



Online Vectorized HD Map Construction using Geometry

Zhixin Zhang¹  Yiyuan Zhang²  Xiaohan Ding³ Fusheng Jin^{1*} Xiangyu Yue²

¹School of Computer Science and Technology, Beijing Institute of Technology

²The Chinese University of Hong Kong ³Tencent AI Lab

zhangzhixin@bit.edu.cn, yiyuanzhang.ai@gmail.com

<https://invictus717.github.io/GeMap/>

Abstract. Online vectorized High-Definition (HD) map construction is critical for downstream prediction and planning. Recent efforts have built strong baselines for this task, however, geometric shapes and relations of instances in road systems are still under-explored, such as parallelism, perpendicular, rectangle-shape, *etc.* In our work, we propose GeMap (**Geometry Map**), which end-to-end learns Euclidean shapes and relations of map instances beyond fundamental perception. Specifically, we design a geometric loss based on angle and magnitude clues, robust to rigid transformations of driving scenarios. To address the limitations of the vanilla attention mechanism in learning geometry, we propose to decouple self-attention to handle Euclidean shapes and relations independently. GeMap achieves new state-of-the-art performance on the nuScenes and Argoverse 2 datasets. Remarkably, it reaches a 71.8% mAP on the large-scale Argoverse 2 dataset, outperforming MapTRv2 by +4.4% and surpassing the 70% mAP threshold for the first time. Code is available at <https://github.com/cnzxx/GeMap>.

Keywords: HD Map Construction · Geometry Representation · Geometry-Decoupled Attention

1 Introduction

Vectorized HD maps provide structured environmental information for autonomous vehicles and have been widely adopted in downstream tasks, such as trajectory forecasting [3, 21, 55] and planning [6, 37]. Online Vectorized HD Map Construction can significantly reduce the need for labor-intensive annotations and facilitate real-time updates in autonomous driving [4, 18, 22, 53].

With the development of Bird’s-Eye-View (BEV) representation and multi-modal learning [5, 11, 19, 28, 50–52], online HD map construction has achieved significant advancements [10, 20, 23]. Early works [14, 32, 35, 47] formulate HD map construction as a dense prediction task. However, these methods generate maps in image format, which is redundant for representing sparse map instances. Then,

* Corresponding author

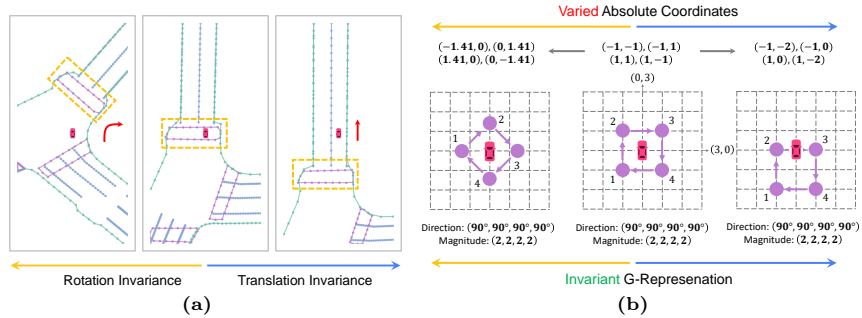


Fig. 1: Geometric Invariance. (a) As the ego vehicle moves, after rotation (\leftarrow) and translation (\rightarrow), the shape of the crossing and the parallelism between lanes remain unchanged, which indicates the invariant property of geometry to rigid transformations. (b) Absolute coordinates are vulnerable to rotation and translation, however, our G-representation is invariant, which is more suitable to capture geometric properties.

a more compact map formulation was introduced to minimize redundancy, albeit at the cost of adding time-consuming post-processing steps [18]. In response to this, recent works [4, 22, 26, 36] attempt to end-to-end construct vectorized HD maps to avoid extensive post-processing. These methods typically sample points from map instances and represent each instance as a polyline [4, 22, 26] or parameterized curve [36].

We observe that transportation road systems exhibit significant geometric characteristics (Figure 2a), such as parallel lanes, perpendicular crossings, equal lane widths, *etc.* However, these geometric properties of shapes and relations between map instances have not been fully explored. Meanwhile, there are two notable limitations among existing methods [4, 22, 26, 36] which can be alleviated by leveraging these geometric properties. **1) Excessive dependency on absolute coordinates:** as the ego-vehicle moves, instances experience rotations and translations, as depicted in Figure 1a. However, widely adopted representations such as polylines and parameterized curves [8, 36] are inherently sensitive to rotation and translation changes. **2) Objective conflicts of vanilla attention mechanism (detail in § 3.5)** struggles to learn diverse shapes and relations, even though geometric properties such as rectangle shape, parallelism, and perpendicular relationships are commonly found in driving scenarios, as illustrated in Figure 2a. We believe that incorporating these geometric properties significantly enhances the precision and efficiency of online HD map construction.

To address the above two limitations, we introduce a novel geometric representation that captures shapes of individual map instances and relations between different instances as illustrated in Figure 2b and 2c, referred to as **G-Representation**. It enhances the vanilla representation of map instances by incorporating a translation- and rotation-invariant representation that effectively leverages instance geometry. The local structures of map features are encoded using *displacement vectors*. These vectors are computed from the absolute coordinates.

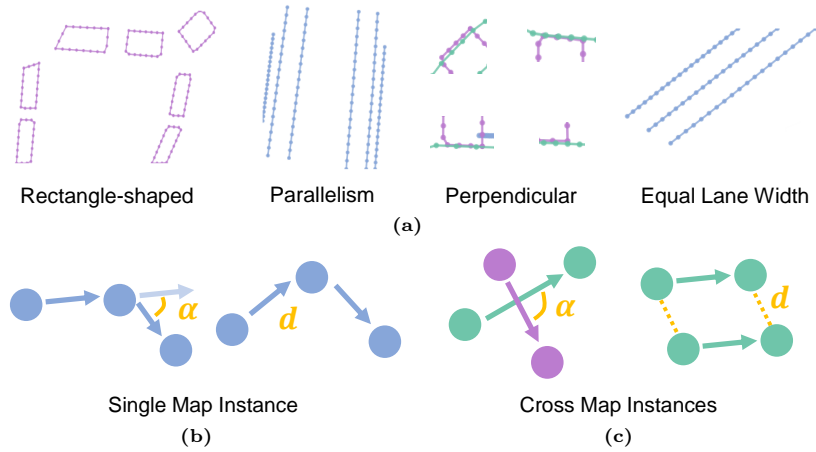


Fig. 2: Geometric properties and G-Representation. (a) Geometry in the transportation road system. (b) (c) We propose to model the geometric properties of a single map instance and multiple instances, with magnitude d and angle α .

ordinates of polyline points to effectively represent the relative positions and orientations of adjacent points. To quantify these features within Euclidean space, we employ both the *magnitudes* of the displacement vectors and the *angles* formed between these vectors. Based on G-Representation, we propose the **Euclidean Shape Clues** to represent the shapes of individual map instances and the **Euclidean Relation Clues** to represent the inter-instance relations in Euclidean space.

G-Representation is a simple yet effective method to address the abovementioned two limitations: **1) Geometric Invariance**, by concentrating on the relative relationships between points within and across map instances, it inherently attains translation and rotation invariance as illustrated in Figure 1b. This enhances its robustness against variations in data collection and equips it to effectively handle different coordinate systems. **2) Euclidean Modeling**, it captures the inherent geometry in transportation road systems. Diverse shapes and relation geometry can be simplified as enumerations of magnitudes and directions. For instance, in Euclidean space, parallel lanes can be easily modeled as approximate directions ($\alpha \approx 0$).

Building upon G-Representation, we propose a framework named **GeMap** for HD map construction. A BEV encoder is used to extract features from multi-view input images, while a geometry-decoupled decoder is employed to focus on geometric aspects. Specifically, we adapt the attention mechanism [42] and propose **Geometry-Decoupled Attention (GDA)**. GDA sequentially applies attention to queries belonging to the same instance and attention to queries across different instances. This can significantly boost the geometry learning of key points of shapes and relations between map instances. Furthermore, we propose an objective function named *Euclidean Loss* to optimize G-Representation.

Specifically, we transform the conventional polylines of the ground truth map into our G-Representation. In this way, the model gets optimized to better understand the magnitude of displacement vectors and the angles between them, thereby facilitating a more effective learning of geometric properties.

Experiments on nuScenes and Argoverse 2 datasets demonstrate the effectiveness of GeMap. We reach new state-of-the-art performances on both datasets. With camera images only, GeMap achieves 69.4% and 71.8% mAP on the nuScenes and Argoverse 2, respectively. Visualization results (§ 4.4) further demonstrate the better perception of shape and relation, alleviation of occlusion, and robustness to rigid transformations of GeMap.

Our contributions are summarized as the following:

- We propose G-Representation which harnesses critical geometric properties of rotation and translation invariance in autonomous driving scenarios, opening up new research avenues within the field.
- We introduce GeMap, a novel framework incorporating geometry-decoupled attention and Euclidean Loss function, specifically designed to learn the intrinsic geometry of online HD maps.
- GeMap achieves new state-of-the-art results in HD map construction on the nuScenes and Argoverse 2 datasets, notably surpassing the 70% mAP on the large-scale Argoverse 2 dataset for the first time.

2 Related Work

2.1 Online HD Map Construction

Traditionally, HD map construction has required labor-intensive manual or semi-automatic annotations [14,33]. To streamline this, recent studies [20,31,34,35,54] have focused on online construction, approaching HD maps as a dense prediction challenge. Innovations such as MetaBEV [9] aim to mitigate sensor issues, while MVNet [45] uses historical data for improved semantic consistency. The trend towards automatic vectorization of HD maps is spearheaded by works such as HDMapNet [18], which fuses camera and LiDAR inputs in BEV space and is advanced by end-to-end solutions such as VectorMapNet [26] and MapTR [22], leveraging Transformer-based models. StreamMapNet [48] and PivotNet [4] build upon this, modifying attention mechanisms for better performance. Our contribution starts from the strengths of the Transformer architecture, optimizing it with a decoupled self-attention block for enhanced geometric processing.

2.2 Cross-view BEV Learning

The conversion of Perspective View (PV) camera images into a unified BEV space is a significant challenge for autonomous driving systems. Previous studies such as [19,35] use depth estimation from monocular images for this transformation, others [2,20,28] have developed PV-to-BEV conversion methods without

explicit depths. However, these methods can distort shapes and geometric relations, leading to inaccuracies. Our approach integrates geometric supervision in prediction, addressing these geometric inconsistencies while acknowledging the value of depth information.

2.3 Geometric Instance Modeling

Traditional vectorization approaches model these instances as polylines [22, 26, 38, 43] or polynomial curves [7, 36, 39, 41], with some incorporating Bézier Curves for improved fitting [7, 36]. However, such methods often overlook geometric properties including shapes, parallelism, perpendicular, and *etc.* In 3D Lane Detection, some works [13, 17, 25] also explore simple geometry priors of lanes such as “equal lane width”. However, they are not generalizable to more complex HD map instances. Existing works on modeling geometry properties of map instances mainly address rasterization-based shape loss [41, 49] or edge loss [22], and they fail to preserve cross-instance relationships. PivotNet [4] offers an architectural advancement with its line-aware point decoder, yet it still does not fully address the complexity of instance geometry.

Given the current under-exploration of diverse geometric properties in shape and relation, we propose utilizing Euclidean shape and relation losses and decoupling the traditional self-attention module to empower the model with a more robust understanding of instance geometry.

3 Method

3.1 Preliminary

The architecture of our method is illustrated in Figure 3. Input images are denoted by $\mathbb{I} = \{\mathbf{I}_i\}_{i=1}^{N_c}$, where $\mathbf{I}_i \in \mathbb{R}^{H \times W \times 3}$ and N_c is the number of cameras. The output is a set of N map instances $\mathbb{M} = \{\mathbf{L}_i\}_{i=1}^N$, where each map instance is represented by a polyline $\mathbf{L}_i \in \mathbb{R}^{N_v \times 2}$, *i.e.*, an ordered sequence of N_v two-dimensional points. As shown in Figure 3, the commonly adopted pipeline comprises three steps. **1)** A BEV feature extractor processes the multi-view images. **2)** The extracted BEV features are fed into a map decoder, which predicts the map instances based on the BEV features. **3)** To optimize the pipeline, the predicted map instances are compared against the ground truth. Beyond point-to-point comparison, G-representations of predictions and the ground truth are computed and the difference is measured by L^1 loss.

3.2 Architecture Overview

BEV Feature Extractor. A shared vision backbone in different views and parameterized PV-to-BEV transformation network are employed to aggregate features from various perspectives. In our default configuration, we utilize ResNet [12] as the PV backbone and GKT [2] as the PV-to-BEV transformation network.

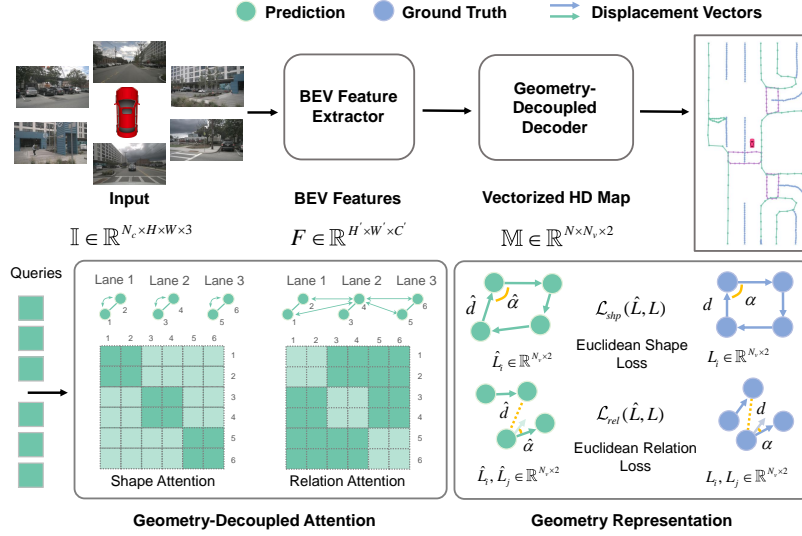


Fig. 3: Illustration of GeMap. First, PV images are transformed into BEV features, then a Geometry-Decoupled Decoder outputs the vectorized HD Map. In each block of the decoder, queries are first processed by Euclidean shape and relation attention, which focuses on geometric relevance. Finally, predictions are enhanced in G-Representations by shape and relation constraint.

Geometry-Decoupled Decoder. We adopt a Transformer decoder to predict polylines. The decoder obtains $N \times N_v$ queries that represent points and processes them via self-attention and aggregate BEV features via deformable cross-attention [56]. Finally, a prediction head is used to convert queries to polylines $\{\hat{L}_i\}_{i=1}^N$. We utilize a multi-layer perceptron as the polyline prediction head.

Geometric Loss. For training, we convert both the predicted polylines and ground truth into the proposed G-Representation. Based on that, we let the model optimize the predicted magnitudes of displacement vectors and angles between displacement vectors to match the ground truth.

In the following subsections, we first introduce Geometric Representation in § 3.3 since it is the core of our framework. Based on that, we introduce Geometric Loss in § 3.4 and Geometry-Decoupled Decoder in § 3.5, respectively.

3.3 Geometric Representation

Euclidean Shape Clues. We first introduce the representation of shapes of individual map instances. For each instance, we describe the local geometry with displacement vectors between neighboring points which are computed as:

$$\mathbf{v}_u^i = \mathbf{L}_{i,u+1} - \mathbf{L}_{i,u} \quad (u \in \{1, 2, \dots, N_v\}), \quad (1)$$

where we define $\mathbf{L}_{i,N_v+1} := \mathbf{L}_{i,1}$ to unify the geometric formulation of closed and open polylines.




Fig. 4: Euclidean Shape Clues. Magnitudes of displacement vectors and angles between neighboring vectors indicate shape clues and are utilized to compute shape loss. The right part shows how to connect Euclidean Shape Clues to shape geometry.

These displacement vectors are sufficient to represent the shape of a map instance and such a representation is invariant to translation transformations. However, we would like to note that this representation is vulnerable to rotations, which might prevent the model from learning robust geometry. To solve this problem, we propose to represent the shape with magnitudes of displacement vectors and angles between consecutive displacement vectors, as illustrated in Figure 4. Specifically, the i -th instance is denoted by N_v angle values and N_v magnitude values. Let u be the index, the u -th angle value and the magnitude value of the i -th instance are computed as:

$$\alpha_u^i = \langle \mathbf{v}_u^i, \mathbf{v}_{u+1}^i \rangle, \quad d_u^i = \|\mathbf{v}_u^i\|_2, \quad (2)$$

where $\langle \cdot \rangle$ denotes the angle between two vectors and similarly $\mathbf{v}_{N_v+1}^i := \mathbf{v}_1^i$.

The proposed representation inherently captures common geometric patterns, such as parallelism, right angles, and proper line width, by translating them into corresponding numerical patterns within this representation. For example, in Figure 4, a rectangle is characterized by one 90° angle and equal magnitudes of opposite displacement vectors; in addition, a straight line is represented by 0° angles. Beyond discussed regular cases, the Euclidean shape clues are capable of handling more complex shapes, as illustrated in scenarios (a) and (c) of Figure 6 and discussed in § 4.4.

Euclidean Relation Clues. Having highlighted the translation and rotation invariance of G-Representation and its advantages in representing the geometric shapes of individual map instances, we further introduce its ability to capture the relations between two map instances, *e.g.*, parallelism and perpendicular, in Euclidean space. Specifically, given the vanilla representations and displacement vectors of the i -th and j -th map instance, we represent the relation between instance i and j with 1) angles between each pair of their respective displacement vectors, and 2) magnitudes of displacement vectors between each pair of points, as illustrated in Figure 5. Formally,

$$\alpha_{u,v}^{i,j} = \langle \mathbf{v}_u^i, \mathbf{v}_v^j \rangle, \quad d_{u,v}^{i,j} = \|\mathbf{L}_{i,u} - \mathbf{L}_{j,v}\|_2. \quad (3)$$

This representation is also translation- and rotation-invariant. Similar to Euclidean Shape Clues, it inherently captures common relations between map instances by translating them into corresponding numerical patterns. For example,




Fig. 5: Euclidean Relation Clues. Angles between pairs of displacement vectors on different polylines, and magnitudes of displacement vectors between point pairs indicate relation clues.

Figure 5 shows that the perpendicular relation can be directly represented by a 90° angle, and the distance between two parallel polylines, which may correspond to the width of a lane, is naturally represented by magnitudes of displacement vector between point pairs. Scenarios (a) and (b) of Figure 6 show the effectiveness of Euclidean Relation Clues and more discussion can be viewed in § 4.4.

3.4 Euclidean Loss and Objectives

We transform both the ground truth data and the model’s predicted polylines from their original format into G-Representation. This conversion allows the model to independently optimize the angles between displacement vectors and the magnitudes of displacement vectors to more accurately align with the ground truth. The proposed Euclidean Loss is composed of two parts that measure how accurately the model predicts the shape of individual map instances and the inter-instance relations, respectively, which are denoted by \mathcal{L}_{shp} and \mathcal{L}_{rel} . The Euclidean Loss can be computed as the following:

$$\mathcal{L}_{\text{Euc}} = \lambda_1 \cdot \mathcal{L}_{\text{shp}} + \lambda_2 \cdot \mathcal{L}_{\text{rel}}, \quad (4)$$

where λ_1 and λ_2 are hyperparameters, whose effects on the model performances are evaluated in Figure 7a. Specifically,

$$w_{i,j} = 1 - \left(\min_{u,v} \{d_{u,v}^{i,j}\} / \sqrt{2} \right)^p, \quad \mathcal{L}_{\text{shp}} = \sum_{i=1}^N \sum_{j=1}^{N_v} |\hat{d}_j^i - d_j^i| + \ell(\hat{\alpha}_j^i, \alpha_j^i), \quad (5)$$

$$\mathcal{L}_{\text{rel}} = \sum_{i=1}^N \sum_{j>i}^N w_{i,j} \sum_{u=1}^{N_v} \sum_{v=1}^{N_v} \left(|\hat{d}_{u,v}^{i,j} - d_{u,v}^{i,j}| + \ell(\hat{\alpha}_{u,v}^{i,j}, \alpha_{u,v}^{i,j}) \right),$$

where $\ell(\cdot)$ is a function based on L^1 loss. We avoid using the inverse trigonometric function to directly compute the angles, but use sine and cosine values instead:

$$\ell(\hat{\alpha}, \alpha) = |\cos(\hat{\alpha}) - \cos(\alpha)| + |\sin(\hat{\alpha}) - \sin(\alpha)|. \quad (6)$$

Moreover, the distance might influence the relation strength of instance pairs. For example, if two instances are far from each other, the relation between them

might be weak. For this reason, we further adopt $w_{i,j}$ to punish weakly related instance pairs and discuss it in § 4.5. According to the best experimental results, we treat all instance pairs equally in other experiments.

Following common practices [22, 48], we also use focal loss [24] \mathcal{L}_{cls} for classification. For polyline regression, we adopt point-to-point loss and edge direction loss, denoted by \mathcal{L}_{pts} and \mathcal{L}_{dir} respectively. Following [23, 36], we further adopt segmentation loss \mathcal{L}_{seg} and depth estimation loss \mathcal{L}_{dep} . These losses are detailed in appendix § A.1. The overall loss function can be written as:

$$\mathcal{L} = \lambda \cdot \mathcal{L}_{\text{Euc}} + \beta_1 \cdot \mathcal{L}_{\text{cls}} + \beta_2 \cdot \mathcal{L}_{\text{pts}} + \beta_3 \cdot \mathcal{L}_{\text{dir}} + \beta_4 \cdot \mathcal{L}_{\text{seg}} + \beta_5 \cdot \mathcal{L}_{\text{dep}}. \quad (7)$$

where hyperparameters are discussed in Figure 7b and appendix Table A1.

3.5 Geometry-Decoupled Decoder

In the Transformer decoder, we obtain the $N \times N_v$ fused queries by adding N instance queries to N_v point queries, which is suggested in [22]. We denote input tokens by $\mathbf{E} \in \mathbb{R}^{N_A \times D}$, where D is the feature dimension and $N_A = N \cdot N_v$ is the number of fused queries. Intuitively, each fused query corresponds to a point on a predicted map instance. Self-Attention (SA) is formulated by:

$$\text{SA}(\mathbf{E}, \mathbf{M}) = \text{Softmax} \left(\frac{(\mathbf{E}\mathbf{W}^q)(\mathbf{E}\mathbf{W}^k)^\top}{\sqrt{D_k}} \odot \mathbf{M} \right) \mathbf{E}\mathbf{W}^v, \quad (8)$$

where $\mathbf{W}^q, \mathbf{W}^k, \mathbf{W}^v \in \mathbb{R}^{D \times D_k}$ are linear projection matrices, $\mathbf{M} \in \mathbb{R}^{N_A \times N_A}$ is the attention mask, and \odot is the Hadamard product. The vanilla SA computes relations between every pair of tokens.

Nevertheless, the geometry of shape and relation pertains to distinct subsets of tokens. For any given map instance, its shape is intimately related to tokens representing that instance’s points. *Precise shape geometry capture requires the model to discern token correlations specific to an instance while avoiding interference from tokens of unrelated instances.* Conversely, for relation geometry modeling, it is beneficial to isolate token correlations that span across different instances, rather than those confined within a single instance.

Therefore, we present Geometry-Decoupled Attention (GDA). First, we multiply a binary mask \mathbf{M} to the computed attention map so that the tokens of a map instance are aggregated according to the tokens within the same map instances only, which allows the model to adjust points’ positions according to learned shape geometry. Denote the index of map instance that the i -th token belongs to as \mathcal{I}_i . For example, tokens $1, 2, \dots, N_v$ belong to the first instance so that $\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_{N_v} = 1$. Tokens $N_v + 1, N_v + 2, \dots, 2 \cdot N_v$ belong to the second instance so that $\mathcal{I}_{N_v + 1}, \mathcal{I}_{N_v + 2}, \dots, \mathcal{I}_{2N_v} = 2$. With this notation, the binary mask \mathbf{M} can be simply constructed by:

$$\mathbf{M}_{i,j}^{\text{shp}} = \begin{cases} 1 & , \mathcal{I}_i = \mathcal{I}_j \\ 0 & , \mathcal{I}_i \neq \mathcal{I}_j \end{cases}, \quad \mathbf{M}_{i,j}^{\text{rel}} = \begin{cases} 1 & , \mathcal{I}_i \neq \mathcal{I}_j \\ 0 & , \mathcal{I}_i = \mathcal{I}_j \end{cases}. \quad (9)$$

The second attention is expected to model relations between tokens of different map instances.

Table 1: Comparison on the nuScenes dataset, GeMap reaches a new state-of-the-art performance. “EB0”, “R50”, “PP”, “Sec”, “Swin-T”, and “V2-99” denote EfficientNet-B0 [40], ResNet50 [12], PointPillars [15], SECOND [46], Swin Transformer Tiny [27], and VoVNetV2-99 [16] respectively. Methods with two backbones utilize both camera and LiDAR inputs. “Dns. Loss” denotes whether any dense prediction (*e.g.* semantic segmentation) loss is adopted. The best result is highlighted in **bold**. We reproduce all methods on a single RTX3090 GPU to test FPS for fair comparison.

Methods	Backbone	Dns. Loss	$AP_{div}(\uparrow)$	$AP_{ped}(\uparrow)$	$AP_{bnd}(\uparrow)$	mAP(\uparrow)	FPS(\uparrow)
VectorMapNet [ICML’23] [26]	R50		42.5	51.4	44.1	46.0	5.3
MapTR [ICLR’23] [22]	R50		59.8	56.2	60.1	58.7	19.8
MapVR [NeurIPS’23] [49]	R50		61.8	55.0	59.4	58.8	19.8
GeMap [Ours]	R50		65.1 \uparrow 3.3	59.8 \uparrow 4.8	63.2 \uparrow 3.8	62.7 \uparrow 3.9	19.0
HDMaNet [ICRA’22] [18]	EB0	✓	14.4	21.7	33.0	23.0	0.7
PivotNet [ICCV’23] [4]	R50	✓	58.8	53.8	59.6	57.4	9.5
BeMapNet [CVPR’23] [36]	R50	✓	66.7	62.6	65.1	64.8	6.6
MapTRv2 [Arxiv’23] [23]	R50	✓	68.3	68.1	69.7	68.7	15.0
GeMap [Ours]	R50	✓	69.8 \uparrow 1.5	67.1	71.4 \uparrow 1.7	69.4 \uparrow 0.7	15.8
HDMaNet [ICRA’22] [18]	EB0 & PP		29.6	16.3	46.7	31.0	0.7
VectorMapNet [ICML’23] [26]	R50 & PP		60.1	48.2	53.0	53.7	-
MapTR [ICLR’23] [22]	R50 & Sec		62.3	55.9	69.3	62.5	6.7
MapVR [NeurIPS’23] [49]	R50 & PP		62.7	60.4	67.2	63.5	-
GeMap [Ours]	R50 & Sec		66.3 \uparrow 3.6	62.2 \uparrow 1.8	71.1 \uparrow 3.9	66.5 \uparrow 3.0	7.3
MapTRv2 [Arxiv’23] [23]	R50 & Sec	✓	65.6	66.5	74.8	69.0	6.6
GeMap [Ours]	R50 & Sec	✓	69.8 \uparrow 4.2	68.0 \uparrow 1.5	73.4	70.4 \uparrow 1.4	6.8
MapTRv2 [Arxiv’23] [23]	V2-99	✓	73.7	71.4	75.0	73.4	10.1
GeMap [Ours]	Swin-T	✓	72.8	70.4	72.8	72.0	11.4
GeMap [Ours]	V2-99	✓	76.0 \uparrow 2.3	74.3 \uparrow 2.9	77.7 \uparrow 2.7	76.0 \uparrow 2.6	10.8

4 Experiments

4.1 Experimental Setups

Datasets. To evaluate GeMap, we conduct experiments on the nuScenes dataset [1], a widely adopted large-scale autonomous driving dataset that includes 1,000 scenes captured by six RGB cameras with a 360-degree field of view, and provides precise annotations from LiDAR point clouds for HD map construction. For comparability with prior research [18, 22], we focus on three static categories of map instances: pedestrian crossings, lane dividers, and road boundaries. With dataset splits provided by BeMapNet [36], we also evaluate GeMaps under three weather conditions: sunny, cloudy, and rainy. Additionally, we use the Argoverse 2 dataset [44] as another benchmark, which consists of approximately 108,000 frames, each providing images from seven cameras. We evaluate the performance by Average Precision (AP) [18, 22] with thresholds at 0.5, 1.0, and 1.5 meters.

Implementation Details. GeMap leverages 8 NVIDIA RTX 3090 GPUs for training. We adopt AdamW [30] as the optimizer and utilize Cosine Annealing with a linear warm-up phase [29] as the learning rate scheduler. For more details on hyperparameters, please refer to appendix § A.2.

Table 2: Following BeMapNet [36], we also compare GeMap on the nuScenes dataset under different weather conditions.

Methods	Backbone	Dns.	Loss	mAP _{sunny} (↑)	mAP _{cloud} (↑)	mAP _{rainy} (↑)	mAP _{avg} (↑)	FPS(↑)
VectorMapNet [ICML'23] [26]	R50			43.8	44.1	36.6	41.5	5.3
MapTR [ICLR'23] [22]	R50			62.1	60.5	52.8	58.4	19.8
GeMap [Ours]	R50			66.0 ↑ 3.9	64.3 ↑ 3.8	54.4 ↑ 1.6	61.5 ↑ 3.1	19.0
BeMapNet [CVPR'23] [36]	R50	✓		67.3	67.5	56.6	63.8	6.6
GeMap [Ours]	R50	✓		73.1 ↑ 5.8	71.0 ↑ 3.5	59.3 ↑ 2.7	67.8 ↑ 4.0	15.8

Table 3: Comparison on Argoverse 2 dataset. GeMap demonstrates significant performance improvements over previous methods.

Methods	Backbone	Dns.	Loss	AP _{div} (↑)	AP _{ped} (↑)	AP _{bnl} (↑)	mAP(↑)	FPS(↑)
VectorMapNet [ICML'23] [26]	R50			36.1	38.3	39.2	37.9	-
GeMap [Ours]	R50			67.6 ↑ 31.5	59.3 ↑ 21.0	64.7 ↑ 25.5	63.9 ↑ 26.0	16.7
HDMaNet [ICRA'22] [18]	E80	✓		5.7	13.1	37.6	18.8	-
MapTRv2 [Arxiv'23] [23]	R50	✓		72.1	62.9	67.1	67.4	13.6
GeMap [Ours]	R50	✓		75.7 ↑ 3.6	69.2 ↑ 6.3	70.5 ↑ 3.4	71.8 ↑ 4.4	13.8

4.2 Main Results

Results on nuScenes. As delineated in Table 1, GeMap delivers a state-of-the-art performance, achieving a mean Average Precision (mAP) of 69.4% using the camera-only setup and ResNet50 [12] backbone. GeMap particularly improves the precision for identifying dividers and boundaries, with enhancements of +1.5% and +1.7%, respectively. Importantly, these performance gains do not come at a high cost of efficiency, as GeMap sustains inference speeds on par with and even exceeds the previously established state-of-the-art, measured in frames per second (FPS). With more powerful vision backbones, GeMap can achieve significantly improved performance. Specifically, VoVNetV2-99 [16] and Swin-T [27] outperform the ResNet50 baseline by +6.6% and +2.6% in mAP, respectively. Also, GeMap achieves significantly improved performance with extra LiDAR inputs, which indicates the generalizability of GeMap to multi-modality settings. Moreover, we also compare performance under different weather conditions in Table 2. GeMap exceeds previous works under sunny, cloudy, and rainy scenarios, which shows the potential of geometry to raise model robustness to varied weather conditions.

Results on Argoverse 2. Extending our evaluation to the Argoverse 2 dataset, as presented in Table 3, GeMap also presents a new SOTA performance of 71.8% mAP, outperforming the highly advanced MapTRv2 model by +4.4%. This achievement not only demonstrates the effectiveness of GeMap on a different dataset but also emphasizes the adaptability and precision of GeMap in the evolving landscape of autonomous driving technologies.

Table 4: Ablation study on the nuScenes dataset. We train the model for 24 epochs. * denotes replacing GDA with 2 layers of vanilla self-attention.

Method	\mathcal{L}_{shp} (§ 3.3)	\mathcal{L}_{rel} (§ 3.3)	GDA (§ 3.5)	$\text{AP}_{\text{div}}(\uparrow)$	$\text{AP}_{\text{ped}}(\uparrow)$	$\text{AP}_{\text{bnd}}(\uparrow)$	mAP(\uparrow)
Baseline				49.5	44.7	53.7	49.3
+ Decoupled Attention			✓	53.4	46.6	53.5	51.2
+ Single Euclidean Loss	✓	✓		51.0	45.4	52.7	49.7
				51.7	43.4	53.0	49.4
+ Euclidean Loss	✓	✓		51.5	43.9	51.1	48.8
+ Single Euclidean Loss	✓		✓	54.0	48.2	53.1	51.8
with Decoupled Attention		✓	✓	54.7	47.3	55.3	52.4
Replace GDA with 2-SA	✓	✓	2-SA*	53.5	46.2	54.4	51.4
Full	✓	✓	✓	53.6 \uparrow 4.1	49.2 \uparrow 4.5	54.8 \uparrow 1.1	52.6 \uparrow 3.3

4.3 Ablation Study

Components Ablation. The ablation experiments, detailed in Table 4, involved training the model for a reduced duration of 24 epochs without dense prediction losses, to facilitate efficient analysis. These experiments confirm the significant role of GDA, which alone increases mAP by +1.9%. Moreover, applying the full suite of components results in a further mAP enhancement of +1.4%. A notable discovery is the detrimental effect on model performance when Euclidean Loss is applied without GDA, which leads to a 0.5% reduction in mAP. This reinforces our position that conventional self-attention mechanisms are insufficient for encoding a variety of geometric properties.

Effectiveness of GDA. In the proposed GDA (3.5), we decouple shape and relation learning, sequentially applying Euclidean shape and relation attention. In addition, we also attempt to intuitively double the self-attention layers in one block, referred to as “2-SA”. As shown in Table 4, GDA outperforms vanilla self-attention by +1.2% mAP, indicating the superiority of GDA when combined with Euclidean Loss.

4.4 Visualization Analysis

More Precise Shape Awareness. As shown in Figure 6(a), the shape geometry allows for an inherent description of the triangle feature, in contrast to baselines which do not discern it as effectively. A more irregularly shaped pedestrian crossing can be viewed in scenario (c) of Figure 6 and GeMap provides more precise construction than other methods.

Better Relation Awareness. With the help of Euclidean Relation Clues, GeMap can better infer parallelism and perpendicular in scenario (a) of Figure 6. More interestingly, in scenario (b), the strong sunlight makes the divider highlighted hard to recognize. However, according to the understanding of lane width patterns, GeMap has the sense that there should be a divider.

Alleviating Occlusion Issues. Figure 6(c) demonstrates that shape geometry alleviates occlusions in partially visible instances. It showcases a black car that obscures part of a pedestrian crossing, while, GeMap successfully deduces the

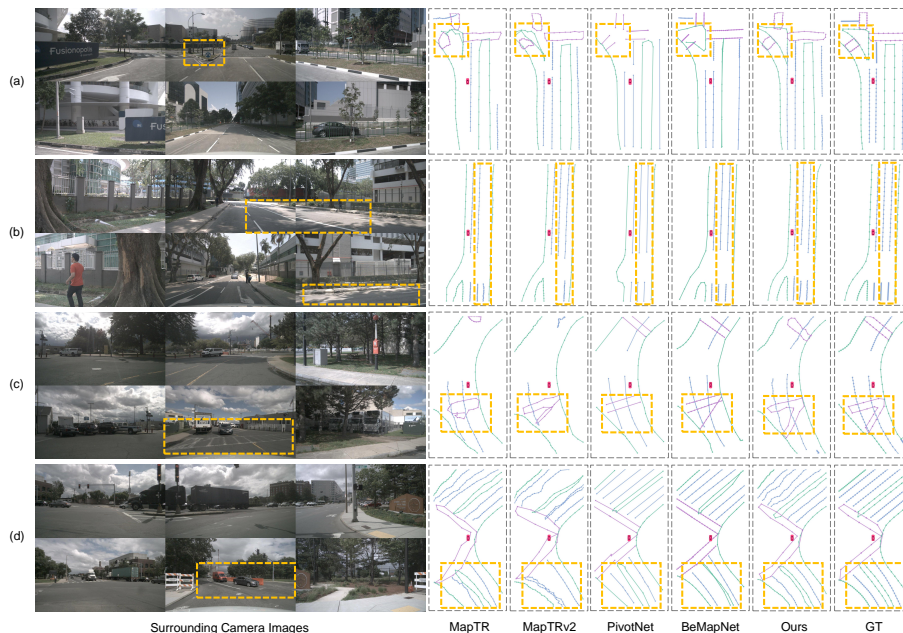


Fig. 6: Visualization results. Hard predictions are highlighted, such as complex triangular road boundary (a), strong sunlight (b), partially observed (c), and obscured situations (d). These challenging cases indicate the superiority and robustness of GeMap.

overall structure of the crossing within a complex shape. In contrast, these baseline models [4, 22, 23] struggle with such complex shape recovery under similar occlusion conditions.

Enhanced Robustness to Rotational Transformations. Figure 6(d) shows that the BEV map is highly tilted by turning, which is extensively obscured by vehicular traffic. GeMap more adeptly maintains the integrity of lane width and parallelism, which are key geometry. This underscores the superior robustness of our geometric constructs against rotational distortions.

4.5 Hyperparameter Experiments

Comparative Impact of Euclidean Losses. Figure 7a shows that GeMap exhibits larger sensitivity to the shape loss than to the relation loss, as evidenced by the steeper change of the performance curve. The optimal result is obtained when the shape and relation losses are relatively balanced by weights.

Optimization of Euclidean Loss Weighting. The performance of the model, as demonstrated in Figure 7b, peaks when the weight of the Euclidean Loss, λ , is set to 5×10^{-3} . Notably, model performance degrades with a too large or too small λ , which informed us to use a λ value of 5×10^{-3} for our experiments.

Distance Weighting. Experimental results of distance weighting are presented in Table 5. We change the order of distance (p) in Equation 5 and $p \in \{1, 2, 4\}$




Fig. 7: Hyper-parameters. The best performance is reached with balanced Euclidean shape loss and relation loss.

corresponds to “Linear”, “Square” and “4th Power” respectively. Moreover, we also try to treat instance pairs equally and it delivers the best performance.

Table 5: Distance weighting on the nuScenes dataset.

Strategy	$AP_{div}(\uparrow)$	$AP_{ped}(\uparrow)$	$AP_{bnd}(\uparrow)$	mAP(\uparrow)
Linear	53.5	46.6	56.2	52.1
Square	53.2	47.9	55.2	52.1
4th Power	53.2	47.4	55.8	52.1
Equal	53.6	49.2	54.8	52.6

5 Conclusion

In this paper, we realize significant shape and relation geometry inherent in HD map instances and propose the GeMap. GeMap includes the integration of Euclidean shape and relation losses for auxiliary supervision. To further refine the model’s awareness of diverse geometry, we introduce the Geometry-Decoupled Attention mechanism. GeMap has achieved state-of-the-art performances on both the nuScenes and Argoverse 2 datasets, underscoring its effectiveness. Despite these promising results, the current application of geometry remains fundamental, and future research could focus on more sophisticated representations or enhanced geometric patterns. Furthermore, the application of geometry extends beyond HD map construction, offering potential solutions to occlusion challenges in other autonomous driving tasks. We anticipate that these findings will inspire further research.

Acknowledgments. This work is partially supported by the National Natural Science Foundation of China (No. 62272045, No. 8326014), The Shun Hing Institute of Advanced Engineering (No. 8115074), and CUHK Direct Grants (No. 4055190).

References

1. Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11621–11631 (2020)
2. Chen, S., Cheng, T., Wang, X., Meng, W., Zhang, Q., Liu, W.: Efficient and robust 2d-to-bev representation learning via geometry-guided kernel transformer. arXiv preprint arXiv:2206.04584 (2022)
3. Deo, N., Wolff, E., Beijbom, O.: Multimodal trajectory prediction conditioned on lane-graph traversals. In: Conference on Robot Learning. pp. 203–212. PMLR (2022)
4. Ding, W., Qiao, L., Qiu, X., Zhang, C.: Pivotnet: Vectorized pivot learning for end-to-end hd map construction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3672–3682 (2023)
5. Ding, X., Zhang, Y., Ge, Y., Zhao, S., Song, L., Yue, X., Shan, Y.: Unireplknet: A universal perception large-kernel convnet for audio video point cloud time-series and image recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5513–5524 (2024)
6. Espinoza, J.L.V., Liniger, A., Schwarting, W., Rus, D., Van Gool, L.: Deep interactive motion prediction and planning: Playing games with motion prediction models. In: Learning for Dynamics and Control Conference. pp. 1006–1019. PMLR (2022)
7. Feng, Z., Guo, S., Tan, X., Xu, K., Wang, M., Ma, L.: Rethinking efficient lane detection via curve modeling. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17062–17070 (2022)
8. Gao, J., Sun, C., Zhao, H., Shen, Y., Anguelov, D., Li, C., Schmid, C.: Vectornet: Encoding hd maps and agent dynamics from vectorized representation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11525–11533 (2020)
9. Ge, C., Chen, J., Xie, E., Wang, Z., Hong, L., Lu, H., Li, Z., Luo, P.: Metabev: Solving sensor failures for 3d detection and map segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8721–8731 (2023)
10. Gu, J., Hu, C., Zhang, T., Chen, X., Wang, Y., Wang, Y., Zhao, H.: Vip3d: End-to-end visual trajectory prediction via 3d agent queries. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5496–5506 (2023)
11. Han, J., Gong, K., Zhang, Y., Wang, J., Zhang, K., Lin, D., Qiao, Y., Gao, P., Yue, X.: Onellm: One framework to align all modalities with language. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 26584–26595 (2024)
12. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
13. Huang, S., Shen, Z., Huang, Z., Ding, Z.h., Dai, J., Han, J., Wang, N., Liu, S.: Anchor3dlane: Learning to regress 3d anchors for monocular 3d lane detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17451–17460 (2023)

14. Jiao, J.: Machine learning assisted high-definition map creation. In: 2018 IEEE 42nd Annual Computer Software and Applications Conference (COMPSAC). vol. 1, pp. 367–373. IEEE (2018)
15. Lang, A.H., Vora, S., Caesar, H., Zhou, L., Yang, J., Beijbom, O.: Pointpillars: Fast encoders for object detection from point clouds. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 12697–12705 (2019)
16. Lee, Y., Hwang, J.w., Lee, S., Bae, Y., Park, J.: An energy and gpu-computation efficient backbone network for real-time object detection. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops. pp. 0–0 (2019)
17. Li, C., Shi, J., Wang, Y., Cheng, G.: Reconstruct from top view: A 3d lane detection approach based on geometry structure prior. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4370–4379 (2022)
18. Li, Q., Wang, Y., Wang, Y., Zhao, H.: Hdmapnet: An online hd map construction and evaluation framework. In: 2022 International Conference on Robotics and Automation (ICRA). pp. 4628–4634. IEEE (2022)
19. Li, Y., Ge, Z., Yu, G., Yang, J., Wang, Z., Shi, Y., Sun, J., Li, Z.: Bevdepth: Acquisition of reliable depth for multi-view 3d object detection. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 1477–1485 (2023)
20. Li, Z., Wang, W., Li, H., Xie, E., Sima, C., Lu, T., Qiao, Y., Dai, J.: Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. In: European conference on computer vision. pp. 1–18. Springer (2022)
21. Liang, M., Yang, B., Hu, R., Chen, Y., Liao, R., Feng, S., Urtasun, R.: Learning lane graph representations for motion forecasting. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16. pp. 541–556. Springer (2020)
22. Liao, B., Chen, S., Wang, X., Cheng, T., Zhang, Q., Liu, W., Huang, C.: MapTR: Structured modeling and learning for online vectorized HD map construction. In: The Eleventh International Conference on Learning Representations (2023), https://openreview.net/forum?id=k7p_YA07yE
23. Liao, B., Chen, S., Zhang, Y., Jiang, B., Zhang, Q., Liu, W., Huang, C., Wang, X.: Maptrv2: An end-to-end framework for online vectorized hd map construction. arXiv preprint arXiv:2308.05736 (2023)
24. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of the IEEE international conference on computer vision. pp. 2980–2988 (2017)
25. Liu, R., Chen, D., Liu, T., Xiong, Z., Yuan, Z.: Learning to predict 3d lane shape and camera pose from a single image via geometry constraints. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 36, pp. 1765–1772 (2022)
26. Liu, Y., Yuan, T., Wang, Y., Wang, Y., Zhao, H.: Vectormapnet: End-to-end vectorized hd map learning. In: International Conference on Machine Learning. pp. 22352–22369. PMLR (2023)
27. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 10012–10022 (2021)
28. Liu, Z., Tang, H., Amini, A., Yang, X., Mao, H., Rus, D.L., Han, S.: Bevfusion: Multi-task multi-sensor fusion with unified bird’s-eye view representation. In: 2023 IEEE International Conference on Robotics and Automation (ICRA). pp. 2774–2781. IEEE (2023)

29. Loshchilov, I., Hutter, F.: Sgdr: Stochastic gradient descent with warm restarts. arXiv preprint arXiv:1608.03983 (2016)
30. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
31. Loukkal, A., Grandvalet, Y., Drummond, T., Li, Y.: Driving among flatmobiles: Bird-eye-view occupancy grids from a monocular camera for holistic trajectory planning. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 51–60 (2021)
32. Lu, C., van de Molengraft, M.J.G., Dubbelman, G.: Monocular semantic occupancy grid mapping with convolutional variational encoder–decoder networks. *IEEE Robotics and Automation Letters* **4**(2), 445–452 (2019)
33. Mi, L., Zhao, H., Nash, C., Jin, X., Gao, J., Sun, C., Schmid, C., Shavit, N., Chai, Y., Anguelov, D.: Hdmapgen: A hierarchical graph generative model of high definition maps. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4227–4236 (2021)
34. Pan, B., Sun, J., Leung, H.Y.T., Andonian, A., Zhou, B.: Cross-view semantic segmentation for sensing surroundings. *IEEE Robotics and Automation Letters* **5**(3), 4867–4873 (2020)
35. Phillion, J., Fidler, S.: Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In: Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16. pp. 194–210. Springer (2020)
36. Qiao, L., Ding, W., Qiu, X., Zhang, C.: End-to-end vectorized hd-map construction with piecewise bezier curve. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13218–13228 (2023)
37. Scheel, O., Bergamini, L., Wolczyk, M., Osiński, B., Ondruska, P.: Urban driver: Learning to drive from real-world demonstrations using policy gradients. In: Conference on Robot Learning. pp. 718–728. PMLR (2022)
38. Tabelini, L., Berriel, R., Paixao, T.M., Badue, C., De Souza, A.F., Oliveira-Santos, T.: Keep your eyes on the lane: Real-time attention-guided lane detection. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 294–302 (2021)
39. Tabelini, L., Berriel, R., Paixao, T.M., Badue, C., De Souza, A.F., Oliveira-Santos, T.: Polylanenet: Lane estimation via deep polynomial regression. In: 2020 25th International Conference on Pattern Recognition (ICPR). pp. 6150–6156. IEEE (2021)
40. Tan, M., Le, Q.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: International conference on machine learning. pp. 6105–6114. PMLR (2019)
41. Van Gansbeke, W., De Brabandere, B., Neven, D., Proesmans, M., Van Gool, L.: End-to-end lane detection through differentiable least-squares fitting. In: Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops. pp. 0–0 (2019)
42. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
43. Wang, J., Ma, Y., Huang, S., Hui, T., Wang, F., Qian, C., Zhang, T.: A keypoint-based global association network for lane detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1392–1401 (2022)

44. Wilson, B., Qi, W., Agarwal, T., Lambert, J., Singh, J., Khandelwal, S., Pan, B., Kumar, R., Hartnett, A., Pontes, J.K., et al.: Argoverse 2: Next generation datasets for self-driving perception and forecasting. arXiv preprint arXiv:2301.00493 (2023)
45. Xie, Z., Pang, Z., Wang, Y.X.: Mv-map: Offboard hd-map generation with multi-view consistency. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8658–8668 (2023)
46. Yan, Y., Mao, Y., Li, B.: Second: Sparsely embedded convolutional detection. Sensors **18**(10), 3337 (2018)
47. Yang, W., Li, Q., Liu, W., Yu, Y., Ma, Y., He, S., Pan, J.: Projecting your view attentively: Monocular road scene layout estimation via cross-view transformation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 15536–15545 (2021)
48. Yuan, T., Liu, Y., Wang, Y., Wang, Y., Zhao, H.: Streammapnet: Streaming mapping network for vectorized online hd map construction. arXiv preprint arXiv:2308.12570 (2023)
49. Zhang, G., Lin, J., Wu, S., Luo, Z., Xue, Y., Lu, S., Wang, Z., et al.: Online map vectorization for autonomous driving: A rasterization perspective. Advances in Neural Information Processing Systems **36** (2024)
50. Zhang, Y., Ding, X., Gong, K., Ge, Y., Shan, Y., Yue, X.: Multimodal pathway: Improve transformers with irrelevant data from other modalities. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6108–6117 (2024)
51. Zhang, Y., Gong, K., Zhang, K., Li, H., Qiao, Y., Ouyang, W., Yue, X.: Meta-transformer: A unified framework for multimodal learning. arXiv preprint arXiv:2307.10802 (2023)
52. Zhang, Y., Li, H., Liu, J., Yue, X.: Explore the limits of omni-modal pretraining at scale. arXiv preprint arXiv:2406.09412 (2024)
53. Zhang, Z., Zhang, Y., Ding, X., Jin, F., Yue, X.: Online vectorized hd map construction using geometry. arXiv preprint arXiv:2312.03341 (2023)
54. Zhou, B., Krähenbühl, P.: Cross-view transformers for real-time map-view semantic segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 13760–13769 (2022)
55. Zhou, Z., Ye, L., Wang, J., Wu, K., Lu, K.: Hivt: Hierarchical vector transformer for multi-agent motion prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8823–8833 (2022)
56. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. arXiv preprint arXiv:2010.04159 (2020)