

# VLP: Vision Language Planning for Autonomous Driving

Chenbin Pan<sup>1\*</sup> Burhaneddin Yaman<sup>2</sup>✉ Tommaso Nesti<sup>2</sup> Abhirup Mallik<sup>2</sup>  
 Alessandro G Allievi<sup>2</sup> Senem Velipasalar<sup>1</sup> Liu Ren<sup>2</sup>

<sup>1</sup>Syracuse University

<sup>2</sup>Bosch Research North America & Bosch Center for Artificial Intelligence (BCAI)

{cpan14, svelipas}@syr.edu,

{burhaneddin.yaman, tommaso.nesti, abhirup.mallik, alessandro.allievi, liu.ren}@us.bosch.com

## Abstract

Autonomous driving is a complex and challenging task that aims at safe motion planning through scene understanding and reasoning. While vision-only autonomous driving methods have recently achieved notable performance, through enhanced scene understanding, several key issues, including lack of reasoning, low generalization performance and long-tail scenarios, still need to be addressed. In this paper, we present VLP, a novel Vision-Language-Planning framework that exploits language models to bridge the gap between linguistic understanding and autonomous driving. VLP enhances autonomous driving systems by strengthening both the source memory foundation and the self-driving car’s contextual understanding. VLP achieves state-of-the-art end-to-end planning performance on the challenging NuScenes dataset by achieving 35.9% and 60.5% reduction in terms of average L2 error and collision rates, respectively, compared to the previous best method. Moreover, VLP shows improved performance in challenging long-tail scenarios and strong generalization capabilities when faced with new urban environments.

## 1. Introduction

Autonomous driving is a complex problem requiring scene understanding and reasoning to ensure safe motion planning. This sophisticated challenge can be broadly divided into three main tasks, namely perception, prediction and planning (also known as P3). Conventional methods adopt a modular approach by developing and optimizing each task in a disjoint manner without a holistic view, leading to compounding errors and safety concerns [15, 16, 18]. End-to-end autonomous driving systems (ADS), unifying all P3 tasks, have garnered attention for their potential to enhance




Figure 1. New-city generalization ability of ADS for planning is evaluated by training on Boston city and testing on Singapore, and vice versa. Our proposed VLP shows strong generalization ability by significantly outperforming state-of-the-art vision-only method, VAD[18], in terms of both L2 error and collision rate.

safe planning. Existing vision-based ADS typically follow a two-stage process: bird’s eye view (BEV) feature extraction and downstream tasks [16, 18]. BEV feature extraction transforms multi-view camera data into a structured top-view representation embedding spatial information around the ego-car. BEV features are further utilized as the information pool by the downstream P3 tasks. While effective, these vision-only methods struggle with out-of-domain generalization, such as maintaining performance in new cities and long-tail scenarios, hindering real-world deployment.

The way humans process and interpret visual driving scenes naturally involve a coherent cognitive framework that maintains a consistent logical flow. This enables humans effortlessly make correct decisions even when faced with previously unseen data or scenarios, eliminating generalization issues in driving tasks. Then, this raises a pertinent question: *How can ADS achieve human-like driving to navigate diverse and dynamic road environments?* Recent advancements in language models (LMs) have led to unprecedented common-sense ability and generalization performance across unseen data and tasks for the natural language processing domain [5, 8, 25]. The superior common-

\*Work done while interned at Bosch Research North America.

✉ Corresponding author.

sense capability of LMs have led to the emergence of multi-modal LMs for diverse applications ranging from medical imaging to robotics [4, 9, 33]. However, incorporation of the reasoning ability of LMs into real-world autonomous driving tasks, to address generalization and long-tail scenarios, is yet to be fully-explored.

To bridge this gap, we propose a **Vision Language Planning (VLP)** framework, which integrates the common-sense capability of LMs into vision-based ADS for safe self-driving. Our VLP consists of two key components: **Agent-centric Learning Paradigm (ALP)** and **Self-driving-car-centric Learning Paradigm (SLP)**, leveraging LMs to enhance the ADS from reasoning and decision-making aspects, respectively.

The BEV feature map serves as the source memory pool in ADS for downstream decoding tasks. It summarizes and encodes the driving environment surrounding the self-driving car, including vehicles, pedestrians, lanes, and more, into a unified feature map. Hence, capturing comprehensive and necessary details in each local position of BEV is critical for safe and precise self-driving performance. To enhance the local semantic representation and reasoning capabilities of BEV, we introduce an innovative **Agent-centric Learning Paradigm (ALP)** module. ALP integrates the consistent feature space of a pretrained LM to revamp the agent features on the BEV, actively shaping semantics, and guiding the BEV reasoning process. Leveraging on the common sense and logic flow embedded in the LM, our ALP equips the ADS with robustness and consistent BEV feature space, enhancing its effectiveness in diverse driving scenarios.

In ADS, the planning module aggregates information from the preceding perception and prediction phases to make the final decisions for self-driving. This global perspective culminates in the formation of a planning query, directly influencing the safety and accuracy of the self-driving navigation. Considering the critical role of the planning module within ADS, we also present a novel **Self-driving-car-centric Learning Paradigm (SLP)** to elevate the decoding and acquiring information ability of the planning query. In the SLP, we align the planning query with intended goals and the ego-vehicle driving status by leveraging the knowledge encoded in the pretrained LM. The language model’s comprehension capabilities contribute to more informed decision-making during the planning phase as well as enabling a more robust planning query formation process.

Through VLP, we bridge the gap between human-like reasoning and autonomous driving, enhancing the model’s contextual awareness and its ability to generalize effectively in complex, ever-changing real-world scenarios, as shown in Fig. 1. The main contributions of this work are summarized as follows:

- We propose VLP, a Vision Language Planning model,

which incorporates reasoning capability of LMs into vision-based autonomous driving systems as an enhancement of motion planning and self-driving safety.

- VLP is composed of novel components ALP and SLP, aiming to improve the ADS from self-driving BEV reasoning and self-driving decision-making aspects, respectively.
- Through extensive experiments in real-world driving scenarios, we show that VLP significantly and consistently outperforms the state-of-the-art vision-based approaches across a spectrum of driving tasks, including open-loop planning, multi-object tracking, motion forecasting, and more.
- We conduct the first new-city generalization study on the nuScenes dataset [6] by training and testing on distinct cities, demonstrating the remarkable zero-shot generalization ability of our VLP approach over vision-only methods.
- To the best of our knowledge, this is the first work introducing LM into multiple stages of ADS to address the generalization ability in new cities and long-tail cases.

## 2. Related Work

### 2.1. End-to-End Autonomous Driving

Early end-to-end approaches have adopted a single neural network to perform motion planning without explicitly designing intermediate tasks, such as prediction [3]. Thus, such approaches not only suffer from the sub-optimal performance but also lack interpretability. Modular end-to-end frameworks introduce intermediate tasks for interpretability and provide improved performance [7, 10, 15, 16, 18, 23, 28]. These interpretable end-to-end frameworks explicitly model perception, prediction and planning (P3) components and train them together with a joint optimization strategy to enhance planning. UniAD [16] and VAD [18] are two prominent state-of-the-art end-to-end frameworks. UniAD leverages rasterized scene representations and explicitly identifies crucial components within the P3 framework. It coordinates all tasks through a unified query design for safe planning. On the other hand, VAD uses vectorized representations for efficient planning and safety improvement. While these approaches generally achieve good performance, their generalization performance have been limited due to lack of common sense and reasoning process.

### 2.2. Vision-Language Models

Large language models (LLMs), which are trained on massive amounts of text data, have shown unprecedented success in language-related tasks [5, 8, 29]. The capability of LLMs for common-sense understanding and versatility in tackling diverse language tasks has driven the exploration of multi-modal LLMs [20, 21, 27]. Vision-language models (VLM) integrate visual and textual information to en-

able common sense understanding of the content by training on large scale open-world vision and text data [2, 17, 20, 27]. CLIP [27] uses a contrastive learning objective to learn a joint embedding space for vision and text data, and has shown immense zero-shot generalization performance, showcasing the importance of VLMs compared to vision-only approaches.

### 2.3. Autonomous Driving with Language Models

In recent years, numerous works have been proposed for extending LMS to the embodied AI domain to improve zero-shot generalization performance [4, 9, 30, 35]. These works incorporate sensory inputs from embodied agents in the form of language with vision and language data for decision making process [4, 9]. While these methods have mainly focused on the robotics domain, few works have been proposed for leveraging embodied LMs for autonomous driving tasks [11, 24, 31, 32, 34]. In particular, DiLu [32] and GPT-Driver [24] propose GPT-based driver agents for closed-loop simulation tasks. In [31], an open loop driving commentator, which combines vision and low-level driving actions with language, is proposed to interpret driving actions and for reasoning. Our work significantly differs from these approaches by leveraging LMs for end-to-end motion planning task for real-world driving applications.

## 3. Methodology

We present a Vision Language Planning (VLP) model, which equips the ADS with the capacity to emulate human common sense and engage in contextual reasoning for safe motion planning. Our proposed VLP model, illustrated in Fig. 2a, comprises two innovative components that leverage LMs in both local and global contexts. The Agent-centric Learning Paradigm (ALP) concentrates on refining local details to enhance source memory reasoning, while the Self-driving-car-centric Learning Paradigm (SLP) focuses on guiding the planning process for the self-driving-car (SDC). Below, we first provide an overview of how general ADS work in Sec.3.1. Subsequently, we describe our proposed ALP and SLP in detail in Sec. 3.2 and Sec. 3.3, respectively.

### 3.1. Preliminary

In vision-based ADS, a sequence of multi-view camera images serves as the source input, offering abundant visual data for downstream tasks. Initially, an image backbone extracts fundamental features for perception. These features are then processed by a BEV encoder to transform multi-view features into a unified 2D representation from a top-down view. This BEV feature map encapsulates crucial perception details like object positions, lane markings, and road boundaries [22]. Acting as an information hub, this BEV feature map supports diverse downstream tasks such

as 3D object tracking, mapping, motion prediction, occupancy prediction, and path planning [16, 18]. Specific task-related information is decoded using a query-based transformer decoder designed for each task to efficiently access the source memory [16]. In the final planning stage, an ego-query models the ego-vehicle’s status and interacts with environmental features to determine the optimal path for autonomous driving.

### 3.2. Agent-centric Learning Paradigm: Enhancing BEV Source Memory

The BEV feature map serves as the foundational source memory pool and holds a pivotal role in ADS. Ensuring that a BEV map provides comprehensive information with intricate details for driving is essential for enabling downstream decoders to make safe, precise, and human-like decisions. However, the BEV in ADS is derived from multi-view camera images rather than a true bird’s-eye-view image, which can introduce discrepancies between the produced BEV and the expected BEV representation. Therefore, to mitigate the discrepancies, we introduce the Agent-centric Learning Paradigm (ALP) to align the produced BEV with a true bird’s-eye-view map. Through direct supervision on the BEV map with LM, our ALP enables the refinement of local details and alignment with the desired human perspective.

**Agent BEV features.** In our proposed ALP, three kinds of BEV agents are considered: ego-vehicle (self-driving-car), foreground (FG) objects, and lane elements. We first align the ground-truth area of each agent with the produced BEV map, and crop the regions of interest. We utilize the 3D bounding box to crop the ego-car and FG object area, and panoptic scene mask to segment the lane area. Subsequently, we perform a pooling operation on the obtained local BEV region, to generate a single feature representation for the corresponding agent. After pooling, the local agent features in each sample along the batch are concatenated to formulate an Agent-BEV tensor denoted as  $\mathcal{A}_{bev} \in \mathbb{R}^{N_B \times C}$ , where  $N_B$  and  $C$  denote total number of agents in the batch and feature dimension, respectively. The process can be formulated as:

$$\begin{aligned} \mathcal{A}_{bev}^{Ego} &= \text{Pool}(\text{Crop}(BEV, \text{bbox}_{3D}^{Ego})), \in \mathbb{R}^C, \\ \mathcal{A}_{bev}^{FG} &= \text{Pool}(\text{Crop}(BEV, \text{bbox}_{3D}^{FG})), \in \mathbb{R}^C, \\ \mathcal{A}_{bev}^{Lane} &= \text{Pool}(\text{Seg}(BEV, \text{SegMask}^{Lane})), \in \mathbb{R}^C, \\ \mathcal{A}_{bev} &= \text{Batch}([\mathcal{A}_{bev}^{Ego}; \mathcal{A}_{bev}^{FG}; \mathcal{A}_{bev}^{Lane}]), \in \mathbb{R}^{N_{batch} \times C}, \end{aligned} \quad (1)$$

where  $\text{Batch}[\cdot]$  denotes the concatenation operation,  $\mathcal{A}_{bev}^{(\cdot)}$  represents the single agent feature of the corresponding type,  $\text{bbox}_{3D}^{(\cdot)}$  and  $\text{SegMask}^{Lane}$  indicate the ground-truth 3D bounding box and the lane segmentation mask, respectively.

**Agent expectation features.** To ensure that local BEV features express the desired information, we conduct a BEV-

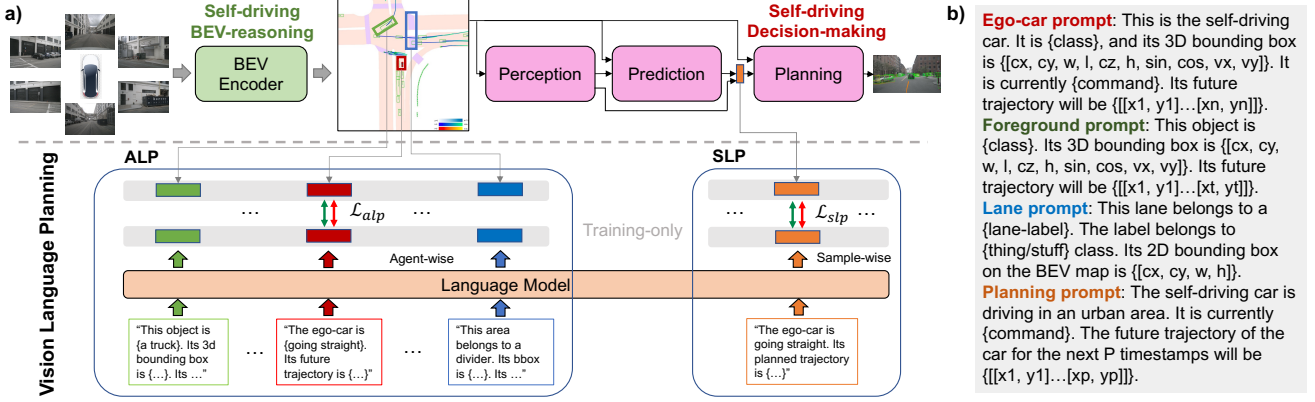


Figure 2. a) The overview of proposed vision language planning (VLP) framework. VLP enhances ADS from self-driving BEV-reasoning and self-driving decision-making aspects, through two innovative modules, ALP and SLP, respectively. Leveraging LM and contrastive learning, ALP conducts agent-wise learning for refining local details on BEV, while SLP engages sample-wise learning for advancing global context understanding ability of the ADS. VLP is only active during training, ensuring no additional parameters or computations are introduced during inference. b) Prompt formats used in VLP.

expectation alignment process by leveraging LM and contrastive learning. We precisely define the perceptual information expected from the corresponding agent, such as agent labels, bounding boxes, and future trajectories. These driving-related ground-truth information, which should also be embedded in the local BEV feature is formulated into a prompt as illustrated in Fig. 2b. The description is then passed to the language encoder,  $LM$ , to generate the corresponding agent expectation feature. We apply an MLP layer  $\mathcal{F}_{bev}$  to adapt the expectation feature to the BEV feature space. Then, the agent expectation features are concatenated along the batch to generate an Agent-Expectation tensor denoted as  $\mathcal{A}_{exp} \in \mathbb{R}^{N_B \times C}$ . The procedure can be formulated as in Eq. (2):

$$\begin{aligned}
 \mathcal{A}_{exp}^{Ego} &= \mathcal{F}_{bev}(LM(\mathbf{T}^{Ego}[y^{Ego}])), \in \mathbb{R}^C, \\
 \mathcal{A}_{exp}^{FG} &= \mathcal{F}_{bev}(LM(\mathbf{T}^{FG}[y^{FG}])), \in \mathbb{R}^C, \\
 \mathcal{A}_{exp}^{Lane} &= \mathcal{F}_{bev}(LM(\mathbf{T}^{Lane}[y^{Lane}])), \in \mathbb{R}^C, \\
 \mathcal{A}_{exp} &= \text{Batch}([\mathcal{A}_{exp}^{Ego}; \mathcal{A}_{exp}^{FG}; \mathcal{A}_{exp}^{Lane}]),
 \end{aligned} \quad (2)$$

where  $\mathcal{A}_{exp}^{(\cdot)}$ ,  $\mathbf{T}^{(\cdot)}$ , and  $y^{(\cdot)}$  represent the single agent expectation feature, the description template, and the ground-truth in ADS for the respective agent. Note that during training, we freeze the  $LM$  and only train the adaptation layer  $\mathcal{F}_{bev}$  to save memory and retain the pre-trained knowledge in the  $LM$ .

**Contrastive Learning Formulation in ALP.** Given the agent BEV feature  $\mathcal{A}_{bev}$  and the agent expectation feature  $\mathcal{A}_{exp}$ , the alignment between the produced BEV and the expected BEV is facilitated through a contrastive learning process [27]. In particular, L2 normalization is applied to both  $\mathcal{A}_{bev}$  and  $\mathcal{A}_{exp}$  to standardize the feature vectors and ensure training stability. Then, matrix multiplication is performed between the normalized two-modality tensors with

a learnable logit scale  $\alpha_{alp}$  to produce a similarity matrix  $\mathcal{S}_{pred} \in \mathbb{R}^{N_B \times N_B}$  in agent-wise. The ground truth for the similarity matrix,  $\mathcal{S}_{gt}$ , is a unit matrix, signifying that each agent's BEV feature should be closest to its corresponding expected feature in the shared space. To optimize this alignment, we apply a cross-entropy loss,  $\mathcal{L}_{CE}$ , along both the Agent-BEV mode axis and Agent-Expectation mode axis of the  $\mathcal{S}_{pred}$ , and compute an average value as the final ALP loss,  $\mathcal{L}_{alp}$ , for the BEV source reasoning part. The entire process can be formally expressed as:

$$\begin{aligned}
 \mathcal{S}_{pred} &= \alpha_{alp} \times \left( \frac{\mathcal{A}_{bev}}{\|\mathcal{A}_{bev}\|_2} \otimes \frac{\mathcal{A}_{exp}}{\|\mathcal{A}_{exp}\|_2} \right), \\
 \mathcal{L}_{alp} &= (\mathcal{L}_{CE}(\mathcal{S}_{pred}, \mathcal{S}_{gt}, \text{dim} = 0) \\
 &\quad + \mathcal{L}_{CE}(\mathcal{S}_{pred}, \mathcal{S}_{gt}, \text{dim} = 1))/2.
 \end{aligned} \quad (3)$$

### 3.3. SDC-centric Learning Paradigm: Enhancing Ego-vehicle Query

**Ego-vehicle query feature.** The ego-vehicle dynamic modeling is the core function of the ADS. In previous ADS, a trainable ego-query  $\mathcal{E}_{query}$  is applied to interact with other agents,  $\mathcal{A}_{query}^{FG}$  and  $\mathcal{A}_{query}^{Lane}$ , on BEV map, to gather the self-driving perception/prediction information. The produced ego-car query feature  $\mathcal{E}_{qfeat}$  is further processed to predict the future waypoints of the self-driving car  $y_{pred}^{plan}$ . The process is formulated as in Eq. (4):

$$\begin{aligned}
 \mathcal{E}_{qfeat} &= \mathcal{M}^{inter}(\mathcal{E}_{query}, \mathcal{A}_{query}^{FG}, \mathcal{A}_{query}^{Lane}) \in \mathbb{R}^{B \times C}, \\
 y_{pred}^{plan} &= \mathcal{M}^{plan}(\mathcal{E}_{qfeat}) \in \mathbb{R}^{B \times P \times 2},
 \end{aligned} \quad (4)$$

where  $\mathcal{M}^{inter}$  and  $\mathcal{M}^{plan}$  denote the intermediate ego information gathering module and the planning head in ADS, respectively.  $P \times 2$  indicate the planned waypoints in the next  $P$  timesteps, and  $B$  is the batch size during training.

Although such mechanism can achieve good planning results, solely relying on numeric ground-truth can make it challenging to understand the rationale behind the ADS’ decisions, which can result in inconsistent feature learning for the planning and lead to limited adaptability to new environments. To address these limitations, we propose a SDC-centric Learning Paradigm (SLP) to equip the ADS with the capability of making decisions from continuous and robust feature space.

**Ego-vehicle planning feature.** We template a language description of the ego-vehicle status with the planning ground truth (GT) including high-level driving command and the future trajectory of the ego-vehicle, as illustrated in Fig. 2b. The prompt is sent to the *LM* to obtain the ground truth planning features of the ego-car embedding driving scenario information and human driving logic. An MLP layer  $\mathcal{F}_{ego}$  is applied to adapt the textual planning features to the ego-query feature space. The process can be expressed as in Eq. (5):

$$\mathcal{E}_{prompt} = \mathcal{F}_{ego}(LM(\mathbf{T}_{slp}^{Ego}[y^{plan}])) \in \mathbb{R}^{B \times C}, \quad (5)$$

where  $\mathbf{T}_{slp}^{Ego}$  indicates the prompt for ego-car description used in SLP, and  $y^{plan}$  denotes the planning ground truth. Note that the *LM* is a shared off-the-shelf model for both ALP and SLP. As in ALP, only the adaptation layer  $\mathcal{F}_{ego}$  is trainable during training to save memory and retain the pre-trained knowledge in *LM*.

**Contrastive Learning Formulation in SLP.** Similar to ALP, we employ a sample-wise contrastive learning approach for SLP to align the ego-vehicle query feature with the ego-vehicle textual planning feature as follows:

$$\begin{aligned} \mathcal{S}_{pred}^{slp} &= \alpha_{slp} \times \left( \frac{\mathcal{E}_{feat}}{\|\mathcal{E}_{feat}\|_2} \otimes \frac{\mathcal{E}_{prompt}}{\|\mathcal{E}_{prompt}\|_2} \right), \\ \mathcal{L}_{slp} &= (\mathcal{L}_{CE}(\mathcal{S}_{pred}^{slp}, \mathcal{S}_{gt}^{slp}, dim = 0) \\ &\quad + \mathcal{L}_{CE}(\mathcal{S}_{pred}^{slp}, \mathcal{S}_{gt}^{slp}, dim = 1))/2, \end{aligned} \quad (6)$$

where  $\alpha_{slp}$  is a learnable logit scale, and  $\mathcal{S}_{pred}^{slp}$  and  $\mathcal{S}_{gt}^{slp}$  are the predicted and ground truth similarity matrices, respectively. Aligning the two modes closely in the feature space, the contrastive process refines the produced ego-vehicle query using the ground truth-embedded linguistic feature. The sample-wise contrastive loss rectifies the relationships between individual samples from a human (linguistic) perspective, establishing a connection for the current data point with the world-wise common sense embedded in the pretrained language model.

### 3.4. Training Loss

The overall loss,  $\mathcal{L}_{vlp}$ , of our VLP training is composed of two parts: the BEV encoder reasoning loss  $\mathcal{L}_{enc}$  and the decoder decision-making loss  $\mathcal{L}_{dec}$  such that:

$$\mathcal{L}_{vlp} = \omega_{enc}\mathcal{L}_{enc} + \omega_{dec}\mathcal{L}_{dec}, \quad (7)$$

where  $\omega_{enc}$  and  $\omega_{dec}$  represent the weights for  $\mathcal{L}_{enc}$  and  $\mathcal{L}_{dec}$ , respectively.

The encoder reasoning loss,  $\mathcal{L}_{enc}$ , is equal to the loss produced in ALP module, i.e.  $\mathcal{L}_{enc} = \mathcal{L}_{alp}$ . The decision-making loss includes the losses in all downstream tasks (perception/prediction/planning) in ADS and the SLP loss:

$$\mathcal{L}_{dec} = \mathcal{L}_{perc} + \mathcal{L}_{pred} + \mathcal{L}_{plan} + \mathcal{L}_{slp}. \quad (8)$$

Note that the off-the-shelf LM head is discarded during inference, and thus, our method introduces no additional parameters and computations to the original ADS.

## 4. Experiments

### 4.1. Implementation Details

**Dataset.** We conduct experiments on the challenging nuScenes dataset [6], which is the first large-scale public dataset to provide data from the entire sensor suite of an autonomous vehicle (6 cameras, 1 LIDAR, 5 RADAR, GPS and IMU sensors). The nuScenes contains 1000 driving scenes from Boston and Singapore, two cities that are known for their dense traffic and highly challenging driving conditions. In our experiments, we utilize 6 camera images as our vision inputs.

**Baselines.** We adopt two recent SOTA models in autonomous driving, namely UniAD [16] and VAD [18], as our baselines to evaluate the effectiveness of our approach.

**Training.** We incorporate our proposed VLP training strategy into vision-only approaches of UniAD [16] and VAD [18], and refer to this enriched versions as VLP-UniAD and VLP-VAD, respectively, for clarity when reporting the results. All methods are trained with the same hyper-parameters reported in the respective baselines [16, 18] for commensurate comparison. For language model, we use the pretrained LM in CLIP [27]. In VLP study, we set T=6 (3 seconds) for planning, which is same as baselines [16] [18]. Experiments are conducted with 8 NVIDIA Tesla A100 GPUs. More details regarding experimental setup can be found in supplementary materials.

### 4.2. Open-loop Planning

Planning constitutes the cornerstone of any ADS, playing a pivotal role in ensuring safety and devising efficient routes for the ego-car. In Tab. 1, we present a series of comparative experiments that showcase the performance of our open-loop planning in comparison to the baseline models. As can be seen in rows 4-6 of the table, the integration of just SLP leads to noticeable reductions in both the L2 error and collision rates for all the baseline models. Moving down the table, rows 7 to 9 demonstrate that the inclusion

of both VLP components (SLP and ALP together) consistently yields further improvements in these planning metrics. In particular, VLP-UniAD shows a 28.1% and 48.4% reduction in terms of average L2 error and collision rate, respectively, compared to baseline UniAD. Similarly, in comparison with VAD, VLP-VAD achieves 35.9% and 60.5% reduction for average L2 error and collision rate, respectively. These significant results underscore the effectiveness of both SLP and ALP, as well as their adaptability across various ADS configurations. The reduced L2 error and collision rate achieved through VLP integration contribute to safer road planning in the realm of self-driving.

### 4.3. Perception and Prediction

In this section, we showcase the consistent effectiveness of our proposed VLP across a spectrum of essential driving tasks. Our proposed approach excels in various perception and prediction tasks, including multi-object tracking, mapping, motion forecasting, occupancy prediction, 3D object detection, and vectorized scene segmentation. These results underscore the versatility and reliability of our VLP framework across a wide range of critical driving tasks.

**Multi-Object Tracking.** The results for multi-object tracking (MOT) are presented in Tab. 2. We apply the standard evaluation protocols of AMOTA (Average Multi-Object Tracking Accuracy), AMOTP (Average Multi-Object Tracking Precision), Recall, and IDS (Identity Switches) to evaluate the MOT performance. As can be seen, with the incorporation of SLP, AMOTA, Recall, and IDS values all surpass those of UniAD, while AMOTP experiences a slight degradation. The combined integration of both SLP and ALP results in improvements across all metrics when compared to the UniAD baseline. The enhanced performance in MOT proves that our VLP can help the system better predict and respond to the movements of other objects on the road, reducing the risk of accidents.

**Online Mapping.** In Tab. 3, we present the results of online mapping, encompassing four key mapping elements: lanes, drivable areas, dividers, and pedestrian crossings. The evaluation employs the Intersection-over-Union (IoU) metric to measure the overlap between the predicted and ground-truth maps. The results reveal that the inclusion of the SLP leads to improvements in mapping accuracy for all four elements. Specifically, it enhances the IoU for drivable areas and crossings by 5.5% and 6.5%, respectively. Furthermore, the integration of both SLP and ALP yields even further improvement, with IoU achieving by 6.6% and 10.2% for drivable areas and crossings, respectively. These enhancements represent substantial gains compared to the SLP-only configuration. These results underscore the valuable contributions of each paradigm in online mapping. They also demonstrate the effectiveness of VLP in bridging the gap between vision-based and language-based information, ul-

timately enhancing the system’s comprehension of various road elements.

**Motion Forecasting.** We present the motion forecasting results in Tab. 4, applying the same evaluation metrics as those in UniAD [16], namely minADE (minimum Average Displacement Error), minFDE (minimum Final Displacement Error), MR (Miss Rate), and EPA (End-to-end Prediction Accuracy). The results demonstrate that the inclusion of the SLP contributes to a reduction in minADE, minFDE, and MR, while simultaneously increasing EPA for the ADS. This outcome signifies that the supervision and refinement applied to the final ego-car feature also has a positive influence on the motion prediction of other vehicles around the ego-car. The improvements are more obvious with the EPA metrics. Equipped with full components of VLP, the performance of the ADS, in motion prediction, experiences further enhancements. This observation underscores the capability of ALP to empower the system in the identification and accurate prediction of the movements of various agents on the road.

**Occupancy Prediction.** We present the occupancy prediction results in Tab. 5, which have been obtained using the IoU and Video Panoptic Quality (VPQ) within two distance ranges around the ego-car, namely near (“-n.”), covering a  $30 \times 30m$  area, and far (“-f.”), spanning a  $50 \times 50m$  area. The results highlight the effectiveness of the SLP in consistently improving all four metrics. SLP, primarily focusing on enhancing the ego-car feature, demonstrates consistent performance gains in occupancy prediction. The inclusion of ALP with SLP shows similar improvement over the baseline as well. Since both configurations of the VLP prioritize attention to the ego-car during training, more improvements are observed within the nearby areas.

**3D Object Detection.** We show the 3D object detection results in Tab. 6 using nuScenes detection metrics of mean Average Precision (mAP), Average Translation Error (ATE), Average Scale Error (ASE), and nuScenes detection score (NDS) as indicated in [6]. As observed in the table, the incorporation of the VLP consistently leads to improvements in NDS. Notably, models incorporating both the SLP and ALP generally perform better than those with SLP alone, providing empirical evidence of the ALP’s effectiveness in enhancing the local BEV representation ability.

**Vectorized Scene Segmentation in VAD.** The VAD [18] framework represents the driving scene as a fully vectorized structure, categorizing map elements into road boundaries, dividers, and pedestrian crossings. Thus, vectorized scene segmentation experiments are conducted within the VAD framework, with IoU serving as the evaluation metric. As depicted in Tab. 7, the introduction of SLP yields substantial improvements. More specifically, the segmentation for boundaries, dividers, and crossings sees improvements of 14.9%, 19.2%, and 23.1%, respectively, contributing to an

ID	Model	SLP	ALP	L2 (m) ↓				Col. Rate (%) ↓			
				1s	2s	3s	Avg.	1s	2s	3s	Avg.
0	NMP [36]			-	-	2.31	-	-	-	1.92	
1	SA-NMP [36]			-	-	2.05	-	-	-	1.59	
2	FF [14]			0.55	1.20	2.54	1.43	0.06	0.17	1.07	0.43
3	EO [19]			0.67	1.36	2.78	1.60	0.04	0.09	0.88	0.33
4	ST-P3 [15]			1.33	2.11	2.90	2.11	0.23	0.62	1.27	0.71
5	UniAD [16]			0.48	0.96	1.65	1.03	0.05	0.17	0.71	0.31
6	VLP-UniAD	✓		<b>0.43</b>	<b>0.86</b>	<b>1.47</b>	<b>0.92</b>	<b>0.03</b>	<b>0.15</b>	<b>0.48</b>	<b>0.22</b>
7	VLP-UniAD	✓	✓	<b>0.36</b>	<b>0.68</b>	<b>1.19</b>	<b>0.74</b>	<b>0.03</b>	<b>0.12</b>	<b>0.32</b>	<b>0.16</b>
8	VAD [18]			0.46	0.76	1.12	0.78	0.21	0.35	0.58	0.38
9	VLP-VAD	✓		<b>0.26</b>	<b>0.47</b>	<b>0.78</b>	<b>0.50</b>	<b>0.12</b>	<b>0.17</b>	<b>0.42</b>	<b>0.23</b>
10	VLP-VAD	✓	✓	<b>0.30</b>	<b>0.53</b>	<b>0.84</b>	<b>0.55</b>	<b>0.01</b>	<b>0.07</b>	<b>0.38</b>	<b>0.15</b>

Table 1. **Open-loop planning performance.** VLP achieves significant end-to-end planning performance improvement over counterpart vision only UniAD and VAD methods on the nuScenes validation dataset [6]. Based on the planning results, we can conclude that both SLP and ALP components plays a vital role to ensure safe motion planning.

Model	SLP	ALP	AMOTA↑	AMOTP↓	Recall↑	IDS↓
ViP3D [10]			21.7	1.625	36.3	-
QD3DT [13]			24.2	1.518	39.9	-
MUTR3D [37]			29.4	1.498	42.7	3822
UniAD [16]			35.9	1.320	46.7	906
VLP-UniAD	✓		<b>36.6</b>	1.332	<b>46.8</b>	<b>820</b>
VLP-UniAD	✓	✓	<b>36.8</b>	<b>1.315</b>	<b>47.3</b>	<b>678</b>

Table 2. **Multi-object tracking.** VLP shows improved performance over vision-only MOT techniques.

Model	SLP	ALP	Lanes↑	Drivable↑	Divider↑	Crossing↑
UniAD [16]			31.3	69.1	25.7	13.8
VLP-UniAD	✓		<b>32.0</b>	<b>69.5</b>	<b>27.1</b>	<b>14.7</b>
VLP-UniAD	✓	✓	<b>32.3</b>	<b>70.2</b>	<b>27.4</b>	<b>15.2</b>

Table 3. **Online mapping.** VLP demonstrates improved segmentation IoU (%) performance over the UniAD.

Model	SLP	ALP	minADE(m)↓	minFDE(m)↓	MR↓	EPA↑
PnPNet [23]			1.15	1.95	0.226	0.222
ViP3D [10]			2.05	2.84	0.246	0.226
UniAD [16]			0.71	1.02	0.151	0.456
VLP-UniAD	✓		0.72	1.04	0.154	0.459
VLP-UniAD	✓	✓	<b>0.68</b>	<b>0.98</b>	<b>0.133</b>	<b>0.460</b>
VAD [18]			0.78	1.07	0.121	0.598
VLP-VAD	✓		<b>0.77</b>	<b>1.03</b>	<b>0.110</b>	<b>0.621</b>
VLP-VAD	✓	✓	<b>0.77</b>	<b>1.05</b>	<b>0.115</b>	<b>0.621</b>

Table 4. **Motion forecasting.** VLP achieves better motion forecasting over counterpart vision based methods.

overall increase of 18.6% in the mean IoU (mIoU). VLP with SLP and ALP also achieves similar improvement over the VAD baseline. This increased accuracy is crucial for developing a safe and efficient self-driving system.

#### 4.4. Generalization

Autonomous vehicles are intended to operate in diverse environments. Generalization allows a system to be deployed in various urban landscapes, suburban areas, or even rural

Model	SLP	ALP	IoU-n.↑	IoU-f.↑	VPQ-n.↑	VPQ-f.↑
FIERY [12]			59.4	36.7	50.2	29.9
StretchBEV [1]			55.5	37.1	46.0	29.0
ST-P3 [15]			-	38.9	-	32.1
UniAD [16]			63.4	40.2	54.7	33.5
VLP-UniAD	✓		<b>64.2</b>	<b>40.7</b>	<b>55.8</b>	<b>34.5</b>
VLP-UniAD	✓	✓	<b>64.1</b>	40.2	<b>55.9</b>	<b>34.1</b>

Table 5. **Occupancy prediction.** VLP shows improvement in both near ( $30 \times 30m$ ) and far ( $50 \times 50m$ ) ranges, denoted as “n.” and “f.”, respectively.

Model	SLP	ALP	mAP↑	mATE↓	mASE↓	NDS↑
VAD [18]			0.27	0.70	0.30	0.389
VLP-VAD	✓		0.27	<b>0.67</b>	0.30	<b>0.394</b>
VLP-VAD	✓	✓	<b>0.28</b>	<b>0.67</b>	0.30	<b>0.406</b>

Table 6. **3D object detection.** VLP achieves improved object detection performance over the baseline.

Model	SLP	ALP	Boundary↑	Divider↑	Crossing↑	mIoU↑
VAD [18]			45.6	42.2	31.6	39.8
VLP-VAD	✓		<b>52.4</b>	<b>50.3</b>	<b>38.9</b>	<b>47.2</b>
VLP-VAD	✓	✓	<b>49.4</b>	<b>48.4</b>	<b>39.2</b>	<b>45.7</b>

Table 7. **Vectorized scene segmentation.** VLP shows improved segmentation IoU (%) performance over the VAD.

settings, making the technology applicable and accessible on a broader scale. Long-tail cases, which are scenarios that occur infrequently, are often underrepresented in training data. Generalization helps the system cope with these rare but critical situations, reducing the risk of biased decision-making based on inadequate exposure during training. In this section, we evaluate the VLP in terms of generalization ability to new cities and long-tail cases.

**New-city Generalization.** To keep the same number of multi-view camera inputs, we construct a multi-city dataset exclusively from nuScenes, by encompassing data from two

Model	VLP	Boston		Singapore	
		Avg.L2↓	Avg.Col↓	Avg.L2↓	Avg.Col↓
UniAD [16]		1.24	0.32	1.05	0.37
VLP-UniAD	✓	<b>1.14</b>	<b>0.26</b>	<b>0.87</b>	<b>0.34</b>
VAD [18]		0.86	0.27	0.78	0.39
VLP-VAD	✓	<b>0.73</b>	<b>0.22</b>	<b>0.63</b>	<b>0.20</b>

Table 8. **New-city generalization.** To evaluate the generalization performance, we train the model on Boston and test it on Singapore, and vice versa. VLP shows remarkable zero-shot generalization performance improvement over the vision-only methods.

distinct urban environments, Boston and Singapore. To comprehensively assess generalizability, we conduct two sets of experiments: (i) training on Boston and testing on Singapore; and (ii) conversely, training on Singapore and testing on Boston. The results, presented in Tab. 8, reveal a noteworthy reduction in planning L2 error and collision rates for the baselines in both scenarios with the integration of VLP. In particular, VLP-VAD achieves strong generalization over counterpart vision-only VAD approach (in Boston: 15.1% and 18.5%, in Singapore: 19.2% and 48.7% reduction in terms of average L2 error and collision rates, respectively). This signifies VLP’s prowess in enhancing the safety and reliability of the ADS as well as its capacity to transcend training confines and excel in diverse real-world conditions.

**Long-tail Generalization.** We assess the performance of UniAD in Multi-Object Tracking (MOT) and VAD in 3D detection concerning long-tail scenarios, as detailed in Tab. 9a and Tab. 9b, respectively. Utilizing the long-tail split methodology from [26], derived from per-class object counts, we observe that VLP consistently enhances the generalization capabilities of both frameworks. Specifically, VLP provides substantial improvements, increasing AMOTA and Recall by 3.7% and 7.4%, respectively, for UniAD. In the case of VAD, VLP improves the mean Average Precision (mAP) by 15.9%. These results highlights the robust feature space cultivated by our proposed VLP, emphasizing its efficacy in handling long-tail scenarios.

#### 4.5. Ablation Studies

**Effectiveness of SLP and ALP.** VLP comprises ALP and SLP components. Extensive experiments presented in prior tables show the importance and necessity of both SLP and ALP for enhanced planning and safety.

**Prompt Format.** The design of the prompt format is pivotal for the success of our VLP, serving as the structured input that guides the learning process. We have performed ablation studies to systematically investigate the impact of different components of the ground truth information included in the prompt. The driving ground truth encompasses labels, bounding boxes, trajectories, and high-level commands. To demonstrate the indispensability of including all aspects of

Model	VLP	AMOTA↑	AMOTP↓	Recall↑	IDS↓
UniAD[16]		29.6	1.446	39.2	68
VLP-UniAD	✓	<b>30.7</b>	<b>1.435</b>	<b>42.1</b>	<b>66</b>

(a) Multi-object tracking on long-tail cases.

Model	VLP	mAP↑	mATE↓	mASE↓	mAOE↓
VAD [18]		17.6	0.79	0.33	0.95
VLP-VAD	✓	<b>20.4</b>	<b>0.75</b>	0.33	<b>0.83</b>

(b) Object detection on long-tail cases.

Table 9. **Long-tail generalization.** VLP shows consistent improvement for long-tail cases across both a) multi-object tracking and b) object detection tasks.

label-gt	bbox-gt	traj-gt	commd-gt	Avg.L2↓	Avg.Col↓
✗				0.64	0.23
	✗			0.56	0.30
		✗		0.59	0.36
			✗	0.59	0.26
✓	✓	✓	✓	<b>0.52</b>	<b>0.17</b>

Table 10. **Ablation for prompt information.** All ground-truth information components contributes to improved planning performance. We provide full task ablation results in the supplementary.

the ground truth in the prompt, we conducted experiments where each element was selectively removed in individual settings. The results, presented in Tab. 10, clearly indicate that the exclusion of any ground truth component leads to a degradation in planning ability. This underscores the necessity of incorporating the entire spectrum of ground truth information in the prompt for optimal performance.

## 5. Conclusion

We have introduced a novel Vision-Language-Planning (VLP) approach to enhance the capabilities of Autonomous Driving Systems (ADS). Our approach leverages both self-driving-car-centric learning paradigm (SLP) and agent-wise learning paradigm (ALP) guided by language prompts to create a comprehensive understanding of the environment. Through a series of experiments on various driving tasks, we have demonstrated the effectiveness of our VLP approach in improving perception, prediction, and planning aspects of ADS. The generalization experiments showcased the robustness of our VLP approach, proving its adaptability to new cities and long-tail cases. By extending the capabilities of ADS beyond the training environment, our VLP approach paves the way for safer and more reliable autonomous driving in real-world conditions.

**Limitations.** Our experiments are currently confined to the nuScenes dataset and camera modality as baseline vision-based approaches. We will assess VLP on a broader range of datasets and sensor modalities in our future work.



## References

- [1] Adil Kaan Akan and Fatma Güney. Stretchbev: Stretching future instance prediction spatially and temporally. In *European Conference on Computer Vision*, pages 444–460. Springer, 2022. 7
- [2] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736, 2022. 3
- [3] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Prasoon Goyal, Lawrence D Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016. 2
- [4] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. *arXiv preprint arXiv:2307.15818*, 2023. 2, 3
- [5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. 1, 2
- [6] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nusenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. 2, 5, 6, 7
- [7] 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
- [8] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022. 1, 2
- [9] Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, et al. Palm-e: An embodied multimodal language model. *arXiv preprint arXiv:2303.03378*, 2023. 2, 3
- [10] 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, 7
- [11] Md Zahid Hasan, Jiajing Chen, Jiyang Wang, Ameya Joshi, Senem Velipasalar, Chinmay Hegde, Anuj Sharma, and Soumik Sarkar. Vision-language models can identify distracted driver behavior from naturalistic videos. *arXiv preprint arXiv:2306.10159*, 2023. 3
- [12] 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. 7
- [13] Hou-Ning Hu, Yung-Hsu Yang, Tobias Fischer, Trevor Darrell, Fisher Yu, and Min Sun. Monocular quasi-dense 3d object tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2):1992–2008, 2022. 7
- [14] Peiyun Hu, Aaron Huang, John Dolan, David Held, and Deva Ramanan. Safe local motion planning with self-supervised freespace forecasting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12732–12741, 2021. 7
- [15] 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. 1, 2, 7
- [16] 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. 1, 2, 3, 5, 6, 7, 8
- [17] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pages 4904–4916. PMLR, 2021. 3
- [18] Bo Jiang, Shaoyu Chen, Qing Xu, Bencheng Liao, Jiajie Chen, Helong Zhou, Qian Zhang, Wenyu Liu, Chang Huang, and Xinggang Wang. Vad: Vectorized scene representation for efficient autonomous driving. *arXiv preprint arXiv:2303.12077*, 2023. 1, 2, 3, 5, 6, 7, 8
- [19] Tarasha Khurana, Peiyun Hu, Achal Dave, Jason Ziglar, David Held, and Deva Ramanan. Differentiable raycasting for self-supervised occupancy forecasting. In *European Conference on Computer Vision*, pages 353–369. Springer, 2022. 7
- [20] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR, 2022. 2, 3
- [21] Liunian Harold Li, Mark Yatskar, D Yin, CJ Hsieh, and KW Chang. Visualbert: A simple and performant baseline for vision and language. arxiv 2019. *arXiv preprint arXiv:1908.03557*. 2
- [22] 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. 3
- [23] Ming Liang, Bin Yang, Wenyuan Zeng, Yun Chen, Rui Hu, Sergio Casas, and Raquel Urtasun. Pnpnet: End-to-end per-

- ception and prediction with tracking in the loop. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11553–11562, 2020. 2, 7
- [24] Jiageng Mao, Yuxi Qian, Hang Zhao, and Yue Wang. Gpt-driver: Learning to drive with gpt. *arXiv preprint arXiv:2310.01415*, 2023. 3
- [25] OpenAI. Introducing chatgpt, 2023. 1
- [26] Neehar Peri, Achal Dave, Deva Ramanan, and Shu Kong. Towards long-tailed 3d detection. In *Conference on Robot Learning*, pages 1904–1915. PMLR, 2023. 8
- [27] 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 *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 2, 3, 4, 5
- [28] 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 *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIII 16*, pages 414–430. Springer, 2020. 2
- [29] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. 2
- [30] Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*, 2023. 3
- [31] Wayve. Lingo-1: Exploring natural language for autonomous driving. <https://wayve.ai/thinking/lingo-natural-language-autonomous-driving/>, 2023. 3
- [32] Licheng Wen, Daocheng Fu, Xin Li, Xinyu Cai, Tao Ma, Pinlong Cai, Min Dou, Botian Shi, Liang He, and Yu Qiao. Dilu: A knowledge-driven approach to autonomous driving with large language models. *arXiv preprint arXiv:2309.16292*, 2023. 3
- [33] Shawn Xu, Lin Yang, Christopher Kelly, Marcin Sieniek, Timo Kohlberger, Martin Ma, Wei-Hung Weng, Attila Kiraly, Sahar Kazemzadeh, Zakkai Melamed, et al. Elixir: Towards a general purpose x-ray artificial intelligence system through alignment of large language models and radiology vision encoders. *arXiv preprint arXiv:2308.01317*, 2023. 2
- [34] 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. 3
- [35] Sherry Yang, Ofir Nachum, Yilun Du, Jason Wei, Pieter Abbeel, and Dale Schuurmans. Foundation models for decision making: Problems, methods, and opportunities. *arXiv preprint arXiv:2303.04129*, 2023. 3
- [36] Wenyuan Zeng, Wenjie Luo, Simon Suo, Abbas Sadat, Bin Yang, Sergio Casas, and Raquel Urtasun. End-to-end interpretable neural motion planner. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8660–8669, 2019. 7
- [37] 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. 7