

Safe and Secure Future AI-Driven Railway Technologies: Challenges for Formal Methods in Railway

Monika Seisenberger¹, Maurice H. ter Beek², Alessio Ferrari², Anne Haxthausen³,
Phillip James¹, Andrew Lawrence⁴, Bas Luttik⁵, Jaco van de Pol⁶, and Simon Wimmer⁶

¹ Swansea University, UK

² ISTI-CNR, Pisa, Italy

³ Technical University of Denmark, Denmark

⁴ Siemens Mobility Chippenham, UK

⁵ Eindhoven University of Technology, The Netherlands

⁶ Aarhus University, Denmark

Abstract. In 2020, the EU launched its sustainable and smart mobility strategy, outlining how it plans to have a 90% reduction in transport emission by 2050. Central to achieving this goal will be the improvement of rail technology, with many new visionary data-driven systems being proposed. AI will be the enabling technology for many of those systems. However, safety and security guarantees will be key for wide-spread acceptance and uptake by Industry and Society. Therefore, suitable verification and validation techniques are needed.

In this article, we argue how formal methods research can contribute to achieving these targets and present several research problems and techniques worth to be further developed.

1 Introduction

In 2020, the EU launched its sustainable and smart mobility strategy, outlining how it plans to have a 90% reduction in transport emission by 2050. This will be key to achieving the European Green deal of becoming carbon neutral by 2050⁷. Central for this reduction will be the improvement of rail technology, as rail is one of the greenest modes of transportation. To address this ambition and support the interoperability and efficiency in the rail domain, new visionary systems based on interdisciplinary approaches in Engineering and Computer Science are being proposed, for example, innovative signalling systems (including moving block technology), smart monitoring and maintenance, optimal scheduling, and automated driving. Underpinning these new systems, data is recognized as a highly valuable asset. For instance, the UK’s Rail Technical Strategy (10/2020) states: ”Data will have fit for purpose governance, access arrangements, systems and technical skills. These building blocks underpin the progression of all the other functional priorities which each have their own specific data requirements and opportunities.”⁸ AI will be the enabling technology for such data driven systems. However, safety and security guarantees will be key for wide-spread acceptance and uptake by industry and society. Therefore, suitable verification and validation techniques are needed.

In the Railway Industry and in verification and validation research, there is a skills gap in terms of knowledge of AI principles, techniques, and practices. The Railway Industry is normally focused on developing software and systems with fully predictable, explainable and verifiable behaviour, while AI techniques are by nature adaptive, and open to different scenarios. More importantly, AI-based systems have explainability problems that collide with the idea of fully controllable and verifiable system behaviour. Therefore, the Railway Industry needs to exploit AI systems to deliver smart and green transport, whilst at the same time maintaining the highest standards in terms of safety and certification.

In the following, we propose several research questions and challenges to the formal methods community, to aid the development as well as the certification of safe and secure next generation rail systems, many of which will make use of AI techniques. Currently, ERTMS/ETCS level 3⁹ is in the process to replace the traditional discrete train separation mechanism (blocks protected by hardware installed along the tracks) by a continuous mechanism (moving block), and use radio communication for the exchange of position information between trains and track-side. Also, Automatic Train

⁷ https://ec.europa.eu/info/strategy/priorities-2019-2024/european-green-deal_en

⁸ <https://railtechnicalstrategy.co.uk/data-driven/>

⁹ www.ertms.net, www.era.europa.eu/activities/european-rail-traffic-management-system-ertms_en

Operation (ATO) will, at the higher grades of automation, eliminate the human driver and require a continuous software-controlled interaction, not only of a train with the track-side system, but also of trains between each other. Future signalling systems will be an order of magnitude more complex than they are today, and formal modelling and verification technology will be essential to cope with that complexity. The three challenges presented in this article will relate directly or indirectly to this situation.

The first challenge concerns certification and the associated verification technology. Automation of the verification is required to address the massive scale and complexity of railway systems. This entails intricate symbolic and parallel verification algorithms. A major problem for certification is that the implementation of the verification tools themselves could contain errors. Needed is a theory to equip verification tools with certificate generation, thus enabling the answers of automated verification tools to be checked independently. Furthermore, there is a need to investigate how AI based systems can be certified as this is not in general possible on the basis of today’s CENELEC standards [Eur17, Eur11, Eur18].

The second challenge concerns the European Rail Traffic Management System (ERTMS) level 3 which is anticipated to be the main railway system in 10 years’ time. By introducing true continuous (moving block) signalling, it allows for both capacity increases and lower energy consumption, with trains effectively and intelligently managed. However, it also introduces elements of a hybrid nature into sub-systems of ERTMS, with components working with both discrete- and continuous-valued data. Along with this, traditional modelling of both safety and security need to change to include this hybrid nature. ERTMS consists of many interconnected components that differ in nature. This provides modelling challenges in both terms of complete models of the system, but also in terms of understanding and modelling the safety and security requirements of the sub systems, both individually and as a whole. In addition, such signalling projects are often large and involve various aspects of backwards compatibility with older signalling systems (e.g., to deal with older rolling stock) which causes scalability issues for current techniques. Modelling and verifying hybrid systems of the size presented by ERTMS level 3 is an open challenge for Computer Science. Needed is a complete model of ERTMS level 3 which includes the safety and security requirements it must uphold. This will also serve as a reference architecture for ERTMS level 3 deployment projects and help with the faster roll-out.

The third challenge concerns the extension of formal methods to include the use of AI techniques. We will present in Section 4 how this can be achieved, and highlight some of the methods we plan utilise, as well as the problems that need to be overcome. Our focus on AI integration aligns well with the planned successor of Shift2Rail, the RAIL European partnership, which will specifically focus on digitalisation and automation. So far, there are a number of research projects and white papers (see e.g. [BDF+22, ADL21]) that look into the integration of AI into the Railway domain. A prominent project is “RAILS” [RAI] which provided a first overview on available AI techniques, as well as application areas and work done in the Railway domain. We complement this by focusing on the specific challenges for formal methods, from a methodological point of view and by means of specific case studies.

2 Certified Verification of Railway Designs

Railway systems are complex and safety-critical, imposing high safety standards and strict certification requirements [FtB22]. The need for extensive testing, verification, and certification imposes an unfortunate barrier to the quick adoption of innovative railway technologies, which are essential to provide interoperability between national systems and increase utilisation of the railways. Therefore, we propose to push automated verification techniques for their certification. The dynamic nature of moving blocks, communication between trains and infrastructure, and the need to anticipate malfunctioning hardware, lead to an explosion of interactions and case distinctions, which can hardly be managed manually. Consequently, to eliminate manual errors and provide the necessary scalability, the verification needs to be *automated*. Moreover, the verification process needs to interact with the overall railway engineering process, in particular certification by regulatory bodies. To cater for the needs of stakeholders in this process (e.g. non-verification engineers, regulators, auditors), verification needs to be *trustworthy* and *explainable* [GMS20, Win21].

In the following, we will discuss the state-of-the-art and open challenges of verification technology in hindsight to railway systems, and on the three axis of automation, trustworthiness, and explainability.

We will also address the challenge of which standards to use for certifying train control (sub-)systems, specifically when they are based on AI technology.

2.1 Automated Verification

Automated verification of railway signalling poses a scientific challenge for several reasons:

1. **Complexity:** As discussed above, the modern railway systems lead to highly complex designs (of software, hardware, systems). Here, we focus primarily on the area of model checking for the automated verification of such designs. Model checking is a purely automatic method, which decides if a given specification (model) satisfies a given requirement (property). Various model-checking algorithms, ranging from exhaustive (probabilistic) model checking [BK08] to statistical model checking [AP18], have been proposed. Model checking is based on discrete enumeration [CHVB18], while statistical model checking involves running a sufficient number of (probabilistic) simulations to obtain statistical evidence (with a predefined confidence level).

2. **Parametricity:** For an effective deployment of signalling systems, there is a need to verify standard components that can be combined and instantiated to particular situations (e.g., track layouts). However, parametric verification in general is a computationally undecidable problem [And19]. This means that some form of abstraction and manual intervention is necessary to achieve verification results of the required generality. Compositional verification techniques [MFH17] can be employed to combine verified standard components.

3. **Continuous reasoning:** Obviously, railway systems need to operate in real-time! Moreover, innovations like moving blocks require reasoning about continuous variables, such as braking curves. Stochastic behaviour (like failure of hardware components) leads to other forms of continuous reasoning. These continuous extensions remain challenging, and sometimes even cross the border of computational decidability [HKPV98]. Moreover, in case of AI-based systems, some components might be machine-learned models. These are of an inherently probabilistic nature and can be considered to be black-box components around which one may need to build a safety shield for runtime enforcement of guarantees [PKPB21].

Model checking is essentially a smart enumeration method of the possible system behaviour. Typically, the properties are checked on-the-fly, during the enumeration. The major threat to model-checking is the state space explosion, exhausting not only time resources, but also the memory of the computer. Statistical model checking scales better, since there is no need to explore the full state space and the required simulations can trivially be distributed and run in parallel, but contrary to model checking, exact results (with 100% confidence) cannot be achieved.

The field of model-checking algorithms has advanced tremendously over the last 25 years. Many improvements are based on symbolic reasoning algorithms, like abstraction, symmetry reduction, partial-order reduction, and the use of BDD, SAT and SMT solvers. Other improvements deploy high-performance computing, like clusters of machines, multi-core hardware, or even many-core GPUs to off-load intensive verification tasks [KLM⁺15]. While model checking is a very active field of research with many open challenges we refrain from addressing them here. Instead, we consider issues more specific to the railway domain in the following.

2.2 Trustworthy and Certifiable Verification

Verification can be seen as an instance of *trust reduction*: when employing verification in the engineering process, trust in the safety of the resulting system will to a large extent be deduced from trust in the correctness of the model. Therefore, on the one hand, the model needs to be a suitable representation of the system and its environment. On the other hand, the results of the verification process need to be trustworthy. We focus on the latter issue, which poses three fundamental challenges:

- Model-checking algorithms and implementations became so advanced and complicated that it is hard to guarantee that these verification tools themselves are free from bugs. Indeed, several cases of bugs in verification tools (and even verification theory [BLR05, Sie19]) have been reported in the literature. This is clearly a threat for the certification of safety-critical systems.
- Model-checkers often run in so-called “bug hunting” mode. This means that they are specialised in finding bugs. To save on computational resources (time, memory, energy), they use various search heuristics. These heuristics give up completeness of the search. Hence, not finding any bugs is no

guarantee of correctness. In this sense, bug hunting with model checking is similar to advanced, automated testing, which is also inherently incomplete (as is statistical model checking).

- Model checking provides an asymmetric method. In principle, the answer to a model-checking query is either yes or no (the model either satisfies the property or violates it). In the “no-case”, the model checker typically returns a counterexample, which can be inspected, tested on the model or even on the real system, and used to debug the model or the property. However, in the yes-case, no further evidence is provided. This is very unfortunate for certification, not only in light of the previous limitations, but also since no information for the safety-case is provided, besides the fact that no bug has been found.

To remedy the situation, one might proceed by formally verifying the model checker. This provides a rigorous mathematical proof, checked in an interactive theorem prover [RPS⁺19, Wie06] (like Coq [BC04] or Isabelle/HOL [NPW02]). This approach would ensure that the “yes-result” of the model-checker is trustworthy. However, although possible in principle, verification of an advanced model checker would be a major undertaking. There is some progress in the formal verification of model-checking algorithms [vdP15, OHJvdP20] and even code [ELN⁺14, BLFS15, BFLW18, BL18, WL18], but so far, this could only be applied to relatively simple algorithms and basic implementations. Consequently, the “verified model-checkers” cannot match the efficiency of high-performance model-checkers.

We propose *certified model checking*, which provides a sweet-spot here: In this approach, one uses high-performance, “unsafe” model checkers, but equips them with the potential to generate some form of certificates in the “yes-case”. One builds a separate, independent certifier, which checks the certificates. Since checking certificates is much simpler than finding certificates, the certifier is a relatively simple tool, which can be formally verified. If the formally verified certifier accepts the certificate generated by the model checker, then we achieve maximal confidence in the safety of the system-under-study. Hence, in certified model checking we combine maximal efficiency (high-performance model checkers) with maximal confidence (formally verified certifiers). The concept of certified model checking was conceived for the μ -calculus [Nam01]. A proof of concept has been provided for SMT-based model checking [KBT⁺16] (where certificates are basically invariants) and more recently for liveness checking of finite-state systems [GRT21] and timed automata [WHvdP20] (where certificates consist of reachability invariants and topological ranking functions). In the latter case, the certifier was fully verified in the theorem prover Isabelle/HOL.

Certified model-checking poses multiple challenges. The generated certificates should contain sufficient information to check the proof independently; yet be more concise than the full state space, and easier to check than to generate. Independent certificate checker need to be constructed, possibly based on an interactive theorem prover. The certificate checker should be simple and amenable for formal verification. This requires a (partial) formalisation of the meta-theory of the verification tools.

The certified model-checking method is in its infancy: It is an open question what certificates for complicated model-checking algorithms, like partial-order order or symmetry reduction, should look like. It is also unclear how to generate certificates from parallel implementations of on-the-fly model-checkers. The formal verification of the certification theory and the certifiers remains a challenge for these novel applications.

Moreover, a *lack of standardization* in modelling formalisms makes it hard to define general-purpose model checking certificates and to provide checkers for them. This is also particular concern with hindsight to verification of the certificate checkers, as a large part of the laborious formalisation process would need to be repeated for each modelling language. Naturally, a lack of standardization is equally challenging for collaboration with regulatory bodies.

Finally, the certified verification approach needs to be extended to AI-intensive systems. In particular, classifiers and schedules generated by game-based AI algorithms should be equipped with certificates, so their essential properties can be checked independently. Interestingly, several model checking techniques, originally designed to analyse systems, can be extended to synthesis tasks. Tools like Uppaal-TIGA [BCD⁺07] and Uppaal-Stratego [DJL⁺15] can solve real-timed games. PRISM-games [KNPS20] can solve stochastic games. The winning strategies generated by such algorithms can (in principle) be converted to safe (and optimal) controllers. The synthesis of safe and optimal driving strategies with Uppaal-Stratego has recently been shown for ERTMS Level 3 moving block railway signalling [BtBL20]. However, certification of synthesis algorithms is still open. In particular, the generated controllers tend to be large and enumerate possibilities, rather than conditions on data.

Recent work [AJK⁺21] proposed to use decision trees to represent winning strategies; this could be a useful approach to synthesise controllers that are not just correct and certifiable but also explainable.

2.3 Explainable Verification

While certifiable verification can significantly increase reliability of the verification process, it is a whole challenge in itself is to transform the computed certificate into an understandable piece of evidence, which can contribute to the safety case for certification authorities. This explainable verification should be the ultimate goal of this line of research. Without aspiring completeness, we identify some open challenges in explainable verification.

Documentation of verification In particular, this should contain the precise claim of what has been verified by the tools, including the modelling assumptions, the assumptions on the environment, and potential approximations and inaccuracies implied by the selected options in the verification tools.

Interactive explainability It could be fruitful to base explainable verification methods on counterfactual explanation techniques for AI [HIA21]. Counterfactuals provide an understanding into AI models by identifying similar inputs with changes in decisive properties that lead to a different model outcome than the one under study. For instance, given a particular dangerous scenario, the verification tool should be able to generate an argument why this particular scenario cannot happen. Novel interactive approaches that allow in-depth investigation of these properties for verification models will be needed. This is important to increase the trust of domain experts and certification authorities in the verification technology.

Practicality Application-oriented research is required to investigate if certified model checking can provide useful evidence for the safety-case of railway systems, so that it significantly speeds-up the (regulatory) certification process for novel railway technology. This requires a careful consideration of the current standards used for certification in railways.

2.4 Standardisation

Certification of autonomous train control systems with Grade of Automation GoA 4 (unattended train operation, neither the driver nor the staff are required) in open railway environments, is a challenge: While conventional train control sub-systems can be certified on the basis of today’s CENELEC standards [Eur17, Eur11, Eur18], this is not the case for all AI-based sub-systems. The certification of such AI-based systems will require extensions/modifications of the current CENELEC standards or additional use of other standards. Therefore, there is a need to investigate how that can be done. In [PHL22], Peleska et al have investigated how the ANSI/UL 4600 pre-standard for Evaluation of Autonomous Products [Und22] can be used as a supplement to the CENELEC standards to certify autonomous freight trains and metro trains based on AI technology.

3 Formal Modelling and Analysis for the Railway Domain

Historically, the application of formal methods in order to verify railway systems is well established within academia and is beginning to see real applications in industry [FtB22]. As early as 1995, formal methods were applied to verify an interlocking for controlling Dutch railway stations [BBV95, GVK95]. Since then, and indeed recently, newer approaches to interlocking verification have also been proposed, also at ISoLA, and have been shown to scale well to modern industrial systems [BFG⁺98, CGM⁺98, Eis02, HKB11, FMGF10, Fan12, Win12, JLM⁺14, JMN⁺14a, JMN⁺14b, MFH16, CLS⁺17, HHP17, BJL⁺18, BJL19]. In spite of this, such approaches still lack widespread use within the Rail industry often due to questions surrounding the usability and expertise required for applying formal methods [FMB⁺20, FMBtB21].

Railway infrastructure managers have started to use semi-formal modelling languages (e.g., UML and SysML) to specify requirements, but they often still have to delegate formal verification activities to academic partners. The culprit is that effectively using state-of-the-art formal verification technology requires a thorough academic background in formal methods. A formal method that can be used by railway engineers needs to facilitate modelling railway systems concisely at the right level of abstraction and it should be easily parametrisable with relevant data (e.g., a track layout). Also, it must allow for

a straightforward specification of relevant safety and security properties, verification algorithms that scale for systems in the railway context, and provide insightful presentation of verification results.

In addition, in the next two decades, European railway infrastructure managers need to sustain an enormous growth in mobility by increasing the capacity of their networks at acceptable costs. Key to the capacity increase will be the introduction of innovative digital systems such as ERTMS level 3, which involves a radically new approach to train separation, and various forms of Automatic Train Operation (ATO). The smooth roll-out of such systems on the dense European railway networks is an enormous challenge. It is, e.g., unacceptable for a railway line to be unavailable for long periods, and it should be possible for the innovative system to coexist with the legacy system. Therefore, it is important to thoroughly prepare roll-out of new systems. Extensive use of formal modelling and analysis techniques in the development process will reduce the need for testing in the field.

3.1 Domain Specific Technology and Usability

Railway infrastructure managers have started to use (semi-)formal languages to model their systems. These models are typically very detailed and use concrete data. For an effective formal analysis, a domain-specific modelling language that supports the appropriate level of abstraction is essential. Furthermore, the modelling language should offer a means to model relevant continuous aspects of railway systems (e.g., braking curves). Safety and liveness properties must be formulated at the same level of abstraction. The latter is non-trivial in the context of the railway domain since normally railway engineers are not used to formulating safety requirements at the appropriate level of abstraction. It will be necessary to develop a property language that is, on the one hand, expressive enough to express relevant safety requirements, and on the other hand can be used by railway engineers. Formal verification in the railway sector has focussed on safety properties. The verification of liveness properties has not received much attention but is highly relevant in view of dependability of railway systems. For the verification of liveness properties, one typically needs to incorporate progress or fairness assumptions in the verification process. Verification technology for the domain-specific modelling language should be built on top of a state-of-the-art general-purpose verification engine.

Another concern that is limiting uptake is the need to model in such a way that requirements can be efficiently verified for all relevant track layouts. With virtual fixed or moving blocks, data-parameterised verification becomes more important because track layouts are dynamically configurable. Currently, formal verification by model checking must often be carried out for specific track layouts. There is a need to develop formal modelling and verification technology that can efficiently verify safety and liveness properties of a signalling system for a class of relevant track layouts. To this end, the modelling language should facilitate data parametrisation. Not all data are realistic and so any approach will need to investigate ways of efficiently expressing assumptions on the parametrisation domain (e.g., using probability distributions). These assumptions must then be considered in the verification activities. Finally, a core aspect of this verification process is that it must be usable by railway engineers, a non-trivial endeavour. For railway-specific modelling and verification technology it is a major concern that railway engineers without extensive formal methods expertise can use it to gain insights in their systems. To this end, the modelling language should be easy to use, and verification results should be visualised (e.g., by running a graphical simulation of a counterexample). Railway engineers verify their designs by considering how they behave with respect to various operational scenarios. Model-checking technology can be used to generate interesting operational scenarios as evidence for certain properties. The idea is to formulate meaningful properties to verify and obtain the operational scenarios as evidence.

3.2 Standardised Reference Architectures

To facilitate interoperability, European railway infrastructure managers and operators and railway supply industry pursue standardisation of command and control systems. The best-known example is the European Rail Traffic Management System which aims to standardise train-trackside communication (GSM-R), the train control system (ETCS), and the train management layer (ERTMS). The standard is formulated mostly in natural language and is therefore inherently ambiguous, which hampers a smooth deployment. Another standardisation project is EULYNX¹⁰; it aims to standardise

¹⁰ <https://eulynx.eu>

the interfaces between the components of signalling systems (interlockings and field elements such as points, level crossings, light signals, etc.). In EULYNX the official standard is still formulated in natural language, but there is also an explicit aim to supplement standard with SysML models. The academic project FormaSig¹¹ develops verification and model-based testing technology by which these SysML models can be formally verified and used for model-based testing purposes [BWL⁺22]. Similar to EULYNX, railway operators have started OCORA in which a standardised modular architecture for on-board command and control equipment is developed. This effort too should be supported with formal methods.

Although parts of the ETCS standard have been formally modelled and analysed [PQ09, CCM⁺10, CCL⁺12, BCPV17, BLW18, BtBC18, BJL⁺18, BtBFL19, BFR21], it would be both challenging and highly desirable to develop a formal model of ETCS level 3. The ongoing development by railway infrastructure managers of an RCA, a common reference architecture for railway command and control systems will be a convenient vehicle. This reference architecture consists of standardised components (e.g., interlocking and field elements with EULYNX-compliant interfaces), which will facilitate the deployment of innovative systems. Here, one particular challenge that is still open is to integrate ETCS level 3 with RCA by developing an integral model including all relevant components, in order to analyse the correctness of the interactions between those components and determine if the safety and security requirements are met. Of course this also requires a systematic analysis of both safety and security requirements for ETCS level 3.

3.3 Digital Railway Innovations

A number of upcoming digital railway innovations bring promises in terms of improved safety, capacity and resilience. With these adaptations to infrastructure come fresh challenges for formal modelling and verification in particular throughout the certification process of these systems.

A significant increase in capacity can be realised by going from blocks protected by train detection equipment to train separation based on more precise position information from the train. The introduction of such a new train separation system entails new challenges for the signalling system. For instance, it needs to be robust against radio connection problems between train and trackside and take into account inaccuracies in positioning information. But most importantly, for a significant period of time such a new system needs to coexist together with the old system. Railway infrastructure managers are designing digital systems to deal with such issues, and these digital systems are an order of magnitude more complex than what inframanagers are used to. An example is the Hybrid ERTMS/ETCS Level 3 concept, which tries to bring the flexibility of moving block train separation to a signalling system based on traditional trackside train detection [FHAB17]. The application of formal modelling and analysis techniques have proved to be beneficial for improving the specification of the concept and for coping with its complexity [BLW18, BHRR20, AKJ20, CM20, DDPS20, HLK⁺20, MFTL20, FFLM20, BFRM21].

A further increase in capacity, energy consumption and reliability is expected to come from Automated Train Operation (ATO). Railway infrastructure managers and operators are currently experimenting with a form of semi-automatic train operation, in which the automatic train operation system is not connected to the safety system meaning that the driver is still responsible for safe movement of the train. The next step is to integrate a form of ATO with a signalling system, and in particular ERTMS/ETCS. This integration serves as the perfect example of a complex hybrid system and thus modelling and verification challenges that exist for hybrid systems apply. In particular, this integration poses challenges in terms of the discrete and continuous nature of the system. Concretely, models need to be developed that focus on the interaction of the ATO system with the safety system, and that consider modelling the influence of braking curves. Following this, suitable abstractions that involve reasoning over continuous data need to be explored and verification processes that scale for such a setting developed.

4 Formal Methods for AI

Formal methods have become a well-established and widely applied technique for ensuring the correctness of fundamental components of safety-critical systems in the railway domain, in particular

¹¹ <https://fsa.win.tue.nl/formasig>

verification techniques based on model checking (cf. Section 2.1). However, while a survey on software engineering for AI-based systems considered 248 studies published during the past decade, of which more than two-thirds since 2018 [MBF⁺22], the application of formal methods to AI-based systems is still in its infancy. A recent paper by Wing [Win21] lists the following three key insights:

- The set of trustworthiness properties for AI systems, in contrast to traditional computing systems, needs to be extended beyond reliability, security, privacy, and usability to include properties such as probabilistic accuracy under uncertainty, fairness, robustness, accountability, and explainability.
- To help ensure their trustworthiness, AI systems can benefit from the scrutiny of formal methods.
- AI systems raise the bar on formal methods for two key reasons: the inherent probabilistic nature of machine-learned models, and the critical role of data in training, testing, and deploying a machine-learned model.

We envision to improve this situation by developing verification techniques that provide explainability or guarantees for AI-based systems in the specific safety-critical domain of railway systems, because, as Bešinović et al. put it, “although AI is still in its very infancy for the railway sector, there is certain evidence showing that its potential should not be underestimated” [BDF⁺22]. To this aim, we will first need to identify the state of the art of formal methods techniques developed for and applied to systems with AI-based components in the specific setting of transport systems (railways, but also automotive [SKOK18, TFP⁺20]) and of safety certification. A key difference with respect to the traditional formal verification approach, i.e. verifying a correctness property specified in some logic over a system model, is the inherent probabilistic nature of the (machine-learned) model in case of AI-based systems. We intend to study how to deal with safety concerns in the presence of uncertainty, and how probabilistic (and statistical) model-checking techniques or correctness-by-construction techniques can be adapted to provide appropriate fail-safe guarantees for AI-based railway systems. Correctness-by-construction, in particular when considering also non-functional properties (X-by-construction), in combination with probability and runtime verification is being studied also at ISoLA [tBCL⁺20, tBCLS22]. The same holds for formal methods for AI [LLSS22].

4.1 Guaranteeing Safety Behaviour

The wide availability of AI technologies and the pace of their evolution makes it hard for industry, with its consolidated processes, to profit from the potential benefits offered by these techniques. This is particularly true for the railway domain, in which the safety culture is strong, thus reinforcing the attachment to traditional, well-established practices that have proven their relevance even for the development of software systems that are not safety-critical. While many railway systems are required to fulfil SIL-4 certification requirements, many others, most notably maintenance systems, could instead make full use of the benefit of AI and process or data mining techniques. At the same time, explainability always must be ensured, to guarantee that the system behaviour can be explained, also for legal reasons, in case a failure occurs. Additionally, correctness is desirable as it would be difficult for a human to intervene and rectify mistakes made by such a system. For example, if a train is incorrectly routed across a junction it may take some time before it can be routed back across the junction in the correct direction and this would have impact on other trains in the area.

A concrete example from industry concerns an AI scheduler. Modern railway control systems are equipped with automatic route setting and traffic management but it is not clear how well they perform, specifically in case of a divergence from operational norms (e.g., if a train breaks down). It is typically the task of a qualified human to ensure trains run according to schedule, and to intervene when problems occur. Initial implementations of an AI scheduler could provide the human with guidance in such cases, assisting the human signaller in an efficient fashion to return the railway to an operational state by rerouting trains around the problem. When assisted by such an AI component, it is natural to want solutions that are presented by the AI component to both be explained or justified, and to not lead to intervention from safety critical components (for example, by providing solutions that violate the rules of the governing interlocking). For the human in the loop case it is essential for the AI system to be explainable and produce a justification that can be manually checked prior to making a decision.

4.2 Learning Formal Models of Railway Behaviour

Formal behavioural models are the building blocks of automated verification techniques in the field of formal methods. Such models define how a system behaves as a result of interacting with its users and

its environment. Traditionally, these models (e.g. variants of automata and state machines) are obtained starting from semi-formal models (e.g. UML and message sequence charts) developed during the initial development phase. However, it frequently occurs that such behavioural models are either unavailable or outdated and thus need to be reconstructed from implementations in order to enable formal analysis. In such cases, model learning is an automated technique that can produce such models. This is a popular research field and much progress has been made since Vaandrager noted that “even though model learning has been applied successfully in several domains, the field is still in its infancy” [Vaa17]. Recent examples include [ABE⁺19, DMdSS21]. However, to the best of our knowledge, specific success stories in the railway domain are missing.

We envision the usage of data or process mining techniques to build digital twins of railway systems that can provide predictive (runtime) behaviour and ultimately enable real-time predictive monitoring and maintenance (cf. Section 4.3). Engineering digital twins is being studied also at ISoLA [FLMW20, MS21, FLM⁺22], including some initial, recent attempts in the railway domain [Lec20]. This requires the use of a variety of techniques from formal methods, in particular probabilistic (and statistical) model checking to deal with the inherent probabilistic nature of the (machine-learned) model, game theory (for instance for controller synthesis), and automata or model learning, but also specific techniques from data or process mining [vdA16]. Railway system models are characterised by the need to deal with real-time aspects and a degree of uncertainty. We will thus have to study how to perform data or process mining on the provided observation data, like execution traces (i.e., logs) of a railway system, and how to use this to learn a digital twin of the railway system. This digital twin is meant to be a formal model that can handle real time and probabilistic or stochastic behaviour (e.g., conform the timed stochastic models accepted by the Uppaal model checker).

4.3 AI for Monitoring and Maintenance

Current railway monitoring and maintenance systems are mostly rule-based and typically do not include AI-based components, which can be particularly useful, e.g., to predict possible failures and to plan specific maintenance actions [CSV⁺19, NHSvN19]. AI systems, combined with existing rules provided by experts, can enable predictive maintenance, by identifying patterns of faults based on systems logs. Model learning for maintenance implies learning of the system model, and learning of a system’s fault model, so that future faults can be predicted based on current system behaviour. The system’s digital twin can also be used to forecast and simulate future long-term scenarios, thus helping to plan for maintenance actions in advance, i.e., before faults occur. Refactoring current rule-based maintenance systems with AI-based components poses numerous challenges. The effective exploitation of AI and process or data mining techniques requires the domain experts to annotate field data, such that the machine can learn from experts. Similarly, experts are needed to assess the correct behaviour and interpretation of possible failures of the AI-based maintenance system. Explainability of these systems becomes crucial, as well as correct communication of the behaviour’s explanation to experts and other railway stakeholders. In the envisioned maintenance process, the behaviour of the onboard system can easily be reconstructed and visualised, and maintenance and improvement actions can be taken in a more flexible and effective way.

4.4 AI for Optimisation in Scheduling and Design

Whilst so far we addressed principal problems that need to be solved when using AI techniques, we now want to look at several specific applications in the railway domain. The first two applications concern problems for which currently no general optimal solutions exist, but where one can hope with the use of AI to achieve better solutions, i.e. solutions which are more efficient in terms of energy/time or which require less track-side equipment. For the solutions in this section it is, of course, essential that they are still safe and fulfil all safety requirements. Therefore it is anticipated that the solutions produced are not only very good, but that they also come with a guarantee or explanation. Overall, optimisations like, saving energy or requiring less track-side equipment contribute to the aim of reducing carbon. Also capacity improvements support the green deal (indirectly), as efficient and safe provision leads to a higher customer satisfaction and a higher uptake of railway use.

Our first application refers to scheduling which has frequently been considered as an optimisation problem in the past [JMN⁺14c, JR14, DH19, MSLF14, PPPW22]. An AI Scheduler (see section 4.1)

could help to support and improve the decision making of a human signaller and optimising the flow of trains through the railway network.

A second standard application concerns the optimisation of railway design and layout. Railways are designed by engineers who create scheme plans with the topology of the railway and the layout of the equipment along the tracks. Publicised solutions for the automatic generation of signalling design seem not to consider optimisation at all. Desired would be an AI solution for the placement of the equipment that is still fulfils all required constraints, such as number of balises in a given area, or a requested distance between balise groups, etc. Model checking/SMT solving can then be used to check the constraints and highlight counter examples in a efficient way (where the visualisation of counterexamples in a domain specific area constitutes an interesting problem on its own.) Various AI techniques as well as Game theory and Explainable AI (XAI) can be used for the optimisation according to a given measure such as energy/material consumption.

5 Conclusion and Further Work

In this article, we discussed several challenges for Formal Methods in Railway linked to the areas of (1) Verification and Certification, (2) Modelling of ETCS related systems, and, (3) the use of AI in the railway domain. Their common aim is a robust development of complex railway technology that is reliable and efficient at the same time. Our specific focus in (1) and (3) was that any (new) techniques need to come with explanations/ guarantees in order to be accepted by the Railway engineers and the general public. Specifically, regarding the question of using AI techniques in Railway, further research is needed about which AI techniques we can apply and how to create explanations and guarantees. This is relatively straightforward in the case of applying, for instance, SMT-solving. Here, SMT solving would provide counterexamples, and the gap to be closed concentrates on making these counterexamples (1) readable in the Railway context and (2) independently verifiable. Conversely, at the other end of the spectrum, if we want to apply machine-learning techniques for, e.g., classification, the situation is completely different: Explainable AI (XAI), with LIME and SHAP as prominent techniques, has become established for providing explanations, however recent work has demonstrated that different XAI techniques do not necessarily coincide on their results [DFB⁺21]. Problems like these prompt research towards a theory of faithful explanations [PYT22] and the idea of Verifiable XAI. More effort and case studies will be needed to develop these techniques and make them usable in the Railway context.

References

- ABE⁺19. Bernhard K. Aichernig, Roderick Bloem, Masoud Ebrahimi, Martin Horn, Franz Pernkopf, Wolfgang Roth, Astrid Rupp, Martin Tappler, and Markus Tranninger. Learning a behavior model of hybrid systems through combining model-based testing and machine learning. In Christophe Gaston, Nikolai Kosmatov, and Pascale Le Gall, editors, *Proceedings of the 31st IFIP WG 6.1 International Conference on Testing Software and Systems (ICTSS'19)*, volume 11812 of *LNCS*, pages 3–21. Springer, 2019.
- ADL21. ADLINK Technology. Transforming the rail industry with artificial intelligence, September 2021. Whitepaper, <https://www.globalrailwayreview.com/whitepaper/127609/transforming-the-rail-industry-with-ai>.
- AJK⁺21. Pranav Ashok, Mathias Jackermeier, Jan Kretínský, Christoph Weinhuber, Maximilian Weininger, and Mayank Yadav. dtControl 2.0: Explainable strategy representation via decision tree learning steered by experts. In *Proceedings of the 27th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'21)*, volume 12652 of *LNCS*, pages 326–345. Springer, 2021.
- AKJ20. Paolo Arcaini, Jan Kofroň, and Pavel Ježek. Validation of the hybrid ERTMS/ETCS level 3 using SPIN. *Int. J. Softw. Tools Technol. Transf.*, 22(3):265–279, 2020.
- And19. Étienne André. What’s decidable about parametric timed automata? *Int. J. Softw. Tools Technol. Transf.*, 21(2):203–219, 2019.
- AP18. Gul Agha and Karl Palmskog. A survey of statistical model checking. *ACM Trans. Model. Comput. Simul.*, 28(1):6:1–6:39, 2018.
- BBV95. Twan Basten, Roland N. Bol, and Marc Voorhoeve. Simulating and analyzing railway interlockings in ExSpect. *IEEE Parallel Distributed Technol. Syst. Appl.*, 3(3):50–62, 1995.

- BC04. Yves Bertot and Pierre Castéran. *Interactive Theorem Proving and Program Development - Coq'Art: The Calculus of Inductive Constructions*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2004.
- BCD⁺07. Gerd Behrmann, Agnès Cougnard, Alexandre David, Emmanuel Fleury, Kim Guldstrand Larsen, and Didier Lime. UPPAAL-Tiga: Time for playing games! In Werner Damm and Holger Hermanns, editors, *Proceedings of the 19th International Conference on Computer Aided Verification (CAV'07)*, volume 4590 of *LNCS*, pages 121–125. Springer, 2007.
- BCPV17. Marco Biagi, Laura Carnevali, Marco Paolieri, and Enrico Vicario. Performability evaluation of the ERTMS/ETCS – level 3. *Transp. Res. C-Emer.*, 82:314–336, 2017.
- BDF⁺22. Nikola Bešinović, Lorenzo De Donato, Francesco Flammini, Rob M. P. Goverde, Zhiyuan Lin, Ronghui Liu, Stefano Marrone, Roberto Nardone, Tianli Tang, and Valeria Vittorini. Artificial intelligence in railway transport: Taxonomy, regulations and applications. *IEEE Trans. Intell. Transp. Syst.*, 2022.
- BFG⁺98. Cinzia Bernardeschi, Alessandro Fantechi, Stefania Gnesi, Salvatore Larosa, Giorgio Mongardi, and Dario Romano. A formal verification environment for railway signaling system design. *Formal Methods Syst. Des.*, 12(2):139–161, 1998.
- BFLW18. Jasmin Christian Blanchette, Mathias Fleury, Peter Lammich, and Christoph Weidenbach. A verified SAT solver framework with learn, forget, restart, and incrementality. *J. Autom. Reason.*, 61(1-4):333–365, 2018.
- BFR21. Davide Basile, Alessandro Fantechi, and Irene Rosadi. Formal analysis of the UNISIG safety application intermediate sub-layer: Applying formal methods to railway standard interfaces. In Alberto Lluch-Lafuente and Anastasia Mavridou, editors, *Proceedings of the 26th International Conference on Formal Methods for Industrial Critical Systems (FMICS'21)*, volume 12863 of *LNCS*, pages 174–190. Springer, 2021.
- BFRM21. Davide Basile, Alessandro Fantechi, Luigi Rucher, and Gianluca Mandò. Analysing an autonomous tramway positioning system with the UPPAAL statistical model checker. *Form. Asp. Comp.*, 33(6):957–987, 2021.
- BHRR20. Michael J. Butler, Thai Son Hoang, Alexander Raschke, and Klaus Reichl. Introduction to special section on the ABZ 2018 case study: Hybrid ERTMS/ETCS level 3. *Int. J. Softw. Tools Technol. Transf.*, 22(3):249–255, 2020.
- BJL⁺18. Ulrich Berger, Phillip James, Andrew Lawrence, Markus Roggenbach, and Monika Seisenberger. Verification of the European Rail Traffic Management System in real-time Maude. *Sci. Comput. Program.*, 154:61–88, 2018.
- BJL19. Mark Bouwman, Bob Janssen, and Bas Luttik. Formal modelling and verification of an interlocking using mCRL2. In Kim Guldstrand Larsen and Tim A. C. Willemse, editors, *Proceedings of the 24th International Conference on Formal Methods for Industrial Critical Systems (FMICS'19)*, volume 11687 of *LNCS*, pages 22–39. Springer, 2019.
- BK08. Christel Baier and Joost-Pieter Katoen. *Principles of Model Checking*. MIT Press, 2008.
- BL18. Julian Brunner and Peter Lammich. Formal verification of an executable LTL model checker with partial order reduction. *J. Autom. Reason.*, 60(1):3–21, 2018.
- BLFS15. Ulrich Berger, Andrew Lawrence, Fredrik Nordvall Forsberg, and Monika Seisenberger. Extracting verified decision procedures: DPLL and resolution. *Log. Methods Comp. Sci.*, 11(1):1–18, 2015.
- BLR05. Patricia Bouyer, François Laroussinie, and Pierre-Alain Reynier. Diagonal constraints in timed automata: Forward analysis of timed systems. In Paul Pettersson and Wang Yi, editors, *Proceedings of the 3rd International Conference on Formal Modeling and Analysis of Timed Systems (FORMATS'05)*, volume 3829 of *LNCS*, pages 112–126. Springer, 2005.
- BLW18. Maarten Bartholomeus, Bas Luttik, and Tim A. C. Willemse. Modelling and analysing ERTMS hybrid level 3 with the mCRL2 toolset. In Falk Howar and Jiri Barnat, editors, *Proceedings of the 23rd International Conference on Formal Methods for Industrial Critical Systems (FMICS'18)*, volume 11119 of *LNCS*, pages 98–114. Springer, 2018.
- BtBC18. Davide Basile, Maurice H. ter Beek, and Vincenzo Ciancia. Statistical model checking of a moving block railway signalling scenario with UPPAAL SMC. In Tiziana Margaria and Bernhard Steffen, editors, *Proceedings of the 8th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation: Verification (ISoLA'18)*, volume 11245 of *LNCS*, pages 372–391. Springer, 2018.
- BtBFL19. Davide Basile, Maurice H. ter Beek, Alessio Ferrari, and Axel Legay. Modelling and Analysing ERTMS L3 Moving Block Railway Signalling with Simulink and UPPAAL SMC. In Kim G. Larsen and Tim Willemse, editors, *Proceedings of the 24th International Conference on Formal Methods for Industrial Critical Systems (FMICS'19)*, volume 11687 of *LNCS*, pages 1–21. Springer, 2019.
- BtBL20. Davide Basile, Maurice H. ter Beek, and Axel Legay. Strategy synthesis for autonomous driving in a moving block railway system with UPPAAL STRATEGO. In Alexey Gotsman and Ana Sokolova, editors, *Proceedings of the 40th IFIP WG 6.1 International Conference on Formal Techniques for*

- Distributed Objects, Components, and Systems (FORTE'20)*, volume 12136 of *LNCS*, pages 3–21. Springer, 2020.
- BWL⁺22. Mark Bouwman, Djurre van der Wal, Bas Luttik, Mariëlle Stoelinga, and Arend Rensink. A case in point: Verification and testing of a EULYNX interface. *Form. Asp. Comput.*, 2022.
- CCL⁺12. Alessandro Cimatti, Raffaele Corvino, Armando Lazzaro, Iman Narasamdya, Tiziana Rizzo, Marco Roveri, Angela Sanseviero, and Andrei Tchaltsev. Formal verification and validation of ERTMS industrial railway train spacing system. In P. Madhusudan and Sanjit A. Seshia, editors, *Proceedings of the 24th International Conference on Computer Aided Verification (CAV'12)*, volume 7358 of *LNCS*, pages 378–393. Springer, 2012.
- CCM⁺10. Angelo Chiappini, Alessandro Cimatti, Luca Macchi, Oscar Rebollo, Marco Roveri, Angelo Susi, Stefano Tonetta, and Bernardino Vittorini. Formalization and validation of a subset of the European Train Control System. In *Proceedings of the 32nd International Conference on Software Engineering (ICSE'10)*, pages 109–118. ACM, 2010.
- CGM⁺98. Alessandro Cimatti, Fausto Giunchiglia, Giorgio Mongardi, Dario Romano, Fernando Torielli, and Paolo Traverso. Model checking safety critical software with SPIN: an application to a railway interlocking system. In Wolfgang D. Ehrenberger, editor, *Proceedings of the 17th International Conference on Computer Safety, Reliability and Security (SAFECOMP'98)*, volume 1516 of *LNCS*, pages 284–295. Springer, 1998.
- CHVB18. Edmund M. Clarke, Thomas A. Henzinger, Helmut Veith, and Roderick Bloem, editors. *Handbook of Model Checking*. Springer, 2018.
- CLS⁺17. Quentin Cappart, Christophe Limbrée, Pierre Schaus, Jean Quilbeuf, Louis-Marie Traonouez, and Axel Legay. Verification of interlocking systems using statistical model checking. In *Proceedings of the 18th International Symposium on High Assurance Systems Engineering (HASE'17)*, pages 61–68. IEEE, 2017.
- CM20. Alcino Cunha and Nuno Macedo. Validating the hybrid ERTMS/ETCS level 3 concept with Electrum. *Int. J. Softw. Tools Technol. Transf.*, 22(3):281–296, 2020.
- CSV⁺19. Thyago Peres Carvalho, Fabrizio Alphonso A. M. N. Soares, Roberto Vita, Roberto da Piedade Francisco, João Pedro Tavares Vieira Basto, and Symone G. S. Alcalá. A systematic literature review of machine learning methods applied to predictive maintenance. *Comput. Ind. Eng.*, 137, 2019.
- DDPS20. Dana Dghaym, Mohammadsadegh Dalvandi, Michael Poppleton, and Colin F. Snook. Formalising the hybrid ERTMS level 3 specification in iUML-B and Event-B. *Int. J. Softw. Tools Technol. Transf.*, 22(3):297–313, 2020.
- DFB⁺21. Jamie Duell, Xiuyi Fan, Bruce Burnett, Gert Aarts, and Shang-Ming Zhou. A comparison of explanations given by explainable artificial intelligence methods on analysing electronic health records. In *Proceedings of the 7th EMBS International Conference on Biomedical and Health Informatics (BHI'21)*, pages 1–4. IEEE, 2021.
- DH19. Stefan Dillmann and Reiner Hähnle. Automated planning of ETCS tracks. In Simon Collart Dutilleul, Thierry Lecomte, and Alexander B. Romanovsky, editors, *Proceedings of the 3rd International Conference on Reliability, Safety, and Security of Railway Systems: Modelling, Analysis, Verification, and Certification (RSSRail'19)*, volume 11495 of *LNCS*, pages 79–90. Springer, 2019.
- DJL⁺15. Alexandre David, Peter Gjøøl Jensen, Kim Guldstrand Larsen, Marius Mikucionis, and Jakob Haahr Taankvist. UPPAAL STRATEGO. In Christel Baier and Cesare Tinelli, editors, *Proceedings of the 21st International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'15)*, volume 9035 of *LNCS*, pages 206–211. Springer, 2015.
- DMdSS21. Carlos Diego Nascimento Damasceno, Mohammad Reza Mousavi, and Adenildo da Silva Simão. Learning by sampling: learning behavioral family models from software product lines. *Empir. Softw. Eng.*, 26(1):4:1–4:46, 2021.
- Eis02. Cindy Eisner. Using symbolic CTL model checking to verify the railway stations of Hoorn-Kersenboogerd and Heerhugowaard. *Int. J. Softw. Tools Technol. Transf.*, 4(1):107–124, 2002.
- ELN⁺14. Javier Esparza, Peter Lammich, René Neumann, Tobias Nipkow, Alexander Schimpf, and Jan-Georg Smaus. A fully verified executable LTL model checker. *Arch. Formal Proofs*, May 2014.
- Eur11. European Committee for Electrotechnical Standardization. CENELEC EN 50128 — Railway applications – Communication, signalling and processing systems – Software for railway control and protection systems, June 2011.
- Eur17. European Committee for Electrotechnical Standardization. CENELEC EN 50126-1 — Railway applications – The specification and demonstration of reliability, availability, maintainability and safety (RAMS) – Part 1: Generic RAMS process, October 2017.
- Eur18. European Committee for Electrotechnical Standardization. CENELEC EN 50129 — Railway applications – Communication, signalling and processing systems – Safety related electronic systems for signalling, November 2018.

- Fan12. Alessandro Fantechi. Distributing the challenge of model checking interlocking control tables. In Tiziana Margaria and Bernhard Steffen, editors, *Proceedings of the 5th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation: Applications and Case Studies (ISoLA'12)*, volume 7610 of *LNCS*, pages 276–289. Springer, 2012.
- FFLM20. Steve Jeffrey Tuono Fotso, Marc Frappier, Régine Laleau, and Amel Mammam. Modeling the hybrid ERTMS/ETCS level 3 standard using a formal requirements engineering approach. *Int. J. Softw. Tools Technol. Transf.*, 22(3):349–363, 2020.
- FHAB17. Nicola Furness, Henri van Houten, Laura Arenas, and Maarten Bartholomeus. ERTMS level 3: the game-changer. *IRSE News*, 232:2–9, April 2017.
- FLM⁺22. John S. Fitzgerald, Peter Gorm Larsen, Tiziana Margaria, Jim Woodcock, and Cláudio Gomes. Digital twin engineering. In Tiziana Margaria and Bernhard Steffen, editors, *Proceedings of the 11th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'22)*, LNCS. Springer, 2022.
- FLMW20. John S. Fitzgerald, Peter Gorm Larsen, Tiziana Margaria, and Jim Woodcock. Engineering of digital twins for cyber-physical systems. In Tiziana Margaria and Bernhard Steffen, editors, *Proceedings of the 9th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation: Tools and Trends (ISoLA'20)*, volume 12479 of *LNCS*, pages 49–53. Springer, 2020.
- FMB⁺20. Alessio Ferrari, Franco Mazzanti, Davide Basile, Maurice H. ter Beek, and Alessandro Fantechi. Comparing formal tools for system design: a judgment study. In *Proceedings of the 42nd International Conference on Software Engineering (ICSE'20)*, pages 62–74. ACM, 2020.
- FMBtB21. Alessio Ferrari, Franco Mazzanti, Davide Basile, and Maurice H. ter Beek. Systematic evaluation and usability analysis of formal tools for railway system design. *IEEE Trans. Softw. Eng.*, 2021.
- FMGF10. Alessio Ferrari, Gianluca Magnani, Daniele Grasso, and Alessandro Fantechi. Model checking interlocking control tables. In Eckehard Schnieder and Géza Tarnai, editors, *Proceedings of the 8th Symposium on Formal Methods for Automation and Safety in Railway and Automotive Systems (FORMS/FORMAT'10)*, pages 107–115. Springer, 2010.
- FtB22. Alessio Ferrari and Maurice H. ter Beek. Formal methods in railways: a systematic mapping study. *ACM Comput. Surv.*, 2022.
- GMS20. Frederik Gossen, Tiziana Margaria, and Bernhard Steffen. Towards explainability in machine learning: The formal methods way. *IT Prof.*, 22(4):8–12, 2020.
- GRT21. Alberto Griggio, Marco Roveri, and Stefano Tonetta. Certifying proofs for SAT-based model checking. *Formal Methods Syst. Des.*, 57(2):178–210, 2021.
- GVK95. Jan Friso Groote, Sebastiaan F.M. Vlijmen, and Jan W.C. Koorn. The safety guaranteeing system at station Hoorn-Kersenboogerd. In *Proceedings of the 10th Annual Conference on Computer Assurance Systems Integrity, Software Safety and Process Security (COMPASS'95)*, pages 57–68. IEEE, 1995.
- HHP17. Linh Vu Hong, Anne E. Haxthausen, and Jan Peleska. Formal modelling and verification of interlocking systems featuring sequential release. *Sci. Comput. Program.*, 133:91–115, 2017.
- HIA21. Frederik Hvilshøj, Alexandros Iosifidis, and Ira Assent. ECINN: Efficient counterfactuals from invertible neural networks. *CoRR*, abs/2103.13701, 2021.
- HKB11. Anne E. Haxthausen, Andreas A. Kjær, and Marie Le Bliguet. Formal development of a tool for automated modelling and verification of relay interlocking systems. In Michael J. Butler and Wolfram Schulte, editors, *Proceedings of the 17th International Symposium on Formal Methods (FM'11)*, volume 6664 of *LNCS*, pages 118–132. Springer, 2011.
- HKPV98. Thomas A. Henzinger, Peter W. Kopke, Anuj Puri, and Pravin Varaiya. What's decidable about hybrid automata? *J. Comput. Syst. Sci.*, 57(1):94–124, 1998.
- HLK⁺20. Dominik Hansen, Michael Leuschel, Philipp Körner, Sebastian Krings, Thomas Naulin, Nader Nayeri, David Schneider, and Frank Skowron. Validation and real-life demonstration of ETCS hybrid level 3 principles using a formal B model. *Int. J. Softw. Tools Technol. Transf.*, 22(3):315–332, 2020.
- JLM⁺14. Phillip James, Andy Lawrence, Faron Moller, Markus Roggenbach, Monika Seisenberger, Anton Setzer, Karim Kanso, and Simon Chadwick. Verification of solid state interlocking programs. In Steve Counsell and Manuel Núñez, editors, *Revised Selected Papers of the Collocated Workshops at the 11th International Conference Software Engineering and Formal Methods (SEFM'13)*, volume 8368 of *LNCS*, pages 253–268. Springer, 2014.
- JMN⁺14a. Phillip James, Faron Moller, Nguyen Hoang Nga, Markus Roggenbach, Steve A. Schneider, and Helen Treharne. Techniques for modelling and verifying railway interlockings. *Int. J. Softw. Tools Technol. Transf.*, 16(6):685–711, 2014.
- JMN⁺14b. Phillip James, Faron Moller, Hoang Nga Nguyen, Markus Roggenbach, Steve Schneider, and Helen Treharne. On modelling and verifying railway interlockings: Tracking train lengths. *Sci. Comput. Program.*, 96:315–336, 2014.

- JMN⁺14c. Phillip James, Faron Moller, Hoang Nga Nguyen, Markus Roggenbach, Steve A. Schneider, Helen Treharne, Matthew Trumble, and David M. Williams. Verification of scheme plans using CSP || B. In Steve Counsell and Manuel Núñez, editors, *Revised Selected Papers of the Collocated Workshops at the 11th International Conference on Software Engineering and Formal Methods (SEFM'13)*, volume 8368 of *LNCS*, pages 189–204. Springer, 2014.
- JR14. Phillip James and Markus Roggenbach. Encapsulating formal methods within domain specific languages: A solution for verifying railway scheme plans. *Math. Comput. Sci.*, 8(1):11–38, 2014.
- KBT⁺16. Guy Katz, Clark W. Barrett, Cesare Tinelli, Andrew Reynolds, and Liana Hadarean. Lazy proofs for DPLL(T)-based SMT solvers. In *Proceedings of the 16th Conference on Formal Methods in Computer-Aided Design (FMCAD'16)*, pages 93–100. IEEE, 2016.
- KLM⁺15. Gijs Kant, Alfons Laarman, Jeroen Meijer, Jaco van de Pol, Stefan Blom, and Tom van Dijk. LTSmin: High-performance language-independent model checking. In Christel Baier and Cesare Tinelli, editors, *Proceedings of the 21st International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'15)*, volume 9035 of *LNCS*, pages 692–707. Springer, 2015.
- KNPS20. Marta Kwiatkowska, Gethin Norman, David Parker, and Gabriel Santos. PRISM-games 3.0: Stochastic game verification with concurrency, equilibria and time. In Shuvendu K. Lahiri and Chao Wang, editors, *Proceedings of the 32nd International Conference on Computer Aided Verification (CAV'20)*, volume 12225 of *LNCS*, pages 475–487. Springer, 2020.
- Lec20. Thierry Lecomte. Digital modelling in the railways. In Tiziana Margaria and Bernhard Steffen, editors, *Proceedings of the 9th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation: Tools and Trends (ISoLA'20)*, volume 12479 of *LNCS*, pages 124–139. Springer, 2020.
- LLSS22. Kim G. Larsen, Axel Legay, Bernhard Steffen, and Mariëlle Stoelinga. Formal methods meet machine learning. In Tiziana Margaria and Bernhard Steffen, editors, *Proceedings of the 11th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'22)*, LNCS. Springer, 2022.
- MBF⁺22. Silverio Martínez-Fernández, Justus Bogner, Xavier Franch, Marc Oriol, Julien Siebert, Adam Trendowicz, Anna Maria Vollmer, and Stefan Wagner. Software engineering for AI-based systems: A survey. *ACM Trans. Softw. Eng. Methodol.*, 31(2):37e:1–37e:59, 2022.
- MFH16. Hugo Daniel Macedo, Alessandro Fantechi, and Anne E. Haxthausen. Compositional verification of multi-station interlocking systems. In Tiziana Margaria and Bernhard Steffen, editors, *Proceedings of the 7th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation: Discussion, Dissemination, Applications (ISoLA'16)*, volume 9953 of *LNCS*, pages 279–293, 2016.
- MFH17. Hugo Daniel Macedo, Alessandro Fantechi, and Anne E. Haxthausen. Compositional model checking of interlocking systems for lines with multiple stations. In Clark W. Barrett, Misty Davies, and Temesghen Kahsai, editors, *Proceedings of the 9th International Symposium on NASA Formal Methods (NFM'17)*, volume 10227 of *LNCS*, pages 146–162, 2017.
- MFTL20. Amel Mammam, Marc Frappier, Steve Jeffrey Tuono Fotso, and Régine Laleau. A formal refinement-based analysis of the hybrid ERTMS/ETCS level 3 standard. *Int. J. Softw. Tools Technol. Transf.*, 22(3):333–347, 2020.
- MS21. Tiziana Margaria and Alexander Schieweck. Towards engineering digital twins by active behaviour mining. In Ernst-Rüdiger Olderog, Bernhard Steffen, and Wang Yi, editors, *Model Checking, Synthesis, and Learning*, volume 13030 of *LNCS*, pages 138–163. Springer, 2021.
- MSLF14. Franco Mazzanti, Giorgio Oronzo Spagnolo, Simone Della Longa, and Alessio Ferrari. Deadlock avoidance in train scheduling: A model checking approach. In Frédéric Lang and Francesco Flammini, editors, *Proceedings of the 19th International Conference on Formal Methods for Industrial Critical Systems (FMICS'14)*, volume 8718 of *LNCS*, pages 109–123. Springer, 2014.
- Nam01. Kedar S. Namjoshi. Certifying model checkers. In Gérard Berry, Hubert Comon, and Alain Finkel, editors, *Proceedings of the 13th International Conference on Computer Aided Verification (CAV'01)*, volume 2102 of *LNCS*, pages 2–13. Springer, 2001.
- NHSvN19. Muhammad Chenariyan Nakhaee, Djoerd Hiemstra, Mariëlle Stoelinga, and Martijn van Noort. The recent applications of machine learning in rail track maintenance: A survey. In Simon Collart Dutilleul, Thierry Lecomte, and Alexander B. Romanovsky, editors, *Proceedings of the 3rd International Conference on Reliability, Safety, and Security of Railway Systems: Modelling, Analysis, Verification, and Certification (RSSRail'19)*, volume 11495 of *LNCS*, pages 91–105. Springer, 2019.
- NPW02. Tobias Nipkow, Lawrence C. Paulson, and Markus Wenzel. *Isabelle/HOL: A Proof Assistant for Higher-Order Logic*, volume 2283 of *LNCS*. Springer, 2002.
- OHJvdP20. Wytse Oortwijn, Marieke Huisman, Sebastiaan J. C. Joosten, and Jaco van de Pol. Automated verification of parallel nested DFS. In Armin Biere and David Parker, editors, *Proceedings of the 26th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'20)*, volume 12078 of *LNCS*, pages 247–265. Springer, 2020.

- PHL22. Jan Peleska, Anne E. Haxthausen, and Thierry Lecomte. Standardisation considerations for autonomous train control, 2022. Submitted.
- PKPB21. Stefan Pranger, Bettina Könighofer, Lukas Posch, and Roderick Bloem. TEMPEST – Synthesis tool for reactive systems and shields in probabilistic environments. In Zhe Hou and Vijay Ganesh, editors, *Proceedings of the 19th International Symposium on Automated Technology for Verification and Analysis (ATVA '21)*, volume 12971 of *LNCS*, pages 222–228. Springer, 2021.
- PPPW22. Tom Peham, Judith Przigoda, Nils Przigoda, and Robert Wille. Optimal railway routing using virtual subsections. In Simon Collart Dutilleul, Anne E. Haxthausen, and Thierry Lecomte, editors, *Proceedings of the 4th International Conference on Reliability, Safety, and Security of Railway Systems: Modelling, Analysis, Verification, and Certification (RSSRail'22)*, volume 13294 of *LNCS*, pages 63–79. Springer, 2022.
- PQ09. André Platzer and Jan-David Quesel. European train control system: A case study in formal verification. In Karin Breitman and Ana Cavalcanti, editors, *Proceedings of the 11th International Conference on Formal Engineering Methods (ICFEM'09)*, volume 5885 of *LNCS*, pages 246–265. Springer, 2009.
- PYT22. Nico Potyka, Xiang Yin, and Francesca Toni. Towards a theory of faithfulness: Faithful explanations of differentiable classifiers over continuous data. *CoRR*, abs/2205.09620, 2022.
- RAI. RAILS: Roadmaps for A.I. integration in the rail Sector. Shift2Rail, 2019–2023, <https://cordis.europa.eu/project/id/881782>.
- RPS⁺19. Talia Ringer, Karl Palmskog, Ilya Sergey, Milos Gligoric, and Zachary Tatlock. QED at large: A survey of engineering of formally verified software. *Found. Trends Program. Lang.*, 5(2-3):102–281, 2019.
- Sie19. Stephen F. Siegel. What’s wrong with on-the-fly partial order reduction. In Isil Dillig and Serdar Tasiran, editors, *Proceedings of the 31st International Conference on Computer Aided Verification (CAV'19)*, volume 11562 of *LNCS*, pages 478–495. Springer, 2019.
- SKOK18. Sina Shafaei, Stefan Kugele, Mohd Hafeez Osman, and Alois C. Knoll. Uncertainty in machine learning: A safety perspective on autonomous driving. In Barbara Gallina, Amund Skavhaug, Erwin Schoitsch, and Friedemann Bitsch, editors, *Proceedings of the Workshops at the 37th International Conference on Computer Safety, Reliability, and Security (SAFECOMP'18)*, volume 11094 of *LNCS*, pages 458–464. Springer, 2018.
- tBCL⁺20. Maurice H. ter Beek, Loek Cleophas, Axel Legay, Ina Schaefer, and Bruce W. Watson. X-by-Construction: Correctness meets probability. In Tiziana Margaria and Bernhard Steffen, editors, *Proceedings of the 9th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation: Verification Principles (ISoLA'20)*, volume 12476 of *LNCS*, pages 211–215. Springer, 2020.
- tBCLS22. Maurice H. ter Beek, Loek Cleophas, Martin Leucker, and Ina Schaefer. X-by-Construction meets runtime verification. In Tiziana Margaria and Bernhard Steffen, editors, *Proceedings of the 11th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'22)*, LNCS. Springer, 2022.
- TFP⁺20. Cumhuri Erkan Tuncali, Georgios Fainekos, Danil V. Prokhorov, Hisahiro Ito, and James Kapinski. Requirements-driven test generation for autonomous vehicles with machine learning components. *IEEE Trans. Intell. Veh.*, 5(2):265–280, 2020.
- Und22. Underwriters Laboratories Inc. *ANSI/UL 4600 Standard for Safety Evaluation of Autonomous Products*, March 2022.
- Vaa17. Frits W. Vaandrager. Model learning. *Commun. ACM*, 60(2):86–95, 2017.
- vdA16. Wil M. P. van der Aalst. *Process Mining: Data Science in Action*. Springer, 2016.
- vdP15. Jaco C. van de Pol. Automated verification of nested DFS. In Manuel Núñez and Matthias Güdemann, editors, *Proceedings of the 20th International Workshop on Formal Methods for Industrial Critical Systems (FMICS'15)*, volume 9128 of *LNCS*, pages 181–197. Springer, 2015.
- WHvdP20. Simon Wimmer, Frédéric Herbretreau, and Jaco van de Pol. Certifying emptiness of timed Büchi automata. In Nathalie Bertrand and Nils Jansen, editors, *Proceedings of the 18th International Conference on Formal Modeling and Analysis of Timed Systems (FORMATS'20)*, volume 12288 of *LNCS*, pages 58–75. Springer, 2020.
- Wie06. Freek Wiedijk, editor. *The Seventeen Provers of the World*. LNCS 3600. Springer, 2006.
- Win12. Kirsten Winter. Optimising ordering strategies for symbolic model checking of railway interlockings. In Tiziana Margaria and Bernhard Steffen, editors, *Proceedings of the 5th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation: Applications and Case Studies (ISoLA'12)*, volume 7610 of *LNCS*, pages 246–260. Springer, 2012.
- Win21. Jeannette M. Wing. Trustworthy AI. *Commun. ACM*, 64(10):64–71, 2021.
- WL18. Simon Wimmer and Peter Lammich. Verified model checking of timed automata. In Dirk Beyer and Marieke Huisman, editors, *Proceedings of the 24th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'18)*, volume 10805 of *LNCS*, pages 61–78. Springer, 2018.