Probabilistic Dialogue Modeling for Speech-Enabled Assistive Technology

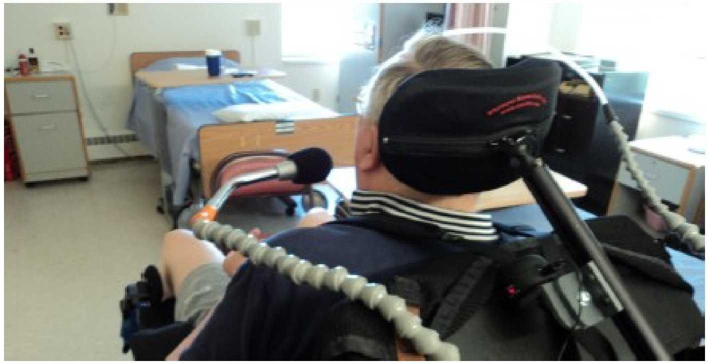

William Li

August 21, 2013

wli@csail.mit.edu
http://people.csail.mit.edu/wli/

Speech Challenges at The Boston Home (TBH)

- Fatigue

  “Chair, what is the activities schedule for Wednesday?”

- Over-nasalization

  “What's Sunday's breakfast?

- Vocal fry

  “Any good gossip today?”

Roadmap

1. Motivation: Spoken dialogue systems for high-error speakers

2. Dialogue system: Partially observable Markov decision process (POMDP) modelling and implementation

3. User study: experimental design and results

Desired Spoken Dialogue System Functions

- Time
- Weather
- Activities schedules
- Breakfast/lunch/dinner menus
- Hands-free phone calls
- Wheelchair navigation
- Nurse call
- Control of bed functions

Desired Spoken Dialogue System Functions

- **Time**
- **Weather**
- **Activities schedules**
- **Breakfast/lunch/dinner menus**
- **Hands-free phone calls**
- Wheelchair navigation
- Nurse call
- Control of bed functions

Challenge: High Speech Recognition Error Rates

Concept error rates for target and control populations
(30 utterances, trigram LM, unadapted acoustic models)

Boston Home users

Speaker (Target)	Concept Error Rate
target01	13.3%
target02	3.3%
target03	33.3%
target04	56.7%
target05	26.7%
target06	9.4%
target07	6.6%
mean	21.4%
std. dev.	18.9%

Lab users

Speaker (Control)	Concept Error Rate
control01	3.3%
control02	10.0%
control03	6.7%
control04	13.3%
control05	3.3%
control06	3.3%
control07	0.0%
mean	5.7%
std. dev.	4.6%

Spoken Dialogue System Components

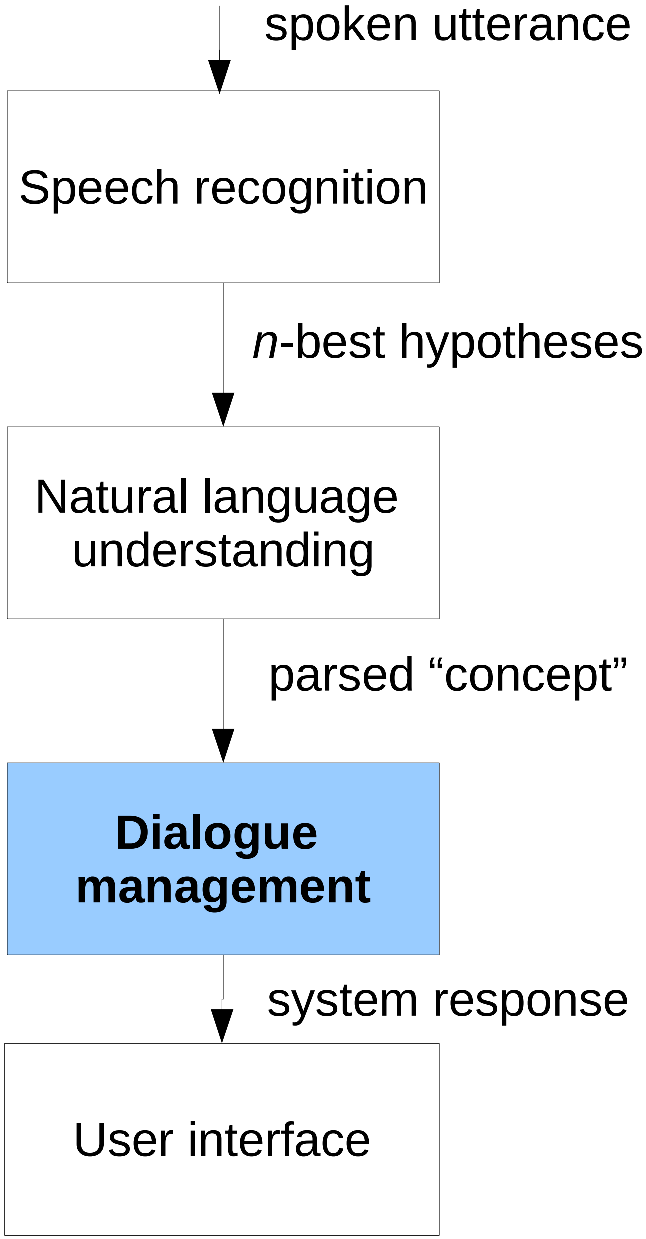

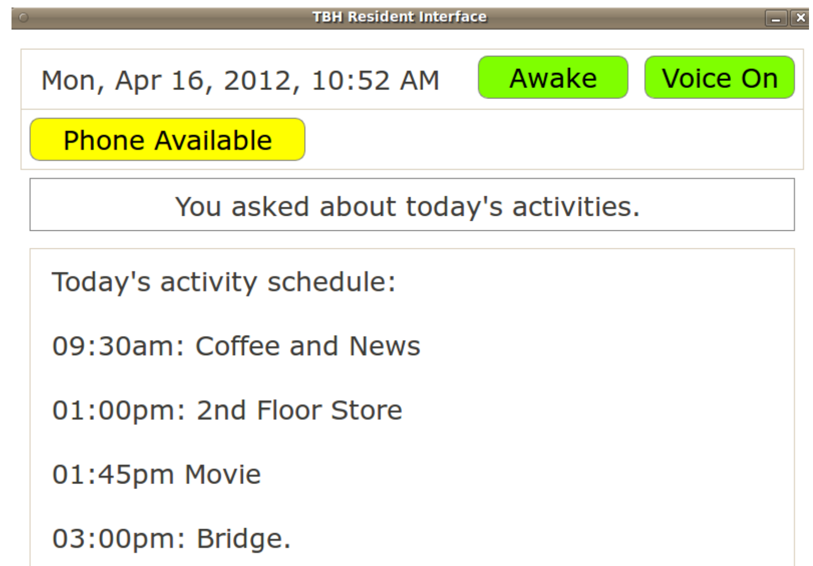

Why Dialogue for Assistive Technology?

- Abstraction: focus on **user intents** instead of words
  - Fewer parameters, shared training data among users
- Handle errors in speech recognition
  - Impaired speech, background noise, inherent ambiguity in spoken interaction
- Natural interaction
  - More acceptable assistive technology?

Partially Observable Markov Decision Process

(POMDP)

Theory and Implementation

Rule-based Dialog Managers

- Large engineering and maintenance effort
- Substantial hand-tuning of parameters (e.g. thresholds, if/then decision statements)

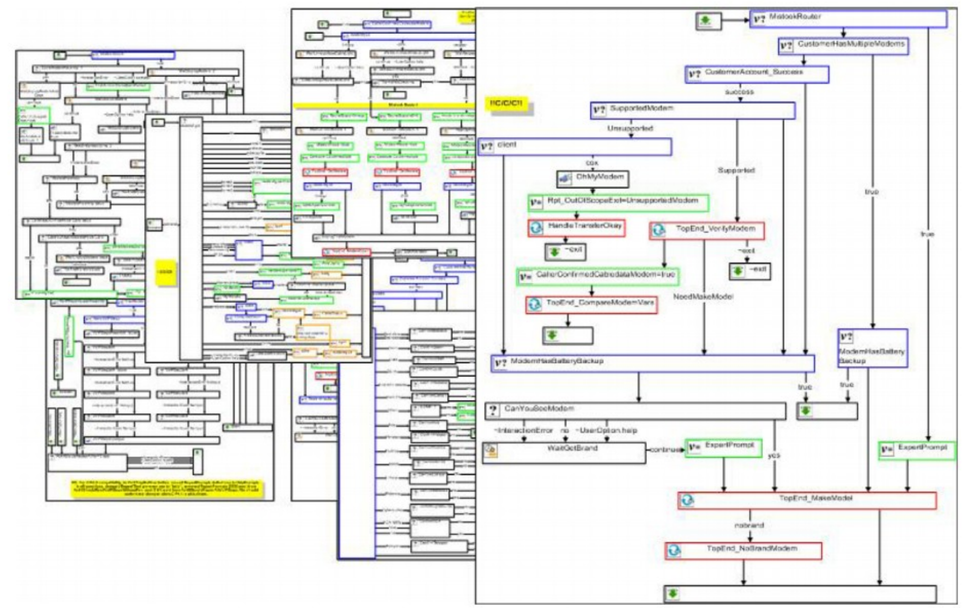

Paek/Pieraccini (2008)

POMDP Definition

- **Partially observable:** state is hidden, as opposed to a fully observable Markov decision process (MDP)
- **Markov:** transition/observation functions depend only on entities in time *t-1*
- **Decision process:** The system infers the state to choose actions
- **Key Terms:**
  - **Belief, b:** probability distribution over states
  - **Policy, f(b)→A:** mapping of beliefs to actions

Spoken Dialog System POMDP (SDS-POMDP)

Intuition: Use dialog to help determine the user's intent

User has a **state (goal/intent)** that is not directly observable

↓

Spoken dialog system (SDS) receives noisy **sensor observations (speech recognition hypotheses)**

↓

SDS updates its **belief (probability distribution over states)** based on observation model

↓

SDS decides, based on its belief, what **action (response)** to take

Spoken Dialog System POMDPs

SYSTEM ACTION

Ready to answer questions.

OBSERVATION
(N-Best List)

1. what's for dinner tuesday
2. what is for dinner
3. what's dinner <noise>

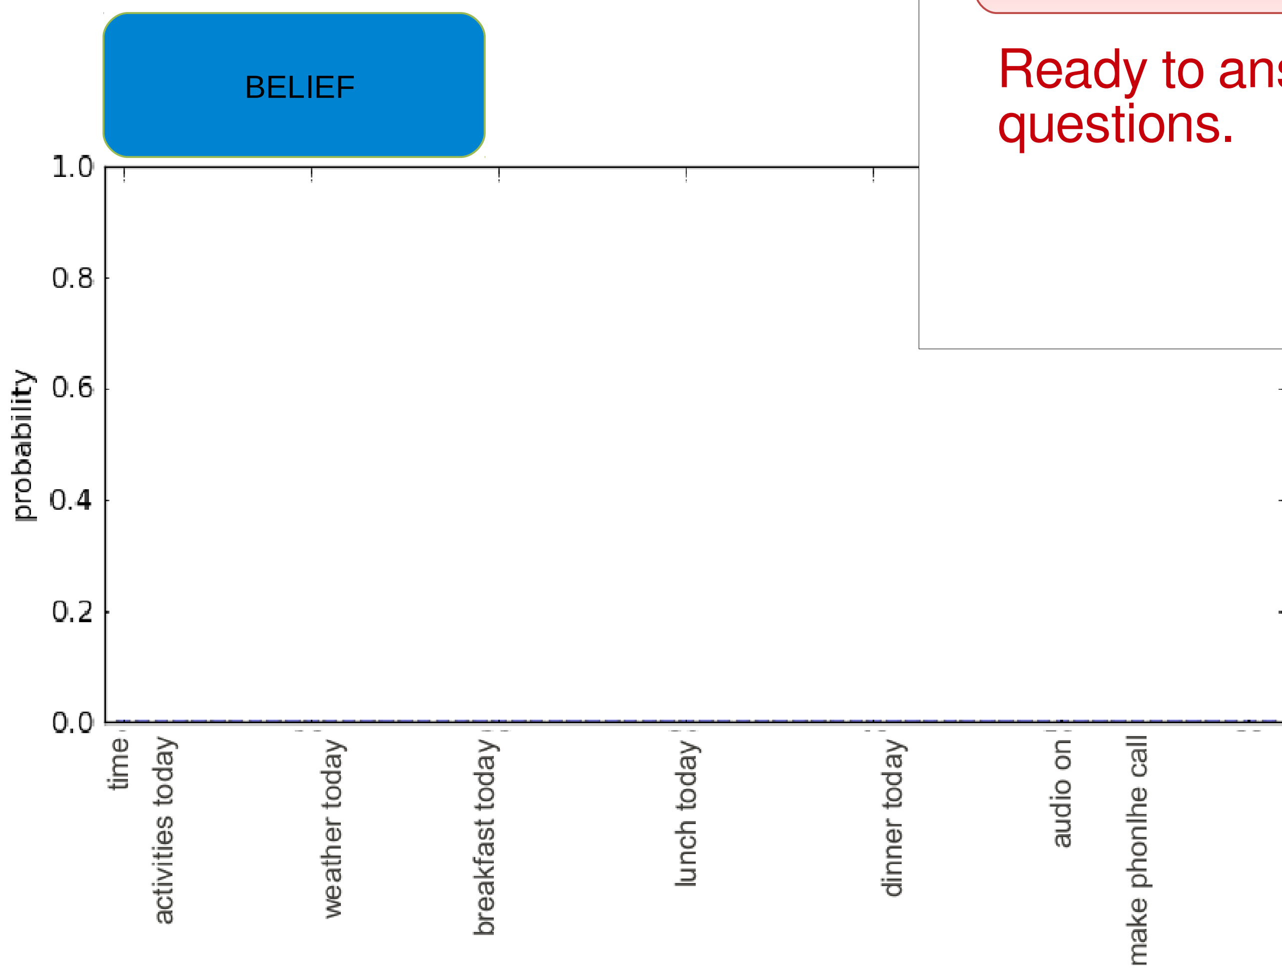

Spoken Dialog System POMDPs

SYSTEM ACTION

OBSERVATION
(N-Best List)

BELIEF

Do you want to know Tuesday's dinner menu?

1. what's for dinner tuesday
2. what is for dinner
3. what's dinner <noise>

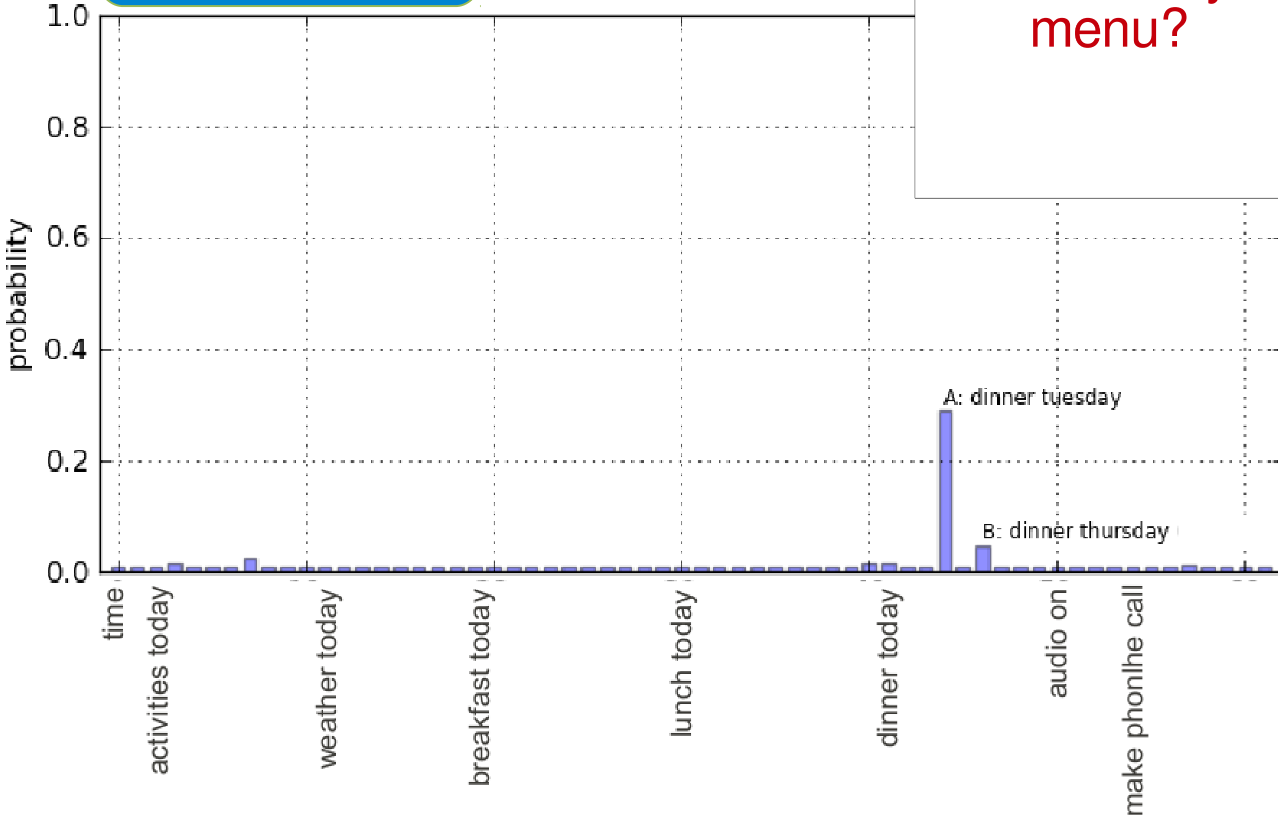

SDS-POMDP Formulation

- **States, S:** User goals
- **Actions, A:** System responses
- **Observations, Z:** Speech recognition hypotheses
- **Transition function, T = P(S'|S,A):** Model of how the user's goal changes
- **Observation function, Ω = P(Z|S,A):** Model of speech recognition “observations” for each user goal/system response
- **Reward function R(S,A):** Function that encodes desirable system responses

Toy Example: 3-State Dialog POMDP

States, S :

$s_1 = \langle \text{time} \rangle,$
$s_2 = \langle \text{weather} \rangle,$
$s_3 = \langle \text{activities} \rangle$

Observations, Z :

$z_1 =$ "time",
$z_2 =$ "weather",
$z_3 =$ "activities"

Actions, A :

$a_{1c} =$ (confirm-time), $a_{1s} =$ (show-time),
$a_{2c} =$ (confirm-weather), $a_{2s} =$ (show-weather),
$a_{3c} =$ (confirm-activities), $a_{2s} =$ (show-activities),
$a_r =$ (greet-user)

Toy Example: 3-State Dialog POMDP

- **Transition function, T = P(S'|S,A):** Assume goal does not change during a single dialog

- **Observation function, P(Z|S,A):** Assume 20% error rate

- **Reward function R(S,A):**
  - +10: correct terminal action
  - -100: incorrect terminal action
  - -5: correct confirmation question
  - -15: incorrect confirmation question
  - -10: greet user/ask to repeat

Updating the Belief

$$\underbrace{b_{n+1}(s')}_{\text{belief in state } s'} \propto \underbrace{P(z|s',a)}_{\text{observation function}} \cdot \underbrace{\sum_s P(s'|s,a)}_{\text{transition function}} \cdot \underbrace{b_n(s)}_{\text{prior belief in state } s}$$

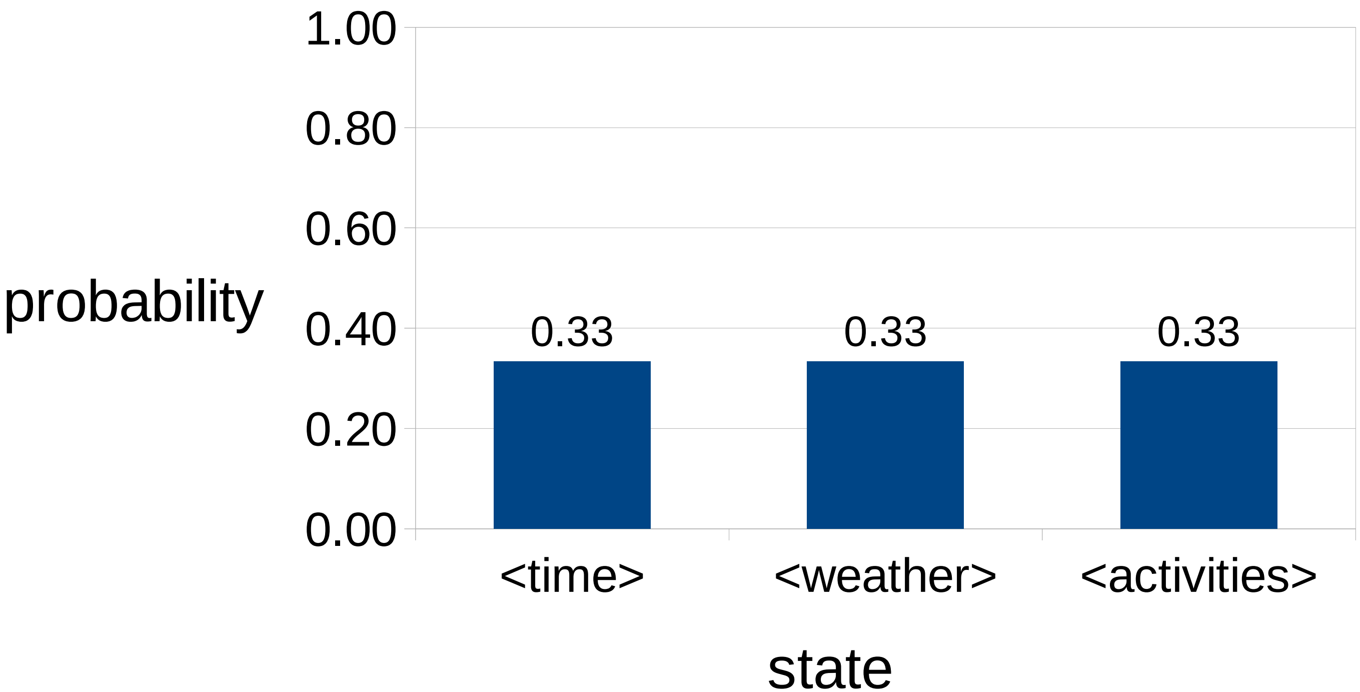

Updating the Belief

$$\underbrace{b_{n+1}(s')}_{\text{belief in state } s'} \propto \underbrace{P(z|s',a)}_{\text{observation function}} \cdot \underbrace{b_n(s')}_{\text{prior belief}}$$

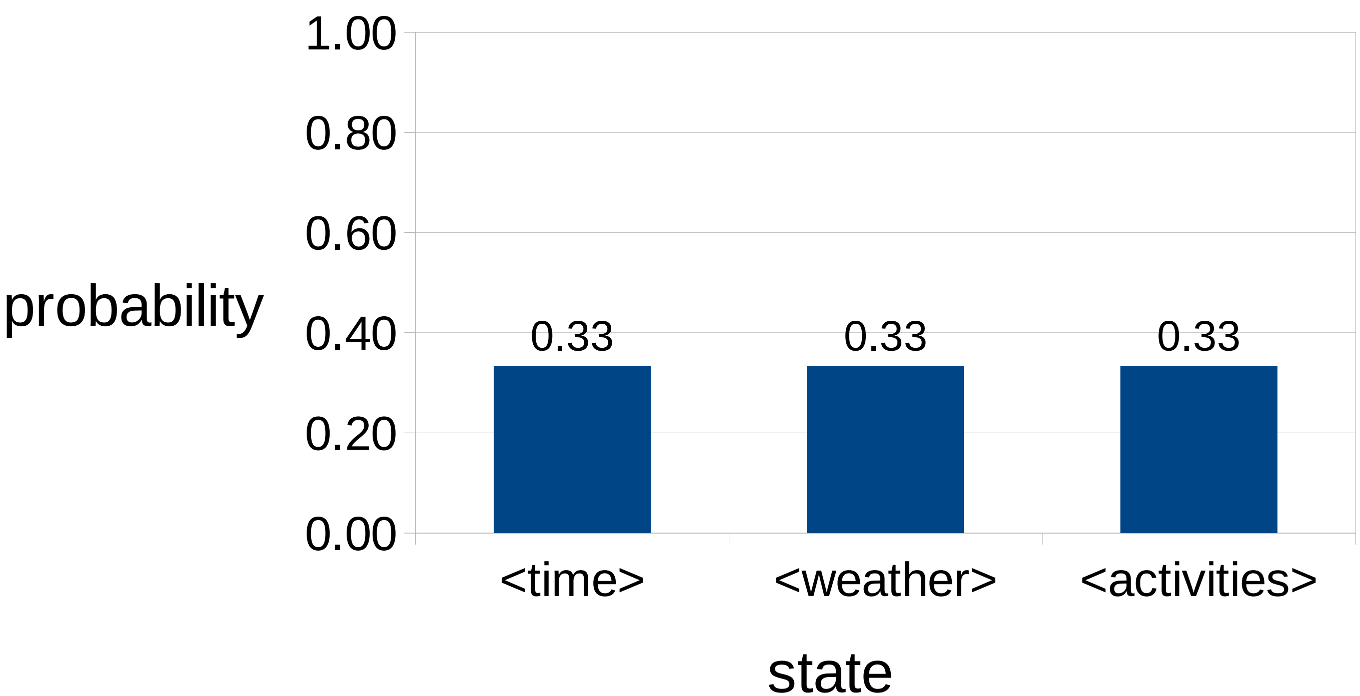

Observation:
“time”

Updating the Belief

$$\underbrace{b_{n+1}(s')}_{\text{belief in state } s'} \propto \underbrace{P(z|s', a)}_{\text{observation function}} \cdot \underbrace{b_n(s')}_{\text{prior belief}}$$

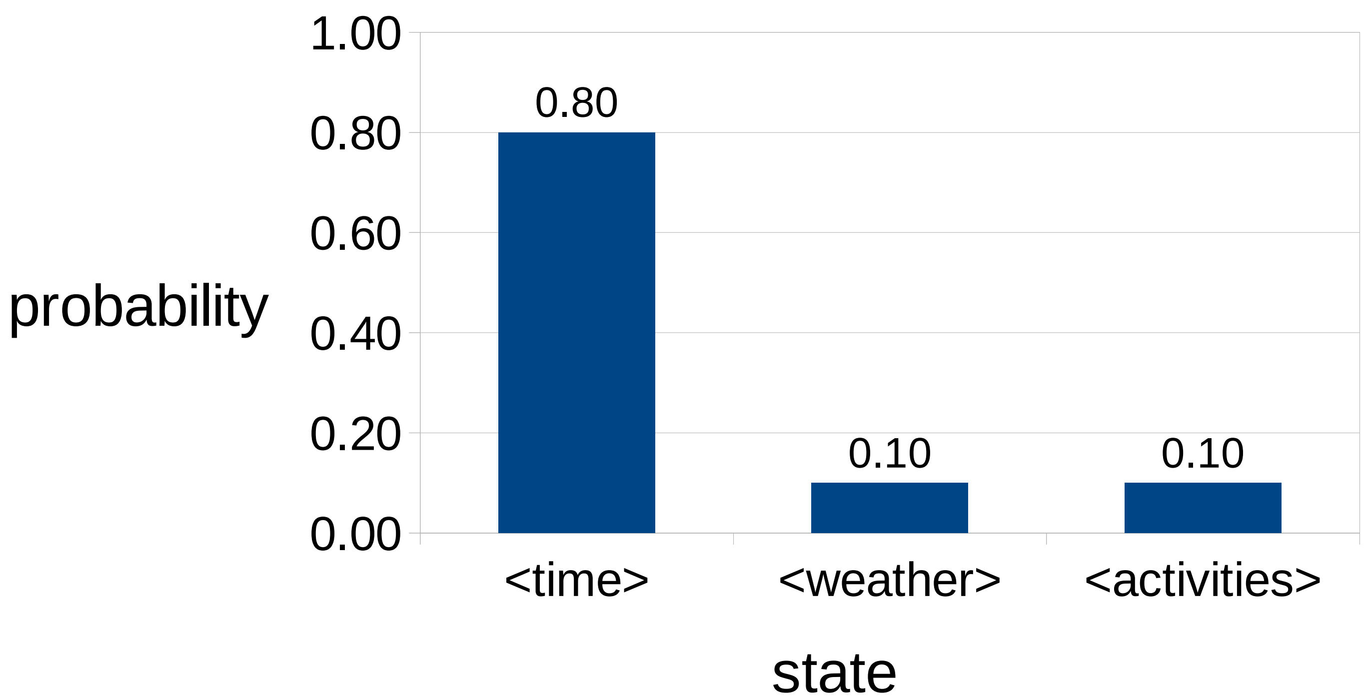

Observation:
“time”

Action:
(confirm-time)

Observation Model, Ω = P(z|s,a)

$$P(z|s,a) \triangleq P(z_d, z_c|s,a)$$

z_d: concept (e.g. “time”, “weather”, “activities”)

z_c: confidence score ($0 < z_c < 1$)

Apply chain rule:

$$P(z_d, z_c|s,a) = \underbrace{P(z_d|s,a)}_{\text{language model}} \cdot \underbrace{P(z_c|z_d,s,a)}_{\text{confidence score model}}$$

Effect of Confidence Score Model

$$\underbrace{b_{n+1}(s')}_{\text{belief in state } s'} \propto \underbrace{P(z_d|s,a)}_{\text{language model}} \cdot \underbrace{P(z_c|z_d,s,a)}_{\text{confidence score model}} \cdot \underbrace{b_n(s')}_{\text{prior belief in state } s'}$$

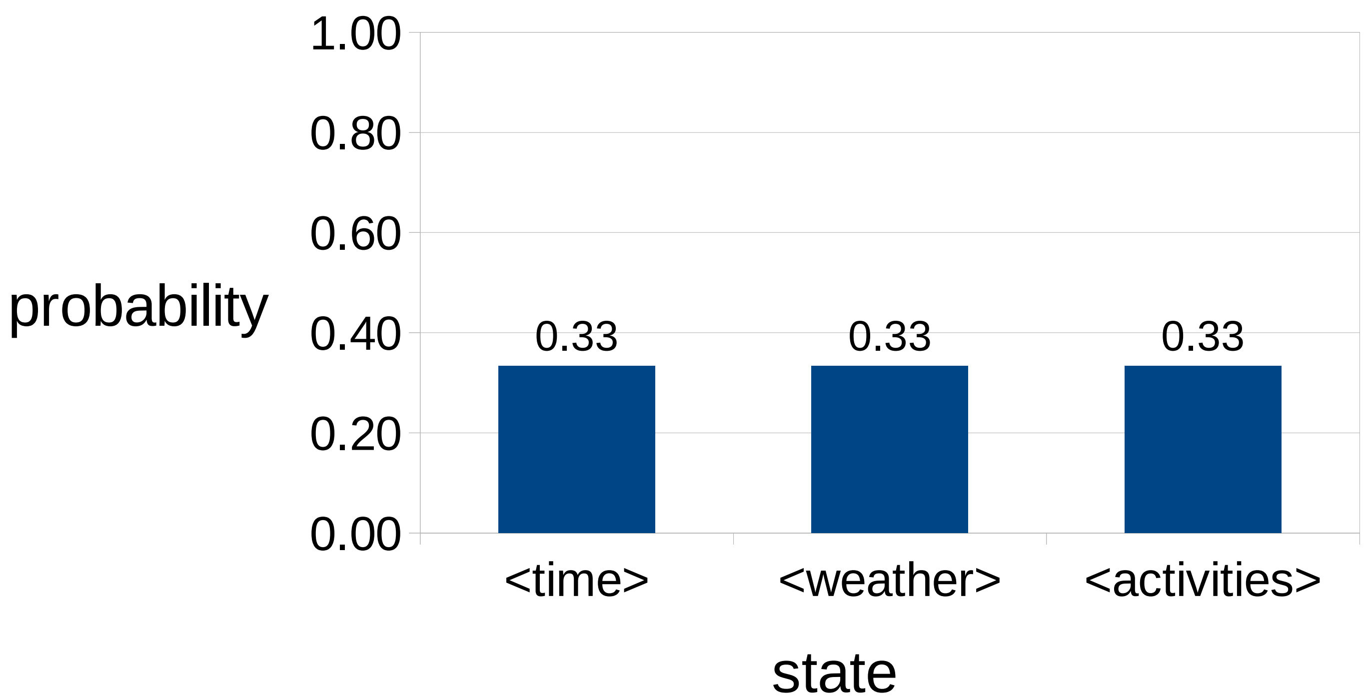

Observation:
z_d: “time”

Updating the Belief

$$\underbrace{b_{n+1}(s')}_{\text{belief in state } s'} \propto \underbrace{P(z_d|s,a)}_{\text{language model}} \cdot \underbrace{P(z_c|z_d,s,a)}_{\text{confidence score model}} \cdot \underbrace{b_n(s')}_{\text{prior belief in state } s'}$$

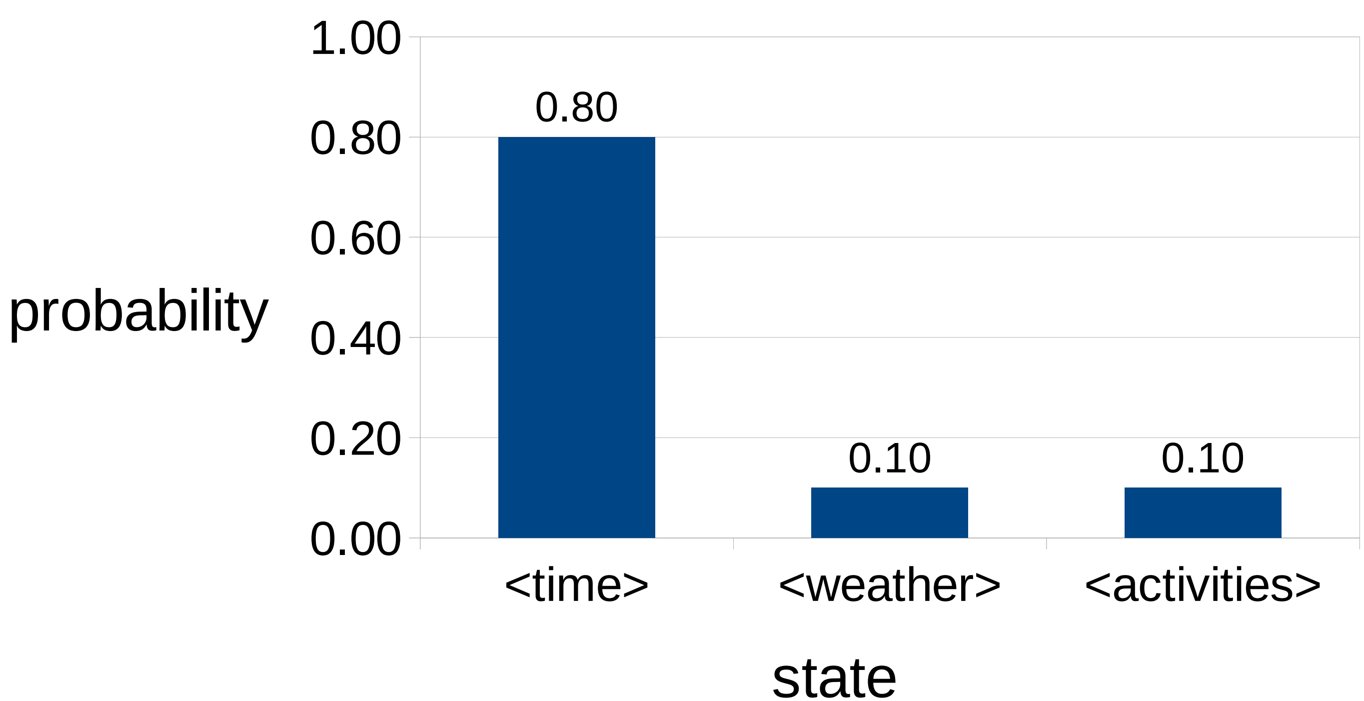

Observation:

z_d: “time”

Updating the Belief

$$\underbrace{b_{n+1}(s')}_{\text{belief in state } s'} \propto \underbrace{P(z_d|s,a)}_{\text{language model}} \cdot \underbrace{P(z_c|z_d,s,a)}_{\text{confidence score model}} \cdot \underbrace{b_n(s')}_{\text{prior belief in state } s'}$$

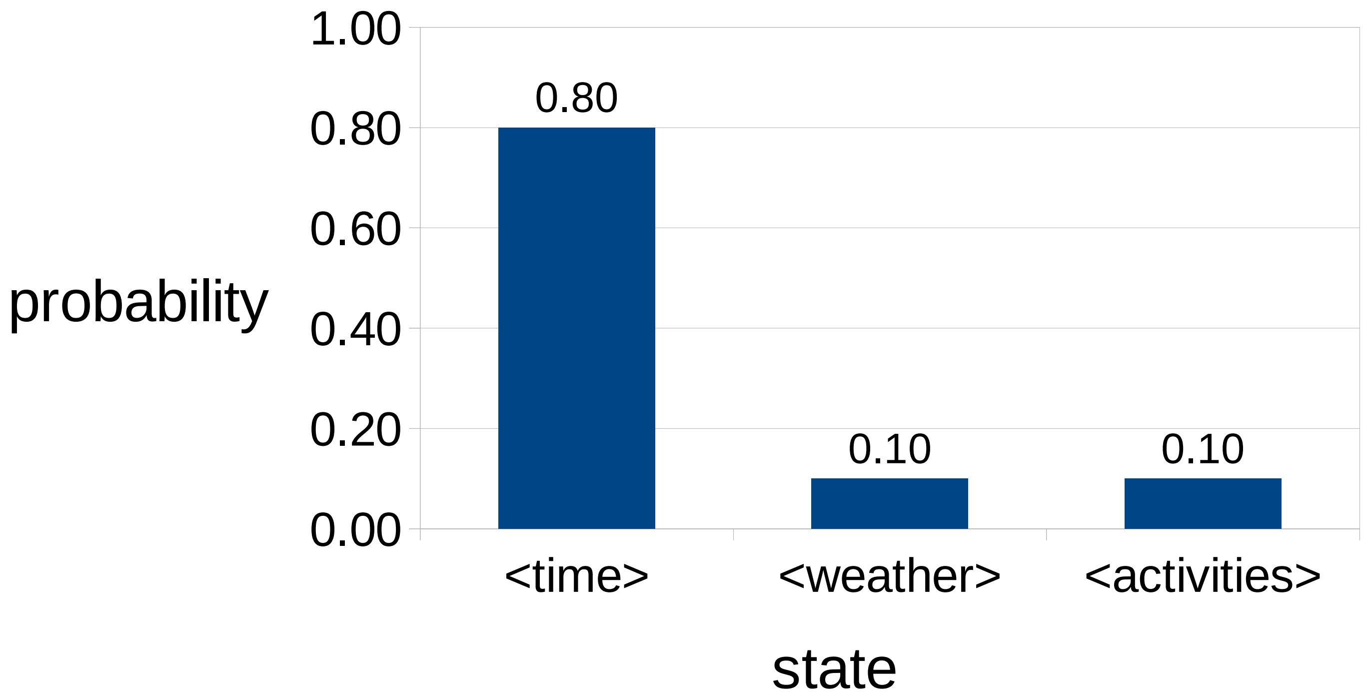

Observation:

z_d: “time”

z_c: 0.95

Updating the Belief

$$\underbrace{b_{n+1}(s')}_{\text{belief in state } s'} \propto \underbrace{P(z_d|s,a)}_{\text{language model}} \cdot \underbrace{P(z_c|z_d,s,a)}_{\text{confidence score model}} \cdot \underbrace{b_n(s')}_{\text{prior belief in state } s'}$$

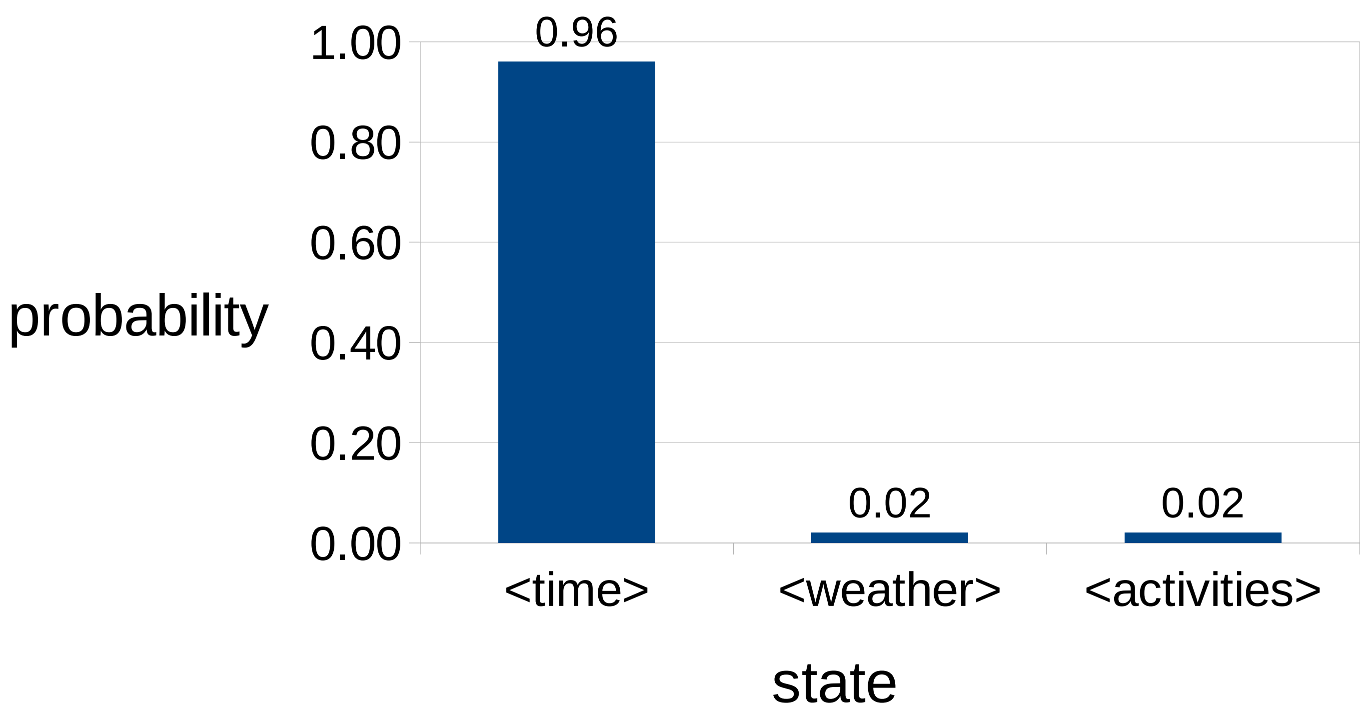

Observation:
z_d: "time"
z_c: 0.95

Action:
(show-time)

Updating the Belief

$$\underbrace{b_{n+1}(s')}_{\text{belief in state } s'} \propto \underbrace{P(z_d|s,a)}_{\text{language model}} \cdot \underbrace{P(z_c|z_d,s,a)}_{\text{confidence score model}} \cdot \underbrace{b_n(s')}_{\text{prior belief in state } s'}$$

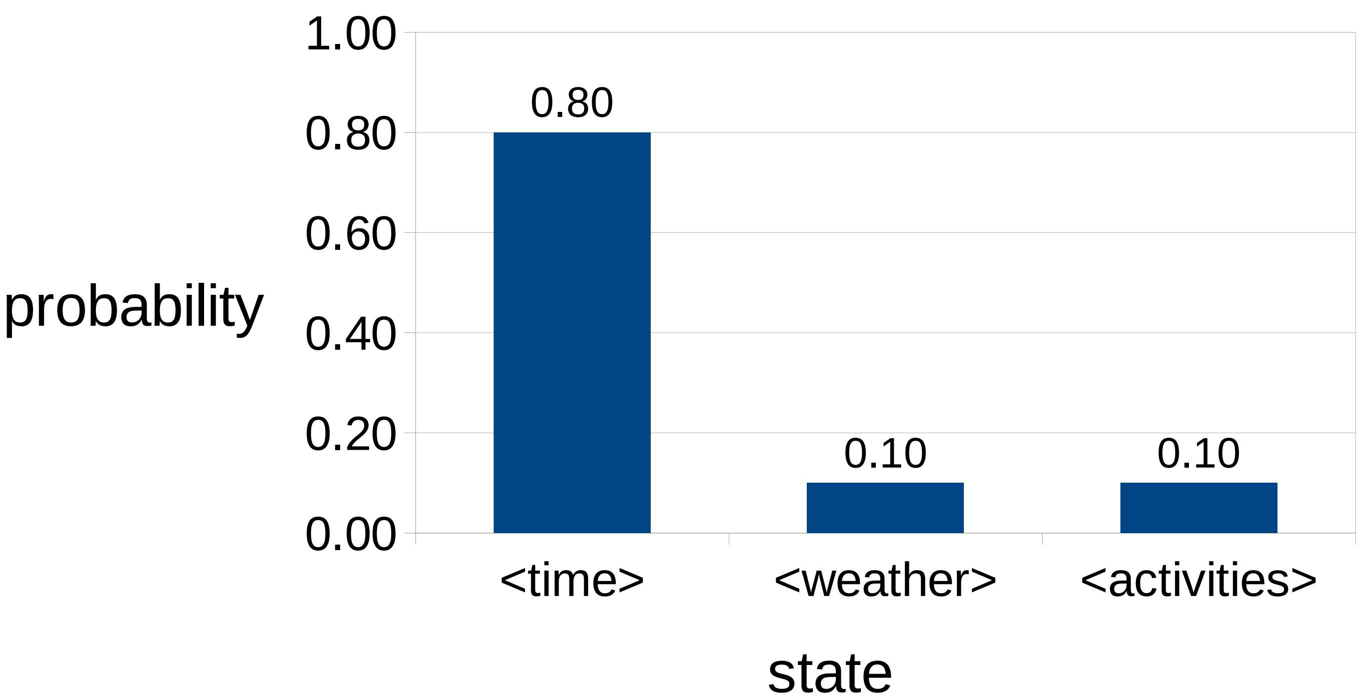

Observation:

z_d: "time"

z_c: 0.15

Updating the Belief

$$\underbrace{b_{n+1}(s')}_{\text{belief in state } s'} \propto \underbrace{P(z_d|s,a)}_{\text{language model}} \cdot \underbrace{P(z_c|z_d,s,a)}_{\text{confidence score model}} \cdot \underbrace{b_n(s')}_{\text{prior belief in state } s'}$$

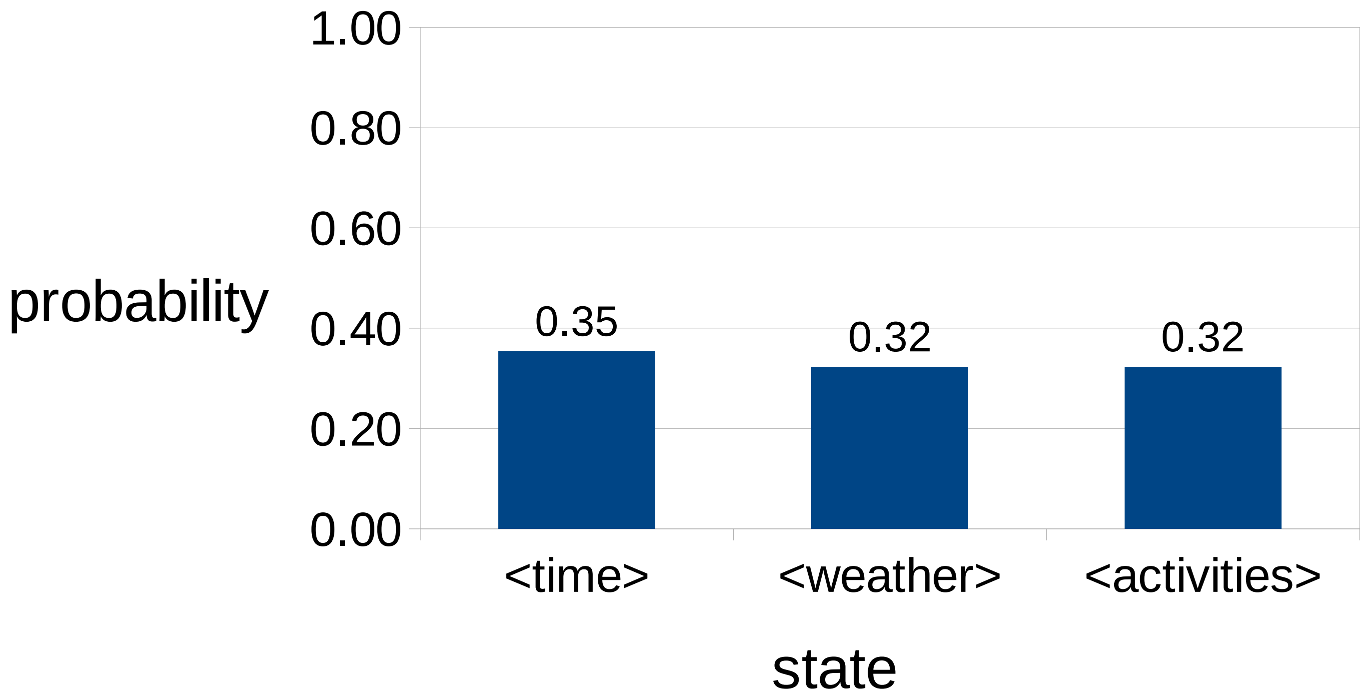

Observation:
z_d: “time”
z_c: 0.15

Action:
(ask-repeat)

Dialog System

Experimental Design and Results

SDS-POMDP Formulation

- **States, S:** 62 (time, weather, activity schedules, menus, phone calls)
- **Actions, A:** 125 (62 “submit-s”, 62 “confirm-s”, ask-initial question)
- **Observations, Z:**
  - 65 discrete concepts (62 possible states, YES, NO, NULL)
  - Confidence score between 0 and 1
- **Transition function, T = P(S'|S,A):** Assume goal does not change during a dialog
- **Observation function, P(Z|S,A):** Learn from hand-labeled training set of 2701 utterances
- **Reward function R(S,A):** Specified similar to toy example

Confidence Scoring of Utterances

- Boosting (AdaBoost) to learn a confidence score function

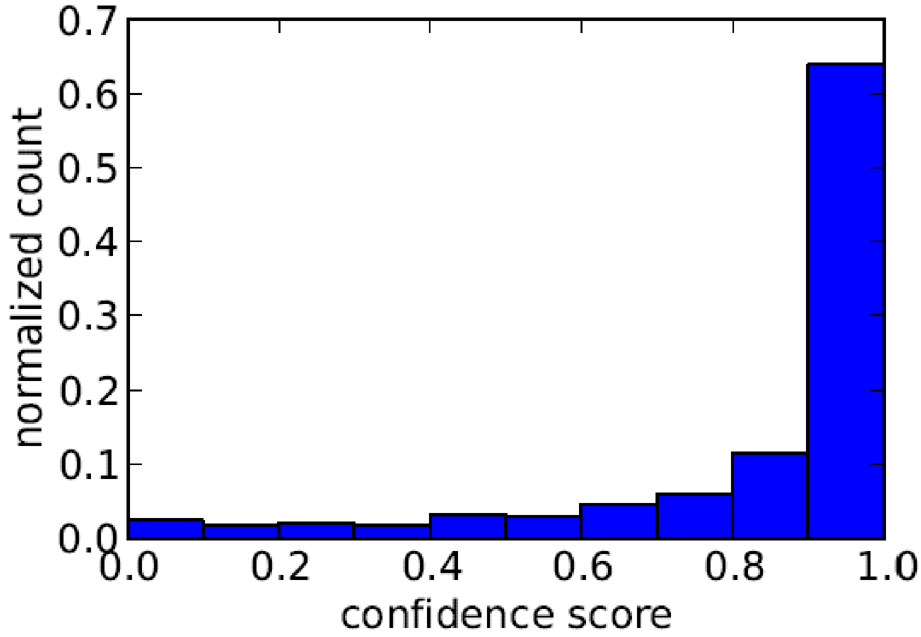

$P(z_c|\text{correct observation})$

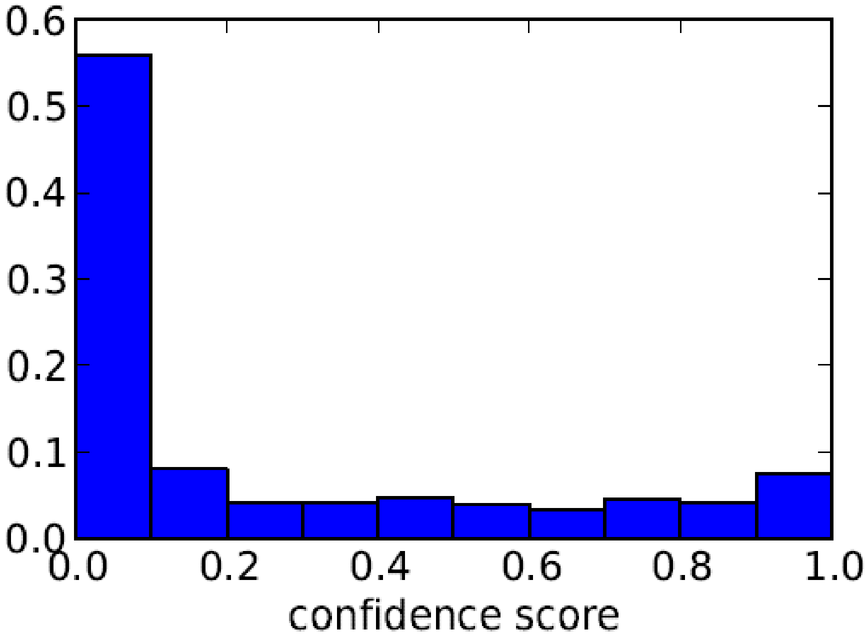

$P(z_c|\text{incorrect observation})$

Confidence Scoring of Utterances

- Boosting (AdaBoost) to learn a confidence score function

Feature Category	**Examples**
Concept-level	parse success
ASR scores	acoustic, language, and total model scores; difference between top score and second-highest hypothesis score
Word-/sentence-level	fraction of stop words; presence of multiple concepts; presence of highly mis-recognized words or often merged/split word pairs
n-best list	concept entropy of *n*-best list

Within-Subjects User Study

- Comparison of two dialog management strategies (20 dialog prompts/dialog manager)
  - Confidence score threshold dialog manager (ask user to repeat if confidence score < 0.7)
  - SDS-POMDP dialog manager

Experimental Setup

- 14 users (7 target, 7 control)
- Users presented with dialog prompts in random order
- 40 dialogs per user (20 with threshold, 20 with POMDP)

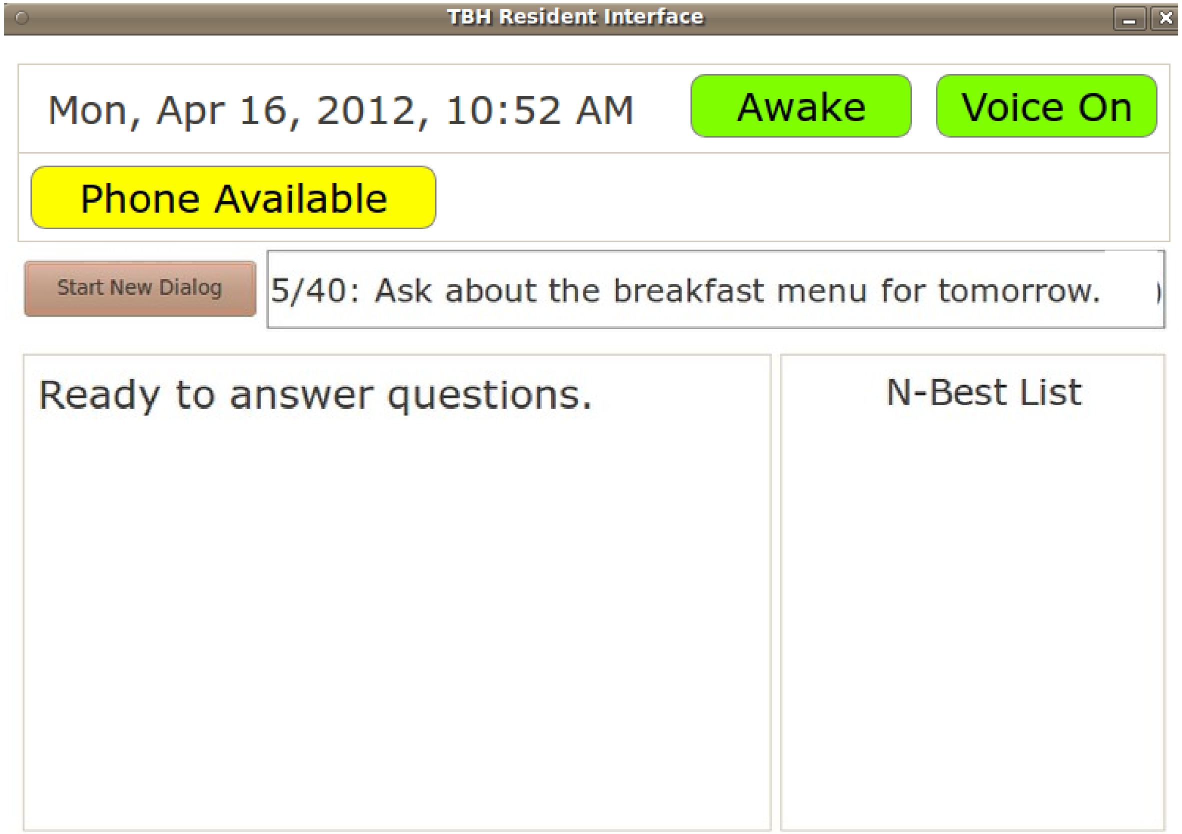

Within-Subjects User Study: Metrics

- Number of dialogs (out of 20) successfully completed
  - “successfully completed”: within one minute
- Average time to complete dialog

Baseline Threshold Dialog Manager vs. POMDP Dialog Manager

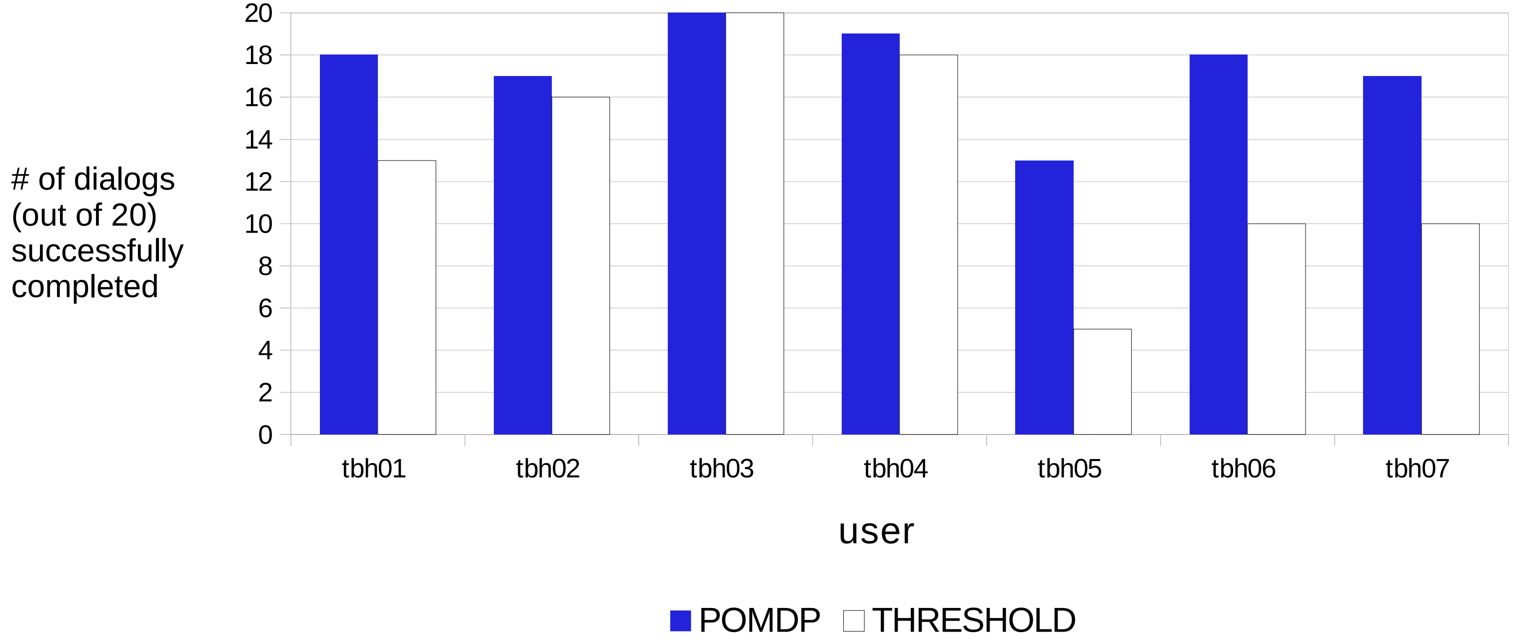

SDS-POMDP: 17.4 ± 0.9
Threshold: 13.1 ± 0.9

One-way repeated measures ANOVA: Significant ($p=.02$) effect of POMDP on dialog completion rates

Baseline Threshold Dialog Manager vs. POMDP Dialog Manager

- Improvements are more pronounced among speakers with high error rates

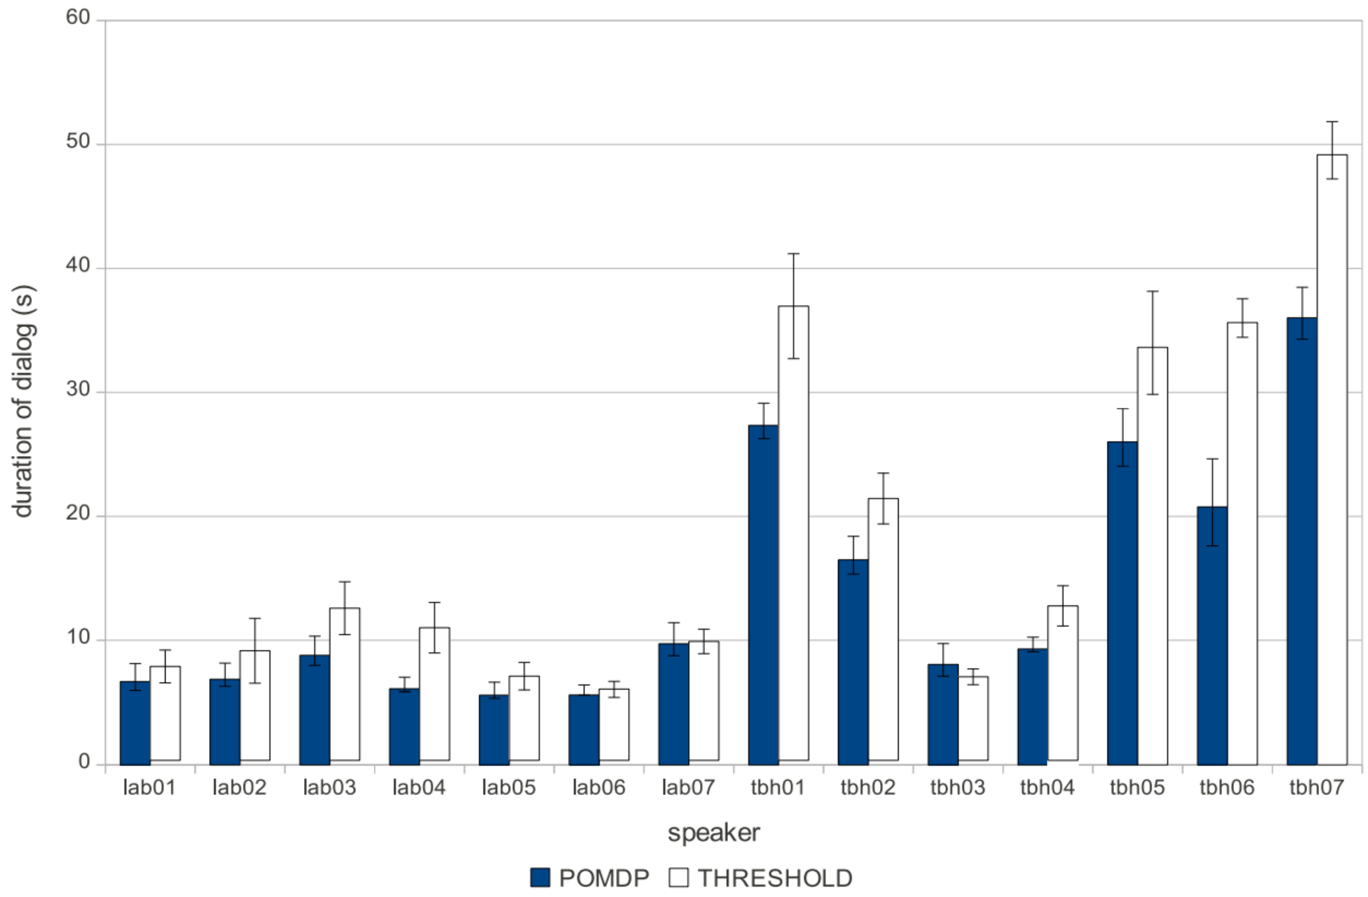

SDS-POMDP Discussion

- Advantages of SDS-POMDP:
  - Belief distribution includes information from past utterances
  - Observation model produces a “variable threshold” for each goal
- Limitations of SDS-POMDP:
  - Off-model errors can cause user to be “stuck” in undesirable belief distributions

Contributions

Problem identification: Understanding the needs of users (residents at The Boston Home)

End-to-end system development: Collecting data, training models, and implementing a partially observable Markov decision process (POMDP) dialogue manager

Experimental evaluation: Validating the POMDP-based spoken dialog system with target users

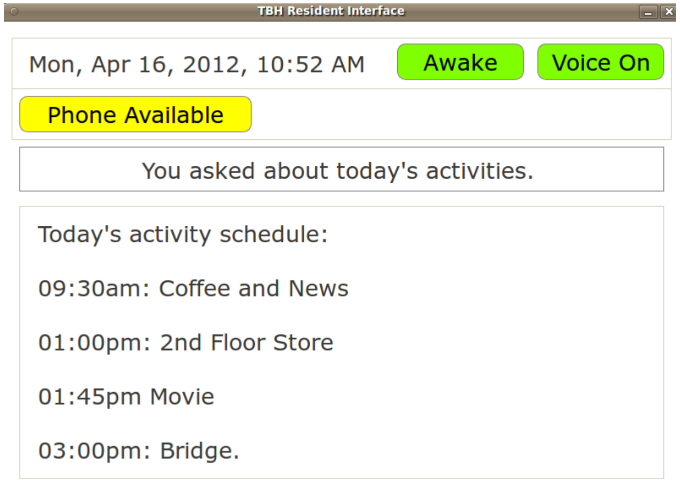

wli@csail.mit.edu
http://people.csail.mit.edu/wli/