The 23rd IEEE International Conference on

Intelligent Transportation Systems

Tutorial Proposal

Deep Reinforcement Learning for Traffic Signal Control

Duration

Half Day

Organizers

Our tutorial is organized by an interdisciplinary team of leading researchers in transportation and computer science research. We hope this team can bring people with different background to discuss the transportation challenges and solutions from a global view:

Hua Wei is currently a Ph.D. candidate of the College of Information Sciences and Technology at the Pennsylvania State University. He got his B.Eng. and M.Eng. degree from Beihang University. His research interests lie in the field of spatio-temporal data mining and reinforcement learning, especially reinforcement learning for intelligent transportation. He has published over 10 papers on top conferences and journals such as KDD, AAAI, CIKM, WWW and TKDD. He has served as a program committee member in major machine learning and data mining conferences such as AAAI, CIKM, WWW and IJCAI.

Zhenhui Li is an Associate Professor of Information Sciences and Technology at the Pennsylvania State University. She is Haile family early career endowed professor. Prior to joining Penn State, she received her Ph.D. degree in Computer Science from University of Illinois Urbana-Champaign in 2012, where she was a member of data mining research group. Her research has been focused on mining spatial-temporal data with applications in ecology, environment, social science, urban computing, and transportation. She took sabbatical leave in Hangzhou China 2018/2019 and has been working on city brain research for the past year.

Vikash Gayah is an assistant professor in the Department of Civil and Environmental Engineering at The Pennsylvania State University. Dr. Gayah's research focuses on urban mobility, traffic operations, traffic flow theory, public transportation, and traffic safety. His research approach includes a combination of analytical models, micro-simulations and empirical analysis of transportation data. He has authored over 35 articles that have either been published or accepted for publications in peer-reviewed journals, 45 conference papers, and 11 research reports to sponsors. Dr. Gayah currently serves as an editorial advisory board member

[1].2019 Urban Mobility Report.
[2].2012 Urban Mobility Report.

of two leading transportation journals, Transportation Research Part B: Methodological and Transportation Research Part C: Emerging Technologies, and is a member of the Transportation Research Board's Committee on Traffic Flow Theory and Characteristics (AHB 45). He has been recognized with multiple awards, including the Dwight D. Eisenhower Transportation Fellowship, Gordon F. Newell Award for Excellence in Transportation Science, University of California Transportation Center Student of the Year Award and most recently the New Faculty Award by the Council of University Transportation Centers.

Aim of the Tutorial

Traffic congestion is a growing problem that continues to plague urban areas with negative outcomes to both the traveling public and society as a whole. These negative outcomes will only grow over time as more people flock to urban areas. In 2017, traffic congestion costs Americans over $166 billion in lost productivity and wasted over 3.3 billion gallons of fuel.[1] Traffic congestion was also attributed to over 56 billion pounds of carbon dioxide emissions in 2011.[2] Mitigating congestion would have significant economic, environmental, and societal benefits. Signalized intersections are one of the most prevalent bottleneck types in urban environments, and thus traffic signal control plays a vital role in urban traffic management.

There have been many promising methods developed for the selection of timings at traffic signals. The typical approach that transportation researchers take is to cast traffic signal control as an optimization problem under certain assumptions about the traffic model, e.g., vehicles come in a uniform and constant rate. Another promising avenue appears to be Reinforcement Learning. Reinforcement Learning can directly learn from the observed data, by first taking actions to change the signal plans and then learning from the outcomes. In essence, an RL-based traffic signal control system observes the traffic condition first, then generates and executes different actions (i.e., traffic signal plans). It will then learn and adjust the strategies based on the feedback from the environment. Recent advances in RL, especially deep RL, offer the opportunity to efficiently work with high dimensional input data (like images), where the agent can learn a state abstraction and a policy approximation directly from its input states. A series of related studies using deep RL for traffic signal control have appeared in the past few years. The advantage for RL applied to signal control is, it can directly learn from the observed data without making unrealistic assumptions about the traffic model.

This tutorial is to provide an overview of the recent development in RL and provide a hands-on experience for RL-based traffic signal control approaches, including both controlling a single intersection and multiple intersections. In this tutorial, we first introduce the formulation of traffic light control problems under RL, and then classify and discuss the current RL control methods from different aspects: agent formulation, policy learning approach, and coordination strategy when facing multiple intersections. In the third section, we provide hands-on experience on how to use the simulators to enable RL for traffic signal control. Specifically, we provide the experimental setups and detailed process for RL-based traffic signal control problems, including both single intersection and multi-intersection control. We then discuss some future research directions.

Topics of Interest

Deep reinforcement learning, traffic signal control, intelligent transportation

Intended Audience

In this tutorial, we aim to answer the question: how can we utilize reinforcement learning towards a more intelligent traffic signal control system? While traffic signal control is not a new topic, especially in the field of transportation research, existing transportation is not effectively exploring the recent advances in reinforcement learning and multi-agent reinforcement learning.

We see active movements for RL in traffic signal control research. In the area of computer science, recent conferences like KDD, AAAI, NeurIPS all witness the relative publications on the RL-based traffic signal control methods. In the area of transportation, traffic optimization and control with traffic signals have always been one of the main topics discussed in ITSC and Transportation Research.

This tutorial would like to bring together the researchers who are interested in the exciting RL techniques to solve transportation problems.

This tutorial is different from previously held tutorials on pure introduction to simulators or reinforcement learning. We plan to have a more focused theme and provide hands-on experiences on RL-based traffic signal control methods. We believe this focused theme can enable more in-depth discussions on related topics and techniques.

Program

We would like to have a novel program setting with both lecture-style introductions and hands-on experiences. We believe such a combination of the tutorial will ensure the quality and diverse participation of the tutorial.

08:00-08:30 am: Introduction to traffic signal control

- Terms and objectives in traffic signal control
- Conventional traffic signal control methods
    - Fixed-time control
    - Adaptive control
    - Practical systems: TRANSYT, SCATS, RHODES

08:30-9:30 am: Introduction to reinforcement learning (RL)

- Introduction of the principles of RL and how they work with the following concepts explained: reward, state, action, Markov decision process. We will also include high-level examples.
- Frequently used methods: value-based RL [32], policy-based RL [1].

- Deep RL. value-based RL (Deep Q-Learning [32]), policy-based RL (Deep deterministic RL [3]) and other frequently used methods [2-7].

09:30-10:00 am: RL for traffic signal control

- Formulating traffic signal control problem as a reinforcement learning problem

10:00-10:30 am: Coffee break

10:30 - 11:30 am: RL for traffic signal control (cont.)

- Design of reward
    - Different measures of reward (e.g., queue length, waiting time) and their pros and cons will be discussed
- Design of state
    - Different measures of state (e.g., image, queue length) and their pros and cons will be discussed
- Design of action
    - Different measures of state (e.g., cycle-based, cycle-free) and their pros and cons will be discussed
- Design of RL models
    - Different models in RL (e.g., Policy-based, value based) and their pros and cons will be discussed
- Coordination of multiple intersections
    - Through reward design (e.g., PressLight[21])
    - Through state design (e.g., CoLight[22])
    - Through model design (e.g., hierarchical control, central control)

11:30-12:00 pm: Hands-on implementation with CityFlow

- Step-by-step instructions on setting up the Python environment with a highly efficient traffic simulator CityFlow[33]. We will introduce how to implement the pipeline of RL with CityFlow.

12:00-12:30 pm: Open problems in RL for traffic signal control

- Simulating environment to real world
    - Building a real simulator (e.g., learning to simulate)
    - Learning from small trials (e.g., meta learning and learning from experts)
- Other issues: benchmarks, interpretability, safety issues, etc.

Tentative papers to cover:

1. Sutton, Richard S., David A. McAllester, Satinder P. Singh, and Yishay Mansour. "Policy gradient methods for reinforcement learning with function approximation." In Advances in neural information processing systems, pp. 1057-1063. 2000.
2. Mnih, Volodymyr, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. "Asynchronous methods for deep reinforcement learning." In International conference on machine learning, pp. 1928-1937. 2016.
3. Lillicrap, Timothy P., Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. "Continuous control with deep reinforcement learning." arXiv preprint arXiv:1509.02971 (2015).
4. Hado V. Hasselt, “Double Q-learning”, NIPS 2010

5. Hado van Hasselt, Arthur Guez, David Silver, “Deep Reinforcement Learning with Double Q-learning”, AAAI 2016
6. Ziyu Wang, Tom Schaul, Matteo Hessel, Hado van Hasselt, Marc Lanctot, Nando de Freitas, “Dueling Network Architectures for Deep Reinforcement Learning”, arXiv preprint, 2015
7. Schaul, Tom, et al. "Prioritized experience replay." arXiv preprint arXiv:1511.05952 (2015).
8. Rashid, T., Samvelyan, M., Schroeder, C., Farquhar, G., Foerster, J., and Whiteson, S. (2018). QMIX: Monotonic value function factorisation for deep multi-agent reinforcement learning. In Proceedings of the 35th International Conference on Machine Learning, pages 4295–4304.
9. Sunehag, P., Lever, G., Gruslys, A., Czarnecki, W. M., Zambaldi, V., Jaderberg, M., Lanctot, M., Sonnerat, N., Leibo, J. Z., Tuyls, K., et al. (2017). Value-decomposition networks for cooperative multi-agent learning. arXiv preprint arXiv:1706.05296.
10. Van der Pol, E. and Oliehoek, F. A. (2016). Coordinated deep reinforcement learners for traffic light control. Proceedings of Learning, Inference and Control of Multi-Agent Systems.
11. Li Li, Ding Wen, and Danya Yao. A survey of traffic control with vehicular communications. IEEE TITS, 2014.
12. Li Li, Yisheng Lv, and Fei-Yue Wang. Traffic signal timing via deep reinforcement learning. IEEE/CAA Journal of Automatica Sinica, 2016.
13. Yuanhao Xiong, Guanjie Zheng, Kai Xu, and Zhenhui Li. Learning traffic signal control from demonstrations. In CIKM, 2019.
14. Kok-Lim Alvin Yau, Junaid Qadir, Hooi Ling Khoo, et al. A survey on reinforcement learning models and algorithms for traffic signal control. ACM Computing Survey, 2017.
15. Xinshi Zang, Huaxiu Yao, Guanjie Zheng, Nan Xu, Kai Xu, and Zhenhui Li. MetaLight: Value-based Meta-reinforcement Learning for Online Universal Traffic Signal Control. In AAAI, 2020.
16. Zhi Zhang, Jiachen Yang, and Hongyuan Zha. Integrating independent and centralized multi-agent reinforcement learning for traffic signal network optimization. arXiv preprint, 2019.
17. Guanjie Zheng, Yuanhao Xiong, Xinshi Zang, Jie Feng, Hua Wei, et al. Learning Phase Competition for Traffic Signal Control. In CIKM, 2019.
18. Guanjie Zheng, Xinshi Zang, Nan Xu, Hua Wei, Zhengyao Yu, et al. Diagnosing Reinforcement Learning for Traffic Signal Control. arXiv preprint, 2019.
19. Yanan Wang, Tong Xu, Xin Niu, Chang Tan, Enhong Chen, and Hui Xiong. STMARL: A SpatioTemporal Multi-Agent Reinforcement Learning Approach for Traffic Light Control. arXiv preprint, 2019.
20. Hua Wei, Guanjie Zheng, Huaxiu Yao, and Zhenhui Li. IntelliLight: A Reinforcement Learning Approach for Intelligent Traffic Light Control. In KDD, 2018.
21. Hua Wei, Chacha Chen, Guanjie Zheng, Kan Wu, Vikash Gayah, Kai Xu, and Zhenhui Li. PressLight: Learning Max Pressure Control to Coordinate Traffic Signals in Arterial Network. In KDD, 2019.
22. Hua Wei, Nan Xu, Huichu Zhang, Guanjie Zheng, et al. CoLight: Learning Network-level Cooperation for Traffic Signal Control. In CIKM, 2019.
23. MA Wiering. Multi-agent reinforcement learning for traffic light control. In ICML, 2000.
24. Tong Thanh Pham, Tim Brys, Matthew E Taylor, Tim Brys, et al. Learning coordinated traffic light control. In AAMAS, 2013.
25. Tomoki Nishi, Keisuke Otaki, Keiichiro Hayakawa, and Takayoshi Yoshimura. Traffic signal control based on reinforcement learning with graph convolutional neural nets. In ITSC. IEEE, 2018.

26. Seyed Sajad Mousavi, Michael Schukat, and Enda Howley. Traffic light control using deep policy-gradient and value-function-based reinforcement learning. Intelligent Transport Systems, 2017.
27. Patrick Mannion, Jim Duggan, and Enda Howley. An experimental review of reinforcement learning algorithms for adaptive traffic signal control. In Autonomic Road Transport Support Systems. 2016.
28. Peter Henderson, Riashat Islam, Philip Bachman, Joelle Pineau, et al. Deep reinforcement learning that matters. In AAAI, 2018.
29. Wade Genders and Saiedeh Razavi. Using a deep reinforcement learning agent for traffic signal control. arXiv preprint, 2016.
30. Chacha Chen, Hua Wei, Nan Xu, Guanjie Zheng, et al. Toward A Thousand Lights: Decentralized Deep Reinforcement Learning for Large-Scale Traffic Signal Control. In AAAI, 2020.
31. Noe Casas. Deep deterministic policy gradient for urban traffic light control. arXiv preprint, 2017.
32. Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, et al. Human-level control through deep reinforcement learning. Nature, 2015.
33. Zhang, Huichu, et al. "Cityflow: A multi-agent reinforcement learning environment for large scale city traffic scenario." *The World Wide Web Conference*. 2019.

Tentative List of Presenters

Hua Wei
Zhenhui Li
Vikash Gayah

Equipment

Each presenter will bring their own laptop. The equipment to be needed is a projector, power socket, and laptop adapters. The attendees are not required to bring any equipment.