Partition-Based Convex Relaxations for Certifying the Robustness of ReLU Neural Networks

Brendon G. Anderson[*] BGANDERSON@BERKELEY.EDU
Ziye Ma[†] ZIYEMA@BERKELEY.EDU
Jingqi Li[†] JINGQILI@BERKELEY.EDU
Somayeh Sojoudi[*†] SOJOUDI@BERKELEY.EDU
[*]*Department of Mechanical Engineering, University of California, Berkeley*
[†]*Department of Electrical Engineering and Computer Sciences, University of California, Berkeley*

Abstract

In this paper, we study certifying the robustness of ReLU neural networks against adversarial input perturbations. To diminish the relaxation error suffered by the popular linear programming (LP) and semidefinite programming (SDP) certification methods, we propose partitioning the input uncertainty set and solving the relaxations on each part separately. We show that this approach reduces relaxation error, and that the error is eliminated entirely upon performing an LP relaxation with an intelligently designed partition. To scale this approach to large networks, we consider courser partitions that take the same form as this motivating partition. We prove that computing such a partition that directly minimizes the LP relaxation error is NP-hard. By instead minimizing the worst-case LP relaxation error, we develop a computationally tractable scheme with a closed-form optimal two-part partition. We extend the analysis to the SDP, where the feasible set geometry is exploited to design a two-part partition that minimizes the worst-case SDP relaxation error. Experiments on IRIS classifiers demonstrate significant reduction in relaxation error, offering certificates that are otherwise void without partitioning. By independently increasing the input size and the number of layers, we empirically illustrate under which regimes the partitioned LP and SDP are best applied.

Keywords: ReLU Neural Networks, Robustness Certification, Adversarial Attacks, Convex Optimization, Partitioning

1. Introduction

It is evident that the data used in real-world engineering systems has uncertainty. This uncertainty takes many forms, including corruptions, random measurement noise, and adversarial attacks (Franceschi et al., 2018; Balunovic et al., 2019; Jin et al., 2020). Recently, researchers have shown that the performance of neural networks can be highly sensitive to these uncertainties in the input data (Szegedy et al., 2014; Fawzi et al., 2016; Su et al., 2019). Clearly, safety-critical systems, such as autonomous vehicles (Bojarski et al., 2016; Wu et al., 2017) and the power grid (Kong et al., 2017; Muralitharan et al., 2018; Pan et al., 2019), must be robust against fluctuations and uncertainty in the inputs to their decision-making algorithms. This fact has motivated a massive influx of research studying methods to certify the robustness of neural networks against uncertainty in their input data (Wong and Kolter, 2018; Raghunathan et al., 2018; Xiang and Johnson, 2018; Weng et al., 2018; Zhang et al., 2018; Royo et al., 2019; Fazlyab et al., 2020; Anderson et al., 2020; Anderson and Sojoudi, 2020a,b; Ma and Sojoudi, 2020; Jin et al., 2021).

The two primary settings taken in the robustness certification literature consider either random input uncertainty or adversarial uncertainty. In the former, the neural network input is assumed to be random and follow a known probability distribution. For instance, the works Weng et al. (2019); Anderson and Sojoudi (2020a,b) derive high-probability guarantees that this randomness causes no misclassification or unsafe output. In the adversarial setting, the input is assumed to be unknown but contained in a prescribed input uncertainty set. The goal here is to certify that all possible inputs from the uncertainty set are mapped to outputs that the network operator deems as safe (Wong and Kolter, 2018; Royo et al., 2019). In this paper, we take the latter, worst-case perspective. We remark that this approach is more generally applicable than the techniques for random inputs. Indeed, certifying that a network is robust against adversarial inputs immediately also certifies it is robust against random inputs distributed over the same uncertainty set; if the worst-case input cannot cause a failure, then neither will randomly selected ones.

The problem of adversarial robustness certification amounts to proving that all possible outputs in the output set, i.e., the image of the input uncertainty set under the mapping of the network, are contained in the safe set. However, this output set is generally nonconvex, even when the input set is convex. Consequently, the certification decision problem has been shown to be NP-complete, and the optimization-based formulation for solving it is an NP-hard, nonconvex optimization problem (Katz et al., 2017; Weng et al., 2018). To make the problem more tractable, researchers have proposed various ways to over-approximate the nonconvex output set with a convex one. Performing the certification over the convex surrogate reduces the problem to an easy-to-solve convex relaxation, and if the relaxation issues a robustness certificate for the outer approximation, it also issues a certificate for the true set of outputs.

Perhaps the simplest and most popular outer approximation technique is based on a linear programming (LP) relaxation of the ReLU activation function (Wong and Kolter, 2018). However, this method has been shown to yield relatively loose outer approximations, making it possible for the approximating set to contain unsafe outputs, even if the true output set is entirely safe. If this occurs, the convex relaxation fails to issue a certificate of robustness, despite the fact the network is indeed robust. A semidefinite programming (SDP) relaxation was proposed in Raghunathan et al. (2018) and shown to yield tighter outer approximations when compared to the LP method. Other methods, such as the quadratically-constrained semidefinite program (Fazlyab et al., 2020), use sophisticated relaxations to tighten the approximation, but these SDP-based methods inevitably gain accuracy at the expense of enlarging the computational cost, and are still susceptible to relaxation error.

1.1 Related Works

Instead of resorting to exorbitantly costly relaxation techniques to lower the approximation error, the approach we take in this paper is based on the efficient method of partitioning the input uncertainty set and solving an LP or SDP relaxation on each input uncertainty subset. Partitioning methods are powerful and have been used to tighten bounds on optimization problems in many areas, including robust optimization (Bertsimas and Dunning, 2016) and deep learning (Montufar et al., 2014). A second benefit of partitioning is the ability to parallelize the optimization, reducing the computational overhead compared to more elegant centralized methods.

We focus our attention in this paper on networks with ReLU activation functions. ReLU networks are extremely popular for their non-vanishing gradient property and their quick training times (Xiang and Johnson, 2018). There exist a handful of works in the literature that utilize partitioning in the robustness certification problem of these networks. For example, Xiang and Johnson (2018) uses partitioning to certify the safety of neural network controllers from the lens of reachability analysis. Their partitioning method cuts the input uncertainty set and the outer approximation of the reachable set into hyperrectangles. Although this approach extends to networks with other activations (e.g., tanh, logistic, etc.), it fails to take into account the special piecewise linear structure of the ReLU function, making the outer approximations unnecessarily loose for the popular ReLU networks. On the other hand, Royo et al. (2019) uses Lagrangian duality to propose a partitioning method specialized to ReLU activations. However, that method is restricted to splitting box-shaped uncertainty sets in half along coordinate axes. Finally, Everett et al. (2020) provides theoretical guarantees for the amount of volume reduction in the outer approximation induced by partitioning, but it only considers axis-aligned gridding of the input space. These recent works demonstrate the increased interest in partition-based certification, and their positive results are evidence that partitioning yields tightened approximations with only a modest increase in computational overhead. However, these works do not fully exploit the special ReLU structure at hand when designing the partition, inevitably leading to conservative bounds.

1.2 Contributions

In this paper, we fully exploit the piecewise linear structure of ReLU networks in order to tighten convex relaxations for robustness certification. A condensed summary of our main contributions is as follows:

1. We prove that partitioning is theoretically guaranteed to tighten both the LP and the SDP relaxation.

2. For the LP relaxation, we show that a methodically designed finite partition attains zero relaxation error. We then use the structure of this motivating partition to derive a computationally efficient two-part partitioning scheme that minimizes worst-case relaxation error. Unlike prior works, this optimal partition is not axis-aligned in general.

3. We prove that computing the two-part partition that minimizes the actual LP relaxation error is NP-hard, in turn theoretically justifying our approach of minimizing the worst-case relaxation error.

4. For the SDP relaxation, we develop a geometrically interpretable measure for how far the solution is from being rank-1. We then show that the partition along a given direction that minimizes this rank-1 gap is indeed given by a uniform grid. Finally, for the specialized case of a two-part partition, we derive the partitioning axis that minimizes the worst-case relaxation error.

This paper is a major extension of the conference paper Anderson et al. (2020). Notably, all of the SDP results of this paper, as well as the third and fourth major theoretical contributions from above, are new additions. Furthermore, two new experiments have been added. The contributions of this paper culminate into two theoretically justified partition-based convex relaxations for ReLU robustness certification. Our experiments demonstrate that the proposed approaches are effective and efficient, and from these results we develop general guidance for which network size and structure regimes the LP, partitioned LP, SDP, and partitioned SDP are best applied.

1.3 Outline

This paper is organized as follows. Section 2 introduces the ReLU robustness certification problem as well as its basic LP and SDP relaxations. In Section 3, we study the effect of partitioning the input uncertainty set for the LP relaxation. After developing formal guarantees for its effectiveness, we develop a two-part partitioning strategy that optimally reduces the worst-case relaxation error. Similarly, in Section 4 we study the partitioned SDP relaxation and propose an optimal two-part partitioning scheme. Section 5 demonstrates the developed methods on numerical examples and studies the effectiveness of partitioning the LP and SDP as the network grows in width and depth. Finally, we conclude in Section 6.

1.4 Notations

The set of n-vectors with real-valued elements is written as $\mathbb{R}^n$. The symbol e_i is reserved to denote the i^{th} standard basis vector of $\mathbb{R}^n$ for $i \in \{1,2,\dots,n\}$, i.e., e_i is a vector in $\mathbb{R}^n$ with i^{th} element equal to one and all other elements equal to zero. We denote the n-vector of all ones by $\mathbf{1}_n$, and for an index set $\mathcal{I} \subseteq \{1,2,\dots,n\}$, the symbol $\mathbf{1}_{\mathcal{I}}$ is used to denote the n-vector with $(\mathbf{1}_{\mathcal{I}})_i = 1$ for $i \in \mathcal{I}$ and $(\mathbf{1}_{\mathcal{I}})_i = 0$ for $i \in \mathcal{I}^c = \{1,2,\dots,n\} \setminus \mathcal{I}$. For $x \in \mathbb{R}^n$, we denote its i^{th} element by x_i, or, when necessary for clarity, by $(x)_i$. We write $\max_{i\in\mathcal{I}} x_i$ to mean $\max\{x_i : i \in \mathcal{I}\}$ (that is, the maximum element of x), and we define $\arg\max_{i\in\mathcal{I}} x_i$ to be the set $\{i^* \in \mathcal{I} : x_{i^*} = \max_{i\in\mathcal{I}} x_i\}$. For a function $f \colon \mathbb{R}^n \to \mathbb{R}^m$ and the point $x \in \mathbb{R}^n$, we denote the i^{th} element of the m-vector $f(x)$ by $f_i(x)$.

The set of $m \times n$ matrices with real-valued elements is written as $\mathbb{R}^{m\times n}$, and we denote the set of symmetric $n \times n$ matrices with real elements by $\mathbb{S}^n$. For a matrix $X \in \mathbb{R}^{m\times n}$, we use two indices to denote an element of X and one index to denote a column of X, unless otherwise stated. In particular, the (i,j) element and the i^{th} column of X are denoted by X_{ij} and X_i respectively, or, when necessary for clarity, by $(X)_{ij}$ and $(X)_i$ respectively. For $X, Y \in \mathbb{R}^{m\times n}$, we write $X \le Y$ to mean $X_{ij} \le Y_{ij}$ for all $i \in \{1,2,\dots,m\}$ and all $j \in \{1,2,\dots,n\}$. The Hadamard (element-wise) product between X and Y is written as $X \odot Y$ and the Hadamard division of X by Y as $X \oslash Y$. We write the rank of the matrix X as $\operatorname{rank}(X)$. Furthermore, for a function $f \colon \mathbb{R} \to \mathbb{R}$, we define $f(X)$ to be an $m \times n$ matrix whose (i,j) element is equal to $f(X_{ij})$ for all $i \in \{1,2,\dots,m\}$ and all $j \in \{1,2,\dots,n\}$. If Z is a square $n \times n$ matrix, we use the notation $\operatorname{diag}(Z)$ to mean the n-vector $(Z_{11}, Z_{22}, \dots, Z_{nn})$ and $\operatorname{tr}(Z)$ to mean the trace of Z. When $Z \in \mathbb{S}^n$, we write $Z \succeq 0$ to mean that Z is positive semidefinite.

The ReLU function is defined as $\mathrm{ReLU}(\cdot) = \max\{0, \cdot\}$. We use $\mathbb{I}$ to denote the indicator function, i.e., $\mathbb{I}(A) = 1$ if event A holds and $\mathbb{I}(A) = 0$ if event A does not hold. For a set $\mathcal{S}$, we denote its cardinality by $|\mathcal{S}|$. Recall that the infimum and supremum of a set $\mathcal{T} \subseteq \mathbb{R}$ are the set's greatest lower bound and least upper bound, respectively. For the set $\mathcal{T}$ we respectively denote its infimum and supremum by $\inf \mathcal{T}$ and $\sup \mathcal{T}$. For a function $f \colon \mathbb{R}^n \to \mathbb{R}$ and a set $\mathcal{X} \subseteq \mathbb{R}^n$, we write $\inf_{x \in \mathcal{X}} f(x)$ to mean $\inf\{f(x) : x \in \mathcal{X}\}$ and similarly for suprema. Finally, we assume that all infima and suprema throughout the paper are attained.

2. Problem Statement

2.1 Description of the Network and Uncertainty

In this paper, we consider a pre-trained K-layer ReLU neural network defined by

$$\begin{aligned} x^{[0]} &= x, \\ \hat{z}^{[k]} &= W^{[k-1]} x^{[k-1]} + b^{[k-1]}, \\ x^{[k]} &= \mathrm{ReLU}(\hat{z}^{[k]}), \\ z &= x^{[K]}, \end{aligned} \tag{1}$$

for all $k \in \{1, 2, \dots, K\}$. Here, the neural network input is $x \in \mathbb{R}^{n_x}$, the output is $z \in \mathbb{R}^{n_z}$, and the k^{th} layer's preactivation is $\hat{z}^{[k]} \in \mathbb{R}^{n_k}$. The parameters $W^{[k]} \in \mathbb{R}^{n_{k+1} \times n_k}$ and $b^{[k]} \in \mathbb{R}^{n_{k+1}}$ are the weight matrix and bias vector applied to the k^{th} layer's activation $x^{[k]} \in \mathbb{R}^{n_k}$, respectively. Without loss of generality,[1] we assume that the bias terms are accounted for in the activations $x^{[k]}$, thereby setting $b^{[k]} = 0$ for all layers k. We let the function $f \colon \mathbb{R}^{n_x} \to \mathbb{R}^{n_z}$ denote the map $x \mapsto z$ defined by (1). In the case that f is a classification network, the output dimension n_z equals the number of classes. The problem at hand is to certify the robustness in the neural network output z when the input x is subject to uncertainty.

To model the input uncertainty, we assume that the network inputs are unknown but contained in a compact set $\mathcal{X} \subseteq \mathbb{R}^{n_x}$, called the *input uncertainty set*. We assume that the set $\mathcal{X}$ is a convex polytope, so that the condition $x \in \mathcal{X}$ can be written as a finite number of affine inequalities and equalities. In the adversarial robustness literature, the input uncertainty set is commonly modeled as $\mathcal{X} = \{x \in \mathbb{R}^{n_x} : \|x - \bar{x}\|_\infty \leq \epsilon\}$, where $\bar{x} \in \mathbb{R}^{n_x}$ is a nominal input to the network and $\epsilon > 0$ is an uncertainty radius (Wong and Kolter, 2018; Raghunathan et al., 2018). We remark that our generalized polytopic model for $\mathcal{X}$ includes this common case. The primary theme of this paper revolves around partitioning the input uncertainty set in order to strengthen convex robustness certification methods. Let us briefly recall the definition of a partition.

[1] Note that $Wx + b = \begin{bmatrix} W & b \end{bmatrix} \begin{bmatrix} x \\ 1 \end{bmatrix} =: \tilde{W}\tilde{x}$, so the bias term b can be eliminated by appending a fixed value of 1 at the end of the activation x. This parameterization can be used throughout this paper by using matching lower and upper activation bounds of 1 in the last coordinate of each layer.

Definition 1 (Partition) *The collection $\{\mathcal{X}^{(j)} \subseteq \mathcal{X} : j \in \{1, 2, \dots, p\}\}$ is said to be a* partition *of the input uncertainty set $\mathcal{X}$ if $\mathcal{X} = \cup_{j=1}^{p} \mathcal{X}^{(j)}$ and $\mathcal{X}^{(j)} \cap \mathcal{X}^{(k)} = \emptyset$ for all $j \neq k$. The set $\mathcal{X}^{(j)}$ is called the j^{th}* input part*.*

Now, in order to describe the robustness of the network (1), we need a notion of permissible outputs. For this, we consider a *safe set*, denoted $\mathcal{S} \subseteq \mathbb{R}^{n_z}$. If an output $z \in \mathbb{R}^{n_z}$ is an element of $\mathcal{S}$, we say that z is *safe*. It is common in the robustness certification literature to consider (possibly unbounded) polyhedral safe sets. We take this same perspective, and assume that $\mathcal{S}$ is defined by

$$\mathcal{S} = \{z \in \mathbb{R}^{n_z} : Cz \leq d\},$$

where $C \in \mathbb{R}^{n_{\mathcal{S}} \times n_z}$ and $d \in \mathbb{R}^{n_{\mathcal{S}}}$ are given. Note that, again, our model naturally includes the classification setting. In particular, suppose that $i^* \in \arg\max_{i \in \{1,2,\dots,n_z\}} f_i(\bar{x})$ is the true class of the nominal input $\bar{x}$. Then, define the i^{th} row of $C \in \mathbb{R}^{n_z \times n_z}$ to be

$$c_i^\top = e_i^\top - e_{i^*}^\top$$

and d to be the zero vector. Then it is immediately clear that an input x (which can be thought of as a perturbed version of $\bar{x}$) is safe if and only if $f_i(x) \leq f_{i^*}(x)$ for all i, i.e., the network classifies x into class i^*. From this classification perspective, the safe set represents the set of outputs without any misclassification (with respect to the nominal input $\bar{x}$). From here on, we consider f, $\mathcal{X}$, and $\mathcal{S}$ in their general forms—we do not restrict ourselves to the classification setting.

2.2 The Robustness Certification Problem

The fundamental goal of the robustness certification problem is to prove that all inputs in the input uncertainty set map to safe outputs, i.e., that $f(x) \in \mathcal{S}$ for all $x \in \mathcal{X}$. If this certificate holds, the network is said to be *certifiably robust*, which of course is a property that holds with respect to a particular input uncertainty set. The robustness condition can also be written as $f(\mathcal{X}) \subseteq \mathcal{X}$, or equivalently

$$\sup_{x \in \mathcal{X}} \left(c_i^\top f(x) - d_i \right) \leq 0 \text{ for all } i \in \{1, 2, \dots, n_{\mathcal{S}}\},$$

where $c_i^\top$ is the i^{th} row of C. Under this formulation, the certification procedure amounts to solving $n_{\mathcal{S}}$ optimization problems. The methods we develop in this paper can be applied to each of these optimizations individually, and therefore in the sequel we focus on a single optimization problem by assuming that $n_{\mathcal{S}} = 1$, namely $\sup_{x \in \mathcal{X}} c^\top f(x)$. We also assume without loss of generality that $d = 0$. If d were nonzero, one may absorb d into the cost vector c and modify the network model by appending a fixed value of 1 at the end of the output vector $f(x)$. Under these formulations, we write the robustness certification problem as

$$f^*(\mathcal{X}) = \sup\{c^\top z : z = f(x),\ x \in \mathcal{X}\}, \tag{2}$$

and recall that we seek to certify that $f^*(\mathcal{X}) \leq 0$.

Since the function f is in general nonconvex, the nonlinear equality constraint $z = f(x)$ makes the optimization (2) a nonconvex problem and the set $f(\mathcal{X})$ a nonconvex set. Furthermore, since the intermediate activations and preactivations of the network generally have a large dimension in practice, the problem (2) is typically a high-dimensional problem. Therefore, computing an exact robustness certificate, as formulated in (2), is computationally intractable. Instead of directly maximizing $c^\top z$ over the nonconvex output set $f(\mathcal{X})$, one can overcome these hurdles by optimizing over a convex outer approximation $\hat{f}(\mathcal{X}) \supseteq f(\mathcal{X})$. Indeed, this new problem is a convex relaxation of the original problem, so it is generally easier and more efficient to solve. If the convex outer approximation is shown to be safe, i.e., $\hat{f}(\mathcal{X}) \subseteq \mathcal{S}$, then the true nonconvex set $f(\mathcal{X})$ is also known to be safe, implying that the robustness of the network is certified. Figure 1 illustrates this idea.

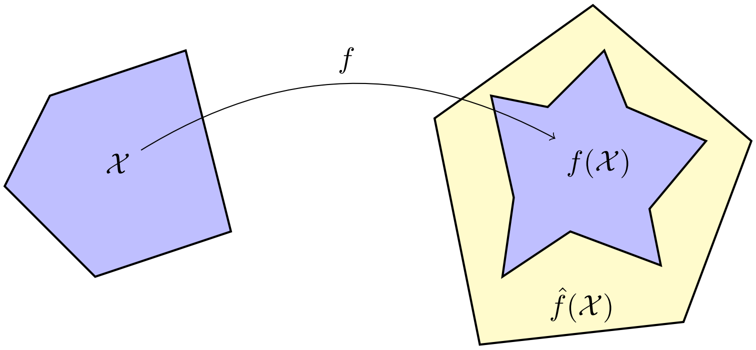

Figure 1: The set $\hat{f}(\mathcal{X})$ is a convex outer approximation of the nonconvex set $f(\mathcal{X})$. If the outer approximation is safe, i.e., $\hat{f}(\mathcal{X}) \subseteq \mathcal{S}$, then so is $f(\mathcal{X})$.

A fundamental drawback to the convex relaxation approach to robustness certification is as follows: if the outer approximation $\hat{f}(\mathcal{X})$ is too loose, then it may intersect with the unsafe region of the output space, meaning $\hat{f}(\mathcal{X}) \nsubseteq \mathcal{S}$, even in the case where the true output set is safe. In this scenario, the convex relaxation fails to issue a certificate of robustness, since the optimal value to the convex relaxation is positive, which incorrectly suggests the presence of unsafe network inputs within $\mathcal{X}$. This situation is illustrated in Figure 2.

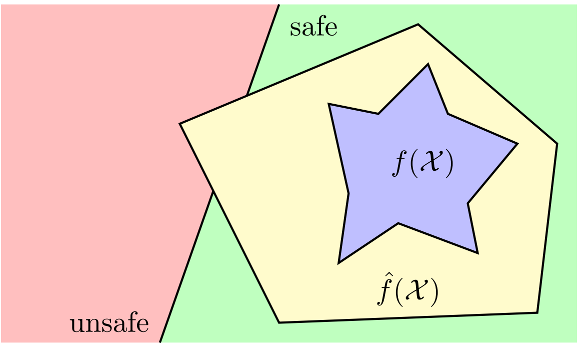

Figure 2: This scenario shows that if the convex outer approximation $\hat{f}(\mathcal{X})$ is too large, meaning the relaxation is too loose, then the convex approach fails to issue a certificate of robustness.

The fundamental goal of this paper is to develop convex relaxation methods for robustness certification such that the outer approximation tightly fits $f(\mathcal{X})$, in effect granting strong certificates while maintaining the computational and theoretical advantages of convex optimization. We focus on two popular types of convex relaxations, namely, LP (Wong and Kolter, 2018) and SDP (Raghunathan et al., 2018) relaxations. It has been shown that the SDP relaxation for ReLU networks is tighter than the LP relaxation, with the tradeoff of being computationally more demanding. Our general approach for both the LP and the SDP relaxations is based on partitioning the input uncertainty set and solving a convex relaxation on each input part separately. Throughout our theoretical development and experiments, we will show that this approach presents a valid, efficient, and effective way to tighten the relaxations of both LP and SDP certifications. We now turn to mathematically formulating these relaxations.

2.3 LP Relaxation of the Network Constraints

We now introduce the LP relaxation. First, we remark that since $\mathcal{X}$ is bounded, the preactivations at each layer are bounded as well. That is, for every $k \in \{1, 2, \dots, K\}$, there exist preactivation bounds $l^{[k]}, u^{[k]} \in \mathbb{R}^{n_k}$ such that $l^{[k]} \leq \hat{z}^{[k]} \leq u^{[k]}$, where $\hat{z}^{[k]}$ is the k^{th} layer's preactivation in (1), for all $x \in \mathcal{X}$. We assume without loss of generality that $l^{[k]} \leq 0 \leq u^{[k]}$ for all k. Although there exist various methods in the literature for efficiently approximating these preactivation bounds, we consider the bounds to be tight, i.e., $\hat{z}^{[k]} = l^{[k]}$ for some $x \in \mathcal{X}$, and similarly for the upper bound $u^{[k]}$. Now, following the approach initially introduced in Wong and Kolter (2018), we can relax the k^{th} ReLU constraint in (1) to its convex upper envelope between the preactivation bounds. This is graphically shown in Figure 3.

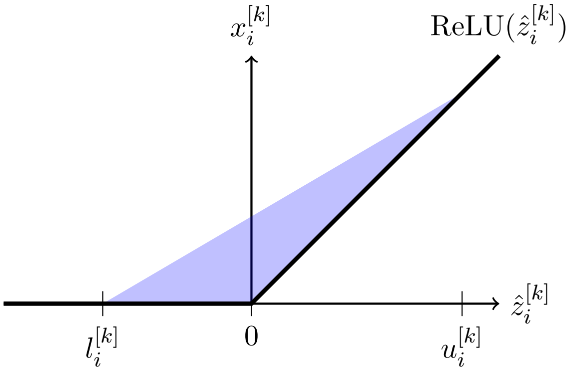

Figure 3: Relaxed ReLU constraint set $\mathcal{N}^{[k]}$ at a single neuron i in layer k of the network.

We call the convex upper envelope associated with layer k the *relaxed ReLU constraint set*, and its mathematical definition is given by three linear inequalities:

$$\begin{aligned}\mathcal{N}^{[k]} = \{(x^{[k-1]}, x^{[k]}) \in \mathbb{R}^{n_{k-1}} \times \mathbb{R}^{n_k} : x^{[k]} \leq u^{[k]} \odot (\hat{z}^{[k]} - l^{[k]}) \oslash (u^{[k]} - l^{[k]}), \\ x^{[k]} \geq 0,\ x^{[k]} \geq \hat{z}^{[k]},\ \hat{z}^{[k]} = W^{[k-1]} x^{[k-1]}\}. \quad (3)\end{aligned}$$

Next, we define the *relaxed network constraint set* to be

$$\mathcal{N} = \{(x,z) \in \mathbb{R}^{n_x} \times \mathbb{R}^{n_z} : (x, x^{[1]}) \in \mathcal{N}^{[1]},\ (x^{[1]}, x^{[2]}) \in \mathcal{N}^{[2]}, \dots, (x^{[K-1]}, z) \in \mathcal{N}^{[K]}\}. \quad (4)$$

In essence, $\mathcal{N}$ is the set of all input-output pairs of the network that satisfy the relaxed ReLU constraints at every layer. Note that, since the bounds $l^{[k]}$ and $u^{[k]}$ are determined by the input uncertainty set $\mathcal{X}$, the set $\mathcal{N}^{[k]}$ is also determined by $\mathcal{X}$ for all layers k.

Remark 2 *For networks with one hidden layer (i.e., $K = 1$), the single relaxed ReLU constraint set coincides with the relaxed network constraint set: $\mathcal{N}^{[1]} = \mathcal{N}$. Therefore, for $K = 1$ we drop the k-notation and simply write z, $\hat{z}$, x, W, l, u, and $\mathcal{N}$.*

Finally, we introduce the LP relaxation of (2):

$$\hat{f}^*(\mathcal{X}) = \sup\{c^\top z : (x,z) \in \mathcal{N},\ x \in \mathcal{X}\}. \quad (5)$$

Notice that, if $x \in \mathcal{X}$ and $z = f(x)$, then $(x,z) \in \mathcal{N}$ by the definition of $\mathcal{N}$. Furthermore, since $\mathcal{X}$ is a bounded convex polytope and $\mathcal{N}$ is defined by a system of linear constraints, we confirm that (5) is a linear program. Therefore, (5) is indeed an LP relaxation of (2), so it holds that

$$f^*(\mathcal{X}) \leq \hat{f}^*(\mathcal{X}). \quad (6)$$

This analytically shows what Figures 1 and 2 illustrate: the condition that $\hat{f}^*(\mathcal{X}) \leq 0$ is sufficient to conclude that the network is certifiably robust, but if $\hat{f}^*(\mathcal{X}) > 0$, the relaxation fails to certify whether or not the network is robust, since it may still hold that $f^*(\mathcal{X}) \leq 0$.

2.4 SDP Relaxation of the Network Constraints

An alternative convex relaxation of the robustness certification problem can be formulated as an SDP. This method was first introduced in Raghunathan et al. (2018). Here, we will introduce the SDP relaxation for a network with a single hidden layer for notational convenience. The extension to multi-layer networks is straightforward. In this formulation, the optimization variable $(x,z) \in \mathbb{R}^{n_x+n_z}$ is lifted to a symmetric matrix

$$P = \begin{bmatrix} 1 \\ x \\ z \end{bmatrix} \begin{bmatrix} 1 & x^\top & z^\top \end{bmatrix} \in \mathbb{S}^{n_x+n_z+1}.$$

We use the following block-indexing for P:

$$P = \begin{bmatrix} P_1 & P_x^\top & P_z^\top \\ P_x & P_{xx} & P_{xz} \\ P_z & P_{zx} & P_{zz} \end{bmatrix},$$

where $P_1 \in \mathbb{R}$, $P_x \in \mathbb{R}^{n_x}$, $P_z \in \mathbb{R}^{n_z}$, $P_{xx} \in \mathbb{S}^{n_x}$, $P_{zz} \in \mathbb{S}^{n_z}$, $P_{xz} \in \mathbb{R}^{n_x \times n_z}$, and $P_{zx} = P_{xz}^\top$. This lifting procedure results in the optimization problem

$$\begin{aligned} \underset{P \in \mathbb{S}^{n_x+n_z+1}}{\text{maximize}} \quad & c^\top P_z \\ \text{subject to} \quad & P_z \geq 0, \\ & P_z \geq W P_x, \\ & \operatorname{diag}(P_{xx}) \leq (l+u) \odot P_x - l \odot u, \\ & \operatorname{diag}(P_{zz}) = \operatorname{diag}(W P_{xz}), \\ & P_1 = 1, \\ & P \succeq 0, \\ & \operatorname{rank}(P) = 1. \end{aligned}$$

Here, there are no preactivation bounds, unlike the LP relaxation. The vectors $l, u \in \mathbb{R}^{n_x}$ are lower and upper bounds on the input, which are determined by the input uncertainty set. For example, if $\mathcal{X} = \{x \in \mathbb{R}^{n_x} : \|x - \bar{x}\|_\infty \leq \epsilon\}$, then $l = \bar{x} - \epsilon \mathbf{1}_{n_x}$ and $u = \bar{x} + \epsilon \mathbf{1}_{n_x}$.

We remark that the above problem is equivalent to the original robustness certification problem; no relaxation has been made yet. The only nonconvex constraint in this formulation is the rank-1 constraint on P. Dropping this rank constraint, we obtain the SDP relaxed network constraint set:

$$\begin{aligned} \mathcal{N}_{\text{SDP}} = \{P \in \mathbb{S}^{n_x+n_z+1} : P_z \geq 0,\ P_z \geq W P_x,\ \operatorname{diag}(P_{zz}) = \operatorname{diag}(W P_{xz}), \\ \operatorname{diag}(P_{xx}) \leq (l+u) \odot P_x - l \odot u,\ P_1 = 1,\ P \succeq 0\}. \end{aligned} \tag{7}$$

Using this relaxed network constraint set as the feasible set for the optimization, we arrive at the SDP relaxation

$$\hat{f}^*_{\text{SDP}}(\mathcal{X}) = \sup\{c^\top P_z : P \in \mathcal{N}_{\text{SDP}},\ P_x \in \mathcal{X}\}. \tag{8}$$

It is clear that by dropping the rank constraint, we have enlarged the feasible set, so again we obtain a viable relaxation of the original problem (2): $f^*(\mathcal{X}) \leq \hat{f}^*_{\text{SDP}}(\mathcal{X})$. In the case the solution P^* to (8) is rank-1, we can factorize it as

$$P^* = \begin{bmatrix} 1 \\ x^* \\ z^* \end{bmatrix} \begin{bmatrix} 1 & x^{*\top} & z^{*\top} \end{bmatrix}$$

and conclude that (x^*, z^*) solves the original nonconvex problem. However, it is generally the case that the SDP solution will be of higher rank, leading to relaxation error and the possibility of a void robustness certificate, similar to the LP relaxation. We now turn to building upon the LP and SDP convex relaxations via input partitioning in order to tighten their relaxations.

3. Partitioned LP Relaxation

3.1 Properties of Partitioned Relaxation

In this section, we investigate the properties and effectiveness of partitioning the input uncertainty set when solving the LP relaxation for robustness certification. We start by validating the approach, namely, by showing that solving the LP relaxation separately on each input part maintains a theoretically guaranteed upper bound on the optimal value of the unrelaxed problem (2). Afterwards, the approach is proven to yield a tighter upper bound than solving the LP relaxation without partitioning.

3.1.1 Partitioning Gives Valid Relaxation

Proposition 3 (Partitioned relaxation bound) *Let $\{\mathcal{X}^{(j)} \subseteq \mathcal{X} : j \in \{1,2,\dots,p\}\}$ be a partition of $\mathcal{X}$. Then, it holds that*

$$f^*(\mathcal{X}) \leq \max_{j \in \{1,2,\dots,p\}} \hat{f}^*(\mathcal{X}^{(j)}). \tag{9}$$

Proof See Appendix A. ■

Despite the fact that Proposition 3 asserts an intuitively expected bound, we remark the importance for its formal statement and proof. In particular, the inequality (9) serves as the fundamental reason for why the partitioned LP relaxation can be used to certify that all inputs in $\mathcal{X}$ map to safe outputs in the safe set $\mathcal{S}$. Knowing that the partitioning approach is valid for robustness certification, we move on to studying the effectiveness of partitioning.

3.1.2 Tightening of the Relaxation

We now show that the bound (6) can always be tightened by partitioning the input uncertainty set. The result is given for networks with one hidden layer for simplicity. However, the conclusion naturally generalizes to multi-layer ReLU networks.

Proposition 4 (Improving the LP relaxation bound) *Consider a feedforward ReLU neural network with one hidden layer. Let $\{\mathcal{X}^{(j)} \subseteq \mathcal{X} : j \in \{1,2,\dots,p\}\}$ be a partition of $\mathcal{X}$. For the j^{th} input part $\mathcal{X}^{(j)}$, denote the corresponding preactivation bounds by $l^{(j)}$ and $u^{(j)}$, where $l \leq l^{(j)} \leq Wx \leq u^{(j)} \leq u$ for all $x \in \mathcal{X}^{(j)}$. Then, it holds that*

$$\max_{j \in \{1,2,\dots,p\}} \hat{f}^*(\mathcal{X}^{(j)}) \leq \hat{f}^*(\mathcal{X}). \tag{10}$$

Proof See Appendix B. ■

Combining Propositions 3 and 4 shows that $f^*(\mathcal{X}) \leq \max_{j \in \{1,2,\dots,p\}} \hat{f}^*(\mathcal{X}^{(j)}) \leq \hat{f}^*(\mathcal{X})$, i.e., that the partitioned LP relaxation is theoretically guaranteed to perform at least as good as the unpartitioned LP when solving the robustness certification problem. The improvement in the partitioned LP relaxation is captured by the difference

$$\hat{f}^*(\mathcal{X}) - \max_{j \in \{1,2,\dots,p\}} \hat{f}^*(\mathcal{X}^{(j)}),$$

which is always nonnegative. We remark that it is possible for the improvement to be null in the sense that $\max_{j \in \{1,2,\dots,p\}} \hat{f}^*(\mathcal{X}^{(j)}) = \hat{f}^*(\mathcal{X})$. This may occur when the partition used is poorly chosen. An example of such a poor choice may be if one were to partition along a direction in the input space that, informally speaking, corresponds to directions near-orthogonal to the cost vector c in the output space. In this case, one would expect all improvements to be nullified, and for the partitioned relaxation to give the same optimal value as the unpartitioned relaxation. Consequently, the following important question arises: *what constitutes a good partition so that the improvement $\hat{f}^*(\mathcal{X}) - \max_{j\in\{1,2,\dots,p\}} \hat{f}^*(\mathcal{X}^{(j)})$ is strictly greater than zero and maximal?* We address this question in Sections 3.2 and 3.3.

3.2 Motivating Partition

In this section, we begin to answer our earlier inquiry, namely, how to choose a partition in order to maximize the improvement $\hat{f}^*(\mathcal{X}) - \max_{j\in\{1,2,\dots,p\}} \hat{f}^*(\mathcal{X}^{(j)})$ in the partitioned LP relaxation. Recall that this is equivalent to minimizing the relaxation error relative to the original unrelaxed problem, since $f^*(\mathcal{X}) \leq \max_{j\in\{1,2,\dots,p\}} \hat{f}^*(\mathcal{X}^{(j)}) \leq \hat{f}^*(\mathcal{X})$. To this end, we construct a partition with finitely many parts, based on the parameters of the network, which is shown to exactly recover the optimal value of the original unrelaxed problem (2). For simplicity, we present the result for a single hidden layer, but the basic idea of partitioning at the "kinks" of the ReLUs in order to collapse the ReLU upper envelope onto the ReLU curve and eliminate relaxation error can be generalized to multi-layer settings. At this point, let us remark that in Proposition 5 below, we use a slight difference in notation for the partition. Namely, we use the set of all n_z-vectors with binary elements, $\mathcal{J} \coloneqq \{0,1\}^{n_z} = \{0,1\} \times \{0,1\} \times \cdots \times \{0,1\}$, to index the parts of the partition. Under this setting, the partition is composed of $p = 2^{n_z}$ parts, so that $\mathcal{X}^{(j)}$ is the part of the partition corresponding to the binary vector j, which is an element of the index set $\mathcal{J}$. This temporary change in notation is chosen to simplify the proof of Proposition 5.

Proposition 5 (Motivating partition) *Consider a feedforward ReLU neural network with one hidden layer and denote the i^{th} row of W by $w_i^\top \in \mathbb{R}^{1\times n_x}$ for all $i \in \{1,2,\dots,n_z\}$. Define $\mathcal{J} = \{0,1\}^{n_z}$ and take the partition of $\mathcal{X}$ to be indexed by $\mathcal{J}$, meaning that $\{\mathcal{X}^{(j)} \subseteq \mathcal{X} : j \in \mathcal{J}\}$, where for a given $j \in \mathcal{J}$ we define*

$$\mathcal{X}^{(j)} = \{x \in \mathcal{X} : w_i^\top x \geq 0 \text{ for all } i \text{ such that } j_i = 1,\ w_i^\top x < 0 \text{ for all } i \text{ such that } j_i = 0\}. \tag{11}$$

Then, the partitioned relaxation is exact, i.e.,

$$f^*(\mathcal{X}) = \max_{j\in\mathcal{J}} \hat{f}^*(\mathcal{X}^{(j)}). \tag{12}$$

Proof We first show that $\{\mathcal{X}^{(j)} \subseteq \mathcal{X} : j \in \mathcal{J}\}$ is a valid partition. Since $\mathcal{X}^{(j)} \subseteq \mathcal{X}$ for all $j \in \mathcal{J}$, the relation $\cup_{j\in\mathcal{J}}\mathcal{X}^{(j)} \subseteq \mathcal{X}$ is satisfied. Now, suppose that $x \in \mathcal{X}$. Then, for all $i \in \{1,2,\dots,n_z\}$, either $w_i^\top x \geq 0$ or $w_i^\top x < 0$ holds. Define $j \in \{0,1\}^{n_z}$ as follows:

$$j_i = \begin{cases} 1 & \text{if } w_i^\top x \geq 0, \\ 0 & \text{if } w_i^\top x < 0, \end{cases}$$

for all $i \in \{1,2,\dots,n_z\}$. Then, by the definition of $\mathcal{X}^{(j)}$ in (11), it holds that $x \in \mathcal{X}^{(j)}$. Therefore, the relation $x \in \mathcal{X}$ implies that $x \in \mathcal{X}^{(j)}$ for some $j \in \{0,1\}^{n_z} = \mathcal{J}$. Hence, $\mathcal{X} \subseteq \cup_{j\in\mathcal{J}}\mathcal{X}^{(j)}$, and therefore $\cup_{j\in\mathcal{J}}\mathcal{X}^{(j)} = \mathcal{X}$.

We now show that $\mathcal{X}^{(j)} \cap \mathcal{X}^{(k)} = \emptyset$ for all $j \neq k$. Let $j,k \in \mathcal{J}$ with the property that $j \neq k$. Then, there exists $i \in \{1,2,\dots,n_z\}$ such that $j_i \neq k_i$. Let $x \in \mathcal{X}^{(j)}$. In the case that $w_i^\top x \geq 0$, it holds that $j_i = 1$ and therefore $k_i = 0$. Hence, for all $y \in \mathcal{X}^{(k)}$, it holds that $w_i^\top y < 0$, and therefore $x \notin \mathcal{X}^{(k)}$. An analogous reasoning shows that $x \notin \mathcal{X}^{(k)}$ when $w_i^\top x < 0$. Therefore, one concludes that $x \in \mathcal{X}^{(j)}$ and $j \neq k$ implies that $x \notin \mathcal{X}^{(k)}$, i.e., that $\mathcal{X}^{(j)} \cap \mathcal{X}^{(k)} = \emptyset$. Hence, $\{\mathcal{X}^{(j)} \subseteq \mathcal{X} : j \in \mathcal{J}\}$ is a valid partition.

We now prove (12). Let $j \in \mathcal{J}$. Since $w_i^\top x \geq 0$ for all i such that $j_i = 1$, the preactivation lower bound becomes $l_i^{(j)} = 0$ for all such i. On the other hand, since $w_i^\top x < 0$ for all i such that $j_i = 0$, the preactivation upper bound becomes $u_i^{(j)} = 0$ for all such i. Therefore, the relaxed network constraint set (4) for the j^{th} input part reduces to

$$\begin{aligned}\mathcal{N}^{(j)} = \{(x,z) \in \mathbb{R}^{n_x} \times \mathbb{R}^{n_z} : z_i = 0 \text{ for all } i \text{ such that } j_i = 0,& \\ z_i = w_i^\top x = (Wx)_i \text{ for all } i \text{ such that } j_i = 1\}.&\end{aligned}$$

That is, the relaxed ReLU constraint envelope collapses to the exact ReLU constraint through the prior knowledge of each preactivation coordinate's sign. Therefore, we find that for all $x \in \mathcal{X}^{(j)}$ it holds that $(x,z) \in \mathcal{N}^{(j)}$ if and only if $z = \text{ReLU}(Wx)$. Hence, the LP over the j^{th} input part yields that

$$\begin{aligned}\hat{f}^*(\mathcal{X}^{(j)}) &= \sup\{c^\top z : (x,z) \in \mathcal{N}^{(j)},\ x \in \mathcal{X}^{(j)}\} = \sup\{c^\top z : z = \text{ReLU}(Wx),\ x \in \mathcal{X}^{(j)}\} \\ &\leq \sup\{c^\top z : z = \text{ReLU}(Wx),\ x \in \mathcal{X}\} = f^*(\mathcal{X}).\end{aligned}$$

Since j was chosen arbitrarily, it holds that

$$\max_{j \in \mathcal{J}} \hat{f}^*(\mathcal{X}^{(j)}) \leq f^*(\mathcal{X}).$$

Since $f^*(\mathcal{X}) \leq \max_{j \in \mathcal{J}} \hat{f}^*(\mathcal{X}^{(j)})$ by the relaxation bound (9), the equality (12) holds, as desired. ■

Although the partition proposed in Proposition 5 completely eliminates relaxation error of the LP, using it in practice may be computationally intractable, as it requires solving 2^{n_z} separate linear programs. Despite this limitation, the result provides two major theoretical implications. First, our input partitioning approach is fundamentally shown to be a simple, yet very powerful method, as the robustness certification problem can be solved exactly via a finite number of linear program subproblems. Second, the partition proposed in Proposition 5 shows us the structure of an optimal partition, namely that the parts of the partition are defined by the half-spaces generated by the rows of W (see Figure 4). This result paves the way to develop a tractable two-part partition that incorporates the optimal reduction in relaxation error endowed by the structure of this motivating partition. In the next section, we explore this idea further, and answer the following question: *if we only partition along a single row of the weight matrix, which one is the best to choose?*

3.3 Partitioning Scheme

In this section, we propose an explicit, computationally tractable, and effective LP partitioning scheme. The partitioning scheme is developed based on analyses for a single hidden layer. However, the resulting partitioning scheme is still applicable to multi-layer networks, and indeed will be shown to remain effective on two-layer networks in the experiments of

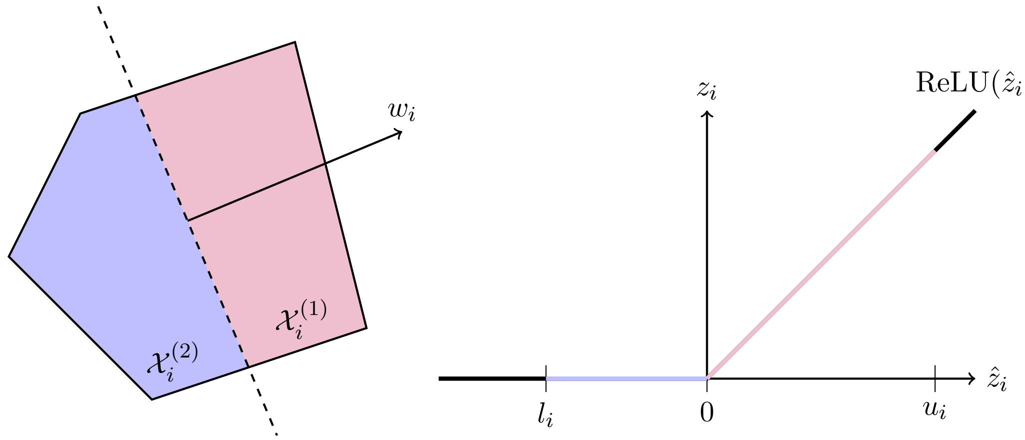

Figure 4: Partitioning based on row $w_i^\top$ of the weight matrix. This partition results in an exact ReLU constraint in coordinate i over the two resulting input parts $\mathcal{X}_i^{(1)} = \{x \in \mathcal{X} : w_i^\top x \geq 0\}$ and $\mathcal{X}_i^{(2)} = \mathcal{X} \setminus \mathcal{X}_i^{(1)}$.

Section 5. The development of the partition boils down to two main ideas. First, we restrict our attention to two-part partitions defined by rows of the weight matrix W, specifically, $\mathcal{X}_i^{(1)} = \{x \in \mathcal{X} : w_i^\top x \geq 0\}$ and $\mathcal{X}_i^{(2)} = \mathcal{X} \setminus \mathcal{X}_i^{(1)}$, as motivated in the previous section. Second, we seek which index $i \in \{1, 2, \dots, n_z\}$ gives the best partition, in the sense that the relaxation error of the resulting partitioned LP is minimized. As will be shown in Section 3.3.3, this second aspect is NP-hard to discern in general. Therefore, to find the optimal row to partition along, we instead seek to minimize the worst-case relaxation error.

3.3.1 Worst-Case Relaxation Bound

We begin by bounding the relaxation error below.

Theorem 6 (Worst-case relaxation bound) *Consider a feedforward ReLU neural network with one hidden layer, with the input uncertainty set $\mathcal{X}$ and preactivation bounds $l, u \in \mathbb{R}^{n_z}$. Consider also the relaxation error $\Delta f^*(\mathcal{X}) \coloneqq \hat{f}^*(\mathcal{X}) - f^*(\mathcal{X})$. Let $(\tilde{x}^*, \tilde{z}^*)$ and (x^*, z^*) be optimal solutions for the relaxation $\hat{f}^*(\mathcal{X})$ and the unrelaxed problem $f^*(\mathcal{X})$, respectively. Given an arbitrary norm $\|\cdot\|$ on $\mathbb{R}^{n_x}$, there exists $\epsilon \in \mathbb{R}$ such that $\|\tilde{x}^* - x^*\| \leq \epsilon$ and*

$$\Delta f^*(\mathcal{X}) \leq \sum_{i=1}^{n_z} \left(\text{ReLU}(c_i) \frac{u_i}{u_i - l_i} (\min\{\epsilon \|w_i\|_*, u_i\} - l_i) + \text{ReLU}(-c_i) \min\{\epsilon \|w_i\|_*, u_i\} \right), \tag{13}$$

where $\|\cdot\|_*$ is the dual norm of $\|\cdot\|$.

Proof First, since $\mathcal{X}$ is bounded, there exists $\epsilon \geq 0$ such that $\|\tilde{x}^* - x^*\| \leq \epsilon$. The definitions of $(\tilde{x}^*, \tilde{z}^*)$ and (x^*, z^*) give that

$$\Delta f^*(\mathcal{X}) = \sum_{i=1}^{n_z} c_i (\tilde{z}_i^* - z_i^*) \leq \sum_{i=1}^{n_z} \Delta f_i^*, \tag{14}$$

where

$$\Delta f_i^* = \sup\Bigg\{c_i(\tilde{z}_i - z_i) : z_i = \mathrm{ReLU}(w_i^\top x),\ \tilde{z}_i \geq 0,\ \tilde{z}_i \geq w_i^\top \tilde{x},\ \tilde{z}_i \leq \frac{u_i}{u_i - l_i}(w_i^\top \tilde{x} - l_i), \\ \|\tilde{x}^* - x^*\| \leq \epsilon,\ x, \tilde{x} \in \mathcal{X}\Bigg\}$$

for all $i \in \{1, 2, \dots, n_z\}$. Note that

$$\Delta f_i^* = \sup\Bigg\{c_i(\tilde{z}_i - z_i) : z_i = \mathrm{ReLU}(\hat{z}_i),\ \tilde{z}_i \geq 0,\ \tilde{z}_i \geq \hat{\tilde{z}}_i,\ \tilde{z}_i \leq \frac{u_i}{u_i - l_i}(\hat{\tilde{z}}_i - l_i), \\ \|\tilde{x}^* - x^*\| \leq \epsilon,\ \hat{z} = Wx,\ \hat{\tilde{z}} = W\tilde{x},\ x, \tilde{x} \in \mathcal{X}\Bigg\}.$$

If $x, \tilde{x} \in \mathcal{X}$ and $\hat{z}, \hat{\tilde{z}}$ satisfy $\hat{z} = Wx$, $\hat{\tilde{z}} = W\tilde{x}$, and $\|\tilde{x} - x\| \leq \epsilon$, then they satisfy $l \leq \hat{z}, \hat{\tilde{z}} \leq u$ and $|\hat{\tilde{z}}_i - \hat{z}_i| = |w_i^\top(\tilde{x} - x)| \leq \|w_i\|_* \|\tilde{x} - x\| \leq \epsilon \|w_i\|_*$ for all $i \in \{1, 2, \dots, n_z\}$ by the Cauchy-Schwarz inequality for dual norms. Therefore,

$$\begin{aligned}\Delta f_i^* \leq\ & \sup\Bigg\{c_i(\tilde{z}_i - z_i) : z_i = \mathrm{ReLU}(\hat{z}_i),\ \tilde{z}_i \geq 0,\ \tilde{z}_i \geq \hat{\tilde{z}}_i,\ \tilde{z}_i \leq \frac{u_i}{u_i - l_i}(\hat{\tilde{z}}_i - l_i), \\ & l \leq \hat{z}, \hat{\tilde{z}} \leq u,\ |\hat{\tilde{z}}_k - \hat{z}_k| \leq \epsilon\|w_k\|_* \text{ for all } k \in \{1, 2, \dots, n_z\},\ \hat{z}, \hat{\tilde{z}} \in \mathbb{R}^{n_z}\Bigg\} \\ =\ & \sup\Bigg\{c_i(\tilde{z}_i - z_i) : z_i = \mathrm{ReLU}(\hat{z}_i),\ \tilde{z}_i \geq 0,\ \tilde{z}_i \geq \hat{\tilde{z}}_i,\ \tilde{z}_i \leq \frac{u_i}{u_i - l_i}(\hat{\tilde{z}}_i - l_i), \\ & l_i \leq \hat{z}_i, \hat{\tilde{z}}_i \leq u_i,\ |\hat{\tilde{z}}_i - \hat{z}_i| \leq \epsilon\|w_i\|_*,\ \hat{z}_i, \hat{\tilde{z}}_i \in \mathbb{R}\Bigg\}.\end{aligned}$$

For $c_i \geq 0$, the above inequality yields that

$$\Delta f_i^* \leq c_i \sup\Bigg\{\tilde{z}_i - z_i : z_i = \mathrm{ReLU}(\hat{z}_i),\ \tilde{z}_i \geq 0,\ \tilde{z}_i \geq \hat{\tilde{z}}_i,\ \tilde{z}_i \leq \frac{u_i}{u_i - l_i}(\hat{\tilde{z}}_i - l_i), \\ l_i \leq \hat{z}_i, \hat{\tilde{z}}_i \leq u_i,\ |\hat{\tilde{z}}_i - \hat{z}_i| \leq \epsilon\|w_i\|_*,\ \hat{z}_i, \hat{\tilde{z}}_i \in \mathbb{R}\Bigg\}.$$

The optimal solution to the above supremum is readily found by comparing the line $\tilde{z}_i = \frac{u_i}{u_i - l_i}(\hat{\tilde{z}}_i - l_i)$ to the function $z_i = \mathrm{ReLU}(\hat{z}_i)$ over $\hat{\tilde{z}}_i, \hat{z}_i \in [l_i, u_i]$. In particular, the maximum distance between $\tilde{z}_i$ and z_i on the above feasible set occurs when $z_i = \hat{z}_i = 0$, $\hat{\tilde{z}}_i = \epsilon\|w_i\|_*$, and $\tilde{z}_i = \frac{u_i}{u_i - l_i}(\epsilon\|w_i\|_* - l_i)$. Therefore, we find that

$$\Delta f_i^* \leq c_i \frac{u_i}{u_i - l_i}(\epsilon\|w_i\|_* - l_i), \tag{15}$$

for all $i \in \{1, 2, \dots, n_z\}$ such that $c_i \geq 0$. We also note the trivial bound that $\tilde{z}_i - z_i \leq u_i$ on the feasible set of the above supremum, so that

$$\Delta f_i^* \leq c_i u_i = c_i \frac{u_i}{u_i - l_i}(u_i - l_i). \tag{16}$$

The inequalities (15) and (16) together imply that

$$\Delta f_i^* \leq c_i \frac{u_i}{u_i - l_i}(\min\{\epsilon\|w_i\|_*, u_i\} - l_i) \tag{17}$$

for all $i \in \{1, 2, \dots, n_z\}$ such that $c_i \geq 0$.

On the other hand, for all $i \in \{1, 2, \dots, n_z\}$ such that $c_i < 0$, we have that

$$\Delta f_i^* \leq c_i \inf \Bigg\{ \tilde{z}_i - z_i : z_i = \mathrm{ReLU}(\hat{z}_i),\ \tilde{z}_i \geq 0,\ \tilde{z}_i \geq \hat{\tilde{z}}_i,\ \tilde{z}_i \leq \frac{u_i}{u_i - l_i}(\hat{\tilde{z}}_i - l_i), \\ l_i \leq \hat{z}_i, \hat{\tilde{z}}_i \leq u_i,\ |\hat{\tilde{z}}_i - \hat{z}_i| \leq \epsilon\|w_i\|_*,\ \hat{z}_i, \hat{\tilde{z}}_i \in \mathbb{R} \Bigg\}.$$

The optimal solution to the above infimum is readily found by comparing the line $\tilde{z}_i = 0$ to the function $z_i = \mathrm{ReLU}(\hat{z}_i)$ over $\hat{\tilde{z}}_i, \hat{z}_i \in [l_i, u_i]$. In particular, the minimum value of $\tilde{z}_i - z_i$ on the above feasible set occurs when $\tilde{z}_i = \hat{\tilde{z}}_i = 0$ and $z_i = \hat{z}_i = \epsilon\|w_i\|_*$. Therefore, we find that

$$\Delta f_i^* \leq -c_i \epsilon \|w_i\|_*, \tag{18}$$

for all $i \in \{1, 2, \dots, n_z\}$ such that $c_i < 0$. We also note the trivial bound that $\tilde{z}_i - z_i \geq -u_i$ on the feasible set of the above infimum, so that

$$\Delta f_i^* \leq -c_i u_i. \tag{19}$$

The inequalities (18) and (19) together imply that

$$\Delta f_i^* \leq -c_i \min\{\epsilon\|w_i\|_*, u_i\} \tag{20}$$

for all $i \in \{1, 2, \dots, n_z\}$ such that $c_i < 0$. Substituting (17) and (20) into (14) gives the desired bound (13). ■

The value Δf_i^* in the proof of Theorem 6 can be interpreted as the worst-case relaxation error in coordinate i. From this perspective, Theorem 6 gives an upper bound on the worst-case relaxation error of the overall network. In the case that $\tilde{x}^* \neq x^*$ and $\epsilon\|w_i\|_* \geq u_i$ for all $\epsilon \in \mathbb{R}$ such that $\|\tilde{x}^* - x^*\| \leq \epsilon$, for all $i \in \{1, 2, \dots, n_z\}$, the bound (13) becomes

$$\Delta f^*(\mathcal{X}) \leq \sum_{i=1}^{n_z} |c_i| u_i.$$

This is the loosest the bound can ever be. On the contrary, if $\tilde{x}^* = x^*$, i.e., the relaxation and the true certification problem share an optimal input (meaning that it is the most adversarial), then Theorem 6 holds for $\epsilon = 0$. Substituting this into (13) gives

$$\Delta f^*(\mathcal{X}) \leq -\sum_{i=1}^{n_z} \mathrm{ReLU}(c_i) \frac{u_i l_i}{u_i - l_i}. \tag{21}$$

Note that in practice we expect the condition $\epsilon \approx 0$ to hold, since a worst-case input to a neural network is likely to also be a worst-case input to the relaxed network. Therefore, for the remainder of this paper, we take the worst-case LP relaxation bound to be that given by (21) to simplify the analysis.

To continue our development of a two-part partitioning scheme that is optimal with respect to the worst-case relaxation error, we use (21) to bound the relaxation error of the partitioned LP in terms of the row $w_i^\top$ that is used to define the partition. This bound is given in the following lemma.

Lemma 7 (Two-part bound) *Let $i \in \{1, 2, \dots, n_z\}$ and consider a two-part partition of $\mathcal{X}$ given by $\{\mathcal{X}_i^{(1)}, \mathcal{X}_i^{(2)}\}$, where $\mathcal{X}_i^{(1)} = \{x \in \mathcal{X} : w_i^\top x \geq 0\}$ and $\mathcal{X}_i^{(2)} = \mathcal{X} \setminus \mathcal{X}_i^{(1)}$. Consider also the partitioned relaxation error $\Delta f^*(\{\mathcal{X}_i^{(1)}, \mathcal{X}_i^{(2)}\}) \coloneqq \max_{j\in\{1,2\}} \hat{f}^*(\mathcal{X}_i^{(j)}) - f^*(\mathcal{X})$. It holds that*

$$\Delta f^*(\{\mathcal{X}_i^{(1)}, \mathcal{X}_i^{(2)}\}) \leq -\sum_{\substack{k=1\\k\neq i}}^{n_z} \mathrm{ReLU}(c_k)\frac{u_k l_k}{u_k - l_k}. \tag{22}$$

Proof Consider the relaxation solved over the first input part, $\mathcal{X}_i^{(1)}$, and denote by $l^{(1)}, u^{(1)} \in \mathbb{R}^{n_z}$ the corresponding preactivation bounds. Since $w_i^\top x \geq 0$ on this input part, the preactivation bounds for the first subproblem $\hat{f}^*(\mathcal{X}_i^{(1)})$ can be taken as

$$l^{(1)} = (l_1, l_2, \dots, l_{i-1}, 0, l_{i+1}, \dots, l_{n_z})$$

and $u^{(1)} = u$. It follows from (21) that

$$\hat{f}^*(\mathcal{X}_i^{(1)}) - f^*(\mathcal{X}_i^{(1)}) \leq -\sum_{k=1}^{n_z} \mathrm{ReLU}(c_k)\frac{u_k^{(1)} l_k^{(1)}}{u_k^{(1)} - l_k^{(1)}} = -\sum_{\substack{k=1\\k\neq i}}^{n_z} \mathrm{ReLU}(c_k)\frac{u_k l_k}{u_k - l_k}. \tag{23}$$

Similarly, over the second input part, $\mathcal{X}_i^{(2)}$, we have that $w_i^\top x < 0$, and so the preactivation bounds for the second subproblem $\hat{f}^*(\mathcal{X}_i^{(2)})$ can be taken as $l^{(2)} = l$ and

$$u^{(2)} = (u_1, u_2, \dots, u_{i-1}, 0, u_{i+1}, \dots, u_{n_z}),$$

resulting in the same bound as in (23):

$$\hat{f}^*(\mathcal{X}_i^{(2)}) - f^*(\mathcal{X}_i^{(2)}) \leq -\sum_{\substack{k=1\\k\neq i}}^{n_z} \mathrm{ReLU}(c_k)\frac{u_k l_k}{u_k - l_k}. \tag{24}$$

Putting the two bounds (23) and (24) together and using the fact that $f^*(\mathcal{X}_i^{(j)}) \leq f^*(\mathcal{X})$ for all $j \in \{1, 2\}$, we find that

$$\begin{aligned}\Delta f^*(\{\mathcal{X}_i^{(1)}, \mathcal{X}_i^{(2)}\}) = \max_{j\in\{1,2\}} \left(\hat{f}^*(\mathcal{X}_i^{(j)}) - f^*(\mathcal{X})\right) &\leq \max_{j\in\{1,2\}} \left(\hat{f}^*(\mathcal{X}_i^{(j)}) - f^*(\mathcal{X}_i^{(j)})\right) \\ &\leq -\sum_{\substack{k=1\\k\neq i}}^{n_z} \mathrm{ReLU}(c_k)\frac{u_k l_k}{u_k - l_k},\end{aligned}$$

as desired. ■

3.3.2 Proposed Partition

Lemma 7 bounds the worst-case relaxation error for each possible row-based partition. Therefore, our final step in the development of our two-part partition is to find which row minimizes the upper bound (22). This optimal two-part partition is now presented.

Theorem 8 (Optimal partition) *Consider the two-part partitions defined by the rows of W: $\{\mathcal{X}_i^{(1)}, \mathcal{X}_i^{(2)}\}$, where $\mathcal{X}_i^{(1)} = \{x \in \mathcal{X} : w_i^\top x \geq 0\}$ and $\mathcal{X}_i^{(2)} = \mathcal{X} \setminus \mathcal{X}_i^{(1)}$ for all $i \in \{1, 2, \dots, n_z\} =: \mathcal{I}$. The optimal partition that minimizes the worst-case relaxation error bound in* (22) *is given by*

$$i^* \in \arg\min_{i \in \mathcal{I}} \operatorname{ReLU}(c_i) \frac{u_i l_i}{u_i - l_i}. \tag{25}$$

Proof Minimizing the bound in (22) of Lemma 7 over the partition i gives rise to

$$\min_{i \in \mathcal{I}} \left(-\sum_{\substack{k=1 \\ k \neq i}}^{n_z} \operatorname{ReLU}(c_k) \frac{u_k l_k}{u_k - l_k} \right) = -\sum_{k=1}^{n_z} \operatorname{ReLU}(c_k) \frac{u_k l_k}{u_k - l_k} + \min_{i \in \mathcal{I}} \operatorname{ReLU}(c_i) \frac{u_i l_i}{u_i - l_i},$$

as desired. ■

Theorem 8 provides the two-part partition that optimally reduces the worst-case relaxation error that we seek. We remark its simplicity: to decide which row to partition along, it suffices to enumerate the values $\operatorname{ReLU}(c_i)\frac{u_i l_i}{u_i - l_i}$ for $i \in \{1, 2, \dots, n_z\}$, then choose the row corresponding to the minimum amongst these values. Note that this optimization over i scales linearly with the dimension n_z, and the resulting LP subproblems on each input part only require the addition of one extra linear constraint, meaning that this partitioning scheme is highly efficient.

We also note that Theorem 8 can be immediately extended to design multi-part partitions in two interesting ways. First, by ordering the values $\operatorname{ReLU}(c_i)\frac{u_i l_i}{u_i - l_i}$, we are ordering the optimality of the rows $w_i^\top$ to partition along. Therefore, by partitioning along the $n_p > 1$ rows corresponding to the smallest n_p of these values, Theorem 8 provides a strategy to design an effective 2^{n_p}-part partition, in the case one prefers to perform more than just a two-part partition. Second, Theorem 8 can be used in a recursive way, similar to branch-and-bound. In particular, by solving the two-part partitioned LP using Theorem 8, we find $\hat{f}^*(\mathcal{X}_{i^*}^{(1)})$ and $\hat{f}^*(\mathcal{X}_{i^*}^{(2)})$. If $\hat{f}^*(\mathcal{X}_{i^*}^{(1)}) > \hat{f}^*(\mathcal{X}_{i^*}^{(2)})$, then we can further partition $\mathcal{X}_{i^*}^{(1)}$ into two more parts, say $\mathcal{X}_{i^*}^{(1,1)}$ and $\mathcal{X}_{i^*}^{(1,2)}$, again according to (25) but this time using the tighter preactivation bounds on $\mathcal{X}_{i^*}^{(1)}$. Then our relaxation bound becomes $f^*(\mathcal{X}) \leq \max\{\hat{f}^*(\mathcal{X}_{i^*}^{(1,1)}), \hat{f}^*(\mathcal{X}_{i^*}^{(1,2)}), \hat{f}^*(\mathcal{X}_{i^*}^{(2)})\} \leq \max\{\hat{f}^*(\mathcal{X}_{i^*}^{(1)}), \hat{f}^*(\mathcal{X}_{i^*}^{(2)})\}$. Choosing the part amongst $\mathcal{X}_{i^*}^{(1,1)}, \mathcal{X}_{i^*}^{(1,2)}, \mathcal{X}_{i^*}^{(2)}$ with the largest LP relaxation value, we can again perform a partition and repeat the process in order to continue reducing the relaxation error.

3.3.3 Optimal Partitioning is NP-Hard

In this section, we show that finding a row-based partition that minimizes the actual LP relaxation error is an NP-hard problem. Recall that this approach is in contrast to our previous approach in the sense that our optimal partition in Theorem 8 minimizes the worst-case relaxation error. Consequently, the results of this section show that the partition given by Theorem 8 is in essence the best tractable LP partitioning scheme.

To start, recall the robustness certification problem for a K-layer ReLU neural network:

$$\begin{aligned} \text{maximize} \quad & c^\top x^{[K]} \\ \text{subject to} \quad & x^{[0]} \in \mathcal{X}, \\ & x^{[k+1]} = \mathrm{ReLU}(W^{[k]} x^{[k]}), \quad k \in \{0, 1, \dots, K-1\}, \end{aligned} \tag{26}$$

where the optimal value of (26) is denoted by $f^*(\mathcal{X})$. Moreover, recall the LP relaxation of (26):

$$\begin{aligned} \text{maximize} \quad & c^\top x^{[K]} \\ \text{subject to} \quad & x^{[0]} \in \mathcal{X}, \\ & x^{[k+1]} \geq W^{[k]} x^{[k]}, && k \in \{0, 1, \dots, K-1\}, \\ & x^{[k+1]} \geq 0, && k \in \{0, 1, \dots, K-1\}, \\ & x^{[k+1]} \leq u^{[k+1]} \odot (W^{[k]} x^{[k]} - l^{[k+1]}) \oslash (u^{[k+1]} - l^{[k+1]}), && k \in \{0, 1, \dots, K-1\}. \end{aligned} \tag{27}$$

As suggested by the motivating partition of Proposition 5, consider partitioning the input uncertainty set into 2^{n_p} parts based on n_p preactivation decision boundaries corresponding to activation functions in the first layer. In particular, for each $j \in \mathcal{J}_p \coloneqq \{j_1, j_2, \dots, j_{n_p}\} \subseteq \{1, 2, \dots, n_1\}$ we partition the input uncertainty set along the hyperplane $w_j^{[0]\top} x^{[0]} = 0$. Note that, for all $j \in \mathcal{J}_p$, the partition implies that the j^{th} coordinate of the first layer's ReLU equality constraint becomes linear and exact on each part of the partition. Therefore, we may write the partitioned LP relaxation as

$$\begin{aligned} \text{maximize} \quad & c^\top x^{[K]} \\ \text{subject to} \quad & x^{[0]} \in \mathcal{X}, \\ & x^{[k+1]} \geq W^{[k]} x^{[k]}, && k \in \{0, 1, \dots, K-1\}, \\ & x^{[k+1]} \geq 0, && k \in \{0, 1, \dots, K-1\}, \\ & x^{[k+1]} \leq u^{[k+1]} \odot (W^{[k]} x^{[k]} - l^{[k+1]}) \oslash (u^{[k+1]} - l^{[k+1]}), && k \in \{0, 1, \dots, K-1\}, \\ & x_j^{[1]} = \mathrm{ReLU}(w_j^{[0]\top} x^{[0]}), && j \in \mathcal{J}_p. \end{aligned} \tag{28}$$

We denote the optimal objective of this problem as $f^*_{\mathcal{J}_p}(\mathcal{X})$. To reiterate, the final equality constraint is linear over each part of the partition, which makes the problem (28) a partitioned linear program. For notational convenience, the restriction of the input to a particular part of the partition, as well as the outer maximization over the parts of the partition, is implicit in the expression (28).

If we now allow the indices used to define the partition, namely $\mathcal{J}_p$, to act as a variable, we can search for the optimal n_p rows of the first layer that result in the tightest partitioned LP relaxation. To this end, the problem of optimal partitioning in the first layer is formulated as

$$\begin{aligned} \underset{\mathcal{J}_p \subseteq \{1,2,\dots,n_1\}}{\text{minimize}} \quad & f^*_{\mathcal{J}_p}(\mathcal{X}) \\ \text{subject to} \quad & |\mathcal{J}_p| = n_p. \end{aligned} \tag{29}$$

In what follows, we prove the NP-hardness of the optimal partitioning problem (29), thereby supporting the use of the worst-case sense optimal partition developed in Theorem 8. To show the hardness of (29), we reduce an arbitrary instance of an NP-hard problem, the Min-$\mathcal{K}$-Union problem, to an instance of (29). The reduction will show that the Min-$\mathcal{K}$-Union problem can be solved by solving an optimal partitioning problem. Before we proceed, we first recall the definition of the Min-$\mathcal{K}$-Union problem.

Definition 9 (Min-$\mathcal{K}$-Union problem (Hochbaum, 1996)) *Consider a collection of n sets $\{\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_n\}$, where $\mathcal{S}_j$ is finite for all $j \in \{1, 2, \dots, n\}$, and a positive integer $\mathcal{K} \leq n$. Find $\mathcal{K}$ sets in the collection whose union has minimum cardinality, i.e., find a solution $\mathcal{J}^*$ of the following optimization problem:*

$$\begin{aligned} \underset{\mathcal{J} \subseteq \{1,2,\dots,n\}}{\text{minimize}} \quad & \left| \bigcup_{j \in \mathcal{J}} \mathcal{S}_j \right| \\ \text{subject to} \quad & |\mathcal{J}| = \mathcal{K}. \end{aligned} \tag{30}$$

Remark the similarities between the optimal partitioning problem and the Min-$\mathcal{K}$-Union problem. In particular, if we think of the convex upper envelopes of the relaxed ReLU constraints as a collection of sets, then the goal of finding the optimal n_p input coordinates to partition along is intuitively equivalent to searching for the $\mathcal{K} = n_1 - n_p$ convex upper envelopes with minimum size, i.e., those with the least amount of relaxation. This perspective shows that the optimal partitioning problem is essentially a Min-$\mathcal{K}$-Union problem over the collection of relaxed ReLU constraint sets. Since the Min-$\mathcal{K}$-Union problem is NP-hard in general, it is not surprising that the optimal partitioning problem is also NP-hard. Indeed, this result is formalized in the following proposition.

Proposition 10 (NP-hardness of optimal partition) *Consider the partitioned LP relaxation* (28) *of the K-layer ReLU neural network certification problem. The optimal partitioning problem in the first-layer, as formulated in* (29)*, is NP-hard.*

Proof See Appendix C. ■

This concludes our development and analysis for partitioning the LP relaxation. In the next section, we follow a similar line of reasoning to develop a partitioning scheme for the other popular convex robustness certification technique, i.e., the SDP relaxation. Despite approaching this relaxation from the same partitioning perspective as the LP, the vastly different geometries of the LP and SDP feasible sets make the partitioning procedures quite distinct.

4. Partitioned SDP Relaxation

4.1 Tightening of the Relaxation

As with the LP relaxation, we begin by showing that the SDP relaxation error is decreased when the input uncertainty set is partitioned. This proposition is formalized below.

Proposition 11 (Improving the SDP relaxation bound) *Consider a neural network with one hidden ReLU layer. Let $\{\mathcal{X}^{(j)} : j \in \{1, 2, \dots, p\}\}$ be a partition of $\mathcal{X}$. For the j^{th} input part $\mathcal{X}^{(j)}$, denote the corresponding input bounds by $l \leq l^{(j)} \leq x \leq u^{(j)} \leq u$, where $x \in \mathcal{X}^{(j)}$. Then, it holds that*

$$\max_{j \in \{1,2,\dots,p\}} \hat{f}^*_{\text{SDP}}(\mathcal{X}^{(j)}) \leq \hat{f}^*_{\text{SDP}}(\mathcal{X}). \tag{31}$$

Proof See Appendix D. ■

Proposition 11 guarantees that partitioning yields a tighter SDP relaxation. However, it is not immediately clear how to design the partition in order to maximally reduce the relaxation error. Indeed, a poorly designed partition may even yield an equality in the bound (31). One notable challenge in designing the SDP partition relates to an inherent difference between the SDP relaxation and the LP relaxation. With the LP relaxation, the effect of partitioning can be visualized by how the geometry of the feasible set changes; see Figure 4. However, with the SDP, the relaxation comes from dropping the nonconvex rank constraint, the geometry of which is more abstract and harder to exploit.

In the next section, we develop a bound measuring how far the SDP solution is from being rank-1, which corresponds to an exact relaxation, where the improvement in (31) is as good as possible. By studying the geometry of the SDP feasible set through this more tractable bound, we find that the partition design for the SDP naturally reduces to a uniform partition along the coordinate axes of the input set.

4.2 Motivating Partition

In this section, we seek the form of a partition that best reduces the SDP relaxation error. By restricting our focus to ReLU networks with one hidden layer, we develop a simple necessary condition for the SDP relaxation to be exact, i.e., for the matrix P to be rank-1. We then work on the violation of this condition to define a measure of how close P is to being rank-1 in the case it has higher rank. Next, we develop a tractable upper bound on this rank-1 gap. Finally, we formulate an optimization problem in which we search for a partition of the input uncertainty set that minimizes our upper bound. We show that an optimal partition takes the form of a uniform division of the input set. The result motivates the use of uniform partitions of the input uncertainty set, and in Section 4.3, we answer the question of which coordinate is best to uniformly partition along. Note that, despite the motivating partition being derived for networks with one hidden layer, the relaxation tightening in Proposition 11 still holds for multi-layer networks. Indeed, the experiments in Section 5 will show that the resulting SDP partition design maintains a relatively constant efficacy as the number of layers increases.

Proposition 12 (Necessary condition for exact SDP) *Let $P^* \in \mathbb{S}^{n_x+n_z+1}$ denote a solution to the semidefinite programming relaxation* (8)*. If the relaxation is exact, meaning that* $\text{rank}(P^*) = 1$*, then the following conditions hold:*

$$\text{tr}(P^*_{xx}) = \|P^*_x\|_2^2, \quad \text{tr}(P^*_{zz}) = \|P^*_z\|_2^2. \tag{32}$$

Proof Since the SDP relaxation is exact, it holds that $\text{rank}(P^*) = 1$. Therefore, P^* can be expressed as

$$P^* = \begin{bmatrix} 1 \\ v \\ w \end{bmatrix} \begin{bmatrix} 1 & v^\top & w^\top \end{bmatrix} = \begin{bmatrix} 1 & v^\top & w^\top \\ v & vv^\top & vw^\top \\ w & wv^\top & ww^\top \end{bmatrix}$$

for some vectors $v \in \mathbb{R}^{n_x}$ and $w \in \mathbb{R}^{n_z}$. Recall the block decomposition of P^*:

$$P^* = \begin{bmatrix} P^*_1 & P^{*\top}_x & P^{*\top}_z \\ P^*_x & P^*_{xx} & P^*_{xz} \\ P^*_z & P^*_{zx} & P^*_{zz} \end{bmatrix}.$$

Equating coefficients, we find that $P^*_{xx} = vv^\top = P^*_x P^{*\top}_x$ and $P^*_{zz} = ww^\top = P^*_z P^{*\top}_z$. Therefore,

$$\text{tr}(P^*_{xx}) = \text{tr}(P^*_x P^{*\top}_x) = \text{tr}(P^{*\top}_x P^*_x) = \|P^*_x\|_2^2,$$

proving the first condition in (32). The second condition follows in the same way. ■

Enforcing the conditions (32) as constraints in the SDP relaxation may assist in pushing the optimization variable P towards a rank-1 solution. However, because the conditions in (32) are nonlinear equality constraints in the variable P, we cannot impose them directly on the SDP without making the problem nonconvex. Instead, we will develop a convex method based on the rank-1 conditions (32) that can be used to motivate the SDP solution to have a lower rank.

In the general case that $\text{rank}(P) = r \geq 1$, P may be written as $P = VV^\top$, where

$$V = \begin{bmatrix} e^\top \\ X \\ Z \end{bmatrix}, \; e \in \mathbb{R}^r, \; X \in \mathbb{R}^{n_x \times r}, \; Z \in \mathbb{R}^{n_z \times r},$$

and where the vector e satisfies the equation $e^\top e = \|e\|_2^2 = 1$. The i^{th} row of X (respectively, Z) is denoted by $X_i^\top \in \mathbb{R}^{1\times r}$ (respectively, $Z_i^\top \in \mathbb{R}^{1\times r}$). Under this expansion, we find that $P_x = Xe$, $P_z = Ze$, $P_{xx} = XX^\top$, and $P_{zz} = ZZ^\top$. Therefore, the conditions (32) can be written as

$$\text{tr}(XX^\top) = \|Xe\|_2^2, \quad \text{tr}(ZZ^\top) = \|Ze\|_2^2.$$

To simplify the subsequent analysis, we will restrict our attention to the first of these two necessary conditions for P to be rank-1. As the simulation results in Section 5 show, this restriction still yields significant reduction in relaxation error. Now, note that

$$\text{tr}(XX^\top) = \sum_{i=1}^{n_x} (XX^\top)_{ii} = \sum_{i=1}^{n_x} \|X_i\|_2^2,$$

where $(XX^\top)_{ii}$ is the (i,i) element of the matrix $XX^\top$, and also that

$$\|Xe\|_2^2 = \sum_{i=1}^{n_x}(Xe)_i^2 = \sum_{i=1}^{n_x}(X_i^\top e)^2,$$

where $(Xe)_i$ is the i^{th} element of the vector Xe. Therefore, the rank-1 necessary condition is equivalently written as

$$g(P) \coloneqq \sum_{i=1}^{n_x}(\|X_i\|_2^2 - (X_i^\top e)^2) = 0,$$

where $g(P)$ serves as a measure of the rank-1 gap. Note that $g(P)$ is solely determined by $P = VV^\top$, even though it is written in terms of X and e, which are blocks of V. In general, $g(P) \geq 0$ when $\operatorname{rank}(P) \geq 1$.

Lemma 13 (Rank-1 gap) *Let $P \in \mathbb{S}^{n_x+n_z+1}$ be an arbitrary feasible point for the SDP relaxation* (8)*. The rank-1 gap $g(P)$ is nonnegative, and is zero if P is rank-1.*

Proof By the Cauchy-Schwarz inequality, we have that $|X_i^\top e| \leq \|X_i\|_2\|e\|_2$ for all $i \in \{1,2,\dots,n_x\}$. Since P is feasible for (8) we also have that $\|e\|_2 = 1$, so squaring both sides of the inequality gives that $(X_i^\top e)^2 \leq \|X_i\|_2^2$. Summing these inequalities over i gives

$$g(P) = \sum_{i=1}^{n_x}(\|X_i\|_2^2 - (X_i^\top e)^2) \geq 0.$$

If P is rank-1, then the dimension r of the vectors e and X_i is equal to 1. That is, $e, X_i \in \mathbb{R}$. Hence, $\|X_i\|_2 = |X_i|$ and $|e| = \|e\|_2 = 1$, yielding $\|X_i\|_2^2 - (X_i^\top e)^2 = X_i^2 - X_i^2 e^2 = 0$. Therefore, $g(P) = 0$ in the case that $\operatorname{rank}(P) = 1$. ■

Since $g(P) = 0$ is necessary for P to be rank-1 and $g(P) \geq 0$, it is desirable to make $g(P^*)$ as small as possible at the optimal solution P^* of the partitioned SDP relaxation. Indeed, this is our partitioning motivation: we seek to partition the input uncertainty set to minimize $g(P^*)$, in order to influence P^* to be of low rank. However, there is a major hurdle with this approach. In particular, the optimal solution P^* depends on the partition we choose, and finding a partition to minimize $g(P^*)$ in turn depends on P^* itself. To overcome this cyclic dependence, we propose first bounding $g(P^*)$ by a worst-case upper bound, and then choosing an optimal partition to minimize the upper bound. This will make the partition design tractable, resulting in a closed-form solution.

To derive the upper bound on the rank-1 gap at optimality, let $\{\mathcal{X}^{(j)} : j \in \{1,2,\dots,p\}\}$ denote the partition of $\mathcal{X}$. For the j^{th} input part $\mathcal{X}^{(j)}$, denote the corresponding input bounds by $l^{(j)}, u^{(j)}$. The upper bound is derived below.

Lemma 14 (Rank-1 gap upper bound) *The rank-1 gap at the solution P^* of the partitioned SDP satisfies*

$$0 \leq g(P^*) \leq \frac{1}{4}\sum_{i=1}^{n_x} \max_{j\in\{1,2,\dots,p\}} (u_i^{(j)} - l_i^{(j)})^2. \tag{33}$$

Proof The left inequality is a direct result of Lemma 13. For the right inequality, note that

$$g(P^*) \leq \max_{j\in\{1,2,\dots,p\}} \sup_{P\in\mathcal{N}_{\text{SDP}}^{(j)},\ P_x\in\mathcal{X}^{(j)}} g(P) \leq \sum_{i=1}^{n_x} \max_{j\in\{1,2,\dots,p\}} \sup_{P\in\mathcal{N}_{\text{SDP}}^{(j)},\ P_x\in\mathcal{X}^{(j)}} (\|X_i\|_2^2 - (X_i^\top e)^2). \tag{34}$$

Let us focus on the optimization over the j^{th} part of the partition, namely,

$$\sup_{P\in\mathcal{N}_{\text{SDP}}^{(j)},\ P_x\in\mathcal{X}^{(j)}} (\|X_i\|_2^2 - (X_i^\top e)^2).$$

To bound this quantity, we analyze the geometry of the SDP relaxation over part j, following the methodology of Raghunathan et al. (2018); see Figure 5.

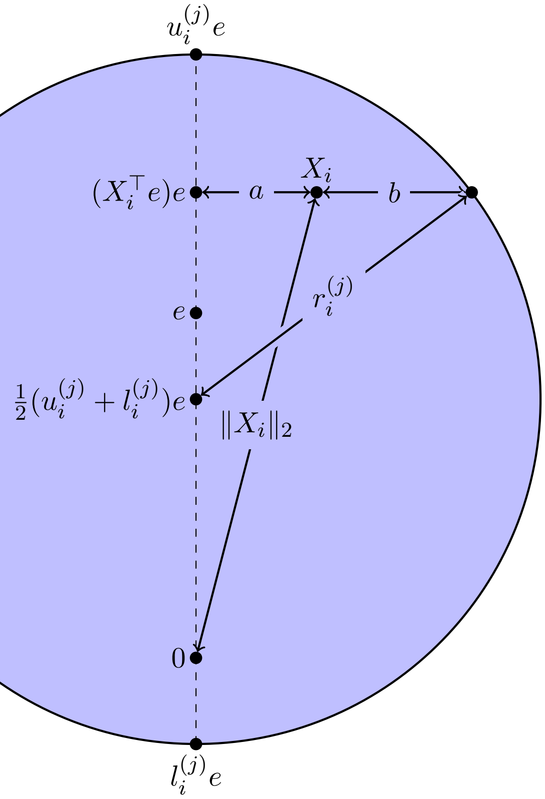

Figure 5: Geometry of the SDP relaxation in coordinate i over part j of the partition. The shaded region shows the feasible X_i satisfying the input constraint (Raghunathan et al., 2018).

The shaded circle represents the set of feasible X_i over part j of the partition, namely, those satisfying the i^{th} coordinate of the constraint $\operatorname{diag}(P_{xx}) \leq (l^{(j)}+u^{(j)})\odot P_x - l^{(j)}\odot u^{(j)}$. To understand this, note that the constraint is equivalent to $\|X_i\|_2^2 \leq (l_i^{(j)}+u_i^{(j)})X_i^\top e -$

$l_i^{(j)}u_i^{(j)}$, or, more geometrically written, that $\|X_i - \frac{1}{2}(u_i^{(j)} + l_i^{(j)})e\|_2^2 \leq \left(\frac{1}{2}(u_i^{(j)} - l_i^{(j)})\right)^2$. This shows that X_i is constrained to a 2-norm ball of radius $r_i^{(j)} = \frac{1}{2}(u_i^{(j)} - l_i^{(j)})$ centered at $\frac{1}{2}(u_i^{(j)} + l_i^{(j)})e$, as shown in Figure 5.

The geometry of Figure 5 immediately shows that $\|X_i\|_2^2 = a^2 + (X_i^\top e)^2$ and $r_i^{(j)2} = (a+b)^2 + (X_i^\top e - \frac{1}{2}(u_i^{(j)} + l_i^{(j)}))^2$, and therefore

$$\|X_i\|_2^2 - (X_i^\top e)^2 = a^2 = r_i^{(j)2} - (X_i^\top e - \tfrac{1}{2}(u_i^{(j)} + l_i^{(j)}))^2 - 2ab - b^2.$$

Since a and b are nonnegative,

$$\begin{aligned}
&\sup_{P \in \mathcal{N}_{\mathrm{SDP}}^{(j)},\ P_x \in \mathcal{X}^{(j)}} \|X_i\|_2^2 - (X_i^\top e)^2 \\
&= \sup_{P \in \mathcal{N}_{\mathrm{SDP}}^{(j)},\ P_x \in \mathcal{X}^{(j)}} (r_i^{(j)2} - (X_i^\top e - \tfrac{1}{2}(u_i^{(j)} + l_i^{(j)}))^2 - 2ab - b^2) \\
&\leq \sup_{P \in \mathcal{N}_{\mathrm{SDP}}^{(j)},\ P_x \in \mathcal{X}^{(j)}} (r_i^{(j)2} - (X_i^\top e - \tfrac{1}{2}(u_i^{(j)} + l_i^{(j)}))^2) \leq r_i^{(j)2} = \frac{1}{4}(u_i^{(j)} - l_i^{(j)})^2.
\end{aligned}$$

Thus, (34) gives

$$g(P^*) \leq \frac{1}{4}\sum_{i=1}^{n_x} \max_{j \in \{1,2,\dots,p\}} (u_i^{(j)} - l_i^{(j)})^2,$$

as desired. ■

With Lemma 14 in place, we now have an upper bound on the rank-1 gap at optimality, in terms of the input bounds $\{l^{(j)}, u^{(j)}\}_{j=1}^p$ associated with the partition. At this point, we turn to minimizing the upper bound over all valid choices of p-part partitions of the input uncertainty set along a given coordinate. Note that, in order for $\{l^{(j)}, u^{(j)}\}_{j=1}^p$ to define valid input bounds for a p-part partition, it must be that the union of the input parts cover the input uncertainty set. In terms of the input bounds, this leads to the constraint that $[l,u] = \cup_{j=1}^p [l^{(j)}, u^{(j)}]$, where $[l,u] := [l_1,u_1] \times [l_2,u_2] \times \cdots \times [l_{n_x}, u_{n_x}]$, and similarly for $[l^{(j)}, u^{(j)}]$. Since we consider the partition along a single coordinate k, this constraint becomes equivalent to $\cup_{j=1}^p [l_k^{(j)}, u_k^{(j)}] = [l_k, u_k]$, because all other coordinates $i \neq k$ satisfy $l_i^{(j)} = l_i$ and $u_i^{(j)} = u_i$ for all j by assumption. We now give the optimal partitioning scheme for the SDP that minimizes the upper bound in Lemma 14.

Theorem 15 (Optimal SDP partition via rank-1 gap) *Let $\mathcal{I}_k = \{1,2,\dots,n_x\} \setminus \{k\}$. Consider the optimization problem of finding the partition to minimize the upper bound*

(33), *namely*

$$
\begin{aligned}
\underset{\mathcal{P}=\{l^{(j)},u^{(j)}\}_{j=1}^{p}\subseteq\mathbb{R}^{n_x}}{\text{minimize}} \quad & h(\mathcal{P}) = \sum_{i=1}^{n_x} \max_{j\in\{1,2,\dots,p\}} (u_i^{(j)} - l_i^{(j)})^2 \\
\text{subject to} \quad & \bigcup_{j=1}^{p} [l_k^{(j)}, u_k^{(j)}] = [l_k, u_k], && i \in \mathcal{I}_k,\ j\in\{1,2,\dots,p\}, \\
& l_i^{(j)} = l_i, && i \in \mathcal{I}_k,\ j\in\{1,2,\dots,p\}, \\
& u_i^{(j)} = u_i, && i \in \mathcal{I}_k,\ j\in\{1,2,\dots,p\},
\end{aligned}
\tag{35}
$$

Consider also the uniform partition defined by $\bar{\mathcal{P}} = \{\bar{l}^{(j)}, \bar{u}^{(j)}\}_{j=1}^{p} \subseteq \mathbb{R}^{n_x}$*, where*

$$
\bar{l}_i^{(j)} = \begin{cases} \frac{j-1}{p}(u_i - l_i) + l_i & \text{if } i = k, \\ l_i & \text{otherwise,} \end{cases}
$$

$$
\bar{u}_i^{(j)} = \begin{cases} \frac{j}{p}(u_i - l_i) + l_i & \text{if } i = k, \\ u_i & \text{otherwise,} \end{cases}
$$

for all $j \in \{1,2,\dots,p\}$*. It holds that* $\bar{\mathcal{P}}$ *is a solution to* (35)*.*

Proof To prove the result, we show that the proposed $\bar{\mathcal{P}}$ is feasible for the optimization, and that $h(\bar{\mathcal{P}}) \le h(\mathcal{P})$ for all feasible $\mathcal{P}$. First, note that it is obvious by the definition of $\bar{\mathcal{P}}$ that $\bar{l}_i^{(j)} = l_i$ and $\bar{u}_i^{(j)} = u_i$ for all $i \in \{1,2,\dots,n_x\} \setminus \{k\}$ and all $j \in \{1,2,\dots,p\}$. Therefore, to prove that $\bar{\mathcal{P}}$ is feasible, it suffices to show that $\cup_{j=1}^{p} [\bar{l}_k^{(j)}, \bar{u}_k^{(j)}] = [l_k, u_k]$. Indeed, since

$$
\bar{u}_k^{(j)} = \frac{j}{p}(u_k - l_k) + l_k = \frac{(j+1)-1}{p}(u_k - l_k) + l_k = \bar{l}_k^{(j+1)}
$$

for all $j \in \{1,2,\dots,p-1\}$,

$$
\bar{l}_k^{(1)} = \frac{1-1}{p}(u_k - l_k) + l_k = l_k,
$$

and

$$
\bar{u}_k^{(p)} = \frac{p}{p}(u_k - l_k) + l_k = u_k,
$$

we have that

$$
\bigcup_{j=1}^{p} [\bar{l}_k^{(j)}, \bar{u}_k^{(j)}] = [\bar{l}_k^{(1)}, \bar{u}_k^{(1)}] \cup [\bar{l}_k^{(2)}, \bar{u}_k^{(2)}] \cup \cdots \cup [\bar{l}_k^{(p)}, \bar{u}_k^{(p)}] = [\bar{l}_k^{(1)}, \bar{u}_k^{(p)}] = [l_k, u_k].
$$

Hence, $\bar{\mathcal{P}} = \{\bar{l}^{(j)}, \bar{u}^{(j)}\}_{j=1}^{p}$ is feasible.

The objective at the proposed feasible point can be computed as

$$\begin{aligned} h(\bar{\mathcal{P}}) &= \sum_{i=1}^{n_x} \max_{j\in\{1,2,\dots,p\}} (\bar{u}_i^{(j)} - \bar{l}_i^{(j)})^2 = \sum_{\substack{i=1\\ i\neq k}}^{n_x} \max_{j\in\{1,2,\dots,p\}} (\bar{u}_i^{(j)} - \bar{l}_i^{(j)})^2 + \max_{j\in\{1,2,\dots,p\}} (\bar{u}_k^{(j)} - \bar{l}_k^{(j)})^2 \\ &= \sum_{\substack{i=1\\ i\neq k}}^{n_x} \max_{j\in\{1,2,\dots,p\}} (u_i - l_i)^2 + \max_{j\in\{1,2,\dots,p\}} \left(\frac{j}{p}(u_k - l_k) + l_k - \frac{j-1}{p}(u_k - l_k) - l_k\right)^2 \\ &= \sum_{\substack{i=1\\ i\neq k}}^{n_x} (u_i - l_i)^2 + \max_{j\in\{1,2,\dots,p\}} \left(\frac{1}{p}(u_k - l_k)\right)^2 = C + \frac{1}{p^2}(u_k - l_k)^2, \end{aligned}$$

where $C \coloneqq \sum_{\substack{i=1\\ i\neq k}}^{n_x} (u_i - l_i)^2$. Now, let $\mathcal{P} = \{l^{(j)}, u^{(j)}\}_{j=1}^p$ be an arbitrary feasible point for the optimization (35). Then by a similar analysis as above, the objective value at $\mathcal{P}$ satisfies

$$\begin{aligned} h(\mathcal{P}) &= \sum_{i=1}^{n_x} \max_{j\in\{1,2,\dots,p\}} (u_i^{(j)} - l_i^{(j)})^2 = \sum_{\substack{i=1\\ i\neq k}}^{n_x} \max_{j\in\{1,2,\dots,p\}} (u_i^{(j)} - l_i^{(j)})^2 + \max_{j\in\{1,2,\dots,p\}} (u_k^{(j)} - l_k^{(j)})^2 \\ &= C + \max_{j\in\{1,2,\dots,p\}} (u_k^{(j)} - l_k^{(j)})^2 = C + \left(\max_{j\in\{1,2,\dots,p\}} (u_k^{(j)} - l_k^{(j)})\right)^2 \\ &\geq C + \left(\frac{1}{p}\sum_{j=1}^{p} (u_k^{(j)} - l_k^{(j)})\right)^2 = C + \frac{1}{p^2}\left(\sum_{j=1}^{p} (u_k^{(j)} - l_k^{(j)})\right)^2. \end{aligned}$$

Since $\mathcal{P}$ is feasible, it holds that $[l_k, u_k] = \bigcup_{j=1}^p [l_k^{(j)}, u_k^{(j)}]$. Therefore, by subadditivity of Lebesgue measure μ on the Borel σ-algebra of $\mathbb{R}$, we have that

$$u_k - l_k = \mu([l_k, u_k]) = \mu\left(\bigcup_{j=1}^p [l_k^{(j)}, u_k^{(j)}]\right) \leq \sum_{j=1}^p \mu([l_k^{(j)}, u_k^{(j)}]) = \sum_{j=1}^p (u_k^{(j)} - l_k^{(j)}).$$

Substituting this into our above expressions, we conclude that

$$h(\bar{\mathcal{P}}) = C + \frac{1}{p^2}(u_k - l_k)^2 \leq C + \frac{1}{p^2}\left(\sum_{j=1}^p (u_k^{(j)} - l_k^{(j)})\right)^2 \leq h(\mathcal{P}).$$

Since $\mathcal{P}$ was an arbitrary feasible point for the optimization, this implies that $\bar{\mathcal{P}}$ is a solution to the optimization. ■

Theorem 15 shows that by choosing the partition of the input set to be uniformly divided amongst the p parts, we obtain an optimal partition that minimizes the worst-case bound on the gap of the rank-1 necessary condition (32). This gives a well-motivated, yet simple way to design a partition of the input uncertainty set in order to push the SDP relaxation towards being rank-1, thereby reducing relaxation error.

4.3 Partitioning Scheme

With the motivating partition of Section 4.2 now established, we turn our attention from the *form* of an optimal SDP partition to the *coordinate* of an optimal partition. In particular, we seek to find the best two-part partition to minimize relaxation error of the SDP. The results of Section 4.2 suggest using a uniform partition, and in this section we seek to find which coordinate to apply the partitioning to. Similar to the LP relaxation, we derive an optimal partitioning scheme by first bounding the relaxation error in the worst-case sense.

4.3.1 Worst-Case Relaxation Bound

In the worst-case relaxation bound of Theorem 18 below, and the subsequent optimal SDP partitioning scheme proposed in Theorem 19, we restrict our attention to a single hidden ReLU layer and make the following assumption on the weight matrix.

Assumption 16 (Normalized rows) *The rows of the weight matrix are assumed to be normalized with respect to the ℓ_1-norm, i.e., that $\|w_i\|_1 = 1$ for all $i \in \{1, 2, \dots, n_z\}$.*

We briefly remark that Assumption 16 imposes no loss of generality, as it can be made to hold for any network by a simple rescaling. In particular, if the assumption does not hold, the network architecture can be rescaled as follows:

$$z = \mathrm{ReLU}(Wx) = \mathrm{ReLU}\left(\begin{bmatrix} w_1^\top \\ w_2^\top \\ \vdots \\ w_{n_z}^\top \end{bmatrix} x\right)$$

$$= \mathrm{ReLU}\left(\mathrm{diag}(\|w_1\|_1, \|w_2\|_1, \dots, \|w_{n_z}\|_1)\begin{bmatrix} \frac{w_1^\top}{\|w_1\|_1} \\ \frac{w_2^\top}{\|w_2\|_1} \\ \vdots \\ \frac{w_{n_z}^\top}{\|w_{n_z}\|_1} \end{bmatrix} x\right) = W_{\mathrm{scale}}\,\mathrm{ReLU}(W_{\mathrm{norm}}x),$$

where $W_{\mathrm{scale}} = \mathrm{diag}(\|w_1\|_1, \|w_2\|_1, \dots, \|w_{n_z}\|_1) \in \mathbb{R}^{n_z \times n_z}$ and

$$W_{\mathrm{norm}} = \begin{bmatrix} \frac{w_1^\top}{\|w_1\|_1} \\ \frac{w_2^\top}{\|w_2\|_1} \\ \vdots \\ \frac{w_{n_z}^\top}{\|w_{n_z}\|_1} \end{bmatrix} \in \mathbb{R}^{n_z \times n_x}$$

are the scaling and normalized factors of the weight matrix W, respectively. The scaling factor can therefore be absorbed into the optimization cost vector c, yielding a problem with normalized rows as desired.

Before introducing the worst-case relaxation bound of Theorem 18, we state a short lemma that will be used in proving the relaxation bound.

Lemma 17 (Bound on elements of PSD matrices) *Let $P \in \mathbb{S}^n$ be a positive semidefinite matrix. Then $|P_{ij}| \leq \frac{1}{2}(P_{ii} + P_{jj})$ for all $i, j \in \{1, 2, \dots, n\}$.*

Proof See Appendix E. ■

Theorem 18 (Worst-case relaxation bound for SDP) *Consider a feedforward ReLU neural network with one hidden layer, and with the input uncertainty set $\mathcal{X}$. Let the network have input bounds $l, u \in \mathbb{R}^{n_x}$ and preactivation bounds $\hat{l}, \hat{u} \in \mathbb{R}^{n_z}$. Consider also the relaxation error $\Delta f^*_{\text{SDP}}(\mathcal{X}) \coloneqq \hat{f}^*_{\text{SDP}}(\mathcal{X}) - f^*(\mathcal{X})$. Let P^* and (x^*, z^*) be optimal solutions for the relaxation $\hat{f}^*_{\text{SDP}}(\mathcal{X})$ and the unrelaxed problem $f^*(\mathcal{X})$, respectively. Given an arbitrary norm $\|\cdot\|$ on $\mathbb{R}^{n_x}$, there exists $\epsilon \in \mathbb{R}$ such that $\|P^*_x - x^*\| \leq \epsilon$ and*

$$\Delta f^*_{\text{SDP}}(\mathcal{X}) \leq \sum_{i=1}^{n_z} \bigg(\text{ReLU}(c_i) q(l, u) + \text{ReLU}(-c_i) \min\{\hat{u}_i, \epsilon \|w_i\|_*\} \bigg), \tag{36}$$

where $\|\cdot\|_*$ is the dual norm of $\|\cdot\|$, and where

$$q(l, u) = \max_{k \in \{1,2,\dots,n_x\}} \max\{|l_k|, |u_k|\}.$$

Proof First, since $\mathcal{X}$ is bounded, there exists $\epsilon \geq 0$ such that $\|P^*_x - x^*\| \leq \epsilon$. The definitions of P^*_x and (x^*, z^*) give that

$$\Delta f^*_{\text{SDP}}(\mathcal{X}) = \sum_{i=1}^{n_z} c_i((P^*_z)_i - z^*_i) \leq \sum_{i=1}^{n_z} \Delta f^*_i, \tag{37}$$

where

$$\Delta f^*_i = \sup \bigg\{ c_i((P_z)_i - z_i) : z_i = \text{ReLU}(w_i^\top x),\ P_z \geq 0,\ P_z \geq W P_x,$$
$$\text{diag}(P_{zz}) = \text{diag}(W P_{xz}),\ P_1 = 1,\ P \succeq 0,\ \|P_x - x\| \leq \epsilon,\ x, P_x \in \mathcal{X} \bigg\}$$

for all $i \in \{1, 2, \dots, n_z\}$. Defining the auxiliary variables $P_{\hat{z}} = W P_x$ and $\hat{z} = Wx$, this is equivalent to

$$\Delta f^*_i = \sup \bigg\{ c_i((P_z)_i - z_i) : z_i = \text{ReLU}(\hat{z}_i),\ P_z \geq 0,\ P_z \geq P_{\hat{z}},\ \text{diag}(P_{zz}) = \text{diag}(W P_{xz}),$$
$$P_1 = 1,\ P \succeq 0,\ \|P_x - x\| \leq \epsilon,\ P_{\hat{z}} = W P_x,\ \hat{z}_i = w_i^\top x,\ x, P_x \in \mathcal{X} \bigg\}.$$

If $x, P_x \in \mathcal{X}$ and $\hat{z}, P_{\hat{z}}$ satisfy $\hat{z} = Wx$, $P_{\hat{z}} = W P_x$, and $\|P_x - x\| \leq \epsilon$, then $|(P_{\hat{z}})_i - \hat{z}_i| = |w_i^\top (P_x - x)| \leq \|w_i\|_* \|P_x - x\| \leq \epsilon \|w_i\|_*$ for all $i \in \{1, 2, \dots, n_z\}$ by the Cauchy-Schwarz inequality for dual norms. Therefore,

$$\Delta f^*_i \leq \sup \bigg\{ c_i((P_z)_i - z_i) : z_i = \text{ReLU}(\hat{z}_i),\ P_z \geq 0,\ P_z \geq P_{\hat{z}},$$
$$\text{diag}(P_{zz}) = \text{diag}(W P_{xz}),\ P_1 = 1,\ P \succeq 0,\ \hat{l} \leq \hat{z}, P_{\hat{z}} \leq \hat{u},$$
$$|(P_{\hat{z}})_k - \hat{z}_k| \leq \epsilon \|w_k\|_* \text{ for all } k \in \{1, 2, \dots, n_z\},\ \hat{z}, P_{\hat{z}} \in \mathbb{R}^{n_z} \bigg\}.$$

We now translate the optimization variables in the above problem from $\hat{z} \in \mathbb{R}^{n_z}$ and $P \in \mathbb{S}^{1+n_x+n_z}$ to the scalars $\hat{z}_i, (P_{\hat{z}})_i \in \mathbb{R}$. To this end, we note that if P is feasible for the above supremum, then

$$\operatorname{diag}(P_{zz})_i = \operatorname{diag}(WP_{xz})_i = w_i^\top (P_{xz})_i \leq \|(P_{xz})_i\|_\infty \|w_i\|_1,$$

where $(P_{xz})_i$ is the i^{th} column of the matrix P_{xz}, and the inequality again comes from Cauchy-Schwarz. By the weight matrix scaling assumption, this yields

$$\operatorname{diag}(P_{zz})_i \leq \|(P_{xz})_i\|_\infty.$$

Now, since P is positive semidefinite, Lemma 17 gives that

$$\begin{aligned}\|(P_{xz})_i\|_\infty &= \max_{k\in\{1,2,\dots,n_x\}} |(P_{xz})_i|_k = \max_{k\in\{1,2,\dots,n_x\}} |(P_{xz})_{ki}| \\ &\leq \max_{k\in\{1,2,\dots,n_z\}} \frac{1}{2}\Big((P_{xx})_{kk} + (P_{zz})_{ii}\Big) = \frac{1}{2}(P_{zz})_{ii} + \frac{1}{2}\max_{k\in\{1,2,\dots,n_x\}} (P_{xx})_{kk}.\end{aligned}$$

Noting that $(P_{zz})_{ii} = \operatorname{diag}(P_{zz})_i$, the bound of interest becomes

$$\operatorname{diag}(P_{zz})_i \leq \max_{k\in\{1,2,\dots,n_z\}} (P_{xx})_{kk}.$$

We now seek to bound $(P_{xx})_{kk}$. Recall that $(P_{xx})_{kk} = \operatorname{diag}(P_{xx})_k \leq (l_k + u_k)(P_x)_k - l_k u_k$. If $(l_k + u_k) \geq 0$, then $(P_x)_k \leq u_k$ implies that $(l_k + u_k)(P_x)_k \leq (l_k + u_k)u_k$, and therefore $(P_{xx})_{kk} \leq (l_k + u_k)u_k - l_k u_k = u_k^2$. On the other hand, if $(l_k + u_k) < 0$, then $(P_x)_k \geq l_k$ implies that $(l_k + u_k)(P_x)_k \leq (l_k + u_k)l_k$, and therefore $(P_{xx})_{kk} \leq (l_k + u_k)l_k - l_k u_k = l_k^2$. Hence, in all cases, it holds that

$$(P_{xx})_{kk} \leq \mathbb{I}(l_k + u_k \geq 0)u_k^2 + \mathbb{I}(l_k + u_k < 0)l_k^2.$$

We can further simplify this bound as follows. If $l_k + u_k \geq 0$, then $u_k \geq -l_k$ and $u_k \geq l_k$, implying $|l_k| \leq u_k$, so $l_k^2 \leq u_k^2$ and therefore $u_k^2 = \max\{l_k^2, u_k^2\}$. On the other hand, if $l_k + u_k < 0$, then an analogous argument shows that $l_k^2 = \max\{l_k^2, u_k^2\}$. Hence, we conclude that the above bound on $(P_{xx})_{kk}$ can be rewritten as

$$(P_{xx})_{kk} \leq \max\{l_k^2, u_k^2\}.$$

Therefore, returning to the bound on $(P_{zz})_i$, we find that

$$\operatorname{diag}(P_{zz})_i \leq \max_{k\in\{1,2,\dots,n_x\}} \max\{l_k^2, u_k^2\},$$

for all $i \in \{1, 2, \dots, n_z\}$. Now, note that since $P \succeq 0$, the Schur complement gives that

$$\begin{bmatrix} P_{xx} - P_x P_x^\top & P_{xz} - P_x P_z^\top \\ P_{xz}^\top - P_z P_x^\top & P_{zz} - P_z P_z^\top \end{bmatrix} \succeq 0,$$

which implies that

$$\operatorname{diag}(P_{zz}) \geq \operatorname{diag}(P_z P_z^\top) = P_z \odot P_z.$$

Therefore, our upper bound on the diagonal elements of P_{zz} yields that

$$(P_z)_i \leq \max_{k \in \{1,2,\dots,n_x\}} \max\{|l_k|, |u_k|\} = q(l,u).$$

Hence, we have derived a condition on the component $(P_z)_i$ that all feasible P must satisfy. The supremum of interest may now be further upper bounded giving rise to

$$\Delta f_i^* \leq \sup \Big\{ c_i((P_z)_i - z_i) : z_i = \mathrm{ReLU}(\hat{z}_i),\ (P_z)_i \geq 0,\ (P_z)_i \geq (P_{\hat{z}})_i,\ (P_z)_i \leq q(l,u), \\ \hat{l}_i \leq \hat{z}_i, (P_{\hat{z}})_i \leq \hat{u}_i,\ |(P_{\hat{z}})_i - \hat{z}_i| \leq \epsilon \|w_i\|_*,\ \hat{z}_i, (P_{\hat{z}})_i \in \mathbb{R} \Big\}, \quad (38)$$

which is now in terms of the scalar optimization variables $\hat{z}_i$ and $(P_{\hat{z}})_i$, as we desired. This reformulation makes it tractable to compute the supremum in (38) in closed-form, which we now turn to do.

First, consider the case that $c_i \geq 0$. Then we seek to maximize the difference $(P_z)_i - z_i$ subject to the given constraints. Noting that $(P_z)_i \leq q(l,u)$ and $z_i \geq 0$ on the above feasible set, we remark that the objective is upper bounded as $c_i((P_z)_i - z_i) \leq c_i q(l,u)$. Indeed, this upper bound is attained at the feasible point defined by $z_i = \hat{z}_i = (P_{\hat{z}})_i = 0$ and $(P_z)_i = q(l,u)$. Hence, we conclude that for all $i \in \{1,2,\dots,n_z\}$ such that $c_i \geq 0$, it holds that

$$\Delta f_i^* \leq c_i q(l,u). \quad (39)$$

Now consider the case that $c_i < 0$. Then we seek to minimize the difference $(P_z)_i - z_i$ subject to the given constraints. In this case, the optimal objective value depends on the relative sizes of $\hat{u}_i$ and $\epsilon\|w_i\|_*$. In particular, when $\hat{u}_i \leq \epsilon\|w_i\|_*$, the constraint $\hat{z}_i \leq \hat{u}_i$ becomes active at optimum, yielding a supremum value of $-c_i u_i$. Alternatively, when $\epsilon\|w_i\|_* \leq \hat{u}_i$, the constraint $|(P_{\hat{z}})_i - \hat{z}_i| \leq \epsilon\|w_i\|_*$ becomes active at optimum, yielding the supremum value of $-c_i\epsilon\|w_i\|_*$. Therefore, we conclude that for all $i \in \{1,2,\dots,n_z\}$ such that $c_i < 0$, it holds that

$$\Delta f_i^* \leq -c_i \min\{\hat{u}_i, \epsilon\|w_i\|_*\}. \quad (40)$$

Substituting (39) and (40) into (37) gives the desired bound. ■

When the x-block P_x^* of the SDP relaxation stays close to the true solution x^*, the bound (36) shows that the worst-case relaxation error scales with the loosest input bound, i.e., the maximum value amongst the limits $|l_k|$ and $|u_k|$. This fact allows us to choose which coordinate to partition along in order to maximally reduce the relaxation bound on the individual parts of the partition. We state our proposed SDP partition next.

4.3.2 Proposed Partition

We now focus on developing an optimal two-part partitioning scheme based on the worst-case relaxation bound of Theorem 18. Similar to the LP relaxation approach, we take $\epsilon \approx 0$ to simplify the analysis (refer to the LP analysis for justifications of this simplification). The bound therefore takes the form

$$\Delta f_{\mathrm{SDP}}^*(\mathcal{X}) \leq q(l,u) \sum_{i=1}^{n_z} \mathrm{ReLU}(c_i), \quad (41)$$

where

$$q(l,u) = \max_{k\in\{1,2,\dots,n_x\}} \max\{|l_k|,|u_k|\}.$$

Since the design of the partition amounts to choosing input bounds l and u for the input parts, the input bounds serve as our optimization variables in minimizing the above worst-case relaxation bound. By restricting the form of our partition to the uniform division motivated in Theorem 15, it follows from the form of q that the best coordinate to partition along is that with the loosest input bound, i.e., along coordinate $i^* \in \arg\max_{k\in\{1,2,\dots,n_x\}} \max\{|l_k|,|u_k|\}$. This observation is formalized below.

Theorem 19 (Optimal SDP partition) *Consider the two-part partitions defined by dividing $\mathcal{X}$ uniformly along the coordinate axes: $\{\mathcal{X}_i^{(1)},\mathcal{X}_i^{(2)}\}$, with $\mathcal{X}_i^{(1)} = \{x\in\mathcal{X} : l_i^{(1)} \le x \le u_i^{(1)}\}$ and $\mathcal{X}_i^{(2)} = \{x\in\mathcal{X} : l_i^{(2)} \le x \le u_i^{(2)}\}$, where $l_i^{(1)} = l$, $u_i^{(1)} = (u_1,u_2,\dots,u_{i-1},\frac{1}{2}(l_i+u_i),u_{i+1},\dots,u_{n_x})$, $l_i^{(2)} = (l_1,l_2,\dots,l_{i-1},\frac{1}{2}(l_i+u_i),l_{i+1},\dots,l_{n_x})$, and $u_i^{(2)} = u$, for all $i\in\{1,2,\dots,n_x\} =: \mathcal{I}$. Let*

$$i^* \in \arg\max_{k\in\{1,2,\dots,n_x\}} \max\{|l_k|,|u_k|\}, \tag{42}$$

and assume that $|l_{i^*}| \neq |u_{i^*}|$. Then the partition $\{\mathcal{X}_{i^*}^{(1)},\mathcal{X}_{i^*}^{(2)}\}$ is optimal in the sense that the upper bound factor $q(l_{i^*}^{(j)},u_{i^*}^{(j)})$ in (41) equals the unpartitioned upper bound $q(l,u)$ on one part j of the partition, is strictly less than $q(l,u)$ on the other part, and $q(l_i^{(j)},u_i^{(j)}) = q(u,l)$ for both $j\in\{1,2\}$ for all other $i \notin \arg\max_{k\in\{1,2,\dots,n_x\}} \max\{|l_k|,|u_k|\}$.

Proof First, consider partitioning along coordinate $i \notin \arg\max_{k\in\{1,2,\dots,n_x\}} \max\{|l_k|,|u_k|\}$. Then

$$q(l_i^{(1)},u_i^{(1)}) = \max_{k\in\{1,2,\dots,n_x\}} \max\{|(l_i^{(1)})_k|,|(u_i^{(1)})_k|\}$$
$$= \max\left\{|l_1|,\dots,|l_{n_x}|,|u_1|,\dots,\left|\frac{l_i+u_i}{2}\right|,\dots,|u_{n_x}|\right\} = \max\{|l_{i^*}|,|u_{i^*}|\} = q(l,u),$$

since $\left|\frac{l_i+u_i}{2}\right| \le \frac{|l_i|+|u_i|}{2} < \max\{|l_{i^*}|,|u_{i^*}|\}$ and $i\neq i^*$ implies that

$$\max\{|l_{i^*}|,|u_{i^*}|\} \in \left\{|l_1|,\dots,|l_{n_x}|,|u_1|,\dots,\left|\frac{l_i+u_i}{2}\right|,\dots,|u_{n_x}|\right\}.$$

In an analogous fashion, it follows that

$$q(l_i^{(2)},u_i^{(2)}) = q(l,u).$$

Now, consider partitioning along coordinate i^*. Note that either

$$\max_{k\in\{1,2,\dots,n_x\}} \max\{|l_k|,|u_k|\} = |l_{i^*}|, \text{ or } \max_{k\in\{1,2,\dots,n_x\}} \max\{|l_k|,|u_k|\} = |u_{i^*}|.$$

Suppose that the first case holds true. Then

$$q(l_{i^*}^{(1)}, u_{i^*}^{(1)}) = \max_{k\in\{1,2,\dots,n_x\}} \max\{|(l_{i^*}^{(1)})_k|, |(u_{i^*}^{(1)})_k|\}$$
$$= \max\left\{|l_1|,\dots,|l_{n_x}|,|u_1|,\dots,\left|\frac{l_{i^*}+u_{i^*}}{2}\right|,\dots,|u_{n_x}|\right\} = |l_{i^*}| = q(l,u),$$

since $|(l_{i^*}+u_{i^*})/2| \leq (|l_i^*|+|u_i^*|)/2 < |l_i^*|$ and

$$|l_{i^*}| \in \max\left\{|l_1|,\dots,|l_{n_x}|,|u_1|,\dots,\left|\frac{l_{i^*}+u_{i^*}}{2}\right|,\dots,|u_{n_x}|\right\}.$$

Over the second part of the partition,

$$q(l_{i^*}^{(2)}, u_{i^*}^{(2)}) = \max_{k\in\{1,2,\dots,n_x\}} \max\{|(l_{i^*}^{(2)})_k|, |(u_{i^*}^{(2)})_k|\}$$
$$= \max\left\{|l_1|,\dots,\left|\frac{l_{i^*}+u_{i^*}}{2}\right|,\dots,|l_{n_x}|,|u_1|,\dots,|u_{n_x}|\right\} < |l_{i^*}| = q(l,u),$$

since $|(l_{i^*}+u_{i^*})/2| < |l_{i^*}|$ and

$$|l_i^*| \notin \left\{|l_1|,\dots,\left|\frac{l_{i^*}+u_{i^*}}{2}\right|,\dots,|l_{n_x}|,|u_1|,\dots,|u_{n_x}|\right\}$$

since $|l_{i^*}| \neq |u_{i^*}|$. In the other case that $\max_{k\in\{1,2,\dots,n_x\}}\max\{|l_k|,|u_k|\} = |u_{i^*}|$, it follows via the same argument that $q(l_{i^*}^{(1)}, u_{i^*}^{(1)}) < q(l,u)$ and $q(l_{i^*}^{(2)}, u_{i^*}^{(2)}) = q(l,u)$. Since partitioning along any other coordinate $i \notin \arg\max_{k\in\{1,2,\dots,n_x\}}\max\{|l_k|,|u_k|\}$ was shown to yield $q(l_i^{(1)}, u_i^{(1)}) = q(l_i^{(2)}, u_i^{(2)}) = q(u,l)$, we conclude that the coordinate i^* is optimal in the sense proposed. ■

Intuitively, the partitioning scheme defined in Theorem 19 is optimal because any other uniform partition along a coordinate axis cannot tighten the worst-case relaxation error bound (41). On the other hand, Theorem 19 guarantees that using the partition coordinate in (42) results in a strict tightening of the worst-case relaxation error on at least one part of the partition.

5. Simulation Results

In this section, we experimentally corroborate the effectiveness of our proposed certification methods. We first perform the partitioned LP relaxation on an IRIS classification network and show that our proposed partitioning scheme is able to generate a robustness certificate which was previously unattainable. We then compute the SDP and partitioned SDP on the same network and compare to the LP results. Finally, we explore the effectiveness of partitioning the LP and SDP as the network grows in size, namely, as the number of inputs and the number of layers independently increase. All experiments are performed using Matlab and CVX on a standard laptop with a 2.9 GHz quad-core i7 processor.

5.1 Partitioned LP Results

In this experiment, we consider a classification network trained on the Iris data set (Fisher, 1936) with a test accuracy of 97%. The network has a single hidden layer, four inputs, and three outputs. We consider 10 different nominal inputs $\bar{x}$ and corresponding uncertainty sets $\mathcal{X} = \{x \in \mathbb{R}^{n_x} : \|x - \bar{x}\|_\infty \leq \epsilon\}$, where $\epsilon = 0.1$. In this setting, a negative optimal objective value of the robustness certification problem proves that no perturbation of $\bar{x}$ within $\mathcal{X}$ results in misclassification.

For each nominal input being tested, we first solve the original, nonconvex certification problem to local optimality using Matlab's `fmincon` local search function with 5 multi-starts. Upon using more than 5 starts, we found no change in the nonconvex objective values. Recall that since the unrelaxed certification problem is nonconvex, the solutions found using the `fmincon` local search are not known to be globally optimal. Consequently, even though the objective values of these local solutions may suggest that the network is robust, they do not provide a formal certificate of robustness, since in general a gap between local and global optimality may exist. Next, we solve the unpartitioned LP relaxation. Finally, we solve three two-part partitioned LP relaxations, one per row of the weight matrix W, to explore the effectiveness of partitioning along different rows, and how the results compare to the optimality given by Theorem 8. The average times (taken over the 10 nominal inputs) to solve the nonconvex approximation, unpartitioned LP relaxation, and partitioned LP relaxations are 0.21, 0.26, and 0.48 seconds, respectively. The computational cost of the two-part partitioned LP is twice that of the unpartitioned LP, both of which are very fast for this network. This factor of 2 is expected, since the two LPs in the partitioned LP are solved sequentially. However, we remark that the two LPs can easily be parallelized in order to speed up the computation.

Figure 6a displays the resulting optimal objective values. The optimally partitioned LP, i.e., that partitioned along row $w_{i^*}^\top$ from Theorem 8, is clearly shown to give the tightest upper bound for every nominal input tested. Two other partitioned LPs are displayed. In particular, the second-best partition (suboptimally partitioned LP 1) defined by

$$i_1 \in \arg \min_{i \in \mathcal{I} \setminus \{i^*\}} \mathrm{ReLU}(c_i) \frac{u_i l_i}{u_i - l_i}$$

is seen to provide the second-best upper bound, and similarly the third-best (suboptimally partitioned LP 2) defined by $i_2 \in \arg\min_{i \in \mathcal{I} \setminus \{i^*, i_1\}} \mathrm{ReLU}(c_i) \frac{u_i l_i}{u_i - l_i}$ gives the worst bound out of the three. Hence, the order of optimality indicated by Theorem 8 holds true, despite the result being derived in a worst-case setting.

This example shows that the suboptimally partitioned LP 2 (partitioned along the worst row $w_{i_2}^\top$) coincides with the unpartitioned LP at every nominal input. This suggests that none of the relaxation error is attributed to the i_2^{th} coordinate of the ReLU layer, and demonstrates the importance of using an intelligent and theoretically justified partitioning scheme, as developed in Theorem 8, instead of partitioning heuristically. It can be observed that the optimally partitioned LP is the only convex certification scheme in this experiment that is able to certify the robustness of this network, and it is able to do so at every input tested.

Finally, we add a 5-neuron layer to the network and re-run the experiment on the resulting two-layer network. The results are displayed in Figure 6b. The optimally partitioned

LP is again seen to yield the best convex upper bound, substantially reducing the relaxation error. However, we remark that the worst-case upper bound from (22) was derived for networks with one hidden layer, so it is violated by a handful of nominal inputs in the two-layer case, unlike the one-layer case where the bound is guaranteed to hold. Despite this, the efficacy of partitioning, and in particular the partitioning scheme of Theorem 8, is still seen to hold empirically for this multi-layer example. Lastly, we note that the computation times for the nonconvex, unpartitioned LP, and partitioned LP rise to 0.94, 0.68, and 1.48 seconds, respectively. The two-part partitioned LP maintains the polynomial-time complexity (with respect to the number of neurons) of linear programming, since it requires solving two instances of the same LP structure (Karmarkar, 1984). This time complexity is further explored in Section 5.3 below.

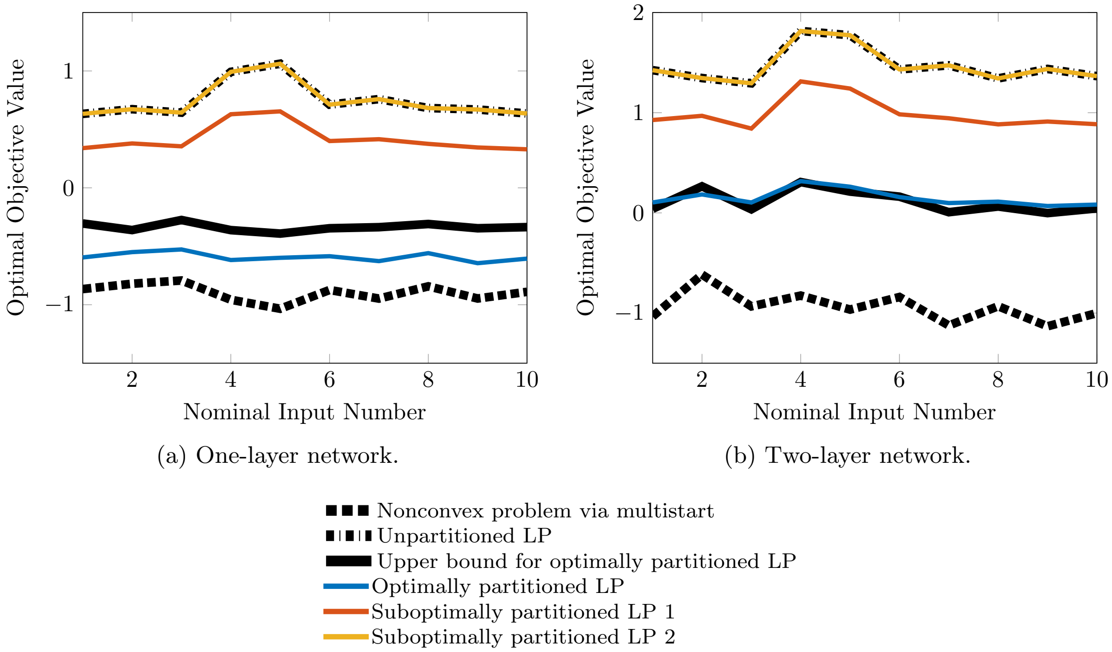

(a) One-layer network.

(b) Two-layer network.

Figure 6: Optimal values of robustness certification for ReLU Iris classifier. For the two-layer network, the optimal partitioning scheme is applied to only the ReLU constraints of the final layer.

5.2 Partitioned SDP Results

In this experiment, we compare the performance of the LP-based certificates against the SDP-based certificates on two classification networks with one hidden layer each. In each case, we test 10 nominal inputs $\bar{x}$ using an input uncertainty set $\mathcal{X} = \{x \in \mathbb{R}^{n_x} : \|x-\bar{x}\|_\infty \leq \epsilon\}$ with $\epsilon = 0.1$, similar to the previous experiment. The first network is the same Iris classifier used in Section 5.1, and the second network is a 5×5 network with randomly generated weights, each element being a standard normal random variable. For each network and nominal input, we approximate the nonconvex certification problem using `fmincon` with 5 multi-starts. We then solve the unpartitioned LP and partitioned LP according to Theorem 8. Finally, we solve the unpartitioned SDP and the partitioned SDP according to Theorem 19. The optimal objective values are displayed in Figure 7. As seen, the SDP consistently outperforms the LP-based approaches, and the partitioning does not have a noticeable effect on this small network.

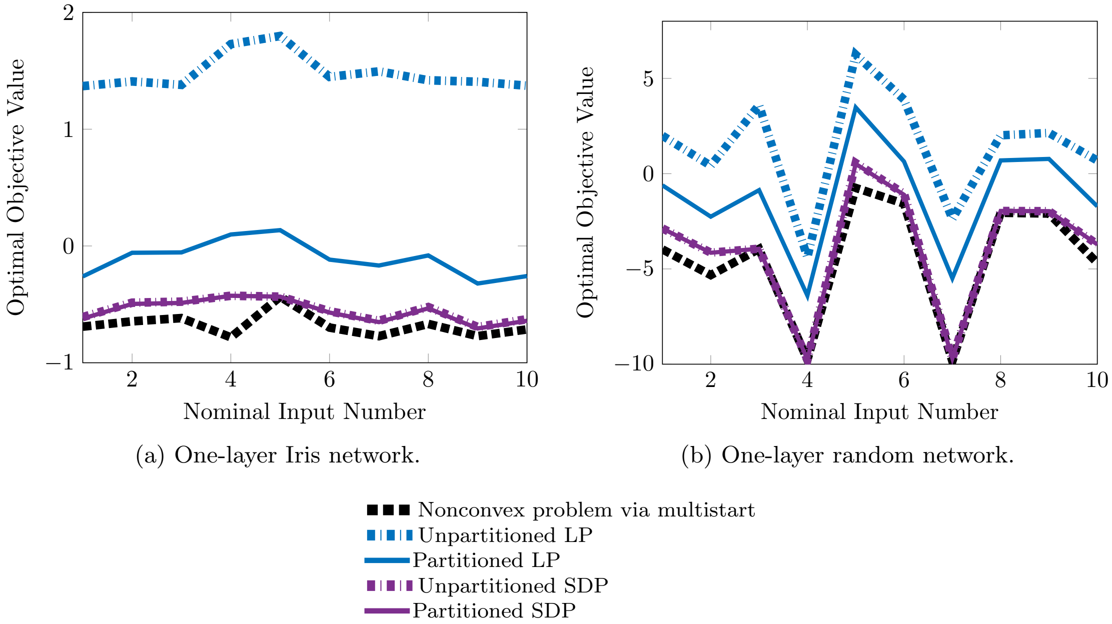

(a) One-layer Iris network.

(b) One-layer random network.

Figure 7: SDP relaxations compared to LP relaxations on one-layer networks with differently generated weights.

The observed SDP and partitioned SDP performance matches the assertion recently made in Zhang (2020), namely, that the SDP provides an exact relaxation (i.e., no relaxation error) of the one-layer ReLU robustness certification problem under only mild assumptions. In this case, if the unpartitioned SDP is exact, then certainly the partitioned SDP will be exact, meaning no further improvement may be made. This is a first indication that for relatively small networks with a single hidden layer, the unpartitioned SDP is sufficient for computing certificates with little to no relaxation error, without needing to resort to partitioning. However, as we will see in the next experiment, partitioning yields a substantial improvement in the certificate once the network becomes two or more layers, which is the case in many practical settings. We now move to this experiment, and propose a general rule of thumb for when LP, SDP, and their partitioned variants are best applied based on the depth and width of the network.

5.3 Effectiveness as Network Grows

In this section, we perform two experiments to test the effectiveness of partitioning as the size and structure of the network changes. First, we consider two-layer networks of structure $n_x \times 100 \times 5$, where n_x is the input dimension. For each input size $n_x \in \{5, 10, 20, 40, 80, 100\}$, we generate one network with standard normal random weights, and another network with uniformly distributed weights (where each element is distributed uniformly on the interval $[0,1]$). The weights are normalized according to Assumption 16. For each network being tested, we compute the LP, partitioned LP, SDP, and partitioned SDP relaxations at a fixed nominal input $\bar{x}$ using the input uncertainty set $\mathcal{X} = \{x \in \mathbb{R}^{n_x} : \|x - \bar{x}\|_\infty \leq \epsilon\}$ with $\epsilon = 0.5$. The optimal values, corresponding computation times, and percentage improvements induced by partitioning are reported in Table 1. The effectiveness of partitioning for the LP remains relatively constant between 5 and 10 percent improvement, whereas the partitioning appears to lose its efficacy on the SDP as the input size grows. As expected, the partitioned convex relaxations take twice as long to solve as their unpartitioned counterparts. Note that, despite the fact that partitioning works better for the LP with wide networks, the actual optimal value of the SDP-based certificates are always lower (tighter) than the LP-based ones. This matches what is known in the literature: the SDP is a tighter relaxation technique than the LP (Raghunathan et al., 2018). However, the computation times of the SDP and partitioned SDP quickly increase as the network size increases, whereas the LP and partitioned LP computation times are seen to slowly increase. All of this suggests the following: in the regime of shallow (i.e., one or two hidden layers) but very wide networks, the partitioned LP should be used, since the partitioning remains effective in tightening the relaxation, yet the method is scalable to large networks where the SDP cannot be feasibly applied.

In the second simulation of this section, we analyze the effectiveness of partitioning as the depth of the network increases. In particular, we consider networks with normal random weights having 5 inputs and 5 outputs, and each intermediate layer having 10 neurons. We run the experiment on networks having 1 through 6 such intermediate layers. Note that when the network has more than one hidden layer, an extra step is needed in order to apply the optimal LP partition from Theorem 8 since the number of rows n_1 of the first layer's weight matrix $W^{[0]}$ (i.e., the rows being partitioned along) may not equal the dimension of the output space n_z. The extra step is to generate a surrogate "c"-vector of size $n_1 \times 1$ so that Theorem 8 can be applied using this surrogate cost vector. There are a few ways of doing this. One such method is to treat the activation at the first hidden layer, $x^{[1]}$, as the output and determine which coordinate $i \in \mathbb{R}^{n_1}$ of the nominal activation $\bar{x}^{[1]} = \mathrm{ReLU}(W^{[0]}\bar{x})$ is maximal. This means that i would be the class assigned to $\bar{x}$ if the output were after the first hidden layer. Then, we find the second-best coordinate $j \neq i$ so that $\bar{x}^{[1]}_i \geq \bar{x}^{[1]}_j \geq \bar{x}^{[1]}_k$ for all other k. Afterwards, the surrogate vector c can be taken as $c = e_j - e_i$, meaning that it serves as a measure of whether the classification after the first hidden layer changes from i to j. In the case of our experiment, we choose the above $j \neq i$ randomly for simplicity. Of course, the surrogate vector c is only used to compute the partitioning coordinate i^* in Theorem 8, and the full network and original cost vector c are used in the resultant partitioned LP.

Unlike the partitioned LP, the SDP partitioning scheme given in Theorem 19 can directly be applied to deep networks, without the need to use the intermediate steps to compute the partition. We compute the LP, partitioned LP, SDP, and partitioned SDP on the networks at hand and report the objective values and computation times in Table 2. In this simulation, we see a stark contrast to the results in Table 1. Specifically, the percentage improvement induced by partitioning the LP reduces quickly to nearly zero percent for networks with 3

Table 1: Varying input size n_x for $n_x \times 100 \times 5$ ReLU network. Optimal values and corresponding computation times reported. P-LP and P-SDP correspond to partitioned LP and partitioned SDP, respectively. %-LP and %-SDP represent the percentage tightening of the optimal values obtained from partitioning.

(a) Normally distributed network weights.

Input size	LP	P-LP	%-LP	SDP	P-SDP	%-SDP
5	126.93	117.92	**7.10%**	16.82	14.83	**11.85%**
	0.71 s	1.46 s	104.18%	1.66 s	3.33 s	101.33%
10	187.57	176.19	**6.07%**	33.62	32.96	**1.98%**
	0.77 s	1.36 s	76.13%	1.54 s	3.16 s	105.57%
20	386.49	364.53	**5.68%**	54.02	54.01	**0.02%**
	0.71 s	1.42 s	100.49%	1.85 s	4.31 s	132.94%
40	874.70	864.56	**1.16%**	104.90	104.38	**0.49%**
	1.27 s	2.68 s	110.93%	4.79 s	9.33 s	95.01%
80	1591.41	1496.23	**5.98%**	310.37	310.31	**0.02%**
	1.76 s	2.97 s	69.00%	9.81 s	17.87 s	82.11%
100	2184.94	2175.87	**0.42%**	383.63	383.50	**0.03%**
	0.78 s	1.84 s	136.93%	5.02 s	10.52 s	109.46%

(b) Uniformly distributed network weights.

Input size	LP	P-LP	%-LP	SDP	P-SDP	%-SDP
5	11.65	10.69	**8.31%**	5.95	5.74	**3.44%**
	0.65 s	1.36 s	109.54%	1.39 s	2.20 s	58.32%
10	34.13	34.13	**0.00%**	12.61	11.92	**5.47%**
	0.68 s	1.36 s	101.35%	1.32 s	2.48 s	87.45%
20	83.74	83.02	**0.86%**	19.20	19.00	**1.06%**
	0.67 s	1.40 s	106.88%	1.31 s	2.85 s	118.39%
40	141.37	133.30	**5.71%**	25.89	25.69	**0.74%**
	0.69 s	1.43 s	106.67%	1.63 s	3.23 s	97.62%
80	260.80	242.19	**7.14%**	21.86	21.68	**0.84%**
	0.71 s	1.42 s	99.25%	2.82 s	5.44 s	92.97%
100	400.73	387.24	**3.37%**	102.87	102.35	**0.51%**
	0.74 s	1.56 s	111.10%	3.33 s	6.89 s	106.64%

or more intermediate 10-neuron hidden layers. Indeed, this is one fundamental drawback behind the LP relaxation: the convex upper envelope is used independently at every neuron, so the relaxation error quickly compounds as the network becomes deeper. On the other hand, the SDP relaxation takes into account the coupling between the layers of the network. This theoretical advantage is demonstrated empirically, as the percentage improvement gained by the partitioned SDP hovers around 10% even for the deep networks tested here.

Moreover, note how the SDP computation time remains relatively close to that of the LP, unlike the rapid increase in computation time seen when increasing the input size. This behavior suggests the following: in the regime of deep but relatively narrow networks, the partitioned SDP should be used, since the partitioning is effective in tightening the relaxation, yet the computational cost grows relatively slowly as more layers are added (compared to the case where more inputs are added).

Table 2: Varying number of hidden layers for a $5 \times 10 \times 10 \times \cdots \times 10 \times 5$ ReLU network with normal random weights. Optimal values and corresponding computation times reported. P-LP and P-SDP correspond to partitioned LP and partitioned SDP, respectively. %-LP and %-SDP represent the percentage tightening of the optimal values obtained from partitioning.

Layers	LP	P-LP	%-LP	SDP	P-SDP	%-SDP
1	10.16	7.03	**30.79**%	4.70	4.65	**1.12**%
	0.59 s	1.21 s	105.06%	0.68 s	1.29 s	91.17%
2	46.29	44.89	**3.03**%	2.42	1.94	**19.94**%
	0.62 s	1.21 s	93.07%	0.71 s	1.49 s	108.81%
3	626.96	626.96	**0.00**%	36.29	34.36	**5.31**%
	0.61 s	1.29 s	110.86%	0.72 s	1.47 s	103.32%
4	5229.32	5229.32	**0.00**%	179.79	167.34	**6.93**%
	0.65 s	1.29 s	97.47%	0.99 s	1.88 s	89.81%
5	37625.91	37625.86	**0.00**%	628.78	561.60	**10.68**%
	0.69 s	1.34 s	94.59%	1.13 s	2.04 s	80.35%
6	326743.55	326743.34	**0.00**%	3245.41	3050.69	**6.00**%
	0.75 s	1.35 s	79.44%	1.19 s	2.35 s	98.01%

6. Conclusions

In this paper, we propose intelligently designed partitioning schemes for linear programming (LP) and semidefinite programming (SDP) robustness certification methods of ReLU neural networks. The partitions are derived by minimizing the worst-case error induced by the corresponding convex relaxations, which is theoretically justified by showing that minimizing the true relaxation error is NP-hard. The proposed techniques are experimentally substantiated by demonstrating significant reduction in relaxation error on real and synthetic networks, with only doubling the computation time. Our experiments show that the LP and SDP partitioning schemes exhibit tradeoffs between different regimes, namely, as the input size and the number of layers are varied. The results conclude that both LP and SDP partitioning schemes yield a reduction in relaxation error on the order of 10%, with LP best applying to shallow but wide networks, and SDP best applying to deep but narrow networks. Consequently, partitioning proves to be a simple yet effective method for obtaining tighter robustness certificates for ReLU neural networks.

Acknowledgments

This work was supported by grants from AFOSR, ONR, and NSF.

Appendix A. Proof of Proposition 3

Proof Assume that $f^*(\mathcal{X}) > \max_{j\in\{1,2,\dots,p\}} \hat{f}^*(\mathcal{X}^{(j)})$. Then,

$$f^*(\mathcal{X}) > \hat{f}^*(\mathcal{X}^{(j)}) \text{ for all } j \in \{1,2,\dots,p\}. \tag{43}$$

Let (x^*, z^*) denote an optimal solution to the unrelaxed problem (2), i.e., $x^* \in \mathcal{X}$, $z^* = f(x^*)$, and

$$c^\top z^* = f^*(\mathcal{X}). \tag{44}$$

Since $\cup_{j=1}^p \mathcal{X}^{(j)} = \mathcal{X}$, there exists $j^* \in \{1,2,\dots,p\}$ such that $x^* \in \mathcal{X}^{(j^*)}$. Since $x^* \in \mathcal{X}^{(j^*)}$ and $z^* = f(x^*)$, it holds that $(x^*, z^*) \in \mathcal{N}^{(j^*)}$, where $\mathcal{N}^{(j^*)}$ is the relaxed network constraint set defined by $\mathcal{X}^{(j^*)}$. Therefore,

$$c^\top z^* \leq \sup\{c^\top z : x \in \mathcal{X}^{(j^*)},\ (x,z) \in \mathcal{N}^{(j^*)}\} = \hat{f}^*(\mathcal{X}^{(j^*)}) < f^*(\mathcal{X}),$$

where the first inequality comes from the feasibility of (x^*, z^*) over the $j^{*\text{th}}$ subproblem and the final inequality is due to (43). This contradicts the optimality of (x^*, z^*) given in (44). Hence, (9) must hold. ■

Appendix B. Proof of Proposition 4

Proof Let $j \in \{1,2,\dots,p\}$. It will be shown that $\mathcal{N}^{(j)} \subseteq \mathcal{N}$. Let $(x,z) \in \mathcal{N}^{(j)}$. Define $u' = u^{(j)}$, $l' = l^{(j)}$, and

$$\begin{aligned} g(x) &= u \odot (Wx - l) \oslash (u - l), \\ g'(x) &= u' \odot (Wx - l') \oslash (u' - l'). \end{aligned}$$

Then, by letting $\Delta g(x) = g(x) - g'(x) = a \odot (Wx) + b$, where

$$\begin{aligned} a &= u \oslash (u - l) - u' \oslash (u' - l'), \\ b &= u' \odot l' \oslash (u' - l') - u \odot l \oslash (u - l), \end{aligned}$$

the following relations are derived for all $i \in \{1,2,\dots,n_z\}$:

$$\begin{aligned} g_i^* &:= \inf_{\{x: l' \leq Wx \leq u'\}} (\Delta g(x))_i \geq \inf_{\{\hat{z}: l' \leq \hat{z} \leq u'\}} (a \odot \hat{z} + b)_i \\ &= \inf_{\{\hat{z}_i: l'_i \leq \hat{z}_i \leq u'_i\}} (a_i \hat{z}_i + b_i) = \begin{cases} a_i l'_i + b_i & \text{if } a_i \geq 0, \\ a_i u'_i + b_i & \text{if } a_i < 0. \end{cases} \end{aligned}$$

In the case that $a_i \geq 0$, we have that

$$g_i^* \geq a_i l'_i + b_i = \left(\frac{u_i}{u_i - l_i} - \frac{u'_i}{u'_i - l'_i}\right) l'_i + \left(\frac{u'_i l'_i}{u'_i - l'_i} - \frac{u_i l_i}{u_i - l_i}\right) = \frac{u_i}{u_i - l_i}(l'_i - l_i) \geq 0,$$

where the final inequality comes from the fact that $u \geq 0$, $l' \geq l$, and $u > l$. On the other hand, if $a_i < 0$, it holds that

$$\begin{aligned} g_i^* \geq a_i u_i' + b_i &= \left(\frac{u_i}{u_i - l_i} - \frac{u_i'}{u_i' - l_i'} \right) u_i' + \left(\frac{u_i' l_i'}{u_i' - l_i'} - \frac{u_i l_i}{u_i - l_i} \right) \\ &= \frac{u_i}{u_i - l_i}(u_i' - l_i) - u_i' = \frac{u_i' - u_i}{u_i - l_i} l_i \geq 0, \end{aligned}$$

where the final inequality comes from the fact that $u' \leq u$, $l \leq 0$, and $u > l$. Therefore,

$$g^* = (g_1^*, g_2^*, \dots, g_{n_z}^*) \geq 0,$$

which implies that $\Delta g(x) = g(x) - g'(x) \geq 0$ for all x such that $l^{(j)} = l' \leq Wx \leq u' = u^{(j)}$. Hence, since $(x, z) \in \mathcal{N}^{(j)}$, it holds that $z \geq 0$, $z \geq Wx$, and

$$z \leq g'(x) \leq g(x) = u \odot (Wx - l) \oslash (u - l).$$

Therefore, we have that $(x, z) \in \mathcal{N}$.

Since $\mathcal{X}^{(j)} \subseteq \mathcal{X}$ (by definition) and $\mathcal{N}^{(j)} \subseteq \mathcal{N}$, it holds that the solution to the problem over the smaller feasible set gives a lower bound to the original solution: $\hat{f}^*(\mathcal{X}^{(j)}) \leq \hat{f}^*(\mathcal{X})$. Finally, since j was chosen arbitrarily, this implies the desired inequality (10). ■

Appendix C. Proof of Proposition 10

Proof We prove the result by reducing an arbitrary instance of the Min-$\mathcal{K}$-Union problem to an instance of the optimal partitioning problem (29). The proof is broken down into steps. In Step 1, we introduce the Min-$\mathcal{K}$-Union problem. We then construct a specific neural network based on the parameters of the Min-$\mathcal{K}$-Union problem in Step 2. In Step 3, we construct the solution to the partitioned LP relaxation for our neural network in the case that the partition is performed along all input coordinates. In Step 4, we construct the solution to the partitioned LP relaxation in the case that only a subset of the input coordinates are partitioned. Finally, in Step 5, we show that the solution to the Min-$\mathcal{K}$-Union problem can be constructed from the solution to the optimal partitioning problem, i.e., by finding the best subset of coordinates to partition along in the fourth step. As a consequence, we show that optimal partitioning is NP-hard.

Step 1: Arbitrary Min-$\mathcal{K}$-Union Problem. Suppose that we are given an arbitrary instance of the Min-$\mathcal{K}$-Union problem, i.e., a finite number of finite sets $\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_n$ and a positive integer $\mathcal{K} \leq n$. Since each set $\mathcal{S}_j$ is finite, the set $\bigcup_{j=1}^n \mathcal{S}_j$ is finite with cardinality $m := \left| \bigcup_{i=j}^n \mathcal{S}_j \right| \in \mathbb{N}$. Therefore, there exists a bijection between the elements of $\bigcup_{j=1}^n \mathcal{S}_j$ and the set $\{1, 2, \dots, m\}$. Hence, without loss of generality, we assume $\mathcal{S}_j \subseteq \mathbb{N}$ for all $j \in \{1, 2, \dots, n\}$ such that $\bigcup_{j=1}^n \mathcal{S}_j = \{1, 2, \dots, m\}$. In this Min-$\mathcal{K}$-Union problem, the objective is to find $\mathcal{K}$ sets $\mathcal{S}_{j_1}, \mathcal{S}_{j_2}, \dots, \mathcal{S}_{j_\mathcal{K}}$ among the collection of n given sets such that $\left| \bigcup_{i=1}^{\mathcal{K}} \mathcal{S}_{j_i} \right|$ is minimized over all choices of $\mathcal{K}$ sets. In what follows, we show that the solution to this problem can be computed by solving a particular instance of the optimal partitioning problem (29).

Step 2: Neural Network Construction. Consider a 3-layer ReLU network, where $x^{[0]}, x^{[1]} \in \mathbb{R}^n$ and $x^{[2]}, x^{[3]} \in \mathbb{R}^m$. Let the weight vector on the output be $c = \mathbf{1}_m$. Take the input uncertainty set to be $\mathcal{X} = [-1, 1]^n$. Let $W^{[0]} = I_n$ and $W^{[2]} = I_m$. In addition, construct the weight matrix on the first layer to be $W^{[1]} \in \mathbb{R}^{m \times n}$ such that

$$W^{[1]}_{ij} = \begin{cases} 1 & \text{if } i \in \mathcal{S}_j, \\ 0 & \text{otherwise.} \end{cases}$$

We remark that, since all entries of $c = \mathbf{1}_m$, $W^{[0]} = I_n$, $W^{[1]}$, and $W^{[2]} = I_m$ are nonnegative, the optimal value of the unrelaxed certification problem (26) is $f^*(\mathcal{X}) = \mathbf{1}_m^\top W^{[1]} \mathbf{1}_n$.

To finish defining the network and its associated LP relaxations, we must specify the preactivation bounds at each layer. Since all weights of the neural network are nonnegative, the largest preactivation at each layer is attained when the input is $x^{[0]} = \mathbf{1}_n$, the element-wise maximum vector in $\mathcal{X}$. The preactivations corresponding to this input are $\hat{z}^{[1]} = \mathbf{1}_n$, $\hat{z}^{[2]} = W^{[1]} \mathbf{1}_n$, and $\hat{z}^{[3]} = W^{[1]} \mathbf{1}_n$. Therefore, setting

$$\begin{aligned} u^{[1]} &= 2\mathbf{1}_n, \\ u^{[2]} &= \frac{3}{2} W^{[1]} \mathbf{1}_n, \\ u^{[3]} &= \frac{5}{4} W^{[1]} \mathbf{1}_n + \frac{1}{8} \mathbf{1}_m, \end{aligned}$$

we obtain valid preactivation upper bounds. Similarly, taking

$$l^{[k]} = -u^{[k]}$$

for all $k \in \{1, 2, 3\}$ defines valid preactivation lower bounds.

Step 3: Densely Partitioned LP Relaxation. With the network parameters defined, we now consider the first variant of our partitioned LP relaxation. In particular, we consider the relaxation where all coordinates of the first layer are partitioned. We denote by $\bar{f}(\mathcal{X})$ the optimal objective value of the problem

$$\begin{aligned} \text{maximize} \quad & c^\top x^{[3]} \\ \text{subject to} \quad & x^{[0]} \in \mathcal{X}, \\ & x^{[k+1]} \geq W^{[k]} x^{[k]}, && k \in \{0, 1, 2\}, \\ & x^{[k+1]} \geq 0, && k \in \{0, 1, 2\}, \\ & x^{[k+1]} \leq u^{[k+1]} \odot (W^{[k]} x^{[k]} - l^{[k+1]}) \oslash (u^{[k+1]} - l^{[k+1]}), && k \in \{0, 1, 2\}, \\ & x^{[1]} = \mathrm{ReLU}(W^{[0]} x^{[0]}). \end{aligned} \tag{45}$$

This problem serves as a baseline; this is the tightest LP relaxation of the certification problem among all those with partitioning along the input coordinates.

We denote by $\bar{x} = (\bar{x}^{[0]}, \bar{x}^{[1]}, \bar{x}^{[2]}, \bar{x}^{[3]})$ an optimal solution of (45). We will now show that

$$\bar{x}^{[3]} = \frac{5}{4} W^{[1]} \mathbf{1}_n + \frac{1}{16} \mathbf{1}_m.$$

To see this, note that since all weights of the network and optimization (45) are nonnegative, the optimal activations $\bar{x}$ will be as large as possible in all coordinates and at all layers. Therefore, since the input is constrained to $\mathcal{X} = [-1,1]^n$, the optimal input for (45) is $\bar{x}^{[0]} = \mathbf{1}_n$. Since the first ReLU constraint in (45) is exact, this implies that the optimal activation at the first layer is

$$\bar{x}^{[1]} = \mathrm{ReLU}(W^{[0]}\bar{x}^{[0]}) = \mathrm{ReLU}(\mathbf{1}_n) = \mathbf{1}_n.$$

Now, for the second layer, the activation attains its upper bound. Since $u^{[2]} = -l^{[2]} = \frac{3}{2}W^{[1]}\mathbf{1}_n$, this implies that

$$\begin{aligned}\bar{x}^{[2]} &= u^{[2]} \odot (W^{[1]}\bar{x}^{[1]} - l^{[2]}) \oslash (u^{[2]} - l^{[2]}) \\ &= u^{[2]} \odot (W^{[1]}\bar{x}^{[1]} + u^{[2]}) \oslash (2u^{[2]}) \\ &= \frac{1}{2}(W^{[1]}\bar{x}^{[1]} + u^{[2]}) \\ &= \frac{1}{2}\left(W^{[1]}\mathbf{1}_n + \frac{3}{2}W^{[1]}\mathbf{1}_n\right) \\ &= \frac{5}{4}W^{[1]}\mathbf{1}_n.\end{aligned}$$

Similarly, for the third layer, we find that the optimal activation attains its upper bound as well. Since $u^{[3]} = -l^{[3]} = \frac{5}{4}W^{[1]}\mathbf{1}_n + \frac{1}{8}\mathbf{1}_m$ and $W^{[2]} = I_m$ this gives that

$$\begin{aligned}\bar{x}^{[3]} &= u^{[3]} \odot (W^{[2]}\bar{x}^{[2]} - l^{[3]}) \oslash (u^{[3]} - l^{[3]}) \\ &= u^{[3]} \odot (\bar{x}^{[2]} + u^{[3]}) \oslash (2u^{[3]}) \\ &= \frac{1}{2}(\bar{x}^{[2]} + u^{[3]}) \\ &= \frac{1}{2}\left(\frac{5}{4}W^{[1]}\mathbf{1}_n + \frac{5}{4}W^{[1]}\mathbf{1}_n + \frac{1}{8}\mathbf{1}_m\right) \\ &= \frac{5}{4}W^{[1]}\mathbf{1}_n + \frac{1}{16}\mathbf{1}_m,\end{aligned}$$

as claimed in (3). It is easily verified that $\bar{x}$ as computed above satisfies all constraints of the problem (45).

Step 4: Sparsely Partitioned LP Relaxation. We now introduce the second variant of the partitioned LP relaxation. In particular, let $\mathcal{J}_p \subseteq \{1,2,\dots,n\}$ be an index set such that $|\mathcal{J}_p| = n_p = n - \mathcal{K}$. Denote the complement of $\mathcal{J}_p$ by $\mathcal{J}_p^c = \{1,2,\dots,n\} \setminus \mathcal{J}_p$. We consider the partitioned LP defined in (28), which partitions along each coordinate in the index set $\mathcal{J}_p$. The optimal value of this problem is denoted by $f^*_{\mathcal{J}_p}(\mathcal{X})$, and we denote an optimal solution by $\hat{x} = (\hat{x}^{[0]}, \hat{x}^{[1]}, \hat{x}^{[2]}, \hat{x}^{[3]})$. We will compute $\hat{x}$ in three steps.

Step 4.1: Upper Bounding the Solution. We start by upper bounding the final layer activation of the solution. In particular, we claim that the optimal solution $\hat{x}$ satisfies

$$\hat{x}^{[3]} \leq t := u^{[3]} - \frac{1}{16}\mathbf{1}_{\mathcal{I}^c}. \tag{46}$$

where $\mathcal{I} = \bigcup_{j \in \mathcal{J}_p^c} S_j \subseteq \{1,2,\dots,m\}$ and $\mathcal{I}^c = \{1,2,\dots,m\} \setminus \mathcal{I}$. Since $\bar{x}^{[3]} = u^{[3]} - \frac{1}{16}\mathbf{1}_m$, the bound (46) is equivalent to

$$\hat{x}_i^{[3]} \leq t_i = \begin{cases} u_i^{[3]} & \text{if } i \in \mathcal{I}, \\ \bar{x}_i^{[3]} & \text{if } i \in \mathcal{I}^c, \end{cases} \tag{47}$$

for all $i \in \{1,2,\dots,m\}$. We now prove the element-wise representation of the bound, (47).

First, by the feasibility of $\hat{x}$ and the definitions of $u^{[3]}, l^{[3]}$, it must hold for all $i \in \{1,2,\dots,m\}$ that

$$\hat{x}_i^{[3]} \leq \frac{u_i^{[3]}}{u_i^{[3]} - l_i^{[3]}}(w_i^{[2]\top}\hat{x}^{[2]} - l_i^{[3]}) = \frac{1}{2}(w_i^{[2]\top}\hat{x}^{[2]} + u_i^{[3]}),$$

and also that

$$\hat{x}_i^{[3]} \geq w_i^{[2]\top}\hat{x}^{[2]}.$$

Combining these inequalities, we find that $\hat{x}_i^{[3]} \leq \frac{1}{2}(\hat{x}_i^{[3]} + u_i^{[3]})$, or, equivalently, that

$$\hat{x}_i^{[3]} \leq u_i^{[3]}.$$

This bound holds for all $i \in \{1,2,\dots,m\}$, and therefore it also holds for $i \in \mathcal{I}$. This proves the first case in the bound (47).

We now prove the second case of the claimed upper bound. For this case, suppose $i \notin \mathcal{I}$. Then $i \notin S_j$ for all $j \in \mathcal{J}_p^c$, which implies that

$$W_{ij}^{[1]} = 0 \text{ for all } j \in \mathcal{J}_p^c,$$

by the definition of $W^{[1]}$. Therefore,

$$w_i^{[1]\top}\hat{x}^{[1]} = \sum_{j=1}^{n} W_{ij}^{[1]}\hat{x}_j^{[1]} = \sum_{j \in \mathcal{J}_p^c} W_{ij}^{[1]}\hat{x}_j^{[1]} + \sum_{j \in \mathcal{J}_p} W_{ij}^{[1]}\hat{x}_j^{[1]} = \sum_{j \in \mathcal{J}_p} W_{ij}^{[1]}\hat{x}_j^{[1]}.$$

Now, note that for $j \in \mathcal{J}_p$, the j^{th} coordinate of the input is being partitioned, and therefore the optimal solution must satisfy

$$\hat{x}_j^{[1]} = \text{ReLU}(w_j^{[0]\top}\hat{x}^{[0]}) = \text{ReLU}(e_j^\top \hat{x}^{[0]}) = \text{ReLU}(\hat{x}_j^{[0]}) \leq 1,$$

since $\hat{x}_j^{[0]} \in [-1,1]$. Therefore,

$$w_i^{[1]\top}\hat{x}^{[1]} \leq \sum_{j \in \mathcal{J}_p^c} W_{ij}^{[1]} \leq \sum_{j=1}^{n} W_{ij}^{[1]} = w_i^{[1]\top}\mathbf{1}_n.$$

It follows from the feasibility of $\hat{x}$ and the definitions of $u^{[2]}, l^{[2]}$ that

$$\begin{aligned} \hat{x}_i^{[2]} &\leq \frac{u_i^{[2]}}{u_i^{[2]} - l_i^{[2]}}(w_i^{[1]\top}\hat{x}^{[1]} - l_i^{[2]}) = \frac{1}{2}(w_i^{[1]\top}\hat{x}^{[1]} + u_i^{[2]}) \\ &\leq \frac{1}{2}\left(w_i^{[1]\top}\mathbf{1}_n + \frac{3}{2}w_i^{[1]\top}\mathbf{1}_n\right) = \frac{5}{4}w_i^{[1]\top}\mathbf{1}_n = \bar{x}_i^{[2]}, \end{aligned}$$

where $\bar{x}$ is the solution computed for the densely partitioned LP relaxation in Step 3. Therefore, we conclude that for all $i \notin \mathcal{I}$, it holds that

$$\hat{x}_i^{[3]} \leq \frac{u_i^{[3]}}{u_i^{[3]} - l_i^{[3]}}(w_i^{[2]\top}\hat{x}^{[2]} - l_i^{[3]}) \leq \frac{u_i^{[3]}}{u_i^{[3]} - l_i^{[3]}}(w_i^{[2]\top}\bar{x}^{[2]} - l_i^{[3]}) = \bar{x}_i^{[3]},$$

by our previous construction of $\bar{x}^{[3]}$. Thus, we have proven the second case in (47) holds. Hence, the claimed bound (46) holds.

Step 4.2: Feasibility of Upper Bound. Let us define $x = (x^{[0]}, x^{[1]}, x^{[2]}, x^{[3]})$, a point in $\mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^m$, by

$$x^{[0]} = \mathbf{1}_n, \quad x^{[1]} = \mathbf{1}_{\mathcal{J}_p} + \frac{5}{4}\mathbf{1}_{\mathcal{J}_p^c}, \quad x^{[2]} = u^{[3]} - \frac{1}{8}\mathbf{1}_{\mathcal{I}^c}, \quad x^{[3]} = u^{[3]} - \frac{1}{16}\mathbf{1}_{\mathcal{I}^c}.$$

Note that $x^{[3]}$ equals the upper bound $t = (t_1, t_2, \dots, t_m)$. We now show that x is feasible for (28).

First, the input uncertainty constraint is satisfied, since $x^{[0]} = \mathbf{1}_n \in \mathcal{X}$. Next, the relaxed ReLU constraints at the first layer are satisfied, since

$$x^{[1]} = \mathbf{1}_{\mathcal{J}_p} + \frac{5}{4}\mathbf{1}_{\mathcal{J}_p^c} \geq 0, \qquad \text{(Layer 1 lower bound.)}$$

$$x^{[1]} - W^{[0]}x^{[0]} = \frac{1}{4}\mathbf{1}_{\mathcal{J}_p^c} \geq 0, \qquad \text{(Layer 1 lower bound.)}$$

$$x^{[1]} - u^{[1]} \odot (W^{[0]}x^{[0]} - l^{[1]}) \oslash (u^{[1]} - l^{[1]}) = -\frac{1}{2}\mathbf{1}_{\mathcal{J}_p} - \frac{1}{4}\mathbf{1}_{\mathcal{J}_p^c} \leq 0. \quad \text{(Layer 1 upper bound.)}$$

The relaxed ReLU constraints are also satisfied in the second layer, since

$$x^{[2]} = \frac{5}{4}W^{[1]}\mathbf{1}_n + \frac{1}{8}\mathbf{1}_{\mathcal{I}} \geq 0, \qquad \text{(Layer 2 lower bound.)}$$

$$x^{[2]} - W^{[1]}x^{[1]} = \frac{1}{4}W^{[1]}\mathbf{1}_{\mathcal{J}_p} + \frac{1}{8}\mathbf{1}_{\mathcal{I}} \geq 0, \qquad \text{(Layer 2 lower bound.)}$$

$$x^{[2]} - u^{[2]} \odot (W^{[1]}x^{[1]} - l^{[2]}) \oslash (u^{[2]} - l^{[2]}) = \frac{1}{8}(\mathbf{1}_{\mathcal{I}} - W^{[1]}\mathbf{1}_{\mathcal{J}_p^c}) \leq 0. \qquad \text{(Layer 2 upper bound.)}$$

The final inequality above follows from the fact that either $(\mathbf{1}_{\mathcal{I}})_i = 0$ or $(\mathbf{1}_{\mathcal{I}})_i = 1$. For coordinates i such that $(\mathbf{1}_{\mathcal{I}})_i = 0$, the inequality obviously holds. For coordinates i such that $(\mathbf{1}_{\mathcal{I}})_i = 1$, we know that $i \in \mathcal{I}$, implying that $i \in \mathcal{S}_j$ for some $j \in \mathcal{J}_p^c$. This in turn implies that $W_{ij}^{[1]} = 1$ for some $j \in \mathcal{J}_p^c$, and therefore $w_i^{[1]\top}\mathbf{1}_{\mathcal{J}_p^c} = \sum_{j \in \mathcal{J}_p^c} W_{ij}^{[1]} \geq 1 = (\mathbf{1}_{\mathcal{I}})_i$.

Continuing to check feasibility of x, the relaxed ReLU constraints in the final layer are also satisfied:

$$x^{[3]} = \frac{5}{4}W^{[1]}\mathbf{1}_n + \frac{1}{16}\mathbf{1}_m + \frac{1}{16}\mathbf{1}_{\mathcal{I}} \geq 0, \qquad \text{(Layer 3 lower bound.)}$$

$$x^{[3]} - W^{[2]}x^{[2]} = \frac{1}{16}\mathbf{1}_{\mathcal{I}^c} \geq 0, \qquad \text{(Layer 3 lower bound.)}$$

$$x^{[3]} - u^{[3]} \odot (W^{[2]}x^{[2]} - l^{[3]}) \oslash (u^{[3]} - l^{[3]}) = 0 \leq 0. \qquad \text{(Layer 3 upper bound.)}$$

Hence, the relaxed ReLU constraints are satisfied at all layers. The only remaining constraints to verify are the exact ReLU constraints for the partitioned input indices $\mathcal{J}_p$. Indeed, for all $j \in \mathcal{J}_p$, we have that

$$x_j^{[1]} - \mathrm{ReLU}(W^{[0]}x^{[0]})_j = (\mathbf{1}_{\mathcal{J}_p})_j + \frac{5}{4}(\mathbf{1}_{\mathcal{J}_p^c})_j - \mathrm{ReLU}(\mathbf{1}_n)_j = 1 + 0 - 1 = 0.$$

Hence, the ReLU equality constraint is satisfied for all input coordinates in $\mathcal{J}_p$. Therefore, our proposed point x is feasible for (28).

Step 4.3: Solution to Sparsely Partitioned LP. As shown in the previous step, the proposed point $x = (x^{[0]}, x^{[1]}, x^{[2]}, x^{[3]})$ is feasible for (28). Recall from the upper bound (46) that our solution $\hat{x} = (\hat{x}^{[0]}, \hat{x}^{[1]}, \hat{x}^{[2]}, \hat{x}^{[3]})$ satisfies $\hat{x}^{[3]} \leq t$. The objective value of the feasible point x gives that

$$c^\top x^{[3]} = \sum_{i=1}^{m} x^{[3]} = \sum_{i=1}^{m} t_i \geq \sum_{i=1}^{m} \hat{x}_i^{[3]} = f^*_{\mathcal{J}_p}(\mathcal{X}).$$

Since $f^*_{\mathcal{J}_p}(\mathcal{X})$ is the maximum value of the objective for all feasible points, it must be that $c^\top x^{[3]} = f^*_{\mathcal{J}_p}(\mathcal{X})$. Hence, the point x is an optimal solution to (28). Therefore, we can write the final activation of our optimal solution $\hat{x}$ to (28) as

$$\hat{x}^{[3]} = x^{[3]} = t = u^{[3]} - \frac{1}{16}\mathbf{1}_{\mathcal{I}^c}. \tag{48}$$

Step 5: Min-$\mathcal{K}$-Union from Optimal Partition. With the solutions constructed in Steps 3 and 4, we compute the difference in the objective values between the two partitioned LP relaxations:

$$\begin{aligned} f^*_{\mathcal{J}_p}(\mathcal{X}) - \bar{f}(\mathcal{X}) &= c^\top \hat{x}^{[3]} - c^\top \bar{x}^{[3]} = c^\top \left(u^{[3]} - \frac{1}{16}\mathbf{1}_{\mathcal{I}^c} - \frac{5}{4}W^{[1]}\mathbf{1}_n - \frac{1}{16}\mathbf{1}_m \right) \\ &= c^\top \left(\frac{5}{4}W^{[1]}\mathbf{1}_n + \frac{1}{8}\mathbf{1}_m - \frac{1}{16}\mathbf{1}_{\mathcal{I}^c} - \frac{5}{4}W^{[1]}\mathbf{1}_n - \frac{1}{16}\mathbf{1}_m \right) = \frac{1}{16}c^\top(\mathbf{1}_m - \mathbf{1}_{\mathcal{I}^c}) \\ &= \frac{1}{16}c^\top \mathbf{1}_{\mathcal{I}} = \frac{1}{16}\sum_{i\in\mathcal{I}} 1 = \frac{1}{16}|\mathcal{I}| = \frac{1}{16}\left| \bigcup_{j\in\mathcal{J}_p^c} S_j \right|. \end{aligned}$$

Therefore,

$$\left| \bigcup_{j\in\mathcal{J}_p^c} \mathcal{S}_j \right| = 16\left(f^*_{\mathcal{J}_p}(\mathcal{X}) - \bar{f}(\mathcal{X}) \right), \tag{49}$$

which holds for all partition index sets $\mathcal{J}_p \subseteq \{1, 2, \dots, n\}$ such that $|\mathcal{J}_p| = n_p = n - \mathcal{K}$.

Now, let $\mathcal{J}_p^*$ be an optimal partition, i.e., a solution to (29) with our specified neural network parameters. Then, by (49), we have $\left|\bigcup_{j\in(\mathcal{J}_p^*)^c} \mathcal{S}_j\right| = 16\left(f^*_{\mathcal{J}_p^*}(\mathcal{X}) - \bar{f}(\mathcal{X})\right) \leq 16\left(f^*_{\mathcal{J}_p}(\mathcal{X}) - \bar{f}(\mathcal{X})\right) = \left|\bigcup_{j\in\mathcal{J}_p^c} \mathcal{S}_j\right|$ for all $\mathcal{J}_p$ with $|\mathcal{J}_p| = n_p$. Since this holds for all $\mathcal{J}_p^c$ with $|\mathcal{J}_p^c| = n - n_p = \mathcal{K}$, this shows that the set $(\mathcal{J}_p^*)^c$ is an optimal solution to the Min-$\mathcal{K}$-Union problem specified at the beginning of the proof. Now, suppose the optimal partitioning problem in (29) could be solved for $\mathcal{J}_p^*$ in polynomial time. Then the optimal solution $(\mathcal{J}_p^*)^c$ to the Min-$\mathcal{K}$-Union problem is also computable in polynomial time. Since this holds for an arbitrary instance of the Min-$\mathcal{K}$-Union problem, this implies that the Min-$\mathcal{K}$-Union problem is polynomially solvable in general, which is a contradiction. Therefore, the problem (29) is NP-hard in general. ■

Appendix D. Proof of Proposition 11

Proof Let $j \in \{1, 2, \dots, p\}$. From the definition of the partition, it holds that $\mathcal{X}^{(j)} \subseteq \mathcal{X}$. What remains to be shown is that $\mathcal{N}_{\text{SDP}}^{(j)} \subseteq \mathcal{N}_{\text{SDP}}$.

Let $P \in \mathcal{N}_{\text{SDP}}^{(j)}$. Define $u' = u^{(j)}$ and $l' = l^{(j)}$. Since $P \in \mathcal{N}_{\text{SDP}}^{(j)}$, it follows that

$$\begin{aligned}
P_z &\geq 0, \\
P_z &\geq W P_x, \\
\text{diag}(P_{zz}) &= \text{diag}(W P_{xz}), \\
\text{diag}(P_{xx}) &\leq (l' + u') \odot P_x - l' \odot u', \\
P_1 &= 1, \\
P &\succeq 0.
\end{aligned}$$

To show that $P \in \mathcal{N}_{\text{SDP}}$, we should show that the above expressions imply that $\text{diag}(P_{xx}) \leq (l + u) \odot P_x - l \odot u$. To do so, define $\Delta l_i \geq 0$ and $\Delta u_i \geq 0$ such that $l'_i = l_i + \Delta l_i$ and $u'_i = u_i - \Delta u_i$ for all $i \in \{1, 2, \dots, n_x\}$. Then we find that

$$\begin{aligned}
((l' + u') \odot P_x - l' \odot u')_i &= (l'_i + u'_i)(P_x)_i - l'_i u'_i \\
&= (l_i + u_i)(P_x)_i - l_i u_i + (\Delta l_i - \Delta u_i)(P_x)_i \\
&\quad - (\Delta l_i u_i - \Delta u_i l_i - \Delta u_i \Delta l_i) \\
&= ((l + u) \odot P_x - l \odot u)_i + (\Delta l_i - \Delta u_i)(P_x)_i \\
&\quad - (\Delta l_i u_i - \Delta u_i l_i - \Delta u_i \Delta l_i) \\
&= ((l + u) \odot P_x - l \odot u)_i + \Delta_i,
\end{aligned}$$

where $\Delta_i \coloneqq (\Delta l_i - \Delta u_i)(P_x)_i - (\Delta l_i u_i - \Delta u_i l_i - \Delta u_i \Delta l_i)$. Therefore, it suffices to prove that $\Delta_i \leq 0$ for all i. Since $-l_i u_i \geq -l'_i u'_i$ by definition, it holds that $\Delta l_i u_i - \Delta u_i l_i - \Delta u_i \Delta l_i \geq 0$. Thus, when $(\Delta l_i - \Delta u_i)(P_x)_i \leq 0$, it holds that $\Delta_i \leq 0$, as desired. On the other hand, suppose that $(\Delta l_i - \Delta u_i)(P_x)_i \geq 0$. Then we find two cases:

1. $(\Delta l_i - \Delta u_i) \geq 0$ and $(P_x)_i \geq 0$. In this case, the maximum value of $(\Delta l_i - \Delta u_i)(P_x)_i$ is $(\Delta l_i - \Delta u_i)u'_i$. Therefore, the maximum value of Δ_i is

$$\begin{aligned}
\Delta_i &= (\Delta l_i - \Delta u_i)u'_i - (\Delta l_i u_i - \Delta u_i l_i - \Delta u_i \Delta l_i) \\
&= \Delta l_i (u'_i - u_i) - \Delta u_i u'_i + \Delta u_i l_i + \Delta u_i \Delta l_i \\
&= \Delta l_i (-\Delta u_i) + \Delta u_i \Delta l_i - \Delta u_i u'_i + \Delta u_i l_i \\
&= -\Delta u_i u'_i + \Delta u_i l_i.
\end{aligned}$$

   Both of the two final terms are nonpositive, and therefore $\Delta_i \leq 0$.

2. $(\Delta l_i - \Delta u_i) \leq 0$ and $(P_x)_i \leq 0$. In this case, the maximum value of $(\Delta l_i - \Delta u_i)(P_x)_i$ is $(\Delta l_i - \Delta u_i)l'_i$. Therefore, the maximum value of Δ_i is

$$\begin{aligned}\Delta_i &= (\Delta l_i - \Delta u_i)l'_i - (\Delta l_i u_i - \Delta u_i l_i - \Delta u_i \Delta l_i) \\ &= -\Delta u_i \Delta l_i + \Delta u_i \Delta l_i + \Delta l_i l'_i - \Delta l_i u_i \\ &= \Delta l_i l'_i - \Delta l_i u_i.\end{aligned}$$

Both of the two final terms are nonpositive, and therefore $\Delta_i \leq 0$.

Hence, we find that $(l' + u') \odot P_x - l' \odot u' \leq (l + u) \odot P_x - l \odot u$ for all $P \in \mathcal{N}_{\mathrm{SDP}}^{(j)}$, proving that $P \in \mathcal{N}_{\mathrm{SDP}}$, and therefore $\mathcal{N}_{\mathrm{SDP}}^{(j)} \subseteq \mathcal{N}_{\mathrm{SDP}}$.

Since $\mathcal{X}^{(j)} \subseteq \mathcal{X}$ and $\mathcal{N}_{\mathrm{SDP}}^{(j)} \subseteq \mathcal{N}_{\mathrm{SDP}}$, it holds that the solution to the problem over the smaller feasible set lower bounds the original solution: $\hat{f}^*_{\mathrm{SDP}}(\mathcal{X}^{(j)}) \leq \hat{f}^*_{\mathrm{SDP}}(\mathcal{X})$. Finally, since j was chosen arbitrarily, this implies the desired inequality (31). ■

Appendix E. Proof of Lemma 17

Proof Let $i, j \in \{1, 2, \ldots, n\}$. Since P is positive semidefinite, the 2^{nd}-order principal minor $P_{ii}P_{jj} - P_{ij}^2$ is nonnegative, and therefore

$$|P_{ij}| \leq \sqrt{P_{ii}P_{jj}}. \tag{50}$$

Furthermore, by the basic inequality that $2ab \leq a^2 + b^2$ for all $a, b \in \mathbb{R}$, we have that $\sqrt{P_{ii}P_{jj}} \leq \frac{1}{2}(P_{ii} + P_{jj})$. Substituting this inequality into (50) gives the desired bound. ■

References

Brendon G. Anderson and Somayeh Sojoudi. Certifying neural network robustness to random input noise from samples. *arXiv preprint arXiv:2010.07532*, 2020a.

Brendon G. Anderson and Somayeh Sojoudi. Data-driven assessment of deep neural networks with random input uncertainty. *arXiv preprint arXiv:2010.01171*, 2020b.

Brendon G. Anderson, Ziye Ma, Jingqi Li, and Somayeh Sojoudi. Tightened convex relaxations for neural network robustness certification. In *Proceedings of the 59th IEEE Conference on Decision and Control*, 2020.

Mislav Balunovic, Maximilian Baader, Gagandeep Singh, Timon Gehr, and Martin Vechev. Certifying geometric robustness of neural networks. In *Advances in Neural Information Processing Systems*, pages 15287–15297, 2019.

Dimitris Bertsimas and Iain Dunning. Multistage robust mixed-integer optimization with adaptive partitions. *Operations Research*, 64(4):980–998, 2016.

Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Prasoon Goyal, Lawrence D. Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, Xin Zhang, Jake Zhao, and Karol Zieba. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.

Michael Everett, Golnaz Habibi, and Jonathan P. How. Robustness analysis of neural networks via efficient partitioning: Theory and applications in control systems. *arXiv preprint arXiv:2010.00540*, 2020.

Alhussein Fawzi, Seyed-Mohsen Moosavi-Dezfooli, and Pascal Frossard. Robustness of classifiers: from adversarial to random noise. In *Advances in Neural Information Processing Systems*, pages 1632–1640, 2016.

Mahyar Fazlyab, Manfred Morari, and George J. Pappas. Safety verification and robustness analysis of neural networks via quadratic constraints and semidefinite programming. *IEEE Transactions on Automatic Control*, 2020. doi: 10.1109/TAC.2020.3046193.

Ronald A. Fisher. The use of multiple measurements in taxonomic problems. *Annals of eugenics*, 7(2):179–188, 1936.

Jean-Yves Franceschi, Alhussein Fawzi, and Omar Fawzi. Robustness of classifiers to uniform ℓ_p and Gaussian noise. In Amos Storkey and Fernando Perez-Cruz, editors, *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, pages 1280–1288. PMLR, April 2018.

Dorit S. Hochbaum. *Approximating Covering and Packing Problems: Set Cover, Vertex Cover, Independent Set, and Related Problems*, pages 94–143. PWS Publishing Co., USA, 1996. ISBN 0534949681.

Ming Jin, Javad Lavaei, Somayeh Sojoudi, and Ross Baldick. Boundary defense against cyber threat for power system state estimation. *IEEE Transactions on Information Forensics and Security*, 16:1752–1767, 2020. doi: 10.1109/TIFS.2020.3043065.

Ming Jin, Heng Chang, Wenwu Zhu, and Somayeh Sojoudi. Power up! Robust graph convolutional network via graph powering. *35th AAAI Conference on Artificial Intelligence*, 2021. to appear.

Narendra Karmarkar. A new polynomial-time algorithm for linear programming. In *Proceedings of the sixteenth annual ACM symposium on Theory of computing*, pages 302–311, 1984.

Guy Katz, Clark Barrett, David L. Dill, Kyle Julian, and Mykel J. Kochenderfer. Reluplex: An efficient smt solver for verifying deep neural networks. In *International Conference on Computer Aided Verification*, pages 97–117. Springer, 2017.

Weicong Kong, Zhao Yang Dong, Youwei Jia, David J. Hill, Yan Xu, and Yuan Zhang. Short-term residential load forecasting based on LSTM recurrent neural network. *IEEE Transactions on Smart Grid*, 10(1):841–851, 2017.

Ziye Ma and Somayeh Sojoudi. Strengthened SDP verification of neural network robustness via non-convex cuts. *arXiv preprint arXiv:2010.08603*, 2020.

Guido F. Montufar, Razvan Pascanu, Kyunghyun Cho, and Yoshua Bengio. On the number of linear regions of deep neural networks. In *Advances in neural information processing systems*, pages 2924–2932, 2014.

K. Muralitharan, Rathinasamy Sakthivel, and R. Vishnuvarthan. Neural network based optimization approach for energy demand prediction in smart grid. *Neurocomputing*, 273:199–208, 2018.

Xiang Pan, Tianyu Zhao, and Minghua Chen. DeepOPF: Deep neural network for DC optimal power flow. In *2019 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*, pages 1–6. IEEE, 2019.

Aditi Raghunathan, Jacob Steinhardt, and Percy S. Liang. Semidefinite relaxations for certifying robustness to adversarial examples. In *Advances in Neural Information Processing Systems*, pages 10877–10887, 2018.

Vicenc Rubies Royo, Roberto Calandra, Dusan M. Stipanovic, and Claire Tomlin. Fast neural network verification via shadow prices. *arXiv preprint arXiv:1902.07247*, 2019.

Jiawei Su, Danilo Vasconcellos Vargas, and Kouichi Sakurai. One pixel attack for fooling deep neural networks. *IEEE Transactions on Evolutionary Computation*, 23(5):828–841, 2019.

Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *International Conference on Learning Representations*, 2014.

Lily Weng, Huan Zhang, Hongge Chen, Zhao Song, Cho-Jui Hsieh, Luca Daniel, Duane Boning, and Inderjit Dhillon. Towards fast computation of certified robustness for ReLU networks. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 5276–5285, Stockholmsmässan, Stockholm Sweden, July 2018. PMLR.

Lily Weng, Pin-Yu Chen, Lam Nguyen, Mark Squillante, Akhilan Boopathy, Ivan Oseledets, and Luca Daniel. Proven: Verifying robustness of neural networks with a probabilistic approach. In *International Conference on Machine Learning*, pages 6727–6736. PMLR, 2019.

Eric Wong and Zico Kolter. Provable defenses against adversarial examples via the convex outer adversarial polytope. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 5286–5295, Stockholmsmässan, Stockholm Sweden, July 2018. PMLR.

Bichen Wu, Forrest Iandola, Peter H. Jin, and Kurt Keutzer. Squeezedet: Unified, small, low power fully convolutional neural networks for real-time object detection for autonomous driving. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 129–137, 2017.

Weiming Xiang and Taylor T. Johnson. Reachability analysis and safety verification for neural network control systems. *arXiv preprint arXiv:1805.09944*, 2018.

Huan Zhang, Lily Weng, Pin-Yu Chen, Cho-Jui Hsieh, and Luca Daniel. Efficient neural network robustness certification with general activation functions. In *Advances in neural information processing systems*, pages 4939–4948, 2018.

Richard Zhang. On the tightness of semidefinite relaxations for certifying robustness to adversarial examples. *Advances in Neural Information Processing Systems*, 33, 2020.