SparseVLR: A Framework for Verified Locally Robust Sparse Neural Networks Search

Sawinder Kaur [1]   Asif Salekin [1]

Abstract

Developing robust sparse models fit for safety-critical and resource-constrained systems such as drones, autonomous robots, etc., has been an issue of longstanding interest. The inability of adversarial training mechanisms to provide a formal robustness guarantee kindles the requirement for verified local robustness mechanisms. This work aims to compute sparse verified locally robust networks which exhibit (benign) accuracy and verified local robustness comparable to their dense counterparts. Towards this objective, we examine several model sparsification approaches and present 'SparseVLR'– a framework to search verified locally robust sparse networks. We empirically investigated SparseVLR's efficacy and generalizability by evaluating various benchmark and application-specific datasets across several models. Above all, we provide an in-depth study and reasoning to unveil the causes for the ascendancy of SparseVLR.

1. Introduction

The application of neural networks (NNs) in safety-critical systems requires a framework for formal guarantees regarding their functioning (Katz et al., 2017). The brittle nature of NNs renders them vulnerable to imperceptible adversarial perturbation ε, which can lead to inaccurate model prediction (Goodfellow et al., 2015), and the actions triggered based on these incorrect predictions may have catastrophic effects (Pereira & Thomas, 2020).

The empirical adversarial training for NN models rely on efficiently finding adversarial perturbations around individual samples leveraging existing adversarial attacks. However, such training cannot guarantee that no other (yet unknown) attacks would find adversarial samples in the ℓ_p-ball of radius ε around the respective samples (Equation (1)) (Zhang et al., 2018a). For instance, recent approaches (Madry

---

[1]Syracuse University, USA.

et al., 2018; Wang et al., 2020; Zhang et al., 2020b; Qin et al., 2019) have proposed adversarial training mechanism which leverage PGD attack (Kurakin et al., 2016) of varying strengths. However, Tjeng et al. (2019) showed that PGD does not necessarily generate the perturbations with maximum loss; thus, the models trained to minimize PGD adversarial loss are still vulnerable to other stronger attacks. On the contrary, verified local robustness training (Katz et al., 2017; Gowal et al., 2018; Zhang et al., 2020a; Xu et al., 2020) guarantees the non-existence of adversarial samples in the vicinity of the benign samples, making it essential for safety-critical systems.

Model Compression - A Necessity. The enormous number of parameters present in dense NNs challenge their deployment in resource-constrained environments such as self-navigation, hazard detection, etc., making it essential to generate sparser NN models. The majority of the existing works that focus on obtaining sparse NN models aim to maintain the prediction (benign) accuracy of the sparse model comparable to the dense models trained with the same objective (Banner et al., 2018; Han et al., 2015b; Zhang et al., 2018b; Han et al., 2015a; Li et al., 2017; Sanh et al., 2020a). Recently, works are aiming at the additional target of adversarial robustness for a certain attack alongside benign accuracy (Sehwag et al., 2019; 2020; Lee et al., 2022). *This first-of-its-kind work extends the scope to obtaining sparse NN which exhibit high accuracy and verified local robustness (definition is provided in Section 2.2) while using significantly fewer parameters (1% of the comparable verified locally robust dense networks evaluated in literature).*

The paper's tasks and contributions are summarized below:

1. This is the first paper demonstrating the existence of *verified locally robust sparse neural networks* and illustrating in-depth investigation, comparisons, and reasoning on sparsification mechanisms to obtain them.

2. *Magnitude-based Pruning* is the most common form of sparsification mechanism (Sehwag et al., 2019; Zhang et al., 2018b; Li et al., 2017; Sanh et al., 2020a; Sehwag et al., 2020), which fails at obtaining sparse networks comparable in generalizability to their denser counterparts at high sparsification. However, pruning followed by 'grow-and-prune' (Dai et al., 2019) based

re-training achieves such sparse networks (Section 3). In this work, we refer to *model generalizability* as the ability of the model to project high (benign) accuracy and verified local robustness simultaneously.

3. We further investigated the feasibility of employing the 'grow-and-prune' training starting from 'randomly initialized sparse networks' (i.e., from scratch, avoiding the training of dense converged models first) to obtain the sparse networks with high generalizability. This adaptation is referred to as SparseVLR—which allows the computation of highly sparse networks (using 1% of original parameters) from scratch having generalizability comparable to their dense counterparts.

4. The paper performs in-depth empirical analysis to uncover the performance of SparseVLR and *pruning followed by 'grow-and-prune'* (Section 5.1 & Appendix E). The 'grow-and-prune' training facilitates higher gradient flows, allowing the models to learn new concepts by exploring new network connections, which leads to lower stabilized loss during training, resulting in higher generalizability in the obtained sparse models (Section 5.1(i)-(iii) & Appendix E).

   However, in the *pruning followed by 'grow-and-prune'* approach, the parameter weights inherited from the fully converged dense network after the pruning step prohibits exploration of new connections through 'grow-and-prune' training, resulting in a relatively inferior performance at high sparsity (Section 5.1(iv)).

5. Our empirical study (Section 5) demonstrates that SparseVLR is effective on benchmark datasets (CIFAR-10, MNIST, & SVHN), application-relevant *Pedestrian Detection* (N J Karthika, 2020) dataset, sentiment analysis NLP dataset SST-2 (Socher et al., 2013), and model architecture combinations used in the literature.

Additionally, Section 2 provides relevant backgrounds and related work discussion, Section 6 presents evaluations on some design choices of SparseVLR, and Section 7 discusses the impact, limitations and future-work scope of this study. Notably, our comprehensive Appendix further discusses the datasets, hyper-parameters, seeds, models, backgrounds, and detailed analysis of different aspects of SparseVLR.

2. Background and Related-Works

2.1. Formal robustness verification

Formal robustness verification of non-linear NNs is an NP-complete problem (discussed in Appendix B). Recent Linear Relaxation based Perturbation Analysis (LiRPA) methods can verify non-linear NNs in polynomial time. These methods use Forward (IBP (Gowal et al., 2018)), Backward (CROWN (Zhang et al., 2018a), Fastened-CROWN (Lyu et al., 2020), α-CROWN (Xu et al., 2021a), β-CROWN (Wang et al., 2021), DeepPoly (Singh et al.,

2019), BaB Attack (Xu et al., 2021b)) or Hybrid (CROWN-IBP (Zhang et al., 2020a)) bounding mechanisms to compute linear relaxation of a model.

Specifically, for each sample x_0, the perturbation neighborhood is defined as ℓ_p-ball of radius ε as

$$\mathbb{B}_p(x_0, \varepsilon) := \{x | \|x - x_0\|_p \le \varepsilon\}. \tag{1}$$

LiRPA aims to compute linear approximations f^ε of the model $\mathcal{M}_\theta$ (θ represents model parameters), and provide lower f_L^ε and upper f_U^ε bounds at the output layer, such that for any sample $x \in \mathbb{B}_p(x_0, \varepsilon)$ and each class j, the model output is guaranteed to follow:

$$f_L^\varepsilon(x)^j \le \mathcal{M}_\theta(x)^j \le f_U^\varepsilon(x)^j. \tag{2}$$

2.2. Verified local robustness.

A NN model, $\mathcal{M}_\theta$ is said to be verified locally robust (Fromherz et al., 2021) for a sample (x_0, y) if: $\forall x \in \mathbb{B}_p(x_0, \varepsilon) \Rightarrow \mathcal{M}_\theta(x) = \mathcal{M}_\theta(x_0)$, scilicet for all the samples $x \in \mathbb{B}_p(x_0, \varepsilon)$, the model is guaranteed to assign the same class as it assigns to x_0. LiRPA approaches measure the local robustness in terms of a margin vector $m(x_0, \varepsilon) = Cf^\varepsilon(x_0)$, where C is a *specification matrix* of size $n \times n$ and n is the number of all possible classes (Gowal et al., 2018; Zhang et al., 2018a; 2020a; Xu et al., 2020). Zhang et al. (2020a) defines C for each sample (x_0, y) as:

$$C_{i,j} = \begin{cases} 1, & j = y \text{ and } i \ne y \text{ (truth class)} \\ -1, & i = j \text{ and } i \ne y \text{ (other classes)} \\ 0, & \text{otherwise.} \end{cases} \tag{3}$$

The entries in matrix C depend on the true class y. In matrix C, the row corresponding to the true class contains 0 at all the indices. All the other rows, contain 1 at index corresponding to the true class ($j = y$) and -1 at index corresponding to current class ($j = i$) and 0 at all the other indices. Thus, the i^{th} value of $m(x_0, \varepsilon)$ is given by $m^i(x_0, \varepsilon) = f^\varepsilon(x_0)^y - f^\varepsilon(x_0)^i$, which is the difference of output values of the true class y with all the other classes.

$\underline{m}(x_0, \varepsilon)$ represents the lower bound of the margin vector. If all the values of the $\underline{m}(x_0, \varepsilon)$ are positive, $\forall_{i \ne y} \underline{m}^i(x_0, \varepsilon) > 0$, the model $\mathcal{M}_\theta$ is said to be locally robust for sample x_0. That implies, the model will always assign the highest output value to the true class label y if a perturbation less than or equal to ε is induced to the sample x_0. Thus, $\forall_j \underline{m}(x_0, \varepsilon)^j > 0$ implies the guaranteed absence of any adversarial sample in the region of interest. Furthermore, to specify the region bounded by $\mathbb{B}_p(x_0, \varepsilon)$, we use $p = \infty$ because ℓ_∞-norm covers the largest region (Wild, 2018).

2.3. Model training to maximize local robustness

Since, a model is considered verified locally robust for a sample (x_0, y) if all the values of $\underline{m}(x_0, \varepsilon) > 0$, the train-

ing schemes proposed by previous approaches (Gowal et al., 2018; Zhang et al., 2020a) aim to maximize robustness of a model by maximizing the lower bound $\underline{m}(x_0, \varepsilon)$. Xu et al. (2020) defines the training objective as minimizing the maximum cross-entropy loss between the model output $\mathcal{M}_\theta(x)$ and the true label y among all samples $x \in \mathbb{B}_p(x_0, \varepsilon)$ (eq. 6 in (Xu et al., 2020)). Thus,

$$\mathcal{L}_{\text{train}}(\mathcal{M}_\theta, \mathcal{D}, \varepsilon) = \sum_{(x_0, y) \in \mathcal{D}} \max_{x \in \mathbb{B}(x_0, \varepsilon)} \mathcal{L}(\mathcal{M}_\theta(x), y), \quad (4)$$

where $\mathcal{D}$ is the dataset and $(x_0, y) \in \mathcal{D}$. Intuitively, $\mathcal{L}_{\text{train}}$ is the sum of maximum cross-entry loss in the neighborhood of each sample in $\mathcal{D}$ for perturbation amount ε. Wong & Kolter (2018) showed that the problem of minimizing $\mathcal{L}_{\text{train}}$ (as defined in Equation (4)) and the problem of maximizing $\underline{m}(x_0, \varepsilon)$ are dual of each other. Thus, *a solution that minimizes $\mathcal{L}_{train}$, maximizes the values of $\underline{m}(x_0, \varepsilon)$, hence maximizes the local robustness for the sample (x_0, y).* Since, $x_0 \in \mathbb{B}_p(x_0, \varepsilon)$, *minimizing $\mathcal{L}_{train}$ also maximizes benign accuracy, ergo maximizes generalizability altogether.*

Xu et al. (2020) computes $\max_{x \in \mathbb{B}(x_0, \varepsilon)} \mathcal{L}(\mathcal{M}_\theta(x), y)$ which requires computing bounding planes f_U^ε and f_L^ε using one of the aforementioned LiRPA bounding techniques. Following Xu et al. (2020), SparseVLR employs CROWN-IBP hybrid approach as the bounding mechanism to optimize $\mathcal{L}_{\text{train}}(\mathcal{M}_\theta, \mathcal{D}, \varepsilon)$ in Equation (4). The perturbation amount ε is initially set to zero and gradually increases to $\varepsilon_{\max}$ according to the perturbation scheduler ε-scheduler$(\varepsilon_{\max}, t, s, l)$ discussed in Appendix B.1.

Metrics for evaluations: Table 1 defines the metrics used to evaluate models which have been used in the previous works for formal robustness verification (Gowal et al., 2018; Zhang et al., 2018a; 2020a; Xu et al., 2020).

Table 1. Evaluation metrics. Here, % implies percentage.

Error	Definition
Standard	% of benign samples classified incorrectly.
Verified	% of benign samples for which at least one value in the $\underline{m}(x_0, \varepsilon)$ is negative.
PGD	% of perturbed samples, generated using 200-step PGD attack, classified incorrectly.

2.4. Model sparsification

Model Sparsification is the most common form of model compression, which constrains the model to use only a subset of model parameters for inference (Sehwag et al., 2019; Zhang et al., 2018b; Han et al., 2015a; Li et al., 2017; Sanh et al., 2020a; Sehwag et al., 2020). Hoefler et al. (2021) categorizes the model sparsification based on the stage at which sparisfication is performed: • *train-and-sparsify*: pruning a fully trained model often followed by re-training (finetuning) the retained parameters (Han et al., 2015a; Zhang et al., 2018b; Sehwag et al., 2019; 2020), • *sparsifying-during-*

training: mutiple rounds of *train-and-sparsify* while increasing sparsity in every round a.k.a. iterative pruning (Tan & Motani, 2020; You et al., 2019), and • *sparse-training*: training a sparse network using either a static (Frankle et al., 2020; Frankle & Carbin, 2019; Frankle et al., 2019; Savarese et al., 2020; Malach et al., 2020) or dynamic (Dai et al., 2019; Liu et al., 2021) mask.

Recently developed *dynamic-mask-based sparse-training* known as Dynamic Sparse Training (DST) (Dai et al., 2019; Liu et al., 2021) approaches iteratively follow *grow-and-prune* paradigm to achieve an optimal sparse network. Starting from a random sparse *seed network*, new connections (i.e., parameters) which minimize the natural loss are added (*grow step*), and the least important connections are removed (*prune step*). DST's higher efficacy than the *static-mask-based sparse-training* is attributed to high gradient flow allowing the model to learn an optimal sparse network for effective inference generation (Evci et al., 2022). *SparseVLR adapts the DST approach to identify and train verified locally robust sparse networks from scratch.*

Parameter removal or pruning: SparseVLR uses global magnitude-based pruning, an unstructured pruning mechanism (Paganini, 2022), as multiple works (See et al., 2016; Ghosh et al., 2022; Liu et al., 2017; 2019; Singh & Alistarh, 2020) suggest its higher efficacy in achieving more generalizable sparse networks. Magnitude-based pruning approaches (Han et al., 2015b; Sehwag et al., 2019; Zhang et al., 2018b; Li et al., 2017; Sanh et al., 2020a; Sehwag et al., 2020; Han et al., 2016) remove the least magnitude parameters suggesting that the importance of a parameter is directly proportional to the magnitude of its weights.

There is rich literature on sparsification mechanisms and parameter removal. A detailed summary is in Appendix C.

3. Motivation: Existence of Verified Locally Robust Sparse Network

Zhang et al. (2020a) noted that the verified local robustness training mechanism proposed by Wong & Kolter (2018) and Wong et al. (2018) induce implicit regularization, thus, penalizing the magnitude of model parameters during training. CROWN-IBP (Zhang et al., 2020a) incurs less regularization and shows an increasing trend in the magnitude of model parameters while training. According to our preliminary analysis (Appendix D), the implicit regularization caused by CROWN-IBP does penalize model's parameters, making them smaller as compared to naturally trained networks, causing a high fraction of the parameters to be close to zero. Removal of such less significant parameters has minimal impact on model generalizability, which indicates *the existence of a sparse sub-network that exhibits verified local robustness comparable to its dense counterpart.*
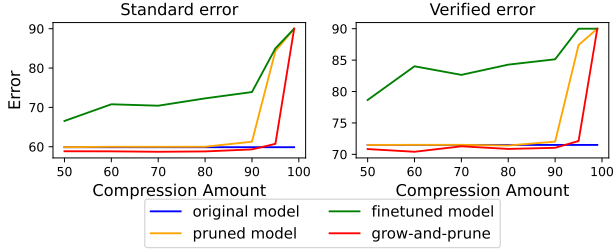
Figure 1. Standard and verified error for dense model (original model) vs. pruned model obtained using (i) global pruning (pruned model), (ii) global pruning followed by fine-tuning (finetuned model), and (iii) re-training using grow-and-prune (grow-and-prune) for 4-layer CNN for CIFAR-10 ($\varepsilon = 8/255$)

Figure 1 demonstrates the *standard* and *verified* errors for 4-layer CNN trained to minimize L_{train} defined in Equation (4) for benchmark dataset CIFAR-10 at various compression amounts. The evaluations are shown for the (i) Global Pruning without any finetuning or re-training, (ii) Conventionally finetuning a globally pruned network using a static mask, and (iii) Re-training the pruned network using DST-based *grow-and-prune* technique (Dai et al., 2019).

Figure 1 shows that as the compression ratio increases ($> 90\%$), weights having higher significance on model inference become the pruning target. Thus, the model generalizability, in terms of natural accuracy and verified local robustness, drops significantly, resulting in a *trivial classifier*[1] in extreme cases. Conventional finetuning ($\#ii$) increases the *standard* and *verified* error even at small compression amounts. However, re-training using *grow-and-prune* ($\#iii$) maintains the generalizability of the sparse models at relatively high compression amounts. The effectiveness of *grow-and-prune* ($\#iii$) in re-training a pruned model can be attributed to the better gradient flow encountered during training as compared to conventional static-mask-based finetuning ($\#ii$), which is in agreement with (Evci et al., 2022). A comparison of ($\#i$) & ($\#iii$) and their gradient flow and loss evolution during training is provided in Appendix E.1.

Moreover, according to our preliminary analysis, locally robust sparse training leveraging static-mask-based sparsifying-during-training approaches such as *iterative pruning* (Tan & Motani, 2020) results in inferior performance. ***In summary,*** this section establishes **(1)** the existence of verified locally robust sparse sub-network in dense verified robustly trained models; and **(2)** *grow-and-prune* training paradigm can help identify such sparse networks.

---

[1] A classifier that always predicts the same class irrespective of the input

4. Verified Locally Robust Sparse Model via Dynamic-Mask-based Sparse-training

Motivated by the above-discussed findings, this paper presents a dynamic-mask-based sparse-training mechanism to achieve a verified locally robust sparse network without training a dense over-parameterized model. SparseVLR is based on *grow-and-prune* paradigm (Dai et al., 2019), which allows network connections to expand under the constraints of an underlying backbone architecture. Each parameter of the backbone model can either be in *active* or *dormant* state based on whether or not it belongs to the sparse network at a particular instant. The weight of a *dormant* parameter is set to zero. The percentage of parameters that are *dormant* is the network's current *sparsity*.

4.1. Problem definition:

For a given randomly initialized backbone architecture $\mathcal{M}_\theta$ with k parameters, the objective is to train a verified locally robust sparse network $\mathcal{M}_{\theta\downarrow}$ which uses only a subset of available parameters (of the backbone), that is, $|\theta^\downarrow|_0 = k'$ and k' is an order of magnitude times smaller than k, $k' << k$. The sparse model thus achieved should exhibit high accuracy and verified local robustness, comparable to a verified locally robust dense network having the same backbone architecture.

Algorithm 1 SparseVLR

Require: $\mathcal{M}_\theta$: Backbone architecture ⋄ p: Target sparsity $\in (0, 100)$ ⋄ $\mathcal{D} = (X_i, y_i)_{i=1}^m$: Dataset ⋄ $\varepsilon_{\max}$: Maximum Perturbation ⋄ T: Number of training epochs ⋄ (s, l): Start and length of perturbation scheduler ⋄ *seed*: A seed value
1: Randomly initialize the backbone architecture using *seed*.
2: $\theta^\downarrow = [m|m \in p^{th}$ percentile of $|\theta|]$ { Select the highest (100-p)% parameters of $\mathcal{M}_\theta$ to form a sparse seed network $\mathcal{M}_{\theta\downarrow}$ while making remaining parameters dormant.}
3: **for** epoch $t = [0, 1, \ldots, T]$ **do**
4:     $\varepsilon_t = \varepsilon$-Scheduler($\varepsilon_{\max}, t, s, l$)
5:     **for** minibatches $d \in \mathcal{D}$ **do**
6:         $\mathcal{M}_{\theta'} = \underset{\theta}{argmin} \, \mathcal{L}_{\text{train}}(\mathcal{M}_{\theta\downarrow}, d, \varepsilon_t)$ {Auxiliary Model $\mathcal{M}_{\theta'}$ is denser than target sparsity, i.e.$|\theta'| > |\theta^\downarrow|$}
7:     **end for**
8:     $\theta^\downarrow = [m|m \in p^{th}$ percentile of $|\theta|]$ { Select the highest (100-p)% parameters of $\mathcal{M}_\theta$, ergo the highest k' parameters of $\mathcal{M}_{\theta'}$ to obtain target sparsity}
9: **end for**

4.2. SparseVLR: The approach

Algorithm 1 describes the SparseVLR procedure. It requires a backbone architecture M_θ; a target sparsity p, the dataset $\mathcal{D}$; and the maximum perturbation $\varepsilon_{\max}$. It also requires hyperparameters that vary for different datasets and includes the number of training epochs T and the inputs for the ε-scheduler: s and l (discussed in Appendix A).

SparseVLR starts with an untrained random sparse (i.e., seed) network while keeping $p\%$ parameters of the backbone as dormant ($p = (1 - \frac{k'}{k}) \times 100$). It is an iterative procedure, and each iteration comprises two phases: ● *Thickening*: aims to explore new parameters in the backbone to maximize verified local robustness resulting in lesser sparsity, and ● *Pruning*: removes parameters that hold lesser importance for model inference to reach the target sparsity at the end of each iteration.

Seed network (lines 1-2 in Algorithm 1) Given a backbone architecture, a sparse seed network is computed by randomly initializing all the backbone parameters, followed by retaining the parameters having only the top $(100 - p)\%$ weight magnitudes. All the other parameters of the backbone are set to zero (made dormant).

Thickening (lines 4-7 in Algorithm 1) This phase aims to densify the sparse network by inspecting all the parameters (within the backbone). The parameter weights, including the *dormant* ones, are updated according to their corresponding gradients. This is also referred to as *gradient based growth phase* in the literature (Dai et al., 2019; Evci et al., 2022; Liu et al., 2021). However, in literature, gradients are computed for the objective of maximizing benign accuracy, that is, minimizing the *standard* error; in contrast, SparseVLR aims to maximize verified local robustness of the model by minimizing the training loss $\mathcal{L}_{\text{train}}$ as defined in Equation (4) for each mini-batch $d \in \mathcal{D}$. The perturbation ε used for computing $\mathcal{L}_{\text{train}}$ is computed using the perturbation scheduler $\varepsilon-\text{scheduler}(\varepsilon_{\max}, t, s, l)$ as discussed in Appendix B.1.

The updated network $\mathcal{M}_{\theta'}$ is called *auxiliary network* . Since there is no restriction on the capacity of the *auxiliary network*, it can have more than k' active parameters.

Pruning (line 8 in Algorithm 1) The auxiliary network obtained after the *Thickening* phase needs to be compressed to the target size k'. This phase aims to reduce the model size in terms of the number of *active* parameters while having minimum impact on the generalizability. As discussed in Section 2, to be consistent with the state-of-the-art, we employ ℓ_1-unstructured global pruning as the model sparsification mechanism (Paszke et al., 2017), to attain the sparse network $\mathcal{M}_{\theta\downarrow}$ with target size k'. Our sparsification uses ℓ_1-norm of individual model parameters as the metric to decide its importance towards model inference. Notably, ℓ_1-norm corresponds to the magnitude of the weight of the individual parameters.

However, Liu et al. (2017) also showed that global pruning could result in a disconnected network. This can be problematic for the approaches which use single-shot pruning, that is, pruning a fully trained network followed by finetuning only the retained parameters. SparseVLR dynamically explores and selects the connections which minimize the training loss, thus, eliminating the chances of receiving a disconnected network.

5. Experiments

Section 3 establishes the superior performance of *pruning followed by retraining using 'grow-and-prune'*, among the pruning based sparsification mechanisms. However, this approach fails at high sparsity, such as 99% (see Table 2). This section's empirical analysis establishes that 'grow-and-prune' training starting from a sparse network that is randomly initialized (i.e., the seed network used in SparseVLR) instead of inheriting parameter weights from a pruned fully converged network, resolves the issue. To compare, this section evaluates *thickening*-and-*pruning*-based sparse model training (lines 3-9 of Algorithm 1) using two different starting/seed sparse network variations:

- Random Sparse Model (RSM): This is the seed network used in SparseVLR in Section 4.2 and is obtained by applying global magnitude-based pruning to the randomly initialized backbone architecture.
- Pruned Model (PM): PM is the same as ($\#i$) model in Section 3, obtained by applying global magnitude-based pruning to a fully converged dense model.

The final models obtained after applying Algorithm 1 (lines 3-9) to PM and RSM are denoted by PM^T and RSM^T. PM^T is similar to the ($\#iii$) in Section 3 and RSM^T denotes the final model obtained using SparseVLR.

Following section's evaluations demonstrate: **1)** The effectiveness of SparseVLR in achieving verified locally robust sparse models (Section 5.1). **(2)** While examining the reasons for SparseVLR's efficacy, we observe that the dynamic masking used by SparseVLR results in high gradient flow (Section 5.1(i)-(ii)), which allows the model to learn new concepts (Section 5.1(iii)), which leads to lower stabilized loss during training. **(3)** During *pruning followed by 'grow-and-prune'* (i.e., $\#iii$ in Section 3), the seed network's weight inheritance from a fully converged dense model prohibits the exploration of new connections (i.e., learning new concepts) of the sparse network during training, resulting in inferior performance at high sparsity (Section 5.1(iv)). Moreover, **(4)** the generalizability of SparseVLR is demonstrated by evaluating several model dataset combinations as discussed below.

Datasets and models: To establish the empirical effectiveness of SparseVLR, we use benchmark dataset CIFAR-10, used in the previous works (Gowal et al., 2018; Zhang et al., 2018a; 2020a; Wong et al., 2018; Xu et al., 2020). To demonstrate the effectiveness and generalizability of SparseVLR, evaluations are shown for four more datasets: MNIST, SVHN, Pedestrian Detection (N J Karthika, 2020), and Sentiment Analysis SST-2 (Socher et al., 2013). To demonstrate
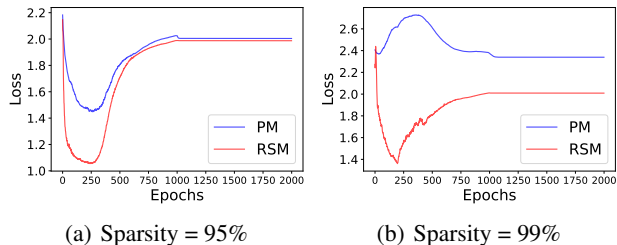
(a) Sparsity = 95%   (b) Sparsity = 99%

Figure 2. Variation of training loss for CIFAR-10 dataset on 4-layer CNN with $\varepsilon_{\text{test}} = 8/255$ at (a) 95% & (b) 99% sparsity amounts

Table 2. Error Metrics for 4-layer CNN, 7-layer CNN, and Resnet trained for CIFAR ($\varepsilon_{\text{train}} = 8.8/255$ and $\varepsilon_{\text{test}} = 8/255$) at different sparsity amounts. Here, PM and PM^T are same as (#i) and (#iii) models in Section 3, and RSM^T denotes the final model obtained using SparseVLR

Model		Sparsity ⇒	0%	80%	90%	95%	99%
		$\mathcal{M}_{\theta\downarrow}$	$\mathcal{M}_\theta$				
	Type	Error					
4-layer CNN	PM	Standard		60.0	61.28	84.33	90.0
		PGD		67.3	68.05	86.43	90.0
		Verified		71.41	72.02	87.4	90.0
	PM^T	Standard	59.88	58.81	59.32	60.73	90.0
		PGD	67.25	68.46	66.01	66.74	90.90
		Verified	71.5	70.86	71.04	72.12	90.0
	RSM^T	Standard		60.83	60.78	61.26	64.14
		PGD		68.45	67.08	67.13	69.09
		Verified		72.15	71.90	72.45	73.43
7-layer CNN	PM	Standard		56.69	70.78	84.05	90.0
		PGD		68.75	82.63	84.88	90.0
		Verified		68.83	79.54	86.19	90.0
	PM^T	Standard	56.08	54.59	55.33	55.24	90.0
		PGD	68.65	68.46	67.59	67.21	90.0
		Verified	68.66	68.5	68.49	69.6	90.0
	RSM^T	Standard		55.11	55.41	56.96	59.14
		PGD		67.17	68.26	68.30	67.04
		Verified		68.60	68.84	69.12	71.21
Resnet	PM	Standard		55.68	88.06	90.0	90.0
		PGD		65.24	89.41	90.0	90.0
		Verified		70.41	89.44	90.0	90.0
	PM^T	Standard	53.6	53.38	55.17	90.0	90.0
		PGD	63.83	64.37	66.25	90.0	90.0
		Verified	68.58	68.57	69.28	90.0	90.0
	RSM^T	Standard		55.36	56.30	57.99	60.31
		PGD		66.55	67.44	67.04	68.88
		Verified		68.95	69.33	70.14	71.38

the versatility of the approach across different architectures, we evaluate (a) Two CNNs with varying capacity: 4-layer CNN and 7-layer CNN (Gowal et al., 2018; Zhang et al., 2020a); (b) Network with skip connections: A 13-layer ResNet (Wong et al., 2018); (c) Two DNNs: DenseNet and ResNext (Xu et al., 2020); (d) Two sequential networks (Xu et al., 2020). Further details about perturbations and hyper-parameters are discussed in Appendix A.

Formal verification mechanisms to compute $\mathcal{L}_{\text{train}}$ and verified error: While training a sparse network, Sparse-VLR employs CROWN-IBP (Zhang et al., 2020a). In Section 6, we show empirically that the sparse models trained using CROWN-IBP exhibit lower errors than sparse models obtained using other bounding mechanisms such as IBP (Gowal et al., 2018), CROWN (Zhang et al., 2018a), and Fastened-CROWN (Lyu et al., 2020). For evaluations, to be consistent with the state-of-the-art (Gowal et al., 2018; Zhang et al., 2020a; Xu et al., 2020), the presented evaluations use IBP to compute *verified error*.

5.1. Establishing the effectiveness of SparseVLR:

Table 2 shows the results for empirical analysis of Sparse-VLR for the benchmark dataset CIFAR-10 for three models: 4-layer CNN, 7-layer CNN, and Resnet at sparsity amounts 80%, 90%, 95%, & 99%. The models are evaluated through *Standard*, *PGD*, and *Verified* errors as described in Table 1. The column corresponding to 0% sparsity represents the error metrics computed for dense models (having the same architecture as the backbone in Algorithm 1) trained using Xu et al. (2020), which aligns well with the state-of-the-art (Zhang et al., 2020a; Wong et al., 2018). The presented results are averaged over five random initializations (seeds).

According to the Table 2, RSM^T achieves similar evaluation metrics (within 3% verified error) as the dense (0% sparsity) models across all architectures and sparsity amounts. *These results establish that (i) SparseVLR can train a verified locally robust sparse model from scratch, i.e., without needing a dense over-parameterized network trained till convergence, and (ii) The sparse model thus obtained is comparable to the verified locally robust dense network*

with the same architecture as the backbone.

Also, RSM^T outperforms PM^T (the best approach from Section 3). Table 2 shows that the error metrics for RSM^T are comparable to PM^T at smaller sparsity and at higher sparsity (e.g., 99%) RSM^T exhibits much less error than PM^T, establishing SparseVLR's (Algorithm 1) higher efficacy compared to the conventional pruning mechanisms. Comparisons for PM, PM^T, and RSM^T are provided in Appendix E.

To investigate the reason for the better performance of SparseVLR, a detailed analysis for 4-layer CNN trained for CIFAR-10 is provided; however, similar patterns were observed for other models as well.

(i) *Why the RSM^T vs. PM^T disparity at higher sparsity?*

Figure 2(a) depicts the loss evolution for PM→PM^T and RSM→RSM^T training at 95% sparsity, and reveals that the losses eventually stabilize at the same level, which in turn is governed by the maximum perturbation $\varepsilon_{\max}$. SparseVLR uses ε-scheduling, which provides 0 perturbation for initial s epochs and then gradually increases the perturbation value. This results in an initial drop in losses, as seen in figures Figure 2(a). As the perturbation amount increases, the loss increases as well.
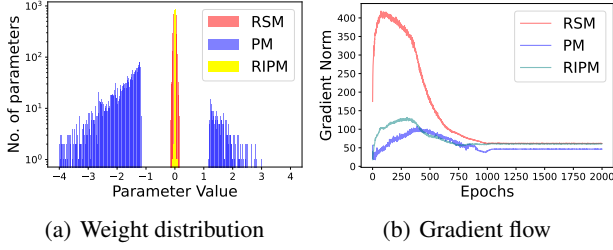
(a) Weight distribution   (b) Gradient flow

Figure 3. (a)Distribution of weights and (b) Gradient flow encounters while training RSM, PM and RIPM for a 4-layer CNN for CIFAR-10 dataset at 99% sparsity
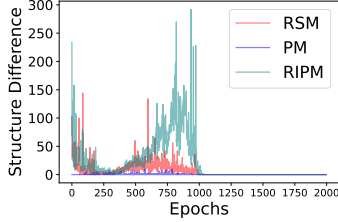

Figure 4. Changes in sparse model structure while training a 4-layer CNN for CIFAR-10 dataset $\varepsilon_{train} = 8.8/255$ at 99% sparsity using RSM, PM and RIPM as seed networks

Notably, at 99% sparsity (Figure 2(b)), the stabilized loss for PM^T is significantly greater than RSM^T. Higher loss leads to higher error, hence, the observation. The evolution of training loss of PM^T is discussed further in Appendix E. The reason for the higher stabilized loss for PM^T at high sparsity is discussed below.

(ii) *Reason for disparity in the stabilized losses for RSM^T and PM^T: Gradient flow over $\mathcal{L}_{train}$*

Evci et al. (2022) claims that different initialization may lead to different gradient flow, which in turn leads to varied accuracy of the sparse model trained via DST. To explain the disparity in stabilized losses for PM^T and RSM^T at 99% sparsity (Figure 2b), we plot the observed gradient flows in Figure 3(b). It depicts that the ℓ_1-norm of gradients observed while training RSM→RSM^T is greater than the gradient norm incurred while training PM→PM^T, thus, allowing the model to learn new concepts (further discussed below), hence incurring a lower stabilized loss.

(iii) *How poor gradient flow hinders the ability of the sparse network to change connections (i.e., learning new concepts) during training?*

The effect of poor gradient flows is two-fold: (i) a lesser number of parameters explored (activated) during the thickening phase, and (ii) relatively small magnitude in the newly activated parameters of the auxiliary model.

The change in the structure of a sparse network obtained after a pruning step can be quantified in terms of the number of parameters of the backbone whose state changed

from active to dormant and vice-versa. Figure 4 depicts the changes in the sparse structure during training. Notably, PM→PM^T training incurs almost no change, whereas the sparse structure of RSM→RSM^T training is quite dynamic. This can be attributed to smaller magnitudes of newly activated parameters which become an easy target of removal during the pruning step, resulting in a similar sparse structure. It further hinders the sparse network's ability to learn new concepts, leading to a higher loss in PM^T than RSM^T.

(iv) *Cause for inefficient learning while PM→PM^T - Is it the structure or the weight initialization of the seed network?*

PM inherits parameter weights from a pre-trained dense model, and the PM structure is learned through global magnitude-based pruning. This section investigates whether the lower performance of PM^T at high sparsity is due to PM's retained weights or learned network structure.

For this investigation, we evaluate the Randomly Initialized Pruned Model (RIPM) as the seed network, which randomly initializes PM's parameters while keeping the learned structure intact. Table 3 compares $RIPM^T$, the sparse model obtained via training RIPM→$RIPM^T$ using Algorithm 1 (lines $3 - 9$), with PM, PM^T, and RSM^T. It can be observed that $RIPM^T$ generalizes comparable to RSM^T. Also, Figure 4 suggests that while training RIPM→$RIPM^T$, the model can learn new connections (high variation in the sparse network structure) in the first half of training.

However, Figure 3(b) shows that the gradient flow while training RIPM→$RIPM^T$ is only marginally better than the gradient flow observed while training PM→PM^T. Then, why $RIPM^T$ is better than PM^T?

At high sparsity (i.e., 99%), the comparison of the initial seed sparse networks' weight distribution shown in Figure 3(a) demonstrates that the magnitude of active weights in RIPM is lower than that of PM. As we know, the poor gradient flows result in low magnitudes among the newly activated parameters after the thickening phase, and the magnitude of parameter weights of RIPM (i.e., initially active) are also low; allowing newly active parameters to be selected (and replace previously active ones) during the pruning step, leading to changes in the structure of sparse network during training, allowing the model to learn new concepts, leading to higher performance in the obtained $RIPM^T$.

In contrast, the high magnitude of initially active parameter weights of PM prohibits changes in sparse network structure during training, leading to inferior performance. *Thus, the parameter weights inherited from a fully converged dense network are responsible for PM→PM^T's inefficient learning at high sparsity amounts.* This finding further motivates the effectiveness of training a verified locally robust sparse network from scratch rather than pruning an over-parameterized pre-trained dense model.

Table 3. Comparison for error metrics for sparse models obtained using different seed networks for a 4-layer CNN at 99% sparsity

Error	PM	PM^T	RSM^T	$RIPM^T$
Standard	90.0	90.0	65.01	65.81
PGD	90.0	90.90	70.43	69.85
Verified	90.0	90.0	73.38	74.97

Table 4. Error metrics for sparse models obtained using SparseVLR using 7 layer CNN as backbone architecture trained for (a) SVHN, (b) Pedestrian Detection and (c) MNIST

Dataset	SVHN			Pedestrian Detection			MNIST		
Sparsity ⇒ Error	0%	95%	99%	0%	95%	99%	0%	95%	99%
Standard	62.55	62.77	65.87	28.40	26.22	28.40	2.27	3.37	2.11
PGD	68.19	68.23	70.52	42.35	39.66	43.03	6.19	4.85	3.61
Verified	74.87	75.21	76.76	63.53	67.39	63.02	12.2	14.51	12.06

5.2. Establishing the generalizability of SparseVLR

• **Evaluation for complex models** To demonstrate the applicability of SparseVLR to more complex networks, we compute sparse models for DenseNet and ResNext (used by (Xu et al., 2020)) backbone architectures trained for MNIST and CIFAR. Appendix F demonstrates that the sparse models having only 1% (at 99% sparsity) of the original parameters generalize comparable to their dense counterparts.

• **Application to sequential model** Appendix G discusses the results for computing a sparse model for a small LSTM (Xu et al., 2020) trained for MNIST. We observe that for sparsity up to 95%, the sparse model generalizes comparable to its dense counterparts.

• **Evaluations for more datasets** Table 4 demonstrates the results for additional image datasets: (a) SVHN, (b) a pedestrian detection dataset that aims at differentiating between people and people-like objects (N J Karthika, 2020), and (c) MNIST. The backbone architecture used for these evaluations is a 7-layer CNN, adapted to the respective datasets. For all datasets, the error metrics of the sparse models are comparable to that of their dense counterparts. The perturbation amount used for SVHN, Pedestrian Detection and MNIST are 8/255, 2/255, and 0.4, respectively. Additionally, sparse models constructed for a sentiment analysis NLP dataset SST-2 (Socher et al., 2013) on an LSTM model (Jia et al., 2019; Xu et al., 2020) exhibit errors within 1% of the dense counterparts, showing SparseVLR's generalizability to non-image domains (details presented in Appendix H).

6. Design Choices for SparseVLR

• **Pruning after every thickening step:** Liu et al. (2021) suggests that allowing the model to *grow* for multiple epochs before *pruning* allows exploring all the parameters at least once. However, the implicit regularisation caused by the $\mathcal{L}_{train}$ (discussed in Section 3) does not allow such exploration. Our empirical evaluations show that *pruning* after every *thickening* step results in the optimal sparse models.

• **Using CROWN-IBP as bounding mechanism:** Appendix K compares the errors exhibited by 4-layer CNN at 99% sparsity when trained by employing different bounding mechanisms to compute $\mathcal{L}_{train}$: IBP (Gowal et al., 2018), CROWN (Zhang et al., 2018a), Fastened-CROWN (Lyu

et al., 2020), and CROWN-IBP (Zhang et al., 2020a). Notably, empirically CROWN-IBP results in the most optimal models for both MNIST and CIFAR. Hence, SparseVLR employs CROWN-IBP to compute $\mathcal{L}_{train}$.

7. Discussion on Impact and Limitations

Impact: SparseVLR will enable resource-constraint platforms to leverage verified robust models. Moreover, the inference time of the SparseVLR obtained sparse models is *five times* lesser, on average, than their dense counterparts on NVIDIA RTX A6000 (detailed results presented in Appendix J). Furthermore, the sparse training reduces the total training time to almost half that of conventional pruning.

Limitations & future-works: •SparseVLR uses CROWN-IBP; thus, its efficacy is restricted by the robustness achieved by CROWN-IBP. Hence, SparseVLR will need to be tailored for further advancement in verified local robustness training mechanisms. •Our preliminary results on applying static-mask-based *sparse-training* paradigm based on 'Lottery Ticket Hypothesis' (Frankle & Carbin, 2019) for verified locally robust sparse NN search show inferior performance for this domain, suggesting the requirement for further investigation leveraging more sophisticated approaches for finding the *lottery ticket*. •Finally, following most studies in literature, this paper addresses up to 99% model sparsification. However, our analysis shows on extreme sparsity (99.9%), generalizability decrease, indicating optimal sparsity is between that range (99-99.9%). A detailed discussion is in Appendix I. Though finding optimal sparsity is out of scope, future work on this scope would be beneficial.

8. Conclusion

This is the first paper demonstrating the existence of *verified locally robust sparse NNs* comparable to their dense counterparts. Presented SparseVLR leverages a dynamic-mask-based mechanism to obtain highly sparse NNs exhibiting high accuracy and verified local robustness simultaneously. Our in-depth empirical analysis establishes SparseVLR's high efficacy. The evaluations on a large variety of datasets and models demonstrate SparseVLR's generalizability, thus, enabling the deployment of verified locally robust models in resource-constrained safety-critical systems.

References

Banner, R., Hubara, I., Hoffer, E., and Soudry, D. Scalable methods for 8-bit training of neural networks. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.

Bastani, O., Ioannou, Y. A., Lampropoulos, L., Vytiniotis, D., Nori, A. V., and Criminisi, A. Measuring neural net robustness with constraints. In *NIPS*, 2016.

Chmiel, B., Ben-Uri, L., Shkolnik, M., Hoffer, E., Banner, R., and Soudry, D. Neural gradients are near-lognormal: improved quantized and sparse training. In *International Conference on Learning Representations*, 2021.

Dai, X., Yin, H., and Jha, N. K. Nest: A neural network synthesis tool based on a grow-and-prune paradigm. *IEEE Transactions on Computers*, 68(10):1487–1497, oct 2019. ISSN 1557-9956. doi: 10.1109/TC.2019.2914438.

Elsken, T., Metzen, J. H., and Hutter, F. Neural architecture search: A survey. *J. Mach. Learn. Res.*, 20(1):1997–2017, jan 2019. ISSN 1532-4435.

Evci, U., Ioannou, Y., Keskin, C., and Dauphin, Y. Gradient flow in sparse neural networks and how lottery tickets win. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(6):6577–6586, Jun. 2022. doi: 10.1609/aaai.v36i6.20611.

Frankle, J. and Carbin, M. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *International Conference on Learning Representations*, 2019.

Frankle, J., Dziugaite, G. K., Roy, D. M., and Carbin, M. Stabilizing the lottery ticket hypothesis, 2019.

Frankle, J., Dziugaite, G. K., Roy, D., and Carbin, M. Linear mode connectivity and the lottery ticket hypothesis. In *International Conference on Machine Learning*, pp. 3259–3269. PMLR, 2020.

Fromherz, A., Leino, K., Fredrikson, M., Parno, B., and Pasareanu, C. Fast geometric projections for local robustness certification. In *International Conference on Learning Representations*, 2021.

Ghosh, S., Prasad, K., Dai, X., Zhang, P., Wu, B., Cormode, G., and Vajda, P. Pruning compact convnets for efficient inference, 2022.

Gong, Y., Liu, L., Yang, M., and Bourdev, L. D. Compressing deep convolutional networks using vector quantization. *ArXiv*, abs/1412.6115, 2014.

Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. *CoRR*, abs/1412.6572, 2015.

Gowal, S., Dvijotham, K., Stanforth, R., Bunel, R., Qin, C., Uesato, J., Arandjelovic, R., Mann, T., and Kohli, P. On the effectiveness of interval bound propagation for training verifiably robust models, 2018.

Han, S., Mao, H., and Dally, W. J. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*, 2015a.

Han, S., Pool, J., Tran, J., and Dally, W. J. Learning both weights and connections for efficient neural networks. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*, NIPS'15, pp. 1135–1143, Cambridge, MA, USA, 2015b. MIT Press.

Han, S., Liu, X., Mao, H., Pu, J., Pedram, A., Horowitz, M. A., and Dally, W. J. Eie: Efficient inference engine on compressed deep neural network, 2016.

Hinton, G., Dean, J., and Vinyals, O. Distilling the knowledge in a neural network. pp. 1–9, 03 2014.

Hoefler, T., Alistarh, D., Ben-Nun, T., Dryden, N., and Peste, A. Sparsity in deep learning: Pruning and growth for efficient inference and training in neural networks. *J. Mach. Learn. Res.*, 22(1), jan 2021. ISSN 1532-4435.

Huang, X., Kwiatkowska, M., Wang, S., and Wu, M. Safety verification of deep neural networks. In Majumdar, R. and Kunčak, V. (eds.), *Computer Aided Verification*, pp. 3–29, Cham, 2017. Springer International Publishing. ISBN 978-3-319-63387-9.

Hubara, I., Courbariaux, M., Soudry, D., El-Yaniv, R., and Bengio, Y. Quantized neural networks: Training neural networks with low precision weights and activations. *J. Mach. Learn. Res.*, 18(1):6869–6898, jan 2017. ISSN 1532-4435.

Jia, R., Raghunathan, A., Göksel, K., and Liang, P. Certified robustness to adversarial word substitutions, 2019.

Katz, G., Barrett, C., Dill, D. L., Julian, K., and Kochenderfer, M. J. Reluplex: An efficient smt solver for verifying deep neural networks. In Majumdar, R. and Kunčak, V. (eds.), *Computer Aided Verification*, pp. 97–117. Springer International Publishing, 2017. ISBN 978-3-319-63387-9.

Kurakin, A., Goodfellow, I., and Bengio, S. Adversarial examples in the physical world, 2016. URL https://arxiv.org/abs/1607.02533.

Lee, B.-K., Kim, J., and Ro, Y. M. Masking adversarial damage: Finding adversarial saliency for robust and sparse network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 15126–15136, June 2022.

Li, H., Kadav, A., Durdanovic, I., Samet, H., and Graf, H. P. Pruning filters for efficient convnets. *ArXiv*, abs/1608.08710, 2017.

Liu, S., Yin, L., Mocanu, D. C., and Pechenizkiy, M. Do we actually need dense over-parameterization? in-time over-parameterization in sparse training. In Meila, M. and Zhang, T. (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 6989–7000. PMLR, 18–24 Jul 2021.

Liu, Z., Li, J., Shen, Z., Huang, G., Yan, S., and Zhang, C. Learning efficient convolutional networks through network slimming. *2017 IEEE International Conference on Computer Vision (ICCV)*, pp. 2755–2763, 2017.

Liu, Z., Sun, M., Zhou, T., Huang, G., and Darrell, T. Rethinking the value of network pruning. In *International Conference on Learning Representations*, 2019.

Lyu, Z., Ko, C.-Y., Kong, Z., Wong, N., Lin, D., and Daniel, L. Fastened crown: Tightened neural network robustness certificates. *ArXiv*, abs/1912.00574, 2020.

Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018.

Malach, E., Yehudai, G., Shalev-Schwartz, S., and Shamir, O. Proving the lottery ticket hypothesis: Pruning is all you need. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 6682–6691. PMLR, 13–18 Jul 2020.

N J Karthika, C. S. Addressing false positives in pedestrian detection. In *International Conference on Electronic Systems and Intelligent Computing (ESIC 2020)*, 3 2020. doi: https://doi.org/10.1007/978-981-15-7031-5_103.

Paganini, M. Pruning tutorial. https://pytorch.org/tutorials/intermediate/pruning_tutorial.html#pruning-tutorial, 2022.

Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. Automatic differentiation in pytorch. 2017.

Pereira, A. and Thomas, C. Challenges of machine learning applied to safety-critical cyber-physical systems. *Machine Learning and Knowledge Extraction*, 2(4):579–602, 2020. ISSN 2504-4990. doi: 10.3390/make2040031. URL https://www.mdpi.com/2504-4990/2/4/31.

Pulina, L. and Tacchella, A. Challenging smt solvers to verify neural networks. *AI Commun.*, 25(2):117–135, apr 2012. ISSN 0921-7126.

Qin, C., Martens, J., Gowal, S., Krishnan, D., Dvijotham, K. D., Fawzi, A., De, S., Stanforth, R., and Kohli, P. *Adversarial Robustness through Local Linearization*. Curran Associates Inc., Red Hook, NY, USA, 2019.

Sainath, T. N., Kingsbury, B., Sindhwani, V., Arisoy, E., and Ramabhadran, B. Low-rank matrix factorization for deep neural network training with high-dimensional output targets. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 6655–6659, 2013. doi: 10.1109/ICASSP.2013.6638949.

Sanh, V., Wolf, T., and Rush, A. M. Movement pruning: Adaptive sparsity by fine-tuning. In *NeurIPS*, 2020a.

Sanh, V., Wolf, T., and Rush, A. M. Movement pruning: Adaptive sparsity by fine-tuning, 2020b.

Savarese, P., Silva, H., and Maire, M. Winning the lottery with continuous sparsification, 2020.

See, A., Luong, M.-T., and Manning, C. D. Compression of neural machine translation models via pruning, 2016.

Sehwag, V., Wang, S., Mittal, P., and Jana, S. Towards compact and robust deep neural networks. *preprint arXiv:1906.06110*, 2019.

Sehwag, V., Wang, S., Mittal, P., and Jana, S. Hydra: Pruning adversarially robust neural networks. *34th Conference on Neural Information Processing Systems (NeurIPS 2020)*, 2020.

Singh, G., Gehr, T., Püschel, M., and Vechev, M. An abstract domain for certifying neural networks. *Proc. ACM Program. Lang.*, 3(POPL), jan 2019. doi: 10.1145/3290354.

Singh, S. P. and Alistarh, D. Woodfisher: Efficient second-order approximation for neural network compression. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 18098–18109. Curran Associates, Inc., 2020.

Sinha, A., Namkoong, H., and Duchi, J. Certifiable distributional robustness with principled adversarial training. In *International Conference on Learning Representations*, 2018.

Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C. D., Ng, A., and Potts, C. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pp. 1631–1642, Seattle, Washington, USA, October 2013. Association for Computational Linguistics. URL https://aclanthology.org/D13-1170.

Tan, C. M. J. and Motani, M. DropNet: Reducing neural network complexity via iterative pruning. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 9356–9366. PMLR, 13–18 Jul 2020.

Tjeng, V., Xiao, K. Y., and Tedrake, R. Evaluating robustness of neural networks with mixed integer programming. In *International Conference on Learning Representations*, 2019.

Wang, S., Zhang, H., Xu, K., Lin, X., Jana, S., Hsieh, C.-J., and Kolter, J. Z. Beta-CROWN: Efficient bound propagation with per-neuron split constraints for neural network robustness verification. In Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*, 2021.

Wang, Y., Zou, D., Yi, J., Bailey, J., Ma, X., and Gu, Q. Improving adversarial robustness requires revisiting misclassified examples. In *International Conference on Learning Representations*, 2020.

Wild, C. M. Know your adversary: Understanding adversarial examples (part 1/2). https://towardsdatascience.com/know-your-adversary-understanding-adversarial-examples-part-1-2-63af4c2f5830, 2018. Published: 2018-01-23.

Wong, E. and Kolter, Z. Provable defenses against adversarial examples via the convex outer adversarial polytope. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 5286–5295. PMLR, 10–15 Jul 2018.

Wong, E., Schmidt, F., Metzen, J. H., and Kolter, J. Z. Scaling provable adversarial defenses. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.

Xu, K., Shi, Z., Zhang, H., Wang, Y., Chang, K.-W., Huang, M., Kailkhura, B., Lin, X., and Hsieh, C.-J. Automatic perturbation analysis for scalable certified robustness and beyond. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS'20, Red Hook, NY, USA, 2020. Curran Associates Inc. ISBN 9781713829546.

Xu, K., Zhang, H., Wang, S., Wang, Y., Jana, S., Lin, X., and Hsieh, C.-J. Fast and complete: Enabling complete neural network verification with rapid and massively parallel incomplete verifiers. In *International Conference on Learning Representations*, 2021a.

Xu, K., Zhang, H., Wang, S., Wang, Y., Jana, S., Lin, X., and Hsieh, C.-J. Fast and complete: Enabling complete neural network verification with rapid and massively parallel incomplete verifiers. In *International Conference on Learning Representations*, 2021b.

You, Z., Yan, K., Ye, J., Ma, M., and Wang, P. Gate decorator: Global filter pruning method for accelerating deep convolutional neural networks, 2019.

Zhang, H., Weng, T.-W., Chen, P.-Y., Hsieh, C.-J., and Daniel, L. Efficient neural network robustness certification with general activation functions. *ArXiv*, abs/1811.00866, 2018a.

Zhang, H., Chen, H., Xiao, C., Gowal, S., Stanforth, R., Li, B., Boning, D. S., and Hsieh, C. Towards stable and efficient training of verifiably robust neural networks. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020a.

Zhang, J., Xu, X., Han, B., Niu, G., Cui, L., Sugiyama, M., and Kankanhalli, M. Attacks which do not kill training make adversarial learning stronger. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 11278–11287. PMLR, 13–18 Jul 2020b.

Zhang, T., Ye, S., Zhang, K., Tang, J., Wen, W., Fardad, M., and Wang, Y. A systematic dnn weight pruning framework using alternating direction method of multipliers. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 184–199, 2018b.

A. Experiment Details

Datasets: The empirical analysis of SparseVLR is done for 5 datasets: MNIST, CIFAR, SVHN, Pedestrian Detection and Sentiment Analysis SST-2.

MNIST is a dataset of hand-written digits with 60000 samples in the training set and 10,000 samples in the testing set. All the images are in greyscale and have a size of 28×28. The training set is generated using samples from approximately 250 writers. The writers for the test set and the training set are disjoint.

CIFAR-10 is a dataset of 60000 images evenly distributed among 10 mutually exclusive classes: airplanes, cars, birds, cats, deer, dogs, frogs, horses, ships, and trucks. The training set and the testing set consists of 50000 and 10000 images, respectively. These RGB images are of size 32×32 each. Additionally, The test set contains 1000 randomly selected images from each class, thus having a uniform distribution over all the classes.

SVHN is a real-world dataset for digit recognition. The images are Street View House Numbers obtained from google street view images. The format used for the evaluations for SparseVLR contains 32×32 images of individual digits with distractions. The dataset contains 73257 images in the training set and 26032 images for testing.

Pedestrian Detection (N J Karthika, 2020) aims to differentiate between the person and person-like objects such as statues, scarecrows, etc., having very similar features to a person. The number of images in train, validate and test set are 944, 160, and 235, respectively, with a total of 1626 persons and 1368 non-human labeling.

SparseVLR is applicable to domains other than image classification. To illustrate this, an empirical study has been done on sentiment analysis NLP dataset SST-2 (Socher et al., 2013) used in the previous formal verification approaches by (Xu et al., 2020; Jia et al., 2019). The dataset contains 10662 sentences evenly distributed between two classes: positive and negative.

Perturbation amount: For image classification datasets, the perturbations are introduced in ℓ_p-ball of radius ε (perturbation amount), $\mathbb{B}_p(x_0, \varepsilon)$ (defined in Section 2). The highest perturbation amounts (i.e., most difficult scenarios) used in the literature addressing formal verification and training (Gowal et al., 2018; Zhang et al., 2020a; Wong et al., 2018; Xu et al., 2020) are 0.4 and 8/255 for MNIST and CIFAR10, respectively. For CIFAR10, Zhang et al. (2020a) uses different perturbation amounts for training ($\varepsilon_{\text{train}}$) and testing ($\varepsilon_{\text{test}}$), such that $\varepsilon_{\text{train}} = 1.1 \times \varepsilon_{\text{test}}$, we follow the same convention. Thus, for the models that need to be tested for 8/255, the training is done for a perturbation amount of 8.8/255. For SVHN, the previous works targetting adversarial robustness (Sehwag et al., 2020), use perturbation amount of 8/255, thus, following the convention of CIFAR-10, we use $\varepsilon_{\text{test}} = 8/255$ and $\varepsilon_{\text{train}} = 8.8/255$. For Pedestrian Detection dataset, we use $\varepsilon_{\text{test}} = 2/255$ and $\varepsilon_{\text{train}} = 2.2/255$.

For Sentiment analysis NLP dataset, the perturbations are introduced in terms of synonym-based word substitution in a sentence. Each word w has a set of synonyms $\mathbb{S}(w)$ with which it can be substituted. $\mathbb{S}(w)$ is computed using 8 nearest neighbors in counter-fitted word embedding, where the distance signifies similarity in meaning (Jia et al., 2019; Xu et al., 2020), lesser distance more similarity and vice-versa. The number of words being substituted is referred to as *budget* (δ). The maximum budget used in state-of-the-art (Xu et al., 2020; Jia et al., 2019) for formal verification is 6, and we use the same amount.

Hyperparameters: The presented approach requires a set of hyperparameters as inputs: (T, s, l), where T is the total number of training epochs, s is the epoch number at which the perturbation scheduler should start to increment the amount of perturbation and l specifies the length of the schedule, that is, the number of epochs in which the perturbation scheduler has to reach the maximum amount of perturbation ε_{max}.

These values are as follows: (1) For 4-layer CNN, 7-layer CNN, and Resnet, we use $(T, s, l) = (100, 10, 60)$ and $(T, s, l) = (2000, 200, 800)$, for MNIST and CIFAR-10, respectively; (2) For LSTM trained for MNIST, we use $(T, s, l) = (20, 1, 10)$; (3) For SVHN, the hyperparameters are same as that of CIFAR-10. (4) For Pedestrian Detection, the hyperparameters $(T, s, l) = (100, 20, 60)$ resulted in the models displaying the least errors empirically; and (5) For Sentiment Analysis NLP dataset, the hyperparameters $(T, s, l) = (25, 1, 10)$ are used. The value of (T, s, l) for MNIST, CIFAR-10, and Sentiment Analysis are in compliance with state-of-the-art (Xu et al., 2020). For Pedestrian Detection and SVHN no prior hyperparameter values are available.

Learning rate: The training approach uses the learning rate scheduler, which gradually decreases the learning rate as the training proceeds after the maximum perturbation is reached. The starting learning rate is set to 10^{-4} as suggested by the state-of-the-art (Xu et al., 2020).

B. Formal Verification

An NP-Complete problem: Linear Programming (Bastani et al., 2016) and Satisfiability Modulo Theory (SMT) (Huang et al., 2017) have been used in literature as formal verification techniques to validate neural networks. However, Pulina and Tacchella (Pulina & Tacchella, 2012) showed that Multi-Layer Perceptrons (MLPs) cannot be verified using SMT solvers, thus, challenging the scalability of SMT solvers to neural networks of realistic sizes. Furthermore, several applications require non-linear activation functions to learn complex decision boundaries, rendering the whole NN non-linear. The large size and non-linearity of NNs make the problem of formal robustness verification non-convex and NP-Complete (Katz et al., 2017; Sinha et al., 2018).

Relaxation based Methods: To address the challenges mentioned above, Katz et al. (2017) used relaxation of ReLU activations, temporarily violating the ReLU semantics, resulting in an efficient query solver for NNs with ReLU activations. The piece-wise linearity of the ReLU activations is the basis for this relaxation. LiRPA based methods such as CROWN (Zhang et al., 2018a), IBP (Gowal et al., 2018), CROWN-IBP (Zhang et al., 2020a), etc. compute linear relaxation of a model containing general activation functions which are not necessarily piece-wise linear.

CROWN (Zhang et al., 2018a) uses a backward bounding mechanism to compute linear or quadratic relaxation of the activation functions. IBP (Gowal et al., 2018) is a forward bounding mechanism that propagates the perturbations induced to the input towards the output layer in a forward pass based on interval arithmetic, resulting in ranges of values for each class in the output layer. CROWN-IBP (Zhang et al., 2020a) is a hybrid approach that demonstrated that IBP can produce loose bounds, and in order to achieve tighter bounds, it uses IBP (Gowal et al., 2018) in the forward bounding pass and CROWN (Zhang et al., 2018a) in the backward bounding pass to compute bounding hyperplanes for the model. CROWN-IBP has been shown to compute tighter relaxations, thus, providing the most optimal results for formal verification.

As discussed in Section 2, the linear relaxation computed using LiRPA methods can be used to train verified locally robust models by minimizing L_{train}. This learning procedure uses a perturbation scheduler to gradually increase the amount of perturbation during training.

B.1. Perturbation Scheduler (ε-**scheduler**$(\varepsilon_{\mathbf{max}}, t, s, l)$)

ε-scheduler aims to provide a perturbation amount for every training epoch $t \geq s$, which gradually increases ε starting at epoch s. The schedule starts with a 0 perturbation and reaches ε_{max} in l epochs (Gowal et al., 2018). Gowal et al. (2018) proposed using $\varepsilon-$scheduler for training based on LiRPA and deems it necessary for an effective learning procedure. Since IBP (Gowal et al., 2018) uses interval propagation arithmetic to propagate input perturbation to the output layer, where the interval size usually keeps on increasing as the propagations reach deeper in the model, the gradual increase of epsilon prevents the problem of intermediate bound explosion while training. Other related approaches also adopted the idea of gradual perturbation increment (Zhang et al., 2020a; Xu et al., 2020). Since the bounding mechanism used in SparseVLR is CROWN-IBP (Zhang et al., 2020a), the use of ε-scheduler is crucial.

C. Model Compression

Several model compression approaches have been proposed to reduce the computational and storage requirements of deep neural networks (DNNs). Quantization aims to convert the parameter values to low-precision approximations requiring less storage bits per parameter (Hubara et al., 2017; Gong et al., 2014; Banner et al., 2018; Chmiel et al., 2021). Model Distillation (Hinton et al., 2014) and Neural Architecture Search (Elsken et al., 2019) train a smaller dense architecture that generalizes comparable to the dense network. Low-Rank Factorization (Sainath et al., 2013) computes a matrix decomposition of the weight matrices, which requires fewer floating point operations. Model Sparsification/Pruning is the most common form of model compression, which constrains the model to use only a subset of original parameters for inference (Han et al., 2015b; Sehwag et al., 2019; Zhang et al., 2018b; Han et al., 2015a; Li et al., 2017; Sanh et al., 2020a; Sehwag et al., 2020).

C.1. Model sparsification.

Hoefler et al. (2021) categorizes the model sparsification according to the stage at which sparsification is performed: *train-and-sparsify* (After training), *sparsifying-during-training* (While Training) and *sparse-training* (Before training).

Train-and-sparsify approaches (Han et al., 2015a; Zhang et al., 2018b; Sehwag et al., 2019; 2020) train a dense model and
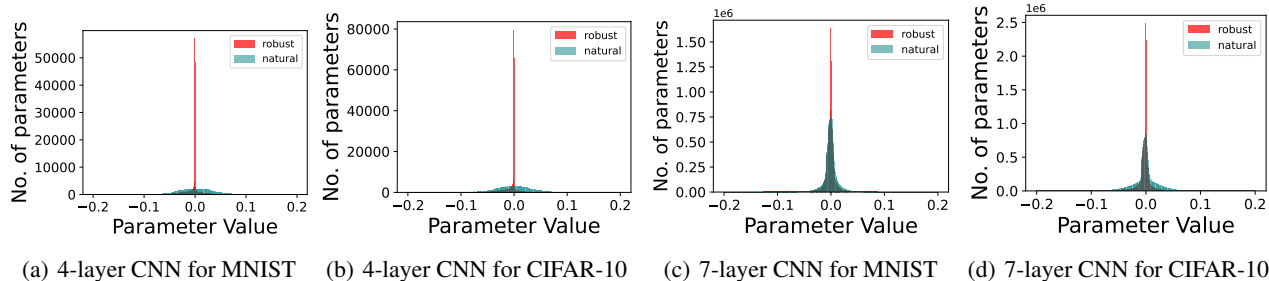
(a) 4-layer CNN for MNIST   (b) 4-layer CNN for CIFAR-10   (c) 7-layer CNN for MNIST   (d) 7-layer CNN for CIFAR-10

Figure 5. Weight distributions for (a) 4-layer CNN (MNIST) (b) 4-layer CNN (CIFAR-10) (b) 7-layer CNN (MNIST)(d) 7-layer CNN (CIFAR-10). The epsilon amounts used for robust training are 0.4 and 8.8/255 for MNIST and CIFAR-10, respectively.

remove (or prune) the parameters contributing the least towards model inference. Since the removal of parameters may lead to loss of information learned by the original dense model, the removal of parameters is often followed by re-training the retained parameters, also known as the finetuning step. *Train-and-sparsify* associates a static binary mask with the parameters, which allows the finetuning step to update only the retained parameters.

The *sparsifying-during-training* mechanism gradually removes NN model parameters in small fractions per round. Each round is followed by a finetuning step allowing the model to recover from the parameter removal step (Tan & Motani, 2020; You et al., 2019). This is often called *iterative pruning*.

The *sparse-training* mechanism involves training a sparse network from scratch by using either a *static* or a *dynamic* mask. Identifying *static* masked sparse network is motivated by *Lottery Ticket Hypothesis (LTH)* (Frankle & Carbin, 2019) which states that "dense, randomly-initialized, feed-forward networks contain subnetworks (winning tickets) that - when trained in isolation - reach test accuracy comparable to the original network in a similar number of iterations". However, finding the *static* mask in LTH (Frankle et al., 2020; Frankle & Carbin, 2019; Frankle et al., 2019; Savarese et al., 2020; Malach et al., 2020) still requires a fully or partially trained dense model to start with.

Recently, *dynamic-mask-based sparse-training* mechanisms known as Dynamic Sparse Training (DST) have been developed. DST approaches (Dai et al., 2019; Liu et al., 2021) start with a randomly initialized compressed sparse network and explore new network connections while learning network weights for the target task during training. SparseVLR adapts a DST approach to identify verified locally robust sparse network from scratch.

C.2. Selecting the parameters to be removed:

While all sparsification approaches mentioned above perform network parameters (i.e., weight) removal, selecting such parameter weights is done based on different criteria such as the magnitude of weights, magnitude of corresponding gradients, minimal impact on Hessian matrix, etc. (Hoefler et al., 2021). Magnitude-based pruning is the most popular pruning mechanism (Han et al., 2015b; Sehwag et al., 2019; Zhang et al., 2018b; Han et al., 2015a; Li et al., 2017; Sanh et al., 2020a; Sehwag et al., 2020; Han et al., 2016), which removes the least magnitude parameters suggesting that the importance of a parameter is directly proportional to the magnitude of their weights.

Moreover, removal of parameters can be done either per-layer or globally, whereas recent studies (See et al., 2016; Ghosh et al., 2022; Liu et al., 2017; 2019; Singh & Alistarh, 2020) suggest the higher efficacy of global unstructured pruning. Sanh et al. (2020b) demonstrated that global magnitude-based pruning tends to prune earlier layers of a network by a lesser amount. Since the earlier layers for most networks are feature extraction layers which hold more importance for a classification task, it can be concluded that global pruning captures the importance of a layer as a whole.

D. CROWN-IBP causes regularization

The motivation for SparseVLR is that a huge fraction of parameters of locally robust models trained using CROWN-IBP have low magnitudes, so removal of such less significant parameters has minimal effect on the model generalizability (Section 3).

Figure 5 shows the weight distribution of models trained to minimize natural, and verified locally robust (CROWN-

IBP ([Zhang et al., 2020a](#))) losses. The trends are shown for 4-layer and 7-layer CNNs trained for MNIST and CIFAR-10 datasets. As evident from the distribution of weights, a higher fraction of weights in verified locally robust models have very low magnitudes ($\approx 10^{-40}$). This observation suggests that models trained to minimize $\mathcal{L}_{\text{train}}$ (Eq. 4) penalize the parameter magnitudes.

E. Explaining the generalizability difference for various sparse training mechanisms

E.1. Conventional Finetuning vs Grow-and-Prune

Conventional finetuning aiming to minimize $\mathcal{L}_{\text{train}}$ (($\#ii$) in Section 3) increases the *standard* and *verified* error even at small compression amounts. Such increment can be attributed to poor gradient flow encountered during static mask-based finetuning, which is in agreement with ([Evci et al., 2022](#)). However, retraining a pruned model using *grow-and-prune* (($\#iii$) in Section 3) does not suffer from poor gradient flow, hence may help maintain the generalizability of the sparse compressed models.

To further investigate, Figure 6 compares the losses and ℓ_1-norm of gradients encountered while re-training a globally pruned network using conventional static mask-based finetuning (while keeping the mask the same throughout the training), and grow-and-prune mechanism ([Dai et al., 2019](#)). The results are shown for a 4-layer CNN trained for CIFAR-10 at 90% sparsity. It is important to note that the training procedure uses ε-scheduling, which provides 0 perturbation for initial s epochs and then gradually increases the perturbation value. This results in an initial drop in losses, as seen in Figure 6(b). As the perturbation amount increases, the loss increases as well and finally stabilizes.

Notably, the losses encountered by finetuning are higher than retraining using grow-and-prune, thus resulting in higher errors in the obtained compressed models (shown

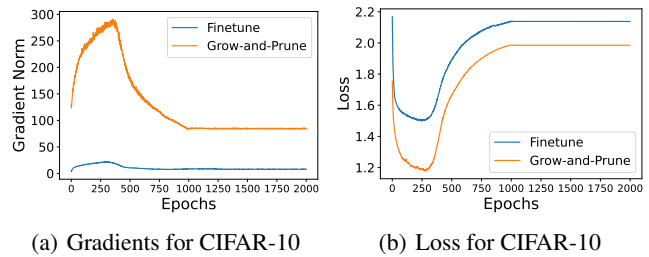(a) Gradients for CIFAR-10   (b) Loss for CIFAR-10

Figure 6. (a) Gradient flow and (b) loss during training a globally pruned 4-layer CNN using fine-tuning, and re-training using grow-and-prune CIFAR-10 datasets for perturbation of 0.4 and 8/255, respectively

in Figure 6(b)). The difference in losses encountered using the two approaches can be attributed to the difference in gradient flows (see Figure 6(a)). Re-training using *grow-and-prune* incurs higher gradient flows than conventional finetuning, and according to [Evci et al. (2022)](#), higher gradient flows lead to efficient training.

E.2. Training PM$\rightarrow$PM^T vs RSM$\rightarrow$RSM^T

As discussed in Section 5, the effectivenes of Sparse-VLR is established by comparing the sparse models obtained by applying Algorithm 1 (steps 3-9) to different seed networks: PM (Pruned Model) and RSM (Random Sparse Model). Table 2 compares the sparse models obatained using three mechanism: PM (Global Pruning), PM^T(Retaining PM using Algorithm 1 steps 3-9.) and RSM^T(obatined using SparseVLR). The results are shown for 3 backbone architectures: 4-layer CNN, 7-layer CNN and Resnet at 4 sparsity amounts: 80%, 90%, 95% and 99%; trained for CIFAR-10. It can be noted that RSM^{T}outperforms both PM and PM^{T}especially at high sparsity amount such at 99%.

(a) Sparsity = 95%   (b) Sparsity = 99%

Figure 7. Variation of training loss for CIFAR-10 dataset on 4-layer CNN with $\varepsilon_{\text{test}} = 8/255$ at two sparsity amounts (a) 95% and (b) 99%

To investigate the difference in losses encountered while training PM$\rightarrow$PM^T and RSM$\rightarrow$RSM^T at 99% sparsity, Figure 7 (same as Figure 2) shows the difference in loss evolution for both the scenarios. It can be noted that PM$\rightarrow$PM^{T}incurs higher loss than RSM$\rightarrow$RSM^T, thus resulting in higher errors.

Interestingly, at 99% sparsity, the loss evolution for PM→PM^T training does not follow the expected pattern, i.e. an initial decrease followed by an increase (see Figure 7). The reason for this can be attributed to poor gradient flow while training PM→PM^T at 99% sparsity (as shown in Fig. 3(b) in Sec. 5). The low magnitude of gradient results in low magnitude in the newly activated parameters, which eventually get removed during the pruning step. On the other hand, the already active parameters, which are significantly high in magnitudes, tend to become more polarized, thus, resulting in over-fitting. However, with an initial high loss, the gradient flow increases enough to generate some newly activated parameters with sufficient magnitude to replace the already activated parameters, which results in some decrease and eventually stabilization of the loss.

F. Evaluation for larger models

Table 5. Error metrics for sparse models obtained using SparseVLR using DenseNet and ResNeXt as backbone architecture trained for (a) MNIST and (b) CIFAR at 99% sparsities

Dataset	MNIST				CIFAR-10			
Model	DenseNet		ResNeXt		DenseNet		ResNeXt	
Sparsity ⇒	0%	99%	0%	99%	0%	99%	0%	99%
Error								
Standard	9.02	7.88	4.65	3.55	63.69	63.43	68.77	64.81
PGD	12.39	8.41	12.29	3.85	70.03	71.4	73.32	71.21
Verified	21.73	21.06	12.66	10.08	71.17	72.76	75.09	72.38

Table 5 demonstrates that sparse models obtained for more complex backbone architectures such as DenseNet and ResNeXt (provided by (Xu et al., 2020)) trained for MNIST and CIFAR-10 generalize comparable to the dense models having same architecture.

G. Application to Sequential Model

To further investigate the application of SparseVLR to sequential networks, we evaluated a small LSTM (Xu et al., 2020) as a backbone architecture to obtain sparse models for MNIST. Table 6 demonstrates the *standard*, *PGD*, and *verified* errors for the sparse models obtained using SparseVLR. It can be noted that up to 95% sparsity amount, the sparse models generalize comparable to their dense counterparts. Additionally, higher error at 99% sparsity can be attributed to the insufficient number of parameters for the target task.

Table 6. Error metrics for sparse models obtained using SparseVLR using LSTM trained for MNIST at various sparsity amounts

Sparsity ⇒	0%	80%	90%	95%	99%
Error					
Standard	4.04	4.42	4.38	5.52	11.55
PGD	9.35	11.60	10.10	13.84	17.22
Verified	14.23	11.85	10.75	13.90	18.70

H. Application to NLP for sentiment analysis

As discussed in Appendix A, for the Sentiment Analysis dataset, the perturbations are introduced in terms of synonym-based word substitution in a sentence, and the number of words substituted is called the *budget* δ. The value of δ used in the empirical analysis is 6.

Notably, this training does not use ℓ_p-norm pertubation, but to keep the perturbation amount continuous, the training procedure consists of an initial warm-up phase (Xu et al., 2020). If the word at index i for a clean sentence and a perturbed sentence are represented by w_i and $\hat{w}_i$, and $e(w)$ is the embedding for the word w, the effective embedding used during training is given by:

$$e(w_i) = \epsilon * e(\hat{w}_i) + (1 - \epsilon) * e(w_i)$$

Table 7. Error metrics for sparse models obtained using our approach using LSTM as the backbone architecture trained for NLP dataset at different sparsity amounts

Sparsity ⇒	0%	80%	90%	95%	99%
Error					
Standard	20.92	20.26	22.24	21.11	21.53
Verified	24.27	22.35	23.94	23.53	23.72

The value of ϵ is gradually increased form 0 to 1 during the training phase. The results thus produced for the original

dense model (sparsity = 0%) are in compliance with state-of-the-art (Xu et al., 2020). Table 7 shows that the sparse models obtained at different compression amounts using the presented approach exhibit *standard* and *verified* errors comparable to the original dense model. PGD attack is based on ℓ_p-norm perturbation which is not applicable to this domain, so the results shown in Table 7 does not include results for *PGD* error.

I. Extreme sparsity:

The results presented in Table 2 show that the 99% sparse models generalize comparable to the dense model, thus, indicating the possiblity of exploring further sparsity. However, at 99.9% sparsity the 7-layer CNN exhibit an increment of approximately 7% in *standard error* and 10% in *verified error* for both MNIST and CIFAR. This suggests that an optimal sparsity exists in the range of (99% - 99.9%) at which the sparse model generalizes comparable to its dense counterpart.

J. Computation time benefit

In addition to reduction of non-zero (active) parameters of a model, we observe that the the sparse models obtained using our approach incur much less computation time. We measure the computation time in terms of average inference time on NVIDIA RTX A6000 required per sample in a dataset. The computation time shown in Table 8 are computed for CIFAR-10 dataset for three models: 4 layer CNN, 7 layer CNN and Resnet at different sparsity amounts. It can be noted that the inference time required by the sparse models obtained using SparseVLR is on an average 5 times less than the inference time of the dense models.

Table 8. Inference Time (in sec) at different sparsity amounts

Compression ⇒ **Model**	0%	80%	90%	95%	99%
4 Layer CNN	1.8×10^{-5}	4.9×10^{-6}	5.2×10^{-6}	7.5×10^{-6}	6.4×10^{-6}
7 Layer CNN	1.1×10^{-4}	1.8×10^{-5}	1.8×10^{-5}	1.8×10^{-5}	1.7×10^{-5}
Resnet	1.1×10^{-4}	2.5×10^{-5}	2.5×10^{-5}	2.9×10^{-5}	2.8×10^{-5}

K. Using Different bounding mechanisms for verified local robustness

Table 9 compares the errors exhibited by 4-layer CNN at 99% sparsity when trained by employing different bounding mechanisms to compute $\mathcal{L}_{\text{train}}$: IBP (Gowal et al., 2018), CROWN (Zhang et al., 2018a), Fastened-CROWN (Lyu et al., 2020), and CROWN-IBP (Zhang et al., 2020a).

Table 9. Error metrics for different formal verification mechanism used to compute $\mathcal{L}_{\text{train}}$ on 4 layer CNN trained for MNIST and CIFAR-10 at 99% sparsity

Dataset	Method ⇒ **Error**	IBP	CROWN	CROWN- Fast	CROWN- IBP
MNIST	Standard	7.44	52.3	13.53	7.23
$\varepsilon = 0.4$	PGD	10.27	68.84	18.20	9.68
	Verified	21.54	93.13	42.18	21.95
CIFAR	Standard	62.85	69.04	68.72	65.01
$\varepsilon = 8/255$	PGD	78.21	71.53	81.69	70.43
	Verified	78.52	88.3	89.39	73.38