

Endowing Robots with capabilities to Learn from non-expert Human Teachers

Carlos Celemin

Cognitive Robotics Department
Delft University of Technology



EUROPEAN
ROBOTICS FORUM

ERF 2020 Workshop, Human Robot Collaboration and AI for
Industrial Applications

Introduction

Imitation Learning



Expert demonstrators required,
but it is not always possible!



Introduction

Interactive Imitation Learning


Human teachers providing feedback while observing the agent-environment interaction

Evaluative Domain

Reinforcements, preferences, rankings.

Action Domain

Corrective Demonstrations (e.g. Active Learning), Data Aggregation, relative corrections.






Teleoperation

4x

Deep COACH

Car Racing Training

Learning Movement Primitives



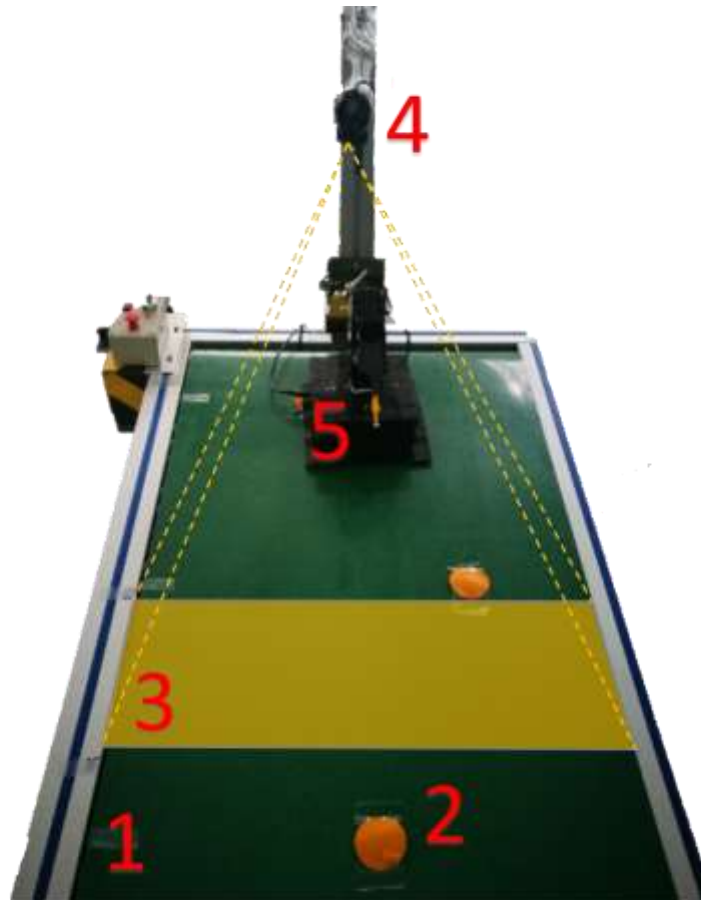
End-to-end Learning

Duckie Racing Training

COACH: COrrective Advice Communicated by Humans

End-to-end policy learning from non-expert teachers


Orange selection in a conveyor belt



Deep COACH




Learning Interactively to solve Ambiguities



Conclusions

- Occasional corrections allow non-experts to teach robots with super-human performance: *Robot learner outperforms human teacher.*
- Workload is decreased with occasional corrections.
- Learning time is reduced and policy performance is improved.
- Non-expert operators in factories and end-users in household environments are able to program robots for complex tasks.





Endowing Robots with capabilities to Learn from non-expert Human Teachers

Carlos Celemin

Cognitive Robotics Department
Delft University of Technology



EUROPEAN
ROBOTICS FORUM