Learning feedback controllers in robotics

Hakan Girgin

09.11.2021

RESEARCH INSTITUTE

Learning in robotics

- What to learn ?
- From pure machine learning perspective towards pure physics perspective

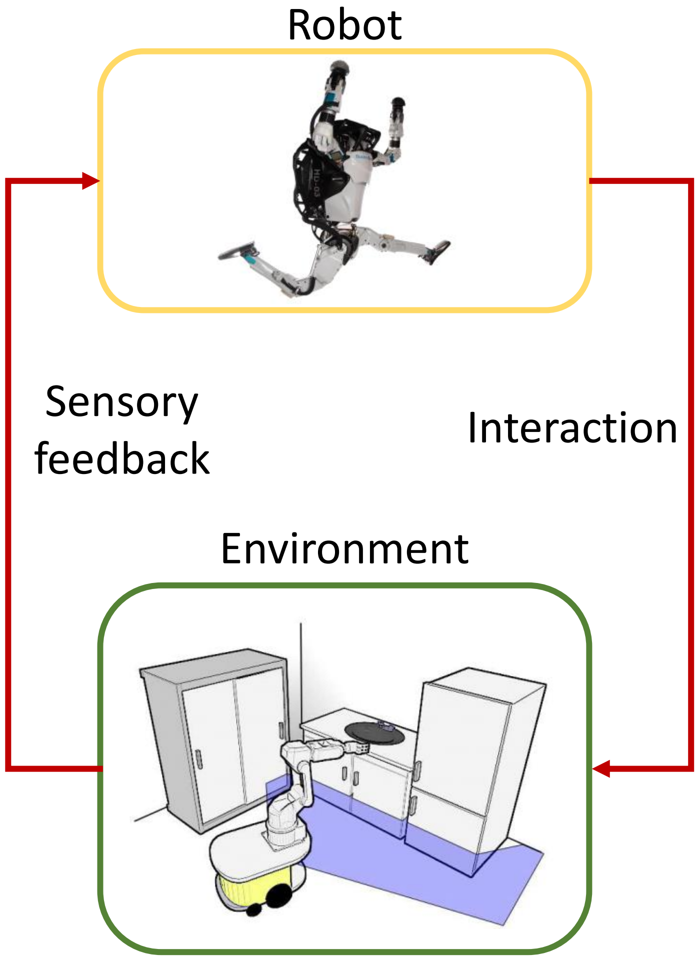

Input	Output	Required data
Camera images	Grasping	800,000 grasp attempts from 14 manipulators*

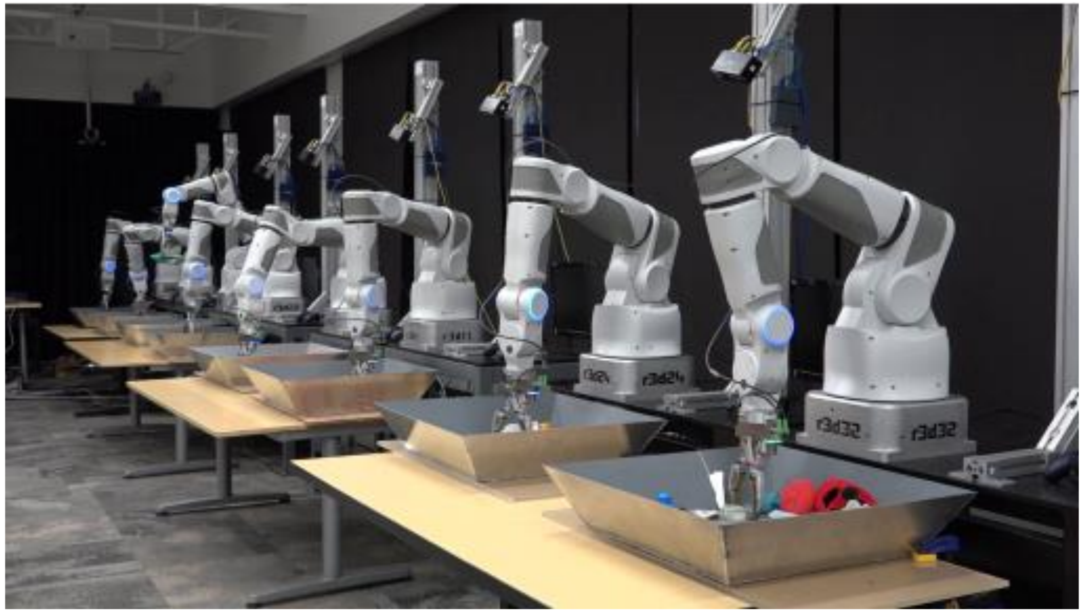

Can physics models be useful?

* Levine, Sergey, Peter Pastor, Alex Krizhevsky, Julian Ibarz, and Deirdre Quillen. “Learning Hand-Eye Coordination for Robotic Grasping with Deep Learning and Large-Scale Data Collection.” The International Journal of Robotics Research 37, no. 4–5 (April 2018)

Learning in robotics – Environment models

- What to learn ?
- From pure machine learning perspective towards pure physics perspective

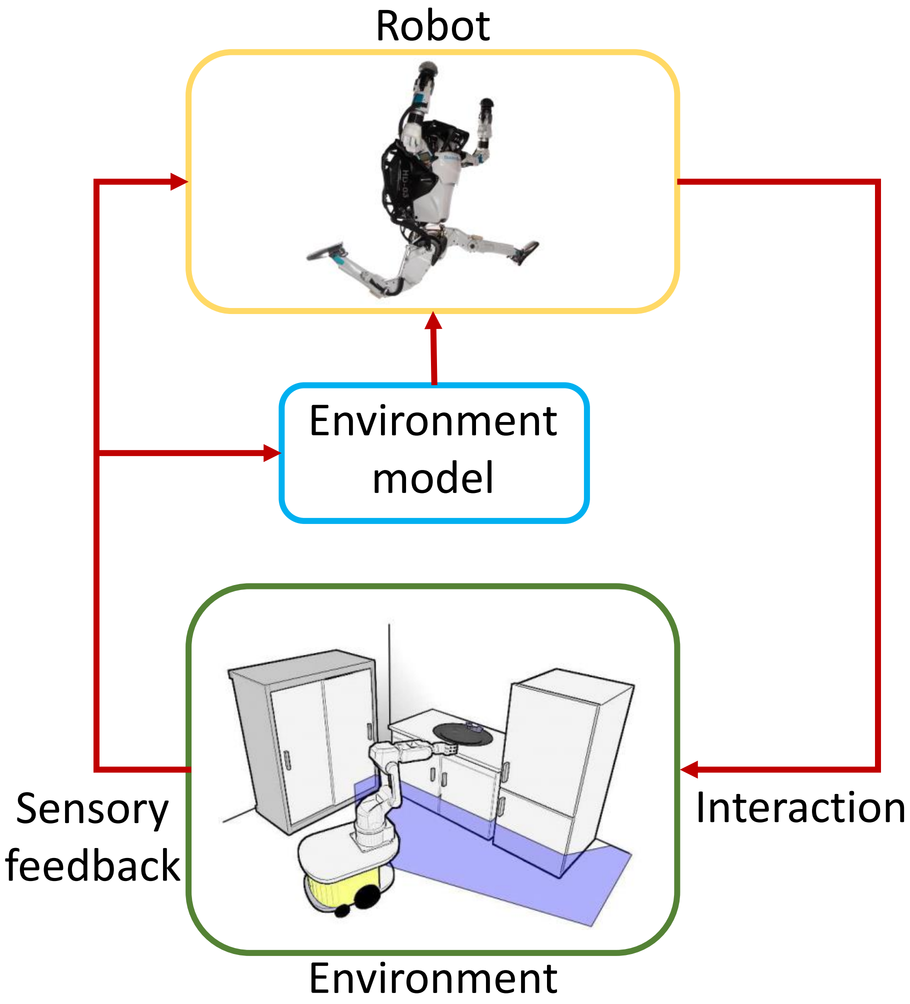

TossingBot*

- Using ballistic physics of an object with a learnable residual part

*Zeng, Andy and Song, Shuran and Lee, Johnny and Rodriguez, Alberto and Funkhouser, Thomas, TossingBot: Learning to Throw Arbitrary Objects with Residual Physics, Proceedings of Robotics: Science and Systems (RSS), 2019.

Learning in robotics – Interaction models

- What to learn ?
- From pure machine learning perspective towards pure physics perspective

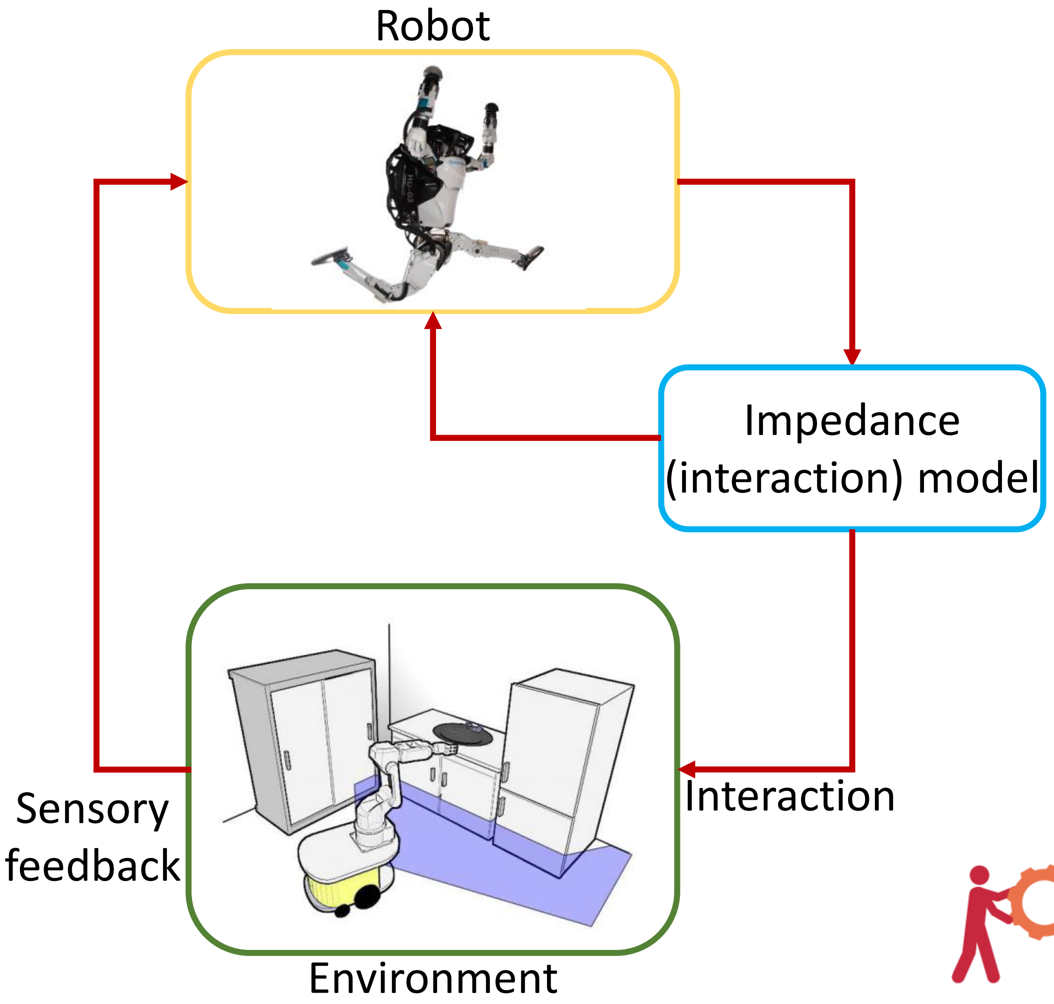

- We can **learn/design interaction models** that can achieve contact-rich tasks:

COLLABORATE

Learning in robotics – Interaction models

- What to learn ?
- From pure machine learning perspective towards pure physics perspective

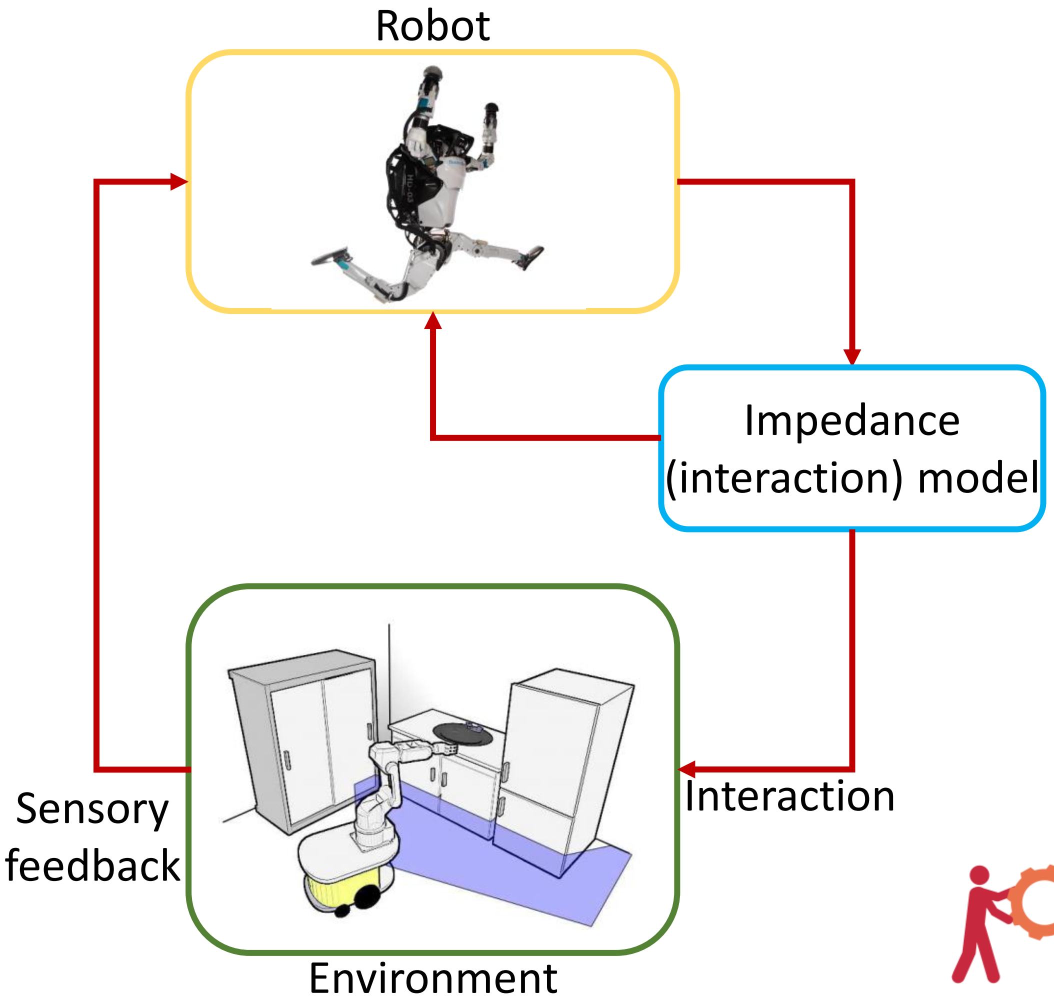

- We can **learn/design interaction models** that can achieve contact-rich tasks:

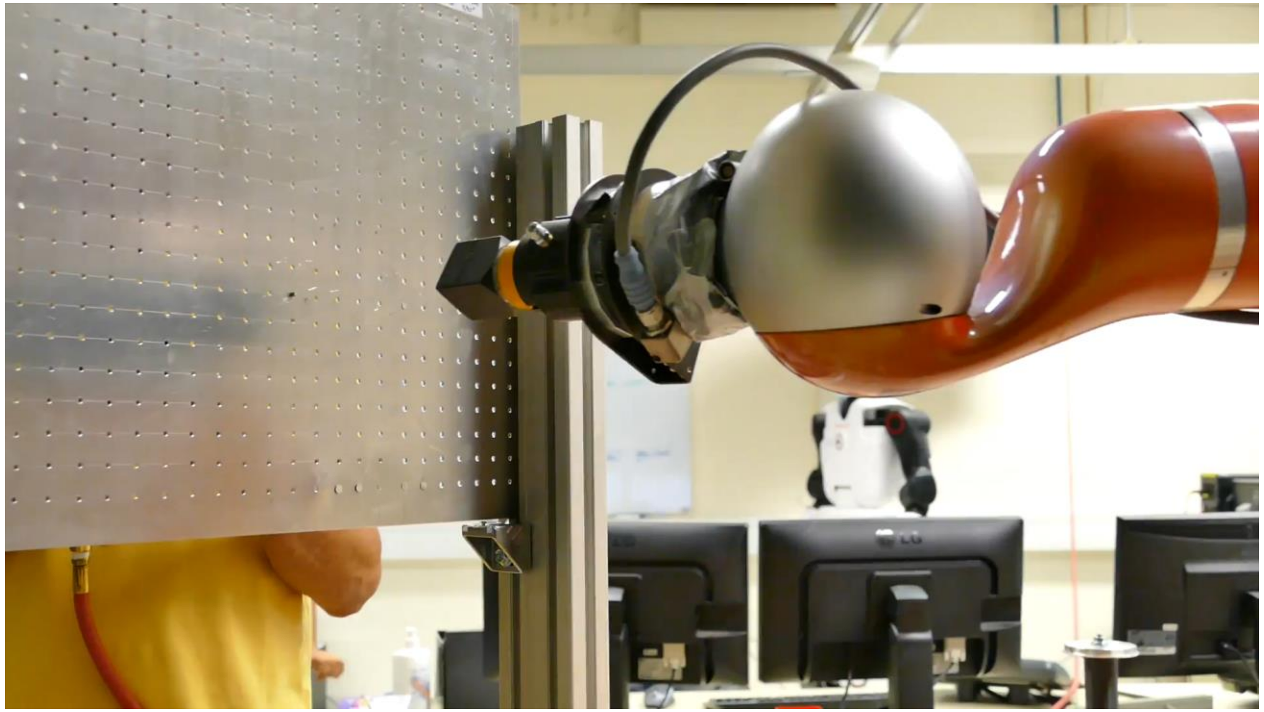

COLLABORATE

Learning in robotics – Collaborative riveting*

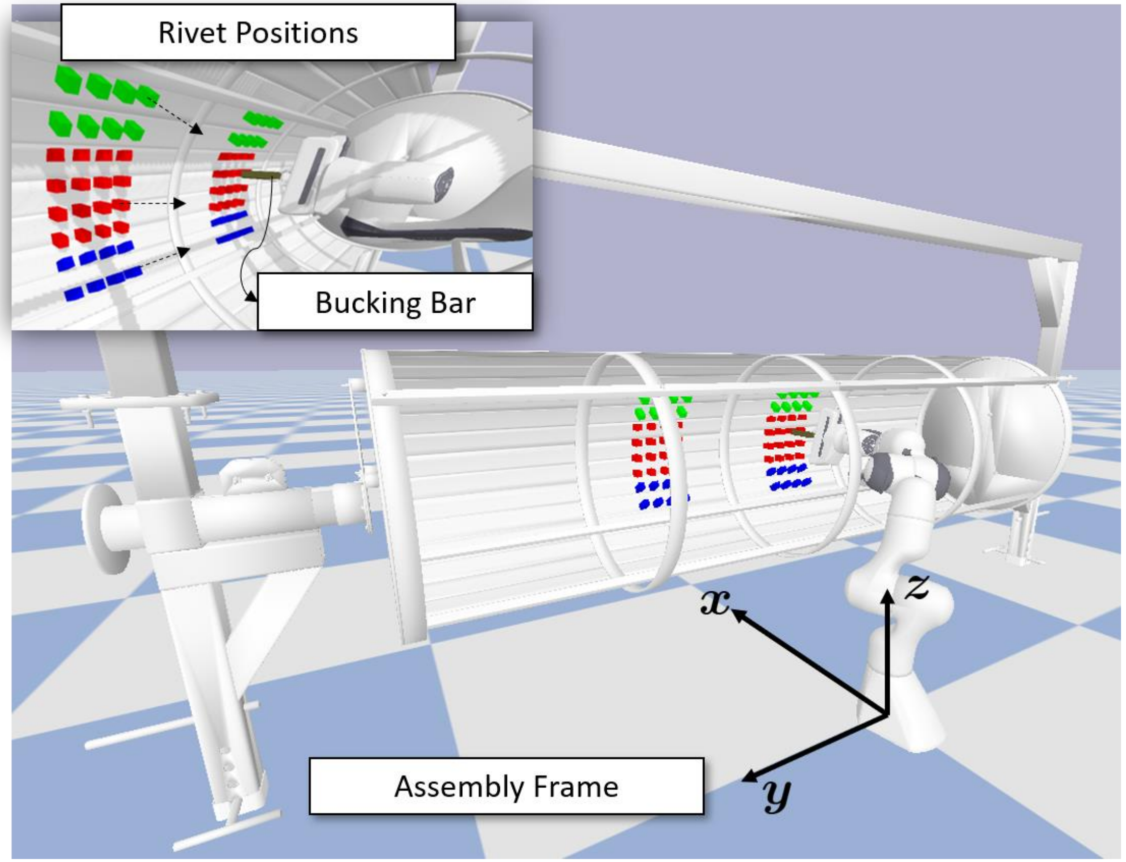

TABLE I: Reduction of the median displacement in %

Frequency (rad/s)	Top	Middle	Bottom
70	82.9	67.79	78.17
100	51.8	54.94	56.4
150	62.52	50.98	50.74

COLLABORATE

* Hakan Girgin, Teguh Santoso Lembono, Radu Cirligeanu and Sylvain Calinon, Optimization of robot configurations for motion planning in industrial riveting, in: Proc. IEEE Intl Conf. on Advanced Robotics (ICAR), 2021

Learning in robotics – towards lower level models

- What to learn ?
- From pure machine learning perspective towards pure physics perspective

Controller = f(sensory feedback, robot states, desired reference states)

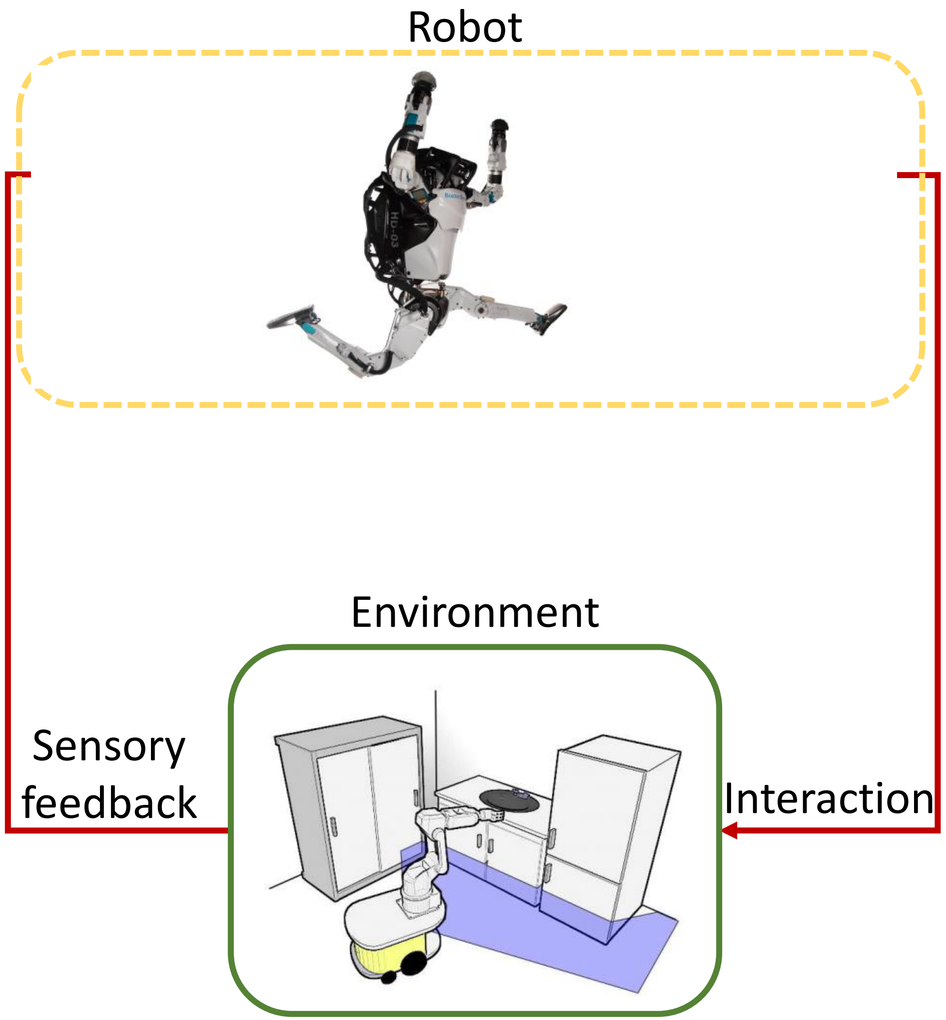

- The control field is all about designing controllers such that it performs the desired task.
- These are parametrized functions of the sensory feedback, robot states and the desired reference.
- What can we gain by designing good controllers?

  - Robustness to uncertainties in the environment, sensory measurements noises, unexpected perturbations
  - **Generalizations** to different environments

Learning in robotics – model-free control

- What to learn ?
- From pure machine learning perspective towards pure physics perspective

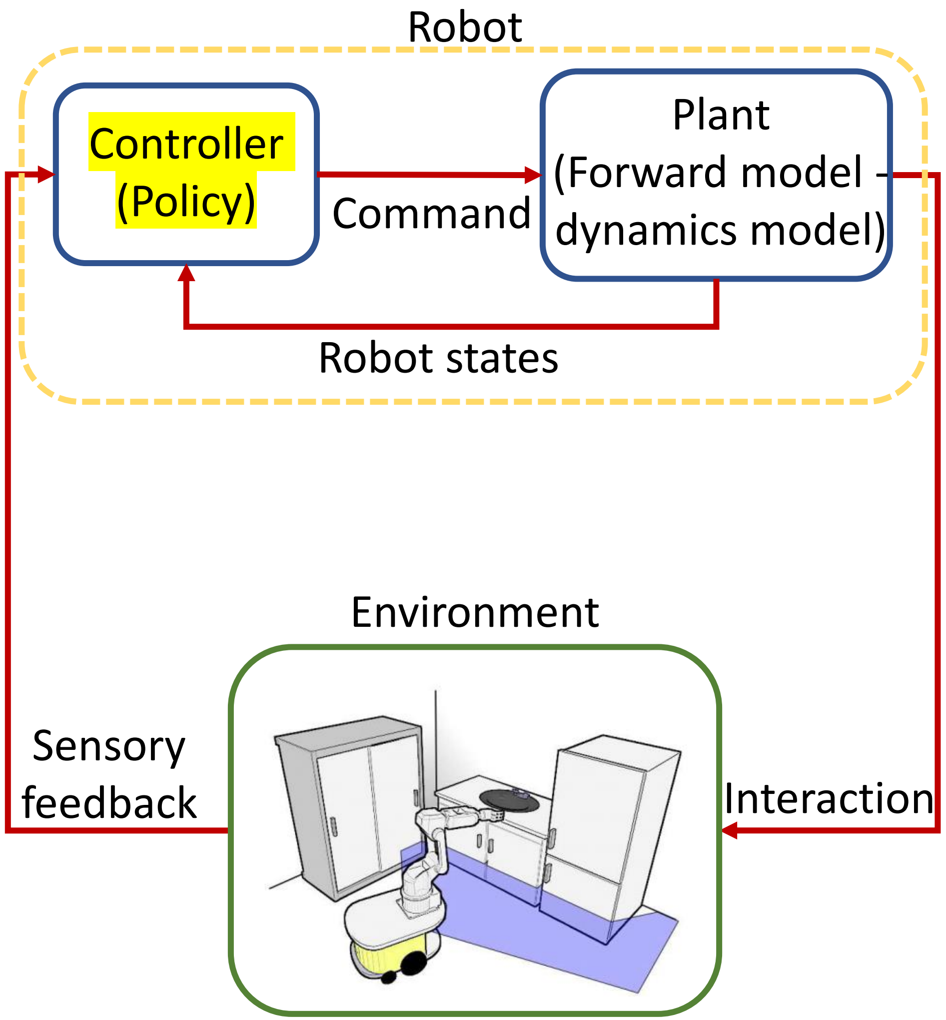

- Controller parameters are determined/learned without using the plant model.
- One typical approach is to gather data from kinesthetic teaching and learn a policy.
- Easier to learn but need more data than model-based counterpart.
- One obvious question is where to get demonstrations so that we minimize this need for "a lot of demonstrations".

Learning in robotics – active learning of model-free controllers

- Actively request demonstrations to learn a better control policy without requiring large amount of data*.

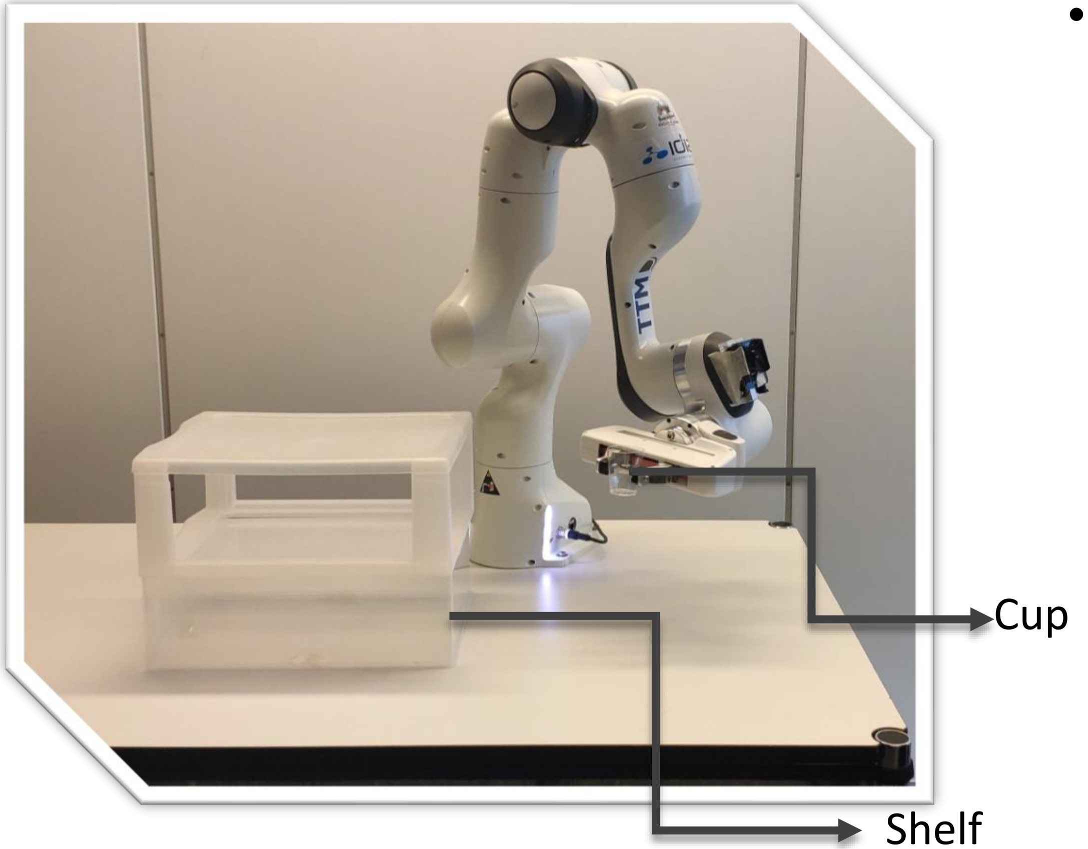

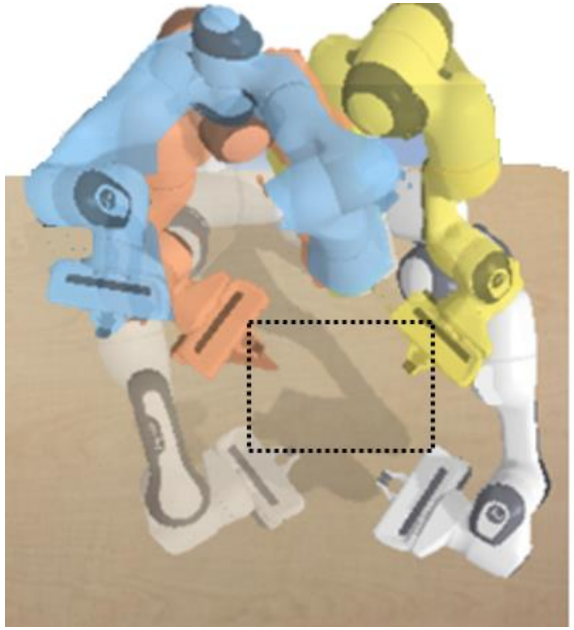

Demonstrations

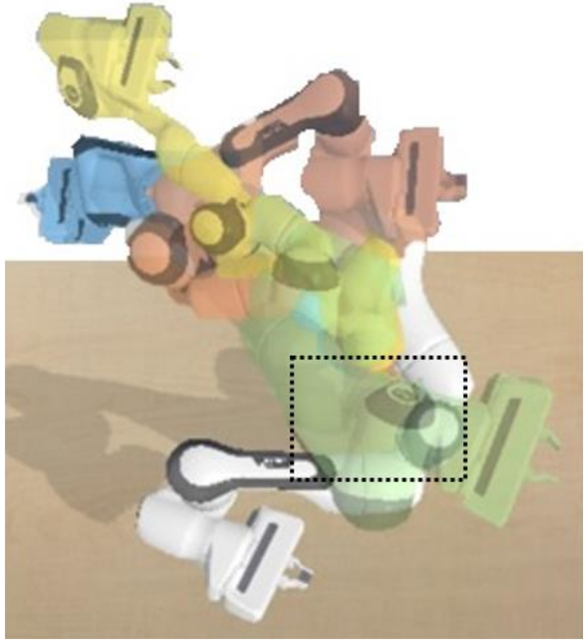

Requested demonstrations

*H. Girgin, E. Pignat, N. Jaquier, S. Calinon, **Active Improvement of Control Policies with Bayesian Gaussian Mixture Model**. In Proc. of IEEE/RSJ Intl Conf. on Intelligent Robots and Systems (IROS), 2020

Learning in robotics – model based control

- What to learn ?
- From pure machine learning perspective towards pure physics perspective

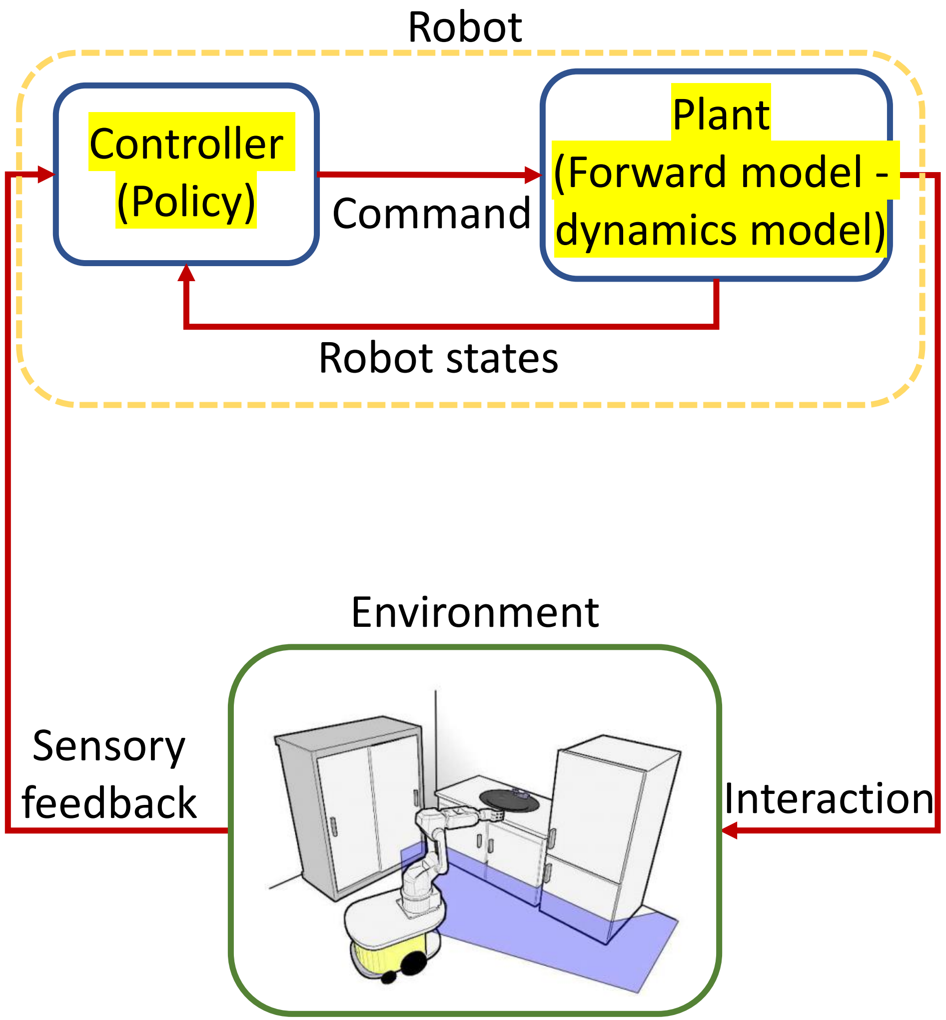

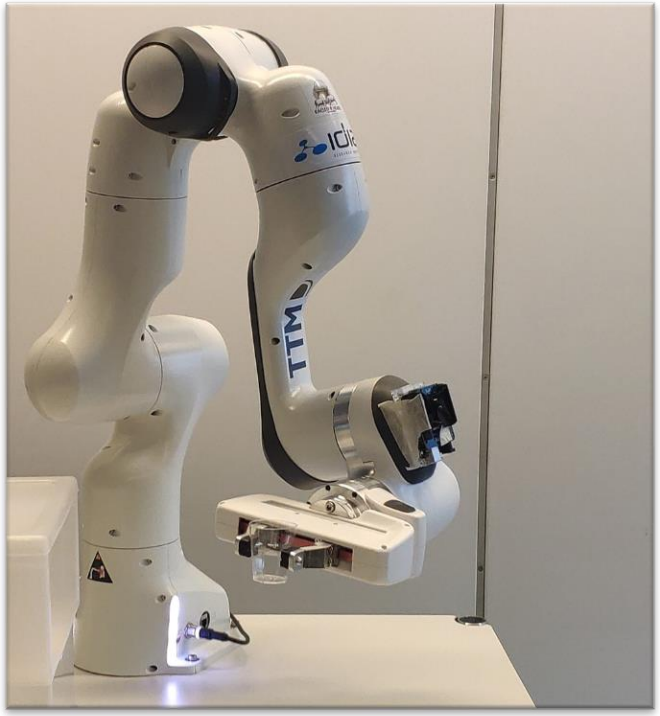

- **Optimal control** optimizes a cost function related to the robotic task with the constraint that the controller satisfies the plant dynamics (LQR, iLQR, MPC) (see last week's TAM)

Learning in robotics – model based control

Feedback controller derived from learning on the trajectory level: learning generalizable reference trajectory generating functions.*

Feedback controller derived from learning on the trajectory level but subject to plant dynamics by imitation learning.**

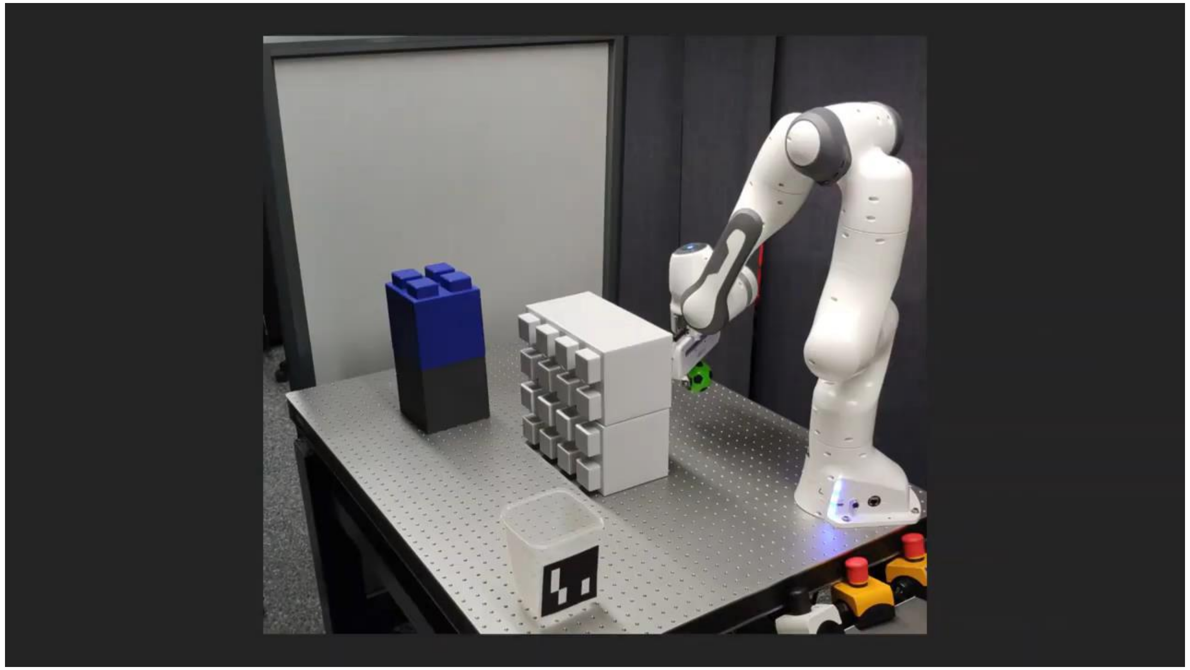

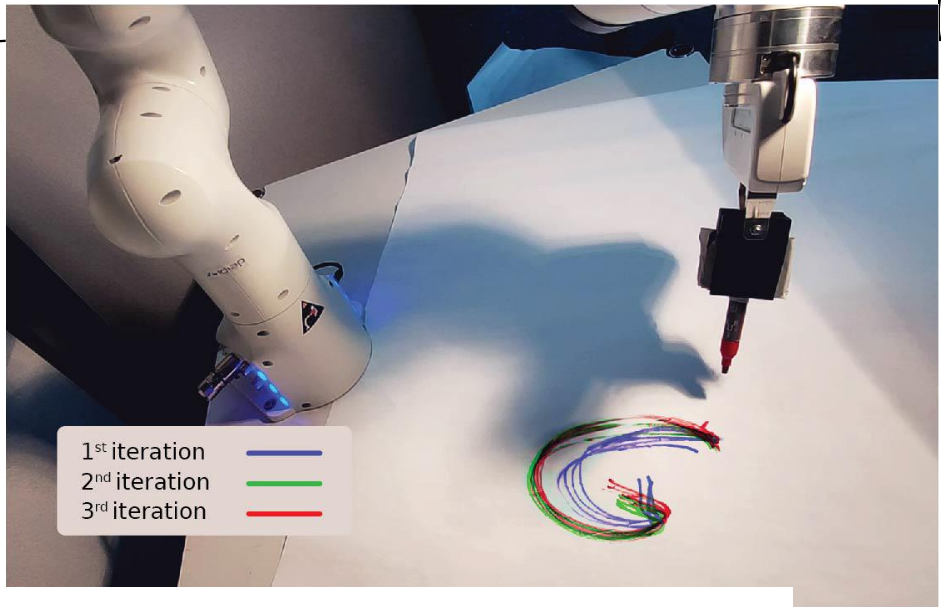

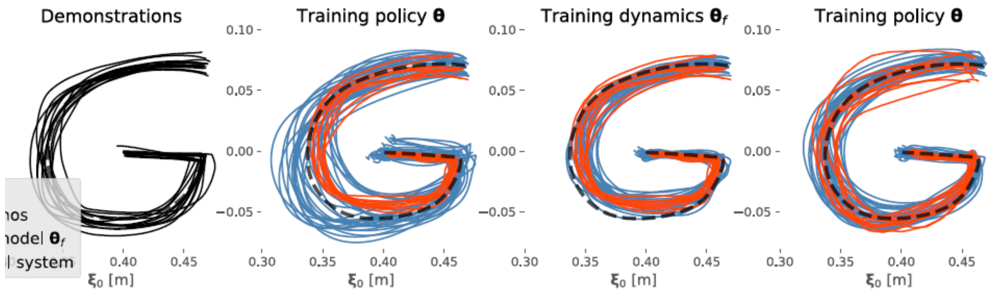

*J. Jankowski, H. Girgin, S. Calinon **Probabilistic Adaptive Control for Robust Behavior Imitation.**, IEEE Robotics and Automation Letters (RA-L), 6:2, pp. 1997-2004, 2021.

E. Pignat, H. Girgin, S. Calinon, **Generative Adversarial Training of Product of Policies for Robust and Adaptive Movement Primitives.. In Proc. Conference on Robot Learning (CoRL), 2020.

Learning in robotics – nullspace controller

- What to learn ?
- From pure machine learning perspective towards pure physics perspective

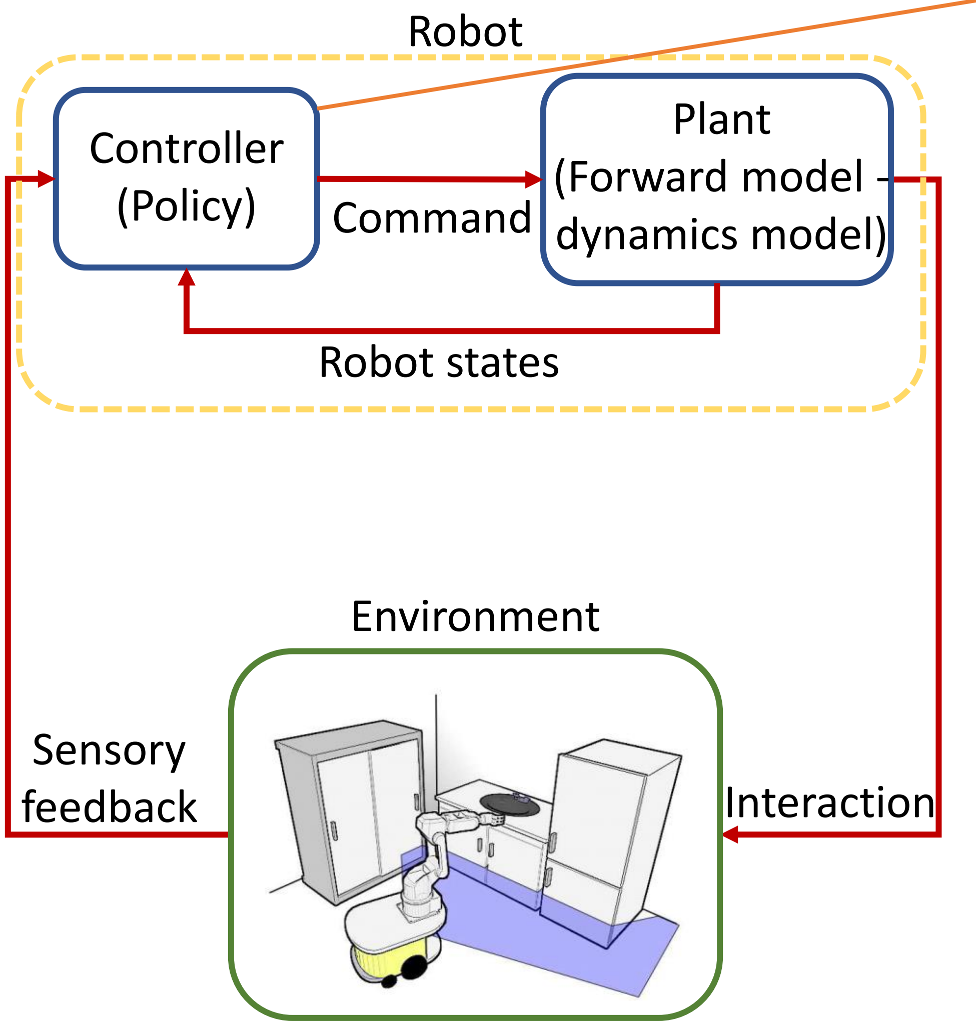

Combination of Controllers
(series, parallel or hierarchical)

- Series combination extensively researched (e.g., movement primitives), parallel combination (e.g., product of experts)
- Designing a feedback controller which can readily achieve hierarchical tasks not available
- Planning approaches combined with high gain(stiff) controllers may result in non-robustness and in bad generalization

Learning in robotics – nullspace controller*

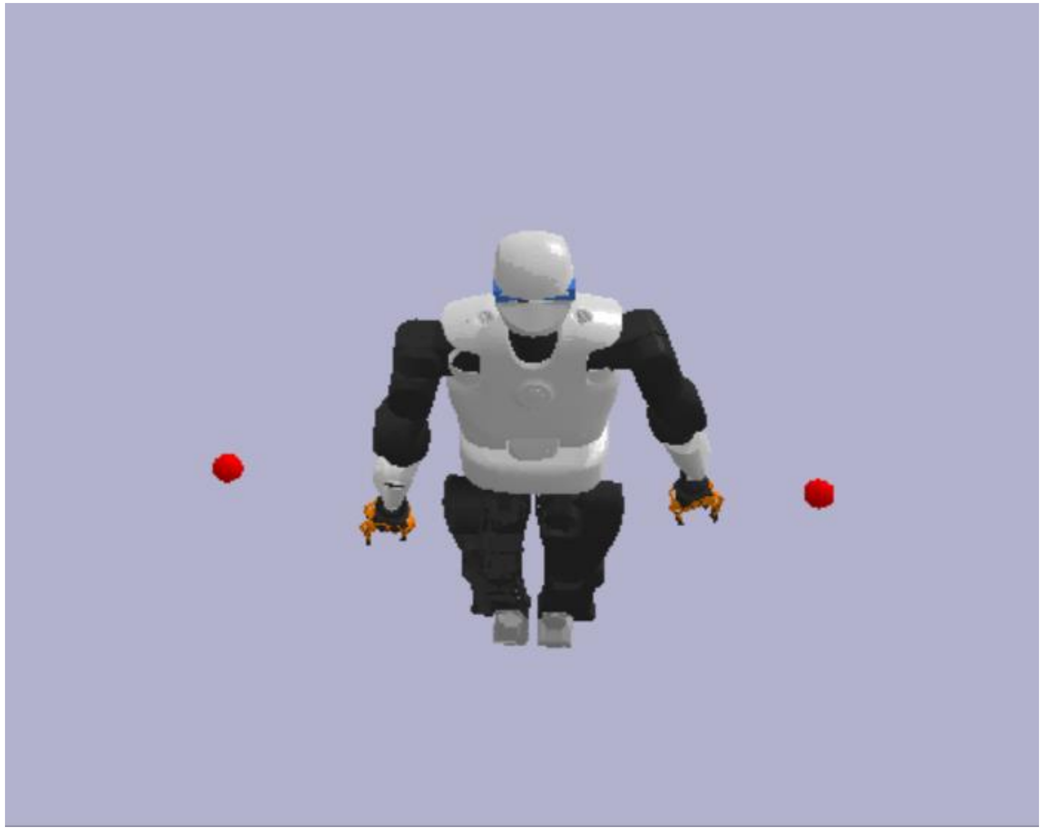

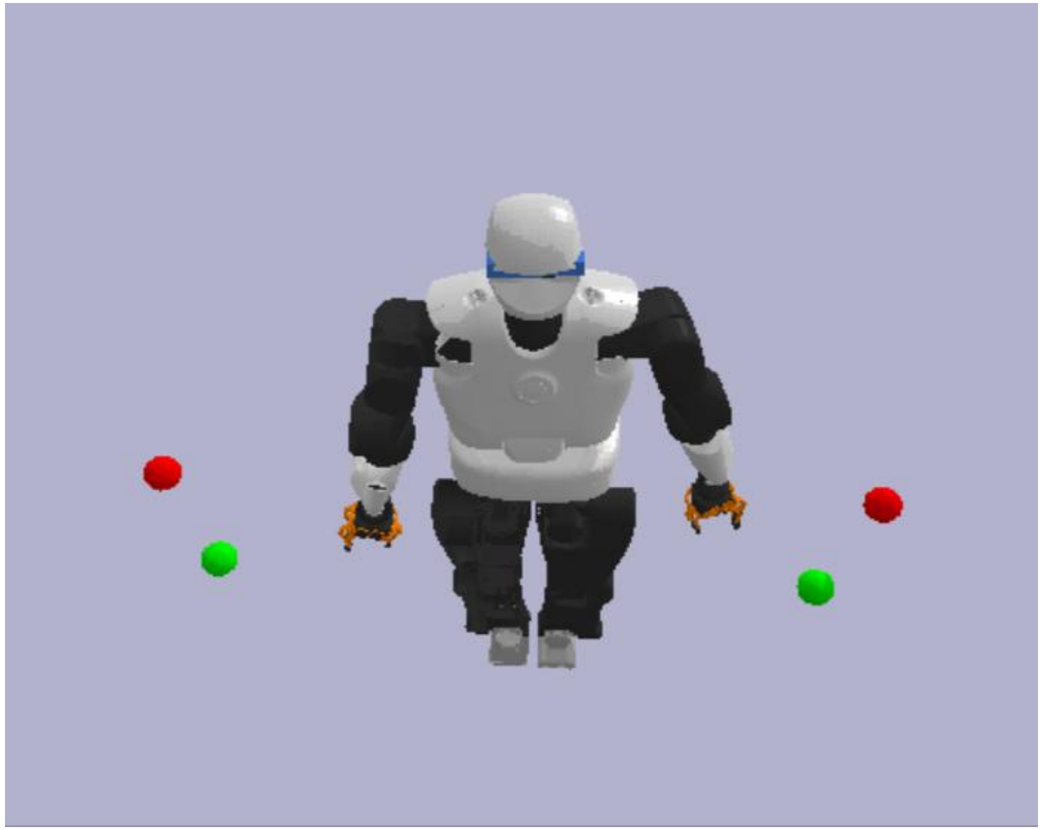

*Hakan Girgin and Sylvain Calinon, Nullspace controllers in system level synthesis with applications in robotics, in preparation.

Learning in robotics – nullspace controller*

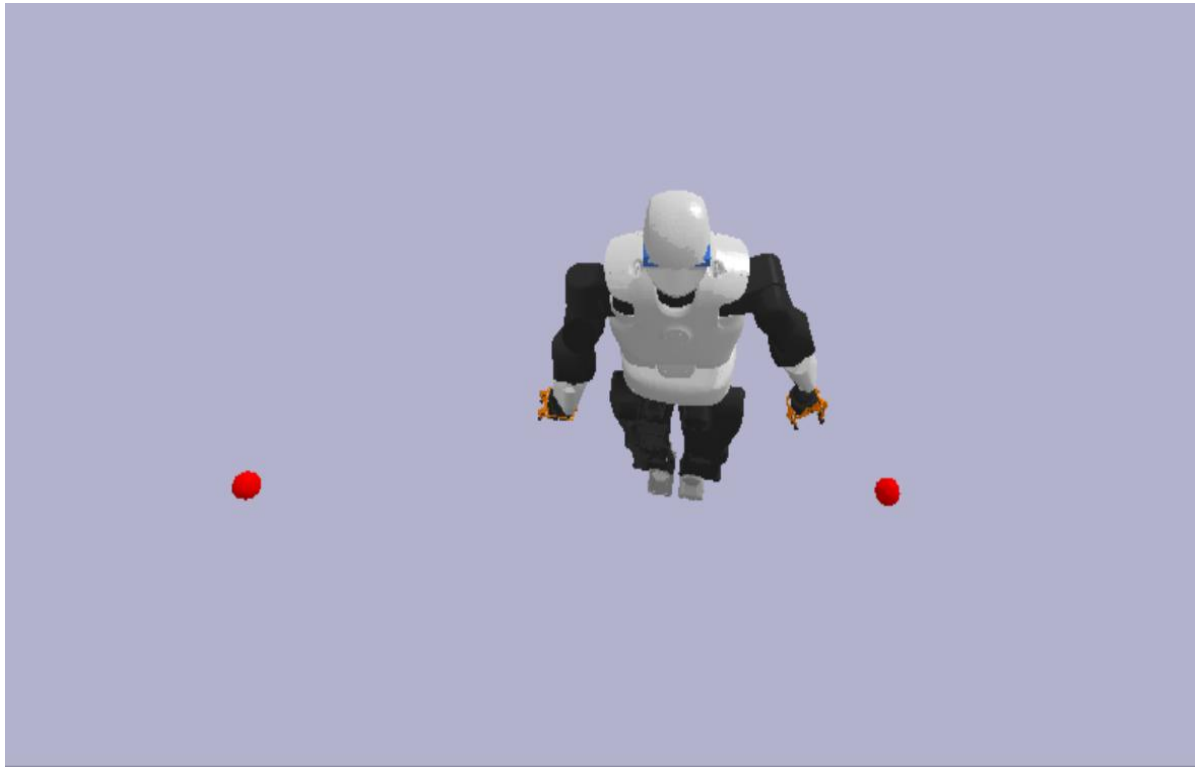

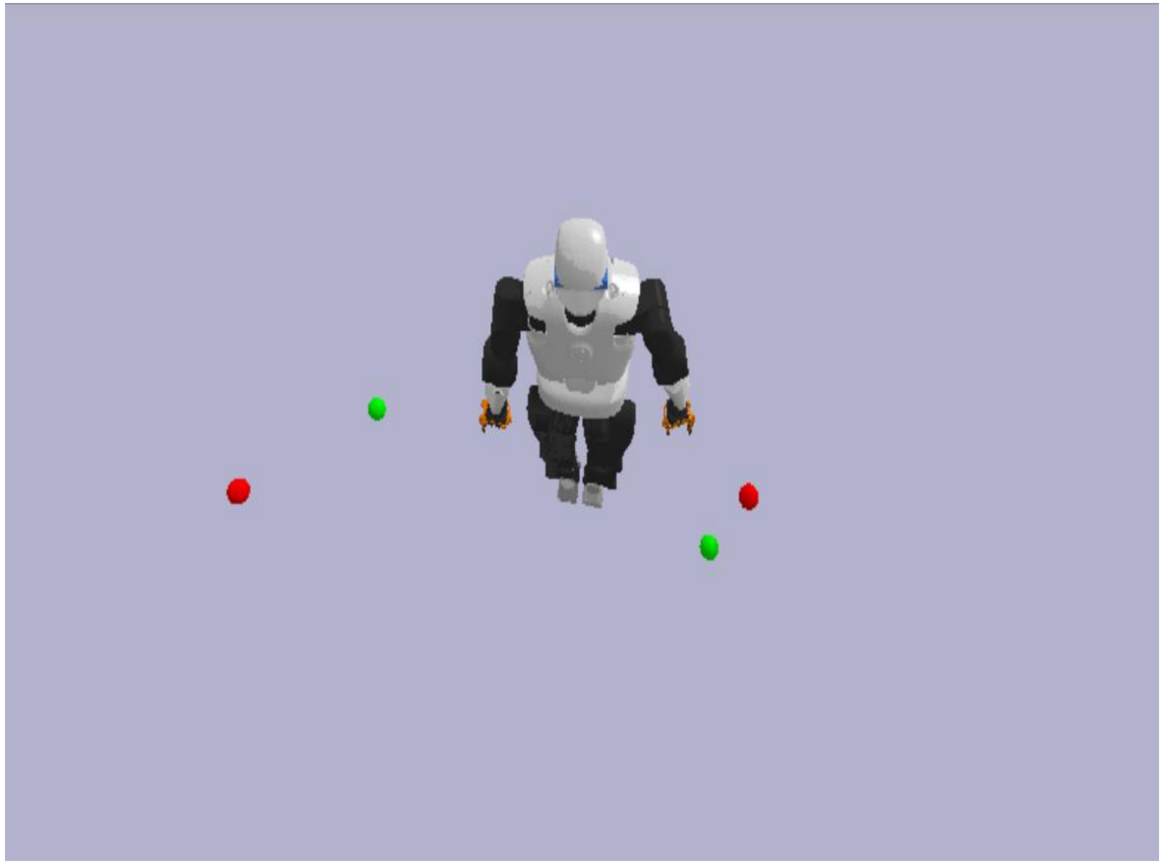

*Hakan Girgin and Sylvain Calinon, Nullspace controllers in system level synthesis with applications in robotics, in preparation.

Learning in robotics – learning forward models (indirect)

- What to learn ?
- From pure machine learning perspective towards pure physics perspective

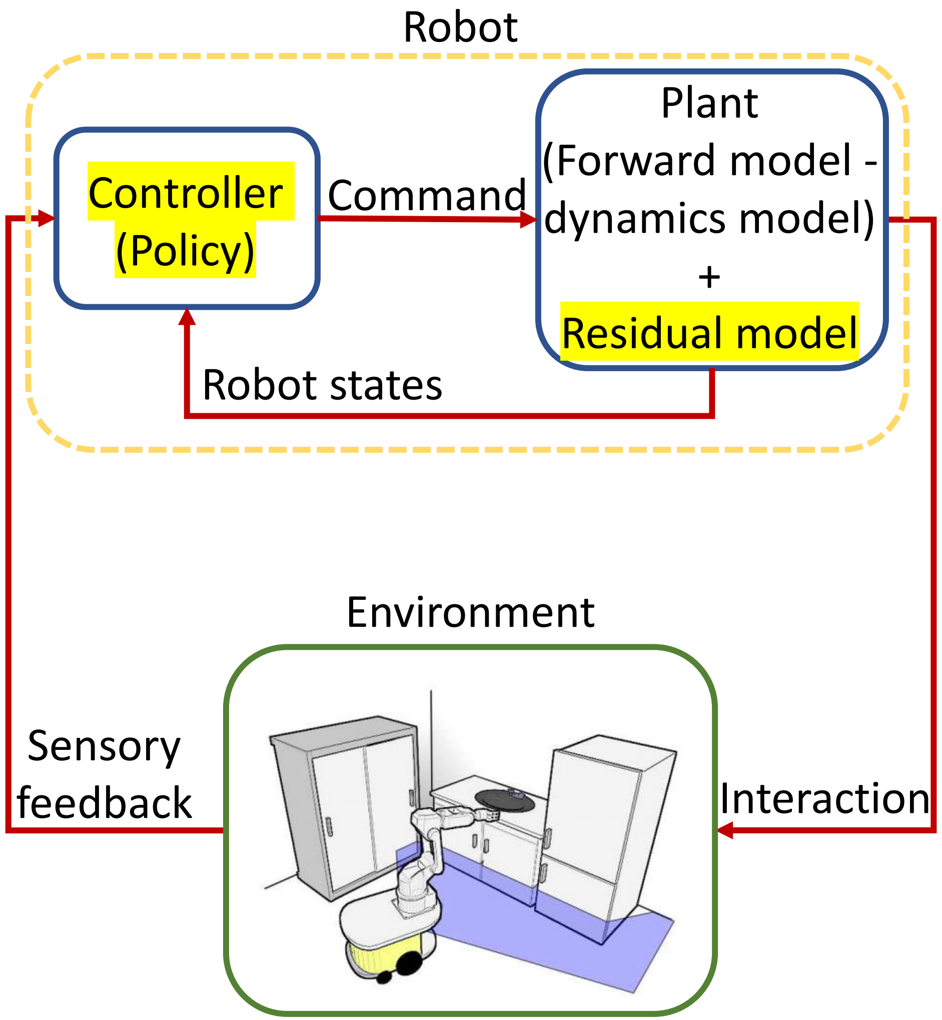

- Sometimes the robot dynamics mathematical model are either not available or imperfect.
- Complex forward models -> harder to get the optimal controller parameters.
- Swingup control of a simple pendulum with an optimal controller with learned linear plant model (Koopman operators) (**ongoing work**)

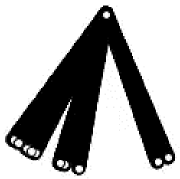

My conclusions

- End-to-end learning applications look much more fascinating than the pure simple model-based approaches, with the trade-off of requiring very large training data, time and parameter tuning.

- There are many components that the learning/designing can take place and these frameworks can be combined to create a more powerful, generalizable framework to create more intelligent robots.

- Robots are physical entities which rely on physics law. There is probably no need to reinvent the Newton’s physical laws from training data using machine learning.

- Model-based approaches combined with (deep) reinforcement learning could be a good way for a new phd student, but with the condition on tackling on an unsolved or partially solved robotic application.