Explainable Mobility Prediction in Urban Transit Zones

Adrian M.P. Braşoveanu[1,2,*], Lyndon J.B. Nixon[1,2], Arno Scharl[1,3], Günther Charwat[4] and Egon Prünster[4]

[1]*Modul University Vienna, Am Kahlenberg 1, 1190, Vienna, Austria*
[2]*Modul Technology GmbH, Am Kahlenberg 1, 1190, Vienna, Austria*
[3]*webLyzard technology GmbH, Liechtensteinstrasse 41/26, 1090 Vienna, Austria*
[4]*ummadum Service GmbH, Prinz-Eugen-Straße 2/5, 1040 Vienna, Austria*

Abstract

Sustainable mobility behavior remains a challenging objective, as individuals rarely alter their transportation habits solely for environmental benefits. Creating effective incentives can prompt change, but first, it requires a comprehensive understanding of mobility patterns. This paper examines the prediction of personal travel activities (including next-trip forecasting and activity classification) utilizing recent journey data collected through a mobility application. A graph-based fusion ensemble that include a graph convolutional network and a statistical user patterns is designed for structured prediction that includes multiple outputs like origin and destination, transportation mode, time. An explainable prediction pipeline is built on top of this ensemble. The results are then converted to a knowledge graph that allows us to run a sophisticated analysis and helps improve our weekly workflows.

Keywords

sustainable transportation, mobility prediction, graph convolutional network, explainable mobility graph,

1. Introduction

Achieving sustainable transportation patterns is challenging since individuals rarely modify established routines solely for environmental benefits. Key considerations encompass personal convenience (overcrowded transit deterring usage), network development (separate cycling corridors), capability requirements, motivation systems (e.g., notifications or rewards, and inclusive design features. If individuals are to change their behavior, the proposed benefits of the motivation systems must exceed the perceived advantages of continuing current practices, including financial savings, comfort zones, and ingrained patterns.

The AI-CENTIVE project's primary goal is to use neural network architectures to understand transportation behaviors in Austrian urban transit zones and identify the most effective reward structures to encourage a shift towards eco-friendly travel options, such as using bicycles or public transit instead of private vehicles for urban commuting. The research involves developing ensembles that combine statistical patterns and graph convolutional networks (GCNs). Our experiments quantify environmental impacts by combining baseline predictions with simulated incentives across mobility networks.

The main contribution presented in this paper is the semantic conversion workflow developed to support explainable AI predictions. The output of the GCNs is converted into a mobility knowledge graph, and the best predictions are analyzed using SPARQL queries.

The paper is organized as follows: Section 2 offers a brief overview of related work, Section 3 describes the general method, while Section 4 presents the explainability pipeline and the associated analysis that can be performed with it. The paper concludes with a brief overview and future work.

SEMANTICS 2025 EU: 21st International Conference on Semantic Systems, September 3-5, 2025, Vienna, Austria
*Corresponding author.
adrian.brasoveanu@modul.ac.at (A. M.P. Braşoveanu); lyndon.nixon@modul.ac.at (L. J.B. Nixon); arno.scharl@modul.ac.at (A. Scharl); guenther.charwat@ummadum.com (G. Charwat); egon.pruenster@ummadum.com (E. Prünster)
0000-0002-3439-4736 (A. M.P. Braşoveanu); 0000-0001-7091-4543 (L. J.B. Nixon); 0000-0001-5346-9521 (A. Scharl)

2. Related Work

This brief section covers only graph mobility network surveys and related work on explainability.

Examining graph-based computational methods, Jiang's survey [1] catalogs traffic applications, and model variations, including standard graph networks, convolutional adaptations, and unsupervised graph encoders. Expanding this scope significantly, Wang et al. [2] provides detailed analysis covering broader transportation applications: vehicle storage systems, safety enhancement, autonomous navigation, and metropolitan design optimization.

Beyond academic surveys identifying research directions, practical implementation insights from the Google Maps arrival prediction system [3] deserve particular attention through detailed experimental variations and qualitative performance evaluation on actual transportation data. Notable mention regarding Google concerns their adaptive approach toward evolving privacy regulations shown by shifting personal movement histories onto individual devices.

Understanding model decision processes remains crucial for discovering inherent biases often stemming from unbalanced training samples. Schwalbe and Finzel [4] deliver extensive examination synthesizing over fifty specialized surveys addressing interpretability challenges across computational domains. Wang et. al [2] focus on knowledge graph related research in smart city domain which was the starting point of our mobility graph idea for fast analysis idea presented in this paper.

3. Method

A core challenge the AI-CENTIVE project addresses is the observation that individuals are often unwilling to change their mobility habits solely for environmental reasons. To overcome this, the project's overarching mission is to develop AI-based incentivization techniques to influence citizens' mobility choices, utilizing multimodal models of mobility activity and data analytics. The ultimate vision is to enable and incentivize Austrian citizens to adopt more sustainable mobility choices.

The work presented here builds upon the Ummadum platform which allows users to log their sustainable mobility activities via a mobile app. This process reports CO2 emissions reductions to raise users' sustainability awareness and rewards sustainable behavior to incentivize environmentally friendly choices. The AI-CENTIVE community within Ummadum serves to engage participants, assess the effectiveness of incentives, conduct pilot tests for integrating AI-driven solutions, and plan rollout strategies across Austria. All trip data is collected with user consent, strictly adhering to GDPR best practices by anonymizing user identities and locations for analysis.

The project's dataset for examining Vienna mobility patterns was provided by Ummadum. This commuter dataset comprises approximately 450,000 user trips collected between January 2024 and April 2025 from users of the Ummadum mobile app. The dataset focuses on activities such as biking, walking, public transport, and car sharing (as a rider or driver), also including data about activity status (e.g., cancelled, finished). To maintain GDPR compliance and route anonymity, user data and trip data are anonymized, with segments added or removed from trips. Locations for origins and destinations are expressed through zip codes, and information about location type (e.g., office, public venue) is also collected. The dataset also includes details on the incentives provided, such as points, rewards, carsharing, community types administering rewards, or activity challenges where users earn additional rewards based on their activity levels.

The project's approach towards incentivization primarily involves rewards and notifications. Rewards include classic elements like: i) challenges; ii) lotteries, or iii) collection of points for sustainable activities. Notifications are displayed directly to the users based on their past history. Three types of notifications were implemented: i) success notifications to motivate participants based on their previous sustainable mobility actions; ii) AI recommendation notifications to suggest future sustainable mobility options using contextual information such as behavior, mobility preferences, location, and time; and iii) weather notifications for notable weather events. Explanations for the AI notifications are included, based on past user history. These notifications are part of a simple workflow that runs automatically each

Monday morning, scheduling notifications after checking eligibility criteria (e.g., high confidence, no more than one notification per user per day).

Initial insights from the first 2025 pilot program indicated a total CO2 savings of 331 tons. These insights were used to refine a second pilot in 2025. This second pilot incorporates improvements such as explainable AI notifications (justifying notifications based on the user's history, including similar trips or partial routes), weather alerts, and enhanced challenges within the rewards system. The overall goal of the second pilot was to test and refine AI-driven incentivization techniques to encourage sustainable mobility choices among Austrian citizens using the Ummadum platform.

We have examined several models designed for a structured prediction task that simultaneously forecasts multiple outputs, including trip origin and destination locations, hour and day of the trip, activity type, distance, and duration. These models, including Transformer and GCN ensembles that include statistical models (e.g., ARIMA, XGBoost) are presented in a previous publication [5]. The best model was an ensemble that combines statistical user patters (e.g., temporal and spatial patterns related to the past trips, including same origin or destination, or time intervals) and a GCN architecture.

To enhance transparency and regulatory compliance of our Graph Neural Network mobility predictions, we developed a semantic conversion pipeline that transforms prediction outputs into structured RDF knowledge graphs. This approach addresses the requirements of the European AI Act for explainable AI systems in transportation applications.

Our system employs a multilayered semantic framework combining established ontologies with custom ontologies [1]. The core structure integrates PROV-O (W3C Provenance Ontology) [6] for comprehensive provenance tracking, allowing full traceability from input data through model inference to final predictions. Temporal relationships are modeled using the W3C TIME [7] ontology, providing a standardized representation of trip scheduling and duration predictions.

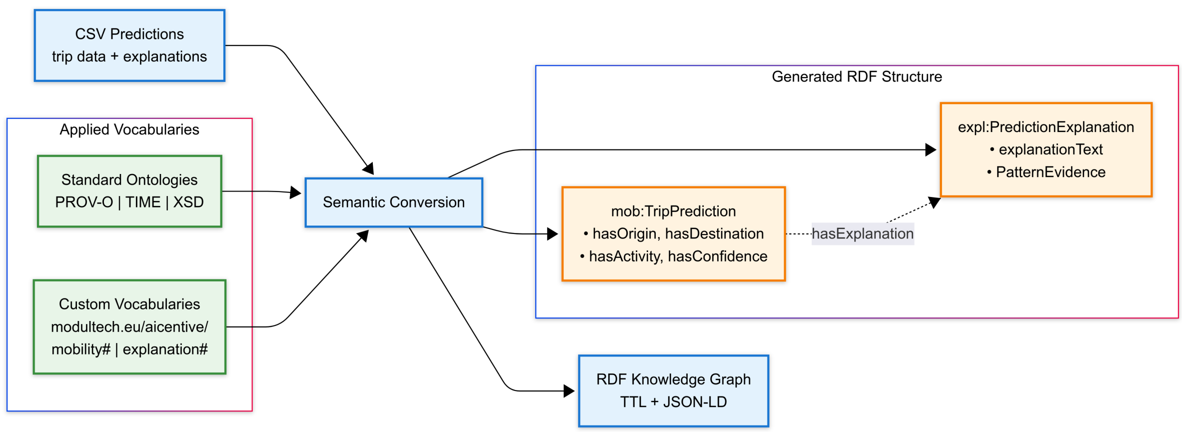

Figure 1: Semantic conversion workflow for Explainable AI predictions.

For mobility-specific concepts, we defined custom RDF vocabularies in the modultech.eu/aicentive namespace, establishing classes such as TripPrediction, PredictionExplanation, and PatternEvidence. These vocabularies maintain compatibility with existing transport standards like GTFS (General Transit Feed Specification) while extending semantic representation for explainable AI applications. The explainability vocabulary captures structured reasoning through properties that include routeFrequency, timeFrequency, and activityFrequency, allowing systematic analysis of prediction confidence factors. Each prediction entity links to explanation instances that contain machine-readable pattern evidence and human-readable justification text.

The conversion pipeline automatically processes GNN outputs into RDF using rdflib, preserving all prediction metadata including confidence scores, historical pattern evidence, and explanatory reasoning. The resulting knowledge graph enables direct SPARQL querying without database infrastructure,

[1]The ontologies were designed for offline use and are publicly available through GitHub - https://github.com/modultechnology/aicentiveontologies

supporting complex analytical queries for model evaluation and regulatory auditing. Semantic queries can systematically identify low-confidence predictions, analyze temporal patterns in model performance, and extract compliance reports for regulatory review. This queryable format facilitates automated detection of prediction biases and supports continuous model improvement through pattern analysis.

4. Evaluation

This semantic conversion enables sophisticated analytical capabilities through SPARQL querying that would be difficult or impossible with traditional CSV formats. Researchers can execute complex queries to identify high-confidence predictions for regulatory audit trails, systematically analyze low-confidence routes that require model improvement, and investigate user behavior patterns by correlating confidence scores with historical pattern evidence (see Table 1).

Table 1
Example SPARQL Queries for Semantic Mobility Dataset Analysis

Query Purpose	SPARQL Query

Top High-Confidence Predictions	PREFIX mob: <http://modultech.eu/aicentive/mobility#>
	SELECT ?trip ?user ?confidence ?activity
	       ?origin ?destination
	WHERE {
	    ?trip a mob:TripPrediction ;
	          mob:hasUser ?user ;
	          mob:hasConfidence ?confidence ;
	          mob:hasActivity ?activity ;
	          mob:hasOrigin ?origin ;
	          mob:hasDestination ?destination .
	}
	ORDER BY DESC(?confidence)
	LIMIT 500

Low-Confidence Routes Requiring Model Improvement	PREFIX mob: <http://modultech.eu/aicentive/mobility#>
	SELECT ?origin ?destination
	       (AVG(?confidence) as ?avg_confidence)
	       (COUNT(*) as ?prediction_count)
	WHERE {
	    ?trip mob:hasOrigin ?origin ;
	          mob:hasDestination ?destination ;
	          mob:hasConfidence ?confidence .
	    FILTER(?confidence < 0.3)
	}
	GROUP BY ?origin ?destination
	ORDER BY ?avg_confidence
	LIMIT 20

User Behavior Pattern Analysis	PREFIX mob: <http://modultech.eu/aicentive/mobility#>
	PREFIX expl: <http://modultech.eu/aicentive/explanation#>
	SELECT ?user ?activity
	       (AVG(?confidence) as ?avg_confidence)
	       (AVG(?route_freq) as ?avg_route_familiarity)
	WHERE {
	    ?trip mob:hasUser ?user ;
	          mob:hasActivity ?activity ;
	          mob:hasConfidence ?confidence ;
	          mob:hasExplanation ?exp .
	    ?exp expl:hasPatternEvidence ?pattern .
	    ?pattern expl:routeFrequency ?route_freq .
	}
	GROUP BY ?user ?activity
	HAVING (COUNT(*) > 5)
	ORDER BY DESC(?avg_confidence)

The semantic structure supports automated quality assurance through queries that filter predictions by confidence thresholds, group routes by performance metrics, and extract explanation metadata for transparency reporting. Beyond basic filtering, the pipeline enables advanced analytical patterns such as identifying users with consistently high prediction confidence, discovering temporal patterns in model performance, and analyzing the relationship between route familiarity and prediction accuracy. The explainability metadata becomes queryable, allowing researchers to systematically study which types of historical evidence correlate with prediction success and identify biases in model performance across different user groups or activity types. This semantic approach transforms static prediction output into a dynamic knowledge graph that supports continuous model evaluation, regulatory compliance reporting, and data-driven insights for model improvement. The standardized format also enables integration with other transportation data sets and supports collaborative research through interoperable semantic web technologies. Most importantly, queryable explanations provide the foundation for automated compliance checking required by emerging AI regulations, while simultaneously supporting research into the relationship between explainability quality and prediction accuracy.

One particular advantage of this approach is speed. It is generally 3x-5x times faster than the pure CSV approach. This is mainly due to the SPARQL engines being optimized for triple pattern matching than traditional SQL or pandas operations on CSV data. The processing efficiency is higher also due to SPARQL's lazy evaluation mechanism (e.g., only process the data needed to answer specific queries) and selective loading (e.g., retrieving only data needed for each analysis). RDFlib itself is also optimized to handle larger datasets more efficiently through streaming.

Table 1 showcases several SPARQL queries that are routinely used for explainable mobility analysis. Table 2 presents the kind of automated analysis that can be done automatically on top of the evaluation results. All the results presented in this table refer to the last week of evaluation from our second pilot.

Table 2
Semantic Explainability Evaluation Summary

Evaluation Category	Metric	Result
High-Confidence Predictions	Retrieved Predictions	500
	Confidence Range	0.733 – 1.000
	Mean Confidence	0.798
	Median Confidence	0.770
User Coverage	Total Users	791
	Users in Top 500	92
	Coverage Percentage	11.6%
Activity Distribution (Top 500)	BIKE	176 (35.2%)
	WALK	157 (31.4%)
	CAR_DRIVER	97 (19.4%)
	PT	43 (8.6%)
	CAR_RIDER	27 (5.4%)

The semantic approach directly addresses EU AI Act mandates for transparency and auditability in AI applications. The structured explainability format enables automated compliance monitoring and systematic verification of AI decision-making processes. Beyond regulatory compliance, the standardized RDF format enhances research reproducibility and enables integration with broader smart city transportation planning systems, positioning our work at the intersection of responsible AI development and practical mobility applications.

5. Conclusion

Our research combines various forecasting approaches for urban transportation pattern analysis across larger city environments such as Vienna. The potential combinations when merging different prediction techniques appears limitless. We have merely scratched the surface regarding possible architectural integrations. Particular attention was paid to blending conventional statistical methods with advanced learning frameworks, recognizing how easily established techniques are overlooked amid technological advancements. Statistical approaches often deliver substantial value despite newer alternatives.

More complex ensembles frequently introduce additional challenges, including extended computation requirements, elaborate technical infrastructure demands, and troubleshooting complexity. Integrating classic statistical models with deep learning architectures allows us to avoid such challenges, as the resulting models are simple and efficient.

The semantic conversation pipeline and the rapid analysis it enables represents one of the highlights of our projects. Due to examining the results through such tools, we were able to quickly improve our best model, and also deliver high quality notifications to our community members.

Future investigations will incorporate emissions calculations from the carbon assessment tool developed by project partner BOKU to evaluate environmental benefits across different reward mechanisms. This enables concrete quantification regarding pollution reduction achieved through altered journey selections. Additionally, data about weather conditions will enhance the ensemble accuracy.

Building widespread participation in eco-friendly journey choices remains equally crucial. Monitoring increased adoption rates of sustainable options helps evaluate progress in behavioral transformation. Furthermore, we intend to target problematic urban locations lacking vehicle storage facilities or convenient transit access within a reasonable walking range. Quantifying improvements within these zones demonstrates enhanced accessibility plus livability improvements. Such enhanced measurements strengthen our predictive capabilities while supporting inclusive, evidence-based approaches toward sustainable transportation development. Through our project initiatives, we aim to empower community members and organizational stakeholders in creating more responsive, equitable, and environmentally responsible urban transit zones.

6. Acknowledgments

All authors are partially funded through the AI-CENTIVE project (FFG G.A. 892238). The AI-CENTIVE project is funded within the RTI Initiative of the ICT of the Future "IKT der Zukunft" – an initiative of the Federal Ministry for Climate Protection, Environment, Energy, Mobility, Innovation and Technology (BMK).

References

[1] W. Jiang, J. Luo, Graph Neural Network for Traffic Forecasting: A Survey, Expert Syst. Appl. 207 (2022) 117921. doi:10.1016/J.ESWA.2022.117921.

[2] Z. Wang, F. Han, S. Zhao, A survey on knowledge graph related research in smart city domain, ACM Transactions on Knowledge Discovery from Data 18 (2024) 1–31. doi:10.1145/3672615.

[3] A. Derrow-Pinion, J. She, D. Wong, O. Lange, T. Hester, L. Perez, M. Nunkesser, S. Lee, X. Guo, B. Wiltshire, P. W. Battaglia, V. Gupta, A. Li, Z. Xu, A. Sanchez-Gonzalez, Y. Li, P. Velickovic, ETA Prediction with Graph Neural Networks in Google Maps, in: G. Demartini, G. Zuccon, J. S. Culpepper, Z. Huang, H. Tong (Eds.), CIKM '21: The 30th ACM International Conference on Information and Knowledge Management, Virtual Event, Queensland, Australia, November 1 - 5, 2021, ACM, 2021, pp. 3767–3776. doi:10.1145/3459637.3481916.

[4] G. Schwalbe, B. Finzel, A Comprehensive Taxonomy for Explainable Artificial Intelligence: A Systematic Survey of Surveys on Methods and Concepts, Data Min. Knowl. Discov. 38 (2024) 3043–3101. doi:10.1007/S10618-022-00867-8.

[5] A. M. Bra¸soveanu, L. J. Nixon, A. Scharl, Hybrid ai models for structured mobility prediction in metropolitan areas, in: IWANN2025, 2013.

[6] P. Missier, K. Belhajjame, J. Cheney, The w3c prov family of specifications for modelling provenance metadata, in: Proceedings of the 16th international conference on extending database technology, 2013, pp. 773–776. doi:10.1145/2452376.2452478.

[7] F. Pan, J. R. Hobbs, Time ontology in owl, W3C working draft, W3C 1 (2006) 1. URL: https://www.w3.org/TR/2022/CRD-owl-time-20221115/.