

Applications of Reinforcement Learning to Combinatorial Optimization

May 29, 2024

Thibaut Cuvelier (tcuvelier@google.com)

Joint work with Bruno De Backer (bdb@google.com), Onno Eberhard (onnoeberhard@gmail.com), Aymane Lotfi (aymanelotfi@google.com)

Operations Research team, Google Research Paris

Can RL be a useful tool for OR?

Two use cases:

- Logistics: middle mile
- Set covering

Hard constraints: the problems have structure

Can RL deal with it?

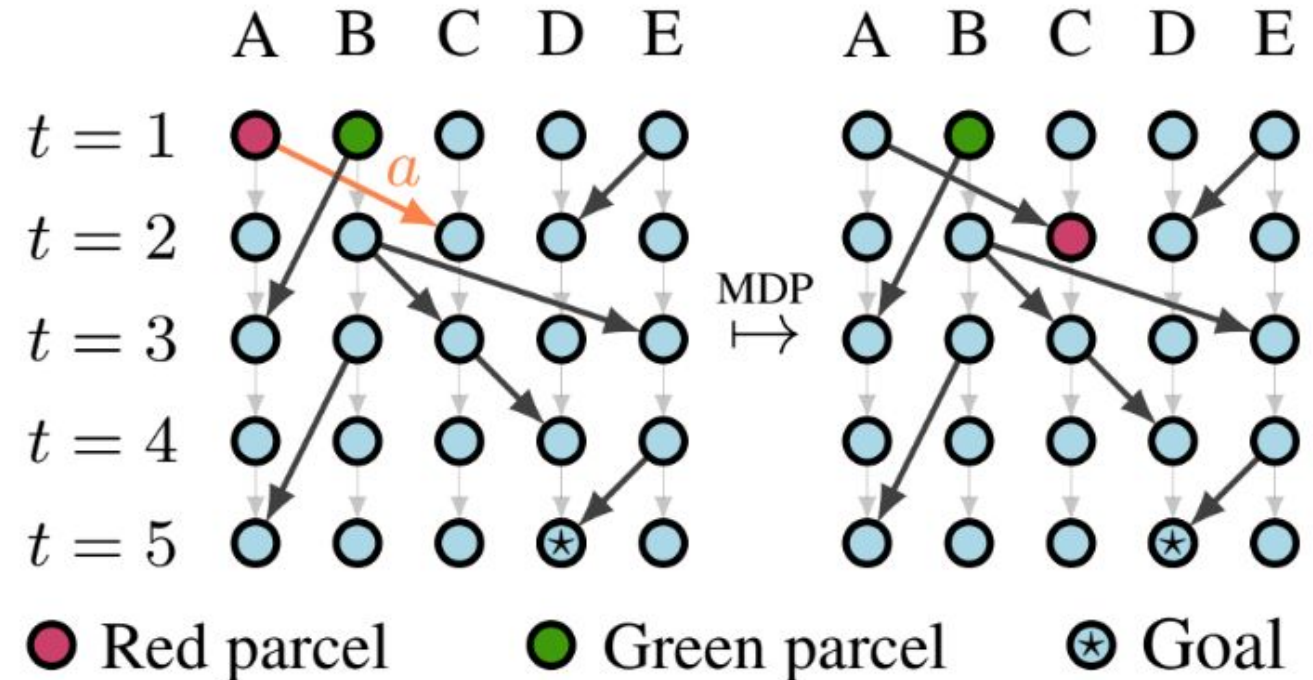
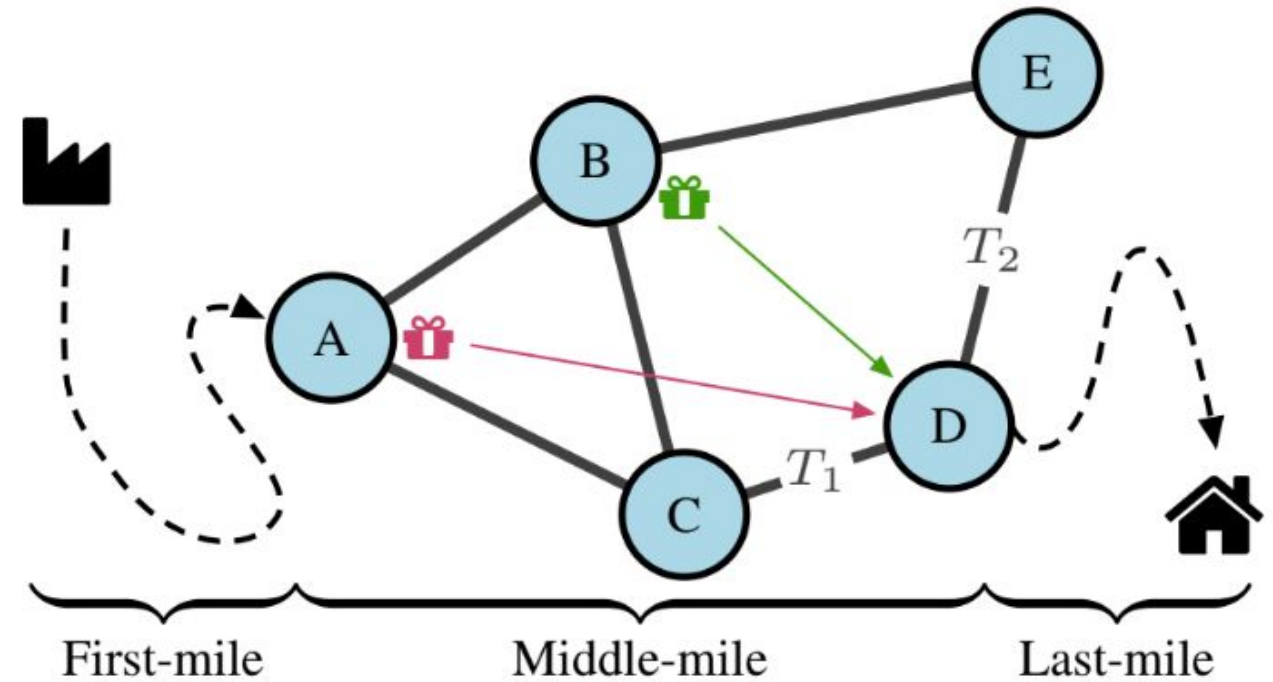
Middle-mile logistics

The problem in a nutshell:

- Hubs, trucks with fixed schedule
- Shipments to move from one hub to another
- Constraints: truck capacity, shipment deadline

Model it as an MDP?

- State: a time-expanded graph
 - Nodes are hubs with time
 - Edges are trucks with their schedule
 - Parcels are located at nodes
- Transition: one shipment moves from one hub to another
- Reward: whenever a shipment reaches its destination



Middle-mile logistics

Use RL to solve the problem

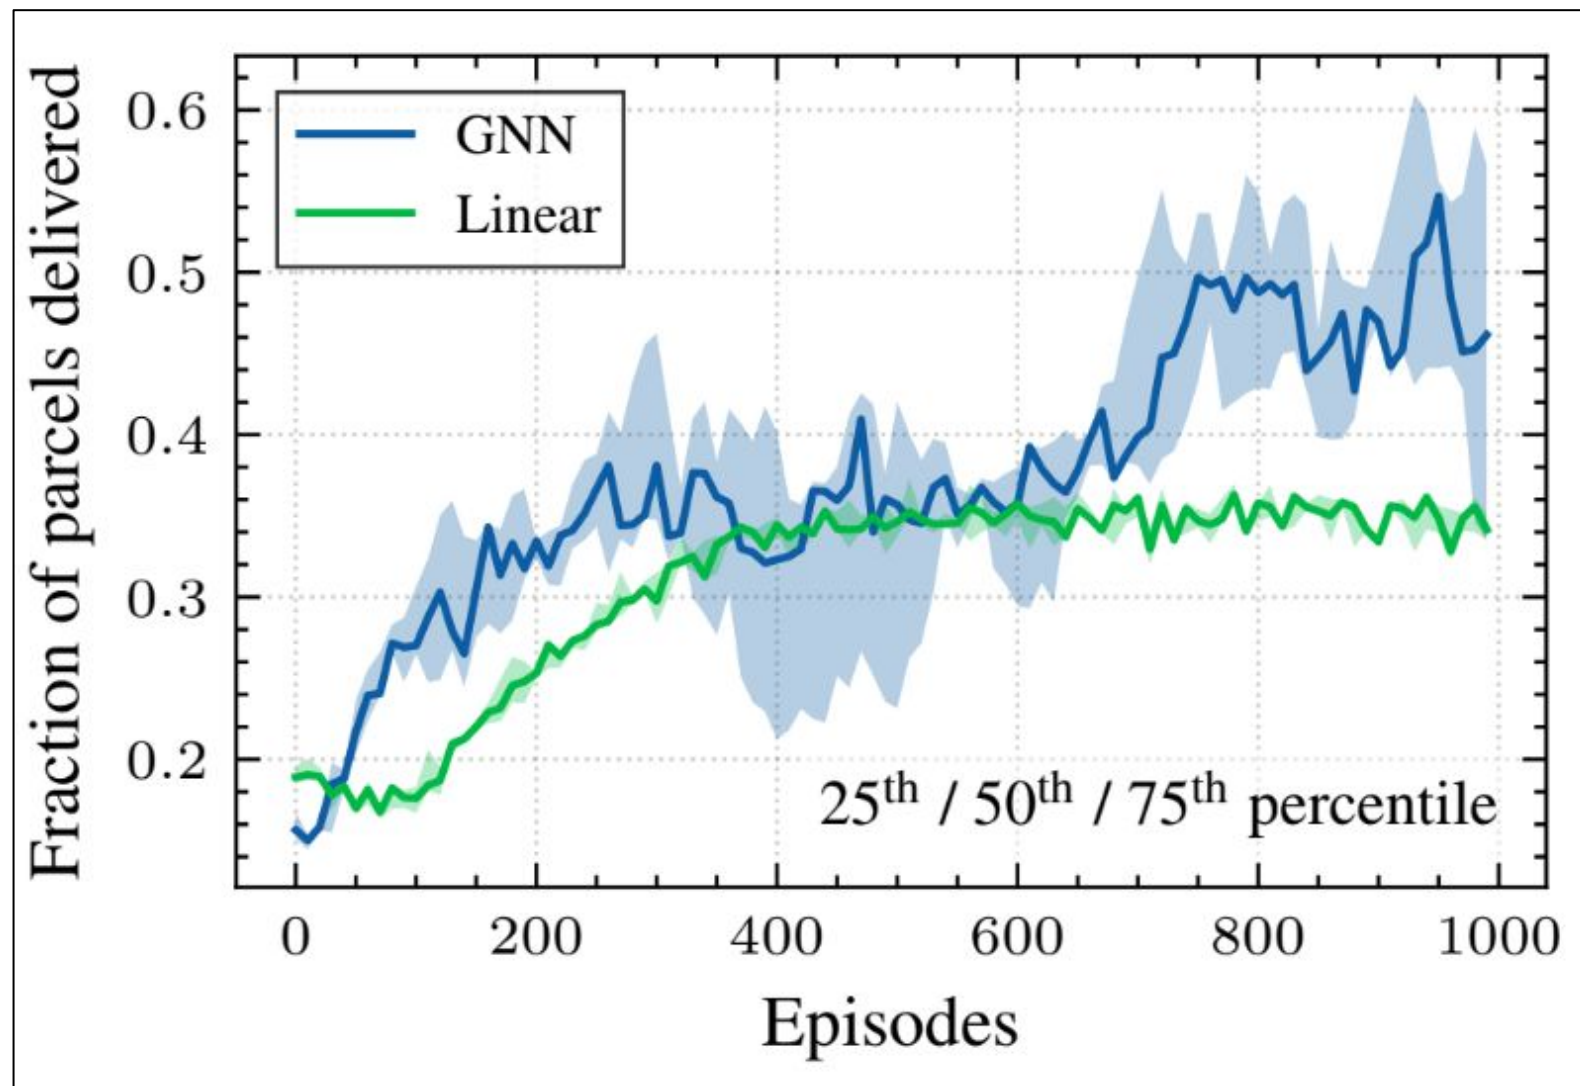
- PPO for RL
- GNN for the Q and policy functions

Results?

- GNNs provide a good fit
- Significantly more parcels delivered than with a linear model!

Feasibility achieved through:

- Modelling for deadlines
- Learning and state for capacities



Set covering

The problem in a nutshell:

- Items to cover (e.g., chessboard square): rows
- Subsets that cover items (e.g., knights): columns

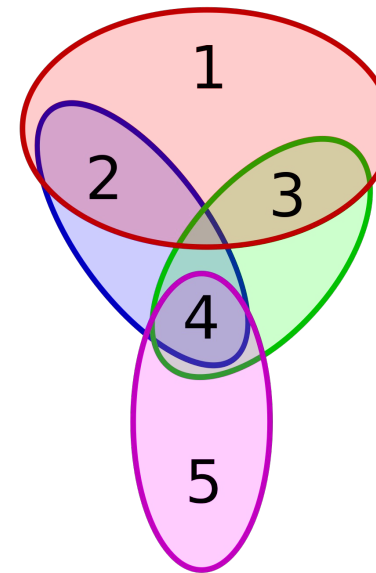
Direct applications in vehicle routing: cover shipments with possible routes

Model it as an MDP?

- State: items covered, subsets with number of items they will cover
- Transition: one subset at a time
- Reward: total cost of solution (number of subsets)

Results?

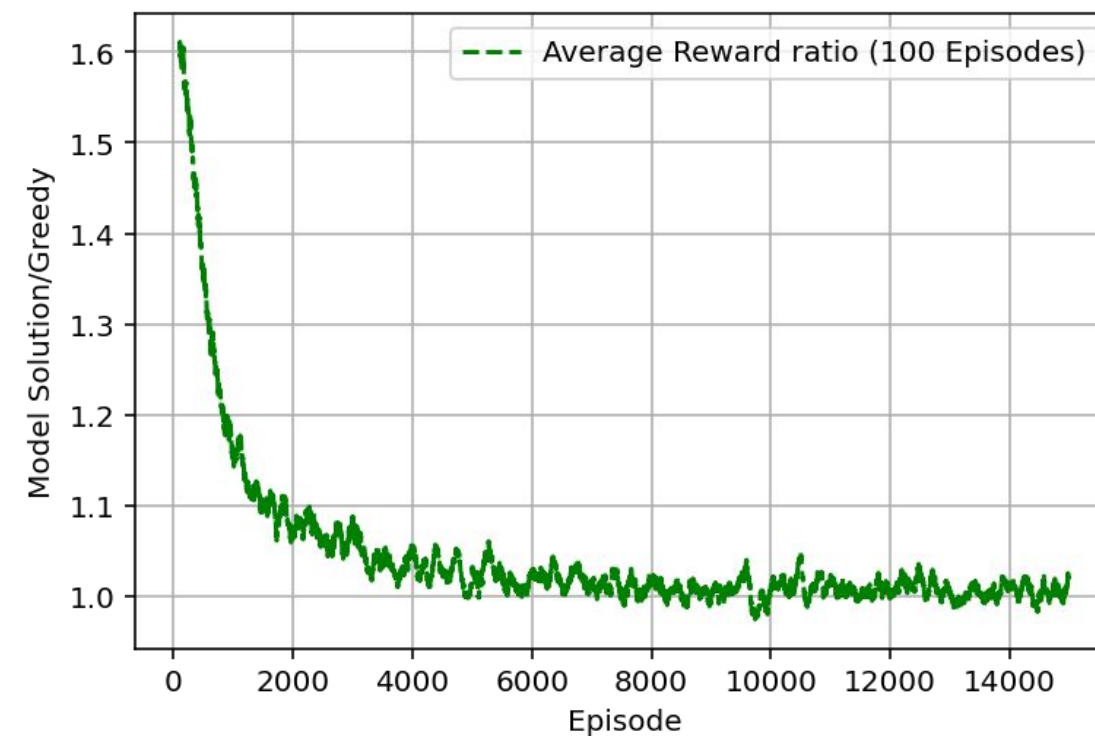
- DQN: works slightly better than a greedy heuristic (Chvátal)
- Not (yet) scaling to any number of items/subsets



By Jochen Burghardt - Own work,
CC BY-SA 4.0,
<https://commons.wikimedia.org/w/index.php?curid=139256108>



Instance of knights covering



Thank
You