Understanding Vulnerable Road User Behavior using Spatio-Temporal Knowledge Graphs

He Tan[1,*], Erick Escandon Bailon[1]

[1]Department of Computing, School of Engineering, Jönköping University, Sweden.

Abstract

Understanding vulnerable road user (VRU) behavior is critical for designing safer and more inclusive urban infrastructure. This paper presents a structured and explainable framework using Spatio-Temporal Knowledge Graphs (STKGs) to model and analyze VRU crossing behaviors. By constructing knowledge graphs from real-world urban traffic datasets, we capture dynamic interactions between pedestrians, cyclists, and vehicles in different crossing scenarios. Through query-based analysis, our approach extracts insights that provide decision support to traffic engineers, addressing key safety-related concerns such as unsafe crossing patterns, pedestrian-vehicle interactions, and speed-related risks. This work contributes to responsible AI by enabling transparent, explainable, and data efficient decision-making support for real-world traffic planning and infrastructure design.

Keywords

Spatio-Temporal Knowledge Graph, Knowledge Graph for AI, Semantic Representation of Behavior, Ontology for Road User Behavior, Vulnerable Road User

1. Introduction

Building an active mobility environment is crucial to foster more sustainable and inclusive urban life [1, 2]. Vulnerable road users (VRUs), such as pedestrians (including children, the elderly, and disabled individuals) and cyclists, are integral to the dynamics of urban traffic. They play an essential role in establishing an active mobility environment. Implementing safer road infrastructure not only safeguards these vulnerable groups but also encourages a broader demographic to participate in active mobility. Municipalities, in particular, have both responsibility and opportunity to enhance road infrastructure design to support this goal. To develop safer and more inclusive road infrastructure while optimizing costs, municipalities must transition from today's subjective, assumption-based decision-making to objective, data-driven strategies. Leveraging artificial intelligence (AI) methods can provide deeper insights and key indicators, establishing a robust basis for informed decision making in infrastructure design choices [3].

Crossing behavior is one of the main aspects of VRUs behavior. It has been studied in numerous research studies. It involves the actions and movements of VRUs while crossing streets or roadways, often guided by traffic signals, road markings, and traffic conditions. However, it is extremely challenging to understand the behavior of VRUs [4]. Unlike vehicles that follow structured traffic rules and exhibit relatively predictable and linear motion patterns,

Knowledge Graphs for Responsible AI (KG-STAR) Workshop at ESWC 2025

*Corresponding author.

VRUs are influenced by a variety of external and internal factors and exhibit highly uncertain and non-deterministic motion.

Traditionally, stochastic models, linear regression models, and discrete choice models are used to understand how pedestrians make crossing decisions based on various factors related to traffic conditions, traffic controls, and traffic regulations [5]. These models were developed using self-reported data from interviews and/or questionnaires and observational data from manually screened video recordings. Another approach involves agent-based models, where road users are represented as intelligent agents making rational decisions in uncertain and complex environments. These studies have focused on modeling pedestrians' collision avoidance mechanism. The rules for building the models are often derived from survey data [6].

More recently, trajectory prediction has become a common approach to understand road user behavior [7]. Trajectory prediction involves predicting the future positions and movement patterns of road users over time [8]. Accurate trajectory prediction allows traffic engineers to gain insights into how road users move through road infrastructure and interact with dynamic environments. Deep neural networks (DNN) based models has revolutionized trajectory prediction [9, 10, 11]. They enable the direct learning of complex representations from large datasets and support long-term future predictions. Recent advancements have focused on capturing spatial features, particularly relationships and interactions within dynamic environments, using graph-based models such as graph neural networks (GNNs), graph convolutional networks (GCNs), and graph attention networks (GATs) (e.g., [12, 13]). Furthermore, the spatial-temporal graph transformer (STGT) combines the strengths of graph-based models and transformer architectures to handle both spatial and temporal aspects of trajectory prediction (e.g., [14, 15, 16]).

Although deep learning (DL) has significantly advanced the modeling of pedestrian crossing behavior, it often requires large-scale and high quality datasets for effective training. Moreover, these methods primarily focus on prediction rather than providing structured, interpretable representations. To enhance model transparency and trustworthiness, explainable AI (XAI) techniques are important, enabling traffic engineers to interpret model outcomes and apply insights in real-world decision making scenarios.

Our work is to address this gap by leveraging a knowledge graph based approach that enable explicit representing and reasoning about crossing behaviors. In this paper, we present a structured and explainable framework for modeling and analyzing VRU crossing behaviors using Spatio-Temporal Knowledge Graphs (STKGs). By employing a structured spatial-temporal representation, we constructed knowledge graphs that capture the dynamic interactions during crossings, based on real-world urban road user dynamics data. Through query-based analysis, we extract insights to support traffic engineers in addressing safety-related concerns. This method offers a structured, explainable, and data-efficient approach to modeling and understanding crossing behavior and can serve as a foundation for a decision support system for traffic engineer.

The remainder of this paper is organized as follows: In Section 2, we introduce the methods employed to understand human activities, particularly VRU behavior, using semantic representation, ontology, and knowledge graphs. In Section 3, we present a crossing scene ontology. We describe the knowledge graphs built using the ontology from road user dynamics data in Section 4. Section 5 presents the insights gained from the knowledge graphs that are relevant to traffic engineer's safety concerns. Finally, Section 6 concludes the paper.

2. Related Work

Human activity is a spatial-temporal evolution of interactions [17]. In 1970, Hägerstrand [18] introduced the concept of a time-space path in understanding human activities, which established a foundation for trajectory modeling. Inspired by Hägerstrand's work, Orellana and Renso [19] developed an interaction ontology. The ontology conceptualizes the characteristics of pedestrian movement behavior. It has focused on identifying various movement patterns from time-space paths, and the different categories of interactions, spatial and temporal contexts, behaviors, and the high-level relations between these concepts. Logic-based reasoning is used to categorize pedestrian movement behavior based on its movement patterns, interactions, and contexts.

Chai et al. [20] utilized fuzzy logic to model the cognition and behavioral patterns of pedestrians, in order to understand the effect of age and gender when pedestrians are crossing a signalized crosswalk and jaywalking. Gharebaghi et al. [21] developed a mobility ontology for people with motor disabilities (PWMD). Specifically, it considers the interactions between people and both the social and physical environment. The ontology was used to support the development of assistive technologies for the mobility of PWMD. Fang et al. [22] developed an ontology defining various kinds of road users, including pedestrians, and describing their relationships. The concepts from the ontology are used to define the rules for describing the interactions between road users and to support rule-based reasoning for predicting road users' behavior.

In cognitive science and neuroscience, it has been recognized that segmentation is a fundamental component of perception, playing a critical role in understanding activities. People tend to perceive ongoing continuous activity as series of discrete events (or called segments) [23, 24, 25]. The relationships between segments are encoded in partonomic hierarchies [26]. Coarse segmentation is often related to objects' locations and their goals, and the causal relations between their actions. Fine segmentation is closely linked to changes in the interactions between objects [27].

Building on these findings in cognitive science and neuroscience, Ji et al. [28] proposed a spatial-temporal scene graph to represent human activity and to improve the performance of action recognition and few-shot action recognition using neural networks. Mlodzian et al. [29] presented an ontology that was tailored for representing entities and their spatial and temporal relations in traffic scenes in the nuScenes dataset[1]. A knowledge graph was constructed from the nuScenes dataset using the ontology and provided as a benchmark dataset for developing advanced trajectory prediction models. These studies within computer vision have suggested that a structured spatial-temporal representation can lead to more accurate human activity understanding and improve the performance of various computer vision tasks.

In our previous work [30], we presented a semantic representation of crossing behavior. The representation captures the dynamic evolution of interactions between road users and objects within the physical environment over time in both spatial and temporal dimensions. The representation is generalizable and can be applied to represent the behavior of road user in various traffic scenarios. We have also demonstrated that knowledge graphs can be constructed from road user dynamics data using the representation, and the queries over the knowledge

[1]https://www.nuscenes.org/

graphs can be constructed to answer safety related questions on pedestrian crossing behavior for traffic engineers.

In this work, we extended the semantic representation to incorporate additional factors relevant to VRUs' crossing behavior. We constructed knowledge graphs from two road user dynamics datasets collected in two school areas in Jönköping municipality, Sweden. In collaboration with traffic engineers from the City planning, Development and Traffic Department of Jönköping municipality, we leveraged insights gained from the knowledge graphs to improve traffic engineers' understanding of crossing behavior.

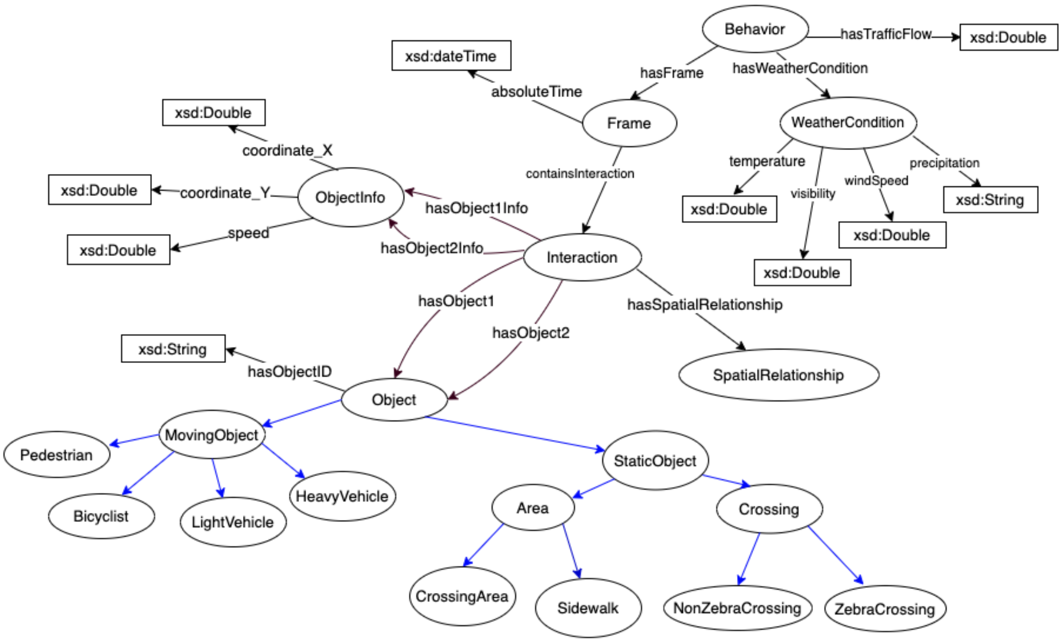

Figure 1: Crossing Scene Ontology (CSO)

3. Crossing Scene Ontology

In this section, we present the ontology describes the semantic representation for crossing behavior (see Fig. 1). A crossing behavior can be seen as a dynamic evolution of interactions between VRUs and other objects within the physical environment over time. Every crossing behavior can be broken down into segments, each representing a distinct phase of the behavior. These segments capture the changes of the interactions between VRUs and objects in both physical and temporal dimensions, and together represent the crossing behavior. For example, Fig. 2 shows a crossing scene, which is extracted from road user behavior measurement performed at a zebra crossing in a school area in Jönköping, Sweden. Fig 2-a displays the trajectories of the VRUs and other moving objects involved in the event. In this event, a cyclist meets a light vehicle at the crossing. The blue trajectory represents a cyclist, and the cyan trajectory

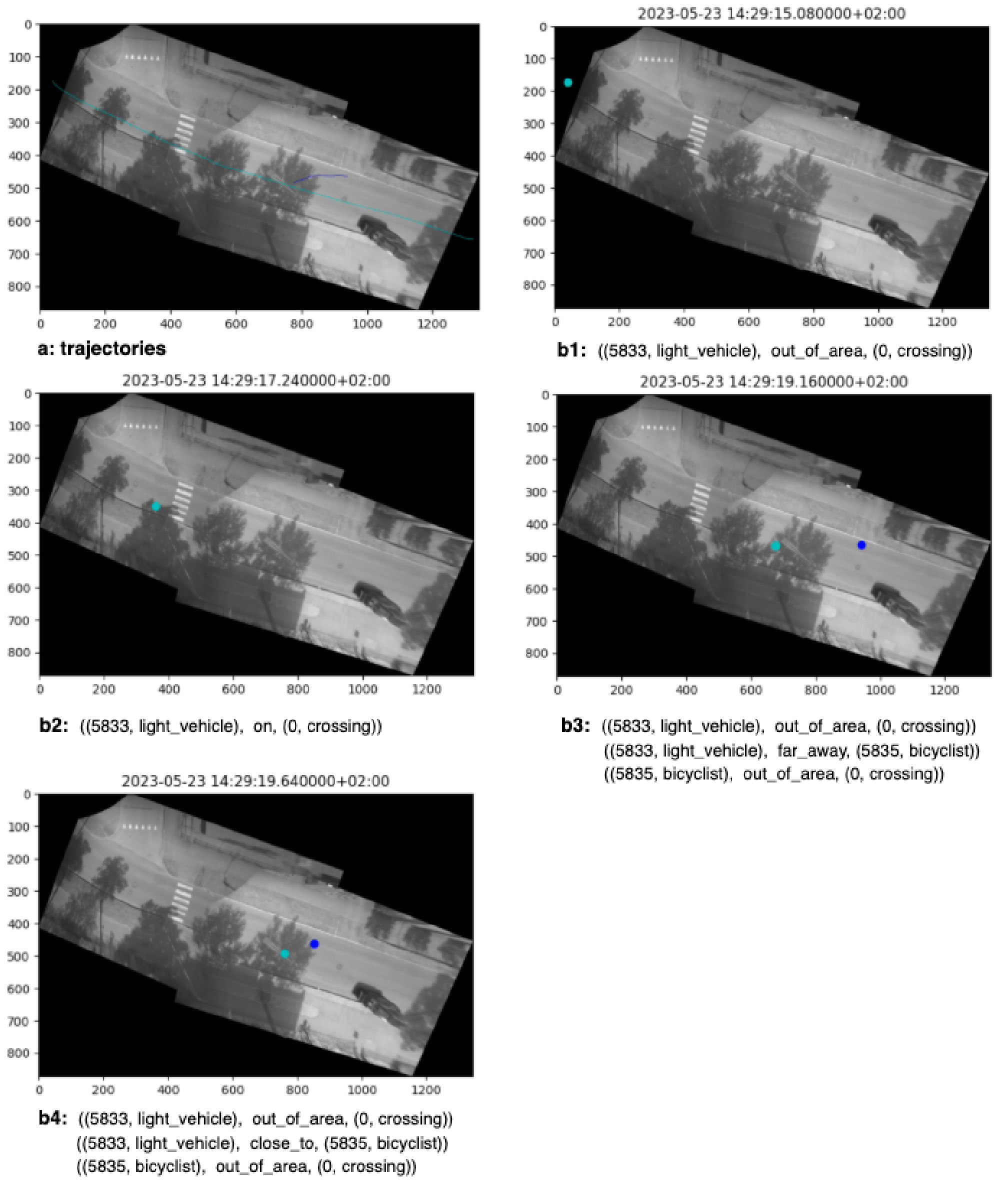

Figure 2: An example of VRU crossing behavior.

represents a light vehicle. Fig 2-b1 to b4 show a sequence of distinct segments that capture the changes in interactions between the cyclist and objects over time during the event.

The triples below express the crossing scene shown in Fig. 2. Each triple follows the format ((id, object_1), hasSpatialRelationship, (id, object_2)), where the object 1 is a moving object such as pedestrian, cyclist, and vehicle, and the object 2 can be a moving object or a static object such as crossing and sidewalk, and id is the unique identifier for each object. This is an unsafe behavior, where the cyclist was trying to cross the street outside the designated crossing area and came very close to the vehicle.

```
b1: ((5833, light_vehicle),  out_of_area, (0, crossing))
b2: ((5833, light_vehicle),  on, (0, crossing))
b3: ((5833, light_vehicle),  out_of_area, (0, crossing))
    ((5833, light_vehicle),  far_away, (5835, bicyclist))
    ((5835, bicyclist),  out_of_area, (0, crossing))
b4: ((5833, light_vehicle),  out_of_area, (0, crossing))
    ((5833, light_vehicle),  close_to, (5835, bicyclist))
    ((5835, bicyclist),  out_of_area, (0, crossing))
```

Fig 1 illustrates the current version of the ontology designed to represent the spatial-temporal evolution of crossing behavior. This version is an extended version of the ontology presented in our previous work [30]. This version also includes the factors relevant to crossing behavior, such as weather condition and traffic flow. The ovals represent the concepts, and the arrows represent the relationships between concepts. Specifically, the blue arrow represents `subclass` relations between concepts. The boxes represent the XSD (XML Schema Definition) data types. This ontology is accessible on GitHub[2]. Since *segment* is often related to regions in an image in computer vision, the term *frame* is used instead. In computer vision, a video can be divided into a sequence of frames. Each frame represents a single still image in the video sequence. For each object appearing in interactions, their coordinates and speed are captured too. Currently, the ontology includes only a limited number of categories for both moving and static objects. However, additional categories will be integrated as the ontology continues to undergo further development.

4. The Spatio-Temporal Knowledge Graphs

In this section, we describe the knowledge graphs constructed from road user dynamic data collected in two school areas in Jönköping, Sweden. The data is described using the semantic representation of crossing behavior mentioned beforehand in Section 3.

The datasets were prepared from the traffic measurement performed by Viscando AB[3] using the 3D&AI based infrastructure sensor OTUS3D. The measurements were carried out over 4 days, from 2023-05-23 to 2023-05-26, at two crossings in school areas in Jönköping. One crossing is a zebra crossing, and the other is a zebra-free crossing. The data contains trajectories of all road users recorded 10 times per second. Trajectories contain the unique track ID for each object, the UTC time stamp, position (i.e. X-coordinate and Y-coordinate), velocity (i.e. object speed in the direction of motion (km/h)) and object type. Currently, the object types include pedestrian, cyclist, light vehicle and heavy vehicle. Vision data are processed in the embedded computational unit and removed within 20 ms from being captured. Thus, the dataset is stored fully anonymously, ensuring compliance with the General Data Protection Regulation (GDPR) of the European Union[4], because personal information is neither stored in the sensors nor transmitted.

Since the application is to support the traffic infrastructure planning and development prioritizing VRUs, the knowledge graph construction has focused on the crossing scenes involving

[2]https://github.com/tanhe-git/crossing_behavior/blob/main/crossing_scene_ontology.owl
[3]www.viscando.com
[4]https://gdpr-info.eu/

	#behaviors	#ave_triples	#ave_nodes
dataset 1	1200	1503	352
dataset 2	312	715	169

Table 1
The structural metrics of the knowledge graphs

both VRU(s) and vehicle(s). The spatial relationship between objects was calculated based on the physical distance between them. The current spatial relationships include the ones between moving objects, i.e. *close_to* and *far_away*, and the ones between a moving object and a static object, i.e., *on*, *close_to*, *far_away*, *out_of_area*. When the information was extracted from the aforementioned datasets, the ontology described in Section 3 was populated, and the knowledge graphs were built.

The structural metrics of the knowledge graphs are provided in Table 1. Each crossing behavior is represented as a graph. Dataset 1 consists of data collected from a zebra crossing in a school area in Jönköping, while dataset 2 contains data from a zebra-free crossing in another school area in Jönköping. The numbers indicate the total number of crossing behaviors, as well as the average number of triples and nodes in the knowledge graphs representing these behaviors. Dataset 1 was collected from a school area in the center of Jönköping, where traffic is typically heavier. This may explain why the behavior KGs are larger in this dataset.

5. Insights Learned from the STKGs

In this section, we present the analysis of the knowledge graphs to gain a better understanding of the crossing behaviors. Moreover, we discuss safety-related concerns collected from traffic engineers, as well as insights queried and derived from the knowledge graphs, which can be used to address these concerns.

5.1. Clustering Analysis

We analyzed the knowledge graphs using clustering techniques. Each STKG represents a crossing scene and is transformed into embeddings. We then apply graph clustering algorithms to generate clusters.

Fig 3 shows the clusters obtained using the K-Means clustering algorithm, along with the embeddings based on the Local Degree Profile (LDP) [31]. LDP encodes the degree distribution of a node's neighbors in a graph and captures the local connectivity patterns. The clusters represent the graph's density. The graphs highlighted in yellow are denser, while those in purple are less dense. Higher density graph, the crossing behavior represented by the graph shows more moving objects involved and more complex interactions between objects during the crossing, which may correlate with higher accident risk. In contrast, sparsely connected graphs represent simpler, low-risk crossing scenarios.

Fig 4 shows the clusters obtained using the K-Means clustering algorithm, based on the embeddings generated by FEATHER [32] and GL2Vec [33]. Both methods capture more global structural patterns. However, they do not directly encode textual node or edge label information of a graph. To assess the similarity between the clustering results, the Adjusted Rand Index (ARI) [34] is used. ARI evaluates clustering similarity by considering all pairs of data points and determining whether they are assigned to the same or different clusters. For dataset 1, the ARI score is 0.51, indicating moderate agreement between clusterings, whereas for dataset 2, the ARI is 0.20, suggesting weak agreement between the clustering results.

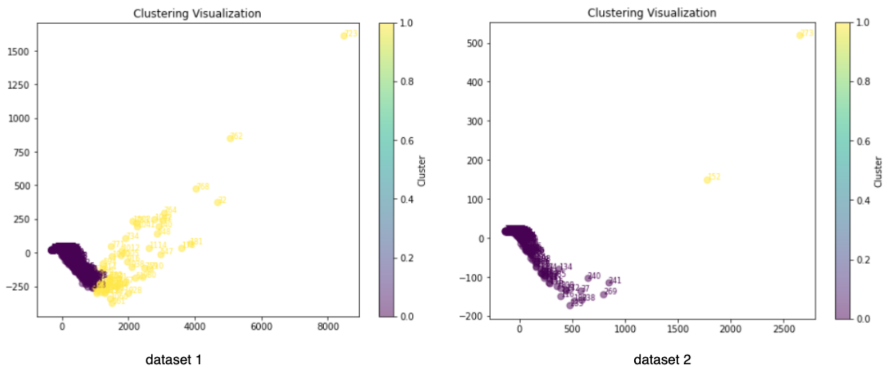

Figure 3: The clustering analysis 1

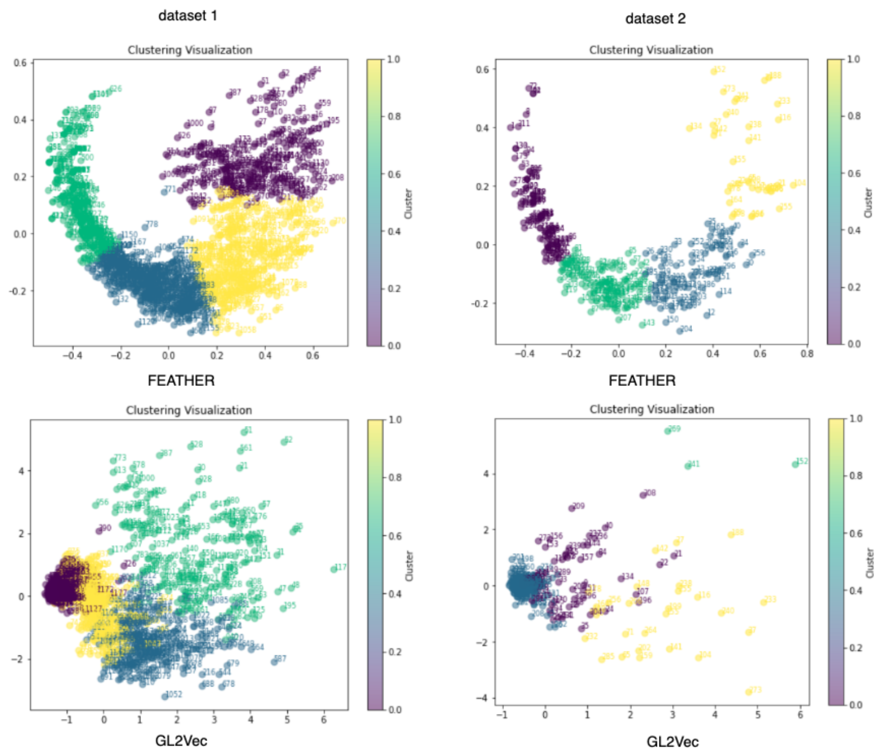

Figure 4: The clustering analysis 2

5.2. Query-based Analysis

Safety-related Concerns

First, we present the safety-related concerns gathered from traffic engineers at the City Planning, Development, and Traffic Department of Jönköping Municipality. These were collected during a workshop with traffic engineers, focusing on factors influencing the safety of crossing behaviors with respect to VRUs (pedestrians and cyclists) and vehicles. The discussions highlighted critical elements affecting behavior and potential risks in mobility.

Traffic engineers raised several concerns regarding pedestrian behavior at crossings, which can impact overall mobility safety:

- P1: Taking shortcuts: Many pedestrians cross diagonally instead of using designated crosswalks, increasing accident risks.
- P2: Walking speed variations: Differences in pedestrians walking speeds may increase collision risks at intersections.
- P3: False sense of security: Some pedestrians cross roads even when a car is approaching, assuming that drivers will stop, which can lead to dangerous situations.

Cyclists are particularly vulnerable road users, and their interactions with vehicles and pedestrians can pose safety challenges. The following concerns were highlighted:

- C1: Swinging out at crossings: Cyclists may unexpectedly change direction or enter vehicle lanes when approaching crossings, creating conflict points with motor vehicles.
- C2: Crossing at high speed: Some cyclists approach crossings at high speeds, reducing their ability to stop in time and increasing the risk of collisions with other VRUs and vehicles.

The behavior of drivers plays a crucial role in ensuring road safety, particularly in varying environmental and traffic conditions. Key factors identified include:

- V1: Speed: How quickly vehicles pass through the crossing area.
- V2: Type of vehicle: Larger vehicles, such as trucks and buses, have longer stopping distances and wider blind spots, increasing accident risks.
- V3: Weather conditions: Rain, snow, and icy roads affect vehicle traction and braking performance.

In this work we have focused on the concerns that can be addressed using datasets collected through performed measurements. Other safety concerns were also raised during the workshop, such as pedestrians crossing while looking at their phones. However, these particular concerns cannot be addressed by analyzing datasets collected from current sensors or by ensuring compliance with GDPR regulations.

Query and Results

The main SPARQL queries are included in Appendix A. For example, *Query 4* is designed to identify patterns of interactions involving cyclists who are not maintaining a safe speed during

pattern	frequency	concerns
pedestrian – close_to – vehicle	124 (1), 127 (2)	P3
pedestrian – cross (via) – out_of_area (crossing_area)	695 (1), 181 (2)	P1, P3
pedestrian – cross (with) – high speed	433 (1), 129 (2)	P2
cyclist – cross (via) – out_of_area (crossing_area)	68 (1), 3 (2)	C1
cyclist – cross (with) – high_speed	127 (1), 14 (2)	C2
vehicle – pass (with) – high speed AND good weather_condition	222 (1), 70 (2)	V1, V3
heavy vehicle – pass (with) – high speed AND good weather_condition	3 (1), 10 (2)	V1, V2, V3
vehicle – close_to – pedestrian AND vehicle – cross (with) – high speed AND good weather_condition	30 (1), 20 (2)	P3, V1, V3

Table 2
Pattern summaries, where (1) indicates dataset 1 and (2) indicates dataset 2.

a crossing scene. In the WHERE clause, it is specified that the spatial relationships for cyclists must include either 'out of crossing area' or 'on', ensuring that they have fully crossed the road rather than merely passing along the sidewalk. The query also uses a UNION to account for scenarios where cyclists are involved in interactions with either moving or static objects during the scene. Finally, a FILTER condition is applied to include only those interactions where the cyclist's speed is greater than 20 km/h.

In Table 2 we present the analysis results related to safety concerns observed in the two datasets. These patterns capture various types of interactions between VRUs and crossing road infrastructure, as well as interactions between vehicles, highlighting potential safety risks. The table provides the occurrence frequency of each pattern in Dataset 1 (zebra crossing area) and Dataset 2 (zebra-free crossing area), offering insights into how crossing behaviors may vary based on infrastructure design. These insights can allow traffic engineers to assess behavioral trends and road infrastructure design regarding improving safety.

6. Conclusions

In this paper, we introduced a structured spatial-temporal representation of VRU crossing behavior using ontology-based knowledge graphs. Our method enables explicit reasoning and explainability, addressing the limitations of deep learning-based trajectory prediction, which often lacks interpretability and structured knowledge representation. By constructing knowledge graphs from real-world urban datasets, we captured crossing behavior patterns and provided query-based insights to support traffic engineers in assessing safety risks and infrastructure planning.

To further enhance the effectiveness of STKG-based analysis, future research will explore additional techniques for pattern mining and causal analysis. Specifically, we aim to identify recurring unsafe crossing behavior patterns using methods such as frequent subgraph mining algorithms [35] and rule mining with LLMs [36], and reinforcement learning [37], and to apply causal inference techniques to perform cause-effect analysis. Moreover, we will use knowledge graph embeddings, rather than simple graph embeddings, to better capture the semantic information within the KGs. Such analyses could provide insights into factors influencing unsafe crossing behaviors and provide decision-making support for road infrastructure planning and development.

Acknowledgments

This work has been conducted in the project "Data and AI for decision Making suppOrt in traffic iNfrastructure Development (DAIMOND)" , which is funded by Vinnova (the Sweden's innovation agency) and AI Sweden (the Swedish national center for applied AI). The authors would like to thank the traffic department in Jönköping municipality for providing traffic safety related use cases and Viscando AB for providing traffic measurement dataset and expertise in traffic measurements and analysis.

Declaration on Generative AI

During the preparation of this work, the author(s) used ChatGPT to improve grammar, check spelling, and reword. After using these tool(s)/service(s), the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

Appendix A: SPARQL Queries

In this appendix, we include the SPARQL queries used to answer the safety related concerns described in Section 5.2.

Query 1: VRU is taking shortcut

```
SELECT (COUNT(DISTINCT ?b) AS ?total)
WHERE {
    ?b rdf:type ts:Behavior .
    ?b ts:hasFrame ?f .
    ?f ts:containsInteraction ?i .
    ?i ts:hasSpatialRelationship ?r .
    FILTER(?r = ts:out_of_area) .
    ?i ts:hasObject1 ?obj1 .
    ?i ts:hasObject2 ?obj2 .
    {?obj1 rdf:type ts:Pedestrian}
    UNION {?obj2 rdf:type ts:Pedestrian}
}
```

Query 2: regarding pedestrian walking speed

```
SELECT (COUNT(DISTINCT ?b) AS ?total)
WHERE {
    ?b rdf:type ts:Behavior .
    ?b ts:hasFrame ?f .
    ?f ts:containsInteraction ?i .
    {
        ?i ts:hasObject1 ?obj1 .
        ?obj1 rdf:type ts:Pedestrian .
        ?i ts:hasObject1Info ?obj1Info .
        ?obj1Info ts:speed ?speed .
        BIND(?obj1 AS ?pedestrian)
```

```
    }
    UNION
    {
        ?i ts:hasObject2 ?obj2 .
        ?obj2 rdf:type ts:Pedestrian .
        ?i ts:hasObject2Info ?obj2Info .
        ?obj2Info ts:speed ?speed .
        BIND(?obj2 AS ?pedestrian)
    }FILTER(?speed > 10)
}
```

Query 3: regarding cyclist swinging out at intersections

```
SELECT (COUNT(DISTINCT ?b) AS ?total)
WHERE {
    {
        SELECT ?b
        WHERE {
            ?b rdf:type ts:Behavior .
            ?b ts:hasFrame ?f .
            ?f ts:containsInteraction ?i .
            ?i ts:hasSpatialRelationship ?r .
            {
                SELECT ?b (COUNT(?r) AS ?totalRelations
                (SUM(IF(?r = ts:out_of_area, 1, 0)) AS ?outOfCrossingCount)
                WHERE {
                    ?b rdf:type ts:Behavior .
                    ?b ts:hasFrame ?f .
                    ?f ts:containsInteraction ?i .
                    ?i ts:hasSpatialRelationship ?r .
                }
                GROUP BY ?b
            }
            FILTER(?outOfCrossingCount > ?totalRelations / 2)
        }
    }
    ?b ts:hasFrame ?f .
    ?f ts:containsInteraction ?i .
    ?i ts:hasSpatialRelationship ?r .
    {
        ?i ts:hasObject1 ?obj1 .
        ?obj1 rdf:type ts:Bicyclist .
        ?i ts:hasObject1Info ?obj1Info .
        ?obj1Info ts:speed ?speed .
        BIND(?obj1 AS ?bicyclist)
    }
    UNION
    {
        ?i ts:hasObject2 ?obj2 .
        ?obj2 rdf:type ts:Bicyclist .
        ?i ts:hasObject2Info ?obj2Info .
        ?obj2Info ts:speed ?speed .
        BIND(?obj2 AS ?bicyclist)
    }
}
```

Query 4: regarding cyclist unsafe speed

```
SELECT (COUNT(DISTINCT ?b) AS ?total)
WHERE {
    ?b rdf:type ts:Behavior .
    ?b ts:hasFrame ?f .
```

```
    ?f ts:containsInteraction ?i .
    ?i ts:hasSpatialRelationship ?r .
    FILTER(?r = ts:out_of_area || ?r = ts:on) .
    {
        ?i ts:hasObject1 ?obj1 .
        ?obj1 rdf:type ts:Bicyclist .
        ?i ts:hasObject1Info ?obj1Info .
        ?obj1Info ts:speed ?speed .
        BIND(?obj1 AS ?object)
    }
    UNION
    {
        ?i ts:hasObject2 ?obj2 .
        ?obj2 rdf:type ts:Bicyclist .
        ?i ts:hasObject2Info ?obj2Info .
        ?obj2Info ts:speed ?speed .
        BIND(?obj2 AS ?object)
    }
    FILTER(?speed > 20)
}
```

Query 5: regarding vehicle's speed

```
SELECT (COUNT(DISTINCT ?b) AS ?total)
WHERE {
    ?b rdf:type ts:Behavior .
    ?b ts:hasFrame ?f .
    ?f ts:containsInteraction ?i .
    {
        ?i ts:hasObject1 ?obj1 .
        {?obj1 rdf:type ts:Pedestrian}
        UNION {?obj1 rdf:type ts:Bicyclist}.
        ?i ts:hasObject2 ?obj2 .
        {?obj2 rdf:type ts:HeavyVehicle}
        UNION {?obj2 rdf:type ts:LightVehicle}.
        ?i ts:hasObject2Info ?obj2Info .
        ?obj2Info ts:speed ?speed .
        BIND(?obj2 AS ?vehicle)
    }
    UNION
    {
        ?i ts:hasObject2 ?obj2 .
        {?obj2 rdf:type ts:Pedestrian}
        UNION {?obj2 rdf:type ts:Bicyclist}.
        ?i ts:hasObject1 ?obj1 .
        {?obj1 rdf:type ts:HeavyVehicle}
        UNION {?obj1 rdf:type ts:LightVehicle}.
        ?i ts:hasObject1Info ?obj1Info .
        ?obj1Info ts:speed ?speed .
        BIND(?obj1 AS ?vehicle)
    }FILTER(?speed > 30)
}
```

Query 6: regarding vehicle's speed and distance to VRUs

```
SELECT (COUNT(DISTINCT ?b) AS ?total)
WHERE {
    ?b rdf:type ts:Behavior .
    ?b ts:hasFrame ?f .
    ?f ts:containsInteraction ?i .
    ?i ts:hasSpatialRelationship ts:close_to .
    {
```

```
        ?i ts:hasObject1 ?obj1 .
        {?obj1 rdf:type ts:Pedestrian}
        UNION {?obj1 rdf:type ts:Bicyclist}.
        ?i ts:hasObject2 ?obj2 .
        {?obj2 rdf:type ts:HeavyVehicle}
        UNION {?obj2 rdf:type ts:LightVehicle}.
        ?i ts:hasObject2Info ?obj2Info .
        ?obj2Info ts:speed ?speed .
        BIND(?obj2 AS ?vehicle)
    }
    UNION
    {
        ?i ts:hasObject2 ?obj2 .
        {?obj2 rdf:type ts:Pedestrian}
        UNION {?obj2 rdf:type ts:Bicyclist}.
        ?i ts:hasObject1 ?obj1 .
        {?obj1 rdf:type ts:HeavyVehicle}
        UNION {?obj1 rdf:type ts:LightVehicle}.
        ?i ts:hasObject1Info ?obj1Info .
        ?obj1Info ts:speed ?speed .
        BIND(?obj1 AS ?vehicle)
    }
    FILTER(?speed > 30)
}
```

References

[1] C. Koszowski, R. Gerike, S. Hubrich, T. Götschi, M. Pohle, R. Wittwer, Active mobility: bringing together transport planning, urban planning, and public health, Towards user-centric transport in europe: challenges, solutions and collaborations (2019) 149–171.

[2] W. H. Organization, Compendium of WHO and other UN guidance on health and environment: version with International Classification of Health Intervention (ICHI) codes, World Health Organization, 2024.

[3] R. Abduljabbar, H. Dia, S. Liyanage, S. A. Bagloee, Applications of artificial intelligence in transport: An overview, Sustainability 11 (2019) 189.

[4] S. Ferguson, B. Luders, R. C. Grande, J. P. How, Real-time predictive modeling and robust avoidance of pedestrians with uncertain, changing intentions, in: Algorithmic Foundations of Robotics XI: Selected Contributions of the Eleventh International Workshop on the Algorithmic Foundations of Robotics, Springer, 2015, pp. 161–177.

[5] E. Papadimitriou, G. Yannis, J. Golias, A critical assessment of pedestrian behaviour models, Transportation research part F: traffic psychology and behaviour 12 (2009) 242–255.

[6] A. Rasouli, I. Kotseruba, Intend-wait-cross: Towards modeling realistic pedestrian crossing behavior, in: 2022 IEEE Intelligent Vehicles Symposium (IV), IEEE, 2022, pp. 83–90.

[7] D. Ridel, E. Rehder, M. Lauer, C. Stiller, D. Wolf, A literature review on the prediction of pedestrian behavior in urban scenarios, in: 2018 21st International Conference on Intelligent Transportation Systems (ITSC), IEEE, 2018, pp. 3105–3112.

[8] E. Schuetz, F. B. Flohr, A Review of Trajectory Prediction Methods for the Vulnerable Road User, Robotics 13 (2023) 1.

[9] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, S. Savarese, Social LSTM:

Human trajectory prediction in crowded spaces, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 961–971.
[10] P. Zhang, W. Ouyang, P. Zhang, J. Xue, N. Zheng, SR-LSTM: State refinement for lstm towards pedestrian trajectory prediction, in: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2019, pp. 12085–12094.
[11] R. Quan, L. Zhu, Y. Wu, Y. Yang, Holistic LSTM for pedestrian trajectory prediction, IEEE transactions on image processing 30 (2021) 3229–3239.
[12] Y. Huang, H. Bi, Z. Li, T. Mao, Z. Wang, STGAT: Modeling spatial-temporal interactions for human trajectory prediction, in: Proceedings of the IEEE/CVF international conference on computer vision, 2019, pp. 6272–6281.
[13] Y. Li, J. Cui, Z. Zhao, L. Li, ST-AGNN: Spatial-Temporal Attention Graph Neural Network for Pedestrian Trajectory Prediction, in: Applied Mathematics, Modeling and Computer Simulation, IOS Press, 2022, pp. 268–275.
[14] C. Yu, X. Ma, J. Ren, H. Zhao, S. Yi, Spatio-temporal graph transformer networks for pedestrian trajectory prediction, in: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XII 16, Springer, 2020, pp. 507–523.
[15] S. Liu, Y. Wang, J. Sun, T. Mao, An efficient spatial-temporal model based on gated linear units for trajectory prediction, Neurocomputing 492 (2022) 593–600.
[16] J. Yang, X. Sun, R. G. Wang, L. X. Xue, PTPGC: Pedestrian trajectory prediction by graph attention network with ConvLSTM, Robotics and Autonomous Systems 148 (2022) 103931.
[17] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L.-J. Li, D. A. Shamma, et al., Visual genome: Connecting language and vision using crowdsourced dense image annotations, International journal of computer vision 123 (2017) 32–73.
[18] T. Hägerstrand, What about people in Regional Science?, Papers of the Regional Science Association 24 (1970) 6–21. URL: http://dx.doi.org/10.1007/bf01936872. doi:10.1007/bf01936872.
[19] D. Orellana, C. Renso, Developing an interactions ontology for characterising pedestrian movement behaviour, in: Movement-aware applications for sustainable mobility: Technologies and approaches, IGI Global, 2010, pp. 62–86.
[20] C. Chai, X. Shi, Y. D. Wong, M. J. Er, E. T. M. Gwee, Fuzzy logic-based observation and evaluation of pedestrians' behavioral patterns by age and gender, Transportation research part F: traffic psychology and behaviour 40 (2016) 104–118.
[21] A. Gharebaghi, M.-A. Mostafavi, G. Edwards, P. Fougeyrollas, S. Gamache, Y. Grenier, Integration of the social environment in a mobility ontology for people with motor disabilities, Disability and Rehabilitation: Assistive Technology 13 (2018) 540–551.
[22] F. Fang, S. Yamaguchi, A. Khiat, Ontology-based reasoning approach for long-term behavior prediction of road users, in: 2019 IEEE Intelligent Transportation Systems Conference (ITSC), IEEE, 2019, pp. 2068–2073.
[23] D. Newtson, Attribution and the unit of perception of ongoing behavior., Journal of personality and social psychology 28 (1973) 28.
[24] L. Spector, J. Grafman, Planning, neuropsychology, and artificial intelligence: cross-fertilization, Handbook of neuropsychology 9 (1994) 377–392.
[25] C. Baldassano, J. Chen, A. Zadbood, J. W. Pillow, U. Hasson, K. A. Norman, Discovering event structure in continuous narrative perception and memory, Neuron 95 (2017) 709–721.

[26] J. M. Zacks, B. Tversky, G. Iyer, Perceiving, remembering, and communicating structure in events., Journal of experimental psychology: General 130 (2001) 29.
[27] N. K. Speer, J. M. Zacks, J. R. Reynolds, Perceiving narrated events, in: Proceedings of the Annual Meeting of the Cognitive Science Society, volume 26, 2004.
[28] J. Ji, R. Krishna, L. Fei-Fei, J. C. Niebles, Action genome: Actions as compositions of spatio-temporal scene graphs, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020, pp. 10236–10247.
[29] L. Mlodzian, Z. Sun, H. Berkemeyer, S. Monka, Z. Wang, S. Dietze, L. Halilaj, J. Luettin, nuScenes Knowledge Graph - A Comprehensive Semantic Representation of Traffic Scenes for Trajectory Prediction, in: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops, 2023, pp. 42–52.
[30] H. Tan, F. Westphal, A semantic representation of pedestrian crossing behavior, in: Joint of the ESWC 2024 Workshops and Tutorials, ESWC-JP 2024 Hersonissos 26 May 2024 through 27 May 2024, volume 3749, CEUR-WS, 2024.
[31] C. Cai, Y. Wang, A simple yet effective baseline for non-attributed graph classification, arXiv preprint arXiv:1811.03508 (2018).
[32] B. Rozemberczki, R. Sarkar, Characteristic functions on graphs: Birds of a feather, from statistical descriptors to parametric models, in: Proceedings of the 29th ACM international conference on information & knowledge management, 2020, pp. 1325–1334.
[33] H. Chen, H. Koga, Gl2vec: Graph embedding enriched by line graphs with edge features, in: Neural Information Processing: 26th International Conference, ICONIP 2019, Sydney, NSW, Australia, December 12–15, 2019, Proceedings, Part III 26, Springer, 2019, pp. 3–14.
[34] D. Steinley, Properties of the hubert-arable adjusted rand index., Psychological methods 9 (2004) 386.
[35] C. Jiang, F. Coenen, M. Zito, A survey of frequent subgraph mining algorithms, The Knowledge Engineering Review 28 (2013) 75–105.
[36] L. Luo, J. Ju, B. Xiong, Y.-F. Li, G. Haffari, S. Pan, Chatrule: Mining logical rules with large language models for knowledge graph reasoning, arXiv preprint arXiv:2309.01538 (2023).
[37] W. Xiong, T. Hoang, W. Y. Wang, Deeppath: A reinforcement learning method for knowledge graph reasoning, arXiv preprint arXiv:1707.06690 (2017).