Dataset: An Indoor Smart Traffic Dataset and Data Collection System

Neiwen Ling*
The Chinese University of Hong Kong
Hong Kong, China
lingnw@link.cuhk.edu.hk

Yuze He*
The Chinese University of Hong Kong
Hong Kong, China
hy019@ie.cuhk.edu.hk

Nan Guan
City University of Hong Kong
Hong Kong, China
nanguan@cityu.edu.hk

Heming Fu
The Chinese University of Hong Kong
Hong Kong, China
hemingfu@link.cuhk.edu.hk

Guoliang Xing†
The Chinese University of Hong Kong
Hong Kong, China
glxing@cuhk.edu.hk

ABSTRACT

Smart traffic is an emerging research area gaining more attention due to a class of emerging applications such as autonomous driving. Most smart traffic scenarios are outdoors, which are hard to collect traffic data and build demanding sensing systems. In this work, an indoor smart traffic testbed with an F1TENTH autonomous driving vehicle is built, allowing the collection of traffic datasets under different scenarios and performing various smart traffic tasks. This novel data collection system and collected dataset can help research teams build various smart traffic systems and evaluate indoor smart traffic datasets. The collected traffic light dataset is publicly available at the link[1].

CCS CONCEPTS

• **Computer systems organization → Sensor networks**.

KEYWORDS

Smart Traffic, Dataset, Autonomous Driving, F1TENTH

ACM Reference Format:
Neiwen Ling*, Yuze He*, Nan Guan, Heming Fu, and Guoliang Xing†. 2022. Dataset: An Indoor Smart Traffic Dataset and Data Collection System. In *The Fifth International Workshop on Data: Acquisition To Analysis (DATA '22), November 6–9, 2022, Boston, MA, USA*. ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3560905.3567772

1 INTRODUCTION

The exploration of smart traffic has drawn growing attention in recent years as it can support a wide range of applications, such as assisting autonomous driving [4], improving public safety and city services [10], etc. Research on smart traffic spans several different important fields, such as autonomous driving [5, 7, 12, 14], real-time AI [9, 18], robotics [3, 8], as well as simultaneous localization and mapping [17]. However, it is difficult for researchers to conduct rapid data collection, system development and implementation, algorithm design and evaluation due to the complex and unpredictable real-world environment [11]. For example, dynamics in the real-world environment not only complicate the data collection procedure but also degrade the quality of the collected data. For example, excessive heat in a hot summer may lead to the malfunction of data collection systems, and heavy rain will affect the quality of the collected visual sensor data. Thus, there is a high demand for reproducible and controllable simulation platforms for smart traffic scenarios. While significant efforts have been made to collect smart traffic datasets from outdoor scenarios [1, 2], building indoor smart traffic testbeds and collecting data from such simulated indoor scenarios are scarce.

In this work, we present the design and implementation of an indoor smart traffic testbed that enables the simulation of multiple road scenarios. We also provide a use case to demonstrate how to utilize our testbed for a specific application. In particular, we mimic real-world road scenarios by simulating real vehicles and roadside infrastructures, such as lane fences, and traffic signs/lights. Our testbed uses F1TENTH autonomous vehicles, which are fully functional and open-source. An F1TENTH vehicle is 1/10th-scale of a real self-driving car but only 1/100th of the cost. An F1TENTH vehicle mainly consists of an embedded compute platform, a motor and servo controller, a power distribution board, a WiFi telemetry, and sensors like LiDARs and cameras, which makes it complex enough to mimic the real vehicle's driving dynamics. Moreover, F1TENTH has a large and vibrant community. It has been used as the main platform in 7 international autonomous driving competitions, and more than a dozen institutions have offered courses based on it. Thus, it is an ideal platform for streamlined algorithm development, testing, and validation in the field of autonomous systems. To demonstrate the utility of our testbed, we develop several real-time traffic light classification applications with this testbed. Specifically, we construct a circular traffic lane with roadside infrastructures and drive the F1TENTH vehicle along the lane fences to collect traffic light data. We then train a popular DNN model (i.e., tinyYOLO [13]) based on the collected data and implement it on the vehicle. With this trained model, our vehicle can classify the traffic lights in real-time.

[1] https://zenodo.org/record/7181314#.Y0a0qXZBxD8

*Co-first author, †Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
DATA '22, November 6–9, 2022, Boston, MA, USA
© 2022 Association for Computing Machinery.
ACM ISBN 978-1-4503-9886-2/22/11…$15.00
https://doi.org/10.1145/3560905.3567772

The major contributions of this paper are as follows. First, we build an indoor smart traffic testbed with a widely-used F1TENTH autonomous vehicle. Second, we run real-time smart traffic tasks such as traffic light detection and road condition collection on the vehicle. The testbed we built can guide researchers to build similar data collection systems and thus evaluate their designed systems/algorithms. For example, researchers in real-time AI can use the dataset collected by the vehicles at different speeds to evaluate the task performance under different deadline settings. The rest of this paper is organized as follows. Section 2 describes the system architecture and implementation. Section 3 describes the collected dataset and experimental results for a specific use case. Section 4 concludes this study with some future works.

2 SYSTEM AND IMPLEMENTATION

2.1 Overall System

We set up our indoor smart traffic testbed in a $5.8m \times 6.8m$ lab. This testbed combines the sensing system from both the roadside and the vehicle side, as well as the infrastructure for building traffic scenarios, such as traffic lanes, traffic signs, and traffic lights. We can simulate various road scenarios by adjusting lane shapes and the locations of traffic signs and traffic lights. The F1Tenth autonomous vehicle is responsible for simulating the driving car and collecting the smart traffic data from the vehicle side. We can perform several autonomous driving tasks such as lane detection, route navigation, traffic sign detection, and traffic light detection on our F1Tenth vehicle. The vision system could record the overall running status of the entire lane from the roadside.

2.1.1 Smart Traffic Infrastructures. In this testbed, we use fences to simulate the road's boundaries and use movable toy traffic signs and traffic lights to simulate those in real-world traffic scenes. The traffic signs are of height $1m$, and the traffic lights come in two heights: $0.7m$ and $1.1m$, which are used to simulate different types of real-world traffic lights with different heights. The traffic lights, traffic signs, and lanes are all movable, which facilitate flexible adjustment of the infrastructures' locations and allow us to construct diverse smart traffic scenarios. For example, we show scenes and layouts of two possible test scenarios in Fig. 1. As shown in Fig. 1(a)1(b), we build a circular runway with a diameter of $4m$, and place traffic signs/lights every $0.5m$ in the inner circle and every $1m$ in the outer circle to simulate the actual traffic scenario. The traffic lights will flash alternately, and we set the traffic light alternating interval to $90s$, which is consistent with the common practice of traffic light duration. We also construct a straight track with one corner as shown at Fig. 1(c)1(d). More traffic scenarios such as crossroads can be built by changing the track layout and traffic light/sign positions.

2.1.2 F1TENTH Autonomous Vehicle. We build our F1TENTH Autonomous Vehicle System from scratch. The designed layout and different views of the vehicle are shown in Fig. 2. As shown in the side view of the F1TENTH vehicle, our car has two chassis. The lower level chassis is mounted with a LIPO battery, a brushless DC motor, a servo steering, and a VESC motor and servo controller, which serve as the foundation of the car. The upper-level chassis is mounted with a small heterogeneous embedded platform (NVIDIA TX2 with Orbitty Carrier board), two RGB cameras, and

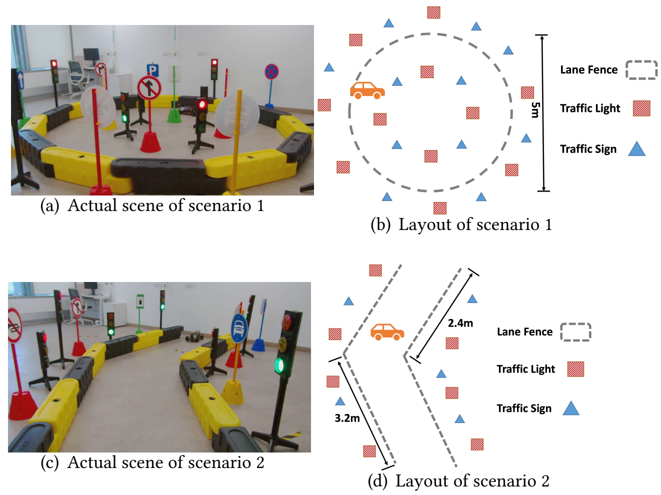

(a) Actual scene of scenario 1
(b) Layout of scenario 1
(c) Actual scene of scenario 2
(d) Layout of scenario 2

Figure 1: Indoor smart traffic test scenario.

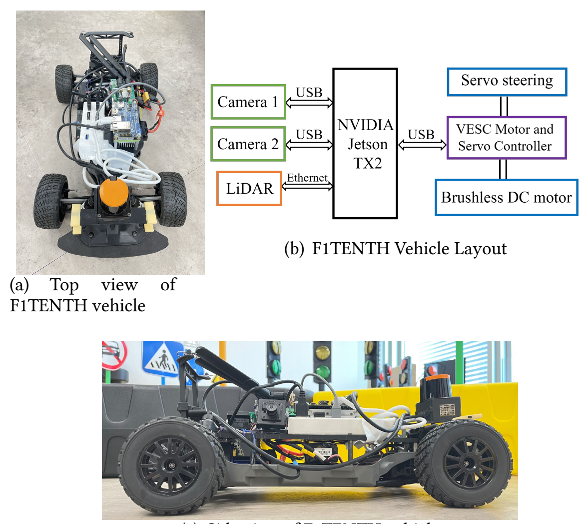

(a) Top view of F1TENTH vehicle
(b) F1TENTH Vehicle Layout
(c) Side view of F1TENTH vehicle

Figure 2: F1TENTH autonomous driving testbed.

a HOKUYO UST-10LX LiDAR [6], which work together to perceive the surrounding environment for multiple smart traffic tasks. For standardized sensor management, we install Linux and ROS (Kinetic) on the NVIDIA TX2.

2.1.3 The Roadside Vision System. Roadside surveillance is a main component of the smart traffic. For example, the camera installed at the smart traffic light can assist autonomous driving by providing the traffic condition of the entire road. In our testbed, we construct a roadside vision system to simulate real-world roadside surveillance. This vision system consists of a Pico Zense ToF camera [15] and a laptop. Fig.3 shows the rear and front views of the vision system. This vision system records the running status of the whole test scenario via capturing $1920 * 1080$ RGB images every $0.1s$.

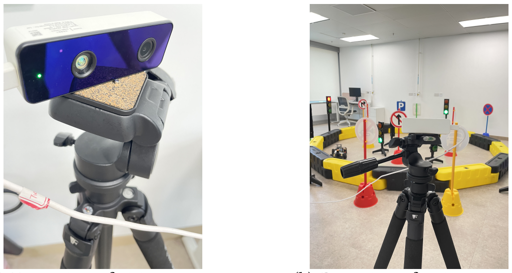

(a) Front view of vision system (b) Rear view of vision system

Figure 3: Roadside vision system.

2.2 Smart Traffic Tasks

This section gives some instances for simulating smart traffic tasks in our smart traffic testbed. To simulate traffic sign detection from the vehicle side, we run a lane detection algorithm on the F1TENTH vehicle and a traffic light detection algorithm based on a popular DNN model (e.g., YOLO [13]). The fence following algorithm allows the car to run along a preset track. The traffic sign detection algorithm processes images captured by the camera installed on the vehicle and thus achieves real-time traffic sign detection. Similar traffic sign/light detection or obstacle detection tasks from the roadside can be performed on the computing unit of the vision system. To simulate roadside-assisted autonomous driving, we can run detection tasks on the computing unit of the vision system and transmit the roadside information to the vehicle by establishing a socket connection. The vehicle can fuse the data from the roadside with its own sensor data to perform downstream autonomous driving tasks such as route planning.

3 EXPERIMENTS AND RESULTS

3.1 Collection Procedure

In this section, we provide a use case of our indoor smart traffic testbed, including data collection, algorithm implementation and evaluation. Specifically, we collect a traffic light dataset, train a popular DNN model, and perform real-time traffic light classification tasks on the vehicle. First, we build a $6m \times 5m$ circular track and place traffic signs and traffic lights on the inner and outer circles of the track at intervals of $1m$ and $0.5m$ respectively. Then, we use the F1TENTH vehicle introduced in Section 2.1.2 to simulate the real vehicle and collect traffic light data. A fence following algorithm is implemented on the car to control its mechanical movement, which makes the car detect the fence using LiDAR and drive along the fence autonomously. Two cameras are mounted on both sides of the car, which capture images of traffic lights and traffic signs on both sides of the track. The frame rate of the camera is $12FPS$. In the data collection process, we remotely launch/stop the F1TENTH car via an SSH session between the car and our laptop. Once launched, the car starts to run along the fence at the speed of $1m/s$, and saves the images captured by the two cameras. Some collected images are shown in Fig 4.

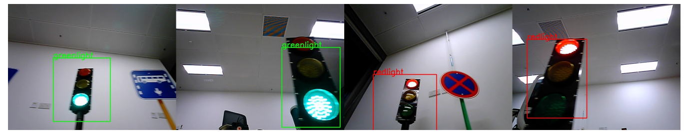

Figure 4: Examples of captured images and detection results.

3.2 Labeling

We use the image annotation tool LabelImg [16] for the data annotation of traffic light detection task, which aims to detect the position and state of traffic lights in images. We manually annotate the bounding boxes that enclose the traffic lights in each image, each with a label indicating the current state of the traffic light, 0 for a green light and 1 for a red light. Annotations are saved as YOLO format files for the training of the traffic light detection model. We summarize the dataset in Table 1. The data with label are available at the link[2]. This dataset has already been used in a published work for conducting a case study[9].

Table 1: Indoor smart traffic dataset.

	# total images	# green lights in images	# red lights in images
Indoor smart traffic dataset	3507	1582	1925

3.3 Benchmark

We choose the Tiny-YOLO [13] object detector for real-time traffic light detection task on the F1TENTH vehicle. The small model size ($< 50MB$) and fast inference speed make the Tiny-YOLO object detector naturally suited for embedded deep learning devices such as the NVIDIA Jetson TX2. We train a Tiny-YOLO model on the annotated image dataset. To evaluate the performance of the model, we deploy two identical trained models on the F1TENTH vehicle, and then launch the vehicle to run along the fence with the two Tiny-YOLO models running on it simultaneously. The two models individually process images from two cameras. As the vehicle moves, live images of both sides of the track are continuously captured by two cameras and fed into two models, which perform traffic light detection in real-time. The inference results of the models are saved to the NVIDIA TX2. In the experiment, a total of 240 images were captured and processed. We use the State Detection Accuracy (SDA) to measure the performance of the models, which are defined as follows:

$$SDA = \frac{TG + TR}{TG + TR + FG + FR} \tag{1}$$

True Green (TG) means that a green traffic light is detected as green, and False Green (FG) means that a red traffic light is detected as green or not detected. The same goes for red traffic lights. This metric indicates the performance of the model in recognizing traffic light states. Experiments show that the average SDA of the two Tiny-YOLO models is 99.43%. Some quantitative results are shown

[2] https://zenodo.org/record/7181314#.Y0a0qXZBxD8

in Fig. 4, which shows that the trained model accurately locates the positions and recognizes the state of the traffic lights in the images.

4 CONCLUSION

This work develops an indoor smart traffic testbed that enables the simulation of diverse real-world road scenarios. We also deploy a real-time traffic application on the tested to illustrate its utility. The data collection from more simulated smart traffic scenarios is on-going and the deployment for other applications is in progress.

Future work will need to investigate how to simulate more smart traffic components and more complex road conditions, such as vehicle collisions. At the same time, we realize that the data collection devices were not diverse enough. Developers can deploy more sensors, such as thermal and radar, to collect data on more modalities for further research.

ACKNOWLEDGEMENT

The work described in this paper was supported by the Innovation and Technology Commission of Hong Kong under Grant CPII-EW01610 (RP4-3).

REFERENCES

[1] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020.

[2] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *International Journal of Robotics Research (IJRR)*, 2013.

[3] Oleg Golubchikov and Mary Thornbush. Artificial intelligence and robotics in smart city strategies and planned smart development. *Smart Cities*, 3(4), 2020.

[4] Dapeng Guo, Melody Moh, and Teng-Sheng Moh. Vision-based autonomous driving for smart city: a case for end-to-end learning utilizing temporal information. In *International Conference on Smart Computing and Communication*, pages 19–29. Springer, 2020.

[5] Yuze He, Li Ma, Zhehao Jiang, Yi Tang, and Guoliang Xing. Vi-eye: semantic-based 3d point cloud registration for infrastructure-assisted autonomous driving. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*, pages 573–586, 2021.

[6] HOKUYO. Ust-10lx. https://hokuyo-usa.com/products/lidar-obstacle-detection/ust-10lx, 2016.

[7] Radoslav Ivanov, Taylor J Carpenter, James Weimer, Rajeev Alur, George J Pappas, and Insup Lee. Case study: verifying the safety of an autonomous racing car with a neural network controller. In *Proceedings of the 23rd International Conference on Hybrid Systems: Computation and Control*, pages 1–7, 2020.

[8] Achin Jain, Matthew O'Kelly, Pratik Chaudhari, and Manfred Morari. Bayesrace: Learning to race autonomously using prior experience. *arXiv preprint arXiv:2005.04755*, 2020.

[9] Neiwen Ling, Kai Wang, Yuze He, Guoliang Xing, and Daqi Xie. Rt-mdl: Supporting real-time mixed deep learning tasks on edge platforms. In *Proceedings of the 19th ACM Conference on Embedded Networked Sensor Systems*, pages 1–14, 2021.

[10] Duncan McLaren and Julian Agyeman. *Sharing cities: A case for truly smart and sustainable cities.* MIT press, 2015.

[11] Matthew O'Kelly, Hongrui Zheng, Dhruv Karthik, and Rahul Mangharam. F1tenth: An open-source evaluation environment for continuous control and reinforcement learning. *Proceedings of Machine Learning Research*, 123, 2020.

[12] Matthew O'Kelly, Hongrui Zheng, Achin Jain, Joseph Auckley, Kim Luong, and Rahul Mangharam. Tunercar: A superoptimization toolchain for autonomous racing. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5356–5362. IEEE, 2020.

[13] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv*, 2018.

[14] Aman Sinha, Matthew O'Kelly, Hongrui Zheng, Rahul Mangharam, John Duchi, and Russ Tedrake. Formulazero: Distributionally robust online adaptation via offline population synthesis. In *International Conference on Machine Learning*, pages 8992–9004. PMLR, 2020.

[15] TEGAKARI. Pico zense sensor module for machine vision. https://www.tegakari.net/en/2019/03/pico_zense/, 2020.

[16] Tzutalin. Labelimg. https://github.com/tzutalin/labelImg, 2015.

[17] Bjorn WM Walk. Towards achieving synergy in multiple slam algorithms via pose-graph optimization.

[18] Zhihe Zhao, Zhehao Jiang, Neiwen Ling, Xian Shuai, and Guoliang Xing. Ecrt: An edge computing system for real-time image-based object tracking. In *Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems*, pages 394–395, 2018.