Ref. Ares(2026)8301273 - 31/08/2026

FEDORA

ADVANCED ROUTING & MITIGATION STRATEGIES

Project deliverable D4.2

HORIZON Research and Innovation Action | 101203465 |
Call HORIZON-CL5-2023-D6-01

Funded by the European Union

Project funded by
Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
Swiss Confederation
Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

DELIVERABLE ADMINISTRATIVE INFORMATION

TITLE OF THE DELIVERABLE	D4.2 - Advanced Routing & Mitigation Strategies
WORK PACKAGE	WP4
TYPE OF DELIVERABLE	R — Document, report
DISSEMINATION LEVEL	PU - Public
STATUS – VERSION, DATE	Final – V1.0, 26/08/2026
DELIVERABLE LEADER	FRONTIER INNOVATIONS (FRONT)
CONTRACTUAL DATE OF DELIVERY	31/08/2026
SUBMISSION DATE	07/08/2026
KEYWORDS	Multimodal, Generative AI, anomaly detection, incident detection, UAV, computer vision, synchromodal transport, freight, last mile, dynamic mitigation plan

Funded by the European Union

Project funded by

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
Swiss Confederation

Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

LIST OF CONTRIBUTORS

Name	Organisation
Kostas Cheliotis, Adel Almohammad	FRONT
Chralambos Menelaou, Tania Panayiotou, Panayiotis Kolios	UCY
Julian Stephens	MJC2
Panagiotis Fafoutellis, Eleni Vlahogianni	NTUA

QUALITY CONTROL

	Name	Organisation	Date
Peer review	Csaba L. Hargitai	BME	24/08/2026
Peer review	Anastasios Kouvelas	ETH	25/08/2026
Quality check	Roozbeh Mohammadi	ERT	31/08/2026

VERSION HISTORY

Version	Date	Author	Summary of changes
0.1	21/05/2026	Kostas Cheliotis	Added Table of Contents
0.2	26/06/2026	Julian Stephens	Added Chapter 3
0.3	21/07/2026	Adel Almohammad	Added Chapter 4
0.4	28/07/2026	Chralambos Menelaou, Tania Panayiotou, Panayiotis Kolios	Added Section 2.3
0.5	06/08/2026	Panagiotis Fafoutellis, Eleni Vlahogianni	Added Section 2.1
0.6	07/08/2026	Kostas Cheliotis	Added Section 1, 5, 2.2
0.7	07/08/2026	Kostas Cheliotis	Proofreading, Formatting
0.8	26/08/2026	Kostas Cheliotis, Panagiotis Fafoutellis, Adel Almohammad	Addressing reviewers' feedback

Funded by the European Union

Project funded by

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

1.0	26/08/2026	Kostas Cheliotis	Proofreading, Formatting

LEGAL DISCLAIMER

Funded by the European Union

Project funded by

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

TABLE OF CONTENTS

TABLE OF CONTENTS 4
LIST OF FIGURES 6
LIST OF TABLES 7
LIST OF ABBREVIATIONS AND ACRONYMS 8
PROJECT EXECUTIVE SUMMARY..10
DELIVERABLE EXECUTIVE SUMMARY..11
1 INTRODUCTION..12
1.1 MAPPING FEDORA OUTPUTS 12
1.2 DELIVERABLE OVERVIEW AND REPORT STRUCTURE 13
1.2.1 Deliverable Overview 13
1.2.2 Report Structure 14
2 COMPLEX EVENT PROCESSING ENGINE 15
2.1 GENAI FOR INCIDENT DETECTION, GENERATION, AND AUTOMATED LABELLING 15
2.1.1 TimeGAN Architecture..15
2.1.2 Evaluation Metrics 17
2.1.3 Dataset Description and Preprocessing..18
2.1.4 Model Training 19
2.1.5 Data Postprocessing..24
2.1.6 Conclusions 26
2.2 NETWORK EVOLUTION FORECASTING AND ANOMALY DETECTION 27
2.2.1 Methodology..27
2.2.2 Network Evolution Forecasting Service 29
2.2.3 Anomaly Detection Service 31
2.2.4 Limitations and Next Steps 32
2.3 COMPUTER VISION FOR EVENT PROCESSING 33
2.3.1 Objectives and role within Task 4.4..33
2.3.2 Operational concept and architecture..33
2.3.3 Altitude-aware detection and performance emulation 34
2.3.4 Traffic indicators and candidate event cues..35
2.3.5 Data interfaces and CEPE integration 35
2.3.6 Operational requirements and current status..36
2.3.7 Assumptions and next steps..37
3 SYNCHROMODAL OPTIMIZATION ALGORITHMS..38
3.1 SUMMARY..38
3.2 CONFLICT AVOIDANCE AND ROUTE OPTIMISATION FOR SYNCHROMODAL NETWORKS..38
3.3 LESSONS LEARNED & NEXT STEPS 47
4 DYNAMIC RESPONSE PLANNER..49

4.1 INTRODUCTION 49
4.2 OBJECTIVES OF THE DYNAMIC RESPONSE PLANNER 49
4.3 PROPOSED APPROACH FOR GENERATING RESPONSE PLANS..51
4.3.1 Network Monitoring Tools 51
4.3.2 Response Plan Generator...52
4.3.3 Simulation Framework 52
4.4 OPERATIONAL WORKFLOW AND TECHNICAL ARCHITECTURE..52
4.4.1 operational workflow 52
4.4.2 Dynamic Response Planner and FEDORA network optimisation services......................................57
4.4.3 Technical aspects of the Dynamic Response Planner...58
4.5 NEXT STEPS 59
5 CONCLUSION 60
6 REFERENCES 61

Funded by the European Union

Project funded by
Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
Swiss Confederation
Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

LIST OF FIGURES

Figure 1. TimeGAN architecture representation Yoon et al. (2019a). Real sequences x are embedded and reconstructed via (E,R) (reconstruction loss L_{rec}). Synthetic latent sequences are generated and refined by (G,S) (supervised loss L_{sup}). The discriminator C distinguishes real h from synthetic ĥ (adversarial loss L_{adv}). 17

Figure 2. Loop Detectors' locations 19

Figure 3. Principal Components Analysis (PCA) Visualization of Single-Sensor Original and Synthetic Data 21

Figure 4. t-SNE Visualization of Single-Sensor Original and Synthetic Data 22

Figure 5. Principal Components Analysis (PCA) Visualization of Multi-Sensor Original and Synthetic Data 23

Figure 6. t-SNE Visualization of Multi-Sensor Original and Synthetic Data 24

Figure 7. Speed comparison of real vs. indicative synthetic sequences 25

Figure 8. Volume comparison of real vs. indicative synthetic sequences 26

Figure 9: Daily curve of hourly vehicle flow for a typical weekday 28

Figure 10: Daily curve statistics for a typical weekday 28

Figure 11: Real-world value compared to the daily curve, with future values projected based on the difference 30

Figure 12: R-squared values for varying daily curve training dataset sizes 31

Figure 13: Daily curve with lower and upper quartiles as the bounds, and actual daily values with highlighted anomalies 32

Figure 14: Conceptual workflow of the UAV-CV component and its interface with the T4.4 Complex Event Processing Engine. 34

Figure 15: Example Synchromodal Network 39

Figure 16: Last-mile Delivery Operations 42

Figure 17: Vehicle-dependent Routing Data 43

Figure 18: Off-peak Case – No Delays at Junctions 44

Figure 19: Moderate Delay at Some Junctions 45

Figure 20: Significant Delay at Some Junctions 46

Figure 21: Input and output data of the Dynamic Response Planner 50

Figure 22: The conceptual architecture of the Dynamic Response Planner 51

Figure 23: The operational workflow of the Dynamic Response Planner 54

Figure 24: A screenshot for intersections with traffic lights in SUMO and a sample of python script for updating the schedules of a traffic light 55

Figure 25: A screenshot and a sample of python script for closing 1 lane and reducing the speed of 2 lanes of a 3-lanes link in SUMO 57

Figure 26: Two examples of response plans generated by the DRP 58

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation

Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

LIST OF TABLES

Table 1: Adherence to FEDORA GA Deliverable & Tasks Descriptions 13

Table 2: Traffic information products generated by UAV-CV 35

Table 3: UAV-CV input data 35

Table 4: UAV-CV output data 36

Table 5: VCR and the transition of a road link from free flow to a blocked road 56

Table 6: An example for rule-based mechanism to select the appropriate mitigation measures for congested road links 56

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation

Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

LIST OF ABBREVIATIONS AND ACRONYMS

Acronym	Meaning
CEPE	Complex Event Processing Engine
CV	Computer Vision
DRP	Dynamic Response Planner
ETA	Estimated Time of Arrival
ETD	Estimated Time of Departure
FPS	Frames Per Second
GA	Grant Agreement
GAN	Generative Adversarial Network
GenAI	Generative AI
GPS	Global Positioning System
GUI	Graphical User Interface
HGV	Heavy Goods Vehicle
LEZ	Low Emission Zone
KPI	Key Performance Indicator
MAE	Mean Absolute Error
NMT	Network Monitoring Tool
OD	Origin-Destination
PCA	Principal Component Analysis
RA	Response Action
ROI	Region-Of-Interest
RP	Response Plan
RPG	Response Plan Generator
SCF-VRO	Synchromodal Conflict-Free Vehicle Routing Optimiser
SUMO	Simulation of Urban MObility

Funded by the European Union
Project funded by
Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
Swiss Confederation
Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

TMO	Traffic Management Operator
TraCI	TRAffic Control Interface
TT	Travel Time
t-SNE	t-distributed Stochastic Neighbor Embedding
UAV	Unmanned Aerial Vehicle
ULEZ	Ultra Low Emission Zone
WP	Work Package
VCR	Volume-to-Capacity Ratio
VMS	Variable Message Sign
VRP	Vehicle Routing Problem
VRU	Vulnerable Road User
ZEZ	Zero Emission Zone

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

PROJECT EXECUTIVE SUMMARY

Lack of orchestration, structured and standardized integration protocols and metadata descriptors, incorporation of real-world traffic complexities and nuances, underutilization of valuable resources, model uncertainties and integration of micro-mobility services and VRUs result in suboptimal performance in addressing complex issues related to the management of mobility services and infrastructure and a divergence from EU's sustainable mobility targets. FEDORA aims to pave the way towards advanced traffic and network management through the development of a federated spaces platform offering a holistic framework of innovative solutions and services that enable precise and pro-acting sensing of supply and demand, facilitate optimal operation of transport services, and advances learning and evolution in complex environments. At the operational level, FEDORA offers a collaborative space of data that can realize advanced data alchemy processes using interconnected services and tools, a space of advanced traffic management optimisation services and a multi-modal simulation space to create and assess future mobility scenarios. The approach is validated in six thematic demonstrations in Vienna, the Basque country,Reggio Emilia, Nicosia, Budapest and Denmark, covering varying EU urban and rural contexts, infrastructure maturity levels, multimodal mobility services availability, organisation/operational structures and social conditions. Interaction with existing programmes on roadmapping and recommendations at national, EU and global level will be promoted, allowing a multiplication effect of project's results.

Social Media link:

@fedora-euproject

For further information please visit https://fedora-horizon.eu/

Funded by the European Union

Project funded by
Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
Swiss Confederation
Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

Funded by the European Union

Project funded by
Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
Swiss Confederation
Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

DELIVERABLE EXECUTIVE SUMMARY

Deliverable D4.2, Advanced Routing and Mitigation Strategies, presents the design and initial development of three interconnected components within FEDORA's Work Package 4 (Next Generation Multimodal Network Management Services Space), covering Tasks 4.4, 4.5, and 4.6. Together, these components form a suite of tools for real-time traffic event detection, multimodal freight routing optimisation, and dynamic incident response planning.

Complex Event Processing Engine (CEPE), developed under Task 4.4, provides a collection of tools for analysing temporal traffic data in real time. Three distinct capabilities are presented: (1) a Generative AI module based on the TimeGAN architecture, which generates realistic synthetic traffic incident sequences and lays the groundwork for automated data labelling using self-supervised learning; (2) a Network Evolution Forecasting and Anomaly Detection service that leverages statistical daily-curve models to predict short-term traffic conditions and flag deviations from expected patterns under normal operating conditions; and (3) a UAV-based Computer Vision module that emulates drone observations to produce structured vehicle detections, traffic indicators, and candidate anomaly cues for the Nicosia in-silico pilot. Collectively, the CEPE enables the detection of incidents and anomalies that trigger the downstream routing and response components.

Synchromodal Conflict-Free Vehicle Routing Optimiser (SCF-VRO), developed under Task 4.5, delivers a logistics planning system for managing freight flows across multimodal networks, with a focus on operational efficiency and resilience to disruptions. The SCF-VRO combines synchromodal freight models with last-mile delivery optimisation, incorporating vehicle-type-specific routing constraints, congestion-aware re-routing, and integration with open data sources such as OpenStreetMap. Designed around an interoperable multi-layered architecture, it supports a wide range of freight types and operational scenarios, including dynamic rerouting in response to real-time incidents such as vessel delays or port congestion.

Dynamic Response Planner (DRP), developed under Task 4.6, acts as an orchestration and decision-support layer that generates comprehensive response plans when traffic degradation is detected or predicted. The DRP integrates outputs from the CEPE, the Social Optimum Models (Task 4.1), the Multimodal Journey Planner (Task 4.3), and the SCF-VRO (Task 4.5) into a rule-based, iteratively refined mitigation strategy. Using SUMO as a simulation backbone, the DRP evaluates network scenarios, identifies congested links via volume-to-capacity ratios, and proposes targeted response actions such as lane closures and speed limit adjustments, until optimal network capacity is restored.

At this stage of the project, the three components have been developed primarily in isolation to establish robust proof-of-concept implementations and initial calibrations. The next phase will focus on integration with real-world pilot data from Vienna and Nicosia, deployment alongside other WP4 and WP5 solutions, and further refinement through practical use case validation across FEDORA's six demonstration cities.

Funded by the European Union

Project funded by

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Federal Department of Economic Affairs,
Education and Research EAER
State Secretariat for Education,
Research and Innovation SERI

1 INTRODUCTION

1.1 MAPPING FEDORA OUTPUTS

The table below shows in detail how this deliverable addresses specific commitments of the FEDORA project as identified in the Grant Agreement (GA). This is achieved by mapping specific sections of this report to the corresponding section in the GA, and by providing a brief justification as to how each GA commitment is addressed in the relevant report section.

FEDORA ga component title	FEDORA GA COMPONENT OUTLINE	RESPECTIVE DOCUMENT CHAPTER(S)	JUSTIFICATION
	DELIVERABLE		
D4.2 – Advanced Routing & Mitigation Strategies	Report and demonstration on conflict-free routing solutions & dynamic mitigation for multimodal networks (T4.4-T4.6)	Whole document	
	TASKS		
Task 4.4 - Usage patterns, bottlenecks, and anomalies detection	Methods such as Generative Adversarial Networks (GANs) will be used to generate synthetic data based on traffic flow theory and simulations, thereby enriching existing datasets. A self-supervised approach will be introduced for non-recurrent event detection to identify complex patterns in the data and generate labels	Section 2.1	Section 2.1 describes how TimeGANs are used to analyse time series traffic data in order to detect incidents and generate synthetic datasets, and lays the pathway for expanding the approach into automated data labelling.
Task 4.4 - Usage patterns, bottlenecks, and anomalies detection	Models will be trained using datasets that represent typical or recurrent traffic patterns, allowing for the identification of anomalies. Historical event sequences will be analysed to uncover underlying patterns and relationships, predicting their evolution on the network.	Section 2.2	Section 2.2 describes the development of tools for the analysis of typical traffic data, and their application in forecasting the evolution of network conditions and detecting anomalies in traffic patterns.
Task 4.4 - Usage patterns, bottlenecks, and anomalies detection	Computer vision techniques for usage patterns, bottlenecks, and anomaly detection using drone data will be developed.	Section 2.3	Section 2.3 presents a detailed description of methods for capturing traffic data using computer vision

Funded by the European Union

Project funded by
Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
Swiss Confederation
Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

			techniques from unmanned aerial vehicles.
Task 4.5 - Conflict-free routing solutions for synchro- and multi-modality	Will provide new routing and scheduling solutions for freight flows in cities and urban areas, enabling dynamic multimodal rerouting, using a conflict-free approach.	Chapter 3	Chapter 3 describes the progress in developing algorithms to enable synchromodal re-routing options for freight in urban areas in response to detected incidents and delays.
Task 4.6 - Dynamic and responsive planning of mitigation measures	Will develop a system that generates real-time optimal response plans for traffic incidents and congestion, using a rule-based engine where different response actions will be integrated together to form a plan. Performance monitoring tools will continuously refine the recommendations.	Chapter 4	Chapter 4 provides a detailed description and architecture for the Dynamic Response Planner which generates network-wide response plans in response to detected traffic degradation, along with the iterative evaluation framework which ensures response plan validity.

Table 1: Adherence to FEDORA GA Deliverable & Tasks Descriptions

1.2 DELIVERABLE OVERVIEW AND REPORT STRUCTURE

1.2.1 DELIVERABLE OVERVIEW

This deliverable titled “Advanced Routing & Mitigation Strategies” presents tools and systems developed within the scope of WP4: “Next Generation Multimodal Network Management Services Space” that, as the report name suggests, aim to address two core processes: vehicle routing in highly constrained scenarios, and the generation of mitigation measures in response to detected traffic degradation. These two correspond directly to the work undertaken as part of Tasks 4.5 and 4.6 respectively as identified in the GA, with Task 4.5 developing algorithms for multi-and synchro-modal freight routing solutions, and Task 4.6 developing an engine to generate and validate comprehensive mitigation plans in response to traffic incidents.

In support of the two core aims presented above, this deliverable also presents the Complex Event Processing Engine (CEPE), which includes work undertaken as part of Task 4.4. While this task is not mentioned explicitly in the title or description of this deliverable, it nevertheless constitutes an integral part that is highly relevant to the other two, as its core aim is develop tools that collect, analyse, and generate traffic data, enabling the detection of traffic incidents. As such the CEPE is considered as a requirement for the other two, as it is used for detecting incidents that both the synchromodal routing and dynamic response plan will need to respond to.

Furthermore, this deliverable is highly linked to additional outputs in the FEDORA project, that are not included in this report. More specifically, additional network optimization systems are being developed as part of Tasks 4.1, 4.2, and 4.3, that are presented in detail in deliverable D4.1 - Optimization and Incentivization Models for Multimodal Networks. Outputs from those systems are relevant both to the synchromodal optimization algorithms, which takes into account multiple modalities that are addressed in those additional Tasks, as well as the dynamic response planner, that acts as a collector of multiple optimization strategies in order to generate comprehensive mitigation plans, and therefore the tools described in this deliverable will receive inputs from systems developed in Tasks 4.1, 4.2, and 4.3.

Funded by the European Union

Project funded by
Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
Swiss Confederation
Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

1.2.2 REPORT STRUCTURE

This report presents the relevant work using the following structure:

This chapter (Chapter 1) presents an overview of the whole report. The following chapter (Chapter 2) presents all work relevant to the Complex Event Processing Engine (CEPE), broken down into three separate sections. Section 2.1 describes the development on Generative AI tools for incident detection, synthetic data generation, and automated data labelling. Section 2.2 focusses on forecasting and anomaly detection tools that operate on typical or recurrent traffic data. Finally Section 2.3 details computer vision techniques for gathering traffic data from UAVs. Chapter 3 outlines the development of algorithms for synchromodal freight routing optimizations. Chapter 4 presents the Dynamic Response Planner that generates and evaluates optimal action plans for traffic incident and congestion response. And finally Chapter 5 concludes this report, offering a summary and future steps.

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

2 COMPLEX EVENT PROCESSING ENGINE

The Complex Event Processing Engine (CEPE) refers to a collection of tools, all related to the process of analysing temporal traffic data, often in real-time. The various tools within CEPE handle data collection, incident and anomaly detection, traffic forecasting, synthetic data generation, and data labelling.

2.1 GENAI FOR INCIDENT DETECTION, GENERATION, AND AUTOMATED LABELLING

While the available traffic datasets provide millions of normal traffic points per day, the number of time windows corresponding to road incidents is very limited and mostly not annotated. As a result, any supervised learning model trained directly on the raw data would be heavily biased towards non-incident behaviour. Moreover, real-world restrictions (privacy, safety, cost) limit the ability to deliberately collect more incident data. To overcome the scarcity of incident data, researchers often turn to synthetic traffic data generation, aiming not to replace real observations, but to augment them with realistic scenarios that preserve the underlying temporal and cross-location dynamics of the traffic system. Synthetic data balance datasets for supervised learning and contribute to forecasting and anomaly detection models. Also, they support the simulation of incident scenarios.

This work aims to fill the gaps by creating a data-driven generative approach for traffic events on urban arterials. The paper investigates the application of the TimeGAN (Yoon et al. 2019)[1], a Generative Adversarial Network (GAN) tailored for time series, to create realistic synthetic sequences of traffic speed and flow during and after incidents. The framework is extended with site conditioning for multi-sensor corridors and evaluated with discriminative/predictive scores, as well as per-site temporal statistics, following best practices. For these purposes we utilize a dataset that consists of high resolution (per 90 seconds) measurements from the sensors on Kifisos Avenue in Athens, Greece to create event-like scenarios to the data and provide researchers with ways to analyze how the system moves when in congested and incident states.

Furthermore, this work will be extended with incident detection and data-labeller modules, based on the principles of Self-Supervised Learning, which will constitute a generic and comprehensive non-recurrent incident processing engine.

2.1.1 TIMEGAN ARCHITECTURE

Trying to apply a standard GAN directly to time series may apparently produce realistic samples but the patterns over time are not sufficiently depicted. TimeGAN extends the foundation of GANs to time series data by combining adversarial training with sequence modeling. The input to TimeGAN is a set of aligned, fixed-length, normalized multivariate time series sequences.

The key differences between a standard GAN and TimeGAN are:

- Data type: instead of static x, the input is a time series sequence (e.g., traffic speed and flow).
- Enriched loss functions: Reconstruction loss (autoencoder) and supervised temporal loss are added to capture sequence dependencies and enforce latent temporal consistency. The embedding at time t is trained to predict the embedding at time t + 1.
- Architecture: instead of MLPs, TimeGAN uses RNNs (GRUs/LSTMs) to handle temporal data.

Specifically, the TimeGAN architecture consists of interacting neural networks (RNNs): an encoder E, a decoder R, a generator G, and a discriminator D. Input data are represented by two types of features: static features s and temporal features x, which are fed into

[1] https://github.com/jsyoon0823/timegan

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

the encoder to obtain latent representations. In this way, TimeGAN captures both the temporal dependencies and cross-variable dynamics of multivariate time series such as traffic data.

- Embedder (E): Encoder, maps real series to the latent space.
- Recovery (R): Decoder, maps latent representations back to reconstructed series.
- Generator (G): creates synthetic latent sequences.
- Discriminator (D): distinguishes real vs. synthetic sequences.
- Supervisor (S): Learns temporal dynamics in latent space.

The generator receives tuples of random static and temporal latent variables (zs , zx) drawn from a known prior distribution and produces synthetic latent sequences $\hat{h}$. The encoder maps real sequences to latent codes h, which, together with the generated $\hat{h}$, are used in both supervised and adversarial training. The discriminator D takes as input pairs of latent sequences (real or synthetic) and outputs the probability of each being real. Formally, the architecture minimizes three complementary loss functions:

$$L_{unsupervised} = E[logD(h)] + E\left[log\left(1 - D(\hat{h})\right)\right], \ (2)$$

$$L_{supervised} = E\left[\left\|h_{t+1} - \hat{h}_{t+1}\right\|_2^2\right] \ (3)$$

$$L_{reconstruction} = E\left[\left\|x - R(E(x))\right\|_2^2\right], \ (4)$$

where Eq. (2) corresponds to the adversarial objective, Eq. (3) enforces temporal consistency between consecutive latent steps, and Eq. (4) ensures faithful sequence reconstruction through the encoder–decoder pair. During training, the model switches between adversarial updates (generator-discriminator) and supervised reconstruction updates, using an Adam optimizer and a batch size of 128.

Compared to classic GANs that learn directly from raw data, TimeGAN initially forms latent representation via the autoencoding framework (Embedder–Recovery). This separates feature extraction from generative learning and enhances temporal consistency. The Supervisor component establishes temporal self-consistency by forecasting the subsequent latent state based on the current one, directing the generator to obtain sequential dependencies. Empirical assessments indicate that TimeGAN generates more authentic and temporally coherent sequences than previous time-series GANs, while enhancing interpretability based on the latent space representation. An illustration of the model's architecture is provided in Figure 1.

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation

Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

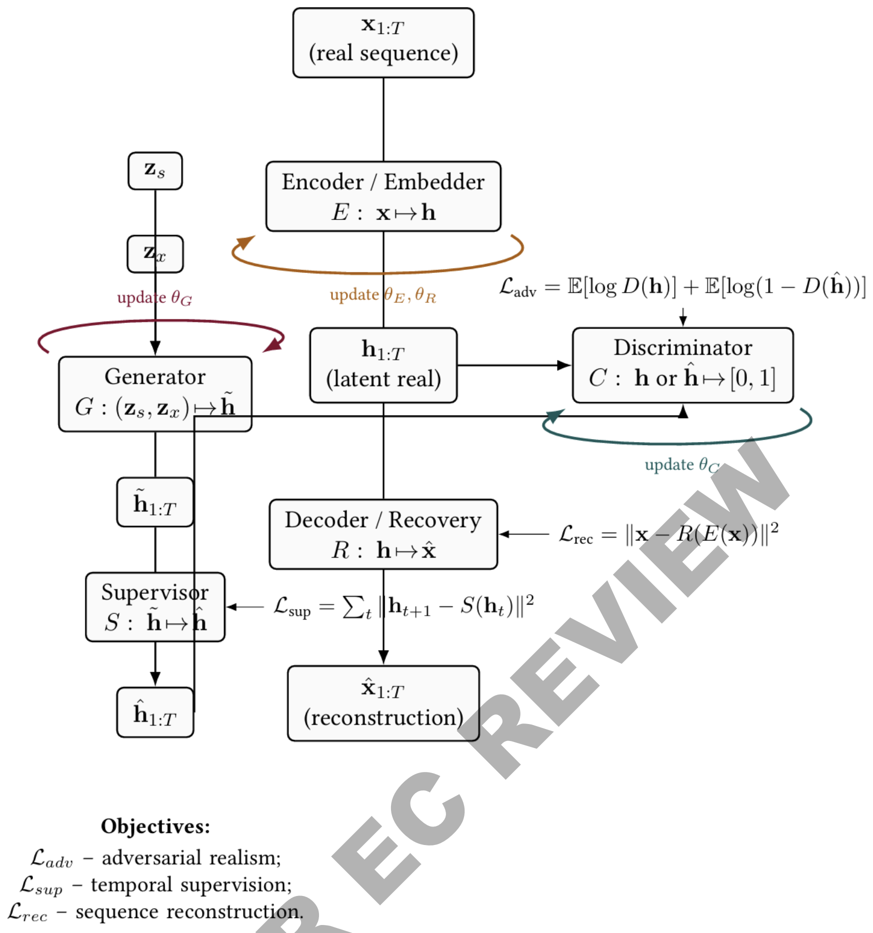

Figure 1. TimeGAN architecture representation Yoon et al. (2019a). Real sequences x are embedded and reconstructed via (E,R) (reconstruction loss L_{rec}). Synthetic latent sequences are generated and refined by (G,S) (supervised loss L_{sup}). The discriminator C distinguishes real h from synthetic $\hat{h}$ (adversarial loss L_{adv}).

2.1.2 EVALUATION METRICS

The evaluation focuses on the discriminative and predictive scores. The discriminative score is a quantitative metric proposed in the original TimeGAN research to depict the similarity between real and synthetic time-series data. A recurrent neural network (RNN) classifier is trained to distinguish between real and generated samples. The metric calculates the absolute deviation of the classifier's accuracy from random chance (|accuracy − 0.5|). Based on the definition, values closer to zero indicate that the classifier cannot reliably distinguish between the two datasets (higher realism of the synthetic data).

The predictive score is also a quantitative metric proposed in the original TimeGAN research to check if the synthetic data have the ability to respond to forecasting tasks. In simple terms, it evaluates whether a model trained only on generated data is capable of predicting future values of real data. A recurrent neural network (RNN) predictor is trained to predict the following step using the synthetic dataset and is then tested on the original data. The results are assessed based on the Mean Absolute Error (MAE) between the predicted and actual real values.

Also, a common way to visualize the outcomes of the applied methods on the data is the dimensionality reduction techniques. In TimeGAN, Principal Component Analysis (PCA) and t-distributed Stochastic Neighbor Embedding (t-SNE) are very useful to present both real and synthetic data in two-dimensional space. Principal Component Analysis (PCA) is a linear technique that projects high-dimensional data to orthogonal components capturing the directions of maximum variance. In contrast, t-SNE is a non-linear, probabilistic algorithm that minimizes the divergence between similarity distributions in high- and low-dimensional spaces. The combination of those two methodologies provides the necessary confidence for the model's success.

2.1.3 DATASET DESCRIPTION AND PREPROCESSING

The data of this work are provided from sensors on Kifisos Avenue, one of the most significant and crowded highways in Athens, Greece, connecting the northern and southern areas of the city. Traffic incidents in Kifisos Avenue can quickly spread throughout the system, causing long delays and heavily congested conditions. Traffic incident modeling in Athens has not been well examined, especially with data-driven methodologies, despite the significance of this topic. Current conventional approaches frequently depend on aggregated averages or thresholding methods, which may not adequately reflect the complexity and diversity of traffic imbalances.

The 4 monitoring sensors of the dataset along Kifisos Avenue, correspond to points where several traffic incidents have been identified. For each location, measurements are available in two time series variables: speed (average vehicle speed, in km/h) and volume (traffic flow rate, in vehicles/hour), both sampled every 90 seconds. Their locations are provided in Figure 2.

Furthermore, an additional preprocessing step was conducted in order to annotate the measurement pairs according to the work of (Papadatou et al., 2025), and detect incident-related time windows. Specifically, following the unsupervised learning framework proposed, four distinct clusters of traffic conditions were observed based on the residual deviations of speed and volume from their expected profiles. Cluster 0 corresponds to recurrent or typical traffic conditions, forming the normal daily traffic pattern. Cluster 1 captures light traffic or undersaturated conditions, where speeds are higher and volumes lower than expected, typical image during weekends or off-peak hours. Cluster 2 represents transitional or boundary states between free-flow and congestion. Finally, Cluster 3 is linked to non-recurrent, incident congestion, showing negative residuals in speed and reduced or unstable volume, commonly associated with accidents, breakdowns, or other unexpected disturbances.

The above step was necessary as the overall objective of this work is to train a model to generate realistic time series of speed and flow for all locations, under the condition that an incident occurs (Cluster 3) at one of them, affecting rest of the road network. In other words, the model should capture dependencies between sites and temporal patterns during extreme conditions, allowing the simulation of multilocation traffic dynamics.

Funded by the European Union

Project funded by
Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
Swiss Confederation
Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

Figure 2. Loop Detectors' locations

2.1.4 MODEL TRAINING

The synthetic traffic incident generation model is built upon the official implementation of the TimeGAN architecture, as provided by Yoon et al. (2019). The repository tutorial[2] served as the baseline pipeline, which was then customized to the specific characteristics of the incident-data traffic set, to match the residual speed/volume representation of incident conditions. After the adjustments (parameter connection to the new dataset), the model was trained on the incident subset and later validated to ensure that the generated synthetic sequences represent consistent with the real incident dataset, temporal dynamics.

2.1.4.1 DATA PREPARATION AND EXPERIMENTAL DESIGN

The initial phase of the implementation focused on identifying the most suitable data structure to serve as input to the TimeGAN model. Several preliminary configurations were tested to ensure compatibility with the model's numerical input requirements. These checks included the inclusion of boolean flags representing the state of each sensor site, the exploration of various data encoding schemes, and the selection of appropriate temporal windows surrounding each incident to capture the complete duration of the event. Through these experiments, the objective was to construct a representation that preserved both spatial relationships between sensors and the temporal evolution of traffic conditions during incident periods, taking into account the input constraints of the TimeGAN architecture.

The investigation led to two final input datasets. Initially, a dataset was created with only the pure traffic variables (speed, volume) and the modified time steps data, without including any sensor site identifiers. This configuration was intended to evaluate the model's capacity to learn and reproduce incident-related dynamics for a specific location, rather than on a network-wide scale.

Furthermore, a second dataset was created to represent incident periods in a system (corridor). A timestep is considered as an incident moment when at least one sensor site indicates an incident (cluster=3). In this configuration, the sites that specifically experienced the incident were not flagged individually; instead, the entire sensor system was treated as jointly reflecting an incident

[2] https://github.com/jsyoon0823/timegan

Funded by the European Union
Project funded by
Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra
Swiss Confederation
Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

momentum. This approach enables an assessment of the TimeGAN model's ability to generalize from local disturbances to collective traffic states, i.e., non-recurrent events in the system.

The sequence length in the TimeGAN model was set to 24 time steps, which corresponds to 36 minutes given the 90-second interval of the traffic dataset. This configuration was chosen after analyzing the duration of the incidents recorded on all sensors, where the median duration ranged between 30 and 45 minutes. Therefore, a 36-minute window was found to effectively capture the core temporal dynamics of most typical duration incidents, while avoiding the instability and vanishing gradient issues that often occur when training generative models on very long sequences. Additionally, this length ensured a balanced trade-off between temporal coverage and computational efficiency, allowing the model to learn multi-sensor dependencies without over-fitting to longer, rare disruptions. The resulting configuration provided stable training and produced synthetic sequences that closely matched the temporal evolution and variability of real incidents.

2.1.4.2 SOLUTION 1 - SINGLE-SENSOR CONFIGURATION

The first application of the TimeGAN method was independent of the site and the tags of each sensor and had as an input only the numerical info of the sensors as presented below:

- Input features: pure traffic variables including speed, volume and modified time steps for each individual sensor. Observations used independently, without cross-sensor information
- Scaling: all numerical features are normalized to the range [0,1] using a min–max normalization function and inverse-transformed after training
- Recurrent module: GRU-based architecture
- Hidden dimensionality: 24 units per recurrent layer
- Network depth: 3 stacked recurrent layers
- Training iterations: 10,000 update steps
- Batch size: 128 samples per mini-batch
- Output: synthetic sequences of speed and flow of an incident at a specific sensor site

The training process followed the three-stage procedure of the TimeGAN framework: the embedding network training, supervised loss pre-training and joint adversarial training. During the embedding stage, embedding loss decreased with a steady pace during runs, from around 0.26 to below 0.01 after 10,000 iterations, indicating that the network successfully learned to reconstruct temporal dependencies within the input sequences. During the supervised training stage, supervised loss stabilized around 0.018, validating that the recurrent generator captured short-term dynamics. Finally, during the joint training stage, the discriminator and generator losses remained within stable ranges, consistent with adversarial learning behavior. The embedding loss term converged toward zero, demonstrating consistency between the embedding and generator networks. After completing 10,000 training rounds, the model successfully generated synthetic time-series data that reflect incident and non-recurrent traffic conditions.

In the implementation, the discriminative score is calculated as proposed in the official paper. The discriminative evaluation is repeated 5 times to ensure precision of the result against stochastic variability in training. At each iteration, the RNN discriminator is retrained on the original and synthetic sequences and the classification accuracy is recorded. The reported discriminative score is the mean of these 5 independent evaluations, computed as

$$DiscriminativeScore = \frac{1}{N}\sum_{i=1}^{N}|Accuracy_i - 0.5|, whereN = 5 \quad (5)$$

For the present solution, the resulting discriminative score was 0.0423, which is close to zero, indicating that the synthetic data generated are very close to the original data and that the model achieved high generative accuracy.

In the implementation correspondingly, the evaluation process is repeated 5 times with independent random iterations to ensure stable and reliable results. For each run, the MAE is computed and stored in a list, and the final predictive score is calculated as the average of these values:

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation

Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

$$PredictiveScore = \frac{1}{N}\sum_{i=1}^{N} MAE_i, whereN = 5 \ (6)$$

The predictive score was found to be 0.0803, a very low mean absolute error and indicating that the synthetic data accurately copy the structure of the real traffic sequences.

Finally, as presented at the theoretical approach section, two very well-known dimensional reduction methods will be used to visualize the results of TimeGAN model. Figure 3 and Figure 4 present the PCA and t-SNE plots for the trained TimeGAN model based on single sensor dataset. In the PCA plot, the distributions of the original (red) and synthetic(blue) data points significantly overlap, indicating that the generated data capture the main variance structure of the real dataset. Similarly, the t-SNE plot demonstrates that the model successfully preserved both local and global temporal dependencies of the original traffic sequences. Both visual results confirm the model's capacity to generate synthetic data with statistical properties closely aligned with those of the real observations. One may additionally notice that the PCA plot exhibits more outliers that were not very accurately captured, whereas the specific issue is less noticeable in the t-SNE plot, which indicates that the linear PCA projection cannot effectively capture the nonlinear structure of the synthetic data.

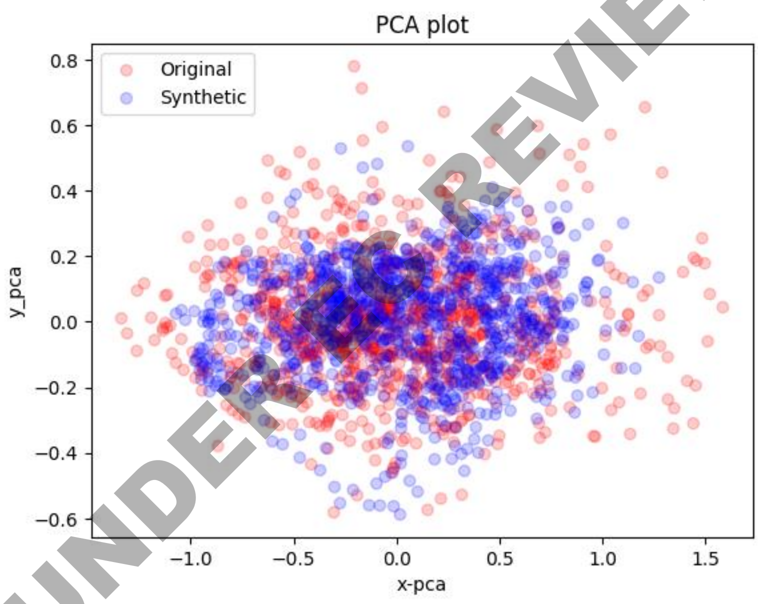

Figure 3. Principal Components Analysis (PCA) Visualization of Single-Sensor Original and Synthetic Data

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

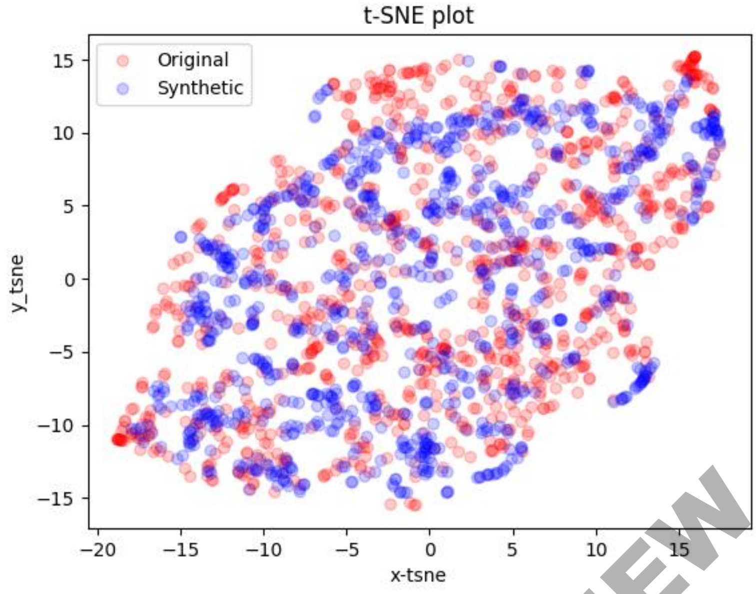

Figure 4. t-SNE Visualization of Single-Sensor Original and Synthetic Data

2.1.4.3 SOLUTION 2 - MULTI-SENSOR CONFIGURATION

The following application of the TimeGAN method treated the data as a system of sensors that explained an incident observation and had as an input only the combined numerical info of the sensors as presented below:

- Input features: traffic variables including speed, volume and modified time steps from all monitored sensors. Each time step represents the overall system state rather than individual site data
- Scaling: all numerical features are normalized to the range [0,1] using a min–max normalization function and inverse-transformed after training
- Recurrent module: GRU-based architecture implementing gated recurrent units for temporal modeling
- Hidden dimensionality: 24 units per recurrent layer
- Network depth: 3 stacked recurrent layers
- Training iterations: 10,000 update steps
- Batch size: 128 samples per mini-batch
- Output: synthetic sequences of speed and flow for all sensors representing incident conditions across the network

The training process for the multisensor configuration followed the same three-stage procedure as defined in the TimeGAN framework: embedding network training, supervised pre-training, and joint adversarial training. During the embedding phase, the embedding loss decreased steadily from approximately 0.26 to 0.017 after 10,000 iterations, demonstrating that the model effectively learned the latent temporal structure across all monitored sensor sites. In the following supervised training stage, the supervised loss showed a smooth decline from 0.205 to around 0.015, indicating that the recurrent generator successfully captured the temporal dependencies across the multivariate space. During the final joint adversarial training phase, the discriminator and generator losses stabilized within the expected ranges, while the embedding reconstruction loss term converged to 0.01. The behavior of all loss functions during training confirms that the model has a stable learning behavior. In conclusion, multisensor TimeGAN successfully generated synthetic time-series data that capture both spatial and temporal correlations of an incident in a system.

In the multisensor configuration, the discriminative score was similarly calculated following the same methodology proposed in the official TimeGAN framework. The discriminative evaluation was repeated 5 times to address the stochastic variability in training. The resulting discriminative score was 0.1682. This value, while slightly higher than that of the single-sensor configuration, still indicates that the generated data are difficult to distinguish from the original sequences, confirming the model's ability to learn realistic patterns across multiple sensor inputs.

The predictive score was also computed through 5 independent evaluations to assess the ability of the synthetic data to generalize temporal dependencies observed in the real data. The resulting predictive score was 0.0578, demonstrating a very low forecasting error and assuring that the synthetic sequences preserve the structure of the real traffic sequences.

Figure 5 and

Figure 6 display the PCA and t-SNE plots for the multisensor TimeGANmodel. The PCA visualization shows that the synthetic data overlap with the original samples, demonstrating that the model captured the main variance structure of the sensors' system. This conclusion is further supported by the t-SNE representation, which shows that the actual and synthetic data points throughout the latent space preserve the relationships. Together, these visual results validate the model's capability to generate realistic, system-level traffic patterns under incident conditions. It is also worth mentioning that the multisensor approach better reproduces the real data compared to the single sensor one, indicating that including measurements from nearby sensor allows the model to generate more accurate synthetic data.

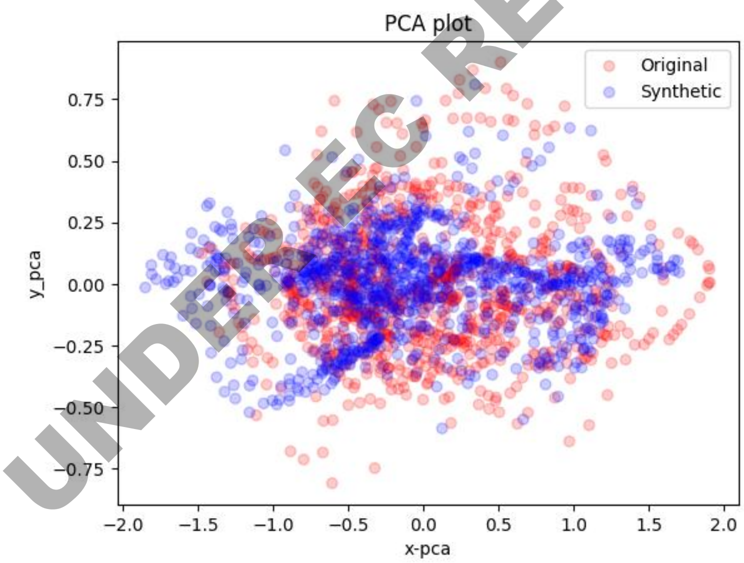

Figure 5. Principal Components Analysis (PCA) Visualization of Multi-Sensor Original and Synthetic Data

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

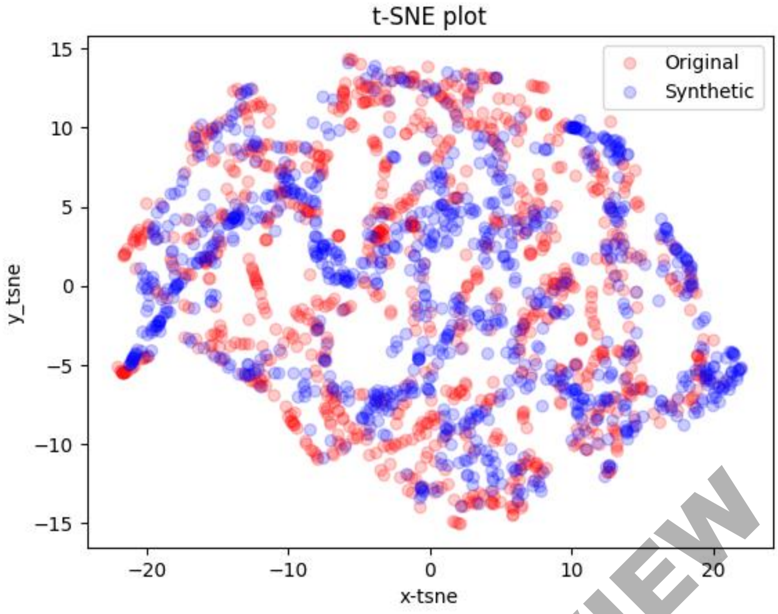

Figure 6. t-SNE Visualization of Multi-Sensor Original and Synthetic Data

2.1.5 DATA POSTPROCESSING

The trained TimeGAN framework has generated synthetic traffic sequences that replicate the behavior of the original dataset. In order to comply with the model's training settings, the output sequences were first generated in normalized form (exported in accordance with the model input). Aiming to restore an interpretable result, all generated data were inverse-transformed using the previously fitted min–max scaling parameters, to form the original numerical range for both speed and volume variables.

To evaluate the realism of the multi-site synthetic incidents generated by the TimeGAN model, a comparison was conducted between synthetic and real multi-site traffic data. Because synthetic sequences are not aligned with specific timestamps, the algorithm searches along the entire real dataset to find the time interval whose pattern most closely resembles the synthetic one.

For every synthetic sequence s of length L_s , a sliding window of equal length is moved across the real dataset and the mean absolute error (MAE) is calculated at each possible position t:

$$MAE_{i,f}(t) = \frac{1}{L_s}\sum_{k=0}^{L_s-1}\left|r_{i,f}(t+k) - s_{i,f}(k)\right|, \quad (7)$$

where $r_{i,f}$ and $s_{i,f}$ represent the real and synthetic values respectively for sensor site i ∈ {MS106, MS114, MS118, MS126} and feature f ∈ {speed, volume}. The overall similarity for each window is given by the average MAE across all sites and features:

$$MAE_{avg}(t) = \frac{1}{8}\sum_i \sum_f MAE_{i,f}(t) \quad (8)$$

The algorithm then identifies the position t where this average error is the smallest, where the two sequences are most similar in shape and magnitude. The corresponding minimum error value MAE represents how well the synthetic sequence matches its best real equivalent. A smaller MAE therefore indicates a more realistic synthetic sequence.

Each synthetic sequence obtains a single score MAE and the corresponding best alignment index t. Sequences are ranked in ascending order of MAE (the top scoring were retained as the most realistic multi-site sequences). For the top sequences, the corresponding real window $(r_{t:t+L_s-1})$ and synthetic sequence s are plotted using a common normalized time axis $\tau \in [0,1]$. Each figure is structured as a 2×2 grid: the four sub-plots correspond to the four sensor sites, while separate figures are generated for the

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

two variables (speed and volume). Solid lines represent real data and dashed lines represent synthetic sequences. Uniform y-limits are applied across sensors for each feature, ensuring consistent scale comparisons.

Figure 7 and Figure 8 depict the comparison between real and synthetic multi-site timeseries for the best-ranked synthetic sequences according to the average MAE metric. The analysis focuses on the model's capacity to reproduce both temporal patterns and inter-site relationships across the four monitored sensors. Of course, the goal here is not to recreate the exact same conditions as in the recorded incident (like duplicating the data), but rather understand the pattern and produce synthetic data with the same features. As it can also be visually observed, the model seems capable of producing realistic incident time series.

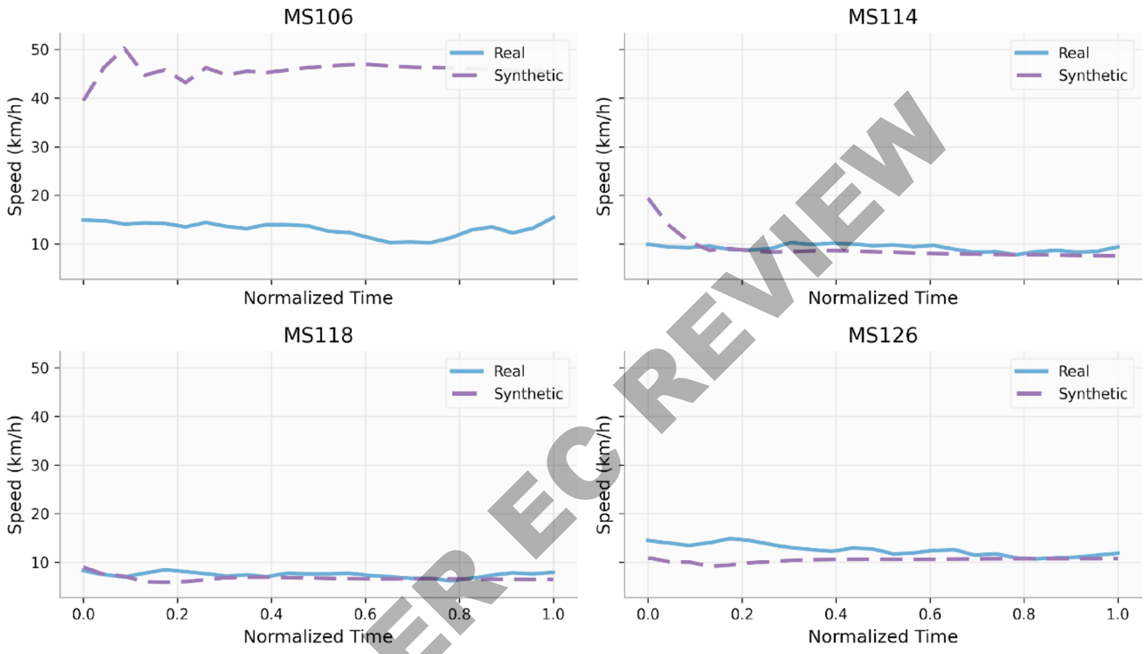

Figure 7. Speed comparison of real vs. indicative synthetic sequences

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

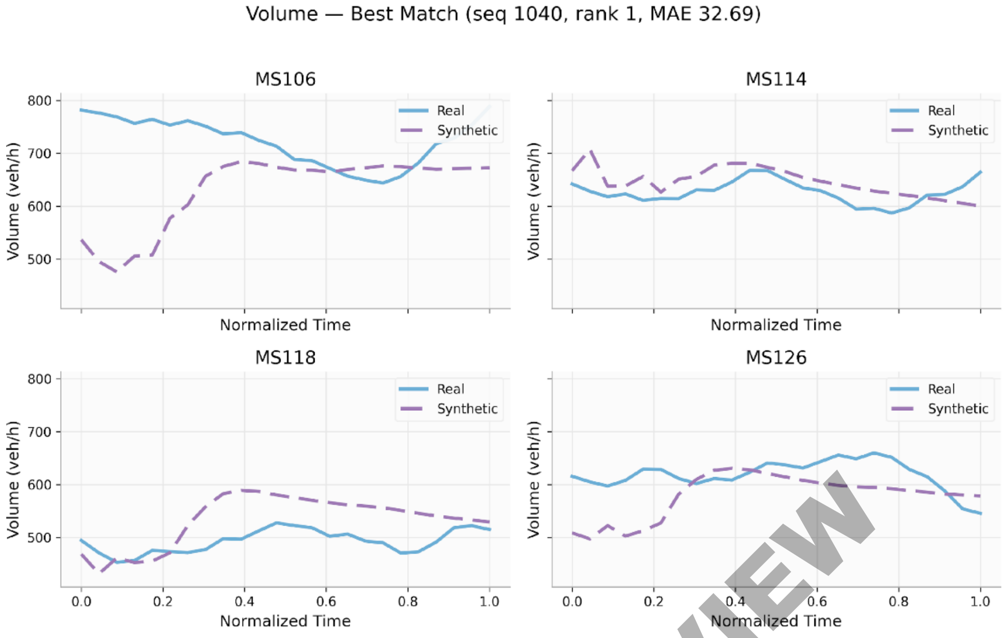

Figure 8. Volume comparison of real vs. indicative synthetic sequences

2.1.6 CONCLUSIONS

This work explored the generation of synthetic traffic sequences using the TimeGAN architecture, applied to a multi-sensor traffic dataset containing both normal and incident conditions. The results were found to be well-structured in feature space, displaying both local sensor behavior and global inter-site relationships. These findings supported the choice of the TimeGAN approach, which was found capable of learning complex temporal and spatial relationships among speed and volume observations. The study also demonstrated that TimeGAN can be effectively applied to real-world traffic data to generate realistic, interpretable, and statistically valid synthetic sequences, laying the groundwork for scalable and data-augmenting methodologies that can strengthen the resilience and adaptability of modern intelligent transportation systems.

More specifically, the generated synthetic data showed a strong similarity to real traffic sequences across multiple evaluation perspectives. In quantitative terms, the discriminative and predictive metrics (approximately 0.17 and 0.06, respectively) indicated that the synthetic data were both statistically realistic and predictive of true traffic evolution. Using a sliding-window matching procedure, the best synthetic sequences achieved low mean absolute errors for both speed and volume, confirming resemblance to real data. The t-SNE and PCA visualizations demonstrated that the synthetic samples overlapped with the real feature distributions, indicating that TimeGAN successfully learned the latent structure of the dataset.

Analysis across the four sensor sites revealed that MS114 and MS118 were best depicted, with smooth temporal representation and consistent correlation between speed and volume. Site MS106 displayed larger range variation. However, the multi-site sequences maintained the joint behavior of real incidents as well as relative sensor hierarchy. The qualitative examination of the top-ranked synthetic sequences further verified the realism of the generated incidents.

The incident processing engine will be further enhanced with techniques for real-time detection and data labelling, based on the principles of self-supervised learning. Moreover, the developed TimeGAN model will be further improved and assessed regarding training instability and other issues inherent to adversarial learning, such as mode collapse.

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation

Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

2.2 NETWORK EVOLUTION FORECASTING AND ANOMALY DETECTION

2.2.1 METHODOLOGY

2.2.1.1 AIMS

The core aim of the network evolution forecasting and anomaly detection tools is to develop a set of tools that support transport network management during typical operating conditions. As such they focus exclusively on capturing and predicting the state of traffic under normal conditions, and furthermore on detecting and predicting when traffic conditions start to deviate and exhibit irregular and unexpected patterns.

Additionally, from a systems-operational point of view, the main aim is to develop a set of tools that can be reliably applied to multiple locations and scenarios, given the number of use cases and pilots that are expected to make use of the Complex Event Processing Engine. For this reason, the core design approach is to develop general-purpose tools, that can operate with a minimal set of inputs (i.e. avoid developing tools that rely on specialized data that is only available in specific locations or through specialized tools).

For the purpose of calibrating and demonstrating the progress at this stage, we have applied the processes described here to a sample dataset of traffic sensor data for a major highway in Luxembourg. The data is publicly available and provided in real-time in DATEX2 format, via the Luxembourg open data platform[3].

2.2.1.2 DAILY CURVES APPROACH

As the main focus is to operate on typical traffic conditions, the first goal is to establish a model that captures such typical network traffic conditions. Typical conditions by definition are not expected to vary significantly under similar demand; Furthermore, traffic demand exhibits strong daily periodicity; therefore, we can analyse past data under normal demand to identify and extract traffic conditions as they evolve throughout a typical day. These traffic conditions are aggregated at a small temporal interval (e.g. every few minutes), and when plotted consecutively demonstrate how traffic conditions evolve throughout the day. We call this sequence of traffic conditions the “daily curve”, as it captures how traffic conditions change continuously throughout a typical day. An example of a daily curve for a typical weekday as calculated using data from multiple weekdays is presented in Figure 9.

The process for generating a daily curve is the following: Given a dataset of past traffic data in an area, we filter and group as needed to produce a dataset that contains days with similar traffic conditions. The parameters for filtering and grouping the dataset include type of day (weekday vs weekend), grouping by traffic sensor location, and grouping by traffic variable to be modelled (e.g. traffic flow, average speed, vehicle density).

Once these filters are applied, we arrive at a dataset that contains rows of traffic observations for the same location, at consistent temporal intervals throughout the day, for multiple days. Using this dataset we then calculate statistics for the same temporal interval across all days in the dataset, essentially generating descriptive statistics for each interval throughout the day, e.g. if using 5-minute interval we calculate average traffic flow during the first 5 minutes across all days, etc.

[3] https://data.public.lu/en/datasets/cita-donnees-trafic-en-datex-ii/

Funded by the European Union

Project funded by
Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
Swiss Confederation
Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

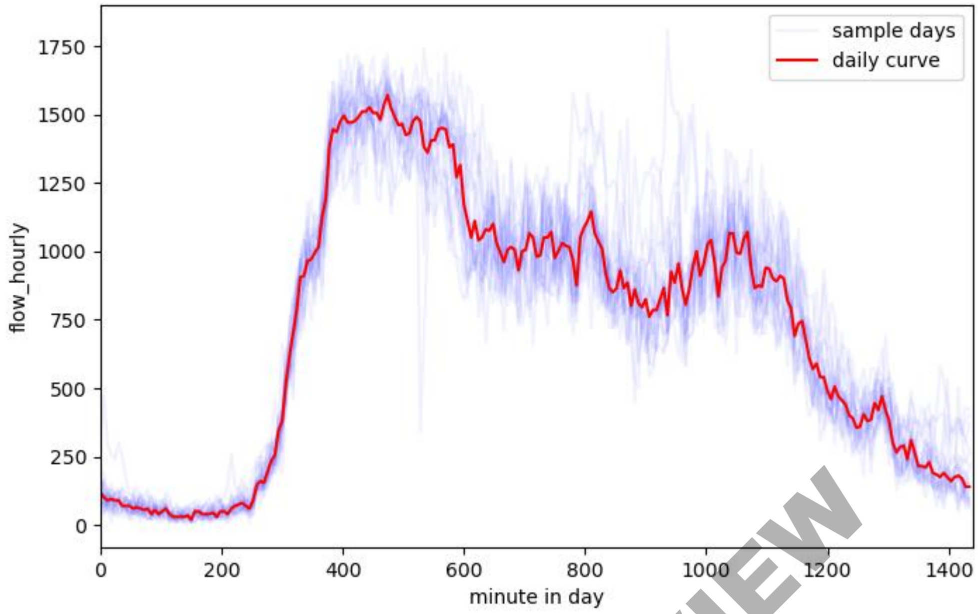

Figure 9: Daily curve of hourly vehicle flow for a typical weekday

This process of calculating the statistics is then repeated for each interval throughout the day, which essentially generates a sequence of average traffic flows per 5 minutes throughout a typical day. This process is further repeated for each combination of traffic variable-location-type of day. This results in a full set of daily curves for each traffic variable per location per type of day that represent typical (i.e. "expected") conditions for that location and time of day. The full set of calculated statistics for a typical daily curve is shown in Figure 10.

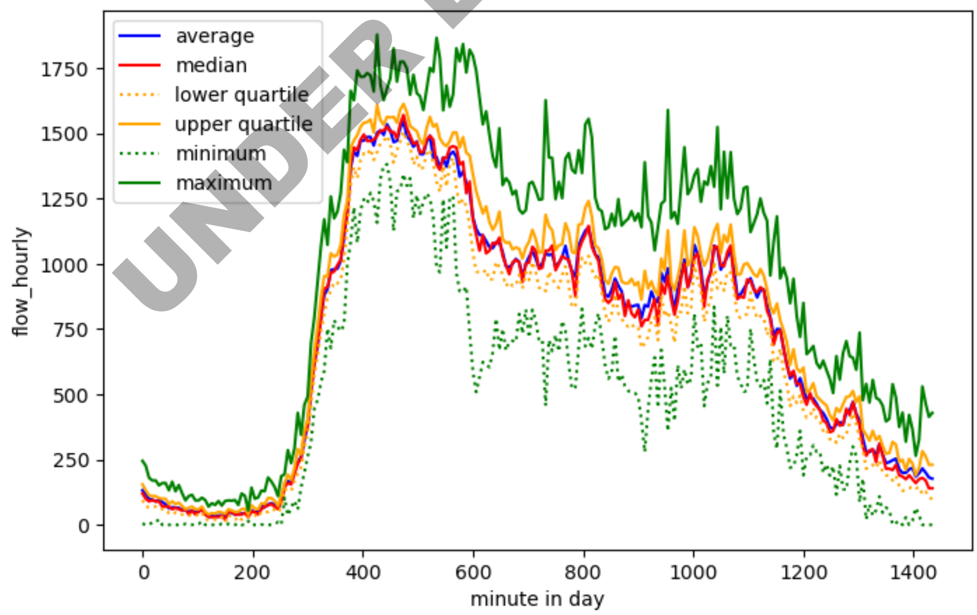

Figure 10: Daily curve statistics for a typical weekday

There are multiple benefits of the daily curves approach: daily curves are easy to generate as they rely on established statistical analysis processes; once generated they are easy to store and utilize (a daily curve model can be easily stored as a dictionary); they produce data of fixed size (each daily curve's size for a specific variable is only defined by the temporal resolution used); and finally

Funded by the European Union

Project funded by
Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
Swiss Confederation
Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

they capture a straightforward relationship between temporal variables and typical traffic conditions, providing a robust model for predicting traffic conditions based on minimal input that is always available(i.e. time of day).

The benefits of daily curves mentioned above make them an ideal fit for the purposes of the network evolution forecasting and anomaly detection tools at this stage: on the one hand their focus is on scenarios that exhibit typical traffic conditions, and daily curves are designed to capture typical daily traffic as averaged from multiple days. And on the other hand, these tools are still in development and testing, and the daily curves provide a highly robust model that can be used reliably during deployment, while still keeping the option available to swap the daily curves models with more detailed ones in the future.

2.2.2 NETWORK EVOLUTION FORECASTING SERVICE

2.2.2.1 CONCEPTUAL APPROACH

The Network Evolution Forecasting Service is a real-time tool designed to provide short-term forecasts of traffic network conditions based on the current state of network traffic. It is designed as a simplified process that requires only a real-time traffic feed as input, and afterwards is able to generate short-term traffic forecasts upon request.

Its primary goals are to be robust and applicable to multiple scenarios, with a secondary consideration to allow for expansion, fine-tuning, and adaptability without considerable overhead. As such its functionality is divided into two main parts: The forecast model, that can be trained and calibrated offline and contains the logic that generates traffic forecasts, and the user interface layer, that allows users and systems to retrieve forecasts. The user interface layer is available as a RestAPI, allowing for deployment both locally and over remote networks. The Network Evolution Forecasting Service has been developed in the python programming language (version 3.14.6), and the RestAPI has been developed using the FastAPI python library.

2.2.2.2 MODEL OVERVIEW

As stated above, the Network Evolution Forecasting Service is designed to support network management services under typical conditions, and its underlying model assumes that a stable relationship exists between input parameters and output traffic conditions under normal demand. Therefore we implement the daily curves approach described above as the underlying forecasting model using the typical daily curve as the baseline for our predictions. This setup provides a robust foundation, as it is able to provide a baseline estimate for any timestep using only a timestamp as input, therefore minimizing dependencies significantly.

In addition to the baseline estimate, we retrieve the latest real-world value from the real-time feed, which we compare to the baseline estimate from the daily curve for the corresponding time of day to calculate the ratio of the actual value to the baseline estimate. This ratio is then applied as a modifier to the baseline estimates for the following timesteps, under the assumption that traffic conditions are generally consistent: if real-world traffic volume is currently higher compared to the same time in the typical daily curve, we expect that the overall trend of the typical daily curve will continue, but at the higher detected volume, until a future reading from the real-world sensor indicates that volume has decreased to the same level as the typical daily curve. An example of this forecast process is shown in Figure 11.

Funded by the European Union

Project funded by
Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
Swiss Confederation
Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

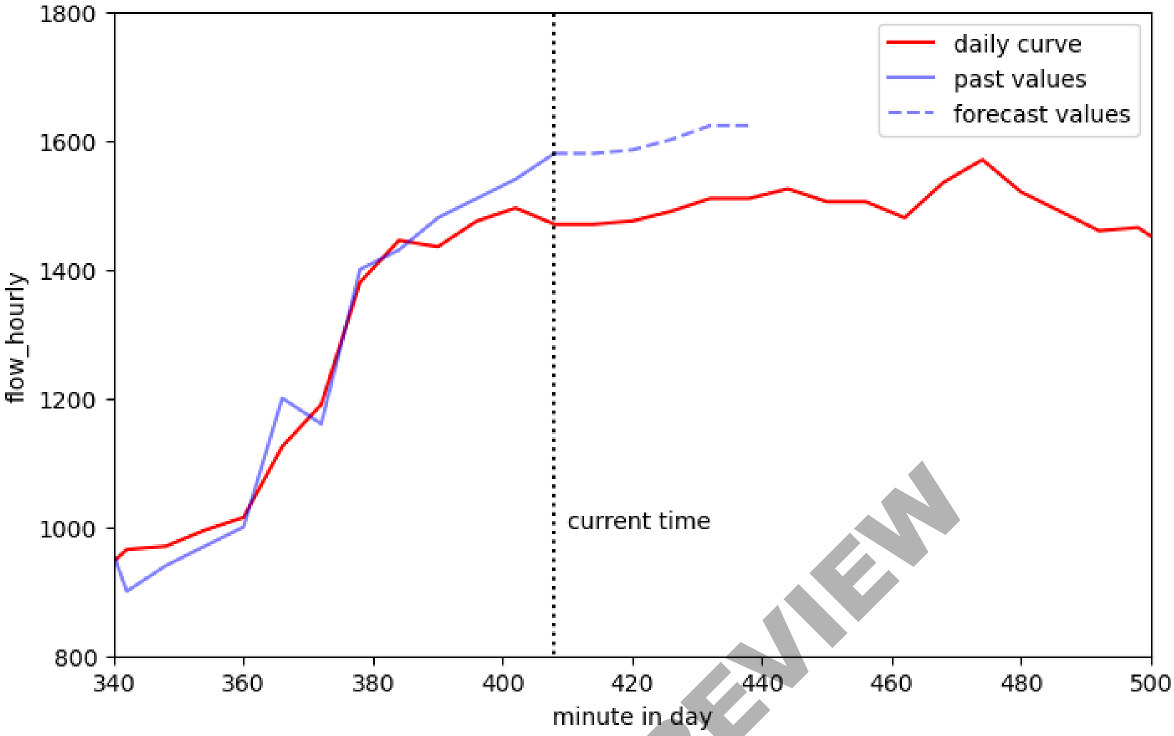

Figure 11: Real-world value compared to the daily curve, with future values projected based on the difference

2.2.2.3 MODEL CALIBRATION

In order to calibrate the daily curve model, given a dataset of traffic sensor data we calculate the average and median value for all available traffic variables during a 5-minute window.

The most recent values are used as the dataset to generate the daily curve, aiming to identify the smallest dataset that generates consistent results. There are two main reasons as to why this constraint on smaller datasets is implemented: First, from an operational perspective, using smaller datasets enables the system to be calibrated faster, both in terms of calculation time, but more importantly in terms of how much data would need to be collected and analysed before the system may come online. This is particularly relevant in cases where historic data is not available, and all data needs to be collected in real-time. And second, while we design the forecasting service for use in normal traffic conditions that are not expected to deviate much from the baseline, we nevertheless expect traffic conditions to exhibit gradual change, over time. As such by always using smaller and more recent datasets to generate the daily curves we ensure that any gradual changes are captured in the daily curve, rather than getting normalized if using a much larger dataset.

We performed the following test to identify the appropriate dataset size for the training set: Using available historic traffic sensor data, we selected a target range of dates to be used as the test dataset, containing 5 weekdays in February 2026 (February 2nd to February 6th). Next we generated daily curves using datasets of varying size, from 1 week up to 6 months, and compared the calculated daily curve values to the actual values from the test dataset using r-squared values. The error scores are shown in Figure 12.

Funded by the European Union

Project funded by
Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
Swiss Confederation
Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

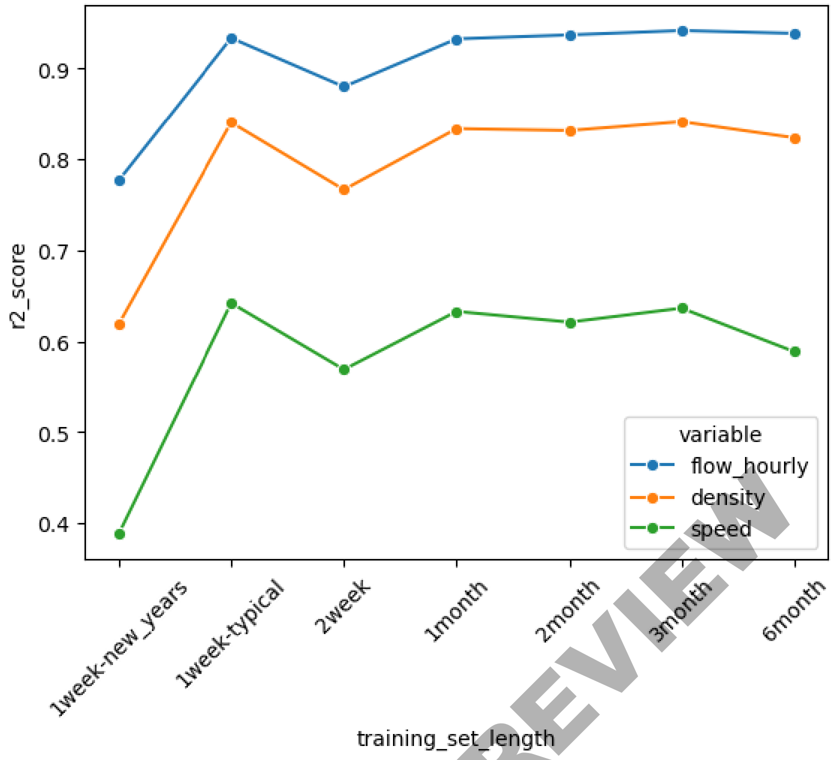

Figure 12: R-squared values for varying daily curve training dataset sizes

As expected, large datasets of 3-6 months provide consistently good results. As the datasets get smaller, while it is possible to generate good daily curves from small datasets (e.g. even using 1 typical week), they are inconsistent in their scores, as any anomalies will affect the output daily curves significantly (e.g. by using the week containing New Year's Day). As such we propose using approximately 20 days' worth of data (e.g. 1 month for weekday daily curves, 8-10 weeks for weekend daily curves), which is small enough to be versatile, but large enough to accommodate any anomalies within the training dataset. Finally, given that training and calibration may be performed automatically and therefore it is unclear whether outliers may have been included in the dataset, we calculate the daily curve using the median as the typical value rather than the average across all days, in order to mitigate any skewness from extreme outliers.

2.2.3 ANOMALY DETECTION SERVICE

2.2.3.1 SERVICE OVERVIEW

The Anomaly Detection Service is a tool for evaluating whether a given set of traffic conditions detected at a specific location should be considered as atypical or out of expected values for the given date and time. The main aim of this service by definition is to operate on recurrent conditions, i.e. under normal demand, in order to detect when current conditions deviate from normal.

The core system essentially operates using upper and lower threshold values for each given timestep. Provided with a traffic variable value and a timestamp, the tool will evaluate whether the given value is within the bounds defined for the day type and time window that corresponds to the provided timestamp. Given this core setup, any model that provides clearly defined upper and lower bounds can be plugged in and be used, allowing future flexibility to introduce more complex models with minimal interruption.

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation

Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

2.2.3.2 MODEL OVERVIEW

Given that any model can be used as long as it conforms to the core principles identified above (upper/lower bounds for each timestep), we demonstrate here an implementation of a straightforward statistical model using the daily curves approach we presented earlier.

Given a dataset of past traffic observations from traffic sensors, we calculate summary statistics (median, lower and upper quartile) for each of the temporal intervals in a daily curve. The lower and upper quartiles are then used as the lower and upper bounds of the anomaly detection system for the given time window. Results of the anomaly detection process applied to a full weekday's data are shown in Figure 13.

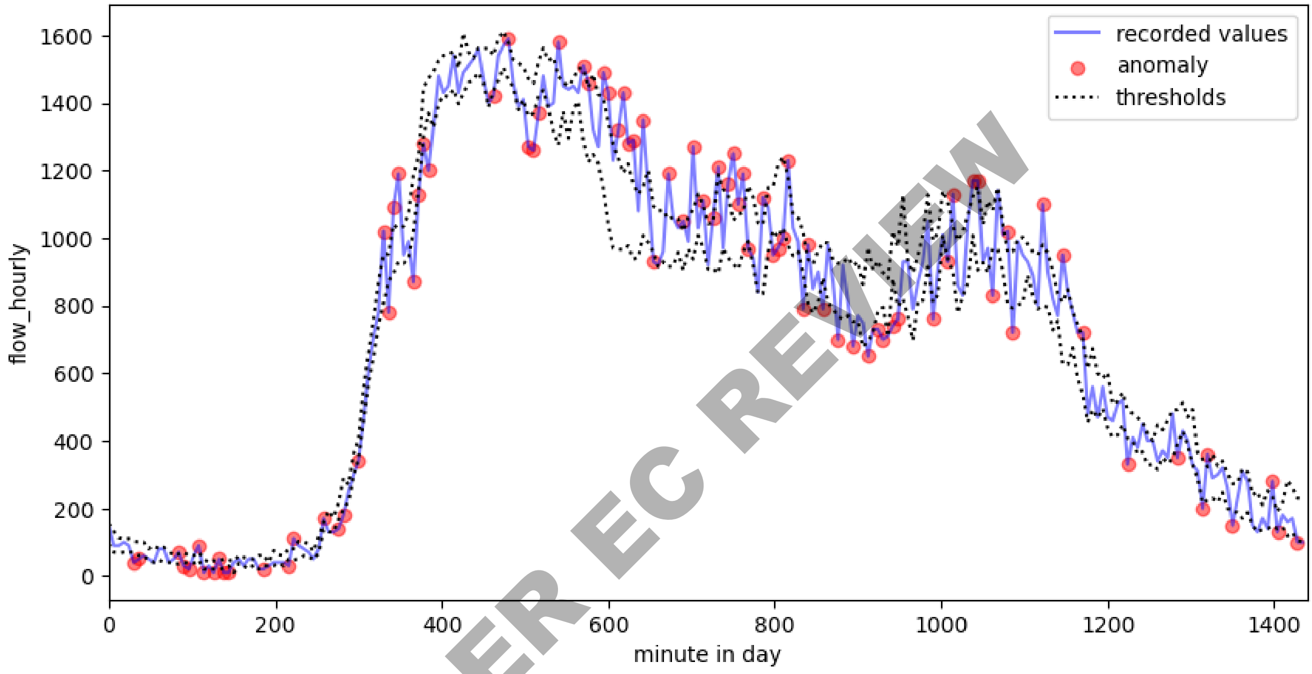

Figure 13: Daily curve with lower and upper quartiles as the bounds, and actual daily values with highlighted anomalies

2.2.3.3 MODEL CALIBRATION

The core goal of the model is to capture a typical day in the daily curve, to be used as the thresholds for detecting whether a particular reading falls outside the expected range of typical values for that time of day, i.e. an anomaly.

We use a dataset of the 17-20 most recent valid days in order to generate the daily curves for the anomaly detection. We categorize the days according to their type (weekday/weekend), as weekday (Monday through Friday) traffic patterns exhibit strong self-similarity, and the same stands for weekend traffic patterns (Saturday/Sunday). Therefore we use 4 weeks' worth of data to generate daily curves for weekdays, and 2 months' worth of data to generate weekend daily curves. This threshold of 20 past days has been found to be the smallest that produces consistent results with daily curves generated from larger datasets, as demonstrated previously. By always using data from the most recent days, we can accommodate any seasonality or gradual trends that may be part of the dataset (e.g. gradually increased demand due to warming weather in spring/summer, commuters gradually returning to daily patterns after summer vacation, etc).

2.2.4 LIMITATIONS AND NEXT STEPS

This report includes the initial versions and first steps in the development work regarding the Network Evolution Forecasting and Anomaly Detection services; as such we recognize that some of the elements described here can benefit from additional calibration

Funded by the European Union

Project funded by
Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
Swiss Confederation
Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

and fine tuning, which will be the focus of this work moving forward. The progress and updates on fine tuning and calibration that will be performed in the following months will be included in the follow-up deliverable D4.3. In this section we aim to address these limitations and highlight the path forward.

Regarding the Network Evolution Forecasting service, by only comparing the latest real-world data point to the baseline in order to predict the evolution of the traffic state, we have used a singular reading as the comparison point for the short-term forecasts. Single-point readings are susceptible to false positives, for example when traffic conditions deviate momentarily from the typical traffic state due to drive behaviour or instrument technical irregularities, but return to normal almost immediately. However as we use only a single data point to generate our forecast, this may lead to incorrect predictions. In the future we aim to introduce a rolling window for detection to include more than one sampling point, and use the average of those to compare to the baseline. In this way we expect that the impact of any momentary fluctuations in detected data points will be minimized, therefore increasing the predictive capabilities of the Network Evolution forecasting Service.

Regarding the Anomaly Detection service, we have established a method that requires an upper and lower bound of typical values for a given point in time during the day. We consider this core methodology to be sound and robust, however we have not performed extensive evaluations as to what the best values are to use for the lower and upper bounds. At this stage of the work we have used the first and third quartiles as the bounds, meaning that we expect 50% of the input data to be within expected values. In the future we will investigate the performance of the tool using additional statistical metrics (e.g. quantiles, standard deviations from the mean) to identify potentially better metrics. Additionally, we are currently implementing a Boolean result, i.e. a value is either within the expected range or outside it. In the future we aim to explore whether a graduated result might be more applicable (e.g. evaluating a value in terms of how far it falls outside the expected range).

2.3 COMPUTER VISION FOR EVENT PROCESSING

Task 4.4 emulates UAV video observations with simulation-derived observations aiming to provide structured vehicle detections and traffic indicators for the FEDORA Complex Event Processing Engine (CEPE). Developed for the Nicosia in-silico pilot, the obtained traffic indicators can be used to support usage-pattern analysis, bottleneck detection and candidate anomaly identification, while final incident classification and prediction remain within the downstream CEPE and GenAI pipeline.

During this reporting period, the operational concept, processing workflow, interfaces and altitude-aware emulation approach were specified. This enables the complete event-processing chain to be evaluated before full pilot deployment without assuming perfect visual sensing.

2.3.1 OBJECTIVES AND ROLE WITHIN TASK 4.4

UAV-computer vision (CV) transforms UAV observations into machine-readable traffic features at road-segment or Region-of-Interest (ROI) level. Its essential functions are to:

- produce time-stamped vehicle detections and confidence information from emulated SUMO observations;
- optionally associate detections over time to derive tracks, movement information and speed proxies;
- aggregate detections and tracks into counts, flow, speed, density or occupancy proxies, and candidate cues such as sudden speed drops, stopped-vehicle clusters or abnormal queue growth;
- emulate altitude-dependent detection quality, and export standardised observations and features to FRONT's CEPE.

This separation preserves a clear architectural boundary: UAV-CV supplies observations, intermediate features and candidate cues; the T4.4 CEPE fuses them with other network information and produces final incident or bottleneck labels and predictions.

2.3.2 OPERATIONAL CONCEPT AND ARCHITECTURE

UAV-CV is planned as an on-demand software module, for example in Python, within the Nicosia city in-silico stack. For each scenario it receives simulation objects or trajectories together with timestamps, UAV altitude and optional camera or pose context. Instead

Funded by the European Union

Project funded by
Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
Swiss Confederation
Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

of exposing perfect SUMO ground truth, the emulation layer produces drone-like observations with missed detections, variable confidence and altitude-dependent processing performance.

The processing chain comprises input and time synchronisation, altitude-aware detection, optional tracking and feature extraction, and an output adapter. It runs once per scenario and exports a self-contained results package. The same architecture can later support small-scale validation with real UAV video when the required telemetry and camera interfaces are available.

The current scope covers road vehicles and configurable ROIs or road segments. Optional field-of-view, UAV pose and camera-error information can be added as the interface matures. The final choice between live TraCI exchange and file-based CSV or JSON input remains an integration decision.

Figure 14 below summarises the component workflow and its boundary with the T4.4 CEPE.

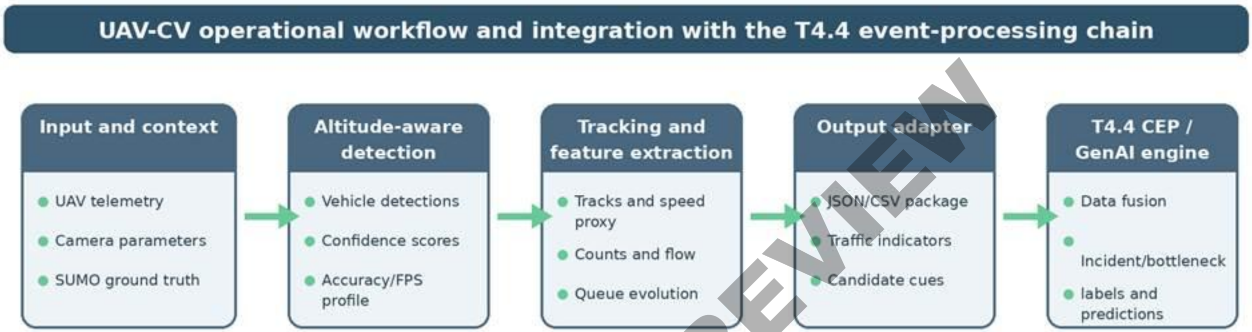

Figure 14: Conceptual workflow of the UAV-CV component and its interface with the T4.4 Complex Event Processing Engine.

At minimum, synchronisation uses timestamp and altitude; optional GPS or pose information relates the camera view to an ROI or road segment. Detection produces confidence-tagged vehicle observations, tracking links them over time, and aggregation creates temporal traffic features.

The output adapter then packages detections, tracks, aggregated indicators, candidate anomaly cues and performance metadata in JSON or CSV for the other FEDORA services.

2.3.3 ALTITUDE-AWARE DETECTION AND PERFORMANCE EMULATION

UAV altitude is treated explicitly because increasing altitude reduces the pixels available per vehicle and can affect both detection accuracy and processing speed. The component therefore does not assume fixed perception performance across a flight profile.

The emulation profile follows the patterns reported by Makrigiorgis, Kyrkou and Kolios (2023) for aerial vehicle detection between 50 m and 500 m, representing detection quality through mAP or precision proxies and processing efficiency through Frames Per Second (FPS). These curves are performance proxies, not Nicosia-specific measured results.

The default is a mixed-altitude, single-model profile, with optional altitude-specific profiles and configurable inference input size for comparative tests. Each run records the selected profile and expected quality-versus-altitude values in its performance metadata.

In simulation, the profile adjusts detection probability, confidence and processing throughput, producing a more realistic observation stream than direct SUMO ground truth. It remains replaceable so that later datasets or Nicosia pilot measurements can recalibrate it.

Funded by the European Union

Project funded by
Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
Swiss Confederation
Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

2.3.4 TRAFFIC INDICATORS AND CANDIDATE EVENT CUES

The main value of UAV-CV is the conversion of individual detections into traffic indicators that evolve over time. Aggregation is performed per ROI or road segment using configurable temporal windows, allowing short-lived fluctuations to be distinguished from persistent deterioration.

Table 2: Traffic information products generated by UAV-CV

INFORMATION PRODUCT	DESCRIPTION	USE IN TASK 4.4
Vehicle detections	Time-stamped vehicle observations with confidence score and ROI or segment association.	Provides the basic visual evidence used for subsequent aggregation and traceability.
Tracks and speed proxies	Optional track identifiers and position evolution from which relative or calibrated speed estimates can be derived.	Supports detection of slow-moving or stopped vehicles and traffic-state changes.
Counts and flow	Number of vehicles observed in an ROI or crossing a virtual line during a defined interval.	Describes usage patterns and temporal demand variation.
Density / occupancy proxies	Vehicle accumulation or occupied-area measures within an ROI.	Identifies sustained build-up and helps distinguish free-flow from congested conditions.
Queue indicators	Estimated queue length, duration and growth or dissipation over time.	Provides direct evidence for bottleneck formation and recovery.
Candidate anomaly cues	Time-stamped flags for sudden speed drop, stopped-vehicle cluster or abnormal queue growth.	Triggers further evaluation by the CEP; it is not a final incident classification.

Queue evolution is particularly important for bottleneck detection: sustained queue growth combined with decreasing speed is stronger evidence than an isolated count. Persistent or unexpected stopped-vehicle clusters can similarly form candidate anomaly cues. Thresholds and persistence rules will be calibrated during scenario testing and CEPE integration.

2.3.5 DATA INTERFACES AND CEPE INTEGRATION

The component uses lightweight JSON and CSV formats for repeatable scenario execution and transparent inspection, while TraCI remains an option for tighter SUMO coupling. The final exchange mechanism will be agreed with the Nicosia simulation and T4.4 CEPE partners.

Table 3: UAV-CV input data

INPUT	MINIMUM CONTENT	FORMAT	STATUS / ROLE
UAV telemetry	Timestamp and altitude in metres; optional GPS and pose.	CSV or JSON	Required for altitude-aware emulation; positioning fields are optional.
Camera parameters	Field of view and, where available, camera or geolocation error estimates.	JSON	Optional contextual input; detailed schema remains to be finalised.
In-silico traffic data	Ground-truth vehicle objects or trajectories and the related network / segment identifiers.	TraCI stream, CSV or JSON	Required in emulation mode; final exchange mechanism is TBD.

Funded by the European Union

Project funded by
Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
Swiss Confederation
Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

Scenario configuration	Run identifier, ROI / segment definitions, temporal aggregation settings and selected performance profile.	JSON	Required to make results reproducible and traceable.

Table 4: UAV-CV output data

OUTPUT	PRINCIPAL FIELDS	FORMAT	PURPOSE
Detection stream	Timestamp, detected vehicle, confidence, ROI or segment reference.	JSON	Detailed observation stream for traceability and optional downstream reprocessing.
Tracking stream	Track ID, time, position or ROI, and speed proxy.	JSON	Optional temporal association of detected vehicles.
Aggregated indicators	Counts, flow, speed proxy, density / occupancy proxy and queue-growth measures.	JSON or CSV	Primary traffic-feature input to the T4.4 CEP.
Candidate anomaly cues	Timestamp, cue type, affected ROI / segment, value, threshold and confidence.	JSON event list	Highlights conditions requiring downstream evaluation.
Performance metadata	Altitude profile, expected quality metrics, processing-speed profile, configuration and warnings.	JSON or CSV	Documents the reliability and reproducibility of each run.

Each output package includes a run identifier, common timestamp convention, data-source mode, ROI or segment identifier, configuration and altitude-performance profile, enabling reliable alignment with other network observations.

Integration logic

The output adapter forms the boundary with FRONT's CEP and GenAI services. UAV-CV exports an event and feature stream, which the CEP combines with other network information to generate the higher-level incident, bottleneck and network-evolution outputs required by the Dynamic Response Planner and other FEDORA services.

Integration follows an evidence hierarchy: detection → track → aggregated indicator → candidate cue → final CEP label or prediction. This preserves explainability because higher-level events remain traceable to the underlying UAV observations.

Confidence, altitude and profile metadata allow the CEP to weight evidence according to sensing reliability and support later analysis of false alarms, missed detections and sensitivity to flight conditions.

2.3.6 OPERATIONAL REQUIREMENTS AND CURRENT STATUS

UAV-CV is an on-demand experimental and pilot-support component. It runs during simulation or validation sessions, executes once per scenario and exports a self-contained results package; continuous operation is not required.

Each run logs its inputs, configuration, warnings and errors, together with the processed interval, selected performance profile, input volume, detections and tracks produced, and output location. This supports troubleshooting and reproducible scenario comparison.

Altitude profiles, detector settings, ROI definitions and feature parameters remain configurable without changing the integration workflow. Access is restricted to FEDORA partners and authorised pilot operators.

Funded by the European Union

Project funded by
Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
Swiss Confederation
Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

2.3.7 ASSUMPTIONS AND NEXT STEPS

The component specification and integration concept are established, including the required information products, architectural boundary, altitude-aware emulation principle and main interfaces. Quantitative Nicosia-specific detection results are not yet available; published altitude-dependent behaviour is currently used as the in-silico performance proxy.

The current design relies on two principal assumptions:

- the reference study's altitude-dependent patterns are sufficiently representative to emulate UAV vehicle detection in Nicosia; this medium-severity assumption requires validation or recalibration; and
- UAV-CV supplies observations, features and candidate cues, while the T4.4 CEPE and GenAI pipeline performs final classification; this low-severity architectural assumption requires an agreed data contract.

Remaining integration decisions concern the simulation interface, camera-context schema, ROI and segment mapping, and the thresholds and persistence rules for candidate cues.

The next development and validation steps are to:

- finalise the JSON data contract, identifiers and timestamp conventions with FRONT;
- implement and test the SUMO/TraCI or file-based adapter and encode the mixed-altitude performance profile;
- implement tracking, aggregation and queue-evolution functions for configurable ROIs and road segments;
- test free-flow, recurrent congestion, sudden speed-drop and abnormal queue-growth scenarios, and compare candidate cues with CEP outputs; and
- use available real UAV traffic video for small-scale validation and, where possible, recalibrate performance profiles and thresholds.

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

3 SYNCHROMODAL OPTIMIZATION ALGORITHMS

3.1 SUMMARY

The Synchromodal Conflict-Free Vehicle Routing Optimiser (SCF-VRO) will provide new logistics routing and scheduling solutions for freight flows, including congestion-aware search and routing techniques to facilitate response to detected bottlenecks and delays.

To achieve FEDORA's goals it is necessary to enhance the capability of current state-of-the-art solutions with new features and components. For example, to enable conflict-avoidance functionality in a synchromodal network it is necessary to integrate powerful dynamic multimodal routing tools with new capability to manage a wider range of freight types. The approach used in FEDORA builds on successful solutions and innovations from previous Horizon projects for unitised cargo and extends them with new algorithm integration layers to model other cargo types.

Similarly, last-mile delivery operations are increasingly more diverse due to the wider range of vehicles and modes of transport available which have different routing constraints and opportunities. In parallel, new data sources are becoming available. For example the innovative application of social optimum models for signalized intersections considered by AIT, ITS-VIE, VIE and ETH in Task T4.1 of FEDORA provides a new opportunity for increased accuracy of routing times and flows at junctions controlled by traffic lights, potentially by road user category (e.g. cargo bike, van, HGV). Open data initiatives such as OpenStreetMap offer an increasingly rich source of valuable routing information which can be used for more sophisticated optimisation of end-to-end multimodal freight operations.

To further exploit and benefit from these opportunities a multi-layered interoperable optimization architecture has been used to address larger-scale industry synchromodal and vehicle routing problems, including new algorithms with higher endurance capabilities with regards to the disruptions for freight movements. The key benefits of this approach are:

- Capability to plan, manage and optimise a wider range of more complex multimodal freight operations.
- Ability to make use of new sources of open data or real-time information in the logistics planning algorithms.
- Increased flexibility to model necessary operational rules and business constraints to achieve workable solutions.

This section summarises the progress so far in Task T4.5 of FEDORA. It should be noted that, as per project plan, at the time of preparing this document Task T4.5 is in progress and will continue for another 14-15 months. Therefore the solutions described in the following pages, although already offering a significant range of new functionality, will be refined and extended as the project progresses.

3.2 CONFLICT AVOIDANCE AND ROUTE OPTIMISATION FOR SYNCHROMODAL NETWORKS

New logistics concepts such as the Physical Internet and synchromodality have been the subject of significant research and innovation activity in Horizon funded projects and many other initiatives. From a routing and scheduling perspective these concepts lead to closer integration between modes. The combination of increased availability of data and a wider range of transport modes, especially for last mile operations, offers opportunities for higher levels of efficiency and resilience in complex freight transport networks. Previous innovation projects have successfully developed powerful multimodal logistics optimisation tools, so the emphasis in FEDORA is on the development of new capability, deployed to support interoperability with existing tools and solutions that significantly enhance the capacity for freight and transport operators to increase efficiency and take further steps towards realising the benefits and operation of future synchromodal options.

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation

Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

The SCF-VRO includes freight flow synchromodal models and last-mile routing optimisation, encouraging modal shift to more sustainable and greener transport modes. Figure 15 below illustrates the synchromodal operations that are considered, including the multimodal hubs (e.g. the Port of Bilbao) forming part of a potential multimodal network in this case comprising ship, rail and road connections.

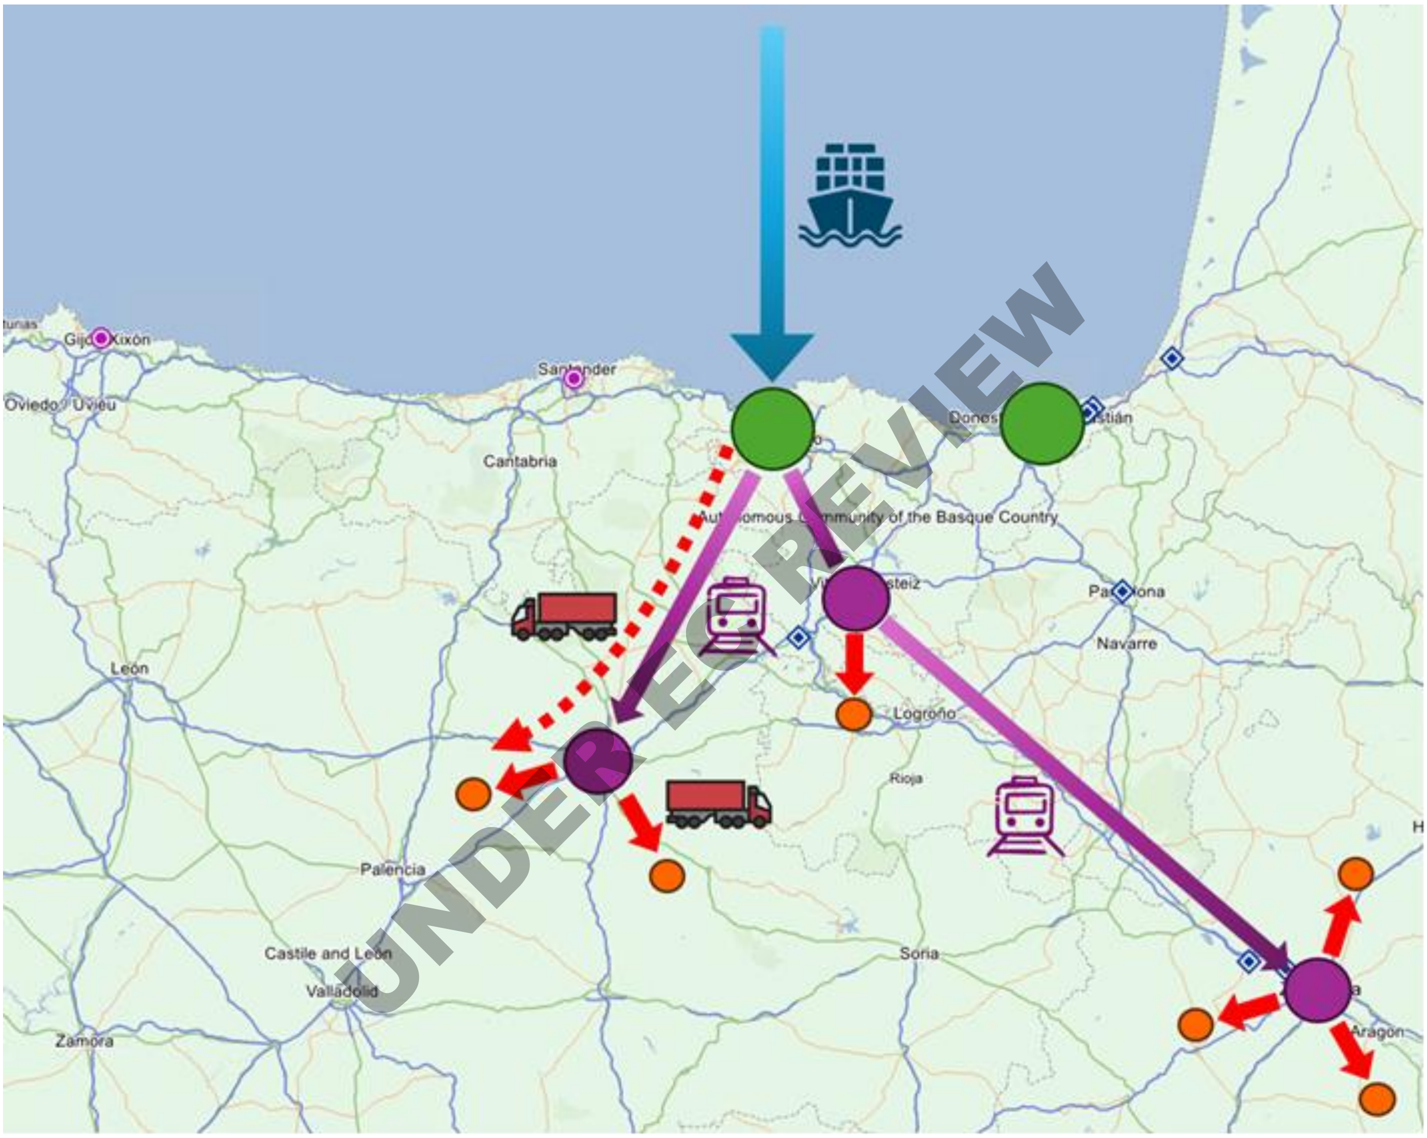

Figure 15: Example Synchromodal Network

Multimodal operations of this nature present a number of challenges from an operational planning and management perspective. Uncertainty, disruptions and congestion can have a major impact. For example, the following list illustrates just some of the possible disruptions that have been identified during early discussions with partners about logistics operations in the Basque region:

- Weather: rain, strong wind or high humidity delay unloading operations in the port, triggering rescheduling of the landside logistics activities.
- Queues in the city or at the entrance to the port delay trucks.
- Unavailability of specialised equipment in the port at the required time.

Funded by the European Union

Project funded by

- Tidal windows constraining vessel berthing and departure times, with cascade effects on downstream operations.
- Vessel queue at the port: multiple ships waiting for the same berth. These vessel arrival delays result in the landside operations being disrupted leading to last-minute replanning.
- Tight train schedules: if the cargo is not ready at the exact departure time, the slot is lost and an alternative transport solution must be found, usually under time constraints.

One of the main aims of the SCF-VRO is to provide a toolset that makes the management of these disruptions easier and more efficient. For example, this might include reacting to a vessel delay, dealing with a train which is insufficiently loaded to justify running, weather disruption delaying loading operations or customs clearance delays leading to a rescheduling of truck operations. These examples should be considered as illustrative at this stage, but the intention is to identify relevant scenarios in the WP6 activities to demonstrate and test the SCF-VRO response.

The main inputs to the SCF-VRO include:

- **Road network data:** OpenStreetMap data has been used for initial tests, and has the benefit of being free to use and highly detailed. However the system is designed to be compatible with most common map data standards such as MIF/MID.
- **Road accessibility and expected speeds**: by vehicle category – see below for details.
- **Junction flow/delay data**: by vehicle category – see below for details.
- **Vessel data:** this includes service schedules, loading/unloading deadlines, updated ETAs and course details if appropriate.
- **Rail transport data:** including service schedules, updated ETD/ETA, capacity, route details, loading/unloading deadlines.
- **Road transport resources:** location and availability of trucks and drivers.
- **Freight orders/volumes**:
    - Cargo type and classification – see below for details
    - Unique cargo reference
    - Size/volume
    - Origin warehouse/location
    - Destination warehouse/location
    - Client identifier
    - Time constraints
    - Currently allocated vessel service
    - Currently allocated rail service
    - Currently allocated truck/haulier
- **Network node details**: hub, terminal, depot, transport centre and customer geocoding and characteristics.

The purpose of the SCF-VRO is to automatically allocate freight movements to multimodal transport resources/capacity and to optimise the routing and use of resources. This is achieved using automated algorithms integrated with an intuitive and easy-to-use GUI. The SCF-VRO can provide synchromodal optimisation capability by reacting to real-time events (e.g. delayed ship arrival), automatically re-routing affected freight. Outputs of the SCF-VRO include:

- The allocation of freight orders/movements to resources and multimodal services.
- Sequencing and scheduling of freight movements in the network, including last-mile optimisation.
- Accurate routing and reaction to congestion – see below for details.
- ETA/ETD at main hubs and origin/destination locations.

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

The SCF-VRO builds on solutions from previous R&I projects, and includes important extensions and new capability to address a wider range of freight operations and scenarios. The new model is designed using an extensible approach that allows additional rules and constraints relating to specific freight types to be modelled. In FEDORA it is proposed to consider steel coils and containerised cargo as example cases of this approach.

For example the steel coils in question are typically very heavy, so that normally only one or two can be carried at any one time on an articulated truck. Trains can transport many times this number, but it is still necessary to respect weight and capacity constraints. In many cases the coils are sensitive to moisture, so weather affects and delays handling and transport.

Rules for repositioning of empty equipment may vary also. In the case of containers, the empty box must be returned to a suitable depot or terminal, often determined by the shipping line, while for trailers carrying coils there is no equivalent box to reposition, but of course the driver and truck must return to their transport centre or positioned at their next job.

The SCF-VRO provides operational planning and scheduling capability for industry organisations such as intermodal operators, freight forwarders and hauliers. It also enables strategic analysis of multimodal networks and traffic management in mobility hubs (e.g. rail nodes or freight hubs). Example use cases could include assessing the benefits of a new rail link/service or analysing the impact of adoption of synchromodal activities on freight traffic flows.

The FEDORA solution provides new tools to reduce delays for trucks moving between terminals, combining innovative multi-layered queue-reduction algorithms with conflict avoidance routing and configurable prioritisation to increase the probability of achieving target e2e transit times.

This is achieved by adopting an interoperable flexible modular algorithm design, enabling different aspects of the e2e logistics operation to be modelled and optimised depending on operational requirements. For example, vessel schedules, ETA/ETD and terminal/customs processing times are contemplated. The delayed arrival of a vessel triggers a replanning process, automatically reallocating freight to other transport resources, and attempting to find productive work for resources which were originally assigned to the delayed consignments. Future synchromodal operations may allow redirection of smaller vessels between ports, and more dynamic allocation of freight to trains. The SCF-VRO provides tools to model these scenarios and assess impacts.

To accommodate more diverse cargo types a structured approach to consignment units is used, taking steps towards the Physical Internet concept. Cargo Units may consist of a conventional ISO container or load units comprising groups of coils or other items, which may be driven by customer and operational requirements at the ultimate destination and therefore need to be respected during the transport process. Each Cargo Unit has a cargo type/category and a unique identifier for traceability.

Available rail services are defined by parameters such as their schedules and terminals visited, cost, capacity, loading deadlines and unloading durations. Haulier and driver capacity are defined in terms of the location and availability of each unit (start base, shift length, vehicle type and scheduled start time). If relevant, vehicles can be flagged as being electric (e.g. those used for relatively short last-mile movements).

The freight flow synchromodal models, contemplating last-mile congestion-aware routing, encourage a modal shift to more sustainable modes, and make it easier to use greener transport solutions such as electric vehicles and cargo bikes. FEDORA will also consider detailed last-mile routing in cities (such as the example below) and other potentially congested areas such as those in the vicinity of large transport hubs.

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

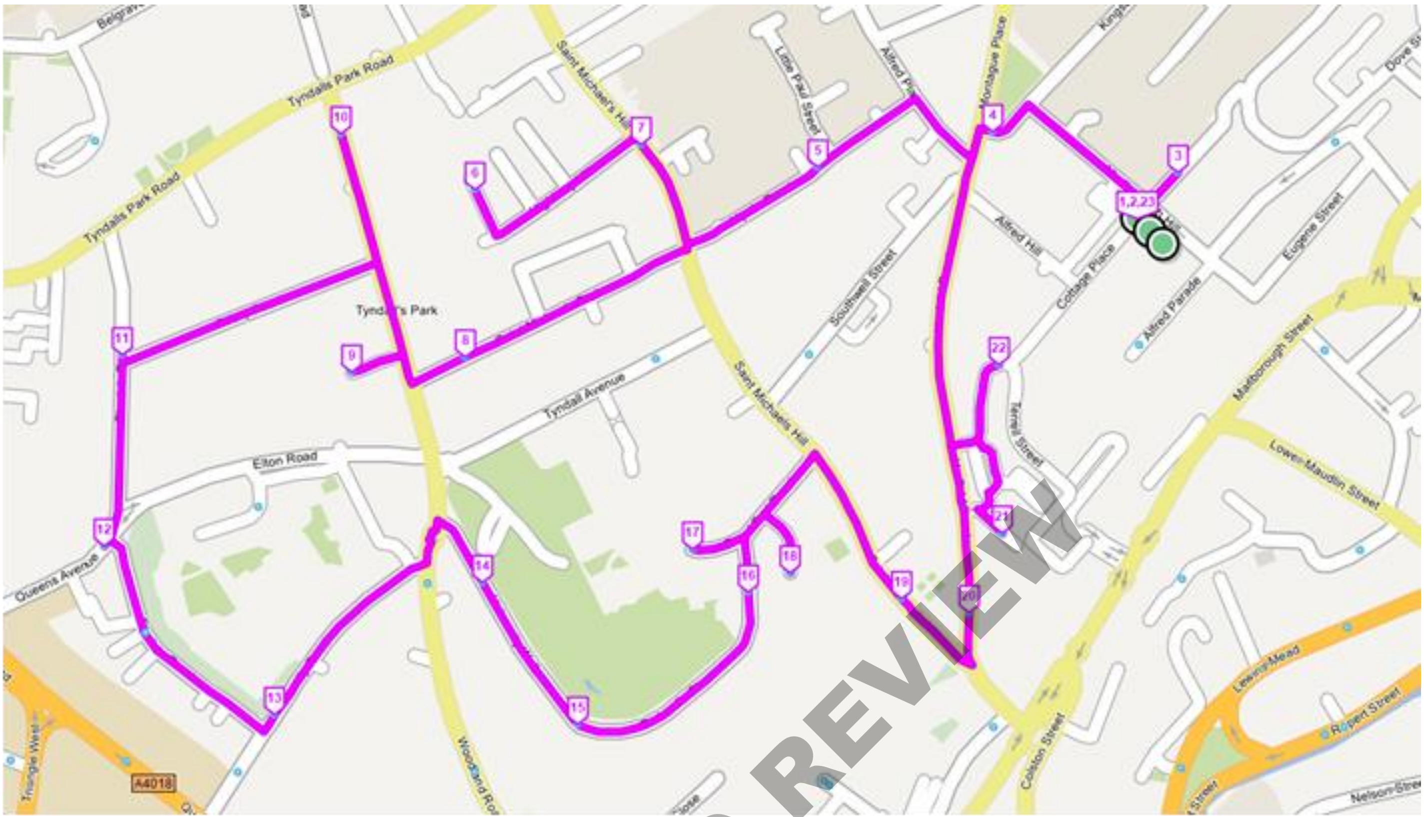

Figure 16: Last-mile Delivery Operations

Increasingly, last-mile routing needs to take into account the possibility of micro-hubs in city centres served by cargo bikes. Cities operate Zero Emission Zones (ZEZ), Ultra Low Emission Zones (ULEZ) or Low Emission Zones (LEZ), and have HGV restrictions and other constraints.

Solutions are available for the Vehicle Routing Problem (VRP) and its many variants for operations such as cargo bikes, electric vans and larger vehicles. One of the aims of the FEDORA development is to bring these closer together and increase the integration with intermodal operations. The vision is to provide a set of interoperable optimisation capability that spans the e2e logistics chain, from first-mile collection at the origin, though multiple modes for mid/long-distance transport (e.g. vessel, rail, linehaul), and last-mile delivery. In practice it is unlikely that every link in this chain would be managed and controlled by a single operator, but the hand-over between operators along the chain happens at different stages depending on each case/commodity type/territory, and this situation is evolving due to the acquisition and consolidation of logistics and shipping companies. Therefore the aim of the FEDORA activity is to increase the interoperability and compatibility between the optimisation of these diverse stages in the logistics chain allowing (for example) intermodal and last-mile HGV operations to be planned in a synchronised way, or large trucks and smaller delivery vehicles to be modelled in the same dataset, applying appropriate constraints to each.

Routing rules and constraints are a key aspect of this capability. Open data sources such as OpenStreetMap provide increasing levels of detail about routes and roads. For example, the image below shows a Clean Air Zone with cycle ways indicated in red, pedestrianised roads in green and footpaths in blue. (The other roads are: primary roads in green; main roads in orange; smaller roads in yellow and streets with vehicular access in white.)

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation

Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

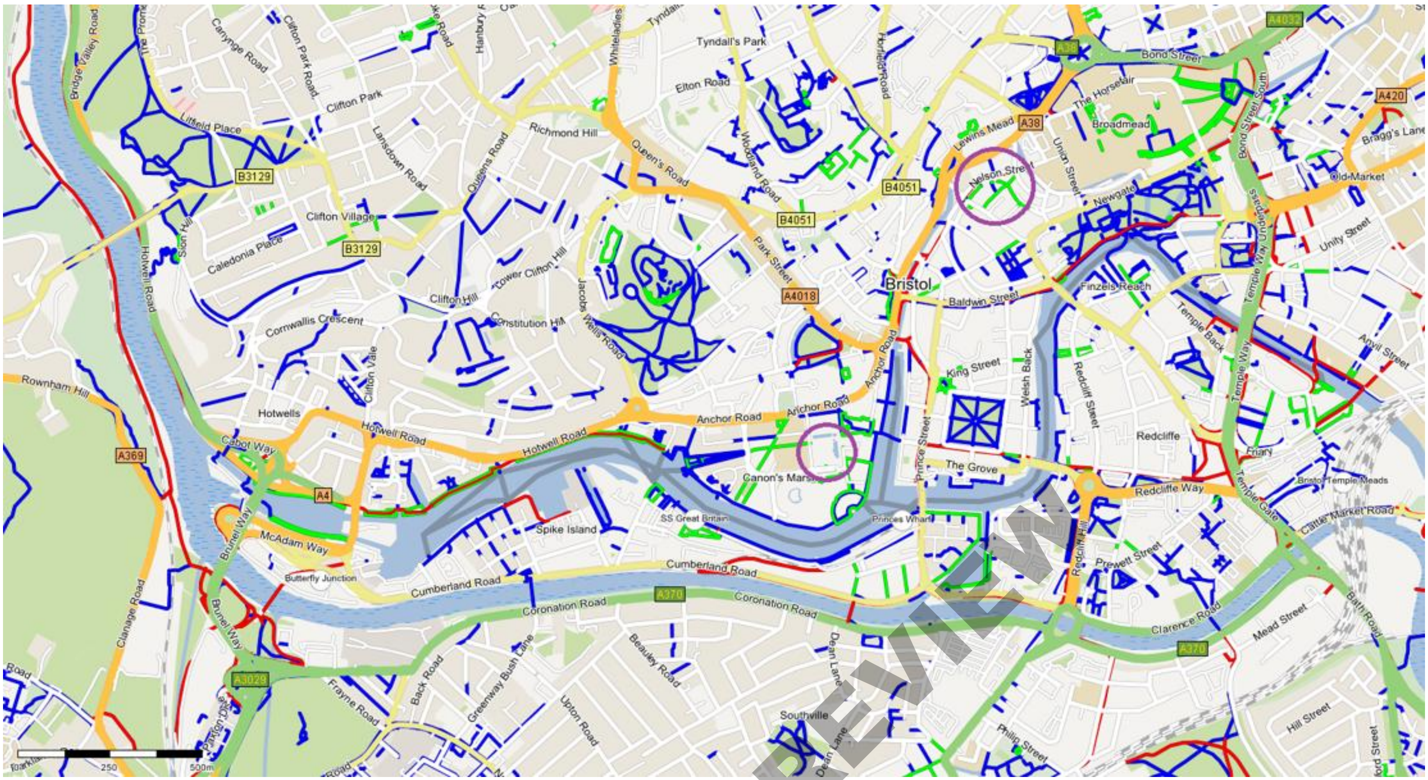

Figure 17: Vehicle-dependent Routing Data

This increased availability of publicly accessible data is welcome, but from a technical perspective it presents some new requirements, opportunities and challenges.

More sophisticated traffic lights turn restriction/delay model. The parallel work undertaken in FEDORA by AIT, ITS-VIE, ETH and VIE relating to social optimum models relating to signalized intersections and signalized corridors is anticipated to realise a more balanced network availability. This innovative work is described in more detail in Deliverable D4.1 of FEDORA, but an interesting outcome from the logistics/freight perspective is the possibility of different flows or transit times at junctions controlled by traffic lights, potentially varying by vehicle type. For example the transit time for cargo bikes for a particular movement across a junction might differ to that for HGVs. The balance and optimisation of these traffic flows is the subject of the work undertaken in Task T4.1, so is not considered in detail here, but the SCF-VRO routing toolsets have been extended to model this variation, and to optimise routes accordingly. It is anticipated that, in the future, information about the optimised traffic light schedules and/or expected flows/delays at junctions could be made available to the route optimisation tools.

This facilitates a conflict-free approach using intelligent risk-aware search and routing techniques to reduce the chance that diverting flows from one congestion area simply causes a second bottleneck in the adjacent road/junction. It also ensures that routing times for smaller, more agile vehicles such as cargo bikes are more accurately estimated. Studies of cargo bike operations suggest that, in densely populated city centres, bikes are in fact faster than vans because they are less affected by congestion. This may well apply to the transit times at signalized junctions also.

A flexible data model is used, allowing complex junctions to be modelled with different times for each inbound and outbound connection, and allowing times per vehicle type to be specified. Results from initial tests are shown in the following images. If the expected transit time is small the optimised route still includes the junction but the total time taken increases. If the delay is large enough the vehicle will avoid that turn altogether. In the examples below some traffic light controlled junctions have been selected (using OpenStreetMap data) and a delivery route through the city has been studied. Varying the expected transit time from zero (no

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

significant delay e.g. off-peak) to 90 seconds results in different routes taken, as expected. Please note that for clarity some detail about road types and other features has been hidden in the images below.

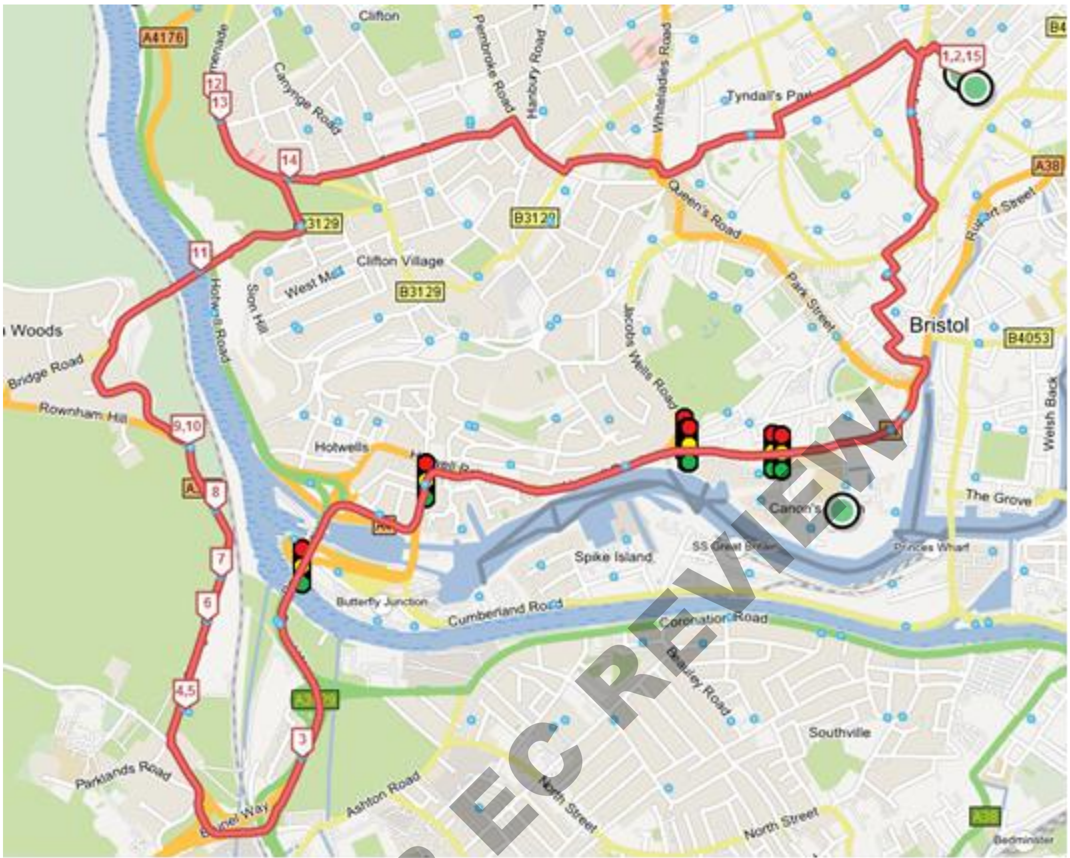

Figure 18: Off-peak Case – No Delays at Junctions

Funded by the European Union

Project funded by

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

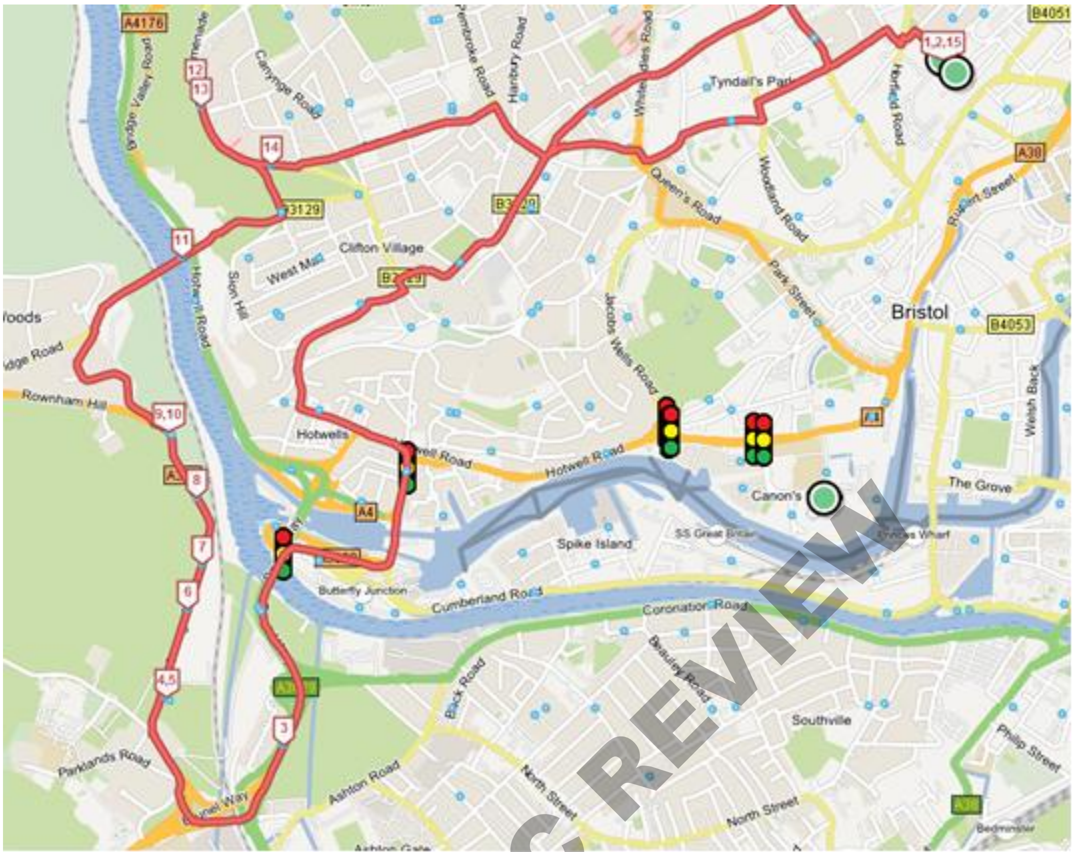

Figure 19: Moderate Delay at Some Junctions

Funded by the European Union

Project funded by
Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
Swiss Confederation
Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

Figure 20: Significant Delay at Some Junctions

Integration aspects. The SCF-VRO is designed to import data using common map data standards such as MIF/MID. OpenStreetMap has been used for initial tests as it is free and offers a number of integration tools and APIs. The SCF-VRO database and algorithms have been adapted to manage new road classifications such as cycle lanes, pedestrianised roads and footpaths. Configurable mapping layers have been defined to allow display and routing rules to be associated with these road types where appropriate.

Data validation and reporting tools have also be enhanced to detect more complex routing errors in the underlying data. For example the user might accidentally define a delivery location in an inaccessible location, or there might be minor connectivity errors in the underlying road network data which only become apparent when used by the routing algorithms. These tools help the planner ensure good, feasible solutions.

Visualisation. The increased level of detail available in the mapping data could lead to map screens becoming cluttered and too complicated to view. Layers for drawing the map in the GUI can be specified in using configuration file, so that additional layers such as cycle ways can be drawn if relevant to the user. Configurable colour-coding indicates the road type as illustrated in Figure 17 above. Similarly the route display algorithms have been adapted to show the planned route based on the vehicle type. Colour coding options are configurable and can be adjusted to suit user preferences and requirements.

Optimised traffic light junctions can be displayed on the map, with icons to indicate the type of junction and signal control system. Free text information can be included and displayed on the map in a tooltip. The design of the visualisation tools allows for further additions or adaptions e.g. if new details emerge from the T4.1 work.

Other routing considerations. Routing algorithms have been extended to allow different vehicle types to be routed differently. In particular this means different routing types use different speeds and are allowed to use different road types. This allows runtimes between sites to be calculated and road routes to be drawn correctly for different routing types. As described above, turn restrictions and junction transit times by vehicle type are also considered. Standard functionality such as ad hoc blocking/restricting roads due to known closures or other problems has also been extended accordingly.

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation

Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

For many applications the user may wish to maintain a database of agreed/average runtimes between locations. Often these are obtained by averaging actual runtimes over an extended period. This typically applies to inter-terminal truck routing or similar cases where the set of locations is relatively static and the journey frequency is high enough to obtain a meaningful average. This runtime override functionality has been extended to allow different overrides for different routing types.

Memory Management and Performance. Multiple route maps for multiple vehicle categories, further extended by complex junction flow data and other information, places a significantly increased burden on system resources, especially memory and processing capability. For large operations with many vehicles, consignments, intermodal connections and other variables, the memory usage can quickly become unmanageable if not carefully designed. Similarly, route calculations through a complex multimodal network, while relatively quick for each one, can occupy a lot of processor capacity for large numbers of journeys. A simple (and relatively small) example illustrates this point. If a delivery operation has to visit 100 different locations, then to optimise the routes it is necessary to know the runtimes and distances between all pairs of locations, i.e. there in principle 100x100=10,000 calculations to do. Usually it is not possible to simply assume that the A->B time is the same as the B->A time, especially in a city with traffic flow restrictions, but even if this is acceptable it only reduces the number by half. For 1000 locations the number is 1 million. Clearly there are strategies that can be adopted to mitigate this problem to some extent, but the underlying pressure is approximately quadratic in the number of locations.

There is increasing concern about the global energy use and availability of computing resources, and in any case such hardware can be expensive and may be out of the reach of smaller operators. Therefore significant attention has been given to reducing computer resource utilisation while ensuring good performance. Caching, intelligent initialisation and incremental recalculations are used to reduce unnecessary processor activity and memory usage.

Control over optimisation process. An additional complexity associated with multimodal planning, especially when considering new vehicle types such as electric vehicles and cargo bikes, is new rules which apply to subsets of the overall transport resources. Furthermore, the system needs to be configurable to adapt to the operational and workflow structure of the organisation. For example different processes and managers might be involved for managing van deliveries compared to the last-metre management of the cargo bike from the microhub.

The SCF-VRO models end-to-end movements as a series of connected transport legs which may be allocated to different modes e.g. ship->rail->truck, or truck->van->bike. Consolidation centres or microhubs typically provide the interface between these legs, and often act as a natural point where control hands over from one planner to another. Operational rules and constraints need to be defined by vehicle type, territory and operation. For example, for a given group of microhubs shift times, breaks and other rules may be agreed with the workforce, while van operators, and certainly HGVs, will have other constraints to comply with such as driver hours rules for large vehicles.

In the first version of the SCF-VRO the underlying database and algorithm architecture have been enhanced to accommodate a wider range of constraint types, and a more flexible model for applying them. Examples have been tested using HGVs, vans and bikes, and also multimodal cases. The intention is that this flexible architecture will be used to model the more challenging scenarios and use cases contemplated in the later project activities in WP6.

3.3 LESSONS LEARNED & NEXT STEPS

The first version of the Synchromodal Conflict-Free Vehicle Routing Optimiser solutions and components have been completed and initial tests proved successful. The emphasis has been on created a flexible toolset that can address a wide range of freight and logistics optimisation problems, targeting operational efficiency and increased resilience to congestion and disruptions in the transport network.

During the work a number of important developments and considerations emerged. Opportunities for modelling a wider range of freight types have been discussed, and an initial approach has been designed which is expected to offer greater flexibility in this

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

respect. This applies particularly to the intermodal case, where the additional complexities of specific handling and loading constraints need to be considered.

Interaction and discussions with the Task T4.1 partners working on the social optimum models for traffic management have been very productive, leading to the further enhancement of the SCF-VRO to contemplate new congestion-avoidance strategies for routing vehicles in urban areas taking into account more refined information about transit times at junctions.

As envisaged in the original project plan, the first version of the SCF-VRO already contains major enhancements and new functionality which can be used for further trials and tests in the next stages of the project. However this first version is intended to be an interim solution which will be extended and refined over the remainder of the project, including further additions to the functionality and refinements in response to the results of further tests and assessment.

Funded by the European Union

Project funded by

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

4 DYNAMIC RESPONSE PLANNER

4.1 INTRODUCTION

Road traffic incidents including recurrent and non-recurrent disruption in traffic flow, such as bottlenecks or collisions, reduce roadway capacity. These may quickly lead to congestion and associated travel delay, increased pollutant emissions and wasted fuel. Therefore, traffic management tools are crucial in minimising the impact of such incidents and their negative consequences on network safety and efficiency. Effectively, incident response management systems operate once an incident is detected and confirmed to maintain a reasonable level of safety for all road users, preserve and protect human life, and minimise travel delay. The Dynamic Response Planner (DRP) represents one of the main FEDORA services for network-wide traffic optimisation.

A response plan (RP) may include one or more response actions (RAs) or mitigation measures which can be used to mitigate the negative impacts of detected or predicted road traffic incidents (i.e., collisions or bottlenecks) and eventually to restore the normal or optimal network capacity. Additionally, the process of generating RPs considers road network topology and characteristics including road types, number of lanes, and maximum speed limit. Moreover, it considers available modes of transportations, applicable and inapplicable response actions. Applicable or valid response actions involve the actions considered and supported by traffic management operators (TMOs) and can be implemented and propagated in the traffic network. However, inapplicable or invalid response actions involve the actions unsupported by TMOs, so these actions cannot be integrated as part of any generated response plan. Generally, a range of response actions can be considered and implemented, including re-directing road users away from the incident area via Variable Message Signs (VMS), motivating multimodality, encouraging Park and Ride schemes, re-routing planned trips, changing traffic lights timings, closing some lanes of a road, and reducing the maximum speed limit of roads.

Effectively, the Complex Event Processing Engine (FEDORA, Task 4.4) will define emerging situations on the transportation network and services and outline predicted scenarios which may unfold soon. It uses various predictive analysis and anomaly detection algorithms to predict deteriorating network conditions and scenarios. Therefore, these predicted scenarios may require a dynamic response plan generation to proactively avoid detected incidents or to mitigate impact of these incidents. However, each response plan may contain one or multiple mitigation measures or response actions. Generally, to minimise the impact of traffic incidents, traffic network actors and road users must be supplied with timely useful and accurate information, guidance, and instructions to plan rapidly, accurately, and orderly during incident response.

Accordingly, the Dynamic Response Planner will generate a response plan based on real time network conditions and scenarios that may be enacted on the network. Commonly, a response plan may include traffic diversion, traffic flow regulation and dynamic information dissemination to road users. However, the Dynamic Response Planner leverages from and integrates several network optimisation strategies provided by different FEDORA network optimisation modules to generate appropriate response plans. In practice, the DRP will be implemented as service, or microservice, and integrated with other project solutions such as FEDORA platform.

4.2 OBJECTIVES OF THE DYNAMIC RESPONSE PLANNER

The Dynamic Response Planner is a service that will support traffic network actors and road users to diminish travel delay, minimise negative impact of road traffic incidents, and enhance traffic network safety and efficiency. It aims to generate a response plan when a traffic incident is detected or predicted. Additionally, it operates by implementing a rule-based engine which uses all available input data from the Complex Event Processing Engine to assess the detected incident and select the most appropriate response actions. These actions will be integrated together to form a response plan. Effectively, the DRP will consider the predictions provided by the Complex Event Processing Engine (Task 4.4) to propose a range of mitigation measures intended to reduce the negative impacts of disruptive events. Additionally, the DRP may optionally integrate the outputs of other transport network optimisation

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation

Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

modules (T4.1, T4.3, and T4.5) subject to availability of these services. Mitigation measures or response actions included within a generated response plan will be transformed and adjusted to a form that they can represent the necessary and actual changes and implementations required in the transport network. Figure 21 shows the interfaces of the Dynamic Response Planner with other FEDORA network optimisation modules and external systems which represent data sources of this component. It also presents the expected input data from these data sources and projected output data of the DRP.

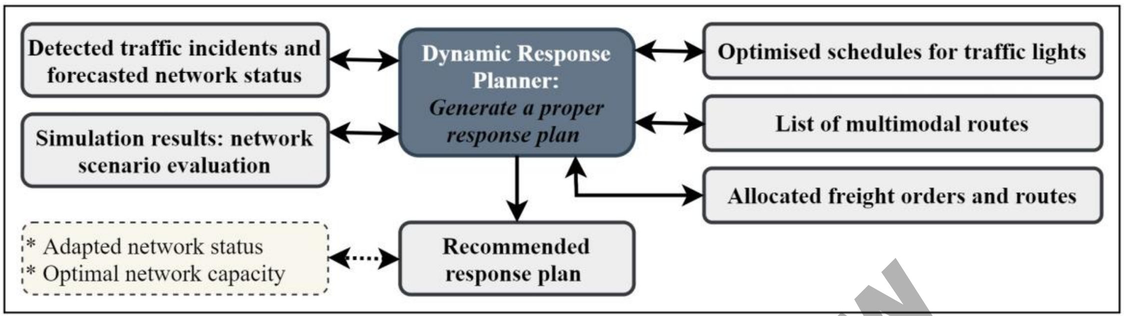

Figure 21: Input and output data of the Dynamic Response Planner

In essence, the DRP is a tool for enhanced traffic management efficiency and safety, designed to act as a proactive safeguard for road users by suggesting real time response plans and supporting the decision-making process. The DRP is part of the network management services space and therefore it connects and interacts with the data space from one side and with the simulation space from the other side. It receives input data and uploads output data to the data space. Additionally, it sends network scenarios (e.g., planned, predicted, or recommended scenario) to be executed in the simulation space and receives the corresponding results from this space. Therefore, the DRP utilises a simulation-based performance monitoring tool to evaluate the network status or scenario and ensure that the generated and recommended response plan is suitable and effective for improving traffic flow, reducing congestion, and enhancing urban mobility. Having the network optimisation strategies considered and used as the output of FEDORA network optimisation services, the DRP integrates a combination of mitigation measures, tests the impact of these measures by using a simulation framework, and then uses the simulation results to refine and recommend an appropriate response plan.

The DRP mainly uses, as input, the output of FEDORA network monitoring and optimisation services including the Complex Event Processing Engine, the Social Optimum Models for Multimodal Nodes and Corridors, the Multimodal Journey Planner Tool, and the Synchromodal Conflict-Free Vehicle Routing Optimiser. So, when the status of a traffic network deteriorates and a response plan is required, the DRP analyses these inputs and provides real-time recommendations as mitigation measures which can be considered and employed to restore normal or optimal network capacity.

Output data of the services and modules developed as part of Task 4.4 represents the main input to the DRP. This includes the forecasted traffic network status (i.e., flow, speed, density, travel time), predicted incidents, and detected critical network anomalies which may require specific remedial actions and response plans to be generated and propagated. However, output data of the Social Optimum Models for Multimodal Nodes and Corridors developed as part of Task 4.1 includes optimised schedules for traffic lights at each network intersection, having the network status as input data. Moreover, the Multimodal Journey Planner Tool developed as part of Task 4.3 generates a list of optimal multimodal routes for each trip request between an origin and a destination (OD), having the network status and origin-destination pair (OD) as input data. Additionally, the Synchromodal Conflict-Free Vehicle Routing Optimiser developed as part of Task 4.5 allocates freight orders and generates optimal multimodal freight routes for a list of delivery orders between a list pick up and drop off points, having the network status and logistics orders as input data. Each of these three network optimisation services (Task 4.1, Task 4.3, and Task 4.5) has been developed to provide an optimal solution in terms of its own objective. Therefore, the DRP combines and considers the behaviour of all these network optimisation modules including the provided solutions to generate one solution which will mitigate the consequences of detected incidents and anomalies.

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

This solution represents the adapted status of the transport network which makes the mentioned network optimisation services work together to restore and maintain normal or optimal network capacity.

4.3 PROPOSED APPROACH FOR GENERATING RESPONSE PLANS

The DRP is designed and developed to generate real-time optimal response plans for traffic network incidents and congestions. The conceptual architecture of the proposed DRP is shown in Figure 22. In this approach, the DRP involves two sub-components: Network Monitoring Tools (NMT) and Response Plan Generator (RPG). The NMT is responsible for:

- Monitoring and analysing the output of T4.4,
- Activating the functionalities of the DRP, if a RP is required,
- Executing FEDORA network optimisation services (i.e., send input and receive output),
- Building network scenarios,
- Assessing the performance of the built scenarios (i.e., via simulation framework).

However, the RPG is responsible for:

- Analysing and examining the evaluated scenarios,
- Identifying appropriate response actions and measures,
- Encoding the response plan into actionable measures (i.e., to identify a new scenario),
- Interacting with the NMM to create and test the new scenarios,
- Deciding if further response plan refinement is required,
- Recommend the appropriate response plan (i.e., to restore optimal network capacity).

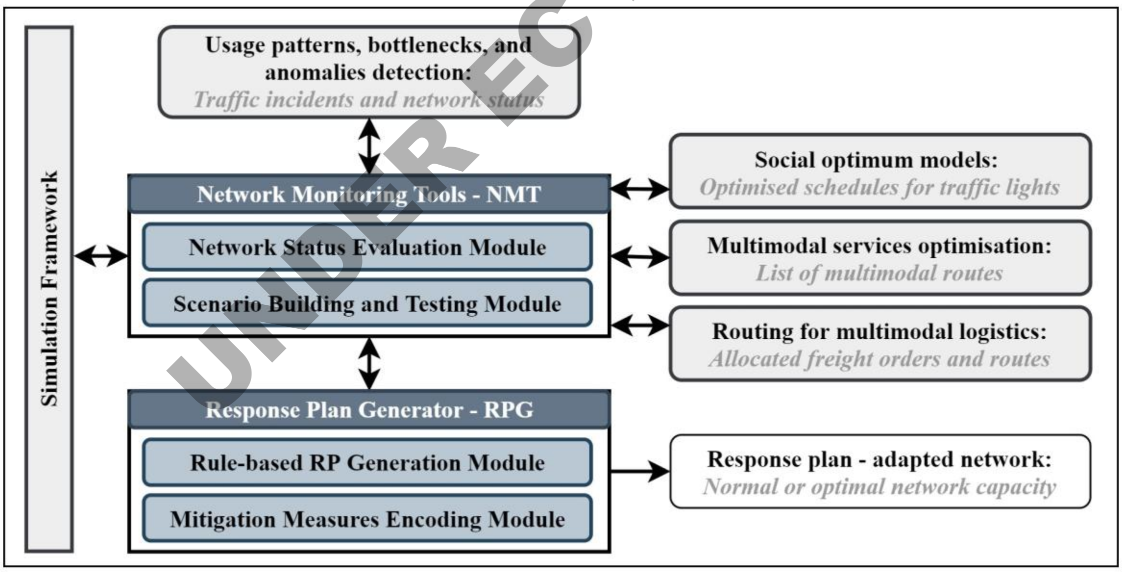

Figure 22: The conceptual architecture of the Dynamic Response Planner

4.3.1 NETWORK MONITORING TOOLS

The Network Monitoring Tools (NMT) frequently interacts with the Complex Event Processing Engine (Task 4.4) and gets input data related to detected incidents and predicted network status. This network status including detected incidents and anomalies represents the network baseline scenario which needs to be adjusted and enhanced through the implementation and propagation of appropriate response plans. Effectively, the output of Complex Event Processing Engine (Task 4.4) may clearly indicate that a response plan is required or alternatively the DRP will decide if a response plan is required. The NMT examines and evaluates the

Funded by the European Union

Project funded by
Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
Swiss Confederation
Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

network status with the support of a simulation framework and may decide whether a response plan generation is required. This decision can be based on a threshold calculated by comparing the network baseline scenario (i.e., congested network) with free flow network conditions. For example, this threshold may represent the increasing percentage in the network-level or section-level travel time (TT).

When a response plan generation is required, the NMT interacts with other relevant and available WP4 network optimisation services, provides input data (i.e., baseline or enhanced network scenario), and gets their outputs. However, the input data required for each of these services and the expected output data have been described above (in section 4.2). Additionally, the NMT integrates some of these outputs (i.e., optimised schedules for traffic lights) with the baseline scenario and creates new network scenarios to be evaluated by the simulation framework. However, other outputs (i.e., (Task 4.3 and Task 4.5) may be used as indicators of network performance to validate the network status enhancement. Therefore, a dedicated simulation framework and model (i.e., Simulation of Urban Mobility – SUMO simulation model for a pilot site) is required for testing and evaluating various network scenarios.

4.3.2 RESPONSE PLAN GENERATOR

The Response Plan Generator (RPG) receives the created and evaluated network scenarios from the NMT when a response plan generation is required. It analyses this data and identifies road links that are still congested in the network depending on a threshold related to these links' traffic condition. For example, a link is considered as congested if the volume to capacity ratio (VCR) of this link is greater than 0.75. Therefore, no congested links means no response plans are required. Effectively, the RPG uses a rule-based engine and considers the congestion severity (i.e., VCR level) at each link to propose appropriate mitigation measures that can increase the travel time of this link. These procedures aim to encourage road users to avoid these congested links and guide the FEDORA network optimisation services to further improve their solutions in terms of planning routes away from these links and increasing the green timing of traffic lights related to these links.

The RPG encodes these proposed measures into actual network management actions such as closing one lane of the link and reducing the maximum speed limit of this link by k%. Then, the RPG sends this created network scenario including the proposed mitigation measures back to the NMT which will iterates the procedures mentioned above (in section 4.3.1). So, the NMT will execute other relevant and available FEDORA network optimisation services, create new scenarios and evaluate the performance of these scenarios. Essentially, this loop of actions (iterations) will continue until the normal or optimal network status is achieved (i.e., the VCR of all network links is less than or equal to 0.75). Therefore, this process refines and recommends the optimal adapted conditions for the traffic network that maintain normal network capacity.

4.3.3 SIMULATION FRAMEWORK

Traffic network simulation frameworks are important to replicate real-world transportation systems, test the impact of infrastructure changes in a risk-free environment, and identify conflict hotspots and bottlenecks before they happen. Effectively, the DRP utilises a simulation framework to test and evaluate the performance of traffic network scenarios. The NMT involves a Scenario Building and Testing Module which is responsible for 1) creating traffic network scenarios based on the data collected and analysed by other NMT services, and 2) evaluating these scenarios with the support of a simulation framework. So, Simulation of Urban Mobility (SUMO) which is an open-source traffic simulation framework has been used for modelling an example transport network and traffic dynamics (i.e., Vienna pilot site). This simulation framework and model are used as tools to analyse and evaluate the forecasted network status (baseline scenario), the impact of FEDORA network optimisation services (i.e., optimised schedules for traffic lights) on the network status, and impact of suggested mitigation measures (response plan) on the network status.

4.4 OPERATIONAL WORKFLOW AND TECHNICAL ARCHITECTURE

4.4.1 OPERATIONAL WORKFLOW

The design and architecture of the Dynamic Response Planner have been developed with a focus on its integration with other FEDORA network management services developed as part of WP4 tasks. So, rather than functioning as a standalone solution, the

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation
Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

DRP receives the network status and detected anomalies from the Complex Event Processing Engine (Task 4.4), and acts as an orchestration and decision-support layer that can receive, analyse, and generate response plans by effectively incorporating and utilising the optimisation strategies generated by the corresponding FEDORA components (Task 4.1, Task 4.3, and Task 4.5). Additionally, it interacts with a simulation framework to evaluate various conditions and scenarios of transport network. Therefore, the DRP has several interfaces with other FEDORA components and external systems.

The development of the DRP and relevant FEDORA tasks and services are progressing in parallel. Therefore, this chapter of the deliverable summarises the progress so far in FEDORA Task 4.6 and the development of the Dynamic Response Planner. Additionally, main aspects of the relevant solutions are still not well defined at the time of preparing this document, and need further clarification (i.e., by the developers of these services). Therefore, the progress will continue, and the solution described here will be refined and extended as the project progresses.

Effectively, the concept and logic of the DRP can be seen as an integrator or orchestrator of the different FEDORA network optimisation services. The DRP mainly monitors the output of the Complex Event Processing Engine (Task 4.4) and starts working as soon as a response plan is required. Each of the relevant network optimisation services developed at part of WP4 (Task 4.1, Task 4.3, and Task 4.5) considers a specific optimisation objective. However, the DRP considers all these objectives collectively. Therefore, the DRP continuously interacts with the Social Optimum Models for Multimodal Nodes and Corridors, the Multimodal Journey Planner Tool, and the Synchromodal Conflict-Free Vehicle Routing Optimiser, provides these services (subject to availability) with input data and gets the corresponding output data. Then, the DRP proposes an appropriate response plan which enables these services to provide optimal collective solutions and restores optimal network capacity. Therefore, the DRP will introduce specific changes (i.e., response actions) to the network supply which mainly focus on reducing the capacity of congested road links. It has been assumed that the capacity reduction will guide the relevant network optimisation services to address network issues, adapt their solutions, and achieve optimal network capacity.

Figure 23 shows the operational workflow of the Dynamic Response Planner which will be used to describe the functionalities of the DRP sub-components. Additionally, this workflow illustrates how these sub-components interact with each other, with other network optimisation services, and with external systems. Therefore, it demonstrates the data exchanged, the decisions to be taken, and the process of generating a response plan when a network anomaly or incident is detected by the Complex Event Processing Engine (Task 4.4).

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation

Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

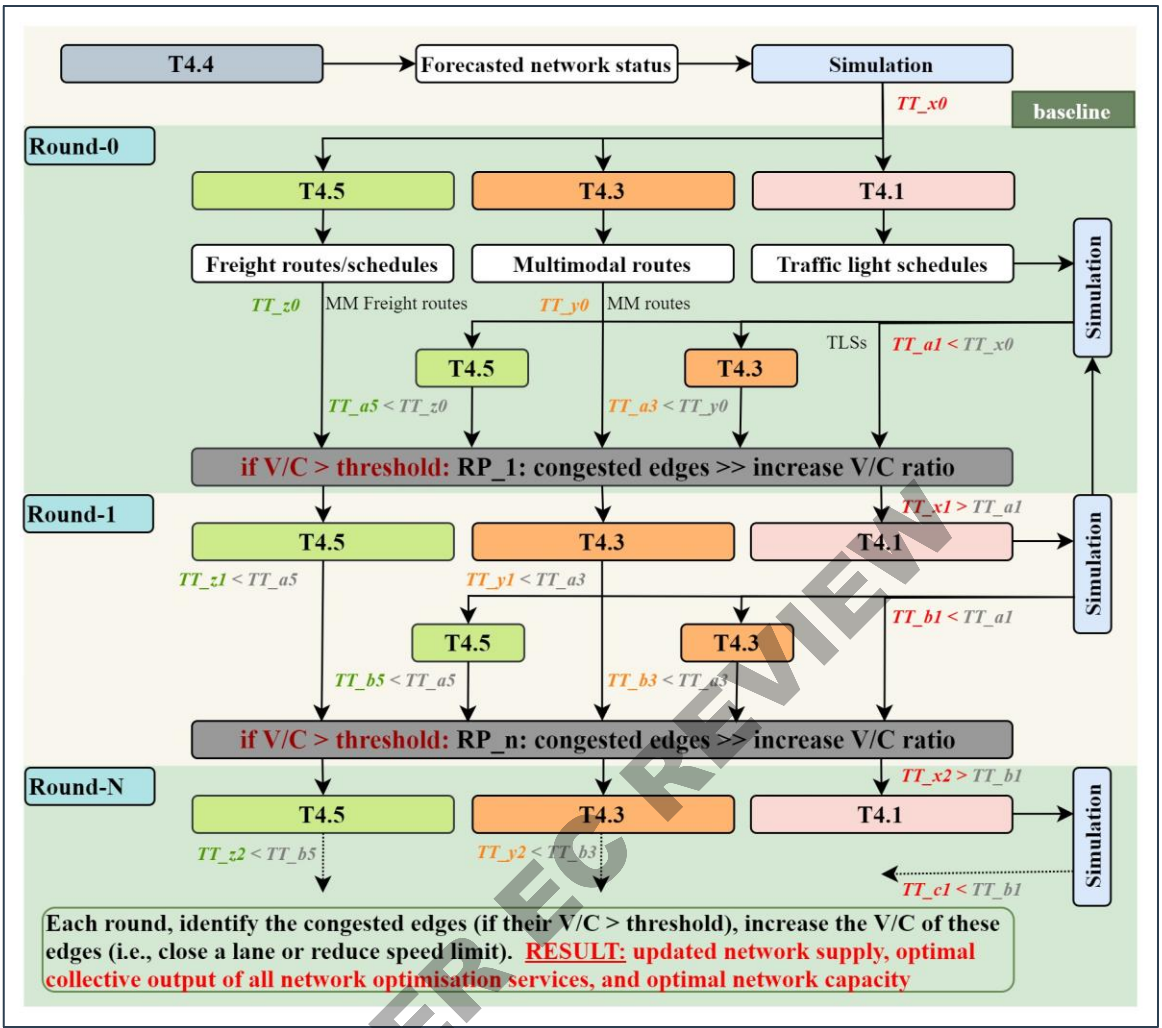

Figure 23: The operational workflow of the Dynamic Response Planner

A. Incidents and anomalies detection

In the proposed approach of the Dynamic Response Planner, it has been assumed that either the Complex Event Processing Engine clearly indicates that a response plan is required or the DRP will analyse and examine the output of this engine (i.e., the baseline scenario) and compare it with the normal network conditions (i.e., free flow scenario) to decide if a response plan is required. Therefore, link-level travel time (TT) of these two network scenarios can be compared and if the difference is greater than a pre-defined threshold (i.e., the travel time is increased by 15%) then the DRP will be triggered to start functioning.

B. Round-0: initial network monitoring and scenario evaluation

When a response plan is required, the DRP analyses the aspects of the deteriorated network status. However, it is not clear yet how this deteriorated network status which represents the output of the Complex Event Processing Engine will be structured and what information is included. So, it has been assumed that this deteriorated network status represents the baseline scenario (i.e., TT_x0 in Figure 23) and involves the travel time (TT) of different network links and the average travel time of the full network. This baseline scenario will be used as input data for the Social Optimum Models for Multimodal Nodes and Corridors which will generate the optimised schedules for traffic lights at each intersection as output data. These optimised schedules will be integrated with the baseline scenario through updating these schedules in the relevant simulation model (Figure 24) to create a new network scenario. Like the baseline scenario (TT_x0), this new network scenario (TT_a1) will be simulated and evaluated. These evaluations will reveal

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

the network status optimisation gained from the Social Optimum Models for Multimodal Nodes and Corridors. However, the application of the Multimodal Journey Planner Tool may use the main input data (i.e., origin and destination coordinates) to generate a list of multimodal routes. Therefore, it will generate TT_y0 if the baseline scenario is considered while it will generate TT_a3 if the new network scenario id considered as input data. Comparing the average travel time of these two sets of routes (TT_y0 and TT_a3) will indicate the impact of the Social Optimum Models for Multimodal Nodes and Corridors on the baseline scenario. The application of the Synchromodal Conflict-Free Vehicle Routing Optimiser is like that of the Multimodal Journey Planner Tool (TT_z0 and TT_a5).

Figure 24: A screenshot for intersections with traffic lights in SUMO and a sample of python script for updating the schedules of a traffic light

C. Evaluating the updated network and checking if another round is required

In the Round-0, the baseline scenario has been simulated, FEDORA network optimisation services have been executed, and the optimisation strategies incorporated with the baseline scenario. Then, this updated network scenario (i.e., TT_a1 represents TT_x0 with optimised schedules for traffic lights integrated) will be evaluated and checked to identify remaining congested links in the network, if there is any. If there is no congested links in the network, then no response plan is required, and the recommended network status or response plan will be (TT_a1). However, when at least one congested road link is identified and detected, then the next round in the workflow (i.e., Round-1) will be executed and implemented. Effectively, the Volume to Capacity Ratio (VCR), sometimes referred to as the v/c ratio or degree of saturation of road links will be calculated and used to identify congested links. So, a road link is considered as congested if the VCR of this link is greater than a threshold (i.e., VCR>0.75). The VCR is a metric used in traffic engineering to measure the level of congestion on a road link. It compares the actual volume of traffic to the maximum capacity of the road. The transition of a road link from free flow to a blocked road is defined in Table 5. It shows that if the VCR is below 0.61, then the traffic is operating freely and stable. When the VCR exceeds 0.85, the road approaches its maximum capacity, speeds drop, and delays increase significantly. Therefore, the threshold of VCR=0.75 may be selected to identify congested links.

Funded by the European Union

Project funded by
Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
Swiss Confederation
Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

Table 5: VCR and the transition of a road link from free flow to a blocked road

VCR	Road link condition	Description
0.00 – 0.20	Free Flow	Travel at or near the speed limit without experiencing delays
0.21 – 0.60	Stable Flow	Traffic volume increases, but it is still a stable, unimpeded flow
0.61 – 0.85	Nearing Congestion	Heavy traffic and speeds begin to drop below free-flow speed
0.86 – 1.00	Saturated Flow	The flow becomes unstable, queues and longer delays
> 1.00	Blocked Road	Long queues form, throughput drops, and travel times spike

D. Generating a response plan to restore normal network capacity

Having the congested road links identified, the VCR of each link will be increased, or its capacity will be decreased. Even though this seems to aggravate the link condition, it will be considered as a congestion mitigation measure because it will not be considered as a solution by itself, but it guides other integrated network optimisation services to enhance the network capacity. For example, decreasing the capacity of a road link will make the Social Optimum Models for Multimodal Nodes and Corridors to further consider the traffic lights controlling the traffic in this link and propose optimal schedules that resolve the congestion in this link. Additionally, this will guide the Multimodal Journey Planner Tool and Synchromodal Conflict-Free Vehicle Routing Optimiser to avoid this link, re-route road users away from this link, and therefore, mitigate the congestion in this link. Practically, closing specific lanes of a road and reducing the maximum speed limit of roads are examples for reducing the capacity of a road link. Mainly, various network conditions related to this round (e.g., TT_x0, TT_a1, TT_y0, TT_a3) determine the type and value of the proposed changes to the congested road links at the end of the round (Round-0). Table 6 shows an example for a rule-based mechanism that support the decision-making process of selecting which mitigation measures to apply, having the VCR of congested links and the number of link lanes. Therefore, the proposed response plan including the changes introduced to all congested road links will be integrated with the TT_a1 network scenario through updating the relevant simulation model (Figure 25) to create a new network scenario (i.e., TT_x1). This new network scenario represents the input to the next round (Round-1). Figure 25 shows an example for reducing the capacity of a road link with three lanes and a VCR of 0.90, so one lane of the link is closed and the speed limit of other two lanes are reduced by 20%.

Table 6: An example for rule-based mechanism to select the appropriate mitigation measures for congested road links

VCR of road link	Number of lanes	Example of mitigation measures
0.75 < VCR < 0.85	1	decrease speed limit by 20%
	>1	close one lane and decrease speed limit by 10%
0.85 < VCR < 0.95	1	decrease speed limit by 40%
	>1	close one lane and decrease speed limit by 20%
0.95 < VCR	1	decrease speed limit by 60%
	>1	close one lane and decrease speed limit by 30%

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation

Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

Figure 25: A screenshot and a sample of python script for closing 1 lane and reducing the speed of 2 lanes of a 3-lanes link in SUMO

E. Round-1: further network monitoring and scenario evaluation

In this round, the procedures and actions mentioned in steps B, C, and D will be repeated and the input network scenario will be TT_x1 instead of TT_x0 (in Round-0). Therefore, the network monitoring and scenario evaluation step (step B) will use this network scenario (TT_x1) as input data for the Social Optimum Models for Multimodal Nodes and Corridors to generate an optimised set of schedules for traffic lights which will be integrated with TT_x1 to create another network scenario (TT_b1). Additionally, the application of the Multimodal Journey Planner Tool with input data of TT_x1 and TT_b1 provides TT_y1 and TT_b3. However, the application of the Synchromodal Conflict-Free Vehicle Routing Optimiser with input data of TT_x1 and TT_b1 provides TT_z1 and TT_b5. In this round, the updated network scenario (TT_b1) will be examined to identify remaining congested links (step C). If there is no congested links in the network, then no response plan is required, and the response plan will be (TT_b1). Otherwise, the next round in the workflow (i.e., Round-2) will be executed and implemented. Effectively, a response plan will be generated for congested road links to restore normal network capacity (step D). Therefore, the VCR of each congested link will be increased, or its capacity will be decreased. The proposed response plan including the changes introduced to all congested road links will be integrated with the TT_b1 network scenario through updating the relevant simulation model to create a new network scenario (i.e., TT_x2). This new network scenario represents the input to the next round (Round-2). Therefore, several rounds may be required to ensure network stabilization and optimal capacity.

4.4.2 DYNAMIC RESPONSE PLANNER AND FEDORA NETWORK OPTIMISATION SERVICES

The DRP conceptual architecture and operational workflow show the relationship between the DRP from one side and other FEDORA network optimisation services (Task 4.1, Task 4.3, and Task 4.5) from the other side. Therefore, the core functionalities of the DRP are based on utilising and integrating these solutions to generate and suggest response plans. However, one or more of these network optimisation services could be missed in some pilot sites or use cases. Therefore, the operational workflow is designed to

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

be dynamic so it can consider the aspects of different use cases. Additionally, it has been assumed that the strategies generated by the Task 4.1 (i.e., optimised schedules for traffic lights) can be integrated with network scenarios to update the network supply. However, the outputs of Task 4.3 (and Task 4.5) which include a list of multimodal routes between origin locations and destinations can be used as indicators to measure the performance of the underlying network (e.g., the average travel time of the generated routes). Accordingly, the DRP focus on manipulating some network supply aspects in such a way to ensure that FEDORA network optimisation services can operate and provide optimal solutions for the deteriorated network and restore normal network capacity. Therefore, when a traffic incident is detected and a response plan is required, the DRP automatically generates and recommends an appropriate response plan. Figure 26 shows two examples of recommended response plans generated by the DRP. Each response plan includes the response actions proposed to avoid the detected congestion or to mitigate the impacts of predicted incidents.


```
{
  "responsePlan": {
    "planId": "RP-001",
    "name": "Congested Road Link Mitigation",
    "status": "Active",
    "createdAt": "2026-07-03T00:00:00Z",
    "link": {
      "linkId": "LINK-12345",
      "roadName": "Example Highway",
      "direction": "Northbound",
      "location": {
        "startJunction": "Junction A",
        "endJunction": "Junction B"
      },
      "lanes": 3,
      "currentSpeedLimitKph": 100
    },
    "mitigationActions": [
      {
        "actionId": "ACT-001",
        "type": "LaneClosure",
        "description": "Close one lane to manage congestion.",
        "targetLane": 3,
        "status": "Scheduled"
      },
      {
        "actionId": "ACT-002",
        "type": "SpeedLimitReduction",
        "description": "Reduce the maximum speed on the 3-lane road.",
        "originalSpeedLimitKph": 100,
        "reducedSpeedLimitKph": 80,
        "appliesToLanes": [1, 2],
        "status": "Scheduled"
      }
    ],
    "expectedOutcome": {
      "improveTrafficFlow": true,
      "reduceCollisionRisk": true,
      "estimatedDurationMinutes": 60
    }
  }
}
```

```
{
  "responsePlan": {
    "planId": "RP-002",
    "name": "Congested Road Link Speed Reduction",
    "status": "Active",
    "createdAt": "2026-07-06T10:30:00Z",
    "link": {
      "linkId": "LINK-67890",
      "roadName": "Example Highway",
      "direction": "Northbound",
      "lanes": 1,
      "lengthMeters": 1250,
      "currentMaxSpeedKph": 60
    },
    "mitigationActions": [
      {
        "actionId": "ACT-001",
        "type": "SpeedLimitReduction",
        "description": "Reduce the maximum speed on the road.",
        "originalSpeedLimitKph": 100,
        "reducedSpeedLimitKph": 60,
        "reason": "Traffic congestion",
        "effectiveImmediately": true
      }
    ],
    "expectedOutcome": {
      "improveTrafficFlow": true,
      "reduceCollisionRisk": true,
      "estimatedDurationMinutes": 60
    }
  }
}
```


Figure 26: Two examples of response plans generated by the DRP

4.4.3 TECHNICAL ASPECTS OF THE DYNAMIC RESPONSE PLANNER

The process of generating response plan may require several operational rounds (Figure 23) and therefore this may involve several runs of the related simulation. So, this introduces some delay and increases the incident-response time. Effectively, the execution time of a simulation run (i.e., SUMO) to compute a scenario depends on numerous factors such as network size and complexity,

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation

Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

vehicle count and density and using a GUI or running in batch mode. In general, these aspects, among others, can be considered to reduce the runtime of simulations but this could be insufficient to realise a reasonable DRP runtime. Therefore, a machine learning technology can be used to train the data collected during the implementation of the operational workflow including the characteristics of detected incidents, deteriorated network status, and produced network scenarios. This data is trained offline to get a machine learning model which can then be used only with the output of the Round-0 to generate a real time response plan. Therefore, this will reduce the number of simulations runs required for generating a response plan and significantly enhance the runtime of the DRP.

The development of the DRP uses Python, leveraging the Flask framework to provide a lightweight and scalable web application architecture. The application exposes a RESTful API that enables seamless communication with external systems through standard HTTP methods. The API is designed to interface with pilot-based network optimisation services, allowing optimisation requests and responses to be exchanged efficiently. Integration with the simulation framework enables automated execution and evaluation of network simulation scenarios. Data exchanged between components is validated to ensure consistency and reliability throughout the workflow. The RESTful endpoints process client requests and return structured responses in JSON format, supporting interoperability across different platforms. This architecture provides a flexible, extensible, and efficient platform for network optimisation and simulation applications.

4.5 NEXT STEPS

The Dynamic Response Planner is highly dependent of the output data of the Complex Event Processing Engine including the forecasted network status and detected anomalies. Additionally, the DRP leverages from and utilises other FEDORA network optimisation modules developed as part of WP4 tasks. Mostly, the plan is to implement and demonstrate the DRP in two FEDORA pilots (i.e., Vienna and Nicosia) which may consider different network optimisation services. Therefore, arrangements have been considered to collect the essential input data and information required for the development of the DRP. Effectively, SUMO simulation framework has been employed to develop an example network model. This simulation framework and model have been used with synthetic data to develop and test the functionalities of the DRP. Additionally, real input data and information including input and output data type and schema, network scope, and used simulation framework have been requested from the task leaders who are developing the relevant network optimisation services and organising the pilot deployments. Therefore, real input data collected from different network optimisation services and actual simulation models developed as part of the simulation space (i.e., FEDORA WP5) will be used instead of the synthetic data and the proposed simulation model. Accordingly, further development and refinements of the DRP are foreseen to complement to the ongoing work.

Funded by the European Union

Project funded by Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra Swiss Confederation

Federal Department of Economic Affairs, Education and Research EAER State Secretariat for Education, Research and Innovation SERI

5 CONCLUSION

This deliverable has presented the design and initial development of a range of services for multimodal network optimizations in the FEDORA project, focussing specifically on complex vehicle routing solutions, incident and anomaly detection and prediction, and response plan generation. This report has provided a comprehensive description of each tool including its rationale and design decisions, its scope, development progress so far, as well as where work will focus in the future.

More specifically, the Complex Event Processing Engine (CEPE) developed a range of tools related to real-time traffic data analysis, including Generative AI tools for incident detection, data generation, and automated data labelling, traffic forecasting and anomaly detection tools, and computer vision tools for collecting traffic data using unmanned aerial vehicles. The Synchromodal Conflict-Free Vehicle Routing Optimiser (SCF-VRO) developed a system that provides freight and logistics optimization solutions for operational efficiency and resilience in response to unexpected disruptions in the transport network. Finally, the Dynamic Response Planner (DRP) presented the architecture and logic for a mitigation plan generator in response to detected disruptions, that combines multiple recommendations into a comprehensive plan.

At this stage in the FEDORA project the majority of the components have been developed mostly in isolation, providing proof-of-concept and initial calibration work, in order to establish a robust foundation. In the next stage the tools and services presented here will be used in practical use cases from real-world pilot sites, where they will be deployed in combination with solutions and work from other Tasks and Work Packages to produce innovative multimodal services as part of the FEDORA platform.

Funded by the European Union

Project funded by

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI

6 REFERENCES

1. “FEDORA Grant Agreement Number 101203465“, 2025-05-05
2. Makrigiorgis, R., Kyrkou, C. and Kolios, P. (2023). How high can you detect? Improved accuracy and efficiency at varying altitudes for aerial vehicle detection. In 2023 International Conference on Unmanned Aircraft Systems (ICUAS), pp. 167–174. IEEE.
3. Papadatou, K. N., Fafoutellis, P., & Vlahogianni, E. I. (2025). Revisiting the non-recurrent traffic incident identification problem for real-time applications using theory-aware unsupervised learning. Data Science for Transportation, 7(2), 12.
4. Yoon, J., Jarrett, D., & Van der Schaar, M. (2019). Time-series generative adversarial networks. Advances in neural information processing systems, 32.

Funded by the European Union

Project funded by

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI