Reflective Learning of Stochastic Physical Models of Objects for Manipulation

Shaojun Zhu and Abdeslam Boularias

Rutgers, The State University of New Jersey

Objectives

- **Robust** robotic manipulation of novel objects through the use of stochastic **friction and mass** models of objects
- Building upon **prior models** instead of starting from scratch with each new object
- **Correcting** the prior models **on the fly**
- Efficient use of **physics engines** for black-box model identification and correction
- Lifelong online learning

Examples

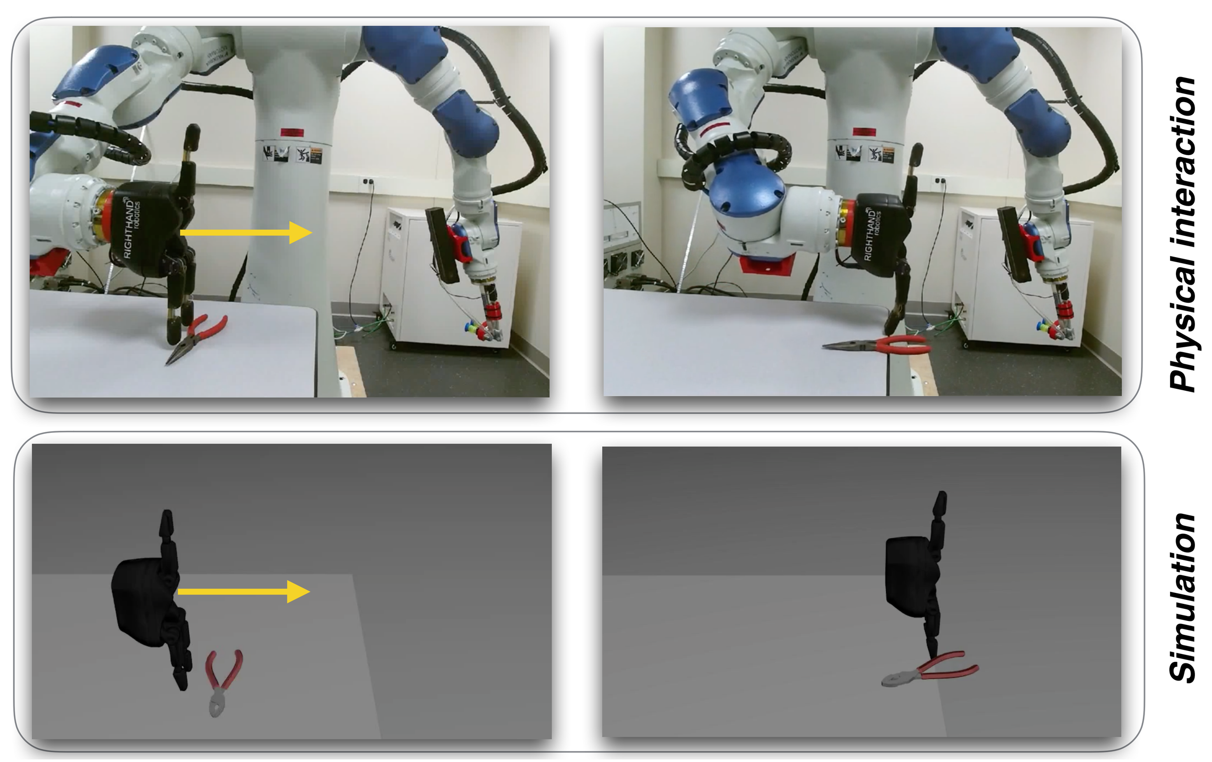

Figure 1: To grasp the pair of pliers from the tabletop, the robot needs to push the object to the table's edge and grab it easily from there. To avoid dropping the pliers, a good model of the mass and friction properties of the clippers need to be learned on the fly. The learned model is used in a physics engine to simulate the motion of the clippers.

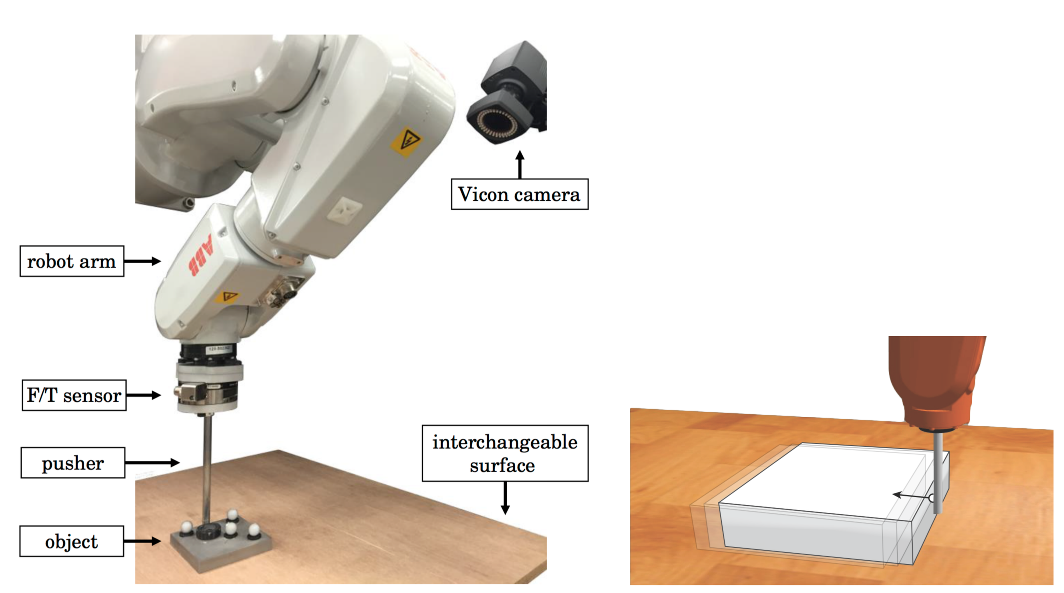

Figure 2: Planar Pushing Dataset [1] (MCUBE lab, MIT). The data consist of recorded poses of a planar object being pushed by a robot. The goal is to identify the unknown friction factors and mass of the object in order to predict its motion.

Notations

- θ is a d-dimensional vector corresponding to the unknown mass and static and kinetic friction coefficients of each subpart of a given object..
- P_t is a probability distribution of θ a time-step t.
- x_t is the observed 6D pose (position and orientation) of the manipulated object at time t.
- μ_t is a vector describing a force applied by the robot's fingertip on the object at time t.
- Applying a force μ_t results in changing the object's pose from x_t to x_{t+1}.
- f is the transition function of a physics engine, such that $f(x_t, \mu_t, \theta) = \hat{x}_{t+1}$.

Problem

- Given P_0, a prior of model θ before starting to interact with the object, and a sequence of actions and observed poses $(x_0, \mu_0, x_1, \mu_1, \dots, x_{t-1}, \mu_{t-1}, x_t)$,
- Calculate P_t, the probability distribution of θ.

Method

- Empirical error on observed data:

$$E(\theta) \stackrel{def}{=} \sum_t \|x_{t+1} - f(x_t, \mu_t, \theta)\|_2$$

- Best model that explains observed data:

$$\theta^* = \arg\min_\theta E(\theta)$$

- We do not know the analytical form of error function E because $E(\theta)$ is obtained from simulation with a physics engine.
- Use black-box Bayesian optimization [2, 3] to find P_t, the probability distribution of θ^*.
- Following the *entropy search technique* [4] ,

$$P_t(\theta) \stackrel{def}{=} P\big(\theta = \arg\min_{\theta^i \in \Theta} E(\theta^i)\big) = \int_{E:\mathbb{R}^d \to \mathbb{R}} p(E) \Pi_{\theta^i \in \Theta - \{\theta\}} H\big(E(\theta^i) - E(\theta)\big) \mathrm{d}E,$$

where H is the Heaviside step, i.e. $H\big(E(\theta^i) - E(\theta)\big) = 1$ if $E(\theta^i) \geq E(\theta)$ and $H\big(E(\theta^i) - E(\theta)\big) = 0$ else, and $p(E)$ is the probability of error function E.
- $p(E)$ is a **Gaussian Process**, and $P_t(\theta)$ is evaluated by **Monte Carlo** samples from $p(E)$.
- $p(E)$ is computed by evaluating E from several simulations with different hypothesized models θ

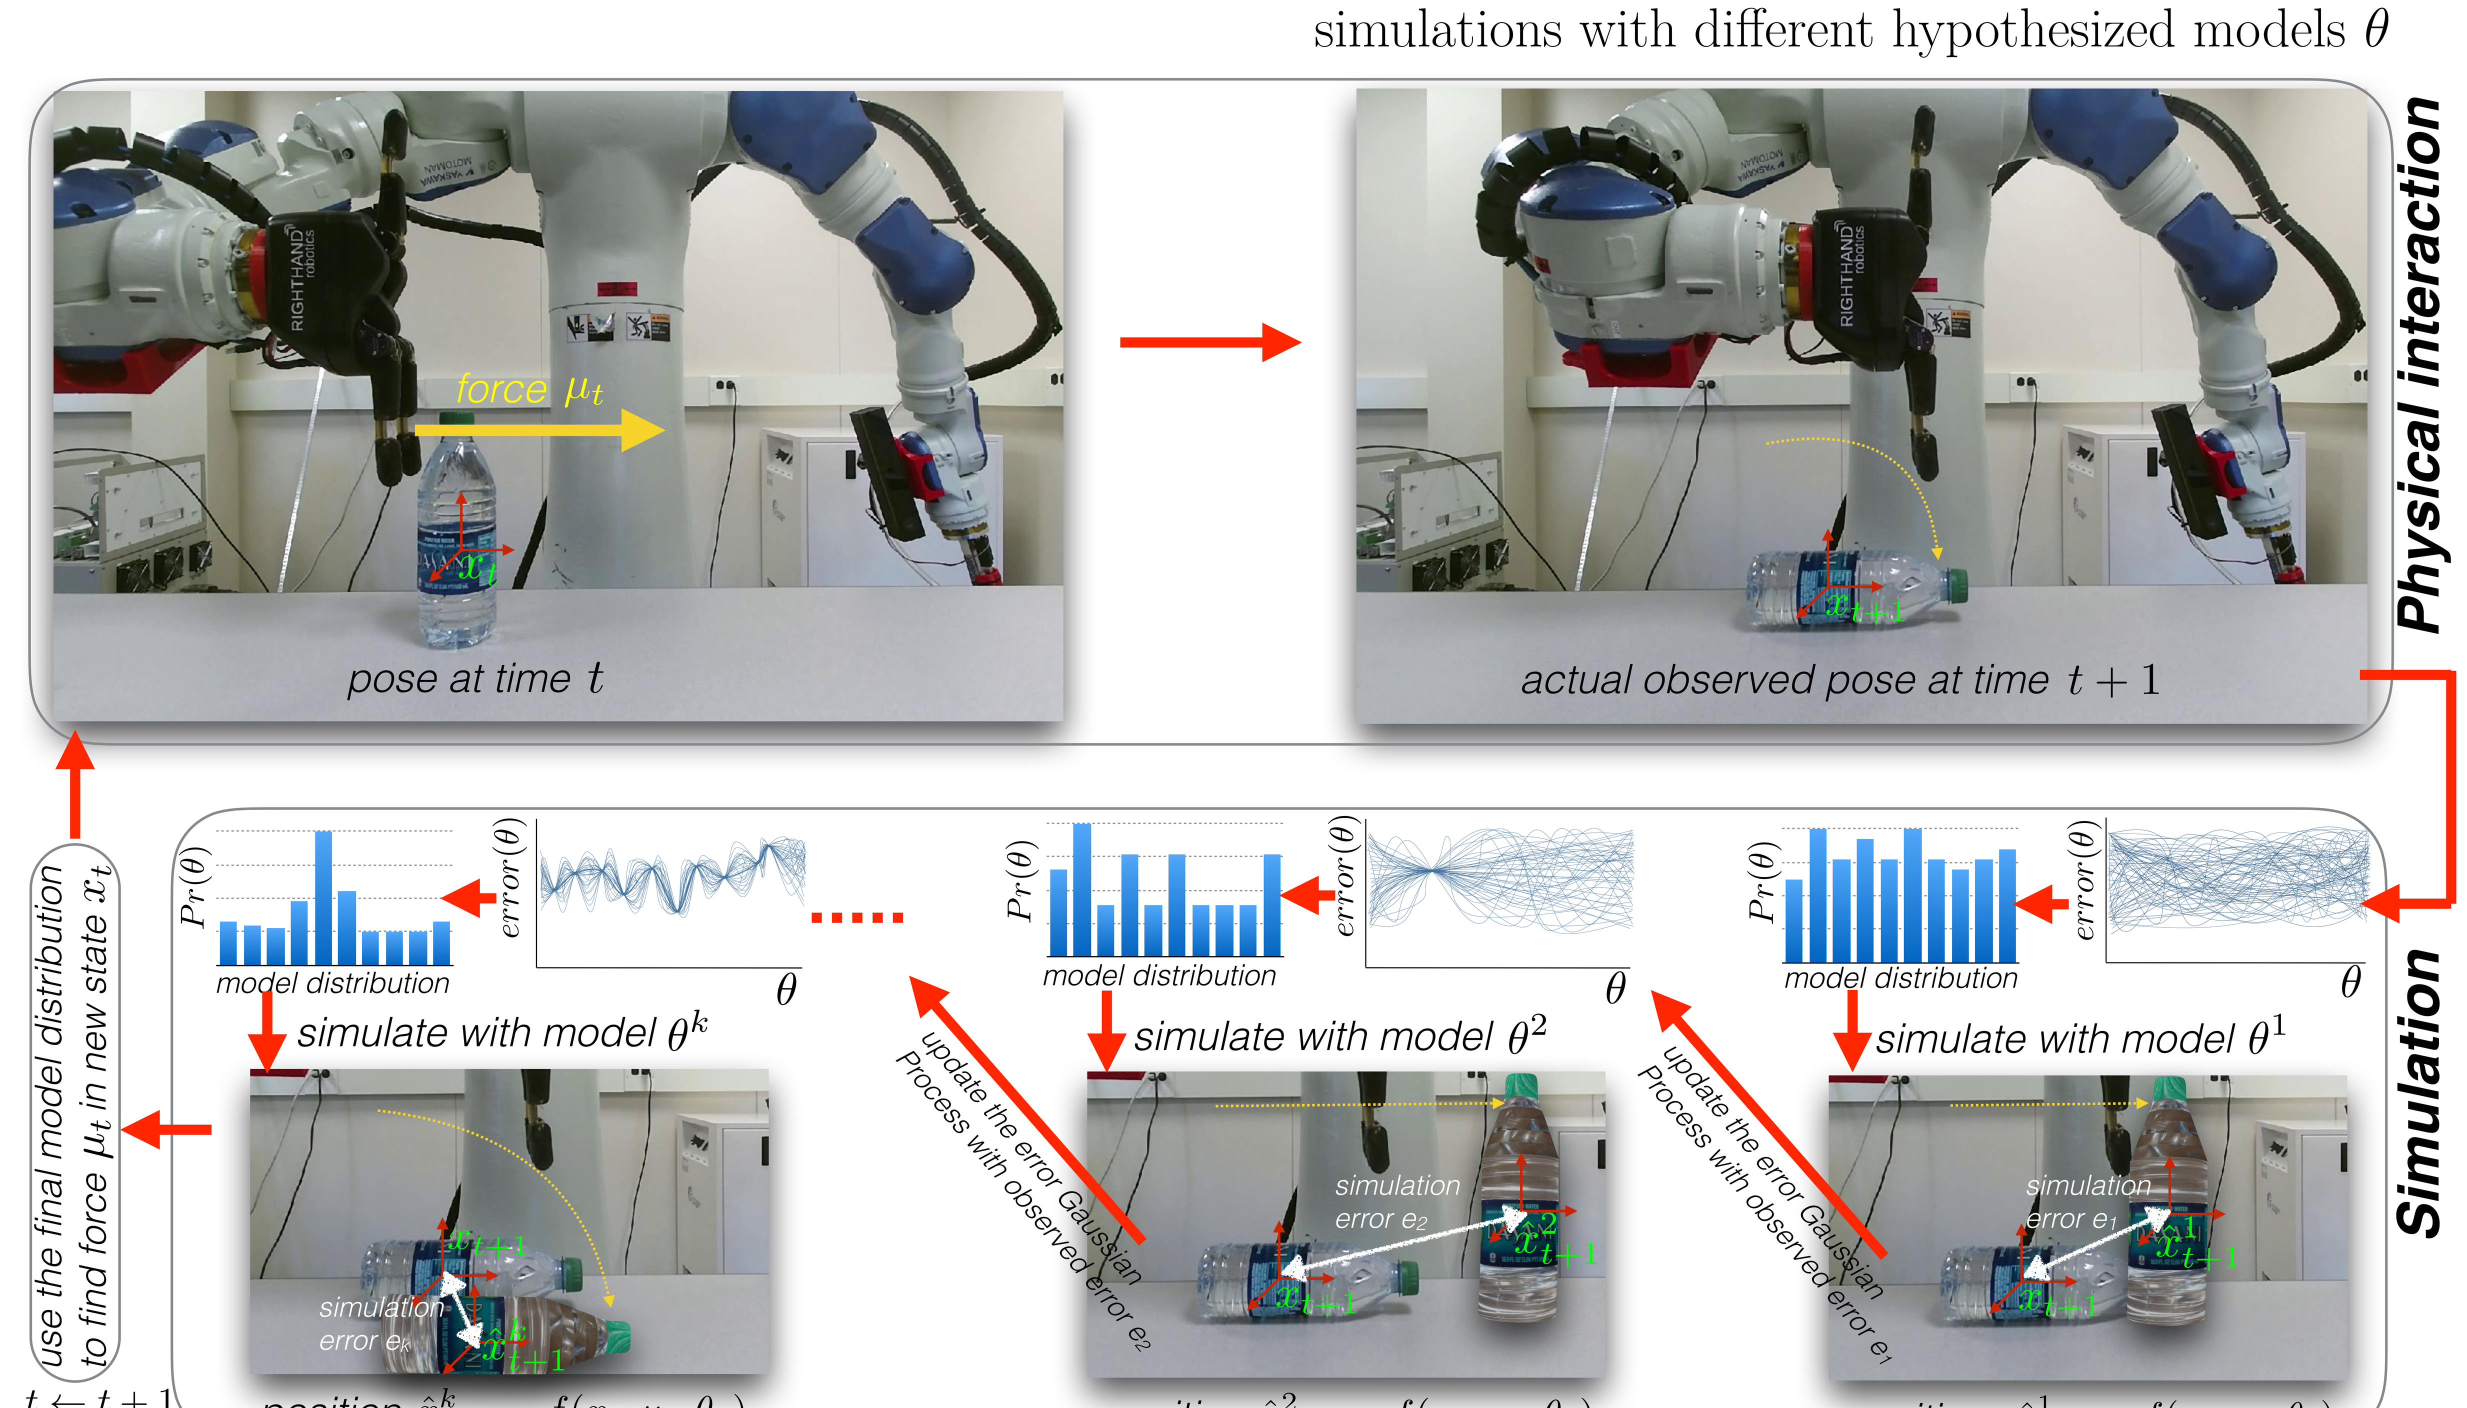

Figure 3: Overview of the proposed approach for learning object models with a physics engine

Greedy Entropy Search

- Physics engine calls, needed to compute the error distribution $p(E)$, are computationally expensive.
- To minimize the number of calls to the physics engine, we choose the model θ that has the highest contribution to the current entropy of P_t, i.e. the one with the highest

$$-P_t(\theta) \log\big(P_t(\theta)\big),$$

as the next model to evaluate in simulation.

Preliminary Results

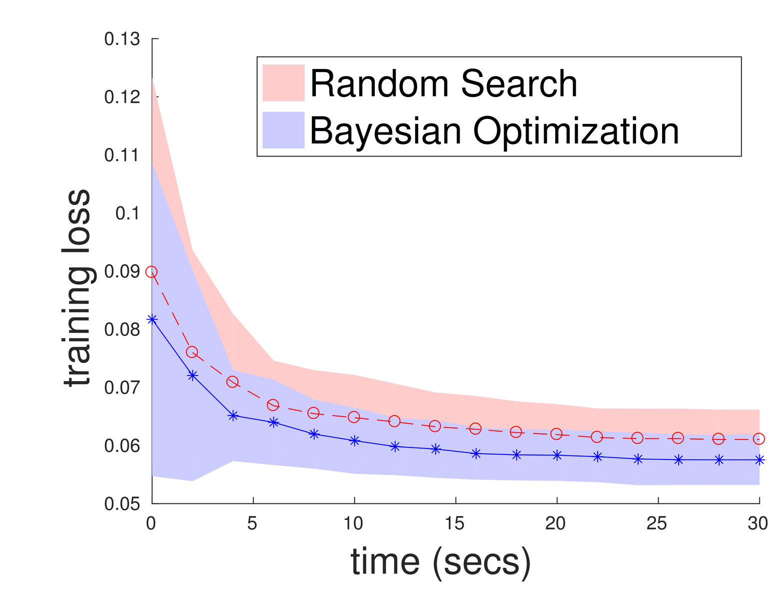

Figure 4: Error in predicting poses of pushed planar objects as a function of simulation time. Bayesian optimization refers to the greedy entropy search approach.

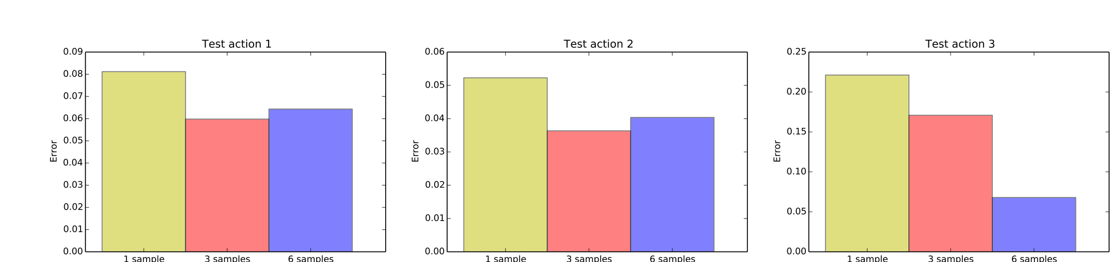

Figure 5: Pose prediction error as a function of the number of training samples with three different objects.

References

[1] Kuan-Ting Yu, Maria Bauzá, Nima Fazeli, and Alberto Rodriguez. More than a million ways to be pushed: A high-fidelity experimental data set of planar pushing. *CoRR*, abs/1604.04038, 2016.

[2] Julien Villemonteix, Emmanuel Vazquez, and Eric Walter. An Informational Approach to the Global Optimization of Expensive-to-evaluate Functions. *Journal of Global Optimization*, 44(4):509–534, 2009.

[3] Daniel J. Lizotte. *Practical Bayesian Optimization.* PhD thesis, University of Alberta, Edmonton, Canada, 2008.

[4] Philipp Hennig and Christian J. Schuler. Entropy Search for Information-Efficient Global Optimization. *Journal of Machine Learning Research*, 13:1809–1837, 2012.

[5] Karl Pauwels and Danica Kragic. Simtrack: A simulation-based framework for scalable real-time object pose detection and tracking. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2015.