

UNIVERSITY OF
BIRMINGHAM


Long-Term Autonomy for Logistics in Dynamic Environments

Nick Hawes

<http://nickhaw.es>

@hawesie

School of Computer Science, University of Birmingham, UK





STRANDS



AKADEMIE FÜR ALTERSFORSCHUNG
AM HAUS DER BARMHERZIGKEIT

<http://strands-project.eu>



Transport Systems Catapult, Milton Keynes, UK




months of
autonomous behaviour

no *expert* interventions

requires **verified**
behaviour planning
and control



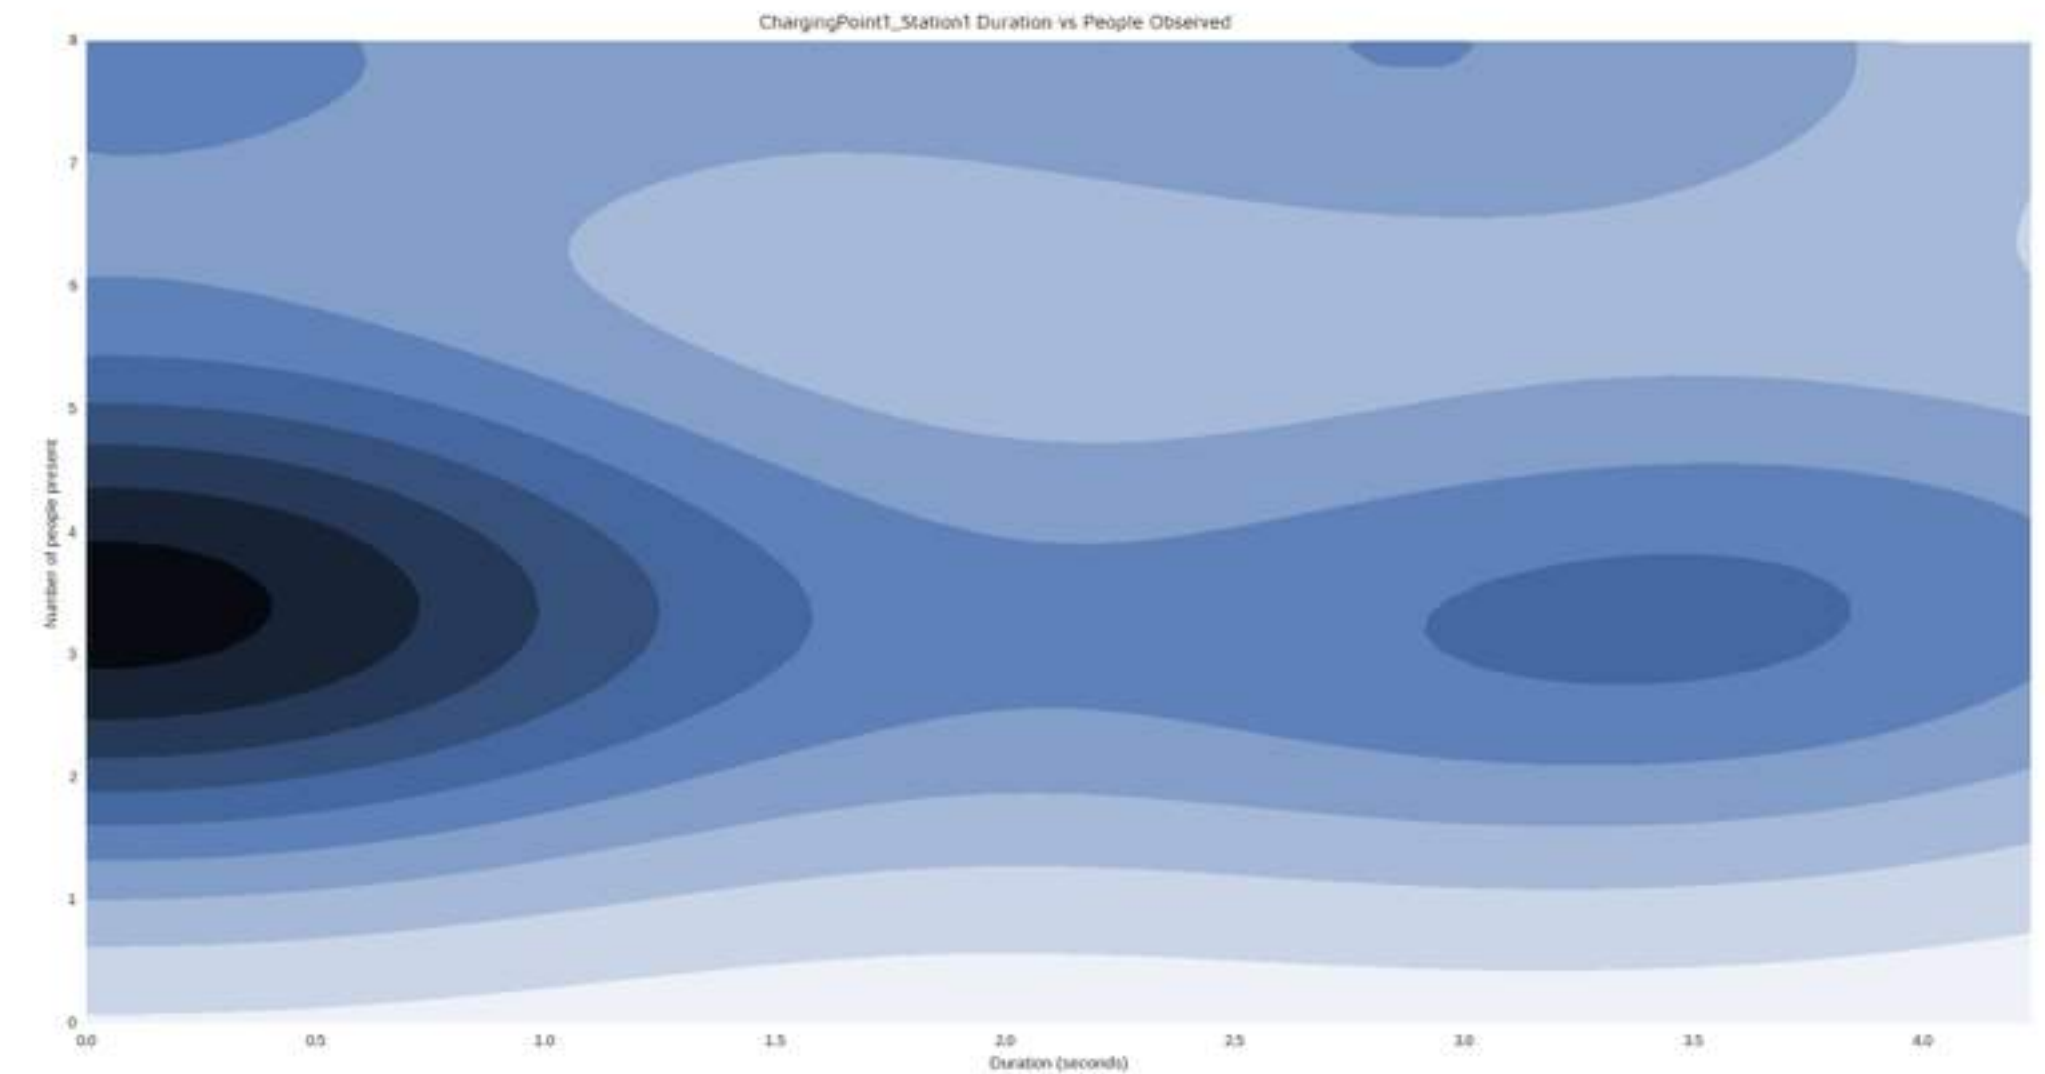
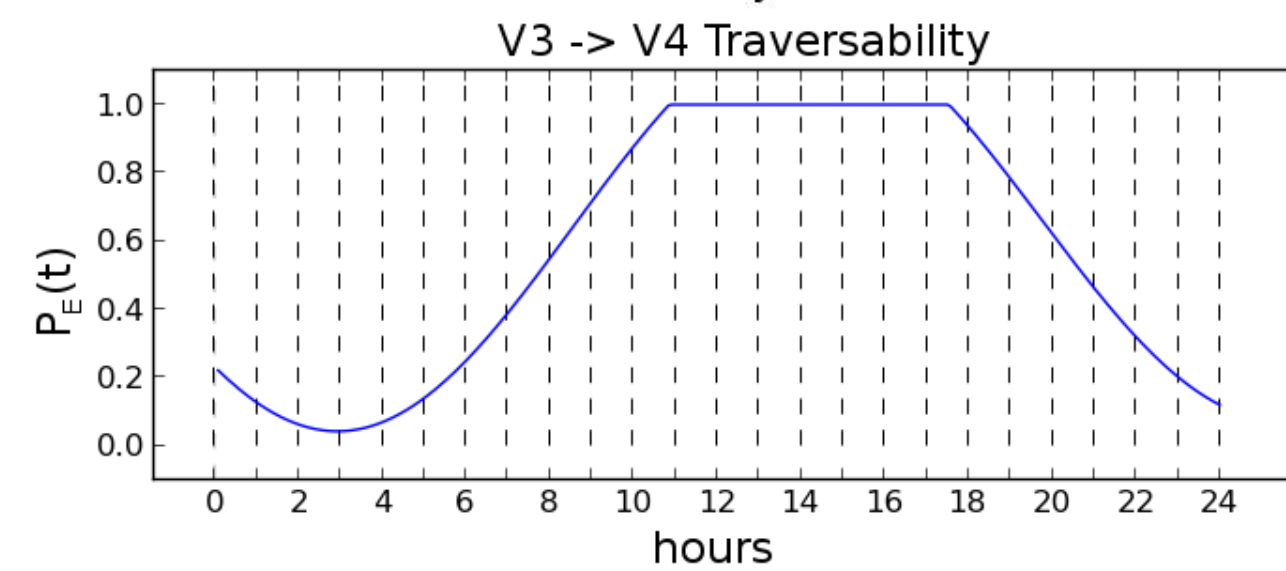
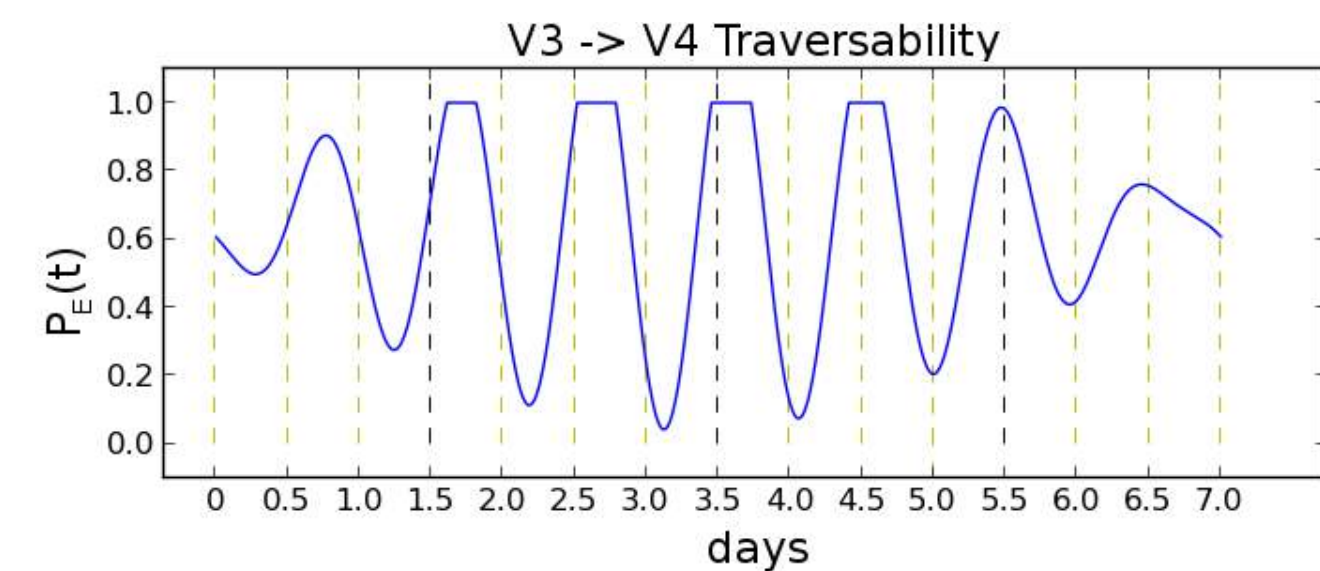
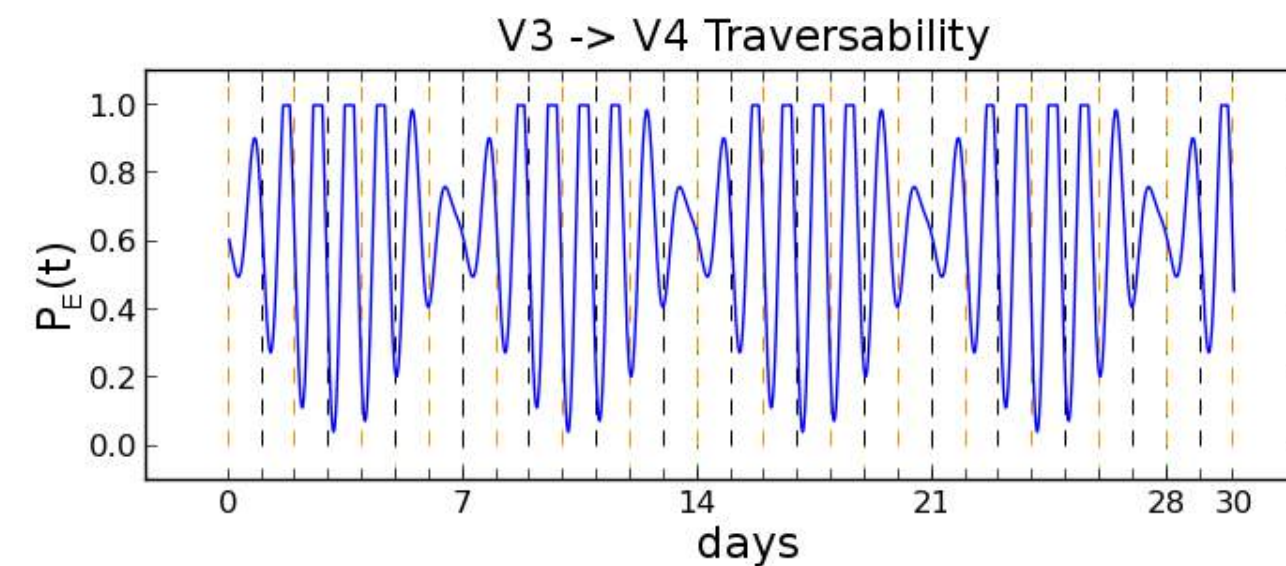
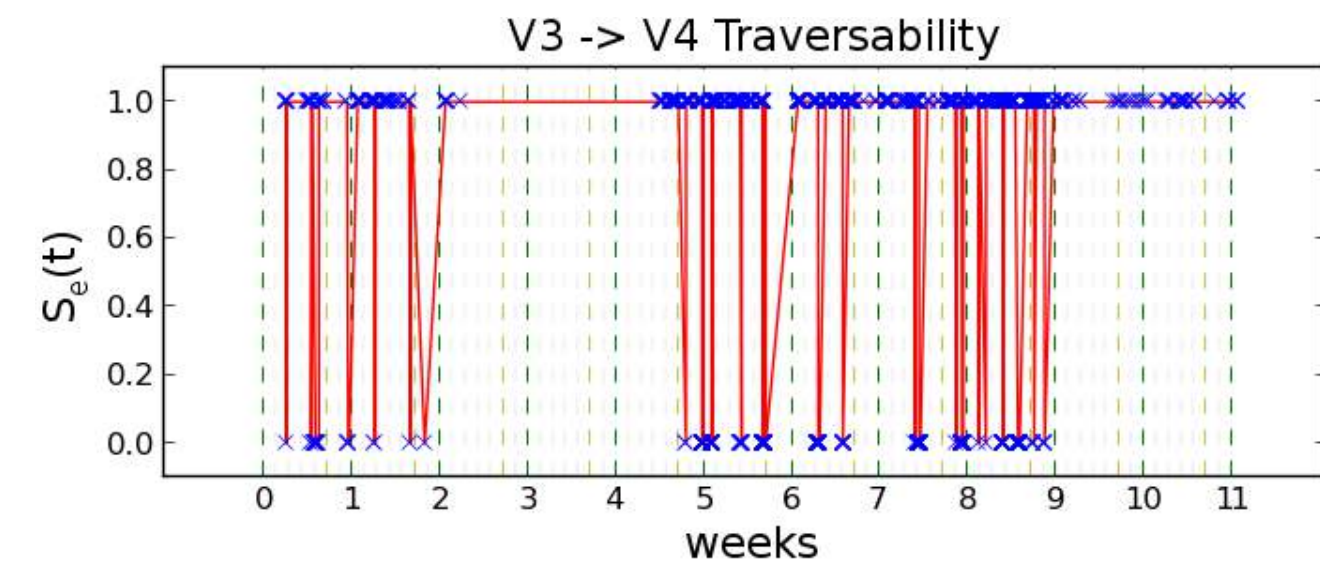
planning world model





Topological Prediction

- With the data, we can build temporal models of:
 - Probability of successful edge traversal
 - Expected time for edge traversal



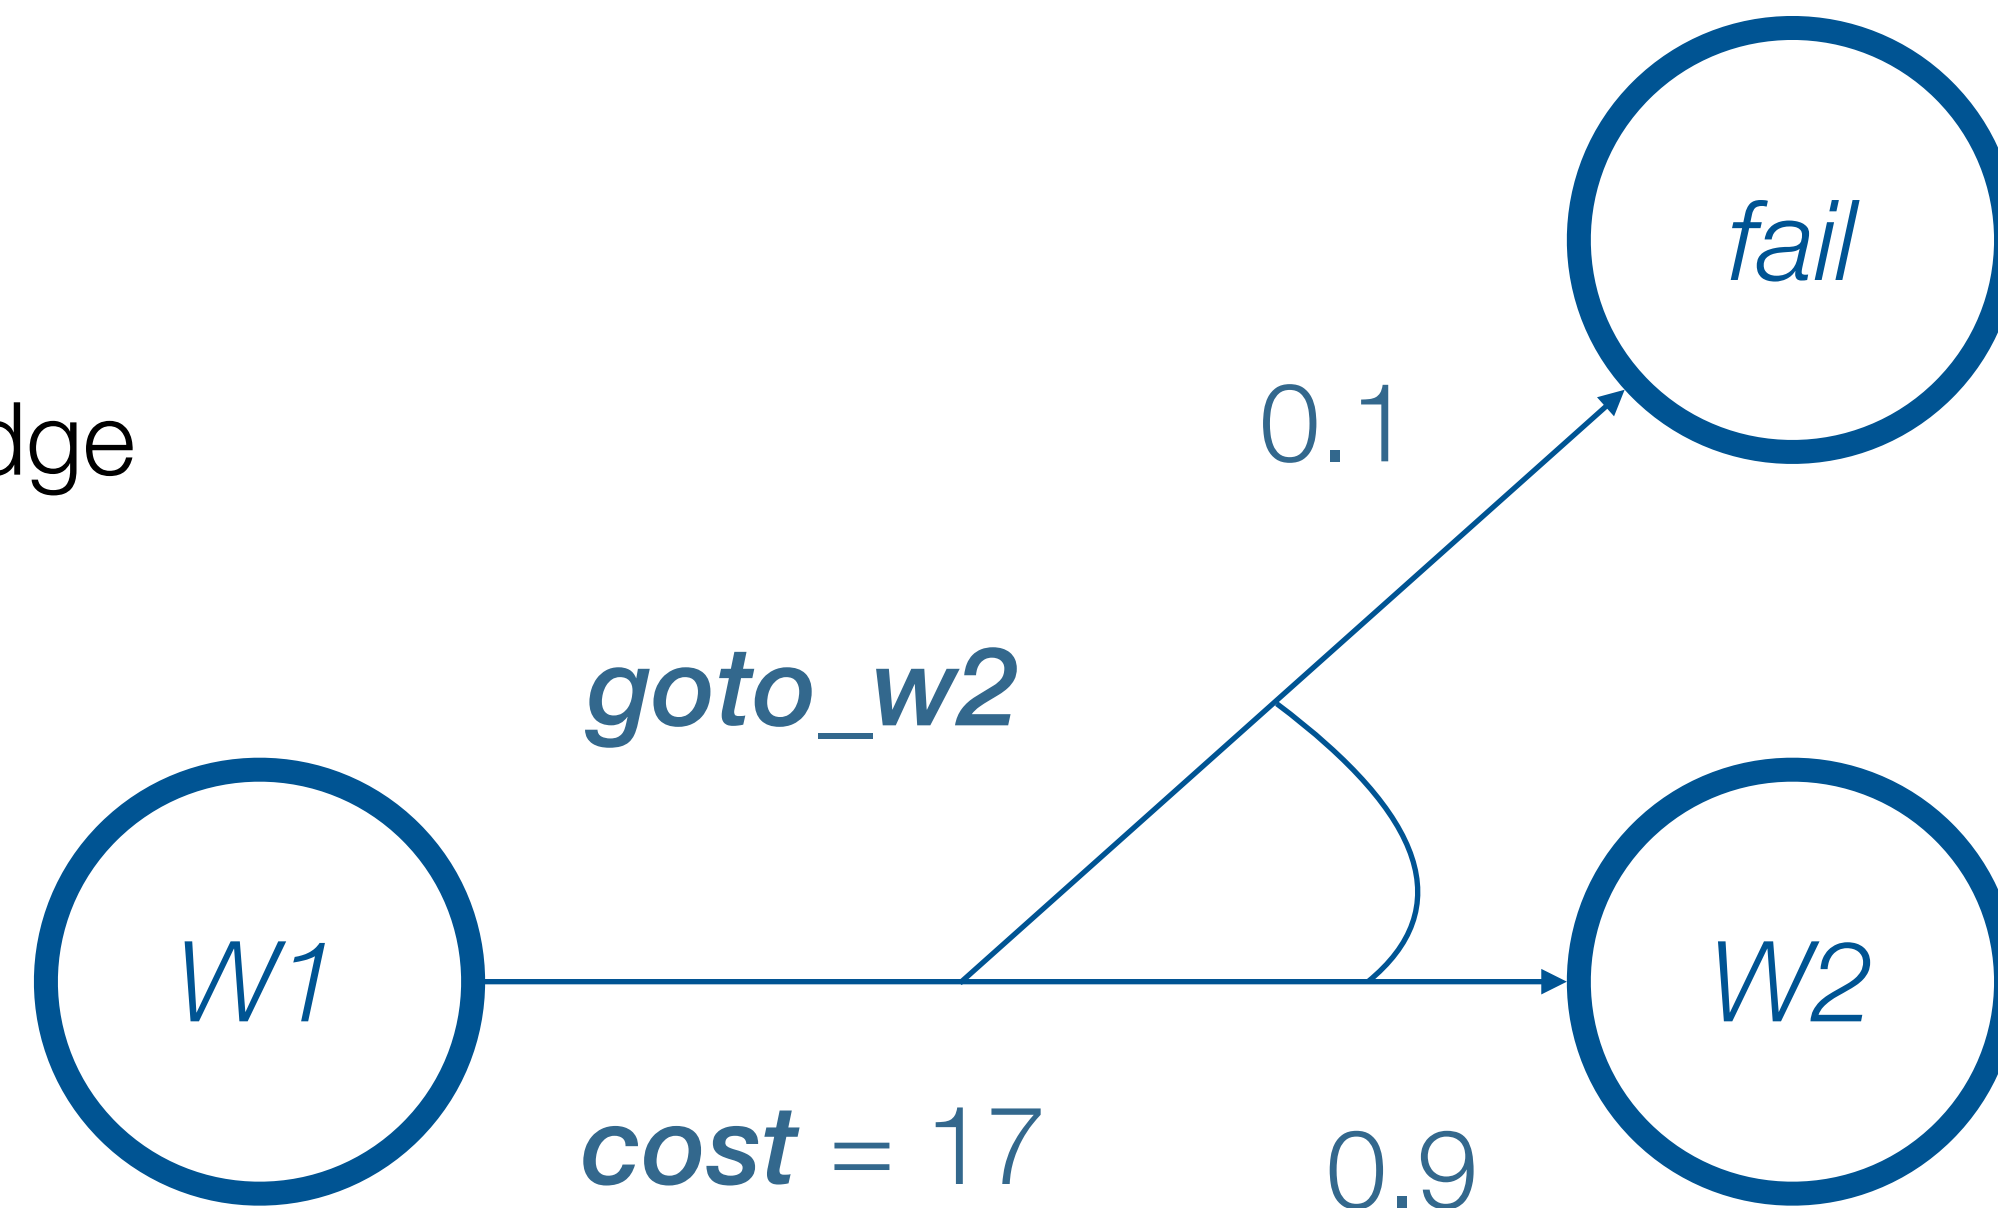
J. Pulido Fentanes, B. Lacerda, T. Krajník, N. Hawes, and M. Hanheide. *Now or later? Predicting and maximising success of navigation actions from long-term experience*. In ICRA, 2015.

MDP Modelling

Navigation Markov Decision Process

- ▶ Use a Markov decision process (MDP) to encode the learned probabilistic models

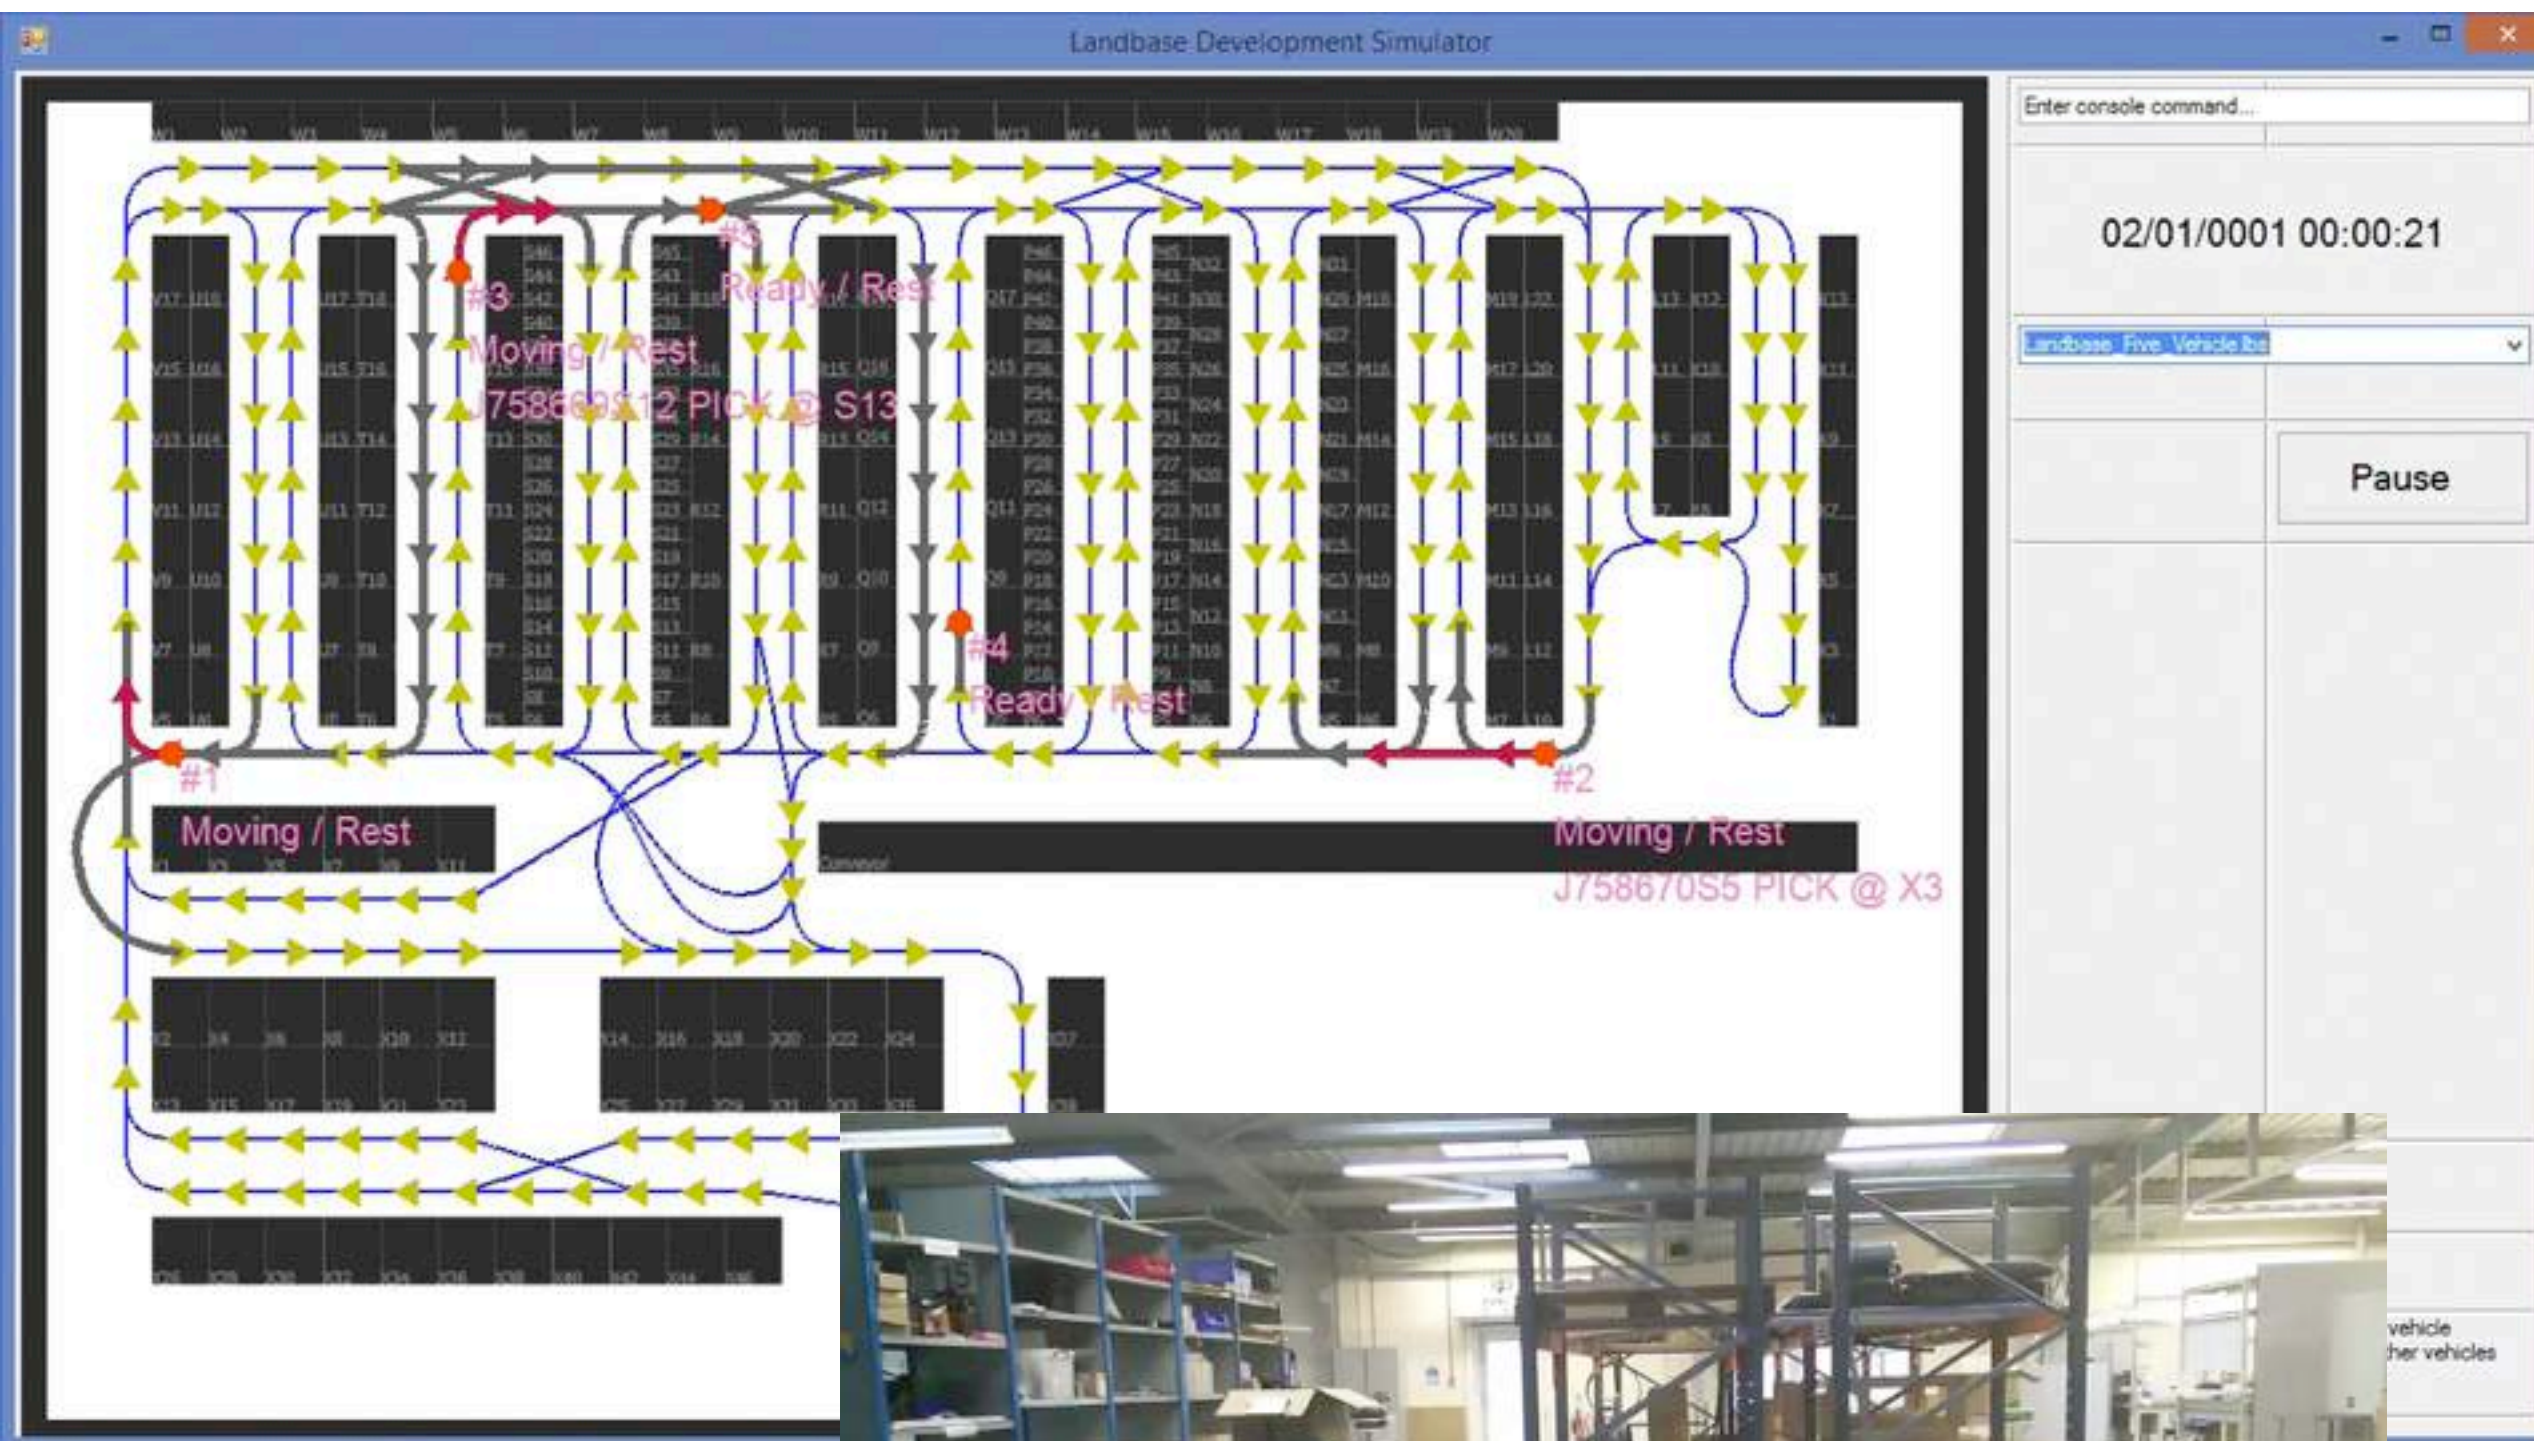
e.g., for a single edge



WHCA* with continuous time and kinematics

TSP and greedy item ordering


FIFO job selection



WHCA* with continuous time and kinematics

TSP and greedy item ordering

FIFO job selection



how can we apply long-term learning in logistics?

- planning to avoid workers
- planning to meet workers

this brings in planning under uncertainty which makes MRPF event harder!

