

The IVORY project



➤ IVORY:



"AI for Vision Zero in Road Safety"
ivory-network.eu

➤ Partners:

- 4 Universities
- 8 Non-academic partners
- 13 Associated Partners
- 10 Countries

➤ Duration of the project:

48 months (November 2023 – October 2027)

➤ Framework Program:

This project has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No 101119590




Objectives

To develop a **new framework**
for the integration of AI in road safety

To create a new generation of leading researchers
to address the EC 'Vision Zero' strategy (eliminating
traffic fatalities by 2050)

- **Responsible and fair AI** for road safety;
- **Safe road users** and human-vehicle-environment interaction by means of AI;
- Scalable and equitable AI technologies for **proactive infrastructure safety management**;
- A sustainable **learning, knowledge sharing and networking platform** on AI for road safety.



Seeing Road Lanes in 3D



1/2 **The challenge** | Aristotelis Tsoutsanis · NTUA & OSeven [Secondment @ Haskoning, NL]

10,500

computer-generated road scenes with exact 3D lanes

9.6 M

trained parameters; the large AI models stay frozen

THE PROBLEM

Lane-keeping assistance and driver-behaviour analysis need to know **where lanes are in 3D**, including on hills, curves and far ahead, from one ordinary camera.

OUR APPROACH

Reuse large ready-made **foundation AI models** as building blocks. One finds the road from the word "road", one estimates distance and one describes the scene. Only a small add-on is trained to turn their output into a bird's-eye view and 3D lane lines.

THE DATA

A synthetic driving dataset (Baidu Apollo) of highways and cities in varied weather and lighting. Next, about 200,000 real-world frames (OpenLane, from Waymo).



Seeing Road Lanes in 3D

2/2 What we found | Aristotelis Tsoutsanis · NTUA & OSeven [Secondment @ Haskoning, NL]

0.60 → 0.93

F-score (lane-matching accuracy) after step-by-step improvements

~0.5 m

sideways error for distant lanes

RESULTS

Testing one design change at a time lifted accuracy from **0.60 to 0.93**.
The model traces lanes at dusk, at night and on banked curves.

CONCLUSION

Frozen foundation models plus a small trained module give **accurate 3D lanes without hand-made road labels**.

THE BIGGER PICTURE

Synthetic success is a first step. The real test is **real roads**, with night, rain, merges and intersections.



Can AI Inspect Our Roads?



1/2 **The challenge** | Júlia Alves Porto · NTUA & iRAP [Secondment @ FRED Engineering, IT]

27

road features we tried to measure automatically

50

repeated train/test runs per model

THE PROBLEM

Road-safety inspectors watch road videos and code features such as guardrails, curves and land use by hand. It takes minutes per kilometre, so most roads are rarely assessed.

OUR APPROACH

Let AI do the coding, using image recognition on street photos plus open map and traffic data. Then test whether these automatic features predict **risky driving** as well as the expert-coded ones.

THE DATA

Harsh braking and acceleration recorded by drivers' smartphones on each road segment, street-level images, OpenStreetMap and TomTom traffic data, and expert inspections (iRAP) for comparison.



Can AI Inspect Our Roads?



2/2 **What we found** | Júlia Alves Porto · NTUA & iRAP [Secondment @ FRED Engineering, IT]

0.173 vs 0.174

prediction error: expert vs. automated features

19 of 27

features AI could measure reliably

RESULTS


Both models predict harsh events **almost equally well**. Speed and traffic volume matter most. Road features add a smaller signal, such as guardrails, speed limits, roadside poles and walls, and sharp curves.

CONCLUSION

Automated road coding **matches, but does not yet beat**, expert inspection, at a fraction of the effort.

THE BIGGER PICTURE

How we drive explains more than what the road looks like, but the road still matters. Powerful image AI is still under-used in road-safety research.



Which Junctions Are Risky?

1/2 **The challenge** | Simone Paradiso · NTUA & OSeven [Secondment @ Verne/P3M, HR]

Every second

driving data from smartphones

Every week

a new risk snapshot of the network

THE PROBLEM

Crashes are rare, recorded late and often incompletely, so waiting for them is a slow way to find dangerous spots. Everyday driving data could flag risk **before** crashes happen.

OUR APPROACH

Treat the city as a network, with junctions as points and roads as links. A **Graph Neural Network** learns a fingerprint of each junction from driving there and at its neighbours, then groups junctions into safer and riskier ones. No crash labels are needed to train it.

THE DATA

Smartphone driving data (speed, harsh braking and acceleration) matched to the road map and summarised weekly, with crash records to check the results.



Which Junctions Are Risky?

2/2 **What we found** | Simone Paradiso · NTUA & OSeven [Secondment @ Verne/P3M, HR]

Clearer groups

than plain driving statistics (higher silhouette score)

Still open

the link between risk groups and recorded crashes

RESULTS


Network-aware fingerprints sort junctions into **clearer, more distinct groups** than raw statistics. A new multiscale version learns how far around each junction to look. The overlap with crash hotspots is still weak.

CONCLUSION

Driving behaviour reveals **structured, city-wide risk patterns**. Proving that they predict crashes needs better crash data.

THE BIGGER PICTURE

Risk at one junction depends on its neighbours, and an AI is only as trustworthy as the real-world data used to check it.



Drowsiness in a Heartbeat



1/2 **The challenge** | Aristotelis Styanidis · NTUA & CardiID [Secondment @ VTI, SE]

40

drivers in the current lab datasets

269

drivers in a new 12-study collection (VTI, Sweden)

THE PROBLEM


Drowsy driving causes many serious crashes, and drivers are poor judges of their own sleepiness. A sensor that reads the heart could **warn them in time**.

OUR APPROACH

Deep learning reads the raw heartbeat signal (ECG) and finds its own warning signs, compared with a model built on classic heart-rate-variability measures. Every model is tested on **people it has never seen**.

THE DATA

Two datasets of 20 participants each (heart, brain and eye signals), with sleepiness self-rated from 1 (very alert) to 9 (fighting sleep). They are now being extended with a much larger collection.



Drowsiness in a Heartbeat

2/2 **What we found** | Aristotelis Styanidis · NTUA & CardioID [Secondment @ VTI, SE]

0.76

AUC on unseen drivers (0.5 = chance, 1 = perfect)

~61%

accuracy, with large person-to-person differences

RESULTS


The heartbeat **does carry a drowsiness signal**, but performance varies a lot between people, and very sleepy moments are rare in the data (accepted at IEEE-EMBS BSN 2026).

CONCLUSION

One-size-fits-all models are not yet reliable. **Personalisation and much larger datasets** are needed.

THE BIGGER PICTURE

In safety systems, an AI that **knows when it is unsure** is as valuable as one that is right. That is the focus of the next step.



Scientific and Societal Impact



- **Unlocking the full potential of AI for road safety**, so that new opportunities for global road safety impact can emerge:
 - **15 highly skilled researchers** meeting industry and policy stakeholders' needs;
 - **A learning and networking platform** on AI for road safety, including open, interdisciplinary and interactive training material, as well as a 'social network' of researchers;
 - **New applications of AI for road safety**, including data protocols and cutting-edge analytical models;
 - **Taxonomies** for designing responsible, equitable and efficient AI for road safety.



Future Challenges

- **Balancing AI developments in road safety**, in order to bridge the gaps between engineering, technology & society, and ensure equal opportunities for all countries around the globe.
- **Meaningful human-AI interaction:** AI can only learn from existing data, but lacks the critical thinking and inspiration needed for making a 'breakthrough'; without steering, AI can induce bias.
- **Systematic and sustainable intersectoral collaboration:** breaking the 'silos' between sectors (academia, industry and policy) and disciplines involving AI (engineering, data science, ethics of technology).



