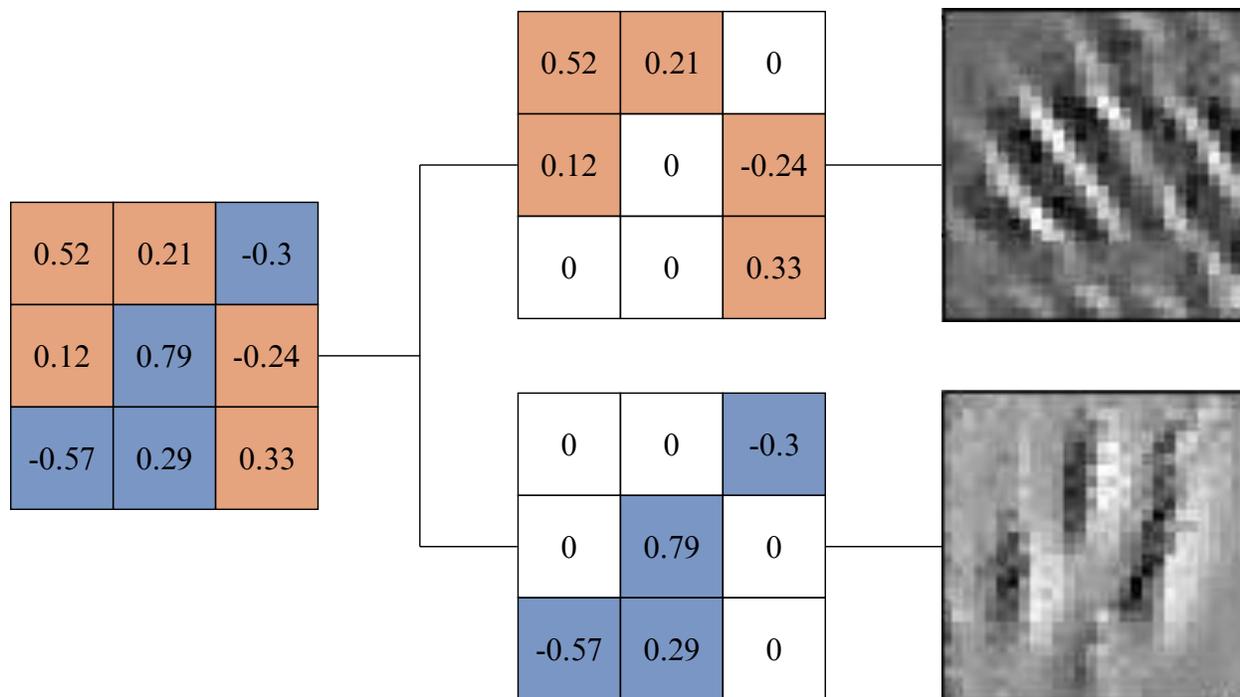
We additionally implement Neural Partitioning when constructing child networks of 50% sparsity. In this case, whenever we generate a child network, we construct the anti-random subnetwork by flipping the polarity of the bits in the subnetwork mask of the original child. This can be repeated  $n$  times to create an ensemble of size  $2n$ . The weights for each child network  $\hat{W}$  and it's anti-random complement  $\hat{W}'$  can be described with.

$$M \sim \text{Bernoulli}(\rho) \tag{5.1}$$

$$\hat{W} = W \circ M \tag{5.2}$$

$$\hat{W}' = W \circ (1 - M) \tag{5.3}$$

This works best with a child sparsity target of 50% such that the size of all the child networks in the ensemble are roughly equivalent. Given a large enough parent network, the children can be partitioned to greater degrees (e.g. three child networks that inherit 33%



**Figure 5.2:** A convolutional parent network filter split into two partitioned child filters. The child networks learn diverse feature representations due to the geometric constraints of their unique topologies.

of the parent’s parameters each). However, children with higher sparsity may require more training time to achieve similar performance.

Given the high dimensionality of deep neural networks, the probability that any two child networks will end up with similar topologies is small since the number of unique graph structures in random and complex graphs grow exponentially with the number of nodes [60]. The unique network topology of each child encourages optimization of diverse feature representations which helps to produce more robust ensembles. Figure 5.2 illustrates how pruning convolutional filters applies unique geometric constraints which inherently forces the production of diverse feature maps.

Each child network is then tuned for a small number of epochs. Since each child inherits its parameters from an accurate parent, the convergence to good solutions often occurs after a relatively small number of epochs [6]. There are several tunable hyperparameters related to optimization that influence the performance of the fine tuning phase. These include the

---

**Algorithm 2:** Sparse Subnetwork Ensembles

---

**Input:**  $N$ , the number of ensemble members  
**Input:**  $(X_{train}, X_{test})$ , the training and test data  
**Input:**  $(t_{parent}, t_{child})$ , the number of parent and child epochs to train for  
**Input:**  $(\eta_{parent}, \eta_{child})$ , the parent and child optimizer hyperparameters  
**Input:**  $\phi(\rho)$ , the subnetwork sampling process with target sparsity  $\rho$

$O = \text{sgd}(\eta_{parent})$   
 $F = \text{initialize}()$   
 $F.w = \text{train}(F, O, X_{train}, t_{parent})$

$ensemble = []$

**for**  $i$  in 1 to  $N$  **do**  
     $f_i = \text{initialize}()$   
     $f_i.w = F.w$   
  
    **for**  $j$  in  $f_i.layers$  **do**  
         $M \sim \phi(\rho)^{f_{i,j}.w}$   
         $f_{i,j}.w \leftarrow f_{i,j}.w \circ M$   
    **end**  
  
     $O = \text{sgd}(\eta_{child})$   
     $f_i.w = \text{train}(f_i, O, X_{train}, t_{child})$   
  
     $ensemble \leftarrow f_i$   
**end**

**for**  $(x, y)$  in  $X_{test}$  **do**  
     $outputs = [F(x)]$   
  
    **for**  $c$  in  $ensemble$  **do**  
         $outputs \leftarrow c(x)$   
    **end**  
  
     $predictions = \text{softmax}(\text{mean}(outputs))$   
     $loss = \text{nll}(predictions, y)$   
     $accuracy = \text{acc}(predictions, y)$   
**end**

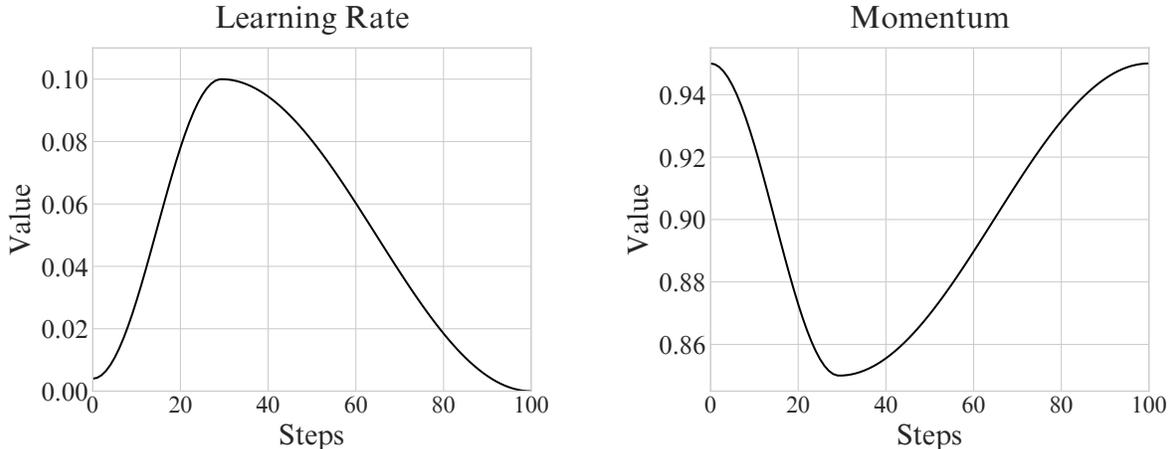
---

number of epochs, learning rate schedule, batch size, weight decay, and data augmentation strategies. These optimization hyperparameters are generally kept consistent with those used for training the parent network [124]. Parent networks are typically trained with a learning rate schedule that starts with a large value and slowly decays to a small value over time. The most common method for tuning subnetworks in pruning literature is to use a small constant learning rate that is consistent with the final phase of training the parent. However, recent work has found that mimicking the decaying schedules used for the parent network can be highly effective for tuning subnetworks [95, 124].

We experiment with a learning rate schedule called the one-cycle policy, which has been shown to be highly effective for rapidly training deep neural networks from scratch [138]. The one-cycle policy consists of a warm-up, cool-down, and annihilation phase, which alters the learning rate during training for each batch of data. The warm up phase begins with a small learning rate that ramps up to a large maximum rate. The cool-down phase gradually decays the maximum rate to a minimum rate that is lower than the initial learning rate. The annihilation phase takes place in the very final stages of training where the learning rate is decayed to 0.

The one-cycle policy is highly effective at encouraging diversity among child networks as the large learning rates enable each child to move greater distances away from the initial weights before converging to optima. This idea has been shown to be effective in other successful low-cost ensemble methods and is the primary insight underlying several popular temporal ensemble methods such as Snapshot Ensembles and Fast Geometric Ensembles [50, 70].

The original introduction of the one-cycle policy utilized linear annealing to modify the learning rates. However, cosine annealing tends to be more popular in subsequent implementations. In our experiments we use a revised one-cycle schedule that disregards the annihilation phase and uses only the warm-up and cool-down. Additionally, we use the inverse of the one-cycle policy for momentum when tuning with stochastic gradient descent [144].



**Figure 5.3:** An example of the one-cycle learning rate and momentum schedule we use to fine-tune the child networks in Sparse Subnetwork Ensembles.

Both of these modifications have been shown empirically in unpublished work to perform better than the original implementations [69].

The value for the learning rate  $\eta$  at iteration  $t$ , where  $\eta_{init}$  is the initial learning rate value,  $\eta_{max}$  is the maximum value,  $\eta_{min}$  is the minimum value and  $T$  is the total number of iterations can be described as:

$$\eta_{warm}(t) = \eta_{init} + \frac{1}{2}(\eta_{max} - \eta_{init}) \left( 1 - \cos \left( \frac{\pi t}{T} \right) \right) \quad (5.4)$$

$$\eta_{cool}(t) = \eta_{min} + \frac{1}{2}(\eta_{max} - \eta_{min}) \left( 1 + \cos \left( \frac{\pi t}{T} \right) \right) \quad (5.5)$$

## 5.3 Experiments

### 5.3.1 Ablations

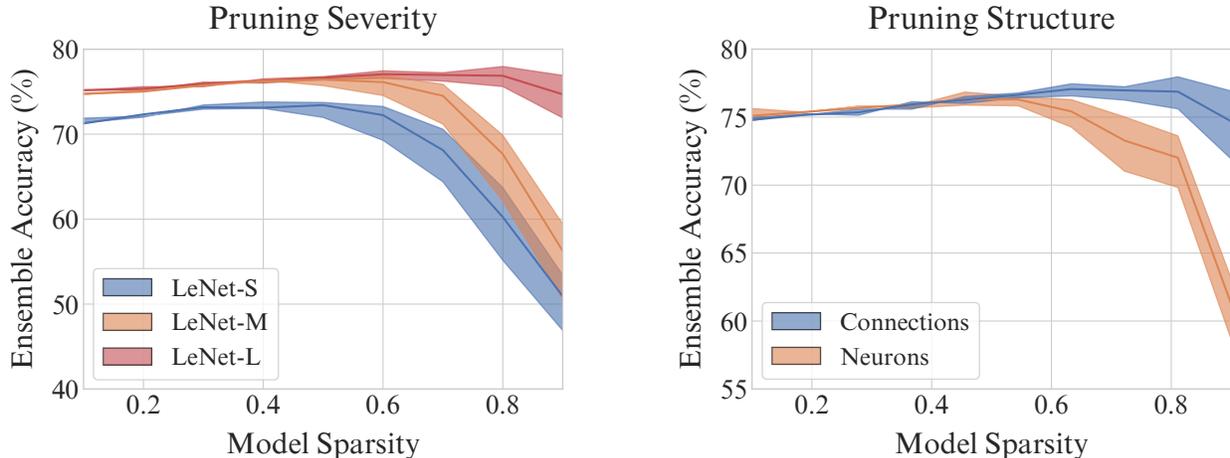
We explore the impact of several hyperparameters that govern how sparse Subnetwork Ensembles are created and optimized. We use three different variations of the LeNet-5 model architecture where we vary the number of neurons in each hidden layer to produce a small, medium and large model that we call *LeNet-S* (253,290 parameters), *LeNet-M* (1,007,306 parameters), and *LeNet-L* (4,017,546 parameters). All ablation experiments use

the Adam optimizer with a learning rate of 0.001 and the benchmark CIFAR-10 and CIFAR-100 datasets, which consist of 60,000 small natural colored images of 32x32 pixels in size. Those 60,000 images are split up into 50,000 training images and 10,000 testing images. CIFAR-10 samples from 10 classes of images, while CIFAR-100 samples from 100 classes of images.

Our first ablation aims to evaluate the effect that the amount of sparsity in the child networks has on both the size of the parent model and the resulting ensemble performance. We train each parent model (*LeNet-S*, *LeNet-M*, and *LeNet-L*) for 10 epochs. An ensemble of 8 sparse subnetwork children are created for each parent by pruning to ablated target sparsity values  $s \in [0.1, 0.9]$ , and subsequently tuned for an additional 3 epochs. Larger models are shown to have more capacity for pruning and learning more efficiently. Ensemble accuracy improves up to a critical region as child networks become more sparse without needing any additional epochs of training. More sparsity in the child networks can yield greater ensemble performance which may suggest better diversity.

Next we explore how structured (neuron-level) and unstructured (connection-level) pruning methods affect ensemble performance at different levels of child network sparsity. We first train the *LeNet-L* parent network for 10 epochs. An ensemble of 8 children is created where each is pruned using either structured or unstructured pruning to ablated target sparsity levels  $s \in [0.1, 0.9]$ , and tuned for an additional 3 epochs. Connection pruning performs much better at higher levels of sparsity compared to neuron pruning. However, we see little difference in performance for child network sparsities up to 50%.

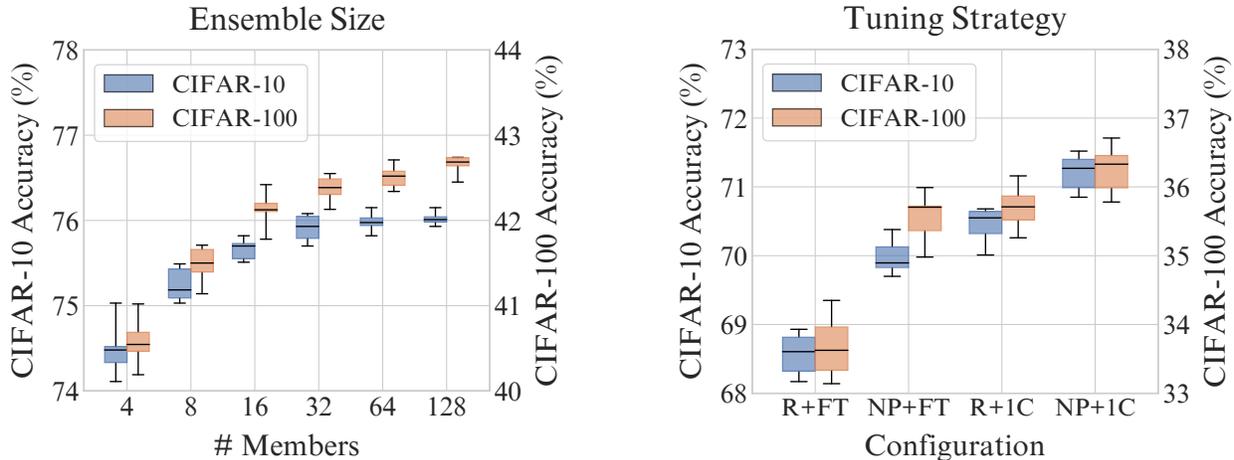
One of the unique advantages of Subnetwork Ensembles is that it is relatively cheap to dynamically generate additional ensemble members. We explore this by taking the *LeNet-M* network and producing up to 128 candidate member networks. We start by training the parent network for 10 epochs. Child networks are created through unstructured pruning with a sparsity target of 0.5. Each child is tuned for one additional epoch. We evaluate ensemble performance over 10 runs as we increase the size of the ensemble from 4 to 128.



**Figure 5.4:** Ablations over pruning sparsity and structure for various model sizes. Each graph reports performance for an ensemble containing 8 models. The left figure reports the effect that pruning sparsity has on models of different sizes. The right figure displays the efficacy of neuron and connection pruning for LeNet-L.

Larger ensembles offer great potential for improved generalization as we see monotonically improved accuracy with reduced variance as the number of ensemble members increases.

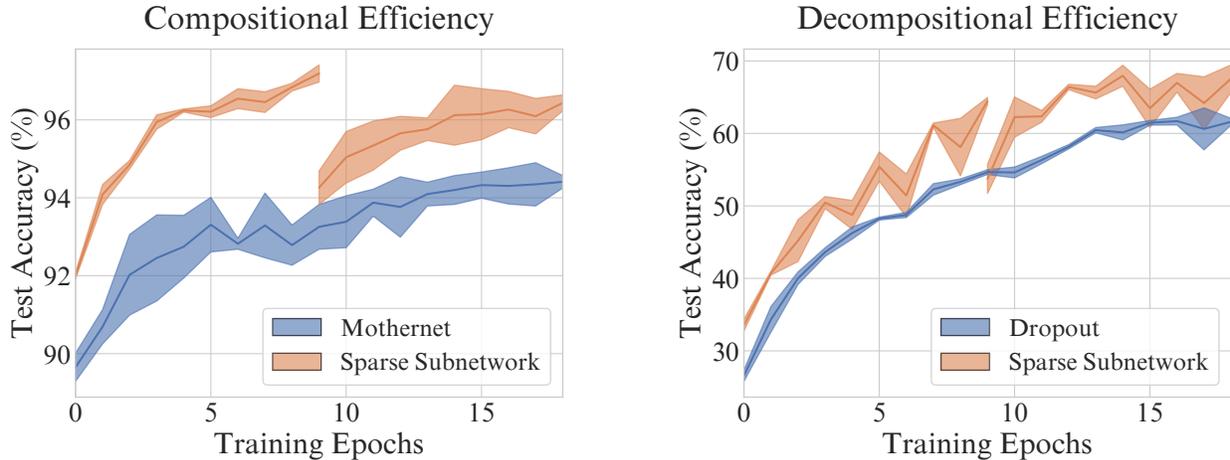
We then aim to evaluate the efficacy of neural partitioning and the one-cycle learning rate policy. We use a more modern training procedure for this experiment where the parent network (*LeNet-L*) is optimized by stochastic gradient descent with Nesterov momentum for 20 epochs. We use an initial learning rate of  $\eta_1 = 0.1$  for 50% of the training budget which decays linearly to  $\eta_2 = 0.001$  at 90% of the training budget. The learning rate is kept constant at  $\eta_2 = 0.001$  for the final 10% of training. Child networks are created with a target sparsity of 0.5 using either random pruning (R) or Neural Partitioning (NP). Child networks are tuned with either a constant fine tuning (FT) learning rate of  $\eta = 0.01$  or with a One-Cycle (1C) schedule that ramps up from  $\eta_1 = 0.001$  to  $\eta_2 = 0.1$  at 10% of tuning, then decaying to  $\eta_3 = 1e - 7$  at the end of 5 epochs. We record results for 10 sample runs of each configuration (R+FT, NP+FT, R+1C, NP+1C), where we observe the best results with the combination of neural partitioning and one-cycle tuning on both CIFAR-10 and CIFAR-100.



**Figure 5.5:** Ablations exploring the effect of ensemble size and various ensemble pruning/tuning strategies. The left graph displays the results for ever increasing ensemble sizes spawned from a single parent network. The right figure contains results for various configurations of random pruning and neural partitioning with constant rate fine tuning and one-cycle fine tuning.

We next explore the early training efficiency of Sparse Subnetwork Ensembles compared to other low-cost methods, MotherNets and Dropout, which are two conceptually competitive approaches [140, 158]. Like Subnetwork Ensembles, MotherNets creates ensembles of child networks from a trained parent network. However, this approach instead uses a small parent network to create larger child networks by adding new neurons and layers around the small shared structure. Meanwhile, Dropout is a popular regularization technique that masks random neurons in the network during training such that a unique subnetwork is active for each forward pass. The purpose of these comparisons is to explore the idea that we get better training dynamics by first training a large parent network and then decomposing it into small networks, rather than using smaller networks throughout training.

The comparison to MotherNets uses the small benchmark dataset MNIST, which is a collection of small handwritten black and white digits [96]. We start by training a MotherNet that contains one hidden layer of 16 neurons. A child network is then created where an additional 16 neurons are added to the original 16 neuron layer of the trained MotherNet. We compare this with a Sparse Subnetwork Ensemble parent that contains a single hidden layer of 64 neurons. A child network is then created by pruning the parent’s hidden layer



**Figure 5.6:** Exploring the early training efficiency of Sparse Subnetwork Ensembles with Mothernets and Dropout. The left figure compares Mothernets on the MNIST dataset with a multilayer perceptron. The right figure compares Dropout on CIFAR-10 with LeNet-L.

to 32 neurons. We train both parents and the resulting children for 10 epochs each using stochastic gradient descent with momentum. The resulting children for both methods are the same size and both approaches train for 20 total epochs. The training graphs of these methods shows that Sparse Subnetwork Ensembles reach higher accuracies than MotherNets in fewer epochs. There is a smooth and consistent training curve with the hatched children of MotherNets, while Subnetwork children undergo a large drop in accuracy before quickly recovering and exceeding the performance of the MotherNet children.

The comparison to Dropout is done on CIFAR-10 with the *LeNet-L* architecture. The Dropout model is identical to the parent network used for Sparse Subnetwork Ensembles, except that it contains dropout layers after every hidden layer with an activation probability of 0.5 during training. We train the parent network for 10 epochs, after which a sparse child network is created by pruning 50% of the neurons and training for an additional 10 epochs. Dropout does offer a smoother training trajectory with reduced variance, however Sparse Subnetwork Ensembles appear to train more efficiently with better accuracy after 20 epochs, despite the final child network containing half the number of neurons of the Dropout model.

### 5.3.2 Landscapes

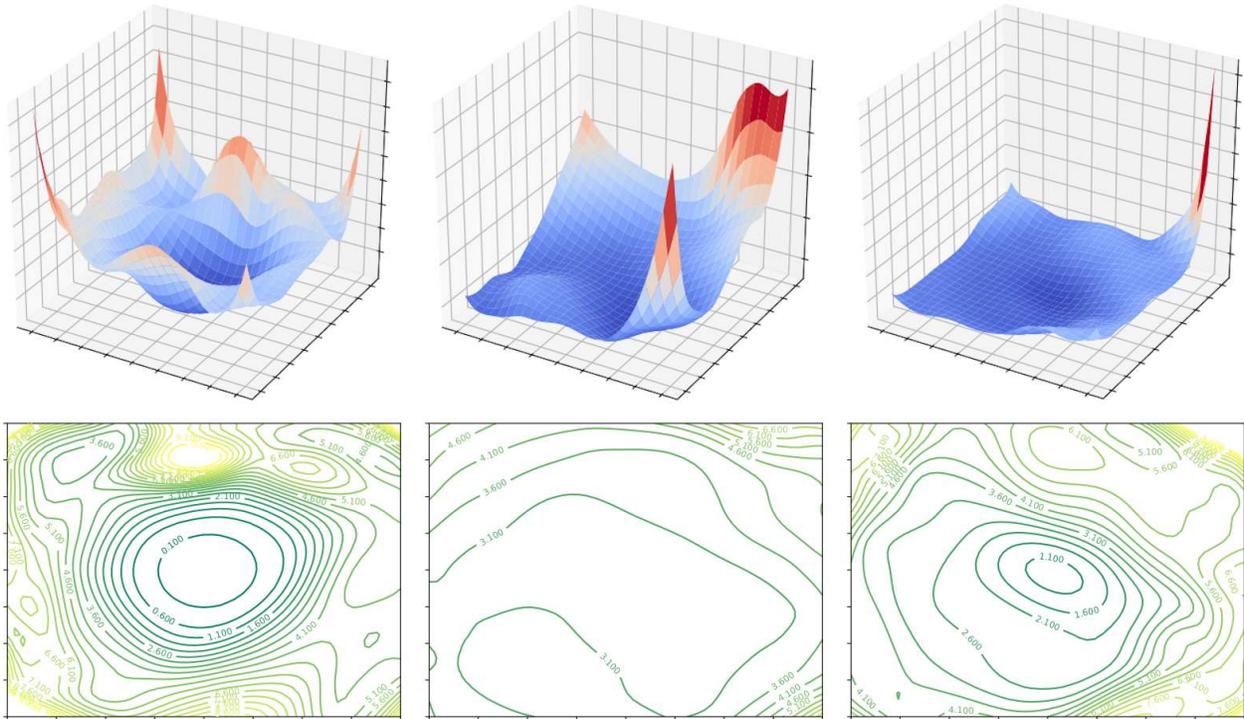
Deep neural networks are known to be vulnerable to adversarial attacks where small manipulations, imperceptible to humans, are applied to input images in order to mislead the network [147]. Previous works have theorized that this is potentially due to excessive over-parameterization which causes spurious correlations to develop that are exploitable. Recent research has demonstrated that pruning can be an effective tool for improving robustness against adversarial attacks [79, 104, 173]. It has also been observed that there is a correlation between robust solutions and flatness of the local optimization landscape [143]. We explore how the loss landscape of a sparse subnetwork evolves as it undergoes both pruning and tuning.

We first train a parent ResNet-56 model for 300 epochs using Stochastic Gradient Descent with momentum on CIFAR-10 with a batch size of 128. The learning rate starts at 0.1 and decays to 0.01 at 150 epochs, 0.001 at 225 epochs and 0.0001 at 275 epochs. We then prune the network by eliminating a random 50% of the weights and continue training the sparse child network for three additional epochs with a constant learning rate 0.001.

The visualizations are created by choosing a center point for the graph with the model weights  $\theta$ . Two orthogonal directional vectors  $\delta$  and  $\eta$  are chosen and the loss landscape is plotted by evaluating the network with weights that are determined by linearly interpolating over the two directional vectors [102].

$$f(\alpha, \beta) = \mathcal{L}(\theta + \alpha\delta + \beta\eta) \tag{5.6}$$

Figure 5.7 displays the results of this visualization. We note the apparent simplification of the loss landscape of the sparse child network as it is tuned. The model appears to settle into an area of the landscape that is shallow and wide, suggesting a good location for robust generalization.



**Figure 5.7:** Visualizations of the loss landscapes of a ResNet-56 model as it undergoes pruning and tuning. The first column is the landscape of the fully trained parent model. The middle figure is the landscape of the model after removing 50% of the parameters. The final figure is the landscape of the pruned model after three additional epochs of fine tuning.

### 5.3.3 Benchmarks

#### Small Budget

Little attention has been given to deep ensemble methods that work well with small training budgets and/or resource constrained environments. While it may seem counter-intuitive to construct ensembles with so few epochs, it is a valuable area of research as insights from these kinds of benchmarks hold the potential to significantly improve training efficiency and reduce computational cost [23, 115].

The experiment presented here limits the total training budget to 16 epochs. We compare the efficiency of Sparse Subnetwork Ensembles to a single independent model, a dropout model, and a traditional bootstrap aggregation (bagged) ensemble. Each method is evaluated with three different model architectures (LeNet-L, ResNet-18, and DenseNet-121) and all models are trained with Adam using a fixed learning rate of 0.001.

**Table 5.1:** Accuracy of various methods for the small training budget benchmark experiment. All ensembles contain 8 models and are trained for a total of 16 epochs. Dropout rate is 50%, Bootstrap Aggregation is trained on 90% subsets.

Model	Method	CIFAR-10	CIFAR-100
LeNet-L	Independent	70.29 $\pm$ 0.21	34.39 $\pm$ 0.18
	Dropout	70.06 $\pm$ 0.32	28.98 $\pm$ 0.33
	Bootstrap Aggregation	70.97 $\pm$ 0.13	32.40 $\pm$ 0.11
	Sparse Subnetwork Ensemble (R)	75.51 $\pm$ 0.06	41.68 $\pm$ 0.14
	Sparse Subnetwork Ensemble (NP)	<b>75.75 <math>\pm</math> 0.06</b>	<b>41.85 <math>\pm</math> 0.13</b>
ResNet-18	Independent	80.99 $\pm$ 0.11	49.26 $\pm$ 0.09
	Dropout	72.97 $\pm$ 0.41	37.69 $\pm$ 0.28
	Bootstrap Aggregation	79.88 $\pm$ 0.05	46.18 $\pm$ 0.06
	Sparse Subnetwork Ensemble (R)	84.61 $\pm$ 0.07	57.08 $\pm$ 0.07
	Sparse Subnetwork Ensemble (NP)	<b>84.68 <math>\pm</math> 0.04</b>	<b>57.26 <math>\pm</math> 0.06</b>
DenseNet-121	Independent	84.14 $\pm$ 0.09	56.62 $\pm$ 0.09
	Dropout	79.59 $\pm$ 0.36	44.02 $\pm$ 0.58
	Bootstrap Aggregation	81.16 $\pm$ 0.04	48.48 $\pm$ 0.06
	Sparse Subnetwork Ensemble (R)	86.26 $\pm$ 0.07	62.02 $\pm$ 0.07
	Sparse Subnetwork Ensemble (NP)	<b>86.35 <math>\pm</math> 0.04</b>	<b>62.09 <math>\pm</math> 0.06</b>

The independent model and the dropout model are identical to each other, except that the dropout model has dropout layers inserted after every hidden layer with a 50% activation probability during training. This parameter is chosen to keep the number of parameters consistent with the child networks generated by Sparse Subnetwork Ensembles. The bagged ensemble contains 8 full size models that are each trained for two epochs on a randomly sampled 90% subset of the training data. Sparse Subnetwork Ensembles trains a single parent network for 8 epochs, then creates 8 children using unstructured pruning with a 50% sparsity target and tunes each of those for an additional epoch. We include results for child networks created with random pruning (R) and with Neural Partitioning (NP), where 4 sets of partitioned pairs of child networks are generated.

Table 5.1 displays the mean accuracy and standard error achieved for each method over 10 runs, except for Sparse Subnetwork Ensembles which have 30 runs. In this small fixed training budget experiment, comparisons show that the Sparse Subnetwork Ensembles are significantly better than all other methods for early training efficiency. While the extent of

**Table 5.2:** Results for ensembles of WideResNet-28-10 models on both CIFAR-10 and CIFAR-100. Methods with \* denote results obtained from [61, 106]. Best low-cost ensemble results are **bold**. cAcc, cNLL, and cECE correspond to corrupted test sets.

Methods (CIFAR-10/WRN-28-10)	Acc $\uparrow$	NLL $\downarrow$	ECE $\downarrow$	cAcc $\uparrow$	cNLL $\downarrow$	cECE $\downarrow$	FLOPs $\downarrow$	Epochs $\downarrow$
Independent Model*	96.0	0.159	0.023	76.1	1.050	0.153	3.6e17	200
Monte Carlo Dropout*	95.9	0.160	0.024	68.8	1.270	0.166	1.00x	200
Snapshot (M=5)	96.3	0.131	0.015	76.0	1.060	0.121	1.00x	200
Fast Geometric (M=12)	96.3	0.126	0.015	75.4	1.157	0.122	1.00x	200
Sparse Subnetworks (M=6)	<b>96.5</b>	<b>0.113</b>	<b>0.005</b>	<b>76.2</b>	<b>0.972</b>	<b>0.081</b>	<b>0.85x</b>	200
TreeNet (M=3)*	95.9	0.258	0.018	75.5	0.969	0.137	1.52x	250
BatchEnsemble (M=4)*	96.2	0.143	0.021	77.5	1.020	0.129	4.40x	250
MIMO (M=3)*	96.4	0.123	0.010	76.6	0.927	0.112	4.00x	250
EDST (M=7)*	96.4	0.127	0.012	76.7	0.880	0.100	0.57x	850
DST (M=3)*	96.4	0.124	0.011	77.6	0.840	0.090	1.01x	750
Dense Ensemble (M=4)*	96.6	0.114	0.010	77.9	0.810	0.087	1.00x	800
Methods (CIFAR-100/WRN-28-10)	Acc $\uparrow$	NLL $\downarrow$	ECE $\downarrow$	cAcc $\uparrow$	cNLL $\downarrow$	cECE $\downarrow$	FLOPs $\downarrow$	Epochs $\downarrow$
Independent Model*	79.8	0.875	0.086	51.4	2.700	0.239	3.6e17	200
Monte Carlo Dropout*	79.6	0.830	0.050	42.6	2.900	0.202	1.00x	200
Snapshot (M=5)	82.1	0.661	0.040	52.2	2.595	0.145	1.00x	200
Fast Geometric (M=12)	82.3	0.653	0.038	51.7	2.638	0.137	1.00x	200
Sparse Subnetworks (M=6)	<b>82.7</b>	<b>0.634</b>	<b>0.013</b>	<b>52.7</b>	<b>2.487</b>	<b>0.131</b>	<b>0.85x</b>	200
TreeNet (M=3)*	80.8	0.777	0.047	53.5	2.295	0.176	1.52x	250
BatchEnsemble (M=4)*	81.5	0.740	0.056	54.1	2.490	0.191	4.40x	250
MIMO (M=3)*	82.0	0.690	0.022	53.7	2.284	0.129	4.00x	250
EDST (M=7)*	82.6	0.653	0.036	52.7	2.410	0.170	0.57x	850
DST (M=3)*	82.8	0.633	0.026	54.3	2.280	0.140	1.01x	750
Dense Ensemble (M=4)*	82.7	0.666	0.021	54.1	2.270	0.138	1.00x	800

hyperparameter tuning is limited for all methods, this experiment is designed to create an initial baseline for small training budgets that can be further explored in future work.

## Large Budget

We next evaluate Sparse Subnetwork Ensembles in the context of a standardized benchmark in low-cost ensemble learning literature. This experiment uses a larger training budget of 200 epochs and the WideResNet-28-10 model architecture [176]. We use the same optimizer configuration and hyperparameter settings reported in other low-cost ensemble studies [50, 70, 158]. We compare Sparse Subnetwork Ensembles with published results of several of the top low-cost ensemble methods including: TreeNets, Snapshot Ensembles, Fast Geometric Ensembles, BatchEnsembles, FreeTickets, and MIMO [50, 61, 70, 100, 106, 159].

All methods train the same WideResNet-28-10 model architecture using stochastic gradient descent with Nesterov momentum  $\mu = 0.9$  and weight decay  $\gamma = 0.0005$  [144]. We

use a batch size of 128 for training and use random crop, random horizontal flip, and mean standard scaling data augmentations for all approaches [50, 61, 70, 106]. The parent network is trained with step-wise decaying learning rate schedule where an initial learning rate of  $\eta_1 = 0.1$  is used for 50% of the training budget and decays linearly to  $\eta_2 = 0.001$  at 90% of the training budget. The learning rate is kept constant at  $\eta_2 = 0.001$  for the final 10% of training.

Table 5.2 includes results for all compared methods. The Independent Model is a single WideResNet-28-10 model trained for 200 Epochs. The Dropout Model includes dropout layers between convolutional layers in the residual blocks at a rate of 30% [176]. Snapshot Ensembles use a cosine annealing learning rate with an initial learning rate  $\eta = 0.1$  for a cycle length of 40 epochs [70]. Fast Geometric Ensembles use a pre-training routine for 156 epochs. A curve finding algorithm then runs for 22 epochs with a cycle length of 4, each starting from checkpoints at epoch 120 and 156. TreeNets, BatchEnsemble and MIMO are all trained for 250 epochs [61, 100, 159]. FreeTickets introduces several configurations for building ensembles. We include their two best configurations for Dynamic Sparse Training (DST, M=3, S=0.8) and Efficient Dynamic Sparse Training (EDST, M=7, S=0.9).

The Sparse Subnetwork ensemble size and training schedule is the same as those used in MotherNets and Fast Geometric Ensembles [50, 158]. Sparse Subnetwork Ensembles train a single parent network for 140 epochs. Six children are created with Neural Partitioning and tuned with a one-cycle learning rate for 10 epochs. The tuning schedule starts at  $\eta_1 = 0.001$ , increases to  $\eta_2 = 0.1$  at 1 epoch and then decays to  $\eta_3 = 1e - 7$  using cosine annealing for the final 9 epochs. We additionally explored the effect of different partition sizes for child networks in the ensemble where  $np = 1/2, 1/3, 1/6$ . Table 5.4 reports these results where we see a very minor degradation in accuracy moving from children with 50% sparsity to 66% sparsity.

We report the mean accuracy (Acc), negative log likelihood (NLL), and expected calibration error (ECE) over 10 runs on both CIFAR-10 and CIFAR-100 along with their corrupted

**Table 5.3:** Results for Sparse Subnetwork Ensembles on CIFAR-10 with different child network partition sizes. The parent network is trained for 140 epochs. 6 children are then created through neural partitioning and each is trained for an additional 10 epochs. Total Params counts all parameters in the ensemble.

Partition Size	Child Acc. $\uparrow$	Ensemble Acc. $\uparrow$	NLL $\downarrow$	ECE $\downarrow$	FLOPs $\downarrow$	Total Params $\downarrow$
1/2	95.5	96.5	0.114	0.005	0.85x	109.5M
1/3	95.3	96.4	0.115	0.004	0.80x	73.0M
1/6	92.9	95.1	0.143	0.013	0.75x	36.5M

variants [63, 118]. We report the total number of floating point operations (FLOPs) and epochs used for training each method. We organize our tables into two groups based on training cost. The first group consists of low-cost training methods that take approximately as long as a single network would take to train. The second group of methods use either significantly more epochs or compute per epoch to achieve comparable performance. MIMO and BatchEnsemble both make use of batch repetition to train on more data while keeping the number of epochs low. FreeTickets (DST and EDST) use very sparse networks to keep FLOP counts low while using many more training epochs.

Sparse Subnetwork Ensembles outperform all low-cost ensemble methods and are competitive with methods that train for significantly longer. Our approach produces well calibrated, robust and diverse ensembles with excellent performance on in distribution and out of distribution datasets.

## 5.4 Discussion

Sparse Subnetwork Ensembles prove to be a highly effective method for constructing diverse, accurate, and efficient ensembles of child networks. We do this by first training a large parent network. We then create child networks by sampling and pruning random subnetworks in the parent to produce a collection of sparse child networks with unique connective structures. We then briefly train each child network for a small number of epochs where they quickly converge to new local optima.

Despite inheriting parameters from the same parent, child networks quickly converge to diverse solutions thanks to their inherited parameters and unique network topologies. We further explore various ways to encourage diversity among child networks through neural partitioning and one-cycle tuning. Neural partitioning is used to generate sets of child networks without any parametric overlap and one-cycle tuning uses a phasic learning rate schedule where large learning rates encourage children to move further apart in weight space before converging to local solutions.

The Sparse Subnetwork Ensemble algorithm is flexible and robust. We conduct several ablation studies where we construct ensembles with various network architectures, pruning strategies, and training configurations. We also visualize the loss landscape of a network throughout the lifecycle of this algorithm, where we see convergence to wide and flat optima that suggest regions of robust generalization.

This work showcases an important concept in Subnetwork Ensembles where it is found to be more efficient to first train a large parent network to a good solution before decomposing it into smaller child networks. The large parent network has better optimization dynamics when trained from scratch, while the final state contains large amounts of unrealized capacity. Sparse Subnetwork children leverage that redundancy to produce accurate and diverse representations with little additional computation.

With the benchmark experiments introduced, Sparse Subnetwork Ensembles outperform other low-cost ensembling methods on benchmark image classification datasets in both low training budget and high training budget regimes. Our approach achieves comparable accuracy to ensembles that are trained on 4 to 5 times more data. Sparse Subnetwork Ensembles not only improve upon the training-time/accuracy trade-off, but also significantly reduces memory and computational cost.

# Chapter 6

## Stochastic Subnetwork Ensembles

### 6.1 Introduction

A critical component of Sparse Subnetwork Ensembles is the creation and tuning of child networks. The standard approach explored in the previous chapter involves pruning the whole subnetwork in one-shot, followed by a tuning phase. Meanwhile, it's long been known that iterative pruning, which consists of alternating pruning/tuning cycles, is one of the most effective techniques to improve performance of extremely sparse networks [6].

The key insight to iterative pruning is that removing too many parameters at once can lead to a drastic performance collapse as subnetworks get stuck in lower performing regions of the optimization landscape. Pruning a small number of parameters over several epochs allows the model to better adapt to the changing network structure. However, these iterative methods can still result in subnetworks that are prone to overfitting local optimization regions.

We introduce a novel regularization technique for fine tuning these subnetworks by leveraging stochasticity for the network structure. Rather than using fixed subnetworks and discrete pruning operations, we instead represent subnetworks with probability matrices that determine how likely it is that a parameter is retained on any given forward pass. The probability matrices are then adjusted during training through the use of dynamic annealing schedules which gradually evolves the network from a stochastic network towards a deterministic subnetwork structure over several epochs. The probabilistic inclusion of extra parameters early in the tuning process allows for gradient information to bleed through into the target subnetwork, encouraging robust adaptation and avoiding the drastic performance collapse observed with one-shot pruning methods. This approach to representing subnet-

works with stochastic masks is highly flexible and offers a lot of fine grained control over the tuning process.

We conduct a large scale ablation study to explore the dynamics of several annealing hyperparameters and their effects on convergence. These include the initial stochasticity, number of annealing epochs, amount of sparsity, learning rate schedules, and parameter selection strategies. Our experiments demonstrate significant improvements over both one-shot and iterative pruning methods with especially large improvements in highly sparse subnetworks (95-98%). We then explore how this technique can be leveraged to tune child networks in Sparse Subnetwork Ensembles, where we observe improved accuracy and generalization on benchmark classification tasks.

## 6.2 Implementation

In pruning literature, sparse network structures are generally static and represented with binary bit masks. We propose a model of representing neural subnetworks with probabilistic masks, where each parameter is assigned a score that determines how likely it is that the parameter will be retained on any given forward pass. This introduces stochasticity into the subnetwork sampling process, which can act as a form of implicit regularization analagous to a reverse dropout, where parameters that would have been pruned have a chance to activate. This technique encourages exploration of a larger space of subnetworks during the fine-tuning phase, preventing it from becoming overly reliant on a single fixed topological configuration, resulting in more robust and generalized subnetworks.

Consider a weight matrix  $W \in \mathbb{R}^{m \times n}$  representing the weights of a particular layer in a neural network. We introduce a probability matrix  $P \in \mathbb{R}^{m \times n}$  containing scores that represent the probability that a parameter  $W_{ij}$  will be masked on any given forward pass. The subnetwork mask  $M \in \{0, 1\}^{m \times n}$  is determined with a Bernoulli realization of the probability matrix  $P$ .

The ideas of neural partitioning naturally extend to probabilistic subnetwork representation where the opposed probability matrix  $P' = 1 - P$  can be used to create a pair of subnetworks. More partitions can be created so long as the probabilities for each parameter sum to one.

$$(P, P') = \left( \begin{bmatrix} 0.6683 & 0.3923 & 0.9669 \\ 0.8343 & 0.6885 & 0.8940 \\ 0.2613 & 0.2249 & 0.9972 \end{bmatrix}, \begin{bmatrix} 0.3317 & 0.6077 & 0.0331 \\ 0.1657 & 0.3115 & 0.1060 \\ 0.7387 & 0.7751 & 0.0028 \end{bmatrix} \right) \quad (6.1)$$

We then anneal these probability values over some number of epochs such that subnetworks are slowly revealed throughout fine-tuning. This is done by introducing an annealing schedule, where the probability values for each parameter slowly move towards 0 or 1 depending on the target subnetwork sparsity. At the beginning of the training process, a high level of stochasticity is desirable as it encourages exploration of the weight space which may help to prevent overfitting and search for more effective subnetwork configurations. As training progresses, the stochasticity should gradually decrease to allow for stable convergence. The annealing function is arbitrary, with some popular examples being the linear, cosine, and exponential decay.

$$y_{lin}(t) = \tau_{max} - (\tau_{max} - \tau_{min}) \left( \frac{t}{T_{max}} \right) \quad (6.2)$$

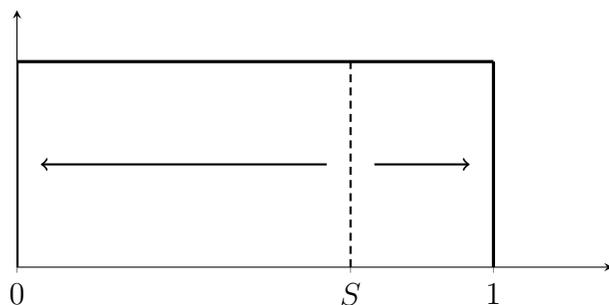
$$y_{cos}(t) = \tau_{min} + \frac{1}{2}(\tau_{max} - \tau_{min}) \left( 1 + \cos \left( \frac{t}{T_{max}} \right) \right) \quad (6.3)$$

$$y_{exp}(t) = \tau_{min} + (\tau_{max} - \tau_{min}) e^{-\frac{t}{T_{max}}k} \quad (6.4)$$

### 6.2.1 Random Annealing

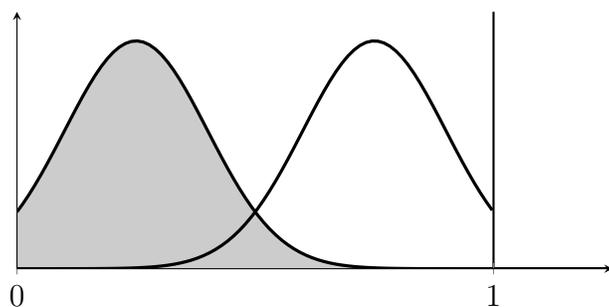
The probability matrix can be generated in arbitrary ways. For example, a uniform distribution  $P \sim U([0, 1])^{m \times n}$  can be used to randomly assign probabilities to each parameter. All parameters with a value less than the sparsity target will then anneal towards 0 while

parameters with a probability value greater than the sparsity target will anneal towards 1. This implementation with a uniform distribution results in a mean parameter activation of approximately 50% at the beginning of tuning regardless of the target subnetwork sparsity.



Mixtures of Gaussians can be used to assign interesting distributions that offer more fine grained control over network structure. For example, assume that a binary matrix  $X \in \{0, 1\}^{m \times n}$  is randomly generated and used to index into a probability matrix  $P \in \mathbb{R}^{m \times n}$ . Using this index matrix  $X$ , we can sample from Gaussian distributions with different means and variances.

$$P = \begin{cases} P_{ij} \sim \mathcal{N}(\mu_1, \sigma_1^2), & \text{if } X_{ij} = 0 \\ P_{ij} \sim \mathcal{N}(\mu_2, \sigma_2^2), & \text{if } X_{ij} = 1 \end{cases} \quad (6.5)$$



This approach to constructing multi-modal distributions can enable many distinct formulations of stochastic subnetworks. Future work may find natural applications to multi-task learning, where certain groups of parameters can be strongly correlated with each other while allowing for overlap with other task specific subnetworks. This stochastic overlap may

encourage shared portions of the network to learn generalized features, avoiding the problem in typical network-splitting approaches where task specific subnetworks become increasingly narrow and lead to degraded performance.

### 6.2.2 Temperature Annealing

Temperature scaling may be applied to binary matrices in order to allow for an even application of stochasticity, reducing some variance relative to random annealing. Assume that a binary matrix  $X \in \{0, 1\}^{m \times n}$  is generated at random. A temperature scaling constant  $\tau$  is introduced such that the values in  $X$  with a 1 are decayed by  $\tau$  and the values with 0 are increased by  $\tau$ . The mask is then determined on every forward pass with an altered probability according to  $\tau$ . Altering the initial value of  $\tau$  allows for a much more controlled approach to noise injection during the early phases of tuning, which can be desirable when the number of total tuning epochs is limited.

A variation of temperature scaling, where  $\tau$  is only applied to parameters that are not a part of the target subnetwork, acts as a highly effective form of regularization analagous to reverse Dropout. That is, the subnetwork is always active, but parameters that would have been pruned away now have a chance to pop back in during tuning. This formulation ensures that the target subnetwork will be optimized on every training step, which can enable faster convergence. Allowing other parameters to become active allows for gradient information to contribute to optimization of the target subnetwork which can help to encourage avoidance of local minima.

$$P = \begin{bmatrix} 1 & 1 & 0 + \tau & 1 \\ 0 + \tau & 1 & 0 + \tau & 0 + \tau \\ 1 & 0 + \tau & 0 + \tau & 1 \\ 0 + \tau & 0 + \tau & 1 & 1 \end{bmatrix} \quad (6.6)$$

---

**Algorithm 3:** Stochastic Subnetwork Ensembles

---

**Input:**  $N$ , the number of ensemble members  
**Input:**  $(X_{train}, X_{test})$ , the training, validation, and test data  
**Input:**  $(t_{parent}, t_{child}, t_{anneal})$ , the number of parent, child, and annealing epochs  
**Input:**  $(\eta_{parent}, \eta_{child})$ , the parent and child optimizer hyperparameters  
**Input:**  $\phi(\rho)$ , the subnetwork sampling process with target sparsity  $\rho$   
**Input:**  $\Theta(M, \tau, e, t)$ , the temperature scaling function applied to a mask  $M$ , with initial temperature  $\tau$ , for current epoch  $e$ , over  $t$  total annealing epochs

$O = \text{sgd}(\eta_{parent})$

$F = \text{initialize}()$

$F.w = \text{train}(F, O, X_{train}, t_{parent})$

$ensemble = []$

**for**  $i$  in 1 to  $N$  **do**

$f_i = \text{initialize}()$

$f_i.w = F.w$

**for**  $j$  in  $f_i.layers$  **do**

$M_j \sim \phi(\rho)^{f_{i,j}.w}$

**end**

$O = \text{sgd}(\eta_{child})$

**for**  $e$  in  $t_{child}$  **do**

$P \leftarrow \Theta(M, \tau, e, t_{anneal})$

**for**  $(x, y)$  in  $X_{train}$  **do**

$\hat{f}_i.w \leftarrow f_i.w \circ \text{Bernoulli}(P)$

$f_i.w = \text{train}(\hat{f}_i, O, (x, y))$

**end**

**end**

$ensemble \leftarrow f_i$

**end**

**for**  $(x, y)$  in  $X_{test}$  **do**

$outputs = [F(x), c(x) \text{ for } c \text{ in } ensemble]$

$predictions = \text{softmax}(\text{mean}(outputs))$

$loss = \text{nll}(predictions, y)$

$accuracy = \text{acc}(predictions, y)$

**end**

---

## 6.3 Experiments

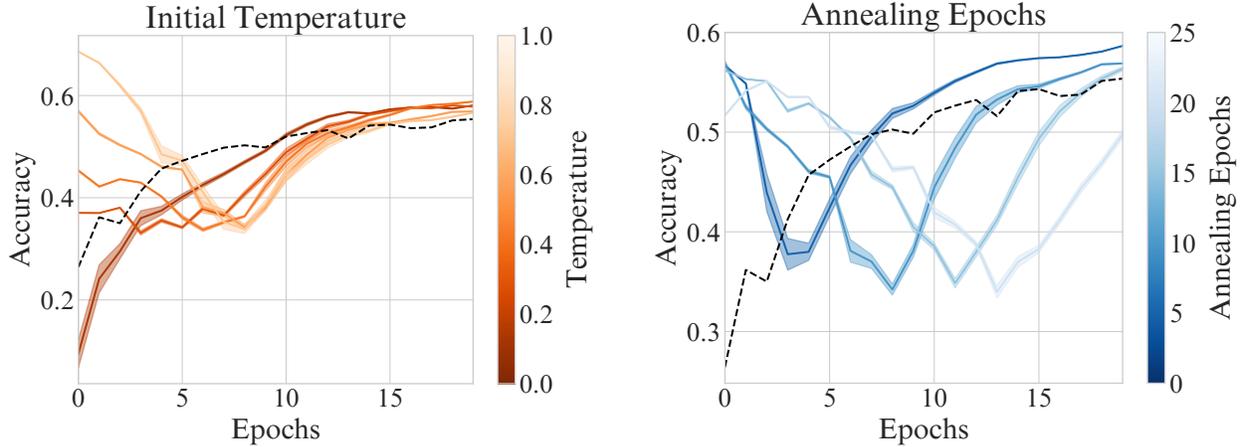
### 6.3.1 Ablations

We begin with an exploration of several stochastic subnetwork annealing configurations with the goal of investigating how different hyperparameters impact the efficiency and quality of subnetwork convergence, compared to established one-shot and iterative pruning techniques.

All ablations use the same ResNet-18 trained for 100 epochs using a standardized optimization configuration [62]. We use PyTorch’s Stochastic Gradient Descent optimizer with an initial learning rate of 0.1 and Nesterov momentum of 0.9 [144]. After 50 epochs, the learning rate is decayed to 0.01 and again to 0.001 for the final 10 epochs. We use the benchmark CIFAR datasets with standard data augmentations including random crop, random horizontal flip, and mean standard normalization.

Pruning is done in a layerwise unstructured fashion, after which each subnetwork is tuned for an additional 20 epochs. We experiment with both a constant learning rate of 0.01 and a one-cycle policy with a max learning rate of 0.1. We explore the results for each configuration with different levels of final subnetwork sparsity  $\rho \in [0.5, 0.7, 0.9, 0.95, 0.98]$ . We additionally include results for subnetworks created through L1 unstructured pruning. In this case, parameters with the smallest magnitudes at each layer are pruned.

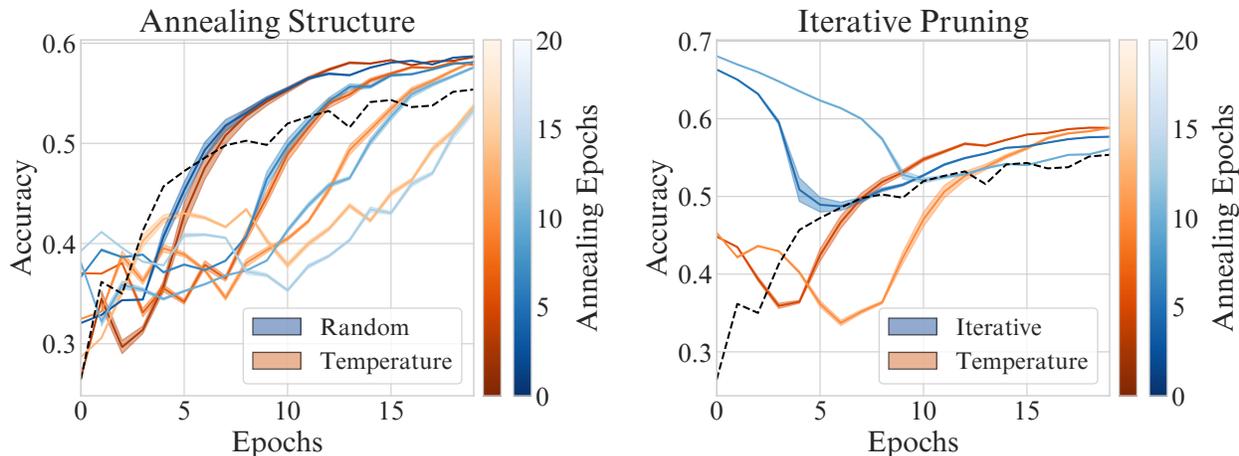
The one-shot baseline prunes the parent network to the target sparsity before we start tuning. Iterative pruning calculates the number of parameters to remove at the beginning of each epoch over some number of pruning epochs  $\varepsilon \in [5, 10, 15, 20]$  such that the final target sparsity is hit. Random Annealing uses a probability matrix that is generated according to a random uniform distribution, where each parameter is assigned a value between 0 and 1. Parameters with a value less than the target sparsity value are linearly annealed to 0 and values greater than the target sparsity are linearly annealed to 1, over some number of annealing epochs  $\varepsilon \in [5, 10, 15, 20]$ . Temperature Annealing uses a randomly generated binary bitmask according to the target sparsity. All parameters with a value of 0 are modified



**Figure 6.1:** Ablations exploring how the initial temperature and the number of annealing epochs affect convergence behavior.

**Table 6.1:** Comparison between various baseline approaches when used with random/magnitude pruning on CIFAR-100. We report the best accuracy for each method at various levels of target sparsity when tuned with both a constant learning rate and a one-cycle learning rate schedule. Our annealing strategies consistently outperform established one-shot and iterative pruning methods.

Random Pruning/Constant LR	50%	70%	90%	95%	98%
One-Shot Baseline	0.640	0.634	0.553	0.490	0.401
Iterative Pruning	0.655	0.642	0.578	0.502	0.422
Random Annealing	0.655	0.636	0.591	0.546	0.457
Temperature Annealing	<b>0.658</b>	<b>0.643</b>	<b>0.595</b>	<b>0.554</b>	<b>0.463</b>
Random Pruning/One-Cycle LR	50%	70%	90%	95%	98%
One-Shot Baseline	0.702	0.685	0.620	0.573	0.493
Iterative Pruning	0.702	0.692	0.639	0.573	0.508
Random Annealing	0.703	0.692	0.646	0.600	0.527
Temperature Annealing	<b>0.706</b>	<b>0.696</b>	<b>0.649</b>	<b>0.602</b>	<b>0.533</b>
Magnitude Pruning/Constant LR	50%	70%	90%	95%	98%
One-Shot Baseline	0.673	0.668	0.661	0.635	0.542
Iterative Pruning	0.682	0.680	0.660	0.633	0.564
Temperature Annealing	<b>0.685</b>	<b>0.682</b>	<b>0.663</b>	<b>0.635</b>	<b>0.566</b>
Magnitude Pruning/One-Cycle LR	50%	70%	90%	95%	98%
One-Shot Baseline	0.709	0.701	0.697	0.675	0.604
Iterative Pruning	0.711	0.711	0.700	0.675	0.606
Temperature Annealing	<b>0.714</b>	<b>0.714</b>	<b>0.702</b>	<b>0.678</b>	<b>0.615</b>

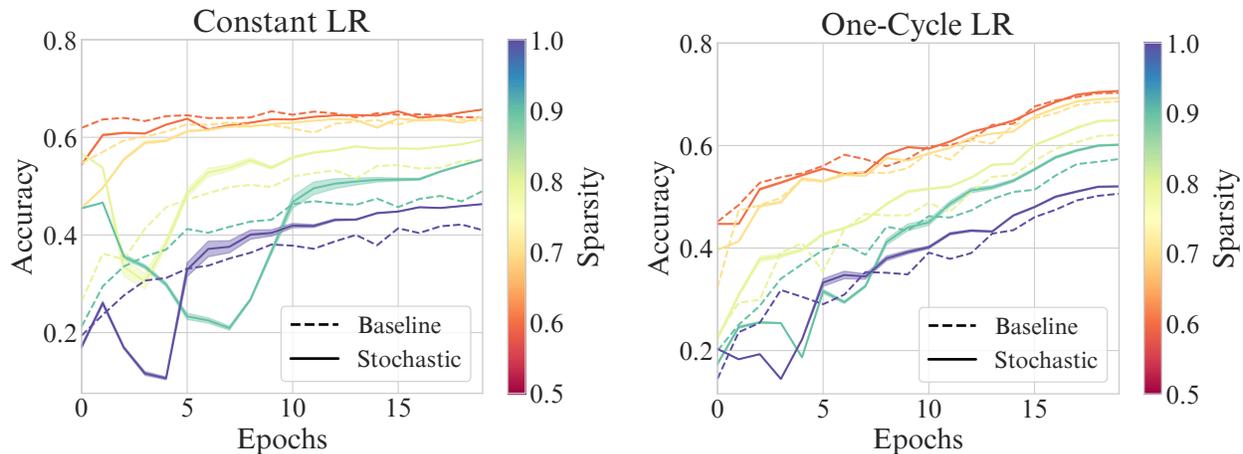


**Figure 6.2:** Ablations over the number of annealing epochs for random annealing vs temperature annealing and temperature annealing vs iterative pruning.

to an initial temperature value of  $\tau \in [0.2, 0.4, 0.6, 0.8, 1.0]$ . Those values are then cosine annealed to 0 over some number of annealing epochs  $\varepsilon \in [5, 10, 15, 20]$ .

Table 6.1 includes the accuracies for the best configurations of each method on CIFAR-100. We see consistent improvement with both of our stochastic annealing methods over the baseline one-shot and iterative pruning techniques across all sparsities and with both a constant learning rate and a one-cycle learning rate policy. As networks become more sparse, the benefits from our annealing approaches become more significant, with a 6% and 4% improvement at 98% sparsity over the one-shot baseline with a constant and one-cycle rate schedule respectively. We also observed improved performance with magnitude pruning, however the differences were smaller at 2% and 1% improvement.

Figure 6.1 includes an exploration of the initial temperature and the number of annealing epochs. The hyperparameter with the most significant impact on all pruning methodologies was the number of epochs that were pruned or annealed over. We saw best results for all subnetwork sparsities with  $e = 5$  or  $e = 10$  annealing epochs. It’s important for the final subnetwork topology to be established for a sufficient number of epochs to allow for optimal convergence behavior. This pattern holds for both constant and one-cycle learning rate schedules. The initial temperature has a smaller impact on performance than the number of annealing epochs. A small value of  $\tau$  means that the target subnetwork is always active



**Figure 6.3:** The rightmost graph displays the best performing annealing configuration for each sparsity and plots it against the one-shot pruning baseline. Stochastic annealing outperformed the other methods for each sparsity, with more drastic improvements appearing in extremely sparse networks.

and other parameters have a small chance to turn on, while a high value of  $\tau$  means that the target subnetwork is always active and other parameters have a high chance to turn on. We saw best results when  $\tau$  was in the 0.4 to 0.6 range. This corresponds to a higher state of entropy regarding network structure which results in a stronger regularization effect for those initial training examples.

Figure 6.2 includes an exploration of random annealing vs temperature annealing and iterative pruning vs temperature annealing. Despite the additional variance associated with random annealing, we found that the performance was very good and nearly approached that of temperature annealing. Temperature annealing consistently outperformed iterative pruning. The graph displays results for annealing epochs  $\varepsilon \in [5, 10]$  and a target subnetwork sparsity of 0.9. While the early accuracy of temperature annealing appears worse, it quickly surpasses iterative pruning as the annealing phase ends.

Figure 6.3 displays the test trajectories for the best models over all sparsities with both a constant learning rate schedule and a one-cycle learning rate schedule. We compare the results of temperature annealing with that of the one-shot baseline. Temperature Annealing results in better performance for both learning rate schedules at all levels of sparsity that we

tested  $\rho \in [0.5, 0.7, 0.9, 0.95, 0.98]$ . The benefits of temperature annealing are most noticeable at higher levels of sparsity with the constant learning rate schedule. The one-cycle schedule does result in much better generalization and a reduced accuracy gap between temperature annealing and the baseline.

### 6.3.2 Benchmarks

We next evaluate the efficacy of Stochastic Subnetwork Annealing in the context Sparse Subnetwork Ensembles. We include results for these ensembles on corrupted versions of the CIFAR datasets as well, in order to test robustness on out-of-distribution images. These corrupted datasets are generated by adding 20 different kinds of image corruptions (Gaussian noise, snow, blur, pixelation, etc.) at five different levels of severity to the original test sets. The total number of images in each of these additional sets is 1,000,000.

We take the training configuration, ensemble size and parameter settings directly from the large scale benchmark experiment described in the Sparse Subnetwork Ensemble chapter.

All methods compared use WideResNet-28-10 and Stochastic Gradient Descent with Nesterov momentum  $\mu = 0.9$  and weight decay  $\gamma = 0.0005$  [144]. We use a batch size of 128 for training and use random crop, random horizontal flip, and mean standard normalization data augmentations for all approaches [50, 61, 70, 106]. The parent learning rate uses a step-wise decay schedule. An initial learning rate of  $\eta_1 = 0.1$  is used for 50% of the training budget which decays linearly to  $\eta_2 = 0.001$  at 90% of the training budget. The learning rate is kept constant at  $\eta_2 = 0.001$  for the final 10% of training.

We train a single parent network for 140 epochs. Six children are then created by randomly pruning 50% of the connections in the parent network. Neural partitioning is implemented where pairs of children are generated with opposite sets of inherited parameters. Each child is tuned with a one-cycle learning rate for 10 epochs. The tuning schedule starts at  $\eta_1 = 0.001$ , increases to  $\eta_2 = 0.1$  at 1 epoch and then decays to  $\eta_3 = 1e - 7$  using cosine annealing for the final 9 epochs.

**Table 6.2:** Results for ensembles of WideResNet-28-10 models on both CIFAR-10 and CIFAR-100. Methods with \* denote results obtained from [61, 106]. Best low-cost ensemble results are **bold**. cAcc, cNLL, and cECE correspond to corrupted test sets.

Methods (CIFAR-10/WRN-28-10)	Acc $\uparrow$	NLL $\downarrow$	ECE $\downarrow$	cAcc $\uparrow$	cNLL $\downarrow$	cECE $\downarrow$	FLOPs $\downarrow$	Epochs $\downarrow$
Independent Model*	96.0	0.159	0.023	76.1	1.050	0.153	3.6e17	200
Monte Carlo Dropout*	95.9	0.160	0.024	68.8	1.270	0.166	1.00x	200
Snapshot (M=5)	96.3	0.131	0.015	76.0	1.060	0.121	1.00x	200
Fast Geometric (M=12)	96.3	0.126	0.015	75.4	1.157	0.122	1.00x	200
Sparse Subnetworks (M=6)	96.5	0.113	0.005	76.2	0.972	0.081	0.85x	200
Stochastic Annealing (M=6)	<b>96.7</b>	<b>0.110</b>	<b>0.005</b>	<b>76.3</b>	<b>0.968</b>	<b>0.079</b>	<b>0.85x</b>	<b>200</b>
TreeNet (M=3)*	95.9	0.258	0.018	75.5	0.969	0.137	1.52x	250
BatchEnsemble (M=4)*	96.2	0.143	0.021	77.5	1.020	0.129	4.40x	250
Multi-Input Multi-Output (M=3)*	96.4	0.123	0.010	76.6	0.927	0.112	4.00x	250
FreeTickets (EDST) (M=7)*	96.4	0.127	0.012	76.7	0.880	0.100	0.57x	850
FreeTickets (DST) (M=3)*	96.4	0.124	0.011	77.6	0.840	0.090	1.01x	750
Dense Ensemble (M=4)*	96.6	0.114	0.010	77.9	0.810	0.087	1.00x	800

Methods (CIFAR-100/WRN-28-10)	Acc $\uparrow$	NLL $\downarrow$	ECE $\downarrow$	cAcc $\uparrow$	cNLL $\downarrow$	cECE $\downarrow$	FLOPs $\downarrow$	Epochs $\downarrow$
Independent Model*	79.8	0.875	0.086	51.4	2.700	0.239	3.6e17	200
Monte Carlo Dropout*	79.6	0.830	0.050	42.6	2.900	0.202	1.00x	200
Snapshot (M=5)	82.1	0.661	0.040	52.2	2.595	0.145	1.00x	200
Fast Geometric (M=12)	82.3	0.653	0.038	51.7	2.638	0.137	1.00x	200
Sparse Subnetworks (M=6)	82.7	0.634	0.013	52.7	2.487	0.131	0.85x	200
Stochastic Annealing (M=6)	<b>83.1</b>	<b>0.633</b>	<b>0.010</b>	<b>52.8</b>	<b>2.440</b>	<b>0.131</b>	<b>0.85x</b>	<b>200</b>
TreeNet (M=3)*	80.8	0.777	0.047	53.5	2.295	0.176	1.52x	250
BatchEnsemble (M=4)*	81.5	0.740	0.056	54.1	2.490	0.191	4.40x	250
Multi-Input Multi-Output (M=3)*	82.0	0.690	0.022	53.7	2.284	0.129	4.00x	250
FreeTickets (EDST) (M=7)*	82.6	0.653	0.036	52.7	2.410	0.170	0.57x	850
FreeTickets (DST) (M=3)*	82.8	0.633	0.026	54.3	2.280	0.140	1.01x	750
Dense Ensemble (M=4)*	82.7	0.666	0.021	54.1	2.270	0.138	1.00x	800

We implement Stochastic Annealing when tuning the child networks by using the procedure illustrated in Section 6.3.1 for Temperature Annealing. We initialize the binary subnetwork masks with a target sparsity of 0.5. We then modify the masks such that the parameters with a value of zero are modified to have an initial temperature of  $\tau = 0.5$ . That value is decayed to  $\tau = 0$  over 3 annealing epochs using a cosine decay. We compare the results of these Stochastic Subnetwork Ensembles tuned with Temperature Annealing to the same low-cost methods described in the Chapter 5 (Sparse Subnetwork Ensembles).

We report the mean accuracy (Acc), negative log likelihood (NLL), and expected calibration error (ECE) over 10 runs on both CIFAR-10 and CIFAR-100 along with their corrupted variants [63, 118]. We also report the total number of floating point operations (FLOPs) and epochs used for training each method. We organize our tables into two groups based on training cost. The first group consists of low-cost training methods that take approximately

as long as a single network would take to train. The second group of methods use either significantly more epochs or compute per epoch to achieve comparable performance.

## 6.4 Discussion

Stochastic Subnetwork Annealing offers a novel approach to tuning pruned models by representing subnetworks with probabilistic masks. Rather than discretely removing parameters, we instead create probability matrices that alter the chance for parameters to be retained on any given forward pass. We then anneal those probability values towards a deterministic binary mask over several epochs. We introduce several approaches for constructing stochastic subnetworks. Random annealing samples from a uniform distribution to assign random probabilities for every parameter, while temperature annealing applies an even amount of stochasticity to all parameters not in the target subnetwork.

The efficacy of Stochastic Subnetwork Annealing is affirmed by recent insights revealed in optimization research investigating the impact of warmup epochs [52, 54, 110]. Early epochs are critical during optimization as they set the foundational trajectory for the rest of the training process. Stochastic Subnetwork Annealing provides effective regularization during the early epochs of subnetwork tuning to promote training stability and encourage robust adaptation.

Our experiments display marked improvement over the established one-shot and iterative pruning benchmarks for subnetworks at various levels of sparsity. This technique is especially effective for very sparse models up to 98%. We conduct an extensive ablation study to explore the dynamics of this technique with regard to different hyperparameters, pruning methodologies, and learning rate schedules. We additionally implement this technique in the context of Sparse Subnetwork Ensembles where we report significantly better performance against benchmark ensemble methods.

# Chapter 7

## Diversity Analysis

### 7.1 Introduction

Diversity is known to be an important consideration in the construction of robust neural network ensembles [13, 91, 154]. A collection of well trained models will generalize better if they are diverse in the patterns they respond to and the predictions they make. Diversity is especially important for low-cost ensemble methods because members often share network structure in order to avoid training several independent models from scratch. Diversity is traditionally analyzed by measuring differences between the outputs of models. However, this gives little insight into how knowledge representations differ between ensemble members.

Meanwhile, neural network interpretability methods continue to grow in popularity and importance in machine learning research [182]. This field is shifting the preconceived notions that deep neural networks act as black box models. By utilizing feature visualization and saliency attribution techniques, it becomes possible to better explore *how* and *why* neural networks make decisions by visualizing the types of patterns that specific neurons within the network respond to. These methods are especially useful for evaluating low-cost ensemble methods that share network structure, because we can meaningfully compare feature representations of the same neuron in two different networks and visualize how they evolve and change over time.

We explore these ideas by conducting a large scale analysis of both traditional diversity metrics and this new interpretable approach in Sparse Subnetwork Ensembles. Our experiments reveal that the unique network topology of Sparse Subnetwork children encourages each ensemble member to learn and respond to unique patterns and feature representations. This insight could be highly effective for building more robust ensemble techniques in the future as our experiments offer a new perspective for analyzing diversity. The application of

**Table 7.1:** Average prediction disagreement ratio and KL divergence between ensemble members for several benchmark methods on CIFAR-10 with WideResNet-28x10. Results reported from [61, 106, 161]

Methods	$d_{PDR} \uparrow$	$d_{KL} \uparrow$	Acc $\uparrow$
Treenet	0.010	0.010	95.9
BatchEnsemble	0.014	0.020	96.2
EDST Ensemble	0.026	0.057	96.4
MIMO	0.032	0.081	96.4
Dense Ensemble	0.032	0.086	96.6
Snapshot Ensemble	0.034	0.068	96.3
DST Ensemble	0.035	0.095	96.4
Sparse Subnetwork Ensemble	0.036	0.090	96.5

interpretability techniques can be invaluable for better understanding how to measure and promote diversity in low-cost ensembles.

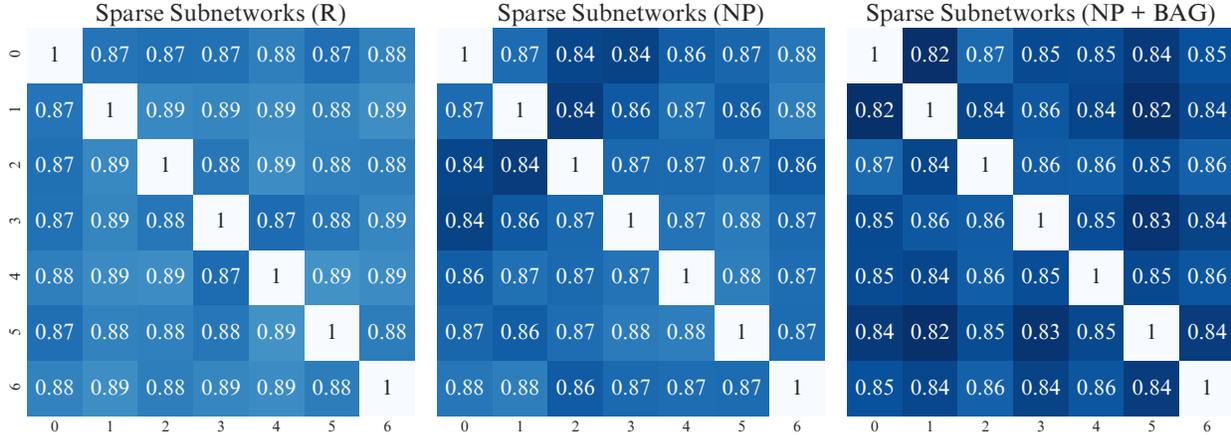
## 7.2 Output Diversity

Several metrics have been used to quantify diversity in classification ensembles, the most popular of which are the pairwise output correlation, Kullback-Leibler divergence, and prediction disagreement ratio [91]. Pairwise Output Correlation measures the average correlation between the outputs for each pair of models in the ensemble (Eqn. 7.1). Kullback-Leibler Divergence measures the relative entropy between two probability distributions (Eqn. 7.2). It approximately describes how different one probability distribution is from one another. Prediction Disagreement Ratio measures the pairwise ratio of disagreements between the predicted classes of models (Eqn. 7.3).

$$d_{corr}(f_1, f_2) = \frac{1}{N} \sum_{i=1}^N \sum_{c=1}^C \frac{\text{cov}(f_1(x_i)_c, f_2(x_i)_c)}{\sigma_{f_1(x_i)_c} \sigma_{f_2(x_i)_c}} \quad (7.1)$$

$$d_{KL}(f_1, f_2) = \frac{1}{N} \sum_{i=1}^N \sum_{c=1}^C f_1(x_i)_c \log \left( \frac{f_1(x_i)_c}{f_2(x_i)_c} \right) \quad (7.2)$$

$$d_{PDR}(f_1, f_2) = \frac{1}{N} \sum_{i=1}^N \text{argmax}(f_1(x_i)) \neq \text{argmax}(f_2(x_i)) \quad (7.3)$$



	Ensemble Accuracy $\uparrow$	Mean Correlation $\downarrow$
Sparse Subnetwork Ensemble (R)	0.840	$0.893 \pm 0.041$
Sparse Subnetwork Ensemble (NP)	0.844	$0.884 \pm 0.045$
Sparse Subnetwork Ensemble (NP + BAG)	0.839	$0.868 \pm 0.051$

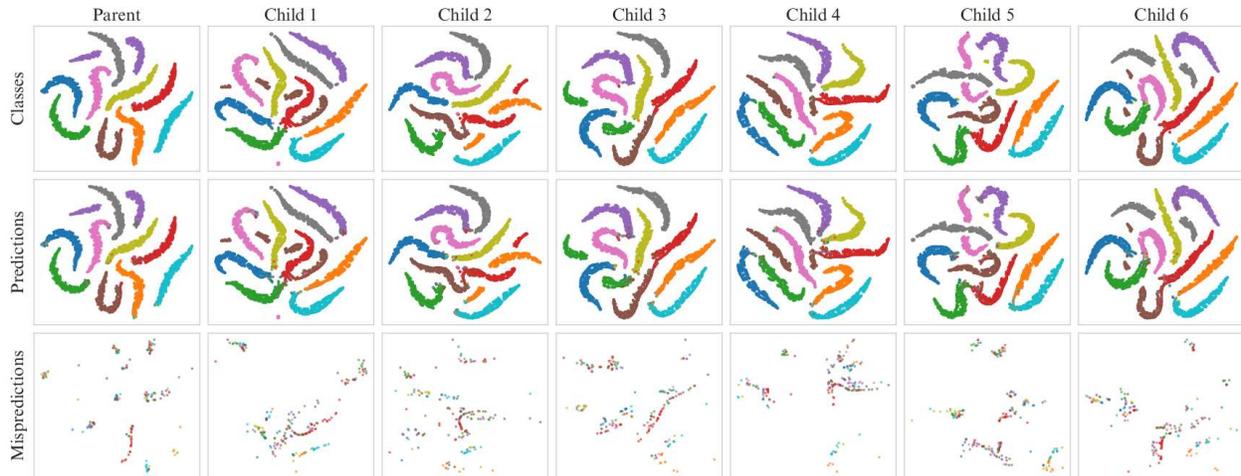
**Figure 7.1:** Pairwise correlation heatmaps for ResNet-18 on CIFAR-10. (R/NP) denotes random pruning/neural partitioning. (BAG) denotes bootstrap aggregation.

where  $N$  is the number of test samples,  $x_i$  in an input sample,  $f_i(x_i)$  is the output of a model,  $\text{argmax}(f_i(x_i))$  is the predicted class label,  $\text{cov}()$  is the covariance function, and  $\sigma_{f(x_i)}$  is the standard deviation of the output vector.

Table 7.1 displays average prediction disagreement ratios and KL divergences for several low-cost ensembles on CIFAR-10 with WideResNet-28x10. These are the same ensembles evaluated in the large scale benchmark experiments described in Chapters 5 and 6. Figure 7.1 displays the pairwise correlation values for three variations of Sparse Subnetwork Ensembles. These include children created with random pruning (R), children created with neural partitioning (NP), and children tuned with bootstrap aggregation (BAG). While bootstrap aggregation does result in more diverse children, the accuracy is lower than methods where the children are trained on the whole dataset.

### 7.2.1 Prediction Clusters

We next use t-distributed stochastic neighbor embedding (T-SNE) to visualize the representations of test data predictions into lower dimensional clusters [155]. T-SNE builds



**Figure 7.2:** T-SNE cluster representations of CIFAR-10 and WideResNet-28x10. The first row is the representation of classes in CIFAR-10 that each of our models produce. The second row displays the results of our model predictions and the final row displays the incorrect predictions.

these cluster visualizations by converting high dimensional outputs into a lower dimensional embedding while preserving the relationships between neighboring data points. It does this by first calculating a pairwise similarity between points in the high dimensional space using a Gaussian kernel. Then the algorithm tries to map the high dimensional outputs to the low dimensional space while preserving the pairwise similarity by minimizing the KL divergence between them.

Figure 7.2 displays the results of this algorithm for the parent and child networks in a Sparse Subnetwork Ensemble. The parent network is a WideResNet-28x10 trained on CIFAR-10 using the training procedure described in the benchmark experiments in Chapters 5 and 6. These cluster visualizations can reveal interesting insights regarding diversity with regard to the types of classes and predictions that are clustered together. The mispredictions reveal significant diversity among child networks as these cluster visualizations illustrate their distinct weaknesses.

## 7.3 Interpretable Diversity

Evaluating diversity through outputs has notable limitations as it can paint an incomplete picture of the underlying diversity within an ensemble. These metrics may be sensitive to data distributions and depend heavily on individual member accuracy. Figure 7.1 provides an example how a group of less accurate models with high diversity can perform worse than a group of more accurate models with lower diversity. Because of this, output diversity metrics can be misleading when comparing the results of different ensemble methods. These diversity metrics alone are not sufficient to fully analyze the generalization capabilities of different ensemble methods.

We explore how we can extend diversity analysis with neural network interpretability methods. These methods can be used to quantify the diversity between models in the learned feature spaces, rather than solely on output predictions. By visualizing the types of patterns that ensemble members respond to, we can gain a more comprehensive look at how diversity can play a role in decision making processes.

Feature visualization uses optimization to create images that maximize or minimize activations of specific parts of the network. This is done by first initializing an image with random noise. This image serves as the starting point for an optimization problem wherein the objective is to maximize or minimize the activation of target neurons. This image is then fed through a pre-trained network, where stochastic gradient descent is used to iteratively update the pixel values of the image to optimize the activation of the designated neuron. The resulting images offer critical insights into the types of patterns and features that particular neurons respond to.

Saliency Maps visualize the gradient of a prediction with respect to the input image, highlighting the parts of the input that a network responds strongly to. This is done by first performing a forward pass through the network with a given input image. The saliency is computed by setting all outputs of the non-predicted class to 0 and back propagating the predicted class score back through to the inputs [132]. This is the gradient  $G(X)$  of the

output of the predicted class label  $Y_c$  with respect to the input image  $X$ . SmoothGrad was later introduced as an extension that reduces noise in the visualizations by averaging the gradient over several Gaussian perturbed inputs  $\gamma \sim \mathcal{N}(0, \sigma^2)$  [137].

$$G(X) = \frac{\partial Y_c}{\partial X} \tag{7.4}$$

$$SG(X) = \frac{1}{N} \sum_{i=1}^N \frac{\partial Y_c}{\partial X + \gamma_i} \tag{7.5}$$

### 7.3.1 Perceptual Hashing

Image similarity metrics are generally used to measure image quality as a response to noise, corruption, or compression against a ground truth image [135]. While these metrics are simple to compute and provide a straightforward quantitative measure, they tend to be highly sensitive and may not correlate well with human perception of image quality. Sensitivity to pixel level details and small spatial transformations can be problematic when applied to images of feature visualizations that are generated via stochastic processes.

We instead introduce a robust image hash similarity based approach to quantifying diversity. Image hashing algorithms compress images into binary bit strings such that images that are visually similar will result in hashes that are similar. The bits of the image hash are spatially significant, so the similarity between two hashes can be described by the Hamming Distance, which is the number of positions in which the bits of two hashes differ. Images that are more similar will therefore have a lower Hamming Distance than images that are diverse [177].

We use several popular image hashing algorithms for measuring the distances between feature visualizations, including: average hash, perceptual hash, difference hash, wavelet hash, and color hash [14]. Average hash, difference hash, wavelet hash and perceptual hash all start by resizing a given image to an 8x8 pixel square and converting it to grayscale. The hash is then constructed by assigning a 1 or 0 for each pixel in the block according to some heuristic. Average Hash assigns a 1 if the pixel value is greater than the mean

pixel value. Perceptual Hash does the same except all values are converted to the frequency domain using a discrete cosine transformation first. Wavelet hash uses a discrete wavelet transform instead of a discrete cosine transform. Difference hash is a gradient hash where the difference between pixels values is compared to the mean difference between pixel values. Finally, color hash maintains color information by skipping the grayscale step and instead using the hue, saturation, and value space of an image separately.

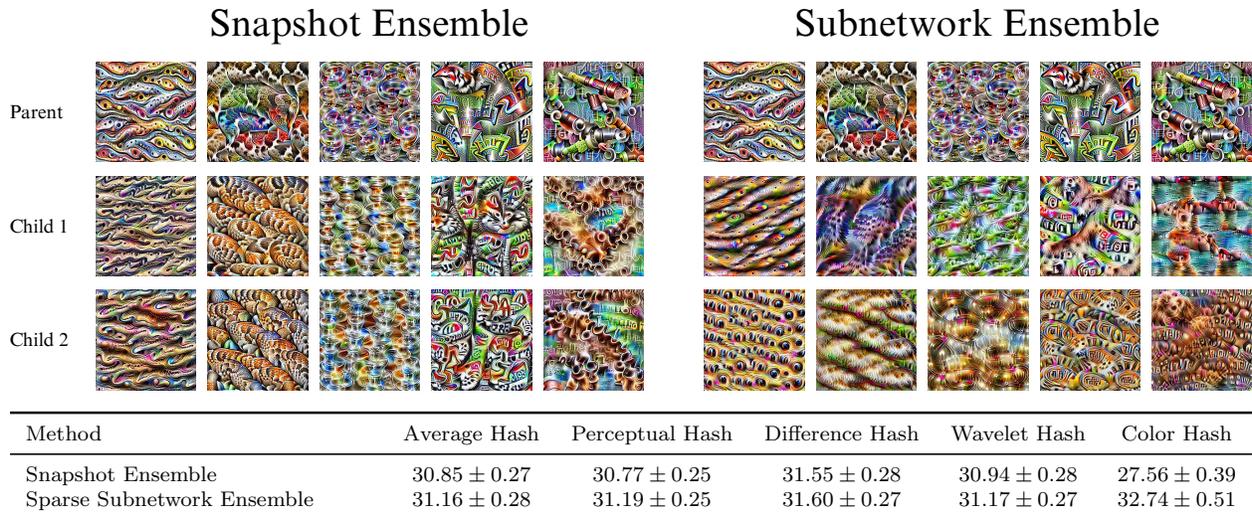
## 7.4 Experiments

We conduct a large scale comparison between child networks in a Snapshot Ensemble and a Sparse Subnetwork Ensemble. We start with an open source Inception network pre-trained on ImageNet [123]. Following the hyperparameters given in the Snapshot Ensemble paper, we continue training of the pre-trained network for two cycles of 40 epochs in length with a cyclic cosine-annealing decay schedule with an initial learning rate of 0.1 and a final learning rate of  $1e-5$  [70].

We use the same pre-trained network as the parent for the Sparse Subnetwork Ensemble as well. Two child networks are created via neural partitioning with an unstructured sparsity target of 50%. Each child is fine tuned for 40 epochs using a one-cycle learning rate schedule with a max learning rate of 0.1 and a final learning rate of  $1e-5$ . We choose these hyperparameters to match the computational budget defined in Snapshot Ensembles [70].

Using the child networks constructed above, we create feature visualizations of identical neurons in each of the networks [83, 151]. Feature visualizations are optimized using ADAM with a learning rate of 0.05 for 1024 steps. Several random data augmentations are applied at each step, including: jittering by up to 8 pixels, scaling by a factor between 0.95 and 1.05, rotating by an angle between -5 and 5 degrees, and jittering a second time by up to 4 pixels.

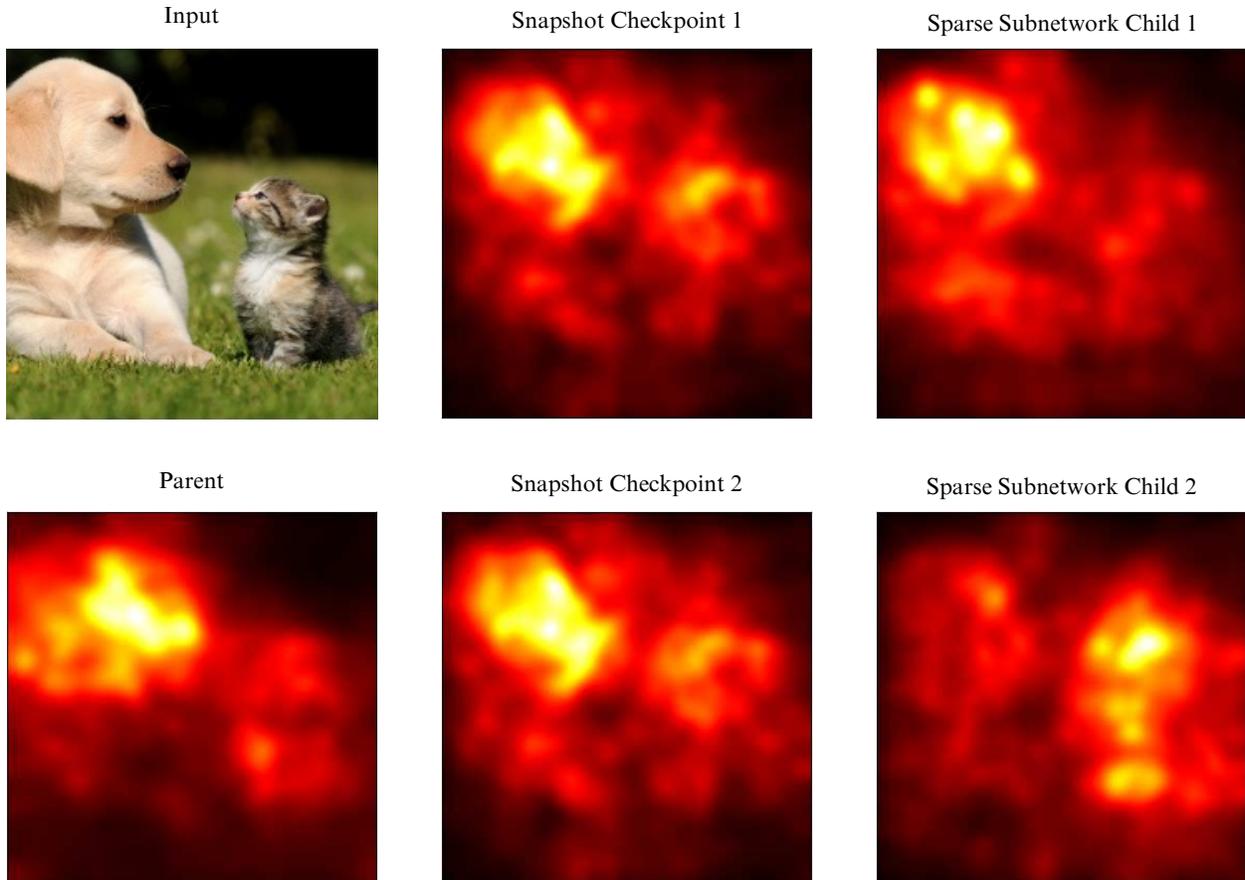
We then compute 64 bit image hashes for each visualization (average hash, perceptual hash, difference hash, wavelet hash, and color hash) and report the Hamming Distance between each pair of child network visualization hashes.



**Figure 7.3:** Feature visualizations of a random selection of neurons in child networks from a Snapshot Ensemble and a Sparse Subnetwork Ensemble. Values reported are the mean hash distances between the two child networks over all 1024 visualizations in the final convolutional layer. Higher values represent greater distance and more diversity between representations. The sparsity from Subnetwork Ensembles results in more diverse feature representations than two subsequent checkpoints in Snapshot Ensembles.

We first present feature visualizations from several neurons selected at random from different layers within each network. Figure 7.3 displays some sample visualizations where each column corresponds to the same neuron in each of the different networks. Subnetwork children often focus on different elements present in the parent network’s visualization, while Snapshot checkpoints appear to be much more highly correlated with the parent.

We extend these feature visualizations by analyzing every channel in the final convolutional layer. Thus, for each pair of child networks, 1024 feature visualizations are created for each. All of these visualizations are hashed using the algorithms described above and then the distance between sibling network hashes are measured. These final neurons are especially relevant as they represent the final stage of feature extraction before being fed into the fully-connected classifier layers. They most closely illustrate how each child network represents concepts and the differences over all the neurons in this layer will provide a comprehensive overview on the diversity between network representations. The table in Figure 7.3 reports the mean Hamming distances between each pair of child networks over



	RMSE	Average Hash	Perceptual Hash	Difference Hash	Wavelet Hash
Snapshot Ensemble	$0.06 \pm 0.001$	$10.01 \pm 0.04$	$18.12 \pm 0.04$	$10.00 \pm 0.04$	$15.13 \pm 0.04$
Sparse Subnetwork Ensemble	$0.07 \pm 0.001$	$11.13 \pm 0.04$	$19.95 \pm 0.04$	$11.03 \pm 0.04$	$15.97 \pm 0.04$

**Figure 7.4:** Mean distances between Snapshot Ensemble and Sparse Subnetwork Ensemble saliency maps for all 50,000 samples in the ILSVRC2012 ImageNet validation set. The visualization includes SmoothGrad saliency maps for a shared parent network, two snapshot ensemble checkpoints, and two sparse subnetwork children. The visualization is an example of how saliency maps can be useful with ambiguous input as both snapshot checkpoints focus on the dog while the subnetworks split their attention between the dog and the cat. The diversity encouraged by sparse network structures help to reduce bias present in the parent network.

all feature visualization hashes. For every hashing algorithm we see more diversity in the Sparse Subnetwork representations than in the Snapshot Ensemble.

We then explore how each ensemble method interprets input by creating SmoothGrad saliency maps for each of the 50,000 samples in the ImageNet ILSVRC2012 validation set [27]. Figure 7.4 reports the average Root Mean Square Error, Average Hash, Perceptual Hash, Difference Hash, and Wavelet Hash distance between each pair of child networks over all

50,000 images. We also include an example visualization of how saliency maps can be useful when dealing with ambiguous inputs. We feed an image containing a puppy and a kitten to each child network in the Snapshot and Subnetwork Ensemble and compute their saliency maps. The parent network and the Snapshot checkpoints focus primarily on the puppy in the input image. The subnetwork children split their attention, where one child focuses on the puppy and the other on the kitten. This diversity can be valuable in building robust ensembles as techniques that encourage diverse representations can help to reduce bias (focus on the puppy) that is inherent in the parent network.

## 7.5 Discussion

Diversity is an important property of robust neural network ensembles. However, traditional measures of diversity focus only on model outputs and give little insight into how ensemble members represent knowledge differently. We introduce a combination of interpretability methods and perceptual hashing as an effective approach for qualitatively analyzing diversity and measuring the similarity between representations.

Diversity is especially important for low-cost ensemble methods as they tend to share information between members in order to reduce computational cost. In particular, temporal ensembles and evolutionary ensembles create members that inherit parameters and network structure from previous iterations, which allows for meaningful comparisons between the same neurons of different networks.

Our experiments explore how feature visualization and saliency maps could be useful in comparing the diversity between members in Subnetwork Ensembles and we compare that diversity to checkpoints taken in Snapshot Ensembles. We find that the sparsity and unique network topology of sparse subnetwork children is highly effective for encouraging diverse feature representations when compared to Snapshot Ensembles. Neural partitioning ensures that the convolutional filters are geometrically opposed which results in convergence

to unique and diverse optima. Despite the significant pruning that these children undergo, Sparse Subnetwork Ensembles maintain high classification accuracy.

Interpretability methods are making deep neural networks more accessible and understandable, and we believe the introduction of these methods to ensemble learning can provide better insights into diversity and aid in the construction of more robust low-cost ensembles.

# Chapter 8

## Conclusions

### 8.1 Summary

We introduce and formalize Subnetwork Ensembling as a general framework for creating ensembles of deep neural networks, where a large parent network is trained and then used to spawn an ensemble of child networks through a perturbative process applied to randomly sampled subnetworks in the parent. Several methods were explored for constructing and optimizing child networks through noise perturbations, network pruning, and stochastic masking. We augment these processes with the introduction of Neural Partitioning, which is an effective variance reduction technique where the parameters of parent networks are divided into sets of children with no parametric overlap. This work is grounded in the landscape of modern deep ensemble learning research where each method was thoroughly evaluated across a number of benchmark tasks and ablation studies.

We frame Noisy Subnetwork Ensembles as a no-cost ensemble approach where a collection of noise perturbed children can reliably improve generalization on the difficult tasks without requiring any additional training. We relate some of the principles behind this work to popular ideas in Neuroevolutionary research. We build on this with Sparse Subnetwork Ensembles where a fine tuning phase is applied to sparse children to significantly improve the predictive capacity of the ensemble. The unique sparse network topology encourages child networks to converge to unique optima. We compare Sparse Subnetwork Ensembles to many recent ensemble learning methods and we observe state-of-the-art performance on fixed training budget benchmarks. We further improve performance with the introduction of Stochastic Subnetwork Ensembles, where subnetworks are represented with probability matrices that govern parameter activation on forward passes. We slowly anneal the probability values of the stochastic mask over time towards a deterministic target subnetwork.

We observe significant improvements over established one-shot and iterative pruning baselines and we further improve performance on benchmarks explored with Sparse Subnetwork Ensembles.

We then conduct a large scale analysis of diversity in Sparse Subnetwork Ensembles. We first explore the diversity of predictions through common diversity metrics used in ensemble learning literature. We then introduce a novel approach to analyzing feature space diversity by leveraging network interpretability methods. We use feature visualization and saliency maps to visualize knowledge representations with ensemble members. We introduce perceptual hashing algorithms as a means to quantify the visual differences between feature representations. We compare the results of these investigations to another popular low-cost ensemble method called Snapshot Ensembles, where we found that Sparse Subnetwork Ensembles display significantly more representational diversity among sibling networks. These results reaffirm the hypothesis that topological structure plays a critical role in encouraging diverse and robust learning.

## 8.2 Discussion

Subnetwork Ensembling proves to be a highly effective technique for improving the training efficiency, parametric utilization, and generalization of deep ensembles while minimizing computational cost. We achieve this by uniting several important concepts and insights observed in deep learning research. Utilizing a large overparameterized parent network facilitates better convergence efficiency and training dynamics. These trained parent networks often contain large amounts of redundancy that we can leverage to construct accurate child networks that train quickly. Utilizing different subspaces of the parent network encourages diverse representations among child networks as topological transformations affect the underlying optimization landscape. The introduction of sparsity, through subnetwork sampling, serves as a regularization effect that leads to more robust solutions. This leads to more

diversity among child networks as distinct network topologies can transform the underlying optimization landscape such that child networks converge to diverse local optima.

This approach also offers several unique advantages over competitive low-cost ensemble methods. Parent networks can be optimized independently of the child networks which allows for flexibility with network architectures, regularization techniques, and training procedures. Child networks can be generated and trained quickly, which enables dynamic growth of ensemble sizes and parallel wide scaling. Child networks can be made to be very sparse with little discernible loss in accuracy which holds great potential for reducing memory and storage costs.

While Subnetwork Ensembles offer numerous advantages for improving training efficiency, generalization, and computational requirements, it is not without its limitations. Subnetwork Ensembles are fundamentally constrained by the size and capability of the parent network. This necessitates a parent network that is large enough to ensure that sampled subnetworks can produce accurate solutions. Furthermore, our research prioritizes the efficient training of these ensembles, which is an increasingly important consideration as the training costs of large scale neural networks are becoming prohibitively expensive and time consuming. However, it is important to note that there is a large body of work that instead looks to optimize the inference cost of evaluating ensembles. Our methodology entails separate forward passes for each subnetwork, rendering it less efficient than pseudo-ensemble techniques that incorporate members into a unified model.

One objective of our investigation into Subnetwork Ensembles is to formulate this process under a generalized framework that could provide a foundation for conducting future research. Emerging network architectures, algorithmic techniques, modalities and problem domains offer new opportunities for implementing these ensembles to improve generalization at minimal cost. Connections between Subnetwork Ensembles and observed structures in biological neural networks may hold interesting prospects for research that aims to bridge the gaps between machine learning and neuroscience. Our approach would also be a nat-

ural fit for multitask learning, which is an important area of research for artificial general intelligence. If the trend of increasing model scale continues, the utilization of subnetworks may become increasingly important for training large models. Colossal super networks may become so large that traditional optimization methods become infeasible and will instead need to rely on the isolation, distribution, and organization of its subnetworks.

As artificial intelligence continues to become a larger part of the cultural conversation, it's interesting to consider how our research impacts the ways in which we view the world. Just as Subnetwork Ensembles benefit from the harmonious interaction of its members, so too does the universe operate through interconnected relationships at all scales. From the diverse set of subatomic quanta that form space-time, to the delicate balance of ecosystems that support life, and even to the gravitational relationships between celestial bodies. We consider how communities at all scales thrive when diverse perspectives and properties are integrated. Subnetwork Ensembles can serve as a compelling metaphorical lens through which to examine the intricate and complex set of relationships that define and sustain the natural world.

# Bibliography

- [1] David Attwell and Simon B. Laughlin. An energy budget for signaling in the grey matter of the brain. *Journal of Cerebral Blood Flow & Metabolism*, 21(10):1133–1145, 2001. PMID: 11598490.
- [2] Philip Bachman, Ouais Alsharif, and Doina Precup. Learning with pseudo-ensembles. *ArXiv preprint:1412.4864*, 2014.
- [3] Thomas Bäck and Hans-Paul Schwefel. An overview of evolutionary algorithms for parameter optimization. *Evolutionary computation*, 1(1):1–23, 1993.
- [4] Roberto Battiti. First-and second-order methods for learning: between steepest descent and newton’s method. *Neural computation*, 4(2):141–166, 1992.
- [5] Hans-Georg Beyer and Hans-Paul Schwefel. Evolution strategies—a comprehensive introduction. *Natural computing*, 1:3–52, 2002.
- [6] Davis Blalock, Jose Javier Gonzalez Ortiz, Jonathan Frankle, and John Gutttag. What is the state of neural network pruning? In *Proceeding of the Machine Learning and Systems Conference*, 2020.
- [7] Hendrik Blockeel. *Hypothesis Space*, pages 511–513. Springer US, Boston, MA, 2010.
- [8] Hamed Bonab and Fazli Can. A theoretical framework on the ideal number of classifiers for online ensembles in data streams. pages 2053–2056, 10 2016.
- [9] Léon Bottou and Olivier Bousquet. The tradeoffs of large scale learning. *Advances in neural information processing systems*, 20, 2007.
- [10] Leo Breiman. Bagging predictors. *Machine learning*, 24(2):123–140, 1996.
- [11] Leo Breiman. Random forests. *Machine learning*, 45:5–32, 2001.

- [12] Dimo Brockhoff, Anne Auger, Nikolaus Hansen, Dirk Arnold, and Tim Hohm. Mirrored sampling and sequential selection for evolution strategies. 04 2010.
- [13] G. Brown, J. Wyatt, R. Harris, and X. Yao. Diversity creation methods: A survey and categorisation. *Information Fusion*, 6:5–20, 03 2005.
- [14] J. Buchner. Imagehash. <https://github.com/JohannesBuchner/imagehash>, 2021.
- [15] Luis Carrillo-Reid and Rafael Yuste. What is a neuronal ensemble? 08 2020.
- [16] Rich Caruana. Multitask learning. *Machine learning*, 28:41–75, 1997.
- [17] Cerebras. Powering extreme-scale hpc with cerebras waferscale accelerators, 2022.
- [18] Gal Chechik, Isaac Meilijson, and Eytan Ruppin. Synaptic Pruning in Development: A Computational Account. *Neural Computation*, 10(7):1759–1777, 10 1998.
- [19] Tianqi Chen and Carlos Guestrin. XGBoost. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, aug 2016.
- [20] Wuyang Chen, Wei Huang, and Zhangyang Wang. No free lunch in neural architectures? a joint analysis of expressivity, convergence, and generalization. In *AutoML Conference 2023*, 2023.
- [21] Anna Choromanska, Mikael Henaff, Michael Mathieu, Gérard Ben Arous, and Yann LeCun. The loss surfaces of multilayer networks, 2015.
- [22] Karl Cobbe, Christopher Hesse, Jacob Hilton, and John Schulman. Leveraging procedural generation to benchmark reinforcement learning, 2020.
- [23] Cody A. Coleman, D. Narayanan, Daniel Kang, T. Zhao, Jian Zhang, Luigi Nardi, Peter Bailis, K. Olukotun, C. Ré, and M. Zaharia. Dawnbench : An end-to-end deep learning benchmark and competition. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, ML Systems Workshop*, 2017.

- [24] Andrew R Conn, Nicholas IM Gould, and Philippe L Toint. *Trust region methods*. SIAM, 2000.
- [25] Moheb Costandi. *Neuroplasticity*. MIT Press, 2016.
- [26] Tim de Bruin, Jens Kober, Karl Tuyls, and Robert Babuška. Fine-tuning deep rl with gradient-free optimization. *IFAC-PapersOnLine*, 53(2):8049–8056, 2020.
- [27] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [28] Thomas G Dietterich. Ensemble methods in machine learning. In *International workshop on multiple classifier systems*, pages 1–15. Springer, 2000.
- [29] James Diffenderfer and Bhavya Kailkhura. Multi-prize lottery ticket hypothesis: Finding accurate binary neural networks by pruning a randomly weighted network, 2021.
- [30] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [31] David A Drachman. Do we have brain to spare?, 2005.
- [32] Harris Drucker and Corinna Cortes. Boosting decision trees. *Advances in neural information processing systems*, 8, 1995.
- [33] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(7), 2011.
- [34] Saso Džeroski and Bernard Ženko. Is combining classifiers with stacking better than selecting the best one? *Mach. Learn.*, 54(3):255–273, March 2004.

- [35] Gerald M. Edelman. *Neural Darwinism : the theory of neuronal group selection*. Basic Books, 1987.
- [36] Erich Elsen, Marat Dukhan, Trevor Gale, and Karen Simonyan. Fast sparse convnets, 2019.
- [37] Lasse Espeholt, Hubert Soyer, Remi Munos, Karen Simonyan, Volodymir Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, Shane Legg, and Koray Kavukcuoglu. Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures, 2018.
- [38] Chrisantha Fernando, Dylan Banarse, Charles Blundell, Yori Zwols, David Ha, Andrei A. Rusu, Alexander Pritzel, and Daan Wierstra. Pathnet: Evolution channels gradient descent in super neural networks, 2017.
- [39] David Ferrucci, Eric Brown, Jennifer Chu-Carroll, James Fan, David Gondek, Aditya Kalyanpur, Adam Lally, J William Murdock, Eric Nyberg, John Prager, Nico Schlaefer, and Christopher Welty. Building watson: An overview of the deepqa project. *AI Magazine*, 31:59–79, 09 2010.
- [40] Dario Floreano, Peter Dürri, and Claudio Mattiussi. Neuroevolution: from architectures to learning. *Evolutionary Intelligence*, 1:47–62, 2008.
- [41] Stanislav Fort, Huiyi Hu, and Balaji Lakshminarayanan. Deep ensembles: A loss landscape perspective, 2020.
- [42] Tiago M. Fragoso, Wesley Bertoli, and Francisco Louzada. Bayesian model averaging: A systematic review and conceptual classification. *International Statistical Review*, 86(1):1–28, Dec 2017.
- [43] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks, 2019.

- [44] Yoav Freund and Robert E Schapire. Game theory, on-line prediction and boosting. In *Proceedings of the ninth annual conference on Computational learning theory*, pages 325–332, 1996.
- [45] Kunihiko Fukushima. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological cybernetics*, 36(4):193–202, 1980.
- [46] Christian Gagné, Michèle Sebag, Marc Schoenauer, and Marco Tomassini. Ensemble learning for free with evolutionary algorithms? 2007.
- [47] Adam Gaier and David Ha. Weight agnostic neural networks, 2019.
- [48] Adam Gaier and David Ha. Weight agnostic neural networks, 2019.
- [49] Edgar Galvan and Peter Mooney. Neuroevolution in deep neural networks: Current trends and future challenges. *IEEE Transactions on Artificial Intelligence*, 2(6):476–493, dec 2021.
- [50] Timur Garipov, Pavel Izmailov, Dmitrii Podoprikin, Dmitry Vetrov, and Andrew Gordon Wilson. Loss surfaces, mode connectivity, and fast ensembling of dnns. *arXiv preprint:1802.10026*, 2018.
- [51] Samanwoy Ghosh-Dastidar and Hojjat Adeli. Spiking neural networks. *International journal of neural systems*, 19(04):295–308, 2009.
- [52] Justin Gilmer, Behrooz Ghorbani, Ankush Garg, Sneha Kudugunta, Behnam Neyshabur, David Cardoze, George Dahl, Zachary Nado, and Orhan Firat. A loss curvature perspective on training instability in deep learning, 2021.
- [53] Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Deep sparse rectifier neural networks. In Geoffrey Gordon, David Dunson, and Miroslav Dudík, editors, *Proceedings of*

- the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, pages 315–323, Fort Lauderdale, FL, USA, 11–13 Apr 2011. PMLR.
- [54] Akhilesh Gotmare, Nitish Shirish Keskar, Caiming Xiong, and Richard Socher. A closer look at deep learning heuristics: Learning rate restarts, warmup and distillation, 2018.
- [55] GraphCore. Introducing 2nd generation ipu systems for ai at scale, 2020.
- [56] J. M. Hammersley and K. W. Morton. A new monte carlo technique: antithetic variates. *Mathematical Proceedings of the Cambridge Philosophical Society*, 52(3):449–475, 1956.
- [57] L. K. Hansen and P. Salamon. Neural network ensembles. *IEEE Trans. Pattern Anal. Mach. Intell.*, 12(10):993–1001, October 1990.
- [58] Nikolaus Hansen. The cma evolution strategy: A tutorial, 2023.
- [59] Nikolaus Hansen and Andreas Ostermeier. Completely derandomized self-adaptation in evolution strategies. *Evolutionary computation*, 9(2):159–195, 2001.
- [60] Frank Harary and Edgar M Palmer. *Graphical enumeration*. Elsevier, 2014.
- [61] Marton Havasi, Rodolphe Jenatton, Stanislav Fort, Jeremiah Zhe Liu, Jasper Snoek, Balaji Lakshminarayanan, Andrew M. Dai, and Dustin Tran. Training independent subnetworks for robust prediction, 2021.
- [62] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- [63] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations, 2019.

- [64] Joel Hestness, Sharan Narang, Newsha Ardalani, Gregory Diamos, Heewoo Jun, Hassan Kianinejad, Md. Mostofa Ali Patwary, Yang Yang, and Yanqi Zhou. Deep learning scaling is predictable, empirically, 2017.
- [65] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint: 1503.02531*, 2015.
- [66] Tin Kam Ho. The random subspace method for constructing decision forests. *IEEE transactions on pattern analysis and machine intelligence*, 20(8):832–844, 1998.
- [67] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.
- [68] Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications, 2017.
- [69] Jeremy Howard et al. The 1cycle policy: Fastai documentation, 2017.
- [70] Gao Huang, Yixuan Li, Geoff Pleiss, Zhuang Liu, John E Hopcroft, and Kilian Q Weinberger. Snapshot ensembles: Train 1, get m for free. *arXiv preprint arXiv:1704.00109*, 2017.
- [71] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. *arXiv preprint:1608.06993*, 2018.
- [72] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Weinberger. Deep networks with stochastic depth, 2016.
- [73] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q. Weinberger. Deep networks with stochastic depth. In *European Conference on Computer Vision (ECCV)*, pages 646–661, 2016.

- [74] Forrest N. Iandola, Song Han, Matthew W. Moskewicz, Khalid Ashraf, William J. Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and 0.5mb model size, 2016.
- [75] Masao Ito. Long-term depression. *Annual review of neuroscience*, 12(1):85–102, 1989.
- [76] Takuya Ito. Icecube - neutrinos in deep ice, 2019.
- [77] Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization, 2019.
- [78] Herbert Jaeger. Echo state network. *scholarpedia*, 2(9):2330, 2007.
- [79] Artur Jordao and Helio Pedrini. On the effect of pruning on adversarial robustness, 2021.
- [80] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. Lightgbm: A highly efficient gradient boosting decision tree. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [81] Helmut Kettenmann, Frank Kirchhoff, and Alexei Verkhratsky. Microglia: New roles for the synaptic stripper. *Neuron*, 77(1):10–18, 2013.
- [82] Shauharda Khadka and Kagan Tumer. Evolution-guided policy gradient in reinforcement learning. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- [83] L. Kiat. Lucent. <https://github.com/greentfrapp/lucent>, 2021.

- [84] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*, 2014.
- [85] Alex Krizhevsky. Learning multiple layers of features from tiny images. *Technical Report, University of Toronto*, May 2012.
- [86] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012.
- [87] Anders Krogh and John Hertz. A simple weight decay can improve generalization. *Advances in neural information processing systems*, 4, 1991.
- [88] Anders Krogh and Jesper Vedelsby. Neural network ensembles, cross validation and active learning. In *Proceedings of the 7th International Conference on Neural Information Processing Systems, NIPS'94*, page 231–238, Cambridge, MA, USA, 1994. MIT Press.
- [89] Jan Kukačka, Vladimir Golkov, and Daniel Cremers. Regularization for deep learning: A taxonomy, 2017.
- [90] Ananya Kumar, Percy Liang, and Tengyu Ma. Verified uncertainty calibration, 2020.
- [91] Ludmila I Kuncheva and Christopher J Whitaker. Measures of diversity in classifier ensembles and their relationship with the ensemble accuracy. *Machine learning*, 51(2):181–207, 2003.
- [92] Mark Laan, Eric Polley, and Alan Hubbard. Super learner. *Statistical applications in genetics and molecular biology*, 6:Article25, 02 2007.
- [93] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning, 2017.

- [94] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles, 2017.
- [95] Duong Hoang Le and Binh-Son Hua. Network pruning that matters: A case study on retraining variants. In *International Conference on Learning Representations*, 2021.
- [96] Y. LECUN. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>.
- [97] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1(4):541–551, 1989.
- [98] Yann LeCun, John Denker, and Sara Solla. Optimal brain damage. In D. Touretzky, editor, *Advances in Neural Information Processing Systems*, volume 2. Morgan-Kaufmann, 1990.
- [99] Namhoon Lee, Thalaiyasingam Ajanthan, and Philip H. S. Torr. Snip: Single-shot network pruning based on connection sensitivity. *arXiv preprint: 1810.02340*, 2019.
- [100] Stefan Lee, Senthil Purushwalkam, Michael Cogswell, David Crandall, and Dhruv Batra. Why m heads are better than one: Training a diverse ensemble of deep networks, 2015.
- [101] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets, 2017.
- [102] Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss landscape of neural nets, 2018.
- [103] Yawei Li, Kamil Adamczewski, Wen Li, Shuhang Gu, Radu Timofte, and Luc Van Gool. Revisiting random channel pruning for neural network compression, 2022.

- [104] Zhangheng Li, Tianlong Chen, Linyi Li, Bo Li, and Zhangyang Wang. Can pruning improve certified robustness of neural networks?, 2022.
- [105] Zhuohan Li, Eric Wallace, Sheng Shen, Kevin Lin, Kurt Keutzer, Dan Klein, and Joseph E. Gonzalez. Train large, then compress: Rethinking model size for efficient training and inference of transformers, 2020.
- [106] Shiwei Liu, Tianlong Chen, Zahra Atashgahi, Xiaohan Chen, Ghada Sokar, Elena Mocanu, Mykola Pechenizkiy, Zhangyang Wang, and Decebal Constantin Mocanu. Freetickets: Accurate, robust and efficient deep ensemble by training with dynamic sparsity, 2021.
- [107] Shiwei Liu, Tianlong Chen, Xiaohan Chen, Li Shen, Decebal Constantin Mocanu, Zhangyang Wang, and Mykola Pechenizkiy. The unreasonable effectiveness of random pruning: Return of the most naive baseline for sparse training, 2022.
- [108] Yong Liu and Xin Yao. Ensemble learning via negative correlation. *Neural networks*, 12(10):1399–1404, 1999.
- [109] Yong Liu, Xin Yao, and T. Higuchi. Evolutionary ensembles with negative correlation learning. *IEEE Transactions on Evolutionary Computation*, 4(4):380–387, 2000.
- [110] Jerry Ma and Denis Yarats. On the adequacy of untuned warmup for adaptive optimization, 2021.
- [111] Rongrong Ma, Jianyu Miao, Lingfeng Niu, and Peng Zhang. Transformed 1 regularization for learning sparse deep neural networks. *Neural Networks*, 119:286–298, 2019.
- [112] Wolfgang Maass, Thomas Natschläger, and Henry Markram. Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural computation*, 14(11):2531–2560, 2002.

- [113] Yashwant Malaiya. Antirandom testing: getting the most out of black-box testing. In *Proceedings of Sixth International Symposium on Software Reliability Engineering. ISSRE'95*, pages 86 – 95, 1995.
- [114] Minsky Marvin and A Papert Seymour. Perceptrons. *Cambridge, MA: MIT Press*, 6:318–362, 1969.
- [115] Peter Mattson, Christine Cheng, Cody Coleman, Greg Diamos, Paulius Micikevicius, David Patterson, Hanlin Tang, Gu-Yeon Wei, Peter Bailis, Victor Bittorf, David Brooks, Dehao Chen, Debojyoti Dutta, Udit Gupta, Kim Hazelwood, Andrew Hock, Xinyuan Huang, Atsushi Ike, Bill Jia, Daniel Kang, David Kanter, Naveen Kumar, Jeffery Liao, Guokai Ma, Deepak Narayanan, Tayo Oguntebi, Gennady Pekhimenko, Lillian Pentecost, Vijay Janapa Reddi, Taylor Robie, Tom St. John, Tsuguchika Tabaru, Carole-Jean Wu, Lingjie Xu, Masafumi Yamazaki, Cliff Young, and Matei Zaharia. Mlperf training benchmark, 2020.
- [116] Warren S McCulloch and Walter Pitts. A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 5:115–133, 1943.
- [117] Gábor Melis. Dissecting the winning solution of the higgsm1 challenge. In *HEPML@NIPS*, 2014.
- [118] Zachary Nado, Neil Band, Mark Collier, Josip Djolonga, Michael W. Dusenberry, Sebastian Farquhar, Qixuan Feng, Angelos Filos, Marton Havasi, Rodolphe Jenatton, Ghassen Jerfel, Jeremiah Liu, Zelda Mariet, Jeremy Nixon, Shreyas Padhy, Jie Ren, Tim G. J. Rudner, Faris Sbahi, Yeming Wen, Florian Wenzel, Kevin Murphy, D. Sculley, Balaji Lakshminarayanan, Jasper Snoek, Yarin Gal, and Dustin Tran. Uncertainty baselines: Benchmarks for uncertainty robustness in deep learning, 2022.
- [119] Nvidia. Nvidia a100 tensor core gpu architecture, 2020.

- [120] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks, 2013.
- [121] Frank Puk Danijel Kivaranovic Ivan Vovkanych Corey Levinson Dimosthenis Karaflos Pascal Pfeiffer Philipp Singer, Dmitry Gordeev. Lanl earthquake prediction, 2019.
- [122] Arnu Pretorius, Surette Bierman, and Sarel Steel. A bias-variance analysis of ensemble learning for classification. 12 2016.
- [123] PyTorch. Torchvision pre-trained models. <https://pytorch.org/vision/stable/models.html>, 2022.
- [124] Alex Renda, Jonathan Frankle, and Michael Carbin. Comparing rewinding and fine-tuning in neural network pruning. In *International Conference on Learning Representations*, 2020.
- [125] Frank Rosenblatt. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological review*, 65(6):386, 1958.
- [126] Sebastian Ruder. An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747*, 2016.
- [127] Alexander Loftus Aina Tersol Ryan Chesler, Ted Kyi. Solution for kaggle vesuvius ink detection challenge, 2023.
- [128] Tim Salimans, Jonathan Ho, Xi Chen, Szymon Sidor, and Ilya Sutskever. Evolution strategies as a scalable alternative to reinforcement learning, 2017.
- [129] Robert E Schapire. The strength of weak learnability. *Machine learning*, 5(2):197–227, 1990.
- [130] John Schulman, Sergey Levine, Philipp Moritz, Michael I. Jordan, and Pieter Abbeel. Trust region policy optimization, 2017.

- [131] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms, 2017.
- [132] A. Shrikumar, P. Greenside, and A. Kundaje. Learning important features through propagating activation differences. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML'17, page 3145–3153. JMLR.org, 2017.
- [133] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- [134] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition, 2015.
- [135] Pawan Sinha and Richard Russell. A perceptually based comparison of image similarity metrics. *Perception*, 40(11):1269–1281, 2011.
- [136] Zuzana Šišková, Marie-Ève Tremblay, et al. Microglia and synapse: interactions in health and neurodegeneration. *Neural plasticity*, 2013, 2013.
- [137] Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda Viégas, and Martin Wattenberg. Smoothgrad: removing noise by adding noise, 2017.
- [138] Leslie N. Smith and Nicholas Topin. Super-convergence: Very fast training of neural networks using large learning rates. In *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications*, 2018.
- [139] Olaf Sporns. *Networks of the Brain*. MIT press, 2016.
- [140] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(56):1929–1958, 2014.

- [141] Kenneth O. Stanley and Risto Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10(2):99–127, 2002.
- [142] Julian Stier and Michael Granitzer. Structural analysis of sparse neural networks. *Procedia Computer Science*, 159:107–116, 2019.
- [143] David Stutz, Matthias Hein, and Bernt Schiele. Relating adversarially robust generalization to flat minima, 2021.
- [144] Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. On the importance of initialization and momentum in deep learning. In Sanjoy Dasgupta and David McAllester, editors, *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 1139–1147, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR.
- [145] A. Swann and N. Allinson. Fast committee learning: preliminary results. *Electronics Letters*, 34:1408–1410, 1998.
- [146] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions, 2014.
- [147] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks, 2014.
- [148] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks, 2020.
- [149] Hidenori Tanaka, Daniel Kunin, Daniel L. K. Yamins, and Surya Ganguli. Pruning neural networks without any data by iteratively conserving synaptic flow, 2020.

- [150] Amirhossein Tavanaei, Masoud Ghodrati, Saeed Reza Kheradpisheh, Timothée Masquelier, and Anthony Maida. Deep learning in spiking neural networks. *Neural networks*, 111:47–63, 2019.
- [151] Tensorflow. Lucid. <https://github.com/tensorflow/lucid>, 2021.
- [152] Timothy J Teyler and P DiScenna. Long-term potentiation. *Annual review of neuroscience*, 10(1):131–161, 1987.
- [153] Jonathan Tompson, Ross Goroshin, Arjun Jain, Yann LeCun, and Christopher Bregler. Efficient object localization using convolutional networks, 2015.
- [154] N. Ueda and R. Nakano. Generalization error of ensemble estimators. In *Proceedings of International Conference on Neural Networks (ICNN'96)*, volume 1, pages 90–95 vol.1, 1996.
- [155] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(86):2579–2605, 2008.
- [156] Li Wan, Matthew Zeiler, Sixin Zhang, Yann Le Cun, and Rob Fergus. Regularization of neural networks using dropconnect. In Sanjoy Dasgupta and David McAllester, editors, *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 1058–1066, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR.
- [157] Liyuan Wang, Xingxing Zhang, Hang Su, and Jun Zhu. A comprehensive survey of continual learning: Theory, method and application, 2023.
- [158] Abdul Wasay, Brian Hentschel, Yuze Liao, Sanyuan Chen, and Stratos Idreos. Mothenets: Rapid deep ensemble learning. *arXiv preprint:1809.04270*, 2018.
- [159] Yeming Wen, Dustin Tran, and Jimmy Ba. Batchensemble: An alternative approach to efficient ensemble and lifelong learning, 2020.

- [160] Paul J Werbos. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, 78(10):1550–1560, 1990.
- [161] T. Whitaker and D. Whitley. Prune and tune ensembles: Low-cost ensemble learning with sparse independent subnetworks. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence*, 2022.
- [162] Darrell Whitley. A genetic algorithm tutorial. *Statistics and computing*, 4:65–85, 1994.
- [163] Darrell Whitley. An overview of evolutionary algorithms: practical issues and common pitfalls. *Information and software technology*, 43(14):817–831, 2001.
- [164] Daan Wierstra, Tom Schaul, Tobias Glasmachers, Yi Sun, and Jürgen Schmidhuber. Natural evolution strategies, 2011.
- [165] Ashia C Wilson, Rebecca Roelofs, Mitchell Stern, Nati Srebro, and Benjamin Recht. The marginal value of adaptive gradient methods in machine learning. *Advances in neural information processing systems*, 30, 2017.
- [166] D.H. Wolpert and W.G. Macready. No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*, 1(1):67–82, 1997.
- [167] Mitchell Wortsman, Gabriel Ilharco, Samir Yitzhak Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S. Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, and Ludwig Schmidt. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time, 2022.
- [168] Stephen J Wright. *Numerical optimization*. 2006.
- [169] Shen Wu, Jandhyala Sridhar, Yashwant Malaiya, and Anura Jayasumana. Antirandom testing: A distance-based approach. *Journal of VLSI Design*, 2008.
- [170] Jingjing Xie, Bing Xu, and Zhang Chuang. Horizontal and vertical ensemble with deep representation for classification. *arXiv preprint: 1306.2759*, 2013.

- [171] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks, 2017.
- [172] Xin Yao and Md. Monirul Islam. Evolving artificial neural network ensembles. *IEEE Computational Intelligence Magazine*, 3(1):31–42, 2008.
- [173] Shaokai Ye, Kaidi Xu, Sijia Liu, Jan-Henrik Lambrechts, Huan Zhang, Aojun Zhou, Kaisheng Ma, Yanzhi Wang, and Xue Lin. Adversarial robustness vs model compression, or both?, 2021.
- [174] Mingjie Wang Vikas Rajpurohit Yuanzhe Zhou Yirun Zhang, Colton Smith. Jane street market prediction, 2021.
- [175] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks?, 2014.
- [176] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint: 1605.07146*, 2016.
- [177] Christoph Zauner. Implementation and benchmarking of perceptual image hash functions. 2010.
- [178] Matthew D. Zeiler. Adadelta: An adaptive learning rate method, 2012.
- [179] Chiyuan Zhang, Oriol Vinyals, Remi Munos, and Samy Bengio. A study on overfitting in deep reinforcement learning, 2018.
- [180] Guodong Zhang, Chaoqi Wang, Bowen Xu, and Roger Grosse. Three mechanisms of weight decay regularization, 2018.
- [181] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices, 2017.

- [182] Y. Zhang, P. Tiño, A. Leonardis, and K. Tang. A survey on neural network interpretability. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 5(5):726–742, 2021.
- [183] Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. A comprehensive survey on transfer learning, 2020.