

# RT-GS2: Real-Time Generalizable Semantic Segmentation for 3D Gaussian Representations of Radiance Fields

Mihnea-Bogdan Jurca\*<sup>1, 2</sup>  
mihnea-bogdan.jurca@vub.be

Remco Royen\*<sup>1</sup>  
remco.royen@vub.be

Ion Giosan<sup>2</sup>  
ion.giosan@cs.utcluj.ro

Adrian Munteanu<sup>1</sup>  
adrian.munteanu@vub.be

<sup>1</sup> Department ETRO  
Vrije Universiteit Brussel  
Brussels, Belgium

<sup>2</sup> Computer Science Department  
Technical University of Cluj-Napoca,  
Cluj-Napoca, Romania

---

## Abstract

Gaussian Splatting has revolutionized the world of novel view synthesis by achieving high rendering performance in real-time. Recently, studies have focused on enriching these 3D representations with semantic information for downstream tasks. In this paper, we introduce RT-GS2, the first generalizable semantic segmentation method employing Gaussian Splatting. While existing Gaussian Splatting-based approaches rely on scene-specific training, RT-GS2 demonstrates the ability to generalize to unseen scenes. Our method adopts a new approach by first extracting view-independent 3D Gaussian features in a self-supervised manner, followed by a novel View-Dependent / View-Independent (VDVI) feature fusion to enhance semantic consistency over different views. Extensive experimentation on three different datasets showcases RT-GS2's superiority over the state-of-the-art methods in semantic segmentation quality, exemplified by a 8.01% increase in mIoU on the Replica dataset. Moreover, our method achieves real-time performance of 27.03 FPS, marking an astonishing 901 times speedup compared to existing approaches. This work represents a significant advancement in the field by introducing, to the best of our knowledge, the first real-time generalizable semantic segmentation method for 3D Gaussian representations of radiance fields. The project page and implementation can be found at <https://mbjurca.github.io/rt-gs2/>

## 1 Introduction

Scene Understanding is a fundamental area of research, essential for facilitating seamless interactions between digital devices and the three-dimensional environment. While 2D representations such as RGB images are traditionally being employed, they fail to fully capture the three-dimensional properties of the scene, resulting in view-dependent outcomes [8]. On

---

\* Both authors contributed equally to the paper.

© 2024. The copyright of this document resides with its authors.

It may be distributed unchanged freely in print or electronic forms.

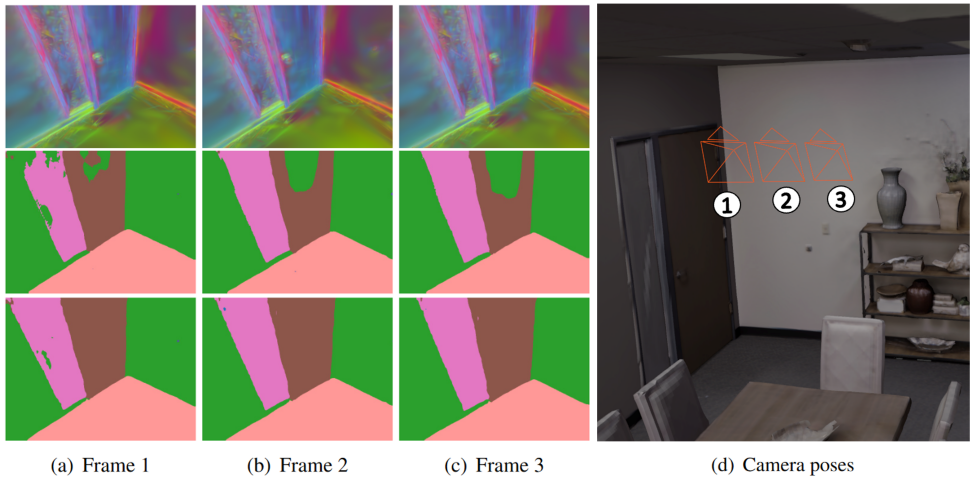


Figure 1: Visualization of the enhanced view-consistency throughout subsequent frames. (Top) Visualization of our view-independent 3D features using PCA, (middle) semantic segmentation without the usage of view-independent 3D features, and (bottom) proposed semantic segmentation when using our view-independent 3D features.

the other hand, 3D representations such as point clouds and polygonal meshes offer the capability to digitally represent 3D scenes, but typically require high-resolution 3D point data to capture fine details [23, 29], often obtained through expensive devices like LiDARs.

In recent years, Neural Radiance Fields (NeRFs) have emerged as a groundbreaking approach for novel view synthesis [10, 12, 18]. By training multi-layered perceptrons (MLPs) on extensive sets of images captured from various viewpoints, NeRFs learn an implicit 3D representation of a scene. Beyond novel view synthesis, this learned 3D representation proves beneficial for downstream tasks as it enhances accuracy and enables view-consistent scene understanding [50]. Moreover, by formulating the loss of rendering and downstream tasks in the 2D domain, the need for time-intensive 3D annotations is eliminated. However, despite their performance, NeRFs exhibit a significant trade-off between visual quality and inference speed [18, 58]. Recently, a novel approach, dubbed Gaussian Splatting, was introduced for novel view synthesis [13]. This method learns 3D Gaussian distributions in space, encompassing not only location and scale but also opacity and color spherical harmonics per 3D Gaussian. By simply splatting these Gaussians during inference, real-time high-quality synthesis of novel views is achieved.

Significant research efforts have been directed towards extending novel view synthesis methods to accommodate downstream tasks. A popular approach is to equip NeRFs with semantic segmentation capabilities [2, 40]. While these methods yield impressive segmentation results, they are trained on a scene-per-scene basis and are thus unable to generalize to unseen scenes. Consequently, methods [6, 12, 14, 15, 16, 28] capable of performing semantic segmentation on unseen 3D NeRF representations were designed. As these methods are built upon NeRFs, real-time capabilities are lacking. Although recent studies have begun exploring downstream applications using 3D Gaussian Splatting [11, 21, 56, 42], to the best of our knowledge, no existing method in the current literature addresses generalizable semantic segmentation for 3D Gaussian Splatting.

In this paper, we present a novel method, dubbed RT-GS2, designed for generalizable semantic segmentation based on 3D Gaussian Splatting. RT-GS2 consists of three distinct stages. Firstly, a self-supervised 3D Gaussian feature extractor learns view-independent 3D features from the complete Gaussian 3D representation. After splatting the features and 3D Gaussians for a specific viewpoint, feature fusion is performed, enhancing view-consistent semantic segmentation, as illustrated in the examples of Figure 1. Results for both semantic segmentation and depth prediction showcase the robustness of the obtained geometric 3D features. RT-GS2 not only surpasses the state of the art in generalizable semantic segmentation but also achieves an impressive speedup, being 901 times faster than existing methods. By doing so, it is the first to meet real-time constraints, marking a significant advancement in the field. In summary, our main contributions include:

- The introduction of a novel method for generalizable semantic segmentation, the first to employ 3D Gaussian splatting. RT-GS2 presents a novel approach by first obtaining view-independent 3D features in a self-supervised manner, followed by a so-called View-Dependent / View-Independent (VDVI) feature fusion to obtain enhanced view-consistency for semantic segmentation.
- A self-supervised feature extractor for 3D Gaussians, enabling the extraction of generic and consistent view-independent 3D features, which prove to be robust for multiple tasks such as semantic segmentation and depth predictions.
- Extensive experimentation on different datasets demonstrating that RT-GS2 not only strongly outperforms the state of the art in segmentation quality but also achieves real-time inference, achieving a notable 901 times speedup compared to existing generalizable semantic segmentation methods.

## 2 Related work

**Novel View Synthesis.** The growing interest in implicit neural representations has greatly advanced the frontier of novel view synthesis in recent years. The seminal NeRF paper [14] has spurred iterative enhancements focusing on faster rendering [22, 31, 38], accelerated training processes [3, 18, 26], and the capability to handle unbounded scenes [1]. Recently, Gaussian Splatting [13] has demonstrated superiority over NeRF-based methods in both rendering quality and inference time. Acknowledging the remarkable potential of 3D Gaussians, we build upon this paradigm.

**Semantic Segmentation.** Traditional semantic segmentation techniques operate on a single modality. While 2D-based techniques [6, 32] benefit from the employment of cost-efficient RGB-camera’s, they have difficulties to fully capture the underlying 3D geometry. 3D semantic segmentation methods [19, 21, 54, 34, 39] on the other hand, achieve high performance but require dense 3D models captured by expensive 3D scanners. The advent of NeRFs allowed to achieve view-consistent results on 2D images for a specific 3D scene by equipping NeRFs with semantic capabilities [7, 40]. In order to achieve real-time constraints, [9, 24, 36] learn semantic features for each 3D Gaussian.

**Generalizable Semantic Radiance Fields.** While the above mentioned NeRF- and Gaussian Splatting-based semantic segmentation papers achieve high performance, their scene-specific training leads to significant overfitting for individual scenes. To address this limitation, generalizable semantic segmentation methods, capable of segmenting unseen scenes,

were proposed for NeRFs [4, 6, 24, 15]. For instance, S-Ray [15] introduces a Cross-Reprojection Attention module for efficient exploitation of semantic information along rays, while GP-NeRF [24] utilizes transformers to aggregate semantic embedding fields. GNeSF [4] employs a soft voting mechanism to aggregate 2D semantic information from different views, and GSNeRF [6] integrates image semantics into the synthesis process for mutual enhancement. However, due to their dependence and build-up on NeRFs, these methods lack real-time execution. To our knowledge, there is no existing method in scientific literature that uses 3D Gaussian Splatting for real-time generalizable semantic segmentation. This is addressed next.

### 3 3D Scene Representations Using 3D Gaussian Splats

In order to render novel views of complex scenes, NeRF-based methods define a continuous volumetric radiance field [14] for each scene, parameterized by  $\mathbf{F}_\theta : \mathbb{R}^5 \rightarrow \mathbb{R}^4$ , where  $\mathbf{F}_\theta$  is implemented by a MLP with learnable parameters  $\theta$ . A differentiable forward mapping function is employed to retrieve discrete 2D views from the radiance field. A per-pixel loss between the synthesized rendering and the ground truth image allows the optimization of parameters  $\theta$  for a specific scene.

Gaussian Splatting [13] takes a different approach and optimizes the training and rendering process while preserving the desirable properties of a radiance field. This is achieved by removing the implicit 3D representation and instead modeling each scene  $k$  as a combination of 3D Gaussian functions  $\mathbf{G}^k = \{\mathbf{g}_1^k, \mathbf{g}_2^k, \dots, \mathbf{g}_{N'}^k\}$ . Each Gaussian  $\mathbf{g}_i^k$  is defined by its world coordinates  $\mathbf{x}_i^k \in \mathbb{R}^3$  and a covariance matrix  $\Sigma_i^k \in \mathbb{R}^{3 \times 3}$ . Additionally, they are further enriched with an opacity  $\alpha_i^k \in \mathbb{R}$ , and a color  $\mathbf{c}_i^k$ , represented by spherical harmonic coefficients with three degrees. Thus, mathematically, each 3D Gaussian can be represented as  $\mathbf{g}_i^k = \{\mathbf{x}_i^k, \Sigma_i^k, \alpha_i^k, \mathbf{c}_i^k\}$ . To render from  $\mathbf{G}^k$  the 2D image  $\hat{\mathbf{I}}_j^k$  for view  $j$  with pose  $\mathbf{p}_j$ , the 3D Gaussians are splatted into a RGB image employing alpha blending, expressed mathematically as follows:

$$\hat{\mathbf{I}}_j^k(u, v) = \sum_{i \in N'} \mathbf{c}_i^k \alpha_i^k \prod_{m=1}^{i-1} (1 - \alpha_m^k), \quad (1)$$

where  $(u, v) \in ([1, H], [1, W])$  represent the pixel coordinates in the 2D image after splatting  $N'$  Gaussians. When repeated for all pixels, the rendered image  $\hat{\mathbf{I}}_j^k$  is obtained. The Gaussian parameters are optimized by employing the following loss function:

$$\mathcal{L} = (1 - \lambda)\mathcal{L}_1 + \lambda \mathcal{L}_{D-SSIM}, \quad (2)$$

with  $\mathcal{L}_1$  a per-pixel  $L_1$ -loss and  $\mathcal{L}_{D-SSIM}$  the structural dissimilarity metric between the rendered  $\hat{\mathbf{I}}_j^k$  and the ground truth image  $\mathbf{I}_j^k$ . The hyperparameter  $\lambda$  is typically set to 0.2. Similar to Equation 1, a feature rendering function is defined, as introduced by [66], allowing the splatting of Gaussian features  $\mathbf{f}_i^k \in \mathbb{R}^D$ , where  $\mathbf{f}_i^k$  are the features attached to  $\mathbf{g}_i^k$ , to the feature image  $\mathbf{Z}_j^k \in \mathbb{R}^{H \times W \times D}$ . This feature rendering function can be expressed mathematically by:

$$\mathbf{Z}_j^k(u, v) = \sum_{i \in N'} \mathbf{f}_i^k \alpha_i^k \prod_{m=1}^{i-1} (1 - \alpha_m^k). \quad (3)$$



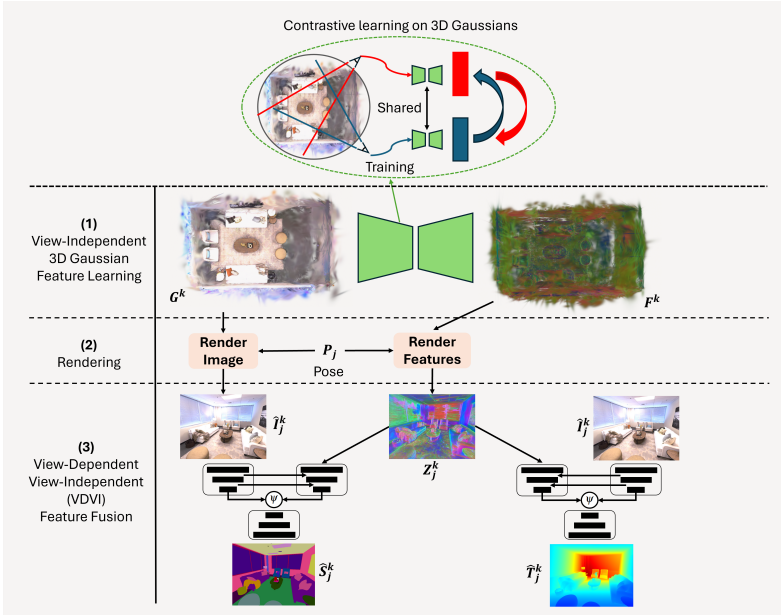


Figure 2: Overview of the proposed method.

## 4 Proposed method

### 4.1 Overview of the proposed method

The architecture of the proposed method is illustrated in Figure 2 and consists out of three main stages: (i) a new self-supervised view-independent 3D Gaussian feature extractor, (ii) the rendering of the 3D information encapsulated in the enhanced 3D Gaussians to a specific view, and (iii) a novel View-Dependent / View-Independent (VDVI) feature fusion. The first stage, described in Section 4.2, transforms a 3D Gaussian representation  $\mathbf{G}^k$  of a scene  $k$ , into a set of features  $\mathbf{F}^k = \{\mathbf{f}_1^k, \mathbf{f}_2^k, \dots, \mathbf{f}_N^k\} \in \mathbb{R}^{N \times D}$ , where  $\mathbf{f}_i^k$  are the features learned for the corresponding 3D Gaussian  $\mathbf{g}_i^k$ , and  $N$  is the number of 3D Gaussians in  $\mathbf{G}^k$ . Important to note is that, since the proposed method operates on the entire 3D Gaussian representation as input, the features  $\mathbf{f}_i^k$  are view-independent. In the second stage,  $\mathbf{F}^k$  is rendered by using alpha blending, described in Equation 3, to a specific view  $j$ , defined by the given pose  $\mathbf{p}_j$ . The feature image obtained by the splatting of  $\mathbf{F}^k$  is denoted by  $\mathbf{Z}_j^k \in \mathbb{R}^{H \times W \times D}$ . In parallel, the novel view  $\hat{\mathbf{I}}_j^k$  is rendered. In the last stage, the novel VDVI feature fusion module, described in Section 4.3, extracts view-dependent features from  $\hat{\mathbf{I}}_j^k$  and fuses them at different scales with  $\mathbf{Z}_j^k$ . A joint decoder is employed to obtain the semantic predictions  $\hat{\mathbf{S}}_j^k$ . By training the model on different scenes  $k$ , RT-GS2 is able to generalize semantic segmentation to unseen scenes.

### 4.2 View-independent 3D Gaussian feature learning

In order to obtain view-independent 3D Gaussian features, suitable for generalization on unseen scenes, we propose to learn 3D features directly from the Gaussian 3D representation.

This is in stark contrast to existing methods [9, 66] which employ loss-terms in 2D to learn per-Gaussian features, leading to view-specific and most importantly, scene-specific Gaussian features. Our fundamentally different approach exploits the spatial distribution of the 3D Gaussians and inter-Gaussian relations to extract 3D features for unseen scenes. More specifically, we employ a point cloud autoencoder as backbone as it allows to process unstructured 3D points and retrieve global and local features. In this autoencoder, each 3D Gaussian is represented by its location while the additional properties are encoded in the channel dimension. Important to note, is that the 3D Gaussian feature extractor operates on the complete Gaussian 3D representation during inference. This ensures the retrieval of view-independent 3D Gaussian features, exploiting information from the entire scene, beyond a specific view. From Figure 1 can be seen that the obtained 3D features are view-consistent and this, by consequence, improves view-consistency of the final semantic segmentation  $\hat{\mathbf{S}}_j^k$ . An additional advantage of processing the entire scene is the possibility to compute the 3D features only once for the whole scene, further reducing required online inference time while navigating through the scene.

Lastly, we have opted to train the view-independent 3D Gaussian feature extractor in a self-supervised manner, following the contrastive learning training procedure of [35]. This not only allows to perform the feature computation entirely in the 3D domain, but also ensures the extraction of robust features, suitable for multiple downstream tasks. In the supplementary, we ablate the robustness of the features by performing depth prediction. The employed self-supervised loss is described in Section 4.4.

### 4.3 View-Dependent / View-Independent (VDVI) feature fusion

The View-Dependent / View-Independent (VDVI) feature fusion module consists out of two parallel encoders:  $Enc_{VD}$  and  $Enc_{VI}$ . The former encodes the image  $\hat{\mathbf{I}}_j^k$  to view-dependent features. The latter takes the splatted view-independent 3D features,  $\mathbf{Z}_j^k$ , as input. At different scales, the view-independent features are fused with the encoded view-dependent features in encoder  $Enc_{VI}$ . At the lowest scale, the output of both encoders are fused by a fusion function  $\psi$ . Hereafter, the resulting features are employed by a decoder  $Dec$  to retrieve the semantic predictions  $\hat{\mathbf{S}}_j^k$ . This can be expressed mathematically as follows:

$$\hat{\mathbf{S}}_j^k = Dec(\psi(Enc_{VD}(\hat{\mathbf{I}}_j^k), Enc_{VI}(\mathbf{Z}_j^k, \hat{\mathbf{I}}_j^k))). \quad (4)$$

VDVI feature fusion improves pure segmentation performance and increases view-consistency of the results.

### 4.4 Loss

The self-supervised 3D Gaussian feature extractor and VDVI feature fusion are optimized with respect to  $L_{cl}$  [35] and  $\mathcal{L}_{sem}$ , respectively, each one defined below. Detailed descriptions on the losses and individual loss-terms can be found in the supplementary material.

$$L_{cl} = - \sum_{(m,n) \in P_k} \log \frac{\exp(f_m^k \cdot f_n^k / \tau)}{\sum_{(l,\cdot) \in P_k} \exp(f_m^k \cdot f_l^k / \tau)} \quad (5)$$

$$\mathcal{L}_{sem} = \mathcal{L}_{CrossEntropy} + \lambda_{CeCo} \mathcal{L}_{CeCo}. \quad (6)$$

	Method	Published	Semantic			Rendering			Time
			mIoU	mAcc	oAcc	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	FPS $\uparrow$
General.	MVSNerF* [10]	CVPR2022	30.21	39.75	69.35	23.68	84.37	28.08	-
	Neuray* [16]	CVPR2022	40.91	50.15	76.23	27.80	89.55	23.68	<0.03
	S-Ray [15]	CVPR2023	43.27	52.85	77.63	26.77	88.54	22.81	0.03
	GSNeRF [9]	CVPR2024	51.23	61.10	83.06	31.71	92.89	12.93	-
	Ours	-	<b>59.24</b>	<b>66.04</b>	<b>93.99</b>	<b>36.02</b>	<b>97.12</b>	<b>4.91</b>	<b>27.03</b>
Finetune	S-Ray [15]	CVPR2023	84.12	88.53	96.36	27.78	84.53	12.88	0.03
	Ours	-	<b>93.75</b>	<b>96.19</b>	<b>99.33</b>	<b>36.02</b>	<b>97.12</b>	<b>4.91</b>	<b>27.03</b>

Table 1: Comparison on REPLICa [25] of the proposed method against the state of the art for generalizable semantic segmentation and after finetuning on a specific scene. \* denotes the addition of a semantic head.

## 5 Experiments

### 5.1 Experimental setup

**Datasets and Metrics.** Experiments were conducted on three datasets: Replica [25], ScanNet [7], and ScanNet++ [57], representing synthetic and real-world indoor scenes. Experimental settings meticulously followed those of [9] for Replica and ScanNet. Details and splits are provided in the supplementary material. As evaluation metrics, we employed mean Intersection over Union (mIoU), mean Accuracy (mAcc), and overall Accuracy (oAcc) for segmentation, and Peak Signal-to-Noise Ratio (PSNR), Structural Similarity (SSIM), and Learned Perceptual Image Patch Similarity (LPIPS) for rendering quality. Inference speed was measured in Frames Per Second (FPS).

**Implementation details.** To train the RT-GS2 model, RGB data and their corresponding semantic masks are required for the training scenes. During testing, the model is able to operate solely on the Gaussian Splatting representation, discarding the need for expensive semantic masks for unseen scenes. Additionally to generalization performance, we also report the performance of our model after fine-tuning the generalized model for a fixed number of iterations using the scene’s semantic labels, allowing an increased performance. The view-independent 3D Gaussian feature extractor employs a PointTransformerV3 [64] with the output feature dimension  $D$  empirically chosen as 32. The input, with channel dimension 10 (xyz, base color, scale information and opacity), is subsampled during training with voxelization (size 0.07). The contrastive loss is computed among 4096 corresponding points from 2 different views. Asymformer [10] is selected as a real-time VDVI feature fusion backbone.  $\lambda_{cCo}$  is chosen 0.4 and LSR [27] is employed for generalization. All experiments were performed using a NVIDIA GeForce RTX 3090 GPU.

### 5.2 Comparison to state of the art

**Semantic segmentation on Replica.** Table 1 showcases our results on Replica [25] alongside a comparison with state-of-the-art methods. Our method significantly outperforms existing approaches in both segmentation quality and inference time. Specifically, for semantic generalization, RT-GS2 outperforms the state of the art across all evaluated metrics, achieving an impressive 8.01% increase in mIoU. Fine-tuning on specific scenes for 20k iterations further improves performance to 93.75% and 99.33% for mIoU and oAcc, respectively. While RT-GS2 does not present rendering generalization, we also present rendering per-

			Semantic			Rendering			Time
	Method	Published	mIoU	mAcc	oAcc	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	FPS $\uparrow$
General.	MVSNerF* [14]	CVPR2022	43.06	53.63	66.90	24.14	80.36	34.63	-
	Neuray* [16]	CVPR2022	46.09	53.79	66.39	25.24	84.39	31.33	<0.11
	S-Ray [15]	CVPR2023	47.69	54.47	64.90	25.13	84.18	30.44	0.11
	GSNeRF [6]	CVPR2024	52.21	60.14	74.71	<b>31.49</b>	<b>90.39</b>	<b>13.87</b>	-
	Ours	-	<b>53.27</b>	<b>62.43</b>	<b>81.20</b>	27.27	89.10	21.77	<b>27.03</b>
Finetune	S-Ray [15]	CVPR2023	91.6	92.2	97.3	27.31	-	-	0.11
	GSNeRF [6]	CVPR2024	93.2	96.8	98.2	<b>30.89</b>	-	-	-
	Ours	-	<b>96.94</b>	<b>98.87</b>	<b>99.01</b>	27.27	89.10	21.77	<b>27.03</b>

Table 2: Comparison on ScanNet [10] of the proposed method against the state of the art for generalizable semantic segmentation and after finetuning on a specific scene. \* denotes the addition of a semantic head.

		Semantic			Rendering			Time
	Method	mIoU	mAcc	oAcc	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	FPS $\uparrow$
Generalization	Ours	<b>66.14</b>	<b>77.32</b>	<b>83.79</b>	<b>26.39</b>	<b>87.31</b>	<b>17.08</b>	<b>27.03</b>
Finetuned	Ours	<b>91.85</b>	<b>95.73</b>	<b>96.24</b>	<b>26.39</b>	<b>87.31</b>	<b>17.08</b>	<b>27.03</b>

Table 3: Quantitative results on ScanNet++ [57] of the proposed method for generalizable semantic segmentation and after finetuning on a specific scene.

formance for completeness. The usage of Gaussian Splatting for rendering allows for an increase in rendering performance. Notably, our method achieves real-time novel view synthesis and segmentation at 27.03 FPS, a remarkable 901 times speedup compared to S-Ray. Qualitative results, depicted in Figure 3, demonstrate compelling segmentation performance. After finetuning, even the fine details are correctly segmented. RT-GS2 consistently outperforms S-Ray in both settings. Additional visual results and a video can be found in the supplementary material.

**Semantic segmentation on ScanNet.** Table 2 displays our results on the real-world dataset ScanNet [10]. While Gaussian Splatting does not surpass NeRF-based GSNeRF [6] in rendering quality on ScanNet, likely due to the high presence of motion blur and other sources of noise, our proposed method is still capable of consistently outperforming existing methods in segmentation quality and inference time, both for generalization and finetuning (5k iterations). The qualitative results in Figure 3 confirm these observations.

**Semantic segmentation on ScanNet++.** Additionally, we present results on ScanNet++ [57]. While this dataset was not publicly available during the publication of prior works [6, 12, 15, 16], ScanNet++ is well-suited for novel view synthesis and generalizable semantic segmentation, exhibiting high-quality images and accurately annotated classes. In Table 3, the proposed method demonstrates solid performance, enabling future comparisons. Qualitative results can be found in the supplementary material.

### 5.3 Ablation study

**Robustness of self-supervised features.** In this study, we investigate the robustness of the self-supervised view-independent 3D Gaussian features for a different task: depth prediction. Quantitative and qualitative results are available in the supplementary material.

**Evaluation of view-dependent and view-independent features.** Table 4 quantifies the influence of the synthesized image  $\hat{I}_j^k$  and rendered view-independent 3D Gaussian features

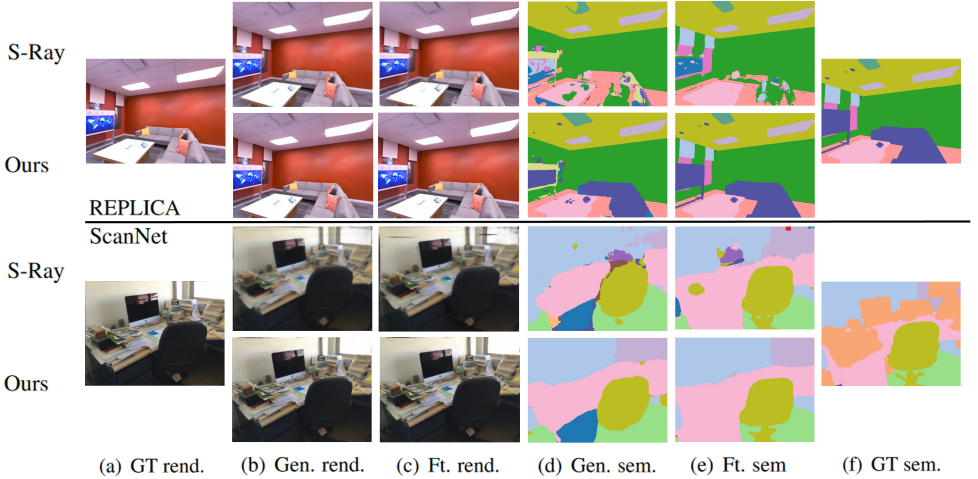


Figure 3: Qualitative results on Replica and ScanNet. The table presents generalizable (gen.) and finetuning (ft.) results on Replica (first two rows) and ScanNet (last two rows) for both rendering (rend.) and semantic segmentation (sem.). Comparisons between RT-GS2 and Semantic-Ray[15] are made.

	VDVI features		Loss		Semantic - top20			Semantic - all		
	$\hat{\mathbf{I}}_j^k$	$\mathbf{Z}_j^k$	LSR	$\mathcal{L}_{CeCo}$	mIoU	mAcc	oAcc	mIoU	mAcc	oAcc
Generalization	✗	✓	✓	✓	37.43	43.87	87.26	27.05	31.49	83.02
	✓	✗	✓	✓	55.57	61.56	93.51	44.41	51.38	89.71
	✓	✓	✗	✓	58.6	65.07	93.92	45.90	54.30	90.13
	✓	✓	✓	✗	58.77	65.38	<b>94.04</b>	47.11	55.59	90.51
	✓	✓	✓	✓	<b>59.24</b>	<b>66.04</b>	93.99	<b>49.11</b>	<b>57.04</b>	<b>90.68</b>
Finetuned	✓	✓	✗	✗	93.11	96.15	99.16	86.05	91.17	99.07
	✓	✓	✓	✓	92.11	94.81	99.25	<b>92.48</b>	<b>94.80</b>	<b>99.25</b>
	✓	✓	✗	✓	<b>94.31</b>	<b>96.19</b>	<b>99.33</b>	89.34	93.13	99.22

Table 4: Ablation study on Replica [25] for both the 20 most frequent and all classes.

$\mathbf{Z}_j^k$ . Removal of  $\hat{\mathbf{I}}_j^k$  and  $\mathbf{Z}_j^k$  results in significant decreases in mAcc by 22.17% and 4.48%, respectively. These findings highlight the vital importance of  $\hat{\mathbf{I}}_j^k$  for semantic segmentation performance. Additionally, the inclusion of view-independent features not only enhances performance but also improves view-consistency, as demonstrated in Figure 1.

**Evaluation of loss-terms.** In Table 4, we examine the impact of the label smoothing regularizer (LSR) [27] and loss-term  $\mathcal{L}_{CeCo}$  [40] through ablation. Results show that both components contribute positively to generalization performance. While the LSR does not provide benefits for the top20 frequent classes when finetuned, it increases performance when all classes are taking into account, suggesting its importance for the less frequent classes.

## 6 Limitations

Despite the significant improvements in both performance and speed achieved by RT-GS2, some limitations exist. While RT-GS2 enhances view-consistency for semantic segmentation by learning and leveraging view-independent 3D features (as qualitatively demonstrated in Figure 1), this approach does not fully guarantee strict view consistency across all perspectives. More specifically, while the primary elements of the scene generally remain stable, minor flickering can occur in smaller regions when the viewpoint changes.

## 7 Conclusion

This paper presents a novel real-time generalizable semantic segmentation method employing Gaussian Splatting. Through extensive experimentation, we have demonstrated its superiority over existing methods in both semantic segmentation quality and real-time performance, achieving significant improvements in mIoU on the Replica dataset and a remarkable 901 times speedup compared to current approaches. RT-GS2 represents a significant advancement in the field, providing the first real-time generalizable semantic segmentation method for 3D Gaussian representations of radiance fields.

## Acknowledgement

This work is funded by Innoviris within the research project SPECTRE and by Research Foundation Flanders (FWO) within the research project G094122N.

## References

- [1] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479, 2022.
- [2] Jiazhong Cen, Zanwei Zhou, Jiemin Fang, Wei Shen, Lingxi Xie, Dongsheng Jiang, Xiaopeng Zhang, Qi Tian, et al. Segment anything in 3d with nerfs. *Advances in Neural Information Processing Systems*, 36:25971–25990, 2023.
- [3] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *European Conference on Computer Vision*, pages 333–350. Springer, 2022.
- [4] Hanlin Chen, Chen Li, Mengqi Guo, Zhiwen Yan, and Gim Hee Lee. Gnesf: Generalizable neural semantic fields. *Advances in Neural Information Processing Systems*, 36, 2024.
- [5] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1290–1299, 2022.

- [6] Zi-Ting Chou, Sheng-Yu Huang, I Liu, Yu-Chiang Frank Wang, et al. Gsnerf: Generalizable semantic neural radiance fields with enhanced 3d scene understanding. *arXiv preprint arXiv:2403.03608*, 2024.
- [7] 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.
- [8] Leandro Di Bella, Yangxintong Lyu, and Adrian Munteanu. Deepkalpose: An enhanced deep-learning kalman filter for temporally consistent monocular vehicle pose estimation. *Electronics Letters*, 60(8):e13191, 2024.
- [9] Bin Dou, Tianyu Zhang, Yongjia Ma, Zhaohui Wang, and Zejian Yuan. Cosseg-gaussians: Compact and swift scene segmenting 3d gaussians. *arXiv preprint arXiv:2401.05925*, 2024.
- [10] Siqi Du, Weixi Wang, Renzhong Guo, and Shengjun Tang. Asymformer: Asymmetrical cross-modal representation learning for mobile platform real-time rgb-d semantic segmentation. *arXiv preprint arXiv:2309.14065*, 2023.
- [11] Xu Hu, Yuxi Wang, Lue Fan, Junsong Fan, Junran Peng, Zhen Lei, Qing Li, and Zhaoxiang Zhang. Semantic anything in 3d gaussians. *arXiv preprint arXiv:2401.17857*, 2024.
- [12] Mohammad Mahdi Johari, Yann Lepoittevin, and François Fleuret. Geonerf: Generalizing nerf with geometry priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18365–18375, 2022.
- [13] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4):1–14, 2023.
- [14] Hao Li, Dingwen Zhang, Yalun Dai, Nian Liu, Lechao Cheng, Jingfeng Li, Jingdong Wang, and Junwei Han. Gp-nerf: Generalized perception nerf for context-aware 3d scene understanding. *arXiv preprint arXiv:2311.11863*, 2023.
- [15] Fangfu Liu, Chubin Zhang, Yu Zheng, and Yueqi Duan. Semantic ray: Learning a generalizable semantic field with cross-reprojection attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17386–17396, 2023.
- [16] Yuan Liu, Sida Peng, Lingjie Liu, Qianqian Wang, Peng Wang, Christian Theobalt, Xiaowei Zhou, and Wenping Wang. Neural rays for occlusion-aware image-based rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7824–7833, 2022.
- [17] 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.



- [18] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)*, 41(4):1–15, 2022.
- [19] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017.
- [20] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017.
- [21] Minghan Qin, Wanhua Li, Jiawei Zhou, Haoqian Wang, and Hanspeter Pfister. Langsplat: 3d language gaussian splatting. *arXiv preprint arXiv:2312.16084*, 2023.
- [22] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 14335–14345, 2021.
- [23] Remco Royen and Adrian Munteanu. Rascal3d: Resolution scalable 3d semantic segmentation of point clouds. In *2023 IEEE International Conference on Image Processing (ICIP)*, pages 2775–2779. IEEE, 2023.
- [24] Myrna C Silva, Mahtab Dahaghin, Matteo Toso, and Alessio Del Bue. Contrastive gaussian clustering: Weakly supervised 3d scene segmentation. *arXiv preprint arXiv:2404.12784*, 2024.
- [25] 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.
- [26] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Superfast convergence for radiance fields reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5459–5469, 2022.
- [27] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.
- [28] Suhani Vora, Noha Radwan, Klaus Greff, Henning Meyer, Kyle Genova, Mehdi SM Sajjadi, Etienne Pot, Andrea Tagliasacchi, and Daniel Duckworth. Nesf: Neural semantic fields for generalizable semantic segmentation of 3d scenes. *arXiv preprint arXiv:2111.13260*, 2021.
- [29] Thang Vu, Kookhoi Kim, Tung M Luu, Thanh Nguyen, and Chang D Yoo. Softgroup for 3d instance segmentation on point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2708–2717, 2022.
- [30] Can Wang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. Clip-nerf: Text-and-image driven manipulation of neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3835–3844, 2022.

- [31] Huan Wang, Jian Ren, Zeng Huang, Kyle Olszewski, Menglei Chai, Yun Fu, and Sergey Tulyakov. R2l: Distilling neural radiance field to neural light field for efficient novel view synthesis. In *European Conference on Computer Vision*, pages 612–629. Springer, 2022.
- [32] Wenhai Wang, Jifeng Dai, Zhe Chen, Zhenhang Huang, Zhiqi Li, Xizhou Zhu, Xiaowei Hu, Tong Lu, Lewei Lu, Hongsheng Li, et al. Internimage: Exploring large-scale vision foundation models with deformable convolutions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14408–14419, 2023.
- [33] Wenxuan Wu, Zhongang Qi, and Li Fuxin. Pointconv: Deep convolutional networks on 3d point clouds. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, pages 9621–9630, 2019.
- [34] 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. *arXiv preprint arXiv:2312.10035*, 2023.
- [35] Saining Xie, Jiatao Gu, Demi Guo, Charles R Qi, Leonidas Guibas, and Or Litany. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 574–591. Springer, 2020.
- [36] Mingqiao Ye, Martin Danelljan, Fisher Yu, and Lei Ke. Gaussian grouping: Segment and edit anything in 3d scenes. *arXiv preprint arXiv:2312.00732*, 2023.
- [37] 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.
- [38] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenoc-trees for real-time rendering of neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5752–5761, 2021.
- [39] 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.
- [40] Shuaifeng Zhi, Tristan Laidlow, Stefan Leutenegger, and Andrew J Davison. In-place scene labelling and understanding with implicit scene representation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15838–15847, 2021.
- [41] Zhisheng Zhong, Jiequan Cui, Yibo Yang, Xiaoyang Wu, Xiaojuan Qi, Xiangyu Zhang, and Jiaya Jia. Understanding imbalanced semantic segmentation through neural collapse. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19550–19560, 2023.
- [42] 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. *arXiv preprint arXiv:2312.03203*, 2023.