

GAUSSIANFUSION: UNIFIED 3D GAUSSIAN REPRESENTATION FOR MULTI-MODAL FUSION PERCEPTION

005 **Anonymous authors**

006 Paper under double-blind review

ABSTRACT

011 The bird’s-eye view (BEV) representation enables multi-sensor features to be
 012 fused within a unified space, serving as the primary approach for achieving com-
 013 prehensive 3D perception. However, the discrete grid representation of BEV leads
 014 to significant detail loss and limits feature alignment and cross-modal information
 015 interaction in multimodal fusion perception. In this work, we break from the
 016 conventional BEV paradigm and propose a new universal framework for multi-
 017 modal fusion based on 3D Gaussian representation. This approach naturally uni-
 018 fies multi-modal features within a shared and continuous 3D Gaussian space, ef-
 019 fectively preserving edge and fine texture details. To achieve this, we design a
 020 novel forward-projection-based multi-modal Gaussian initialization module and a
 021 shared cross-modal Gaussian encoder that iteratively updates Gaussian properties
 022 based on an attention mechanism. GaussianFusion is inherently a task-agnostic
 023 model, with its unified Gaussian representation naturally supporting various 3D
 024 perception tasks. Extensive experiments demonstrate the generality and robust-
 025 ness of GaussianFusion. On the nuScenes dataset, it outperforms the 3D object
 026 detection baseline BEVFusion by 2.6 NDS. Its variant surpasses GaussFormer on
 027 3D semantic occupancy with 1.55 mIoU improvement while using only 30% of
 028 the Gaussians and achieving a 450% speedup.

1 INTRODUCTION

030 Fusing complementary signals captured by different sensors is essential for autonomous driving
 031 perception systems. Leveraging the distinct characteristics of each sensor helps reduce prediction
 032 uncertainty, leading to more accurate and robust perception outcomes [Liu et al. \(2023b\)](#); [Bai et al. \(2022\)](#);
 033 [Yan et al. \(2023\)](#). Since different sensors present data in varying formats, such as cameras
 034 providing perspective semantic data and Lidar capturing 3D spatial information, multi-modal fusion
 035 faces significant challenges due to these view discrepancies. To address this, some methods [Vora et al. \(2020\)](#);
 036 [Bai et al. \(2022\)](#); [Li et al. \(2024\)](#); [Wang et al. \(2024b\)](#) achieve multi-modal 3D object
 037 detection through point-level fusion. However, point-level fusion strategies are generally unsuitable
 038 for semantics-oriented 3D perception tasks like 3D semantic occupancy prediction. Consequently,
 039 recent approaches aim to construct unified representations for multi-modal feature fusion.

040 Recently, the shared Bird’s Eye View
 041 (BEV) space has emerged as a promis-
 042 ing direction for fusing cross-modal
 043 features to enable task-agnostic learn-
 044 ing. Several existing fusion methods
 045 [Liu et al. \(2023b\)](#); [Wang et al. \(2023a\)](#),
 046 such as BEVFusion [Liu et al. \(2023b\)](#),
 047 integrate multimodal information via
 048 CNNs and feature concatenation, while
 049 MetaBEV [Ge et al. \(2023\)](#) mitigates
 050 cross-modal feature misalignment by in-
 051 troducing meta-BEV queries. However, despite its widespread adoption in multimodal 3D percep-
 052 tion, the BEV representation inherently suffers from limitations in information expression.

Table 1: Impact of BEV size on model performance

Method	BEV size	Grid size	Memory↓	NDS↑
BEVFusion	100×100	1.008m	3228 M	70.5
	200×200	0.504m	5140 M	71.4
	400×400	0.252m	20560 M	72.7
GaussianFusion	100×100	1.008m	3576 M	73.1
	200×200	0.504m	5418 M	74.0
	400×400	0.252m	6151 M	74.4

054 BEV directly discretizes and quantizes data, leading to inevitable information loss. During 055 feature extraction, perception data are projected onto a fixed-resolution BEV grid, which 056 compresses spatial information. This issue 057 becomes particularly severe when the BEV 058 resolution is low, as it directly impacts model 059 performance by failing to adequately preserve 060 fine-grained scene structures. While increasing 061 the BEV resolution will bring unacceptable 062 computational overhead, as shown in Table 063 1. Additionally, BEV fusion strategies often rely 064 on simple feature concatenation or weighted 065 summation, which are insufficient for effective 066 cross-modal feature interaction and alignment, 067 ultimately leading to suboptimal fusion performance, 068 as illustrated in Fig. 1(a). 069

070 To address these challenges, we introduce a 071 fusion approach based on 3D Gaussian Splatting 072 (3DGS) [Kerbl et al. (2023)] to achieve more 073 fine-grained information modeling and more 074 natural multimodal alignment. As shown in 075 Fig. 1(b), 3DGS employs continuous 076 Gaussian distributions to represent the scene, 077 preserving rich geometric and semantic 078 information in the Gaussian stage and preventing 079 the early quantization-induced information 080 loss seen in BEV-based methods. Unlike direct BEV 081 quantization, 3DGS aggregates information before 082 its final projection onto the BEV grid, allowing 083 cross-modal features to interact at a higher-dimensional 084 level and capturing finer spatial structures prior 085 to quantization, Table 1 shows the effectiveness of 086 this strategy. Moreover, the covariance matrices 087 of Gaussians enable adaptive modeling of 088 uncertainty, enhancing the representation of 089 object shapes and boundaries. 090

091 Specifically, inspired by [Phlion & Fidler (2020)], we propose a forward projection Gaussian 092 initialization strategy to better initialize camera Gaussian 093 representations in 3D space rather than using 094 random initialization [H et al. (2024)]. To further 095 achieve continuous alignment and cross-modal 096 feature enhancement, we construct a shared Gaussian 097 encoder. The shared Gaussian encoder supports 098 cross-feature learning of 3D Gaussian features 099 from different modalities, where the covariance 100 matrix of each 3D Gaussian adaptively captures 101 feature differences between modalities and 102 iteratively updates the Gaussian parameters. Camera 103 and LiDAR Gaussians are naturally fused via a 104 Gaussian mixture model, and a high-performance 105 Gaussian-to-voxel fusion module aggregates 106 surrounding Gaussians to generate voxel features, 107 enabling task-agnostic 3D perception. We conduct 108 extensive experiments on BEV object detection and 109 3D occupancy prediction tasks to validate the 110 generality of the GaussianFusion. Main contributions 111 are as follows:

- We propose the first unified 3D Gaussian representation multi-modal fusion framework, where cross-view and cross-modal Gaussian representations are naturally aggregated through the Gaussian mixture model.
- A progressive update strategy is designed to optimize the multi-modal 3D Gaussian properties iteratively.
- The shared 3D Gaussian encoder enables alignment and complementary enhancement of cross-modal features, allowing Gaussian representations from both modalities to achieve consistent uncertainty within a unified space.
- Our GaussianFusion achieves state-of-the-art benchmarks in task-agnostic methods on various 3D perception tasks within the nuScenes dataset.

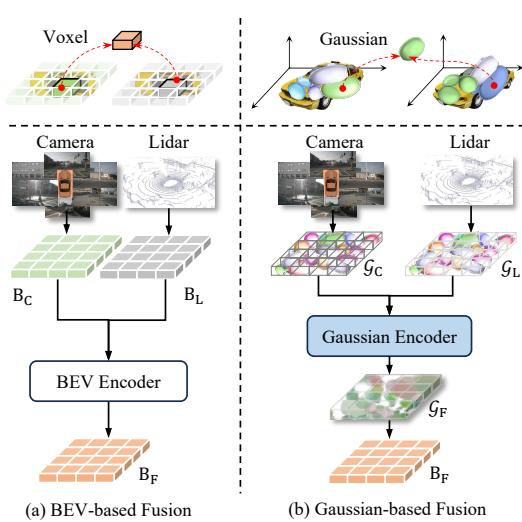


Figure 1: Comparison of the discrete BEV representation fusion paradigm [Liu et al. (2023b)] and our proposed continuous Gaussian representation fusion paradigm. B, \mathcal{G} , C, L, and F denote BEV, Gaussian, Camera, Lidar, and Fusion.

Unlike direct BEV quantization, 3DGS aggregates information before its final projection onto the BEV grid, allowing cross-modal features to interact at a higher-dimensional level and capturing finer spatial structures prior to quantization, Table 1 shows the effectiveness of this strategy. Moreover, the covariance matrices of Gaussians enable adaptive modeling of uncertainty, enhancing the representation of object shapes and boundaries.

Specifically, inspired by [Phlion & Fidler (2020)], we propose a forward projection Gaussian initialization strategy to better initialize camera Gaussian representations in 3D space rather than using random initialization [H et al. (2024)]. To further achieve continuous alignment and cross-modal feature enhancement, we construct a shared Gaussian encoder. The shared Gaussian encoder supports cross-feature learning of 3D Gaussian features from different modalities, where the covariance matrix of each 3D Gaussian adaptively captures feature differences between modalities and iteratively updates the Gaussian parameters. Camera and LiDAR Gaussians are naturally fused via a Gaussian mixture model, and a high-performance Gaussian-to-voxel fusion module aggregates surrounding Gaussians to generate voxel features, enabling task-agnostic 3D perception. We conduct extensive experiments on BEV object detection and 3D occupancy prediction tasks to validate the generality of the GaussianFusion. Main contributions are as follows:

- We propose the first unified 3D Gaussian representation multi-modal fusion framework, where cross-view and cross-modal Gaussian representations are naturally aggregated through the Gaussian mixture model.
- A progressive update strategy is designed to optimize the multi-modal 3D Gaussian properties iteratively.
- The shared 3D Gaussian encoder enables alignment and complementary enhancement of cross-modal features, allowing Gaussian representations from both modalities to achieve consistent uncertainty within a unified space.
- Our GaussianFusion achieves state-of-the-art benchmarks in task-agnostic methods on various 3D perception tasks within the nuScenes dataset.

108
109

2 RELATED WORK

110
111

2.1 MULTIMODAL 3D PERCEPTION

112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
We categorize current multimodal fusion methods into object-centric methods and dense BEV methods. Object-centric methods Vora et al. (2020); Chen et al. (2023); Zhou & T (2023); Yin et al. (2024); Li et al. (2024); Wang et al. (2024b) are specifically designed for tasks such as 3D object detection or tracking. Advanced object-centric Li et al. (2024); Wang et al. (2024b) methods typically use 2D detection results on camera to enhance multi-modal fusion 3D detection. Additionally, some works Yang et al. (2022); Yan et al. (2023) use query-based 3D detection decoders to learn features from perspective images and Lidar BEV features directly. However, these object-centric methods cannot easily generalize to dense semantic tasks such as BEV map segmentation and 3D occupancy prediction. Dense BEV methods Li et al. (2022a); Liu et al. (2023b); Liang et al. (2022); Chen et al. (2022); Zhao et al. (2024b); Ge et al. (2023); Jiao et al. (2023); Wang et al. (2023a) naturally adapt to various tasks. Both BEVFusion Liu et al. (2023b) and UniTR Wang et al. (2023a) achieve multi-modal BEV fusion perception through CNN and feature concatenation in the BEV space. In addition, MetaBEV Ge et al. (2023) proposed a learnable cross-attention mechanism to generate unified BEV features. BEV or 3D voxel also provides a unified representation for 3D occupancy prediction. Some methods design fusion modules Pan et al. (2024); Wang et al. (2023c); Ming et al. (2024) based on voxel, such as adaptive fusion Wang et al. (2023c), etc. for multi-modal 3D occupancy prediction. There are also many camera-only methods Zhao et al. (2024a); Wang et al. (2024a); Li et al. (2023); Lu et al. (2023); Tian et al. (2024); Li et al. (2025a); Ma et al. (2024b); Cao et al. (2024); Liu et al. (2024); Ma et al. (2024a) for 3D semantic occupancy prediction based on voxel representation. However, discrete voxel representations may result in significant detail loss and hinder effective multimodal complementary fusion.

2.2 3D GAUSSIAN SPLATTING

132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
3D Gaussian Splatting (3DGS) Kerbl et al. (2023) combines the advantages of implicit neural radiance fields Mildenhall et al. (2021) and voxel-based explicit radiance fields Fridovich-Keil et al. (2022); Müller et al. (2022) and is widely applied in 3D reconstruction of 2D image. 3DGS uses a set of Gaussian functions to capture the geometric shapes and semantic of different objects or regions within a scene, effectively representing the scene. Based on this, some works Ye et al. (2025); Hu et al. (2024) leverage 3DGS’s multi-view synthesis capabilities to achieve 3D scene segmentation.155
156
157
158
159
160
161
Recent studies Gan et al. (2024); Chabot et al. (2024); H et al. (2024); Zuo et al. (2024); Liu et al. (2025) have applied 3DGS to vision-only 3D semantic occupancy prediction, BEV segmentation, and end-to-end autonomous driving. However, these methods rely on randomly initialized Gaussians. For example, GaussianFormer H et al. (2024) randomly initializes and re-predicts Gaussian parameters in each iteration, making fine-tuning difficult and limiting accurate scene representation. Moreover, these methods do not fully exploit the advantages of Gaussian Mixture Models (GMM) for seamless multimodal Gaussian fusion. In contrast, we propose a forward-projection-based 3DGS parameter initialization and a shared optimization model, leveraging GMM to fuse multi-view camera and LiDAR features within a shared space, ultimately enabling dense semantic understanding and object-centric multitask perception.

3 METHODS

3.1 OVERALL ARCHITECTURE

155
156
157
158
159
160
161
The overall pipeline of GaussianFusion is illustrated in Fig. 2, with the goal of fusing multimodal features through 3D Gaussian representations, which naturally preserve both geometric and semantic information. We first initialize separate 3D Gaussian representations for camera and Lidar, denoted as $\mathcal{G}_c \sim Q_c$ and $\mathcal{G}_L \sim Q_L$, within a unified space. Then, the multimodal Gaussian sets are processed through a shared Gaussian Encoder, enabling the integration of semantic and geometric information from both modalities. Finally, the learned 3D Gaussian sets $\hat{\mathcal{G}}_c \sim \hat{Q}_c$ and $\hat{\mathcal{G}}_L \sim \hat{Q}_L$ are fused within the unified Gaussian space and fed into task-specific heads to perform 3D perception. Fig. 2 shows the single-task training setup, while multi-task joint training is detailed in the appendix.

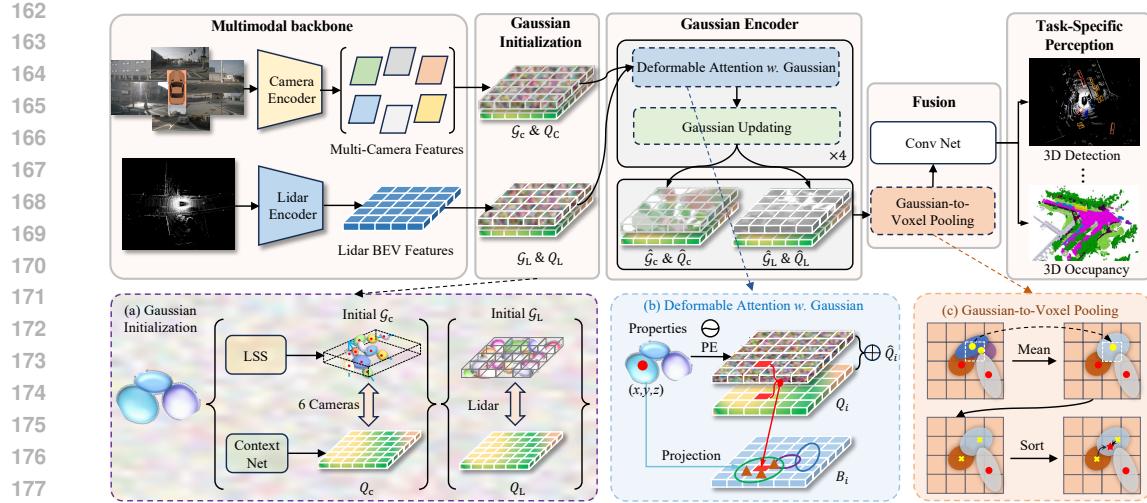


Figure 2: Overview of the GaussianFusion framework. Initial Gaussians are refined by a shared encoder and fused in Gaussian space, followed by task-specific heads for 3D perception.

3.2 GAUSSIAN INITIALIZATION

BEVFusion [Liu et al. (2023b)] projects multimodal features into discrete BEV space for fusion. In contrast, inspired by the Gaussian mixture model paradigm of 3DGS [Kerbl et al. (2023)] for modeling scene geometry and semantics, we utilize two Gaussian sets, \mathcal{G}_c and \mathcal{G}_L , within a shared space to represent surround-view camera and Lidar information, respectively, to achieve a seamless multimodal fusion.

Camera Gaussian Initialization with Forward Projection. The properties of every single 3D Gaussian function are defined by a mean $\mu \in \mathbb{R}^3$, scale $\mathbf{s} \in \mathbb{R}^3$, and rotation vectors $\mathbf{r} \in \mathbb{R}^4$. Given N surround camera view, each camera's view can be represented by a set of 3D Gaussian distributions $G_{c,i} \in \mathbb{R}^{D_g \times D_c \times H_c \times W_c} | i = 1, 2, \dots, N$, where $D_g = \mu + \mathbf{s} + \mathbf{r}$, D_c denotes the number of discrete depths for camera, as shown in Fig. 2(a), a Gaussian function is assigned to a depth point. It is worth noting that this is completely different from GausianFormer [H et al. (2024)], which randomly initializes a set of Gaussians in space, which makes model optimization more difficult. Specifically, inspired by [Huang & Huang (2022b)a]; [Liu et al. (2023b)]; [Philion & Fidler (2020)], given surround camera input features $F_{c,i} \in \mathbb{R}^{C \times H_c \times W_c}, i = 1, 2, \dots, N$, where C , H_c and W_c represent the channel, height, and width of the camera features. $F_{c,i}$ is fed to LSS [Philion & Fidler (2020)] to obtain the depth distribution $D_i \in \mathbb{R}^{D_c \times H_c \times W_c}, i = 1, 2, \dots, N$, D_i are then used as the initial mean μ of the Gaussian, which is the location of every single Gaussian center. For \mathbf{s} and \mathbf{r} , we initialize them randomly. And $\mathcal{G}_c = \{G_{c,i}\}, \mathcal{G}_c \in \mathbb{R}^{N \times D_g \times D_c \times H_c \times W_c}$.

The image features $F_{c,i}$ are processed through a context network composed of multiple convolutional layers to obtain the semantic features $F'_{c,i}$. Next, an inner product is computed between D_i and $F'_{c,i}$ to derive the features at each depth point in 3D space, denoted as initial query features $Q_c = \{Q_{c,i} \in \mathbb{R}^{C \times D_c \times H_c \times W_c} | i = 1, 2, \dots, N\}$. Then, Q_c are associated with Gaussian ($\mathcal{G}_c \sim Q_c$). For a given Gaussian set ($g_c \sim q_c, g_c \in \mathcal{G}_c, q_c \in Q_c$), the feature at a point $\mathbf{p} = (x, y, z)$ within its elliptical space are:

$$g_c(\mathbf{p}; \mu, \mathbf{s}, \mathbf{r}) = \exp\left(-\frac{1}{2}(\mathbf{p} - \mu)^T \Sigma^{-1}(\mathbf{p} - \mu)\right) q_c, \quad (1)$$

where $\Sigma = \mathbf{R} \mathbf{S}^T \mathbf{R}^T$, $\mathbf{S} = \text{diag}(\mathbf{s})$, and $\mathbf{R} = \text{q2r}(\mathbf{r})$. $\Sigma, \text{diag}(\cdot)$, and $\text{q2r}(\cdot)$ represent the covariance matrix, the function that constructs a diagonal matrix from a vector, and the function that transforms a quaternion into a rotation matrix, respectively.

Lidar Gaussian initialization. Lidar's BEV space naturally provides an initialization for the Gaussian mean μ . And, \mathbf{s} and \mathbf{r} are initialized randomly. The Gaussian initialization for Lidar is formulated as $\mathcal{G}_L \in \mathbb{R}^{C \times H_L \times W_L}$. Since directly extracting information from the massive raw lidar

216 point clouds to construct 3D Gaussian representations is both difficult and computationally intensive,
 217 grid-based representations offer an effective approach to alleviate these challenges. For query,
 218 given the Lidar features $B_L \in \mathbb{R}^{C \times H_L \times W_L}$, where C , H_L , and W_L represent the channel, height,
 219 and width of the BEV features, respectively. Then, we fed the BEV feature B_L into a multilayer
 220 perceptron (MLP) to obtain Lidar query Q_L , which is associated with the initial query features for
 221 each Gaussian function ($\mathcal{G}_L \sim Q_L$).
 222

223 3.3 GAUSSIAN ENCODER

225 The 3D Gaussian distribution effectively represents the scene, and we have designed a Gaussian
 226 encoder to optimize both the properties and query features. This encoder includes a deformable
 227 attention with Gaussian module and a Gaussian Updating module. The Gaussian encoder is stacked
 228 multiple times to update the Gaussian properties in an iterative refinement paradigm. Additionally,
 229 to better fuse multi-modal information, we employ a shared Gaussian encoder to simultaneously pro-
 230 cess the Gaussian distributions \mathcal{G}_c and \mathcal{G}_L , as the two modalities are ultimately intended to converge
 231 towards similar Gaussian distributions. Specifically, we merge \mathcal{G}_c and \mathcal{G}_L into the batch dimension.
 232

233 **Deformable Attention with Gaussian.** As shown in Fig. 2(b), after obtaining the Gaussian distri-
 234 bution sets and corresponding query features ($\mathcal{G}_c \sim Q_c$, $\mathcal{G}_L \sim Q_L$), we first encode $\mathcal{G}_i \sim Q_i$, $i = c, L$,
 235 into a new query \hat{Q} . Specifically, to capture the position and geometric information of the Gaussian
 236 functions within the feature maps, we employ an MLP to encode the Gaussian properties. The new
 237 query \hat{Q} is then obtained by adding the encoded properties P_q to the original query Q_i :
 238

$$\hat{Q}_i = \text{MLP}(\mathcal{G}) + Q_i, i = c, L, \quad (2)$$

240 where the MLP as the position embedding (PE) transforms the dimensions from \mathbb{R}^{D_g} to \mathbb{R}^C . We
 241 obtain sets $\mathcal{G}_i \sim \hat{Q}_i$.
 242

243 Furthermore, we adopt deformable
 244 attention with Gaussian [Zhu et al.
 245 (2020); H et al. (2024)] to extract fea-
 246 tures, as shown in Fig. 3. Vanilla de-
 247 formable attention (Fig. 3(b)) initial-
 248 izes sampling locations with an ap-
 249 proximately “square/kernel-like” re-
 250 gion and learns offsets to cover the
 251 regions of interest. However, this initial
 252 distribution lacks inherent geometric
 253 priors about object shape. In con-
 254 trast, our deformable attention with
 255 Gaussian (Fig. 3(a)) directly inherits
 256 and leverages the shape properties
 257 of Gaussians: by projecting the 3D
 258 Gaussian distributions onto the BEV
 259 feature map (projection in Figure 2b),
 260 we obtain a prior sampling distribu-
 261 tion that encodes orientation, scale,
 262 and covariance structure. In other
 263 words, the initial sampling points are
 264 not uniformly spaced on a grid, but
 265 instead follow a Gaussian distribution aligned with the underlying object geometry—such as aspect
 266 ratio, orientation, and spatial uncertainty. This Gaussian prior enables better alignment of cross-
 267 modal features to the “likely object extent,” thereby enhancing fusion effectiveness—a capability
 268 absent in conventional square-shaped initialization.
 269

Evidently, the properties of the Gaussian functions effectively describe the shape of the potential objects or regions. For each Gaussian function g , we calculate a set of offsets $\Delta\mu = (\Delta x, \Delta y, \Delta z)$ based on the covariance matrix. These offsets, combined with the mean μ , yield the corresponding reference points $\mu + \Delta\mu$. We then project the 3D reference points onto the BEV feature map, where

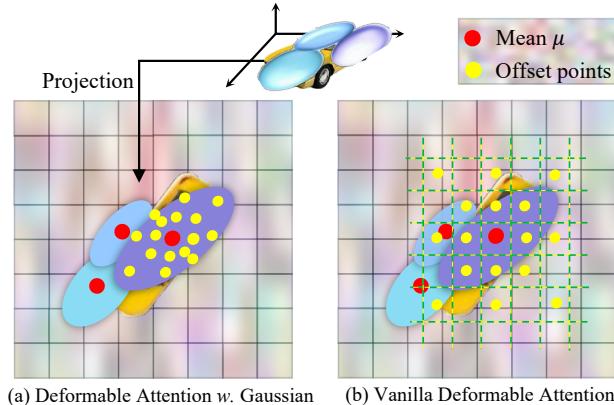


Figure 3: Comparison of the vanilla deformable attention [Zhu et al. (2020)] and our proposed deformable attention with Gaussian.

270 each Gaussian query $q_i \in Q_i$ is updated through deformable attention, expressed as:
 271

$$272 \quad \text{DeformAtt}(q_i, B_i) = \sum_{k=1}^K A_k \cdot W_k B_i(\boldsymbol{\mu} + \Delta\boldsymbol{\mu}), \quad (3)$$

273 where B_i is the camera BEV feature B_c or Lidar BEV feature B_L , W_k is the weights obtained
 274 by linear layer, A_k is the attention weights and $A_k \in [0, 1]$, K is the number of sampling points,
 275 $B(\boldsymbol{\mu} + \Delta\boldsymbol{\mu})$ are the sampled features.
 276

277 **Gaussian Updating.** To update the Gaussian properties, we propose an iterative optimization strategy
 278 of predicting offsets instead of predicting a set of new Gaussian distributions as adopted in
 279 GaussianFormer [H et al. (2024)]. In Lidar and camera fusion perception, incrementally updating the
 280 Gaussian parameters allows for better handling of the discrepancies caused by different modalities
 281 when perceiving the same object. This approach is particularly effective in handling fusion uncer-
 282 tainties caused by such as signal attenuation, depth prediction uncertainties, or multi-modal signal
 283 discrepancies, as demonstrated by the ablation experiments. Incremental updates across layers allow
 284 the model to gradually reduce the disparity between modalities, improving fusion accuracy. Specif-
 285 ically, by predicting the offsets $\Delta\boldsymbol{\mu}$, $\Delta\mathbf{s}$, and $\Delta\mathbf{r}$ for the Gaussian mean $\boldsymbol{\mu}$, scale \mathbf{s} , and rotation \mathbf{r}
 286 using an MLP, we refine the Gaussian distribution without having to predict a completely new set of
 287 properties. The updated Gaussian \hat{g}_i as follows:
 288

$$289 \quad \hat{g}_i = \text{MLP}(\hat{Q}) + \mathcal{G}_i = (\Delta\boldsymbol{\mu} + \boldsymbol{\mu}, \Delta\mathbf{s} + \mathbf{s}, \Delta\mathbf{r} + \mathbf{r}). \quad (4)$$

290 3.4 MULTI-SENSOR FUSION

291 By the Gaussian Encoder model, we obtain the multimodal 3D Gaussian representations $\hat{\mathcal{G}}_c$ and $\hat{\mathcal{G}}_L$
 292 within the shared 3D space, respectively. We merge $\hat{\mathcal{G}}_c$ and $\hat{\mathcal{G}}_L$ into a unified set $\hat{\mathcal{G}}$, and we can easily
 293 fuse them. Although the 3D Gaussian distributions can effectively represent the scene, to handle the
 294 irregular distribution of Gaussian points, we need to voxelize these Gaussian distributions to achieve
 295 task-independent 3D perception.
 296

297 Specifically, given the unified Gaussian sets $\hat{\mathcal{G}} \sim \hat{Q}$, we divide the Gaussian space into a voxel grid
 298 $H \times W$. For a non-empty voxel V that contains M Gaussian means $\boldsymbol{\mu}$, the Gaussian set for that voxel
 299 is $V = \{(\hat{g}_1, \hat{g}_2, \dots, \hat{g}_M) \sim (\hat{q}_1, \hat{q}_2, \dots, \hat{q}_M)\}$. To ensure real-time performance and the receptive
 300 field of the voxel, we use MeanVFE [Zhou & Tuzel (2018)] to downsample the Gaussian within
 301 the voxel, as illustrated in Fig. 2c). After Gaussian pooling, each voxel contains one Gaussian
 302 distribution ($\hat{g} \sim \hat{q}$):
 303

$$304 \quad \hat{g} = \frac{1}{M} [\sum \boldsymbol{\mu}_m, \sum \mathbf{s}_m, \sum \mathbf{r}_m], \hat{q} = \frac{1}{M} \sum \hat{q}_m. \quad (5)$$

305 Furthermore, the Gaussian mixture model [Kerbl et al. (2023)] can naturally aggregate multiple Gaus-
 306 sian distributions into a finer-grained distribution, unifying multi-modal Gaussian representations
 307 and elegantly capturing the complexity of autonomous driving scenes. Thus, if the total number
 308 of Gaussian distributions covering point \mathbf{p} in the entire scene is J , the feature $f(\mathbf{p})$ at point \mathbf{p} is
 309 composed of the cumulative contributions of each individual Gaussian:
 310

$$311 \quad f(\mathbf{p}) = \sum_{i=1}^J \hat{g}_i(\mathbf{p}; \boldsymbol{\mu}, \mathbf{s}, \mathbf{r}) \hat{q}_i. \quad (6)$$

312 Since each voxel may be associated with multiple 3D Gaussian distributions, following the strategy
 313 in [H et al. (2024)], we calculate the neighborhood radius based on the scale property of each Gaus-
 314 sian. The indices of the Gaussians and the voxels within their neighborhood are paired as tuples
 315 and appended to a list. This list is then sorted by voxel indices, determining which 3D Gaussians
 316 each voxel should focus on. Furthermore, for each voxel using 6, we can get the fused feature B_F .
 317 Finally, a simple convolutional network is used to further optimize B_F .
 318

319 3.5 PERCEPTION TASK SETUP

320 Without loss of generality, we follow BEVFusion [Liu et al. (2023b)], GaussianFusion can be applied
 321 to most 3D perception tasks based on B_F . We evaluate the performance of GaussianFusion on 3D
 322

324
 325 Table 2: Comparisons with state-of-the-art 3D object detection methods on nuScenes dataset. C
 326 denote Camera, L denote Lidar. All methods construct BEV-based feature maps instead of object-
 327 centric fusion based on proposals, which means these methods can also be naturally used for seman-
 328 tic tasks. UniTR uses a unified backbone for both the camera and Lidar.

329 Methods	330 Modality	331 Resolution	332 Backbone		333 validation set		334 test set	
			335 Camera	336 Lidar	337 NDS	338 mAP	339 NDS	340 mAP
333 BEVFormer [Li et al. (2022b)]	334 C	335 1600×900	336 ResNet-101	337 -	338 51.7	339 41.6	340 56.9	341 48.1
333 PETRv2 [Liu et al. (2023a)]	334 C	335 1600×640	336 VoV-99	337 -	338 -	339 -	340 59.1	341 50.8
333 FB-BEV [Li et al. (2023)]	334 C	335 1600×640	336 VoV-99	337 -	338 -	339 -	340 62.4	341 53.7
333 AutoAlignV2 [Chen et al. (2022)]	334 C+L	335 1600×640	336 CSPNet	337 VoxelNet	338 71.2	339 67.1	340 72.4	341 68.4
333 BEVFusion(M) [Liu et al. (2023b)]	334 C+L	335 704×256	336 Swin-T	337 VoxelNet	338 71.4	339 68.5	340 72.9	341 70.2
333 MetaBEV [Ge et al. (2023)]	334 C+L	335 704×256	336 Swin-T	337 VoxelNet	338 71.5	339 68.0	340 -	341 -
333 MSMDFusion [Jiao et al. (2023)]	334 C+L	335 800×448	336 ResNet-50	337 VoxelNet	338 72.1	339 69.3	340 74.0	341 71.5
333 FusionFormer-S [Hu et al. (2023a)]	334 C+L	335 1600×640	336 VoV-99	337 VoxelNet	338 73.2	339 70.0	340 -	341 -
333 EA-LSS [Hu et al. (2023b)]	334 C+L	335 704×256	336 Swin-T	337 VoxelNet	338 73.1	339 71.2	340 74.4	341 72.2
333 UniTR [Wang et al. (2023a)]	334 C+L	335 704×256	336 -	337 -	338 73.3	339 70.5	340 74.5	341 70.9
333 GaussianFusion(Ours)		335 C+L	336 704×256	337 Swin-T	338 VoxelNet	339 74.0	340 71.7	341 74.9
								342 72.4

343 object detection and 3D semantic occupancy prediction tasks. We adopt the same Transformer-
 344 based detection head as BEVFusion [Bai et al. (2022); Liu et al. (2023b)] and the occupancy head
 345 consistent with BEVDet [Huang & Huang (2022a)]. We also constructed a camera-only version,
 346 GaussianFusion-C, containing only the camera branch shown in Fig. 2.

4 EXPERIMENTS

4.1 DATASET

350 The nuScenes dataset [Caesar et al. (2020)] provides annotation data for tasks such as semantic
 351 segmentation, object detection, and 3D occupancy (Occ) prediction. It is a large-scale multimodal
 352 dataset officially split into 700/150/150 scenes for training, validation, and testing, respectively.
 353 Each scene includes annotated Lidar point cloud data captured by a 32-beam scanner, along with
 354 6 perspective camera views, offering comprehensive 360-degree coverage at each timestamp. We
 355 evaluate our method on the 3D object detection and Occ. In our task, we down-sample the input
 356 camera images to 704×256 and voxelize the point cloud to 0.075m for detection, following BEV-
 357 Fusion [Liu et al. (2023b)] and UniTR [Wang et al. (2023a)]. For 3D object detection, the perception
 358 range of the point cloud is set to $[-51.2m, 51.2m]$ along the X and Y axes, and $[-5m, 3m]$ along
 359 the Z axis. For 3D Occ, we evaluate within the region of $[-50m, 50m] \times [-50m, 50m]$ around the
 360 ego vehicle, following [Liu et al. (2023b); Wang et al. (2023a); Wei et al. (2023)].

4.2 IMPLEMENTATION DETAILS

363 We adopt VoxelNet [Zhou & Tuzel (2018)] and Swin-T [Liu et al. (2021)] as the Lidar and camera
 364 backbone to extract the multimodal features following BEVFusion [Liu et al. (2023b)]. The dimen-
 365 sions of Lidar features, image features, and 3D Gaussian query features are all set to 128. Depth of
 366 image $D_c=41$. Image feature $F_{c,i}$ dimensions: $8 \times 22 \times 6$. We set Gaussian Encoder blocks to 4, see
 367 the Appendix for experiments. The BEV size $H \times W$ is set to 200×200 .

368 During training, we follow BEVFusion to adapt the aligned multimodal data augmentation strategy
 369 and the class-balanced sampling strategy from CBGS [Zhu et al. (2019)]. GaussianFusion is trained
 370 on 8 NVIDIA A800 GPUs. We use AdamW [Loshchilov & Hutter (2017)] optimizer with a weight
 371 decay 0.01. We adopt the one-cycle learning rate policy [Smith (2017)] with a maximum learning
 372 rate of $2e^{-4}$. Both BEV object detection and 3D semantic occupancy prediction are trained for 20
 373 epochs, following the same settings as BEVFusion and GaussianFormer [H et al. (2024)], respectively.

4.3 3D OBJECT DETECTION

375 **Setting.** We utilize the official evaluation metric nuScenes Detection Score (NDS) and mean Average
 376 Precision (mAP) for 3D detection.

378 **Results.** To highlight the effect of Gaussian representation, we only compare the BEV-based
 379 method. As shown in Table 2, GaussianFusion achieves SOTA results compared to previous discrete
 380 BEV representation multimodal fusion methods [Liu et al. (2023b); Ge et al. (2023); Wang et al.
 381 (2023a); Hu et al. (2023b)] on nuScenes dataset, achieving 74.0 NDS and 71.7 mAP on the val split.
 382 Specifically, compared with BEVFusion [Liu et al. (2023b)], our GaussianFusion achieves +2.6 NDS
 383 and +3.2 mAP on val split by exploring a more natural continuous Gaussian cross-modal comple-
 384 mentary fusion. In addition, compared with recent SOTA fusion works, such as UniTR [Wang et al.
 385 (2023a)], EA-LSS [Hu et al. (2023b)], and FusionFormer-S [Hu et al. (2023a)], GaussianFusion shows
 386 superior performance, outperforming them by 1.2, 0.5, and 1.7 respectively in mAP.

387 Additionally, we provide a comparison
 388 with other open-source state-of-the-art
 389 (SOTA) methods in inference latency
 390 and performance accuracy in Table 3.
 391 Benefiting from the unified architecture,
 392 it achieves an excellent performance of
 393 71.7 mAP while maintaining lower in-
 394ference latency (132 ms) and memory consumption (4271 MB) compared to BEVFusion.

395 Here, we design a simple temporal
 396 extension, termed GaussianFusion-T:
 397 historical Gaussian representations are
 398 warped to the current timestamp and
 399 then fused via Equations (5) and (6). In
 400 Equation (5), the mean μ_m , scale s_m ,
 401 rotation r_m , and query q^m encompass
 402 not only the multi-modal (image and Li-
 403 DAR) Gaussians and queries, but also
 404 those from history frames. Experi-
 405 mental results show that, compared to BEV-
 406 Fusion4D [Liu et al. (2023b)], our temporal variant GaussianFusion-T achieves significant im-
 407 provements. Moreover, even without sophis-
 408 ticated temporal modeling, GaussianFusion-T achieves com-
 409 petitive NDS against advanced temporal fusion methods such as SparseLIF-T [Zhang et al. (2024a)].

410 Advanced Detection Head Ex- 411 periment.

412 For a fair comparison, we initially adopt the same de-
 413 tection head as BEVFusion [Liu
 414 et al. (2023b)] and UniTR to high-
 415 light the advantages of Gaussian
 416 representation. However, GaussianFusion can seamlessly adopt
 417 advanced query-based detection
 418 heads [Bai et al. (2022); Yan et al.
 419 (2023)], which bring additional
 420 performance gains and enable it to achieve detection metrics comparable to state-of-the-art sparse
 421 detection specialists. Specifically, query-based head [Yan et al. (2023)] treats image Gaussians and
 422 LiDAR Gaussians as *Key (K)* and *Value (V)*. The position encoding is consistent with the PE men-
 423 tioned in this paper. The queries then interact with the *K* and *V* through a Transformer decoder
 424 structure, followed by a feed-forward-network to predict the final 3D bounding boxes.

425 As shown in Tab. 5, our GaussianFusion (query-based) achieves NDS and mAP on par with the
 426 SOTA Mv2DFusion [Wang et al. (2024b)], further demonstrating the superiority of Gaussian repre-
 427 sentation over BEV representation for perception tasks.

428 4.4 3D SEMANTIC OCCUPANCY PREDICTION

429 **Setting.** We report the Intersection-over-Union (IoU) of occupied voxels as the evaluation metric
 430 of the class-agnostic scene completion task and the mIoU of all semantic classes for the Occ task
 431 following SurroundOcc [Wei et al. (2023)].

Table 3: Latency and performance on nuScenes *val.* set.

Method	Latency ↓	Memory ↓	NDS ↑	mAP ↑
BEVFusion	156 ms	5140 M	71.4	68.5
GaussianFusion	132 ms	4271 M	74.0	71.7

Table 4: Comparison with temporal methods.

Method	NDS ↑	mAP ↑
BEVFusion4D [Liu et al. (2023b)]	73.5	72.0
FusionFormer [Hu et al. (2023a)]	74.1	71.4
SparseLIF-T [Zhang et al. (2024a)]	77.5	74.7
GaussianFusion-T	77.6	75.0

Table 5: Expansion experiment of query-based detection head.

Method	Head	NDS ↑	mAP ↑
SparseFusion [Zhou & Ti (2023)]	Query-based	73.1	71.0
GaussianFusion(ours)	Heatmap-based	74.0	71.7
IS-FUSION [Yin et al. (2024)]	Query-based	74.0	72.8
SparseLIF [Zhang et al. (2024a)]	Query-based	74.6	71.2
Mv2dfusion [Wang et al. (2024b)]	Query-based	74.7	72.8
GaussianFusion(ours)	Query-based	74.5	72.9

432
 433 Table 6: Semantic scene completion results on nuScenes [Wei et al. (2023); Caesar et al. (2020)] val
 434 set. \dagger represents trained on nuScenes. For Camera-only and C+L, the top performance is indicated
 435 in **bold black** and **bold blue**, respectively.

Method	Modality	IoU												mIoU					
		barrier	bicycle	bus	car	const. veh.	motorcycle	pedestrian	traffic cone	trailer	truck	drive. suf.	other flat	sidewalk	terrain	mannade	vegetation		
BEVFormer [Li et al. (2022b)]	C	30.50	16.75	14.22	6.58	23.46	28.28	8.66	10.77	6.64	4.05	11.20	17.78	37.28	18.00	22.88	22.17	13.80	22.21
TPVFormer [H et al. (2023)]	C	30.86	17.10	15.96	5.31	23.86	27.32	9.79	8.74	7.09	5.20	10.97	19.22	38.87	21.25	24.26	23.15	11.73	20.81
FB-Occ(1F) [Li et al. (2023)]	C	31.55	20.17	20.31	12.29	26.33	31.07	10.78	15.95	13.31	11.14	13.24	22.13	39.56	22.26	25.14	23.59	13.92	21.64
SurroundOcc [Wei et al. (2023)]	C	31.49	20.30	20.59	11.68	28.06	30.86	10.70	15.14	14.09	12.06	14.38	22.26	37.29	23.70	24.49	22.77	14.89	21.86
GaussianFormer [H et al. (2024)]	C	29.83	19.10	19.52	11.26	26.11	29.78	10.47	13.83	12.58	8.67	12.74	21.57	39.63	23.28	24.46	22.99	9.59	19.12
GaussianFusion-C	C	32.48	20.65	21.09	10.95	29.01	31.65	10.03	15.64	14.31	12.56	13.82	23.19	40.06	22.49	25.80	23.49	14.36	22.14
BEVFusion \dagger [Liu et al. (2023b)]	C+L	39.11	24.65	23.78	12.29	30.67	34.95	14.62	17.23	20.76	15.75	19.83	26.3	40.01	23.34	25.47	26.41	25.15	37.92
M-CONet [Wang et al. (2023c)]	C+L	39.20	24.70	24.80	13.00	31.60	34.80	14.60	18.00	20.00	14.70	20.00	26.60	39.20	22.80	26.10	26.00	26.00	37.10
CO-Occ [Pan et al. (2024)]	C+L	41.10	27.10	28.10	16.10	34.00	37.20	17.00	21.60	20.80	15.90	21.90	28.70	42.30	25.40	29.10	28.60	28.20	38.00
OccFusion [Ming et al. (2024)]	C+L	43.53	27.55	25.15	19.87	34.75	36.21	20.03	23.11	25.25	17.50	22.70	30.06	39.47	23.26	25.68	27.57	29.54	40.60
GaussianFusion	C+L	44.75	28.65	28.92	18.31	34.87	37.43	19.45	23.53	26.71	17.96	23.38	29.89	43.32	25.96	30.51	28.35	30.32	39.67

448
 449 **Results.** As shown in Table 6, our
 450 GaussianFusion achieves SOTA perfor-
 451 mance at 28.65 mIoU among all single-
 452 frame models. GaussianFusion outper-
 453 forms the multi-modal SOTA method
 454 OccFusion [Ming et al. (2024)], which is
 455 based on multi-scale voxel fusion, by
 456 +1.11 mIoU and significantly surpasses
 457 camera-only methods [H et al. (2023); Wei et al. (2023); Li et al. (2023)]. More importantly, benefiting
 458 from our proposed Gaussian initialization strategy and iterative update mechanism, GaussianFusion-
 459 C achieves a 1.55 mIoU improvement and nearly 4.5 \times computational efficiency compared to Gaus-
 460 sianFormer, while using only 30% of the Gaussians, as shown in Table 7. GaussianFormer randomly
 461 initializes a set of Gaussians in 3D space and predicts new Gaussian parameters for these Gaussians
 462 in an update. Extensive experiments demonstrate the effectiveness of our 3D Gaussian representa-
 463 tion across multiple tasks, including both object-centric and dense semantic perception.

464 Table 7: Comprehensive comparison with Gaussian-
 465 Former on nuScenes val set.

Method	Num. Gaussians	mIoU \uparrow	Latency \downarrow
GaussianFormer	140,000	19.10	475 ms
GaussianFusion-C	43,296	20.65	105 ms

464 Table 8: Ablation of Gaussian initializa-
 465 tion strategy.

Gaussian Initialization	NDS	mAP
Random Initialization	71.2	68.3
Backward Projection	72.4	70.0
Lidar Projection	73.6	71.1
Forward Projection	74.0	71.7

464 Table 9: Ablation of the proposed Gaussian Encoder.
 465 DA.G means Deformable Attention with Gaussian.

Share	Separate	DA.G	PE	Offset	NDS	mAP
✓			✓	✓	74.0	71.7
✓			✓	✓	73.6	71.1
		✓	✓	✓	73.4	71.0
✓			✓	✓	73.6	71.2
✓			✓	✓	73.2	70.8

4.5 ABLATION STUDIES

475 **Effect of Gaussian Initialization.** Table 8 provides a detailed analysis of the performance im-
 476 pact of different Gaussian initialization strategies, including the classic random initialization, our
 477 proposed forward projection strategy, the backward projection BEVFormer-based strategy [Li et al.
 478 (2023) (2022b)], and the strategy of projecting Lidar points onto the image. The initialization details
 479 and corresponding Gaussian encoder for the latter two projection strategies are in Appendix. The
 480 forward projection and Lidar projection strategies show comparable performance (74.0 NDS v.s
 481 73.6 NDS), both outperforming the backward projection method (72.4 NDS). Notably, the forward
 482 projection strategy brings a significant improvement of +2.8 NDS over random initialization.

483 **Effect of Gaussian Encoder.** In Table 9, we first compare the shared and separate Gaussian En-
 484 coders. We find that the shared Gaussian Encoder provides a slight performance improvement of
 485 +0.7 mAP. We attribute this to the unified Gaussian space, which helps the model learn uncertain
 486 cross-modal complementary features. And the sharing strategy makes the model leaner. We then

486 conduct an ablation study on the deformable attention module. Results show that deformable
 487 attention with Gaussian priors outperforms the vanilla variant by +0.4 NDS, demonstrating that the
 488 shape prior encoded by Gaussians facilitates model convergence and enhances detection accuracy.
 489 For the deformable attention-based query updating module, the results show that encoding the Gaus-
 490 sian properties as PE into the query leads to a gain of +0.5 mAP. For the Gaussian updating module,
 491 predicting the offsets of the properties, rather than the properties themselves, improves multimodal
 492 fusion perception by +0.9 mAP. This validates the multimodal Gaussian Encoder theory proposed
 493 in [\[3.3\] Gaussian Encoder](#).

494 4.6 VISUALIZATIONS

495 As shown in Fig. [4](#) in BEV object detection, compared to previous BEV-based SOTA methods
 496 like UniTR [\[Wang et al. \(2023a\)\]](#) and BEVFusion [\[Liu et al. \(2023b\)\]](#), GaussianFusion achieve higher
 497 accuracy for distant or small objects (white marks). Furthermore, UniTR exhibits significant object
 498 yaw errors (yellow marks). For Occ, GaussianFusion-C produces sharper object boundaries (red
 499 marks) and better class separation (yellow marks) compared to GaussianFormer. See Appendix for
 500 more visualizations.

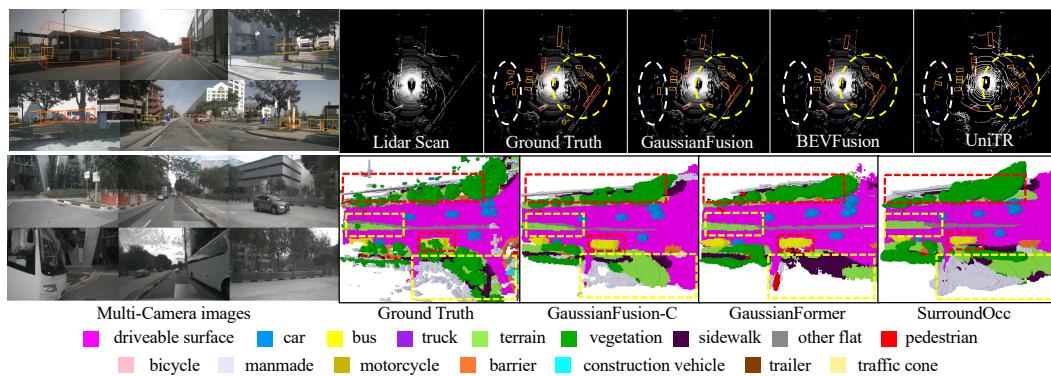


Figure 4: Qualitative results of GaussianFusion on BEV object detection and 3D semantic occupancy prediction.

547 4.7 LIMITATIONS

548 Several approaches—covering both detection [\[Wang et al. \(2023b\)\]](#) and Occ [\[Zhang et al. \(2024b\)\]](#)—employ carefully designed temporal fusion modules to enhance performance. While our
 549 method naturally extends to multi-frame settings through simple temporal alignment and already
 550 achieves performance comparable to such multi-frame methods, this is likely suboptimal. A promising
 551 direction for future work is to explore motion-aware Gaussian updates, for instance by predicting
 552 velocity-guided offsets, enabling more coherent 4D scene modeling over time.

557 5 CONCLUSION

558 We present GaussianFusion, a novel multi-modal fusion perception framework grounded in a unified
 559 3D Gaussian representation that seamlessly integrates camera and LiDAR features in a continuous
 560 spatial domain, effectively preserving fine-grained scene details. A shared Gaussian encoder is
 561 introduced to facilitate adaptive cross-modal interaction and alignment, with Gaussian properties it-
 562 eratively refined through optimization. To support task-agnostic applications, we design an efficient
 563 Gaussian-to-voxel transformation module incorporating Gaussian pooling and aggregation mech-
 564 anisms. Extensive experiments across multiple 3D perception tasks on the nuScenes dataset validate
 565 the effectiveness of GaussianFusion, achieving state-of-the-art performance among task-agnostic
 566 baselines. Although it may slightly lag behind certain task-specialized approaches, our work rep-
 567 presents a meaningful step toward generalizable and principled multi-modal fusion. We believe that
 568 the proposed Gaussian representation paradigm offers a promising direction for future research in
 569 multimodal 3D perception.

540 REFERENCES
541

542 Xuyang Bai, Zeyu Hu, Xinge Zhu, Qingqiu Huang, Yilun Chen, Hongbo Fu, and Chiew-Lan Tai.
543 Transfusion: Robust lidar-camera fusion for 3d object detection with transformers. In *Proceedings*
544 *of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 1090–1099, 2022.

545 Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liang, Qiang Xu, Anush
546 Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for
547 autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern*
548 *recognition*, pp. 11621–11631, 2020.

549 Anh-Quan Cao, Angela Dai, and Raoul de Charette. Pasco: Urban 3d panoptic scene completion
550 with uncertainty awareness. In *Proceedings of the IEEE/CVF Conference on Computer Vision*
551 *and Pattern Recognition*, pp. 14554–14564, 2024.

553 Florian Chabot, Nicolas Granger, and Guillaume Lapouge. Gaussianbev: 3d gaussian representation
554 meets perception models for bev segmentation. *arXiv preprint arXiv:2407.14108*, 2024.

555 Yilun Chen, Zhiding Yu, Yukang Chen, Shiyi Lan, Anima Anandkumar, Jiaya Jia, and Jose M
556 Alvarez. Focalformer3d: focusing on hard instance for 3d object detection. In *Proceedings of the*
557 *IEEE/CVF International Conference on Computer Vision*, pp. 8394–8405, 2023.

559 Zehui Chen, Zhenyu Li, Shiquan Zhang, Liangji Fang, Qinhong Jiang, and Feng Zhao. Deformable
560 feature aggregation for dynamic multi-modal 3d object detection. In *European conference on*
561 *computer vision*, pp. 628–644. Springer, 2022.

562 Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo
563 Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF*
564 *conference on computer vision and pattern recognition*, pp. 5501–5510, 2022.

566 Wanshui Gan, Fang Liu, Hongbin Xu, Ningkai Mo, and Naoto Yokoya. Gaussianocc: Fully
567 self-supervised and efficient 3d occupancy estimation with gaussian splatting. *arXiv preprint*
568 *arXiv:2408.11447*, 2024.

569 Chongjian Ge, Junsong Chen, Enze Xie, Zhongdao Wang, Lanqing Hong, Huchuan Lu, Zhenguo
570 Li, and Ping Luo. Metabev: Solving sensor failures for 3d detection and map segmentation.
571 In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 8721–8731,
572 2023.

573 Yuanhui H, Wenzhao Zheng, Yunpeng Zhang, Jie Zhou, and Jiwen Lu. Tri-perspective view for
574 vision-based 3d semantic occupancy prediction. In *Proceedings of the IEEE/CVF conference on*
575 *computer vision and pattern recognition*, pp. 9223–9232, 2023.

577 Yuanhui H, Wenzhao Zheng, Yunpeng Zhang, Jie Zhou, and Jiwen Lu. Gaussianformer: Scene as
578 gaussians for vision-based 3d semantic occupancy prediction. *arXiv preprint arXiv:2405.17429*,
579 2024.

580 Chunyong Hu, Hang Zheng, Kun Li, Jianyun Xu, Weibo Mao, Maochun Luo, Lingxuan Wang,
581 Mingxia Chen, Kaixuan Liu, Yiru Zhao, et al. Fusionformer: A multi-sensory fusion in bird’s-
582 eye-view and temporal consistent transformer for 3d objection. *arXiv preprint arXiv:2309.05257*,
583 2023a.

584 Haotian Hu, Fanyi Wang, Jingwen Su, Yaonong Wang, Laifeng Hu, Weiye Fang, Jingwei Xu, and
585 Zhiwang Zhang. Ea-lss: Edge-aware lift-splat-shot framework for 3d bev object detection. *arXiv*
586 *preprint arXiv:2303.17895*, 2023b.

588 Xu Hu, Yuxi Wang, Lue Fan, Junsong Fan, Junran Peng, Zhen Lei, Qing Li, and Zhaoxiang Zhang.
589 Semantic anything in 3d gaussians. *arXiv preprint arXiv:2401.17857*, 2024.

590 Junjie Huang and Guan Huang. Bevdet4d: Exploit temporal cues in multi-camera 3d object detec-
591 tion. *arXiv preprint arXiv:2203.17054*, 2022a.

593 Junjie Huang and Guan Huang. Bevpoolv2: A cutting-edge implementation of bevdet toward de-
ployment. *arXiv preprint arXiv:2211.17111*, 2022b.

594 Yang Jiao, Zequn Jie, Shaoxiang Chen, Jingjing Chen, Lin Ma, and Yu-Gang Jiang. Msmdfusion:
 595 Fusing lidar and camera at multiple scales with multi-depth seeds for 3d object detection. In
 596 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 21643–
 597 21652, 2023.

598 Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splat-
 599 ting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023.

600 Jinke Li, Xiao He, Chonghua Zhou, Xiaoqiang Cheng, Yang Wen, and Dan Zhang. Viewformer: Ex-
 601 ploring spatiotemporal modeling for multi-view 3d occupancy perception via view-guided trans-
 602 formers. In *European Conference on Computer Vision*, pp. 90–106. Springer, 2025a.

603 Yanwei Li, Yilun Chen, Xiaojuan Qi, Zeming Li, Jian Sun, and Jiaya Jia. Unifying voxel-based rep-
 604 resentation with transformer for 3d object detection. *Advances in Neural Information Processing
 605 Systems*, 35:18442–18455, 2022a.

606 Yiheng Li, Yang Yang, and Zhen Lei. Rctrans: Radar-camera transformer via radar densifier and
 607 sequential decoder for 3d object detection. In *Proceedings of the AAAI Conference on Artificial
 608 Intelligence*, volume 39, pp. 5048–5056, 2025b.

609 Yingyan Li, Lue Fan, Yang Liu, Zehao Huang, Yuntao Chen, Naiyan Wang, and Zhaoxiang Zhang.
 610 Fully sparse fusion for 3d object detection. *IEEE Transactions on Pattern Analysis and Machine
 611 Intelligence*, 2024.

612 Zhiqi Li, Wenhui Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Yu Qiao, and Jifeng Dai.
 613 Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotempo-
 614 ral transformers. In *European conference on computer vision*, pp. 1–18. Springer, 2022b.

615 Zhiqi Li, Zhiding Yu, Wenhui Wang, Anima Anandkumar, Tong Lu, and Jose M Alvarez. Fb-
 616 bev: Bev representation from forward-backward view transformations. In *Proceedings of the
 617 IEEE/CVF International Conference on Computer Vision*, pp. 6919–6928, 2023.

618 Tingting Liang, Hongwei Xie, Kaicheng Yu, Zhongyu Xia, Zhiwei Lin, Yongtao Wang, Tao Tang,
 619 Bing Wang, and Zhi Tang. Bevfusion: A simple and robust lidar-camera fusion framework.
 620 *Advances in Neural Information Processing Systems*, 35:10421–10434, 2022.

621 Haisong Liu, Yang Chen, Haiguang Wang, Zetong Yang, Tianyu Li, Jia Zeng, Li Chen, Hongyang
 622 Li, and Limin Wang. Fully sparse 3d occupancy prediction. *arXiv preprint arXiv:2312.17118*,
 623 2024.

624 Shuai Liu, Quanmin Liang, Zefeng Li, Boyang Li, and Kai Huang. Gaussianfusion: Gaussian-based
 625 multi-sensor fusion for end-to-end autonomous driving. *arXiv preprint arXiv:2506.00034*, 2025.

626 Yingfei Liu, Junjie Yan, Fan Jia, Shuailin Li, Aqi Gao, Tiancai Wang, and Xiangyu Zhang. Petrv2: A
 627 unified framework for 3d perception from multi-camera images. In *Proceedings of the IEEE/CVF
 628 International Conference on Computer Vision*, pp. 3262–3272, 2023a.

629 Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo.
 630 Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the
 631 IEEE/CVF international conference on computer vision*, pp. 10012–10022, 2021.

632 Zhijian Liu, Haotian Tang, Alexander Amini, Xinyu Yang, Huizi Mao, Daniela L Rus, and Song
 633 Han. Bevfusion: Multi-task multi-sensor fusion with unified bird’s-eye view representation. In
 634 *2023 IEEE international conference on robotics and automation (ICRA)*, pp. 2774–2781. IEEE,
 635 2023b.

636 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint
 637 arXiv:1711.05101*, 2017.

638 Yuhang Lu, Xinge Zhu, Tai Wang, and Yuexin Ma. Octreeocc: Efficient and multi-granularity
 639 occupancy prediction using octree queries. *arXiv preprint arXiv:2312.03774*, 2023.

648 Junyi Ma, Xieyuanli Chen, Jiawei Huang, Jingyi Xu, Zhen Luo, Jintao Xu, Weihao Gu, Rui Ai,
 649 and Hesheng Wang. Cam4docc: Benchmark for camera-only 4d occupancy forecasting in au-
 650 tonomous driving applications. In *Proceedings of the IEEE/CVF Conference on Computer Vision*
 651 and *Pattern Recognition*, pp. 21486–21495, 2024a.

652 Qihang Ma, Xin Tan, Yanyun Qu, Lizhuang Ma, Zhizhong Zhang, and Yuan Xie. Cotr: Compact oc-
 653 cupancy transformer for vision-based 3d occupancy prediction. In *Proceedings of the IEEE/CVF*
 654 *Conference on Computer Vision and Pattern Recognition*, pp. 19936–19945, 2024b.

655 Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and
 656 Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications*
 657 of the ACM, 65(1):99–106, 2021.

658 Zhenxing Ming, Julie Stephany Berrio, Mao Shan, and Stewart Worrall. Occfusion: A straightfor-
 659 ward and effective multi-sensor fusion framework for 3d occupancy prediction. *arXiv e-prints*,
 660 pp. arXiv–2403, 2024.

661 Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics prim-
 662 itives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)*, 41(4):1–15,
 663 2022.

664 Jingyi Pan, Zipeng Wang, and Lin Wang. Co-occ: Coupling explicit feature fusion with volume
 665 rendering regularization for multi-modal 3d semantic occupancy prediction. *IEEE Robotics and*
 666 *Automation Letters*, 2024.

667 Jonah Philion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs
 668 by implicitly unprojecting to 3d. In *Computer Vision–ECCV 2020: 16th European Conference,*
 669 *Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16*, pp. 194–210. Springer, 2020.

670 Leslie N Smith. Cyclical learning rates for training neural networks. In *2017 IEEE winter conference*
 671 *on applications of computer vision (WACV)*, pp. 464–472. IEEE, 2017.

672 Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui,
 673 James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for au-
 674 tonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF conference on com-*
 675 *puter vision and pattern recognition*, pp. 2446–2454, 2020.

676 Xiaoyu Tian, Tao Jiang, Longfei Yun, Yucheng Mao, Huitong Yang, Yue Wang, Yilun Wang, and
 677 Hang Zhao. Occ3d: A large-scale 3d occupancy prediction benchmark for autonomous driving.
 678 *Advances in Neural Information Processing Systems*, 36, 2024.

679 Sourabh Vora, Alex H Lang, Bassam Helou, and Oscar Beijbom. Pointpainting: Sequential fusion
 680 for 3d object detection. In *Proceedings of the IEEE/CVF conference on computer vision and*
 681 *pattern recognition*, pp. 4604–4612, 2020.

682 Haiyang Wang, Hao Tang, Shaoshuai Shi, Aoxue Li, Zhenguo Li, Bernt Schiele, and Liwei Wang.
 683 Unitr: A unified and efficient multi-modal transformer for bird’s-eye-view representation. In
 684 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 6792–6802,
 685 2023a.

686 Shihao Wang, Yingfei Liu, Tiancai Wang, Ying Li, and Xiangyu Zhang. Exploring object-centric
 687 temporal modeling for efficient multi-view 3d object detection. In *Proceedings of the IEEE/CVF*
 688 *International Conference on Computer Vision*, pp. 3621–3631, 2023b.

689 Xiaofeng Wang, Zheng Zhu, Wenbo Xu, Yunpeng Zhang, Yi Wei, Xu Chi, Yun Ye, Dalong Du, Ji-
 690 wen Lu, and Xingang Wang. Openoccupancy: A large scale benchmark for surrounding semantic
 691 occupancy perception. In *Proceedings of the IEEE/CVF International Conference on Computer*
 692 *Vision*, pp. 17850–17859, 2023c.

693 Yuqi Wang, Yuntao Chen, Xingyu Liao, Lue Fan, and Zhaoxiang Zhang. Panoocc: Unified
 694 occupancy representation for camera-based 3d panoptic segmentation. In *Proceedings of the*
 695 *IEEE/CVF conference on computer vision and pattern recognition*, pp. 17158–17168, 2024a.

702 Zitian Wang, Zehao Huang, Yulu Gao, Naiyan Wang, and Si Liu. Mv2dfusion: Leveraging modality-
 703 specific object semantics for multi-modal 3d detection. *arXiv preprint arXiv:2408.05945*, 2024b.
 704

705 Yi Wei, Linqing Zhao, Wenzhao Zheng, Zheng Zhu, Jie Zhou, and Jiwen Lu. Surroundocc: Multi-
 706 camera 3d occupancy prediction for autonomous driving. In *Proceedings of the IEEE/CVF Interna-*
 707 *tional Conference on Computer Vision*, pp. 21729–21740, 2023.

708 Junjie Yan, Yingfei Liu, Jianjian Sun, Fan Jia, Shuailin Li, Tiancai Wang, and Xiangyu Zhang. Cross
 709 modal transformer: Towards fast and robust 3d object detection. In *Proceedings of the IEEE/CVF*
 710 *International Conference on Computer Vision*, pp. 18268–18278, 2023.

711 Zeyu Yang, Jiaqi Chen, Zhenwei Miao, Wei Li, Xiatian Zhu, and Li Zhang. Deepinteraction: 3d
 712 object detection via modality interaction. *Advances in Neural Information Processing Systems*,
 713 35:1992–2005, 2022.

714 Mingqiao Ye, Martin Danelljan, Fisher Yu, and Lei Ke. Gaussian grouping: Segment and edit
 715 anything in 3d scenes. In *European Conference on Computer Vision*, pp. 162–179. Springer,
 716 2025.

717 Junbo Yin, Jianbing Shen, Runnan Chen, Wei Li, Ruigang Yang, Pascal Frossard, and Wenguan
 718 Wang. Is-fusion: Instance-scene collaborative fusion for multimodal 3d object detection. *arXiv*
 719 *preprint arXiv:2403.15241*, 2024.

720 Hongcheng Zhang, Liu Liang, Pengxin Zeng, Xiao Song, and Zhe Wang. Sparselif: High-
 721 performance sparse lidar-camera fusion for 3d object detection. In *European Conference on*
 722 *Computer Vision*, pp. 109–128. Springer, 2024a.

723 Shuo Zhang, Yupeng Zhai, Jilin Mei, and Yu Hu. Fusionocc: Multi-modal fusion for 3d occupancy
 724 prediction. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pp. 787–
 725 796, 2024b.

726 Xiao Zhao, Bo Chen, Mingyang Sun, Dingkang Yang, Youxing Wang, Xukun Zhang, Mingcheng
 727 Li, Dongliang Kou, Xiaoyi Wei, and Lihua Zhang. Hybridocc: Nerf enhanced transformer-based
 728 multi-camera 3d occupancy prediction. *IEEE Robotics and Automation Letters*, 2024a.

729 Xiao Zhao, Xukun Zhang, Dingkang Yang, Mingyang Sun, Mingcheng Li, Shunli Wang, and Lihua
 730 Zhang. Maskbev: Towards a unified framework for bev detection and map segmentation. *arXiv*
 731 *preprint arXiv:2408.09122*, 2024b.

732 Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection.
 733 In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4490–
 734 4499, 2018.

735 Zhizhuo Zhou and Shubham T. Sparsefusion: Distilling view-conditioned diffusion for 3d recon-
 736 struction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recog-*
 737 *nition*, pp. 12588–12597, 2023.

738 Benjin Zhu, Zhengkai Jiang, Xiangxin Zhou, Zeming Li, and Gang Yu. Class-balanced grouping
 739 and sampling for point cloud 3d object detection. *arXiv preprint arXiv:1908.09492*, 2019.

740 Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr:
 741 Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020.

742 Sicheng Zuo, Wenzhao Zheng, Yuanhui Huang, Jie Zhou, and Jiwen Lu. Gaussianworld: Gaussian
 743 world model for streaming 3d occupancy prediction. *arXiv preprint arXiv:2412.10373*, 2024.

744

745

746

747

748

749

750

751

752

753

754

755