

# Analyzing Neural Network Information Flow Using Differential Geometry

Anonymous authors

Paper under double-blind review

## Abstract

This paper provides a fresh view of the neural network (NN) data flow problem, i.e., identifying the NN connections that are most important for the performance of the full model, through the lens of graph theory. Understanding the NN data flow provides a tool for symbolic NN analysis, e.g., robustness analysis or model repair. Unlike the standard approach to NN data flow analysis, which is based on information theory, we employ the notion of graph curvature, specifically Ollivier-Ricci curvature (ORC). The ORC has been successfully used to identify important graph edges in various domains such as road traffic analysis, biological and social networks. In particular, edges with negative ORC are considered bottlenecks and as such are critical to the graph's overall connectivity, whereas positive-ORC edges are not essential. We use this intuition for the case of NNs as well: we 1) construct a graph induced by the NN structure and introduce the notion of neural curvature (NC) based on the ORC; 2) calculate curvatures based on activation patterns for a set of input examples; 3) aim to demonstrate that NC can indeed be used to rank edges according to their importance for the overall NN functionality. We evaluate our method through pruning experiments and show that removing negative-ORC edges quickly degrades the overall NN performance, whereas positive-ORC edges have little impact. The proposed method is evaluated on a variety of models trained on three image datasets, namely MNIST, CIFAR-10 and CIFAR-100. The results indicate that our method can identify a larger number of unimportant edges as compared to state-of-the-art pruning methods.

## 1 Introduction

Neural networks (NN) have proven to be effective in a variety of areas, such as computer vision (Dosovitskiy et al., 2020), natural language processing (Achiam et al., 2023), and autonomous systems (Hwang et al., 2024). Despite their success, NNs are known to suffer from a number of vulnerabilities, including robustness issues (Szegedy et al., 2013) and distribution shifts (Recht et al., 2019). These limitations make it difficult for trained NNs to generalize beyond training sets and impede the development of models in real system applications. To address these problems, much work has already been done, including adversarial training (Madry et al., 2017b) and out-of-distribution (OOD) detection (Hendrycks & Gimpel, 2018; Yang et al., 2022). However, these methods are usually either based on complex optimization or strong distributional assumptions. While these techniques may alleviate some of the above challenges, ultimately, their black-box nature limits our understanding of NNs' behavior and reliability.

This paper provides a fresh view of NN analysis by making use of graph theory and network science, which would provide a well-understood tool for compositional analysis of the NN graph structure and corresponding data patterns. In particular, we view a NN as a weighted directed graph, where the neurons are the graph nodes, and the connections are the graph edges. The features that transfer from layer to layer carry the information flow in this graph. Analyzing such an information graph would allow us to identify the main data flow through the NN, which would enable follow-on symbolic analysis such as NN architecture search or model repair.

In the related literature, one approach to analyze the data flow in a NN is through information theory (Khadivi et al., 2016; Shwartz-Ziv & Tishby, 2017; Nasiri, 2020), such as the seminal work on the information bottleneck principle (Tishby & Zaslavsky, 2015). These works use tools such as entropy and mutual information to analyze the information transport between NN layers and the information change during training. While such approaches provide

an interpretable way to understand the learning process and its limitations, they cannot be directly used to analyze a trained model and the corresponding data paths.

NN pruning is another method for locating the important information paths in a NN. Classic pruning methods are based on weight magnitude (Han et al., 2015), weight sensitivity (Lee et al., 2018; Santacroce et al., 2023), and loss change (Ma et al., 2023; van der Ouderaa et al., 2023). Some of these techniques can be used to prune a trained NN to a smaller model with similar performance to that of the original – this smaller model can be used as a representation of the main data paths through the original NN as well. In contrast, our paper does not aim to reduce a NN’s size but directly estimates the data flows using concepts from graph theory and differential geometry.

To identify NN data flows, we propose to employ the notion of graph curvature, in particular Ollivier-Ricci Curvature (ORC), as introduced by Ollivier (2009), which provides an effective way to quantify the importance of an edge in the graph. The ORC of a given edge is an indication of the local graph connectivity – edges with positive ORC are typically parts of a strongly connected component, whereas edges with negative ORC normally serve as bridges between connected components. The ORC has been effectively applied to a number of domains that can be modeled as graphs, e.g., road traffic analysis (Wang et al., 2022), internet routing (Ni et al., 2015), biological networks (Znaidi et al., 2023), and deep learning (Znaidi et al., 2023), including graph NNs (Liu et al., 2023). We aim to demonstrate that a similar approach can be used for feedforward NNs used for classification purposes.

Inspired by a recent paper, which demonstrates the utility of graph curvature in analyzing NN robustness (Tan et al., 2024), in this paper, we aim to develop a general graph-curvature-based method to identify the main NN data flows. We propose a novel neural curvature (NC) metric that can be used to rank NN connections according to their importance, and thus, identify the main data flows. In particular, given a trained NN and a calibration set, we 1) construct a neural graph induced by the NN structure; 2) propose a NC metric to quantify the importance of each NN connection for a given example; 3) rank NN connections according to their NC as observed in the calibration set. Identifying the most/least important connections in this way would provide a great tool for the development of curvature-based methods for NN analysis in future work.

We evaluate the proposed method on three image classification datasets, namely MNIST, CIFAR-10 and CIFAR-100. Since there is no standard way to evaluate the importance of NN connections, we perform an edge removal experiment by gradually pruning edges based on their NC. We compare our method to state-of-the-art pruning approaches, including magnitude-based pruning (Han et al., 2015), SNIP (Lee et al., 2018), and SynFlow (Tanaka et al., 2020), so as to demonstrate the effectiveness of NC. The magnitude-based method is very effective for small models and heavily regularized models but it often suffers from layer collapse, i.e., removing an entire layer, on larger models where weight magnitudes may differ significantly across layers. While SynFlow and SNIP are able to avoid layer collapse, ultimately they are still constrained by their limited ability to capture edge importance across layers. In contrast, our method provides a continuous measure of edge importance, and is, thus, able to rank edges across the entire model in a robust manner that avoids layer collapse and is better able to identify the main NN data flows.

In summary, this paper makes three contributions: 1) we introduce the notion of using graph curvature for analyzing NN data flow; 2) we provide a ranking algorithm for the importance of NN connections, through observing each connection’s neural curvature over a calibration dataset; 3) we evaluate our method in three image classification datasets so as to demonstrate the effectiveness of the proposed method.

## 2 Related Work

**NN Information Flow Analysis.** Information bottleneck is a famous tool (Tishby et al., 2000; Tishby & Zaslavsky, 2015; Shwartz-Ziv & Tishby, 2017) in NN information flow analysis. It treats a NN as a Markov chain and uses mutual information to analyze the change in information across layers. Based on the information bottleneck principle, Khadivi et al. (2016) study how the entropy of information changes between consecutive layers and develop an optimization problem that can be used in the training process for a feedforward NN. Chaddad et al. (2017) and Nasiri (2020) utilize conditional entropy to analyze the information flow in a trained convolutional NN (CNN), which can help improve the original CNN’s accuracy. Achille et al. (2019) analyzes the information encoded in a trained NN’s weights and how it affects its future performance on the test set based on Shannon and Fisher information theory. Neural Information Flow (Davis et al., 2019; 2020) was proposed as a new metric for codifying information flow, which combines mutual information with network science. There also exists work on the estimation of mutual information in noisy NNs, to

detect clusters in a trained NN (Goldfeld et al., 2019). Wang et al. (2018; 2020) find the critical subnet of a NN by finding nodes that are important for NN performance, and understand how the NN makes the decision based on those important neurons. These methods give an explicit interpretation for NNs’ decision behavior; however, they cannot be directly used to identify the main information flows in a trained model. In this paper, we use graph curvature as a local analysis tool that can quantify the importance of each edge in terms of propagating information to the next layer.

**NN Pruning Methods.** Pruning based on weight magnitude was first proposed by Han et al. (2015). They show that weights with large magnitude are more important and can contain more useful information than smaller weights. The weight magnitude-based approach has been applied effectively to either structured or unstructured pruning, and also works for large language models (Frankle & Carbin, 2018; Li et al., 2016; Sun et al., 2023; Dery et al., 2024). Weight sensitivity-based pruning methods (Lee et al., 2018; Zhao et al., 2019; Tanaka et al., 2020; Santacroce et al., 2023) provide a weight sensitivity score based on the gradient of the loss and weight. Loss change is also a widely used metric for pruning (Liu et al., 2021; Nonnenmacher et al., 2021; van der Ouderaa et al., 2023), which calculates the weight importance based on the loss change with or without that weight. There is also another group of pruning methods that use learning, such as sparsity regularization (Voita et al., 2019; Li et al., 2020; Dery et al., 2024), graph NNs (Zhang et al., 2022), and reinforcement learning (Rao et al., 2018; He et al., 2018; Zhong et al., 2018).

**Applications of Graph Curvature Methods.** The ORC (Ollivier, 2007; 2009) is the discrete-space extension of the standard Ricci curvature (Bochner, 1946), which is defined over continuous metric spaces. The ORC is a useful tool to analyze local graph connectivity and identify communities (Sia et al., 2019) in network science. The ORC has also been widely used in different graph-based domains as follows. Gao et al. (2019b) shows that ORC can help locate the vulnerable parts in the road network systems. Wang et al. (2022) utilizes ORC to improve transit network design by analyzing the mismatch of travel demand and supply. Ni et al. (2015) and Salamatian et al. (2022) apply the ORC for internet topologies and network path connectivity analysis. Ricci curvature has also shown effectiveness in biological networks robustness analysis (Simhal et al., 2025). Recently, Ricci curvature has been applied to deep learning as well, including supervising the phase transitions within the dynamics of a time-varying complex network (Znaidi et al., 2023) as well as integrating graph curvature into graph NN design (Wang et al., 2021; Sun et al., 2022; Liu et al., 2023).

### 3 Problem Statement

This section formalizes the problem considered in this paper. We focus on *trained* NN classifiers of the kind  $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ , where  $\theta$  are the NN parameters,  $\mathcal{X} \subseteq \mathbb{R}^n$  is the input space, and  $\mathcal{Y} \subseteq \mathbb{R}^p$  is the output space. In particular, we consider feedforward NNs, i.e., fully-connected NNs and CNNs, parameterized by  $\theta = \{(W_1, b_1), \dots, (W_L, b_L)\}$ , where  $L$  is the total number of layers, and  $W_i$  and  $b_i$  are the weights and biases of layer  $i$ . We assume  $f_\theta$  has either ReLU or Tanh activations after each hidden layer.

Given a trained NN  $f_\theta$  and a calibration set  $\mathcal{D}$ , the problem considered in this paper is to rank the NN connections according to their effect on classification accuracy. We aim to design a graph-curvature-based method where the goal is to rank NN parameters according to their importance for the overall NN accuracy, such that the most important parameters represent the main information paths through the NN. We aim to evaluate the proposed approach using NN parameter pruning experiments.

**Notation.** A weighted directed graph  $G(V, E, w)$  is defined as a set of vertices  $V = \{v_1, \dots, v_N\}$  and a set of directed edges  $E \subseteq \{(u, v) \mid u, v \in V\}$ , where each edge has a positive weight given by  $w : E \rightarrow \mathbb{R}_+$ . For a directed edge  $(u, v)$ ,  $u$  and  $v$  denote the source and target vertices, respectively. We denote by  $\Gamma(v)$  the neighborhood of a vertex  $v \in V$ ;  $\Gamma^{in}(v)$  is the incoming neighborhood and  $\Gamma^{out}(v)$  the outgoing neighborhood. For a NN  $f_\theta$  and an input example  $x$ , we use  $l_{i,j}(x)$  and  $n_{i,j}(x)$  to denote the logit and activation of neuron  $i$  in layer  $j$ , i.e.,  $n_{i,j}(x) = \sigma(l_{i,j}(x))$  for  $\sigma \in \{\text{Tanh}, \text{ReLU}\}$ .<sup>1</sup> We use  $K_i$  to denote the number of neurons in layer  $i$ , whereas  $[W_{k+1}]_{ij}$  denotes the weight of the NN edge connecting  $n_{i,k}$  to  $n_{j,k+1}$ . We use  $[n]$  to denote the list of integers from 1 to  $n$ .

<sup>1</sup>To obtain neuron ordering in a convolutional layer, we unroll convolution and treat convolutional layers as sparse fully-connected layers.

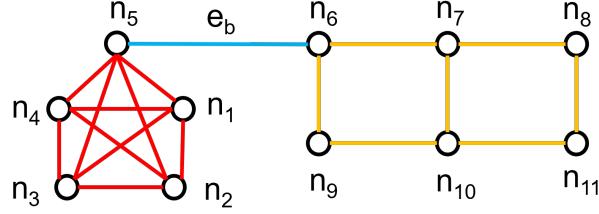


Figure 1: ORC illustration on an undirected graph with unit edge weights. Red-colored edges (left) have a positive curvature; blue edge (bridge) has a negative curvature; yellow edges (right) have a zero curvature.

## 4 Background

This section presents the definition of ORC and its modified version,  $\alpha$ -Ricci Curvature, as well as their applications. The concept of Ricci curvature (Bochner, 1946) was first introduced in Riemannian geometry to capture the degree to which a given space deviates from being flat (with Euclidean space being the standard example of a flat space). Positive curvature indicates an elevated, spherical shape, whereas negative curvature corresponds to hyperbolic shapes (with a “valley” in between two “hills”). Thus, negative curvature between two points indicates a “bottleneck” such that all shortest paths within a small enough radius include that bottleneck. ORC (Ollivier, 2009) is defined on discrete metric spaces equipped with a Markov chain and can thus be naturally extended to graphs. It aims to capture the notion of a bottleneck edge, as defined below. Further extensions of ORC have been developed, e.g.,  $\alpha$ -Ricci Curvature (Lin et al., 2011), that aim to provide a better graph approximation of the original (continuous) version of Ricci curvature.

**Remark 1** *ORC is most commonly applied to undirected graphs in network science, but several works have extended the definition to directed and weighted graphs (Yamada, 2016; Samal et al., 2018; Eidi & Jost, 2020). In the rest of this paper, we will focus on directed and weighted graphs. We follow a similar mechanism when applying ORC to a directed graph as Eidi & Jost (2020).*

We begin with a graph-based definition of Wasserstein distance, as the building block of ORC.

**Definition 1 (Wasserstein Distance (Optimal Transport Distance))** *Given a directed and weighted graph  $G(V, E, w)$ , let  $u$  and  $v$  be the source and target nodes of edge  $(u, v)$ . Suppose  $m_u$  and  $m_v$  are two probability distributions supported on the incoming-edge neighborhoods of  $u$  and outgoing-edge neighborhoods of  $v$ , respectively. The Wasserstein distance<sup>2</sup>  $W(m_u, m_v)$  between distributions  $m_u$  and  $m_v$  is given by*

$$W(m_u, m_v) = \inf_{\mu_{u,v} \in \Pi_{u,v}} \sum_{(u',v') \in \Gamma^{in}(u) \times \Gamma^{out}(v)} d(u', v') \mu_{u,v}(u', v'), \quad (1)$$

where  $\Pi_{u,v}$  is the set of probability measures  $\mu_{u,v}$  that capture all possible ways of transporting mass from  $m_u$  to  $m_v$ , i.e.,

$$\sum_{v' \in \Gamma^{out}(v)} \mu_{u,v}(u', v') = m_u(u'), \quad \sum_{u' \in \Gamma^{in}(u)} \mu_{u,v}(u', v') = m_v(v').$$

In Equation 1,  $d(u', v')$  denotes the cost from  $u'$  to  $v'$ , e.g.,  $d(u, v) = w(u, v)$ .

**Definition 2 (ORC (Ollivier, 2009))** *Given a directed and weighted graph  $G(V, E, w)$ , the ORC along a directed edge  $(u, v)$  is defined as*

$$\kappa_{ORC}(u, v) = 1 - \frac{W(m_u, m_v)}{d(u, v)}. \quad (2)$$

ORC effectively treats the graph as a Markov chain, with transition probabilities derived from edge weights (in a proportional or exponential fashion). Figure 1 provides an example undirected graph with corresponding ORCs, illustrating the relationship between curvature and graph connectivity. Intuitively, the ORC is positive for edges between

<sup>2</sup>Typically, the Wasserstein distance with 1-finite moments is written  $W_1$  – we drop the subscript to disambiguate from the NN weights notation.

nodes in a strongly connected component, and is negative for edges that serve as bridges between connected components (since mass needs to be transported from one component to the other, at a cost that is larger than the edge weight). Finally, the ORC is zero for edges in plane-like components. Thus, the ORC is an effective tool for identifying bottleneck edges in a graph. Since such bottleneck edges are essential for the overall graph connectivity, in the next section we argue that negative-curvature edges constitute the main data flow in the case when this graph encodes a corresponding NN.

**Example 1** Consider the example undirected graph in Figure 1 where all edges have weight 1. We will illustrate how to calculate the ORC of the bridge edge,  $e_b$ , connecting nodes  $n_5$  and  $n_6$ . To calculate the Wasserstein distance, we first need the distributions  $m_{n_5}$  and  $m_{n_6}$ , which assign probabilities to neighbors of  $n_5$  and  $n_6$ , respectively, depending on their weights. Since all edges have the same weight, we have  $m_{n_5} = [0.2 \ 0.2 \ 0.2 \ 0.2 \ 0.2 \ 0 \ 0 \ 0 \ 0 \ 0]$  and  $m_{n_6} = [0 \ 0 \ 0 \ 0 \ 1/3 \ 0 \ 1/3 \ 0 \ 1/3 \ 0 \ 0]$ . The optimal transport cost from  $m_{n_5}$  to  $m_{n_6}$  is  $0.2 + 3 * (2/3 - 0.2) + 1 * (1/3)$ , i.e.,  $W(m_{n_5}, m_{n_6}) \approx 1.93$ , hence  $\kappa_{ORC}(e_b) \approx -0.93$ .

Although ORC can be applied to graphs, its standard formulation assumes that the probability measure  $m_u$  around a node  $u$  is induced by a simple random walk and that the underlying metric assigns uniform distances between adjacent nodes. These assumptions are often not true in weighted graphs, where edge weights may vary a lot, leading to a more discrete and non-uniform neighborhood geometry and potentially large variations in local distances. As a result, a direct application of the standard ORC definition may fail to accurately capture the geometric structure of weighted graphs. To address this limitation, Lin et al. (2011) modified Definition 2 and proposed  $\alpha$ -Ricci Curvature, denoted by  $\kappa_\alpha(u, v)$ , where  $\alpha$  is a parameter that allows additional probability mass to remain at the source and target nodes, thereby better approximating a continuous neighborhood structure on weighted graphs.

**Definition 3 ( $\alpha$ -ORC (Lin et al., 2011))** Given a weighted and directed graph  $G(V, E, w)$ , for any  $\alpha \in [0, 1]$ , the probability distribution  $m_x$  over the neighborhood of a node  $x$  is modified to  $m_x^\alpha$ , which assigns an  $\alpha$ -fraction of the probability mass to  $x$  itself and distributes the remaining mass among its neighbors. The definition is given by:

$$m_x^\alpha(n) = \begin{cases} \alpha, & \text{if } n = x, \\ (1 - \alpha) \frac{w_n}{\sum_{n_i \in \bar{\Gamma}(x)} w_{n_i}}, & \text{if } n \in \bar{\Gamma}(x), \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

If  $x$  is a source node,  $\bar{\Gamma}(x) := \Gamma^{in}(x)$ ; if  $x$  is a target node,  $\bar{\Gamma}(x) = \Gamma^{out}(x)$ .  $w_n$  represents the weight magnitude of the edge between node  $n$  and  $x$ .

Then the  $\alpha$ -ORC curvature,  $\kappa_\alpha(u, v)$ , between two vertices  $u$  and  $v$  is given by

$$\kappa_\alpha(u, v) = 1 - \frac{W(m_u^\alpha, m_v^\alpha)}{d(u, v)}. \quad (4)$$

Note that  $\kappa_{ORC}(u, v) = \kappa_0(u, v)$  and  $\kappa_1(u, v) = 0$ , since  $W(m_u^1, m_v^1) = d(u, v)$ . Thus, it is not clear what value of  $\alpha$  is most suitable in Definition 3. To address this issue, Lin et al. (2011) further define  $h(\alpha) = \kappa_\alpha(u, v)/(1 - \alpha)$  and prove that  $h(\alpha)$  is an increasing function and remains bounded as  $\alpha \rightarrow 1$ . This implies that the limit

$$\lim_{\alpha \rightarrow 1} \frac{\kappa_\alpha(u, v)}{1 - \alpha}$$

exists. This limit is denoted by  $\kappa(u, v)$ , the Ricci curvature on graphs.

**Definition 4 (Ricci Curvature on Graphs (Lin et al., 2011))** Given a directed graph  $G(V, E, w)$ , the graph Ricci curvature of a directed edge  $(u, v)$  is defined as:

$$\kappa(u, v) = \lim_{\alpha \rightarrow 1} \frac{\kappa_\alpha(u, v)}{1 - \alpha}. \quad (5)$$

**Remark 2** In this paper, we build on Definition 4 for the proposed concept of neural Ricci curvature.

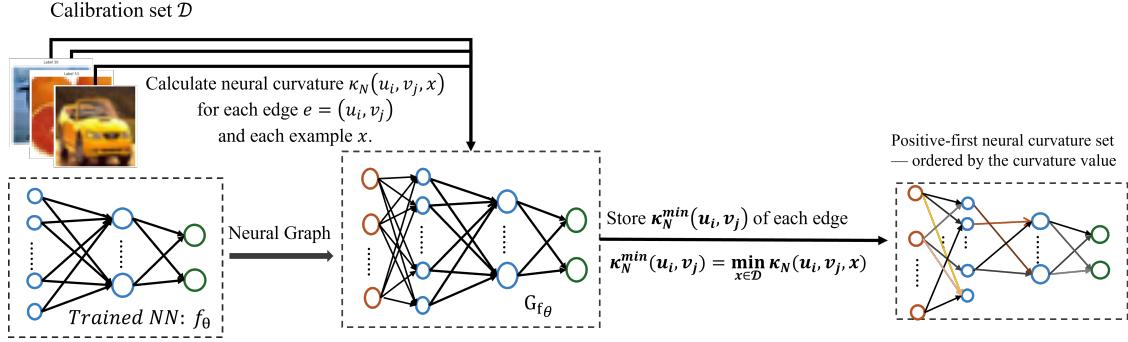


Figure 2: Approach overview. Given a trained NN  $f_\theta$  and a calibration set  $\mathcal{D}$ , we aim to identify the importance of NN connections, respectively. Specifically, we 1) construct a neural graph  $G_{f_\theta}$  induced by the NN architecture (including an the input layer); 2) for each edge  $(u, v) \in G_{f_\theta}$  and each example  $x \in \mathcal{D}$ , we calculate the neural curvature  $\kappa_N(u, v, x)$ , as defined in Definition 8; 3) finally we order the whole edge set by positive-first-curvature edges.

## 5 Approach

This section describes the proposed approach of using graph curvature to rank NN edges according to their importance. As discussed in the introduction, graph curvature has been used to analyze a number of domains that can be modeled as graphs. For example, in road and transit traffic analysis (Gao et al., 2019a; Wang et al., 2022), researchers used graph curvature to identify bottleneck segments which could be upgraded or which require better alternatives. Analyzing NNs has clear analogies and differences with road analysis. In particular, NN edges can be considered as roads that transport data, with NN weights determining the capacity of each road. Thus, it is natural to focus on bottleneck NN edges as potentially important for the overall NN functionality, similar to the traffic analysis case. The main challenge is how to encode the NN as a weighted graph, specifically how to handle 1) activation functions and 2) changing data patterns depending on the input example.

### 5.1 Overview

A high-level overview of the proposed approach is shown in Figure 2 and summarized in Algorithm 1. Given a trained NN  $f_\theta$  and a calibration set  $\mathcal{D}$ , for each NN connection  $e$  we aim to provide a curvature value that quantifies this edge's importance. Similar to the standard ORC setting, edges with lower curvature are considered more important for the overall NN functionality. We quantify this process through the following three steps: 1) create a neural graph  $G_{f_\theta}$ , as introduced in prior work (Tan et al., 2024; Xiao et al., 2024); 2) for each example  $x$  and each NN connection  $e$ , calculate the neural curvature of  $e$ , as described in this section; 3) for each NN connection  $e$ , store the minimum curvature value  $e$  over all the examples. Since more negative curvature is associated with a stronger contribution to information flow and model performance, the minimum curvature captures the most influential role that the connection may play across the dataset. The rest of this section formalizes the notions of a neural graph as well as the novel definition of neural curvature.

### 5.2 Neural Graph and Neural Curvature

This subsection introduces the concept of neural curvature. We begin with a definition of a neural graph, i.e., a weighted graph induced by the NN architecture, as introduced by prior work (Tan et al., 2024; Xiao et al., 2024).

**Definition 5 (Neural Graph (Tan et al., 2024; Xiao et al., 2024))** Consider a NN  $f_\theta$  with  $L$  layers, where  $V_l = \{v_{1,l}, \dots, v_{K_l,l}\}$  denotes the neuron set of layer  $l$  and  $E_{l,l+1} = \{(v_{1,l}, v_{1,l+1}), \dots, (v_{K_l,l}, v_{K_{l+1},l+1})\}$  denotes

---

**Algorithm 1** Rank NN connections using neural curvature

---

**Input:** Trained neural network  $f_\theta$ , calibration set  $\mathcal{D}$   
**Output:** Curvature-ranked parameter set  $edge\_set$

```

1:  $(V, E, w) \leftarrow \text{neural\_graph}(f_\theta)$  // Def. 5
2: Initialize  $edge\_cur[e] \leftarrow +\infty \quad \forall e \in E$ 
3: Initialize  $para\_set \leftarrow \emptyset$ 
4: for  $e \in E$  do
5:   for  $x \in \mathcal{D}$  do
6:     Compute  $\kappa_N(e, x)$  // Def. 8
7:      $edge\_cur[e] \leftarrow \min(edge\_cur[e], \kappa_N(e, x))$ 
8:   end for
9: end for
10: // Fully-connected parameters correspond to single edges
11: for fully-connected parameter  $p \in \theta$  do
12:   //Let  $e_p \in E$  be the edge corresponding to  $p$ 
13:    $\kappa(p) \leftarrow edge\_cur[e_p]$ 
14:    $para\_set[p] \leftarrow \kappa(p)$ 
15: end for
16: // For convolutional parameters, assign the minimum curvature among all incoming edges
17: for convolutional parameter  $p \in \theta$  do
18:    $\mathcal{E}_p \leftarrow \{e \in E \mid e \text{ is induced by } p\}$ 
19:    $\kappa(p) \leftarrow \min_{e \in \mathcal{E}_p} edge\_cur[e]$ 
20:    $para\_set[p] \leftarrow \kappa(p)$ 
21: end for
22: // Rank parameters by curvature
23:  $edge\_set \leftarrow \text{sort}(para\_set, \text{descending})$  // highest curvature first

```

---

the NN connections from layer  $l$  to  $l + 1$ . A neural graph  $G_{f_\theta} = (V_{f_\theta}, E_{f_\theta}, w_{f_\theta})$  is defined as follows:

$$\begin{aligned} V_{f_\theta} &= V_0 \cup V_1 \cup V_2 \cup \dots \cup V_L \\ E_{f_\theta} &= E_{0,1} \cup E_{1,2} \cup \dots \cup E_{L-1,L} \\ w_{f_\theta}(v_{i,l}, v_{j,l+1}) &= \frac{1}{|[W_{l+1}]_{ij}|}, \end{aligned}$$

where  $V_0 = \{v_{1,0}, \dots, v_{n,0}\}$  adds one vertex per input dimension (i.e., the NN input layer). Thus,  $G_{f_\theta}$  has one vertex per input dimension, one vertex per neuron in the NN and one edge per NN connection. The weight of each edge in the graph is the inverse of the magnitude of the corresponding NN weight.<sup>3</sup>

Definition 5 describes the graph induced by the NN structure. Note that setting the graph weights to be the inverse of the NN weights signifies that the “cost” of sending data through an edge is inversely proportional to the NN weight; intuitively, the larger the weight, the more data is sent. Other weightings will be considered in future work as well.

**Remark 3** Definition 5 implicitly treats convolutional layers as linear maps, followed by an activation function, similar to a fully-connected layer. While this transformation is valid, it creates multiple neural graph edges per NN convolutional parameter. As described at the end of this section, we calculate a single curvature value per convolution parameter, by taking the minimum curvature over all edges that correspond to a single NN parameter.

Note that Definition 5 does not depend on the input example and as such cannot be easily used to capture the NN information flow. We incorporate the input example in the curvature calculation as follows. First note that the standard  $\alpha$ -Ricci curvature in Definition 3 derives the neighbor distributions,  $m_u$  and  $m_v$ , directly from the graph weights.

<sup>3</sup>Although prior work Tan et al. (2024) considers an additional normalization of the graph weights depending on their sign, we found it to be unnecessary.

Unlike the standard definition, which is static, we incorporate the input examples into the curvature calculation. Specifically, we modify the neighbor distribution depending on the magnitude of each neuron activation for a given input example. By incorporating these node values, the curvature captures how the input propagates through the NN, making it data-dependent and providing a more meaningful measure of the NN's behavior for a specific example.

**Definition 6 (Neural Neighbor Distribution)** Consider a neural graph  $G_{f_\theta}$  and a corresponding vertex  $v_{i,l}$  for  $l \geq 1$ , and a given input example  $x$ . The incoming neural neighbor distribution of  $v_{i,l}$  is defined as follows:

$$m_{v_{i,l}}(x) = [norm(n_{1,l-1}(x)) \dots norm(n_{K_{l-1},l-1}(x))],$$

where  $norm$  is a normalization function applied in three steps:

0. Let  $n_{min} = \min\{|n_{1,l-1}(x)|, \dots, |n_{K_{l-1},l-1}(x)|\}$  and  $n_{max} = \max\{|n_{1,l-1}(x)|, \dots, |n_{K_{l-1},l-1}(x)|\}$ .

1. **Per-layer neuron normalization:**  $n_{i,l-1}^{norm}(x) = \frac{|n_{i,l-1}(x)| - n_{min}}{n_{max} - n_{min}}$ ;

2. **Inversion:** take the reciprocal of the value from step 1, i.e.,  $n_{i,l-1}^{inv}(x) = \frac{1}{n_{i,l-1}^{norm}(x)}$ ;

3. **Exponential normalization** (Ni et al., 2019):

$$norm(n_{i,l-1}(x)) = \frac{\exp(- (n_{i,l-1}^{inv}(x))^2)}{\sum_j \exp(- (n_{j,l-1}^{inv}(x))^2)}.$$

The outgoing neighbor distribution is defined similarly, but using the values in layer  $l + 1$ . If  $v_{i,l}$  is in the input layer, then the input neighbor distribution is just a mass of 1 on  $v_{i,l}$  itself. Finally, if  $v_{i,l}$  is in the output layer, then the outgoing neighbor distribution is just a mass of 1 on  $v_{i,l}$ .

Intuitively, Definition 6 captures the fact that the neighbor distribution is determined by the value of each neuron. The inversion in Step 2 ensures that larger neuron values get a weight that is closer to 0. In turn, this ensures large neurons get a larger density after being sent through a standard normal distribution in Step 3, as proposed in prior work on normalizing neighbor distributions (Ni et al., 2019). Finally, the normalization in Step 1 just ensures that all neuron values are normalized within the range  $(0, 1]$ , to prevent numeric issues.<sup>4</sup> We emphasize again that, unlike the original neighbor distribution in Definitions 2 and 4, which is calculated based on the weights connecting those neighbors, Definition 6 weights neighbors according to their corresponding neuron (activation) values.

To further capture the effect of activation functions on an edge cost, we note that some NN connections' importance may change depending on the activations of the neurons they go in/out of. In particular, connections which go into a saturated neuron (in the case of Tanh) or in/out of a non-activated neuron (in the case of ReLU) have reduced importance for the NN's output on the current example. Thus, we define an edge's neural cost as follows.

**Definition 7 (Neural Edge Cost.)** Consider a directed edge  $(v_{i,l}, v_{j,l+1})$  between layers  $l$  and  $l + 1$  in a neural graph. Let  $\sigma \in \{\text{Tanh}, \text{ReLU}\}$  be the activation function in the corresponding NN. The neural cost of  $(v_{i,l}, v_{j,l+1})$  is defined as follows:

$$d_{ReLU}(v_{i,l}, v_{j,l+1}, x) = \frac{w_{f_\theta}(v_{i,l}, v_{j,l+1})}{\min\{\beta(l_{i,l}(x)), \beta(l_{j,l+1}(x))\}}$$

$$d_{Tanh}(v_{i,l}, v_{j,l+1}, x) = \frac{w_{f_\theta}(v_{i,l}, v_{j,l+1})}{\beta(l_{j,l+1}(x))}.$$

where  $\beta(x) = \sigma(x)/x$  is the fraction of information that passes through the activation function. If  $\beta(x) = 0$ , then the cost is set to  $\infty$ .

**Remark 4** In the case of ReLU,  $\beta(x) \in \{0, 1\}$ . In the case of Tanh,  $\beta(x) \in (0, 1]$ , where we define  $\beta(0) = 1$  in the case of Tanh since  $\lim_{x \rightarrow 0} \text{Tanh}(x)/x = 1$ .

<sup>4</sup>Note that a value of zero occurs only when  $n_{min} = 0$ ; the lower bound of the interval is effectively open to prevent division by zero – we achieve this in the implementation by adding a small positive constant.

Intuitively, Definition 7 increases an edge's cost based on the fraction of data that is removed by the activation function, which we use as a proxy for the reduced importance of the corresponding NN connection. Note that in the case of ReLU, we consider both the source and target neurons – in either case, we set the edge's cost to  $\infty$  if the ReLU is not activated. In the case of Tanh, we only consider the target neuron since even if the source neuron is saturated, that connection may still carry information.

**Remark 5** Note that the neural edge cost does not update the weights in the neural graph  $G_{f_\theta}$ , i.e., the Wasserstein distance is based purely on the neural neighbor distribution in Definition 6 and the weight function in Definition 5. The neural edge cost is used only in the neural curvature calculation in the following definition.

We are now ready to state the definition of neural curvature, which is the main technical contribution of this paper.

**Definition 8 (Neural Curvature)** Consider a trained NN  $f_\theta$ , an input example  $x$ , and the corresponding neural graph  $G_{f_\theta}(V_{f_\theta}, E_{f_\theta}, w_{f_\theta})$ . The neural curvature of an edge  $(v_{i,l}, v_{j,l+1}) \in E_{f_\theta}$  is defined as

$$\kappa_N(v_{i,l}, v_{j,l+1}, x) = \lim_{\alpha \rightarrow 1} \frac{\kappa_{\alpha,N}(v_{i,l}, v_{j,l+1}, x)}{1 - \alpha}, \quad (6)$$

where the  $\alpha$ -neural-curvature is defined as:

$$\kappa_{\alpha,N}(v_{i,l}, v_{j,l+1}, x) = 1 - \frac{W(m_{v_{i,l}}^\alpha(x), m_{v_{j,l+1}}^\alpha(x))}{d_\sigma(v_{i,l}, v_{j,l+1}, x)}. \quad (7)$$

In equation 7,  $\sigma \in \{\text{Tanh}, \text{ReLU}\}$ ; note that  $m_{v_{i,l}}^\alpha(x)$  and  $m_{v_{j,l+1}}^\alpha(x)$  are calculated according to Definition 6 and equation 3. Finally, if  $d_\sigma(v_{i,l}, v_{j,l+1}, x) = \infty$ , then  $\kappa_{\alpha,N}(v_{i,l}, v_{j,l+1}, x) = 1$  if  $l = 0$  or  $l = L - 1$ , and  $\kappa_{\alpha,N}(v_{i,l}, v_{j,l+1}, x) = 2$ , otherwise, as shown in Proposition 1 next.

Note that  $\beta(x)$  is often 0 in the case of ReLU activations, which in turn causes  $d_\sigma$  to be set to  $\infty$  in Definition 8. It turns out that, as the cost gets large, the  $\alpha$ -neural-curvature converges to either 1 or 2, depending on the edge's position in the graph, as shown next.

**Proposition 1 (Neural Curvature with Infinite Costs)** Consider a neural graph  $G_{f_\theta}$  and consider the neural curvature  $\kappa_N(v_{i,l}, v_{j,l+1}, x)$  for a directed edge  $(v_{i,l}, v_{j,l+1})$ . When the edge cost gets large, i.e.,  $d_\sigma(v_{i,l}, v_{j,l+1}, x) \rightarrow \infty$ , the neural curvature converges as follows

$$\lim_{d_\sigma(v_{i,l}, v_{j,l+1}, x) \rightarrow \infty} \kappa_N(v_{i,l}, v_{j,l+1}, x) = \begin{cases} 1, & \text{if } l = 0 \text{ or } l = L - 1, \text{ i.e., } (v_{i,l}, v_{j,l+1}) \text{ connects an input/output layer} \\ 2, & \text{otherwise.} \end{cases}$$

**Proof 1** To simplify notation, in what follows we define  $u := v_{i,l}$  and  $v := v_{j,l+1}$ .

Part 1: Edges in the input or output layers. Without loss of generality, assume  $u$  is in the input layer. This means the probability mass assigned to  $u$  equals 1. Thus, the Wasserstein distance between  $m_u^\alpha(x)$  and  $m_v^\alpha(x)$  is equal to the cost of transferring a mass of  $\alpha$  from  $u$  to  $v$  and a mass of  $(1 - \alpha)$  from  $u$  to  $\Gamma^{\text{out}}(v)$ , i.e., the neighbors of  $v$  in layer 2. When  $d_\sigma$  is sufficiently large, all shortest paths from  $u$  to  $\Gamma^{\text{out}}(v)$  are cheaper than  $d_\sigma$ , hence:

$$W(m_u^\alpha(x), m_v^\alpha(x)) = d_\sigma(u, v, x) \alpha + r(1 - \alpha),$$

where  $r = W_{-(u,v)}(m_u^1(x), m_v^0(x))$ , i.e., the transport cost from  $u$  to  $\Gamma^{\text{out}}(v)$ , without using the edge  $(u, v)$ . Substituting into the definition of neural curvature gives

$$\begin{aligned} \kappa_N(u, v, x) &= \lim_{\alpha \rightarrow 1} \frac{1 - \frac{W(m_u^\alpha(x), m_v^\alpha(x))}{d_\sigma(u, v, x)}}{1 - \alpha} \\ &= \lim_{\alpha \rightarrow 1} \frac{1 - \alpha - \frac{r(1 - \alpha)}{d_\sigma(u, v, x)}}{1 - \alpha} \\ &= \lim_{\alpha \rightarrow 1} \left( 1 - \frac{r(1 - \alpha)}{d_\sigma(u, v, x)(1 - \alpha)} \right) \\ &= 1 - \frac{r}{d_\sigma(u, v, x)}. \end{aligned}$$

Taking the limit yields

$$\lim_{d_\sigma(u,v,x) \rightarrow \infty} \kappa_N(u, v, x) = 1.$$

Part 2: Edges in hidden layers. *In this case, as  $d_\sigma$  gets large, the Wasserstein distance simplifies to three terms: i) the cost of transporting a mass of  $(1 - \alpha)$  from  $\Gamma^{in}(u)$  to  $v$ , without using edge  $(u, v)$ ; ii) the cost of transporting a mass of  $\alpha - (1 - \alpha)$  from  $u$  to  $v$ ; (iii) and the cost of transporting a mass of  $(1 - \alpha)$  from  $u$  to  $\Gamma^{out}(v)$ , without using edge  $(u, v)$ . Thus, the Wasserstein distance can be expressed as*

$$W(m_u^\alpha(x), m_v^\alpha(x)) = r_1(1 - \alpha) + d_\sigma(u, v, x)(\alpha - (1 - \alpha)) + r_2(1 - \alpha),$$

where  $r_1 = W_{-(u,v)}(m_u^0(x), m_v^1(x))$  and  $r_2 = W_{-(u,v)}(m_u^1(x), m_v^0(x))$ , with  $W_{-(u,v)}$  interpreted as in Part 1. Substituting into the neural curvature definition yields

$$\begin{aligned} \kappa_N(u, v, x) &= \lim_{\alpha \rightarrow 1} \frac{1 - \frac{W(m_u^\alpha(x), m_v^\alpha(x))}{d_\sigma(u, v, x)}}{1 - \alpha} \\ &= \lim_{\alpha \rightarrow 1} \frac{1 + 1 - 2\alpha - \frac{(r_1 + r_2)(1 - \alpha)}{d_\sigma(u, v, x)}}{1 - \alpha} \\ &= \lim_{\alpha \rightarrow 1} \left( \frac{2(1 - \alpha)}{1 - \alpha} - \frac{r_1 + r_2}{d_\sigma(u, v, x)} \right) \\ &= 2 - \frac{r_1 + r_2}{d_\sigma(u, v, x)}. \end{aligned}$$

Taking the limit gives

$$\lim_{d_\sigma(u,v,x) \rightarrow \infty} \kappa_N(u, v, x) = 2.$$

### 5.3 Treatment of Convolutional Parameters

Note that each NN convolutional parameter results in multiple edges in the corresponding neural graph. While our neural curvature framework can, in principle, compute a separate curvature value for each edge in a convolutional layer (i.e., each connection between input and output feature maps at specific spatial locations), we adopt a simplified approach so as to rank NN parameters as opposed to resulting edges in the neural graph.

Specifically, for each convolutional parameter, we take the minimum curvature across all edges corresponding to a given parameter, capturing the strongest contribution that the parameter can make to information propagation. This approach allows us to assign a single curvature value to each NN parameter such that we can compare convolutional and fully-connected parameters.

### 5.4 NN Connection Ranking

Given Definition 8, we rank NN connections according to the minimum neural curvature value that each edge attains across all examples in the dataset. Specifically, we combine all the edges and construct the curvature set as follows:

$$\mathcal{C}_{full} = \left\{ \min_{x \in \mathcal{D}} \min_{(u,v) \in \mathcal{E}(p)} \kappa_N(u, v, x) \mid p \in \theta \right\},$$

where  $\mathcal{E}(p)$  denotes all edges in the neural graph that correspond to NN parameter  $p$ . We interpret the curvature values such that edges with the most negative curvature are considered the most important for information propagation, while edges with the most positive curvature are considered the least important. Accordingly,  $\mathcal{C}_{full}$  can be ranked from highest to lowest curvature to prioritize the least critical connections, as summarized in Algorithm 1.

In the next section, we evaluate the importance of these two edge sets and provide strong evidence that negative-curvature edges carry the essential NN data flow, whereas the majority of positive edges have no impact on NN performance.

## 6 Neural Ricci Curvature Implementation

The overall implementation of neural Ricci curvature is based on the `GraphRicciCurvature` library (Ni et al., 2019), which efficiently handles graph construction and curvature computations on the CPU using Python. To better integrate with PyTorch and leverage GPU acceleration, which is widely used in NN training, we reimplemented parts of the library in PyTorch, enabling curvature computations directly on the GPU. Specifically, while the Wasserstein distance computation still needs to be performed on the CPU (effectively solving a linear program for each edge), the shortest-path calculation can be performed on the GPU, as follows.

Since the neural graph  $G_{f_\theta}$  is directed and has a layer-by-layer structure, we can apply dynamic programming to compute the shortest-path distances required for the Wasserstein distance in the curvature calculation. For a neural graph  $G_{f_\theta}$ , we define an adjacency cost matrix of layer  $l$  as a matrix  $C^{(l)}$ , where each entry  $C_{ij}^{(l)} = w_{f_\theta}(v_{i,l-1}, v_{j,l})$  is the cost associated with the edge from neuron  $v_{i,l}$  to neuron  $v_{j,l+1}$ .

Let  $D^{(k,l)}$  denote the minimum accumulated path cost between neurons in layer  $t$  and neurons in layer  $l$ , for any  $k < l$ . For  $l = k + 1$ , we set  $D^{(k,l)} = C^{(l)}$ . For  $l > k + 1$ , the distance matrix is computed via the dynamic programming recursion

$$D_{ij}^{(k,l)} = \min_m \left( D_{im}^{(k,l)} + C_{mj}^{(l)} \right).$$

This procedure can be efficiently implemented on a GPU, thereby allowing us to scale our approach to large CNNs.

Finally, to approximate the limit of the neural curvature as Definition 4, we set  $\alpha = 0.9$ . Based on our empirical observations, the  $\alpha$ -curvature does not change significantly beyond that point, so it is a good compromise between approximation accuracy and numerical stability.

## 7 Experiments

This section provides an experimental evaluation of the proposed ranking algorithm for NN connections. We aim to show that our approach can effectively identify both the least and the most important connections in a given NN. We evaluate the proposed approach on three image classification tasks, namely MNIST (LeCun et al., 1998), CIFAR-10, and CIFAR-100 (Krizhevsky & Hinton, 2009) datasets. All experiments were run on a 95-core machine with an NVIDIA A40 unit.

### 7.1 Experiment Design

Given a trained model  $f_\theta$  and a calibration dataset  $\mathcal{D}$ , the evaluation has the following structure:

- Generate a sorted edge set following Algorithm 1.
- For each model, compute curvature edge sets using a validation subset: 10 examples (one per label) for CIFAR-10, 100 examples (one per label) for CIFAR-100, and 100 examples (10 per label) for MNIST.
- Evaluate the impact of the ordered edge set via an edge removal experiment: starting from the full network, progressively remove edges according to the sorted set and measure the corresponding effect on test accuracy.
- For evaluation purposes, we compare our method with standard pruning techniques, namely magnitude-based pruning (Han et al., 2015), SNIP (Lee et al., 2018), and SynFlow (Tanaka et al., 2020) on the pre-trained models (without re-training); the goal is to demonstrate that our approach is more effective at identifying the least- and most-important NN connections. For SynFlow, we perform 100 pruning iterations, while for magnitude-based pruning and SNIP, we perform a single iteration, following the protocol described in Tanaka et al. (2020).<sup>5</sup>

<sup>5</sup>We use the same implementation as provided in <https://github.com/ganguli-lab/Synaptic-Flow> (Tanaka et al., 2020).

**Model description.** To demonstrate the effectiveness of our method, we train a variety of NN architectures with three different losses and two activation functions. On MNIST, we use a CNN architecture, based on the LeNet5 architecture (LeCun et al., 1998); on CIFAR-10 and CIFAR-100, we trained a CNN that adopts a VGG-style convolutional architecture (Simonyan & Zisserman, 2015), with the main differences being the removal of padding in convolutional layers, pooling after blocks, and a reduction in fully connected layer size. We refer to this variant as VGG9-lite. In each task, the training set is separated into a 50k training set and a 10k validation set. Three different training schemes are applied: 1) cross-entropy (CE) loss; 2) cross-entropy loss with weight decay regularization (WD); 3) adversarial training (AT) (Madry et al., 2017a), as well as two activation functions: 1) ReLU and 2) Tanh. In total, we have trained six models on MNIST, six models on CIFAR-10, as well as one model on CIFAR-100 (ReLU). All model hyperparameters are shown in the Appendix A.1. We emphasize that our results are equally effective across all considered combinations.

## 7.2 Pruning-Based Comparison

For evaluation purposes, we compare our approach with standard pruning techniques, including magnitude-based pruning (Han et al., 2015), SNIP (Lee et al., 2018), and SynFlow (Tanaka et al., 2020). Magnitude-based pruning is based on the assumption that the smallest-magnitude weights are least important for a NN’s performance; SNIP (Lee et al., 2018) and SynFlow (Tanaka et al., 2020) are based on the weight gradient and assume that the weights with the smallest gradient scores are least important. While SNIP and SynFlow were originally introduced for pruning at initialization, they can similarly be applied to trained models. In the pruning procedure for the various methods, edges are removed based on weight magnitudes, SNIP/SynFlow scores or neural curvature (in our case), and the model’s performance is evaluated after each pruning step. We demonstrate that our method is able to identify the least important NN connections more effectively, as it can integrate local structure information and capture relationships across layers.

The results are presented in Figure 3. On MNIST, the network size is modest and therefore less over-parameterized, leaving little redundancy to exploit. Consequently, pruning performance saturates, and our method performs comparably to magnitude pruning. By comparison, SNIP and SynFlow consistently underperform across all model architectures. Note that SynFlow was not designed to handle Tanh activations – we include its performance in all graphs for completeness.

On CIFAR-10, although pruning-based methods generally provide a meaningful separation between more and less important connections, our results indicate that our method is substantially more effective across most model configurations, particularly for CE- and AT-trained networks. Magnitude pruning performs well for WD-regularized models, likely because weight decay naturally pushes many parameters toward zero, increasing the ease of identifying redundant weights; nevertheless, our method achieves performance that is comparable to magnitude pruning across all experimental settings, and a deeper discussion about the effect of the WD parameter is provided in an ablation study in Section 7.3.5. SNIP behaves similarly to a magnitude-times-gradient (a first-order Taylor approximation) criterion when used on a pre-trained network, tending to eliminate parameters that are effectively inactive but struggling to distinguish among weights that already contribute substantially to predictive performance. SynFlow, by contrast, is data-free and therefore produces pruning dynamics similar to magnitude pruning, yet it consistently initiates pruning in the largest layers. This behavior helps prevent catastrophic layer collapse but may not be effective when layer size is not perfectly correlated with the number of redundant parameters. Overall, our method achieves the best or on-par performance with state-of-the-art approaches across all training configurations and activation function regimes, highlighting its effectiveness and broad applicability.

We additionally report results on CIFAR-100, where our method outperforms all other baselines by a substantial margin, similar to the CIFAR-10 CE model. Although CIFAR-100 is similar to CIFAR-10 in structure, its input examples and associated data pathways are more complex due to the increased number of classes and visual variability. Consequently, despite the training protocol including weight decay, our method remains more effective than magnitude pruning, and continues to consistently outperform SNIP and SynFlow. This observation provides further evidence that the algorithm proposed in this work is able to identify the primary data flow paths within NNs.

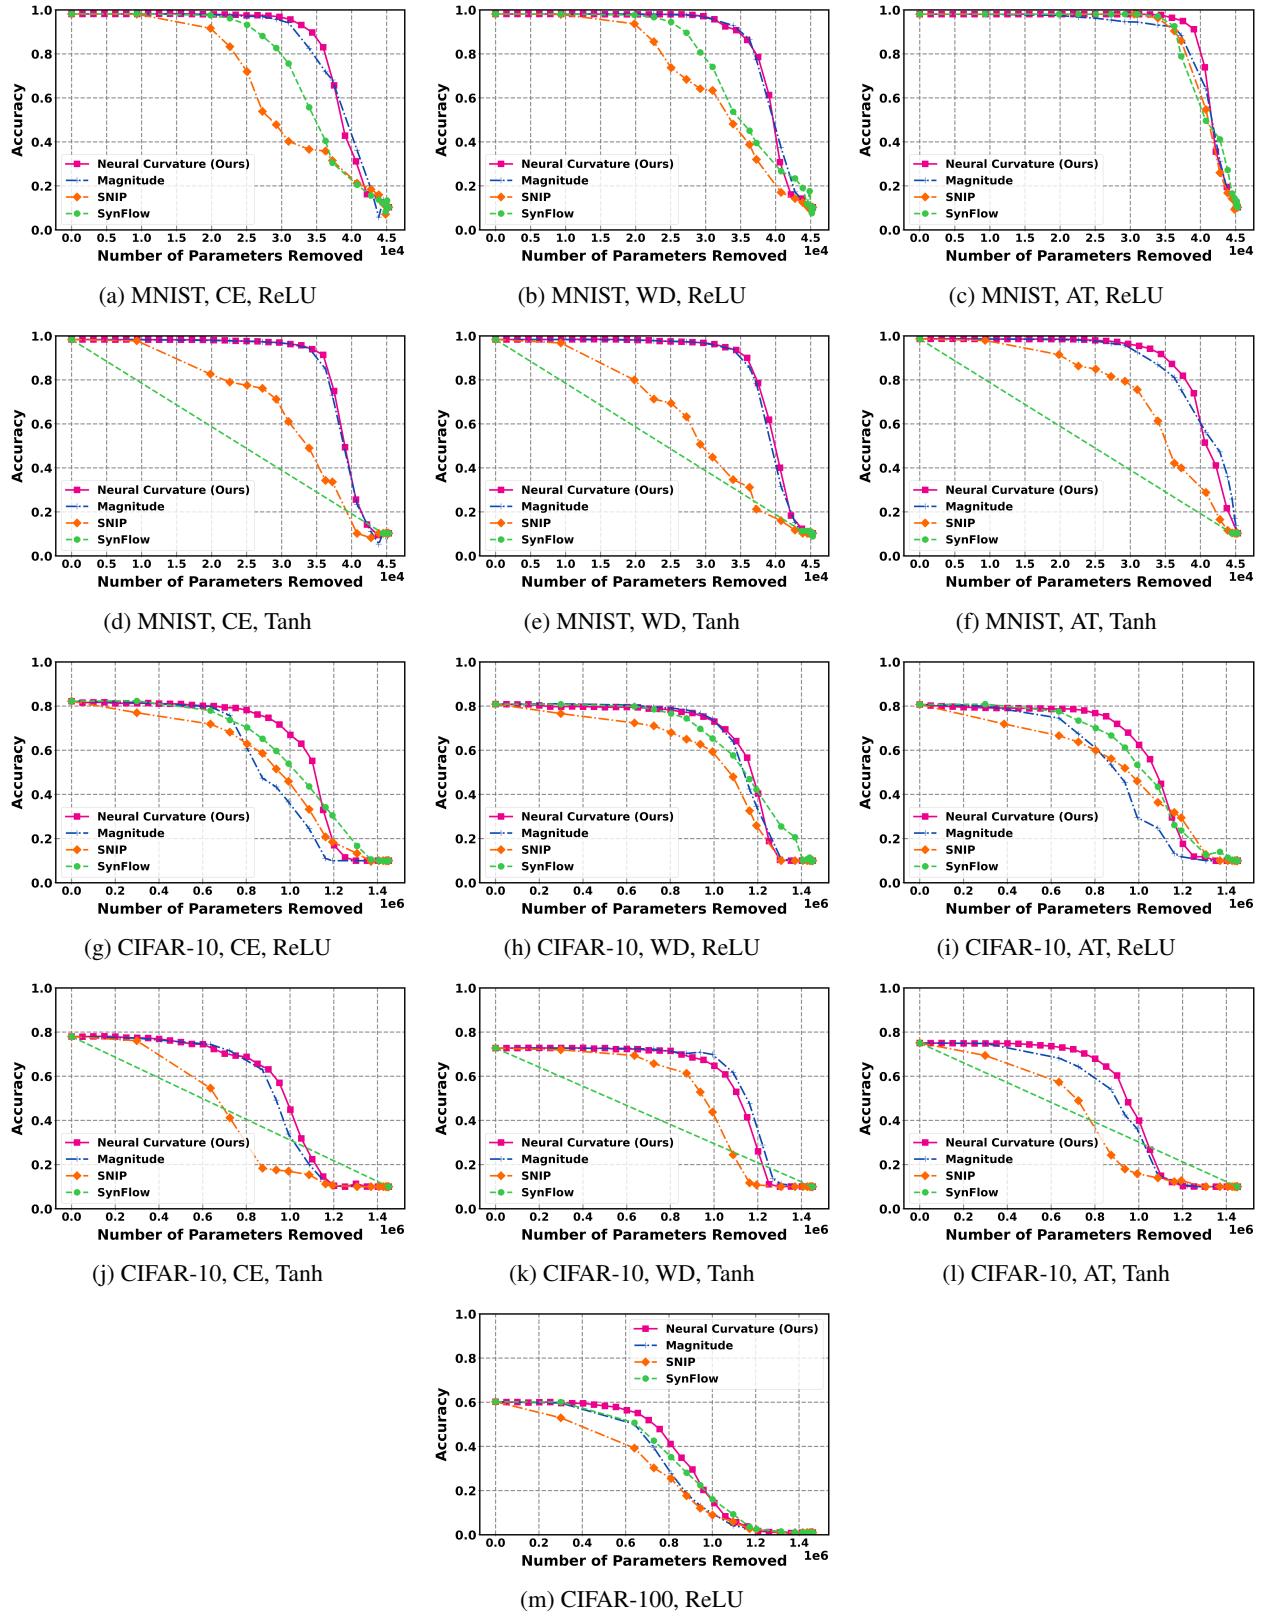


Figure 3: Edge removal evaluation on MNIST, CIFAR-10, and CIFAR-100. Each subfigure shows a comparison of our neural curvature algorithm with magnitude, SNIP, and SynFlow pruning methods.

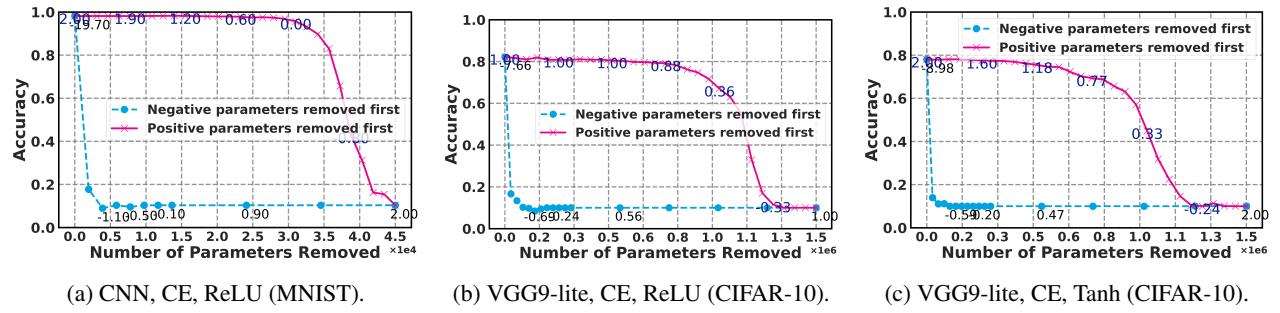


Figure 4: Edge removal evaluation on MNIST and CIFAR-10 using our neural curvature algorithm. Each subfigure shows a comparison of the impact of removing edges in order of positive-curvature-first versus negative-curvature-first. Numbers on the curves indicate the minimum curvature per removed edge over the validation set.

### 7.3 Ablation Study

For a deeper analysis of the proposed method, we also provide an ablation study where we investigate the results per layer, the effect of the proposed neural modules for neural curvature calculation, the data requirements of our method, as well as the effect of the WD hyper-parameter.

#### 7.3.1 Effect of Negative-Curvature Edge Removal

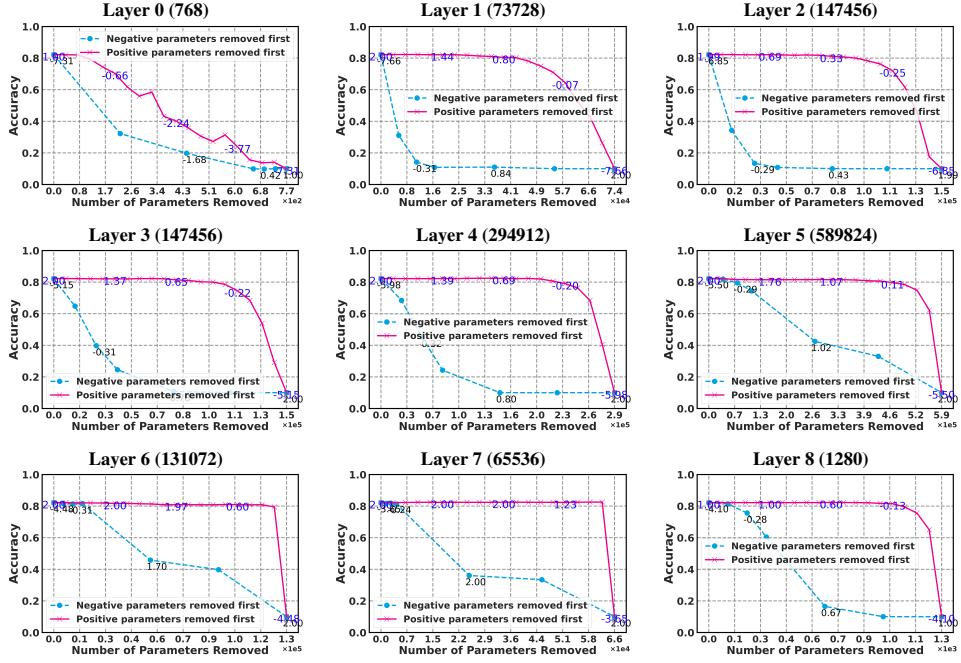
While in the previous section we demonstrated that our method is very effective at identifying the least important NN connections as the ones with the highest neural curvature (positive-first pruning), we now also show that negative-curvature edges tend to be the most important. We show the positive-first and negative-first edge removal results for one MNIST, CE model (ReLU), and two CIFAR-10 CE models (ReLU and Tanh). As shown in Figure 4, the results show a very convincing tendency that the negative edges constitute the primary data flow, whereas the majority of the positive edges have no impact on the NN accuracy. In particular, the accuracy degrades sharply when negative edges are removed first, while it remains almost the same when the majority of positive edges are removed first. We emphasize that the results show the same trends across all models and training algorithms – please consult the Appendix A.5 for the full results.

This demonstrates both the effectiveness of our approach as well as the fact that standard models do not utilize the vast majority of the available weights. In addition, the curvature annotations above the pruning curves reveal a consistent trend: model accuracy decreases as the removed edges include those with curvature values closer to zero. This provides further evidence that highly negative-curvature edges are structurally important for maintaining predictive performance. For example, in the CNN model shown in Figure 4a, accuracy remains stable while pruning edges with positive curvature, and begins to noticeably drop only once the removed edges include those with curvature values near or below zero. A similar pattern appears in the VGG9-lite experiments (Figures 4b and 4c), where the curvature value at which accuracy first declines is substantially smaller than the value associated with the initial flat region of the curve. Moreover, as illustrated by the blue curve, pruning edges with strongly negative curvature can abruptly collapse accuracy, effectively breaking the network. These observations indicate that curvature is not merely correlated with edge importance, but serves as a reliable and fine-grained signal for determining which edges should be preserved during pruning.

#### 7.3.2 Per-Layer Edge Removal Results

Figure 5 presents the per-layer removal experiments for the VGG9-lite ReLU model trained with CE (the same model as in Figure 3g) on CIFAR-10, comparing our algorithm against magnitude-based pruning. An interesting observation is that when pruning is restricted to a single layer at a time, our method behaves similarly to magnitude pruning. This suggests that within a given layer, the largest-magnitude weights are typically the most important. However, magnitude alone cannot be meaningfully compared across layers, which explains why our method demonstrates significantly stronger performance when pruning the full network: the neural curvature provides an effective normalization across layers and enables globally coherent pruning decisions.

## Neural Curvature Method (Ours)



## Magnitude-based

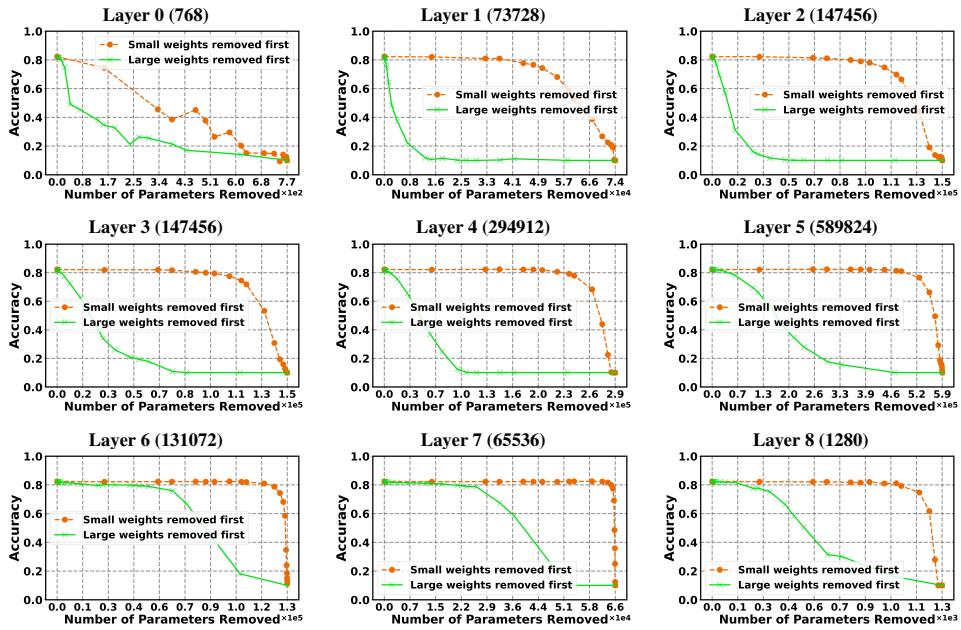


Figure 5: Ablation experiment on per-layer edge removal. Each subfigure shows the edge removal analysis for each layer based on neural curvature value (top 9 figures) and magnitude-based pruning (bottom 9 figures) for the VGG9-lite, CE ReLU configuration. The number above the figure is the total number of parameters in that layer.

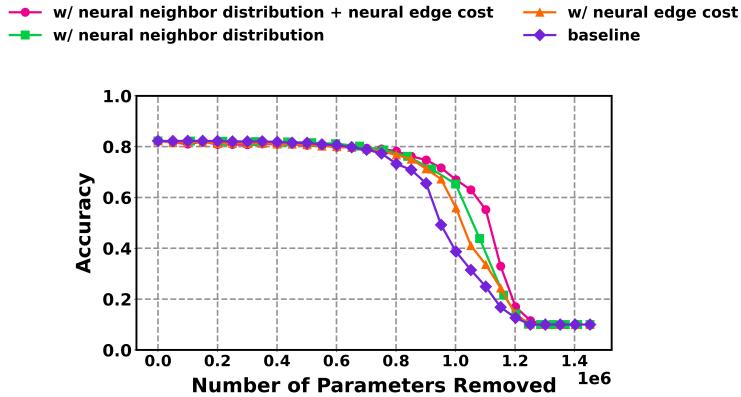


Figure 6: Ablation experiment for edge removal with different neural curvature calculation modules (VGG9-lite, CE ReLU, on CIFAR-10). Baseline is neural graph + the graph-based Ricci curvature (Lin et al., 2011).

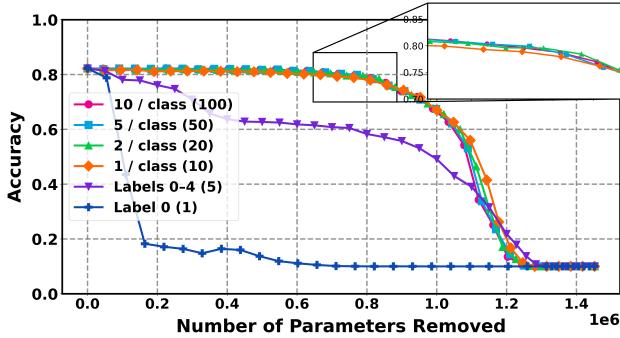


Figure 7: Ablation experiment for edge removal using different numbers of calibration examples (VGG9-lite, CE ReLU, on CIFAR-10).

Moreover, the curvature annotations along our pruning curves reveal that model accuracy typically begins to degrade only once the edges being removed have curvature values near or below zero. This layer-wise pattern again validates curvature as a meaningful importance indicator. Finally, this experiment highlights that our approach is capable of reliably identifying both the most and least important edges across the network. Additional per-layer edge removal results on four more models are provided in Appendix A.6, further demonstrating the generality and effectiveness of our method.

### 7.3.3 Effect of Neural Modules on Curvature Calculation

Our algorithm for neural curvature computation is largely aligned with the standard formulation Lin et al., 2011, with two modifications. First, we use the neural neighbor distribution as in Definition 6. Second, we apply a neural edge cost to factor in the activation function effect as Definition 7. In particular, for ReLU-based models, we assign an infinite cost to a neural edge when either the source or target node has  $\beta = 0$ . To assess the importance of these two components, we perform an ablation study (Figure 6). As shown in the results, the baseline Lin et al., 2011, which is static and only uses the weight magnitudes of the NN, is significantly outperformed by either modification based on incorporating the input examples in addition to the NN structure.

### 7.3.4 Data Requirements

While the baselines we compare against (e.g., SynFlow) are data-free, our curvature computation requires data to estimate neuron activations. It is therefore important to assess how much data is needed for our method to produce reliable pruning decisions. Figure 7 provides an illustration of the data requirements of our method, for the same model as in Figure 5. We compare the results with using only one example of label 0, five examples of labels 0 – 4, 10

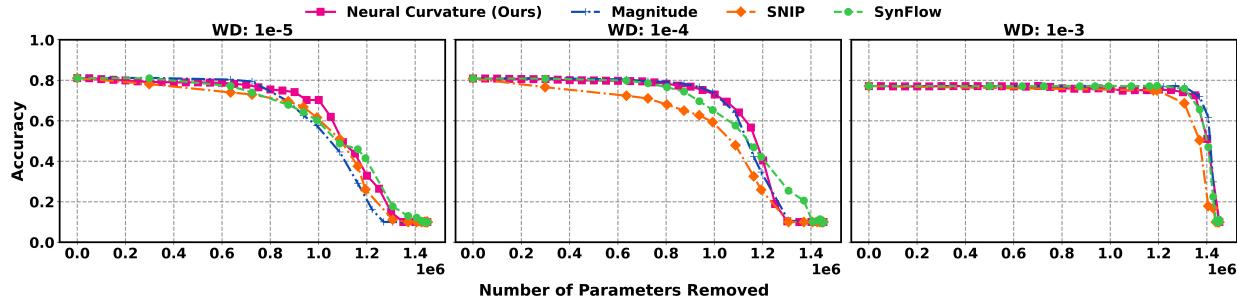


Figure 8: Full model edge removal comparison between magnitude-based and ours (VGG9-lite, WD ReLU, on CIFAR-10). Models are trained with different WD parameters.

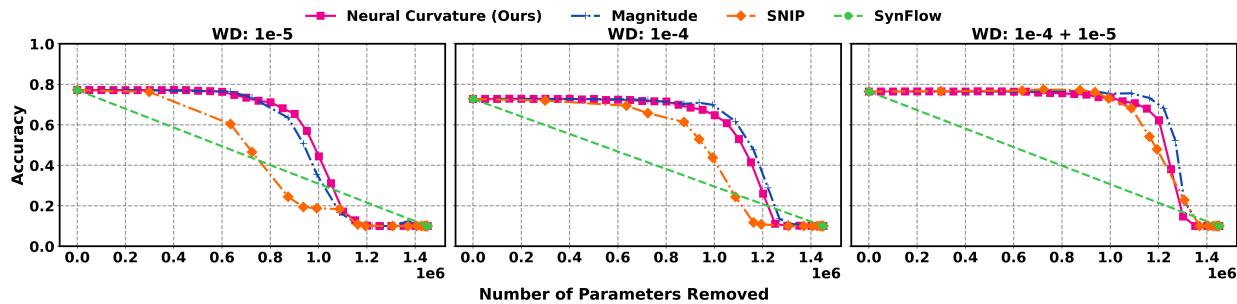


Figure 9: Full model edge removal comparison between magnitude-based and ours (VGG9-lite, WD Tanh, on CIFAR-10). Models are trained with different WD parameters.

examples (one example per label), 20 examples (two examples per label), 50 examples (five examples per label), and 100 examples (10 examples per label). As shown in the figure, one example per label is sufficient to recover most of the pruning behavior observed with the full dataset. While adding more examples per label enables a larger number of parameters to be pruned, it leads to a faster decline near the end. Notably, our method is highly data-efficient and can effectively separate the two sets even with only 10 calibration examples. Increasing the number of examples does not substantially change the results and, in rare cases, may slightly alter the separation, likely due to noise in the activation statistics that perturbs the curvature distribution.

### 7.3.5 Effect of Weight Decay on Neural Curvature and Magnitude

Weight decay (WD) is known to influence the distribution of weights in neural networks, which can in turn affect pruning behavior. The results in the previous sections showed that when WD training is used, magnitude-based pruning improves and can approach or even exceed the performance of our method. In this section, we systematically investigate models trained with different WD hyper-parameters to examine how WD affects pruning outcomes and to further illustrate its impact on different pruning methods. We compare full-model edge removal across models trained with different weight decay (WD) hyper-parameters on CIFAR-10. Figure 8 shows results for VGG9-lite ReLU trained with WD parameters  $1e-5$ ,  $1e-4$ , and  $1e-3$ , while Figure 9 presents results for VGG9-lite Tanh, trained with WD parameters  $1e-5$ ,  $1e-4$ , and a mixed  $1e-4/1e-5$  setup to achieve higher accuracy. The results indicate that larger WD values substantially improve the effectiveness of magnitude pruning. Effectively, WD helps eliminate unimportant weights by pushing them toward very small values, thus compressing the valid data paths within the network. As the WD parameter increases, magnitude pruning becomes more efficient; however, our method still achieves comparable performance on ReLU models and only slightly underperforms at the end for Tanh models with larger WD hyper-parameters. Finally, it is important to note that increasing the WD hyper-parameter typically results in lower accuracy and requires significantly more training effort. Thus, although it is possible to obtain models which are effectively pruned through magnitude-based pruning, one may need to trade off other model properties, such as accuracy or computation.

## 8 Discussion and Conclusion

This paper proposed a differential-geometry-based approach to identify the main NN data flow. In particular, we introduced the notion of neural curvature that can be used to rank NN connections and separate them according to their importance. We provided an extensive evaluation over three image classification datasets.

The symbolic nature of our approach provides a well-understood tool that can be used in a variety of applications such as robustness analysis, model repair and alignment. The method is general and can be used in a number of other domains, including natural language processing, reinforcement learning and magnetic resonance image analysis. As immediate future work, we intend to apply this technique to the adversarial training problem, e.g., by focusing on the edges that are currently not used by the NN. A similar approach can be used for model repair and alignment, e.g., by only re-training connections that are not currently used.

In terms of scalability, the main challenge is calculating the Wasserstein distance, which requires solving a linear program. At the same time, this computation can be greatly simplified in a variety of cases, e.g., in input and output layers where the Wasserstein distance reduces to an inner product. For further improvements in scalability, we will investigate approximations to the Wasserstein distance, which can be computed on the GPU, e.g., through Sinkhorn approximation (Luise et al., 2018).

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Alessandro Achille, Giovanni Paolini, and Stefano Soatto. Where is the information in a deep neural network? *arXiv preprint arXiv:1905.12213*, 2019.
- S. Bochner. Vector fields and Ricci curvature. *Bulletin of the American Mathematical Society*, 52(9):776 – 797, 1946.
- Ahmad Chaddad, Behnaz Naisiri, Marco Pedersoli, Eric Granger, Christian Desrosiers, and Matthew Toews. Modeling information flow through deep neural networks. *arXiv preprint arXiv:1712.00003*, 2017.
- Brian Davis, Umang Bhatt, Kartikeya Bhardwaj, Radu Marculescu, and José Moura. Nif: A framework for quantifying neural information flow in deep networks. In *Proc. Workshop Netw. Interpretability Deep Learn. AAAI Conf. Artif. Intell.(AAAI)*, pp. 1–4, 2019.
- Brian Davis, Umang Bhatt, Kartikeya Bhardwaj, Radu Marculescu, and José MF Moura. On network science and mutual information for explaining deep neural networks. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 8399–8403. IEEE, 2020.
- Lucio Dery, Steven Kolawole, Jean-François Kagy, Virginia Smith, Graham Neubig, and Ameet Talwalkar. Everybody prune now: Structured pruning of llms with only forward passes. *arXiv preprint arXiv:2402.05406*, 2024.
- 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.
- Marzieh Eidi and Jürgen Jost. Ollivier ricci curvature of directed hypergraphs. *Scientific Reports*, 10(1):12466, 2020.
- Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. *arXiv preprint arXiv:1803.03635*, 2018.
- Lei Gao, Xingquan Liu, Yu Liu, Pu Wang, Min Deng, Qing Zhu, and Haifeng Li. Measuring road network topology vulnerability by ricci curvature. *Physica A: Statistical Mechanics and Its Applications*, 527, 2019a.
- Lei Gao, Xingquan Liu, Yu Liu, Pu Wang, Min Deng, Qing Zhu, and Haifeng Li. Measuring road network topology vulnerability by ricci curvature. *Physica A: Statistical Mechanics and Its Applications*, 527:121071, 2019b.

Ziv Goldfeld, Ewout Van Den Berg, Kristjan Greenewald, Igor Melnyk, Nam Nguyen, Brian Kingsbury, and Yury Polyanskiy. Estimating information flow in deep neural networks. In *International Conference on Machine Learning*, pp. 2299–2308. PMLR, 2019.

Song Han, Jeff Pool, John Tran, and William Dally. Learning both weights and connections for efficient neural network. *Advances in neural information processing systems*, 28, 2015.

Yihui He, Ji Lin, Zhijian Liu, Hanrui Wang, Li-Jia Li, and Song Han. Amc: Automl for model compression and acceleration on mobile devices. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 784–800, 2018.

Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks, 2018. URL <https://arxiv.org/abs/1610.02136>.

Jyh-Jing Hwang, Runsheng Xu, Hubert Lin, Wei-Chih Hung, Jingwei Ji, Kristy Choi, Di Huang, Tong He, Paul Covington, Benjamin Sapp, et al. Emma: End-to-end multimodal model for autonomous driving. *arXiv preprint arXiv:2410.23262*, 2024.

Pejman Khadivi, Ravi Tandon, and Naren Ramakrishnan. Flow of information in feed-forward deep neural networks. *arXiv preprint arXiv:1603.06220*, 2016.

Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical Report 0, University of Toronto, Toronto, Ontario, 2009. URL <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>.

Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

Namhoon Lee, Thalaiyasingam Ajanthan, and Philip HS Torr. Snip: Single-shot network pruning based on connection sensitivity. *arXiv preprint arXiv:1810.02340*, 2018.

Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. *arXiv preprint arXiv:1608.08710*, 2016.

Yawei Li, Shuhang Gu, Kai Zhang, Luc Van Gool, and Radu Timofte. Dhp: Differentiable meta pruning via hypernetworks. In *European conference on computer vision*, pp. 608–624. Springer, 2020.

Yong Lin, Linyuan Lu, and Shing-Tung Yau. Ricci curvature of graphs. *Tohoku Mathematical Journal, Second Series*, 63(4):605–627, 2011.

Liyang Liu, Shilong Zhang, Zhanghui Kuang, Aojun Zhou, Jing-Hao Xue, Xinjiang Wang, Yimin Chen, Wenming Yang, Qingmin Liao, and Wayne Zhang. Group fisher pruning for practical network compression. In *International Conference on Machine Learning*, pp. 7021–7032. PMLR, 2021.

Yang Liu, Chuan Zhou, Shirui Pan, Jia Wu, Zhao Li, Hongyang Chen, and Peng Zhang. Curvdrop: A ricci curvature based approach to prevent graph neural networks from over-smoothing and over-squashing. In *Proceedings of the ACM Web Conference 2023*, pp. 221–230, 2023.

Giulia Luise, Alessandro Rudi, Massimiliano Pontil, and Carlo Ciliberto. Differential properties of sinkhorn approximation for learning with wasserstein distance. *Advances in Neural Information Processing Systems*, 31, 2018.

Xinyin Ma, Gongfan Fang, and Xinchao Wang. Llm-pruner: On the structural pruning of large language models. *Advances in neural information processing systems*, 36:21702–21720, 2023.

Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv:1706.06083*, 2017a.

Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017b.

- Behnaz Nasiri. *Modeling information flow through deep convolutional neural networks*. PhD thesis, École de technologie supérieure, 2020.
- Chien-Chun Ni, Yu-Yao Lin, Jie Gao, Xianfeng David Gu, and Emil Saucan. Ricci curvature of the internet topology. In *2015 IEEE conference on computer communications (INFOCOM)*, pp. 2758–2766. IEEE, 2015.
- Chien-Chun Ni, Yu-Yao Lin, Feng Luo, and Jie Gao. Community detection on networks with ricci flow. *Scientific reports*, 9(1):1–12, 2019.
- Manuel Nonnenmacher, Thomas Pfeil, Ingo Steinwart, and David Reeb. Sosp: Efficiently capturing global correlations by second-order structured pruning. *arXiv preprint arXiv:2110.11395*, 2021.
- Yann Ollivier. Ricci curvature of metric spaces. *Comptes Rendus Mathematique*, 345(11):643–646, 2007.
- Yann Ollivier. Ricci curvature of Markov chains on metric spaces. *J. Functional Analysis*, 256(3):810–864, 2009.
- Yongming Rao, Jiwen Lu, Ji Lin, and Jie Zhou. Runtime network routing for efficient image classification. *IEEE transactions on pattern analysis and machine intelligence*, 41(10):2291–2304, 2018.
- Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do imagenet classifiers generalize to imagenet? In *International conference on machine learning*, pp. 5389–5400. PMLR, 2019.
- Loqman Salamatian, Scott Anderson, Joshua Matthews, Paul Barford, Walter Willinger, and Mark Crovella. Curvature-based analysis of network connectivity in private backbone infrastructures. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 6(1):1–32, 2022.
- Areejit Samal, RP Sreejith, Jiao Gu, Shiping Liu, Emil Saucan, and Jürgen Jost. Comparative analysis of two discretizations of ricci curvature for complex networks. *Scientific reports*, 8(1):8650, 2018.
- Michael Santacroce, Zixin Wen, Yelong Shen, and Yuanzhi Li. What matters in the structured pruning of generative language models? *arXiv preprint arXiv:2302.03773*, 2023.
- Ravid Shwartz-Ziv and Naftali Tishby. Opening the black box of deep neural networks via information. *arXiv preprint arXiv:1703.00810*, 2017.
- Jayson Sia, Edmond Jonckheere, and Paul Bogdan. Ollivier-ricci curvature-based method to community detection in complex networks. *Scientific reports*, 9(1):9800, 2019.
- Anish K Simhal, Corey Weistuch, Kevin Murgas, Daniel Grange, Jiening Zhu, Jung Hun Oh, Rena Elkin, and Joseph O Deasy. Orco: Ollivier-ricci curvature-omics—an unsupervised method for analyzing robustness in biological systems. *Bioinformatics*, 41(3):btaf093, 2025.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition, 2015.
- Li Sun, Zhongbao Zhang, Junda Ye, Hao Peng, Jiawei Zhang, Sen Su, and Philip S Yu. A self-supervised mixed-curvature graph neural network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 4146–4155, 2022.
- Mingjie Sun, Zhuang Liu, Anna Bair, and J Zico Kolter. A simple and effective pruning approach for large language models. *arXiv preprint arXiv:2306.11695*, 2023.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- Shuhang Tan, Jayson Sia, Paul Bogdan, and Radoslav Ivanov. Analyzing neural network robustness using graph curvature. *arXiv:1312.6199*, 2024.
- Hideyori Tanaka, Daniel Kunin, Daniel L Yamins, and Surya Ganguli. Pruning neural networks without any data by iteratively conserving synaptic flow. *Advances in neural information processing systems*, 33:6377–6389, 2020.
- Naftali Tishby and Noga Zaslavsky. Deep learning and the information bottleneck principle. In *2015 ieee information theory workshop (itw)*, pp. 1–5. Ieee, 2015.

- Naftali Tishby, Fernando C Pereira, and William Bialek. The information bottleneck method. *arXiv preprint physics/0004057*, 2000.
- Tycho FA van der Ouderaa, Markus Nagel, Mart Van Baalen, Yuki M Asano, and Tijmen Blankevoort. The llm surgeon. *arXiv preprint arXiv:2312.17244*, 2023.
- Elena Voita, David Talbot, Fedor Moiseev, Rico Sennrich, and Ivan Titov. Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned. *arXiv preprint arXiv:1905.09418*, 2019.
- Shen Wang, Xiaokai Wei, Cicero Nogueira Nogueira dos Santos, Zhiguo Wang, Ramesh Nallapati, Andrew Arnold, Bing Xiang, Philip S Yu, and Isabel F Cruz. Mixed-curvature multi-relational graph neural network for knowledge graph completion. In *Proceedings of the web conference 2021*, pp. 1761–1771, 2021.
- Yaoli Wang, Zhou Huang, Ganmin Yin, Haifeng Li, Liu Yang, Yuelong Su, Yu Liu, and Xv Shan. Applying ollivier-ricci curvature to indicate the mismatch of travel demand and supply in urban transit network. *International Journal of Applied Earth Observation and Geoinformation*, 106, 2022.
- Yulong Wang, Hang Su, Bo Zhang, and Xiaolin Hu. Interpret neural networks by identifying critical data routing paths. In *proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8906–8914, 2018.
- Yulong Wang, Hang Su, Bo Zhang, and Xiaolin Hu. Interpret neural networks by extracting critical subnetworks. *IEEE Transactions on Image Processing*, 29:6707–6720, 2020.
- Xiongye Xiao, Chenyu Zhou, Heng Ping, Defu Cao, Yaxing Li, Yizhuo Zhou, Shixuan Li, and Paul Bogdan. Exploring neuron interactions and emergence in llms: From the multifractal analysis perspective. *arXiv preprint arXiv:2402.09099*, 2024.
- Taiki Yamada. The ricci curvature on directed graphs. *arXiv preprint arXiv:1602.07779*, 2016.
- Jingkang Yang, Pengyun Wang, Dejian Zou, Zitang Zhou, Kunyuan Ding, WENXUAN PENG, Haoqi Wang, Guangyao Chen, Bo Li, Yiyu Sun, Xuefeng Du, Kaiyang Zhou, Wayne Zhang, Dan Hendrycks, Yixuan Li, and Ziwei Liu. Openood: Benchmarking generalized out-of-distribution detection. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 32598–32611. Curran Associates, Inc., 2022. URL [https://proceedings.neurips.cc/paper\\_files/paper/2022/file/d201587e3a84fc4761eadc743e9b3f35-Paper-Datasets\\_and\\_Benchmarks.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/d201587e3a84fc4761eadc743e9b3f35-Paper-Datasets_and_Benchmarks.pdf).
- Mingyang Zhang, Xinyi Yu, Jingtao Rong, and Linlin Ou. Graph pruning for model compression. *Applied Intelligence*, 52(10):11244–11256, 2022.
- Chenglong Zhao, Bingbing Ni, Jian Zhang, Qiwei Zhao, Wenjun Zhang, and Qi Tian. Variational convolutional neural network pruning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 2780–2789, 2019.
- Zhao Zhong, Junjie Yan, Wei Wu, Jing Shao, and Cheng-Lin Liu. Practical block-wise neural network architecture generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2423–2432, 2018.
- Mohamed Ridha Znaidi, Jayson Sia, Scott Ronquist, Indika Rajapakse, Edmond Jonckheere, and Paul Bogdan. A unified approach of detecting phase transition in time-varying complex networks. *Scientific reports*, 13(1), 2023.

## A Appendix

### A.1 Model Training

This section presents the model training details, including the training hyperparameters and model architectures.

### A.2 Training hyper-parameters

1. Training hyperparameters of CNN on MNIST and VGG9-lite on CIFAR-10:

- learning rate: 2e-3 for all CE, WD, and AT models
- weight decay parameter for WD models: 0.0001 (unless mentioned otherwise)
- MNIST PGD parameter:  $\text{eps} = 0.1$ ,  $\alpha = 2/255$ , iterations = 40
- CIFAR-10 PGD parameter:  $\text{eps} = 2/255$ ,  $\alpha = 2/255$ , iterations = 20
- optimizer: Adam

2. Training hyperparameters of VGG9-lite on CIFAR-100

- learning rate: 0.1
- weight decay parameter: 0.0001
- optimizer: SGD (momentum = 0.9)

### A.3 CNN Architecture

The CNN model is a modified convolutional neural network inspired by the original LeNet-5 architecture. The network is designed for classification tasks and consists of two convolutional layers followed by three fully connected layers. The activation function (ReLU or Tanh) is applied after each convolution and linear transformation.

- **Input:** A grayscale image of size  $28 \times 28$  (1 channel).

- **Layer 1: Convolution**

- Convolutional layer with 6 output channels
- Kernel size:  $6 \times 6$
- Stride: 2
- Output size:  $[6 \times 12 \times 12]$
- Followed by ReLU/Tanh activation

- **Layer 2: Convolution**

- Convolutional layer with 16 output channels
- Kernel size:  $6 \times 6$
- Stride: 2
- Output size:  $[16 \times 4 \times 4]$
- Followed by ReLU/Tanh activation

- **Flattening:** The  $16 \times 4 \times 4$  output is flattened into a vector of size 256.

- **Layer 3: Fully Connected**

- Linear layer from 256 to 120 units
- Followed by ReLU/Tanh activation

- **Layer 4: Fully Connected**

- Linear layer from 120 to 84 units
- Followed by ReLU/Tanh activation

- **Layer 5: Output Layer**

- Linear layer from 84 to 10 units (representing 10 classes)

- **Output:** Raw logits for each class.

Note: The implementation uses custom CNN and linear methods for the forward computation, possibly to track graph-based properties such as edge connectivity and activation flows. Max-pooling operations are commented out in the current version.

#### A.4 VGG9-lite Architecture

The VGG9-lite model is a customized version of the VGG architecture, adapted for the CIFAR-10/CIFAR-100 dataset. It contains 6 convolutional layers, followed by 3 fully connected layers. Batch normalization and ReLU/Tanh activations are used throughout. All the convolution and linear layers are computed using custom CNN and linear functions for graph-based analysis.

- **Input:** RGB image of size  $32 \times 32$  (3 channels), normalized with CIFAR-10/CIFAR-100 dataset statistics.

- **Conv Block 1:**

- One convolutional layers:  $3 \rightarrow 64$  (kernel  $3 \times 3$ , padding 0, stride 2)
- BatchNorm + ReLU/Tanh after each
- Output size:  $64 \times 16 \times 16$

- **Conv Block 2:**

- One convolutional layers:  $64 \rightarrow 128$  (kernel  $3 \times 3$ , padding 0, stride 1)
- BatchNorm + ReLU/Tanh after each
- Output size:  $128 \times 14 \times 14$

- **Conv Block 3:**

- Two convolutional layers:  $128 \rightarrow 128$  (kernel  $3 \times 3$ , padding 0)
- First convolutional layers with a stride 1; second convolutional layers with a stride 2
- BatchNorm + ReLU/Tanh after each
- Output size:  $128 \times 5 \times 5$

- **Conv Block 4:**

- Two convolutional layers:  $128 \rightarrow 256$  (kernel  $3 \times 3$ , padding 0, stride 1)
- BatchNorm + ReLU/Tanh after each
- Output size:  $256 \times 1 \times 1$

- **Flatten:** The  $256 \times 1 \times 1$  tensor is flattened to a vector of length 256.

- **Fully Connected Layers:**

- $fc1$ : Linear layer from 256 to 512 (with ReLU/Tanh)
- $fc2$ : Linear layer from 512 to 128 (with ReLU/Tanh)
- $fc3$ : Linear layer from 128 to 10 (raw logits for 10 classes for CIFAR-10 (100 classes for CIFAR-100))

The use of custom CNN and linear functions enables explicit tracking of information flow and layer-wise graph structures, which is beneficial for graph-theoretic or curvature-based network analysis.

## A.5 Edge Removal Evaluation

In this section, we present the results for all the models, including CNN on MNIST, VGG9-lite on CIFAR-10 and CIFAR-100 (Figure 10). Each subfigure contains the full model edge removal results based on neural curvature values using positive-first and negative-first edge sets. The number on the curve shows the curvature value at that data point.

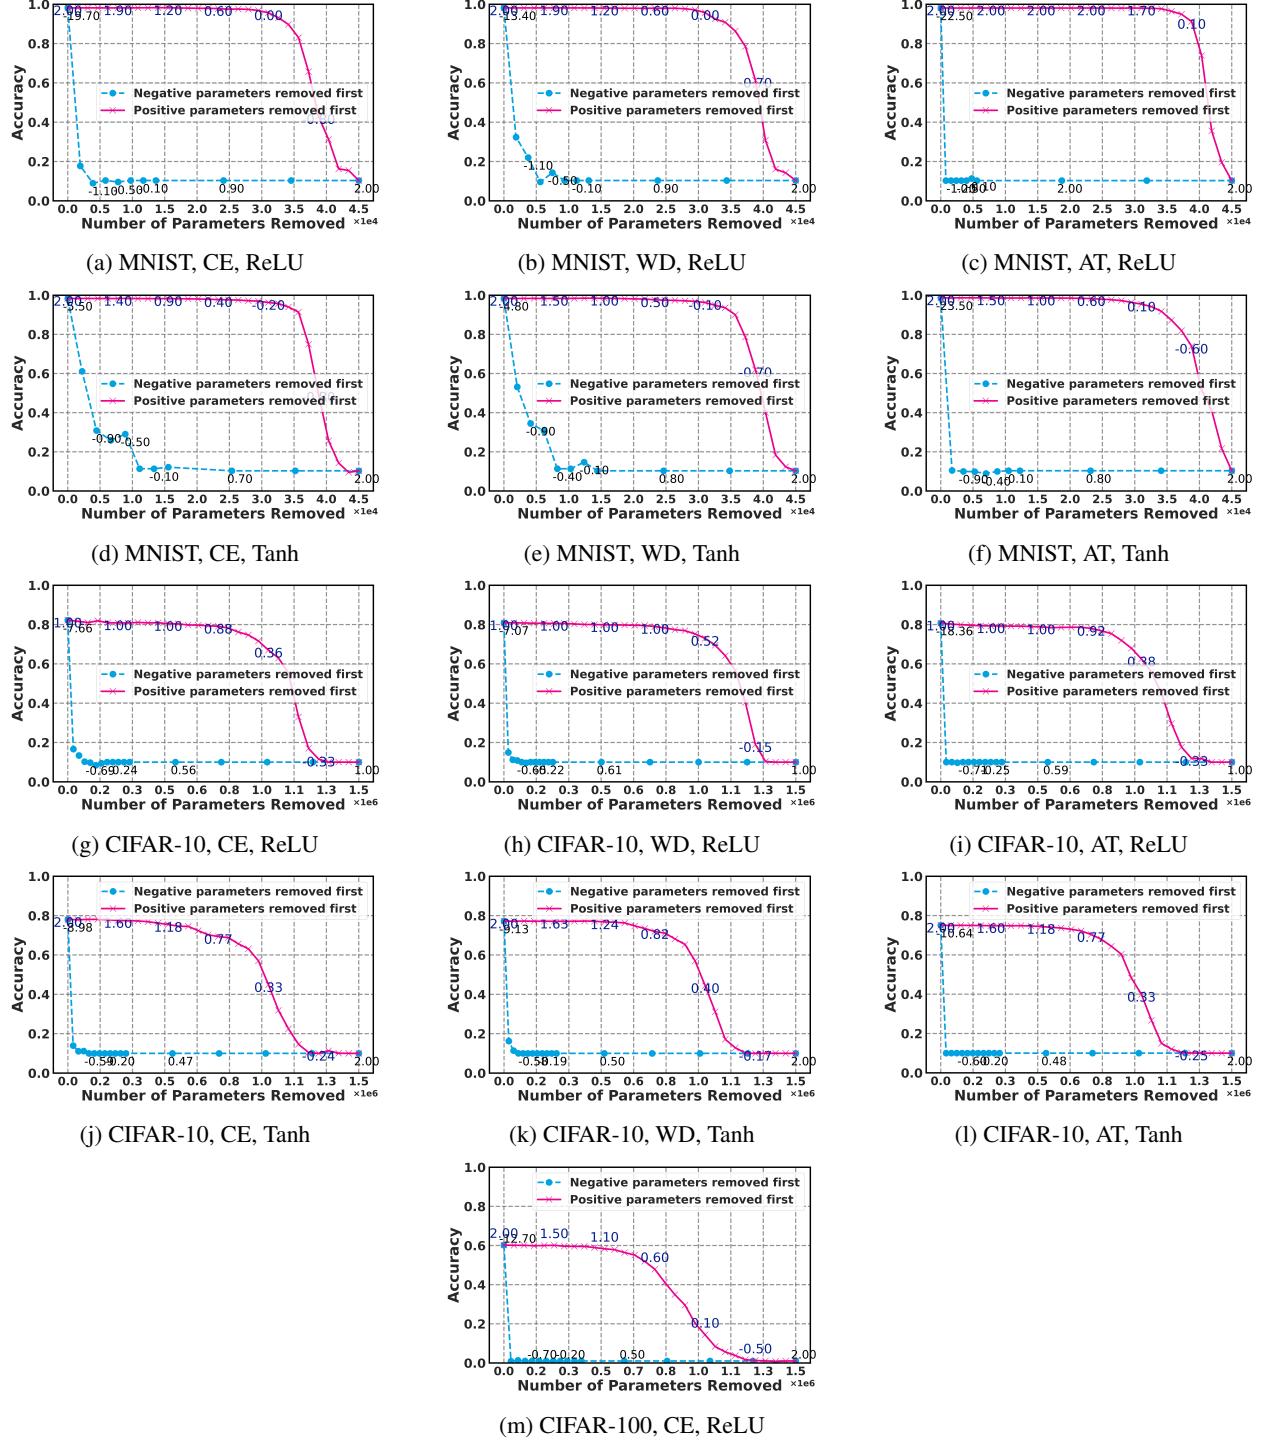


Figure 10: Edge removal evaluation on MNIST, CIFAR-10, and CIFAR-100.

## A.6 Ablation Study of Per-layer edge removal

In this section, we presented the results of the ablation study on per-layer edge removal, including CNN on MNIST (Figure 11, and Figure 12), VGG9-lite on CIFAR-10 (Figure 13, Figure 14). Each figure contains the per-layer edge removal results based on neural curvature values (top) and based on weight magnitude (bottom), for comparison.

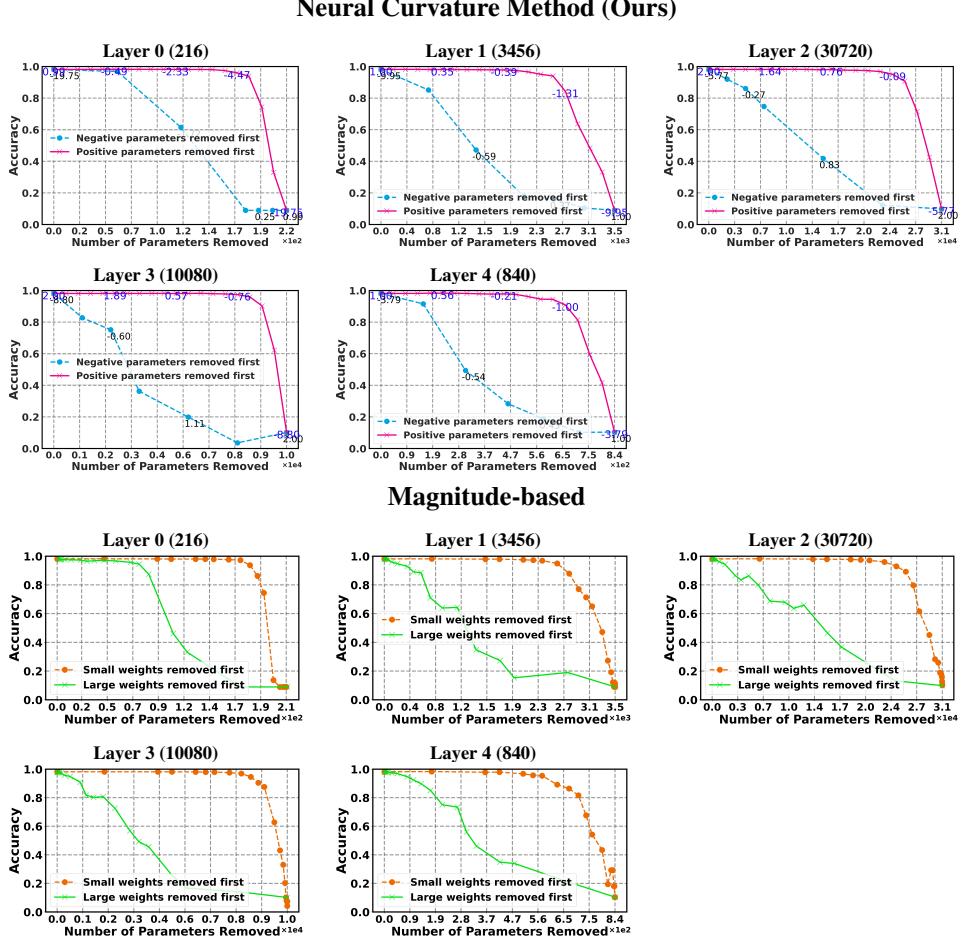


Figure 11: Ablation experiment on per-layer edge removal. Each subfigure shows the edge removal analysis for each layer based on neural curvature value (top 5 figures) and magnitude-based pruning (bottom 5 figures) for the CNN, CE ReLU configuration. The number above the figure is the total number of parameters in that layer.

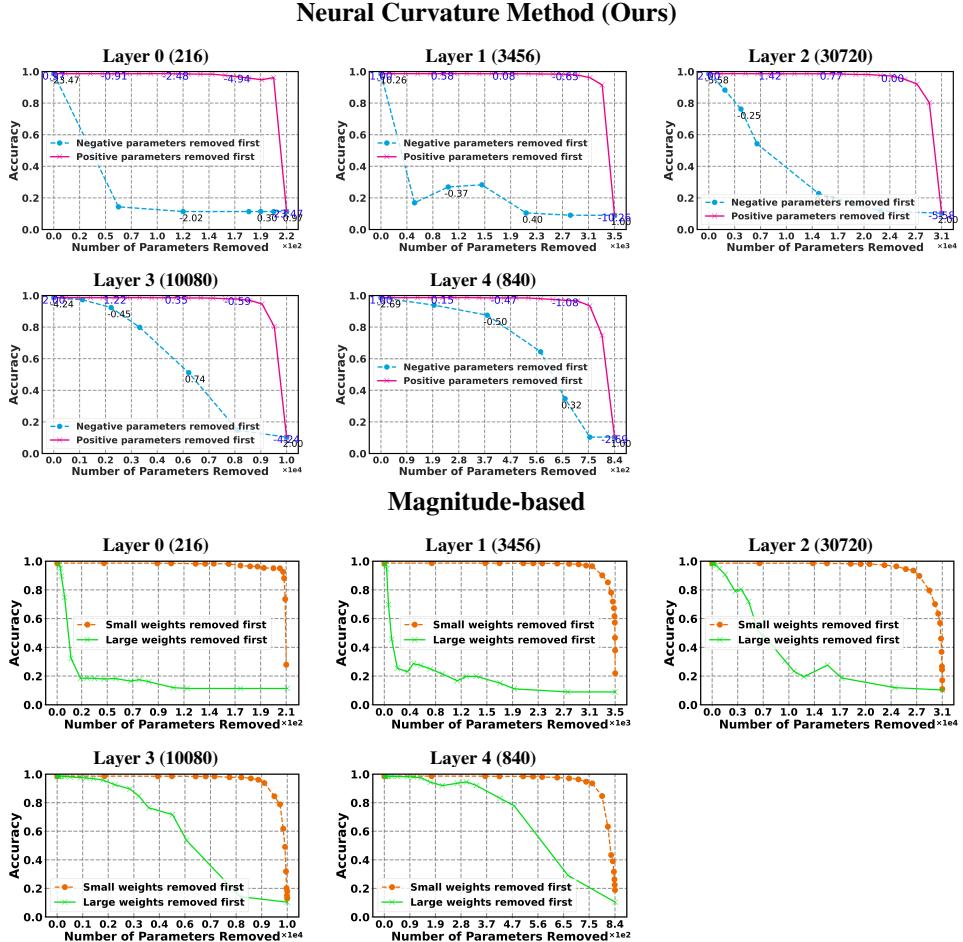
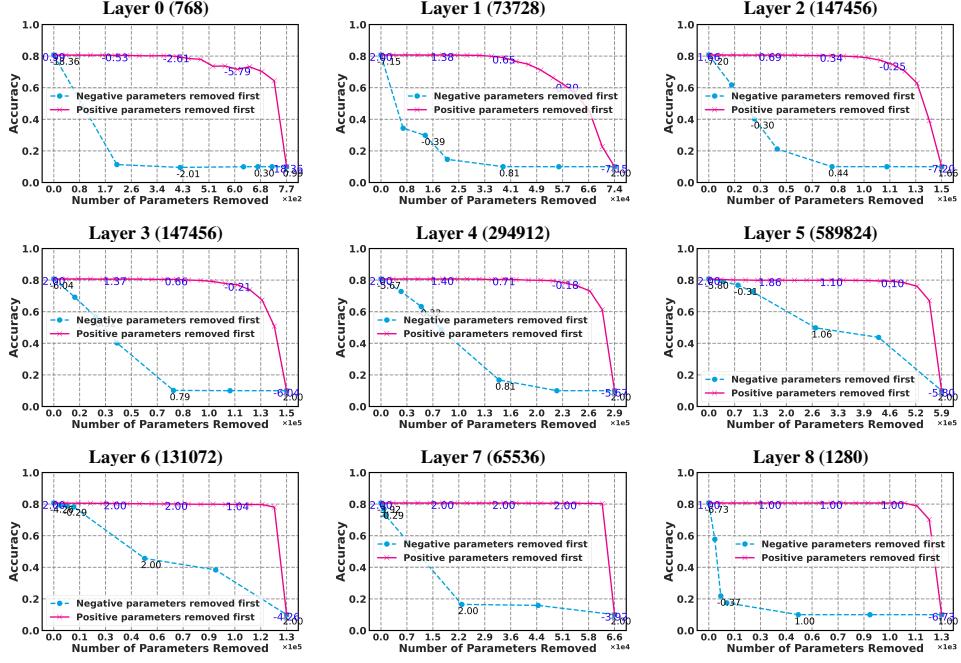


Figure 12: Ablation experiment on per-layer edge removal. Each subfigure shows the edge removal analysis for each layer based on neural curvature value (top 5 figures) and magnitude-based pruning (bottom 5 figures) for the CNN, AT Tanh configuration. The number above the figure is the total number of parameters in that layer.

### Neural Curvature Method (Ours)



### Magnitude-based

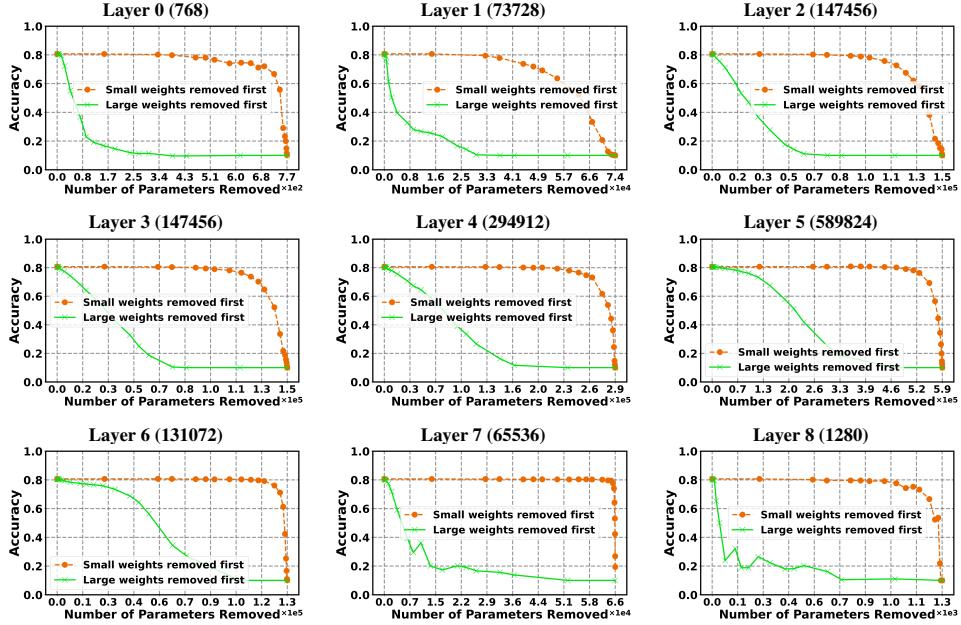
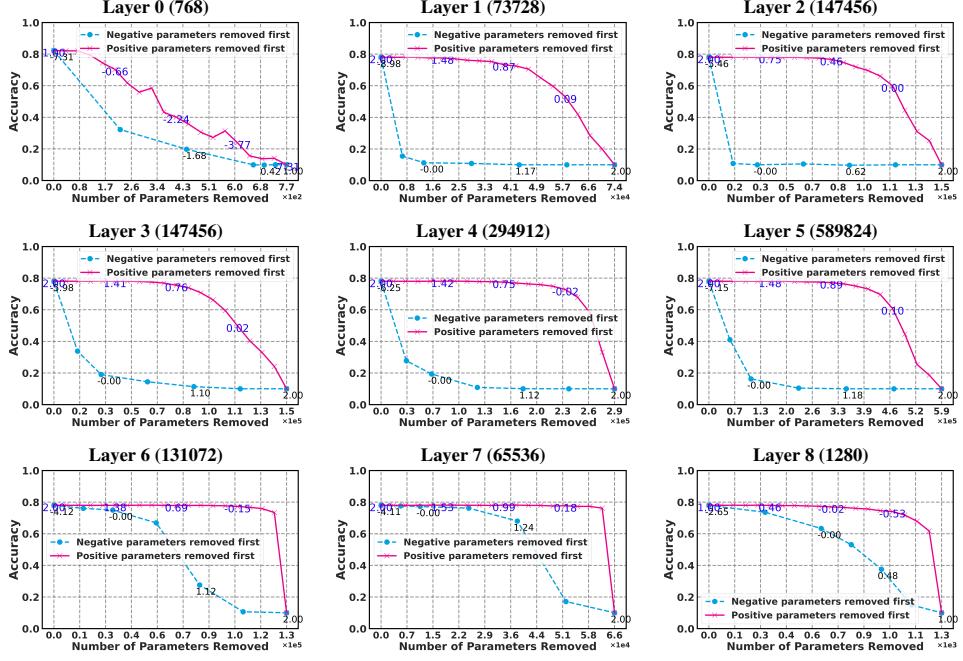


Figure 13: Ablation experiment on per-layer edge removal. Each subfigure shows the edge removal analysis for each layer based on neural curvature value (top 9 figures) and magnitude-based pruning (bottom 9 figures) for the VGG9-lite, AT ReLU configuration. The number above the figure is the total number of parameters in that layer.

## Neural Curvature Method (Ours)



## Magnitude-based

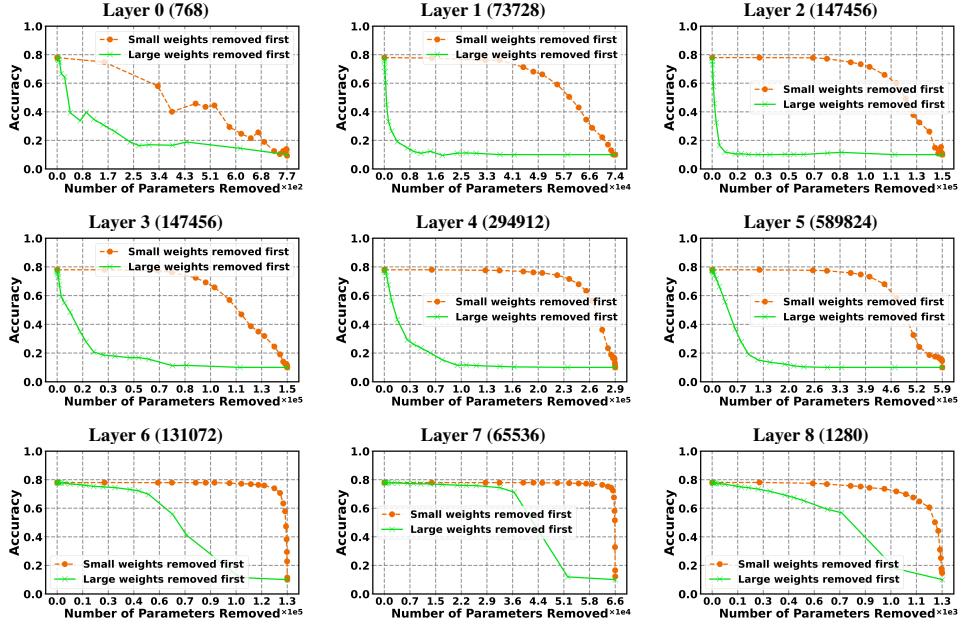


Figure 14: Ablation experiment on per-layer edge removal. Each subfigure shows the edge removal analysis for each layer based on neural curvature value (top 9 figures) and magnitude-based pruning (bottom 9 figures) for the VGG9-lite, CE Tanh configuration. The number above the figure is the total number of parameters in that layer.