Thesis for the degree of Licentiate of Engineering

Towards Data-Driven Railway Maintenance Supported by Autonomous Inspection Robots

VIVIEN LACORRE

Department of Mechanical Engineering
CHALMERS UNIVERSITY OF TECHNOLOGY
Göteborg, Sweden 2026

Towards Data-Driven Railway Maintenance Supported by Autonomous Inspection Robots
VIVIEN LACORRE

Thesis for Licentiate of Engineering

Department of Mechanical Engineering
Chalmers University of Technology
SE–412 96 Göteborg
Sweden
Telephone: +46 (0)31–772 1000

Chalmers digitaltryck
Göteborg, Sweden 2026

Towards Data-Driven Railway Maintenance Supported by Autonomous Inspection Robots
Vivien Lacorre
Department of Mechanical Engineering
Division of Vehicle Engineering and Autonomous Systems
Chalmers University of Technology

Abstract

Europe aims to double the use of railway transport by 2050 to support efficient and low-carbon transportation. However, increasing the pressure on an already stressed system calls for innovative and modern solutions. Integrating rail-bound autonomous robots in daily inspection, diagnosis, maintenance and repair operations would be a major leap forward. Benefits include high flexibility and efficiency in terms of cost and workforce needs. In this thesis, the concept and design of a versatile inspection, diagnostic and autonomous repair robot is presented, with a focus on onboard localization. While recent technological advances escalated the deployment of robotic solutions in numerous fields, the railway ecosystem is characterized by extensive and constraining safety regulations regarding vehicles and operations. To satisfy current requirements, special attention must be paid to achieving a robust and reliable tracking of the robot at all times, regardless of weather conditions, satellite coverage or network connectivity. This work uses custom digital maps of the French and Swedish railway networks to enable real-time path planning and navigation onboard a robot. Such maps also strongly improve the reliability of positioning approaches by constraining the problem to one-dimensional tracks. Beyond improving the quality of georeferenced data collected during inspection, a precise localization of the robot allows a safe and efficient navigation throughout the railway network. An obstacle detection system was developed to reduce collision risk. Finally, preliminary results are obtained with a prototype during field tests on a real track. These results demonstrate the feasibility of map-enabled autonomous inspection and underscore the critical role of resilient localization for real-world deployment.

Keywords: railway infrastructure, railway maintenance, railway inspection, field robotics, navigation, autonomous inspection robot

The most exciting phrase to hear in science,
the one that heralds the most discoveries,
is not "Eureka!" (I found it!) but "That's funny..."

— Isaac Asimov

Acknowledgements

I am sincerely grateful to my main supervisor, Professor Krister Wolff, for this opportunity to learn and grow in a research field that I have always been passionate about. Your continuous patience and thoughtful guidance have been a source of inspiration and motivation. I want to thank my co-supervisors, Professor Marco L. Della Vedova and Professor Peter Forsberg. I greatly value our discussions and look forward to our future collaborations. I also express my profound gratitude to my examiner, Professor Mattias Wahde, for his insightful feedback and direction.

I would like to acknowledge all my colleagues at VEAS for their shared enthusiasm about science and for contributing to a pleasant and supportive work environment.

I am profoundly thankful to my family, my friends, and all the people who participated in their own unique ways to bring joy to my life throughout this journey.

This work has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement No. 101101966 (IAM4RAIL). The project is supported by Europe's Rail Joint Undertaking and its members. Trafikverket (Swedish Transport Administration) is gratefully acknowledged for its financial co-funding and for providing the test track and technical support.

List of included papers

Paper A L. -R. Joly, V. Lacorre and K. Wolff, *VERNE: A Spatial Data Structure Representing Railway Networks for Autonomous Robot Navigation.* In: IEEE Open Journal of Vehicular Technology, vol. 7, pp. 1-14, 2026.

Paper B V. Lacorre, K. Wolff and L. -R. Joly, *Autonomous Robots Enabling Smart, Data-Driven Railway Maintenance.* To appear in: Proceedings of the Transport Research Arena (TRA) Conference, May 18-21, 2026, Budapest, Hungary. Lecture Notes in Mobility, Springer.

Table of contents

Chapter 1

Introduction and motivation

This thesis presents the development and initial field testing of a VIDAR[1] (Versatile Inspection, Diagnostic and Autonomous Repair Robot). The VIDAR concept aims to support flexible, modular inspection and maintenance operations on rail infrastructure. Here, the focus is put on the specific case of railway tracks and their surrounding areas (e.g., overhead line equipment, trackside signaling and power, tunnels, bridge decks...), excluding rolling stock.

Efficient transportation is key to enabling the exchange of resources, goods, and labor. Historically, railways have been the cornerstone of industrialization and economic development as a result of its efficiency and high capacity. In 2019 they accounted for 8 % of the world's motorized passengers and 7 % of its freight, and traffic is expected to double between 2020 and 2050 [34].

As global populations urbanize and the demand for mobility intensifies, expanding rail transport offers a sensible solution, especially since the reduction of carbon emissions is at the heart of the ecological transition. As a matter of fact, trains generate substantially less greenhouse gases per passenger-kilometer compared to cars and aircraft [2] and high-speed railway lines are less land-intensive than highways [35].

The European Union's strategy regarding sustainability and smart mobility aims to double rail freight traffic and triple high-speed rail traffic between 2020 and 2050 [18]. Importantly, from 2050, the single European Rail Traffic Management System (ERTMS) will have replaced national systems, leading to an increased capacity of existing railway lines [68] and aiding the integration of new digital systems.

[1] In Paper B, the prototype is designated under the name "Multi-Purpose Inspection Robot (MPIR)"

As global demand grows, existing infrastructure becomes increasingly worn and stressed both physically and in terms of scheduling. An increase in vehicle speeds, passenger-kilometers (pkm) and tonne-kilometers (tkm) raises the risk of accidents and calls for more frequent inspection and maintenance. These processes, in turn, further disrupt traffic.

In the United Kingdom, according to Network Rail's 2023-24 report, maintenance and renewals accounted for 63% of core operational expenditure [54]. Moreover, in 2000, track maintenance represented 25-35% of the total operating costs of Australian freight railways [67]. Even modest improvements in maintenance efficiency and flexibility can yield significant cost savings, reduce downtime and promote more streamlined operations.

Traditional routines include the use of nondestructive testing (NDT) via diverse sensors installed on inspection rail-road vehicles (i.e., vehicles that can operate on both rail tracks and roads) or full-sized rail cars (e.g., track geometry cars which are capable of traveling up to 335 kilometers per hour). Upon the discovery of a defect, a separate maintenance team will follow with appropriate tools (tamping machine, ballast regulator, and so on).

While these methods are more efficient than historical ones relying on hand-held devices, they still require a crew of humans to operate the machines and interpret detections. Since the 2010s, most regular traffic trains are equipped with sensors to collect data used for predictive maintenance and real-time monitoring. However, this comes at the expense of flexibility.

First, high-speed lines (and slower lines when conditions demand it) often require an initial inspection before the traffic of the day can start. The current solution is to send a train (equipped with the sensors) without passengers to verify the absence of risks (obstacles, defects, slippery rails, etc.) but it is inefficient. Apart from this, inspection is not on-demand, but fixed by the schedule of regular traffic.

Second, regular trains are heavy and cannot easily stop to perform slow, high-resolution measurements, or reverse to collect additional data in a specific area.

Third, if a minor defect or wear is found, an entire maintenance machine has to be scheduled on the track, while a small repair tool mounted on the train could have sufficed.

Fourth, several parts of the infrastructure are not directly accessible from a regular rail vehicle. Specialized machines are required to access bridges, for example [22, 55].

Since the introduction of the first steam locomotives of the nineteenth century, rail transport has undergone multiple technological revolutions. The pantograph, the continuous welded rail, the ballast tamping machine, the track circuit... These inventions profoundly changed railway vehicles, infrastructure, maintenance, and operations, each more transformative than the last. In the last decade, promising AI-driven predictive maintenance methods have been studied in the literature for numerous components of the railway infrastructure [8, 70].

On the other hand, recent research identified the availability and completeness of sensor data as a current limitation, highlighting the need for more frequent, extensive and diversified inspection [8, 70].

The twenty-first century presents significant progress in the fields of digitalization and machine learning, paving the way for fully integrated autonomous robots. With limited deployment cost and exceptional scheduling flexibility, a fleet of robots may continuously collect data along the tracks. Such a large volume of frequently updated data is required to fuel a digital twin of the infrastructure used for predictive maintenance. This is expected to increase the use of the rail network with less frequent interruptions of operations and more targeted maintenance activities while simultaneously curbing the probability of defects.

Similarly, if equipped with a manipulator arm, a robot may perform cost-effective on-site maintenance and repair with optimal precision and repeatability, thus alleviating human fatigue and reducing the risk of errors.

Relevant inspection sensors that can be embarked on a rail robot include but are not limited to: light detection and ranging (lidar), (thermal) camera, acoustic sensors, vibration sensors, ground-penetrating radar, eddy current and magnetic flux sensors [7, 8, 72]. While equipping robots with such sensors is technically feasible, deploying them autonomously in railway environments requires solving fundamental challenges in navigation, localization, and safety. The following section reviews the state of the art in these areas and identifies the gaps this thesis targets.

1.1 Related work

Research in the field of applied robotics for autonomous railway inspection and maintenance is still at an early stage. Essential aspects of a fully operational autonomous robot are: (1) real-time, robust onboard localization and navigation, and (2) reliable obstacle detection.

Although robust simultaneous localization and mapping (SLAM) has been substantially studied for autonomous vehicles [83], the rail environment presents unique characteristics that deserve dedicated experiments (vibrations, slippery rails, monotonous or featureless non-urban landscapes, high speeds, and so on). Indeed, a few studies specifically investigated the special case of railway [43, 56, 61], and some innovative solutions appeared [29, 30, 41]. Nevertheless, while most localization methods for rail vehicles rely on digital maps of the network, the concrete software implementation of such maps is not made explicit [27, 37, 49, 64]. As of now, it is unclear if existing solutions are usable in practice for onboard navigation with limited computing power. In contrast, a large share of railway models focus on describing elements of the infrastructure and are not designed to enable railway navigation [9, 20, 48, 77, 80].

The majority of existing robotic prototypes in the literature depend entirely on a global navigation satellite system (GNSS) receiver for localization. To cope with satellite coverage outage, some experiments combine a mix of different sensors under favorable weather conditions [13, 38, 44, 62].

The study of obstacle detection methods in the railway context is rare and primarily applied to trains [63]. Hence, there is a need for deeper analysis focused on smaller robotic systems.

To tackle these gaps, and clear the way for the future adoption of autonomous rail-bound robots, this thesis develops and validates a real-time navigation stack supported by a custom digital map (built with the "vectorial railway network" (VERNE) framework [39]) and demonstrates its feasibility through field experiments with a prototype of VIDAR. Additionally, this preliminary work sets the ground for advanced research on onboard robust localization and obstacle detection, as well as the collection of valuable datasets.

1.2 Research questions

To support the development of autonomous railway robots, this thesis investigates the following research questions:

RQ1 *How can a digital map of a railway network be built from openly available spatial data and enable autonomous navigation?*

RQ2 *What are the hardware requirements to store and process a digital map onboard a lightweight railway robot for real-time navigation?*

RQ3 *How does a digital map of a railway network contribute to improve localization accuracy and robustness?*

RQ4 *Which sensors and safety mechanisms are needed for an autonomous robot to reliably detect an obstacle on its track and stop preemptively under favorable weather conditions?*

RQ1 and RQ2 examine the feasibility of using digital maps for onboard autonomous navigation on railway, and RQ3 investigates potential benefits for localization. RQ4 explores solutions for enforcing safe operations.

1.3 Scope

Beyond technical aspects, the evolution of railway systems needs to be done in a systemic way. Notably, the compliance with functional safety norms (e.g., safety integrity levels (SIL) defined by IEC 61508) and European standards (e.g., EN50128 and EN50129) is mandatory for the certification of safety-critical railway systems in Europe. However, these frameworks currently lack the necessary methodologies for verifying and validating AI-driven components [1, 19, 59]. On that topic, future directions are outside of the scope of this thesis and are instead better suited in fields connected to regulatory science, safety assurance, and safety of the intended functionality.

Another challenge will be the seamless integration of such robots into railway operations, asset management and traffic planning workflows. Noteworthy concerns are the definition of the operational domain (OD) and operational design domain (ODD) for these new systems. This falls within the field of Systems Engineering, such as operations research, prognostics and health management, and human-systems integration.

Finally, software and hardware deployed for such a critical application require in-depth cybersecurity audits and strict adherence to "secure by design" principles [17].

Technical solutions presented in this thesis are deployed in a controlled and isolated rail environment, and are not tested for industrial use. The railway tracks employed for field experiments are clear of incoming traffic (either on parallel tracks or near level crossings) and traffic rules are not required to be followed. Human operators are supervising the robot at all times and are ready to stop it if needed.

The real-time localization of the robot is achieved onboard, without any trackside equipment. Rail switches included in the path planned by the robot have to be manually operated (the robot waits for clearance before crossing them).

Once an autonomous robot can navigate safely, its effectiveness depends on the quality and reliability of its inspection and maintenance tools, from both the software and hardware aspects. Modern machine learning approaches represent the majority of the current literature on the detection of defects. The research on the realization of maintenance/repair tasks, however, is still emerging and at an early stage [12, 15, 42].

1.4 Outline and author's contributions

Chapter 2 presents the general concept of VIDAR (Versatile Inspection, Diagnostic and Autonomous Repair robot), as well as design choices for the prototype used in this work. Chapter 3 focuses on the construction and exploitation of railway digital maps with concrete deployment on a VIDAR prototype. Chapter 4 details the sensor suite of the robot, addressing safety and obstacle detection. Chapter 5 demonstrates the capabilities of the map-based navigation system of the VIDAR, including localization, path planning and safe inspection of geographic zones. A conclusion to the thesis is given in Chapter 6.

As a major contributor in Paper A, the author created an optimized implementation of the proposed data structure and developed extensive tests to validate the algorithms on real network data. This paper provides a framework to generate and exploit an interpretable representation of railway networks designed for autonomous robots to navigate and plan paths in a fast and efficient way.

The author's contribution in Paper B concerned the implementation of the robot's architecture and the processing of data collected during field experiments. By openly sharing design choices, challenges and results obtained with a VIDAR prototype in real-world scenarios, the author aims to contribute to the railway and robotics communities. First, this allows infrastructure managers to prepare for the future potential adoption of robots and their integration into existing workflows. Second, roboticists working in the field of railway can benefit from the experience of each other to avoid duplicate work and, for example, create generic software components.

Chapter 2

The VIDAR system: concept and design

2.1 Concept

This section explains the concept of the rail-bound VIDAR (versatile inspection, diagnostic and autonomous repair robot).

2.1.1 An autonomous robot for railways

Before presenting the robot in detail, the definition of an "autonomous" robot must be clarified. The grade of automation (GoA) is used to characterize the level of autonomy of a rail vehicle. However, since the capabilities of a VIDAR are more diverse than regular trains, a VIDAR will require an extension of the GoA4, which is currently the highest GoA existing for trains. GoA4 rail vehicles are capable of operating automatically without on-train staff, including door closing, obstacle detection, and emergency situations, but are usually deployed in controlled subway environments.

How a VIDAR will be integrated to operations highly depends on how the operations are run on a given line. A line where all trains are GoA4 will require automatic and frequent communication with other active rail vehicles, for example, via communications-based train control (CBTC) technologies. Deploying VIDARs on a line with a lower GoA requires extended communication with the operations control center (OCC) so that it takes the responsibility for handling the cohabitation of all rail entities. Regardless of how traffic is handled, this thesis considers that mission-related tasks can be performed without any human intervention, from fault detection and diagnosis to repair. A human supervisor would only be required to monitor the status and outcome of current and past missions and execute appropriate actions if needed, such as dispatching heavy-duty machines to complete tasks the robot cannot perform.

The hypothetical operation scenario that guided the robot's design is as follows. A VIDAR can be deployed manually by placing it on the track or automatically from a docking station connected to the rail network. As a starting point, a mission plan must be given. It contains geographic waypoints and a purpose (e.g., data collection, routine fault detection or repair intervention). From there, the robot plans its path to reach all waypoints efficiently throughout the mission.

The spatial aspect of the path is computed using an onboard digital map of the rail network.

The temporal aspect of the path – handling potential track occupancy conflicts with other rail vehicles – can be solved in two ways. One way could be to run the robot on reserved tracks, isolated from other traffic. Another way could be to send a traffic timetable to the robot so that it generates a path that safely reaches and vacates tracks in an efficient manner. The former approach should be prioritized for near-term deployment, but the latter can be accomplished by building on existing multi-agent path finding algorithms [4, 65] and research on the train routing problem [10, 73].

For safety, the robot's planned path must be communicated to the operations control center (OCC) so that humans can perform potential adjustments to avoid conflicts with other traffic. In addition, the robot's current geographic position must be transmitted in real time. If rail switches must be crossed, a VIDAR stops at a safe distance and waits for operators to set the switch according to the planned path, and give clearance to continue the mission. Alternatively, the robot could be given control of switches it is about to cross.

How a VIDAR would interact with railroad signaling systems, notably level crossings, is currently an open question since the materials it is made of and its lower weight compared to usual rail vehicles could make it invisible to track circuits or axle counters. On the other hand, it should be mentioned that thanks to its light weight, a VIDAR approaching an unprotected level crossing could quickly lower its speed and use an advanced obstacle detection system to safely go through it. Additionally, via the use of onboard cameras, visual signals can be identified and interpreted by the robot.

During its mission, a VIDAR has full control over its sensors and the data they collect. Sensor data is processed onboard, for example to detect faults and their position in real time, or extract high-level information. This limits the amount of information the robot needs to send to human operators: only the status of the mission or urgent events need to be transmitted. For missions involving large data recordings, the robot stores data onboard until it docks or connects to a high-throughput mobile network, at which point the data is transferred.

2.1.2 Foreseen capabilities

Now that the hypothetical operating conditions are established, the capabilities of a VIDAR can be discussed. Essentially, the purpose of a VIDAR is to support infrastructure managers (IMs)[1] and railway undertakings (RUs)[2] in planning and performing maintenance activities to increase safety and rail utilization.

In this context, the first major contribution of a VIDAR is its ability to collect georeferenced data, which is currently lacking in terms of quantity, quality and diversity [8]. Solving this issue would enable the extensive deployment and use of digital twins to access the numerous benefits of predictive maintenance [70]. Indeed, an accurate and reliable localization system is required for the safe operations of a VIDAR, and the quality of its positioning will directly impact the quality of the recorded data. Thanks to modern computing hardware, it is possible to handle multiple sensors and large quantities of data with limited space and energy consumption onboard the robot.

The second major role a VIDAR could play is that of an inexpensive reconnaissance vehicle, evaluating the condition of a track before daily passenger traffic. This is particularly relevant in cold or windy areas where snow, tree leaves or branches risk covering the track. As a matter of fact, a project[3] is currently under development in France to use unmanned rail vehicles on high-speed lines for that exact purpose instead of sending empty trains.

Finally, a VIDAR equipped with appropriate repair tools would represent an essential asset in daily railway maintenance operations. In particular, a VIDAR on an inspection mission could provide quick repairs on minor new-found defects to avoid dispatching on site a fully equipped maintenance vehicle with a crew. While robotic manipulation is still an active research field, recent breakthroughs have brought increased resilience, efficiency and the ability to adapt to a wide variety of tasks [47, 60, 74].

[1] An IM is a body or undertaking responsible for establishing, managing, and maintaining railway infrastructure, including traffic management and control-command and signaling systems.

[2] An RU is a company that provides rail transport services, meaning it operates trains for either passengers or freight.

[3] `https://www.groupe-sncf.com/en/innovation/emerging-technologies/network-monitoring`

2.2 Prototype architecture and design choices

With the support of the Swedish Transport Administration (Trafikverket), Chalmers University of Technology undertook the development of the VIDAR prototype presented in this thesis (see Fig. 2.1).

The existing Alumicart Bance trolley (Bance Track Ltd 2025) built for manual analogue control was modified to allow digital control. Its lightweight aluminium frame, rail wheels, and battery-powered electric drive allow a simple deployment on tracks with or without catenary poles for continuous operation of more than four hours. Speeds of up to 15 km/h can be achieved, which is considered to be representative of a scenario where slow speeds are required to obtain high-resolution measurements via sensors with low sampling frequency or high sensitivity to vibrations.

Within the scope of this thesis, further modifications have been made to the robot to introduce onboard localization, navigation, and autonomous behaviors. This work, presented in the following chapters, focuses on the robot's autonomy and navigation systems. It does not cover real-world inspection, diagnosis, or repair, but provides the foundation for future development.

The design adopts a distributed modular architecture for both hardware and software. This allows the independent development, replacement and validation of each component, and avoids having a single point of failure thanks to redundancies in the system. A central onboard industrial computer (IXH-AL35 Intel) runs the majority of the calculations. It is connected to a mobile broadband wifi router (4G/5G) to receive commands and mission plans from a remote base station. Four separate Arduino Zero microcontrollers are given dedicated tasks:

One receives target speeds and the distance to the closest front and back obstacles. Its role is to control the motors of the robot, gradually reaching the target speeds to avoid too strong accelerations. It will also automatically stop the robot if an obstacle is too close.

Two are connected to single-point lidars pointing either forward or backward. After applying a lightweight preprocessing filter, they report the shortest distance to any front or back obstacle within range.

One is in charge of configuring and receiving data from an inertial measurement unit (IMU) and two magnetometers.

The robot operating system (ROS 2) [51] was chosen as the cornerstone of the modular onboard software. Each software component runs on individual processes and exchanges data in real time. If a component crashes, the others continue to run and can adapt. The Micro-ROS stack [5] is used in all microcontrollers to enable communication with the central computer regardless of hardware differences.

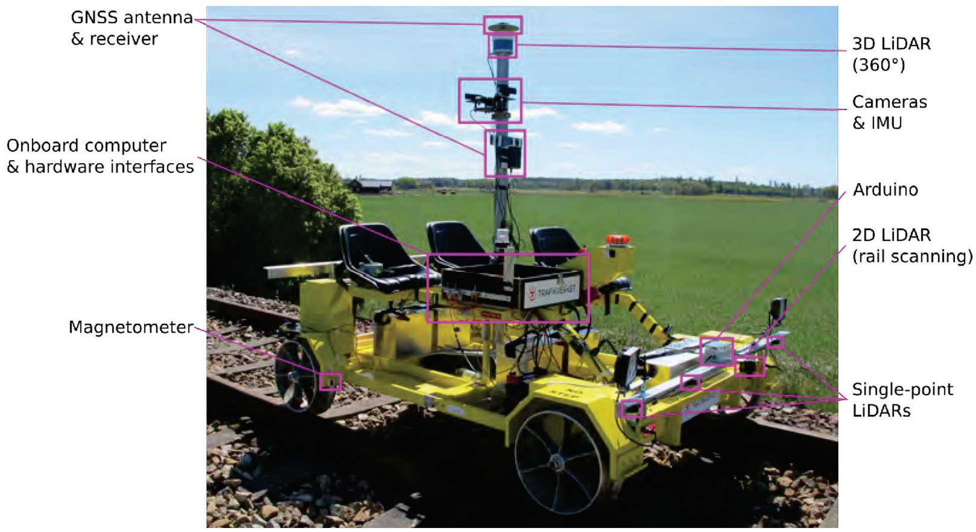

Figure 2.1: The VIDAR prototype and its components (reproduced from Paper B).

The hypothetical scenario presented in Section 2.1.1 was a determining factor in the development of the prototype. As a result, the robot is now functioning fully autonomously after receiving a mission plan. It is only sending signals to a remote computer so that a human can monitor its current position, planned path, and mission status. However, it is necessary to add a few precisions on the choices taken during the experiments presented in Section 5.2. In practice, during tests, the robot operates is manually placed on an isolated track with no interaction with railroad signaling systems or visual signals, as they are absent from the test track. Rail switches are manually operated when needed, and human operators are present at level crossings to guarantee safety despite the current lack of advanced obstacle detection system in the robot. Finally, for the time being, no data processing aiming at supporting predictive maintenance is performed on the robot with the data it collects.

2.3 Emergency stopping and safety behaviors

To facilitate adoption and deployment and respond to RQ4, backup software functionalities must be discussed. Solutions presented here rely on the distributed architecture presented above.

The VIDAR prototype is remotely controlled through the mobile network via a computer in a fixed control station. To ensure secure operations, fail-safe mechanisms engage the brakes upon any of these failures:

- The onboard computer loses contact with the control station (outdated health check).
- The microcontroller in charge of the motor loses contact with the onboard computer.
- An obstacle is detected closer than a configured distance.
- The data coming from the short-range or long-range obstacle detection system is outdated (i.e., no new data has been received since a configured duration).
- The current position estimation is outdated or too imprecise.

It should be noted that losing contact with the control station is expected in some tunnels or rural areas. This edge case can be handled by setting a timed meeting point in a zone where mobile network coverage is good. If the robot fails to contact the control stating within the set duration, the robot should perform an emergency stop. From there, appropriate actions must be taken from the control station's side to prevent collision with external traffic and plan the manual retrieval of the robot.

To enhance safety, the robot slows down preventively as soon as it detects an obstacle, such that the braking distance diminishes as the obstacle gets closer.

Moreover, thanks to a map of the rail network, the robot is aware of all rail switches and level crossings it will encounter along its path. Before reaching any of them, the robot stops at a safe distance and waits for a clearance message sent remotely by an operator.

2.4 Remote operations

As previously discussed, the VIDAR prototype is remotely controlled and supervised by a control station via the mobile network. Fig. 2.2 clarifies the current workflow and highlights the distinct roles of the robot and the remote operator. While this setup represents a crucial stepping stone toward long-term operations, systematic real-world deployment requires addressing several unresolved issues. Specifically: which operational entity should assume responsibility for the control station? Should the system's autonomy be increased, or does safe deployment necessitate the involvement of more humans? Furthermore, establishing the maximum number of robots assigned to a single operator and determining the maximum acceptable wireless communication latency remain critical questions for future investigation.

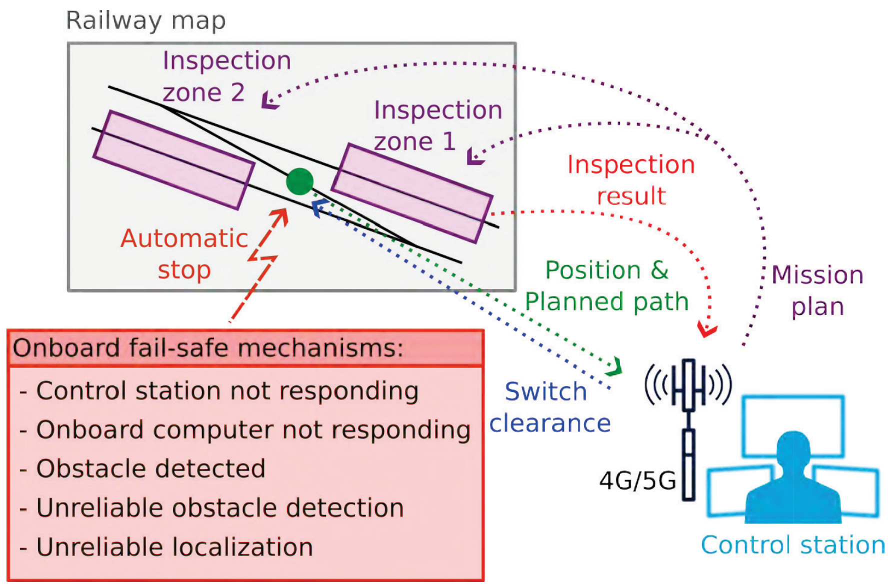

Figure 2.2: Foreseen operating scenario with one VIDAR (green dot) and expected interactions.

Chapter 3

Autonomous navigation: mapping, planning, and control

3.1 A data structure for railway digital maps

The localization of rail vehicles can be improved via the use of digital maps, [27, 37, 49, 64], notably via the use of multiple track hypotheses, [25, 26, 43, 50]. However, the literature does not clarify how such hypotheses can be fetched from a map, nor does it indicate the data structure employed and its associated query algorithms. Most importantly, the requirements in terms of computational efficiency and memory overhead have not been studied in this context.

Therefore, research conducted within the scope of this thesis was published in Paper A to answer RQ1 and RQ2. Digital maps of the railway networks of Sweden and France respectively were built with the proposed "vectorial railway network" (VERNE) spatial data structure.

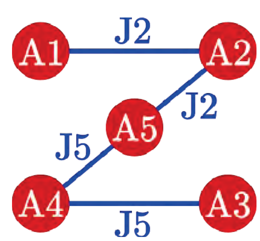

Figure 3.1: Representation of a railway network with atoms (red) as vertices and junctions (blue) as edges of an undirected graph (reproduced from Paper A).

The VERNE framework is built around "atoms" which are indivisible groups of consecutive rail segments forming a continuous path. They are separated from other atoms by railroad junctions, or exist as isolated units with two dead ends. This representation allows the construction of an undirected graph with junctions as edges and atoms as vertices. An example is displayed in Fig. 3.1. The proposed "containing atom query" returns all atoms whose bounding boxes intersect the uncertainty disk, providing a set of candidate tracks for the robot's position.

To answer RQ1, two open access databases, one for each country, were examined. Despite the fact that they present a few limitations in terms of completeness and uncertainties, the maps were successfully built and verified extensively. A valuable outcome of this research is the quantification of the amount of computer memory needed to load and exploit the map. This provides a response to RQ2 by showing that under 20 MiB of memory is sufficient for any of these two countries, which can be easily handled by any standard computer.

Results regarding the query speed on the VERNE data structure must be carefully interpreted before drawing a conclusion on its computational cost. Using an AMD EPYC 9354 processor capable of running at up to 3.8 GHz, the presented custom k-d tree LR with a "containing atom query" consistently obtained shorter query times (between 1 µs and 2 µs) compared to an STRtree [45] with a distance query, despite the latter being a well-optimized and widely used spatial index. Contrary to the processor used for testing, a low-to-mid-tier CPU will not have enough cache memory to hold the entire 20 MiB map on-chip. As a result, such a CPU would need to access its random-access memory repeatedly, thus incurring latency penalties. While this assumption remains to be quantified, the algorithm is likely to remain highly viable and performant on consumer-grade hardware. To provide additional evidence, Section 5.2.2 demonstrates a real-world deployment where a VERNE map is used to navigate between inspection zones. The onboard computer includes an Intel® Core™ i5-1235U processor with 12 MB cache and running at up to 4.4 GHz.

3.2 Path planning algorithm

Since the robot is constrained to move along a known railway network, the path planning problem is fundamentally one-dimensional: every feasible trajectory is a sequence of atoms joined at switches.

Graph formulation In a graph like the example shown in Fig. 3.1, each atom carries information on its geometry, so the cost of traversing an atom is equal to its arc length. A query is defined by two points p_s and p_t (the robot's starting position and its target position). The graph is augmented with two virtual nodes, one for each query point p_s and p_t. Each virtual node is connected to every neighbor of the atom it belongs to, with weight equal to the along-track distance from p_s (resp. p_t) to the shared junction. Atom endpoints that connect to no other atom (e.g. at a buffer stop) thus contribute no augmentation edge.

Planning The A* search algorithm [24] is run on the augmented graph, using the straight-line Euclidean distance from an atom's closer endpoint to p_t as the heuristic. While using the exact along-track distance would provide a perfect heuristic, calculating it dynamically would essentially require solving the routing problem itself. Instead, the straight-line Euclidean distance serves as a fast, easily computable alternative. This heuristic is guaranteed to be an admissible lower bound on the remaining along-track distance, since any path from an atom to the target must pass through one of the atom's two endpoints, and the straight line represents the shortest possible physical path.

A* returns an optimal atom sequence under the chosen cost. The cost includes the arc length of each traversed atom but does not separately model switch reversals. In practice, at points where the plan requires crossing the same switch twice to branch into an adjacent track, the robot performs an overshoot-and-reverse maneuver that physically travels only a short, fixed distance into the transition atom before reversing. To keep the cost model computationally efficient, the transition atom's full arc length is directly added to the path's cost. This overestimation serves a dual purpose: it avoids the complexity of modeling precise overshoot distances and naturally biases the search away from inefficient paths with excessive reversals.

Memoization During a mission, the robot is dispatched to inspect a set of n areas of interest (AOIs) specified as a polygon on the map. Each AOI is represented in the planner by the two points where the track polyline crosses its boundary (i.e., the entry and exit points). The mission planner enumerates every candidate AOI ordering and every entry-direction choice. This is essentially a traveling salesman problem that is tractable at the scale of a single inspection mission with $n = 4$, giving $4! \cdot 2^4 = 384$ candidate tours.

A naïve implementation would invoke A* once per leg of every candidate tour. Each tour visits $2n + 1$ points (the robot's current pose plus the two boundary crossings of each AOI), giving $2n$ legs, hence $2n \cdot n! \cdot 2^n = 3{,}072$. A* calls for $n = 4$. However, every candidate tour is a permutation over the same small set of query points, so the number of distinct unordered point pairs that ever need a length query is bounded by $\binom{2n+1}{2}$ ($= 36$ for $n = 4$). Therefore, the along-track length of each pair is cached the first time it is computed and reused across all tours that revisit it. This reduces A* invocations from 3,072 to 36 for $n = 4$, and scales favorably with mission size.

3.3 Localization enhanced by railway maps

The map and concepts introduced in Paper A enable several promising strategies to improve the accuracy and robustness of an onboard localization system. At the current stage of this research, experiments conducted in Chapter 5 purely rely on GNSS for positioning, and only use a VERNE map for path planning. Therefore, this section presents theoretical insights to answer RQ3 and to guide the future development of the VIDAR prototype.

Information on the spatial geometry of the current track of the robot sets a boundary on the cross-track positioning error. Indeed, it is possible to make the robot stop if the current position estimation deviates too much from the expected path. As a matter of fact, results discussed in Section 5.2.1 indicate that the Swedish digital map has an accuracy roughly equivalent to that obtained with an RTK GNSS receiver in the countryside. In addition, the map allows the detection of the catastrophic scenario where the robot crosses a rail switch and ends up on a different track than the intended one.

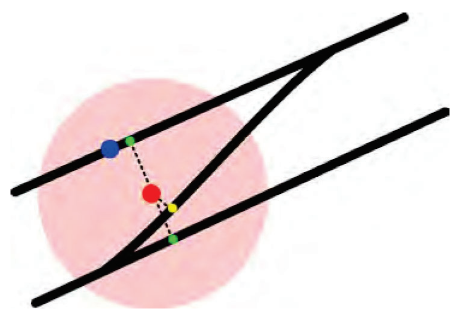

Figure 3.2: Initial localization scenario (reproduced from Paper A). The black lines represent the railway tracks, the blue dot indicates the true location. The transparent red circle represents the uncertainty region of the current position estimate, and the red dot indicates its center. The yellow dot results from the projection of the red dot on the closest track. The green dots denote projections on two other candidate tracks.

Depending on the deployment context, the robot might not be aware of the exact track it is starting on (or the initial hypothesis might be wrong). By using the “containing atom query”, a VERNE map enables the identification of situations where multiple tracks fall within the area of uncertainty of the current position estimation (Fig. 3.2). In this case, several position hypotheses on different tracks need to be maintained simultaneously, for example via the use of a particle filter [16]. This has not been implemented on the robot yet, and shall be explored in the future.

Chapter 4 introduces odometry algorithms capable of locating the robot without GNSS coverage. However, these methods accumulate error over time. To cope with this, future work will use onboard sensors like cameras or magnetometers to detect the presence of specific infrastructure elements nearby (switches, landmarks, etc.). If the positions of such elements are contained in the digital map, it is possible to deduce the position of the robot relatively to their known geographic coordinates.

As opposed to VERNE which is mainly focused on spatial representation, the descriptive model "RailTopoModel" [32] and the RailML [53] data exchange format focus on associating metadata to railway tracks. A promising future direction is to use these tools to enrich the VERNE spatial data structure with information about landmarks or other elements detectable by onboard sensors.

Chapter 4

Perception systems for self-localization and obstacle detection

To safely accomplish its mission, the VIDAR must be equipped with multiple sensors to track its position and detect obstacles. This chapter introduces the set of sensors chosen and the algorithms used.

4.1 Sensor suite overview

Accurate real-time localization is required to safely navigate the railway network, and to produce geospatial data with accurate spatial information. Relying solely on global navigation satellite systems (GNSS) is not a valid option for railway vehicles due to the risk of intermittence, outage and large errors in crucial parts of the network like tunnels and railway stations [40, 52]. Hence, to supplement GNSS receivers, modern methods fuse data coming from multiple sensors of different natures [14, 36, 46, 66]. Nevertheless, recent works still highlight the need to make localization algorithms more robust [11, 23, 58, 83]. As a matter of fact, railway tracks present many challenges, such as wheel slippage, fog, seasonal changes and featureless environments. These degeneracies each affect different sensors typically used in this context (wheel encoders, Doppler radars, cameras and lidars) [57, 61].

Bearing this in mind, a complementary set of sensors is selected for the VIDAR prototype and summarized in Table 4.1 It should be noted that the lidar, cameras and IMU are placed on an anti-vibration platform. However, occasional jolts can still disturb recordings, for example, when passing a rail switch.

On top of their contribution to the localization system, the data collected

Sensor	Contribution to localization	Sensitivity to external conditions
GNSS receiver	Absolute geographic positioning. Real-time kinematic (RTK) corrections can be provided by a nearby base stations to achieve centimeter-level precision.	Accuracy strongly depends on the current satellite coverage and the distance to the nearest RTK base station.
3D lidar	Provides a 3D point cloud indicating the relative distance to surrounding objects.	Performance degrades in heavy precipitation (rain, snow) and dense fog due to signal attenuation and scattering. Illumination invariant.
Monocular camera	Enables the identification of visual features which can be tracked to infer the robot's ego-motion. Also allows semantic interpretation of images and long-term landmark recognition.	Vulnerable to rain, snow and fog, as well as poor or varying lighting conditions (darkness, glare). Risk of lens occlusion from environmental debris. Subject to motion blur.
Inertial measurement unit (IMU)	Gives high-frequency inertia-based estimations of the instantaneous 3D acceleration and 3D angular velocity. Typical noise levels do not allow the direct integration of accelerations to correctly estimate position.	Affected by vibrations and temperature variations.
Tachometer (or wheel encoder)	Indirect measurement of traveled distance and longitudinal speed by counting the number of revolutions per minute.	Prone to wheel slip/slide, especially on tracks covered with leaves, water or ice, and when accelerating or braking.
Magnetometer	When placed near the top of a rail, a 3D magnetometer measures the magnetic field of the Earth after it has been uniquely distorted by the ferromagnetic elements of the track at this position. This can be used to deduce the current spatial position of the robot [28].	Highly susceptible to interference from catenaries and train traffic on parallel tracks.

Table 4.1: Comparison of the sensors selected for onboard localization.

by these sensors can be used for inspection [7, 8, 21]. For instance, the combination of lidar and camera allows the generation of a 3D representation of the track and surrounding infrastructure (structure from motion) [85]. Similarly, data from an onboard accelerometer can help recognize the health condition of the track [78].

4.2 Self-localization algorithms

To track its own position in real time, a robot can benefit from one of the sensors listed in Table 4.1 or a combination of them via sensor fusion [82]. Common fusion strategies include recursive Bayesian filters [6] and factor graph optimization [79], often supplemented with a dynamic model of the robot's motion.

Self-localization can be achieved with or without the mapping. Simultaneous localization and mapping (SLAM) approaches build and update a map of the environment by processing raw sensor data together with the latest location estimates. This enables the correction of accumulated drift via the detection of "loop closures", for example, when matching current camera recordings with previously saved georeferenced images [31, 33].

However, within the scope of this thesis, it is considered that the environment being currently explored is new, such that loop closures are not possible. Instead, "odometry" algorithms are evaluated in the rail context. The term "odometry" can be employed regardless of what sensor is used to track the relative distance traveled from a starting position. The most well-known sensor used for odometry is the wheel encoder, but many more sensors can provide such information.

4.2.1 Lidar-based odometry algorithms

This section introduces the existing open-source odometry algorithms that were adapted for the VIDAR prototype and its lidar. They were chosen for their remarkable precision in the field of robotics.

KISS-ICP

KISS-ICP [76] is a minimalist, lidar-only approach. It achieves robust, real-time performance without complex feature extraction by operating directly on raw point clouds stored within a voxel grid (i.e., lidar points are stored in 3D cubes indexed by a spatial hash table). For state estimation, the algorithm relies on a constant velocity model to predict and deskew the next scan. It then snaps the scan into place using point-to-point registration to directly find the closest matches.

FAST-LIO2

FAST-LIO2 [81] is a complex, highly efficient framework that tightly couples lidar and IMU data to achieve fast and accurate mapping. At its core, it utilizes an iterated error-state Kalman filter (IESKF) to fuse the high-frequency IMU measurements with raw lidar data for precise state estimation. Instead of extracting features, it employs point-to-plane registration, aligning new raw points to flat surfaces calculated from the map. This continuous mapping process is handled by an incremental k-d tree (ikd-Tree) data structure, which allows rapid updates.

4.2.2 Evaluation methodology and results

This work aims to verify if state-of-the-art odometry methods are directly applicable for the railway environment, and if their precision is competitive with traditional tachometer dead reckoning when GNSS is not available.

Results were obtained offline via the playback of previously recorded real sensor data. To serve as a ground truth arc-length path, a 2D cubic B-spline approximation of a test track is built from GNSS points collected during multiple runs (see Section 5.2.1). During evaluations, the received GNSS positions are projected on the reference B-spline to serve as a real-time reference ground truth. The positions estimated by the tested SLAM algorithms are relative to their own 3D frames which have an arbitrary origin and orientation. Therefore, after a run is complete, each trajectory is projected on the 2D plane and aligned to the reference B-spline via Umeyama's method [75]. The aligned metric coordinates are then reprojected onto the spline to derive a scalar distance-along-track from a shared starting position. Finally, the one-dimensional absolute trajectory error (ATE) of every tested odometry system are computed [71, 84].

It is expected that KISS-ICP will experience greater drift compared to FAST-LIO2 since the latter receives additional information from an IMU, allowing it to better stabilize trajectories even when the lidar faces degeneracies. Trajectories are shown in Fig. 4.1, and the ATE of each odometry algorithm is indicated, indicating a relatively higher precision of FAST-LIO2 compared to KISS-ICP on this dataset.

Fig. 4.2 offers a visual comparison of the errors accumulated by the traditional tachometer-based odometry and the lidar-based odometry methods. The results are revealing of the intrinsic differences between them. During

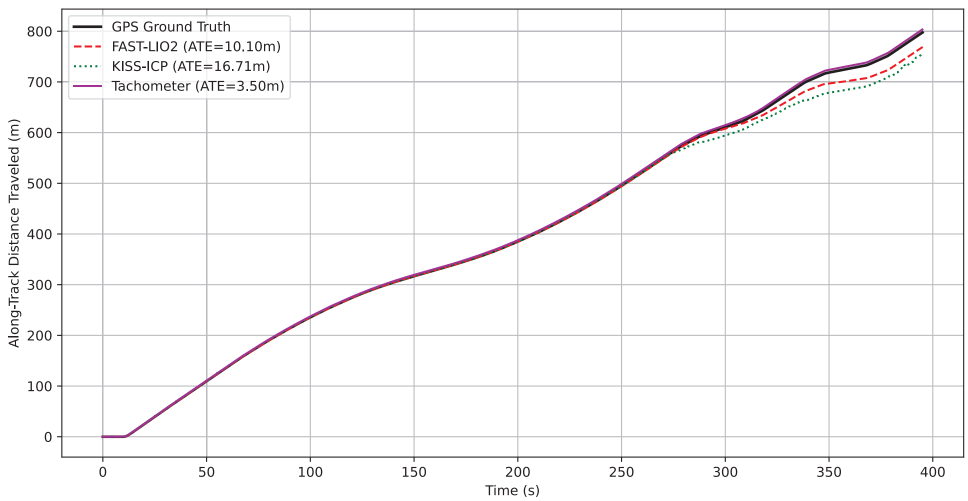

Figure 4.1: Along-track trajectories estimated by different odometry sources.

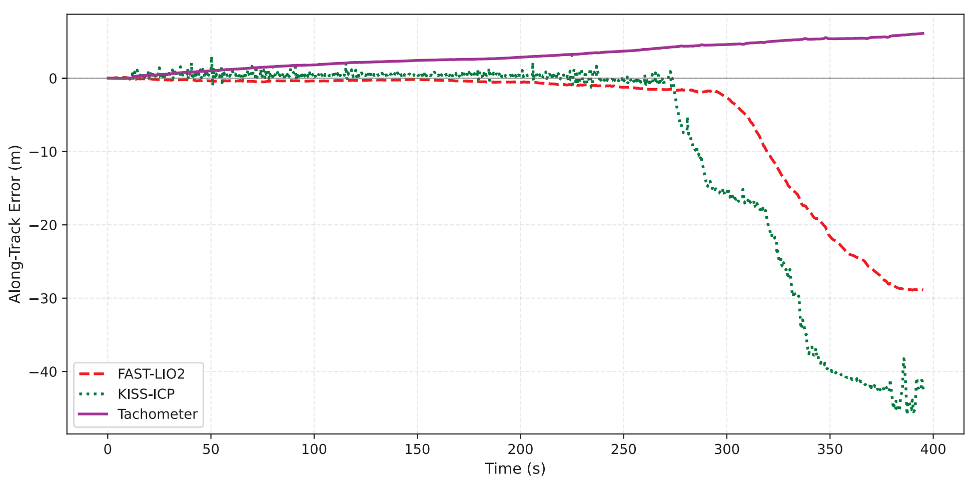

Figure 4.2: Along-track trajectory error over time for different odometry sources.

the first half of the data set, the robot is moving in a forest with many geometric features (Fig. 4.3a). In such an environment, lidar-based algorithms are capable of reaching high levels of precision by using the distance to external objects as a reference, thus outperforming the tachometer. It should be noted that although the tachometer suffers from low and constant drift caused by imperfect calibration, achieving a perfect calibration is impossible in practice due to wear and changes in temperature affecting the wheel radius.

On the other hand, in the second part of the data set, the robot leaves the forest and continues on a track surrounded by agricultural land with significantly fewer geometric features (Fig. 4.3b). From there, the accuracies of KISS-ICP and FAST-LIO2 degrade rapidly, reaching errors well beyond those returned by the tachometer odometry.

These results emphasize the need to fuse complementary sensors to achieve robust and precise localization. Future work will investigate different ways to fuse the VIDAR sensors by building on existing multi-sensor odometry algorithms. Moreover, map-matching [69] will be used to align the estimates of the robot's position with the tracks of a digital map.

(a) Start of the data set in a feature-rich forest. *(b) End of the data set with few geometric features.*

Figure 4.3: Camera images from the data set used for evaluating lidar odometry algorithms.

4.3 Obstacle detection algorithms

To fully cover the scope of RQ4 the issue of obstacle detection must be answered in case a person or animal suddenly appears in front of the robot. This section details the theory behind the proposed obstacle detection system. Preliminary field experiments mentioned in Paper B show the viability of the solution by placing a humanoid wooden figure on the track (Fig. 4.4). The prototype systematically started to slow down for obstacles closer than 50 m and stopped around 3 m before them. Further experiments need to be conducted to quantify and validate the proper response of each component.

Figure 4.4: Obstacle detection experiment.

4.3.1 Long range detection

A Velodyne VLP-32C lidar is used to detect obstacles from afar. The lidar produces thirty-two different horizontal lidar scans simultaneously at a rate of 10 Hz, each at different elevation angles spanning from $+15°$ to $-25°$ (they are highly concentrated around the $0°$ horizon, but begin to spread significantly near the extremes, such that the angular gaps between them vary from $0.333°$ to $9°$). On the other hand, the lasers are shot at azimuth angles evenly spaced by $0.2°$.

Raw point clouds are received every 100 ms, and before being used for obstacle detection, they are preprocessed as follows:

1. Points belonging to the front and the back areas of the robot are stored in two separate point clouds.

2. A radius outlier removal filter is applied. It removes a point if its number of neighbors in a search radius $R = 40\,\text{cm}$ is smaller than $K = 4$.

3. For each filtered point cloud, its centroid and its distance to the lidar are computed.

4. A moving average filter of size 10 is applied to the distance of each centroid, resulting in two values representing the distance of obstacles in front and at the back of the robot.

For illustration purposes, let us consider a flat obstacle of width 40 cm and height 2 m. If at least four points land on the obstacle within a sphere of radius $R = 0.4$, it can be detected despite the outlier filter.

For simplification and to give an approximation of the maximum detection distance, a lidar point P at 0° elevation is considered here, so its vertical neighbors are shot 0.333° under and above it, respectively, and land at a distance d_v from P. Likewise, its horizontal neighbors are shot at 0.2° to the left and right of it, respectively, and land at a distance d_h from P, see Fig. 4.5.

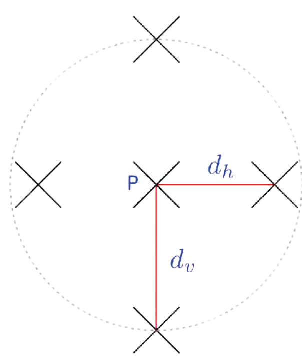

Figure 4.5: Projection of a sphere of radius R centered on the lidar point P such that its four neighbors land on a flat obstacle at a distance D.

A distance D to the obstacle such that a sphere of radius R centered on the lidar point contains 4 points is given by the formula:

$$\max(d_h, d_v) = \max(D\tan(0.2°), D\tan(0.333°)) = R \tag{4.1}$$

Where d_h and d_v are the respective distances between P and its immediate horizontal and vertical neighbors. Therefore, with $R = 40\,\mathrm{cm}$, this system can detect an obstacle at an approximate maximal distance $D \approx 70\ m$, which is about one third of the sensor's maximum range.

4.3.2 Short range detection

To act as a backup system and to detect smaller obstacles that could lie under the line of sight of the lidar, two sets of four single-point lidar sensors (TFMini Plus) are placed at the front and at the back of the chassis respectively. They are mounted every 40 cm horizontally, at a height of 50 cm from the ground. According to their data sheet, these sensors can detect a human closer than 7 m and at least 10 cm away. Every 10 ms, the shortest distance recorded by the four sensors is stored into a moving average filter of size 10.

Chapter 5

Experimental validation and field tests

5.1 Small-scale indoor testbed

To support the early stages of development toward a first VIDAR prototype ready for field testing, a small-scale indoor model of railway tracks was used (Fig. 5.2a).

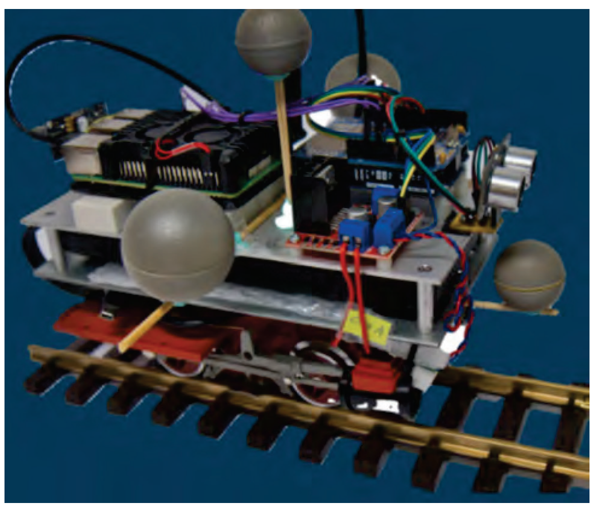

Figure 5.1: Miniature prototype of the VIDAR.

A modified miniature train, shown in Fig. 5.1, travels on the network. To emulate the decentralized architecture of the VIDAR, the motor of the chassis is controlled by an Arduino board connected to a Raspberry Pi 3B. Likewise, an ultrasonic distance sensor allows the detection of frontal obstacles. Speed is estimated via a hall effect sensor measuring rotational frequency of the

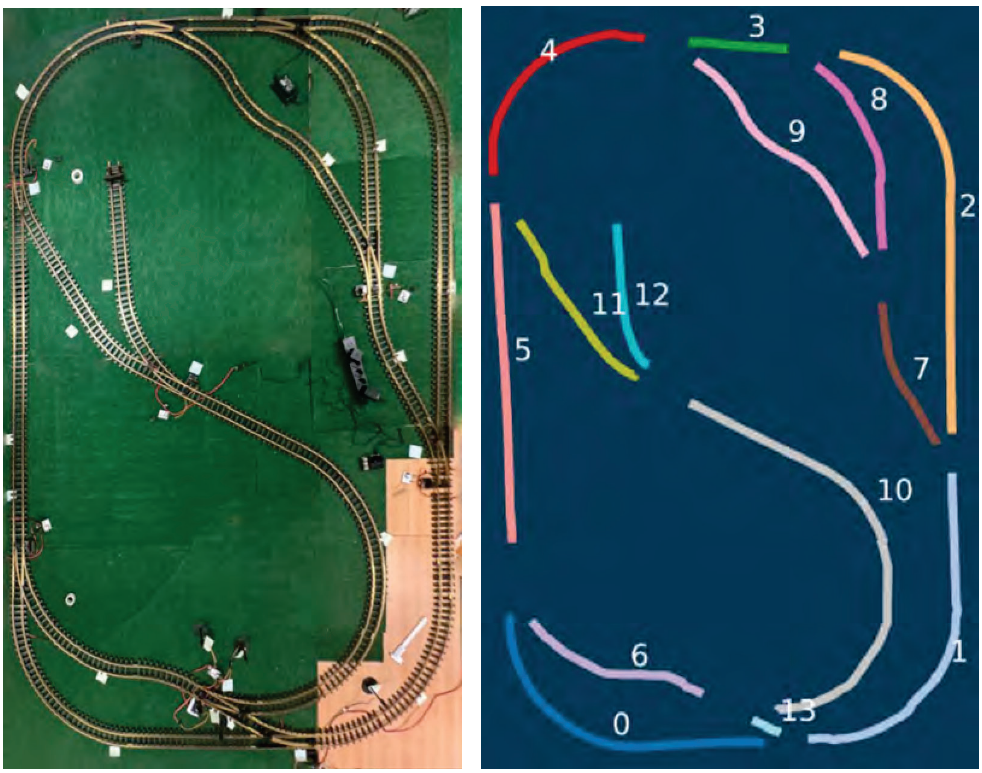

(a) Top view of the network.

(b) VERNE atoms of the network.

Figure 5.2: Network of the small-scale indoor testbed.

wheels. Successful tests on this preliminary version provided insights that helped build the full-scale VIDAR prototype.

The subsequent goal was to verify the VERNE-based navigation system. To this end, a related Bachelor's thesis project [3] upgraded the testbed with remotely-operated switches to control them in real-time upon request. Furthermore, a Qualisys infrared camera tracking system was configured to (1) create a VERNE map of the network by segmenting the tracks into atoms, as seen in Fig. 5.2b, and (2) track the absolute position of the robot on that map.

This improved testbed enabled the validation of the path planning software currently used on the VIDAR prototype. The miniature train received a target destination on the network and successfully planned its path across switches to reach it.

5.2 Field testing

Trafikverket provided a railway in Sweden isolated from regular train traffic and dedicated to testing. This track is characterized by surrounding grassland or forest and features a few switches, no overhead catenary wire, and several utility poles. The total accumulated track length is approximately 3.5 km, including two parallel tracks spanning 1 km. All tests were performed via a control station connected to the mobile network.

5.2.1 GNSS trajectory repeatability

An experiment was conducted to assess the stability of GNSS fixes received when the robot is traveling along the same path multiple times with several hours, days or months between runs. Simultaneously, a B-spline was fitted on the fixes to build a high-resolution reference trajectory that could potentially be more accurate than the digital map. In fact, the data used to build the digital map comes from a national database where an uncertainty of roughly ± 2 m has been empirically observed with a 90% confidence interval.

Fig. 5.3 is a zoomed satellite view[1] of the start of the testing area where the GNSS fixes, the fitted B-spline, and the digital map have been superposed. By plotting the cumulative distribution of the fixes' distances to the map or the B-spline (Fig. 5.4), it can be seen that the cloud of GNSS fixes is contained within 3.5 m of the two track representations, which is aligned with their expected accuracy.

[1] The projected coordinate system used in this work is SWEREF 99 TM (EPSG:3006)

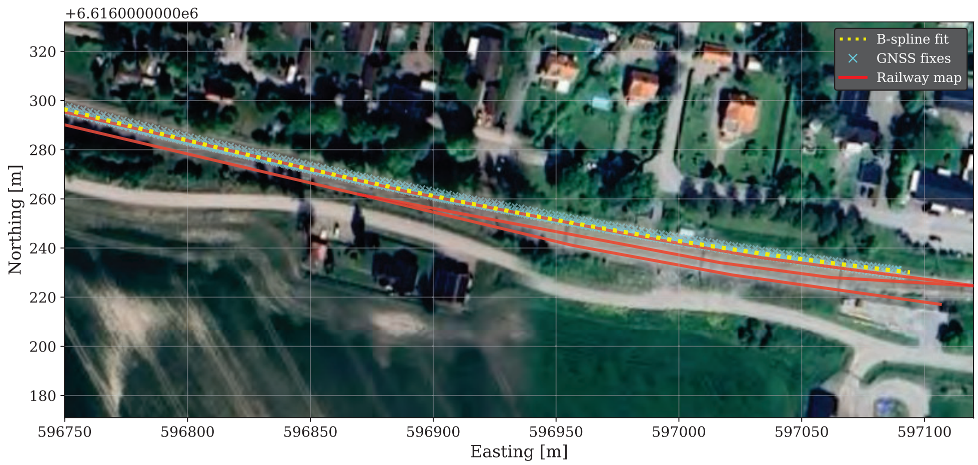

Figure 5.3: B-spline fitted on GNSS fixes accumulated from independent replicates of the same trajectory on the top track. The figure shows the first 300 m *of the full* 2 km *B-spline built from 5609 GNSS fixes. Background imagery: Google, Airbus, Maxar Technologies (2026).*

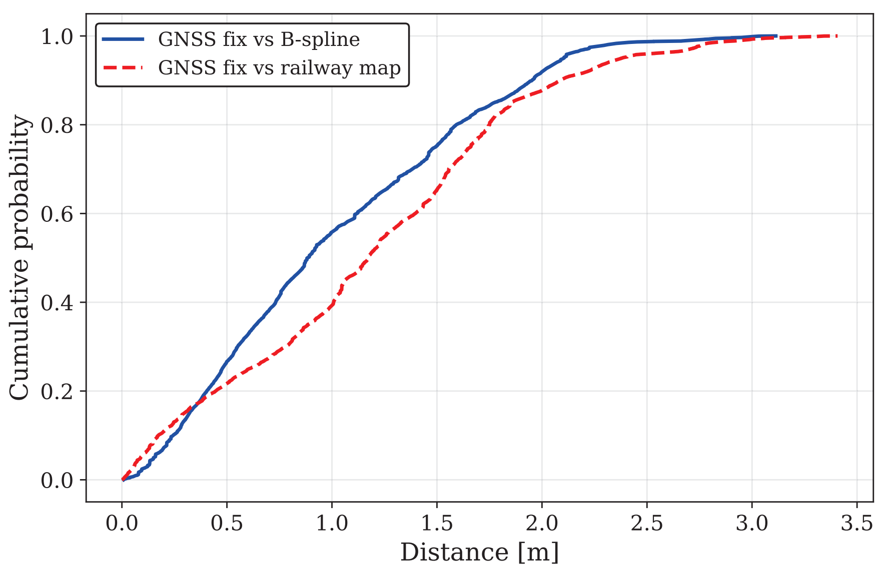

Figure 5.4: Cumulative distribution of the distances from GNSS fixes to either the B-spline or the railway map.

5.2.2 Simulated fault detection: unexpected small object on rail

The typical operating conditions considered in the field tests are as follows. First, an operator deploys the robot on an isolated track and provides multiple target inspection zones via geographic coordinates. Second, the robot uses a digital map and its localization system to plan an inspection path, taking into account potential track switches. For the time being, the robot relies on a GNSS receiver to track its position. Finally, once an inspection zone is reached, the robot goes at a lower speed and records data, while saving the geographic position of any faults detected along the way.

In this experiment, faults are physically simulated by placing metal plates on the rail with duct tape. They are detected with the onboard accelerometer by using a threshold trigger value on the vertical acceleration.

To perform a diagnostic on detected faults, a camera aiming toward the rail is placed above one of the wheels. Images from a 4 Hz video feed are continuously placed in a circular buffer, but the buffered images are only saved after a fault has been detected. This allows the identification of faults, as shown in Fig. 5.5.

Figure 5.5: Image of simulated fault B (metal plate taped to the rail), taken in real-time by the onboard track camera.

To assess the reliability and precision of the fault detection system, data from six independent runs was collected on a 200 m path with four faults (A–D). Results are displayed in Fig. 5.6 and compiled in Table 5.1. σ_{along} and σ_{cross} denote the standard deviation of the detected positions decomposed into the along-track and cross-track directions, respectively. DRMS (distance root mean square), computed as $\sqrt{\sigma_{\text{along}}^2 + \sigma_{\text{cross}}^2}$, represents the radius of a circle that contains ∼65% of the detections around the mean position.

All faults were detected with a standard deviation smaller than 2 m, which is deemed precise enough to facilitate future repair missions. It should be noted that the cross-track standard deviation for fault D is notably smaller than for others. This discrepancy can be attributed to the anisotropic nature of GNSS errors, which are typically characterized by a covariance ellipse. Because fault D was positioned on a track with a noticeably different heading than the others, the projection of this error ellipse yielded a reduced cross-track component.

It should be noted that the statistics shown in this table were calculated using the B-spline of the track presented in Fig. 5.3. Since both tracks are parallel, the results are slightly less accurate but remain valid. The reason behind this choice is that only a limited amount of data was available for the track where faults were placed, contrary to the track used to build the B-spline.

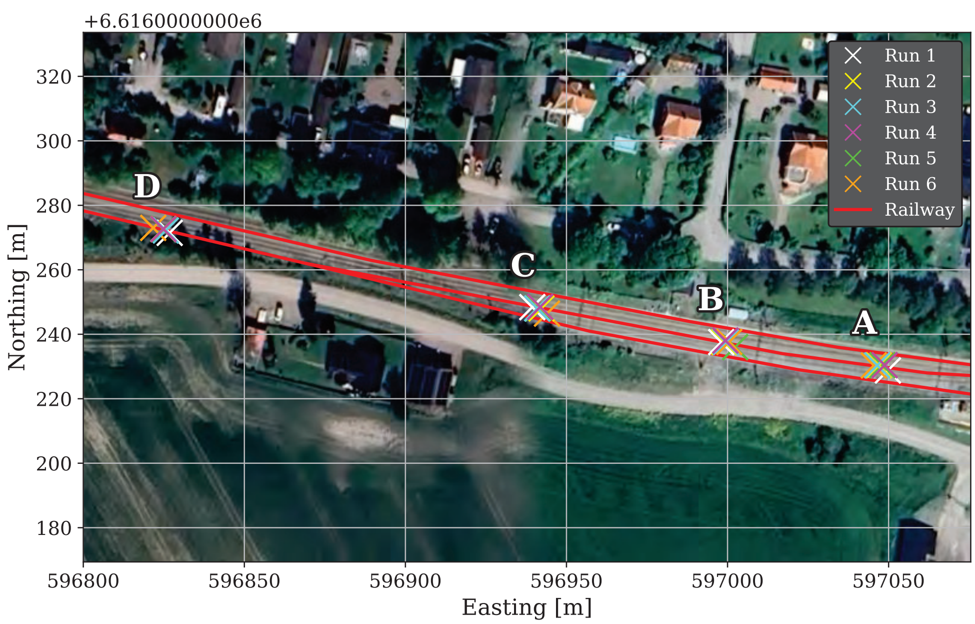

Figure 5.6: Evaluation of the detection repeatability of each simulated fault (A–D) across multiple runs. Background imagery: Google, Airbus, Maxar Technologies (2026).

Variations in the estimated position of the faults are due to GNSS uncertainties (timewise and distance-wise, since GNSS fixes are received every one

Fault	σ_{along} [m]	σ_{cross} [m]	DRMS [m]
A	1.50	0.49	1.58
B	1.43	0.49	1.51
C	1.46	0.50	1.54
D	1.60	0.26	1.62

Table 5.1: Localization repeatability of four simulated track faults, each detected across six independent measurement runs.

second with an approximate precision of three meters) and software delays (time between the bump, the detection of the abnormal vertical acceleration, and the storage of the current position in the onboard computer's memory).

5.2.3 Track health estimation from 2D lidar scans

To serve as an illustrative inspection capability, a low-cost 2D lidar (UST-10LX) is used to estimate the health condition of a track by collecting statistics on track gauge measurements.

The lidar is placed at the front of the VIDAR and is oriented toward the ground. Every scan provides the distances detected by each light ray at different angles, which can be transformed into a point cloud, as presented in Fig. 5.7. Two narrow angular windows ($\approx 5°$ wide, centered around -65° and +52°) are defined so that each window intersects one rail. Within each window, the smallest measured distance (range) is taken as the estimated rail-edge point, and the gauge d is recovered as the chord between the two points via the law of cosines:

$$d = \sqrt{r_1^2 + r_2^2 - 2 r_1 r_2 \cos(|\theta_1| + |\theta_2|)}$$

Where r_i and θ_i are the range and angle of the estimated rail-edge point of rail $i \in \{1, 2\}$.

Per-scan gauge estimates deviating from the standard gauge value of 1435 mm by more than ±50 mm are rejected as outliers. The sequence of estimates is then compared between the two segments in terms of mean, extrema and standard deviation. A subset of these estimates is presented in Fig. 5.8. The contrast in dispersion and bias between the two test tracks provides an indicator of track-health degradation. Results are summarized in Table 5.2.

To contextualize the severity of these geometric deviations, the estimated gauge extrema can be compared against the immediate action limits (IAL)

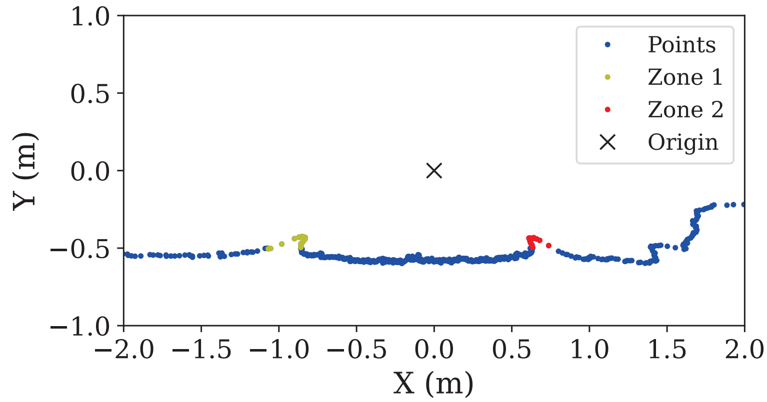

Figure 5.7: 2D point cloud of a degraded test track.

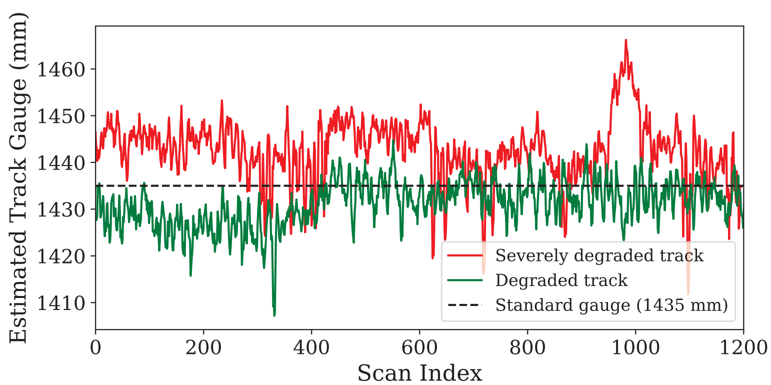

Figure 5.8: Track gauge estimation of two test tracks (5-point uniform moving-average filter).

Metric	Severely degraded track	Degraded track
Mean [m]	1.443	1.432
Std [m]	0.011	0.009
Min [m]	1.387	1.389
Max [m]	1.473	1.459

Table 5.2: Track health metrics collected on two test tracks.

defined by European Standard EN 13848-5 for isolated defects. For speeds up to 80 km/h, the standard dictates a maximum allowable gauge widening of 35 mm and a maximum narrowing of −11 mm relative to the nominal 1435 mm gauge.

As shown in Table 5.2, the severely degraded track exhibits both extreme widening (38 mm) and extreme narrowing (−48 mm), violating both normative IAL thresholds. In contrast, the degraded track remains within the widening safety limit with a maximum peak at 24 mm, though it also exceeds narrowing limits with a minimum at −46 mm. Furthermore, the higher standard deviation of the severely degraded track ($\sigma = 0.011$ m) compared to the degraded track ($\sigma = 0.009$ m) indicates more local irregularities, confirming its more advanced state of structural deterioration.

Chapter 6

Conclusion and future work

6.1 Conclusion

This thesis provides the foundation needed for flexible, data-driven maintenance powered by autonomous rail-bound robots to support the growth of efficient and low-carbon railway transportation across Europe. The research was guided by four primary research questions (RQs), each of which has been answered through the development of the VIDAR concept and the validation of a modular prototype through real-world experiments facilitated by VERNE, a specialized spatial data structure. This work demonstrates the feasibility of deploying lightweight, autonomous robots capable of safely navigating on a real railway network and performing simulated inspection tasks.

Regarding the construction and utility of digital maps (RQ1 and RQ2), the VERNE framework was successfully implemented from open access data and extensively verified. By decomposing the network into “atoms” (indivisible track segments) and ‘junctions”, VERNE maps enable efficient graph-based path planning. Experimental validations confirmed that the entire railway networks of Sweden and France can be represented in less than 20 MiB of memory. Extrapolating from the benchmark, queries on this structure should complete within one millisecond on standard hardware. Therefore, complex, country-scale railway navigation is computationally feasible on typical hardware available for embedded applications.

RQ3 tackles the connection between localization accuracy and spatial data structures. Digital maps provide multiple benefits to onboard localization systems, notably in terms of reliability and accuracy. By constraining the robot’s position to a one-dimensional track, the cross-track error of any odometry source can be bounded, and 3D odometry sources sensitive to rotational errors can be anchored to the linear direction of the track, thus mitigating drift due to heading errors. In Chapter 4, two state-of-the-art lidar-based odometry algorithms were found to achieve lower drift than a GNSS-calibrated tachometer in areas rich in geometric features, but in environments with few features their precision quickly degrades. These results are encouraging and will guide future work where a railway map and different sensors will be combined with lidar to achieve levels of precision without GNSS.

The central question of safety (RQ4) has been answered to a level deemed sufficient for experiments conducted in a controlled environment. The modular architecture of the VIDAR allows two obstacle detection systems to run in parallel. A long-range 3D lidar is employed to detect obstacles at distances up to 70 m and initiate controlled deceleration early. Short-range single-point lidars handle obstacles closer than 7 m at lower elevations. A theoretical analysis and repeated qualitative tests described in Chapter 4 were run under favorable weather conditions. These results prove the architectural viability of the presented safety stack. Furthermore, the addition of redundant fail-safe mechanisms safely handles situations where communication is lost, sensors are failing or localization is uncertain.

In summary, this thesis establishes that autonomous railway inspection is technically achievable with current technology. The VIDAR prototype serves as a functional proof of concept, bridging the gap between theoretical autonomous navigation and practical railway operations in a controlled environment.

6.2 Future directions: toward deployment and scalability

While the VIDAR prototype validates the core concepts of autonomous railway navigation, significant challenges remain before such systems can be systematically deployed at full scale. Moving beyond controlled environments requires addressing three interconnected pillars: advancing robust perception, ensuring operational integration, and achieving system scalability.

6.2.1 Advancing robust perception and localization

The most critical technical barrier preventing real-world deployment is achieving reliable, real-time localization independent of GNSS to allow operations in degraded environments such as tunnels and deep cuttings. Future research must evaluate the sensitivity of SLAM algorithms in these environments. A potential path is to enhance existing methods like FAST-LIO2 that tightly couple lidar and IMU. Specifically, integrating data from a digital map and a tachometer will help bounding drifting errors, even in feature-poor areas. Moreover, it shall be studied how cameras and magnetometers can provide absolute positioning information to compensate the drift accumulated during extended GNSS outage. The next milestone is to run experiments under difficult weather conditions, and quantify the impact on each sensor individually and on the position estimation resulting from the fusion of all these sensors.

Beyond localization, the perception pipeline must be improved to cover more complex operational conditions. The current obstacle detection system is likely vulnerable to weather-related degeneracies, such as rain, fog, and the presence of uncontrolled vegetation. Furthermore, its effectiveness when traveling at higher operational speeds is uncertain. Future iterations could also integrate deep learning-based semantic segmentation to reliably distinguish static infrastructure from dynamic entities. Moreover, such models could be utilized to automatically classify specific defect types. Finally, future work must include testing specialized inspection sensors to validate real-time onboard processing.

6.2.2 Operational integration and safety certification

Deploying autonomous robots into the railway ecosystem requires navigating a complex regulatory and infrastructural landscape. Currently, the prototype operates on isolated tracks and requires a full stop at switches, relying on manual human intervention to proceed. Future systems must be integrated with the signaling infrastructure in place, whether it be through communications-based train controls (CBTCs) or operations control centers (OCCs).

Regulatory frameworks must also evolve alongside the technology. The current grade of automation 4 (GoA4) standard is insufficient for the unique capabilities of non-standard vehicles like VIDAR (e.g., unplanned changes of behavior, physical interventions on the infrastructure and exploration of unsafe zones). Future work should contribute to alternative frameworks tailored for lightweight rail vehicles, defining strict protocols for interaction with human operators and traditional rolling stock. Furthermore, compliance with European functional safety standards (EN 50128, EN 50129) is mandatory. Research must therefore address the verification and validation methodologies for AI-driven components, alongside rigorous cybersecurity audits to protect the command and control links from malicious interference.

6.2.3 Scalability and fleet management

To fully realize the economic benefits of autonomous railway maintenance, individual robots must evolve into coordinated, autonomous fleets. As a critical preliminary step, initial future work should focus on simulated fleet coordination to validate complex interactions. This foundation will support the long-term development of distributed algorithms for dynamic multi-agent path planning, optimizing fleet efficiency while actively avoiding track occupancy conflicts in real time.

The high-frequency data collected by these fleets should ultimately feed into a live digital twin of the railway infrastructure to reflect the real-world state of the railway network at any given moment. To maximize the versatility of this fleet, future prototypes should be designed with standardized interfaces for modular payloads. This approach will allow a single chassis to be dynamically equipped for inspections, cleaning, or minor repair tasks depending on immediate mission requirements. In addition, potential synergies with other modern inspection technologies such as legged robots or unmanned aircraft should be investigated.

Bibliography

1. Adler, R. & Klaes, M. *Assurance cases as foundation stone for auditing AI-enabled and autonomous systems: Workshop results and political recommendations for action from the ExamAI project* in *International Conference on Human-Computer Interaction* (2022), 283–300.
2. Ahsan, N., Hewage, K., Razi, F., Hussain, S. A. & Sadiq, R. A critical review of sustainable rail technologies based on environmental, economic, social, and technical perspectives to achieve net zero emissions. *Renewable and Sustainable Energy Reviews* **185,** 113621 (2023).
3. Andreasson, A., Engström, J., Eweman, G., Nilsson, A., Sitell, E. & Åkerlund, I. Realtidsnavigering av autonom inspektionsrobot i komplexa rälsnätverk: En småskalig implementation (2025).
4. Bartak, R. *Tackling train routing via multi-agent pathfinding and constraint-based scheduling* in *Proceedings of the 14th International Conference on Agents and Artificial Intelligence* (2022).
5. Belsare, K., Rodriguez, A. C., Sánchez, P. G., Hierro, J., Kołcon, T., Lange, R., Lütkebohle, I., Malki, A., Losa, J. M., Melendez, F., *et al.* in *Robot Operating System (ROS) The Complete Reference (Volume 7)* 3–55 (Springer, 2023).
6. Bergman, N. Recursive bayesian estimation. *Department of Electrical Engineering, Linköping University, Linköping Studies in Science and Technology. Doctoral dissertation* **579** (1999).
7. Bianchi, G., Fanelli, C., Freddi, F., Giuliani, F. & La Placa, A. Systematic review railway infrastructure monitoring: From classic techniques to predictive maintenance. *Advances in Mechanical Engineering* **17,** 16878132241285631 (2025).

8. Binder, M., Mezhuyev, V. & Tschandl, M. Predictive maintenance for railway domain: A systematic literature review. *IEEE Engineering Management Review* **51,** 120–140 (2023).
9. Böhringer, F. & Geistler, A. Location in railway traffic: Generation of a digital map for secure applications. *WIT Trans. on The Built Environment,* 459–468 (2006).
10. Borndörfer, R., Klug, T., Schlechte, T., Fügenschuh, A., Schang, T. & Schülldorf, H. The freight train routing problem for congested railway networks with mixed traffic. *Transportation Science* **50,** 408–423 (2016).
11. Bujanca, M., Shi, X., Spear, M., Zhao, P., Lennox, B. & Luján, M. *Robust slam systems: Are we there yet?* in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (2021), 5320–5327.
12. Chenariyan Nakhaee, M., Hiemstra, D., Stoelinga, M. & van Noort, M. *The recent applications of machine learning in rail track maintenance: A survey* in *International conference on reliability, safety, and security of railway systems* (2019), 91–105.
13. Daniyan, I., Mpofu, K. & Nwankwo, S. Design of a robot for inspection and diagnostic operations of rail track facilities. *International Journal of Quality & Reliability Management* **40,** 653–673 (2023).
14. Datar, A., Pokhrel, A., Nazeri, M., Rao, M. B., Rangwala, H., Pan, C., Zhang, Y., Harrison, A., Wigness, M., Osteen, P. R., *et al. M2p2: A multi-modal passive perception dataset for off-road mobility in extreme low-light conditions* in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (2025), 13690–13696.
15. De Donato, L., Flammini, F., Marrone, S., Mazzariello, C., Nardone, R., Sansone, C. & Vittorini, V. A survey on audio-video based defect detection through deep learning in railway maintenance. *IEEE Access* **10,** 65376–65400 (2022).
16. Djuric, P. M., Kotecha, J. H., Zhang, J., Huang, Y., Ghirmai, T., Bugallo, M. F. & Miguez, J. Particle filtering. *IEEE signal processing magazine* **20,** 19–38 (2003).
17. Durlik, I., Miller, T., Kostecka, E., Zwierzewicz, Z. & Łobodzińska, A. Cybersecurity in autonomous vehicles—are we ready for the challenge? *Electronics* **13,** 2654 (2024).

18. European Commission. *Sustainable and Smart Mobility Strategy – putting European transport on track for the future* Communication from the Commission COM(2020) 789 final (European Commission, Brussels, 2020). `https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52020DC0789`.

19. Filip, A. Synergies between road and rail transport in the development of safe self-driving vehicles. *International journal of transport development and integration* **6,** 313–325 (2022).

20. Gerlach, K. & zu Hörste, M. M. *A precise digital map for GALILEO-based train positioning systems* in *2009 9th International Conference on Intelligent Transport Systems Telecommunications,(ITST)* (2009), 343–347.

21. Gonzalo, A. P., Entezami, M., Weston, P., Roberts, C. & Marquez, F. P. G. *Railway track and vehicle onboard monitoring: A review* in *E3S Web of Conferences* **409** (2023), 02014.

22. Guo, W., Liang, G., Ren, S. & Zeng, C. *Automatic crack detection method for high-speed railway box girder based on deep learning techniques and inspection robot* in *Structures* **68** (2024), 107116.

23. Hao, X., Zhao, Y., Ji, Y., Dai, L., Hao, P., Li, D., Cheng, S. & Yin, R. *What Really Matters for Robust Multi-Sensor HD Map Construction?* in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (2025), 1298–1304.

24. Hart, P. E., Nilsson, N. J. & Raphael, B. A formal basis for the heuristic determination of minimum cost paths. *IEEE transactions on Systems Science and Cybernetics* **4,** 100–107 (1968).

25. Heirich, O. Bayesian train localization with particle filter, loosely coupled GNSS, IMU, and a track map. *Journal of Sensors* **2016,** 2672640 (2016).

26. Heirich, O., Robertson, P., García, A. C., Strang, T. & Lehner, A. *Probabilistic localization method for trains* in *2012 IEEE Intelligent Vehicles Symposium* (2012), 482–487.

27. Heirich, O., Robertson, P. & Strang, T. *RailSLAM-Localization of rail vehicles and mapping of geometric railway tracks* in *2013 IEEE International Conference on Robotics and Automation* (2013), 5212–5219.

28. Heirich, O., Siebler, B. & Hedberg, E. Study of Train-Side Passive Magnetic Measurements with Applications to Train Localization. *Journal of Sensors* **2017,** 8073982 (2017).

29. Hensel, S. & Hasberg, C. *Probabilistic landmark based localization of rail vehicles in topological maps* in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems* (2010), 4824–4829.

30. Hensel, S. & Marinov, M. B. Comparison of Time Warping Algorithms for Rail Vehicle Velocity Estimation in Low Speed Scenarios. *Metrology and Measurement Systems* (2017).

31. Hess, W., Kohler, D., Rapp, H. & Andor, D. *Real-time loop closure in 2D LIDAR SLAM* in *2016 IEEE international conference on robotics and automation (ICRA)* (2016), 1271–1278.

32. Hlubuček, A. RailTopoModel and RailML 3 in overall context. *Acta Polytechnica CTU Proceedings* **11,** 16–21 (2017).

33. Ho, K. L. & Newman, P. Detecting loop closure with scene sequences. *International journal of computer vision* **74,** 261–286 (2007).

34. IEA. *The Future of Rail* tech. rep. Licence: CC BY 4.0 (IEA, Paris, 2019). `https://www.iea.org/reports/the-future-of-rail`.

35. Janić, M. A multidimensional examination of performances of HSR (High-Speed Rail) systems. *Journal of Modern Transportation* **24,** 1–21 (2016).

36. Jeong, S., Kim, H. & Cho, Y. *MARSCalib: Multi-robot, Automatic, Robust, Spherical Target-based Extrinsic Calibration in Field and Extraterrestrial Environments* in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (2025), 21062–21069.

37. Jiang, W., Chen, S., Cai, B., Wang, J., ShangGuan, W. & Rizos, C. A multi-sensor positioning method-based train localization system for low density line. *IEEE Transactions on Vehicular Technology* **67,** 10425–10437 (2018).

38. Jing, G., Qin, X., Wang, H. & Deng, C. Developments, challenges, and perspectives of railway inspection robots. *Automation in construction* **138,** 104242 (2022).

39. Joly, L.-R., Lacorre, V. & Wolff, K. VERNE: A Spatial Data Structure Representing Railway Networks for Autonomous Robot Navigation. *IEEE Open Journal of Vehicular Technology* **7,** 1–14 (2025).

40. Kliman, A., Grosch, A. & Garcia Crespillo, O. Permanent train-side GNSS multipath characterization considering time-correlation for safe railway localization. *Engineering Proceedings* **88,** 71 (2025).

41. Kröper, B., Lauer, M. & Spindler, M. *Using the ferromagnetic fingerprint of rails for velocity estimation and absolute localization of railway vehicles* in *2020 European Navigation Conference (ENC)* (2020), 1–10.

42. Kumar, A. & Harsha, S. A systematic literature review of defect detection in railways using machine vision-based inspection methods. *International Journal of Transportation Science and Technology* **18,** 207–226 (2025).

43. Lauer, M. & Stein, D. A train localization algorithm for train protection systems of the future. *IEEE Transactions on Intelligent Transportation Systems* **16,** 970–979 (2014).

44. Lee, J., Kim, B., Kim, D. & Cho, S. Real-time urban railway damage detection-based cascade deep learning networks using automated inspection robot. *Journal of Civil Structural Health Monitoring,* 1–21 (2025).

45. Leutenegger, S. T., Lopez, M. A. & Edgington, J. *STR: A simple and efficient algorithm for R-tree packing* in *Proceedings 13th international conference on data engineering* (1997), 497–506.

46. Li, C., Zhang, Y., Yu, Z., Liu, X. & Shi, Q. *A robust visual slam system for small-scale quadruped robots in dynamic environments* in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (2024), 321–326.

47. Li, D., Jin, Y., Sun, Y., A, Y., Yu, H., Shi, J., Hao, X., Hao, P., Liu, H., Li, X., *et al.* What foundation models can bring for robot learning in manipulation: A survey. *The International Journal of Robotics Research,* 02783649251390579 (2024).

48. Li, H., Zhu, Q., Zhang, L., Ding, Y., Guo, Y., Wu, H., Wang, Q., Zhou, R., Liu, M. & Zhou, Y. Integrated representation of geospatial data, model, and knowledge for digital twin railway. *International Journal of Digital Earth* **15,** 1657–1675 (2022).

49. Liu, J., Cai, B.-g. & Wang, J. *A GNSS/trackmap cooperative train positioning method for satellite-based train control* in *17th International IEEE Conference on Intelligent Transportation Systems (ITSC)* (2014), 2718–2724.

50. Löffler, W. & Bengtsson, M. *Evaluating the Impact of Map Inaccuracies on Path Discrimination Behind Railway Turnouts* in *2022 IEEE 95th Vehicular Technology Conference:(VTC2022-Spring)* (2022), 1–5.

51. Macenski, S., Foote, T., Gerkey, B., Lalancette, C. & Woodall, W. Robot operating system 2: Design, architecture, and uses in the wild. *Science robotics* **7,** eabm6074 (2022).

52. Mikhaylov, D., Amatetti, C., Polonelli, T., Masina, E., Campana, R., Berszin, K., Moatti, C., Amato, D., Vanelli-Coralli, A., Magno, M., *et al.* Toward the future generation of railway localization exploiting RTK and GNSS. *IEEE Transactions on Instrumentation and Measurement* **72,** 1–10 (2023).

53. Nash, A., Huerlimann, D., Schütte, J. & Krauss, V. P. Railml† a standard data interface for railroad applications. *WIT Transactions on The Built Environment* **74** (2004).

54. Office of Rail and Road. *Cost Benchmarking of Network Rail's Maintenance and Renewals Expenditure: Annual Report 2023–2024* tech. rep. (Office of Rail and Road, London, UK, 2024). `https://www.orr.gov.uk/monitoring-regulation/rail/networks/network-rail/monitoring-performance/benchmarking-network-rails-costs` (2025).

55. Oh, J.-K., Lee, A.-Y., Oh, S. M., Choi, Y., Yi, B.-J. & Yang, H. W. *Design and control of bridge inspection robot system* in *2007 International Conference on Mechatronics and Automation* (2007), 3634–3639.

56. Otegui, J., Bahillo, A., Lopetegi, I. & Díez, L. E. A survey of train positioning solutions. *IEEE Sensors Journal* **17,** 6788–6797 (2017).

57. Palmer, A. W. & Nourani-Vatani, N. *Robust odometry using sensor consensus analysis* in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (2018), 3167–3173.

58. Patel, M., Yang, F., Qiu, Y., Cadena, C., Scherer, S., Hutter, M. & Wang, W. *Tartanground: A large-scale dataset for ground robot perception and navigation* in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (2025), 20524–20531.

59. Peleska, J., Haxthausen, A. E. & Lecomte, T. *Standardisation considerations for autonomous train control* in *International Symposium on Leveraging Applications of Formal Methods* (2022), 286–307.

60. Peng, K., Li, Q., He, Z., Zhang, B., Fu, X., Li, B., Wang, X., Cheng, Z.-Q., Yan, Y. & Peng, X. A survey of robotic manipulation: From bottom-up approaches to end-to-end paradigms with LLMs. *Neurocomputing,* 132921 (2026).

61. Rahimi, M., Liu, H., Cardenas, I. D., Starr, A., Hall, A. & Anderson, R. A review on technologies for localisation and navigation in autonomous railway maintenance systems. *Sensors* **22,** 4185 (2022).

62. Rahman, M., Liu, H., Masri, M., Durazo-Cardenas, I. & Starr, A. A railway track reconstruction method using robotic vision on a mobile manipulator: A proposed strategy. *Computers in Industry* **148,** 103900 (2023).

63. Ristić-Durrant, D., Franke, M. & Michels, K. A review of vision-based on-board obstacle detection and distance estimation in railways. *Sensors* **21,** 3452 (2021).

64. Saab, S. S. A map matching approach for train positioning. I. Development and analysis. *IEEE Transactions on vehicular technology* **49,** 467–475 (2000).

65. Salerno, M., E-Martín, Y., Fuentetaja, R., Gragera, A., Pozanco, A. & Borrajo, D. *Train route planning as a multi-agent path finding problem* in *Conference of the Spanish Association for Artificial Intelligence* (2021), 237–246.

66. Sha, Y., Zhu, S., Guo, H., Wang, Z. & Wang, H. *Towards autonomous indoor parking: A globally consistent semantic slam system and a semantic localization subsystem* in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (2025), 21151–21157.

67. Simson, S. A., Ferreira, L. & Murray, M. H. Rail Track Maintenance Planning: An Assessment Model. *Transportation Research Record* **1713,** 29–35 (2000).

68. SNCF Réseau. *ERTMS: High-performance, European-standard signalling* Updated August 6, 2025. Accessed: 2026-03-10. `https://www.sncf-reseau.com/en/tomorrows-network/ertms-high-performance-european-standard-signalling`.

69. Sobreira, H., Costa, C. M., Sousa, I., Rocha, L., Lima, J., Farias, P., Costa, P. & Moreira, A. P. Map-matching algorithms for robot self-localization: a comparison between perfect match, iterative closest point and normal distributions transform. *Journal of Intelligent & Robotic Systems* **93,** 533–546 (2019).

70. Sresakoolchai, J. & Kaewunruen, S. Railway infrastructure maintenance efficiency improvement using deep reinforcement learning integrated with digital twin based on track geometry and component defects. *Scientific reports* **13,** 2439 (2023).

71. Sturm, J., Engelhard, N., Endres, F., Burgard, W. & Cremers, D. *A benchmark for the evaluation of RGB-D SLAM systems* in *2012 IEEE/RSJ international conference on intelligent robots and systems* (2012), 573–580.
72. Stypułkowski, K., Gołda, P., Lewczuk, K. & Tomaszewska, J. Monitoring system for railway infrastructure elements based on thermal imaging analysis. *Sensors* **21,** 3819 (2021).
73. Sun, Y., Cao, C. & Wu, C. Multi-objective optimization of train routing problem combined with train scheduling on a high-speed railway network. *Transportation Research Part C: Emerging Technologies* **44,** 1–20 (2014).
74. Suomalainen, M., Karayiannidis, Y. & Kyrki, V. A survey of robot manipulation in contact. *Robotics and Autonomous Systems* **156,** 104224 (2022).
75. Umeyama, S. Least-squares estimation of transformation parameters between two point patterns. *IEEE Transactions on pattern analysis and machine intelligence* **13,** 376–380 (2002).
76. Vizzo, I., Guadagnino, T., Mersch, B., Wiesmann, L., Behley, J. & Stachniss, C. Kiss-icp: In defense of point-to-point icp–simple, accurate, and robust registration if done the right way. *IEEE Robotics and Automation Letters* **8,** 1029–1036 (2023).
77. Wandelt, S., Wang, Z. & Sun, X. Worldwide railway skeleton network: Extraction methodology and preliminary analysis. *IEEE Transactions on Intelligent Transportation Systems* **18,** 2206–2216 (2016).
78. Wang, S., Zheng, H., Tang, L., Li, Z., Zhao, R., Lu, Y. & Aw, K. C. Vibration-based and computer vision-aided nondestructive health condition evaluation of rail track structures. *Journal of Civil Structural Health Monitoring* **13,** 1–14 (2023).
79. Wen, W., Pfeifer, T., Bai, X. & Hsu, L.-t. Factor graph optimization for GNSS/INS integration: A comparison with the extended Kalman filter. *Navigation* **68,** 315–331 (2021).
80. Winter, H., Luthardt, S., Willert, V. & Adamy, J. *Generating compact geometric track-maps for train positioning applications* in *2019 IEEE Intelligent Vehicles Symposium (IV)* (2019), 1027–1032.
81. Xu, W., Cai, Y., He, D., Lin, J. & Zhang, F. Fast-lio2: Fast direct lidar-inertial odometry. *IEEE Transactions on Robotics* **38,** 2053–2073 (2022).

82. Yeong, D. J., Velasco-Hernandez, G., Barry, J. & Walsh, J. Sensor and sensor fusion technology in autonomous vehicles: A review. *Sensors* **21,** 2140 (2021).

83. Zhang, D., Zhang, J., Sun, Y., Li, T., Yin, H., Xie, H. & Yin, J. *Towards robust sensor-fusion ground SLAM: A comprehensive benchmark and a resilient framework* in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (2025), 8894–8901.

84. Zhang, Z. & Scaramuzza, D. *A tutorial on quantitative trajectory evaluation for visual (-inertial) odometry* in *2018 IEEE/RSJ international conference on intelligent robots and systems (IROS)* (2018), 7244–7251.

85. Zhen, W., Hu, Y., Yu, H. & Scherer, S. *LiDAR-enhanced structure-from-motion* in *2020 IEEE International Conference on Robotics and Automation (ICRA)* (2020), 6773–6779.