

Decision Theory is Great (in theory)

(... but can we make it great in practice?)

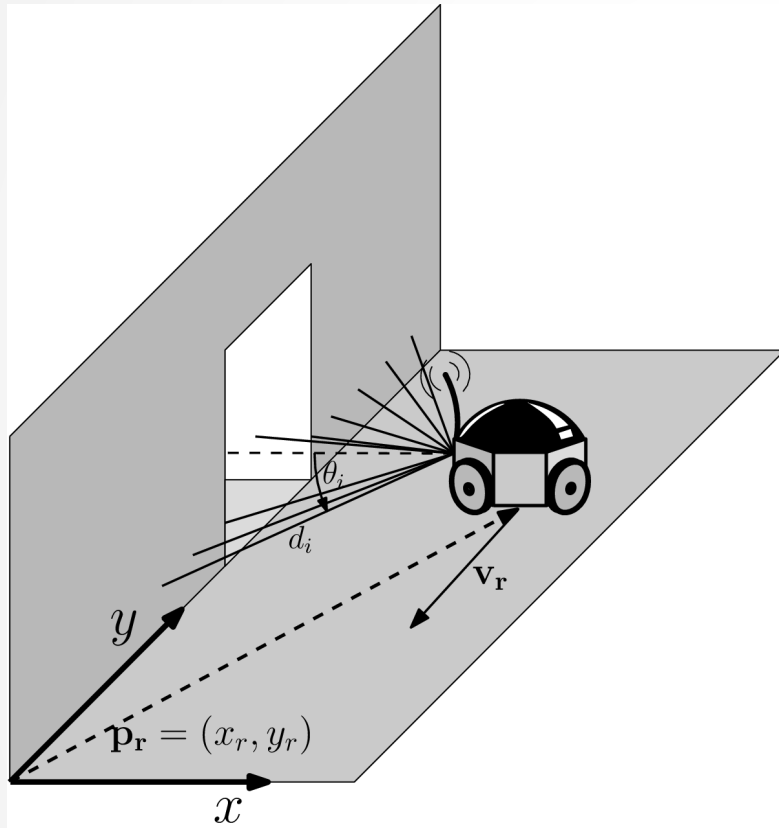
João Messias

University of Amsterdam

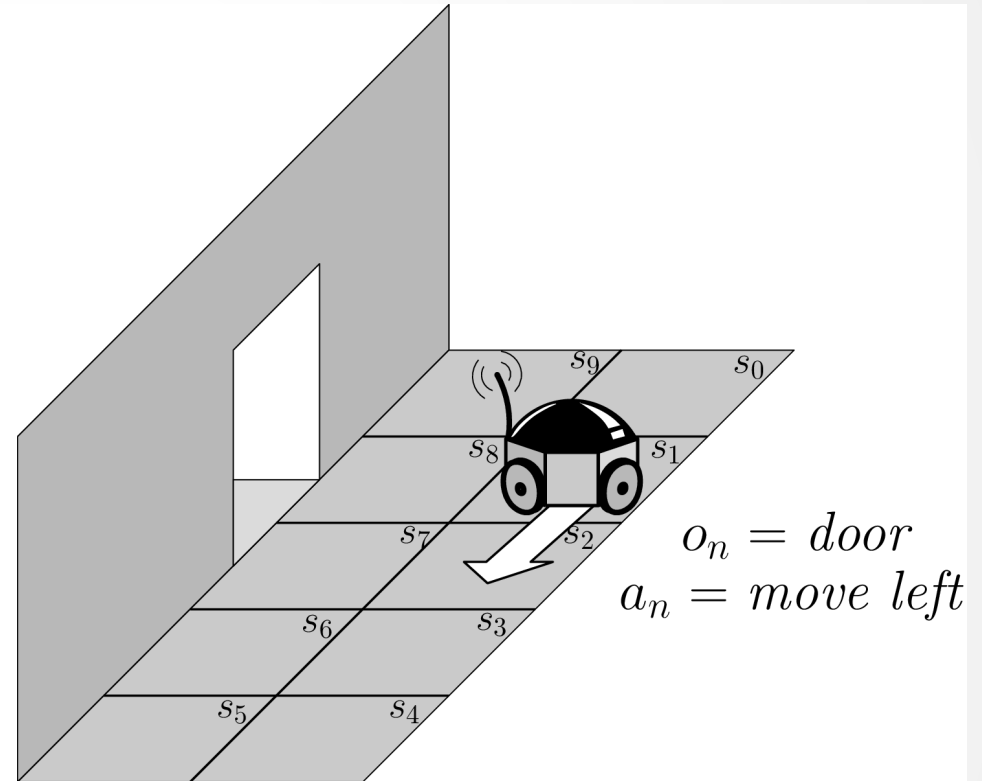
Decision Theory is Great (in theory)

- Powerful methods for planning under uncertainty and reinforcement learning;
- In (social) robotics, abstraction / discretisation usually necessary;
- Requires domain-specific modeling work;

Decision Theory is Great (in theory)



Real problem

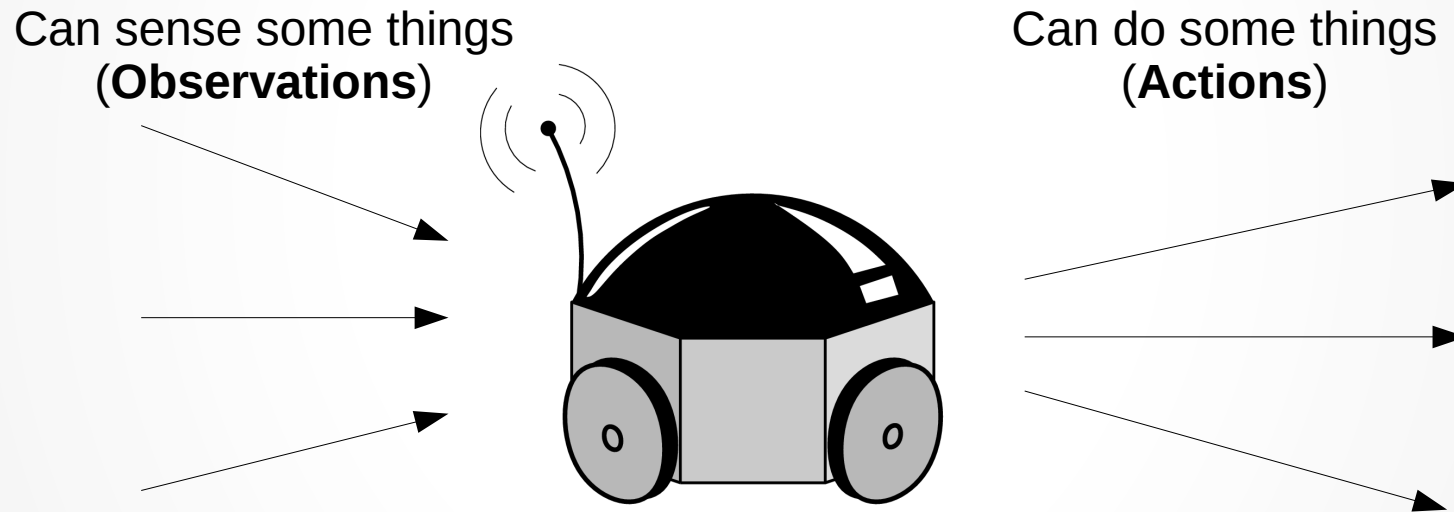


DT Model

Deployment across different domains is difficult

Decision Theory is Great (in theory)

An agent-centric abstraction is still OK



Does not depend on the environment – easy to deploy in the real world!

Decision Theory is Great (in theory)

Learning how to map from (sequences of) observations directly to actions?

General RL

e.g. AIXI (Hutter, 2005)

Decision Theory is Great (in theory)

End-to-end approaches e.g. RNNs would be great...
..but usually require lots of data to train

Can't run robots for many trials

Can't run humans for many trials

Are there practical approaches to this problem?