SymAttack: Symmetry-aware Imperceptible Adversarial Attacks on 3D Point Clouds *(Supplementary Material)*

1 MORE QUANTITATIVE RESULTS

Feasibility of Enforcing Symmetry in Whole Shapes. Due to the importance of maintaining symmetry for the imperceptibility of adversarial attacks, we investigate the application of this symmetry constraint to whole shapes. Specifically, after identifying the plane of symmetry, we perform symmetry-aware adjustment on each pair of points that are symmetric with respect to this plane. The experimental results are shown in Tab. 1. It is evident that, due to the imposition of overly strict symmetry constraints, the attack success rate does not reach 100%. Moreover, performance across various imperceptibility metrics is suboptimal. Therefore, these findings underscore the significance of using symmetric element detection to sample part- and patch-level symmetric elements.

Undefendability Performance. We evaluate the robustness of SymAttack against various defense solutions, including simple random sampling (SRS), statistical outlier removal (SOR), denoiser and upsampler network (DUP-Net) [2], and IF-Defense [1]. The results, presented in Tab. 2, show a significant decrease in the success rates for most methods. However, SymAttack maintains a success rate of over 90% against SRS, SOR, and DUP-Net defenses in all cases. Even under the strongest IF-Defense, it still achieves a minimum success rate of 82%. These results confirm SymAttack's effectiveness against these defenses.

Parameter Tuning Experiments. In this section, we investigate the impact of various parameters on the results. Specifically, we analyze the SymAttack on PointNet, which is trained on the ModelNet40 dataset, as illustrated in Fig. 1. For the distance threshold for defining patches, denoted as τ_r, SymAttack achieves optimal attack performance when set to 0.1. A threshold lower than this value results in the sampling process failing to identify sufficiently large patches, whereas a higher threshold leads to decreased performance due to an increased number of perturbed points. Regarding the threshold for the number of points in a part, denoted as τ_P, setting it to 64 yields the best results. For the step size for direction adjustment, denoted as α, this parameter controls the degree of perturbation direction. The most effective setting is at 0.5. If α is set below 0.5, the symmetry constraint becomes insufficient, causing the attack to revert to the performance of traditional methods. Conversely, excessively stringent constraints result in a significant drop in performance. Therefore, in this paper, we set τ_r to 0.1, τ_P to 64, and α to 0.5 for all experiments.

2 MORE VISUALIZATION RESULTS

Visualization of Adversarial Point Clouds. To illustrate how our approach enhances imperceptibility, we present visualizations of adversarial point clouds generated using various attack strategies on ShapeNetPart aimed at fooling PointNet and DGCNN, as depicted in Fig. 2 and Fig. 3. Point clouds modified by PGD and 3d-Adv exhibit noticeable outliers due to their less restrictive deformation techniques. In contrast, methods like GeoA^3, SI-Adv, and ITA, which leverage the geometric properties of shapes for modifications, result in significantly fewer visible outliers. Notably, SymAttack, by preserving the symmetry of the shape, produces adversarial point clouds that are virtually free of outliers, thereby underscoring the effectiveness and superiority of our method in generating imperceptible attacks.

REFERENCES

[1] Ziyi Wu, Yueqi Duan, He Wang, Qingnan Fan, and Leonidas J Guibas. 2020. If-defense: 3d adversarial point cloud defense via implicit function based restoration. *arXiv preprint arXiv:2010.05272* (2020).

[2] Hang Zhou, Kejiang Chen, Weiming Zhang, Han Fang, Wenbo Zhou, and Nenghai Yu. 2019. Dup-net: Denoiser and upsampler network for 3d adversarial point clouds defense. In *ICCV*. 1961–1970.

Table 1: Comparison on the perturbation sizes required by a variant of SymAttack that applies symmetric constraints on the whole shapes (SymAttack-W) and SymAttack to reach their highest achievable ASR in attacking PointNet.

Model	Attack	ModelNet40							ShapeNet Part						
		ASR (%)	CD (10^{-4})	HD (10^{-2})	l_2	GR	Curv (10^{-2})	EMD (10^{-2})	ASR (%)	CD (10^{-4})	HD (10^{-2})	l_2	GR	Curv (10^{-2})	EMD (10^{-2})
PointNet	SymAttack-W	97.9	123.574	38.561	11.564	0.897	5.247	98.56	98.5	471.632	79.857	16.924	1.165	5.641	131.145
	SymAttack	**100**	**0.451**	**0.915**	**0.228**	**0.086**	**0.109**	**0.425**	**100**	**0.589**	**0.998**	**0.334**	**0.110**	**0.317**	**0.340**

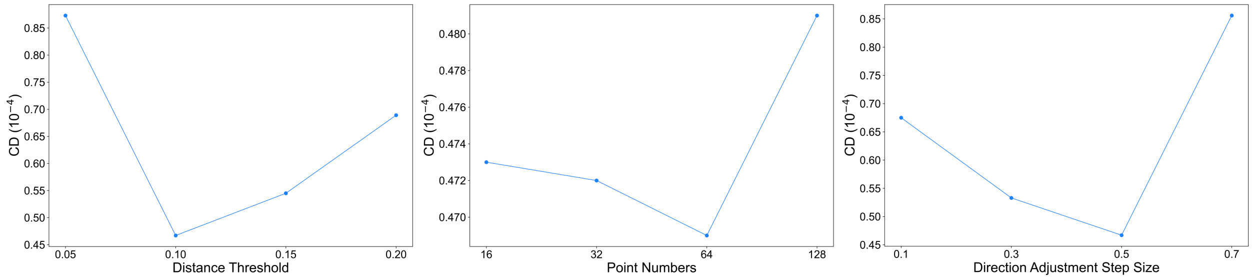

Figure 1: Imperceptibility performance of SymAttack measured by Chamfer distance (CD) under different parameter settings: distance threshold for defining patches (τ_r), threshold for the number of points in a part (τ_P), and the step size for direction adjustment (α).

Table 2: Attack success rate (%) ↑ of different attack methods with and without defense on ModelNet40 and ShapeNet Part.

Model	Defense	ModelNet40							ShapeNet Part						
		Attack Method							Attack Method						
		IFGM	3d-Adv	AdvPC	GeoA^3	ITA	SI-Adv	Ours	IFGM	3d-Adv	AdvPC	GeoA^3	ITA	SI-Adv	Ours
PointNet	-	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
	SOR	21.20	17.19	33.6	62.47	90.37	97.4	**100.0**	4.53	15.23	48.05	12.11	91.37	98.34	**98.98**
	SRS	91.69	22.53	98.87	72.65	91.85	85.78	**99.84**	76.34	19.67	99.60	72.65	86.34	82.17	**98.13**
	DUP-Net	16.29	12.30	29.00	73.70	85.41	95.80	**99.80**	3.30	12.24	29.49	8.20	82.67	93.27	**98.95**
	IF-Defense	13.80	13.70	16.77	6.04	69.32	80.30	**91.68**	5.18	8.79	17.38	9.76	48.65	76.58	**82.37**
DGCNN	-	100.0	100.0	100.0	100.0	100.0	100.0	100	100.0	100.0	100.0	100.0	100.0	100.0	100.0
	SOR	30.39	19.64	57.08	52.50	82.63	**98.65**	97.96	10.19	6.03	11.04	56.25	84.61	91.27	**95.22**
	SRS	55.71	30.51	70.63	77.71	83.48	87.62	**98.65**	44.99	11.39	48.33	86.46	85.17	88.64	**93.26**
	DUP-Net	21.15	15.64	48.12	33.34	80.65	88.67	**93.54**	3.06	9.67	8.33	39.16	81.74	89.84	**91.51**
	IF-Defense	21.13	12.35	25.00	28.75	71.26	79.64	**82.45**	4.73	3.435	12.50	3.174	53.49	78.65	**80.61**
PointConv	-	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
	SOR	59.36	24.60	91.25	18.35	87.61	92.67	**98.14**	26.20	5.46	72.92	19.14	86.69	91.55	**96.06**
	SRS	95.06	86.22	93.54	21.48	88.74	91.63	**98.27**	54.39	53.06	90.42	33.84	84.21	92.49	**97.28**
	DUP-Net	29.9	24.00	76.04	9.38	78.69	89.46	**95.73**	11.80	4.03	55.00	17.97	71.64	88.69	**94.18**
	IF-Defense	25.61	9.99	35.62	4.29	40.35	72.64	**82.89**	5.15	8.56	26.04	7.42	29.37	70.36	**83.54**

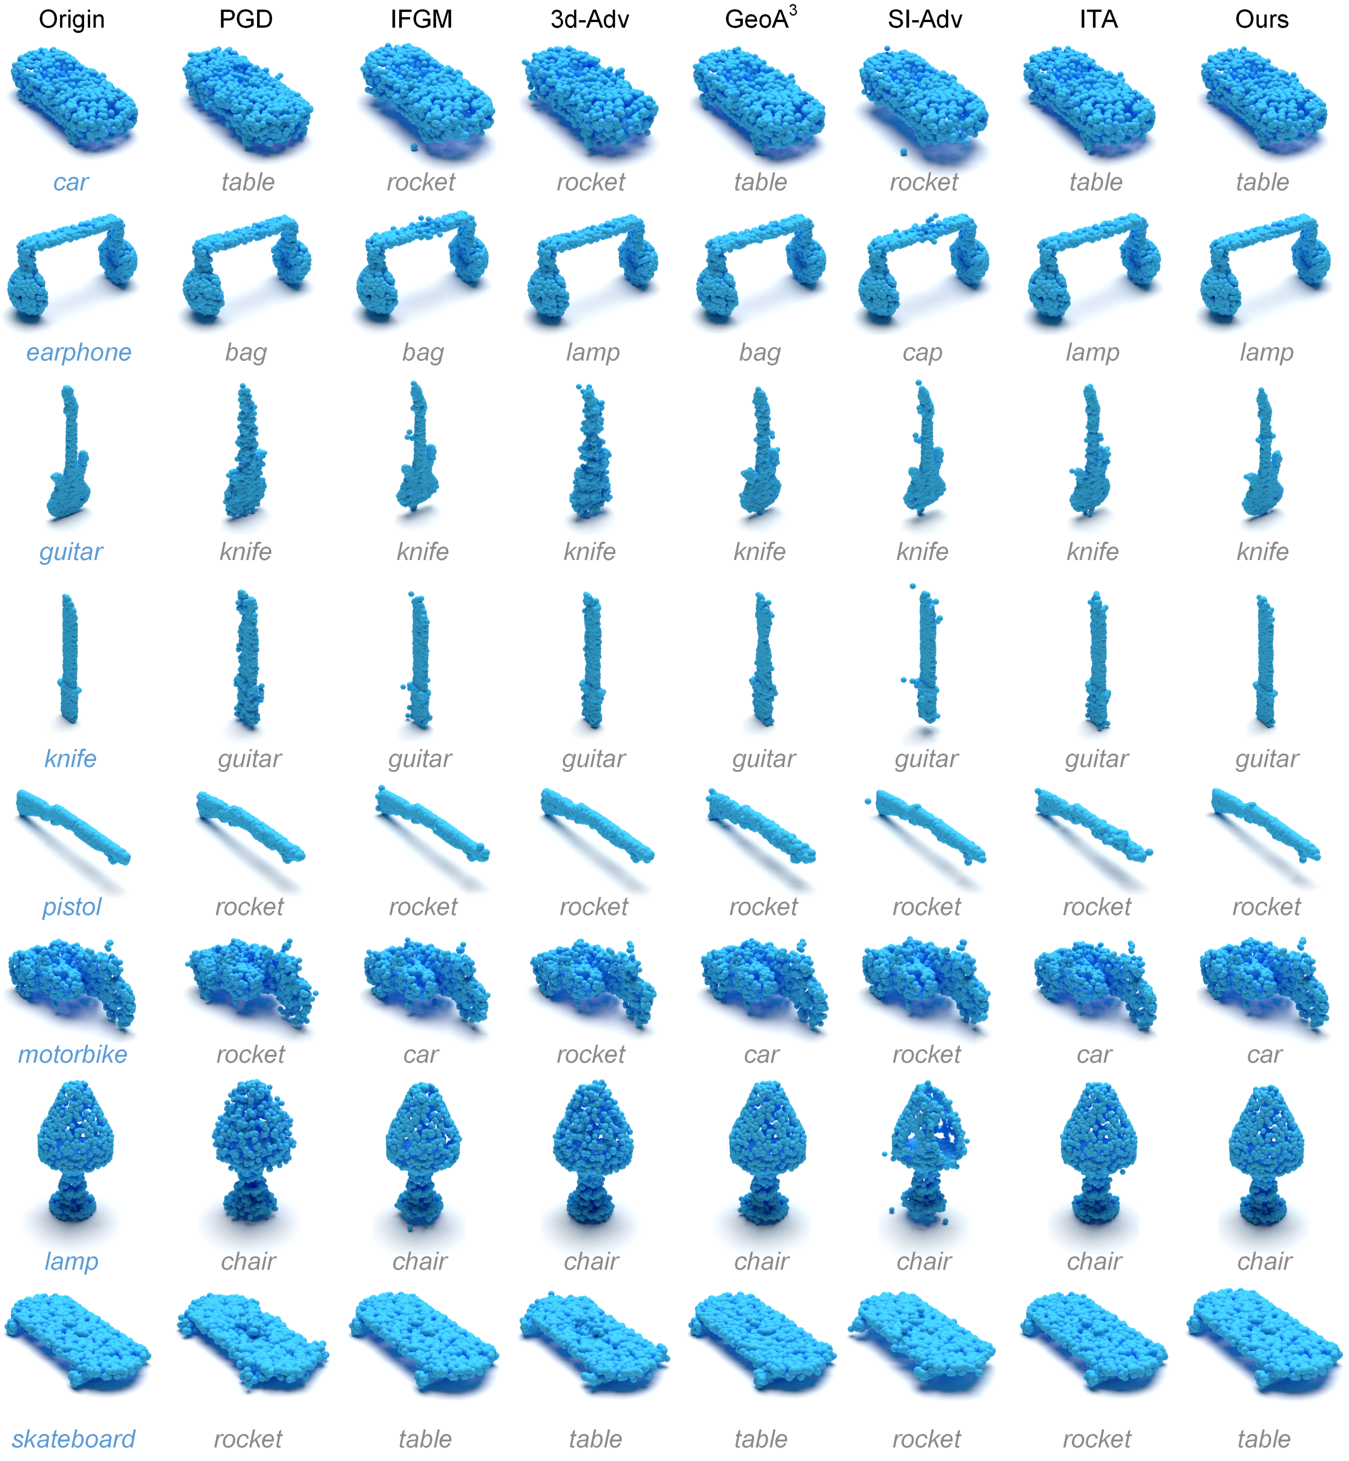

Figure 2: Visualizations of original and adversarial point clouds generated to fool PointNet on the ShapeNetPart dataset using different adversarial attack methods.

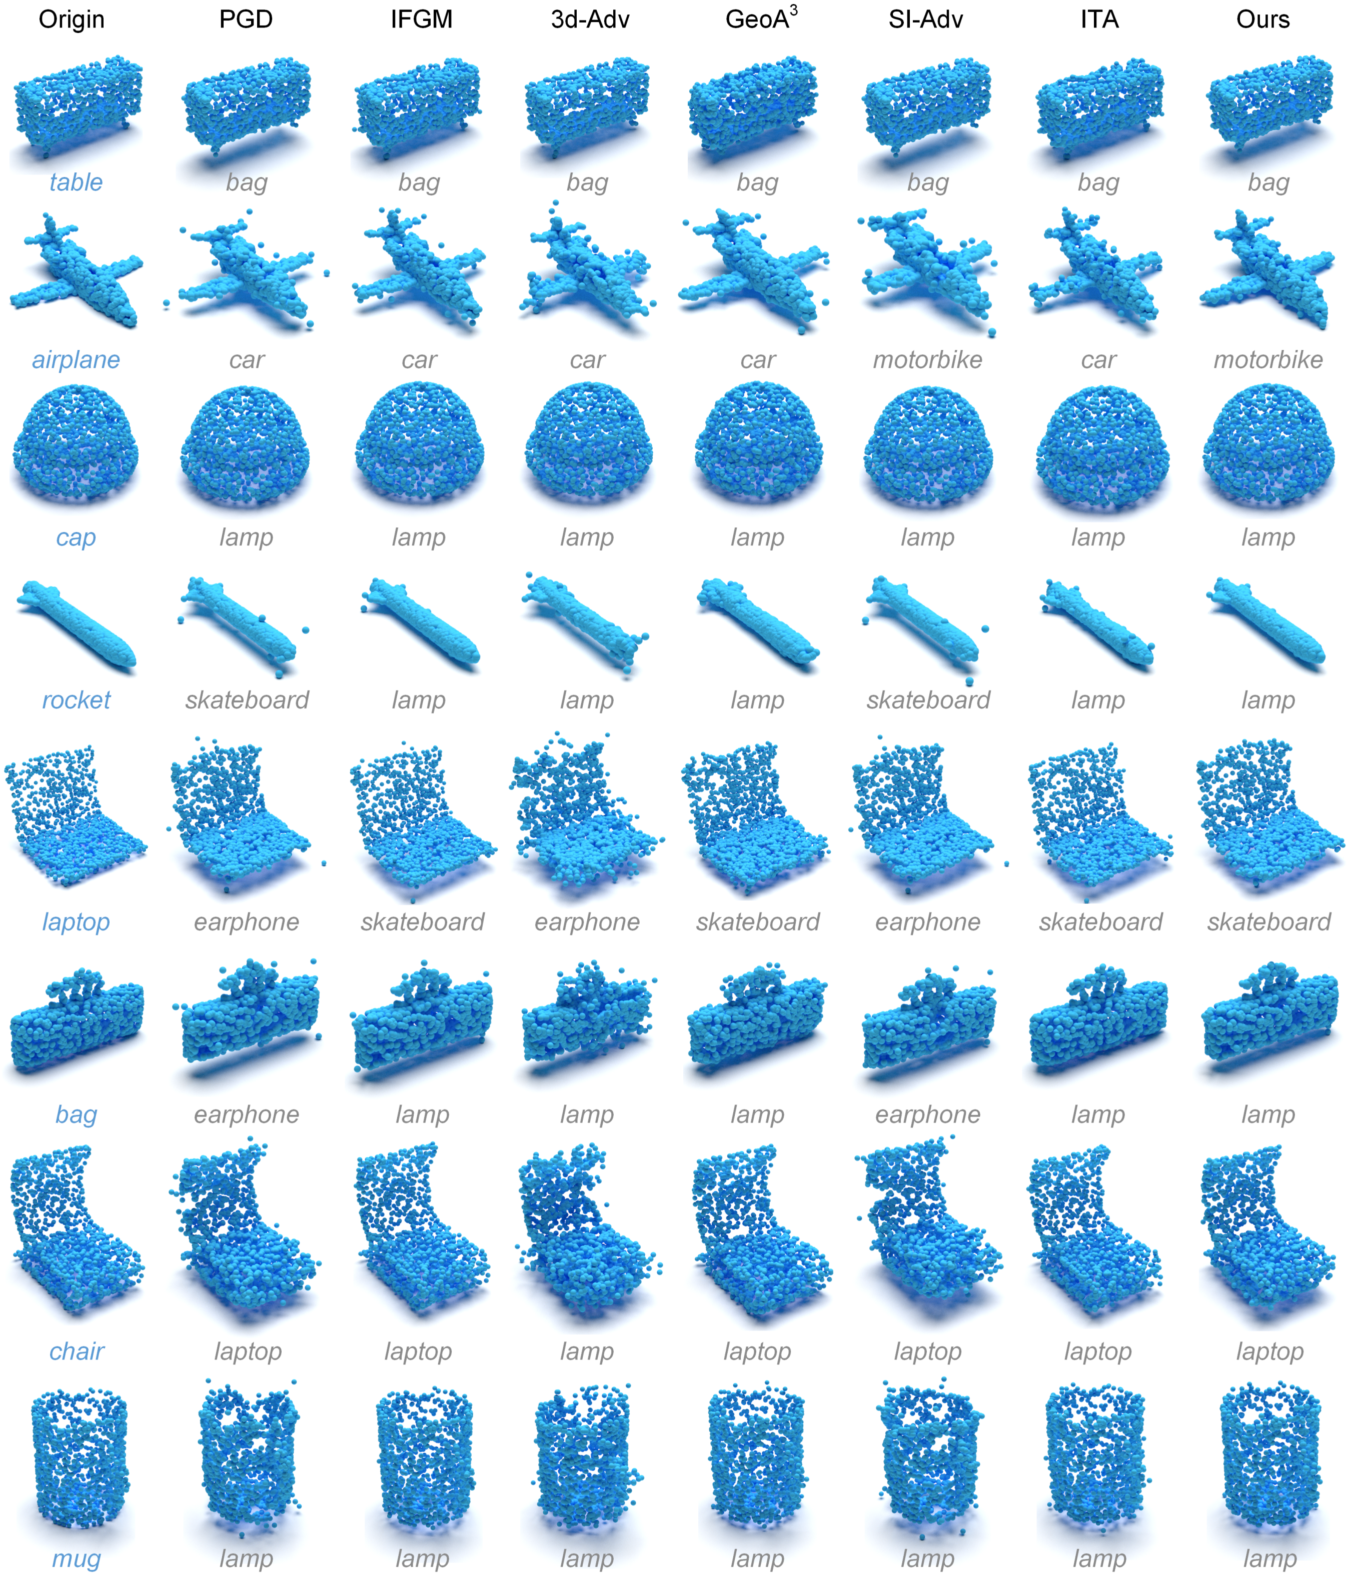

Figure 3: Visualizations of original and adversarial point clouds generated to fool DGCNN on the ShapeNetPart dataset using different adversarial attack methods.