



Traffic3D: A Rich 3D-Traffic Environment to Train Intelligent Agents

Deepeka Garg^(✉), Maria Chli, and George Vogiatzis

Aston University, Birmingham B4 7ET, UK
{gargd,m.chli,g.vogiatzis}@aston.ac.uk

Abstract. The last few years marked a substantial development in the domain of Deep Reinforcement Learning. However, a crucial and not yet fully achieved objective is to devise intelligent agents which can be successfully taken out of the laboratory and employed in the real world. Intelligent agents that are successfully deployable in true physical settings, require substantial prior exposure to their intended environments. When this is not practical or possible, the agents benefit from being trained and tested on powerful test-beds, effectively replicating the real world. To achieve traffic management at an unprecedented level of efficiency, in this paper, we introduce a significantly richer new traffic simulation environment; Traffic3D. Traffic3D is a unique platform built to effectively simulate and evaluate a variety of 3D-road traffic scenarios, closely mimicking real-world traffic characteristics including faithful simulation of individual vehicle behavior, precise physics of movement and photo-realism. We discuss the merits of Traffic3D in comparison to state-of-the-art traffic-based simulators. Along with deep reinforcement learning, Traffic3D facilitates research across various domains such as object detection and segmentation, unsupervised representation learning, visual question answering, procedural generation, imitation learning and learning by interaction.

Keywords: Virtual reality 3D traffic simulator · Intelligent transportation systems · Machine learning · Deep learning

1 Introduction

Training agents to autonomously act in real-world settings involves challenges that transcend beyond the competency of commonly-used annotated data in a typical supervised learning setting. In control optimization problems such as traffic infrastructure optimization, the control strategies that lead to the optimal solution are not obvious at the outset. Equipping an agent with the ability to learn and discover effective solutions, enables it to be adaptable to a variety of settings. However, learning often depends on aggressive exploration of unknown space; as to generalize well in an unseen 3D environment, an agent is required to demonstrate robustness to different variations of the environment it is to be

exposed to. Consequently, to be able to perform well within a dynamic physical environment, an agent is required to observe and learn from a considerable set of environment states, where the agent is also able to perceive the outcome of its behavior by receiving feedback from the environment it is interacting with. Real-world physical environments fulfill these requirements, but they are expensive and unsafe to train in and are hard to scale. Especially in the traffic optimization case, among others, safety and economic cost concerns in obtaining sufficient amount of diversity (different traffic conditions and different seasons) in the training data preclude this vital training from being carried out in situ.

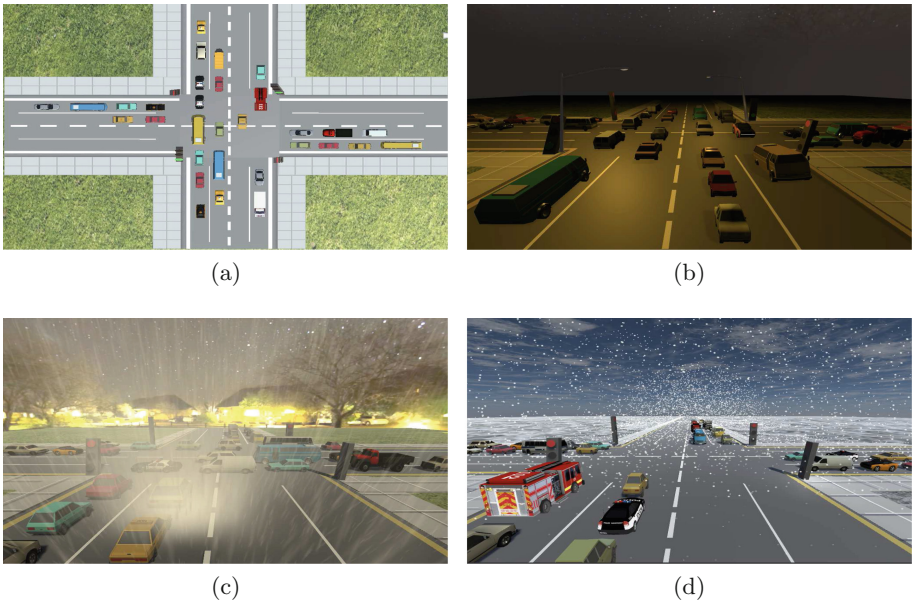


Fig. 1. Different views of a junction in Traffic3D. (a) Top-down view. (b) A night scene. (c) A rainy day. (d) A snowy day.

In addition, deep learning solutions (such as deep neural networks), the state-of-the-art paradigm used to effectively train agents to autonomously accomplish tasks (such as autonomous traffic infrastructure optimization) are known to depend on a large amount of training data to achieve peak performance, which further makes training an autonomous agent in the physical world impractical. An alternative is to train agents in simulations, providing a viable environment for protocol development.

Agents trained in more realistic settings are known to learn features which are generalizable to their corresponding real-world environments [1]. However, the most-widely used state-of-the-art traffic simulators fail to deliver crucial functionalities that are indispensable for authentic traffic simulation [2]. They lack in realism, diversity, stochasticity and perception challenges of the real-world.

To bridge the gap between simulations and real-world traffic specifications, we created a traffic micro-simulation tool; Traffic3D. Traffic3D is an interactive 3D-environment designed to precisely simulate traffic entities, their interactions and their emergent physical and visual properties.

Table 1. Comparison between different traffic-based and deep learning-based simulation environments.

Environment	Suitable for traffic simulation	Photo-realistic	3D	Physics	Customizable
SUMO [2]	Yes	No	No	Yes (with restrictions ^a)	Yes (with restrictions ^b)
VISSIM [2]	Yes	Yes	Yes	Yes (with restrictions ^c)	Yes (with restrictions ^d)
TORCS [3]	Yes	Yes	Yes	Yes (with restrictions ^e)	Yes (with restrictions ^f)
Virtual KITTI [1]	Yes	Yes	Yes	No ^g	Yes
CHALET [4]	No	Yes	Yes	Yes	Yes
AI2-THOR [5]	No	Yes	Yes	Yes	Yes
ATARI [6]	No	No	No	No	No
DeepMind Lab [7]	No	No	Yes	No	Yes
Traffic3D	Yes	Yes	Yes	Yes (with high degree of realism)	Yes (fully)

^a no proper reactive control to random incidents like collisions between vehicles.

^b does not support simulation of autonomous vehicles and does not prioritize public transport.

^c unrealistic lane-closing behavior.

^d restrictions in customizing delay.

^e limited sensor suite.

^f does not support road intersection simulation.

^g information not available.

2 Related Work

Traffic optimization being an established field, a number of traffic simulators have been developed. Also, a handful of deep learning platforms are available for training and benchmarking purposes of deep learning agents, mostly built around computer game environments. Pell et al. [2] intensively reviewed a comprehensive set of traffic simulation platforms. The review acknowledges that none of the

currently-used traffic simulation tools are able to adequately deliver important functionalities that are fundamental to faithful traffic simulation. The existing traffic models lack in detail and flexibility, and we believe that a detailed network model with efficient real-time traffic information collection capabilities is necessary to simulate heterogeneous transportation networks. In Table 1, we summarize the capabilities of a few widely-used traffic and deep learning-based simulation environments over important simulation characteristics; photo-realistic graphical rendering, 3D nature of objects, simulation physics and flexibility to customize the simulation environment according to the requirements of the application. It is evident that no single simulator supports comprehensive traffic-based research and analysis. In contrast, Traffic3D provides a flexible platform to facilitate research for novel traffic optimization protocol development and analysis in an environment that is as close as possible to a physical one.

3 Our Simulation Environment: Traffic3D

The main contribution of this paper is our gamified simulator; Traffic3D. Traffic3D consists of a diverse range of traffic scenes including a variety of photo-realistic vehicles (such as emergency vehicles, personal and public transport vehicles) and street furniture (sidewalks and traffic lights etc.). Scenes include, from the clear view of a sunny day, to the blurred scene of a dimly-lit night, a rainy and a snowy day (as shown in Fig. 1). We use real-world traffic images as a reference to create 3D-traffic scenes with photo-realistic lighting and texture. To facilitate control optimization, we developed a framework to enable an agent perform actions in a 3D-traffic environment and subsequently, perceive the outcomes of its actions. Our framework includes integration of Unity engine with deep learning python support. Figure 1 shows different views Traffic3D’s junction’s graphical display. In the next couple of paragraphs we discuss architectural design and performance capabilities of Traffic3D.

3.1 Traffic3D’s Design and Architecture

Traffic3D is built of two elements; (1) the traffic environment built within the Unity Game Development Platform¹ and (2) a light-weight python API that supports online learning. We designed this architecture to facilitate implementation of different control algorithms for various traffic-based and general applications. The traffic simulation environment and the learning agent (implemented in python) form a client-server system. The traffic environment acts as a client while the learning agent acts as server. The learning agent listens for http requests from the traffic environment and responds immediately as soon as a request is made. Figure 2 illustrates Traffic3D’s traffic environment and the learning agent’s bilateral interaction pipeline.

¹ <https://unity3d.com/>.

3.2 Traffic3D’s Performance

The socket data exchange mechanism we implemented to facilitate interaction between the simulation and the deep learning python module works accurately and seamlessly. However, the deep learning training is typically long. In order to accelerate training, Traffic3D allows control of the frame rate and speed of the simulation without affecting its quality. Collecting data using an accurate simulator is still faster and safer than doing so directly from the physical world [8]. At the same time, to boost Traffic3D’s performance in computationally-intensive machine learning tasks, we are currently exploring an alternative architecture version for Traffic3D. Under this architecture, the simulation engine will be running as a powerful server application, interacting with another server which is responsible for the machine learning. Visualization software, such as Unity, will act as a client to the simulation server.

4 Traffic3D’s Properties

Here, we discuss some of the key simulation properties of Traffic3D. This section highlights Traffic3D’s potential in creating natural-looking and realistically-operating traffic environment with believable visuals and physics.

4.1 Traffic3D’s Physical Properties

To ensure high-precision simulation, Traffic3D allows effective modelling of complex physical interactions between the transportation entities based on mass, friction and other forces such as gravity. Within the Traffic3D environment, vehicles are independently tuned to react appropriately accordingly to their input parameters and also to the presence of other vehicles and traffic infrastructure (such as their progressive slow-down and acceleration). Vehicles in Traffic3D are well-equipped to behave realistically. They have appropriate sensors (i.e. ray-cast sensors which can be customized to use as an equivalent to Lidar and microwaves) and programmed behaviour to avoid collisions, overtake and give way to emergency vehicles.

4.2 Traffic3D’s Visual Properties

Unlike most of the state-of-the-art traffic simulators (such as SUMO), Traffic3D provides high-definition graphical rendering facilitating transfer of developed models to the real-world with no or minimum pre-training. In Traffic3D, light, material, texture and scale work in synergy to make digital content look as close to a real scene as possible. The lighting techniques used in Traffic3D create immersive scene lighting, mimicking the real-world effects. Objects’ shadows can be dynamically cast on the scene, adding further realism to the environment.

4.3 Traffic3D’s Diversity of Parameters

To ensure stability and generalizability of agents trained using Traffic3D to different variations of the environment, Traffic3D supports plug-n-play architecture such that distinct scenes can be reliably incorporated. Essentially, Traffic3D facilitates seamless creation and switching between different types of scenes during runtime without interrupting the simulation or causing frame-rate hiccups. For example, traffic density can be dynamically varied independently, or around different times of the day (rush and quiet hours), as well as different weather conditions (e.g. rain, snow) during run-time, to study quantitatively the impact of different times and weather conditions on the learning agent’s behavior. The current level of diversity offered by Traffic3D includes: (1) vehicle number, types, models, colors and sizes, (2) vehicles’ trajectory configuration and speed, (3) different road layouts and surface textures, (4) variety of street furniture, (5) different lighting and weather conditions and (6) different camera position and orientation to comprehensively capture a wide range of possible aspects of a traffic scene.

4.4 Traffic3D’s Sensory Data Collection Capabilities

Traffic3D facilitates fast, inexpensive, limitless, diverse and photo-realistic collection of traffic data for various research purposes such as online reinforcement learning and supervised learning. Cameras in Traffic3D replicate the operation of real-world cameras; supporting capture of photo-realistic images and videos with proper field-of-view and depth-of-view. To study the effects of occlusions in the traffic environment, which is common in real-world traffic scenarios; vehicles and street furniture that are not currently being seen by the camera can have their rendering disabled. Multiple cameras can be deployed at the same time within a scene to perceive different aspects of the scene and their views can be combined in different ways. For example, in addition to the main view of the complete scene, when simulating a car, camera output to show a rear view mirror footage can also be simulated separately, if needed.




Fig. 2. The design architecture of Traffic3D encompassing the simulation environment and a python deep learning agent.

4.5 Traffic3D's Reusability

Traffic3D consists of a library of pre-created traffic environments. In addition, it offers complete flexibility over the deployment of new simulation scenarios; allowing complete customization over vehicle spawning, vehicle routing, road layouts and street furniture placement. Apart from creating simulation scenes manually, which can be tedious and less intuitive, Traffic3D also allows its users to programmatically generate scenes, even during run-time. Any of the traffic elements (such as vehicles) available within the project can be programmatically placed at any position and their physical (such a vehicle's size, mass and its speed) and visual behavior (such as a vehicle's texture) can also be programmatically configured.

5 Conclusion

This paper introduced our novel traffic simulation tool; Traffic3D. Traffic3D is created to build visually and physically intelligent traffic simulation models to build and test new technology with the goal to its eventual deployment in the real-world. It supports unique simulation features such as complex physical phenomenon, creation of relevant content such as traffic objects with appropriate background and photo-realism, comprehensibility, robustness, adaptability, partial observability challenges and inexpensive collection of diverse training data. Going forward, to further improve realism and preserve objects' geometry and other aspects, we intend to incorporate photogrammetry to enable automatic creation of traffic environments within Traffic3D from real-world traffic images and videos.

References

1. Gaidon, A., Wang, Q., Cabon, Y., Vig, E.: Virtual worlds as proxy for multi-object tracking analysis. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4340–4349 (2016)
2. Pell, A., Meingast, A., Schauer, O.: Trends in real-time traffic simulation. *Transp. Res. Procedia* **25**, 1477–1484 (2017)
3. Wymann, B., Espié, E., Guionneau, C., Dimitrakakis, C., Coulom, R., Sumner, A.: TORCS, the open racing car simulator. *Software*, vol. 4, no. 6 (2000). <http://torcs.sourceforge.net>
4. Yan, C., Misra, D., Bennett, A., Walsman, A., Bisk, Y., Artzi, Y.: CHALET: cornell house agent learning environment. arXiv preprint [arXiv:1801.07357](https://arxiv.org/abs/1801.07357) (2018)
5. Kolve, E., Mottaghi, R., Gordon, D., Zhu, Y., Gupta, A., Farhadi, A.: AI2-THOR: an interactive 3D environment for visual AI. arXiv preprint [arXiv:1712.05474](https://arxiv.org/abs/1712.05474) (2017)
6. Bellemare, M.G., Naddaf, Y., Veness, J., Bowling, M.: An evaluation platform for general agents: the arcade learning environment. *J. Artif. Intell. Res.* **47**, 253–279 (2013)
7. Beattie, C., et al.: Deepmind lab. arXiv preprint [arXiv:1612.03801](https://arxiv.org/abs/1612.03801) (2016)
8. Zhu, Y., et al.: Target-driven visual navigation in indoor scenes using deep reinforcement learning. In: 2017 IEEE International Conference on Robotics and Automation (ICRA), pp. 3357–3364. IEEE (2017)