High-definition road map generation from mobile mapping data: a case study on the Tangenziale di Napoli

Enrico Breggion[1], Maria Alicandro[2], Caterina Balletti[1], Francesco Guerra[1], Andrea Martino[1], Nicole Pascucci[2], Giovanni Pugliano[3], Sara Zollini[2], Donatella Dominici[2]

[1] Università Iuav di Venezia, Santa Croce 191, Venezia, Italy – ebreggion; balletti; guerra2; amartino @iuav.it
[2]Università degli Studi dell’Aquila, Piazzale E. Pontieri, 1, Monteluco di Roio, L’Aquila, Italy - maria.alicandro; nicole.pascucci; sara.zollini; donatella.dominici@univaq.it.
[3]Università degli Studi di Napoli Federico II, Via Claudio 21, Napoli, Italy - giovanni.pugliano@unina.it.

Keywords: HD map, MMS, LiDAR, segmentation, OpenDRIVE

1. Abstract

High-definition (HD) maps provide lane-level road geometry and semantics for vehicle localization, scenario-based simulation, and safety-oriented validation in connected and automated driving contexts. In Italy, these objectives are consistent with the MOST – Centro Nazionale per la Mobilità Sostenibile programme, whose Spoke 7 addresses connected and automated mobility and smart infrastructures.

This paper presents a structured workflow for generating an ASAM OpenDRIVE HD map from mobile mapping system data in a geometrically and operationally complex highway environment. The pipeline was tested on an approximately 10 km segment of the Tangenziale di Napoli, including dual carriageways, ramps, interchanges, and multiple tunnels. Data were acquired through repeated runs with the GAIA M1 MMS, integrating dense LiDAR point clouds and spherical imagery.

Given the size and limited direct usability of raw MMS outputs for corridor-scale annotation, the workflow focuses on upstream preprocessing: trajectory management, georeferencing, DTM and orthophoto generation, and point-cloud classification/segmentation of road-related elements. Automated extraction strategies were explored as support tools, while the final OpenDRIVE network was compiled through the controlled integration of raster underlays, segmented point clouds, and vector candidates in a dedicated authoring environment.

The contribution is methodological: the paper documents a practical and reproducible workflow for HD map generation in a tunnel-rich and topologically complex motorway scenario. It highlights the role of preprocessing, the limits of automation for schema-complete map production, and the quality-control steps required for consistent OpenDRIVE restitution.

2. Introduction

HD maps have become a core enabling layer for automated driving because they encode lane-level road geometry and semantic relationships that can act as strong priors for localization and behavior planning—particularly in conditions where on-board sensing and GNSS observability are challenged, such as tunnel segments and complex junctions (Andert et al., 2025; Geiger et al., 2012). This is especially relevant for highway corridors in dense metropolitan regions, where merges/diverges, ramps, and tunnels coexist and where robust, repeatable digital representations are needed for both engineering testing and infrastructure-level monitoring.

Within the Italian research and innovation framework, MOST Spoke 7 explicitly targets CCAM and smart infrastructure technologies intended to improve safety and traffic efficiency through advanced sensing and digital integration (Centro Nazionale per la Mobilità Sostenibile, 2021). In such multi-stakeholder contexts (public institutions, research labs, industrial partners), open standards are not merely a distribution format; they are a practical requirement to enable heterogeneous toolchains and long-term reuse (Greffrath et al., 2017). The adoption of ASAM OpenDRIVE follows this rationale: OpenDRIVE is a widely used, XML-based standard for static road network descriptions and supports representation of reference-line geometry, lane structures, road linkages, junctions, and objects in a structured schema (ASAM e.V.; Diaz-Diaz et al., 2022).

This case study focuses on generating an OpenDRIVE HD map from MMS data for a highway-like scenario in the Naples area (Italy), where the field conditions (traffic, safety constraints, tunnels) impose both logistical and methodological constraints on acquisition and validation. In particular, acquisition was planned to maximize completeness while minimizing traffic-induced occlusions: repeated runs were conducted in low-traffic early-morning windows, including one run per carriageway and one dedicated run for ramps and interchanges. The restitution phase was designed around a realistic production constraint: raw MMS datasets are too large to be efficiently exploited “as-is” for lane-level authoring, and preprocessing must therefore compress and structure the data into products that support fast, reliable interpretation and editing (Pang et al., 2025).

The contribution of the paper is primarily methodological and application-oriented: it documents a reproducible acquisition strategy for a tunnel-rich highway environment; a pragmatic preprocessing chain combining classical classification with exploratory machine-learning tests for feature extraction; and a transparent account of the integration steps required to compile a consistent OpenDRIVE network using orthophoto/DTM/segmented point clouds as mapping underlays, including the quality-control actions needed to meet topology and semantics expectations.

3. State of the art

In contemporary automated driving practice, “HD maps” typically refer to lane-level representations with centimeter-level ambition, combining geometric fidelity (lane boundaries, widths, curvature, elevation) with semantic content (markings, signs, signals, traffic rules, and relational/topological links) (Kokkinogenis et al., 2019). Surveys consistently position HD maps as a complementary information layer that can improve localization and downstream decision-making, while also serving simulation and validation workflows. This dual role—runtime support and offline verification—explains why HD mapping has split into two partially different ecosystems: industrial pipelines oriented to scalable production and frequent updates, and research and infrastructure pipelines oriented to

openness, reproducibility, and integration with public-sector constraints (Schott et al., 2021).
From the production standpoint, the literature largely agrees that HD-map construction is best described as a multi-stage pipeline: raw data collection, followed by offline/post-processed data cleaning, alignment and georeferencing, semantic feature extraction, and finally map compilation in a formal schema with topology and semantics enforced (Park et al., 2020; Boyko et al., 2011). Importantly, multiple sources emphasize that the hardest part is not the initial detection of individual elements but the “completion step”: turning heterogeneous detections into a complete and consistent map model with correct connectivity, junction logic, and rule-relevant semantics. This challenge becomes sharper in OpenDRIVE, where networks must satisfy explicit linkage rules and where junction modeling requires coherent structure, not just geometry (Althoff et al., 2018).
A second, equally practical bottleneck is data scale. Point clouds are attractive because each point carries 3D information, potentially enabling more precise detection than purely image-based approaches; however, large-scale point-cloud pipelines are constrained by data transport, memory usage, and annotation complexity. This is why many production-grade workflows prioritize early-stage classification/segmentation and thinning, so that later steps are performed on structured subsets rather than on monolithic raw datasets (Bardak et al., 2025). In addition, image-derived annotations are known to be vulnerable to illumination variability and shadows, which can degrade detection and annotation quality—an issue that becomes acute in tunnel portals where lighting transitions are abrupt.
Automation strategies in HD mapping span a spectrum. Classical pipelines frequently exploit intensity cues for road marking extraction on LiDAR/mobile mapping point clouds, using thresholding and filtering steps to create vector candidates; these remain relevant baselines and appear repeatedly in academic and applied work. More recently, learning-based approaches have expanded both offline extraction and online “local HD map” construction using vehicle sensor streams. Surveyed trends include semantic segmentation, lane graph construction, and incremental updating from crowdsourced fleets. These approaches can reduce post-processing load in constrained scenarios, but they often deliver partial HD map content or require significant training data, domain adaptation, and careful system integration for production robustness (Chiang et al., 2022).
At the industrial level, vertically integrated solutions exist in which HD map layers are produced and maintained at scale, often leveraging fleet data and near-continuous update mechanisms. Here, the model shifts from “survey once, process offline” to “detect changes continuously and update incrementally,” with change detection and freshness highlighted as differentiators. Examples of commercial offerings include Orbis Maps for automated driving by TomTom and HD map products from ZENRIN, both marketed around lane-level positioning and automated driving support (HERE Technologies, 2019; TomTom, n.d.; ZENRIN Co., n.d.). While such ecosystems can reduce internal production burden, they are often expensive, partially opaque, and not necessarily aligned with open, reproducible academic workflows or public-sector procurement constraints. Consequently, research and infrastructure projects frequently implement heterogeneous toolchains, combining commercial and open components, and rely on open standards to preserve interoperability despite pipeline fragmentation. Despite these advances, several challenges remain open. In particular, the integration of heterogeneous data sources into schema-compliant HD maps (e.g., OpenDRIVE) still requires significant manual intervention, especially in complex highway scenarios with tunnels and interchanges. Moreover, while automation techniques can support feature extraction, they rarely address the full pipeline from raw MMS data to a topologically consistent and semantically complete map representation (Eisemann et al., 2023). In this context, this work focuses on bridging this gap by proposing a structured and reproducible workflow tailored to large-scale MMS datasets and complex infrastructure environments.

4. Materials and methods

4.1 Study corridor and experimental constraints

The surveyed corridor is a ~10 km highway segment in the Naples metropolitan area, characterized by dual carriageways, multiple ramps/interchanges, and several tunnels (Figure 1). From a scientific-test perspective, this setting is representative of real operational constraints: traffic management limitations restrict prolonged stoppages or full closures, the environment contains dynamic objects that produce occlusions in LiDAR, and tunnel segments create extended GNSS-challenged intervals that stress trajectory estimation and strip consistency.
A key methodological implication concerns the validation strategy, which had to be adapted to these constraints. Since the corridor could not be closed to deploy dense ground control or to perform repeated static surveys with higher-precision instruments , the study emphasized: redundancy via multi-pass acquisition, enabling overlap-based internal consistency checks; and trajectory-solution quality metrics from the integrated GNSS/INS processing, treated as the primary quantitative indicator of absolute positioning performance in the absence of external control. This approach is coherent with the broader recognition that HD map production at scale is constrained by logistics and by the high cost of “gold-standard” ground truth acquisition (Kalenjuk et al., 2022).

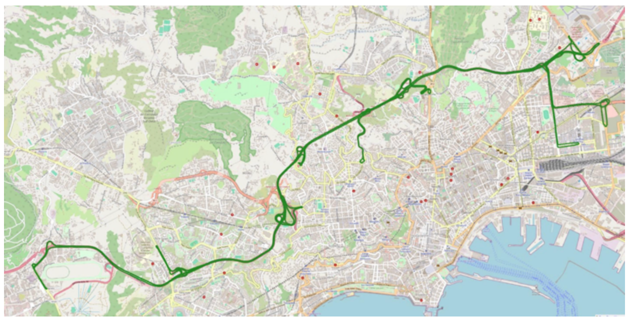

Figure 1. Napoli highway segment surveyed using the MMS.

4.2 Acquisition system and field strategy

Data acquisition was performed using the GAIA M1 mobile mapping system developed by MicroGeo (Figure 2). The platform is designed to support mobile laser scanning with a FARO Focus-class laser scanner, a Ladybug5 spherical camera system for 360° imagery, GNSS positioning and an inertial platform based on Applanix components, plus auxiliary sensors such as an odometer. Manufacturer documentation reports operational vehicle speed capability up to 100 km/h and LiDAR acquisition at about 1 million points per second (MicroGeo, 2022).
For this corridor, the driving speed during acquisition was intentionally kept around ~40 km/h to increase point density and uniformity on road-surface markings and to reduce motion irregularities in complex ramps. This choice is a practical trade-off: while high-speed acquisition is feasible according to system

https://doi.org/10.5194/isprs-archives-XLIX-B4-2026-361-2026 |

specifications, the combination of required lane-level detail and corridor complexity benefits from more conservative kinematics (Ma et al., 2018).
The survey design consisted of multiple acquisition runs: one for each carriageway and an additional run dedicated to ramps and interchanges. Surveys were scheduled in low-traffic early-morning windows to reduce congestion, minimize occlusions and moving-object artifacts, and improve operational safety. Spherical imagery was acquired jointly with LiDAR to support interpretation, especially for ambiguous or degraded markings and for fast navigation during restitution.

Figure 2. The GAIA M1 mobile mapping system installed on the survey vehicle.

4.3 Georeferencing

Mobile mapping outputs were referenced to the Italian realization of ETRS89 via ETRF2000 at epoch 2008.0 (IGMI/IGM, n.d.). GNSS differential corrections were provided through the HxGN NTRK SmartNet Services (Hexagon, 2023).
For vertical referencing, ellipsoidal heights from the GNSS/INS processing were converted to geoid-referenced heights using the Italian official geoid ITALGEO2005 (ISGeoid, n.d.). This conversion is essential for elevation profiles in HD mapping, because inconsistent height datums can introduce artificial discontinuities that propagate into OpenDRIVE elevation definitions and ramp/tunnel approach modelling (Schwab et al., 2022).
A forward-looking issue for HD-map deployment is the potential mismatch between the map reference frame and the positioning solution used by vehicles. ETRS89 (and its national realizations) is fixed to the Eurasian plate; the difference between ETRS89 and the global ITRF grows over time due to plate motion, commonly approximated at ~2.5 cm/year relative drift. PPP-based positioning-based services often provide coordinates in an ITRF realization at an epoch close to the observation date, which can create a systematic offset if the HD map is stored in a fixed-epoch ETRF realization and the vehicle localizes in a near-current ITRF epoch without proper transformation (Eurocontrol, 2017; Baasch et al., 2022). For future operational alignment, a practical recommendation is to enforce a common frame/epoch in the full stack—either transforming vehicle positions to the HD-map frame using time-dependent ETRF/ITRF transformations and velocities or maintaining the map in a time-tagged framework with explicit epoch handling.

4.4 Data management and preprocessing products

The raw dataset acquired for this case study comprised 3.526 billion points, ~72 GB of raw data, >91 GB of exported LAS point clouds, approximately 12,200 spherical images, and roughly ~1 TB of intermediate processing outputs. These figures are representative of why “raw MMS outputs” are difficult to exploit directly: storage, memory, and navigation burdens scale quickly, and point-cloud annotation is explicitly recognized as constrained by hardware limitations and data-volume effects in the HD mapping literature (Qi et al., 2017).
The preprocessing workflow generated three primary products used in restitution: an orthophoto at 4 cm GSD, derived from oriented imagery and used as a visual underlay for lane markings and traffic sign interpretation; a DTM with 5 cm grid spacing, used for elevation consistency and as support for road-surface modeling; a classified/segmented point cloud with macro-classes tailored to road environments, including road surface, terrain, guardrails and barriers, poles and sign supports, tunnel-related surfaces where relevant, and static/dynamic vehicles as separate classes.

4.5 Automated extraction tests

To assess how much of the HD-map production effort could be shifted from manual interpretation to algorithmic processing, we carried out a staged set of preliminary extraction experiments before consolidating the final preprocessing chain. The analysis focused on two families of targets that most directly impact OpenDRIVE restitution: horizontal markings (lane boundaries, edge lines, arrows, and other road-surface symbols) and a subset of roadside assets that act as geometric constraints and semantic anchors (e.g., guardrails and poles/sign supports).
A first feasibility round investigated state-of-the-art learning-based segmentation paradigms for large point clouds and LiDAR projections—specifically point-based hierarchical networks (PointNet++), efficiency-oriented large-scale semantic segmentation (RandLA-Net), kernel point convolutions (KPConv), and projection/range-image approaches exemplified by SqueezeSeg (Esmaeili et al., 2017; Roddick et al., 2020; Wu et al., 2019; Zhou et al.,2020). These experiments were intentionally treated as exploratory: while such models can deliver strong semantic labeling, they typically require non-trivial dataset preparation, labeling effort, and domain adaptation to maintain reliability across local pavement materials, marking styles, and lighting regimes. In a corridor like the Tangenziale di Napoli—where tunnels, ramps, and heterogeneous road surfaces introduce repeated domain shifts—this customization and integration burden would have dominated the schedule compared to the objectives of this paper.
After this preliminary testing, the operational workflow was consolidated around two commercial toolchains—LiDAR360MLS and TerraScan—because they provide more general-purpose, corridor-scale modules that reduce the need for custom development while still supporting the core preprocessing requirements of HD mapping. In practice, LiDAR360MLS was used to organize the large MMS dataset into manageable blocks, perform automated classification/segmentation with road-scene classes (including ground/road surface, guardrail, pole, and vehicle-related classes during testing), and refine consistency across repeated acquisitions by adjusting overlapping/revisited areas so that

multi-pass data could be treated as a coherent base for interpretation and feature extraction (GreenValley International, n.d.). In parallel, TerraScan was used to reinforce the road-surface extraction with a hard-surface-oriented classification step suited to paved-road MLS (Mobile Laser Scanner) data, improving the separation of the dominant road surface from surrounding structures and reducing ambiguity in downstream vectorization tasks (TerraSolid, n.d.).
Within this consolidated preprocessing chain, we also extracted and vectorized selected road-surface elements—most notably horizontal road markings—in those corridor sections where the point-cloud signal (density/intensity contrast) was sufficiently clear and stable (Figure 3). This choice follows established MMS-based HD mapping practice where intensity-driven candidates provide an efficient scaffold for lane-level modeling, while complex or degraded areas are deferred to more controlled interpretation during restitution.
Overall, the tests confirmed a pattern widely reported in HD-map generation research: automation is most effective when applied upstream to structure the raw MMS data, while the production of an OpenDRIVE-complete representation remains constrained by the need to enforce connectivity, junction logic, and schema-consistent semantics across heterogeneous road situations.

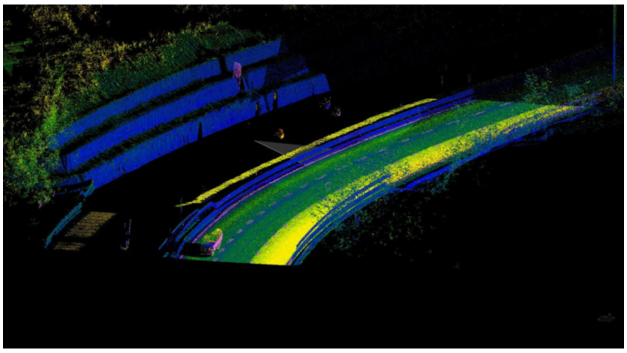

Figure 3. Representative section of the point cloud acquired by the MMS platform.

4.6 OpenDRIVE restitution and quality control

The OpenDRIVE map was compiled in MathWorks' RoadRunner environment (MathWorks, 2023) with all data expressed in metric units (m) within the UTM zone 33N projection (Geo-spatial CEN/ISO, 2019). In the adopted workflow, the pre-processed orthophoto served as a textured underlay and the DTM as a terrain base, providing visual and geometric guidance for tracing road alignments.
Although MathWorks provides automated-driving tools that can generate RoadRunner HD maps from sensor data, these machine-learning methods were used primarily for preprocessing and candidate generation. The final OpenDRIVE network was constructed through careful integration of the underlays and segmented point clouds: lane centerlines and boundaries were drawn and constrained to the DTM profile and imagery cues, and lane connectivity was defined to match the real-world interchange layout. By working in meters, we ensured that all elements (geometry, elevation, lane offsets) remained in a consistent unit system.
To ensure map quality, geometric and logical consistency checks were performed throughout the workflow. Road segments were verified to be continuous, and lane counts were checked for consistency through merges and splits. The elevation values in the OpenDRIVE file were cross-checked against the 5 cm DTM to catch any abrupt jumps. Where possible, we also ran automated validators: the open-source QC-OpenDRIVE checker library was used to examine the final XODR file for schema compliance, connectivity, and other rule-based inconsistencies (ASAM e.V., n.d.). Any issues flagged (for example, dangling lane references or missing junction links) were corrected by adjusting the geometry in RoadRunner. The outcome was a clean, 1.8.0-compliant OpenDRIVE map with all lanes and road marks correctly attributed, suitable for downstream simulation and localization purposes.

5. Results and discussion

5.1 Derived products and dataset characteristics

The preprocessing chain produced a DTM at 5 cm spacing, an orthophoto at 4 cm GSD, and segmented point clouds supporting targeted extraction and faster navigation than raw data. The average road-surface point density was approximately 6,300 points/m², consistent with the design choice of a reduced acquisition speed to enhance lane-level detail. The segmentation macro-classes enabled practical downstream workflows: road/bare pavement supported marking candidate extraction; guardrails and poles supported constraint-based modeling; and tunnel-related subsets supported interpretation where markings were degraded or partially absent. The integration of the LiDAR point cloud with spherical imagery also supported visual interpretation and verification of road elements during restitution, especially in complex or partially occluded sections (Figure 4)
The main geometric quality indicator available without external control points was the trajectory-solution estimate: approximately ±4 cm in open-sky conditions, with larger peaks in long tunnel segments where GNSS observations degrade and trajectory propagation relies primarily on inertial integration. This behavior is aligned with the general expectation that GNSS-constrained environments stress absolute alignment and can introduce drift unless robust constraints and overlaps are available.

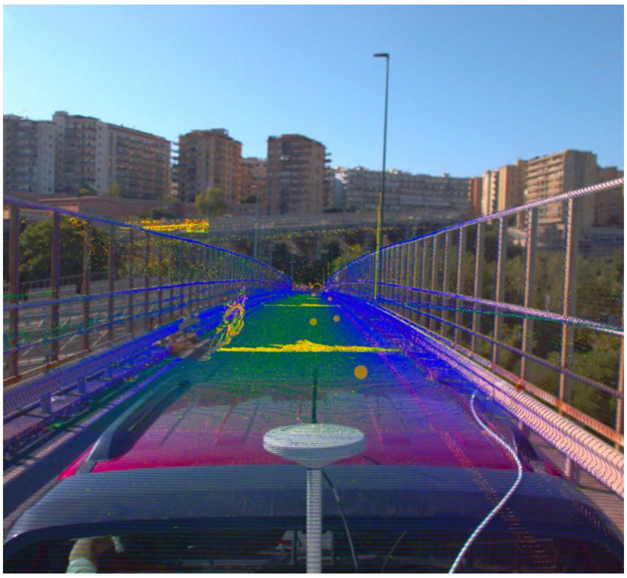

Figure 4. Frame acquired by the MMS, showing the LiDAR point cloud overlaid onto the spherical image.

5.2 Automated extraction tests

Rule-based marking extraction provided usable candidates on open-sky highway segments where markings were intact and intensity contrast was stable. Performance degraded in multiple corridor-specific cases: worn paint, strong occlusions near guardrails, complex ramp geometry, and tunnel approaches. These limits are not surprising: lighting and surface conditions affect both imagery and LiDAR intensity responses, and highway junction logic requires coherent semantic connection rules beyond local marking geometry.

The tunnel portals introduced an additional image-specific limitation observed during restitution: camera exposure adapted abruptly at tunnel entrances/exits, producing overexposed imagery for roughly ~50 m before/after the transition. In these stretches, visual identification of lane markings and signs became more difficult, and any image-assisted automated extraction degraded accordingly. This corridor-specific observation is consistent with broader remarks in HD mapping literature that illumination and shadows can materially degrade image-based extraction and annotation quality.

Learning-based segmentation tests were promising for coarse semantic separation when using pretrained models designed for road scenes, but their outputs still required careful integration into a complete OpenDRIVE model. In practice, the limiting factor was not only per-point classification accuracy, but “engineering completeness”: OpenDRIVE requires explicit and coherent road and lane linkage, junction definitions, and rule-relevant connectivity that cannot be guaranteed by segmentation alone in a complex motorway interchange without additional logic and validation layers.

5.3 OpenDRIVE compilation and interoperability implications

The OpenDRIVE map was successfully compiled in RoadRunner by integrating the raster underlays (DTM and orthophoto) and segmented point clouds as contextual constraints for authoring. RoadRunner’s OpenDRIVE import/visualize/export support enabled iterative refinement of lane geometry, lane counts across merges/diverges, junction connectivity, and object placement, resulting in an OpenDRIVE 1.8-compatible output.

From a workflow perspective, the most time-effective automation occurred upstream: classification/segmentation, DTM generation, orthophoto generation, and candidate extraction. This distribution of automation is consistent with the literature view that “manual annotation remains the most reliable method” when high precision and semantic correctness are required, even as automation can reduce the amount of manual work needed. The remaining bottleneck is interoperability: outputs from segmentation and extraction tools exist as point labels, rasters, and vector candidates, but assembling them into a fully consistent OpenDRIVE network still requires strong schema-aware integration and validation logic.

6. Conclusions

This paper updated and clarified a two-phase workflow for generating an OpenDRIVE HD map of a tunnel-rich highway corridor: multi-pass MMS acquisition and an offline restitution pipeline centered on structured preprocessing and controlled OpenDRIVE compilation. GAIA M1’s combined LiDAR acquisition and spherical imagery provides complementary geometric and interpretive support, but raw datasets at this scale are intrinsically difficult to handle directly, making early-stage preprocessing essential.

The preprocessing chain proved decisive to make restitution tractable and to support experiments on automated extraction. Classical routines (progressive TIN densification and hard-surface classification) provided robust baselines for isolating pavement and terrain in mobile laser data, while ML-based segmentation tests showed potential for coarse semantic labeling. The tunnel portal imagery limitation was a notable practical factor reducing the reliability of image-assisted interpretation and automation in localized segments; this reinforces the broader lesson that illumination variability remains a key weakness for image-driven map feature extraction.

From a strategic standpoint, two future needs emerge clearly. First, stronger end-to-end automation for “schema-complete” HD maps, i.e., automation that goes beyond detecting elements and can reliably reconstruct lane graphs, junction logic, and semantic associations in standards such as OpenDRIVE. Second, explicit reference-frame/epoch management for deployment: when HD maps are stored in a fixed-epoch national ETRF realization, and vehicles localize using PPP with ITRF solutions at near-current epochs, time-dependent transformations are required to prevent systematic offsets that can reach the decimeter range over years, an effect comparable to or larger than the tolerance expected by lane-level localization.

Finally, while integrated commercial ecosystems can deliver HD map layers with strong update pipelines, their cost structures and opacity can be misaligned with open, reproducible, or public-sector-oriented workflows. In this context, open standards and open validation tooling are practical enablers, but they do not yet remove the need for robust cross-toolchain integration engineering.

Overall, the present work contributes a structured operational workflow for HD map generation in complex motorway environments, showing how acquisition design, large-scale data preprocessing, and OpenDRIVE compilation must be addressed as closely interconnected steps within a coherent production chain. In this sense, the case study highlights the value of a pragmatic integration framework capable of supporting map restitution under realistic survey and infrastructure constraints, rather than treating each processing stage as an isolated task. Future developments will extend the experimentation to more controlled scenarios, where external references, more stable acquisition conditions, and simplified geometric configurations may support broader quantitative analyses and more explicit comparative testing of alternative processing strategies.

References

Althoff, M., Urban, S., Koschi, M., 2018. Automatic Conversion of Road Networks from OpenDRIVE to Lanelets, in: 2018 IEEE International Conference on Service Operations and Logistics, and Informatics (SOLI). Presented at the 2018 IEEE International Conference on Service Operations and Logistics, and Informatics (SOLI), IEEE, Singapore, pp. 157–162. https://doi.org/10.1109/SOLI.2018.8476801

Andert, F., Böttcher, O., Mushyam, A., Schmälzle, P., 2025. Semantic LiDAR Point Cloud Mapping and Cloud-Based SLAM for Autonomous Driving, in: 2025 IEEE/ION Position, Location and Navigation Symposium (PLANS). Presented at the 2025 IEEE/ION Position, Location and Navigation Symposium (PLANS), IEEE, Salt Lake City, UT, USA, pp. 853–860. https://doi.org/10.1109/PLANS61210.2025.11028497

ASAM e.V., 2023. OpenDRIVE 1.8.x Specification. asam.net/standards/detail/opendrive/ (January 2026).

ASAM e.V., n.d. qc-opendrive (ASAM Quality Checker for OpenDRIVE). github.com/asam-ev/qc-opendrive (January 2026).

Baasch, K.-N., Icking, L., Ruwisch, F., Schön, S., 2022. Coordinate Frames and Transformations in GNSS Ray-Tracing for Autonomous Driving in Urban Areas. Remote Sensing 15, 180. https://doi.org/10.3390/rs15010180

Bardak, G., Sodano, M., Scholz, M., 2025. Integration of HD Maps and Point Clouds: An Efficient 3D Reconstruction Framework for Autonomous Driving Applications. Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci. XLVIII-4/W13-2025, 49–56. https://doi.org/10.5194/isprs-archives-XLVIII-4-W13-2025-49-2025

Boyko, A., Funkhouser, T., 2011. Extracting roads from dense point clouds in large scale urban environment. ISPRS Journal of Photogrammetry and Remote Sensing 66, S2–S12. https://doi.org/10.1016/j.isprsjprs.2011.09.009

Centro Nazionale per la Mobilità Sostenibile (MOST), 2021. Spoke 7: Cooperative, Connected, and Automated Mobility and Smart Infrastructure. centronazionalemost.it/Spoke%207.html (January 2026).

Chiang, K.-W., Pai, H.-Y., Zeng, J.-C., Tsai, M.-L., El-Sheimy, N., 2022. Automated Modeling of Road Networks for High-Definition Maps in OpenDRIVE Format Using Mobile Mapping Measurements. Geomatics 2, 221–235. https://doi.org/10.3390/geomatics2020013

Diaz-Diaz, A., Ocana, M., Llamazares, A., Gomez-Huelamo, C., Revenga, P., Bergasa, L.M., 2022. HD maps: Exploiting OpenDRIVE potential for Path Planning and Map Monitoring, in: 2022 IEEE Intelligent Vehicles Symposium (IV). Presented at the 2022 IEEE Intelligent Vehicles Symposium (IV), IEEE, Aachen, Germany, pp. 1211–1217. https://doi.org/10.1109/IV51971.2022.9827297

Eisemann, L., Maucher, J., 2023. Automatic Odometry-Less OpenDRIVE Generation From Sparse Point Clouds, in: 2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC). pp. 681–688. https://doi.org/10.1109/ITSC57777.2023.10421842

Esmaeili, S.A., Singh, B., Davis, L.S., 2017. Fast-At: Fast Automatic Thumbnail Generation Using Deep Neural Networks, in: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). Presented at the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE, Honolulu, HI, pp. 4178–4186. https://doi.org/10.1109/CVPR.2017.445

Eurocontrol, 2017. Positioning in ETRS89 and GPS/GLONASS. Technical Note 7.

Geiger, A., Lenz, P., Urtasun, R., 2012. Are we ready for autonomous driving? The KITTI vision benchmark suite, in: 2012 IEEE Conference on Computer Vision and Pattern Recognition. Presented at the 2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE, Providence, RI, pp. 3354–3361. https://doi.org/10.1109/CVPR.2012.6248074

Geo-spatial CEN/ISO, 2019. ETRS89 and UTM Projections. epsg.org/Transformations.aspx (January 2026).

GreenValley International, n.d. LiDAR360MLS Manual: Ground Classification and Deep Learning Modules. greenvalleyintl.com/docs/lidar360mls/Main/Classification/ClassifyGroundPoints.html (January 2026).

Greffrath, G., Hecker, S., 2017. Distribution of OpenDRIVE Map Files. ASAM Technical Report.

HERE Technologies, 2019. HERE HD Live Map: Self-Healing Technology. White Paper. go.engage.here.com/rs/142-UEL-347/images/Tech%20Brief%20HERE%20Technologies%20Self-healing%20Map.pdf (January 2026).

Hexagon, 2023. HxGN SmartNet Network-RTK Service. hxgnsmartnet.com/ (January 2026).

IGMI/IGM, n.d. Rete Dinamica Nazionale (RDN) – ETRF2000 (Italy). igmi.org/en/direzione-geodetica/progetto-rdn-rete-dinamica-nazionale (28 March 2026).

ISGeoid (Politecnico di Milano), n.d. ITALGEO05: Italian Geoid Model. isgeoid.polimi.it/Geoid/Europe/Italy/italgeo05_g.html (January 2026).

Kalenjuk, S., Lienhart, W., 2022. A Method for Efficient Quality Control and Enhancement of Mobile Laser Scanning Data. Remote Sensing 14, 857. https://doi.org/10.3390/rs14040857

Kokkinogenis, Z., Teixeira, M., d'Orey, P.M., Rossetti, R.J.F., 2019. Tactical Level Decision-Making for Platoons of Autonomous Vehicles Using Auction Mechanisms, in: 2019 IEEE Intelligent Vehicles Symposium (IV). Presented at the 2019 IEEE Intelligent Vehicles Symposium (IV), IEEE, Paris, France, pp. 1632–1638. https://doi.org/10.1109/IVS.2019.8814122

Ma, L., Li, Y., Li, J., Wang, C., Wang, R., Chapman, M.A., 2018. Mobile Laser Scanned Point-Clouds for Road Object Detection and Extraction: A Review. Remote Sensing 10, 1531. https://doi.org/10.3390/rs10101531

MathWorks, 2023. RoadRunner. MathWorks, Natick, USA. mathworks.com/products/roadrunner.html (January 2026).

MicroGeo, 2022. GAIA M1 Mobile Mapping System. Brochure.

Pang, E., Davletshina, D., Marie d'Avigneau, A., Park, N., De Silva, L., Brilakis, I., 2025. Digitalizing Road Maintenance: A Novel Approach for the Preparation and Integration of 2D and 3D Road Data through Image Shadow Removal and Point Cloud Densification. J. Comput. Civ. Eng. 39, 04025080. https://doi.org/10.1061/JCCEE5.CPENG-6246

Park, J., Kim, S., Lee, J., Ham, J., Oh, K., 2020. Method of operating a GIS-based autopilot drone to inspect ultrahigh voltage power lines and its field tests. Journal of Field Robotics 37, 345–361. https://doi.org/10.1002/rob.21916

Qi, X., Luo, Y., Wu, G., Boriboonsomsin, K., Barth, M.J., 2017. Deep reinforcement learning-based vehicle energy efficiency autonomous learning system, in: 2017 IEEE Intelligent Vehicles Symposium (IV). Presented at the 2017 IEEE Intelligent

Vehicles Symposium (IV), IEEE, Los Angeles, CA, USA, pp. 1228–1233. https://doi.org/10.1109/IVS.2017.7995880

Roddick, T., Cipolla, R., 2020. Predicting Semantic Map Representations From Images Using Pyramid Occupancy Networks, in: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). Presented at the 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), IEEE, Seattle, WA, USA, pp. 11135–11144. https://doi.org/10.1109/CVPR42600.2020.01115

Schott, M., Grinberger, A.Y., Lautenbach, S., Zipf, A., 2021. The Impact of Community Happenings in OpenStreetMap—Establishing a Framework for Online Community Member Activity Analyses. IJGI 10, 164. https://doi.org/10.3390/ijgi10030164

Schwab, B., Kolbe, T.H., 2022. VALIDATION OF PARAMETRIC OPENDRIVE ROAD SPACE MODELS. ISPRS Ann. Photogramm. Remote Sens. Spatial Inf. Sci. X-4/W2-2022, 257–264. https://doi.org/10.5194/isprs-annals-X-4-W2-2022-257-2022

TerraSolid, n.d. TerraScan User Guide: Hard Surface Routine. terrasolid.com/guides/tscan/crhardsurface.html (January 2026).

TomTom, n.d. TomTom Orbis Maps for Automated Driving. tomtom.com/products/orbis-maps-for-automated-driving/ (January 2026).

Wu, H., Yao, L., Xu, Z., Li, Y., Ao, X., Chen, Q., Li, Z., Meng, B., 2019. Road pothole extraction and safety evaluation by integration of point cloud and images derived from mobile mapping sensors. Advanced Engineering Informatics 42, 100936. https://doi.org/10.1016/j.aei.2019.100936

ZENRIN Co., n.d. ZENRIN HD Maps. zenrin-europe.com/automotive-solutions (January 2026).

Zhou, M., Sui, H., Chen, S., Wang, J., Chen, X., 2020. BT-RoadNet: A boundary and topologically-aware neural network for road extraction from high-resolution remote sensing imagery. ISPRS Journal of Photogrammetry and Remote Sensing 168, 288–306. https://doi.org/10.1016/j.isprsjprs.2020.08.019