A Differentiable Recipe for Learning Visual Non-Prehensile Planar Manipulation

Bernardo Aceituno, Alberto Rodriguez, Shubham Tulsiani, Abhinav Gupta, Mustafa Mukadam

contact: aceituno@mit.edu

Problem Statement

Can we infer how to manipulate an object given a video description of the desired motion?

Specifying tasks with videos is a powerful technique towards acquiring novel and general robot skills. However, reasoning over mechanics and dexterous interactions can make it challenging to scale learning contact-rich manipulation.

In this work, we focus on the problem of visual non-prehensile planar manipulation: given a video of an object in planar motion, find contact-aware robot actions that reproduce the same object motion

Fundamental Challenges

1 Mechanical Parameters

The robot finger-trajectories are dictated by the mechanics that describe the task: object facets, friction cones, and external contact location.

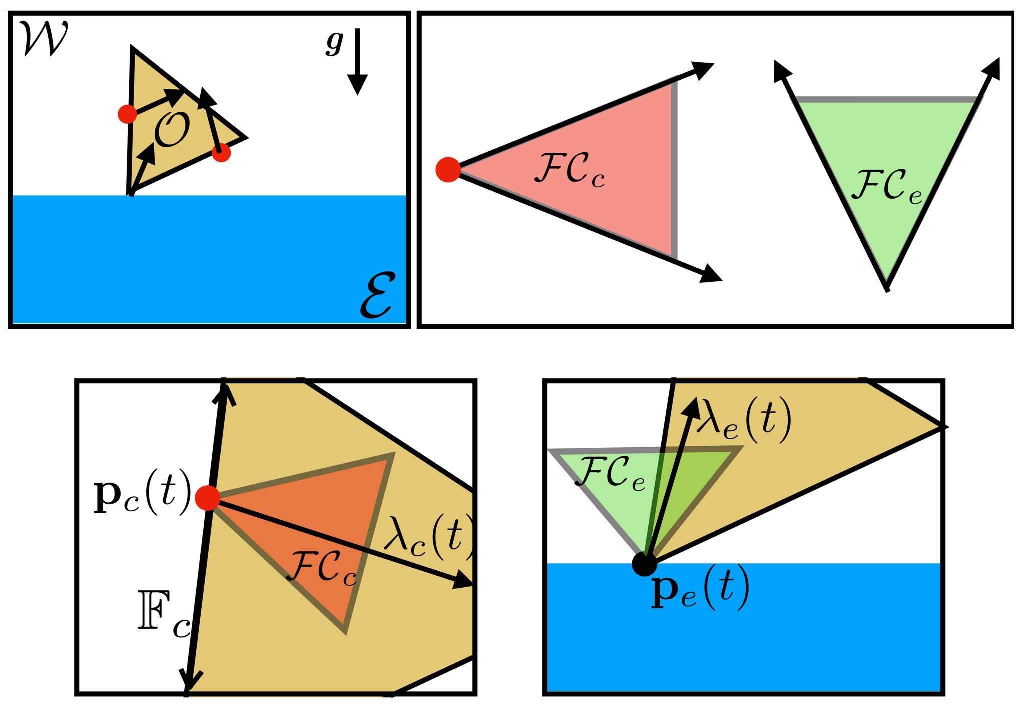

2 Non-uniqueness of the solution

There is no direct mapping from an object motion to a robot finger-trajectory, since different robot interactions can lead to a single object motion.

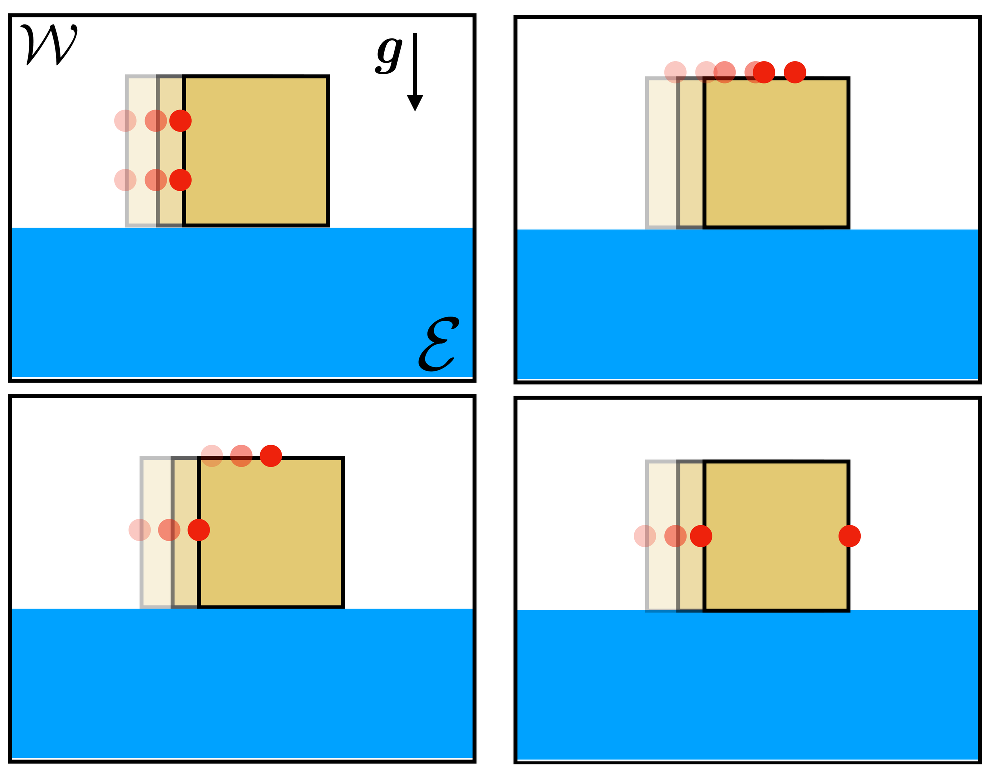

Proposed Architecture

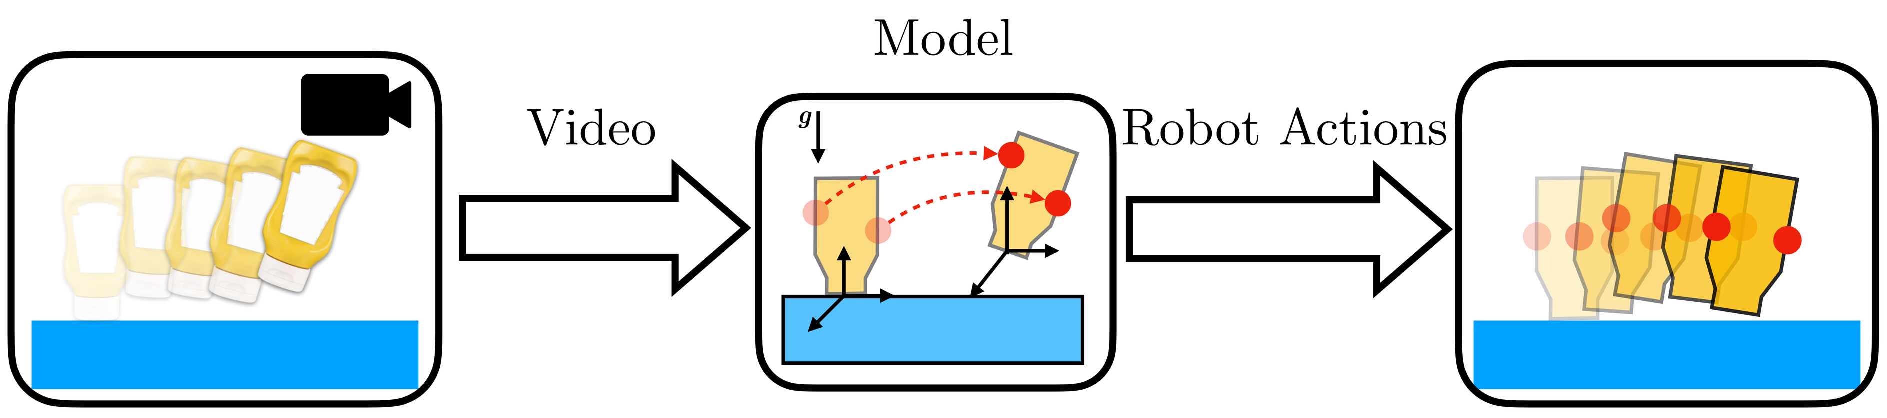

Problem overview

To address the fundamental challenges of learning a policy in this setup, we propose the following differentiable architecture:

1. We encode the video of the desired object motion in a latent vector $L^{\mathcal{V}}$ and decode it to infer a set of mechanical parameters $\mathcal{P}$ that describe the task, trained under a cost function $\mathcal{L}_P$.
2. The parameters $\mathcal{P}$ are used to solve a linearized inverse dynamics problem, posed as differentiable QP. We train the network such that QP outputs robot finger-trajectories p, λ, under a cost function $\mathcal{L}_{QP}$.
3. The robot finger-trajectory is fed to a differentiable simulator of the scene. We train the full model such that the simulation output matches the input video, under a cost function $\mathcal{L}_{sim}$.

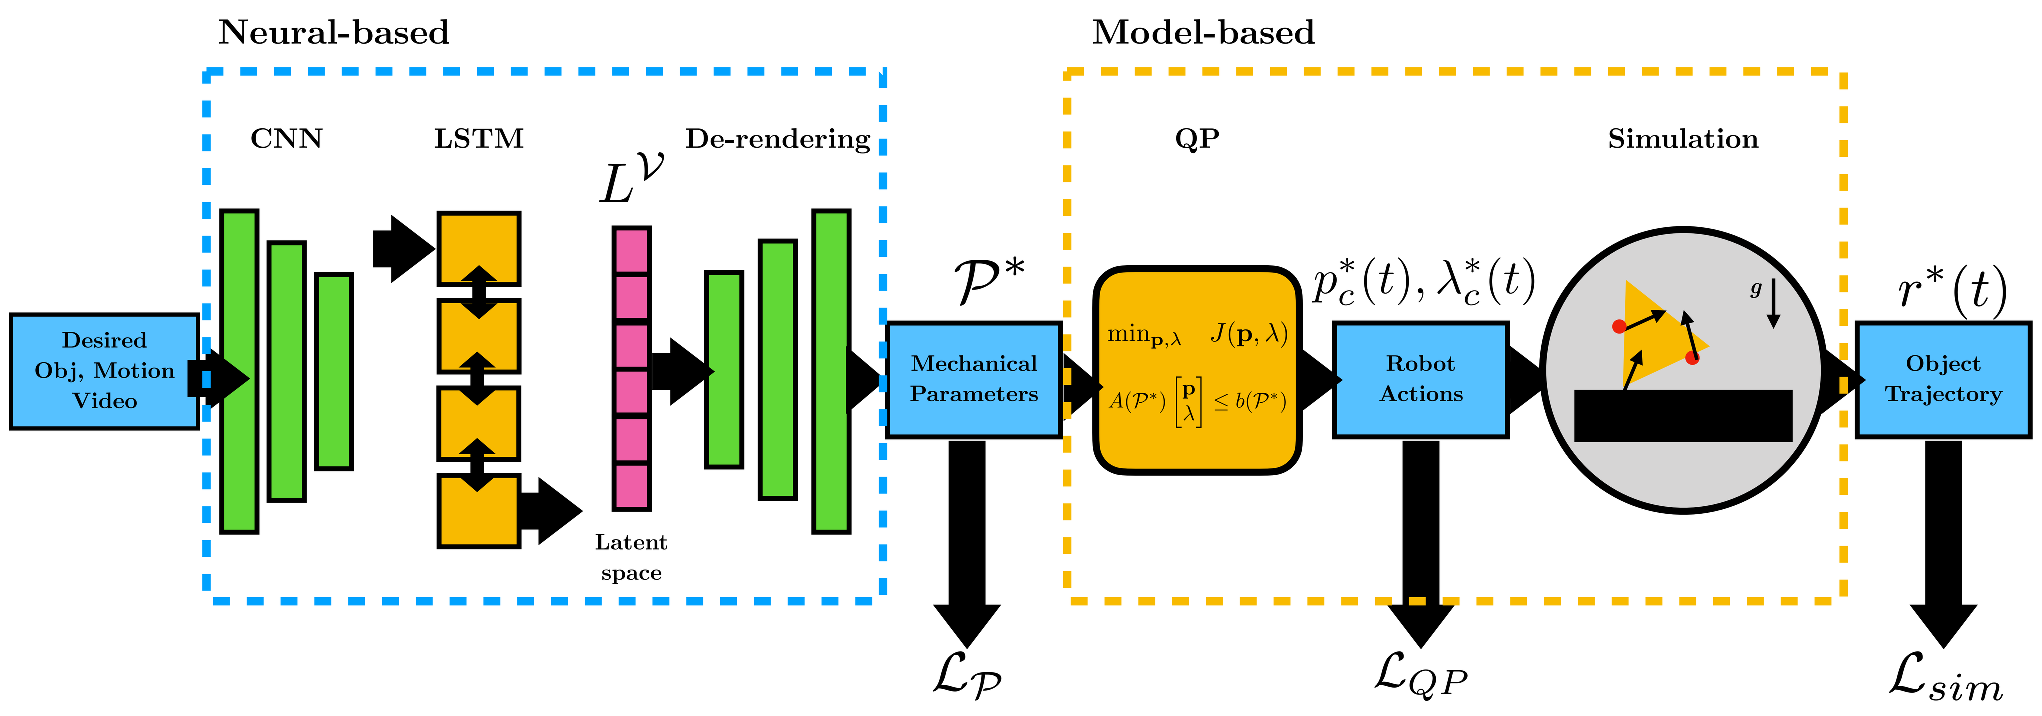

Proposed architecture: Differentiable Learning for Manipulation (DLM)

Quantitative Comparison

We demonstrate the capabilities of this model on four different baselines of this problem with different degrees of intermediate training and differentiable structure. We test the architectures:

1. Fully-connected networks trained over finger-trajectories (NN) or simulator output (NNM).
2. Our architecture trained solely over the mechanical parameters (MDR).
3. Our architecture trained over the mechanical parameters and the QP output (CVX).
4. Our full architecture trained over mechanical parameters, QP, and simulator output (DLM).

We train and validate these models over videos with randomized shapes illustrated as:

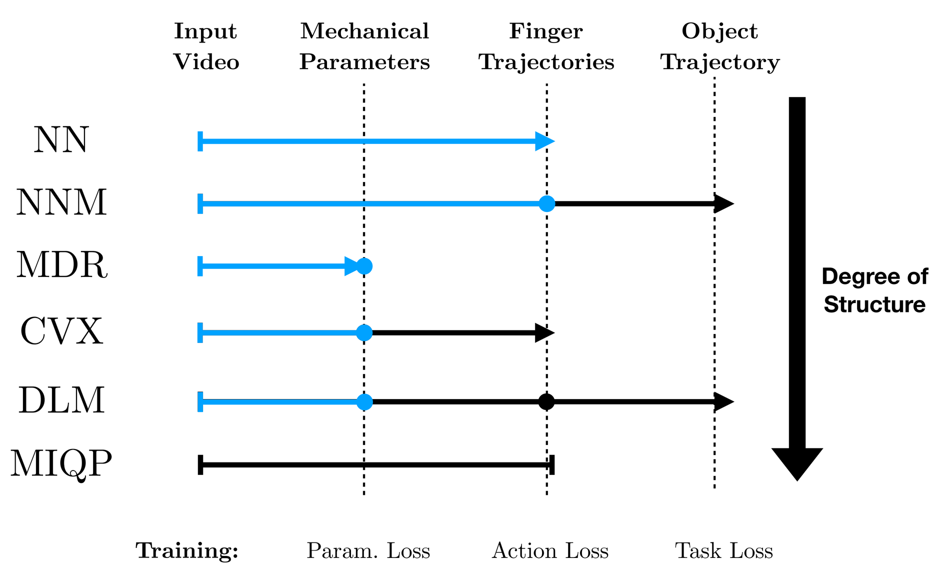

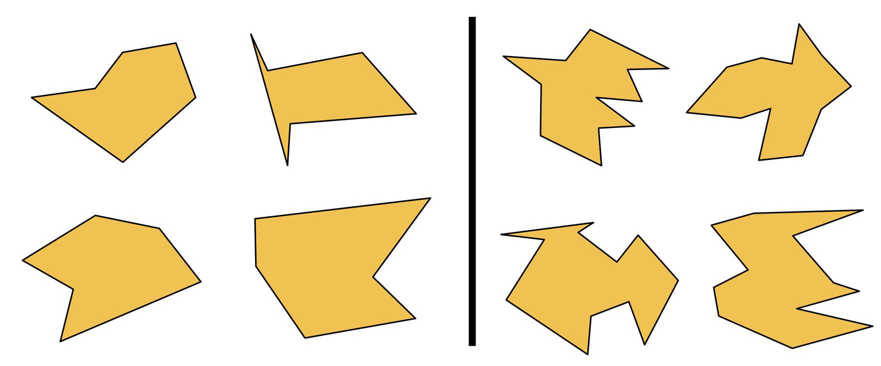

Training, fine-tuning and family-test sets **Test set**

We report our qualitative results with videos from randomized shapes below:

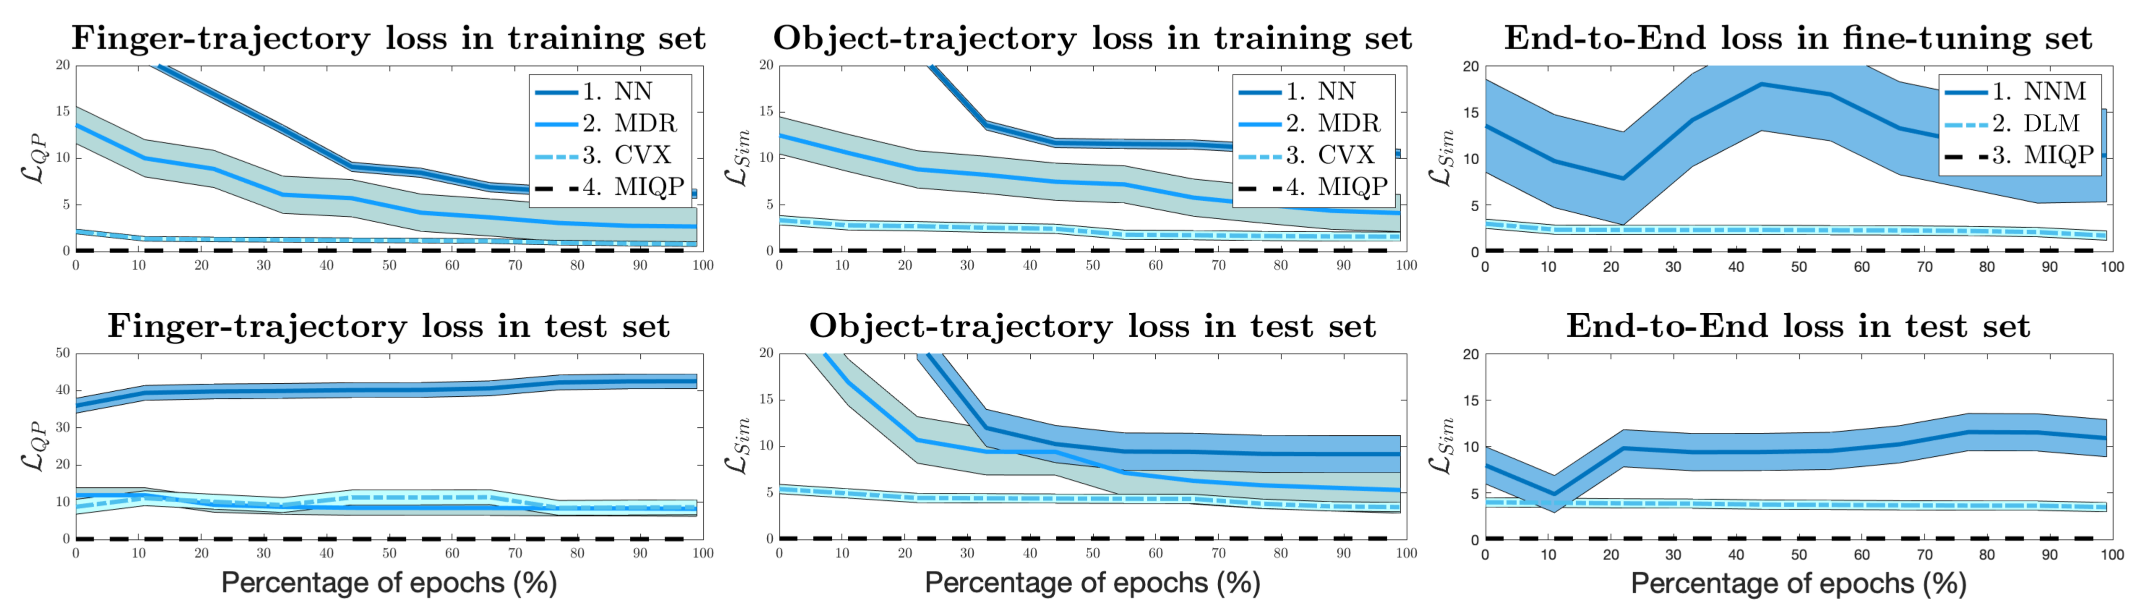

References

[1] de Avila Belbute-Peres, Filipe, et al. "End-to-end differentiable physics for learning and control." NeurIPS (2018).

[2] Agrawal, Akshay, et al. "Differentiable Convex Optimization Layers." NeurIPS (2019).

[3] Cheng, Xianyi, et al. "Contact mode guided sampling-based planning for quasistatic dexterous manipulation in 2d. ICRA (2021).

[4] Aceituno, Bernardo, and Alberto Rodriguez. "A global quasi-dynamic model for contact-trajectory optimization." RSS (2020).