

MoST: Multi-modality Scene Tokenization for Motion Prediction

Norman Mu* Jingwei Ji* Zhenpei Yang* Nate Harada* Haotian Tang* Kan Chen*
 Charles R. Qi Runzhou Ge Kratarth Goel Zoey Yang Scott Ettinger
 Rami Al-Rfou Dragomir Anguelov Yin Zhou†
 Waymo LLC

Abstract

Many existing motion prediction approaches rely on symbolic perception outputs to generate agent trajectories, such as bounding boxes, road graph information and traffic lights. This symbolic representation is a high-level abstraction of the real world, which may render the motion prediction model vulnerable to perception errors (e.g., failures in detecting open-vocabulary obstacles) while missing salient information from the scene context (e.g., poor road conditions). An alternative paradigm is end-to-end learning from raw sensors. However, this approach suffers from the lack of interpretability and requires significantly more training resources. In this work, we propose tokenizing the visual world into a compact set of scene elements and then leveraging pre-trained image foundation models and LiDAR neural networks to encode all the scene elements in an open-vocabulary manner. The image foundation model enables our scene tokens to encode the general knowledge of the open world while the LiDAR neural network encodes geometry information. Our proposed representation can efficiently encode the multi-frame multi-modality observations with a few hundred tokens and is compatible with most transformer-based architectures. To evaluate our method, we have augmented Waymo Open Motion Dataset with camera embeddings. Experiments over Waymo Open Motion Dataset show that our approach leads to significant performance improvements over the state-of-the-art.

1. Introduction

In order to safely and effectively operate in complex environments, autonomous systems must model the behavior of nearby agents. These motion prediction models now often rely on symbolic perception outputs such as 3D bounding box tracks to represent agent states, rather than directly processing sensor inputs. These representations reduce input

*Equal contribution

†Corresponding author

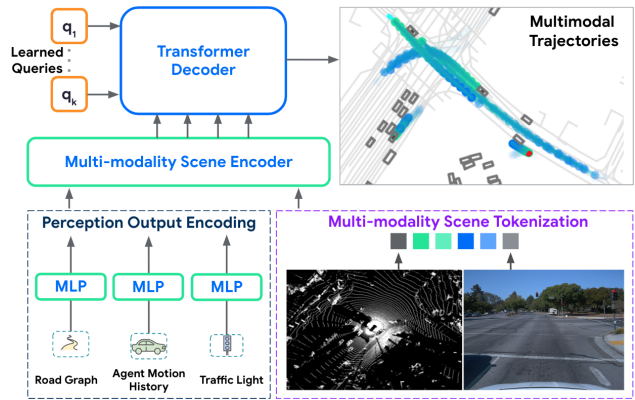


Figure 1. Overview of the proposed motion prediction paradigm. It fuses symbolic perception output and our multi-modality scene tokens. While symbolic representation offers a convenient world abstraction, the multi-modality scene tokens links behavior models directly to sensor observations via token embeddings.

dimensionality, facilitating computationally efficient model training. Additionally, since inputs such as 3D boxes are easily rearranged and manipulated, it is possible to construct many hypothetical scenarios leading to efficient simulation and testing. Yet in order to continue improving the accuracy and robustness of behavior models, it may be necessary to feed the models higher-fidelity sensor features. For instance, pedestrian pose and gaze offer richer cues than mere bounding boxes for motion prediction. Moreover, many scene elements like lane markings cannot be well represented by boxes. Furthermore, scene context (e.g., road surface conditions, hazardous locations) is difficult to characterize with symbolic representations. Manually crafting representation for diverse concepts demands considerable engineering effort in implementation, training, and evaluation. Instead, we want the behavior model to directly access the raw sensor data and determine what and how to encode.

Deep learning models’ performance generally improves when we replace hand-crafted features, designed to encode inductive bias according to expert domain knowledge, with the directly observed feature as long we scale compute and data accordingly. But learning to predict complex patterns

such as agent behavior directly from very high-dimensional sensor inputs (e.g. many high-resolution LiDAR and camera sensors all operating at high frequency) is an extremely challenging learning problem. It requires learning to organize many hundreds of thousands of points and pixels across time into meaningful representations. Moreover, the intermediate representations of fully end-to-end systems are far more difficult to validate and inspect.

Rather than choosing strictly between the two approaches, we instead propose combining existing symbolic representations with learned tokens encoding scene information. We first decompose the scene into a compact set of disjoint elements representing ground regions, perception-detected agents and open-set objects, based on ground plane fitting and connected component analysis. We then leverage large pre-trained 2D image models and 3D point cloud models to encode these scene elements into “tokens”. The 2D image models are trained on Internet-scale data, and show impressive capabilities in understanding the open visual world. These tokens encapsulate relevant information for reasoning about the environment, such as object semantics, object geometry as well as the scene context. We compactly represent multi-modality information about ground, agents and open-set objects into a few hundred tokens, which we later feed to Wayformer-like network [48] alongside tokens encoding agent position and velocity, road graph, and traffic signals. All tokens are processed via a linear projection into same dimension and self-attention layers.

To evaluate our method, we introduce camera embeddings to the Waymo Open Motion Dataset (WOMD) [16]. With LiDAR points [10] and camera embeddings, WOMD has become a large-scale multi-modal dataset for motion prediction. On the WOMD, our model, which combines learned and symbolic scene tokens, brings 6.6% relative improvement on soft mAP or 10.3% relative improvement on minADE. While we obtain the strongest results with the recently released image backbone from [36], other pre-trained image models [51, 55] also yield considerable gains. We further analyze the performance of our trajectory prediction model under challenging scenarios. Notably, we discover that even in the presence of imperfect symbolic perception outputs and incomplete road graph information, our model maintains exceptional robustness and accuracy.

Our contributions are three-fold:

- We have augmented WOMD into a large-scale multi-modal dataset to support research in end-to-end learning. Camera embeddings are released to the community.
- We have conducted a thorough study of modeling ideas of varying complexity to demonstrate the value of those sensory inputs in motion prediction.
- We have proposed a novel method MoST, which effectively leverages the multi-modality data and leads to significant performance improvement.

2. Related Works

Motion Prediction for Autonomous Driving The increasing interest in autonomous driving has led to a significant focus on motion prediction [8, 22, 48, 59, 60, 65]. Early methods [1, 4, 6, 7, 13, 20, 27, 47, 52] rasterize the input scene into a 2D image, followed by processing using convolutional neural networks (CNNs). However, as a result of the inherent lossiness in the rasterization process, contemporary research has shifted its focus towards representing road elements, such as object bounding boxes, road graphs, and traffic light signals, as discrete graph nodes [19]. These elements are then directly processed using graph neural networks (GNNs) [5, 21, 35, 39]. Another stream of research also employs this discrete set representation for scene elements but processes them using recurrent neural networks [2, 25, 47, 58, 63, 65], rather than GNNs. Thanks to the rapid advancement of transformer-based architectures in natural language processing and computer vision, the latest state-of-the-art motion predictors also extensively incorporate the attention mechanism [32, 48, 49, 59, 60]. More recently, the community has also started to study interactive behavior prediction, which jointly models the future motion of multiple objects [45, 61, 62, 64].

End-to-end Autonomous Driving The concept of end-to-end learning-based autonomous driving systems started in the late 1980s [53]. Since then, researchers have developed differentiable modules that connect perception and behavior [14, 18, 23, 28, 40, 44, 73], behavior and planning [24, 34, 41, 56, 61], or span from perception to planning [6, 29, 57, 71]. Building on the inspiration from [38, 72], Hu *et al.* introduced UniAD [30], which leverages transformer queries and a shared BEV feature map to facilitate end-to-end learning of perception, prediction, and planning [11, 31, 33]. More recently, there has been a growing interest in achieving end-to-end motion planning using large language models (LLMs) [46, 67, 69].

Challenges of Existing Methods While substantial advancements have been achieved in standard motion prediction benchmarks [3, 9, 10, 17, 68], the deployment of existing behavior models in real-world scenarios remains challenging. Many motion prediction models heavily rely on pre-processed, symbolic data from perception models [19, 54, 74, 75], and therefore are vulnerable to potential failures. Moreover, the manually-engineered interface greatly restrict the flexibility and scalability of the models in handling long-tail and novel categories of objects. In contrast, end-to-end learning [11, 26, 30, 31, 33] from raw sensors, while overcoming some limitations, encounters challenges in interpretability and scaling up batch size due to computational constraints.

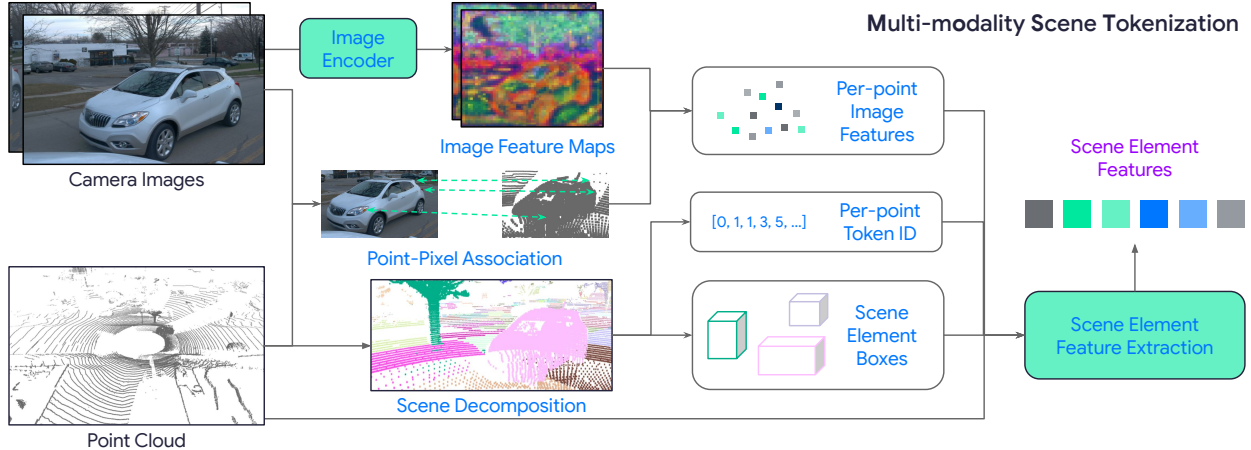


Figure 2. Overview of the proposed Multi-modality Scene Tokenization. Our method takes as input multi-view camera images and a full scene point cloud. We leverage a pre-trained image foundation model to obtain descriptive feature maps and decompose the scene into disjoint elements via clustering. Based on the sensor calibration information between camera and LiDAR, we obtain point-wise image features. From scene decomposition, we assign each point with a token/cluster id and derive box information for each element. Finally, we extract one feature embedding for each scene element.

3. Multi-modality Scene Tokenization

We propose a novel method, MoST (Multi-modality Scene Tokenization), to enrich the information fed to Transformer-based motion prediction models, by efficiently combining existing symbolic representations with scene tokens that encode multi-modality sensor information. In this section, we focus on how we obtain these scene tokens, each represented by a scene element feature enriched with semantic and geometric knowledge extracted from both image and LiDAR data. Figure 2 shows an overview of MoST.

3.1. Image Encoding and Point-Pixel Association

We start by extracting image feature maps for each camera and subsequently associating these features to the corresponding 3D LiDAR points using sensor calibration information. At each time step, we have a set of images $\{\mathbf{I}^k \in \mathbb{R}^{H_k \times W_k \times 3}\}_k$ captured by a total number of K cameras, where H_k and W_k represent the image dimensions. Additionally, we have a LiDAR point cloud $\mathbf{P}_{xyz} \in \mathbb{R}^{N_{pts} \times 3}$, with N_{pts} denoting the number of points. Using a pre-trained 2D image encoder E^{img} , we obtain a feature map of each image, denoted as $\{\mathbf{V}^k \in \mathbb{R}^{H'_k \times W'_k \times D}\}_k$. Subsequently, we leverage camera and LiDAR calibrations to establish a mapping between 3D LiDAR points and their corresponding 2D coordinates on the image feature map of size $H'_k \times W'_k$. This mapping associates each 3D point with the corresponding image feature vector. As a result, we obtain image features for all N_{pts} 3D points, represented as $\mathbf{F}_{pts} \in \mathbb{R}^{N_{pts} \times D}$. Note that for points projecting outside of any image plane, we set their image features as zeros and mark their image features as invalid.

To harness a wider range of knowledge, we utilize large

pre-trained image models trained on a diverse collections of datasets and tasks, capturing a richer understanding of the real world. We experiment with several image encoder candidates: SAM ViT-H [36], VQ-GAN [15], CLIP [55] and DINO v2 [51]. Different from others, VQ-GAN uses a codebook to build the feature map. To derive \mathbf{V}^k from VQ-GAN, we bottom-crop and partition each input image into multiple 256×256 patches. Subsequently, we extract 256 tokens from each patch and convert them into a 16×16 feature map through querying the codebook. Finally, these partial feature maps are stacked together according to their original spatial locations to produce \mathbf{V}^k .

3.2. Scene Decomposition

Next, our approach groups full scene LiDAR point cloud into three element types: **ground**, **agents**, and **open-set objects** (see illustration in Figure 3). We use the term “**scene element**” to denote the union of these three element types. We denote the number of elements of each type to be N_{elem}^{gnd} , N_{elem}^{agent} , $N_{elem}^{open-set}$ respectively, and we define $N_{elem} = N_{elem}^{gnd} + N_{elem}^{agent} + N_{elem}^{open-set}$ as the total number of scene elements.

- **Ground elements:** These are segmented blocks of the ground surface, obtained through either a dedicated ground point segmentation model or a simple RANSAC algorithm. Since the ground occupies a large area, we divide it into disjoint $10m \times 10m$ tiles, following [42].
- **Agent elements:** These correspond to the points within the bounding boxes of agents, detected by established perception pipelines for a pre-defined set of categories.
- **Open-set object elements:** These capture the remaining objects not included in the agent categories. Examples




Figure 3. Visualization of scene decomposition. We decompose a scene into agent elements, open-set elements and ground elements. We also visualize the perception bounding boxes for agents.

include novel categories of traffic participants and obstacles beyond the training data, long-tail instances that a perception model suppresses due to low confidence. We extract these elements by first removing ground and agent elements from the scene point cloud and then using connected component analysis to group points into instances.

Per-point Token ID Based on scene decomposition of each LiDAR frame, we can assign a unique token id to each LiDAR point. Points within the same scene element share one token id. With the point-pixel association, we can scatter each scene token ID to a set of camera pixels and/or locations on image feature maps. As a result, we can obtain features from both LiDAR and camera for each scene element. Based on point-wise token id, we can pool per-point image features into three sets of cluster-wise embedding vectors, *i.e.*, $\mathbf{F}_{\text{img}}^{\text{gnd}} \in \mathbb{R}^{N_{\text{elm}}^{\text{gnd}} \times D}$, $\mathbf{F}_{\text{img}}^{\text{agent}} \in \mathbb{R}^{N_{\text{elm}}^{\text{agent}} \times D}$, $\mathbf{F}_{\text{img}}^{\text{open-set}} \in \mathbb{R}^{N_{\text{elm}}^{\text{open-set}} \times D}$.

Scene Element Boxes We propose to encode each scene element with a combination of image features, coarse-grained geometry features, and fine-grained geometry features. Here we describe how we construct scene element boxes \mathbf{B} to represent coarse-grained geometry. For agent elements, coarse-grained geometry feature are derived from perception pipelines, capturing information of agent positions, sizes, and heading. For open-set object elements, we compute the tightest bounding boxes covering the point cluster, and these bounding boxes are also represented by box centers, box sizes and headings. For ground elements, we have divided the ground into fixed size tiles and simply use the tile center coordinates as position information. These box representations will be further encoded with a MLP and combined with image features and fine-grained features. We will dive into this combination in Sec 3.3.

3.3. Scene Element Feature Extraction

We finally extract scene element features with a neural network module. Multi-frame information are first compressed in an efficient way, then fed into this feature extraction module, which generate a single feature vector for each scene element. The feature extraction module is connected with the downstream Transformer-based motion prediction models, formulating an end-to-end trainable paradigm.

3.3.1 Efficient Multi-frame Data Representation

While we’ve compiled valuable information for each element within a single-frame scene – LiDAR points, per-point image features, and a bounding box – collecting this data across multiple frames leads to a large increase in memory usage. Considering that self-attention layers have quadratic complexity with respect to the number of tokens, naively concatenating tokens across all history frames will also lead to significantly increased memory usage. We propose an efficient data representation to reduce the amount of data sent to the model with the following three ingredients.

Open-set element tracking We compress the representation of open-set elements by associating open-set elements across frames using a simple Kalman Filter. For each open-set element, we only store its box information for all T frames with a tensor of shape $(N_{\text{elem}}^{\text{open-set}} \times T \times 7)$ and we apply average pooling across T frames of its image features resulting in an image feature tensor of shape $(N_{\text{elem}}^{\text{open-set}} \times 1 \times D)$.

Ground-element aggregation Instead of decomposing the ground into tiles for each frame, we apply decomposition after combining the ground points from all frames.

Cross-frame LiDAR downsampling Directly storing LiDAR points for all frames is computationally prohibitive. Simply downsampling LiDAR points in each frame still suffers from high redundancy over static parts of the scene. Therefore, we employ different downsampling schemes for ground elements (which are always static), and open-set/agent elements (which could be dynamic). For ground elements, we first merge the ground points across all frames then uniformly subsample to a fixed number $N_{\text{pts}}^{\text{gnd}}$. For open-set/agent elements, we subsample them to the fixed number $N_{\text{pts}}^{\text{open-set}}$ and $N_{\text{pts}}^{\text{agent}}$ respectively. The final LiDAR points $N_{\text{pts}} = N_{\text{pts}}^{\text{gnd}} + N_{\text{pts}}^{\text{agent}} + N_{\text{pts}}^{\text{open-set}}$. We also create a tensor $\mathbf{P}_{\text{ind}} \in \mathbb{R}^{N_{\text{pts}} \times 2}$ that stores the frame id and scene-element id for each point. Note that in this representation, the number of points from each frame is a variable, which is more efficient compared to storing a fixed number of points for all frames with padding.

3.3.2 Network Architecture

With the efficient multi-frame data representation, the inputs to the scene element feature extraction module are

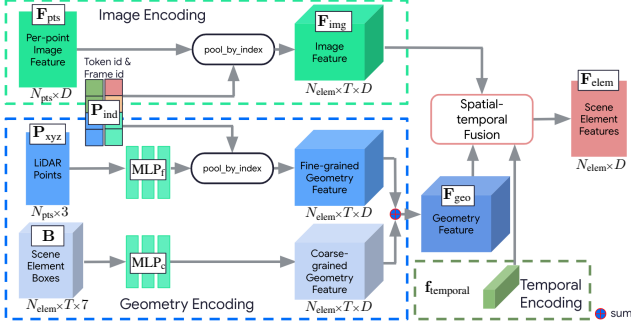


Figure 4. Scene element feature extraction. Scene-element feature is derived from a spatial-temporal module that fusing together image feature, geometry feature and temporal embedding. Image feature contains pooled feature from large pre-trained image encoder, and characterize the appearance and semantic attribute of the scene element. Geometry feature, on the other hand, characterizes the spatial location as well as the detailed geometry. Temporal information is injected through a learned temporal embedding.

summarized as following:

- $\mathbf{F}_{\text{pts}} \in \mathbb{R}^{N_{\text{pts}} \times D}$, point-wise image embeddings derived for all the LiDAR points across T frames.
- $\mathbf{B} \in \mathbb{R}^{N_{\text{elem}} \times T \times 7}$, bounding boxes of different scene elements across time, where $N_{\text{elem}} = N_{\text{elem}}^{\text{gnd}} + N_{\text{elem}}^{\text{agent}} + N_{\text{elem}}^{\text{open-set}}$. For ground elements, we only encode the tile center, leaving the rest four attributes as zeros.
- $\mathbf{P} = \{\mathbf{P}_{\text{xyz}}, \mathbf{P}_{\text{ind}}\}$, where $\mathbf{P}_{\text{xyz}} \in \mathbb{R}^{N_{\text{pts}} \times 3}$ collects multi-frame LiDAR points, and $\mathbf{P}_{\text{ind}} \in \mathbb{R}^{N_{\text{pts}} \times 2}$ stores frame id and token id for each point respectively.

For each tracked element across T time steps, our network (as shown in Figure 4) will process the previously listed multi-modality information into one embedding for each scene element, denoted as \mathbf{F}_{elem} .

As shown in the top branch of Figure 4, the network leverages \mathbf{P}_{ind} to group point-wise image embeddings \mathbf{F}_{pts} according to the token id and frame id, which results in the image feature tensor for all scene elements across frames, \mathbf{F}_{img} . In the bottom branch of Figure 4, we aim to derive geometry information \mathbf{F}_{geo} by encoding two pieces of information, *i.e.*, fine-grained geometry information from point clouds \mathbf{P}_{xyz} and coarse-grained shape information from 3D boxes \mathbf{B} . The fine-grained geometry is encoded by first mapping point xyz coordinates into a higher dimensional space and grouping high-dimensional features according to the token id and frame id. The coarse-grained shape encoding is derived by projecting box attributes to the same high dimensional space. Formally, \mathbf{f}_{geo} is defined as

$$\mathbf{f}_{\text{geo}}^i = \text{pool_by_index}(\text{MLP}_f(\mathbf{P}_{\text{xyz}}), \mathbf{P}_{\text{ind}})[i, :, :] + \text{MLP}_c(\mathbf{B})[i, :, :] \quad (1)$$

where i is the token id, function *pool_by_index* pools point-

	ViT-VQGAN [70]	SAM ViT-H [36]
Sensors	8 cameras (front, front left front right, side left, side right, rear left, rear right, rear) and LiDAR	
Temporal	1.0 s, 11 Frames	
Pre-trained Dataset	WebLi [12]	SA-1B [36]
Format	Token & Embedding	Embedding

Table 1. Details of the WOMD camera embeddings.

wise features based on token id and frame id.

Spatial-temporal Fusion Our spatial-temporal fusion module (Figure 4 right) takes as input the image feature \mathbf{F}_{img} , the geometry feature \mathbf{F}_{geo} , and a trainable temporal embedding $\mathbf{f}_{\text{temporal}} \in \mathbb{R}^{T \times D}$ that corresponds to T frames. It produces a temporally aggregated feature \mathbf{F}_{elem} for all scene elements. Under the hood, the spatial-temporal fusion module adds up the two input tensors, and then conducts axial attention across the temporal and element axes of the tensor, which is followed by the final average pooling across the temporal axis, as listed below:

$$\begin{aligned} \mathbf{F}_{\text{elem}} &\leftarrow \mathbf{F}_{\text{img}} + \mathbf{F}_{\text{geo}} + \mathbf{f}_{\text{temporal}} \in \mathbb{R}^{N_{\text{elem}} \times T \times D} \\ \mathbf{F}_{\text{elem}} &\leftarrow \text{AttnAlongAxis}(\mathbf{F}_{\text{elem}}, \text{axis} = \text{time}) \\ \mathbf{F}_{\text{elem}} &\leftarrow \text{AttnAlongAxis}(\mathbf{F}_{\text{elem}}, \text{axis} = \text{scene element}) \\ \mathbf{F}_{\text{elem}} &\leftarrow \text{mean}(\mathbf{F}_{\text{elem}}, \text{axis} = \text{time}) \end{aligned}$$

The final \mathbf{F}_{elem} uses a single vector to describe each scene element. This tensor can be fed as the additional inputs to scene encoding module of transformer-based motion prediction models, such as recently published [48, 60].

4. Experiments

4.1. The Release of WOMD Camera Embeddings

To advance research in sensor-based motion prediction, we have augmented Waymo Open Motion Dataset (WOMD) [16] with camera embeddings. WOMD contains the standard perception output, *e.g.* tracks of bounding boxes, road graph, traffic signals, and now it also includes synchronized LiDAR points [16] and camera embeddings. Given one scenario, a motion prediction model is required to reason about 1 second history data and generate predictions for the future 8 seconds at 5Hz. Our LiDAR can reach up to 75 meters along the radius and the cameras provide a multi-view imagery for the environment. WOMD characterize each perception-detected objects using a 3D bounding box (3D center point, heading, length, width, and height), and the object’s velocity vector. The road graph is provided as a set of polylines and polygons with semantic types. WOMD is divided into training, validation and testing subsets according to the ratio 70%, 15%, 15%. In

Method	Reference	Sensor	# Decoders	minADE↓	minFDE↓	Miss Rate↓	mAP↑	soft-mAP↑
MotionCNN [37]	CVPRW 2021	-	-	0.7383	1.4957	0.2072	0.2123	-
MultiPath++ [50]	ICRA 2022	-	-	0.978	2.305	0.440	-	-
SceneTransformer [50]	ICLR 2022	-	-	0.9700	2.0700	0.1867	0.2433	-
MTR [60]	NeurIPS 2022	-	-	0.6046	1.2251	0.1366	0.4164	-
Wayformer [48]	ICRA 2023	-	3	0.5512	1.1602	0.1208	0.4099	0.4247
MotionLM* [59]	ICCV 2023	-	1	0.5702	1.1653	0.1327	0.3902	0.4063
Wayformer	Reproduced	-	1	0.5830	1.2314	0.1347	0.3995	0.4110
MoST-SAM_H-6	Ours	C+L	1	0.5228	1.0764	0.1303	0.4040	0.4207
MoST-SAM_H-64	Ours	C+L	1	0.5487	1.1355	0.1238	0.4230	<u>0.4380</u>
Wayformer	Reproduced	-	3	0.5494	1.1386	<u>0.1190</u>	0.4052	0.4239
MoST-VQGAN-64	Ours	C	3	<u>0.5391</u>	<u>1.1099</u>	0.1172	<u>0.4201</u>	0.4396

Table 2. Performance comparison on WOMD validation set. MoST leads to significant performance gain to the Wayformer baselines and achieves state-of-the-art results in all compared metrics. MoST-SAM_H- $\{6, 64\}$: our method using SAM ViT-H feature and predicting based on 6 or 64 queries. MoST-VQGAN-64: our method using VQGAN feature with 64 queries. Bold font highlights the best result in each metric and underline denotes the second best. For methods with multiple decoders, results are based on ensembling of predictions. MotionLM* is based on contacting authors for their 1 decoder results, which was not reported in the original publication.

this paper, we report results over the validation set. We’ll reserve the test set for future community benchmarking.

Due to the data storage issue and risk of leakage of sensitive information (e.g., human faces, car plate numbers, etc.), we will not release the raw camera images. Instead, the released multi-modality dataset will be in two formats:

- **ViT-VQGAN Tokens and Embeddings:** We apply a pre-trained ViT-VQGAN [70] to extract tokens and embeddings for each camera image. The number of tokens per camera is 512, where each token corresponds to a 32 dimensional embedding in the quantized codebook.
- **SAM ViT-H Embeddings:** We apply a pre-trained SAM ViT-H [36] model to extract dense embeddings for each camera image. We release the per-scene-element embedding vectors, each being 256 dimensional.

In the released dataset, we have 1 LiDAR and 8 cameras (front, front-left, front-right, side-left, side-right, rear-left, rear-right, rear). Please see details in Table 1.

Task and Metrics Based on the augmented WOMD, we investigate the standard marginal motion prediction task, where a model is required to generate 6 mostly likely future trajectories for each of the agents independently of other agents futures. We report results for various methods under commonly adopted metrics, namely minADE, minFDE, miss rate, mAP and soft-mAP [16]. For fair comparison, we only compare results based on single model prediction.

4.2. Experimental Results

Our MoST is a general paradigm applicable to most transformer-based motion prediction architectures. Without losing generality, we adopt a state-of-the-art architecture, Wayformer [48], as our motion prediction backbone and we

augment it with our new design by fusing multi-modality tokens. In the following sections, we use Wayformer as the baseline and show the performance improvement by MoST. Please refer to appendix for implementation details.

4.2.1 Baseline Comparison

In Table 2, we evaluate the proposed approach and compare it with recently published models, *i.e.*, MTR [60], Wayformer [48], MultiPath++ [66], MotionCNN [37], MotionLM [59], SceneTransformer [50]. Specifically, we study our approach in two settings, 1) using LiDAR + camera tokens with single decoder and 2) using camera tokens with 3 decoders. During inference, we fit a Gaussian Mixture Model by merging predictions from the decoder(s) and draw 2048 samples, which are finally aggregated into 6 trajectories through K-means [48]. In both settings, the introduction of sensory tokens leads to a clear performance gain over the corresponding Wayformer baselines based on our re-implementation. Moreover, our approach achieves state-of-the-art performance across various metrics.

Figure 5 illustrates two comparisons between our MoST and the baseline model. The upper example shows that with tokenized sensor information, our MoST rules out the possibility that a vehicle runs onto walls after a U-turn. The lower example makes a prediction that a cyclist may cross the street which is safety critical for the autonomous vehicle to take precaution regarding this behavior.

4.2.2 Ablation Study

We find that applying MoST to the current frame can also lead to significant improvement over the baseline. For ef-




Figure 5. Qualitative comparison. The agent boxes are colored by their types: gray for vehicle, red for pedestrian, and cyan for cyclist. The predicted trajectories are ordered temporally from green to blue. For each modeled agent, the models predict 6 trajectory candidates, whose confidence scores are illustrated by transparency: the more confident, the more visible. Ground truth trajectory is shown as red dots. In the upper example, MoST rules out the possibility that a vehicle runs onto a wall after U-turn; in the lower example, MoST correctly predicts that a cyclist could suddenly cross the street.

Image Encoder	minADE↓	mAP↑	soft-mAP↑
DINO-v2 [51]	0.5597	0.4154	0.4285
CLIP [55]	0.5590	0.4138	0.4272
VQ-GAN [15]	0.5670	0.4058	0.4192
SAM ViT-H [36]	0.5483	0.4162	0.4321

Table 3. Ablation study of different image features. All these image features improves the motion prediction performance, while we observe SAM ViT-H [36] leads to the most improvement. We uses single-frame multi-modal feature for these study.

efficient experimentation, we perform ablation study by employing single frame MoST and SAM ViT-H feature.

Effects of Different Pre-trained Image Encoders To investigate different choices of the image encoder for our model, we have conducted experiments comparing the performance of using image feature encoders from various pre-trained models: SAM ViT-H [36], CLIP [55], DINO v2 [51], and VQ-GAN [15]. As show in Table 3, the SAM ViT-H encoder yields the highest performance across all behavior prediction metrics. We hypothesize that this performance advantage likely stems from SAM’s strong capability to extract comprehensive and spatially faithful feature

Open-set	Agent	Ground	M-frame	minADE↓	soft-mAP↑
✗	✓	✓	✗	0.5654	0.4112
✓	✗	✓	✗	0.5520	0.4273
✓	✓	✗	✗	0.5514	0.4241
✓	✓	✓	✗	0.5483	0.4321
✓	✓	✓	✓	0.5487	0.4380

Table 4. Ablation study on how different scene element affects the performance. The first four rows shows that all types of scene element brings benefits to the model. The last row shows that aggregating scene element across frame considerably improves the soft-mAP, though leads to slight regression of minADE.

Perception	Camera	LiDAR	minADE↓	mAP↑	soft-mAP↑
✓	✗	✗	0.5830	0.3995	0.4110
✓	✓	✗	0.5483	0.4118	0.4265
✓	✗	✓	0.5486	0.4040	0.4212
✓	✓	✓	0.5483	0.4162	0.4321

Table 5. Ablation study on different input modality. The first row corresponds to the setting of Wayformer [48]. The second row adds camera image feature. The last row further adds LiDAR. We can see using both Camera and LiDAR yields to the best results.

Sensory Token	minADE↓	mAP↑	soft-mAP↑
None (Wayformer)	0.5830	0.3995	0.4110
Image-grid Token (Ours)	0.5495	0.4109	0.4261
Scene Cluster Token (Ours)	0.5483	0.4162	0.4321

Table 6. Comparing variant of scene tokenization strategy with single-frame sensor data. Both token strategy leads to improvement the vanilla Wayformer [48], which does not use sensor data.

maps, as it is trained with a large and diverse dataset for the dense understanding task of image segmentation. Other large pre-trained image models also demonstrate notable capability, outperforming the Wayformer baseline on mAP and soft mAP, albeit inferior to SAM.

Ablation on Input Modality To understand how different input modalities affect the final model performance, we conduct ablation experiments and summarize results in Table 5 where we remove image feature or remove LiDAR feature of our single-frame model. We can see image feature and LiDAR feature are both beneficial, and combining both modality leads to the biggest improvement.

Ablation on Scene Element To gain deeper insights into the contribution of each type of scene element, we conduct ablation studies in Table 4 by removing specific element types and evaluate the impact on behavior prediction metrics. Combining all types of scene elements leads to the best soft-mAP metric. Note that only associating image features

Method	minADE↓	mAP↑	soft-mAP↑
Wayformer [48]	0.9002	0.2312	0.2382
MoST-SAM_H-64	0.8720	0.2615	0.2677

Table 7. Evaluation on hard scenarios. We curate a set of hard scenarios based on the performance of MoST and Wayformer on them. MoST consistently shows improved performance.

to agents give the smallest improvement. We hypothesize the reasons to be two-fold: 1) in most cases, the agent box is sufficient to characterize the motion of the object; 2) there are only a handful of agents in the scene and very few image features are included in the model.

Alternative Scene Tokenizer We use SAM ViT-H in this experiment. We design another baseline tokenizer, denoted as Image-grid token, which tokenizes each image feature as $16 \times 16 = 256$ image embeddings by subsampling 4X along column and row axes. The feature from all camera images are flattened and concatenated to form scene tokens. In Table 6 we can see the Image-grid tokenizer also leads to improvement compared to Wayformer baseline, though inferior to our cluster-based sparse tokenizer which utilizes point cloud to derive accurate depth information and leverages the intrinsic scene sparsity to get compact tokens.

4.2.3 Evaluation on Challenging Scenarios

While the improvement shown above demonstrates overall improvement across all driving scenarios, we are also interested in investigating the performance gain in the most challenging cases. Here we present how our model performs in challenging scenarios, specifically on (a) a mined set of hard scenarios, (b) situations where perception failures happen, and (c) situations where roadgraph is inaccurate.

Mined Hard Scenarios To assess the effectiveness of our method in complex situations, we have curated a set of hard scenarios. We conduct a per-scenario evaluation throughout the entire validation set, identifying the 1000 scenarios with the lowest minADE across vehicle, pedestrian, and cyclist categories for the baseline and MoST-SAM_H-64, respectively. In this way, we ensure the mining is symmetric and fair for both methods. Then we combine these 6000 scenarios, resulting in 4024 unique scenarios, forming our curated challenging evaluation dataset. As shown in Table 7, MoST demonstrates more pronounced relative improvement in mAP and soft-mAP, *i.e.*, 13.1% and 12.4% respectively, compared to the baseline in these hardest scenarios, confirming its effectiveness of enhanced robustness and resilience in complex situations. We also find that improving minADE in these hard scenarios is still a challenge.

Perception Failure Most motion prediction algorithms [48, 59] assume accurate perception object boxes as inputs. It is critical to understand how such a system will perform when this assumption breaks due to various reasons, such as long-

Failure type	Failure rate	minADE↓	soft-mAP↑
None	0%	0.5515	0.4396
	10%	0.5560	0.4302
	30%	0.5625	0.4235
Perception	50%	0.5712	0.4164
	10%	0.5647	0.4217
	30%	0.6020	0.4010
Roadgraph	50%	0.6707	0.3499

Table 8. Evaluation on simulated perception and roadgraph failure. We vary the ratio of miss detected agent boxes and miss detected roadgraph segments in scenes as 10%, 30% and 50%, respectively. With multi-modal features, MoST performs on par with baselines even with 50% perception or 30% roadgraph failure.

tail and novel categories of object beyond training supervision, occlusion, long-range, etc. Thus, we propose to additionally evaluate our method against the baseline method in the case of perception failure. Concretely, we simulate perception failure of not detecting certain object boxes by randomly removing agents according to a fixed ratio of agents in the scene. The boxes dropped out are consistent for MoST-SAM_H-64 and the baseline. As shown in Table 8, MoST shows robustness against perception failures: even when failure rate raises to 50%, our model still performs on par with the Wayformer baseline (soft mAP 0.4121).

Roadgraph Failure Motion prediction models often exhibit a strong reliance on roadgraphs, leading to potential vulnerabilities in situations where the roadgraph is incomplete or inaccurate. Our proposed model, MoST, tackles this issue by incorporating multi-modality scene tokens as additional inputs, thereby enhancing its robustness against roadgraph failures. We demonstrate this advantage by simulating various levels of roadgraph errors, similar to the aforementioned perception failure simulation. Specifically, we evaluate MoST-SAM_H-64 under scenarios with 10%, 30%, and 50% missing roadgraph segments in the validation set. Notably, as showcased in Table 8, even with a 30% of the roadgraph missing, MoST performs on par with baseline models that assume perfect roadgraph information.

5. Conclusions

To promote sensor-based motion prediction research, we have enhanced WOMD with camera embeddings, making it a large-scale multi-modal dataset for benchmarking. To efficiently integrate multi-modal sensor signals into motion prediction, we propose a method that represents the multi-frame scenes as a set of scene elements and leverages large pre-trained image encoders and 3D point cloud networks to encode rich semantic and geometric information for each element. We demonstrate that our approach leads to significant improvements in motion prediction task.

References

- [1] Yuriy Biktairov, Maxim Stebelev, Irina Rudenko, Oleh Shli-azhko, and Boris Yangel. Prank: motion prediction based on ranking. *Advances in neural information processing systems*, 33:2553–2563, 2020. 2
- [2] Thibault Buhet, Emilie Wirbel, Andrei Bursuc, and Xavier Perrotton. Plop: Probabilistic polynomial objects trajectory planning for autonomous driving. *arXiv preprint arXiv:2003.08744*, 2020. 2
- [3] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *CVPR*, 2020. 2
- [4] Sergio Casas, Wenjie Luo, and Raquel Urtasun. Intentnet: Learning to predict intention from raw sensor data. In *CoRL*, 2018. 2
- [5] Sergio Casas, Cole Gulino, Renjie Liao, and Raquel Urtasun. Spaggn: Spatially-aware graph neural networks for relational behavior forecasting from sensor data. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9491–9497. IEEE, 2020. 2
- [6] Sergio Casas, Abbas Sadat, and Raquel Urtasun. Mp3: A unified model to map, perceive, predict and plan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14403–14412, 2021. 2
- [7] Yuning Chai, Benjamin Sapp, Mayank Bansal, and Dragomir Anguelov. Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction. In *CoRL*, 2019. 2
- [8] Yuning Chai, Benjamin Sapp, Mayank Bansal, and Dragomir Anguelov. Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction. In *CoRL*, 2019. 2
- [9] Ming-Fang Chang, John Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew Hartnett, De Wang, Peter Carr, Simon Lucey, Deva Ramanan, et al. Argoverse: 3d tracking and forecasting with rich maps. In *CVPR*, 2019. 2
- [10] Kan Chen, Runzhou Ge, Hang Qiu, Rami Al-Rfou, Charles R Qi, Xuanyu Zhou, Zoey Yang, Scott Ettinger, Pei Sun, Zhaoqi Leng, et al. Womd-lidar: Raw sensor dataset benchmark for motion forecasting. *arXiv preprint arXiv:2304.03834*, 2023. 2
- [11] Li Chen, Penghao Wu, Kashyap Chitta, Bernhard Jaeger, Andreas Geiger, and Hongyang Li. End-to-end autonomous driving: Challenges and frontiers. *arXiv preprint arXiv:2306.16927*, 2023. 2
- [12] Xi Chen, Xiao Wang, Soravit Changpinyo, AJ Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyer, et al. Pali: A jointly-scaled multilingual language-image model. *arXiv preprint arXiv:2209.06794*, 2022. 5
- [13] Henggang Cui, Vladan Radosavljevic, Fang-Chieh Chou, Tsung-Han Lin, Thi Nguyen, Tzu-Kuo Huang, Jeff Schneider, and Nemanja Djuric. Multimodal trajectory predictions for autonomous driving using deep convolutional networks. In *ICRA*, 2019. 2
- [14] Nemanja Djuric, Henggang Cui, Zhaoen Su, Shangxuan Wu, Huahua Wang, Fang-Chieh Chou, Luisa San Martin, Song Feng, Rui Hu, Yang Xu, et al. Multixnet: Multiclass multistage multimodal motion prediction. In *2021 IEEE Intelligent Vehicles Symposium (IV)*, pages 435–442. IEEE, 2021. 2
- [15] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12873–12883, 2021. 3, 7
- [16] Scott Ettinger, Shuyang Cheng, Benjamin Caine, Chenxi Liu, Hang Zhao, Sabeek Pradhan, Yuning Chai, Ben Sapp, Charles R. Qi, Yin Zhou, Zoey Yang, Aurélien Chouard, Pei Sun, Jiquan Ngiam, Vijay Vasudevan, Alexander McCauley, Jonathon Shlens, and Dragomir Anguelov. Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset. In *ICCV*, 2021. 2, 5, 6
- [17] Scott Ettinger, Shuyang Cheng, Benjamin Caine, Chenxi Liu, Hang Zhao, Sabeek Pradhan, Yuning Chai, Ben Sapp, Charles R Qi, Yin Zhou, et al. Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset. In *ICCV*, 2021. 2
- [18] Sudeep Fadadu, Shreyash Pandey, Darshan Hegde, Yi Shi, Fang-Chieh Chou, Nemanja Djuric, and Carlos Vallespi-Gonzalez. Multi-view fusion of sensor data for improved perception and prediction in autonomous driving. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2349–2357, 2022. 2
- [19] Jiyang Gao, Chen Sun, Hang Zhao, Yi Shen, Dragomir Anguelov, Congcong Li, and Cordelia Schmid. Vectornet: Encoding hd maps and agent dynamics from vectorized representation. In *CVPR*, 2020. 2
- [20] Thomas Gilles, Stefano Sabatini, Dzmityr Tsishkou, Bogdan Stanculescu, and Fabien Moutarde. Home: Heatmap output for future motion estimation. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, pages 500–507. IEEE, 2021. 2
- [21] Thomas Gilles, Stefano Sabatini, Dzmityr Tsishkou, Bogdan Stanculescu, and Fabien Moutarde. Gohome: Graph-oriented heatmap output for future motion estimation. In *2022 international conference on robotics and automation (ICRA)*, pages 9107–9114. IEEE, 2022. 2
- [22] Junru Gu, Chen Sun, and Hang Zhao. Densetnt: End-to-end trajectory prediction from dense goal sets. In *ICCV*, 2021. 2
- [23] Junru Gu, Chenxu Hu, Tianyuan Zhang, Xuanyao Chen, Yilun Wang, Yue Wang, and Hang Zhao. Vip3d: End-to-end visual trajectory prediction via 3d agent queries. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5496–5506, 2023. 2
- [24] Cole Gulino, Justin Fu, Wenjie Luo, George Tucker, Eli Bronstein, Yiren Lu, Jean Harb, Xinlei Pan, Yan Wang, Xianguyu Chen, et al. Waymax: An accelerated, data-driven simulator for large-scale autonomous driving research. *arXiv preprint arXiv:2310.08710*, 2023. 2
- [25] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. Social gan: Socially acceptable trajectories with generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2255–2264, 2018. 2

- [26] Steffen Hagedorn, Marcel Hallgarten, Martin Stoll, and Alexandru Condurache. Rethinking integration of prediction and planning in deep learning-based automated driving systems: A review. *arXiv preprint arXiv:2308.05731*, 2023. **2**
- [27] Joey Hong, Benjamin Sapp, and James Philbin. Rules of the road: Predicting driving behavior with a convolutional model of semantic interactions. In *CVPR*, 2019. **2**
- [28] Anthony Hu, Zak Murez, Nikhil Mohan, Sofia Dudas, Jeffrey Hawke, Vijay Badrinarayanan, Roberto Cipolla, and Alex Kendall. Fiery: Future instance prediction in bird’s-eye view from surround monocular cameras. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15273–15282, 2021. **2**
- [29] Shengchao Hu, Li Chen, Penghao Wu, Hongyang Li, Junchi Yan, and Dacheng Tao. St-p3: End-to-end vision-based autonomous driving via spatial-temporal feature learning. In *European Conference on Computer Vision*, pages 533–549. Springer, 2022. **2**
- [30] Yihan Hu, Jiazhi Yang, Li Chen, Keyu Li, Chonghao Sima, Xizhou Zhu, Siqi Chai, Senyao Du, Tianwei Lin, Wenhai Wang, et al. Planning-oriented autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17853–17862, 2023. **2**
- [31] Xiaosong Jia, Yulu Gao, Li Chen, Junchi Yan, Patrick Langechuan Liu, and Hongyang Li. Driveadapter: Breaking the coupling barrier of perception and planning in end-to-end autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7953–7963, 2023. **2**
- [32] Xiaosong Jia, Penghao Wu, Li Chen, Yu Liu, Hongyang Li, and Junchi Yan. Hdgt: Heterogeneous driving graph transformer for multi-agent trajectory prediction via scene encoding. *IEEE transactions on pattern analysis and machine intelligence*, 2023. **2**
- [33] Xiaosong Jia, Penghao Wu, Li Chen, Jiangwei Xie, Conghui He, Junchi Yan, and Hongyang Li. Think twice before driving: Towards scalable decoders for end-to-end autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21983–21994, 2023. **2**
- [34] Alexey Kamenev, Lirui Wang, Ollin Boer Bohan, Ishwar Kulkarni, Bilal Kartal, Artem Molchanov, Stan Birchfield, David Nistér, and Nikolai Smolyanskiy. Predictionnet: Real-time joint probabilistic traffic prediction for planning, control, and simulation. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 8936–8942. IEEE, 2022. **2**
- [35] Siddhesh Khandelwal, William Qi, Jagjeet Singh, Andrew Hartnett, and Deva Ramanan. What-if motion prediction for autonomous driving. *arXiv preprint arXiv:2008.10587*, 2020. **2**
- [36] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023. **2, 3, 5, 6, 7, 13**
- [37] Stepan Konev, Kirill Brodt, and Artsiom Sanakoyeu. Motioncnn: A strong baseline for motion prediction in autonomous driving. In *CVPRW*, 2021. **6**
- [38] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Yu Qiao, and Jifeng Dai. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. In *European conference on computer vision*, pages 1–18. Springer, 2022. **2**
- [39] Ming Liang, Bin Yang, Rui Hu, Yun Chen, Renjie Liao, Song Feng, and Raquel Urtasun. Learning lane graph representations for motion forecasting. In *ECCV*, pages 541–556. Springer, 2020. **2**
- [40] Ming Liang, Bin Yang, Wenyuan Zeng, Yun Chen, Rui Hu, Sergio Casas, and Raquel Urtasun. Pnpnet: End-to-end perception and prediction with tracking in the loop. In *CVPR*, 2020. **2**
- [41] Jerry Liu, Wenyuan Zeng, Raquel Urtasun, and Ersin Yumer. Deep structured reactive planning. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4897–4904. IEEE, 2021. **2**
- [42] Minghua Liu, Yin Zhou, Charles R Qi, Boqing Gong, Hao Su, and Dragomir Anguelov. Less: Label-efficient semantic segmentation for lidar point clouds. In *European conference on computer vision*, pages 70–89. Springer, 2022. **3**
- [43] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2017. **14**
- [44] Wenjie Luo, Bin Yang, and Raquel Urtasun. Fast and furious: Real time end-to-end 3d detection, tracking and motion forecasting with a single convolutional net. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 3569–3577, 2018. **2**
- [45] Wenjie Luo, Cheol Park, Andre Cornman, Benjamin Sapp, and Dragomir Anguelov. Jfp: Joint future prediction with interactive multi-agent modeling for autonomous driving. In *Conference on Robot Learning*, pages 1457–1467. PMLR, 2023. **2**
- [46] Jiageng Mao, Yuxi Qian, Hang Zhao, and Yue Wang. Gpt-driver: Learning to drive with gpt. *arXiv preprint arXiv:2310.01415*, 2023. **2**
- [47] Francesco Marchetti, Federico Becattini, Lorenzo Seidenari, and Alberto Del Bimbo. Mantra: Memory augmented networks for multiple trajectory prediction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7143–7152, 2020. **2**
- [48] Nigamaa Nayakanti, Rami Al-Rfou, Aurick Zhou, Kratarth Goel, Khaled S Refaat, and Benjamin Sapp. Wayformer: Motion forecasting via simple & efficient attention networks. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2980–2987. IEEE, 2023. **2, 5, 6, 7, 8, 13**
- [49] Jiquan Ngiam, Benjamin Caine, Vijay Vasudevan, Zhengdong Zhang, Hao-Tien Lewis Chiang, Jeffrey Ling, Rebecca Roelofs, Alex Bewley, Chenxi Liu, Ashish Venugopal, et al. Scene transformer: A unified architecture for predicting multiple agent trajectories. *arXiv preprint arXiv:2106.08417*, 2021. **2**

- [50] Jiquan Ngiam, Benjamin Caine, Vijay Vasudevan, Zhengdong Zhang, Hao-Tien Lewis Chiang, Jeffrey Ling, Rebecca Roelofs, Alex Bewley, Chenxi Liu, Ashish Venugopal, David Weiss, Ben Sapp, Zhifeng Chen, and Jonathon Shlens. Scene transformer: A unified architecture for predicting multiple agent trajectories. In *ICLR*, 2022. 6
- [51] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 2, 3, 7
- [52] Seong Hyeon Park, Gyubok Lee, Jimin Seo, Manoj Bhat, Minseok Kang, Jonathan Francis, Ashwin Jadhav, Paul Pu Liang, and Louis-Philippe Morency. Diverse and admissible trajectory forecasting through multimodal context understanding. In *ECCV*, pages 282–298. Springer, 2020. 2
- [53] Dean A Pomerleau. Alvin: An autonomous land vehicle in a neural network. *Advances in neural information processing systems*, 1, 1988. 2
- [54] Charles R Qi, Yin Zhou, Mahyar Najibi, Pei Sun, Khoa Vo, Boyang Deng, and Dragomir Anguelov. Offboard 3d object detection from point cloud sequences. In *CVPR*, 2021. 2
- [55] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 2, 3, 7
- [56] Nicholas Rhinehart, Rowan McAllister, Kris Kitani, and Sergey Levine. Precog: Prediction conditioned on goals in visual multi-agent settings. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2821–2830, 2019. 2
- [57] Abbas Sadat, Sergio Casas, Mengye Ren, Xinyu Wu, Pranaab Dhawan, and Raquel Urtasun. Perceive, predict, and plan: Safe motion planning through interpretable semantic representations. In *ECCV*, pages 414–430. Springer, 2020. 2
- [58] Tim Salzmann, Boris Ivanovic, Punarjay Chakravarty, and Marco Pavone. Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data. In *ECCV*, pages 683–700. Springer, 2020. 2
- [59] Ari Seff, Brian Cera, Dian Chen, Mason Ng, Aurick Zhou, Nigamaa Nayakanti, Khaled S Refaat, Rami Al-Rfou, and Benjamin Sapp. Motionlm: Multi-agent motion forecasting as language modeling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8579–8590, 2023. 2, 6, 8
- [60] Shaoshuai Shi, Li Jiang, Dengxin Dai, and Bernt Schiele. Motion transformer with global intention localization and local movement refinement. *Advances in Neural Information Processing Systems*, 35:6531–6543, 2022. 2, 5, 6
- [61] Haoran Song, Wenchao Ding, Yuxuan Chen, Shaojie Shen, Michael Yu Wang, and Qifeng Chen. Pip: Planning-informed trajectory prediction for autonomous driving. In *ECCV*, pages 598–614. Springer, 2020. 2
- [62] Qiao Sun, Xin Huang, Junru Gu, Brian C Williams, and Hang Zhao. M2i: From factored marginal trajectory prediction to interactive prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6543–6552, 2022. 2
- [63] Charlie Tang and Russ R Salakhutdinov. Multiple futures prediction. *Advances in neural information processing systems*, 32, 2019. 2
- [64] Ekaterina Tolstaya, Reza Mahjourian, Carlton Downey, Balakrishnan Vadarajan, Benjamin Sapp, and Dragomir Anguelov. Identifying driver interactions via conditional behavior prediction. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3473–3479. IEEE, 2021. 2
- [65] Balakrishnan Varadarajan, Ahmed Hefny, Avikalp Srivastava, Khaled S. Refaat, Nigamaa Nayakanti, Andre Cornman, Kan Chen, Bertrand Douillard, Chi-Pang Lam, Dragomir Anguelov, and Benjamin Sapp. Multipath++: Efficient information fusion and trajectory aggregation for behavior prediction. *CoRR*, abs/2111.14973, 2021. 2
- [66] Balakrishnan Varadarajan, Ahmed Hefny, Avikalp Srivastava, Khaled S. Refaat, Nigamaa Nayakanti, Andre Cornman, Kan Chen, Bertrand Douillard, Chi Pang Lam, Dragomir Anguelov, and Benjamin Sapp. Multipath++: Efficient information fusion and trajectory aggregation for behavior prediction. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 7814–7821, 2022. 6
- [67] Pengqin Wang, Meixin Zhu, Hongliang Lu, Hui Zhong, Xianda Chen, Shaojie Shen, Xuesong Wang, and Yin Hai Wang. Bvgpt: Generative pre-trained large model for autonomous driving prediction, decision-making, and planning. *arXiv preprint arXiv:2310.10357*, 2023. 2
- [68] Benjamin Wilson, William Qi, Tanmay Agarwal, John Lambert, Jagjeet Singh, Siddhesh Khandelwal, Bowen Pan, Ratnesh Kumar, Andrew Hartnett, Jhony Kaesemodel Pontes, et al. Argoverse 2: Next generation datasets for self-driving perception and forecasting. *arXiv preprint arXiv:2301.00493*, 2023. 2
- [69] Zhenhua Xu, Yujia Zhang, Enze Xie, Zhen Zhao, Yong Guo, Kenneth KY Wong, Zhenguo Li, and Hengshuang Zhao. Drivegpt4: Interpretable end-to-end autonomous driving via large language model. *arXiv preprint arXiv:2310.01412*, 2023. 2
- [70] Jiahui Yu, Xin Li, Jing Yu Koh, Han Zhang, Ruoming Pang, James Qin, Alexander Ku, Yuanzhong Xu, Jason Baldridge, and Yonghui Wu. Vector-quantized image modeling with improved vqgan. *arXiv preprint arXiv:2110.04627*, 2021. 5, 6, 13
- [71] Wenyuan Zeng, Wenjie Luo, Simon Suo, Abbas Sadat, Bin Yang, Sergio Casas, and Raquel Urtasun. End-to-end interpretable neural motion planner. In *CVPR*, 2019. 2
- [72] Tianyuan Zhang, Xuanyao Chen, Yue Wang, Yilun Wang, and Hang Zhao. Mutr3d: A multi-camera tracking framework via 3d-to-2d queries. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4537–4546, 2022. 2
- [73] Yunpeng Zhang, Zheng Zhu, Wenzhao Zheng, Junjie Huang, Guan Huang, Jie Zhou, and Jiwen Lu. Bverse: Unified perception and prediction in birds-eye-view for vision-centric

autonomous driving. *arXiv preprint arXiv:2205.09743*, 2022. 2

- [74] Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *CVPR*, 2018. 2
- [75] Yin Zhou, Pei Sun, Yu Zhang, Dragomir Anguelov, Jiyang Gao, Tom Ouyang, James Guo, Jiquan Ngiam, and Vijay Vasudevan. End-to-end multi-view fusion for 3d object detection in lidar point clouds. In *CoRL*, 2020. 2



Figure 6. Additional qualitative comparison between MoST and Wayformer [48] baseline. The agent boxes are colored by their types: gray for vehicle, red for pedestrian, and cyan for cyclist. The predicted trajectories are ordered temporally from green (+0s) to blue (+8.0s). For each modeled agent, the models predict 6 trajectory candidates, whose confidence scores are illustrated by transparency: the more confident, the more visible. Ground truth trajectory is shown as red dots. Note that the vehicle indicated by the red arrow is entering a plaza which has no map coverage. Since our model has access to the rich visual signals, it correctly predicts the vehicle’s possible trajectory which includes follows the arrow and turn right. Wayformer, on the other hand, completely missed this possibility due to the lack of road graph information in that region.

Appendix

A. Additional Qualitative Results

An additional qualitative comparison can be found in Figure 6. In this scenario, the model is asked to predict the future trajectory of a vehicle entering a plaza which is not mapped by the road graph. Our model with access to visual information correctly predicts several trajectories following the arrow painted on the ground and turning right.

B. WOMD Camera Embeddings

VQGAN Embedding To extract VQGAN embedding for an image, we first resize the image into shape of 256×512 . Then we horizontally split the image into two patches and apply pre-trained ViT-VQGAN [70] model on each patch respectively. Each patch contains 16×16 tokens so each camera image can be represented as 512 tokens. The codebook size is 8192.

SAM-H Embedding For each camera we extract SAM ViT-H [36] embedding of size $64 \times 64 \times 256$. Compared to VQGAN embeddings, SAM features are less spatially compressed due to its high-resolution feature map. The visualization of SAM Embedding can be found in Figure 8. We release the SAM features pooled per-scene-element.




Figure 7. Examples of reconstructed driving images from ViT-VQGAN codes. We show 3 cameras at 3 consecutive timestamps. We are able to decode high quality images from VQGAN codes.




Figure 8. Examples of SAM feature. The first row shows camera images and the second row illustrates the SAM feature map visualized by PCA reduction from 256 to 3 dimensions.

C. Implementation Details

Model Detail We use $N_{\text{elem}}^{\text{agent}} = 128$, $N_{\text{elem}}^{\text{open-set}} = 384$, $N_{\text{elem}}^{\text{gnd}} = 256$, and $N_{\text{pts}} = 65536$ in our experiments. We use sensor data from past 10 frames that correspond to the 1 second history and the current frame (i.e. $T = 11$). Following Wayformer [48], we train our model to output K modes for the Gaussian mixture, where we experiment with $K = \{6, 64\}$. During inference, we draw 2048 samples from the predicted Gaussian mixture distribution, and use K-Means clustering to aggregate those 2048 samples into 6 final trajectory predictions.

Variable name	Description	Tensor Shape
N_{pts}	The total number of LiDAR points after down-sampling.	1
N_{elem}	The total number of scene elements.	1
T	The total number of frames.	1
D	The feature dimension.	1
\mathbf{P}_{xyz}	The aggregated LiDAR points from all frames after downsampling	$N_{\text{pts}} \times 3$
\mathbf{P}_{ind}	The scene element index and frame index for each LiDAR point	$N_{\text{pts}} \times 2$
\mathbf{F}_{pts}	The per point image feature.	$N_{\text{pts}} \times D$
\mathbf{B}	The box attributes, including box center, box size, and box heading.	$N_{\text{elem}} \times T \times 7$
\mathbf{F}_{img}	The per scene-element image feature.	$N_{\text{elem}} \times T \times D$
\mathbf{F}_{geo}	The per scene-element geometry feature.	$N_{\text{elem}} \times T \times D$
$\mathbf{f}_{\text{temporal}}$	The learnable temporal embedding.	$1 \times T \times D$

Table 9. Descriptions for variables used in the main paper.

Training Detail For all experiments, we train our model using AdamW [43] on 64 Google Cloud TPUv4 cores¹ with a global batch size of 512. We use a cosine learning rate schedule, where the learning rate is initialized to 3×10^{-4} and ramps up to 6×10^{-4} after 1,000 steps. The training finishes after 500,000 steps.

Notations Please refer to Table 9 for a summary of the notations used in the main paper.

¹<https://cloud.google.com/tpu>