

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 OGGSPLAT: OPEN- VOCABULARY GAUSSIAN GROW- ING FOR EXPANDED FIELD-OF- VIEW

Anonymous authors

Paper under double-blind review

ABSTRACT

Reconstructing open-vocabulary 3D scenes from sparse views is both challenging and crucial, driven by the demands of emerging applications such as virtual reality and embodied AI. However, existing generalizable open-vocabulary 3D Gaussian Splatting methods struggle to reconstruct semantically enriched regions outside the input view cone. To address this limitation, we introduce OGGSplat, an open-vocabulary Gaussian growing method that extends the field-of-view for generalizable, semantically-enriched 3D scene reconstruction. Our key insight is that the semantic attributes of open-vocabulary Gaussians serve as strong priors for image extrapolation, ensuring both semantic consistency and visual plausibility. Specifically, once Gaussians with semantic attributes are initialized from sparse views, we introduce an RGB-semantic consistent inpainting module to selected rendered views. This module enables bidirectional control between an image diffusion model and a semantic diffusion model. The inpainted regions integrated with semantics are then lifted back into 3D space for efficient, progressive optimization of Gaussian parameters. To evaluate our method, we propose the Open-Vocabulary Gaussian Outpainting (OVGO) benchmark, which measures both the semantic and generative quality of the reconstructed open-vocabulary scenes. OGGSplat also demonstrates promising semantic-aware reconstruction capabilities when provided with two views captured directly from a smartphone camera.

1 INTRODUCTION

Building realistic and semantically meaningful 3D representations of the world has become a crucial goal in computer vision, driven by applications in robotics, virtual reality, and embodied AI. Beyond reconstructing vivid textures and accurate geometry, modern systems increasingly demand semantic awareness to support high-level understanding and interaction within 3D environments. This dual demand for geometric fidelity and semantic interpretability introduces new challenges for scene representation. Recent researches typically address this by combining open-vocabulary features with 3D reconstructive representations like 3D Gaussians (Kerbl et al., 2023). Approaches based on per-scene optimization (Qin et al., 2024; Shi et al., 2024; Qu et al., 2024; Qiu et al., 2024; Wu et al., 2024; Ye et al., 2024), which leverage dense multi-view inputs, achieve well-structured 3D geometry with fine-grained semantic alignment. In contrast, newly emerging feed-forward methods (Wang et al., 2024b; Hu et al., 2024) offer improved scalability and generalization across scenes by predicting semantic-aware 3D representations directly from sparse input views via a trained neural network.

Although recent feed-forward open-vocabulary 3D Gaussian reconstruction methods enable fast inference and efficiently handle sparse input views, their performance is often constrained by the limited scope of sparse inputs. For extremely sparse views, such as those with only two input perspectives, the semantic-aware 3D reconstructions can exhibit distorted geometry and semantically implausible content when extrapolated to novel viewpoints. For example, in Figure 1, the black regions in the novel view arise from the absence of Gaussians, resulting in incomplete reconstructions for both 3D scene and semantics. This highlights an urgent need for a generalizable open-vocabulary 3D reconstruction framework that can reliably expand the field-of-view while maintaining geometric coherence and semantic consistency. We argue that incorporating semantic cues from open-vocabulary features can provide valuable guidance in imagining plausible content for unseen regions, thus extending the application of generalizable reconstruction.

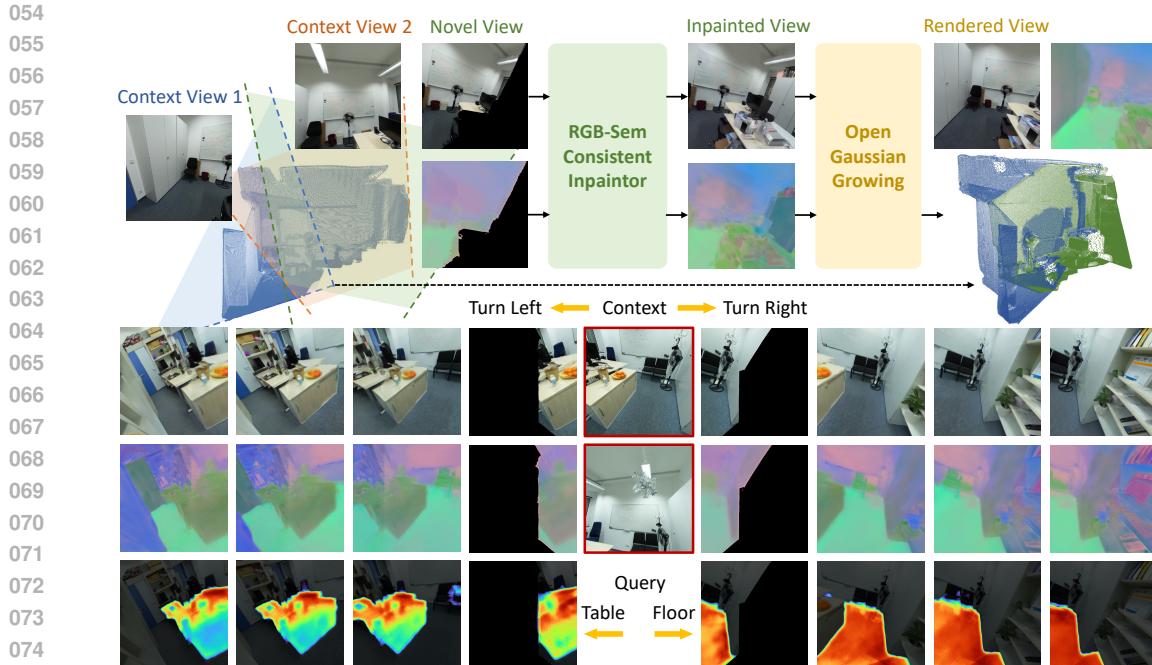


Figure 1: We propose **OGGSplat**, an open-vocabulary Gaussian growing method that expands the field-of-view of semantic-aware generalizable Gaussian reconstruction. The last three rows visualize the rendered images, semantic maps, and heatmaps obtained through open-vocabulary queries.

In this paper, we address the aforementioned challenge of generalizable open-vocabulary 3D reconstruction by introducing **OGGSplat**, an **Open-Vocabulary Gaussian Growing** framework designed to extrapolate semantically meaningful 3D Gaussians beyond the input view coverage. Our goal is to enhance open-vocabulary Gaussian representations with the capacity to grow new, semantic-aware Gaussians, thereby recovering occluded parts or expanding the field-of-view in open-vocabulary scenes reconstructed from sparse inputs. A key insight of our approach is that the semantic attributes inherent in open-vocabulary Gaussians provide a strong prior for semantically plausible extrapolation. To exploit this, OGGSplat uses a progressive Gaussian growing strategy that builds on the initial reconstruction from sparse views. Central to this process is a novel RGB-semantic consistent inpainting module, which enables bidirectional interaction between image and semantic inpainting: semantic maps guide image completion, while inpainted images refine the semantic features in return, ensuring pixel-level alignment. The synthesized RGB images and semantic maps are then used to efficiently optimize the newly introduced Gaussians.

We conduct extensive experiments on ScanNet++ (Yeshwanth et al., 2023) and introduce a novel Open-Vocabulary Gaussian Outpainting (OVGO) benchmark. OVGO evaluates both visual fidelity and semantic plausibility in extrapolated regions, enabling quantitative assessment via segmentation mean Intersection-over-Union (mIoU) alongside generative metrics such as Fréchet Inception Distance (FID) (Heusel et al., 2017). We further demonstrate OGGSplat on context images captured with a smartphone, highlighting its potential for portable applications.

In conclusion, the contributions can be summarized as: (1) We propose OGGSplat, the first work to expand the field-of-view for generalizable open-vocabulary Gaussian reconstructions. (2) We design an RGB-semantic consistent inpainting module that enforces bidirectional interaction between image and semantic map inpainting, and introduce a progressive Gaussian growing strategy to optimize new Gaussians from the inpainted content. (3) We establish the Open-Vocabulary Gaussian Outpainting (OVGO) benchmark, enabling semantic-aware evaluation on expanded regions.

2 RELATED WORK

3D Gaussian Splatting. Existing 3D Gaussian Splatting (3DGS) approaches can be broadly categorized by their optimization strategy and the number of input views. Early methods (Yu et al.,

108 2024; Lu et al., 2024; Fan et al., 2024a; Fu et al., 2024) rely on per-scene optimization with
 109 hundreds of images, achieving high-fidelity reconstructions but suffering from high computational cost
 110 and limited scalability. Subsequent works (Xiong et al., 2023; Paliwal et al., 2024; Chung et al.,
 111 2024; Zhu et al., 2024) aim to reconstruct scenes from only a few input views. More recently, generalizable
 112 3D reconstruction methods such as PixelSplat (Charatan et al., 2024) and MVSSplat (Chen
 113 et al., 2024) have been proposed to avoid per-scene optimization by leveraging feed-forward neural
 114 networks trained on large-scale datasets. Splatt3R (Smart et al., 2024) directly infers point clouds
 115 and Gaussian parameters from unposed image pairs, eliminating the need for camera calibration.

116 **Open-Vocabulary 3DGS with Per-Scene Optimization.** LangSplat (Qin et al., 2024) pioneers
 117 open-vocabulary 3D Gaussian Splatting (3DGS) by distilling knowledge from vision-language
 118 models such as CLIP (Radford et al., 2021) and DINO (Caron et al., 2021). Similarly, Fmgs (Zuo et al.,
 119 2025) and Feature 3DGS (Zhou et al., 2024) explicitly inject vision-language features into the 3DGS
 120 pipeline. Building upon these approaches, LEGaussians (Shi et al., 2024) and GOI (Qu et al., 2024)
 121 introduce quantization techniques to compress high-dimensional semantic embeddings into compact
 122 Gaussian representations. Ji et al. (2025) proposes a Feature Grid Mapping strategy to accelerate
 123 open-vocabulary queries for high-resolution reconstruction. Methods such as OpenGaussian (Wu
 124 et al., 2024) and Gaussian Grouping (Ye et al., 2024) leverage 2D open-vocabulary segmentation
 125 tools like SAM (Kirillov et al., 2023) to assign semantic labels to rendered images, without explicitly
 126 embedding semantics into the Gaussians themselves.

127 **Generalizable Open-Vocabulary 3DGS.** Building on generalizable 3DGS, GSemSplat (Wang
 128 et al., 2024b) extends DUSt3R (Wang et al., 2024a) and Splatt3R (Smart et al., 2024) by incorporating
 129 semantic prediction heads to jointly estimate open-vocabulary features together with Gaussian
 130 parameters, thereby enabling feed-forward scene reconstruction. Similarly, SparseLGS (Hu et al.,
 131 2024) leverages MAST3R (Leroy et al., 2024) and introduces a multi-view semantic alignment strategy
 132 to achieve generalizable Gaussian semantic reconstruction from sparse input images. LSM (Fan
 133 et al., 2024b) further explores semantic anisotropic Gaussians for reconstructing explicit radiance
 134 fields from only two unposed images, supporting real-time 3D perception and scene understanding.
 135 More recent works, such as GaussTR (Jiang et al., 2025) and Uni3R (Sun et al., 2025), advance
 136 open-vocabulary 3DGS by achieving higher-quality occupancy prediction and scene understanding.
 137 Nevertheless, a major limitation remains: these open-vocabulary generalizable methods often
 138 struggle to reconstruct regions outside the narrow visual field covered by the sparse input views.

3 APPROACH

141 As illustrated in Figure 2, OGGSplat comprises three main stages. First, in Section 3.1, we initialize
 142 an open-vocabulary 3D Gaussian representations from the unposed image pairs. Next, Section 3.2
 143 introduces the RGB-semantic consistent inpaintor, where we propose a bidirectional control mech-
 144 anism to ensure pixel-level alignment between inpainted semantic maps and RGB images. The
 145 semantic map guides the image completion process, while the inpainted image, in turn, refines the
 146 semantic features. Finally, to allow the semantic-aware 3D Gaussian structure to grow consistently
 147 with the generated content, we design a progressive open-vocabulary Gaussian growing strategy,
 148 detailed in Section 3.3. The second and third stages are applied iteratively to gradually expand the
 149 Gaussian representation beyond the initial field-of-view. In practical usage, OGGSplat takes as input
 150 any two uncalibrated images and processes them through the above three stages to produce an ex-
 151 panded open-vocabulary 3D Gaussian scene. This enables real-time rendering of both RGB images
 152 and their corresponding semantic feature maps from wider arbitrary viewpoints.

3.1 GENERALIZABLE OPEN- VOCABULARY GAUSSIAN INITIALIZATION

154 **Gaussian Reconstruction.** Given any two uncalibrated but overlapping images $I_1, I_2 \in \mathbb{R}^{H \times W \times 3}$
 155 with height H and width W , we adopt Splatt3R (Smart et al., 2024) to reconstruct an initial Gaussian
 156 $\mathcal{G}_0 \in \mathbb{R}^{N \times d}$ via a shared backbone, cross-attention interactions and multiple Gaussian heads. The
 157 number of Gaussian primitives $N = 2 \times H \times W$ corresponds to the total number of image pixels,
 158 while each Gaussian feature of dimension d is composed of the following components: (1) a 3D
 159 point position $p \in \mathbb{R}^3$, (2) a position offset $p_\Delta \in \mathbb{R}^3$, defining the Gaussian center $\mu = p + p_\Delta$, (3) a
 160 rotation quaternion $q \in \mathbb{R}^4$ and a scale vector $s \in \mathbb{R}^3$, together determining the covariance matrix Σ ,

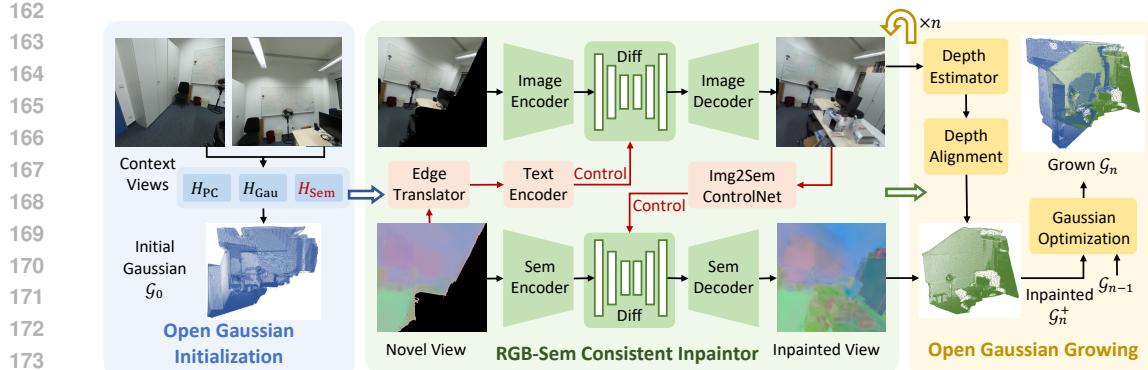


Figure 2: **OGGSplat Architecture.** We first initialize an open-vocabulary Gaussian reconstruction, injecting semantics via an additional semantic head. Then, the RGB-semantic consistent inpaintor applies bidirectional controls between images and semantic maps to ensure semantic plausibility and spatial alignment. Finally, the inpainted regions are lifted back to 3D and optimized to expand the Gaussians. The last two stages are performed iteratively to progressively grow the Gaussians.

(4) an opacity scalar $\alpha \in \mathbb{R}$, controlling the transparency of the Gaussian, and (5) a view-dependent appearance embedding represented by spherical harmonics $\mathbf{S} \in \mathbb{R}^{3 \times d_{\text{color}}}$ of d_{color} degrees.

Open Feature Injection. To incorporate open-vocabulary clues, we introduce an additional semantic head H_{sem} to predict semantic parameters $f \in \mathbb{R}^{d_{\text{sem}}}$ for each Gaussian primitive, inspired by GSemSplat (Wang et al., 2024b). Following common practice (Ye et al., 2024; Wang et al., 2024b), we set the semantic embedding dimension $d_{\text{sem}} = 16$ to reduce the computational overhead during Gaussian rendering. To supervise the predicted semantic features f , we adopt the well-optimized vision-language APE (Shen et al., 2024) model to efficiently obtain pixel-dense open-vocabulary semantic supervision signals $F^{\text{gt}} \in \mathbb{R}^{H \times W \times d_{\text{APE}}}$, where the APE semantic feature dimension $d_{\text{APE}} \gg d_{\text{sem}}$. To align the dimensionality, we train an autoencoder composed of a down-projection encoder \mathcal{E}_{\downarrow} that maps the APE features to d_{sem} , and a corresponding decoder \mathcal{D}_{\uparrow} that reconstructs the original features with minimal information loss. The semantic learning objective in this stage is formulated as a cosine similarity loss:

$$\mathcal{L}_{\text{sem}} = \sum_v \sum_{h,w} \left(1 - \cos \left(f_{v,h,w}, \mathcal{E}_{\downarrow} \left(f_{v,h,w}^{\text{gt}} \right) \right) \right), \quad (1)$$

where $h \in [0, H)$, $w \in [0, W)$ denote pixel coordinates and v represents target view index. The semantic feature $f_{v,h,w}$ is computed with α -blending, analogous to that used for RGB rendering.

3.2 RGB-SEMANTIC CONSISTENT INPAINTOR

Once the initial Gaussian \mathcal{G}_0 is reconstructed, we render RGB images I_v and their corresponding semantic maps F_v from novel viewpoints v . However, when rendering outside the vision cone of the context views, hollow regions often appear due to out-of-view areas and occlusion variations, as illustrated in Figure 1. While pre-trained inpainting diffusion models (Rombach et al., 2022; Lugmayr et al., 2022; Xie et al., 2023) can partially address this issue, maintaining pixel-wise consistency between inpainted images and their semantic maps remains challenging. This spatial misalignment will be inherited by the following Gaussian growing process and can lead to increasingly severe semantic inconsistencies as the scene expands. Fortunately, we observe that although the semantic modality introduces challenges, it also offers valuable guidance: the partial semantic information, especially around the boundaries of incomplete regions, can be translated into explicit textual prompts to guide image inpainting. Symmetrically, inpainted RGB images can provide pixel-wise appearance cues to control semantic map completion. Therefore, we propose bidirectional controls between the RGB branch Diff_{rgb} and the semantic branch Diff_{sem} , allowing them to mutually enhance each other during the inpainting process.

Semantic-to-RGB Control. To define the inpainting mask that determines whether a pixel should be inpainted, we rely on the rendered opacity α of each pixel. Similar to color rendering, we render an opacity map A , and then derive the inpainting mask M_v for each view v by applying a pre-defined

216 threshold τ . For simplicity, we omit the view subscript v in the following discussion.
 217

$$218 \quad 219 \quad 220 \quad A_{h,w} = \sum_{i \in \Theta_{h,w}} \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j), \quad M_{h,w} = \mathbb{1}[A_{h,w} < \tau], \quad (2)$$

221 where $\Theta_{h,w}$ denotes the set of Gaussians contributing to the pixel at coordinate (h, w) .
 222

223 Then we design an *Edge Translator* to extract semantic concepts near the inpainting boundaries
 224 defined by the mask M , providing clearer guidance for hollow region inpainting. Specifically, we
 225 first identify pixels along the boundary as Ω_{edge} . The corresponding semantic features f_{edge} of these
 226 boundary pixels are then decoded into a higher-dimensional space using our pre-trained decoder \mathcal{D}_{\uparrow} :
 227

$$227 \quad g_{\text{edge}} = \mathcal{D}_{\uparrow}(f_{\text{edge}}), \text{ for pixels in } \Omega_{\text{edge}} \quad (3)$$

228 Simultaneously, we prepare a set of candidate classes $\mathcal{C}_{\text{cand}}$, consisting of the top 100 semantic
 229 categories in our training dataset. These categories are encoded into the same feature space as g_{edge} .
 230 We then compute the cosine similarity between g_{edge} and g_{cand} to perform pixel-wise segmentation:
 231

$$232 \quad c_{\text{edge}} = \text{argmax}_{c_i \in \mathcal{C}_{\text{cand}}} \cos(g_{\text{edge}}, g_{c_i}), \quad (4)$$

233 In this way, we can obtain a set of semantic categories $\mathcal{C}_{\text{edge}}$ that are most relevant to the inpainting
 234 region. Based on these categories, we generate a prompt text T in the format of “*a room with cate₁,*
 235 *cate₂, ..., and cate_i*”, which is used to guide the diffusion-based RGB image inpainting model:
 236

$$237 \quad I^{\text{inp}} = \text{Diff}_{\text{rgb}}(I, M, T), \quad (5)$$

238 **RGB-to-Semantic Control.** Inspired by ControlNet (Zhang et al., 2023), we also design an RGB-
 239 to-Semantic control module to ensure that the generated semantic content aligns well with the cor-
 240 responding regions in the RGB image. Formally, the completed semantic map is computed as:
 241

$$242 \quad F^{\text{inp}} = \text{Diff}_{\text{sem}}(F, M, T, \text{ControlNet}(I^{\text{inp}})), \quad (6)$$

243 where F is the incomplete rendered semantic feature map, and $\text{ControlNet}(I^{\text{inp}})$ denotes the con-
 244 trol module conditioned on the inpainted image I^{inp} . Please refer to the ControlNet paper or Ap-
 245 pendix B.4 for further details. This module guides the semantic generation process, ensuring both
 246 structural and appearance consistency between the predicted semantic features and the RGB content.
 247

248 3.3 OPEN-VOCABULARY GAUSSIAN GROWING

250 Obtaining the inpainted RGB images and semantic feature maps from selected views is not the
 251 final step of our pipeline. These results must be aggregated back into the initial Gaussian \mathcal{G}_0
 252 to enable real-time rendering from arbitrary novel viewpoints. For selected anchor views $V =$
 253 $\{v_3, v_4, \dots, v_a\}$, we perform iterative inpainting and progressively incorporate the newly completed
 254 regions into the Gaussian. At each iteration n , a new view is rendered based on the currently aggre-
 255 gated Gaussians \mathcal{G}_{n-1} and the newly inpainted content \mathcal{G}_n^+ is fused into this representation. Below,
 256 we break down a single iteration and describe the Gaussian growing process in detail.
 257

258 The inpainted image I^{inp} and semantic map F^{inp} will serve as supervision targets for the newly
 259 grown Gaussians. However, establishing 3D geometry from a single novel view is inherently ill-
 260 posed, especially in regions that are newly generated during inpainting. To enrich these views with
 261 structural knowledge, we adopt custom depth estimation model (Piccinelli et al., 2024; Yang et al.,
 262 2024a;b) to predict an absolute depth map D^{inp} from I^{inp} . This depth map is then used to lift
 263 pixels back into 3D space, forming a point cloud in the global coordinate system. The resulting 3D
 264 points are used to initialize the position of the incremental Gaussian set \mathcal{G}^+ , which is progressively
 265 integrated into the scene representation.
 266

$$267 \quad 268 \quad 269 \quad P^+ = \text{proj}(D^{\text{inp}}, v^{\text{inp}}, v_1, K) \cdot \beta, \text{ where } \beta = \frac{\sqrt{\frac{1}{M} \sum_{i=1}^M \|p_i^{\text{ori}}\|_2^2}}{\sqrt{\frac{1}{N} \sum_{i=1}^N \|p_i^{\text{new}}\|_2^2}} \quad (7)$$

270 where v^{inp} and v_1 are the camera poses corresponding to the images I^{inp} and I_1 , respectively, and K
 271 denotes the intrinsic camera parameters. The scale factor β is introduced to align the newly projected
 272

270 point cloud with the original 3D space in terms of depth. $p^{\text{ori}}, p^{\text{new}}$ denote the original and newly
 271 projected 3D points within the overlapping regions, while M and N represent the respective number
 272 of points in each set. It is worth noticing that scaling point coordinates alone does not ensure perfect
 273 alignment. Nonetheless, it offers an efficient and approximate initialization, since the entire scene is
 274 constructed with respect to the normalized coordinate system of the first view.

275 At the n^{th} iteration, after merging \mathcal{G}_{n-1} with the newly initialized Gaussians \mathcal{G}_n^+ , we perform effi-
 276 cient per-scene optimization to update the grown Gaussian \mathcal{G}_n . This optimization is supervised by
 277 the original sparse context views, previously and newly inpainted views. The objective function is:
 278

$$279 \quad \mathcal{L} = \lambda_{\text{rgb}} \cdot \mathcal{L}_{\text{rgb}} + \lambda_{\text{feat}} \cdot \mathcal{L}_{\text{feat}}, \quad (8)$$

$$280 \quad \text{where } \mathcal{L}_{\text{rgb}} = \lambda_1 \cdot \mathcal{L}_{\text{L1}}(I^{\text{r}}, I^{\text{inp}}) + \lambda_2 \cdot \mathcal{L}_{\text{SSIM}}(I^{\text{r}}, I^{\text{inp}}), \text{ and } \mathcal{L}_{\text{feat}} = 1 - \cos(F^{\text{r}}, F^{\text{inp}}) \quad (9)$$

281 where λ_1 and λ_2 balance pixel-wise accuracy and perceptual similarity, while λ_{rgb} and λ_{feat} control
 282 the overall contributions of the photometric and semantic losses, respectively. $I^{\text{r}}, F^{\text{r}}$ denote the
 283 rendered RGB images and semantic features from the optimizing Gaussian from v^{inp} .
 284

285 4 EXPERIMENTS

288 4.1 THE OPEN-VOCABULARY GAUSSIAN OUTPAINTING (OVGO) BENCHMARK

289 To effectively evaluate both the visual fidelity and semantic plausibility of OGGSplat in extrapolated
 290 regions, we introduce a novel Open-Vocabulary Gaussian Outpainting benchmark based on the
 291 ScanNet++ (Yeshwanth et al., 2023) dataset. Detailed information can be found in the Appendix B.1.
 292

293 **Data Composition.** The OVGO benchmark covers all 50 validation scenes from ScanNet++. For
 294 each scene, we select *1 image pair* as the context views of inputs. To ensure consistency in data sam-
 295 pling and maintain temporal coherence, the context views are chosen as the 1st and 10th frames. This
 296 selection introduces moderate viewpoint variation while preserving semantic continuity, enabling a
 297 more meaningful evaluation of extrapolated content. For evaluation, we uniformly sample *16 novel*
 298 *camera poses* within a horizontal range of $[-60^\circ, 60^\circ]$ and a vertical range of $[-20^\circ, 20^\circ]$ around the
 299 pose of the context image I_1 . Novel RGB images and semantic maps are directly rendered from the
 300 reconstructed Gaussians at these poses and used as evaluation samples. Considering randomness,
 301 we repeat the experiment five times and report the average results.

302 **Visual Fidelity Evaluation.** We adopt the Fréchet Inception Distance (FID) (Heusel et al., 2017)
 303 to evaluate the statistical similarity between rendered and real images. For FID computation, all
 304 images from the validation split of the ScanNet++ dataset are used as the reference distribution.
 305 FID is then calculated between this reference distribution and the distribution of the newly rendered
 306 images. However, we observe that the limited number of generated images can negatively affect the
 307 stability of the FID metric. To address this, we increase the context views from *one pair* to *ten pairs*
 308 per scene, while maintaining a frame interval of 10 within each pair. This expands the number of
 309 newly rendered images by a factor of ten, resulting in a more stable and reliable FID evaluation.

310 **Semantic Plausibility Evaluation.** We assess semantic plausibility by performing open-vocabulary
 311 semantic segmentation on novel views. To enable a more comprehensive analysis, we separately
 312 evaluate performance in low-confidence (mIoU_L) and high-confidence (mIoU_H) regions. Low-
 313 confidence regions are defined as pixels in novel views where the accumulated opacity of the ini-
 314 tial Gaussians falls below 0.3, corresponding to occluded or out-of-view areas. Since evaluating
 315 these regions emphasizes semantic consistency in extrapolated areas, mIoU_L serves as our primary
 316 metric. For reference, we additionally report mIoU_H , which measures semantic plausibility in high-
 317 confidence regions, defined as pixels where the accumulated opacity of the initial Gaussians exceeds
 318 0.3, thus indicating areas reliably rendered by the original representation.

319 Since ground truth semantic annotations are unavailable for extrapolated regions, we generate them
 320 using five state-of-the-art open-vocabulary 2D semantic segmentation models (Xu et al., 2023; Shen
 321 et al., 2024; Zeng et al., 2024; Yu et al., 2023; Jiao et al., 2024). Their predictions are aggregated via
 322 a majority voting scheme. To assess the quality of semantic segmentation, we follow the protocol
 323 in (Kerr et al., 2023; Shi et al., 2024; Qin et al., 2024) by computing a relevancy score for each text
 query. More details about relevancy score are provided in Appendix B.5. To ensure generality, we
 324 retain only those predicted mask regions with a relevancy score exceeding 50% as the final binary

324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377

Table 1: **Open-Vocabulary Gaussian Outpainting (OVGO) benchmark results.** We compare generative metric FID and semantic metric mIoU (%) between OGGSplat and previous methods.

Methods	Generation		Segmentation (IoU↑)										
	FID↓	mIoU _H	mIoU _L	wall	ceiling	floor	table	door	(s)cabinet	chair	(b)shelf	box	bed
LangSplat	50.4	13.5	6.9	29.0	13.4	15.8	1.8	4.0	1.3	2.5	0.0	0.8	0.0
Splatt3R	46.4	24.9	6.0	10.1	2.1	18.9	5.1	0.0	1.6	13.8	0.3	0.0	2.3
OGGSplat	37.5	25.4	17.6	45.6	0.1	58.3	13.3	5.4	3.7	21.4	7.4	3.1	18.0

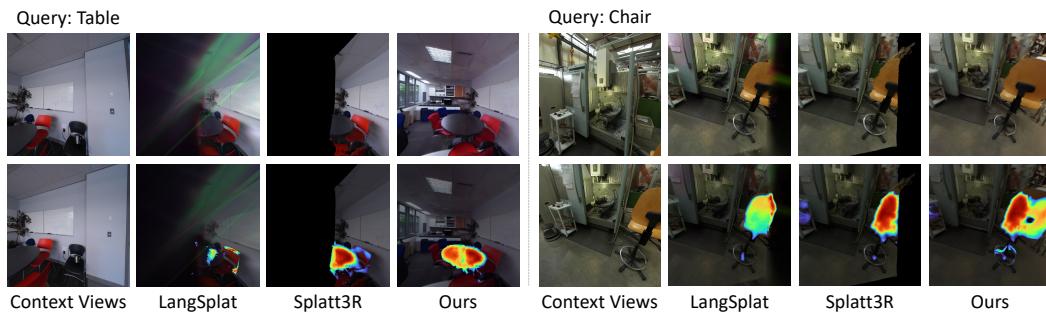


Figure 3: **Qualitative comparisons between LangSplat, Splatt3R, and OGGSplat on the OVGO benchmark.** The first row presents RGB images rendered from novel, out-of-scope viewpoints. The second row visualizes the heatmap when querying different text concepts.

mask. This filtering mechanism makes our evaluation suitable even for images where a specific category may be absent. During evaluation, we focus on 10 commonly used semantic categories selected from the top 20 classes in ScanNet++, such as *wall*, *floor*, *chair*, *table*, and others.

4.2 MAIN RESULTS

Baseline Methods for Comparison. We select two representative open-vocabulary Gaussian baselines for comparison: LangSplat (Qin et al., 2024), a per-scene optimization model, and Splatt3R (Smart et al., 2024), a generalizable model. LangSplat relies heavily on accurate initialization via COLMAP (Schonberger & Frahm, 2016), which becomes unreliable when only two input images are available. To address this limitation and enable fair comparison, we initialize LangSplat using point cloud positions predicted by Splatt3R, allowing the model to focus more effectively on learning semantic representations. Meanwhile, as vanilla Splatt3R does not support open-vocabulary semantic prediction in its original form, we extend it with a semantic head trained in our first stage in Section 3.1. During evaluation, for all models, we consider only the regions rendered by Gaussians with an accumulated opacity greater than 0.01 as valid predictions for computing the IoU scores. This threshold filters out low-confidence regions and ensures consistency across models.

Quantitative Comparisons. In Table 1, we compare LangSplat (Qin et al., 2024), Splatt3R (Smart et al., 2024), and OGGSplat on the OVGO benchmark. OGGSplat consistently outperforms the baselines by a significant margin on both visual fidelity (FID) and semantic plausibility (mIoU). It’s worth noticing that the overall FID remains relatively high across all methods. The main reason is the limited number of context pairs available in the validation set, which constrains data diversity. We are unable to sample more pairs because some scenes in the ScanNet++ validation set are relatively small. To maintain a consistent sampling ratio across all validation scenes, we limit the number of context pairs to 10 per scene. Regarding semantic plausibility, in the low-confidence regions, which are our primary focus, OGGSplat achieves notably better performance on common large objects such as *chair*, *table*, and *bed*. However, the model performs relatively worse on the *ceiling* class. We attribute this to the limitations of the APE encoding, as well as the difficulty of the Splatt3R backbone in distinguishing between the *ceiling* and *wall* with similar appearance in color and texture. We believe this limitation can be addressed in future work by leveraging more powerful vision-language models and more superior generalizable Gaussian reconstruction methods. For high-confidence regions, although these are not our primary focus, it is still notable that OGGSplat slightly outperforms the Splatt3R backbone. This improvement is mainly due to our method’s ability to recover small chunks within the view cone that initially have relatively low confidence.

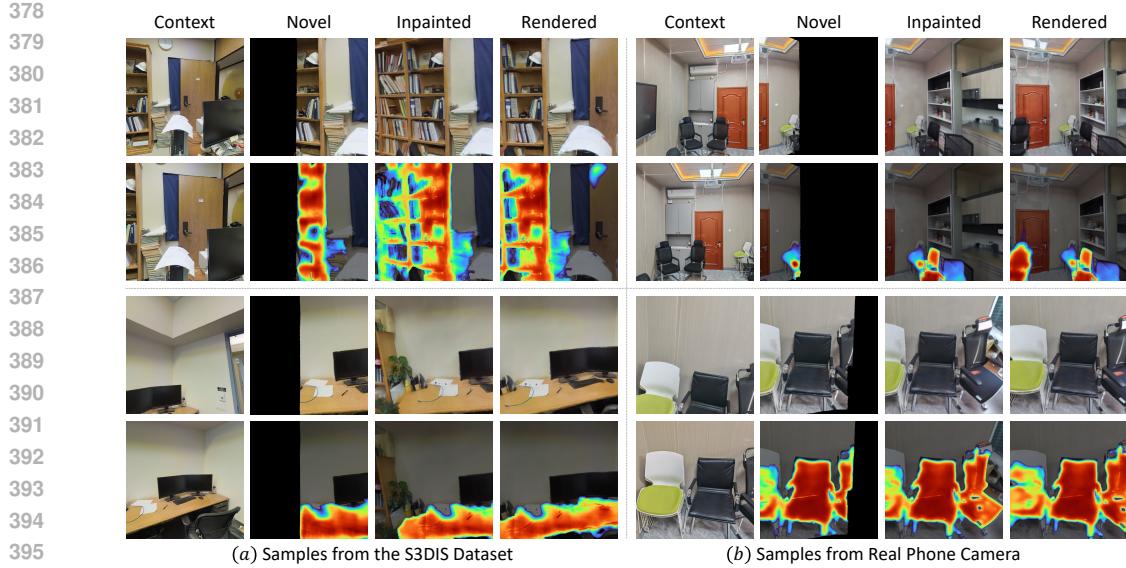


Figure 4: **Zero-shot generalization to out-of-distribution data.** (a) Context views are taken from the S3DIS dataset, with queries *bookshelf* and *table*. (b) Context views are captured using a **phone camera**, with the query *chair*. In both cases, we directly apply the model trained on ScanNet++.

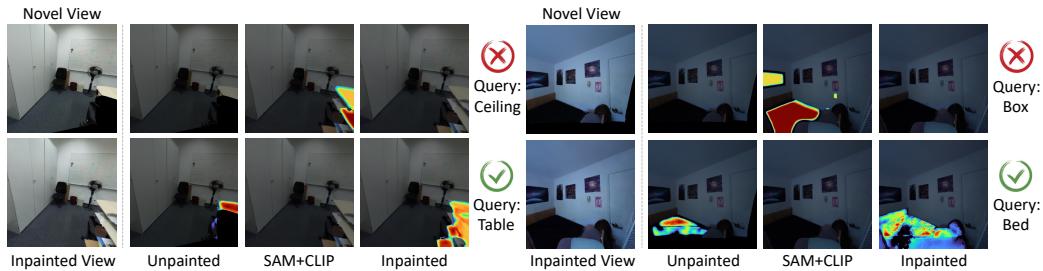


Figure 5: **Ablations on the effect of semantic diffusion model.** We compare open-vocabulary predictions between the SAM+CLIP offline method and our semantic diffusion inpainting module.

Qualitative Comparisons. We conduct extensive qualitative comparisons with baseline methods and illustrate them in Figure 3. OGGSplat performs better in both novel rendered images and open-vocabulary querying. Regarding rendered images, LangSplat tends to overfit the context views, resulting in blurry renderings from novel viewpoints, even when the Gaussian positions have been initialized. Splatt3R, on the other hand, exhibits large black regions in areas outside the input views. In contrast, OGGSplat reasonably extrapolates unseen regions by leveraging semantic information. Regarding open-vocabulary querying, both LangSplat and Splatt3R are limited to input vision cones. OGGSplat, however, is capable of accurately identifying and querying objects even in previously unseen regions, demonstrating stronger generalization and semantic understanding capabilities.

Zero-Shot Generalization to Out-of-Distribution Data. Apart from ScanNet++ used for training, we also test OGGSplat’s zero-shot ability on different distributions. As shown in Figure 4, OGGSplat successfully reconstructs expanded semantic-aware scenes on the S3DIS (Armeni et al., 2016) dataset. We further demonstrate the practicality of OGGSplat on portable devices in column (b), where the inputs are captured by a phone camera. The inpainted image and category-specific query show promising results, highlighting OGGSplat’s potential for applications in daily life.

4.3 ABLATION STUDIES

In Section 3.2, we introduced the RGB-semantic consistent inpainting module. In this section, we first highlight the importance of the semantic diffusion branch, followed by comprehensive ablations on the OVG0 benchmark to evaluate the effectiveness of the proposed bidirectional control strategy.

Semantic Diffusion Model. To obtain reliable semantics for inpainted regions, we train a semantic diffusion module. A simple alternative is using an offline open-vocabulary semantic segmentation

432 **Table 2: Ablations on the OVGO benchmark evaluating the impact of the bidirectional control**
 433 **strategy.** The performance is measured by mIoU (%) across various semantic categories.

Control Type		Segmentation Results (IoU \uparrow)										
S \rightarrow RGB	RGB \rightarrow S	mIoU _L	wall	ceiling	floor	table	door	(s)cabinet	chair	(b)shelf	box	bed
\times	\checkmark	16.6	45.8	0.1	56.8	12.3	4.6	2.8	19.3	6.3	3.6	15.1
\checkmark	\times	14.4	43.0	0.1	47.6	10.3	5.0	3.5	16.6	2.4	2.5	12.7
\checkmark	\checkmark	17.6	45.6	0.1	58.3	13.3	5.4	3.7	21.4	7.4	3.1	18.0

	Context Views			Rendered View			Context Views			Rendered View			Context Views			Rendered View		
	Inpainted Sem	Query: Door	Inpainted RGB	Inpainted Sem	Query: Bed	Inpainted RGB	Inpainted Sem	Query: Chair	Inpainted RGB	Inpainted Sem	Query: Door	Inpainted RGB	Inpainted Sem	Query: Bed	Inpainted RGB	Inpainted Sem	Query: Chair	Inpainted RGB
w/o Sem \rightarrow RGB Control																		
OGGSplat																		

455 **Figure 6: Qualitative comparison of bidirectional control.** Row 1 shows the context images and
 456 the incomplete renderings from novel views. Rows 2 to 4 correspond to the ablation settings in
 457 Table 2, where each variant removes one of the control mechanisms to examine its individual effect.

458
 459 model, e.g., SAM (Kirillov et al., 2023)+CLIP (Radford et al., 2021). However, this often causes
 460 semantic inconsistency with the original Gaussian, especially for partially visible objects (Figure 5).
 461 In contrast, our semantic diffusion model preserves semantic consistency in unpainted areas and
 462 leverages context priors to improve accuracy in the inpainted regions, ensuring new content aligns
 463 with existing scene semantics and enhancing overall reconstruction quality.

464 **Semantic-to-RGB Control.** With access to open-vocabulary semantics, we propose an edge trans-
 465 lator to extract semantic cues from the Gaussian boundaries and guide the completion. In the first
 466 row of Table 2, we remove the edge translator and instead use a generic description (“a room”) as the
 467 text prompt. As a result, semantic segmentation performance across most categories decreases. This
 468 degradation is also evident in the qualitative comparison in Figure 6, where the generated content
 469 appears more ambiguous and less semantically grounded. These results validate the effectiveness of
 470 our semantic-to-RGB control in guiding high-fidelity, semantically consistent Gaussian growth.

471 **RGB-to-Semantic Control.** In OGGSplat, the semantic inpainting model is explicitly controlled
 472 by inpainted images. We remove it in the second row of Table 2 and the third row of Figure 6.
 473 Without RGB-to-semantic control, the generated RGB images and semantic maps exhibit poor spa-
 474 tial alignment, leading to significantly degraded segmentation accuracy. In contrast, introducing the
 475 RGB-to-semantic control clearly improves spatial consistency and yields much better performance.

477 5 CONCLUSION

478
 479 In this paper, we present OGGSplat, a generalizable open-vocabulary 3D Gaussian growing method
 480 for expanded field-of-view. By leveraging semantic cues and introducing RGB-semantic consis-
 481 tent inpainting with bidirectional control, OGGSplat effectively expands the view while maintain-
 482 ing visual fidelity and semantic coherence. Out-of-view regions are then progressively refined through
 483 efficient Gaussian optimization. We also propose the Open-Vocabulary Gaussian Outpainting bench-
 484 mark for semantic-aware evaluation on expanded regions. Extensive experiments show that OG-
 485 GSplat effectively extrapolates beyond the input view cone while keeping RGB-semantic alignment,
 marking a significant step forward in generalizable and flexible open-vocabulary 3D reconstruction.

486 REFERENCES
487

488 Iro Armeni, Ozan Sener, Amir R Zamir, Helen Jiang, Ioannis Brilakis, Martin Fischer, and Silvio
489 Savarese. 3d semantic parsing of large-scale indoor spaces. In *CVPR*, 2016.

490 Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang,
491 Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*,
492 2025.

493 Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and
494 Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021.

495 David Charatan, Sizhe Lester Li, Andrea Tagliasacchi, and Vincent Sitzmann. pixelsplat: 3d gaus-
496 sian splats from image pairs for scalable generalizable 3d reconstruction. In *CVPR*, 2024.

497 Yuedong Chen, Haofei Xu, Chuanxia Zheng, Bohan Zhuang, Marc Pollefeys, Andreas Geiger, Tat-
498 Jen Cham, and Jianfei Cai. Mvsplat: Efficient 3d gaussian splatting from sparse multi-view
499 images. In *ECCV*, 2024.

500 501 Jaeyoung Chung, Jeongtaek Oh, and Kyoung Mu Lee. Depth-regularized optimization for 3d gaus-
502 sian splatting in few-shot images. In *CVPR*, 2024.

503 504 Zheng Ding, Jieke Wang, and Zhuowen Tu. Open-vocabulary panoptic segmentation maskclip.
505 2022.

506 507 Zhiwen Fan, Kevin Wang, Kairun Wen, Zehao Zhu, Dejia Xu, Zhangyang Wang, et al. Lightgaus-
508 sian: Unbounded 3d gaussian compression with 15x reduction and 200+ fps. *NeurIPS*, 2024a.

509 510 Zhiwen Fan, Jian Zhang, Wenyan Cong, Peihao Wang, Renjie Li, Kairun Wen, Shijie Zhou, Achuta
511 Kadambi, Zhangyang Wang, Danfei Xu, et al. Large spatial model: End-to-end unposed images
512 to semantic 3d. *Advances in neural information processing systems*, 37:40212–40229, 2024b.

513 514 Yang Fu, Sifei Liu, Amey Kulkarni, Jan Kautz, Alexei A Efros, and Xiaolong Wang. Colmap-free
515 3d gaussian splatting. In *CVPR*, 2024.

516 517 Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter.
518 Gans trained by a two time-scale update rule converge to a local nash equilibrium. *NeurIPS*,
519 2017.

520 Jun Hu, Zhang Chen, Zhong Li, Yi Xu, and Juyong Zhang. Sparselgs: Sparse view language
521 embedded gaussian splatting. *arXiv preprint arXiv:2412.02245*, 2024.

522 523 Yuzhou Ji, He Zhu, Junshu Tang, Wuyi Liu, Zhizhong Zhang, Xin Tan, and Yuan Xie. Fastlgs:
524 Speeding up language embedded gaussians with feature grid mapping. In *Proceedings of the
AAAI Conference on Artificial Intelligence*, volume 39, pp. 3922–3930, 2025.

525 526 Haoyi Jiang, Liu Liu, Tianheng Cheng, Xinjie Wang, Tianwei Lin, Zhizhong Su, Wenyu Liu, and
527 Xinggang Wang. Gausstr: Foundation model-aligned gaussian transformer for self-supervised 3d
528 spatial understanding. In *Proceedings of the Computer Vision and Pattern Recognition Confer-
ence*, pp. 11960–11970, 2025.

529 530 Siyu Jiao, Hongguang Zhu, Jiannan Huang, Yao Zhao, Yunchao Wei, and Humphrey Shi. Collabo-
531 rative vision-text representation optimizing for open-vocabulary segmentation. In *ECCV*, 2024.

532 533 Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splat-
534 ting for real-time radiance field rendering. *ACM Trans. Graph.*, 2023.

535 Justin Kerr, Chung Min Kim, Ken Goldberg, Angjoo Kanazawa, and Matthew Tancik. Lerf: Lan-
536 guage embedded radiance fields. In *ICCV*, 2023.

537 Diederik P Kingma, Max Welling, et al. Auto-encoding variational bayes, 2013.

538 539 Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete
540 Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *ICCV*,
2023.

540 Vincent Leroy, Yohann Cabon, and Jérôme Revaud. Grounding image matching in 3d with mast3r.
 541 In *ECCV*, 2024.

542 Tao Lu, Mulin Yu, Lining Xu, Yuanbo Xiangli, Limin Wang, Dahua Lin, and Bo Dai. Scaffold-gs:
 543 Structured 3d gaussians for view-adaptive rendering. In *CVPR*, 2024.

544 Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool.
 545 Repaint: Inpainting using denoising diffusion probabilistic models. In *CVPR*, 2022.

546 Avinash Paliwal, Wei Ye, Jinhui Xiong, Dmytro Kotovenko, Rakesh Ranjan, Vikas Chandra, and
 547 Nima Khademi Kalantari. Coherentgs: Sparse novel view synthesis with coherent 3d gaussians.
 548 In *ECCV*, 2024.

549 Luigi Piccinelli, Yung-Hsu Yang, Christos Sakaridis, Mattia Segu, Siyuan Li, Luc Van Gool, and
 550 Fisher Yu. Unidepth: Universal monocular metric depth estimation. In *CVPR*, 2024.

551 Minghan Qin, Wanhua Li, Jiawei Zhou, Haoqian Wang, and Hanspeter Pfister. Langsplat: 3d lan-
 552 guage gaussian splatting. In *CVPR*, 2024.

553 Ri-Zhao Qiu, Ge Yang, Weijia Zeng, and Xiaolong Wang. Feature splatting: Language-driven
 554 physics-based scene synthesis and editing. *arXiv preprint arXiv:2404.01223*, 2024.

555 Yansong Qu, Shaohui Dai, Xinyang Li, Jianghang Lin, Liujuan Cao, Shengchuan Zhang, and Ron-
 556 grong Ji. Goi: Find 3d gaussians of interest with an optimizable open-vocabulary semantic-space
 557 hyperplane. In *ACM MM*, 2024.

558 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
 559 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual
 560 models from natural language supervision. In *ICML*, 2021.

561 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
 562 resolution image synthesis with latent diffusion models. In *CVPR*, 2022.

563 Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *CVPR*, 2016.

564 Yunhang Shen, Chaoyou Fu, Peixian Chen, Mengdan Zhang, Ke Li, Xing Sun, Yunsheng Wu,
 565 Shaohui Lin, and Rongrong Ji. Aligning and prompting everything all at once for universal visual
 566 perception. In *CVPR*, 2024.

567 Jin-Chuan Shi, Miao Wang, Hao-Bin Duan, and Shao-Hua Guan. Language embedded 3d gaussians
 568 for open-vocabulary scene understanding. In *CVPR*, 2024.

569 Brandon Smart, Chuanxia Zheng, Iro Laina, and Victor Adrian Prisacariu. Splatt3r: Zero-shot
 570 gaussian splatting from uncalibrated image pairs. *arXiv preprint arXiv:2408.13912*, 2024.

571 Xiangyu Sun, Haoyi Jiang, Liu Liu, Seungtae Nam, Gyeongjin Kang, Xinjie Wang, Wei Sui,
 572 Zhizhong Su, Wenyu Liu, Xinggang Wang, et al. Uni3r: Unified 3d reconstruction and seman-
 573 tic understanding via generalizable gaussian splatting from unposed multi-view images. *arXiv
 574 preprint arXiv:2508.03643*, 2025.

575 Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Ge-
 576 ometric 3d vision made easy. In *CVPR*, 2024a.

577 Xingrui Wang, Cuiling Lan, Hanxin Zhu, Zhibo Chen, and Yan Lu. Gsem3plat: Generalizable
 578 semantic 3d gaussian splatting from uncalibrated image pairs. *arXiv preprint arXiv:2412.16932*,
 579 2024b.

580 Yanmin Wu, Jiarui Meng, Haijie Li, Chenming Wu, Yahao Shi, Xinhua Cheng, Chen Zhao,
 581 Haocheng Feng, Errui Ding, Jingdong Wang, et al. Opengaussian: Towards point-level 3d
 582 gaussian-based open vocabulary understanding. *arXiv preprint arXiv:2406.02058*, 2024.

583 Shaoan Xie, Zhifei Zhang, Zhe Lin, Tobias Hinz, and Kun Zhang. Smartbrush: Text and shape
 584 guided object inpainting with diffusion model. In *Proceedings of the IEEE/CVF conference on
 585 computer vision and pattern recognition*, pp. 22428–22437, 2023.

594 Haolin Xiong, Sairisheek Muttukuru, Rishi Upadhyay, Pradyumna Chari, and Achuta Kadambi.
 595 Sparsegs: Real-time 360 $\{\backslash\deg\}$ sparse view synthesis using gaussian splatting. *arXiv preprint*
 596 *arXiv:2312.00206*, 2023.

597

598 Jiarui Xu, Sifei Liu, Arash Vahdat, Wonmin Byeon, Xiaolong Wang, and Shalini De Mello. Open-
 599 vocabulary panoptic segmentation with text-to-image diffusion models. In *CVPR*, 2023.

600 Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth
 601 anything: Unleashing the power of large-scale unlabeled data. In *CVPR*, 2024a.

602

603 Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang
 604 Zhao. Depth anything v2. *NeurIPS*, 2024b.

605

606 Mingqiao Ye, Martin Danelljan, Fisher Yu, and Lei Ke. Gaussian grouping: Segment and edit
 607 anything in 3d scenes. In *ECCV*, 2024.

608 Chandan Yeshwanth, Yueh-Cheng Liu, Matthias Nießner, and Angela Dai. Scannet++: A high-
 609 fidelity dataset of 3d indoor scenes. In *ICCV*, 2023.

610 Qihang Yu, Ju He, Xueqing Deng, Xiaohui Shen, and Liang-Chieh Chen. Convolutions die hard:
 611 Open-vocabulary segmentation with single frozen convolutional clip. *NeurIPS*, 2023.

612

613 Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-splatting: Alias-
 614 free 3d gaussian splatting. In *CVPR*, 2024.

615

616 Quan-Sheng Zeng, Yunheng Li, Daquan Zhou, Guanbin Li, Qibin Hou, and Ming-Ming Cheng.
 617 Maskclip++: A mask-based clip fine-tuning framework for open-vocabulary image segmentation.
arXiv preprint arXiv:2412.11464, 2024.

618

619 Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image
 620 diffusion models. In *ICCV*, 2023.

621

622 Shijie Zhou, Haoran Chang, Sicheng Jiang, Zhiwen Fan, Zehao Zhu, Dejia Xu, Pradyumna Chari,
 623 Suya You, Zhangyang Wang, and Achuta Kadambi. Feature 3dgs: Supercharging 3d gaussian
 624 splatting to enable distilled feature fields. In *Proceedings of the IEEE/CVF Conference on Com-
 puter Vision and Pattern Recognition*, pp. 21676–21685, 2024.

625

626 Zehao Zhu, Zhiwen Fan, Yifan Jiang, and Zhangyang Wang. Fsgs: Real-time few-shot view synthe-
 627 sis using gaussian splatting. In *ECCV*, 2024.

628

629 Xingxing Zuo, Pouya Samangouei, Yunwen Zhou, Yan Di, and Mingyang Li. Fmgs: Foundation
 630 model embedded 3d gaussian splatting for holistic 3d scene understanding. *International Journal*
of Computer Vision, 133(2):611–627, 2025.

631

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646

647

648 **A ADDITIONAL EXPERIMENTAL RESULTS**
649650 **A.1 VIDEO RESULTS**
651652 To provide a more comprehensive and intuitive visualization of our method, we include video re-
653 sults in the supplementary ZIP file. Specifically, we present visualizations across five different
654 scenes. For each scene, we showcase the rendering results of both Splatt3R (Smart et al., 2024) and
655 OGGSplat under continuous camera views. Additionally, we provide the corresponding relevance
656 score heatmaps under a specific open-vocabulary query, enabling a direct comparison of semantic
657 understanding across the two methods. As clearly demonstrated, our model effectively extrapolates
658 to unseen regions while maintaining both high visual fidelity and semantic plausibility.
659660 **A.2 MORE ANALYSIS ON VISUAL FIDELITY**
661662 In Table 1, we use FID to evaluate the consistency and diversity between the distributions of rendered
663 novel views and the original dataset. Here, we supplement this evaluation with additional metrics,
664 including CLIP score, DINO score, and LLM score, providing a more comprehensive assessment of
665 the rendered image quality.666 **CLIP Score.** We use the CLIP (Radford et al., 2021) score to measure the semantic similarity
667 between the novel (outpainted) regions and the original contextual regions. For each sample in
668 the ScanNet++ validation set, we randomly generate masks and use MaskCLIP (Ding et al., 2022)
669 to extract features from both the masked and unmasked regions. The cosine similarity between
670 these features is computed and averaged across all samples to obtain the oracle score, serving as an
671 upper bound for contextual alignment. We further compute CLIP scores for our method as well as
672 LangSplat and Splatt3R, comparing the novel regions against the original context regions.
673674 **DINO Score.** While the CLIP score emphasizes semantic coherence at the boundary between ex-
675 panded and original content, it is less sensitive to style differences. To address this, we evaluate edge
676 regions using the DINO score (Caron et al., 2021). DINO, as a well-established visual representa-
677 tion model, captures both semantic and texture features, including style, providing a more holistic
678 assessment of visual fidelity. The measurement procedure for all models and the oracle is kept fully
679 consistent with that of the CLIP score, ensuring a fair and comparable evaluation across metrics.
680681 **LLM Score.** To provide a more global assessment, we employ the well-trained vision-language
682 model Qwen2.5-VL-7B-Instruct (Bai et al., 2025) as an expert evaluator. The model is prompted to
683 score inpainted images based on authenticity, style continuity, and structural integrity. Specifically,
684 we use the following prompt: “Please assess the image’s authenticity, style continuity, and structural
685 integrity, and assign a score between 1 and 100. Output your result strictly in the format: The score is
686 XXX.” This approach allows for a holistic evaluation that complements the CLIP and DINO scores,
687 capturing both visual fidelity and stylistic coherence.
688689 Table 3 presents the results for all three metrics. We observe that our method consistently approaches
690 the oracle scores across all measures, demonstrating superior performance in semantic coherence,
691 style continuity, and structural integrity.
692693 Table 3: Comparison of CLIP score, DINO score, and LLM score across different models.
694695

Model	CLIP Score	DINO Score	LLM Score
Oracle	53.7	84.9	76.5
LangSplat	47.9	82.3	66.9
Splat3R	44.6	74.7	71.8
OGGSplat	50.8	83.4	73.6

697 **A.3 ABLATION ON SEPARATE DIFFUSION UNET**
698699 To enable the generation of both spatially consistent RGB images and semantic content, we train two
700 separate diffusion models: Diff_{rgb} and Diff_{sem} , and enforce spatial consistency between them using
701 a ControlNet (Zhang et al., 2023) based approach. A simpler alternative would be to employ a single
702

702
703
704 Table 4: Ablation on semantic type for inpainting.
705
706
707

Methods	mIoU _L	wall	ceiling	floor	table	door	cabinet	chair	shelf	box	bed
Implicit Condition	17.0	44.7	0.1	58.8	13.5	3.5	2.7	22.2	5.8	3.6	14.8
Explicit Condition	17.6	45.6	0.1	58.3	13.3	5.4	3.7	21.4	7.4	3.1	18.0

708
709 Table 5: Ablation study of the depth alignment
710 module evaluated by Chamfer Distance (CD).
711

Alignment Strategy	CD1↓	CD2↓
No Depth Alignment	0.28	0.48
Bounding Box Alignment	0.21	0.32
EMS Alignment (Ours)	0.13	0.17

712 Table 6: Comparison of time consumption.
713
714
715

Stage	LangSplat	OGGSplat
Gaussian Init.	0.6s	0.6s
Inpainting	-	5.8s
Gaussian Opt.	157s	27.5s
Total	157.6s	33.9s

716
717 shared diffusion UNet based on an image diffusion model (Rombach et al., 2022), modified to allow
718 additional semantic inputs and outputs by adjusting the input and output convolutional channels.
719 However, our experiments show that this approach fails to produce meaningful RGB and semantic
720 outputs. As illustrated in Fig. 7, using a hybrid (shared) diffusion UNet leads to severe distortions in
721 both RGB images and semantic content. We think that this failure is due to the significant differences
722 between the latent spaces of the RGB image VAE and the semantic VAE, which makes it difficult for
723 a single UNet to learn consistent mappings in both domains. These results highlight the effectiveness
724 and necessity of our separate Diff_{sem} model and the corresponding control module design.
725

726 A.4 ABLATION ON SEMANTIC TYPE FOR INPAINTING
727

728 We utilize explicit text representations as conditions rather than edge features since our inpainting
729 model is fine-tuned from a pre-trained image diffusion model, which is originally conditioned on
730 text prompts via the CLIP text encoder. Maintaining the same conditioning modality enables us
731 to better leverage the prior knowledge acquired during pretraining, leading to more controllable
732 and semantically accurate inpainting. In contrast, using implicit edge features as conditioning can
733 result in less effective control due to potential embedding space misalignment, particularly given the
734 relatively small scale of our fine-tuning dataset compared to the pretraining corpus.
735

736 Moreover, explicit text conditioning provides an additional advantage over implicit feature-based
737 conditioning: user-specified descriptions can be directly appended to the original text prompt, al-
738 lowing for more flexible and intuitive control over the generated content.
739

740 To validate this design choice, we perform a quantitative ablation study. In the first row of Table 4,
741 we replace the explicit text prompt with implicit APE features as the inpainting condition. The
742 results demonstrate that explicit text prompts consistently achieve higher segmentation mIoU, with
743 notable improvements observed for categories such as *cabinet*, *shelf*, and *bed*.
744

745 A.5 ABLATION ON DEPTH ALIGNMENT MODULE
746

747 The design of the depth alignment module is a crucial component for minimizing the distance be-
748 tween newly added Gaussian points and the original Gaussian points. To demonstrate its effec-
749 tiveness, we conduct an ablation study. Specifically, we employ Chamfer Distance L1 (CD1) and
750 Chamfer Distance L2 (CD2) to quantify the spatial discrepancy between the reprojected Gaussian
751 points and the original points within overlapping regions across multiple views. The experimental
752 results are presented in Table 5, where the rows from top to bottom correspond to: no depth align-
753 ment, bounding box diagonal ratio-based alignment, and our proposed EMS ratio-based alignment.
754

755 The results highlight the importance of the depth alignment module. Introducing the alignment
756 method significantly reduces CD values compared to the case without alignment. Notably, our
757 EMS-based alignment consistently outperforms the bounding box-based strategy. We attribute this
758 improvement to the robustness of EMS alignment against outliers, which can otherwise negatively
759 impact the accuracy of alignment estimation.
760

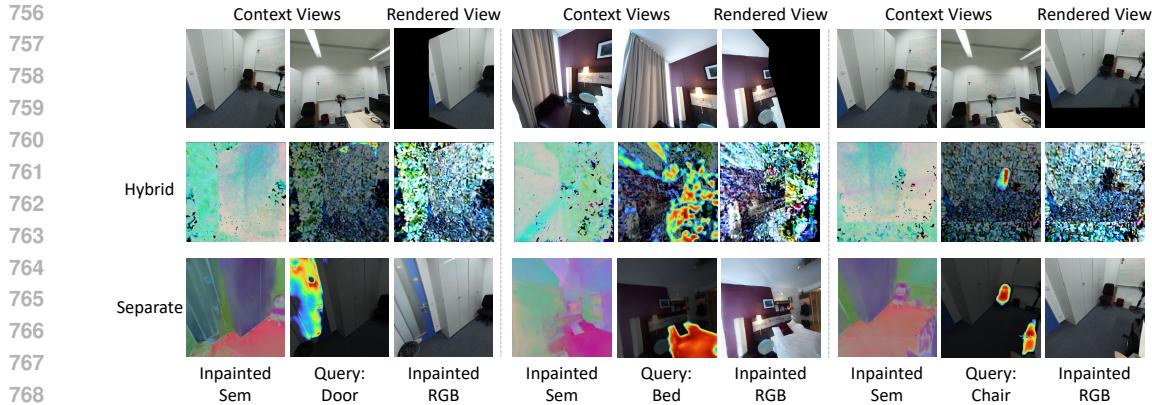


Figure 7: Qualitative comparison between hybrid (shared-weight) and separate diffusion UNet architectures. Row 1 shows the context images along with the incomplete renderings from novel views. Row 2 presents the results by using a hybrid UNet that jointly predicts RGB image and semantic content using shared weights. Row 3 shows the results from our proposed architecture with two separate UNets: one for RGB image synthesis and the other for semantic prediction.

A.6 DISCUSSION ON TIME CONSUMPTION

We provide a detailed comparison of the runtime between our method and LangSplat for a single user-defined novel view in Table 6. Although our pipeline introduces an additional inpainting stage, which is absent in LangSplat, this stage significantly accelerates the subsequent Gaussian optimization process. By leveraging the inpainted content together with UniDepth, we achieve more effective Gaussian initialization and optimization. Specifically, while the inpainting step incurs an extra cost of 5.8s, the overall Gaussian optimization time is drastically reduced (27.5s vs. 157s in LangSplat), leading to a substantial reduction in the total runtime (33.9s vs. 157.6s).

Furthermore, our pipeline naturally supports an arbitrary number of user-defined novel views, with the computational cost scaling linearly with the number of views, offering controllability and flexibility. After optimization with the inpainted view(s), our method further enables real-time rendering for additional novel views, consistent with the runtime characteristics of existing Gaussian Splatting-based approaches.

B IMPLEMENTATION DETAILS

B.1 SCANNET++ DATASET

ScanNet++ dataset (Yeshwanth et al., 2023) provides high-quality 3D geometry along with high-resolution RGB images of various indoor environments. Following the protocol introduced by Splatt3R, originally designed for 3D reconstruction, we adopt the standard training split comprising 230 scenes and the validation split containing 50 scenes. Following (Smart et al., 2024; Wang et al., 2024b), we also discard frames missing reliable depth information. All selected frames are uniformly cropped and resized to a spatial resolution of 512×512 .

B.2 TRAINING SETTINGS

To provide a clearer overview of the experimental configurations used at different training stages, we summarize the details in Table 8. The table includes the settings for all key components that need to be trained in our method, namely the generalizable open-vocabulary Gaussian initialization module, RGB UNet, semantic VAE (Kingma et al., 2013), semantic UNet, ControlNet (Zhang et al., 2023), and the open-vocabulary Gaussian growing process.

Generalizable Open-Vocabulary Gaussian Initialization. We adopt the pretrained Splatt3R model and freeze its backbone, which is responsible for predicting the basic Gaussian attributes. We then train only the newly added semantic head, denoted as H_{sem} . During training, we use two context images as input and supervise the model by rendering three target views from the training

Table 8: Experiment settings for different training stages.

Config	Gaussian Init.	RGB-Semantic Consistent Inpaintor				Gaussian Growing
		RGB UNet	Sem. VAE	Sem. UNet	ControlNet	
optimizer	Adam	AdamW8bit	AdamW	AdamW8bit	AdamW8bit	Adam
learning rate	1e-5	1e-5	6e-6	1e-5	1e-5	hybrid (Table 7)
weight decay	5e-2	1e-2	1e-2	1e-2	1e-2	0
scheduler	multi-step	constant	cosine	constant	constant	exponential
batch size	12	4	2	4	4	4
accumulation steps	1	2	4	2	2	1
training iterations	500,000	50,000	45,000	20,000	10,000	600
GPU device	8 RTX 3090	8 RTX 3090	8 RTX 3090	8 RTX 3090	8 RTX 3090	1 RTX 3090
image size	512×512	512×512	512×512	512×512	512×512	512×512

split. Following the setup in Splatt3R (Smart et al., 2024), the context images are selected such that at least 30% of the pixels in the second image have direct correspondences in the first image. Similarly, target images are chosen such that at least 30% of their content is visible in at least one of the context images.

RGB-Semantic Consistent Inpaintor. For RGB image inpainting model Diff_{sem} , we fine-tune a stable diffusion inpainting model (Rombach et al., 2022) to better align the generated appearance with realistic indoor scenes. In addition to standard RGB inpainting, we propose a novel diffusion-based feature inpainting model, denoted as Diff_{sem} , which consists of both a Variational Autoencoder (Kingma et al., 2013) (VAE) and a UNet architecture. This model enables semantic-aware inpainting in the feature space while maintaining consistency with the RGB domain. To ensure spatial consistency between the RGB and semantic contents, we train an auxiliary RGB control module inspired by ControlNet (Zhang et al., 2023) that guides the inpainting process in the feature space.

Open-Vocabulary Gaussian Growing. We set the horizontal and vertical outpainting angles to lie within the ranges of $[-60^\circ, 60^\circ]$ and $[-20^\circ, 20^\circ]$, respectively. To simplify this stage, we decouple the horizontal and vertical rotations: when the horizontal angle is non-zero, the vertical angle is set to zero, and vice versa. For each optimization round, to improve efficiency, we use two inpainted images and their corresponding semantic maps under symmetrical camera poses to provide the supervision signal. Moreover, the selected camera view pairs are arranged to exhibit progressively increasing angular differences, thereby enabling a gradual and progressive Gaussian growing process. Specifically, denoting the camera rotation angles in the horizontal and vertical directions as (θ_h, θ_v) , the sampled camera angles are selected in the following order: $(0^\circ, 0^\circ)$, $(0^\circ, \pm 20^\circ)$, $(\pm 30^\circ, 0^\circ)$, and $(\pm 60^\circ, 0^\circ)$. It is worth noting that during actual optimization, camera poses can be arbitrary. This sampling strategy is adopted purely to facilitate a simpler, more consistent, and computationally efficient optimization process. We conduct a total of four optimization rounds. In the first round, we perform inpainting without changing the camera poses, i.e., using poses of the original context views. This step focuses on refining low-confidence regions through inpainting to enhance rendering quality under the original views. In subsequent rounds, we fix the batch size to 4 and include supervision signals from the originally inpainted context views, previously inpainted views, and newly generated inpainted views. For the optimization of Gaussian parameters, we adopt parameter-specific learning rates following the setting proposed in (Qin et al., 2024). The detailed learning rates for each type of parameter are summarized in Table 7. Empirically, we observe that each optimization round converges efficiently within 600 training iterations.

B.3 OVG0 BENCHMARK

For evaluation on our proposed OVG0 benchmark, we uniformly sample 16 novel camera poses around the context image I_1 , covering a horizontal angular range of $[-60^\circ, 60^\circ]$ and a vertical angular range of $[-20^\circ, 20^\circ]$. To simplify the evaluation setup, we decouple horizontal and vertical

Table 7: Learning rates for different Gaussian parameters.

Parameter	Learning Rate
point position μ	1e-2
rotation quaternion q	1e-3
scale vector s	5e-3
opacity scalar α	5e-2
spherical harmonics \mathbf{S}	2.5e-2
semantic feature f	2.5e-3

864 rotations, following the same strategy described in Appendix B.2. The IoU score for every query is
 865 computed by averaging over a total of 50×16 images. If the union of predicted and ground-truth
 866 regions in an image is empty, that image is excluded from the IoU computation. To ensure robust-
 867 ness, we repeat the inpainting, growing, and evaluation process five times with the same settings and
 868 report the mean IoU as the final benchmark result.
 869

870 B.4 RGB-TO-SEMANTIC CONTROLNET MODULE

872 To ensure spatial alignment between the in-
 873 painted RGB image and its corresponding se-
 874 mantic map, we adopt a control mechanism
 875 inspired by ControlNet (Zhang et al., 2023),
 876 where the RGB image serves as guidance
 877 for the generation of the semantic map. An
 878 overview of the ControlNet architecture is il-
 879 lustrated in Fig. 8. Specifically, our control
 880 module comprises the encoder and bottleneck
 881 components of the stable diffusion UNet ar-
 882 chitecture, with their weights initialized from
 883 the corresponding layers of a pretrained stable
 884 diffusion UNet. Conditional signals are then
 885 injected into the bottleneck and decoder parts
 886 via zero convolutions and element-wise addi-
 887 tion. To accelerate training and enhance the
 888 effectiveness of control learning, we initialize
 889 the control module with pretrained parameters
 890 from a ControlNet model (Zhang et al., 2023)
 891 conditioned on image segmentation. This ini-
 892 tialization strategy provides a strong prior for
 893 spatially consistent generation and significantly
 894 improves both training efficiency and overall
 895 performance. Details of the training settings for
 896 this module are provided in Table 8.
 897

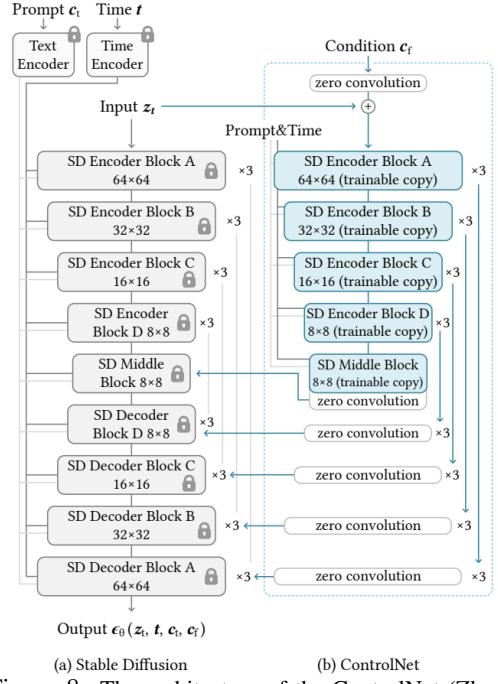
898 B.5 RELEVANCE SCORE FOR EVALUATION

900 During open-vocabulary querying, we select regions with a relevance score greater than 0.5 as the fi-
 901 nal predicted category mask. The computation of the relevance score is inspired by prior works (Kerr
 902 et al., 2023; Qin et al., 2024; Shi et al., 2024), and is defined as follows for each query:
 903

$$904 \text{Relevance} = \min_i \frac{\exp(g_{\text{img}} \cdot g_{\text{qry}})}{\exp(g_{\text{img}} \cdot g_{\text{qry}}) + \exp(g_{\text{img}} \cdot g_{\text{canon}}^i)}, \quad (10)$$

906 where g_{img} denotes the image semantic feature, g_{qry} is the query APE embedding, and g_{canon}^i repre-
 907 sent the APE embedding of a predefined canonical phrase such as “object”, “things”, “stuff”, or
 908 “texture”.

909 In contrast to the mentioned prior works, which typically focus on a limited set of categories in a
 910 single scene and require the set of possible scene categories to be known in advance, we adopt a
 911 more general strategy. These prior methods often normalize the relevance score and select masks
 912 based on a threshold over the normalized values. However, this approach may incorrectly force the
 913 prediction of masks even for categories absent in the scene. To address this limitation and enhance
 914 generalizability, we directly apply a fixed threshold of 0.5 to the raw (unnormalized) relevance scores
 915 and select pixels with scores exceeding this threshold as the final predicted mask. This ensures
 916 that only queries with truly high relevance scores produce predictions, avoiding false positives in
 917 irrelevant categories. As a result, we are able to compute per-category prediction masks from a
 918 predefined query set without requiring manual query specification for each individual scene.



919 Figure 8: The architecture of the ControlNet (Zhang
 920 et al., 2023).
 921

(a) Stable Diffusion

(b) ControlNet

922

923

924

925

926

927

928

929

930

931

932

933

934

935

936

937

938

939

940

941

942

943

944

945

946

947

948

949

950

951

952

953

954

955

956

957

958

959

960

961

962

963

964

965

966

967

968

969

970

971

972

973

974

975

976

977

978

979

980

981

982

983

984

985

986

987

988

989

990

991

992

993

994

995

996

997

998

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

999

918 **C ACKNOWLEDGMENT OF LLM USAGE**
919920 We acknowledge the use of Large Language Models (LLMs) in the process of refining certain sec-
921 tions of this manuscript. The LLMs were employed to assist in polishing and improving the clarity
922 of the text, but all content and intellectual contributions remain the result of the authors' research.
923 We emphasize that the LLMs were utilized only as a tool for language enhancement and did not
924 influence the research findings or the overall scientific content presented herein.
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971