

Semantics for Transport and Logistics 2026 – Sem4Tra

Shahrom Sohi¹, Julian Rojas², Ghislain Ateazing³ and Cornelis Bouter⁴

¹Vienna University of Economics and Business, Welthandelsplatz 1, Vienna, 1020, Austria

²Ghent University – IDLab – imec, Technologiepark-Zwijnaarde 126, 9052 Gent, Belgium

³ERA, European Union Agency for Railways, 120 Rue Marc Lefrancq, 59307 Valenciennes, France

⁴TNO, The Netherlands Organisation for Applied Scientific Research, 2595 DA, The Hague, The Netherlands

Abstract

This is a summary of the Seventh International Workshop on Semantics for Transport and Logistics (Sem4Tra 2026), co-located with the 22nd International Conference on Semantic Systems (SEMANTiCS 2026) in Ghent, Belgium.

Summary

Transport and logistics data in Europe is spread across National Access Points, public transport operators, railway infrastructure managers, regulatory registers and the message models of individual logistics companies, each with its own identifiers, granularity and modelling practices. European regulation increasingly requires this data to be published in standardised and transparent ways, from the ITS Directive and Delegated Regulation (EU) 2017/1926 for multimodal travel information, to the Railway Interoperability Directive for vehicle authorisation and registration, and the Data Act for the semantic assets used in data sharing. This year's workshop examines what it takes, in practice, to turn these requirements into data that can actually be found, exchanged and combined.

Its five contributions, read together, tell a single story: shared standards and reference models are necessary but not sufficient for interoperability, and what ultimately decides whether data can be exchanged and combined is how the artefacts connecting heterogeneous systems are governed, maintained and validated over time. The story begins at the level of metadata. Martineau, Scrocca, Lubrich, Bureš and Beernaerts show how the controlled vocabularies of mobilityDCAT-AP, the metadata specification used by European National Access Points, are maintained within NAPCORE through a reuse-first strategy, SKOS representations with persistent URIs and a lightweight spreadsheet-driven workflow that lets domain experts contribute while producing both RDF and human-readable documentation. Vocabularies, in their account, are not a supporting detail but first-class semantic assets whose lifecycle must be managed deliberately. Mulder, van den Berg, Wijbenga and Bouter carry the same idea from vocabularies to mappings. Starting from logistics SMEs that cannot converge on a common model such as the Open Trip Model, and building on the notion of Eventual Interoperability, they propose an architecture in which conceptual alignment is captured at design time in SSSOM⁺, executable transformation logic is expressed at run time in JSONata⁺, and an LLM-assisted component connects the two through elicitation, generation and validation, as demonstrated in a transport-order matching use case.

The railway sector offers two complementary perspectives on the same theme. Vanden Auweele, Olsen, Berge and Korsmo report on building a production knowledge graph for the Norwegian infrastructure manager Bane NOR, with roughly 13 million triples across seventeen ontologies and over five thousand kilometres of track, from which a European Register of Infrastructure (RINF) export is

Sem4Tra 2026: Seventh International Workshop on Semantics for Transport and Logistics, co-located with SEMANTiCS 2026: 22nd International Conference on Semantic Systems, September 15, 2026, Ghent, Belgium

EMAIL: shahrom.sohi@wu.ac.at (S. Sohi); julianandres.rojasmelendez@ugent.be (J. Rojas); ghislain.ateazing@era.europa.eu (G. Ateazing); cornelis.bouter@tno.nl (C. Bouter)

ORCID: 0009-0000-0735-1645 (S. Sohi); 0000-0002-6645-1264 (J. Rojas); 0000-0003-1562-6922 (G. Ateazing);

0000-0002-5448-0543 (C. Bouter)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

generated directly. Their experience with SPARQL Anything, SHACL rules and GeoSPARQL shows that the Semantic Web stack can carry much of this integration, but that railway infrastructure is fundamentally topology- and linear-referencing-centric, so a significant part of the spatial logic still has to live outside the graph. Where their work concerns the physical railway, Duhoux and Atemezing turn to its regulatory processes. Their VPA ontology models how verified permissions come into being, from requests, cases and evidence to compliance checks against sections of legal requirements, restrictions and time-bounded permissions, generalising vehicle authorisation, registration and safety certification across the ERA registers. Validated with competency questions and SHACL shapes and already reused in the ERA ontology, VPA sits between policy languages such as ODRL, which state what a permission allows, and the administrative reality of how it is granted.

The programme closes by testing the premise directly on published data. Miclaus and Sohi ask whether GTFS and NeTEx feeds, both rooted in Transmodel and mandated through National Access Points, are actually comparable in practice. Using a Minimum Viable Denominator of stops, routes and service calendars across fifteen European countries, they find that stops align best, routes less well and calendars least, with close agreement where one organisation publishes both formats and breakdowns caused by fragmented publication, inconsistent identifiers, repurposed fields and differences in scope. Their conclusion that the binding constraints are organisational rather than representational echoes the thread running through the whole workshop: from vocabularies and mappings to ontologies, knowledge graphs and data feeds, interoperability in transport and logistics depends less on agreeing on a single model than on governing identifiers, provenance, alignments and publishing practice well.

The website of the workshop is: <https://semantic-transportation.github.io/sem4tra-kg-website/>