Apprenticeship Learning for Robotic Control

(FA9550-XX-X-XXXX, your grant number)

PI: Pieter Abbeel (UC Berkeley)

AFOSR Program Review:
Mathematical and Computational Cognition Program
Computational and Machine Intelligence Program
Robust Decision Making in Human-System Interface Program
(Jan 28 – Feb 1, 2013, Washington, DC)

Apprenticeship Learning for Robotic Control (Pieter Abbeel)

Research Objectives:

- Make robots easier to program for new tasks through enabling them to learn as apprentices
- Primary domain: robotic manipulation

Technical Approach:

- A geometry-oriented approach to representing manipulation primitives and how they depend on context

DoD Benefits:

- Significant advances in robotic manipulation, and more generally robotic control
- Example applications include robotic surgery on the battlefield, service robots, and robotic interventions in disaster scenarios

Budget ($k):

YR 1	YR 2	YR 3	YR 4
120k	120k	120k	n/a

Project Start Date: 7/1/2012

Project End Date: 6/30/2015

List of Project Goals

1. Develop a representation of robotic control primitives that allow for generalization to new situations

2. Develop the capability for autonomous learning to improve performance

3. Enable learning from lower quality demonstrations

4. Connect language (verbs) to the primitives

Progress Towards Goals (or New Goals)

Our focus thus far:

Goal 1: Develop a representation of robotic control primitives that allow for generalization to new situations

We have developed the foundations of such a representation. It has enabled quickly teaching a robot new manipulation primitives that it can generalize to new situations. Results include enabling robots to tie knots.

Your Presentation Slides Go Here

(No more than 25 slides, not including the five mandatory slides)

Overview

- Motivation:
  - Apprenticeship learning for autonomous helicopter flight

- Apprenticeship learning for robotic manipulation

Challenges in Helicopter Control

- Unstable
- Nonlinear
- Complicated dynamics
  - Air flow
  - Coupling
  - Blade dynamics
- Noisy estimates of position, orientation, velocity, angular rate (and perhaps blade and engine speed)

Success Stories: Hover and Forward Flight

- Just a few examples:
  - Bagnell & Schneider, 2001;
  - LaCivita, Papageorgiou, Messner & Kanade, 2002;
  - Ng, Kim, Jordan & Sastry 2004a (2001); Ng et al., 2004b;
  - Roberts, Corke & Buskey, 2003;
  - Saripalli, Montgomery & Sukhatme, 2003;
  - Shim, Chung, Kim & Sastry, 2003;
  - Doherty et al., 2004;
  - Gavrilets, Martinos, Mettler and Feron, 2002.
- Varying control techniques: inner/outer loop PID with hand or automatic tuning, H1, LQR, …

One of our first attempts at autonomous flips [using similar methods to what worked for ihover]

Target trajectory: meticulously hand-engineered
Model: from (commonly used) frequency sweeps data

Target Trajectory

- Difficult to specify by hand:
  - Required format: position + orientation over time
  - Needs to satisfy helicopter dynamics

- Our solution:
  - Collect demonstrations of desired maneuvers
  - Challenge: extract a clean target trajectory from many suboptimal/noisy demonstrations

Abbeel, Coates, Ng, IJRR 2010

Expert Demonstrations

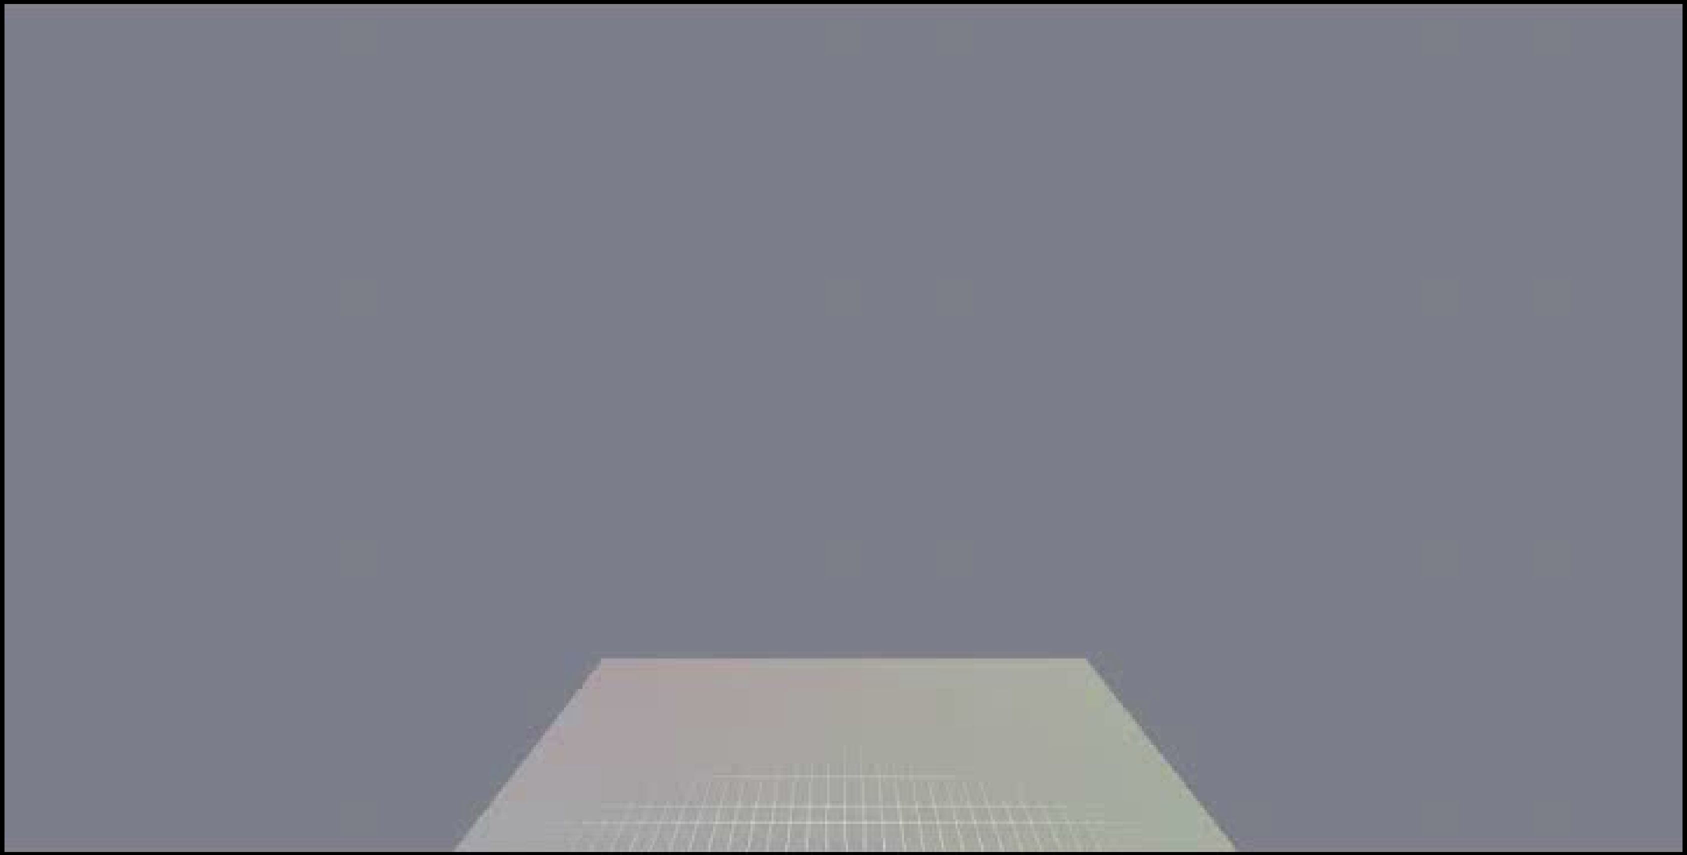

Results: Time-Aligned Demonstrations

- White helicopter is inferred “intended” trajectory.

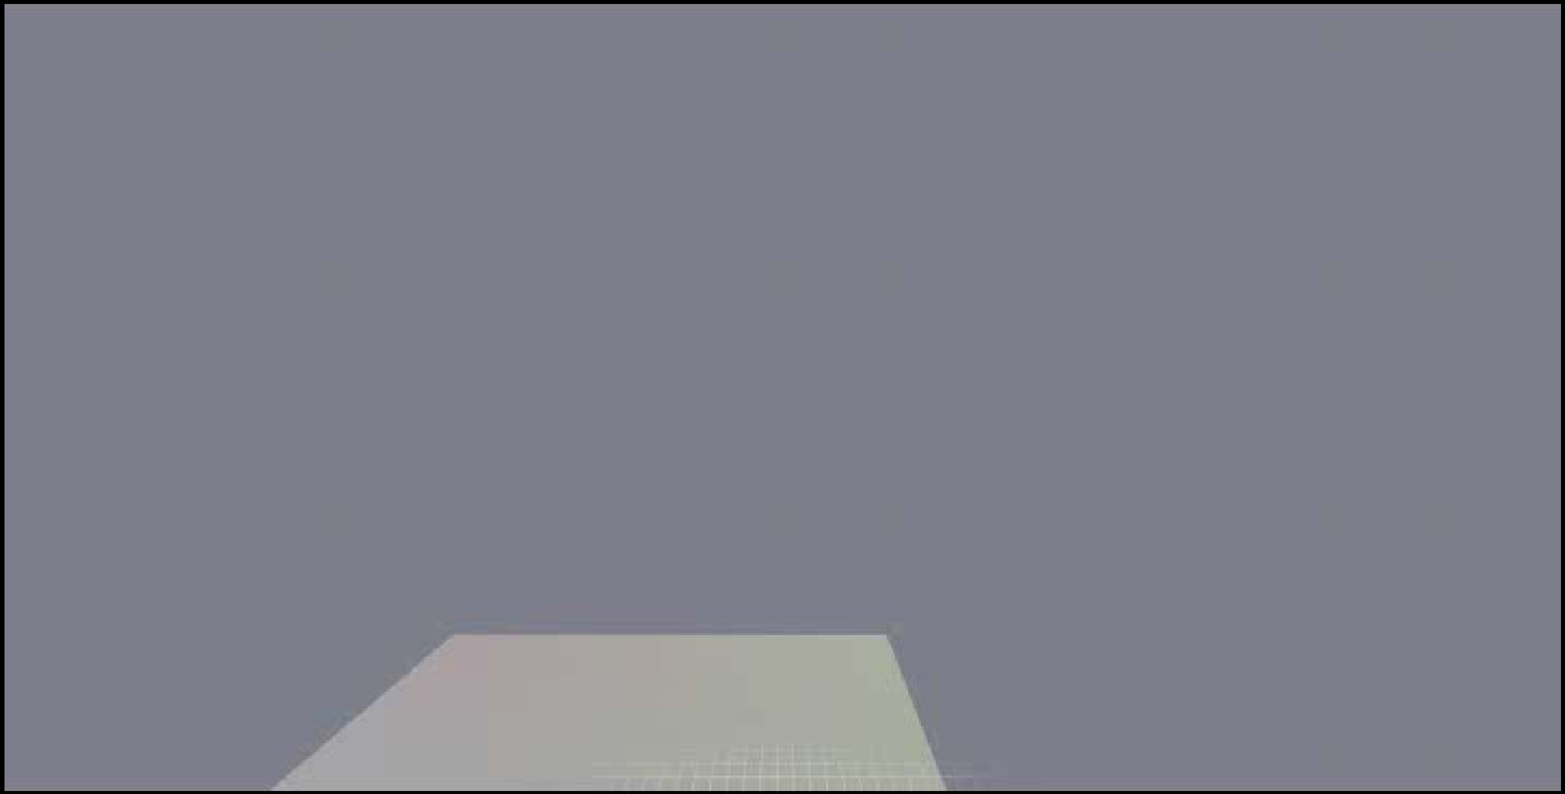

Experimental Setup

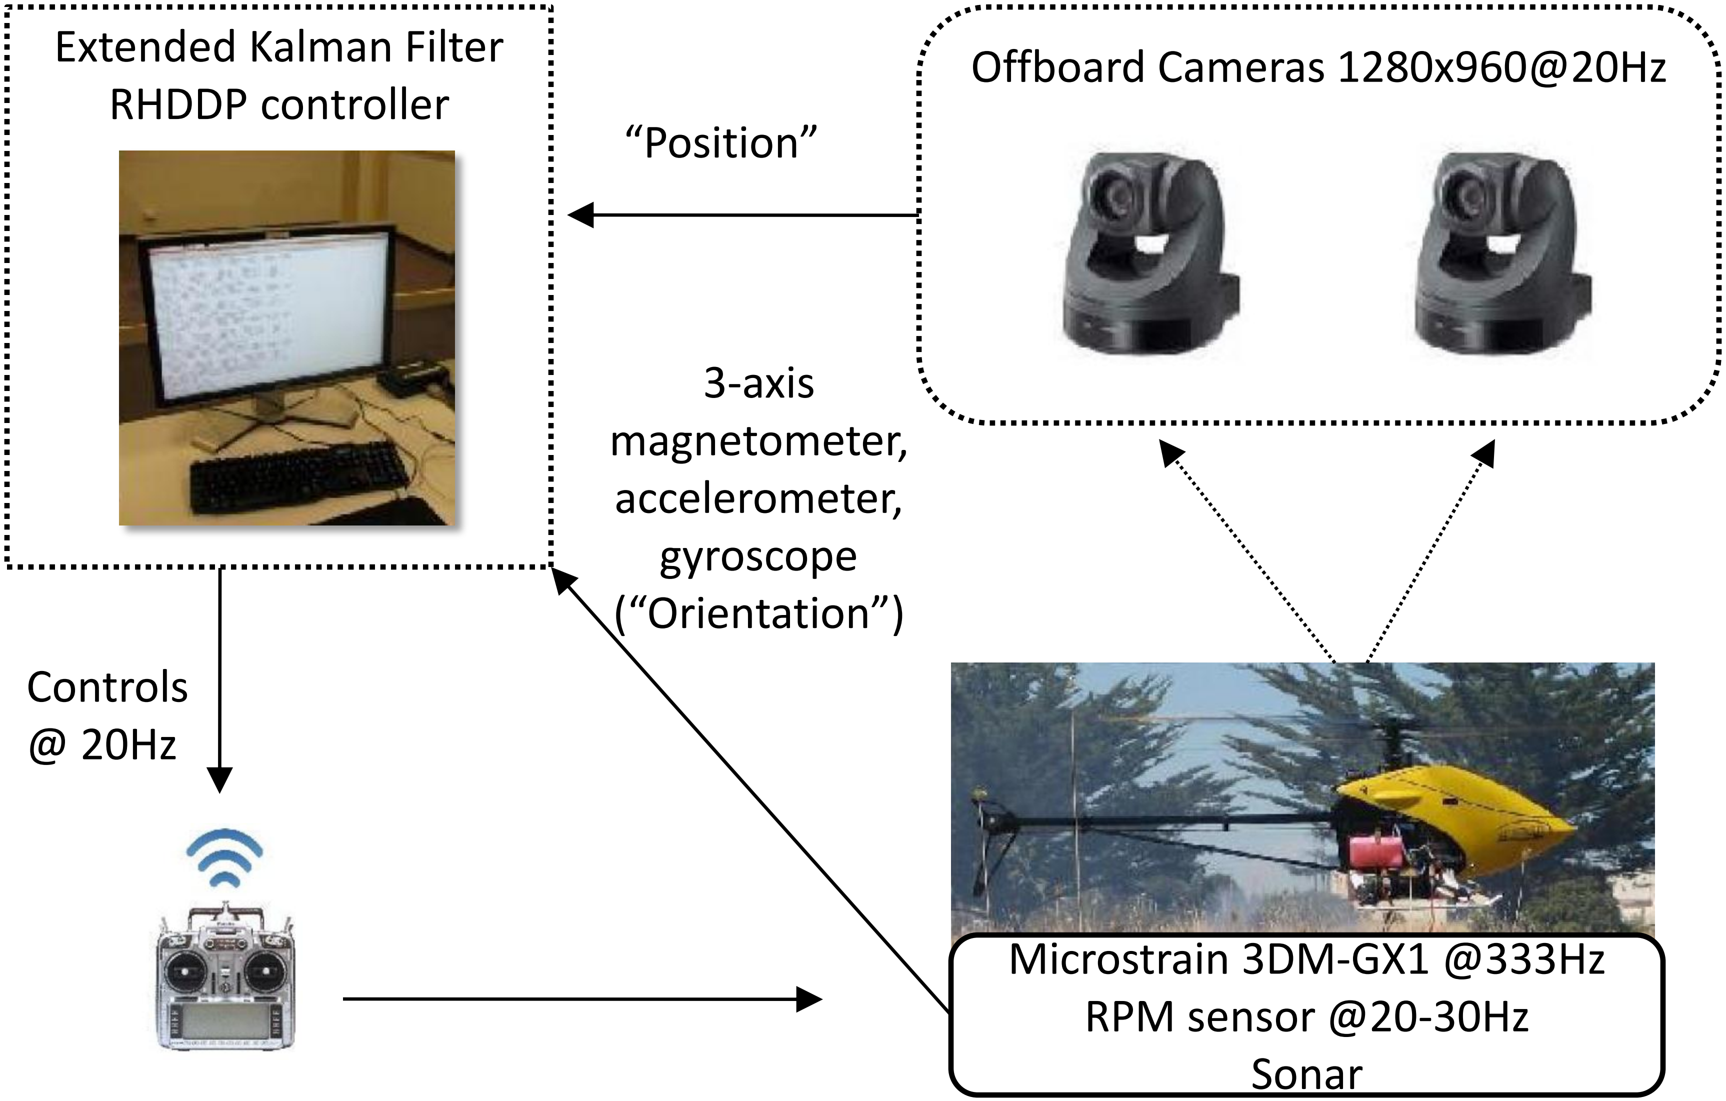

Abbeel, Coates, Quigley, Ng, NIPS 2007

Experimental Procedure

1. Collect sweeps to build a baseline dynamics model
2. Our expert pilot demonstrates the airshow several times.

3. Learn a target trajectory.
4. Learn a dynamics model.
5. Find the optimal control policy for learned target and dynamics model.
6. Autonomously fly the airshow

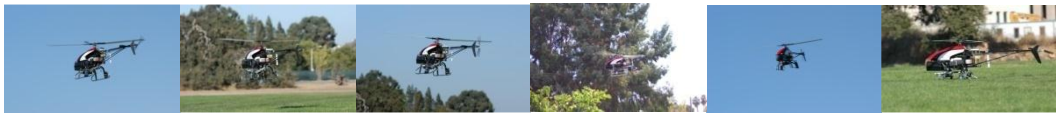

7. Learn an improved dynamics model. Go back to step 4.

→ **Learn to fly new maneuvers in < 1hour.**

Abbeel, Coates, Ng, IJRR 2010

Results: Autonomous Airshow

Thus Far

- Apprenticeship learning
  - Learn to perform task from expert demonstrations
  - Enabled by far most advanced helicopter aerobatics
- How about:

Surgical Knot Tie

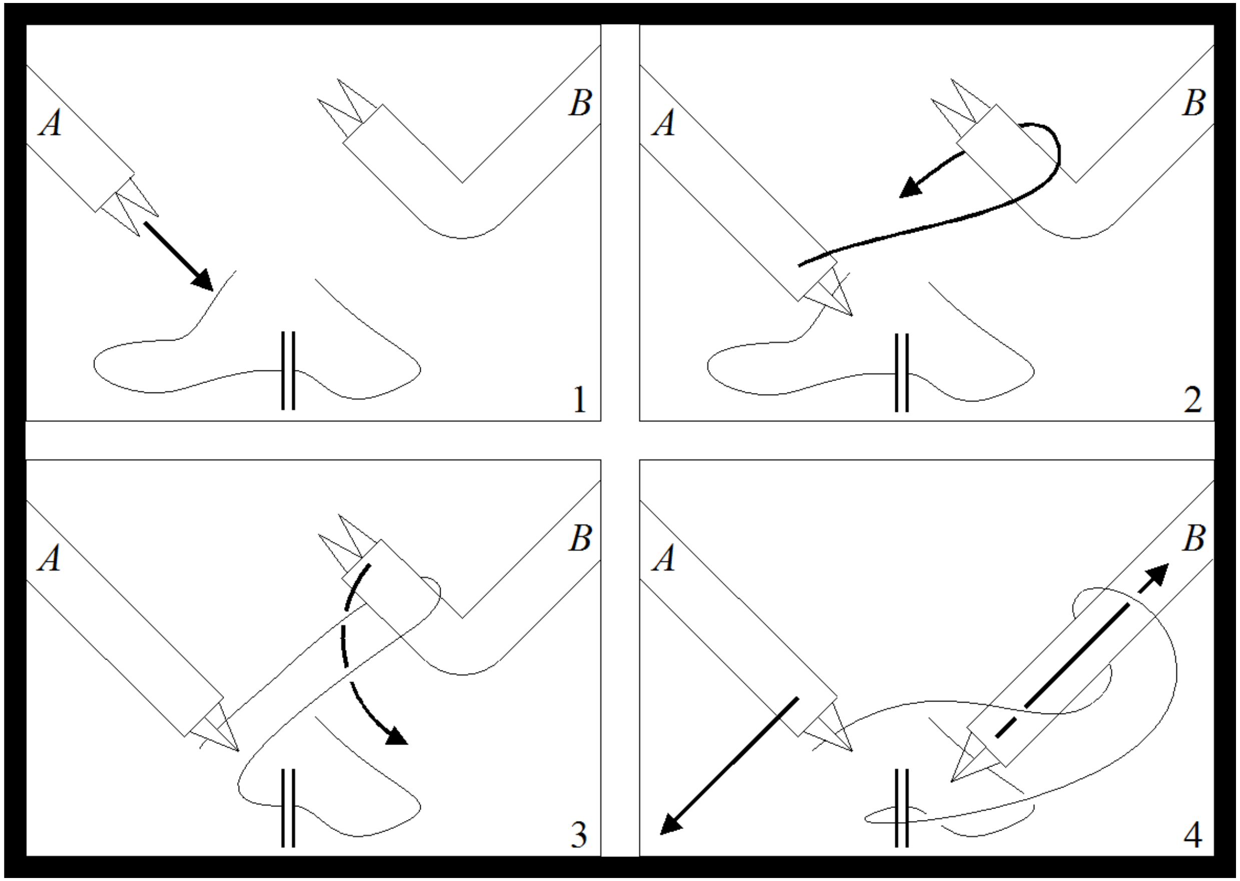

Surgical Knot Tie

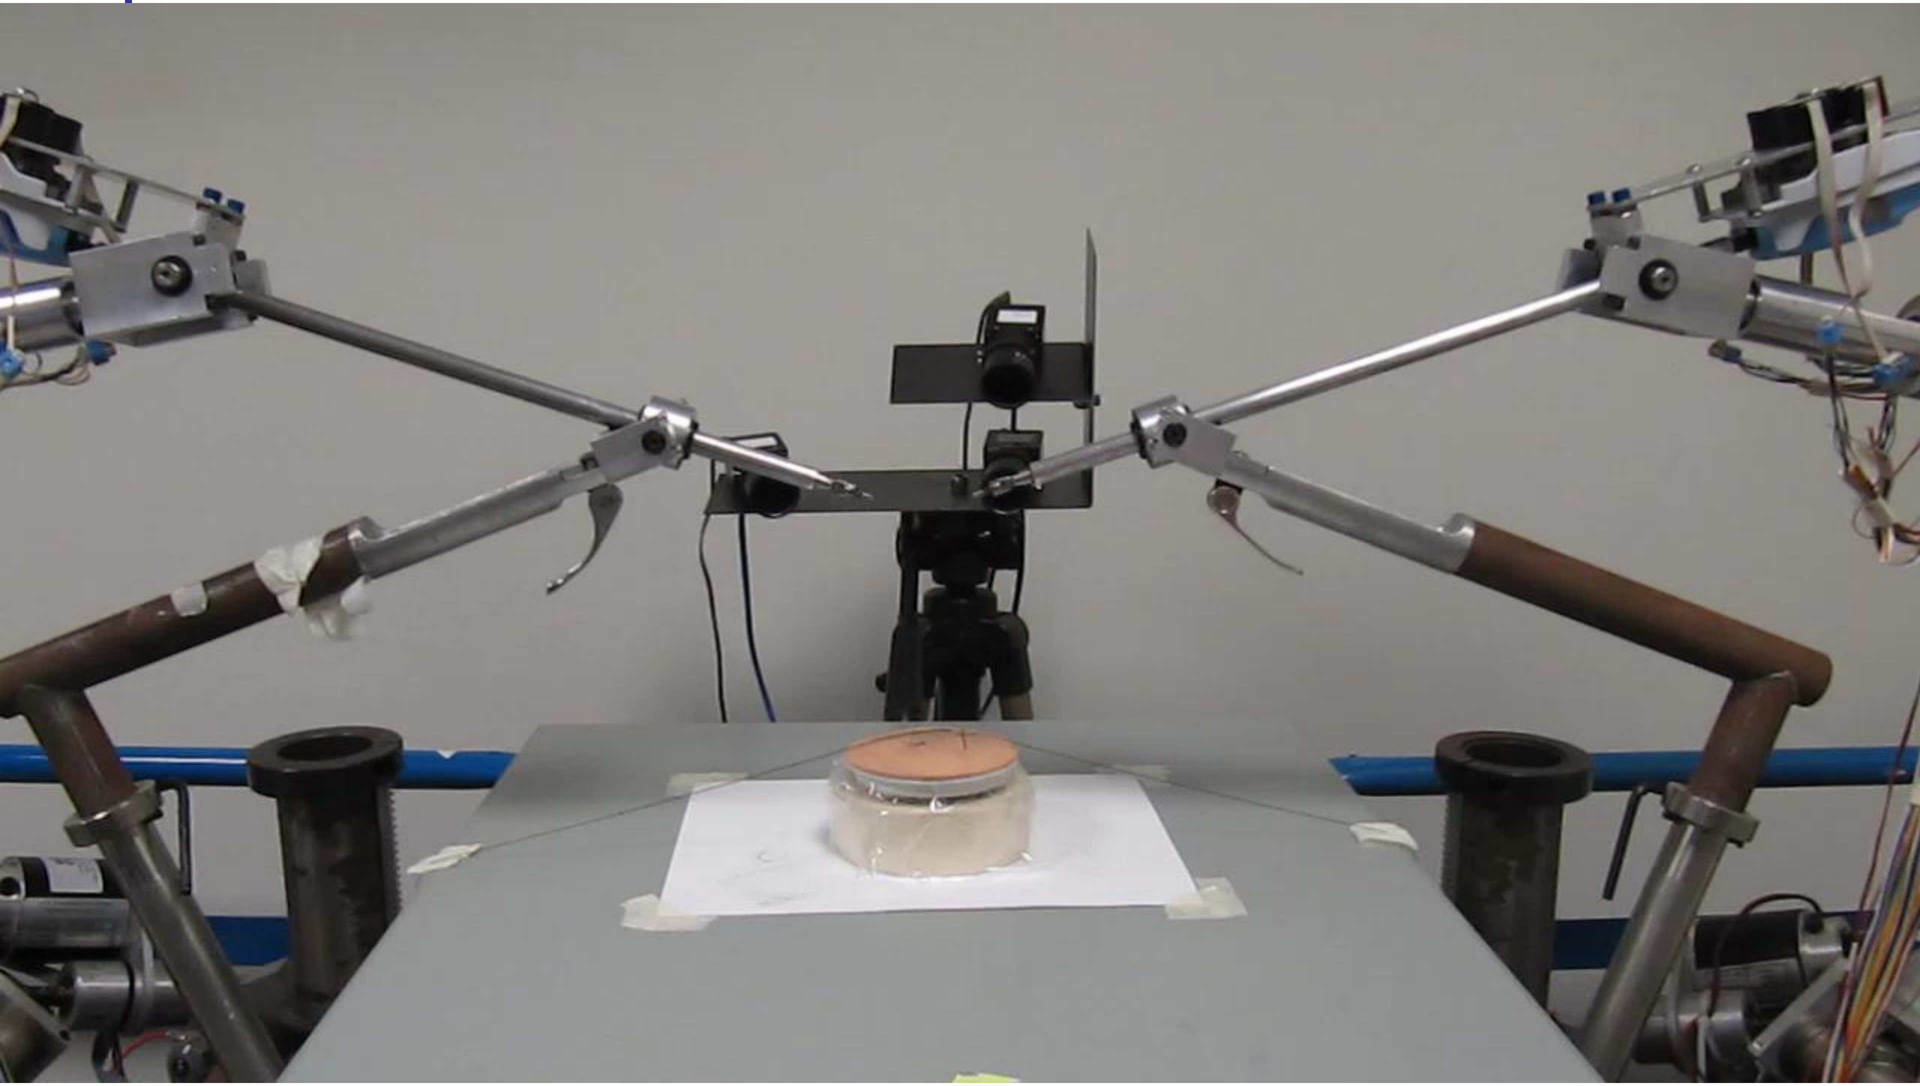

van den Berg, Miller, Duckworth, Humphrey, Wan, Fu, Goldberg, Abbeel, ICRA 2010

Surgical Knot Tie

- Open loop
- If careful about initial conditions
  - 50% success rate

Generalizing Trajectories

- The problem
  - Human demonstrated knot-tie in this rope

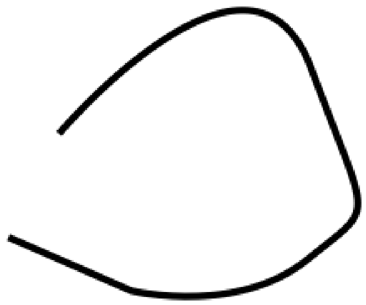

  - Robot has to tie a knot in this rope

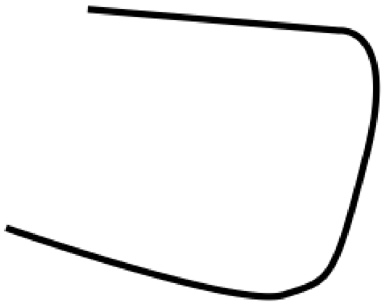

Cartoon Problem Setting

Train situation:

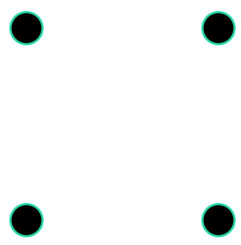

Verb demonstration: --- trajectory

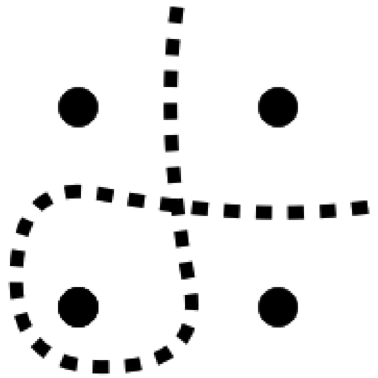

Test situation:

How to perform verb here?

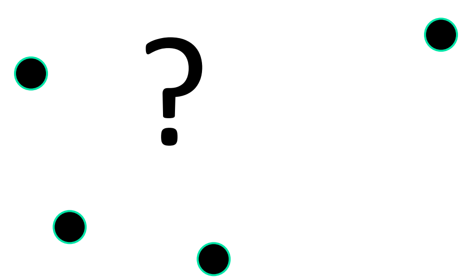

Cartoon Problem Setting

Train situation:

Verb demonstration: --- trajectory

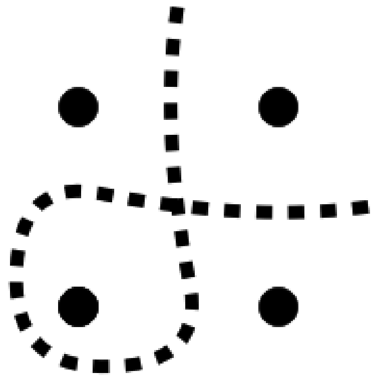

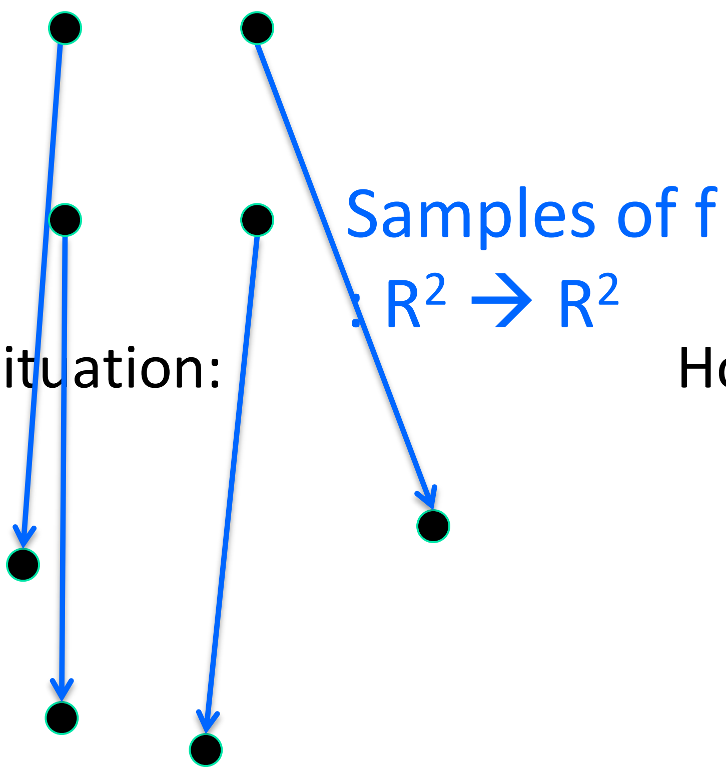

Test situation:

How to perform verb here?

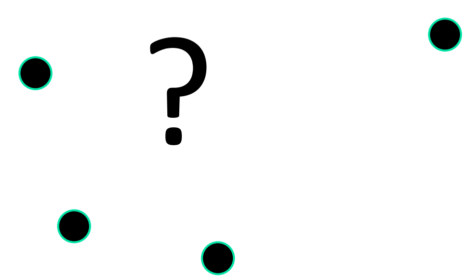

Cartoon Problem Setting

Train situation:

Verb demonstration: --- trajectory

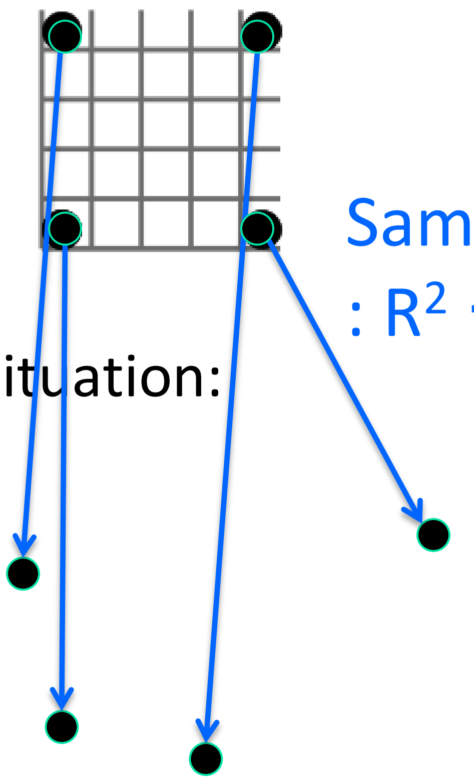

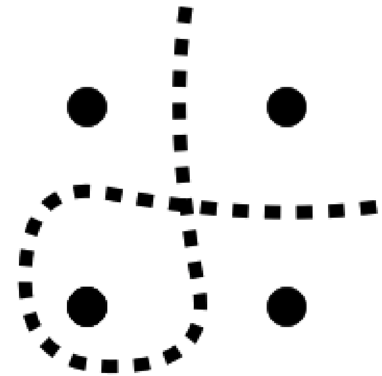

Samples of f $: R^2 \rightarrow R^2$

Test situation:

How to perform verb here?

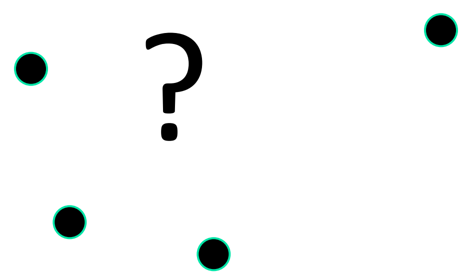

Cartoon Problem Setting

Train situation:

Verb demonstration: --- trajectory

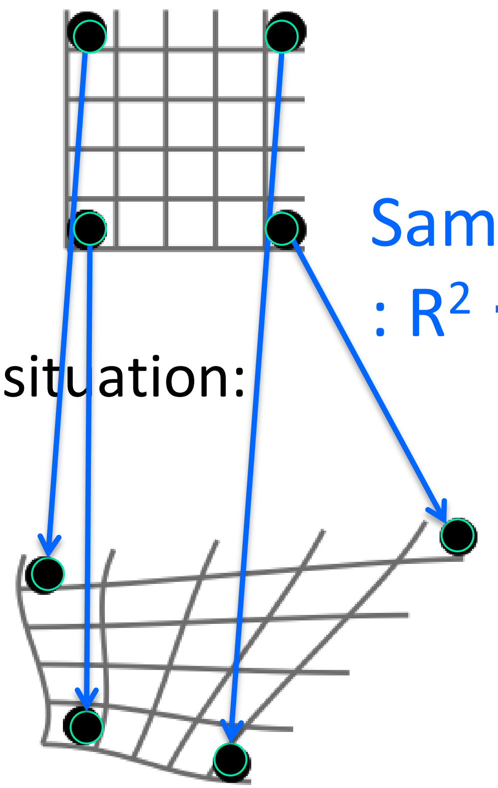

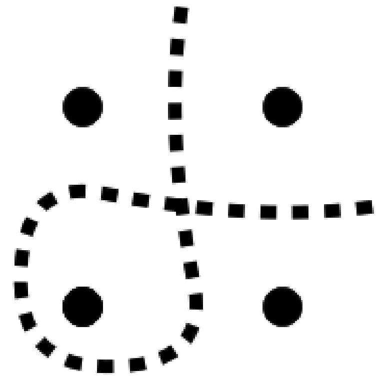

Samples of f

$: R^2 \rightarrow R^2$

Test situation:

How to perform verb here?

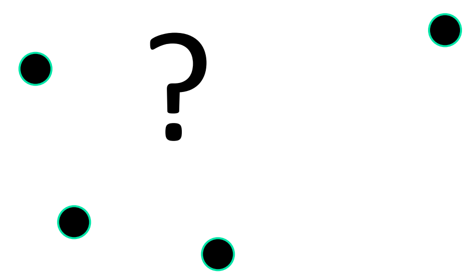

Cartoon Problem Setting

Train situation:

Verb demonstration: --- trajectory

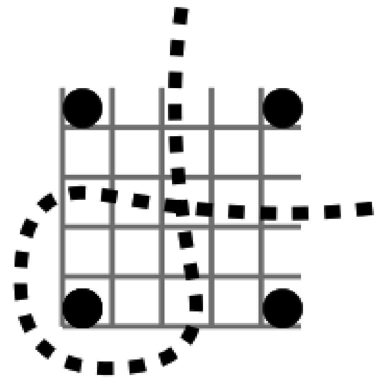

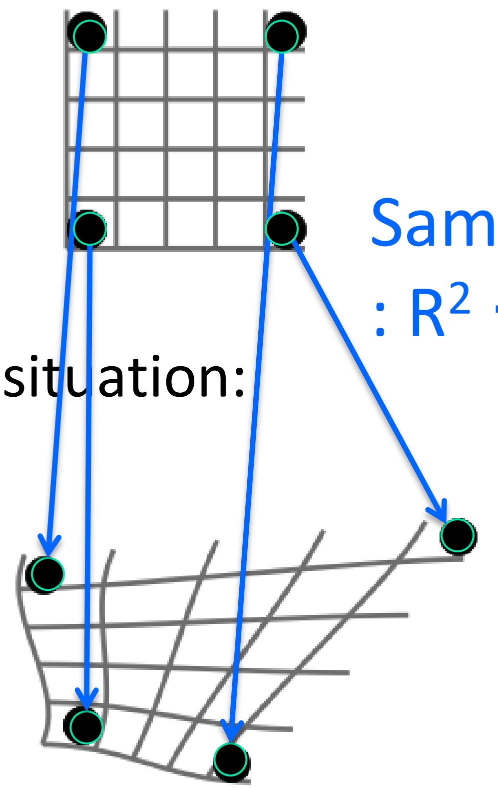

Samples of f : $R^2 \rightarrow R^2$

Test situation:

How to perform verb here?

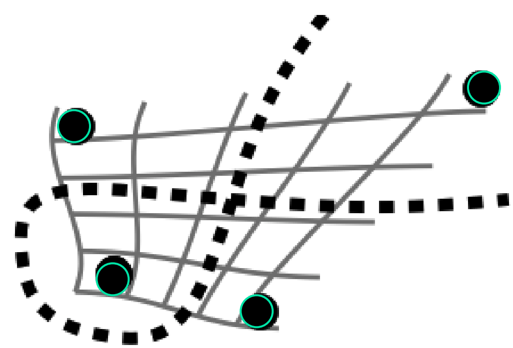

Learning f : R³ → R³ from Samples

$$\min_{f \in \{\mathbb{R}^3 \rightarrow \mathbb{R}^3\}} \quad \int_{x \in \mathbb{R}^3} \|D^2 f\|^2_{\text{Frob}}(x)dx$$

$$\text{s.t.} \quad f(x^{(i)}_{\text{train}}) = x^{(i)}_{\text{test}} \quad \forall i \in 1, \ldots, m$$

- Observations
  - Translations, rotations and scaling are FREE
  - Can be solved efficiently manipulating matrices of size of number of examples

Learning f : R³ → R³ from Samples

$$\min_{f \in \{\mathbb{R}^3 \rightarrow \mathbb{R}^3\}} \quad \int_{x \in \mathbb{R}^3} \|D^2 f\|^2_{\text{Frob}}(x)dx$$

$$\text{s.t.} \quad f(x^{(i)}_{\text{train}}) = x^{(i)}_{\text{test}} \quad \forall i \in 1, \ldots, m$$

- Solution has form:

$$f(x) = \sum_{i=1}^{m} a_i K(x_i, x) + b^\top x + c,$$

$$K(x,y) = \begin{cases} c_0 r^{4-d} \ln r, & d = 2 \text{ or } d = 4 \\ c_1 r^{4-d}, & \text{otherwise} \end{cases} \qquad \text{with} \quad r = \|x - y\|_2.$$

Wahba, Spline models for observational data. Philadelphia: Society for Industrial and Applied Mathematics. 1990.
Evgeniou, Pontil, Poggio, Regularization Networks and Support Vector Machines. Advances in Computational Mathematics. 2000.
Hastie, Tibshirani, Friedman, Elements of Statistical Learning, Chapter 5. 2008.

Experiment: Knot-Tie

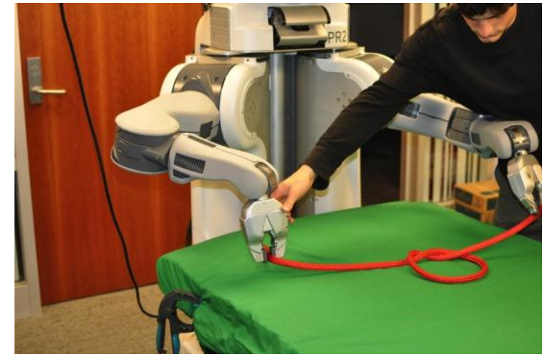

Knot-Tie in New Situation

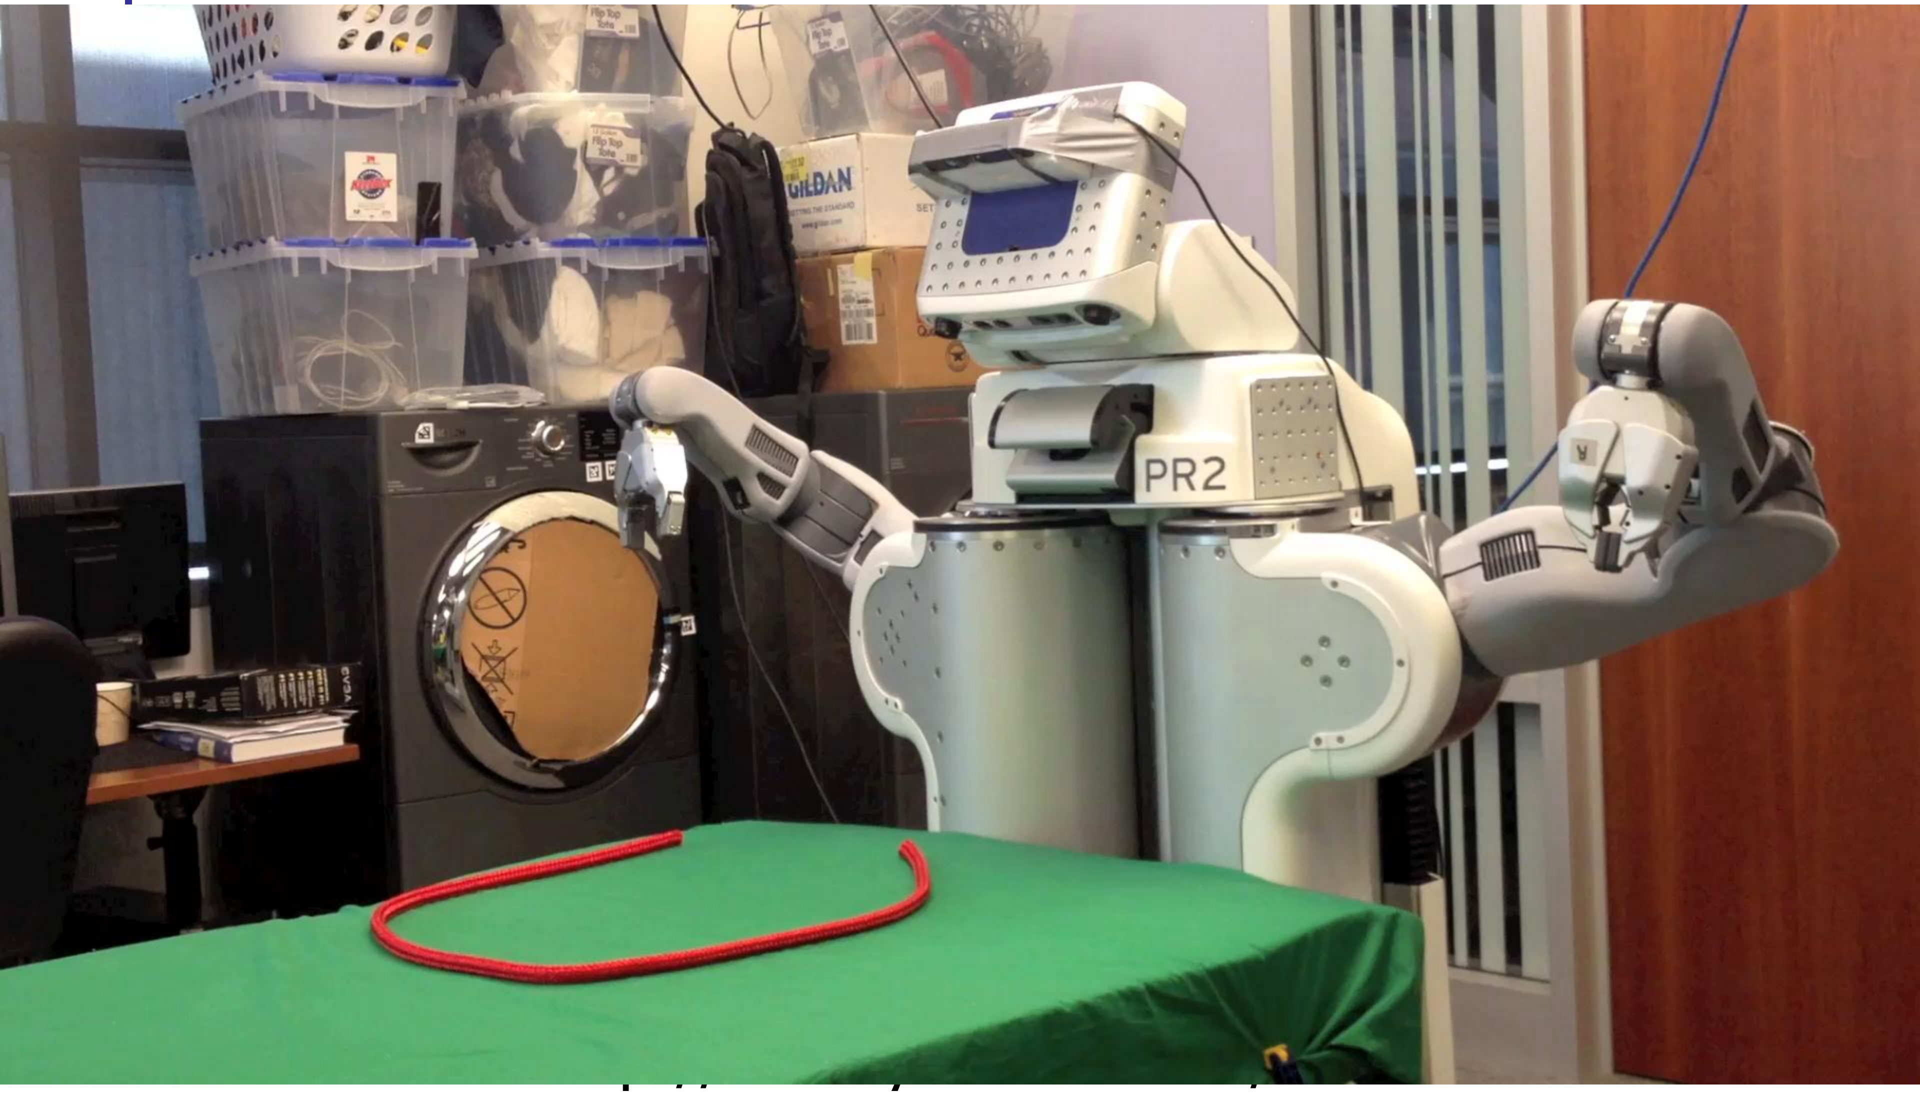

Schulman & Abbeel, TBD 2012

Experiments: Plate Pick-Up

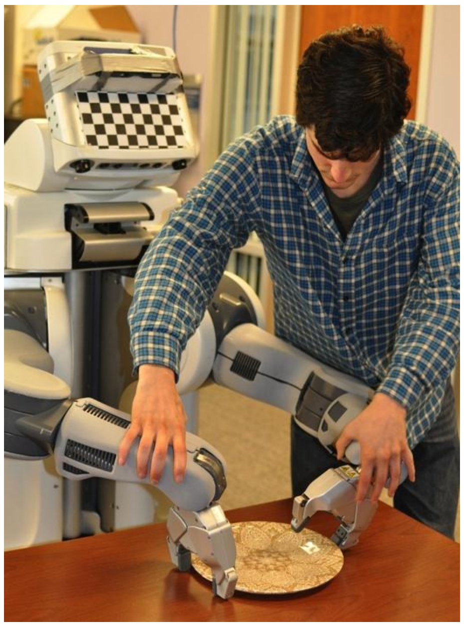

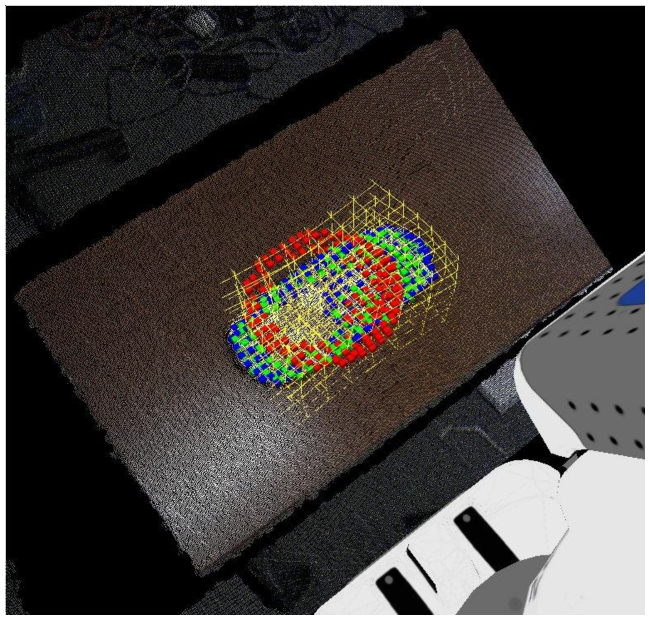

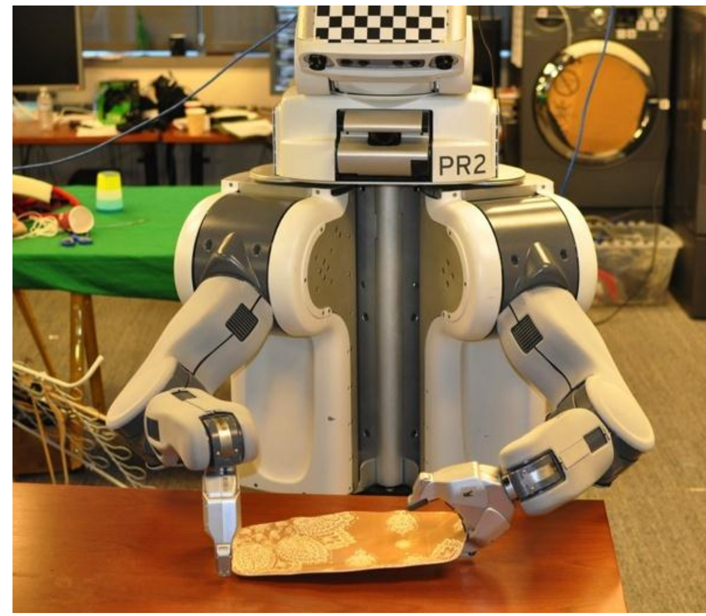

Experiments: Scooping

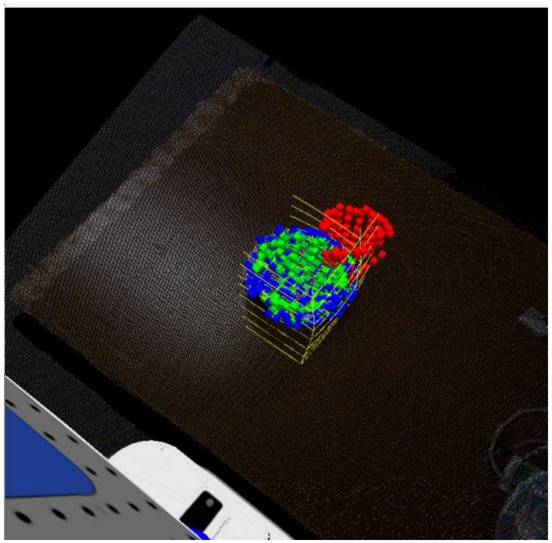

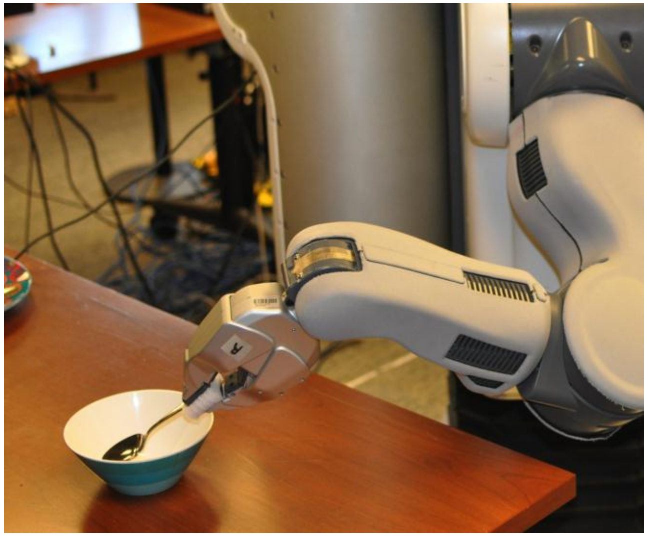

List of Publications Attributed to the Grant

A Geometry-Oriented Approach to Learning from Demonstrations for Robotic Manipulation

John Schulman, Jonathan Ho, Cameron Lee, Pieter Abbeel

Under review for Robotics: Science and Systems (RSS) 2013.