KEPT: Knowledge-Enhanced Prediction of Trajectories from Consecutive Driving Frames with Vision-Language Models

Yujin Wang[a], Tianyi Wang[b], Quanfeng Liu[a], Wenxian Fan[a], Junfeng Jiao[c], Christian Claudel[b],Yunbing Yan[d], Bingzhao Gao[a,*], Jianqiang Wang[e], and Hong Chen[f]

[a] *School of Automotive Studies, Tongji University, Shanghai, 201804, China*

[b] *Department of Civil, Architectural, and Environmental Engineering, The University of Texas at Austin, Austin, Texas, 78712, USA*

[c] *School of Architecture, The University of Texas at Austin, Austin, Texas, 78712, USA*

[d] *School of Automotive and Traffic Engineering, Wuhan University of Science and Technology, Wuhan, 430081, China*

[e] *School of Vehicle and Mobility, Tsinghua University, Beijing, 100084, China*

[f] *College of Electronic and Information Engineering, Tongji University, Shanghai, 201804, China*

[*] *Corresponding author.*

E-mail address: gaobz@tongji.edu.cn

Abstract

Accurate short-horizon trajectory prediction is pivotal for safe and reliable autonomous driving, yet existing vision-language models (VLMs) often fail to effectively ground their reasoning in scene dynamics and domain knowledge. To address this challenge, this paper introduces KEPT, a knowledge-enhanced VLM framework that predicts ego trajectories directly from consecutive front-view driving frames. KEPT couples a temporal frequency–spatial fusion (TFSF) video encoder—trained via self-supervised learning with hard-negative mining—with a scalable k-means + HNSW retrieval stack that supplies scene-aligned exemplars. Retrieved priors are embedded into chain-of-thought (CoT) prompts with explicit planning constraints, while a triple-stage fine-tuning schedule incrementally aligns the language head to metric

spatial cues, physically feasible motion, and temporally conditioned front-view planning. Evaluated on nuScenes dataset, KEPT achieves state-of-the-art performance across open-loop protocols: under NoAvg, it achieves 0.70m average L2 with a 0.21% collision rate; under TemAvg with lightweight ego status, it attains 0.31m average L2 and a 0.07% collision rate. Ablation studies show that all three fine-tuning stages contribute complementary benefits, and that using Top-2 retrieved exemplars yields the best accuracy–safety trade-off. The k-means–clustered HNSW index delivers sub-millisecond retrieval latency, supporting practical deployment. These results indicate that retrieval-augmented, CoT-guided VLMs offer a promising, data-efficient pathway toward interpretable and trustworthy autonomous driving.

Keywords

Autonomous Driving · Trajectory Prediction · Vision-Language Model · Retrieval-Augmented Generation · Chain-of-Thought Prompt

Nomenclature

Symbol	**Meaning (units)**
Sets, indices, and dimensions	
H, W, P	Image height, width, patch size (pixels)
T	Number of frames per clip (here $T = 7$)
N_p	Number of patches per frame ($N_p = \frac{H}{P}\frac{W}{P}$)
$\mathbb{R}, \mathbb{C}$	Real and complex domains
Images, patches, and transforms	
$I_{RGB} \in \mathbb{R}^{3\times H\times W}$	Input RGB image (normalized intensity)
$I_g \in \mathbb{R}^{1\times H\times W}$	Grayscale image
$X_{\text{patch}}^{(i)} \in \mathbb{R}^{P\times P}$	i-th grayscale patch
$\mathscr{F}\{\cdot\}$	2D FFT operator
$F^{(i)} = \mathscr{F}\{X_{\text{patch}}^{(i)}\} \in \mathbb{C}^{P\times P}$	Patch spectrum
$A^{(i)} = \lvert F^{(i)}\rvert \in \mathbb{R}^{P\times P}$	Amplitude spectrum
GAP$(\cdot)$	Global average pooling

Symbol	**Meaning (units)**
$a^{(i)} = \mathrm{GAP}(A^{(i)}) \in \mathbb{R}$	Amplitude statistic of patch i
$\mathbf{w}_{\text{freq}} \in [0,1]^{N_p}$	Frequency-domain attention weights
$X_f^{(i)} = w_{\text{freq}}^{(i)} X_{\text{patch}}^{(i)}$	Re-weighted patch
$X_f \in \mathbb{R}^{1\times H\times W}$	Frequency-attended image
$X_{\text{spa}}^{(\ell)}$	Swin-Tiny stage-ℓ spatial feature map
X_s	Upsampled / Concatenated spatial representation
$X_{fs} = [X_f; X_s]$	Channel concatenation of X_f and X_s
$X_{\text{feat}} \in \mathbb{R}^{C'\times H\times W}$	Fused representation after 1×1 conv, BN, ReLU
Temporal encoding	
$\mathbf{h}_t = \mathrm{GAP}(X_{\text{feat},t}) \in \mathbb{R}^{C'}$	Frame-t feature
$\mathbf{H}_{\text{seq}} = [\mathbf{h}_1; \ldots; \mathbf{h}_T] \in \mathbb{R}^{T\times C'}$	Temporal feature matrix
$\mathbf{H}_{\text{in}}, \mathbf{H}_{\text{out}}$	Transformer input / output sequences
$h_{\text{TFSF}} \in \mathbb{R}^{C'}$	CLS-token feature used as final vector
Self-supervised TFSF training (InfoNCE)	
$\mathscr{E}_\theta$	TFSF backbone with parameters θ
$\mathscr{G}_\phi$	Projection head with parameters ϕ
$\mathbf{z} = \dfrac{\mathscr{G}_\phi(\mathbf{h})}{\lVert\mathscr{G}_\phi(\mathbf{h})\rVert_2} \in \mathbb{R}^m$	ℓ_2-normalized projection on unit hypersphere
$\mathrm{sim}(\mathbf{u},\mathbf{v}) = \mathbf{u}^\top \mathbf{v}$	Inner product (cosine similarity numerator)
τ_{temp}	Temperature for InfoNCE
$\mathscr{N}_i^B$, $\mathscr{Q}$	In-batch negatives; memory queue (capacity M)
$\mathrm{TopK}(\cdot, K)$	Hard-negative mining (keep top-K)
$\mathscr{L}_i$	InfoNCE loss of sample i; $L = (1/B)\sum_i \mathscr{L}_i$
B, E, M, K	Batch size; #epochs; queue capacity; #hard negatives
Embedding & retrieval	
$\mathbf{e}_i \in \mathbb{R}^{128}$	Normalized embedding of clip i
n_c, $\mathbf{c}_j$	#Clusters; k-means centroid of cluster j
$d_{\cos}(\mathbf{x},\mathbf{y}) = 1 - \dfrac{\mathbf{x}^\top \mathbf{y}}{\lVert\mathbf{x}\rVert_2 \lVert\mathbf{y}\rVert_2}$	Cosine distance

Symbol	**Meaning (units)**
$\mathrm{HNSW}(M,\ \mathrm{ef}_{\mathrm{construction}})$	HNSW index (connectivity M, build ef)
k	#Nearest neighbors for retrieval
Triple-stage fine-tuning (Sec. 3.4)	
Stage A: Spatial perception (VQA multi-task)	
$\mathscr{L}_{\mathrm{A}}$	Stage-A multi-task loss (classification + size + distance)
K	Number of semantic categories for classification
$\eta_{k,t},\ \pi_{k,t}$	One-hot GT label and predicted probability for class k in sample t
$s_{m,t},\ \hat{s}_{m,t}$	Predicted / Reference object size along dimension $m \in \{1,2,3\}$
$d_t,\ \hat{d}_t$	Predicted / reference distance to ego vehicle
$\alpha_t,\ \beta_t,\ \gamma_t$	Binary indicators activating cls. / size / dist. terms for sample t
Stage B: Surround-view trajectory regression	
$\mathscr{L}_B$	Stage-B trajectory regression objective (coords + speed; with smoothness / curvature regularization)
Q	Number of prediction horizons (here $Q=3$, at 1s / 2s / 3s)
$q,\ r$	Horizon index $q \in \{1,\ldots,Q\}$; sample index $r \in \{1,\ldots,B\}$
$x_{q,r},\ y_{q,r},\ v_{q,r}$	Predicted ego-centric x/y displacements and speed at horizon q for sample r
$\tilde{x}_{q,r},\ \tilde{y}_{q,r},\ \tilde{v}_{q,r}$	Ground-truth counterparts of $x_{q,r}, y_{q,r}, v_{q,r}$
Stage C: Front-view end-to-end trajectory prediction	
$\mathscr{L}_C$	Stage-C loss, reusing Stage-B objective ($\mathscr{L}_C = \mathscr{L}_B$)
Kinematics and evaluation	
(x,y)	Ego-centric coordinates (m)
v	Speed (km/h)
$\hat{\tau} = \{(\hat{x}_i,\hat{y}_i)\}_{i=1}^{6},\ \tau = \{(x_i,y_i)\}_{i=1}^{6}$	Predicted / Ground-truth trajectory at 2Hz over 3s (m)

Symbol	Meaning (units)
$l_i = \sqrt{(x_i - \hat{x}_i)^2 + (y_i - \hat{y}_i)^2}$	Per-step Euclidean error (m)
$\bar{l}_{2,i}$	Mean l_2 at step i (m)
$L_{2,k}^{\text{UniAD}}$, $L_{2,k}^{\text{VAD}}$	L2 error metrics (UniAD at step $2k$; VAD cumulative to k) (m)
C_i, $\bar{C}_i$	Collision indicator (0 / 1) and its average
C_k^{UniAD}, C_k^{VAD}	Collision metrics (UniAD at $2k$; VAD cumulative)

1 Introduction

Accurate prediction of vehicle trajectories is a fundamental capability of autonomous driving systems, crucial for safe and efficient navigation in dynamic traffic environments Chai et al. (2020); Liu et al. (2021a); Ngiam et al. (2022). Recently, end-to-end autonomous driving methods have emerged as a promising paradigm by directly taking sensor data as input for perception and output planning results with one holistic model Chib and Singh (2023); Chen et al. (2024b). Through extensive data training, end-to-end approaches have demonstrated impressive planning capabilities, providing a streamlined and competitive alternative to traditional modular pipelines Casas et al. (2021); Chen and Krähenbühl (2022).

While end-to-end approaches have been widely adopted and constantly making breakthroughs on challenging benchmarks, they solely rely on fixed-format inputs, which restricts the agent's ability to comprehend multimodal information and interact with the environment and human users Shao et al. (2023a). Moreover, certain methods skip explicit scene understanding and directly predict driving commands from sensor data Hu et al. (2022b); Weng et al. (2024); Huang et al. (2024b), which sacrifices interpretability and introduces challenges in optimization and safety validation. Along this line of work, UniAD Hu et al. (2023) introduced a query-based design that integrates perception and prediction components, enabling an end-to-end planning scheme. VAD Jiang et al. (2023a) improved computational efficiency and interpretability by adopting a vectorized scene representation, replacing the dense rasterized representations from UniAD. Despite the remarkable progress of end-to-end models in autonomous driving, such approaches inherently struggle in long-tail scenarios, where prediction errors compound,

severely degrading downstream planning performance Chitta et al. (2021); Wu et al. (2022).

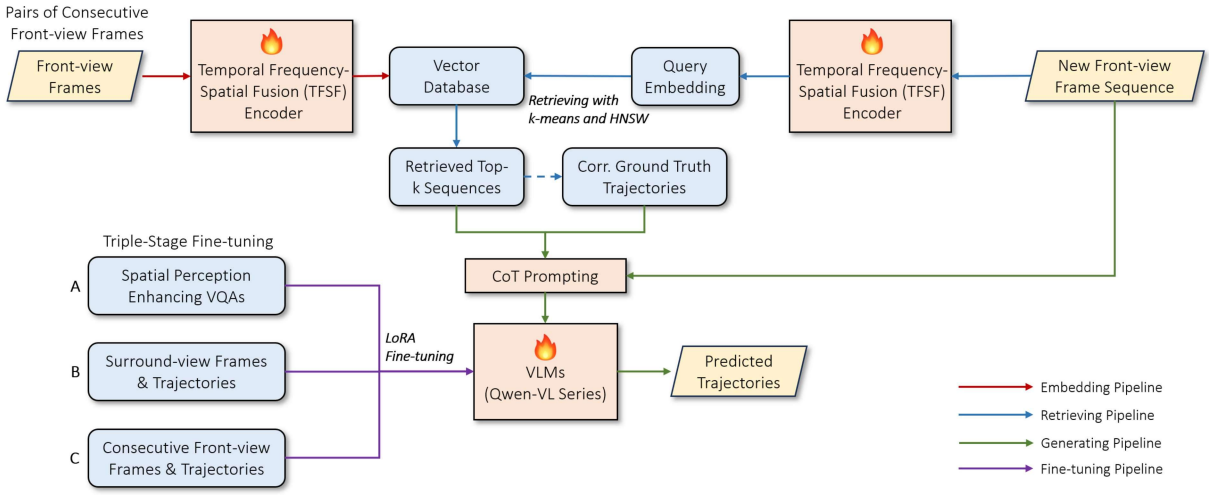

Fig. 1 The overview of our KEPT (Knowledge-Enhanced Prediction of Trajectories) method. The method consists of four pipelines: **(1) The embedding pipeline** encodes front-view frame sequences into a vector database via a self-designed and trained temporal frequency-spatial fusion (TFSF) encoder. **(2) The retrieving pipeline** retrieves the most similar scene from the database via k-means clustering and HNSW searching. **(3) The generating pipeline** guides vision language models (VLMs) in predicting trajectories in the future 3 seconds according to the new frame sequence and retrieved information via a chain-of-thought (CoT) prompting strategy. **(4) The fine-tuning pipeline** involves a triple-stage fine-tuning paradigm, which aims to equip VLMs with spatial grounding, motion feasibility, and temporally conditioned planning capabilities, enabling trajectory prediction from raw perceptual inputs.

Concurrently, large language models (LLMs) have demonstrated impressive capabilities in language comprehension and reasoning, showing potential to solve this problem Wei et al. (2022); Hegde et al. (2025). Going beyond text-based prompting, multimodal large language models (MLLMs) integrate image and video inputs to the LLM, enabling tasks such as visual question answer (VQA) and dense captioning Li et al. (2023a); Liu et al. (2023). Nevertheless, existing MLLM datasets for autonomous driving remain limited in scale, typically containing fewer than one million VQAs Wu et al. (2024). Consequently, without careful model and training scheme designs, these MLLMs exhibit poor performance in reasoning and planning tasks due to lacking scene understanding and grounding capability Ma et al. (2024b).

In this work, we propose **KEPT**, a knowledge-enhanced trajectory prediction framework, introducing a novel approach to assist an MLLM in predicting future trajectories according to consecutive driving

frames, as depicted in Figure 1. The main research contributions of this work are outlined as follows:

- **Temporal Frequency–Spatial Fusion (TFSF) Encoder for Consecutive Driving Frames:** We design a video encoder that fuses fast fourier transform (FFT)-based frequency attention with multi-scale Swin-Tiny spatial features and a temporal transformer over 7 frames (2Hz), yielding robust clip embeddings that capture dynamics salient for planning.
- **Self-Supervised Representation Learning with Hard-Negative Mining and Memory Queue:** We train TFSF via a contrastive InfoNCE objective using in-batch negatives, a fixed-capacity queue, and Top-K hard-negative selection, producing normalized embeddings on the unit hypersphere for downstream retrieval.
- **Scalable Embedding–Retrieval Pipeline for Trajectory Priors:** We construct a vector database of scene embeddings and perform two-stage approximate nearest-neighbor search—k-means clustering followed by HNSW under cosine distance-to retrieve top-K similar scenarios and their ground-truth trajectories as informative priors.
- **Triple-Stage Fine-Tuning of Vision-Language Models (VLMs) Bridging Perception to Planning:** (i) We enhance spatial grounding via VQA-style supervision on nuScenes (entity class, size, distance) with LoRA; (ii) We supervise surround-view trajectory regression from six cameras plus compact kinematics, enforcing smooth, collision-averse motion; (iii) We finalize end-to-end prediction from seven front-view frames using the same regression objective, focusing the language head on temporal reasoning.
- **Chain-of-Thought (CoT) Prompting with Exemplar Comparison and Planning Constraints:** We provide the VLMs with multiple retrieved reference scenes (images, kinematics, waypoint, and future trajectory) before the target scene, and explicitly instruct comparative reasoning under constraints (collision avoidance, speed smoothness, drivable-area adherence) to elicit transparent, controllable predictions.

The remainder of this paper is organized as follows: In Section 2, the literature review is conducted. In Section 3, details of the proposed KEPT method are introduced. In Section 4, comparative experiments and ablation studies are designed. Section 5 summarizes the work and discusses future research directions.

2 Related Work

2.1 Large Language Models in Autonomous Driving

Recent advancements in LLMs have demonstrated their ability to reason over complex contextual information, interpret human intent, and generate logically consistent outputs, sparking growing interest in their application to autonomous driving systems Wei et al. (2022); Mao et al. (2023); Ma et al. (2024c); Yang et al. (2025a). These works investigated the capabilities of LLMs to generalize to novel scenarios as well as their ability to reason about the scene in the form of text. "Drive as you speak" Cui et al. (2024) enriched LLMs with comprehensive environmental data from different vehicle modules, supporting safer and more context-aware driving decisions. "Driving with LLMs" Chen et al. (2024a) introduced LLMs that generated 10,000 driving scenarios for agent training. "Drive like a human" Fu et al. (2024) demonstrated LLMs' capabilities of understanding and interacting with environments in closed-loop systems, effectively navigating long-tail autonomous driving scenarios. Dilu Wen et al. (2024) first leveraged knowledge-driven capability in decision-making for autonomous vehicles.

Meanwhile, recent works have investigated the incorporation of MLLMs into end-to-end frameworks Wang et al. (2024); Xing et al. (2025); Brandstätter et al. (2025). Autonomous driving systems integrating multi-source information-such as vision, language, and rules-tend to exhibit greater generalization and safety, particularly in complex or long-tail scenarios Chen et al. (2025); Wang et al. (2025b); Ma et al. (2025). VLP Pan et al. (2024) incorporated linguistic descriptions into the training process and aligned them with visual features, significantly improving cross-city and cross-scenario generalization. DriveGPT4 Xu et al. (2024) utilized VLMs to predict control commands and simultaneously generate textual justifications using an iterative VQA format. To enhance embodied intelligence, LMDrive Shao et al. (2024) integrated a vision encoder with LLMs, enabling multimodal scene understanding and natural language command execution. Omnidrive Wang et al. (2025a) introduced a 3D VLM architecture to strengthen planning and reasoning capabilities. DriveVLM Tian et al. (2025) incorporated traditional 3D perception into a multistage reasoning chain, combining scene description, dynamic analysis, and hierarchical planning to bridge cognitive depth with real-time control. DiMA Hegde et al. (2025) employed a knowledge distillation approach to transfer planning capabilities from MLLMs to lightweight

vision-based planners via surrogate tasks.

2.2 End-to-End Trajectory Prediction

Early end-to-end trajectory prediction models primarily used LiDAR point clouds Casas et al. (2020); Liang et al. (2020); Wu et al. (2020). However, due to the dependence on accurate bounding box detection, these approaches often failed to generalize in the presence of occluded or unclassified objects. To solve this problem, recent studies have shifted toward vision-centric methods based on bird's-eye view (BEV) representations, which provide a unified spatial-temporal understanding of the driving scene Hu et al. (2021a); Akan and Güney (2022); Liang et al. (2022). For instance, PowerBEV Li et al. (2023b) realized semantic and instance-level trajectory prediction by simply forecasting segmentation and centripetal backward flow. TPV Huang et al. (2023) extended BEV with two additional perpendicular planes and utilized attention mechanism to aggregate the image features corresponding to each query in each TPV plane. LiDAR-camera fusion approaches have also emerged. For example, DAL Huang et al. (2024a) contained an attention-free predicting pipeline and an easy training process to achieve trade-off between speed and accuracy of 3D object detection and prediction through efficient multimodal fusion. BEVFormer Li et al. (2025) leveraged a unified multi-modality spatial attention mechanism for integrating camera and LiDAR data with historical BEV features. Despite these advances, BEV-based methods still face limitations in capturing fine-grained 3D structures and object-level semantics. To address these shortcomings, vision-centric 4D perception emerges as a promising alternative to effectively extend temporal occupancy prediction with camera images as input but also broaden both semantic and instance prediction beyond BEV format and predefined categories Tian et al. (2023); Tong et al. (2023); Wang et al. (2023); Wei et al. (2023). Cam4DOcc Ma et al. (2024a) was the first to achieve vision-centric 4D occupancy forecasting task, allowing simultaneous prediction of future trajectory for both general movable and static objects. Building upon this, Drive-OccWorld Yang et al. (2025b) investigated potential applications of the camera-based 4D occupancy forecasting world model by injecting action conditions and integrating this generative capability with end-to-end planners for safe driving.

2.3 Retrieval-Augmented Generation in Vision-Language Models

In vision-language tasks, retrieval-augmented generation (RAG) mitigates knowledge limitations by leveraging external knowledge bases, enabling models to extract insights from images while supplementing them with retrieved contextual data Shao et al. (2023b); Wang et al. (2025b). Prior studies Jiang et al. (2023b); Ram et al. (2023) on VLMs have shown that retrieval can improve contextual integration and multi-step reasoning under knowledge deficits, strengthen performance on complex question answering when paired with strong pre-trained encoders, and yield better downstream results when large external corpora are incorporated during pre-training and fine-tuning. Follow-up works Zheng et al. (2024); Hussien et al. (2025) further highlighted RAG's advantages for adaptive, multimodal generation in data-scarce domains and its ability to tighten cross-modal associations between images and text for more faithful grounding.

For driving-specific applications, dynamic retrieval policies that adjust what to fetch at run time have been proposed to match task requirements and latency budgets Bandyopadhyay et al. (2025). Domain knowledge can also be queried explicitly: a traffic-regulation agent retrieves applicable rules and guideline snippets conditioned on ego state and scene context to inform compliance-aware planning Gao et al. (2025). RAG offers a principled way to compensate for the limited world knowledge and long-tail corner cases that arise in autonomous driving Atakishiyev et al. (2024). For instance, RAG-Driver Yuan et al. (2024) introduced a retrieval-augmented in-context learning mechanism through a curated multimodal driving in-context instruction tuning dataset and a vector similarity-based retrieval engine specifically tailored for driving applications. VistaRAG Dai et al. (2024) positioned RAG as a safety-and-trustworthiness layer that dynamically retrieves prior driving experience, live road-network state, and other contextual evidence to ground decisions and make the reasoning auditable. SenseRAG Luo et al. (2025) constructed an LLM-readable environmental knowledge base from multimodal sensor/V2X streams and issued proactive queries to retrieve time-critical facts, yielding measurable gains in perception and prediction under latency constraints.

In summary, recent advances show that VLMs and RAG substantially enhance grounding, contextual reasoning, and data efficiency for autonomous driving; however, most pipelines (i) retrieve generic

knowledge rather than planning-specific cues, (ii) model temporal dynamics weakly for long-horizon forecasts, and (iii) offer limited safety calibration and auditability when reasoning under distribution shift. Moreover, training is often misaligned with metric, physics-aware supervision, leading to brittle behavior in crowded, rule-constrained scenes. To close these gaps, we introduce KEPT, a retrieval-conditioned VLM planner that couples a temporally aware perception encoder with scene-aligned, CoT-style retrieval and a three-stage fine-tuning recipe that grounds language reasoning in metric trajectories and collision risks. KEPT is designed to turn externally retrieved priors into actionable, long-horizon plans, improving both accuracy and safety while remaining lightweight and interpretable.

3 Methodology

3.1 Temporal Frequency-Spatial Fusion Encoder for Consecutive Driving Frames

In this subsection, we introduce a novel TFSF encoder for consecutive driving frames. The TFSF encoder integrates frequency-domain and spatial-domain representations at multiple granularities and employs a temporal transformer for encoding temporal dependencies. The overall architecture is illustrated in Figure 2.

For an input RGB image $I_{RGB} \in \mathbb{R}^{3\times H\times W}$, we first convert it into a grayscale representation $I_g \in \mathbb{R}^{1\times H\times W}$ by standard luminance-based weighting:

$$I_g = 0.299 I_R + 0.587 I_G + 0.114 I_B. \tag{1}$$

We partition the grayscale image into non-overlapping patches of size $P \times P$, generating a tensor $X_{patch} \in \mathbb{R}^{N_p\times P\times P}$, where $N_p = \frac{H}{P} \times \frac{W}{P}$. Each patch $X_{patch}^{(i)}$ is independently transformed into the frequency domain via a 2D FFT:

$$F^{(i)} = \mathscr{F}\{X_{\text{patch}}^{(i)}\}, \quad F^{(i)} \in \mathbb{C}^{P\times P}. \tag{2}$$

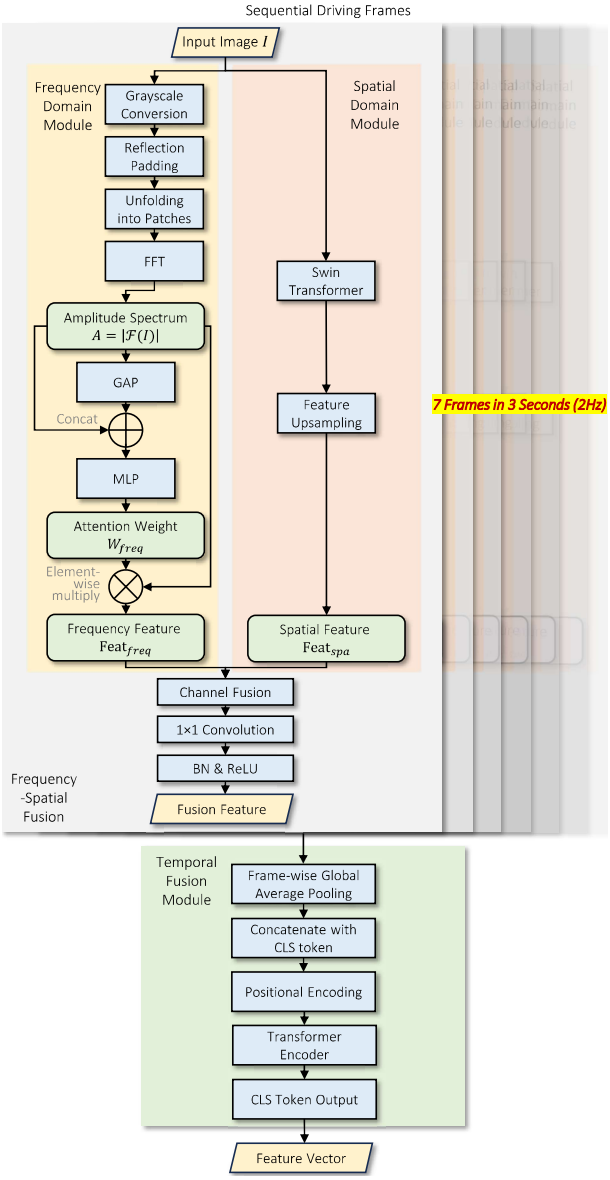

Fig. 2 The architecture of our proposed TFSF encoder for consecutive driving frames.

We subsequently compute the amplitude spectrum $A^{(i)}$:

$$A^{(i)} = |F^{(i)}|, \quad A^{(i)} \in \mathbb{R}^{P \times P}. \tag{3}$$

Then, a global average pooling (GAP) is applied to $A^{(i)}$:

$$a^{(i)} = \text{GAP}(A^{(i)}) = \frac{1}{P^2} \sum_{u=1}^{P} \sum_{v=1}^{P} A_{u,v}^{(i)}, \quad a^{(i)} \in \mathbb{R}. \tag{4}$$

These amplitude statistics across all patches form a vector $a \in \mathbb{R}^{N_p}$. A multilayer perceptron (MLP),

parameterized by weights W_f and biases b_f, predicts frequency-domain attention weights $w_f \in \mathbb{R}^{N_p}$:

$$w_f = \sigma(\text{MLP}(a; W_f, b_f)), \quad w_f \in [0,1]^{N_p}, \tag{5}$$

where σ denotes the sigmoid activation function. The patches are then re-weighted in the spatial domain by multiplying each original grayscale patch $X^{(i)}_{patch}$ with the corresponding scalar attention weight $w^{(i)}_f$:

$$X^{(i)}_f = w^{(i)}_f \cdot X^{(i)}_{\text{patch}}, \quad X^{(i)}_f \in \mathbb{R}^{P \times P}. \tag{6}$$

Finally, we reconstruct the frequency-attended image $X_f \in \mathbb{R}^{1 \times H \times W}$ by folding patches back to their original spatial positions.

To effectively capture hierarchical spatial representations, we employ a pre-trained hierarchical Swin Transformer Liu et al. (2021b) (Swin-Tiny) backbone. Given the original RGB image I_{RGB}, we feed it into the Swin Transformer:

$$X^{(\ell)}_{\text{spa}} = \text{SwinTransformer}_\ell(I_{RGB}), \quad \ell = 1,2,3,4, \tag{7}$$

where each stage ℓ yields feature maps at progressively coarser scales, thus providing multi-scale spatial representations. To align these multi-scale features with the original spatial resolution, we apply bilinear interpolation-based upsampling, resulting in the refined spatial representation $S \in \mathbb{R}^{C \times H \times W}$:

$$X_s = \text{Upsample}\left(\text{Concat}\left(X^{(1)}_{\text{spa}}, X^{(2)}_{\text{spa}}, X^{(3)}_{\text{spa}}, X^{(4)}_{\text{spa}}\right)\right). \tag{8}$$

To integrate complementary information from frequency and spatial domains, we concatenate the frequency-attended image X_f with the spatial features X_s along the channel dimension, yielding the fused tensor $X_{fs} \in \mathbb{R}^{(C+1) \times H \times W}$:

$$X_{fs} = [X_f; X_s]. \tag{9}$$

Subsequently, a 1×1 convolution followed by batch normalization (BN) and ReLU activation generates the unified spatial-frequency representation X_{feat}:

$$X_{\text{feat}} = \text{ReLU}\big(\text{BN}\big(\text{Conv}_{1\times1}(X_{\text{fs}})\big)\big), \quad X_{\text{feat}} \in \mathbb{R}^{C' \times H \times W}. \tag{10}$$

Given a sequence of T (in this work, T is set to 7 according to the sampling frequency of 2Hz in 3 seconds) consecutive image frames, we encode temporal dependencies among frame-level representations. First, we perform GAP on each spatial-frequency feature map $X_{feat,t}$ (frame t), obtaining temporal feature vectors $\mathbf{h}_t \in \mathbb{R}^{C'}$:

$$\mathbf{h}_t = \mathrm{GAP}(X_{\mathrm{feat},t}), \quad t = 1,2,\ldots,T. \tag{11}$$

The resulting feature vectors across all frames form a temporal feature matrix $\mathbf{H}_{\mathrm{seq}} \in \mathbb{R}^{T\times C'}$:

$$\mathbf{H}_{\mathrm{seq}} = [\mathbf{h}_1;\mathbf{h}_2;\ldots;\mathbf{h}_T]. \tag{12}$$

We prepend a learnable classification token $\mathbf{H}_{\mathrm{cls}} \in \mathbb{R}^{1\times C'}$ and add positional embeddings $\mathbf{P}_{\mathrm{pos}} \in \mathbb{R}^{(T+1)\times C'}$:

$$\mathbf{H}_{\mathrm{in}} = [\mathbf{h}_{\mathrm{cls}};\mathbf{h}_1;\mathbf{h}_2;\ldots;\mathbf{h}_T] + \mathbf{P}_{\mathrm{pos}}. \tag{13}$$

This sequence is then processed through an L-layer Transformer encoder (multi-head self-attention (MHSA) and MLP blocks):

$$\mathbf{H}_{\mathrm{out}} = \mathrm{TransformerEncoder}^{(L)}(\mathbf{H}_{\mathrm{in}}), \quad \mathbf{H}_{\mathrm{out}} \in \mathbb{R}^{(T+1)\times C'}. \tag{14}$$

We take the output corresponding to the classification token as the final feature representation f_{TFSF}:

$$\mathbf{h}_{TFSF} = \mathbf{H}_{out}^{(0)}, \quad \mathbf{h}_{TFSF} \in \mathbb{R}^{C'}. \tag{15}$$

3.2 Self-Supervised Training of the TFSF Encoder

The TFSF encoder is trained through a label-free contrastive learning paradigm that synthesizes in-batch negatives, a fixed-capacity memory queue and targeted hard-negative mining, as shown in Algorithm 1.

Let $\mathscr{V} = \{V_i\}_{i=1}^{N}$ denote an unlabeled video corpus where each clip $V_i = \{I_{i,t}\}_{t=1}^{T}$ comprises T RGB frames $I_{i,t} \in \mathbb{R}^{3\times H\times W}$. Two independently drawn stochastic augmentations $\mathscr{A}_1$ and $\mathscr{A}_2$ (random-resized crop, horizontal flip, color jitter, Gaussian blur) yield a positive pair:

$$\tilde{V}_i^{(1)} = \mathscr{A}_1(V_i), \quad \tilde{V}_i^{(2)} = \mathscr{A}_2(V_i). \tag{16}$$

Algorithm 1 Self-supervised training of the TFSF encoder

Inputs:

$\mathscr{V}$ — unlabeled video corpus

$\mathscr{E}_\theta, \mathscr{G}_\phi$ — TFSF backbone and projection head

E, B — #epochs, mini-batch size

τ_{temp} — temperature

K — #hard negatives

M_Q — queue capacity

Outputs:

θ^*, ϕ^* — trained parameters

procedure TRAIN($\mathscr{V}, \mathscr{E}_\theta, \mathscr{G}_\phi, E, B, \tau_{temp}, K, M_Q$)
  initialise θ, ϕ ▷ AdamW defaults
  $Q \leftarrow \emptyset$ ▷ FIFO memory queue
  for epoch $\leftarrow 1$ **to** E **do**
    for mini-batch $\mathbb{B} \leftarrow$ SAMPLE($\mathscr{V}, B$) **do**
      for each clip $V_i \in \mathbb{B}$ **do**
        $V_i^1 \leftarrow \mathscr{A}_1(V_i)$; $V_i^2 \leftarrow \mathscr{A}_2(V_i)$
        $z_i^1 \leftarrow \ell_2\_\mathtt{norm}(\mathscr{G}_\phi(\mathscr{E}_\theta(V_i^1)))$
        $z_i^2 \leftarrow \ell_2\_\mathtt{norm}(\mathscr{G}_\phi(\mathscr{E}_\theta(V_i^2)))$
      end for
      $N \leftarrow \{z_j^k \mid j \neq i \text{ or } k = 2\}_{\forall i} \cup Q$
      for $i \leftarrow 1$ **to** B **do**
        $H_i \leftarrow \text{TOP-K}_{\mathtt{roll}}(\text{sim}(z_i^1, \cdot), N, K)$ ▷ K most similar
        $s_i^+ \leftarrow \text{sim}(z_i^1, z_i^2)/\tau_{temp}$
        $\Delta_i \leftarrow \sum_{z \in H_i} \exp(\text{sim}(z_i^1, z)/\tau_{temp})$
        $\mathscr{L}_i \leftarrow -\log\left(\frac{\exp(s_i^+)}{\exp(s_i^+)+\Delta_i}\right)$
      end for
      $\mathscr{L} \leftarrow (1/B)\sum_i \mathscr{L}_i$
      $(\theta, \phi) \leftarrow \text{ADAMW}((\theta, \phi), \nabla\mathscr{L})$
      $Q \leftarrow \text{ENQUEUE}(Q, \{z_i^2\})$; TRIM($Q, M_Q$)
    end for
  end for
  return θ^*, ϕ^*
end procedure

Notation: $\text{sim}(u, v) = u^\top v$; $\ell_2\_\mathtt{norm}(x) = x/\|x\|_2$; TOP-K$_{\mathtt{roll}}$ returns the K largest inner-products.

The TFSF encoder $f_\theta : \mathbb{R}^{T\times 3\times H\times W} \to \mathbb{R}^d$ maps a clip to a d-dimensional representation $\mathbf{h} = f_\theta(\tilde{V})$. A three-layer projection head $\mathscr{G}_\phi : \mathbb{R}^d \to \mathbb{R}^m$ projects $\mathbf{h}$ onto the unit hypersphere:

$$\mathbf{z} = \frac{\mathscr{G}_\phi(\mathbf{h})}{\|\mathscr{G}_\phi(\mathbf{h})\|_2}, \quad \mathbf{z} \in \mathbb{R}^m, \quad d = 64, m = 128. \tag{17}$$

For mini-batch index i, we denote the anchor as $\mathbf{z}_i^{(1)}$ and its augmented counterpart as $\mathbf{z}_i^{(2)}$. Their similarity (positive logits) is calculated as follows:

$$s_i^+ = \frac{\mathbf{z}_i^{(1)\top}\mathbf{z}_i^{(2)}}{\tau_{temp}}, \quad \tau_{temp} > 0, \tag{18}$$

where two negative pools are employed, namely the in-batch negatives $\mathscr{N}_i^{\mathrm{B}} = \left\{\mathbf{z}_j^{(k)} \mid j \neq i \text{ or } k = 2\right\}$ and the memory-queue negatives $\mathscr{Q} = \{\mathbf{q}_\ell\}_{\ell=1}^M$ with fixed size $M = 1024$, storing projections from recent mini-batches.

Since most negatives are easily separable, we retain only the K most similar hard negatives:

$$\mathscr{N}_i^* = \mathrm{TopK}\left(\mathscr{N}_i^{\mathrm{B}} \cup \mathscr{Q}, \mathbf{z}_i^{(1)\top}\mathbf{z}\right), \quad K = 10. \tag{19}$$

The InfoNCE loss could be computed as the following formula:

$$\mathscr{L}_i = -\log \frac{\exp(s_i^+)}{\exp(s_i^+) + \sum_{\mathbf{z}\in\mathscr{N}_i^*} \exp\left(\frac{\mathbf{z}_i^{(1)\top}\mathbf{z}}{\tau_{temp}}\right)}, \tag{20}$$

and the batch loss could be defined as $\mathscr{L} = B^{-1}\sum_{i=1}^B \mathscr{L}_i$.

After each iteration the vectors $\left\{\mathbf{z}_i^{(2)}\right\}_{i=1}^B$ are en-queued and the oldest B items are dequeued, keeping $|\mathscr{Q}| = M$. Both the backbone parameters θ and the projection parameters ϕ receive direct gradient updates:

$$\theta \leftarrow \theta - \eta_e \nabla_\theta \mathscr{L}, \quad \phi \leftarrow \phi - \eta_p \nabla_\phi \mathscr{L}, \tag{21}$$

with learning rates $\eta_e = 1\times 10^{-5}$, $\eta_p = 1\times 10^{-4}$ and weight decay 1×10^{-4}. Training is conducted for $E = 50$ epochs on 7-frame clips, batch size $B = 8$ and temperature $\tau_{temp} = 0.07$ on a single 3090 GPU. The model with minimum validation loss is preserved for downstream embedding and retrieval tasks.

3.3 The Embedding and Retrieval Pipeline

The overview of the embedding and retrieval pipeline is shown in Algorithm 2. Given a collection of unlabeled video sequences, a database embedding set could be constructed firstly using the trained TFSF encoder coupled with a projection head. Specifically, for each video clip $V_i = \{I_{i,t}\}_{t=1}^{T}$, containing $T = 7$ consecutive frames, we perform deterministic image preprocessing, including resizing to 224×224, normalization by ImageNet mean and standard deviation, and tensor stacking. These preprocessed frames form an input tensor $X_i \in \mathbb{R}^{1\times 7\times 3\times 224\times 224}$.

Each video sequence embedding is computed as:

$$\mathbf{h}_i = \mathscr{E}_\theta(X_i), \quad \mathbf{e}_i = \frac{\mathscr{G}_\phi(\mathbf{h}_i)}{\|\mathscr{G}_\phi(\mathbf{h}_i)\|_2} \in \mathbb{R}^{128}, \tag{22}$$

where $\mathscr{E}_\theta$ denotes the trained TFSF encoder and $\mathscr{G}_\phi$ represents a multilayer projection head consisting of linear layers, batch normalization, and ReLU activations. This produces an embedding $\mathbf{e}_i$ that resides on the 128-dimensional unit hypersphere. Embeddings for all sequences are subsequently stored in a database along with unique identifiers for retrieval tasks.

To facilitate efficient retrieval, we adopt a two-stage indexing strategy combining unsupervised clustering with hierarchical navigable small-world graphs (HNSW) Malkov and Yashunin (2018).

Given the database embeddings $\mathscr{E}_{\text{db}} = \{\mathbf{e}_i\}_{i=1}^{N_{\text{db}}}$, we first partition them into n_c clusters via k-means clustering:

$$\arg\min_{\{\mathbf{c}_j\}} \sum_{i=1}^{N_{\text{db}}} \min_j \|\mathbf{e}_i - \mathbf{c}_j\|_2^2, \tag{23}$$

where $\{\mathbf{c}_j\}_{j=1}^{n_c}$ denote cluster centroids. Each embedding is thus associated with a cluster label.

Within each cluster, we construct a separate HNSW index designed for efficient approximate nearest-neighbor search under cosine distance:

$$d_{\cos}(\mathbf{x},\mathbf{y}) = 1 - \frac{\mathbf{x}^\top \mathbf{y}}{\|\mathbf{x}\|_2 \|\mathbf{y}\|_2}. \tag{24}$$

Specifically, for cluster j, an HNSW index is initialized and built using all embeddings within that cluster. Parameters including $M = 16$ and ef_construction $= 200$ are set to balance retrieval speed and accuracy.

Algorithm 2 Embedding and Retrieval Pipeline

Inputs:

$V_{\text{db}}, V_{\text{val}}$ — Database and validation video clips

$\mathscr{E}_\theta, \mathscr{G}_\phi$ — Trained TFSF encoder and projection head

nc, M — Number of clusters, HNSW connectivity

procedure EMBEDDING AND RETRIEVAL

  1. Embedding Generation

  for each clip $V_i \in V_{\text{db}}$ **do**

    $X_i \leftarrow \text{Preprocess}(V_i)$ ▷ deterministic transformations

    $h_i \leftarrow \mathscr{E}_\theta(X_i)$

    $e_i \leftarrow \text{Normalize}(\mathscr{G}_\phi(h_i))$

  end for

  Save embeddings $\{e_i\} \rightarrow$ database

  2. Clustered HNSW Indexing

  $\{c_j\}, \text{labels} \leftarrow \text{k-means}(\{e_i\}, \text{n\_clusters=nc})$

  for each cluster j **do**

    Build $\text{HNSW}(\{e_i \mid \text{labels}[i] = j\}, \text{connectivity} = M)$

  end for

  3. Retrieval

  for each validation clip $V_k \in V_{\text{val}}$ **do**

    $X'_k \leftarrow \text{Preprocess}(V_k)$

    $h'_k \leftarrow \mathscr{E}_\theta(X'_k)$, $e'_k \leftarrow \text{Normalize}(\mathscr{G}_\phi(h'_k))$

    $c_j \leftarrow \text{k-means.predict}(e'_k)$

    $\text{retrieved\_indices} \leftarrow \text{HNSW\_query}(e'_k, \text{cluster} = c_j, \text{topk} = k)$

    $\text{retrieved\_trajectories} \leftarrow \text{Map indices} \rightarrow \text{trajectories}$

    Output retrieved_trajectories

  end for

end procedure

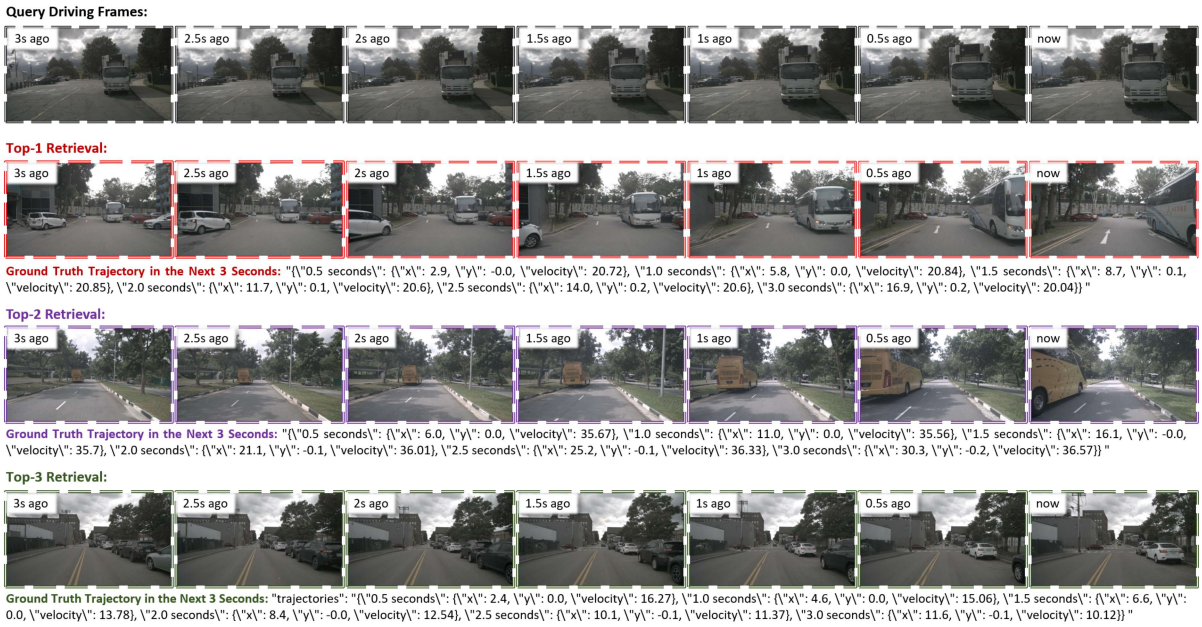

Fig. 3 Illustration of the embedding-based retrieval pipeline in action.

Given validation embeddings $\mathscr{E}_{\text{val}} = \{\mathbf{e}'_k\}_{k=1}^{N_{\text{val}}}$, retrieval is conducted by predicting cluster assignments via the pre-fitted k-means model, and querying the corresponding cluster-specific HNSW index to identify the k nearest database embeddings:

$$\{\mathbf{e}_l\}_{l=1}^{k} = \text{HNSW\_knn\_query}(\mathbf{e}'_k). \tag{25}$$

Each retrieved embedding index corresponds to an entry in the original database JSON file, from which associated ground truth trajectories are extracted. Formally, given validation embedding index k and retrieved indices $\{l_1, \ldots, l_k\}$, the output structure is:

$$R_k = \left\{\mathscr{T}(l_j)\right\}_{j=1}^{k}, \tag{26}$$

where $\mathscr{T}(l_j)$ denotes the trajectory annotation corresponding to database embedding index l_j. This is a pipeline designed for large-scale validation datasets. In practical applications, a new scene sequence can be independently embedded and used for one-to-one retrieval, retaining a single trajectory.

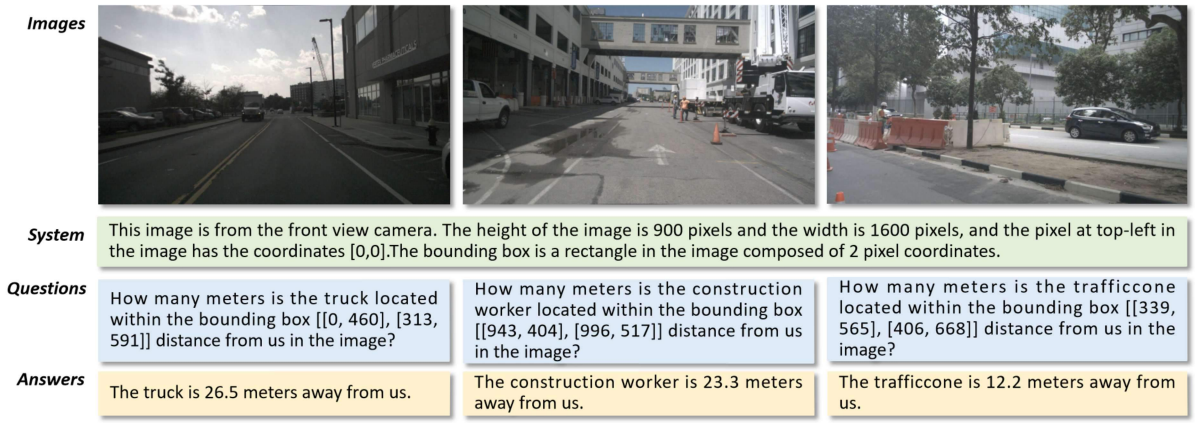

Fig. 4 Examples of distance inference tasks in the training dataset.

Representative retrieval examples are illustrated in Figure 3. The top row displays a sequence of query driving frames sampled at uniform intervals over the past 3 seconds, capturing a scenario involving a truck ahead. The subsequent three rows illustrate the top-3 retrieval results obtained by the proposed embedding and retrieval method. Each retrieval result consists of a visually similar sequence of frames, retrieved from the training database using the learned embeddings. The ground truth trajectories corresponding to each retrieved sequence are provided, including position coordinates (x, y) and velocities (in km/h) for future timestamps at intervals of 0.5 seconds, spanning the next 3 seconds. These retrieved cases demonstrate the ability of the embedding pipeline to effectively capture visual similarity and retrieve semantically relevant driving scenarios, thereby providing useful trajectory priors for downstream prediction tasks.

3.4 The Triple-Stage Fine-Tuning Strategy

Reliable trajectory prediction in autonomous driving hinges on VLMs' ability to accurately perceive spatial relationships in complex traffic scenes. Although VLMs have already demonstrated strong reasoning capabilities, their depth perception and scale estimation from monocular inputs often remain suboptimal, particularly for distant targets. Such deficiencies can cascade through the prediction pipeline, leading to unsafe motion proposals. Therefore, we adopt a triple-stage fine-tuning framework for VLMs, in which the first stage is dedicated to enhancing spatial perception from front-view frames, while the second stage directly optimizes trajectory prediction using paired camera frames and ground-truth trajectories.

3.4.1 Stage A: Towards Enhancement of Spatial Perception Capabilities of VLMs

In this stage, the visual encoder is frozen, and only the LLM component of the VLM is adapted to strengthen spatial grounding. We construct a spatial perception dataset from nuScenes in a VQA format, containing over 100,000 verified question–answer pairs that cover three perception tasks essential for planning: entity categorization, physical size estimation, and distance inference Wang et al. (2025b). Examples of distance inference tasks in the training dataset are illustrated in Figure 4. All labels are cross-checked against ground-truth annotations to ensure both uniqueness and precision.

The multi-task loss for this stage is formulated as:

$$\mathscr{L}_{\mathrm{A}} = \frac{1}{B}\sum_{t=1}^{B}\left[\alpha_t \cdot \left(-\sum_{k=1}^{K}\eta_{k,t}\log\pi_{k,t}\right) + \beta_t \cdot \frac{1}{3}\sum_{m=1}^{3}(s_{m,t}-\hat{s}_{m,t})^2 + \gamma_t \cdot (d_t-\hat{d}_t)^2\right], \tag{27}$$

where B denotes the batch size, K is the number of semantic categories, $\eta_{k,t}$ and $\pi_{k,t}$ represent the one-hot ground-truth label and predicted probability for category k in sample t, $s_{m,t}$ and $\hat{s}_{m,t}$ indicate the predicted and reference object size along dimension m, d_t and $\hat{d}_t$ refer to the predicted and actual distance to the ego vehicle, and α_t, β_t, γ_t are binary indicators specifying whether the classification, size regression, and distance regression terms are active for that sample.

We adopt Qwen-series VLMs as the backbone, and parameter-efficient adaptation is achieved using LoRA Hu et al. (2022a) over three epochs. This stage equips the model with refined metric spatial awareness, forming a robust perceptual foundation for the subsequent trajectory-supervised fine-tuning.

3.4.2 Stage B: Towards Scene-Conditioned Trajectory Prediction from Surround-View Images

In this stage, panoramic perception is coupled with short-horizon motion synthesis by supervising VLMs to output ego-centric trajectories directly from six synchronized cameras together with compact kinematic context. Each training instance supplies front, front-left, front-right, rear-left, rear-right, and rear images, along with the ego vehicle's instantaneous state in the vehicle frame, the historical trajectory over the preceding three seconds, and a guiding waypoint that indicates the intended direction without requiring arrival within the prediction window. VLMs are prompted in natural language to respect collision avoidance, velocity smoothness, and drivable-area constraints, and must produce a JSON triplet

of positions and speeds at 1, 2, and 3 seconds ahead, expressed in an ego-centric coordinate system (x forward, y left) with values serialized to one decimal place. Quantization is applied only at serialization so that learning remains fully differentiable.

Visual tokens from the six views are extracted by a frozen image encoder and fused with structured kinematics (current speed, the past three-second trajectory, and the waypoint) before entering the language head. Training minimizes a regression objective over coordinates and speed with lightweight regularization on temporal smoothness and curvature:

$$\mathscr{L}_{\mathrm{B}} = \frac{1}{B}\sum_{r=1}^{B}\frac{1}{Q}\sum_{q=1}^{Q}\left[(x_{q,r}-\tilde{x}_{q,r})^2+(y_{q,r}-\tilde{y}_{q,r})^2+(v_{q,r}-\tilde{v}_{-q,r})^2\right], \tag{28}$$

where B denotes the batch size, $Q=3$ denotes the prediction horizons (1s, 2s, 3s), $(x_{q,r}, y_{q,r}, v_{q,r})$ are the predicted ego-centric x/y coordinates and speed for sample r at horizon q, tildes indicate their ground-truth counterparts. The LoRA fine-tuning is still applied in this stage over three epochs.

This stage trains VLMs to transform surround-view evidence and recent motion history into physically realizable, short-term plans that remain collision-free, speed-smooth, and road-aligned. The resulting motion-aligned representation forms the bridge from Stage A's perception enhancement to Stage C's end-to-end trajectory prediction.

3.4.3 Stage C: Towards End-to-End Trajectory Prediction from Consecutive Front-View Frames

Stage C finalizes the triple-stage pipeline by training VLMs to produce short-horizon ego trajectories directly from a temporally ordered front-view frame sequence and compact kinematic context. Each training instance supplies seven front-view frames timestamped at $-3.0, -2.5, -2.0, -1.5, -1.0, -0.5, 0.0$ seconds relative to the present, a short history of ego positions and speeds expressed in the ego frame, and an optional soft waypoint that provides a directional prior without imposing arrival within the prediction horizon. The supervisory target is the same canonical JSON triplet used elsewhere in the pipeline, containing ego-centric (x, y) displacements and speed at the prescribed horizons.

Architecturally, the seven frames are tokenized by the frozen visual encoder and augmented with relative-time embeddings to preserve temporal order. These visual tokens are concatenated with structured

kinematic features (recent positions / speeds and the waypoint) and consumed by the language head. Consistent with earlier stages, we retain a Qwen-series VLM and adapt only the LLM parameters via LoRA, which keeps adaptation parameter-efficient while focusing the optimization on temporal reasoning from consecutive frames.

To avoid mathematical duplication, Stage C reuses the trajectory regression objective introduced in Stage B ($\mathscr{L}_{\mathrm{C}} = \mathscr{L}_{\mathrm{B}}$), which jointly penalizes pointwise coordinate and speed errors and contains the same smoothness and curvature regularizers. In practice, Stage C minimizes $\mathscr{L}_{\mathrm{C}}$ without introducing new loss terms, only modest task-specific tuning of the regularization weights is performed to accommodate the reduced field-of-view and the stronger reliance on temporal cues.

Stage C therefore completes this triple-stage fine-tuning paradigm by converting the perception capabilities from Stage A and the surround-view priors from Stage B into a front-view–centric temporal planner. Because the loss, decoding constraints and regularizers are shared with Stage B, Stage C integrates seamlessly into the unified optimization framework while contributing front-view–specific temporal reasoning essential for practical, deployable single-camera planning.

3.5 The Chain-of-Thought Prompting Paradigm

To enhance the reasoning ability of VLMs in trajectory prediction for autonomous driving with the help of historical reference and knowledge, we adopt a CoT prompting paradigm that guides the model through stepwise, exemplar-based reasoning. Unlike conventional prompting that directly requests a prediction based on the current scene, our approach leverages structured multimodal context and comparative reasoning to emulate human-like decision-making.

As shown in Figure 5, our CoT prompting is structured around a multi-scene framework, where the model is sequentially presented with multiple similar reference scenes (retrieved by the embedding and retrieval pipeline, and the number is determined by the top-K parameter during retrieval) followed by a target scene for inference. Each reference scene includes:

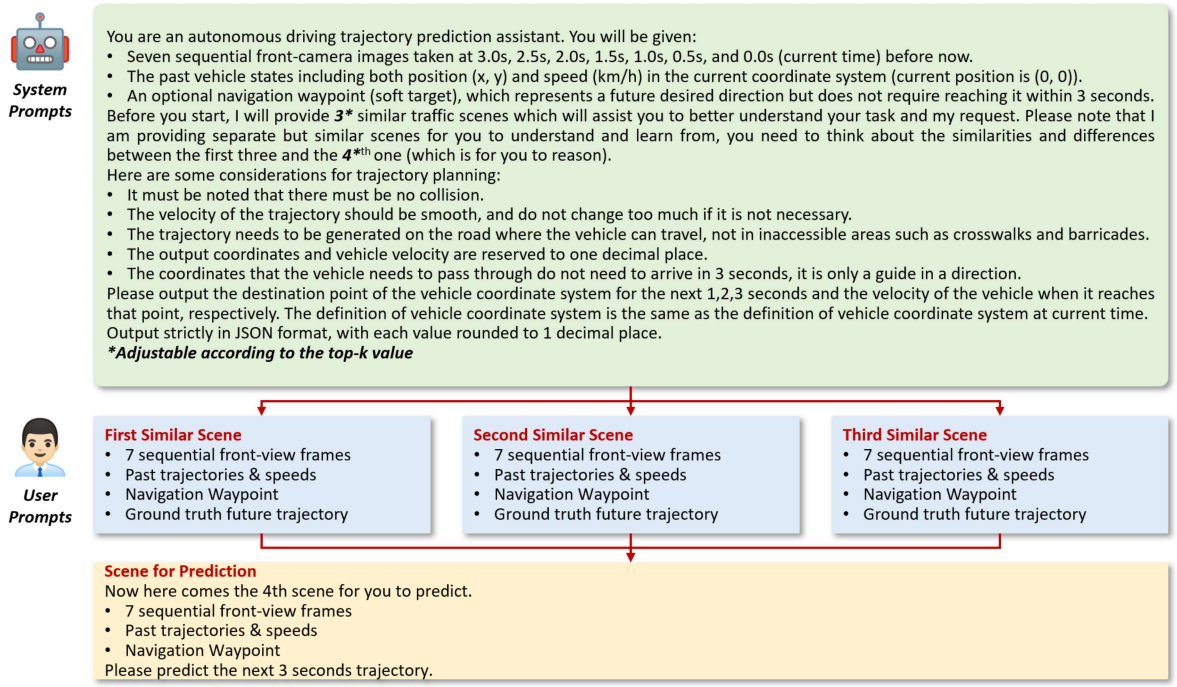

Fig. 5 An example of the CoT prompting paradigm.

- A sequence of seven front-view images captured at time steps {3.0, 2.5, 2.0, 1.5, 1.0, 0.5, 0.0} seconds before the present,
- historical trajectory data including position and velocity in the ego-vehicle coordinate system,
- an optional navigation waypoint representing a desired future direction,
- the ground truth future-trajectory of the ego vehicle.

This exemplar-guided context enables the model to understand the expected mapping between scene dynamics, vehicle motion, and feasible future trajectories. By observing multiple instances, the model is encouraged to identify patterns, discern scene-specific constraints, and generalize to novel but related driving situations. To ensure coherent and safe trajectory generation, the system prompt includes a set of explicit planning constraints, such as avoiding collisions, maintaining speed smoothness, and adhering to drivable regions. These constraints help enforce output controllability and facilitate seamless integration with downstream evaluation pipelines. A key feature of our CoT paradigm is the emphasis on relational reasoning. The model is explicitly instructed to analyze the similarities and differences between the reference scenes and the target scene. This enhances knowledge transfer across

scenes and fosters contextual adaptation. Furthermore, all spatial data are expressed in a consistent vehicle-centric coordinate system, and temporal cues are clearly indicated, enabling precise temporal-spatial grounding of the trajectory prediction task. Likewise, the structure of the prompt is designed to allow seamless integration of the ego vehicle's current dynamic state (e.g., velocity, acceleration, and yaw angle) during inference. These parameters can be directly input into the system without requiring additional preprocessing, enabling more accurate and context-aware decision making.

4 Experiment

4.1 Experimental Setup

4.1.1 Data Preparation

We divide the 34,000 scenes from the nuScenes dataset Caesar et al. (2020) into four subsets:

- The first subset consists of 10,000 scenes. After removing redundant frames that lack either the first 3 seconds of historical context or the last 3 seconds of future trajectory information, a total of 9,062 scene sequences remain. We use this subset to train the TFSF encoder.
- The second subset also contains 10,000 driving scenes. In Stage A, we use the surround-view images from this subset to build a dataset for fine-tuning VLMs' spatial perception abilities, enabling it to better interpret multi-camera inputs. In Stage B, we further fine-tune VLMs using both the surround-view images and corresponding 3-second future trajectory ground truths, enhancing its temporal reasoning and motion prediction performance. Following the same filtering strategy as in the first subset, front-view frames of these 10,000 scenes are divided into 1,117 scene sequences. These sequences, paired with their 3-second future trajectory ground truths, are utilized in Stage C, serving as the final stage to jointly optimize spatiotemporal reasoning and trajectory prediction.
- The third subset also includes 10,000 driving scenes. As with the first subset, we extract 9,062 temporally complete scene sequences. Each sequence is paired with its corresponding 3-second future trajectory and encoded into a joint vector representation with the TFSF encoder. These representations are stored in a vector-based knowledge base, which supports RAG by enabling

efficient similarity search over historical scenarios for trajectory prediction.

- The remaining 4,000 scenes are used to construct the validation set, from which 3,871 valid scene sequences are obtained. This subset is used to evaluate the overall performance of VLMs.

4.1.2 Evaluation Metrics

The prediction of trajectories is evaluated using the L2 distance from ground truth trajectories and the object collision rate. More specifically, we adopt the L2 and collision rate based on UniAD Hu et al. (2023) and ST-P3 Hu et al. (2022b). UniAD's **NoAvg** protocol evaluates each horizon using only the current-frame prediction at that exact horizon (no temporal aggregation), whereas ST-P3's **TemAvg** protocol reports the score at a horizon as the cumulative arithmetic mean of the per-sub-horizon scores from 0.5s up to that horizon, effectively smoothing performance over time. In the comparative experiments, to enable comparison with baselines that adopt different evaluation protocols, we report KEPT's best results under both protocols; in the ablation study, since no comparison with other methods is involved, we evaluate the results using the NoAvg protocol.

4.1.3 The Vision-Language Model Backbone

In our experiments, we adopt the Qwen2-VL-2B model as the default vision-language backbone. As previously noted, this component is modular and can be substituted with other models in the Qwen family, including Qwen2-VL-7B, Qwen2.5-VL-3B, Qwen2.5-VL-7B, or larger-scale variants. Our preference for the 2B model is based on its significantly reduced computational demands, both during training and inference, which makes it a practical option for real-time applications and potential deployment on resource-constrained edge platforms such as vehicle-mounted systems. All evaluation is conducted on four NVIDIA A800 GPUs.

4.2 Comparative Experiments

4.2.1 Results under NoAvg

As shown in Table 2, averaged over 1-3s, KEPT attains 0.70m L2 and 0.21% collision rate, improving over UniAD (1.03m / 0.31%) by 32.0% and 32.3%, respectively, and over VAD-Base (1.22 m / 0.43%) by

42.6% and 51.2%. Against the recent camera-only world-model Drive-OccWorld (0.85m / 0.29%), KEPT reduces average error by 17.6% and collision rate by 27.6%. At the longest horizon (3s), KEPT's 1.23m L2 is 17.5-37.9% lower than competing UniAD/VAD-Base/Drive-OccWorld variants, with a concomitant drop in 3-s collisions (0.44% vs. 0.64-0.96%).

4.2.2 Results under TemAvg without Ego Status

Under temporal averaging, KEPT continues to deliver the lowest L2 at all horizons, achieving 0.37m on average, which is 21.3% lower than Drive-OccWorld (0.47m). Collision rates are competitive at short horizons (1s: 0.03%, on par with the best baselines) but slightly higher at longer horizons (average: 0.16% vs. 0.11% for Drive-OccWorld), revealing a small safety-accuracy trade-off when ego priors are withheld.

4.2.3 Results under TemAvg with Ego Status

When lightweight ego signals are provided, matching many recent VLM-for-driving setups, KEPT achieves the best overall numbers: 0.31m average L2 and 0.07% average collisions. This improves over OmniDrive (0.33m / 0.30%) by 6.1% / 76.7%, over Drive-OccWorld (0.32m / 0.29%) by 3.1% / 75.9%, and over RDA-Driver (0.40m / 0.10%) by 22.5% / 30.0%. Notably, KEPT's 3-s L2 (0.51m) is on par with the best competitors (0.49-0.55m) while reducing average collisions by 2-4$\times$ relative to most baselines.

4.2.4 Discussion

The consistent L2 gains, especially at 2-3s, stem from (i) the TFSF encoder that preserves motion-relevant dynamics, (ii) retrieval-augmented CoT prompting that supplies scene-aligned priors, and (iii) triple-stage fine-tuning that grounds the VLM in metric spatial cues and physically feasible motion. When ego status is available, KEPT converts these advantages into simultaneous accuracy and safety improvements, indicating better long-horizon stability rather than early-horizon over-fitting. Overall, Table 2 evidences that KEPT is a strong, data-efficient VLM planner that advances various baselines across protocols while remaining lightweight enough for practical deployment.

Methods	L2 (m)↓				Collision (%)↓			
	1s	2s	3s	Avg.	1s	2s	3s	Avg.
FF Hu et al. (2021b)	0.55	1.20	2.54	1.43	0.06	0.17	1.07	0.43
EO Khurana et al. (2022)	0.67	1.36	2.78	1.60	0.04	0.09	0.88	0.33
ST-P3† Hu et al. (2022b)	1.72	3.26	4.86	3.28	0.44	1.08	3.01	1.51
OccNet† Ahmed et al. (2023)	1.29	2.13	2.99	2.14	0.21	0.59	1.37	0.72
UniAD† Hu et al. (2023)	0.48	0.96	1.65	1.03	**0.05**	0.17	0.71	0.31
VAD-Base† Jiang et al. (2023a)	0.54	1.15	1.98	1.22	0.10	0.24	0.96	0.43
Drive-OccWorld† Yang et al. (2025b)	0.32	0.75	1.49	0.85	**0.05**	0.17	0.64	0.29
KEPT† (ours, on Qwen2-VL-2B)	**0.23**	**0.63**	**1.23**	**0.70**	**0.05**	**0.14**	**0.44**	**0.21**
ST-P3‡ Hu et al. (2022b)	1.33	2.11	2.90	2.11	0.23	0.62	1.27	0.71
UniAD‡ Hu et al. (2023)	0.44	0.67	0.96	0.69	0.04	**0.08**	0.23	0.12
VAD-Base‡ Jiang et al. (2023a)	0.41	0.70	1.05	0.72	0.07	0.17	0.41	0.22
Drive-OccWorld‡ Yang et al. (2025b)	0.25	0.44	0.72	0.47	**0.03**	**0.08**	**0.22**	**0.11**
KEPT‡ (ours, on Qwen2-VL-2B)	**0.17**	**0.35**	**0.59**	**0.37**	**0.03**	0.10	0.34	0.16
GPT-Driver†,* Mao et al. (2023)	0.27	0.74	1.52	0.84	0.07	0.15	1.10	0.34
RDA-Driver†,* Huang et al. (2024b)	0.23	0.73	1.54	0.80	**0.00**	0.13	0.83	0.32
DME-Driver†,* Han et al. (2025)	0.43	0.91	1.58	0.97	0.04	0.14	0.64	0.27
VLP†,* Pan et al. (2024)	0.36	0.68	1.19	0.74	0.03	0.12	0.32	0.16
KEPT†,* (ours, on Qwen2-VL-2B)	**0.19**	**0.52**	**1.12**	**0.61**	0.02	**0.11**	**0.29**	**0.14**
UniAD‡,* Hu et al. (2023)	0.20	0.42	0.75	0.46	0.02	0.25	0.84	0.37
VAD-Base‡,* Jiang et al. (2023a)	0.17	0.34	0.60	0.37	0.04	0.27	0.67	0.33
GPT-Driver‡,* Mao et al. (2023)	0.20	0.40	0.70	0.44	0.04	0.12	0.36	0.17
RDA-Driver‡,* Huang et al. (2024b)	0.17	0.37	0.69	0.40	0.01	**0.05**	0.26	0.10
BEV-Planner‡,* Li et al. (2024)	0.16	0.32	0.57	0.35	**0.00**	0.29	0.73	0.34
VLP‡,* Pan et al. (2024)	0.30	0.53	0.84	0.55	0.01	0.07	0.38	0.15
DriveVLM‡,* Tian et al. (2025)	0.18	0.34	0.68	0.40	0.10	0.22	0.45	0.27
OmniDrive‡,* Wang et al. (2025a)	**0.14**	0.29	0.55	0.33	**0.00**	0.13	0.78	0.30
Drive-OccWorld‡,* Yang et al. (2025b)	0.17	0.31	**0.49**	0.32	0.02	0.24	0.62	0.29
KEPT‡,* (ours, on Qwen2-VL-2B)	**0.14**	**0.28**	0.51	**0.31**	0.01	0.06	**0.13**	**0.07**

Table 2 Performance comparison with representative published baselines and our KEPT. Results from preprints are excluded. ↓ indicates "lower is better". † denotes the **NoAvg** protocol (UniAD evaluation), ‡ represents the **TemAvg** protocol (ST-P3 evaluation), and ∗ indicates the use of ego status as inputs.

4.3 Ablation Studies

4.3.1 Ablation Study on Fine-Tuning Stages

Fine-Tuning Stages			Ego Status	L2 (m)↓				Collision (%)↓			
Stage A	**Stage B**	**Stage C**	**Current** (v, a, yaw)	1s	2s	3s	Avg.	1s	2s	3s	Avg.
×	✓	✓	×	0.23	0.65	1.40	0.76	0.15	0.39	0.89	0.48
✓	×	✓	×	0.23	0.69	1.41	0.78	0.13	0.42	0.93	0.49
✓	✓	✓	×	**0.23**	**0.64**	**1.25**	**0.71**	**0.06**	**0.16**	**0.51**	**0.24**
×	✓	✓	✓	0.23	0.64	1.39	0.75	0.14	0.37	0.83	0.45
✓	×	✓	✓	0.22	0.66	1.36	0.75	0.11	0.44	0.97	0.51
✓	✓	✓	✓	**0.21**	**0.56**	**1.13**	**0.63**	**0.03**	**0.13**	**0.31**	**0.16**

Table 3 Ablation study on fine-tuning stages. Only the NoAvg protocol is adopted. ↓ indicates lower is better. The Top-K value of RAG is set to 1 accordingly.

We evaluate the contribution of the three fine-tuning stages by toggling Stage A and Stage B while keeping Stage C active, and report results under two usage modes, namely without and with ego status, using the NoAvg protocol (Top-K=1 for RAG). The numerical results are listed in Table 3. As mentioned in Section 3.4, Stage A strengthens spatial perception from multi-camera inputs, Stage B introduces trajectory-supervised motion reasoning, and Stage C jointly optimizes spatiotemporal planning on front-view sequences.

(i) Without ego status. Activating A+B+C gives the best accuracy and safety with an average L2 of 0.71m and an average collision rate of 0.24%. Disabling Stage A increases Avg L2 to 0.76m and doubles collisions to 0.48%, while disabling Stage B further raises Avg L2 to 0.78m with 0.49% collisions. The largest gains appear at longer horizons: 3s L2 drops from 1.40-1.41m (w/o A or w/o B) to 1.25m with all stages, while 1s L2 remains similar at 0.23m, indicating that the synergy of A and B mainly benefits long-horizon stability.

(ii) With ego status. The ordering is unchanged and the margins widen. A+B+C achieves 0.63m L2 and 0.16% collisions, whereas removing Stage A or Stage B yields 0.75m Avg L2 with 0.45% and 0.51% collisions, respectively. Improvements again concentrate at longer horizons (3s L2 from 1.39/1.36 m to 1.13m, 2s L2 from 0.64/0.66 m to 0.56 m), while short-horizon differences are small.

Across both usage modes, the joint presence of Stage A and Stage B (on top of Stage C) is essential: removing either stage degrades long-horizon displacement and substantially elevates collision risk, whereas A+B+C consistently provides the most accurate and safest trajectories. This matches the design intent in Section 3.4: Stage A improves metric spatial grounding, Stage B injects physically feasible motion priors, and Stage C consolidates them for temporally coherent planning.

4.3.2 Ablation Study on the Top-K Value of RAG

Top-K Retrieval	L2 (m)↓				Collision (%)↓			
	1s	2s	3s	Avg.	1s	2s	3s	Avg.
Top-0 (w/o RAG)	0.24	0.66	1.30	0.73	0.06	0.18	0.55	0.26
Top-1	0.23	0.64	1.25	0.71	0.06	0.16	0.51	0.24
Top-2	**0.23**	**0.63**	**1.23**	**0.70**	**0.05**	**0.14**	**0.44**	**0.21**
Top-3	0.25	0.68	1.27	0.73	0.07	0.19	0.59	0.28
Top-4	0.31	0.78	1.41	0.83	0.09	0.23	0.74	0.35
Top-0* (w/o RAG)	0.21	0.58	1.15	0.65	0.05	0.15	0.38	0.19
Top-1*	0.21	0.56	1.13	0.63	0.03	0.13	0.31	0.16
Top-2*	**0.19**	**0.52**	**1.12**	**0.61**	**0.02**	**0.11**	**0.29**	**0.14**
Top-3*	0.22	0.58	1.17	0.66	0.07	0.19	0.44	0.23
Top-4*	0.24	0.62	1.26	0.71	0.08	0.21	0.53	0.27

Table 4 Ablation study on the top-K value of RAG. ↓ indicates lower is better. ∗ indicates the use of ego status as inputs. Only the NoAvg protocol is adopted. ↓ indicates lower is better. All VLMs used in this ablation study are fine-tuned through the same triple-stage training process.

Table 4 examines how many retrieved evidences (K) to feed into the CoT-guided planner under the NoAvg protocol. All variants share the same triple-stage fine-tuning to isolate the effect of retrieval depth. Consistent with our embedding-retrieval design in Section 3.3 and CoT integration in Section 3.5, performance forms a concave curve: modest K supplies complementary priors, whereas larger K introduces distractors that dilute reasoning.

(i) Without ego status. Moving from Top-0 (no retrieval) to Top-2 reduces average L2 from 0.73m to 0.70m, and average collisions from 0.26% to 0.21%, with the largest gains at the 3s horizon (L2 from

1.30m to 1.23m, collision rate from 0.55% to 0.44%). Increasing K beyond 2 reverses the trend, indicating over-retrieval noise.

(ii) With ego status. The ordering remains the same and margins widen slightly: Top-2* achieves 0.61m Avg L2 and 0.14% Avg collisions, improving over Top-0* (0.65m Avg L2 and 0.19% Avg collisions) by approximately 6% and 26%, respectively, and long-horizon metrics also benefit (L2 from 1.15m to 1.12m, collision rate from 0.38% to 0.29%). Larger K hurts again.

A small but non-trivial retrieval set is optimal for our pipeline: $K = 2$ consistently offers the best accuracy-safety trade-off across usage modes, aligning with the intent of Section 3.3 and Section 3.5, to inject just enough scene-aligned evidence for the CoT planner without overwhelming it with off-topic neighbors. Therefore, we adopt Top-2 as the default for all main results.

4.3.3 Ablation Study on Different Retrieval Strategies

Retrieval Strategy	Top-1 Avg.	Top-2 Avg.	Top-3 Avg.	Top-4 Avg.	Top-5 Avg.	Overall Avg.
Simple Search	0.766	0.783	0.748	0.796	0.756	0.770
k-means only	0.162	0.159	0.149	0.150	0.150	0.154
HNSW only	0.032	0.033	0.031	0.032	0.034	0.032
k-means & HNSW	**0.014**	**0.014**	**0.015**	**0.013**	**0.014**	**0.014**

Table 5 Ablation study on the retrieval time (ms) of different retrieval strategies. $*$ uses ego status as inputs. Only the NoAvg protocol is adopted.

To assess the effectiveness of our retrieval strategy, we conduct a ablation study among the proposed k-means & HNSW pipeline, a basic similarity-based search, search with k-means clustering only, and standard HNSW. As clustering and HNSW indexing contribute little to retrieval accuracy but substantially accelerate retrieval, we focus on evaluating their impact through efficiency metrics. Specifically, in this ablation study, we report the average time required to retrieve results for a single test scene sequence as the primary indicator of retrieval performance.

For every retrieval strategy, we perform evaluations with Top-K set to 1 through 5. The final

performance metric is computed as the arithmetic mean over these Top-K results, providing a balanced assessment of retrieval speed across different thresholds. As detailed in Table 5, although simple cosine similarity search achieves millisecond-level performance even over a database of near 10,000 scene embeddings, applying k-means clustering reduces the average query latency to approximately 0.1ms. Replacing exhaustive search with HNSW yields a further improvement, reducing the retrieval time to about 0.03ms. Our proposed k-means & HNSW retrieval strategy achieves the best performance, bringing the average per-query retrieval time down to 0.014ms. All retrieval operations are executed on a single CPU (Intel i7-13700) with no GPU acceleration.

5 Conclusion

In this paper, we introduce KEPT (Knowledge-Enhanced Prediction of Trajectories), a retrieval-augmented vision-language model (VLM) framework for short-horizon trajectory prediction from consecutive driving frames. KEPT couples a temporal frequency-spatial fusion (TFSF) encoder for robust, motion-aware embeddings, a scalable k-means+HNSW retrieval stack, chain-of-thought (CoT) prompting tailored to planning constraints, and a triple-stage fine-tuning recipe that progressively grounds the VLM in metric spatial cues, physically feasible motion, and temporally coherent front-view planning.

Comprehensive experiments on nuScenes dataset demonstrate that KEPT advances various baseline methods across open-loop protocols. Under NoAvg, KEPT attains markedly lower displacement errors and collision rates than recent modular and VLM-based planners; under TemAvg, it remains competitive without ego priors and achieves the strongest overall accuracy-safety trade-off when lightweight ego status is available. Ablation studies clarify why KEPT works: (i) Stages A+B+C are complementary, removing either spatial grounding or trajectory-supervised motion reasoning disproportionately harms long-horizon stability and safety; and (ii) a small retrieval set is optimal, since Top-2 retrieved exemplars consistently provide the best balance between informative priors and distraction noise, aligning with our CoT design. Finally, we show that clustered HNSW indexing offers sub-millisecond-level retrieval latency, supporting practical deployment.

Despite the gains, KEPT has several limitations. Our retrieval priors are built from nuScenes-like

distributions; coverage and clustering granularity bound performance, and long-tail or out-of-domain scenes (e.g., extreme weather, unusual road rules) may yield suboptimal priors. The TFSF encoder currently uses seven 2 Hz front-view frames and predicts short horizons; richer sensors and longer temporal context could improve stability. CoT prompting adds variability and modest latency, and our constraint set is partly hand-crafted. Finally, we report open-loop metrics; closed-loop, real-time evaluations on a vehicle are required to verify safety and robustness.

Future work will therefore (i) extend KEPT to multi-sensor inputs and cross-city generalization with domain adaptation, (ii) develop active and safety-aware retrieval that updates the knowledge base on-the-fly, and (iii) integrate uncertainty estimation and formal safety monitoring to further reduce rare-event risk. We believe these directions, together with the lightweight nature of the Qwen-series backbones used here, make KEPT a promising foundation for trustworthy, interpretable, and scalable VLM planning in autonomous driving.

Replication and Data Sharing

The codes are available at https://github.com/yjwangtj/KEPT.

The JSONs of datasets used are available at https://huggingface.co/datasets/larswang/kept_datasets/tree/main.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

This research was supported by National Key R&D Program of China (2023YFB2504400), the National Nature Science Foundation of China (No. 62373289, No. 62273256 and No. 62473291) and the Fundamental Research Funds for the Central Universities.

References

Ahmed, A., Xiaoyang, Z., Tunio, M.H., Butt, M.H., Shah, S.A., Chengxiao, Y., Pirzado, F.A., Aziz, A., 2023. Occnet: Improving imbalanced multi-centred ovarian cancer subtype classification in whole slide images, in: 2023 20th International Computer Conference on Wavelet Active Media Technology and Information Processing (ICCWAMTIP), IEEE. pp. 1–8.

Akan, A.K., Güney, F., 2022. Stretchbev: Stretching future instance prediction spatially and temporally, in: European Conference on Computer Vision, Springer. pp. 444–460.

Atakishiyev, S., Salameh, M., Yao, H., Goebel, R., 2024. Explainable artificial intelligence for autonomous driving: A comprehensive overview and field guide for future research directions. IEEE Access .

Bandyopadhyay, S., Cole, J., Goldstein, T., Jacobs, D., 2025. Multimodal agentic model predictive control, in: Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems, pp. 2844–2848.

Brandstätter, F., Schütz, E., Winter, K., Flohr, F.B., 2025. Bev-llm: Leveraging multimodal bev maps for scene captioning in autonomous driving, in: 2025 IEEE Intelligent Vehicles Symposium (IV), IEEE. pp. 345–350.

Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O., 2020. Nuscenes: A multimodal dataset for autonomous driving, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 11621–11631.

Casas, S., Gulino, C., Liao, R., Urtasun, R., 2020. Spagnn: Spatially-aware graph neural networks for relational behavior forecasting from sensor data, in: 2020 IEEE International Conference on Robotics and Automation (ICRA), IEEE. pp. 9491–9497.

Casas, S., Sadat, A., Urtasun, R., 2021. Mp3: A unified model to map, perceive, predict and plan, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 14403–14412.

Chai, Y., Sapp, B., Bansal, M., Anguelov, D., 2020. Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction, in: Conference on Robot Learning, PMLR. pp. 86–99.

Chen, D., Krähenbühl, P., 2022. Learning from all vehicles, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 17222–17231.

Chen, L., Sinavski, O., Hünermann, J., Karnsund, A., Willmott, A.J., Birch, D., Maund, D., Shotton, J., 2024a. Driving with llms: Fusing object-level vector modality for explainable autonomous driving, in: Proceedings of the IEEE International Conference on Robotics and Automation, pp. 14093–14100.

Chen, L., Wu, P., Chitta, K., Jaeger, B., Geiger, A., Li, H., 2024b. End-to-end autonomous driving: Challenges and frontiers. IEEE Transactions on Pattern Analysis and Machine Intelligence .

Chen, X., Huang, L., Ma, T., Fang, R., Shi, S., Li, H., 2025. Solve: Synergy of language-vision and end-to-end networks for autonomous driving, in: Proceedings of the Computer Vision and Pattern Recognition Conference, pp. 12068–12077.

Chib, P.S., Singh, P., 2023. Recent advancements in end-to-end autonomous driving using deep learning: A survey. IEEE Transactions on Intelligent Vehicles 9, 103–118.

Chitta, K., Prakash, A., Geiger, A., 2021. Neat: Neural attention fields for end-to-end autonomous driving, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 15793–15803.

Cui, C., Ma, Y., Cao, X., Ye, W., Wang, Z., 2024. Receive, reason, and react: Drive as you say, with large language models in autonomous vehicles. IEEE Intelligent Transportation Systems Magazine 16, 81–94.

Dai, X., Guo, C., Tang, Y., Li, H., Wang, Y., Huang, J., Tian, Y., Xia, X., Lv, Y., Wang, F.Y., 2024. Vistarag: Toward safe and trustworthy autonomous driving through retrieval-augmented generation. IEEE Transactions on Intelligent Vehicles 9, 4579–4582.

Fu, D., Li, X., Wen, L., Dou, M., Cai, P., Shi, B., Qiao, Y., 2024. Drive like a human: Rethinking autonomous driving with large language models, in: Proceedings of the Winter Conference on Applications of Computer Vision Workshops, IEEE. pp. 910–919.

Gao, X., Wu, Y., Wang, R., Liu, C., Zhou, Y., Tu, Z., 2025. Langcoop: Collaborative driving with language, in: Proceedings of the Computer Vision and Pattern Recognition Conference, pp. 4226–4237.

Han, W., Guo, D., Xu, C.Z., Shen, J., 2025. Dme-driver: Integrating human decision logic and 3d scene perception in autonomous driving, in: Proceedings of the AAAI Conference on Artificial Intelligence,

pp. 3347–3355.

Han, X., Meng, Z., Xia, X., Liao, X., He, B.Y., Zheng, Z., Wang, Y., Xiang, H., Zhou, Z., Gao, L., Fan, L., Li, Y., Ma, J., 2024. Foundation intelligence for smart infrastructure services in transportation 5.0. IEEE Transactions on Intelligent Vehicles 9, 39–47.

Hegde, D., Yasarla, R., Cai, H., Han, S., Bhattacharyya, A., Mahajan, S., Liu, L., Garrepalli, R., Patel, V.M., Porikli, F., 2025. Distilling multi-modal large language models for autonomous driving, in: Proceedings of the Computer Vision and Pattern Recognition Conference, pp. 27575–27585.

Hu, A., Murez, Z., Mohan, N., Dudas, S., Hawke, J., Badrinarayanan, V., Cipolla, R., Kendall, A., 2021a. Fiery: Future instance prediction in bird's-eye view from surround monocular cameras, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 15273–15282.

Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W., 2022a. Lora: Low-rank adaptation of large language models, in: Proceedings of the International Conference on Learning Representations.

Hu, P., Huang, A., Dolan, J., Held, D., Ramanan, D., 2021b. Safe local motion planning with self-supervised freespace forecasting, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 12732–12741.

Hu, S., Chen, L., Wu, P., Li, H., Yan, J., Tao, D., 2022b. St-p3: End-to-end vision-based autonomous driving via spatial-temporal feature learning, in: European Conference on Computer Vision, Springer. pp. 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: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 17853–17862.

Huang, J., Ye, Y., Liang, Z., Shan, Y., Du, D., 2024a. Detecting as labeling: Rethinking lidar-camera fusion in 3d object detection, in: European Conference on Computer Vision, Springer. pp. 439–455.

Huang, Y., Zheng, W., Zhang, Y., Zhou, J., Lu, J., 2023. Tri-perspective view for vision-based 3d semantic occupancy prediction, in: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 9223–9232.

Huang, Z., Tang, T., Chen, S., Lin, S., Jie, Z., Ma, L., Wang, G., Liang, X., 2024b. Making large

language models better planners with reasoning-decision alignment, in: European Conference on Computer Vision, Springer. pp. 73–90.

Hussien, M.M., Melo, A.N., Ballardini, A.L., Maldonado, C.S., Izquierdo, R., Sotelo, M.A., 2025. Rag-based explainable prediction of road users behaviors for automated driving using knowledge graphs and large language models. Expert Systems with Applications 265, 125914.

Jiang, B., Chen, S., Xu, Q., Liao, B., Chen, J., Zhou, H., Zhang, Q., Liu, W., Huang, C., Wang, X., 2023a. Vad: Vectorized scene representation for efficient autonomous driving, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 8340–8350.

Jiang, Z., Xu, F.F., Gao, L., Sun, Z., Liu, Q., Dwivedi-Yu, J., Yang, Y., Callan, J., Neubig, G., 2023b. Active retrieval augmented generation, in: Proceedings of the Conference on Empirical Methods in Natural Language Processing, pp. 7969–7992.

Khurana, T., Hu, P., Dave, A., Ziglar, J., Held, D., Ramanan, D., 2022. Differentiable raycasting for self-supervised occupancy forecasting, in: European Conference on Computer Vision, Springer. pp. 353–369.

Li, J., Li, D., Savarese, S., Hoi, S., 2023a. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models, in: Proceedings of the International Conference on Machine Learning, pp. 19730–19742.

Li, P., Ding, S., Chen, X., Hanselmann, N., Cordts, M., Gall, J., 2023b. Powerbev: A powerful yet lightweight framework for instance prediction in bird's-eye view, in: Elkind, E. (Ed.), Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23, International Joint Conferences on Artificial Intelligence Organization. pp. 1080–1088. URL: https://doi.org/10.24963/ijcai.2023/120. main Track.

Li, Z., Wang, W., Li, H., Xie, E., Sima, C., Lu, T., Yu, Q., Dai, J., 2025. Bevformer: Learning bird's-eye-view representation from lidar-camera via spatiotemporal transformers. IEEE Transactions on Pattern Analysis and Machine Intelligence 47, 2020–2036.

Li, Z., Yu, Z., Lan, S., Li, J., Kautz, J., Lu, T., Alvarez, J.M., 2024. Is ego status all you need for open-loop end-to-end autonomous driving?, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 14864–14873.

Liang, M., Yang, B., Zeng, W., Chen, Y., Hu, R., Casas, S., Urtasun, R., 2020. Pnpnet: End-to-end perception and prediction with tracking in the loop, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 11553–11562.

Liang, T., Xie, H., Yu, K., Xia, Z., Lin, Z., Wang, Y., Tang, T., Wang, B., Tang, Z., 2022. Bevfusion: A simple and robust lidar-camera fusion framework. Advances in Neural Information Processing Systems 35, 10421–10434.

Liu, H., Li, C., Wu, Q., Lee, Y.J., 2023. Visual instruction tuning. Advances in neural information processing systems 36, 34892–34916.

Liu, Y., Zhang, J., Fang, L., Jiang, Q., Zhou, B., 2021a. Multimodal motion prediction with stacked transformers, in: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 7577–7586.

Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B., 2021b. Swin transformer: Hierarchical vision transformer using shifted windows, in: Proceedings of the IEEE/CVF international conference on computer vision, pp. 10012–10022.

Luo, X., Liu, C., Ding, F., Yang, F., Zhou, Y., Loo, J., Tew, H.H., 2025. Senserag: Constructing environmental knowledge bases with proactive querying for llm-based autonomous driving, in: Proceedings of the Winter Conference on Applications of Computer Vision, pp. 989–996.

Ma, J., Chen, X., Huang, J., Xu, J., Luo, Z., Xu, J., Gu, W., Ai, R., Wang, H., 2024a. Cam4docc: Benchmark for camera-only 4d occupancy forecasting in autonomous driving applications, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 21486–21495.

Ma, Y., Abdelraouf, A., Gupta, R., Moradipari, A., Wang, Z., Han, K., 2025. Video token sparsification for efficient multimodal llms in driving visual question answering, in: 2025 IEEE Intelligent Vehicles Symposium (IV), IEEE. pp. 2235–2242.

Ma, Y., Cao, Y., Sun, J., Pavone, M., Xiao, C., 2024b. Dolphins: Multimodal language model for driving, in: European Conference on Computer Vision, Springer. pp. 403–420.

Ma, Y., Cui, C., Cao, X., Ye, W., Liu, P., Lu, J., Abdelraouf, A., Gupta, R., Han, K., Bera, A., et al., 2024c. Lampilot: An open benchmark dataset for autonomous driving with language model programs, in: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition,

pp. 15141–15151.

Malkov, Y.A., Yashunin, D.A., 2018. Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. IEEE transactions on pattern analysis and machine intelligence 42, 824–836.

Mao, J., Qian, Y., Ye, J., Zhao, H., Wang, Y., 2023. Gpt-driver: Learning to drive with gpt, in: NeurIPS 2023 Foundation Models for Decision Making Workshop.

Ngiam, J., Vasudevan, V., Caine, B., Zhang, Z., Chiang, H.T.L., Ling, J., Roelofs, R., Bewley, A., Liu, C., Venugopal, A., et al., 2022. Scene transformer: A unified architecture for predicting future trajectories of multiple agents, in: International Conference on Learning Representations.

Pan, C., Yaman, B., Nesti, T., Mallik, A., Allievi, A.G., Velipasalar, S., Ren, L., 2024. Vlp: Vision language planning for autonomous driving, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 14760–14769.

Ram, O., Levine, Y., Dalmedigos, I., Muhlgay, D., Shashua, A., Leyton-Brown, K., Shoham, Y., 2023. In-context retrieval-augmented language models. Transactions of the Association for Computational Linguistics 11, 1316–1331.

Shao, H., Hu, Y., Wang, L., Song, G., Waslander, S.L., Liu, Y., Li, H., 2024. Lmdrive: Closed-loop end-to-end driving with large language models, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 15120–15130.

Shao, H., Wang, L., Chen, R., Waslander, S.L., Li, H., Liu, Y., 2023a. Reasonnet: End-to-end driving with temporal and global reasoning, in: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 13723–13733.

Shao, Z., Gong, Y., Huang, M., Duan, N., Chen, W., et al., 2023b. Enhancing retrieval-augmented large language models with iterative retrieval-generation synergy, in: The 2023 Conference on Empirical Methods in Natural Language Processing.

Tian, X., Gu, J., Li, B., Liu, Y., Wang, Y., Zhao, Z., Zhan, K., Jia, P., Lang, X., Zhao, H., 2025. Drivevlm: The convergence of autonomous driving and large vision-language models, in: Conference on Robot Learning, PMLR. pp. 4698–4726.

Tian, X., Jiang, T., Yun, L., Mao, Y., Yang, H., Wang, Y., Wang, Y., Zhao, H., 2023. Occ3d: A large-

scale 3d occupancy prediction benchmark for autonomous driving. Advances in Neural Information Processing Systems 36, 64318–64330.

Tong, W., Sima, C., Wang, T., Chen, L., Wu, S., Deng, H., Gu, Y., Lu, L., Luo, P., Lin, D., et al., 2023. Scene as occupancy, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 8406–8415.

Wang, S., Yu, Z., Jiang, X., Lan, S., Shi, M., Chang, N., Kautz, J., Li, Y., Alvarez, J.M., 2025a. Omnidrive: A holistic vision-language dataset for autonomous driving with counterfactual reasoning, in: Proceedings of the Computer Vision and Pattern Recognition Conference, pp. 22442–22452.

Wang, T.H., Maalouf, A., Xiao, W., Ban, Y., Amini, A., Rosman, G., Karaman, S., Rus, D., 2024. Drive anywhere: Generalizable end-to-end autonomous driving with multi-modal foundation models, in: 2024 IEEE International Conference on Robotics and Automation (ICRA), IEEE. pp. 6687–6694.

Wang, X., Zhu, Z., Xu, W., Zhang, Y., Wei, Y., Chi, X., Ye, Y., Du, D., Lu, J., Wang, X., 2023. Openoccupancy: A large scale benchmark for surrounding semantic occupancy perception, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 17850–17859.

Wang, Y., Liu, Q., Jiang, Z., Wang, T., Jiao, J., Chu, H., Gao, B., Chen, H., 2025b. Rad: Retrieval-augmented decision-making of meta-actions with vision-language models in autonomous driving, in: Proceedings of the Computer Vision and Pattern Recognition Conference, pp. 3838–3848.

Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Le, Q.V., Zhou, D., et al., 2022. Chain-of-thought prompting elicits reasoning in large language models. Advances in neural information processing systems 35, 24824–24837.

Wei, Y., Zhao, L., Zheng, W., Zhu, Z., Zhou, J., Lu, J., 2023. Surroundocc: Multi-camera 3d occupancy prediction for autonomous driving, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 21729–21740.

Wen, L., Fu, D., Li, X., Cai, X., MA, T., Cai, P., Dou, M., Shi, B., He, L., Qiao, Y., 2024. Dilu: A knowledge-driven approach to autonomous driving with large language models, in: The Twelfth International Conference on Learning Representations.

Weng, X., Ivanovic, B., Wang, Y., Wang, Y., Pavone, M., 2024. Para-drive: Parallelized architecture for real-time autonomous driving, in: Proceedings of the IEEE/CVF Conference on Computer Vision and

Pattern Recognition, pp. 15449–15458.

Wu, J., Gao, B., Gao, J., Yu, J., Chu, H., Yu, Q., Gong, X., Chang, Y., Tseng, H.E., Chen, H., et al., 2024. Prospective role of foundation models in advancing autonomous vehicles. Research 7, 0399.

Wu, P., Chen, S., Metaxas, D.N., 2020. Motionnet: Joint perception and motion prediction for autonomous driving based on bird's eye view maps, in: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 11385–11395.

Wu, P., Jia, X., Chen, L., Yan, J., Li, H., Qiao, Y., 2022. Trajectory-guided control prediction for end-to-end autonomous driving: A simple yet strong baseline. Advances in Neural Information Processing Systems 35, 6119–6132.

Xing, S., Qian, C., Wang, Y., Hua, H., Tian, K., Zhou, Y., Tu, Z., 2025. Openemma: Open-source multimodal model for end-to-end autonomous driving, in: Proceedings of the Winter Conference on Applications of Computer Vision, pp. 1001–1009.

Xu, Z., Zhang, Y., Xie, E., Zhao, Z., Guo, Y., Wong, K.Y.K., Li, Z., Zhao, H., 2024. Drivegpt4: Interpretable end-to-end autonomous driving via large language model. IEEE Robotics and Automation Letters .

Yang, K., Guo, Z., Lin, G., Dong, H., Huang, Z., Wu, Y., Zuo, D., Peng, J., Zhong, Z., Wang, X., Guo, Q., Jia, X., Yan, J., Lin, D., 2025a. Trajectory-llm: A language-based data generator for trajectory prediction in autonomous driving, in: Proceedings of the International Conference on Learning Representations.

Yang, Y., Mei, J., Ma, Y., Du, S., Chen, W., Qian, Y., Feng, Y., Liu, Y., 2025b. Driving in the occupancy world: Vision-centric 4d occupancy forecasting and planning via world models for autonomous driving, in: Proceedings of the AAAI Conference on Artificial Intelligence, pp. 9327–9335.

Yuan, J., Sun, S., Omeiza, D., Zhao, B., Newman, P., Kunze, L., Gadd, M., 2024. Rag-driver: Generalisable driving explanations with retrieval-augmented in-context learning in multi-modal large language model. CoRR .

Zheng, J., Liang, M., Yu, Y., Li, Y., Xue, Z., 2024. Knowledge graph enhanced multimodal transformer for image-text retrieval, in: Proceedings of the IEEE International Conference on Data Engineering, IEEE. pp. 70–82.