Learning Approaches for Grasping

Jan Peters, Oliver Kroemer
Max Planck Institute for Biological Cybernetics

Rendaud Detry, Justus Piater
Universite de Liege

Motivation

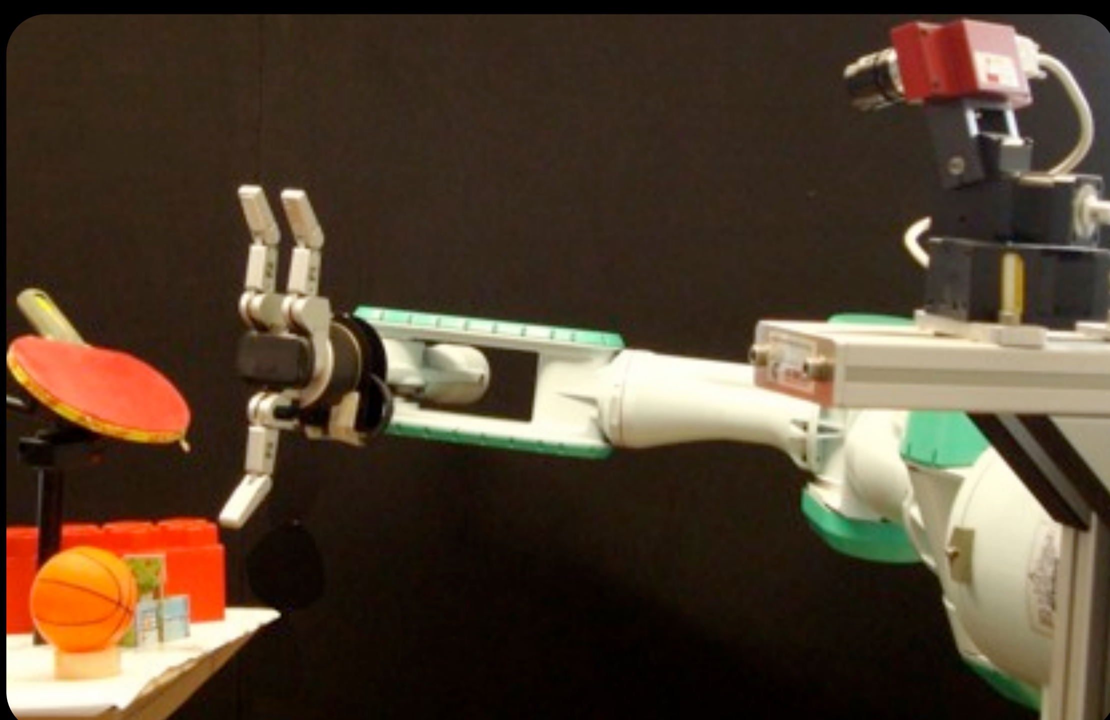

Why learning for grasping?

- Grasping is generically data-driven!
- New objects often need to be explored
- Motions need to be adapted to the object

➡ Two important topics are:
1. "Smart" *exploration* of new grasps
2. *generalizing* grasping movements

Collaborators

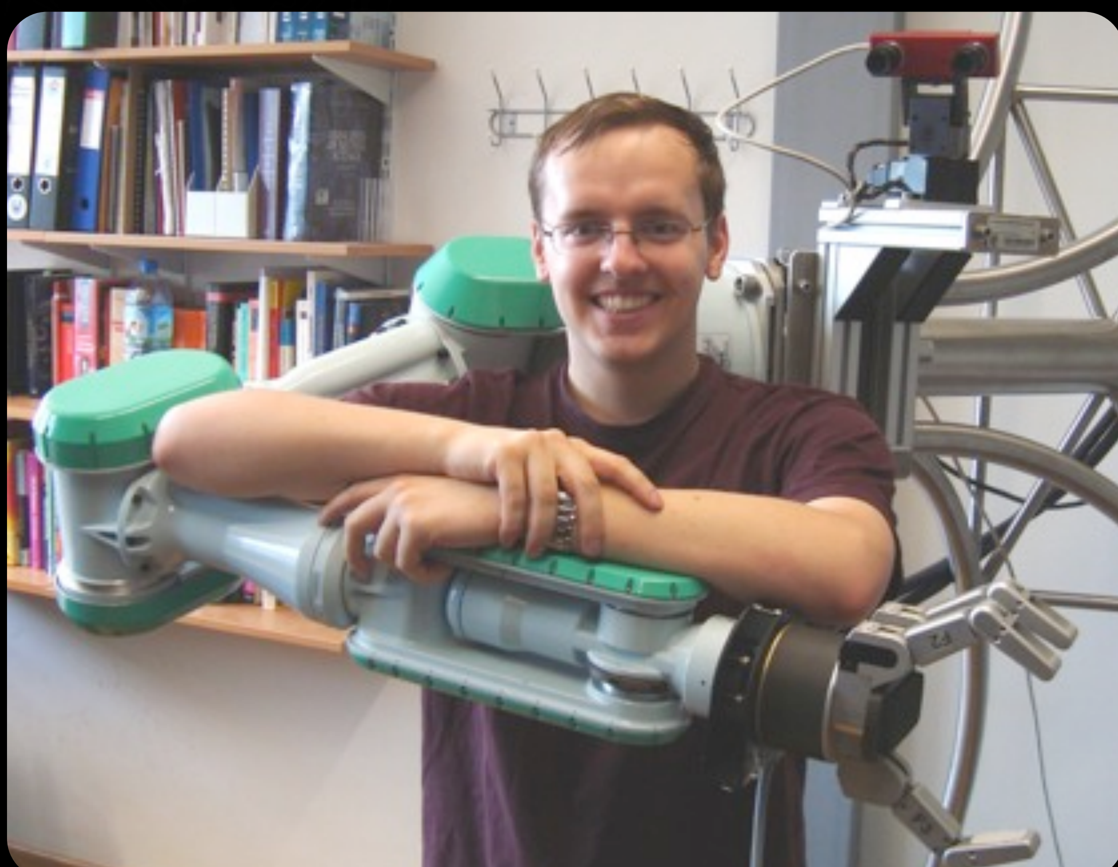

Outline

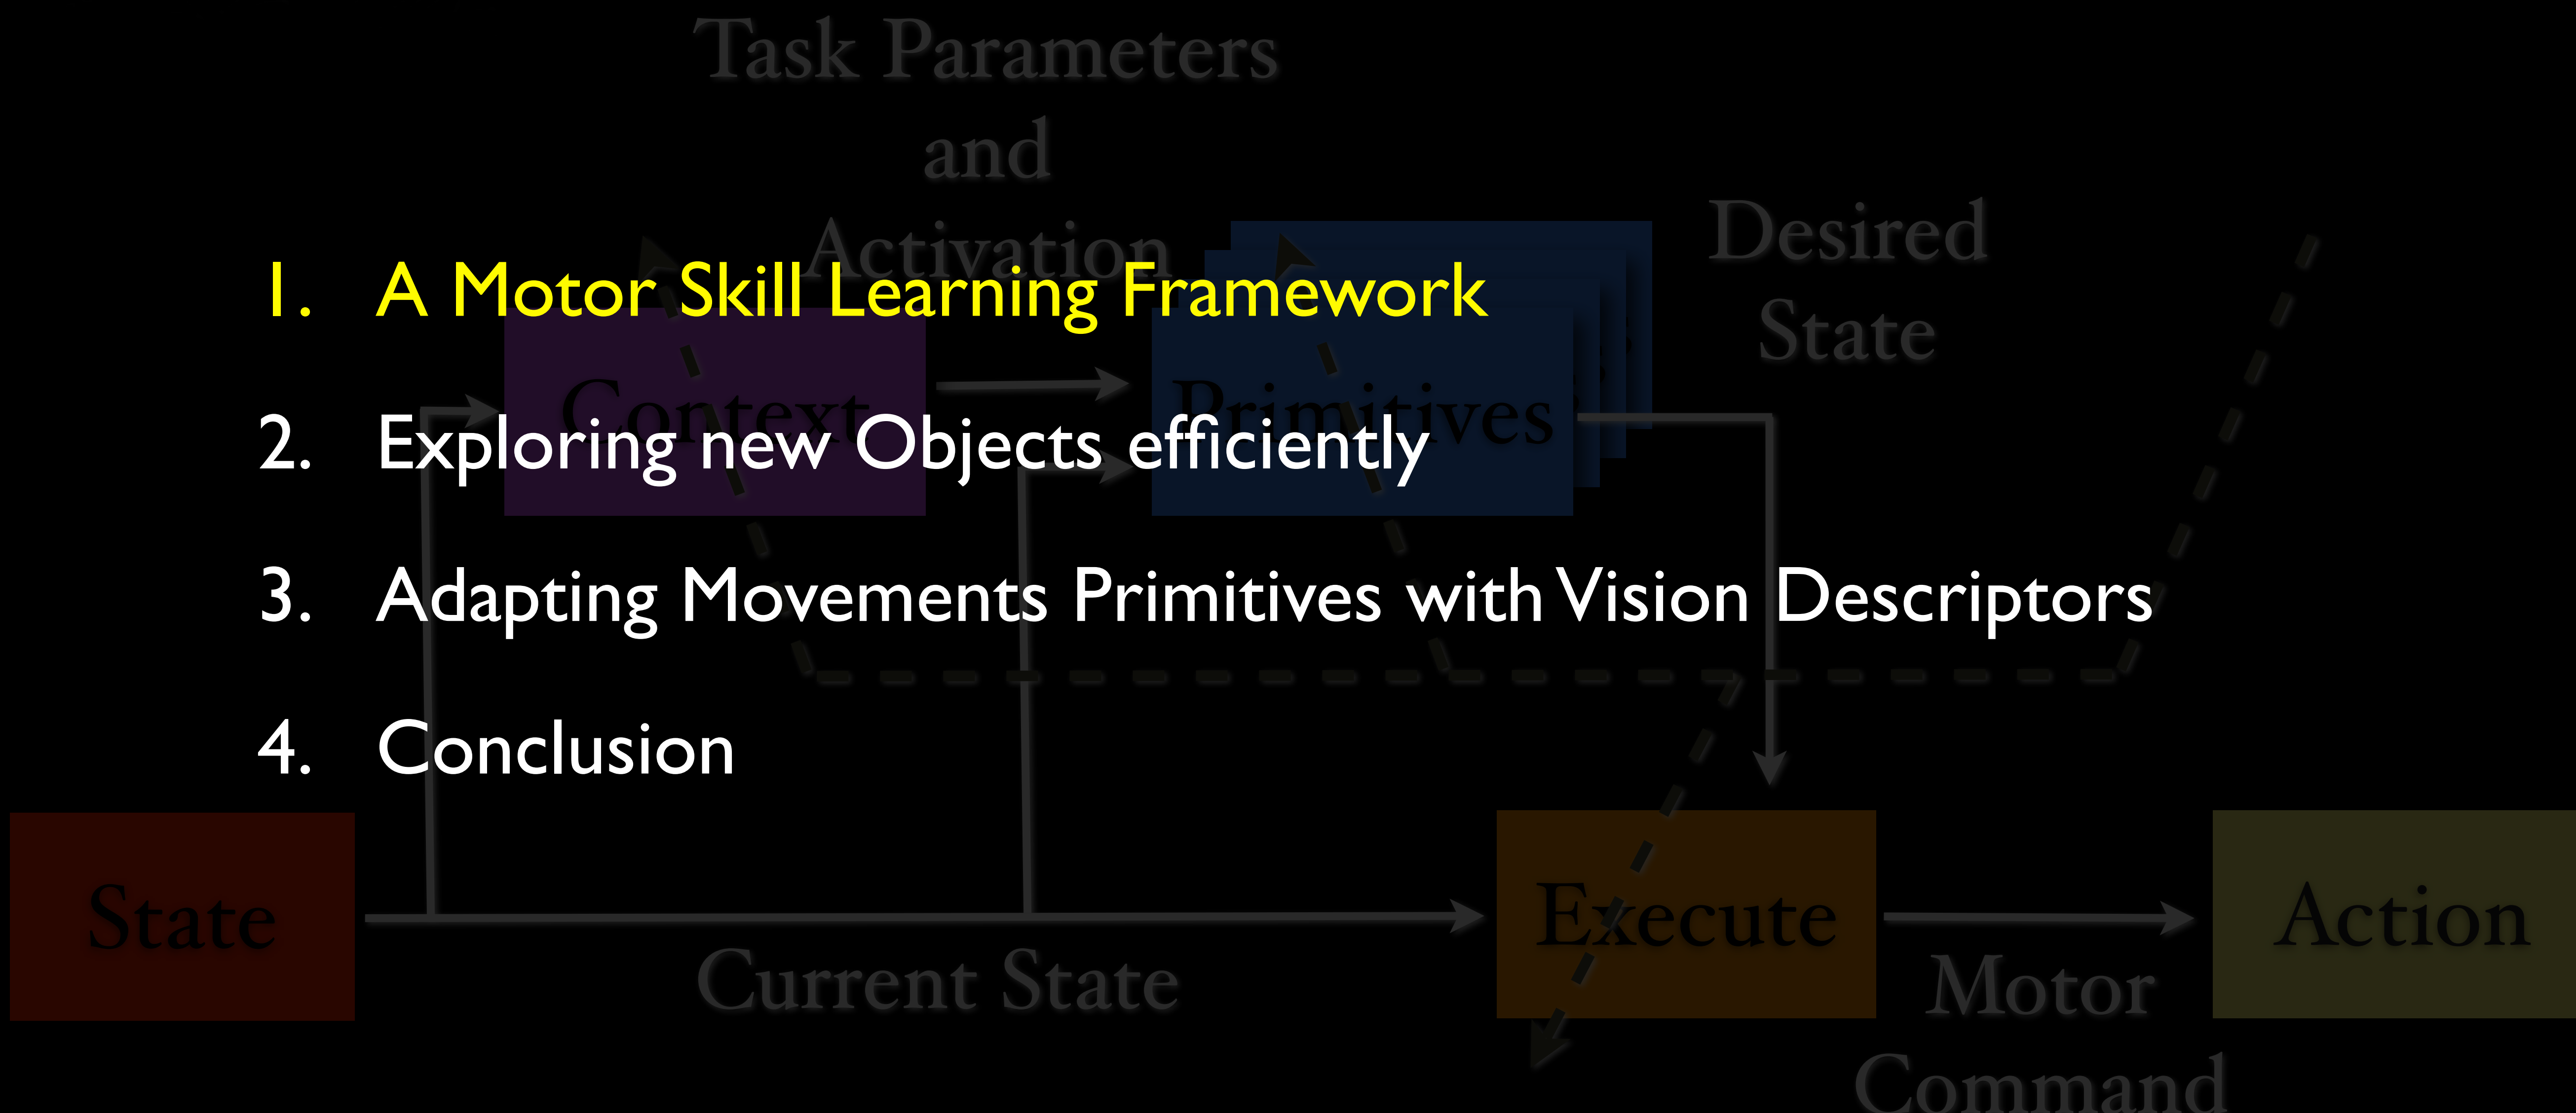

1. A Motor Skill Learning Framework
2. Exploring new Objects efficiently
3. Adapting Movements Primitives with Vision Descriptors
4. Conclusion

Motor Primitives

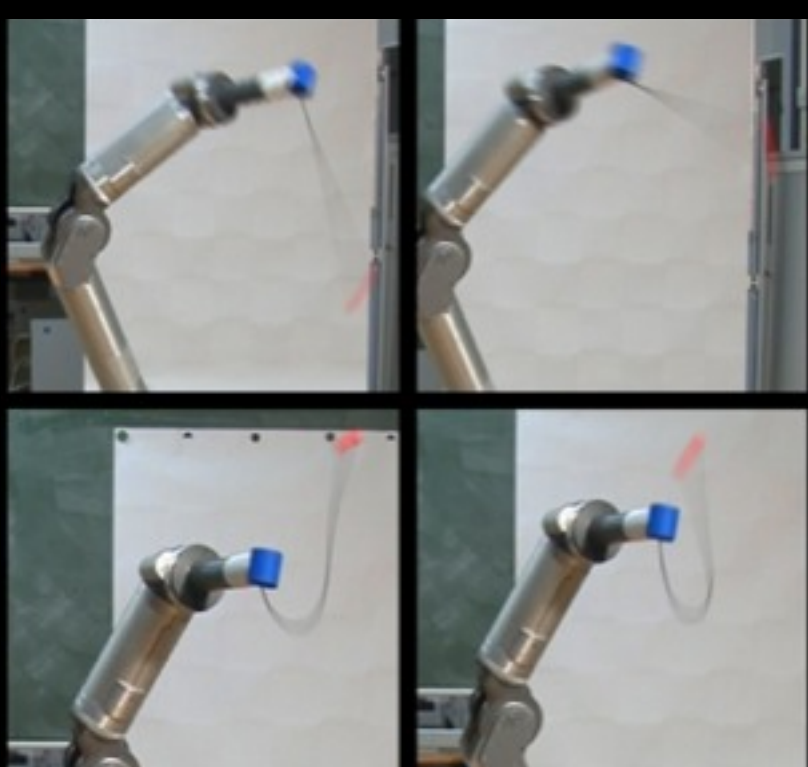

How can we learn a set of skills?

- Humans appear to rely on context-driven motor primitives (Flash & Hochner, 2005)
- Dynamic system-based motor primitives (Schaal, Peters, Nakanishi, Ijspeert, ISRR2003) offer a computational alternative.
- Primitives need to be suitable for fast learning.

➡ *Resulting approach:*

- Initialize by Imitation Learning.
- Improve by trial and error on the real system with Reinforcement Learning.
- Adjust primitives using context information.

Peters et al. (2009). Towards Motor Skill Learning for Robotics, ISRR

Motor Primtives

Task/Hyperparameter

Trajectory Plan Dynamics

$$\dot{z} = \alpha_z(\beta_z(g-y)-z)$$
$$\dot{y} = \alpha_y(f(x,v)+z)$$

where

Canonical Dynamics

$$\dot{v} = \alpha_v(\beta_v(g-x)-v)$$
$$\dot{x} = \alpha_x v$$

Linear in learnable Policy Parameters

Local Linear Model Approx.

$$f(x,v) = \frac{\sum_{i=1}^{k} w_i b_i v}{\sum_{i=1}^{k} w_i}$$

$$w_i = \exp\left(-\frac{1}{2} d_i (\bar{x}-c_i)^2\right) \text{ and } \bar{x} = \frac{x-x_0}{g-x_0}$$

(Schaal, Peters, Nakanishi, Ijspeert, ISRR 2003)

Acquisition by Imitation

Teacher shows the task and the student reproduces it.

- maximize similarity

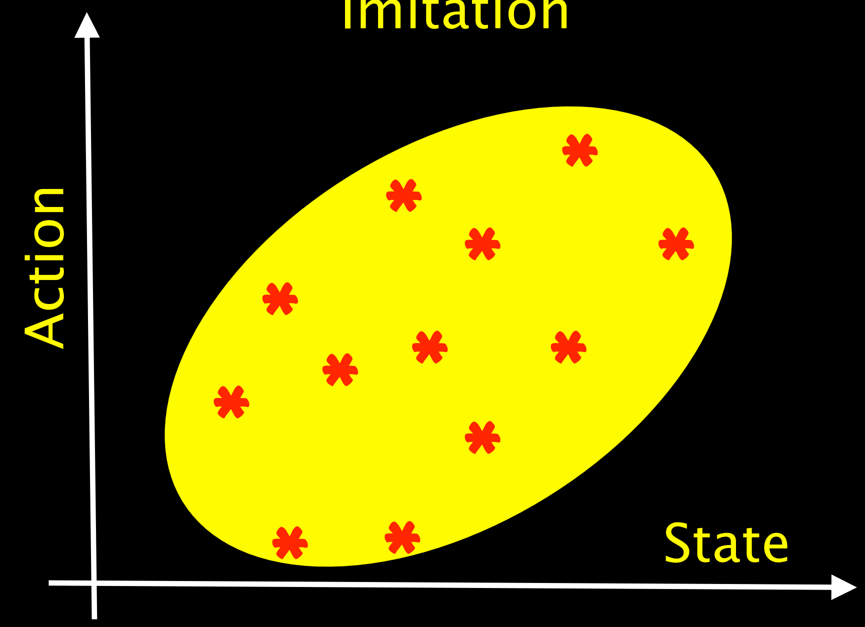

Kober & Peters (2009). Learning Motor Primitives, ICRA

Self-Improvement by Reinforcement Learning

Student improves by reproducing his successful trials.

- maximize reward-weighted similarity

Kober & Peters (2009). Policy Search for Motor Primitives in Robotics, NIPS

Motor Primtives

Task/Hyperparameter

Trajectory Plan Dynamics

$$\begin{cases} \dot{z} = \alpha_z\left(\beta_z\left(g - y\right) - z\right) \\ \dot{y} = \alpha_y\left(f\left(x, v\right) + z\right) \end{cases}$$

where

Linear in learnable Policy Parameters

Canonical Dynamics

$$\begin{cases} \dot{v} = \alpha_v\left(\beta_v\left(g - x\right) - v\right) \\ \dot{x} = \alpha_x v \end{cases}$$

Local Linear Model Approx.

$$\begin{cases} f\left(x, v\right) = \dfrac{\sum_{i=1}^{k} w_i b_i v}{\sum_{i=1}^{k} w_i} \\ w_i = \exp\left(-\dfrac{1}{2} d_i \left(\bar{x} - c_i\right)^2\right) \text{ and } \bar{x} = \dfrac{x - x_0}{g - x_0} \end{cases}$$

(Schaal, Peters, Nakanishi, Ijspeert, ISRR 2003)

Task Context: Goal Learning

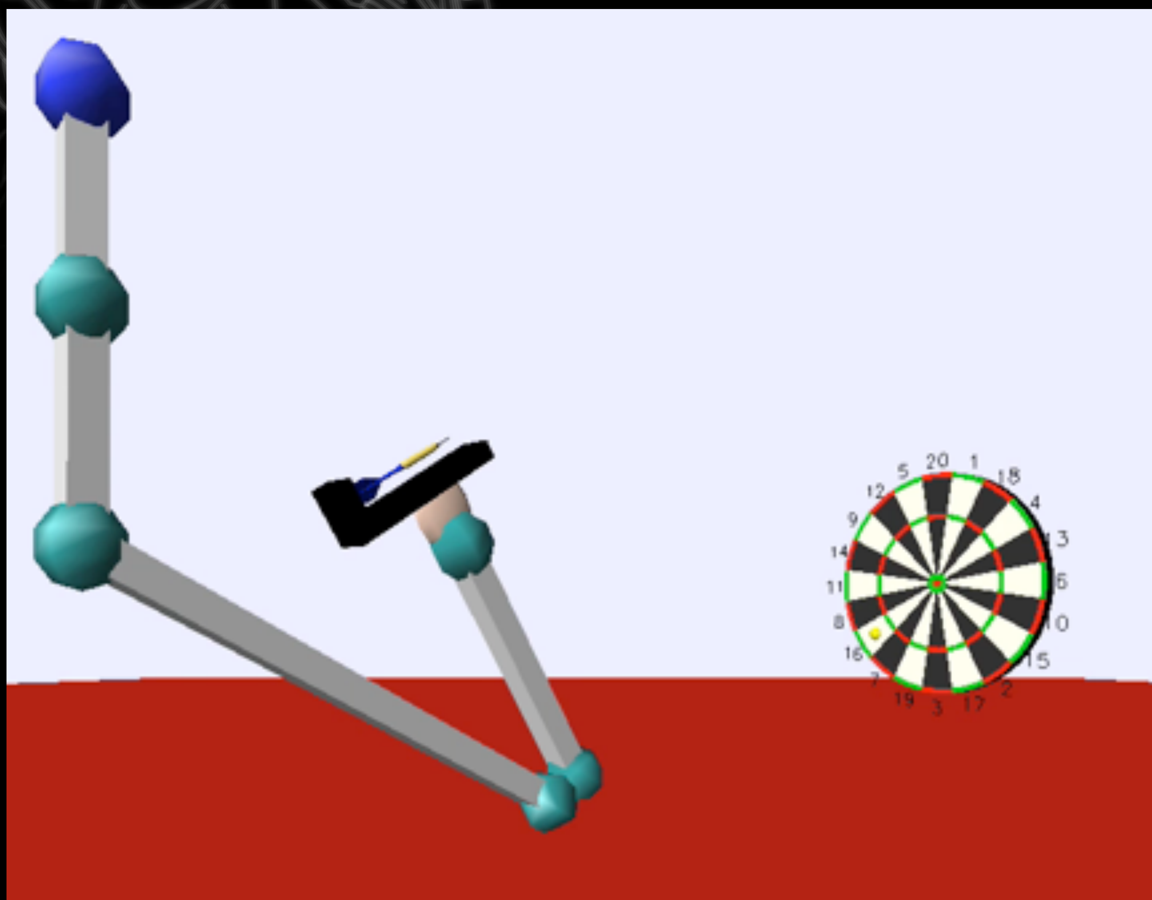

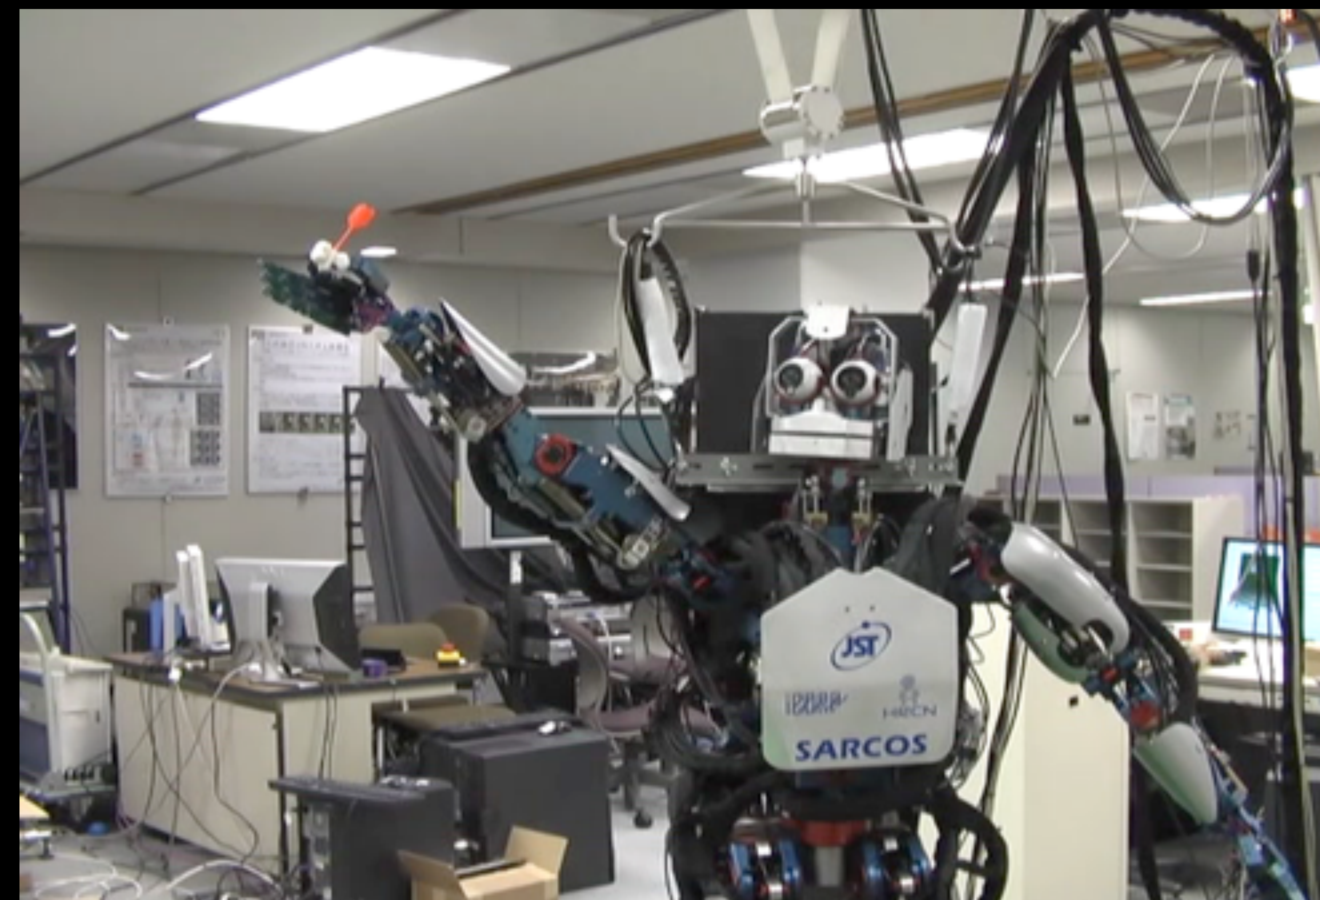

Adjusting Motor Primitives through their Hyperparameters:

- learn a single motor primitive using imitation and reinforcement learning
- DMPs are goal and timing invariant, hence single example suffices
- learn policies for the goal parameter and timing parameters by reinforcement learning

Kober, Oztop & Peters (2010). Reinforcement Learning to adjust Robot Movements to New Situations, R:SS

All of the above?

Preliminary Work

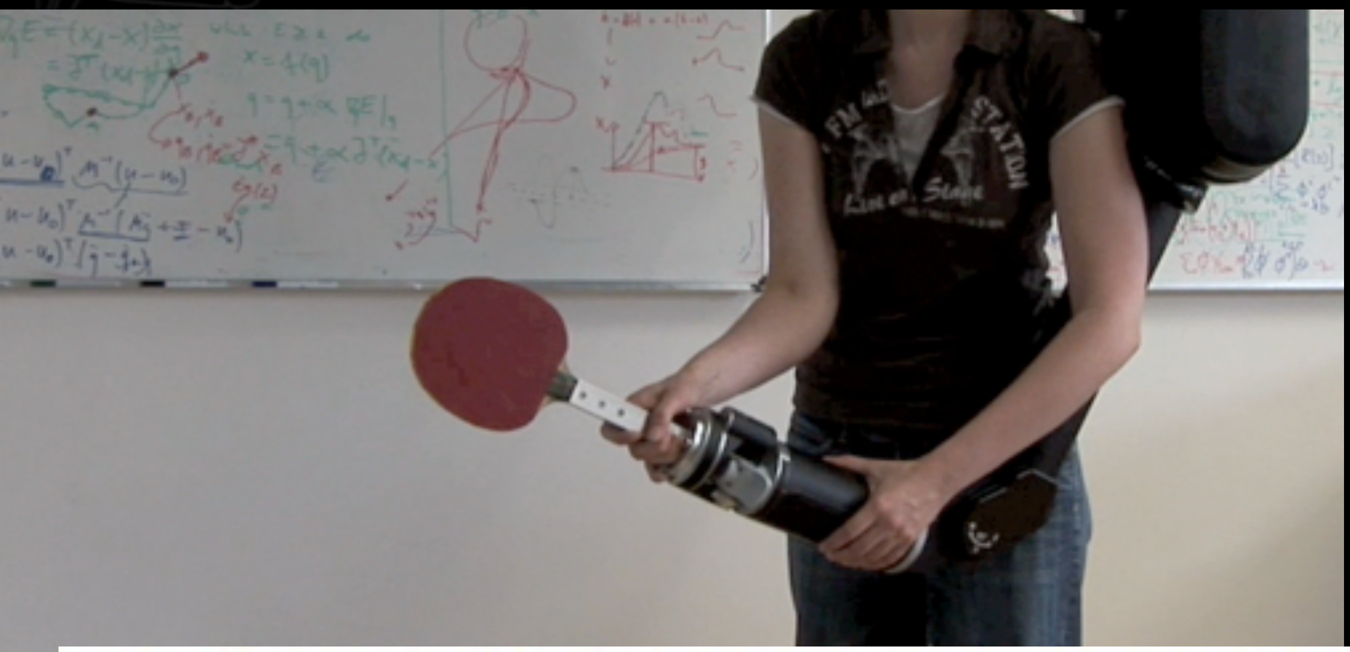

TuD7 Regular Sessions, Egan Center Lower Level Room 11/12

Behavior Learning

14:50-15:05, Paper TuD7.3 — Add to My Program

Movement Templates for Learning of Hitting and Batting ‡

Kober, Jens — Max-Planck Inst. for Biological Cybernetics

Mülling, Kober & Peters (unpublished). Learning to Play Ping.

Kober et al. (2010). Movement Templates for Learning of Hitting and Batting, ICRA

Kober, Oztop & Peters (2010). Reinforcement Learning to adjust Robot Movements to New Situations, R:SS

Blue Print for Skill Learning

Does it apply to grasping?

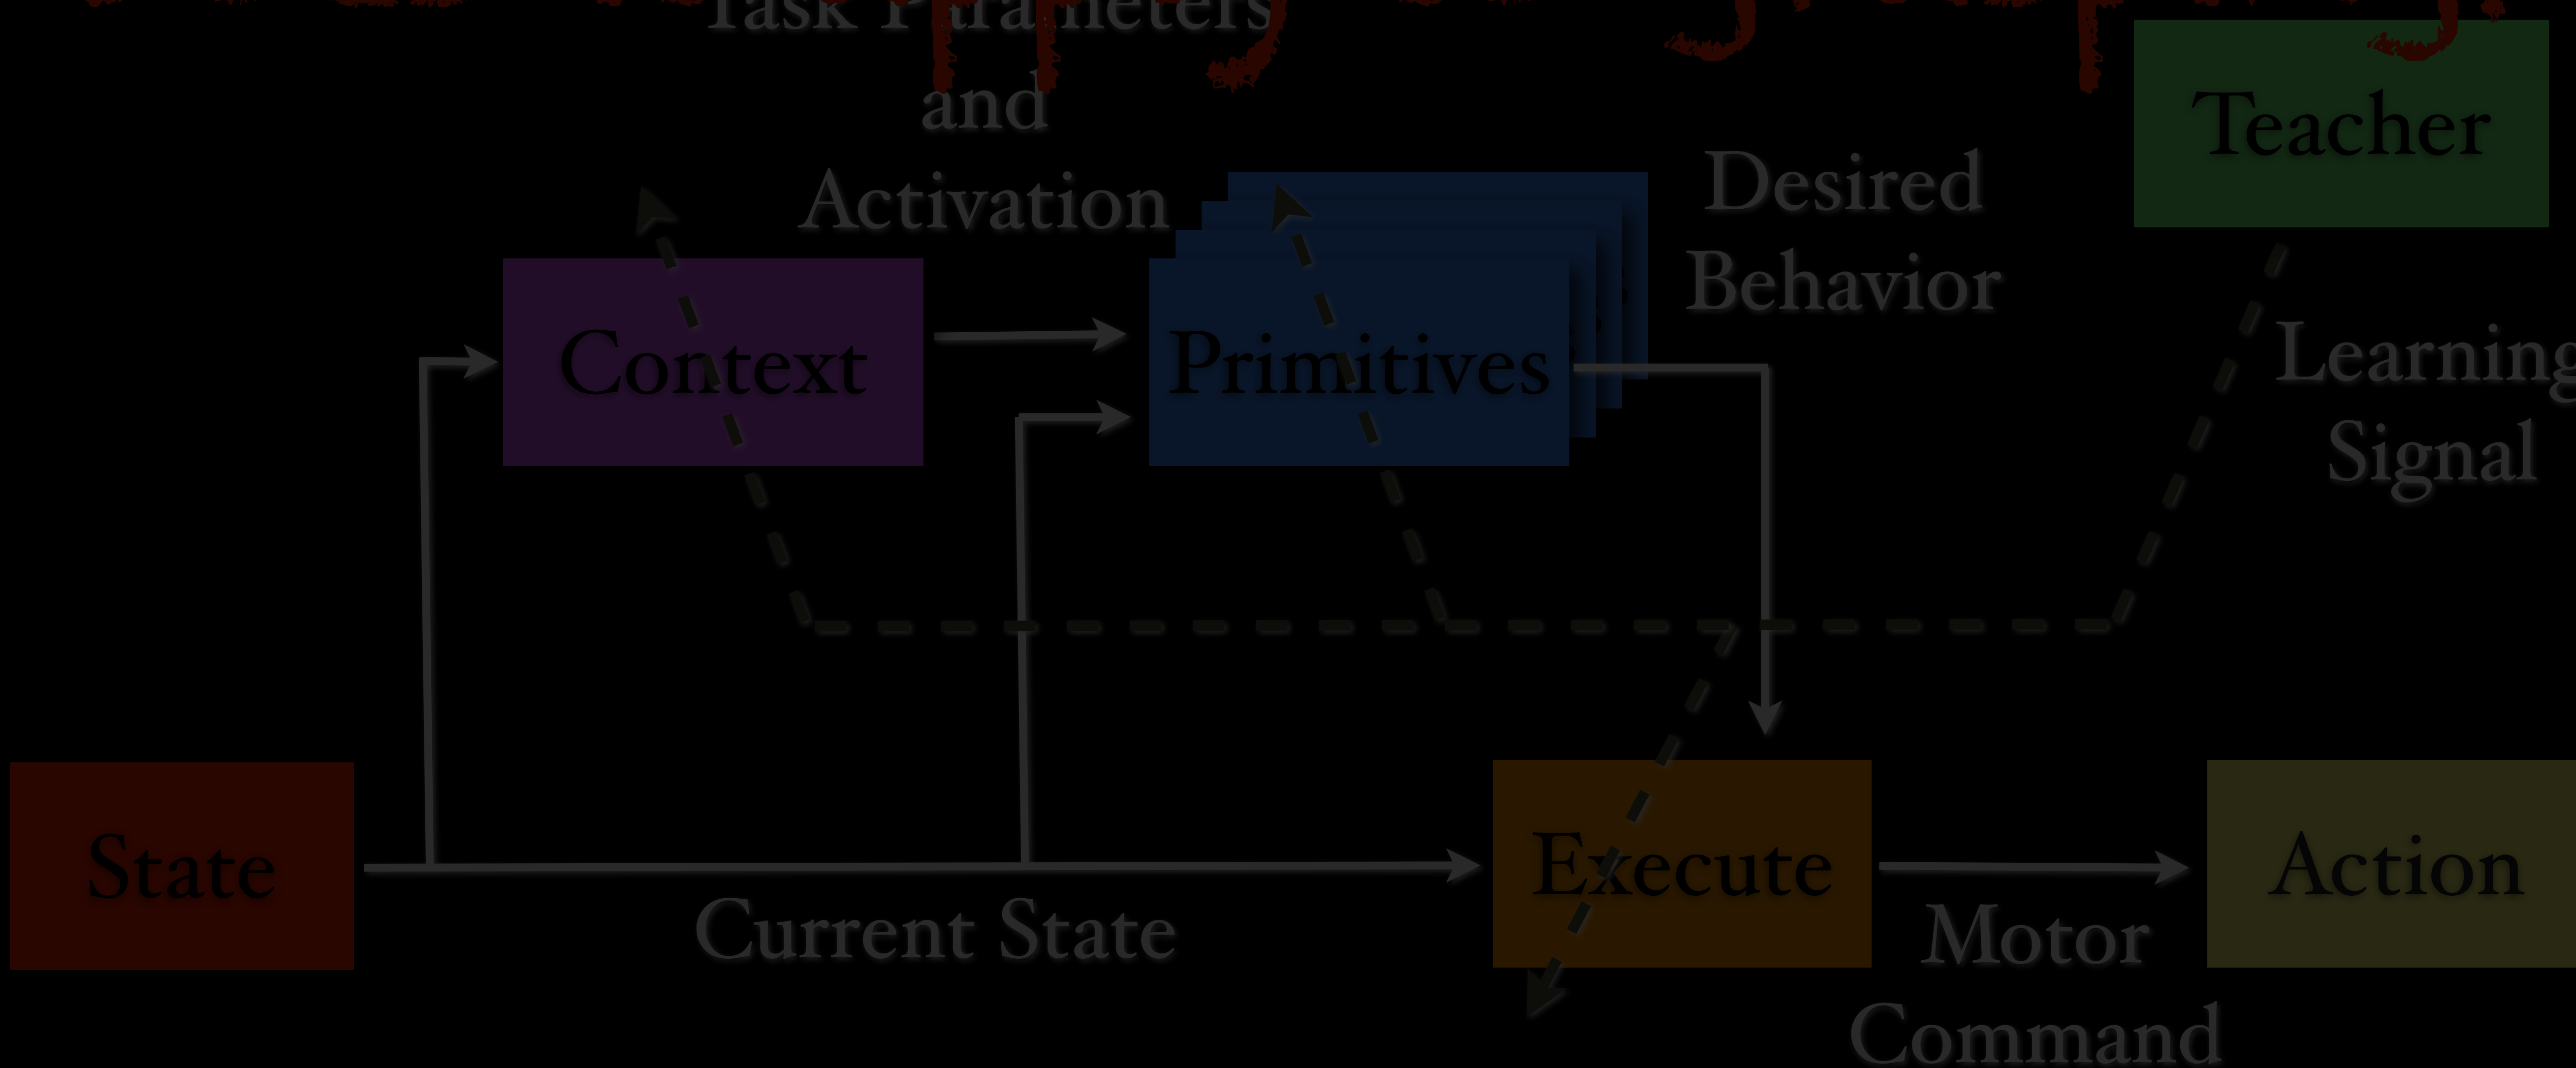

Peters et al. (2009). Towards Motor Skill Learning for Robotics, ISRR

Today's Questions

How can we adapt primitives to new grasps?

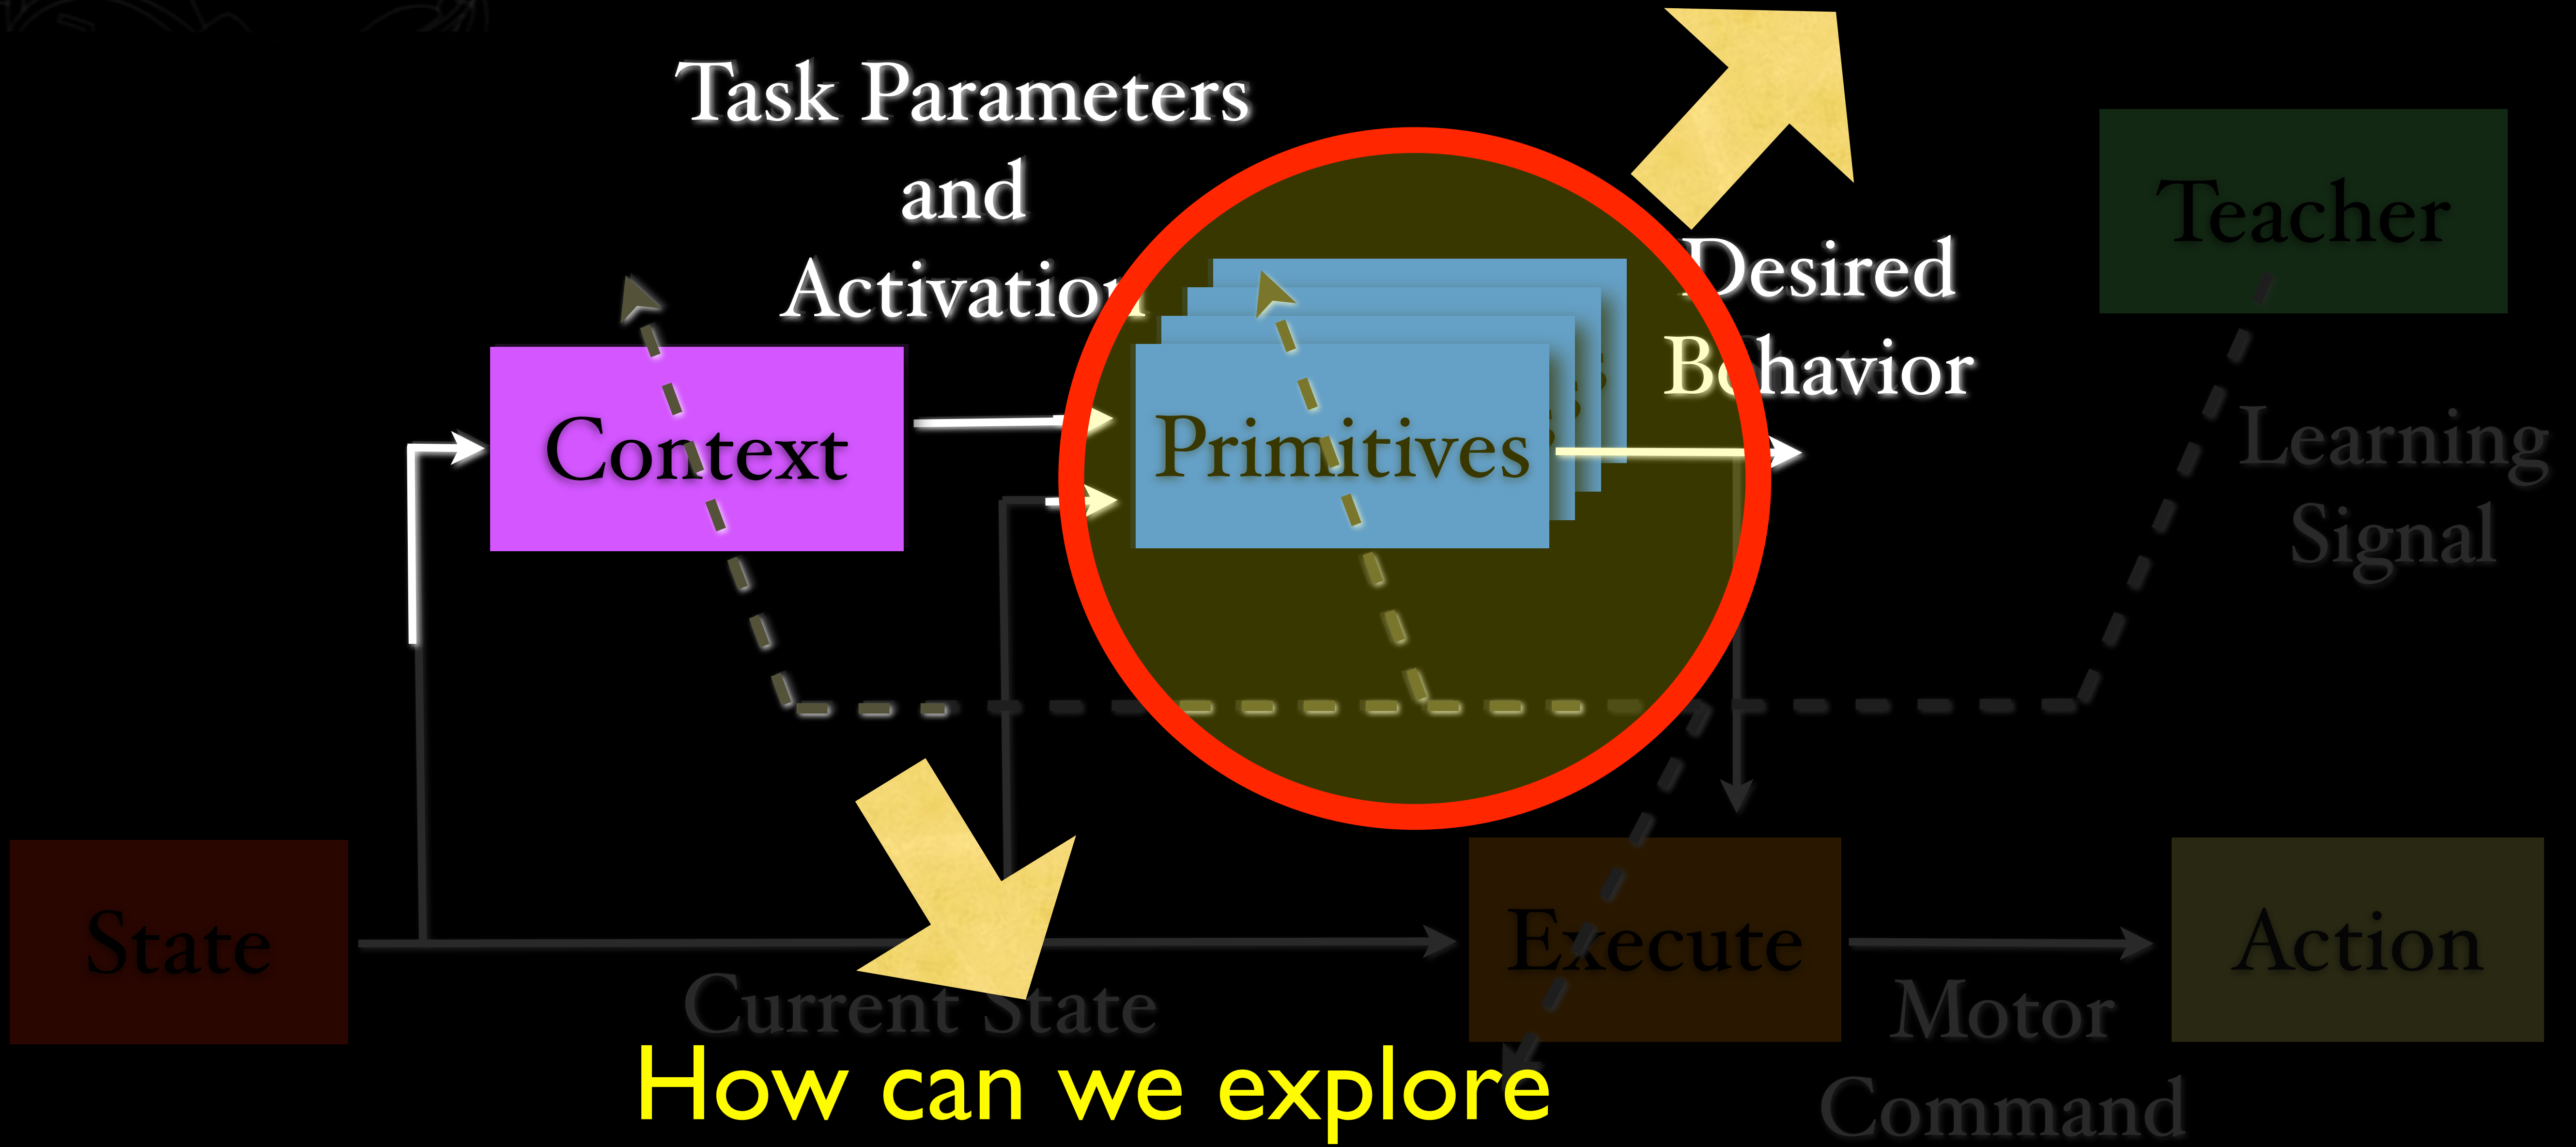

How can we explore an object efficiently?

Outline

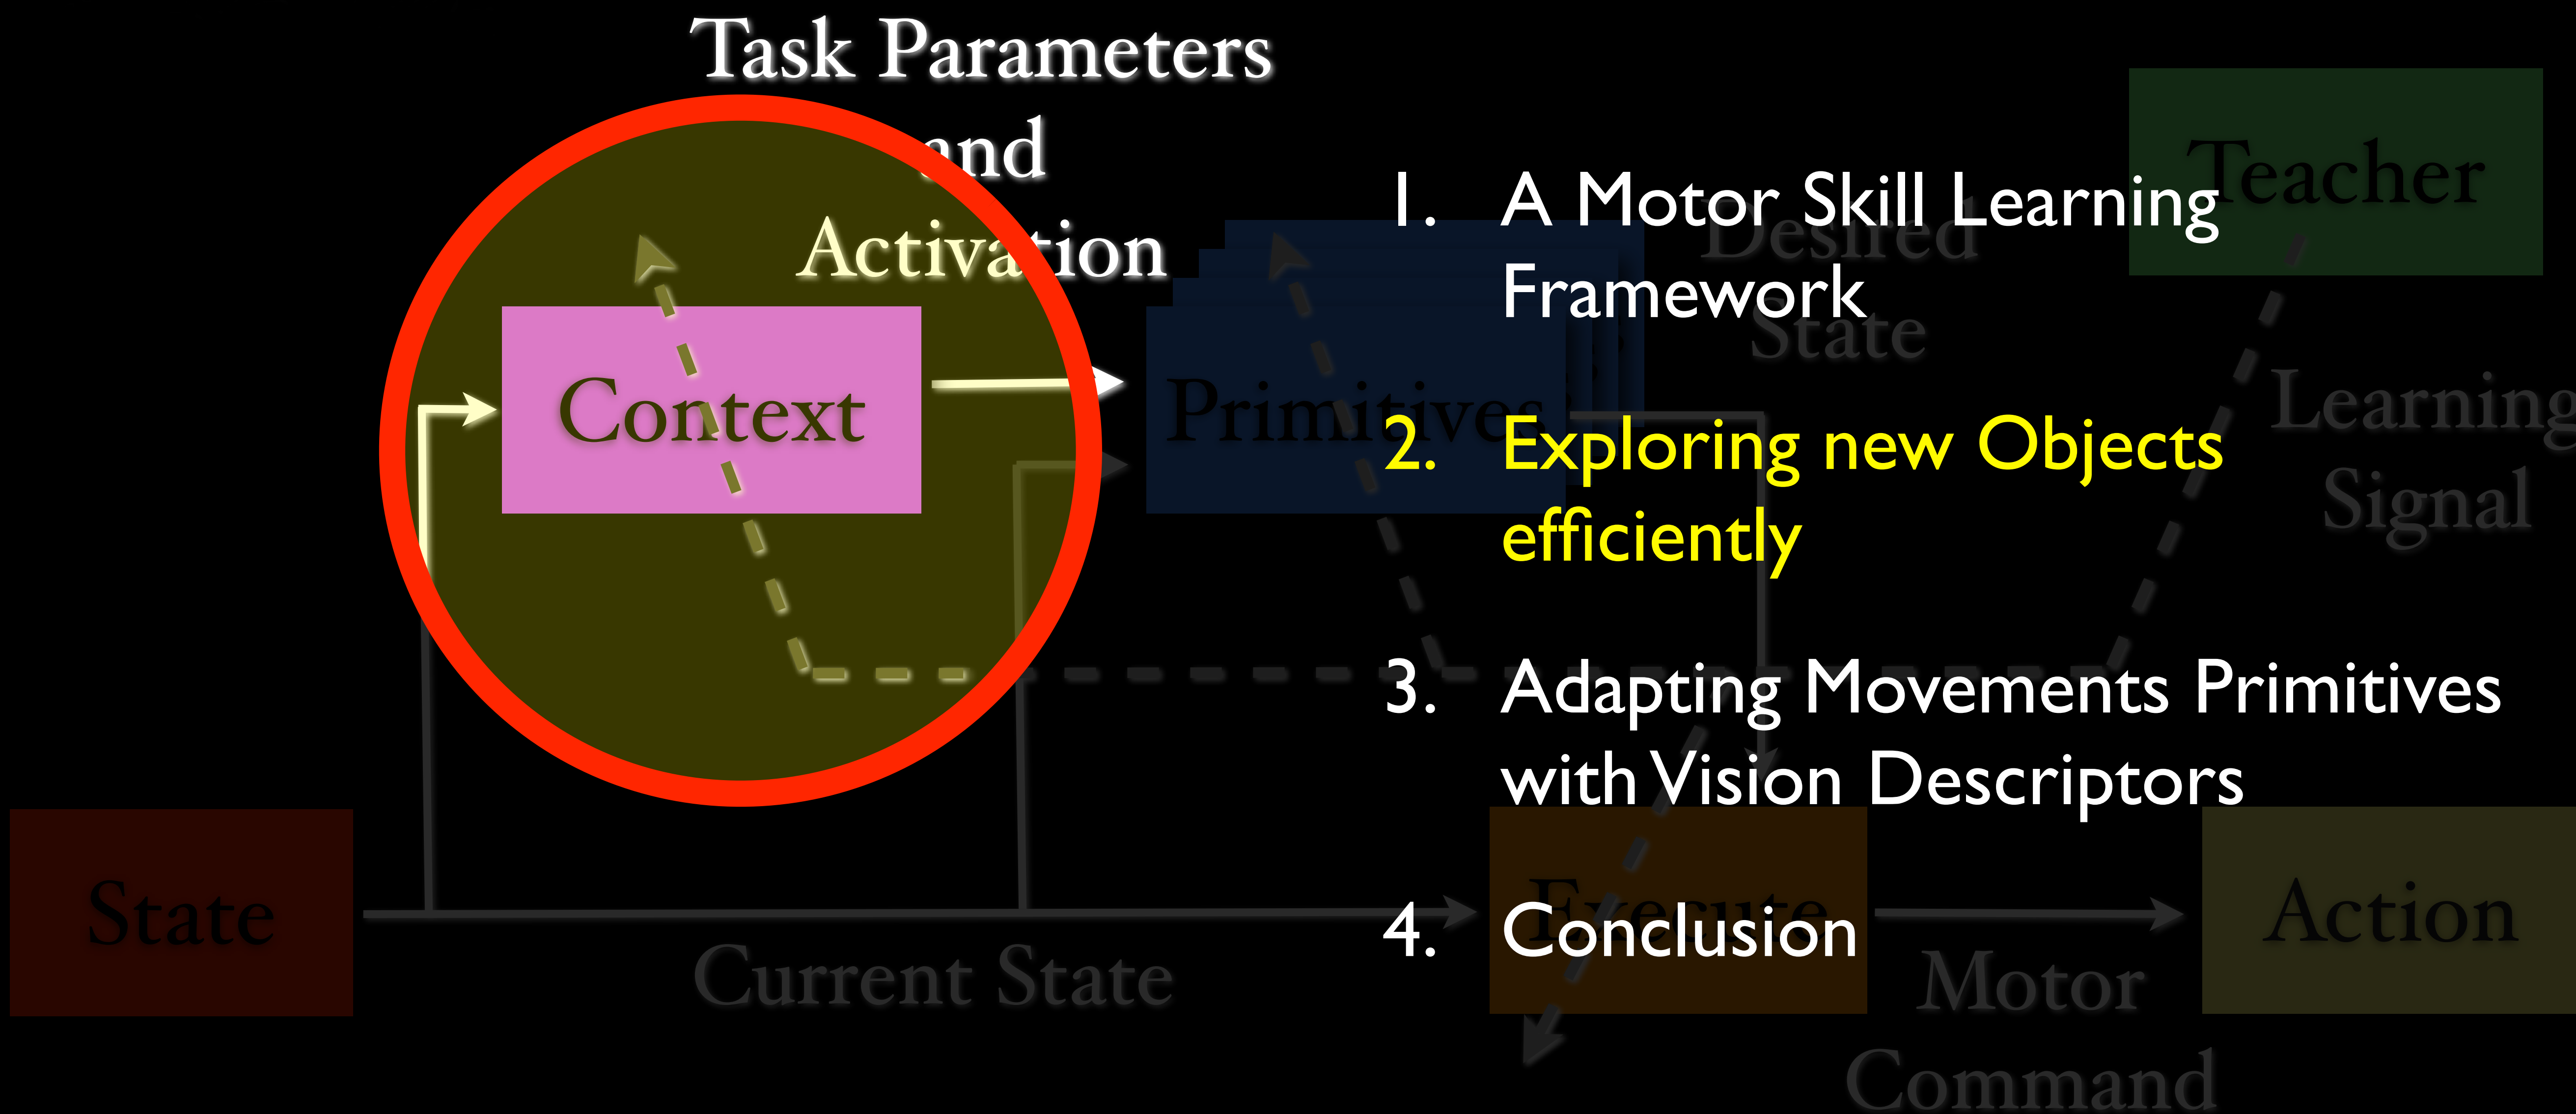

1. A Motor Skill Learning Framework
2. Exploring new Objects efficiently
3. Adapting Movements Primitives with Vision Descriptors
4. Conclusion

How to Explore Efficiently?

Observed
Object

Vision
Descriptors

Possible Grasps

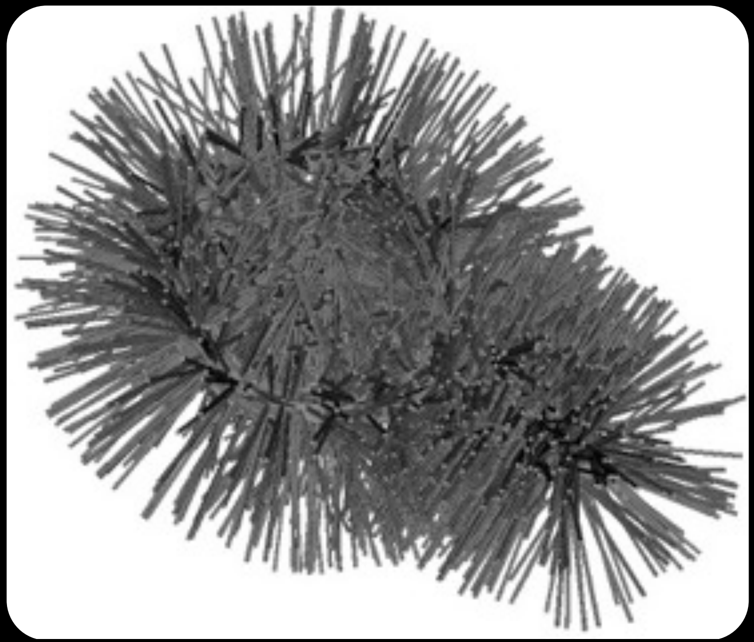

Successful Grasps

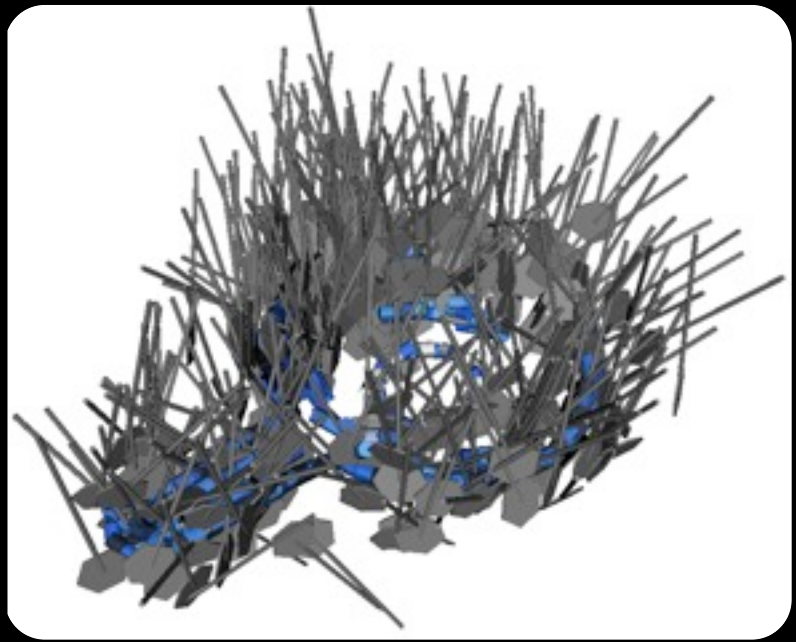

Detry, Baseski, Popovic, Touati, Krueger, Kroemer, Peters, Piater (2009).
Learning Object-specific Grasp Affordance Densities, ICDL

Grasping as a Bandit Problem?

Goal: find good grasps...

- ▪▪▪ fast
- ▪▪▪ by trial and error
- ▪▪▪ without a simulated model
- ▪▪▪ do not stick to grasps

Appropriate Approach:

- Upper confidence bound trades *exploration* & *exploitation*
- Gaussian Process Regression-based quality estimation
- Mean-shift inspired maxima detection

➡ New Efficient Algorithms!

Krömer, Detry, Piater, Peters (submitted). Combining Active Learning and Reactive Control for Robot Grasping, Robotics and Autonomous Systems

Modeling Success

Observed Data

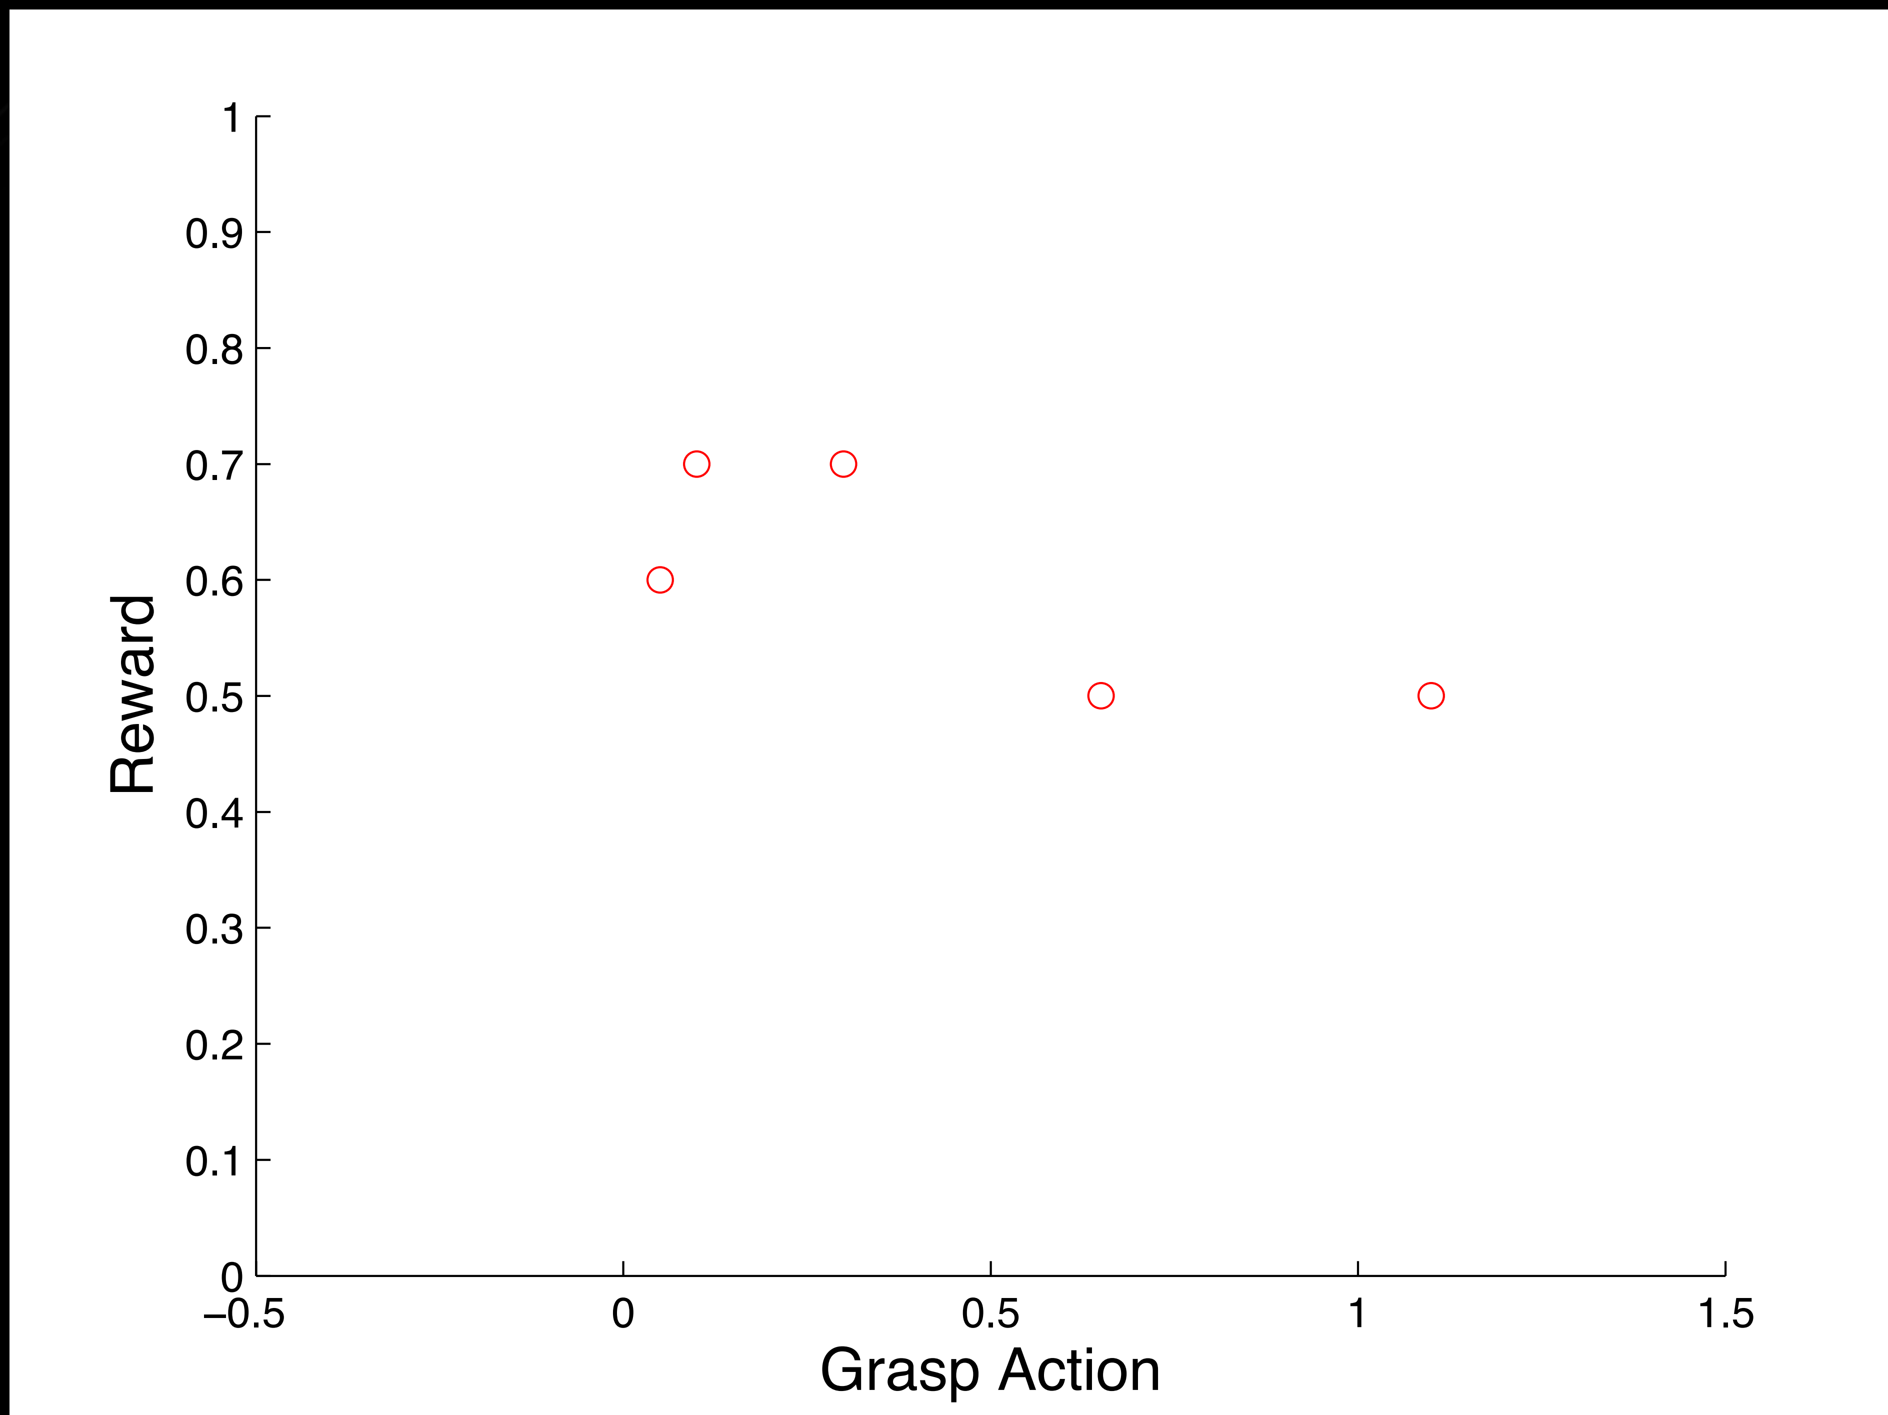

Krömer, Detry, Piater, Peters (submitted). Combining Active Learning and Reactive Control for Robot Grasping, Robotics and Autonomous Systems

Modeling Success

Gaussian Process Model

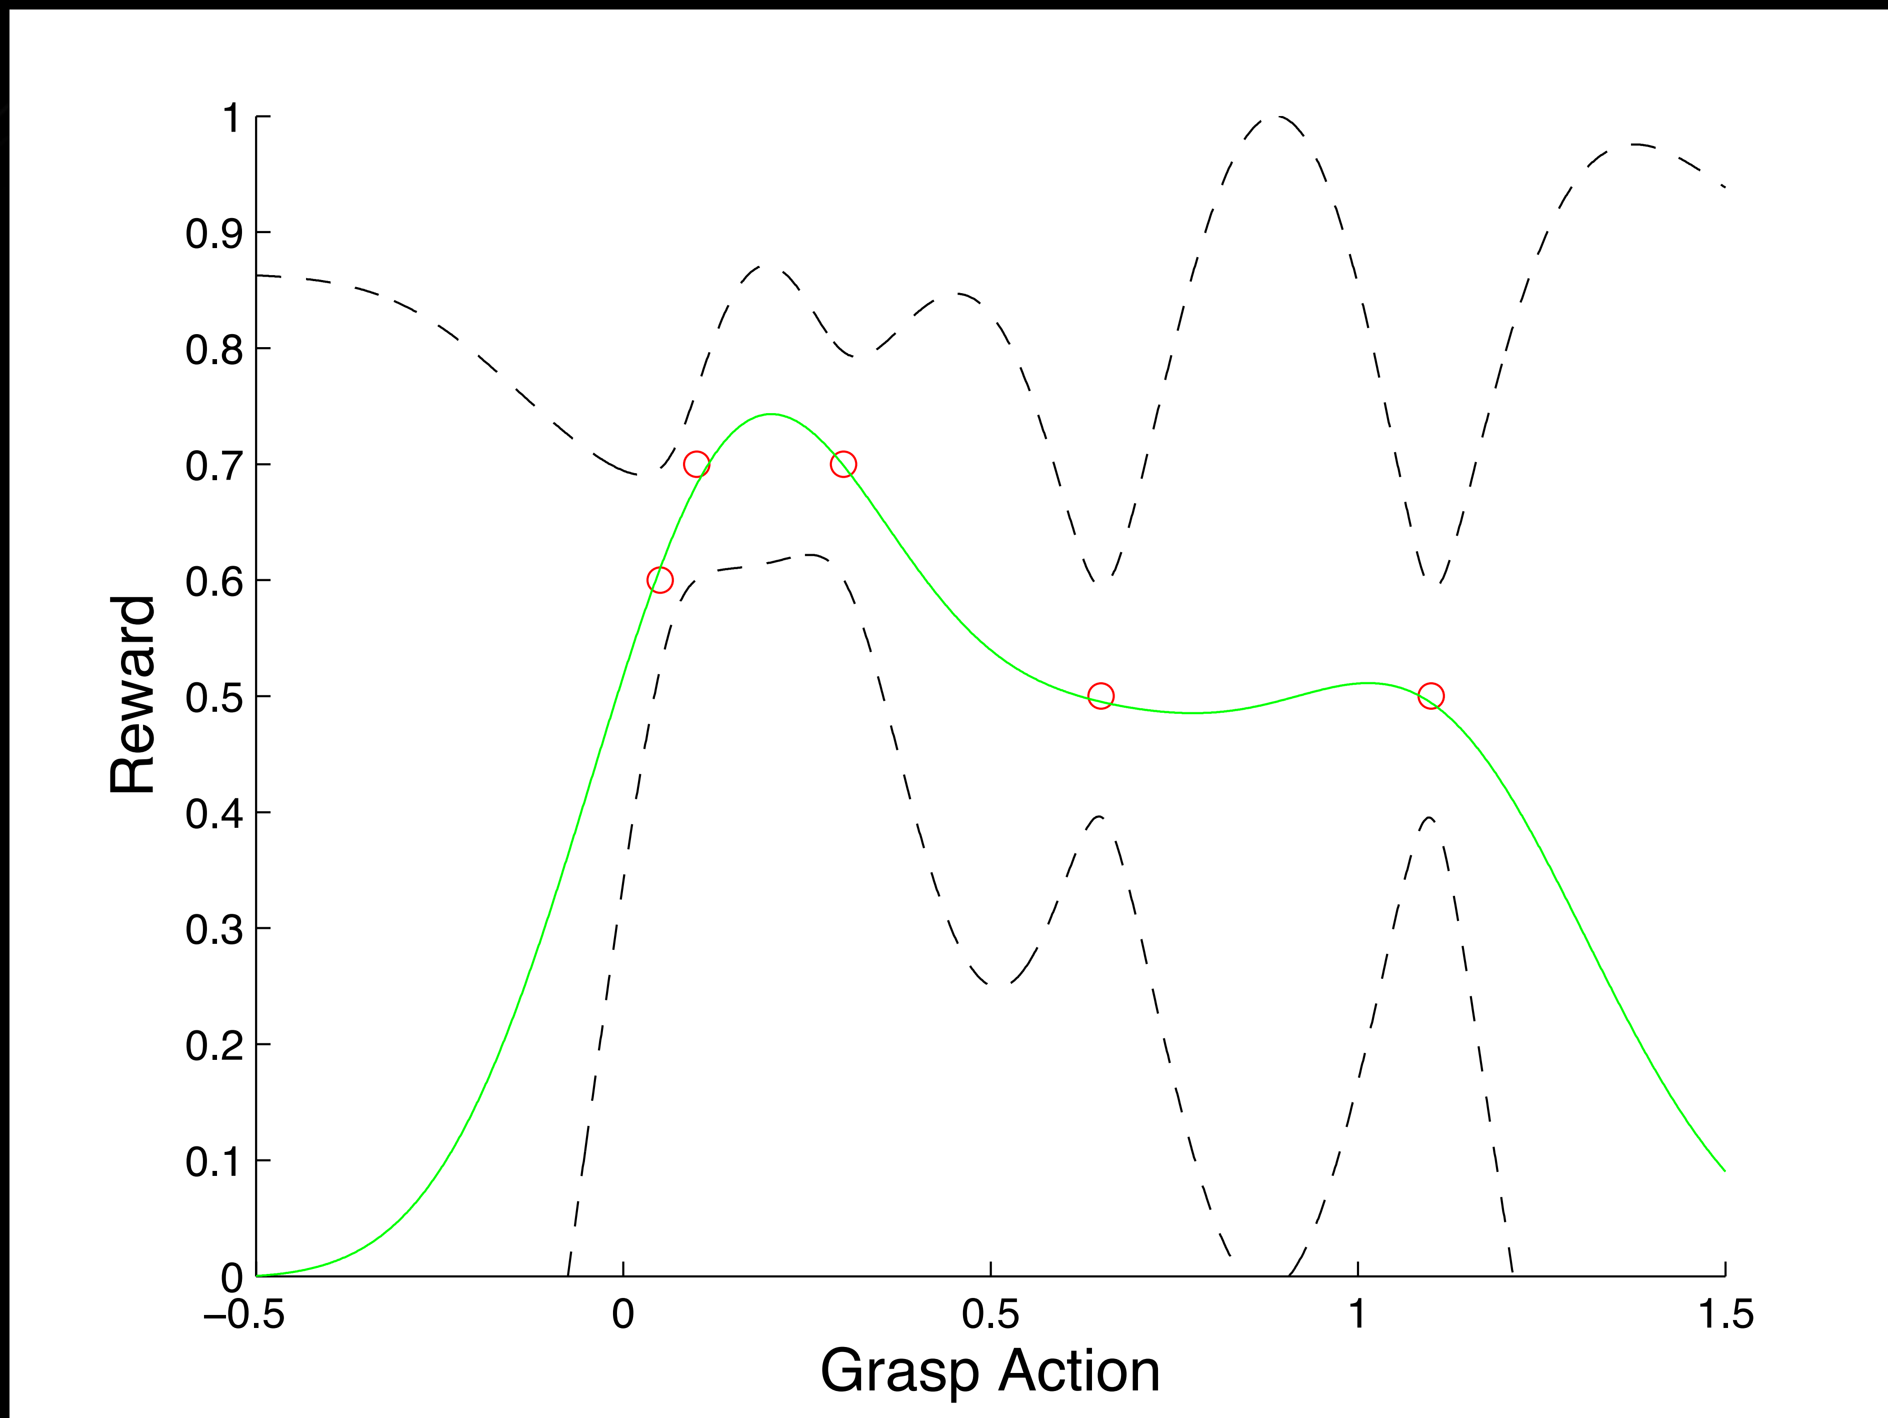

Krömer, Detry, Piater, Peters (submitted). Combining Active Learning and Reactive Control for Robot Grasping, Robotics and Autonomous Systems

Modeling Success

UCB Merit Function Model

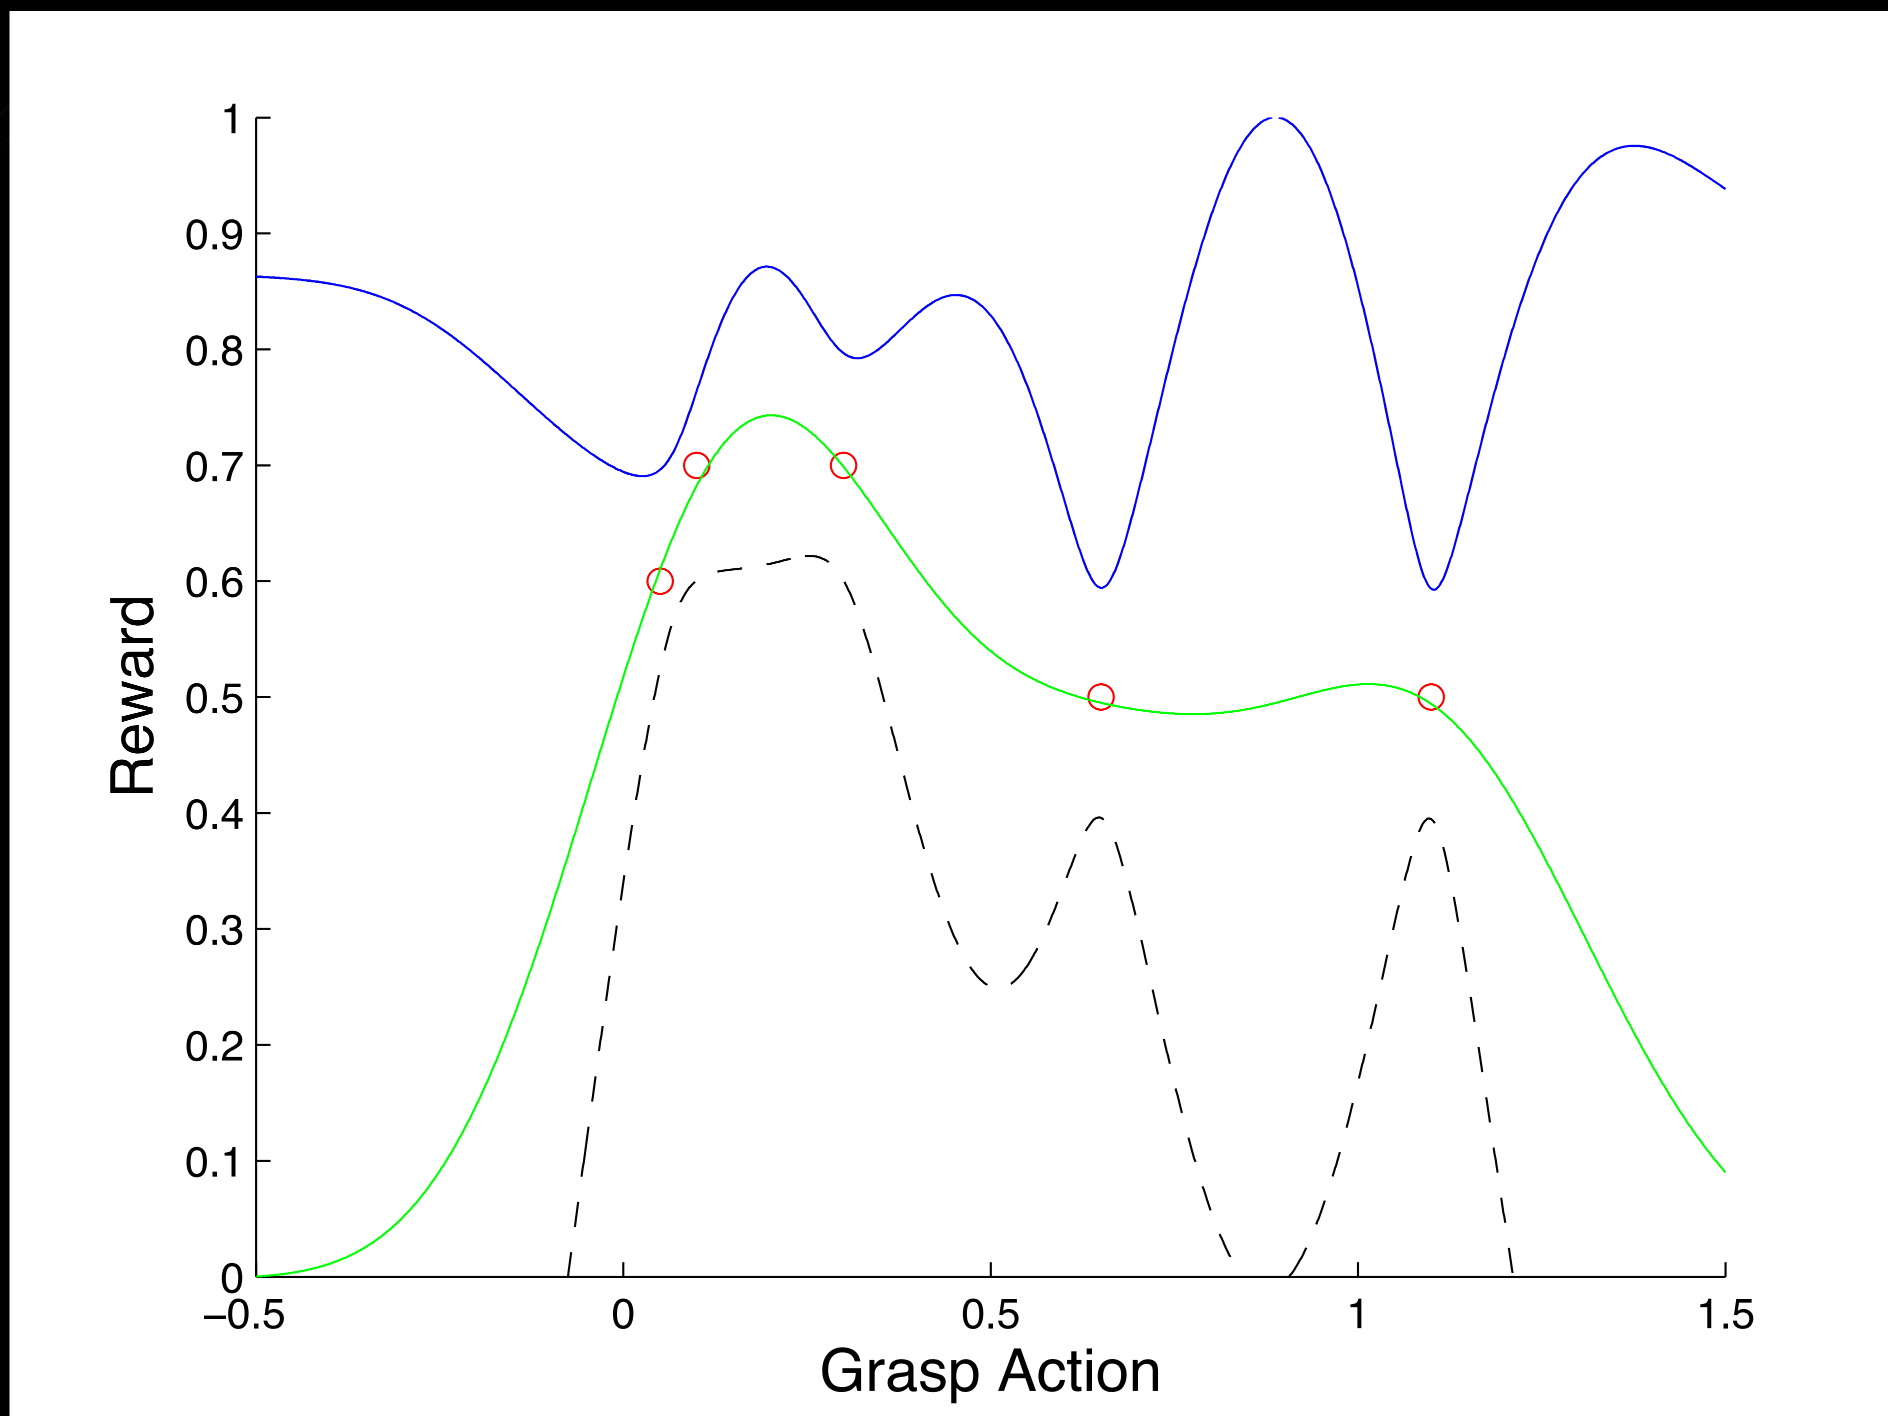

Krömer, Detry, Piater, Peters (submitted). Combining Active Learning and Reactive Control for Robot Grasping, Robotics and Autonomous Systems

Choosing the Next Action

Insights

- There may be infinitely many maxima (=grasps).
- The maxima are guaranteed to be near the data points
- Find *most* local maxima close to the data points!

▸ Back to our example...

Krömer, Detry, Piater, Peters (submitted). Combining Active Learning and Reactive Control for Robot Grasping, Robotics and Autonomous Systems

Detecting Maxima

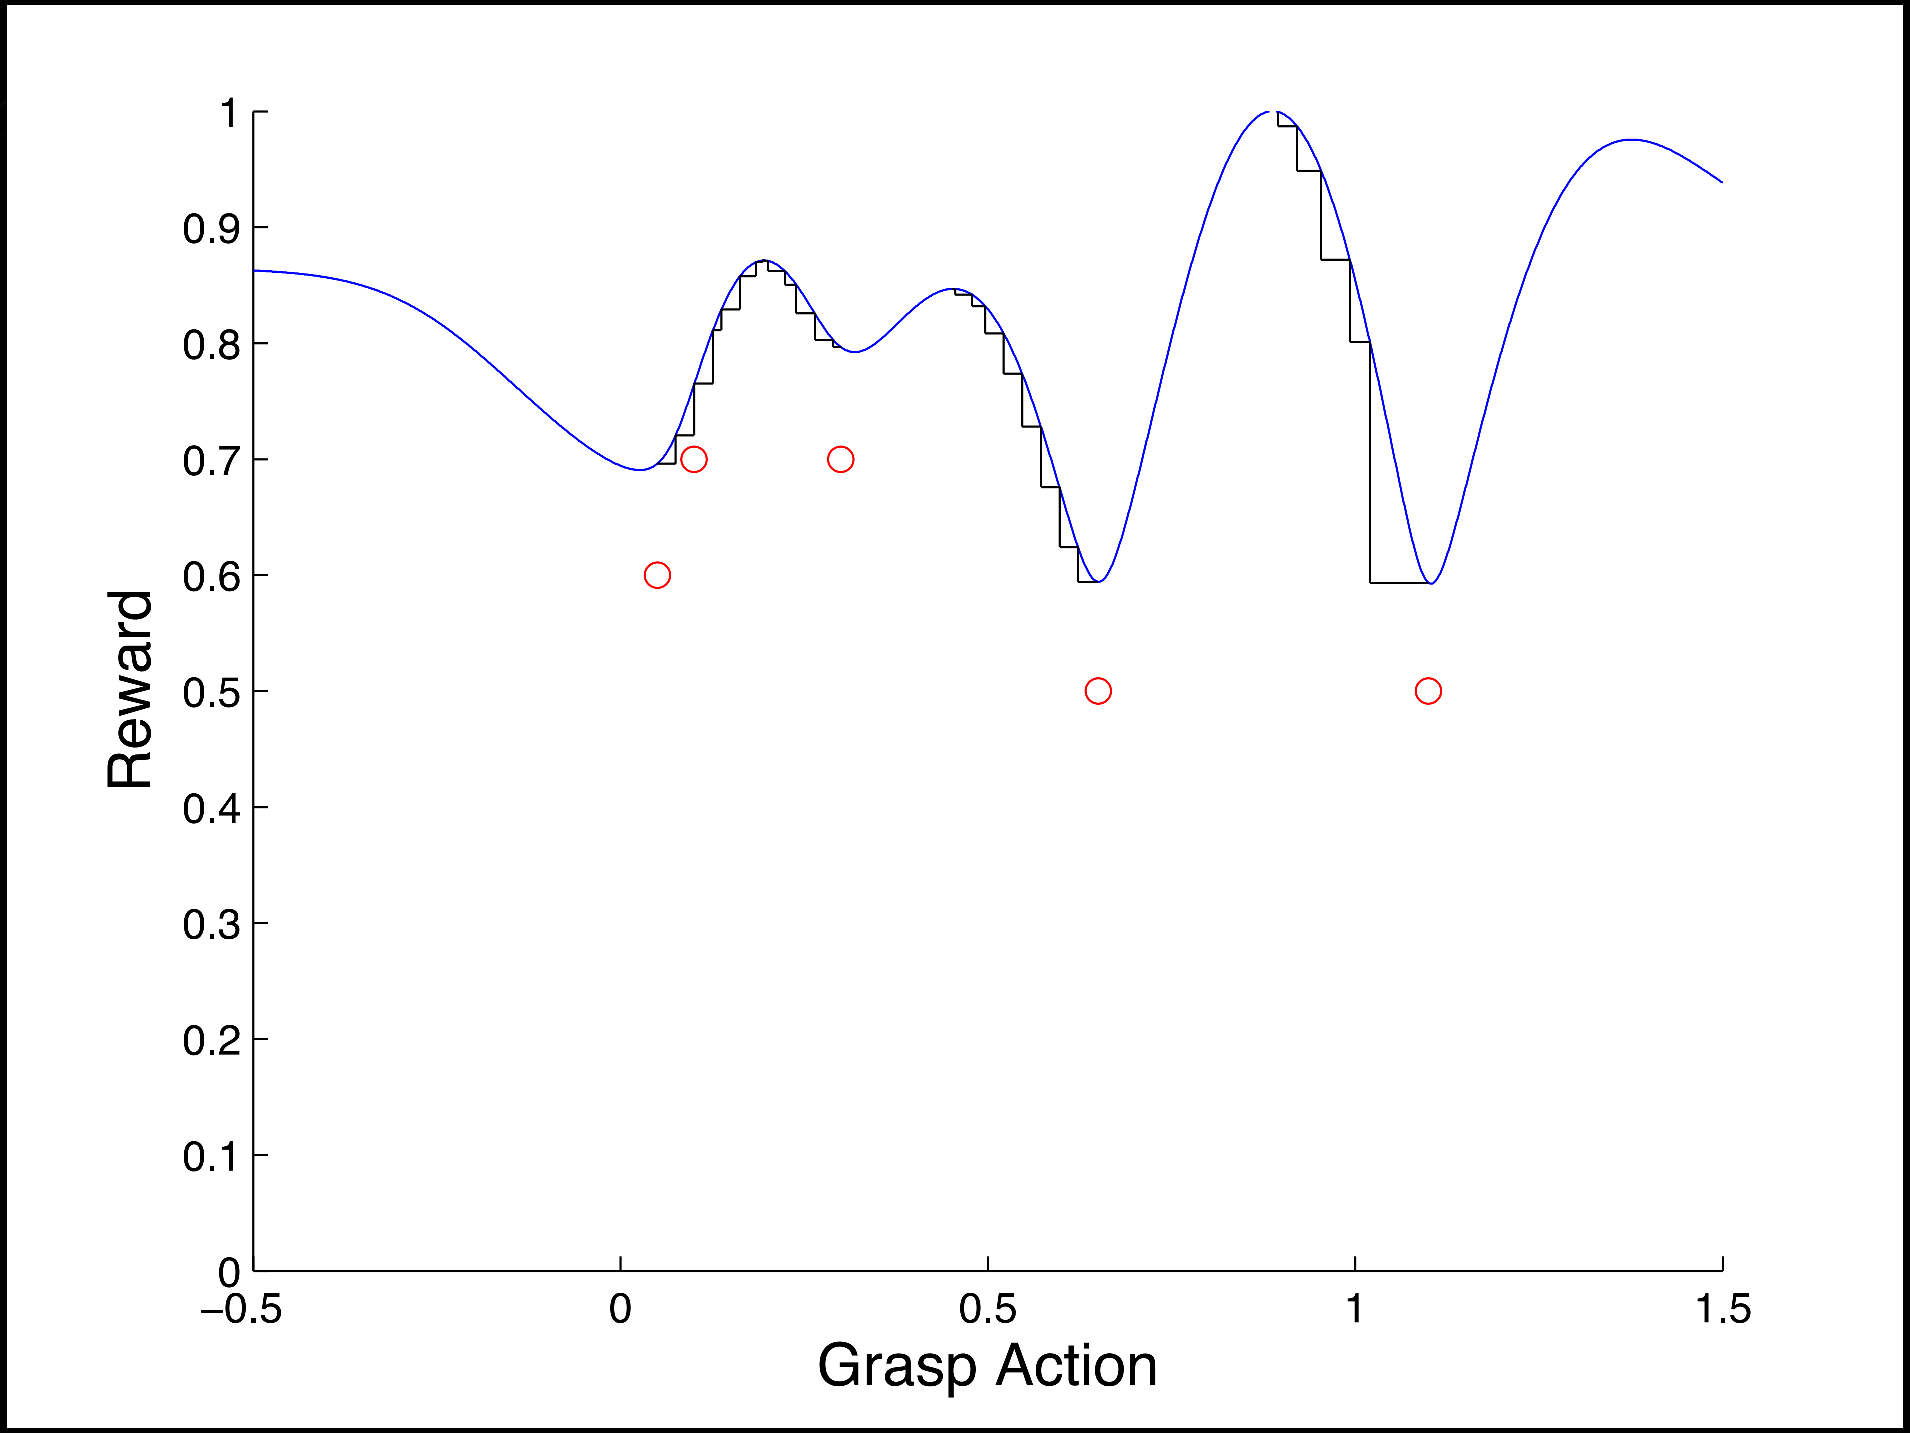

Krömer, Detry, Piater, Peters (submitted). Combining Active Learning and Reactive Control for Robot Grasping, Robotics and Autonomous Systems

Detecting Maxima

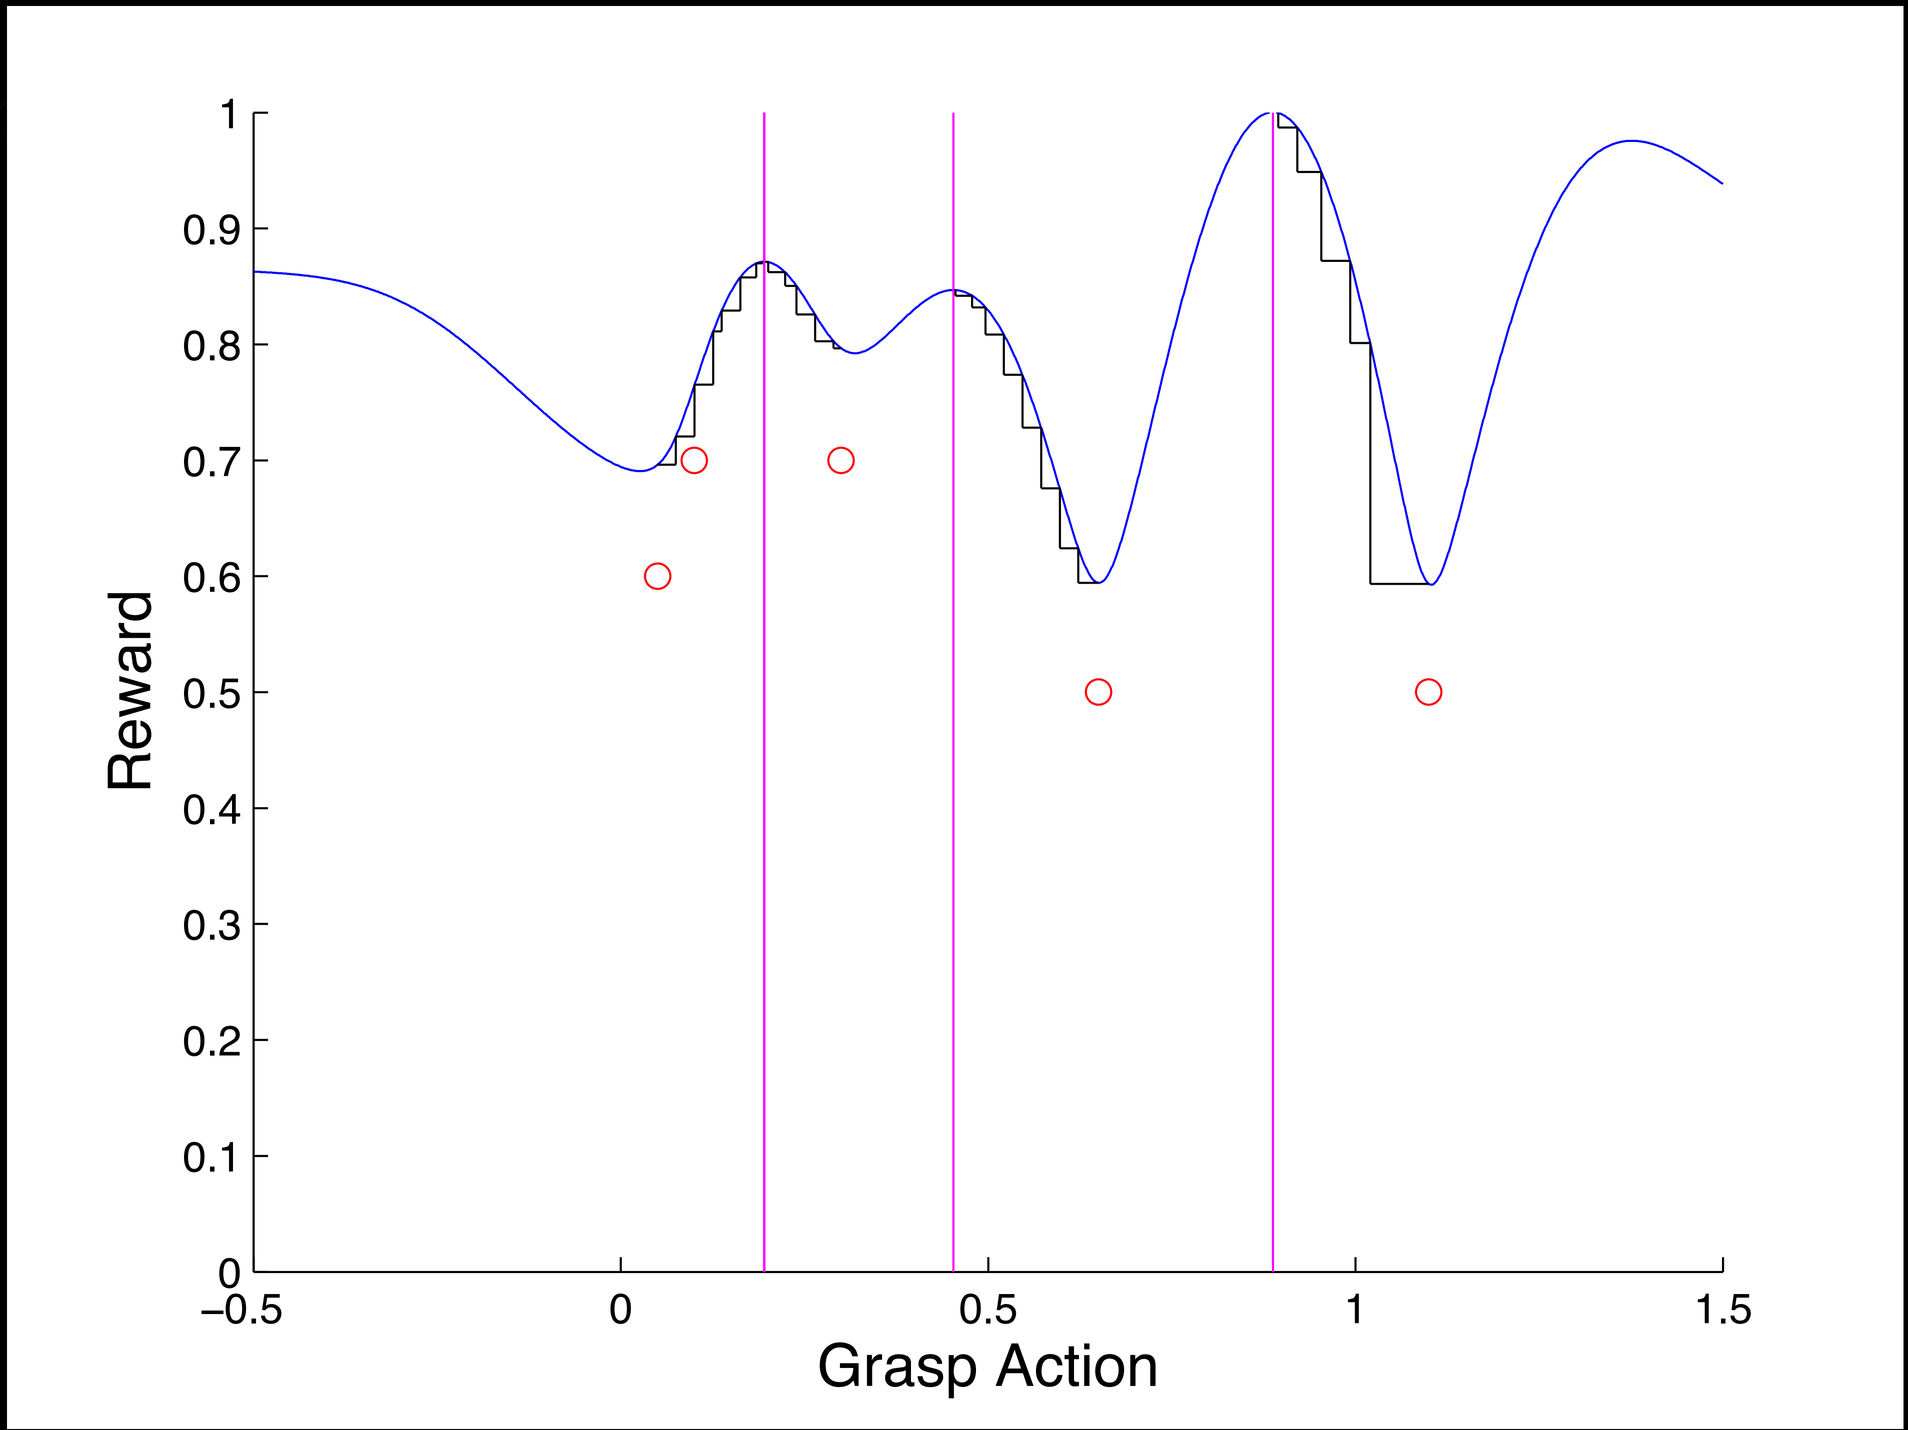

Krömer, Detry, Piater, Peters (submitted). Combining Active Learning and Reactive Control for Robot Grasping, Robotics and Autonomous Systems

Evaluate Candidate

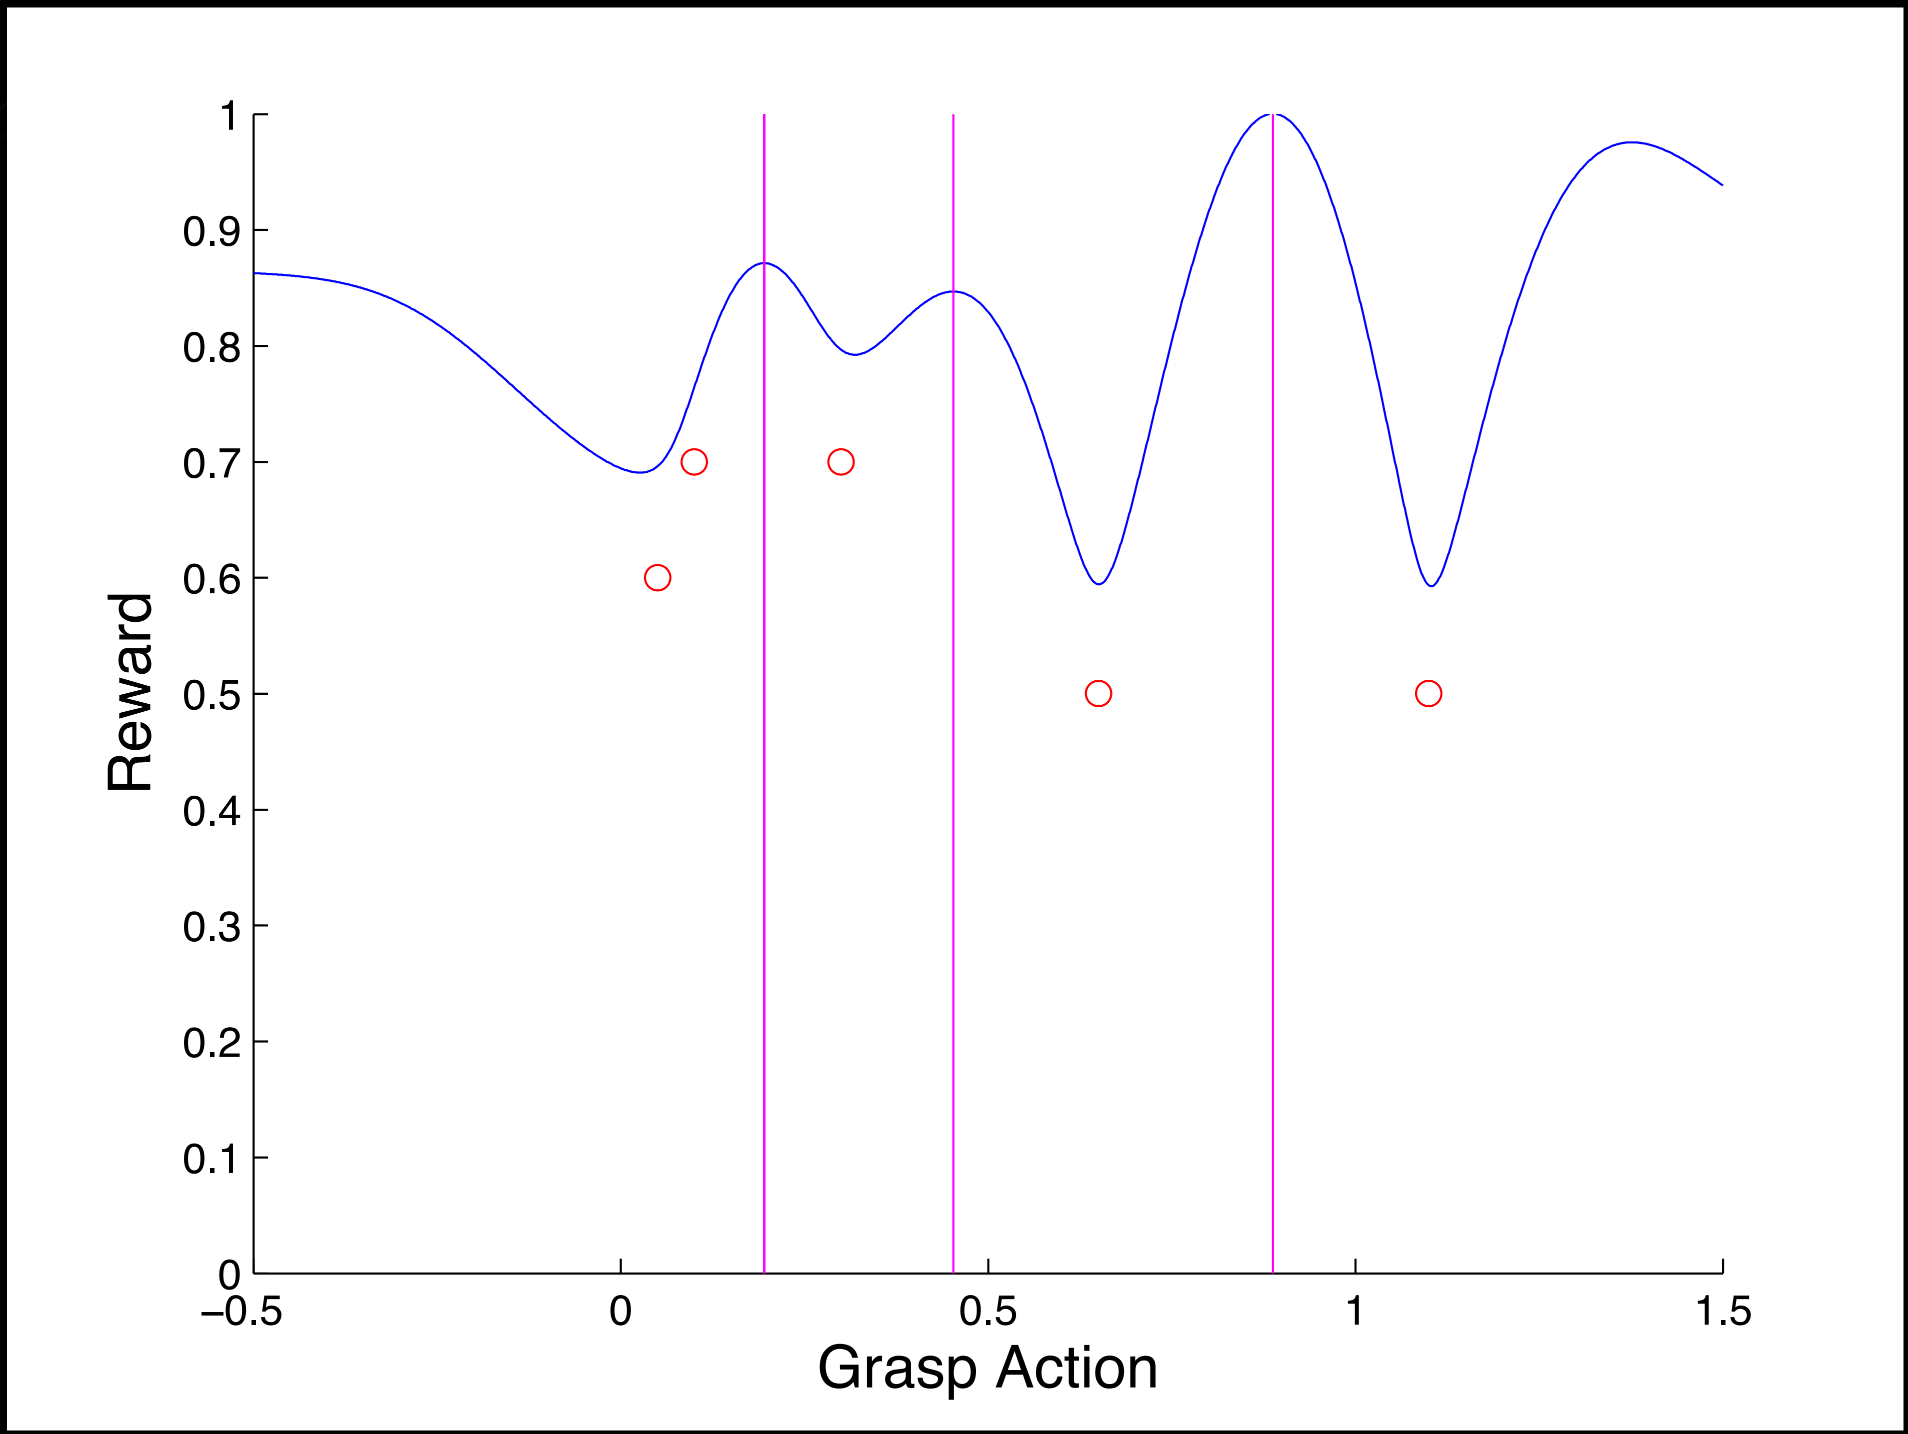

Krömer, Detry, Piater, Peters (submitted). Combining Active Learning and Reactive Control for Robot Grasping, Robotics and Autonomous Systems

Another attempt...

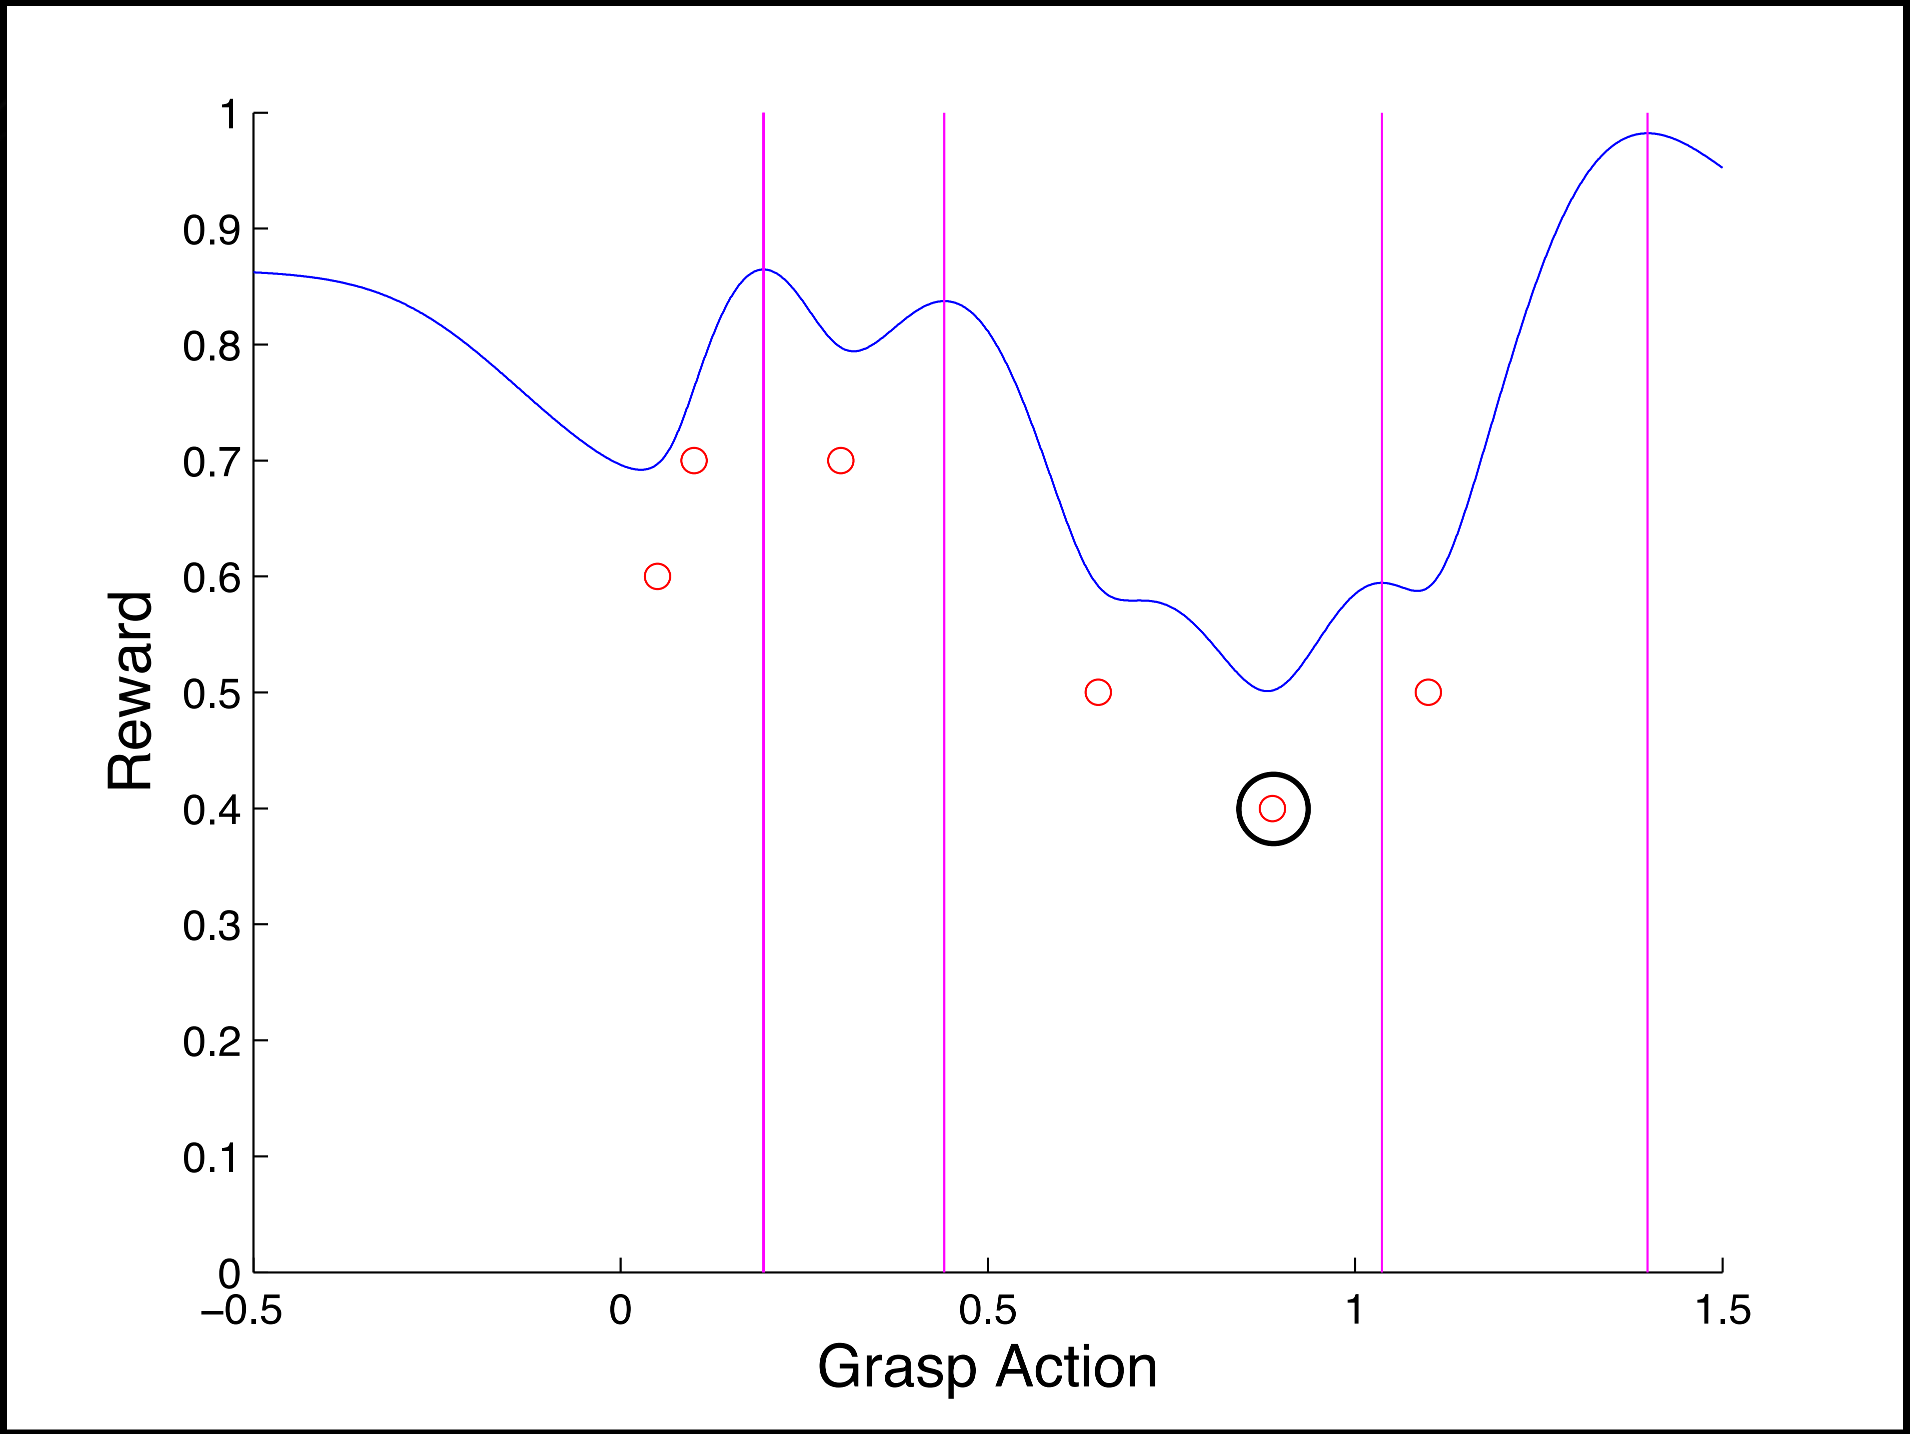

Krömer, Detry, Piater, Peters (submitted). Combining Active Learning and Reactive Control for Robot Grasping, Robotics and Autonomous Systems

Another attempt...

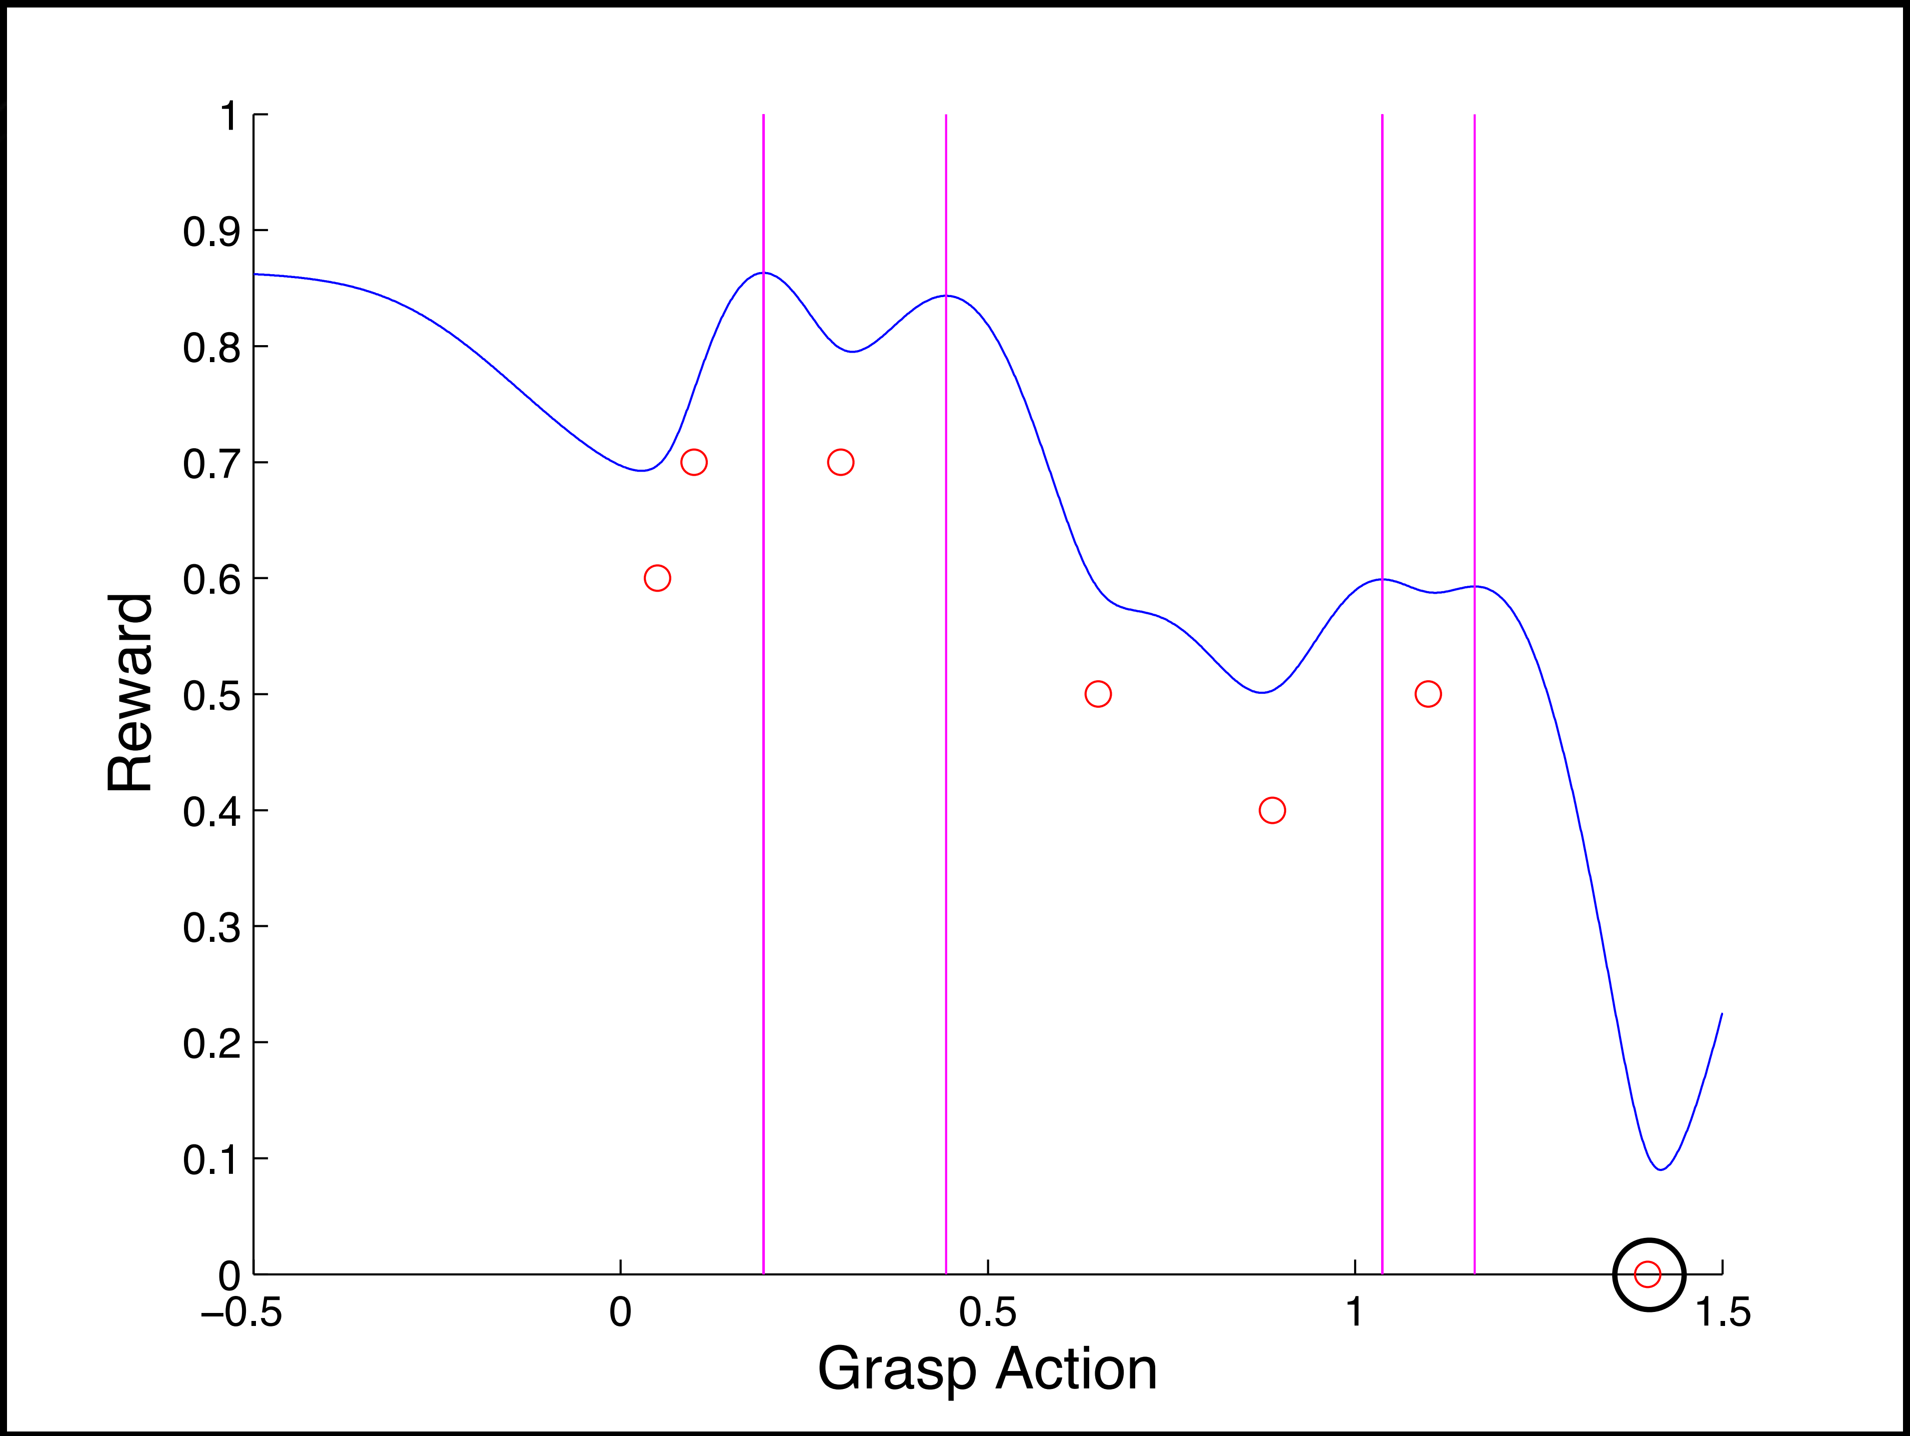

Krömer, Detry, Piater, Peters (submitted). Combining Active Learning and Reactive Control for Robot Grasping, Robotics and Autonomous Systems

Performance

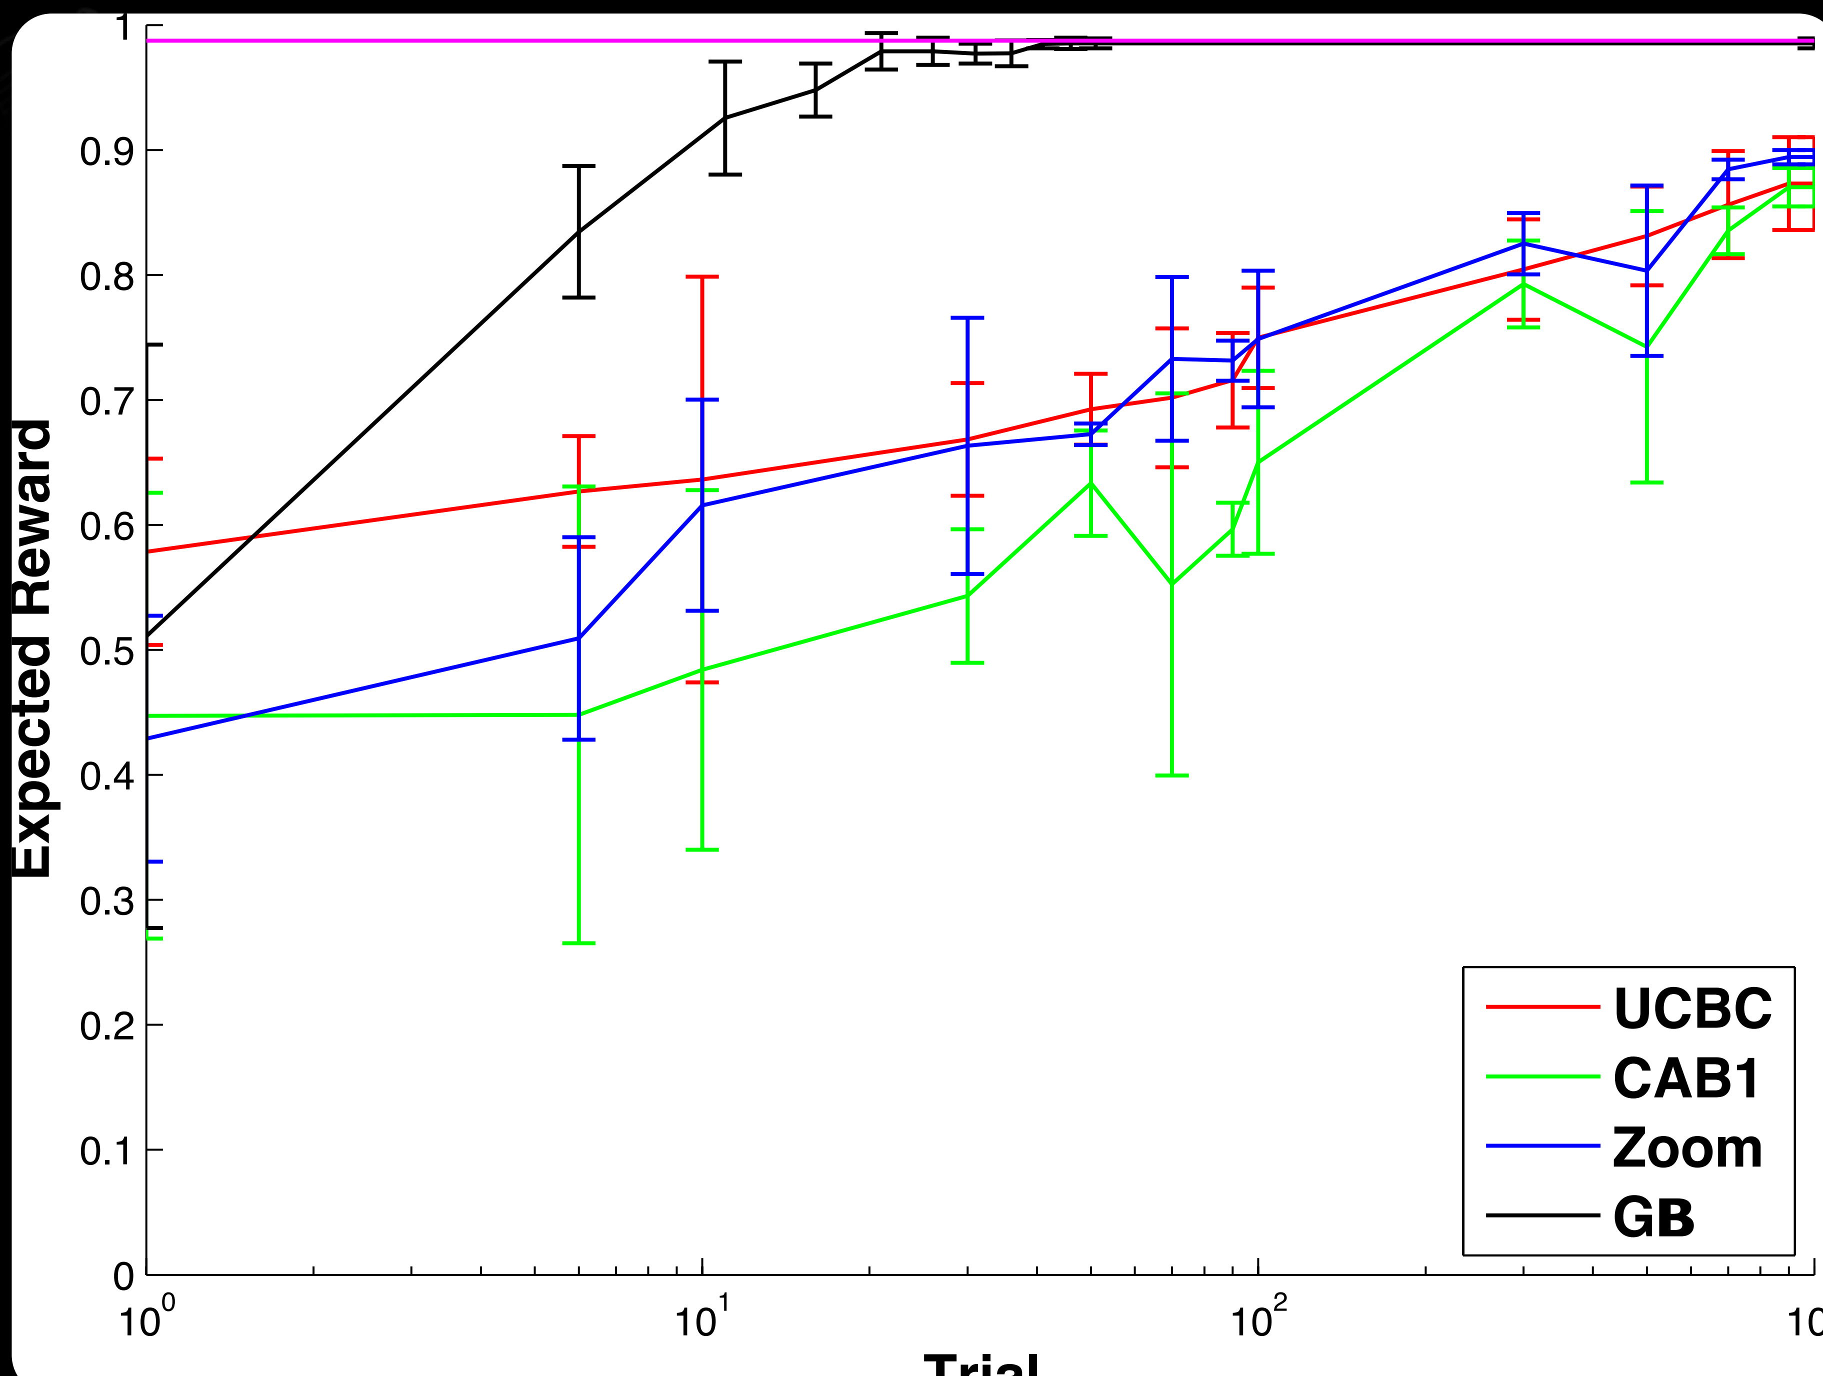

0.6419
0.4987
0.6065
0.9122

Krömer, Detry, Piater, Peters (submitted). Combining Active Learning and Reactive Control for Robot Grasping, Robotics and Autonomous Systems

Rewards...

Krömer, Detry, Piater, Peters (submitted). Combining Active Learning and Reactive Control for Robot Grasping, Robotics and Autonomous Systems

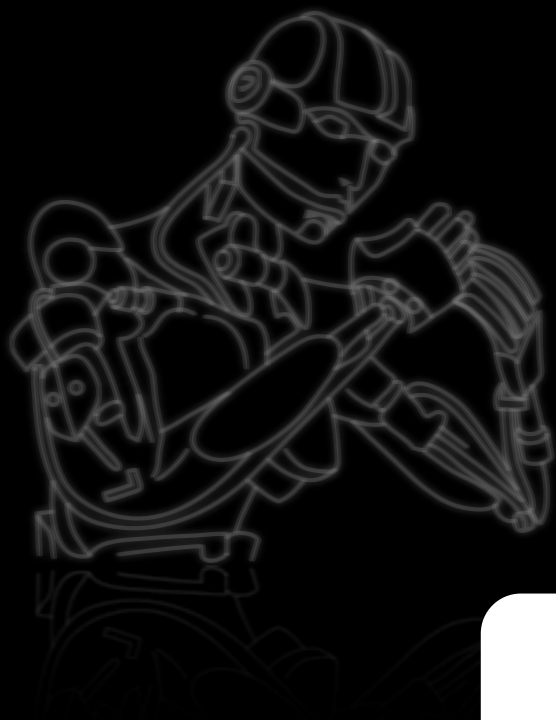

Learning Performance

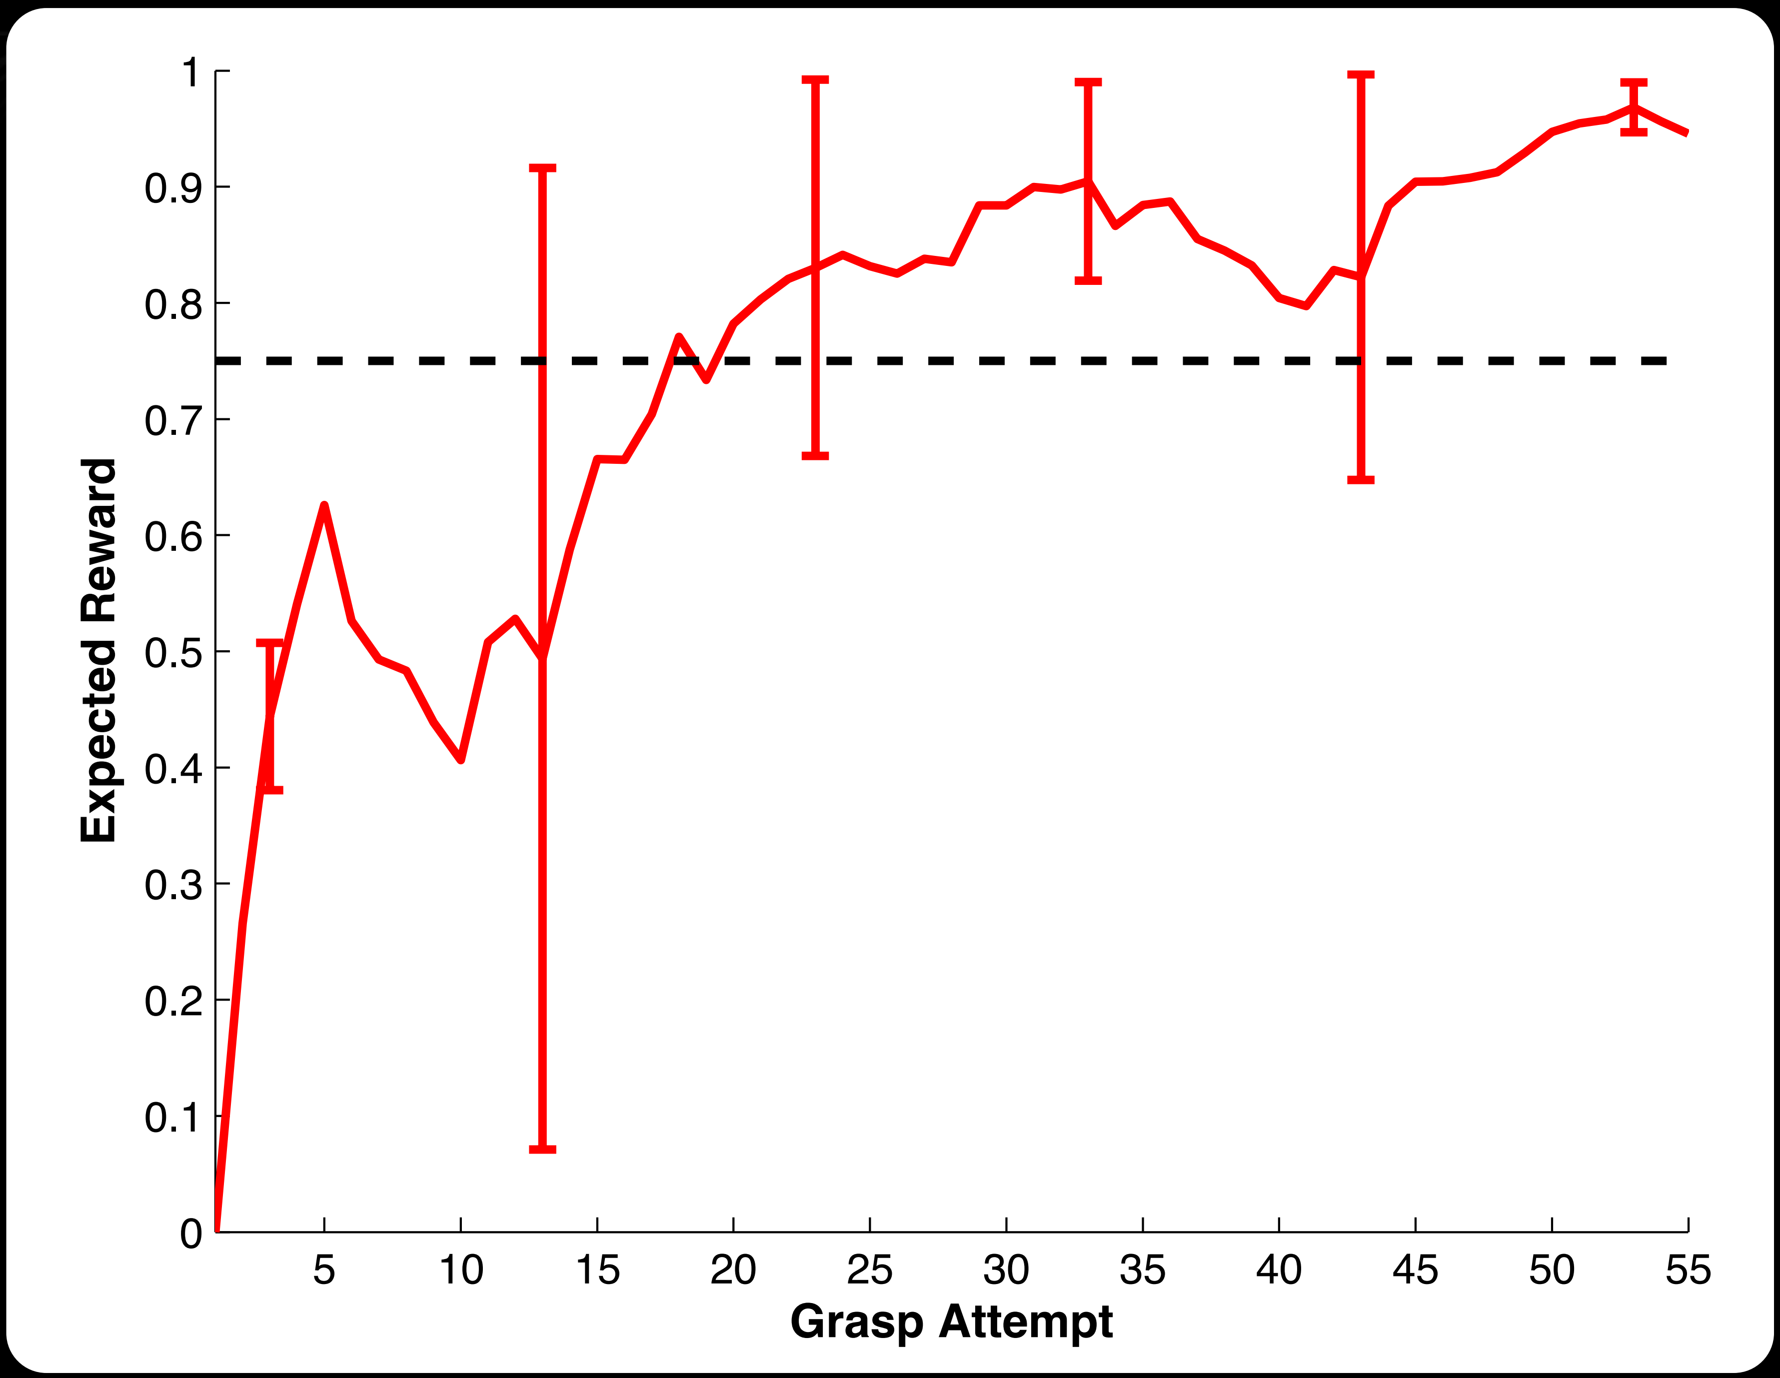

Krömer, Detry, Piater, Peters (submitted). Combining Active Learning and Reactive Control for Robot Grasping, Robotics and Autonomous Systems

Newly found grasps...

Krömer, Detry, Piater, Peters (submitted). Combining Active Learning and Reactive Control for Robot Grasping, Robotics and Autonomous Systems

Outline

1. A Motor Skill Learning Framework
2. Exploring efficiently
3. Adapting Movements Primitives with Vision Descriptors
4. Conclusion

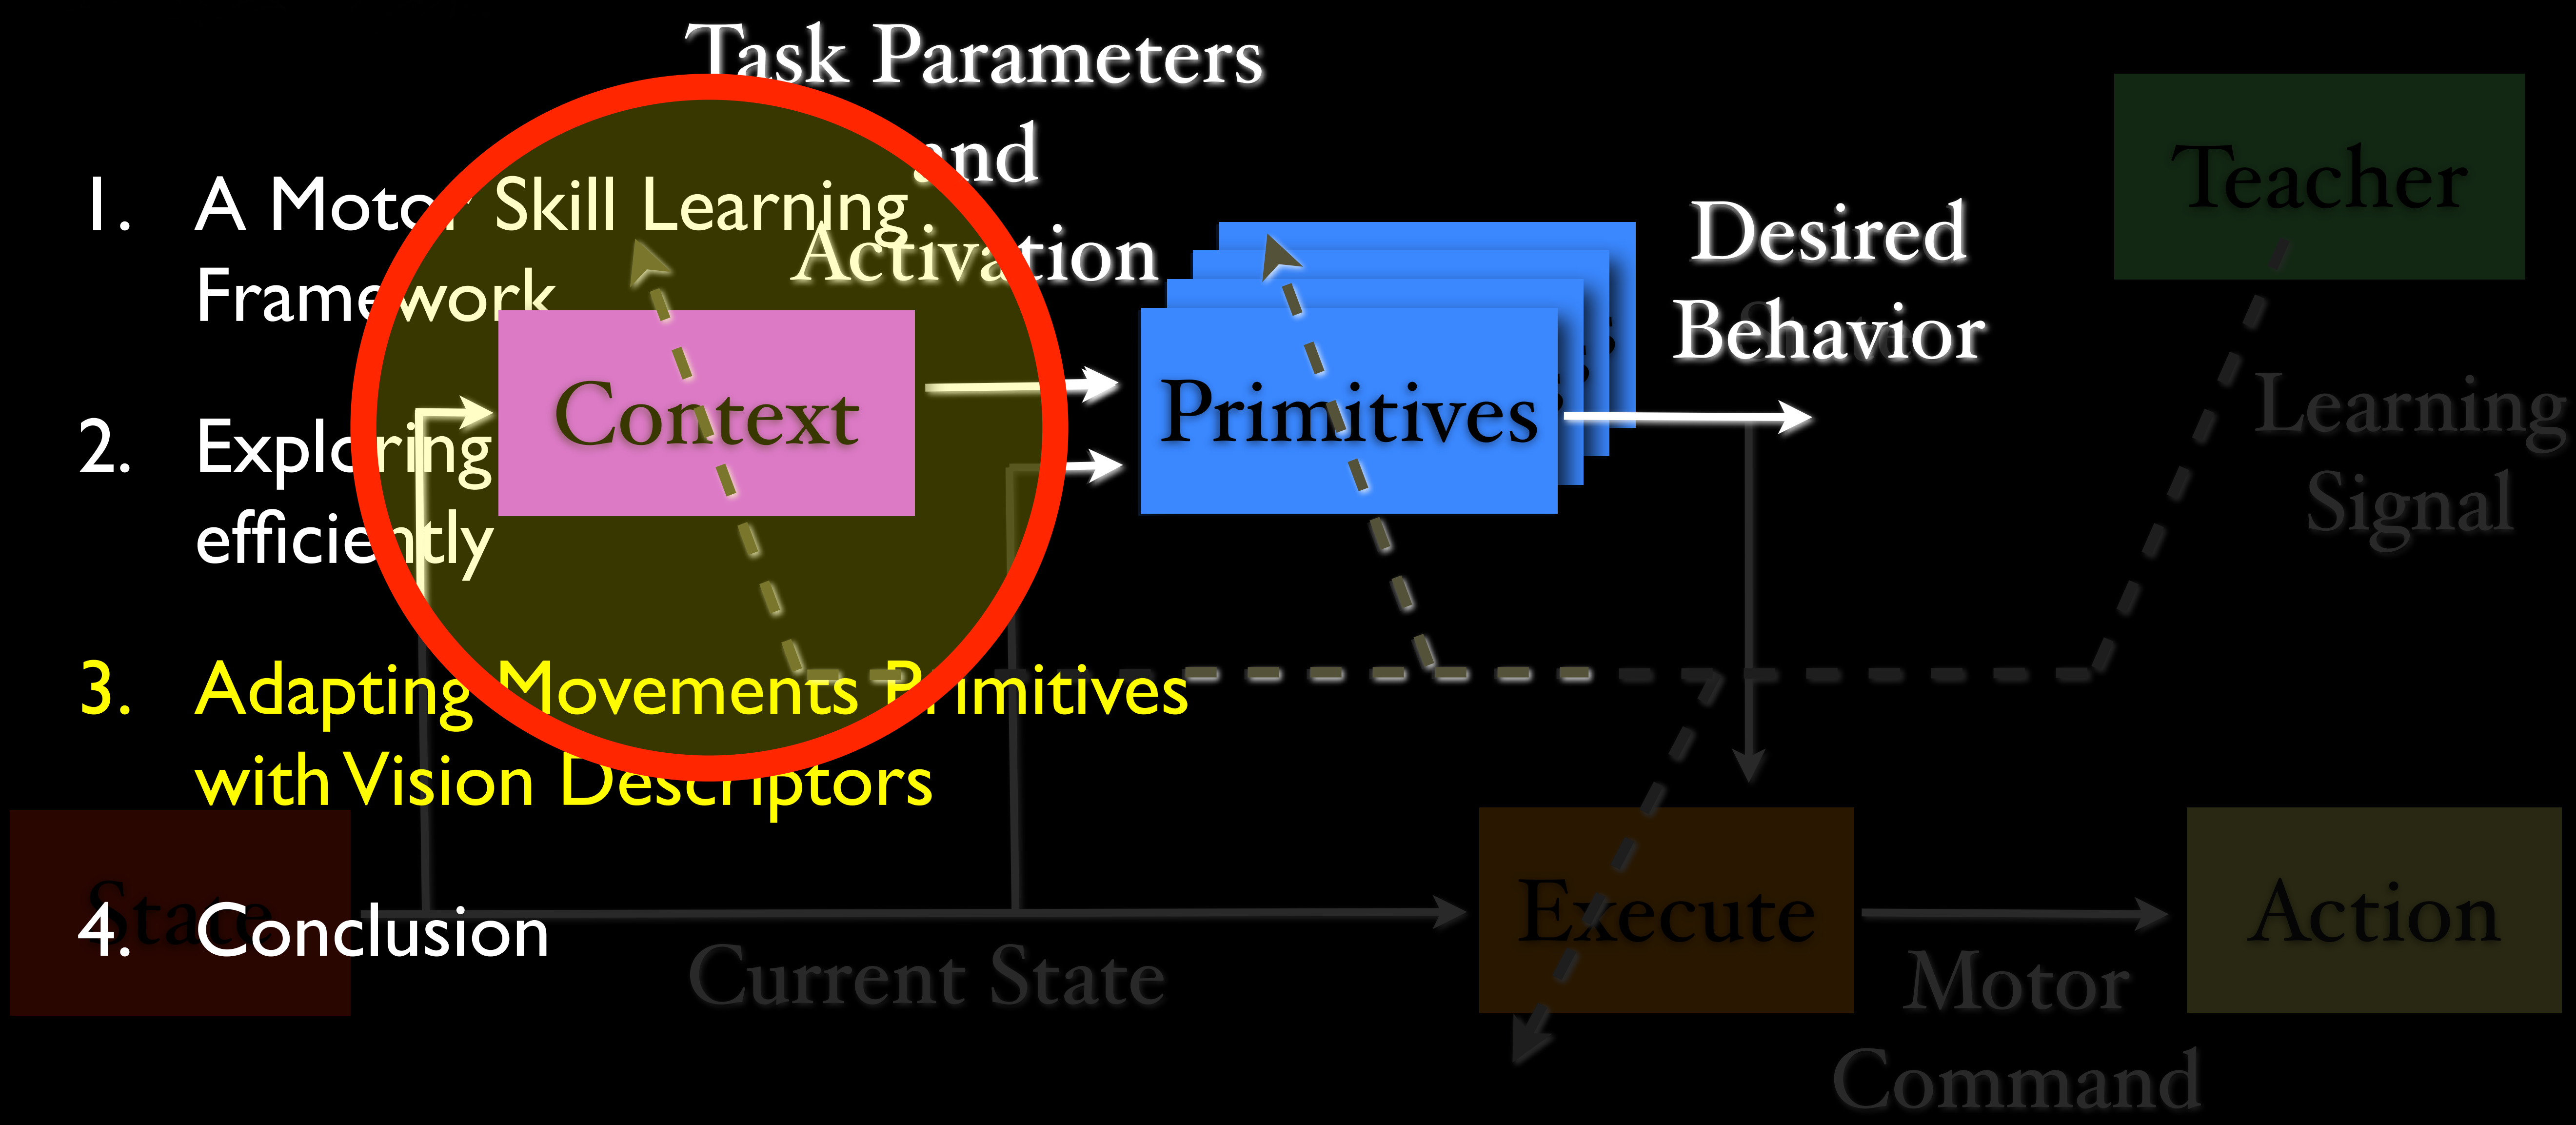

Vision Descriptors

Scene

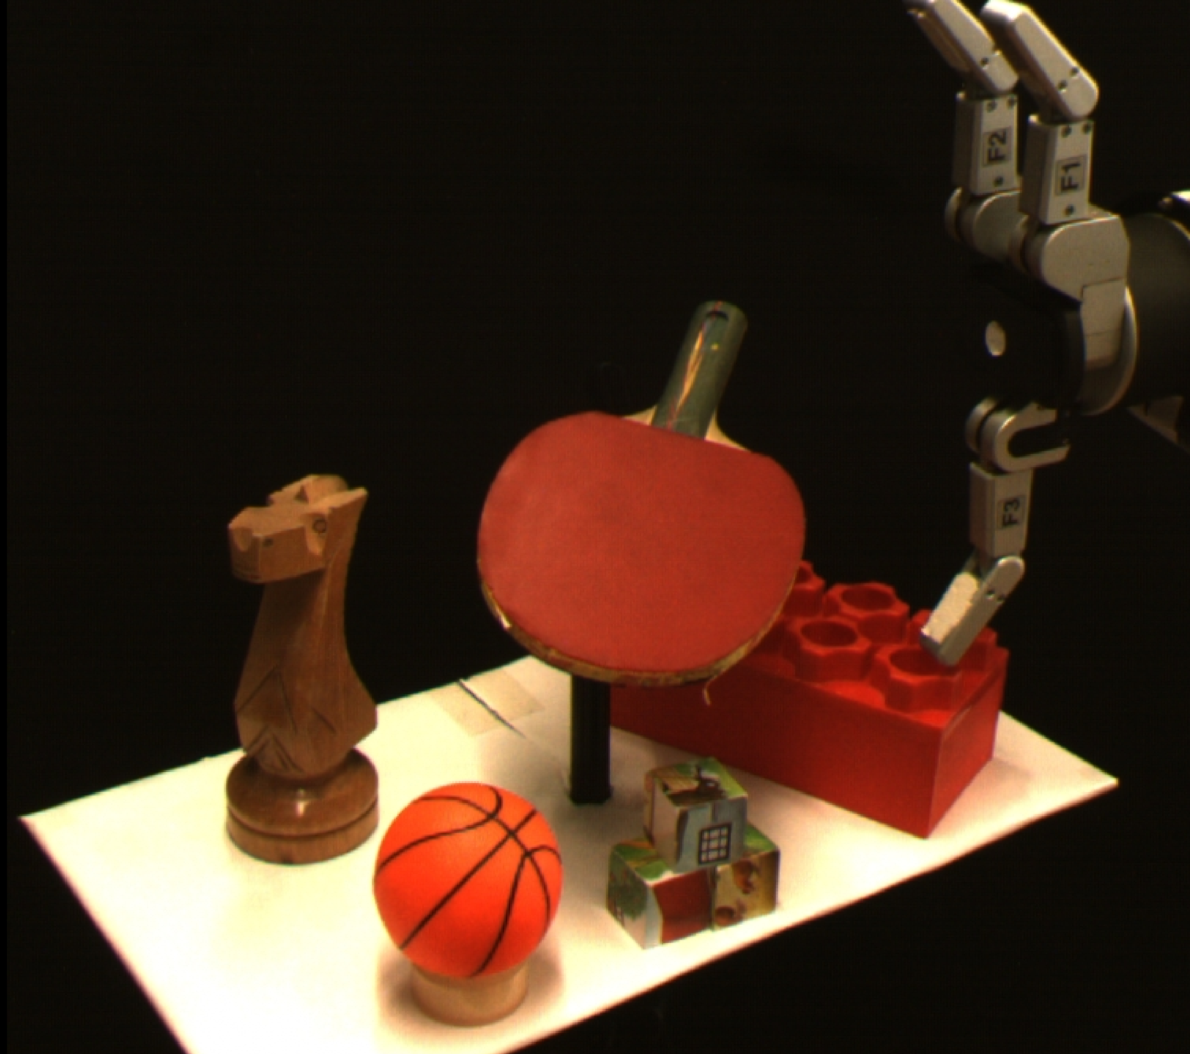

Vision Descriptor Representation

N. Krüger, M. Lappe, F. Wörgötter. (2004) Biologically Motivated MultimodalProcessing of Visual Primitives. J.AI&SB.

Basic Insights

Motor primitives (DMP) allow...

- Initialization by demonstrations
- Acquire benefits of human motions: regular, smooth motions, small overshoots, etc. (Jeannerod, 1996)

Local scene geometry in form of vision descriptors allow...

- incorporating proximity to objects
- preshaping the hand to the object
- avoid obstacles

Krömer, Detry, Piater, Peters (submitted). Combining Active Learning and Reactive Control for Robot Grasping, Robotics and Autonomous Systems

Generalizing with DMPs

Imitation allows reproducing observed movements

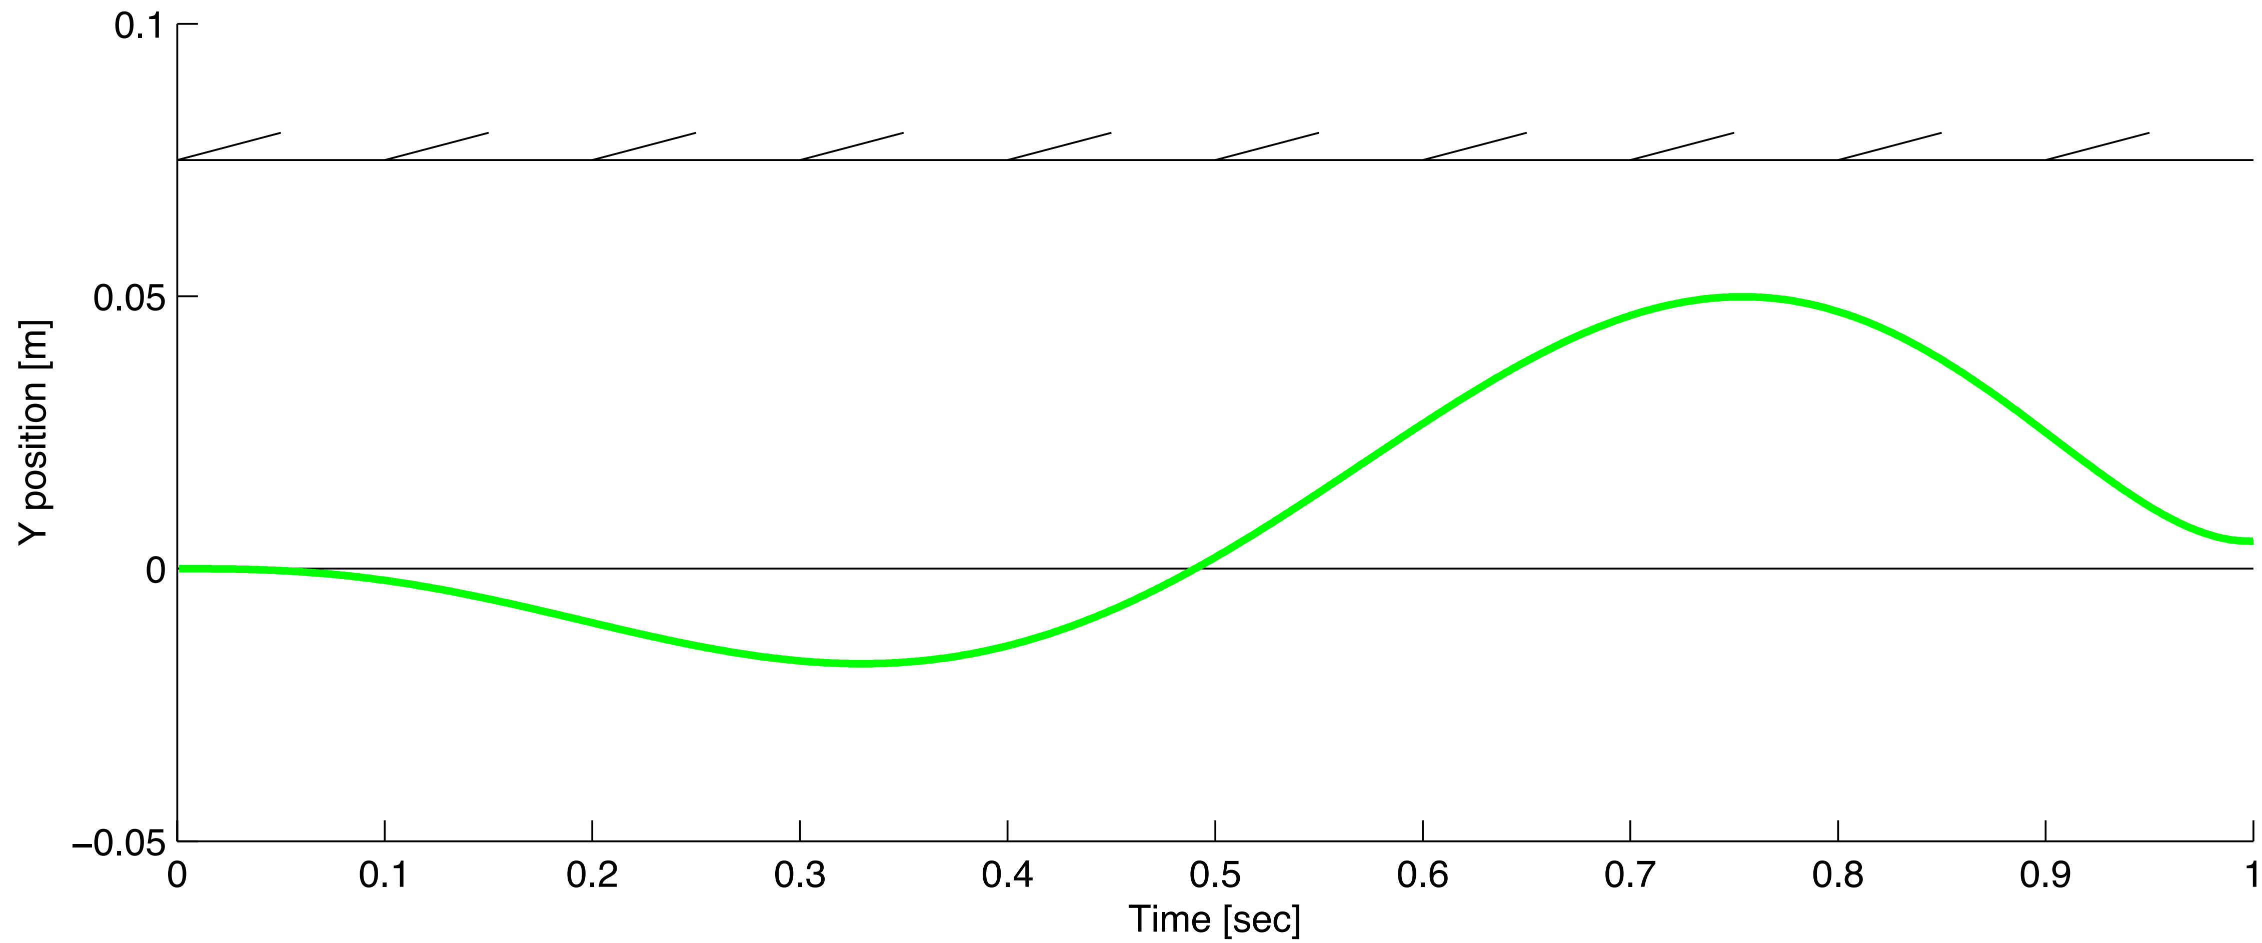

Krömer, Detry, Piater, Peters (submitted). Combining Active Learning and Reactive Control for Robot Grasping, Robotics and Autonomous Systems

Generalizing with DMPs

Generalization may cause collisions with objects...

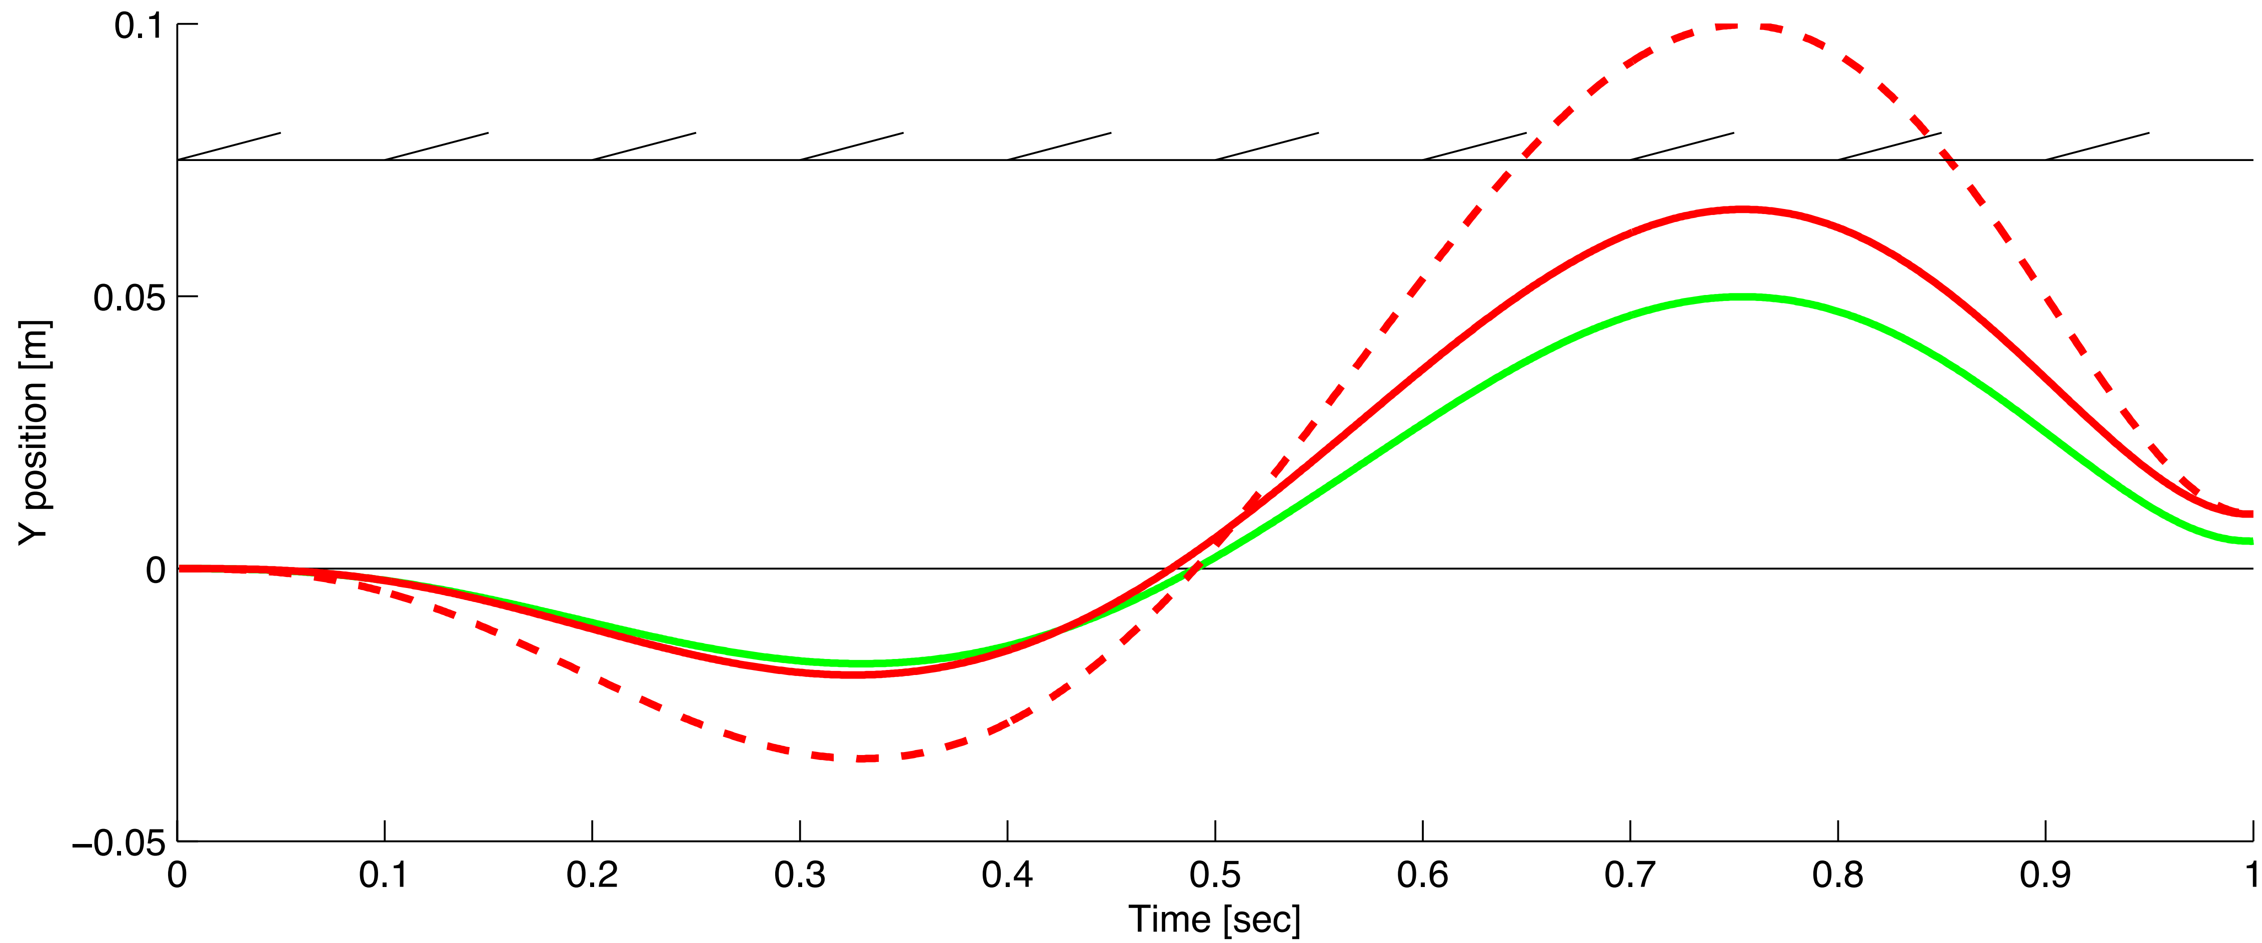

How can we fix this?

Krömer, Detry, Piater, Peters (submitted). Combining Active Learning and Reactive Control for Robot Grasping, Robotics and Autonomous Systems

How can we deal with this?

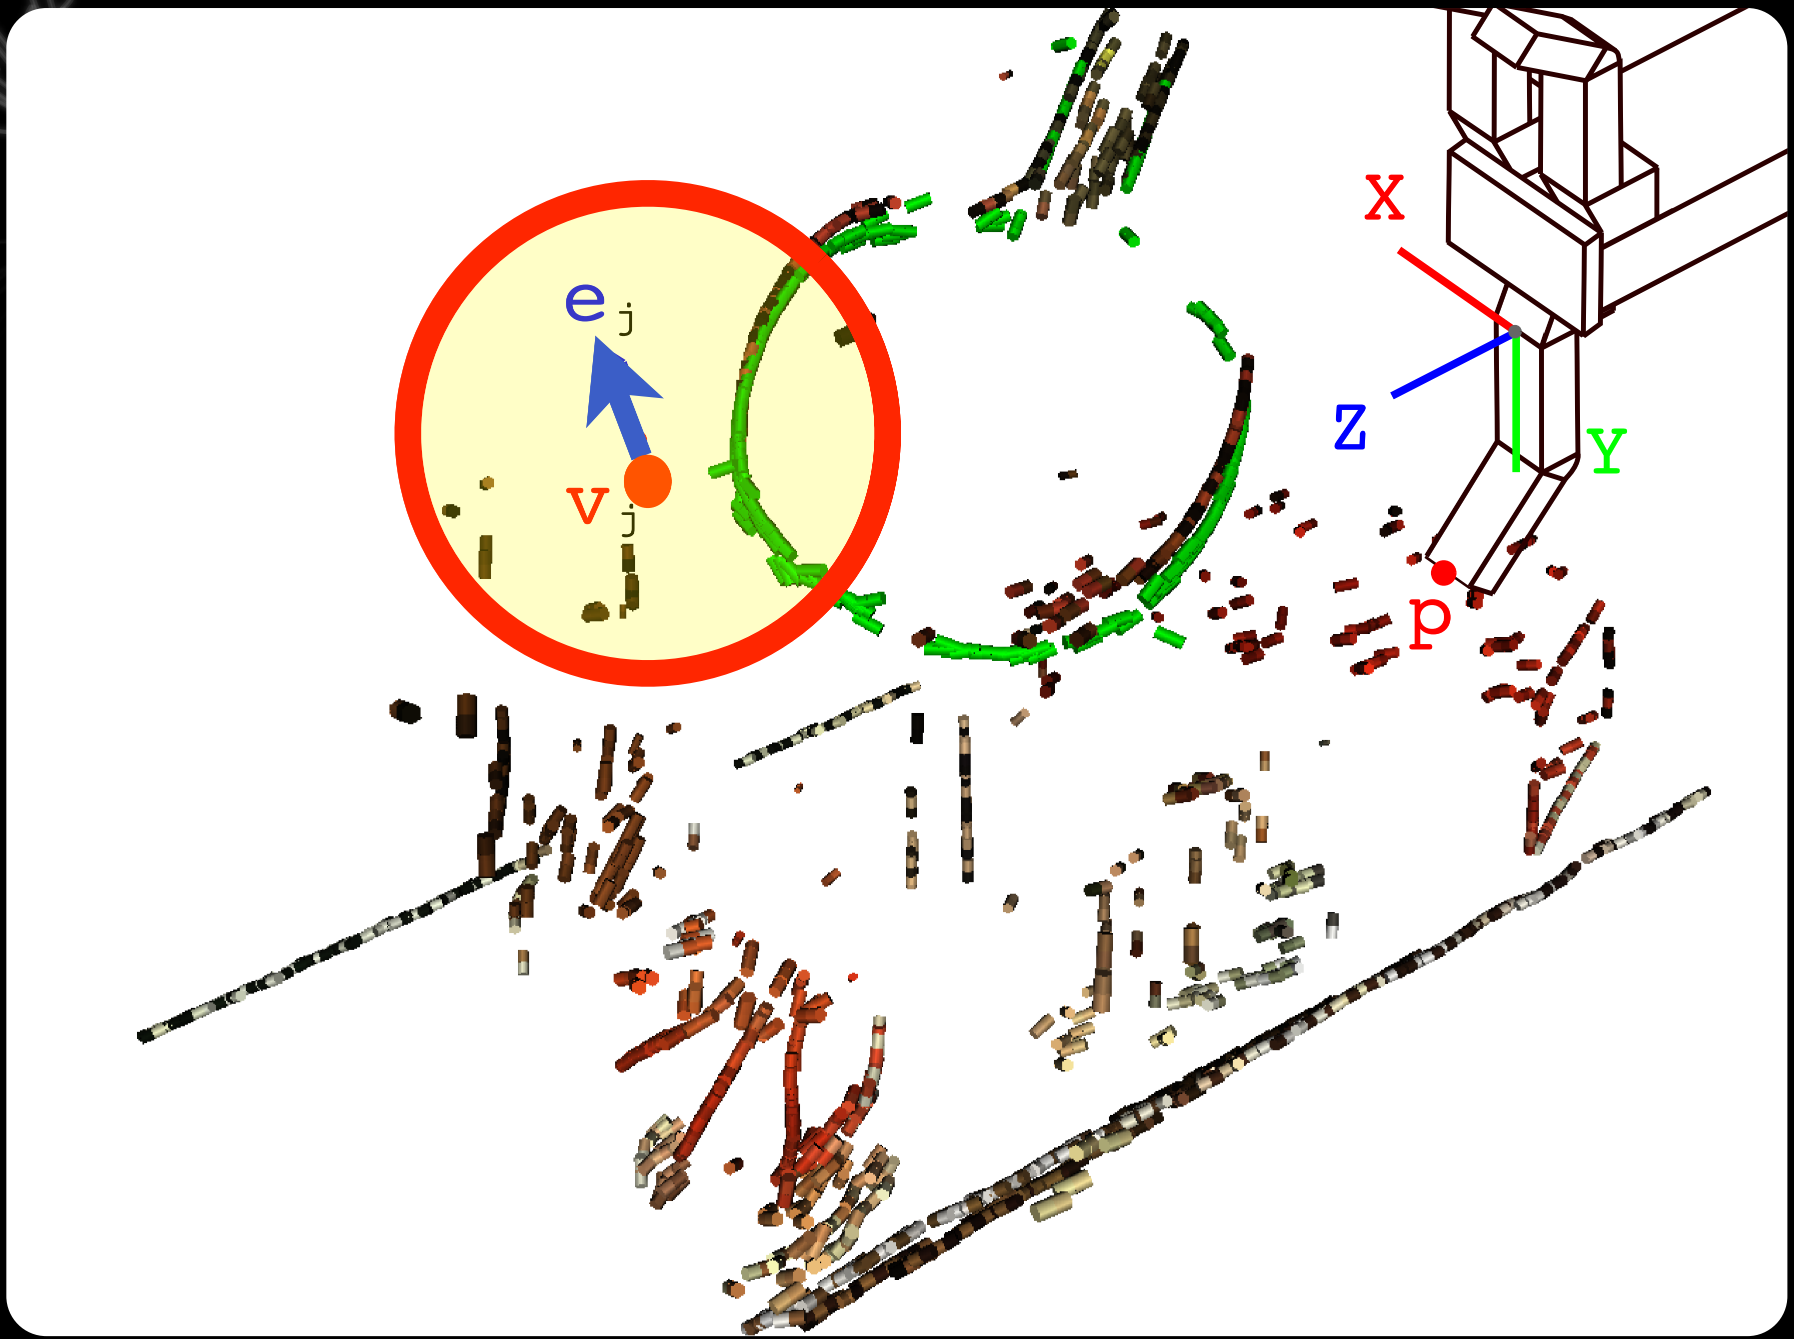

DMPs are dynamic systems and can be modified straightforwardly by force fields!

Krömer, Detry, Piater, Peters (submitted). Combining Active Learning and Reactive Control for Robot Grasping, Robotics and Autonomous Systems

ECVD-based Potential Fields

Hand Detracting Field

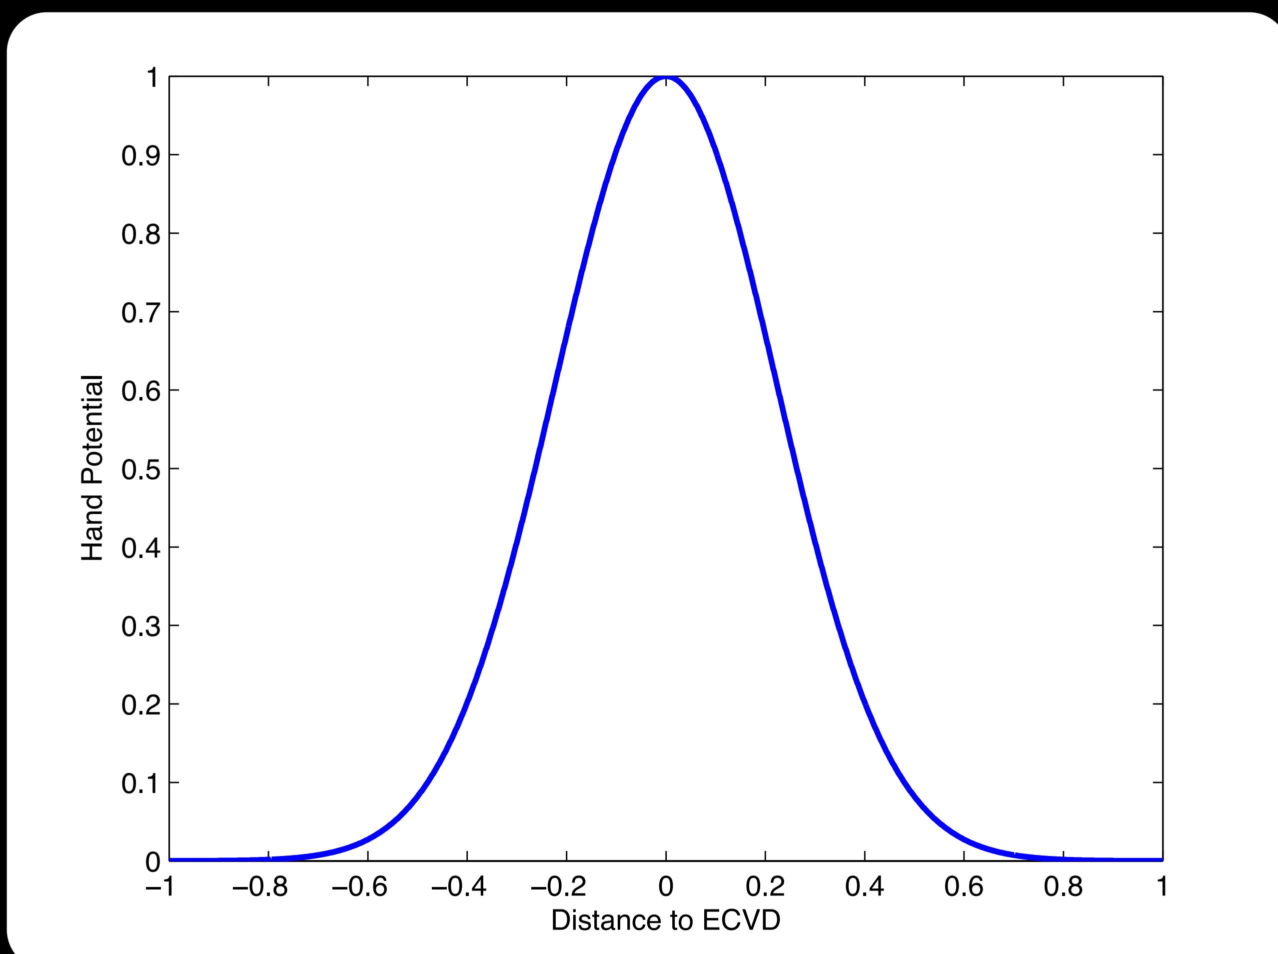

Finger Dectracting Field

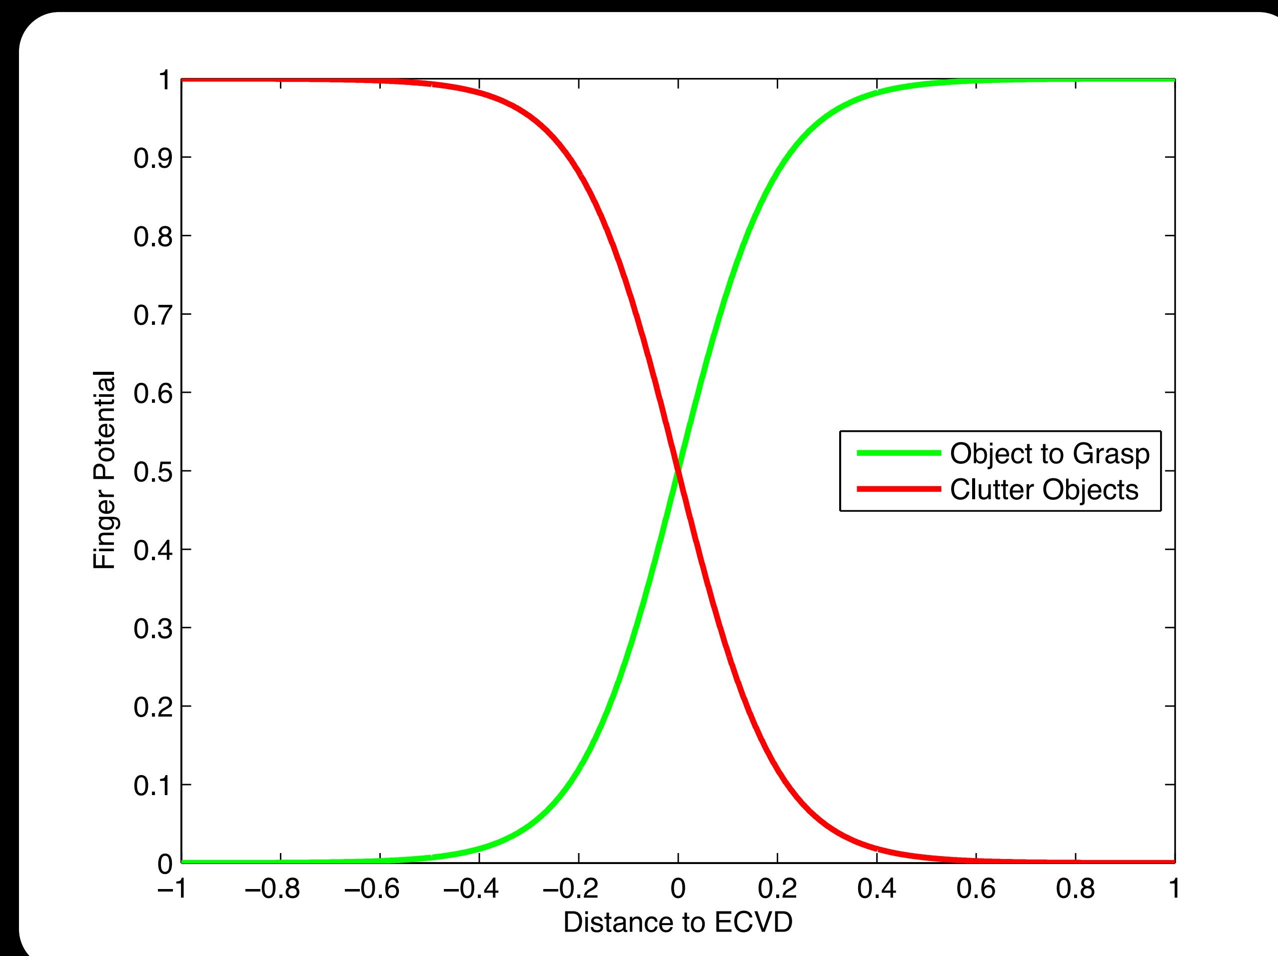

Krömer, Detry, Piater, Peters (submitted). Combining Active Learning and Reactive Control for Robot Grasping, Robotics and Autonomous Systems

Resulting Hand Preshaping

Plane

Handle

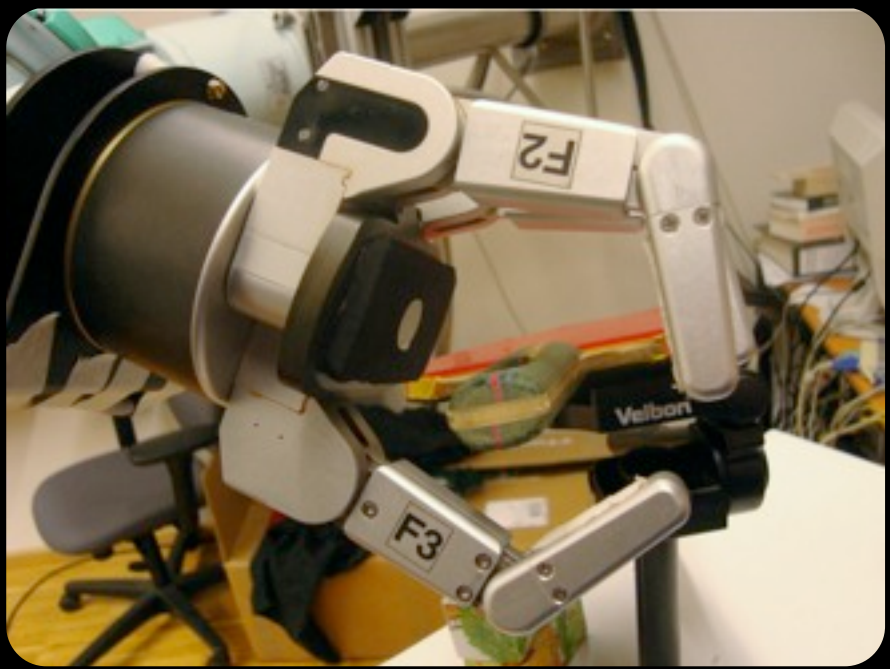

Slanted

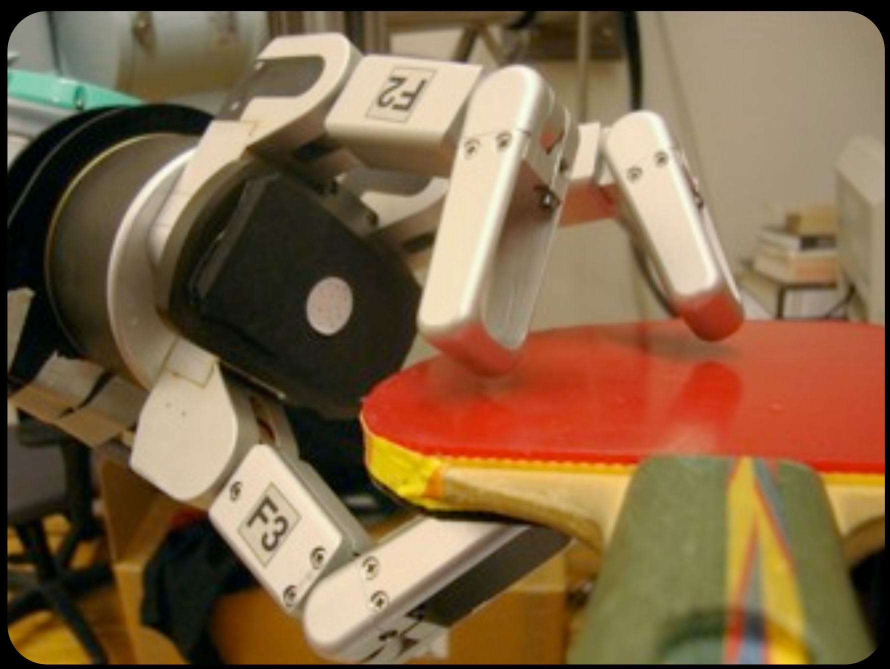

Krömer, Detry, Piater, Peters (submitted). Combining Active Learning and Reactive Control for Robot Grasping, Robotics and Autonomous Systems

Outline

1. A Motor Skill Learning Framework
2. Exploring new Objects efficiently
3. Adapting Movements Primitives with Vision Descriptors
4. Conclusion

Conclusion

- Quick Intro to our view on *Motor Skill Learning for Robotics*
- Representations:
  - Motor primitives for actions
  - Vision descriptors for modifying actions and context
- Resulting New Methods:
  - Grasp-Point Exploration with Continuum-armed GB
  - A straightforward modification to make Imitation work
- Our results appear promising!

Supported by

Towards Closing the Loop: Active Learning for Robotics

Call for Contributions

We invite submission of extended abstracts to the workshop. Extended abstracts should be 1-2 pages in length, formatted in according to RSS style. However, submissions should not be blind. Extended abstracts should be sent in PDF or PS file format by email to alrss10@gmail.com

The selected submission may be accepted either as an oral presentation or as a poster presentation. We encourage participants who can contribute in the following areas:

- Active learning
- Active filtering
- Sequential experimental design
- Adaptive sensing
- Optimal information gathering
- Autonomous exploration
- Bayesian optimization
- Active cognitive development
- Attention systems or gaze control
- Sensor placement
- Active vision
- Online decision making
- Selection criteria/Utility functions
- Information theoretic metrics in the context of robotics.

The above list is not exhaustive, and we welcome submissions on highly related topics too. Accepted extended abstracts will be made available online at the workshop website.

COLLOCATED WITH:

ROBOTICS SCIENCE AND SYSTEMS

SPONSOR:

PASCAL2 Pattern Analysis, Statistical Modelling and Computational Learning

SUPPORTED BY:

IEEE Robotics and Automation Society Technical Committee on Robot Learning

Submission Deadline: May 19, 2010

Thanks

Thanks for your attention!

Stereo Camera

- New head (two identical eyes)
- Pan-tilt unit
- Semi-automatic calibration
- Software based head protection
- 3D model acquisition

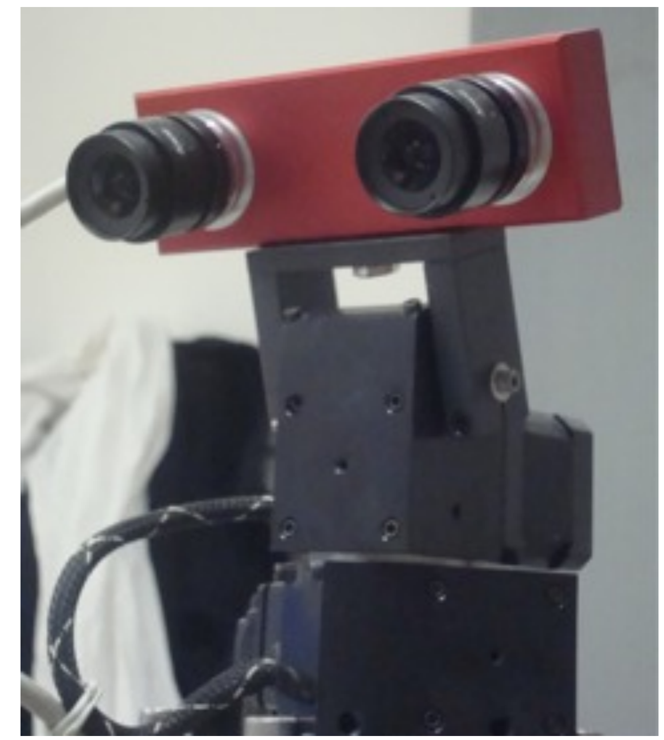

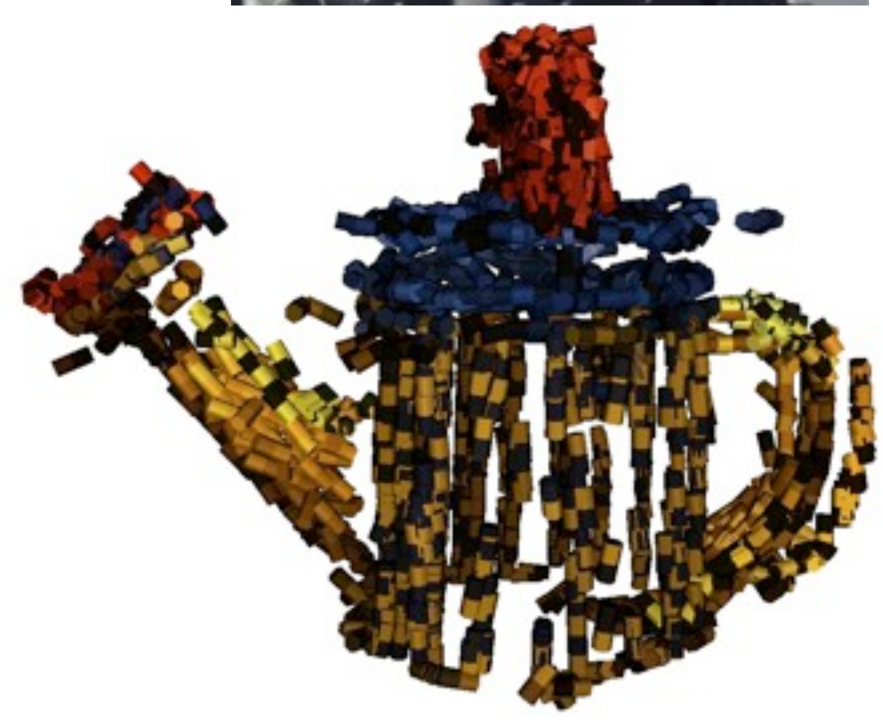

Simulation

- 17 degree of freedom model
- Incorporated control of hand, arm, and pan-tilt unit
- Test movements for joint limits before execution
- Simulated eye sight
- Open Dynamics Engine

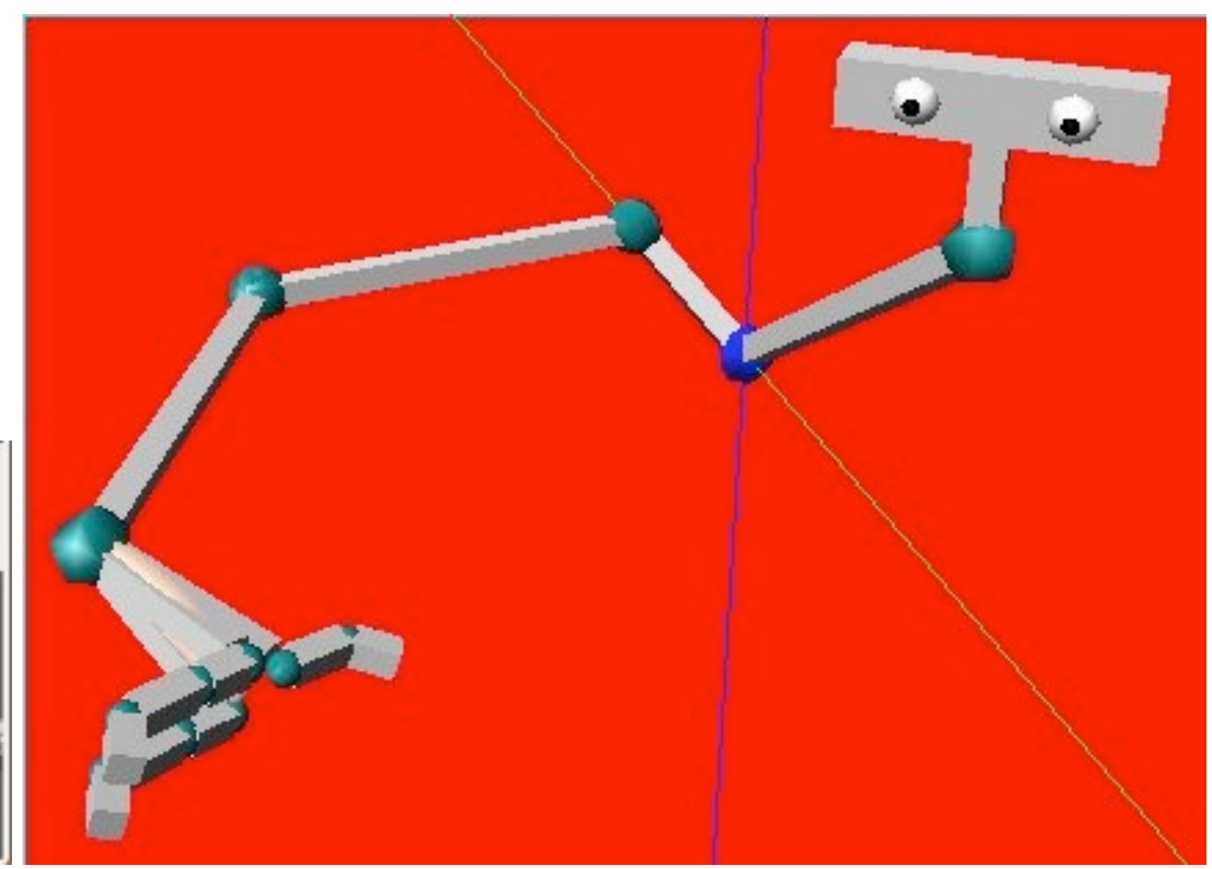

Grasping Experiment System

- Modular experiment setup
  - Observing scene
  - Object Pose Estimation
  - Grasp Selection
  - Robot Execution
- Automated data collection
- Early Cognitive Vision system
- Collaborated on experiments

Collaborators

Renaud Detry

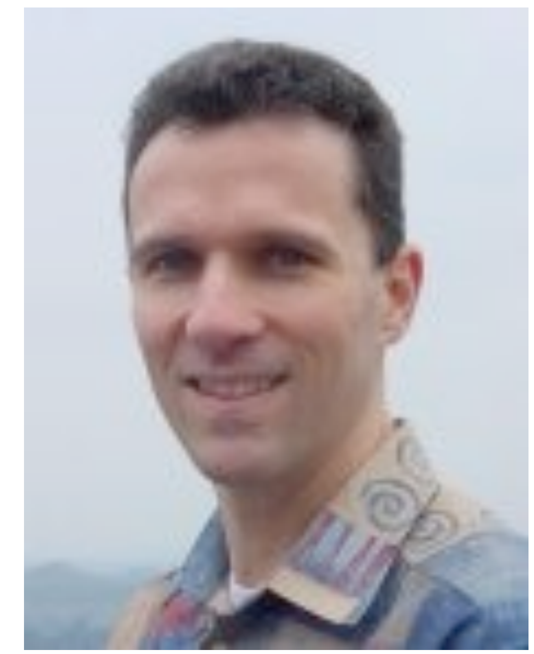

Justus Piater

Dirk Kraft

Norbert Krueger

Robot Grasping

- Where to grasp sets context for how to grasp
  - want to reflect this hierarchical structure
- Current Setup:

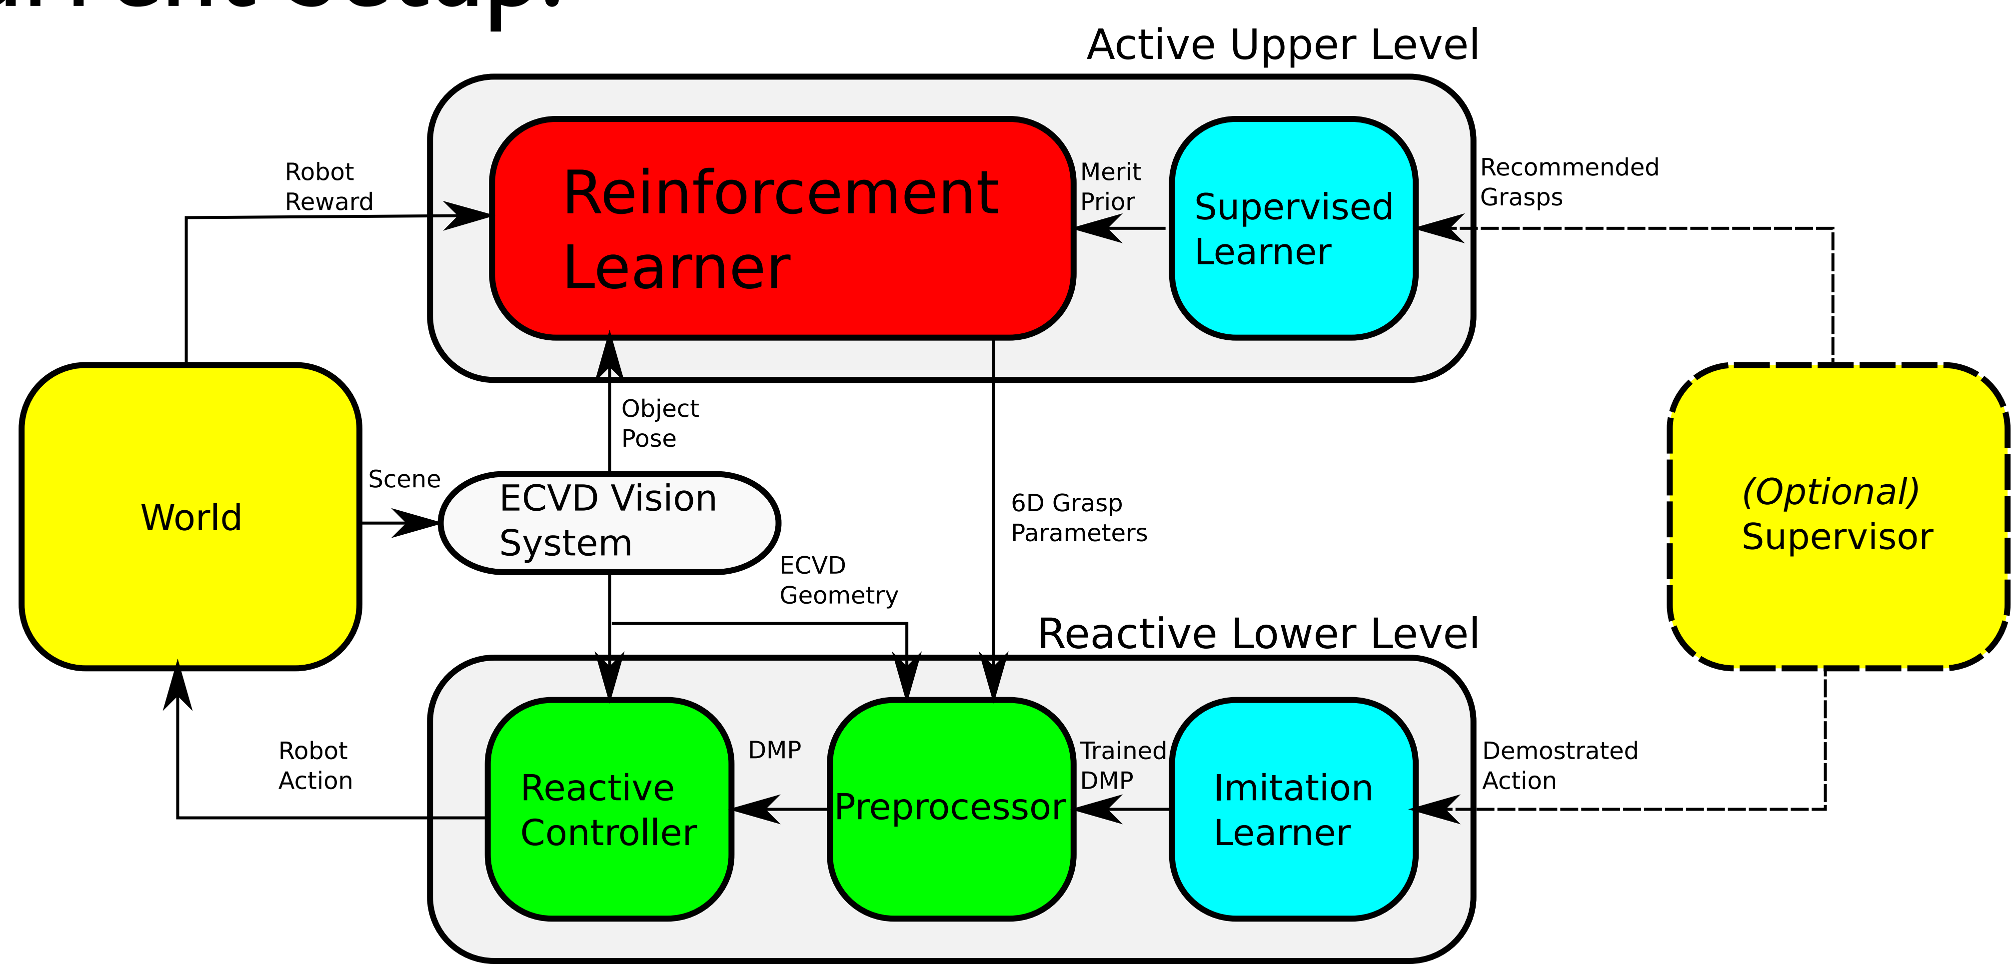

Future Directions

- Generalize grasps between objects
  - Supervised Learning
  - Incorporate into architecture as initialization
- Learning for how to grasp
  - Adapt grasp execution by learning
- Sequencing DMPs together for complex tasks
  - Reinforcement learning
  - Planning with learned operators

Outline

I. Introduction

II. Robot System Integration

III. Grasping System:

  I. Where to grasp

  II. How to grasp

  III. Experiments

IV. Future Directions

V. Conclusion

Summary

- Autonomously improve grasp performance using reinforcement learning
- Grasp execution using motor primitives augmented by vision information
- Future plans:
  - Supervised learning to generalize between objects
  - Reinforcement learning for multi-action manipulation tasks
- Happy Holidays!