

Collaborative picking guided by DRL

A collaboration between industry and academia

Kasper Hendriks | Vanderlande

Luca Begnardi | TU/e

Contents

01

About Vanderlande

02

A brief history of picking

03

Orchestrating chaos

04

Reinforcement Learning

05

Modelling chaos

06

Industrial scale solution

07

Results

About Vanderlande: Company profile



About Vanderlande: Local and global presence



Business segment: Warehousing

FOCUS AREAS

Food



General Merchandise



Fashion

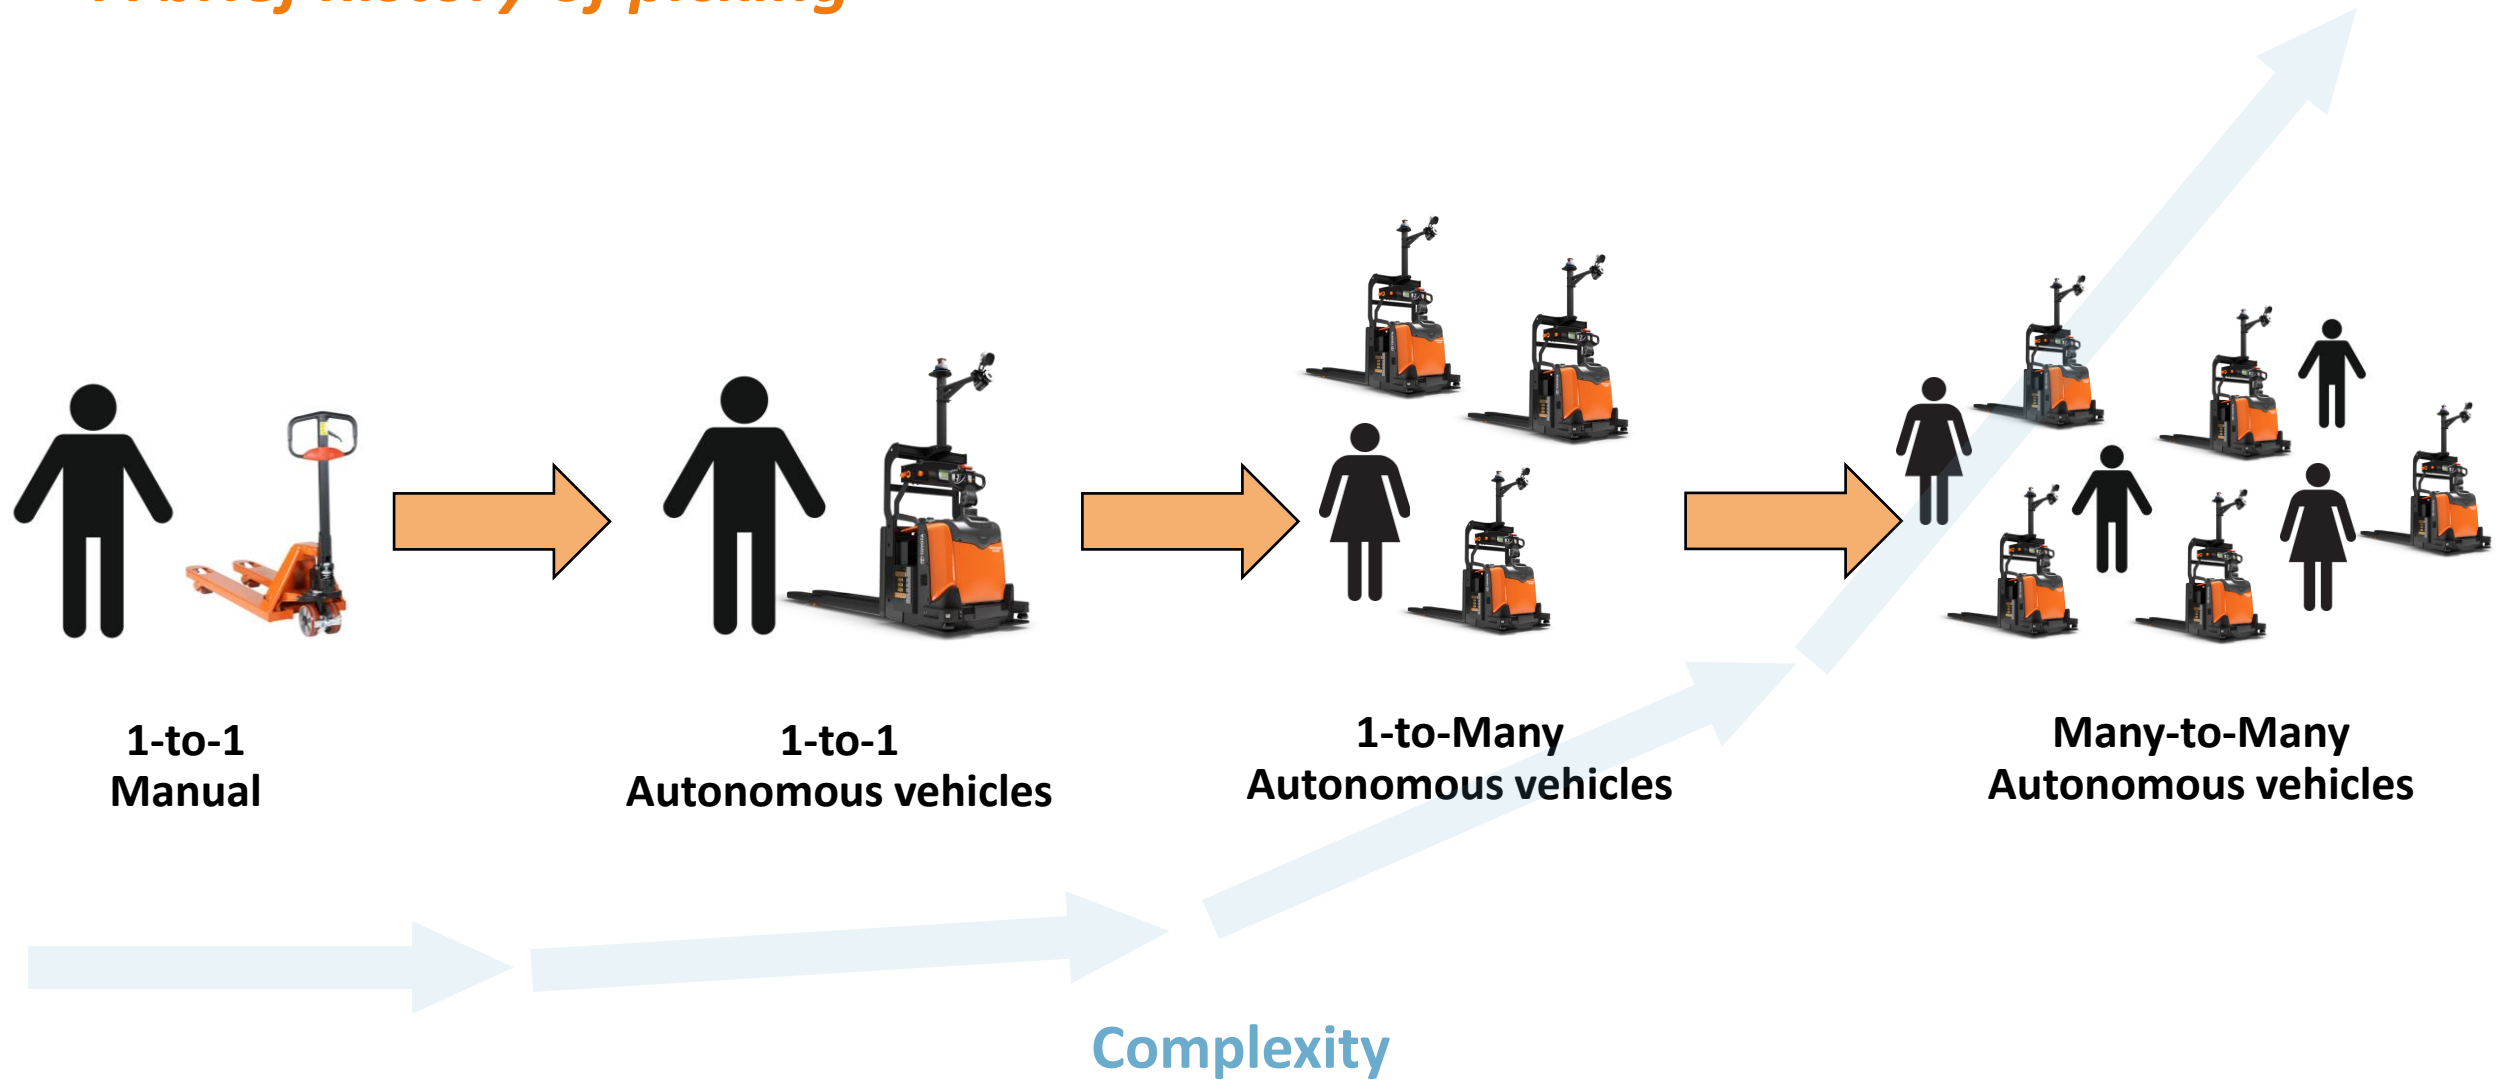


Warehousing impressions




<https://www.youtube.com/watch?v=3d6fo9TIBV0>

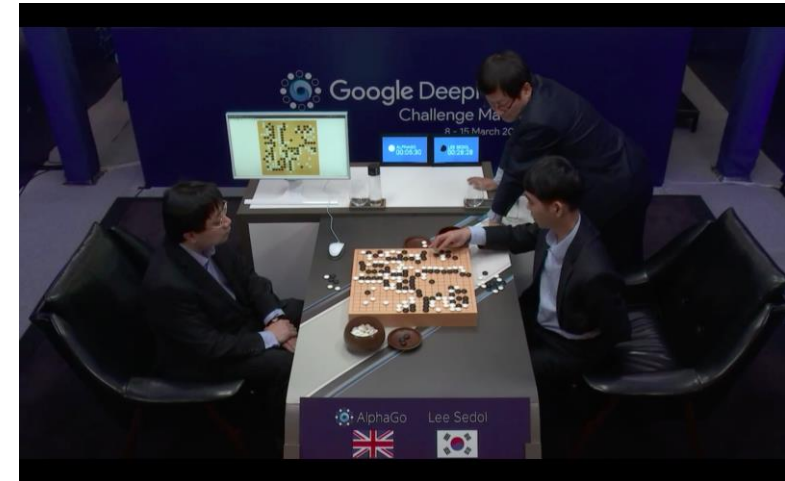

A brief history of picking



Orchestrating chaos




Reinforcement Learning



- › Goal: **learn** a Policy (action to take given the State) that maximizes the expected reward obtained over time
- › RL algorithm can generalize well among problem instances and deal well with **dynamic (stochastic) environments**.
- › In **Deep Reinforcement Learning (DRL)**, **artificial neural networks** are used to approximate the policy, allowing to deal with very large problem instances.


Modelling chaos

- › Grid-world setting where multiple agents move to pick orders that appear at random places.
- › Idle agents are **assigned with a destination** and then follow the **shortest path** to reach it.
- › The objective is to **minimize** the sum of two costs:
 - **Holding costs (10)** for orders that are active on the grid;
 - **Lost costs (100)** for tardy orders, i.e. for orders not picked within the due time.
- › Information about current and possible future picks is in a **flexible lifecycle** model.



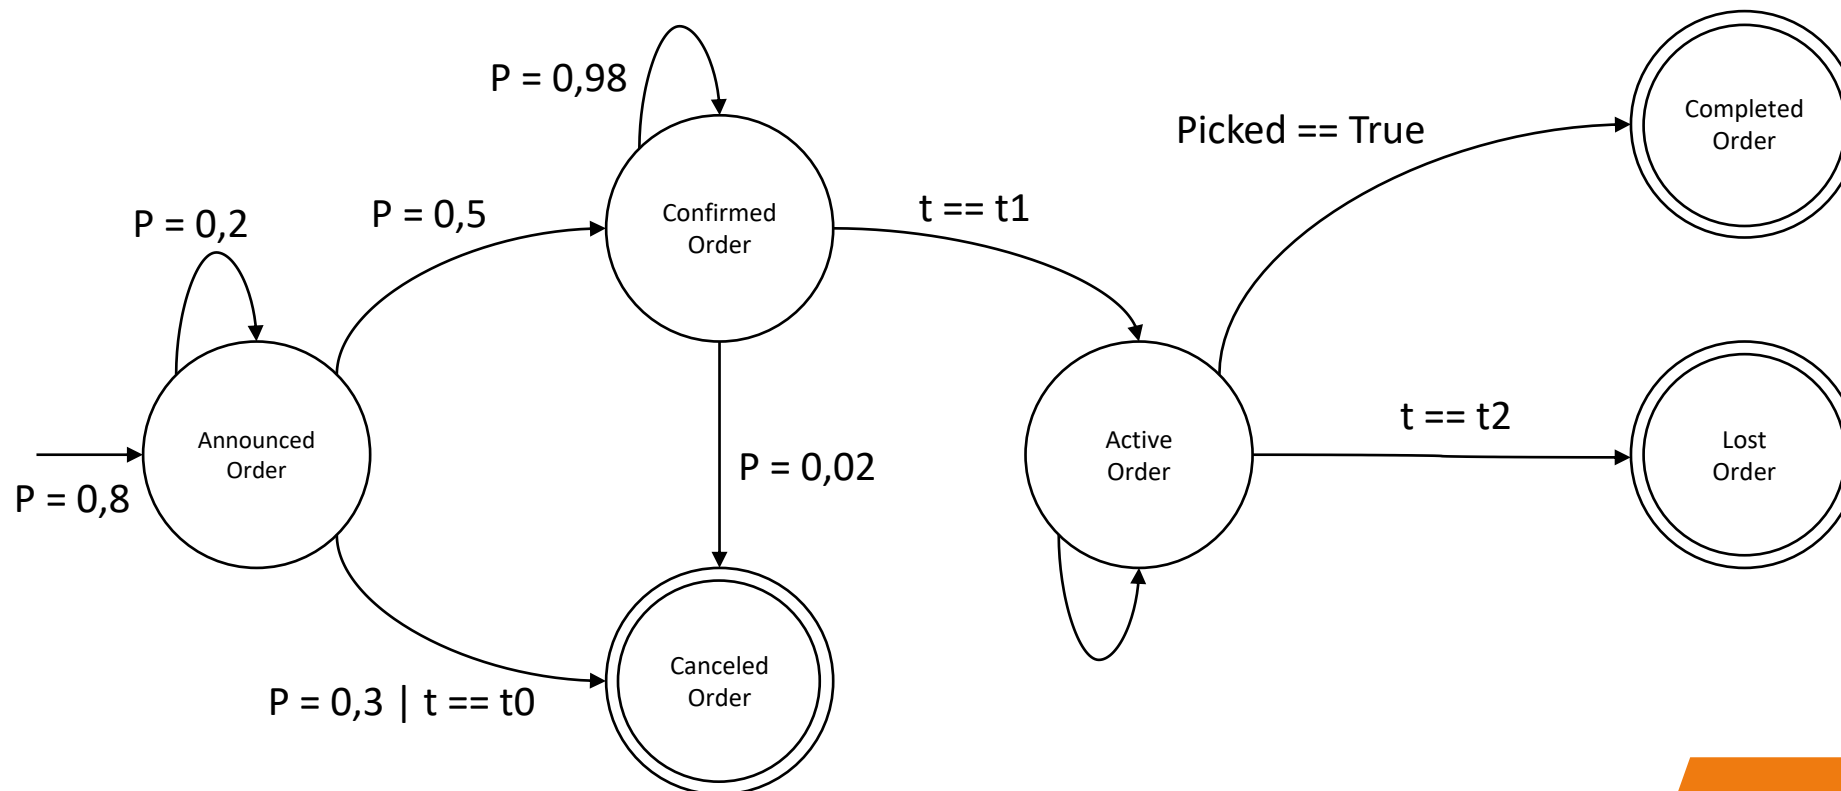
Modelling chaos: order lifecycle

- › **Announced:** not pickable, can become Confirmed, stay Announced or be cancelled (significant probability). Uncertain predictions of possible future locations of AMRs.
- › **Confirmed:** not pickable, will become Active at t1 if not cancelled first (much lower probability). More reliable predictions of possible future locations of AMRs.
- › **Active:** pickable. Must be fulfilled within t2, otherwise they are lost.



Modelling chaos: order lifecycle

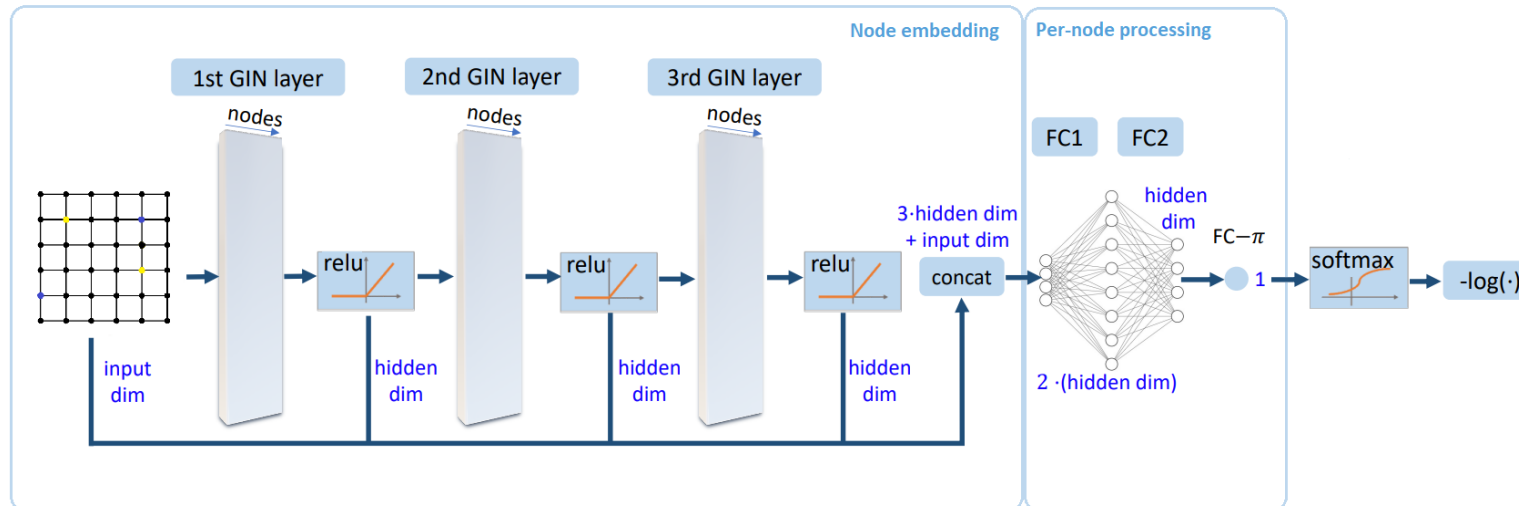
- › **Announced:** not pickable, can become Confirmed, stay Announced or be cancelled (significant probability). Uncertain predictions of possible future locations of AMRs.
- › **Confirmed:** not pickable, will become Active at t_1 if not cancelled first (much lower probability). More reliable predictions of possible future locations of AMRs.
- › **Active:** pickable. Must be fulfilled within t_2 , otherwise they are lost.



Industrial scale solution

Two main components:

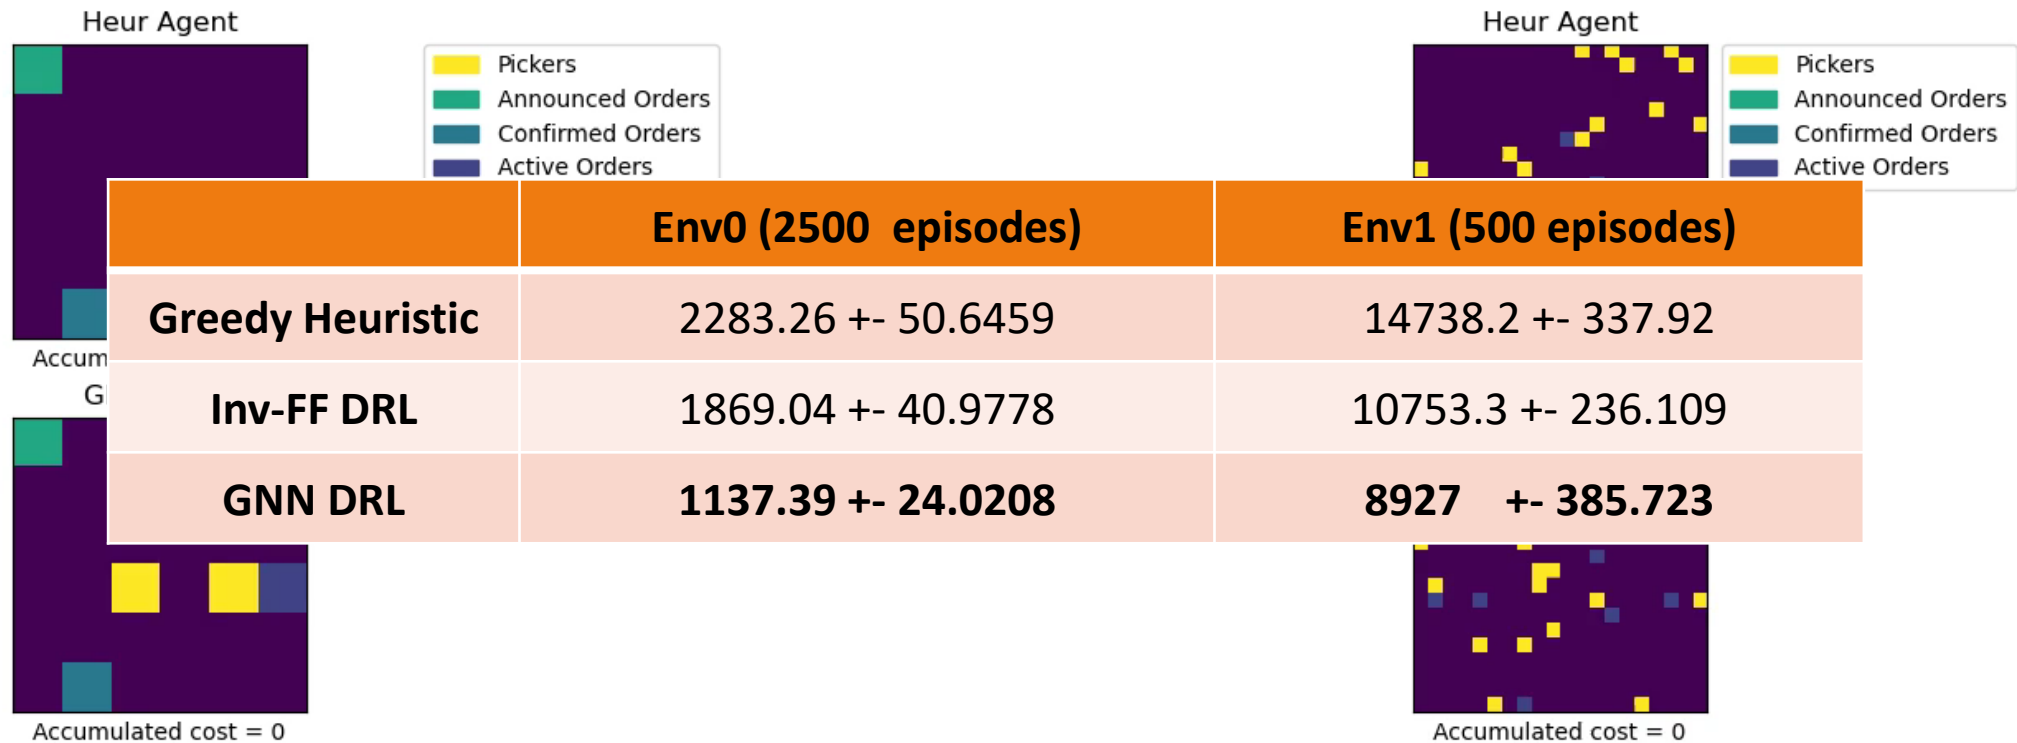
- **Dynaplex:**
 - Framework to easily implement a very fast environment
 - Support for High Performance Computing, allowing to run many simulations in parallel.
- **Graph Neural Networks:**
 - Allow to learn policies for small problem instances that can successfully be transferred to larger ones.



Results

Env0: 6x6, 2 agents, 1 order per step

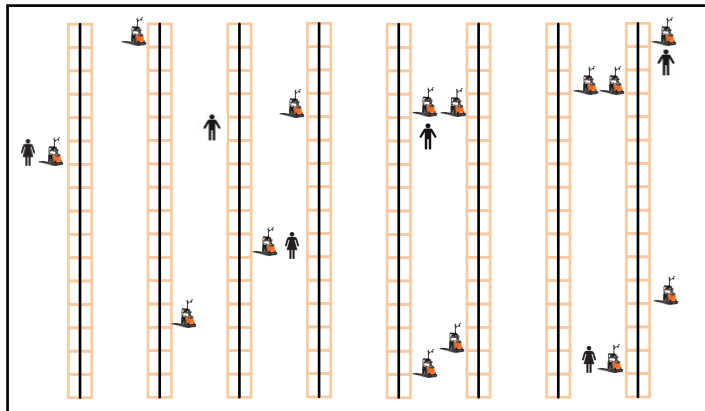
Env1: 20x20, 24 agents, 16 orders per step



Collaboration with academia



Collaborative picking



VANDERLANDE

VANDERLANDE

TU/e EINDHOVEN
UNIVERSITY OF
TECHNOLOGY



Dynaplex



TU/e EINDHOVEN
UNIVERSITY OF
TECHNOLOGY