

LangCoop: Collaborative Driving with Language

Xiangbo Gao¹, Yuheng Wu², Rujia Wang¹, Chenxi Liu³, Yang Zhou¹, Zhengzhong Tu^{1*}

¹Texas A&M University, ²KAIST, ³University of Utah

{xiangbog, tzz}@tamu.edu

* Corresponding author

Abstract

*Multi-agent collaboration holds great promise for enhancing the safety, reliability, and mobility of autonomous driving systems by enabling information sharing among multiple connected agents. However, existing multi-agent communication approaches are hindered by limitations of existing communication media, including high bandwidth demands, agent heterogeneity, and information loss. To address these challenges, we introduce **LangCoop**, a new paradigm for collaborative autonomous driving that leverages natural language as a compact yet expressive medium for inter-agent communication. LangCoop features two key innovations: Mixture Model Modular Chain-of-thought (M^3CoT) for structured zero-shot vision-language reasoning and Natural Language Information Packaging (LangPack) for efficiently packaging information into concise, language-based messages. Through extensive experiments conducted in the CARLA simulations, we demonstrate that LangCoop achieves a remarkable 96% reduction in communication bandwidth ($< 2KB$ per message) compared to image-based communication, while maintaining competitive driving performance in the closed-loop evaluation. Our project page and code are at <https://xiangbogaobarry.github.io/LangCoop/>.*

1. Introduction

Recent advances in autonomous driving have demonstrated that multi-agent collaboration [30] significantly enhances both safety and efficiency compared to single-vehicle operations, primarily through real-time information sharing and intention communication. This collaborative approach has become increasingly crucial as autonomous vehicles navigate complex environments where interaction with other traffic participants is inevitable and constant. However, the selection of an appropriate collaboration medium remains a critical challenge that has attracted substantial research attention.

A key element of multi-agent collaboration is the medium used for inter-vehicle communication. Researchers have proposed various modalities for exchanging information, including: raw sensor data, neural network features, and downstream task results. Despite their utility, each of these communication media suffers from one or more critical drawbacks. Specifically, they often: **(1)** Require high bandwidth, placing a heavy load on communication infrastructures and increasing the risk of latency or packet loss. **(2)** Fail to accommodate the inherent heterogeneities across agents, which may use different sensor configurations, model architectures, or targeting on different downstream tasks. **(3)** Lose critical contextual information when data are overly compressed, abstracted, or otherwise transformed into a limited representation. **(4)** Does not support planning-level or control-level collaboration.

To address these issues, we propose that **human natural language** can serve as an effective communication medium for multi-agent collaborative driving. Unlike conventional sensor-based or feature-based communications, natural language is inherently flexible and capable of conveying a broad range of contextual and semantic cues, therefore offering additional advantages. **First**, it bridges the gap between machine-readable modalities [4] (e.g., numbers, features, embeddings) and human-spoken language, making the reasoning [25, 56], communication [22], negotiation [7], and decision-making process more transparent. Such transparency benefits research, development, and debugging by enabling human operators to understand and verify the messages being exchanged among autonomous vehicles. **Second**, ongoing research in leveraging LVLMs within autonomous driving has already demonstrated their utility in understanding [42], reasoning [50], decision-making [39, 54], and even low-level vehicle control [5]. Consequently, natural language collaboration can synergistically exploit the general intelligence of LVLMs to achieve more robust, versatile, and explainable multi-agent collaboration. **Third**, natural language enables high-level fusion or negotiation

at the planning and prediction levels, allowing agents—including automated vehicles, human drivers, and roadside units—to communicate intention and decision rationale rather than just perception data. This capability simplifies the coordination process, allowing agents to reach mutual understanding and agreements rapidly and clearly, ultimately promoting smoother, safer, and more socially acceptable driving behaviors. **Lastly**, natural language naturally provides scalability and generalization across diverse scenarios and heterogeneous vehicle platforms. Using standardized language-based communication seamlessly integrates autonomous and human-driven vehicles, regardless of sensor suites or underlying technologies. Moreover, natural language communication is inherently model-agnostic, compatible with both open-source (e.g. LLAMA [17], DeepSeek [18]) and commercial LLMs (e.g. ChatGPT [34], Gemini [41]), enabling easy adoption and interoperability across diverse autonomous vehicle systems.

Leveraging these insights, we introduce **LangCoop**, a novel framework for collaborative autonomous driving that uses natural language as the primary medium for inter-vehicle communication. Our framework consists of three key components: **(1)** a Mixture Model Modular Chain-of-thought (M3CoT) module that structures reasoning into distinct stages for comprehensive scene understanding; **(2)** a Natural Language Information Packaging (LangPack) system that compresses rich semantic information into compact messages; and **(3)** multiple driving signal generation approaches that translate natural language reasoning into actionable controls. Our experimental results in closed-loop evaluations using the Carla simulator [10] show that, by using zero-shot LVLMs, LangCoop achieves driving scores of up to 48.8 and route completion rates of up to 90.3%, significantly outperforming non-collaborative baselines while maintaining exceptional communication efficiency (<2 KB). The framework also operates effectively with heterogeneous agent capabilities, demonstrating the viability of natural language as a medium for autonomous vehicle collaboration.

2. Related Works

2.1. LVLMs in Autonomous Driving

The integration of Vision-Language Large Models (LVLMs) into autonomous driving has enabled a unified approach to perception, reasoning, and decision-making, offering enhanced interpretability and adaptability [8, 23, 49]. Early studies have explored LVLMs for closed-loop driving, where multimodal sensor data is processed alongside natural language instructions to generate vehicle control outputs. Shao *et al.* [37] introduced one of the first LVLM-based end-to-end driving

models, while Wang *et al.* [47] focused on translating language instructions into high-level driving commands. Xu *et al.* [54] and Sima *et al.* [39] further emphasized explainability, using question-answering and graph-based reasoning to interpret scene dynamics and decision rationales, making autonomous systems more transparent and human-interpretable. Hwang *et al.* [24] used LVLMs to directly output the future planning waypoints. Xing *et al.* [49] proposed a comprehensive benchmark for evaluating the truthfulness, safety, fairness, security, and generalizability of LVLMs in the autonomous driving scenes.

Beyond perception, LVLMs have demonstrated robustness in out-of-distribution (OOD) scenarios, addressing challenges that conventional deep-learning models struggle with in unseen environments. Wang *et al.* [46] showed that LVLMs could simulate novel situations through latent space editing, improving generalization. Mei *et al.* [32] introduced a dual-process framework, combining slow but rigorous reasoning from an LVLM with fast real-time execution from a smaller model, mimicking human cognitive processes. Additionally, Dong *et al.* [9] and Xing *et al.* [50] explored zero-shot prompting, demonstrating how LLMs can guide autonomous systems without extensive re-training.

LVLMs also play a pivotal role in multi-agent collaboration and human-centric driving by improving vehicular communication [48] and personalized decision-making [8, 48]. Liang *et al.* [28] and Zhang *et al.* [60] explored how generative AI models enable semantic-rich, context-aware inter-vehicle communication, surpassing traditional bandwidth-intensive numeric exchanges. In personalized driving, Li *et al.* [27] highlighted that LVLMs improve context understanding and human-like reasoning, while Lan *et al.* [26] and Duan *et al.* [11] demonstrated their ability to simulate human driving behaviors and dynamically adjust trajectories. As LVLMs continue evolving, their integration into autonomous systems paves the way for more interpretable, adaptable, and collaborative driving solutions that better align with human expectations and real-world challenges.

2.2. Collaboration Medium in Multi-agent Driving

Effective collaboration among autonomous agents in multi-agent driving scenarios hinges on the choice of communication medium. Several approaches have been explored, including the exchange of raw sensor data [1, 3, 14], neural network features [2, 6, 15, 19, 29, 43–45, 51, 53, 58], and perception results [13, 16, 33, 38, 59]. Specifically, raw sensor data (such as LiDAR point clouds or camera images) offers comprehensive envi-

ronmental perception but demands high communication bandwidth and latency. Meanwhile, neural network features (intermediate embeddings, BEV feature maps, or feature queries) can reduce bandwidth usage yet introduce incompatibility when agents rely on heterogeneous feature extraction networks. Another approach is sharing perception results, such as predicted depth maps [21], object detection outputs [52], occupancy grids [40], or BEV map segmentations [53]. While enumerating all possible perception outputs can strain communication bandwidth, limiting the shared set risks losing critical semantic details.

Given these challenges, natural language has emerged as a promising alternative for communication in multi-agent driving. Unlike numeric-based representations, natural language is compact, human-interpretable, and adaptable to heterogeneous agents. It also supports planning or control interactions. Recent studies in robotics and autonomous driving have begun to explore language-based communication, leveraging its ability to capture rich contextual information with minimal overhead. For instance, Hu *et al.* [20], Yao *et al.* [55], and Fang *et al.* [12] use Large Language Models (LLMs) for driving-scenario reasoning on highly abstract traffic descriptions but overlook pedestrians, cyclists, unknown obstacles, and environmental conditions that are pivotal in real-world driving. Another approach, V2V-LLM [4], augments an LLM backbone with pre-trained perception features (such as object detections) to incorporate environmental cues. However, it does not exploit the vision-based reasoning capabilities of LVLMs. V2X-VLM [57] is the first work to combine perception and reasoning within a LVLM framework, yet it essentially treats multi-agent collaboration as a multi-sensor fusion problem, neglecting important factors like cross-sensor coordination transformations and collaboration at the planning or control level. Moreover, its evaluation remains limited to open-loop benchmarks, and its model is not open-sourced.

In this work, we advance the field by harnessing both the perception and reasoning capabilities of LVLMs, enabling planning- and control-level collaboration among autonomous vehicular agents. Unlike previous approaches, we conduct closed-loop evaluations to assess real-time performance and provide open-source code for the research community to facilitate further exploration and benchmarking.

3. Methodology

3.1. Framework Overview

In this section, we present LangCoop, a novel framework that natively leverages Large Vision Language Models (LVLMs) for collaborative driving among Con-

nected Autonomous Vehicles (CAVs). As illustrated in Fig. 1, our framework establishes a systematic pipeline for information extraction, processing, exchange, and decision-making in collaborative driving scenarios. Each CAV initially captures front-view images through its onboard cameras, which serve as the primary sensory input. These images are passed through our Mixture Model Modular Chain-of-thought (M^3CoT) module (detailed in Section 3.2), which systematically extracts environmental and object-level information as well as process goal-oriented information, and behavioral intentions.

The extracted information is then packaged into a compact, structured natural language format via our Natural Language Information Packaging (LangPack) module. This standardized format facilitates information exchange between connected vehicles while minimizing bandwidth requirements. Concurrently, each vehicle receives packets from other CAVs within the communication range. Upon receiving the packets, each vehicle integrates the messages with its own and feeds them into the LVLMs to generate appropriate driving signals. The driving signals are formulated as discrete trajectories, continuous trajectories, or direct control commands depending on the specific implementation context (detailed in Section 3.4). These signals guide the vehicle’s planning and control systems to execute safe and efficient maneuvers.

3.2. Mixture Model Modular Chain-of-thought

The Mixture Model Modular Chain-of-thought (M^3CoT) module forms the cognitive foundation of our LangCoop framework, expanding upon the chain-of-thought reasoning process introduced by OpenEmma [50]. M^3CoT systematically decomposes the complex task of driving scene understanding into four distinct prompting stages, each addressing a specific aspect of the driving context: **driving scene description** that focuses on holistic environmental understanding, **interactive object description** that identifies and characterizes specific objects relevant to the driving task, **navigation goal prompting** that informs the agent about its next navigational goal’s relative location, shifting the agent’s perspective from mere trajectory prediction to goal-oriented planning, and finally **future intent description** that articulates the vehicle’s intended actions and decision rationale.

A key innovation in our approach is the flexibility to employ different specialized LVLMs for each prompting stage. This design choice offers several significant advantages: First, it acknowledges that different prompting tasks demand distinct capabilities—driving scene and object description rely predominantly on visual understanding capabilities, while navigation goal

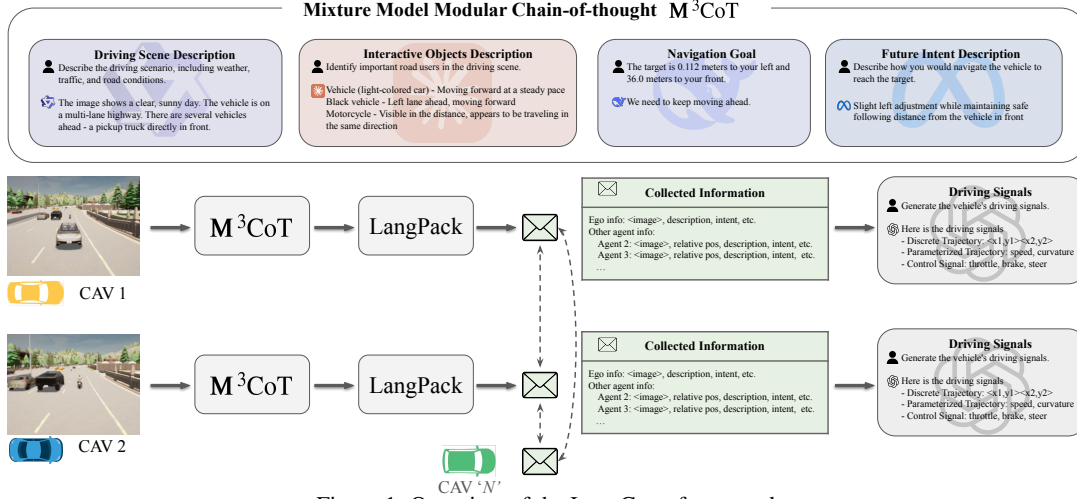


Figure 1. Overview of the LangCoop framework.

interpretation and future intent formulation necessitate stronger logical reasoning skills. By selecting models optimized for these specific competencies, our system potentially outperforms monolithic approaches that use a single model for all tasks. Second, this modular design offers practical benefits related to computational efficiency and cost management. Given that zero-shot LVLM inference can be resource-intensive, particularly for high-performance models, our approach allows for strategic resource allocation—deploying more powerful (and potentially more expensive) models only for the stages that critically require their capabilities. This alleviates the need for a single large generalist model, potentially reducing inference time and operational costs without compromising performance.

3.3. Natural Language Information Packaging

Our framework introduces Natural Language Information Packaging (LangPack) as an innovative medium for information sharing. LangPack gathers diverse information sources into a standardized, human-readable, and machine-processable format that balances comprehensiveness with transmission efficiency. Upon completing the M³CoT processing stages, each vehicle constructs a LangPack packet that integrates prompting results with agent metadata, including location, velocity, acceleration, etc.

The LangPack approach offers several distinct advantages for collaborative driving systems. First, the inherent compactness of natural language representation allows for information-dense communication with minimal bandwidth requirements—typical LangPack packages require less than 2KB of data, making them suitable for transmission even in bandwidth-constrained V2X communication environments. Furthermore, natural language provides a flexible and extensible medium that

Table 1. Nature Language Information Packaging Structure.

✉ Nature Language Information Packaging
Agent Metadata: <i>location, velocity, acceleration, etc.</i>
Scene Description: <i>The image shows ...</i>
Objects Description: <i>Vehicle (light-colored car) - Moving forward ...</i>
Navigation Goal: <i>We need to keep moving ahead ...</i>
Intent Description: <i>Slight left adjustment while maintaining safe ...</i>
...
Total Package Size: < 2KB

can accommodate diverse information types without requiring rigid structural redesigns. This adaptability is particularly valuable for autonomous driving systems that must process heterogeneous and sometimes unexpected environmental elements.

Upon receiving LangPack packages from other connected vehicles, each CAV performs essential post-processing operations including coordinate transformation and temporal alignment. The processed information is then aggregated with the vehicle’s own perceptions and prompting results to create a comprehensive knowledge ready to be passed into the following decision-making module.

3.4. Driving Signal Generation

The final component of our LangCoop framework involves translating the aggregated, multi-vehicle understanding into actionable driving signals. We propose three driving signal formulations, each offering specific advantages depending on the implementation context and downstream control requirements:

Discrete Trajectory Generation: The LVLM outputs a sequence of waypoints (x_i, y_i) for the future n sec-

onds. This high-precision path representation is suitable for complex maneuvers and enables straightforward validation against environmental boundaries.

Continuous Trajectory Generation: Rather than discrete positions, this approach defines vehicle motion through speed and turning curvature parameters over time. It produces smoother motion profiles that better align with vehicle dynamics for natural-feeling behavior.

Direct Control Signal Generation: In this most direct formulation, the LVLM outputs low-level control signals—specifically steering angle, throttle position, and brake pressure—for each time step. A key advantage of this approach is that outputs can be explicitly constrained within physically feasible ranges (e.g., steering angle limits, maximum acceleration rates), ensuring generated commands never exceed the vehicle’s operational capabilities.

In Section 4.2, we present a comparative analysis of all three driving signal formulations across diverse driving scenarios.

4. Experiments

In this section, we present comprehensive experimental evaluations of our LangCoop framework through closed-loop simulations in the CARLA environment [10]. We first outline our experimental setup and evaluation metrics (§ 4.1), followed by a systematic assessment of key components within our framework, including driving signal formulations (§ 4.2), prompting methods (§ 4.3), communication strategies (§ 4.4), LVLM selection (§ 4.5), and modular design approaches (§ 4.6). We investigate the framework’s performance under heterogeneous agent configurations [15, 31] (§ 4.7). Finally, we display some visualization results and analysis in § 4.8.

4.1. Experimental Setup

In this work we conduct closed-loop evaluations using the CARLA simulation platform. We use 10 testing scenarios in Town05 with each scenario involves two CAVs controlled by our LangCoop framework while interacting with various dynamic actors including other vehicles, pedestrians, and cyclists controlled by CARLA’s traffic manager. The two CAVs are initialized at different positions within the same general vicinity. We implement V2V communication with a simulated range of 200 meters. For perception, each vehicle receives front-view RGB camera images at 800×600 resolution.

We employ three primary evaluation metrics to assess performance comprehensively: **Driving Score (DS)**: Calculated as $DS = RC \times (1 - IP)$, where RC is route completion and IP is infraction penalty. Infractions in-

clude collisions, traffic light violations, and lane invasions, each weighted according to severity. **Route Completion (RC)**: The percentage of the predefined route successfully traversed by the vehicle, measured from 0 to 100. **Time Consumed (TC)**: The total time in seconds required to complete the route or until a terminal failure. For communication efficiency assessment, we additionally track: **Transmission Bandwidth (TB)**: The average data size in KB transmitted between vehicles.

Unless otherwise specified, our baseline configuration employs GPT-4o-mini [35] as the LVLM, utilizes a concise version of the M³CoT module described in Section 3.2, and exchanges both front-view images (compressed JPEG) and LangPack messages between vehicles.

4.2. Driving Signal Comparison

As described in Section 3.4, our framework supports three distinct driving signal formulations: discrete trajectory, continuous trajectory, and direct control signals. We first compare these formulations to identify the most effective approach for subsequent experiments.

Table 2. Performance comparison of different driving signal formulations. The discrete trajectory approach performs poorly while continuous trajectory and direct control signals show much stronger performance.

Driving Signal	Vehicle 1		Vehicle 2		TC(s)↓
	DS ↑	RC ↑	DS ↑	RC ↑	
Discrete Traj.	5.0	23.1	1.3	19.4	139.9
Continuous Traj.	33.1	74.9	48.8	90.3	124.6
Control Signal	33.7	89.0	18.1	70.2	124.8

Table 2 reveals that the discrete trajectory approach performs poorly for both vehicles. This underperformance can be attributed to the poor capability of LVLMs towards discrete waypoints understandings—it is hard for zero-shot LVLMs to output discrete waypoints that are smooth and dynamically feasible. In comparison, both continuous trajectory and direct control signal approaches demonstrate better performance. The continuous trajectory formulation achieves better performance for Vehicle 2 (DS: 48.8, RC: 90.3), while the direct control signal approach has better performance for Vehicle 1 (DS: 33.7, RC: 89.0). The continuous trajectory approach also finish the route slightly faster than other methods. We postulate that the strong performance of the continuous trajectory and direct control signal approaches stems from a more natural action space that better aligns with vehicle dynamics and control systems. Based on these results, we adopt the continuous trajectory approach as our default driving signal formulation for subsequent experiments for its balance of performance across both vehicles.

4.3. Prompting Methods Comparison

Next, we evaluate three different prompting strategies to assess the impact of reasoning structure on driving performance: Naive Prompting, which directly asks the LVLM to generate driving signals without structured reasoning, Chain-of-thought (CoT), and Concise CoT. The concise CoT variation is inducing LVLMs to output a more concise description by simply adding “Please be very concise” at the end of each prompt.

Table 3. Performance comparison of different prompting methods. The naive approach performs poorly, while both CoT approaches demonstrate strong performance.

Prompting	Vehicle 1		Vehicle 2		TC _(s) ↓
	DS ↑	RC ↑	DS ↑	RC ↑	
Naive	2.7	23.0	0.7	21.1	248.7
CoT	37.0	85.2	41.1	80.3	105.2
CoT (concise)	33.1	74.9	48.8	90.3	124.6

Table 3 demonstrates that the naive prompting approach performs poorly for both vehicles. This underscores the critical importance of structured reasoning for the autonomous driving task. Both CoT approaches substantially outperform the naive method, where there is no prominent performance priority between standard and concise CoT. The standard CoT approach achieves the highest performance for Vehicle 1 (DS: 37.0, RC: 85.2) and completes navigation in the shortest time (105.2 seconds). Meanwhile, the concise CoT variation achieves the best performance for Vehicle 2 (DS: 48.8, RC: 90.3). The performance differences between standard and concise CoT prompting highlight an interesting tradeoff. The standard CoT provides more comprehensive reasoning, potentially allowing for more nuanced decision-making, while the concise version reduces computational overhead and may focus the model on the most critical aspects of the driving task. For subsequent experiments, we adopt the concise CoT method as our default prompting strategy, as it provides strong overall performance while maintaining computational efficiency.

4.4. Communicative Message Comparison

A central aspect of our collaborative driving approach is the mechanism and content of inter-vehicle communication. We compare four different communication strategies: no collaboration (baseline), image-only sharing, LangPack-only sharing, and combined image+LangPack sharing.

As shown in Table 4, the non-collaborative baseline performs poorly with driving scores, which affirms the importance of multi-vehicular collaboration. The image-only strategy shows modest improvements over the non-collaborative baseline but falls significantly

Table 4. Performance comparison of different communication strategies. LangPack provides substantial performance gains with minimal bandwidth usage, while the combined approach achieves the highest overall performance.

Message	Vehicle 1		Vehicle 2		TC _(s) ↓	TB _(KB) ↓
	DS ↑	RC ↑	DS ↑	RC ↑		
Non-collab	13.5	33.1	11.35	29.44	200.1	0
Image (JPEG)	15.3	38.9	31.3	60.7	65.8	43.1
LangPack	35.1	71.6	42.8	80.1	<u>114.6</u>	1.8
Image+LangPack	<u>33.1</u>	74.9	48.8	90.3	124.6	44.9

short of the LangPack-based methods. This suggests that raw visual data, while information-rich, may not be optimally structured for inter-vehicle understanding without additional processing. The LangPack-only approach achieves remarkable performance (Vehicle 1: DS 35.1, RC 71.6; Vehicle 2: DS 42.8, RC 80.1) while requiring minimal bandwidth (1.8 KB), demonstrating the exceptional efficiency of our natural language packaging approach. This represents a bandwidth reduction of over 96% compared to image sharing while delivering substantially better performance. The combined Image+LangPack approach achieves the highest overall performance, particularly for Vehicle 2 (DS: 48.8, RC: 90.3), but has highest bandwidth consumption (44.9 KB).

These results demonstrate that LangPack offers an exceptional balance between performance and communication efficiency, highlighting the information density and semantic richness of structured natural language representations. For bandwidth-constrained applications, LangPack-only communication provides near-optimal performance with minimal data requirements. When bandwidth constraints are less severe, the combined approach offers incremental performance improvements at the cost of substantially higher data transmission.

4.5. LVLM Performance Comparison

The choice of LVLM significantly impacts collaborative driving performance. We evaluate six popular vision-language models (GPT-4o, Claude-3.7 Sonnet, GPT4o-mini, Gemini Flash Lite 2.0, Qwen-2.5-VL-7B, and Llama 3.2 11B Vision Instruct) to determine their effectiveness within our framework. In the following, we refer these models as GPT-4o, Claude-3.7, GPT4o-mini, Gemini-2.0, Qwen-2.5, and Llama-3.2 respectively.

Table 5 shows that GPT-4o, Claude-3.7, and GPT4o-mini consistently outperform other options across both vehicles, suggesting these models possess superior capabilities for understanding complex driving scenes and generating appropriate driving actions in collaborative contexts. The remaining models Gemini-2.0, Qwen-2.5, and Llama-3.2 demonstrate lower performance. In-

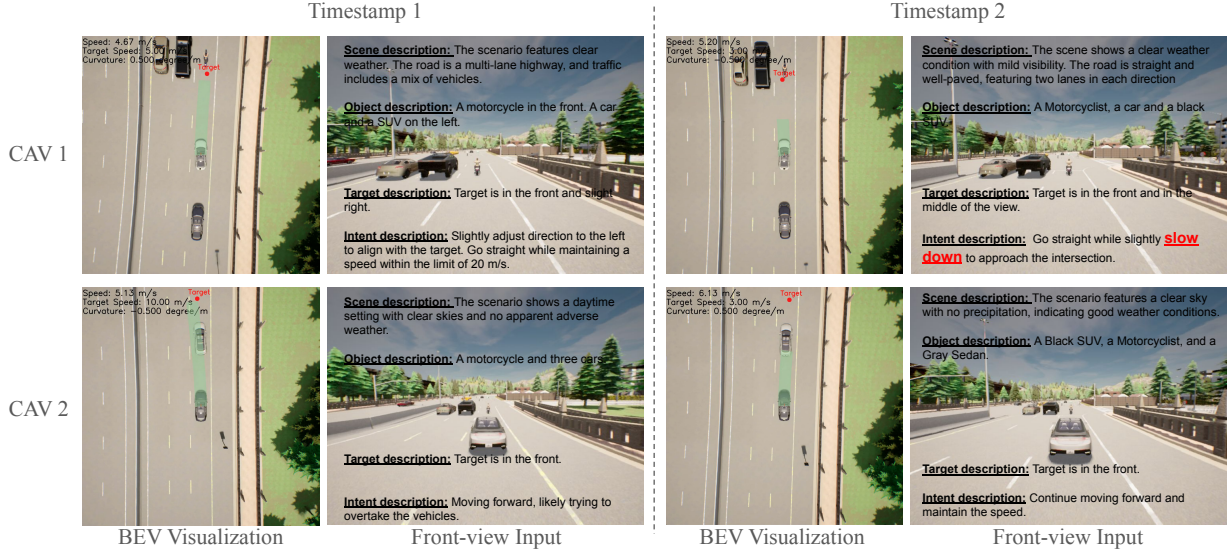


Figure 2. Visualization of a natural-language-based collaborative driving scenario. CAV 2 slows down upon receiving the ‘slow down’ intent description from CAV 1. The context is slightly paraphrased for better visualization.

Table 5. Performance comparison of different LVLMS. The top-tier commercial models (GPT-4o, Claude-3.7) demonstrate the strongest performance, with GPT-4o-mini offering competitive capabilities at lower computational cost.

Model	Vehicle 1		Vehicle 2		TC _(s) ↓
	DS ↑	RC ↑	DS ↑	RC ↑	
GPT-4o	41.3	70.0	47.7	<u>91.0</u>	79.0
Claude-3.7	32.0	<u>67.0</u>	72.1	94.1	88.5
GPT-4o-mini	<u>33.1</u>	74.9	<u>48.8</u>	90.3	124.6
Gemini-2.0	12.1	33.7	25.6	58.0	46.5
Qwen-2.5	15.5	32.2	19.4	28.8	<u>70.7</u>
Llama-3.2	11.6	31.1	19.0	42.2	102.5

terestingly, Gemini-2.0 completes routes in the shortest time (46.5 seconds), suggesting more aggressive driving behavior that may prioritize speed over safety or adherence to traffic rules.

4.6. Mixture Model Modular Design

Our M³CoTarchitecture enables the use of different specialized LVLMS for distinct reasoning stages. To evaluate the potential benefits of this modular approach, we implement two experimental configurations with varying model assignments for each prompting stage. In Experiment 6.A, we use Gemini-2.0 for driving scene and interactive objects description, Llama-3.2 for navigation goal and feature intent description, and use GPT4o-mini for driving signal generation. In Experiment 6.B, we use Qwen-2.5 for driving scene and interactive objects description, Llama-3.2 for navigation goal and feature intent description, and use GPT4o-mini for driving signal generation.

From Table 6, in experiments 6.A and 6.B, we observe that replacing the reasoning modules with LVLMS

Table 6. Performance comparison of different Mixture Model Modular (M³CoT) configurations.

M ³ CoTSetup	Vehicle 1		Vehicle 2		TC _(s) ↓
	DS ↑	RC ↑	DS ↑	RC ↑	
GPT4o-mini	33.1	74.9	48.8	90.3	124.6
Exp 6.A	31.4	67.9	37.2	71.3	144.6
Exp 6.B	35.2	68.5	42.1	82.6	119.3

other than GPT4o-mini results in slightly lower but still competitive performance compared to the pure GPT4o-mini model. Given that the API costs of Gemini-2.0 and Llama-3.2 are lower than that of GPT4o-mini, these experimental results suggest that in practical scenarios with limited computational budgets, our Mixture Model Modular Chain-of-thought module supports the possibility of replacing reasoning modules with a mixture of models.

4.7. Heterogeneous Agents Evaluation

In real-world deployments, collaborative driving systems will likely operate in environments where different vehicles utilize AI models with varying capabilities. To assess our framework’s effectiveness in such heterogeneous settings, we conduct two experiments with vehicle pairs using different LVLMS. In experiment 7.A, the vehicles are equipped with GPT-4o-mini and Gemini-2.0, while in experiment 7.B, they are equipped with GPT-4o-mini and Llama-3.2.

As shown in Table 7, collaboration improves both driving scores and route completion rates across both experiments. In experiment 7.A, pairing GPT-4o-mini with Gemini-2.0, and in experiment 7.B, pairing GPT-4o-mini with Llama-3.2, both vehicles benefit from the

Table 7. Performance in heterogeneous agent configurations where vehicles use different LVLMS.

Heterogeneous Setup			DS \uparrow	RC \uparrow	TC(s) \downarrow
Exp 7.A	Non-collab	GPT-4o-mini	18.2	56.1	167.3
		Gemini-2.0	12.6	61.1	167.3
	Image+LangPack	GPT-4o-mini	59.1	73.2	126.8
		Gemini-2.0	45.3	70.2	126.8
Exp 7.B	Non-collab	GPT-4o-mini	16.7	70.2	142.0
		Llama-3.2	11.5	51.0	142.0
	Image+LangPack	GPT-4o-mini	51.9	96.1	144.5
		Llama-3.2	12.6	40.1	144.5

collaborative setup. This demonstrates that our framework is adaptable not only to homogeneous settings but also to heterogeneous environments.

4.8. Visualization

Figure 2 displays a scenario where a leading CAV approaches an intersection and decides to slow down. After sharing its intent ‘slow down’ with other CAVs, the following vehicle also decides to slow down despite originally intending to continue forward. This demonstrates effective collaborative decision-making, as the follower vehicle appropriately adjusts its behavior based on the other CAV’s communicated intent. The example illustrates how language-based communication enables real-time adaptive driving behaviors, enhancing overall traffic safety through multi-agent decision-level collaboration. Furthermore, this interaction highlights the practical value of our framework in translating natural language intents into concrete driving decisions across multiple autonomous vehicles. For more visualization results, please refer to our anonymous project page <https://xiangbogaobarry.github.io/LangCoop/>.

5. Discussion, Limitations, and Future Work

Our experiments with LangCoop reveal several key insights that inform future research directions:

Advantage of Zero-shot LVLMS. Despite benefits of domain-specific training for LVLMS, zero-shot approaches offer clear advantages. They eliminate costly dataset collection and training while maintaining adaptability across diverse driving scenarios. Additionally, proprietary models like GPT and Gemini series cannot be fine-tuned by third parties. A zero-shot pipeline that leverages all LVLMS without domain-specific fine-tuning provides flexibility and accessibility for resource-limited institute.

Computational and Latency Concerns. Regarding computational concerns, we note that LVLMS efficiency is rapidly improving, and large models can generate trajectories for training more compact deployment models. Some novel dual-system designs[32, 42] may also alle-

viate the computational intensity. The conceptual advantages of language-based collaboration outweigh current computational demands, opening new possibilities for interpretable, efficient, and adaptable multi-agent driving systems.

Prompting Strategies for Driving. We observed significant sensitivity to prompt formulation in driving contexts. For example, we observed that explicitly instructing the model to “avoid collisions” (which might seem obvious in driving) substantially improved performance. This suggests that current LVLMS may not fully internalize driving-specific common knowledge. This indicates potential for improvement through specialized prompts or fine-tuning approaches focused on autonomous driving scenarios.

Physical-Informed Control Integration. Our current implementation does not fully incorporate detailed vehicle dynamics into the planning pipeline. Future extensions could address this by integrating physical vehicle models (e.g., bicycle model). Using techniques like quintic polynomial trajectory planning could ensure physically realizable motion while preserving the high-level reasoning capabilities of language models.

Expanding V2X Communication Paradigms. While we currently focus on vehicle-to-vehicle communication, the approach naturally extends to broader V2X ecosystems [36]. The unstructured nature of messages like Emergency Vehicle Alerts, Signal Phase and Timing, and Roadside Safety Alerts[61] aligns well with natural language representations. Future research could incorporate these additional message types, leveraging language models’ ability to process diverse information streams within a unified framework.

6. Conclusion

This work introduces LangCoop, a novel framework that leverages natural language as a communication medium for multi-agent collaborative driving. We introduce Mixture Model Modular Chain-of-thought (M3CoT) for reasoning and the Natural Language Information Packaging (LangPack) for efficient data sharing. Extensive closed-loop experiments in simulation environments demonstrate that language-based collaboration not only reduces bandwidth requirements but also enhances driving performance and interpretability by including decision-level communication. Looking forward, further optimization of prompting strategies and deeper integration of vehicle dynamics promise to extend the capabilities of language-driven autonomous systems, marking a significant step toward safer and more efficient collaborative driving.

References

- [1] Eduardo Arnold, Mehrdad Dianati, Robert de Temple, and Saber Fallah. Cooperative perception for 3d object detection in driving scenarios using infrastructure sensors. *IEEE Transactions on Intelligent Transportation Systems*, 23(3):1852–1864, 2020. 2
- [2] Zhengwei Bai, Guoyuan Wu, Matthew J Barth, Yongkang Liu, Emrah Akin Sisbot, and Kentaro Oguchi. Pillargrid: Deep learning-based cooperative perception for 3d object detection from onboard-roadside lidar. In *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, pages 1743–1749. IEEE, 2022. 2
- [3] Qi Chen, Sihai Tang, Qing Yang, and Song Fu. Cooper: Cooperative perception for connected autonomous vehicles based on 3d point clouds. In *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*, pages 514–524. IEEE, 2019. 2
- [4] Hsu-kuang Chiu, Ryo Hachiuma, Chien-Yi Wang, Stephen F Smith, Yu-Chiang Frank Wang, and Min-Hung Chen. V2v-llm: Vehicle-to-vehicle cooperative autonomous driving with multi-modal large language models. *arXiv preprint arXiv:2502.09980*, 2025. 1, 3
- [5] Can Cui, Zichong Yang, Yupeng Zhou, Juntong Peng, Sung-Yeon Park, Cong Zhang, Yunsheng Ma, Xu Cao, Wenqian Ye, Yiheng Feng, et al. On-board vision-language models for personalized autonomous vehicle motion control: System design and real-world validation. *arXiv preprint arXiv:2411.11913*, 2024. 1
- [6] Jiaxun Cui, Hang Qiu, Dian Chen, Peter Stone, and Yuke Zhu. Coopernaut: End-to-end driving with cooperative perception for networked vehicles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17252–17262, 2022. 2
- [7] Jiaxun Cui, Chen Tang, Jarrett Holtz, Janice Nguyen, Alessandro G Allievi, Hang Qiu, and Peter Stone. Talking vehicles: Cooperative driving via natural language, 2025. 1
- [8] Longchao Da, Tiejun Chen, Zhuoheng Li, Shreyas Bachiraju, Huaiyuan Yao, Xiyang Hu, Zhengzhong Tu, Yue Zhao, Dongjie Wang, Ram Pendyala, et al. Generative ai in transportation planning: A survey. *arXiv preprint arXiv:2503.07158*, 2025. 2
- [9] Zeyu Dong, Yimin Zhu, Yansong Li, Kevin Mahon, and Yu Sun. Generalizing end-to-end autonomous driving in real-world environments using zero-shot llms. *arXiv preprint arXiv:2411.14256*, 2024. 2
- [10] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator. In *Conference on robot learning*, pages 1–16. PMLR, 2017. 2, 5
- [11] Yiqun Duan, Qiang Zhang, and Renjing Xu. Prompting multi-modal tokens to enhance end-to-end autonomous driving imitation learning with llms. *arXiv preprint arXiv:2404.04869*, 2024. 2
- [12] Shiyu Fang, Jiaqi Liu, Mingyu Ding, Yiming Cui, Chen Lv, Peng Hang, and Jian Sun. Towards interactive and learnable cooperative driving automation: a large language model-driven decision-making framework. *arXiv preprint arXiv:2409.12812*, 2024. 3
- [13] Chen Fu, Chiyu Dong, Christoph Mertz, and John M Dolan. Depth completion via inductive fusion of planar lidar and monocular camera. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 10843–10848. IEEE, 2020. 2
- [14] Hongbo Gao, Bo Cheng, Jianqiang Wang, Keqiang Li, Jianhui Zhao, and Deyi Li. Object classification using cnn-based fusion of vision and lidar in autonomous vehicle environment. *IEEE Transactions on Industrial Informatics*, 14(9):4224–4231, 2018. 2
- [15] Xiangbo Gao, Runsheng Xu, Jiachen Li, Ziran Wang, Zhiwen Fan, and Zhengzhong Tu. Stamp: Scalable task and model-agnostic collaborative perception. *arXiv preprint arXiv:2501.18616*, 2025. 2, 5
- [16] Nathaniel Moore Glaser and Zsolt Kira. We need to talk: Identifying and overcoming communication-critical scenarios for self-driving. *arXiv preprint arXiv:2305.04352*, 2023. 2
- [17] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv e-prints*, pages arXiv–2407, 2024. 2
- [18] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025. 2
- [19] Jingda Guo, Dominic Carrillo, Sihai Tang, Qi Chen, Qing Yang, Song Fu, Xi Wang, Nannan Wang, and Paparao Palacharla. Coff: Cooperative spatial feature fusion for 3-d object detection on autonomous vehicles. *IEEE Internet of Things Journal*, 8(14):11078–11087, 2021. 2
- [20] Senkang Hu, Zhengru Fang, Zihan Fang, Yiqin Deng, Xianhao Chen, and Yuguang Fang. Agentscodriver: Large language model empowered collaborative driving with lifelong learning. *arXiv preprint arXiv:2404.06345*, 2024. 3
- [21] Yue Hu, Yifan Lu, Runsheng Xu, Weidi Xie, Siheng Chen, and Yanfeng Wang. Collaboration helps camera overtake lidar in 3d detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9243–9252, 2023. 3
- [22] Yue Hu, Xianghe Pang, Xiaoqi Qin, Yonina C Eldar, Siheng Chen, Ping Zhang, and Wenjun Zhang. Pragmatic communication in multi-agent collaborative perception. *arXiv preprint arXiv:2401.12694*, 2024. 1
- [23] Yue Huang, Chujie Gao, Siyuan Wu, Haoran Wang, Xiangqi Wang, Yujun Zhou, Yanbo Wang, Jiayi Ye, Jiawen Shi, Qihui Zhang, et al. On the trustworthiness of generative foundation models: Guideline, assessment, and perspective. *arXiv preprint arXiv:2502.14296*, 2025. 2
- [24] Jyh-Jing Hwang, Runsheng Xu, Hubert Lin, Wei-Chih Hung, Jingwei Ji, Kristy Choi, Di Huang, Tong He,

- Paul Covington, Benjamin Sapp, et al. Emma: End-to-end multimodal model for autonomous driving. *arXiv preprint arXiv:2410.23262*, 2024. 2
- [25] Bo Jiang, Shaoyu Chen, Qian Zhang, Wenyu Liu, and Xinggang Wang. Alphadrive: Unleashing the power of vlms in autonomous driving via reinforcement learning and reasoning, 2025. 1
- [26] Zhengxing Lan, Lingshan Liu, Bo Fan, Yisheng Lv, Yilong Ren, and Zhiyong Cui. Traj-llm: A new exploration for empowering trajectory prediction with pre-trained large language models. *IEEE Transactions on Intelligent Vehicles*, 2024. 2
- [27] Yun Li, Kai Katsumata, Ehsan Javanmardi, and Manabu Tsukada. Large language models for human-like autonomous driving: A survey. *arXiv preprint arXiv:2407.19280*, 2024. 2
- [28] Chengsi Liang, Hongyang Du, Yao Sun, Dusit Niyato, Jiawen Kang, Dezong Zhao, and Muhammad Ali Imran. Generative ai-driven semantic communication networks: Architecture, technologies and applications. *IEEE Transactions on Cognitive Communications and Networking*, 2024. 2
- [29] Genjia Liu, Yue Hu, Chenxin Xu, Weibo Mao, Junhao Ge, Zhengxiang Huang, Yifan Lu, Yinda Xu, Junkai Xia, Yafei Wang, et al. Towards collaborative autonomous driving: Simulation platform and end-to-end system. *arXiv preprint arXiv:2404.09496*, 2024. 2
- [30] Si Liu, Chen Gao, Yuan Chen, Xingyu Peng, Xianghao Kong, Kun Wang, Runsheng Xu, Wentao Jiang, Hao Xiang, Jiaqi Ma, et al. Towards vehicle-to-everything autonomous driving: A survey on collaborative perception. *arXiv preprint arXiv:2308.16714*, 2023. 1
- [31] Yifan Lu, Yue Hu, Yiqi Zhong, Dequan Wang, Siheng Chen, and Yanfeng Wang. An extensible framework for open heterogeneous collaborative perception. *arXiv preprint arXiv:2401.13964*, 2024. 5
- [32] Jianbiao Mei, Yukai Ma, Xueming Yang, Licheng Wen, Xinyu Cai, Xin Li, Daocheng Fu, Bo Zhang, Pinlong Cai, Min Dou, et al. Continuously learning, adapting, and improving: A dual-process approach to autonomous driving. *arXiv preprint arXiv:2405.15324*, 2024. 2, 8
- [33] Gledson Melotti, Cristiano Premebida, and Nuno Gonçalves. Multimodal deep-learning for object recognition combining camera and lidar data. In *2020 IEEE International Conference on Autonomous Robot Systems and Competitions (ICARSC)*, pages 177–182. IEEE, 2020. 2
- [34] OpenAI. Introducing chatgpt. <https://openai.com/index/chatgpt/>, 2022. [Accessed 13-03-2025]. 2
- [35] OpenAI. Gpt-4o-mini, 2024. Large language model, available at <https://openai.com>. 5
- [36] SAE International. V2X Communications Message Set Dictionary. Technical Report SAE J2735-202409, SAE International, 2024. Revised September 2024. 8
- [37] Hao Shao, Yuxuan Hu, Letian Wang, Guanglu Song, Steven L Waslander, Yu Liu, and Hongsheng Li. Lm-drive: Closed-loop end-to-end driving with large language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15120–15130, 2024. 2
- [38] Shuyao Shi, Jiahe Cui, Zhehao Jiang, Zhenyu Yan, Guoliang Xing, Jianwei Niu, and Zhenchao Ouyang. Vips: Real-time perception fusion for infrastructure-assisted autonomous driving. In *Proceedings of the 28th annual international conference on mobile computing and networking*, pages 133–146, 2022. 2
- [39] Chonghao Sima, Katrin Renz, Kashyap Chitta, Li Chen, Hanxue Zhang, Chengen Xie, Jens Beißwenger, Ping Luo, Andreas Geiger, and Hongyang Li. Drivelm: Driving with graph visual question answering. In *European Conference on Computer Vision*, pages 256–274. Springer, 2024. 1, 2
- [40] Rui Song, Chenwei Liang, Hu Cao, Zhiran Yan, Walter Zimmer, Markus Gross, Andreas Festag, and Alois Knoll. Collaborative semantic occupancy prediction with hybrid feature fusion in connected automated vehicles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17996–18006, 2024. 3
- [41] Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023. 2
- [42] Xiaoyu Tian, Junru Gu, Bailin Li, Yicheng Liu, Yang Wang, Zhiyong Zhao, Kun Zhan, Peng Jia, Xianpeng Lang, and Hang Zhao. Drivevlm: The convergence of autonomous driving and large vision-language models. *arXiv preprint arXiv:2402.12289*, 2024. 1, 8
- [43] Binglu Wang, Lei Zhang, Zhaozhong Wang, Yongqiang Zhao, and Tianfei Zhou. Core: Cooperative reconstruction for multi-agent perception. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8676–8686. IEEE Computer Society, 2023. 2
- [44] Rujia Wang, Xiangbo Gao, Hao Xiang, Runsheng Xu, and Zhengzhong Tu. Cocmt: Communication-efficient cross-modal transformer for collaborative perception, 2025.
- [45] Tsun-Hsuan Wang, Sivabalan Manivasagam, Ming Liang, Bin Yang, Wenyuan Zeng, and Raquel Urtasun. V2vnet: Vehicle-to-vehicle communication for joint perception and prediction. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 605–621. Springer, 2020. 2
- [46] Tsun-Hsuan Wang, Alaa Maalouf, Wei Xiao, Yutong Ban, Alexander Amini, Guy Rosman, Sertac Karaman, and Daniela Rus. Drive anywhere: Generalizable end-to-end autonomous driving with multi-modal foundation models. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6687–6694. IEEE, 2024. 2
- [47] Wenhai Wang, Jiangwei Xie, ChuanYang Hu, Haoming Zou, Jianan Fan, Wenwen Tong, Yang Wen, Silei Wu, Hanming Deng, Zhiqi Li, et al. Drivelm: Aligning multi-modal large language models with behavioral

- planning states for autonomous driving. *arXiv preprint arXiv:2312.09245*, 2023. 2
- [48] Keshu Wu, Pei Li, Yang Zhou, Rui Gan, Junwei You, Yang Cheng, Jingwen Zhu, Steven T Parker, Bin Ran, David A Noyce, et al. V2x-llm: Enhancing v2x integration and understanding in connected vehicle corridors. *arXiv preprint arXiv:2503.02239*, 2025. 2
- [49] Shuo Xing, Hongyuan Hua, Xiangbo Gao, Shenzhe Zhu, Renjie Li, Kexin Tian, Xiaopeng Li, Heng Huang, Tianbao Yang, Zhangyang Wang, et al. Autotrust: Benchmarking trustworthiness in large vision language models for autonomous driving. *arXiv preprint arXiv:2412.15206*, 2024. 2
- [50] Shuo Xing, Chengyuan Qian, Yuping Wang, Hongyuan Hua, Kexin Tian, Yang Zhou, and Zhengzhong Tu. Openemma: Open-source multimodal model for end-to-end autonomous driving. In *Proceedings of the Winter Conference on Applications of Computer Vision*, pages 1001–1009, 2025. 1, 2, 3
- [51] Runsheng Xu, Hao Xiang, Zhengzhong Tu, Xin Xia, Ming-Hsuan Yang, and Jiaqi Ma. V2x-vit: Vehicle-to-everything cooperative perception with vision transformer. In *European conference on computer vision*, pages 107–124. Springer, 2022. 2
- [52] Runsheng Xu, Jinlong Li, Xiaoyu Dong, Hongkai Yu, and Jiaqi Ma. Bridging the domain gap for multi-agent perception. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6035–6042. IEEE, 2023. 3
- [53] Runsheng Xu, Zhengzhong Tu, Hao Xiang, Wei Shao, Bolei Zhou, and Jiaqi Ma. Cobevt: Cooperative bird’s eye view semantic segmentation with sparse transformers. In *Conference on Robot Learning*, pages 989–1000. PMLR, 2023. 2, 3
- [54] Zhenhua Xu, Yujia Zhang, Enze Xie, Zhen Zhao, Yong Guo, Kwan-Yee K Wong, Zhenguo Li, and Hengshuang Zhao. Drivegpt4: Interpretable end-to-end autonomous driving via large language model. *IEEE Robotics and Automation Letters*, 2024. 1, 2
- [55] Huaiyuan Yao, Longchao Da, Vishnu Nandam, Justin Turnau, Zhiwei Liu, Linsey Pang, and Hua Wei. Comal: Collaborative multi-agent large language models for mixed-autonomy traffic. *arXiv preprint arXiv:2410.14368*, 2024. 3
- [56] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations (ICLR)*, 2023. 1
- [57] Junwei You, Haotian Shi, Zhuoyu Jiang, Zilin Huang, Rui Gan, Keshu Wu, Xi Cheng, Xiaopeng Li, and Bin Ran. V2x-vlm: End-to-end v2x cooperative autonomous driving through large vision-language models. *arXiv preprint arXiv:2408.09251*, 2024. 3
- [58] Haibao Yu, Yingjuan Tang, Enze Xie, Jilei Mao, Jirui Yuan, Ping Luo, and Zaiqing Nie. Vehicle-infrastructure cooperative 3d object detection via feature flow prediction. *arXiv preprint arXiv:2303.10552*, 2023. 2
- [59] Wenyuan Zeng, Shenlong Wang, Renjie Liao, Yun Chen, Bin Yang, and Raquel Urtasun. Dsdnet: Deep structured self-driving network. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16*, pages 156–172. Springer, 2020. 2
- [60] Ruichen Zhang, Ke Xiong, Hongyang Du, Dusit Niyato, Jiawen Kang, Xuemin Shen, and H Vincent Poor. Generative ai-enabled vehicular networks: Fundamentals, framework, and case study. *IEEE Network*, 2024. 2
- [61] Walter Zimmer, Ross Greer, Xingcheng Zhou, Rui Song, Marc Pavel, Daniel Lehmberg, Ahmed Ghita, Akshay Gopalkrishnan, Mohan Trivedi, and Alois Knoll. Enhancing highway safety: Accident detection on the a9 test stretch using roadside sensors. *arXiv preprint arXiv:2502.00402*, 2025. 8

LangCoop: Collaborative Driving with Language

Supplementary Material

Example of Natural Language Information Packaging (LangPack) – Sent by Agent 1, Received by Agent 0

1.5 KB

Agent 1, located at: [2.69048, 69.03092], current speed: 4.518m/s.

It's **scene description**:

The driving scenario shows a clear, daytime environment with sunny weather. The road appears to be well-maintained, featuring multiple lanes in both directions. There is moderate traffic, including vehicles such as cars and a motorcycle. Trees line the road-side, indicating a suburban or semi-rural area. Overall, conditions are favorable for driving.

It's **object description**:

1. Motorcycle
 - Location: Center lane, mid-distance
 - Status: Riding in the same direction as the vehicle
 - Intent: Likely continuing straight ahead on the road.
2. Black SUV
 - Location: Right lane, close to the front of the vehicle
 - Status: Stationary
 - Intent: Appears to be waiting or preparing to merge into the lane.
3. Red Car
 - Location: Left lane, further ahead
 - Status: Moving slowly, closer to the center divider
 - Intent: Preparing to turn or change lanes.

It's **target description**:

The target is 0.15402 meters to my right and 32.39753 meters to my front.

It's **intent description**:

1. Check Environment: Identify the surrounding vehicles and road conditions.
2. Speed Control: Maintain a speed and adhere to the speed limit.
3. Adjust Direction: Slightly adjust to the right to align with the target location, ensuring no obstacles are in the way.
4. Avoid Collisions: The car ahead is too close, slow down while adjusting to ensure safe distance.

Example of Natural Language Information Packaging (LangPack) – Sent by Agent 0, Received by Agent 1

1.4 KB

Agent 0, located at: [2.25215 80.50174], current speed: 5.630m/s.

It's **scene description**:

The scenario depicts a clear, sunny day with good visibility. The road is a multi-lane paved highway with light to moderate traffic. Several vehicles are moving in both directions, and there are trees lining the road-side, indicating a suburban or rural environment. The road condition appears smooth and well-maintained.

It's **object description**:

1. Car (front center): Positioned in the right lane, traveling at a moderate speed. Intent appears to be straightforward driving.
2. Motorcycle (left center): Located in the left lane, approaching the vehicle. Likely intent is to overtake the car in the center lane.
3. Car (right lane, further ahead): Positioned in the right lane, moving in the same direction as others. Status indicates normal driving behavior.

It's **target description**:

The target is 0.86387 meters to by right and 36.0 meters to by front. The target is not an endpoint—continue moving forward after reaching it.

It's **intent description**:

1. Decelerate: Reduce speed to stay within the speed limit of 20 m/s.
2. Slightly Adjust Direction: Steer right towards the target (0.15402 meters to your right).
3. Monitor Traffic: Vehicles are ahead. To ensure a safe distance, slow down or change lanes if necessary.
4. Continue Forward: Maintain forward motion, adjusting as needed for further navigation.