

ARTICLE TYPE

Spoken Language Interaction with Model Uncertainty: An Adaptive Human-Robot Interaction System

Finale Doshi^{a*} and Nicholas Roy^b

^aCambridge University, Cambridge, UK; ^bMassachusetts Institute of Technology, Cambridge MA, USA

Spoken language is one of the most intuitive forms of interaction between humans and agents. Unfortunately, agents that interact with people using natural language often experience communication errors and do not correctly understand the user's intentions. Recent systems have successfully used probabilistic models of speech, language, and user behavior to generate robust dialog performance in the presence of noisy speech recognition and ambiguous language choices, but decisions made using these probabilistic models are still prone to errors due to the complexity of acquiring and maintaining a complete model of human language and behavior.

In this paper, we describe a decision-theoretic model for human-robot interaction using natural language. Our algorithm is based on the Partially Observable Markov Decision Process (POMDP), which allows agents to choose actions that are robust not only to uncertainty from noisy or ambiguous speech recognition but also unknown user models. Like most dialog systems, a POMDP is defined by a large number of parameters that may be difficult to specify *a priori* from domain knowledge, and learning these parameters from the user may require an unacceptably long training period. We describe an extension to the POMDP model that allows the agent to acquire a linguistic model of the user online, including new vocabulary and word choice preferences. Our approach not only avoids a training period of constant questioning as the agent learns, but also allows the agent to actively query for additional information when its uncertainty suggests a high risk of mistakes. We demonstrate our approach both in simulation and on a natural language interaction system for a robotic wheelchair application.

Keywords: dialog management, human-computer interface, adaptive systems, online learning, partially observable Markov decision processes

1. Introduction

Spoken language is one of the most intuitive forms of interaction between humans and agents. The ability for an agent to understand natural language commands can be especially useful when interacting with people who may not be trained to use more conventional robotic interfaces. For example, when intelligent robots are used as a medical assistive technology, both care-givers and care recipients rarely have experience with robots. Furthermore, as most people reach a comfortable level with technology, they often begin interacting with the technology as they would another person Reeves and Nass (1996). As a result, natural language is an attractive way for robots to interact with users, minimizing the need for training and increasing the likelihood that the robot will be accepted by the intended users.

Spoken language interaction with a mobile robot is most often implemented as a process consisting of the following systems:

*Corresponding author. Email: finale@alum.mit.edu

- A speech recognition system to recover the sequence of words uttered by the user.
- A dialog management system to infer the user intent (e.g., a command for the robot to perform some action).
- A robot control for responding to the user intent (e.g., go to some location).

For a robot to be able to use natural language to interact with people, it is critical that the robot react to the user's language in ways that are consistent with the user's expectations of reasonable behavior. In most human-robot interaction domains, and especially in medical assistive domains, a robot that violates user expectations and appears to make errors periodically will be quickly ignored as untrustworthy by the human users. As a result, natural language interaction must be robust to various errors in the models it uses to predict user intent.

One challenge is simply correctly identifying the spoken words, since speech recognition errors will make the input to the dialog system a noisy process. For example, the system may hear the words "coffee machine" when the user asks the robot to go to "copy machine." While speech recognition errors may be mitigated to some extent by asking the user to use only acoustically distinct keywords when speaking to the system, a system that does not model the likelihood of recognition errors and act accordingly will be brittle; a robust system must be able to infer user intent under uncertainty. An ideal dialog manager will also adapt to the user, instead of requiring the user to adapt to the robot's vocabulary set—for example, it should learn that words such as "kitchen," "kettle," and "microwave," may refer to the same location.

Another set of challenges arise because people may have different preferences in how a system responds to the interaction. Unfortunately, humans are challenging to model, and real-world interaction domains typically require a large number of parameters that are difficult to specify *a priori*. Instead of requiring a lengthy training period, it is desirable that the dialog manager be able to take rough estimate of the model and improve it online. A dialog manager that can represent its confidence in its model of the user and can choose actions appropriately will not only to avoid errors due to model uncertainty, but will also be able take actions to learn more about the user and improve the model for the future.

If the dialog manager has an accurate probabilistic model of word error rates and user preferences, the planning framework known as the Partially Observable Markov Decision Process (POMDP) can be used to compute decision-theoretic policies for the interaction. The optimal POMDP policy trades between asking questions to reduce uncertainty (thus avoiding errors), and fulfilling the user's request within a reasonable amount of time. The ability to reason about actions under uncertainty has led to the successful use of POMDPs in several assistive health-care Roy et al. (2000), Hoey et al. (2005), Fern et al. (2007) and dialog management domains (Litman et al., 2000, Williams and Young, 2005).

Although POMDPs have been used previously for human-robot interaction, the problem of learning these models online has not been addressed. After formally describing the POMDP dialog model in section 2, we extend it to include model uncertainty in section 3. In section 4 we describe an approach to passive learning of the interaction model, in which the system receives feedback after every action. We present an algorithm that allows the dialog manager to interact with the user while improving its model, an approach that is particularly important for a robotic wheelchair application that must adapt to users and new environments. This passive learner can easily be run in real-time settings, and we show in simulation that this approach is usually effective in converging to the optimal policy.

Unfortunately, the approach of passively learning a model from reward feedback



Figure 1. Our dialog manager allows for more natural human communication with a robotic wheelchair.

suffers from several drawbacks. First, the dialog manager must make mistakes to learn about the consequences of a poor decision, leading to user dissatisfaction. Second, requiring the user to supply reward feedback after each action may be tedious, leading to frustration and inaccurate results. Finally, without being explicitly aware of the quality of its knowledge, the dialog manager cannot know if its decisions are based on a solid or unclear understanding of the world. We address all three issues in section 5, where we show how the dialog manager can explicitly maintain an estimate of the quality its knowledge about the world. We also introduce the concept of a meta-query, that is, a question about an action that the dialog manager should take. These meta-queries take an intuitive form:

“I think you definitely want me to go to the printer. Should I go to the printer?”

By introducing these meta-queries, the dialog manager can discover new vocabulary, word choices, and user preferences such as risk aversion.

Throughout this paper, we focus on natural language interaction for assistive technology such as the wheelchair shown in Figure 1. Our goal is to design an adaptable human-robot interaction system, or dialog manager, that allows both the user of the wheelchair and a caregiver to give natural instructions to the wheelchair, as well as ask the wheelchair computer for general information that may be relevant to the user’s daily life. In particular, we focus on the scenario in which the wheelchair user is trying to tell the wheelchair where to drive next, and the dialog manager must overcome an unfamiliar, noisy speech recognition system to discern the user’s intent.¹ We demonstrate the active-learning approach both in simulation and in several dialogs with human subjects and a real speech-recognition system.

2. The POMDP Model

Formally, a POMDP consists of the n-tuple $\{S, A, O, T, \Omega, R, \gamma\}$. The set S is a set of states, which, in the dialog management setting, correspond to what the user actually desires from the interaction. In our robotic wheelchair scenario, the states correspond to locations to which the user wishes to go. The dialog manager cannot directly observe this desired location and must infer it from a set of observations O . For our basic dialog manager, we take the observations to be the number of times certain keywords occur in an n-best list from the speech recognition system. However, the POMDP can easily accommodate observations that are comprised of more sophisticated natural language processing of the speech recognition output. The set A represents actions that the dialog manager may take.

¹We will not discuss the navigation task of driving to the desired location once it is chosen; this is outside the scope of this article.

The actions may include asking a question, commanding the robot to perform some physical movement, or doing nothing.

In Figure 2, we show a cartoon of a simple dialog model. Initially, we model the user as being in a “start” state. Then, at some point in time, the user speaks to the robot to indicate that he or she wants it to perform a task. We denote this step by the set of vertical stack of nodes in the center of the model. Each node represents a different task. The dialog manager must now interact with the user to determine what is wanted. Once the task is successfully completed, the user transitions to the right-most “end” node, in which he or she again does not desire anything from the robot. We note that it can be easily augmented to handle more complex scenarios. For example, by including the time of day as part of the state, we can model the fact that the user may usually wish to go to certain locations in the morning and other locations in the afternoon.

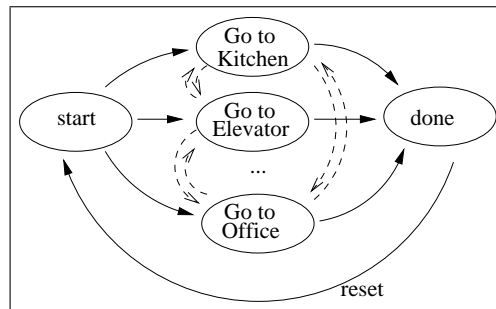


Figure 2. A toy example of a dialog POMDP. The nodes in the graph are different states of the dialog (i.e., user requests). Solid lines indicate likely transitions; we assume that the user is unlikely to change their request before their original request is fulfilled. The system automatically resets once we reach the end state.

The transition function T and observation function O of the POMDP model describe how the sets of states, actions, and observations behave. The transition function $T(s'|s, a)$ places a probability distribution over the states to which the user in state s may transition if the dialog manager takes action a . Similarly, the observation function $\Omega(o|s, a)$ encodes what keywords the dialog manager is likely to hear given its most recent action. For example, the model might encode that the keyword “coffee” is commonly heard if the dialog manager asks “Where would you like to go?” and the user wishes to go to the coffee machine. The POMDP model allows for different observations to be received in the same state (for example, both “coffee” and “machine” may be commonly heard if the user wishes to go to the coffee machine) but the assumption of the standard POMDP is that we do know the probabilities of hearing each keyword.

The reward function $R(s, a)$ specifies the immediate reward the dialog manager receives for taking action a when the user is in state s . In the wheelchair scenario, where the dialog manager is trying to determine where the user wishes to go, a clarification action such as “Do you want to go to the coffee machine?” may carry a small negative reward for causing a minor inconvenience, while incorrectly driving to the coffee machine may carry a large penalty for taking the user to the wrong location. To maximize the likelihood of gaining a high positive reward for completing a task and avoiding a large negative reward for taking an incorrect movement, the dialog manager may decide to accept small penalties by asking clarification questions. Finally, the discount factor $\gamma \in [0, 1)$ measures the relative importance of current and future rewards.

In making decisions, the dialog manager can only choose actions based on the knowledge of its past actions and observations it has received. In general, the optimal action to take now will depend on *all* prior actions and observations; however,

it can be shown that when it comes to making a decision, what really matters is the dialog manager’s belief of what the user may desire. More formally, the belief is a probability distribution over states. If the agent takes some action a and hears observation o from an initial belief b , it can easily update its belief using Bayes rule:

$$b^{a,o}(s) = \frac{\Omega(o|s', a) \sum_{s \in S} T(s'|s, a)b(s)}{\sum_{\sigma \in S} \Omega(o|\sigma, a) \sum_{s \in S} T(\sigma|s, a)b(s)} \quad (1)$$

This probability distribution will evolve as the dialog manager asks clarification questions and receives responses.

Intuitively, we can see how the belief can be used to select an appropriate action. For example, if the dialog manager believes that the user may wish to go to either the coffee machine or the copy machine (but not the printer), then it may ask the user for clarification before commanding the wheelchair to one of the locations. More formally, we call the mapping from beliefs to actions a policy. We represent this mapping using the concept of a value function $V(b)$. The value of a belief is defined to be the expected long-term reward the dialog manager will receive if it starts a user interaction in belief b . The optimal value function is piecewise-linear and convex, so we represent V with the vectors V_i ; $V(b) = \max_i V_i \cdot b$. The optimal value function satisfies the Bellman equation Bellman (1957):

$$V(b) = \max_{a \in A} Q(b, a),$$

$$Q(b, a) = R(b, a) + \gamma \sum_{o \in O} \Omega(o|b, a)V(b_a^o), \quad (2)$$

where $Q(b, a)$ represents the expected reward for starting in belief b , performing action a , and then acting optimally. The belief b_a^o is b after a Bayesian update of b using equation 1, and $\Omega(o|b, a)$, the probability of seeing o after performing a in belief b ($\sum_{s \in S} \Omega(o|s, a)b(s)$).

The Bellman equation may be solved iteratively:

$$V_n(b) = \max_{a \in A} Q_n(b, a), \quad (3)$$

$$Q_n(b, a) = R(b, a) + \gamma \sum_{o \in O} \Omega(o|b, a)V_{n-1}(b_a^o). \quad (4)$$

Each iteration, or *backup*, brings the value function closer to its optimal value Gordon (1995). Once the value function has been computed, it is used to choose actions. After each observation, we update the belief using equation 1 and then choose the next action using $\arg \max_{a \in A} Q(b, a)$ with $Q(b, a)$ given in equation 2. The exact solution to equation 3 using an iterative backup approach is exponentially expensive, so we approximate the true backup operation by backing up at only a small set of beliefs Pineau et al. (2003).

3. Modeling POMDP Uncertainty

The behavior of the dialog manager derived from solving equations (3) and (4) depends critically on accurate choices of the transition probabilities, observation probabilities and the reward. For example, the observation parameters affect how the system associates particular keywords with particular requests. Similarly, the

reward function affects how aggressive the dialog manager will be in assuming it understands a user’s request, given limited and noisy information. An incorrect specification of the dialog model may lead to behavior that is either overly optimistic or conservative, depending on how accurately the model captures the user’s expectations on the interaction.

At the same time, learning all the parameters required to specify a complete dialog model can require a prohibitively large amount of data. However, while the model parameters may be difficult to specify exactly, we can often provide the dialog manager with an initial estimate of the model parameters that will generate a reasonable policy that can be executed while the model is improved. For example, even though we may not be able to attach an exact numerical value to driving a wheelchair user to the wrong location, we can at least specify that this behavior is undesirable. Similarly, we can specify that the exact numerical value is initially uncertain. As data about model parameters accumulate, the parameter estimates should converge to the correct underlying model with a corresponding reduction in uncertainty.

In order to formally specify such an adaptive interaction model, we extend the POMDP model to include the dialog parameters. Recall from the previous section we introduced the POMDP as a model of decision making conditioned on the current belief, that is, the probability of how likely a user is to be making a particular request. Figure 3(a) depicts this conventional model, where the arrows in the graph show which parts of the model affect each other from time t to $t + 1$. Although the variables below the “hidden” line in Figure 3(a) are not directly observed by the dialog manager, the parameters defining the model (i.e., the parameters in the function giving the next state) are fixed and known *a priori*. For instance, the reward at time t is a function of the state at the previous time and the action chosen by the dialog manager.

If the model parameters are not known *a priori* because the model is uncertain—for example, how much reward is received by the agent given the previous state and the action selected—then we extend the concept of the belief to also include the agent’s uncertainty over possible models. In this new representation, which we call the *model-uncertainty POMDP*, both the user’s request and the model parameters are hidden. Figure 3(b) shows this extended model, in which the reward at time t is still a function of the state at the previous time and the action chosen by the dialog manager, but the parameters are not known *a priori* and are therefore hidden model variables that must be estimated along with the user state. As the system designer, we can encode our knowledge of the system in the dialog manager’s initial belief over what dialog models it believes are likely—a Bayesian prior over models—and let the agent improve upon this belief with experience.

We will restrict our discussion to inference of unknown transition, observation and reward function parameters as it is often reasonable to assume that the sets S , A , and O are known. For example, the states S correspond to the particular requests the system was designed to service. Similarly, we can pre-specify what kinds of statements and movements a robot can make during an interaction. Specifying the observation space is more tricky, but we assume that the designer has sufficient knowledge of the application to create a list of keywords. We also note that we can always add additional states, actions, or observations. The techniques for doing so are outside the scope of this article, but it can be done without having to resolve the entire policy.

The need to represent the beliefs over models raises the question of how to represent beliefs over model parameters. The user state space is a discrete state space, so a standard histogram or multinomial distribution could be used. However,

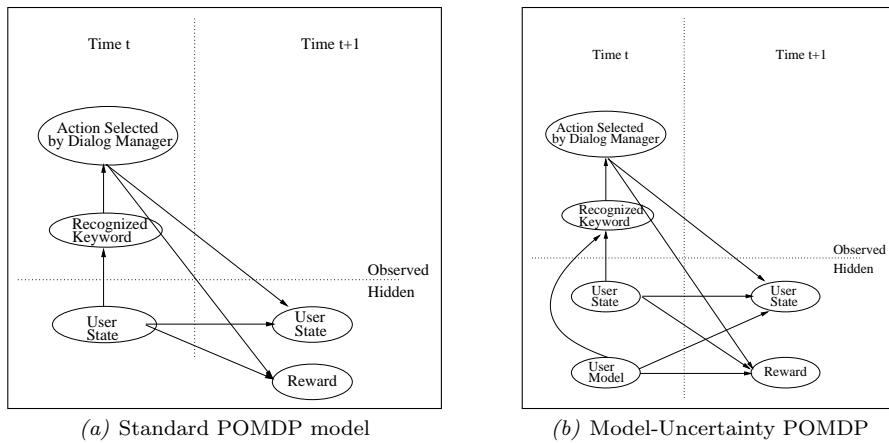


Figure 3. (a) The standard POMDP model. (b) The extended POMDP model. In both cases, the arrows show which parts of the model are affected by each other from time t to $t+1$. Not drawn are the dependencies from time $t+1$ onwards, such as the user state and user model's effect on the recognized keyword at time $t+1$.

the rewards are continuous-valued and T and O are distributions themselves; a distribution over T is effectively a distribution over distributions. We describe how we model the uncertainty in different model parameters below.

3.1. Reward Uncertainty

Because each reward is a single scalar value conditioned on the state and action, we have considerable flexibility over distributions to use. To set an absolute scale for the reward values, we fix a large positive reward value for correctly satisfying a user's request, and a small penalty for confirming a correct request with the user (for the minor inconvenience of causing the user to communicate with the robot, instead of immediately satisfying the user's request). These two values set a scale and spread for the remaining reward values.

Our initial algorithm requires knowledge only of the expected reward values, so in theory, any distribution may be used. Our first approach employs Gaussian distributions, as they are easy to update. While it may be reasonable in some cases for a system designer to specify to the agent reasonably accurate expected reward values for each action, we relax this requirement in section 5, where we require the designer only to specify a range for each reward parameter, and we assume that the reward values are uniformly distributed between these ranges.

3.2. Transition and Observation Uncertainty

Since the state space is discrete, the T and Ω parameters define collections of histograms or multinomial distributions. A prior over the transition and observation parameters must therefore be a distribution over possible histograms. We choose a Dirichlet prior because it provides a measure over valid multinomial distributions and, as a conjugate prior, can be readily updated with observed data. Given a set of parameters $\alpha_1 \dots \alpha_m$, where m is the number of values that a discrete variable may take, the likelihood of the discrete probability distribution $p_1 \dots p_m$ under the Dirichlet distribution is given by

$$P(\underline{p}; \underline{\alpha}) = \eta(\underline{\alpha}) \prod_i^m p_i^{\alpha_i - 1} \delta(1 - \sum_i^m p_i),$$

where η , the normalizing constant, is the multinomial beta function. The expected values of the Dirichlet distribution are given by

$$E[p_i|\underline{\alpha}] = \frac{\alpha_i}{\sum_j^m \alpha_j}, \quad (5)$$

Intuitively, the α parameters correspond to counts of how often each value of the variable has been observed. Thus, an expert can initially specify an educated guess of the distribution by specifying the α values. These values are updated online each time the agent observes the variable take on a particular value. The counts grow monotonically with time, which is consistent with our notion that the system should get more confident about its model of the world as it gains experience.

Unfortunately, a problem occurs when applying this approach in a real-world setting. Recall that the specific task that the user wants the agent to perform is hidden; the agent only receives observations that provide clues as to what the task may be. However, the α parameters correspond to what observations are likely to occur in a specific state s . To update the observation probability of an observation o , we must therefore know what the state was when o was observed. We accomplish this task by constructing an HMM to represent the completed dialog and applying a modified form of the Expectation-Maximization algorithm. A description of an HMM is beyond the scope of this article, but the reader may refer to Rabiner [1989] for more details.

4. Passive Learning for Natural Language Interaction

We have introduced the POMDP dialog model that allows the agent to interact in noisy conditions and described how the agent may express uncertainty in the dialog model itself. As a baseline for more sophisticated algorithms, we now consider a scenario in which the dialog manager knows that its knowledge of the world is imperfect, although it cannot deliberately take any actions to improve its knowledge of the world. It can learn only through observation: for example, it can infer what are the most popular user requests. However, it cannot actively determine the user's frustration due to a poor action or repeated questions. In the traditional reinforcement learning framework employed in this section, the dialog manager gains information about the user's preferences by asking the user to provide it with a reward ("reinforcement") after each interaction.

Recall that the standard dialog POMDP solution maximizes expected reward given the uncertainty over the user's request (the state); we now desire a solution that maximizes the expected reward given uncertainty in both the user's request and the user model. We rewrite equation 4, which describes the value of performing action a given a belief b over the user's request, to include the uncertainty over the model parameters M :

$$\begin{aligned} q_{a,i}(s) &= E_M[R(s, a) + \gamma \sum_{o \in O} \sum_{s' \in S} T(s'|s, a) \Omega(o|s', a) V_{n-1,i}(s)] \\ &= E_M[R(s, a)] + \gamma \sum_{o \in O} \sum_{s' \in S} E_M[T(s'|s, a) \Omega(o|s', a) V_{n-1,i}(s)] \\ &= E_M[R(s, a)] + \gamma \sum_{o \in O} \sum_{s' \in S} E_M[T(s'|s, a)] E_M[\Omega(o|s', a)] V_{n-1,i}(s), \end{aligned} \quad (6)$$

where $E_M[R(s, a)]$, $E_M[T(s'|s, a)]$ and $E_M[\Omega(o|s', a)]$ are the means of the Dirichlet

Table 1. Passive dialog model learning.

PASSIVE DIALOG MODEL LEARNING

- | |
|---|
| <ol style="list-style-type: none"> (1) Initialize a prior distribution over dialog models and determine the policy of the mean model by solving equation 6. (2) Interact with the user using the policy in (1) until the dialog is complete, asking for numerical feedback at each step. (3) Use the utterances from the user to update the prior over the transition and observation models. (4) Use the numerical reinforcement signal from the user to update the prior over the reward model. (5) Refine the policy based on the new mean model. |
|---|

distributions. The second line follows from the linearity of expectations, while the third line follows because the transitions T , the likelihood that the users will change their intended request, are independent of the noisy dialog observations Ω . The function $V_{n-1,i}$ is a fixed value from the previous iteration and does not require averaging. Thus, to find the optimal dialog policy given an uncertain user model, it is sufficient to solve the POMDP with the expected values of the model parameters.

Given a history of rewards, observations, and actions, the agent can update its belief over models as described in section 3. We do not expect individual parameters to change greatly during an interaction, so the dialog manager refines its current value function via additional backups instead of solving an entirely new policy (admissible because equation 4 is a contraction Gordon (1995)). Table 1 summarizes the passive learning algorithm.

A key question is how much refinement is required after each model update—that is, how many backup operations should be performed—after updating the user model. Performing enough backups so that value function fully converges to a new value may require significant time. Especially early in the learning process, when the agent has very little knowledge of the true dialog model, such effort may not be useful because the updated average model may still be far from the true model. We compare three approaches to determine the how much to replan after each successful dialog:

- (1) Replan to completion. After each completed dialog, we perform enough backups for the new value function to converge. (For simulation speed, we capped the number of backups to 50 per iteration.) This should lead to the best expected performance given the uncertainty in the user model. However, computing a POMDP policy to convergence can be also a slow process, leading to latencies in dialog manager responses.
- (2) Replan k times. Perform k backups, regardless of how much the updated model estimate differs from the previous estimate. This approach may prevent latencies in the dialog manager performance, but does not give the dialog manager time to compute better plans once it has confidence in the model.
- (3) Replan proportionally to model confidence. The sum of the variances on each parameter is a rough measure of the overall uncertainty in the dialog model. If an update reduces the overall variance by d_v , the agent backs up $\min(50, \lfloor k * d_v \rfloor)$ times. The intuition is that the dialog manager should expend the most computational effort when it has the most reliable knowledge.

4.1. Results

We used simulation to test the learning approach against a known “ground truth” optimal policy. In our test scenario, the dialog manager initially believed that its observations were more accurate than they actually were and that the consequences of going to the wrong place were not particularly severe. There were seven total states: a start state, an end state, and five goal states. The agent had the option of confirming each goal state, driving to each goal state, asking a general query, or doing nothing. (Since the focus was on dialog management, no navigation was involved in the simulation; we assumed that once the agent determined the correct goal location, it would be able to navigate to it successfully.) Each goal state had an observation associated with it, and the model included additional observations for the keywords ‘yes,’ ‘no,’ and ‘done.’ Table 2 contains the initial guesses and true values for the transition, observation, and reward parameters. The simulator used the ground truth model to produce observations and rewards given the underlying state (hidden from the agent) and the agent’s most recent action.

Figure 4 plots performance and planning time over a period 60 user iterations or dialogs. In all approaches, the performance improves as the agent learns the dialog model. However, replanning proportionally to the change in model confidence achieved that result with almost as little computation as replanning once per update (see Table 3; note that updates occur *after* a dialog is completed, while the robot is performing the requested task, and thus do not affect the flow of the conversation.) The model confidence approach focused replanning near the beginning of the trial when parameters first become certain. Each curve represents the average 100 trials of 60 user interactions.¹ We note that the $k_{d_v}=0.01$ approach completed fewer total backups than backing up once per iteration, but the larger number of backups near the beginning of the learning process helped stabilize the solution.²

Table 2. Initial and True Parameter Values for Simulation Tests. The initial model parameters assume a higher speech recognition accuracy and a user that is unforgiving of confirmation actions. This is an example of a very demanding user relative to the initial model specification.

	Initial	True
P(self-transition)	.95	.95
P(correct obs if ask-which)	0.7	0.5
P(correct obs if confirm)	0.9	0.7
R(complete task)	100	100
R(ask-which)	-1	-10
R(correct confirm)	-1	-1
R(incorrect confirm)	-10	-2
R(incorrect destination)	-50	-500

4.2. Discussion

The benefit of the simple approach presented in this section is that it requires very little computational effort; the backups are fast enough that the dialog manager can learn and interact with the user in real time. However, it suffers from several

¹Except for ‘convergence’—in the interests of time, we completed only 37 trials).

²These tests used an oracle that provided the dialog manager with the user’s true request to estimate the user state for model updating after the dialog was completed. We present these results because they show a clearer picture of the algorithm’s performance. Tests using HMM-based state estimation for model updating performed slightly but not significantly worse.

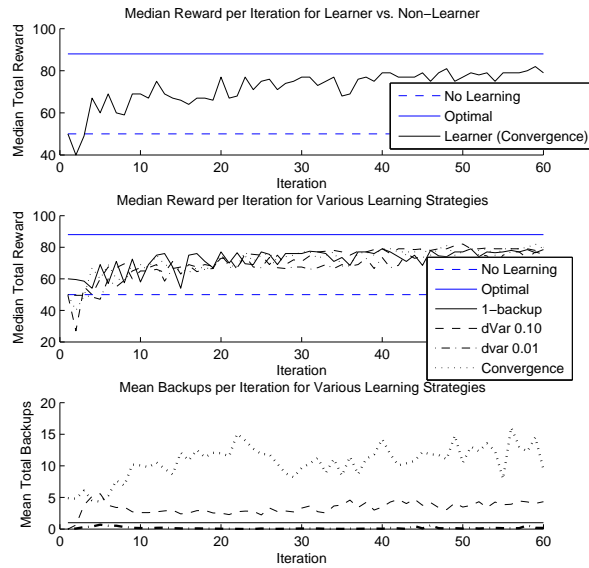


Figure 4. Performance and computation graphs. The learner outperforms the non-learner (top graph), and all of the replanning approaches have roughly the same increase in performance (middle graph), but replanning proportionally to the confidence of the model achieves that performance with much less computation (and therefore faster response) than replanning to convergence (bottom graph, same legend as the middle graph).

Table 3. Mean update times for each of the four approaches. Note that updates take place only after a dialog has been completed; when the dialog policy does not require any updates, the dialog manager’s average response time is 0.019 seconds.

Approach	Time (sec)
Convergence	135.80
1-backup	13.91
0.10-var	42.55
0.01-var	18.99

drawbacks. First, the user is forced to provide reinforcement after every interaction within a dialog. The agent must also make mistakes to learn about negative reinforcement. Furthermore, by only considering a point estimate of the dialog model when choosing actions, the agent is not fully aware of the limitations in its understanding of the dialog model. In the worst case scenario, these limitations can lead it to choose actions that will never allow it to learn the true dialog model. For example, consider a scenario in which there are only two actions: *ask* and *confirm*. Suppose that under some prior belief M over reward parameters, we have the following relationship between the true rewards and their expected values under the prior:

$$R_{ask} > R_{conf} = E_M[R_{conf}] > E_M[R_{ask}], \tag{7}$$

where R_{ask} is the reward for asking a general query and R_{conf} is the reward for asking a confirmation question. If the dialog manager attempts action *ask*, it will discover that its belief about R_{ask} is incorrect. However, if the dialog manager only makes decisions based on the rewards it expects to receive, $E_M[R_{conf}]$ and $E_M[R_{ask}]$, it will never try the action *ask*. Thus, the dialog manager will never discover the mistake in its understanding. This situation will occur if the domain expert estimates the reward means incorrectly, even if the expert states that he is very unsure about some of the values he chose. We remedy these issues in the next section.

5. Active Learning for Natural Language Interaction

In this section, we introduce a more sophisticated algorithm which, by using meta-queries and explicitly considering uncertainty of its model, behaves more robustly and can correct poorly initialized parameter values. The full state space of this planning problem is continuous and high-dimensional; to solve the problem tractably our agent also assumes that while it may use a series of actions to satisfy a user’s request, only its next action can potentially improve its knowledge of the dialog model. This approximation simplifies planning under model uncertainty because the agent only needs to think one step ahead. (Of course, the next time it needs to select an action, it will also think that it has one more chance to learn about the true dialog model.)

In figure 5, we show typical simulation results for the model-uncertainty POMDP for a variety of priors. In each case, the overall performance is about the same, showing (as we expect) that the parameter POMDP approach is robust to poor initialization. What is most interesting, however, is that if we start out with a conservative prior, that is, a prior that puts most of its weight on a “tough” set of rewards, the initial policy performance is relatively close to the optimal policy during the learning process. In this case, the system is robust because it is very cautious not to do the wrong thing, and the actions that allow it to actively learn about the model are relatively low-cost.




Figure 5. The three different priors—a uniform prior, a correct prior, and a conservative prior—converge to the same steady-state performance, but choosing a conservative prior allows us to do as well as if we had known the user’s preference state initially. All solutions were backed up to convergence.

We use meta-queries to avoid the need for tedious numerical feedback. Just as we augmented the state space of our original dialog POMDP with the model parameters to create the model-uncertainty POMDP, we now augment the the action space to allow for an alternative form of feedback: the meta-query. The meta-query asks the user to help the system by suggesting what action it should take next. For example, suppose that the dialog manager on our wheelchair is fairly certain that the user wishes to go to the printer. Then it might ask:

“I think you definitely want me to go to the printer. Should I go to the printer in this situation?”

On the contrary, if the wheelchair thinks that the user may want to go to the printer but is not very certain, it might ask:

“I think it is possible that you may want me to go to the printer, but I’m not really




Figure 6. Boxplot of dialog manager performance with a discrete set of four possible models. In this case, the user is very intolerant to errors, but the learner does not initially know this. Although the medians of the two policies are not so different, the active learner (right) makes fewer mistakes than the passive learner (left), leading to overall much less user-annoying behavior.

sure. Should I go to the printer in this situation?”

The initial phrase gives the user an intuitive sense of the dialog manager’s uncertainty. Thus, users can advise the robotic wheelchair based on their internal preferences. For example, if a user is risk averse, they may respond “yes” to the first question—since the system is sure of the user’s goal. In contrast, they may respond “no” to the second question—if system is uncertain, perhaps it should really confirm first with the user. If the user answers a question to the negative, the wheelchair might follow up with further questions such as,

“In that case, I think I should confirm that you want to go to printer first. Is that correct?”

until it receives an affirmative response.¹

From a user’s perspective, meta-queries simplify model-learning because only desired behaviors for the agent need to be specified. The meta-query approach is also active in that the dialog manager chooses when it needs help instead of asking the user for feedback after each interaction. However, meta-queries place an additional computational burden on the agent because it must now consider what parameter values for the observation, transition, and reward models are consistent with the policy information provided by the user.

5.1. Solving the Model-Uncertainty POMDP directly

In general, the transition, observation, and reward parameters are continuous-valued, with an infinite number of possible models, leading to computational intractability for conventional solvers. In special situations, however, uncertainty in the dialog model may be expressed as a small, discrete set of possible models rather than a continuous distribution, making the model-uncertainty POMDP much easier to solve.

For example, consider a scenario where we already have accurate transition and observation models, but a new user’s preference model is unknown. In an extreme case, we may decide to only characterize the user’s frustration with an incorrect movement as *low* or *high*, and similarly characterize the user’s frustration with an incorrect confirmation as *low* or *high*, leading to four possible dialog models.

¹In our tests, we used an abbreviated form of the meta-queries for simulation speed.

Table 4. Dialog model learning approach using Bayes risk and meta-queries.

DIALOG MODEL LEARNING WITH BAYES RISK
(1) Sample POMDPs from a prior distribution over dialog models.
(2) Solve for the optimal policy for each sample.
(3) Interact with the user:
a) Use the dialog model samples to compute the action with (approximately) minimum Bayes risk.
b) If the risk is larger than a given ϵ , perform a meta-query.
c) Update each dialog model sample's belief based on the observation received from the user.
(4) Periodically resample from an updated prior over dialog models and return to step 2.

We show simulated results with this very simple scenario in Figure 6. The figure compares the performance of the policy without using meta-queries (left column) to the performance of the policy *with* meta-queries. As expected, the system which has the ability to ask meta-queries can use the questions to gain information about the user's internal preference model. The reduction in large negative mistakes is substantial—which is particularly important in dialog management, where users will likely find a system that regularly makes mistakes annoying. Unfortunately, our approach of representing the user model as discrete values (such as *Wrong-QuestionPenalty = low*) does not scale and is difficult to apply to learning other parts of the model. We therefore turn to approximation techniques which will allow us to represent a larger class of models with continuous parameters.

5.2. Approximately Solving for a Dialog Policy

We separate the problem of finding the dialog policy into steps. In the first step, the dialog manager establishes a representative set of dialog models it believes are consistent with the user interactions it has experienced. Next, it solves for the optimal policy in each of these models. Thus, after any user utterance, it has a collection of actions—votes from each of the candidate models—to take next. These votes can be used to compute the risk of taking each action; if the safest action is too risky, the dialog manager will ask a meta-query. By providing information about the best action to take, the meta-query allows the dialog manager to prune away inconsistent models from its representative set. Table 4 outlines this approach.

Minimum Risk Action Selection

Given a dialog model m , the dialog manager can compute the loss of taking a particular action a compared to the best action a_{opt} . The loss is simply the difference $Q(b, a) - Q(b, a_{opt})$, where the Q -function denotes the value of taking action a in belief b (equation 2). Given a belief over models, the Bayes risk of an action is defined to be the loss the dialog manager expects to incur on average by choosing action a :

$$BR(a) = \int_M (Q_m(b_m, a) - Q_m(b_m, a_{opt,m}))p_M(m), \quad (8)$$

where M is the space of all possible dialog models, b_m is the current belief over possible user requests according to dialog model m , and $a_{opt,m}$ is the optimal action for the current belief b_m according to dialog model m . Let $a^* = \arg \max_{a \in A} BR(a)$ be the action with the least risk. If the risk $BR(a^*)$ is less than fixed cost of a

meta-query, that is, if the least expected loss is still more than a certain threshold, the dialog manager will ask the meta-query. Intuitively, equation 8 computes the potential loss due to taking action a instead of the optimal action a_{opt} according to each possible dialog model m and weights that loss by the probability of model m . When the agent is sufficiently sure about the model, the risk will be low. When it is unsure about the model, the risk may be high but the series of meta-queries will lead it to choose the correct action.

We cannot compute the Bayes risk in closed form. Thus, we represent the probability over models $p(m)$ by a set of sampled models and apply the following approximation:

$$BR(a) = \sum_i (Q_i(b_i, a) - Q_i(b_i, a_{opt,i}))w_i, \quad (9)$$

where Q_i provides the value of taking actions from belief according to dialog sample i . If model samples are drawn from $p(m)$, the weight w_i of each model is simply $\frac{1}{N}$, where N is the number of samples. As the dialog manager's belief over models changes, the w_i values are adjusted to be proportional to the ratio of the previous likelihood of the sample and its likelihood given the new information. While it is possible to bound the risk to a specified degree of accuracy, these bounds are loose and in practice we found that fifteen samples sufficed for our dialog management application.

POMDP Resampling

After several human-robot interactions, the system may discover that none of the dialog models it initially thought were representative are in fact reasonable. The need for resampling may also arise if one model's weight becomes very close to one and the remaining models have weights close to zero, causing the risk to appear smaller than it is. We have two sources of information to update our sample set of dialog models: the history h of action-observation pairs from the most recent dialog and the set Q of meta-query-response pairs. Given h and Q , the posterior probability $p_{M|h,Q}$ over models is:

$$p_{M|h,Q}(m|h, Q) = \eta p(Q|m)p(h|m)p_M(m), \quad (10)$$

where η is a normalizing constant.

If p_M is a Dirichlet distribution, then $\eta'p(h|m)p_M(m)$ is also a Dirichlet distribution since the likelihood $p(h|m)$ is product of multinomials and can be updated using the approach in 3. Incorporating meta-query information requires a different approach, since each specific meta-query response provides information about how the agent should behave, not the dialog model parameters. We do not have a closed-form expression for $p_{M|h,Q}$, so we use sampling instead. We first generate models from the updated Dirichlet priors. Next, we solve for the optimal policy of each model and check if each dialog model's policy is consistent with the previous meta-query responses stored in Q .

5.3. Results

We now present results for this approach from simulated trials which allow us to test against a ground truth model, and also a small set of evaluations on real users.

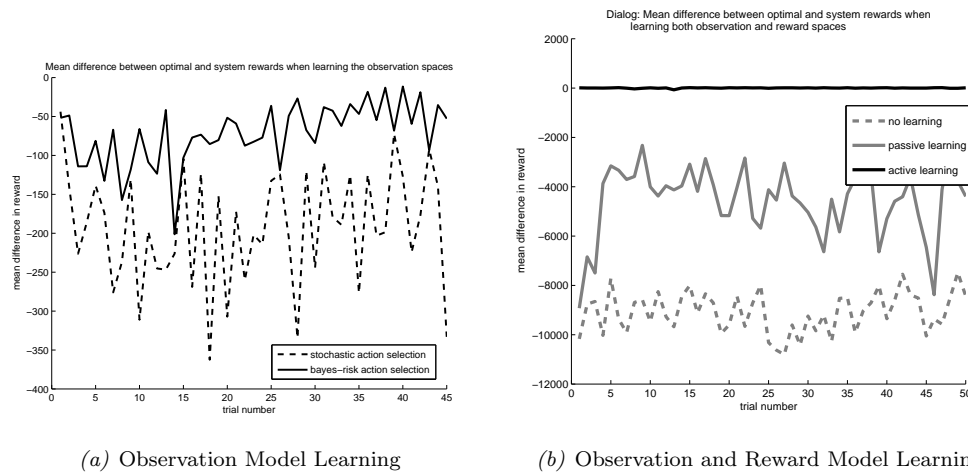


Figure 7. Dialog manager simulation results. (a) Results from learning only the observation model. (b) Benefits of active learning when learning both the observation and reward model.

Simulation

Figure 7 shows results from a simulated dialog manager for our robot wheelchair control application. Both figures show the difference between the learning policies and the ground truth optimal policy over time. In figure 7(a), we see the usefulness of the Bayes-risk approach (compared to stochastic actions selection based on the weights of the sampled models) when the reward model is known. In this case, the Bayes risk action selection allows us to choose non-risky actions. The usefulness of our approach is even more dramatic when the reward prior is uninformative (figure 7b). In this case, the dialog manager can improve somewhat by passively updating its priors based on what it has heard (solid gray line). However, simply listening cannot provide the dialog manager with information about the user’s preferences. Moreover, since the active learning system asks the user for help whenever it is confused, this system does not suffer from poor performance while the model estimate converges. The meta-queries allow the active learner to learn while avoiding mistakes, thus maintaining a high level of performance.

Robotic Wheelchair

We also validated our approach on a dialog manager for a robotic wheelchair with simple user evaluations. The underlying POMDP, with 10 states, 38 observations, and 21 actions, used keywords from a speech recognition system output as observations. Initially, each state had one strongly mapped observation (such as ‘printer’ for the printer location). The remaining observations received uniform initial priors. Four users conducted 12-15 interactions (20-25 minutes) with the system.

By asking meta-queries, the dialog manager was able to successfully complete all 57 interactions without ever trying to drive the user to an incorrect location. Table 5 shows that the proportion of dialogs with meta-queries decreased significantly from the user’s initial interactions to his or her final interactions with the system. The proportions are relatively high because the users, being new with the system, asked for a variety of locations or experimented with different vocabulary. In Table 6, we see that the system often asked a meta-query the first or second time a user asked to go to a particular location, but there was a significant drop in the number of meta-queries after the third time the user requested to go to the same place. We note that our active-learning approach allowed our dialog manager to automatically determine the length of its training phase. Our trials were too small to draw strong

Table 5. Proportion of dialogs with meta-queries by location. The decrease in the number of meta-queries from the second to the third time the location was asked for is statistically significant at the $p = 0.05$ level.

First half of interactions	.79
Second half of interactions	.48

Table 6. Proportion of dialogs with meta-queries. The decrease in the proportion of queries is significant at the $p = 0.05$ level.

How often place was requested	First time	Second time	Third Time
Total number of requests	29	15	8
Number of requests with meta-queries	22	11	2
Proportion of requests with meta-queries	.76	.73	.25

conclusions from these results (especially since we were unable to account for the effect of the novelty factor in the interactions), but our quantitative analysis does show statistically significant decreases in the number of meta-queries asked as the dialog manager adapted to each user.

Next we present some sample dialogs to illustrate our system’s learning. Table 7 contains a sample dialog where the system used a meta-query to establish the meaning of a new word. In the initial dialog, the word ‘forecast’ is new to the system, that is, it has no prior on what that word may mean. The system determines that the Bayes risk is great enough to warrant asking for a policy clarification. As a result, future dialogs are biased toward choosing POMDP samples that satisfy the initial policy query request, which allows later dialogs to avoid the same kind of training.

Table 7. Learning a new word: The system learns the meaning of a new word through a meta-query.

<p>Early Conversation ROBOT: How can I help you? USER: Give me the forecast. ROBOT: I’m confused. What action should I take now? <i>The system pops up a meta-query window, and user indicates action “show weather” is correct.</i> ROBOT: Showers.</p>
<p>Later Conversation (with the same user) ROBOT: How can I help you? USER: What’s the forecast for today? ROBOT: Do you want the weather? <i>The system has partially the word “forecast” and confirms the correct state.</i> USER: Yup. ROBOT: Showers.</p>

In Table 8, the system adapts to a situation where two words, ‘coffee’ and ‘copy’ are often mistaken for each other in the speech recognition system. We stress that at the beginning of the dialog, the system was initialized to believe that the word ‘coffee’ was most likely if the user wished to go to the coffee machine, while the word ‘copy’ was most likely if the user wished to go to the copy machine. The other key words—“tea” and “kitchen”—in this segment began with uniform priors. During the first conversation, the dialog manager is confused by the many new words. Even after asking a meta-query, it is confused by hearing the word “copy” since the user seems to want to go to the kitchen. After several interactions, the dialog manager

goes to the correct location. We note that even in this difficult scenario, the use of the meta-queries helped the dialog manager focus on to the kitchen when the only word it recognized—“copy”—seemed to indicate that the user wished to go to another location. In the later dialog, we see the effect of the learning: upon hearing the words “kitchen” and “copy,” the dialog manager correctly infers that the user wishes to drive to the kitchen.

Table 8. Learning about observation noise: the system learns new words for a noisy pair of states.

<p>Early Conversation ROBOT: How can I help you? USER: I want some tea from the kitchen. ROBOT: How can I help you? USER: The tea is in the kitchen near the coffee machine. <i>The system heard “copy” instead of “coffee” and is confused about the new word “kitchen.”</i> ROBOT: How can I help you? USER: Go to the kitchen. ROBOT: I’m confused. What action should I take now? <i>The user indicates “Confirm the coffee machine area” is the correct action.</i> ROBOT: Do you want to go to the coffee machine area? USER: Yup. <i>... The system asks a few more (non-meta) questions before re-confirming the coffee machine state.</i> ROBOT: Do you want to go to the coffee machine area? USER: Yup. ROBOT: Going to the coffee machine area.</p>
<p>Later Conversation (with same user) ROBOT: How can I help you? USER: Take me to the kitchen. <i>The system hears a noisy combination of “nick” and “kitchen” (mostly “kitchen”).</i> ROBOT: How can I help you? USER: I want some coffee. <i>The system hears “copy” instead of “coffee,” but with its learned observation model, it confirms correctly.</i> ROBOT: Do you want to go to the coffee machine area? USER: Yup. ROBOT: Going to the coffee machine area.</p>

5.4. Related Work

The approach presented in this paper builds upon prior work in POMDP learning that has also considered incorporating the parameters of the POMDP as additional hidden states into a larger model-uncertainty POMDP. Closest to our work are the Bayesian approaches in the Medusa Jaulmes et al. (2005) and Beetle Poupart et al. (2006) algorithms. Applying observations from Dearden et. al. (Dearden et al. (1999)), the Medusa algorithm uses a candidate set of samples to represent a distribution over models. As in our approach, the agent updates the weight it places on each sample as new information arrives and changes its knowledge of the world. However, the Medusa algorithm chooses actions stochastically based on the weights of the samples. By choosing actions to minimize Bayes risk, our dialog manager can act robustly even with high uncertainty in the dialog model.

POMDP-based dialog systems are faced with two difficulties. The first is that simply solving a POMDP can be quite computationally expensive. Fortunately, much research has focused on finding tractable dialog manager policies. For example, Williams and Young (2005) derive methods to incorporate an explicit confidence output from the speech recognition system as an additional measurement,

allowing the agent to reason about the possibility of speech recognition errors. In situations where only certain actions are relevant to certain states, the POMDP can be factored into a hierarchy that reduces the overall amount of computation required Pineau et al. (2001). Certain dialogs also contain symmetries that lend themselves to efficient approximations Williams and Young (2005).

The Beetle algorithm takes a decision-theoretic approach to solving MDPs with uncertain parameters. Much like our model-uncertainty POMDP, Poupart et. al. treat the unknown MDP parameters as hidden state in a larger POMDP and derive an analytic solution (based on Porta et al. (2006)) for a policy that will trade optimally between learning the MDP and maximizing reward. Unfortunately, these techniques do not extend tractably to the model-uncertainty POMDP, which is continuous in both the POMDP parameters (like the MDP) and the belief state (unlike the MDP). While approximate, our Bayes risk action selection criterion allows the dialog manager to function in this complex space of dialog models.

We also note that there are also non-Bayesian approaches for acting in uncertain environments. Many interaction systems provide the dialog manager with a set of rules to follow given particular outputs from a speech recognition system. For example, the Mercury system Seneff and Polifroni (2000) builds a network in which a flight reservation system keeps track of what information has already been provided and for what the user needs to be prompted. These rules can also help the agent monitor the quality of the interaction Paek and Horvitz (2004) and adapt to new users Litman and Pan (2002).

The drawback to rule-based systems is that they often have difficulty managing the many uncertainties that stem from noisy speech recognition or linguistic ambiguities. The ability to manage the trade-off between gathering additional information and servicing a user's request have made (PO)MDP planners particularly useful in dialog management; applications include a nursebot robot, designed to interact with the elderly in nursing homes Roy et al. (2000), a vision-based system that aids Alzheimer's patients with basic tasks such as hand-washing Hoey et al. (2005), an automated telephone operator Williams and Young (2005), and a tourist information kiosk Litman et al. (2000).

Some implementations, targeted at industrial applications Nilim and Ghaoui (2004), Xu and Mannor (2007), aim to make the agent robust to (usually small) variations in its environment. While they provide guarantees for worst-case behavior, they do not describe how the agent should act if it has the ability to adapt to some of these variations through learning. Finally, a class of model-free approaches attempt to bypass the issue of model uncertainty by directly learning a policy Littman et al. (1995), Even-Dar et al. (2005), James et al. (2004)). The model-free approaches typically require a large amount of data to produce reasonable behaviors, and without a prior notion of possible pitfalls, they are likely to make many mistakes during the learning process. Our Bayesian approach avoids this problem by allowing the system designer to provide an indication of where potential pitfalls may be; by also letting the agent query the human for help, we ensure that the agent's policy is robust from the start.

6. Conclusion

Many approaches have been developed for human-robot interaction. In our work, we focused on robustly learning decision-theoretic models for effective human-robot interaction, in contrast to information-theoretic models such as Lopes and Teixeira (2000) and Lopes and Chauhan (2007) that try explicitly to learn new vocabulary. Our approach to decision-theoretic models of interaction relies on the partially

observable Markov decision process, in which decisions are made with respect to a distribution over user intents, with the goal of minimizing expected costs. Unfortunately, the POMDP approach typically assumes a reasonably accurate user model. In domains where large amounts of data are available—for example, automated telephone operators—the user model may be relatively easy to obtain. For human-robot interaction, however, collecting sufficient user data to learn a statistically accurate model is usually expensive: trials take a lot of time from human volunteers. We therefore developed an approach that allows the agent to leverage the benefits of POMDP-based dialog managers while addressing the difficulty of specifying the POMDP model parameters. An expert need only specify a prior belief describing his or her knowledge—and uncertainty—regarding the dialog model, and the agent refines this knowledge through interactions with the user.

To make the learning process robust, we introduced a risk-averse action selection criterion that allows our dialog manager to behave robustly even when its knowledge of the true dialog model is uncertain. Instead of asking for feedback after every interaction, as many traditional reinforcement learning approaches require, our dialog manager uses an active learning scheme to only ask for help when it feels it does not have enough information to make a safe decision. Framed as meta-queries, or questions about actions the dialog manager is thinking of taking, our active learning scheme is also more intuitive than asking for numerical feedback. We demonstrated our approach both in simulation and on a real dialog manager for a robotic wheelchair.

While our active learning approach addresses many of the limitations in previous research and resolves the drawbacks of the passive learning approach, there are two key areas for improvement and future work. First, the most computationally expensive part of this approach is finding dialog models that are consistent with the meta-query information that the user has provided; better sampling strategies will be required to allow the algorithm to scale to more complex dialogs. Second, we are still seeking effective ways for the user to aid the dialog management system when it is confused; we believe that meta-queries are only the first step toward more natural (and accurate) human-robot feedback. Resolution of these issues will lead to even more robust and effective human-robot interactions.

References

- Reeves, B., and Nass, C., *The Media Equation: How People Treat Computers, Television and New Media Like Real People and Places*, Cambridge Press (1996).
- Roy, N., Pineau, J., and Thrun, S. (2000), “Spoken Dialogue Management Using Probabilistic Reasoning,” in *Proceedings of the 38th Annual Meeting of the ACL*, Hong Kong.
- Hoey, J., Poupart, P., Boutilier, C., and Mihailidis, A. (2005), “POMDP models for assistive technology,” *IATSL Technical Report*.
- Fern, A., Natarajan, S., Judah, K., and Tedepalli, P. (2007), “A Decision-Theoretic Model of Assistance,” *IJCAI*.
- Bellman, R., *Dynamic Programming*, NJ: Princeton University Press (1957).
- Gordon, G.J. (1995), “Stable function approximation in dynamic programming,” in *Proceedings of the Twelfth International Conference on Machine Learning*, San Francisco, CA: Morgan Kaufmann.
- Pineau, J., Gordon, G., and Thrun, S. (2003), “Point-based value iteration: An anytime algorithm for POMDPs,” *IJCAI*.
- Jaulmes, R., Pineau, J., and Precup, D. (2005), “Learning in non-stationary Partially Observable Markov Decision Processes,” in , ECML Workshop.
- Poupart, P., Vlassis, N., Hoey, J., and Regan, K. (2006), “An analytic solution to discrete Bayesian reinforcement learning,” in *ICML*, Pittsburgh, Pennsylvania, New York, NY, USA: ACM Press, pp. 697–704.
- Dearden, R., Friedman, N., and Andre, D. (1999), “Model based Bayesian Exploration,” in , pp. 150–159.
- Williams, J., and Young, S. (2005), “Scaling Up POMDPs for Dialogue Management: The ”Summary POMDP” Method,” in *Proceedings of the IEEE ASRU Workshop*.
- Pineau, J., Roy, N., and Thrun, S. (2001), “A Hierarchical Approach to POMDP Planning and Execution,” in *Workshop on Hierarchy and Memory in Reinforcement Learning (ICML)*, June.

- Porta, J.M., Vlassis, N., Spaan, M., and Poupart, P. (2006), "An Point-Based Value Iteration for Continuous POMDP," *Journal of Machine Learning Research*.
- Seneff, S., and Polifroni, J. (2000), "Dialogue management in the Mercury flight reservation system," in *ANLP/NAACL 2000 Workshop on Conversational systems*, Seattle, Washington, Morristown, NJ, USA: Association for Computational Linguistics, pp. 11–16.
- Paek, T., and Horvitz, E. (2004), "Optimizing Automated Call Routing by Integrating Spoken Dialog Models with Queuing Models.," in *HLT-NAACL*, pp. 41–48.
- Litman, D.J., and Pan, S. (2002), "Designing and Evaluating an Adaptive Spoken Dialogue System," *User Modeling and User-Adapted Interaction*, 12(2-3), 111–137.
- Litman, D., Singh, S., Kearns, M., and Walker, M. (2000), "NJFun: a reinforcement learning spoken dialogue system," in *Proceedings of the ANLP/NAACL 2000 Workshop on Conversational Systems*, Seattle.
- Nilim, A., and Ghaoui, L. (2004), "Robustness in markov decision problems with uncertain transition matrices," *NIPS*.
- Xu, H., and Mannor, S. (2007), "The Robustness-Performance Tradeoff in Markov Decision Processes," in *Advances in Neural Information Processing Systems 19* eds. B. Schölkopf, J. Platt and T. Hoffman, Cambridge, MA: MIT Press.
- Littman, M.L., Cassandra, A.R., and Kaelbling, L.P. (1995), "Learning policies for partially observable environments: Scaling up," in *Proceedings of the Twelfth International Conference on Machine Learning*, eds. A. Prieditis and S. Russell, San Francisco, CA, USA: Morgan Kaufmann publishers Inc.: San Mateo, CA, USA, pp. 362–370.
- Even-Dar, E., Kakade, S.M., and Mansour, Y. (2005), "Reinforcement Learning in POMDPs Without Resets.," in *IJCAI*, pp. 690–695.
- James, M.R., Singh, S., and Littman, M. (2004), "Planning with Predictive State Representations," in .
- Lopes, L.S., and Teixeira, A. (2000), "Human-Robot Interaction through Spoken Language Dialogue," *Proceedings of the International Conference on Intelligent Robots and Systems*.
- Lopes, L.S., and Chauhan, A. (2007), "How many Words can my Robot learn? An Approach and Experiments with One-Class Learning," *Interaction Studies*, 8(1).