CS 188: Artificial Intelligence
Spring 2010

Advanced Applications: Robotics

Pieter Abbeel – UC Berkeley
A few slides from Sebastian Thrun, Dan Klein

Announcements

- Project 5 due Thursday --- Classification!

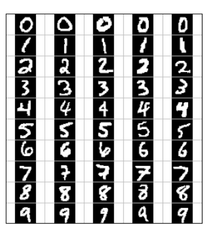

- Contest!!
  - Tournaments every night.
  - **Final tournament: We will use submissions received by Thursday May 6, 11pm.**

Estimation: Laplace Smoothing

- Laplace's estimate (extended):
  - Pretend you saw every outcome k extra times

$$P_{LAP,k}(x) = \frac{c(x) + k}{N + k|X|}$$

  - What's Laplace with k = 0?
  - k is the strength of the prior

- Laplace for conditionals:
  - Smooth each conditional independently:

$$P_{LAP,k}(x|y) = \frac{c(x,y) + k}{c(y) + k|X|}$$

H H T

$$P_{LAP,0}(X) =$$

$$P_{LAP,1}(X) =$$

$$P_{LAP,100}(X) =$$

So Far: Foundational Methods

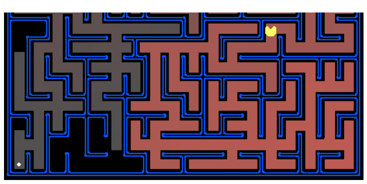

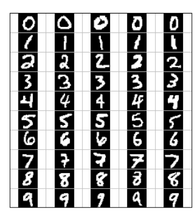

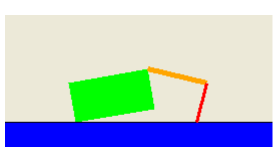

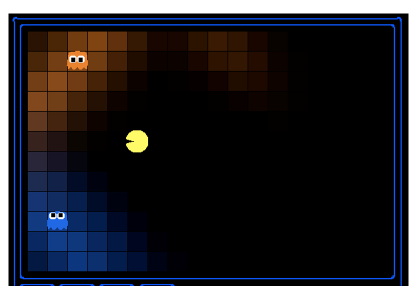

4

Now: Advanced Applications

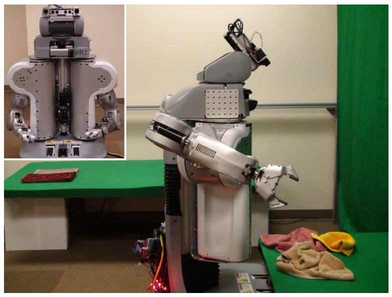


5

Robotic Control Tasks

- Perception / Tracking
  - Where exactly am I?
  - What's around me?
- Low-Level Control
  - How to move the robot and/or objects from position A to position B
- High-Level Control
  - What are my goals?
  - What are the optimal high-level actions?

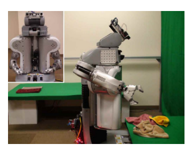


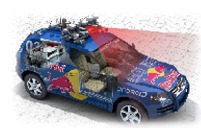

Robot folds towels

- [pile of 5 video]

[Maitin-Shepard, Cusumano-Towner, Lei & Abbeel, 2010]

Low-Level Planning

- Low-level: move from configuration A to configuration B

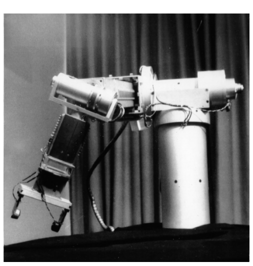

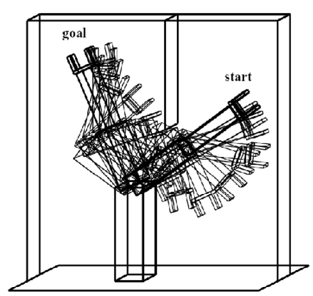

A Simple Robot Arm

- Configuration Space
  - What are the natural coordinates for specifying the robot's configuration?
  - These are the *configuration space* coordinates
  - Can't necessarily control all degrees of freedom directly
- Work Space
  - What are the natural coordinates for specifying the effector tip's position?
  - These are the *work space* coordinates

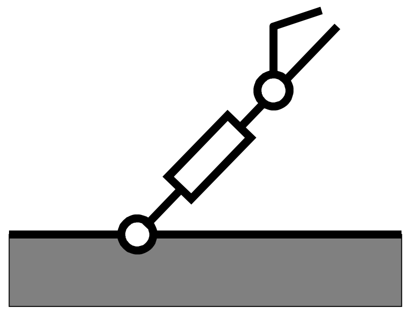

Coordinate Systems

- Workspace:
  - The world's (x, y) system
  - Obstacles specified here
- Configuration space
  - The robot's state
  - Planning happens here
  - Obstacles can be projected to here

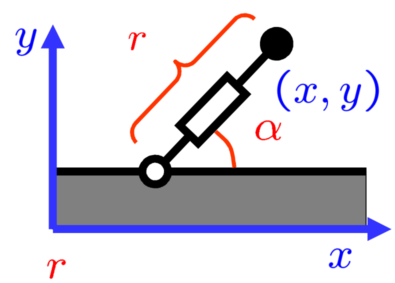

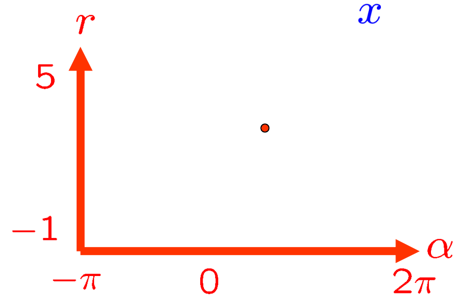

Obstacles in C-Space

- What / where are the obstacles?
- Remaining space is *free space*

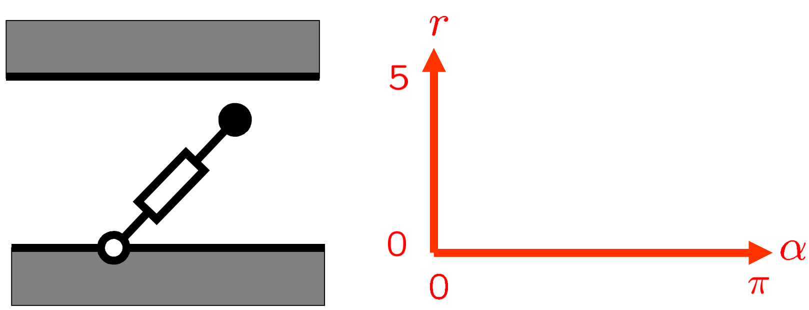

Two-link manipulator

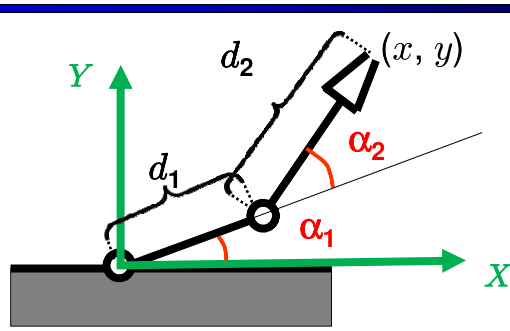

$$x = d_1 \cos\alpha_1 + d_2 \cos(\alpha_1 + \alpha_2)$$
$$y = d_1 \sin\alpha_1 + d_2 \sin(\alpha_1 + \alpha_2)$$

Example Obstacles in C-Space

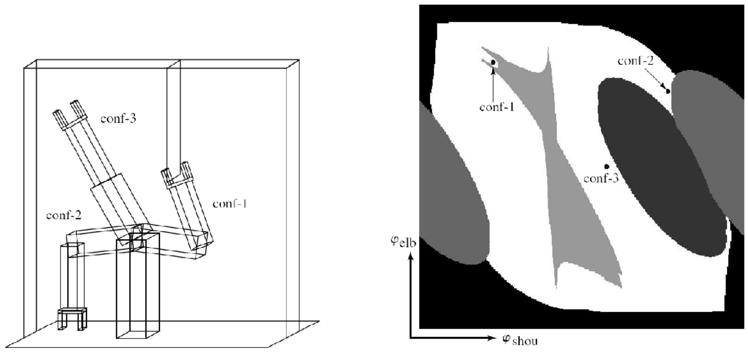

Two-link manipulator

- Demo

http://www-inst.eecs.berkeley.edu/~cs188/fa08/demos/robot.html

Probabilistic Roadmaps

- Idea: sample random points as nodes in a visibility graph
- This gives *probabilistic roadmaps*
  - Very successful in practice
  - Lets you add points where you need them
  - If insufficient points, incomplete or weird paths

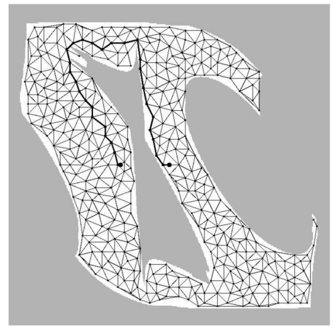

Perception

1. Find a point see in two camera views
2. Find 3D coordinates by finding the intersection of the rays

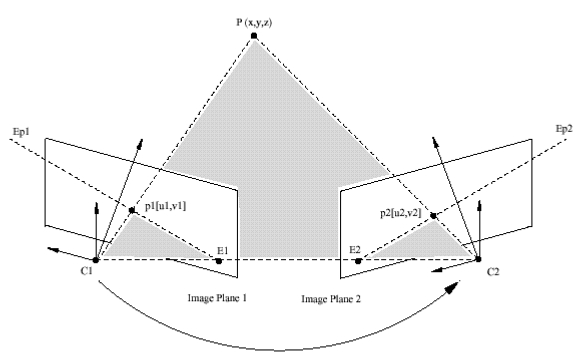

18

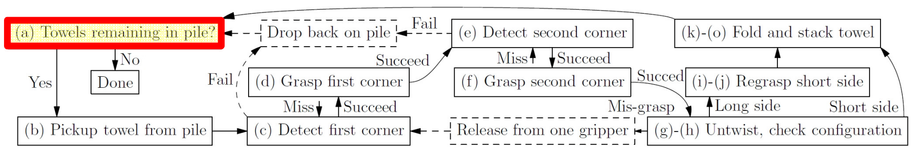
(a) Towels remaining in pile?
Drop back on pile
Fail
(e) Detect second corner
(k)-(o) Fold and stack towel
No
Done
Yes
Fail
Succeed
(d) Grasp first corner
Miss
Succeed
(f) Grasp second corner
Succeed
(i)-(j) Regrasp short side
Miss
Succeed
Mis-grasp
Long side
Short side
(b) Pickup towel from pile
(c) Detect first corner
Release from one gripper
(g)-(h) Untwist, check configuration

19

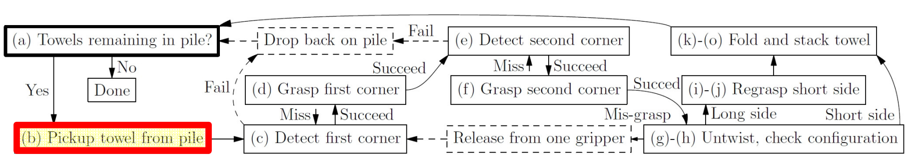
(a) Towels remaining in pile?
Drop back on pile
Fail
(e) Detect second corner
(k)-(o) Fold and stack towel
No
Done
Yes
Fail
Succeed
(d) Grasp first corner
Miss
Succeed
(f) Grasp second corner
Succeed
(i)-(j) Regrasp short side
Miss
Succeed
Mis-grasp
Long side
Short side
(b) Pickup towel from pile
(c) Detect first corner
Release from one gripper
(g)-(h) Untwist, check configuration

20

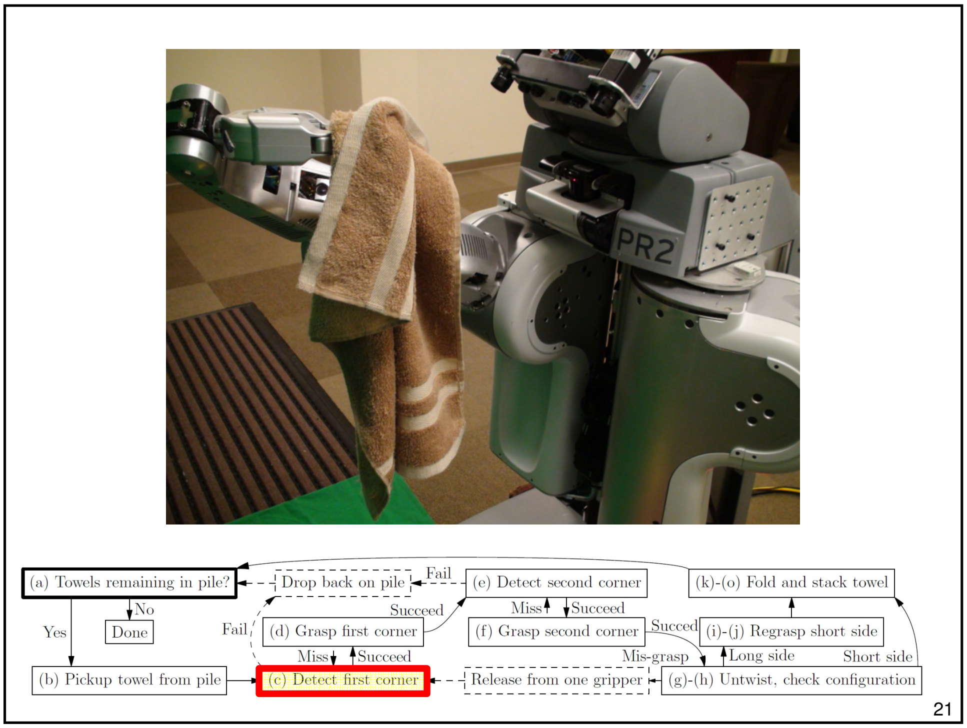
PR2
(a) Towels remaining in pile?
Drop back on pile
Fail
(e) Detect second corner
(k)-(o) Fold and stack towel
No
Done
Yes
Fail
(d) Grasp first corner
Succeed
Miss
Succeed
(f) Grasp second corner
Succeed
(i)-(j) Regrasp short side
Miss
Succeed
Mis-grasp
Long side
Short side
(b) Pickup towel from pile
(c) Detect first corner
Release from one gripper
(g)-(h) Untwist, check configuration
21

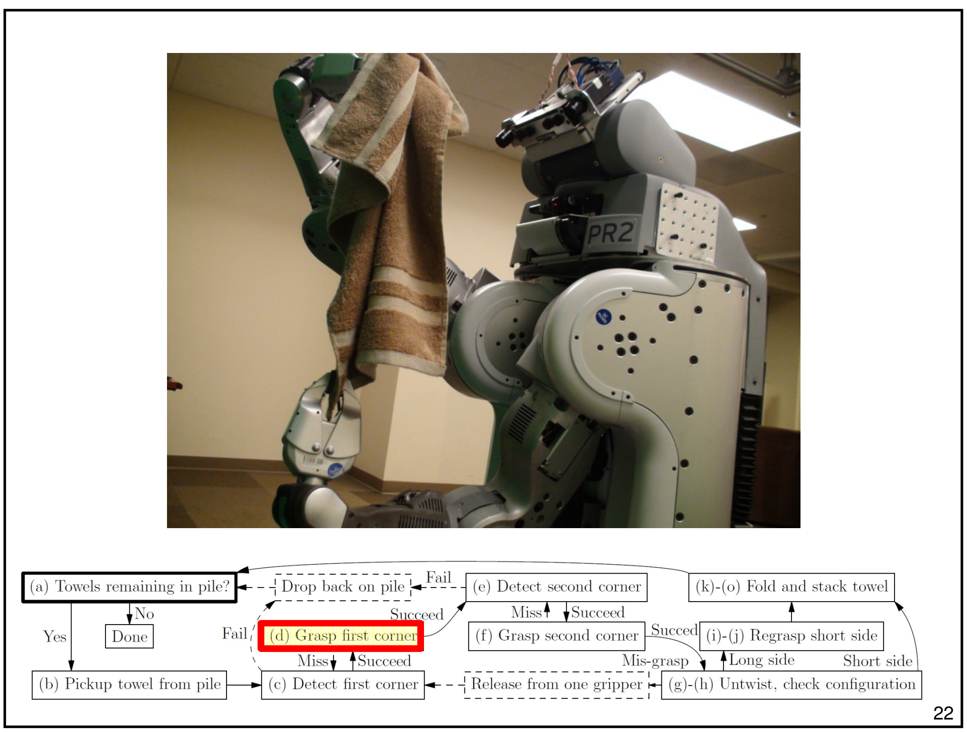
PR2
(a) Towels remaining in pile?
Drop back on pile
Fail
(e) Detect second corner
(k)-(o) Fold and stack towel
No
Done
Yes
Fail
(d) Grasp first corner
Succeed
Miss
Succeed
(f) Grasp second corner
Succeed
(i)-(j) Regrasp short side
Miss
Succeed
Mis-grasp
Long side
Short side
(b) Pickup towel from pile
(c) Detect first corner
Release from one gripper
(g)-(h) Untwist, check configuration
22

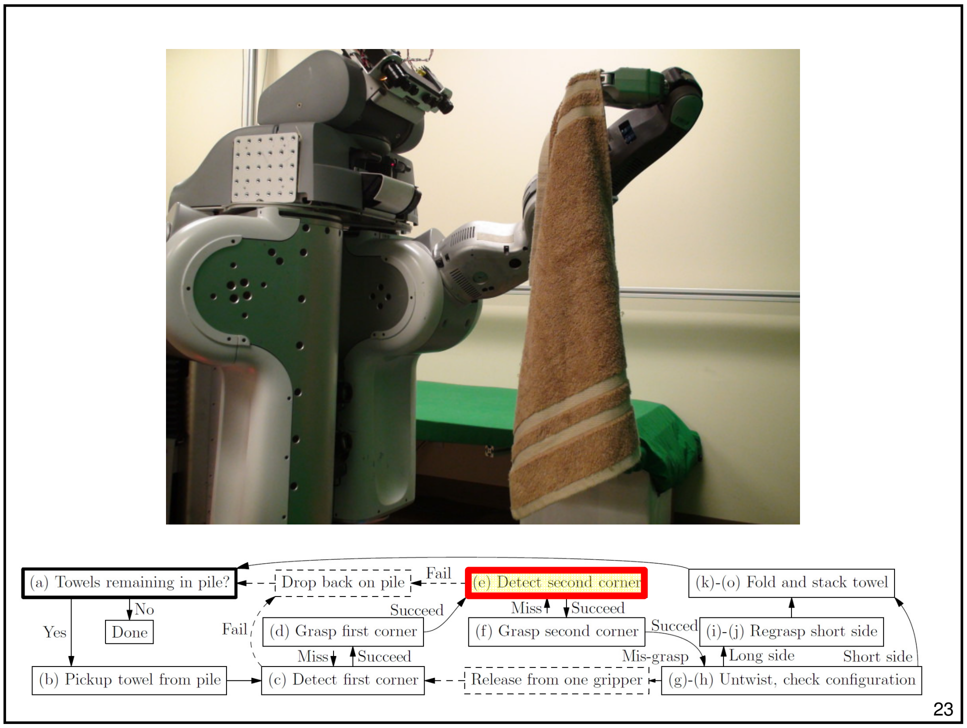
(a) Towels remaining in pile?
Drop back on pile
Fail
(e) Detect second corner
(k)-(o) Fold and stack towel
No
Done
Yes
Succeed
Miss
Succeed
Fail
(d) Grasp first corner
(f) Grasp second corner
Succeed
(i)-(j) Regrasp short side
Miss
Succeed
Mis-grasp
Long side
Short side
(b) Pickup towel from pile
(c) Detect first corner
Release from one gripper
(g)-(h) Untwist, check configuration
23

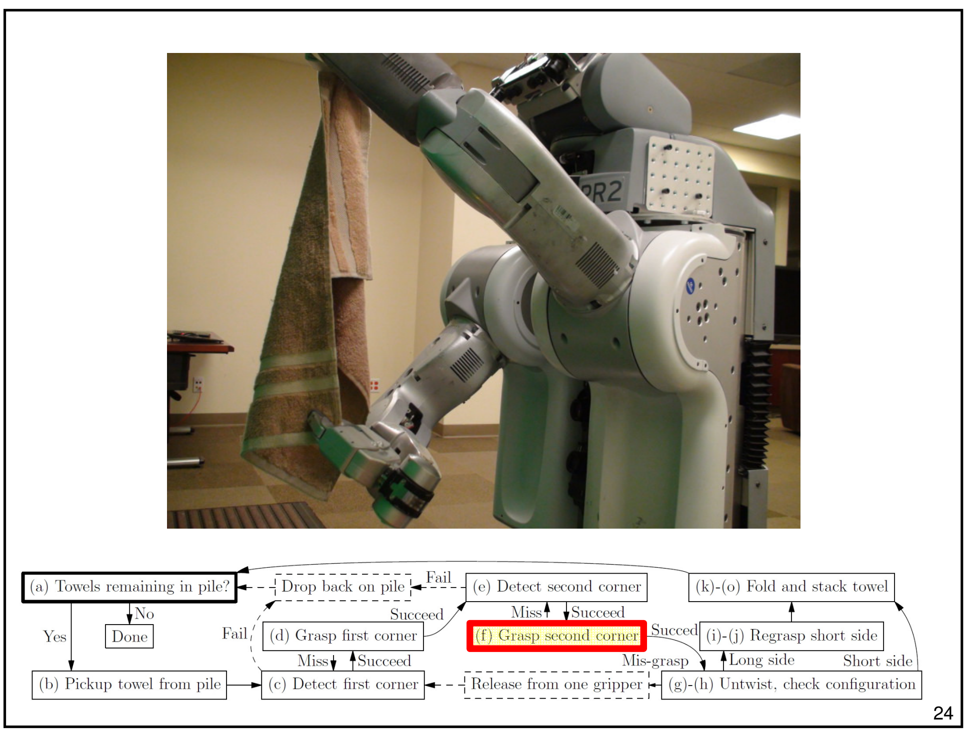
PR2
(a) Towels remaining in pile?
Drop back on pile
Fail
(e) Detect second corner
(k)-(o) Fold and stack towel
No
Done
Yes
Succeed
Miss
Succeed
Fail
(d) Grasp first corner
(f) Grasp second corner
Succeed
(i)-(j) Regrasp short side
Miss
Succeed
Mis-grasp
Long side
Short side
(b) Pickup towel from pile
(c) Detect first corner
Release from one gripper
(g)-(h) Untwist, check configuration
24

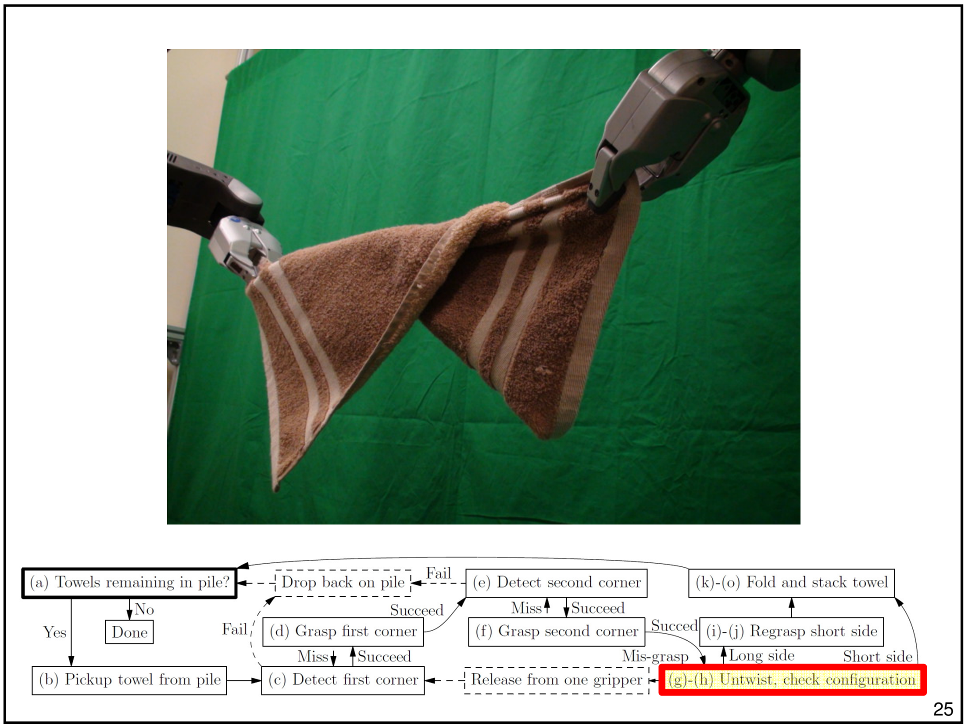
(a) Towels remaining in pile?
Drop back on pile
Fail
(e) Detect second corner
(k)-(o) Fold and stack towel
No
Done
Yes
Succeed
Fail
(d) Grasp first corner
Miss
Succeed
(f) Grasp second corner
Succeed
(i)-(j) Regrasp short side
Miss
Succeed
Mis-grasp
Long side
Short side
(b) Pickup towel from pile
(c) Detect first corner
Release from one gripper
(g)-(h) Untwist, check configuration
25

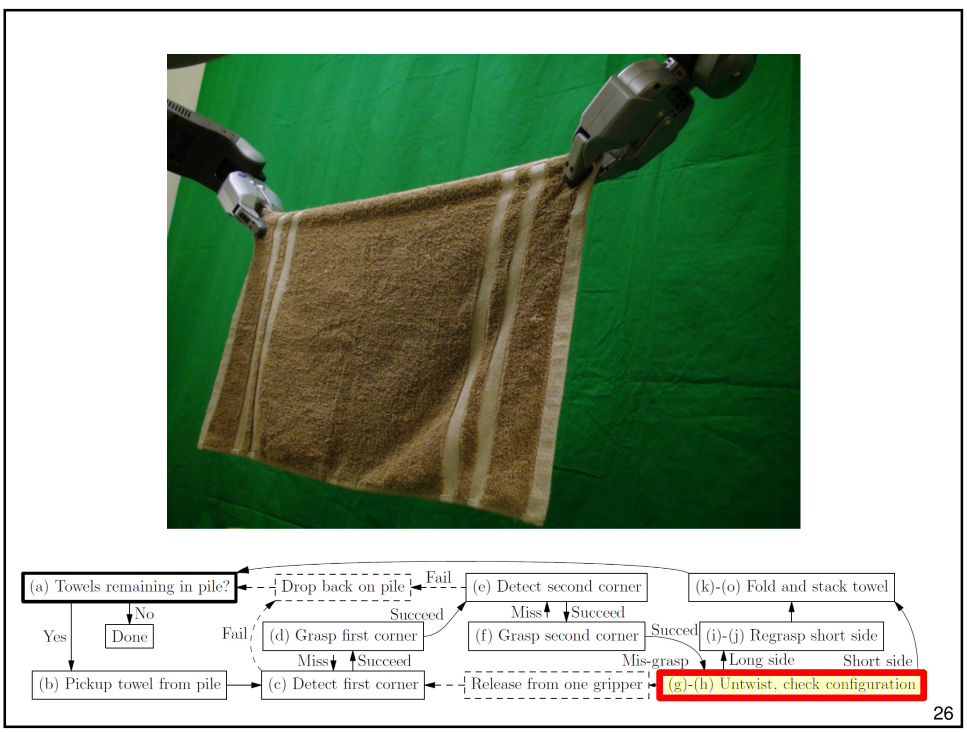
(a) Towels remaining in pile?
Drop back on pile
Fail
(e) Detect second corner
(k)-(o) Fold and stack towel
No
Done
Yes
Succeed
Fail
(d) Grasp first corner
Miss
Succeed
(f) Grasp second corner
Succeed
(i)-(j) Regrasp short side
Miss
Succeed
Mis-grasp
Long side
Short side
(b) Pickup towel from pile
(c) Detect first corner
Release from one gripper
(g)-(h) Untwist, check configuration
26

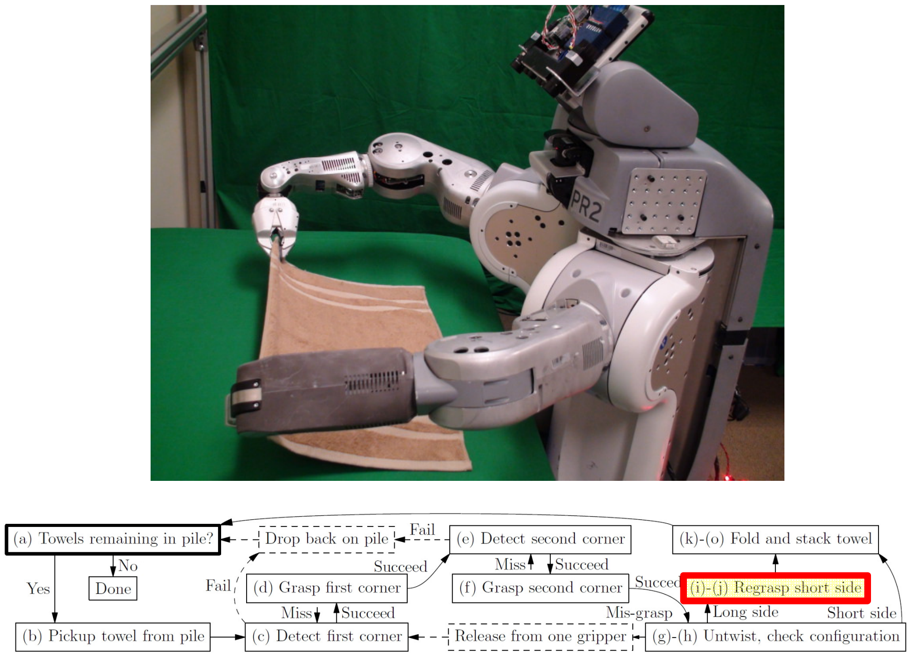
(a) Towels remaining in pile?
Drop back on pile
Fail
(e) Detect second corner
(k)-(o) Fold and stack towel
No
Succeed
Miss
Succeed
Yes
Done
Fail
(d) Grasp first corner
(f) Grasp second corner
Succeed
(i)-(j) Regrasp short side
Miss
Succeed
Mis-grasp
Long side
Short side
(b) Pickup towel from pile
(c) Detect first corner
Release from one gripper
(g)-(h) Untwist, check configuration

27

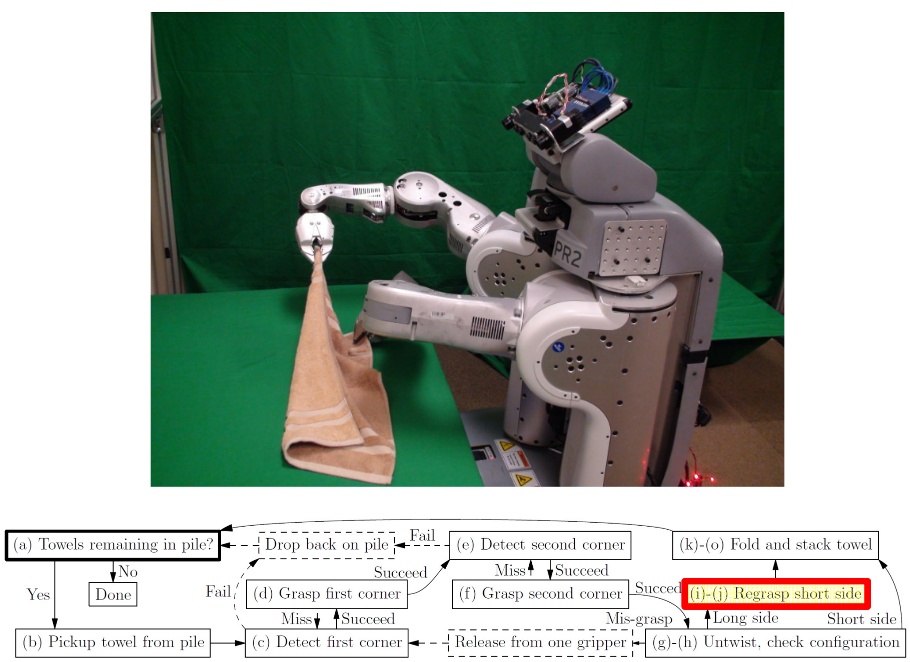
(a) Towels remaining in pile?
Drop back on pile
Fail
(e) Detect second corner
(k)-(o) Fold and stack towel
No
Succeed
Miss
Succeed
Yes
Done
Fail
(d) Grasp first corner
(f) Grasp second corner
Succeed
(i)-(j) Regrasp short side
Miss
Succeed
Mis-grasp
Long side
Short side
(b) Pickup towel from pile
(c) Detect first corner
Release from one gripper
(g)-(h) Untwist, check configuration

28

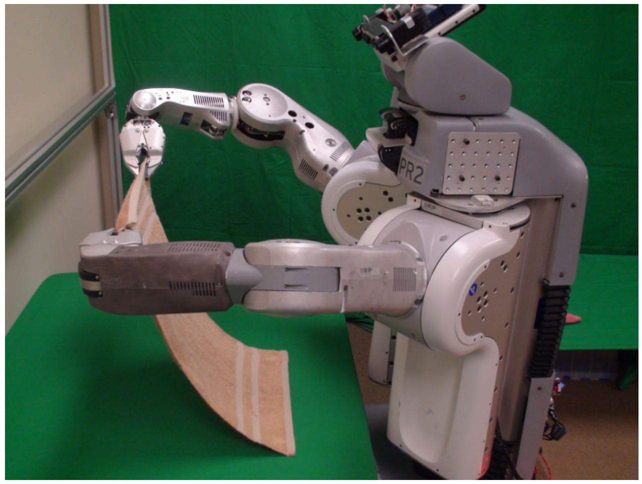

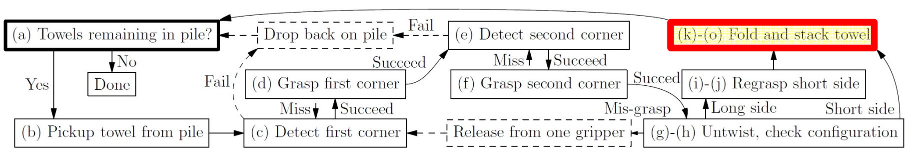
(a) Towels remaining in pile?
Drop back on pile
Fail
(e) Detect second corner
(k)-(o) Fold and stack towel
No
Done
Yes
Succeed
Fail
(d) Grasp first corner
Miss
Succeed
(f) Grasp second corner
Succeed
(i)-(j) Regrasp short side
Miss
Succeed
Mis-grasp
Long side
Short side
(b) Pickup towel from pile
(c) Detect first corner
Release from one gripper
(g)-(h) Untwist, check configuration

29

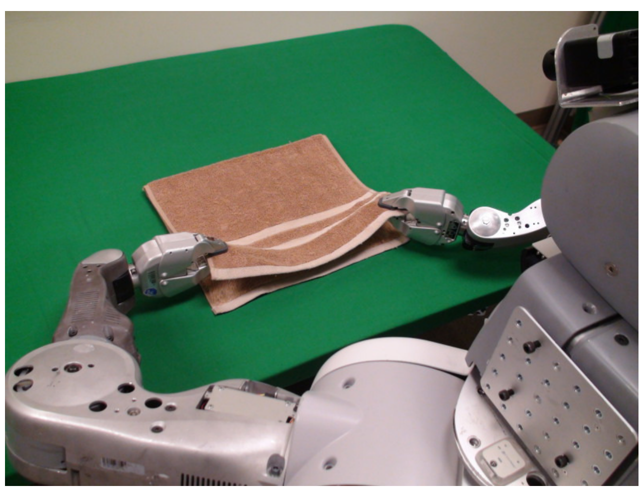

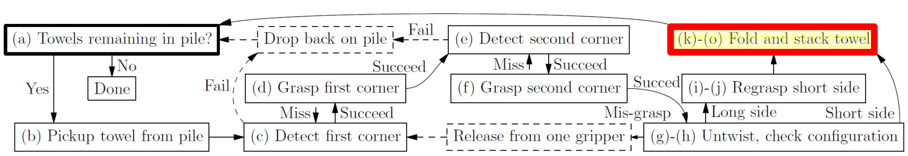
(a) Towels remaining in pile?
Drop back on pile
Fail
(e) Detect second corner
(k)-(o) Fold and stack towel
No
Done
Yes
Succeed
Fail
(d) Grasp first corner
Miss
Succeed
(f) Grasp second corner
Succeed
(i)-(j) Regrasp short side
Miss
Succeed
Mis-grasp
Long side
Short side
(b) Pickup towel from pile
(c) Detect first corner
Release from one gripper
(g)-(h) Untwist, check configuration

30

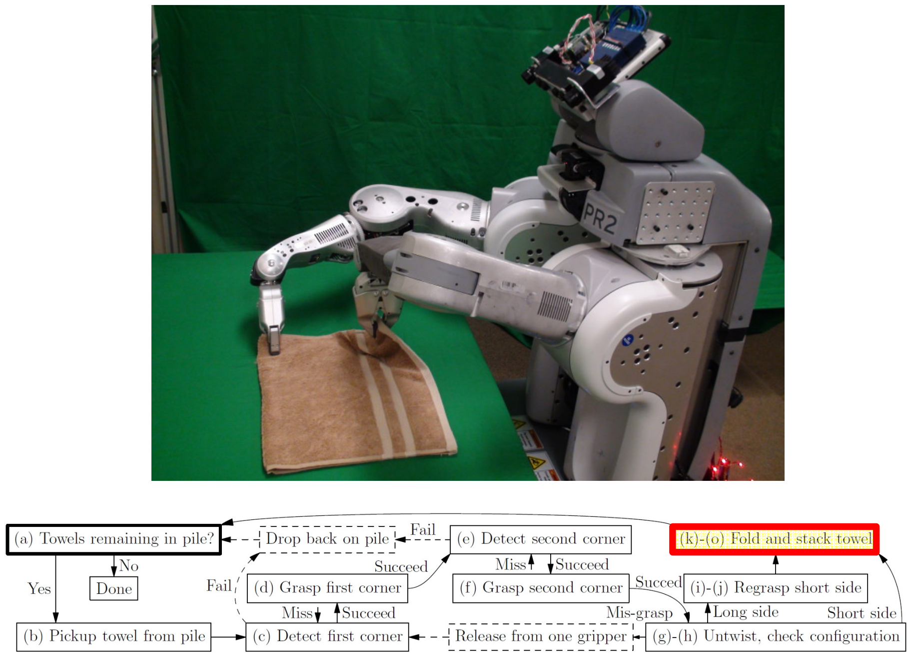
PR2
(a) Towels remaining in pile?
Drop back on pile
Fail
(e) Detect second corner
(k)-(o) Fold and stack towel
No
Done
Yes
Succeed
Miss
Succeed
Fail
(d) Grasp first corner
(f) Grasp second corner
Succeed
(i)-(j) Regrasp short side
Miss
Succeed
Mis-grasp
Long side
Short side
(b) Pickup towel from pile
(c) Detect first corner
Release from one gripper
(g)-(h) Untwist, check configuration

31

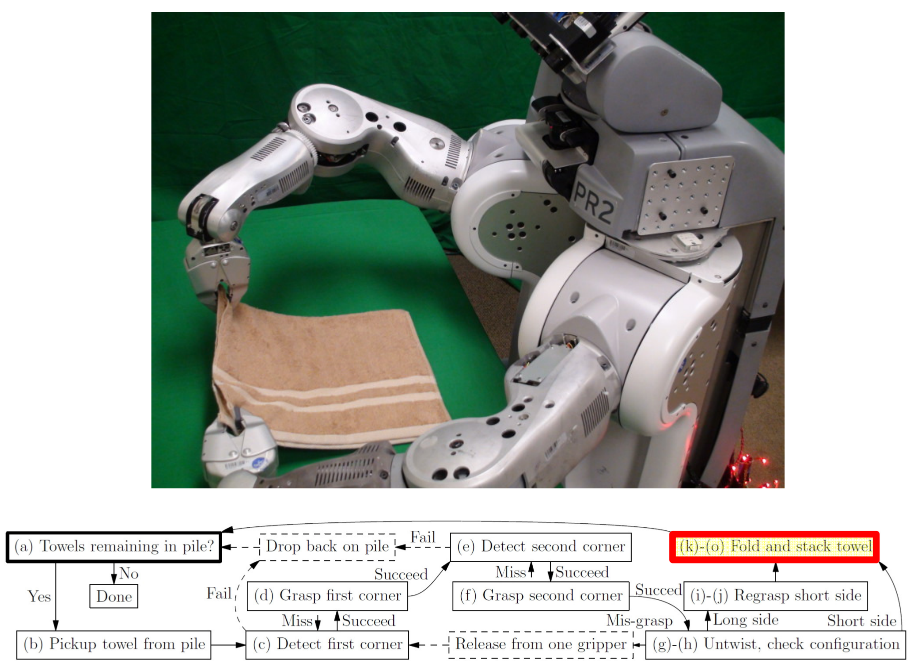
PR2
(a) Towels remaining in pile?
Drop back on pile
Fail
(e) Detect second corner
(k)-(o) Fold and stack towel
No
Done
Yes
Succeed
Miss
Succeed
Fail
(d) Grasp first corner
(f) Grasp second corner
Succeed
(i)-(j) Regrasp short side
Miss
Succeed
Mis-grasp
Long side
Short side
(b) Pickup towel from pile
(c) Detect first corner
Release from one gripper
(g)-(h) Untwist, check configuration

32

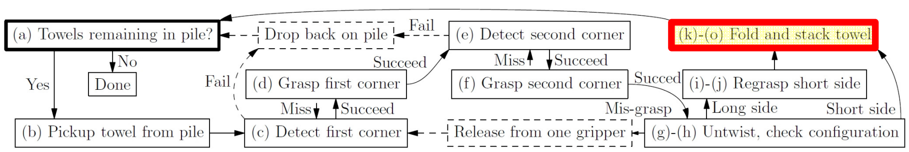
(a) Towels remaining in pile?
Drop back on pile
Fail
(e) Detect second corner
(k)-(o) Fold and stack towel
No
Succeed
Miss
Succeed
Yes
Done
Fail
(d) Grasp first corner
(f) Grasp second corner
Succeed
(i)-(j) Regrasp short side
Miss
Succeed
Mis-grasp
Long side
Short side
(b) Pickup towel from pile
(c) Detect first corner
Release from one gripper
(g)-(h) Untwist, check configuration

33

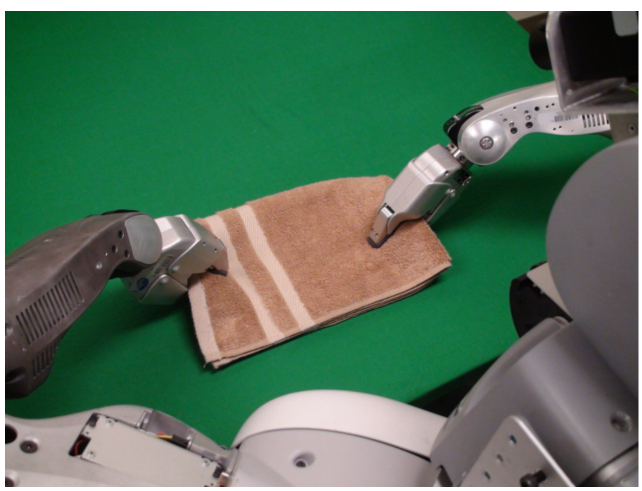

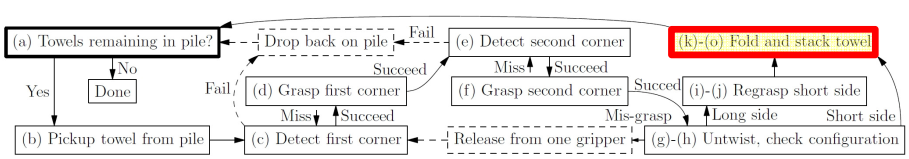
(a) Towels remaining in pile?
Drop back on pile
Fail
(e) Detect second corner
(k)-(o) Fold and stack towel
No
Succeed
Miss
Succeed
Yes
Done
Fail
(d) Grasp first corner
(f) Grasp second corner
Succeed
(i)-(j) Regrasp short side
Miss
Succeed
Mis-grasp
Long side
Short side
(b) Pickup towel from pile
(c) Detect first corner
Release from one gripper
(g)-(h) Untwist, check configuration

34

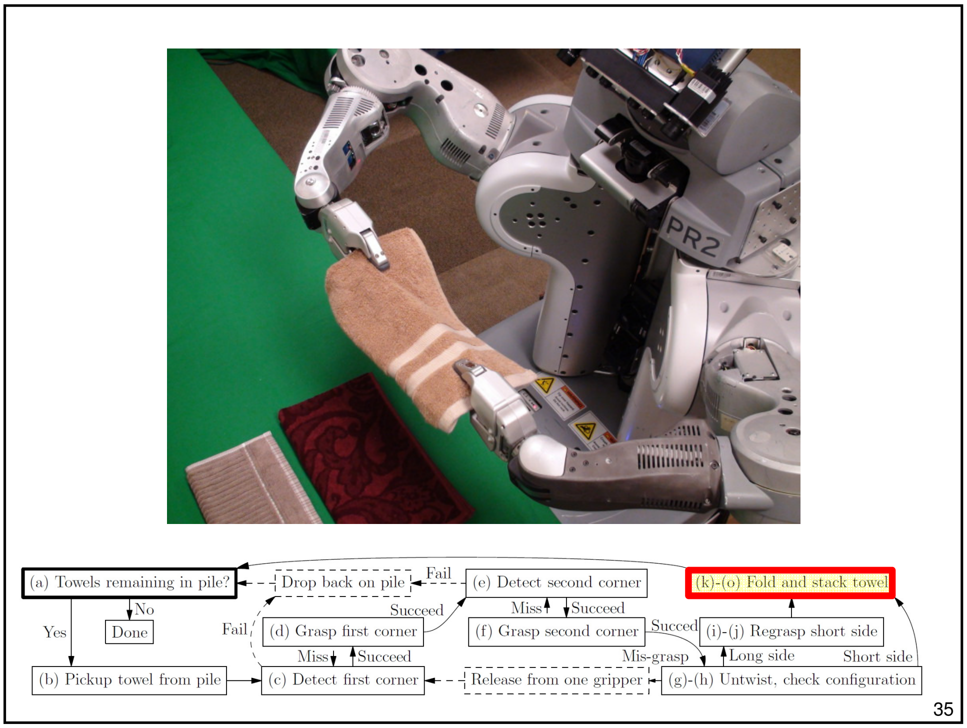
PR2
(a) Towels remaining in pile?
Drop back on pile
Fail
(e) Detect second corner
(k)-(o) Fold and stack towel
No
Done
Yes
Succeed
Fail
(d) Grasp first corner
Miss
Succeed
(f) Grasp second corner
Succeed
(i)-(j) Regrasp short side
Miss
Succeed
Mis-grasp
Long side
Short side
(b) Pickup towel from pile
(c) Detect first corner
Release from one gripper
(g)-(h) Untwist, check configuration
35

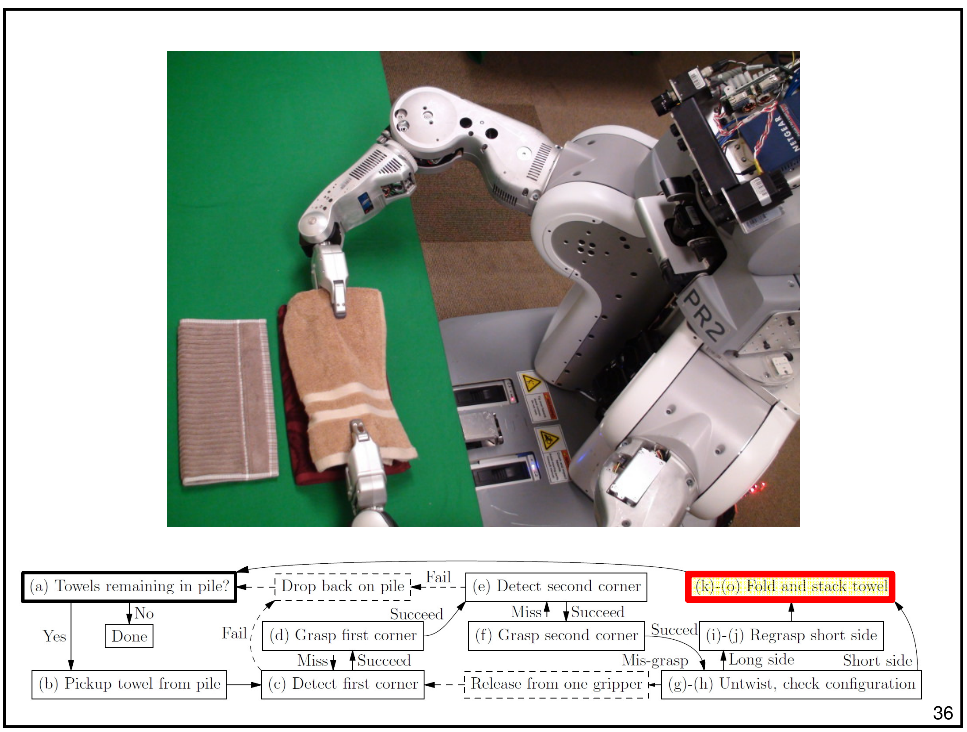
PR2
(a) Towels remaining in pile?
Drop back on pile
Fail
(e) Detect second corner
(k)-(o) Fold and stack towel
No
Done
Yes
Succeed
Fail
(d) Grasp first corner
Miss
Succeed
(f) Grasp second corner
Succeed
(i)-(j) Regrasp short side
Miss
Succeed
Mis-grasp
Long side
Short side
(b) Pickup towel from pile
(c) Detect first corner
Release from one gripper
(g)-(h) Untwist, check configuration
36

Glanced over

- Calibration of camera and robot
- Recognition of corners
- More generally: visual feedback during all manipulations
- How should we move the corners such that we obtain the desired result?

37

Now: Advanced Applications

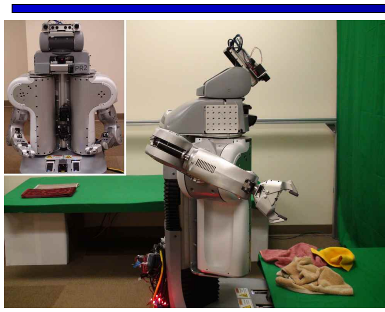

38

Motivating Example

- How do we specify a task like this?

[demo: autorotate / tictoc]

Autonomous Helicopter Flight

- Control inputs:
  - a_{lon} : Main rotor longitudinal cyclic pitch control (affects pitch rate)
  - a_{lat} : Main rotor latitudinal cyclic pitch control (affects roll rate)
  - a_{coll} : Main rotor collective pitch (affects main rotor thrust)
  - a_{rud} : Tail rotor collective pitch (affects tail rotor thrust)

Autonomous Helicopter Setup

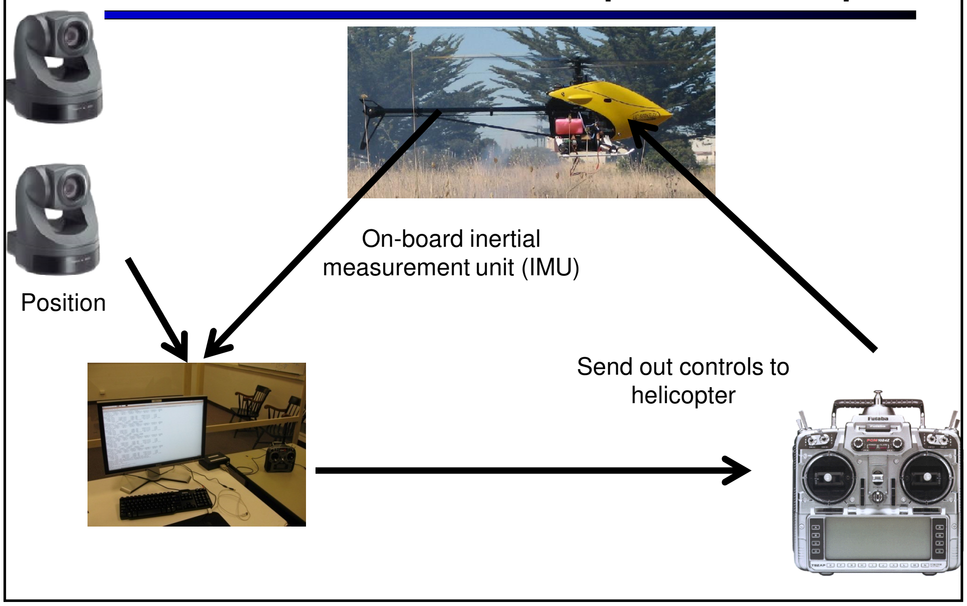

HMM for Tracking the Helicopter

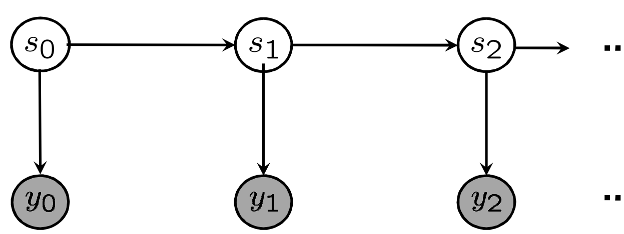

- State: $s = (x, y, z, \phi, \theta, \psi, \dot{x}, \dot{y}, \dot{z}, \dot{\phi}, \dot{\theta}, \dot{\psi})$
- Measurements:
  - 3-D coordinates from vision, 3-axis magnetometer, 3-axis gyro, 3-axis accelerometer
- Transitions (dynamics): [time elapse update]
  - $s_{t+1} = f(s_t, a_t) + w_t$

  [f encodes helicopter dynamics]

  [w is a probabilistic noise model]

42

Helicopter MDP

- State: $s = (x, y, z, \phi, \theta, \psi, \dot{x}, \dot{y}, \dot{z}, \dot{\phi}, \dot{\theta}, \dot{\psi})$
- Actions (control inputs):
  - a_{lon} : Main rotor longitudinal cyclic pitch control (affects pitch rate)
  - a_{lat} : Main rotor latitudinal cyclic pitch control (affects roll rate)
  - a_{coll} : Main rotor collective pitch (affects main rotor thrust)
  - a_{rud} : Tail rotor collective pitch (affects tail rotor thrust)
- Transitions (dynamics):
  - $s_{t+1} = f(s_t, a_t) + w_t$

  [f encodes helicopter dynamics]

  [w is a probabilistic noise model]

- Can we solve the MDP yet?

Problem: What's the Reward?

- Rewards for hovering:

[demo: hover]

$$R(s) = -(\alpha_x (x - x^*)^2 + \alpha_y (y - y^*)^2 + \alpha_z (z - z^*)^2 + \alpha_{\dot{x}} (\dot{x} - \dot{x}^*)^2 + \alpha_{\dot{y}} (\dot{y} - \dot{y}^*)^2 + \alpha_{\dot{z}} (\dot{z} - \dot{z}^*)^2)$$

- Rewards for "Tic-Toc"?
  - Problem: what's the target trajectory?
  - Just write it down by hand?

[demo: bad]

44

Helicopter Apprenticeship?

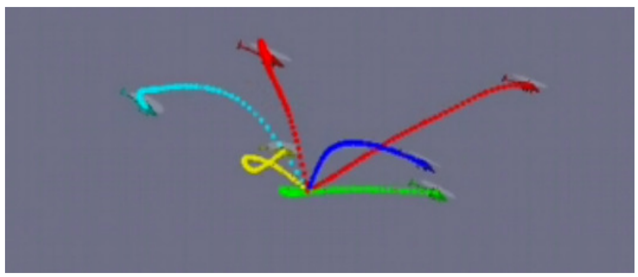

47

Probabilistic Alignment

Intended trajectory

$z_{t+1} = f(z_t) + \omega_t$

Expert demonstrations

$y_j = z_{\tau_j} + \nu_j$

Time indices

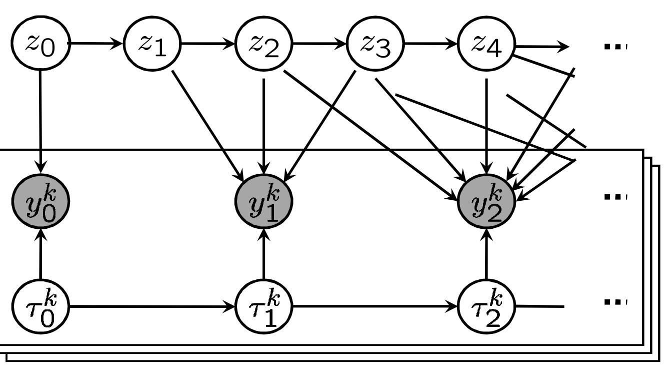

- Intended trajectory satisfies dynamics.
- Expert trajectory is a noisy observation of one of the hidden states.
  - But we don't know exactly which one.

[Coates, Abbeel & Ng, 2008]

[demo: alignment]

Alignment of Samples

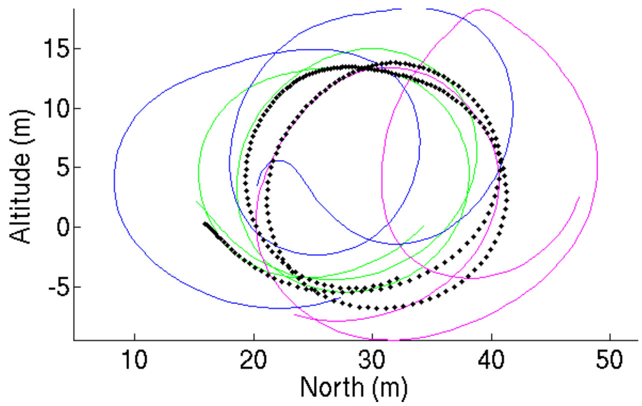

- Result: inferred sequence is much cleaner!

[demo: airshow]

Final Behavior