

SceneSplat++: A Large Dataset and Comprehensive Benchmark for Language Gaussian Splatting

*Mengjiao Ma^{1,2}, *Qi Ma^{1,3}, *Yue Li⁴, Jiahuan Cheng⁵, Runyi Yang¹, †Bin Ren^{1,6,7}, Nikola Popovic¹, Mingqiang Wei², Nicu Sebe⁷, Luc Van Gool¹, Theo Gevers⁴, Martin R. Oswald⁴, Danda Pani Paudel¹

¹INSAT, Sofia University “St. Kliment Ohridski” ²Nanjing University of Aeronautics and Astronautics ³ETH Zürich

⁴University of Amsterdam ⁵Johns Hopkins University ⁶University of Pisa ⁷University of Trento




Figure 1: **(a) SceneSplat-49K Dataset.** We introduce SceneSplat-49K, a large-scale 3DGS dataset comprising approximately 49 K diverse indoor and outdoor scenes. **(b) SceneSplat-Bench.** We introduce a comprehensive benchmark for evaluating *Language Gaussian Splatting* (LGS) methods at scale. **Benchmark Scale Comparison** shows SceneSplat-Bench significantly exceeds existing evaluation protocols with $3.7\times$ more semantic classes and $50.5\times$ more scenes. **Performance Comparison on SceneSplat-Bench** using 4 datasets across 9 methods demonstrates that the generalizable approach (SceneSplat) consistently outperforms the per-scene methods in most of the benchmark variants. Our codes, benchmark, and datasets are released at <https://scenesplatpp.gaussianworld.ai/>.

Abstract

3D Gaussian Splatting (3DGS) serves as a highly performant and efficient encoding of scene geometry, appearance, and semantics. Moreover, grounding language in 3D scenes has proven to be an effective strategy for 3D scene understanding. Current *Language Gaussian Splatting* line of work falls into three main groups: (i) per-scene optimization-based, (ii) per-scene optimization-free, and (iii) generalizable approach. However, most of them are evaluated only on rendered 2D views of a handful of scenes and viewpoints close to the training views, limiting their ability and insight into holistic 3D understanding. To address this gap, we propose the first large-scale benchmark that systematically assesses these three groups of methods

* indicates equal contribution. † indicates the corresponding author: Bin Ren <bin.ren@insait.ai>.

directly in 3D space, evaluating on 1060 scenes across three indoor datasets and one outdoor dataset. SceneSplat-Bench results demonstrate a clear advantage of the generalizable paradigm, particularly in relaxing the scene-specific limitation, enabling fast feed-forward inference on novel scenes, and achieving superior segmentation performance. We further introduce SceneSplat-49K, a carefully curated 3DGS dataset comprising around 49K diverse indoor and outdoor scenes obtained from multiple sources, with which we demonstrate the generalizable approach could harness strong data priors. Our codes, benchmark, and datasets are available at <https://scenesplatpp.gaussianworld.ai/>.

1 Introduction

The field of 3D computer vision is highly focused on capturing the geometry [68, 21, 18, 42, 60, 45, 62, 54] and the visual appearance [14, 27, 13, 39, 22] of a scene, as well as understanding its content [6, 70, 46, 88, 59]. Recently, 3D Gaussian Splatting (3DGS) [22] has emerged as the most desirable 3D representation due to its unique ability to jointly encode geometry, appearance, and scene understanding properties [77, 61, 19] in a compact form that can be effectively optimized from 2D posed images [78]. Furthermore, vision-language reasoning represents the most promising direction for 3D scene understanding [58, 11, 75, 37, 19, 73, 16], as it bridges the visual and geometric attributes of a scene with the language we use to define, describe, and reason about concepts. Therefore, this work focuses exclusively on vision-language scene understanding with 3DGS.

The most relevant methods for language Gaussian Splatting (LGS) can be categorized into three groups. The first two groups begin by extracting 2D features from all training images using vision-language foundation models (e.g., CLIP [50]). The first group then performs gradient-based per-scene optimization [47, 90, 86, 72, 48], assigning feature vectors to each 3D Gaussian primitive and optimizing them such that their renderings align with the corresponding 2D feature maps. The second group also operates on a per-scene basis but adopts an optimization-free approach [20, 38, 9]. Instead of iterative refinement, these methods directly lift 2D features to 3D primitives via weighted feature aggregation schemes. Lastly, we consider the generalizable approach. To date, only a single method has been trained on 2K indoor scenes [29], learning to directly predict a vision-language feature vector for each 3D Gaussian primitive. Although 2D vision-language features are required during training, inference relies solely on 3D Gaussians, eliminating the need to prepare additional 2D feature maps. Furthermore, all features are produced in a single forward pass.

While these methods represent important advances in integrating vision-language reasoning with 3DGS, their evaluation protocols still suffer from several critical limitations. First, most methods are evaluated on only a small number of selected scenes [47, 90, 86, 72, 48, 20, 9]. As a result, such evaluations carry a high risk of producing scene-specific rather than generalizable conclusions. Moreover, the absence of a standardized benchmark further limits systematic and comparable evaluation. Second, most methods are evaluated close to the training views [47, 38, 72]. Consequently, there is no guarantee that the results generalize well to novel viewpoints. Finally, most evaluations are conducted in 2D rather than directly in 3D space [47, 90, 86, 72, 48, 20, 38, 9]. However, *3D scene understanding fundamentally concerns performance in 3D space, which cannot be fully inferred from 2D projections*. Therefore, we introduce SceneSplat-Bench, an evaluation benchmark for 3DGS vision-language scene understanding that includes 1060 scenes spanning 325 unique semantic classes. It evaluates performance in 3D for each segment of the scene, and we use it to assess representative methods from all three groups, as presented in Tab. 1.

Among the evaluated approaches, the generalizable 3D scene understanding paradigm demonstrates the strongest potential. It eliminates the need for substantial per-scene computation at inference and enables a single forward pass per 3D scene. In addition, this approach enables the application of 3D foundation models for scene understanding without the need for task- or scene-specific retraining. This concept aligns with established practices in 2D computer vision, where pretrained foundation models have been widely adopted [17, 12, 43, 53, 50, 31]. Foundation models tailored to 3DGS have already emerged across key areas, including reconstruction [5, 7, 82, 89], scene understanding [29], and generation [74, 57]. Lastly, benchmark results demonstrate the state-of-the-art performance of the generalizable approach, highlighting its capacity to leverage data priors and extract meaningful vision-language features—effectively mitigating the noise inherent in weak supervision.

Methods	Generalizable	Optimization-free	Runtime / Scene	Benchmark Dataset	Scene Number	Class Number
<i>Per-Scene Optimization Methods</i>						
LangSplat [47]	✗	✗	621 min	LERF [24], 3DOVS [32]	9	88
FMGS [90]	✗	✗	76 min	LERF [24], 3DOVS [32]	11	123
Feature3DGS [86]	✗	✗	235 min	Replica [63]	4	7
OpenGaussian [72]	✗	✗	297 min	LERF [24], ScanNet [10](10scenes)	14	76
GOI [48]	✗	✗	252min	Mip-NeRF360 [2], Replica [63]	8	30
<i>Per-Scene Optimization-Free Methods</i>						
Gradient-Weighted 3DGS [20]	✗	✓	4 min	LERF [24], 3DOVS [32]	9	88
LUDVIG [38]	✗	✓	5.6 min	LERF [24], SPIn-NeRF [41], NVOS [55]	21	80
OccamLGS [9]	✗	✓	3.2 min	LERF [24], 3DOVS [32]	9	88
<i>Generalizable Method</i>						
SceneSplat [29]	✓	✓	0.24 min	ScanNet [10], ScanNet++ [79], Matterport3D [4]	732	321
SceneSplat-Bench	—	—	—	ScanNet [10], ScanNet++ [79], Matterport3D [4], HoliCity [87].	1060	325

Table 1: **Language Gaussian Splatting (LGS) Benchmark Overview.** **Left:** Grouped methods and their properties. **Right:** Benchmark characteristics. Our proposed SceneSplat-Bench benchmark evaluates the LGS methods at scale, across three indoor datasets and one outdoor dataset.

To facilitate the development of generalizable 3DGS scene understanding, we introduce SceneSplat-49K, a carefully curated 3DGS dataset comprising diverse indoor and outdoor scenes collected from multiple sources. Existing large-scale 3DGS datasets primarily focus on single objects [36, 33, 67], whereas scene-level datasets [29] remain limited in scale. To the best of our knowledge, SceneSplat-49K represents the most extensive open-source dataset of complex, high-quality 3DGS reconstructions at the full-scene level. The preparation of the dataset required approximately 891 GPU days and considerable human involvement. Furthermore, we demonstrate that training a generalizable 3DGS scene understanding method on a larger subset from SceneSplat-49K leads to improved performance, achieving state-of-the-art results on the benchmark.

Our contributions can be summarized as follows:

- We introduce SceneSplat-Bench, the first *benchmark* for systematically evaluating vision-language scene understanding methods in the 3DGS domain.
- We release SceneSplat-49K, a high-quality, carefully curated 3DGS *dataset* comprising 49K indoor and outdoor scenes.
- We scale up the training and evaluation of a generalizable VLM for 3DGS-based scene understanding across both indoor and outdoor environments, offering new insights.

2 Related Works

Gaussian Splatting Datasets. Most existing 3DGS datasets remain object-centric, synthetic, or limited in scale and scope. Moreover, recent works [5, 57, 89, 8] showcase 3DGS generalization, yet do not fully release their datasets, thus limiting reproducibility and scalability. Among publicly available datasets, ShapeSplat [36] transforms 65K object-level CAD models into 3DGS. 3DAffordSplat [67] focuses on functionality, providing 23K object instances labeled with affordances across 21 categories. uCO3D [33] expands the object-centric landscape with over 1,000 annotated categories and 360° coverage. Furthermore, scene-level 3DGS remains underexplored, with SceneSplat [29] being the only available dataset providing 6.8K indoor scenes to support open-vocabulary 3D understanding, though its scope is limited to the indoor domain. In summary, existing datasets fall short of enabling generalizable, multi-modal learning across diverse scenes. To address this, we introduce SceneSplat-49K, a large-scale dataset of 49K 3DGS scenes spanning indoor and outdoor environments, enriched with vision-language embeddings to support scene-level understanding.

Open-Vocabulary Scene Understanding. Recent advancements in open-vocabulary models have significantly expanded the capabilities of vision-language understanding. Foundation models [81, 43, 26, 52, 49, 80, 65] enabled visual feature extraction at scale, effectively supporting tasks like detection, segmentation and open-vocabulary understanding. Recent works have also introduced 2D foundation models into 3D by utilizing NeRF [40] or 3DGS [23]. Existing 3DGS approaches fall into three categories: per-scene optimization-based [86, 84, 47, 15, 48, 90, 46, 72], per-scene optimization-free [20, 38, 9] and generalizable [29]. Per-scene optimization-based LGS emerged first, with LERF [24, 51] and LangSplat [47] integrating 2D foundation model features into NeRF- and GS-based models, a design adopted by many later methods. Per-scene optimization-free LGS, such as OccamLGS [9], LUDVIG [38] and Gradient Weighted 3DGS [20] further optimized this process

Property	SceneSplat-7K [29]	DL3DV-10K [30]	HoliCity [87]	Aria-ASE [1]	Crowdsourced	Total
<i>Data Statistics</i>						
Raw Scenes	9114	10K	6 253	25K	603	48 856 ~49K
GS Scenes	7916	6 376	5 940	25K	603	45 835 ~46K
Data Type	Indoor	Indoor/Outdoor	Outdoor	Indoor	Indoor/Outdoor	Indoor/Outdoor
RGB Frames	4.72M	3.09M	48K	14.79M	N/A	22.65M
Avg. GS / Scene	1.42M	1.48M	800K	1.27M	1.02M	1.26M
3DGS GPU Time [h]	3.59K	2.1K	0.79K	6.25K	N/A	12.73K
Scenes w/ VLM Embedding	6 121	-	5 940	-	-	12 061
Embedding GPU Time [h]	8 161	-	495	-	-	8 656
Storage [TB]	2.76	2.91	0.26	2.2	0.23	8.36
Total 3DGS	11.27B	9.42B	4.75B	3.18B	615M	29.24B
<i>Evaluation Metrics (Avg.)</i>						
PSNR [dB] $\uparrow$	29.64	25.70	29.85	27.32	N/A	27.83
Depth Loss [m] $\downarrow$	0.035	-	0.010	0.082	N/A	0.061
SSIM $\uparrow$	0.897	0.845	0.921	0.906	N/A	0.898
LPIPS $\downarrow$	0.212	0.203	0.122	0.230	N/A	0.209

Table 2: **SceneSplat-49K Statistics and Quality Metrics.** Our proposed dataset includes outdoor and indoor scenes drawn from established datasets and publicly available sources, including our own collected data. In total, it comprises approximately 49 K raw scenes, 46 K curated 3D Gaussian Splatting (3DGS) scenes, and 29.24 B Gaussian splats. Generating all 3DGS required ~ 530 GPU days on an NVIDIA L4 GPU. Appearance and geometry reconstruction quality are evaluated using four metrics: PSNR, SSIM, LPIPS and depth ℓ_1 error. Additionally, 12 K of the 3DGS scenes are annotated with per-primitive vision-language embeddings to support the training of generalizable scene understanding methods. Computing these embeddings further required ~ 361 GPU days.

by employing feature lifting, directly projecting all 2D features into 3DGS, reducing computation time from hours to seconds. SceneSplat [29] is the first generalizable LGS method, trained and evaluated directly on 3D data, achieving significant gains in open-vocabulary scene understanding with substantially faster inference on novel scenes. However, most of these methods have been evaluated in constrained 2D settings, with a limited number of scenes and semantic classes. We therefore introduce a comprehensive benchmark and use it to systematically evaluate recent language Gaussian Splatting methods.

3D Generalizability. Large-scale training and generalization in 2D have become increasingly mature [50, 65, 34, 35, 52, 53, 85]. In contrast, achieving semantic and language generalization in 3D remains challenging and relatively underexplored [8, 44, 83, 71, 70, 28, 69]. One major obstacle is the limited quantity and fragmented nature of 3D datasets, as no single dataset provides the needed combination of (i) large-scale coverage, (ii) multi-scale diversity (from object-centric to indoor and outdoor scenes), and (iii) rich multimodal information (geometry, appearance, and semantics). As a result, even state-of-the-art 3D models struggle to generalize broadly. Mosaic3D-5.6M compiles over 30K RGB-D scenes with 5.6 million mask-text pairs by leveraging 2D vision-language models, yielding an significant resource for open-vocabulary 3D segmentation [28]. The SceneSplat-7K dataset supports open-vocabulary training but with limited scale [29]. These efforts represent early steps toward 3D foundation models, aiming to enable robust, optimization-free, and generalizable 3D reasoning across vision-language tasks. Therefore, we introduce SceneSplat-49K, a large and comprehensive dataset designed to unlock scaling laws in generalizable 3DGS tasks.

3 Dataset

We present SceneSplat-49K, a large-scale 3D Gaussian Splatting dataset comprising approximately 49 K raw scenes and 46 K curated 3DGS scenes, aggregated from multiple established sources, including SceneSplat-7K [29], DL3DV-10K [30], HoliCity [87], Aria Synthetic Environments [1], and our own crowdsourced data. This diverse dataset consists of both indoor and outdoor environments, spanning from rooms, apartments to streets. To support 3DGS scene understanding model training, 12 K scenes are further enriched with per-primitive vision-language embeddings extracted with state-of-the-art vision-language models [65]. The comprehensive statistics of the introduced dataset are presented in Tab. 2.

Fig. 2 visualizes the distribution of evaluation metrics of our SceneSplat-49K dataset. Across ~ 49 K 3DGS scenes, the mean photometric quality is 27.8 dB PSNR and 0.90 SSIM, with a perceptual LPIPS of 0.20—values that are of high quality renders. Importantly, the geometry reconstruction is




Figure 2: **Appearance, Geometry, and Scale Statistics of the SceneSplat-49K Dataset.** Distributions of photometric (PSNR, SSIM, LPIPS) and geometric (depth ℓ_1) reconstruction errors show consistently high-quality renders across scenes, while the wide spread in total Gaussian number and indoor/outdoor scene floor area demonstrates the dataset’s diversity. The curves are convolved from the bucket values and vertical dotted lines mark the mean of each metric.

equally reliable, with a mean depth ℓ_1 error of 0.061 m. The distribution of the per-scene Gaussian number spans two orders of magnitude, indicating the dataset’s complexity. The scene footprint areas also show clear diversity: indoor environments cluster all across 25 m² to 250 m², whereas outdoor scenes extend over a square kilometer, with a long-tailed distribution. This diversity in spatial extent, together with high-quality appearance and geometry, underpins the suitability of the dataset for training a generalizable 3DGS scene understanding model.

3.1 Data Collection and Processing

To ensure high-quality 3DGS scenes, we implement several quality control measures throughout the optimization pipeline. Starting with the training views, we select scenes with at least 400 frames to ensure sufficient multi-view coverage. When depth information is available, we initialize the Gaussians at fused point cloud locations and apply depth supervision to maximize geometry quality. We use gsplat [78] for 3DGS optimization. We filter blurry frames using the variance of the Laplacian as a sharpness metric. To efficiently compress the 3DGS scene, we employ the Markov Chain Monte Carlo strategy [25] and add opacity and scale regularization. Once optimized, the dataset supports filtering 3DGS scenes based on PSNR and depth quality before using them as inputs for model training. We introduce the processing of each data source in the following.

SceneSplat-7K contains 7,916 curated 3DGS indoor scenes derived from ScanNet [10], ScanNet++ [79], Replica [63], Hypersim [56], 3RScan [66], ARKitScenes [3], and Matterport3D [4]. We further enrich over 6,000 of these scenes with vision–language embeddings to support open-vocabulary scene queries.

Aria Synthetic Environments (ASE) is a large-scale synthetic dataset of procedurally-generated multi-room indoor scenes, each populated with 3D object models drawn from a digital asset library. We select the first 25 K scene subset of it. For our 3DGS optimization, we undistort the fisheye image and depth frames and apply the devignetnetting mask for brightness correction. The sensor depths are fused and transformed to yield point clouds for the 3DGS initialization.

DL3DV-10K dataset covers 65 everyday environments, featuring both indoor and outdoor locations. It includes 10,510 high-quality videos captured by mobile devices and drones, along with sparse COLMAP reconstructions. For 3DGS training, we use an evenly sampled set of every 10 views for novel view evaluation, while the remaining views are used for training.

HoliCity is a city-scale 3D dataset comprising 6,300 real-world panoramas, each accurately aligned to a downtown London CAD model covering over 20 km², providing depth frames and semantic

labels. We used eight perspective views from each panorama generated by evenly sampling the yaw angles at 45° intervals to optimize the 3DGS.

Crowdsourced Data. To diversify our dataset, we collected high-quality 3DGS scenes from various sources, including Polycam and Sketchfab. We have also distributed our questionnaire to the community to collect data.

3.2 Vision-Language Embedding Collection

Vision-language embeddings are extracted from the frames used to optimize the 3DGS. Unlike existing pretraining methods that align 3D primitives with text embeddings from vision-language models, we align each Gaussian directly in the image-embedding space, thereby preserving richer latent semantics. This approach avoids the information loss inherent in textual descriptions.

We follow the fusion strategy in [29] and employ a dynamic weighting mechanism that adaptively combines three distinct features: global context from the entire frame, local features with background, and masked features without background. This mechanism automatically balances contributions based on each segment’s contextual relationship—emphasizing background-inclusive features for integrated objects (e.g., keyboard with monitor) and object-specific features for isolated objects. This adaptive approach provides more nuanced semantic understanding than fixed weighting strategies.

4 Benchmark

We formulate the assigning of Gaussian language feature as follows. A scene is represented by a set of N Gaussian primitives $\mathcal{G} = \{G_j\}_{j=1}^N$, where each $G_j \in \mathbb{R}^{59}$ encodes its position, covariance, color, and opacity. For every primitive we compute a language embedding $E_j^G \in \mathbb{R}^d$, obtained by per-scene processing or by inference with a learned encoder. The user provides a set of free-form text queries $C = \{c_i\}_{i=1}^n$. A frozen text encoder P_T (e.g., CLIP) maps each query to an embedding $E_i^T = P_T(c_i)$. Each Gaussian is labeled with the query whose embedding is most similar to its own:

$$\text{label}(G_j) = \arg \max_i \cos(E_j^G, E_i^T). \quad (1)$$

This label field assigns every primitive to the most relevant query.

4.1 Benchmark settings

Evaluation Protocols. We evaluate 9 methods from 3 categories on our SceneSplat-Bench benchmark. We report two key metrics in 3D: foreground mean Intersection over Union (f-mIoU) and foreground mean accuracy (f-mAcc), which address object size imbalances and reduce viewpoint dependency compared with traditional 2D-only evaluation. We unified the evaluation pipeline by first converting the language fields into a unified format $\mathcal{E}_G = \{E_j^G\}_{j=1}^N$, then computing the cosine similarity between each language feature and the queried text embeddings. Each Gaussian primitive is assigned the semantic label corresponding to the most similar text embedding. For every ground-truth 3D point used for evaluation, we identify its K-nearest Gaussian neighbors (K=25) and assign the majority-voted label as the semantic prediction for that point. We exclude floor, wall, and ceiling classes as background for indoor scenes and the sky class for outdoor scenes. To make our evaluation more complete, we additionally evaluate semantics on 2D renders, following the protocol of LERF [24]. The 2D evaluation is conducted on subsets of the benchmark datasets with two key 2D metrics: mean Intersection over Union (mIoU) and mean accuracy (mAcc) on all classes (background and foreground classes). All runtime measurements were obtained on NVIDIA RTX A6000 GPUs, except for FMGS, whose training requires 50–80 GB of GPU memory; its runtime was therefore measured on NVIDIA H200 GPUs.

Benchmark Dataset. We evaluate the methods’ performance on indoor scenes using ScanNet (with 20 and 200 classes), ScanNet++ (100 classes), and Matterport3D (21 classes). Queries cover a wide range of object granularity—from small items like a mouse to larger objects such as a bed—reflecting the method’s capability to varying scales understanding. HoliCity (4 classes) is employed to evaluate the performance on outdoor scenes, emphasizing large structures under complex surrounding environment. Performance results are reported on randomly sampled 10-scenes subsets

Method	ScanNet20 (10 scenes)		ScanNet20 (312 scenes)		ScanNet200 (10 scenes)		ScanNet200 (312 scenes)	
	f-mIoU	f-mAcc	f-mIoU	f-mAcc	f-mIoU	f-mAcc	f-mIoU	f-mAcc
<i>Per-Scene Optimization Methods</i>								
LangSplat [47]	0.0153	0.0741	-	-	0.0057	0.0106	-	-
OpenGaussian [72]	0.1186	0.2340	-	-	0.0545	0.1086	-	-
Feature-3DGS [86]	0.1711	0.2512	-	-	0.0541	0.0876	-	-
FMGS [90]	0.0916	0.2028	0.1169	0.2684	0.0641	0.1374	0.0579	0.1443
GOI [48]	0.2295	0.4638	0.2115	0.4286	0.1193	0.2297	0.1001	0.2240
<i>Per-Scene Optimization-Free Methods</i>								
Gradient-Weighted 3DGS [20]	0.4184	0.5174	0.3573	0.4638	0.1638	0.2502	0.0601	0.1100
LUDVIG [38]	0.1406	0.2885	0.1337	0.2969	0.0887	0.1833	0.0602	0.1701
OccamLGS [9]	0.3559	0.4587	0.2308	0.4148	0.1936	0.2458	0.1204	0.2503
SceneSplat [29] (Pseudo Label)	0.3521	0.5210	0.3500	0.5550	0.2520	0.3038	0.2280	0.3590
<i>Generalizable Method</i>								
SceneSplat [29]	0.2968	0.5711	0.3540	0.5780	0.1387	0.4198	0.1648	0.3573

Table 3: **Zero-shot 3D Semantic Segmentation Experiments on ScanNet20 [10] (20 classes) and ScanNet200 [10] (200 classes).** All methods are evaluated on a 10-scene mini-validation set, with the 312-scene full validation set run only for selected methods due to runtime limitations (Tab. 1).

for each of the four benchmarks across all methods. For methods demonstrating strong efficiency, we conduct a full evaluation on the complete validation set.

Baseline Methods. We categorize the methods based on their optimization paradigm and generalizability to novel scenes. *(i) Per-Scene Optimization-based Methods* assign learnable attributes to Gaussian splats and optimize them via backpropagation. However, rendering additional features with alpha compositing is computationally expensive and requires per-scene optimization. *(ii) Per-Scene Optimization-Free methods* use a training-free approach, lifting 2D features into 3D space as a weighted sum of 2D representations, requiring only a single forward rendering pass. *(iii) Generalizable Method* SceneSplat trains a feed-forward 3DGS encoder, learning to predict representations aligned with image embeddings. We further extend its training based on our dataset. SceneSplat (Pseudo Label) denotes the VLM embeddings used for its vision-language pretraining. We refer to the supplement for details on running each baseline.



Figure 3: **Qualitative Results of Zero-Shot 3D Semantic Segmentation.** The semantic classes "bicycle" and "kitchen table" are highlighted, which are not labeled in Ground Truth.

4.2 Key Insights

The primary experimental results from our SceneSplat-Bench benchmark across all selected methods are presented in Tabs. 3 to 5 for indoor scenes and in Tab. 6 for outdoor scenes. The corresponding runtime statistics for individual methods can be found in Tab. 1. SceneSplat – representing the generalizable paradigm – stands out as the clear winner in terms of both performance and efficiency (see Fig. 3 for qualitative examples). Interestingly, in 75% of the experiments, it even outperforms its own pseudo labels used for pretraining (SceneSplat (Pseudo Label)). This result highlights its

Method	ScanNet++ (10 scenes)		ScanNet++ (50 scenes)		Matterport3D (10 scenes)		Matterport3D (370 scenes)	
	f-mIoU	f-mAcc	f-mIoU	f-mAcc	f-mIoU	f-mAcc	f-mIoU	f-mAcc
<i>Per-Scene Optimization Methods</i>								
LangSplat [47]	0.0131	0.0332	-	-	0.0219	0.1257	-	-
OpenGaussian [72]	0.0695	0.1186	-	-	0.1387	0.2698	-	-
Feature-3DGS [86]	0.0798	0.1765	-	-	0.3043	0.4732	-	-
FMGS [90]	0.0984	0.1892	0.1004	0.2300	0.2105	0.3664	-	-
GOI [48]	0.1593	0.2790	0.1631	0.3296	0.1631	0.3296	0.1652	0.3130
<i>Per-Scene Optimization-Free Methods</i>								
Gradient-Weighted 3DGS [20]	0.0925	0.1391	0.0900	0.1328	0.3083	0.4586	0.2762	0.4187
LUDVIG [38]	0.0995	0.2099	0.1099	0.2547	0.2348	0.3888	0.1960	0.3738
OccamLGS [9]	0.1580	0.2861	0.1502	0.3312	0.1803	0.3263	0.1725	0.3151
SceneSplat [29] (Pseudo Label)	0.2340	0.3709	0.2243	0.4667	0.2392	0.4368	0.2748	0.4384
<i>Generalizable Method</i>								
SceneSplat [29]	0.2263	0.4916	0.2836	0.4992	0.2732	0.5379	0.3384	0.5745

Table 4: **Zero-Shot 3D Semantic Segmentation Experiments On ScanNet++ [79] and Matterport3D [4].** All methods are evaluated on a 10-scene mini-validation set, with the full validation set evaluated only for selected methods due to runtime limitations (Tab. 1).

Method	ScanNet20 (10 scenes)		ScanNet200 (10 scenes)		ScanNet++ (10 scenes)	
	mIoU	mAcc	mIoU	mAcc	mIoU	mAcc
<i>Per-Scene Optimization Methods</i>						
LangSplat [47]	0.0193	0.0700	0.0133	0.0272	0.0111	0.0272
OpenGaussian [72]	0.1958	0.3294	0.0740	0.1273	0.1290	0.2133
Feature-3DGS [86]	0.5054	0.6474	0.1867	0.2798	0.1605	0.2368
FMGS [90]	0.1413	0.2629	0.0848	0.1825	0.1134	0.2244
GOI [48]	0.2866	0.4536	0.1134	0.1806	0.1635	0.2767
<i>Per-Scene Optimization-Free Methods</i>						
Gradient-Weighted 3DGS [20]	0.4533	0.5785	0.1639	0.2367	0.1210	0.1760
LUDVIG [38]	0.1673	0.3058	0.0958	0.1828	0.1239	0.2588
OccamLGS [9]	0.2937	0.4645	0.1300	0.2146	0.1670	0.2866
<i>Generalizable Method</i>						
SceneSplat [29]	0.4458	0.6727	0.2599	0.4049	0.3165	0.5094

Table 5: **Zero-Shot 2D Semantic Segmentation Experiments on ScanNet20 [10] (20 classes), ScanNet200 [10] (200 classes) and ScanNet++ [79].** All methods are evaluated on a 10-scene validation set.

ability to generalize by leveraging large-scale data priors and learning to predict meaningful vision-language features, effectively mitigating the noise inherent in weak supervision. In addition, Tab. 5 shows that the performance on the 2D metric is aligned with the zero-shot performance measured with the 3D metric, with a generalizable method leading the performance. As illustrated in Fig. 4, visualizations of text query results using enhanced NeRFView tools [64, 76, 29] further demonstrate that the generalizable method outperforms the others. Moreover, unlike other methods that rely on extensive feature extraction, the generalizable approach requires only 3DGS as input at inference time and predicts per-primitive vision-language features in a single forward pass.

Method	HoliCity (10 scenes)		Holicity (328 scenes)	
	f-mIoU	f-mAcc	f-mIoU	f-mAcc
<i>Per-Scene Optimization Methods</i>				
LangSplat [47]	0.0918	0.1950	-	-
OpenGaussian [72]	0.1853	0.2567	-	-
Feature-3DGS [86]	0.2479	0.5646	-	-
FMGS [90]	0.1774	0.2513	-	-
GOI [48]	0.1126	0.3773	0.1587	0.3772
<i>Per-Scene Optimization-Free Methods</i>				
Gradient-Weighted 3DGS [20]	0.2484	0.4808	0.2762	0.4187
LUDVIG [38]	0.1593	0.2296	0.1843	0.2775
OccamLGS [9]	0.2255	0.4754	0.2234	0.5056
SceneSplat [29] (Pseudo Label)	0.2464	0.4780	0.2522	0.5159
<i>Generalizable Method</i>				
SceneSplat [29]	0.3081	0.5647	0.2880	0.6045

Table 6: **Zero-Shot 3D Outdoor Segmentation On HoliCity [87] Dataset.** All methods are evaluated on mini-validation sets, with the full validation set evaluated only for selected methods due to runtime limitations (Tab. 1).




Figure 4: **Text-Based Scene Query.** Given the prompt "These are fruits" to different LGS methods, the queried parts are highlighted in red.

Method	ScanNet (10 scenes)		ScanNet++ (10 scenes)	
	Accuracy (BBox)	Accuracy (Seg)	Accuracy (BBox)	Accuracy (Seg)
<i>Per-Scene Optimization Methods</i>				
LangSplat [47]	0.0663	0.0415	0.0655	0.0337
OpenGaussian [72]	0.4565	0.3483	0.5138	0.3394
Feature-3DGS [86]	0.5861	0.4759	0.5046	0.2783
FMGS [90]	0.6082	0.4639	0.5494	0.3303
GOI [48]	0.5269	0.4414	0.5107	0.3568
<i>Per-Scene Optimization-Free Methods</i>				
Gradient-Weighted 3DGS [20]	0.5486	0.4669	0.4322	0.2661
LUDVIG [38]	0.6750	0.5481	0.4801	0.2875
OccamLGS [9]	0.2074	0.1599	0.1896	0.1182
<i>Generalizable Method</i>				
SceneSplat [29]	0.7701	0.6928	0.7339	0.5464

Table 7: **3D Object Localization Experiments on ScanNet [10] and ScanNet++ [79].** Accuracy is reported with bounding box-based and segmentation-based evaluation.

Training Source	#Scene	ScanNet++		ScanNet200		Matterport3D		HoliCity	
		f-mIoU	f-mAcc	f-mIoU	f-mAcc	f-mIoU	f-mAcc	mIoU	mAcc
Indoor Training-Data Scaling									
ScanNet++ (v1)	280	0.1683	0.3230	0.1081	0.1650	0.1010	0.2108	0.1888	0.3003
ScanNet++ (v2)	906	0.2383	0.4324	0.1180	0.1920	0.1491	0.2756	0.2339	0.3524
ScanNet++ (v2) + ScanNet	2107	0.2779	0.5055	0.1427	0.3434	0.2289	0.3803	0.2633	0.3504
ScanNet++ (v2) + ScanNet + M3D	3503	0.2836	0.4992	0.1648	0.3573	0.3384	0.5745	0.2026	0.3600
Matterport3D only	1396	—	—	—	—	0.3244	0.5354	0.2240	0.3889
Outdoor-Domain Training									
HoliCity only	3000	—	—	—	—	—	—	0.2880	0.6045

Table 8: **Impact of Training-Data Scaling on Indoor Benchmarks and Cross-Domain Generalization to HoliCity [87].** Results show that more training data consistently improves indoor performance. Furthermore, models trained on indoor data only surprisingly transfer to outdoor scenes.

Among the per-scene methods, optimization-free approaches clearly outperform optimization-based ones, both in terms of segmentation accuracy, as shown in Tabs. 3 to 6, and in runtime efficiency, as reported in Tab. 1. This is likely due to the fact that the objective function in optimization-based methods is typically designed to perform as well as possible on the training views, which may not be ideal to novel viewpoints. Taking this into account, along with the strong performance of the generalizable approach, we conclude that per-scene optimization is not necessary for effective vision-language reasoning in 3DGS. However, within both the optimization-based and optimization-free groups, no single method consistently dominates across all datasets, indicating that performance remains sensitive to dataset-specific characteristics.

A key requirement for reliable benchmarking is having a sufficient number of scenes and appropriately challenging evaluation settings. As shown in Tabs. 3, 4 and 6, the performance of various methods varies noticeably when scaling from a small subset of 10 scenes to the full validation set, highlighting the importance of the benchmark scale in assessing performance. Furthermore, Tab. 3 demonstrates

a clear drop in accuracy when the task complexity increases from 20 to 200 semantic classes, emphasizing the need for demanding benchmarks to reveal the limitations of competing approaches.

To further diversify the open-vocabulary evaluation, we include a 3D object localization task that assesses how well each method retrieves objects in 3D space using its constructed 3DGS field. Following the protocol of LangSplat [47], we directly select the point with the highest relevance score for each query. Our evaluation is conducted on 2,895 objects from ScanNet and ScanNet++ validation scenes, which is over 12× more objects compared to LERF [24], offering a broader and more realistic assessment. This evaluation is presented in Tab. 7 and shows similar conclusion as previously discussed evaluations. The generalizable paradigm shows dominant performance, followed by the per-scene optimization-free methods, which outperform per-scene optimization method group.

Tab. 8 presents our scaling study, which offers two key takeaways. First, training-data scaling consistently improves performance on indoor benchmarks: expanding the training set from 280 to 3503 scenes lifts ScanNet++ f-mIoU from 0.168 $\rightarrow$ 0.284 and ScanNet200 from 0.108 $\rightarrow$ 0.165. A similar effect appears when the model jointly trained on three indoor datasets outperforms that trained on Matterport3D alone. Second—and more surprisingly—an indoor-only model transfers to outdoor scenes, but domain-specific data is still required to close the gap. Without ever seeing outdoor data, the largest indoor model reaches 0.263 mIoU on HoliCity, an improvement of 7.4% over the 280-scene baseline. This suggests that large-scale indoor training induces representations that partially generalize across the indoor–outdoor boundary, see Fig. 5 for zero-shot examples. Nevertheless, a model trained only on HoliCity still leads (0.288 mIoU), indicating that cross-domain capability ultimately needs outdoor supervision. These findings underscore the value of our large-scale indoor–outdoor 3DGS dataset as a powerful training resource.




Figure 5: **Zero-Shot Predictions of Indoor-Trained SceneSplat on Outdoor Scenes.** The results highlight the cross-domain capability. Color palette denotes buildings, roads, terrains, and trees.

5 Conclusion and Limitations

This work introduces the first comprehensive benchmark for 3DGS-based scene understanding methods, along with a large-scale 3DGS dataset of diverse indoor and outdoor scenes. Our evaluation demonstrates that generalizable approach consistently outperforms per-scene methods, establishing a new paradigm for scalable scene understanding through pretrained models. Despite these contributions, limitations remain. Our dataset scale, while substantial, could be further expanded. The outdoor benchmark is constrained by limited semantic classes, and not all 3DGS scenes include precomputed vision-language embeddings. We leave these as important directions for future work and hope our benchmark and dataset will be helpful to advance language-grounded 3DGS scene understanding.

Acknowledgments and Disclosure of Funding

Yue Li is financially supported by TomTom, the University of Amsterdam and the allowance of Top consortia for Knowledge and Innovation (TKIs) from the Netherlands Ministry of Economic Affairs and Climate Policy. This work used Dutch national e-infrastructure with the support of the SURF Cooperative under grant no. NWO-2024.035. This work was partially supported by INSAIT, Sofia University “St. Kliment Ohridski” (Partially funded by the Ministry of Education and Science of Bulgaria’s support for INSAIT as part of the Bulgarian National Roadmap for Research Infrastructure), EU Horizon projects ELIAS (No.101120237), and ELLIOT (No.101214398). Some computational resources were provided by the Google Cloud Platform (GCP).

References

- [1] Armen Avetisyan, Christopher Xie, Henry Howard-Jenkins, Tsun-Yi Yang, Samir Aroudj, Suvam Patra, Fuyang Zhang, Duncan Frost, Luke Holland, Campbell Orme, et al. Scene-script: Reconstructing scenes with an autoregressive structured language model. In *European Conference on Computer Vision*, pages 247–263. Springer, 2024. 4
- [2] Jonathan T. Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P. Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields, 2021. 3
- [3] Gilad Baruch, Zhuoyuan Chen, Afshin Dehghan, Tal Dimry, Yuri Feigin, Peter Fu, Thomas Gebauer, Brandon Joffe, Daniel Kurz, Arik Schwartz, et al. Arkitscenes: A diverse real-world dataset for 3d indoor scene understanding using mobile rgb-d data. *arXiv preprint arXiv:2111.08897*, 2021. 5
- [4] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. *arXiv preprint arXiv:1709.06158*, 2017. 3, 5, 8
- [5] David Charatan, Sizhe Li, Andrea Tagliasacchi, and Vincent Sitzmann. pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. In *CVPR*, 2024. 2, 3
- [6] R. Qi Charles, Hao Su, Mo Kaichun, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 77–85, 2017. 2
- [7] Yuedong Chen, Hao-fei Xu, Chuanxia Zheng, Bohan Zhuang, Marc Pollefeys, Andreas Geiger, Tat-Jen Cham, and Jianfei Cai. Mvsplat: Efficient 3d gaussian splatting from sparse multi-view images. *arXiv preprint arXiv:2403.14627*, 2024. 2
- [8] Yutong Chen, Marko Mihajlovic, Xiyi Chen, Yiming Wang, Sergey Prokudin, and Siyu Tang. Splatformer: Point transformer for robust 3d gaussian splatting, 2025. 3, 4
- [9] Jiahuan Cheng, Jan-Nico Zaeche, Luc Van Gool, and Danda Pani Paudel. Occam’s lgs: A simple approach for language gaussian splatting. *arXiv preprint arXiv:2412.01807*, 2024. 2, 3, 7, 8, 9
- [10] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017. 3, 5, 7, 8, 9
- [11] Runyu Ding, Jihan Yang, Chuhui Xue, Wenqing Zhang, Song Bai, and Xiaojuan Qi. Pla: Language-driven open-vocabulary 3d scene understanding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7010–7019, 2023. 2
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. 2
- [13] Yasutaka Furukawa and Jean Ponce. Accurate, dense, and robust multiview stereopsis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(8):1362–1376, 2010. 2
- [14] Eric L. W. Grimson. *From Images to Surfaces: A Computational Study of the Human Early Visual System*. MIT Press, Cambridge, MA, USA, 1981. 2
- [15] Jun Guo, Xiaojian Ma, Yue Fan, Huaping Liu, and Qing Li. Semantic gaussians: Open-vocabulary scene understanding with 3d gaussian splatting. *arXiv preprint arXiv:2403.15624*, 2024. 3
- [16] Anna-Maria Halacheva, Jan-Nico Zaeche, Xi Wang, Danda Pani Paudel, and Luc Van Gool. Gaussianvlm: Scene-centric 3d vision-language models using language-aligned gaussian splats for embodied reasoning and beyond. *arXiv preprint arXiv:2507.00886*, 2025. 2

- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2
- [18] Heiko Hirschmüller. Stereo processing by semiglobal matching and mutual information. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(2):328–341, 2008. 2
- [19] Li Jiang, Shaoshuai Shi, and Bernt Schiele. Open-vocabulary 3d semantic segmentation with foundation models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21284–21294, 2024. 2
- [20] Joji Joseph, Bharadwaj Amrutar, and Shalabh Bhatnagar. Gradient-weighted feature back-projection: A fast alternative to feature distillation in 3d gaussian splatting, 2024. 2, 3, 7, 8, 9
- [21] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Proceedings of the Fourth Eurographics Symposium on Geometry Processing*, SGP ’06, page 61–70, 2006. 2
- [22] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), July 2023. 2
- [23] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023. 3
- [24] Justin Kerr, Chung Min Kim, Ken Goldberg, Angjoo Kanazawa, and Matthew Tancik. Lrf: Language embedded radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19729–19739, 2023. 3, 6, 10
- [25] Shakiba Kheradmand, Daniel Rebain, Gopal Sharma, Weiwei Sun, Yang-Che Tseng, Hossam Isack, Abhishek Kar, Andrea Tagliasacchi, and Kwang Moo Yi. 3d gaussian splatting as markov chain monte carlo. *Advances in Neural Information Processing Systems*, 37:80965–80986, 2025. 5
- [26] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4015–4026, 2023. 3
- [27] K.N. Kutulakos and S.M. Seitz. A theory of shape by space carving. In *Proceedings of the Seventh IEEE International Conference on Computer Vision*, volume 1, pages 307–314 vol.1, 1999. 2
- [28] Junha Lee, Chunghyun Park, Jaesung Choe, Yu-Chiang Frank Wang, Jan Kautz, Minsu Cho, and Chris Choy. Mosaic3d: Foundation dataset and model for open-vocabulary 3d segmentation. *arXiv preprint arXiv:2502.02548*, 2025. 4
- [29] Yue Li, Qi Ma, Runyi Yang, Huapeng Li, Mengjiao Ma, Bin Ren, Nikola Popovic, Nicu Sebe, Ender Konukoglu, Theo Gevers, Luc Van Gool, Martin R. Oswald, and Danda Pani Paudel. Scenesplat: Gaussian splatting-based scene understanding with vision-language pretraining, 2025. 2, 3, 4, 6, 7, 8, 9
- [30] Lu Ling, Yichen Sheng, Zhi Tu, Wentian Zhao, Cheng Xin, Kun Wan, Lantao Yu, Qianyu Guo, Zixun Yu, Yawen Lu, et al. D13dv-10k: A large-scale scene dataset for deep learning-based 3d vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22160–22169, 2024. 4
- [31] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *NeurIPS*, 2023. 2
- [32] Kunhao Liu, Fangneng Zhan, Jiahui Zhang, Muyu Xu, Yingchen Yu, Abdulmoteleb El Saddik, Christian Theobalt, Eric Xing, and Shijian Lu. Weakly supervised 3d open-vocabulary segmentation, 2024. 3

- [33] Xingchen Liu, Piyush Tayal, Jianyuan Wang, Jesus Zarzar, Tom Monnier, Konstantinos Tertikas, Jiali Duan, Antoine Toisoul, Jason Y. Zhang, Natalia Neverova, Andrea Vedaldi, Roman Shapovalov, and David Novotny. Uncommon objects in 3d. In *arXiv*, 2024. 3
- [34] Yuanhuiyi Lyu, Xu Zheng, Lutao Jiang, Yibo Yan, Xin Zou, Huiyu Zhou, Linfeng Zhang, and Xuming Hu. Realrag: Retrieval-augmented realistic image generation via self-reflective contrastive learning. *arXiv preprint arXiv:2502.00848*, 2025. 4
- [35] Yuanhuiyi Lyu, Xu Zheng, Jiazhou Zhou, and Lin Wang. Unibind: Llm-augmented unified and balanced representation space to bind them all. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26752–26762, 2024. 4
- [36] Qi Ma, Yue Li, Bin Ren, Nicu Sebe, Ender Konukoglu, Theo Gevers, Luc Van Gool, and Danda Pani Paudel. Shapesplat: A large-scale dataset of gaussian splats and their self-supervised pretraining. In *3DV*, 2024. 3
- [37] Qi Ma, Runyi Yang, Bin Ren, Nicu Sebe, Ender Konukoglu, Luc Van Gool, and Danda Pani Paudel. Cityloc: 6dof pose distributional localization for text descriptions in large-scale scenes with gaussian representation. *arXiv preprint arXiv:2501.08982*, 2025. 2
- [38] Juliette Marrie, Romain Menegaux, Michael Arbel, Diane Larlus, and Julien Mairal. Ludvig: Learning-free uplifting of 2d visual features to gaussian splatting scenes. *arXiv preprint arXiv:2410.14462*, 2024. 2, 3, 7, 8, 9
- [39] 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 *ECCV*, 2020. 2
- [40] 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. *Communications of the ACM*, 65(1):99–106, 2021. 3
- [41] Ashkan Mirzaei, Tristan Aumentado-Armstrong, Konstantinos G. Derpanis, Jonathan Kelly, Marcus A. Brubaker, Igor Gilitschenski, and Alex Levinshtein. SPIn-NeRF: Multiview segmentation and perceptual inpainting with neural radiance fields. In *CVPR*, 2023. 3
- [42] Richard A. Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J. Davison, Pushmeet Kohi, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *2011 10th IEEE International Symposium on Mixed and Augmented Reality*, pages 127–136, 2011. 2
- [43] Maxime Oquab, Timothée Darcet, Theo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Russell Howes, Po-Yao Huang, Hu Xu, Vasu Sharma, Shang-Wen Li, Wojciech Galuba, Mike Rabbat, Mido Assran, Nicolas Ballas, Gabriel Synnaeve, Ishan Misra, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. Dinov2: Learning robust visual features without supervision, 2023. 2, 3
- [44] Yatian Pang, Wenxiao Wang, Francis EH Tay, Wei Liu, Yonghong Tian, and Li Yuan. Masked autoencoders for point cloud self-supervised learning. In *European conference on computer vision*, pages 604–621. Springer, 2022. 4
- [45] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 165–174, 2019. 2
- [46] Songyou Peng, Kyle Genova, Chiyu Jiang, Andrea Tagliasacchi, Marc Pollefeys, Thomas Funkhouser, et al. Openscene: 3d scene understanding with open vocabularies. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 815–824, 2023. 2, 3

- [47] Minghan Qin, Wanhua Li, Jiawei Zhou, Haoqian Wang, and Hanspeter Pfister. Langsplat: 3d language gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20051–20060, June 2024. [2](#), [3](#), [7](#), [8](#), [9](#), [10](#)
- [48] Yansong Qu, Shaohui Dai, Xinyang Li, Jianghang Lin, Liujuan Cao, Shengchuan Zhang, and Rongrong Ji. Goi: Find 3d gaussians of interest with an optimizable open-vocabulary semantic-space hyperplane. *arXiv preprint arXiv:2405.17596*, 2024. [2](#), [3](#), [7](#), [8](#), [9](#)
- [49] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR, 2021. [3](#)
- [50] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763, 2021. [2](#), [4](#)
- [51] Adam Rashid, Satvik Sharma, Chung Min Kim, Justin Kerr, Lawrence Yunliang Chen, Angjoo Kanazawa, and Ken Goldberg. Language embedded radiance fields for zero-shot task-oriented grasping. In *7th Annual Conference on Robot Learning*, 2023. [3](#)
- [52] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. [3](#), [4](#)
- [53] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, Eric Mintun, Junting Pan, Kalyan Vasudev Alwala, Nicolas Carion, Chao-Yuan Wu, Ross Girshick, Piotr Dollár, and Christoph Feichtenhofer. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. [2](#), [4](#)
- [54] Bin Ren, Guofeng Mei, Danda Pani Paudel, Weijie Wang, Yawei Li, Mengyuan Liu, Rita Cucchiara, Luc Van Gool, and Nicu Sebe. Bringing masked autoencoders explicit contrastive properties for point cloud self-supervised learning. In *Proceedings of the Asian Conference on Computer Vision*, pages 2034–2052, 2024. [2](#)
- [55] Zhongzheng Ren, Aseem Agarwala, Bryan Russell, Alexander G Schwing, and Oliver Wang. Neural volumetric object selection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6133–6142, 2022. [3](#)
- [56] Mike Roberts, Jason Ramapuram, Anurag Ranjan, Atulit Kumar, Miguel Angel Bautista, Nathan Paczan, Russ Webb, and Joshua M Susskind. Hypersim: A photorealistic synthetic dataset for holistic indoor scene understanding. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10912–10922, 2021. [5](#)
- [57] Barbara Roessle, Norman Müller, Lorenzo Porzi, Samuel Rota Bulò, Peter Kotschieder, Angela Dai, and Matthias Nießner. L3dg: Latent 3d gaussian diffusion. In *SIGGRAPH Asia 2024 Conference Papers*, December 2024. [2](#), [3](#)
- [58] David Rozenberszki, Or Litany, and Angela Dai. Language-grounded indoor 3d semantic segmentation in the wild. In *European Conference on Computer Vision*, pages 125–141. Springer, 2022. [2](#)
- [59] Sayan Deb Sarkar, Ondrej Miksik, Marc Pollefeys, Daniel Barath, and Iro Armeni. Crossover: 3d scene cross-modal alignment. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. [2](#)
- [60] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. [2](#)

- [61] Jin-Chuan Shi, Miao Wang, Hao-Bin Duan, and Shao-Hua Guan. Language embedded 3d gaussians for open-vocabulary scene understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5333–5343, June 2024. [2](#)
- [62] Yongliang Shi, Runyi Yang, Zirui Wu, Pengfei Li, Caiyun Liu, Hao Zhao, and Guyue Zhou. City-scale mapping system with three-layer sampling and panoptic representation. *Available at SSRN 4568335*, 2023. [2](#)
- [63] Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, et al. The replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*, 2019. [3](#), [5](#)
- [64] Matthew Tancik, Ethan Weber, Evonne Ng, Ruilong Li, Brent Yi, Justin Kerr, Terrance Wang, Alexander Kristoffersen, Jake Austin, Kamyar Salahi, Abhik Ahuja, David McAllister, and Angjoo Kanazawa. Nerfstudio: A modular framework for neural radiance field development. In *ACM SIGGRAPH 2023 Conference Proceedings, SIGGRAPH '23*, 2023. [8](#)
- [65] Michael Tschannen, Alexey Gritsenko, Xiao Wang, Muhammad Ferjad Naeem, Ibrahim Alabdulmohsin, Nikhil Parthasarathy, Talfan Evans, Lucas Beyer, Ye Xia, Basil Mustafa, et al. Siglip 2: Multilingual vision-language encoders with improved semantic understanding, localization, and dense features. *arXiv preprint arXiv:2502.14786*, 2025. [3](#), [4](#)
- [66] Johanna Wald, Armen Avetisyan, Nassir Navab, Federico Tombari, and Matthias Nießner. Rio: 3d object instance re-localization in changing indoor environments. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7658–7667, 2019. [5](#)
- [67] Zeming wei, Junyi Lin, Yang Liu, Weixing Chen, Jingzhou Luo, Guanbin Li, and Liang Lin. 3daffordspat: Efficient affordance reasoning with 3d gaussians, 2025. [3](#)
- [68] Andrew P. Witkin. Recovering surface shape and orientation from texture. *Artificial Intelligence*, 17(1):17–45, 1981. [2](#)
- [69] Xiaoyang Wu, Daniel DeTone, Duncan Frost, Tianwei Shen, Chris Xie, Nan Yang, Jakob Engel, Richard Newcombe, Hengshuang Zhao, and Julian Straub. Sonata: Self-supervised learning of reliable point representations. In *CVPR*, 2025. [4](#)
- [70] Xiaoyang Wu, Li Jiang, Peng-Shuai Wang, Zhijian Liu, Xihui Liu, Yu Qiao, Wanli Ouyang, Tong He, and Hengshuang Zhao. Point transformer v3: Simpler, faster, stronger. In *CVPR*, 2024. [2](#), [4](#)
- [71] Xiaoyang Wu, Yixing Lao, Li Jiang, Xihui Liu, and Hengshuang Zhao. Point transformer v2: Grouped vector attention and partition-based pooling. *Advances in Neural Information Processing Systems*, 35:33330–33342, 2022. [4](#)
- [72] Yanmin Wu, Jiarui Meng, Haijie Li, Chenming Wu, Yahao Shi, Xinhua Cheng, Chen Zhao, Haocheng Feng, Errui Ding, Jingdong Wang, and Jian Zhang. Opengaussian: Towards point-level 3d gaussian-based open vocabulary understanding. In *Advances in Neural Information Processing Systems*, volume 37, 2024. [2](#), [3](#), [7](#), [8](#), [9](#)
- [73] Zirui Wu, Tianyu Liu, Liyi Luo, Zhide Zhong, Jianteng Chen, Hongmin Xiao, Chao Hou, Haozhe Lou, Yuantao Chen, Runyi Yang, Yuxin Huang, Xiaoyu Ye, Zike Yan, Yongliang Shi, Yiyi Liao, and Hao Zhao. Mars: An instance-aware, modular and realistic simulator for autonomous driving. *CICAI*, 2023. [2](#)
- [74] Jianfeng Xiang, Zelong Lv, Sicheng Xu, Yu Deng, Ruicheng Wang, Bowen Zhang, Dong Chen, Xin Tong, and Jiaolong Yang. Structured 3d latents for scalable and versatile 3d generation. *arXiv preprint arXiv:2412.01506*, 2024. [2](#)
- [75] Jihan Yang, Runyu Ding, Weipeng Deng, Zhe Wang, and Xiaojuan Qi. Regionplc: Regional point-language contrastive learning for open-world 3d scene understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19823–19832, 2024. [2](#)

- [76] Runyi Yang, Zhenxin Zhu, Zhou Jiang, Baijun Ye, Xiaoxue Chen, Yifei Zhang, Yuantao Chen, Jian Zhao, and Hao Zhao. Spectrally pruned gaussian fields with neural compensation. *arXiv preprint arXiv:2405.00676*, 2024. 8
- [77] Mingqiao Ye, Martin Danelljan, Fisher Yu, and Lei Ke. Gaussian grouping: Segment and edit anything in 3d scenes. In *ECCV*, 2024. 2
- [78] Vickie Ye, Ruilong Li, Justin Kerr, Matias Turkulainen, Brent Yi, Zhuoyang Pan, Otto Seiskari, Jianbo Ye, Jeffrey Hu, Matthew Tancik, and Angjoo Kanazawa. gsplat: An open-source library for gaussian splatting. *Journal of Machine Learning Research*, 26(34):1–17, 2025. 2, 5
- [79] Chandan Yeshwanth, Yueh-Cheng Liu, Matthias Nießner, and Angela Dai. Scannet++: A high-fidelity dataset of 3d indoor scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12–22, 2023. 3, 5, 8, 9
- [80] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 11975–11986, 2023. 3
- [81] Hao Zhang, Feng Li, Shilong Liu, Lei Zhang, Hang Su, Jun Zhu, Lionel M Ni, and Heung-Yeung Shum. Dino: Detr with improved denoising anchor boxes for end-to-end object detection. *arXiv preprint arXiv:2203.03605*, 2022. 3
- [82] Kai Zhang, Sai Bi, Hao Tan, Yuanbo Xiangli, Nanxuan Zhao, Kalyan Sunkavalli, and Zexiang Xu. Gs-irm: Large reconstruction model for 3d gaussian splatting. *European Conference on Computer Vision*, 2024. 2
- [83] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip HS Torr, and Vladlen Koltun. Point transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 16259–16268, 2021. 4
- [84] Yuhang Zheng, Xiangyu Chen, Yupeng Zheng, Songen Gu, Runyi Yang, Bu Jin, Pengfei Li, Chengliang Zhong, Zengmao Wang, Lina Liu, et al. Gaussiangrasper: 3d language gaussian splatting for open-vocabulary robotic grasping. *IEEE Robotics and Automation Letters*, 2024. 3
- [85] Ding Zhong, Xu Zheng, Chenfei Liao, Yuanhuiyi Lyu, Jialei Chen, Shengyang Wu, Linfeng Zhang, and Xuming Hu. Omnisam: Omnidirectional segment anything model for uda in panoramic semantic segmentation. *arXiv preprint arXiv:2503.07098*, 2025. 4
- [86] Shijie Zhou, Haoran Chang, Sicheng Jiang, Zhiwen Fan, Zehao Zhu, Dejia Xu, Pradyumna Chari, Suyu You, Zhangyang Wang, and Achuta Kadambi. Feature 3dgs: Supercharging 3d gaussian splatting to enable distilled feature fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21676–21685, 2024. 2, 3, 7, 8, 9
- [87] Yichao Zhou, Jingwei Huang, Xili Dai, Shichen Liu, Linjie Luo, Zhili Chen, and Yi Ma. Holicity: A city-scale data platform for learning holistic 3d structures, 2021. 3, 4, 8, 9
- [88] Xiangyang Zhu, Renrui Zhang, Bowei He, Ziyu Guo, Ziyao Zeng, Zipeng Qin, Shanghang Zhang, and Peng Gao. Pointclip v2: Prompting clip and gpt for powerful 3d open-world learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2639–2650, 2023. 2
- [89] Chen Ziwen, Hao Tan, Kai Zhang, Sai Bi, Fujun Luan, Yicong Hong, Li Fuxin, and Zexiang Xu. Long-irm: Long-sequence large reconstruction model for wide-coverage gaussian splats. *arXiv preprint 2410.12781*, 2024. 2, 3
- [90] Xingxing Zuo, Pouya Samangouei, Yunwen Zhou, Yan Di, and Mingyang Li. Fmgs: Foundation model embedded 3d gaussian splatting for holistic 3d scene understanding. *arXiv preprint arXiv:2401.01970*, 2024. 2, 3, 7, 8, 9

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: We will publicly release our proposed dataset, SceneSplat-49K, mentioned in Section 3 along with all benchmark codes. A complete summary of the benchmark experiments and key findings was mentioned in Section 4.1 and 4.2.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: Limitations are discussed in Section 5.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: The paper does not include theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We provide detailed implementations of both the proposed and benchmark methods in the main paper and supplementary materials. All code, benchmarks, and datasets will be released publicly.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: Our code, benchmark and datasets will be made public.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: The experimental settings are mentioned in Section 4.1, as well as the supplementary materials.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: We follow the standard practices of Language Gaussian Splatting for reporting zero-shot semantic segmentation metrics. In large-scale validation scenarios, the reported mean metric serves as a reliable performance indicator. Furthermore, performing all experiments was very computationally expensive.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We mention the computer resources used for all the experiments in Section 4.1 and for the dataset in Tab. 2, as well as the supplementary materials.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: All authors have thoroughly reviewed the full NeurIPS Code of Ethics and confirm that our work adheres to every guideline therein, no personally identifiable information is processed.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: Broader impacts are discussed in the supplementary materials.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [\[Yes\]](#)

Justification: Models and code are hosted behind a license that prohibits harmful uses. A brief “Acceptable Use” document accompanies all downloads. We reserve the right to revoke repository access if misuse is detected. A dedicated email address is provided for the community to report vulnerabilities or harmful behavior.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [\[Yes\]](#)

Justification: The original authors and owners of the data used to construct our dataset are mentioned in Section 3. The licence details for the mentioned data is also mentioned in the supplementary materials.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.

- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: The dataset creation workflow is detailed in Section 3 and the supplementary materials, and the benchmark methodology is fully described in 4 and the supplementary materials. The benchmark code will also be released publicly.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowd-sourcing or research with human participants.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowd-sourcing or research with human participants.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.