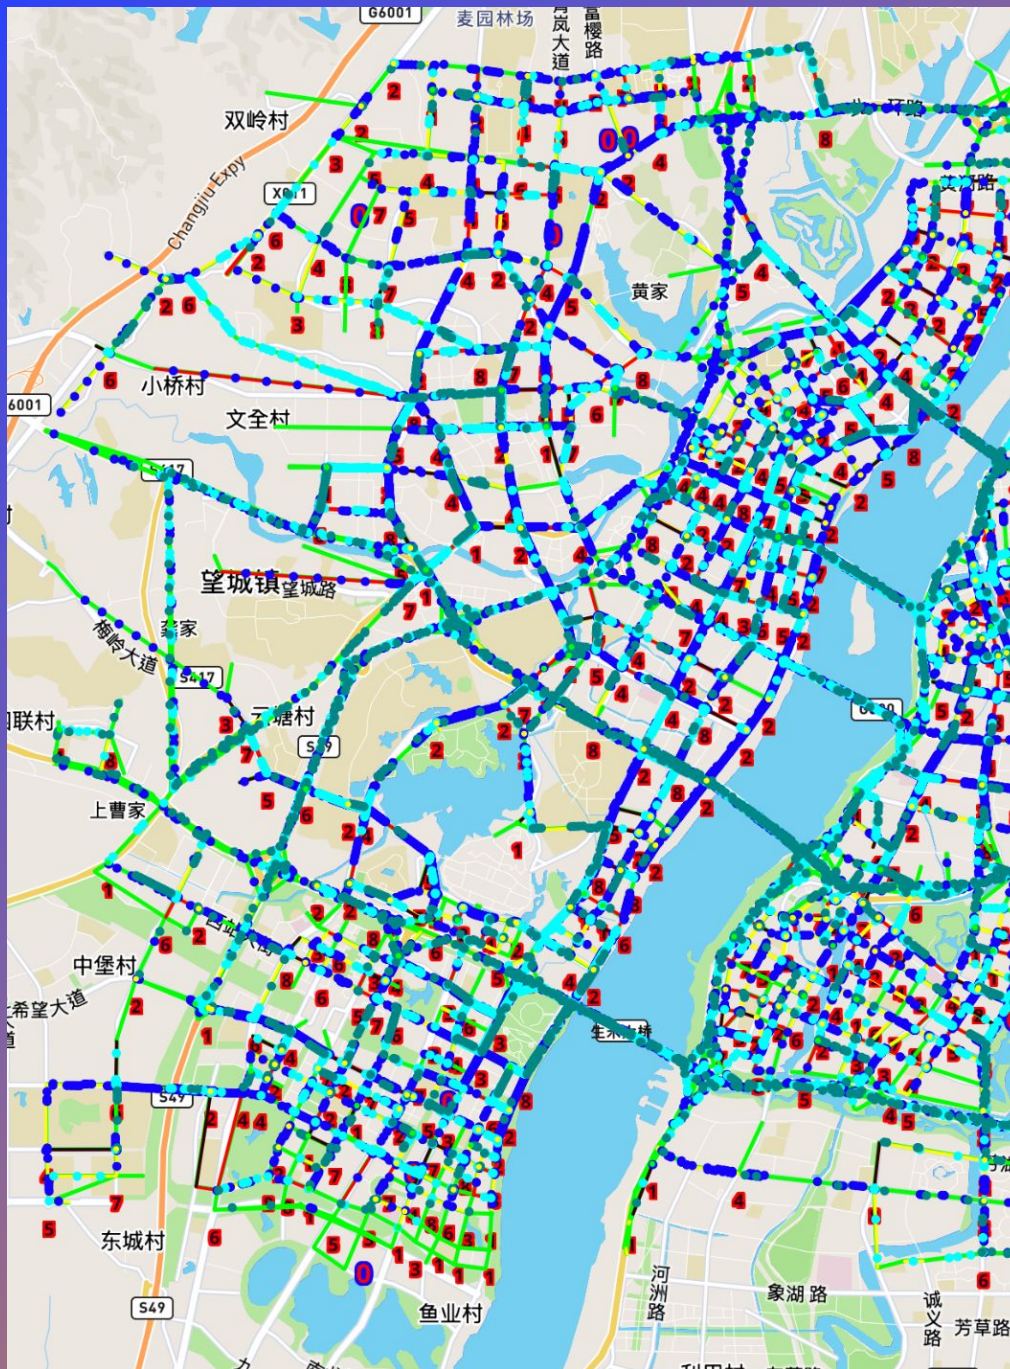


# CITY-BRAIN CHALLENGE

Phase-traffickers (GroupFour):  
Charlotte  
Faheem  
Max

---



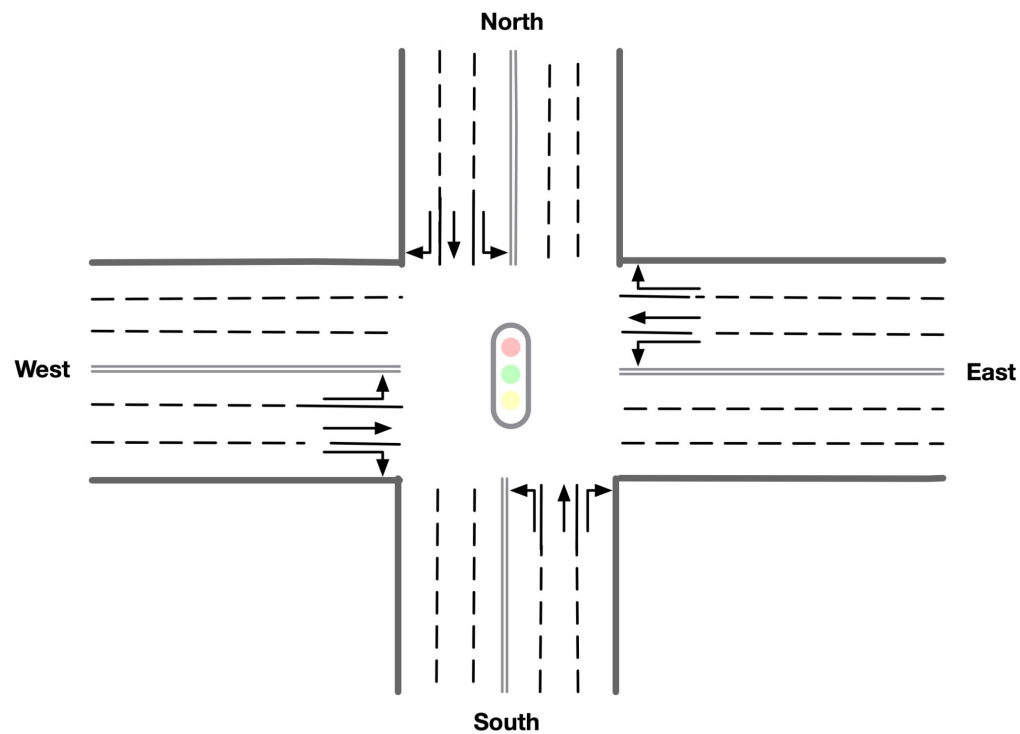


# KDD CUP

## GOAL

Build a multi-agent system that minimizes travel time of vehicles

# City Brain Challenge - KDD Cup 2021



roadnet\_1x1.txt 286 Bytes

```

1 5
2 30 120 0 1
3 31 120 1 0
4 30 121 2 0
5 29 120 3 0
6 30 119 4 0
7 4
8 0 1 30 20 3 3 1 2
9 1 0 0 0 1 0 0 1 1
10 1 0 0 0 1 0 0 1 1
11 0 2 30 20 3 3 3 4
12 1 0 0 0 1 0 0 1 1
13 1 0 0 0 1 0 0 1 1
14 0 3 30 20 3 3 5 6
15 1 0 0 0 1 0 0 1 1
16 1 0 0 0 1 0 0 1 1
17 0 4 30 20 3 3 7 8
18 1 0 0 0 1 0 0 1 1
19 1 0 0 0 1 0 0 1 1
20 1
21 0 1 3 5 7
    
```

flow\_1x1.txt 194 Bytes

```

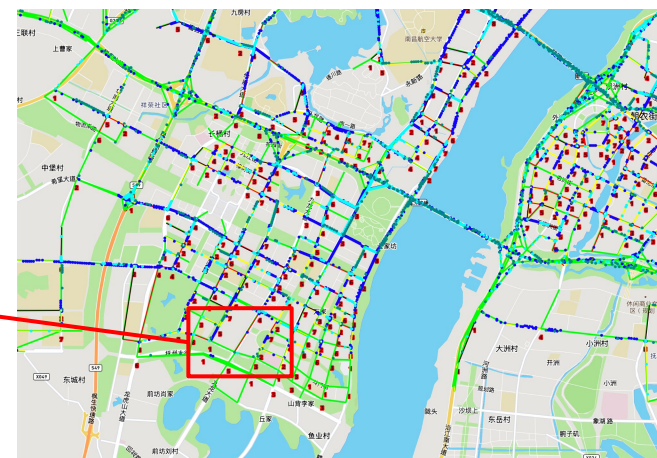
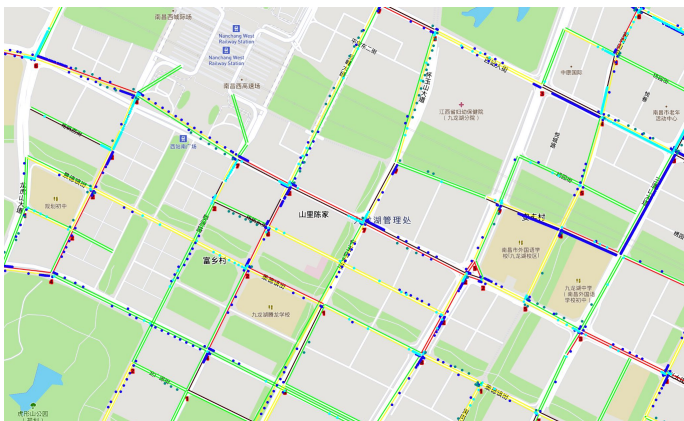
1 12
2 0 1800 40
3 2
4 2 3
5 0 1800 60
6 2
7 2 5
8 0 1800 70
9 2
10 2 7
11 0 1800 60
12 2
13 4 5
14 0 1800 50
15 2
    
```

Total Served Vehicles

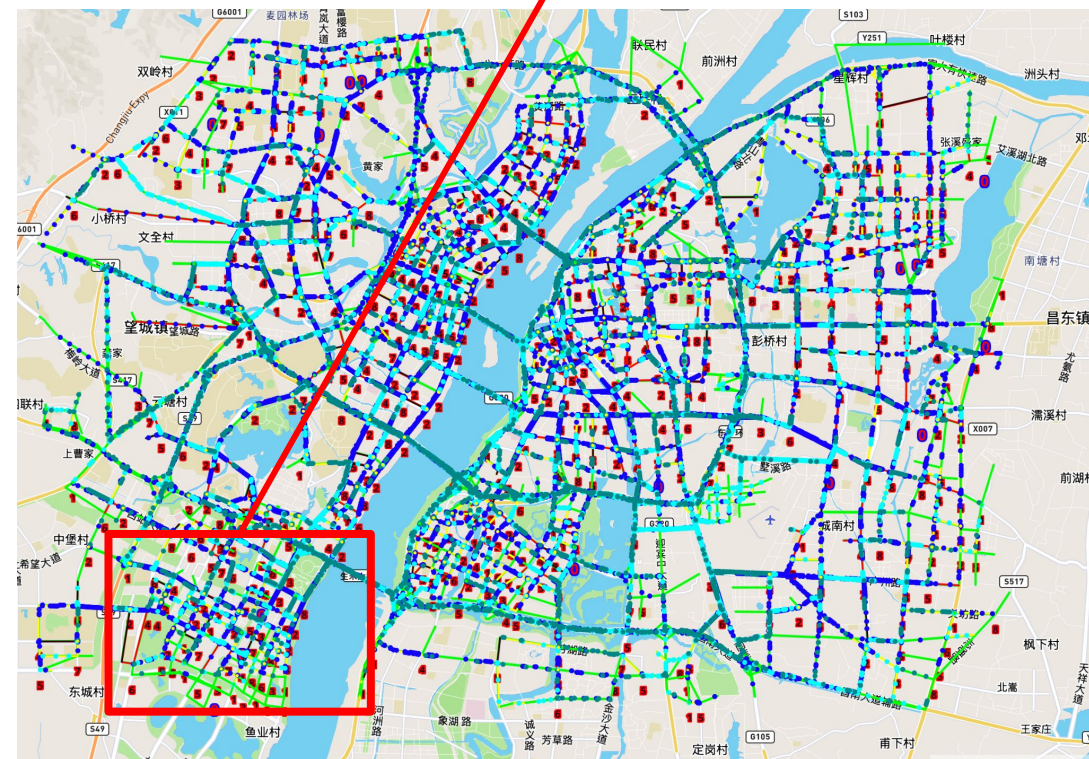
Warm Up Data Set  
= max. 1047 vehicles

Delay Index

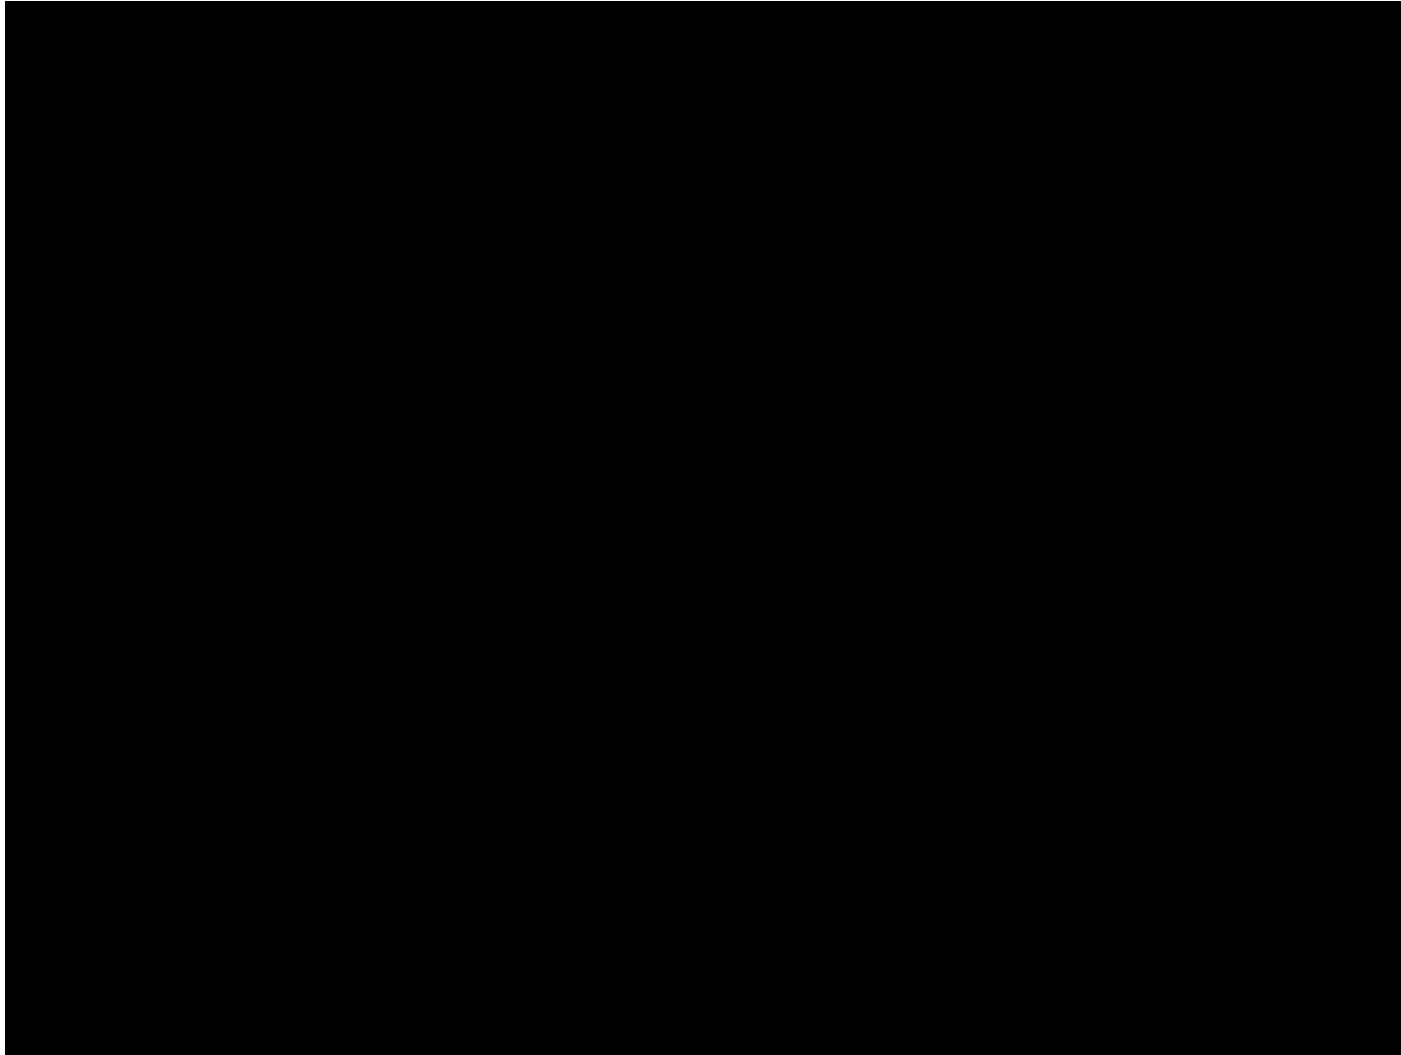
= 1 (best)



Data Set	Intersection
Warm Up	36
Round 2	2048
Round 3	92344



# Visualization



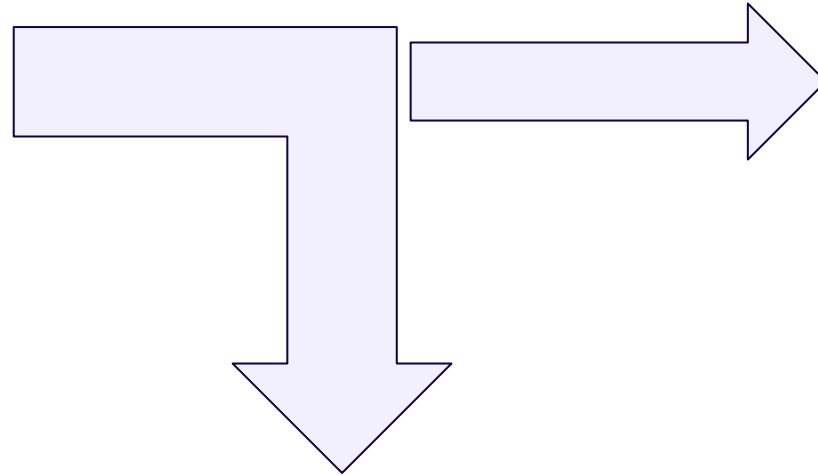
# Tasks and Goals in our Project

- Infrastructure Set Up - MLFlow & Slurm
- Baselines
  - CoLight
  - Presslight
  - FRAP
- Extensions to Baselines
  - Ape-X
  - QMix

# Infrastructure Setup



Set up LRZ CC VM  
and Slurm



SLURM:

- Image + Container for Environment Simulation  
→ Installation of Cuda
- Training

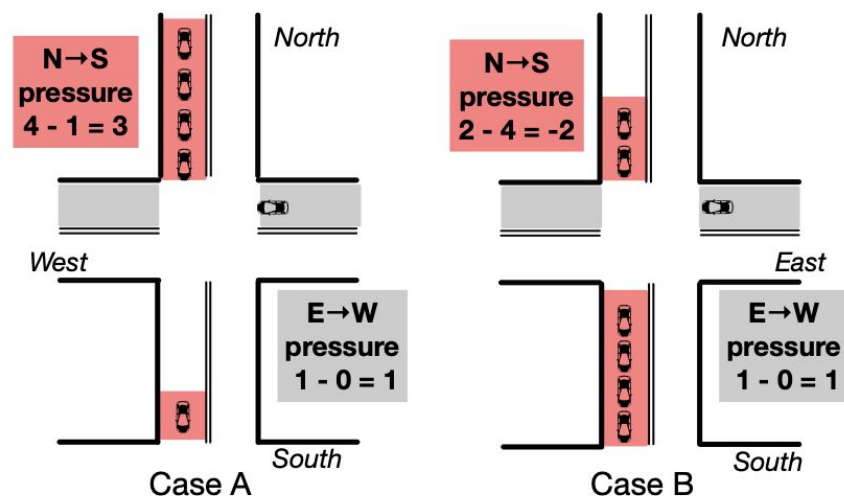
dashboard      mlflow server

	Parameters >			Metrics >			Tags
Start Time	env	env_config	multiagent	agent_timesteps_to	episode_len_mean	episode_reward_ma	trial_name
✓ 2021-07-15 12:24:16	<class '__main__...	{'simulator_cfg_f...	{'policies': {'defa...	120480	120	-810344	DQN_MultiFlowC...
✓ 2021-07-15 12:16:11	<class '__main__...	{'simulator_cfg_f...	{'policies': {'defa...	120480	120	-923140	DQN_MultiFlowC...

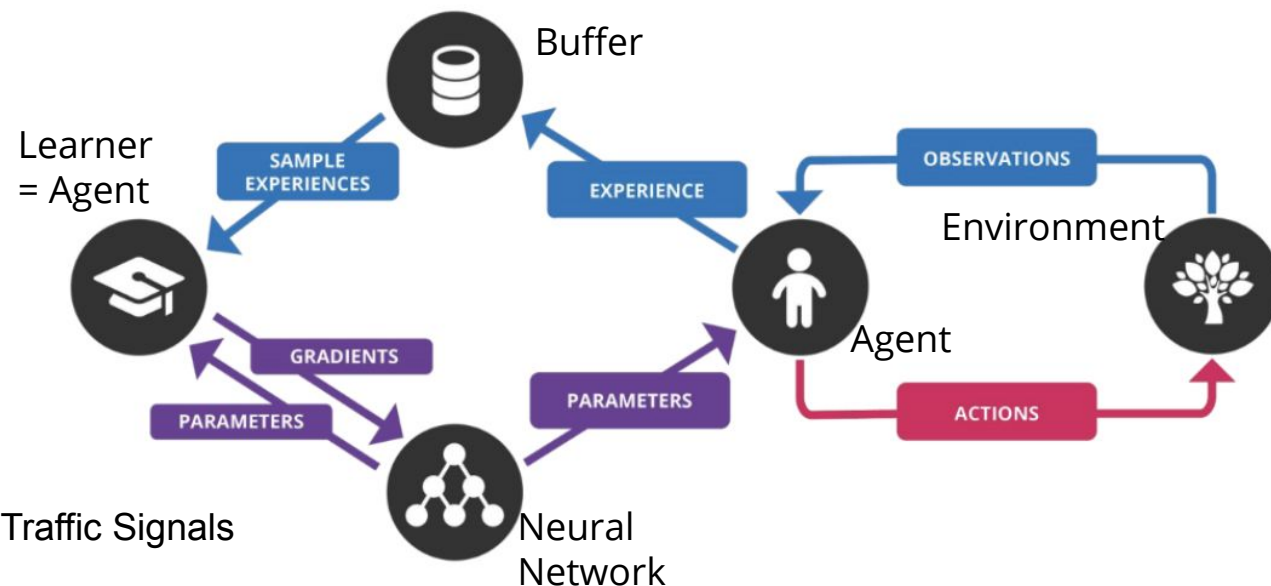
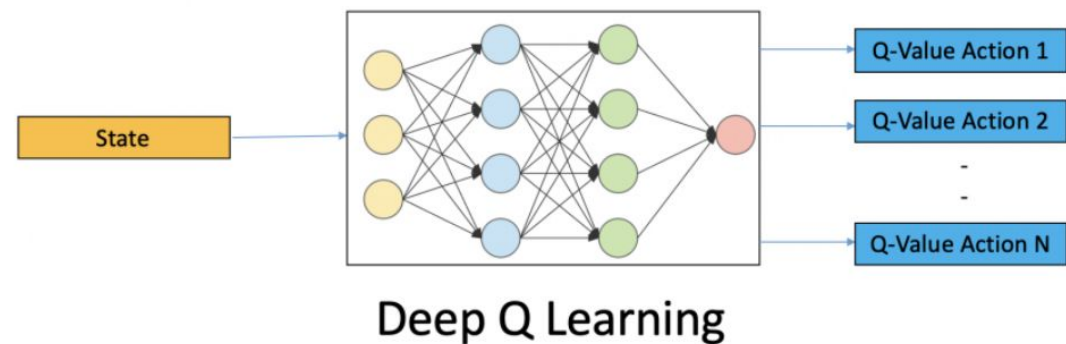
- log artifacts, hyperparameters, metrics

# Baseline: Presslight

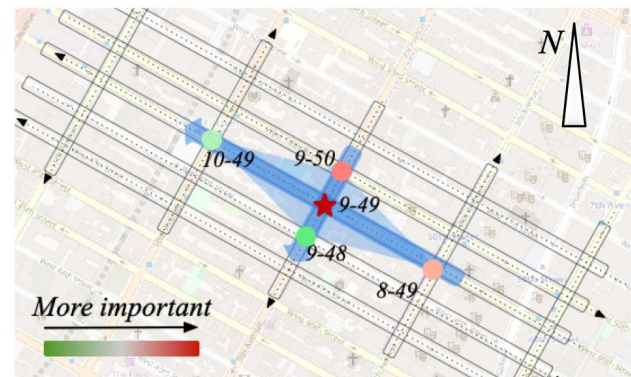
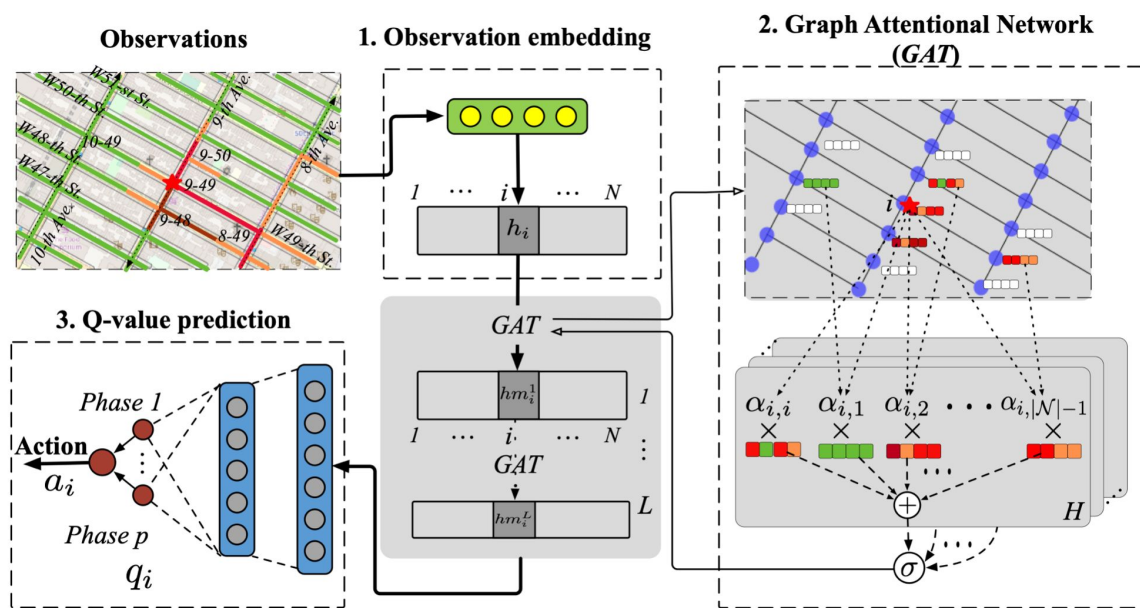
- Reward function based on pressure



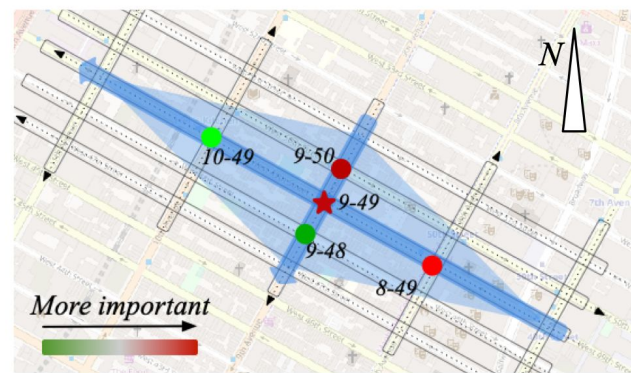
PressLight: Learning Max Pressure Control to Coordinate Traffic Signals in Arterial Network



# Baseline: Colight



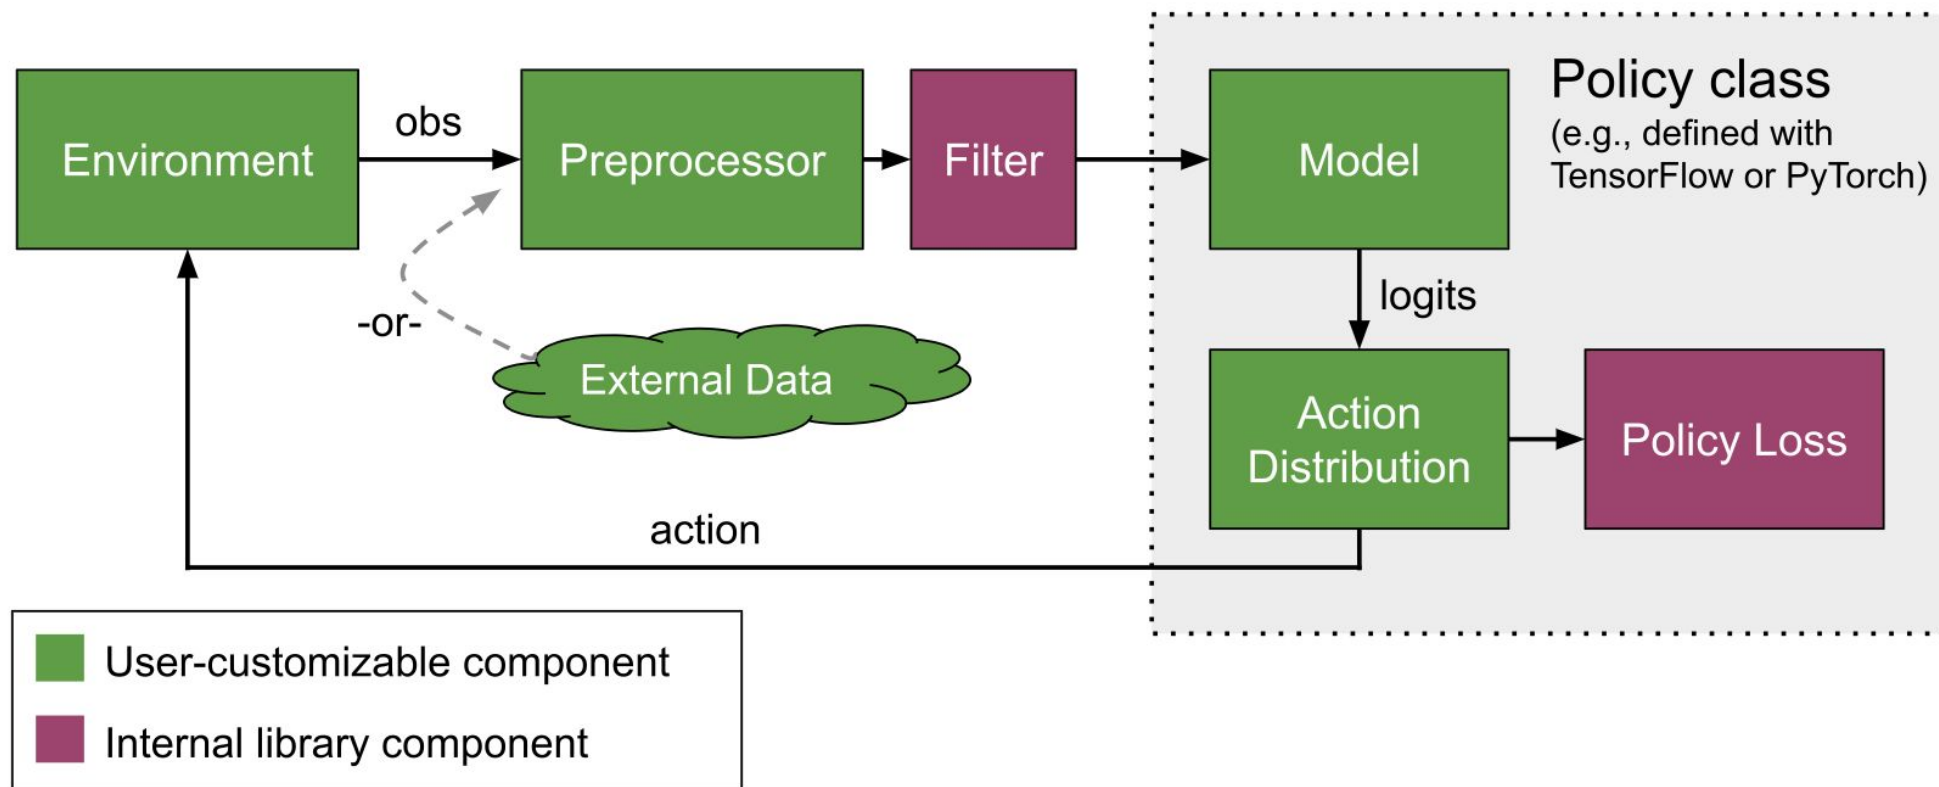
(b) 1st GAT layer (one-hop view)



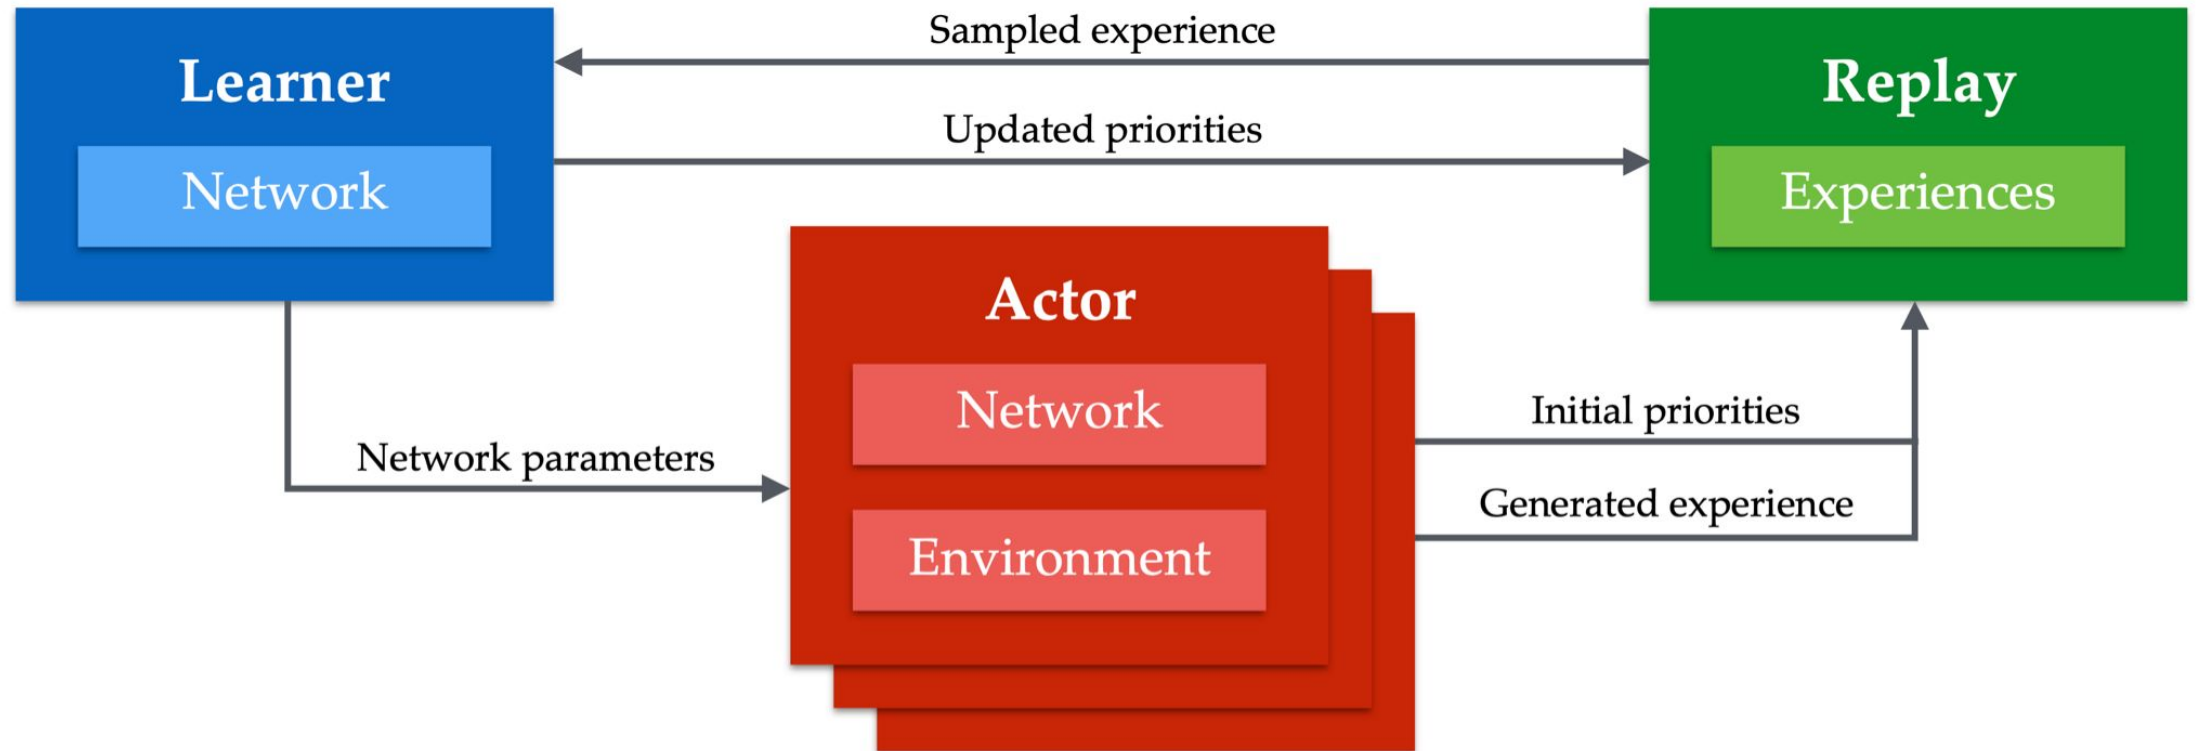
(c) 2nd GAT layer (two-hop view)

CoLight: Learning Network-level Cooperation for Traffic Signal Control

# RLLib

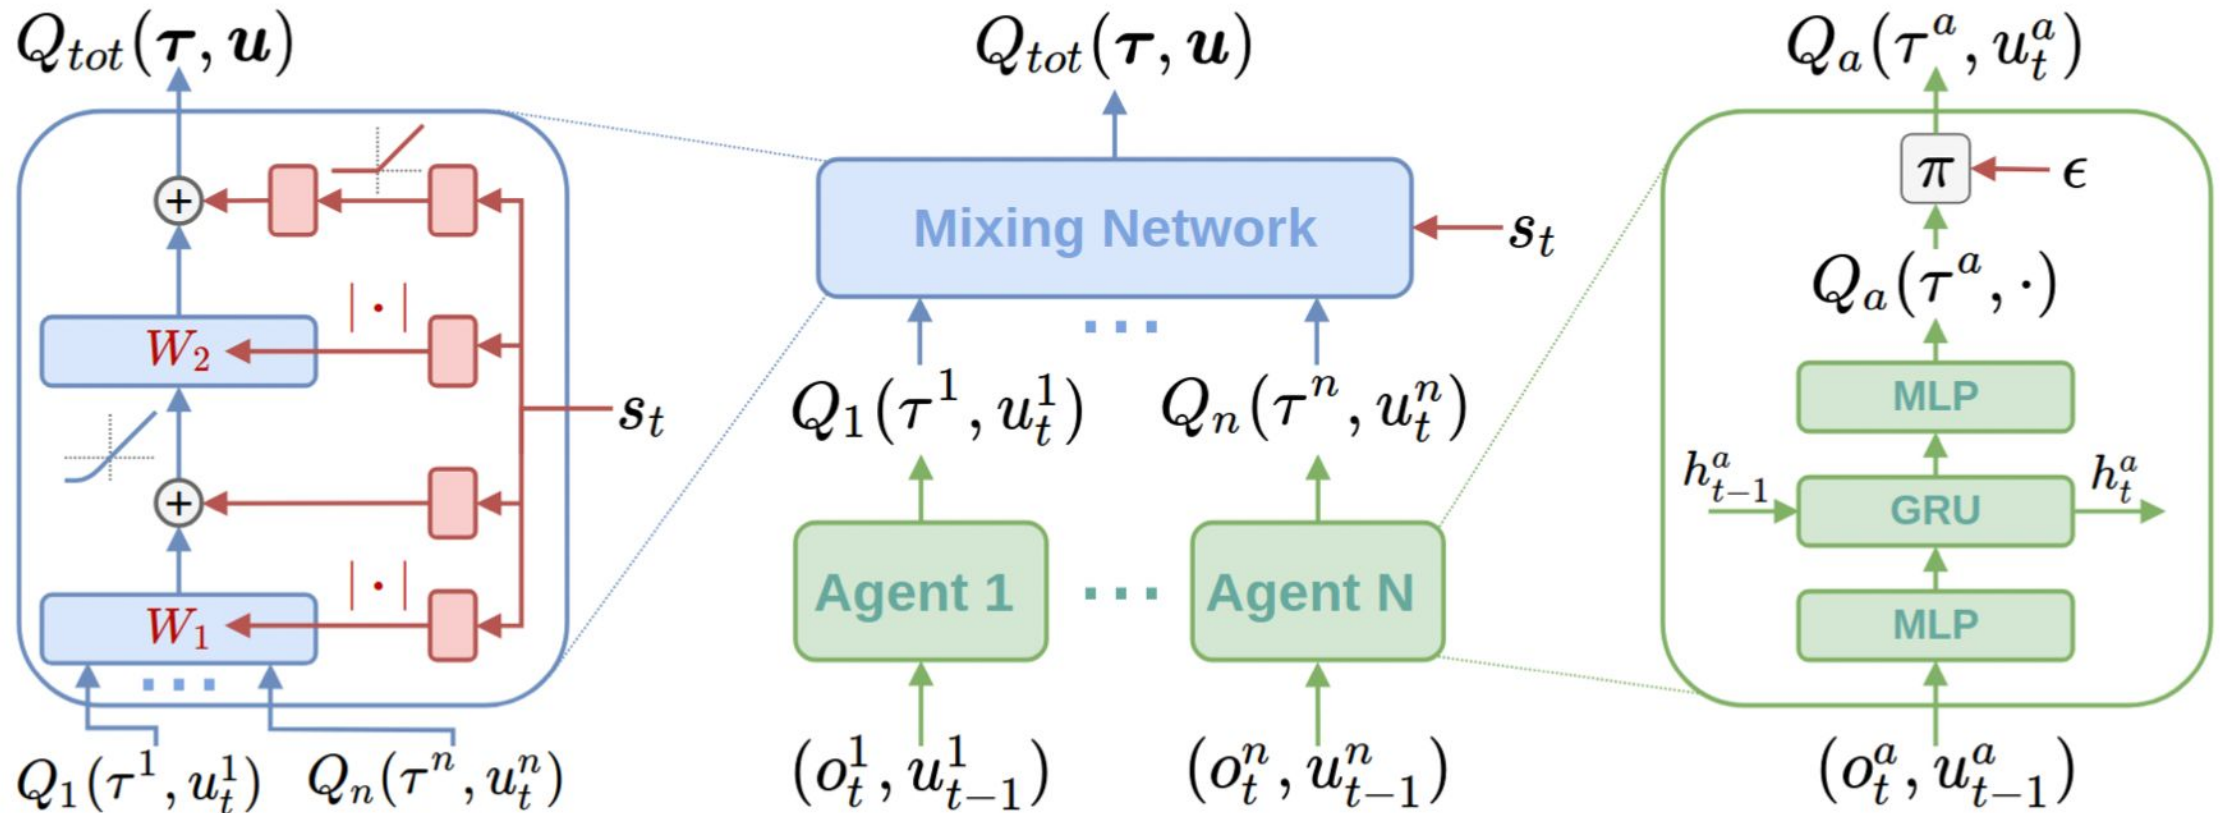


# Extension: Ape-X



## Distributed Prioritized Experience Replay

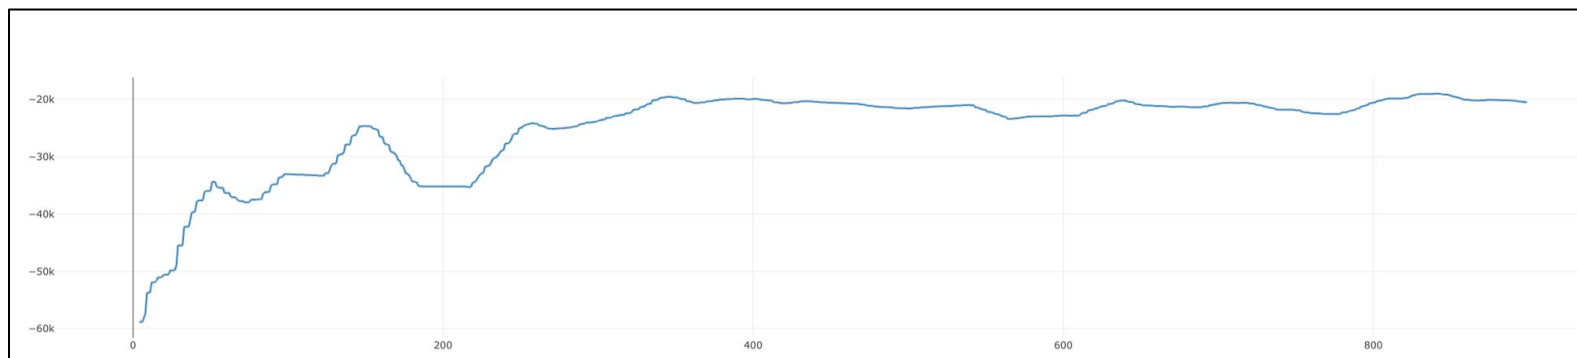
# Extension: QMix



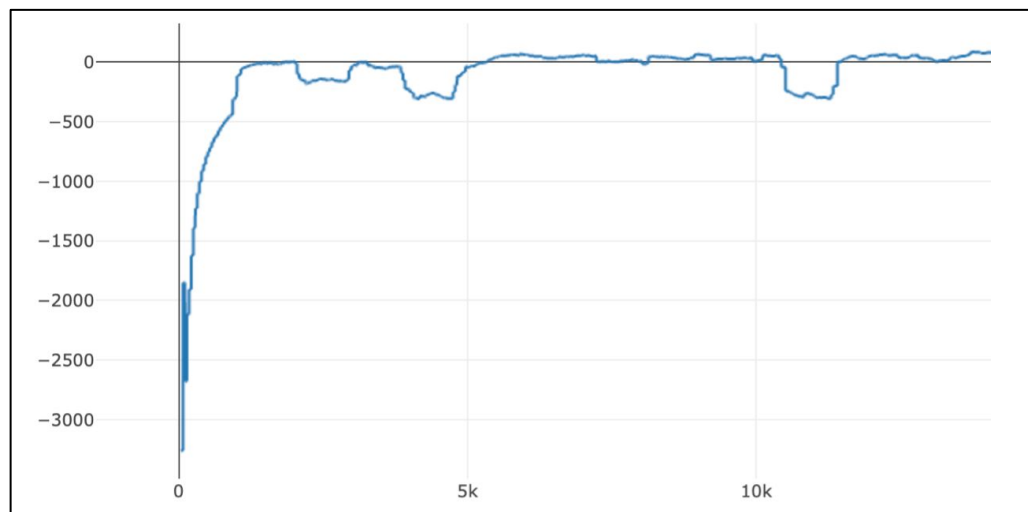
QMIX: Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning

# Training Results

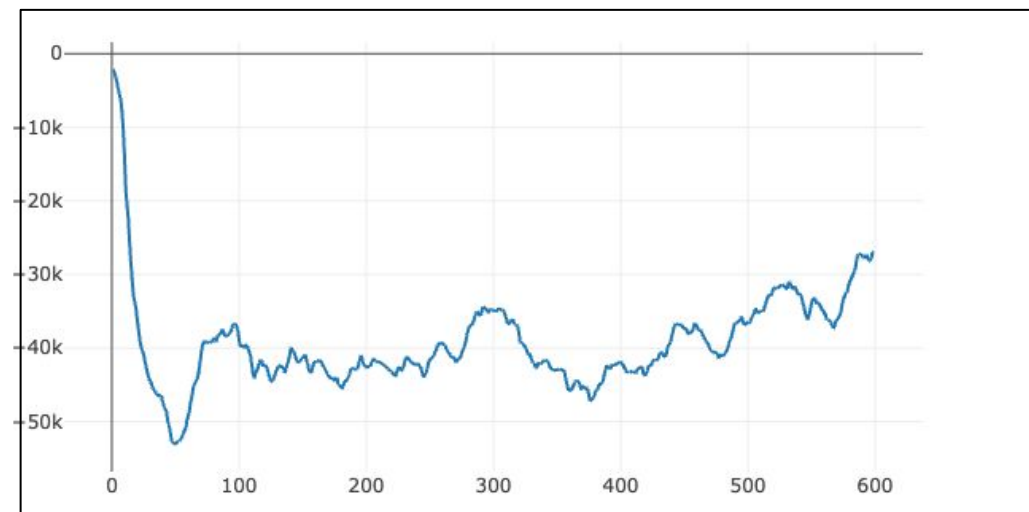
Mean Episode Reward vs Iterations



CoLight





Presslight



QMix with LSTM

# Evaluation Results

Total vehicles served, total delay

Model	Dataset	Total vehicles served	Total delay
Max Pressure (RLLib)	Warm Up	1047	1.3185
Presslight (RLLib)	Warm Up	1047	1.2609
CoLight (RLLib)	Warm Up	1047	16.307
QMix (RLLib)	Warm Up		
Presslight (Original, single hidden)	Round2	22962	1.6034
Presslight (Original, extra hidden)	Round2	29128	1.7507
Starter-kit DQN	Round2	70026	1.6566

# Challenges

- Familiarisation with RL concepts under a very tight deadline
- Task redistribution when 2 members left our team
- Switching to new starter-kit with RLlib
- Poor documentation of starter-kit & RLlib
- RLlib's low GPU utilization on PyTorch
- Limited resources on LRZ AI System

# What have we learned?

- Fundamentals of multi-agent reinforcement learning
- Coding RL models with PyTorch, later using RLLib
- Setting up ML environment
- Team collaboration

What would we have done differently in retrospect

- Stick to TensorFlow for (GPU) facilitation
- Clear prioritization of tasks around submission deadline

# Questions?

## - Phase-traffickers

