

# Logic-RAG: Augmenting Large Multimodal Models with Visual-Spatial Knowledge for Road Scene Understanding

Imran Kabir<sup>1</sup>, Md Alimoor Reza<sup>2</sup>, and Syed Billah<sup>1</sup>

**Abstract**—Large multimodal models (LMMs) are increasingly integrated into autonomous driving systems for user interaction. However, their limitations in fine-grained spatial reasoning pose challenges for system interpretability and user trust. We introduce Logic-RAG, a novel Retrieval-Augmented Generation (RAG) framework that improves LMMs’ spatial understanding in driving scenarios. Logic-RAG constructs a dynamic knowledge base (KB) about object-object relationships in first-order logic (FOL) using a perception module, a query-to-logic embedder, and a logical inference engine. We evaluated Logic-RAG on visual-spatial queries using both synthetic and real-world driving videos. When using popular LMMs (GPT-4V, Claude 3.5) as proxies for an autonomous driving system, these models achieved only 55% accuracy on synthetic driving scenes and under 75% on real-world driving scenes. Augmenting them with Logic-RAG increased their accuracies to over 80% and 90%, respectively. An ablation study showed that even without logical inference, the fact-based context constructed by Logic-RAG alone improved accuracy by 15%. Logic-RAG is extensible: it allows seamless replacement of individual components with improved versions and enables domain experts to compose new knowledge in both FOL and natural language. In sum, Logic-RAG addresses critical spatial reasoning deficiencies in LMMs for autonomous driving applications. Code and data are available at: <https://github.com/Imran2205/LogicRAG>.

## I. INTRODUCTION

End-to-end AI systems such as autonomous vehicles are becoming mainstream, yet their black-box nature makes their decisions difficult to interpret [1]. These systems typically communicate through either visualization interfaces or natural language interactions powered by large multimodal models (LMMs) [2], [3], [4]. While visualizations show system state, they offer limited user interaction. LMM interfaces, in contrast, enable users to converse about and reason through system decisions, potentially improving interpretability [5].

However, LMMs struggle with fine-grained visual-spatial reasoning (VSR) – they fail to accurately represent relationships between objects (e.g., in front of, behind) [6], [7], confuse nuanced action relationships (e.g., distinguishing between *vehicle01* chased *vehicle02* vs. *vehicle02* chased *vehicle01*) [8], and exhibit weak logical reasoning [9]. Moreover, LMMs can generate content that they do not fully understand [10] and are prone to hallucination [11]. Fig. 2 demonstrates these limitations, showing how popular LMMs (GPT-4V [12] and Claude 3.5 [13]) struggle with even simple VSR questions like “Does the white car

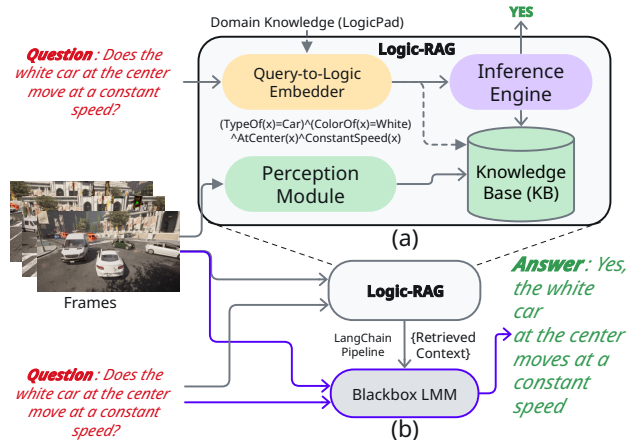


Fig. 1. (a) Components of the Logic-RAG framework. It takes  $N$  frames and visual-spatial reasoning questions as inputs. The perception module analyzes the frames to generate facts about properties and visual-spatial relationships of objects in the scene and constructs a knowledge base (KB) in the form of First-Order Logic (FOL). The Query-to-Logic embedder parses the natural language question into an FOL query predicate, which is then passed to the Inference Engine that performs the query resolution. (b) The integration of Logic-RAG into a black box LMM, which receives the inference output of our framework while generating the response.

at the center move at a constant speed?”. Such deficiencies hinder the integration of LMM interfaces into autonomous driving systems and undermine users’ trust.

Some of these limitations can be addressed by retrieval-augmented generation (RAG) techniques [14], which provide LMMs with relevant information retrieved from vector storage as “context” to reduce factual errors in their responses.

Building on this concept, we present **LOGIC-RAG**, a novel RAG framework that constructs a knowledge base (KB) of facts about road scene objects and their relationships using first-order logic (FOL). It comprises four key components (Fig. 1.a): 1) a Perception module, 2) a Query-to-Logic Embedder, 3) the Knowledge Base (KB), and 4) an Inference Engine that performs query resolution by applying inference rules to derive new facts until a contradiction is found or the desired query is proven. The facts retrieved from the inference are fed into LMMs for augmentation (Fig. 1.b), which enhances their spatial reasoning abilities in dynamic scenes while generating responses. Like any expert system, human experts provide a small set of declarative rules and facts in FOL to bootstrap the framework [15].

We evaluated Logic-RAG on visual-spatial relationship questions using both synthetic (Fig. 2) and real-world (KITTI [16]) driving videos. We used popular commercial LMMs as proxies for black-box autonomous driving systems. On synthetic data, our results show that while these LMMs

<sup>1</sup>College of Information Sciences and Technology, Pennsylvania State University, State College, PA 16801, USA {ibk5106, skb5969}@psu.edu

<sup>2</sup>Department of Mathematics and Computer Science, Drake University, Des Moines, IA 50311, USA md.reza@drake.edu



**Prompt:** Here is a sequence of  $N$  frames extracted from a video at  $M$  fps. The frames are increasing in time and thus contain temporal information. The camera is positioned at the center in first-person view. Answer the following questions:

**Question 1:** Does the white car at the center move at a constant speed?

**Human:** Yes.  
**Our Framework:** Yes.  
**Statement:**  
 $ConstantSpeed(vehicle01)?$

**GPT-4V:** The white car at the center does not move at a constant speed (No).

**Claude-3.5:** The white car at the center remains stationary throughout the sequence (No).

**Question 2:** Does a white van appear from the left of the scene?

**Human:** Yes.  
**Our Framework:** Yes.  
**Statement:**  $AtLeft(vehicle02)$   
and  $Appears(vehicle02)?$

**GPT-4V:** A white van appears from the left of the scene (Yes).

**Claude-3.5:** Yes, a white van enters the scene from the left.

**Question 3:** Does the white van come close to the white car at the center?

**Human:** Yes.  
**Our Framework:** Yes.  
**Statement:**  
 $GettingCloser(vehicle0,vehicle02)?$

**GPT-4V:** The white van does not come close to the white car at the center (No).

**Claude-3.5:** No, the white van does not come particularly close to the stationary white car at the center.

**Question 4:** Does a pedestrian near the left increase pace and disappear from the scene?

**Human:** Yes.  
**Our Framework:** Yes.  
**Statement:**  $IncreasePace(pedestrian01)$   
and  $AtLeft(pedestrian01)$   
and  $CloseToCamera(pedestrian01)?$

**GPT-4V:** The pedestrian near the left does appear to increase pace and then disappears from the scene (Yes).

**Claude-3.5:** No, there does not appear to be a pedestrian near the left that increases pace and disappears from the scene.

Fig. 2. An illustration of the limitations of current commercial LMM models in visual-spatial reasoning (VSR). (Top) A sample of 4 consecutive frames from a synthetic driving video where the VSR task will be performed. (Middle) The original prompts and four representative questions. The values of  $N$  and  $M$  are 10 for GPT-4V and 5 for Claude-3.5, respectively, due to their current limits. (Bottom) Four responses from: 1) human oracle (accuracy: 100%), 2) Logic-RAG, 3) GPT-4V, and 4) Claude-3.5. The correct answers are colored green, and the incorrect answers are colored red. Note that Logic-RAG’s accuracy is higher than black-box LMMs, and its response is generated by logical inference in our FOL system.

achieve an average accuracy of around 55% (baseline), augmenting them with Logic-RAG increases accuracy by more than 25%, exceeding 80%. An ablation study revealed that augmenting LMMs with facts represented as template sentences (without the logical inference module) still improved accuracy by over 10% compared to the baseline condition. On real-world data, Logic-RAG achieved 91% accuracy, outperforming the baseline LMMs by 17%.

These findings strongly suggest that when LMMs learn from joint embedding of image-text pairs [17], they may have bypassed traditional computer vision pipelines (e.g., semantic segmentation, depth estimation, multi-object tracking), leading to deficiencies in spatial reasoning and object-object relationship understanding [18]. By recovering this missing information and augmenting LMMs with it, our framework significantly improves their accuracy in visual-spatial reasoning tasks.

In sum, we make the following **contributions**:

- **Logic-RAG:** We introduce the Logic-RAG framework, which encodes detailed scene facts into a dynamic knowledge base using First-Order Logic.
- **Extensibility and Adaptability:** Our framework allows the incorporation of new domain knowledge from human experts, using existing predicates to tailor the framework to specific tasks and utility.

## II. BACKGROUND AND RELATED WORK

### A. Logical Reasoning using First-Order Logic

First-order logic (FOL) is an expressive formal system for representing and reasoning about objects, their properties, and relationships. FOL has been fundamental in AI for knowledge representation and inference [19]. Neuro-symbolic approaches like DeepProbLog [20] and Neural Theorem Provers [21]

combine deep learning with symbolic reasoning. LOGIC-LM [9] demonstrated the integration of large language models with symbolic solvers for enhanced logical reasoning.

a) *Components of FOL:* FOL consists of *predicates*, *functions*, *constants*, *variables*, and *quantifiers*. Predicates describe properties of objects or relationships between objects. They take objects as arguments and produce a proposition that is either *true* or *false*. Functions, like predicates, can take any number of arguments but always return a single value rather than a proposition. For instance,  $ColorOf(vehicle01)$  is a function that returns the color of the vehicle represented by  $vehicle01$ . Constants represent specific objects, while variables are placeholders. Quantifiers, such as the universal ( $\forall$ ) and existential ( $\exists$ ) quantifiers, enable reasoning about multiple objects. For example, if  $Vehicle(vehicle01)$  is a specific segment of an object type vehicle in a scene and  $Vehicle(x)$  is a base predicate indicating that variable  $x$  is a vehicle, then  $vehicle01$  is a constant in FOL.

b) *FOL Statements and Operators:* FOL statements are constructed by combining predicates, logical connectives (e.g., conjunction  $\wedge$ , disjunction  $\vee$ , negation  $\neg$ , equality  $=$ , and implication  $\rightarrow$ ), and quantifiers. For example, the statement  $\forall x, y : (DistanceDecreases(x, y)) \rightarrow GettingCloser(x, y)$  indicates that for all objects  $x$  and  $y$ , if the distance between objects  $x$  and  $y$  decreases in a set of consecutive video frames, then  $x$  and  $y$  are getting closer. The equality ( $=$ ) operator determines whether two objects or object properties are identical. For example, if  $vehicle01$  and  $vehicle02$  are constants representing two black vehicles, the statement  $ColorOf(vehicle01) = ColorOf(vehicle02)$  is *true*.

### B. Retrieval Augmented Generation (RAG)

Recent advances in RAG have reduced the hallucination problem of LMMs by combining the strengths of language

```

Vehicle(vehicle01)
Vehicle(vehicle02)
AtLeft(vehicle02)
Appears(vehicle02)
DistanceDecreases(vehicle01, vehicle02)
Pedestrian(pedestrian01)
CloseToCamera(pedestrian01)
SpeedUp(pedestrian01)
Disappears(pedestrian01)
...

```

---

```

ConstantSpeed(vehicle01)?
AtLeft(vehicle02)?
Appears(vehicle02)?
GettingCloser(vehicle01, vehicle02)?
IncreasePace(pedestrian01)?

```

Fig. 3. The predicates (in black) show a portion of KB that we construct using the output of the perception module for frames in Fig. 2. Text in orange shows the queries that we make in the KB to learn the facts. For instance, for the frames in Fig. 2, *ConstantSpeed(vehicle01)* is `true`.

models with information retrieval systems [14]. This allows the model to generate content based on accurate and up-to-date information [11]. RAG works as follows:  $y = LLM(q, R(q))$ , where  $y$  is the output generated by the  $LLM(\cdot, \cdot)$ ,  $q$  is the input (in this case, the question or query), and  $R(q)$  is the relevant information retrieved from an external KB. In Logic-RAG, we employ RAG using a formal method. Our KB is composed of FOL facts, and  $R(q)$  is the result of the resolution of the query  $q$ .

### III. METHOD

We now describe the four key components of Logic-RAG (Fig. 1): the *Knowledge Base* (KB), the *Perception* module, the *Query-to-Logic Embedder*, and the *Inference Engine*. Logic-RAG can be integrated with any LMM via LangChain (www.langchain.com) interface, a popular RAG pipeline.

#### A. Knowledge Base (KB)

Our Knowledge Base (KB) is comprised of base predicates, functions, and statements that represent object properties, spatial relationships, and events in driving contexts (Table I). The KB operates in a sliding window of  $N$  consecutive video frames ( $N = 10$ ), where facts are generated by instantiating predicates with specific instances. For example, substituting *vehicle01* for  $x$  in *Vehicle*( $x$ ) produces the fact *Vehicle*(*vehicle01*). The KB updates dynamically as the sliding window progresses, with our multistage perception module (§III-B) continuously recognizing and incorporating new facts (see rows 2<sup>nd</sup> and 3<sup>rd</sup> in Table I). Fig. 3 shows some facts generated for the frames in Fig. 2.

Importantly, our KB is extensible; it allows domain experts to compose predicates through *LogicPad*, a text file written in YAML format that our system dynamically loads and interprets. The capabilities of the perception module are exposed in LogicPad as atomic predicates and functions, similar to library methods in traditional programming languages. Experts create new compound predicates (e.g., *Collide*( $x, y$ )) by combining atomic predicates and functions in FOL (see the *Derived* rows in Table I). Alternatively, they can create compound predicates using natural language that our Query-to-Logic Embedder (§III-C) translates into FOL. As new capabilities are added in the perception module—either by swapping internal models




Fig. 4. Internal block diagram of our perception module. It takes video frames as input and generates semantic, depth, and optical flow maps, which are then utilized to track object instances and estimate relative distances.

or adding new ones—these capabilities automatically appear as new atomic predicates (or functions) in LogicPad. All predicates are available in the project [repository](#).

#### B. Perception Module

Given an input driving video, we first generate dense prediction tasks for each video frame, which include (i) semantic segmentation, (ii) depth estimation, and (iii) optical flow prediction (see Fig. 4). We then track each semantic object using an object tracking model. Simultaneously, we compute the relative distance between objects within two adjacent frames inside the sliding window. These distances enable spatial reasoning among objects in the scene, which we subsequently use in our FOL predicates, such as *GettingCloser*( $x, y$ ), as discussed in §II-A.

1) *Dense Prediction Modules*: Let us denote the input video comprising  $M$  image frames. We also define a sliding window of  $N$  consecutive frames, i.e.,  $\mathbf{I} = I_1, I_2, \dots, I_N$ . For each frame  $I_i$ , we need to identify various foreground and background objects as well as their semantic instances. Furthermore, to describe inter-object relationships and the motion dynamics of different objects over multiple frames, we need to identify 3D spatial relationships between these objects, particularly depth and motion flow. We describe these steps in the following sections.

a) *Semantic Segmentation*: We start by densely predicting semantic labels for every pixel in an image frame  $I_i$ . This involves pixel-wise labeling of foreground objects, such as *pedestrian* and *vehicle*, as well as background objects, including *sidewalk*, *vegetation*, and *building*. Let us assume that there are  $L$  semantic categories  $\{c_1, c_2, \dots, c_L\}$  in an image frame  $I_i$ , and each category  $c_j$  can have  $K_i$  instances, represented as  $\{s_{j1}^i, s_{j2}^i, \dots, s_{jK_i}^i\}$ . Each semantic category is an atomic predicate in our framework (e.g., *Pedestrian*( $x$ ), *Vehicle*( $x$ )). Instances of each category can then be encoded as *constants* in FOL. For instance, we denote two instances obtained in the current image frame  $I_i$  as *pedestrian01* and *vehicle01*. These instances serve as *constants* that are bound to *atomic predicates*, in this case, *Pedestrian*(*pedestrian01*) and *Vehicle*(*vehicle01*). The complete list of *atomic predicates* used in our experiments is shown in Table I. For semantic segmentation, we employed HRNet-v2 [22] model, which has demonstrated effective generalization across multiple datasets trained from a composite one [23]. Furthermore, to determine the color (e.g., ‘black’, ‘blue’, or ‘white’) and type (e.g., ‘car’, ‘bus’, or ‘SUV’) of *vehicle01*, the values of the functions *ColorOf*(*vehicle01*) and *TypeOf*(*vehicle01*) are predicted using a learned classifier.



b) *Optical Flow*: To approximate the scene motion between two adjacent image frames in the video, we compute the optical flow for each pixel. These computed optical flows are then utilized to track important semantic instance segments, enabling us to understand the movement of objects within the scene. We employ the state-of-the-art pretrained Rigid Mask model [24], which excels at estimating optical flow and scene motion.

c) *Depth Estimation*: Alongside semantic segmentation and optical flow, we estimate dense depth from an input image. Our depth estimation module provides per-pixel 3D depth values, which are crucial for identifying 3D spatial relationships between various semantic instance segments. These relationships encompass both intra- and inter-segment connections that evolve over time, as discussed in §III-B.3. To estimate the depth of various segments in each frame  $I_i$ , we employ the recent state-of-the-art depth estimation model called PixelFormer [25], which has shown impressive results in predicting depth information from a single image.

2) *Object Tracking for Semantic Instances*: We track the movement trajectory of each instance  $s_{jk}^i$  belonging to all semantic categories across all image frames within our specified window by framing it as a bipartite matching problem. These trajectories are implicitly used to describe single-object and multi-object interactions, which we will describe in detail in §III-B.3. To construct the bipartite graph, we treat each instance as a vertex, and edges are the dense connections between the vertices of two frames. We assign edge weights using two factors: an optical flow-based factor derived from accumulated flow vectors of instance  $s_{jk}^i$  at position  $s_{jk}^{i+1}$  [26], and a bounding box-based factor using the intersection over union (IoU) of boxes encompassing two vertices connected by an edge. We address the challenges from under-segmentation and occlusion by inserting hypothetical vertices as described in [27].

This process provides trajectories of individual instances within the specific sliding window. Based on these trajectories, we encode predicates, such as  $(Appears(x))$  and  $(Disappears(x))$ , to represent the appearance and disappearance of instances. Consider an instance  $(vehicle01)$  of the category  $Vehicle(x)$ . If  $(vehicle01)$  is first detected in the current sliding window, the predicate  $(Appears(vehicle01))$  becomes *true* within that window, indicating that the instance has just appeared. However, if  $(vehicle01)$  was already present at the beginning of the current sliding window,  $(Appears(vehicle01))$  will be false (i.e.,  $\neg Appears(vehicle01)$ ), indicating that it existed prior to the current window.

3) *Inter-Object Relative Distance Estimator*: Using instance trajectories, we also determine the spatial displacement of inter-instance and intra-instance segments between two consecutive frames  $I_i$  and  $I_{i+1}$ . These displacements are used to encode predicates such as  $Moves(x)$ ,  $SpeedUp(x)$ ,  $SpeedDown(x)$ ,  $AtLeft(x)$ ,  $CloseToCamera(x)$ , and  $DistanceIncreases(x, y)$ , which represent various spatial relationships and movements of objects.

To calculate the spatial displacement, we first utilize the estimated depth map to obtain the 3D coordinates of each

pixel in the current frame’s camera coordinate space. Given a pixel with 2D image coordinates  $(x^j, y^j)$ , we compute its corresponding 3D coordinates  $(X^j, Y^j, Z^j)$  using standard perspective projection:  $X^j = \frac{(x^j - c_x)z^j}{f_x}$ ,  $Y^j = \frac{(y^j - c_y)z^j}{f_y}$ , and  $Z^j = z^j$ . Here,  $z^j$  is the depth value for the  $j^{th}$  pixel retrieved from the estimated depth map,  $(c_x, c_y)$  is the optical center of the camera, and  $(f_x, f_y)$  is the focal length of the camera. For our visual-spatial reasoning, we consider the  $X$ -axis and  $Z$ -axis planes (bird’s eye view) by discarding the  $Y$ -axis component along the height. This means that we only use the  $(X, Z)$  coordinates when reasoning about object-object relationships and their spatial displacements.

To estimate the relative displacement of objects across frames, we need to establish a common world coordinate system origin for all frames within the window. Traditional approaches for this estimation include the SLAM [28] or Structure from Motion (SfM) [29] techniques. However, these techniques are computationally expensive. As an alternative, we propose an approximation strategy that leverages the relative stability of static objects compared to dynamic objects in two consecutive frames. Our approach assumes that the locations of static objects (e.g., trees, buildings, and poles) remain nearly constant between two consecutive frames, while dynamic objects (e.g., vehicles, persons) move more significantly. We feed the locations of static objects from two consecutive frames into a *multilateration* [30] algorithm, which yields a temporary origin that remains mostly constant between the two frames. This temporary origin serves as a reference point for estimating the relative displacement of moving objects between consecutive frames. This enables efficient computation of object trajectories and spatial relationships.

We use the displacement calculated from the temporary origin to derive the velocity and acceleration of moving objects. We then utilize these values to establish various predicates in our KB. For example, we evaluate the predicate  $Moves(vehicle01)$  as true if the displacement of the object  $vehicle01$  within a sliding window exceeds 0. Similarly, we evaluate the predicate  $SpeedUp(vehicle01)$  as true if the velocity of  $vehicle01$  at the end of a sliding window exceeds its initial velocity.

### C. Query-to-Logic Embedder

We employ an open-source large language model (e.g., LLaMa 3 [31] or Phi-3 [32]) with few-shot prompting. Our prompt includes: a) Our predefined predicates; b) Pairs of natural language statements and their corresponding FOL predicates; and c) A few pairs of natural language questions and their corresponding FOL queries. This concise yet expressive prompt enables the conversion of natural language queries into FOL predicates in conjunctive normal form (CNF), which we term (CNF) embedding. For example, given the natural language question “Does the white car at the center move at a constant speed?”, our Query-to-Logic Embedder generates the logic:  $(TypeOf(x) = Car) \wedge (ColorOf(x) = White) \wedge AtCenter(x) \wedge ConstantSpeed(x)$  During our evaluation of 645 natural language questions, we

Submodule	Atomic Predicates and Statements
<b>Semantic Segmentation</b>	$Road(x), LaneMarking(x), TrafficSign(x), Sidewalk(x), Fence(x), Pole(x), Wall(x), Building(x), Vegetation(x), Vehicle(x), Pedestrian(x), Other(x)$
<b>Tracker</b>	$Appears(x), Disappears(x), Moves(x), SpeedUp(x), SpeedDown(x), CloseToCamera(x), AtRight(x), AtLeft(x), AtCenter(x), DistanceIncreases(x, y), DistanceDecreases(x, y), DistanceDecreasesToZero(x, y) On(x, y)$
<b>Functions</b>	$ColorOf(x), TypeOf(x)$
	Stopped $\forall x : (\neg Moves(x)) \rightarrow Stopped(x)$
	Walk $\forall x : ((Pedestrian(x) \wedge Moves(x)) \rightarrow Walk(x)$
	Stand $\forall x : ((Pedestrian(x) \wedge \neg Moves(x)) \rightarrow Stand(x)$
<b>Derived</b>	Accelerate $\forall x : (Vehicle(x) \wedge SpeedUp(x)) \rightarrow Accelerate(x)$
	Constant Speed $\forall x : (Vehicle(x) \wedge \neg SpeedUp(x) \wedge \neg SpeedDown(x)) \rightarrow ConstantSpeed(x)$
	Increase Pace $\forall x : ((Pedestrian(x) \wedge SpeedUp(x)) \rightarrow IncreasePace(x)$
	Fixed Pace $\forall x : ((Pedestrian(x) \wedge \neg SpeedUp(x) \wedge \neg SpeedDown(x)) \rightarrow FixedPace(x)$
	Getting closer $\forall x, y : (DistanceDecreases(x, y)) \rightarrow GettingCloser(x, y)$
	Collide $\forall x, y : (DistanceDecreases(x, y) \wedge DistanceDecreasesToZero(x, y)) \rightarrow Collide(x, y)$
...	...

TABLE I

A PARTIAL LIST OF BASE PREDICATES, FUNCTIONS, AND STATEMENTS REPRESENTING OBJECT PROPERTIES, SPATIAL RELATIONSHIPS, AND EVENTS.

observed no errors in the embedder’s output, demonstrating its robustness. The resulting FOL query is then passed to the Inference Engine for processing.

#### D. Inference Engine

We built our inference system with two established tools: *miniKanren* [33], a symbolic computation framework, and *aima-python* [34], which implements FOL representation and inference based on Russell and Norvig’s work [19]. To derive new facts from our KB, we apply standard logical principles such as modus ponens and resolution. For example, to identify objects approaching a specific vehicle, we can query  $GettingCloser(vehicle01, y)$ , which binds variable  $y$  to all objects moving closer to *vehicle01* within the current window. Our inference is sound; it ensures that we derive only valid conclusions from the given premises.

### IV. VISUAL-SPATIAL REASONING EXPERIMENTS

We conducted the experiment by first creating a synthetic dataset with the CARLA [35] simulator, to ensure we have access to ground truth to evaluate our perception module. We then composed questions from this dataset to assess the effectiveness of our framework against two state-of-the-art commercial models, LMM-1: GPT-4V [12] and LMM 2: Claude 3.5 [13]. Lastly, we integrate our framework into these LMMs with and without logical inference and measure performance differences. We also conduct an experiment on KITTI object tracking dataset [16] to evaluate the effectiveness of our framework on real-world scenarios.

#### A. Reasoning on Synthetic Dataset

*a) Evaluation Dataset:* We generated distinct environmental conditions in Carla Simulator by altering various properties, including *Day, Night, Rain, Snow, and Sunset*. We captured 5,256 video frames at 10 fps, divided into 525 sequences of 10 consecutive frames each. For each frame, we extracted the corresponding depth map, semantic segmentation map with 12 categories, and instance segmentation map.

#### *b) Implementation Details of the Perception Module:*

We fine-tuned several models in our perception module (Fig. 4) using a separate training dataset of 9,706 synthetic RGB images with pixel-level annotations for semantic segmentation and depth. For the semantic segmentation submodule, we fine-tuned HRNetv2 [22] using its official PyTorch implementation. We used the pre-trained model of Rigid Mask [36] for the optical flow submodule without fine tuning. For the depth estimation submodule, we fine-tuned PixelFormer [25] using its official implementation. We conducted all training and inference on a single Nvidia RTX A6000 GPU with 48GB memory. Note that input and output interfaces of these models are well-defined so that we can seamlessly replace a model with a better one within the perception module.

*c) Evaluation Criteria:* From the evaluation dataset, we selected 95 sequences of 10 consecutive frames and composed 3-9 visual-spatial reasoning (VSR) questions per sequence. These questions included object attributes, actions, object-object relationships, and events. We ensure question validity by focusing on objects present for at least one-third of the sliding window, referring to them by color and position. We composed all questions manually. All authors collaboratively analyzed each frame sequence to ensure questions were answerable from visual information in the frames, without introducing bias toward any specific model architecture. In total, we prepared 645 such questions.

We manually categorize questions into *Unary* (about one object) and *Binary* (involving two objects). Unary subcategories include *object query (U1)*, *velocity (U2)*, *change in velocity (U3)*, and *object appearance/disappearance (U4)*. The binary subcategories are *relative position (B1)* and *relative distance change (B2)*.

For each frame sequence and question set, we generated responses under three conditions:

C1 Logic-RAG, LMM-1, and LMM-2 in standalone.

C2 LMM-1 and LMM-2 with Logic-RAG via LangChain.

Relationship	Question Type	No. of Questions	Logic-RAG (Oracle) Accuracy/ $F_1$	Accuracy $\uparrow$					$F_1$ $\uparrow$				
				LR/ LMM-1 (w LR)/ LMM-2 (w LR)	LMM-1 (w LR $^\dagger$ )	LMM-1 (w/o LR)	LMM-2 (w LR $^\dagger$ )	LMM-2 (w/o LR)	LR/ LMM-1 (w LR)/ LMM-2 (w LR)	LMM-1 (w LR $^\dagger$ )	LMM-1 (w/o LR)	LMM-2 (w LR $^\dagger$ )	LMM-2 (w/o LR)
Unary	<b>U1</b>	193	1.0/1.0	<b>0.82</b>	0.70	0.78	0.77	0.74	<b>0.90</b>	0.82	0.87	0.87	0.85
	<b>U2</b>	79	1.0/1.0	<b>0.82</b>	0.77	0.43	0.71	0.53	<b>0.90</b>	0.87	0.60	0.83	0.69
	<b>U3</b>	129	1.0/1.0	<b>0.84</b>	0.74	0.35	0.68	0.51	<b>0.87</b>	0.79	0.44	0.75	0.60
	<b>U4</b>	98	1.0/1.0	<b>0.83</b>	0.63	0.19	0.63	0.39	<b>0.90</b>	0.76	0.31	0.76	0.55
Binary	<b>B1</b>	84	1.0/1.0	<b>0.79</b>	0.70	0.77	0.73	0.63	<b>0.88</b>	0.83	0.87	0.84	0.77
	<b>B2</b>	62	1.0/1.0	<b>0.77</b>	0.58	0.42	0.56	0.45	<b>0.87</b>	0.73	0.59	0.72	0.62
Aggregated Result		645	1.0/1.0	<b>0.82</b>	0.69	0.54	0.72	0.57	<b>0.88</b>	0.79	0.68	0.81	0.70

TABLE II

ACCURACY AND  $F_1$  SCORES OF LOGIC-RAG AND BASELINES FOR VARIOUS VSR QUESTION TYPES.  $LR$  REPRESENTS LOGIC-RAG, WHILE  $LR^\dagger$  DENOTES LOGIC-RAG’S KB EXPORTED AS TEMPLATE SENTENCES WITHOUT LOGICAL INFERENCE.

C3 LMM-1 and LMM-2 provided with template sentences of facts from Logic-RAG (ablation, no logical inference).

We manually verified all responses under each condition, compared results with the oracle, and reported accuracy and  $F_1$  scores in Table II.

*d) Results: Condition C1:* When we populate the KB using ground truth for the perception module, Logic-RAG unsurprisingly achieves 100% accuracy. This suggests that our FOL predicates can express and infer the correct response. However, our current implementation of the perception module reaches 82% average accuracy ( $F_1$  score: 0.88). Although it outperforms standalone LMM-1 and LMM-2 in all question categories (overall accuracy: 0.54 and 0.57; overall  $F_1$  score: 0.68 and 0.70, respectively), it indicates room for improvement, which is feasible by replacing current models with newer, better ones.

*e) Results: Condition C2:* LMM-1 and LMM-2 integrated with Logic-RAG are the best-performing models that yields the same performance as the standalone Logic-RAG. This is because LMMs’ responses are conditioned by the inference output of Logic-RAG.

*f) Results: Condition C3:* LMMs with only the kb (in template sentences) of Logic-RAG still outperform LLMs without it. For example, LMM-1 achieves an accuracy of 0.69 and an  $F_1$  score of 0.79, while LMM-2 reaches 0.72 and 0.81, respectively. Compared to standalone LMMs, this represents a 15% improvement in accuracy and an 11% increase in  $F_1$  score. The integration of template sentences significantly enhances LMMs’ performance, particularly for complex reasoning questions. For **U3**, we observe  $F_1$  score improvements of 0.35 and 0.15 for LMM-1 and LMM-2, respectively. Similarly, for **B2**, improvements of 0.14 and 0.10 are observed in the  $F_1$  score for LMM-1 and LMM-2, respectively (Table II).

*g) Results: VSR question types:* LMM-1 and LMM-2 in standalone condition perform reasonably well in simple *object queries* ( $F_1$  scores: 0.87 and 0.85), but struggle with deeper spatial understanding tasks. Their performance decreases for *changes in velocity (U3)* ( $F_1$  scores: 0.44 and 0.60) and *relative distance change (B2)* ( $F_1$  scores: 0.59 and 0.62), indicating limitations in complex spatial reasoning.

### B. Testing the Replacibility of Components within Perception Module and Performance on Real-World Driving Dataset

In this experiment, we evaluated how easily one could replace various components in the perception module by

substituting them with newer ones. Specifically, we substituted Mask2Former [37] with HRNetv2 [22] and CoTracker3 [38] with Rigid Mask [36]. We used these models without fine-tuning or modifying other system components.

We used the test set of the KITTI object tracking dataset [16]. It contained 29 real-world video sequences (84 to 1,175 frames each, totaling approximately 11,000 frames). Similar to our previous experiment, we cut these sequences into 10-frame clips, yielding 1,100 clips total. From this collection, we randomly selected 19 clips and manually created 100 visual-spatial reasoning questions with corresponding ground-truth answers. We then generated responses under condition C1 to compare our reconfigured system’s real-world performance directly against baseline methods (LMM-1 and LMM-2) operating in standalone mode.

Method	Accuracy $\uparrow$	$F_1$ $\uparrow$
<b>Logic-RAG</b>	<b>0.91</b>	<b>0.95</b>
LMM-1 (w/o LR)	0.71	0.83
LMM-2 (w/o LR)	0.74	0.84

TABLE III

ACCURACY AND  $F_1$  SCORES OF LOGIC-RAG WITH UPDATED PERCEPTION MODULE AND BASELINES FOR VSR QUESTIONS ON KITTI DATASET.

*a) Results:* Logic-RAG with its updated perception module achieved 91% accuracy ( $F_1$  score: 0.95), outperforming both baseline methods. Compared to LMM-1 (accuracy: 71%;  $F_1$  score: 0.83) and LMM-2 (accuracy: 74%;  $F_1$  score: 0.84), our system demonstrated improvements of 20 and 17 percentage points respectively. Table III presents these comparative results in detail.

## V. CONCLUSION

We introduce Logic-RAG, a first-order logic-based framework that enhances visual-spatial reasoning capabilities of popular LMMs through the retrieval-augmented generation mechanism. Logic-RAG constructs a dynamic knowledge base about object-object relationships and spatial dynamics in driving scenes. This work addresses crucial spatial reasoning deficiencies in LMMs for autonomous driving applications, improving system interpretability and user trust. Future research could focus on extending the framework to more complex scenarios, incorporating temporal reasoning, and exploring real-time optimizations for practical deployment in autonomous vehicles.

## VI. ACKNOWLEDGEMENT

This work was supported in part by the National Science Foundation under grant #2326406.



## REFERENCES

- [1] D. Omeiza, H. Webb, M. Jirotko, and L. Kunze, "Explanations in autonomous driving: A survey," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 8, pp. 10 142–10 162, 2021.
- [2] Z. Yang, X. Jia, H. Li, and J. Yan, "A survey of large language models for autonomous driving," *arXiv preprint arXiv:2311.01043*, 2023.
- [3] X. Zhou, M. Liu, B. L. Zagar, E. Yurtsever, and A. C. Knoll, "Vision language models in autonomous driving and intelligent transportation systems," *arXiv preprint arXiv:2310.14414*, 2023.
- [4] Z. Xu, Y. Zhang, E. Xie, Z. Zhao, Y. Guo, K.-Y. K. Wong, Z. Li, and H. Zhao, "Drivep4: Interpretable end-to-end autonomous driving via large language model," *IEEE Robotics and Automation Letters*, 2024.
- [5] L. Chen, O. Sinavski, J. Hünermann, A. Karnsund, A. J. Willmott, D. Birch, D. Maund, and J. Shotton, "Driving with llms: Fusing object-level vector modality for explainable autonomous driving," *arXiv preprint arXiv:2310.01957*, 2023.
- [6] F. Liu, G. Emerson, and N. Collier, "Visual spatial reasoning," *Transactions of the Association for Computational Linguistics*, vol. 11, pp. 635–651, 2023.
- [7] A. Kamath, J. Hessel, and K.-W. Chang, "What's" up" with vision-language models? investigating their struggle with spatial reasoning," *arXiv preprint arXiv:2310.19785*, 2023.
- [8] T. Thrush, R. Jiang, M. Bartolo, A. Singh, A. Williams, D. Kiela, and C. Ross, "Winoground: Probing vision and language models for visio-linguistic compositionality," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 5238–5248.
- [9] L. Pan, A. Albalak, X. Wang, and W. Wang, "Logic-LM: Empowering large language models with symbolic solvers for faithful logical reasoning," in *Findings of the Association for Computational Linguistics: EMNLP 2023*, H. Bouamor, J. Pino, and K. Bali, Eds. Singapore: Association for Computational Linguistics, Dec. 2023, pp. 3806–3824. [Online]. Available: <https://aclanthology.org/2023.findings-emnlp.248>
- [10] P. West, X. Lu, N. Dziri, F. Brahman, L. Li, J. D. Hwang, L. Jiang, J. Fisher, A. Ravichander, K. Chandu, *et al.*, "The generative ai paradox: "what it can create, it may not understand"," in *The Twelfth International Conference on Learning Representations*, 2023.
- [11] S. Banerjee, A. Agarwal, and S. Singla, "Llms will always hallucinate, and we need to live with this," *arXiv preprint arXiv:2409.05746*, 2024.
- [12] OpenAI, "GPT-4," 2024.
- [13] Anthropic, "Claude 3 opus," 2024. [Online]. Available: <https://claude.ai/chats>
- [14] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, *et al.*, "Retrieval-augmented generation for knowledge-intensive nlp tasks," in *Advances in Neural Information Processing Systems*, 2020.
- [15] K. Lu, S. Zhang, P. Stone, and X. Chen, "Robot representation and reasoning with knowledge from reinforcement learning," *arXiv preprint arXiv:1809.11074*, 2018.
- [16] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *2012 IEEE conference on computer vision and pattern recognition*. IEEE, 2012, pp. 3354–3361.
- [17] T. Gupta, A. Kamath, A. Kembhavi, and D. Hoiem, "Towards general purpose vision systems: An end-to-end task-agnostic vision-language architecture," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 16 399–16 409.
- [18] N. Rajabi and J. Kosecka, "Towards grounded visual spatial reasoning in multi-modal vision language models," *ICLR*, 2024.
- [19] S. J. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*. Pearson Education Limited, 2016.
- [20] R. Manhaeve, S. Dumancic, A. Kimmig, T. Demeester, and L. De Raedt, "Deepproblog: Neural probabilistic logic programming," in *Advances in Neural Information Processing Systems*, 2018.
- [21] T. Rocktäschel and S. Riedel, "End-to-end differentiable proving," in *Advances in Neural Information Processing Systems*, 2017.
- [22] J. Wang, K. Sun, T. Cheng, B. Jiang, C. Deng, Y. Zhao, D. Liu, Y. Mu, M. Tan, X. Wang, *et al.*, "Deep high-resolution representation learning for visual recognition," *IEEE transactions on pattern analysis and machine intelligence*, vol. 43, no. 10, pp. 3349–3364, 2020.
- [23] J. Lambert, Z. Liu, O. Sener, J. Hays, and V. Koltun, "Mseg: A composite dataset for multi-domain semantic segmentation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 2879–2888.
- [24] G. Yang and D. Ramanan, "Learning to segment rigid motions from two frames," in *CVPR*, 2021.
- [25] A. Agarwal and C. Arora, "Attention attention everywhere: Monocular depth prediction with skip attention," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2023.
- [26] M. A. Reza, H. Zheng, G. Georgakis, and J. Košecká, "Label propagation in rgb-d video," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 4917–4922.
- [27] A. Roshan Zamir, A. Dehghan, and M. Shah, "Gmcp-tracker: Global multi-object tracking using generalized minimum clique graphs," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2012.
- [28] H. Durrant-Whyte and T. Bailey, "Simultaneous localization and mapping: part i," *IEEE Robotics & Automation Magazine*, vol. 13, no. 2, pp. 99–110, 2006.
- [29] Y. Furukawa and C. Hernández, "Multi-view stereo: A tutorial," *Foundations and Trends® in Computer Graphics and Vision*, vol. 9, no. 1-2, pp. 1–148, 2015.
- [30] H. B. Lee *et al.*, "Accuracy limitations of range-range (spherical) multilateration systems." Lincoln Laboratory, Tech. Rep., 1973.
- [31] A. Dubey, A. Jauhri, A. Pandey, A. Kadian, A. Al-Dahle, A. Letman, A. Mathur, A. Schelten, A. Yang, A. Fan, *et al.*, "The llama 3 herd of models," *arXiv preprint arXiv:2407.21783*, 2024.
- [32] M. Abdin, S. A. Jacobs, A. A. Awan, J. Aneja, A. Awadallah, H. Awadalla, N. Bach, A. Bahre, A. Bakhtiari, H. Behl, *et al.*, "Phi-3 technical report: A highly capable language model locally on your phone," *arXiv preprint arXiv:2404.14219*, 2024.
- [33] B. T. Willard, "minikanren as a tool for symbolic computation in python," *arXiv preprint arXiv:2005.11644*, 2020.
- [34] Aimaicode, "GitHub - aimaicode/aima-python: Python implementation of algorithms from Russell And Norvig's "Artificial Intelligence - A Modern Approach"." [Online]. Available: <https://github.com/aimaicode/aima-python>
- [35] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "Carla: An open urban driving simulator," in *Conference on robot learning*. PMLR, 2017, pp. 1–16.
- [36] G. Yang and D. Ramanan, "Learning to segment rigid motions from two frames," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 1266–1275.
- [37] B. Cheng, I. Misra, A. G. Schwing, A. Kirillov, and R. Girdhar, "Masked-attention mask transformer for universal image segmentation," 2022.
- [38] N. Karaev, I. Makarov, J. Wang, N. Neverova, A. Vedaldi, and C. Rupprecht, "Cotracker3: Simpler and better point tracking by pseudo-labelling real videos," *arXiv preprint arXiv:2410.11831*, 2024.