

GraphSCENE: On-Demand Critical Scenario Generation for Autonomous Vehicles in Simulation

Efimia Panagiotaki^{1,2}, Georgi Pramatarov¹, Lars Kunze^{2,3}, Daniele De Martini¹

¹Mobile Robotics Group (MRG) and ²Cognitive Robotics Group (CRG), University of Oxford, UK

³Bristol Robotics Laboratory (BRL), University of the West of England, UK

{efimia, georgi, lars, daniele}@robots.ox.ac.uk

Abstract—Testing and validating Autonomous Vehicle (AV) performance in safety-critical and diverse scenarios is crucial before real-world deployment. However, manually creating such scenarios in simulation remains a significant and time-consuming challenge. This work introduces a novel method that generates dynamic temporal scene graphs corresponding to diverse traffic scenarios, *on-demand*, tailored to user-defined preferences, such as AV actions, sets of dynamic agents, and criticality levels. A temporal Graph Neural Network (GNN) model learns to predict relationships between ego-vehicle, agents, and static structures, guided by real-world spatiotemporal interaction patterns and constrained by an ontology that restricts predictions to semantically valid links. Our model consistently outperforms the baselines in accurately generating links corresponding to the requested scenarios. We render the predicted scenarios in simulation to further demonstrate their effectiveness as testing environments for AV agents.

Index Terms—Scenarios generation, graph generation, dynamic temporal scene graphs, heterogeneous graph learning

I. INTRODUCTION

Ensuring the safety of passengers and road users through extensive testing and validation of Autonomous Driving (AD) systems is an important task. Despite the availability of diverse datasets [1], corner-case scenarios – i.e. rare or unique situations [2] – remain a significant challenge for Autonomous Vehicle (AV) systems [3]. These safety-critical scenarios often fall within the long tail of real-world dataset distributions [4], as their unexpected occurrence makes them difficult to capture. Simulation has become an essential tool to identify potential failure points before real-world deployment through extensive AV testing in controllable, complex, diverse, and realistic traffic situations.

Effective scenario generation must capture complex, unsafe, and unpredictable interactions that AVs might encounter. Recent works range from leveraging real-world data to synthesising rare events in simulation, aiming to generate targeted scenarios for systematic safety validation, thus reducing the likelihood of unforeseen failures in AV deployment [5], [6]. However, despite their promising results, data-driven approaches don’t allow situation control, while simulation-based methods lack plausibility [7]. Aiming to

The project was supported by the EPSRC Programme Grant “From Sensing to Collaboration” (EP/V000748/1), the EPSRC project RAILS (EP/W011344/1), and the Oxford Robotics Institute research project Robot-Cycle. Efimia Panagiotaki was supported by the DeepMind Engineering Science Scholarship.

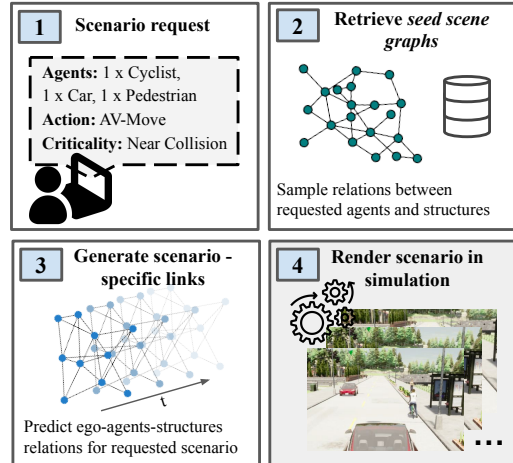


Fig. 1: *On-demand* scenario generation pipeline: (1) User requests a scenario. (2) GraphSCENE samples relations between requested agents and static structures from a database, to create a scene graph used as *seed* to (3) the temporal scene graph generation model, predicting scenario-specific links between the ego-vehicle node and all other graph entities. (4) The output scenario is rendered in simulation.

overcome these limitations, this work proposes a hybrid approach that combines data-driven insights with logical constraints to generate synthetic scenarios in simulation and create more comprehensive testing frameworks.

Addressing the challenges of plausibility and diversity in synthetic datasets, our method generates safety-critical traffic scenarios *on-demand* from user inputs, leveraging real-world data. The user simply requests a scenario and our model generates dynamic temporal scene graphs corresponding to the desired scenario, which are then rendered in CARLA Simulation [8], as seen in fig. 1. In particular, the user defines the dynamic agents, the AV action (move, turn left/right, or overtake) and the scenario’s criticality level (visible, near, or near collision). Then, the relations between the agents – excluding the AV, namely *ego-vehicle* – are sampled from our *database* consisting of a collection of interactions extracted from the ROad event Awareness Dataset (ROAD) [9]. These serve as *seeds* for our graph generation model, which learns to generate temporal links corresponding to relations between the AV and the agents and static structures.

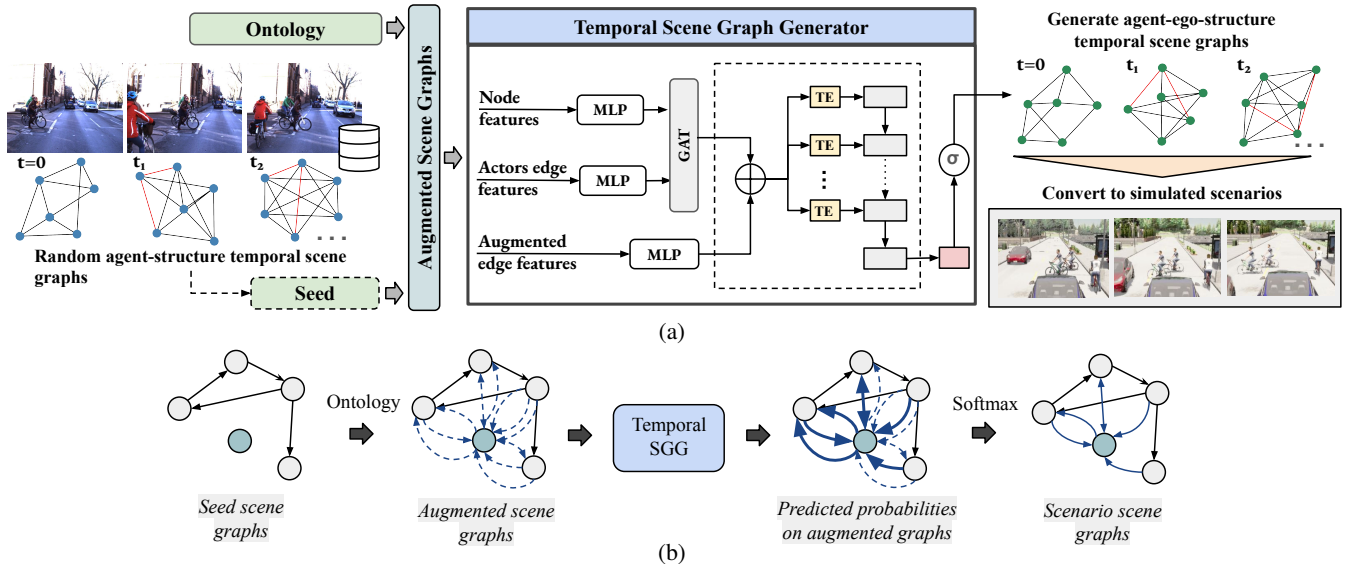


Fig. 2: Overview of our methodology. (a) We first extract temporal scene graphs from the ROAD dataset, capturing relations between agents and static structures, excluding the ego-vehicle. These serve as *seeds* for a temporal scene graph generation model, which encodes heterogeneous features and, through a temporal encoder (TE), predicts relationships that semantically link the ego-vehicle with agents and static structures. The output, consisting of predicted *scenario scene graphs*, models dynamic interactions and is converted into simulation scenarios for AV testing. (b) The underlying progression of link prediction, shown in a sample graph. Here, the blue nodes correspond to the ego-vehicle and grey nodes to the agents and static structures. We first augment the input *seed scene graphs* with all possible relations from the ontology, allowing the temporal scene graph generation model to predict the likelihood of each edge in the augmented graphs. These edges correspond to edges in *scenario scene graphs*.

This process is guided and constrained by an *ontology* to ensure the plausibility of interactions and scene evolution. Our method provides a structured approach in generating complex traffic situations for evaluating AV navigation performance in diverse and complex scenarios.

Our key contributions are summarised as follows:

- A novel end-to-end ontology-based and data-driven approach for *on-demand* generation of safety-critical traffic scenarios in simulation;
- A novel learning methodology for temporal graph generation, capturing semantic, topological, and historical patterns between scene elements to generate realistic dynamic temporal scene graphs;
- A seamless integration of the generated scenarios into simulation, creating complex and diverse testing scenarios for evaluating AV navigation performance.

II. RELATED WORK

Our method tackles *traffic scenario generation* as a *graph generation* task through *temporal graph learning*. This section reviews related work in these three key areas.

Traffic scenario generation. Scenario generation is a key component of applications such as extracting scenarios from naturalistic driving data or adjusting scenario parameters [10]. However, existing methods are often limited in applicability and lack plausibility. ChatScene [7] employs a Large Language Model (LLM) to generate safety-critical scenarios in simulation, but is restricted by user prompts and lacks

real-world grounding. CC-SGG [11] perturbs scene graphs to generate corner cases but relies on a small, simulation-based dataset and doesn't capture temporal patterns. TrafficGen [6] augments scenarios on HD maps and loads them in simulation. However, it relies on artificially generated agents, reducing scenario plausibility. Similarly, real-to-synthetic approaches [12] represent scenes as graphs and generate actors' positions on HD maps to construct traffic scenarios, but do not directly render them in simulation. Rule-based methods [13], [14] generate scenarios via ontologies but often lack relevance to real traffic events. Our method overcomes these limitations by leveraging real-world data to capture complex interactions between scene participants and generate ontology-constrained temporal scene graphs, which are then rendered in simulation.

Graph Generation. Recent advancements in graph generation employ generative, autoregressive, or recurrent models to construct nodes and edges from incomplete input graphs. Variational Autoencoders (VAEs) have been employed to sequentially generate graph elements through a series of probabilistic decisions [15], but have applicability and scalability limitations. In protein structure prediction, [16] introduces an attention-based encoder with an autoregressive decoder, capturing node and edge features but not evolving relations over time. For scene graph generation, SceneGraphGen [17] uses a *seed* object in an autoregressive model but lacks temporal modeling, limiting its suitability for dynamic graphs. GraphRNN [18] treats graph generation as

a sequential task, while GRAN [19] uses recurrent attention mechanisms to predict edges in augmented graphs. Despite strong performance, neither method captures evolving edge attributes over time. In comparison, our approach combines seed initialisation with edge prediction on augmented graphs, while uniquely capturing evolving spatiotemporal patterns. As such, our model captures temporal dependencies and intrinsic structural variations by generating links guided by the input seed graphs.

Temporal Graph Learning. Temporal graph learning has been applied to agent interactions modelling in traffic prediction. STGCN [20] and ST-LGSL [21] combine temporal and spatial aggregations for timeseries prediction, while ST-GAT [22] captures spatiotemporal dependencies using graph attention and an LSTM. Although effective, these methods mainly focus on specific traffic prediction tasks limiting their broader applicability. Additionally, spatiotemporal scene graphs with LSTMs have been employed for explainable action prediction [23], but they primarily capture temporal information at the node level rather than modeling the sequential evolution of events. Building on these foundations, our approach employs temporal Graph Neural Networks (GNNs) to encode traffic agents and model their evolving interactions, effectively generating plausible traffic scenarios.

III. OVERVIEW

Our approach, visualised in fig. 2, uses scene graphs extracted from the ROAD dataset to generate new scenarios that satisfy user-defined requirements, such as configuration of agents, AV action, and scenario criticality. To achieve this, we first extract semantic elements from distinct frames in the ROAD dataset to build two types of graphs: (1) *seed scene graphs* containing evolving relations between dynamic agents (excluding the ego-vehicle) and static structures, forming our *database*, and (2) *scenario scene graphs* including also relational information about the ego-vehicle. The latter serve as ground truth and a weak supervision signal, as the model learns to generate links corresponding to scenario graphs, using the former as *seeds*. Ultimately, our method predicts the relationships between the ego-vehicle and other scene entities to output dynamic scene graphs containing the desired temporal relations between scene elements, facilitating scenario generation in CARLA.

A. Notation and Problem Definition

We define a traffic scenario as $\mathbb{S} : \{S_t \mid t \in [1, N]\}$ consisting of N sequential scenes S_t . Each S_t is represented as a scene graph $G_t = \langle V_t, E_t \rangle$, where nodes V_t represent key dynamic and static entities (e.g. cars, pedestrians, or lanes) and the edges E_t denote their relationships (e.g. distance or topology). Each node $n_t^i \in V_t$ corresponds to a semantic category (e.g., ego-vehicle, road, pedestrian, or cyclist) with a node feature representation h_t^i . Similarly, each edge $e_t^{ij} \in E_t$ represents a semantic relationship (e.g., distance, topology) between two nodes $n_t^i, n_t^j \in V_t$ with an associate edge feature representation p_t^{ij} . Edges can connect

	Type
(1) Agents	EGO, Pedestrian, Car, Cyclist, Motorbike, Bus, TrafficLight
(2) Locations	VehicleLane, OutgoingLane, OutgoingCycleLane, IncomingLane, IncomingCycleLane, Pavement, Junction, PedestrianCrossing, BusStop, Parking
(3) Actions	Move, Brake, Stop, IndicateLeft, IndicateRight, TurnLeft, TurnRight, Cross
(4) AV Actions	AV-Move, AV-MoveLeft, AV-MoveRight, AV-Overtake, AV-Stop, AV-TurnLeft, AV-TurnRight
(5) Criticality	Near Collision, Near, Visible

TABLE I: Entities (1)-(2) and relations (3)-(5) in our ontology forming *node-edge-node* triplets.

two entities or link an entity to itself, forming self-edges. This creates a collection of N graphs, $\mathbb{G} : \{G_t \mid t \in [1, N]\}$ for each scenario \mathbb{S} . In our proposed method, the goal is to learn to generate *scenario scene graphs*, \mathbb{G} , from *seed scene graphs* $\tilde{\mathbb{G}} : \{\tilde{G}_t \mid t \in [1, N]\}$, where each $\tilde{G}_t \subsetneq G_t$. The nodes V_t and predicted relationships in E_t from the generated scene graphs are used to create scenarios in simulation.

IV. METHODOLOGY

Our methodology consists of three main components: the extraction of the *seed* and *scenario scene graphs*, the temporal scene graph generation model, and the scenario integration in simulation. The following subsections detail our system.

A. Scene Graph Extraction

Ontology Definition. We establish a traffic domain *ontology*, shown in table I, structured around the Automotive Urban Traffic Ontology (A.U.T.O.) [24] and adapted for the ROAD dataset [9]. The ontology is used to retrieve nodes and edge labels for each graph G_t , integrating concepts such as object types, traffic rules, relative positions, actions, and behaviours in the graphs. The node labels are derived directly from the ROAD dataset [9] and consist of agents and locations – (1) and (2) in table I, respectively. Agent positions are represented by an `ISIN` relation linking to the locations in (2). Agent states are represented by self-relations, where for dynamic agents the states are in (3), for traffic lights the state is its colour `{red, amber, green}`, and for the ego-vehicle the state is in (4). Relative motion is captured by a relation in `{MovingAway, MovingTowards}` to the ego node. We also include a `MustStop` traffic rule edge for agents, if the traffic light is red. Finally, we model the distance between agents and the ego-vehicle as a discrete proximity relation in (5).

Scenario Graphs. For each frame in ROAD, we extract a graph G_t where nodes represent entities and edges model the relations from the ontology. Their states and relationships are implemented as one-hot-encoded node and edge features. We extract a scenario \mathbb{S} from a sequence of $N = 5$ frames, where each frame corresponds to a scene S_t , that captures a single AV action, as seen in table I (4). For each traffic scenario \mathbb{S} and its corresponding scene graphs \mathbb{G} , we construct a temporal graph $\mathbb{G}^T = (\mathbb{V}^T, \mathbb{E}^T)$ by combining the nodes

and edges from each $G_t \in \mathbb{G}$. The graphs \mathbb{G}^T correspond to the temporal representation of the *scenario scene graphs* \mathbb{G} . Since objects – i.e. agents and static structures – can appear in multiple frames, we assign unique IDs and track them in order to only include unique nodes from \mathbb{G} in \mathbb{G}^T . As multiple edges can be connected to the same node in \mathbb{G}^T , each edge is assigned a temporal label $\tau \in [0, 4]$ depending on the graph G_t it originated from, within the sequence \mathbb{G} . This process is illustrated in fig. 3.

We then label the graph \mathbb{G}^T with the ego-vehicle AV action and determine the scenario’s *criticality* level based on the proximity relations of all agents. Specifically, we assign the criticality level by selecting the most severe proximity relation present in any graph $G_t \in \mathbb{G}$. The ego-vehicle AV action is derived from the last recorded AV action in \mathbb{G}^T . For example, if the ego-vehicle moves forward and then turns left, the scenario is labeled as ‘AV-TurnLeft’. A criticality node with type in table I (5) is added to the graphs linking to the ego-vehicle node through a *criticality* relation. We also add a self-loop to the ego-vehicle node representing the scenario’s AV action, with type in table I (4).

Seed Graphs. From the graphs \mathbb{G}^T , we construct the temporal representation of the *seed scene graphs*, $\tilde{\mathbb{G}}^T$, where $\tilde{\mathbb{G}}^T = (\tilde{\mathbb{V}}^T, \tilde{\mathbb{E}}^T)$. In particular, we are pruning all edges and relations between ego-vehicle and objects in \mathbb{G}^T to generate $\tilde{\mathbb{G}}^T$. The temporal graphs $\tilde{\mathbb{G}}^T$ contain all the nodes and edges from each *seed scene graph* $\tilde{G}_t \in \tilde{\mathbb{G}}$, where each \tilde{G}_t can be easily retrieved from \mathbb{G}^T using the edge temporal labels τ . During training, these graphs are used as *seeds* for our scene graph generation model which learns to generate the *scenario scene graphs*. To build our agents-structures relations *database*, we also remove the AV action and criticality components from the temporal *seed scene graphs* before storing them in our database.

B. Temporal Scene Graph Generation

We design a temporal GNN model that learns to generate *scenario scene graphs* by predicting edges connecting the ego-vehicle to agents and static structures, conditioned on a requested AV action and criticality level. We approach this as a link prediction task, where our model learns to predict edges corresponding to the requested scenario in ontology-augmented scene graphs. Our method builds upon recent link prediction studies, introducing a more refined approach that predicts the existence of multiple temporal edges and relationships between nodes, guided by a structured ontology, instead of relying on exhaustive or broad search strategies that find single links between node pairs in graphs [25].

Scene Graphs Augmentation. At the first step of our method, we augment the temporal *seed scene graphs* by adding all possible links and associated edge attributes between the ego-vehicle and the agents and static structures, from our ontology. These edges are added for every timestamp and are labelled with a corresponding temporal label τ in the augmented $\tilde{\mathbb{G}}^T$. The links between agents-structures remain unchanged. Our model is explicitly trained to predict only the links involving the ego-vehicle in order to construct

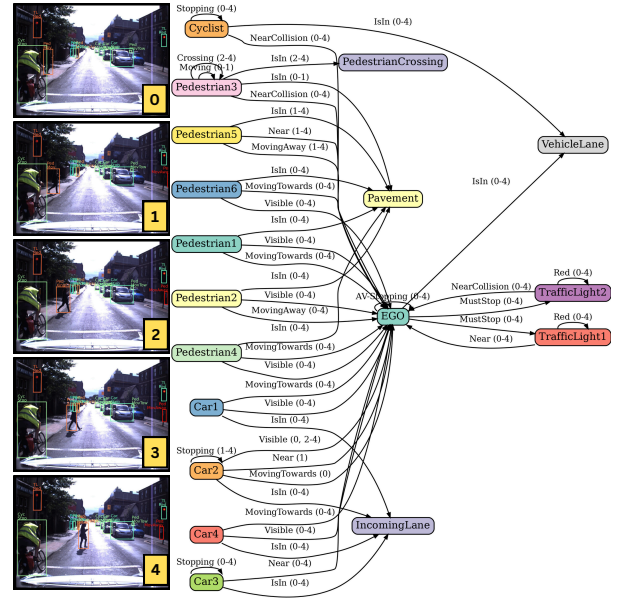


Fig. 3: Example of scene graph extraction. A sequence of frames \mathbb{S} is annotated (left) and a temporal scene graph \mathbb{G}^T is constructed by linking the entities (right) using the ontology.

the desired *scenario scene graphs* by learning to estimate the likelihood of each edge. The new edges are denoted as $e_e \in \tilde{\mathbb{E}}^T$ with associated feature representations p_e . The original agents-structures links are denoted as $e_a \in \mathbb{E}^T$ with associated feature representations p_a . We note that both the *criticality* node, connected via the criticality edge, and the *action* edge, represented by a self-loop, are linked to the ego-vehicle. As special cases within the set of edges in e_a , both the criticality and action feature vectors are encoded into the ego-vehicle’s feature vector, but the corresponding edges are not being predicted. The augmented scene graphs are passed as input to our model.

Feature Encoders. To preserve the structural and spatial information about the nodes, each of them is assigned a sinusoidal positional encoding [26], calculated using its unique ID, to capture its global position within the temporal graphs. This allows the network to distinguish between entities of the same class and track their evolving patterns across time.

Each input graph entity is encoded to generate compact feature vector representations. We apply a set of Multi-Layer Perceptrons (MLPs) to both the node and edge features of the augmented scene graphs to learn their corresponding embedding representations. In particular, node features are encoded as $z_n = MLP_n(h)$, actors-structures edge features as $z_a = MLP_e(p_a)$, and the ego-vehicle-related edge features as $z_e = MLP_e(p_e)$. We then apply two layers of the Graph Attention Network (GAT) [27] to aggregate the feature vectors of neighbouring nodes and corresponding edge features in z_a to generate query node embeddings, $z'_n = GAT(z_n, z_a, e_a)$.

Temporal Layers. The model then sequentially processes the information, guided by the temporal labels τ in the edges, to capture complex temporal relationships and the evolution of interactions between all scene participants.

Our temporal encoder receives *node-edge-node* triplets in the input, grouped in sets based on the label τ in each corresponding edge in the form of $\langle z'_{nk}, z_e, z'_{nl} \rangle_\tau$, where k, l are node indices guided by the respective edges. The triplet sets are sequentially passed to the triplet encoder, retaining the order from τ . The encoder consists of a single MLP layer that encodes the concatenated feature vectors, $g_{tr} = MLP_{tr}(z'_{nk}|z_e|z'_{nl})_\tau$ for each set sequentially. In between each sequence, we apply a Graph Convolutional Network (GCN) layer to further aggregate node feature vectors, capturing both temporal and relational patterns. The updated representations are then propagated to the next step, expressed as $z'_n = GCN(z'_n, e_e)$.

Link Prediction. After the last iteration of the temporal module, we apply a Sigmoid function to the triplet embeddings g_{tr} to extract the probability of each triplet, and consequently the likelihood of each edge and edge feature existing in the query *scenario scene graph*, as $\hat{y} = \sigma(g_{tr})$. During training, we use a Binary Cross Entropy (BCE) loss averaged over all samples as follows:

$$\mathcal{L}_{BCE}(y, \hat{y}) = -(y \cdot \log(\hat{y}) + (1 - y) \cdot \log(1 - \hat{y})) \quad (1)$$

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N \mathcal{L}_{BCE}(y_i, \hat{y}_i) \quad (2)$$

where $y \in \{0, 1\}$ is the ground-truth label and $\hat{y} \in [0, 1]$ is the predicted triplet likelihood. In the final step, we apply softmax filtering to retain only the highest probability relation from mutually exclusive edges as defined in the ontology, such as `{MovingAway}` and `{MovingTowards}`. This ensures that the generated graphs remain plausible by preventing contradictions like an entity being classified as *both* moving towards and moving away from the ego-vehicle.

C. Scenario Integration in Simulation

Inference. During inference, a user requests a scenario specifying dynamic agents, AV action, and criticality level. Based on the requested agents, we sample agents-structures relations from our database to build *seed scene graphs*. The specified AV action and criticality are then reintroduced into the graphs, ensuring they retain a structure similar to the *seed scene graphs* used during training. In particular, we add the AV action as ego-vehicle self-edge and add the criticality as a node linking to the ego-vehicle node. During inference, these graphs are passed in the input of our temporal scene graph generation model which, also constrained by the ontology, generates links corresponding to the requested scenario.

Simulation. The nodes and predicted relationships are then extracted from the graphs at each τ to build the requested scenarios in CARLA [8] using the OpenSCENARIO¹ standard. The behaviour of the scene participants and the scenario parameters are defined by the scene graphs. We use the relative position and action edges to define the lane and orientation of each agent in the simulation with respect to the ego-vehicle. The initial position of each actor

is defined by the scenario’s criticality and corresponds to the distance from the ego-vehicle. For each actor, the start and goal positions and orientations are determined by the lane position edges in the graphs, ensuring realistic movement and positioning throughout the scenario.

V. EXPERIMENTAL SETUP

A. Dataset

For this work, we used the ROAD dataset [9], an extension of the Oxford RobotCar Dataset [28], to train and evaluate our approach, as well as build our agents-structures database. This dataset is designed to test an AV’s ability to detect various road events. Video sequences were downsampled 5-fold from 12 frames per second (fps) to increase variation across 5-frame sequences, similar to [23]. Each sequence corresponds to a specific action of the ego-vehicle, defined as AV Action in table I (4). We extract the proximity relations between actors and the ego-vehicle by computing and thresholding the box-wise median Relative Inverse Depth (RID) obtained from an off-the-shelf Midas (DPT Large) network [29], with Near Collision ($< 5 m$), Near ($5 - 10 m$), and Visible ($> 10 m$). We generate 6,771 temporal graphs grouped by action, where each action group is split into training, validation, and testing sets using a 70-20-10 allocation for balanced evaluation.

B. Training

To ensure stable learning and effective generalisation, we carefully tuned hyperparameters and employed several training techniques. For hyperparameters tuning, we used k-fold cross-validation with $k = 5$. The dimensions of the MLP layers in our model were configured as follows: $MLP_n : [22, 8, 1]$, $MLP_e : [27, 8, 1]$, $MLP_{tr} : [3, 16, 1]$, and $MLP_{gat} = [64, 128, 256]$. The first GAT layer used dimensions $[1, 64]$, followed by another GAT layer $[256, 1]$. Finally, the GCN was set with dimensions $[1, 1]$ to capture recurrent node aggregations. This layered structure allowed to process efficiently both node and edge features and generate the required triplet embedding representations. We additionally used a sinusoidal positional encoding of length 2, the Adam optimiser with 0.01 learning rate, and weight decay with $\lambda = 10^{-5}$. Gradient clipping at ± 1 prevented exploding gradients, while Xavier initialisation helped stabilise weight distributions across layers.

C. Baselines

We evaluated our method against three key baselines.

- 1) F-S LLM: Inspired by ChatScene [7], we evaluated our method against a Few-Shot In-Context Learning (ICL) setup using GPT-4 [30]. We first converted our ontology along with the nodes, edges, and attributes of the *seed scene graphs* in our testing dataset into lists. We then provided the model with these lists and an example response within the prompt, asking it to find plausible relationships between scene entities to construct query scenarios, given our ontology and existing relationships. This comparison assessed the

¹asam.net/standards/detail/openscenario/

strengths and limitations of LLMs in structured data prediction, particularly in scene graphs.

- 2) GRAN⁺: We employ the model architecture proposed in GRAN [19] and adapt it for our task by incorporating node and edge feature encoders while preserving the attention encoder and GRU. We adjust the prediction mechanism to our triplet prediction task and train the model in our pipeline. Among graph generation models, GRAN is the most comparable to our approach as it predicts links from augmented graphs and its architecture facilitates guiding the generation process through ontology-based constraints.
- 3) CC-SGG [11]: A non-temporal GNN model that learns to predict links corresponding to static corner case scene graphs. We adapt it for our task and train the model on our dataset to predict the entire sequence in a single step. This comparison highlights the advantage of our temporal GNN method over static GNNs in capturing evolving patterns in the input.

VI. RESULTS

A. Classification Performance

We evaluate our link prediction method using standard classification metrics: F1 score, accuracy, precision, and recall, as seen in table II. Overall, our model outperforms the baselines achieving the highest **F1 score 0.706**, and **recall 0.859**, and comparable accuracy and precision. Our model achieves the optimal balance in its predictions, effectively minimising false negatives while maintaining strong overall performance.

	F1	Accuracy	Precision	Recall
F-S LLM	0.218	0.436	0.212	0.262
GRAN ⁺	0.647	0.752	0.663	0.704
CC-SGG	0.583	0.801	0.729	0.510
Ours	0.706	0.777	0.623	0.859

TABLE II: Link prediction accuracy across the baselines.

B. Contextual Evaluation

We additionally conduct a contextual evaluation of the output predictions using the *answer correctness* metric from the Retrieval Augmented Generation Assessment (RAGAs) framework [31], designed to measure the accuracy of an LLM’s answer to a prompt compared to a ground-truth. To achieve this, we convert the predicted and ground-truth graphs into textual descriptions of the respective scenes and compute the corresponding *answer correctness*, obtaining a score between 0 and 1. This metric combines both semantic and factual similarity between the predicted and ground truth descriptions without requiring additional context. The former captures the overall resemblance between them, while the latter computes their factual overlap, i.e. comparing the individual statements. To obtain such descriptions of our temporal graphs, we produce a paragraph for each constituent graph, including its timestamp and a statement for each relation. For example, the graph at frame 42 with

relations $\{(Pedestrian, Near, EGO), \dots\}$ becomes "At time 42: Pedestrian 1 is near the ego-vehicle. ...". This relation-to-statement mapping aligns well with the factual similarity component, providing a contextual evaluation of our predicted graphs’ structure.

The *answer correctness* results per AV action and criticality are presented in tables III and IV. Our model achieves the highest *answer correctness* scores across almost all evaluations, demonstrating its strong ability to generate both accurate and semantically meaningful relations. In contrast, the lower scores of the LLM-based method justify the benefits of model-based approaches, such as ours, even in the context of semantic understanding.

Method	AV Actions							Avg
	Mov	MovL	MovR	Ovtk	Stop	TurnL	TurnR	
F-S LLM	0.668	0.721	0.478	0.761	0.55	0.619	0.689	0.64
GRAN ⁺	0.735	0.823	0.667	0.715	0.763	0.733	0.878	0.746
CC-SGG	0.709	0.615	0.814	0.732	0.805	0.771	0.883	0.735
Ours	0.758	0.862	0.972	0.776	0.769	0.761	0.831	0.765

TABLE III: Contextual evaluation - per AV action.

This approach, alongside the classification evaluation described above, allowed us to measure not only how well each model performed but also how accurate and contextually appropriate the predictions were.

Method	Criticality		
	Visible	Near	Near Collision
F-S LLM	0.674	0.631	0.596
GRAN ⁺	0.735	0.823	0.667
CC-SGG	0.794	0.694	0.675
Ours	0.835	0.736	0.68

TABLE IV: Contextual evaluation - per scenario criticality.

C. Scenario Generation

To evaluate the suitability of our generated simulation scenarios for AV testing, we assess the performance of 3 CARLA agents – *normal*, *cautious*, and *aggressive* – provided by the simulator library. Each agent is tested in 30 randomly selected scenarios with evenly distributed criticality levels. We compute the Scenario Consistency Rate (SCR) (%), measuring whether the agents’ final states align with each scenario’s criticality, as shown in table V. Our results demonstrate that our method effectively generates diverse scenarios. The behaviour of the CARLA agents in those scenarios mostly leads to situations matching the intended criticality, especially for *visible* and *near* scenarios. Interestingly, the behaviour of the *aggressive* agent consistently led to more dangerous situations than expected, even in lower-criticality scenarios like *visible* and *near*. We also note a lower SCR percentage for *near collision* scenarios, likely due to the fact that the agents in CARLA have access to the perfect traffic state – i.e. knowing the state of all vehicles and pedestrians – making it easier to avoid collisions.

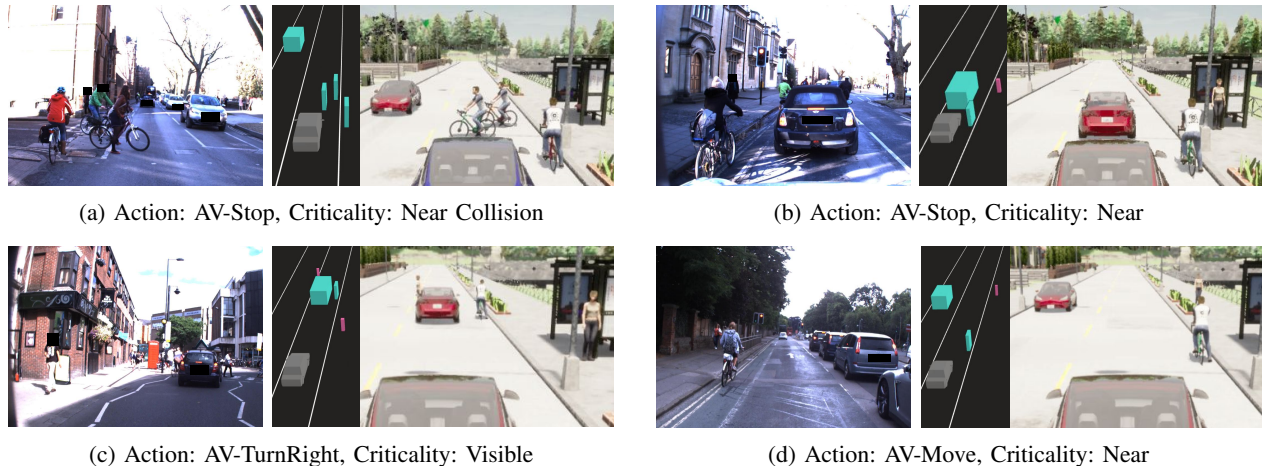


Fig. 4: Qualitative evaluation of the generated scenarios in CARLA using sampled *seed scene graphs* from the ROAD dataset [9] for various actions and criticalities. As the dataset was recorded in a left-hand drive country, we mirrored the topological information to align with the configuration used in CARLA, a right-hand drive environment.

Agents	Criticality		
	Visible	Near	Near Collision
Normal	70	80	40
Cautious	80	80	50
Aggressive	60	60	50

TABLE V: Performance evaluation of 3 CARLA agents, measured in SCR (%) for assessing alignment with scenario criticality. Predicted scene graphs define the behaviours and positions of non-ego-vehicle entities, as the ego-vehicle navigates the generated scenario.

D. Qualitative Results

We qualitatively assess our system’s ability to generate novel scenarios from random user inputs. Here, we compare the agent’s positions and behaviours in the generated scenarios in CARLA against the agent’s positions and behaviours in the scenarios in the ROAD dataset from which the sampled *seed scene graphs* originated, depicted in fig. 4 which also illustrates the actors’ positions on the map. In our scenarios, actors and associated relational information are successfully generated in simulation and actor behaviour generally aligns with expected criticality levels, as seen in fig. 4 (a,c,d). We observe that, in some cases, the ‘Near’ collision criticality results in more severe situations than expected, e.g. in fig. 4 (b), due to small distance threshold margins for *near* and *near collision* scenarios. To improve distinction, we manually adjust these thresholds in simulation.

VII. ABLATION

We evaluate several models as candidate triplet encoders, exploring both temporal and non-temporal architectures.

Temporal Triplet Encoders. The goal of this ablation study is to find the optimal model for the temporal triplet encoder. Our results, shown in table VI, indicate that, from the candidate architectures, the GRU with GCN demonstrates promising performance but imbalance between precision and

recall. The MLP with GCN encoder consistently outperforms the others and thus we select it for our method.

Model	F1	Accuracy	Precision	Recall
GRU	0.431	0.773	0.504	0.401
GRU/GCN	0.655	0.732	0.559	0.835
GAT	0.518	0.380	0.380	1.000
Ours (MLP/GCN)	0.706	0.777	0.623	0.859

TABLE VI: Ablation study for various triplet encoders.

Non-Temporal Triplet Encoders. We also assess the performance of non-temporal triplet encoders for this task. Our temporal encoder achieves higher F1 score than non-temporal encoders, indicating a well-balanced trade-off between precision and recall as seen in table VII. We also note that the non-temporal MLP, which also showed promising results, is very similar to the encoder used in CC-SGG, and, as such, its performance was previously evaluated in section VI.

	F1	Accuracy	Precision	Recall
GRU	0.319	0.420	0.246	0.502
GAT	0.518	0.380	0.380	1.000
MLP	0.675	0.784	0.644	0.745
Ours	0.706	0.777	0.623	0.859

TABLE VII: Non-Temporal Triplet Encoders compared to our temporal encoder.

Positional Encodings. We also evaluate the impact of the positional encodings in our graphs. For this, we employ a dataset containing actions with lower variability, *AV-Move*, *AV-Stop*, *AV-TurnLeft*, *AV-TurnRight*, and build graphs with a simpler architecture, where same-type actors were mapped to the same node in the graph, and without positional encodings. We evaluate the performance of all baselines on these graphs in table VIII. While our model achieves strong recall, effectively identifying true positive instances, its low F1 score indicates unbalanced predictions. Across almost all metrics, baseline models exhibit consistently poor

performance. These results support the argument for more descriptive node features to improve differentiation and tracking across actors.

	F1	Accuracy	Precision	Recall
F-S LLM	0.207	0.427	0.204	0.249
GRAN ⁺	0.524	0.605	0.506	0.605
CC-SGG	0.610	0.795	0.679	0.586
Ours	0.669	0.702	0.587	0.818

TABLE VIII: Performance without positional encodings.

VIII. CONCLUSION

Our proposed methodology effectively generates diverse scenarios in simulation by modelling complex semantic interactions through temporal scene graphs. We evaluated the accuracy and contextual relevance of our model’s predictions, achieving superior performance to the baseline methods. Additionally, we assessed the performance of 3 navigation agents in our generated scenarios, demonstrating their suitability as testing environments. Our results further indicate that our approach can generate scenarios *on-demand*, achieving the desired criticality levels. Overall, our method establishes a robust framework for dynamic scenario generation, leveraging temporal graphs and GNNs to enhance AV testing and evaluation in diverse scenarios.

ACKNOWLEDGEMENTS

We thank Valentina Muşat for reviewing the paper prior to submission, and Harry Way and Pawit Kochakarn for their work on the early stages of the scene graph pipelines.

REFERENCES

- [1] W. Liu, Q. Dong, P. Wang, G. Yang, L. Meng, Y. Song, Y. Shi, and Y. Xue, “A survey on autonomous driving datasets,” in *2021 8th International Conference on Dependable Systems and Their Applications (DSA)*. IEEE, 2021, pp. 399–407.
- [2] J. Breitenstein, J.-A. Termöhlen, D. Lipinski, and T. Fingscheidt, “Systematization of corner cases for visual perception in automated driving,” in *2020 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2020, pp. 1257–1264.
- [3] H. Sun, S. Feng, X. Yan, and H. X. Liu, “Corner case generation and analysis for safety assessment of autonomous vehicles,” *Transportation research record*, vol. 2675, no. 11, pp. 587–600, 2021.
- [4] W. Zhou, Z. Cao, Y. Xu, N. Deng, X. Liu, K. Jiang, and D. Yang, “Long-tail prediction uncertainty aware trajectory planning for self-driving vehicles,” in *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, 2022, pp. 1275–1282.
- [5] S. Tan, B. Ivanovic, X. Weng, M. Pavone, and P. Kraehenbuehl, “Language conditioned traffic generation,” in *7th Annual Conference on Robot Learning*, 2023. [Online]. Available: <https://openreview.net/forum?id=PK2debCKaG>
- [6] L. Feng, Q. Li, Z. Peng, S. Tan, and B. Zhou, “Trafficgen: Learning to generate diverse and realistic traffic scenarios,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 3567–3575.
- [7] J. Zhang, C. Xu, and B. Li, “Chatscene: Knowledge-enabled safety-critical scenario generation for autonomous vehicles,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 15 459–15 469.
- [8] 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.
- [9] G. Singh, S. Akrigg, M. Di Maio, V. Fontana, R. J. Alitappeh, S. Saha, K. Jeddissaravi, F. Yousefi, J. Culley, T. Nicholson *et al.*, “Road: The road event awareness dataset for autonomous driving,” *IEEE Transactions on Pattern Analysis & Machine Intelligence*, feb.

- [10] B. Schütt, J. Ransiek, T. Braun, and E. Sax, “1001 ways of scenario generation for testing of self-driving cars: A survey,” in *2023 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2023, pp. 1–8.
- [11] G. Drayton, E. Panagiotaki, D. Omeiza, and L. Kunze, “Cc-sgg: Corner case scenario generation using learned scene graphs,” 2024. [Online]. Available: <https://arxiv.org/abs/2309.09844>
- [12] Y. Tian, A. Carballo, R. Li, and K. Takeda, “Real-to-synthetic: Generating simulator friendly traffic scenes from graph representation,” in *2022 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2022, pp. 1615–1622.
- [13] F. Klück, Y. Li, M. Nica, J. Tao, and F. Wotawa, “Using ontologies for test suites generation for automated and autonomous driving functions,” in *2018 IEEE International symposium on software reliability engineering workshops (ISSREW)*. IEEE, 2018, pp. 118–123.
- [14] D. Bogdoll, S. Guneshka, and J. M. Zöllner, “One ontology to rule them all: Corner case scenarios for autonomous driving,” 2022.
- [15] Y. Li, O. Vinyals, C. Dyer, R. Pascanu, and P. Battaglia, “Learning deep generative models of graphs,” 2018. [Online]. Available: <https://arxiv.org/abs/1803.03324>
- [16] J. Ingraham, V. Garg, R. Barzilay, and T. Jaakkola, “Generative models for graph-based protein design,” in *Advances in Neural Information Processing Systems*, vol. 32. Curran Associates, Inc., 2019.
- [17] S. Garg, H. Dhama, A. Farshad, S. Musatian, N. Navab, and F. Tombari, “Unconditional scene graph generation,” in *IEEE International Conference on Computer Vision (ICCV)*, 2021.
- [18] J. You, R. Ying, X. Ren, W. L. Hamilton, and J. Leskovec, “Graphrnn: Generating realistic graphs with deep auto-regressive models,” vol. 80. Proceedings of Machine Learning Research, 2018, pp. 5694–5703.
- [19] R. Liao, Y. Li, Y. Song, S. Wang, C. Nash, W. L. Hamilton, D. Duvenaud, R. Urtasun, and R. Zemel, “Efficient graph generation with graph recurrent attention networks,” in *NeurIPS*, 2019.
- [20] B. Yu, H. Yin, and Z. Zhu, “Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting,” in *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*, ser. IJCAI-2018. International Joint Conferences on Artificial Intelligence Organization, Jul. 2018, p. 3634–3640.
- [21] J. Tang, T. Qian, S. Liu, S. Du, J. Hu, and T. Li, “Spatio-temporal latent graph structure learning for traffic forecasting,” in *2022 International Joint Conference on Neural Networks (IJCNN)*, 2022, pp. 1–8.
- [22] C. Zhang, J. J. Q. Yu, and Y. Liu, “Spatial-temporal graph attention networks: A deep learning approach for traffic forecasting,” *IEEE Access*, vol. 7, pp. 166 246–166 256, 2019.
- [23] P. Kochakarn, D. De Martini, D. Omeiza, and L. Kunze, “Explainable action prediction through self-supervision on scene graphs,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 1479–1485.
- [24] L. Westhofen, C. Neurohr, M. Butz, M. Scholtes, and M. Schuldes, “Using ontologies for the formalization and recognition of criticality for automated driving,” *IEEE Open Journal of Intelligent Transportation Systems*, vol. 3, pp. 519–538, 2022.
- [25] X. Wang, H. Ji, C. Shi, B. Wang, Y. Ye, P. Cui, and P. S. Yu, “Heterogeneous graph attention network,” in *The world wide web conference*, 2019, pp. 2022–2032.
- [26] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017, p. 6000–6010.
- [27] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, “Graph Attention Networks,” *International Conference on Learning Representations*, 2018, accepted as poster. [Online]. Available: <https://openreview.net/forum?id=rJXMpikCZ>
- [28] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, “1 Year, 1000km: The Oxford RobotCar Dataset,” *The International Journal of Robotics Research (IJRR)*, vol. 36, no. 1, pp. 3–15, 2017.
- [29] R. Ranftl, K. Lasinger, D. Hafner, K. Schindler, and V. Koltun, “Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 3, 2022.
- [30] OpenAI *et al.*, “Gpt-4 technical report,” 2024. [Online]. Available: <https://arxiv.org/abs/2303.08774>
- [31] S. Es, J. James, L. Espinosa Anke, and S. Schockaert, “RAGAs: Automated evaluation of retrieval augmented generation,” in *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, Mar. 2024, pp. 150–158.