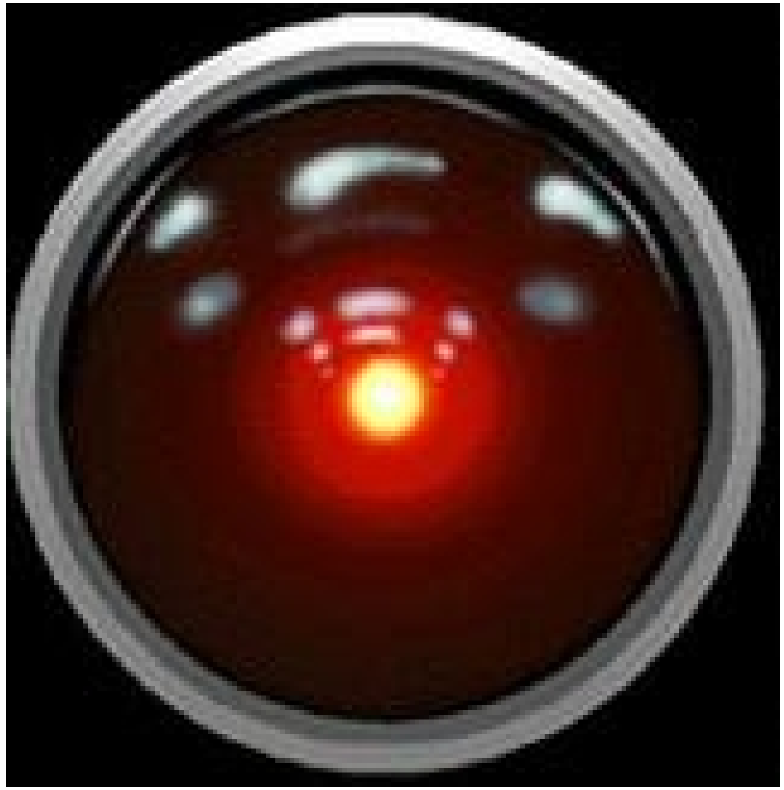

CS 378: Autonomous Intelligent Robotics

Instructor: Jivko Sinapov

http://www.cs.utexas.edu/~jsinapov/teaching/cs378/

Tool Use in Animals and Robots

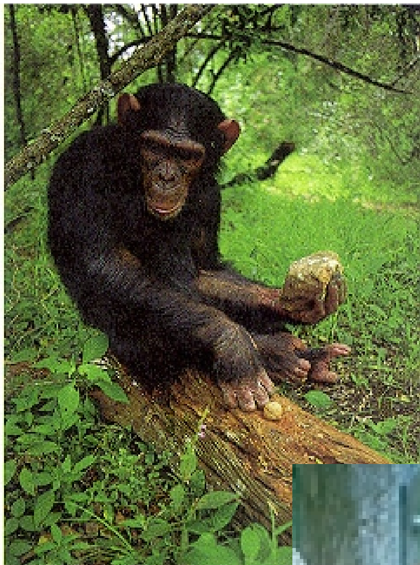

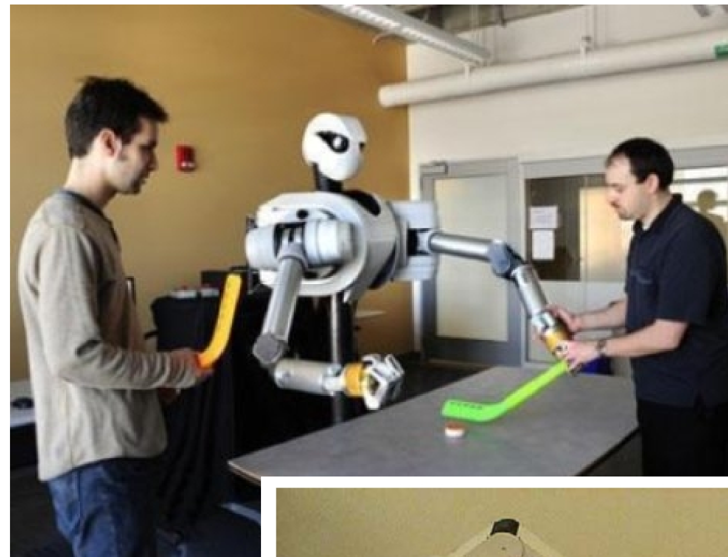

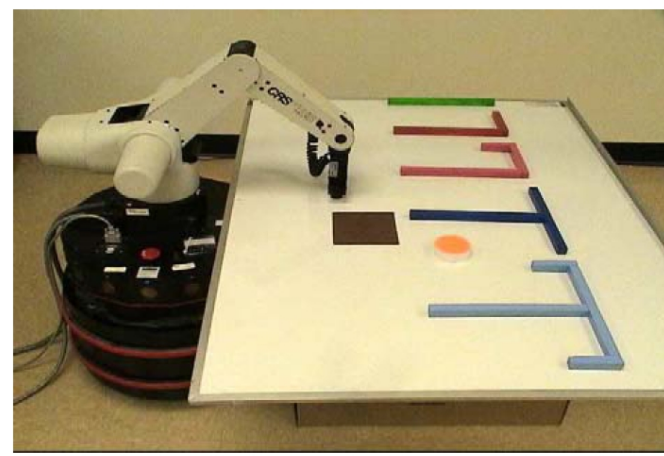

Announcements

FRI Survey – please take the time to respond

Announcements

Final Projects Presentation Date:

Thursday, May 12, 9:00-12:00 noon

Tool Use in Animals and Robots

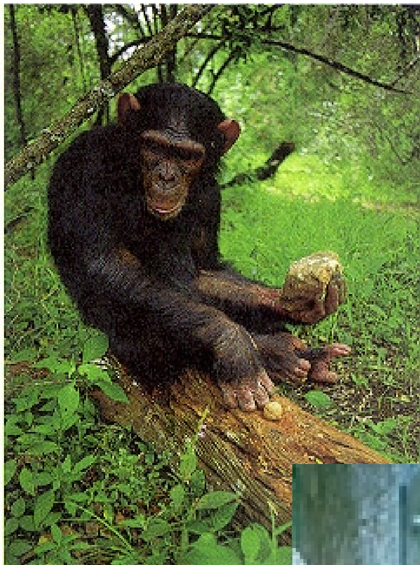

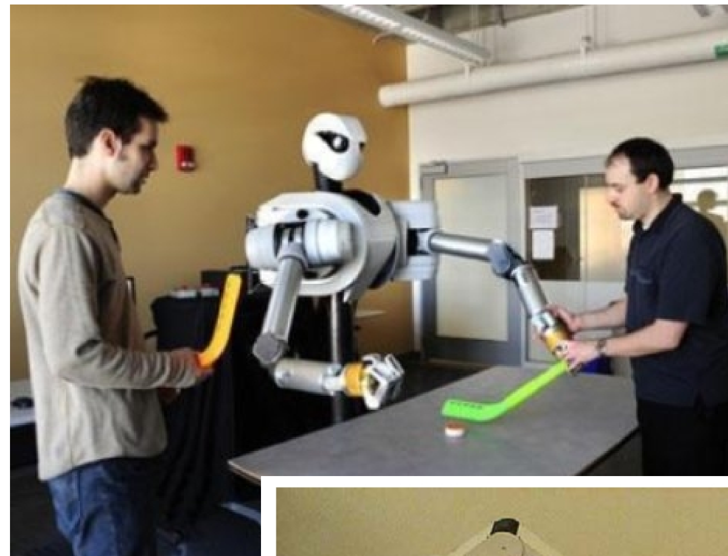

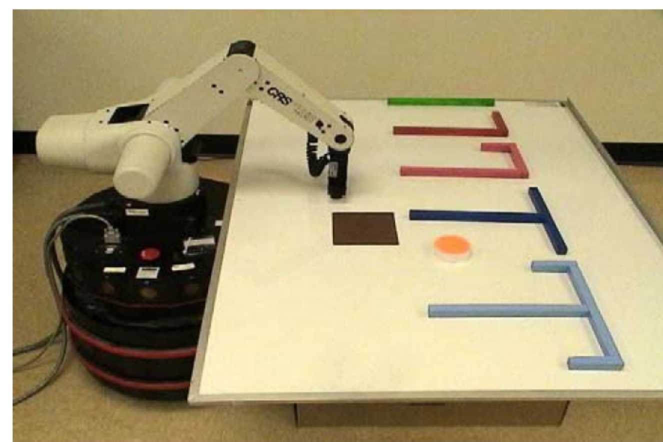

Main References

- Sinapov, J., and Stoytchev, A. (2007). *Learning and Generalization of Behavior-Grounded Tool Affordances.* In proceedings of the IEEE International Conference on Development and Learning (ICDL 2007)
- Sinapov, J., and Stoytchev, A. (2008). *Detecting the Functional Similarities Between Tools Using a Hierarchical Representation of Outcomes.* In proceedings of the IEEE International Conference on Development and Learning (ICDL 2008)

Tool Use in Animals




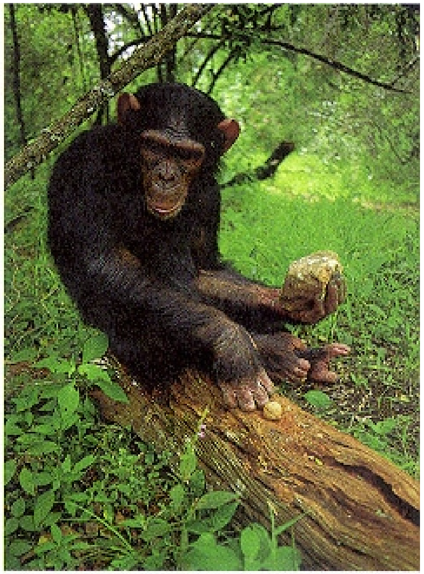

Origins of Human Tool Use (according to Hollywood)

Tool Use in the Wild

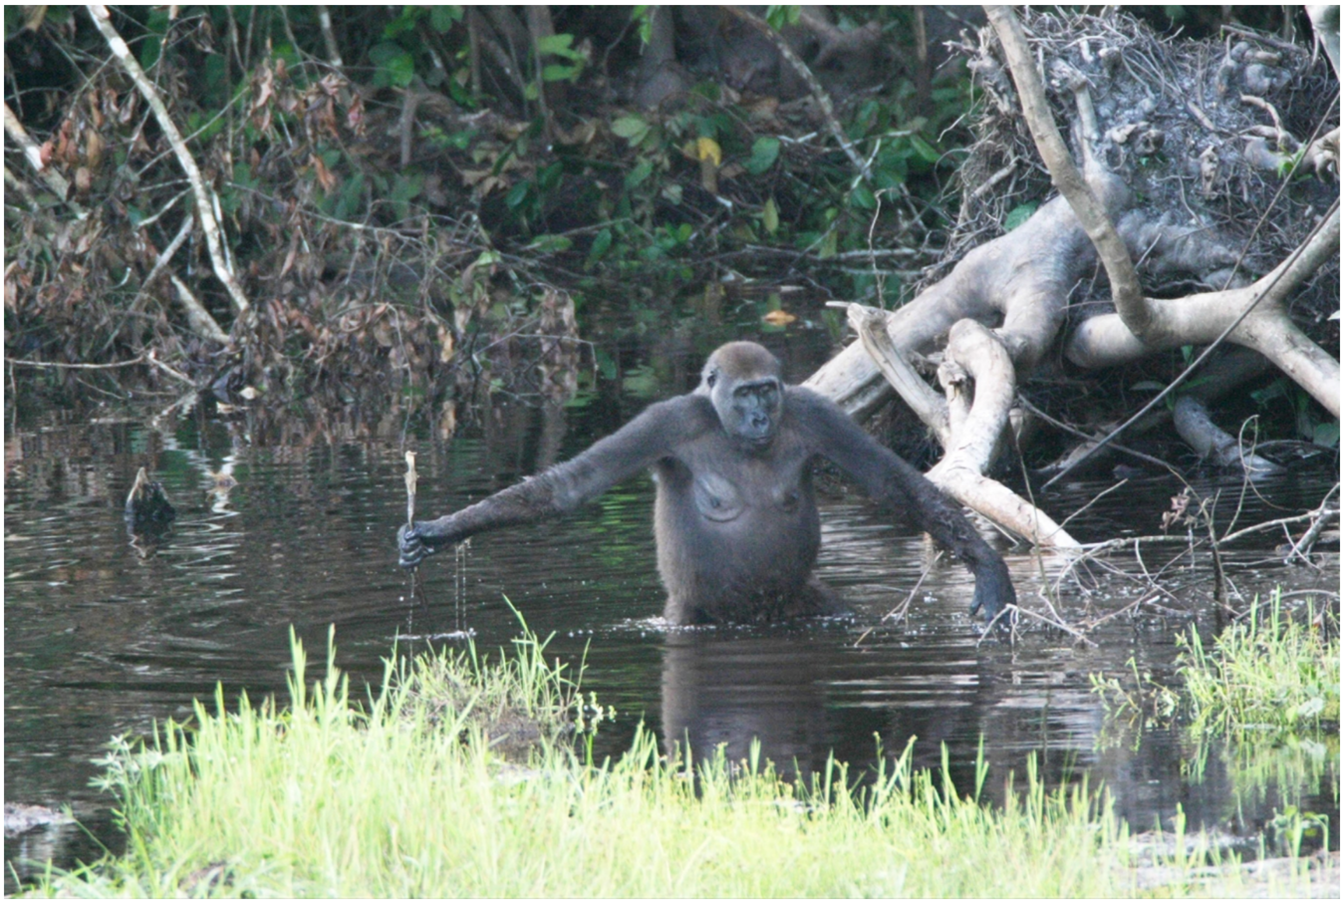

Tool Use in the Wild

Tool Use in the Wild

"A young corvide bird, confronted with an object it has never seen, runs through practically all of its behavioral patterns, except social and sexual ones."

-Konrad Lorenz

(Nobel Prize Winner, 1973)

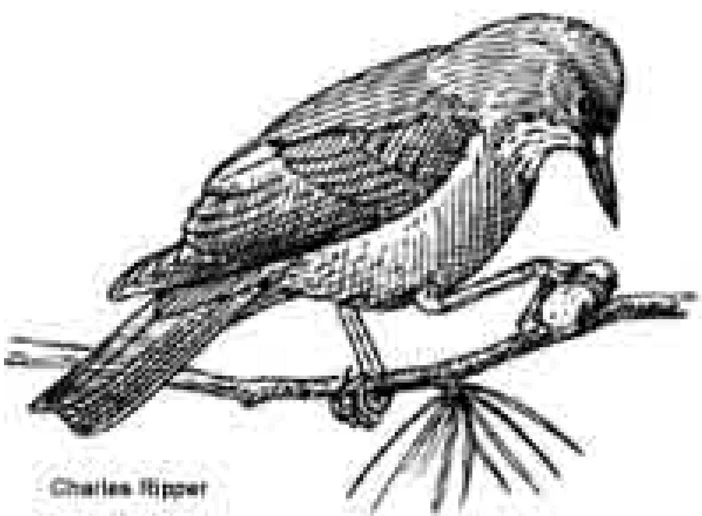

Tool Use in the Lab

Tool Use in the Lab

Ishibashi et al. (2000) show that monkeys can generalize tool-related knowledge from one tool to a novel tool, as long as the novel tool shared similar features with the one to which they had been previously exposed to.

Studies by Povinelli et al. (2000) conclude that monkeys infer simple rules from their experience regarding tool use, e.g., “visual contact leads to movement.”

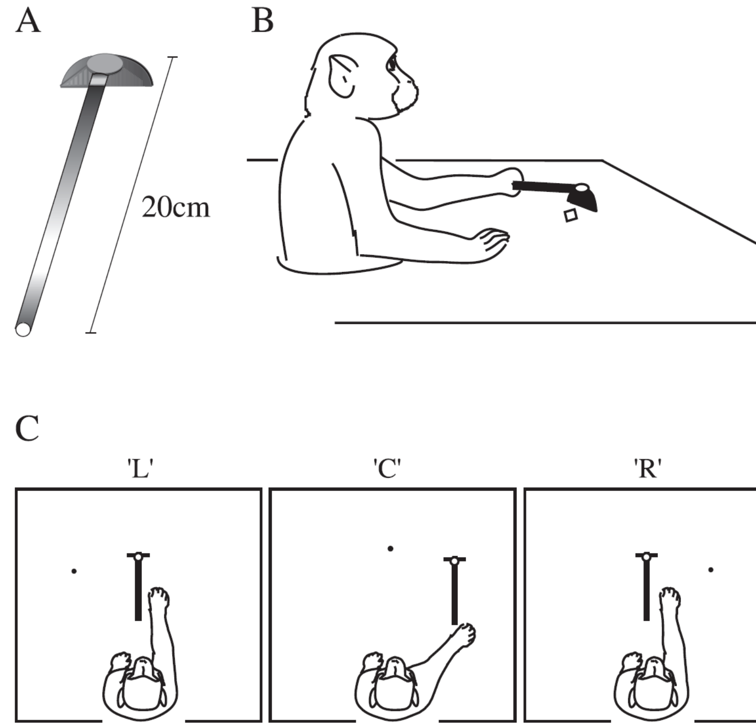

[Ishibashi, Hihara, Iriki (2000)]

Adaptive Body Representation

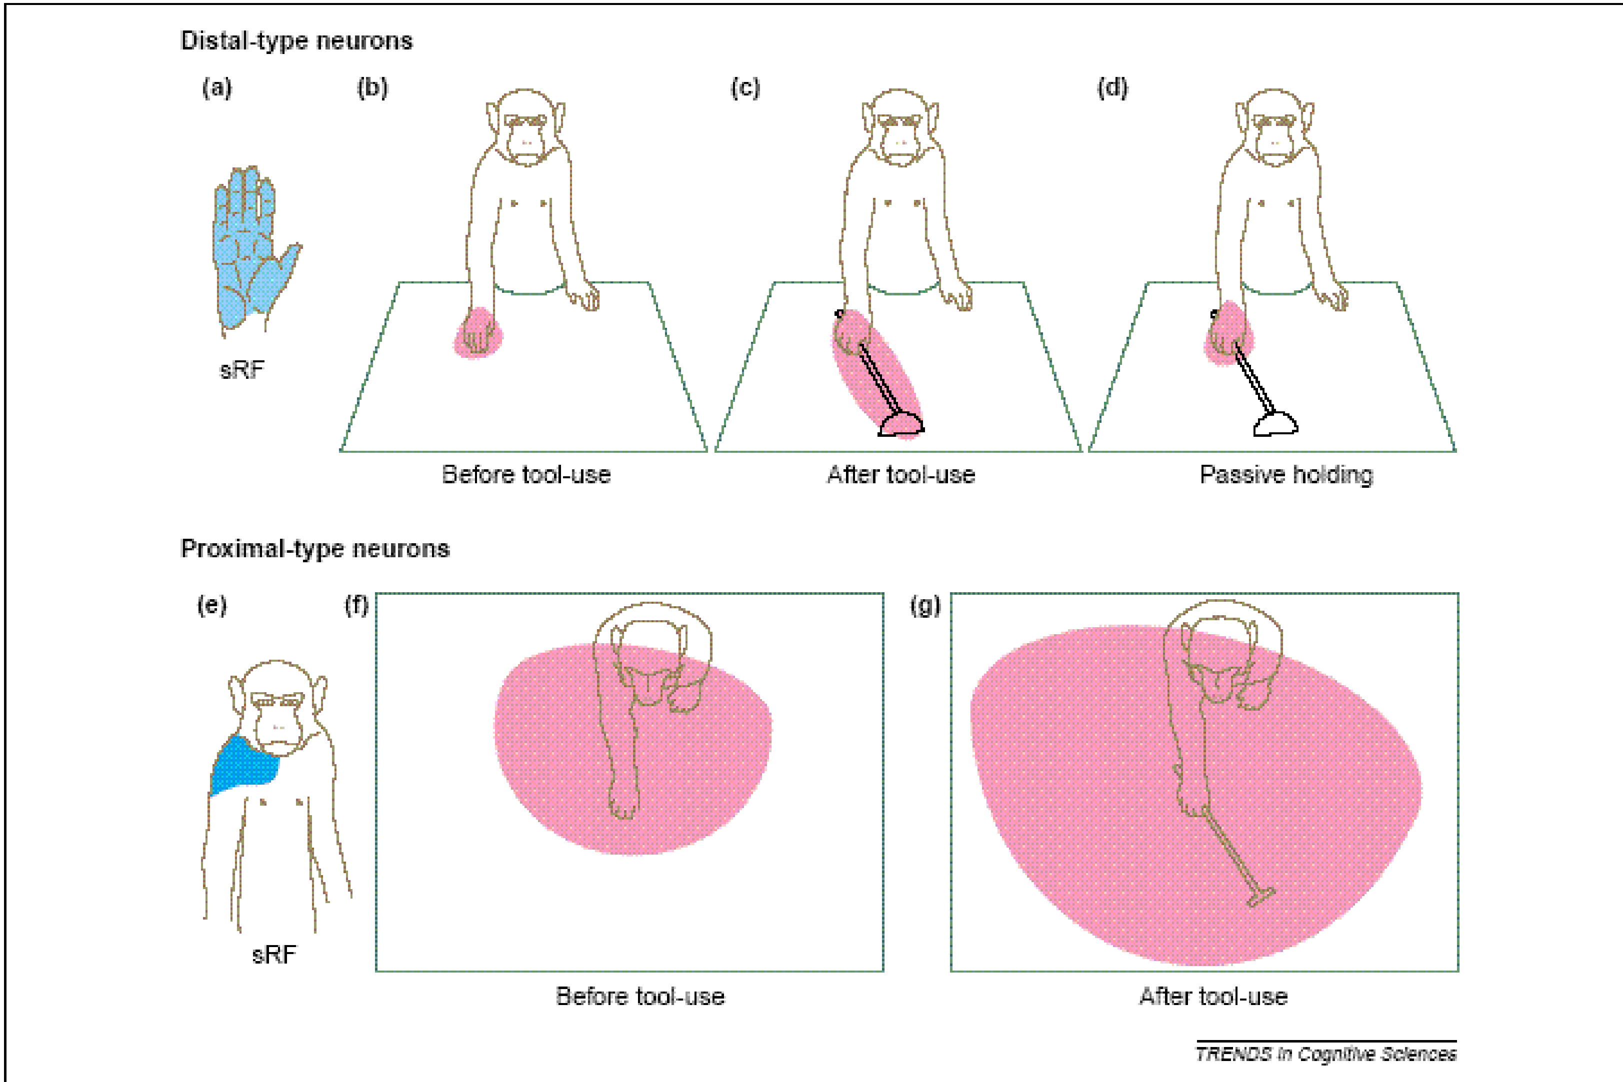

Figure 1. Changes in bimodal receptive field properties following tool-use. The somatosensory receptive fields (sRF) of cells in this region were identified by light touches, passive manipulation of joints or active hand-use. The visual RF (vRF) was defined as the area in which cellular responses were evoked by visual probes (the most effective ones being those moving towards the sRF). (a) sRF (blue area) of the 'distal type' bimodal neurons and their vRF (pink areas) (b) before tool-use, (c) immediately after tool-use, and (d) when just passively grabbing the rake. (e) sRF (blue area) of 'proximal type' bimodal neurons, and their vRF (pink areas) (f) before and (g) immediately after tool-use.

Tool Use by Robots

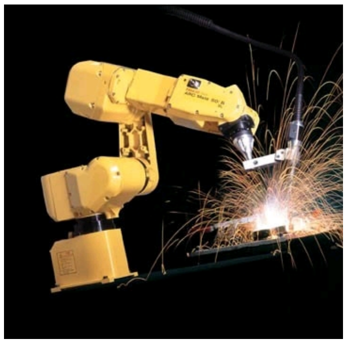


“Specialists”


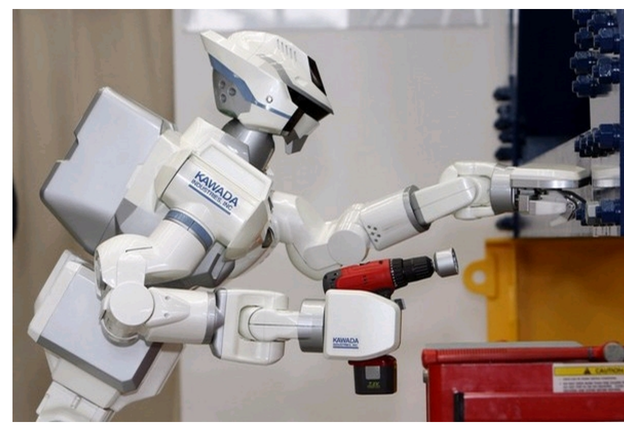

“Generalists”

An Example Tool-using Robot

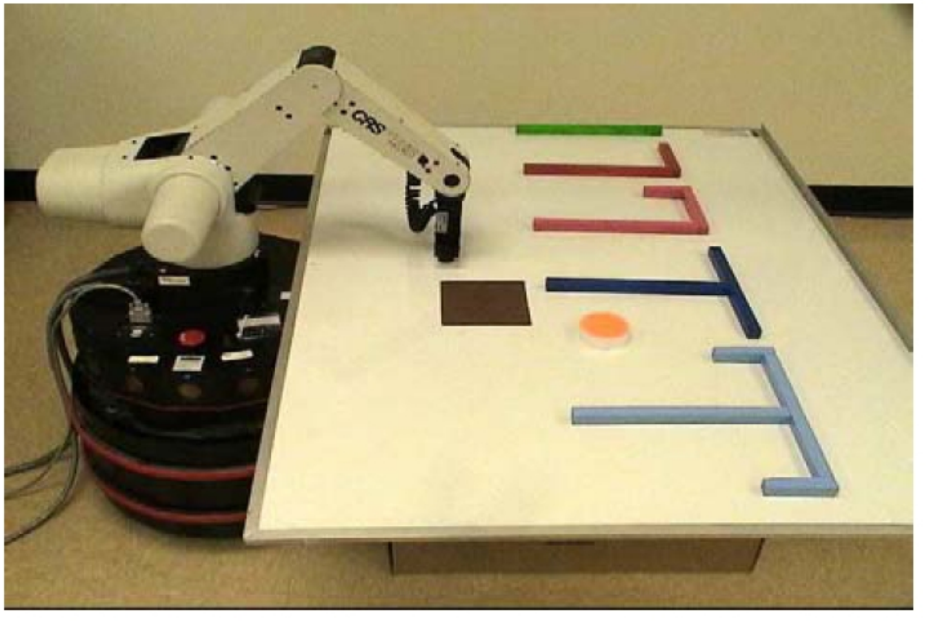

Stoytchev, A., "Behavior-Grounded Representation of Tool Affordances", In Proceedings of IEEE International Conference on Robotics and Automation (ICRA), Barcelona, Spain, April 18-22, 2005.

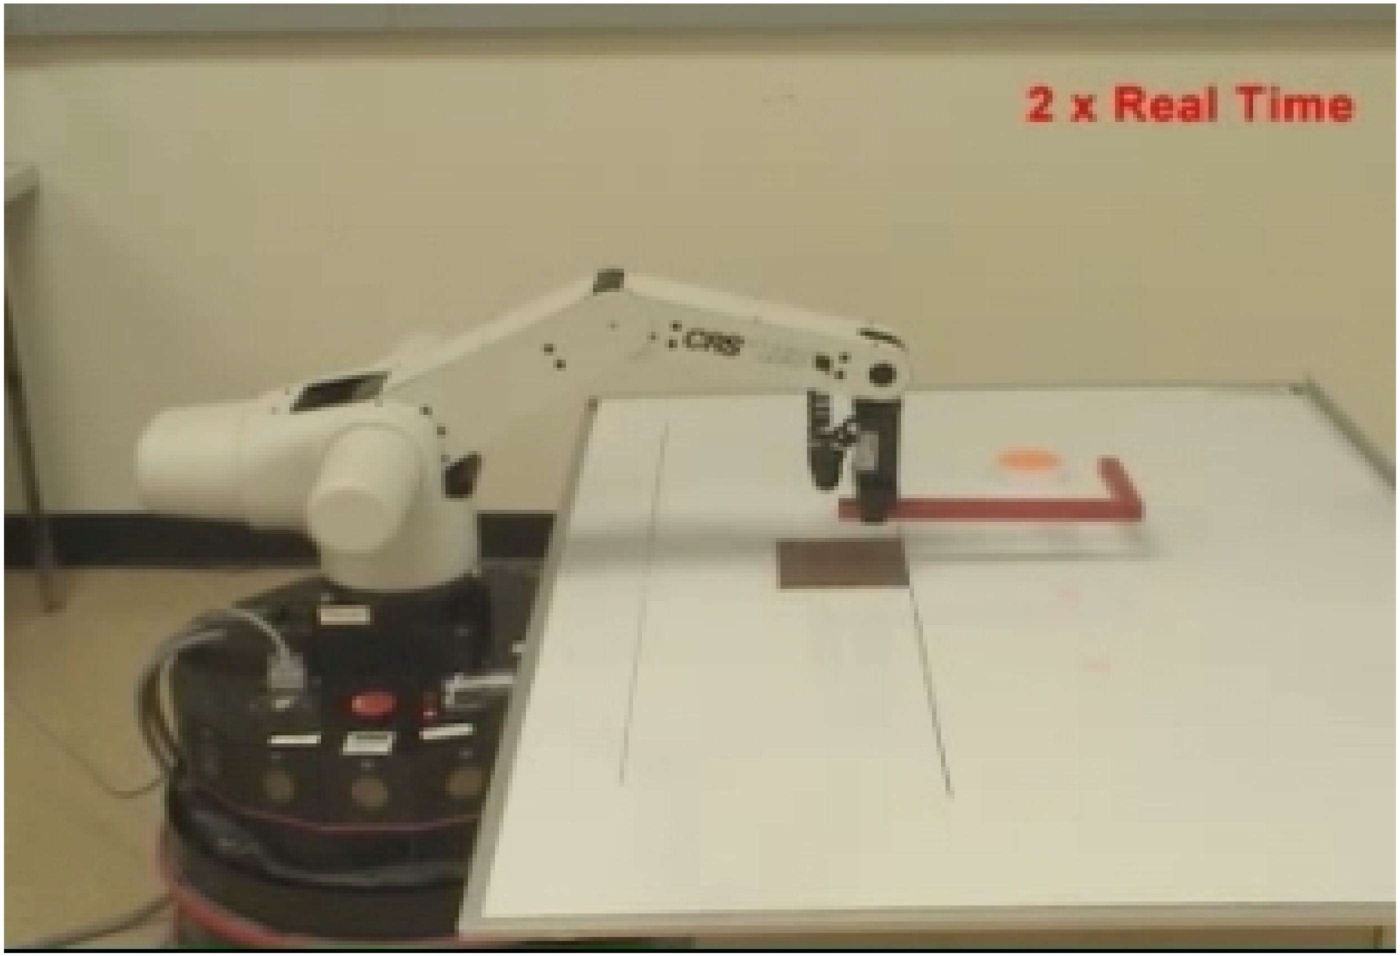
2 x Real Time

What did the robot actually learn?

Grasping Behavior and its Parameters	Exploratory Behavior and its Parameters	O_{start}	O_{end}	Replication Probability
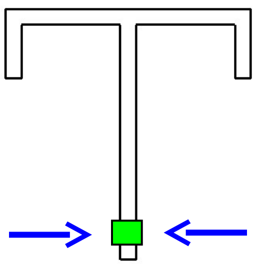 Grasp point	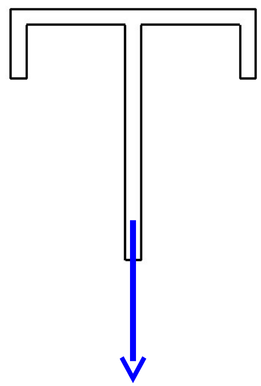 Contract Arm 5 inches	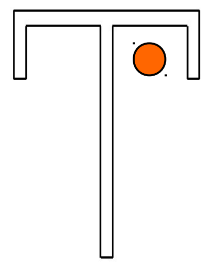	detected movement 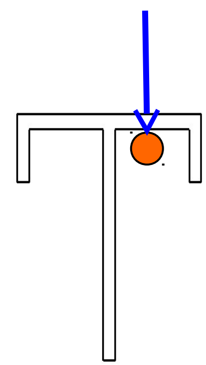	$\frac{\text{Times Successful}}{\text{Times Used}}$

[Stoytchev (2005)]

What did the robot actually learn?

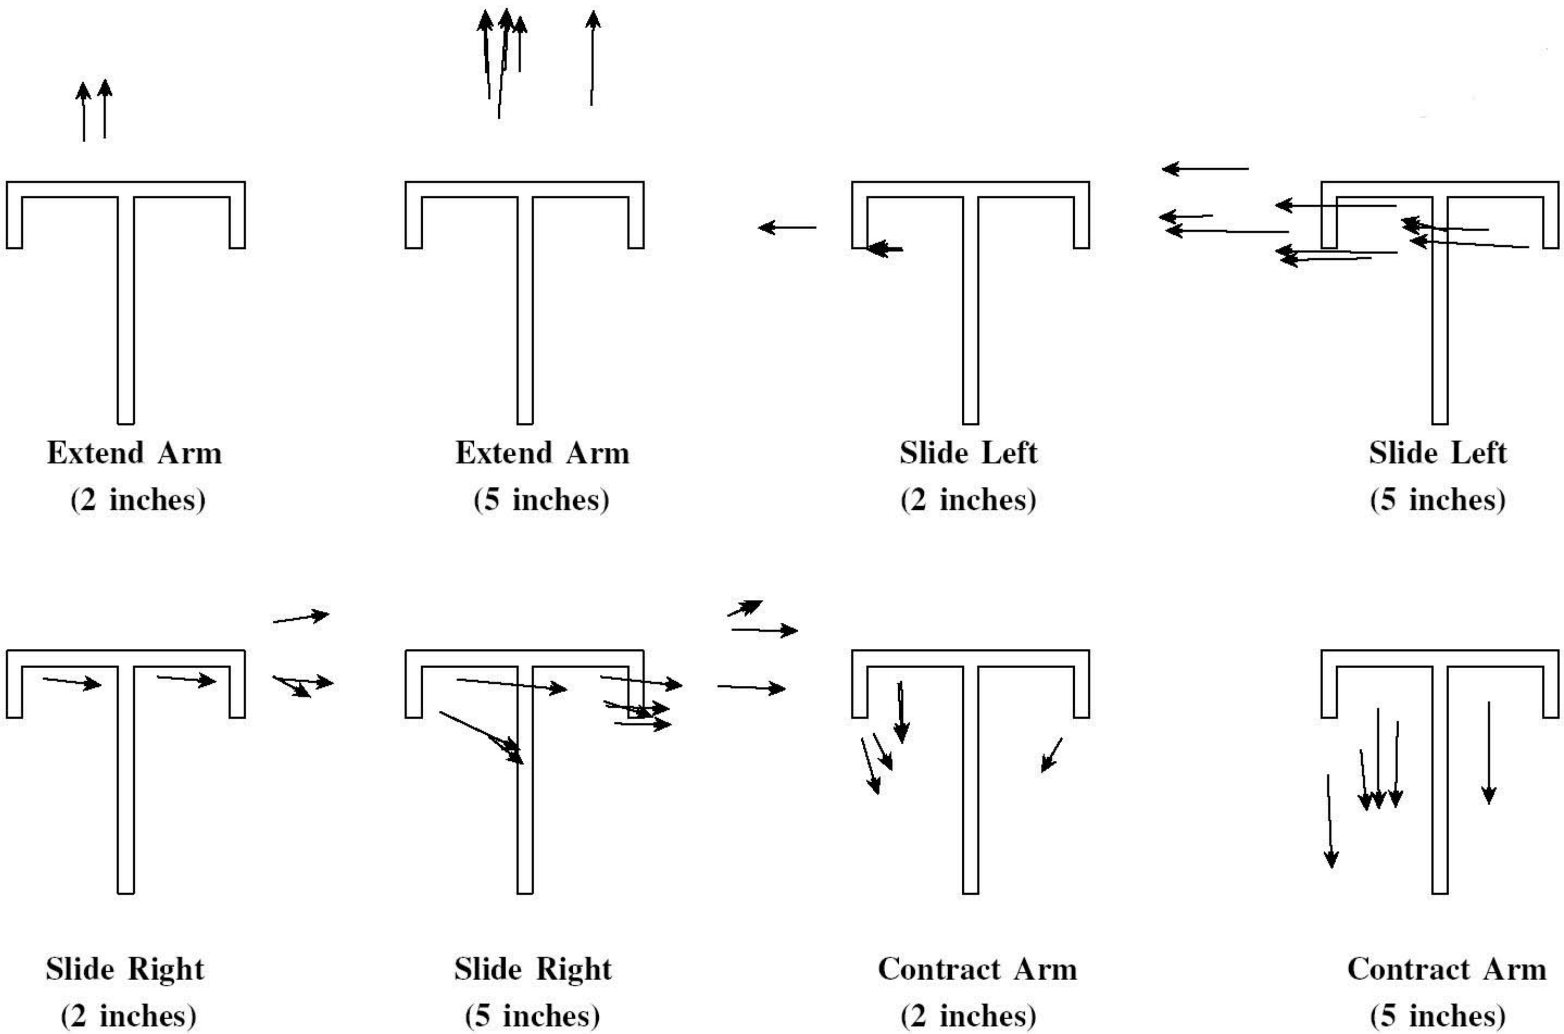

Adaptation to “Broken” Tool

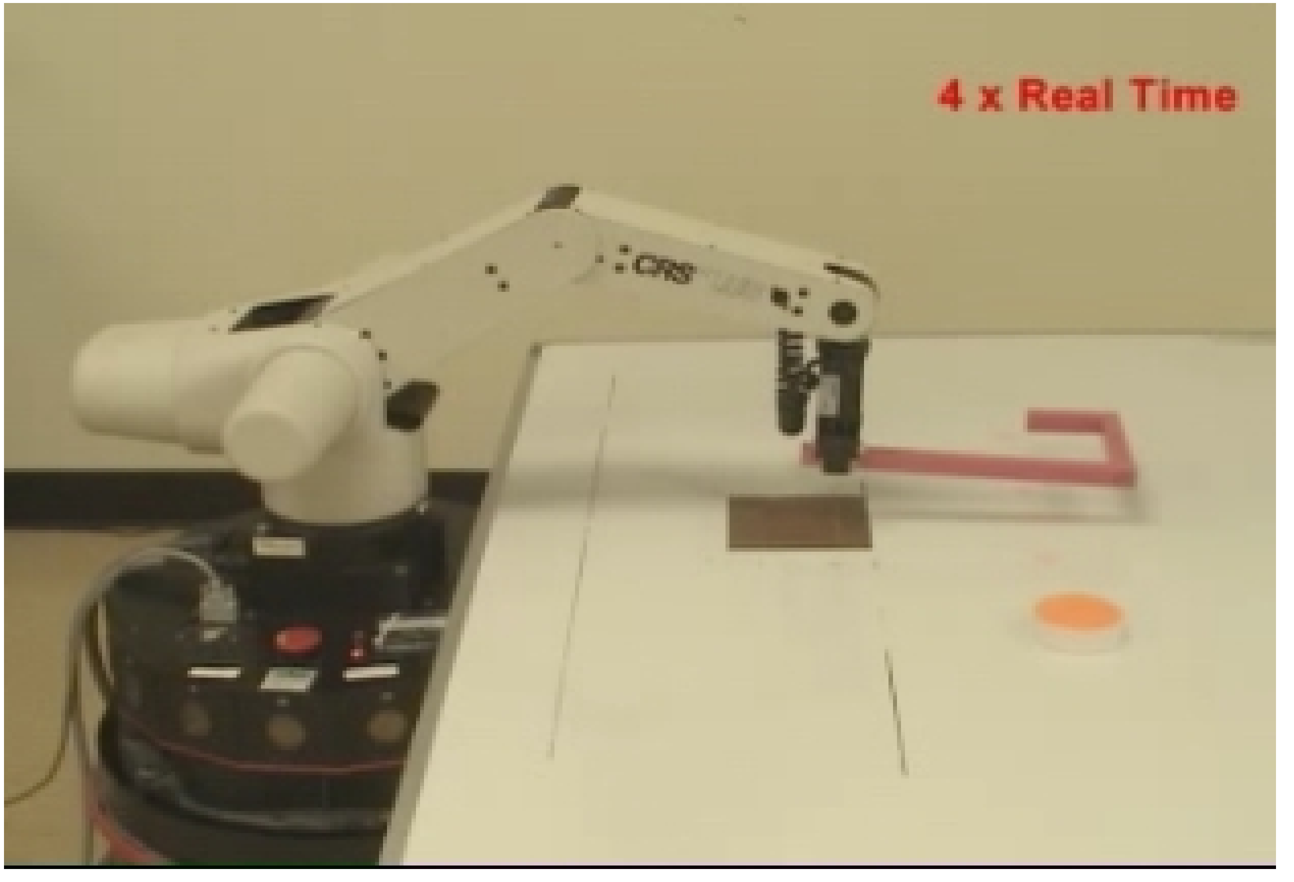

Limitations

Learned affordances are kept in a look-up table, difficult to predict consequences of new actions with the tool if the data is not already included.

Knowledge from experience with one tool cannot be applied to a novel tool.

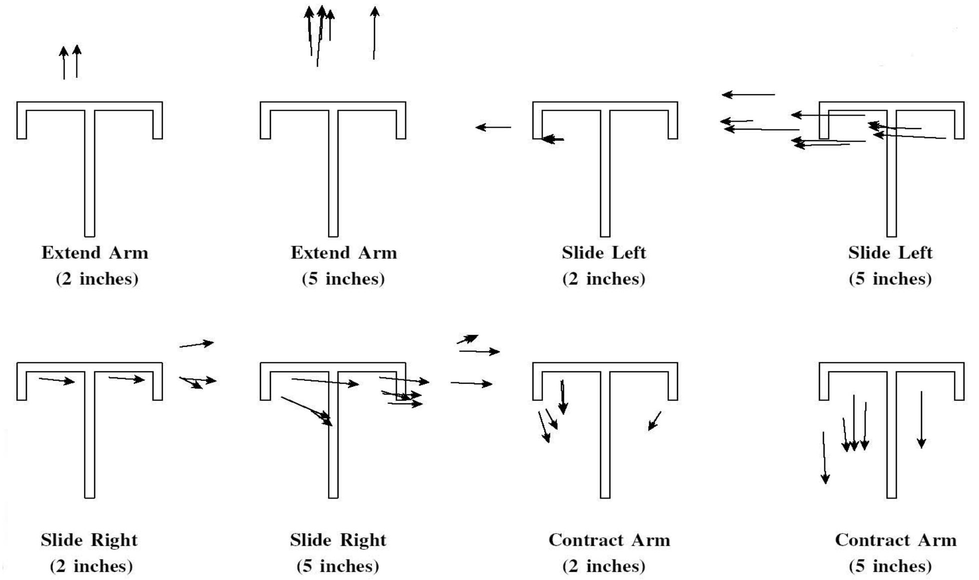

Main Reference

Sinapov, J., Stoytchev, A. (2007).

Learning and Generalization of Behavior-Grounded Tool Affordances

In proceedings of the IEEE International Conference on Development and Learning (ICDL 2007)

Experimental Setup

The robot is a 6-DOF arm simulated in the BREVE robot simulator.

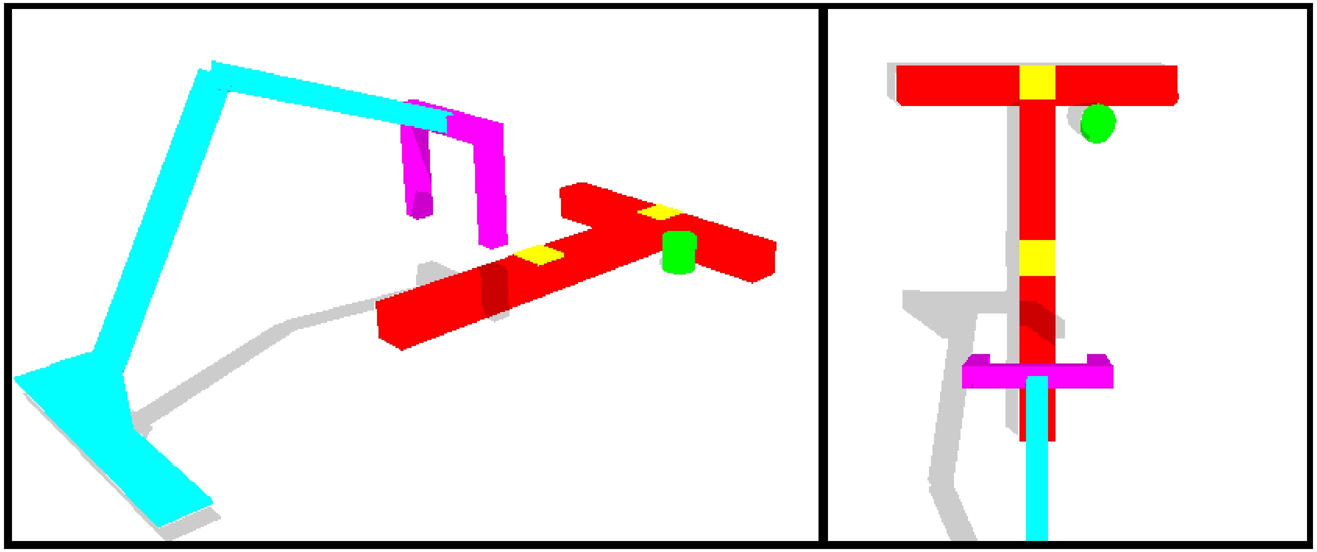

6 tools: T-Stick, L-Stick, L-Hook, Stick, T-Hook, Paddle

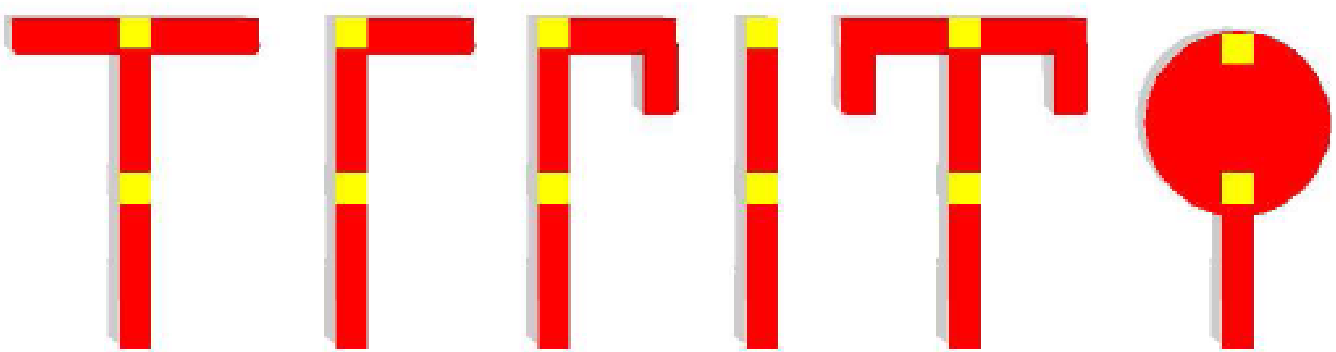

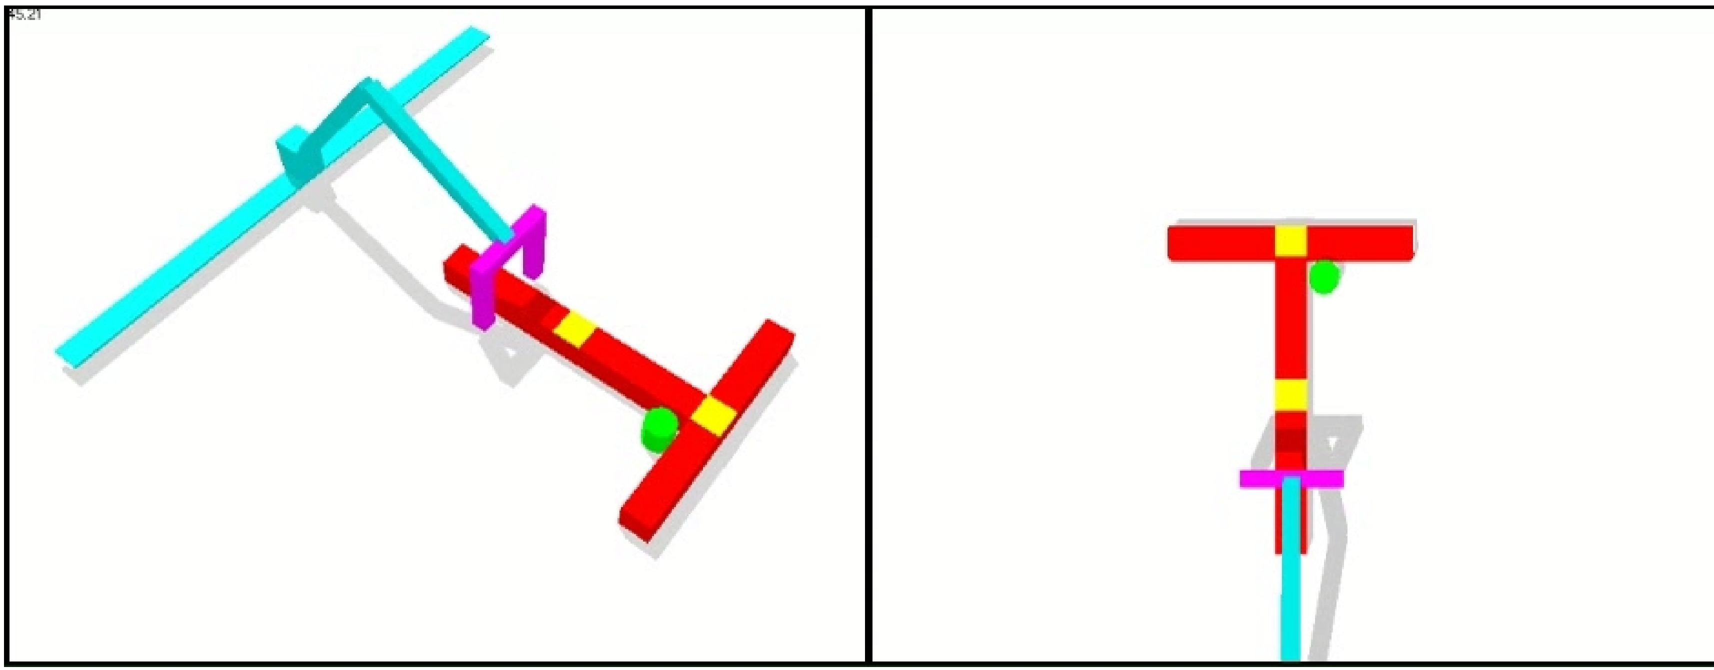

45.21

44.88

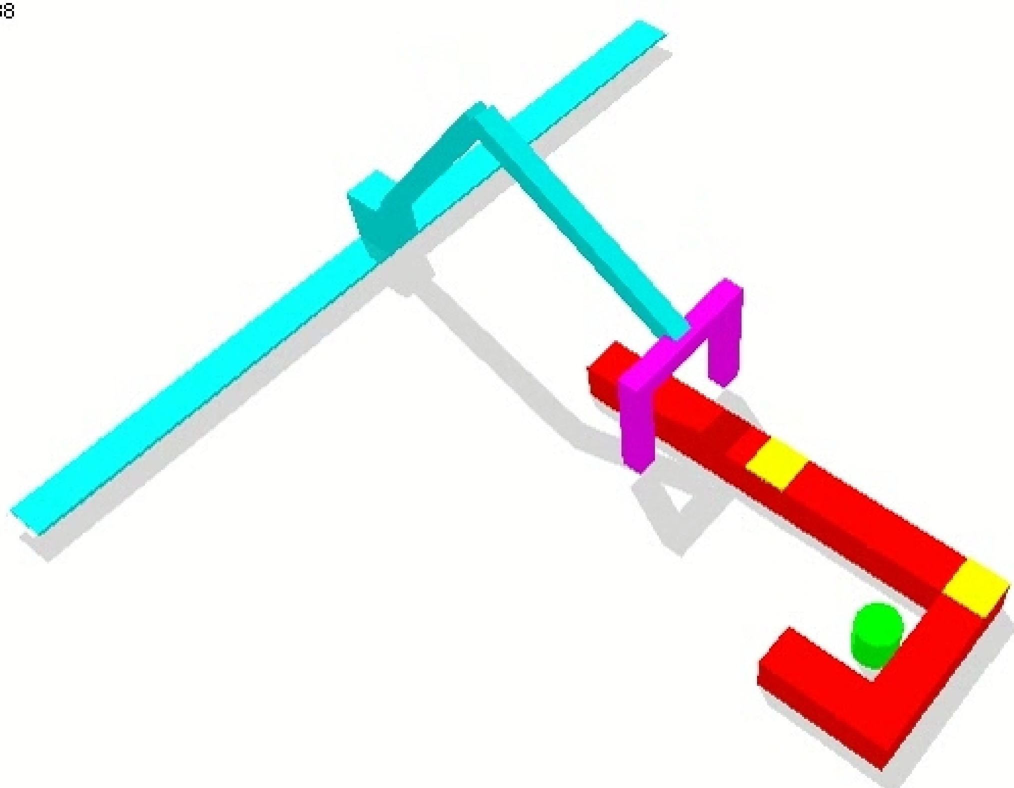

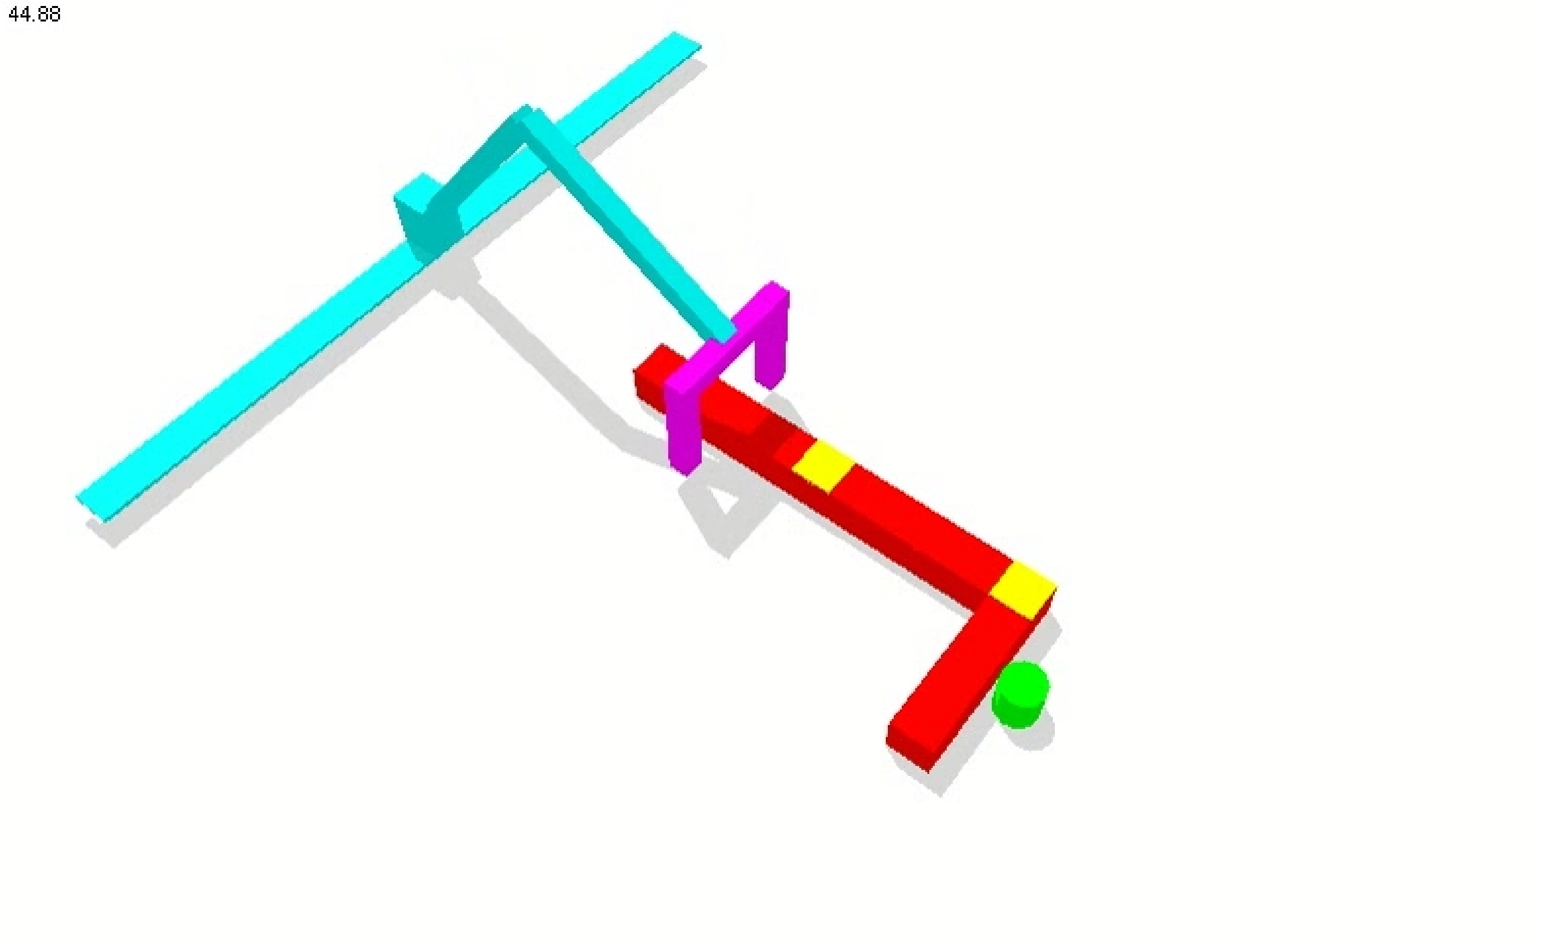
44.88

The robot’s action, A_t , consists of grabbing the tool and sliding it by *x* and *y* in the horizontal plane

The robot’s sensory input, S_t , is extracted from a camera overlooking the robot:

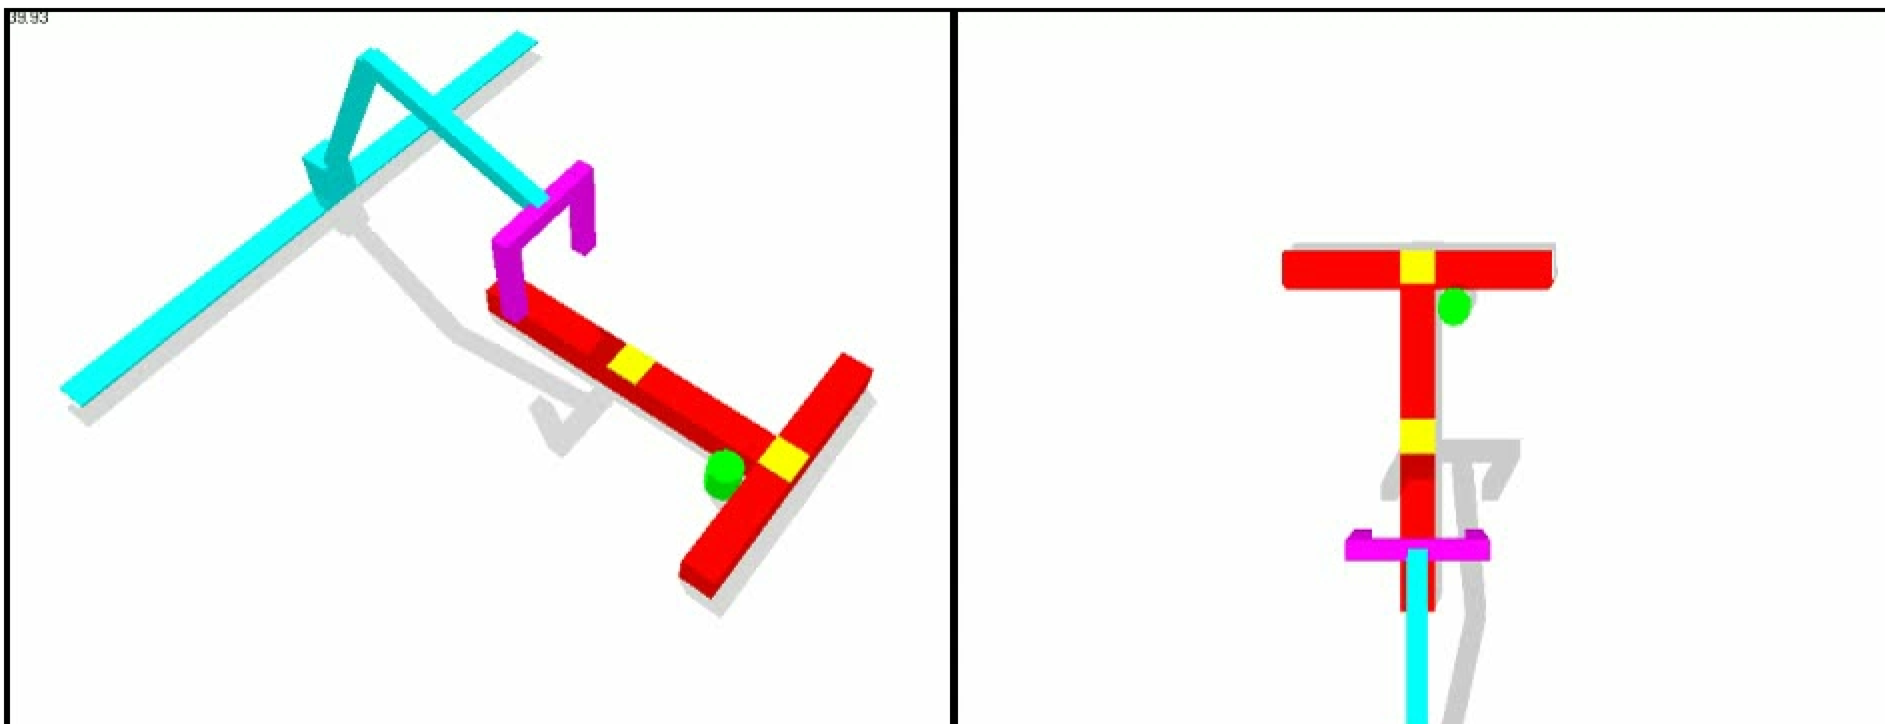

Change Detection

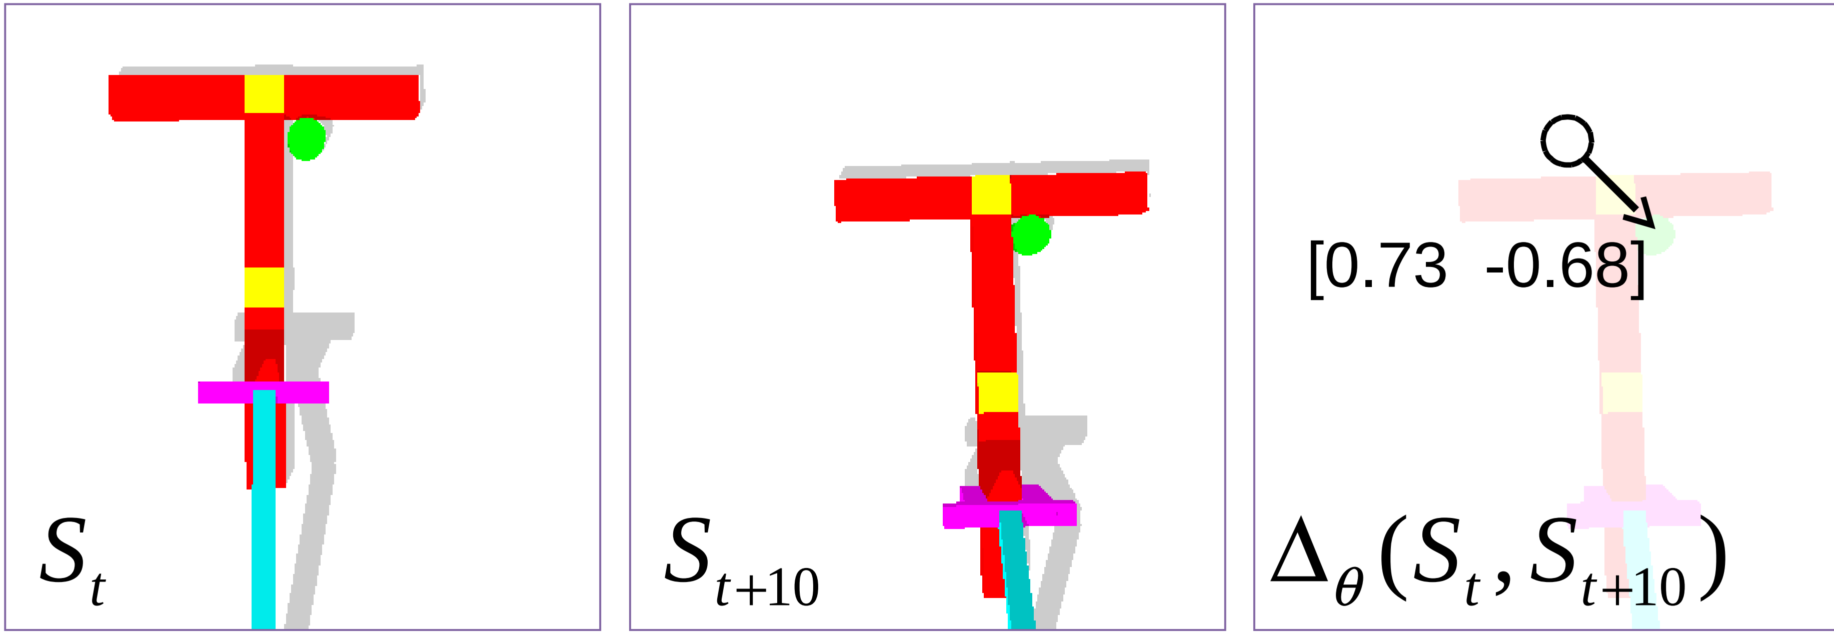

The task of the robot is to learn a predictive,model, M_θ such that given the robot's action and visual features, the model can predict the future outcomes of the action as measured by the change detection function:

$$M_\theta\left(A_t, \Phi_i\left(S_t\right)\right) \to \Delta_\theta\left(S_t, S_{t+k}\right)$$

Feature Extraction and Frames of Reference

Robot can extract features from its sensory input with a set of *perceptual functions*:

$$\Phi_1 \ldots \Phi_m \text{ such that } \Phi_i(S_t) \rightarrow U_t$$

$$\text{where } U_t = [u_1 \ldots u_k] \text{ and } k << n$$

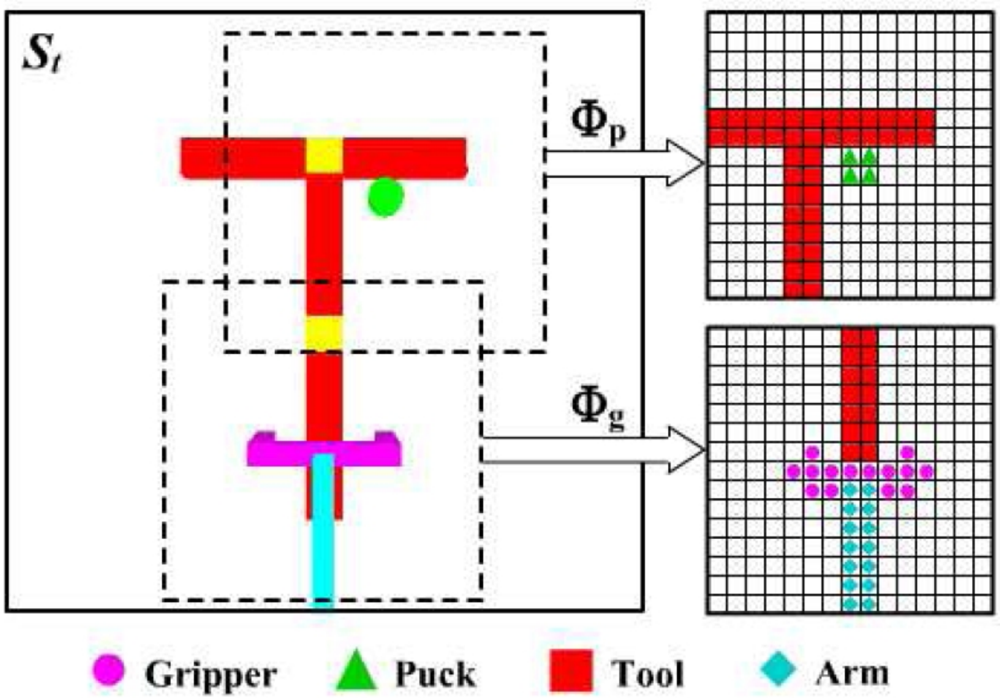

Feature Extraction and Frames of Reference

The robot extracts features from sensory input with the help of five *perceptual functions*:

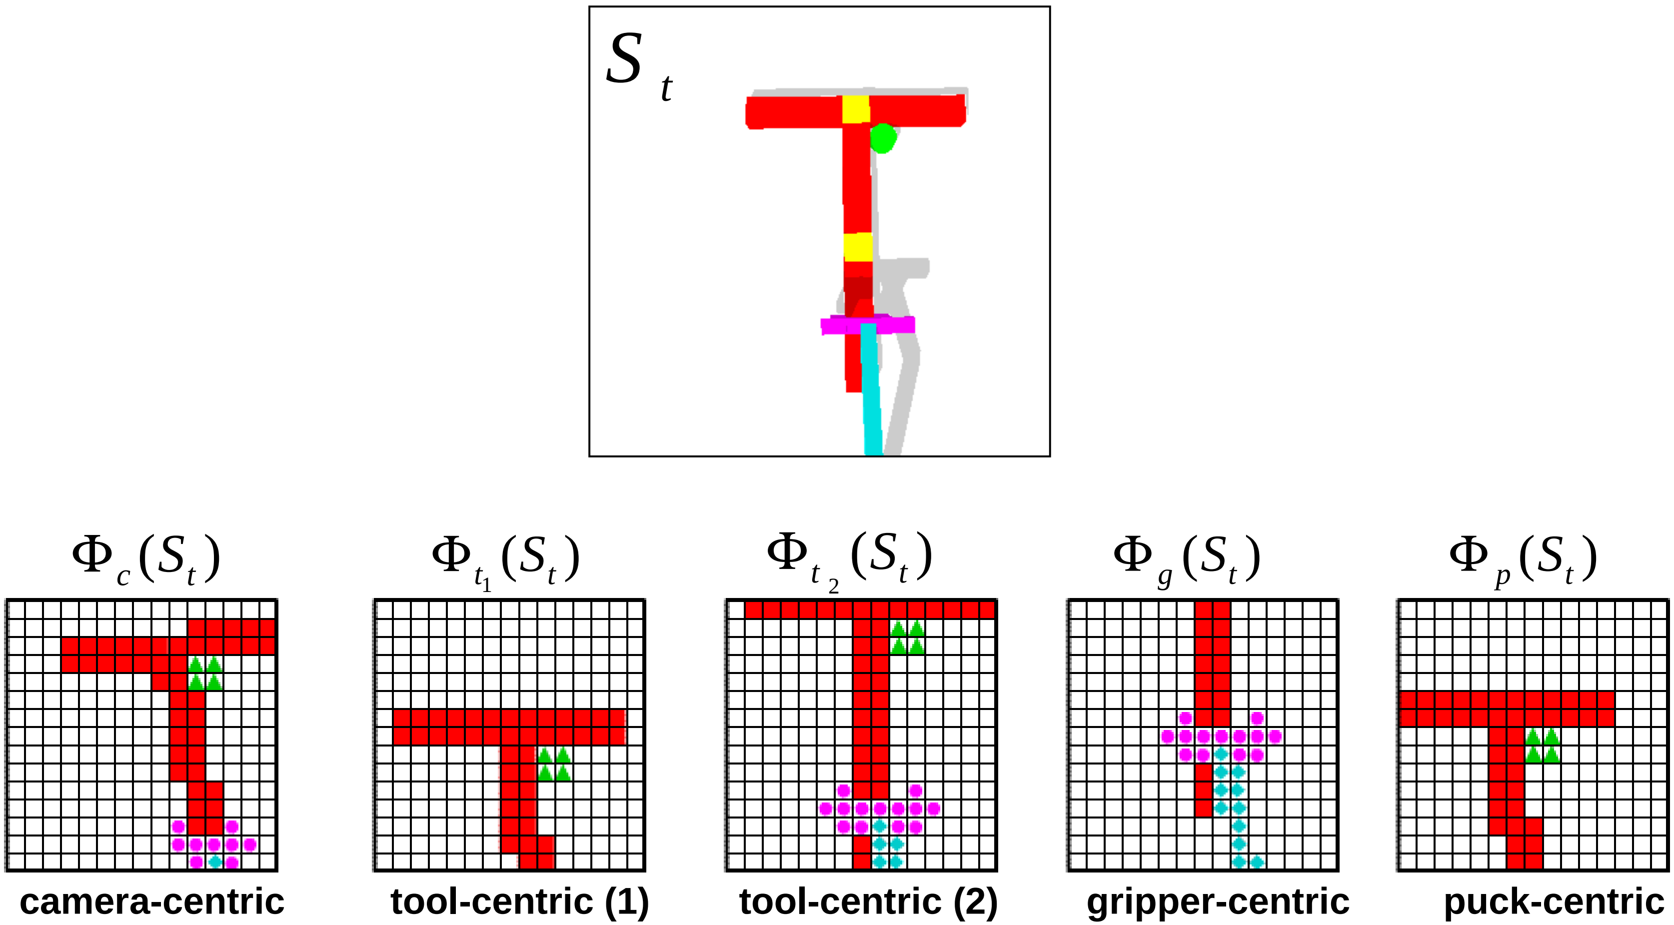

Behavior Babbling

The robot explores the tool through behavior-babbling. During this stage it collects data points of the form:

$$\left\langle A_t, \Phi_p\left(S_t\right), \Delta_\theta\left(S_t, S_{t+k}\right)\right\rangle$$

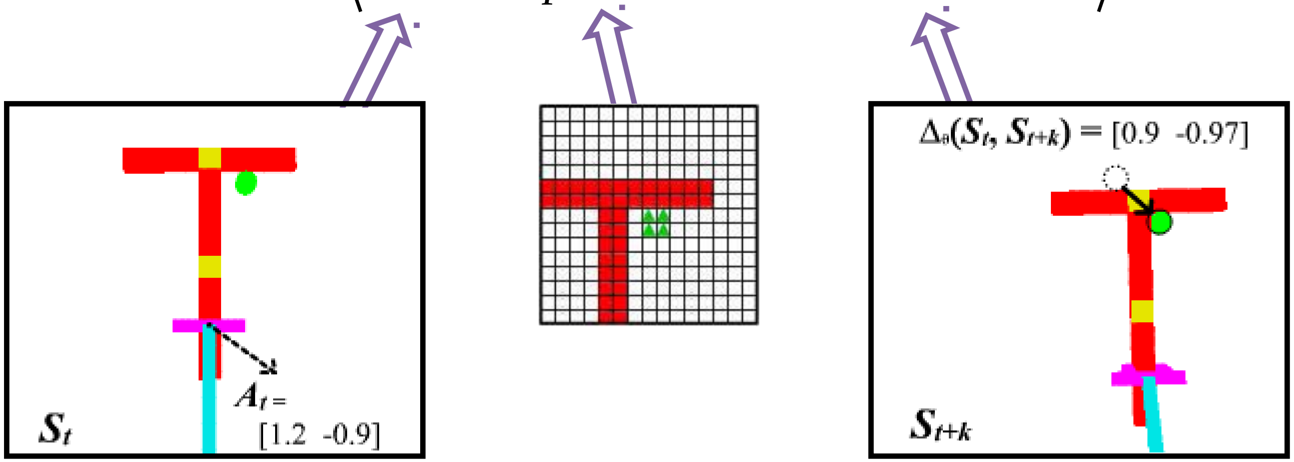

600 trials with each tool are performed, for a total of 3600. During each trial and the tool's and puck's starting positions, as well as the action's parameters are randomly chosen.

Verification

Prediction:

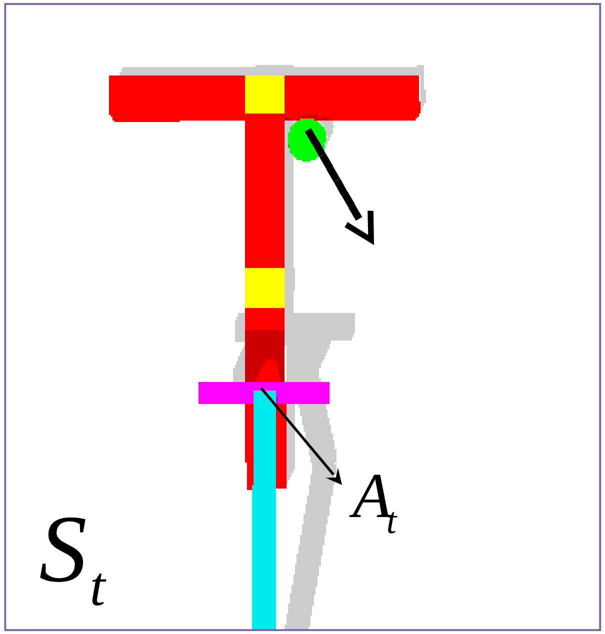

Observation:

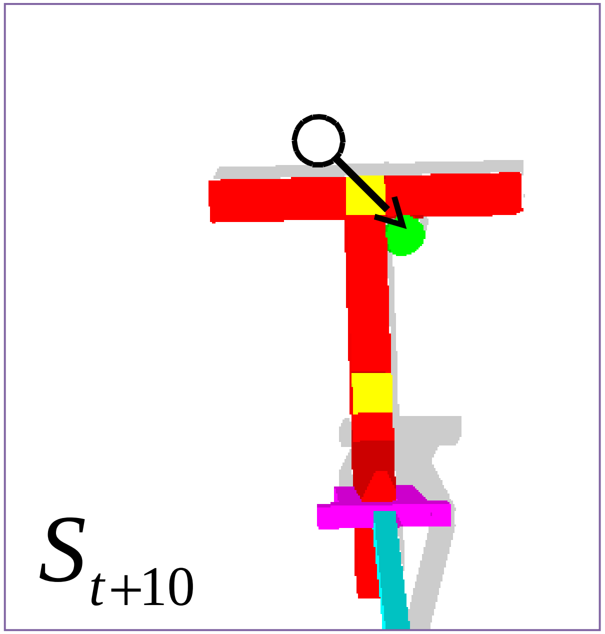

Evaluation

Initial State | After Executing Action | Prediction Evaluation

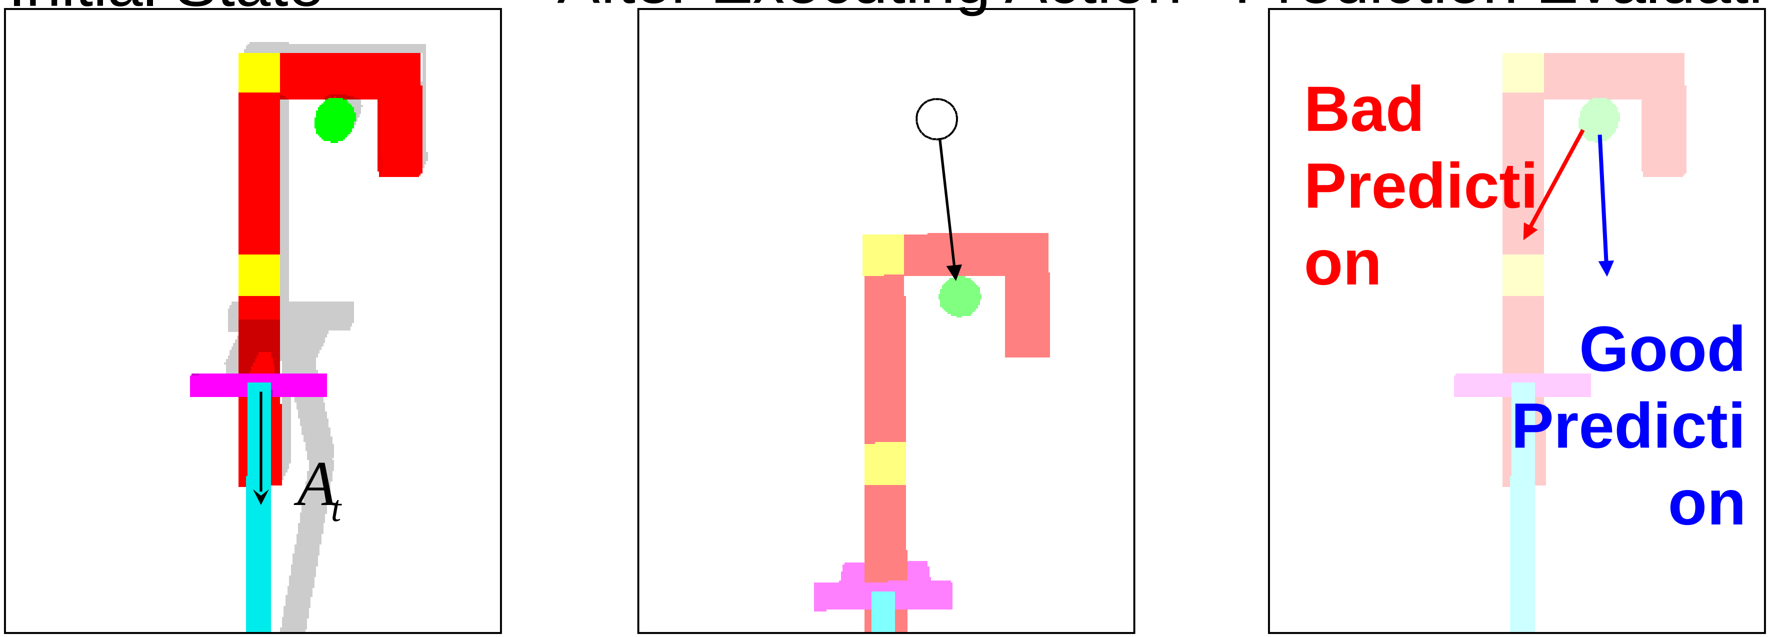

A prediction is *good* if the difference between the predicted and actual angles of the puck's motion is less than 20 degrees.

Evaluation Methodology

- Three Test Scenarios
  - Familiar Tool
  - Novel Tool
  - Larger Version of Familiar Tool
- Two learning algorithms:
  - k-Nearest Neighbors
  - Decision Tree with Linear Regression Leaf Nodes
- Five Frames of Reference

Results

Results

Best Frame of Reference with Decision Tree:

$$\Phi_p(S_t)$$

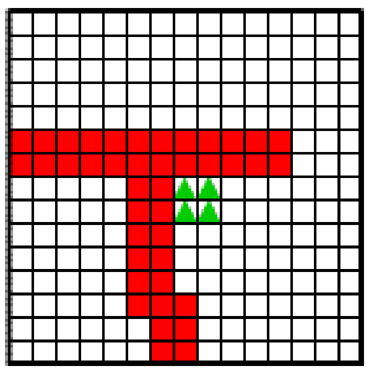

puck-centric

Performance is worst with circular shaped tool

Tool	% Good Predictions with Decision Tree
	94.4%
	92.7%
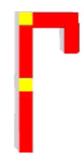	87.2%
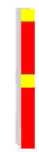	90.2%
	88.1%
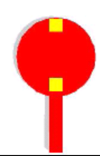	85.2%

TABLE I

PERCENTAGE OF GOOD PREDICTIONS FOR Δ_θ ACROSS DIFFERENT FRAMES OF REFERENCE (DECISION TREE ALGORITHM)

Tool	Frame of Reference				
	Φ_c	Φ_g	Φ_{t_1}	Φ_{t_2}	Φ_p
T-stick	51.3%	44.1%	71.4%	57.6%	**94.4%**
L-stick	40.6%	52.1%	78.8%	68.6%	**92.7%**
L-hook	47.9%	50.7%	67.3%	56.1%	**87.2%**
Stick	72.1%	69.3%	86.2%	85.8%	**90.2%**
T-hook	63.4%	43.2%	57.4%	49.4%	**88.1%**
Paddle	41.7%	50.5%	59.4%	53.7%	**85.2%**

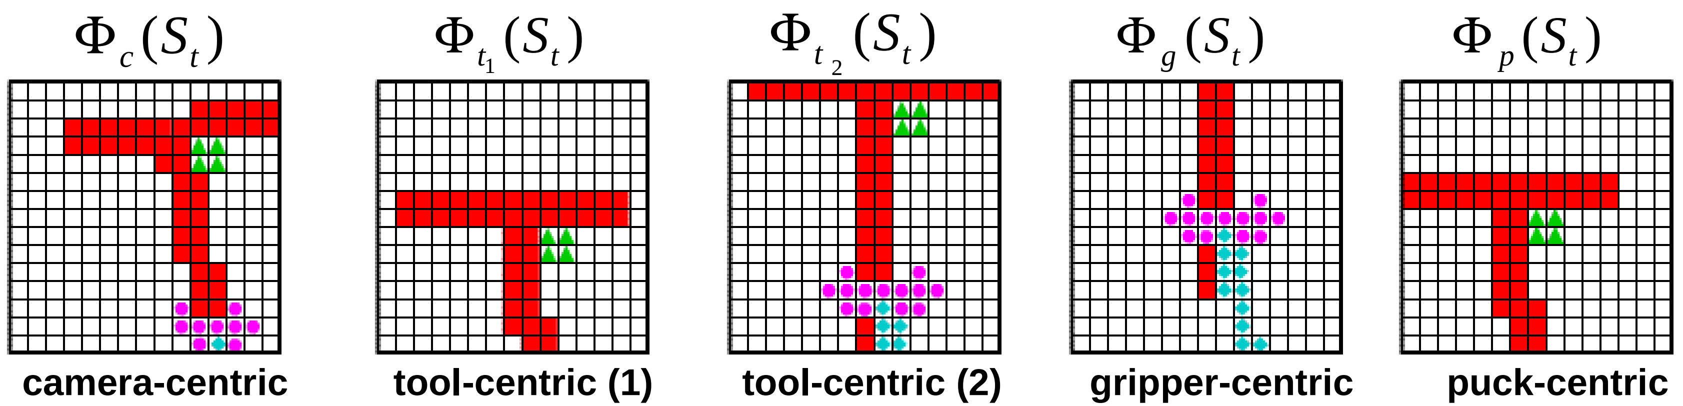

When does the model make mistakes?

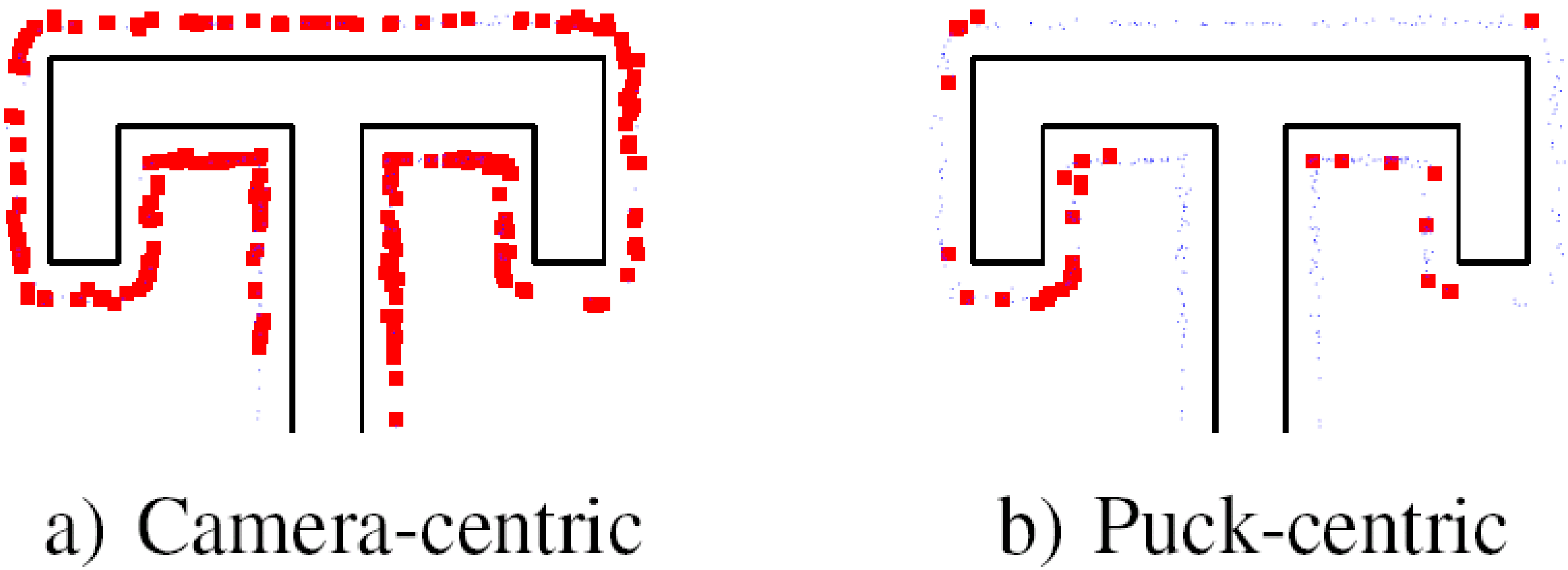

With puck-centric frame of reference, the errors are distributed around the corners of the tool, i.e. if the puck is positioned near a corner at the start of a trial, there is greater chance of error.

With the other frames of reference, the errors are distributed uniformly around the tool

Results with k-NN classifier

Best Frame of Reference with k-Nearest Neighbors:

$\Phi_{t_1}(S_t)$

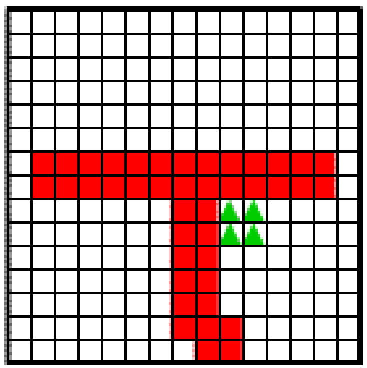

tool-centric (1)

Decision Tree model outperforms k-NN model on all tools

Tool	% Good Predictions with k-NN
	88.8%
	90.4%
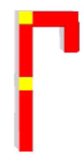	80.0%
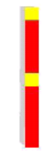	85.0%
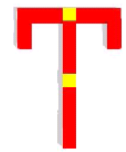	72.4%
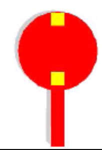	68.5%

Experience vs. Model Performance

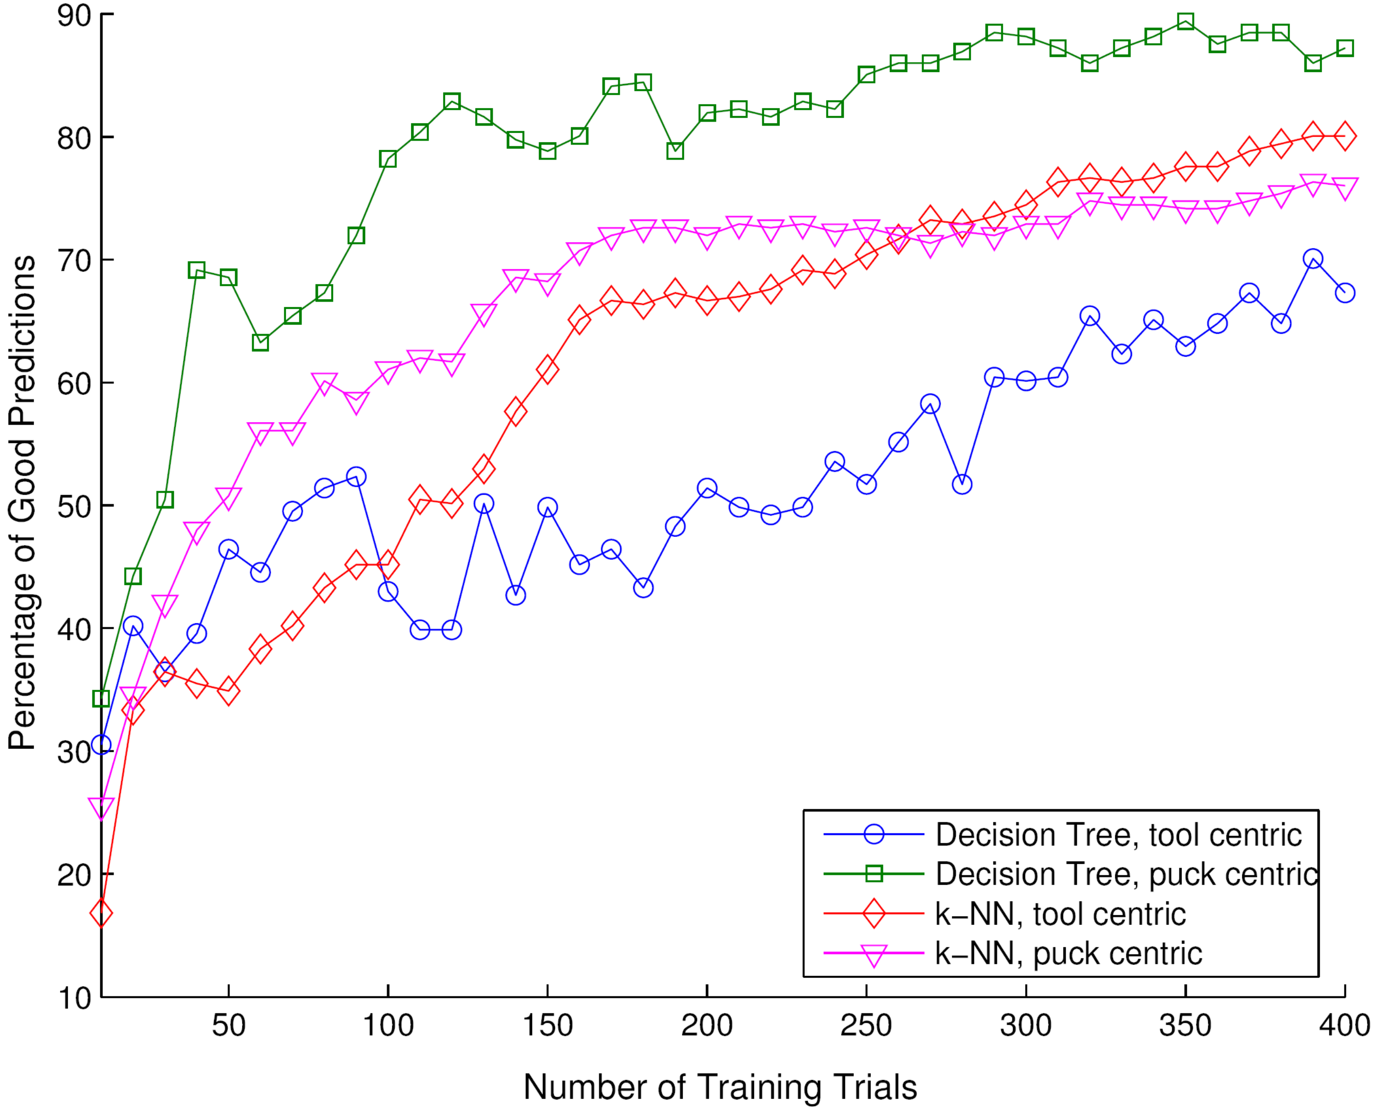

Evaluation on a Novel Tool

Train Tool	**Test Tool**	% Good Predictions	
		Decision Tree	k-NN
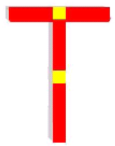		**86.6%**	61.1%
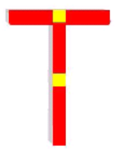	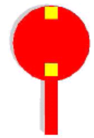	**76.7%**	45.6%
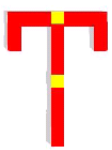		**85.4%**	58.7%
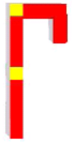	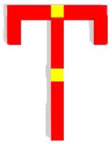	**80.0%**	50.9%
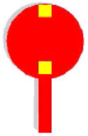		**44.2%**	19.9%

Evaluation on a Larger Tool

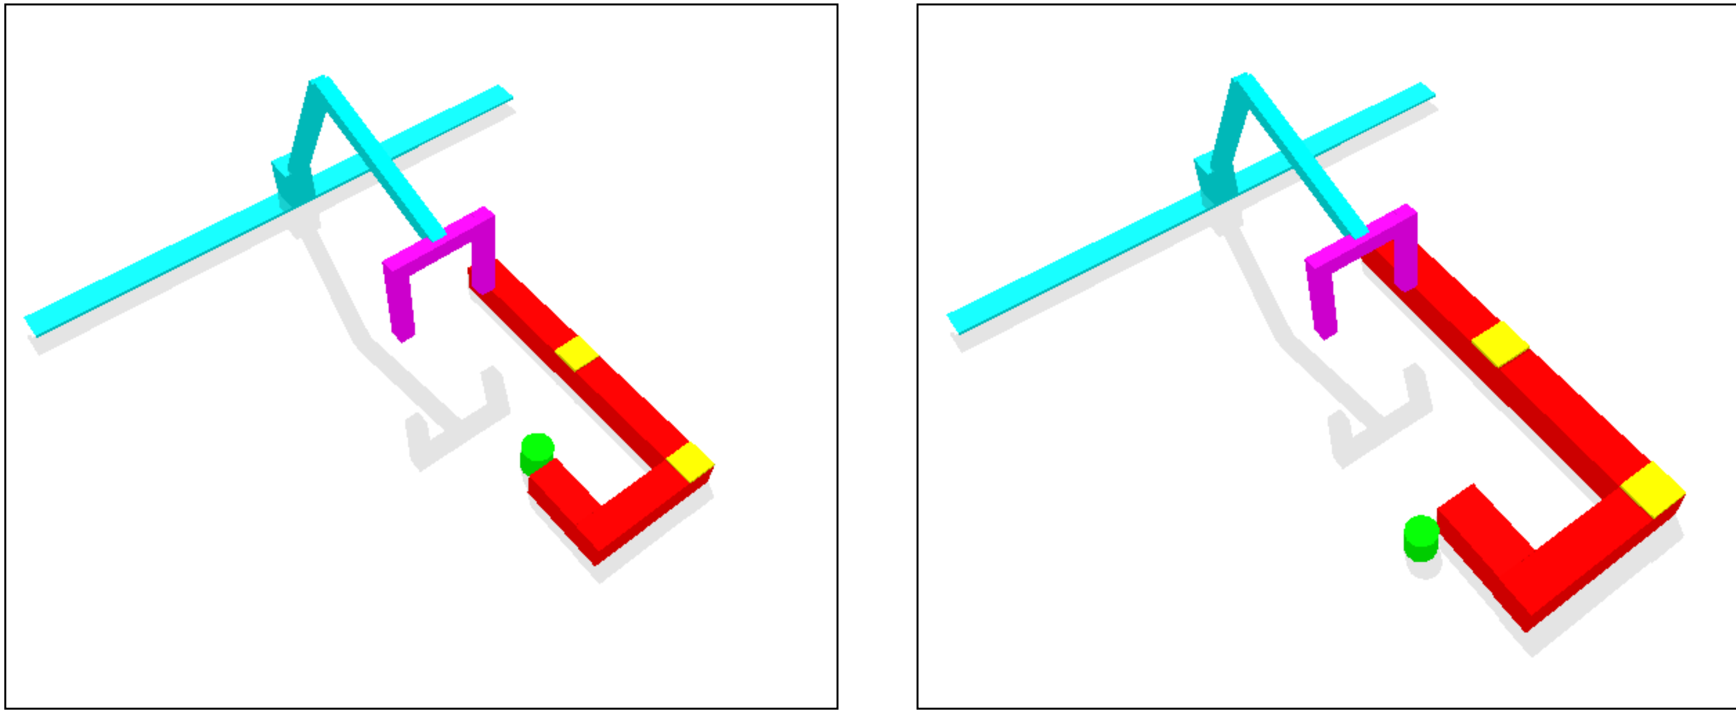

In this experiment the robot is trained on the L-hook tool, but then tested on a larger version of the tool

Decision tree with puck-centric frame of reference achieves the highest performance: 83.8% good predictions

Evaluation on a Larger Tool

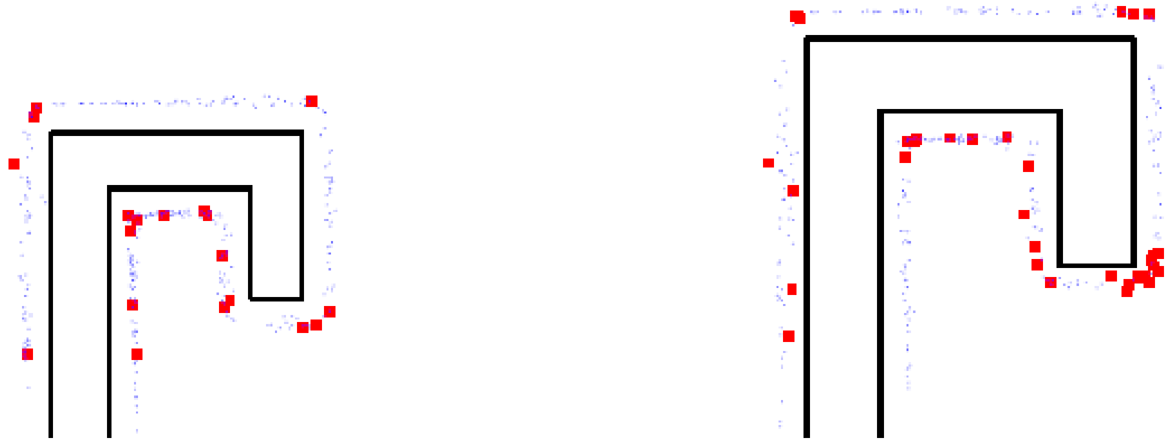

a) Regular tool b) Larger tool

Fig. 7. Visualization of the prediction errors made by decision tree model using a puck-centered frame of reference for the regular (left) and the enlarged (right) L-hook tools. In both scenarios, the model is trained on the regular-sized tool. Each point in the plot represents the puck's starting position relative to the tool during some particular trial. The points represented by the large squares indicate cases in which the prediction error is greater than 20°.

Questions and Discussion

- What are some of the limitations of this experiment? Would it work on a real robot?

- What are some cases where a robot may need to use a tool in our environment?

Detecting Functional Similarity Between Tools

- Jump to ICDL slides...

THE END