

Generative Planning with 3D-vision Language Pre-training for End-to-End Autonomous Driving

Tengpeng Li^{1,2}, Hanli Wang^{1,2*}, Xianfei Li³, Wenlong Liao³, Tao He^{3,4}, Pai Peng^{3*}

¹College of Electronic and Information Engineering, Tongji University, Shanghai, China

²School of Computer Science and Technology, Tongji University, Shanghai, China

³COWAROBOT, China

⁴School of Electronic Engineering, University of South China, Hunan, China

hanliwang@tongji.edu.cn, pengpai_sh@163.com

Abstract

Autonomous driving is a challenging task that requires perceiving and understanding the surrounding environment for safe trajectory planning. While existing vision-based end-to-end models have achieved promising results, these methods are still facing the challenges of vision understanding, decision reasoning and scene generalization. To solve these issues, a generative planning with 3D-vision language pre-training model named GPVL is proposed for end-to-end autonomous driving. The proposed paradigm has two significant aspects. On one hand, a 3D-vision language pre-training module is designed to bridge the gap between visual perception and linguistic understanding in the bird's eye view. On the other hand, a cross-modal language model is introduced to generate holistic driving decisions and fine-grained trajectories with perception and navigation information in an autoregressive manner. Experiments on the challenging nuScenes dataset demonstrate that the proposed scheme achieves excellent performances compared with state-of-the-art methods. Besides, the proposed GPVL presents strong generalization ability and real-time potential when handling high-level commands in various scenarios. It is believed that the effective, robust and efficient performance of GPVL is crucial for the practical application of future autonomous driving systems.

Introduction

Autonomous driving is a challenging task that requires a profound understanding about the surrounding environment of autonomous vehicles to ensure safety and efficient real-world deployment. An excellent autonomous driving system must have the capacity to comprehensively perceive the driving environment and precisely make the planning decision based on the road information.

Recently, several end-to-end autonomous driving frameworks are proposed and achieve promising results, by understanding the driving scenes with sensor data and outputting planning decisions. Some early methods (Codevilla et al. 2019; Prakash, Chitta, and Geiger 2021) directly obtain the predicted planning trajectories without thorough understanding of driving scenes, resulting in limited

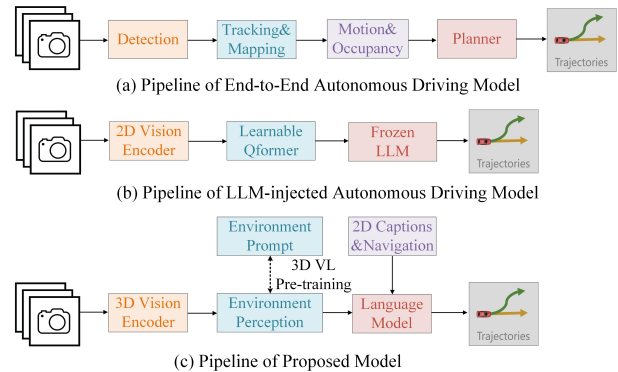


Figure 1: (a) The existing end-to-end autonomous driving framework only utilizes visual information to complete perception, prediction and planning tasks. (b) The emerging LLM-injected autonomous driving models merely introduce 2D visual features and use time-consuming LLM for planning decision. (c) The designed scheme focuses on 3D-vision language pre-training and conducts the planning by a language generation style.

interpretability and difficulty in convergence. As illustrated in Fig. 1(a), most schemes (Hu et al. 2022, 2023; Jiang et al. 2023) utilize the sensor information to integrate multiple vision tasks (*e.g.*, 3D object detection and motion prediction) in one network to enhance planning performance. However, these vision-only approaches struggle to incorporate contextual cues for safe decision-making and require further optimization. With the rapid development of large language models (LLM), the outstanding reasoning ability has attracted much interest and promoted applications in the autonomous driving area. As presented in Fig. 1(b), several recent methods (Xu et al. 2023; Wang et al. 2023b) introduce multi-modal large language models (MLLM) to produce driving explanations and language-injected planning results. Nonetheless, these models pre-trained on 2D vision-language datasets struggle with understanding the complex 3D spatial relations in driving scenarios. Therefore, several effective strategies are designed to overcome above challenges, such as directly using the detected 3D bounding box-

*Corresponding authors.

es of crucial objects (Tian et al. 2024) and integrating the bird’s eye view (BEV) feature with learnable queries (Ding et al. 2024). However, these advanced strategies merely integrate 2D and 3D visual information without adequately addressing the alignment of 3D visual features with linguistic representations.

To address the aforementioned challenges, a novel generative planning with 3D-vision language pre-training (GPVL) paradigm is proposed for autonomous driving, as shown in Fig. 1(c). GPVL first utilizes the pre-trained BEV-former (Li et al. 2022d) to extract the BEV feature map of multi-view images, which covers the essential semantic elements (*e.g.*, vehicles, pedestrians and traffic lines) of driving scenes. Motivated by VAD (Jiang et al. 2023), the vectorized detection, motion and mapping transformers are introduced to learn the crucial perception information. Then, a 3D-vision language pre-training module is developed to align BEV features with linguistic representations, enabling comprehensive 3D scene understanding and textual reasoning in a shared feature space, where the group-wise alignment establishes the multi-level associations between vision and language features. Furthermore, GPVL designs a 2D scene captioning model based on the pretrained BLIP (Li et al. 2022a), which is trained and tested on the Ominidrive-nuScenes dataset (Wang et al. 2024) to generate scene-level descriptions. Finally, the 2D visual captioning, the aligned 3D perception feature and navigation instruction are fed into a language model to generate the holistic driving decision and fine-grained trajectory in an auto-regressive manner.

The major contributions of this work are summarized in the following three aspects.

- A 3D-vision language pre-training module is presented to establish group-wise correlations between vision and language features, facilitating a thorough understanding of the driving environment.
- A cross-modal language model is developed to generate decisions and trajectories using captioning, perception and navigation information in an auto-regressive manner, endowing the model with reasoning and generation capabilities.
- We propose a generative planning with 3D-vision language pre-training framework, which learns language-guided perception features and generates contextual trajectories, thereby enhancing system safety.

Related Work

End-to-End Autonomous Driving

The goal of end-to-end autonomous driving is to construct a completely differentiable system that can convert sensor data into vehicle control commands (Prakash, Chitta, and Geiger 2021; Wu et al. 2022). This system integrates all modules (*i.e.*, perception, prediction, planning and control) for interpretable end-to-end optimization, which can reduce the accumulated errors of all modules and enhance the safety of the autonomous driving system. Recently, several effective models are proposed for this challenging task and can be categorized into open-loop autonomous driving methods and closed-loop autonomous driving methods. Many

schemes (Hu et al. 2023; Gu et al. 2023; Jiang et al. 2023; Pan et al. 2024) address this task in the open-loop manner, whose training and evaluation stages are conducted on real-world datasets. For example, Hu *et al.* (Hu et al. 2023) propose an innovative planning-oriented approach to interactively optimize all blocks. Jiang *et al.* (Jiang et al. 2023) develop a vectorized scene representation technology to meet real-time requirements. In contrast, closed-loop methods use simulators like CARLA (Dosovitskiy et al. 2017) to create virtual environments for the feedback of other agents on the road. Specifically, Jia *et al.* (Jia et al. 2023) optimize future position and action features twice to improve model safety. Chen *et al.* (Chen et al. 2024) utilize the environmental features to learn the probabilistic distribution of actions. Nevertheless, these methods have difficulty to generate robust trajectories in a reasonable way for efficient vehicle control. In this work, the proposed GPVL is capable of performing reasonable analysis and generating contextual trajectory.

Large Language Models for Autonomous Driving

Recently, many researches (Driess et al. 2023; Wang et al. 2023a; Brohan et al. 2023) employ LLMs to endow autonomous driving systems with the embodied intelligence capability and achieve favorable results in the decision stage. Inspired by the success in the field of robotics, several works (Wen et al. 2023; Mao et al. 2023) attempt to use LLMs for autonomous driving. In particular, Dilu (Wen et al. 2023) and GPT-driver (Mao et al. 2023) utilize LLM to generate reasonable trajectories by regarding the perception information of agents as language prompts. Furthermore, recent M-LLMs (Driess et al. 2023; Li et al. 2023; Bai et al. 2023; Liu et al. 2023) introduce both visual encoder and LLM to successfully establish the semantic correlations between vision and language, which have attracted much interest in the autonomous driving area. For example, Shao *et al.* (Shao et al. 2024) learn an integrated representation of multi-modal sensor data and human navigation instruction to output reliable control commands. Ding *et al.* (Ding et al. 2024) incorporate the BEV representation into multi-view video features to enhance the understanding of driving scenes. However, the fused multi-modal features of these schemes either lack 3D spatial awareness of all targets or fail to align with linguistic representations. In this work, the introduced BEV-level visual transformer and the 3D-vision language pre-training module are devised to address these challenges.

Proposed Method

The overall framework of the proposed GPVL is presented in Fig. 2, including three key components: (1) the backbone generates the supervised detection, map and motion features based on the extracted BEV feature, (2) the 3D-vision language pre-training module aims to align the visual and linguistic features in a common semantic space, and (3) the cross-modal language model produces the reliable decision and trajectory in a generative fashion.

Method Overview

Suppose $\mathcal{X} = \{\mathbf{X}_n\}_{n=1}^N$ denotes the sampled N multi-view images, the goal of the proposed scheme is to generate the




Figure 2: Pipeline of GPVL for autonomous driving. The framework is divided into three parts: (1) the backbone includes a 3D-vision encoder to obtain the basic BEV feature, then it is decoded into constrained detection, motion and map features, (2) the 3D-vision language pre-training module establishes the associations between vision and language features with the group-wise alignment, (3) the cross-modal language model generates the future planning decision in an auto-regressive manner based on aligned visual feature and navigation prompt. Note that in (a), the env prompt represents the detailed descriptions of 3D bounding boxes, agent motions and BEV map on the road, and the navigation means the detailed description of high-level instruction for the self-driving car. Four kinds of 3D-vision language (VL) alignments are presented in (b) for explanation.

safe driving trajectory of ego vehicle. To acquire the satisfying planning result, GPVL is proposed to establish semantic associations between 3D visual features and linguistic representations, and generates the high-quality route in an auto-regressive manner. Figure 2 illustrates the detailed structure of the proposed model. Specifically, for the input multi-view images \mathcal{X} , we first use the BEV encoder in the pre-trained BEVformer (Li et al. 2022d) to extract the BEV feature $\mathbf{B} = \mathcal{F}_{bev}(\mathcal{X})$. Afterwards, \mathbf{B} is fed into the detection, motion and map transformers to simultaneously learn the 3D object boxes, agent motions and map elements, resulting in constrained detection feature \mathbf{F}_{det} , motion feature \mathbf{F}_{motion} and map feature \mathbf{F}_{map} . Then, \mathbf{F}_{det} , \mathbf{F}_{motion} and \mathbf{F}_{map} are sent into a visual attention block \mathcal{F}_{vatt} composed of multiple transformer layers, generating the attentive visual features \mathbf{V}_{det} , \mathbf{V}_{motion} , \mathbf{V}_{map} and their concatenation \mathbf{V}_{global} . Meanwhile, the environment ground-truth descriptions are integrated into a textual attention block \mathcal{F}_{tatt} based on the BERT (Devlin et al. 2019) structure to obtain textual representations \mathbf{T}_{det} , \mathbf{T}_{motion} , \mathbf{T}_{map} and their concatenation \mathbf{T}_{global} . A group-wise alignment module is further designed to project these vision and language features in a shared semantic space. Finally, \mathbf{V}_{global} together with the navigation prompt of ego vehicle are sent into the language model to generate the dependable trajectory in a natural language format.

3D-vision Language Pre-training

The 3D-vision language pre-training module is developed to establish the multi-level alignment between vision and lan-

guage modalities through contrastive learning (Radford et al. 2021). Several previous methods (Shao et al. 2024; Wang et al. 2023b; Ding et al. 2024) have introduced pre-trained LLMs (e.g., LLaMA (Touvron et al. 2023)) to enhance perception and understanding of driving scenes. However, these approaches either lack 3D spatial information or exhibit semantic gaps between different representations, which hinder accurate target localization and trajectory inference based on extracted features. In contrast, the proposed model utilizes the supervised detection, motion and map features to perceive the 3D spatial distributions of targets and minimizes the semantic gap between different features through the 3D-vision language pre-training module.

BEV-injected Visual Features. In the proposed GPVL, three kinds of targets in the driving scene are introduced, including foreground objects, agent motions and map elements. We first utilize the visual embedding layer to encode $\mathbf{F}_{det} \in \mathbb{R}^{N_d \times d_d}$, $\mathbf{F}_{motion} \in \mathbb{R}^{N_o \times d_o}$ and $\mathbf{F}_{map} \in \mathbb{R}^{N_m \times d_m}$ into the same channel dimension d_c , where N_d , N_o and N_m represent the numbers of bounding boxes, future trajectories and map elements, while d_d , d_o and d_m denote the feature dimensions of boxes, motions and maps, respectively. Then, the transformer structure (Vaswani et al. 2017) is introduced to establish compact interactions via long-range attention for enhancing visual features. The whole function can be formulated as

$$\begin{aligned} \mathbf{V}_{det} &= \text{VisAtt}(\text{Embed}(\mathbf{F}_{det})), \\ \mathbf{V}_{motion} &= \text{VisAtt}(\text{Embed}(\mathbf{F}_{motion})), \\ \mathbf{V}_{map} &= \text{VisAtt}(\text{Embed}(\mathbf{F}_{map})), \end{aligned} \quad (1)$$

where $\text{Embed}(\cdot)$ means the feature embedding layer and $\text{VisAtt}(\cdot)$ represents the visual attention block. Therefore, $\mathbf{V}_{det} \in \mathbb{R}^{N_d \times d_c}$, $\mathbf{V}_{motion} \in \mathbb{R}^{N_o \times d_c}$ and $\mathbf{V}_{map} \in \mathbb{R}^{N_m \times d_c}$, where d_c means the common channel dimension. Afterwards, the attentive detection, motion and map features are integrated along the spatial-wise dimension to obtain the global visual feature $\mathbf{V}_{global} \in \mathbb{R}^{(N_d+N_o+N_m) \times d_c}$ as

$$\mathbf{V}_{global} = \text{Cat}[\mathbf{V}_{det}, \mathbf{V}_{motion}, \mathbf{V}_{map}], \quad (2)$$

where $\text{Cat}(\cdot)$ denotes the concatenation operation.

Environmental Linguistic Features. In order to endow the model with language awareness, a textual attention module is presented by utilizing the BERT structure. This module processes perception and planning labels, such as bounding boxes, future trajectories and map elements, which are transformed into driving-specific language prompts using predefined templates. The descriptions related to detection, motion and map features are subsequently fed into the textual attention block to generate corresponding linguistic representations. Now, the entire procedure can be formulated as

$$\begin{aligned} \mathbf{T}_{det} &= \text{TxtAtt}(\text{Embed}(\mathbf{P}^d[\mathbf{L}_{det}])), \\ \mathbf{T}_{motion} &= \text{TxtAtt}(\text{Embed}(\mathbf{P}^o[\mathbf{L}_{motion}])), \\ \mathbf{T}_{map} &= \text{TxtAtt}(\text{Embed}(\mathbf{P}^m[\mathbf{L}_{map}])), \end{aligned} \quad (3)$$

where \mathbf{L}_{det} , \mathbf{L}_{motion} and \mathbf{L}_{map} represent the ground-truth labels of foreground objects, future motions and map elements, respectively. \mathbf{P}^d , \mathbf{P}^o and \mathbf{P}^m denote the template descriptions of foreground objects, future motions and map elements, respectively. $\text{TxtAtt}(\cdot)$ represents the BERT-based textual attention block. $\mathbf{T}_{det} \in \mathbb{R}^{L_d \times d_c}$, $\mathbf{T}_{motion} \in \mathbb{R}^{L_o \times d_c}$ and $\mathbf{T}_{map} \in \mathbb{R}^{L_m \times d_c}$, where L_d , L_o and L_m are sentence lengths. Subsequently, these descriptive features are concatenated to form the global textual representation $\mathbf{T}_{global} \in \mathbb{R}^{(L_d+L_o+L_m) \times d_c}$ as

$$\mathbf{T}_{global} = \text{Cat}[\mathbf{T}_{det}, \mathbf{T}_{motion}, \mathbf{T}_{map}]. \quad (4)$$

Group-wise Alignment. In order to fully exploit the multi-level associations between different features, a group-wise alignment strategy is designed by using contrastive learning (Radford et al. 2021). Four types of 3D-vision language (VL) alignment groups are considered, including detection-related VL group, motion-related VL group, map-related VL group and global VL group. For example, in a batch of K training samples, the alignment function aims to treat matched VL samples as positives and mismatched VL samples as negatives. Given visual feature \mathbf{X} and textual feature \mathbf{T} as inputs, the contrastive loss is defined as

$$\begin{aligned} \mathcal{L}_{align}(\mathbf{V}, \mathbf{T}) &= -\log \frac{\exp(s(\mathbf{V}_i, \mathbf{T}_i)/\tau)}{\sum_{j=1}^K \exp(s(\mathbf{V}_i, \mathbf{T}_j)/\tau)} \\ &\quad -\log \frac{\exp(s(\mathbf{T}_i, \mathbf{V}_i)/\tau)}{\sum_{j=1}^K \exp(s(\mathbf{T}_i, \mathbf{V}_j)/\tau)}, \end{aligned} \quad (5)$$

where τ is a learnable temperature coefficient, and $s(\cdot, \cdot)$ is a similarity function. The function $s(\cdot, \cdot)$ is obtained by computing the similarity score between visual and textual features processed by global average pooling. In this module,

a learnable weight is introduced to explore fine-grained relations between different representations. For instance, the similarity function of detection-related VL group can be formulated as

$$\begin{aligned} s(\mathbf{V}_{det}, \mathbf{T}_{det}) &= \frac{1}{2} \max_{j=1}^{L_d} (\mathbf{W}_1 \mathbf{V}_{det} \mathbf{T}_{det}^j) + \\ &\quad \frac{1}{2} \max_{j=1}^{N_d} (\mathbf{W}_2 \mathbf{T}_{det} \mathbf{V}_{det}^j), \end{aligned} \quad (6)$$

where $\mathbf{W}_1 \in \mathbb{R}^{1 \times N_d}$ and $\mathbf{W}_2 \in \mathbb{R}^{1 \times L_d}$. Therefore, the total group-wise alignment loss is defined as

$$\begin{aligned} \mathcal{L}_{ga} &= \mathcal{L}_{align}(\mathbf{V}_{det}, \mathbf{T}_{det}) + \mathcal{L}_{align}(\mathbf{V}_{motion}, \mathbf{T}_{motion}) \\ &\quad + \mathcal{L}_{align}(\mathbf{V}_{map}, \mathbf{T}_{map}) + \mathcal{L}_{align}(\mathbf{V}_{global}, \mathbf{T}_{global}). \end{aligned} \quad (7)$$

Planning via Cross-modal Language Model

The ego-agent dynamic interaction is an essential issue in the autonomous driving system. Previous studies (Hu et al. 2023; Jiang et al. 2023) have attempted to introduce learnable queries to model ego-agent relations for the query feature of ego vehicle, and it is sent into a multi-layer perceptron (MLP) to acquire expected future trajectories. Although this strategy can present favorable performance on the specific benchmark dataset, directly generating the trajectory using MLP can lead to overfitting and have difficulty in contextual relation inference among generated waypoints. Facing these challenges, a cross-modal language model for generative planning is developed to empower GPVL with the capability to make the safe decision in a rational and robust manner.

Ego-agent Cross-modal Decoder. The proposed model formulates a language prompt for the self-driving car’s current status, which includes its high-level driving command and location. The informative prompt is sent to the text embedding layer to yield the initial linguistic representation of the self-driving car. Meanwhile, as shown in Fig. 2, the designed scene captioning model generates visual descriptions (e.g., traffic lights, signs, crucial objects and weather). Subsequently, the environmental visual feature \mathbf{V}_{global} and the prompt feature are fed into a language model to learn the planning feature with rich driving scenario and navigation information. The process can be formulated as

$$\mathbf{F} = \text{LM}(\text{Embed}(\text{Cap}), \text{Embed}(\text{Nav}), \mathbf{V}_{global}), \quad (8)$$

where $\text{LM}(\cdot)$ denotes the pre-trained language model, Cap is the 2D scene captioning, Nav means the high-level navigation and its detailed template is “The box of ego-car is [cx, cy, cz, cw, ch, cl]. It is currently {command} and can not collide with other vehicles or the boundaries of the BEV map. Please predict the trajectory for the next p timestamps.” Eventually, $\mathbf{F} \in \mathbb{R}^{N \times d_c}$ is sent into the Linear (LeCun, Bengio, and Hinton 2015) layer and the Softmax (Bishop and Nasrabadi 2006) layer to yield the word vector as

$$p(w_t | w_{1:t-1}) = \text{Softmax}(\text{Linear}(\mathbf{F})), \quad (9)$$

where p means the probability prediction of vocabulary.

Generative Planning. Similar to most visual captioning tasks (Pan et al. 2020; Li et al. 2022b,c), a cross-entropy loss

is introduced to output the trajectory in a language modeling fashion. During training, each sample \mathcal{X} is equipped with a prompt captioning Gt as a reference, which includes high-level command and future trajectory of ego vehicle, and the loss function is expressed as

$$\mathcal{L}_{cap} = - \sum_{t=1}^T \log(p_{\theta}(w_t | (Gt)_{1:t-1})), \quad (10)$$

where θ represents all trainable parameters during training and w_t means the t -th predicted word.

Training Loss

The overall loss \mathcal{L}_{gpvl} of the proposed GPVL comprises three training components, including visual perception loss, 3D-vision language alignment loss and trajectory captioning loss, which can be defined as

$$\mathcal{L}_{gpvl} = \mathcal{L}_{vis} + \mathcal{L}_{ga} + \mathcal{L}_{cap}, \quad (11)$$

where \mathcal{L}_{vis} means the losses associated with visual downstream tasks, including 3D object detection, map construction and motion prediction.

Experiment

Dataset and Automatic Metric Evaluation

Dataset. Substantial experiments are conducted on the complex public nuScenes (Caesar et al. 2020) dataset, which comprises 1,000 traffic scenarios, and the duration of each video is around 20 seconds. This dataset offers over 1.4 million 3D bounding boxes across 23 different object categories. The scene images are recorded by six cameras, and the key frames are labeled 2 times per second.

Evaluation Metric. Two widely used objective metrics are employed to comprehensively validate the planning results, including displacement error and collision rate, which are denoted as L2 and Collision for short in the following result presentation, respectively. In particular, L2 measures the distance between predicted trajectory and ground-truth trajectory, and Collision evaluates the frequency of collisions that occur in real-world driving scenarios. Additionally, the Latency and FPS metrics are introduced to evaluate the real-time performance of models.

Implementation Details

The proposed model aims to predict the trajectory for the future 3 seconds. The input image size is 1280×720 . GPVL utilizes ResNet50 (He et al. 2016) to extract the multi-view image features. The numbers of BEV queries, bounding boxes and map points are 200×200 , 200 and 100×20 , respectively. The feature dimension and hidden size are 768 and 512, respectively. The model utilizes the AdamW (Loshchilov and Hutter 2017) optimizer and weight decay 0.01 in the training process. The learning rates in three training stages are 2×10^{-4} , 1×10^{-4} and 5×10^{-6} , respectively. The BERT (Devlin et al. 2019) structure is used by 3D-vision language pre-training and cross-modal language model. As for testing, the size of greedy search is set to 1 to generate the trajectory caption. The proposed model is trained on PyTorch framework with 8 NVIDIA RTX A6000 cards.

Comparison with State-of-the-art Methods

The proposed scheme is compared with 13 autonomous driving methods, including 8 traditional non-autoregressive methods and 5 LLM-injected autoregressive methods.

Quantitative Result Table 1 shows the statistical comparisons of the proposed GPVL with other state-of-the-art methods. Generally, the statistical performance of GPVL is obviously better than other approaches. Specifically, GPVL obtains the lowest scores in L2 distance metrics and obviously reduces the planning displacement errors by 0.18m at 1s, 0.28m at 2s, 0.34m at 3s and 0.27m on average compared to VAD, indicating the superior accuracy of trajectory prediction. Moreover, GPVL exhibits the best performance in most collision rate metrics, highlighting its exceptional safety and robustness in avoiding collisions. Compared to the approaches that introduce auxiliary tasks or utilize LLMs, the proposed method maintains comparable planning performance and achieves a latency of 188.7 milliseconds and an inference speed of 5.3 fps, presenting the potential for practical applications. According to the report of BEV-Planner (Li et al. 2024), the proposed scheme excludes ego status information to prevent shortcut learning.

Table 2 shows the evaluated L2 and Collision scores in terms of turn left, turn right and go straight commands. It is evident that GPVL significantly outperforms UniAD and VAD across all metrics. In the nuScenes dataset, 87.7% training and 88.2% validation samples consist of simple go straight scenes. Therefore, UniAD and VAD are more prone to overfitting and learning shortcuts on these samples, resulting in poor performances in more complex turning scenarios. In contrast, the proposed method acquires favorable results across all scenarios, demonstrating strong generalization capabilities in diverse driving situations.

Qualitative Result The planning results generated by GPVL, compared to VAD and ground-truth, are illustrated in Fig. 3. To provide a comprehensive understanding of driving scenes, multi-view camera images are included and planning trajectories are visualized in the front camera image. Generally, benefiting from the proposed 3D-vision language pre-training module and cross-modal language model, the designed model generates accurate and reasonable trajectories. For example, in the first scenario, where the ego vehicle is instructed to go straight, the trajectory generated by GPVL guides the vehicle to safely navigate through the urban road. Conversely, VAD’s trajectory has a collision risk with the roadside. In the third scenario, the result of GPVL closely matches the ground-truth, enabling the vehicle to safely turn left at intersections. Nonetheless, the planning decision of VAD is more aggressive, potentially leading to collision with the oncoming vehicle.

Ablation Study

The ablation study in Table 3 systematically investigates the contributions of the key components of GPVL on the nuScenes dataset. Without the perception module, GPVL struggles with detecting foreground objects, predicting motions and constructing maps, leading to higher L2 and Colli-

Method	L2 (m) ↓				Collision (%) ↓				Latency (ms)	FPS
	1s	2s	3s	Avg.	1s	2s	3s	Avg.		
Non-Autoregressive										
NMP [†] (Zeng et al. 2019)	-	-	2.31	-	-	-	1.92	-	-	-
SA-NMP [†] (Zeng et al. 2019)	-	-	2.05	-	-	-	1.59	-	-	-
FF [†] (Hu et al. 2021)	0.55	1.20	2.54	1.43	0.06	0.17	1.07	0.43	-	-
ST-P3 (Hu et al. 2022)	1.33	2.11	2.90	2.11	0.23	0.62	1.27	0.71	628.3	1.6
Ego-MLP (Zhai et al. 2023)	0.46	0.76	1.12	0.78	0.21	0.35	0.58	0.38	-	-
UniAD (Hu et al. 2023)	0.48	0.96	1.65	1.03	0.05	0.17	0.71	0.31	555.6	1.8
VAD (Jiang et al. 2023)	0.41	0.70	1.05	0.72	0.07	0.17	0.41	0.22	224.3	4.5
BEV-Planner (Li et al. 2024)	0.30	0.52	0.83	0.55	0.10	0.37	1.30	0.59	-	-
Autoregressive										
LLaVA (Liu et al. 2023)	1.04	1.74	2.57	1.79	0.58	1.17	1.74	1.16	2532.5	0.4
Vicuna (Chiang et al. 2023)	1.06	1.80	2.54	1.80	0.60	1.21	1.78	1.20	2498.9	0.4
Merlin (Yu et al. 2023)	1.03	1.71	2.40	1.71	0.48	1.05	1.77	1.10	-	-
Ominidrive (Wang et al. 2024)	0.40	0.80	1.32	0.84	0.04	0.46	2.32	0.94	-	-
Atlas (Bai et al. 2024)	0.52	0.97	1.53	1.00	0.15	0.31	0.70	0.38	-	-
GPVL	0.21	0.39	0.69	0.43	0.07	0.09	0.27	0.14	198.2	5.1

Table 1: **Open-loop planning performance.** GPVL achieves the highest score on most evaluation metrics on the nuScenes (Caesar et al. 2020) val dataset. LiDAR-based methods are denoted with †. In open-loop evaluation, the ego status information of GPVL is deactivated for a fair comparison.

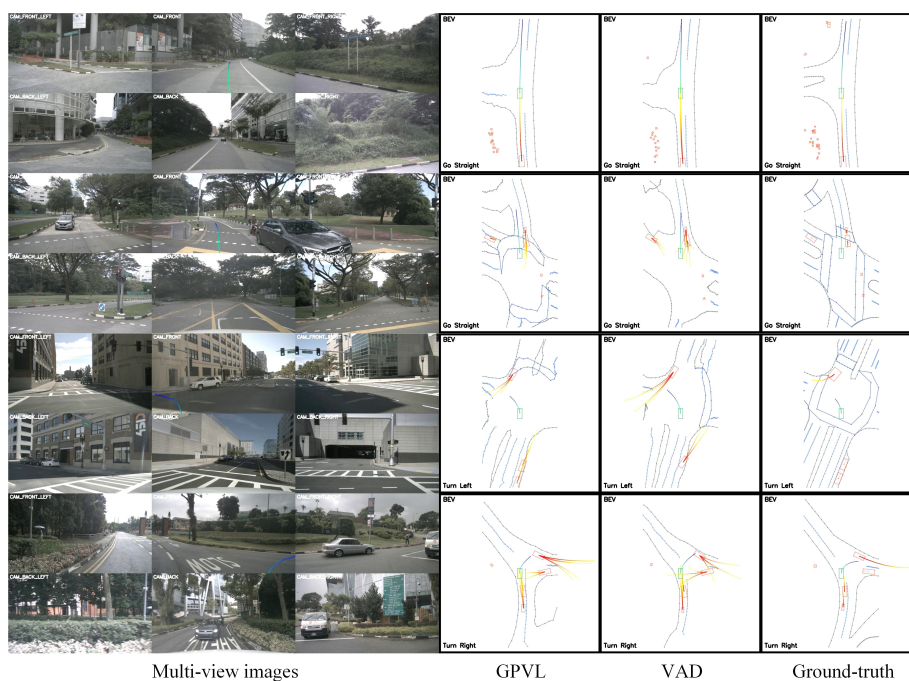


Figure 3: Visualized comparison of the proposed GPVL, VAD and the ground-truth on the nuScenes dataset.

Methods	Turn Left				Turn Right				Go Straight				
	1s	2s	3s	Avg.	1s	2s	3s	Avg.	1s	2s	3s	Avg.	
L2	UniAD	0.37	0.62	0.95	0.65	0.42	0.77	1.23	0.80	0.46	0.70	1.03	0.73
	VAD	0.49	0.83	1.24	0.85	0.49	0.88	1.34	0.90	0.40	0.67	1.02	0.70
	GPVL	0.29	0.52	0.88	0.56	0.27	0.51	0.85	0.54	0.19	0.38	0.66	0.41
Col.	UniAD	0.15	0.35	0.75	0.42	0.09	0.93	1.06	0.69	0.10	0.15	0.55	0.27
	VAD	0.13	0.21	0.68	0.34	0.18	0.85	1.19	0.74	0.12	0.14	0.31	0.19
	GPVL	0.12	0.18	0.37	0.22	0.09	0.23	0.39	0.24	0.07	0.08	0.24	0.13

Table 2: Statistical results of L2 distance and collision rate (Col.) with turn left, turn right and go straight commands.

ID	Perc	Cap	VLP	GA	CLM	L2 (m) ↓				Collision (%) ↓			
						1s	2s	3s	Avg.	1s	2s	3s	Avg.
1	-	✓	✓	✓	✓	0.99	1.27	1.88	1.38	0.41	0.72	1.13	0.75
2	✓	-	✓	✓	✓	0.25	0.43	0.76	0.48	0.10	0.12	0.34	0.19
3	✓	✓	-	-	✓	0.49	0.75	1.12	0.79	0.25	0.33	0.57	0.38
4	✓	✓	✓	-	✓	0.36	0.59	0.91	0.62	0.16	0.24	0.48	0.29
5	✓	✓	✓	✓	-	0.30	0.54	0.85	0.56	0.11	0.17	0.39	0.22
6	✓	✓	✓	✓	✓	0.21	0.39	0.69	0.43	0.07	0.09	0.27	0.14

Table 3: Ablation study of GPVL on nuScenes, where Perc, Cap, VLP, GA and CLM represent perception modules, captioning model, 3D-vision language pre-training, group-wise alignment and cross-modal language model, respectively.

sion scores. Disabling the VLP and GA components significantly degrades performance, highlighting model’s strong ability to bridge visual and linguistic understanding. The absence of GA results in obvious performance degradation, indicating its importance in fine-grained feature association. The exclusion of CLM increases L2 and Collision scores, emphasizing its role in generating reasonable planning decisions. Finally, as shown in the sixth row of Table 3, the integration of all modules yields the best performance, showcasing the synergistic effect of the combined system.

Zero-shot Generalization

Methods	L2 (m) ↓				Collision (%) ↓				
	1s	2s	3s	Avg.	1s	2s	3s	Avg.	
Group1	UniAD	0.49	0.91	1.66	1.02	0.25	0.44	0.93	0.54
	VAD	0.44	0.78	1.19	0.80	0.22	0.34	0.69	0.41
	GPVL	0.28	0.49	0.81	0.53	0.11	0.15	0.32	0.19
Group2	UniAD	0.53	0.97	1.71	1.07	0.28	0.41	0.99	0.56
	VAD	0.49	0.72	1.08	0.76	0.25	0.33	0.72	0.43
	GPVL	0.31	0.55	0.84	0.57	0.09	0.17	0.33	0.20

Table 4: To evaluate the zero-shot performance on the new city, the models are trained on Boston and tested on Singapore in Group1, and the models are trained on Singapore and tested on Boston in Group2.

Methods	L2 (m) ↓				Collision (%) ↓			
	1s	2s	3s	Avg.	1s	2s	3s	Avg.
UniAD	0.59	1.12	1.89	1.20	0.21	0.39	0.92	0.51
VAD	0.48	0.77	1.15	0.80	0.19	0.31	0.53	0.34
GPVL	0.27	0.50	0.89	0.55	0.11	0.18	0.37	0.22

Table 5: To validate the robustness of models in unseen scenarios, four types of noise are introduced into test images, including rain, fog, snow and darkness.

To validate the model’s generalization ability, we train and test the models on datasets constructed from two different urban environments (*i.e.*, Boston and Singapore). Specifically, two groups of experiments are introduced: (1) training on Boston and testing on Singapore, (2) training on Singapore and testing on Boston. As shown in Table 4, the evaluated scores of GPVL in both groups are obviously better than UniAD and VAD. Moreover, four kinds of noise (*i.e.*, rain,

fog, snow and darkness) are introduced into test images to validate the robustness of GPVL, as presented in Table 5, these noise conditions have a significant negative impact on UniAD and VAD, while they have a minor effect on GPVL. Therefore, the outstanding performance of GPVL in various real-world scenarios demonstrates its capacity to improve the robustness and safety of autonomous driving system.

Conclusion

In this work, a novel generative planning with 3D-vision language model is proposed for end-to-end autonomous driving. Specifically, the 3D-vision language pre-training module is designed to integrate valuable textual information and establish a rich 3D-vision language relation, where group-wise alignment aims to exploit multi-level associations between different representations, facilitating a better understanding and reasoning of driving scenes. The cross-modal language model is developed to serve as the generative engine, which utilizes the aligned feature and navigation to produce the future trajectory in an auto-regressive manner. This generative style enables the model to make correct decisions similar to natural language modeling. The proposed GPVL constructs a unified framework that not only performs reliable planning but also exhibits superior generalization capacity across various driving scenarios. Extensive experiments on the nuScenes dataset demonstrate that GPVL significantly outperforms state-of-the-art methods. In future works, the proposed GPVL is expected to promote the development of safer and more reliable autonomous driving technology.

Acknowledgments

This work was supported in part by National Natural Science Foundation of China under Grant 62371343, and in part by Shanghai Municipal Science and Technology Major Project under Grant 2021SHZDZX0100.

References

- Bai, J.; Bai, S.; Yang, S.; Wang, S.; Tan, S.; Wang, P.; Lin, J.; Zhou, C.; and Zhou, J. 2023. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*.
- Bai, Y.; Wu, D.; Liu, Y.; Jia, F.; Mao, W.; Zhang, Z.; Zhao, Y.; Shen, J.; Wei, X.; Wang, T.; et al. 2024. Is a 3D-

- Tokenized LLM the Key to Reliable Autonomous Driving? *arXiv preprint arXiv:2405.18361*.
- Bishop, C. M.; and Nasrabadi, N. M. 2006. *Pattern recognition and machine learning*, volume 4. Springer.
- Brohan, A.; Brown, N.; Carbajal, J.; Chebotar, Y.; Chen, X.; Choromanski, K.; Ding, T.; Driess, D.; Dubey, A.; Finn, C.; et al. 2023. Rt-2: Vision-language-action models transfer web knowledge to robotic control. *arXiv preprint arXiv:2307.15818*.
- Caesar, H.; Bankiti, V.; Lang, A. H.; Vora, S.; Liong, V. E.; Xu, Q.; Krishnan, A.; Pan, Y.; Baldan, G.; and Beijbom, O. 2020. nuscenes: A multimodal dataset for autonomous driving. In *CVPR*, 11621–11631.
- Chen, S.; Jiang, B.; Gao, H.; Liao, B.; Xu, Q.; Zhang, Q.; Huang, C.; Liu, W.; and Wang, X. 2024. VADv2: End-to-End Vectorized Autonomous Driving via Probabilistic Planning. *arXiv preprint arXiv:2402.13243*.
- Chiang, W.-L.; Li, Z.; Lin, Z.; Sheng, Y.; Wu, Z.; Zhang, H.; Zheng, L.; Zhuang, S.; Zhuang, Y.; Gonzalez, J. E.; et al. 2023. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. See <https://vicuna.lmsys.org> (accessed 14 April 2023), 2(3): 6.
- Codevilla, F.; Santana, E.; López, A. M.; and Gaidon, A. 2019. Exploring the limitations of behavior cloning for autonomous driving. In *ICCV*, 9329–9338.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, 4171–4186.
- Ding, X.; Han, J.; Xu, H.; Liang, X.; Zhang, W.; and Li, X. 2024. Holistic Autonomous Driving Understanding by Bird’s-Eye-View Injected Multi-Modal Large Models. In *CVPR*, 13668–13677.
- Dosovitskiy, A.; Ros, G.; Codevilla, F.; Lopez, A.; and Koltun, V. 2017. CARLA: An open urban driving simulator. In *CoRL*, 1–16.
- Driess, D.; Xia, F.; Sajjadi, M. S.; Lynch, C.; Chowdhery, A.; Ichter, B.; Wahid, A.; Tompson, J.; Vuong, Q.; Yu, T.; et al. 2023. Palm-e: An embodied multimodal language model. *arXiv preprint arXiv:2303.03378*.
- Gu, J.; Hu, C.; Zhang, T.; Chen, X.; Wang, Y.; Wang, Y.; and Zhao, H. 2023. Vip3d: End-to-end visual trajectory prediction via 3d agent queries. In *CVPR*, 5496–5506.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*, 770–778.
- Hu, P.; Huang, A.; Dolan, J.; Held, D.; and Ramanan, D. 2021. Safe local motion planning with self-supervised freespace forecasting. In *CVPR*, 12732–12741.
- Hu, S.; Chen, L.; Wu, P.; Li, H.; Yan, J.; and Tao, D. 2022. St-p3: End-to-end vision-based autonomous driving via spatial-temporal feature learning. In *ECCV*, 533–549.
- Hu, Y.; Yang, J.; Chen, L.; Li, K.; Sima, C.; Zhu, X.; Chai, S.; Du, S.; Lin, T.; Wang, W.; et al. 2023. Planning-oriented autonomous driving. In *CVPR*, 17853–17862.
- Jia, X.; Wu, P.; Chen, L.; Xie, J.; He, C.; Yan, J.; and Li, H. 2023. Think twice before driving: Towards scalable decoders for end-to-end autonomous driving. In *CVPR*, 21983–21994.
- Jiang, B.; Chen, S.; Xu, Q.; Liao, B.; Chen, J.; Zhou, H.; Zhang, Q.; Liu, W.; Huang, C.; and Wang, X. 2023. Vad: Vectorized scene representation for efficient autonomous driving. In *ICCV*, 8340–8350.
- LeCun, Y.; Bengio, Y.; and Hinton, G. 2015. Deep learning. *Nature*, 521(7553): 436–444.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*, 19730–19742.
- Li, J.; Li, D.; Xiong, C.; and Hoi, S. 2022a. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, 12888–12900.
- Li, Q.; Li, T.; Wang, H.; and Chen, C. W. 2022b. Taking an emotional look at video paragraph captioning. *arXiv preprint arXiv:2203.06356*.
- Li, T.; Wang, H.; He, B.; and Chen, C. W. 2022c. Knowledge-enriched attention network with group-wise semantic for visual storytelling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8634–8645.
- Li, Z.; Wang, W.; Li, H.; Xie, E.; Sima, C.; Lu, T.; Qiao, Y.; and Dai, J. 2022d. Bevformer: Learning birds-eye-view representation from multi-camera images via spatiotemporal transformers. In *ECCV*, 1–18.
- Li, Z.; Yu, Z.; Lan, S.; Li, J.; Kautz, J.; Lu, T.; and Alvarez, J. M. 2024. Is ego status all you need for open-loop end-to-end autonomous driving? In *CVPR*, 14864–14873.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023. Visual instruction tuning. In *NeurIPS*, 12021–12036.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Mao, J.; Qian, Y.; Zhao, H.; and Wang, Y. 2023. Gpt-driver: Learning to drive with gpt. *arXiv preprint arXiv:2310.01415*.
- Pan, C.; Yaman, B.; Nesti, T.; Mallik, A.; Allievi, A. G.; Velipasalar, S.; and Ren, L. 2024. VLP: Vision Language Planning for Autonomous Driving. In *CVPR*, 14760–14769.
- Pan, Y.; Yao, T.; Li, Y.; and Mei, T. 2020. X-linear attention networks for image captioning. In *CVPR*, 10971–10980.
- Prakash, A.; Chitta, K.; and Geiger, A. 2021. Multi-modal fusion transformer for end-to-end autonomous driving. In *CVPR*, 7077–7087.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*, 8748–8763.
- Shao, H.; Hu, Y.; Wang, L.; Song, G.; Waslander, S. L.; Liu, Y.; and Li, H. 2024. Lmdrive: Closed-loop end-to-end driving with large language models. In *CVPR*, 15120–15130.
- Tian, X.; Gu, J.; Li, B.; Liu, Y.; Hu, C.; Wang, Y.; Zhan, K.; Jia, P.; Lang, X.; and Zhao, H. 2024. DriveVLM: The Convergence of Autonomous Driving and Large Vision-Language Models. *arXiv preprint arXiv:2402.12289*.

Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is all you need. In *NeurIPS*, 6000–6010.

Wang, G.; Xie, Y.; Jiang, Y.; Mandlekar, A.; Xiao, C.; Zhu, Y.; Fan, L.; and Anandkumar, A. 2023a. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*.

Wang, S.; Yu, Z.; Jiang, X.; Lan, S.; Shi, M.; Chang, N.; Kautz, J.; Li, Y.; and Alvarez, J. M. 2024. OmniDrive: A Holistic LLM-Agent Framework for Autonomous Driving with 3D Perception, Reasoning and Planning. *arXiv preprint arXiv:2405.01533*.

Wang, W.; Xie, J.; Hu, C.; Zou, H.; Fan, J.; Tong, W.; Wen, Y.; Wu, S.; Deng, H.; Li, Z.; et al. 2023b. Drivemlm: Aligning multi-modal large language models with behavioral planning states for autonomous driving. *arXiv preprint arXiv:2312.09245*.

Wen, L.; Fu, D.; Li, X.; Cai, X.; Ma, T.; Cai, P.; Dou, M.; Shi, B.; He, L.; and Qiao, Y. 2023. Dilu: A knowledge-driven approach to autonomous driving with large language models. *arXiv preprint arXiv:2309.16292*.

Wu, P.; Jia, X.; Chen, L.; Yan, J.; Li, H.; and Qiao, Y. 2022. Trajectory-guided control prediction for end-to-end autonomous driving: A simple yet strong baseline. In *NeurIPS*, 6119–6132.

Xu, Z.; Zhang, Y.; Xie, E.; Zhao, Z.; Guo, Y.; Wong, K. K.; Li, Z.; and Zhao, H. 2023. Drivegpt4: Interpretable end-to-end autonomous driving via large language model. *arXiv preprint arXiv:2310.01412*.

Yu, E.; Zhao, L.; Wei, Y.; Yang, J.; Wu, D.; Kong, L.; Wei, H.; Wang, T.; Ge, Z.; Zhang, X.; et al. 2023. Merlin: Empowering multimodal llms with foresight minds. *arXiv preprint arXiv:2312.00589*.

Zeng, W.; Luo, W.; Suo, S.; Sadat, A.; Yang, B.; Casas, S.; and Urtasun, R. 2019. End-to-end interpretable neural motion planner. In *CVPR*, 8660–8669.

Zhai, J.-T.; Feng, Z.; Du, J.; Mao, Y.; Liu, J.-J.; Tan, Z.; Zhang, Y.; Ye, X.; and Wang, J. 2023. Rethinking the open-loop evaluation of end-to-end autonomous driving in nuscenes. *arXiv preprint arXiv:2305.10430*.