Proceedings of the Seventh International Conference on
Railway Technology: Research, Development and Maintenance
Edited by: J. Pombo
Civil-Comp Conferences, Volume 15, Paper 8.9
Civil-Comp Press, Edinburgh, United Kingdom, 2026
ISSN: 2753-3239, doi: 10.4203/ccc.15.8.9

Keeping the Driver in the Loop: How Artificial Intelligence Can Support Railway Safety Without Replacing Human Judgement

D. V. Katkoria and Y. Idmalek

Logiicdev, Graz, Austria

Abstract

Railway traffic is rising while the systems that keep trains apart and stop them in time rest on a safety base whose core logic was largely settled decades ago. Over the same period road transport has absorbed learning-based perception and driver assistance as a matter of course, yet the railway has been far slower to bring such methods into its safety-critical functions. This paper argues that the right response is neither to resist artificial intelligence nor to import it wholesale, but to adopt it wherever it strengthens human judgement and to avoid adopting it in ways that replace that judgement. We review where artificial intelligence already pays its way, chiefly in maintenance and inspection, and contrast this with the much harder problem of real-time safety and perception, where the margin for error vanishes. We then examine why current standards cannot yet certify learning-based components for safety-critical use, and why this favours a human-in-the-loop design rather than autonomy. An edge perception platform is presented as one practical embodiment of this principle, pairing a verifiable networking substrate with a learned layer that informs rather than overrules the driver. We conclude with the assurance steps still required.

Keywords: railways, artificial intelligence, physical AI, predictive maintenance, condition monitoring, obstacle detection, human oversight, functional safety, edge computing

1 Introduction

Watch how transport has changed over the last two decades and you notice an odd asymmetry. The road vehicle has been quietly transformed by machine learning. Lane-keeping, pedestrian detection, automatic emergency braking, the whole apparatus of driver assistance now ships in ordinary cars, and an entire regulatory and research culture has grown up around vehicles that perceive and decide for themselves. The railway, which is in most respects a far more controlled and better-instrumented environment, has not followed at anything like the same pace, at least not in the functions that matter most for safety. Its core protection systems remain deterministic, its safety-critical perception remains largely human, and learning-based artificial intelligence (AI) sits mostly at the edges of the operation rather than its heart.

There are good reasons for that caution, and this paper does not treat the railway's conservatism as a failing to be corrected. But the caution is now in tension with two pressures that are not going away. Traffic is rising, sharply on the passenger side, while the safety base carrying it was largely designed for a slower era. And the tools that could help, mature, inexpensive, proven elsewhere, are sitting right there. The temptation is to import the automotive approach wholesale and let the machine take over. The argument of this paper is that this would be a mistake, and that the more defensible path is also the more useful one.

Our position is simple to state. Artificial intelligence should be adopted in railway maintenance and safety wherever it strengthens human judgement, and it should not be adopted in ways that quietly replace it. Keeping a competent human in the decision loop is not a temporary scaffold to be removed once the technology matures; for a safety-critical, heavily regulated system carrying the public, it is the responsible design. The sections that follow build this case in stages. We look first at the ageing safety base and the rising demand pressing on it, then at the places where AI already pays its way in maintenance and inspection, then at the harder frontier of real-time safety where the same tools become much riskier. We then examine, critically, why current AI standards are not yet able to certify learning-based components for safety-critical rail, and why that pushes towards a human-in-the-loop design rather than away from it. Finally, we introduce our own platform, KiiGen™ by Logiicdev, as one concrete attempt to put that principle into hardware, and we are candid about both what it can and cannot yet claim.

2 A safety base built for a slower era

Train detection by track circuit goes back to William Robinson's closed rail circuit of 1872, as shown in Figure 1, a piece of engineering so durable that its basic principle still sits underneath a great deal of modern signalling [1,2]. Mechanical and then relay interlockings, which lock points and signals into safe routes so that a signaller cannot set two conflicting paths at once, were refined across the first half of the twentieth century. Colour-light signals replaced semaphores between roughly the 1920s and the

1960s. Automatic train protection, the layer that actually intervenes when a driver overspeeds or runs past a danger signal, only became widespread in the 1980s and 1990s, and the first genuinely standardised digital version, the European Train Control System (ETCS) within the European Rail Traffic Management System (ERTMS), was conceived in the 1990s and is still being rolled out across the network today [2, 3].

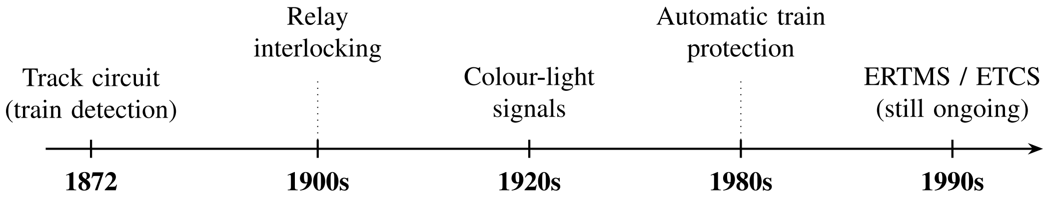

Figure 1: A simplified timeline of railway protection technology. More than a century separates the first track circuit from the digital train control still being deployed today.

None of that is a criticism, or at least it should not be read as one too quickly. There is a good reason the railway moves slowly here. A signalling function is expected to meet a tolerable hazard rate on the order of 10^{-9} dangerous failures per hour, and a relay either picks up or it does not [4]. The behaviour is deterministic, inspectable, and provable in a way that engineers and regulators have spent more than a century learning to trust. When the cost of a wrong answer is a derailment, conservatism is not laziness; it is the rational response to the stakes. So the long service life of these technologies reflects, in part, the fact that they work.

The control-command and signalling (CCS) subsystem, where train protection actually lives, is under active revision: the technical specification for interoperability for CCS, Commission Implementing Regulation (EU) 2023/1695, sits within the framework of the Interoperability Directive (EU) 2016/797 and is steadily folding in automated train operation and a new radio layer [5, 6]. There is a deeper point hiding in that revision. Automatic train protection delivered through ETCS was designed to stop a train that has already gone too far, not to perceive what lies ahead. The demand now taking shape is for something faster and more anticipatory, closer to collision avoidance and continuous monitoring than to after-the-fact intervention, and that is a job the deterministic base was never meant to do on its own.

The difficulty is that the traffic running over this base keeps growing while the base itself changes only gradually. In 2024 rail passenger transport in the European Union reached 443 billion passenger-kilometres, the highest figure since data collection began in 2004, up 5.8% on 2023 and already 10.7% above the pre-pandemic peak of 2019 [7]. Some national figures are stranger still: Hungary, host to this conference, recorded a 60% jump in passengers carried in a single year after a tariff reform [7]. Freight has not followed the same upward line, drifting slightly down to 375 billion tonne-kilometres in 2024, but the broader political ambition across Europe is to move a great deal more freight onto rail for decarbonisation reasons, not less [8]. More trains, running closer together, often faster, on infrastructure whose safety logic was largely settled decades ago. That is the squeeze.

Set this against what has happened on the road in the same period, and the contrast is hard to ignore. A mid-range car bought today ships with learned perception as a matter of course: cameras and radar that read lane markings, recognise pedestrians, and brake without the driver touching the pedal. An entire research and regulatory culture has grown up around machines that perceive and decide [9]. The railway, by comparison, has held its safety-critical perception almost entirely in human hands. Drivers still look out of the window. Lookouts still stand at the lineside. This is not because the rail sector is unaware of the technology, but because dropping a probabilistic, learning-based component into a system certified on deterministic guarantees is genuinely hard, and nobody has wanted to be the first to get it wrong.

Therein lies the tension this paper is concerned with. The pressure to do more, and to do it more safely, is real and rising. The tools that have transformed perception elsewhere are mature and cheap. Yet the path by which they enter the railway cannot simply copy the automotive playbook, because the railway's relationship with proof is different and, arguably, should stay different. The question is not whether artificial intelligence belongs in rail safety. It plainly does, and as the next sections show, it already earns its keep in several places.

3 Where AI pays off, and where it must: monitoring, inspection and safety

The biggest prize for artificial intelligence in rail is safety: a system that watches the line ahead in real time, for collision avoidance and obstacle detection, where the stakes are measured in lives rather than maintenance hours. Yet that is exactly where the technology has made the least headway. A recent review of the field makes the point bluntly: the great majority of AI research in rail has gone into maintenance and inspection, well ahead of traffic management, and far ahead of the safety and autonomous-driving applications that attract the headlines [10]. The real-time safety task has been comparatively ignored, and we return to why in the next section. For now the honest picture is that AI has earned its keep in rail mostly in the less glamorous, more forgiving work of maintenance and condition monitoring. The reason it got there first is not hard to see. Maintenance is data-rich and, crucially, it is forgiving in a way that real-time safety is not. A predictive maintenance model that gets it wrong usually still leaves a human inspection step before anything fails. The blast radius of a mistake is small.

The clearest published example comes from Deutsche Bahn's work within the Europe's Rail FP3-IAM4Rail programme, the ARGO project, which set out to partially automate the visual inspection of trains [9]. Today that inspection is manual, slow, and puts staff in awkward positions under the vehicle. ARGO instead uses a camera-equipped robot and machine learning to score the imagery. Because defects are rare and the data badly imbalanced, the project leans on anomaly detection, learning what 'normal' looks like and flagging departures from it [9]. A high anomaly score is not the machine declaring a fault; it is the machine saying this one is worth a human's attention.

The AI narrows the search. The person decides. Table 1 sketches a few such maintenance tasks and, deliberately, names the step that stays human in each. What stands out, looking at the list, is what is absent from it. Every entry watches the asset; none watches the line ahead. The same sensing and inference that has matured so quickly for maintenance has barely been turned towards the safety-critical task of seeing a hazard in time to act, and closing that gap is going to take new technology rather than a rearrangement of the old.

Maintenance task	**Typical AI method**	**What the human still decides**
Underfloor and shell inspection	Anomaly detection (autoencoders) on captured imagery	Whether a flagged anomaly is a real defect and what to do about it
Rail surface and wheel defects	Image classification, signal processing	Severity grading and repair scheduling

Table 1: A few maintenance tasks where AI is already useful, and the decision that remains with a qualified person in each.

4 The harder frontier: real-time safety and perception

Real-time safety is a different problem from maintenance, and a far less forgiving one. The data is no longer a tidy archive of inspection images; it is a live, messy stream from a vehicle travelling at speed, in weather, at night. The decision is no longer "schedule this for a closer look" but "brake, now, or do not". Here the margin for error that made maintenance such a comfortable place for AI is simply gone. This, we suspect, is why the research has lagged here. The same survey that found maintenance dominating rail-AI work also noted that safety, security and autonomous-driving applications have received comparatively little attention, held back by the heavier burden of testing, validation and regulatory approval that anything safety-critical must carry [10].

The functions themselves are easy enough to name. Obstacle and intrusion detection on the track ahead. Risk assessment at level crossings, which remain one of the more stubborn sources of fatalities on the network. Perception in conditions where the human eye and conventional cameras struggle, which is to say darkness, fog, driving rain, low winter sun. These are not exotic problems. A car already does versions of them. The catch is that rail asks for this at ranges a car never needs. Even at the moderate speeds of inter-city, urban and freight services, a train carries far more momentum than a car and cannot stop on the same terms, so a perception system earns its place only if it sees far enough ahead to give the driver real time to act, and does so when there is no ambient light to work with and no second chance if it misreads the scene. High-speed lines are a different matter: there the speeds and stopping distances outrun what any forward sensor at a few hundred metres could meaningfully change, and they are not the target here.

It helps to remember what is at stake. In 2024, 750 people were killed in railway accidents in the European Union, and the great majority were struck while on or crossing the track: unauthorised people on the line and level-crossing users together accounted for over nine in ten deaths, passengers for a small minority [11]. That is precisely the hazard a forward-looking sensor is meant to catch, and Figure 2 shows the breakdown. Animals add a quieter toll: Swedish railways alone record on the order of five thousand collisions with ungulates such as moose each year, and per kilometre the railway can kill more of them than the road does [12]. None of this says a sensor would have prevented any particular death; it says the problem is large, real, and concentrated where perception at a distance, in the dark, would help.

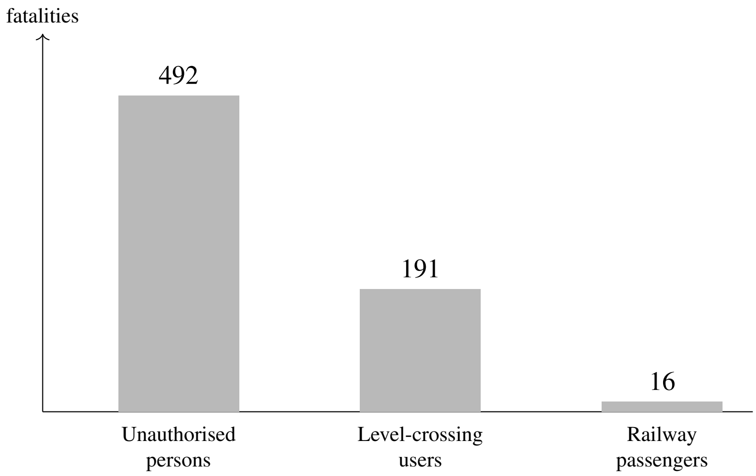

Figure 2: People killed in railway accidents in the European Union in 2024, by category of victim, out of 750 in total. The two largest groups are people on the track and at level crossings. Source: Eurostat.

The conventional answers leave gaps. Lineside cameras see little after dark. Human lookouts, where they are still used, tire and lose concentration, and there is a real literature within the rail community on just how quickly a driver's capacity to perceive and process information degrades under load [9]. So the appeal of an always-on, weather-agnostic sensor that can flag a hazard before the driver could possibly see it is obvious. The trouble is equally obvious once you sit with it. An obstacle-detection system that misses a real hazard at line speed fails in the worst possible way. One that cries wolf too often, triggering unnecessary braking, is not merely annoying; it teaches drivers to distrust and eventually ignore it, and it gums up a network that is already running near capacity. Both failure modes are unacceptable for different reasons, and a learning-based perception system, by its nature, will produce some of each.

That is the frontier. It is where AI could do the most good in rail safety and also where it is hardest to trust, and those two facts are not a coincidence. Before looking at one concrete attempt to work in this space, it is worth being precise about why the trust is so hard to earn, because the answer is not really about the cleverness of the algorithm. It is about whether the rules we use to certify safety can say anything meaningful about a component that learns.

5 A critical look at AI standards and the case for the human in the loop

Railway safety does not run on good intentions; it runs on a stack of standards that tell engineers how to prove a system is safe before it carries anyone. Any significant change has to pass through the Common Safety Method for risk assessment set out in Implementing Regulation (EU) No 402/2013, which forces the proposer to identify hazards, classify them, and demonstrate that the residual risk is acceptable [13]. Underneath that sits the EN 5012x family for reliability, availability, maintainability and safety, and, for software specifically, EN 50716, which in 2024 consolidated the older EN 50128 and EN 50657 [14]. The whole edifice is built around a safety integrity level, a way of saying how confident we need to be that a given function will not fail dangerously [4].

Here is the problem, and it is not a small one. That framework assumes you can analyse a function's failure behaviour. For a relay, or for conventionally written software, you largely can; the logic is fixed and inspectable. A trained neural network is a different kind of object. Its behaviour is statistical, shaped by data rather than written by hand, and it can fail in ways that are hard to predict and harder still to enumerate exhaustively. Assigning a safety integrity level to a convolutional network that classifies what is on the track is not a solved problem; it may not even be the right question. Deutsche Bahn's own engineers, working on the relatively gentle case of automated visual inspection, found that the current rulebooks were insufficient to cover the AI-related risks they had logged, that no comparable reference system existed to lean on, and that it was genuinely unsettled whether their system should even count as "high-risk" under the new Artificial Intelligence Act [9, 15]. If that is the state of play for a robot inspecting a parked train, the bar for a system that brakes a moving one is considerably higher.

The research community has started to name what is missing. Work on AI trustworthiness for railways points out that the sector has always prized predictable, tractable, verifiable behaviour, and that the black-box character of modern AI, its complexity and its autonomy, cut directly against that grain [10]. The response taking shape across Europe is not to lower the bar but to widen it, extending the familiar risk vocabulary to cover attributes that traditional software never had to worry about: transparency, explainability, robustness, accountability. It is worth noticing that when the European High-Level Expert Group set out its requirements for trustworthy AI, the very first one was human agency and oversight [16]. That ordering is not an accident.

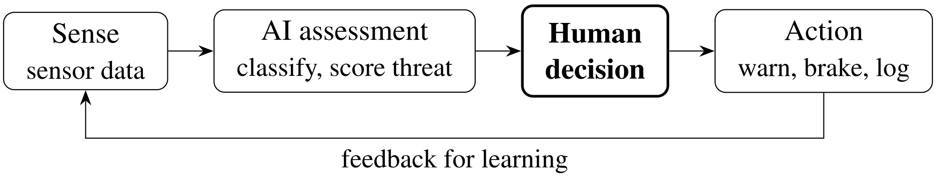

Figure 3: The human-in-the-loop arrangement. The AI narrows and scores; the qualified person remains the decision gate before any safety action; outcomes feed back to improve the model offline.

So what does keeping a human in the loop actually buy, beyond reassurance? Three things. It guards against automation bias, the well-documented tendency of operators to defer to a confident-looking machine even when it is wrong; a driver who is the decision-maker, not just a monitor, stays engaged. It absorbs the false positives, because a person can dismiss an obvious nuisance alert that a fully automatic system would have acted on, which keeps the network moving and keeps trust intact. And it handles the long tail, the rare and weird situations that fall outside the conditions the model was trained for, the so-called operational design domain, where machine confidence is least deserved [9]. Figure 3 sketches the arrangement: the model does the seeing and the scoring, the person does the deciding, and the outcomes flow back to improve the model in slow time rather than in the moment.

There is a principled version of this too. The GAME principle, globally at least equivalent, drawn from the EN 5012x tradition, holds that any change to the system must be at least as safe as what it replaces [4]. Deutsche Bahn operationalised exactly this through a shadow experiment, running the automated inspector alongside a human worker on the same task to compare their error rates before trusting the machine [9]. The logic generalises cleanly to perception on a moving train. You do not remove the human until you can demonstrate, on the operator's own terms and in their own conditions, that the combined system is safer than the human alone. Until then, the responsible design is not AI instead of the driver, but AI in the service of the driver. That is not a hedge or a transitional compromise to be engineered away later. For a safety-critical, heavily regulated system carrying the public, it is arguably the correct end state.

This also reframes how a vendor should present a perception platform. The interesting claim is not "our system can drive the train". It is "our system can be verified where it must be, and it defers to a human where it cannot yet be". That distinction, between the parts of a platform that admit proof and the parts that do not, turns out to be a useful lens for looking at an actual product, which is where we turn next.

6 Introducing Logiicdev and KiiGen™

Logiicdev is a company based in Graz, Austria, working on edge artificial intelligence and industrial connectivity. Its KiiGen™ platform is the example we want to hold up against the argument of the previous section, partly because it was built around exactly the split we just described, and partly, in fairness, because the present authors can speak to it directly. The name comes from a word meaning "origin" or "new birth", which captures the intent: not to rip out and replace legacy equipment, but to give existing assets a second life by adding intelligence at the edge.

At its core, KiiGen™ is a deterministic, retrofittable bridge with an embedded AI engine. On the connectivity side it integrates AI on the industrial protocols that rarely speak to each other on their own, among them the Robot Operating System (ROS and ROS 2), OPC UA, Modbus and EtherNet/IP with time-sensitive net-working (TSN) so that the timing is predictable rather than best-effort, and it does so over power over ethernet (PoE). This is the part of the platform that can be held to a high standard of proof, and it has been: the protocol-translation and runtime-verification approach underlying the bridge was published and peer-reviewed at ETFA 2025, with correctness checked against a formal interaction model [17]. That matters for the present discussion. It is a concrete instance of the deterministic, verifiable substrate that Section 5 argued for. The networking layer is the kind of thing a safety case can actually get its arms around.

On top of that substrate sits the learned part, an edge AI engine that runs perception and analytics locally rather than in the cloud. Logiicdev frames this layer as physical AI: FPGA-based perception that lets a machine sense, interpret and act on the physical world in real time at the edge, with no cloud round-trip. Its stated live use case for rail is precisely that, three-dimensional perception for track safety, simultaneously detecting, classifying and tracking obstacles, whether people, animals or other hazards, from fused laser-based sensing, or lidar [18]. In practice this means reading the track ahead in conditions where ordinary cameras fail, including total darkness and poor weather. The intended response pipeline mirrors Figure 3 almost exactly: the system senses and classifies the obstacle, scores the threat, and then escalates through graded responses, an audible warning, high-power lights, and only as a last resort automatic emergency braking, while logging every event for later review [18]. By design, the driver is warned first and the machine takes irreversible action only at the extreme. The human is built into the loop rather than bolted on. Figure 4 shows the graded sequence: the alerts to the cab come first, and automatic braking sits at the end as a last resort, not a first response.

Two practical details matter for a train fit. The unit draws its power over the network cable itself, using IEEE 802.3at (PoE+, Type 2), which keeps a retrofit's wiring to a minimum. And because anything mounted on rolling stock has to survive shock, vibration and a temperature range that office hardware never meets, the equipment has to satisfy the mandatory EN 50155 standard for electronic equipment on rolling stock before it can promise uninterrupted service in that environment [19]. Figure 5 sets out

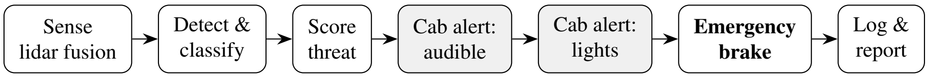

Figure 4: The KiiGen™ railway-safety response pipeline. The driver is alerted first, audibly then visually; automatic emergency braking is the final, last-resort step, and every event is logged for later review.

the same idea at the level of architecture, and it is worth dwelling on the split it shows, because that split is the whole point.

It is worth being concrete about how a box this small manages the computation at all, because that is much of why KiiGen™ belongs in this discussion rather than a rack of servers. A lidar throws off a dense three-dimensional point cloud, hundreds of thousands of points every tenth of a second, and turning that torrent into “a person at two hundred metres” quickly enough to matter is not a light task. Done the obvious way, on a general-purpose processor, it either misses the latency budget or demands a power-hungry graphics card that no cab will tolerate. The platform instead pushes the heavy, repetitive stages of the pipeline onto the FPGA fabric of its accelerator. Point-cloud filtering, ground-plane removal, voxel downsampling and clustering are fixed and highly parallel operations, and an FPGA can stream them as the points arrive rather than buffering whole frames, at a latency that is bounded and predictable instead of being subject to the jitter of an operating system. The neural network that classifies what remains is quantised to integer arithmetic and mapped onto the same fabric, so that its weights and intermediate results sit largely in on-chip memory rather than making constant journeys to external DRAM or cloud. Logiicdev packages this as a set of ultra-lean, low-memory FPGA IP modules, which is the supply-side version of the same small-footprint argument.

That is what keeps both the memory and the power small. Streaming dataflow and on-chip storage mean the device needs neither gigabytes of fast external memory nor the bandwidth that feeds a desktop card, which is how a perception engine of this kind fits inside roughly a ten-watt envelope, without a fan, in an enclosure that can still meet the EN 50155 temperature classes noted above; the accelerator’s throughput is a vendor-stated, configuration-dependent figure [18]. And because the whole chain, from raw points to threat score, runs on the device, nothing depends on a network link or a remote service. A tunnel, a remote stretch of track, a cutting with no signal: none of it changes the response time, because nothing leaves the box. For a safety function that has to answer in well under a second, every time, that locality is a requirement rather than a convenience. The reconfigurability of the fabric is a quieter advantage, since the same hardware can take an improved model after a field trial without a hardware swap, which matters for a learned component that, as Section 5 argued, has to keep earning its trust.

The vendor-reported figures for the rail configuration are collected in Table 2. They are manufacturer claims rather than independently certified results, but the headline numbers at least hang together: at the roughly 400 m detection range, and at the

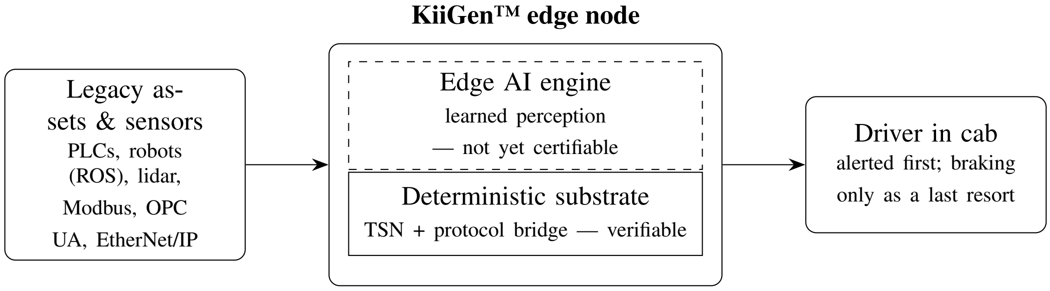

Figure 5: The KiiGen™ architecture, simplified. A deterministic networking substrate, the part that admits formal proof and has been peer-reviewed, carries a learned perception layer that does not yet admit the same proof. The output informs the driver rather than acting over their head.

moderate speeds of inter-city, urban and freight services, between 50 and 80 km/h, the warning window works out to the eighteen to twenty-nine seconds shown, time enough for a driver to react. The platform is aimed squarely at these slower services rather than high-speed lines, where the speeds and stopping distances would outrun what a sensor at this range could usefully change.

Parameter	**Vendor-reported value**
Detection range, person	up to about 400 m
Detection range, small object	greater than 100 m
Field of view (configurable)	360 degrees (fused point cloud)
Detection-to-alert latency	under 20 ms
Operating illumination	0 lux to full daylight
Warning time at 50–80 km/h	about 18–29 s
Power	about 10 W over PoE

Table 2: Vendor-reported performance figures for the KiiGen™ rail configuration. These are manufacturer claims for the intended operating envelope, not independently certified results.

A rail deployment is not only about detection. The cab is a hostile place for electronics, and a device earns its place there only by meeting a stack of environmental standards. Table 3 sets out the main ones, a reminder that a credible rail product is defined as much by these unglamorous requirements as by its detection range.

Two honest caveats belong here, and leaving them out would weaken rather than strengthen the case. First, KiiGen™ is fundamentally a general-purpose industrial edge platform; its proven deployments to date are in manufacturing and similar settings, and the railway obstacle-detection application is exactly that, an application,

Requirement	Governing standard	What it demands
Temperature extremes	EN 50155 [19]	Operation across wide cab and roof temperature classes
Environmental tolerance	EN 50125-1	Humidity, altitude derating, dust and water ingress
Shock and vibration	EN 61373	Random vibration and shock by mounting category
Electromagnetic compatibility	EN 50121-3-2	Bounded emissions and immunity in the rail environment
Fire safety	EN 45545-2	Fire behaviour of materials and components

Table 3: Rail environmental requirements an on-board device must meet, the standard that governs each, and what each entails. Meeting them is necessary before any detection claim becomes relevant in service.

not a fielded and independently validated rail product. Second, the figures in Table 2 are vendor-reported rather than independently certified. They are plausible targets, and the underlying sensing technology supports the general shape of the claim, but they are precisely the sort of numbers that the assurance machinery of Section 5 exists to scrutinise. A shadow experiment comparing the system against human detection, a carefully bounded operational design domain, an explicit hazard analysis under the Common Safety Method: these are the steps that would turn a promising platform into a trustworthy one. Presenting KiiGen™ as decision support that keeps the driver in command is not only the safer engineering posture, it is also the more credible claim to make in front of a regulator.

7 Conclusions

The railway is being asked to do more, on a safety base largely settled decades ago, at a moment when the perception and prediction tools that reshaped road transport are cheap and mature. Artificial intelligence has a place in the response. In maintenance and inspection it already earns it, triaging floods of sensor and image data so that scarce human expertise lands where it is needed.

The harder lesson concerns the safety-critical frontier. The same techniques become dangerous when asked to decide, in real time, whether to brake a moving train, and the reason is not a shortage of cleverness: the standards we use to certify railway safety were built for deterministic systems and cannot yet say anything rigorous about a component that learns. Until they can, removing the human is not progress but the quiet transfer of a decision to a system that cannot be held to the same proof. Keep the person deciding, and let the machine make them faster, sharper and better informed. Edge physical AI platforms such as KiiGen™ show one practical shape this can take, pairing a verifiable networking substrate with a learned perception layer that informs the driver rather than overruling them. What remains is demanding but not mysterious: standards that can certify learned perception, bounded operational design domains, and honest field validation against human performance. A short demonstration of the platform is available at `https://www.youtube.com/watch?v=jS_OTiTzfbE`.

Acknowledgements

The authors thank colleagues at Logiicdev GmbH for technical discussions, and gratefully acknowledge the research community whose published work on railway safety assurance and AI trustworthiness informed this paper.

References

[1] G. Theeg, S. Vlasenko (Editors), "Railway Signalling and Interlocking: International Compendium", 3rd Edition, Edition Eurailpress, Hamburg, Germany, 2019.

[2] J. Pachl, "Railway Signalling Principles", Edition 2.0, Universitatsbibliothek Braunschweig, Braunschweig, Germany, 2021, DOI:10.24355/dbbs.084-202110181429-0.

[3] European Union Agency for Railways, "The European Train Control System (ETCS): Overview", https://www.era.europa.eu, 2024.

[4] European Committee for Electrotechnical Standardization, "EN 50126-1 – Railway Applications: The Specification and Demonstration of Reliability, Availability, Maintainability and Safety (RAMS) – Part 1: Generic RAMS Process", Brussels, Belgium, 2017.

[5] European Commission, "Commission Implementing Regulation (EU) 2023/1695 on the Technical Specification for Interoperability Relating to the Control-Command and Signalling Subsystems of the Rail System in the European Union", Brussels, Belgium, 2023.

[6] European Parliament and Council, "Directive (EU) 2016/797 on the Interoperability of the Rail System within the European Union", Brussels, Belgium, 2016.

[7] Eurostat, "Rail passenger transport increased by 5.8% in 2024", https://ec.europa.eu/eurostat, 2025.

[8] Eurostat, "Railway freight transport statistics", https://ec.europa.eu/eurostat, 2025.

[9] D. Walter, R. Hegemann, “Safety Considerations for Use of Artificial Intelligence in Train Maintenance: Partial Automation of Visual Inspection Tasks”, FP3-IAM4Rail Whitepaper, Europe’s Rail Joint Undertaking, 2024.

[10] A. Ladj, A. Boussif, “Towards AI Trustworthiness Assessment Framework for Railway Applications”, in Proceedings of the 35th European Safety and Reliability Conference (ESREL SRA-E 2025), Research Publishing, Singapore, 1081-1088, 2025, DOI:10.3850/978-981-94-3281-3_ESREL-SRA-E2025-P7834-cd.

[11] Eurostat, “Railway Safety Statistics in the EU”, Statistics Explained, https://ec.europa.eu/eurostat, 2025.

[12] L. Borda-de-Agua, R. Barrientos, P. Beja, H.M. Pereira (Editors), “Railway Ecology”, Springer, Cham, Switzerland, 2017.

[13] European Commission, “Commission Implementing Regulation (EU) No 402/2013 on the Common Safety Method for Risk Evaluation and Assessment”, Brussels, Belgium, 2013.

[14] European Committee for Electrotechnical Standardization, “EN 50716 – Railway Applications: Requirements for Software Development”, Brussels, Belgium, 2023.

[15] European Parliament and Council, “Regulation (EU) 2024/1689 Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act)”, Brussels, Belgium, 2024.

[16] High-Level Expert Group on Artificial Intelligence, “Ethics Guidelines for Trustworthy AI”, European Commission, Brussels, Belgium, 2019.

[17] Q.-D. Nguyen, D. Rammal, C. Gaston, D.V. Katkoria, A. Lapitre, et al., “Towards Bridging Industrial Ethernet Networks: Protocol Translation and Runtime Verification”, in ETFA 2025 – IEEE 30th International Conference on Emerging Technologies and Factory Automation, 2025, DOI:10.1109/ETFA65518.2025.11205755.

[18] Logiicdev GmbH, “KiiGen™: Industrial Edge AI Platform – Datasheet and Use Cases”, Graz, Austria, 2025.

[19] European Committee for Electrotechnical Standardization, “EN 50155:2021 – Railway Applications: Rolling Stock – Electronic Equipment”, Brussels, Belgium, 2021.