

# Tight Robustness Certification Through the Convex Hull of $\ell_0$ Attacks

Yuval Shapira, Dana Drachler-Cohen

Technion - Israel Institute of Technology  
shapirayuval@campus.technion.ac.il, ddana@ee.technion.ac.il

## Abstract

Few-pixel attacks mislead a classifier by modifying a few pixels of an image. Their perturbation space is an  $\ell_0$ -ball, which is not convex, unlike  $\ell_p$ -balls for  $p \geq 1$ . However, existing local robustness verifiers typically scale by relying on linear bound propagation, which captures convex perturbation spaces. We show that the convex hull of an  $\ell_0$ -ball is the intersection of its bounding box and an asymmetrically scaled  $\ell_1$ -like polytope. The volumes of the convex hull and this polytope are nearly equal as the input dimension increases. We then show a linear bound propagation that precisely computes bounds over the convex hull and is significantly tighter than bound propagations over the bounding box or our  $\ell_1$ -like polytope. This bound propagation scales the state-of-the-art  $\ell_0$  verifier on its most challenging robustness benchmarks by 1.24x-7.07x, with a geometric mean of 3.16.

**Code** — <https://github.com/YuvShap/top-t-CoVerD>

## Introduction

Image network classifiers are a central part of many safety-critical systems in healthcare (Miller 2023), autonomous driving (Hawkins 2025), and automated visual inspection (Marshall 2023). However, neural networks are susceptible to adversarial example attacks (Szegedy et al. 2013). An adversarial attack injects a small perturbation, often confined to the  $\ell_p$ -ball of a correctly classified input, to mislead the classifier. Local robustness (Bastani et al. 2016) is the main safety property for showing the resilience of networks to adversarial attacks.

Many robustness verifiers analyze the local robustness of neural networks to perturbations in  $\ell_\infty$ -balls (Tjeng, Xiao, and Tedrake 2019; Zhang et al. 2018; Müller et al. 2021; Gehr et al. 2018; Katz et al. 2017),  $\ell_2$ -balls (Leino, Wang, and Fredrikson 2021; Huang et al. 2021), and  $\ell_1$ -balls (Behl et al. 2021; Wu and Zhang 2021). To scale, most incomplete verifiers rely on linear bound propagation that overapproximates the network’s computations with convex polytopes. Since  $\ell_p$ -balls are convex for  $p \geq 1$ , such analysis does not introduce overapproximation error in the perturbation space. However,  $\ell_0$ -balls are not convex. This raises the question: can linear bound propagation avoid introducing overapproximation error in the non-convex  $\ell_0$  perturbation space?

We mathematically characterize the convex hull of an  $\ell_0$ -ball around an input  $x$  (single- or multi-channel). We show

that the convex hull is the intersection of the bounding box of the  $\ell_0$ -ball and an asymmetrically scaled  $\ell_1$ -like polytope. We show that the relative excess volume of this polytope compared to the convex hull converges exponentially to zero, suggesting it may be a good overapproximation for linear bound propagation.

We then present a linear bound propagation that precisely computes the minimum and maximum of a linear function over an  $\ell_0$ -ball, which coincide with those over its convex hull. We show that the minimum and maximum of a linear function  $f$  over an  $\ell_0$ -ball of radius  $t$  depend on the sum of the  $t$  lowest (for the minimum) or highest (for the maximum) input entry contributions to  $f$ . This bound propagation generalizes prior bound propagations for few-pixel attacks (Chiang et al. 2020; Xu et al. 2020) to any box input domain and to multi-channel inputs. We also present a bound propagation for our  $\ell_1$ -like polytope and show that the minimum and maximum depend on the product of  $t$  and the lowest or highest input entry contribution to  $f$ . That is, this polytope is a looser overapproximation than the convex hull, although their volumes are very close.

We integrate our bound propagation in GPUPoly (Müller et al. 2021), which is repeatedly called by CoVerD (Shapira et al. 2024), the state-of-the-art complete (exact)  $\ell_0$  robustness verifier. We evaluate it over fully-connected and convolutional classifiers, for MNIST, Fashion-MNIST, and CIFAR-10. Our bound propagation boosts CoVerD on its most challenging robustness benchmarks by 1.24x-7.07x, with a geometric mean of 3.16.

In summary, our main contributions are:

- A characterization of the convex hull of  $\ell_0$  perturbations.
- A linear bound propagation that precisely computes the minimum and maximum of a linear function over an  $\ell_0$ -ball, which is significantly tighter than bound propagations over its bounding box or our  $\ell_1$ -like polytope.
- An integration of our bound propagation in GPUPoly, for boosting  $\ell_0$  robustness verification.

## Preliminaries

In this section, we provide our notation and setting.

An input  $x$  is in a bounded box domain  $\mathcal{D}_v = \prod_{i=1}^v [a_i, b_i]$  if  $x$  is single-channel, or  $\mathcal{D}_v = \prod_{i=1}^v (\prod_{j=1}^d [a_i^{(j)}, b_i^{(j)}])$  if  $x$  is multi-channel. For example, an RGB image  $x$  is in  $\mathcal{D}_v =$

$([0, 1]^3)^v$ . We denote an entry by  $x_i$  for  $i \in [v] = \{1, \dots, v\}$  and a channel by  $x_i^{(j)}$  for  $i \in [v]$  and  $j \in [d]$ .

We focus on the local robustness of a network classifier to few-pixel attacks (Croce et al. 2022). A classifier maps an input to a score vector over  $c$  labels:  $N : \mathbb{R}^v \rightarrow \mathbb{R}^c$ . The classification of an input is the label with the maximal score:  $\text{class}_N(x) = \text{argmax}(N(x))$ . An adversarial attack considers an input  $\bar{x} \in \mathcal{D}_v$  and computes a small perturbation  $r$  such that  $\bar{x} + r \in \mathcal{D}_v$  and the perturbed input is classified differently:  $\text{class}_N(\bar{x}) \neq \text{class}_N(\bar{x} + r)$ . In few-pixel attacks, the attacker is given a bound  $t \in [v]$  on the number of entries that can be perturbed, which is often very small, and computes  $r$  whose  $\ell_0$  norm is at most  $t$ :  $|\{i \in [v] \mid r_i \neq 0\}| \leq t$ .

We consider a generalized setting in which the attacker is limited to perturbing a subset of pixels,  $\mathcal{K} \subseteq [v]$ . Note that  $\mathcal{K}$  can be arbitrarily large, and in particular it can be  $[v]$ , as in our experiments. Since the attacker cannot perturb the pixels in  $[v] \setminus \mathcal{K}$ , given  $\bar{x} \in \mathcal{D}_v$ , the effective input space is  $\mathcal{D} = \prod_{i \in \mathcal{K}} [a_i, b_i]$  (and similarly for multi-channel inputs), i.e., pixels not in  $\mathcal{K}$  are treated as constants. Formally, the extension of  $y \in \mathcal{D}$  to  $y' \in \mathcal{D}_v$  sets  $y'_i = y_i$ , for  $i \in \mathcal{K}$ , and  $y'_i = \bar{x}_i$ , for  $i \in [v] \setminus \mathcal{K}$ . We abuse notation: for  $\bar{x} \in \mathcal{D}_v$ , we also write  $\bar{x}$  for its projection onto  $\mathcal{K}$ , and for  $y \in \mathcal{D}$ , we also write  $y$  for its extension to  $\mathcal{D}_v$ . The space of all perturbed inputs is the  $\ell_0$ -ball of  $\bar{x}$  with radius  $t$ :

$$\mathcal{B}_0^t(\bar{x}) = \{y \in \mathcal{D} \mid |\{i \in \mathcal{K} \mid y_i \neq \bar{x}_i\}| \leq t\}$$

A network  $N$  is locally robust in  $\mathcal{B}_0^t(\bar{x})$  if it classifies all inputs the same:  $\forall y \in \mathcal{B}_0^t(\bar{x}). \text{class}_N(\bar{x}) = \text{class}_N(y)$ . Without loss of generality and to simplify notations, we assume  $\mathcal{K} = [k]$ , for  $k \in [t, v]$  (though our formulation, theorems, and bound propagations hold for any  $\mathcal{K} \subseteq [v]$ ). In particular, we write  $\mathcal{D} = \prod_{i=1}^k [a_i, b_i]$ . Figure 1 (left) shows the  $\ell_0$ -balls of three single-channel inputs, for  $t = 2$ . The figure shows that the  $\ell_0$ -ball is the union of  $\binom{3}{2}$  planes. Generally, an  $\ell_0$ -ball is a union of  $\binom{k}{t} t \cdot d$ -dimensional flats.

Although our focus is on robustness of image classifiers, our formulation, theorems, and bound propagations apply to other classifiers. For example, to text classifiers, which are susceptible to word replacement attacks (Alzantot et al. 2018). Given an input sentence, the attacker replaces up to  $t$  words. Technically, if the words are independently embedded (e.g., with Word2Vec (Mikolov et al. 2013)), the problem setting is similar: given an input sentence  $\bar{x}$  with  $v$  words  $\bar{x}_1, \dots, \bar{x}_v$ , the input domain is  $\mathcal{D}_v = \prod_{i=1}^v (\prod_{j=1}^d [a_i^{(j)}, b_i^{(j)}])$ , where  $\prod_{j=1}^d [a_i^{(j)}, b_i^{(j)}]$  is the bounding box of the embedding of  $\bar{x}_i$  and its replacements.

## The Convex Hull of $\ell_0$ -Balls

In this section, we characterize the convex hull of an  $\ell_0$ -ball.

### The Convex Hull for Single-Channel Inputs

In this section, we focus on a single-channel input  $\bar{x}$ .

**Illustration** We illustrate by considering the  $\ell_0$ -ball of  $\bar{x} = (0, 0, 0)$  that lies in the center of  $\mathcal{D} = [-1, 1]^3$ , shown in Figure 1a. The figure illustrates that the  $\ell_0$ -ball is not

convex, unlike  $\ell_p$ -balls for  $p \geq 1$ . The corners  $E_{\mathcal{B}_0^t(\bar{x})}$  of the  $\ell_0$ -ball are all inputs  $y \in \mathbb{R}^k$  such that  $k - t$  of their entries are zeros, and the other entries  $y_i$  are one of the bounds  $-1$  or  $1$  (blue dots in Figure 1a). The convex hull of the  $\ell_0$ -ball is the smallest convex set that contains it. It is obtained by connecting the corners  $E_{\mathcal{B}_0^t(\bar{x})}$  (Figure 1b). We shortly provide a characterization of this convex hull. One may wonder whether it is crucial, or whether it suffices to overapproximate the  $\ell_0$ -ball by a containing convex set that can be analyzed by existing robustness verifiers. For example, most verifiers are designed for box neighborhoods. Thus, one may attempt verifying local robustness in an  $\ell_0$ -ball by overapproximating it with its bounding box, whose boundary contains the corners of the  $\ell_0$ -ball, and checking robustness in this box. However, this box equals the input space  $\mathcal{D}$ . For large  $k$ , the network is not locally robust in  $\mathcal{D}$  (e.g., for  $k = v$ , it is not robust since the network does not classify all inputs the same). Thus, this approach usually fails for large  $k$ . Instead, one may overapproximate the  $\ell_0$ -ball with the tightest convex  $\ell_p$ -ball, which is the  $\ell_1$ -ball:  $\{y \in \mathbb{R}^k \mid \|y\|_1 \leq t\}$  (Zass and Shashua 2006; Baninajjar, Rezine, and Aminifar 2024), and submit it to an  $\ell_1$ -ball robustness verifier (Behl et al. 2021). However, the  $\ell_1$ -ball has sharp corners whose  $\ell_\infty$  distance from  $\mathcal{D}$  can be very high (Figure 1c). This overapproximation error can fail the verification. Our first theorem shows that, in the case where  $\bar{x}$  is in the center of  $\mathcal{D}$ , the convex hull of the  $\ell_0$ -ball of  $\bar{x}$  is the intersection of  $\mathcal{D}$  and the  $\ell_1$ -ball. Figure 1c shows  $\mathcal{D}$  (in green) and the  $\ell_1$ -ball with radius  $t = 2$  (in blue).

**General case** We next present our characterization for the general case where  $\bar{x}$  may not be in the center of  $\mathcal{D}$ , e.g., the  $\bar{x}$  in Figures 1d and 1g (the black dot). In this case, the corners  $E_{\mathcal{B}_0^t(\bar{x})}$  (the blue dots) are inputs  $y \in \mathbb{R}^k$  such that  $k - t$  of their entries  $y_i$  are  $\bar{x}_i$ , and the other entries  $y_i$  are one of the bounds  $a_i$  or  $b_i$  (which may be equal to  $\bar{x}_i$ , e.g., Figure 1g). We note that, unlike the simpler case, it can happen that an input  $y$  is not a corner even if it has  $k - t$  entries  $y_i = \bar{x}_i$  and its other entries  $y_i$  are  $a_i$  or  $b_i$  (red dots and  $\bar{x}$  itself in Figure 1g). The convex hull is obtained by connecting the corners  $E_{\mathcal{B}_0^t(\bar{x})}$  (Figures 1e and 1h). We show that it equals the intersection of  $\mathcal{D}$  and an *asymmetrically scaled*  $\ell_1$ -ball around  $\bar{x}$ . To define this polytope, we first define the *asymmetrically scaled  $i$  distance*, for  $i \in [k]$ . It measures how far  $y_i$  is from  $\bar{x}_i$ : if  $y_i = \bar{x}_i$ , this distance is zero, otherwise if  $y_i$  is one of the bounds  $a_i$  or  $b_i$ , this distance is one, and otherwise this distance is proportional to the distance between  $\bar{x}_i$  and the bound that is closer to  $y_i$  (than to  $\bar{x}_i$ ). Formally, the asymmetrically scaled  $i$  distance from input  $y$  to  $\bar{x}$  is the difference of their  $i^{\text{th}}$  entries normalized by the distance of  $\bar{x}_i$  and the bound  $b_i$ , if  $y_i > \bar{x}_i$ , or  $a_i$ , if  $y_i < \bar{x}_i$ :

$$\delta_{\bar{x}, a_i, b_i}^i(y) = \frac{y_i - \bar{x}_i}{\mathbf{1}_{\{y_i > \bar{x}_i\}}(b_i - \bar{x}_i) + \mathbf{1}_{\{y_i < \bar{x}_i\}}(a_i - \bar{x}_i)} \quad (1)$$

We omit  $a_i$  and  $b_i$  to simplify notation. We define  $\delta_{\bar{x}}^i(y) = 0$ , if  $y_i = \bar{x}_i$  (when the numerator and denominator are 0) and  $\delta_{\bar{x}}^i(y) = \infty$ , if  $\bar{x}_i = b_i$  and  $y_i > \bar{x}_i$  or if  $\bar{x}_i = a_i$  and  $y_i < \bar{x}_i$  (when the denominator is 0). Thus,  $\delta_{\bar{x}}^i(y) \in [0, \infty]$ . If  $y \in \mathcal{D}$ , then  $\delta_{\bar{x}}^i(y) \in [0, 1]$ , since for  $y_i = \bar{x}_i$ ,  $\delta_{\bar{x}}^i(y) = 0$ , and for

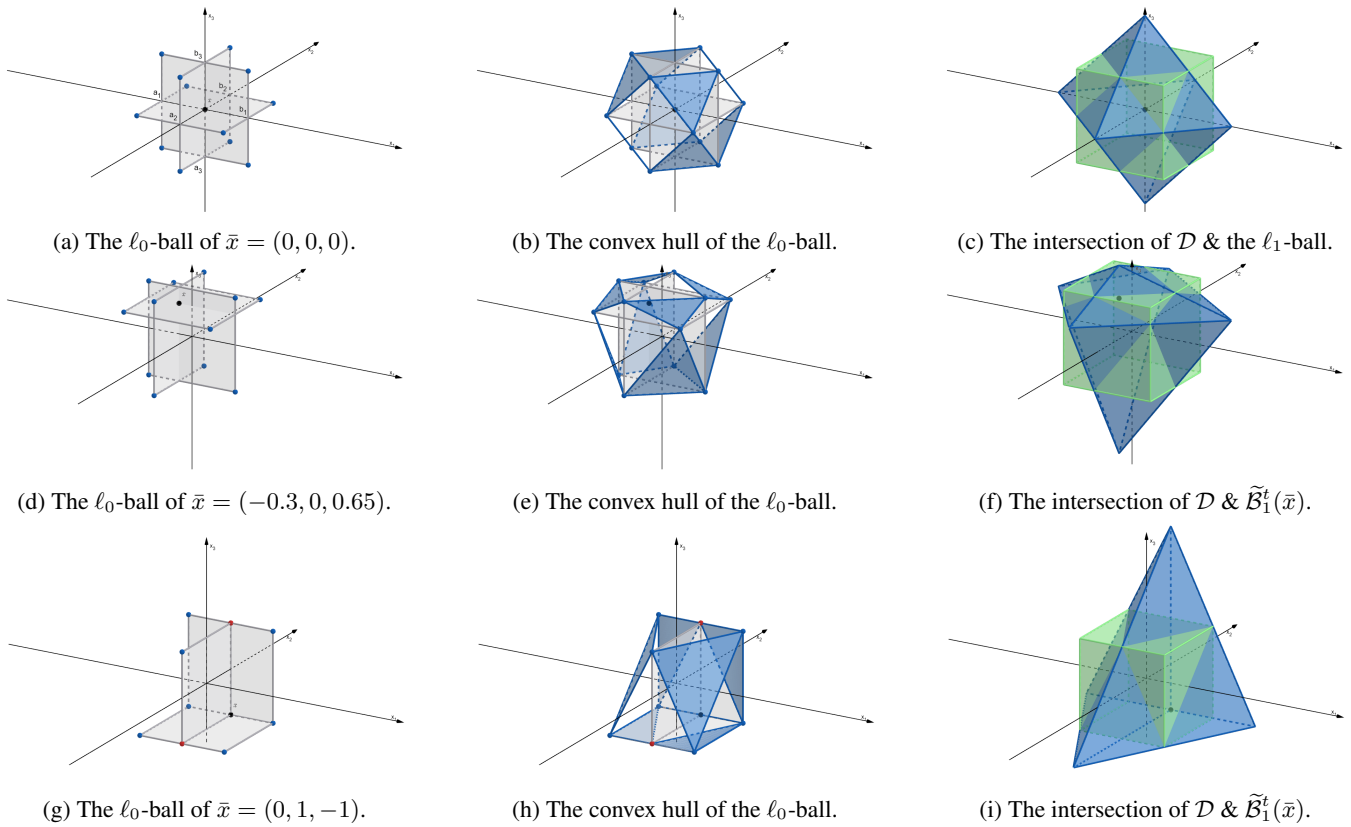


Figure 1: Illustration of the perturbation space, for  $\mathcal{D} = \prod_{i=1}^k [a_i, b_i] = [-1, 1]^3$  and  $t = 2$ . Left: The  $\ell_0$ -ball of an input  $\bar{x}$ ,  $\mathcal{B}_0^t(\bar{x})$ . Middle: Its convex hull. Right: The convex hull as the intersection of  $\mathcal{D}$  and an  $\ell_1$ -like polytope (Theorem 1),  $\tilde{\mathcal{B}}_1^t(\bar{x}) = \{y \in \mathbb{R}^k \mid \sum_{i=1}^k \delta_{\bar{x}}^i(y) \leq t\}$  (see Equation (1)). The plots show  $\bar{x}$  as a black dot and the corners of the  $\ell_0$ -ball as blue dots.

$y_i \in \{a_i, b_i\} \setminus \{\bar{x}_i\}$ ,  $\delta_{\bar{x}}^i(y) = 1$ . As an example, consider  $[a_i, b_i] = [-1, 1]$  and  $y_i = 0.3$ . For  $\bar{x}_i = 0$  (the center),  $\delta_{\bar{x}}^i(y) = 0.3$ . For  $\bar{x}_i = 0.7$ ,  $\delta_{\bar{x}}^i(y) = \frac{0.3-0.7}{0+1-(-1-0.7)} = 0.235$ .

The *asymmetrically scaled*  $\ell_1$ -ball  $\tilde{\mathcal{B}}_1^t(\bar{x})$  is the set of inputs whose sum over all asymmetrically scaled  $i$  distances from  $\bar{x}$  (Equation (1)) is at most  $t$ :

$$\tilde{\mathcal{B}}_1^t(\bar{x}) = \{y \in \mathbb{R}^k \mid \sum_{i=1}^k \delta_{\bar{x}}^i(y) \leq t\}$$

Note that  $\tilde{\mathcal{B}}_1^t(\bar{x})$  is defined over  $y \in \mathbb{R}^k$  (unlike  $\mathcal{B}_0^t(\bar{x})$ , defined over  $y \in \mathcal{D}$ ), because it generalizes the  $\ell_1$ -ball, defined over  $y \in \mathbb{R}^k$ . Figures 1f and 1i exemplify  $\tilde{\mathcal{B}}_1^t(\bar{x})$  (in blue).

**Theorem** We next state and prove our main theorem, which characterizes the convex hull of an  $\ell_0$ -ball. In the following, we denote by  $\text{Conv}(S)$  the convex hull of a set  $S$ .

**Theorem 1.**  $\text{Conv}(\mathcal{B}_0^t(\bar{x})) = \mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x})$ .

*Proof.*

•  $\text{Conv}(\mathcal{B}_0^t(\bar{x})) \subseteq \mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x})$ : By definition,  $\mathcal{B}_0^t(\bar{x}) \subseteq \mathcal{D}$  and  $\mathcal{B}_0^t(\bar{x}) \subseteq \tilde{\mathcal{B}}_1^t(\bar{x})$ . Hence,  $\mathcal{B}_0^t(\bar{x}) \subseteq \mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x})$ . Since  $\text{Conv}(\cdot)$  is monotone with respect to set inclusion,  $\text{Conv}(\mathcal{B}_0^t(\bar{x})) \subseteq$

$\text{Conv}(\mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x}))$ . The set  $\mathcal{D}$  is convex (as a box) and the set  $\tilde{\mathcal{B}}_1^t(\bar{x})$  is convex (by Lemma 1, appendix) and thus so is  $\mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x})$ . Thus,  $\text{Conv}(\mathcal{B}_0^t(\bar{x})) \subseteq \mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x})$ .

•  $\mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x}) \subseteq \text{Conv}(\mathcal{B}_0^t(\bar{x}))$ : The sets  $\mathcal{D}$  and  $\tilde{\mathcal{B}}_1^t(\bar{x})$  are convex and compact (as a box or by Lemma 1, appendix), and thus so is  $\mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x})$ . By the Krein-Milman Theorem (Rockafellar 1997), a compact convex set is the convex hull of its extreme points. Thus,  $\mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x}) = \text{Conv}(E_\cap)$ , where  $E_\cap$  is the set of extreme points of  $\mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x})$ . Since  $E_\cap \subseteq \mathcal{B}_0^t(\bar{x})$  (Lemma 2, appendix),  $\text{Conv}(E_\cap) \subseteq \text{Conv}(\mathcal{B}_0^t(\bar{x}))$ .  $\square$

**Volumes** We next study the volumes of the convex hull,  $\mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x})$ , and  $\tilde{\mathcal{B}}_1^t(\bar{x})$  (the volume of  $\mathcal{D}$  is  $\prod_{i=1}^k (b_i - a_i)$ ). We show that the relative excess volume of  $\tilde{\mathcal{B}}_1^t(\bar{x})$  compared to  $\mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x})$  converges exponentially to 0 as  $k$  increases. We begin by providing closed-form expressions for these volumes, which depend only on  $\mathcal{D}$  and  $t$  (not on the input  $\bar{x}$ ). In particular, this implies that the volumes of the convex hulls shown in Figures 1b, 1e and 1h are equal.

**Theorem 2.** 1.  $\text{vol}(\tilde{\mathcal{B}}_1^t(\bar{x})) = \text{vol}(\mathcal{D}) \frac{t^k}{k!}$ .  
2.  $\text{vol}(\mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x})) = \text{vol}(\mathcal{D}) \frac{t^k}{k!} \sum_{r=0}^{t-1} (-1)^r \binom{k}{r} (1 - \frac{r}{t})^k$ .

*Proof outline (proof is in the appendix).* We partition  $\tilde{\mathcal{B}}_1^t(\bar{x})$  and  $\mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x})$  into  $2^k$  parts, one for each orthant. Each part is a scaled version of  $\Delta_{k,t} = \{z \in [0, \infty)^k \mid \sum_{i=1}^k z_i \leq t\}$  for  $\tilde{\mathcal{B}}_1^t(\bar{x})$  and  $[0, 1]^k \cap \Delta_{k,t}$  for  $\mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x})$ . It is known that  $\text{vol}(\Delta_{k,t}) = \frac{t^k}{k!}$ . Thus, summing over the scaled volumes yields  $\text{vol}(\tilde{\mathcal{B}}_1^t(\bar{x})) = \text{vol}(\mathcal{D}) \frac{t^k}{k!}$  (proving bullet 1) and  $\text{vol}(\mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x})) = \text{vol}(\mathcal{D}) \cdot \text{vol}(\Delta_{k,t} \cap [0, 1]^k)$ . We show that  $\text{vol}(\Delta_{k,t} \cap [0, 1]^k)$  is equal to the cumulative distribution function (CDF) of the Irwin–Hall distribution (Irwin 1927; Hall 1927) evaluated at  $t$ , which is  $\frac{t^k}{k!} \sum_{r=0}^{t-1} (-1)^r \binom{k}{r} (1 - \frac{r}{t})^k$  (proving bullet 2). Technically, to compute the volume, we compute the probability that a uniformly sampled input from  $[0, 1]^k$  lies in  $\Delta_{k,t}$ , i.e., that the sum of its entries is at most  $t$ , and multiply it by  $\text{vol}([0, 1]^k) = 1$ . This probability is obtained by evaluating the Irwin-Hall CDF at  $t$ .  $\square$

**Excess volumes** We next compare the volume of the convex hull to  $\mathcal{D}$  and to  $\tilde{\mathcal{B}}_1^t(\bar{x})$ . To shorten notation, we write  $\text{vol}(\mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x})) = \text{vol}(\tilde{\mathcal{B}}_1^t(\bar{x})) \sum_{r=0}^{t-1} c_{r,k,t}$ , where  $c_{r,k,t} = (-1)^r \binom{k}{r} (1 - \frac{r}{t})^k$ , for  $r \in \{0, \dots, t-1\}$ . Note that  $c_{r,k,t} = 1$ , for  $r = 0$ , and  $c_{r,k,t}$  converges exponentially to 0 as  $k$  increases, for  $r > 0$ . The relative excess volume of  $\tilde{\mathcal{B}}_1^t(\bar{x})$ :

$$\frac{\text{vol}(\tilde{\mathcal{B}}_1^t(\bar{x}) \setminus (\mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x})))}{\text{vol}(\mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x}))} = \frac{-\sum_{r=1}^{t-1} c_{r,k,t}}{1 + \sum_{r=1}^{t-1} c_{r,k,t}}$$

This expression is asymptotically negligible because the numerator approaches 0 and the denominator approaches 1. The relative excess volume of  $\mathcal{D}$  is:

$$\frac{\text{vol}(\mathcal{D} \setminus (\mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x})))}{\text{vol}(\mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x}))} = \frac{\frac{k!}{t^k} - \sum_{r=0}^{t-1} c_{r,k,t}}{1 + \sum_{r=1}^{t-1} c_{r,k,t}}$$

This expression goes to infinity, since  $\frac{k!}{t^k}$  diverges with  $k$ . Figure 2 illustrates this behavior for different values of  $t$ . It shows that for very small values of  $k$ , the relative excess volume of  $\mathcal{D}$  is very small, but for  $k \geq 20$  it is very large (as context, for MNIST images  $k = 784$ ). In contrast, the relative excess volume of  $\tilde{\mathcal{B}}_1^t(\bar{x})$  converges to zero very quickly.

## The Convex Hull for Multi-Channel Inputs

In this section, we focus on a multi-channel input  $\bar{x}$ . In this setting, the attacker chooses  $t$  entries  $\bar{x}_i$  to perturb across all their channels (overall,  $t \cdot d$  values). To characterize the convex hull, we extend our definitions to multi-channel inputs. The *asymmetrically scaled*  $(i, j)$  distance from  $y$  to  $\bar{x}$  is:

$$\delta_{\bar{x}}^{i,j}(y) = \frac{y_i^{(j)} - \bar{x}_i^{(j)}}{\mathbf{1}_{y_i^{(j)} > \bar{x}_i^{(j)}} (b_i^{(j)} - \bar{x}_i^{(j)}) + \mathbf{1}_{y_i^{(j)} < \bar{x}_i^{(j)}} (a_i^{(j)} - \bar{x}_i^{(j)})}$$

This definition is identical to  $\delta_{\bar{x}}^i(y)$ , but over a channel  $y_i^{(j)}$ , for  $i \in [k], j \in [d]$ . We extend the asymmetrically scaled  $\ell_1$ -ball to the set of inputs whose sum, over all entries  $i$ , of the *maximal* asymmetrically scaled  $(i, j)$  distances across

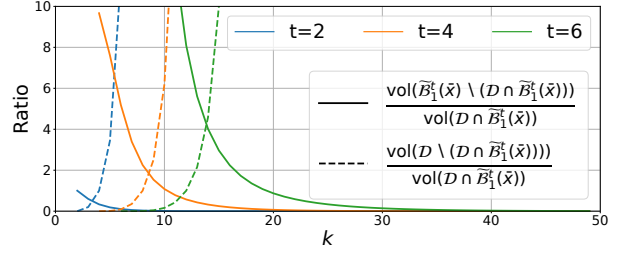


Figure 2: The relative excess volumes.

all channels  $j$  is at most  $t$ . Since the sum is over maximal values, we denote it  $\tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})$ :

$$\tilde{\mathcal{B}}_{1,\infty}^t(\bar{x}) = \{y \in (\mathbb{R}^d)^k \mid \sum_{i=1}^k \max_{j \in [d]} \delta_{\bar{x}}^{i,j}(y) \leq t\}$$

This definition is identical to  $\tilde{\mathcal{B}}_1^t(\bar{x})$ , except that it is over the maximal  $\delta_{\bar{x}}^{i,j}(y)$ . Like Theorem 1, we show that the convex hull of  $\mathcal{B}_0^t(\bar{x})$  is the intersection of  $\mathcal{D}$  and  $\tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})$ :

**Theorem 3.**  $\text{Conv}(\mathcal{B}_0^t(\bar{x})) = \mathcal{D} \cap \tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})$

The proof is identical to the proof of Theorem 1, except that it relies on Lemmas 3 and 4 (provided in the appendix), which extend Lemmas 1 and 2 to multi-channel inputs. The appendix also shows Theorem 4, which extends Theorem 2 to multi-channel inputs and implies a similar trend for the relative excess volumes.

## Bound Propagations

In this section, we introduce linear bound propagation methods for local robustness verification over an  $\ell_0$ -ball. We begin with background, including bound propagation over a box domain. Then, we present an equivalent formulation for the box domain, which we build on in our subsequent bound propagations. Next, we present bound propagations over an  $\ell_0$ -ball and over an asymmetrically scaled  $\ell_1$ -ball. Lastly, we extend them to multi-channel inputs and compare them.

### Background

Bound propagation is a central component in many neural network verifiers (e.g., Zhang et al. (2018); Müller et al. (2021); Singh et al. (2019)). This technique bounds every neuron  $n$  within a real-valued interval  $[l, u] \in \mathbb{R}^2$ . These bounds allow the verifier to efficiently check properties (i.e., constraints over the network’s inputs and outputs) as well as to tighten linear relaxations of the network’s computations, which reduce the verification time.

To obtain a lower bound  $l$  for a neuron  $n$ , a verifier computes an affine function that lower bounds the computation of  $n$  and computes its minimum. To shorten notations, we assume that the affine function is linear. Formally, for single-channel inputs, a verifier computes a linear function over the network’s inputs  $y \mapsto \sum_{i=1}^k w_i y_i$ , for  $w_i \in \mathbb{R}$ , lower bounding  $n$  in the property’s input subspace  $\mathcal{I} \subseteq \mathbb{R}^k$ :

$\forall y \in \mathcal{I}. n(y) \geq \sum_{i=1}^k w_i y_i$ . It then computes the minimum:  $l = \min_{y \in \mathcal{I}} \sum_{i=1}^k w_i y_i$ , based on the shape of  $\mathcal{I}$ . Commonly,  $\mathcal{I}$  is a box domain (shortly described). Similarly, an upper bound  $u$  of  $n$  is obtained by computing a linear function that upper bounds the computation of  $n$  and computing its maximum. The linear functions are computed by linear relaxation to the nonlinear activation functions (e.g., ReLU). Our bound propagation is orthogonal to the linear relaxation computation. In particular, it is applicable to relaxations computed by a backward pass (e.g., Zhang et al. (2018); Singh et al. (2019)), a forward pass (e.g., Wang et al. (2018a,b)), or a hybrid pass (e.g., Kouvaros et al. (2025)).

**Bound propagation for a box domain** Consider a neuron  $n$  whose computation is lower bounded by  $y \mapsto \sum_{i=1}^k w_i y_i$  for  $y \in \mathcal{I} = \mathcal{D} = \prod_{i=1}^k [a_i, b_i]$ . In the box domain, it is easy to define an input  $y$  obtaining the minimum: every entry  $y_i$  is on the boundaries,  $a_i$  or  $b_i$ , depending on the sign of  $w_i$ . Formally,  $y_i = a_i$  if  $w_i \geq 0$ , or  $y_i = b_i$  otherwise.

### Equivalent Bound Propagation for a Box Domain

We begin with an equivalent formulation of bound propagation over the box domain, to allow a unified formulation for all bound propagations and simplify their extension to multi-channel inputs (shown later). Our formulations are defined over the *minimum and maximum input entry contributions*.

Given  $\bar{x} \in \mathcal{D} = \prod_{i=1}^k [a_i, b_i]$  and  $y \mapsto \sum_{i=1}^k w_i y_i$ , we write:  $\sum_{i=1}^k w_i y_i = \sum_{i=1}^k w_i \bar{x}_i + \sum_{i=1}^k w_i (y_i - \bar{x}_i)$ . Note that the minimum and maximum of this linear function over  $y \in \mathcal{I}$  depend only on the second sum (the first sum is fixed). We define the *minimum input entry contribution* of  $i \in [k]$ :

$$d_i^- = \min(w_i(b_i - \bar{x}_i), w_i(a_i - \bar{x}_i)),$$

which lower bounds the term  $w_i \cdot (y_i - \bar{x}_i)$ , for  $y_i \in [a_i, b_i]$ . We define the *maximum input entry contribution* of  $i \in [k]$ :

$$d_i^+ = \max(w_i(b_i - \bar{x}_i), w_i(a_i - \bar{x}_i)),$$

which upper bounds  $w_i \cdot (y_i - \bar{x}_i)$ , for  $y_i \in [a_i, b_i]$ . Note that  $d_i^- \leq 0$  and  $d_i^+ \geq 0$ .

Given this notation, if  $n$ 's function is lower bounded by  $y \mapsto \sum_{i=1}^k w_i y_i$  for  $y \in \mathcal{I} = \mathcal{D}$ , the lower bound depends on the sum of all  $d_i^-$ :

$$l = \min_{y \in \mathcal{D}} \sum_{i=1}^k w_i y_i = \sum_{i=1}^k w_i \bar{x}_i + \sum_{i=1}^k d_i^-$$

If  $n$ 's function is upper bounded by  $y \mapsto \sum_{i=1}^k w_i y_i$  for  $y \in \mathcal{D}$ , the upper bound depends on the sum of all  $d_i^+$ :

$$u = \max_{y \in \mathcal{D}} \sum_{i=1}^k w_i y_i = \sum_{i=1}^k w_i \bar{x}_i + \sum_{i=1}^k d_i^+$$

**Example** Figure 3 shows a network with inputs  $y_1, y_2, y_3$  and an output  $o_1$ . The weights are on the edges and the bias of  $o_1$  is above it. The neurons  $n_1, n_2, o_1$  compute the sum of their bias and the products of their inputs and respective

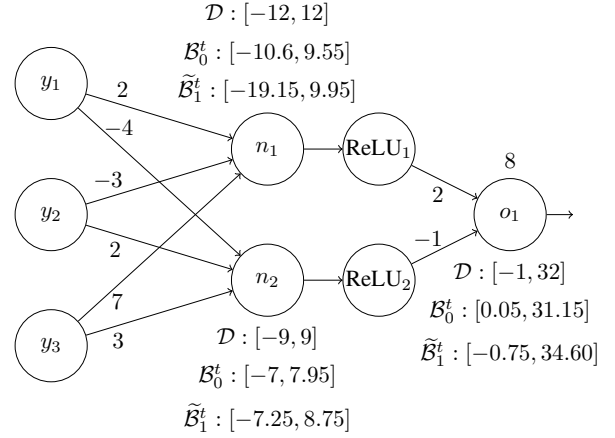


Figure 3: The three approaches for bound propagation over  $\mathcal{B}_0^t(\bar{x})$ , where  $\bar{x} = (-0.3, 0, 0.65)$ ,  $\mathcal{D} = [-1, 1]^3$ , and  $t = 2$ .

weights. The neurons  $\text{ReLU}_1, \text{ReLU}_2$  compute the activation function  $\text{ReLU}_i(n_i) = \max(0, n_i)$ . Our goal is to prove  $o_1 \geq 0$  for every input in the  $\ell_0$ -ball of  $\bar{x} = (-0.3, 0, 0.65)$ , for  $\mathcal{D} = [-1, 1]^3$  and  $t = 2$  (Figure 1d). The box bound propagation overapproximates the  $\ell_0$ -ball by its bounding box  $\mathcal{D}$  (although  $\mathcal{B}_t^0(\bar{x})$  is significantly smaller). That is,  $y_i \in [-1, 1]$  for  $i \in [3]$ . Then, a linear relaxation lower and upper bounds every neuron by linear functions. In this example, we use the linear relaxations of ReLU proposed by Zhang et al. (2018); Singh et al. (2019). Then, the box bound propagation computes the minimum and maximum of these linear functions, as described. This results in the intervals marked by  $\mathcal{D}$  in Figure 3. For example, the linear functions bounding  $n_1$  are exactly its function (since it is linear):  $2 \cdot y_1 - 3 \cdot y_2 + 7 \cdot y_3$ . Since  $\bar{x} = (-0.3, 0, 0.65) \in [-1, 1]^3$ , we get  $d_1^- = \min(2 \cdot (1 + 0.3), 2 \cdot (-1 + 0.3)) = -1.4$  and similarly  $d_2^- = -3, d_3^- = -11.55$ . Thus, its lower bound is:

$$\min_{(y_1, y_2, y_3) \in [-1, 1]^3} 2 \cdot y_1 - 3 \cdot y_2 + 7 \cdot y_3 = \underbrace{(2 \cdot (-0.3) - 3 \cdot 0 + 7 \cdot 0.65)}_{3.95} - 1.4 - 3 - 11.55 = -12$$

Since the box bound propagation computes  $o_1 \in [-1, 32]$ , it cannot show that the property  $o_1 \geq 0$  holds.

### Bound Propagation for an $\ell_0$ -Ball

We next show a bound propagation that precisely captures an  $\ell_0$ -ball input subspace. That is, it computes the minimum and maximum of a linear function over an  $\ell_0$ -ball, which coincide with those over its convex hull. We call it *top-t*, since it relies on the *top-t* input entry contributions. This idea has been proposed by Chiang et al. (2020); Xu et al. (2020), however they assume  $\mathcal{D} = [0, 1]^v$  and single-channel inputs.

As before, the goal is to compute the minimum of a neuron  $n$  whose function is lower bounded by  $y \mapsto \sum_{i=1}^k w_i y_i$ . Again, we write  $\sum_{i=1}^k w_i \bar{x}_i + \sum_{i=1}^k w_i (y_i - \bar{x}_i)$ , however at least  $k - t$  of the terms in the second sum are 0 since  $y \in \mathcal{B}_0^t(\bar{x})$ . Recall that  $w_i (y_i - \bar{x}_i) \geq d_i^-$ , for  $y_i \in [a_i, b_i]$ .

Since at most  $t$  entries are not 0, the second sum is lower bounded by the sum of the  $t$  lowest  $d_i^-$ . Thus, we sort them in a non-decreasing order,  $d_{l_1}^- \leq d_{l_2}^- \leq \dots \leq d_{l_k}^-$ . The minimum is obtained by  $y$  where  $y_i$  is: (1)  $a_i$  if  $i \in [t]$  and  $w_{l_i} \geq 0$ , (2)  $b_i$  if  $i \in [t]$  and  $w_{l_i} < 0$ , and (3)  $\bar{x}_i$  if  $i \notin [t]$ . The lower bound is:

$$l = \min_{y \in \mathcal{B}_0^t(\bar{x})} \sum_{i=1}^k w_i y_i = \left( \sum_{i=1}^k w_i \bar{x}_i \right) + (d_{l_1}^- + \dots + d_{l_t}^-)$$

Similarly, the upper bound depends on the  $t$  highest  $d_i^+$ :

$$u = \max_{y \in \mathcal{B}_0^t(\bar{x})} \sum_{i=1}^k w_i y_i = \left( \sum_{i=1}^k w_i \bar{x}_i \right) + (d_{u_1}^+ + \dots + d_{u_t}^+)$$

for the non-increasing series  $d_{u_1}^+ \geq d_{u_2}^+ \geq \dots \geq d_{u_k}^+$  of  $d_i^+$ .

**Example** The intervals marked by  $\mathcal{B}_0^t$  in Figure 3 show the top- $t$  bound propagation. For  $n_1$ , we showed  $d_1^- = -1.4$ ,  $d_2^- = -3$ ,  $d_3^- = -11.55$ . Thus, its  $t = 2$  lowest values are  $d_2^-$  and  $d_3^-$  and its lower bound is:  $\min_{y \in \mathcal{B}_0^t(\bar{x})} \sum_{i=1}^k w_i y_i = 3.95 - 3 - 11.55 = -10.6$ . In this example,  $o_1 \in [0.05, 31.15]$ , thus top- $t$  successfully proves  $o_1 \geq 0$ .

**Observation 1.** *The top- $t$  bound propagation computes the minimum and maximum over  $\mathcal{B}_0^t(\bar{x})$  and over its convex hull.*

This follows since the minimum and maximum of a linear function over a compact set (in our setting,  $\mathcal{B}_0^t(\bar{x})$ ) are equal to those over its convex hull. Namely, linear bound propagation over a non-convex perturbation space or over its convex hull are the same. Thus, the top- $t$  bound propagation over  $\mathcal{B}_0^t(\bar{x})$  coincides with the same propagation over  $\mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x})$ .

### Bound Propagation for $\tilde{\mathcal{B}}_1^t(\bar{x})$

We next define the  $t$ -times-top bound propagation over a  $\tilde{\mathcal{B}}_1^t(\bar{x})$  input subspace. The goal is to contrast it with the top- $t$  bound propagation, since Theorem 2 shows that the volumes of  $\tilde{\mathcal{B}}_1^t(\bar{x})$  and  $\mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x})$  are almost the same, for large  $k$ .

As before, we write  $\sum_{i=1}^k w_i \bar{x}_i + \sum_{i=1}^k w_i (y_i - \bar{x}_i)$ , for  $y \in \tilde{\mathcal{B}}_1^t(\bar{x})$ , and define its minimum. By definition of  $\delta_{\bar{x}}^i(y)$ :

$$w_i (y_i - \bar{x}_i) = \underbrace{w_i (\mathbf{1}_{\{y_i > \bar{x}_i\}} (b_i - \bar{x}_i) + \mathbf{1}_{\{y_i < \bar{x}_i\}} (a_i - \bar{x}_i))}_{\geq d_i^-} \delta_{\bar{x}}^i(y)$$

Thus,  $\sum_{i=1}^k w_i (y_i - \bar{x}_i) \geq \sum_{i=1}^k d_i^- \delta_{\bar{x}}^i(y)$ . Since  $\delta_{\bar{x}}^i(y) \geq 0$ ,  $\sum_{i=1}^k d_i^- \delta_{\bar{x}}^i(y) \geq d_{l_1}^- \sum_{i=1}^k \delta_{\bar{x}}^i(y)$ , for  $d_{l_1}^- = \min_{i \in [k]} d_i^-$ . Since  $d_{l_1}^- \leq 0$  and  $y \in \tilde{\mathcal{B}}_1^t(\bar{x})$ :  $d_{l_1}^- \sum_{i=1}^k \delta_{\bar{x}}^i(y) \geq d_{l_1}^- \cdot t$ . This minimum is obtained by  $y$  where  $y_i = \bar{x}_i$ , for  $i \neq 1$ , and  $y_{l_1} = \bar{x}_{l_1} + t \cdot (c - \bar{x}_{l_1})$ , for  $c = a_{l_1}$  if  $w_{l_1} \geq 0$ , and  $c = b_{l_1}$  otherwise. Thus, the lower bound depends on the product of  $t$  and the lowest  $d_i^-$  and is equal to:

$$l = \min_{y \in \tilde{\mathcal{B}}_1^t(\bar{x})} \sum_{i=1}^k w_i y_i = \left( \sum_{i=1}^k w_i \bar{x}_i \right) + t \cdot d_{l_1}^-$$

Similarly, the upper bound depends on the product of  $t$  and the highest  $d_i^+$ :

$$u = \max_{y \in \tilde{\mathcal{B}}_1^t(\bar{x})} \sum_{i=1}^k w_i y_i = \left( \sum_{i=1}^k w_i \bar{x}_i \right) + t \cdot d_{u_1}^+$$

**Example** The intervals marked by  $\tilde{\mathcal{B}}_1^t$  in Figure 3 show the  $t$ -times-top bound propagation. For  $n_1$ , the minimum of  $d_1^- = -1.4$ ,  $d_2^- = -3$ ,  $d_3^- = -11.55$  is  $d_3^-$ . Thus, its lower bound is  $\min_{y \in \tilde{\mathcal{B}}_1^t(\bar{x})} \sum_{i=1}^k w_i y_i = 3.95 - 2 \cdot 11.55 = -19.15$ . This bound propagation cannot prove  $o_1 \geq 0$ .

### Multi-Channel Inputs

Our formulation of the bound propagations as functions of the input entry contributions  $d_i^-$  and  $d_i^+$  naturally extends to multi-channel inputs. In this setting, given  $i \in [k]$ ,  $d_i^-$  and  $d_i^+$  sum over all channels' contributions. Given  $\bar{x} \in \mathcal{D}$  and a linear function  $y \mapsto \sum_{i=1}^k \sum_{j=1}^d w_i^{(j)} y_i^{(j)}$ , we define  $d_i^- = \sum_{j=1}^d \min(w_i^{(j)} (b_i^{(j)} - \bar{x}_i^{(j)}), w_i^{(j)} (a_i^{(j)} - \bar{x}_i^{(j)}))$  and  $d_i^+ = \sum_{j=1}^d \max(w_i^{(j)} (b_i^{(j)} - \bar{x}_i^{(j)}), w_i^{(j)} (a_i^{(j)} - \bar{x}_i^{(j)}))$ . The lower and upper bounds of all bound propagations are the same as for single-channel inputs, except that the first sum is  $\sum_{i=1}^k \sum_{j=1}^d w_i^{(j)} \bar{x}_i^{(j)}$ .

### Comparison

Our formulation explains why a bound propagation over the  $\ell_0$ -ball is tighter than over our  $\ell_1$ -like polytope or  $\mathcal{D}$ . All bound propagations have a shared term (i.e.,  $\sum_{i=1}^k w_i \bar{x}_i$ , for single-channel inputs) and they differ in the term they add: for the lower bound, top- $t$  adds the sum of the  $t$  lowest minimum input entry contributions, while  $t$ -times-top adds the product of  $t$  and the lowest minimum contribution, and the box bound propagation adds the sum of all minimum contributions (and similarly for the upper bound). The  $t$ -times-top and box bound propagations are incomparable for  $t > 1$ , because  $t$  times the top-1 value can be smaller or greater than the sum of all values. For example, in Figure 3, the box bound propagation has a tighter lower bound of  $n_1$  than  $t$ -times-top, but a looser upper bound. The theoretical time complexity of our bound propagations is linear in  $k$ , for single-channel inputs, or  $k \cdot d$ , for multi-channel inputs, like the common box bound propagation. In practice, the bound propagations run on a GPU, which affects the complexity (described in the appendix).

### Evaluation

In this section, we evaluate the top- $t$  bound propagation. We empirically show that top- $t$  is more precise than the other bound propagations for  $\ell_0$ -balls and that it significantly boosts CoVerD, the state-of-the-art  $\ell_0$  robustness verifier.

**Setup** We extended GPUPoly (Müller et al. 2021), which CoVerD calls a large number of times (sometimes  $10^6$ ). GPUPoly verifies using box bound propagation with the linear relaxation of Singh et al. (2019). We extended it with

Dataset	Network	# CC	# R	$\text{top-}t\text{-GP}$
MNIST	$6 \times 200\text{-PGD}$	99	97	94
	ConvSmall	98	95	36
	ConvSmallPGD	94	93	57
	ConvMedPGD	99	94	43
	ConvBig	97	94	0
Fashion	ConvSmallPGD	83	75	0
	ConvMedPGD	86	83	0
CIFAR-10	ConvSmallPGD	70	61	0

Table 1: The number of verified  $\ell_0$ -balls for  $t = 1$ .

our bound propagations:  $\text{top-}t$  (called  $\text{top-}t\text{-GP}$ ) and  $t\text{-times-top}$  ( $t\text{-times-top-GP}$ ). The implementation is described in the appendix, which also shows that the running time of  $\text{top-}t\text{-GP}$  is similar to GPUPoly. We ran experiments on a dual AMD EPYC 7713 server with 2TB RAM and eight A100 GPUs. We consider the same image classifiers as CoVerD, for MNIST (LeCun et al. 1998), Fashion-MNIST (Xiao, Rasul, and Vollgraf 2017), and CIFAR-10 (Krizhevsky 2009). MNIST and Fashion-MNIST consist of single-channel images (grayscale) and CIFAR-10 consists of three-channel images (RGB). The network architectures are provided in ERAN (ETH SRI Lab 2019).

**Bound propagation only** GPUPoly cannot directly prove robustness for  $\ell_0$ -balls, since it runs box bound propagation and thus requires to overapproximate an  $\ell_0$ -ball with its bounding box. However, the bounding box is  $\mathcal{D}_v$  and networks are not robust in  $\mathcal{D}_v$  since not all inputs are classified the same. Although  $\text{top-}t\text{-GP}$  improves GPUPoly’s precision for  $\ell_0$ -balls, we next show that it is still typically not precise enough to prove robustness in an  $\ell_0$ -ball where all pixels can be perturbed (i.e.,  $\mathcal{K} = [v]$ ). We show that even for  $t = 1$  (the smallest  $\ell_0$ -ball), it often loses too much precision to prove robustness. In this experiment, for each network, we consider the first 100 images from its test set. For each image, we check whether the network classifies it correctly. If so, we run  $\text{top-}t\text{-GP}$  for its  $\ell_0$ -ball with  $t = 1$ . Table 1 shows the number of images that are correctly classified (# CC), the number of robust  $\ell_0$ -balls (# R), and the number of  $\ell_0$ -balls that  $\text{top-}t\text{-GP}$  verifies. We obtain the number of robust  $\ell_0$ -balls from Calzone (Shapira, Avneri, and Drachslor-Cohen 2023), a complete (exact) verifier for  $\ell_0$ -balls. The table shows that except for MNIST classifiers,  $\text{top-}t\text{-GP}$  cannot verify even a single  $\ell_0$ -ball. These results imply that  $\text{top-}t\text{-GP}$  cannot, on its own, prove robustness in  $\ell_0$ -balls with  $\mathcal{K} = [v]$ . However, we later show that it boosts CoVerD, which does prove robustness in  $\ell_0$ -balls with  $\mathcal{K} = [v]$ . CoVerD decomposes the verification task into multiple robustness verification tasks over  $\ell_0$ -balls defined over subsets of pixels ( $\mathcal{K} \subseteq [v]$ ) whose size is at most 200 (a more detailed description is provided later). The next experiment explains why  $\text{top-}t\text{-GP}$  boosts CoVerD.

**Bound propagation comparison** We next show that  $\text{top-}t\text{-GP}$  is more precise than the other bound propagations for verifying robustness in  $\ell_0$ -balls over subsets of pixels ( $|\mathcal{K}| < v$ ). We consider three networks, three respective images,  $k \in [200]$ , and  $t \in [6]$ . For each network, image, and  $k$ , we randomly sample 500 subsets  $S \subseteq [v]$  of size  $k$  and let each bound propagation verify robustness in the neighborhood  $\mathcal{B}_0^t(\bar{x})$  where  $\mathcal{K} = S$ , for each  $t$ . Figure 4 shows the success rate (the ratio of  $\mathcal{B}_0^t(\bar{x})$  proven robust out of all 500 analyzed) as a function of  $k$ , for each approach and each  $t$ . The results show that  $\text{top-}t\text{-GP}$ ’s success rate is never lower than GPUPoly (with the box bound propagation) and often improves it, especially for small  $t$  or large  $k$ . It is also significantly better than  $t\text{-times-top-GP}$  for  $t > 1$ , despite the negligible volume difference shown in Theorem 2. In fact, for  $t > 4$ , even the box bound propagation is better than  $t\text{-times-top-GP}$ , although the volume of  $\tilde{\mathcal{B}}_1^t(\bar{x})$  is significantly smaller than  $\mathcal{D}$  (for large  $k$ ). This shows that bound propagation highly depends on the shape of the property’s input subspace and not only on its volume.

**Boosting  $\ell_0$  certification** Lastly, we show that  $\text{top-}t\text{-GP}$  boosts the performance of CoVerD (Shapira et al. 2024), which verifies robustness in  $\ell_0$ -balls where all pixels can be perturbed ( $\mathcal{K} = [v]$ ). CoVerD is a sound and complete local robustness verifier for  $\ell_0$ -balls. It is the state-of-the-art, improving Calzone (Shapira, Avneri, and Drachslor-Cohen 2023), which in turn outperforms a MILP-based verification. A naive verifier would submit  $\binom{v}{t}$  box neighborhoods to a verifier (e.g., GPUPoly) and determine robustness if all are robust. However, verifying  $\binom{v}{t}$  neighborhoods is infeasible for  $t > 2$ . Instead, CoVerD defines sets  $S_1, \dots, S_R$  of size  $k$  that cover all subsets of  $t$  pixels over  $[v]$ . For each  $S_r$ , it submits the box neighborhood  $\{y \in \mathcal{D}_v \mid \forall i \notin S_r : y_i = \bar{x}_i\}$  to GPUPoly. If it is robust, CoVerD continues to  $S_{r+1}$ ; otherwise, it refines  $S_r$  to smaller subsets and submits their neighborhoods to GPUPoly. Generally, the larger the sets  $S_r$  whose neighborhoods are proven robust, the fewer calls to GPUPoly. With  $\text{top-}t\text{-GP}$ , CoVerD can submit  $\mathcal{B}_0^t(\bar{x})$  where  $\mathcal{K} = S_r$ . Also, since  $\text{top-}t\text{-GP}$  has higher success rate than GPUPoly for large  $k$  (Figure 4), CoVerD can verify significantly larger sets  $S_r$ , which boosts its performance.

We evaluate on the most challenging benchmarks of CoVerD: for each network, we focus on the largest evaluated  $t$ . The appendix includes experiments over less challenging benchmarks of CoVerD. We ran the latest version of CoVerD, with the original hyper-parameters, except that we ran with five MILP instances. We used eight GPUs, as in CoVerD’s experiments (the running time decreases roughly linearly with the number of GPUs). The timeout of each verifier is five hours for each  $\ell_0$ -ball. Since CoVerD relies on random choices, we repeat each experiment three times with different seeds to show that our speedup is independent of the choice of seed. Figure 5 reports the average execution time and standard deviation. Each plot corresponds to a network. A pair of bars corresponds to an  $\ell_0$ -ball where each bar shows the execution time in minutes (the y-axis), for CoVerD and ours (CoVerD +  $\text{top-}t\text{-GP}$ ). The horizontal line marks the timeout. Bars reaching this line in-

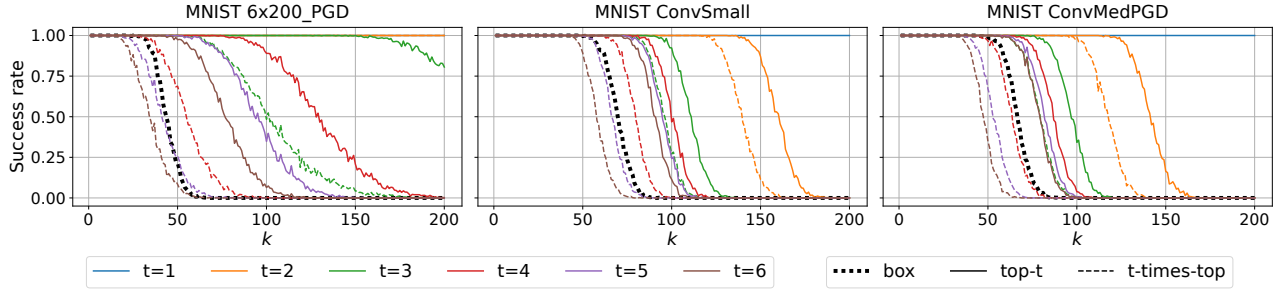


Figure 4: The success rate of the three bound propagations over  $\mathcal{B}_0^t(\bar{x})$  for different choices of  $\mathcal{K} \subseteq [v]$ , as a function of  $k = |\mathcal{K}|$ .

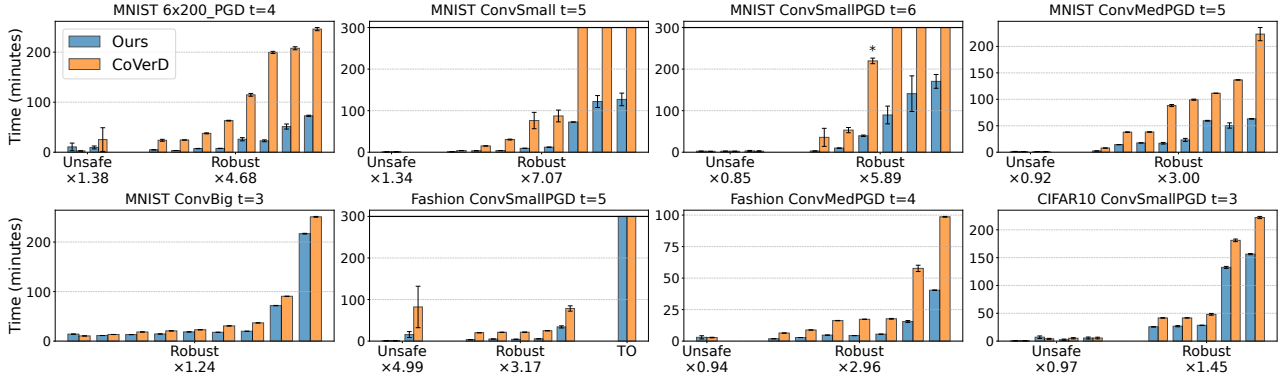


Figure 5: CoVerD +  $t_{\text{top-}t}\text{-GP}$  compared to CoVerD on its most challenging benchmarks.

dicating timeouts, in all three runs. For a single  $\ell_0$ -ball, one run of CoVerD out of three reached the timeout (denoted by \*) and is excluded from its average and standard deviation. The x-axis begins with the *unsafe*  $\ell_0$ -balls (non-robust), if they exist, followed by the *robust*  $\ell_0$ -balls. One plot also has a timeout instance (*TO*). Within each group, the  $\ell_0$ -balls are sorted by the execution time of CoVerD to better visualize the time differences. Below each group, we show the ratio between the sums of the average execution times of CoVerD and ours, excluding  $\ell_0$ -balls where CoVerD timed out in all three runs. The results show that  $t_{\text{top-}t}\text{-GP}$  boosts CoVerD’s performance on its robustness benchmarks by 1.24x-7.07x, with a geometric mean of 3.16. It is sometimes slower than CoVerD, but only when both verifiers terminate within minutes (as opposed to hours in the more challenging  $\ell_0$ -balls). In the appendix, we demonstrate that CoVerD +  $t\text{-times-top-GP}$  is substantially slower than our approach, consistent with the results in Figure 4.

## Related Work

Many verifiers analyze the local robustness of neural networks over convex perturbation spaces, e.g.,  $\ell_p$ -balls for  $p \geq 1$  (Tjeng, Xiao, and Tedrake 2019; Zhang et al. 2018; Müller et al. 2021; Gehr et al. 2018; Katz et al. 2017; Leino, Wang, and Fredrikson 2021; Huang et al. 2021; Behl et al. 2021; Wu and Zhang 2021), where many rely on bound propagation and convex relaxations to scale. Chiang et al. (2020); Xu et al. (2020) propose bound propagation for  $\ell_0$ -balls, which is similar to ours but is limited to single-channel

inputs over  $[0, 1]^v$ . Our experiments show that this bound propagation often fails to prove robustness even when only a single pixel can be perturbed (i.e.,  $t = 1$ ). Behl et al. (2021) propose a tight bound propagation for simplex input spaces. Though it can precisely capture our asymmetrically scaled  $\ell_1$ -ball, it currently does not support the intersection with a box, required to precisely capture the convex hull. Jia et al. (2019) verify robustness to word substitutions using interval bound propagation (IBP), where  $t = v$ . Huang et al. (2019) propose IBP for  $t < v$ . They overapproximate the convex hull of the  $\ell_0$ -ball with a simplex and say that manipulating their convex hull is infeasible (whereas we explicitly characterize ours). For a similar setting, Xu et al. (2020) verify using dynamic programming. Chiang et al. (2020); Levine and Feizi (2020) propose certified defenses against patch attacks, where the attacker perturbs a subset of pixels within a contiguous region. In contrast, we focus on few-pixel attacks, where the attacker can perturb any subset of pixels, and the allowed number of perturbed pixels is typically smaller than that in patch attacks.

## Conclusion

We show that the convex hull of an  $\ell_0$ -ball is the intersection of its bounding box and an  $\ell_1$ -like polytope. We present a bound propagation over an  $\ell_0$ -ball, tighter than bound propagations over its bounding box or the  $\ell_1$ -like polytope. It boosts the analysis of a neural network robustness verifier for  $\ell_0$ -balls by 1.24x-7.07x, with a geometric mean of 3.16, on its most challenging robustness benchmarks.

## References

- Alzantot, M.; Sharma, Y.; Elgohary, A.; Ho, B.-J.; Srivastava, M.; and Chang, K.-W. 2018. Generating Natural Language Adversarial Examples. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2890–2896. Association for Computational Linguistics.
- Baninajjar, A.; Rezine, A.; and Aminifar, A. 2024. VNN: Verification-Friendly Neural Networks with Hard Robustness Guarantees. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, 2846–2856. PMLR.
- Bastani, O.; Ioannou, Y.; Lampropoulos, L.; Vytiniotis, D.; Nori, A. V.; and Criminisi, A. 2016. Measuring Neural Net Robustness with Constraints. In *Advances in Neural Information Processing Systems 29 (NIPS)*, 2613–2621.
- Behl, H. S.; Kumar, M. P.; Torr, P.; and Dvijotham, K. 2021. Overcoming the Convex Barrier for Simplex Inputs. In *Advances in Neural Information Processing Systems*, volume 34, 4871–4882. Curran Associates, Inc.
- Chiang, P.; Ni, R.; Abdelkader, A.; Zhu, C.; Studer, C.; and Goldstein, T. 2020. Certified Defenses for Adversarial Patches. In *International Conference on Learning Representations*.
- Croce, F.; Andriushchenko, M.; Singh, N. D.; Flammarion, N.; and Hein, M. 2022. Sparse-RS: A Versatile Framework for Query-Efficient Sparse Black-Box Adversarial Attacks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(6): 6437–6445.
- ETH SRI Lab. 2019. ERAN: ETH Robustness Analyzer for Neural Networks. <https://github.com/eth-sri/eran>.
- Gehr, T.; Mirman, M.; Drachler-Cohen, D.; Tsankov, P.; Chaudhuri, S.; and Vechev, M. T. 2018. AI2: Safety and Robustness Certification of Neural Networks with Abstract Interpretation. In *IEEE Symposium on Security and Privacy, SP*.
- Gradshteyn, I. S.; and Ryzhik, I. M. 2014. *Table of Integrals, Series, and Products*. Academic Press, 8th edition. ISBN 978-0-12-384933-5. Edited by D. Zwillinger and V. Moll.
- Hall, P. 1927. The Distribution of Means for Samples of Size  $N$  Drawn from a Population in which the Variate Takes Values Between 0 and 1, All Such Values Being Equally Probable. *Biometrika*, 19(3/4): 240–245.
- Hawkins, A. J. 2025. Waymo is still good at avoiding serious distraction and death after 56.7 million miles. <https://www.theverge.com/news/658952/waymo-injury-prevention-human-benchmark-study>. Accessed: 2025-06-24.
- Huang, P.-S.; Stanforth, R.; Welbl, J.; Dyer, C.; Yogatama, D.; Goyal, S.; Dvijotham, K.; and Kohli, P. 2019. Achieving Verified Robustness to Symbol Substitutions via Interval Bound Propagation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 4083–4093. Association for Computational Linguistics.
- Huang, Y.; Zhang, H.; Shi, Y.; Kolter, J. Z.; and Anandkumar, A. 2021. Training Certifiably Robust Neural Networks with Efficient Local Lipschitz Bounds. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems NeurIPS*.
- Irwin, J. O. 1927. On the Frequency Distribution of the Means of Samples from a Population Having any Law of Frequency with Finite Moments, with Special Reference to Pearson’s Type II. *Biometrika*, 19(3/4): 225–239.
- Jia, R.; Raghunathan, A.; Göksel, K.; and Liang, P. 2019. Certified Robustness to Adversarial Word Substitutions. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 4129–4142. Association for Computational Linguistics.
- Katz, G.; Barrett, C. W.; Dill, D. L.; Julian, K.; and Kochenderfer, M. J. 2017. Reluplex: An Efficient SMT Solver for Verifying Deep Neural Networks. In *Computer Aided Verification - 29th International Conference, CAV*.
- Kouvaros, P.; Brückner, B.; Henriksen, P.; and Lomuscio, A. 2025. Dynamic Back-Substitution in Bound-Propagation-Based Neural Network Verification. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(26): 27383–27391.
- Krizhevsky, A. 2009. Learning multiple layers of features from tiny images. Technical report, University of Toronto.
- LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE*, 86(11): 2278–2324.
- Leino, K.; Wang, Z.; and Fredrikson, M. 2021. Globally-Robust Neural Networks. In *Proceedings of the 38th International Conference on Machine Learning, ICML*, volume 139 of *Proceedings of Machine Learning Research*. PMLR.
- Levine, A.; and Feizi, S. 2020. (De)Randomized Smoothing for Certifiable Defense against Patch Attacks. In *Advances in Neural Information Processing Systems*, volume 33, 6465–6475. Curran Associates, Inc.
- Marengo, J. E.; Farnsworth, D. L.; and Stefanic, L. 2017. A Geometric Derivation of the Irwin-Hall Distribution. *International Journal of Mathematics and Mathematical Sciences*, 2017(1): 3571419.
- Marshall, A. 2023. Amazon’s AI-Powered Van Inspections Give It a Powerful New Data Feed. <https://www.wired.com/story/amazons-ai-van-inspections-powerful-data-feed>. Describes deployment of AI-powered vehicle inspections (UV-eye AVI) across Amazon fleet centers.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed Representations of Words and Phrases and their Compositionality. In *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc.
- Miller, K. 2023. Could Self-Supervised Learning Be a Game-Changer for Medical Image Classification? <https://hai.stanford.edu/news/could-self-supervised-learning-be-game-changer-medical-image-classification>. Accessed: 2025-06-24.

- Müller, C.; Serre, F.; Singh, G.; Püschel, M.; and Vechev, M. T. 2021. Scaling Polyhedral Neural Network Verification on GPUs. In *Proceedings of Machine Learning and Systems MLSys*.
- Rockafellar, R. T. 1997. *Convex analysis*, volume 28. Princeton university press.
- Shapira, Y.; Avneri, E.; and Drachler-Cohen, D. 2023. Deep Learning Robustness Verification for Few-Pixel Attacks. *Proc. ACM Program. Lang.*, 7(OOPSLA1).
- Shapira, Y.; Wiesel, N.; Shabelman, S.; and Drachler-Cohen, D. 2024. Boosting Few-Pixel Robustness Verification via Covering Verification Designs. In *Computer Aided Verification - 36th International Conference, CAV 2024*, volume 14682 of *Lecture Notes in Computer Science*, 377–400. Springer.
- Singh, G.; Gehr, T.; Püschel, M.; and Vechev, M. T. 2019. An abstract domain for certifying neural networks. *Proc. ACM Program. Lang.*, 3(POPL): 41:1–41:30.
- Szegedy, C.; Zaremba, W.; Sutskever, I.; Bruna, J.; Erhan, D.; Goodfellow, I.; and Fergus, R. 2013. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*.
- Tjeng, V.; Xiao, K. Y.; and Tedrake, R. 2019. Evaluating Robustness of Neural Networks with Mixed Integer Programming. In *7th International Conference on Learning Representations, ICLR*.
- Wang, S.; Pei, K.; Whitehouse, J.; Yang, J.; and Jana, S. 2018a. Efficient Formal Safety Analysis of Neural Networks. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Wang, S.; Pei, K.; Whitehouse, J.; Yang, J.; and Jana, S. 2018b. Formal Security Analysis of Neural Networks using Symbolic Intervals. In *27th USENIX Security Symposium (USENIX Security 18)*, 1599–1614. USENIX Association. ISBN 978-1-939133-04-5.
- Wu, Y.; and Zhang, M. 2021. Tightening Robustness Verification of Convolutional Neural Networks with Fine-Grained Linear Approximation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(13): 11674–11681.
- Xiao, H.; Rasul, K.; and Vollgraf, R. 2017. Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms. *CoRR*, abs/1708.07747.
- Xu, K.; Shi, Z.; Zhang, H.; Wang, Y.; Chang, K.-W.; Huang, M.; Kailkhura, B.; Lin, X.; and Hsieh, C.-J. 2020. Automatic Perturbation Analysis for Scalable Certified Robustness and Beyond. In *Advances in Neural Information Processing Systems*, volume 33, 1129–1141. Curran Associates, Inc.
- Zass, R.; and Shashua, A. 2006. Nonnegative Sparse PCA. In Schölkopf, B.; Platt, J.; and Hoffman, T., eds., *Advances in Neural Information Processing Systems*, volume 19. MIT Press.
- Zhang, H.; Weng, T.; Chen, P.; Hsieh, C.; and Daniel, L. 2018. Efficient Neural Network Robustness Certification with General Activation Functions. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc.

## Proofs

**Lemma 1.**  $\tilde{\mathcal{B}}_1^t(\bar{x})$  is convex and compact.

*Proof. Convexity:* Let  $y, y' \in \tilde{\mathcal{B}}_1^t(\bar{x})$  and  $\lambda \in (0, 1)$ . We prove that for  $z = \lambda y + (1 - \lambda)y'$ , we have  $z \in \tilde{\mathcal{B}}_1^t(\bar{x})$ . We show it by proving:

$$\sum_{i=1}^k \delta_{\bar{x}}^i(z) \leq \lambda \sum_{i=1}^k \delta_{\bar{x}}^i(y) + (1 - \lambda) \sum_{i=1}^k \delta_{\bar{x}}^i(y') \leq t$$

To this end, we prove that  $\delta_{\bar{x}}^i(z) \leq \lambda \delta_{\bar{x}}^i(y) + (1 - \lambda) \delta_{\bar{x}}^i(y')$  for all  $i \in [k]$ . We split to cases. If  $z_i = \bar{x}_i$ , then  $\delta_{\bar{x}}^i(z) = 0$  and the claim holds. Otherwise,  $z_i \neq \bar{x}_i$ . If  $y_i = y'_i$ , then  $z_i = y_i = y'_i$  and so  $\delta_{\bar{x}}^i(z) = \lambda \delta_{\bar{x}}^i(y) + (1 - \lambda) \delta_{\bar{x}}^i(y')$ , and the claim holds. Otherwise,  $y_i \neq y'_i$ . Assume without loss of generality that  $z_i > \bar{x}_i$  and  $y_i > y'_i$  (and thus  $y_i > z_i > y'_i$ , since  $\lambda \in (0, 1)$ ). If  $y'_i > \bar{x}_i$ , we have:

$$z_i - \bar{x}_i = \lambda y_i + (1 - \lambda)y'_i - \bar{x}_i = \lambda(y_i - \bar{x}_i) + (1 - \lambda)(y'_i - \bar{x}_i)$$

By dividing by  $b_i - \bar{x}_i$ , we get:

$$\delta_{\bar{x}}^i(z) \leq \lambda \delta_{\bar{x}}^i(y) + (1 - \lambda) \delta_{\bar{x}}^i(y')$$

because  $y_i, z_i, y'_i > \bar{x}_i$  and their denominators in  $\delta_{\bar{x}}^i(y)$ ,  $\delta_{\bar{x}}^i(z)$ ,  $\delta_{\bar{x}}^i(y')$  are all  $b_i - \bar{x}_i$ . This proves the claim. Otherwise, in case  $y'_i \leq \bar{x}_i$ , we have:

$$\begin{aligned} z_i - \bar{x}_i &= \lambda y_i + (1 - \lambda)y'_i - \bar{x}_i \\ &= \lambda(y_i - \bar{x}_i) + (1 - \lambda)(y'_i - \bar{x}_i) \leq \lambda(y_i - \bar{x}_i) \end{aligned}$$

By dividing by  $b_i - \bar{x}_i$ , we get:  $\delta_{\bar{x}}^i(z) \leq \lambda \delta_{\bar{x}}^i(y)$ , because  $y_i, z_i > \bar{x}_i$  and their denominators in  $\delta_{\bar{x}}^i(y)$  and  $\delta_{\bar{x}}^i(z)$  are both  $b_i - \bar{x}_i$ . This proves the claim.

**Compactness:** We decompose  $\tilde{\mathcal{B}}_1^t(\bar{x})$  into  $2^k$  orthants:

$$\bigcup_{s_1, \dots, s_k \in \{\pm 1\}} \{y \in \tilde{\mathcal{B}}_1^t(\bar{x}) \mid \forall i \in [k] : s_i(y_i - \bar{x}_i) \geq 0\}$$

Since a finite union of compact sets is compact, we prove that every orthant is compact. We denote the orthant of fixed  $s_1, \dots, s_k$  by:

$$\tilde{\mathcal{B}}_1^t(\bar{x})_{s_1, \dots, s_k} \triangleq \{y \in \tilde{\mathcal{B}}_1^t(\bar{x}) \mid \forall i \in [k] : s_i(y_i - \bar{x}_i) \geq 0\}$$

We remind that  $\tilde{\mathcal{B}}_1^t(\bar{x}) = \{y \in \mathbb{R}^k \mid \sum_{i=1}^k \delta_{\bar{x}}^i(y) \leq t\}$  and  $\delta_{\bar{x}}^i(y) = \frac{y_i - \bar{x}_i}{\mathbf{1}_{\{y_i > \bar{x}_i\}}(b_i - \bar{x}_i) + \mathbf{1}_{\{y_i < \bar{x}_i\}}(a_i - \bar{x}_i)}$ . Consider the set of inputs within the first orthant of the  $\ell_1$ -ball with radius  $t$ :  $\Delta_{k,t} = \{z \in [0, \infty)^k \mid \sum_{i=1}^k z_i \leq t\}$ , which is compact. We show that  $\Delta_{k,t}$  can be transformed to  $\tilde{\mathcal{B}}_1^t(\bar{x})_{s_1, \dots, s_k}$  with an affine transformation (which in particular is continuous) and thus  $\tilde{\mathcal{B}}_1^t(\bar{x})_{s_1, \dots, s_k}$  is compact as a continuous image of a compact set. The affine transformation from  $\Delta_{k,t}$  to  $\tilde{\mathcal{B}}_1^t(\bar{x})_{s_1, \dots, s_k}$  is:  $y_i = (b_i - \bar{x}_i) \cdot z_i + \bar{x}_i$ , if  $s_i = 1$ , and  $y_i = (a_i - \bar{x}_i) \cdot z_i + \bar{x}_i$ , if  $s_i = -1$ . This concludes the proof.  $\square$

**Lemma 2.**  $E_{\cap} \subseteq \mathcal{B}_0^t(\bar{x})$ , where  $E_{\cap}$  is the extreme points of  $\mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x})$ .

*Proof.* We show that  $E_{\cap}$  is contained in a subset of  $\mathcal{B}_0^t(\bar{x})$  consisting of all inputs  $z$  with at least  $k - t$  entries  $z_i = \bar{x}_i$  and the other entries are the bounds  $a_i$  or  $b_i$ :

$$\{z \in \prod_{i=1}^k \{a_i, \bar{x}_i, b_i\} \mid |\{i \in [k] \mid z_i \neq \bar{x}_i\}| \leq t\}$$

Let  $y \in E_{\cap}$ . We show: (1)  $y \in \prod_{i=1}^k \{a_i, \bar{x}_i, b_i\}$  and (2) if  $y \in \prod_{i=1}^k \{a_i, \bar{x}_i, b_i\}$ , then  $|\{i \in [k] : y_i \neq \bar{x}_i\}| \leq t$ .

*Proof for (2):* Since  $y \in E_{\cap} \subseteq \tilde{\mathcal{B}}_1^t(\bar{x})$ , we get  $\sum_{i=1}^k \delta_{\bar{x}}^i(y) \leq t$ . Since  $y \in \prod_{i=1}^k \{a_i, \bar{x}_i, b_i\}$ , for every  $i$ , either  $y_i = \bar{x}_i$  and  $\delta_{\bar{x}}^i(y) = 0$  or  $y_i \in \{a_i, b_i\}$  and then  $\delta_{\bar{x}}^i(y) = 1$ . Thus,  $|\{i \in [k] : y_i \neq \bar{x}_i\}| \leq t$ .

*Proof for (1):* Assume to the contrary that there exists  $l$  such that  $y_l \notin \{a_l, \bar{x}_l, b_l\}$ . Since  $y \in E_{\cap} \subseteq \mathcal{D}$ , we have  $\delta_{\bar{x}}^i(y) \in [0, 1]$  for every  $i \in [k]$ . As  $y_l \notin \{a_l, \bar{x}_l, b_l\}$  we also have  $\delta_{\bar{x}}^l(y) \in (0, 1)$ . We show that in this case  $y$  is a convex combination of two different inputs  $y', y'' \in \mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x})$  and thus  $y$  cannot be an extreme point, in contradiction. We define these two points symmetrically, i.e.,  $y = 0.5(y' + y'')$ . We remind that since  $y \in E_{\cap} \subseteq \tilde{\mathcal{B}}_1^t(\bar{x})$ , we have  $\sum_{i=1}^k \delta_{\bar{x}}^i(y) \leq t$ . Accordingly, we split to cases:

- $\sum_{i=1}^k \delta_{\bar{x}}^i(y) < t$ : In this case,  $y$  is within  $\tilde{\mathcal{B}}_1^t(\bar{x})$  (not on its boundary). We identify a small value that can be added and subtracted from  $y_l$  such that the new inputs do not cross  $\bar{x}_l$  (possible since  $y_l \neq \bar{x}_l$  and thus  $\delta_{\bar{x}}^l(y) > 0$ ), do not cross the bound that is closer to  $y_l$  than to  $\bar{x}_l$  (possible since  $y_l \neq a_l, b_l$  and thus  $\delta_{\bar{x}}^l(y) < 1$ ) and do not cross the boundary of  $\tilde{\mathcal{B}}_1^t(\bar{x})$  (possible since  $\sum_{i=1}^k \delta_{\bar{x}}^i(y) < t$ ). Formally,  $y' = y + \epsilon \cdot v_l$  and  $y'' = y - \epsilon \cdot v_l$  for:

$$v_l = (\mathbf{1}_{\{y_l > \bar{x}_l\}}(b_l - \bar{x}_l) + \mathbf{1}_{\{y_l < \bar{x}_l\}}(a_l - \bar{x}_l))e_l$$

where  $e_l$  is the standard unit vector with 1 at entry  $l$  and 0 elsewhere, and

$$\epsilon = \min\{\delta_{\bar{x}}^l(y), 1 - \delta_{\bar{x}}^l(y), t - \sum_{i=1}^k \delta_{\bar{x}}^i(y)\}$$

As  $\epsilon > 0$  and  $v_l$  is not the zero vector (since  $y_l$  lies strictly between  $\bar{x}_l$  and one of the bounds  $a_l$  or  $b_l$ ) we have  $y' \neq y''$  and by the construction  $y', y'' \in \mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x})$ .

- $\sum_{i=1}^k \delta_{\bar{x}}^i(y) = t$ : In this case,  $y$  is on the boundary of  $\tilde{\mathcal{B}}_1^t(\bar{x})$ . We show there is another index  $m$  such that  $y_m \notin \{a_m, \bar{x}_m, b_m\}$  and identify small values that can be added and subtracted from  $y_l, y_m$  such that the new inputs do not cross  $\bar{x}_l, \bar{x}_m$  (possible since  $y_l \neq \bar{x}_l, y_m \neq \bar{x}_m$  and thus  $\delta_{\bar{x}}^l(y), \delta_{\bar{x}}^m(y) > 0$ ), do not cross the bound that is closer to  $y_l$  and  $y_m$  than to  $\bar{x}_l$  and  $\bar{x}_m$  (possible since  $y_l \neq a_l, b_l$  and  $y_m \neq a_m, b_m$  and thus  $\delta_{\bar{x}}^l(y), \delta_{\bar{x}}^m(y) < 1$ ) and remain on the boundary of  $\tilde{\mathcal{B}}_1^t(\bar{x})$ . We have  $\delta_{\bar{x}}^l(y) \in (0, 1)$  and since  $\sum_{i=1}^k \delta_{\bar{x}}^i(y) = t$  for an integer  $t$ , there must be  $m \neq l$  such that  $\delta_{\bar{x}}^m(y) \in (0, 1)$ , and  $y_m \notin \{a_m, \bar{x}_m, b_m\}$ . We define  $y' = y + \epsilon \cdot v_{lm}$  and  $y'' = y - \epsilon \cdot v_{lm}$ , where  $v_{lm}$  changes coordinates  $l$  and  $m$  in a scaled complementary manner to preserve

$\sum_{i=1}^k \delta_{\bar{x}}^i(y) = t$ , formally  $v_{lm} = v_l - v_m$  for:

$$v_l = (\mathbf{1}_{\{y_l > \bar{x}_l\}}(b_l - \bar{x}_l) + \mathbf{1}_{\{y_l < \bar{x}_l\}}(a_l - \bar{x}_l)) \cdot e_l$$

and similarly:

$$v_m = (\mathbf{1}_{\{y_m > \bar{x}_m\}}(b_m - \bar{x}_m) + \mathbf{1}_{\{y_m < \bar{x}_m\}}(a_m - \bar{x}_m)) \cdot e_m$$

and:

$$\epsilon = \min\{\delta_{\bar{x}}^l(y), 1 - \delta_{\bar{x}}^l(y), \delta_{\bar{x}}^m(y), 1 - \delta_{\bar{x}}^m(y)\}$$

As  $\epsilon > 0$  and  $v_{lm}$  is not the zero vector we have  $y' \neq y''$  and by the construction  $y', y'' \in \mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x})$ .  $\square$

**Theorem 2.** 1.  $\text{vol}(\tilde{\mathcal{B}}_1^t(\bar{x})) = \text{vol}(\mathcal{D}) \frac{t^k}{k!}$ .  
2.  $\text{vol}(\mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x})) = \text{vol}(\mathcal{D}) \frac{t^k}{k!} \sum_{r=0}^{t-1} (-1)^r \binom{k}{r} (1 - \frac{r}{t})^k$ .

*Proof.* We decompose  $\tilde{\mathcal{B}}_1^t(\bar{x})$  into  $2^k$  orthants:

$$\bigcup_{s_1, \dots, s_k \in \{\pm 1\}} \{y \in \tilde{\mathcal{B}}_1^t(\bar{x}) \mid \forall i \in [k] : s_i(y_i - \bar{x}_i) \geq 0\}$$

where the intersection of difference sets in this union has zero volume. We next determine the volume of a single orthant, with a fixed  $s_1, \dots, s_k$ , denoted by  $\tilde{\mathcal{B}}_1^t(\bar{x})_{s_1, \dots, s_k} \triangleq \{y \in \tilde{\mathcal{B}}_1^t(\bar{x}) \mid \forall i \in [k] : s_i(y_i - \bar{x}_i) \geq 0\}$ . We remind that  $\tilde{\mathcal{B}}_1^t(\bar{x}) = \{y \in \mathbb{R}^k \mid \sum_{i=1}^k \delta_{\bar{x}}^i(y) \leq t\}$  and  $\delta_{\bar{x}}^i(y) = \frac{y_i - \bar{x}_i}{\mathbf{1}_{\{y_i > \bar{x}_i\}}(b_i - \bar{x}_i) + \mathbf{1}_{\{y_i < \bar{x}_i\}}(a_i - \bar{x}_i)}$ . Consider the set of inputs within the first orthant of the  $\ell_1$ -ball with radius  $t$ :  $\Delta_{k,t} = \{z \in [0, \infty)^k \mid \sum_{i=1}^k z_i \leq t\}$ . It is known that  $\text{vol}(\Delta_{k,t}) = \frac{t^k}{k!}$ . We next transform  $\Delta_{k,t}$  to  $\tilde{\mathcal{B}}_1^t(\bar{x})_{s_1, \dots, s_k}$  by a change of variables:  $y_i = (b_i - \bar{x}_i) \cdot z_i + \bar{x}_i$ , if  $s_i = 1$ , and  $y_i = (a_i - \bar{x}_i) \cdot z_i + \bar{x}_i$ , if  $s_i = -1$ . The change of variables shows that every axis  $z_i$  is scaled by  $|b_i - \bar{x}_i|$ , if  $s_i = 1$ , or by  $|a_i - \bar{x}_i|$ , if  $s_i = -1$ , thus:

$$\begin{aligned} \text{vol}(\tilde{\mathcal{B}}_1^t(\bar{x})_{s_1, \dots, s_k}) &= \\ \text{vol}(\Delta_{k,t}) \prod_{i=1}^k (\mathbf{1}_{\{s_i=1\}}|b_i - \bar{x}_i| + \mathbf{1}_{\{s_i=-1\}}|a_i - \bar{x}_i|) \end{aligned}$$

As mentioned, the total volume is their summation. We observe that the summation of the scaling is  $\text{vol}(\mathcal{D})$ :

$$\begin{aligned} \text{vol}(\mathcal{D}) &= \prod_{i=1}^k (b_i - a_i) = \prod_{i=1}^k (|b_i - \bar{x}_i| + |a_i - \bar{x}_i|) = \\ &\sum_{s_1, \dots, s_k \in \{\pm 1\}} \prod_{i=1}^k (\mathbf{1}_{\{s_i=1\}}|b_i - \bar{x}_i| + \mathbf{1}_{\{s_i=-1\}}|a_i - \bar{x}_i|) \end{aligned}$$

Overall, we get  $\text{vol}(\tilde{\mathcal{B}}_1^t(\bar{x})) = \text{vol}(\mathcal{D}) \frac{t^k}{k!}$ .

The proof for  $\mathcal{D} \cap \tilde{\mathcal{B}}_1^t(\bar{x})$  is similar. First, we split to the orthants:  $\tilde{\mathcal{B}}_1^t(\bar{x})_{s_1, \dots, s_k} \cap \mathcal{D}$ . Second, we consider  $\Delta_{k,t} \cap [0, 1]^k$

(since  $\delta_{\bar{x}}^i(y) \in [0, 1]$  for  $y \in \mathcal{D}$ ). Third, we use the same variable change and obtain the same scaling factor:

$$\begin{aligned} \text{vol}(\tilde{\mathcal{B}}_1^t(\bar{x})_{s_1, \dots, s_k} \cap \mathcal{D}) &= \\ \text{vol}(\Delta_{k,t} \cap [0, 1]^k) \prod_{i=1}^k (\mathbf{1}_{\{s_i=1\}}|b_i - \bar{x}_i| + \mathbf{1}_{\{s_i=-1\}}|a_i - \bar{x}_i|) \end{aligned}$$

Fourth, again, the summation of the scaling is  $\text{vol}(\mathcal{D})$ , thus we get  $\text{vol}(\tilde{\mathcal{B}}_1^t(\bar{x}) \cap \mathcal{D}) = \text{vol}(\mathcal{D}) \cdot \text{vol}(\Delta_{k,t} \cap [0, 1]^k)$ . Fifth, the second term can be computed by the Irwin–Hall distribution (Irwin 1927; Hall 1927). Technically, to compute the volume, we compute the probability that a uniformly sampled input from  $[0, 1]^k$  lies in  $\Delta_{k,t}$ , that is its sum of entries is at most  $t$ . The sum of entries of a uniformly sampled input in  $[0, 1]^k$  is distributed as the Irwin-Hall distribution. Therefore, the probability we are interested in is obtained by evaluating the Irwin-Hall CDF at  $t$ . Since the volume of  $[0, 1]^k$  is 1, we get that  $\text{vol}(\Delta_{k,t} \cap [0, 1]^k)$  is exactly this probability:  $\text{vol}(\Delta_{k,t} \cap [0, 1]^k) = \frac{t^k}{k!} \sum_{r=0}^{t-1} (-1)^r \binom{k}{r} (1 - \frac{r}{t})^k$ .  $\square$

**Lemma 3.**  $\tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})$  is convex and compact.

*Proof.* **Convexity:** Let  $y, y' \in \tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})$  and  $\lambda \in (0, 1)$ . We prove that for  $z = \lambda y + (1 - \lambda)y'$ , we have  $z \in \tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})$ . We show it by proving:

$$\begin{aligned} \sum_{i=1}^k \max_{j \in [d]} \delta_{\bar{x}}^{i,j}(z) &\leq \\ \lambda \sum_{i=1}^k \max_{j \in [d]} \delta_{\bar{x}}^{i,j}(y) + (1 - \lambda) \sum_{i=1}^k \max_{j \in [d]} \delta_{\bar{x}}^{i,j}(y') &\leq t \end{aligned}$$

To this end, we prove that for all  $i \in [k]$ :

$$\max_{j \in [d]} \delta_{\bar{x}}^{i,j}(z) \leq \lambda \max_{j \in [d]} \delta_{\bar{x}}^{i,j}(y) + (1 - \lambda) \max_{j \in [d]} \delta_{\bar{x}}^{i,j}(y')$$

We prove it by showing that for all  $i \in [k]$  and  $j \in [d]$ :  $\delta_{\bar{x}}^{i,j}(z) \leq \lambda \delta_{\bar{x}}^{i,j}(y) + (1 - \lambda) \delta_{\bar{x}}^{i,j}(y')$ . The proof is identical to the proof of the corresponding inequality in Lemma 1 with an additional index.

**Compactness:** We decompose  $\tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})$  into  $(2^d)^k$  orthants:

$$\begin{aligned} \bigcup_{s_1, \dots, s_k \in \{\pm 1\}^d} \\ \{y \in \tilde{\mathcal{B}}_{1,\infty}^t(\bar{x}) \mid \forall i \in [k], j \in [d] : s_i^{(j)}(y_i^{(j)} - \bar{x}_i^{(j)}) \geq 0\} \end{aligned}$$

Since a finite union of compact sets is compact, we prove that every orthant is compact. We denote the orthant of fixed  $s_1, \dots, s_k$  by:

$$\begin{aligned} \tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})_{s_1, \dots, s_k} \triangleq \\ \{y \in \tilde{\mathcal{B}}_{1,\infty}^t(\bar{x}) \mid \forall i \in [k], j \in [d] : s_i^{(j)}(y_i^{(j)} - \bar{x}_i^{(j)}) \geq 0\} \end{aligned}$$

We remind that:

$$\tilde{\mathcal{B}}_{1,\infty}^t(\bar{x}) = \{y \in (\mathbb{R}^d)^k \mid \sum_{i=1}^k \max_{j \in [d]} \delta_{\bar{x}}^{i,j}(y) \leq t\}$$

and

$$\delta_{\bar{x}}^{i,j}(y) = \frac{y_i^{(j)} - \bar{x}_i^{(j)}}{\mathbf{1}_{y_i^{(j)} > \bar{x}_i^{(j)}}(b_i^{(j)} - \bar{x}_i^{(j)}) + \mathbf{1}_{y_i^{(j)} < \bar{x}_i^{(j)}}(a_i^{(j)} - \bar{x}_i^{(j)})}$$

Let  $\Delta_{k,d,t} \triangleq \{z \in ([0, \infty)^d)^k \mid \sum_{i=1}^k \max_{j \in [d]} z_i^{(j)} \leq t\}$ . We show it is compact as a finite union of compact sets. To this end, we decompose it into sets based on the channel  $j_i \in [d]$  of each entry  $i \in [k]$  with the maximum value:

$$\Delta_{k,d,t} = \bigcup_{j_1, \dots, j_k \in [d]} \{z \in ([0, \infty)^d)^k \mid \sum_{i=1}^k z_i^{(j_i)} \leq t \wedge \forall i \in [k], j \in [d] : z_i^{(j_i)} \geq z_i^{(j)}\}$$

Note that each of these sets is bounded by linear constraints, and is thus compact. Thus,  $\Delta_{k,d,t}$  is compact. Next, we show that  $\Delta_{k,d,t}$  can be transformed to  $\tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})_{s_1, \dots, s_k}$  with an affine transformation (which in particular is continuous) and thus  $\tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})_{s_1, \dots, s_k}$  is compact as a continuous image of a compact set. The affine transformation is:  $y_i^{(j)} = (b_i^{(j)} - \bar{x}_i^{(j)}) \cdot z_i^{(j)} + \bar{x}_i^{(j)}$ , if  $s_i^{(j)} = 1$ , and  $y_i^{(j)} = (a_i^{(j)} - \bar{x}_i^{(j)}) \cdot z_i^{(j)} + \bar{x}_i^{(j)}$ , if  $s_i^{(j)} = -1$ .  $\square$

**Lemma 4.**  $E_\cap \subseteq \mathcal{B}_0^t(\bar{x})$ , where  $E_\cap$  is the extreme points of  $\mathcal{D} \cap \tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})$ .

*Proof.* We show that  $E_\cap$  is contained in a subset of  $\mathcal{B}_0^t(\bar{x})$ , consisting of all inputs  $z$  with at least  $k - t$  entries  $z_i = \bar{x}_i$  and for the other entries all the channels are either  $\bar{x}_i^{(j)}$  or the bounds  $a_i^{(j)}$  or  $b_i^{(j)}$ :

$$\{z \in \prod_{i=1}^k \left( \prod_{j=1}^d \{a_i^{(j)}, \bar{x}_i^{(j)}, b_i^{(j)}\} \right) \mid |\{i \in [k] : z_i \neq \bar{x}_i\}| \leq t\}$$

Let  $y \in E_\cap$ . We show:

1.  $y \in \prod_{i=1}^k \left( \prod_{j=1}^d \{a_i^{(j)}, \bar{x}_i^{(j)}, b_i^{(j)}\} \right)$  and
2. if  $y \in \prod_{i=1}^k \left( \prod_{j=1}^d \{a_i^{(j)}, \bar{x}_i^{(j)}, b_i^{(j)}\} \right)$ , then:  $|\{i \in [k] \mid y_i \neq \bar{x}_i\}| \leq t$ .

*Proof for 2:* Since  $y \in E_\cap$  and  $E_\cap \subseteq \tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})$ , we get  $\sum_{i=1}^k \max_{j \in [d]} \delta_{\bar{x}}^{i,j}(y) \leq t$ . We are given  $y \in \prod_{i=1}^k \left( \prod_{j=1}^d \{a_i^{(j)}, \bar{x}_i^{(j)}, b_i^{(j)}\} \right)$ . Thus, for every  $i$ , either  $y_i = \bar{x}_i$  and thus  $\max_{j \in [d]} \delta_{\bar{x}}^{i,j}(y) = 0$  or  $\max_{j \in [d]} \delta_{\bar{x}}^{i,j}(y) = 1$ . Thus,  $|\{i \in [k] \mid y_i \neq \bar{x}_i\}| \leq t$ .

*Proof for 1:* Assume to the contrary that there are  $l \in [k]$  and  $l' \in [d]$  such that  $y_l^{(l')} \notin \{a_l^{(l')}, \bar{x}_l^{(l')}, b_l^{(l')}\}$ . Since  $y \in E_\cap \subseteq \mathcal{D}$ , we have  $\delta_{\bar{x}}^{i,j}(y) \in [0, 1]$  for every  $i \in [k]$  and  $j \in [d]$ . Since  $y_l^{(l')} \notin \{a_l^{(l')}, \bar{x}_l^{(l')}, b_l^{(l')}\}$ , we get  $\delta_{\bar{x}}^{l,l'}(y) \in (0, 1)$ . We show that in this case  $y$  is a convex combination of two different inputs  $y', y'' \in \mathcal{D} \cap \tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})$  and thus  $y$  cannot be an extreme point, in contradiction. We define these two points

symmetrically, i.e.,  $y = 0.5(y' + y'')$ . We remind that since  $y \in E_\cap \subseteq \tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})$ , we have  $\sum_{i=1}^k \max_{j \in [d]} \delta_{\bar{x}}^{i,j}(y) \leq t$ . Accordingly, we split to cases:

- $\sum_{i=1}^k \max_{j \in [d]} \delta_{\bar{x}}^{i,j}(y) < t$ : In this case,  $y$  is within  $\tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})$  (not on its boundary). The proof for this case is similar to the proof of the corresponding case in Lemma 2, with an additional index. We repeat it here for convenience. We identify a small value that can be added and subtracted from  $y_l^{(l')}$  such that the new inputs do not cross  $\bar{x}_l^{(l')}$  (possible since  $y_l^{(l')} \neq \bar{x}_l^{(l')}$  and thus  $\delta_{\bar{x}}^{l,l'}(y) > 0$ ), do not cross the bound that is closer to  $y_l^{(l')}$  than to  $\bar{x}_l^{(l')}$  (possible since  $y_l^{(l')} \neq a_l^{(l')}, b_l^{(l')}$  and thus  $\delta_{\bar{x}}^{l,l'}(y) < 1$ ) and do not cross the boundary of  $\tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})$  (possible since  $\sum_{i=1}^k \max_{j \in [d]} \delta_{\bar{x}}^{i,j}(y) < t$ ). Formally,  $y' = y + \epsilon \cdot v_l^{(l')}$  and  $y'' = y - \epsilon \cdot v_l^{(l')}$  for:

$$v_l^{(l')} = (\mathbf{1}_{y_l^{(l')} > \bar{x}_l^{(l')}}(b_l^{(l')} - \bar{x}_l^{(l')}) + \mathbf{1}_{y_l^{(l')} < \bar{x}_l^{(l')}}(a_l^{(l')} - \bar{x}_l^{(l')}))e_l^{(l')}$$

where  $e_l^{(l')}$  is the standard unit vector with 1 at the  $(l')$ th channel of entry  $l$  and 0 elsewhere, and:

$$\epsilon = \min\{\delta_{\bar{x}}^{l,l'}(y), 1 - \delta_{\bar{x}}^{l,l'}(y), t - \sum_{i=1}^k \max_{j \in [d]} \delta_{\bar{x}}^{i,j}(y)\}$$

As  $\epsilon > 0$  and  $v_l^{(l')}$  is not the zero vector (since  $y_l^{(l')}$  lies strictly between  $\bar{x}_l^{(l')}$  and one of the bounds  $a_l^{(l')}$  or  $b_l^{(l')}$ ) we have  $y' \neq y''$  and by the construction  $y', y'' \in \mathcal{D} \cap \tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})$ .

- $\sum_{i=1}^k \max_{j \in [d]} \delta_{\bar{x}}^{i,j}(y) = t$ : In this case,  $y$  is on the boundary of  $\tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})$ . Similarly to the respective case in Lemma 2, we look for another entry and identify small values that can be added and subtracted in a scaled complementary manner to obtain  $y'$  and  $y''$ . Compared to Lemma 2, this case is more challenging, since  $\tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})$  considers the maximal  $\delta_{\bar{x}}^{i,j}(y)$  over all  $d$  channels of the  $i$ th entry. Thus, we may need to add or subtract from multiple channels of the same entry. Since  $y$  is an extreme point (by our assumption), for all  $i \in [k]$ , the non zero values among  $\delta_{\bar{x}}^{i,j}(y)$  for  $j \in [d]$  are equal (the proof is given at the end of this proof and marked with  $\star$ ). Given  $i \in [k]$ , we denote by  $\mathcal{J}_i$  the set of all channels  $j \in [d]$  such that  $\delta_{\bar{x}}^{i,j}(y) \neq 0$ . Next, we consider  $i = l$ . Since  $\delta_{\bar{x}}^{l,l'}(y) \in (0, 1)$ , we have  $\delta_{\bar{x}}^{l,j}(y) = \delta_{\bar{x}}^{l,l'}(y) \in (0, 1)$  for all  $j \in \mathcal{J}_l$ . Thus,  $y_l^{(j)} \notin \{a_l^{(j)}, \bar{x}_l^{(j)}, b_l^{(j)}\}$  for all  $j \in \mathcal{J}_l$ . We show that there is another entry  $m \neq l$  such that  $\max_{j \in [d]} \delta_{\bar{x}}^{m,j}(y) \in (0, 1)$  and  $y_m^{(j)} \notin \{a_m^{(j)}, \bar{x}_m^{(j)}, b_m^{(j)}\}$  for all  $j \in \mathcal{J}_m$  and identify small values that can be added and subtracted from  $y_l^{(j)}, y_m^{(j)}$  (for all  $j \in \mathcal{J}_l$  and  $j' \in \mathcal{J}_m$ ) such that the new inputs do not cross  $\bar{x}_l^{(j)}, \bar{x}_m^{(j')}$  (possible since  $y_l^{(j)} \neq \bar{x}_l^{(j)}, y_m^{(j')} \neq \bar{x}_m^{(j')}$ ) and

thus  $\delta_{\bar{x}}^{l,j}(y), \delta_{\bar{x}}^{m,j'}(y) > 0$ , for all  $j \in \mathcal{J}_l$  and  $j' \in \mathcal{J}_m$ , do not cross the bound that is closer to  $y_l^{(j)}$  and  $y_m^{(j')}$  than to  $\bar{x}_l^{(j)}$  and  $\bar{x}_m^{(j')}$  (possible since  $y_l^{(j)} \neq a_l^{(j)}, b_l^{(j)}$  and  $y_m^{(j')} \neq a_m^{(j')}, b_m^{(j')}$ ) and thus  $\delta_{\bar{x}}^{l,j}(y), \delta_{\bar{x}}^{m,j'}(y) < 1$ , for all  $j \in \mathcal{J}_l$  and  $j' \in \mathcal{J}_m$  and remain on the boundary of  $\tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})$ . Since  $\sum_{i=1}^k \max_{j \in [d]} \delta_{\bar{x}}^{i,j}(y) = t$  for an integer  $t$  and since  $\max_{j \in [d]} \delta_{\bar{x}}^{l,j}(y) \in (0, 1)$ , there must be  $m \neq l$  such that  $\max_{j \in [d]} \delta_{\bar{x}}^{m,j}(y) \in (0, 1)$ . As explained before, for all  $j \in \mathcal{J}_m$ , we have  $y_m^{(j)} \notin \{a_m^{(j)}, \bar{x}_m^{(j)}, b_m^{(j)}\}$ . We define  $y' = y + \epsilon \cdot v_{lm}$  and  $y'' = y - \epsilon \cdot v_{lm}$ , where  $v_{lm}$  changes coordinates in  $l$  and  $m$  in a scaled complementary manner to preserve  $\sum_{i=1}^k \max_{j \in [d]} \delta_{\bar{x}}^{i,j}(y) = t$ . Formally,  $v_{lm} = v_l - v_m$  where:

$$v_l = \sum_{j \in \mathcal{J}_l} (\mathbf{1}_{\{y_l^{(j)} > \bar{x}_l^{(j)}\}} (b_l^{(j)} - \bar{x}_l^{(j)}) + \mathbf{1}_{\{y_l^{(j)} < \bar{x}_l^{(j)}\}} (a_l^{(j)} - \bar{x}_l^{(j)})) \cdot e_l^{(j)}$$

and similarly:

$$v_m = \sum_{j \in \mathcal{J}_m} (\mathbf{1}_{\{y_m^{(j)} > \bar{x}_m^{(j)}\}} (b_m^{(j)} - \bar{x}_m^{(j)}) + \mathbf{1}_{\{y_m^{(j)} < \bar{x}_m^{(j)}\}} (a_m^{(j)} - \bar{x}_m^{(j)})) \cdot e_m^{(j)}$$

and:

$$\epsilon = \min[\min_{j \in \mathcal{J}_l} \{\delta_{\bar{x}}^{l,j}(y)\}, \min_{j \in \mathcal{J}_l} \{1 - \delta_{\bar{x}}^{l,j}(y)\}, \min_{j \in \mathcal{J}_m} \{\delta_{\bar{x}}^{m,j}(y)\}, \min_{j \in \mathcal{J}_m} \{1 - \delta_{\bar{x}}^{m,j}(y)\}]$$

As  $\epsilon > 0$  and  $v_{lm}$  is not the zero vector, we have  $y' \neq y''$  and by the construction  $y', y'' \in \mathcal{D} \cap \tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})$ .

Proof for  $\star$ : We prove that, if  $y$  in an extreme point, then for every  $i \in [k]$  all the non zero values of  $\delta_{\bar{x}}^{i,j}(y)$  are equal. Assume to the contrary that there are  $l \in [k]$  and  $l' \in [d]$  such that  $0 < \delta_{\bar{x}}^{l,l'}(y) < \max_{j \in [d]} \delta_{\bar{x}}^{l,j}(y)$ . We can slightly increase or decrease this channel without crossing  $\bar{x}_l^{(l')}$  and the bound that is closer to  $y_l^{(l')}$  than to  $\bar{x}_l^{(l')}$  and without changing  $\max_{j \in [d]} \delta_{\bar{x}}^{l,j}(y)$  to obtain two different inputs  $y', y'' \in \mathcal{D} \cap \tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})$  such that  $y = 0.5(y' + y'')$  and thus  $y$  is not an extreme point. Formally, we define  $y' = y + \epsilon \cdot v_l^{(l')}$  and  $y'' = y - \epsilon \cdot v_l^{(l')}$  for

$$v_l^{(l')} = (\mathbf{1}_{\{y_l^{(l')} > \bar{x}_l^{(l')}\}} (b_l^{(l')} - \bar{x}_l^{(l')}) + \mathbf{1}_{\{y_l^{(l')} < \bar{x}_l^{(l')}\}} (a_l^{(l')} - \bar{x}_l^{(l')})) e_l^{(l')}$$

and  $\epsilon = \min\{\delta_{\bar{x}}^{l,l'}(y), \max_{j \in [d]} \delta_{\bar{x}}^{l,j}(y) - \delta_{\bar{x}}^{l,l'}(y)\}$ . As  $\epsilon > 0$  and  $v_l^{(l')}$  is not the zero vector, we have  $y' \neq y''$  and by the construction  $y', y'' \in \mathcal{D} \cap \tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})$ .  $\square$

**Theorem 4.** 1.  $\text{vol}(\tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})) = \text{vol}(\mathcal{D}) \frac{t^{dk} (d!)^k}{(dk)!}$ .

2.  $\text{vol}(\mathcal{D} \cap \tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})) = \text{vol}(\mathcal{D}) \frac{t^{dk} (d!)^k}{(dk)!} (1 + \sum_{r=1}^{t-1} c_{r,k,d,t})$ , where

$$c_{r,k,d,t} = (-1)^r \binom{k}{r} \sum_{m_1, \dots, m_r \in [d]} \frac{(dk)! (1 - \frac{r}{t})^{dk} (t-r)^{-dr + \sum_{i=1}^r m_i}}{(d(k-r) + \sum_{i=1}^r m_i)! \prod_{i=1}^r (d - m_i)!}$$

Like Theorem 2, given fixed  $d, t$  and  $r \in [1, t-1]$ , we have that  $c_{r,k,d,t}$  exponentially converges to 0 as  $k$  increases.

*Proof.* Like Theorem 2, we decompose  $\tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})$  and  $\mathcal{D} \cap \tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})$  by the  $2^{d \cdot k}$  orthants, where each is a scaled version of  $\Delta_{k,d,t} = \{z \in ([0, \infty)^d)^k \mid \sum_{i=1}^k \max_{j \in [d]} z_i^{(j)} \leq t\}$  for  $\tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})$  and  $\Delta_{k,d,t} \cap ([0, 1]^d)^k$  for  $\mathcal{D} \cap \tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})$ . Like Theorem 2, the summation of these scalings is  $\text{vol}(\mathcal{D})$ , thus  $\text{vol}(\tilde{\mathcal{B}}_{1,\infty}^t(\bar{x})) = \text{vol}(\mathcal{D}) \cdot \text{vol}(\Delta_{k,d,t})$  and  $\text{vol}(\tilde{\mathcal{B}}_{1,\infty}^t(\bar{x}) \cap \mathcal{D}) = \text{vol}(\mathcal{D}) \cdot \text{vol}(\Delta_{k,d,t} \cap ([0, 1]^d)^k)$ .

Next, we compute  $\text{vol}(\Delta_{k,d,t})$ :

$$\begin{aligned} \text{vol}(\Delta_{k,d,t}) &= \int_{\substack{\sum_{i=1}^k \max_{j \in [d]} z_i^{(j)} \leq t \\ \forall i \in [k], j \in [d]: 0 \leq z_i^{(j)} \leq z_i^{(1)}}} \prod_{i=1}^k \prod_{j=1}^d dz_i^{(j)} \\ &= d^k \int_{\substack{\sum_{i=1}^k z_i^{(1)} \leq t \\ \forall i \in [k], j \in [d]: 0 \leq z_i^{(j)} \leq z_i^{(1)}}} \prod_{i=1}^k \prod_{j=1}^d dz_i^{(j)} \\ &= d^k \int_{\substack{\sum_{i=1}^k z_i^{(1)} \leq t \\ \forall i \in [k]: 0 \leq z_i^{(1)} \leq z_i^{(2)}}} \prod_{i=1}^k dz_i^{(1)} \prod_{j=2}^d \int_0^{z_i^{(1)}} dz_i^{(j)} \\ &= d^k \int_{\substack{\sum_{i=1}^k z_i^{(1)} \leq t \\ \forall i \in [k]: 0 \leq z_i^{(1)} \leq z_i^{(2)}}} \prod_{i=1}^k (z_i^{(1)})^{d-1} dz_i^{(1)} \\ &= d^k \frac{t^{dk} ((d-1)!)^k}{(dk)!} = \frac{t^{dk} (d!)^k}{(dk)!} \end{aligned}$$

The second equality follows by partitioning  $\Delta_{k,d,t}$  into  $d^k$  symmetric options (one for each possibility of a maximal channel  $j \in [d]$  for each  $i \in [k]$ ). Note that the intersection of any two options has a zero volume. The third equality follows by Fubini's theorem. The fourth equality follows by computing the inner integrals, each contributing  $z_i^{(1)}$ . The fifth equality follows by Gradshteyn and Ryzhik (2014, Eq. 4.634). The last transition unifies  $d^k$  with  $(d-1)!^k$ .

We continue with  $\text{vol}(\Delta_{k,d,t} \cap ([0, 1]^d)^k)$ , which is identical to the first four transitions above, except that  $z_i^{(j)} \leq 1$ . We continue from the fourth transition, followed by a transition that omits the superscript  $(1)$  to simplify notations for

the following equalities:

$$\begin{aligned} & \text{vol}(\Delta_{k,d,t} \cap ([0,1]^d)^k) \\ &= d^k \int_{\substack{\sum_{i=1}^k z_i^{(1)} \leq t \\ \forall i \in [k]: 0 \leq z_i^{(1)} \leq 1}} \prod_{i=1}^k (z_i^{(1)})^{d-1} dz_i^{(1)} \\ &= d^k \int_{\substack{\sum_{i=1}^k z_i \leq t \\ \forall i \in [k]: 0 \leq z_i \leq 1}} \prod_{i=1}^k z_i^{d-1} dz_i \end{aligned}$$

To compute this integral, we remove the requirement of  $z_i \leq 1$  and then repair it by excluding all the cases of  $z_i \geq 1$ , using the inclusion-exclusion principle (similarly to a geometric derivation of Irwin-Hall by Marengo, Farnsworth, and Stefanic (2017)).

$$\begin{aligned} &= d^k \int_{\substack{\sum_{i=1}^k z_i \leq t \\ \forall i \in [k]: 0 \leq z_i}} \prod_{i=1}^k z_i^{d-1} dz_i \\ &+ d^k \sum_{r=1}^k (-1)^r \sum_{1 \leq i_1 < \dots < i_r \leq k} \int_{\substack{\sum_{i=1}^k z_i \leq t \\ \forall i \in [k]: 0 \leq z_i \\ 1 \leq z_{i_1}, \dots, z_{i_r}}} \prod_{i=1}^k z_i^{d-1} dz_i \\ &= \frac{t^{dk} (d!)^k}{(dk)!} \left( 1 + \sum_{r=1}^k c_{r,k,d,t} \right) \end{aligned}$$

where the last transition extracts the left term and denotes the rest of the terms  $c_{r,k,d,t}$ , where:

$$\begin{aligned} c_{r,k,d,t} &= \frac{(-1)^r (dk)!}{t^{dk} ((d-1)!)^k} \\ &\sum_{1 \leq i_1 < \dots < i_r \leq k} \int_{\substack{\sum_{i=1}^k z_i \leq t \\ \forall i \in [k]: 0 \leq z_i \\ 1 \leq z_{i_1}, \dots, z_{i_r}}} \prod_{i=1}^k z_i^{d-1} dz_i \\ &= \frac{(-1)^r (dk)!}{t^{dk} ((d-1)!)^k} \binom{k}{r} \int_{\substack{\sum_{i=1}^k z_i \leq t \\ \forall i \in [k]: 0 \leq z_i \\ 1 \leq z_1, \dots, z_r}} \prod_{i=1}^k z_i^{d-1} dz_i \end{aligned}$$

The above transition follows since the integral is the same for every choice of  $i_1, \dots, i_r$ . Note that if  $r \geq t$ , then  $c_{r,k,d,t} = 0$ , because of the conditions  $\sum_{i=1}^k z_i \leq t$  and  $1 \leq z_1, \dots, z_r$ . Thus, the sum is up to  $t-1$  instead of  $k$ :

$$\text{vol}(\Delta_{k,d,t} \cap ([0,1]^d)^k) = \frac{t^{dk} (d!)^k}{(dk)!} \left( 1 + \sum_{r=1}^{t-1} c_{r,k,d,t} \right)$$

Lastly, we wish to compute the integral of  $c_{r,k,d,t}$  using Gradshteyn and Ryzhik (2014, Eq. 4.634). To this end, we change variables and rely on the binomial expansion. We first change the variables  $z_i \mapsto z_i - 1$  for  $i \in [r]$ :

$$\begin{aligned} & \int_{\substack{\sum_{i=1}^k z_i \leq t \\ \forall i \in [k]: 0 \leq z_i \\ 1 \leq z_1, \dots, z_r}} \prod_{i=1}^k z_i^{d-1} dz_i \\ &= \int_{\substack{\sum_{i=1}^k z_i \leq t-r \\ \forall i \in [k]: 0 \leq z_i}} \left( \prod_{i=1}^r (z_i + 1)^{d-1} dz_i \right) \left( \prod_{i=r+1}^k z_i^{d-1} dz_i \right) \end{aligned}$$

We then rely on the binomial expansion for  $\prod_{i=1}^r (z_i + 1)^{d-1}$ :

$$\begin{aligned} \prod_{i=1}^r (z_i + 1)^{d-1} &= \prod_{i=1}^r \sum_{m=0}^{d-1} \binom{d-1}{m} z_i^m = \\ \prod_{i=1}^r \sum_{m=1}^d \binom{d-1}{m-1} z_i^{m-1} &= \sum_{m_1, \dots, m_r \in [d]} \prod_{i=1}^r \binom{d-1}{m_i-1} z_i^{m_i-1} \end{aligned}$$

The second transition follows by changing the sum over 1 to  $d$ . The third transition switches the multiplication and the sum. Substituting it into the integral:

$$\begin{aligned} & \int_{\substack{\sum_{i=1}^k z_i \leq t \\ \forall i \in [k]: 0 \leq z_i \\ 1 \leq z_1, \dots, z_r}} \prod_{i=1}^k z_i^{d-1} dz_i \\ &= \sum_{m_1, \dots, m_r \in [d]} \left( \prod_{i=1}^r \binom{d-1}{m_i-1} \right) \\ & \int_{\substack{\sum_{i=1}^k z_i \leq t-r \\ \forall i \in [k]: 0 \leq z_i}} \prod_{i=1}^r z_i^{m_i-1} dz_i \prod_{i=r+1}^k z_i^{d-1} dz_i \\ &= \sum_{m_1, \dots, m_r \in [d]} \left( \prod_{i=1}^r \binom{d-1}{m_i-1} \right) \\ & \frac{(t-r)^{d(k-r) + \sum_{i=1}^r m_i} ((d-1)!)^{k-r} \prod_{i=1}^r (m_i-1)!}{(d(k-r) + \sum_{i=1}^r m_i)!} \\ &= \sum_{m_1, \dots, m_r \in [d]} \frac{(t-r)^{d(k-r) + \sum_{i=1}^r m_i} ((d-1)!)^k}{(d(k-r) + \sum_{i=1}^r m_i)! \prod_{i=1}^r (d-m_i)!} \end{aligned}$$

The second equality follows by Gradshteyn and Ryzhik (2014, Eq. 4.634). The last equality is algebraic rewrites. We substitute into  $c_{r,k,d,t}$  and simplify with algebraic rewrite:

$$\begin{aligned} c_{r,k,d,t} &= \frac{(-1)^r (dk)!}{t^{dk} ((d-1)!)^k} \binom{k}{r} \\ &\sum_{m_1, \dots, m_r \in [d]} \frac{(t-r)^{d(k-r) + \sum_{i=1}^r m_i} ((d-1)!)^k}{(d(k-r) + \sum_{i=1}^r m_i)! \prod_{i=1}^r (d-m_i)!} \\ &= (-1)^r \binom{k}{r} \\ &\sum_{m_1, \dots, m_r \in [d]} \frac{(dk)! (1 - \frac{r}{t})^{dk} (t-r)^{-dr + \sum_{i=1}^r m_i}}{(d(k-r) + \sum_{i=1}^r m_i)! \prod_{i=1}^r (d-m_i)!} \end{aligned}$$

Note that  $(1 - \frac{r}{t})^{dk}$  causes the exponential decay of  $c_{r,k,d,t}$ , since  $\frac{(dk)!}{(d(k-r) + \sum_{i=1}^r m_i)!}$  and  $\binom{k}{r}$  are polynomial in  $k$  and the rest are independent of  $k$ .  $\square$

## Implementation

In this section, we describe our extension to GPUPoly to support our bound propagations. We focus on single-channel inputs. The extension to multi-channel inputs is very similar.

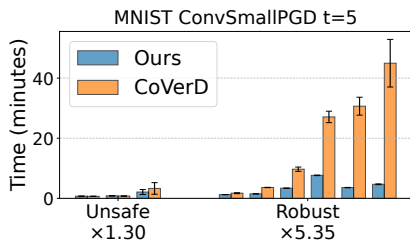


Figure 6: CoVerD +  $t$ -times-top-GP compared to CoVerD on a less challenging benchmark.

Recall that GPUPoly computes bounds over the box domain, where the lower bound is  $l = \min_{y \in \mathcal{D}} \sum_{i=1}^k w_i y_i = \sum_{i=1}^k w_i \bar{x}_i + \sum_{i=1}^k d_i^-$  (the upper bound is analogous). The CUDA kernel distributes the set of indices  $[k]$  across threads. Each thread computes, for every assigned index  $i$ , the value  $\min(w_i a_i, w_i b_i) = w_i \bar{x}_i + d_i^-$  and sums these values over its assigned indices. These sums from all threads are then combined using a tree-based reduction scheme to obtain the lower bound. The GPU’s shared memory is used to share the intermediate sums between threads.

For  $\tilde{\mathcal{B}}_1^t(x)$ , we remind that the lower bound is  $l = \min_{y \in \tilde{\mathcal{B}}_1^t(\bar{x})} \sum_{i=1}^k w_i y_i = \left( \sum_{i=1}^k w_i \bar{x}_i \right) + t \cdot d_{l_1}^-$ . To efficiently compute the sum  $w_i \bar{x}_i$  and add  $t$  times the smallest  $d_i^-$ , we modify the CUDA kernel so that each thread computes the sum  $w_i \bar{x}_i$  over its assigned indices and records the minimum  $d_i^-$  over its assigned indices. In the tree-based reduction scheme, it combines the sums and propagates the minimum  $d_i^-$  observed so far. Finally, the lower bound is obtained by adding  $t$  times the minimum  $d_i^-$  to the accumulated sum. Synchronizing the minimum  $d_i^-$  across threads doubles the required shared memory in the reduction scheme.

For the  $\ell_0$ -ball, the lower bound is  $l = \min_{y \in \mathcal{B}_0^t(\bar{x})} \sum_{i=1}^k w_i y_i = \left( \sum_{i=1}^k w_i \bar{x}_i \right) + (d_{l_1}^- + \dots + d_{l_t}^-)$ . As in the previous case, each thread records the  $t$  smallest  $d_i^-$  values it encounters and selects the values to propagate at each reduction step. This computation increases the complexity of each synchronization step to  $\mathcal{O}(t)$  (instead of  $\mathcal{O}(1)$ ). If we kept double precision, this would increase the shared memory usage by a factor of  $t + 1$ . Instead, our implementation propagates the  $d_i^-$  values as floats (instead of doubles).

We note that, besides the aforementioned changes, our implementation does not modify GPUPoly, thus benefiting from its optimizations. Additionally, the resource requirements remain comparable to those of GPUPoly.

## Additional Results

In this section, we provide additional results.

Figure 6 and Figure 8 extend Figure 5 with CoVerD’s less challenging benchmarks (smaller values of  $t$  than in Figure 5). For a single  $\ell_0$ -ball, two runs of CoVerD out of three reached the timeout (denoted by \*\*) and as in figure Fig-

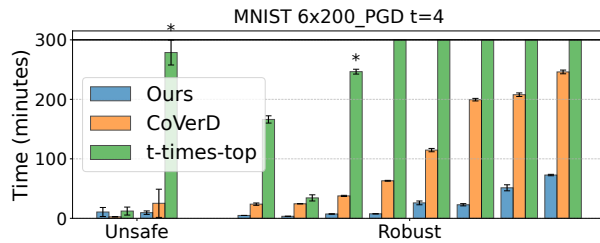


Figure 7: CoVerD with  $t$ -times-top-GP.

ure 5, these runs are excluded. The figures show that our approach significantly boosts CoVerD.

Figure 7 compares CoVerD +  $t$ -times-top-GP with our approach and CoVerD on the MNIST 6x200\_PGD network. For two  $\ell_0$ -balls, one run of CoVerD +  $t$ -times-top-GP out of three reached the timeout (denoted by \*). It shows that not only our approach is significantly better, but also CoVerD is better than CoVerD +  $t$ -times-top-GP, which is consistent with Figure 4.

Figure 9 shows the average running time of GPUPoly and  $t$ -times-top-GP for the experiments of Figure 4. It shows that our  $t$ -times-top-GP does not introduce significant overhead and is sometimes faster than GPUPoly, for our benchmarks.

Table 2 and Table 3 provide a detailed list of our experimental results shown in Figure 5. For every dataset, network,  $t$ , seed (for fixing CoVerD’s random choices), and image index (from the dataset’s testset), it shows the analysis time of CoVerD and CoVerD +  $t$ -times-top-GP in minutes. The notation ‘-’ indicates that an incorrectly classified image (thus not analyzed), ‘\*’ denotes that the verification determined that the  $\epsilon$ -ball is not robust, and ‘TO’ indicates a timeout.

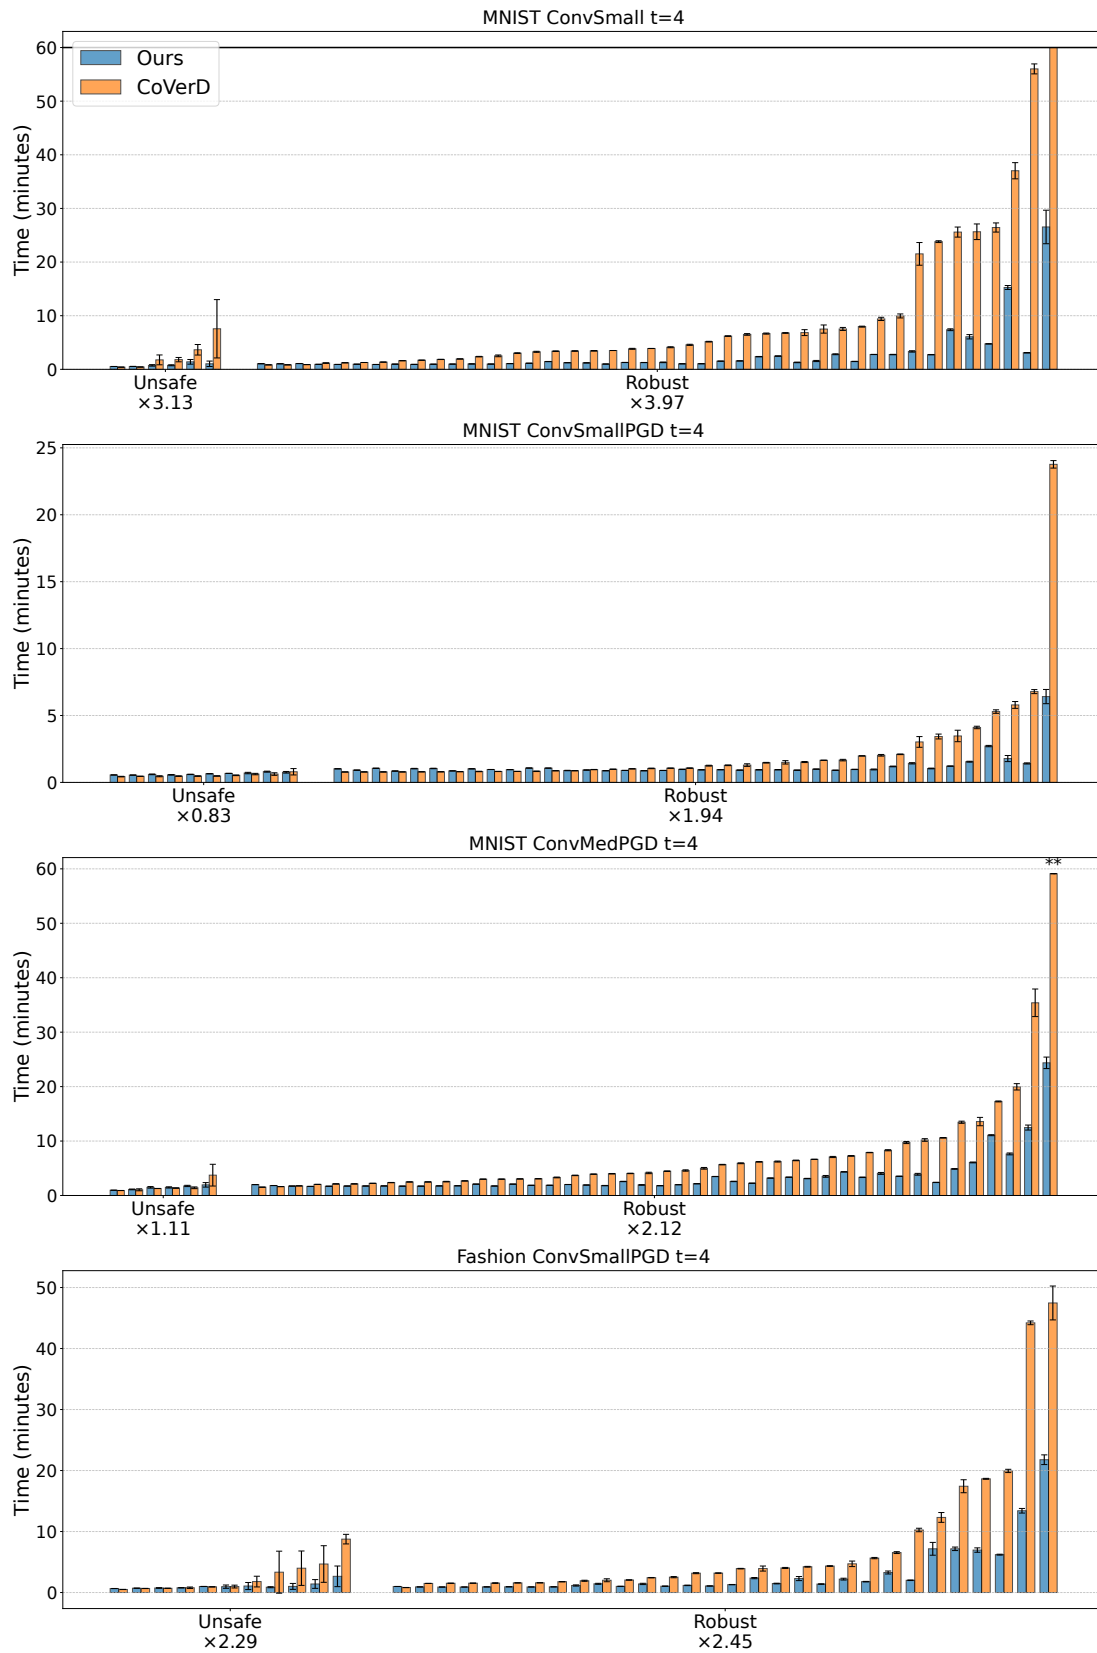


Figure 8: CoVerD +  $\text{top-}t\text{-GP}$  compared to CoVerD on less challenging benchmarks.

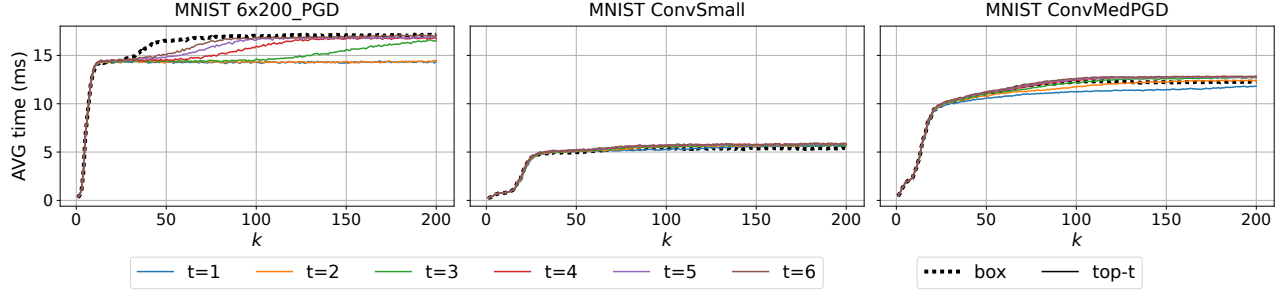


Figure 9: The mean running time of GPUPoly and  $t_{\text{op-t-GP}}$  on neighborhoods capturing perturbations to subsets of  $k$  pixels.

Dataset	Network	$t$	Seed	Image Index										
				0	1	2	3	4	5	6	7	8	9	
MNIST	$6 \times 200\_PGD$	4	42	CoVerD	24.3	36.5	205.7	23.8	3.0*	62.1	112.9	246.8	198.0	12.8*
			Ours	3.4	6.8	53.5	4.8	2.2*	7.8	22.0	74.4	25.5	13.5*	
		41	CoVerD	24.5	38.9	212.1	21.5	2.7*	63.6	112.5	241.9	202.5	4.6*	
			Ours	3.5	7.6	44.3	4.8	9.0*	7.7	29.8	72.2	21.4	6.7*	
		40	CoVerD	24.5	37.7	205.6	26.3	2.2*	63.6	118.6	249.5	197.8	58.4*	
			Ours	3.4	7.5	56.2	4.9	20.7*	7.3	26.1	71.7	21.6	8.3*	
	ConvSmall	5	42	CoVerD	15.1	29.4	TO	3.8	0.9*	TO	TO	104.7	-	104.0
			Ours	3.5	3.9	72.7	1.2	0.7*	128.0	117.9	12.2	-	9.3	
		41	CoVerD	15.3	30.2	TO	4.0	0.7*	TO	TO	87.0	-	62.5	
			Ours	3.4	3.9	72.8	1.2	0.8*	102.3	114.8	12.0	-	9.5	
		40	CoVerD	15.0	31.9	TO	3.8	1.5*	TO	TO	69.9	-	62.1	
			Ours	3.5	4.0	72.3	1.2	0.9*	135.8	148.4	12.1	-	9.4	
ConvSmallPGD	6	42	CoVerD	47.3	TO	TO	27.1	4.0*	213.0	TO	2.2*	-	2.2*	
		Ours	9.9	59.4	193.5	3.0	2.6*	42.0	116.0	2.5*	-	2.6*		
	41	CoVerD	50.3	TO	TO	65.6	1.5*	226.4	TO	2.3*	-	2.0*		
		Ours	10.5	104.7	154.8	2.9	2.0*	38.7	105.5	2.4*	-	2.3*		
	40	CoVerD	61.9	TO	TO	14.6	1.6*	TO	TO	2.4*	-	2.1*		
		Ours	9.3	104.6	163.1	2.9	4.7*	38.1	201.2	2.4*	-	2.4*		
ConvMedPGD	5	42	CoVerD	37.5	91.0	111.8	8.1	1.2*	38.7	136.4	213.4	1.2*	100.2	
		Ours	14.7	15.9	58.5	2.4	1.4*	17.6	47.9	62.4	1.4*	21.6		
	41	CoVerD	38.7	87.0	111.3	7.8	1.2*	37.8	136.5	215.6	1.3*	97.8		
		Ours	14.2	15.8	60.2	2.3	1.4*	17.6	46.2	63.0	1.3*	27.3		
	40	CoVerD	37.6	86.8	111.2	7.8	1.2*	38.2	137.0	240.8	1.3*	99.9		
		Ours	14.0	19.1	59.8	2.3	1.5*	17.5	57.5	64.0	1.3*	20.9		
ConvBig	3	42	CoVerD	22.9	20.8	252.0	10.5	36.7	90.8	30.4	18.7	-	13.4	
		Ours	18.6	14.5	216.3	14.0	19.6	72.2	17.8	13.1	-	11.2		
	41	CoVerD	22.9	20.5	250.8	10.5	36.7	90.4	31.0	18.3	-	13.3		
		Ours	18.5	14.4	218.0	14.1	20.3	71.2	17.6	13.2	-	11.3		
	40	CoVerD	23.0	20.7	250.1	10.4	37.0	90.3	30.8	18.3	-	13.3		
		Ours	18.7	14.4	216.6	14.1	19.6	71.7	17.9	13.2	-	11.4		

Table 2: A detailed list of our experimental results on the MNIST dataset.

Dataset	Network	$t$	Seed	Image Index										
				0	1	2	3	4	5	6	7	8	9	
Fashion	ConvSmallPGD	5	42	CoVerD	19.6	69.5	TO	-	15.0*	-	21.6	0.8*	22.0	25.3
				Ours	3.5	30.8	TO	-	18.6*	-	5.0	0.9*	4.7	5.5
		41	CoVerD	20.9	81.8	TO	-	96.8*	-	21.1	0.8*	20.7	25.0	
			Ours	3.4	37.2	TO	-	22.7*	-	4.9	1.0*	4.1	5.5	
		40	CoVerD	20.3	84.2	TO	-	134.5*	-	20.9	0.8*	21.6	24.7	
			Ours	3.4	34.5	TO	-	5.9*	-	5.0	0.9*	4.1	5.5	
	ConvMedPGD	4	42	CoVerD	17.5	8.8	98.4	3.1*	54.4	-	6.3	-	16.4	17.4
				Ours	4.4	2.8	40.7	3.9*	16.6	-	1.9	-	4.8	5.5
		41	CoVerD	17.3	8.9	98.4	2.5*	60.5	-	6.5	-	16.3	17.7	
			Ours	4.4	2.8	40.1	1.2*	15.2	-	1.9	-	4.6	5.5	
		40	CoVerD	17.2	8.8	99.1	2.9*	58.1	-	6.4	-	16.1	18.1	
			Ours	4.4	2.8	40.3	3.9*	14.8	-	1.9	-	4.8	5.6	
CIFAR-10	ConvSmallPGD	3	42	CoVerD	41.4	42.0	50.3	3.5*	5.7*	184.1	0.6*	224.9	5.3*	-
				Ours	28.3	26.3	28.5	3.9*	5.1*	131.5	0.5*	157.8	7.8*	-
		41	CoVerD	41.7	41.9	46.9	5.7*	3.5*	180.6	0.4*	220.7	3.3*	-	
			Ours	26.0	25.4	27.8	0.9*	3.6*	135.0	0.6*	155.6	9.3*	-	
		40	CoVerD	41.9	41.0	47.0	6.1*	7.3*	178.5	0.5*	220.8	3.6*	-	
			Ours	26.1	25.3	28.6	2.7*	7.9*	130.9	0.5*	155.6	4.0*	-	

Table 3: A detailed list of our experimental results on Fashion-MNIST and CIFAR-10 datasets.