

Explainable Lane Change Prediction for Near-Crash Scenarios Using Knowledge Graph Embeddings and Retrieval Augmented Generation

M. Manzour¹, A. Ballardini¹, R. Izquierdo¹, and M. A. Sotelo¹

¹Department of Computer Engineering, University of Alcalá, Madrid, Spain
[ahmed.manzour, augusto.ballardini, ruben.izquierdo, miguel.sotelo]@uah.es

Abstract—Lane-changing maneuvers, particularly those executed abruptly or in risky situations, are a significant cause of road traffic accidents. However, current research mainly focuses on predicting safe lane changes. Furthermore, existing accident datasets are often based on images only and lack comprehensive sensory data. In this work, we focus on predicting risky lane changes using the CARLA Risky-lane-change Anticipation in Simulated Highways (CRASH) dataset (our own collected dataset specifically for risky lane changes), and safe lane changes (using the HighD dataset). Then, we leverage Knowledge Graphs (KGs) and Bayesian inference to predict these maneuvers using linguistic contextual information, enhancing the model's interpretability and transparency. The model achieved a 91.5% f1-score with anticipation time extending to four seconds for risky lane changes, and a 90.0% f1-score for predicting safe lane changes with the same anticipation time. We validate our model by integrating it into a vehicle within the CARLA simulator in scenarios that involve risky lane changes. The model managed to anticipate sudden lane changes, thus providing automated vehicles with further time to plan and execute appropriate safe reactions. Finally, to enhance the explainability of our model, we utilize Retrieval Augmented Generation (RAG) to provide clear and natural language explanations for the given prediction.

Index Terms—Risky Lane Change Prediction, Near-Crash, Knowledge Graph Embeddings, Bayesian Inference, CARLA, Retrieval Augmented Generation.

I. INTRODUCTION

Lane-changing maneuvers (especially abrupt lane changes) are one of the causes of vehicle crashes, as a report indicated that 33% of all road crashes occur due to driver decision errors including lane changes [1]. A risky lane change occurs when drivers switch lanes suddenly without adequate warning, which leaves insufficient time for other vehicles to react, leading to frequent near-crashes or collisions. In contrast, a safe lane change usually comes with early signals or clear intentions to change lanes, giving other drivers sufficient time to react. While recent research has mainly focused on predicting safe lane changes, there is a critical gap when it comes to addressing risky lane-changing scenarios. Figure 1 shows a risky lane change in which the white vehicle was involved in a risky situation, it was forced to change lanes to the left with little warning which led to a severe collision. Ignoring these unsafe lane-changing events makes it difficult for automated systems to anticipate and prevent near-crash situations. Moreover, many models that use numerical inputs function like black boxes, making it tough to interpret or

explain their outputs. This lack of transparency poses challenges when trying to justify predictions to users who may not be familiar with the underlying algorithms. Additionally, existing accident datasets do not provide sufficient input for designing explainable prediction models. These datasets typically consist of video recordings without accompanying numerical or linguistic data, such as the velocities or positions of vehicles. This absence of crucial information makes it hard to develop models that provide clear and comprehensive predictions. To address these challenges, this work focuses




Figure 1. Target vehicle made a risky left lane change that led to a collision.

on addressing the following points:

- 1) A customized dataset (CRASH dataset) of risky lane-changing scenarios is developed using the CARLA simulator, reconstructing real-world scenes with detailed numerical data.
- 2) An interpretable and transparent prediction model for safe and risky lane changes is developed by combining KGs with Bayesian inference and incorporating contextual linguistic information to improve interpretability.
- 3) The model is validated in the CARLA simulator, demonstrating that it can predict unsafe maneuvers in advance and provide the vehicle with sufficient time to react accordingly.
- 4) The model's explainability is increased by using RAG to provide clear natural language explanations for the predictions.

The rest of this article is organized as follows. Section II presents the state of the art. Section III contains a brief introduction to the CRASH dataset. Then, our proposed methodology is discussed in detail in section IV. In section V, results will be presented. Finally, section VI concludes the work and provides some recommendations for future work.

II. STATE OF THE ART

Recently, different works have focused on predicting vehicle lane changes using different types of inputs and using different models. In 2019, [2] used two machine learning models; Support Vector Machine (SVM) and Artificial Neural Network (ANN) to predict lane changes of surrounding vehicles on highways. They used the Next Generation Simulation (NGSIM) dataset, their input features included vehicle's speeds and accelerations in both longitudinal and lateral directions, distances to adjacent lane boundaries, yaw angles, and yaw rates. Also in 2019, [3] focused on predicting the lane changes of surrounding vehicles using visual information from the PREVENTION dataset. They explored two methods: the first was a Motion History Image (MHI) combined with a Convolutional Neural Network (CNN), where temporal and visual data were fed into the CNN. The second method involved a GoogleNet-LSTM model, where features extracted by a GoogleNet CNN were passed to an Long Short-Term Memory (LSTM) to learn temporal patterns. Inputs included RGB images, dimensions and center coordinates (x, y) of the bounding boxes around vehicles.

In 2020, [4] trained Recurrent Neural Network (RNN) and LSTM models using the PREVENTION dataset to predict lane-changing intentions by tracking the positions of surrounding vehicles, specifically the centers of their bounding boxes. In 2018, [5] employed an LSTM model to predict lane changes by considering both the vehicle's past trajectory and the states of surrounding vehicles. Using data from the NGSIM dataset, they incorporated features such as the vehicle's lateral and longitudinal positions relative to the lane, its acceleration, the presence of vehicles to the right or left, and the distances to surrounding vehicles.

In 2022, [6] utilized eXtreme Gradient Boosting (XGBoost) and LSTM models to predict lane change decisions and future trajectories using scenarios from the HighD dataset. Their approach considered factors like traffic density, vehicle type, and the relative movements between the target vehicle and the surrounding vehicles. Initially, they built a traffic flow model using the target vehicle's longitudinal speed and acceleration, headway distance, and relative velocities to the surrounding vehicles.

In 2023, a dual transformer model was developed by [7], consisting of two parts: one for predicting lane changes and another for forecasting trajectories. The first transformer used the target vehicle's historical lateral movements and the states of surrounding vehicles, including distances and velocities relative to the target vehicle. The output from this lane change prediction was then combined with the target vehicle's past lateral movements and fed into the second transformer to link intentions with future trajectories. This model was trained and tested on the HighD and NGSIM datasets.

In 2024, [8] used Knowledge Graph Embedding (KGE) followed by Bayesian inference which acted as a downstream task on the grounds of the learned embedding to predict safe lane changes on the HighD dataset. They considered linguistic inputs to enhance the model's interpretability. The considered features are the vehicle lateral velocity and acceleration,

Time To Collision (TTC) risk with the surrounding vehicles. Work [9] used the same methodology. However, it extends the architecture by adding the RAG mechanism which leverages the power of Large Language Models (LLMs) to generate context-specific explanations that are grounded in external, verifiable knowledge, and provide explanations of the scene provided with the reason for the lane change in natural language which is easy to understand by end-users.

The work in [10] addressed predicting safe lane changes and trajectories in the HighD dataset by leveraging the strong reasoning capabilities and self-explanation abilities of LLMs. Data is processed as numerical values in natural language prompts to be the input to the LLMs. Then, they utilized Chain-of-Thought (COT) reasoning to enhance prediction explainability. The advantage of this work is that it provides high-performance metrics and fine explanations to the end users with only the powers of the tuned LLMs. The challenge is that even though the LLMs provides explanations and employs COT reasoning to interpret input prompts, there remains a gap in the transparency and grounding of the LLM reasoning process as these models rely on internal representations (the knowledge and patterns that the LLMs has learned internally during its training). This reliance can make the LLMs reasoning process less transparent, and the explanations might be difficult to verify or to be traced back to the inputs. So, the way it processes information internally is not easily visible or understandable to the users. Also, users and experts still cannot trace the prediction process starting from having the numerical input data till the prediction, they can't see how a change in the input can affect the probability of the final prediction, and it is often considered a "black box" in that case. On the other side, integrating KGs and Bayesian inference improves interpretability and transparency by creating a structured and interpretable framework where users can trace the process of obtaining the prediction and can show the effect of each input on the probability of the final prediction through the priors given by the KG for every event (input), which is something hard to trace in the tuned LLM approach. Moreover, without grounding in external knowledge, LLMs explanations may be less reliable. However, by using RAG, we ensure that explanations are grounded in external, verifiable knowledge, enhancing reliability.

Finally, Table I presents a comparison of the literature discussed previously to highlight the identified research gaps. The first column indicates whether the models in the related works use numerical inputs or linguistic explainable inputs. The second and third columns specify whether the work addresses safe or risky lane changes respectively. The fourth column notes whether the related work includes experimental validation for simple decision-making (like gentle braking). The fifth column indicates whether the architecture provides a natural language explanation. The last row summarizes the aspects that our work addresses. After a thorough examination of the existing literature, we have identified the following research gaps, which are addressed in this work:

- 1) **Specialized dataset for risky lane changes:** There is a lack of datasets focusing specifically on risky

TABLE I
Comparison of related works on lane change prediction based on input types, focus, and explainability.

Work	Linguistic Explainable Inputs	Safe Lane Changes	Risky Lane Changes	Decision Making	Natural Language Explanation
[2], [3], [4] [5], [6], [7]	✗	✓	✗	✗	✗
[8]	✓	✓	✗	✗	✗
[9]	✓	✓	✗	✗	✓
[10]	✗	✓	✗	✗	✓
Ours	✓	✓	✓	✓	✓

lane-change (e.g. abrupt cut-in) scenarios. We address this gap by collecting the CRASH dataset using the CARLA simulator.

- 2) **Prediction based on contextual linguistic information:** Current models rely on trajectory and numerical data from past experiences. We propose a method to predict both safe and risky lane-change maneuvers using contextual linguistic information, moving beyond mere numerical inputs.
- 3) **Interpretable and transparent predictions:** Many existing models lack interpretability and transparency. Our approach utilizes KGE and Bayesian inference to provide interpretable and transparent predictions.
- 4) **Validation in simulation environment:** There is a need for models to be validated in realistic simulation settings. We validate our model within the CARLA simulator, demonstrating its ability to predict unsafe maneuvers in advance and provide the vehicle with sufficient time to respond appropriately.
- 5) **Enhanced explainability through RAG integration:** To further improve the model's reliability and explainability, we integrate RAG, which provides clear explanations for the predictions based on external, reliable knowledge sources. This integration enriches the architecture's explainability, making the target vehicle's decision-making process more understandable to users.

Addressing these gaps will not only give the ability to provide explanations but also make the architecture more suitable for safety-critical applications where input-to-output traceability, understanding, and transparency are essential.

III. DATASET

This section briefly introduces the CRASH dataset which focuses on risky lane-changing maneuvers. The dataset is created by examining videos from various datasets containing different real accident and near-crash scenarios, some of these scenarios include lane changes [11][12]. These datasets were collected from front-view cameras mounted on vehicles. However, they only provide images and lack numerical data about the vehicle's state and the state of the surrounding vehicles. Therefore, CARLA simulator is utilized to recreate similar scenes [13]. By integrating CARLA, it is possible to generate this scenarios and obtain numerical values for all required inputs concerning the ego vehicle, target vehicle (which does the risky maneuver), and all surrounding vehicles in the scene. The dataset contains detailed information about every vehicle in the scene. For each vehicle, data are provided on its (x, y) location within the map, velocities, and accelerations in the lateral and longitudinal directions, braking status

(whether the vehicle is braking or not), rear braking light and turning signal status, lane position, position relative to the center of the lane, gap distance and time to collision with the surrounding vehicles. Additionally, if there is an adjacent left or right vehicle, the lateral distance between the vehicle and the adjacent vehicle is included. The dataset also includes lane densities for the left, center, and right lanes, all measured relative to the vehicle's position, as well as the average speeds of these lanes. Furthermore, the lane attraction score is provided, which is calculated as each lane's average speed divided by its density, it gives an intuition about which lane will be more attractive to the vehicle based on the average velocity and density of that lane.

The dataset contains a total of 50 samples, including 25 left lane changes, and 25 right lane changes, all these samples are classified as risky maneuvers that end up with a collision or a near-crash. Safe maneuvers are not included because existing datasets, such as the HighD dataset, already cover safe lane-changing maneuvers. The scenarios are recorded from two different perspectives. The first is the windshield view from the ego vehicle's perspective. The second is the bird's eye view, which allows the observer to see the entire situation, including how the ego vehicle interacts with the target vehicle's risky maneuver and all surrounding vehicles.

IV. METHODOLOGY

A. Architecture Overview

Our architecture is divided into four stages as indicated in fig. 2. The first stage is Linguistic Input Generation, where we take our numerical input values and convert it to linguistic categories. This is obtained based on literature and the data distribution along the lane-changing categories. For the second stage (Knowledge Graph Embedding), we use the linguistic inputs to create our KG, which is in the form of triples inside a CSV file. Then, this KG is embedded using the Ampligraph library. In the third stage, we carry out Bayesian inference on the grounds of these learned embeddings to get our final prediction. Finally, we feed this prediction along with the linguistic inputs to the Retrieval Augmented Generation stage in order to get an explanation of these predictions and the target vehicle decision-making process.

B. Input and Knowledge Graph Ontology Definition

From the dataset, the features selected include the y-velocity and y-acceleration of the target vehicle, the vehicle's lane position, distance from the center of the vehicle to the center of the lane, the time headway between the target vehicle and the preceding vehicle, time to collision with the (left preceding, preceding, right preceding, left following,

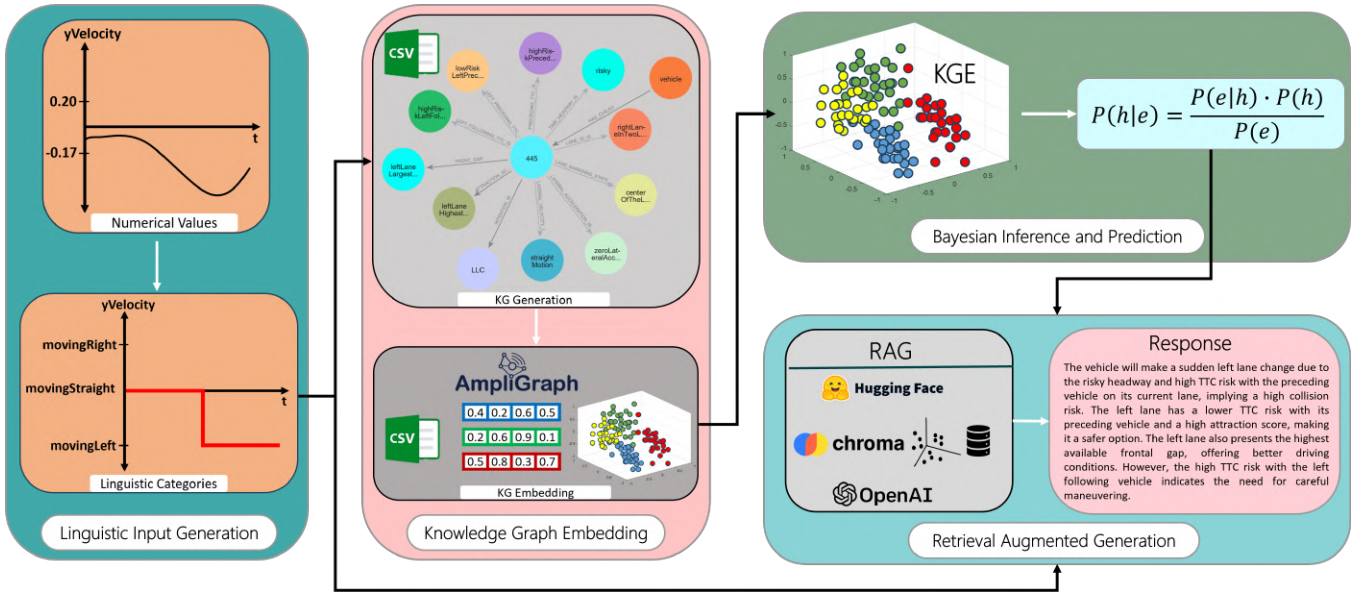


Figure 2. The pipeline for anticipating and explaining lane changes consists of four phases: Linguistic Input Generation, Knowledge Graph Embedding, Bayesian Inference and Prediction, and Retrieval Augmented Generation.

and right following) vehicles, lane with the highest attraction score and lane with the highest frontal gap. These features are extracted in numerical format from the HighD [14] and CRASH datasets. The numerical features are converted into linguistic categories based on literature and the data distribution along the lane-changing categories. For instance, the y-velocity is classified into *leftMotion*, *straightMotion*, and *rightMotion*. Similarly, the y-acceleration is classified into *leftAcceleration*, *zeroLateralAcceleration*, and *rightAcceleration*. These thresholds are determined by analyzing the normal distribution of each feature category, following the methodology outlined in [8]. The lane position is already a categorical feature. For the distance from the center of the lane and time headway, thresholds are determined based on the normal distribution following the same approach as in the lateral velocity and acceleration. Thresholds for the TTC are defined following previous studies [15], [16]: (0, 4] seconds (high risk), (4, 10] seconds (medium risk), and above 10 seconds or negative values (low risk). For the lane attraction and frontal gap, thresholds are not applied; instead, lanes with the highest attraction score and frontal gap are selected directly. Table II introduces the ontology of our KG which provides a formal and general representation of the entities and their relationships within a KG. Ontologies are essential in KGs because they serve as a schema for constructing the graph, ensuring consistency and improving its interpretability.

C. Knowledge Graph Embedding Phase

A KG is a directed heterogeneous multigraph, where nodes can have different types, and each pair of nodes can be connected by multiple types of relationships. KGs are used to represent the data in the form of triples inside a CSV file. Each feature was expressed in one triple: subject, relation, and object. For example, the triple $\langle \text{vehicle } 445, \text{LATERAL_ACCELERATION_IS, leftAcceleration} \rangle$ is mapped to $\langle \text{subject, relation, object} \rangle$. There is a generic entity called

vehicle which has several children such as *vehicle 445*, *vehicle 446*. Each vehicle ID is connected to its corresponding features. Figure 3 shows KG instant for one vehicle from the full KG. Figure 4 presents KG instances for 80 vehicles, given that the total number of instances used for training is 40K, making it hard to fit in a single figure. It is important to note that in our KG representation, *vehicle 445* and *vehicle 446* are considered different vehicles, although they could represent the same vehicle in a real scenario. So, during training, each vehicle will have a new child ID in each frame, even if it is the same vehicle. During validation, vehicle IDs are not used. However, the parent entity *vehicle* points to all the other features. After that, we start to train/embed



Figure 3. A KG instance for a single vehicle at one frame within the dataset.

this knowledge graph. KGE is a supervised machine learning technique that learns to represent (embed) the knowledge graph entities and relations into a low-dimensional vector space while preserving semantic meaning [17]. In this stage, we use the AmpliGraph library [18]. The scoring function used is the TransE function. The embedding size is set to

TABLE II

Ontology table which includes the definition of all entities (classes), their instances, and possible relations that can be connected to them.

Class	Class Description	Instance	Possible Relation
intention	Lane change intention of the vehicle	LLC LK RLC	INTENTION_IS
lateralVelocity	Lateral velocity	{left, straight, right} Motion	LATERAL_VELOCITY_IS
lateralAcceleration	Lateral acceleration	{left, right} Acceleration zeroLateralAcceleration	LATERAL_ACCELERATION_IS
laneID	Vehicle's lane position	{left, right} LaneInTwoLanesRoad {left, center, right} LaneInThreeLanesRoad	LANE_ID_IS
laneMarkingOffset	Lateral distance from lane marking	near {left, right} Marking centerOfTheLane	LANE_MARKING_STATE_IS
timeHeadway	Time headway	safe risky collision	TIME_HEADWAY_IS
ttcLeftPreceding	TTC with the left preceding vehicle	{low, medium, high} RiskLeftPreceding	LEFT_PRECEDING_TTC_IS
ttcPreceding	TTC with the preceding vehicle	{low, medium, high} RiskPreceding	PRECEDING_TTC_IS
ttcRightPreceding	TTC with the right preceding vehicle	{low, medium, high} RiskRightPreceding	RIGHT_PRECEDING_TTC_IS
ttcLeftFollowing	TTC with the left following vehicle	{low, medium, high} RiskLeftFollowing	LEFT_FOLLOWING_TTC_IS
ttcRightFollowing	TTC with the right following vehicle	{low, medium, high} RiskRightFollowing	RIGHT_FOLLOWING_TTC_IS
attractionScore	Lane with the highest attraction score	{left, center, right} LaneHighestAttraction	ATTRACTION_SCORE
frontGap	Lane with the largest front gap	{left, center, right} LaneLargestFrontGap	FRONT_GAP
vehicleID	Child vehicle ID which changes every frame	vehicle ID number (e.g. '445')	HAS_CHILD
vehicle	Generic entity pointing to every child vehicle	—	Any




Figure 4. A sample KG representation with 80 instances from a total of 40K instances.

100. The model parameters include using the Adam optimizer with a learning rate of 0.0005, a batch size of 10,000, and employing SelfAdversarialLoss. Additionally, for each positive triple in our dataset, we generate five negative triples, maintaining a corruption ratio of 5 : 1. An early stopping criterion is used to monitor the Mean Reciprocal Rank (MRR) metric during validation with a patience threshold of five validation epochs.

D. Bayesian Inference and Prediction Phase

In this stage, Bayesian inference is performed as a downstream task using the learned embeddings from the previous stage. This stage utilizes the Bayesian inference eq. (1), which involves calculating the probabilities of a hypothesis (h), an event given the hypothesis ($e|h$), and an event (e), to determine the probability of the hypothesis given the event, as shown in fig. 5. Here, events are data obtained from sensors. After converting the data to the linguistic format, we formulate the triples related to these events and propose a hypothesis, for instance, the vehicle will make a left lane

change. We then calculate $P(LLC)$ and $P(e|LLC)$ as indicated in fig. 5. For example, we ask the model to calculate the probability of a vehicle moving straight, given that it intends to make a left lane change. In case of the presence of multiple events, eq. (2) and eq. (3) are applied to determine the final values of $P(e)$ and $P(e|h)$, respectively.

$$P(h|e) = \frac{P(h)P(e|h)}{P(e)} \quad (1)$$

$$P(e) = P(e_1) \times \dots \times P(e_n) \quad (2)$$

$$P(e|h) = P(e_1, \dots, e_n|h) = P(e_1|h) \times \dots \times P(e_n|h) \quad (3)$$

These probabilities are evaluated using the embeddings obtained in the previous stage and are then fed into the Bayesian inference equation. This allows us to calculate $P(LLC|e)$ which is the probability of the hypothesis (left lane change) given the events. However, this is just one hypothesis, we have to consider other possible maneuvers, which are lane keep or a right lane change. So, the probability for each maneuver is calculated, and the hypothesis with the highest probability is considered as the model's final prediction. Once the final prediction is obtained, its reasoning can be traced by comparing $P(e|h)$ for each event across all hypotheses. As shown in the figure, $P(highRiskPreceding|LLC)$ is higher than $P(highRiskPreceding|LK)$ and $P(highRiskPreceding|RLC)$. However, this trend does not hold for the *movingStraight* event. This indicates that the *highRiskPreceding* event is the primary factor influencing the prediction in this situation and highlights the transparency of the model.

E. RAG Explanation Phase

After generating the prediction, we need our architecture to justify why our model provided this prediction and reason about the target vehicle decision making in natural language using an LLM. However, the challenge is that LLMs have access to a wide public database, which does not fit our purposes. We require the justifications to be based only on the provided KGs derived from the HighD and the CRASH datasets, which are private databases, not public ones. We also need the justifications to be more reliable and verifiable.

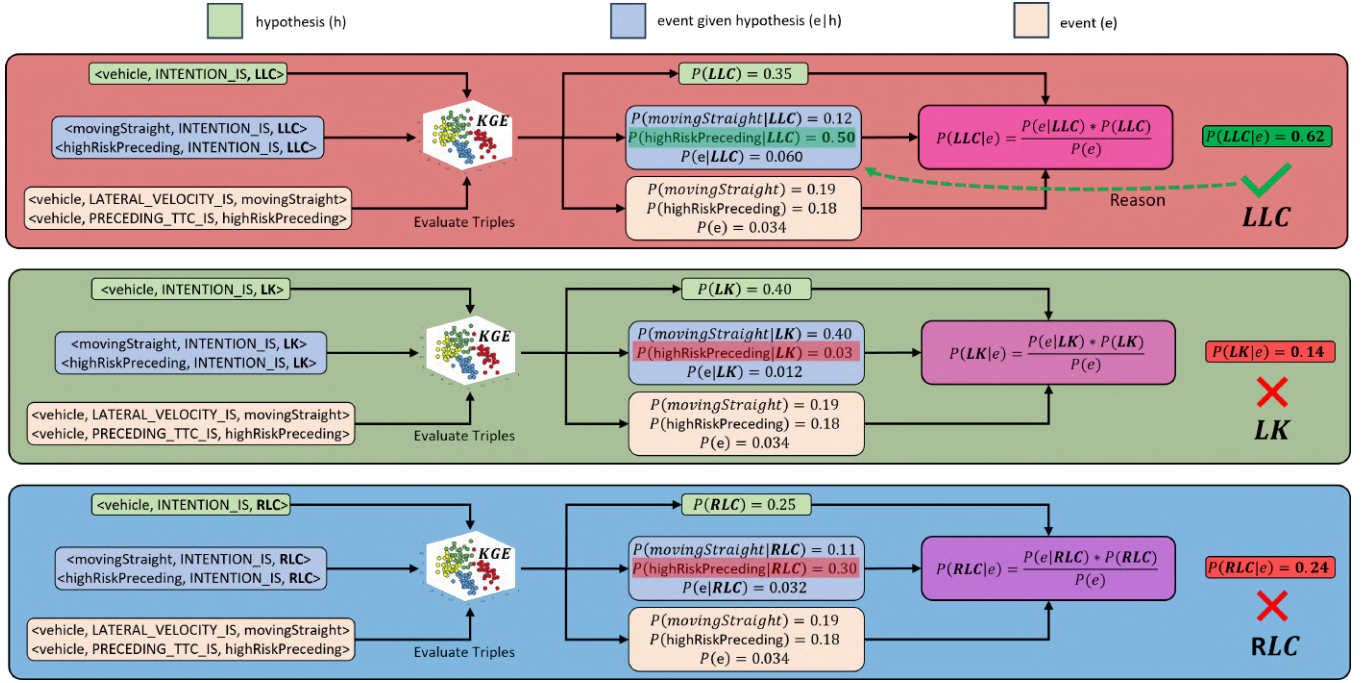


Figure 5. Example of Bayesian inference reasoning for a left lane change prediction that shows the model traceability and transparency. The target vehicle will do the maneuver because $P(\text{highRiskPreceding}|\text{LLC})$ is higher than $P(\text{highRiskPreceding}|\text{LK})$ and $P(\text{highRiskPreceding}|\text{RLC})$.

This can be achieved by grounding them in external, independently verifiable knowledge sources rather than relying solely on internal knowledge acquired during the training process.

RAG is used to address this issue. First, we organize our KG triples into chunks, with each chunk representing one sample from our data. We then embed these chunks using an embedding model and store the resulting embeddings in a vector database. *all-MiniLM-L6-v2* from Hugging Face is the used embedding model and *Chroma* is the used vector database. When a user submits a query, the system embeds the query and retrieves chunks with embeddings that are most similar to the query. These retrieved chunks are then augmented to the query, and to a prompt that guides our model to the type of needed responses to ensure that the responses are accurate and based solely on the private KG data. The final augmented query is then fed into an *OpenAI GPT-4* LLM to generate the required response.

V. RESULTS

In this section, we present the results of our work, which aims to predict both safe and risky lane changes using a unified model and consistent labels. We compare our approach with existing studies that focus solely on predicting safe lane-changing maneuvers. Furthermore, we validate our model by integrating it into an IDM within the CARLA simulation environment to test the IDM reactions with and without our prediction model. Finally, we demonstrate the explainability of our approach by utilizing the RAG module to provide natural language explanations for the model's predictions.

A. Model Performance and Comparative Analysis

Regarding safe lane-changing maneuvers, our model achieved an average f1-score of 90% over the interval

[0,4] seconds. We identified four relevant articles addressing safe lane-changing maneuvers using the HighD dataset, and compared their results with ours in table III. The first two

TABLE III
Performance comparison of models predicting **safe** lane changes on the HighD dataset using the f1-score metric.

Pred. time (s)	0.5s	1s	1.5s	2s	2.5s	3s	3.5s	4s
[6]	98.2	97.1	96.6	95.2	—	—	—	—
[7]	99.2	99.0	97.6	91.8	—	—	—	—
[8]	97.7	97.9	98.1	98.0	97.2	93.6	82.8	66.5
Ours	95.0	95.5	95.5	95.0	94.0	90.0	80.5	65.5
Interval (s)	[0,1]	(1,2]	(2,3]	(3,4]	[0,4]	—	—	—
[10]	98.5	98.9	98.1	93.0	97.1	—	—	—
Ours	95.5	95.8	92.4	76.0	90.0	—	—	—

articles ([6], [7]) report results at specific times. Their models exhibit higher performance than ours at time points (0.5 to 2 seconds). Also, no results are indicated for longer horizons. Additionally, their models employ deep learning techniques, such as LSTMs and transformers, which are considered "black-boxes" and lack transparency and interpretability. In contrast, our model is interpretable and provides transparency to users, as indicated in fig. 5. The third article [8] also demonstrates consistently high-performance scores. However, like the previous two models, this model was trained only to predict safe lane changes, whereas our model is designed to predict both safe and risky lane changes. Another study reports remarkable results for predicting safe lane changes on the HighD dataset, presenting performance over intervals from [0,1], (1,2], (2,3], (3,4], and [0,4] seconds. Their model utilizes an LLM with a chain-of-thought approach. While the model's results are impressive, the model is based on numerical inputs and lacks transparency and verifiability, similar to other "black-box" models where the internal workings of

the model are not accessible, making it difficult to understand which inputs influence the predictions. In contrast, our model provides high transparency as mentioned previously and it also deals with risky situations (as mentioned in table I) which are not considered in state-of-the-art systems.

For risky lane changes, results are shown in table IV. The model attained an f1-score of 91.5% over the interval [0,4] seconds. To the best of our knowledge, no prior studies have addressed these specific types of maneuvers. Therefore, we present our results accordingly without comparison to other studies. To demonstrate the significance of the risky dataset, we trained another model (model 2) only on the HighD dataset and tested it on the risky lane changes included in the CRASH dataset. The model is unable to predict risky lane changes effectively as shown in the last row of table IV, highlighting the importance of including a dataset that contains risky lane changes in the training process.

TABLE IV

Results of our model on **risky and safe** lane changes based on the f1-score metric. Model 1 is the original model trained on both CRASH and HighD datasets. Model 2 is trained on HighD only.

Pred. time (s)	[0,1]	(1,2]	(2,3]	(3,4]	[0,4]
Model 1 tested on CRASH (risky)	99.3	91.9	91.5	81.8	91.5
Model 1 tested on (CRASH+HighD) (mixed)	95.7	94.9	91.7	75.7	89.4
Model 2 tested on CARLA	89.3	36.1	30.6	21.9	50.7

B. Validation in Simulation and Real-World Scenarios

Initially, some scenarios are validated using the Intelligent Driver Model (IDM) integrated on the ego vehicle in CARLA without integrating our prediction model. In these cases, the IDM didn't have enough time to brake, ending up with a collision. For example, fig. 6 shows a scenario where the ego vehicle moves forward and the target vehicle is on the right front side. The IDM detects no obstacles ahead and increases its velocity. However, the target vehicle faces a risk with the vehicle in front of it and needs to change lanes. In this situation, the IDM does not know the intention of this vehicle or that the target vehicle is at risk which will lead to the execution of a risky left lane change. Consequently, the ego vehicle suddenly encounters another vehicle changing lanes directly ahead, leaving insufficient time to brake. After




Figure 6. The ego vehicle is moving forward safely (first two images). Then, the target vehicle unexpectedly changed lanes to the left to be in the ego vehicle's lane to avoid a slower preceding vehicle (third image). This leads to a collision due to insufficient braking time (fourth image).

integrating the prediction model, it anticipates that the target vehicle will execute a risky/sudden left lane change, as shown in fig. 7. This anticipation allowed the ego vehicle to decelerate smoothly, rather than braking abruptly. The model

provides the ego vehicle with awareness of the upcoming risky lane change, enabling it to decelerate in anticipation of the target vehicle's maneuver. Eventually, the target vehicle changed lanes to the left, while the ego vehicle continued moving safely.




Figure 7. In the first image, the ego vehicle is moving forward normally and predicts a lane-keep behavior for the target vehicle. Then, the target vehicle is predicted to make a safe lane change to the left, at that moment, the ego vehicle started to decelerate for more safety (second image). The third image shows the anticipation of a risky lane change. Finally, both vehicles continue to move safely without collision (fourth image).

Then, our model is tested on real-world videos that contain risky lane changes. The purpose of these scenarios is to demonstrate that our model functions effectively not only with simulation data but also in real-world situations. In fig. 8, we have some captures from a real-world video where our model anticipated the risky left lane change in advance. By integrating our model into the vehicle, we can potentially avoid near-crashes or accidents resulting from risky lane changes and understand the target vehicle's decisions. Finally,



Figure 8. The model anticipates a risky left lane change three seconds in advance: In the first image, the white target vehicle is moving straight as expected. In the second image, the model predicts a safe left lane change. However, the target vehicle continues moving forward, resulting in a risky left lane change shown in the third image, leading to a collision with the ego vehicle in the fourth image. If the ego vehicle had reacted based on the model's early predictions, this collision could have been avoided.

two links are attached in the bottom of this page. The first one directs to a Website¹ to access the CRASH dataset, and shows an interactive graph for a sample from our designed KG. The second link directs to a YouTube² playlist that contains some multimedia videos of different scenes including the scenes discussed in this section.

C. Natural Language Explanation Through RAG

After validation of the model, we analyze one of the captions extracted from the same real-world scenario, then we feed the linguistic inputs obtained from this capture along with the prediction to the RAG model. fig. 9 shows the formed query obtained from the linguistic inputs and the predicted maneuver. Then, the query is fed to the RAG model

¹<https://butternut-clef-a48.notion.site/Explainable-Lane-Change-Prediction-for-Near-Crash-Scenarios-Using-Knowledge-Graph-Embeddings-and-Ret-162eb56a8de78064966cc55607d5b11a?pvs=4>

²https://www.youtube.com/playlist?list=PLAeK3AuwxeFawQut8umkxfUXf6qoB_s7

which outputs an explanation that illustrates the model's ability to provide clear, reasonable, and precise natural language explanations for the target vehicle's predicted maneuver.

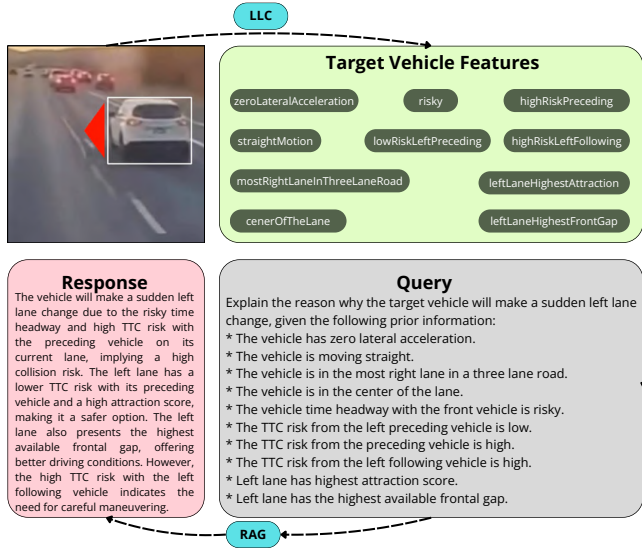


Figure 9. Query formation by merging the prediction with the linguistic inputs. Then, the query is fed to the RAG model to provide the explanation of the predicted risky left lane change.

VI. CONCLUSION

In this work, we anticipated safe and risky lane changes using KGE and Bayesian inference. We trained our model on the HighD dataset for safe lane changes and on our own created CRASH dataset for risky lane changes. Our model achieved an f1-score of 90% for predicting safe lane changes with an anticipation interval of 4 seconds and an f1-score of 91.5% for risky lane changes within the same anticipation window. In terms of numerical performance, our model is competitive with other state-of-the-art systems in anticipating safe lane changes and contributes by anticipating risky lane changes. Moreover, the model's performance goes beyond numerical results by enhancing transparency and interpretability of the prediction process. We validated our model by integrating it to an IDM in the CARLA simulator, and it is proven that predicting risky lane changes provides the ego vehicle with more time to react and avoid near-crash situations. Additionally, we employed RAG to enhance the model's explainability by providing natural language explanations for its predictions. Future work can focus on the decision-making process of the ego vehicle by incorporating advanced actions based on the predicted lane changes.

ACKNOWLEDGMENT

This research has been funded by the HEIDI project of the European Commission under Grant Agreement 101069538 and the SEGVAUTO5G-CM project of the Regional Government of Madrid under reference 2024/ECO-277.

REFERENCES

- [1] U. NHTSA, "Critical reasons for crashes investigated in the national motor vehicle crash causation survey," *DOT HS*, vol. 812, p. 115, 2015.
- [2] A. Benterki, M. Boukhifir, V. Judalet, and M. Choubeila, "Prediction of surrounding vehicles lane change intention using machine learning," in *2019 10th IEEE International Conference on Intelligent Data Acquisition and Advanced Computing Systems: Technology and Applications (IDAACS)*, vol. 2. IEEE, 2019, pp. 839–843.
- [3] R. Izquierdo, A. Quintanar, I. Parra, D. Fernández-Llorca, and M. A. Sotelo, "Experimental validation of lane-change intention prediction methodologies based on cnn and lstm," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2019, pp. 3657–3662.
- [4] O. Laimona, M. A. Manzour, O. M. Shehata, and E. I. Morgan, "Implementation and evaluation of an enhanced intention prediction algorithm for lane-changing scenarios on highway roads," in *2020 2nd Novel Intelligent and Leading Emerging Sciences Conference (NILES)*. IEEE, 2020, pp. 128–133.
- [5] S. Su, K. Muelling, J. Dolan, P. Palanisamy, and P. Mudalige, "Learning vehicle surrounding-aware lane-changing behavior from observed trajectories," in *2018 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2018, pp. 1412–1417.
- [6] Q. Xue, Y. Xing, and J. Lu, "An integrated lane change prediction model incorporating traffic context based on trajectory data," *Transportation research part C: emerging technologies*, vol. 141, p. 103738, 2022.
- [7] K. Gao, X. Li, B. Chen, L. Hu, J. Liu, R. Du, and Y. Li, "Dual transformer based prediction for lane change intentions and trajectories in mixed traffic environment," *IEEE Transactions on Intelligent Transportation Systems*, 2023.
- [8] M. Manzour, A. Ballardini, R. Izquierdo, and M. Sotelo, "Vehicle lane change prediction based on knowledge graph embeddings and bayesian inference," in *2024 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2024, pp. 1893–1900.
- [9] M. M. Hussien, A. N. Melo, A. L. Ballardini, C. S. Maldonado, R. Izquierdo, and M. Ángel Sotelo, "Rag-based explainable prediction of road users behaviors for automated driving using knowledge graphs and large language models," *Expert Systems with Applications*, vol. 265, p. 125914, 2025. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0957417424027817>
- [10] M. Peng, X. Guo, X. Chen, M. Zhu, K. Chen, X. Wang, Y. Wang *et al.*, "Lc-llm: Explainable lane-change intention and trajectory predictions with large language models," *arXiv preprint arXiv:2403.18344*, 2024.
- [11] Y. Yao, X. Wang, M. Xu, Z. Pu, Y. Wang, E. Atkins, and D. Crandall, "Dota: unsupervised detection of traffic anomaly in driving videos," *IEEE transactions on pattern analysis and machine intelligence*, 2022.
- [12] L. Kezebou, V. Oludare, K. Panetta, J. Intriligator, and S. Agaian, "Highway accident detection and classification from live traffic surveillance cameras: a comprehensive dataset and video action recognition benchmarking," in *Multimodal Image Exploitation and Learning 2022*, vol. 12100. SPIE, 2022, pp. 240–250.
- [13] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "CARLA: An open urban driving simulator," in *Proceedings of the 1st Annual Conference on Robot Learning*, 2017, pp. 1–16.
- [14] R. Krajewski, J. Bock, L. Kloecker, and L. Eckstein, "The highd dataset: A drone dataset of naturalistic vehicle trajectories on german highways for validation of highly automated driving systems," in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, 2018, pp. 2118–2125.
- [15] M. Saffarzadeh, N. Nadimi, S. Naseralavi, and A. R. Mamdoohi, "A general formulation for time-to-collision safety indicator," in *Proceedings of the Institution of Civil Engineers-Transport*, vol. 166, no. 5. Thomas Telford Ltd, 2013, pp. 294–304.
- [16] E. Ramezani-khansari, F. M. Nejad, and S. Moogeh, "Comparing time to collision and time headway as safety criteria," *Pamukkale Üniversitesi Mühendislik Bilimleri Dergisi*, vol. 27, no. 6, pp. 669–675, 2020.
- [17] D. Zheng, X. Song, C. Ma, Z. Tan, Z. Ye, J. Dong, H. Xiong, Z. Zhang, and G. Karypis, "Dgl-ke: Training knowledge graph embeddings at scale," in *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2020, pp. 739–748.
- [18] L. Costabello, S. Pai, C. L. Van, R. McGrath, N. McCarthy, and P. Tabacof, "AmpliGraph: a Library for Representation Learning on Knowledge Graphs," Mar. 2019. [Online]. Available: <https://doi.org/10.5281/zenodo.2595043>