

GaussianBeV : 3D Gaussian Representation meets Perception Models for BeV Segmentation

Florian Chabot

Nicolas Granger

Guillaume Lapoue

CEA, List, F-91120, Palaiseau, France

{firstname.lastname}@cea.fr

Abstract

The Bird’s-eye View (BeV) representation is widely used for 3D perception from multi-view camera images. It allows to merge features from different cameras into a common space, providing a unified representation of the 3D scene. The key component is the view transformer, which transforms image views into the BeV. However, actual view transformer methods based on geometry or cross-attention do not provide a sufficiently detailed representation of the scene, as they use a sub-sampling of the 3D space that is non-optimal for modeling the fine structures of the environment. In this paper, we propose GaussianBeV, a novel method for transforming image features to BeV by finely representing the scene using a set of 3D gaussians located and oriented in 3D space. This representation is then splattered to produce the BeV feature map by adapting recent advances in 3D representation rendering based on gaussian splatting [12]. GaussianBeV is the first approach to use this 3D gaussian modeling and 3D scene rendering process in an optimization free manner, i.e. without optimizing it on a specific scene and directly integrated into a single stage model for BeV scene understanding. Experiments show that the proposed representation is highly effective and place GaussianBeV as the new state-of-the-art on the BeV semantic segmentation task on the nuScenes dataset [2].

1. Introduction

Multi-camera 3D perception tasks, such as semantic segmentation, are crucial for autonomous navigation applications. A common strategy involves projecting and merging features from different cameras into a bird’s-eye view (BeV) representation, which is then analyzed by perception heads. The primary challenge in these approaches lies in addressing the loss of 3D information during the projection of the physical world into camera images, thus solving the inverse problem of transforming image views into the BeV.

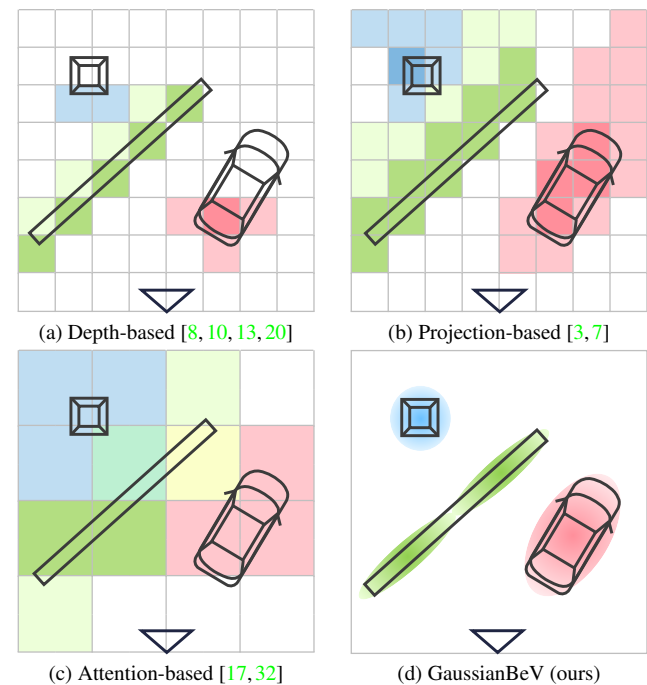


Figure 1. Illustration of multiple BeV representations for BeV semantic segmentation. A camera is represented by the triangle at the bottom of each BeV. Features are represented by colors where blue, red and green represent the streetlight, car and lane marking respectively. (a) Depth-based methods place image features along the optical ray on the surface of objects. (b) In projection-based methods, 3D points on the optical ray receive the same feature. (c) Attention-based methods use downsampled dense spatial queries to keep memory costs down. (d) In GaussianBeV, the scene is represented by a set of rotated gaussians that finely describes the semantic structures in the scene.

Recent literature identifies three main subsets of methods for image-to-BeV transformation. First, depth-based methods [8, 10, 13, 20] achieve view transformation geometrically by filling a 3D grid with features extracted from images

based on the prediction of discrete depth distribution. The key idea is to roughly localize the 3D position of each image feature and then accumulate them through a voxelization step. However, in these approaches, 3D feature localization depends on the depth discretization granularity and is sub-optimal, as features are placed at the level of the visible faces of objects. Second, projection-based methods [3, 7] also use a geometric approach, project 3D grid points into the cameras and gathers corresponding features from them. While straightforward, these methods do not yield accurate 2D-to-3D back-projection, as all grid points along the same camera ray receive the same feature. Third, transformer-based methods [15, 17, 32] utilize cross-attention to merge multi-view features. Though effective for 3D object detection, their application to dense tasks like BeV semantic segmentation incurs a high computational cost due to dense spatial queries needed for BeV representation in the attention process. Some works [17, 32] address this by reducing the BeV resolution, leading to inevitable information loss.

In this article, we propose a novel view transformation method called GaussianBeV, which enables fine 3D modeling of scenes. Drawing on recent advances in explicit 3D representation for rendering novel views based on Gaussian Splatting (GS) [12], our method represents a scene using a set of 3D gaussians, each parameterized by a center, scale, rotation, opacity and *semantic features* (instead of colors in GS). Furthermore, unlike the original GS method, which uses offline optimization for a specific scene to determine the 3D gaussian representation, we propose to train a neural network to directly generate an *optimization free* 3D gaussian representation of the scene from multi-view images. This representation is then rendered into a BeV feature map which is analyzed by semantic segmentation heads.

The representation of a scene by a set of 3D gaussians allows to model its entire content. Indeed, the geometrical properties of the gaussians (position, size and rotation) enable to cover 3D space with varying level of detail depending on the structures encountered in the scene. Intuitively, a gaussian representing a lane marking will be rotated and elongated along its length. A gaussian representing a vehicle will be placed in the center of the vehicle and will follow its shape. Figure 1d illustrates the intuition behind the representation proposed in this paper. Our contributions can be summarized as follows. (1) Introduction of GaussianBeV for BeV feature map generation from images through an *optimization free* image-to-3D gaussian representation for any scene, allowing fine 3D content modeling. This representation is then splattered in BeV using a rasterizer module. To our knowledge, this is the first time that a gaussian splatting representation that is not scene-specific is proposed and integrated into a BeV perception model. (2) Experiments demonstrating the effectiveness of our method, establishing it as the new state-of-the-art in BeV semantic segmentation.

2. Related work

Depth-based. A series of models were built on the explicit prediction of pixel-wise depth estimations along with image features. Combined with camera calibration parameters, this enables the back-projection of 2D features into 3D feature point cloud which is finally aggregated in the BeV grid. To accommodate for uncertainty in depth estimation, the features are actually propagated all along the ray that traverses their pixel and modulated by a discrete depth probability estimation [8, 10, 20]. To improve depth prediction, an explicit depth supervision scheme is proposed [13] using LiDAR data during model training. However, depth-based methods are sensitive to the ray sampling strategy, usually back-projecting features along the ray and on object surfaces (see Figure 1a).

Projection-based. Using a thorough comparative study, [7] proposes to discard depth estimation in favor of a simpler projection scheme: a predefined set of 3D points is used to describe the scene and camera features are probed by projecting the point on the camera feature maps using calibration data. This projection disregards actual objects and background placement but returns a denser representation of the scene without void beyond the depth of objects surfaces. The computational and memory overhead of the generating the BeV grid is reduced by opting for a sparse grid representation [3]. Projection-based view transformation methods are simple but result in a coarse BeV representation because all voxels along the optical ray receive the same features (see Figure 1b).

Attention-based. Capitalising on recent advances in Transformer models, depth estimation is replaced by an attention-based feature modulation scheme [15–17, 25, 32]. Several optimization schemes are proposed to resolve computational complexity of a pairwise matching between image and BeV grid tokens: factorization of spatial and temporal attentions, deformable attention [15], injection of calibration and timestamp priors [17]. For the segmentation task, attention-based view transformation is computationally and memory intensive, due to the need to define a dense query map [15]. This is why some methods [17, 32] predict a low-resolution BeV (see Figure 1c), which is then upsampled by successive deconvolutions.

Gaussian splatting. Gaussian splatting (GS) [12] is a 3D scene rendering technique which uses 3D gaussians to describe a scene. Each gaussian is parameterized by its position, scale, rotation, opacity and a Spherical Harmonics color model. The entire rendering pipeline is differentiable allowing the optimization of the gaussian parameters to a particular scene based on a set of images. GS is both fast and parallel, allowing real-time operation on GPUs. In comparison to sparse voxel grids, gaussians offer more efficient representations of a scene since individual gaussians can describe large volumes while smaller ones can accu-




Figure 2. **Overview of GaussianBeV.** The network takes as input a set of multiview images and extracts features for each of them. The 3D gaussian generator module (Sec 3.2) predicts a 3D gaussian representation G of the scene which is then sent to the BeV rasterizer module (Sec 3.3) to performs BeV rendering. The resulting BeV feature map B is passed through a BeV backbone and segmentation heads to obtain the segmentation prediction. G and B are represented with colors only for visualization purpose.

rately encode finer details with arbitrary resolution. Several extensions have been proposed, allowing the management of dynamic objects [14, 28] or the distillation of semantic features from foundation models in the representation [21, 22, 30, 34]. Although GS rasterization itself is fast, the optimization of gaussian parameters to match a scene usually requires many iterations. In order to address this issue for real-time SLAM, [18] leverages ground-truth 3D information from a depth sensor to initialize gaussian positions and [27] uses temporal consistency in order to minimize the optimization workload between successive frames. To perform scene reconstruction with stereo images [4] or object-centric single image [23], optimization-free methods are introduced allowing the model to directly learn to predict the gaussian parameters.

In our work, we propose to overcome the drawbacks of view transformer methods by modelling scenes as set of gaussians which are subsequently rasterized into a feature map. Unlike previous gaussian splatting works which optimize the gaussian representations to each scene, we propose to learn a neural network capable of directly predicting a gaussian representation of any scene. Compared to [4, 23] which targets scene reconstruction, our model is designed to perform real-time outdoor BeV semantic segmentation from multiple calibrated cameras.

3. GaussianBeV

3.1. Overview

Figure 2 presents an overview of GaussianBeV. The model takes as input a set of multiview images $I \in \mathbb{R}^{N \times H \times W \times 3}$ with N the number of cameras, H and W the dimensions of the images. These images are passed sequentially through four modules, leading to BeV segmentation.

The first module extracts image features using an image backbone and a neck to obtain feature maps $F \in \mathbb{R}^{N \times H_F \times W_F \times C_F}$, with C_F the number of channels, H_F and W_F the dimensions of the feature maps.

The second module is the 3D gaussian generator (Sec 3.2) that predicts for each pixel in the feature maps

the parameters of the corresponding gaussian in the world reference frame. The output of this module is a set of 3D gaussians $G \in \mathbb{R}^{NH_F W_F \times (C+11)}$ with C the number of channels of the embedding associated to each gaussian. More specifically, G contains the following parameters : positions $P \in \mathbb{R}^{NH_F W_F \times 3}$, scales $S \in \mathbb{R}^{NH_F W_F \times 3}$, rotations as unit quaternions $Q \in \mathbb{R}^{NH_F W_F \times 4}$, opacities $O \in \mathbb{R}^{NH_F W_F \times 1}$ and embeddings $E \in \mathbb{R}^{NH_F W_F \times C}$. First, the module predicts a set of 3D gaussian for each camera in its own camera reference frame. Next, camera extrinsic parameters are applied to transform 3D gaussians from the camera to the world reference frame to finally concatenate all the gaussians into the single set G .

The third module is the BeV rasterizer (Sec 3.3) that performs a BeV rendering of the 3D gaussian set G to produce the BeV feature map $B \in \mathbb{R}^{H_B \times W_B \times C}$, with H_B and W_B the dimensions of the BeV map.

Finally, in the last module, a BeV backbone and segmentation heads are sequentially applied to the BeV features to provide the final prediction.

3.2. 3D Gaussian generator

Given the input feature maps F , the 3D gaussian generator predicts the 3D gaussian representation of the scene using several prediction heads. Figure 3 illustrates how it operates on the feature maps.

Gaussian centers. The 3D positions of the gaussians in the scene are estimated by a depth head and a 3D offset head applied to F . The first predicts an initial position of the 3D centers along the optical rays. The second refines this 3D position by adding a small 3D displacement to it, giving more flexibility in the positioning of gaussians by not freezing them along the optical rays.

More precisely, for a pixel i in the feature map of the camera n with coordinates $(u_{n,i}, v_{n,i})$, the depth head predicts the disparity $d_{n,i} \in [0, 1]$ as in previous works dealing with monocular depth map estimation [6, 29]. To compensate the influence of focal length diversity from one camera to another on depth prediction, disparity is predicted up to

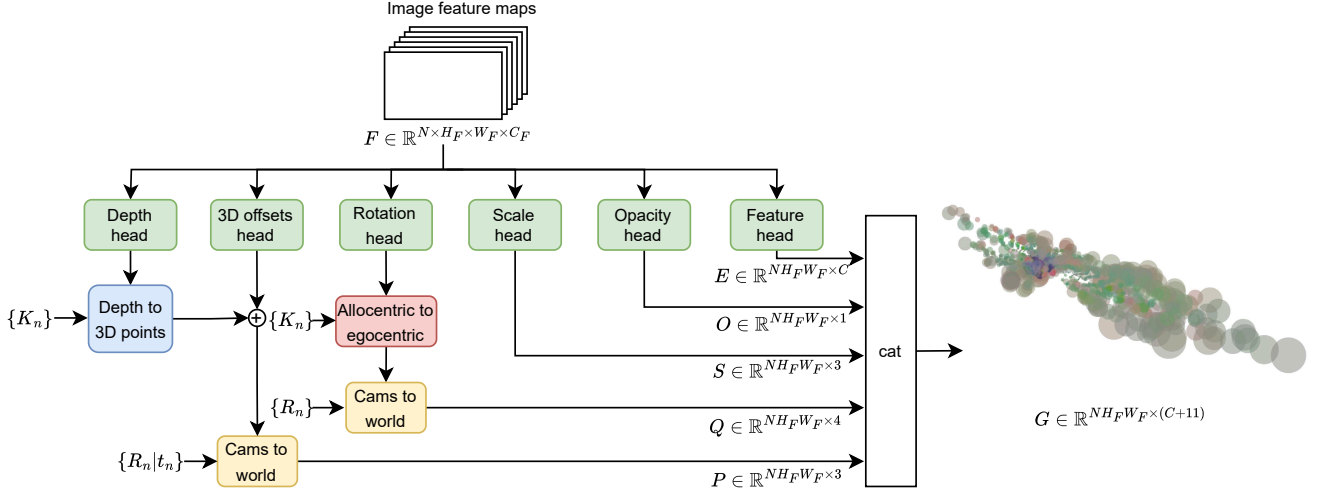


Figure 3. **3D Gaussian generator.** This module takes as input the feature map extracted from each camera, as well as the set of intrinsic $\{K_n\}$ and extrinsic $\{R_n|t_n\}$ parameters. For each pixel, it calculates the corresponding 3D gaussian by passing through prediction heads (green boxes). Some of these predictions are then decoded (blue and red boxes) and transformed to be expressed in the world reference frame (yellow boxes). All predicted gaussian parameters are then concatenated to produce the 3D gaussian representation G of the scene.

a scaling factor, in a reference focal length f as proposed in [31]. Knowing the true focal length f_n associated to the camera n , the metric depth $z_{n,i}$ is then decoded as follows:

$$z_{n,i} = \frac{f_n}{f} \left(\frac{1}{d_{n,i}} - 1 \right) \quad (1)$$

The corresponding 3D point $p_{n,i}^c$ in the camera reference frame is then deduced using the intrinsic matrix K_n of the n -th camera:

$$p_{n,i}^c = K_n^{-1} \begin{bmatrix} z_{n,i} u_{n,i} \\ z_{n,i} v_{n,i} \\ z_{n,i} \end{bmatrix} \quad (2)$$

The resulting 3D points are constrained to lie along the optical ray passing through the pixel under consideration. Because of this constraint, their positioning is not necessarily optimal. To overcome this problem, we propose to use the 3D offset prediction head. It aims to provide a small displacement $\Delta_{n,i} = (\Delta x_{n,i}, \Delta y_{n,i}, \Delta z_{n,i})^T$ to be applied to the 3D center of the gaussian $p_{n,i}^c$ to refine its position in all three directions. The refined 3D point $\bar{p}_{n,i}^c$ is simply obtained by:

$$\bar{p}_{n,i}^c = p_{n,i}^c + \Delta_{n,i} \quad (3)$$

At this stage, the 3D gaussian centers calculated for each camera are expressed in the corresponding camera reference frame. To express these points in the world reference frame, the extrinsic parameter matrices $[R_n|t_n]$ are applied, allowing the camera-to-world transformation :

$$p_{n,i}^w = [R_n|t_n] \bar{p}_{n,i}^c \quad (4)$$

The result is the gaussian center set $P = \{p_{n,i}^w\} \in \mathbb{R}^{N H_F W_F \times 3}$.

Gaussian rotations. The 3D rotations of the gaussians in the scene are estimated by a rotation head applied to F . For a given pixel in the feature map of the camera n , it outputs an allocentric rotation in the form of a unit quaternion $q_{n,i}^a$. The allocentric rotation of a pixel corresponds to a rotation relative to the 3D optical ray passing through it. This modelization makes it easier for the rotation head to learn, as it has no knowledge of the optical ray corresponding to the pixel it is processing. To take an example, two objects placed at two different locations in the scene and with different absolute (egocentric) rotations in the camera reference frame may have the same appearance in the image. In this case, the allocentric rotation predicted by the rotation head will be the same. The intrinsic parameters of the camera are then used to retrieve the egocentric rotation information.

For that purpose, the quaternion $q_{n,i}$ representing the rotation between the optical ray passing through the pixel i of the camera n and the axis $[0, 0, 1]^T$ is calculated. The quaternion $q_{n,i}^e$ representing the egocentric rotation in the camera reference frame is then recovered by:

$$q_{n,i}^e = q_{n,i} q_{n,i}^a \quad (5)$$

Finally, as for gaussian centers, the quaternion $q_{n,i}^w$ representing the rotation of the gaussian in the world reference frame is calculated using q_{R_n} , the quaternion modeling the camera-to-world rotation of the camera n :

$$q_{n,i}^w = q_{R_n} q_{n,i}^e \quad (6)$$

The quaternions thus calculated form the set of gaussian rotations $Q = \{q_{n,i}^w\} \in \mathbb{R}^{NH_F W_F \times 4}$.

Gaussian scales, opacities and features. The last three gaussian parameters do not depend on optical properties and camera positioning, but rather encode semantic properties. Therefore, three heads are simply used to predict the sets, S , O and E required to render the gaussian set G by the BeV rasterizer module.

3.3. BeV rasterizer

The BeV rasterizer module is used to obtain the BeV feature map $B \in \mathbb{R}^{H_B \times W_B \times C}$ from the set of gaussians G predicted by the 3D gaussian generator. To this end, the differentiable rasterization process proposed in gaussian splatting [12] has been adapted to perform this rendering. The first adaptation, already proposed in other offline semantic reconstruction works [21], consists in rendering C -dimensional features rather than colors. In our case, this produces a rendering containing semantic features essential for the perception task to solve. The second adaptation concerns the type of projection used. We have parameterized the rendering algorithm to generate orthographic rather than perspective renderings, more suited for BeV representation of the scene.

3.4. GaussianBeV training

Our model is trained end-to-end using the loss functions used in [3, 7, 8]. In these previous works, the semantic segmentation loss L_{sem} is defined as follows:

$$L_{sem} = \lambda_{bce} L_{bce} + \lambda_{ctr} L_{ctr} + \lambda_{off} L_{off} \quad (7)$$

With L_{bce} the binary cross-entropy loss. L_{ctr} and L_{off} correspond to the centerness and offset losses respectively, used as auxiliary losses to regularize training. λ are weights to balance the three losses.

Gaussian regularization losses. Although GaussianBeV can be trained efficiently with the aforementioned losses, the addition of regularization functions acting directly on the gaussian representation improves its representational qualities. In particular, two regularization losses are added during training.

First, a depth loss aims to regularize the position of the gaussians using the depth information provided by the projection of the LiDAR in the images. This loss adds constraints on the depth head predictions to obtain an initial 3D position, which is then refined by the 3D offsets (see Sec 3.2). With z the ground truth depth and z^* the predicted depth, the depth loss L_{depth} is defined as follows:

$$L_{depth}(z, z^*) = |\log(z) - \log(z^*)| \quad (8)$$

Second, an early supervision loss aims to optimize the gaussian representation before BeV backbone + heads are applied. The idea is to constrain the BeV features to directly provide all necessary information for the semantic segmentation task. In practice, segmentation heads are added and connected directly to the output of the BeV rasterizer module. The early supervision loss L_{sem}^{early} is defined similarly to L_{sem} (see equation 7). The total loss function is therefore defined by:

$$L = L_{sem} + L_{sem}^{early} + \lambda_{depth} L_{depth}, \quad (9)$$

The influence of these learning strategies is analyzed in section 4.2.

4. Experiments

Dataset. We use the nuScenes dataset [2] consisting in a set of multiview sequences obtained with a system of $N=6$ surrounding cameras. It is divided into 750 sequences (28,130 images) for training and 150 sequences (6,019 images) for validation. It provides 3D bounding box annotations (vehicles and pedestrians) and semantized HDmap annotations of the road surface (drivable area and lane boundaries).

Architecture details. For image feature extraction, the EfficientNet-b4 backbone [24] and the Simple-BeV neck [7] are used. For the BeV backbone, unless specified, the model of LSS [20] and the segmentation heads of Simple-BEV [7] are connected to the output of the BeV rasterizer. We use a channel size of $C=128$ for all experiments. For gaussian prediction heads, we use a light convolutional neural network composed by two blocks of Conv-Batchnorm-ReLU followed by the output layer. The sigmoid activation function is applied on both disparity and opacity outputs. L2 normalization and absolute value activation are applied on the rotation and scale outputs, respectively.

Implementation details. Following previous work [3], GaussianBeV is trained on a maximum of 100 epochs using an AdamW optimizer with learning rate of $3e-4$, a weight decay of $1e-7$ and a one-cycle linear learning rate scheduler. As in [3, 8], loss functions are balanced using uncertainty weighting as proposed in [11]. We use random scaling, random crop and random flip for images augmentation and random translation and rotation for BeV augmentation. Two input image resolutions $H \times W$ are tested in the experiments, 224×480 and 448×800 . The resulting image feature maps have a size $H_F \times W_F$ of 28×60 and 56×100 , respectively. GaussianBeV training on the lower resolution uses a batch size of 11 on an A100 GPU and for the higher resolution a batch size of 9, distributed over 3 A100 GPUs. The BeV rasterizer output a BeV feature map of size 200×200 representing a $100m \times 100m$ BeV representation with a 50cm resolution. The IoU metric is used on the validation set for evaluation.

	Image backbone	all vehicles		vehicles with vis > 40%	
		224×448	448×800	224×448	448×800
FIERY [8]	EN-b4	35.8	-	39.8	-
CVT [32]	EN-b4	31.4	32.5	36.0	37.7
LaRa [1]	EN-b4	35.4	-	38.9	-
BEVFormer [15]	RN-50	35.8	39.0	42.0	45.5
Simple-BEV [7]	RN-50	36.9	40.9	43.0	46.6
BAEFormer [19]	EN-b4	36.0	37.8	38.9	41.0
PointBeV [3]	EN-b4	38.7	42.1	44.0	47.6
GaussianBeV	EN-b4	41.6	43.9	47.5	50.3

Table 1. GaussianBeV IoU results on vehicle segmentation on the nuScenes validation set using different input resolutions and different visibility filtering. Our method remarkably outperforms previous works on all experimental settings.

4.1. State-of-the-art comparison

Vehicle segmentation. We compare GaussianBeV with previous works on vehicle semantic segmentation using different input resolutions (224×448 and 448×800) and different visibility filtering : (1) considering all vehicles and (2) only keeping vehicles with visibility > 40%. Results are given in Table 1. It shows that GaussianBeV clearly outperforms previous methods on all experimental settings. For instance, it surpasses the previous state-of-the-art method PointBeV [3] by +3.5 IoU for the experiments using 224×448 input resolution and visibility filtering.

Pedestrian segmentation. We also compare GaussianBeV to other previous methods on the pedestrian segmentation task. For this evaluation, we use the 224×448 input resolution and visibility filtering. Results are available in Table 2. Once again, GaussianBeV outperforms previous state-of-the-art method PointBeV [3] by +2.7 IoU.

Ground surface segmentation. We train GaussianBeV for joint segmentation of drivable area and lane boundaries using an input resolution of 448×800. Results are given in Table 3. Compared to previous state-of-the-art method MatrixVT [33], GaussianBeV gives superior results for lane boundary segmentation (+2.6 IoU). However, our method is slightly less effective in segmenting the drivable area (-0.9 IoU) than MatrixVT. The ability of GaussianBeV to model the scene in details enables better segmentation of fine structures, but does not improve performance in larger areas that are easier to segment.

Inference time. We compared for both input resolutions the inference time of GaussianBeV and the previous method, PointBeV [3], on an A100 GPU. GaussianBeV runs at 24 fps and 13 fps while PointBeV runs at 19 fps and 15 fps. Our method is therefore comparable in terms of computation time, but future research may enable to speed up the model by optimizing the gaussian representation.

	Pedestrians
LSS [20]	15.0
FIERY [8]	17.2
ST-P3 [9]	14.5
TBP-Former static [5]	17.2
PointBeV [3]	18.5
GaussianBeV	21.2

Table 2. GaussianBeV IoU results on pedestrian segmentation on the nuScenes validation set. GaussianBeV is trained with 224×448 input image resolution and visibility filtering. Our method increases performances compared to previous works.

	Drivable area	Lane boundaries
LSS [20]	72.9	20.0
M ² BEV [26]	75.9	38.0
BEVFormer [15]	80.1	25.7
PETRv2 [17]	83.3	44.8
MatrixVT [33]	83.5	44.8
PointBeV [3] [†]	72.5	37.6
GaussianBeV	82.6	47.4

Table 3. GaussianBeV IoU results on both drivable area and lane boundary segmentation on the nuScenes validation set. GaussianBeV is trained with 448×800 input image resolution. † means that we trained the model by ourselves. Our method outperforms previous works on lane boundary segmentation and gives competitive results on drivable area segmentation.

4.2. Ablations

Table 4 illustrates the influence of different design choices and training strategies. The Figure 4 shows qualitative results for vehicle and ground surface segmentation. It also gives a visualization of the output of the BeV rasterizer module. To visualize the BeV feature map, we applied Principal Component Analysis (PCA) to the 3D gaussian

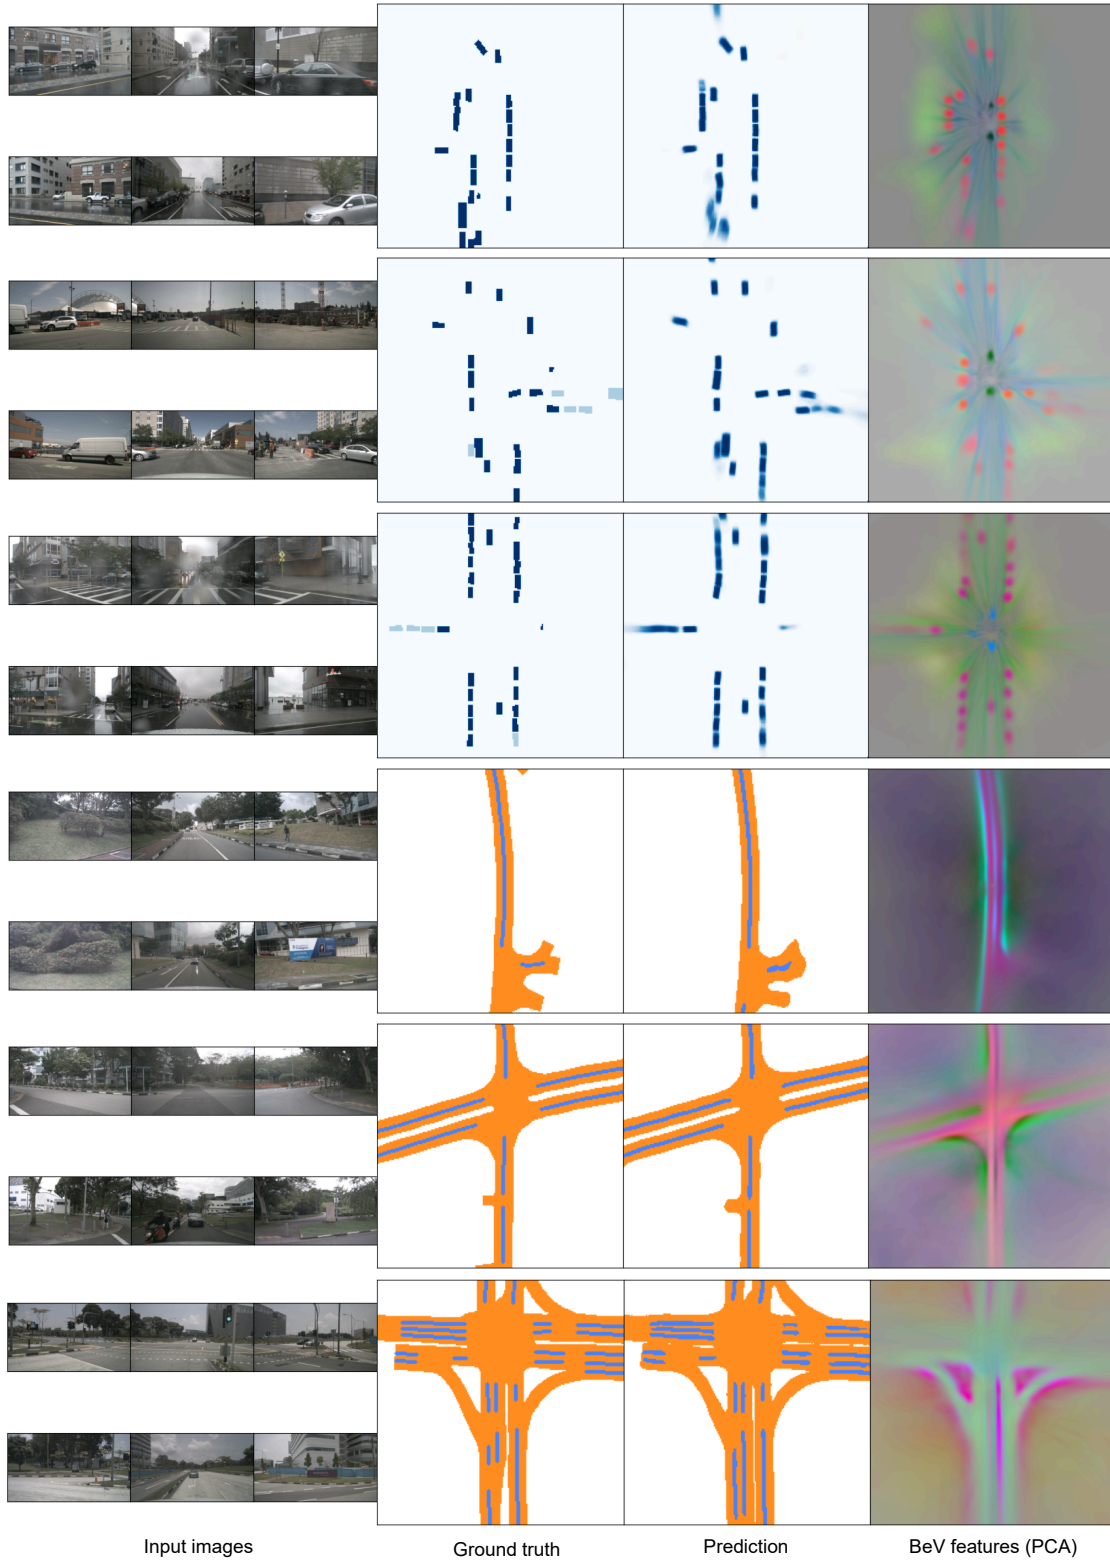


Figure 4. Visualization of predicted vehicle segmentation in the first three rows and drivable area (blue) / lane boundary (orange) segmentation in the last three rows, on the nuScenes validation set. PCA is used to visualize the BeV feature map.

Depth sup	Early sup	BeV backbone	rot	off	IoU
	✓	none	✓	✓	45.4
✓	✓	none	✓	✓	45.9
		LSS [20]	✓	✓	47.0
✓		LSS [20]	✓	✓	46.9
	✓	LSS [20]	✓	✓	46.7
✓	✓	LSS [20]			47.1
✓	✓	LSS [20]	✓		47.3
✓	✓	LSS [20]		✓	47.1
✓	✓	SimpleBeV [7]	✓	✓	47.0
✓	✓	LSS [20]	✓	✓	47.5

Table 4. Influence of BeV backbone, depth and early supervision, rotation and 3D offset prediction. Models are trained using 224×448 input resolution and visibility filtering on the vehicle segmentation task.

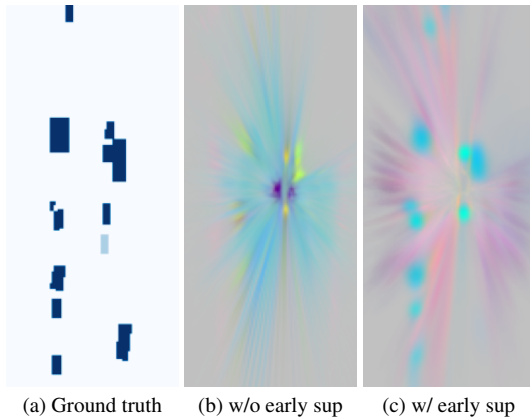


Figure 5. Influence of the early supervision. (a) Vehicle segmentation ground truth. (b) and (c) visualization of the BeV feature map for GaussianBeV trained without and with early supervision, respectively. Early supervision helps to achieve a representation closer to the content of the 3D scene. We observe that gaussians corresponding to the vehicles follow their shapes.

features before BeV rendering. This allows us to reduce the size of the features to render the BeV features in color.

BeV backbone. We conducted experiments to evaluate the influence of the choice of the BeV backbone. In particular, GaussianBeV is trained (1) without using a BeV backbone, by connecting the BeV feature directly to the segmentation heads (row 2 in the Table 4) and (2) with the Simple-BeV backbone [7] to compare performance with the standard LSS backbone [20] (rows 9 and 10 in the Table 4). Remarkably, GaussianBeV trained without BeV backbone is already yielding good results (45.9 IoU that is effectively superior to previous methods). This shows that GaussianBeV is able to provide a geometrically and seman-

tically relevant gaussian representation that can be directly exploited for segmentation. With this observation, the BeV backbone is used to refine the representation using multi-scale information. We observe that the BeV backbone from SimpleBeV [7] achieves very good results (47.0 IoU), even if they are slightly inferior to those obtained with the BeV backbone from LSS [20] (47.5 IoU).

Depth and early supervision. We investigate the influence of auxiliary losses on the gaussian representation as detailed in Sec 3.4. As shown in rows 1 and 2 of the Table 4, the use of the depth loss increases performance when GaussianBeV is trained without a BeV backbone. This shows that this loss regularizes the model to help it outputs a more geometrically consistent gaussian representation. Note that early supervision is used by default if the BeV backbone is not used, as this becomes the main segmentation loss. With regard to rows 3, 4 and 5 of the table using the entire BeV backbone, depth supervision and early supervision do not bring any performance gains when used separately. However, the last line of the table shows that, used together, these two losses improve results. In Figure 5, we provide a visualization of the impact of early supervision on the 3D gaussian representation. We observe that using early supervision results in a more coherent BeV feature map. Indeed, the gaussians better fit the shape of the vehicles.

3D offsets and rotations. Table 4 (rows 6 to 8) shows the impact of learning 3D offsets and rotation. When GaussianBeV is trained without these two heads, it gives slightly poorer results than when it is learned with (last row).

5. Conclusion

In this paper, we introduced GaussianBeV, a novel image-to-BeV transformation method that is the new state-of-the-art on BeV semantic segmentation. Based on an *optimization free* 3D gaussian generator, it transforms each pixel of an image feature map into semantized 3D gaussians. Gaussians are then splattered to obtain the BeV feature map. We have shown that the gaussian representation enables fine 3D modeling thanks to its ability to adapt to the different geometric structures present in the scene. Future work will aim to generate a more compact 3D Gaussian representation to improve computing time. We hope that this initial work will open the door to further research in 3D perception using *optimization free* gaussian splatting representation.

Acknowledgements. This work was carried out as part of the CEA-Valeo joint laboratory and partially financed by BPI France as part of the AVFS project. This work benefited from the FactoryIA supercomputer financially supported by the Ile-de-France Regional Council.

References

- [1] Florent Bartoccioni, Eloi Zablocki, Andrei Bursuc, Patrick Perez, Matthieu Cord, and Karteek Alahari. Lara: Latents and rays for multi-camera bird's-eye-view semantic segmentation. In *Conference on Robot Learning*, 2022. 6
- [2] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nusenes: A multi-modal dataset for autonomous driving. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. 1, 5
- [3] Loick Chambon, Eloi Zablocki, Mickael Chen, Florent Bartoccioni, Patrick Perez, and Matthieu Cord. Pointbev: A sparse approach to bev predictions. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2024. 1, 2, 5, 6
- [4] David Charatan, Sizhe Li, Andrea Tagliasacchi, and Vincent Sitzmann. pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2024. 3
- [5] Shaoheng Fang, Zi Wang, Yiqi Zhong, Junhao Ge, and Siheng Chen. Tbp-former: Learning temporal bird's-eye-view pyramid for joint perception and prediction in vision-centric autonomous driving. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2023. 6
- [6] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J. Brostow. Digging into self-supervised monocular depth prediction. In *Int. Conf. Comput. Vis.*, 2019. 3
- [7] Adam W. Harley, Zhaoyuan Fang, Jie Li, Rares Ambrus, and Katerina Fragkiadaki. Simple-BEV: What really matters for multi-sensor bev perception? In *IEEE International Conference on Robotics and Automation*, 2023. 1, 2, 5, 6, 8
- [8] Anthony Hu, Zak Murez, Nikhil Mohan, Sofia Dudas, Jeffrey Hawke, Vijay Badrinarayanan, Roberto Cipolla, and Alex Kendall. FIERY: future instance prediction in bird's-eye view from surround monocular cameras. In *Int. Conf. Comput. Vis.*, 2021. 1, 2, 5, 6
- [9] Shengchao Hu, Li Chen, Penghao Wu, Hongyang Li, Junchi Yan, and Dacheng Tao. St-p3: End-to-end vision-based autonomous driving via spatial-temporal feature learning. In *Eur. Conf. Comput. Vis.*, 2022. 6
- [10] Junjie Huang, Guan Huang, Zheng Zhu, Ye Yun, and Dalong Du. Bevdet: High-performance multi-camera 3d object detection in bird-eye-view. In *arXiv preprint arXiv:2112.11790*, 2021. 1, 2
- [11] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018. 5
- [12] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. In *ACM Trans. Graph.*, 2023. 1, 2, 5
- [13] Yin hao Li, Zheng Ge, Guanyi Yu, Jinrong Yang, Zengran Wang, Yukang Shi, Jianjian Sun, and Zeming Li. Bevdepth: Acquisition of reliable depth for multi-view 3d object detection. In *Thirty-Seventh AAAI Conference on Artificial Intelligence*, 2022. 1, 2
- [14] Zhan Li, Zhang Chen, Zhong Li, and Yi Xu. Spacetime gaussian feature splatting for real-time dynamic view synthesis. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2024. 3
- [15] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Yu Qiao, and Jifeng Dai. Bevformer: Learning bird's-eye-view representation from multi-camera images via spatiotemporal transformers. In *Eur. Conf. Comput. Vis.*, 2022. 2, 6
- [16] Haisong Liu, Yao Teng, Tao Lu, Haiguang Wang, and Limin Wang. Sparsebev: High-performance sparse 3d object detection from multi-camera videos. In *Int. Conf. Comput. Vis.*, 2023. 2
- [17] Yingfei Liu, Junjie Yan, Fan Jia, Shuailin Li, Qi Gao, Tiancai Wang, Xiangyu Zhang, and Jian Sun. Petrv2: A unified framework for 3d perception from multi-camera images. In *Int. Conf. Comput. Vis.*, 2023. 1, 2, 6
- [18] Hidenobu Matsuki, Riku Murai, Paul H. J. Kelly, and Andrew J. Davison. Gaussian splatting slam. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2024. 3
- [19] Cong Pan, Yonghao He, Junran Peng, Qian Zhang, Wei Sui, and Zhaoxiang Zhang. Baeformer: Bi-directional and early interaction transformers for bird's eye view semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2023. 6
- [20] Jonah Philion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *Eur. Conf. Comput. Vis.*, 2020. 1, 2, 5, 6, 8
- [21] Minghan Qin, Wanhua Li, Jiawei Zhou, Haoqian Wang, and Hanspeter Pfister. Langsplat: 3d language gaussian splatting. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2024. 3, 5
- [22] Jin-Chuan Shi, Miao Wang, Hao-Bin Duan, and Shao-Hua Guan. Language embedded 3d gaussians for open-vocabulary scene understanding. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2024. 3
- [23] Stanislaw Szymanowicz, Christian Rupprecht, and Andrea Vedaldi. Splatter image: Ultra-fast single-view 3d reconstruction. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2024. 3
- [24] Mingxing Tan and Quoc Le. EfficientNet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning, ICML*, 2019. 5
- [25] Shihao Wang, Yingfei Liu, Tiancai Wang, Ying Li, and Xiangyu Zhang. Exploring object-centric temporal modeling for efficient multi-view 3d object detection. In *Int. Conf. Comput. Vis.*, 2023. 2
- [26] Enze Xie, Zhiding Yu, Daquan Zhou, Jonah Philion, Anima Anandkumar, Sanja Fidler, Ping Luo, and Jose M Alvarez. M²bev: Multi-camera joint 3d detection and segmentation with unified birds-eye view representation. In *arXiv preprint arXiv:2204.05088*, 2022. 6
- [27] Chi Yan, Delin Qu, Dan Xu, Bin Zhao, Zhigang Wang, Dong Wang, and Xuelong Li. Gs-slam: Dense visual slam with 3d gaussian splatting. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2024. 3
- [28] Yunzhi Yan, Haotong Lin, Chenxu Zhou, Weijie Wang, Haiyang Sun, Kun Zhan, Xianpeng Lang, Xiaowei Zhou, and Sida Peng. Street gaussians for modeling dynamic urban scenes. In *Eur. Conf. Comput. Vis.*, 2024. 3

- [29] Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything: Unleashing the power of large-scale unlabeled data. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2024. 3
- [30] Mingqiao Ye, Martin Danelljan, Fisher Yu, and Lei Ke. Gaussian grouping: Segment and edit anything in 3d scenes. In *Eur. Conf. Comput. Vis.*, 2024. 3
- [31] Wei Yin, Chi Zhang, Zhipeng Cai Hao Chen, Gang Yu, Kaixuan Wang, Xiaozhi Chen, and Chunhua Shen. Metric3d: Towards zero-shot metric 3d prediction from a single image. In *Int. Conf. Comput. Vis.*, 2023. 4
- [32] Brady Zhou and Philipp Krähenbühl. Cross-view transformers for real-time map-view semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2022. 1, 2, 6
- [33] Hongyu Zhou, Zheng Ge, Zeming Li, and Xiangyu Zhang. Matrixvt: Efficient multi-camera to bev transformation for 3d perception. In *Int. Conf. Comput. Vis.*, 2023. 6
- [34] Shijie Zhou, Haoran Chang, Sicheng Jiang, Zhiwen Fan, Zehao Zhu, Dejia Xu, Pradyumna Chari, Suyu You, Zhangyang Wang, and Achuta Kadambi. Feature 3dgs: Supercharging 3d gaussian splatting to enable distilled feature fields. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2024. 3