Traffic3D: A New Traffic Simulation Paradigm

Deepeka Garg
Aston University
Birmingham, West Midlands
gargd@aston.ac.uk

Maria Chli
Aston University
Birmingham, West Midlands
m.chli@aston.ac.uk

George Vogiatzis
Aston University
Birmingham, West Midlands
g.vogiatzis@aston.ac.uk

ABSTRACT

The field of Deep Reinforcement Learning has evolved significantly over the last few years. However, an important and not yet fully-attained goal is to produce intelligent agents which can be successfully taken out of the laboratory and employed in the real-world. Intelligent agents that are successfully deployable in real-world 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 work, we demonstrate a significantly richer new traffic simulation environment; Traffic3D, a platform 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. In addition to deep reinforcement learning, Traffic3D also facilitates research in several other domains such as imitation learning, learning by interaction, visual question answering, object detection and segmentation, unsupervised representation learning and procedural generation.

KEYWORDS

Virtual Reality 3D-Traffic Simulator; Intelligent Transportation Systems.

ACM Reference Format:

Deepeka Garg, Maria Chli, and George Vogiatzis. 2019. Traffic3D: A New Traffic Simulation Paradigm. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,* IFAAMAS, 3 pages.

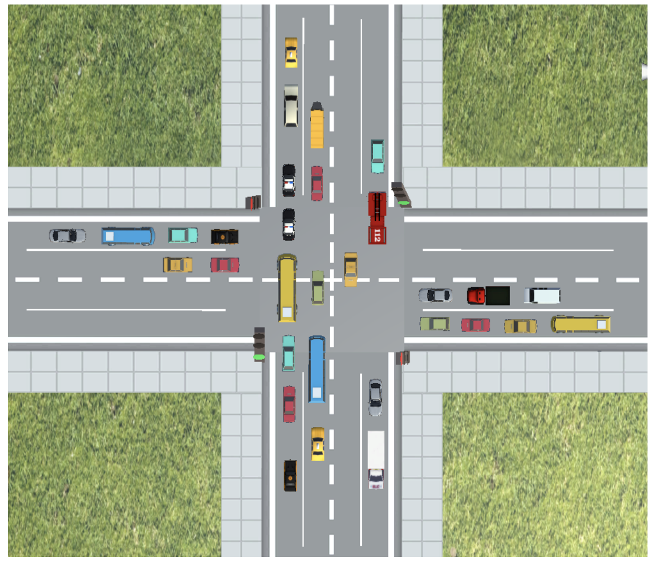

Figure 1: A view of Traffic3D's Graphical Display

1 INTRODUCTION

Training autonomous agents to act in a real world setting entails challenges that transcend beyond the commonly-used annotated data in a supervised learning setting. The considerable set of environmental states an agent may observe and learn from require interactive training environments, where the agent is able to observe the outcome of its behavior by receiving feedback from the environment it is interacting with. Real-world physical environments satisfy these requirements, but they are expensive, unsafe and hard to scale. In addition, deep learning solutions; deep neural networks, the state-of-the-art paradigm used to effectively train agents to autonomously accomplish tasks (such as autonomous driving, autonomous traffic infrastructure control [1] etc.) are known to depend on a large amount of training data to achieve peak performance, which further makes training an agent in the physical world infeasible [2–4]. An alternative is to train agents in simulations which provide a safe, controlled and accelerated environment for protocol development. We believe that the more faithful the simulation environment is to the real world, the less pre-training will be required for the autonomous agent prior to deploying it, and the more effective the agent is likely to be in its intended setting. This has already been acknowledged in the literature, where an autonomous robot is trained within a simulation environment and the trained model is effectively transferred to a real-world robot [5]. However, the most prominent state-of-the-art traffic simulators fail to deliver important functionalities that are fundamental to realistic traffic simulation. To address the discrepancy between simulations and real-world, we have created a traffic micro-simulation tool; Traffic3D, an interactive environment, designed to conduct rigorous, realistic and high quality traffic simulation. Traffic3D, rich in its content and structure, effectively reproduces real-world dynamic and diverse traffic environments. Along with offering 3D photo-realistic graphics and realistic physics of movement of transportation entities, Traffic3D supports a real-time interface with a learning agent to ensure seamless interaction between the agent and the traffic environment. Figure 1 shows Traffic3D's visual display. Traffic3D allows a deep learning agent to be exposed to a comprehensive set of dynamic traffic situations, which it will potentially encounter in

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada.

Environment	Suitable for Traffic Simulation	Photo-Realistic	3D	Physics	Customisable
SUMO [6]	Yes	No	No	Yes (with restrictions [1])	Yes (with restrictions [2])
VISSIM [7]	Yes	Yes	Yes	Yes (with restrictions [3])	Yes (with restrictions [4])
TORCS [8]	Yes	Yes	Yes	Yes (with restrictions [5])	Yes (with restrictions [6])
Virtual KITTI [9]	Yes	Yes	Yes	No [7]	Yes
CHALET [10]	No	Yes	Yes	Yes	Yes
AI2-THOR [11]	No	Yes	Yes	Yes	Yes
ATARI [12]	No	No	No	No	No
DeepMind Lab [13]	No	No	Yes	No	Yes
Traffic3D	**Yes**	**Yes**	**Yes**	**Yes**	**Yes(fully)**

Table 1: Comparison between the most widely-known traffic-based and deep learning-based simulation environments.

the physical world and enables it to rapidly and safely learn the best set of policies to optimize the performance of traffic entities (such as traffic signal control) across a range of metrics traffic throughput, travel time and delay.

2 RELATED WORK

In Pell et al. [14], a comprehensive set of traffic simulators is thoroughly reviewed. The review acknowledges that none of the presently-used traffic simulation tools are capable of delivering important functionalities that are fundamental to realistic traffic simulation. The existing traffic models lack in flexibility and a detailed network model with efficient real-time traffic data collection capabilities, which are necessary to simulate heterogeneous transportation networks. In Table 1, we summarize the capabilities of the most widely-used traffic simulators and a handful of popular deep learning platforms available for training and bench-marking purposes of deep learning agents, mostly built around computer game environments. It is clear that no single simulation platform delivers all the functionality which is paramount for comprehensive traffic research and analysis.

3 ENVIRONMENT

Traffic3D gives the learning agent a natural and unstructured environment to operate in. Traffic3D is built using Unity [15], a unique game development platform used to create state-of-the-art 3D photo-realistic graphics and simulate realistic physics. Traffic3D provides a photo-realistic urban traffic environment including a variety of road junction layouts (2-way, 4-way and 5-way junctions), roads, side-walks, lane markings and traffic light poles, among others. To ensure visual variability, the roads are populated with realistic vehicle models (that encompass vehicle weight, acceleration, etc.) including hatchbacks, sedans, SUVs and emergency vehicles such as police cars, fire engines and ambulances. In addition, Traffic3D facilitates simulation of different times of the day and different seasons with distinct illumination characteristics such as sunny, cloudy, rainy and snowy days. To further add realism, the shadows rendered by different objects are dynamically cast on the surfaces within the traffic scene. The pertaining traffic situation can be captured as raw pixels from multiple viewpoints with each pixel containing a high precision depth value. This is useful when certain effects need the scene's depth to be available such as soft particles like snow and screen space ambient occlusions.

To ensure reusability, Traffic3D offers complete flexibility over its design. Traffic3D embodies a generic design of different traffic entities including vehicles and traffic infrastructure, facilitating its applicability to various traffic-related applications. Users can freely place any of the above mentioned traffic elements in a scene. To specifically evaluate the stability and generalizability of a learning agent, users can programatically create different traffic scenes with the available traffic elements depending on the application under consideration such as autonomous driving and traffic infrastructure optimization.

At the same time, we are currently exploring an alternative architecture version for Traffic3D that is likely to boost its performance in computationally-intensive machine learning scenarios. The simulation engine will be running as a server application, interacting with another server which is responsible for the machine learning. A visualization software, such as Unity, will act as a client to the simulation server.

4 CONCLUSION

The goal of Traffic3D is to facilitate building physically and visually intelligent traffic models and accelerate research in the area of traffic and transportation. It supports unique traffic-specific 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. In addition, we believe that Traffic3D has allowed us to take a step forward towards training autonomous agents using deep learning methods in more realistic settings, as the state-of-the-art deep learning agents are generally trained and validated on computer games such as the Atari suite. As a potential direction for future work, to further improve traffic objects' realism and preserve their geometry and other aspects, we intend to use photogrammetry to create real-world emulating significantly richer traffic scenarios using real-world traffic images and videos.

[1] no proper reactive control to random incidents like collisions between vehicles.
[2] does not support simulation of autonomous vehicles and does not prioritize public transport.
[3] unrealistic lane-change behavior.
[4] restrictions in customizing delay.
[5] limited sensor suite
[6] does not support road intersection simulation
[7] information not available

REFERENCES

[1] Deepeka Garg, Maria Chli, and George Vogiatzis. Deep reinforcement learning for autonomous traffic light control. In *2018 3rd IEEE International Conference on Intelligent Transportation Engineering (ICITE)*, pages 214–218. IEEE, 2018.

[2] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.

[3] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.

[4] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436, 2015.

[5] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *Intelligent Robots and Systems (IROS), 2017 IEEE/RSJ International Conference on*, pages 23–30. IEEE, 2017.

[6] Daniel Krajzewicz. Traffic simulation with sumo–simulation of urban mobility. In *Fundamentals of traffic simulation*, pages 269–293. Springer, 2010.

[7] Martin Fellendorf and Peter Vortisch. Microscopic traffic flow simulator vissim. In *Fundamentals of traffic simulation*, pages 63–93. Springer, 2010.

[8] Bernhard Wymann, Eric Espié, Christophe Guionneau, Christos Dimitrakakis, Rémi Coulom, and Andrew Sumner. Torcs, the open racing car simulator. *Software available at http://torcs. sourceforge. net*, 4:6, 2000.

[9] Adrien Gaidon, Qiao Wang, Yohann Cabon, and Eleonora Vig. Virtual worlds as proxy for multi-object tracking analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4340–4349, 2016.

[10] Claudia Yan, Dipendra Misra, Andrew Bennnett, Aaron Walsman, Yonatan Bisk, and Yoav Artzi. Chalet: Cornell house agent learning environment. *arXiv preprint arXiv:1801.07357*, 2018.

[11] Eric Kolve, Roozbeh Mottaghi, Daniel Gordon, Yuke Zhu, Abhinav Gupta, and Ali Farhadi. Ai2-thor: An interactive 3d environment for visual ai. *arXiv preprint arXiv:1712.05474*, 2017.

[12] Marc G Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013.

[13] Charles Beattie, Joel Z Leibo, Denis Teplyashin, Tom Ward, Marcus Wainwright, Heinrich Küttler, Andrew Lefrancq, Simon Green, Víctor Valdés, Amir Sadik, et al. Deepmind lab. *arXiv preprint arXiv:1612.03801*, 2016.

[14] Andreas Pell, Andreas Meingast, and Oliver Schauer. Trends in real-time traffic simulation. *Transportation research procedia*, 25:1477–1484, 2017.

[15] Unity3d.Com. Image, build and succeed with unity.