Ref. Ares(2025)6197324 - 30/07/2025

EVOROADS

INFRASTRUCTURE MONITORING TOOLS V1

Project deliverable D2.1

HORIZON Research and Innovation Action | 101147850 | 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	D2.1 – Infrastructure monitoring tools V1
WORK PACKAGE	WP2
TYPE OF DELIVERABLE	DEM — Demonstrator, pilot, prototype
DISSEMINATION LEVEL	PU - Public
STATUS – VERSION, DATE	Final - V1.0, 28/07/2025
DELIVERABLE LEADER	INLECOM
CONTRACTUAL DATE OF DELIVERY	31/07/2025
SUBMISSION DATE	29/07/2025
KEYWORDS	AI tools, infrastructure, road monitoring, deep learning, machine learning, computer vision, defect identification, interfaces, dashboards

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
Stathis Antoniou, Konstantinos Loupos, Despina Brasinika, Semina Mertzani	INLECOM
Leonidas Katelaris, Charalambos Kyriakou, Christos Kyrkou	UCY
Jānis Bērziņš, Andrejs Varfolomejevs, Edgars Starķis, Kristīne Tiltiņa, Liene Cīrule, Larisa Naumova	DOTS
Angela-Maria Despotopoulou, Panos Georgakis, Vivian Kiousi, Konstantinos N. Christidis, Babis Magoutas	FRONT
Hassane Essafi, Mathieu Carrier	CEA
Duc Trung Hoang, Dimitrios Papageorgiou	DTU
Djemre Ozmendji, Raul Villalba, Xavier Sol	IDIADA
Martín Rivas Caneiro	INDRA
Stefan Stefanescu	BEIA

QUALITY CONTROL

	Name	Organisation	Date
Peer review	Duc Trung Hoang, Gonzalo de Leon, Dimitrios Papageorgiou, Roberto Galeazzi	DTU	10/07/2025
Peer review	Kristīne Tiltiņa, Liene Cīrule, Larisa Naumova	DOTS	11/07/2025
Peer review	Despina Brasinika	INLE	16/07/2025
Peer review	Iain Macbeth	ERT	14/07/2025
Peer review	Angela-Maria Despotopoulou	FRONT	23/07/2025

VERSION HISTORY

Version	Date	Authors	Summary of changes
01	28/01/2025	Stathis Antoniou	First draft
02	27/06/2025	All partners	Version sent for review

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

03	22/07/2026	DTU, DOTS reviewers	Reviewed version
04	24/07/2026	Despina Brasinika, Iain Macbeth, Angela-Maria Despotopoulou	Reviewed version
05	25/07/2026	Stathis Antoniou	Final version

LEGAL DISCLAIMER

TABLE OF CONTENTS

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

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

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
AI	Artificial Intelligence
API	Application Programming Interface
AoI	Area of Interest
CCTV	Closed Circuit Television
CFD	Chicago Face Database (face-image dataset)
CMX	Connection Matrix Memory
CNN	Convolutional Neural Networks
CPU	Central Processing Unit
CSG	Constructive Solid Geometry (solid-modeling technique)
CVAT	Computer Vision Annotation Tool
DFG	Data Flow Graph
DFL	Distribution Focal Loss
DMP	Data Management Plan
ESA	European Space Agency
FN	False Negatives
FP	False Positives
GAP	Global Average Pooling layer
GFLOP	Giga Floating-Point Operations
GPS	Global Positioning System
GIS	Geographic Information System
GMM	Gaussian Mixture Model

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

GPR	Ground Penetration Radar
GTSDB	German Traffic Sign Detection Benchmark
HSV	Hue-Saturation-Value colour space
INSAR	Interferometric Synthetic Aperture Radar
IoU	Intersection over Union
IR	Infrared
IRI	International Roughness Index
ISAD	Infrastructure Support Assessment for Driving Automation
JSON-LD	JavaScript Object Notation for Linked Data
KPI	Key Performance Indicators
LiDAR	Light Detection and Ranging
LLs	Living Labs
LSTM	Long short-term memory
mAP	mean Average Precision
MDP	Mean Depth Profile
ML	Machine Learning
NaN	Not a Number
OBU	On-Board Unit
OEM	Original Equipment Manufacturer
OAK-D	OpenCV AI Kits specifically the Depth (Spatial AI stereo camera)
PSA	Partial Spatial Attention
PSI	Persistent Scatterer Interferometry
RDD	Road Damage Detection
RNNs	Recurrent Neural Networks

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

qcut	quantile-based binning
QGIS	Quantum Geographic Information System
RVC2	Robotics Vision Core 2
SAR	Synthetic Aperture Radar
SHAVE	Streaming Hybrid Architecture Vector Engine
SMDS	Safe Mobility Data Space
SMP	Symmetric Multiprocessing architecture
SNAP	Sentinel Application Platform
SPPF	Spatial Pyramid Pooling Fast
TN	True Negatives
TP	True Positives
UAV	Unmanned Aerial Vehicle
uid	Unique indentifier
VRUs	Vulnerable Road Users
WBCEL	Weighted Binary Cross-Entropy Loss
YOLO	You Only Look Once – Computer Vision Tool
ZOD	Zenseact Open Dataset homepage

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

The EC-funded EvoRoads (Evolutionary Solutions for Realising a Holistic Safe System Approach for All Road Users) is committed to advancing the European Union's Vision Zero initiative by implementing a comprehensive framework that integrates innovative models, tools, and services for data-driven safety assessments.

The project involves the development of a connectivity platform that digitalizes transport infrastructure assets and ensures the seamless integration of safety assessment services. By leveraging advanced artificial intelligence, EvoRoads analyses infrastructure monitoring data at various geospatial levels, enabling proactive risk warnings and supporting road operators in managing maintenance more effectively, which enhances both safety and operational efficiency. It also defines safety criteria and quantification methods for key performance indicators (KPIs) to monitor safety performance as part of the "Safe System" approach.

The results are developed alongside 5 axes: (1) an **Evolutionary Safety Assessment Framework** setting the basis by defining how to measure road safety in a multilayered catalogue; (2) a **Road Asset & Safety Management Digital Twin** collecting and distributing safety data; (3) **Advanced Monitoring Safety Systems** using that data that enable infrastructure owners to obtain clear insights on their roads' level of safety; (4) **Proactive Safety Warning Systems** ensuring that critical safety issues are raised pro-actively and in real-time to prevent safety hazards; and finally, all these results are combined in (5) **Solutions Integration, Augmentation and Impact Assessment tools and products** allowing for take-up of the results in the wider market. These axes converge into a **Safe Mobility Data Space** where all data related to EvoRoads' safety criteria and solutions are combined.

EvoRoads' methodologies and technologies will be validated in four Living Labs (LLs) addressing diverse road user scenarios across urban and rural environments to ensure broad applicability and effectiveness. EvoRoads LLs are situated in **Spain** (focusing on **Hazard-aware assets for better infrastructure safety level for existing and future roads users**), **Italy** (focusing on **Dynamic Road infrastructure safety diagnosis enhanced by emerging vehicle and digital technologies**), **Latvia** (focusing on **Automated (advanced) low-cost infrastructure monitoring and diagnostics: road hot spots, bridges, tunnels**), and Romania (focusing on **Remote sensing and warning technologies for better infrastructure and road conditions**).

The EvoRoads project runs from May 2024 until April 2027. A scale-up event is foreseen towards the end of the project to allow for the wider market to consider EvoRoads' results and to help provide ways forward.

Social Media link:

@evoroads-project

For further information please visit evoroads-project.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

DELIVERABLE EXECUTIVE SUMMARY

Deliverable D2.1 consolidates the outcomes of Task 2.2: AI-based tools for road infrastructure monitoring, which is focused on the development of infrastructure monitoring tools. D2.1 goal is to present the developed AI tools and their applicability to the pilots of the EvoRoads project and beyond.

More precisely, eight AI tools are documented in this deliverable: (i) a risk quantification model for road segments (INLECOM); (ii) a deep-learning crack and pothole detection model for Unmanned Aerial Vehicle imagery (UCY); (iii) a lightweight YOLOv11 model running on OAK-D (OpenCV AI Kits specifically the Depth) stereo sensors for real-time defect spotting (DOTS); (iv) a vision module that detects traffic signs despite occlusion and wear (CEA); (v) a vibration-plus-image anomaly detector for predictive maintenance (DTU); (vi) a ride-quality estimator that turns connected-vehicle telemetry into pavement status scores (INDRA); (vii) a tool that generates interferograms and flags areas of interests (FRONT); and (viii) a hazard-aware Closed Circuit Television (CCTV) pipeline that recognises debris, animals and pedestrians (EUT). All tools are described in Section 2 with details concerning used datasets, data preprocessing, model choices and validation metrics.

This collection of tools collectively addresses all types of roads (urban, semi-urban and rural) and are particularly useful in RURAL roads where traffic is less dense and citizens less likely to report and observe defects by themselves. The tools have been developed aligned with the project pilots' challenges as well as indicative domain challenges. Some tools have already been applied to pilots and are already using their data, while some others are currently discussing with pilots. The applicability of all tools at the time of writing of this deliverable has been assessed in a dedicated section.

Collectively, the developed tools now span macro-scale crash analytics down to millimetre-level surface inspection, and every Living-Lab pilot (Italy, Spain, Latvia, Romania) has at least two candidate tools mapped to its local needs. The next phase will require working together with the pilots to extend the tools applicability while, at the same time, integrating the tools to the EvoRoad's 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

1 INTRODUCTION

1.1 PURPOSE

This deliverable summarises the work done in T2.2 - "AI-based tools for cyber-physical road infrastructure monitoring". It presents the development of an array of algorithms for defect identification using AI tools and Computer Vision. Additional effort has been made to maximise the applicability of these tools to as many pilots as possible (or discuss their potential applicability). The developed tools can be used as services integrated into the Digital Twin (T3.1) and will be exposed to authorised users through the dashboards (T1.5).

This deliverable contributes to Objective 2 (O2): To facilitate holistic monitoring of urban and rural road infrastructures through the deployment of a federated platform of interconnected services and tools that can (a) detect infrastructure deterioration and distress conditions at different geospatial resolutions; (b) fuse heterogeneous data from crowdsourcing, satellite and aerial imagery, drone- and vehicle-mounted video capture technologies and multi-vehicle on-board sensors for optimal infrastructure monitoring; (c) proactively warn road-users of anticipated hazards through situational aware traffic management and on-the-edge computing and (d) enhance road maintenance operations through the construction of advanced diagnosis and prognosis roadmaps based on predictive analytics services.

More precisely, its Expected Innovation (EI) is EI2.2: An array of AI algorithms for the implementation of novel defect identification techniques to offer automated and efficient detection of various types of defects with higher consistency and efficiency compared to manual inspections.

1.2 MAPPING EVOROADS OUTPUTS

The table below shows how and which sections of this deliverable answer to both the deliverable and the task descriptions of the GA.

Table 1: Adherence to EvoRoads GA Deliverables & Tasks Descriptions

EVOROADS GA COMPONENT TITLE	EVOROADS GA COMPONENT OUTLINE	SECTIONS	JUSTIFICATION
	DELIVERABLE		
D2.1 - Infrastructure Monitoring Tools V1	The first release of the tools and models for the collection, processing and analysis of data for infrastructure monitoring	*2.1 - 2.8*	*Each section describes a tool/model developed. Every partner has his own section following a predefined structure including datasets description, data processing as well as tool development and validation.*
		3.1 - 3.3	*These sections are focused on the application of the developed tools.*

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 document the needed EU compliance, discuss their integration to the EvoRoads platform and their applicability to the project's pilots for infrastructure monitoring.
	TASKS		
Task 2.2: AI-based tools for cyber-physical road infrastructure monitoring	This task is responsible for developing and implementing cutting-edge tools for road infrastructure monitoring that leverage Street View and satellite imagery and computer vision/AI-based defect identification technology to provide real-time information on the condition of road equipment and infrastructure.	*2.1 - 2.8*	*These sections describe all AI tools developed for infrastructure monitoring.* *More precisely, we have timeseries models (2.1, 2.5 & 2.7), computer vision models (2.2, 2.3, 2.4) and a monitoring tool (2.6, 2.8). Satellite imagery is used in 2.8 while UAV imagery is used in 2.2. Tools of sections 2.2, 2.3, 2.4, 2.5, 2.6 & 2.8 are applicable real-time.*
	Using data feeds from cameras mounted on vehicles (e.g., ZED-2, FARO 3D Laser scanner, high-resolution smartphone cameras and more), the system captures high-resolution images up to a millimetre accuracy of the road network and using advanced computer vision algorithms and annotation tools is able to detect and label signs of distress or deterioration (i.e., Co-DETR, CVAT). 3D reconstruction of the road infrastructure will be attained through data fusion of the images and point clouds along with the geolocation.	*2.2, 2.3, 2.4 & 2.8*	*Onboard cameras are used in 2.2, 2.3 & 2.8. All computer vision models (2.2, 2.3, 2.4) used annotations tools (CVAT or similar) as well as advanced computer vision algorithms with deep learning architectures such as YOLO, UNet and other CNNs.*
	Moreover, to increase the robustness and performance of the AI models, data augmentations techniques will be deployed. (i.e.,CycleGANs). For object detection, the solutions will employ established architectures (e.g., YOLOv8, MaskRCNN, EfficientDet) aimed at localising and classifying road damage/deterioration.	*2.2, 2.3, 2.4 & 2.5*	*Data augmentation techniques were only used in 2.5 as the other object detection algorithms relied on existing datasets which did not require augmentation. Moreover, all object detection models (2.2, 2.3, 2.4) relied on established architectures such as YOLO, UNet and other CNNs.*

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

	Once a verified defect is identified, the system will generate an alert that will be communicated in near real-time to the relevant stakeholders, including maintenance crews and road operators. This enables prompt and targeted responses to any potential safety hazards, such as a damaged sign or a pothole in the road. To strengthen the analysis, deep learning models based on convolutional neural networks techniques (CNNs or SWIN transformers) will be utilized to detect and classify objects in the images. A set of interfaces will enable authorized users to access, through the dashboards (Task 1.5) the results of the defect identification system, such as defect reports, maps, predictive maintenance suggestions and other relevant data.	*3.3*	*The dashboard integrations of Task 1.5 and its corresponding alerting mechanisms are presented in dedicated section 3.3.*

1.3 RELATION TO THE OTHER WORK PACKAGES

This deliverable is a milestone of WP2 - “Advanced infrastructure monitoring and predictive maintenance tools” and it relates to WP1 - “Evolutionary safety assessment framework and solutions enablers”, WP3 - “Advanced tools for implementing ‘Safe System’ approach “and WP4 - "Demonstration of EvoRoads Solutions”.

The relation to WP1 is through T1.5 - “EvoRoads integrated platform for safer urban and rural environments” since part of T2.2 is to develop the interfaces to dashboards of T1.5 to trigger alerts on the findings of the developed tools.

The relation to WP3 is trough T3.1 - “Safe Roads: Digitalised ecosystem for transport infrastructure and operation” which exposes the data and maps dashboard calls with the corresponding algorithm. Also, the DOTS algorithm presented in section 2.3 is related to the equipment of Task 3.5 - “Safe MicroVehicles and VRUs: Micro-mobility Safety Solutions”.

The relation to WP4 is via the application of the developed T2.2 tools to each of the pilots. Hence the link is with the following tasks: T4.2 - Spanish pilot, T4.3 - Italian pilot, T4.4 - Latvian pilot and T4.5 - Romanian pilot.

Lastly, D2.1 is related to the following deliverables: the discussion on the tools continues in D2.3 - “Advanced infrastructure monitoring and predictive maintenance tools”, the integrated platform appears on D1.3 and D1.4 - “KPIs quantification methodologies, data space, user interfaces and integrated platform” and the comments from the stakeholders will appear in WP4 deliverables.

1.4 DELIVERABLE OVERVIEW AND REPORT STRUCTURE

The document is organised per contribution of each partner responsible for each of the outcomes. All the AI tools are presented in section 2 as subsections written by the partner who developed it. Then, section 3 discusses the applicability of AI tools in the pilots and how these tools will be integrated to the EvoRoad’s platform.

1.5 AI ACT COMPLIANCE

To guarantee the responsible development and use of AI, the European Union has defined the concept of Trustworthy AI based on three fundamental pillars: compliance with existing laws and regulations (Lawful), adherence to fundamental human values (Ethical) and technical security and resilience against risks (Robust). As outlined in the **Ethics Guidelines**

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 Trustworthy AI, these principles have been translated into practical requirements, such as transparency, human oversight, data protection, and fairness.

The **AI Act** is a European regulation that integrates the Ethics Guidelines for Trustworthy AI into a binding regulatory framework, establishing concrete rules, classifying AI-based systems, and imposing stringent requirements for high-risk applications. The regulation introduces a risk-based classification system for AI-based systems, dividing them into four main categories:

- *Unacceptable risk*: Prohibited systems, such as those employing subliminal manipulative techniques or social scoring;
- *High risk*: Applications impacting critical sectors (e.g., healthcare, infrastructure, mobility) that are subject to strict compliance requirements;
- *Limited risk*: Systems that require transparency obligations, such as chatbots or deepfakes;
- *Minimal risk*: Unregulated systems, such as spam filters or AI-powered video games.

The AI Act seeks to balance technological innovation with the protection of fundamental rights, promoting the responsible use of AI. It includes human oversight, data security, and non-discrimination provisions while supporting AI-based research and development in strategic sectors.

EvoRoads Project is committed to fully complying with the EU AI Act, its core principles, and the Trustworthy AI criteria, ensuring that all developed AI-based technologies are lawful, ethical, robust, and respectful of human rights and values.

The EvoRoads project involves multiple entities, meaning any subject (company, individual, or public body) engaged in developing, distributing, or using an AI-based system. Depending on their role, each entity has specific obligations and responsibilities to ensure compliance with European AI regulations. The key entities involved in the EvoRoads Project include:

- **Provider**: A company, public entity, or individual that develops an AI-based system or a general-purpose AI model and places it on the market or into service under its own name or brand;
- **Deployer**: Any entity using an AI system under its authority (e.g., companies, public bodies), except for personal, non-professional use.

To correctly classify the role of each Evoroads beneficiary and the characteristics of the AI-based systems being developed, used, or distributed, the project will leverage the **EU AI Act Compliance Checker.**

By answering a series of questions, this tool will support EvoRoads beneficiaries in identifying the risk level of EvoRoads AI-based system and being informed about specific articles of the AI Act to be considered to be fully compliant with trustworthy AI principles. The results provided by the tool will be documented in project deliverables dedicated to the EvoRoads AI-based solutions.

By leveraging the EU AI Act Compliance Checker, the EvoRoads project will ensure full compliance with European AI regulations, fostering the development and adoption of safe, ethical, and reliable AI-based systems in road transport.

More details on the answers of the presented tools to the EU AI Act Compliance Checker can be found in section 3.2.

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 AI TOOLS

2.1 ROAD ACCIDENT RISK QUANTIFICATION USING HISTORICAL DATA [INLECOM]

2.1.1 INTRODUCTION

This section outlines the technical methodology and implementation of a Machine Learning (ML) pipeline developed under the EvoRoads project. The goal of this work is to design an AI tool that assesses and classifies road segments based on accident risk and its evolution over time, supporting evidence-based decision-making for road safety interventions. The AI tool offers a range of functionalities to support road safety analysis and decision-making across different spatial levels, including municipalities, administrative extensions, and grouped road segments. These functionalities allow users to extract insights and prioritize interventions flexibly, depending on their specific geographic focus or preferred grouping. For example, three use cases exemplifying the tool's functionality are:

- Prioritization of High-Risk Road Segments per Municipality: *identifying which road segments should be prioritized for safety interventions. This is particularly relevant in contexts where budgets are limited and must be allocated efficiently to maximize the impact on reducing road accidents and improving public safety.*
- Identifying High-Need Municipalities for Budget Allocation: *stakeholders responsible for regional or municipal budget planning may wish to prioritize funding for communities with the most urgent road safety needs. To support such decisions, the tool offers multiple perspectives on risk (i.e. static, dynamic, and combination of the two) across municipalities.*
- Risk Overview by Road Group Across the Regional Network: *allows a broad visibility into accident patterns across the entire road network. Specifically, users may wish to review accident statistics for road segments that belong to large road groups.*

For more details on these use cases, please refer to section 3.3.3 of D2.2 as D2.1 is focused on the tool's development.

Accordingly, this section is structured as follows. Section 2.2. describes the available dataset describing the accidents in Piedmont region of Italy over seven years. Section 2.3. overviews the preprocessing steps undertaken in order to provide a comprehensive and complete set of data to the AI tool. Section 2.4 navigates the reader to the development of the tool which comprises three functionalities: risk clustering, trend estimation and analysis, and trend clustering. Finally, Section 2.5 summarises the results and provides a demonstration of this AI tool for road accident prediction using historical data.

2.1.2 DATASETS

The analysis is based on a comprehensive dataset provided by project partner CSI, detailing road accidents that occurred in Piedmont between 2016 and 2022. The dataset comprises 3,516,774 entries across 68 columns, capturing various variables including:

- Geolocation of accidents
- Number of injuries and fatalities
- Vehicle flow characteristics
- Aggregate accident counts per road segment

The data is available on a yearly basis, with each entry describing the total number of accidents occurring on a given road segment during the corresponding year. Through a combination of spatial aggregation, risk clustering, and trend analysis,

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

a tool that allows stakeholders to interpret historical accident patterns and monitor how road segment risk evolves over time, was developed.

2.1.3 PREPROCESSING ANALYSIS

The preprocessing analysis briefly introduces the data transformation and engineering performed, which will be extensively analysed in the following deliverables of this project, while also it highlights the process of spatial aggregation.

2.1.3.1 DATA ANALYSIS AND ENGINEERING

The initial phase of data preprocessing involves analysing the data and making the necessary transformation of the raw accident dataset provided by CSI Piemonte. This process includes tasks such as parsing and structuring accident data across multiple years, unifying geographic identifiers, and engineering features relevant for clustering and trend detection. This process will be elaborated further in D2.3 - "Advanced infrastructure monitoring and predictive maintenance tools".

2.1.3.2 SPATIAL AGGREGATION: THICKENED UIDS

In the original dataset, each accident is linked to a road segment (identified via the uid column), represented by a LineString geometry, i.e., a series of GPS points. However, to support reliable risk analysis, the raw accident data needs to be spatially aggregated in a way that captures relevant neighbouring incidents to account for junctions, parallel lanes, and interconnecting segments.

To do so, two main approaches for spatial aggregation were evaluated, both implemented using spatial indexing and buffering techniques. Each approach aims to quantify the accident risk around a given road segment (uid) by also considering its spatial neighbours.

Approach 1: Circular Buffer Around the Segment's Centroid

This method uses a fixed-radius circular buffer around the centroid of each road segment geometry. The centroid is a single point that summarises the spatial location of the segment. The average length of a road segment is 197,061 meters and therefore the radius was set to 100 meters to cover the whole area in which the road segment spans. The steps are as follows:

1. For each uid, a circular buffer is created around its centroid using buffer (100).
2. A spatial index (sindex) is used to efficiently find potential overlapping segments within the buffer's bounding box.
3. The number of neighbouring segments and the total number of accidents (including the central one) and related information (e.g. number of injured/dead, number of accidents with injuries/deaths) are computed and stored.

This approach simplifies the geometry to a point and focuses on identifying segments located within a specific radius from that point. It is computationally efficient and works well in open, less-dense road networks. However, it may miss longer but adjacent segments, especially in curved or complex geometries, and is sensitive to the placement of the centroid.

Approach 2: Fixed Buffer Around the Entire LineString Geometry

In this technique, a buffer is applied directly to the entire geometry of the road segment (i.e., its LineString shape) rather than just its centroid. Specifically:

1. Each road segment is buffered by a fixed distance, producing a polygon that encapsulates the entire segment with a certain spatial margin. The distance parameter was finetuned so that the extended road segment includes adjacent road segments while minimising the chance of aggregating segments that are not related with the given road segment. Specifically, this was set to 20 meters which produces an average number of neighbours for each road segment of value 2.
2. A bounding box is created for this polygon, and sindex is used to retrieve candidate neighbouring segments.

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. These candidates are filtered based on the actual distance between geometries — only segments within 20 metres are retained.
4. As with the previous method, the number of neighbours and cumulative accident statistics are computed.

This approach captures neighbouring segments more comprehensively, especially in dense or interconnected urban areas where roads may overlap or run parallel. By working with the full geometry, it better reflects physical proximity between segments.

Methodological Choice: Thickened LineStrings

After comparing both methods, the second approach was selected (e.g. applying a buffer around the full LineString geometry) for the following reasons:

- It reflects the full extent of a road segment, not just a central point, making it more appropriate for identifying real-world spatial proximity.
- It is robust in dense urban environments where centroids may fall close together despite segments not being physically adjacent.

The result of this process is a "thickened UID", which is a road segment enriched with information about nearby segments and their associated accident data. This aggregation forms the basis for subsequent clustering and trend analysis, ensuring that the tool captures not just isolated risks, but also local accident hotspots along contiguous road infrastructure.

2.1.4 TOOL DEVELOPMENT

This subsection describes the process of development of the AI tool. Specifically, it describes the three steps of the development, each of which resulted in an additional functionality. The three functionalities are risk clustering, trend estimation, and trend clustering, which independently and in combination increase the decision-making leverage of the infrastructure maintenance officials in Piedmont.

2.1.4.1 RISK CLUSTERING

The objective of this functionality is to categorise each road segment (uid) into one of four risk levels based on historical accident data aggregated over the period 2016–2022. This enables stakeholders to distinguish between road segments with varying degrees of accident likelihood and severity, and to prioritise interventions accordingly.

Three clustering approaches of increasing complexity were assessed:

1. Quantile-based binning (baseline)
2. K-Means clustering
3. Gaussian Mixture Model (GMM) clustering

Each technique segments the road network into four risk categories: 0 – No/Very Low Risk, 1 – Low Risk, 2 – Intermediate Risk, and 3 – High Risk.

To compute a meaningful and stable risk score, aggregated accident data were used, over all seven years for each thickened uid. For the first approach, the number of accidents in the thickened road segment was considered. For the two later approaches, the total number of accidents was used, along with the total number of accidents with injuries, the total number of accidents with fatalities, the total number of injured individuals and the total number of deaths, all for the thickened road segment.

The three techniques operate as described below.

Quantile-based Binning

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 method divides the distribution of total accidents into four equally sized groups (quartiles). Each segment is assigned to a cluster depending on which quartile its value falls into. Mathematically, this is a form of univariate discretisation defined as

$$risk(x) = \begin{cases} 0 if x < Q_1 \\ 1 if Q_1 \leq x < Q_2 \\ 2 if Q_2 \leq x < Q_3 \\ 3 if x \geq Q_3 \end{cases}$$,where Q_1, Q_2 ,Q_3 are the first three quartiles of the distribution.

This technique is intuitive and easy to implement but has significant limitations. These limitations are that it considers only one variable (e.g. accidents in the thickened uid), it assumes an even distribution of values, which may not hold in skewed data, and it is sensitive to outliers.

Graphically, the clusters produced using this approach for all the road segments, and the ones having at least one accident can be found below in turn.

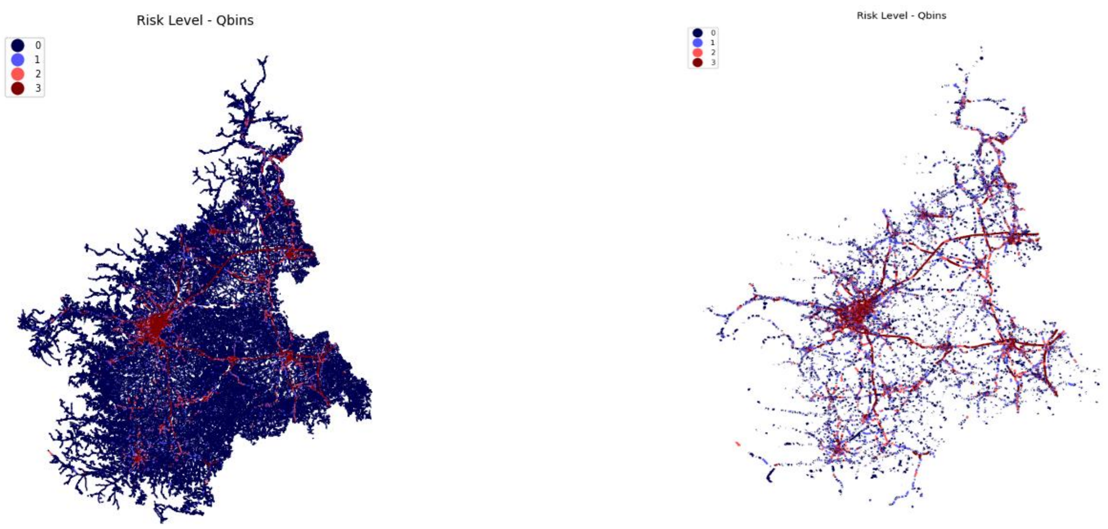

Figure 1: Risk Clusters assigned with Quantiles-based Binning with all the road segments (left) and the ones having accidents (right)

<u>K-Means Clustering</u>

K-Means is a distance-based clustering algorithm that partitions the data into clusters by minimising the within-cluster sum of squared distances. The algorithm operates as follows:

1. Randomly initialise centroids.
2. Assign each data point to the nearest centroid (using Euclidean distance).
3. Update centroids as the mean of points assigned to each cluster.
4. Repeat until convergence.

Mathematically, the algorithm minimises $\sum_{i=1}^{k} \sum_{x_j \in C_i} \parallel x_j - \mu_i \parallel^2$ where μ_i is the centroid of cluster C_i, and x_j are the feature vectors.

K-Means considers all five features simultaneously and is computationally efficient. However, it assumes that clusters are spherical in shape (due to Euclidean distance), equal in size and density, and linearly separable in the feature space. These assumptions may not hold in real-world accident data, which can have overlapping or skewed distributions.

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 clusters produced using K-Means Clustering for all the road segments, and the ones having at least one accident can be found below in turn.

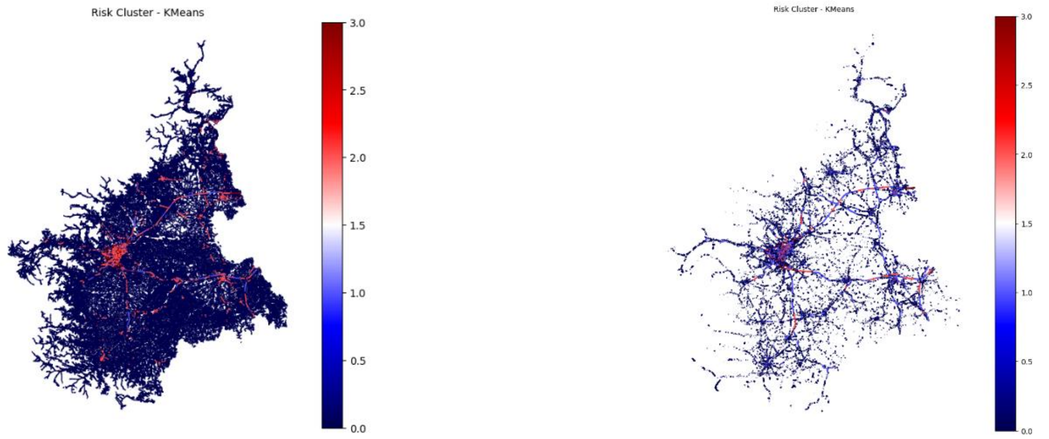

Figure 2: Risk Clusters assigned with K-Means Clustering with all the road segments (left) and the ones having accidents (right)

<u>Gaussian Mixture Models (GMM)</u>

GMM is a probabilistic clustering algorithm that models the data as a mixture of multiple multivariate Gaussian distributions. Unlike K-Means, which assigns each point to a single cluster, GMM estimates the *probability* that each point belongs to each cluster and then assigns it to the most likely one.

Formally, GMM assumes $p(x) = \sum_{i=1}^{k} \pi_i N(x \vee \mu_i, \Sigma_i)$ where π_i is the mixing coefficient of the *i-th* Gaussian, μ_i, Σ_i are the mean and covariance of the *i-th* Gaussian, and $N(x \vee \mu_i, \Sigma_i)$ is the multivariate normal distribution.

GMM is trained using the Expectation-Maximisation (EM) algorithm, which iteratively:

1. Estimates cluster membership probabilities (E-step)

2. Updates the parameters of the Gaussians (M-step)

Here the resulting clusters for all the road segments and the ones having at least one accident are:

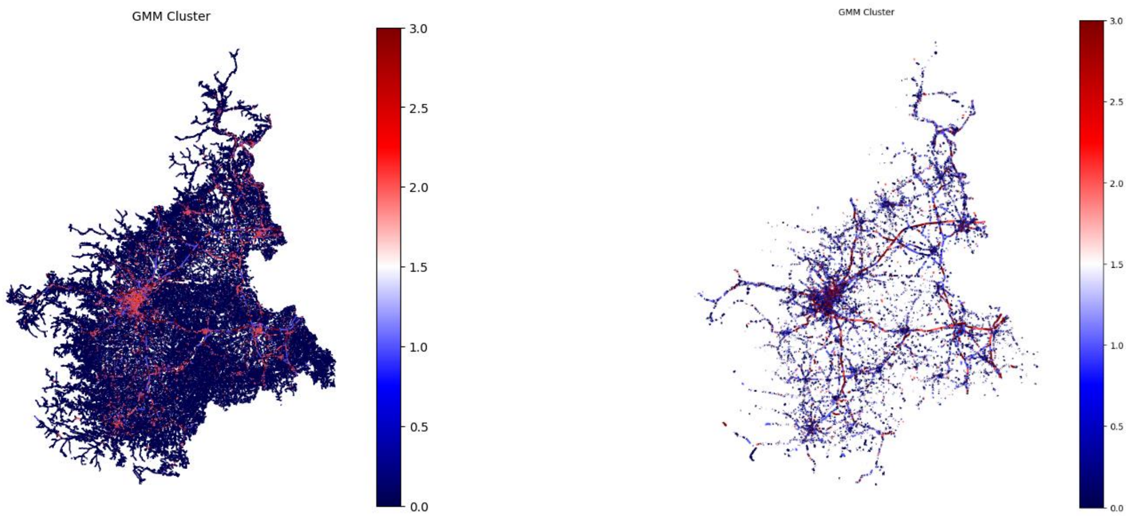

Figure 3: Risk Clusters assigned with GMM Clustering with all the road segments (left) and the ones having accidents (right)

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

Methodological Choice: GMM

After evaluating the results of all three methods, GMM clustering was selected for the final risk classification. This decision is motivated by both empirical and theoretical considerations:

- Smooth boundaries: Unlike K-Means, GMM does not impose hard cluster borders based solely on distance. It allows for overlapping clusters, better capturing the continuum of risk.
- Flexibility in shape: GMM accommodates ellipsoidal cluster shapes and varied densities, which is crucial for modelling complex real-world risk factors.
- Probabilistic assignment: This adds interpretability, as each segment's risk profile is not strictly binary but can be viewed in terms of likelihood, supporting more nuanced decision-making.

In practice, GMM produced more coherent and spatially contiguous clusters across the road network, aligning better with observed accident hotspots and expert intuition.

2.1.4.2 TREND ESTIMATION

In addition to clustering road segments by accident risk aggregated across all years, the functionality of trend estimation is developed, which adds a crucial temporal dimension to the risk assessment. While clustering reveals the cumulative severity of risk, trend estimation and analysis enables the stakeholders of this project to determine whether accident rates in a given road segment are increasing, decreasing, or stable. This distinction is essential for proactive planning, resource allocation, and the prioritisation of interventions.

This dual approach was directly informed by consultations with stakeholders at CSI Piemonte, who emphasised the value of understanding not only the current level of risk but also its trajectory over time. Combining both perspectives offers a richer understanding of road safety dynamics and will be further developed in Deliverable D2.2.

Initially, we were uncertain whether the most appropriate analytical goal was to quantify trends and acceleration in historical accident data or to attempt a forecast of future accident counts. To explore this, we implemented and tested a range of time-series techniques, namely polynomial regression, smoothing methods, time-series decomposition, and exponential forecasting. However, due to the limited number of yearly data points per road segment (a maximum of seven years, often fewer), we determined that forecasting future values would not be robust. Hence, we focused on characterising the historical trend in accident counts. Below, we summarise the steps of the whole process, and ultimately we describe and justify the approach selected.

2.1.4.2.1 DATA PREPARATION

To ensure robust and meaningful trend estimates, several preprocessing decisions are required:

i) Data from 2020 (COVID-19 impact)

The year 2020 posed challenges due to pandemic-related traffic restrictions. However, we decided to retain data from 2020 because firstly lockdowns did not span the entire year, with significant portions showing normal traffic conditions, and secondly manual inspection revealed no drastic deviation in accident numbers compared to adjacent years.

ii) Handling missing values

Three types of uid time series exist in the dataset:

1. Complete data for all years (2016–2022)
2. Late-starting series (for instance from 2019 onwards)
3. Series with missing values between years

To preserve data integrity and avoid distorting the trend signal, we exclude uids with gaps between years, resulting in a loss of 3.742% of uids.

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

iii) Minimum data points for trend estimation

We need a minimum number of data points to ensure stability. After testing, we set a threshold of 3 years per uid. This balances statistical validity with data coverage: 2 points are too sensitive to noise, while >3 would exclude too many segments. This step resulted in a further 3.743% reduction in valid uids.

iv) Final dataset coverage

After filtering we obtain trend values for 511,165 uids, which correspond to 257,842 year-wise rows. This captures 99.692% of all uids with at least one recorded accident.

2.1.4.2.2 EVALUATED TECHNIQUES

To assess how accident counts evolve over time across different road segments, we explore several time series techniques, each with different assumptions and mathematical formulations. These methods are tested on the same dataset of road segments with accident counts between 2016 and 2022, with the goal of selecting the most reliable and interpretable estimate of trend under data sparsity.

Polynomial Trend (Second-Degree Fit)

The first method we test involves fitting a second-degree polynomial to each time series of accident counts. The mathematical form of this model is $y = a + bx + cx^2$, where x denotes the centred year variable (for numerical stability), and y is the recorded number of accidents. The linear coefficient b represents the primary trend direction (indicating whether accidents are increasing or decreasing) while the quadratic term c captures any acceleration or deceleration in that trend.

Moving Average Slope

The second approach is based on a simple moving average. Here, we compute a centred rolling mean over a three-year window and use the difference between the smoothed values at the final and initial years, divided by the number of time steps, as an estimate of the trend.

STL Decomposition

The third technique uses STL (Seasonal-Trend decomposition via Loess), a robust and widely used method for separating a time series into trend, seasonal, and residual components. STL applies locally weighted regression to iteratively extract the trend, offering flexibility and robustness to outliers. In our case, we used the estimated trend component to compute the slope between the start and end of the series.

Holt's Linear Trend Method

The fourth method is Holt's exponential smoothing technique, which assumes that the data follow a linear trend and updates both the level and slope of the series recursively over time. It estimates two parameters: one for the level (intercept) and another for the trend (slope), which are updated in each time step using smoothing coefficients.

Linear Regression Trend (Least Squares Fit)

The final approach evaluated is a simple linear regression using least squares. It is similar to the first but here we consider a polyonymy of first order, so we fit a line of the form $y = a + bx$ directly to the accident counts for each road segment across years.

Representative examples for each method

Below we include representative examples of trend estimation results using each technique on selected road segments. These plots demonstrate the advantages and limitations of each method.

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

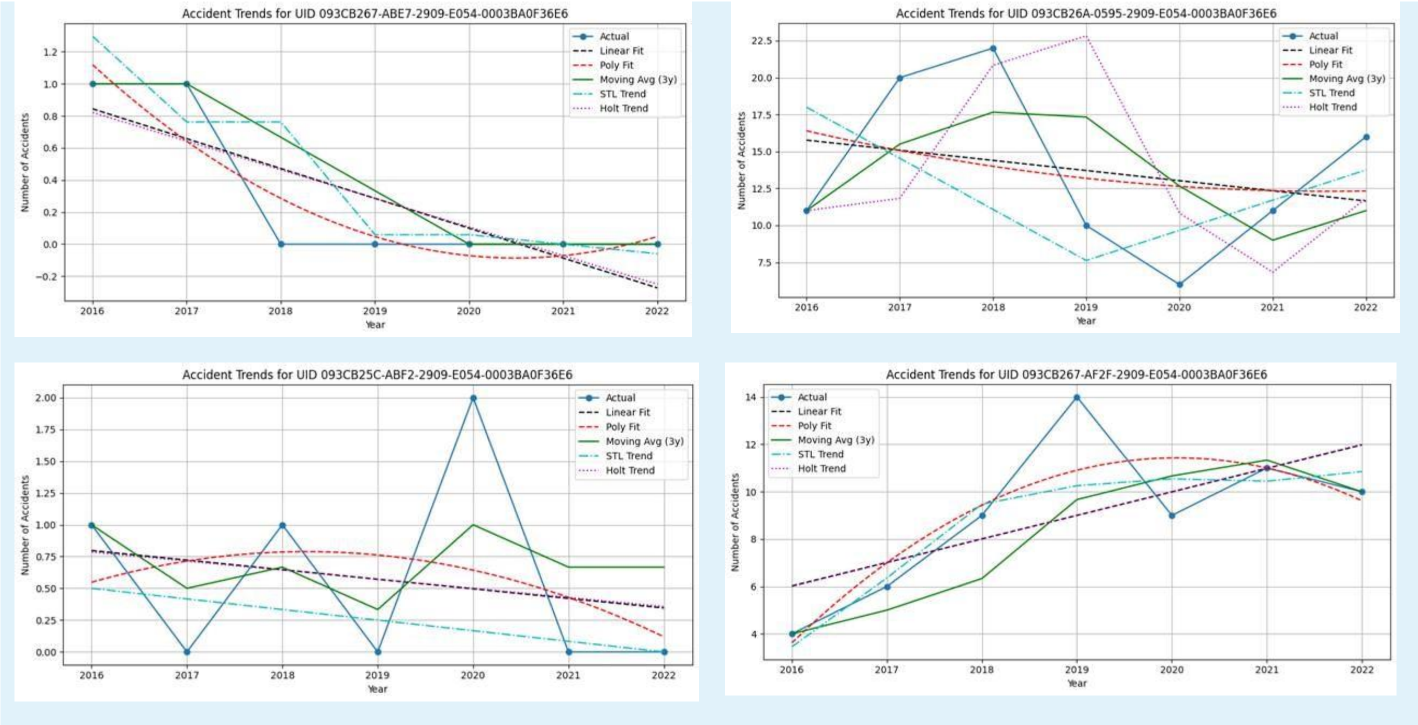

Figure 4: Representative examples of fitting to the trend

Final Selection: Linear Trend

To identify trends in accident frequency over time, we evaluated several candidate techniques, each with distinct methodological characteristics (see Figure 4 for some representative examples). Non-linear approaches such as polynomial regression and STL decomposition can capture complex patterns and smooth trends, but both are sensitive to noise and require a higher temporal resolution than our data offers. With only 3–7 annual observations per road segment, these methods are prone to overfitting and produce unstable estimates. Smoother alternatives like the moving average slope and Holt's linear method offer noise reduction and adaptability but come with limitations: the moving average lacks explicit temporal modelling and performs poorly near data boundaries, while Holt's method showed convergence issues and instability on flat or irregular time series. In contrast, linear regression offers a statistically sound, robust, and interpretable solution. It provides a clear estimate of the trend through the slope coefficient and is well suited to small datasets. From a statistical perspective, it minimises residual error efficiently and avoids overfitting. Based on these considerations—and in consultation with stakeholders prioritising clarity and stability—we selected linear regression as the primary technique for trend estimation, balancing scientific validity with operational relevance.

2.1.4.3 TREND CLUSTERING

After estimating the accident trend for each road segment, the third functionality of the tool involves classifying the estimated trends into meaningful categories to support interpretability and downstream analysis. The goal here is to distinguish between segments with stable accident rates (trend ≈ 0), increasing trends (positive trend values, ranging from mild to steep), and decreasing trends (negative trend values, similarly ranging in intensity).

Evaluated Techniques: To cluster data based on the trend, we assess three different clustering techniques based on the trend value computed per uid: quantile-based binning (qcut), Gaussian Mixture Models (GMM), and a fixed-threshold categorisation.

To understand the distribution of trends across all road segments, we first visualise the trend values and overlay their quantile boundaries. This allows for evaluating the skewness and concentration of values before applying any grouping method.

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

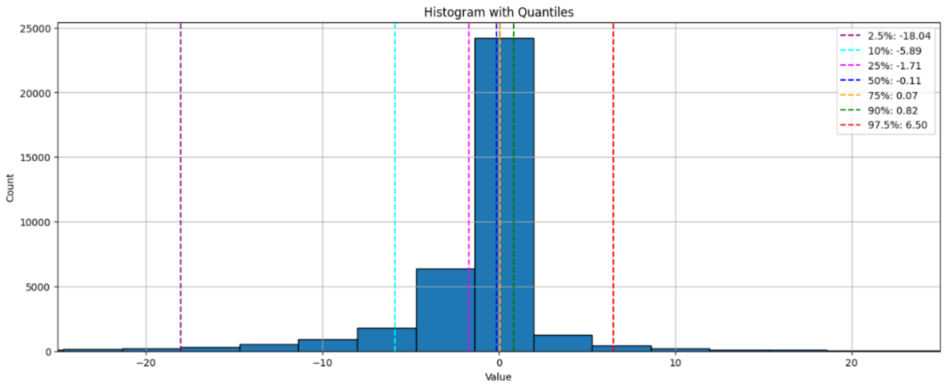

Figure 5: Histogram depicting the distribution of trends with quantiles

For the first two methods (qcut and GMM) we opt for 10 clusters each. This granularity is chosen to provide sufficient resolution to stakeholders, enabling them to identify subtle variations in risk evolution. Quantile binning assigns roughly equal numbers of segments to each category, while GMM, described in a previous section, uses probabilistic modelling to fit the underlying distribution and assign clusters based on likelihood. The mathematical background for both methods is already introduced in the risk clustering part of this section, so they will not be analysed further here. In addition, we implement a fixed-threshold classification to produce six interpretable trend categories, defined by manually specified breakpoints, which are i) strongly decreasing (trend < -0.5) – labelled as 0, ii) mildly decreasing (-0.5 ≤ trend < 0) – labelled as 2, iii) stable (trend = 0) – labelled as 4, iv) mildly increasing (0 < trend ≤ 0.5) – labelled as 6, v) strongly increasing (trend > 0.5) – labelled as 9, and vi) unclassified (NaN, for segments with insufficient data) – no label assigned just a NaN value. The results for the three techniques are shown in the graphs below in turn:

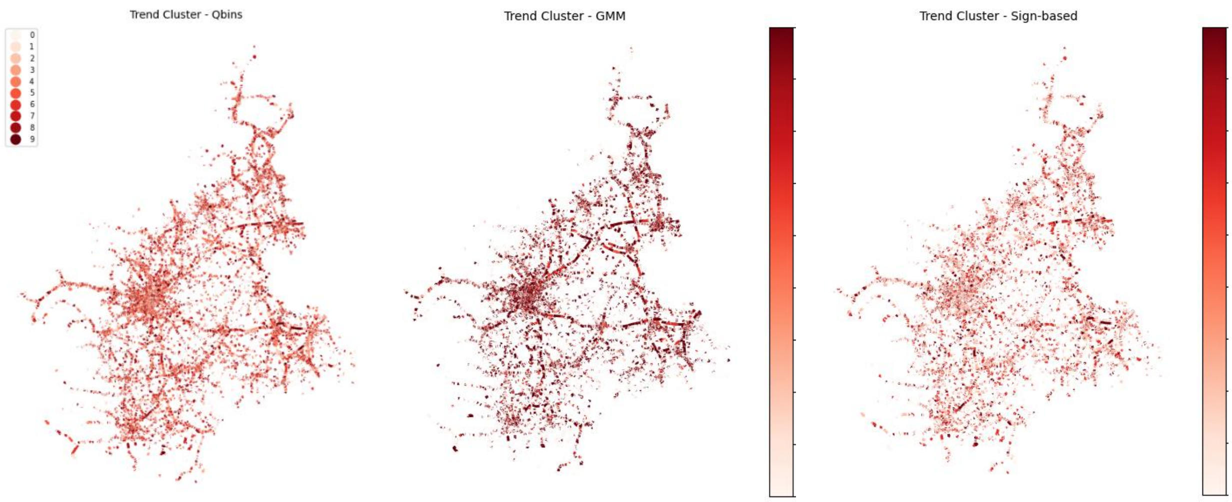

Figure 6: Trend Clusters with three techniques (left: quantiles-based binning, center: gaussian mixture models, right: sign-based clustering

Final Selection: GMM

The Qbins method produces a relatively uniform spatial distribution, but this regularity also results in a somewhat artificial segmentation of trends. Since Qbins imposes equally sized quantiles without considering the underlying distributional shape of the data, it may mask subtle but meaningful variations between accident trends, especially around critical thresholds of change. Visually, while the map offers some variation, it does not prioritise continuity or similarity between nearby segments.

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 Sign-based clustering results in a rather noisy and less discriminative pattern across the road network. Because the threshold categories (e.g., stable, mildly increasing, strongly increasing) are based on fixed arbitrary cut-offs (like ±0.5), the result is overly coarse. Furthermore, the number of segments falling in each category is imbalanced, with certain regions looking either over- or under-represented. This can reduce the interpretability and usefulness of the classification in practice, particularly for stakeholders seeking to prioritise interventions.

The GMM-based clustering, by contrast, delivers a smoother and more coherent spatial distribution of trend classes. This is particularly evident in the continuity of colours across contiguous road segments, suggesting that GMM captures underlying statistical patterns in how accident trends evolve spatially. GMM has the advantage of fitting the distribution of the data, rather than slicing it mechanically. The probabilistic nature of GMM allows for overlapping trend regions to be separated based on likelihood, which leads to a more natural and nuanced classification. It avoids the rigidity of Qbins while offering more granularity than the sign-based approach.

For these reasons, we chose to proceed with the GMM-based clustering of accident trends. It combines smoothness, granularity, and statistical validity, offering stakeholders a more interpretable and actionable overview of how accident patterns are evolving over time across the network.

2.1.5 CONCLUSIONS

In this section, we presented a pipeline of AI-based tools for the monitoring and analysis of cyber-physical road infrastructure, aimed at supporting proactive risk management. The approach included data preprocessing, focusing on the grouping of neighbouring road segments. This was followed by the description of the development of the AI tool. In more details, this outlined the three functionalities by describing the process of clustering of accident risk using historical data, estimating the trend of accident occurrences, and clustering road segments based on these trends. Together, these steps enable the identification of not only currently high-risk areas but also those where risk is increasing over time, thus offering a comprehensive view for strategic intervention. The upcoming Deliverable D2.2 will expand on how these analytical insights can be effectively integrated into stakeholder decision-making processes and operational planning. A demo of the tool’s functionalities can be found here: **Demo**.

2.2 ROAD DAMAGE DETECTION & CLASSIFICATION [UCY]

The solution described in this section deals with recognising crack damages on the road based on deep learning models. The presented solution covers aspects of data generation, crack segmentation, and crack localisation all of which can be used complimentary to provide a holistic approach to road damage detection, segmentation and classification.

2.2.1 APPLICATION

Contributions of the application to Objective 2:

The developed model contributes to the detection of infrastructure deterioration at varying image resolutions. The model can automatically detect cracks or potholes, significant indicators of structural degradation, from high-resolution images obtained by UAVs (Unmanned Aerial Vehicle). UAVs provide adaptable spatial coverage, allowing for precise and localized inspections of infrastructure, such as urban highways. These aerial evaluations can effectively complement data from other sources, including satellite imagery and vehicle-mounted video feeds, enhancing the completeness of monitoring efforts.

In terms of data integration, the model—though currently focused on UAV imagery—has interoperability within its scope. This makes it suitable for integration into broader federated platforms that combine multiple heterogeneous data streams. By incorporating inputs from sources like onboard vehicle cameras used for continuous surveillance, the system can support the integration of data across various modalities and scales. This enhances the overall accuracy and spatial-temporal coverage of infrastructure monitoring systems.

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 application also enables proactive hazard warnings. With real-time crack detection capabilities during UAV surveillance missions, the system can trigger early alerts for high-risk areas. Such capability is crucial for improving situational awareness in traffic management systems, particularly when UAVs are deployed in emergency scenarios or during post-disaster infrastructure assessments.

Lastly, the tool contributes to optimising maintenance operations. By supporting systematic and repeatable inspections, it allows infrastructure authorities to prioritise the maintenance of road infrastructure based on the severity of the defects, and the observed progression of deterioration over time. This facilitates more effective and timely decision-making infrastructure management.

Contribution to EI2.2: AI Algorithms for Novel Defect Identification:

The core of the application is a deep learning-based algorithm that was designed to automatically detect defects using UAV-collected imagery. This approach minimises the dependency on visual inspections, thereby reducing operational costs and the risk of human error. The model identifies defects reliably as it performs consistently in diverse environmental conditions and various road types. It has moreover the flexibility to be retrained or fine-tuned for detecting other types of defects beyond cracks, such as surface erosion, making it a scalable and extensible solution for broader infrastructure monitoring applications.

The deployment of the object detection model enables more rapid, more consistent, and more scalable infrastructure monitoring. It reduces operational costs through automation, enhances accuracy and consistency in crack detection, and improves decision-making with reliable data for maintenance planning. The solution can be employed on various platforms. For the pilot application (Romanian Pilot), we created a small build recipe that packages everything needed to run the object-detection model in a clean, repeatable way (using Dockerfile). Overall, the model can assess the road condition efficiently, cost-effectively, and in a data-driven way. Nevertheless, the practical deployment of UAV-based inspections is subject to operational limitations, such as flight height restrictions and battery endurance, which may limit continuous coverage of extended road networks.

2.2.2 DATASETS

To support the training of deep learning models for road damage detection, classification and segmentation, two distinct datasets were developed for this purpose. The first dataset consists of artificially generated (synthetic) images. It was created by overlaying crack damage sourced from publicly available datasets onto road background images that originally contained no natural road damage. This synthetic dataset was specifically used for training crack segmentation models. The second dataset comprises annotations (bounding boxes) of cracked roads in Cyprus (as gathered by UCY) and was used for object detection tasks. It includes four damage categories: horizontal crack, vertical crack, alligator crack, and pothole. Due to the limited number of images in the original dataset, the final object detection dataset was augmented by incorporating images from the publicly available RDD2022[1] dataset (specifically the China Drone subset of the dataset). This resulted in a combined dataset consisting of 1,600 images for training and 400 images for validation (see Table 2).

Table 2: Class instances of the object detection dataset.

Object Detection Dataset	Images	Horizontal Cracks	Vertical Cracks	Alligator Cracks	Potholes
Train Set	1628	1854	1910	248	81
Validation Set	403	479	432	21	68

[1] D. Arya, H. Maeda, S. K. Ghosh, D. Toshniwal, and Y. Sekimoto, “RDD2022: A multi-national image dataset for automatic road damage detection,” Geoscience Data Journal, Jun. 2024, doi: 10.1002/gdj3.260.

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 PREPROCESSING ANALYSIS

For the object detection dataset, crack damage annotations were created using the Computer Vision Annotation Tool (CVAT[2]) annotation tool (Figure 7). In contrast, the synthetic crack segmentation dataset was generated through a custom image synthesis pipeline. This pipeline overlays externally sourced crack damage onto road background images captured by an UAV. The process of generating these synthetic images is as follows:

Initially, top-down realistic road infrastructure images were captured by a UAV (with an altitude of 120m) to provide a comprehensive view of the surface. Along with these, images of cracks and their corresponding binary masks were sourced from existing annotated datasets, though these originated from different contexts. The binary masks were used to isolate the cracks, allowing their integration onto our own road images. To ensure diversity in crack damages we combined multiple crack images sourced from different publicly available datasets such as Crack500[3], DeepCrack[4], CrackTree200[5], GAPs[6], Volker, Rissbilder[7], CFD[8], and the Crack Segmentation Dataset by Roboflow[9].

The process of overlaying cracks onto background images starts with first isolating the road segment. This ensures that cracks are applied only to the relevant road area, resulting in a more realistic outcome. A segmented mask of the road from the background images is used to guide the synthetic crack deployment. Valid positions (e.g., all (x, y) pixel points of the image that belong to the road area) for crack placement are determined to avoid unrealistic positioning, such as cracks appearing on non-road areas. The crack image is adjusted to match the colour and scale of the road, ensuring a natural integration into the background. In detail, an inverse mask is created from the original binary mask of the crack. This inverse mask is used to black out the area where the crack will be applied, preserving the rest of the region in the background image. The inverse mask is applied to the region of interest (ROI)[10] in the background, ensuring that only the background area outside the crack region remains visible, while the area where the crack will be placed is blacked out. The crack image is further isolated by applying the binary mask to it, keeping only the crack region visible and blacking out the rest. The crack region is then resized to match the size of the background ROI, ensuring consistent dimensions between the crack and the background areas. Finally, the crack region is added to the blacked-out ROI from the background, effectively overlaying the crack on the background. This combined region is then placed back into the original background image at the specified position, resulting in a seamless integration of the crack into the background (Figure 7).

---

[2] https://www.cvat.ai/

[3] L. Zhang, F. Yang, Y. D. Zhang, and Y. J. Zhu, "Road crack detection using deep convolutional neural network," *2022 IEEE International Conference on Image Processing (ICIP)*, pp. 3708–3712, Aug. 2016, doi: 10.1109/icip.2016.7533052.

[4] Y. Liu, J. Yao, X. Lu, R. Xie, and L. Li, "DeepCrack: A deep hierarchical feature learning architecture for crack segmentation," *Neurocomputing*, vol. 338, pp. 139–153, Jan. 2019, doi: 10.1016/j.neucom.2019.01.036.

[5] Q. Zou, Y. Cao, Q. Li, Q. Mao, and S. Wang, "CrackTree: Automatic crack detection from pavement images," *Pattern Recognition Letters*, vol. 33, no. 3, pp. 227–238, Nov. 2011, doi: 10.1016/j.patrec.2011.11.004.

[6] M. Eisenbach *et al.*, "How to get pavement distress detection ready for deep learning? A systematic approach," *2022 International Joint Conference on Neural Networks (IJCNN)*, pp. 2039–2047, May 2017, doi: 10.1109/ijcnn.2017.7966101.

[7] M. Pak and S. Kim, "Crack detection using fully convolutional network in Wall-Climbing robot," in *Lecture notes in electrical engineering*, 2021, pp. 267–272. doi: 10.1007/978-981-15-9343-7_36.

[8] Y. Shi, L. Cui, Z. Qi, F. Meng, and Z. Chen, "Automatic Road crack detection using random structured forests," *IEEE Transactions on Intelligent Transportation Systems*, vol. 17, no. 12, pp. 3434–3445, May 2016, doi: 10.1109/tits.2016.2552248.

[9] Roboflow Team, "Crack Instance Segmentation Dataset and Pretrained Model." *Roboflow*. https://universe.roboflow.com/ university-bswxt/crack-bphdr?ref=ultralytics.

[10] The region of interest refers to the specific portion of the background image where the crack will be placed. It is defined by the coordinates (x, y), which represent the top-left corner of the area, and the dimensions h (height) and w (width) of the crack image. By slicing the background image from y to y+h (vertical range) and from x to x+w (horizontal range), the function extracts the area of the background where the crack will be overlaid, ensuring that the crack fits within this region.

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

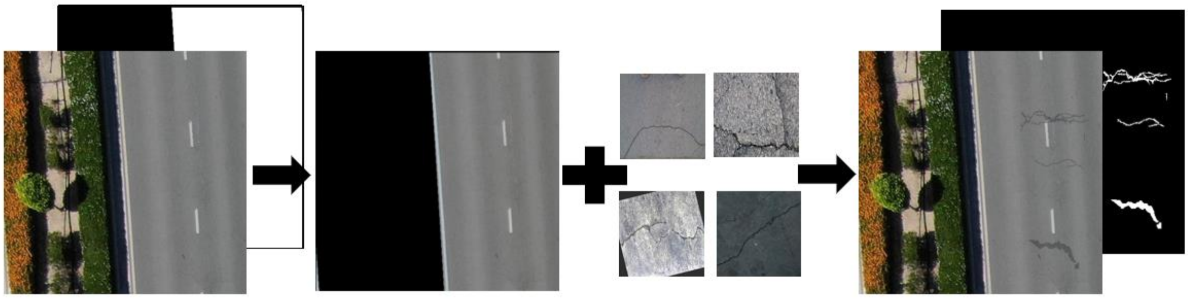

Figure 7: Process of overlaying crack damage onto road backgrounds

During this process, certain images containing vehicles on the road were identified. Since overlaying cracks on the cars would result in unrealistic damage of the road, these images were removed from the dataset so that the accuracy is maintained. Additionally, cracks were superimposed on road images where the roads had no natural cracks. The synthetic dataset can be used to develop crack segmentation models. Examples of the images can be seen in Figure 8.

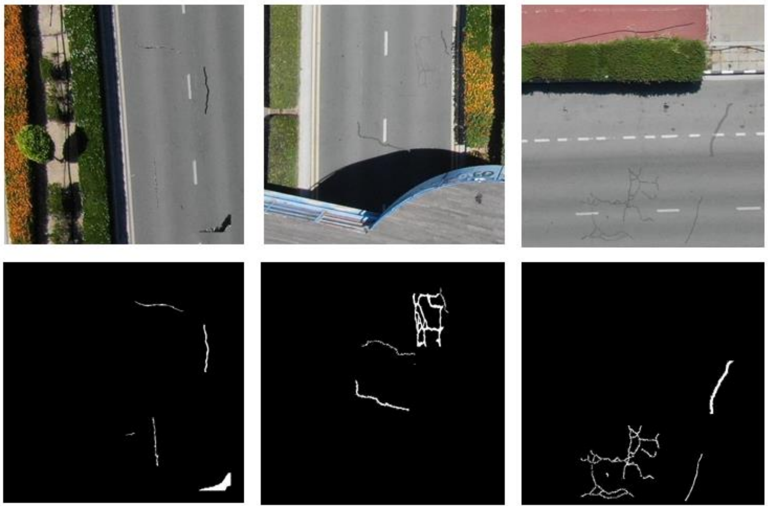

Figure 8: Generated synthetic images with their corresponding segmentation mask

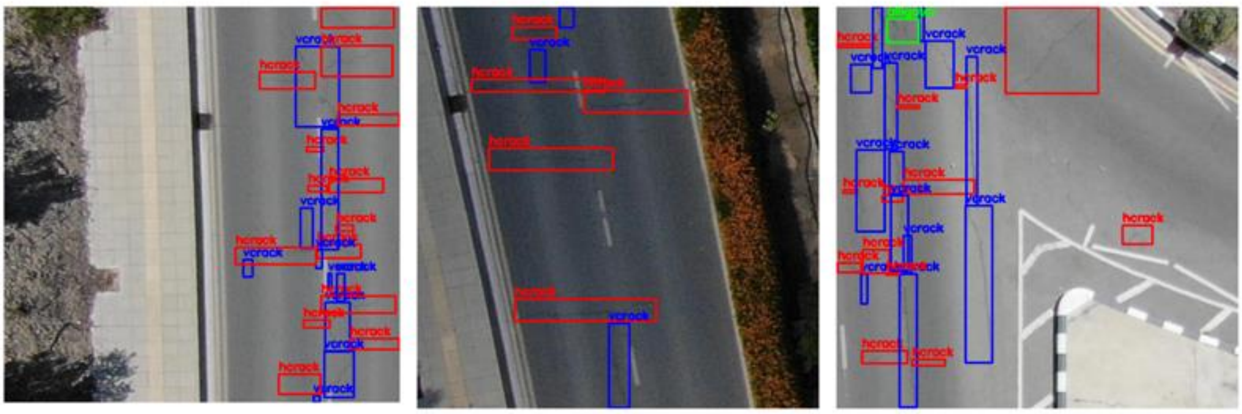

Figure 9: Images from object detection dataset with annotations as bounding boxes

2.2.4 TOOL DEVELOPMENT

The datasets described above were used to train either a segmentation model or an object detection model. For crack segmentation, a U-Net[11] architecture was employed, using either MobileNet[12] or EfficientNet[13] as the encoder backbone. To address the class imbalance issues in the dataset, various loss functions were used during training, including Weighted Binary Cross-Entropy Loss (WBCEL), Focal Loss, and Dice Loss. The trained models were subsequently evaluated on real-world road images containing crack damage to assess their performance. In parallel, object detection models based on YOLO11[14] (Figure 10) of varying sizes (YOLO11n, YOLO11s, YOLO11m, YOLO11l) were trained using the annotated detection dataset.

The best-performing segmentation model tested on real-world crack-damaged road images was the U-Net with a MobileNet encoder, trained using the WBCEL. This model achieved a crack pixel accuracy of 63.5%. In contrast, the YOLO11 models of various sizes yielded similar performance in the object detection task. However, when model size is a limiting factor, the lightweight versions of the model such as YOLO11n and YOLO11s are preferred, achieving a mean Average Precision (mAP@50) of approximately 57%.

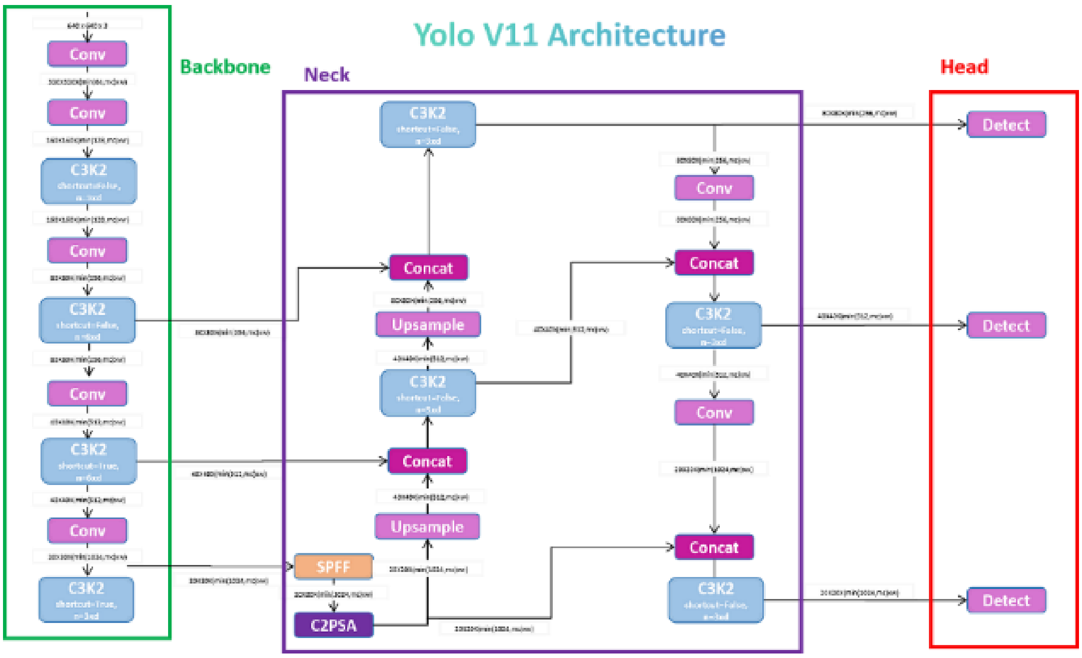

Figure 10: YOLO11 architecture

2.2.5 TOOL VALIDATION

The predictions of the best-performing segmentation model, U-Net with a MobileNet encoder trained using WBCEL, are shown below as a quantitative validation of the algorithm. The model was tested on real-world images of crack-damaged

---

[11] O. Ronneberger, P. Fischer, and T. Brox, "U-NET: Convolutional Networks for Biomedical Image Segmentation," in *Lecture notes in computer science*, 2015, pp. 234–241. doi: 10.1007/978-3-319-24574-4_28.

[12] A. Howard *et al.*, "Searching for MobileNetV3," *arXiv.org*, May 06, 2019. https://arxiv.org/abs/1905.02244

[13] M. Tan and Q. Le V., "EfficientNet: Rethinking model scaling for convolutional neural networks," *arXiv.org*, May 28, 2019. https://arxiv.org/abs/1905.11946

[14] Ultralytics, "GitHub - ultralytics/ultralytics: Ultralytics YOLO11," *GitHub*. https://github.com/ultralytics/ultralytics

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

roads. Additionally, the predictions produced by the YOLO11n and YOLO11s object detection models are also presented below in Figure 11.

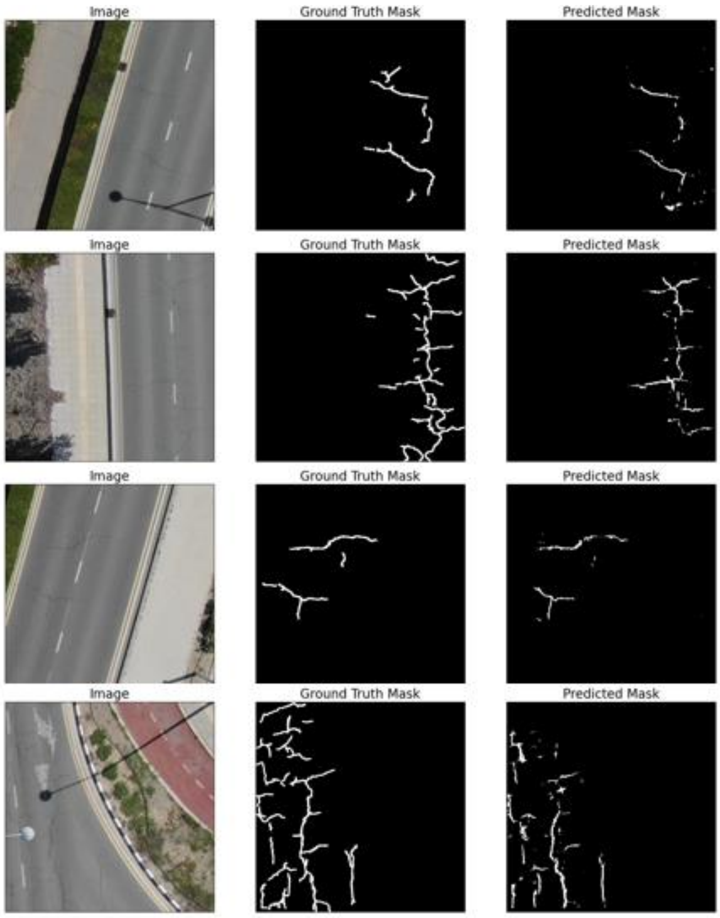

Figure 11: Predictions of UNet-MobileNet with WBCEL on real crack images

Figure 12: Predictions of the YOLO11s model

Additionally, Figure 13 presents a visual comparison of the performance metrics mAP@50 and mAP@50:95 for the YOLOv11n and YOLOv11s models. As discussed earlier, both models can achieve a mAP at IoU (Intersection over Union) threshold 0.50 of approximately 57%, indicating good object localization accuracy under a relatively moderate matching criterion. However, when evaluating the models across a stricter and more comprehensive range of IoU thresholds (from 0.50 to 0.95 in steps of 0.05), the mAP@50:95 drops to about 30%. This decrease reflects the models' reduced precision

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

under more demanding localization conditions and highlights the challenge of maintaining high accuracy across all levels of IoU overlap.

Figure 13: From left to right, a) mAP@50 b) mAP@50:95 of YOLO11n and YOLO11s trained on the crack object detection dataset

To summarise two deep learning models were developed to detect cracks on road surfaces. The first model focused on crack segmentation using a synthetic dataset, where cracks were overlaid onto background images to simulate real-world conditions. The segmentation model achieved a crack pixel accuracy of 63.5%, highlighting the potential and limitations of training synthetic data for pixel-level crack detection. The second model employed the YOLOv11 architecture for object detection, trained on a dataset annotated with four distinct crack types: horizontal cracks, vertical cracks, alligator cracks, and potholes. This model reached a mean Intersection over Union (mIoU) of 57%, demonstrating effective multi-class crack detection capabilities. Together, these models offer complementary approaches to automated crack analysis, one focusing on fine-grained segmentation and the other on broader categorical detection. The video demonstration of the YOLO11 model making predictions on images is in the following **link**.

2.3 REAL-TIME INFRASTRUCTURE DEFECTS DETECTION [DOTS]

This section presents the AI technology, employed by the sensor kit (Task 3.5) to detect potholes and monitor road surface quality for micromobility infrastructure. This innovation consists of an edge-deployed AI model that provides real-time object detection capabilities with no additional postprocessing involved. The combination of kit/algorithm is intended to be validated in the context of the Riga pilot. Additionally, the system utilises a unique dataset contributed for this purpose — the first such effort in Riga and Latvia.

The kit's goal is to ensure safer, more efficient, and cost-effective infrastructure monitoring for the benefit of both businesses and the community. Therefore, the kit is designed to be mounted on a variety of micromobility vehicles (e.g. e-scooters, bicycles, e-bikes, and e-skateboards) to continuously monitor road infrastructure quality. Special attention is given to monitoring infrastructure used by micromobility vehicles due to the increasing number of traffic accidents involving these users in Riga. The data collected and analysed will help identify hazardous areas and prioritise maintenance efforts to improve safety for this growing segment of urban travellers.

Developed according to the use case of Riga city municipality, the sensor kit follows key architectural principles detailed in deliverable D3.2 under "Sensor kit use case" and "Principles of the sensor kit architecture" sections. The technology will undergo validation in the Riga Living Lab (Latvia).

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

2.3.1 DATASETS

To address imaging quality concerns with the OAK-D-S2[15] camera in outdoor environments, a new dataset has been prepared to contribute imaging qualities specifically for the OAK-D-S2 capabilities. It provides imaging options:

- RGB central camera image,
- Stereo depth camera left or right image,
- Stereo depth camera image pair with disparity maps for depth estimation (no IR illuminator, no active stereo[16]).

Two key technical constraints governed the data collection methodology:

- Edge device inference in real-time. Training data does not rely on time-consuming 2D image analysis or processing algorithms to enhance the image quality.
- Camera image quality changes minimised. Regarding noise levels and sensitivity to variable outdoor environmental conditions.

Initial field tests using OAK-D-S2 cameras revealed significant limitations. The RGB central camera produced unacceptable motion blur even at e-scooter speeds below 6 km/h. While the stereo cameras with fixed focus provided clearer grayscale images, these images could not be combined with disparity maps for depth estimation. Further evaluation of disparity map quality identified critical limitations:

- Insufficient sensitivity to detect potholes (1-10 cm deep), as these depth differentials frequently fell within the noise threshold of the disparity estimation.
- Minor variations in camera positioning relative to the road surface resulted in systematically incorrect disparity calculations.

Based on technical evaluation and field test results, the fixed focus left stereo camera was selected as the sole data source for AI model training. This configuration provided the acceptable image clarity, processing efficiency, and implementation feasibility within the device's operational constraints.

The pothole detection dataset comprises 2,640 images, 4,696 pothole instances, representing the largest pothole dataset in Latvia collected throughout Riga's urban areas. This dataset focuses exclusively on pothole-type road damage, as one of the main reasons of Micro-mobility accidents compared to other road hazards.

Data was gathered under diverse daylight conditions, including both cloudy and sunny days, with inclusion of dark shadow patterns. The collection focused on asphalt and clay paver road surfaces. Potholes of varying sizes and depths were selected, including those that developed from longitudinal cracks, transverse cracks, alligator cracks. Different unique pothole was captured from multiple distances and angles, providing diverse perspectives as an e-scooter approaches the hazard.

The data acquisition process employed an electric unicycle equipped with a Raspberry Pi 5 and an OAK-D-S2 camera. The camera mounted at the front of the vehicle at a height of 62cm above ground level and positioned at a 60-degree angle from the vertical axis. The stereo left camera image sensor was configured to capture at 400p (640 * 400 pixels resolution), with output images scaled to 640 * 640 pixels.

[15] OAK-D-S2. 2025. Source. [Online] Available: https://shop.luxonis.com/products/oak-d-s2?variant=42455432265951

[16] OAK-D-Pro. 2025. Source. [Online] Available: https://shop.luxonis.com/products/oak-d-pro?variant=42455252369631

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

To optimise the recorded data volume, an initial object detection model was developed using the YOLOv11 [17]neural network architecture trained on publicly available pothole datasets N-RDD2024[18]. This model was deployed directly on the OAK-D-S2 device, where image capture was triggered only when pothole instances were detected by the model.

Annotations were performed using the know data annotation platform CVAT[19] with auto-labelling conducted as a preliminary step. For the auto-labelling, a pretrained semantic segmentation model[20] was used with the following specifications: UNet++ architecture featuring an EfficientNet-B5 encoder, implemented using Segmentation Models PyTorch (SMP[21]), utilizing ImageNet pretrained weights. The generated segmentation masks were subsequently converted into bounding boxes and filtered to retain only pothole class instances.

The dataset is divided for cross-validation - training (70%), validation (10%), and testing (20%). The total number of images and pothole instances as shown in Table 3.

Table 3: Dataset split for cross validation

Data subset	Images	Pothole instances
Train	1848	3249
Validation	264	471
Test	528	976

To reduce number of False Positive model pothole detections, negative samples (images without pothole instance) were included, comprising 2.95% (54 images) of the training data. These samples include dark shadows, manhole covers that can be detected as pothole because of radial shape. Additionally, out-of-distribution data samples were excluded. Samples with unique lighting conditions (over exposed or too dark), unique road surface conditions (asphalt with heavy crocodile cracks).

In this dataset, potholes were defined as portions of asphalt missing or crumbled to the point of having significant displacement in the road surface (where the inside of a pothole is lower than the surrounding road) and/or where the terrain under the road surface was clearly visible. Given their irregular form and unpredictable depth, a consistent annotation protocol was followed. The bounding box contours were created to minimise the area where surface displacement was not present, capturing only the actual depression. In cases where what appeared to be a single damaged area contained multiple distinct depressions, each was marked as a separate pothole instance with its own bounding box. For potholes with radiating cracks around, included cracks within a narrow margin of where significant displacement began, while excluding extended cracking that showed no displacement. This methodology ensured more consistent representation of pothole instances throughout the dataset.

2.3.2 PREPROCESSING ANALYSIS

The task formulation requires choosing between instance segmentation and object detection approaches for pothole identification. Due to limited availability of publicly available instance segmentation datasets for potholes, object detection

[17] G. Jocher, J. Qiu, and A. Chaurasia, "Yolov11 by ultralytics.", 2024. Source. [Online]. Available: https://github.com/ultralytics/ultralytics

[18] Ömer Kaya, 08.01.24, "N-RDD2024: Road damage and defects," Source. [Online]. Available: https://data.mendeley.com/datasets/27c8pwsd6v/3

[19] CVAT, Source. [Online]. Available: https://www.cvat.ai/

[20] Pothole-Mix Segmentation, 05.07.2022, Source. [Online]. Available: https://gitlab.com/4ndr3aR/pothole-mix-segmentation

[21] Segmentation Models PyTorch, 2020, Source. [Online] Available: https://segmentation-modelspytorch.readthedocs.io/en/latest/

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

became the preferred approach. Object detection requires bounding boxes around objects of interest rather than complicated pixel-level segmentation masks.

The architecture selection focused specifically on object detection models whose deployment environment targeted edge devices with specialised CPU inference. Edge deployment demands lightweight models that can process images in real time while maintaining acceptable detection accuracy.

Other aspects of the architecture selection involve several key criteria. First, model inference speed emerged as the primary constraint, requiring real-time inference capabilities on resource-limited hardware. Second, the ability to handle large intra-class variation, as potholes exhibit diverse appearances while sharing similar surrounding environments. Additionally, partial occlusion handling, where potholes can be obscured by water, debris, or shadows. Third, local and limited global context capturing capabilities were required for robustness against diverse weather conditions, requiring distinguishing actual potholes from surface stains and shadows.

Model inference speed optimisation involved two primary objectives. First, the neural network architecture must have scalable performance across different architecture layers (not just configurable input node) to contribute resolution agnostic capability. Therefore, model architectures with fixed fully connected layers or hard-coded feature pyramids would create constraints when retraining on smaller resolution images becomes required. Second, architecture operational complexity needs to be minimized.

Context limitations significantly influenced the architecture selection. Pothole detection primarily relies on local features such as edge boundaries, surface texture variations, and immediate surrounding context rather than comprehensive scene understanding. The surrounding environment typically consists of relatively uniform asphalt or clay surfaces with predictable lighting variations, making extensive global context analysis unnecessary. This characteristic eliminated Transformer[22] architectures from consideration, despite their superior contextual reasoning capabilities. Transformers excel at capturing long-range dependencies and global scene relationships, but these strengths become computational overhead.

By considering these objectives and limitations, YOLO[23] family architectures, specifically YOLOv11 (with YOLOv10[24] also being a choice), were chosen as suitable for baseline model development. These latest architectures employ a one-stage detection approach, directly predicting bounding boxes and class probabilities in a single forward pass without the two-stage process of region proposal generation followed by classification refinement. Additionally, these architectures feature improvements in global and local context capturing capabilities. Starting with the YOLOv10 architecture, Partial Spatial Attention (PSA) modules have enabled the identification of discriminative parts within objects by enhancing local context capturing capabilities. Furthermore, the YOLOv11 architecture introduced the C3k2 Bottleneck Structure, consisting of multiple 3×3 convolutions with skip connections and concatenations. These components contribute to extracting fine-grained intra-object features and textures essential for distinguishing between different object appearances. The Spatial Pyramid Pooling Fast (SPPF) module, present since YOLOv5[25] architecture, captures surrounding scene information at multiple scales by using pooling operations. It gathers information from larger spatial regions around objects of interest. However, the limited pooling kernel sizes and fixed receptive fields of SPPF prevent dynamic expansion of long-range dependencies, which benefits pothole detection by avoiding over-reliance on global scene context.

[22] Tahira Shehzadi, "Object Detection with Transformers: A Review" arXiv preprint arXiv: arXiv:2306.04670

[23] Ultralytics YOLO Docs, Source. [Online]. Available: https://docs.ultralytics.com

[24] A. Wang, H. Chen, "Yolov10: Real-time end-to-end object detection," arXiv preprint arXiv:2405.14458, 2024.

[25] Ultralytics, "YOLOv5: A state-of-the-art real-time object detection system.", 2020. Source. [Online]. Available: https://docs.ultralytics.com.

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 very last consideration in favour of YOLO family is the Ultralytics and DepthAI [26]API support for model deployment. Inference speed benchmarks were reviewed as shown in Table 4. Performed on Robotic Vision Core 2 (RVC2) device, with models compiled to 8 shaves and 2 NN inference threads, including USB3 latency[27].

Table 4: YOLO nano model performance on OAK-D camera device

Architecture	Image resolution	FPS
YOLOv10n	640×640	12.62
YOLOv10s	416×416	14.03
YOLOv11n	640×640	12.80
YOLOv11n	416×415	28.08

The minimum 10 FPS processing rate allows for operation at speeds up to 36 km/h, providing sufficient flexibility for various driving behaviours including urban navigation, parking lot monitoring, and controlled inspection scenarios where moderate speeds are typical.

2.3.3 DEVELOPMENT

The baseline model neural network architecture selected YOLOv11 nano with 2.6M parameters, validated on OAK-D-S2 device to achieve optimal balance between detection performance and computational constraints given the small dataset size. Hardware performance evaluation revealed that model inference achieved 99.6% Connection Matrix Memory (CMX) utilization across 6 SHAVE (Streaming Hybrid Architecture Vector Engine) cores and 2 processing threads, while maintaining efficient LeonRT processing with only 12.93% CPU usage during inference operations.

The inference computational complexity remains predictable due to the single-image batch size. This configuration ensures consistent computational demand of 6.5 GFLOPs for 640×640 image resolution.

Model training and validation were performed on the dataset described in previous section. The model was trained from scratch without using pretrained weights. Used Automatic Mixed Precision with single and half precision preserved throughout training and subsequent model conversion to OpenVino IR[28] and DepthAI blob formats. All training experiments implemented early stopping to control overfitting, which terminated training when box loss increased on the validation set for 15 consecutive epochs.

Initial training was conducted at two different image resolutions to determine the optimal input size for model performance. Selected resolution options — 480×480 and 640×640—both compatible with the stride-32 architecture requirements. Training at 480×480 resolution exhibited early overfitting, with validation metrics showing increasing box loss, classification loss, and Distribution Focal Loss (DFL) beginning around epoch 100. The box loss reached at 0.808 during epochs 80-100, indicating suboptimal learning convergence.

In contrast, training at 640×640 image resolution demonstrated sustained improvement throughout the training process, achieving a significantly lower box loss of 0.643 at epoch 160 without signs of premature overfitting. The performance and training stability confirmed that 640×640 image resolution for training provides sufficient spatial detail for effective pothole

[26] DepthAI, 2025. Source. [Online]. Available: https://github.com/luxonis/depthai/tree/main

[27] RVC2 NN Performance, 2025. Source. [Online]. Available: https://docs.luxonis.com/hardware/platform/rvc/rvc2/#Robotics%20Vision%20Core%202%20(RVC2)-RVC2%20NN%20Performance

[28] OpenVINO IR format. 2025. Source. [Online]. Available: https://docs.openvino.ai/2023.3/openvino_ir.html

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

feature extraction. Consequently, 640×640 pixels were selected as the final image resolution, and image size was excluded from subsequent hyperparameter optimization to focus on other parameters.

Due to the complex properties of pothole objects, several hyperparameters were selected to enhance model regularization. The weight decay regularization factor was increased three times from the default 0.0005 to 0.0015, resulting in consistently better evaluation metrics and smaller box loss values across multiple experiments with different random seed values. As the learning curves appeared unstable during the initial training phase until epochs 15-20, a warm-up period of 15 epochs was implemented to stabilize early training dynamics. Additionally, the learning rate was initially set to 0.01 and gradually decreased to 0.001 to balance training speed with convergence stability.

Due to the small training data, additional training diversity was created through an augmentation pipeline optimized for grayscale imagery. Given the grayscale nature of the training images, HSV-Hue and HSV-Saturation augmentations were eliminated, while HSV-Value augmentation was increased to 0.4 to provide additional brightness variation. Mosaic augmentation—a technique that combines several training images into a single composite image to expose the model to multiple objects and contexts simultaneously—had its probability increased from original 0.1 to 0.8, as multiple experiments demonstrated significantly better evaluation results and learning curves when the probability was set to at least 0.4. Copy-paste augmentation was applied with a 0.1 probability, where pothole objects were flipped about vertical and, or horizontal axes and pasted back onto the original image, creating additional object variations while maintaining realistic spatial relationships. The final model architecture training setup were selected considering the following critical parameters as shown in Table 5 and the model training evaluation parameters as shown in Table 6.

Table 5: Final model training parameters

Parameter	**Value**
Training image resolution	640×640
Training period	160 epochs
Warm-up period	15 epochs
Weight decay regularization factor	0.0015
HSV-Hue probability	0.0
HSV-Saturation probability	0.0
HSV-Value probability	0.4
Mosaic augmentation probability	0.8
Copy-Paste augmentation probability	0.0

Table 6: Model training evaluation metrics

Metric	**Value**
Training Loss (Final Epoch)	
Box Loss	1.179
Classification Loss	0.606
DFL Loss	1.333
Validation Loss (Final Epoch)	

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

Box Loss	1.304
Classification Loss	0.604
DFL Loss	1.234
Bounding box detection Performance	
mAP50	0.777
Precision	0.807
Recall	0.671

2.3.4 VALIDATION

Model validation was performed using five classical evaluation metrics to assess detection accuracy for pothole detection applications. Measurements were conducted on the test subset of the dataset described in previous section. A bounding box proposal is considered valid if the model prediction confidence score is higher than or equal to the specified threshold.

Intersection Over Union (IoU) quantifies the spatial similarity between ground-truth and predicted bounding boxes to evaluate prediction accuracy, calculated as the ratio of overlapping area to total combined area of both boxes. The IoU value is used to categorize each model proposal into performance classes as listed below:

- **True Positives (TP)** - correctly detected potholes with at least 0.1 IoU overlap between predicted and ground truth bounding boxes, representing successful detections.
- **False Positives (FP)** - incorrect detections where predicted bounding boxes have less than 0.1 IoU overlap with any ground-truth bounding box in the specific image, indicating detection errors.
- **False Negatives (FN)** - missed detections where actual potholes present in ground-truth annotations were not detected by the model.
- **True Negatives (TN)** - excluded from object detection evaluation metrics as they represent correctly identified background.

Based on performance classification, four major evaluation metrics established as listed below:

- **Recall** - measures the model's ability to detect all actual potholes present in the test set.
- **Precision** - measuring the proportion of correct predictions among all positive detections.
- **F1 score –** geometric mean between precision and recall.

For the pothole detection application, model selection was primarily driven by Precision metric optimisation to minimise FP rates while accepting increased FN rate. This trade-off is justified because the same pothole can be detected across multiple consecutive frames as the vehicle approaches the pothole. The dataset described in Chapter 2.3.1 is designed so that the model trained on it, is capable to detect potholes from various viewing angles and distances, ensuring that critical road hazards are identified through multiple detection. Therefore, prioritising detection reliability aligns with practical deployment requirements, where false alarms cause more operational disruption than temporarily missed detections. Table 7 demonstrates the precision-recall trade-off across varying model confidence thresholds (range 0 to 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

Table 7: Final model evaluation results

Confidence Threshold	Precision	Recall	F1 Score	TP	FP	FN
0.3	0.866	0.748	0.803	730	113	246
0.4	0.894	0.709	0.791	692	82	284
0.5	0.940	0.628	0.753	613	39	363
0.6	0.960	0.541	0.692	528	22	448

By prioritising reliability, model deployment was configured with confidence threshold 0.5, providing 94% correct predictions, 6% incorrect predictions, and 37% missed detections. Table 8 demonstrates different challenging cases where the model prediction confidence threshold was set to 0.4.

Table 8: Final model hard samples

A.	B.	C.
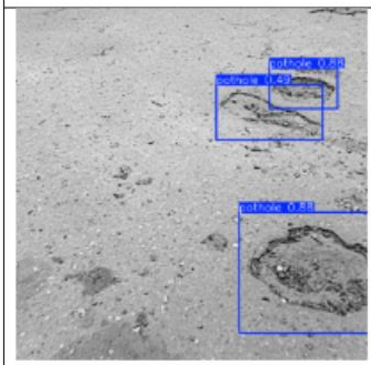	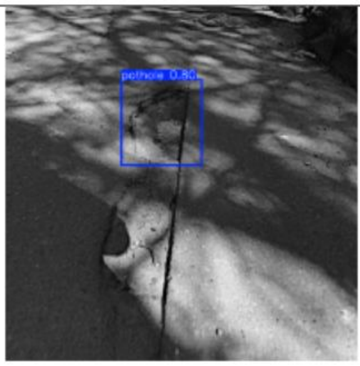	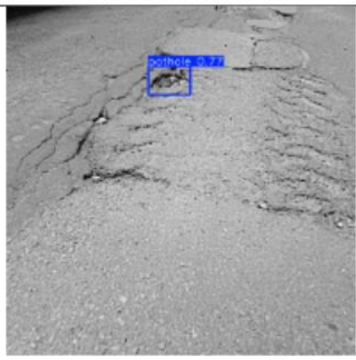
D.	E.	F.
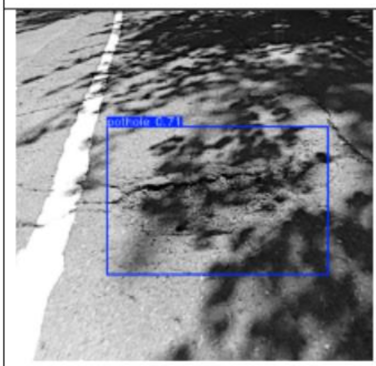	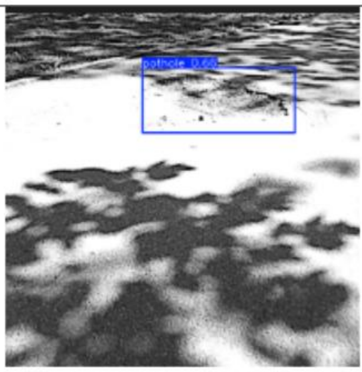	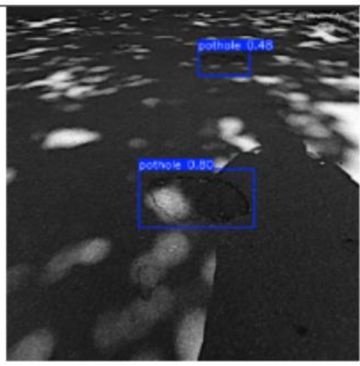

Case A demonstrates the model's capability to distinguish between radial-shaped dirty patches and actual potholes, avoiding false positive predictions that could mislead the detection system.

Case B demonstrates the model's difficulty in detecting two potholes under very similar environmental conditions, highlighting detection limitations in challenging scenarios.

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

Case C demonstrates the model's ability to detect the largest pothole while remaining robust to asphalt surface inconsistencies. This detection capability is particularly critical as such large potholes would pose serious hazards for micromobility drivers.

Cases D, E, and F demonstrate challenging detection scenarios where potholes are obscured by shadows and adverse lighting conditions, representing real-world deployment complexities.

The above algorithm and the kit will be validated in the Riga Living Lab (Latvia) as part of Task 4.4, where it will be installed on vehicles from the Riga City fleet to detect defects in urban infrastructure. Additional details will also be provided in subsequent deliverable D2.3 “Advanced infrastructure monitoring and predictive maintenance tools”

2.4 VISION-BASED AI ALGORITHM FOR TRAFFIC SIGNS [CEA]

2.4.1 INTRODUCTION

The primary objective of this work is to support road authorities and maintenance operators by automatically detecting occluded, problematic or degraded traffic signs, which pose safety risks for drivers and other road users. The developed AI-based detection tool is particularly valuable for secondary and rural roads, where manual inspections are infrequent and reports from citizens are sparse.

This section presents the development of the Vision-based AI algorithms for detecting and classifying traffic signs with occlusions and defects. It covers thorough exploration of the problem by explaining the choice of the dataset with its preprocessing, a well-known detection model, and validation of the tool with the results obtained at this point.

The proposed Vision-based AI tools could be applicable to any of the pilots. However, data from pilots representing specific situations could be useful for the validation phase.

2.4.2 DATASETS

Developing an AI model that detects / classifies traffic signs and distinguishes between their different states requires annotated data for training a Deep Learning model in a supervised manner. There exists a number of public datasets in the literature to solve traffic sign detection or classification task:

1. **German Traffic Sign Detection Benchmark** (GTSDB[29]) introduced in 2013 contains 900 images, 600 for training and 300 for evaluation. In addition to being a fairly small dataset, the images are in PPM format. It is relatively a less popular format compared to JPG or PNG which makes it less convenient to use.
2. **Belgian Traffic Sign Dataset** (BTSD[30]) is a bit larger than GTSDB with PPM images and 62 different classes for traffic signs. Although such a variety on the type of traffic sign is good, our final goal of differentiating between actual states of traffic signs is not a considered aspect in this dataset.
3. **Mapillary and DFG**[31] contains 19346 images and 76 classes for different categories of traffic signs ensuring at least 200 instances for each. It is literally a mix of two open-source datasets, however specifically refined for Africa region. Moreover, the documentation in Kaggle (online platform containing shared datasets) clearly states missing annotations for some samples.

[29] https://www.kaggle.com/datasets/safabouguezzi/german-traffic-sign-detection-benchmark-gtsdb

[30] https://www.kaggle.com/datasets/abhi8923shriv/belgium-ts

[31] https://www.kaggle.com/datasets/nomihsa965/traffic-signs-dataset-mapillary-and-dfg

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

There are several other datasets to mention such as **Chinese Traffic Sign Detection Benchmark** (CTSDB2017[32]), **TT100K**[33] focusing on Traffic Sign Detection in the Wild, etc. However, almost none of them addresses the problem of degradation in the traffic sign. There are few papers tackling this problem by generating a small internal dataset via self-collection or synthetic approach[34]. They can be used alone or merged with one or more traffic sign benchmarks mentioned above. For example, there is Damaged Signs Datasets[35] which has collected Portuguese traffic signs in various regular and unfavourable conditions. The latter corresponds to defective states of the traffic signs, but the labels consider only the type of the traffic sign and not its actual state.

Furthermore, the datasets above focus on 2D detection. By assuming possible usage of traffic signs in autonomous driving, we investigated several 3D dataset benchmarks in this field as well. A popular benchmark **nuScenes**[36] does not contain traffic signs at all. **Waymo**[37] considers the annotation of traffic signs, but it is in 3D and as a starting point, we focus our work more on 2D detection.

Our final choice stood on **Zenseact Open Dataset** (ZOD[38]), which is a recent autonomous driving dataset. It is quite large with 100k frames, where for each frame there are 2D annotations of different objects, including 446k instances of traffic signs across 156 classes with large variations in viewing angle, distance, lighting condition, etc. It is worth noting that there are no labelled advertisement boards among these 156 classes. The annotated objects have also different properties such as *occlusion_ratio* and *clearness*, which could potentially help for distinction between occluded / degraded signs and the normal ones. ZOD is very well documented and complete, contains almost no flaws of acquisition and according to authors, it is the largest commercial dataset for traffic sign annotations. Another advantage of using ZOD is robustness as data collection is performed in 14 different European counties. Last, but not least, the current work can easily be extended to 3D detection as a future perspective since ZOD provides 3D annotations as well.

In conclusion, considering the strengths of the ZOD and the aforementioned datasets, such as a limed instances, less used image format, incomplete and poor organization of the dataset and diversity of region of acquisition, ZOD emerges as the most suitable choice for our needs.

For the moment, we use ZOD without making any extra changes on it, i.e., it does not take into account the defects on traffic signs, and we did not apply certain techniques for its extension. We decided to have the first baseline using the existing annotations in ZOD, and then the second step will be a model that distinguishes between degraded and non-degraded traffic signs.

Based on the bibliographical review, since there are very few and less known available datasets for degraded traffic signs, specific traffic sign images with partial defects can be found in existing datasets and further augmented with different approaches. Synthetic simulation of sign degradation with augmentations can be applied via adding blur, noise or fake occlusion. Several existing datasets can be concatenated as well to benefit both from degraded and non-degraded aspect, but one should also pay attention to avoid unbalancing issue in the final merged dataset. Given that the pilot countries'

---

[32] https://github.com/csust7zhangjm/CCTSDB

[33] https://www.kaggle.com/datasets/braunge/tt100k

[34] Floros, Georgios, Kyritsis, Konstantinos, et Potamianos, Gerasimos. Database and baseline system for detecting degraded traffic signs in urban environments. 2014 5th European Workshop on Visual Information Processing (EUVIP). IEEE, 2014. p. 1-5 & SOO, Yin-Yi, WANG, Yujie, XIANG, Haoxiang, et al. A novel transfer learning model for battery state of health prediction based on driving behavior classification. Journal of Energy Storage, 2025, vol. 111, p. 115409 & CHEN, Tengyang et REN, Jiangtao. MFL-YOLO: An object detection model for damaged traffic signs. arXiv preprint arXiv:2309.06750, 2023.

[35] https://www.kaggle.com/datasets/danielvareta/damaged-signs-dataset

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

[37] https://waymo.com/

[38] https://zod.zenseact.com/

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

traffic signs may differ in shape, colour, or design from those in the open datasets used for training, additional fine-tuning is planned to ensure reliable detection performance.

2.4.3 PREPROCESSING ANALYSIS

For experimental purposes described in this and following sections, we take advantage from Ultralytics[39], a powerful Python package which eases the process of using AI in production. It also provides pre-trained models for fine-tuning and this is one of the reasons why we decided to work with this package.

Ultralytics requires certain dataset format for training and inference, which is why preprocessing on ZOD was a must. Originally, annotations for 2D bounding boxes are defined via 4 points for each corner of the box in JSON files. Below one can see an example of the generated annotation file after preprocessing:

0 0.526989 0.4930764 0.522621 0.4902214 0.5257027 0.5023579 0.5221413 0.4969514
1 0.439121 0.4629339 0.444504 0.4651291 0.4434745 0.4744655 0.4381546 0.4651823
...

Every image in the dataset has a corresponding annotation file organised in the format above. Each row corresponds to an object in that image and consists of 9 values separated by blank space. The first value indicates the class of the given object, whilst the following 8 numbers define the coordinates of the bounding box as *x1y1x2y2x3y3x4y4*. Normalization of these coordinates is another requirement in Ultralytics, which is why they are between 0 and 1. Images themselves should be normalized as well, but within the package, it is automatically done during training or inference.

There are 156 unique object types in ZOD and we keep only « Traffic Sign » among them for our final pre-processed dataset. The objects with bounding box coordinates beyond the image size or not following the convention (e.g. containing less than 4 points) are excluded. Out of 100k frames, there are only 3,427 images not containing any Traffic Sign at all, which is negligible. Those are avoided for training purposes and not kept in the final dataset after pre-processing. The dataset is splitt with a ratio of 0.80 / 0.15 / 0.05 among training, validation and testing data respectively.

All the preprocessing described above is done with a Python script on a raw data of ZOD.

2.4.4 TOOL DEVELOPMENT

Considering recent advancements on YOLO models, our plan was to use YOLO11[40] as an object detector. According to documentation of Ultralytics, this version introduces significant improvements in architecture and performance compared to its predecessors. For real-time object detection, the fastest version of YOLO11 (i.e. yolo11n) is used as a pre-trained model.

The experiments are conducted with one A100 GPU of 40Go, batch size of 16, 50 epochs, varying image sizes (640, 1280, 2000) as an input and automatically chosen optimizer based on model design. The model eventually performs one-class, i.e. traffic sign detection.

2.4.5 TOOL VALIDATION

For the moment, the best model is trained with the image size of 2,000 and it shows *70 mean Average Precision (mAP)* for traffic sign detection. The inference time per image is measured as *20ms* in a P5000 GPU of 16Go. Some visual results are illustrated in Figure 14:

---

[39] https://www.ultralytics.com/

[40] Khanam, R. et Hussain, M. Yolov11: An overview of the key architectural enhancements. arXiv 2024. arXiv preprint arXiv:2410.17725.

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

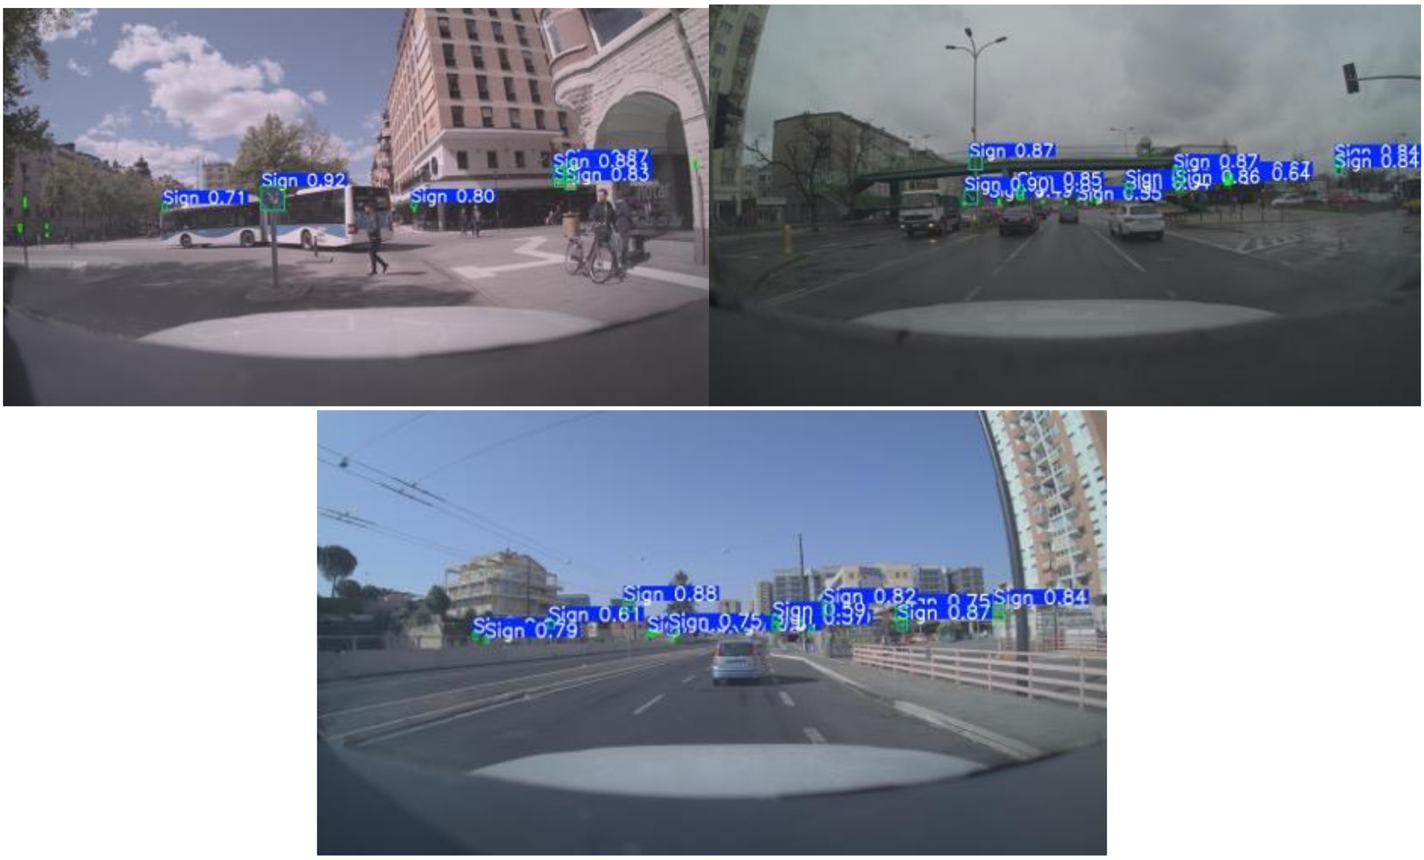

Figure 14: Visual results of the tool

A demo video can be found **here.**

2.5 AUTOMATED ROAD DISTRESS DETECTION [DTU]

The goal of the **Anomaly Detection Tool** is to provide an AI-based solution for detecting road distress and anomalies through the use of multisensory data, specifically by leveraging **vibrations (time series data)** and **images**. This tool is designed to automatically detect road surface conditions such as cracks, potholes, and other forms of degradation, enabling faster and more accurate decision-making in road infrastructure management. By analysing vibration data from sensors (capturing vehicle-induced vibrations) and visual data (via road surface imaging), the tool helps identify potential anomalies and degradation patterns that may not be immediately visible or detected through traditional inspection methods.

We utilise open-source data (see section 2.5.1 for details), which provides both vibration data (time series) and image data collected from road inspections. These datasets are key to training the AI model for effective anomaly detection, allowing it to learn patterns that correspond to road distress. Currently, due to the lack of both raw and image data for the same pilot, the model will be trained separately for **time series anomaly detection** using vibration data and **image anomaly detection** using road surface images. Once both data types become available for the same pilot, it is planned to merge these two models into a **multi-modal model** that can analyse both vibration and image data together. The application of this tool significantly enhances the efficiency of road monitoring, supporting real-time damage detection and prioritisation of repairs. These capabilities directly contribute to optimising road maintenance processes as outlined in Task 2.3 and the broader objectives in Task 3.4, further advancing the automation and effectiveness of infrastructure maintenance management.

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

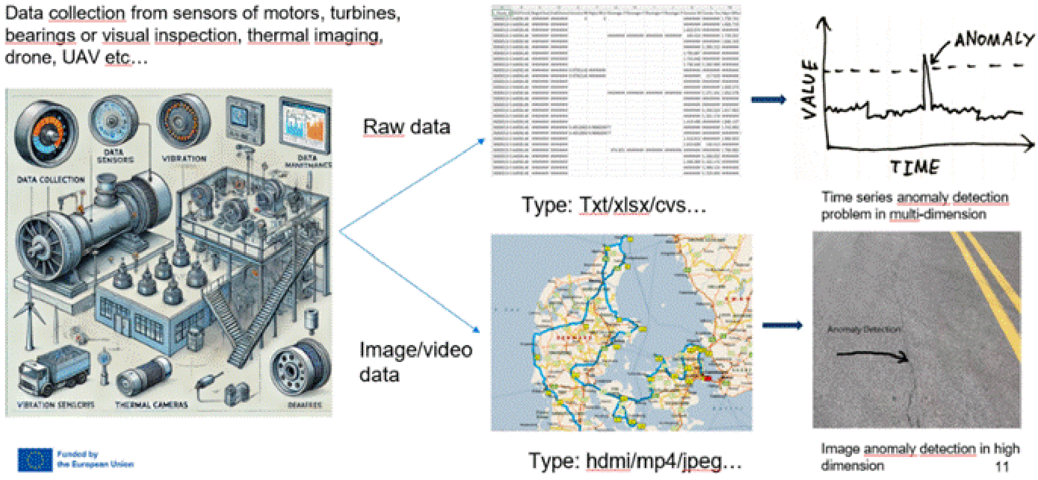

Figure 15: A general framework for anomaly detection using vibration and image data

The business value of this tool is the following:

1. **Operational Efficiency**: By automating the detection process, the tool reduces the need for manual inspections, lowers operational costs, and speeds up the detection of road distress.
2. **Cost Savings**: Early detection of road degradation helps prioritize repairs, thereby minimising long-term maintenance costs by addressing issues before they escalate into more expensive repairs or traffic disruptions.

The tool might be piloted in **Italy** where the focus is on direct distress detection or **detecting high-risk areas**. This allows for more focused maintenance and safety interventions, ensuring that resources are directed toward the most critical areas. Furthermore, the tool directly contributes to the objectives of the O2 work package by enhancing infrastructure resilience and sustainability. It enables a more targeted and data-driven approach to maintenance planning, supporting the long-term goal of optimizing road networks' performance while minimising maintenance costs.

2.5.1 DATASETS

The data used in this project is sourced from the **Lira Project**[41] and the **Road Damage Detection (RDD)** dataset[42]. These open-source datasets are critical for monitoring road distress through sensors and imaging technologies. Specifically, the datasets include:

- **Vibration Data**: This data comes from road vibration sensors placed along road segments. It captures the vibrations caused by vehicles moving over distressed road surfaces, enabling anomaly detection related to road degradation. The vibration data used in this project is sourced from the **Lira Project**, specifically for the **Copenhagen area in Denmark**.
- **Image Data**: High-resolution images captured from road inspections provide visual evidence of distress such as cracks, potholes, and surface deterioration. These images are essential for detecting road surface anomalies and are sourced from the **Road Damage Detection (RDD)** dataset, which includes data from six countries: **Japan,**

[41] https://lira-project.dk/

[42] https://datasetninja.com/road-damage-detector

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

India, the Czech Republic, Norway, the United States, and China. The data has four annotated damage classes such as Longitudinal cracks, Transverse cracks, Alligator cracking, and Potholes.

2.5.2 PREPROCESSING ANALYSIS

Before feeding the data into the ML models, the raw data undergoes extensive preprocessing. The preprocessing steps for both **vibration** and **image data** are as follows:

1. **Vibration Data:**
    a. **Noise Removal**: The vibration data is filtered to remove noise and irrelevant signals using advanced signal processing techniques such as **wavelet transforms** or **Fourier analysis**. This ensures that only the relevant vibrations, indicative of road distress, are retained.
    b. **Feature Extraction**: Key features are extracted from the time series data to identify potential road distress. These features may include **frequency anomalies**, **amplitude changes**, and **irregular patterns** that correspond to road degradation, such as the International Roughness Index (IRI), Mean Depth Profile (MDP), or surface irregularities.
    c. **Normalization**: The vibration data is normalized to ensure that all features are on a comparable scale, allowing for consistent analysis by the AI model. This step ensures the data is ready for input into the model for anomaly detection.
2. **Image Data**:
    a. **Image Enhancement**: Image processing techniques are applied to improve the quality of the images, such as **contrast adjustment**, **noise reduction**, and **sharpness enhancement**. This step ensures that the visual features of road distress (cracks, potholes, etc.) are clearly visible and detectable.
    b. **Segmentation**: Road surface anomalies, such as cracks, potholes, and surface deterioration, are detected and segmented using techniques like **convolutional neural networks (CNNs)** or other **deep learning-based approaches**. This allows the AI model to identify specific regions of distress in the images.
    c. **Data Augmentation**: Augmentation techniques, such as **rotation**, **flipping**, and **scaling**, are applied to create a more diverse and robust dataset. This enhances the generalizability of the model, allowing it to perform well on various types of road conditions and distress patterns.

2.5.3 TOOL DEVELOPMENT

The **Automated Road Distress Detection** tool integrates both **vibration** and **image data** to create a unified AI solution for road monitoring. The tool is based on advanced ML algorithms that analyse multisensory data to detect road distress. The development process includes:

1. **Model Selection**: The tool uses a hybrid model that combines **time series anomaly detection** techniques for vibration data, such as **ML algorithms for classification** (e.g., **random forest** and **logistic regression**) with **probability prediction**. While current implementations focus on machine learning approaches due to limited data for deep learning, **recurrent neural networks (RNNs)** or **long short-term memory (LSTM) networks** could potentially be explored in the future as more data becomes available. For image data, the tool uses a generalized family of **convolutional neural networks (CNNs)** for detecting road distress. This hybrid approach allows the model to effectively leverage both vibration and image data for robust anomaly detection.

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

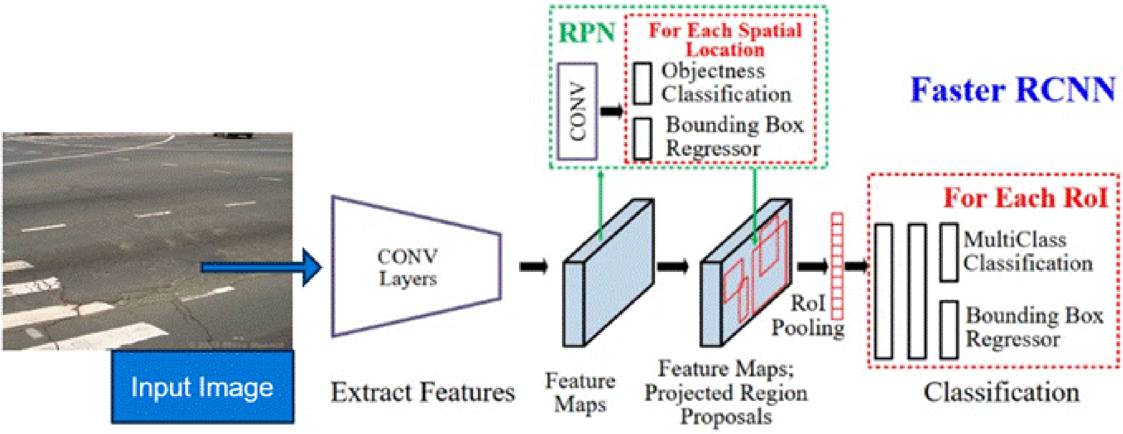

Figure 16: A general framework for object detection using faster RCNN architecture

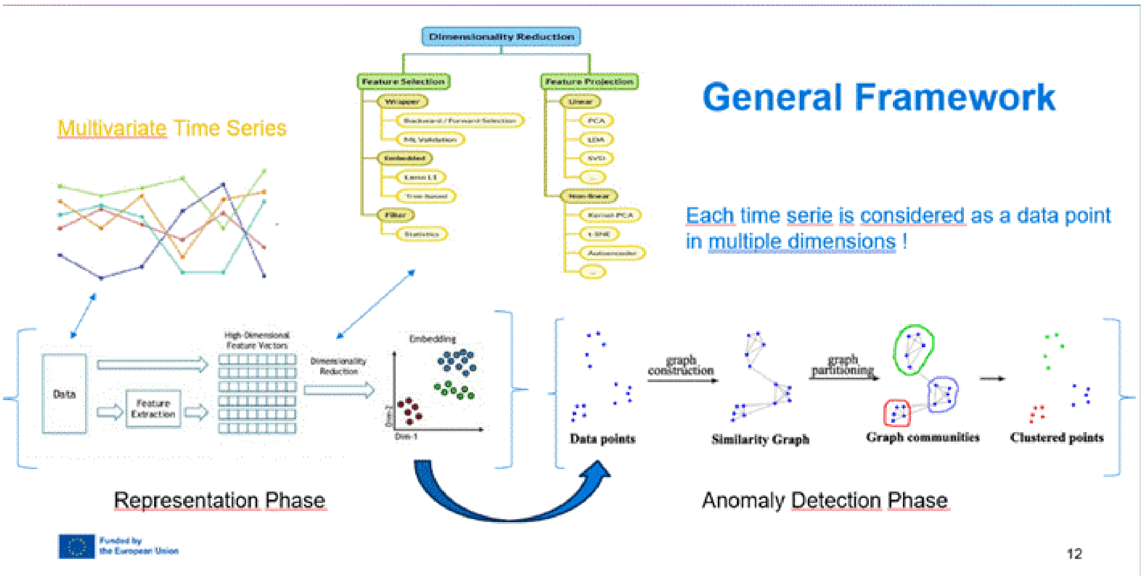

Figure 17: A general framework for time series detect anomaly via dimensional reduction and graph decomposition

2. **Training and Optimisation**: The models are trained using labelled datasets[43] from the **Lira Project** for vibration data and the **Road Damage Detection (RDD[44]) 2024** dataset for image data. Hyperparameter optimisation and **cross-validation** are employed to maximise model accuracy and ensure the model generalises well without overfitting. These techniques help improve the model's robustness and performance on unseen data.
3. **Integration**: The tool is integrated into a **user-friendly interface** with a GPS map that allows infrastructure managers to upload data, process it, and receive actionable insights on road conditions. The tool automatically flags areas requiring immediate attention and provides recommendations for maintenance priorities based on detected anomalies, contributing to more efficient decision-making in road infrastructure management.

---

[43] A. Skar, A. M. Vestergaard, T. Br¨usch, S. Pour, E. Kindler, T. S. Alstrøm, U. Schlotz, J. E. Larsen, and M. Pettinari, “Liracd: An open-source dataset for road condition modelling and research,” in Data in Brief. IEEE, 2023, pp. 1–15.

[44] Arya, D., Maeda, H., Ghosh, S. K., Toshniwal, D. & Sekimoto, Y. RDD2022: A multi-national image dataset for automatic Road Damage Detection. Preprint at https://arxiv.org/abs/2209.08538 (2022).

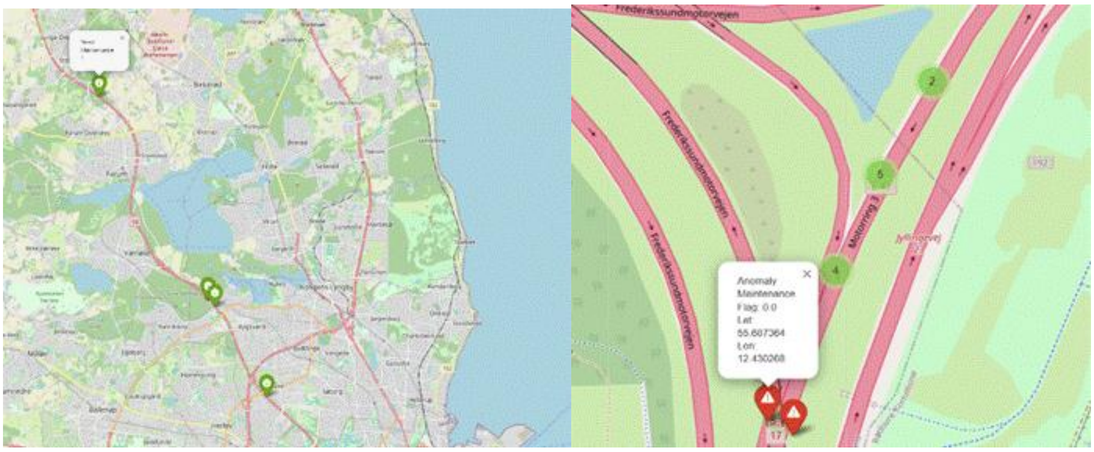

Figure 18: An integrated dynamic platform with GPS coordinates for anomaly places

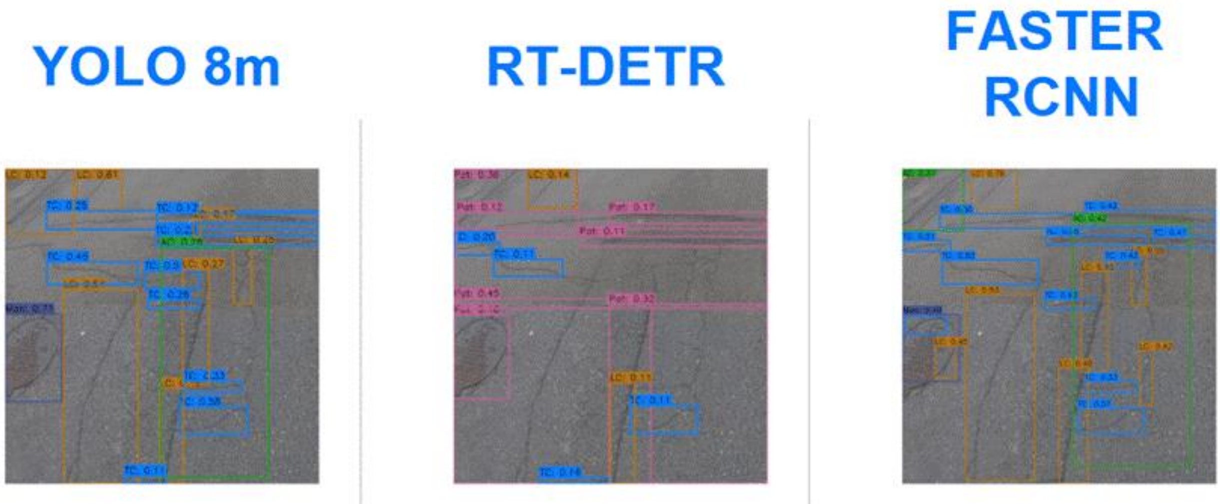

Figure 19: Results of model training from several architectures with probability of object detection

2.5.4 TOOL VALIDATION

The effectiveness of the tool is validated through a series of rigorous tests, including:

1. **Pilot Testing:** The tool needs to be confirmed in pilot regions such as Latvia and Italy to assess its performance under real-world conditions. In these regions, vibration sensors and image capture systems are installed on road sections to gather data for analysis.
2. **Performance Metrics:** The model's accuracy in detecting road distress is evaluated using standard performance metrics, including precision, recall, F1 score, and the area under the receiver operating characteristic curve (AUC-ROC).
3. **Comparison with Traditional Methods:** The tool's performance is compared to traditional manual inspection methods to demonstrate its added value in terms of speed, cost efficiency, and accuracy.

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

2.6 ROAD PAVEMENT STATUS USING CONNECTED VEHICLES DATA [INDRA]

Unlike the other sections, this section presents a monitoring tool hence the section is smaller, does not follow the same structure and is not as AI-focused as the others.

The monitoring tool is designed to assess the status of road pavement by leveraging data collected from connected vehicles. To estimate pavement conditions accurately, a range of advanced AI techniques will be explored. These methods aim to infer both the current state and the temporal evolution of road surfaces by analysing data streams, primarily vertical accelerometery and precise positional information (1m-10m approximately).

While the ultimate implementation of the pilot will rely on data sourced from a fleet of connected vehicles, Indra is currently utilising embedded GPS and accelerometery data obtained from public buses to train and validate the AI models. This preliminary step is critical as it provides a cost-effective large and diverse dataset, which enhances the robustness of the ML algorithms. Additionally, using buses as data sources allows for the evaluation of how different vehicle types and their dynamic behaviours impact the quality and reliability of the collected data.

This approach, combining with future connected vehicles data, will ensure that the monitoring tool is capable of handling variations in vehicle characteristics, which is essential for the successful deployment of a scalable, real-time pavement monitoring system based on connected vehicle data. By leveraging continuous streams of vibration and telemetry data processed through AI-based algorithms, this solution will be able to detect pavement anomalies in near real-time, thus ensuring timely and accurate insights for maintenance planning and road safety improvement. To date, data from a public transport operator in the city of Alicante (Spain) has been collected and meticulously processed through a series of well-defined steps to ensure its quality and usability for pavement condition monitoring:

- **Data ingestion and preprocessing:** The raw data collected, including GPS and vertical accelerometery signals, is first ingested and systematically organised into a database. This step involves cleaning the data, handling missing or inconsistent entries, and structuring it for efficient querying and subsequent processing. A relational database was chosen as the volume of data is moderate and structured, making relational queries efficient.
- **Data fusion:** The vertical accelerometery data is combined with GPS positional information through spatial and temporal alignment techniques. This fusion process is critical to accurately correlate vehicle vibrations with precise locations on the road network.
- **Road network extraction:** Road infrastructure data is obtained from a publicly available source, OpenStreetMap, which provides the basic map layout. This data is further refined by segmenting the roads into smaller sections, limiting the segment length to no more than 100 meters. This granularity allows for more detailed analysis of pavement conditions at a localized level.
- **Trajectory smoothing:** To reduce the noise and irregularities in the GPS position data, Kalman filtering is applied. This mathematical technique helps smooth the trajectories, resulting in more reliable vehicle location estimates over time.
- **Map matching:** The smoothed GPS coordinates are then matched to the corresponding road segments in the extracted network. This step ensures that each position data point is accurately assigned to a specific part of the road infrastructure, enabling precise spatial analysis.
- **Accelerometery filtering:** To isolate relevant pavement-related vibrations, the accelerometery signals undergo filtering through both high-pass and low-pass filters. These filters remove exogenous effects such as vehicle dynamics unrelated to road surface conditions, ensuring that the data reflects pavement irregularities as accurately as possible.
- **Final dataset preparation:** After all these processing stages, the data is carefully structured, cleaned, and formatted to create a comprehensive dataset. This final dataset is ready for use in training AI models and conducting detailed analysis to estimate road surface conditions effectively.

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

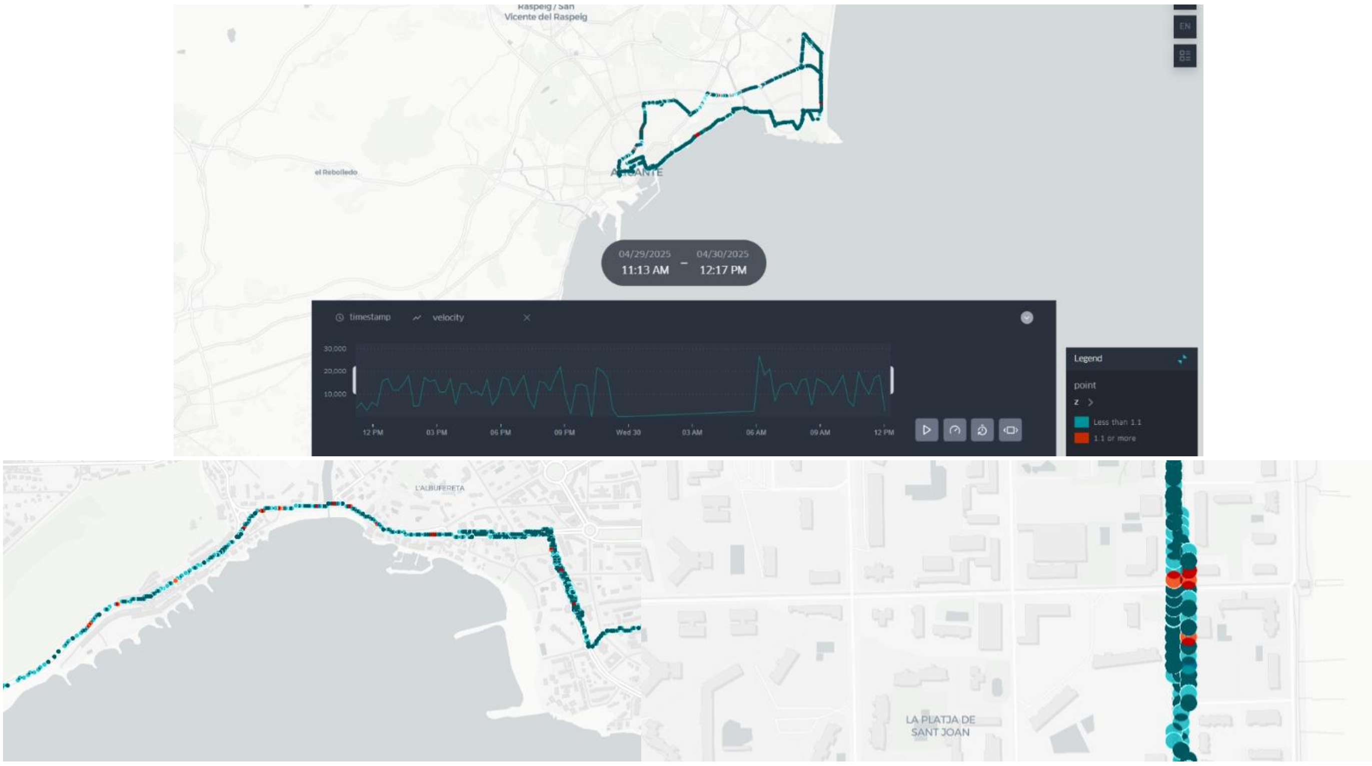

Figure 20: Initial representation of the pavement status with public buses data

Currently, the development of algorithms is underway, focusing on the detection of pavement damage (by using the accelerometers of the vehicles) and the modelling of its progression over time. These algorithms leverage the processed accelerometery and positional data to identify signs of deterioration, such as cracks, potholes, or uneven surfaces, and to track how these defects evolve under varying traffic and environmental conditions.

In future phases of the project, data collected from connected vehicles will be integrated into the system. This integration will enable the validation and refinement of the previously trained models across diverse spatial and temporal contexts, encompassing different road types, traffic patterns, and environmental scenarios. By doing so, the robustness and generalizability of the models will be thoroughly tested and enhanced, ensuring their effectiveness for widespread deployment in real-world pavement monitoring applications.

2.7 INTERFEROMETRIC SYNTHETIC APERTURE RADAR FOR ROAD CONDITION EVALUATION [FRONT]

2.7.1 INTRODUCTION

The use of Interferometric Synthetic Aperture Radar satellite data to identify areas where road conditions have deteriorated was proposed in Sections 1.1.1, 1.2.1.1, and 1.2.1.2 of the Description of Action of the Project. See the Figure below as to the hierarchy of infrastructure defect detection technologies.

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

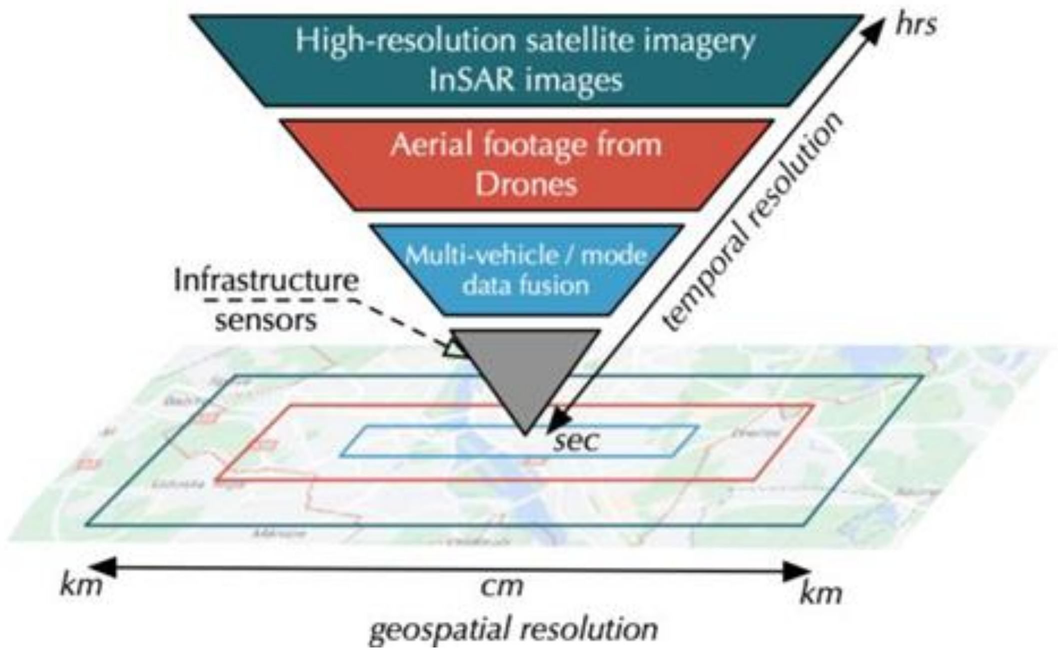

Figure 21: Position of InSAR in the Infrastructure Monitoring Solution Hierarchy

2.7.2 THEORETICAL BACKGROUND

2.7.2.1 INTERFEROMETRIC SYNTHETIC APERTURE RADAR (INSAR)

Synthetic Aperture Radar (SAR) signals contain information on amplitude, i.e., the strength of the radar response, and phase, which is determined by the distance between the antenna and target. Interferometric SAR uses the phase difference between two complex SAR observations of the same area, taken from slightly different positions, to generate distance information. More specifically, in the methodology called Differential Interferometry (DInSAR) the following methodology is used:

• Interferograms are generated from two observations of the same area by combining the phase of the two observations.

• The difference between interferograms from different time periods can be used to obtain deformation patterns by removing the phase shift related to topography.

Persistent Scatterer Interferometry (PSI) uses point scatterers with strong radar backscatter over a long time period (years) to provide a phase history of the point target over time. Persistent scatterers can be small, usually manmade, features that remain very correlated over time. Such scatterers may be purpose-made targets on embankment or cut slopes to monitor ground movement, or regard to road pavements and structures (safety barriers, manhole covers, bridge or culvert abutments, etc.) themselves. This allows overcoming some of the limitations of conventional DInSAR as it maximises the number of usable interferograms which can then be used to calculate mean trends over time, see[45] for a comprehensive introduction to InSAR techniques.

[45] A. Ferretti, A. Monti-Guarnieri, C. Prati, F. Rocca, and D. Massonnet, InSAR Principles: Guidelines for SAR Interferometry Processing and Interpretation. in ESA Technical Manual, no. TM-19. Noordwijk: European Space Agency, 2007. [Online]. Available: https://esamultimedia.esa.int/multimedia/publications/TM-19/TM-19_InSAR_web.pdf

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

2.7.2.2 INFRASTRUCTURE CONDITION MONITORING WITH INSAR

Research on infrastructure condition monitoring using InSAR leverages the availability of high-quality data and sometimes combines it with other sensors, such as Ground Penetration Radar (GPR), or conventional land-surveying techniques. Indicatively, in[46] uses PSI to evaluate the condition of urban infrastructure using Homogeneous Distributed Scatterers, such as road pavements. In a similar vein, the authors[47] examines the difference working with C-Band or X-Band data and, in[48], links the technique with Geographic Information System (GIS) technologies, whereas, in[49], the authors examine the important question of pixel selection. Subsistence rates along a road network were examined in subsequent works[50]. Monitoring bridge conditions was examined in Valencia[51]. Efforts to supplement or replace more traditional pavement condition monitoring techniques, such as profilometric surveys, include[52], which discusses the techniques, opportunities, and challenges when using medium-resolution data, such as the Sentinel-1.

[46] J. Eppler and B. Rabus, "Monitoring Urban Infrastructure With An Adaptive Multilooking INSAR Technique," in Proceedings of Fringe 2011, Frascati: ESA, 2012, p. 68. Accessed: Apr. 20, 2025. [Online]. Available: https://ui.adsabs.harvard.edu/abs/2012ESASP.697E..68E/abstract

[47] L. Bianchini Ciampoli, V. Gagliardi, and A. Benedetto, "InSAR analysis of C-band data for transport infrastructure monitoring," in Eleventh International Conference on the Bearing Capacity of Roads, Railways and Airfields, Trondheim, 2022, p. 10.

[48] V. Macchiarulo, P. Milillo, C. Blenkisopp, and G. Giardina, "Monitoring deformations of infrastructure networks: A fully automated GIS integration and analysis of InSAR time-series," Struct. Health Monit., vol. 21, no. 4, pp. 1849–1878, 2022, doi: 10.1177/14759217211045912.

[49] A. Piter, M. H. Haghighi, and M. Motagh, "Evaluation of Pixel Selection Methods for Traffic Infrastructure Monitoring Using Sentinel-1 InSAR," in The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Nice: ISPRS, 2022. doi: 10.5194/isprs-archives-XLIII-B3-2022-333-2022.

[50] B. Yu et al., "Monitoring subsidence rates along road network by persistent scatterer SAR interferometry with high-resolution TerraSAR-X imagery," J. Mod. Transp., vol. 21, no. 4, pp. 236–246, 2013, doi: 10.1007/s40534-013-0030-y.
L. M. G. Jaya, J. Safani, A. Okto, M. Hasbi, and A. Kadir, "Ground deformation analysis on road infrastructure in north buton Indonesia using interferometric synthetic aperture radar (Insar)," in IOP Conference Series: Earth and Environmental Science, Bali: IOP Publishing, 2019, p. 012095. doi: 10.1088/1755-1315/419/1/012095.
N. Fiorentini, M. Maboudi, P. Leandri, M. Losa, and M. Gerke, "Surface Motion Prediction and Mapping for Road Infrastructures Management by PS-InSAR Measurements and Machine Learning Algorithms," Remote Sens., vol. 12, no. 23, p. 3976, 2020, doi: 10.3390/rs12233976.
S. Karimzadeh and M. Matsuoka, "Remote Sensing X-Band SAR Data for Land Subsidence and Pavement Monitoring," Sensors, vol. 20, no. 17, p. 4751, 2020, doi: 10.3390/s20174751.
L. Zhu, X. Xing, Y. Zhu, W. Peng, Q. Xia, and Z. Yuan, "An Advanced Time-Series InSAR Approach Based on Poisson Curve for Soft Clay Highway Deformation Monitoring," IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens., vol. 14, pp. 7682–7698, 2021, doi: 10.1109/JSTARS.2021.3100086.
O. Pritchard, L. Andrews, J. Gates, M. Willis, M. Free, and J. Codd, "Application of InSAR for geotechnical asset management on England's Strategic Road Network," presented at the Geo-Resilience 2023, Cardiff: British Geotechnical Association, 2023. doi: 10.53243/Geo-Resilience-2023-7-2.

[51] M. Del Soldato, R. Tomás, J. Pont Castillo, G. Herrera Garcia, L.-D. Garcia, and O. Mora, "A multi-sensor approach for monitoring a road bridge in the Valencia harbor (SE Spain) by SAR Interferometry (InSAR)," Rendiconti Online Della Soc. Geol. Ital., vol. 41, pp. 235–238, 2016, doi: 10.3301/ROL.2016.137.
D. Cusson, K. Trischuk, P. Ghuman, D. Hébert, G. Hewus, and M. Gara, "Satellite-Based InSAR Monitoring of Highway Bridges: Validation Case Study on the North Channel Bridge in Ontario, Canada," Transp. Res. Rec. J. Transp. Res. Board, vol. 2672, no. 45, 2018, doi: 10.1177/0361198118795013.
M. Lyu, Y. Ke, X. Li, L. Zhu, L. Guo, and H. Gong, "Detection of Seasonal Deformation of Highway Overpasses Using the PS-InSAR Technique: A Case Study in Beijing Urban Area," Remote Sens., vol. 12, no. 18, p. 3071, 2020, doi: 10.3390/rs12183071.
S. Selvakumaran, C. Rossi, A. Marinoni, G. Webb, J. Bennetts, and E. Barton, "Combined InSAR and Terrestrial Structural Monitoring of Bridges," IEEE Trans. Geosci. Remote Sens., vol. 58, no. 10, pp. 7141–7153, 2020, doi: 10.1109/TGRS.2020.2979961.
D. Tonelli et al., "Interpretation of Bridge Health Monitoring Data from Satellite InSAR Technology," Remote Sens., vol. 15, no. 21, p. 5242, 2023, doi: 10.3390/rs15215242.

[52] N. Fiorentini, M. Maboudi, P. Leandri, and M. Losa, "Can Machine Learning and PS-InSAR Reliably Stand in for Road Profilometric Surveys?," Sensors, vol. 21, no. 10, p. 3377, 2021, doi: 10.3390/s21103377.
A. Piter, M. H. Haghighi, and M. Motagh, "Challenges and Opportunities of Sentinel-1 InSAR for Transport Infrastructure Monitoring," PFG – J. Photogramm. Remote Sens. Geoinformation Sci., vol. 92, pp. 609–627, 2024, doi: 10.1007/s41064-024-00314-x.

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

2.7.3 DATASETS

2.7.3.1 COPERNICUS SENTINEL-1

Sentinel-1 is a SAR satellite system that operates in the C-band[53], which uses interferometry and polarimetry for a range of applications including monitoring ground movement, identifying vegetation and tracking ships on water bodies. Sentinel-1 datasets are available through the Copernicus Browser[54], and the example screenshot below:

Figure 22: Screenshot of the Copernicus Browser

2.7.3.2 COSMO-SKYMED

COSMO-SkyMed are a series of SAR satellites that has been developed over two generations (CSK and CSG) that operate in the X-Band with a resolution of up to 1m[55]. Its primary missions include defence and security, risk management, agriculture and food management, marine and coastal monitoring & meteorological applications. COSMO SkyMED data are available for research purposes after an application through the European Space Agency's (ESA's) Earth Online website[56].

2.7.3.3 TERRASAR-X AND TANDEM-X

TerraSAR-X are SAR satellites with very high geometric accuracy that operate in the X-Band with a resolution of 1m, see[57] for a description of the mission and system design, whereas TanDEM-X is an add-on for Digital Elevation Measurements for TerraSAR-X, see[58] for more details. TerraSAR-X and TanDEM-X data are available for research purposes after an application through ESA's Earth Online website[59].

---

[53] R. Torres et al., "GMES Sentinel-1 mission," Remote Sens. Environ., vol. 120, pp. 9–24, 2012, doi: 10.1016/j.rse.2011.05.028.

[54] European Space Agency, "Sentinel-1." Copernicus Browser. [Online]. Available: https://browser.dataspace.copernicus.eu/

[55] F. Covello et al., "COSMO-SkyMed an existing opportunity for observing the Earth," J. Geodunamics, vol. 49, no. 3–4, pp. 171–180, 2010, doi: 10.1016/j.jog.2010.01.001.

[56] European Space Agency, "Earth Online." [Online]. Available: https://earth.esa.int/

[57] R. Werninghaus and S. Buckreuss, "The TerraSAR-X Mission and System Design," IEEE Trans. Geosci. Remote Sens., vol. 48, no. 2, pp. 606–614, 2009, doi: 10.1109/TGRS.2009.2031062.

[58] G. Krieger, A. Moreira, H. Fiedler, I. Hajnsek, M. Werner, and M. Younis, "TanDEM-X: A Satellite Formation for High-Resolution SAR Interferometry," IEEE Trans. Geosci. Remote Sens., vol. 45, no. 11, pp. 3317–3341, 2007, doi: 10.1109/TGRS.2007.900693.

[59] European Space Agency, "Earth Online." [Online]. Available: https://earth.esa.int/

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

2.7.4 PREPROCESSING

2.7.4.1 GENERATING INTERFEROGRAMS

Preprocessing for the Sentinel-1, Sky-Med, and TerraSAR-X raw data, i.e. generating interferograms from observations over different times, is carried out using the Sentinel Application Platform (SNAP) software provided gratis by the European Space Agency, see[60]. SNAP offers an environment that allows a multitude of functions and is relatively intuitive. See the screenshot below with test raw data:

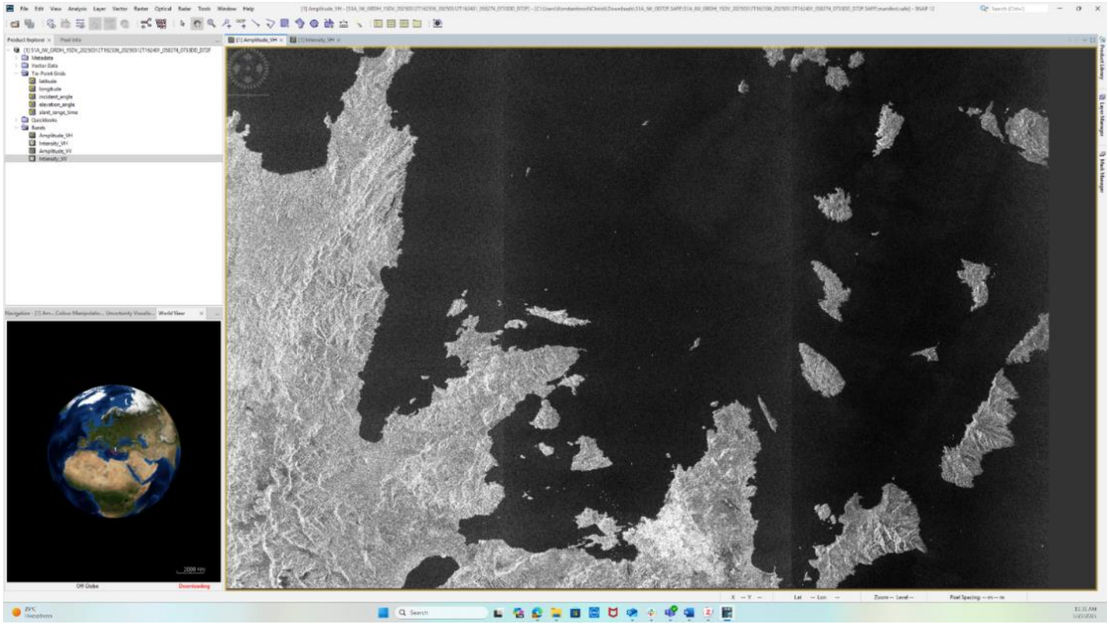

Figure 23: Screenshot of SNAP Opening Raw Sentinel-1 Data. Observe how the image is rotated by 180° relative to mapping conventions

2.7.4.2 CREATING AN AREA OF INTEREST BOUNDARY

Using the entirety of InSAR observations for the evaluation of road infrastructure conditions is counterproductive, as it means spending computational effort and storage on non-road elements. As such, creating an outline of the road infrastructure to act as a boundary of the Area of Interest (AoI) that will identify data that falls in the scope of the process is a necessary step. This can be done using QGIS, an open-source GIS tool, see[61], to create boundary polygons. QGIS can use OpenStreetMap[62] as background or satellite imagery, such as from GoogleEarth[63]. Given that InSAR data is typically registered in the Earth (2015) reference system (latitude and longitude in degrees – IAU_2015:39902), whereas it is more convenient to create bounding boxes in projection reference systems, such as WGS 84 / World Mercator (easting and northing in meters – EPGS: 3395), it is necessary to translate from one reference system to the other.

[60] European Space Agency, Brockmann Consult, Skywatch, Sensar, and C-S, Sentinel Application Platform. (May 08, 2025). [Online]. Available: https://earth.esa.int/eogateway/tools/snap

[61] QGIS Development Team, QGIS Geographic Information System. (2025). QGIS Association. [Online]. Available: https://www.qgis.org

[62] https://www.openstreetmap.org/#map=7/38.359/23.810

[63] https://earth.google.com/web/

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

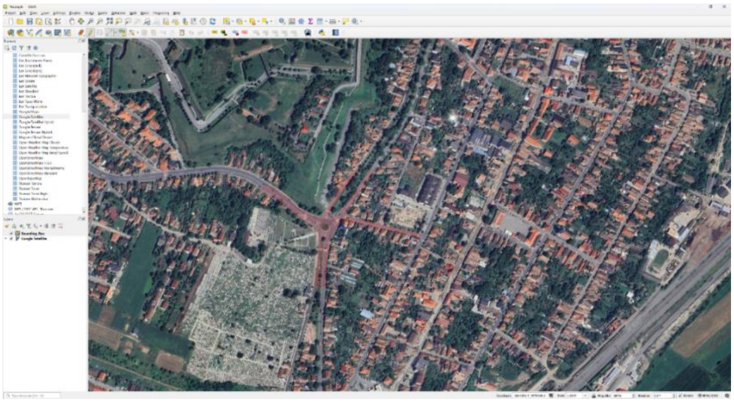

Figure 24: Example of Drawing an AoI Bounding Box in GGIS

2.7.5 TOOL DEVELOPMENT

Once a temporal set of interferograms for the AoI has been obtained, an appropriate tool for a time-series analysis to identify ground movements will need to be applied. This will be based on the StaMPS/MTI solution, see[64] for a more detailed discussion.

2.7.6 TOOL VALIDATION

To ensure practical relevance, the InSAR-based road condition monitoring tool will undergo validation through selected pilot sites within the EvoRoads project. Currently, the tool has been chosen for deployment in the Alba Iulia (Romania) pilot, where it will be applied to monitor structural road conditions via satellite-derived deformation analysis. Discussions are ongoing to expand validation to the Madrid and Galicia pilots, providing coverage across urban, peri-urban, and rural road types. This diverse site selection will help assess the tool's performance under varying topographical and infrastructural conditions. Importantly, the tool is designed to be location-agnostic and applicable anywhere in Europe, provided that suitable satellite data (e.g., Copernicus Sentinel-1) are available. The pilot-based testing will focus on evaluating the precision and usability of the tool in generating actionable insights for road authorities and validating satellite-detected deformations against local ground truth data.

For its technical development and in-house validation, the tool will rely on historical ground-truth data to assess the accuracy of satellite-derived insights. The three pilots - Alba Iulia, Madrid, and Galicia - have been asked to liaise with their respective municipal authorities to identify documented incident sites where substantial road repairs were conducted in the past decade. These known problem areas will serve as control cases: the tool will analyse archived InSAR data to retrospectively observe any surface deformation trends preceding the repairs. By comparing these deformation signals with verified ground interventions, developers can calibrate the detection algorithms and improve predictive reliability. This process not only supports scientific validation but also helps establish practical thresholds for triggering alerts.

---

[64] A. Hooper, D. Bekaert, K. Spaans, and M. Arikan, "Recent advances in SAR interferometry time series analysis for measuring crustal deformation," Tectonophysics, vol. 514–517, pp. 1–13, 2012, doi: 10.1016/j.tecto.2011.10.013. A. Hooper, D. Bekaert, E. Hussain, and K. Spaans, StaMPS/MTI Manual, 4.1b. University of Leeds, 2018. [Online]. Available: https://homepages.see.leeds.ac.uk/~earahoo/stamps/StaMPS_Manual_v4.1b1.pdf

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

2.8 HAZARD AWARE INFRASTRUCTURE MONITORING TOOL [EUT]

It is worth pointing out that the tool presented in this section is developed by EUT who is not involved in T2.1 and current deliverable. Therefore, this section only presents a high-level description and does not contain the technical details of previous sections. For an extensive review of the tool, the reader is refered to D2.2.

The Hazard aware infrastructure monitoring tool aims at identifying objects (vehicles, humans, animals, foreign objects such as rocks, plastic pieces, tree branches, traffic cones etc.) on the road and assess asphalt quality. To achieve it, EUT makes use of a combination of CNN and computer vision techniques.

The tool will be tested in Santa Oliva pilot in terms of its achievement in Debris Detection (using IDIADA cameras) and pedestrian and animal detection. For more details on the tool's application to the Spanish pilot (Santa Olivia), the reader is referred to the Annex.

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 PILOT'S APPICABILITY

In this section we discuss the applicability of the AI tools discussed in previous section. We start by presenting a mapping between tools and pilots summarising the current status of the tools. Then, we present the EU AI Act compliance of our tools and, finally, discuss their integration into the EvoRoads platform.

3.1 TOOLS MAPPED TO PILOTS

In the table below, for each tool developed we present the status of applicability (or potential applicability) to the four EvoRoads pilots. The video demonstrating each tool can be found **here**.

Table 9: Tools mapped to pilots

Tool	Partner	Pilot application
Road accident prediction using historical data	INLECOM	**Applied** to Italian pilot
Automated road distress detection	DTU	Applicable to the Italian pilot
Real-time infrastructure defects detection	DOTS	**Applied** to Latvian pilot
Road damage detection & classification	UCY	**Applied** to Romanian pilot
INSAR for Road Condition Evaluation	FRONT	**Applied** to Romanian pilot; Spanish pilot sites still under consideration
Road pavement status using connected vehicles data	INDRA	**Applied** to Spanish pilot (Alicante)
Hazard aware infrastructure monitoring tool	EUT	**Applied** to Spanish pilot (Santa Oliva)
Vision-based AI algorithm for traffic signs	CEA	Applicable to any of the pilots (under consideration)

3.2 EU AI ACT COMPLIANCE

All eight partners developing the AI tools of previous section have assess the risk of their tool according to the EU AI compliance checker here. This online checker was developed by the team at the Future of Life Institute to aid the effective implementation of the AI Act.

In most cases, the tools were either out of scope of the AI Act or Excluded as RnD or Open-source systems, see figure below (in yellow). The only exceptions are FRONT & IDIADA who have an AI literacy obligation, see figure below (in green).

> Excluded: Research & development
>
> AI systems and models with the sole purpose of scientific research and development are excluded. For all other systems, research & development activities are likely excluded until your AI system is **placed on the market** or **put into service**. Systems and activities that are excluded are not subject to any obligations. For more information see Article 2 points 6 and 8.

> Excluded: Open source systems
>
> Until your AI system is **placed on the market** or **put into service** by a provider as part of an AI system that is high-risk, prohibited, general purpose, or has transparency obligations, your open source system is likely excluded from the EU AI Act, which means it is not subject to any obligations. For more information see Article 2 point 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

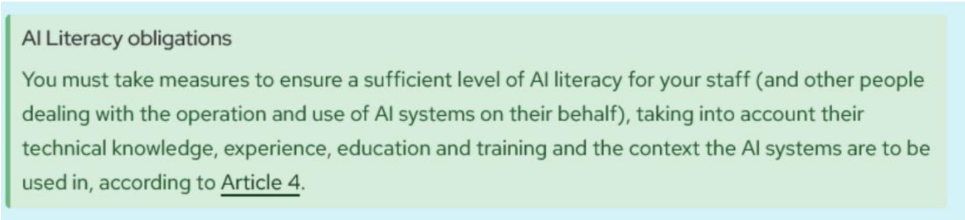
AI Literacy obligations

You must take measures to ensure a sufficient level of AI literacy for your staff (and other people dealing with the operation and use of AI systems on their behalf), taking into account their technical knowledge, experience, education and training and the context the AI systems are to be used in, according to Article 4.

Figure 25: AI compliance checker results

For more details, analytical results of all tools can be found here.

3.3 INTEGRATION OF AI TOOLS INTO THE EVOROADS PLATFORM

3.3.1 INTRODUCTION

This chapter provides an overarching analysis of how the diverse AI tools developed within EvoRoads, with emphasis on those described in this Deliverable, will be validated across the project's pilot sites and progressively integrated into a coherent and user-centric platform. It begins by outlining the validation approach to be followed, detailing the objectives, pilot selection criteria, and contextual challenges involved in ensuring the tools' real-world applicability.

The chapter then transitions to preliminary integration and deployment considerations, with emphasis on establishing seamless interfaces between back-end AI services and front-end components. This includes the design of mechanisms that support end-to-end data flows - originating at the edge and culminating in the dashboards developed in Task 1.5 - through which findings from the AI tools can be surfaced and actioned. Particular attention is paid to cases where algorithms are too resource-intensive for on-device processing, necessitating remote hosting, orchestration, and secure connectivity solutions.

While technical implementation specifics - including system architecture and interface specifications - are the focus of Deliverable D1.3 (due in October 2025), this chapter sets the stage by presenting the initial concepts, requirements, and strategic direction. The concluding section offers a roadmap outlining the necessary steps toward full integration, including progressive deployment, testing, and user interface unification. Overall, the chapter acts as a bridge between the tool-level developments and their operational embedding into the EvoRoads platform.

3.3.2 EVOROADS PILOT VALIDATION APPROACH

Validation within real-world pilot environments is a key component of EvoRoads, ensuring that AI-enabled tools and other technological developments respond directly to operational needs. Unlike feasibility studies or lab-based tests, pilot validation in EvoRoads is rooted in the principle that solutions are built for actual users - not theoretical models. Developers are not guided by assumption, but by direct engagement with those who will ultimately benefit from and maintain the tools.

This principle is embedded in the Living Lab (LL) structure established across the four pilot countries - Italy, Spain, Latvia, and Romania. Each LL brings together local stakeholders including municipalities, infrastructure operators, and citizens. These actors help define requirements, provide feedback during testing, and support the evaluation of the tools in their respective environments. Co-creation and open communication channels are essential to ensure that the solutions developed are usable, relevant, and aligned with local priorities.

All technologies, including the AI components described in this document, will be assessed with measurable results. Each tool is associated with Key Performance Indicators (KPIs) derived from the Road Safety Criteria Framework developed in Task 1.1 and enhanced dynamically under Task 1.3. These KPIs offer a standardised basis for evaluating performance across dimensions such as infrastructure condition, road user behaviour, and safety-related outcomes. Each pilot will test

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

a selected subset of the EvoRoads technologies. The current set of algorithms has been distributed evenly across the four sites to ensure a balanced validation effort. Pilots will evaluate both the technical performance of the tools and their usefulness in supporting real-time decision-making, maintenance planning, or safety interventions. The alignment between each tool and its validation context, including the selection of appropriate KPIs, will be fully detailed in Deliverable D4.1.

This structured and metrics-based validation strategy ensures not only technical verification, but also operational relevance. It also supports the long-term scalability and replicability of the solutions across different urban and regional settings. By grounding development in local needs and verifying results through defined KPIs, EvoRoads ensures that its innovations move beyond the prototype stage to become trusted tools in the field of road infrastructure safety and management.

The selection of pilot sites and the matching of EvoRoads technologies were guided by a structured framework that ensures operational relevance, diversity of conditions, and alignment with the project's validation objectives. Four pilot sites - Turin (Italy), Galicia/Madrid/Santa Oliva (Spain), Riga (Latvia), and Alba Iulia (Romania) - were chosen to represent a spectrum of environments: dense urban centres, secondary rural roads, mixed-use areas, and high-risk pedestrian zones. Importantly, Italy, Latvia, and Romania were selected due to their high road fatality rates and long-standing safety challenges, making them ideal testbeds for demonstrating impact in adverse conditions. In contrast, Spain - with a more mature safety landscape - serves as a blueprint pilot, helping to benchmark practices and validate transferability across different levels of infrastructural development.

Matching technologies to pilots was based on four key criteria: (1) **Local road safety priorities**, such as pedestrian safety in Turin or connectivity mapping in Galicia; (2) **Infrastructure and sensor readiness**, including access to roadside units, drones, or connected vehicles; (3) **Stakeholder engagement**, with strong Living Lab participation used as a prerequisite for co-creation and feedback loops; and (4) **Logistical feasibility**, such as the ability to deploy edge devices or support cloud processing.

Each site tests a distinct subset of AI tools - covering pavement monitoring, signage detection, behavioural nudging, and predictive maintenance - mapped against pilot-specific KPIs selected from the Road Safety Criteria Framework, ensuring both site-specific insights and cross-pilot comparability.

3.3.3 INTEGRATION DATAFLOW: FROM EDGE TO USER INTERFACE

The validation and deployment of AI tools within EvoRoads rely on a robust and adaptable dataflow architecture that ensures seamless integration from field data collection to end-user interaction. This process is designed to support diverse hardware environments and computational requirements while maintaining responsiveness and data fidelity.

3.3.3.1 BASIC VALIDATION DATAFLOW

At a foundational level, the EvoRoads dataflow follows a modular sequence:

1. **Data collection at the edge** – Real-time sensing is performed via connected vehicles, drones, fixed sensor installations, or micro-mobility kits.
2. **Data engineering pipeline** – Initial data pre-processing includes filtering, normalisation, tagging, and compression.
3. **Data storage** – Processed data are temporarily stored locally or forwarded to cloud infrastructure, depending on system architecture.
4. **AI algorithm execution** – Pre-processed data are directed to AI modules that detect anomalies, identify road hazards, or generate predictions.
5. **Aggregation by orchestrator** – Outputs from AI modules are collected and contextualised by a central orchestrator (e.g., Digital Twin).
6. **User interface output** – Results are presented via dashboards or alerting mechanisms tailored to maintenance operators, planners, or public authorities.

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 data lifecycle ensures traceability and consistency across use cases, while allowing adaptation to local constraints. For example, in resource-constrained environments, only key metadata (rather than full sensor feeds) may be transmitted, ensuring operational efficiency without sacrificing decision value.

In a typical use case, a drone patrols a rural road segment and captures high-resolution images. Onboard or nearby edge hardware runs an AI model to detect potholes. Once detected, only geo-tagged event metadata (e.g., location, severity) is sent to the cloud. The Digital Twin aggregates this input with other spatial data, updates the road status, and triggers a visual alert on the operator dashboard. This strategy minimises bandwidth requirements and enhances responsiveness.

3.3.3.2 EXPANDED TECHNICAL FLOW

In many cases, engineering and storage occur in two stages. A lightweight pre-processing step is performed on edge devices (e.g., Raspberry Pi or embedded sensor logic), followed by additional engineering and archival operations in the cloud. This two-tiered approach balances local processing efficiency with the need for centralised storage, analytics, and cross-pilot interoperability.

Similarly, algorithm execution may be either edge-based or cloud-based, depending on computational load and transmission feasibility. Lightweight classifiers (e.g., vibration anomaly detection) may operate entirely on-device. In contrast, large neural networks, such as those used for high-resolution image analysis, may require GPU-backed cloud infrastructure. When bandwidth constraints exist, inference is performed locally and only results are sent upstream. Conversely, if raw data can be uploaded efficiently, centralised processing may enable multi-sensor fusion and batch optimisation.

This adaptability is essential for EvoRoads' objective to deploy AI tools in both dense urban environments and remote rural contexts, ensuring a common operational framework without imposing uniform technical requirements.

3.3.3.3 APPLIED EXAMPLES FROM EVOROADS TECHNOLOGIES

The EvoRoads pilots offer diverse scenarios that illustrate how this layered architecture is applied in practice, based on varying resource, latency, and data-size considerations.

In **Latvia**, the micro-mobility monitoring kit installed on scooters combines a Raspberry Pi with camera and accelerometer modules. The accelerometer readings (used to infer road surface roughness) are processed in real-time on the device, and minimal JSON messages can be transmitted near real-time to the cloud. These messages include timestamped location and vibration values. The lightweight nature of the data and the low computational overhead allows all processing to occur on the edge. In contrast, images collected by the camera are not streamed due to size and bandwidth limitations. Instead, selected images - processed locally using an embedded AI model - generate pothole detection events, which are then forwarded to the Digital Twin for mapping and historical analysis.

In **Italy (Turin)**, high-resolution cameras and LiDAR sensors are used in urban areas near hospitals and pedestrian crossings. Because of the density and volume of raw data, images are processed locally to extract pedestrian and vehicle trajectories. However, due to the complexity of predictive behaviour models (e.g. identifying risk-prone crossings), some components are offloaded to cloud servers where aggregated pedestrian-vehicle interactions are assessed. These outputs feed into nudging displays or dashboards, which notify users or recommend infrastructural interventions.

In **Spain (indicatively in Santa Oliva)**, fixed sensors with LiDAR and temperature modules are deployed at road edges. These sensors detect objects and ice risk via real-time local inference. Given their fixed nature and proximity to power and network infrastructure, they synchronise periodically with the central platform. The results are then embedded into messages that may inform connected vehicles or trigger alerts via local signage.

In **Romania**, where drones monitor rural infrastructure, the imagery is too heavy to transfer continuously. AI models trained to detect cracking, or potholes are deployed on an intermediate local server that connects to the drone base station. Only tagged event metadata is sent to the cloud, saving bandwidth and enabling efficient dashboard updates.

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

These examples showcase how EvoRoads adapts data routing, processing, and user interaction according to operational context, maintaining real-time awareness while respecting infrastructure constraints.

3.3.4 FROM CLOUD TO USER: INFORMATION FLOW VIA THE EVOROADS DATASPACE AND USER INTERFACES

Once AI-driven insights are generated - either locally at the edge or centrally in the cloud - they enter a critical phase: transformation into actionable information for human users. In EvoRoads, this process is orchestrated through a federated infrastructure comprising the Safe Mobility Data Space (SMDS) and the front-end interfaces developed in Task 1.5. Unlike linear data pipelines, the information flow is often initiated by users themselves. When a stakeholder - such as a maintenance operator or municipal planner - interacts with a dashboard or mobile application, a request is sent to the backend of the interface, which communicates with the Digital Twin orchestrator. The orchestrator identifies the relevant AI component to consult, gathers the necessary data inputs from the SMDS or live sensor feeds, and feeds these into the selected algorithm. Once the result is returned, the orchestrator aggregates and reformats it into a human-readable output and redirects it to the requesting interface. This responsive, user-driven architecture ensures that data insights are not only computed but delivered in the right form, to the right people, at the right time.

The SMDS plays a central role in enabling this interaction. It supports both the archival of AI outputs for retrospective analysis and the streaming of live, event-driven information to subscribing components. Data moving through the SMDS are structured, standardised, and semantically enriched to allow interoperability between platform modules. Once AI components generate outputs - such as pothole detections, infrastructure deterioration scores, or behavioural risk alerts - these are pushed to the front-end interfaces: either dashboards or mobile applications. Dashboards provide layered visualisations, adapted to both user roles and pilot areas, supporting operational focus and geographic relevance. For example, a city operator in Turin might view real-time alerts on crossing safety, while a traffic analyst in Riga examines risk cluster heatmaps. In parallel, mobile applications provide behavioural nudging and real-time alerts to vulnerable road users (VRUs), including pedestrians, cyclists, and micro-mobility users, based on their movement and proximity to identified risks.

EvoRoads has been structured around a clear understanding of its end-user profiles, ensuring that each group receives tailored insights. Infrastructure maintainers access maintenance alerts and task scheduling tools that help allocate repair crews efficiently. Municipal and regional authorities view summary indicators and policy-relevant dashboards that support long-term planning and budget allocation. Traffic engineers and safety analysts engage with geospatial tools and behavioural data to understand where interventions are most needed. Finally, VRUs receive contextual feedback directly on their mobile devices to promote safer choices in real time. In some scenarios, AI-derived outputs also feed into external systems, such as roadside variable message signs or automated maintenance logs, thus completing the loop from sensing and interpretation to action. By linking user interfaces, cloud intelligence, and orchestrated logic through a semantically harmonised dataspace, EvoRoads transforms complex data flows into precise, usable knowledge embedded in real-world safety operations.

3.3.5 RELATED WORK AND NEXT IMPLEMENTATION STEPS

While this deliverable has outlined the preliminary validation strategy and integration vision for the AI components developed in EvoRoads, further architectural and technical detailing will be explored in the upcoming project deliverables. Specifically, **Deliverables 3.1 (M15) and 1.3 (M18)** will provide a more in-depth analysis of the data pipelines that enable the journey from edge-collected sensor data to end-user intelligence. These deliverables will address the design and deployment of the **Data Acquisition Platform**, the integration logic within the **Digital Twin**, the semantic structure and interoperability mechanisms of the **Safe Mobility Data Space**, and the user-facing architecture of the **Dashboards**. These four layers form the operational backbone of EvoRoads, and their cohesive interaction will determine how effectively raw inputs are transformed into actionable, trusted information. In parallel, **Deliverables D1.3 (M18) and D1.4 (M34)** will focus specifically on the integration of the AI artefacts described in the present document, establishing the APIs, data formats, and orchestration logic needed to embed each model into a unified and modular platform. These integration points will be the technical enablers that bind together the intelligence layer of EvoRoads and ensure functional, operational deployment in real-world pilot conditions.

Looking ahead, a critical milestone in the implementation phase will be the definition of user stories for the dashboards and mobile interfaces. Each user story will describe a scenario in which a specific question - relevant to infrastructure maintenance, safety evaluation, or road user behaviour - is posed by an end-user. These questions will then be mapped to one or more of the developed AI components, identifying whether they can be addressed in isolation or require composite inference across multiple models. This mapping will provide the foundation for dashboard interaction logic that is both intuitive and operationally aligned with user needs. Alongside this, the exploration of available datasets, which has been ongoing throughout the first year of the project, will continue. This activity supports both the refinement of AI models and the validation of their performance in real-world settings. The upcoming second version of the Data Management Plan (DMP) will capture these dataset-related developments, including data sources, processing workflows, ethical considerations, and provisions for reuse. Together, these next steps mark the transition from experimentation to implementation, setting the stage for a fully integrated, data-driven road safety platform.
Further on, these integrated assets will provide the seeds for potential post-project commercialisation. The transition from integrated platform assets to business assets will be further developed in WP5: Dissemination, scale-up, standardization and exploitation.

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 CONCLUSION

Deliverable D2.1 has documented the development of the eight AI-enabled infrastructure-monitoring tools now forming the modelling backbone of EvoRoads (excluding BEIA's pothole detection tool, LINKS' algorithms, and CTAG's network connectivity classification algorithm, which will be reported in subsequent deliverables). Together they operationalise Objective 2 by detecting deterioration at multiple spatial scales, fusing heterogeneous data streams, and issuing actionable warnings. Moreover, the applicability of these tools to EvoRoads pilots and their integration to the EvoRoads platform is discussed. All tools will be revisited in D2.3 - "Advanced infrastructure monitoring and predictive maintenance tools".

The presented tools illustrate how EvoRoads' offerings ultimately translate into fewer crashes—directly supporting the EU's Vision Zero ambition. This is achieved through a combination of innovative models giving both long-term insights and tactical, on-the-edge alerts that can proactively warn road users.

Across partners, several best practices were followed: (i) synthetic data generation to bootstrap models where labelled imagery is scarce; (ii) lightweight YOLO/UNet backbones for on-device inference; (iii) probabilistic clustering (GMM) to provide confidence scores for planners; and (iv) early alignment with the EU AI Act compliance checker.

Looking ahead, the focus will shift from tool development to operational validation within the four pilots. Integration into a unified, user-centric platform will enable seamless data flow from sensor capture to actionable dashboard insights, completing the loop from detection to intervention. Ultimately, these advances set the foundation not only for safer roads but also for post-project commercial exploitation, ensuring that EvoRoads' innovations can scale and contribute to systemic, long-term improvements in road safety across Europe.

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

ANNEX

APPLYING THE HAZARD INFRASTRUCTURE TOOL TO THE SANTA OLIVA PILOT [IDIADA]

TOOL & PILOT OVERVIEW

The Santa Oliva pilot takes place within IDIADA facilities located near Barcelona, where IDIADA´s clients (OEMs and Tier 1) test autonomous and connected vehicle's solutions. As a global homologation and system development partner, we maintain our test tracks to be always at the highest quality. Any foreign objects or severe weather conditions (heavy rain, fog) damaging the roads creates a risk for IDIADA road users (clients, employees as well as visitors).

IDIADA together with EUT are working on establishing and implementation of a hazard aware infrastructure monitoring (HAIM) tool. EUT following the needs and specifics of Santa Oliva pilot, has been developing a low-cost device designed to measure environmental road conditions, identify objects on the road and assess asphalt quality. Environmental road conditions measurement includes detection of hazards such as debris, water puddles/accumulation and spotting dangerous weather conditions such as fog or rain. Identifying objects aims to detect vehicles, humans, animals, foreign objects (e.g. rocks, plastic pieces, tree branches, traffic cones etc.). The device includes a thermal sensor for measuring asphalt and ambient temperature as well as an ultrasonic sensor to determine vehicle travel lane. For asphalt characteristics evaluation, the tool integrates a LiDAR sensor to scan the road section, generating 2D maps that will allow asphalt wear and/or damage recognition.

To achieve above mentioned measurements and analysis, EUT makes use of a combination of artificial intelligence (AI) algorithms, convolutional neural networks (CNN) and computer vision technologies. Along with the necessary electronics-sensors, the Jetson Nano Orin (JNO) system is being used as the main controller of the tool. JNO system is programmable in several popular languages, including Python and C++, and is fully compatible with AI and computer vision frameworks like (YOLO) used by TensorFlow. This way, hazard aware advanced infrastructure monitoring tool will be equipped to make precise measurements and process the data efficiently, contributing to the improvement of road infrastructure and maintenance.

The AI algorithms used in hazard aware advanced infrastructure monitoring tool to be tested in Santa Oliva pilot are:

- Debris Detection (IDIADA Cameras): *EUT will make use of already installed security cameras in IDIADA premises, Given the restrictive nature of these cameras, EUT will utilize original images captured by the camera itself for the creation of datasets for object detection on the road functions. EUT will showcase an initial model that is capable of recognizing three types of objects using datasets specifically created for IDIADA cameras.*
- Pedestrian and Animal Detection.

These algorithms to be implemented in the Jetson Nano device are being programmed in Python to process and interpret the data collected by various sensors, ensuring efficient real-time information handling. Additionally, communication protocols will be established to enable data exchange with IDIADA's external systems, facilitating real-time monitoring and analysis. As the project progresses, further optimizations will be made to enhance performance and ensure reliable integration with the external platforms.

Steps required for the tool applicability

A specific road segment of the IDIADA road's network has been selected for the EvoRoads pilot prior to project start, which is frequently used by passenger cars, special vehicles and trucks, as well as vulnerable road users (pedestrians), to move between the offices/workshops and the IDIADA Proving Grounds. This road segment has been selected since it is being used 24/7 by different vehicle types at a normal driving behaviour, as well as being under coverage of the private 5G network deployed at IDIADA. It is a bidirectional somehow curvy road with thick trees which has possibilities in terms

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 infrastructure deployment's locations (e.g. sensors, dynamic road panels, etc.), use cases and scenarios. This road is marked as two (2) in the Figure 26.

However, when EUT visited IDIADA facilities, together we have examined another possible road to implement the Hazard aware infrastructure monitoring tool. This road, which is marked as one (1) in the Figure 26, is also being used by passenger cars, special vehicles and trucks as well as vulnerable road users (pedestrians). Road number 1 is located just in front of the offices and has an exit to "Dynamic Platform" test track used by our clients. The road is active 24/7 used by different type of vehicles and similar to the road segment 2, it is under coverage of the private 5G network deployed at IDIADA. The location of the HAIM sensors shall be close to the selected road segment, pointing to the asphalt from both sides. Unlike road segment 2, road segment 1 has a larger sidewalk and have easier access to electricity source to deploy Hazard aware infrastructure monitoring tool. Larger sidewalk allows better and more flexible implementation still allowing pedestrians to walk. Together with EUT, it is decided to implement the HAIM tool primarily in road segment 1 instead of 2. Yet, we will still implement dynamic signals from EUT (WP3) in road 2, leveraging the curvature of the road for better visualization.

Figure 26: IDIADA selected road segments 1 and 2

EUT will process all sensor´s data locally in the HAIM tool and the output will be transmitted via 4G/5G (as a plan B with Wifi) connection to the IDIADA EDGE server. From IDIADA Edge server, this information will be communicated with EVOROADS platform.

Other potential AI tools from EvoRoads Partners

IDIADA had preliminary meetings on the applicability of the other potential AI tools from EvoRoads partners to Santa Oliva pilot. Two possible feasibility studies are identified in collaboration with DTU´s Predictive Maintenance and post-impact repair analysis solution as well as INDRA´s (Madrid pilot) GIS-based digital twin tool to monitor urban and rural road conditions in real time.

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

INDRA (Madrid Pilot): Madrid pilot establishes a digital ecosystem using a GIS-based digital twin to monitor urban and rural road conditions in real time. Their system enables rapid and accurate detection of pavement deterioration, improving maintenance planning and execution. As IDIADA we need to check our CAV (Connected and Autonomous Vehicle) and verify our ability of capturing and providing the necessary data to INDRA to execute their solutions in Santa Oliva pilot. This assessment is currently being executed prior to committing to any formal collaboration for the upcoming months. If the outcome is favourable, we will determine the appropriate road or test track for implementation.

DTU (Predictive maintenance and post-impact repair analysis solutions): DTU is working on a predictive road maintenance system that is AI based which prioritizes maintenance interventions analysing road conditions, maintenance cost, ride quality. Together with DTU we have conducted several meetings and shared information on how currently the maintenance is being done in IDIADA test tracks. We have exchanged insights on maintenance frequency, workforce needed, what triggers a maintenance. We have detected a use-case where DTU tools can be possibly implemented in our "Straight Line breaking Surface" track. In this track we have three different road textures available as Concrete, Asphalt and Ceramic (see Figure 27). In the upcoming months we will do a feasibility study with fast pass/fail approach to assess whether we can successfully implement DTU predictive maintenance and post impact repair analysis solution to Santa Oliva pilot.

Figure 27: IDIADA Straight Line Breaking Surfaces Test Track

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