

3DeformRS: Certifying Spatial Deformations on Point Clouds

Gabriel Pérez S.^{*,1,2}

Juan C. Pérez^{*,2}

Motasesm Alfarra^{*,2}

Silvio Giancola²

Bernard Ghanem²

¹Universidad Nacional de Colombia, ²King Abdullah University of Science and Technology (KAUST)

¹gaperezs@unal.edu.co

²{juan.perezsantamaria,motasesm.alfarra,silvio.giancola,bernard.ghanem}@kaust.edu.sa

Abstract

3D computer vision models are commonly used in security-critical applications such as autonomous driving and surgical robotics. Emerging concerns over the robustness of these models against real-world deformations must be addressed practically and reliably. In this work, we propose 3DeformRS, a method to certify the robustness of point cloud Deep Neural Networks (DNNs) against real-world deformations. We developed 3DeformRS by building upon recent work that generalized Randomized Smoothing (RS) from pixel-intensity perturbations to vector-field deformations. In particular, we specialized RS to certify DNNs against parameterized deformations (e.g. rotation, twisting), while enjoying practical computational costs. We leverage the virtues of 3DeformRS to conduct a comprehensive empirical study on the certified robustness of four representative point cloud DNNs on two datasets and against seven different deformations. Compared to previous approaches for certifying point cloud DNNs, 3DeformRS is fast, scales well with point cloud size, and provides comparable-to-better certificates. For instance, when certifying a plain PointNet against a 3° z -rotation on 1024-point clouds, 3DeformRS grants a certificate $3\times$ larger and $20\times$ faster than previous work¹.

1. Introduction

Perception of 3D scenes plays a critical role in applications such as autonomous navigation [14] and robotics [32]. The success of autonomous driving cars and surgical robots largely depends on their ability to understand the surrounding scenes. Recent works [36, 37, 46, 51] demonstrated the capability of Deep Neural Networks (DNNs) to process 3D point clouds. These advances allowed DNNs to achieve exceptional performance in challenging 3D tasks,

* Denotes Equal Contribution. The order is random.

Work done during Gabriel’s internship at KAUST.

¹Code: <https://github.com/gaperezs/3DeformRS>

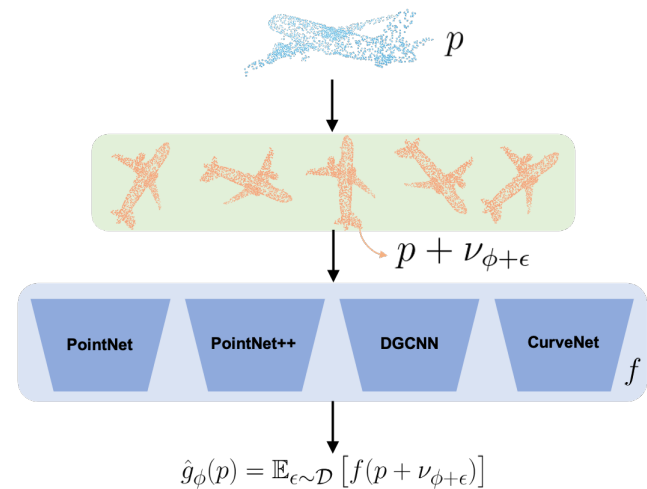


Figure 1. **Randomized Smoothing for Point Cloud DNNs.** We certify the prediction of a DNN on the input point cloud via randomized smoothing by constructing a *smooth* DNN \hat{g}_ϕ around the original DNN f . For an input p , a parametric transformation $\nu_{\phi+\epsilon}$, and a parameter σ , the Smooth DNN predicts the expected value of the predictions from transformed versions of p .

such as shape classification [45, 49], object detection [17] and semantic segmentation [3, 12]. Despite the superior performance of DNNs over traditional approaches, explaining the logic behind their decisions and understanding their limitations is extremely challenging. For instance, earlier works [22, 50] elucidated the limitations of point cloud DNNs to withstand tiny and meaningless perturbations in their input, known as adversarial attacks. The vulnerability of point cloud DNNs against adversarial attacks underscores the importance of considering their robustness for security-critical applications.

The study of adversarial robustness in the image domain demonstrated that defenses [20] are often later broken by more sophisticated attacks [4, 8, 11]. This phenomenon poses a difficulty for evaluating robustness [7]. Such difficulty found response in the domain of certified robustness,



Figure 2. **Visualization of Spatially-Deformed Point Clouds.** In blue the original point cloud and in red its deformed version. *From left to right:* Rotation (plane), Translation (plane), Shearing (bottle), Tapering (bottle) and Twisting (table).

wherein a defense’s robustness is theoretically guaranteed and thus independent of the attack’s sophistication. Recent methods for certifying DNNs on images have spurred interest in extending such methods to point cloud DNNs. In particular, the seminal work of Liu *et al.* [27] certified point cloud DNNs against modification, addition, and deletion of points. Furthermore, Lorenz *et al.* [28] recently proposed a verification approach to certify against 3D transformations.

In this work, we propose 3DeformRS, a certification approach for point cloud DNNs against spatial and deformable transformations. Figure 1 provides an overview of 3DeformRS’s inner workings. Our approach builds upon the theoretical background of Randomized Smoothing (RS) [10]. Specifically, we build 3DeformRS by leveraging DeformRS [1], an RS reformulation that generalized from pixel-intensity perturbations to vector field deformations, and specializing it to point cloud data. In contrast to previous approaches, our work considers spatial deformations on any point cloud DNN, providing efficiency and practicality. Compared to previous work, we find that 3DeformRS produces comparable-to-better certificates, while requiring significantly less computation time. We thus build on 3DeformRS’ virtues and provide a comprehensive empirical study on the certified robustness of four representative point cloud DNNs (PointNet [36], PointNet++ [37], DGCNN [46], and CurveNet [51]) on two datasets (ModelNet40 [49] and ScanObjectNN [45]), and against seven different types of spatial deformations (Rotation, Translation, Affine, Twisting, Tapering, Shearing, and Gaussian Noise).

Contributions. In summary, our contributions are three-fold: **(i)** We propose 3DeformRS by extending Randomized Smoothing (RS) from certifying image deformations to point cloud transformations. We further show that RS’ classical formulation for certifying input perturbations can be seen as a special case of 3DeformRS. **(ii)** We conduct a comprehensive empirical study with 3DeformRS, where we assess the certified robustness of four point cloud DNNs, on two classification datasets and against seven spatial deformations. **(iii)** We compare 3DeformRS with an earlier point cloud certification approach in certification and runtime. We show that 3DeformRS delivers consistent im-

provements while inheriting RS’ scalability and efficiency.

2. Related Work

3D Computer Vision. Images enjoy a canonical representation as fixed-sized matrices; in contrast, several representations exist for 3D data. Such representations include point clouds [36, 37, 51], meshes [15, 25, 40], voxels [9, 38, 56], multi-view [21, 41, 47] and implicit representations [33–35]. Given the prevalence and practicality of point clouds, we focus our attention on this representation and the associated DNNs that process it. PointNet [36] was the first successful attempt at using DNNs on point clouds. This architecture introduced a point-wise MLP with a global set pooling, and achieved remarkable performance in shape classification and semantic segmentation. PointNet++ [37] then introduced intermediate pooling operations in point clouds for local neighborhood aggregation. Afterwards, DGCNN [46] modeled convolutional operations in point clouds based on dynamically-generated graphs between closest point features. Recently, CurveNet [51] learned point sequences for local aggregation, achieving state-of-the-art performance on 3D computer vision tasks. In this work, we conduct a comprehensive empirical study on certified robustness by analyzing the robustness of these four point cloud DNNs.

Robustness. Szegedy *et al.*’s [43] seminal work exposed the vulnerability of DNNs against small input modifications, now known as adversarial examples. Later works observed the pervasiveness of this phenomenon [8, 18], spurring an arms race between defenses that enhanced DNNs’ adversarial robustness [20, 31, 52] and attacks that could break such defenses [4, 8]. The conflict between ever-more complex defenses and attacks also incited interest towards “certified robustness” [48], wherein defenses enjoy theoretical guarantees about the inexistence of adversarial examples that could fool them. A set of works focused on exact verification [19, 54], while others considered probabilistic certification [24, 26]. Randomized Smoothing (RS) [10] has emerged as a certification approach from the probabilistic paradigm that scales well with models and datasets. Notably, RS has been successfully combined with adversarial training [39], regularization [53], and

smoothing-parameters’ optimization [2, 13]. Recently, DeformRS [1] reformulated RS to consider general parameterized vector-field deformations. In this work, we develop 3DeformRS by specializing and extending DeformRS to certify point cloud DNNs against spatial deformations.

Certification on 3D Point Clouds. Since the seminal work of Xiang *et al.* [50] attacked point cloud DNNs, several works studied the robustness of such DNNs [22, 23, 29, 30, 42, 44, 50, 55]. Despite growing interest in the robustness of point cloud DNNs, only two works have addressed their certification. PointGuard [27] provided tight robustness guarantees against modification, addition, and deletion of points. 3DCertify [28] generalized DeepG [5] to 3D point clouds and proposed 3DCertify, a verification approach to certify robustness against common 3D transformations. PointGuard has the benefit of low computational cost (compared to that of exact verification), but does not allow for spatial deformations. Conversely, 3DCertify considers such transformations, but suffers from impractical computational costs. In contrast, our 3DeformRS approach combines the best of both worlds, thus allowing for spatial deformations while enjoying low computational cost.

3. Our approach: 3DeformRS

We present 3DeformRS, a probabilistic certification for point cloud DNNs against spatial deformations.

Preliminaries. Our methodology builds on Randomized Smoothing (RS) [10], arguably the most scalable DNN-certification method. Given a classifier $f : \mathbb{R}^d \rightarrow \mathcal{P}(\mathcal{Y})$ that maps an input $x \in \mathbb{R}^d$ to the probability simplex $\mathcal{P}(\mathcal{Y})$, RS constructs a smooth classifier g that outputs the most probable class when f ’s input is subjected to additive perturbations sampled from a distribution \mathcal{D} . While RS certified against additive pixel perturbations in images, DeformRS [1] extended this formulation to certify against image *deformations* by proposing a *parametric-domain smooth classifier*. Specifically, given image x with coordinates p , a parametric deformation function ν_ϕ with parameters ϕ (e.g. if ν is a rotation, then ϕ is the rotation angle), and an interpolation function I_T , DeformRS defined a parametric-domain smooth classifier

$$g_\phi(x, p) = \mathbb{E}_{\epsilon \sim \mathcal{D}} [f(I_T(x, p + \nu_{\phi+\epsilon}))]. \quad (1)$$

In a nutshell, g outputs f ’s average prediction over transformed versions of x . Note that, in contrast to RS, this formulation deforms pixel *locations* instead of their intensity. DeformRS showed that parametric-domain smooth classifiers are certifiably robust against perturbations to the deformation function’s parameters via the following corollary.

Corollary 1 (restated from [1]). *Suppose g assigns class c_A to an input x , i.e. $c_A = \arg \max_i g_\phi^i(x, p)$ with*

$$p_A = g_\phi^{c_A}(x, p) \text{ and } p_B = \max_{c \neq c_A} g_\phi^c(x, p)$$

Then $\arg \max_c g_{\phi+\delta}^c(x, p) = c_A$ for all parametric perturbations δ satisfying:

$$\begin{aligned} \|\delta\|_1 &\leq \lambda (p_A - p_B) && \text{for } \mathcal{D} = \mathcal{U}[-\lambda, \lambda], \\ \|\delta\|_2 &\leq \frac{\sigma}{2} (\Phi^{-1}(p_A) - \Phi^{-1}(p_B)) && \text{for } \mathcal{D} = \mathcal{N}(0, \sigma^2 I). \end{aligned}$$

In short, Corollary (1) states that, as long as the parameters of the deformation function (e.g. rotation angle) are perturbed by a quantity upper bounded by the certified radius, g_ϕ ’s prediction will remain constant. This result allowed certification against various image deformations. In this work, we build upon DeformRS and specialize it to certify against spatial deformations in point clouds.

3.1. Parametric Certification for 3D DNNs

We specialize the result in Corollary (1) to point clouds. In this setup, $p \in \mathbb{R}^{N \times 3}$ is a point cloud consisting of N 3-dimensional points. We highlight two key differences between certifying images and point clouds. **(i)** The interpolation function I_T , while essential in images for the pixels’ discrete locations, is irrelevant for 3D coordinates and so can be omitted. **(ii)** Most recent DNNs neglect the points’ color information and *exclusively* rely on the points’ location. We combine these observations and modify the parametric domain smooth classifier from Eq. (1) to propose

$$\hat{g}_\phi(p) = \mathbb{E}_{\epsilon \sim \mathcal{D}} [f(p + \nu_{\phi+\epsilon})]. \quad (2)$$

Our \hat{g}_ϕ , inheriting g_ϕ ’s structure, is certifiable against parametric perturbations via Corollary (1). Note that, since point cloud DNNs exclusively rely on location, our *parametric* certification is thus equivalent to *input* certification.

To consider input deformations, let us study the general case where $\nu_\phi = \phi \in \mathbb{R}^{N \times 3}$. Under this setup, and by setting $p' = p + \phi$, our smooth classifier becomes:

$$\tilde{g}_\phi(p) = \mathbb{E}_{\epsilon \sim \mathcal{D}} [f(p' + \epsilon)]. \quad (3)$$

The classifier in Eq. (3) is a general case of both the domain smooth classifier [1] and the input smooth classifier proposed earlier in [10]. Note that this expression elucidates how parametric certification of point cloud DNNs is equivalent to input certification. At last, we note that the smooth classifier in Eq. (3) is also certifiable via Corollary 1. We highlight here that directly certifying \tilde{g} against parametric transformations (e.g. rotation) will perform poorly, as empirically observed in [27]. However, such deficient performance is not an inherent weakness of RS for certifying point cloud transformations, but rather due to a sub-optimal formulation of RS for spatial deformations. We argue that our formulation, i.e. the parametric domain smooth classifier in Eq. (2), is more suitable for modeling spatial deformations than the one presented in [27].

Next, we outline various spatial deformations we consider for assessing the robustness point cloud DNNs.

Name	Flow (\tilde{p})	ϕ	Name	Flow (\tilde{p})	ϕ	Name	Flow (\tilde{p})	ϕ
Translation	$\tilde{x} = t_x$ $\tilde{y} = t_y$ $\tilde{z} = t_z$	$\begin{bmatrix} t_x \\ t_y \\ t_z \end{bmatrix}$	z - Rotation	$\tilde{x} = (c_\gamma - 1)x - s_\gamma y$ $\tilde{y} = s_\gamma x + (c_\gamma - 1)y$ $\tilde{z} = 0$	$[\gamma]$	Affine	$\tilde{x} = ax + by + cz + d$ $\tilde{y} = ex + fy + gz + h$ $\tilde{z} = ix + jy + kz + l$	$\begin{bmatrix} a \\ \vdots \\ l \end{bmatrix}$
z - Shearing	$\tilde{x} = az$ $\tilde{y} = bz$ $\tilde{z} = 0$	$\begin{bmatrix} a \\ b \end{bmatrix}$	z - Twisting	$\tilde{x} = (c_{\gamma z} - 1)x - s_{\gamma z} y$ $\tilde{y} = s_{\gamma z} x + (c_{\gamma z} - 1)y$ $\tilde{z} = 0$	$[\gamma]$	z - Tapering	$\tilde{x} = (\frac{1}{2}a^2 + b)zx$ $\tilde{y} = (\frac{1}{2}a^2 + b)zy$ $\tilde{z} = 0$	$\begin{bmatrix} a \\ b \end{bmatrix}$

Table 1. **Per-point Deformation Flows** for semantically-viable spatial deformations. Convention: $c_\alpha = \cos(\alpha)$ and $s_\alpha = \sin(\alpha)$. Without loss of generality, we only show the rotation around the z -axis, and leave the formulation of other rotations to the **Appendix**.

3.2. Modeling Spatial Deformations

We now detail the spatial deformations we consider to assess the robustness of point cloud DNNs. Our formulation from Eq. (3) requires modeling spatial deformations as additive perturbations on the points’ coordinates. Thus, given a point cloud p that is transformed to yield p' , we define the flow that additively perturbs p as $\tilde{p} = p' - p$, where $\tilde{p} \in \mathbb{R}^{N \times 3}$. Modeling transformations is then equivalent to modeling the induced per-point flow. Thus, we are required to model each deformation via a *parametric flow*, whose parameters are those of the corresponding transformation.

We consider six parameterizable flows, corresponding to four linear and two nonlinear transformations. In particular, we consider four linear, rigid and deformable transformations: (1) rotation, (2) translation, (3) shearing, and (4) the general affine transformation. Furthermore, we follow prior work [28] and consider two nonlinear transformations: (5) tapering and (6) twisting. As a summary, we report all the deformation flows and their corresponding parameters in Table 1, and visualize some of their effects in Figure 2.

Note that the affine transformation is the most general transformation we consider, capable of modeling any linear transformation. Presumably, a DNN enjoying certified robustness against affine transformations would also enjoy robustness against combinations of the other spatial deformations. We leave formulations of the aforementioned transformations in homogeneous coordinates to the **Appendix**.

Gaussian Noise. In addition to the above transformations, we also consider Gaussian noise perturbations. Notably, this noise deforms the underlying vector field and, thus, is the most general perturbation. To certify against Gaussian noise, we construct the smooth classifier in Eq. (3) with $\epsilon \in \mathbb{R}^{N \times 3}$ and $\mathcal{D} = \mathcal{N}(0, \sigma^2 I)$. While this deformation is very general, the high dimensionality of the certified parameters limits its applicability to imperceptible perturbations. Nevertheless, as adversaries may take such form [50], we consider this perturbation in our experiments.

Design Choices. For each spatial deformation and point cloud DNN, we assess certified robustness by constructing the smooth classifier from Eq. (2). In deformations whose

parameter space is bounded (e.g. rotations, where angles beyond $\pm\pi$ radians are redundant), we smooth with a uniform distribution and thus obtain an ℓ_1 certificate. For the remaining deformations, we employ Gaussian smoothing and thus obtain an ℓ_2 certificate.

4. Experiments

4.1. Setup

Datasets. We experiment on ModelNet40 [49] and ScanObjectNN [45]. ModelNet40 comprises 12,311 3D CAD models from 40 classes. ScanObjectNN has 2,902 real-world 3D scans of 15 classes. While ModelNet40 has synthetic objects, ScanObjectNN used 3D cameras and so contains natural and challenging self-occlusions. We use the ScanObjectNN variant that omits background data. For both datasets, the shapes are dimension-normalized, origin-centered and sampled with 1024 points.

Models. We study four point cloud DNNs: PointNet [36], PointNet++ [37], DGCNN [46] and CurveNet [51]. For PointNet, we used 3DCertify’s implementation, which uses z -rotation augmentation on ModelNet40. For CurveNet, we used the official implementation’s weights, trained on ModelNet40 with axis-independent $[.66, 1.5]$ scaling and ± 0.2 translation. On ScanObjectNN, PointNet and CurveNet are trained without augmentation. For PointNet++ and DGCNN, we used the PyTorch Geometric library [16] with default hyper-parameters and without augmentations.

Certification. In our experiments, following [1, 39, 53], we construct the hard version of the smooth classifier in Eq. (2) for each deformation. Moreover, we follow common practice and adapt the public implementation from [10] for estimating the certified radius via Monte Carlo sampling. In particular, we use 1,000 samples to estimate the certified radius with a probability of failure of 10^{-3} . For all experiments, we provide envelope certified accuracy curves cross validated at several values of smoothing parameters, detailed in the **Appendix**. Since all deformations used Gaussian smoothing, the certificates we find are in the ℓ_2 sense (except for rotation, which used uniform smoothing and so its certificate is in the ℓ_1 sense).

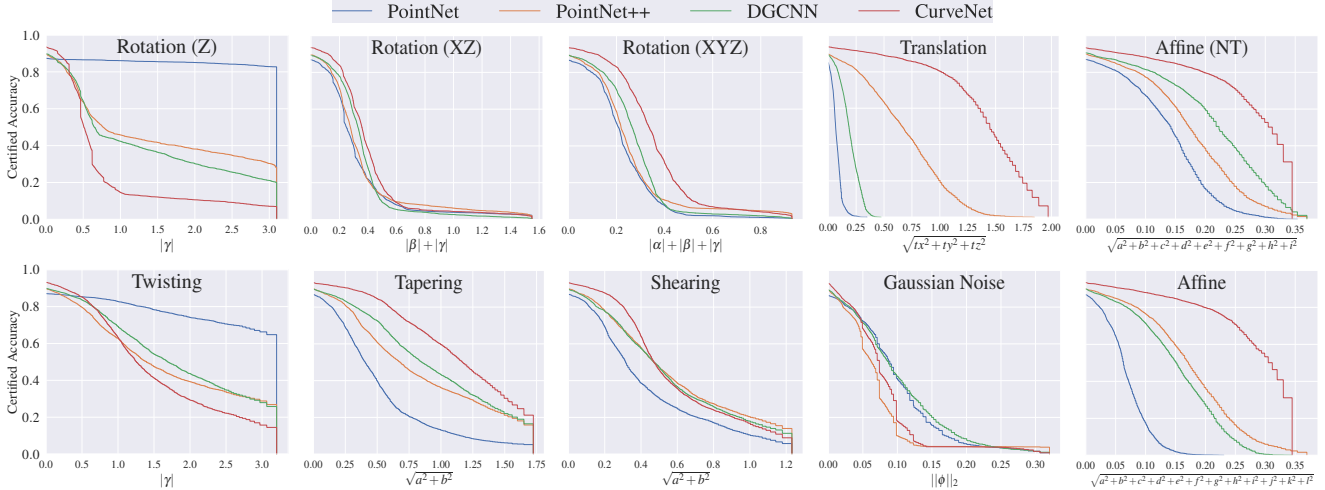


Figure 3. 3DDeformRS Certification on ModelNet40 of four point cloud DNNs against 10 transformations.

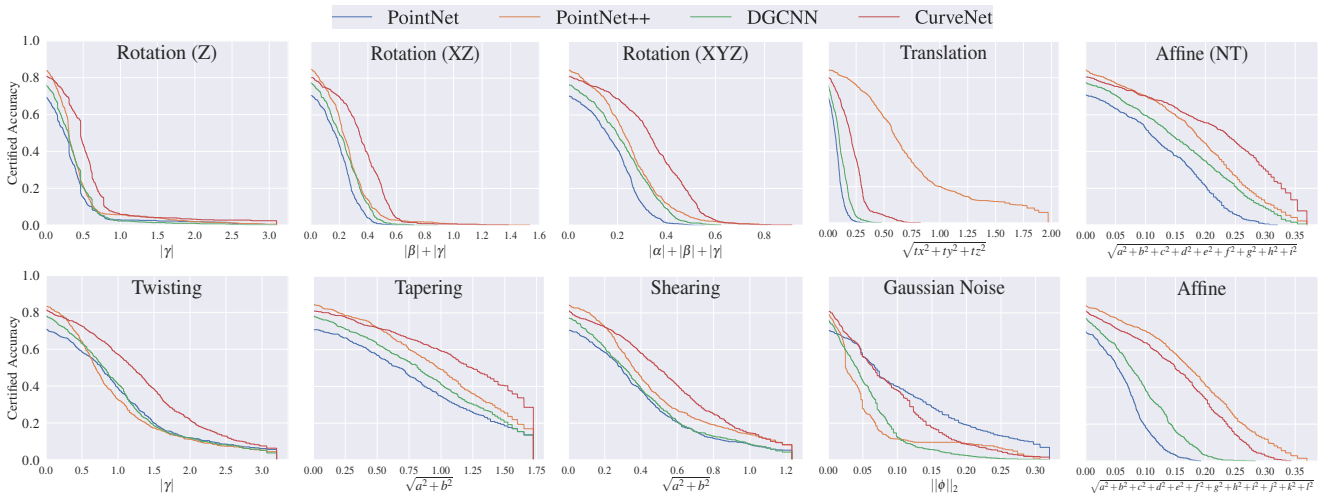


Figure 4. 3DDeformRS Certification on ScanObjectNN of four point cloud DNNs against 10 transformations.

4.2. Benchmarking 3D Networks

We present certification curves in Figures 3 and 4 for ModelNet40 and ScanObjectNN, respectively. We also report each curve’s associated Average Certification Radius (ACR) in Table 2 as a summary metric. We further highlight nine main observations from these results, and leave a detailed analysis, together with ablations on RS’ hyperparameters, to the **Appendix**.

Vulnerability against Rigid Transformations. Our analysis considers, among others, rotation and translation. Remarkably, we find that DNNs are significantly vulnerable even against these rigid perturbations. For most DNNs, the certified accuracy plots in Figures 3 and 4 show that performance drops dramatically as the perturbation’s magnitude increases. This phenomenon supports further research on increasing the robustness of point cloud DNNs against sim-

ple transformations that could happen in the real world.

Deformation Complexity. Each spatial deformation is parameterized by a certain number of values. The number of parameters can be associated with the deformation’s complexity: “simple” deformations, *i.e.* rotation around the x axis, require few parameters, while “complex” deformations, *i.e.* affine, require several parameters. Under these notions, we notice that as the deformation’s complexity increases, the DNNs’ certified accuracies drop more rapidly. This observation agrees with intuition: complex transformations should be harder to resist than simple ones.

Gaussian Noise. This deformation is arguably the most general and information-destroying, as it may not preserve distances, angles nor parallelism. Experimentally, we indeed observe that DNNs are rather brittle against this noise: even small magnitudes can break the DNNs’ performance.

ModelNet40	z -Rot.	xz -Rot.	xyz -Rot.	Shearing	Twisting	Tapering	Translation	Affine	Affine (NT)	Gaussian Noise
PointNet [36]	2.64	<u>0.31</u>	<u>0.21</u>	0.42	2.48	<u>0.50</u>	<u>0.07</u>	<u>0.06</u>	<u>0.13</u>	0.09
PointNet++ [37]	1.45	0.34	0.25	0.55	1.66	0.81	0.64	0.16	0.17	<u>0.07</u>
DGCNN [46]	1.29	0.33	0.26	0.54	1.78	0.90	0.18	0.14	0.20	0.10
CurveNet [51]	<u>0.76</u>	0.39	0.34	0.56	<u>1.51</u>	1.09	1.32	0.26	0.26	0.08

ScanObjectNN	z -Rot.	xz -Rot.	xyz -Rot.	Shearing	Twisting	Tapering	Translation	Affine	Affine (NT)	Gaussian Noise
PointNet [36]	<u>0.30</u>	<u>0.16</u>	<u>0.15</u>	<u>0.35</u>	0.90	<u>0.72</u>	<u>0.06</u>	<u>0.05</u>	<u>0.11</u>	0.10
PointNet++ [37]	0.38	0.23	0.22	0.45	<u>0.89</u>	0.93	0.67	0.16	0.17	<u>0.05</u>
DGCNN [46]	0.32	0.20	0.19	0.37	0.92	0.81	0.10	0.08	0.14	<u>0.05</u>
CurveNet [51]	0.51	0.31	0.29	0.51	1.24	1.02	0.19	0.14	0.20	0.08

Table 2. **Certified Robustness Assessment on ModelNet40 and ScanObjectNN.** We report the Average Certification Radius (ACR) of all DNNs against 10 deformations, on each dataset. For each deformation, we **embolden** the best performance and underline the worst.

CurveNet Performs Remarkably. In terms of ACR, CurveNet displays larger robustness than competitors across the board. In particular, Table 2 (top), shows that, on ModelNet40, CurveNet achieves the best ACR for seven of the 10 deformations we present, while scoring last only in two deformations (z -Rot and Twisting). Analogously, for ScanObjectNN (Table 2, bottom), CurveNet is the best performer, displaying the best robustness for seven of the 10 deformations we consider, while never scoring last.

PointNet Performs Poorly. On ModelNet40, reported in Table 2 (top), PointNet achieves the lowest ACR values for nine of the 10 deformations, while only achieving the best score in two deformations (z -Rot and Twisting). Similarly, for ScanObjectNN, in Table 2 (bottom), PointNet consistently scores last across all transformations, holding the last position in eight out of 10 transformations, and only holding the first position in only one (Gaussian noise).

Training-time Augmentations Boost Certified Robustness. On ModelNet40, PointNet was trained with z -rotation augmentations. Indeed, we find that, when PointNet is evaluated on ModelNet40, it displays superior performance against rotation-based transformations along the z axis, *i.e.* z -Rot and Twisting. In particular, its ACR is > 0.8 more than the runner-up in z -Rot and ~ 0.4 in Twisting, while its robust accuracy is mostly maintained across the entire rotation regime (from $-\pi$ to $+\pi$ radians, *i.e.* all possible rotations). That is, PointNet correctly classifies most objects, independent of whether they are rotated or twisted around the z axis. More interestingly, we also remark how PointNet’s dominance is *not* observed in ScanObjectNN, where it was *not* trained with z -rotation augmentations. We thus attribute PointNet’s robustness to the training-time augmentations it enjoyed, and so we further study this phenomenon in the next subsection.

Certified vs. Regular Accuracy. Overall, our analysis finds that the best performer is CurveNet, while the worst performer is PointNet. Notably, this fact agrees with each DNN’s plain test performance: CurveNet has an advantage over PointNet of about 7% and 10% in ModelNet40 and

	ModelNet40	ScanObjectNN
PointNet [36]	85.94%	71.35%
PointNet++ [37]	90.03%	83.53%
DGCNN [46]	90.03%	78.04%
CurveNet [51]	93.84%	81.47%

Table 3. **Test Set Accuracy** on ModelNet40 and ScanObjectNN.

ScanObjectNN, respectively. Thus, our results find a correlation between regular and certified performance.

ModelNet40 and ScanObjectNN. CurveNet is, arguably, the best performer on both datasets in terms of ACR. However, we note that *all* of CurveNet’s certified accuracies drop significantly from ModelNet40 to ScanObjectNN. Thus, our results agree with how ModelNet40’s synthetic nature (compared to ScanObjectNN’s realistic nature) implies that ModelNet40 is a “simpler” dataset than ScanObjectNN.

Sizable Variations in Robustness. For a given deformation and dataset, we note sizable robustness differences across DNNs. Specifically, we observe that one DNN may obtain an ACR even $10\times$ larger than other DNN. This phenomenon happens even when the plain accuracy of the DNNs being certified is mostly comparable, as reported in Table 3. Thus, we argue that certified robustness should be a design consideration when developing DNNs. That is, our results suggest that (i) plain accuracy may not provide the full picture into a DNN’s performance, and (ii) certified accuracy may be effective and efficient for assessing DNNs.

4.3. Simple but Effective: Augmented Training

Previously, in Figure 3, we observed PointNet’s dominant performance on z -axis rotation and twisting. We attributed this superiority to augmentations PointNet enjoyed during training. Here we investigate the effect of other training augmentations on certified robustness. In particular, we train two DNNs with four different deformations and then assess their certified robustness against such deformations. We report the results of this experiment in Figure 5.




Figure 5. **Relative effects of training augmentations on Average Certified Accuracy.** We conduct augmentations during training and record relative improvements on certified accuracy. Most training augmentations improve certification across the board.

	Plain	z-Rot.	xyz-Rot.	Trans.	Twist.
PointNet++	90.1%	+0.2%	-0.3%	-0.4%	-0.1%
DGCNN	89.8%	+1.3%	+0.8%	+1.3%	-0.2%

Table 4. **Augmentation impact on accuracy on ModelNet40.**

We draw the following three conclusions from these results. **(i)** Conducting augmentation with one deformation significantly increases the robustness against that same deformation. This is an expected phenomenon, as the model trained on deformed versions of inputs. **(ii)** Training on some deformations yields robustness against other deformations (e.g., augmenting with twisting results in robustness against rotation). This observation aligns with our earlier result where PointNet displayed superiority in z -axis rotation and twisting. This result further suggests that simple training augmentations strategies are effective for “robustifying” models against deformations. **(iii)** The larger certificates come at virtually no cost on clean accuracy: each model’s clean accuracy varied less than $\pm 1.5\%$ w.r.t. the model trained without augmentations (refer to Table 4).

4.4. Comparison with 3DCertify

We further compare 3DeformRS and 3DCertify [28], which is, to the best of our knowledge, the current state-of-the-art certification approach against spatial deformations. 3DCertify’s analysis focuses exclusively on PointNet and ModelNet40. We used their official implementation and pre-trained weights, and note that 3DCertify certifies a 100-instance subset of the test set. We note that the certified accuracies reported in [28] are w.r.t. the correctly-classified samples from such subset. Moreover, we underscore that 3DCertify’s analysis focuses exclusively on PointNet trained on ModelNet40.

Technical Details. We first discuss the intricacies of providing a fair comparison between 3DCertify (an exact verification approach) and 3DeformRS (a probabilistic certifi-

cation approach). We underscore fundamental differences regarding the input each approach receives when certifying **(i)** a given DNN, e.g. PointNet, **(ii)** w.r.t. a desired transformation, e.g. Rotation, on **(iii)** a selected point cloud input, e.g. a plane. In particular, 3DCertify receives as additional input the exact magnitude of the transformation. 3DCertify then uses the transformation and its magnitude, e.g. 3° , to return a boolean value stating if the given DNN’s output is certifiable against a transformation of such magnitude. On the other hand, 3DeformRS mechanism is fundamentally different: it computes a tight certification radius for the same (DNN-transformation-point cloud) tuple, but leveraging a random distribution with which the input was smoothed. The two approaches can be compared if the certified radius provided by 3DeformRS is evaluated against the transformation magnitude received by 3DCertify.

However, we note that the certification radius computed by 3DeformRS depends on the hyper-parameter σ (or λ for uniform smoothing) from Corollary 1. That is, 3DeformRS will compute a different radius for each σ considered. While there exists a σ^* with which 3DeformRS provides the *largest* certified radius for the transformation being considered, σ^* is not known *a priori*. Thus, we are required to run a grid search of σ values to compute each instance’s largest certified radius. While this procedure is computationally intensive, the output of *each* experiment—considering one σ over the whole dataset—provides an entire certified accuracy *curve*, in contrast to a *single point* in this curve, as provided by exact verification approaches. We underscore how this curve can provide insights into a DNN’s robustness. In practice, when comparing with 3DCertify on specific transformation magnitudes, we run a grid search of at most 18 σ values to obtain each instance’s largest certified radius. Hence, we report certified accuracy envelope curves, *i.e.* certified accuracy curves that consider each instance’s largest certified radius. Additionally, to circumvent the problem arisen from considering different norms (3DCertify’s ℓ_∞ norm vs. 3DeformRS’s ℓ_1 or ℓ_2), we limit our comparison to single-parameter transformations.

Results. We compare against 3DCertify, in their reported experimental setup, along three dimensions: **(i)** certificate magnitude, **(ii)** point cloud cardinality and **(iii)** speed. Overall, we find that 3DeformRS provides comparable-to-better certificates, scales well with point cloud size, and has manageable computational cost.

Certificate magnitude. We consider rotations w.r.t. each axis, and evaluate 3DCertify’s official 64-points DNN against rotations of $\gamma \in \{1, 2, 3, 4, 5, 6, 7, 8, 10, 15\}$ degrees. Whenever possible, we used 3DCertify’s reported results, otherwise we used their public implementation to run certification. We find that 3DeformRS achieves comparable-to-better certificates, while providing the additional benefit of full certified accuracy curves, instead of

Points	16	32	64	128	256	512	1024
Boopathy <i>et al.</i> [6]	3.7	3.6	3.3	2.2	4.4	5.6	6.7
DeepPoly	95.1	94.0	91.3	72.2	51.1	39.3	28.1
3DCertify (Taylor 3D)	97.5	94.0	93.5	81.1	66.7	49.4	37.1
3DeformRS (ours)	98.8	97.6	97.8	100	100	97.8	100

Table 5. $3^\circ z$ - **Rotation Certificates when varying Point Cloud Cardinality.** The three previous methods consider linear relaxations and use the DeepPoly verifier, while 3DeformRS is based on Randomized Smoothing. Baselines taken from [28].

individual points in these plots. Refer to the **Appendix** for such certified accuracy curves.

Point Cloud Cardinality. 3DCertify provides exact verification on point clouds of limited size (64 points reported in [28]). However, DNNs may deal with point clouds of at least 1024 points in practice. We experiment on this setup and compare with previous approaches by varying the point clouds’ size and certifying against a $3^\circ z$ -rotation. For this experiment, we follow 3DCertify’s setup [28] and use the same DNN weights for each certification method. Table 5 shows that our approach provides a better certificate for DNNs when trained and tested on large point clouds. In particular, we mark three main observations. First, 3DeformRS performs up to 60% better on the realistic setup of 1024 points. Second, 3DeformRS provides a certified accuracy of over 97% across the board; while large robustness is expected from a model augmented with z -rotations, our approach shows that providing such certificates *is* possible. Third, 3DeformRS breaks the trend of decreasing certificates from which other approaches suffered when handling larger point clouds. Thus, we find that 3DeformRS enjoys scalability and invariance to larger input size.

Speed. Certification via 3DCertify [28] is computationally expensive, and its current official implementation cannot benefit from GPU hardware accelerators. Furthermore, 3DCertify’s computational cost scales with the perturbation’s magnitude, hindering certification against large perturbations. For example, certifying PointNet with 64 points against an z - rotation on a CPU requires ~ 13 k seconds for 1° , and up to ~ 89 k seconds for 10° .

Unlike 3DCertify, 3DeformRS enjoys a virtually-constant computation cost, as it requires a fixed number of forward passes and σ values. In addition, our native implementation can leverage GPUs for accelerating certification. Certifying against z - rotation with 3DeformRS on the *same* CPU requires only ~ 40 seconds per σ (independent of σ ’s magnitude, as reported in Table 6). Even the extreme case of certifying with 100 values for σ , arguably an unnecessary amount, still requires only ~ 4 k seconds. That is, certifying small perturbations with 3DeformRS attains a $\sim 3\times$ speed boost, while large perturbations can even enjoy a $\sim 20\times$ boost. Moreover, accelerating 3DeformRS via an

Device	σ		
	0.01	0.2	0.4
CPU	38.8	39.9	40.4
GPU	7.3	7.4	7.7

Table 6. **Runtimes for 3DeformRS.** We compare certification runtime for a single σ value both in CPU and GPU (values in seconds). 3DeformRS’s CPU version enjoys reasonable certification times, and leveraging a GPU lowers the runtime by $\sim 5\times$.

Nvidia V100 GPU provides a $5\times$ boost over its CPU counterpart, further improving runtime over 3DCertify.

4.5. Ablations

The stochastic nature of 3DeformRS requires a Monte Carlo method followed by a statistical test to bound the probability of returning an incorrect prediction. This statistical test is parameterized by a failure ratio α which, throughout our experiments, was set to the default value of 10^{-3} , following [10]. Here, we analyze the sensitivity of 3DeformRS to other failure probabilities α . We show in the **Appendix** the certified accuracy curves for failure probabilities $\alpha \in \{10^{-2}, 10^{-4}, 10^{-5}\}$. We underscore that, in our assessment, 3DeformRS shows negligible variation in certified accuracy w.r.t. changes in α . In particular, we notice that all ACRs are ~ 0.22 for all the α values we considered.

5. Conclusions and Limitations

In this work, we propose 3DeformRS, a method for certifying point cloud DNNs against spatial deformations. Our method provides comparable-to-better certificates than earlier works while scaling better to large point clouds and enjoying practical computation times. These virtues or 3DeformRS allow us to conduct a comprehensive empirical study of the certified robustness of point cloud DNNs against semantically-viable deformations. Furthermore, 3DeformRS’ practical runtimes may enable its usage in real-world applications. While our stochastic approach is practical with its faster and top-performing certification, its stochasticity may also raise concerns when comparing against exact verification approaches. Moreover, we note that our work solely focused on assessing the 3D robustness of point cloud DNNs against input deformations, disregarding other types of perturbations. Possible avenues for future work include incorporating better training algorithms such as MACER [53] and SmoothAdv [39] for further robustness improvements.

Acknowledgements. This publication is based upon work supported by the King Abdullah University of Science and Technology (KAUST) Office of Sponsored Research (OSR) under Award No. OSR-CRG2019-4033. We would also like to thank Jesús Zarzar for the help and discussions.

References

- [1] Motasem Alfarra, Adel Bibi, Naeemullah Khan, Philip H. S. Torr, and Bernard Ghanem. Deformrs: Certifying input deformations with randomized smoothing. *CoRR*, abs/2107.00996, 2021.
- [2] Motasem Alfarra, Adel Bibi, Philip H. S. Torr, and Bernard Ghanem. Data dependent randomized smoothing, 2020.
- [3] Iro Armeni, Ozan Sener, Amir R. Zamir, Helen Jiang, Ioannis Brilakis, Martin Fischer, and Silvio Savarese. 3d semantic parsing of large-scale indoor spaces. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition*, 2016.
- [4] Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *International Conference on Machine Learning (ICML)*, 2018.
- [5] Mislav Balunovic, Maximilian Baader, Gagandeep Singh, Timon Gehr, and Martin Vechev. Certifying geometric robustness of neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [6] Akhilan Boopathy, Tsui-Wei Weng, Pin-Yu Chen, Sijia Liu, and Luca Daniel. Cnn-cert: An efficient framework for certifying robustness of convolutional neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3240–3247, 2019.
- [7] Nicholas Carlini, Anish Athalye, Nicolas Papernot, Wieland Brendel, Jonas Rauber, Dimitris Tsipras, Ian Goodfellow, Aleksander Madry, and Alexey Kurakin. On evaluating adversarial robustness. *arXiv preprint arXiv:1902.06705*, 2019.
- [8] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, 2017.
- [9] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3075–3084, 2019.
- [10] Jeremy Cohen, Elan Rosenfeld, and Zico Kolter. Certified adversarial robustness via randomized smoothing. In *International Conference on Machine Learning (ICML)*, 2019.
- [11] Francesco Croce and Matthias Hein. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *ICML*, 2020.
- [12] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2017.
- [13] Francisco Eiras, Motasem Alfarra, M. Pawan Kumar, Philip H. S. Torr, Puneet K. Dokania, Bernard Ghanem, and Adel Bibi. ANCER: anisotropic certification via sample-wise volume maximization. *CoRR*, abs/2107.04570, 2021.
- [14] Francesca M Favarò, Nazanin Nader, Sky O Eurich, Michelle Tripp, and Naresh Varadaraju. Examining accident reports involving autonomous vehicles in california. *PLoS one*, 12(9):e0184952, 2017.
- [15] Yutong Feng, Yifan Feng, Haoxuan You, Xibin Zhao, and Yue Gao. Meshnet: Mesh neural network for 3d shape representation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8279–8286, 2019.
- [16] Matthias Fey and Jan E. Lenssen. Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.
- [17] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3354–3361. IEEE, 2012.
- [18] Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations (ICLR)*, 2015.
- [19] Sven Gowal, Krishnamurthy Dvijotham, Robert Stanforth, Rudy Bunel, Chongli Qin, Jonathan Uesato, Relja Arandjelovic, Timothy Mann, and Pushmeet Kohli. On the effectiveness of interval bound propagation for training verifiably robust models. *arXiv preprint arXiv:1810.12715*, 2018.
- [20] Chuan Guo, Mayank Rana, Moustapha Cisse, and Laurens van der Maaten. Countering adversarial images using input transformations. In *International Conference on Learning Representations (ICLR)*, 2018.
- [21] Abdullah Hamdi, Silvio Giancola, and Bernard Ghanem. Mvtn: Multi-view transformation network for 3d shape recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1–11, 2021.
- [22] Abdullah Hamdi, Sara Rojas, Ali Thabet, and Bernard Ghanem. Advpc: Transferable adversarial perturbations on 3d point clouds. In *European Conference on Computer Vision (ECCV)*, 2020.
- [23] Jaeyeon Kim, Binh-Son Hua, Thanh Nguyen, and Sai-Kit Yeung. Minimal adversarial examples for deep learning on 3d point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (CVPR)*, 2021.
- [24] Mathias Lecuyer, Vaggelis Atlidakis, Roxana Geambasu, Daniel Hsu, and Suman Jana. Certified robustness to adversarial examples with differential privacy. In *2019 IEEE Symposium on Security and Privacy (SP)*, 2019.
- [25] Huan Lei, Naveed Akhtar, and Ajmal Mian. Picasso: A cuda-based library for deep learning over 3d meshes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13854–13864, 2021.
- [26] B Li, C Chen, W Wang, and Lawrence Carin. Second-order adversarial attack and certifiable robustness. *arXiv preprint arXiv: 1809.03113*, 2018.
- [27] Hongbin Liu, Jinyuan Jia, and Neil Zhenqiang Gong. Point-guard: Provably robust 3d point cloud classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6186–6195, 2021.
- [28] Tobias Lorenz, Anian Ruoss, Mislav Balunović, Gagandeep Singh, and Martin Vechev. Robustness certification for point cloud models. *arXiv preprint arXiv:2103.16652*, 2021.
- [29] Chengcheng Ma, Weiliang Meng, Baoyuan Wu, Shibiao Xu, and Xiaopeng Zhang. Efficient joint gradient based attack

- against sor defense for 3d point cloud classification. In *Proceedings of the 28th ACM International Conference on Multimedia*, 2020.
- [30] Chengcheng Ma, Weiliang Meng, Baoyuan Wu, Shibiao Xu, and Xiaopeng Zhang. Towards effective adversarial attack against 3d point cloud classification. In *2021 IEEE International Conference on Multimedia and Expo (ICME)*, 2021.
- [31] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations (ICLR)*, 2018.
- [32] Jacques Marescaux, Joel Leroy, Michel Gagner, Francesco Rubino, Didier Mutter, Michel Vix, Steven E Butner, and Michelle K Smith. Transatlantic robot-assisted telesurgery. *Nature*, 413(6854):379–380, 2001.
- [33] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4460–4470, 2019.
- [34] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European conference on computer vision*, pages 405–421. Springer, 2020.
- [35] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2019.
- [36] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017.
- [37] Charles R Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *arXiv preprint arXiv:1706.02413*, 2017.
- [38] Gernot Riegler, Ali Osman Ulusoy, and Andreas Geiger. Octnet: Learning deep 3d representations at high resolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3577–3586, 2017.
- [39] Hadi Salman, Jerry Li, Ilya P Razenshteyn, Pengchuan Zhang, Huan Zhang, Sébastien Bubeck, and Greg Yang. Provably robust deep learning via adversarially trained smoothed classifiers. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [40] Jonas Schult, Francis Engelmann, Theodora Kontogianni, and Bastian Leibe. Dualconvmesh-net: Joint geodesic and euclidean convolutions on 3d meshes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8612–8622, 2020.
- [41] Hang Su, Subhransu Maji, Evangelos Kalogerakis, and Erik Learned-Miller. Multi-view convolutional neural networks for 3d shape recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 945–953, 2015.
- [42] Jiachen Sun, Yulong Cao, Christopher Choy, Zhiding Yu, Anima Anandkumar, Zhuoqing Mao, and Chaowei Xiao. Adversarially robust 3d point cloud recognition using self-supervisions. In *Thirty-Fifth Conference on Neural Information Processing Systems (NeurIPS)*, 2021.
- [43] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *International Conference on Learning Representations (ICLR)*, 2014.
- [44] Tzungyu Tsai, Kaichen Yang, Tsung-Yi Ho, and Yier Jin. Robust adversarial objects against deep learning models. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2020.
- [45] Mikaela Angelina Uy, Quang-Hieu Pham, Binh-Son Hua, Thanh Nguyen, and Sai-Kit Yeung. Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1588–1597, 2019.
- [46] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *Acm Transactions On Graphics (tog)*, 38(5):1–12, 2019.
- [47] Xin Wei, Ruixuan Yu, and Jian Sun. View-gcn: View-based graph convolutional network for 3d shape analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1850–1859, 2020.
- [48] Eric Wong and Zico Kolter. Provable defenses against adversarial examples via the convex outer adversarial polytope. In *International Conference on Machine Learning (ICML)*, 2018.
- [49] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015.
- [50] Chong Xiang, Charles R Qi, and Bo Li. Generating 3d adversarial point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [51] Tiange Xiang, Chaoyi Zhang, Yang Song, Jianhui Yu, and Weidong Cai. Walk in the cloud: Learning curves for point clouds shape analysis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 915–924, October 2021.
- [52] Cihang Xie, Yuxin Wu, Laurens van der Maaten, Alan L. Yuille, and Kaiming He. Feature denoising for improving adversarial robustness. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [53] Runtian Zhai, Chen Dan, Di He, Huan Zhang, Boqing Gong, Pradeep Ravikumar, Cho-Jui Hsieh, and Liwei Wang. Macer: Attack-free and scalable robust training via maximizing certified radius. In *International Conference on Learning Representations (ICLR)*, 2019.
- [54] Huan Zhang, Hongge Chen, Chaowei Xiao, Sven Gowal, Robert Stanforth, Bo Li, Duane Boning, and Cho-Jui Hsieh. Towards stable and efficient training of verifiably robust neu-

ral networks. In *International Conference on Learning Representations (ICLR)*, 2020.

- [55] Hang Zhou, Kejiang Chen, Weiming Zhang, Han Fang, Wenbo Zhou, and Nenghai Yu. Dup-net: Denoiser and up-sampler network for 3d adversarial point clouds defense. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (CVPR)*, 2019.
- [56] Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4490–4499, 2018.