

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 TALK UNTIL YOU BURN OUT: ESCALATING 3D-LLM OVERGENERATION VIA SEMANTIC MANIPULATION

005 **Anonymous authors**

006 Paper under double-blind review

ABSTRACT

011 The rise of 3D large language models (3D-LLMs) has unlocked new potential in
012 multimodal reasoning over unstructured 3D data, powering applications such as
013 robotics and autonomous driving. However, these models also face new security
014 risks, particularly during the inference-time computation. In this work, we present
015 **Exhaust3D**, the first targeted energy-oriented adversarial framework against 3D-
016 LLMs. Exhaust3D performs a **resource exhaustion attack** by injecting imper-
017 ceptible yet strategically structured semantic perturbations into 3D point clouds,
018 causing the model to overgenerate outputs and inflate inference latency. Speci-
019 fically, we design two key components: (1) a *semantic-aware adversarial manipu-*
020 *lation strategy* that leverages internal model representations to selectively perturb
021 semantically critical point regions while preserving geometric structure, and (2)
022 a *trajectory disruption mechanism* that maintains high-entropy token predictions
023 to prolong auto-regressive decoding and induce verbose outputs. Experiments on
024 widely-used 3D-LLM benchmarks show that Exhaust3D increases decoding steps
025 and energy consumption by up to **6.45 \times** with negligible degradation in functional
026 performance. These results expose a previously underestimated vulnerability of
027 3D-LLMs to resource exhaustion attacks, highlighting the urgent need for energy-
028 aware robustness in future multimodal foundation models.

1 INTRODUCTION

031 Recent advances in 3D large language models (3D-LLMs) such as PointLLM (Xu et al., 2024;
032 2025) and X-InstructBLIP (Panagopoulou et al., 2023) have significantly improved spatial-language
033 reasoning, enabling robust performance in a wide range of 3D tasks including object understanding,
034 scene-level decision-making, and multimodal navigation. These models typically adopt a language
035 foundation model (LLM) as their backbone and extend it to the 3D domain by aligning textual and
036 geometric modalities, thereby generalizing the LLM’s linguistic capabilities to visual and spatial
037 perception. Consequently, 3D-LLMs are increasingly deployed in high-stakes applications such as
038 autonomous driving, embodied AI, and digital reconstruction.

039 However, the powerful generative and reasoning capabilities of 3D-LLMs come at the cost of enor-
040 mous model sizes and expensive inference pipelines. In particular, their autoregressive decoding
041 requires substantial computational resources per query. Moreover, due to the limited availability
042 and high cost of constructing 3D datasets, users often rely on publicly available 3D assets (e.g.,
043 meshes, point clouds) sourced from the internet to support downstream reasoning tasks. This opens
044 up a new vulnerability: an adversary may subtly manipulate these 3D inputs to induce 3D-LLMs to
045 generate abnormally long outputs at inference time—causing excessive energy usage or cloud token
046 consumption, thereby launching a resource exhaustion attack.

047 Several prior works have explored resource exhaustion attacks in text and image modalities. In the
048 text domain, Sponge examples (Shumailov et al., 2021) increase inference cost by maximizing ac-
049 tivation norms across layers. Meanwhile, Engorgio Prompt (Dong et al., 2024) demonstrates that
050 carefully crafted prompts can suppress end-of-sequence (EOS) token generation in LLMs, forcing
051 abnormally long outputs without degrading semantic quality. In the vision-language setting,
052 NICGSlowdown (Chen et al., 2022) manipulates logit dynamics to delay EOS emission, and Ver-
053 bose Images (Gao et al., 2024) introduce imperceptible perturbations that promote diverse outputs
in VLMs, thereby amplifying computational overhead.

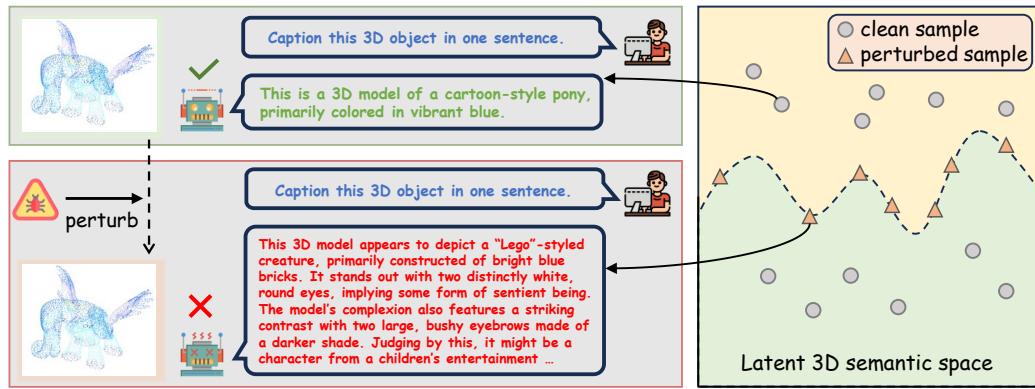


Figure 1: Demonstration of Exhaust3D’s effects. **Left (output level):** Adversarial perturbations significantly prolong auto-regressive decoding while preserving input fidelity. **Right (model level):** Perturbations shift samples into regions of semantic ambiguity, blurring distinctions from clean examples in semantic space.

Despite these efforts, existing computation-based attacks have primarily focused on LLMs and VLMs, whereas the study of 3D-LLMs remains largely unexplored. The unique characteristics of 3D domains pose distinctive challenges for effective attacks, such as maintaining geometric fidelity, ensuring robust cross-modal alignment, and addressing the semantic sparsity inherent in 3D spaces. To bridge this gap, we present **Exhaust3D**, the first adversarial framework that performs semantic-aware resource exhaustion attacks against 3D-LLMs through purely 3D-modal perturbations. Exhaust3D injects imperceptible but strategically crafted changes into point cloud inputs to manipulate the LLM’s generation process and escalate inference-time energy consumption. The key components of our method are two-fold: (1) **Semantic-aware adversarial manipulation**: We leverage the model’s internal hidden states to identify semantically important tokens and selectively perturb their corresponding points in the input space, disrupting the model’s reasoning process while preserving the overall geometric structure. (2) **Trajectory disruption mechanism**: We maintain high-entropy token predictions and suppress premature EOS emission to prolong auto-regressive decoding, triggering verbose outputs and significantly increasing inference-time energy consumption. Figure 1 illustrates the effect of Exhaust3D. Experiments on widely-used 3D-LLM benchmarks show that Exhaust3D substantially increases decoding steps and energy consumption, exposing a critical vulnerability in 3D multimodal reasoning systems. By combining semantic-aware adversarial manipulation and a trajectory disruption mechanism, the attack induces excessively long sequences in a highly stealthy and resource-draining manner. Our contributions are summarized as follows:

- To the best of our knowledge, we propose the first attack framework named Exhaust3D, which performs semantic-aware **resource exhaustion attacks** on 3D-LLMs via imperceptible perturbations crafted entirely in the 3D input space.
- We design a novel ambiguity-driven trajectory disruption mechanism that effectively manipulates token dispersion and persistence to prolong decoding and induce verbose, energy-expensive outputs from various 3D-LLMs.
- We conduct comprehensive experiments on the Objaverse (Deitke et al., 2023) and ModelNet40 (Wu et al., 2015) benchmark, showing that Exhaust3D increases output length and energy cost by up to 6.45 \times and 6.12 \times respectively, demonstrating consistent and powerful attack performance.

2 RELATED WORK

2.1 RECENT ADVANCES IN 3D-LLMs

Recent advances in 3D large language models (3D-LLMs) have substantially enhanced the ability of language models to perceive and reason over 3D data. A first line of research leverages 2D vision-language models for 3D understanding. 3D-LLM (Hong et al., 2023) renders 3D objects into

108 multi-view images and applies CLIP-like encoders (Radford et al., 2021) together with BLIP (Li
 109 et al., 2023a) to bridge 3D perception and language understanding.
 110

111 Beyond explicit multi-view rendering, another line of work focuses on directly aligning 3D point
 112 clouds with language models. Point-Bind (Guo et al., 2023) constructs a shared multimodal em-
 113 bedding space under the guidance of ImageBind (Girdhar et al., 2023), using image features to as-
 114 sist point cloud alignment and enabling generative capabilities through 2D multimodal models like
 115 ImageBind-LLM (Han et al., 2023). PointLLM (Xu et al., 2024; 2025) encodes 3D point clouds into
 116 latent point tokens and concatenates them with textual tokens for auto-regressive LLM processing,
 117 enabling end-to-end spatial-language reasoning from 3D input. LEO (?) extends this framework by
 118 incorporating additional modalities to enhance cross-modal reasoning for downstream tasks.
 119

120 A number of other models explore advanced 3D reasoning and scene-level understanding.
 121 ShapeLLM (Qi et al., 2024) combines contrastive pretraining with cross-modal alignment to enable
 122 zero-shot generalization. MiniGPT-3D (Tang et al., 2024) improves scalability through a multi-stage
 123 projection pipeline from 3D point clouds to token representations. GreenPLM (Tang et al., 2025)
 124 achieves data-efficient learning by mapping point clouds into text space, requiring only 12% of 3D
 125 training data. Video-3D LLM (Zheng et al., 2025) introduces 3D positional encoding into video
 126 representations for reasoning over dynamic scenes. X-InstructBLIP (Panagopoulou et al., 2023)
 127 aligns multiple modalities with a frozen LLM using complementary Q-Former and projection mod-
 128 ules. Finally, LSceneLLM (Zhi et al., 2025) targets large-scale scene understanding with adaptive
 129 preference identification and scene magnification.
 130

131 2.2 RESOURCE EXHAUSTION ATTACKS

132 In the field of computer security, Denial-of-Service (DoS) attacks are a classic and pervasive threat.
 133 Denial-of-Service (DoS) attacks aim to exhaust system resources or bandwidth, preventing legiti-
 134 mate users from accessing services (Elleithy et al., 2005; Aldhyani & Alkahtani, 2023; Bhatia et al.,
 135 2018). Typical strategies include overwhelming servers with massive request floods or exploiting
 136 vulnerabilities (Mirkovic & Reiher, 2004; Long & Thomas, 2001). Such attacks pose long-standing
 137 risks to real-world platforms, as they directly undermine service availability.

138 With the emergence of adaptive neural architectures, including adaptive computation time networks
 139 and large-scale models such as Large Language Models (LLMs) and Vision-Language Models
 140 (VLMs), inference-time computation has become input-dependent. Unlike conventional networks
 141 with fixed costs, these models dynamically adjust decoding length or intermediate computation ac-
 142 cording to input content. This paradigm shift has enabled a new family of resource exhaustion
 143 attacks, which drain computational resources by inducing excessive energy consumption and lat-
 144 ency through adversarial inputs (Hong et al., 2020; Liu et al., 2023). For example, sponge sam-
 145 ples (Shumailov et al., 2021) maximize the L2 norm of intermediate activations to trigger redundant
 146 computation, while NICGSlowDown (Chen et al., 2022) manipulates end-of-sequence (EOS) logits
 147 to prolong sequence generation and increase decoding overhead.

148 In the era of generative models, resource exhaustion attacks have also emerged against LLMs and
 149 VLMs. Engorgio Prompt (Dong et al., 2024) shows that carefully crafted prompts can suppress EOS
 150 token generation in LLMs, forcing abnormally long outputs without degrading semantic quality.
 151 Verbose Images (Gao et al., 2024) extends this idea to vision-language models by delaying EOS
 152 generation, increasing output uncertainty, and enhancing token diversity to induce energy-intensive
 153 decoding. LLMEffiChecker (Feng et al., 2024) further identifies critical tokens and constructs input
 154 variants to systematically reduce LLM inference efficiency across layers. In the 3D domain, Poison-
 155 splat (Lu et al., 2024) perturbs multi-view data to substantially increase the computational overhead
 156 of 3D Gaussian Splatting (Kerbl et al., 2023), highlighting how input-level perturbations translate
 157 into heavy resource consumption.

158 Despite this growing body of work, **resource exhaustion attacks targeting 3D-LLMs remain**
159 largely unexplored. These models combine multi-view perception, geometric reasoning, and cross-
 160 modal text generation, leading to complex and input-dependent inference pipelines. Such character-
 161 istics create new vulnerabilities but also raise unique challenges for attack design. Our work takes
 162 the first step toward systematically exploring and exploiting these vulnerabilities, providing insights
 163 that can guide the development of more robust and resource-efficient 3D-LLMs.

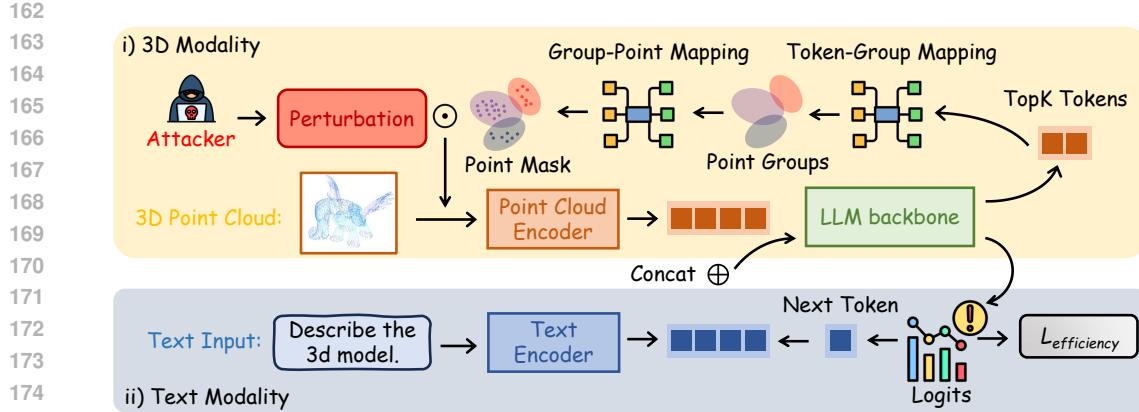


Figure 2: Overall pipeline of Exhaust3D. The input 3D point cloud P is first tokenized, and subtle perturbations are injected only at semantically critical points indicated via the token–group–point mapping, creating targeted semantic ambiguity. Point tokens are then combined with textual embeddings to drive auto-regressive decoding in the LLM. By evaluating the predicted logits, the efficiency-oriented loss $\mathcal{L}_{\text{efficiency}}$ directs the perturbations to extend output sequences and amplify computational cost.

3 METHODOLOGY

3.1 METHOD OVERVIEW

To extend generated sequences while keeping perturbations imperceptible, **Exhaust3D** leverages 3D point cloud structure through two modules. **Semantic-aware Adversarial Manipulation** identifies semantically critical points based on token importance and selectively perturbs them to preserve local geometry. **Trajectory Disruption Mechanism** prolongs generated outputs using a *dispersion loss* to increase token uncertainty and a *persistence loss* to delay EOS emission. To balance two losses, we adopt a projection-based adjustment optimization method. Together, these components generate longer outputs with minimal geometric distortion, amplifying the inference-time cost of 3D-LLMs. An overview of the attack pipeline is shown in Figure. 2.

3.2 PROBLEM FORMULATION

3D-LLM Inference Procedure. We consider a 3D large language model (3D-LLM) composed of an encoder and an auto-regressive decoder. Let $P \in \mathbb{R}^{N \times 3}$ denote the input point cloud and $T = \{t_1, \dots, t_m\}$ the textual instruction. The encoder maps P into latent point tokens H_p and the textual input into tokens H_t . These are concatenated to form the initial context for the decoder. The decoder then generates an output sequence $Y = \{y_1, \dots, y_k\}$ auto-regressively. At decoding step s , the model predicts token y_s conditioned on the encoded inputs and the prefix $Y_{<s} = \{y_1, \dots, y_{s-1}\}$. Concretely, the decoder logits at step s are:

$$\mathbf{z}_s(P, T, y_{<s}) = f_{\text{dec}}([H_p; H_t], y_{<s}), \quad (1)$$

and the conditional probability of generating token $v \in V$ at step s is:

$$p(v | P, T, y_{<s}) = \text{softmax}(\mathbf{z}_s(P, T, y_{<s}))_v. \quad (2)$$

Attack Objective. We aim to craft an imperceptible perturbation δ applied to the point cloud P to obtain a perturbed input $\tilde{P} = P + \delta$, such that the 3D-LLM generates abnormally long output sequences \tilde{Y} , increasing inference-time computation. This is formalized by minimizing an efficiency-oriented objective:

$$\min_{\|\delta\|_\infty \leq \epsilon} \mathcal{L}_{\text{efficiency}}(f_{\text{dec}}([f_{\text{enc}}(\tilde{P} + \delta); H_t])), \quad (3)$$

where ϵ bounds the perturbation magnitude. The perturbation δ should preserve the geometric integrity of the original point cloud (See details in Section 3.3) while the minimization of $\mathcal{L}_{\text{efficiency}}$

216 aims to reduce the decoding efficiency, encourages high energy consumption during inference (See
 217 details in Section 3.4).

219 **3.3 SEMANTIC-AWARE ADVERSARIAL MANIPULATION**

221 To preserve the geometric fidelity of the input, we introduce a semantic-aware adversarial manipulation
 222 strategy that perturbs only a subset of semantically critical points. This strategy consists of
 223 three steps: (1) estimating token importance, (2) mapping tokens back to their corresponding points,
 224 and (3) applying masked perturbations under perceptual constraints.

225 We first measure the importance of each point token using the ℓ_2 norm of its final-layer hidden state.
 226 The top- k tokens are then selected according to a ratio ρ :

$$228 \quad \mathcal{I} = \text{TopK}(\{\|\mathbf{h}_j^{(L)}\|_2\}_{j=1}^n, k = \lfloor \rho n \rfloor). \quad (4)$$

229 This step highlights the tokens with the strongest influence on the next-token logits and therefore
 230 the model’s generation behavior. Each selected token j corresponds to a local group of raw points
 231 \mathcal{G}_j , as determined by the encoder’s grouping stage. We expand the token set \mathcal{I} into a point-level
 232 set $\mathcal{P} = \bigcup_{j \in \mathcal{I}} \mathcal{G}_j$, which specifies all candidate points eligible for perturbation. In practice, we
 233 construct a **semantic-aware mask** $M \in \{0, 1\}^N$ to mark these points, ensuring that gradients
 234 outside \mathcal{P} are suppressed during optimization. Finally, the adversarial perturbation is applied in
 235 masked form:

$$236 \quad \tilde{P} = P + \delta \odot M, \quad (5)$$

237 where δ denotes the additive perturbation constrained by an ℓ_∞ budget to maintain imperceptibility.
 238 During optimization, only masked points are updated, focusing perturbations on semantically influ-
 239 ential regions while leaving the majority of the point cloud untouched. This semantic-aware mask
 240 design enables strong attack effectiveness with minimal perceptual distortion.

242 **3.4 TRAJECTORY DISRUPTION MECHANISM**

244 Building upon the semantic-aware perturbation strategy described above to maintain imperceptibil-
 245 ity, we further introduce our trajectory disruption mechanism. It aims to prolong the auto-regressive
 246 decoding trajectory, which maximizes the inference-time energy cost of 3D-LLMs.

248 **From Adversarial Examples to Semantic Ambiguity.** In classification tasks, adversarial exam-
 249 ples (AEs) are known to probe the decision boundary by applying small perturbations that induce
 250 misclassification (He et al., 2018; Ilyas et al., 2019; Rice et al., 2020). However, large language
 251 models, particularly 3D-LLMs, operate in an open-ended generative setting rather than a closed la-
 252 bel space. In this context, the notion of a decision boundary must be reinterpreted. Instead of forcing
 253 high-confidence misclassifications, we transfer the functionality of AEs into the semantic space of
 254 generative models: perturbations are designed to push samples toward regions of high uncertainty,
 255 where token predictions become more ambiguous.

256 **Definition of SAI.** Formally, let $\tilde{P} = P + \delta$ denote the perturbed point cloud, and let T denote the
 257 input text. Let $\tilde{\mathbf{z}}^i \in \mathbb{R}^{|V|}$ denote the model’s logits for the i -th token, and define

$$259 \quad \hat{p}_i(\cdot) := \text{softmax}(\tilde{\mathbf{z}}^i), \quad (6)$$

261 where $\hat{p}_i(v)$ gives the predicted probability for token $v \in V$. We call \tilde{P} a **spatially ambiguous**
 262 **instance (SAI)** if its predictive distribution exhibits sufficiently high entropy:

$$263 \quad \mathcal{H}(\hat{p}_i) \geq \log |V| - \tau, \quad (7)$$

265 where $\mathcal{H}(\hat{p}_i) = -\sum_{v \in V} \hat{p}_i(v) \log \hat{p}_i(v)$, $\log |V|$ is the entropy of the uniform distribution over the
 266 vocabulary, and $\tau \geq 0$ is a slack threshold. Intuitively, a smaller τ enforces a distribution closer
 267 to uniform (higher ambiguity); in the extreme $\tau = 0$ the condition requires near-uniformity. SAIs
 268 therefore correspond to perturbed samples that lie near the semantic boundary of the model’s latent
 269 space, producing ambiguous predictions that prolong auto-regressive decoding and inflate inference-
 time energy consumption.

270 **Dispersion loss.** This geometric interpretation provides intuition for why SAIs exhibit unstable
 271 decoding behavior, and it naturally motivates the need for an objective that can deliberately induce
 272 such high-entropy states. To actively construct SAIs, we define a dispersion loss that encourages the
 273 predicted distributions to approach the maximum entropy distribution across the full sequence:
 274

$$275 \quad \mathcal{L}_{\text{dispersion}} = \mathbb{E}_{i=1}^L \left[\text{KL}(\hat{p}_i \parallel \mathcal{U}) \right], \quad (8)$$

277 where \mathcal{U} is the uniform distribution over the vocabulary V , and the expectation is taken over se-
 278 quence positions $i = 1, \dots, L$. Minimizing $\mathcal{L}_{\text{dispersion}}$ drives the predicted distributions toward
 279 high entropy, thereby encouraging SAIs as defined in Eq. 7.

280 **Persistence loss.** Entropy maximization alone only diversifies token predictions but leaves the
 281 stopping behavior uncontrolled. To regulate the termination dynamics, we design a persistence loss
 282 that discourages premature EOS emission:
 283

$$284 \quad \mathcal{L}_{\text{persistence}} = \mathbb{E}_{i=1}^L \left[-\log(1 - \hat{p}_i(\text{EOS})) \right], \quad (9)$$

286 where $\hat{p}_i(\text{EOS})$ denotes the predicted EOS probability at position i . This penalizes early EOS
 287 predictions and enforces a more persistent decoding trajectory. By combining the dispersion and
 288 persistence terms, we obtain the overall efficiency-oriented objective:

$$289 \quad \mathcal{L}_{\text{efficiency}} = \mathcal{L}_{\text{dispersion}} + \mathcal{L}_{\text{persistence}}. \quad (10)$$

291 This joint formulation drives predictions toward high-entropy distributions while simultaneously
 292 suppressing EOS emission, thereby prolonging auto-regressive trajectories and inflating inference-
 293 time energy consumption in 3D-LLMs.

294 **Gradient conflict optimization.** When jointly optimizing $\mathcal{L}_{\text{dispersion}}$ and $\mathcal{L}_{\text{persistence}}$, their
 295 gradients may conflict. To stabilize the optimization, we adopt a projection-based adjustment inspired
 296 by PCGrad (Yu et al., 2020). Given gradients g_{disp} and g_{pers} , if their inner product is negative, we
 297 remove the conflicting component:
 298

$$299 \quad g_{\text{disp}} \leftarrow g_{\text{disp}} - \frac{\langle g_{\text{disp}}, g_{\text{pers}} \rangle}{\|g_{\text{pers}}\|^2 + \epsilon} g_{\text{pers}}. \quad (11)$$

301 The final update direction is then

$$302 \quad g_{\text{final}} = g_{\text{disp}} + g_{\text{pers}}. \quad (12)$$

304 This adjustment mitigates destructive interference and improves optimization stability.

306 4 EXPERIMENTS

308 4.1 EXPERIMENTAL CONFIGURATIONS

310 **Datasets and Models.** We conduct experiments on two point-cloud datasets: Objaverse (Deitke
 311 et al., 2023), a large-scale collection of diverse 3D objects with rich annotations, and Model-
 312 Net40 (Wu et al., 2015), a widely used benchmark of CAD-based 3D models. For each dataset,
 313 we uniformly sample 8,192 points per object and randomly select 500 point clouds as the eval-
 314 uation set. For 3D-LLMs, we consider four representative models that take raw point clouds as
 315 input: PointLLM (Xu et al., 2024), X-InstructBLIP (Panagopoulou et al., 2023), GreenPLM (Tang
 316 et al., 2025), and MiniGPT-3D (Tang et al., 2024), ensuring that the evaluation focuses purely on
 317 point-cloud understanding without interference from auxiliary modalities.

318 **Comparison Baselines.** To the best of our knowledge, we are the first to investigate resource ex-
 319 haustion attacks against 3D-LLMs. Since no prior work explicitly targets resource-oriented degra-
 320 dation, we compare Exhaust3D against Gaussian Noise, which adds random perturbations to point
 321 coordinates, and Random Drop, which removes 10% of points. These baselines simulate realistic
 322 point cloud corruptions commonly encountered in practice and generally affect 3D understand-
 323 ing. Comparing with them highlights that our method specifically targets inference-time computation
 rather than recognition accuracy.

Table 1: Resource overhead of four 3D-LLMs on Objaverse and ModelNet40 under different attack settings. Best results are in **bold**.

Model	Method	Objaverse			ModelNet40		
		Length	Latency	Energy	Length	Latency	Energy
PointLLM	Original	19.77	0.84	66.45	14.76	0.65	52.45
	Gaussian Noise	20.58	1.08	104.78	22.84	0.89	70.62
	Random Drop	25.27	1.32	127.95	19.26	0.78	62.32
	Exhaust3D	127.52	5.65	406.50	45.90	1.88	139.27
X-InstructBLIP	Original	16.87	1.76	98.65	10.04	1.62	88.52
	Gaussian Noise	11.86	1.22	72.68	7.91	1.67	92.34
	Random Drop	16.62	1.72	100.35	10.12	1.62	89.29
	Exhaust3D	39.03	2.83	272.43	32.85	2.40	193.54
MiniGPT-3D	Original	27.84	6.92	183.96	11.34	6.93	239.10
	Gaussian Noise	12.98	6.91	181.70	13.00	6.92	243.54
	Random Drop	26.28	6.87	190.14	11.34	6.92	237.48
	Exhaust3D	74.71	6.79	553.37	65.87	6.71	383.17
GreenPLM	Original	16.52	0.92	43.03	13.69	0.82	14.01
	Gaussian Noise	16.25	1.04	45.84	16.69	0.91	15.26
	Random Drop	17.67	1.09	47.80	13.93	0.84	14.46
	Exhaust3D	41.48	1.52	67.10	35.58	1.50	59.82

Evaluation Metrics. Our evaluation considers efficiency-oriented metrics, including average output sequence length, inference latency (s), and energy consumption (J), which directly reflect the ability of an attack to exhaust computational resources.

Implementation Details. For all models, the default prompts are applied, and the maximum output length is set to 2,048 tokens. Exhaust3D is optimized for 100 iterations with a perturbation bound of $\ell_\infty \leq 0.1$. Gaussian Noise is also applied with the same $\ell_\infty \leq 0.1$ bound, while Random Drop directly removes 10% of the points. Additional architectural settings and implementation details are provided in Appendix A.

4.2 MAIN RESULTS

Table. 1 evaluates the resource overhead induced by different perturbations on four 3D-LLMs across Objaverse and ModelNet40. Exhaust3D consistently produces the most severe increases, while Gaussian Noise and Random Drop have limited effects.

On **Objaverse**, Exhaust3D significantly inflates PointLLM’s response length (from 20 to 128 tokens, $6.5\times$), latency ($6.7\times$), and energy ($512\% \uparrow$). MiniGPT-3D also shows notable overhead, whereas X-InstructBLIP exhibits moderate growth and GreenPLM remains relatively robust (costs roughly double).

On **ModelNet40**, the same pattern holds with slightly smaller magnitudes: PointLLM and MiniGPT-3D are most affected, X-InstructBLIP sees moderate inflation, and GreenPJM is stable.

Overall, PointLLM and MiniGPT-3D are consistently the most vulnerable to Exhaust3D, X-InstructBLIP shows moderate sensitivity, and GreenPLM demonstrates strong resilience, highlighting efficiency robustness as a key consideration for 3D-LLMs.

To further illustrate these effects, we analyze the distribution of output length, latency, and energy across all perturbations (Figure 3). Interestingly, all three metrics under Exhaust3D exhibit a bimodal distribution: one peak remains close to the baseline responses of the original inputs, while another peak emerges at significantly higher values. This pattern suggests that the attack not only shifts the average cost upward but also creates highly variable responses, with some inputs remaining

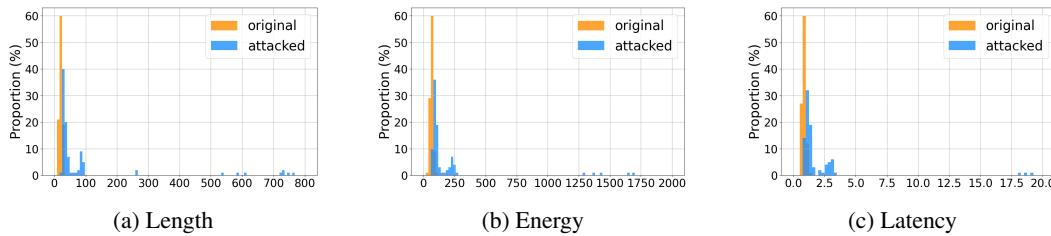


Figure 3: Distribution of PointLLM results on Objaverse: (a) Sequence Length, (b) Energy Consumption, and (c) Inference Latency.

relatively stable while others trigger extreme resource escalation. Such bimodality highlights the instability of 3D-LLMs under targeted perturbations and the difficulty of predicting computational demands in adversarial settings.

We further conduct t-SNE analyses of the learned embeddings to gain an interpretable perspective. At the **token level**, Exhaust3D induces clear structural changes in the feature space: clusters become more concentrated and compact compared to the dispersed structure in the original setting, reflecting that token semantics are significantly altered by the perturbation. At the **instance level**, however, the global layout of embeddings remains relatively close to the original distribution, indicating that the overall point cloud appearance is preserved. Taken together, these results reveal that Exhaust3D achieves a favorable trade-off: it maintains a degree of visual imperceptibility at the instance level while substantially disrupting token-level semantic alignment, thereby driving the observed escalation in sequence length, latency, and energy.

Table 2: Ablation study of loss objectives for PointLLM on the Objaverse dataset. Best results are in **bold**.

$\mathcal{L}_{\text{dispersion}}$	$\mathcal{L}_{\text{persistence}}$	PCGrad	Length	Latency	Energy
✓			101.85	4.62	320.79
	✓		31.42	1.53	116.79
✓	✓		124.25	5.21	336.17
✓	✓	✓	127.52	5.65	406.50

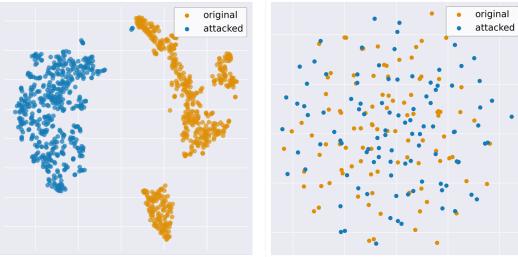


Figure 4: T-SNE visualization of learned embeddings under Exhaust3D. Left: token-level embeddings become more compact. Right: instance-level embeddings retain the overall layout.

4.3 ABLATION STUDY

Effect of Loss Objectives. We first investigate the contribution of each loss objective in Exhaust3D. Table 2 reports the results for $\mathcal{L}_{\text{dispersion}}$, $\mathcal{L}_{\text{persistence}}$, and PCGrad, individually and in combination. Both $\mathcal{L}_{\text{dispersion}}$ and $\mathcal{L}_{\text{persistence}}$ independently improve attack effectiveness, increasing sequence length and energy consumption, with $\mathcal{L}_{\text{dispersion}}$ having a more decisive impact due to its substantially higher numerical gains. Combining the two further enhances performance, while the addition of PCGrad yields the most pronounced effect, achieving the longest outputs, highest latency, and largest energy overhead. This confirms that all three components contribute positively and synergistically to Exhaust3D’s efficiency attack.

Effect of Masking Strategy. We next assess the importance of the **semantic-aware point mask** by comparing it with a random mask of identical sparsity (Table 3). While random masking can still produce moderate perturbation effects, our semantic-aware mask consistently focuses on points associated with top-ranked latent tokens, yielding slightly stronger attack effects. Visual inspection

432
433

Table 3: Ablation study on the mask method.

434
435
436
437
438

Method	Length	Latency	Energy
Original	14.76	0.65	52.45
Random Mask	42.85	1.88	139.27
Ours	45.90	2.25	148.46

439
440
441

Table 5: Ablation study on the effect of point number.

442
443
444
445
446
447
448

Points	Objaverse			ModelNet40		
	Length	Energy	Latency	Length	Energy	Latency
2048	52.10	2.15	154.04	46.38	2.44	179.23
4096	140.71	7.71	503.01	45.91	1.88	138.54
8192	127.52	5.65	406.50	45.90	1.88	139.27

449
450
451

further confirms that the mask tends to cover informative regions, enhancing the attack impact while maintaining imperceptibility (see Appendix B for detailed visualization).

452
453
454
455
456
457
458

Effect of Mask Ratio. We conduct an ablation study on the mask ratio to examine its effect on attack behavior (Table 4). The results reveal a clear trade-off: smaller mask ratios yield more imperceptible perturbations with reduced ℓ_2 norms, while moderate ratios enhance attack effectiveness with only a slight loss in stealthiness. However, overly large ratios do not bring proportional gains and instead increase perturbation magnitude, suggesting that balanced sparsity is key to achieving both efficiency and imperceptibility.

459
460
461
462
463
464
465
466
467

Effect of Point Cloud Size. We further study the impact of point count on Exhaust3D (Table 5). On Objaverse, the attack weakens significantly at 2,048 points, with much shorter responses and lower overhead compared to denser inputs. This is likely because Objaverse objects contain fine-grained structures that become underrepresented at lower resolutions, limiting the effectiveness of semantic-aware perturbations. In contrast, ModelNet40 exhibits relatively stable results across all point counts. Since its CAD-style objects are already coarse and lack detailed geometry, reducing the number of points does not substantially change the attack outcome.

471
472
473
474
475
476

Summary. Overall, these ablation studies demonstrate that (i) the combination of loss objectives with PCGrad maximizes attack efficiency, (ii) semantic-aware masking plays a key role in selecting high-impact points while preserving imperceptibility, (iii) smaller mask ratios yield more stealthy perturbations without fully sacrificing attack strength, and (iv) the effect of point count depends on dataset granularity: attacks weaken on Objaverse at low resolutions due to loss of fine-grained details, whereas ModelNet40 remains stable given its inherently coarse geometry. These findings highlight the design choices that underpin Exhaust3D’s effectiveness and adaptability across diverse scenarios.

477
478

5 CONCLUSION

479
480
481
482
483
484
485

In this work, we presented **Exhaust3D**, the first adversarial framework that performs semantic-aware resource exhaustion attacks against 3D-LLMs through purely 3D-modal perturbations. By integrating semantic-aware adversarial manipulation with a trajectory disruption mechanism, Exhaust3D induces verbose decoding and substantially escalates inference-time energy consumption while remaining imperceptible to human observers. Our experiments across multiple datasets, point cloud resolutions, and masking strategies confirm both the effectiveness and generalizability of the method, exposing a critical vulnerability in 3D multimodal reasoning systems and underscoring the need for future research on efficiency-aware defense mechanisms.

Table 4: Ablation study on the mask ratio.

Mask ratio	Length	Energy	Latency	ℓ_2
30%	32.68	1.39	108.52	0.38
50%	45.90	1.88	139.27	0.45
100%	44.06	1.79	141.61	0.70

486 REFERENCES
487

488 Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen
489 Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, et al. Phi-3 technical report: A
490 highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*, 2024.

491 THH Aldhyani and H Alkahtani. Cyber security for detecting distributed denial of service attacks
492 in agriculture 4.0: Deep learning model. *mathematics*, 11 (233), 2023.

493

494 Sajal Bhatia, Sunny Behal, and Irfan Ahmed. Distributed denial of service attacks and defense
495 mechanisms: current landscape and future directions. In *Versatile Cybersecurity*, pp. 55–97.
496 Springer, 2018.

497

498 Simin Chen, Zihe Song, Mirazul Haque, Cong Liu, and Wei Yang. Nicgslowdown: Evaluating the
499 efficiency robustness of neural image caption generation models. In *Proceedings of the IEEE/CVF*
500 *Conference on Computer Vision and Pattern Recognition*, pp. 15365–15374, 2022.

501

502 Wei-Lin Chiang, Zhuohan Li, Ziqing Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng,
503 Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, et al. Vicuna: An open-source chatbot
504 impressing gpt-4 with 90%* chatgpt quality. See <https://vicuna.lmsys.org> (accessed 14 April
505 2023), 2(3):6, 2023.

506

507 Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig
508 Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of anno-
509 tated 3d objects. In *Proceedings of the IEEE/CVF conference on computer vision and pattern*
510 *recognition*, pp. 13142–13153, 2023.

511

512 Jianshuo Dong, Ziyuan Zhang, Qingjie Zhang, Tianwei Zhang, Hao Wang, Hewu Li, Qi Li, Chao
513 Zhang, Ke Xu, and Han Qiu. An engorgio prompt makes large language model babble on. *arXiv*
514 *preprint arXiv:2412.19394*, 2024.

515

516 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas
517 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An
518 image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint*
519 *arXiv:2010.11929*, 2020.

520

521 Khaled M Elleithy, Drazen Blagovic, Wang K Cheng, and Paul Sideleau. Denial of service attack
522 techniques: analysis, implementation and comparison. 2005.

523

524 Xiaoning Feng, Xiaohong Han, Simin Chen, and Wei Yang. Llmeffichecker: Understanding and
525 testing efficiency degradation of large language models. *ACM Transactions on Software Engi-
526 neering and Methodology*, 33(7):1–38, 2024.

527

528 Kuofeng Gao, Yang Bai, Jindong Gu, Shu-Tao Xia, Philip Torr, Zhifeng Li, and Wei Liu. Induc-
529 ing high energy-latency of large vision-language models with verbose images. *arXiv preprint*
530 *arXiv:2401.11170*, 2024.

531

532 Rohit Girdhar, Alaaeldin El-Nouby, Zhuang Liu, Mannat Singh, Kalyan Vasudev Alwala, Armand
533 Joulin, and Ishan Misra. Imagebind: One embedding space to bind them all. In *Proceedings of*
534 *the IEEE/CVF conference on computer vision and pattern recognition*, pp. 15180–15190, 2023.

535

536 Ziyu Guo, Renrui Zhang, Xiangyang Zhu, Yiwen Tang, Xianzheng Ma, Jiaming Han, Kexin Chen,
537 Peng Gao, Xianzhi Li, Hongsheng Li, et al. Point-bind & point-llm: Aligning point cloud
538 with multi-modality for 3d understanding, generation, and instruction following. *arXiv preprint*
539 *arXiv:2309.00615*, 2023.

540

541 Jiaming Han, Renrui Zhang, Wenqi Shao, Peng Gao, Peng Xu, Han Xiao, Kaipeng Zhang, Chris Liu,
542 Song Wen, Ziyu Guo, et al. Imagebind-llm: Multi-modality instruction tuning. *arXiv preprint*
543 *arXiv:2309.03905*, 2023.

544

545 Warren He, Bo Li, and Dawn Song. Decision boundary analysis of adversarial examples. In *Inter-
546 national Conference on Learning Representations*, 2018.

540 Sanghyun Hong, Yiğitcan Kaya, Ionuț-Vlad Modoranu, and Tudor Dumitraș. A panda? no, it's
 541 a sloth: Slowdown attacks on adaptive multi-exit neural network inference. *arXiv preprint*
 542 *arXiv:2010.02432*, 2020.

543 Yining Hong, Haoyu Zhen, Peihao Chen, Shuhong Zheng, Yilun Du, Zhenfang Chen, and Chuang
 544 Gan. 3d-llm: Injecting the 3d world into large language models. *Advances in Neural Information*
 545 *Processing Systems*, 36:20482–20494, 2023.

546 Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Logan Engstrom, Brandon Tran, and Aleksander
 547 Madry. Adversarial examples are not bugs, they are features. *Advances in neural information*
 548 *processing systems*, 32, 2019.

549 Mojan Javaheripi, Sébastien Bubeck, Marah Abdin, Jyoti Aneja, Sébastien Bubeck, Caio
 550 César Teodoro Mendes, Weizhu Chen, Allie Del Giorno, Ronen Eldan, Sivakanth Gopi, et al.
 551 Phi-2: The surprising power of small language models. *Microsoft Research Blog*, 1(3):3, 2023.

552 Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splat-
 553 tting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023.

554 Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image
 555 pre-training with frozen image encoders and large language models. In *International conference*
 556 *on machine learning*, pp. 19730–19742. PMLR, 2023a.

557 Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image
 558 pre-training with frozen image encoders and large language models. In *International conference*
 559 *on machine learning*, pp. 19730–19742. PMLR, 2023b.

560 Han Liu, Yuhao Wu, Zhiyuan Yu, Yevgeniy Vorobeychik, and Ning Zhang. Slowlidar: Increasing
 561 the latency of lidar-based detection using adversarial examples. In *Proceedings of the IEEE/CVF*
 562 *Conference on Computer Vision and Pattern Recognition*, pp. 5146–5155, 2023.

563 Neil Long and Rob Thomas. Trends in denial of service attack technology. *CERT Coordination*
 564 *Center*, 648(651):569, 2001.

565 Jiahao Lu, Yifan Zhang, QiuHong Shen, Xinchao Wang, and Shuicheng Yan. Poison-splat: Compu-
 566 tation cost attack on 3d gaussian splatting. *arXiv preprint arXiv:2410.08190*, 2024.

567 Jelena Mirkovic and Peter Reiher. A taxonomy of ddos attack and ddos defense mechanisms. *ACM*
 568 *SIGCOMM Computer Communication Review*, 34(2):39–53, 2004.

569 Artemis Panagopoulou, Le Xue, Ning Yu, Junnan Li, Dongxu Li, Shafiq Joty, Ran Xu, Silvio
 570 Savarese, Caiming Xiong, and Juan Carlos Niebles. X-instructblip: A framework for aligning
 571 x-modal instruction-aware representations to llms and emergent cross-modal reasoning. *arXiv*
 572 *preprint arXiv:2311.18799*, 2023.

573 Zekun Qi, Runpei Dong, Shaochen Zhang, Haoran Geng, Chunrui Han, Zheng Ge, Li Yi, and
 574 Kaisheng Ma. Shapellm: Universal 3d object understanding for embodied interaction. In *Euro-
 575 pean Conference on Computer Vision*, pp. 214–238. Springer, 2024.

576 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
 577 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual
 578 models from natural language supervision. In *International conference on machine learning*, pp.
 579 8748–8763. PMLR, 2021.

580 Leslie Rice, Eric Wong, and Zico Kolter. Overfitting in adversarially robust deep learning. In
 581 *International conference on machine learning*, pp. 8093–8104. PMLR, 2020.

582 Ilia Shumailov, Yiren Zhao, Daniel Bates, Nicolas Papernot, Robert Mullins, and Ross Anderson.
 583 Sponge examples: Energy-latency attacks on neural networks. In *2021 IEEE European sym-
 584 posium on security and privacy (EuroS&P)*, pp. 212–231. IEEE, 2021.

585 Quan Sun, Yuxin Fang, Ledell Wu, Xinlong Wang, and Yue Cao. Eva-clip: Improved training
 586 techniques for clip at scale. *arXiv preprint arXiv:2303.15389*, 2023.

594 Yuan Tang, Xu Han, Xianzhi Li, Qiao Yu, Yixue Hao, Long Hu, and Min Chen. Minigpt-3d:
 595 Efficiently aligning 3d point clouds with large language models using 2d priors. In *Proceedings*
 596 *of the 32nd ACM International Conference on Multimedia*, pp. 6617–6626, 2024.

597

598 Yuan Tang, Xu Han, Xianzhi Li, Qiao Yu, Jinfeng Xu, Yixue Hao, Long Hu, and Min Chen. More
 599 text, less point: Towards 3d data-efficient point-language understanding. In *Proceedings of the*
 600 *AAAI Conference on Artificial Intelligence*, volume 39, pp. 7284–7292, 2025.

601 Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong
 602 Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE*
 603 *conference on computer vision and pattern recognition*, pp. 1912–1920, 2015.

604 Runsen Xu, Xiaolong Wang, Tai Wang, Yilun Chen, Jiangmiao Pang, and Dahua Lin. Pointllm:
 605 Empowering large language models to understand point clouds. In *European Conference on*
 606 *Computer Vision*, pp. 131–147. Springer, 2024.

607

608 Runsen Xu, Shuai Yang, Xiaolong Wang, Tai Wang, Yilun Chen, Jiangmiao Pang, and Dahua Lin.
 609 Pointllm-v2: Empowering large language models to better understand point clouds. *IEEE Trans-*
 610 *actions on Pattern Analysis and Machine Intelligence*, 2025.

611 Le Xue, Ning Yu, Shu Zhang, Artemis Panagopoulou, Junnan Li, Roberto Martín-Martín, Jiajun
 612 Wu, Caiming Xiong, Ran Xu, Juan Carlos Niebles, et al. Ulip-2: Towards scalable multimodal
 613 pre-training for 3d understanding. In *Proceedings of the IEEE/CVF Conference on Computer*
 614 *Vision and Pattern Recognition*, pp. 27091–27101, 2024.

615

616 Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn.
 617 Gradient surgery for multi-task learning. *Advances in neural information processing systems*, 33:
 618 5824–5836, 2020.

619 Xumin Yu, Lulu Tang, Yongming Rao, Tiejun Huang, Jie Zhou, and Jiwen Lu. Point-bert:
 620 Pre-training 3d point cloud transformers with masked point modeling. In *Proceedings of the*
 621 *IEEE/CVF conference on computer vision and pattern recognition*, pp. 19313–19322, 2022.

622

623 Duo Zheng, Shijia Huang, and Liwei Wang. Video-3d llm: Learning position-aware video represen-
 624 tation for 3d scene understanding. In *Proceedings of the Computer Vision and Pattern Recognition*
 625 *Conference*, pp. 8995–9006, 2025.

626 Hongyan Zhi, Peihao Chen, Junyan Li, Shuailei Ma, Xinyu Sun, Tianhang Xiang, Yinjie Lei,
 627 Mingkui Tan, and Chuang Gan. Lscenellm: Enhancing large 3d scene understanding using
 628 adaptive visual preferences. In *Proceedings of the Computer Vision and Pattern Recognition*
 629 *Conference*, pp. 3761–3771, 2025.

630 Junsheng Zhou, Jinsheng Wang, Baorui Ma, Yu-Shen Liu, Tiejun Huang, and Xinlong Wang. Uni3d:
 631 Exploring unified 3d representation at scale. *arXiv preprint arXiv:2310.06773*, 2023.

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646

647

648 APPENDIX
649650 A MORE EXPERIMENTAL DETAILS
651652 In this work, we focus on **3D point clouds** as the input modality (instead of projecting them into
653 images or multi-view representations). Accordingly, we adopt several open-source and competi-
654 tive 3D-to-text models as baselines. Below we describe their architectural details and the common
655 implementation protocols.
656657 **PointLLM Settings** The adopted PointLLM (Xu et al., 2024) is composed of a Point-BERT (Yu
658 et al., 2022) encoder pre-trained with ULIP-2 (Xue et al., 2024), a multi-layer projector with GeLU
659 activations to map point features into the embedding space, and a Vicuna-7B (Chiang et al., 2023)
660 LLM as the language backbone. With two additional special tokens, the vocabulary size is 32003.
661 Prompt template used in PointLLM is ‘‘Describe the 3D model.’’.
662663 **X-InstructBLIP Settings** The adopted X-InstructBLIP (Panagopoulou et al., 2023) integrates a
664 Point-BERT (Yu et al., 2022) encoder pre-trained with ULIP-2 (Xue et al., 2024) for 3D repre-
665 sentation learning, a Q-Former module that learns query tokens, and a Vicuna-7B (Chiang et al.,
666 2023) as the language backbone. All Q-Formers are initialized with BLIP-2 (Li et al., 2023b) stage-
667 1 weights to ensure stable training and effective multimodal alignment. Prompt template used in
668 X-InstructBLIP is ‘‘Describe the 3D model.’’.
669670 **MiniGPT-3D Settings** The adopted MiniGPT-3D (Tang et al., 2024) is composed of a Q-Former
671 initialized from BLIP-2 (Li et al., 2023b), a Mixture of Query Experts (MQE) to enhance semantic
672 representation, a modality projector for aligning point queries with the text embedding space, and
673 the Phi-2-2.7B backbone (Jawaheripi et al., 2023) as the language model. Prompt template used in
674 MiniGPT-3D is ‘‘Describe the 3D model in short’’.
675676 **GreenPLM Settings** The adopted GreenPLM (Tang et al., 2025) consists of a ViT (Dosovitskiy
677 et al., 2020) point encoder and an EVA-CLIP-E (Sun et al., 2023) text encoder, both trained by
678 Uni3D (Zhou et al., 2023), a two-layer MLP projector with GeLU activation to align encoder outputs
679 with the language embedding space, and the Phi-3 (Abdin et al., 2024) model as the language back-
680 bone. Prompt template used in GreenPLM is ‘‘Describe the 3D model in detail’’.
681682 **Implementation Notes**
683684

- 685 • All experiments are conducted with FP16 precision.
- 686 • Each point cloud is uniformly sampled into 8192 points (except for ablation studies on
687 point number).
- 688 • ModelNet40 objects are assigned a fixed black color to compensate for missing texture.
- 689 • All models are evaluated in their default inference mode (no fine-tuning for our attack).
- 690 • We ensure that all models receive the same point sampling (e.g., fixed 8,192 points unless
691 testing point-number ablations).
- 692 • For metrics like latency and energy, we run each model under the same hardware setting
(single H20 GPU) and average over multiple runs.

693 B VISUALIZATION OF OUR SAMANTIC-AWARE MASK
694695 To qualitatively illustrate the behavior of our semantic-aware mask, we visualize its effect on rep-
696 resentative point clouds from both Objaverse (Deitke et al., 2023) (Figure 5) and ModelNet40 (Wu
697 et al., 2015) (Figure 6). In each figure, the top row shows the point cloud with the semantic-aware
698 mask applied, highlighting the points selected for perturbation, while the bottom row shows the
699 original, unmasked point cloud for reference.
700701 As can be observed, the semantic-aware mask consistently covers the core regions of the objects
702 that are strongly associated with their semantic meaning. For example, in the Objaverse samples

(Figure 5), the mask predominantly highlights the body and characteristic parts of objects, leaving peripheral or less informative points unaltered. Similarly, in ModelNet40 (Figure 6), masked points focus on the geometrically and semantically salient regions, such as the roof and legs of chairs or the wings and fuselage of airplanes. This visualization confirms that our mask selectively targets points that are most influential for the model’s understanding, ensuring that perturbations concentrate on semantically critical areas while minimally affecting irrelevant points.

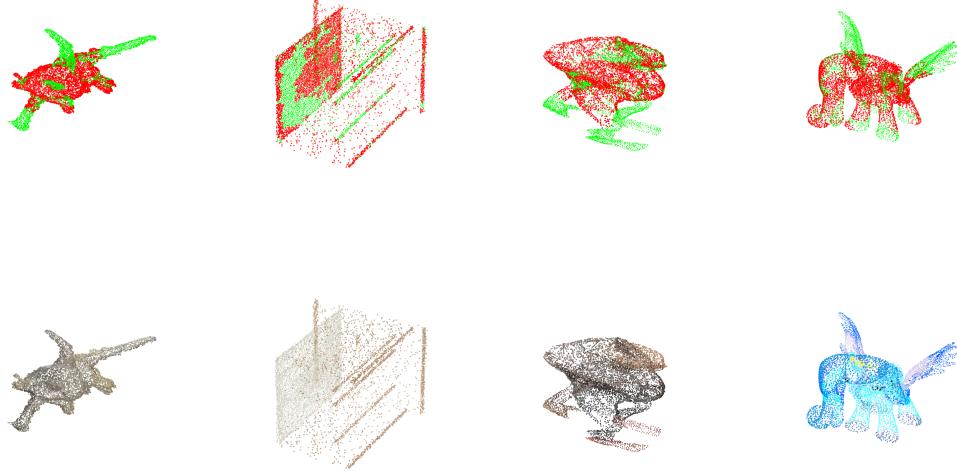


Figure 5: The visualization of our semantic-aware mask on Objaverse. The top row shows the masked object, while the bottom row shows the original object without mask. In the masked objects, points covered by the semantic-aware mask are highlighted in red, while the remaining unmasked points are shown in green.



Figure 6: The visualization of our semantic-aware mask on ModelNet40. The top row shows the masked object, while the bottom row shows the original object without mask. In the masked objects, points covered by the semantic-aware mask are highlighted in red, while the remaining unmasked points are shown in green.

756 **C LLM USAGE**
757758 We used an OpenAI LLM (GPT-5) as a writing and formatting assistant. In particular, it helped refine
759 grammar and phrasing, improve clarity, and suggest edits to figure/table captions and layout (e.g.,
760 column alignment, caption length, placement). The LLM did not contribute to research ideation,
761 experimental design, implementation, data analysis, or technical content beyond surface-level edits.
762 All outputs were reviewed and edited by the authors, who take full responsibility for the final text
763 and visuals.

764

765

766

767

768

769

770

771

772

773

774

775

776

777

778

779

780

781

782

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809