Connectivity matters! Traffic flow optimisation with (deep) reinforcement learning

Renata Borovica-Gajic*

*joint work with Egemen Tanin, Shanika Karunasekera, Majid Sarvi, Hairuo Xie, and Udesh Gunarathna

THE UNIVERSITY OF MELBOURNE

Why traffic flow optimisation?

Traffic congestion costs the Australian economy $19 billion annually and is rising to an expected $40 billion by 2031*

Reduced congestion: positive impact on economy, environment and wellbeing

*https://www.infrastructureaustralia.gov.au/publications/urban-transport-crowding-and-congestion

Are autonomous vehicles solution?

- Connectivity provides useful traffic telemetry
- Ability to communicate with a smart/intelligent traffic infrastructure
- Can provide more complex instructions to vehicles

Connected Autonomous Vehicles (CAVs) are not just about self-driving: connectivity is what we need…

Intelligent transportation system (ITS) to the rescue

- Intelligent transportation system (ITS) is a combination of CAVs and intelligent traffic infrastructure
- Vehicles can send GPS position and other vehicle telematics via the communication infrastructure
- Our industry partner (Telstra) has a 4G network covering 99.4% of Australian population
- The 4G network is used by Sydney Coordinated Adaptive Traffic System (SCATS)

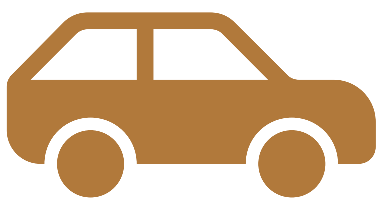

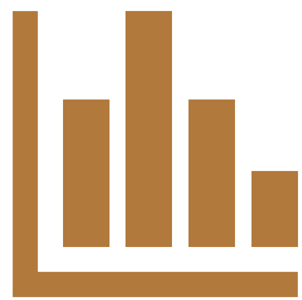

Leverage data to optimise traffic flow…

Data-driven traffic flow optimisation

- **Autonomous intersection management**
  - [IEEE IVS’22, SIGSPATIAL’22]
- **Dynamic lane reversal and road network optimisation**
  - [ECML/PKDD’20, TIST’23]
- **Adaptive route guidance**
  - [SIGSPATIAL’20, SSDBM’20, SIGSPATIAL’22]

Data-driven traffic flow optimisation

- **Autonomous intersection management**
  - [IEEE IVS'22, SIGSPATIAL'22]

- **Dynamic lane reversal and road network optimisation**
  - [ECML/PKDD'20, TIST'23]

- **Adaptive route guidance**
  - [SIGSPATIAL'20, SSDBM'20, SIGSPATIAL'22]

Autonomous Intersection Management (AIM)

- Avoid traffic signals and cross the intersection collaboratively
- Connected (autonomous) vehicles and smart infrastructure present such opportunities
- Vehicles coordinate speed and arrival times for collision free traversal through the intersection

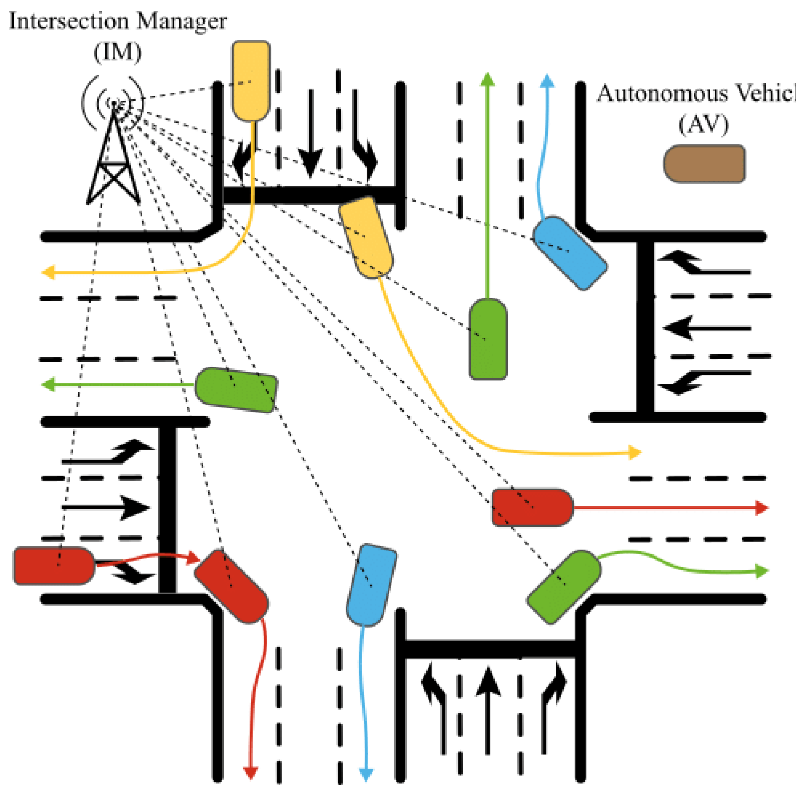

www.mdpi.com/1424-8220/22/6/2217

Learning vehicle schedule and their speed

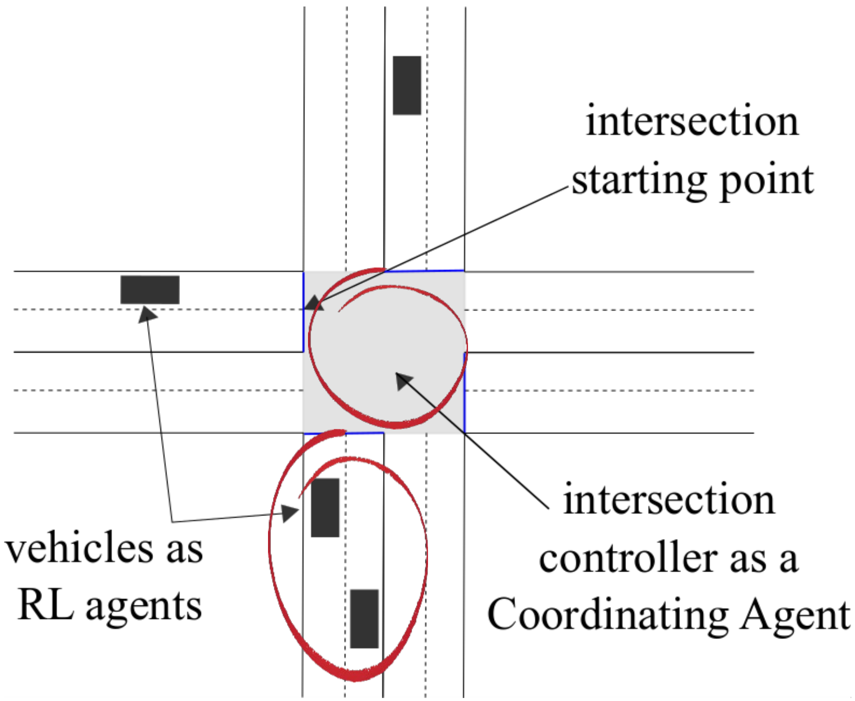

Models the problem as a **multi-agent** solution to overcome the computational complexity

1. **Coordinating Agent** uses a modified polling-based system to find an optimized order to travel and assign a schedule for each vehicle

2. **RL Agents** use a novel reinforcement learning algorithm to compute the best trajectory to adhere to the schedule (**trajectory** and **cruise** control)

Implementation testbed: e-SMARTS

[SIGSPATIAL'22]

An easily extendable traffic simulation system named e-SMARTS to allow researchers to experiment with novel data-driven traffic management algorithms in a setup that mimics real-world traffic conditions.

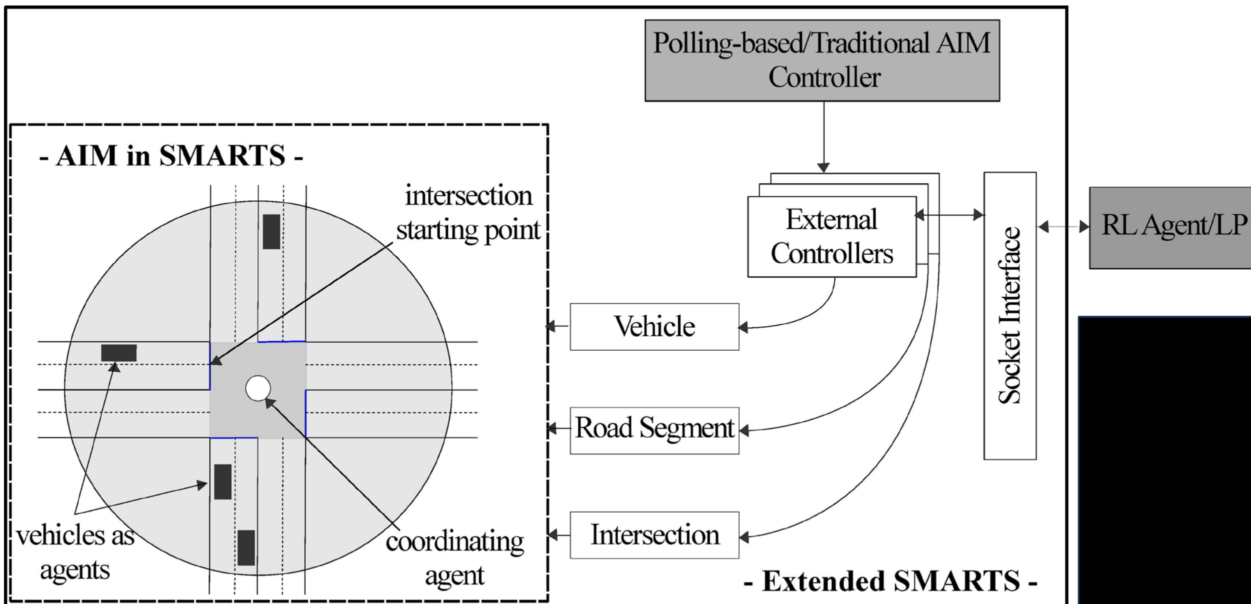

Overview of e-SMARTS

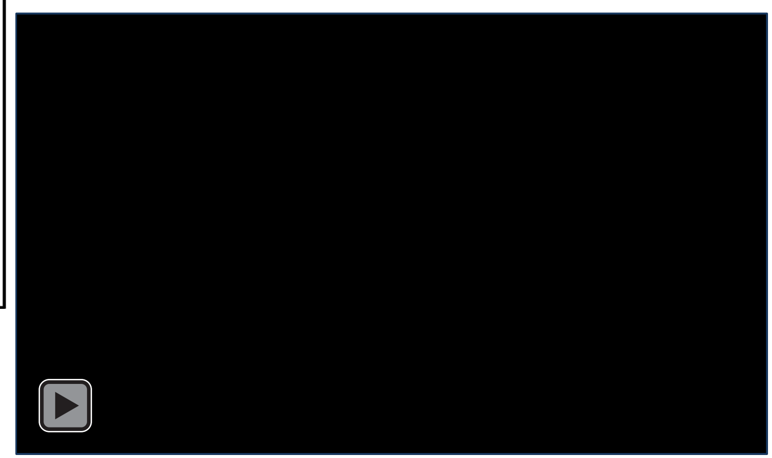

*https://projects.eng.unimelb.edu.au/smarts/

Experimental results: performance and safety

[IEEE IVS'22]

Experiments were carried out in a four-legged intersection.
The travel time is shown in seconds

Traffic Level	Low	Mid	High
Dynamic Traffic Signals	77.59	235.73	543.18
FCFS-AIM [Dresner'08]	15.21	131.46	388.57
H-AIM [Miculescu'20]	15.25	98.09	313.34
LP-AIM [Au'12]	13.85	94.92	304.44
CMQ-AIM	**13.66**	**92.76**	**302.63**

> 1000 faster computation

The number of vehicles that violated the arrival times

Traffic Level	Low	Mid	High
H-AIM [Miculescu'20]	65	49	69
LP-AIM [Au'12]	1	66	51
CMQ-AIM	**0**	**0**	**0**

Travel times halved, with 0 violations

Data-driven traffic flow optimisation

- Autonomous intersection management
  - [IEEE IVS’22, SIGSPATIAL’22]
- **Dynamic lane reversal and road network optimisation**
  - [ECML/PKDD’20, TIST’23]
- Adaptive route guidance
  - [SIGSPATIAL’20, SSDBM’20, SIGSPATIAL’22]

Real-time lane reversal

Changing the travelling direction of lanes in road segments based on real-time traffic information

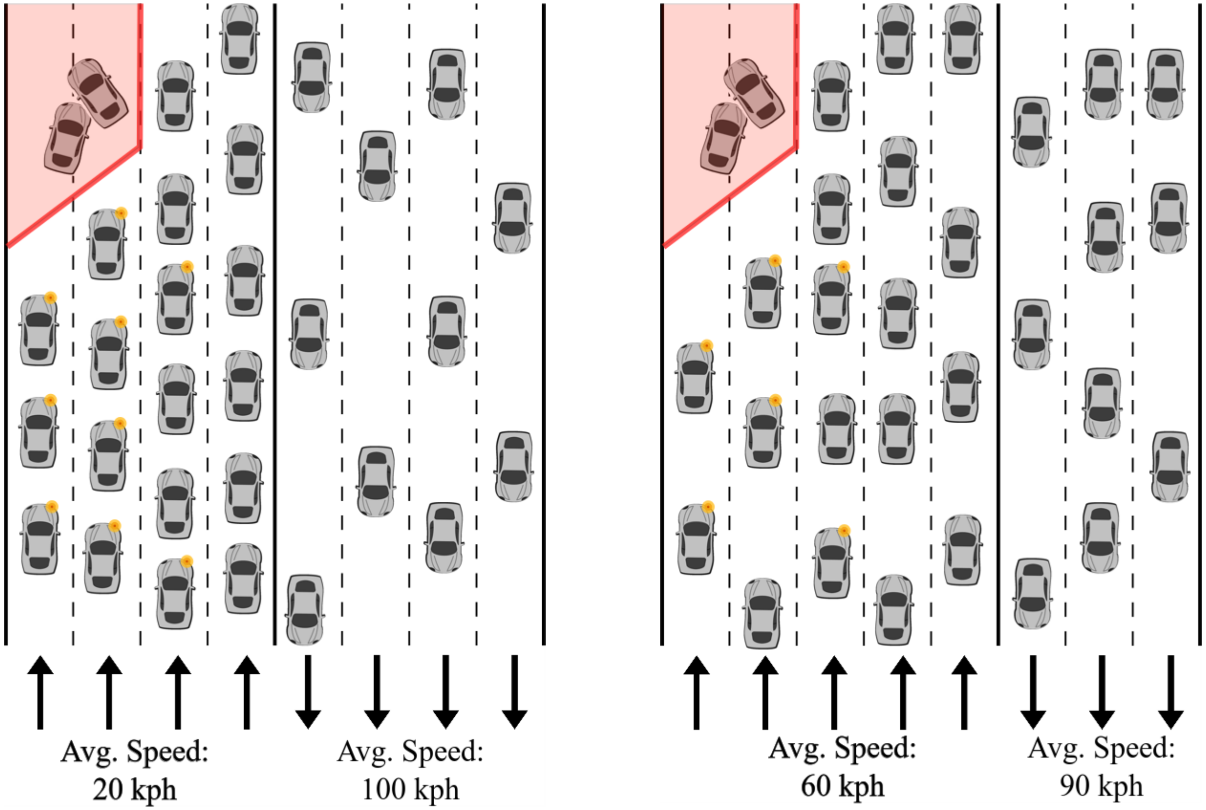

Challenges

- Existing approaches use mathematical programming to compute lane-direction allocation based on *pre-known* traffic patterns
- Inability to work with real-time data
- Computation cost is very high

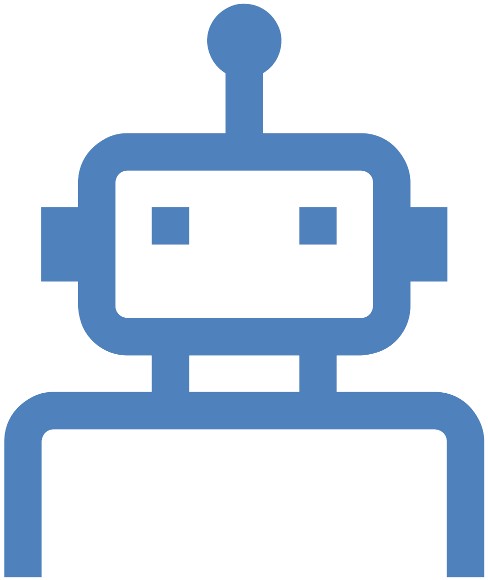

Coordinated Learning-based Lane Allocation

[ECML/PKDD'20]

Architecture of CLLA consists of RL Agents that operate at the intersection level and Coordinating Agents that evaluate the global impact of local lane-direction changes.

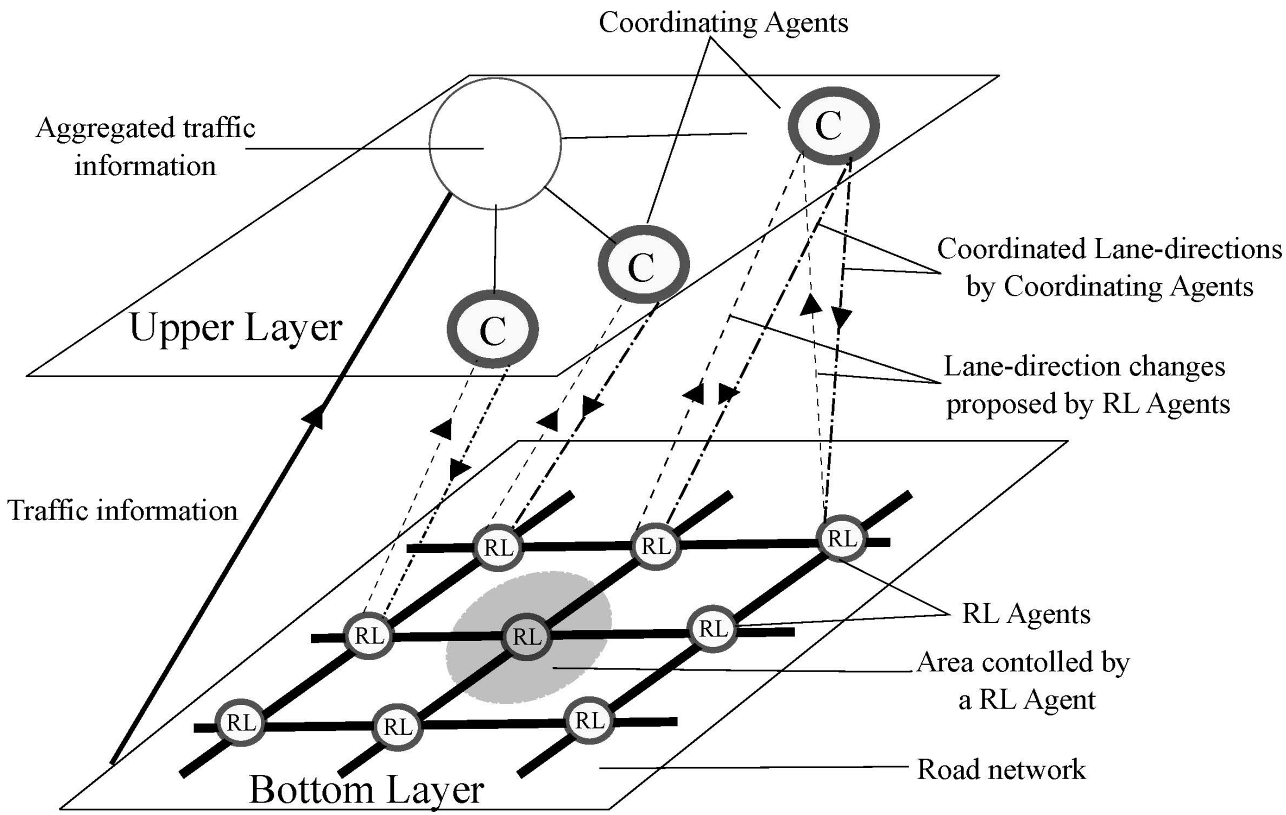

CLLA in practice

Travel times achieved in an area of the Manhattan road network using New York taxi data. The area consists of 656 road segments and 366 intersections

Baseline	Travel Time(s)		Congested Travel Time(s)		Imbalanced Travel Time(s)		DFFT	
	7am	12pm	7am	12pm	7am	12pm	7am	12pm
FTS_FR	1313.30	1350.70	1831.80	1905.36	1420.68	2021.20	13.08	16.94
DTS_FR	1237.65	1169.48	2107.30	1944.68	1253.64	1575.32	12.49	15.17
noLA	954.71	1007.81	1231.93	1414.07	1051.01	1554.80	9.19	12.18
LLA	915.76	938.46	1142.08	1276.99	759.84	1183.38	8.36	10.07
DLA	923.24	1002.55	1165.19	1386.44	859.37	1450.40	8.46	12.05
CLLA+	893.79	925.42	1019.34	1214.69	726.04	834.13	8.04	10.05
DTS	815.03	849.19	1325.81	1320.26	957.10	1124.46	7.65	9.43
DTS_CLLA+	**764.95**	**802.79**	**1014.94**	**1196.21**	**666.17**	**706.26**	**7.12**	**9.01**

Travel times halved for imbalanced and congested traffic

Highest benefit: combining dynamic traffic signals and lane reversal

Impact of vehicle connectivity

CLLA in a mix of human-driven and fully autonomous vehicles

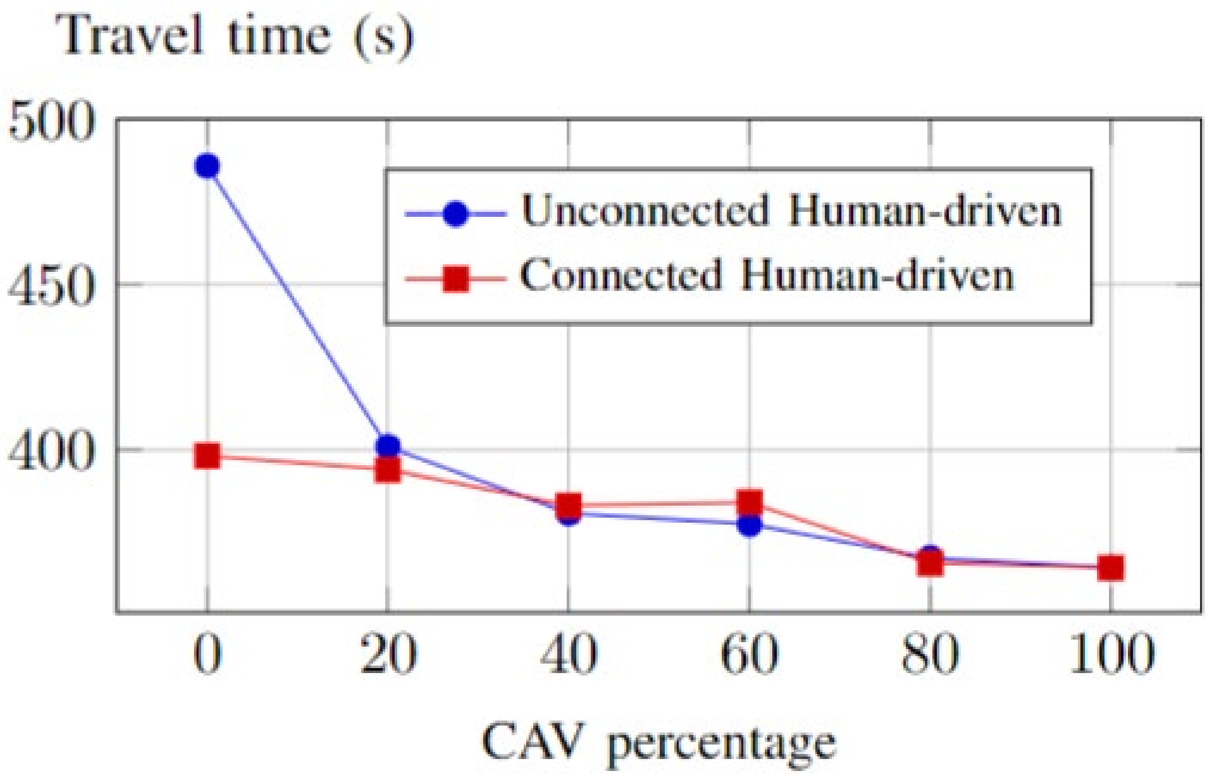

Even a small portion of connected vehicles is sufficient to observe the benefits

Data-driven traffic flow optimisation

- **Autonomous intersection management**
  - [IEEE IVS’22, SIGSPATIAL’22]
- **Dynamic lane reversal and road network optimisation**
  - [ECML/PKDD’20, TIST’23]
- **Adaptive route guidance**
  - [SIGSPATIAL’20, SSDBM’20, SIGSPATIAL’22]

Common Optimization Problems in Transport

- Traffic optimization applications:
  - Autonomous ride-sharing
  - Autonomous goods delivery systems
  - Logistic planning/supply chains
  - Autonomous route planning

All these applications require solving a version of dynamic combinatorial optimization on graphs

Why "dynamicity" matters?

- Typical combinatorial optimization problems:
  - Travelling Salesmen Problem
  - Vehicle Routing Problem
- Usually consider the graph as a static graph (NP-hard)
- The dynamic version has more real-world applications but *more challenging* to solve

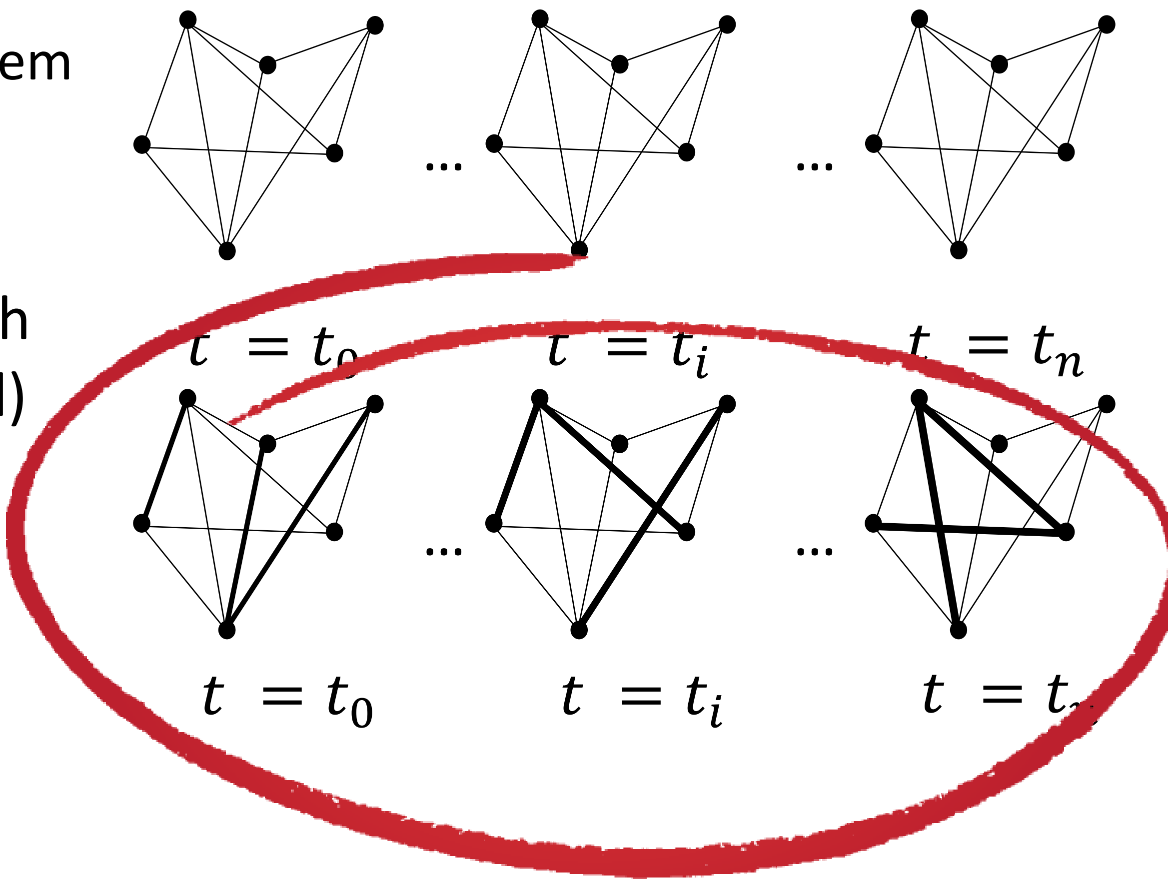

Learning-based infrastructure for dynamic combinatorial optimisation over graphs

- Caters for changes in space and time
- Lightweight and effective

Graph Temporal Attention with Reinforcement Learning (GTA-RL)

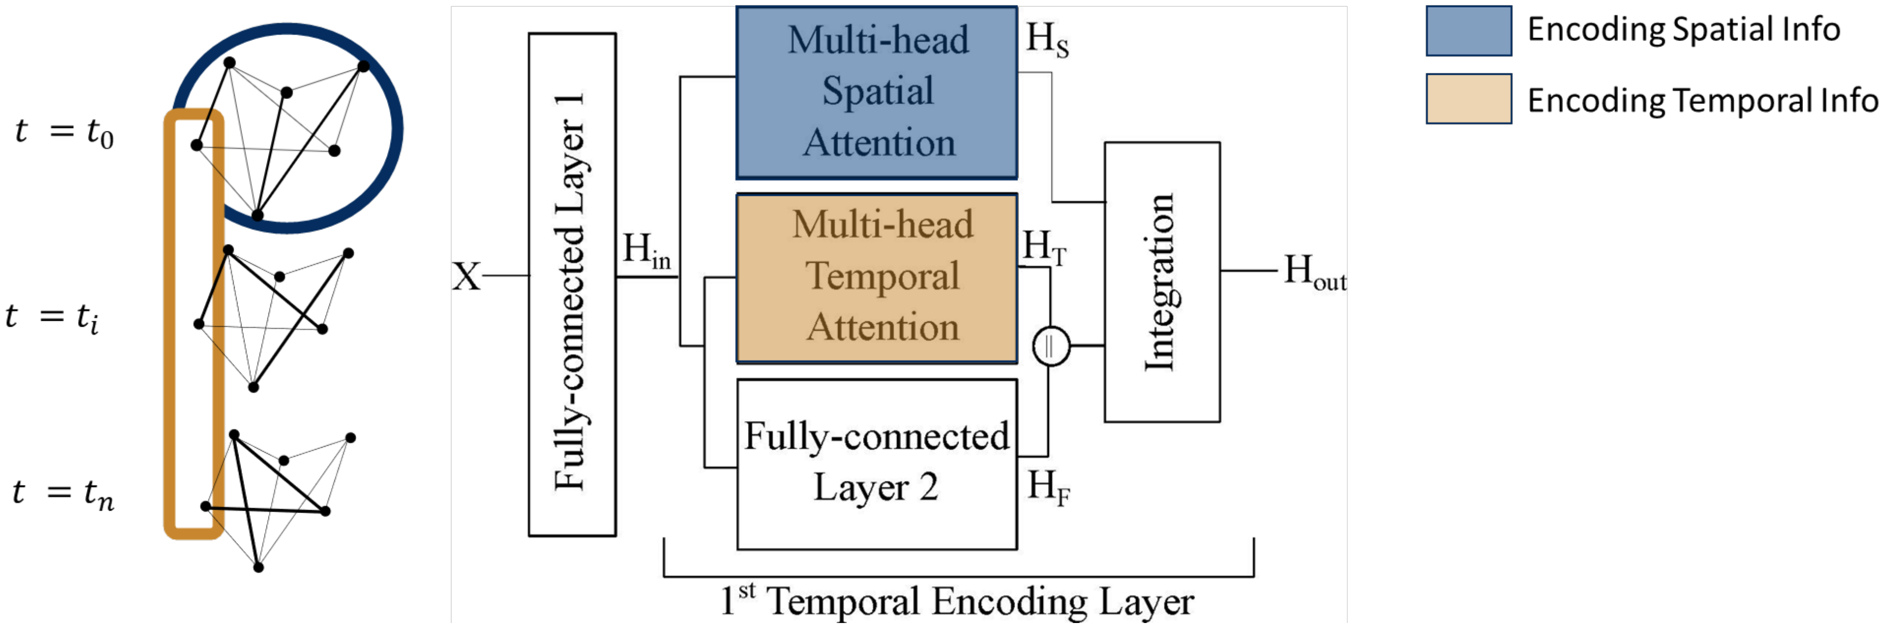

GTA-RL: Encoder

Graph Temporal Attention with Reinforcement Learning (GTA-RL)

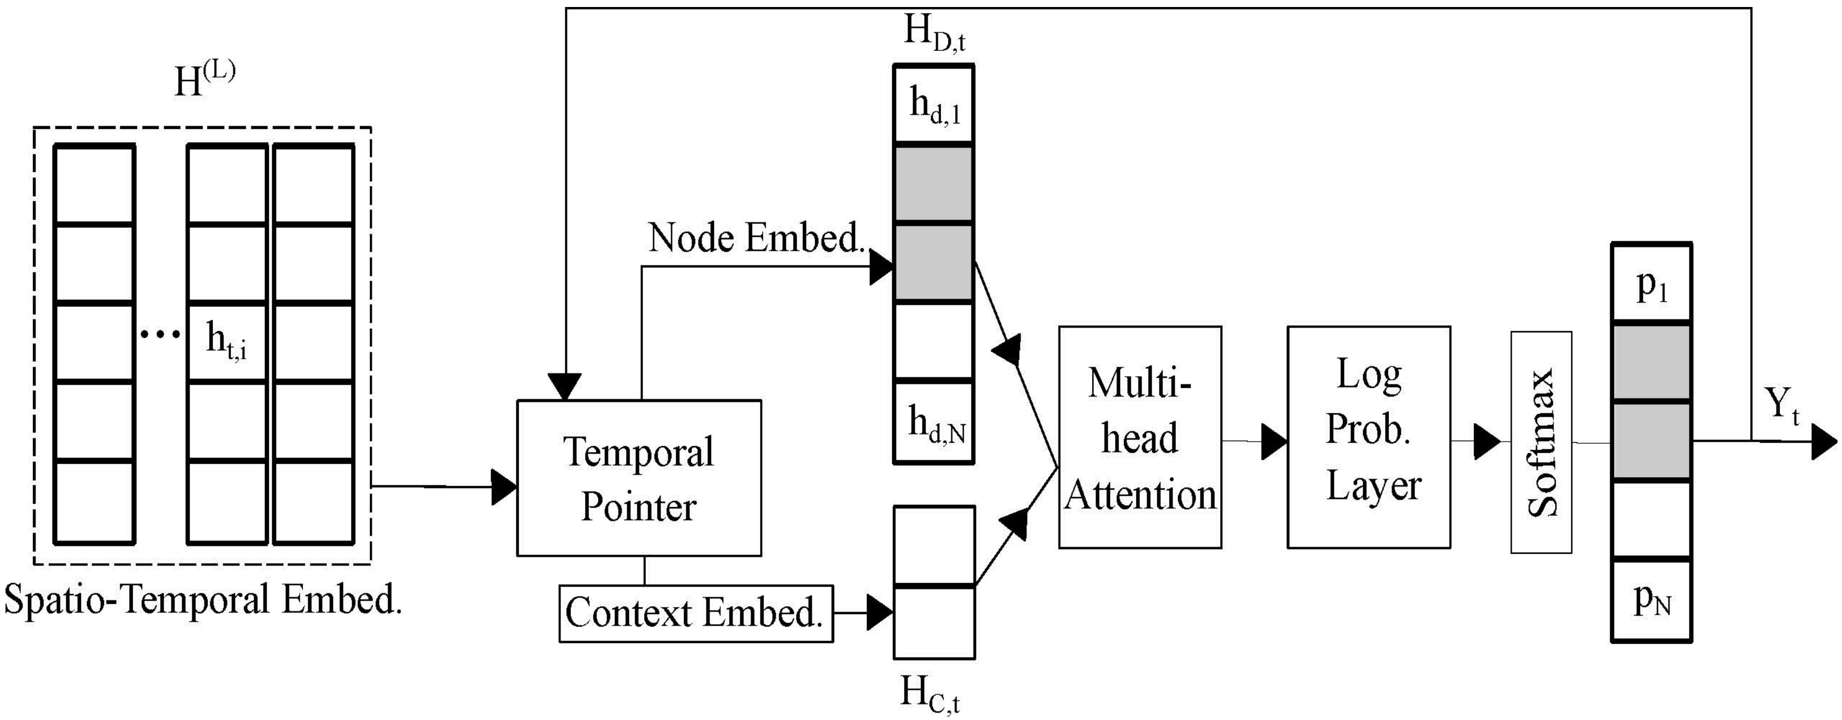

GTA-RL: Decoder

Dynamic traveling salesman problem

Gurobi optimal solver cannot scale to 50 nodes

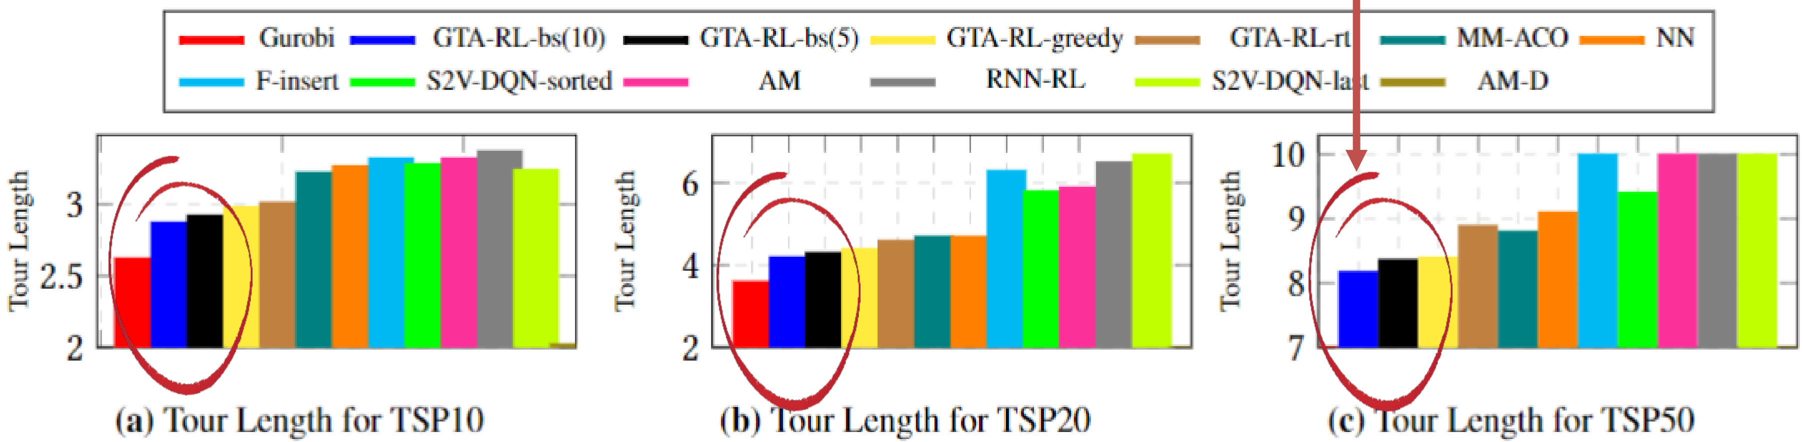

Tour length of dynamic TSP

Close to optimal solution, yet much more scalable (900sec for one instance vs 0.45sec for 100 instances)

Conclusions

- Connectivity opens up opportunities for data-driven traffic flow optimization
- Physical infrastructure is already there (SCATS, telemetry, etc)
- Recent advancements in (Deep) RL great fit for traffic
- Proposed methods needs to be **scalable** and **lightweight**
  - Optimal solvers inappropriate since traffic conditions change
- Often calls for innovative approaches in terms of computation
  - **Dynamicity**, and **multiple-objectives** need to be accommodated

THANK YOU

Q&A