Assisting Persons With Dementia During Handwashing

Using a Partially Observable Markov Decision Process

Jesse Hoey

ICVS - March 21, 2007

School of Computing
UNIVERSITY OF DUNDEE

System Overview

An intelligent cognitive device that tracks a user through handwashing, providing cues when necessary.

The COACH

Cognitive Orthosis for Assisting Activities in the Home

System Features

- **Monitors** a user with dementia during handwashing,
- **Prompts** a user only when necessary,
- **Calls** for human assistance if needed,
- **Vision based**: only video inputs,
- **Robust**: able to recover from tracker failures,
- **Real-Time**: 320×240 @ 40fps,
- **Probabilistic**: Explicit modeling of uncertainty,
- **Partially observable**: estimates user internal states,

Integrated System

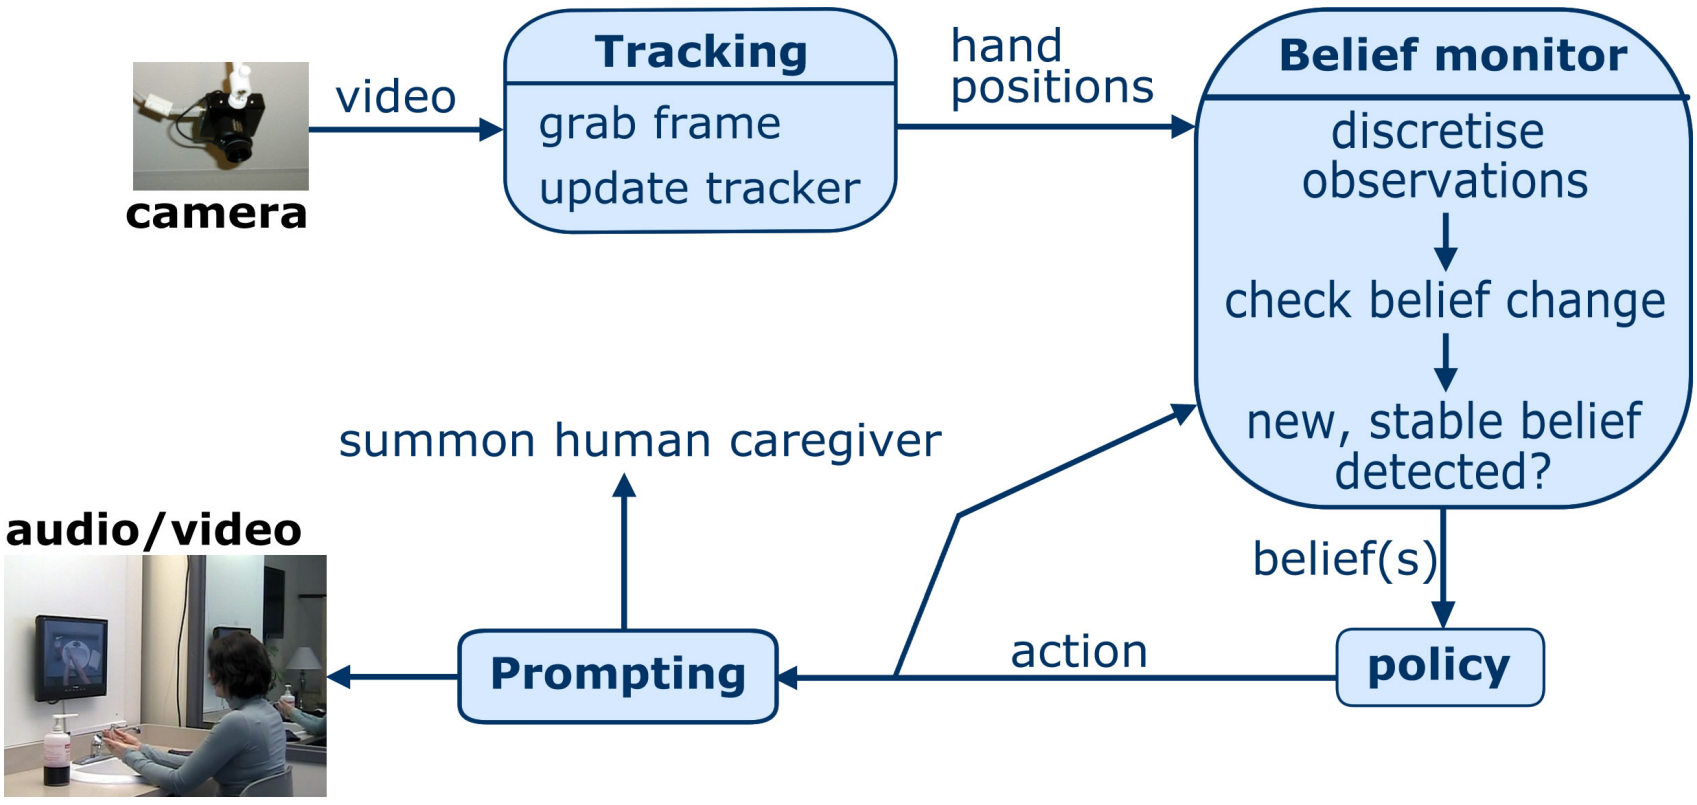

Tracking

Tracking through occlusions and changes in shape

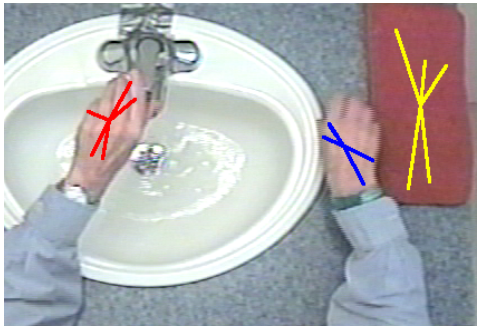

(a) 4974

(b) 5096

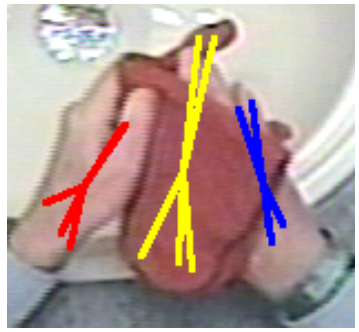

(c) 5502

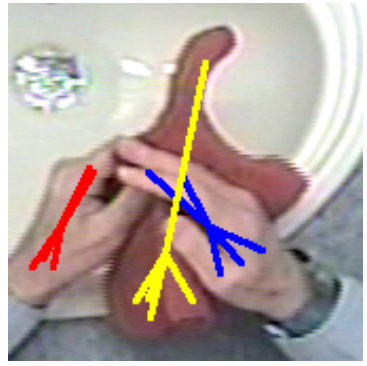

(d) 5576

- Objects modeled as **flocks of features**
- Simple **color features**, Gaussian distributions
- **Bayesian sequential** estimation
- Monte Carlo approximation (**particle filter**)
- **Three interacting** filters: 2 hands + towel
- Dynamics: **Constant mean velocity**
- **Mixed-state dynamics**
- **Data-driven proposal** (e.g. ICondensation)

Data Driven Proposal

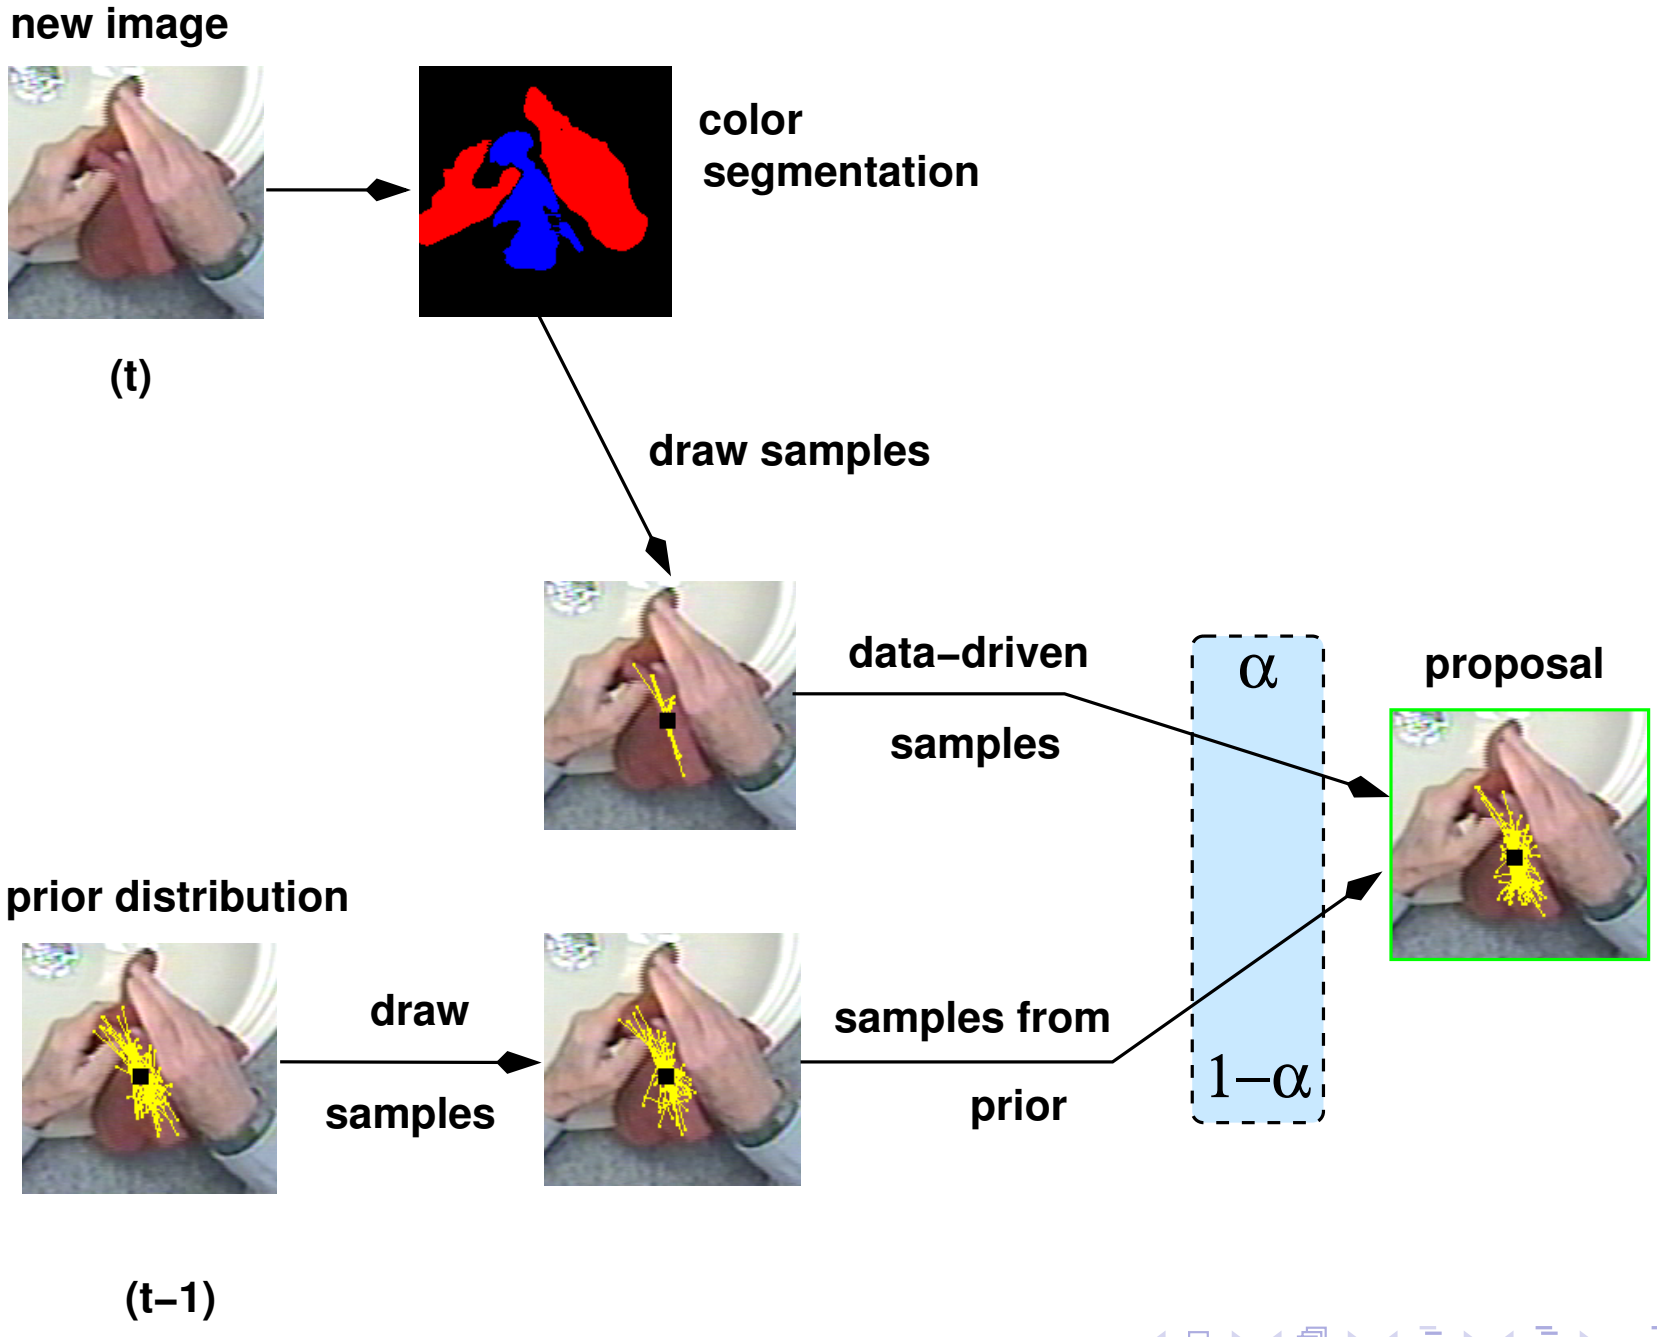

Partially Observable Markov Decision Process (POMDP)

A POMDP is a probabilistic temporal model of an agent interacting with its environment : a tuple $\langle S, A, T, R, O, B \rangle$

- S: finite set of unobservable states
- O: set of observations
- $B : S \times A \rightarrow O$ observation function
- A: finite set of agent actions
- $T : S \times A \rightarrow S$ transition function
- $R : S \times A \rightarrow \mathcal{R}$ reward function

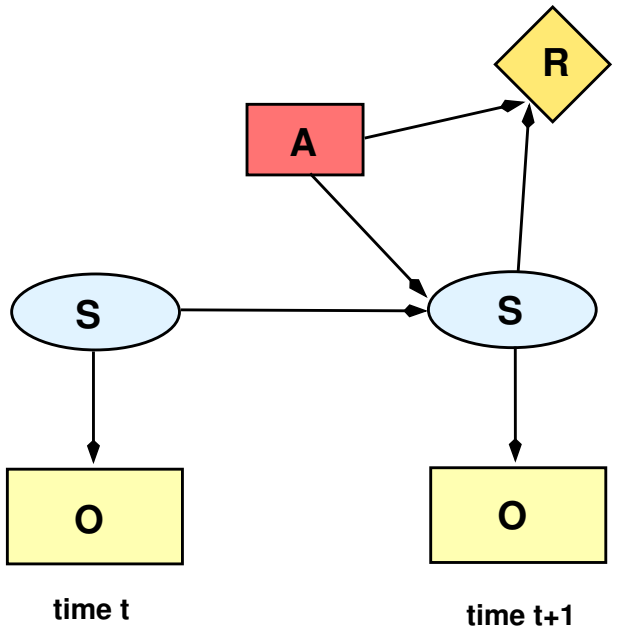

What can you do with a POMDP?

Everything you can do with a dynamic Bayes net (DBN)...

- monitor belief(s)
- compute data likelihood given model

What can you do with a POMDP?

Everything you can do with a dynamic Bayes net (DBN)...

- monitor belief(s)
- compute data likelihood given model

... **plus!** compute a *policy of action* , π

- $\pi(b(s))$: belief states $\rightarrow$ actions
- policy actions maximise *long-term* reward
- tradeoff multiple, competing objectives

- Policy **does not imitate** a human - **optimises** for system

Handwashing POMDP

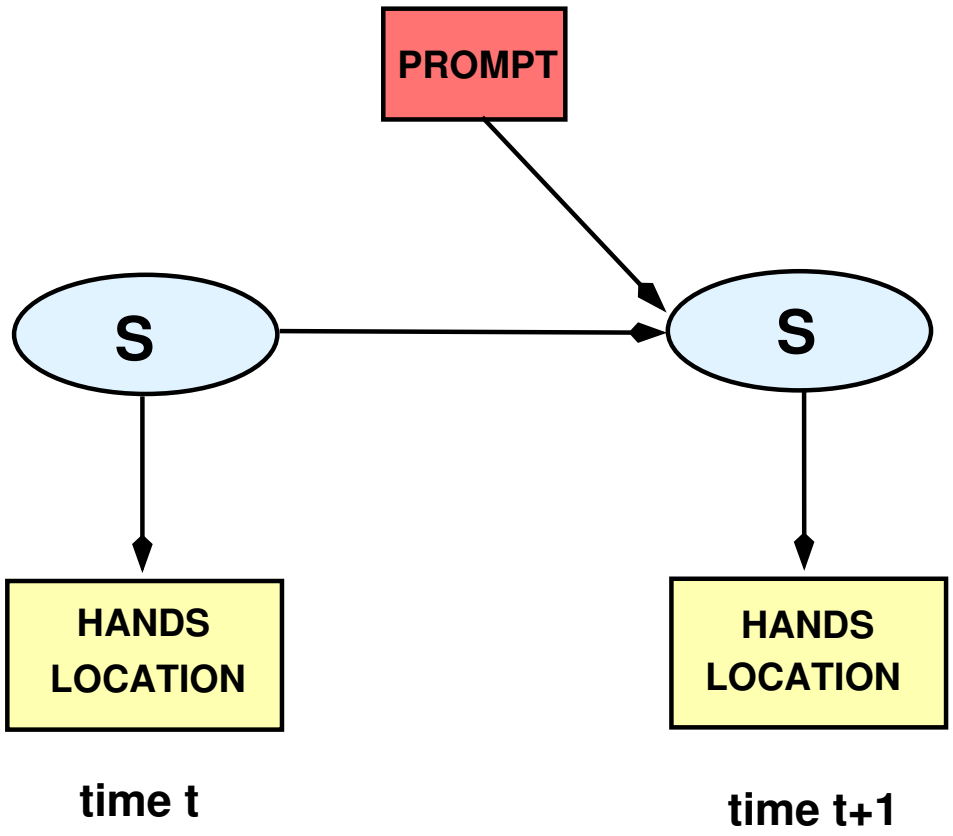

Handwashing POMDP

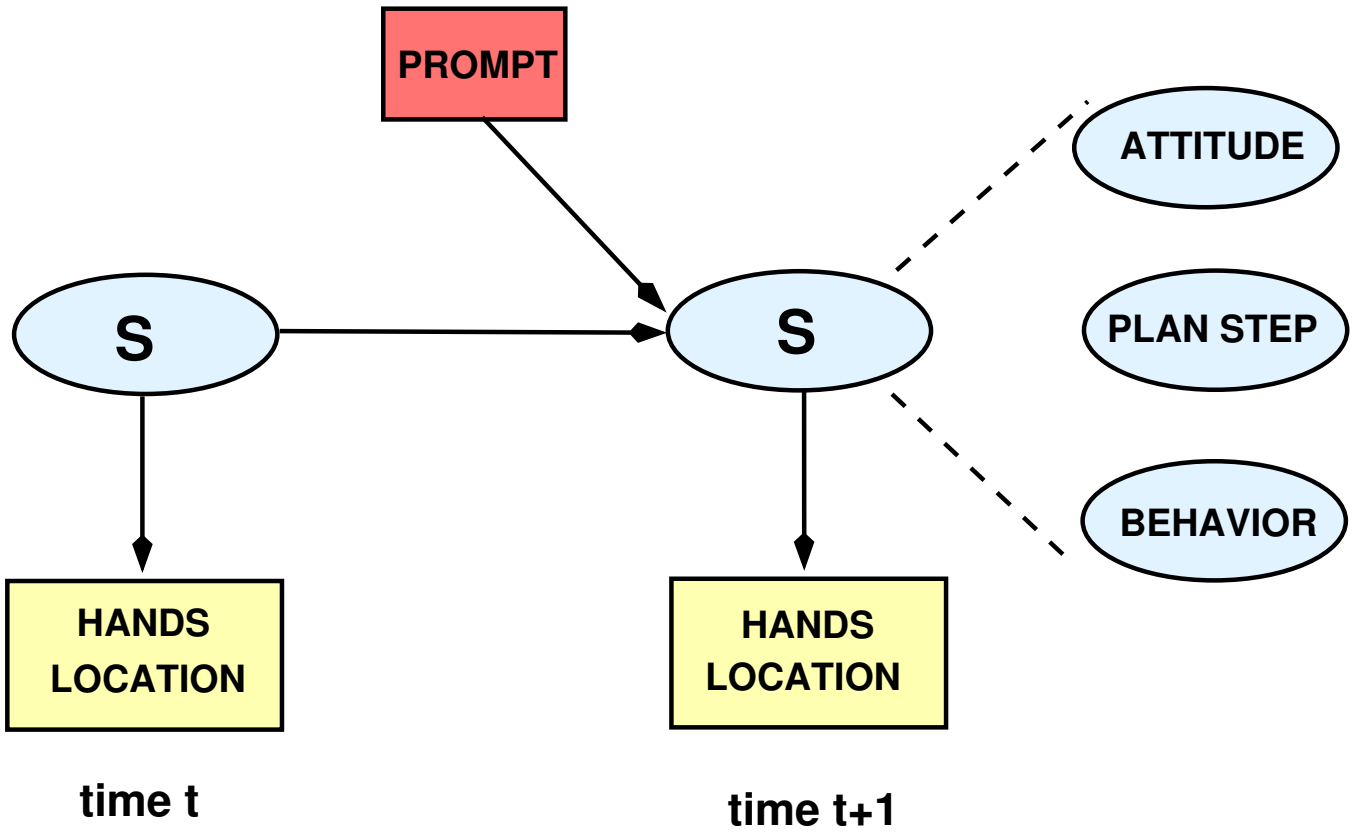

Handwashing POMDP: Actions

- **do nothing**: system waits
- **call caregiver**: system calls for **single step** assistance
- **prompts**:
  - audio/video
  - male voice
  - preceded with reminder : *"you're washing your hands"*
  - 3 levels of specificity:

low : basic prompt with few details
"Use the soap"

medium : include person's name, more details
"John, use the soap in the pink bottle"

high : medium prompt with video demonstration

Handwashing POMDP: Plansteps

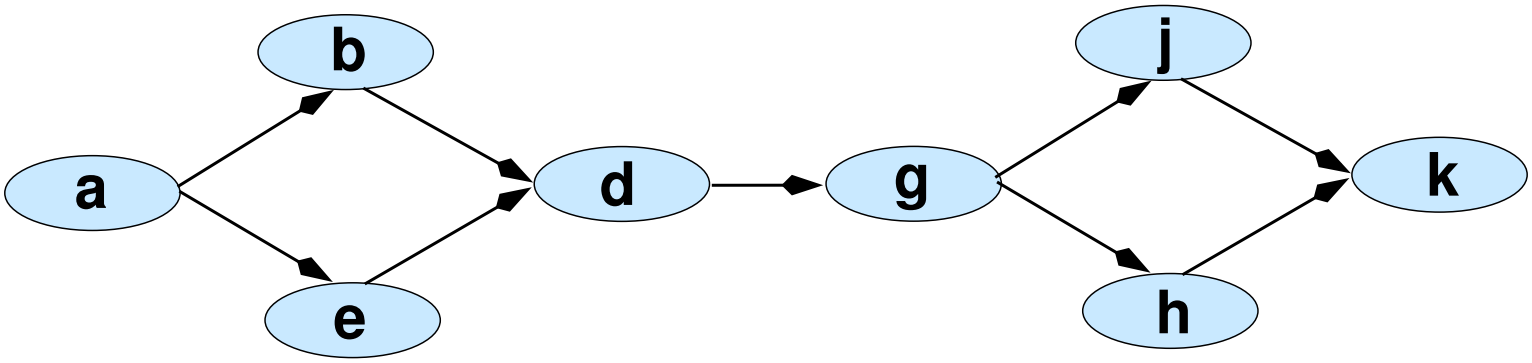

Handwashing POMDP: Plansteps

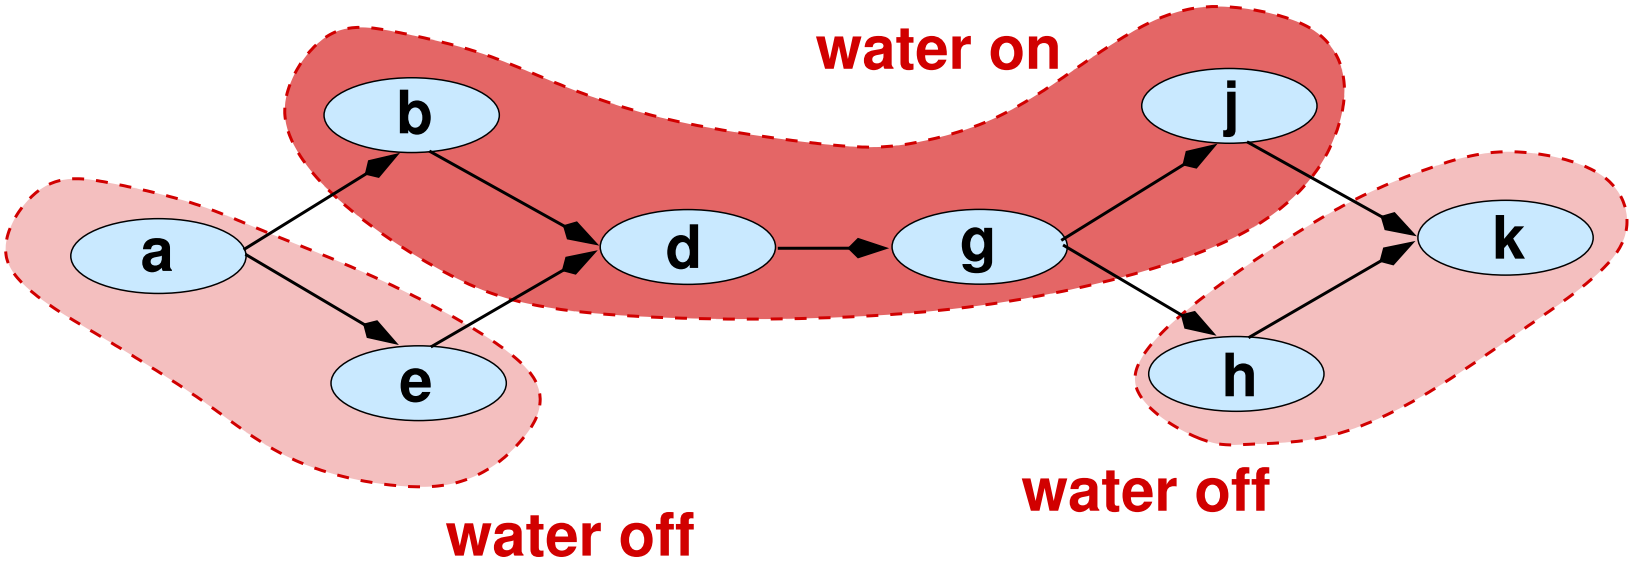

Handwashing POMDP: Plansteps

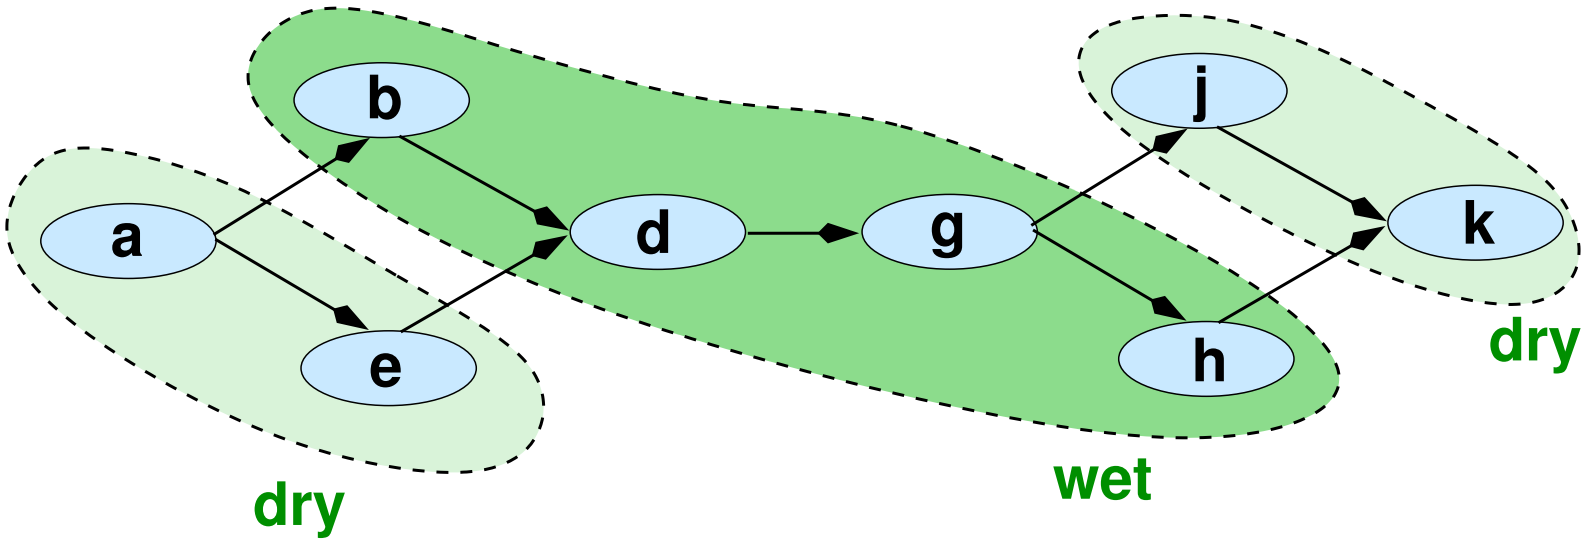

Handwashing POMDP: Plansteps

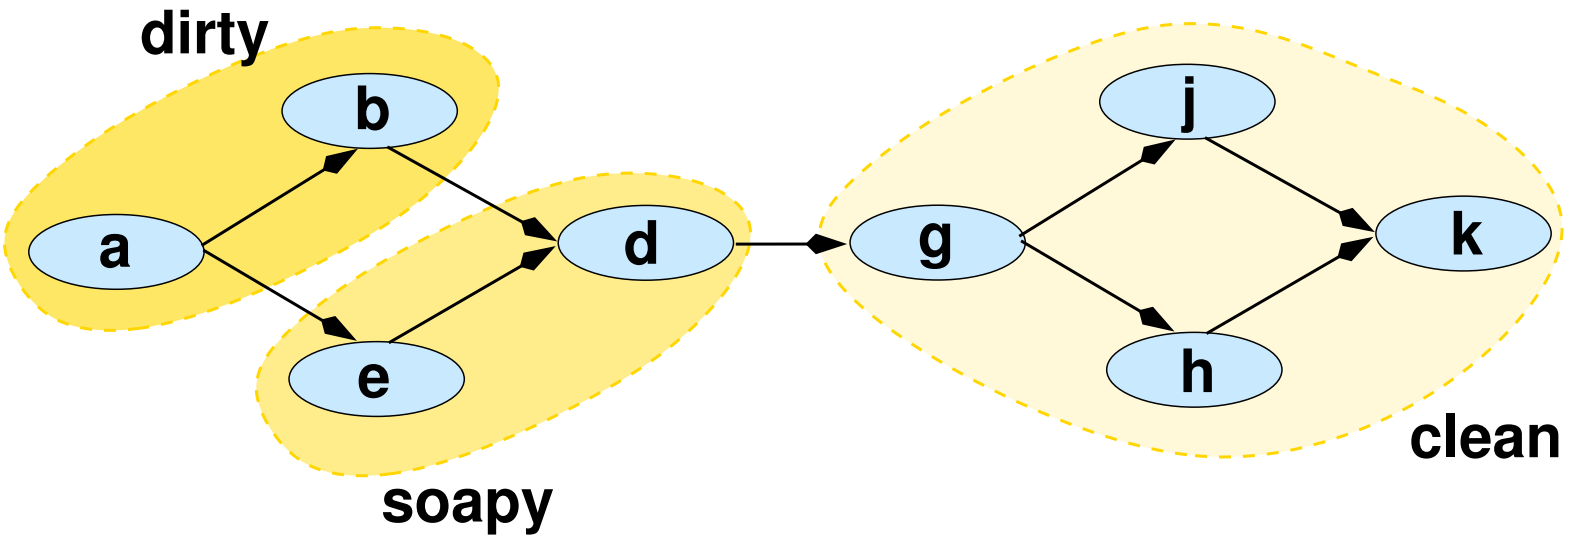

Handwashing POMDP: Plansteps

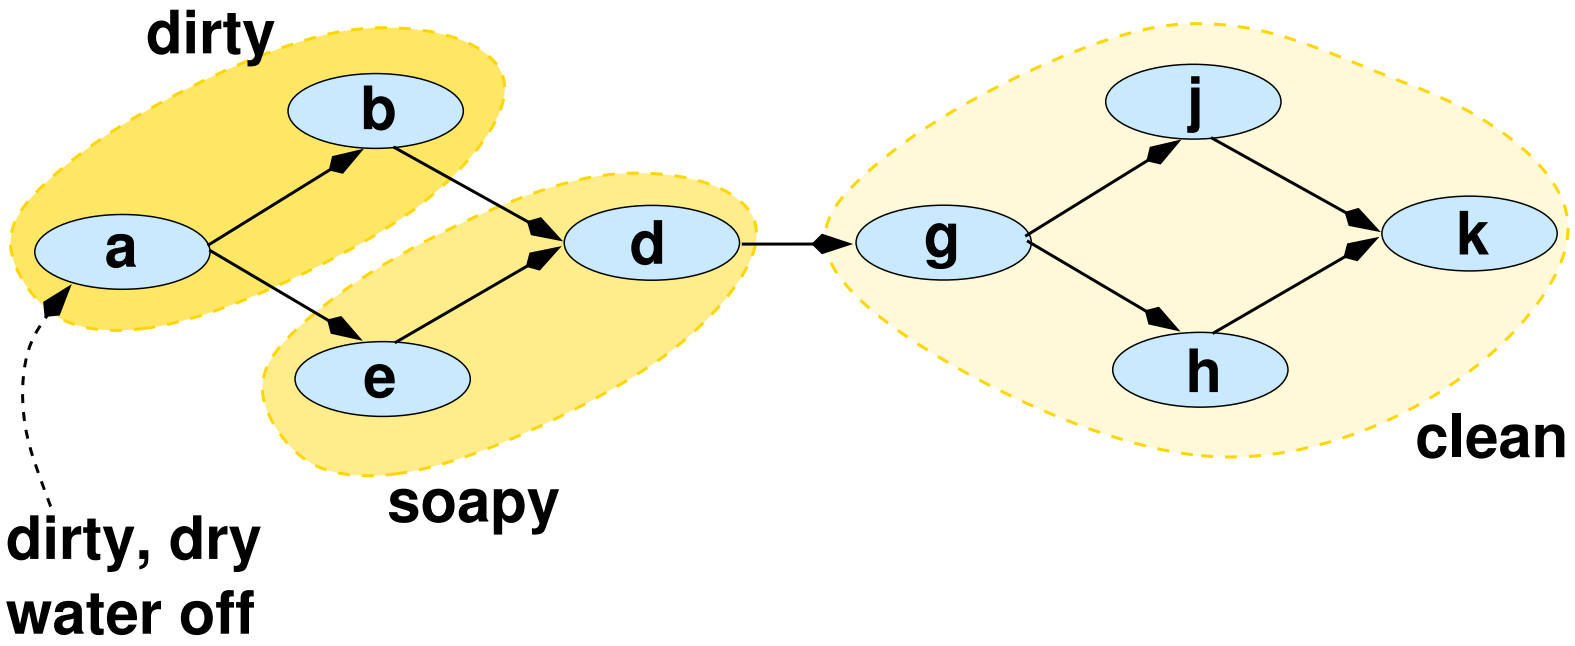

Handwashing POMDP: Plansteps

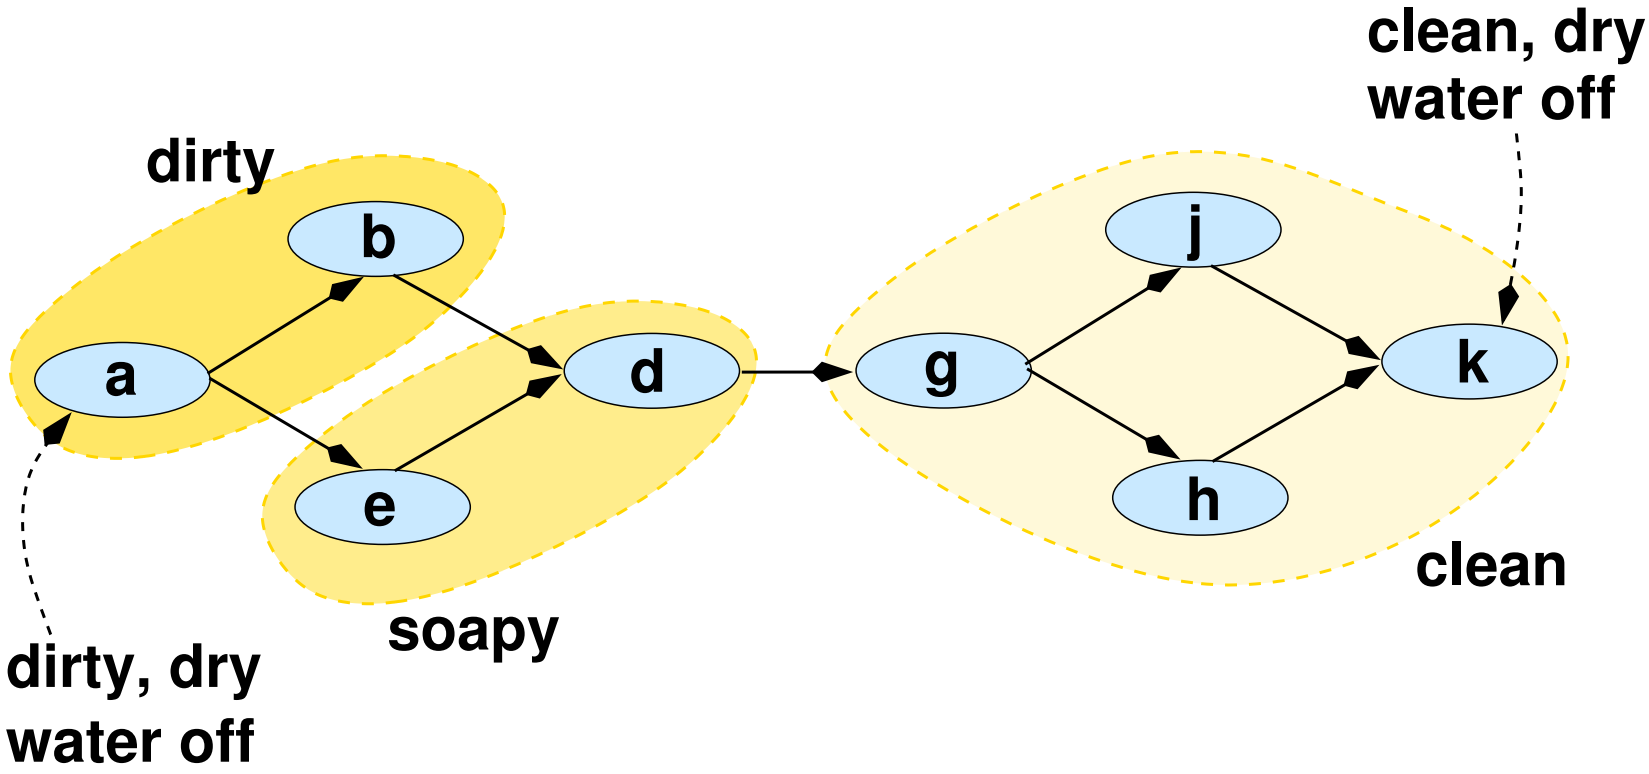

Handwashing POMDP: Plansteps

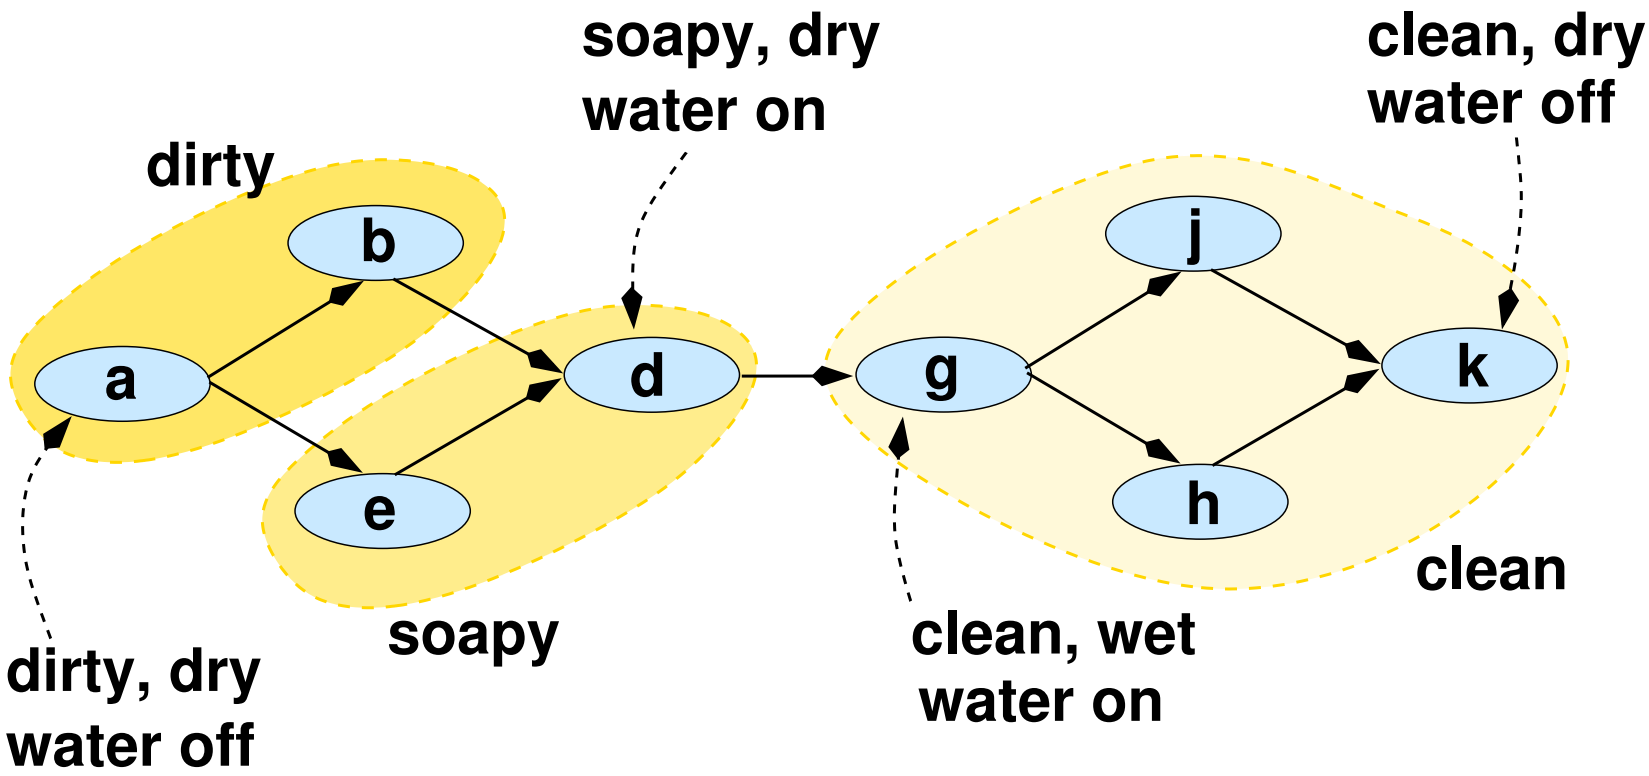

Handwashing POMDP: Plansteps

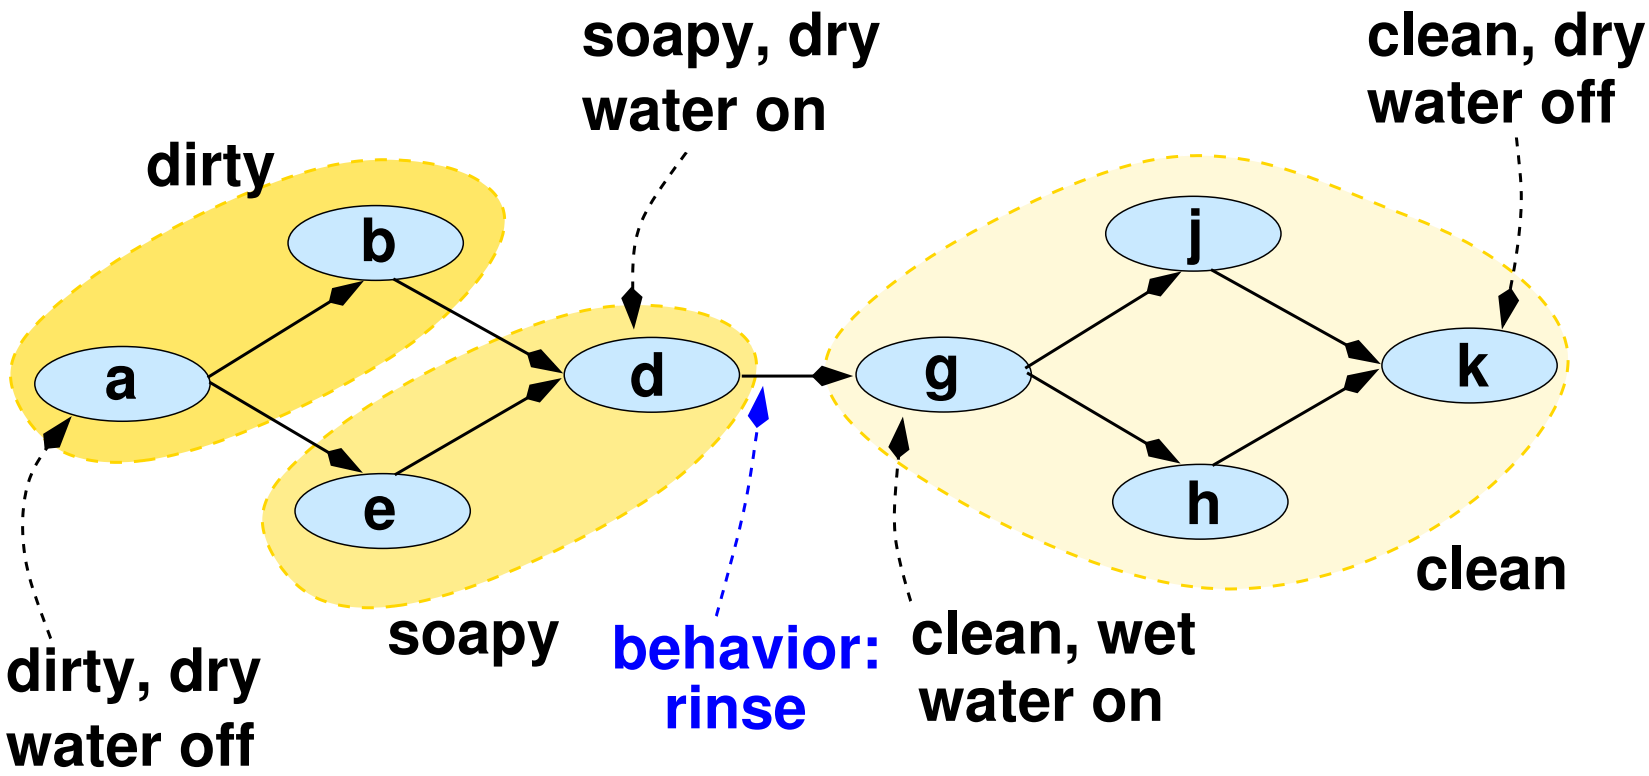

Handwashing POMDP: Behaviors

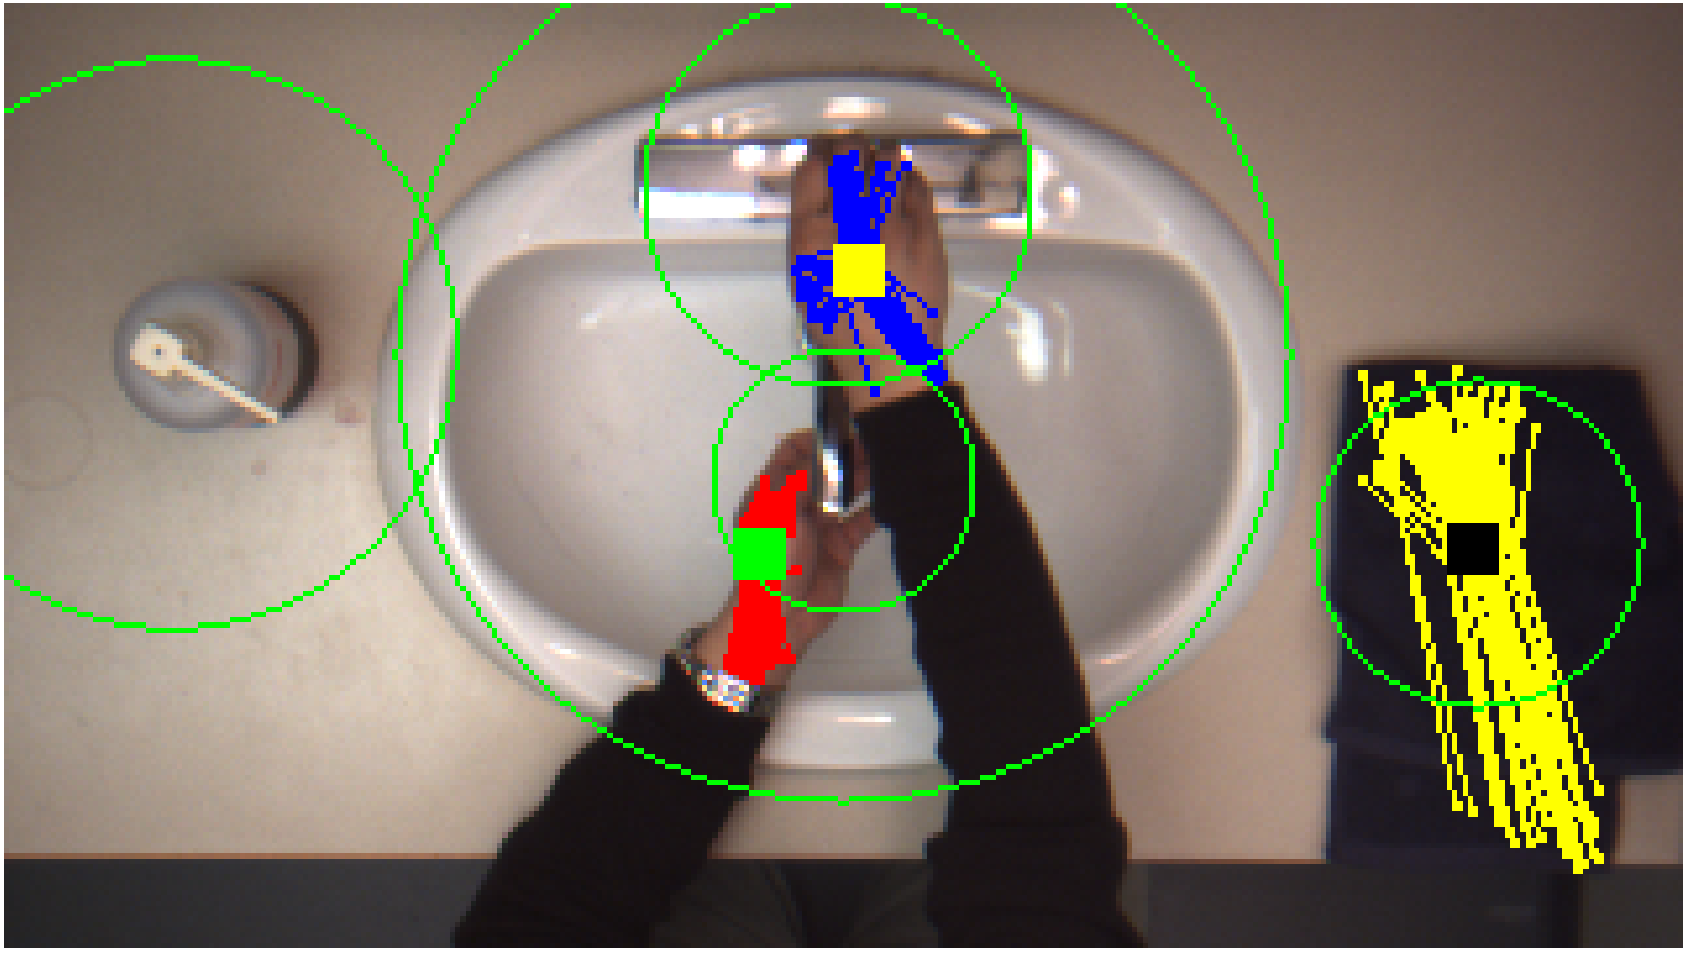

Handwashing POMDP: Attitude

Factor:	Models:	Dynamics:
Awareness	need for assistance	changes quickly
Responsiveness	response to assistance	changes from day to day
Dementia Level	likelihood user will be aware and responsive	does not change

Handwashing POMDP

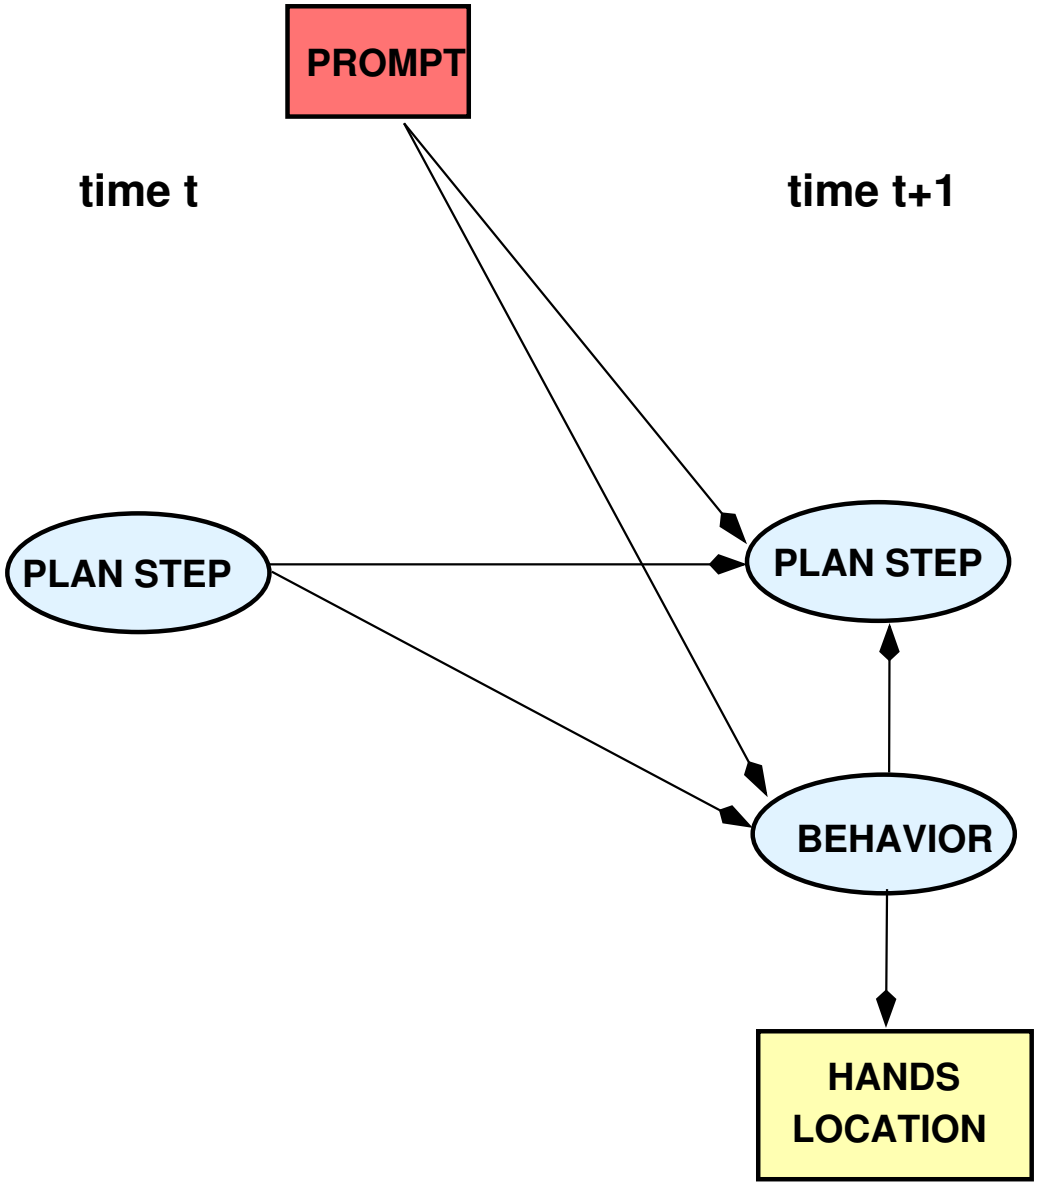

Handwashing POMDP

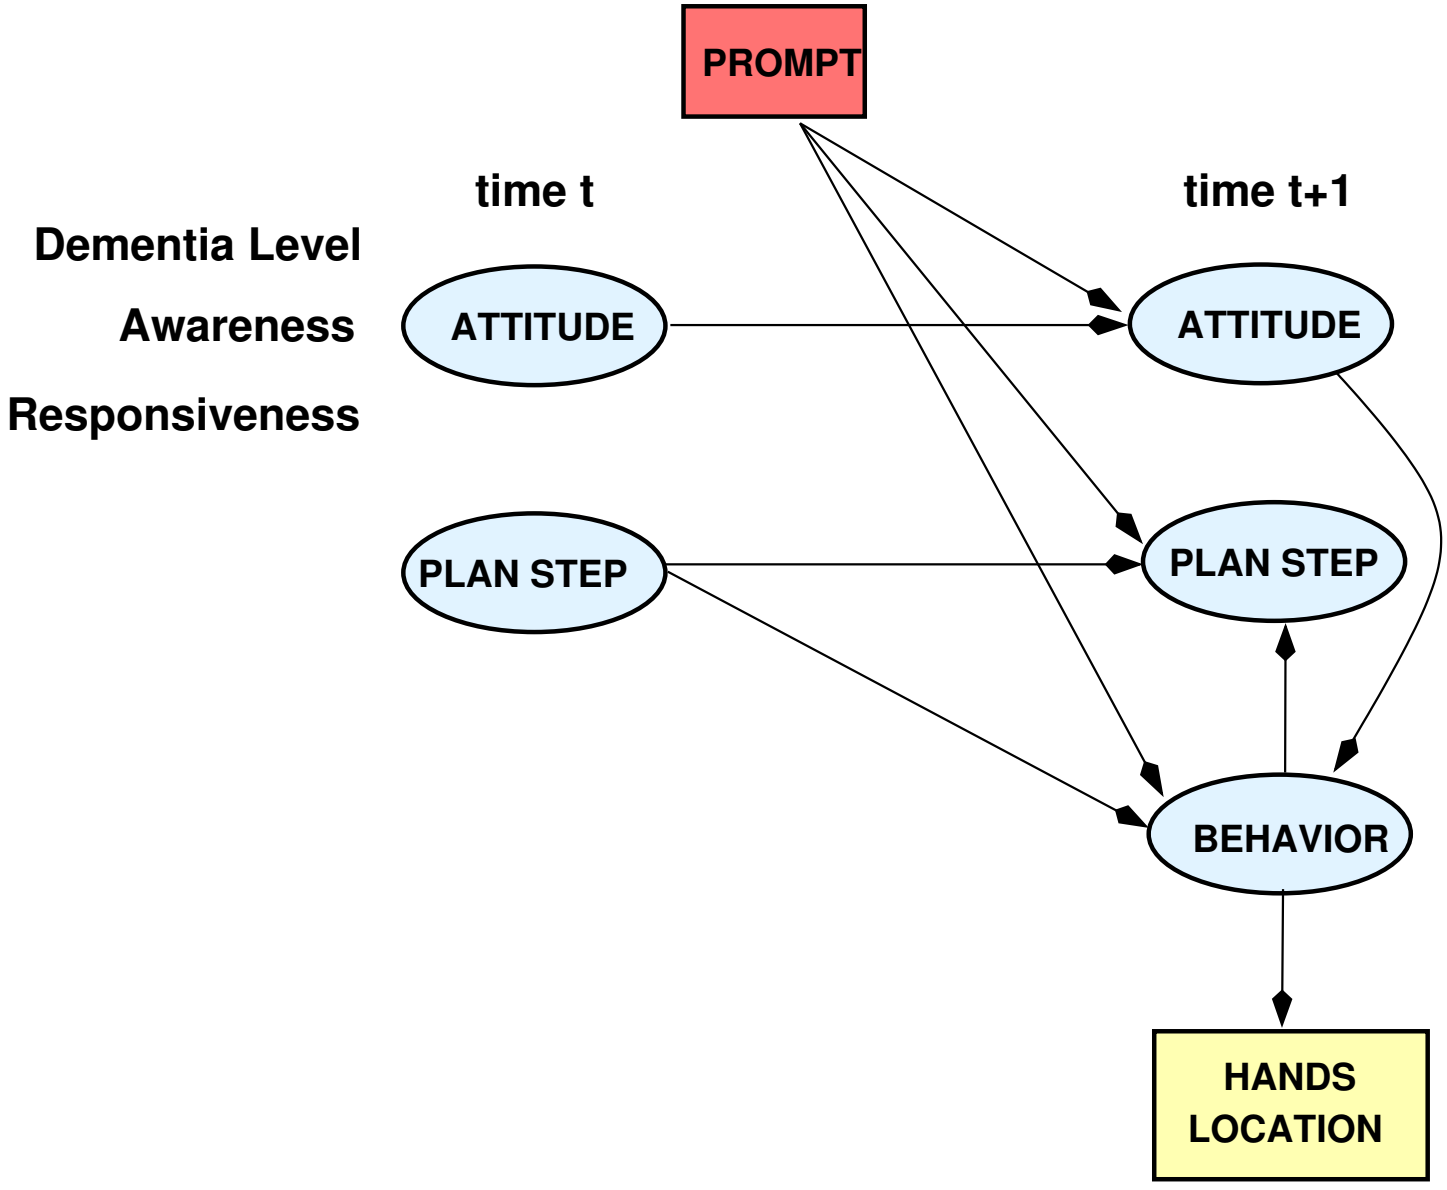

Rewards and Solutions

Rewards and Costs:

- Task completion: large reward
- Prompts: costly
- Cost for prompting twice in a row (pestering)

Solution:

- 200K states $\times$ 198 observations $\times$ 26 actions
- Optimal solution **intractable**
- **Approximations** used:
  - Point-based (**Perseus** - Vlassis & Spaan, 2005)
  - Structured solution (**SPUDD** - Hoey & St. Aubin, 1999)
  - Restrict size of solution (Poupart, 2005)
  - Disregard observations (Hoey & Poupart, 2005)
  - Merge states (St. Aubin & Hoey, 2000)

Overall System

- Runs on a single laptop Intel core 2 duo 2Gb Ram
- LCD screen for prompts
- Point Grey Research Dragonfly II
- 320x240 at 40Hz
- Modular processes
- UDP IPC communication through central broker

Simulation Results

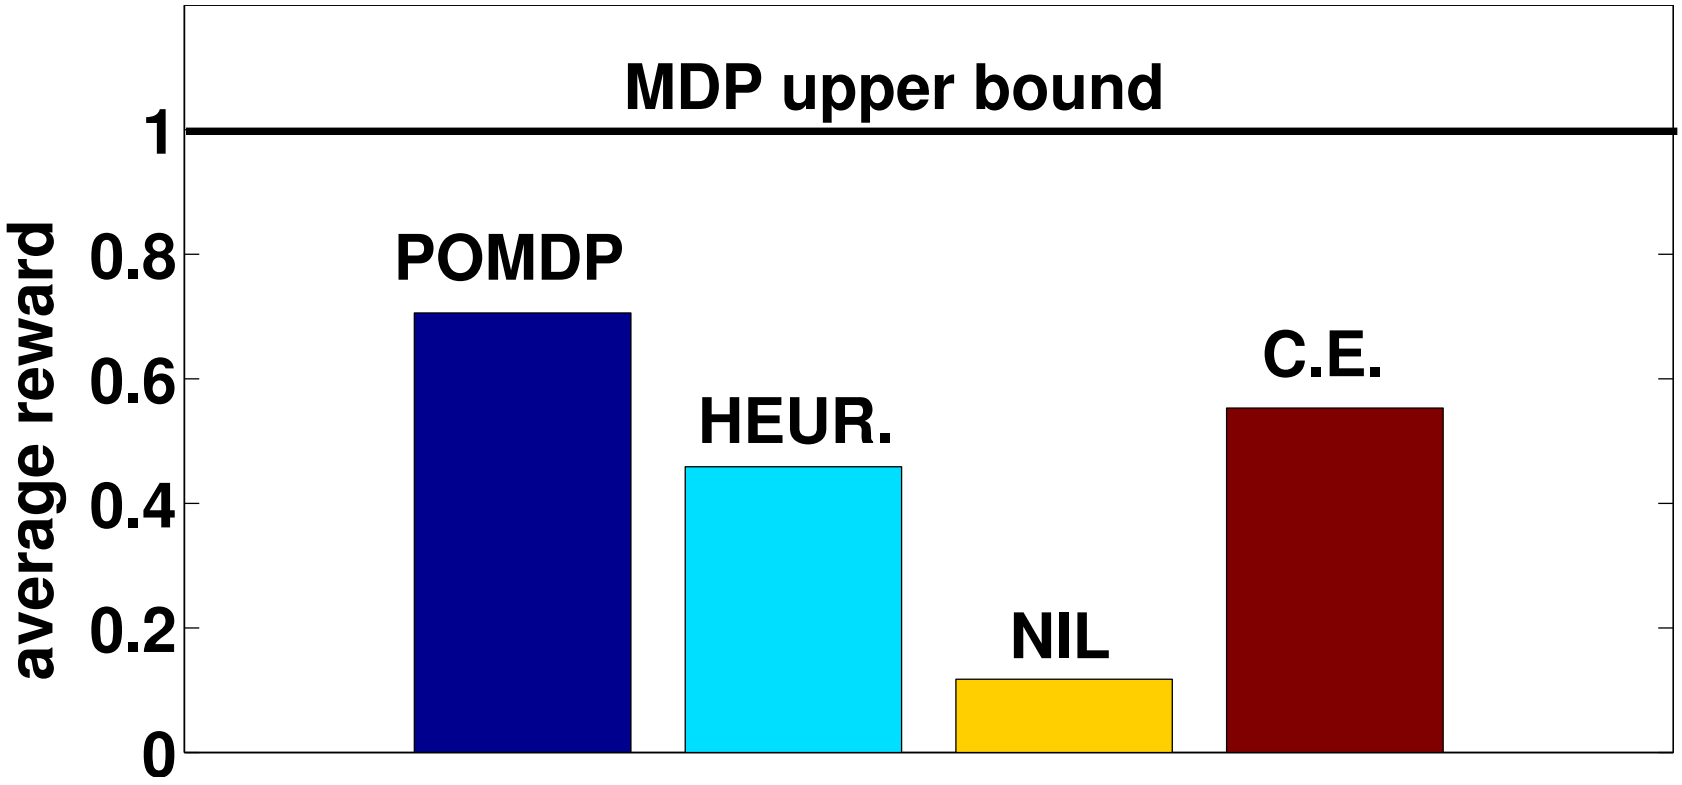

Average over

- 20 simulation trials (50 steps each)
- 10 experiments
- all true user types (attitudes)

Simulation Results

Learning the user's dementia level over a clinical trial
(20 handwashing simulations)

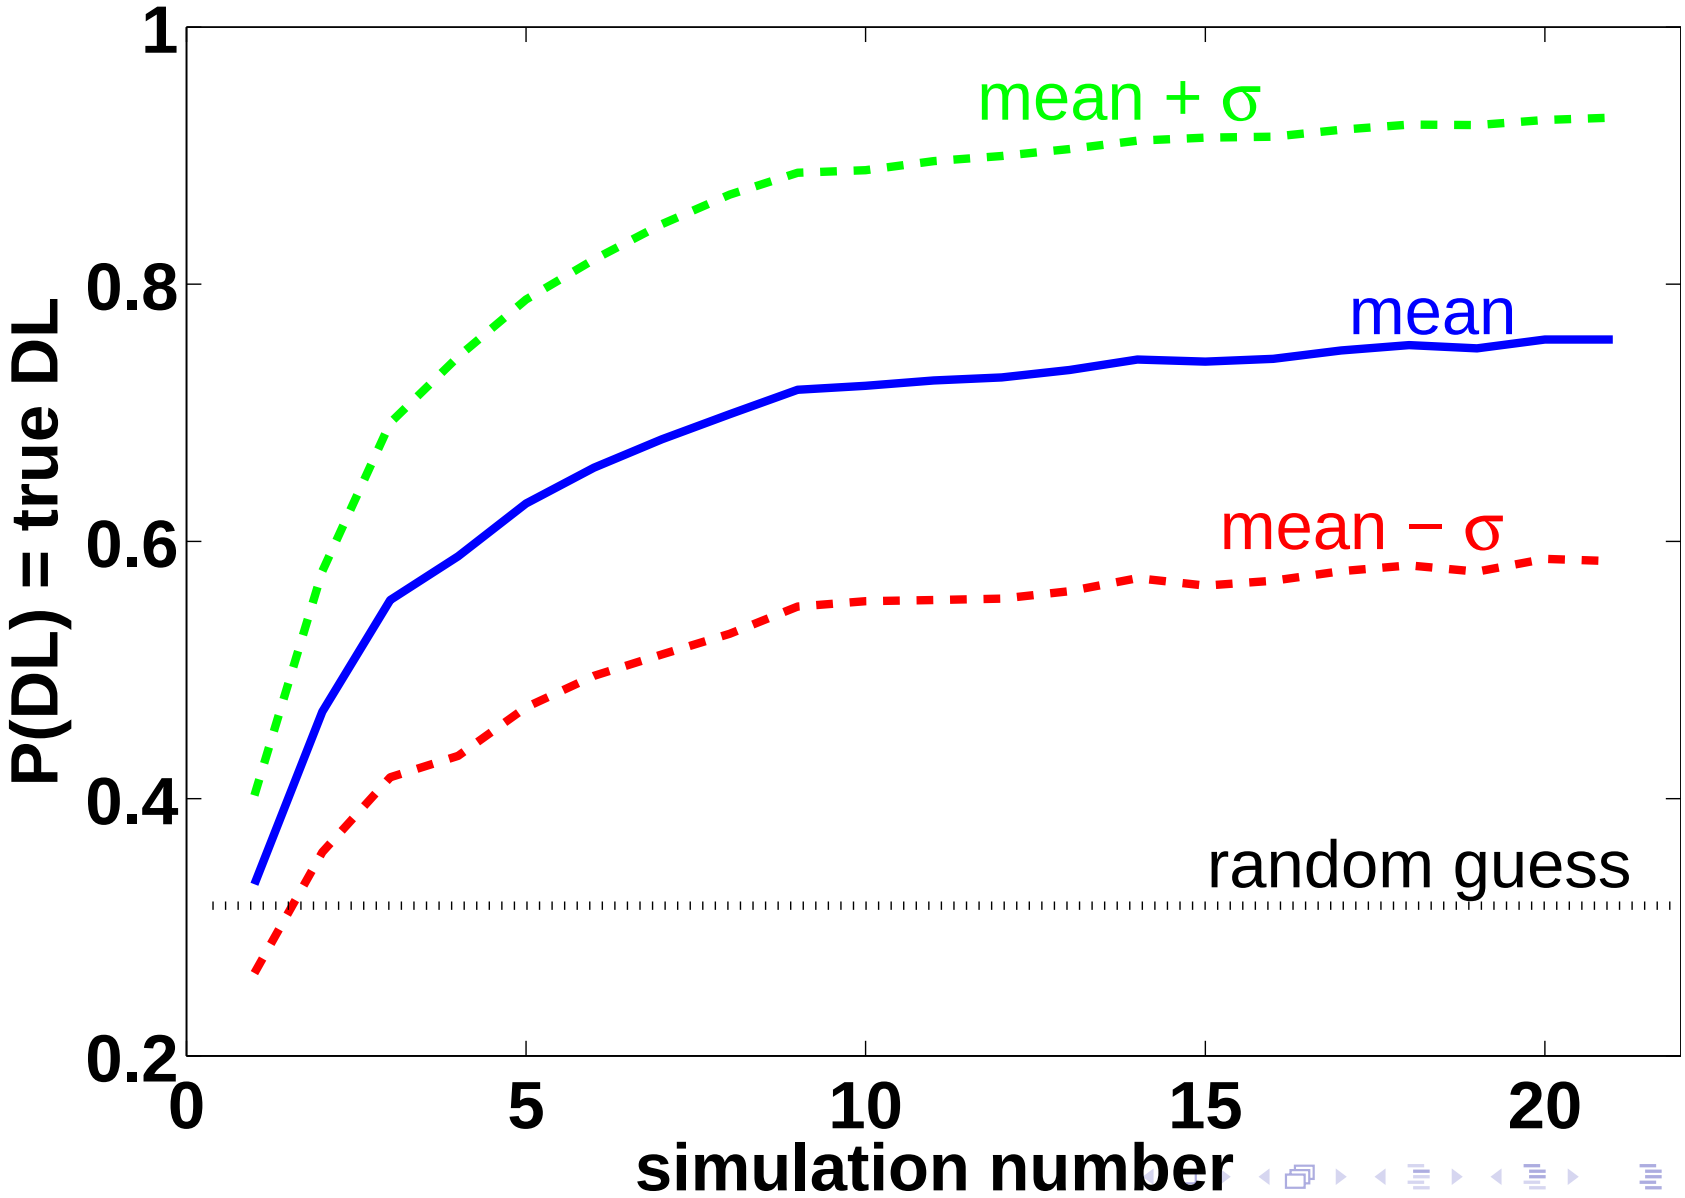

Actor Trials

Scenario 3:

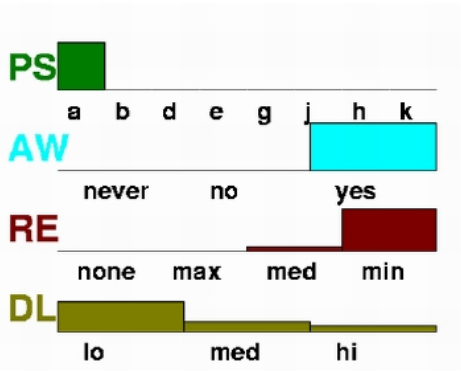

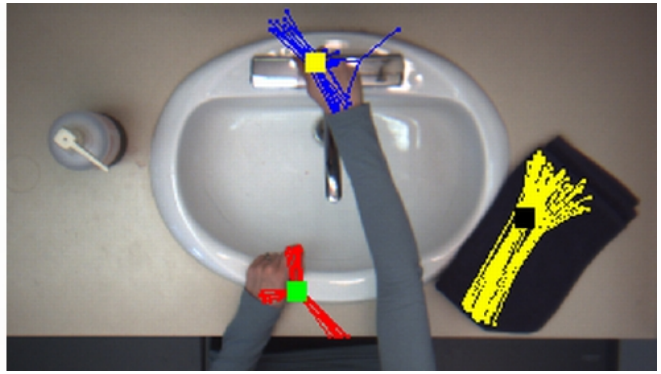

Scenario 4:

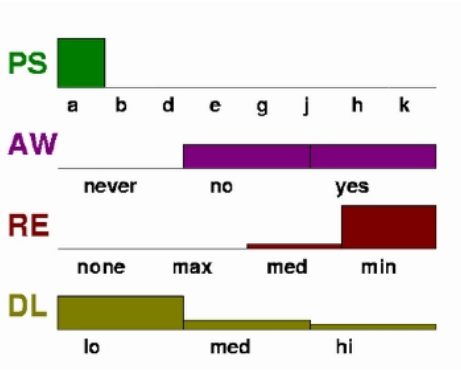

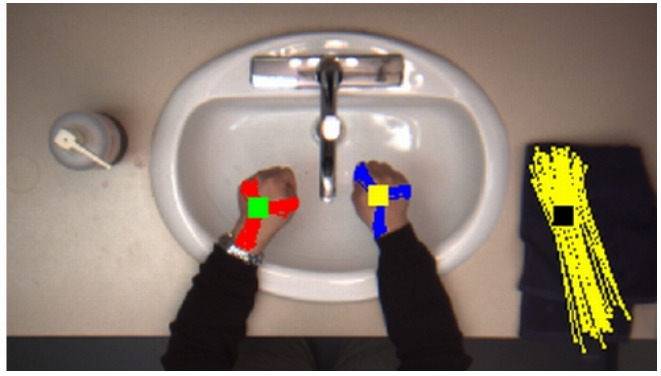

Current System - Clinical Trials

- Full system to be deployed in clinical trials spring 2007
- Long-term care facility in Toronto, Canada
- 8 week trial (A-B-A-B baseline/system alternating 2 weeks)
- 10 participants moderate to severe dementia
- Measure: reduction in caregiver burden

Future Work

- Hierarchical modeling for full integration of tracking with POMDP
- Apply to other activities
- Learn POMDP parameters
- Bayesian Reinforcement Learning (online)

Thanks to...

Collaborators

- Jesse Hoey, University of Dundee
- Axel von Bertoldi, University of Toronto
- Alex Mihailidis, University of Toronto
- Pascal Poupart, University of Waterloo

Other help

- Craig Boutilier, University of Toronto
- Kate Fenton, University of Toronto
- Zhonghai Li, University of Toronto
- Rozanne Wilson, University of Toronto
- Jennifer Boger, University of Toronto

Thanks to...

Support

- American Alzheimer's Assocation
- Intel Corporation

Institutions

IATSL Intelligent Assistive Technology and Systems Lab

School of Computing
UNIVERSITY OF DUNDEE

University of Waterloo

UNIVERSITY of TORONTO