

# Feature Splatting for Better Novel View Synthesis with Low Overlap

Tomas Berriel Martins

tberriel@unizar.es

Javier Civera

jcivera@unizar.es

I3A

University of Zaragoza

Zaragoza, Spain

---

## Abstract

3D Gaussian Splatting has emerged as a very promising scene representation, achieving state-of-the-art quality in novel view synthesis significantly faster than competing alternatives. However, its use of spherical harmonics to represent scene colors limits the expressivity of 3D Gaussians and, as a consequence, the capability of the representation to generalize as we move away from the training views. In this paper, we propose to encode the color information of 3D Gaussians into per-Gaussian feature vectors, which we denote as Feature Splatting (FeatSplat). To synthesize a novel view, Gaussians are first "splatted" into the image plane, then the corresponding feature vectors are alpha-blended, and finally the blended vector is decoded by a small MLP to render the RGB pixel values. To further inform the model, we concatenate a camera embedding to the blended feature vector, to condition the decoding also on the viewpoint information. Our experiments show that these novel model for encoding the radiance considerably improves novel view synthesis for low overlap views that are distant from the training views. Finally, we also show the capacity and convenience of our feature vector representation, demonstrating its capability not only to generate RGB values for novel views, but also their per-pixel semantic labels. Code is available on the [project page](#).

Keywords: Gaussian Splatting, Novel View Synthesis, Feature Splatting

## 1 Introduction

Finding appropriate 3D scene representations is a key challenge in order to enable robotics, VR and AR applications. In general, representations should not only encode a scene's geometry. Radiance properties are very relevant for novel view synthesis, and high-level aspects such as semantics or affordances are essential to define robotics tasks. In the last years, Neural Radiance Fields (NeRFs) [21] have been the dominant approach in the research literature to learn implicit scene representations from images. However, NeRFs are computationally demanding at training and rendering and scale badly with the size of the scene, which has limited their use in practical applications.

3D Gaussian Splatting (3DGS) [14] appeared very recently as a promising alternative to NeRFs. 3DGS is able to achieve similar quality for novel view synthesis significantly faster, although with a much higher memory footprint. Representing a scene with 3D Gaussians, however, does not come without limitations. As the most relevant to our work, a explicit

scene representation composed of 3D Gaussians and spherical harmonics lacks the capacity to model complex image textures. This limitation is typically alleviated by an increase of the number of Gaussians and a reduction of their size, so that each tiny Gaussian models a smaller and simpler part of the texture. Such strategy, in addition to being intensive in memory, leads to overfitting to the training views and to a poor generalization.

In this paper we propose a novel approach, that we denote as Feature Splatting (Feat-Splat), that replaces spherical harmonics by a feature vector representation for each 3D Gaussian. Using such representations, we show in ScanNet++ [58] a significant improvement in novel view synthesis with low overlap. This comes as a result of the higher capacity of our feature vector representation compared to spherical harmonics. Figure 1 illustrates our method. We splat the 3D Gaussians and alpha-blend the feature vectors. We then concatenate to the splatted and blended feature vector dimensions to camera intrinsics and extrinsics channels, and finally decode them into colors with a MLP. As an additional use case illustrating the flexibility of our representation, we also show its capability to render accurate semantic labels.

## 2 Related Work

In the last years, fully implicit scenes representations have been typically modelled as a function that maps scene coordinates to the desired scene attributes (*e.g.*, color), and a neural network is trained on the scene data so that it learns such function. One of the most successful examples are NeRFs [21]. A NeRF consists on a Multi Layer Perceptron (MLP) that maps a 3D point and a viewing direction to a RGB color and an occupancy or density value, coupled with a differentiable rendering equation to aggregate values of 3D points along a ray casted through each pixel to compute the final RGB color. Given a set of posed images, the MLP is optimized to perform novel view synthesis afterwards. To compensate for the initial high training and rendering times, several approaches have been developed. Some proposed to use depth information to reduce sampling in 3D [0, 23, 27, 55]; others combined the MLP with explicit 3D representations like Octrees [8, 40], Voxels [19], 3D point clouds [36] and 3D hierarchical mesh grids [22] and encoded into it learned features that where then quickly decoded using a smaller MLP.

Due to its ability to encode the scenes’ radiance and geometry, the original NeRF formulation has been extended with other capabilities and representations like depth estimation [8, 5, 6, 18]; Signed Distance Functions (SDFs) [32, 33]; semantics, for both panoptic [16] and open vocabulary [10] segmentation; and have been also integrated in a wide variety of use cases like scene editing [31, 41], robotics [29], and SLAM [28, 30, 42].

Recently, 3DGS [9] has emerged as a substantially faster yet still high-quality alternative. It relies on the explicit modelling of the 3D geometry by a set of 3D Gaussians, each of them with a set of spherical harmonics (SHs) [40] to represent their color. Rather than using ray tracing like NeRFs, Gaussians are rasterized by being projected into the image plane, sorted by the Radix algorithm [20], and combined using alpha-blending to give the RGB color for every pixel. The rendered image is then used to compute a photometric loss, and the Gaussians’ parameters are optimized using backpropagation. By avoiding the MLP training and relying in an efficient sorting algorithm and CUDA implementation of the rasterizer, 3DGS is the state of the art for the combined rendering performance, training time, and inference rendering speed. Nevertheless, 3DGS is also characterized by 1) a high memory usage; 2) poor surfaces’ reconstructions; and 3) the limited capacity of SHs, which

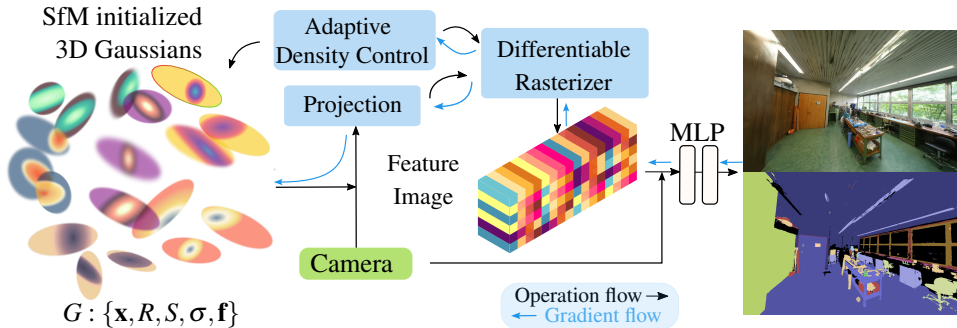


Figure 1: Overview of our approach. *Left*: We augment the 3D Gaussian Splatting [2] representation with learned feature vectors  $\mathbf{f}$  to encode colors, removing the spherical harmonics. *Center*: To render a point of view, 3D Gaussians are projected to the image plane, where the differentiable rasterizer alpha-blends corresponding feature vectors. Resulting vectors are concatenated with a camera embedding, and a tiny MLP renders the final RGB values, with potentially also higher level information such as semantic labels. *Right*: Illustrative result of novel view synthesis result with low overlap, with RGB values and semantic labels.

increases the number of Gaussians needed for high quality results, worsening the previous two problems. Several approaches augment Gaussians adding latent vectors which are later decoded. LangSplat [25] combines CLIP [26] features with SAM [12] segmentation masks to generate a language field enabling open-vocabulary 2D semantic segmentation. Similarly, both Feature 3DGS [43] and Gaussian Grouping [67], use SAM [12] as a teacher to embed in each Gaussian a semantic feature. FMGS [45] uses a hash grid based on iNGP [24] to encode CLIP and DINOv2 [24] features, to also generate an intrinsic representation capable of open-vocabulary 2D segmentation. Nevertheless, none of them neither address the SHs capacity nor show improvement in novel view synthesis, as we do. Based also on iNGP, Compact-3DGS [17] adds a hash grid that encodes a latent vector for each Gaussian and uses a small MLP conditioned by the viewing direction to recover the corresponding color for each point-of-view. Similarly, our FeatSplat relies on using a latent vector to encode each Gaussian’s color. But, differently from it, we directly attach a feature vector to each Gaussian, and, during the rasterization, rather than first decoding the colors and then alpha-blending, we first alpha-blend the feature vectors of corresponding Gaussians, and then decode the final vector to RGB, which allows us to condition on the intrinsic and extrinsic camera parameters. In our experiments, we show that FeatSplat, due to these differences, outperforms Compact-3DGS at novel view synthesis.

### 3 Preliminaries: 3D Gaussian Splatting

3DGS [2] is an explicit scene representation that relies on a set of 3D Gaussians to encode a scene, which is optimized by differentiable rasterization from known viewpoints.

**Representation.** The geometry of a 3D Gaussian  $G$  is parameterized by its 3D center position  $\mathbf{x}_G \in \mathbb{R}^3$  and a 3D covariance  $\Sigma_G = R_G S_G S_G^\top R_G^\top \succeq 0$  that can be decomposed into a rotation matrix  $R_G \in SO(3)$  and a scaling matrix  $S_G \in \text{diag}(\mathbb{R}_{\geq 0}^3)$ . Each Gaussian has a

density value  $\sigma_G \in \mathbb{R}$ ; and a set of spherical harmonics (SHs) coefficients  $\mathbf{k}_G \in \mathbb{R}^h$  that parameterize the function  $\mathcal{S}_{\mathbf{k}_G}(\mathbf{d}_G)$  to encode RGB values  $\mathbf{c}_G \in \mathbb{R}^3$  depending on the viewing direction  $\mathbf{d}_G \in S^2$  relative to each Gaussian.

**Differentiable rendering.** To render a RGB image  $\hat{I} \in \mathbb{R}^{w \times h \times 3}$  of the scene, first SHs are converted to RGB,  $\mathbf{c}_G = \mathcal{S}_{\mathbf{k}_G}(\mathbf{d}_G)$ . Then, Gaussians are projected, or splatted, into 2D Gaussians in the image plane [46] and sorted using the Radix algorithm. Finally, the color  $\mathbf{c}_p$  for each pixel  $\mathbf{p} \in \Omega$ ,  $\Omega$  standing for the image domain, comes from the alpha blending of the  $\mathcal{N} \subset G$  2D Gaussians falling into it

$$\mathbf{c}_p = \sum_{i \in \mathcal{N}} T_i \alpha_i \mathbf{c}_i, \text{ with } T_i = \prod_{j=i}^{i-1} (1 - \alpha_j), \quad (1)$$

where  $\alpha_i \in \mathbb{R}$  is obtained by multiplying the learned occupancy value  $\sigma_i$  by the evaluation of the projected 2D Gaussian [44, 45, 49] for each pixel  $\mathbf{p}$ .

**Optimization.** The representation is fitted as follows. The 3D Gaussians are initialized either randomly or from a sparse Structure from Motion point cloud. At each optimization step, the Gaussians are used to render a viewpoint from the training set and compute the loss

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

between the rendered image  $\hat{I}$  and the ground truth  $I$ .  $\mathcal{L}_1(\cdot, \cdot)$  stands for the  $L_1$ -norm of the difference of its two arguments,  $\mathcal{L}_{D-SSIM}(\cdot, \cdot)$  is the Structure Similarity Index Measure (SSIM) [44] between its two arguments, and  $\lambda_{SSIM}$  weights both terms. Adam [45] is used to update the Gaussians' parameters by backpropagating the loss, and Gaussians are pruned or cloned adaptively according to certain policies.

Despite SHs encoding direction-dependent RGB values, their variability is limited, and are only able to encode one color per viewing direction. To represent complex textures in natural scenes, 3DGS over-densifies scenes by optimizing Gaussians at different depths, visible only from certain viewpoints. In a manner, this strategy abuses the geometric primitive to overfit the photometry, achieving high-quality rendering for views close to the training distribution, but causing artifacts observable from viewpoints far from the training views.

## 4 Feature Splating (FeatSplat)

Our FeatSplat relies on learned feature vectors instead of SHs to better encode information and increase the capacity and flexibility of 3D Gaussians. Each 3D Gaussian is initialized with a feature vector  $\mathbf{f}_G$  sampled from a normal distribution  $\mathcal{N}(0, 1)$ . When rendering an image  $\hat{I}$ , the  $\mathcal{N}$   $\mathbf{f}_G$  vectors splatted in a pixel  $\mathbf{p}$  are alpha-blended

$$\mathbf{f}_p = \sum_{i \in \mathcal{N}} T_i \alpha_i \mathbf{f}_i, \quad (3)$$

similarly to what was done with colors in Eq. (1). After that, a small MLP transforms  $\mathbf{f}_p$ , into the corresponding RGB color  $\mathbf{c}_p$  of the target image. The per-pixel feature vectors  $\mathbf{f}_p$  are concatenated with camera parameter embeddings, in order to condition the MLP not only on the information of the 3D Gaussians, but also on the viewpoint. The camera extrinsics

are represented by the 3D camera position of the target image in global coordinates, and the intrinsics by an embedding of the relative position of the rendered pixel in the target image plane. The final rendered color is then

$$\mathbf{c}_p = \text{MLP}(\mathbf{f}_p \oplus \mathbf{x}_{cam} \oplus \mathbf{e}_p), \quad (4)$$

where  $\oplus$  is a concatenation operation,  $\mathbf{x}_{cam} \in \mathbb{R}^3$  is the 3D camera position, and  $\mathbf{e}_p \in [-1, 1] \times [-1, 1]$  the two-dimensional pixel embedding. We also considered adding camera orientation information, but our ablations (App. A.2) did not show a benefit.

**Semantic segmentation.** Learning per-Gaussian feature vectors  $\mathbf{f}_p$  increases the capability of a 3DGS representation. In this paper we illustrate how FeatSplat can be adapted for 2D semantic segmentation. All is needed is changing the output dimension of the MLP from 3 channels to  $3 + \#\text{classes}$  and training with an additional semantic cross-entropy loss  $\mathcal{L}_{ce}$ .

$$\mathcal{L} = (1 - \lambda) \mathcal{L}_1(\hat{\mathbf{I}}, \mathbf{I}) + \lambda \mathcal{L}_{D-SSIM}(\hat{\mathbf{I}}, \mathbf{I}) + \lambda_{sem} \mathcal{L}_{ce}(\hat{S}, S), \quad (5)$$

where  $\lambda_{sem}$  is a hyperparameter, and  $\hat{S}$  and  $S$  are the predicted and ground truth segmentation labels.




Figure 2: Comparison of 3D position of test (red) and training (blue) points of view on one scene for each of the four datasets.

## 5 Experiments

**Implementation details.** We implemented FeatSplat with both 16– and 32–dimensional feature vectors for novel view synthesis, and only 32–dimensional ones for novel view synthesis plus semantic segmentation. For each scene, we train a MLP with one hidden layer of 64 neurons, common for all Gaussians. For semantic segmentation, the model is trained on 64 classes: the 63 most frequent classes in ScanNet++, plus an unknown category label. Further implementation and optimization details are included in App. A.

**Baselines.** We compare FeatSplat against 3DGS [9], the current state of the art for real-time novel view synthesis, Compact-3DGS [17] which also augmented 3DGS with learned feature vectors instead of SHs, and iNGP [22], which was the previous state of the art. Finally, we also include Zip-NeRF and Mip-NeRF 360 metrics on the datasets they have been reported, as a reference. Nevertheless, FeatSplat should not be compared against them due to their offline nature (they both render at less than 1 frame per second). Finally, on

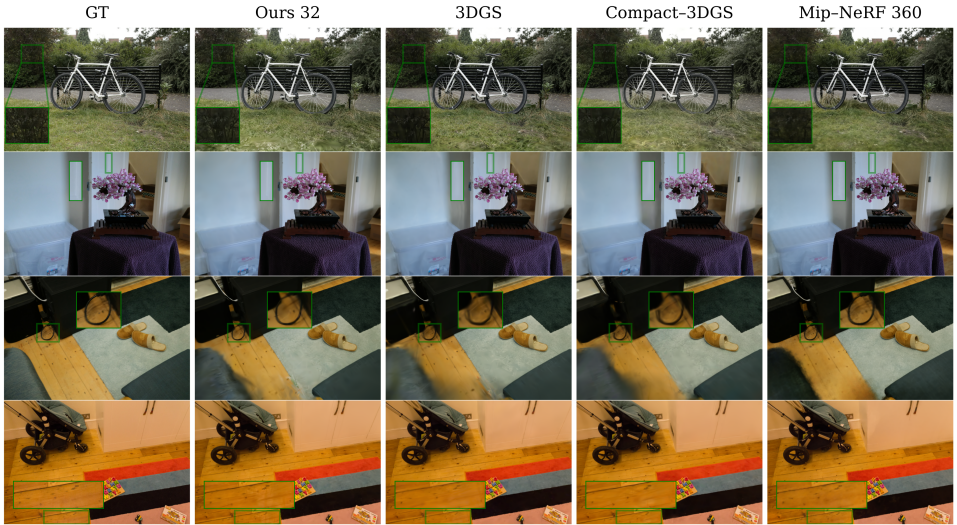


Figure 3: Comparison of our method trained with features with 16 (FeatSplat-16) and 32 dimensions (FeatSplat-32), against prior work. From top to down the scenes are *Bicycle*, *Bonsai*, and *Room* from Mip-360 [10], and *Playroom* from DB [12].

the semantic segmentation experiment, we do not compare against other models that learn specific semantic features for open-vocabulary semantic segmentation. Given that our model is trained on a closed-vocabulary, we would have an unfair advantage.

**Datasets and metrics.** Rendering performance is evaluated using PSNR, LPIPS [14], and SSIM [54], computational aspects by the representation size in megabytes (MB), rendering speed in Frames per Second (FPS) on a GTX 3090, and semantic segmentation by the mIOU weighted over classes’ support among all scenes. We evaluate on three common datasets in the literature [10, 12, 17]: Mip360 [10], Tanks and Temples [13] (T&Ts); and Deep Blending [12] (DB), for a total of 12 scenes. In these datasets the test set is selected by sampling one out of every eight frames of the image sequences, and the other 7 frames are used to train the model. We also evaluate FeatSplat in a subset of 21 scenes from the novel ScanNet++ [58]. ScanNet++ is a large-scale dataset with high-resolution RGB and semantic annotations of indoor scenes, and a test set generated from an independent camera trajectory specifically designed to evaluate novel view synthesis with low overlap and separated views.

The independent test set from ScanNet++ represents a considerable increase in difficulty with respect to commonly used datasets. Sampling one out of 8 frames from a video stream (as done in Mip-360, T&Ts and DB) means that at least the previous and following frames would have a significant overlap with the test frames. In contrast, an independent sequence from a different trajectory results in test images looking at unexplored parts of the scene or at the same areas but from rather different viewpoints. In Figure 2 we can visualize, for each of the four datasets, the scene with the highest average angular distance from each test viewpoint (red) to its 10 closest training viewpoints (blue). The contrast between datasets is clear. While on Mip-360, T&Ts, and DB, cameras tend to be nearby each other and looking in the same direction, on ScanNet++ there are several test views that are isolated and looking



Figure 4: Novel view synthesis with low overlap of FeatSplat-32 against 3DGS. From left to right the scenes are *Train* from T&Ts [13], *Treehill*, and *Bicycle* from Mip-360 [10].

at mostly unexplored areas. Further details on the datasets can be found in App. C.

	Mip-360					Tanks and Temples					Deep Blending				
	SSIM	PSNR	LPIPS	FPS	Mem	SSIM	PSNR	LPIPS	FPS	Mem	SSIM	PSNR	LPIPS	FPS	Mem
INGP-Big	0.699	25.59	0.331	9.4	48MB	0.745	21.92	0.305	14.4	48MB	0.817	24.96	0.390	2.8	48MB
3D Gaussian Splatting	0.813	27.39	0.218	115	880MB	0.844	23.65	0.179	156	431MB	0.899	29.46	0.247	126	662MB
Compact-3DGS	0.798	26.97	0.244	140	49MB	0.831	23.32	0.201	185	39MB	0.901	29.79	0.258	181	43MB
<b>FeatSplat-16 (Ours)</b>	0.808	27.48	0.230	96	275MB	0.848	24.48	0.185	171	149MB	0.904	29.78	0.249	111	212MB
<b>FeatSplat-32 (Ours)</b>	0.808	27.53	0.227	82	420MB	0.848	24.53	0.183	147	227MB	0.904	29.85	0.247	98	327MB
Mip-NeRF 360	0.792	27.29	0.237	0.06	9MB	0.759	22.22	0.257	0.14	9MB	0.901	29.4	0.245	0.09	9MB
ZipNeRF	0.836	28.54	0.177	0.25	-	-	-	-	-	-	-	-	-	-	-

Table 1: Evaluation on Mip-360 [10], T&Ts [13] and DB [10] for novel view synthesis.

## 5.1 Novel View Synthesis Results

**Mip-360, T&Ts and DB.** The quantitative results on Mip-360 dataset, T&Ts and DB in Table 1 show how both FeatSplat models achieve similar rendering performance. Although our 32-dimensions variant has a slight advantage on rendering quality (equal or better PSNR, LPIPS and SSIM), the 16-dimensional version has higher rendering speed and a smaller representation size. We conclude that 16 dimensions is enough to encode a scene’s RGB information. Comparing FeatSplat against 3DGS and Compact-3DGS, our method consistently achieves the best PSNR on all three datasets, with a better SSIM on two of them. Nevertheless, 3DGS has a better LPIPS overall. The qualitative results in Figure 3 show how all methods render high quality images, although with slight differences. 3DGS generates more accurate high-frequency details due to its tendency to over-densify scenes, e.g. the grass on the *Bicycle* scene and the trees on the *Train* scene. Instead, FeatSplat-32 is also able to capture details missed by 3DGS (e.g. the bush in *Bicycle*, the door in *Bonsai*, and the floor in *Playroom*) while reducing the amount of artifacts thanks to its more compact and less overfitted representation. This compactness can also be observed in the memory used, that is half for FeatSplat-32 and a fourth for FeatSplat-16. It can be noticed that the rendering speed is reduced for our FeatSplat in comparison to 3DGS and Compact-3DGS. However, we are still well within real-time limits, so we believe this is not a significant limitation.




Figure 5: Global illumination change modifying MLP input values of the pixel embedding ( $PE$ ), and the 3D coordinates ( $X$ ,  $Y$ ,  $Z$ ), on scene  $03f7a0e617$  from *ScanNet++*.

	ScanNet++				
	SSIM	PSNR	LPIPS	FPS	Mem
3D Gaussian Splatting	0.867	23.83	0.250	177	217
Compact 3DGS	0.854	22.76	0.272	170	35
<b>FeatSplat-16 (Ours)</b>	0.875	24.62	0.245	82	74
<b>FeatSplat-32 (Ours)</b>	0.875	24.62	0.245	71	112

Table 2: Features Splatting evaluation on 21 scenes from *ScanNet++*

**Out of distribution.** In order to analyze points of view far from the training cameras, we sample four points of view from simple linear trajectories on outdoor scenes, see Figure 4. Observe how, in *Train*, 3DGS generates floating 3D Gaussians to model far away clouds, which obfuscate the visibility for closer views. Observe instead in FeatSplat-32 that, as the camera gets closer to the train, the Gaussians that before were white, first switch to blue and finally also to green, to better represent the background sky and trees. A similar effect can be observed on the *Bicycle* scene in the right top corner of the image. Finally, the *Treehill* data reveals how FeatSplat learned to adapt the Gaussians’ brightness for slightly different positions of the camera while 3DGS did not.

**Light encoding.** Furthermore, we analyze the pixel and position embeddings impact on the MLP by individually evaluating different input values for each of them, see Figure 5. Apparently, the MLP learns to use these embeddings to codify global lightning information. As a consequence, our model allows to modify the scene lightning at inference time without any addition training.

**ScanNet++.** Table 2 shows that FeatSplat significantly outperforms both 3DGS and Compact-3DGS on all SSIM, PSNR, and LPIPS, with both FeatSplat-16 and FeatSplat-32 achieving the same performance. The gap between FeatSplat and 3DGS is due to a more challenging test set, as previously mentioned. To properly learn view-dependent effects, 3DGS overfits





Figure 6: Novel view synthesis of FeatSplat–16 and FeatSplat–32, against 3DGS [9] and Compact-3DGS [17] on a subset of scenes from ScanNet++ [38].

to the training cameras. Some Gaussians are placed in such positions that are seen only in the optimized direction, while not being visible from other training directions without the same effect. This effect is unnoticed from points of view that are close to the training samples, or with a high visual overlap. However, for viewpoints far-away from the training ones, or with lower visual overlap, they become easily noticeable. The ability of FeatSplat to encode multiple colors in the same Gaussian, and to condition the decoding on the point of view, reduces the impact of this problem. Figure 6 shows how FeatSplat renders more accurate colors, captures smaller details, and reduces the number of artifacts.

**Semantic FeatSplat** Our results so far suggest that 16 dimensions suffice to encode color. We then chose FeatSplat–32 to encode both RGB and semantic labels and perform semantic segmentation. FeatSplat’s segmentation masks for novel views achieve weighted mIOU of 0.629, maintaining the rendering performance with a PSNR of 24.64, and SSIM of 0.875

and an LPIPS of 0.244, and still rendering in real-time at 56 FPS. The segmentation maps generated, see Figure 7, show that the model produces coherent segmentations even for small elements. However, the edges of the masks are noisy, as they follow Gaussians’ shapes. Nevertheless, it proves the versatility of these features and their potential to encode different kind of information.




Figure 7: Semantic FeatSplat with 32-dimensional vectors and 64 classes on ScanNet++.

## 5.2 Limitations

The small size of our MLP allows real-time rendering but limits the complexity of textures encoded in each Gaussian. Therefore, our FeatSplat still generates small artifacts in unobserved areas. Furthermore, our model sometimes over-smooths some highly textured areas, like the grass or some tree leaves. This trade-off between capacity and speed can have a negative impact in applications focused on rendering around the training cameras, for which generalization is less relevant. Finally, FeatSplat is slower than 3DGS in part due to the CUDA kernel implementation designed for RGB channels, with rendering time scaling linearly with the dimension of the alpha-blended feature vector. Nevertheless, the increased capacity of FeatSplat vectors halves the number of Gaussians per scene, limiting the training and rendering slow-down to just a half of 3DGS at a resolution of  $1752 \times 1168$ .

## 6 Conclusions

In this paper we showed how the use of feature vectors to encode color information surpasses the performance of spherical harmonics when linked to small 3D primitives such as Gaussians for novel view synthesis. We argue that SHs lack capacity, and as a consequence an overpopulation of 3D Gaussians is required in order to approximate complex textures, which in turns results in the model overfitting the training viewpoints. Our FeatSplat overcomes this issue by using feature vectors, alpha blended, concatenated to camera parameters and decoded by an MLP, that render a higher variety of more complex textures and hence reduce the Gaussians’ overpopulation. FeatSplat outperforms previous methods on standard datasets in the literature and on the challenging ScanNet++ [58]. In addition, we keep a competitive real-time rendering and enable the extension to other tasks. Finally, it is worth mentioning that FeatSplat should also benefit from any 3DGS optimizations, like the quantization proposed by Compact-3DGS [17], which would speed up its rendering.

## References

- [1] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. *CVPR*, 2022.
- [2] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12882–12891, 2022.
- [3] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5501–5510, 2022.
- [4] Vitor Guizilini, Igor Vasiljevic, Jiading Fang, Rare Ambru, Greg Shakhnarovich, Matthew R Walter, and Adrien Gaidon. Depth field networks for generalizable multi-view scene representation. In *European Conference on Computer Vision*, pages 245–262. Springer, 2022.
- [5] Vitor Guizilini, Igor Vasiljevic, Dian Chen, Rareş Ambruş, and Adrien Gaidon. Towards zero-shot scale-aware monocular depth estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9233–9243, 2023.
- [6] Vitor Guizilini, Igor Vasiljevic, Jiading Fang, Rares Ambrus, Sergey Zakharov, Vincent Sitzmann, and Adrien Gaidon. Delira: Self-supervised depth, light, and radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17935–17945, 2023.
- [7] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. *ACM Transactions on Graphics (ToG)*, 37(6):1–15, 2018.
- [8] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus), 2023.
- [9] 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.
- [10] Justin Kerr, Chung Min Kim, Ken Goldberg, Angjoo Kanazawa, and Matthew Tancik. Lerf: Language embedded radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19729–19739, 2023.
- [11] Diederik P. Kingma and Jimmy Ba. Adam: a method for stochastic optimization. In *International Conference for Learning Representations*,, 2015.
- [12] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. *arXiv:2304.02643*, 2023.
- [13] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)*, 36(4):1–13, 2017.

- [14] Georgios Kopanas, Julien Philip, Thomas Leimkühler, and George Drettakis. Point-based neural rendering with per-view optimization. In *Computer Graphics Forum*, volume 40, pages 29–43. Wiley Online Library, 2021.
- [15] Georgios Kopanas, Thomas Leimkühler, Gilles Rainer, Clément Jambon, and George Drettakis. Neural point catacaustics for novel-view synthesis of reflections. *ACM Transactions on Graphics (TOG)*, 41(6):1–15, 2022.
- [16] Abhijit Kundu, Kyle Genova, Xiaoqi Yin, Alireza Fathi, Caroline Pantofaru, Leonidas J Guibas, Andrea Tagliasacchi, Frank Dellaert, and Thomas Funkhouser. Panoptic neural fields: A semantic object-aware neural scene representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12871–12881, 2022.
- [17] Joo Chan Lee, Daniel Rho, Xiangyu Sun, Jong Hwan Ko, and Eunbyung Park. Compact 3d gaussian representation for radiance field. *arXiv preprint arXiv:2311.13681*, 2023.
- [18] Jiaxin Li, Zijian Feng, Qi She, Henghui Ding, Changhu Wang, and Gim Hee Lee. Mine: Towards continuous depth mpi with nerf for novel view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12578–12588, 2021.
- [19] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *Advances in Neural Information Processing Systems*, 33:15651–15663, 2020.
- [20] Duane G Merrill and Andrew S Grimshaw. Revisiting sorting for gpgpu stream architectures. In *Proceedings of the 19th international conference on Parallel architectures and compilation techniques*, pages 545–546, 2010.
- [21] 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.
- [22] 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.
- [23] Thomas Neff, Pascal Stadlbauer, Mathias Parger, Andreas Kurz, Joerg H Mueller, Chakravarty R Alla Chaitanya, Anton Kaplanyan, and Markus Steinberger. Donerf: Towards real-time rendering of compact neural radiance fields using depth oracle networks. In *Computer Graphics Forum*, volume 40, pages 45–59. Wiley Online Library, 2021.
- [24] Maxime Oquab, Timothée Darcet, Theo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Russell Howes, Po-Yao Huang, Hu Xu, Vasu Sharma, Shang-Wen Li, Wojciech Galuba, Mike Rabbat, Mido Assran, Nicolas Ballas, Gabriel Synnaeve, Ishan Misra, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. Dinov2: Learning robust visual features without supervision, 2023.

- [25] Minghan Qin, Wanhua Li, Jiawei Zhou, Haoqian Wang, and Hanspeter Pfister. Langsplat: 3d language gaussian splatting. *arXiv preprint arXiv:2312.16084*, 2023.
- [26] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [27] Barbara Roessle, Jonathan T Barron, Ben Mildenhall, Pratul P Srinivasan, and Matthias Nießner. Dense depth priors for neural radiance fields from sparse input views. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12892–12901, 2022.
- [28] Erik Sandström, Yue Li, Luc Van Gool, and Martin R Oswald. Point-slam: Dense neural point cloud-based slam. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18433–18444, 2023.
- [29] William Shen, Ge Yang, Alan Yu, Jansen Wong, Leslie Pack Kaelbling, and Phillip Isola. Distilled feature fields enable few-shot language-guided manipulation. In *Conference on Robot Learning*, pages 405–424. PMLR, 2023.
- [30] Edgar Sucar, Shikun Liu, Joseph Ortiz, and Andrew J Davison. imap: Implicit mapping and positioning in real-time. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6229–6238, 2021.
- [31] Bing WANG, Lu Chen, and Bo Yang. DM-neRF: 3d scene geometry decomposition and manipulation from 2d images. In *The Eleventh International Conference on Learning Representations*, 2023.
- [32] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *Advances in Neural Information Processing Systems*, 34:27171–27183, 2021.
- [33] Yiming Wang, Qin Han, Marc Habermann, Kostas Daniilidis, Christian Theobalt, and Lingjie Liu. Neus2: Fast learning of neural implicit surfaces for multi-view reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023.
- [34] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004. doi: 10.1109/TIP.2003.819861.
- [35] Yi Wei, Shaohui Liu, Yongming Rao, Wang Zhao, Jiwen Lu, and Jie Zhou. Nerfingmvs: Guided optimization of neural radiance fields for indoor multi-view stereo. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5610–5619, 2021.
- [36] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf: Point-based neural radiance fields. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5438–5448, 2022.

- [37] Mingqiao Ye, Martin Danelljan, Fisher Yu, and Lei Ke. Gaussian grouping: Segment and edit anything in 3d scenes. In *European Conference on Computer Vision*. Springer, 2024.
- [38] Chandan Yeshwanth, Yueh-Cheng Liu, Matthias Nießner, and Angela Dai. Scannet++: A high-fidelity dataset of 3d indoor scenes. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2023.
- [39] Wang Yifan, Felice Serena, Shihao Wu, Cengiz Öztireli, and Olga Sorkine-Hornung. Differentiable surface splatting for point-based geometry processing. *ACM Transactions on Graphics (TOG)*, 38(6):1–14, 2019.
- [40] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenotrees for real-time rendering of neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5752–5761, 2021.
- [41] Jiakai Zhang, Xinhang Liu, Xinyi Ye, Fuqiang Zhao, Yanshun Zhang, Minye Wu, Yingliang Zhang, Lan Xu, and Jingyi Yu. Editable free-viewpoint video using a layered neural representation. *ACM Transactions on Graphics*, 40(4):1–18, 2021.
- [42] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [43] Shijie Zhou, Haoran Chang, Sicheng Jiang, Zhiwen Fan, Zehao Zhu, Dejjia Xu, Pradyumna Chari, Suyu You, Zhangyang Wang, and Achuta Kadambi. Feature 3dgs: Supercharging 3d gaussian splatting to enable distilled feature fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21676–21685, 2024.
- [44] Zihan Zhu, Songyou Peng, Viktor Larsson, Weiwei Xu, Hujun Bao, Zhaopeng Cui, Martin R Oswald, and Marc Pollefeys. Nice-slam: Neural implicit scalable encoding for slam. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12786–12796, 2022.
- [45] Xingxing Zuo, Pouya Samangouei, Yunwen Zhou, Yan Di, and Mingyang Li. Fmgs: Foundation model embedded 3d gaussian splatting for holistic 3d scene understanding. *arXiv preprint arXiv:2401.01970*, 2024.
- [46] Matthias Zwicker, Hanspeter Pfister, Jeroen Van Baar, and Markus Gross. Ewa volume splatting. In *Proceedings Visualization, 2001. VIS'01.*, pages 29–538. IEEE, 2001.