

TrainSim: A Railway Simulation Framework for LiDAR and Camera Dataset Generation

Gianluca D'Amico¹, Mauro Marinoni¹, *Member, IEEE*, Federico Nesti, Giulio Rossolini¹,
Giorgio Buttazzo², *Fellow, IEEE*, Salvatore Sabina, and Gianluigi Lauro

Abstract—The railway industry is investigating new ways to improve the safety and the performance of signalling functions (e.g., train localization) and automate other complex train functions, such as signal and sign recognition, obstacle detection, and track discrimination. Such tasks require the artificial perception of the railway environment through the data acquired from different types of sensors, including cameras, LiDARs, wheel encoders, GNSS receivers, and inertial measurement units. However, testing new algorithms and solutions that use such sensory data requires the availability of a large amount of labeled data, acquired in different scenarios and operating conditions, which are difficult to obtain in a real railway setting, due to strict regulations and practical constraints in accessing the trackside infrastructure and equipping a train with the required sensors. To cope with such difficulties, this paper presents a visual simulation framework able to generate realistic railway scenarios in a virtual environment and automatically produce a variety of labeled datasets from different types of emulated sensors, including cameras, LiDARs, and inertial measurement units. Such scenarios and datasets can be used for testing innovative algorithms, as well as for training and testing deep neural networks for a variety of tasks, as image segmentation, object detection, visual odometry, track discrimination, etc. The proposed framework is particularly relevant for the railway domain, considered the lack of similar datasets and the difficulty of reproducing critical situations in a real environment. A set of experimental results are reported to show the effectiveness of the proposed approach.

Index Terms—Railway simulator, dataset generation, LiDAR simulation, LiDAR modeling.

I. INTRODUCTION

RAILWAY stakeholders have started many initiatives to accelerate the migration of European railways towards a unique sustainable and safe railway system without frontiers, and increase the integration with the multi-modal European transportation system. To this end, they are working on solutions to increase the capacity of the lines, reduce the

setup phases of European Rail Traffic Management System (ERTMS) lines, and accelerate the integration of Automatic Train Operations over European Train Control System (ETCS), as demonstrated by the activities of the ERTMS User Group [1]. However, reaching these goals requires not only improving the performance of some railway functions, such as train odometry and train localization, but also automating other functions, such as track discrimination, which consists of identifying the track where the train is running with respect to the other tracks present on the line, track elements recognition, and obstacle detection.

All these functions require the execution of sophisticated perceptual tasks that process and integrate data from different heterogeneous sensors, including cameras and LiDARs. Given the high criticality of such railway functions, any solution has to follow a rigorous development process, analysis, and tests to show that the tolerable hazard rates associated with these functions and with the whole system are below given targets (e.g., 10^{-9} per hour and per function).

Unfortunately, carrying out an exhaustive field test campaign is impractical, due to the difficulty of replicating a large variety of critical scenarios in a real environment, especially those involving faults and degraded conditions. In addition, planning acquisition campaigns is complex and quite expensive for a number of reasons, such as overcoming the strict regulations and practical constraints in accessing the track-side infrastructure, equipping a train with the required sensors, and properly labeling the huge amount of data acquired.

Simulation is a crucial technique to speed up the development and test of algorithms for perception and control, because it has the great advantage of creating many types a scenarios, while automatically generating labeled datasets for different types of sensors. This is why many simulations tools have been developed in the last years for different application domains, as CARLA [2] for self-driving cars, AirSim [3] for unmanned aerial vehicles, Open AI Gym [4] for robot control, and so on.

The railway domain, however, lacks of a similar simulation tool, as well as public datasets properly labeled for training and testing deep neural models and vision-based algorithms for camera and LiDAR sensors. This is probably due to the fact that railway systems evolved more slowly with respect to other applications domains, given the strong regulations they are subject to for safety reasons that normally lead to the use of consolidated technologies and methodologies.

To fill such a gap, this paper proposes TrainSim, a simulation framework that can automatically generate synthetic datasets for training and testing neural networks in a variety of

Manuscript received 26 January 2023; revised 24 May 2023; accepted 30 June 2023. Date of publication 2 August 2023; date of current version 29 November 2023. This work was supported in part by the Hitachi Rail STS under Industrial Research Grant for Developing this Simulation Framework. The Associate Editor for this article was J. Felez. (*Corresponding author: Gianluca D'Amico.*)

Gianluca D'Amico, Mauro Marinoni, Federico Nesti, Giulio Rossolini, and Giorgio Buttazzo are with the Department of Excellence in Robotics and AI, Scuola Superiore Sant'Anna, 56124 Pisa, Italy (e-mail: gianluca.damico@santannapisa.it).

Salvatore Sabina is with the Department of Excellence in Robotics and AI, Scuola Superiore Sant'Anna, 56124 Pisa, Italy, and also with Hitachi Rail STS, 16151 Genoa, Italy.

Gianluigi Lauro is with Hitachi Rail STS, 16151 Genoa, Italy.

This article has supplementary downloadable material available at <https://doi.org/10.1109/TITS.2023.3297728>, provided by the authors.

Digital Object Identifier 10.1109/TITS.2023.3297728

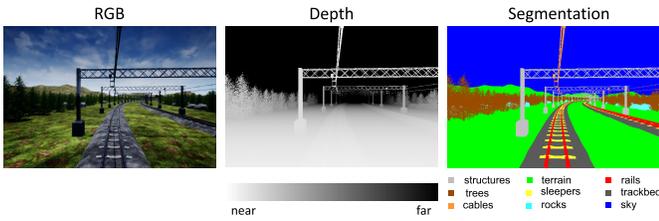


Fig. 1. Examples of images from the same scene produced by the simulator. Left: RGB camera; center: depth camera (each pixel value encodes the distance between the object represented by the pixel and the camera); right: segmented image (each pixel color encodes a different object class).

railway scenarios and operating conditions for verifying new solutions based on sensors as cameras, LIDAR, and inertial measurement units.

In particular, the paper provides the following contributions:

- 1) It presents a highly configurable and extendable environment generator to create a wide range of realistic railway scenarios controlled by a set of user-defined parameters.
- 2) It proposes a method for generating arbitrarily large labeled datasets from such virtual environments using a set of simulated sensors (LiDARs, cameras, and IMUs) that can produce data similar to their physical counterparts.
- 3) It provides a method for exporting the obtained datasets in a standard format for training deep neural networks or streaming them to ROS [5] for real-time visualization.

At present, TrainSim, can generate the following datasets:

- RGB images: taken from one or multiple cameras placed on the train, in positions specified by the user.
- Depth images: taken from one or multiple cameras placed in user-defined positions. They can be used as ground truth data for depth estimation algorithms.
- Segmented images: taken from one or more cameras placed in the same position of the RGB or Depth cameras, where each pixel value encodes the class of the object corresponding to the pixel. They are used as labels for semantic segmentation and other tasks.
- Point Clouds: taken from a LiDAR sensor placed on the locomotive in a position specified by the user. A point cloud includes a set of 3D points acquired according to the scanning pattern of the user-defined LiDAR configuration.
- Segmented Point Clouds: taken from the same LiDAR. Each point is associated with a label that identifies the object hit by the LiDAR beam. They are used as ground truth data for the point cloud segmentation task.
- IMU dataset: 6-axes accelerometric and gyroscopic data computed with user-defined IMU models from the ground-truth acceleration and angular velocity data.

In addition, the tool provides the ground truth for the vehicle position, attitude, speed, and acceleration data, computed according to the kinematics of the train trajectory.

Figure 1 shows an example of images produced by the RGB and depth cameras, along with the segmented image from the same simulated scene.

Figure 2 shows two point clouds of the same scene produced by the virtual LiDAR: in Figure 2a each color encodes a

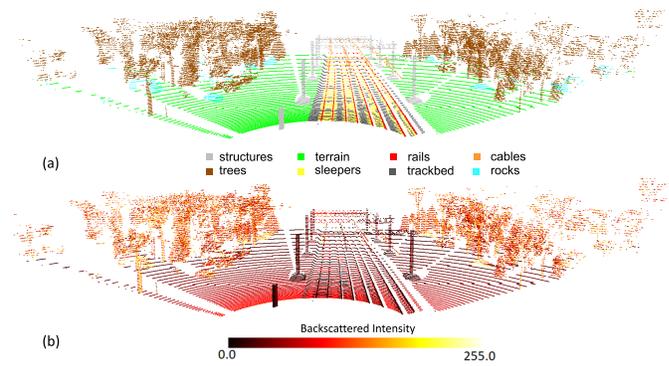


Fig. 2. (a) Example of a segmented point cloud generated by TrainSim, where each color encodes a different object class; (b) the same point cloud where each color encodes the back-scattered intensity value computed by the LiDAR model.

different object class, whereas in Figure 2b the color encodes the back-scattered intensity.

To summarize, the TrainSim framework aims at providing datasets and ground truth data for the following tasks: Visual Odometry, LiDAR Odometry, Image Segmentation, Point Cloud Segmentation, Image Depth Estimation, and Inertial navigation. The generation of datasets for tasks like 2D and 3D object detection is in progress and will be part of future work.

The rest of the paper is organized as follows. Section II discusses the related works; Section III presents TrainSim; Section IV reports some experimental results; and Section V states the conclusions and future work.

II. RELATED WORKS

The design of proper datasets for training and testing purposes is crucial for developing and verifying effective perception algorithms. The tools developed for the automotive domain typically use benchmarks that provide several visual frames captured in different environments, such as the KITTI benchmark [6] and its semantic segmentation variant [7], or the Cityscapes dataset [8]. Most of such datasets are focused on urban scenarios, and the vast amount of images required for training is typically obtained by data augmentation [9], mixing real and virtual images. The lack of open datasets in the railway domain represents a severe obstacle to testing and verifying novel algorithms. Zendel et al. [10] pointed out that, excluding the thousands of labeled images taken from street-view or spectator-view, image datasets of railway environments taken directly from the train are nearly nonexistent. Many solutions presented in the literature for the railway domain are tested and verified on private datasets that only include a few hundred data samples for camera and LiDAR frames, as declared by the authors [10], [11], [12].

Simulators offer the possibility to test perception and control algorithms in a variety of situations that would be hard to reproduce in the real world. For this reason, several synthetic generation tools have been presented in the last years to overcome the lack of real datasets, as CARLA [2] for automotive simulation and *AirSim* [3] for unmanned aerial vehicles (UAV), both based on the Unreal Engine 4 (UE4) [13] graphic engine.

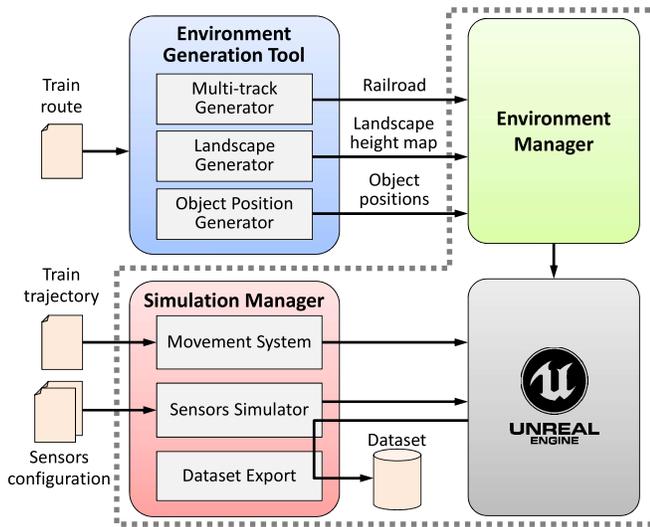


Fig. 3. Architecture of TrainSim (the gray dashed line indicates the modules integrated in the UE4 Editor). More details are reported in Sections II and III of the supplementary material.

Another tool is *AutonoVi-Sim* [14], which supplies LiDAR frames gathered into a virtual world.

Unfortunately, most of the existing simulators have been developed for self-driving cars and drones, and there is a lack of tools for the railway domain that support the integration of the LiDAR, Camera, IMU and GNSS technologies. This work presents *TrainSim*, a train simulation framework for generating realistic datasets of images, point clouds, and inertial data to test and validate novel algorithms for tasks such as inertial navigation, object detection, and semantic segmentation in the railway domain. In particular, the camera model is naturally derived from the graphic engine frame, producing RGB, semantic, and depth images directly from the graphic environment of UE4. On the other hand, the emulation of the LiDAR sensor exploits the ray-casting system of UE4, which allows the detection of objects between two endpoints, making the LiDAR emulation straightforward. Unlike *CARLA* and *AirSim*, the proposed approach also generates the backscattered intensity of the LiDAR sensor by exploiting a simplified version of the Lambertian-Beckmann model [15] that describes how different surfaces reflect light rays. More details on the images and point clouds generation models are described in Section III-F and Section III-E, respectively.

III. SIMULATION FRAMEWORK

The architecture of *TrainSim* is depicted in Figure 3 and is composed of three main modules: the *Environment Generation Tool* (EGT), which manages the creation of the rail-track surrounding environment, the *Environment Manager* (EM), which instantiates the created environment in Unreal Engine, and the *Simulation Manager* (SM), which simulates the train movement, emulates the sensors working principles, and generates the various datasets.

The environment generation is based on the *GeoGen* project of Matěj Záborský [16], which is a tool for creating realistic terrains with desired height maps. The virtual environment is generated starting from the train route specified in a file, which contains set points that are either sampled from a real

trajectory or synthetically generated by a separate trajectory generator, described in Section III-H.

In the following, we refer to a *track* as the physical structure (a pair of rails) where a train can run, and to the *railway environment* as all the ensemble of tracks placed in the environment. A track is defined as a sequence of 3D waypoints (referred to the track centerline), called *track points* and is divided into *blocks*, where each block identifies a specific type of railway structure, namely a straight line, a curve, a station, a tunnel, or a bridge. The *main track*, also called the *route*, is the one traveled by the train, whereas the remaining tracks are referred to as *auxiliary*. Then, a *trajectory* refers to a specific train journey on a route (i.e., the sequence of positions, velocities, and accelerations of rear and front bogies of the front vehicle sampled at a given frequency). The *Environment Generation Tool* is responsible for managing the creation of the railway environment and consists of three main modules:

- *Multi-track generator*. It creates a number of *auxiliary tracks* that run parallel to the main track, but can also join it or depart from it with different given rules. It produces a *Railroad.json* file that contains the list of tracks (the main and the auxiliary ones). Each track is divided into blocks (e.g., straight, curve, bridge, etc.) and it is associated with a 3D point sequence and other information needed in the generation of the virtual environment. Refers to Section III-B for further details.
- *Landscape Generator*. It creates the area surrounding the tracks, including the terrain and the mountains. It produces a height-map (i.e., a grid of vertices) in which each vertex is associated with a defined height that derives from the elevation of the tracks. Section III-C describes the height-map generation in more detail.
- *Object Position Generator*. It generates random spawn points near the tracks, where different types of objects can be placed (e.g., trees, rocks, buildings). Thanks to its modularity, the placement algorithm can easily be extended to add other types of objects to the scene. In particular, positions are selected taking object size into account to avoid reciprocal overlapping.

The outputs of the modules are sent to the *Environment Manager*, responsible for creating the virtual environment within Unreal Engine by placing the landscapes, the environmental objects, and the railway building structures into the simulated world. It exploits the 3D object models (meshes, materials and textures) of the *TrainTemplate* plugin [17], which provides high-fidelity models for railways objects, vehicles, stations, tunnels, and bridges. The meshes for other objects (e.g., trees and rocks) are randomly chosen from a set of different meshes to diversify the simulated environment, both for images and point clouds. Furthermore, ballast and landscape materials can be randomly drawn at the start of each simulation, avoiding reusing the same texture in each generated dataset.

The *Simulation Manager* takes as input the train trajectory and a set of files (described in Section III-A) containing the configuration parameters needed to emulate the working principles of specific sensors. It tightly interacts with Unreal Engine sending the sequence of train positions and receiving

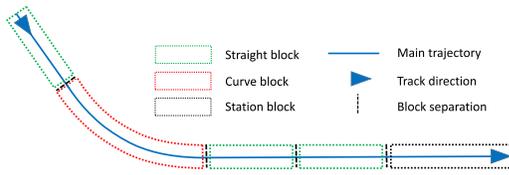


Fig. 4. Example of a railroad section divided into blocks.

the data produced by the simulated sensors. It includes three main modules:

- *Movement System*. It manages the train movement, advancing the train in each point of the specified trajectory.
- *Sensors Simulator*. It consists of a set of blocks, each responsible for emulating a specific sensor.
- *Dataset Export*. It exports the generated dataset saving it on the disk. The generated datasets can also be transferred in real-time to a ROS Bridge application for visualization or test purposes by the ROS communication system.

The following sections describe the details of the main architecture components, whereas more details about the remaining modules (e.g., Object Position Generator, Environment Manager) are described in the supplementary material.

A. Input Files

The train route is specified in a file as a sequence of 3D points $\mathcal{P} = \{P_k | k = 1, \dots, N\}$ in local north-east-down (NED) coordinates [18]. This sequence is used to generate the corresponding main track, which is divided into construction blocks of different types (e.g., straight, curve, station, bridge, etc.). Each block type has specific characteristics that constrain the construction of the relative way-point sequence and the velocity profile (e.g., curve blocks have a minimum curvature radius and constrain the maximum travelling speed of the train). An example of a railroad section divided into blocks is illustrated in Figure 4.

The train trajectory file specifies the position, the velocity, and the acceleration of both front and rear of the vehicle at each timestamp. The trajectory can either be sampled from real IMU and GNSS sensors, or it can be synthetically generated by a proper tool, briefly described in Section III-H.

Each configuration file provides information on a specific sensor, describing its type, features, parameters, and noise models. These data allow the Sensors Simulator to produce a realistic output by applying the noise models to the data acquired in UE4.

B. Multi-Tracks Generation

The Multi-track generator randomly creates a number of additional tracks next to the main track to populate the railway environment. This can be useful to test the performance of track discrimination algorithms. The user can also decide to duplicate the point sequence of the main track to have double track instead of a single one. The duplicated track is generated to the right of the main track, since the train hand of drive is on the left at a fixed inter-track distance defined by the user.

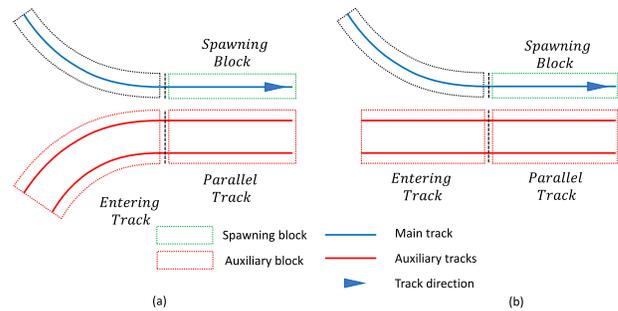


Fig. 5. Entering part examples with 2 parallel tracks, a single line (blue or red) represents a single track): (a) the auxiliary track is generated in correspondences of a straight block, hence it is composed of a straight and a curve blocks; (b) the auxiliary track is generated in correspondence of a curve block, hence it can be composed of a straight block only.

To generate additional tracks, the module parses the list of railroad blocks of the main track to decide where to begin or end auxiliary tracks, following the constraints imposed by each track block. Figure 4 shows an example of a railroad blocks division, given as input to the Multi-tracks generator. An auxiliary track has three parts: an entering part, a parallel part, and an outgoing part. The entering part is composed of a straight dead-end block, and a curve that joins it to block parallel to the main track. If the railroad block is a curve, the entering part can be composed only of a straight block, as depicted in Figure 5.

Some of the building rules are derived from the railway standards [19], [20], such as the inter-track distance or the minimum curve radius, whereas others need to be user-defined, like the number of auxiliary tracks. The creation of auxiliary tracks follows pseudo-random decisions based on a user-defined probability. In this way, the user can manage the auxiliary tracks generation, creating different scenarios from the same train route.

C. Landscape Generation

Once the auxiliary tracks have been generated, the landscape generation module creates the terrain of the virtual environment, containing the ground, mountains, and valleys.

In UE4, a landscape is defined from a height map, which is a matrix M of vertices, referred to as *main map*, in which each vertex $v_{i,j}$ has its own height value $M_{i,j}$. The main map has a rectangular size that includes all track points. The final height map is produced by GeoGen [16], a library that allows manipulating height maps using different operations, such as noising, scalar multiplication, and composition.

The landscape is partitioned into three different sections with respect to its distance from the outermost tracks, as illustrated in Figure 6. In particular, if the minimum distance d of vertex $v_{i,j}$ to the track is less than or equal to d_{near} , its height is set equal to the one of the most immediate track point. If the distance to the track is greater than d_{far} , the height is sampled from a noise function $N_{i,j}$, using the method proposed in [16]. Finally, if d is between the two thresholds, the assigned height grows linearly between the two values according to the distance function $f(d)$ illustrated in Figure 7.

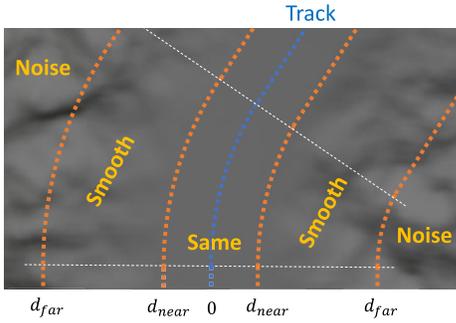


Fig. 6. Example of UE4 generated landscape showing how the terrain surrounding the track is partitioned in three areas: the area denoted as *Same* has a height equal to the closest track point; the area denoted as *Noise* has a noisy height; while the area in the middle (*Smooth*) smoothly changes the height between the two values.

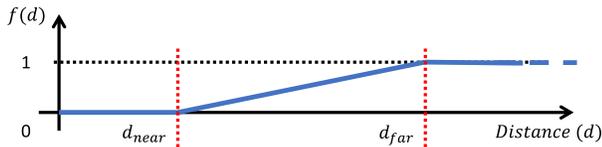


Fig. 7. Distance function used to set the height of the vertices located in the smooth region at a distance from the track between d_{near} and d_{far} .

The two distance bounds d_{near} and d_{far} can be set by the user. Please note that d_{near} has a minimum value imposed by the railway construction standards [19], [20], namely 1.5 m. This solution is needed to safely place a number of environmental objects (e.g., trees) around the railway structure.

More specifically, given a vertex $v_{i,j}$ in the main map, let $P_n(v_{i,j})$ be the track point with the minimum distance to $v_{i,j}$ and let $d_{i,j}$ be such a distance. For the sake of clarity, the point coordinates are expressed in the East-North-Up (ENU) reference system, and P_n^U represents the Up component (i.e., the height). Then, the height value $M_{i,j}$ associated with vertex $v_{i,j}$ is computed as

$$M_{i,j} = P_n^U(v_{i,j}) * [1 - f(d_{i,j})] + N_{i,j} * f(d_{i,j}). \quad (1)$$

Furthermore, the sub-module generates a valley patch (i.e., a sub-matrix of vertices) that is superimposed to the main map every time a railway bridge is present in the trajectory, so allowing to possibly create a river in that specific position. In station blocks, the value of d_{near} is increased to accommodate buildings and other structures.

Once the whole main map is generated, it is divided into sub-maps, so that the user can save memory by storing only those sub-matrices near the track, as shown in Figure 8. In the proposed implementation, a sub-matrix has 1009×1009 vertices. A scale factor on the E-N axes equal to 100 is required since the UE4 measurement unit is expressed in *cm*.

With this setting, the distance between two vertices in the horizontal and vertical direction results to be 1 m.

D. Movement System

The movement system is responsible for updating the position of the train along the route, following the train trajectory specified in the corresponding input file. Each trajectory point

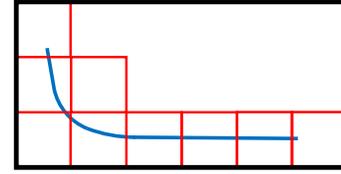


Fig. 8. Example of sub-matrices (red squares) that can be selected based on their distance from the track (blue line).

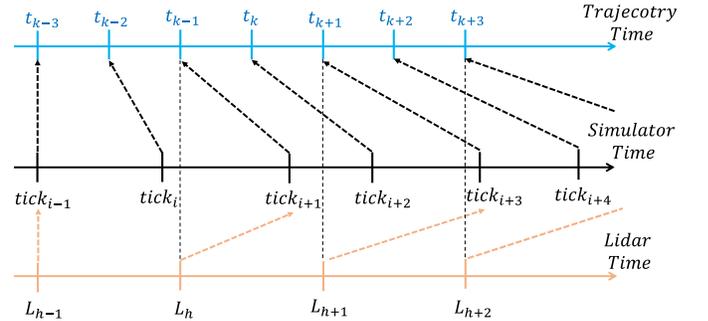


Fig. 9. Timelines corresponding to the train trajectory, the simulator frames, and LiDAR acquisitions. Black dashed arrows show the timestamps associated with the graphic frames, while brown dashed arrows show the graphic frames associated with the LiDAR acquisitions.

includes the position, speed, and acceleration of the front and rear bogies, as well as the corresponding timestamp.

To reproduce the train motion according to the given trajectory, the positions of the bogies have to be computed at each frame by interpolating the position of the two consecutive points in the trajectory that are before and after the tick absolute time. This solution, however, gives rise to two distinct problems:

- 1) The interpolation introduces an error on the virtual position of the train that increases with the train speed (the higher the speed, the higher the distance between trajectory points).
- 2) If graphic frames are visualised at a time that is different from the timestamps associated with the trajectory points, then the data produced by virtual IMU are not synchronized with those produced by visual sensors (cameras and LiDARs), hence they are not consistent.

To address these problems, the *real time* associated with the train trajectory has been decoupled from the *simulated time* at which UE4 produces a visual frame. While the difference between the time stamps of any two consecutive trajectory points is constant and equal to the sampling period $T_S = t_k - t_{k-1}$, the time elapsed between two consecutive frames can vary depending on the machine running the graphic engine. Hence, each frame produced by the graphic engine must be associated with a trajectory point and its corresponding timestamp, so ignoring the simulation time. Figure 9 compares the timelines associated with the trajectory, the simulator frames, and the acquisitions from a LiDAR sensor, visualizing the time stamps associated with each frame.

In the example shown in Figure 9, the LiDAR is acquired with a period that is twice the one used for the trajectory. Black dashed arrows show the timestamps associated with the

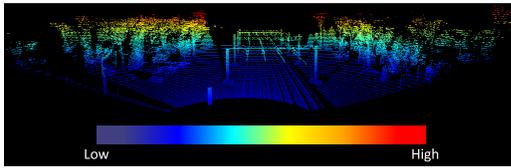


Fig. 10. Example of a point cloud captured from the simulated environment, where each point is colored according to the height value of the point itself.

graphic frames, while brown dashed arrows show the graphic frames associated with the LiDAR acquisitions. In the current implementation, the acquisition period of a visual sensor must be a multiple of the sampling time of the trajectory points.

E. LiDAR Sensor Model

Light Detection And Ranging (LiDAR) sensors are active devices that emit light rays and compute the time needed for the rays to be backscattered to the sensor receivers, or the phase change of the backscattered ray. If the emitted rays are backscattered, the distance between the LiDAR and the object hit by the ray can be computed from the travelling time and the speed of light, or from the phase difference between the emitted ray and the backscattered one. Namely, a LiDAR sensor emits several laser beams (or a flash light) and uses a matrix of receivers to create a depth map of the surrounding environment, referred to as a *point cloud*.

Common LiDARs emit a vertical array of laser beams that rotates around the vertical axis to acquire the surrounding scene. The angular inclination of each laser beam defines the vertical resolution of the sensor, whereas the horizontal angular step at which the rays are emitted defines the horizontal resolution. The `json` input file for a LiDAR specifies the horizontal resolution, the horizontal and vertical field of view (FOV), the number of beams (from which it is possible to derive the vertical resolution knowing the vertical FOV), the range of the laser beams, and the frame rate.

The LiDAR working principle is emulated by exploiting the ray tracing system of UE4. In particular, the `SingleLineTraceByChannel` function of UE4 generates a ray from a starting point to an ending point given as inputs. If there are objects along the ray, this function returns the closest 3D point in which the ray intersects the first object surface. It also returns a reference to the object hit, from which it is possible to retrieve other object features stored in the system. Hence, the relative position of the hit point is computed by subtracting the absolute position of the starting point. To make the data more realistic, the distance associated with each LiDAR beam is perturbed by adding a Gaussian noise with zero mean and user-defined variance.

The sensor is implemented as a UE4 actor component positioned on the front vehicle of the train in a user-defined location. A point cloud is produced with a user-specified period that must be a multiple of the period at which the trajectory points are generated. Figure 10 depicts a sample point cloud captured in the simulated environment.

Note that, in the real world, if the LiDAR is moving, the 3D points belonging to a full scan refer to different LiDAR positions. To compensate for such a mis-alignment

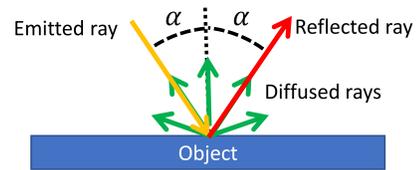


Fig. 11. Object (blue rectangle) response to an emitted light ray (yellow arrow). Part of the light energy is diffused in all directions (green arrows), while another part (red arrow) is reflected in the opposite direction of the emitted ray.

due to motion, most of LiDAR sensors incorporate an IMU that automatically compensates for the LiDAR movements during a scan. For this reason, TrainSim does not take this phenomenon into account and acquires the point cloud in a single graphic frame, also providing velocity and position information for considering motion distortion effects by a post-processing algorithm. For high-speed trains, flash LiDARs are more suitable than laser-scanner LiDARs, because they are less prone to motion blur effects. Note that each object in the virtual environment is associated with an identifier that specifies its class and its instance (e.g., `Rock_0`). Such an identifier is associated with each 3D point to create labeled datasets for semantic segmentation.

Real LiDAR sensors also provide the intensity of each backscattered ray, in terms of light energy, which can be used to discern objects in the environment. As shown in Figure 11, a fraction of the incident ray is reflected in the opposite direction (red arrow) with an angle equal to the incidence angle, but opposite with respect to the normal to the object surface, whereas another fraction is diffused in all directions (green arrows). The backscattered ray detected by the LiDAR is the diffused ray reflected back in the same direction of the incident ray.

As proposed by Tian et al. [15], we used the Labertian-Beckmann model to compute the backscattered intensity as a function of three factors:

- 1) The distance between the sensor and the object;
- 2) The incident angle between the emitted ray and the normal to the object surface;
- 3) The material of the object.

In particular, each object is associated with a diffusive and reflective coefficient and a maximum incidence angle, over which the object results to be completely diffusive. The model can also be extended to consider other parameters of the environment, as air density and humidity, which, at present, are not taken into account.

To reproduce the backscattering effect in the LiDAR simulation, the incident angle and the material of the object are needed. The object reference gathered by the `SingleLineTraceByChannel` function provides the normal to the object surface, which is used to compute the ray incident angle. Furthermore, a mapping between each object class in the virtual environment and the material parameters needed in the Labertian-Beckmann model is defined and exploited to compute the response of each single object to the LiDAR rays. In this work, the parameters of different material are taken from the study presented by Tian et al. [15].



Fig. 12. Examples of RGB images generated by TrainSim at different daytimes: left, morning; middle, evening; right, night.

Figure 2a shows an example of a point cloud where each point is colored with the class of the corresponding object, whereas Figure 2b shows the same point cloud where each color encodes the backscattered intensity value normalized as an integer in the range $[0, 255]$.

F. Camera Sensor Model

Unreal Engine 4 allows the user to create cameras to gather images of the virtual environment from user defined-locations. In particular, a camera can be placed in front-top of the locomotive to capture an image at each tick of the graphic engine. As for the LiDAR sensor, the camera capture period can be specified as a multiple of the one used for the trajectory. The user can also define a number of parameters of the camera, such as shutter speed, aperture, ISO, and resolution. It is worth to notice that UE4 allows to set post processing parameters to handle motion effects, such as blur effects, and shutter speed. Such effects can then be filtered out, exploiting well-known computer vision algorithms. Additionally, UE4 allows generating a depth image of the scene, where each pixel value encodes the distance between the camera and the object represented by the pixel. The distance is normalized into a range $[0, depth_{max}]$, where all the values above $depth_{max}$ are cut off and set equal to $depth_{max}$. The value of $depth_{max}$ is set to 100 m by default and can be redefined by the user.

UE4 also allows defining post-processing routines that exploit custom stencils to create a segmented version of an RGB image. Each type of object placed in the virtual environment is assigned a specific custom stencil value used to distinguish each object in a segmented image. Figure 1 shows an RGB image (left image) with the corresponding depth image (middle image) taken from the same scene, and along with the corresponding segmented one (right image), where each object class is identified by a different color.

At last, TrainSim allows defining different ambient aspects and weather conditions used to test visual based-algorithms in a wide range of operating conditions. In particular, it is possible to define:

- the Sun position, for creating images with different shadows and light intensity. The framework defines three different time slots, morning, evening, and night, as shown in Figure 12;
- the fog, with a desired intensity, by inserting the *ExponentialHeightFog* UE4 actor in the environment.

G. IMU Model

The proposed simulation framework includes a model of a 9-axis inertial measurement unit (IMU) with accelerometers,

gyroscopes, and magnetometers, for estimating the current position, velocity, and orientation of the train by means of inertial navigation algorithms [18]. To reproduce realistic data with high fidelity, the IMU model allows the user to specify noise properties, calibrated bias, and other parameters that affect the quality of the measures.

The simulated measured quantity \tilde{a} in the IMU reference frame is computed from the ground-truth quantity a in the NED frame by means of the accelerometer model \mathcal{A} to obtain $\tilde{a} = \mathcal{A}(a, \theta)$, where θ is the orientation of the IMU, necessary to return readings in the IMU frame.

Function \mathcal{A} depends on the following factors:

- 1) The gravitational acceleration g , added along the Down component of a and converted to the IMU frame using the rotation matrix $C_{NED}^{IMU}(\theta)$ computed from the orientation θ .
- 2) The misalignment matrix Mis (due to geometrical imperfections of the orientation of the individual accelerometer axes) and a constant calibrated bias ϵ , used to alter the ground-truth acceleration.
- 3) A drift term δ , which depends on noise parameters, such as Bias Instability, Noise Density, Random Walk, and environmental causes, as temperature-induced bias.
- 4) A quantization factor Q , used to replicate the resolution of the sensor.

The resulting function for the accelerometer is then:

$$\tilde{a} = \mathcal{A}(a, \theta) = Q \left(Mis C_{NED}^{IMU}(\theta)(a + g) + \epsilon + \delta \right). \quad (2)$$

Similar formulations are used for simulating the outputs of gyroscopes and magnetometers. Changing the noise, bias, resolution, or any other parameter in the datasheet of a specific sensor allows simulating different IMUs.

H. Trajectory Generator

This tool generates pseudo-random train routes and related journey trajectories to be used as input files for the simulator.

As a first step, the route is generated as a sequence of curves and straight blocks that follow the constraints imposed by the construction standards. The output is a sequence of spatially evenly distributed points $\mathcal{P} = \{P_k | k = 1, \dots, N\}$. Then, among the straight blocks, some of them are randomly selected as bridges, tunnels, and stations, using constraints and probability provided by the user. It is possible to interpolate the position between points by fitting three smoothing splines [21] on such points \mathcal{P} , one for each axis (north, east, down). This approach allows importing the set \mathcal{P} from the digital map of an actual route.

The second step defines the maximum train velocity on each block of the track, e.g., inside tunnels, on bridges, within stations, and in each curve as a function of its curvature. The details on the generation of such synthetic routes and velocity profiles are omitted for space limitations, also considering that the algorithm could be extended to better adapt to the final dataset produced by the simulator.

Then, the tool exploits the kinematic model and the control law of the train, combined with the geometry of the route and the maximum velocity profile, to yield the final trajectory of

the front and rear bogies of the vehicle (position, velocity, acceleration, and orientation) with the desired sampling time T_S . Note that the orientation at each sampling instant is computed only using the geometry of the track. In fact, since the motion of the train is heavily constrained by the track, the only allowed orientation of a vehicle is the orientation of the track itself. Hence, the yaw, pitch, and roll angles are computed from the line tangent to the track in the current bogie position. At the same time, the angular velocity is obtained by kinematics as a function of the orientation and orientation rates on the three axes.

The trajectory of a generic point on a vehicle can be computed from the trajectories of its front and rear bogies.

I. Dataset Export

The generated dataset can be transmitted for online usage or stored to be employed offline.

The proposed online method allows the user to directly connect the UE4 simulator to a ROS network, creating a sensor node that exposes the frame data right after the acquisition, providing a simulation system that can be tested and evaluated online. The ROSIntegration plugin [22] for UE4 is used to create distinct topics for images, point clouds, and inertial data that are transmitted through a TCP connection to a ROS [5] bridge node.

Datasets are saved on the disk with the same data format used by other urban open datasets that can be found in the scientific community, such as the KITTI dataset [6]. In this way, most of the automotive algorithms can be tested on the saved train dataset to evaluate their performance in a railway environment without additional pre-processing.

Along with the (x, y, z) coordinate values for each point in the LiDAR frames, TrainSim also provides the surface normal related to each object hit by the simulated LiDAR laser (the relative object which each point belong to). It is worth to remarking that, even so a LiDAR sensor does not provide such information, and that it must be computed with proper computer vision algorithms, it is an important local feature exploited by different point cloud processing methods.

IV. EXPERIMENTAL RESULTS

This section presents some experimental results aimed at testing the realism of the simulated datasets. Section IV-A compares a real point cloud gathered in a static environment with a point cloud generated by TrainSim on a similar static scene re-created on the graphic engine. Section IV-B compares the performance of a state-of-the-art LiDAR odometry algorithm on a sequence LiDAR frames generated by TrainSim and taken from the KITTI dataset. Finally, Section IV-C compares the results of an image semantic segmentation algorithm applied to both the TrainSim generated data and the RailSem19 [10] dataset.

A. LiDAR Working Principle Analysis

This section aims at evaluating the emulation of the working principles of a LiDAR sensor in TrainSim, considering both distance and backscattered intensity measurements. Real point



Fig. 13. Real-world reference static scene (left) and similar scene re-created in the simulation framework (right).

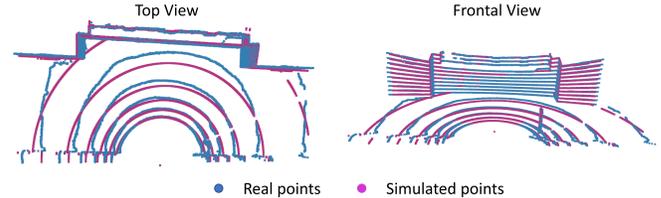


Fig. 14. Top and frontal view of real and simulated point clouds.

clouds were acquired using a Scout Mini Robot¹ equipped with a Velodyne VLP-16² LiDAR sensor. As illustrated in Figure 13, the robot was positioned on an courtyard in front of a wall (Figure 13, left) and a similar scenario has been re-created in TrainSim (Figure 13, right).

The VLP-16 is a 360° rotating LiDAR with 16 vertically aligned laser beams covering a vertical Field Of View (FoV) of 30°, vertical resolution of 2°, and horizontal resolution of 0.2° for the default rotation speed. In the following, θ denotes the yaw orientation angle, while ϕ_i and ρ_i denote the vertical angle displacement and the distance reading of the i^{th} beam, respectively.

The Root Mean Square Error (RMSE) between the real and the simulated cloud points was computed to evaluate the realism of the LiDAR simulation, as done in [23]. To be more consistent with some restrictions in the reconstructed environment, the point clouds have been cropped to reduce the horizontal FoV, setting $\theta \in [-\frac{\pi}{2}, \frac{\pi}{2}]$. It is worth noting that the simulated virtual scene has been re-constructed manually, introducing position errors due to measurement errors and shape misalignment imprecisions, which increased the resulting RMSE. Figure 14 illustrates the top and frontal view of the two point clouds (real data are presented in blue and simulated ones in purple). Note that the largest misalignment between points is due to the irregular shape of the real sidewalk, which is simply represented by a flat surface in the simulation. Excluding the sidewalk points from the comparison, the RMSE resulted of 0.035 m, which is in accordance with the precision of the VLP-16 LiDAR.

The presented results were obtained with a simulated point cloud without considering the noise, since the datasheet of the LiDAR device reports the precision only for the distance ρ and not for the ray angular displacements. Adding to the measured distance ρ a Gaussian white noise comparable with the VLP-16 precision (i.e., zero mean and variance 0.015) did not significantly change the RMSE, which resulted of 0.081 m

¹<https://global.agilex.ai/products/scout-mini>

²<https://velodynelidar.com/products/puck/>

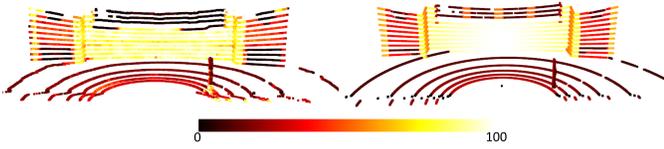


Fig. 15. Frontal view of the real (left) and simulated (right) point clouds. The color associated with each point encodes the backscattered intensity values.

considering the whole point cloud and 0.04 m excluding the points belonging to the sidewalk.

Concerning the backscattered intensity, there is not a standard way to process data, thus LiDAR manufacturers use different methods to compensate the measurements with respect to various parameters, as distance and incidence angle (see Section III-E). Such compensation methods are frequently unknown, making it challenging to precisely reproduce the output backscattered intensity values for a specific LiDAR device. In particular, VLP-16 divides the intensity values into two sub-ranges: values in $[0, 100]$ map diffuse reflectors with a reflectance in the range 0–100%, while values in $[101, 255]$ represent retro reflectors with an ideal reflection. Unfortunately, the calibration mechanism is not precisely described in the VLP-16 User Manual, restricting the possibility of reproducing an exact representation of the backscattered intensity in TrainSim. For this reason, a qualitative comparison between the intensity values is presented, showing the distribution of the intensities with respect to incidence angles, distances, object materials, roughness, and reflectance. The model to account for the last three parameters has been taken from [15]. In the reference scene, retro-reflected elements are not present, and the intensity values are scaled in the $[0, 100]$ range to match the VLP-16 specifications. Figure 15 shows the point clouds gathered from the reference scene and the simulated one, showing three different effects that can be underlined:

- the backscattered values of the sidewalk in front of the LiDAR decrease by increasing the angle of incidence;
- the metal pole has high-intensity values in the frontal part of the pole, rapidly decreasing on the pole boundaries;
- the effect of the angle of incidence on concrete material such as the wall is lower than the effect on metallic or plastic material.

B. LiDAR Odometry Analysis

This experiment compares the results obtained with simulated point clouds against real ones on an odometry task. Due to the lack of public point cloud datasets acquired from a train, the KITTI [6] urban automotive dataset was selected for comparison.

The purpose of LiDAR odometry is to predict the motion of the LiDAR sensor from consecutive LiDAR frames. The ego-motion estimation is done by iteratively computing the homogeneous transformation matrix T_k between two consecutive frames F_k and F_{k+1} that maximize the alignment between the two frames. Formally, the transformation matrix is defined as $T_k = \begin{bmatrix} R_k & t_k \\ \bar{0} & 1 \end{bmatrix}$, where R_k is a rotation matrix, t_k is a translation vector, and $\bar{0}$ is a vector of zeros. The best

alignment can be defined as an optimization process aimed at minimizing the following distance function:

$$d_m(T_k) = \sum_{i=1}^{N_k} ((R_k \cdot p_i + t_k - q_i)), \quad (3)$$

where N_k is the number of points in frame F_k , $p_i \in F_k$ is a point in frame F_k , and $q_i \in F_{k+1}$ is the point closest to p_i after applying transformation T_k to F_k . From the estimated transformation T_{k+1} , computed at time $k+1$, it is possible to predict the ego-motion of the LiDAR sensor in terms of orientation R_k and translation t_k .

In this work, the LiDAR Odometry And Mapping (LOAM) [24] algorithm was used for the odometry task. In particular, the LOAM algorithm is divided into two consecutive modules: (i) an odometry algorithm that is computed at a high frequency with low precision, and (ii) a mapping algorithm that is executed at a lower frequency but with a higher accuracy. By default, the odometry algorithm extracts 24 features, whereas the mapping algorithm extracts 240 features to have higher precision, with a ratio of 1:2 between corner and planar features. Both algorithms extract a fixed number of corner and planar features from frame F_{k+1} , find and matches the same features in the frame F_k , and iteratively minimize the distance presented in Equation 3 to compute the best alignment transformation T_{k+1} .

To distinguish between planar and corner features, a feature factor c is computed [24] for each point, where a low c value indicates a planar feature, whereas a high c value indicates a corner feature. Then, features are ordered based on the c values and N corner points are selected taking the highest c values, and $2N$ planar points are selected taking the lowest c values, where N is a user-defined parameter (set to 8 for the odometry step and to 80 for the mapping step).

The estimation error of the LOAM depends on the quality of the extracted features: environments containing repetitive features, such as tunnels or highways (hard to be detected in different frames), or with a low number of peculiar features leads to higher estimation errors. Moreover, since train and car motion mostly evolve in the X-Y plane, with low variations of the Z values, the Z evolution is not observable in such environments, unless the terrain presents substantial Z variations during motion.

In this experiment, three different sequences of the KITTI dataset have been chosen. The sequence with the identifier 00 is completely gathered in a urban environment with abundant high-quality features. The second, with identifier 01, is gathered on a highway, which has low-quality features. The third sequence, identified as 09, is divided into two parts: the first is collected in a leaning street with a lot of vegetation on the sides, and the second is gathered in a urban environment. The comparison is made against three different sequences generated by TrainSim, where the environment is composed of vegetation, railway structures (e.g., poles, electrified structures, and rails), fences, and stations.

Three different metrics have been chosen to evaluate the LOAM algorithm on the selected sequences: the estimation error of the translation along the X and the Y axis on the

TABLE I

LOCALIZATION RESULTS OF THE LOAM ALGORITHM [24] APPLIED TO THREE SEQUENCES GENERATED WITH TRAINSIM AND THREE SIMILAR SEQUENCES FROM THE KITTI DATASET [6]. THE VALUES INDICATE THE MEAN \pm THE VARIANCE AND THE MAXIMUM ERROR VALUE OF THE TRANSLATION ERROR ALONG THE X AXIS (TEX), THE TRANSLATION ERROR ALONG THE Y AXIS (TEY) COMPUTED IN A SINGLE ESTIMATED TRANSFORMATION BETWEEN TWO CONSECUTIVE LiDAR FRAMES EXPRESSED IN METERS, AND THE OVERALL ERROR OVER DISTANCE (EOD) IN PERCENTAGE

Sequence	TEX		TEY		EOD	
	$\mu \pm \sigma$	Max	$\mu \pm \sigma$	Max	$\mu \pm \sigma$	Max
Sim 1	0.06 \pm 0.04	0.18	0.08 \pm 0.05	0.17	0.25 \pm 0.15%	1.45%
Sim 2	0.34 \pm 0.18	0.69	0.54 \pm 0.28	1.18	2.47 \pm 2.06%	22.80%
Sim 3	0.15 \pm 0.10	0.38	0.09 \pm 0.15	0.89	0.83 \pm 1.26%	7.64%
Kitti 00	0.14 \pm 0.18	1.11	0.17 \pm 0.24	1.40	0.44 \pm 0.80%	13.2%
Kitti 01	2.17 \pm 0.71	3.57	1.04 \pm 0.51	2.76	11.40 \pm 3.68%	45.81%
Kitti 09	0.26 \pm 0.19	0.70	0.34 \pm 0.2	0.86	1.65 \pm 0.86%	7.22%

single transformations between each two consecutive LiDAR frames (TEX and TEY), and the cumulative position error in the X - Y plane computed over the traveled distance (EOD). The translation over Z was not taken into account due to the low variability of the Z coordinates, whereas the orientation estimations were not reported because the estimation error resulted to be below 1° . Table I shows the results of the LOAM algorithm applied to the six sequences.

The results indicate that the estimation error obtained on the simulated environment is comparable with the one over the KITTI sequences. In particular, the first simulated sequence is surrounded by vegetation and buildings, whereas the second and the third sequences present some fence stripes that introduce repetitiveness in the distribution of the features; in particular, the second sequence contains some heathland that is comparable with the highway environment. The same trend can be seen in the KITTI sequence, where the highest error occurs in the KITTI 01 sequence gathered in the highway, while the lowest error is achieved in the KITTI 00, which is entirely acquired in a urban environment.

C. Image Segmentation Analysis

In several works in the autonomous driving domain (e.g., [25], [26], [27], [28]), synthetic scenarios are used with *domain adaptation* (DA) techniques to improve the accuracy of a neural model whenever there is a scarce availability of real-world annotated samples, which is particularly true for the railway domain. In particular, during training, such techniques help to select learnable features from synthetic images that enhance the model outcome in real-world testing scenarios. Therefore, this section presents an experiment aimed at evaluating the improvement obtained on a neural model when augmenting the training set with synthetic images generated by TrainSim.

To do that, we evaluated the performance of a neural network on a real-world test set by comparing two different training modes: *semi-supervised* (SS) and *semi-supervised with domain adaptation* (SSDA) [28]. More specifically, in SS mode, the neural model is trained using only real-world

TABLE II

PERFORMANCE OF THE BiSeNET MODEL [29] ACHIEVED BY A SEMI-SUPERVISED (SS) APPROACH (WITH ONLY RAILSEM19 SAMPLES) AND A SEMI-SUPERVISED DOMAIN ADAPTATION (SSDA) WITH BOTH REAL-WORLD AND SYNTHETIC SAMPLES (RAILSEM19 + TrainSim). THE VALUES DENOTE THE INTERSECTION OVER UNION (IoU) AND $std \times 10$ OF EACH CLASS AMONG A 4-FOLD CROSS-VALIDATION ON RailSEM19. '10' AND '20' DENOTE THE NUMBER OF REAL-WORLD ANNOTATED SAMPLES, RANDOMLY EXTRACTED FROM RailSEM19

Class	SS-10	SSDA-10	SS-20	SSDA-20
pole	0.104 \pm 0.11	0.095 \pm 0.13	0.117 \pm 0.16	0.102 \pm 0.11
vegetation	0.344 \pm 0.26	0.298 \pm 0.79	0.407 \pm 0.17	0.366 \pm 0.26
terrain	0.102 \pm 0.78	0.156 \pm 0.14	0.164 \pm 0.35	0.194 \pm 0.16
sky	0.740 \pm 0.67	0.740 \pm 0.10	0.773 \pm 0.13	0.780 \pm 0.14
trackbed	0.325 \pm 0.14	0.366 \pm 0.41	0.391 \pm 0.44	0.415 \pm 0.05
rail-track	0.175 \pm 0.06	0.192 \pm 0.18	0.197 \pm 0.26	0.208 \pm 0.22

images, following supervised and unsupervised paradigms for labeled and unlabeled samples, respectively. In SSDA mode, instead, the model is trained using the same paradigms for real-world samples, but the training set is augmented with labeled synthetic images.

In the experiment presented here, SSDA was performed via a discriminator approach [26] and, for consistency, the SS mode was also implemented by a discriminator approach [27] using the real-world annotated samples as the source dataset.

Real-world images were taken from the RailSem19 dataset [10], containing more than 8000 annotated samples collected from both railway and urban scenarios. In our tests, 6000 samples were used for the training set: 6000 - k with annotations and k without annotations, setting $k = 10$ and $k = 20$ to observe the difference in performance. Other 2000 samples were used for the real-world test set.

In SSDA mode, the training set was augmented with 6700 annotated synthetic images collected from different simulated scenarios, similar to those described in Section IV-B, where different materials were randomly applied to the trackbed and the landscapes, and various lighting conditions were used to add some variability to the gathered images. Since RailSem19 and TrainSim define two different sets of object classes, the analysis was conducted on a subset of RailSem19 classes also present in TrainSim (see Table II and Figure 16), while all the remaining classes were considered as 'background'.

The neural architecture selected for the semantic segmentation task is a BiSeNetX39 [29], trained by the Adam optimizer [30] with its default settings and a learning rate of 0.003. Batch size and training steps were set to 30 and 8000, respectively. The training was performed by using the classic pixel-wise cross-entropy loss. Input images were resized to ($H=680$, $W=720$) to reduce the computational cost, while random crop (scale 1/2) and random horizontal flip were used for training set augmentation.

Table II reports the performance achieved on the RailSem19 dataset (details in the caption), showing that the use of TrainSim improves the IoU on crucial classes (rail-track, trackbed, and terrain), whereas the performance on other classes is reduced, most likely due to a more accentuated domain shift

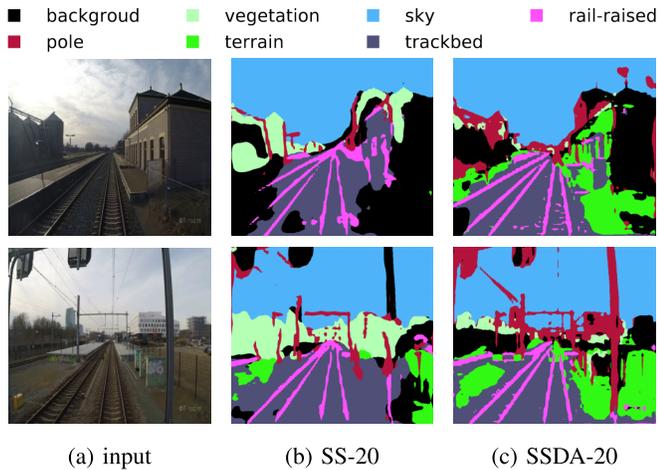


Fig. 16. Output predictions of RailSem19 real-world images.

between synthetic and real-world textures. Figure 16 shows two real-world images taken from RailSem19 (a) and the corresponding segmented images produced by SS-20 (b) and SSDA-20 (c). In accordance with Table II, the model trained using synthetic samples (SSDA-20) produces more accurate segmentation maps.

Despite the benefits discussed above, we also noticed that increasing the number of annotated real-world samples (i.e., more than 50) SSDA yields lower performance than SS (without TrainSim images). We believe this is due to a semantic domain shift between the simulated and real-world images, which become more relevant for a higher number of real-world samples. For instance, TrainSim does not account for complex textures contained in RailSem19 (e.g., crowded urban and driving scenarios). This forces the model to learn a constrained subset of visual patterns during SSDA, forgetting those that are not well-represented in TrainSim but still useful in real-world scenarios. Please also note that such a domain shift is more accentuated when running an unsupervised DA or without any DA strategy. This motivated us to explore SSDA, where a small subset of real-world data helps alleviate the domain shift.

These points open interesting future works for investigating novel DA approaches for railway scenarios. Finally, it is also worth remarking that, to the best of our knowledge, this work is the first one that proposes an SSDA approach for semantic segmentation in railway scenarios.

V. CONCLUSION

This paper presented TrainSim, a visual simulation framework designed to automatically generate a number of realistic railway scenarios and produce labeled datasets from emulated sensors, as LiDARs, cameras, and inertial measurement units. Such datasets are exported in a format suitable for training deep neural models for object detection, semantic segmentation, and depth estimation for camera data, or for processing 3D point clouds from a LiDAR. For each 3D point, the LiDAR model also provides the intensity of the backscattered ray, which can be used to simplify the discrimination of the tracks from other objects with higher diffusion coefficients.

The preliminary experiments carried out on the simulated sensors showed the effectiveness of the proposed approach, making the simulator a useful tool for investigating, training, and testing new perception algorithms for railway applications.

As a future work, we plan to extend the simulator by adding several new features as: (i) effects of adverse weather conditions (e.g., rain and snow) on visual sensors; (ii) railroad switches and turnouts; and (iii) moving objects in the environment. Other objectives will be the refinement of the train dynamics with a more accurate physical model, a better photorealism in terms of meshes, and the optimization of the code to reach a real-time performance, useful for using the tool in hardware-in-the-loop simulations.

Finally, we plan to conduct future investigations in domain adaptation methods for point-cloud data, since the current lack of proper real-world annotated samples limits a comprehensive evaluation of models trained with simulated LiDAR data in railway scenarios.

REFERENCES

- [1] *Localisation Working Group (LWG)—Railways Localisation System High Level Users' Requirements*, ERTMS User Group, Brussels, Belgium, 2019.
- [2] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "CARLA: An open urban driving simulator," in *Proc. 1st Annu. Conf. Robot Learn.*, 2017, pp. 1–16.
- [3] S. Shah, D. Dey, C. Lovett, and A. Kapoor, "AirSim: High-fidelity visual and physical simulation for autonomous vehicles," in *Proc. 11th Int. Conf. Field Service Robot.* Switzerland: Springer, 2018, pp. 621–635.
- [4] G. Brockman et al., "OpenAI gym," 2016, *arXiv:1606.01540*.
- [5] Stanford Artificial Intelligence Laboratory. *Robotic Operating System*. Accessed: May 23, 2018. [Online]. Available: <https://www.ros.org>
- [6] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? The KITTI vision benchmark suite," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2012, pp. 3354–3361.
- [7] J. Behley et al., "SemanticKITTI: A dataset for semantic scene understanding of LiDAR sequences," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 9297–9307.
- [8] M. Cordts et al., "The cityscapes dataset," in *Proc. CVPR Workshop Future Datasets Vis.*, vol. 2, 2015.
- [9] J. Fang et al., "Augmented LiDAR simulator for autonomous driving," *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 1931–1938, Apr. 2020.
- [10] O. Zendel, M. Murschitz, M. Zeilinger, D. Steininger, S. Abbasi, and C. Beleznai, "RailSem19: A dataset for semantic rail scene understanding," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2019, pp. 1–9.
- [11] F. Tschopp et al., "Experimental comparison of visual-aided odometry methods for rail vehicles," *IEEE Robot. Autom. Lett.*, vol. 4, no. 2, pp. 1815–1822, Apr. 2019.
- [12] S. Sahebdivani, H. Arefi, and M. Maboudi, "Rail track detection and projection-based 3D modeling from UAV point cloud," *Sensors*, vol. 20, no. 18, p. 5220, Sep. 2020.
- [13] Epic Games. *Unreal Engine*. Accessed: Apr. 25, 2019. [Online]. Available: <https://www.unrealengine.com>
- [14] A. Best, S. Narang, L. Pasqualin, D. Barber, and D. Manocha, "AutonoVi-Sim: Autonomous vehicle simulation platform with weather, sensing, and traffic control," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2018, pp. 1048–1056.
- [15] W. Tian et al., "Analysis and radiometric calibration for backscatter intensity of hyperspectral LiDAR caused by incident angle effect," *Sensors*, vol. 21, no. 9, p. 2960, Apr. 2021.
- [16] M. Záborský. *Geogen Studio*. Accessed: 2011. [Online]. Available: <https://code.google.com/archive/p/geogen/>
- [17] ZerstorenGames. *Train Template*. Accessed: 2019. [Online]. Available: <https://www.unrealengine.com/marketplace/en-US/product/train-template>
- [18] P. D. Groves, *Principles of GNSS, Inertial, and Multisensor Integrated Navigation Systems*. Norwood, MA, USA: Artech House, 2008.

- [19] Council of European Union. (2014). *Council Regulation (EU) No 1299/2014 on the Technical Specifications for Interoperability Relating to the Infrastructure Subsystem of the Rail System in the European Union*. [Online]. Available: <http://data.europa.eu/eli/reg/2014/1299/2019-06-16>
- [20] *TCRP Report 155—Track Design Handbook for Light Rail Transit Second Edition*, Transp. Res. Board Executive Committee, Transp. Res. Board, Washington, DC, USA, 2012.
- [21] C. Gu and C. Gu, *Smoothing Spline ANOVA Models*, vol. 297. New York, NY, USA: Springer, 2013.
- [22] P. Mania and M. Beetz, “A framework for self-training perceptual agents in simulated photorealistic environments,” in *Proc. Int. Conf. Robot. Autom. (ICRA)*, Montreal, QC, Canada, May 2019, pp. 4396–4402.
- [23] G. F. Gusmão, C. R. H. Barbosa, and A. B. Raposo, “Development and validation of LiDAR sensor simulators based on parallel raycasting,” *Sensors*, vol. 20, no. 24, p. 7186, Dec. 2020.
- [24] J. Zhang and S. Singh, “LOAM: LiDAR odometry and mapping in real-time,” in *Proc. Robot., Sci. Syst.*, vol. 2, 2014, pp. 1–9.
- [25] S. R. Richter, V. Vineet, S. Roth, and V. Koltun, “Playing for data: Ground truth from computer games,” in *Proc. Eur. Conf. Comput. Vis.* Amsterdam, The Netherlands: Springer, 2016, pp. 102–118.
- [26] Z. Wang et al., “Alleviating semantic-level shift: A semi-supervised domain adaptation method for semantic segmentation,” in *Proc. IEEE Comput. Vis. Pattern Recognit. Workshops*, Jun. 2020, pp. 936–937.
- [27] Y.-H. Tsai, W.-C. Hung, S. Schuster, K. Sohn, M.-H. Yang, and M. Chandraker, “Learning to adapt structured output space for semantic segmentation,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 7472–7481.
- [28] M. Wang and W. Deng, “Deep visual domain adaptation: A survey,” *Neurocomputing*, vol. 312, pp. 135–153, Oct. 2018.
- [29] C. Yu, J. Wang, C. Peng, C. Gao, G. Yu, and N. Sang, “BiSeNet: Bilateral segmentation network for real-time semantic segmentation,” in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 325–341.
- [30] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” 2014, *arXiv:1412.6980*.



Federico Nesti received the degree (cum laude) in robotics and automation engineering from the University of Pisa and the Ph.D. degree in emerging digital technologies from Scuola Superiore Sant’Anna, Pisa, Italy. He investigated adversarial examples in the physical world and trustworthiness of AI-based computer vision systems with Scuola Superiore Sant’Anna Pisa.



Giulio Rossolini received the degree (cum laude) in embedded computing systems engineering from Scuola Superiore Sant’Anna, Pisa, Italy, and the joint M.S. degree from Scuola Superiore Sant’Anna and the University of Pisa. He is currently pursuing the Ph.D. degree with the ReTiS Laboratory, Scuola Superiore Sant’Anna. His current research interests include the design and implementation of software tools to increase the trustworthiness of machine learning algorithms in computer vision applications and safety-critical systems.



Giorgio Buttazzo (Fellow, IEEE) received the degree (cum laude) in electronic engineering from the University of Pisa, the M.S. degree in computer science from the University of Pennsylvania, and the Ph.D. degree in computer engineering from Scuola Superiore Sant’Anna, Pisa, Italy. He is currently a Full Professor with Scuola Superiore Sant’Anna. He has authored seven books on real-time systems and over 300 papers in the field of real-time systems and neural networks. He was a recipient of the 13 best paper awards.



Gianluca D’Amico received the joint M.S. degree in computer science and networking from the Scuola Superiore Sant’Anna, Pisa, Italy, and the University of Pisa. He is currently pursuing the Ph.D. degree with the ReTiS Laboratory, Scuola Superiore Sant’Anna. His current research interests include computer graphics, LiDAR technology, and simulation tools for testing algorithms and deep neural networks in railway and automotive environments.



Salvatore Sabina received the M.S. degree in electronic engineering from the University of Pisa in 1983. From 1985 to 1998, he involved in satellite applications, hardware designs, and HRT operating systems. In 1999, he joined Ansaldo Segnalamento Ferroviario (now Hitachi Rail STS) as responsible for the Product Development Department. He is currently the Vice President with the STS Innovation Department, Hitachi Rail STS. He is also responsible for setting up railway innovative projects and collaborations with research and development organizations.



Mauro Marinoni (Member, IEEE) received the M.S. and Ph.D. degrees in computer engineering from the University of Pavia, Italy, in 2003 and 2007, respectively. Since 2007, he has been with the ReTiS Laboratory, Scuola Superiore Sant’Anna, Pisa, where he was an Assistant Professor from 2009 to 2020. He is currently a Researcher with Scuola Superiore Sant’Anna. He coordinated several European and industrial projects in different application fields from e-Health devices to autonomous and distributed systems.



Gianluigi Lauro received the degree (cum laude) in electronic engineering from the University of Naples Federico II in 2009. He is currently responsible for the development of Innovative train positioning and odometry solutions and the RAILGAP European Project at Hitachi Rail STS. He also coordinates project VOLIERA of the European Space Agency. His current research interests include sensor fusion for developing accurate innovative train positioning systems.