Thursday, April 30, 2026

TReC 2026 Project Proposal Submission Form

Submit your project proposal for the 7th TRAIL Research Camp (August 24th - September 4th, 2026, Lausanne, Switzerland). Please complete all required sections and submit your proposal before April 30th, 01:00 PM (CET).

Administrative Data

Full Name of Team Leader	Augustin Crespin
Contact Email	augustin.crespin@uclouvain.be

Project Information

Project Title	Open Synthetic Simulation for Railway Defect Detection: Parametric Modeling and Texture Extraction from Real Imagery

Profile of the Team Leader(s) & Expected Team Composition

Augustin Crespin and Ioannis Kostis have co-authored an extended abstract presented at BNAIC/Benelearn 2025 [1], developed in the context of the Track Sentry Vision (TSV) project, a Walloon Region-funded initiative to build a low-cost, autonomous preventative maintenance system for railway sectors in developing countries, which forms the direct foundation for the work proposed here.

Augustin Crespin develops defect detection algorithms for the TSV project. He brings expertise in Artificial Intelligence, with a particular focus on Constraint Programming. His work is relevant for modeling and simulating complex systems under constraints, which is a core aspect of digital twins for the railway infrastructure. In the present project, his expertise can be used to model scenarios of defect progression, enforce structural and operational constraints, and search large spaces of possible system states in an interpretable manner.

Ioannis Kostis contributes expertise in Computer Vision, Generative AI, and Multimodal Learning. His work on Activity Recognition under challenging conditions, including settings where inputs are partial, occluded, or degraded [2], maps directly onto the inspection scenarios this project targets, where objects of interest may be partially visible or embedded in visually complex environments. In addition, his broader work on Stereo Vision (Novel View Synthesis, Neural Radiance Fields, Pose Estimation), as well as generative methods in general, positions him to contribute to the texture extraction pipeline and the generation of realistic synthetic data.

Have you already identified potential team members for your project?

Yes

List the team members you have identified and briefly describe their profiles/roles (e.g., expertise, affiliation, expected contribution).

Otmane Amel and Tanguy Vansnick from the UMONS ILIA (Informatique, Logiciel et Intelligence Artificielle) research center bring expertise in Visual Data Processing, Deep Learning-based feature extraction, and large-scale data management. Their work on Multimodal Learning in complex railway environments, illustrated by their paper on heterogeneous data fusion for dangerous action recognition in railway construction sites [3], and on occlusion-robust architectures [4] is directly relevant to the challenges of processing and structuring the large volumes of annotated data this simulation pipeline will produce.

In the context of this project, their contribution focuses on the data management layer: organising simulated defect scenarios, structuring outputs for efficient retrieval, and ensuring the pipeline remains compatible with the downstream requirements of detection model training.

Domain of Application Mobility & Logistics

Scientific Theme Physics-based Models & Digital Twins

Proposal Content

Abstract

The Track Sentry Vision (TSV) project, funded by the Walloon Region and carried out in collaboration with academic (UCLouvain, UMons) and industrial partners (Pepps, Cegelec, Vecturis, Future Resources), aims to develop a low-cost, autonomous system for detecting railway defects using onboard visual and gyroscopic sensors. In developing its detection pipeline, TSV exposed a fundamental bottleneck that this present project sets out to resolve: reliable automated detection depends on large, annotated datasets covering a wide range of failure types such as geometric deviations in track alignment, ballast degradation, vegetation contamination or settlement, sleeper damage including cracking or displacement, and rail surface faults including wear and tear, squats, shelling, and cracks, as well as the absence of required safety or signage elements. Such datasets are difficult and costly to obtain in real-world conditions. Some defect categories, rail surface faults in particular, are relatively well represented in the literature, while rarer failure modes, which are precisely the cases where reliable detection matters most, are often absent. Most notably, track geometry deviations remain, to the best of our knowledge, uncovered. Railway operators across Europe are reluctant to share inspection data given its operational sensitivity, and while simulation environments do exist, they tend to be proprietary and closed, leaving the research community with limited options for publicly available, high-quality training data across the full spectrum of failure types.

We propose building an open digital twin of railway infrastructure, constructed from open-source data and made freely accessible. Our proposed work has two main components. The first involves the modeling of accurate, parametric 3D track assemblies in a 3D creation software such as Blender, covering rail profiles, sleeper types and spacing, fastening systems, and ballast beds, with and without geometric defects such as lateral displacement or differential settlement. The second consists of a texture extraction pipeline that transforms images of real track components into 3D assets with physically accurate surface properties, which are then applied as coverings of the geometric models within the simulation environment. Together, these two components would allow the simulation to generate annotated video sequences of defective track across varied configurations, lighting conditions, and environments.

Since the goal is to train models capable of general defect detection, drawing from varied environments and infrastructure types is effectively an advantage. A model trained on data derived from diverse simulated conditions is better positioned for deployment across different countries and network configurations. The resulting framework offers the community a scalable, transparent resource for producing the training data that safety-critical diagnostic systems require.

Background Information & Problem Statement

The development of open-source automated defect detection systems for railway infrastructure is fundamentally constrained by the uneven coverage of defect types in publicly available datasets, and the inaccessibility of the inspection data that operators do hold. A substantial body of annotated imagery does exist, but it is heavily skewed toward rail surface defects: for example, squats, shellings, spallings, and head checkings are well documented, in part because they are visible, relatively common, and have attracted sustained research attention. More serious failure categories, such as track geometry deviations and sleeper degradation, are barely represented. To our knowledge, no publicly available dataset addresses geometry problems at all.

This absence has several causes. In well-maintained networks equipped with different infrastructure quality monitoring systems, such as vibration-based detection, Eddy Current Testing (ECT) or Ultrasonic Testing (UT), defects are typically identified and treated before they develop to a stage visible to the naked eye. There is simply little imagery to collect. Beyond that, railway operators are understandably protective of their inspection data, and the simulation tools that do exist are largely proprietary, making it difficult for the broader research community to work on these failure modes. Moreover, producing a convincing digital twin requires going beyond generic textures that fail to capture the physical nuances of real-world materials like metallic oxidation, concrete weathering, or the granular heterogeneity of ballast, as well as different lighting and environmental conditions. Without that level of detail, synthetic data loses the visual fidelity that makes it a viable training source.

Track geometry problems are thus a particularly compelling target for simulation: they are operationally significant, structurally tractable to model, and absent from current training pipelines. The proposed project aims to address this by building an open digital twin of railway track infrastructure, capable of representing a realistic range of defect types and conditions. The pipeline extracts surface material properties from real photographs of track components and uses them to produce geometrically and visually accurate 3D assets. The result is a flexible, publicly available framework for generating the kind of diverse, balanced synthetic data that current detection systems lack.

Project Objectives & Concrete Implementation

The primary objective for the two-week camp is to build a functional pipeline that takes real-world images of railway components and produces high-fidelity, rendered annotated sequences of defective tracks.

The pipeline has two interconnected parts, roughly corresponding to the two weeks of the event. The first is the construction of accurate 3D track models in a 3D creation software such as Blender. A railway track is an assembly of distinct components, each with its own geometry, material properties, and failure modes. At a minimum, a representative model must account for the rail profiles themselves, the sleepers and their spacing, the fastening systems that hold the rail to the sleeper, and the ballast bed that supports the whole structure. The construction of these assets will draw on publicly available technical specifications and open infrastructure datasets to ensure dimensional accuracy. Where possible, parametric modeling will be used so that track configurations can be varied systematically, covering different rail profiles, sleeper types (wooden, concrete, …) and forms (conventional or not, such as Y-shaped, twin ties, ...) and ballast grades. This parametric approach also means that geometric variations associated with degradation can be introduced in a controlled and reproducible way.

The second part is the texture extraction pipeline. Starting from real photographs of track components, a segmentation step using Vision Foundation Models such as SAM [5] isolates individual material regions: rail steel, concrete or wooden sleepers, ballast [6]. From these patches, the pipeline derives tileable Physically Based Rendering (PBR) texture maps that feed directly into Blender's shader system. Through this approach, the pipeline includes albedo (capturing the base color and tone of a material), roughness (controlling how light scatters across its surface), and surface normals (encoding fine geometric detail without increasing mesh complexity). Together they allow the simulation of the way light interacts with oxidized steel, weathered concrete, or loose stone in a way that a simple photograph mapped onto a flat surface cannot. Critically, because the maps are tileable, the same extracted material can be applied across track segments of arbitrary length without visible repetition artifacts. The pipeline will use railway datasets (see "Dataset" section) to identify infrastructure components and establish a material baseline. The result is a set of render-ready assets whose surface appearance faithfully reflects the actual material properties of field-deployed infrastructure.

With geometry and materials in place, the pipeline introduces defects into the scene. Rather than targeting a fixed set of anomalies, the framework is designed to be modular and extensible, allowing end-users to add failure types according to their needs. All outputs are produced with annotations, since the downstream goal is training detection algorithms that require labeled data. The final deliverable is a proof-of-concept that takes raw imagery as input and produces annotated, high-fidelity video sequences of defective track as output.

Do you plan to deliver, as an outcome of your project, a reusable "brick" for the TRAIL Factory (https://factory.trail.ac/en/home_page) that could later be transferred and converted into a company process?

Yes

Briefly describe what the brick would be and its intended users.

The primary deliverable is an open-source Python library that implements the image-to-texture conversion pipeline at the core of the project. Given a set of photographs of railway components, it segments individual material regions, extracts PBR texture maps for each, and outputs render-ready assets compatible with Blender. This library is the reusable core: it can be run independently of the rest of the pipeline, adopted by anyone working on railway simulation, and extended to new component types or material classes without modifying the underlying architecture.

The intended users are computer vision researchers working on infrastructure diagnostics, railway simulation engineers, and potentially infrastructure operators looking to generate training data for their own detection systems. The framework is specifically designed to lower the cost of entry for this kind of work: rather than building a simulation from scratch, users get a functional, documented pipeline they can adapt to their own track configurations, defect types, or environmental conditions. The open-source release is also a direct response to the closed-ecosystem problem described earlier. By making both the tool and the underlying data pipeline publicly available, the project creates a shared resource that the research community can build on rather than duplicate.

Project Dataset

The project draws on three open-access datasets to establish both the environmental context and material-specific defect morphology needed for high-fidelity simulation.

RailSem19 [7] (available at https://www.wilddash.cc/railsem19) provides the foundation for modeling healthy structural components and diverse operational environments. Developed by VITRO and the Austrian Institute of Technology, it comprises 8,500 annotated sequences captured from the ego-perspective of trains and trams across 38 countries, spanning all four seasons and a range of weather conditions including rain, snow, and mud. These sequences are used to extract textures for intact rail track components and background terrains under varied lighting conditions.

For material-level defect signatures, two high-resolution rail surface defect datasets were selected based on their high image quality [8]. The first (https://universe.roboflow.com/railway-defects-osolv/rail-surface-defects-flrty-mbj1y) contains 2,684 images across five classes: cracks, scars, breaks, rails, and lightbands. The second (https://universe.roboflow.com/ee-6wq9x/rail-uykcc) contains 4,110 images across six classes: clipping, lightbands, perlage, rails, seams, and sharp rails.

The Railway Track Surface Faults Dataset [9] (available at https://data.mendeley.com/datasets/8hxtgyyxrw/2), provided by NCRA MUET in Pakistan, further supports the simulation of material degradation. Captured at 120 frames per second with a focused 14-inch field of view on the railhead, it includes over 5,200 annotated images spanning seven defect classes: cracks, flakings, joints, shellings, spallings, squats, and grooves.

Together, these datasets supply macro-scale environmental conditions from RailSem19 and micro-scale defect signatures from the surface fault datasets, ensuring the resulting synthetic data captures both realistic infrastructure context and authentic failure morphology.

Detailed Work Plan

Week 1: 3D Track Modeling and Scene Infrastructure

The first week focuses on building the geometric foundation of the simulation in Blender (or equivalent).

Day 1: Setup and specification

The first day is dedicated to environment setup and a review of publicly available technical specifications for railway track standards. This covers standard rail profiles, sleeper dimensions and spacing, fastener types, and ballast grading. These specifications will allow the models to be built at correct dimensions, ensuring the resulting assets reflect real infrastructure rather than approximations. The parametric modeling strategy is also defined at this stage: which parameters will be variable, what ranges are realistic, and how the parameter space maps onto the diversity of track configurations the simulation should cover.

Days 2-3: Parametric asset construction

Construction of the core track assembly in Blender. Rail profiles, sleeper types and forms, fastening systems, and ballast beds are modeled as parametric components that can be combined and varied programmatically. By the end of day 3, the pipeline should be able to generate geometrically valid track configurations from a parameter set without manual intervention.

Day 4-5: Geometric defect modeling

With a clean baseline model in place, the focus shifts to introducing geometric degradations. Priority targets are track alignment deviations (longitudinal leveling/cross-leveling, lateral deviation, cant defect, spacing defect, ...), displacement of sleepers and differential settlement of the ballast bed, as these represent the most underrepresented failure category in existing datasets. Each defect type is implemented as a modular perturbation to baseline geometry. Camera paths, lighting rigs, and basic environmental variation (time of day, weather approximations) are also set up at this stage.

End of week 1 milestone:

A parametric pipeline capable of generating varied, geometrically accurate track configurations with and without geometric defects, renderable from configurable camera paths and lighting conditions.

Week 2: Texture Extraction, Integration, and Defect Rendering

The second week focuses on the texture extraction pipeline and its integration with the geometric models built in week 1.

Days 6-7: Segmentation and material extraction

The texture pipeline begins with automated segmentation of real track photographs using a Vision Foundation Model. Individual material regions, rail steel, concrete or wooden sleepers, ballast, are isolated and passed to the PBR extraction stage, which derives albedo, roughness, and surface normal maps for each. Outputs are validated visually for tileability and material plausibility. The pipeline will draw on available railway image datasets to build an initial material library covering the core component types.

Days 8-9: Blender integration and surface defect modeling

Extracted PBR maps are plugged into Blender's shader system and applied to the parametric models from week 1. This is where the two halves of the pipeline meet: geometric skeletons are dressed with physically accurate surface materials. Surface-level defect appearances, rail oxidation, concrete weathering, crack textures are introduced as material variations within the shader system. By the end of day 9, the pipeline should be able to produce a fully textured, defective track render from raw input imagery and a parameter configuration.

Day 10: End-to-end assembly, annotation, and packaging

The final day is dedicated to connecting the full pipeline into a coherent, documented workflow and verifying that all outputs carry correct annotations. Because defect positions, types, and parameters are known at generation time, annotations can be derived automatically from the simulation's ground truth. The Python library is packaged with a command-line interface and basic documentation. A short set of rendered sequences is produced as demonstration output.

End of week 2 milestone:

A functional, documented, open-source proof-of-concept that takes raw track imagery as input and produces annotated, high-fidelity video sequences of defective track as output.

Bibliographic References

[1] Crespin, A., Kostis IA., Schaus, P. (2025) Railway Track Defect Detection, The 37th Benelux Conference on Artificial Intelligence and the 34th Belgian Dutch Conference on Machine Learning

[2] Kostis, IA., Mathe, E., Spyrou, E., Mylonas, P. (2022) *Human Activity Recognition Under Partial Occlusion*. In: Iliadis, L., Jayne, C., Tefas, A., Pimenidis, E. (eds) Engineering Applications of Neural Networks. EANN 2022. Communications in Computer and Information Science, vol 1600. Springer ,Cham. https://doi.org/10.1007/978-3-031-08223-8_25

[3] Amel, O.,Siebert, X., Mahmoudi, SA. (2024) *Comparison Analysis of Multimodal Fusion for Dangerous Action Recognition in Railway Construction Sites*. *Electronics*; 13(12):2294. https://doi.org/10.3390/electronics13122294

[4] Ouardirhi, Z., Amel, O., Zbakh, M., Mahmoudi, SA. (2024). *FuDensityNet: Fusion-Based Density-Enhanced Network for Occlusion Handling*. 632-639.https://doi.org/10.5220/0012425400003660.

[5] Kirillov, A., Mintun, E., Ravi, N., Mao, H. et al. (2023) *Segment anything*, in Proceedings of the IEEE/CVF international conference on computer vision, pp. 4015–4026

[6] Wu, X., Zhang, R., Qin, J., Ma, S, Liu, C (2024) *Wps-sam: Towards weakly-supervised part segmentation with foundation models* in European Conference on Computer Vision, pp. 314–333, Springer.

[7] Zendel, O., Murschitz, M., Zeilinger, M., Steininger, D., Abbasi, S., and Beleznai,C. (2019). RailSem19: A dataset for semantic rail scene understanding. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops

[8] Dwyer, B., Nelson, J., Hansen, T., et al. (2026). Roboflow (Version 1.0) [Software]. Available from https://roboflow.com. Computer vision.

[9] Arain, A., Mehran, S., Shaikh, M. Z., Kumar, D., Chowdhry, B. S., and Hussain, T. (2024). Railway track surface faults dataset. Data in Brief, 52:110050.

Eligibility & Evaluation

Does the project include multidisciplinarity between STEM & SSH?

No

We confirm that the Team Leader will be present for the full duration of TReC'26 if the project is selected (August 24th - September 4th, 2026, Lausanne, Switzerland)

I/We agree and confirm