Shunting Trains with Deep Reinforcement Learning

Wan-Jui Lee

R&D Hub Logistics

Dutch Railways (NS)

Public Transport by Rail

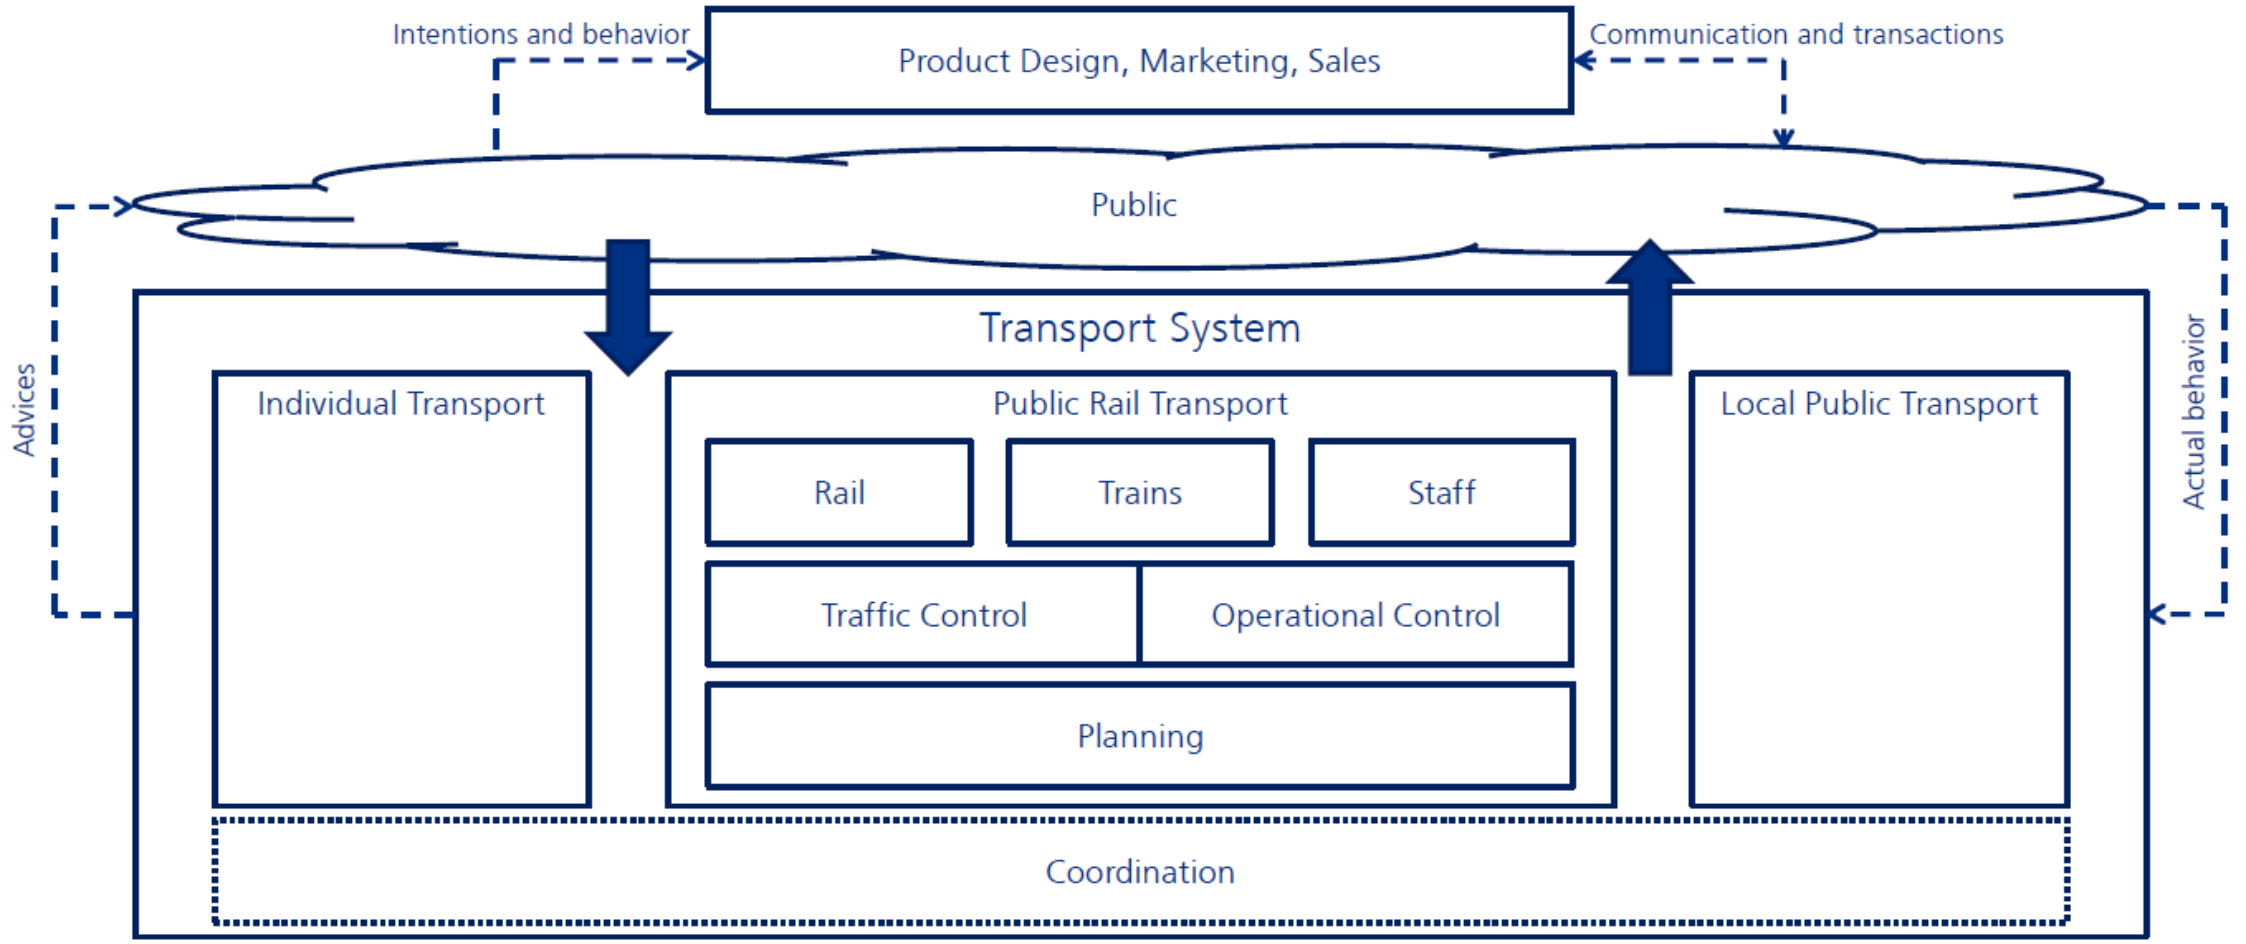

Logistic Planning

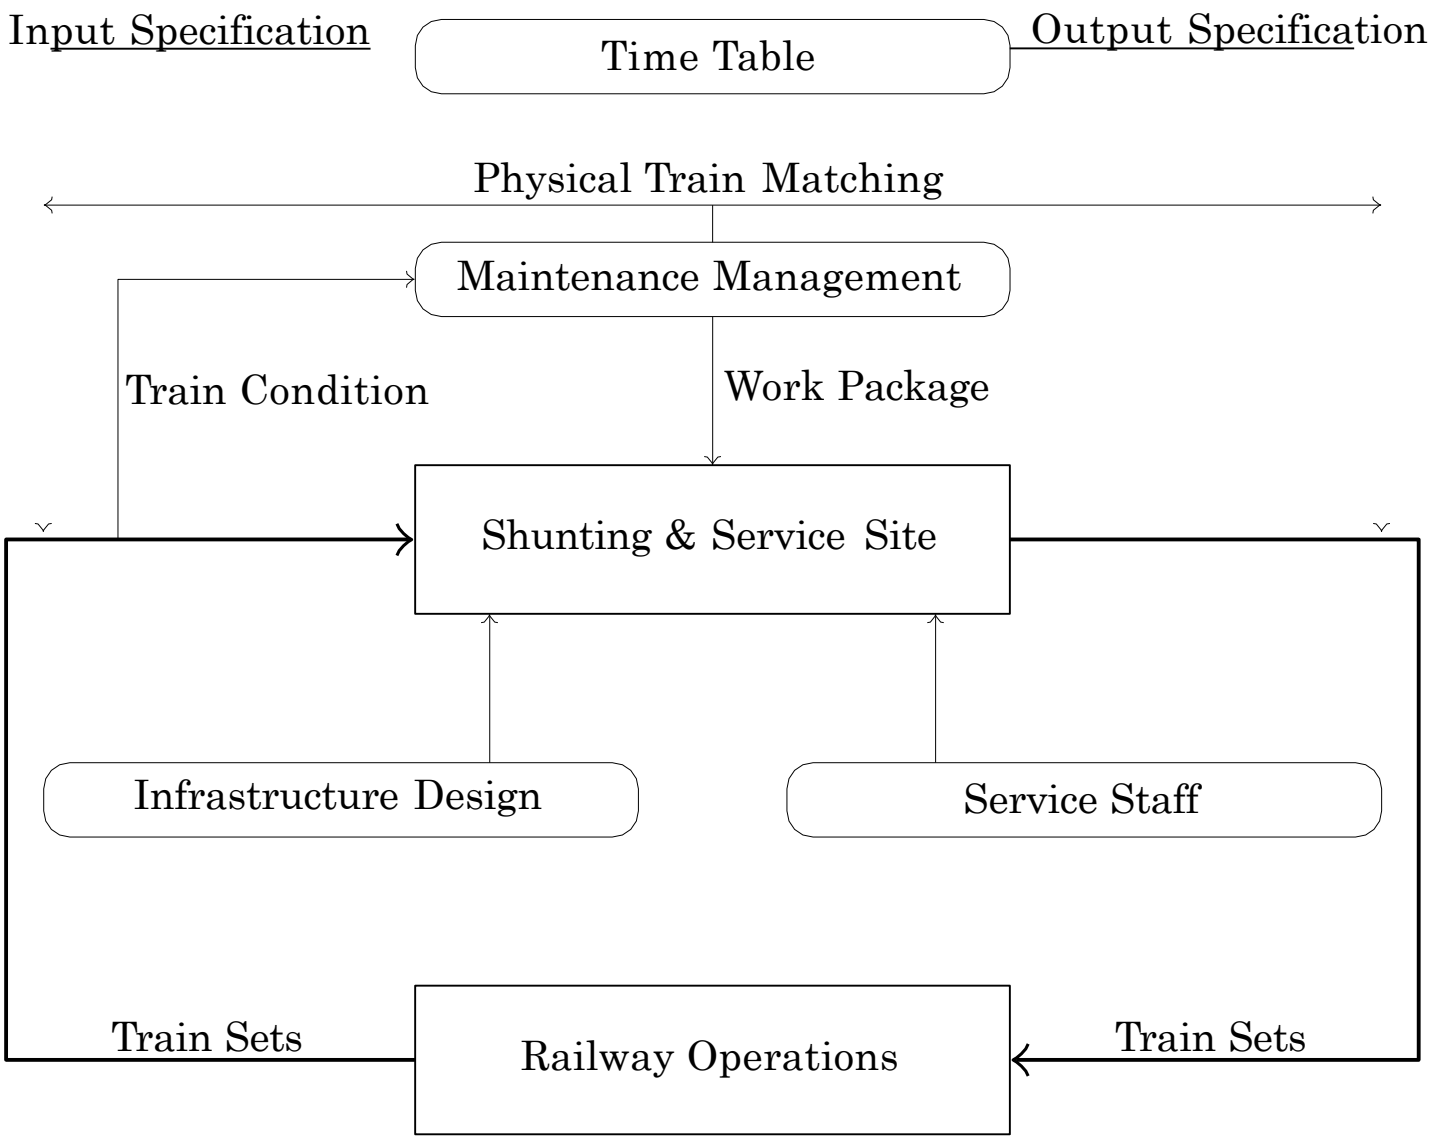

Logistic Planning of Train Shunting at Railway Hubs

- Hub infrastructure is a bottleneck for further growth of rail transportation due to space limitations
- Logistic recovery capability needs to be increased to maintain robustness on high density network
- Anticipating predicted passenger flows requires dynamic fleet assignment
- Computationally hard problem to find feasible solution
- Long-term R&D cooperation with several universities
- Hybrid approach
  - Heuristics and Local Search
  - (Deep) Machine Learning
  - Optimization

Train Unit Shunting Problem

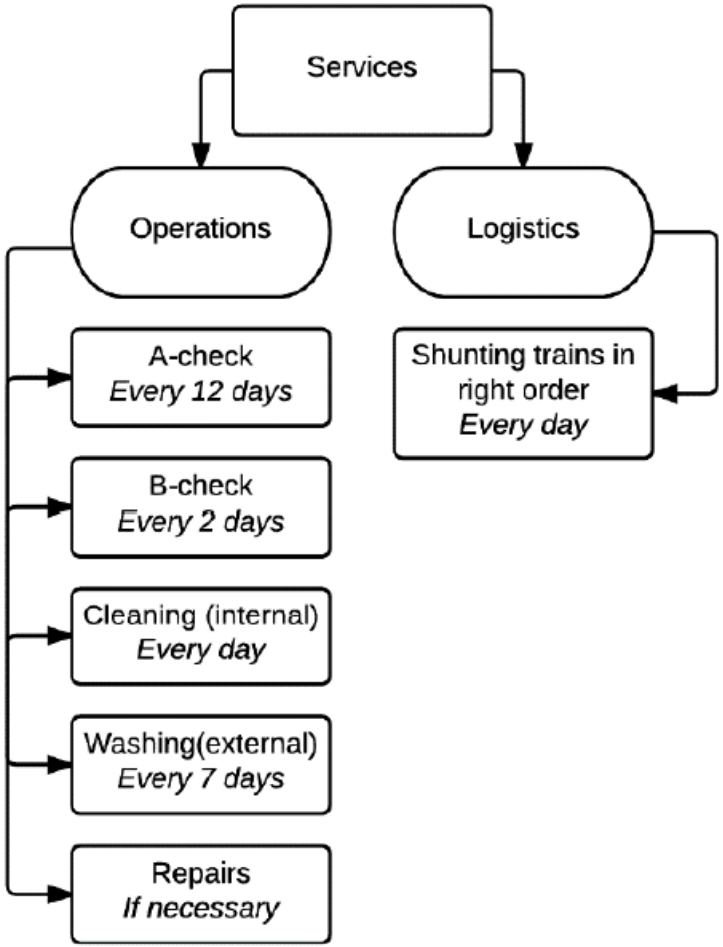

Train Unit Shunting Problem

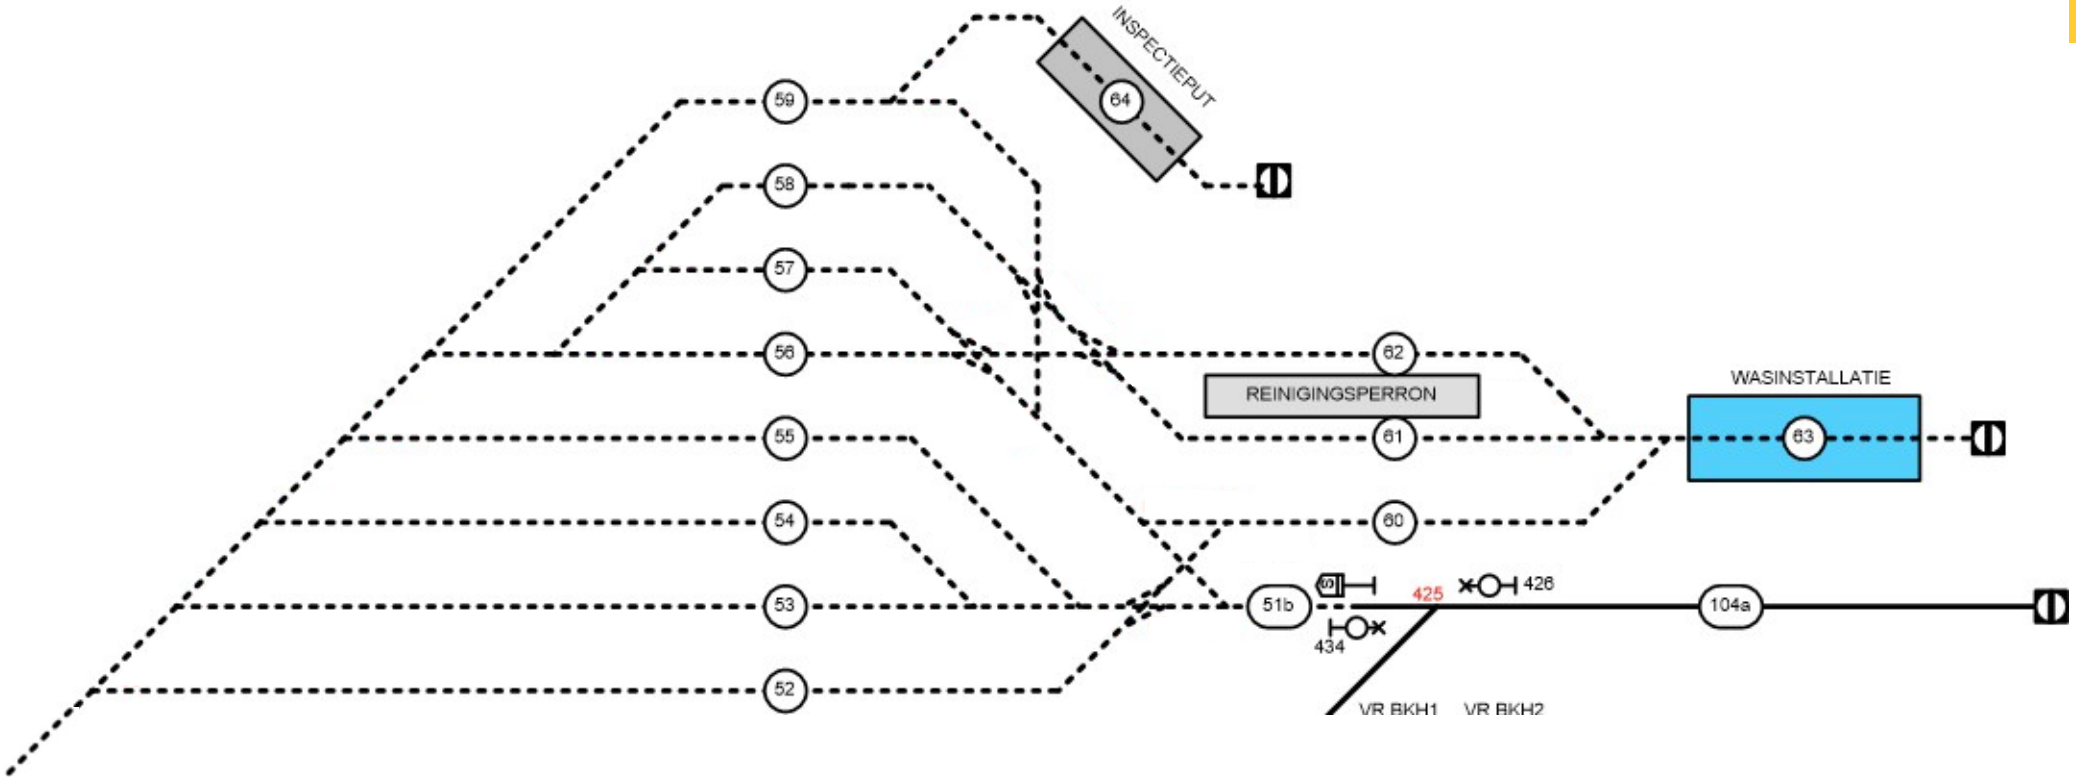

Service Location with carousel layout (Den Haag Kleine Binckhorst)

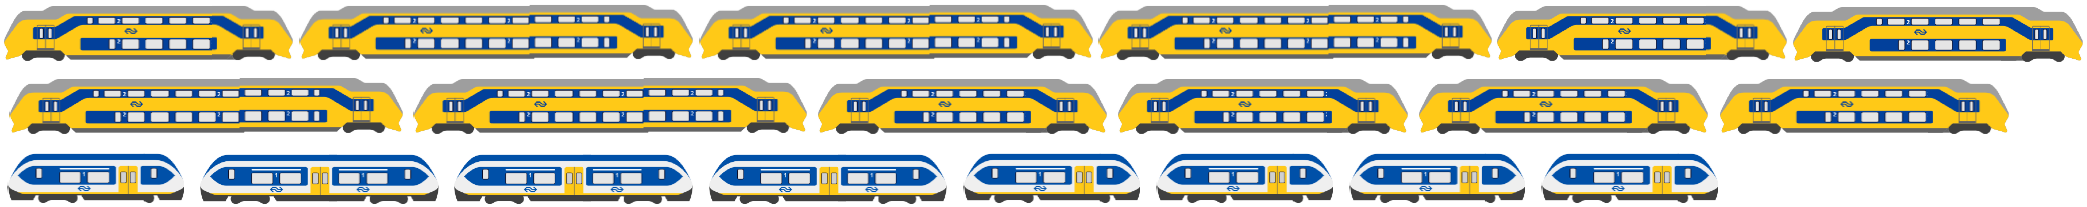

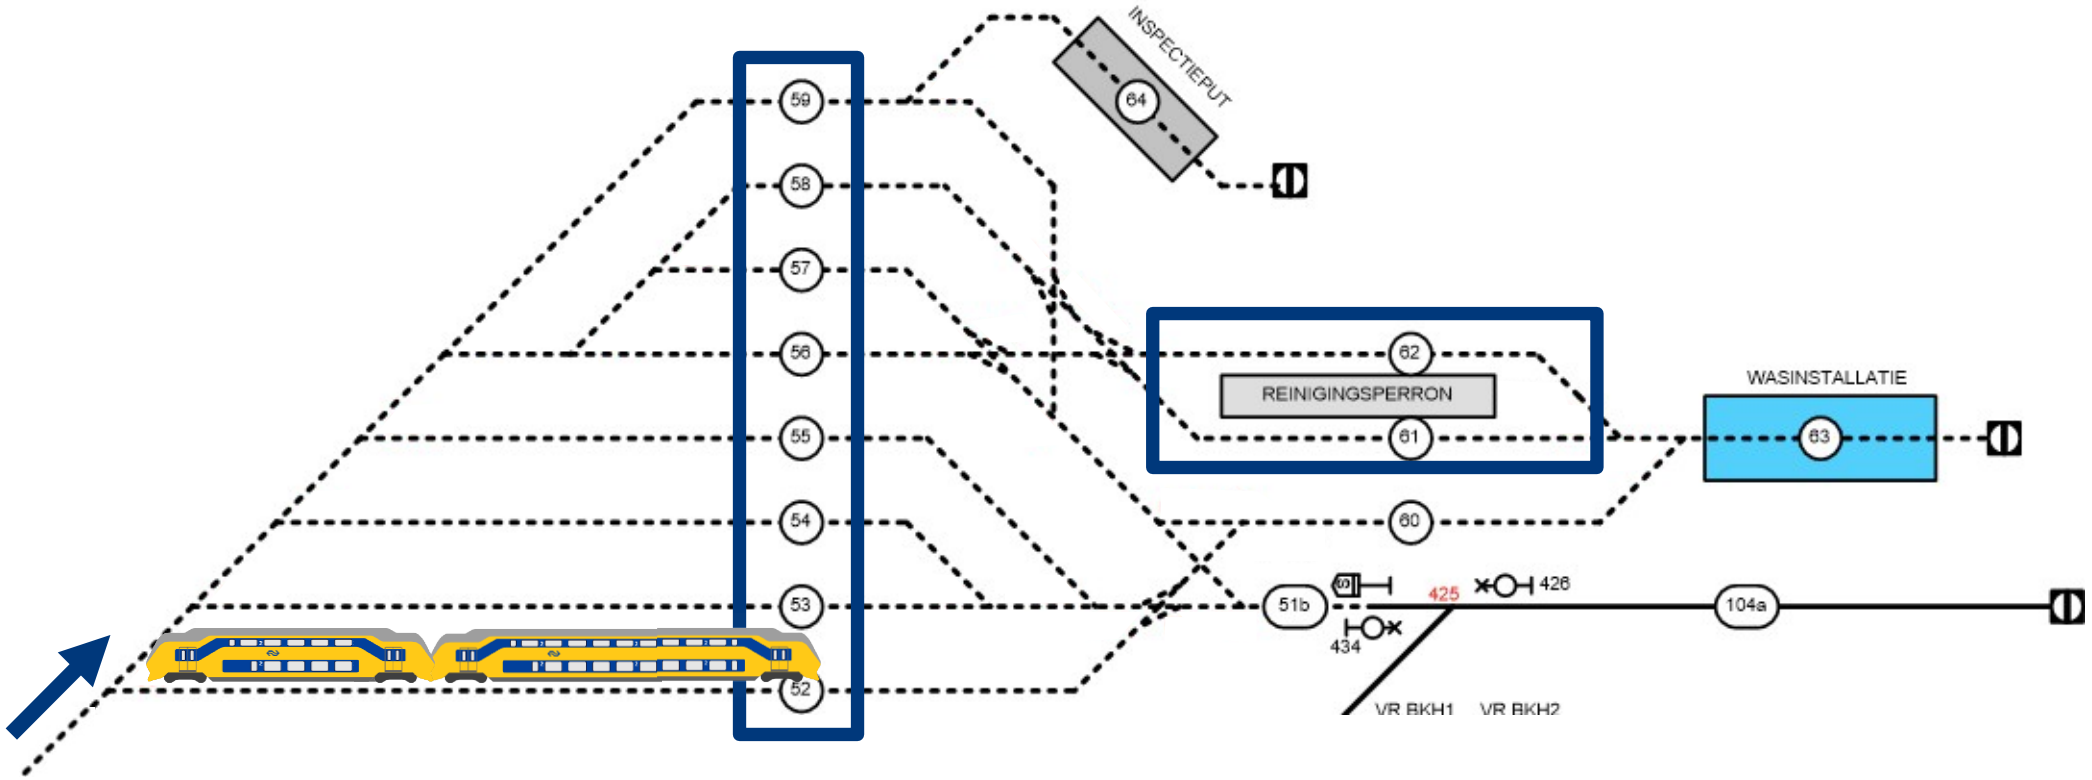
INSPECTIEPUT
64
59
58
57
56
55
54
53
52
62
REINIGINGSPERRON
61
WASINSTALLATIE
63
60
51b
425
426
434
104a
VR BKH1
VR BKH2

Capacity Determination & Daily Planning

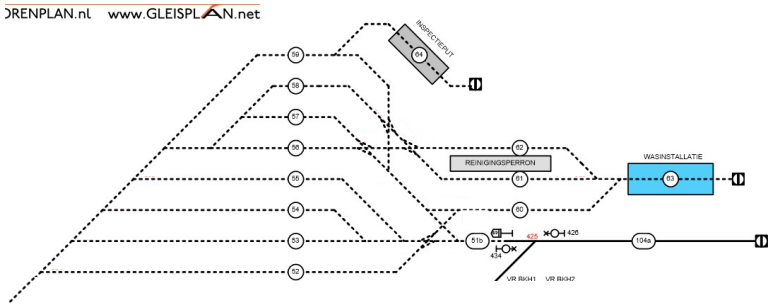

Heuristic approach

Feasible / Infeasible

Capacity Determination & Daily Planning

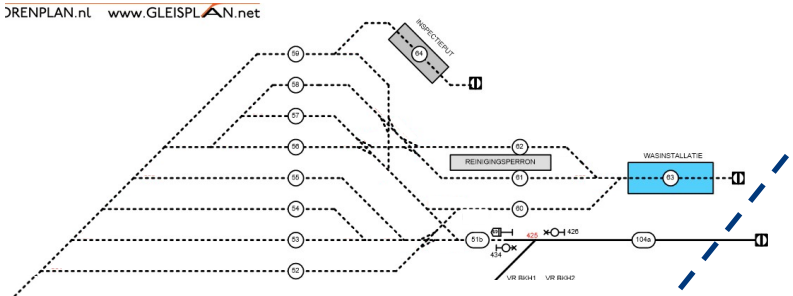

Heuristic approach

Feasible / Infeasible

Capacity Determination & Daily Planning

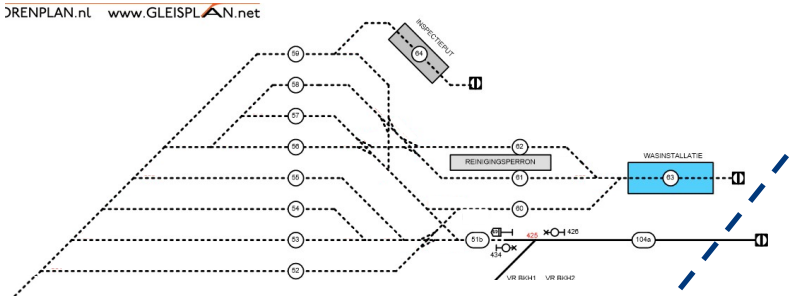

Capacity Determination & Daily Planning

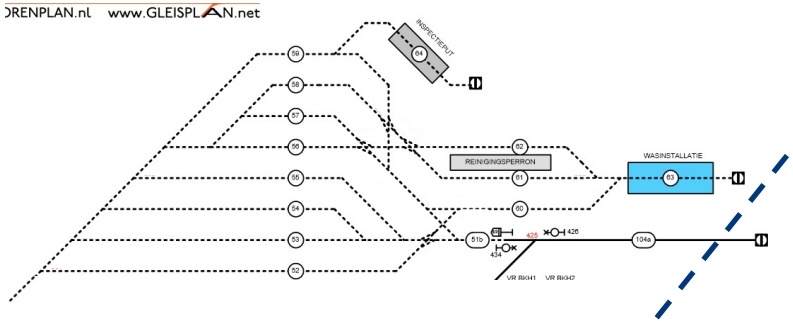

Capacity Determination & Daily Planning

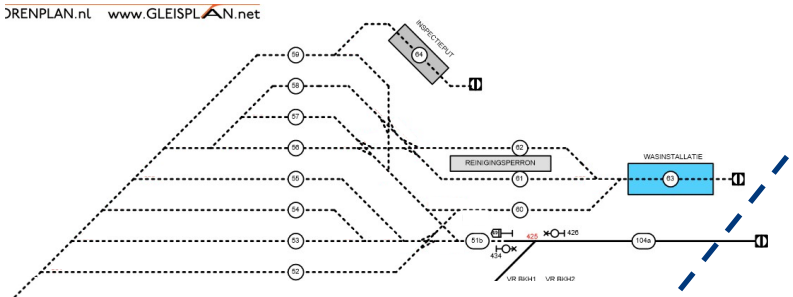

Logistical Analysis on Hubs

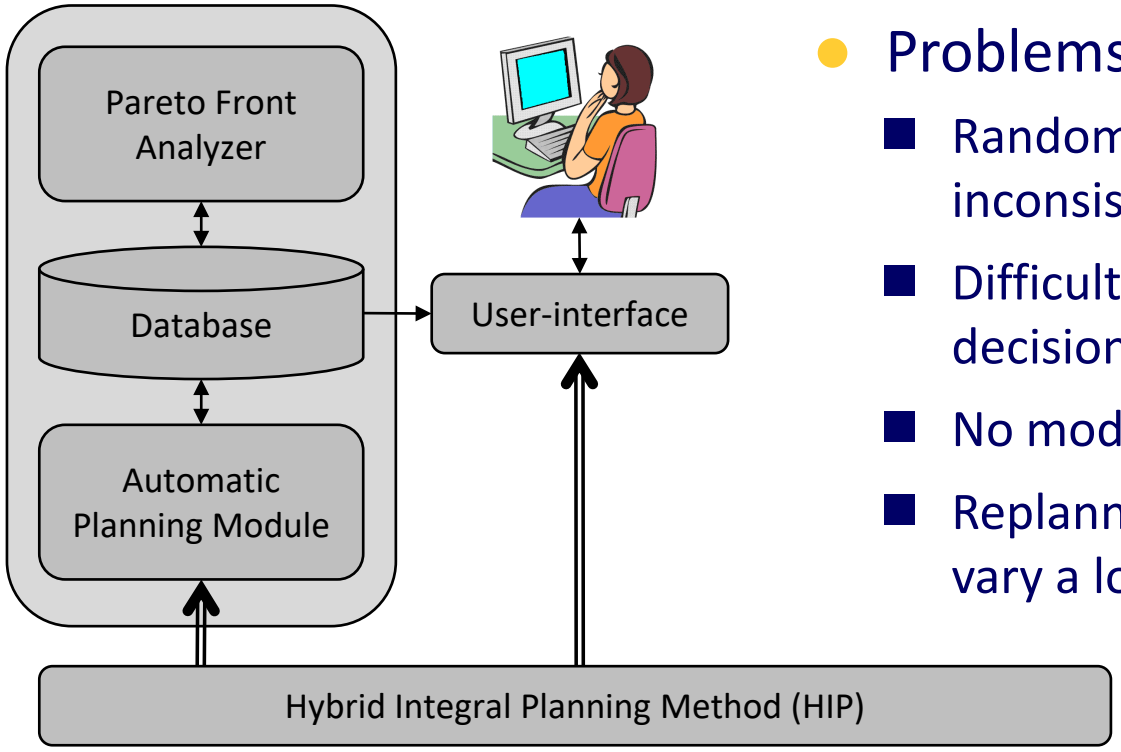

- Problems of Local Search
  - Random initial solutions result in inconsistent plans
  - Difficult to decipher how the planning decision is made
  - No model of planning strategy
  - Replanning due to disturbance may vary a lot from the original plan

Capacity Determination & Daily Planning

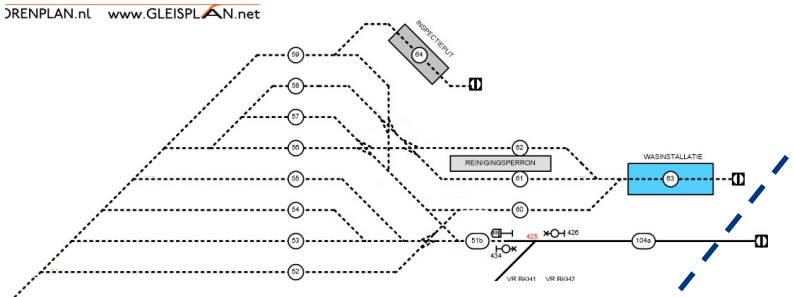

Can machines learn to plan?

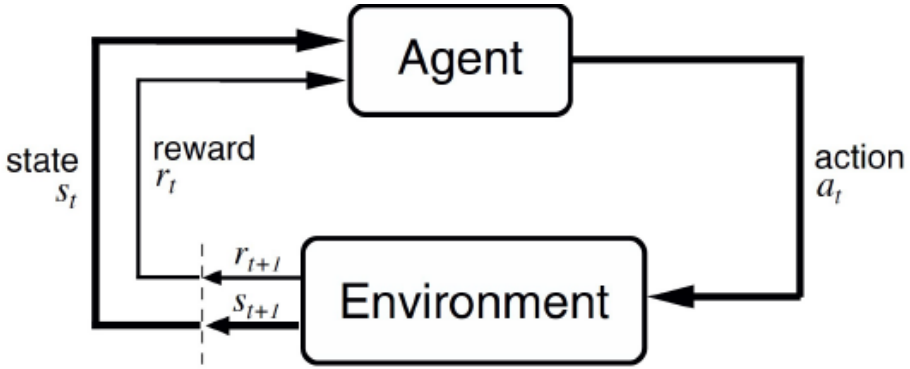

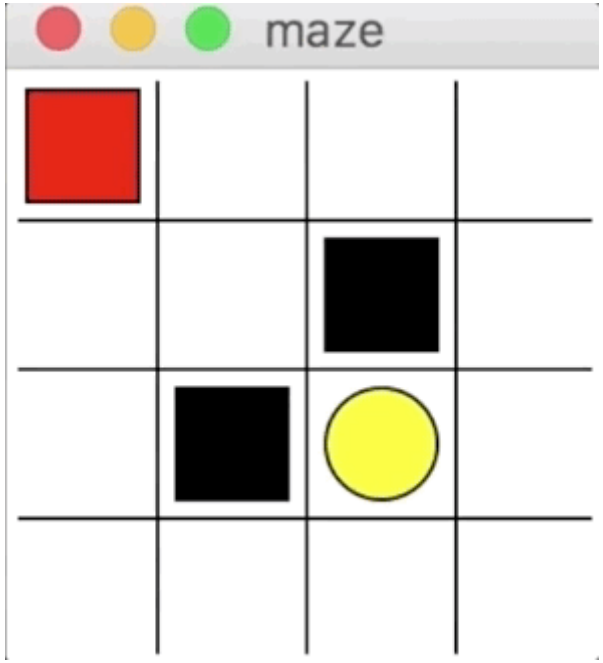

Reinforcement Learning learns to play a game by gaining experience, just like a human player:

- Try various actions in different situations (explore)
- Learn/store information about the game that can be generalized to potentially unseen scenarios
- Learn the most valuable actions by using the reward signal (exploit)

Deep Reinforcement Learning for TUSP

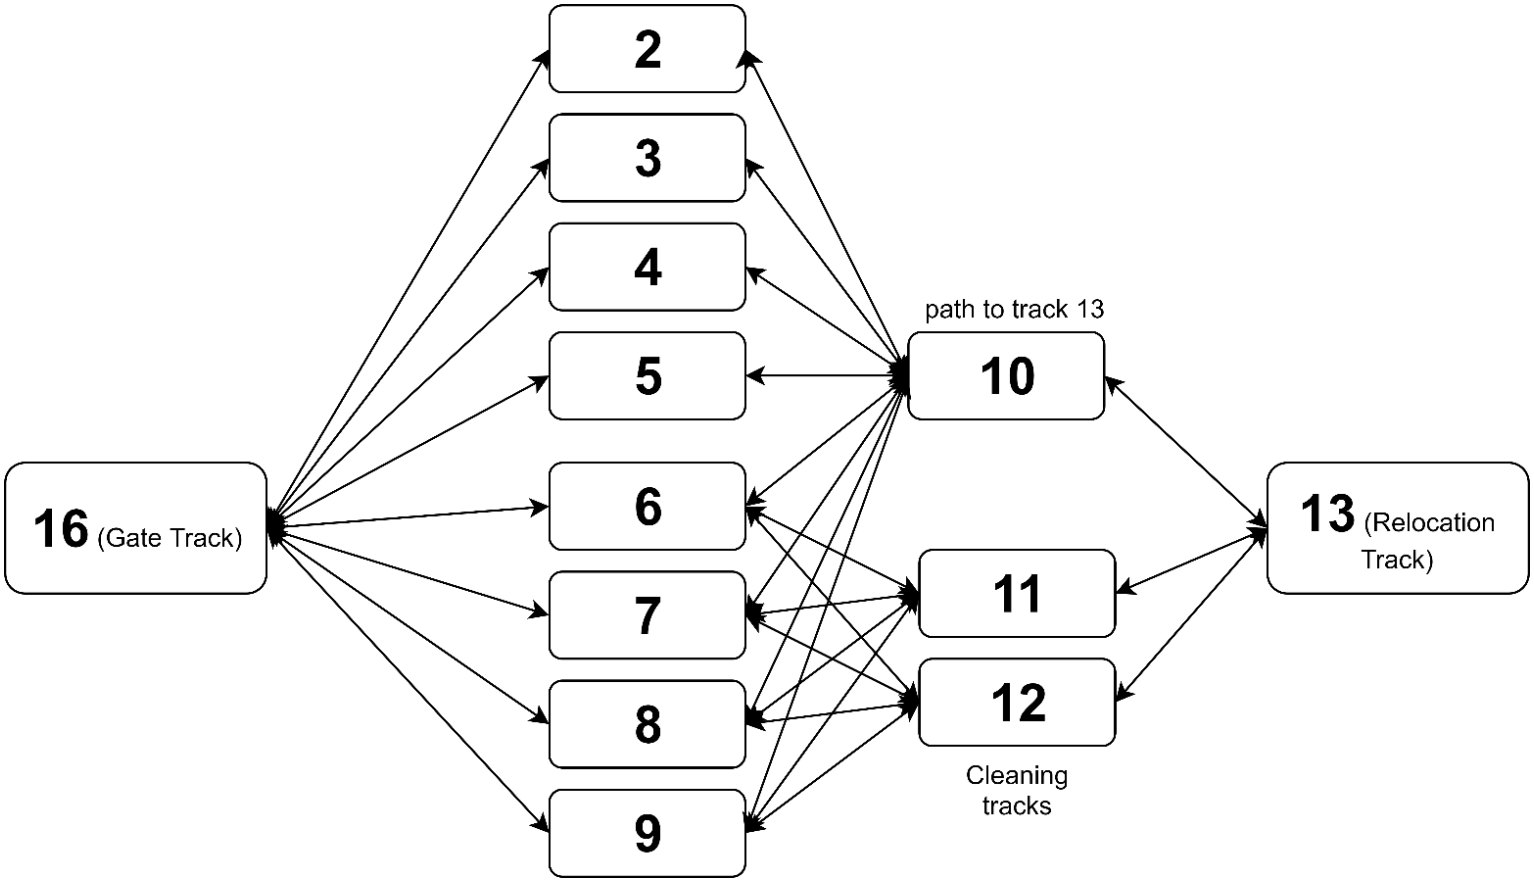

A simplified view of Den Haag Kleine Binckhorst

- **Design of**
  - **State**
  - **Action**
  - **Reward**

Multi-agent DRL for TUSP

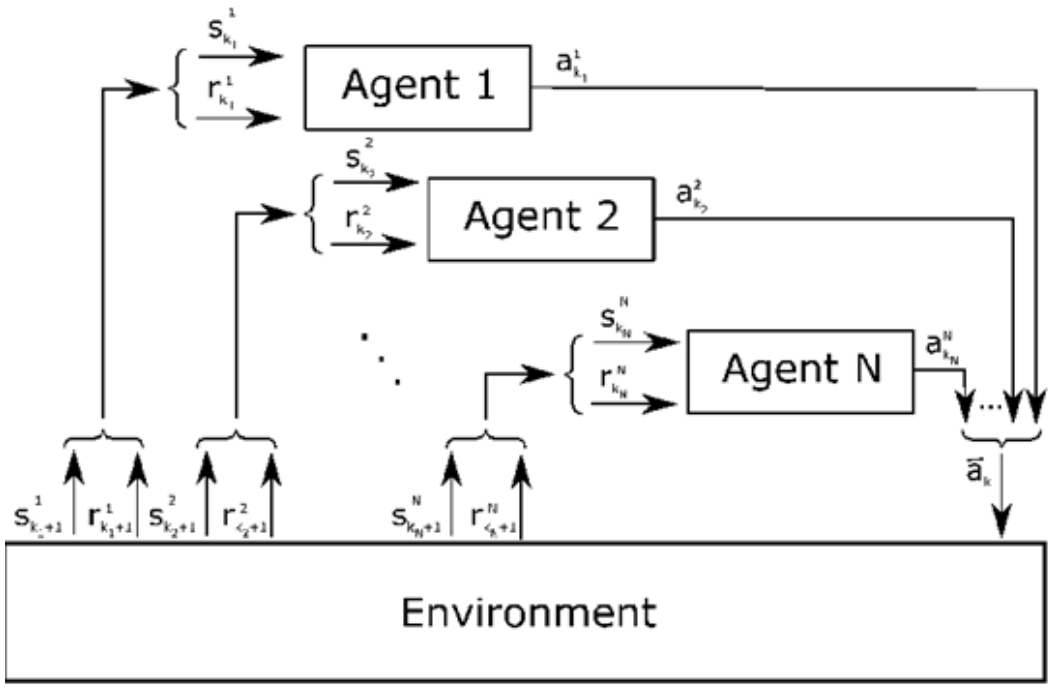

- Each train unit is an agent
- Possible to take combine/split actions
- Agents are distinguished by their input to the DNN model
- Agents learn independently to reach their own optimal V values
- One weight-sharing DNN model for all agents

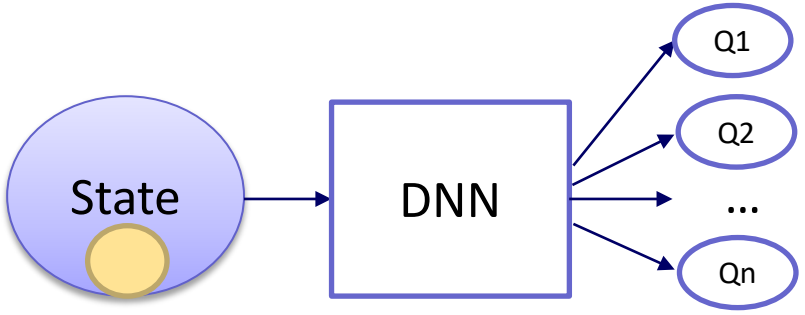

DRL for Multi-Agent Path Finding

- A routing decision for a train can impact the yard in the future, not just for one train, but for many. The potential impact increases as the yard becomes more congested.
- Combining Graph Convolutional Networks and Multi-Agent Reinforcement Learning (MARL) gives us a scalable method for optimizing routing decisions
- Agents observe only their local environment, trading off less relevant information for increased scalability

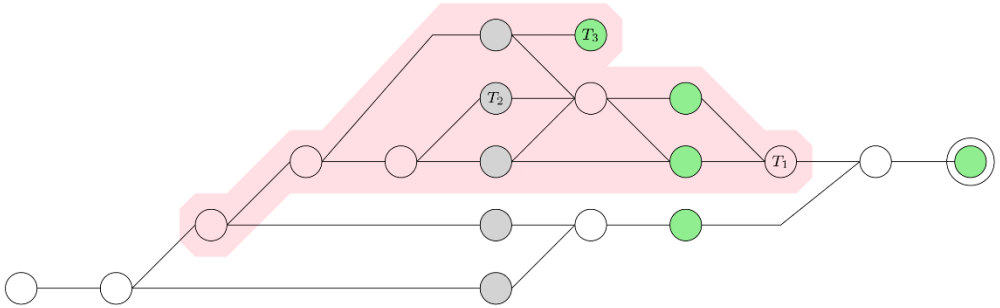

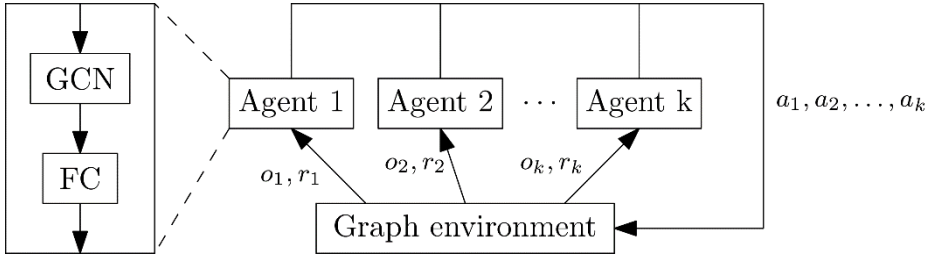

Performance: Problem Solving Capability

Train units	14	15	16	17	18
DRL solvability	93%	85%	51%	54%	28%

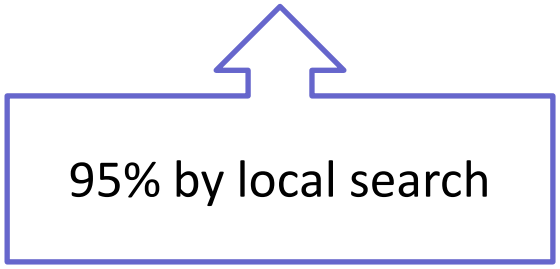

Improve the consistency of local search using DRL to generate initial solutions

Generated Solutions Includes Exact Routing

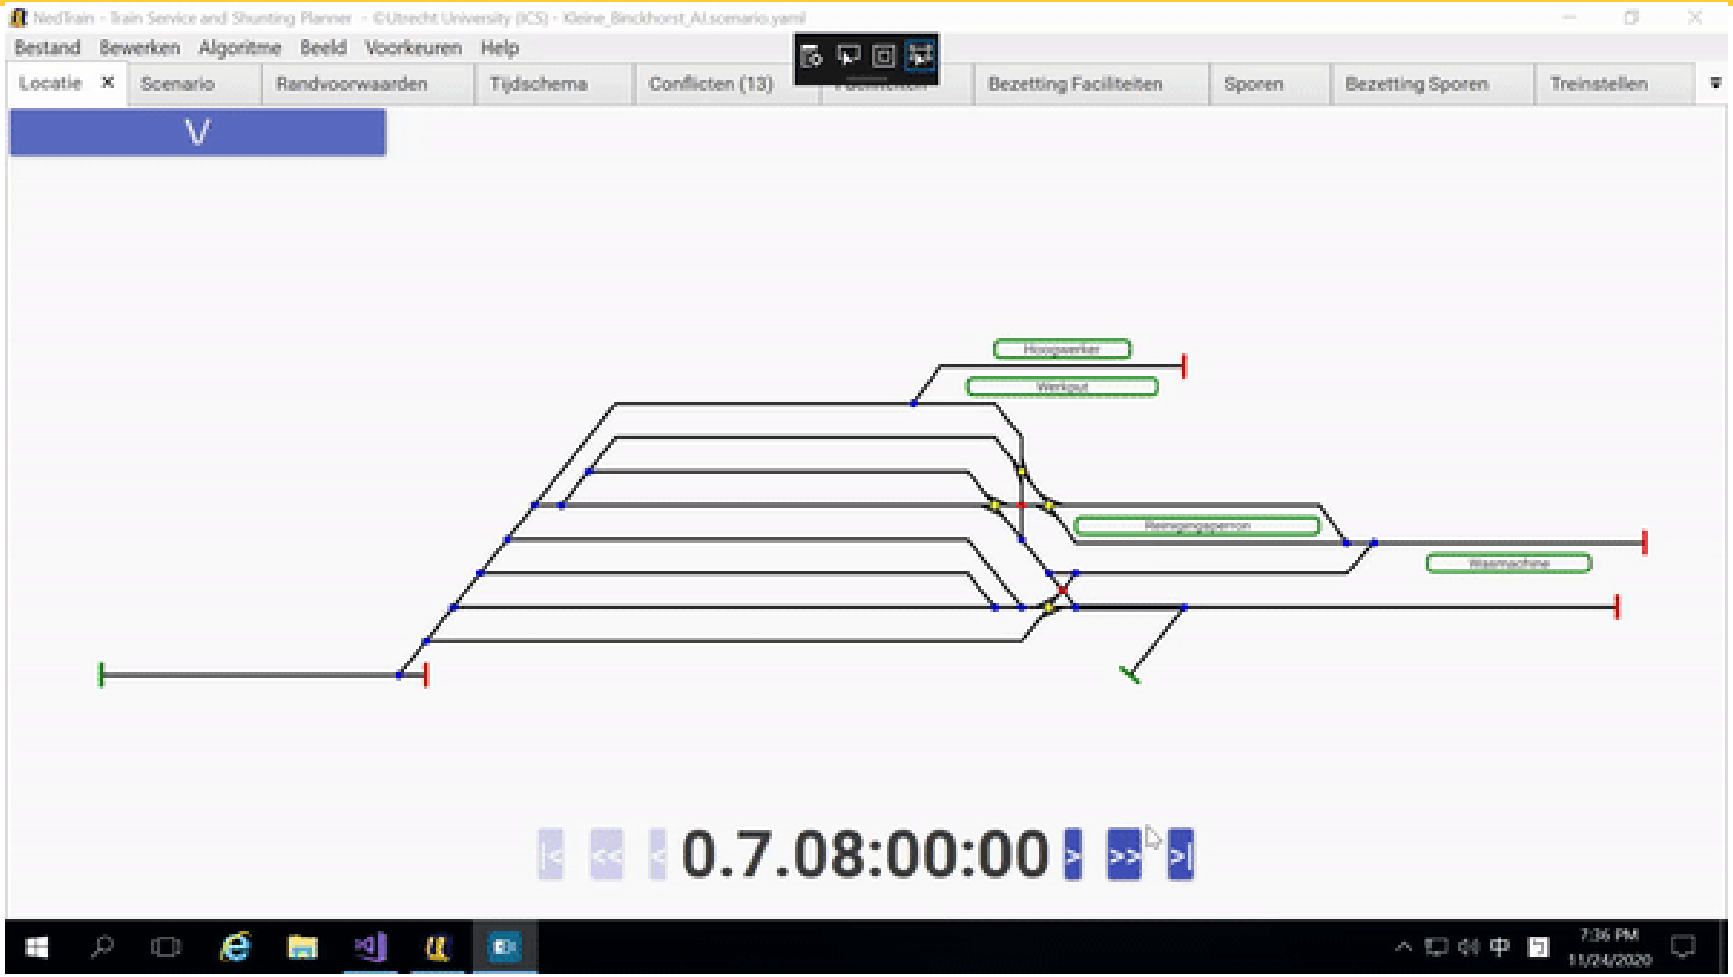

Cooperation: Internal and External

- In-house research & development groups, closely linked to the business
- Long-term research cooperation, partly funded by NWO (Dutch Research Council)
  - Delft University of Technology
  - Utrecht University
  - Eindhoven University of Technology
  - University of Amsterdam
  - VU Amsterdam
  - Erasmus University Rotterdam
- AI & Mobility Lab at Utrecht University

Cooperation: Internal and External

- In-house research & development groups, closely linked to the business
- Long-term research cooperation, partly funded by NWO (Dutch Research Council)
  - Delft University of Technology
  - Utrecht University
  - Eindhoven University of Technology
  - University of Amsterdam
  - VU Amsterdam
  - Erasmus University Rotterdam

  - Making the difficulties of instances and the properties of solutions explainable
    - Ya Song, Laurens Bliek, Yingqian Zhang
  - Multi-agent pathfinding with DRL
    - Marijn van Knippenberg, Vlado Menkovski
  - Finding Robust & Multi-objective Train Shunting Heuristics using DRL
    - Matthew Macfarlane, Herke van Hoof
- AI & Mobility Lab at Utrecht University

Q&A

Further interets/questions: wan-jui.lee@ns.nl