

# UNCAP: Uncertainty-Guided Neurosymbolic Planning Using Natural Language Communication for Cooperative Autonomous Vehicles

Neel P. Bhatt\*

The University of Texas at Austin  
Austin, Texas, United States  
npbhatt@utexas.edu

Po-han Li\*

The University of Texas at Austin  
Austin, Texas, United States  
pohanli@utexas.edu

Kushagra Gupta\*

The University of Texas at Austin  
Austin, Texas, United States  
kushagrag@utexas.edu

Rohan Siva

The University of Texas at Austin  
Austin, Texas, United States  
rohansiva@utexas.edu

Daniel Milan

The University of Texas at Austin  
Austin, Texas, United States  
dm10787@utexas.edu

Alexander T. Hogue

The University of Texas at Austin  
Austin, Texas, United States  
alex.hogue@utexas.edu

Sandeep P. Chinchali

The University of Texas at Austin  
Austin, Texas, United States  
sandeepc@utexas.edu

David Fridovich-Keil

The University of Texas at Austin  
Austin, Texas, United States  
dfk@utexas.edu

Zhangyang Wang

The University of Texas at Austin  
Austin, Texas, United States  
atlaswang@utexas.edu

Ufuk Topcu

The University of Texas at Austin  
Austin, Texas, United States  
utopcu@utexas.edu

## ABSTRACT

Safe large-scale coordination of multiple cooperative connected autonomous vehicles (CAVs) hinges on communication that is both efficient and interpretable. Existing approaches either rely on transmitting high-bandwidth raw sensor data streams or neglect perception and planning uncertainties inherent in shared data, resulting in systems that are neither scalable nor safe. To address these limitations, we propose **Uncertainty-Guided Natural Language Cooperative Autonomous Planning (UNCAP)**, a vision-language model-based planning approach that enables CAVs to communicate via lightweight natural language messages while explicitly accounting for perception uncertainty in decision-making. UNCAP features a two-stage communication protocol: (i) an ego CAV first identifies the subset of vehicles most relevant for information exchange, and (ii) the selected CAVs then transmit messages that quantitatively express their perception uncertainty. By selectively fusing messages that maximize mutual information, this strategy allows the ego vehicle to integrate only the most relevant signals into its decision-making, improving both the scalability and reliability of cooperative planning. Experiments across diverse driving scenarios show a 63% reduction in communication bandwidth with a 31% increase in driving safety score, a 61% reduction in decision uncertainty, and a four-fold increase in collision distance margin during near-miss events. Project website: <https://uncap-project.github.io/>.

## KEYWORDS

Natural Language Communication; Planning Under Uncertainty; Cooperative Autonomous Vehicles; Vision-Language Models; Safety and Efficiency

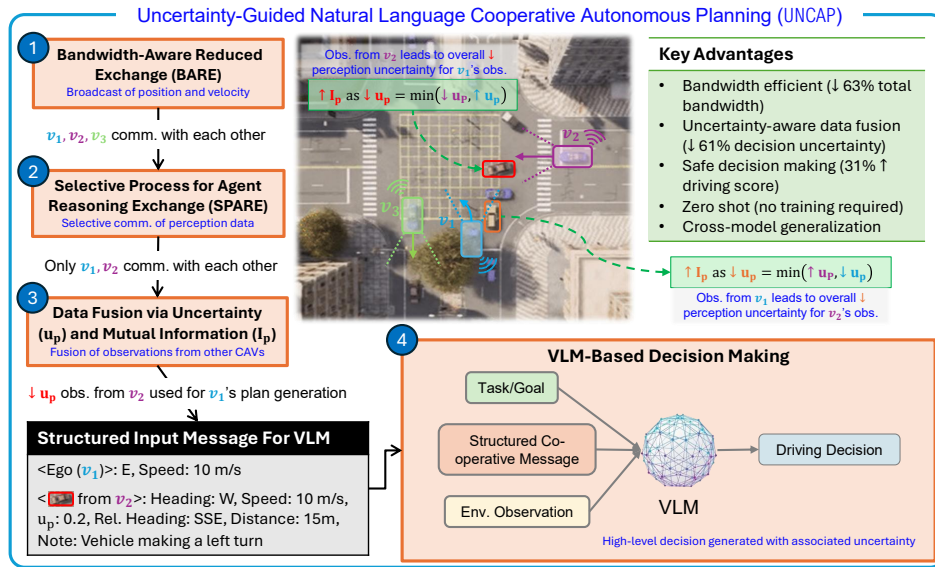
### ACM Reference Format:

Neel P. Bhatt\*, Po-han Li\*, Kushagra Gupta\*, Rohan Siva, Daniel Milan, Alexander T. Hogue, Sandeep P. Chinchali, David Fridovich-Keil, Zhangyang Wang, and Ufuk Topcu. 2026. UNCAP: Uncertainty-Guided Neurosymbolic Planning Using Natural Language Communication for Cooperative Autonomous Vehicles. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 9 pages. <https://doi.org/10.65109/WHRE3513>

## 1 INTRODUCTION

We study cooperation among connected autonomous vehicles (CAVs) using *natural language* communication as the medium of mutual interaction. While inter-vehicle cooperation has been extensively studied, most existing approaches focus on transmitting raw sensor data or features extracted from deep neural networks [13, 34, 35]. These modalities require substantial communication bandwidth, impose high computational costs for inference, and implicitly assume homogeneous sensing and processing capabilities across cooperative vehicles. Although a few recent works have explored cooperative communication through natural language in CAVs equipped with vision-language models (VLMs), these methods are typically restricted to pairwise interactions and do not explicitly account for uncertainties in vehicle perception, raising the question of how language-based communication and planning

\*These authors contributed equally to this work.



**Figure 1: Overview of UNCAP.** To minimize bandwidth, vehicles ( $v_1, v_2, v_3$ ) first engage in BARE and share essential state information. SPARE then enables selective communication among relevant agents ( $v_1, v_2$ ) to focus on critical interactions. Observations are fused based on their contribution to reducing perception uncertainty ( $u_p$ ), prioritizing information that maximizes mutual information ( $I_p$ ) for the ego vehicle. The fused messages are then put in a structured format and used by a VLM to produce driving decisions with an associated decision uncertainty score ( $u_d$ ) for safe and interpretable planning.

frameworks can scale to larger CAV fleets while accounting for perception uncertainties [11, 14, 26].

The use of VLMs opens new opportunities for cooperative autonomous driving. By bridging raw perceptual inputs with high-level reasoning, VLMs allow vehicles to interpret their surroundings, explain their behaviors, and incorporate human feedback. In particular, they enable CAVs to exchange semantic information through text rather than raw sensor data, supporting scalable, low-bandwidth coordination. For instance, a vehicle might report “pedestrian entering crosswalk at 3 m, speed 1.2 m/s” or broadcast “changing lane left in 2 s”, allowing nearby planners to adjust trajectories and avoid conflicts. Moreover, VLMs can enrich such messages with confidence scores and relevance tags, so that downstream planning modules selectively integrate only the messages that reduce decision uncertainty.

Beyond coordination, VLMs can adapt to human driving preferences and traffic regulations expressed in natural language, e.g. “maintain a safe distance on wet roads” or “drive more conservatively near schools”. This capability enhances both flexibility and interpretability in autonomous driving. Semantic natural language messages also reduce transmission latency relative to full sensor streams, preserve privacy by avoiding the exchange of raw images, and provide interpretable explanations that can be audited by human operators and regulators.

Natural language has been previously explored as a medium for collaborative autonomous driving in works such as LangCoop [11, 14]. However, these frameworks face key challenges in scalability and robustness for real-world deployment. First, they evaluate performance primarily with traditional metrics such as driving score or time-to-collision, which reflect driving quality but fail

to capture perception and planning certainty and bandwidth constraints, all of which are crucial for decision-making in autonomous driving. Second, communication is restricted to pairwise exchanges; this scales poorly to larger fleets and risks redundancy or miscoordination, highlighting the need for selective communication. Third, all received messages are treated as equally useful, without mechanisms to filter for information that meaningfully reduces decision uncertainty. Finally, approaches such as LangCoop are in particular vulnerable to communication errors or VLM misinterpretations, and lack guarantees for safe decision-making under operational constraints. To address these challenges, we make the following three contributions:

- (1) **Efficient and Robust Communication framework:** We propose **Uncertainty-Guided Natural Language Cooperative Autonomous Planning (UNCAP)**: a zero-shot, two-stage natural language-based communication and planning framework for CAVs that explicitly incorporates perception uncertainties into decision-making. In the first stage, we introduce **Bandwidth-Aware Reduced Exchange (BARE)** followed by **Selective Process for Agent Reasoning Exchange (SPARE)** in the second stage, which together allow an ego CAV to select relevant communication partners in an online fashion, thereby *improving bandwidth efficiency*. Further, we propose a mechanism that enables communicating CAVs to quantify and share their perception uncertainty. The ego CAV then strategically fuses only the most informative messages in a *zero-shot* manner, *improving robustness and safety in cooperative planning*.
- (2) **Accounting for Uncertainty and Mutual Information:** We introduce **Information Gain (IG)**, an uncertainty-guided

metric that extends beyond conventional metrics such as driving score and time-to-collision, which do not account for uncertainties in transmitted communication. This metric evaluates the perception and planning uncertainties in VLMs for multi-agent communication and quantifies the value of shared observations via the mutual information between communicating CAVs.

- (3) **Enhanced Empirical Performance:** We evaluate UNCAP on OPV2V [35], a diverse dataset for CAV driving scenarios in the CARLA simulator [12], to demonstrate that UNCAP reduces bandwidth cost by 63%, increases driving score by 31%, lowers planning uncertainty by 61%, and improves safety with a 4× increase in collision distance margin in near-miss scenarios.

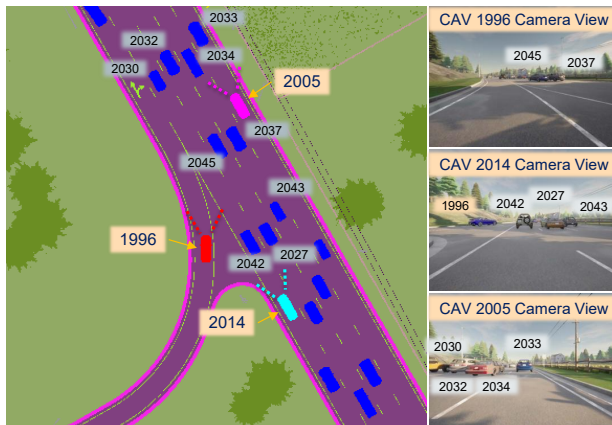
Together, these contributions establish a decision-making framework for VLM-based CAV coordination that is both uncertainty-guided and communication-efficient, and demonstrate measurable improvements in safety and planning performance.

## 2 RELATED WORKS

**Language for Autonomous Driving.** The capabilities of Large Language Models (LLMs) for decision-making are of increasing interest in the field of autonomous driving, especially in the *non-cooperative* setting. Prior non-cooperative works have focused on incorporating (possibly evolving) knowledge bases and common sense human-style reasoning into high-level driving decision-making [9, 21, 22, 26, 33], and using language as a medium for conflict resolution and negotiation at intersection-like scenarios, where agents have to coordinate the order in which they cross [20]. These approaches remain non-cooperative in nature and primarily address decision-making for a CAV, without leveraging inter-agent communication or coordination.

**Cooperative Autonomous Driving.** Research in the field of cooperative CAV communication and decision-making has primarily focused on *cooperative perception*, where CAVs share sensor data [3, 5, 10, 13, 16] and/or network features [30, 32, 34]. However, such raw sensor data is often expensive to store, transmit, and infer, and the latency of these operations can be unsuitable for large-scale cooperative driving scenarios. Recent works have investigated the use of text as a medium of communication between CAVs [11, 14, 18]. However, it is not clear how these works scale beyond the immediate two-CAV scenario, where the questions of “who to talk to” and “what information to filter out” naturally arise.

**Uncertainty Calibration in Decision Making.** Several post-processing calibration techniques have been proposed to improve the reliability of predictive models. Calibration ensures that the model’s predicted probabilities accurately reflect the true likelihood of correctness. We focus on methods that quantify uncertainty in classification tasks by estimating the conditional probability that a given sample belongs to each class, conditioned on the observed input. Platt scaling [25] fits a logistic regression to model outputs to produce calibrated probabilities but is primarily restricted for binary classification. For multiclass settings, more expressive methods such as temperature scaling [15] and Dirichlet calibration [19] provide greater flexibility and improved calibration performance. Temperature scaling introduces a single temperature parameter



**Figure 2: (Left) Labeled bird’s-eye view depicting CAV 1996’s intention to merge into highway traffic. CAV 2014 detects relevant vehicles such as 2042 and shares this information with CAV 1996, aiding in decision making. (Right) Front camera views of CAVs.**

to soften logits, offering a simple calibration method for neural network-based approaches without affecting accuracy, while vector scaling learns class-specific scaling and bias parameters [15]. Dirichlet calibration fits a Dirichlet distribution over class probabilities. Beyond such probabilistic calibration methods, conformal prediction [27, 29] is a popular approach and offers a model-agnostic framework to construct prediction sets or intervals with probabilistic guarantees; conformal prediction has been extensively utilized in safety-critical autonomous system frameworks [4, 17, 23].

## 3 METHODOLOGY

We present an uncertainty-guided communication and planning approach for VLM-enabled CAVs. Our method is designed to answer three central questions: (i) who to communicate with, (ii) how to quantify uncertainty and the value of communication, and (iii) how to guarantee safety of decisions made under uncertainty.

**Running Example.** Throughout Section 3, we use a highway merging scenario as a running example to illustrate our approach. As shown in Figure 2, this scenario involves CAV 1996 (assumed as ego CAV) merging onto a highway where CAV 2014 provides relevant information about 2042, which CAV 2005 cannot.

### 3.1 Problem Setup

We consider a multi-CAV driving scenario involving a set of CAVs  $\mathcal{V} = \{v_1, v_2, \dots, v_N\}$  operating in a shared environment. Each vehicle  $v_i \in \mathcal{V}$  perceives its surroundings through onboard sensors and obtains a local observation  $o_i$ . From  $o_i$  and onboard sensors, the vehicle generates a textual semantic state representation  $s_i = (p_{ego}, p'_{ego})$ , where  $p_{ego}$  and  $p'_{ego}$  denote its position and velocity, respectively. The textual representation  $s_i$  captures high-level contextual information that may be used for selective communication. Examples of  $s_i$  have been provided in Appendix ??.

Given the goal to minimize the communication bandwidth budget while preserving safety, each CAV  $v_i$  must decide (i) which peers  $v_j \in \mathcal{V} \setminus \{v_i\}$  to communicate with, (ii) what subset of semantic content from  $v_j$  is valuable, and (iii) what plan to execute

to drive safely. The objective of all CAVs is to minimize bandwidth consumption while ensuring that they generate a plan above a specified confidence threshold. We formulate this as a constrained optimization problem:

$$\begin{aligned} \min_{\forall b_{ij} \in \{0,1\}, \pi_i} & \sum_{i,j \in [1,\dots,N], i \neq j} b_{ij}, \\ \text{s.t.} & \max_j b_{ij} I_p(o_j; \pi_i | o_i) \geq \tau_{\text{safety}}, \forall i, \end{aligned} \quad (1)$$

where  $b_{ij}$  is a binary integer indicating the bandwidth cost of communicating with  $v_j$ ,  $I_p(o_j; \pi_i | o_i)$  is the pointwise mutual information between the set of shared observations  $o_j$  and the ego CAV’s plan  $\pi_i$ , representing the reduction in planning uncertainty, and  $\tau_{\text{safety}}$  is a threshold ensuring that the communicated information is sufficient for safe planning.

In realistic deployments, communication bandwidth is not binary as in Equation (1); CAVs can reduce transmission costs through selective and compressed messaging. For instance, in the proposed method, broadcasting minimal textual descriptions can replace high-bandwidth image sharing while preserving essential context, further reducing communication costs. Also, semantic messages may still suffer from noise or delay. The formulated problem simplifies this realistic setting, where each vehicle must adaptively prioritize uncertainty-reducing and safety-critical information (e.g., “pedestrian crossing ahead”) while suppressing redundant or low-impact updates (e.g., “lane keeping stable”) to maintain safe and robust planning under dynamic network and sensing conditions.

### 3.2 Overview of UNCAP

We now describe the overall proposed method, UNCAP. As shown in Figure 1, UNCAP has four stages. In stage 1, all CAVs  $\mathcal{V} = \{v_1, \dots, v_N\}$  broadcast a lightweight message containing their position and velocity, without sharing their sensor observations. The message exchange has complexity  $O(N^2)$  as it occurs pairwise and is not selective. In the running example, it corresponds to CAVs 1996, 2005, and 2014 sending their positions and velocities to others.

After broadcasting, in stage 2, each CAV  $v_i$  reasons to determine whether a peer CAV  $v_j$  provides relevant observations and selects a smaller subset of CAVs to initiate textual communication with. In stage 3, the ego CAV then evaluates the mutual information between its own observation and each selected CAV  $v_j$ ’s observation, measuring how relevant the shared information is for planning. A higher mutual information value indicates a more useful observation from CAV  $v_j$ , allowing prioritization of information that best enhances the ego vehicle’s awareness. The fused messages, observations, and task specifications then serve as input to a VLM in stage 4, which outputs driving decisions with uncertainty scores to support safe and interpretable planning.

In the running example, CAV 1996 selects to communicate about semantic information with only CAV 2014 and not with CAV 2005, as it has passed beyond the merging area. To arrive at this decision, CAV 1996 reasons on the mutual information between its observation and CAV 2014’s observation.

UNCAP emphasizes bandwidth efficiency, training-free use, uncertainty-aware decisions, and scalability. Examples of all exchanged messages and VLM prompts are shown in Appendix ??.

### 3.3 BARE to SPARE: Selective Communication via Natural Language

For an ego CAV, choosing which CAVs to communicate with is crucial. In general, for  $N > 2$  vehicles, as  $N$  increases, communicating with all CAVs quickly becomes expensive, redundant, and even unsafe (since contributions from irrelevant communicating vehicles can lead to spurious scene understanding). Moreover, pairwise semantic communication, as in LangCoop [14], scales poorly with combinatorial dependence. This suggests the clear need for an autonomous vehicle to reason about identifying a *smaller yet relevant subset of vehicles* to communicate with.

To address this challenge, we develop **bandwidth-aware reduced exchange** (BARE) and **selective process for CAV reasoning** (SPARE). BARE consists of broadcasting an initial, low-bandwidth package with minimal text-based information from the semantic message that every CAV transmits globally. It consists of the transmitting CAVs position  $p_{\text{cav}}$  and heading angle, indicated by their velocity  $\dot{p}_{\text{cav}}$ . In the running example, BARE comprises of CAVs 1996, 2005, and 2014 all communicating their positions and velocities to each other. Following this global broadcast, SPARE enables an ego CAV to efficiently select communication partners using a heuristic scheme based on intuitive geometric reasoning, thereby improving overall bandwidth efficiency and communication relevance. Specifically, the ego vehicle with position  $p_{\text{ego}}$  and goal position  $p_{\text{goal,ego}}$  for the near future receives the BARE packet from every CAV, and decides to communicate with a certain CAV if *all* the following conditions are true:

$$\|p_{\text{ego}} - p_{\text{cav}}\|_2 \leq d, \quad (2)$$

$$(p_{\text{goal,ego}} - p_{\text{cav}})^\top \dot{p}_{\text{cav}} > 0, \text{ and} \quad (3)$$

where  $d > 0$  in Equation (2) represents a “relevant” distance parameter, that can be set according to the driving environment (e.g., a block distance in urban traffic). Equations (2) and (3) ensure the ego CAV communicates only with the CAVs which (i) are close enough to it and (ii) are heading towards the ego CAV’s goal (and have observations relevant to the ego CAV).

In the running example, these rules will ensure that 1996 selects to communicate only with 2014, and 2005 decides not to communicate with either 1996 or 2014, as neither offers it relevant information. These simple geometric arguments allow a CAV to use BARE and SPARE to narrow down the set of relevant CAVs to continue communicating with. Pseudocode for this stage is presented in Algorithm 1.

Once CAV observations have been selectively communicated, we generate a bird’s eye view (BEV) image that captures fused observations with uncertainty quantification and mutual information. An example BEV is presented in Figure 4.

### 3.4 Data Fusion via Uncertainty Quantification and Mutual Information

To determine the value of communication, we introduce metrics grounded in uncertainty quantification and information theory. In UNCAP, communication is grounded in perception uncertainty and quantified at the object level. This allows agents to exchange only those object detections that reliably reduce uncertainty in the ego vehicle’s scene understanding.

**Algorithm 1** Selecting CAVs to communicate with (BARE & SPARE)

---

**Input:** Ego CAV current and goal positions  $p_{\text{ego}}, p_{\text{goal,ego}}$ , set of BARE packets for other CAVs with positions and headings  $\mathcal{D}_{\text{BARE}} = \{i, p_{\text{cav}}^i, \hat{p}_{\text{cav}}^i\}_{i \neq \text{ego}}$ , distance threshold  $d > 0$

**Output:** Set  $\mathcal{C}$  of relevant CAVs to communicate with

- 1: Initialize  $\mathcal{C} \leftarrow \emptyset$
- 2: **for**  $i = 1, \dots, |\mathcal{D}_{\text{BARE}}|$  **do** ▷ SPARE process
- 3:   **if**  $p_{\text{cav}}^i, \hat{p}_{\text{cav}}^i$  satisfies Equations (2) and (3) **then**
- 4:      $\mathcal{C} \leftarrow \mathcal{C} \cup \{i\}$  ▷ Add  $i$  to communicating CAV set  $\mathcal{C}$
- 5:   **end if**
- 6: **end for**
- 7: **return**  $\mathcal{C}$

---

**Per-Object Perception Uncertainty.** Each CAV  $v_i$  observes an image  $o_i$ , yielding a set of detected objects  $Y_i = \{y_1, \dots, y_K\}$  and the corresponding raw confidence scores  $\{\hat{p}(y_k | o_i)\}_{k=1}^K$  (i.e., class conditional probabilities) obtained from a vision-based detector, such as YOLOv9 [31] consisting of a projection head applied to VLM embeddings. For each detected object  $y_k \in Y_i$ , the detector produces a *class-wise confidence vector*

$$\hat{p}(y_k | o_i) = [\hat{p}(c_1 | y_k, o_i), \dots, \hat{p}(c_L | y_k, o_i)], \quad (4)$$

where each entry denotes the probability that the object  $y_k$  belongs to class  $c_l$  out of  $L$  possible classes. The top-1 predicted class and its corresponding confidence are given by  $\hat{c}_k = \arg \max_l \hat{p}(c_l | y_k, o_i)$  and  $\hat{p}_{\max}(y_k | o_i) = \max_l \hat{p}(c_l | y_k, o_i)$ , respectively. Using SPARE, these class-wise confidence vectors are transmitted selectively to the ego CAV. These raw confidence scores  $\{\hat{p}(y_k | o_i)\}_{k=1}^K$  are calibrated using conformal prediction [2, 27, 29], providing probabilistic guarantees on these class predictions matching the ground truth labels  $Y_k^*$ . Note, however, that other confidence calibration methods can also be employed, as UNCAP is *agnostic to the specific calibration technique*. We use conformal prediction in this work solely for the sake of illustration.

This *perception uncertainty score* associated with the raw confidence scores is defined as

$$u_p(y_k | o_i) = 1 - p(y_k | o_i), \quad k = 1, \dots, K, \quad (5)$$

where  $\{p(y_k | o_i)\}_{k=1}^K$  denotes the calibrated confidence scores. The calibration process yields one calibrated uncertainty score per object detection, representing the CAV's uncertainty about its detection of object  $y_k$ ; higher scores indicate greater uncertainty. We call a confidence score *calibrated* when it reflects the true likelihood of correctness. We now describe the calibration process.

**Conformal Confidence–Uncertainty Calibration.** The observation  $o_i$  is passed through a vision encoder  $V$ , which operates on the entire image to produce a feature embedding. A projection or detection head  $H$  then maps these embeddings to a set of per-object confidence vectors  $c_k = H(V(o_i))_{y_k} = [p(c_1 | y_k, o_i), \dots, p(c_L | y_k, o_i)]$ , where each vector represents the model's softmax confidence scores over the  $L$  object classes for the detected object  $y_k$ . Given a calibration set  $\{(o_i, Y_i^*)\}_{i=1}^N$ , where  $Y_i^*$  is a set of ground truth labels for the set of objects  $Y_i$ , we compute nonconformity scores using  $S_{nc} = \{1 - H(V(o_i))_{y_k}\}_{i=1}^N$ , which quantify how much each prediction deviates from the ground truth. Using  $S_{nc}$ , we empirically

estimate a probability density function  $f_{nc}$  from these nonconformity scores.

Given this distribution and an error tolerance  $\epsilon$  on the probability of correctly classifying an object, we can obtain a corresponding calibrated confidence score  $c^*$ , which according to the theory of conformal prediction, is the  $1 - \epsilon$  quantile of  $f_{nc}$ . Conversely, given  $c^*$ , we can compute the corresponding error bound as  $\epsilon = 1 - \int_0^{c^*} f_{nc}(x) dx$ . Given a new image  $o_{N+1}$  (not in the calibration set), let  $H(V(o_i))_l$  return a confidence vector  $c \in \mathbb{R}^L$ . Using  $f_{nc}$ , we can find a *prediction band* for each object  $y_k$   $\widehat{C}(o_{N+1}) = \{l : H(V(o_{N+1}))_{y_k, l} > 1 - c^*\}$ , such that the probability of the ground-truth label  $y_{N+1}^*$  being contained in  $\widehat{C}(o_{N+1})$  is bounded by  $1 - \epsilon$ . Formally, the conformal prediction coverage guarantee [29] ensures  $P[y_{N+1}^* \in \widehat{C}(o_{N+1})] \geq 1 - \epsilon$ . For downstream planning, we want this prediction band to be a singleton set containing the top-1 class prediction. To enforce this, given a confidence vector  $c$ , we define  $c^* = 1 - \text{sort}(c)_{-2}$ , where  $\text{sort}(c)_{-2}$  denotes the second-highest confidence value. Substituting this into the expression for  $\widehat{C}$  yields

$$\begin{aligned} \widehat{C}(o_{N+1}) &= \{l : H(V(o_{N+1}))_l > 1 - c^*\} \\ &= \{l : H(V(o_{N+1}))_l > \text{sort}(c)_{-2}\} = \{\arg \max_l H(V(o_{N+1}))_l\}. \end{aligned} \quad (6)$$

This results in a singleton prediction set containing only the most confident class.

By conformal prediction, there exists an  $\epsilon \in [0, 1]$  such that

$$P[y_{N+1}^* \in \widehat{C}(o_{N+1})] = P[y_{N+1}^* = \arg \max H(V(o_{N+1}))] \geq 1 - \epsilon.$$

Thus, the calibrated perception uncertainty, which is the lower bound on the probability of incorrectly identifying an object is expressed as

$$u_p = \int_0^{c^*} f_{nc}(x) dx, \quad \text{where } c^* = 1 - \text{sort}(c)_{-2}.$$

This formulation provides a rigorous, probabilistically grounded estimate of perception uncertainty, ensuring that  $u_p$  faithfully reflects the likelihood of correctness.

**Uncertainty Fusion.** When ego CAV  $v_i$  and neighbor CAV  $v_j$  both detect the same object  $y_k$ , we define the fused probability as

$$p(y_k | o_i, o_j) = \max(p(y_k | o_i), p(y_k | o_j)). \quad (7)$$

The resulting fused uncertainty is therefore

$$u_p(y_k | o_i, o_j) = 1 - p(y_k | o_i, o_j) = \min(u_p(y_k | o_i), u_p(y_k | o_j)), \quad (8)$$

which corresponds to adopting the detections from the least uncertain view. This uncertainty-fusion method ensures that joint perception is never worse than the best observer, while avoiding overconfidence from a single uncertain calibrated detection.

**Perception-Based Mutual Information.** The utility of incorporating object  $y_k$  from CAV  $v_j$  is quantified by pointwise mutual information (PMI) [7]:

$$I_p(y_k; o_j | o_i) = \log \frac{p(y_k | o_i, o_j)}{p(y_k | o_i)}. \quad (9)$$

The fused probability  $p(y_k | o_i, o_j)$  is computed using the min-uncertainty rule. Objects with high PMI values correspond to cases where a neighbor's observation reduces ego uncertainty about  $y_k$ . Intuitively, the denominator measures ego confidence in  $y_k$  using only

the ego CAV’s observation, while the numerator reflects confidence when both ego and neighbor observations are available. PMI thus represents the *marginal perception confidence gain* provided by CAV  $j$ . If  $I_p(o_j; \pi_i | o_i) > 0$ , the neighbor contributes positively to ego’s plan confidence; if  $I_p(o_j; \pi_i | o_i) \leq 0$ , the contribution is uninformative or harmful.

Communication is performed at the *object level*: for each detection, ego CAV evaluates its calibrated uncertainty score and computes the PMI contribution. Only detections that both reduce uncertainty and yield positive PMI are selected for fusion. The resulting fused object sets are then passed to VLM for planning.

**Running Example.** In the running example, CAV 2014 calibrates the perception confidence of detected objects and transmits the selected results to CAV 1996 according to SPARE. CAV 1996 then compares its locally observed objects with those reported by CAV 2014. In this example, the two vehicles observe no common objects, reflecting highest PMI. However, if overlap occurs, CAV 1996 computes the confidence boost using Equation (9) to improve planning certainty.

### 3.5 VLM-Based Planning with Uncertainty

Unlike perception models that offer real-time inference, given the higher inference times of existing VLMs, we query them for high-level decisions that are consumed by low-level control for real-time planning. Each VLM output plan is accompanied by a confidence score:

$$u_d(o, \pi) = -\log p_{\text{vlm}}(\pi | o), \quad (10)$$

where  $p_{\text{vlm}} : \mathcal{O} \times \Pi \rightarrow [0, 1]$  is the VLM next token likelihood of an output plan  $\pi \in \Pi$  given input observation  $o \in \mathcal{O}$ .  $o$  may represent the ego vehicle’s camera observation and the shared observation from other CAVs, and a natural language prompt to explain the driving scenario, while  $\pi$  denotes the corresponding high-level plan, e.g., “wait, then proceed after the vehicle on the right”. This negative log-likelihood serves as an uncertainty measure and may be calibrated using conformal prediction [4].

Similar to perception uncertainty, we adopt a formulation inspired by [6] to compute the PMI between the ego CAV  $v_i$ ’s plan  $\pi_i$  and the shared observation from CAV  $v_j$ :

$$I_p(\pi_i; o_j | o_i) = \log \frac{p_{\text{VLM}}(\pi_i | o_i, o_j)}{p_{\text{VLM}}(\pi_i | o_i)}. \quad (11)$$

Here,  $p_{\text{VLM}}(\pi_i | o_i)$  denotes the planning decision based solely on ego’s local observation  $o_i$ , while  $p_{\text{VLM}}(\pi_i | o_i, o_j)$  fuses both ego and neighbor observations.

This formulation ensures *decision safety* since  $I_p(\pi_i; o_j | o_i)$  represents the *marginal decision confidence gain* provided by CAV  $j$ . If  $I_p(\pi_i; o_j | o_i) > 0$ , the additional observation  $o_j$  provides a meaningful confidence gain for the ego CAV’s plan, strengthening the reliability of the decision. If  $I_p(\pi_i; o_j | o_i) \leq 0$ , the neighbor CAV’s information does not increase or possibly even decreases the ego CAV’s confidence, and the neighbor CAV’s information is therefore excluded from fusion. By filtering shared information in this way, we ensure that only uncertainty-reducing and safety-preserving observations influence the ego CAV’s final planning actions.




Figure 3: Illustration of 4 scenes featuring highway merging and intersection handling with No-Comm vs. UNCAP.

## 4 EXPERIMENTAL RESULTS

We now evaluate UNCAP in diverse driving scenarios. Our objectives are to: (i) assess the driving quality resulting from using natural language instead of raw images for communication; (ii) examine the effect of selective communication via BARE and SPARE on bandwidth savings; (iii) evaluate the reduction in perception and planning uncertainty; (iv) analyze safety performance in terms of distance margin in near-miss events; and (v) test the robustness of UNCAP to different VLMs.

### 4.1 Experiment Setup

We leverage the OPV2V dataset [35], which is built upon the CARLA simulator [12]. OPV2V serves as a challenging benchmark specifically curated to evaluate cooperative perception and communication strategies among CAVs. It features diverse, realistic driving scenarios, making it well-suited for testing CAV coordination frameworks such as ours. We evaluate UNCAP on several scenarios consisting of 4 diverse challenging layouts from the OPV2V dataset [35]. The scenarios range from highway merging and intersection turning to urban driving with occluded vehicles. For low-level perception, we use YOLOv9 [31] in all experiments. For high-level planning, we use GPT-4o [24]. We also test UNCAP using different VLMs. For selective communication in SPARE, we set the safety distance threshold  $d$  as 50 meters. Unless noted otherwise, all methods share identical detection outputs and route definitions to isolate the effect of communication. Perception and communication runs at 10 Hz.

### 4.2 Evaluation Metrics

We evaluate UNCAP using multiple metrics. The *Driving Score* (DS) is a normalized composite of progress, rule compliance, and comfort, with higher values indicating better driving quality. *Route Completion* (RC) measures the percentage of the route completed without failure. The *Infraction Penalty* (IP), starts at 1.0 and decreases with each traffic violation, down to 0. DS, RC, and IP are standard CARLA metrics widely used in literature [28]. *Total Bandwidth* (TB) records the total communication volume (KB) per episode, where lower values indicate higher efficiency. *Information Gain* (IG) captures the average increase in ego plan confidence (PMI) given received messages. We also analyze inter-vehicle distance margin in near-miss events to evaluate plan safety.




Figure 4: Data fusion illustration. Vehicle 2530 occludes 2541, causing high ego perception uncertainty, which is reduced by fusing the helper CAV’s low uncertainty observation.




Figure 5: Inter-vehicular distance in the near-miss scenario. Communication enables CAVs to stop within a safe distance.

### 4.3 Baselines

We compare UNCAP against several baselines and present an ablation of all key stages in UNCAP. No-Comm refers to using only local sensing without any inter-vehicle communication. UNCAP w/o SPARE & Fusion, refers to broadcasting all observations (i.e. without selective communication) without data fusion, and represents existing language-based cooperative planning approaches such as LangCoop [14]. UNCAP w/o SPARE uses all messages for data fusion but no selective communication. UNCAP w/ Images refers to sharing raw visual inputs in addition to natural language. UNCAP represents the full framework with BARE, SPARE, and data fusion.

### 4.4 Qualitative Results

Figure 3 compares highway merging and intersection scenarios. Without communication (No-Comm), ego vehicles rely on local perception, often leading to unsafe maneuvers or collisions due to occlusions. In contrast, UNCAP enables proactive decision-making: helper CAVs send semantic, uncertainty-calibrated messages via SPARE, allowing the ego vehicle to defer merges for fast traffic or enter intersections safely. As shown in the BEV visualizations (Figure 4), UNCAP uses mutual information-driven fusion to reduce uncertainty for occluded objects (e.g., vehicle 2530) by leveraging high-confidence detections from selected helpers (e.g., CAV 2506). By focusing communication on agents directly influencing the ego plan, UNCAP complements VLM reasoning and significantly reduces overall planning uncertainty.

Across all qualitative cases, UNCAP consistently exhibits: (1) **Selective awareness**: only uncertainty-reducing detections are fused;




Figure 6: Control signals for the near-miss scenario.

Table 1: Performance metrics for different communication strategies. UNCAP outperforms all baselines across all metrics.

Method	DS ↑	RC ↑	IP ↑	TB (KB) ↓
No-Comm	48.9%	39.7%	62%	–
UNCAP w/o SPARE & Fusion (represents LangCoop [14])	52.4%	79.6%	65%	89
UNCAP w/o SPARE	78.8%	87.2%	90%	89
UNCAP w/ Images	69.5%	88.3%	78%	33,600
UNCAP	<b>80.3%</b>	<b>89.2%</b>	<b>90%</b>	<b>33</b>

(2) **Proactive planning**: ego CAVs anticipate conflicts several frames earlier and initiates proactive breaking, maintaining a higher inter-vehicle distance margin. We illustrate these attributes via Figure 5 and 6; and (3) **Interpretable cooperation**: messages remain human-readable and auditable. Overall, these results illustrate that uncertainty-guided natural language communication yields interpretable, bandwidth-efficient, and safer cooperative behavior compared to unstructured or purely visual exchanges.

### 4.5 Quantitative Results

**UNCAP Enhances Driving Quality Compared to Baselines.** Across diverse scenarios (merges, intersections, occlusions), UNCAP (BARE+SPARE) consistently yields superior decision quality. As shown in Table 1, UNCAP achieves the highest driving score, exceeding No-Comm by 31.4% and naive broadcasting (i.e., LangCoop [14]) by 27.9%. Route completion rates show a similar trend, with UNCAP outperforming all baselines, and it attains the best IP score, indicating the fewest traffic violations. Notably, the use of image observations in communication degrades planning performance, likely because VLMs struggle to integrate multiple visual inputs, often confusing scene content across views. Bearing this in mind, UNCAP avoids this pitfall by leverages language-only communication, translating one image per CAV into text locally, which can be seen as a reasoning step. We present further results for UNCAP on a total of 9 scenarios from both OPV2V and a custom dataset, which have 2 to 4 CAVs in each scenario in Appendix ??.

**SPARE in UNCAP Boosts Bandwidth Efficiency Compared to Baselines.** A key motivation for language-based communication methods such as LangCoop [14] is their substantial bandwidth savings compared to approaches that exchange images or raw sensor data. As shown in Table 1 and Figure 7, the BARE and SPARE processes enable UNCAP to achieve state-of-the-art bandwidth efficiency,

**Table 2: Model confidence and information gain across communication methods. IG scores are presented on log scale.**

Method	Perception		Decision	
	Confidence	IG	Confidence	IG
No-Comm	0.25	—	0.43	—
UNCAP w/o SPARE	0.71	1.04	0.44	0.02
UNCAP w/ Images	0.71	1.04	0.48	0.11
UNCAP	0.71	1.04	<b>0.78</b>	<b>0.60</b>

**Table 3: Comparison of different VLMs across driving and safety metrics, showing UNCAP’s robustness to model choice. We omit IG for GPT-5 as the API does not provide access to token probabilities.**

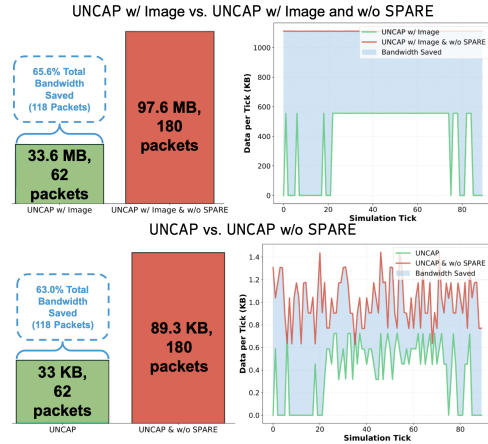
Model	DS↑	RC↑	IP↑	IG↑
GPT-4o-mini	64.7%	88.2%	73%	0.42
GPT-4o	<b>80.3%</b>	<b>89.2%</b>	<b>90%</b>	<b>0.60</b>
GPT-4.1	<b>80.3%</b>	<b>89.2%</b>	<b>90%</b>	0.40
GPT-5	74.5%	<b>89.2%</b>	83%	—

requiring only 37% of the bandwidth used by LangCoop (equivalent to UNCAP without SPARE and fusion) when using textual communication alone, and only 34% when using both text and image communication. It underscores the importance of selective communication done in UNCAP and reveals the poor scalability of conventional frameworks that broadcast all information indiscriminately.

**UNCAP Enhances Safety via Uncertainty Quantification and Mutual Information.** Per Equations (9) and (11), PMI measures IG. A positive IG indicates boosted confidence from shared data, while negative values suggest noise. UNCAP achieves a significant confidence boost of 0.60. In perception, UNCAP matches the IG of full broadcasting because our selective fusion identifies the most informative objects without redundant data. For planning, broadcasting without selection introduces noise from irrelevant CAVs, resulting in marginal IG gains over no communication. Furthermore, image-based exchange drops planning IG to 0.11, confirming that multi-view images are less scalable for VLMs. In contrast, UNCAP’s selective language-based communication reaches an IG of 0.60, demonstrating superior decision-making confidence and safety.

**UNCAP Enables Real-Time Planning.** A major concern for VLM-based planning is latency, and we verify that UNCAP makes real-time planning feasible. In UNCAP, VLM is queried only at critical points, when intentions change, and actions must be executed; it avoids the high latency that querying the VLM at every tick would incur. Communication with KB-scale messages achieves ~ 200ms transmission time, with an estimated broadcast speed of 1.05Mbps, a groupcast speed of 1.52Mbps [1], and 10ms overhead [8]. In contrast, communicating with images incurs higher latency, increasing total transmission time by 1000× to around 200s, even with selective communication. UNCAP achieves an average VLM planning latency of 1.33s, with an overall latency of ~ 1.5s per decision step.

**UNCAP is Robust to the Underlying VLM.** Table 3 evaluates UNCAP across various VLMs. The framework maintains consistent performance regardless of the specific model, with larger VLMs



**Figure 7: SPARE enables substantial bandwidth savings in UNCAP, whether the communication uses images (top) or text (bottom). (Top) Since we selectively communicate with other CAVs, the bandwidth fluctuates per tick. (Bottom) The bandwidth fluctuates as the text messages vary in length.**

like GPT-4o achieving peak scores and smaller models like GPT-4o-mini remaining highly competitive. While Information Gain (IG) varies based on model-specific confidence scores, the consistently positive trend confirms that our communication and fusion framework—rather than a specific VLM—is the primary driver of performance. This highlights the adaptability and robustness of UNCAP to different VLM architectures.

## 5 CONCLUSION

We propose UNCAP, a vision–language model-based planning approach for CAVs using natural language communication. Unlike traditional methods that broadcast raw data or unfiltered text, UNCAP employs a two-stage process: (i) identifying relevant CAVs and (ii) exchanging messages containing explicit quantitative perception uncertainty. By fusing only the most informative messages, UNCAP achieves state-of-the-art driving performance and bandwidth efficiency. Our results demonstrate UNCAP’s potential for reliable, scalable planning in complex environments.

**Future Work.** We plan to extend UNCAP to mixed-traffic scenarios involving human-driven vehicles and natural language instructions. Additionally, we will investigate the impact of realistic network conditions (e.g., packet loss) and evaluate performance on physical CAV fleets to account for real-world sensor noise and vehicle dynamics.

## ACKNOWLEDGMENTS

This material is based upon work supported in part by the Lockheed Martin Corporation under grant No. MRA16-005-RPP023, the Office of Naval Research (ONR) under Grant No. N00014-22-1-2254, the Defense Advanced Research Projects Agency (DARPA) contract FA8750-23-C-1018, DARPA ANSR: RTXCW2231110, and the Army Research Laboratory under Cooperative Agreement Number W911NF-25-2-0021, W911NF-24-1-0220, and W911NF-23-2-0011.

## REFERENCES

- [1] Sanghoon An and Kyunghi Chang. 2023. Enhancing reliability in 5G NR V2V communications through priority-based groupcasting and IR-HARQ. *IEEE Access* 11 (2023), 72717–72731.
- [2] Anastasios N. Angelopoulos and Stephen Bates. 2022. A Gentle Introduction to Conformal Prediction and Distribution-Free Uncertainty Quantification. arXiv:2107.07511 [cs.LG] <https://arxiv.org/abs/2107.07511>
- [3] Eduardo Arnold, Mehrdad Dianati, Robert De Temple, and Saber Fallah. 2020. Cooperative perception for 3D object detection in driving scenarios using infrastructure sensors. *IEEE Transactions on Intelligent Transportation Systems* 23, 3 (2020), 1852–1864.
- [4] Neel P. Bhatt, Yunhao Yang, Rohan Siva, Daniel Milan, Ufuk Topcu, and Zhangyang Wang. 2025. Know Where You’re Uncertain When Planning with Multimodal Foundation Models: A Formal Framework. arXiv:2411.01639 [cs.RO] <https://arxiv.org/abs/2411.01639>
- [5] Qi Chen, Sihai Tang, Qing Yang, and Song Fu. 2019. Cooper: Cooperative Perception for Connected Autonomous Vehicles Based on 3D Point Clouds. In *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, Dallas, USA, 514–524. <https://doi.org/10.1109/ICDCS.2019.00058>
- [6] Shenghui Chen, Po-han Li, Sandeep Chinchali, and Ufuk Topcu. 2025. VIBE: Video-to-Text Information Bottleneck Evaluation for TL; DR.
- [7] Kenneth Church and Patrick Hanks. 1990. Word association norms, mutual information, and lexicography. *Computational linguistics* 16, 1 (1990), 22–29.
- [8] Baldomero Coll-Perales, M Carmen Lucas-Estañ, Takayuki Shimizu, Javier Gozalvez, Takamasa Higuchi, Sergei Avedisov, Onur Altintas, and Miguel Sepulcre. 2022. End-to-end V2X latency modeling and analysis in 5G networks. *IEEE Transactions on Vehicular Technology* 72, 4 (2022), 5094–5109.
- [9] Can Cui, Yunsheng Ma, Zichong Yang, Yupeng Zhou, Peiran Liu, Juanwu Lu, Lingxi Li, Yaobin Chen, Jitesh H. Panchal, Amr Abdelraouf, Rohit Gupta, Kyungtae Han, and Ziran Wang. 2025. Large Language Models for Autonomous Driving (LLM4AD): Concept, Benchmark, Experiments, and Challenges. arXiv:2410.15281 [cs.RO] <https://arxiv.org/abs/2410.15281>
- [10] Jiaxun Cui, Hang Qiu, Dian Chen, Peter Stone, and Yuke Zhu. 2022. Coopernaut: End-to-end driving with cooperative perception for networked vehicles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, New Orleans, USA, 17252–17262.
- [11] Jiaxun Cui, Chen Tang, Jarrett Holtz, Janice Nguyen, Alessandro G. Allievi, Hang Qiu, and Peter Stone. 2025. Towards Natural Language Communication for Cooperative Autonomous Driving via Self-Play. arXiv:2505.18334 [cs.RO] <https://arxiv.org/abs/2505.18334>
- [12] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. 2017. CARLA: An Open Urban Driving Simulator. In *Proceedings of the 1st Annual Conference on Robot Learning (Proceedings of Machine Learning Research, Vol. 78)*, Sergey Levine, Vincent Vanhoucke, and Ken Goldberg (Eds.). PMLR, Mountain View, USA, 1–16. <https://proceedings.mlr.press/v78/dosovitskiy17a.html>
- [13] Hongbo Gao, Bo Cheng, Jianqiang Wang, Keqiang Li, Jianhui Zhao, and Deyi Li. 2018. Object classification using CNN-based fusion of vision and LIDAR in autonomous vehicle environment. *IEEE Transactions on Industrial Informatics* 14, 9 (2018), 4224–4231.
- [14] Xiangbo Gao, Yuheng Wu, Rujia Wang, Chenxi Liu, Yang Zhou, and Zhengzhong Tu. 2025. LangCoop: Collaborative Driving with Language. arXiv:2504.13406 [cs.RO] <https://arxiv.org/abs/2504.13406>
- [15] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. 2017. On calibration of modern neural networks. In *International conference on machine learning*. PMLR, PMLR, Sydney, Australia, 1321–1330.
- [16] Songyang Han, Shanglin Zhou, Lynn Pepin, Jiangwei Wang, Caiwen Ding, and Fei Miao. 2023. Shared Information-Based Safe And Efficient Behavior Planning For Connected Autonomous Vehicles.
- [17] Po han Li, Yunhao Yang, Mohammad Omama, Sandeep Chinchali, and Ufuk Topcu. 2024. Any2Any: Incomplete Multimodal Retrieval with Conformal Prediction. arXiv:2411.10513 [cs.CV] <https://arxiv.org/abs/2411.10513>
- [18] Senkang Hu, Zhengru Fang, Zihan Fang, Yiqin Deng, Xianhao Chen, and Yuguang Fang. 2024. Agentscodriver: Large language model empowered collaborative driving with lifelong learning.
- [19] Meelis Kull, Miquel Perello-Nieto, Markus Kängsepp, Telmo Silva Filho, Hao Song, and Peter Flach. 2019. *Beyond temperature scaling: obtaining well-calibrated multiclass probabilities with Dirichlet calibration*. Curran Associates Inc., Red Hook, NY, USA, Chapter 1, 1–11.
- [20] Changxing Liu, Genjia Liu, Zijun Wang, Jinchang Yang, and Siheng Chen. 2025. CoLMDriver: LLM-based Negotiation Benefits Cooperative Autonomous Driving.
- [21] Jiageng Mao, Yuxi Qian, Hang Zhao, and Yue Wang. 2023. GPT-Driver: Learning to Drive with GPT.
- [22] Jiageng Mao, Junjie Ye, Yuxi Qian, Marco Pavone, and Yue Wang. 2024. A Language Agent for Autonomous Driving. <https://openreview.net/forum?id=UPE6WYE8vg>
- [23] Zhenjiang Mao, Mrinal Eashan Umasudhan, and Ivan Ruchkin. 2025. How Safe Will I Be Given What I Saw? Calibrated Prediction of Safety Chances for Image-Controlled Autonomy. arXiv:2508.09346 [cs.RO] <https://arxiv.org/abs/2508.09346>
- [24] OpenAI. 2024. GPT-4o System Card. arXiv:2410.21276 [cs.CL] <https://arxiv.org/abs/2410.21276>
- [25] John Platt et al. 1999. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in large margin classifiers* 10, 3 (1999), 61–74.
- [26] Hao Sha, Yao Mu, Yuxuan Jiang, Li Chen, Chenfeng Xu, Ping Luo, Shengbo Eben Li, Masayoshi Tomizuka, Wei Zhan, and Mingyu Ding. 2025. LanguageMPC: Large Language Models as Decision Makers for Autonomous Driving. arXiv:2310.03026 [cs.RO] <https://arxiv.org/abs/2310.03026>
- [27] Glenn Shafer and Vladimir Vovk. 2008. A tutorial on conformal prediction. *Journal of Machine Learning Research* 9, 3 (2008), 371–421.
- [28] CARLA Team. 2025. Evaluation Criteria for the Leaderboard 2.0. [https://leaderboard.carla.org/evaluation\\_v2\\_0/](https://leaderboard.carla.org/evaluation_v2_0/). [Accessed 09-10-2025].
- [29] Vladimir Vovk, Alexander Gammerman, and Glenn Shafer. 2005. *Algorithmic learning in a random world*. Springer, Heidelberg, Germany.
- [30] Binglu Wang, Lei Zhang, Zhaozhong Wang, Yongqiang Zhao, and Tianfei Zhou. 2023. Core: Cooperative reconstruction for multi-agent perception. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. IEEE, Vancouver, Canada, 8710–8720.
- [31] Chien-Yao Wang and Hong-Yuan Mark Liao. 2024. YOLOv9: Learning What You Want to Learn Using Programmable Gradient Information.
- [32] Tsun-Hsuan Wang, Sivabalan Manivasagam, Ming Liang, Bin Yang, Wenyuan Zeng, and Raquel Urtasun. 2020. V2vnet: Vehicle-to-vehicle communication for joint perception and prediction. In *European conference on computer vision*. Springer, Springer, Glasgow, United Kingdom, 605–621.
- [33] Licheng Wen, Daocheng Fu, Xin Li, Xinyu Cai, Tao MA, Pinlong Cai, Min Dou, Botian Shi, Liang He, and Yu Qiao. 2024. DiLu: A Knowledge-Driven Approach to Autonomous Driving with Large Language Models. In *International Conference on Learning Representations*. PMLR, Vienna, Austria, 0. <https://openreview.net/forum?id=OqTMUPuLuC>
- [34] Runsheng Xu, Hao Xiang, Zhengzhong Tu, Xin Xia, Ming-Hsuan Yang, and Jiaqi Ma. 2022. V2x-vit: Vehicle-to-everything cooperative perception with vision transformer. In *European conference on computer vision*. Springer, Springer, Tel Aviv, Israel, 107–124.
- [35] Runsheng Xu, Hao Xiang, Xin Xia, Xu Han, Jinlong Li, and Jiaqi Ma. 2022. OPV2V: An Open Benchmark Dataset and Fusion Pipeline for Perception with Vehicle-to-Vehicle Communication. In *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, Philadelphia, PA, 2583–2589. <https://doi.org/10.1109/ICRA46639.2022.9812038>