# Modeling Brief Alcohol Intervention Dialogue with MDPs for Delivery by ECAs

Ugan Yasavur, Christine Lisetti, and Napthali Rishe

School of Computing and Information Sciences
Florida International University
Miami, FL, 33199, USA
ugan.yasavur@fiu.edu, {lisetti,rishen}@cs.fiu.edu

**Abstract.** This paper describes the design of a multimodal spoken dialogue system using Markov Decision Processes (MDPs) to enable embodied conversational virtual health coach agents to deliver brief interventions for lifestyle behavior change - in particular excessive alcohol consumption. Its contribution is two fold. First, it is the first attempt to-date to study stochastic dialogue policy optimization techniques in the health dialogue domain. Second, it provides a model for longer branching dialogues (in terms of number of dialogue turns and number of slots) than the usual slot filling dialogue interactions currently available (e.g. tourist information domain). In addition, the model forms the basis for the generation of a richly annotated dialogue corpus, which is essential for applying optimization methods based on reinforcement learning.

**Keywords:** spoken dialogue system, markov decision Processes, reinforcement learning, embodied conversational agent (ECA), intelligent virtual agents, brief intervention, behavior change, alcoholism, at-risk drinking.

## 1 Introduction

Excessive alcohol use, with approximately 85,000 of directly or indirectly attributable deaths per year, is the 3rd leading lifestyle-related cause of death in the United States [1]. In 2006, there were more than 1.2 million emergency room visits and 2.7 million physician office visits due to excessive drinking [2]. The economic costs of excessive alcohol consumption in 2006 were approximately $223.5 billion [2].

*Brief interventions (BI)* are short, well structured, one-on-one counseling sessions, focused on specific aspects of problematic lifestyle behavior, and are ideally suited for people who drink in ways that are harmful or abusive. BIs can be delivered in 3-5 minutes [3] and (for alcohol consumption as a target) aim to moderate a person's alcohol consumption to reasonable levels and to eliminate harmful drinking behaviors. BIs are the top ranked out of 87 treatment styles in terms of efficiency [4]. It is reported that even a few minutes of advice and discussion about behavioral problems can be as effective as more extended

R. Aylett et al. (Eds.): IVA 2013, LNAI 8108, pp. 92–105, 2013.

counseling [5]. Many challenges are involved in delivering BIs to people in need, such as finding the time to administer them in busy doctors' offices, obtaining the extra training that helps staff become comfortable providing these interventions, and managing the cost of delivering the interventions [6]. Patients are often encouraged to use computer programs developed based on BI content in the doctor's waiting room or at home, or to access the intervention through the Internet, which not only offers privacy but also the ability to complete the program anywhere, any time of the day [7–9]. Although computer-based interventions adapted from one-on-one brief interventions are reported to have positive effect on reducing patients' drinking level [7, 8, 10], these programs interact with patients with menu-based text-only user interfaces [10, 11], and are less attractive to some users than one-on-one interventions. We posit that these challenges on the adoption of BIs delivered by computers can be overcome by delivering these interventions with spoken dialogue systems (SDS) integrated with multimodal interfaces and embodied conversational agents (ECAs) [12].

ECAs are animated anthropomorphic characters which is an emerging technology in multi-modal interfaces [13] that have become increasingly interesting user interfaces for a wide range of applications, such as tutoring systems [14], health behavior change systems [15, 16], and health applications [17]. ECAs can provide users with a natural anthropomorphic interface which can deliver verbal and nonverbal modalities similar to those found in face-to-face human interaction (e.g., facial expressions, hand and body gestures). The presence of non-verbal communication is shown to have different types of positive effects such as greater feelings of rapport [18] and greater feelings of trustworthiness [19] about the agent. In our current system, we use an ECA system (discussed in section 3) which can convey basic non-verbal behaviors with facial animation and lip-synchronization. However, because we currently focus on verbal communication performed by the spoken dialogue system, we do not exercise the option of controlling its non-verbal behaviors such as facial animations (i.e. its default animation engine generates facial expressions with lip-synchronization) and body gestures.

We have concentrated on the specific brief intervention which is prepared by National Institute on Alcohol Abuse and Alcoholism (NIAAA) [20] for alcohol screening and intervention. In this article, we survey related research on techniques used to-date to develop dialogue systems; we discuss the overview of our dialogue system and its integration with an Embodied Conversational Agent (ECA) and a multimodal interface; we describe our approach to modeling dialogue for brief interventions based on Markov Decision Processes (MDPs), our state-based unoptimized baseline system, and the nature of our annotated dialogue corpus that our system generates.

## 2 Related Research

Although, there exist no spoken dialogue systems (SDS) for the alcohol consumption domain, there has been growing interest to develop multimodal SDS which can converse, guide, assist or motivate users for different health related topics [21, 22, 17]. Dialogue management for health-related dialogue systems have so far been mostly designed based on *finite state* dialogue management mechanisms such as hierarchical transition networks [16, 21, 23]. These systems usually do not have speech recognition integration. Interaction is conducted based on menu-based choices but the system utterances is delivered vocally via text-to-speech or prerecorded voice.

*Plan-based* [24, 25] and *Information State Update (ISU)* based [26, 27] approaches are also employed in health-related dialogue systems. Dialogue management adapted from the existing plan-based TRIPS framework [24] has been used in the personal health assistance domain to help users with heart failure related problems [22]. SimCoach, designed to provide support and healthcare information about post-traumatic stress disorder, incorporates traditional information-state approach [26] with dialogue moves with assigned reward values [17]. While plan-based [24, 25] and ISU-based [26] approaches have been shown to provide a basis for flexible dialogue interaction, these approaches have a number of general limitations which stem from a design methodology based on the designer's intuition. These approaches require manual specification of update or inference rules which define an action for all possible dialogue situations. It is not practically possible, however, for the designer to anticipate all the possible situations of a dynamic dialogue environment. The main drawback of ISU-based approaches is that it is difficult for the dialogue designer to track the combined effect of sequentially applied updates to the information state. Since plan-based approaches highly depend on domain-dependent empirical design approach, system development can become opaque, and have high development and deployment costs.

In our system, we model brief interventions as *Markov Decision Processes (MDPs)*, which provide a stochastic data-driven framework for optimizing dialogue strategies. Optimization of dialogue strategies is usually performed by applying reinforcement learning algorithms [28]. Potential advantages of this approach in dialogue management are: **1)** data-driven development cycle, **2)** provably optimal dialogue actions, **3)** precise mathematical model for action selection, **4)** possibilities for generalization to unseen states, **5)** reduced development and deployment costs [29].

The *Reinforcement Learning (RL)* paradigm, in conjunction with fully observable and partially observable MDP dialogue models, are usually used in dialogue systems which use speech as communication medium [30–32], or which involve learning and optimization under noisy environments [31, 33]. Experiments showed that RL-based optimized approaches outperforms handcrafted dialogue management approaches [31]. So far, they have been mostly used in the tourist information domain, e.g. finding fun things to do in New Jersey [30], finding out about restaurants, hotels and bars [34], serving as a museum guide [32], with few exceptions such a system in the tutoring domain [35]. The main reason which limits the usage of RL-based dialogue management in different domains is the lack of *training dialogue corpus* for different domains. Most of the current work developed is based on annotated human-machine spoken dialogues corpora called

*Communicator* [36] which is used in user simulations to learn dialogue strategies [37, 38]. Versions of the *Communicator* corpora have been used by many researchers and have led to new technologies for speech and language processing. Therefore, annotated dialogue corpus is essential for performant RL-based systems.

Alternative to user simulation-based learning and using existing corpora is the *model-based learning* approach via collecting data from real user interactions [30]. In model-based approaches, the RL agent learns partial strategies from exploratory data generated by dialogues with real users. In model-based approaches a model represents the dynamics of the dialogue to compute an approximate value of taking each action in a particular state. With a model, the problem of learning a good dialogue strategy is reduced to computing the optimal policy for choosing a dialogue action in a dialogue state. We follow the model-based approach based on fully observable MDPs with some differences from previous systems [30, 31]. Our model of the problem is represented by interconnected separate MDPs with local sub-goals and global goals (details discussed in section 4.5). The system in the tourist information domains, the model of the problem is represented by a single MDPs and the optimization is performed based on a single global goal (e.g. task completion). Our approach helps to compute local optimal dialogue strategies.

The computation of optimal dialogue strategy can be achieved with standard RL algorithms [28]. This approach requires to build initial training system which can deliver basic but unoptimized functionality, and to specify performance criteria and estimates of dialogue states.

## 3 System Overview

Although this article is focused on the *Dialog Manager* of our system, we give a brief overview of the overall system in which it operates (see Fig. 1). Our multimodal spoken dialogue system has a multimodal ECA-based interface where: user's speech is recognized by the *Automatic Speech Recognition (ASR)* engine[1], user's facial information processing is performed by the *Facial Processing* third-party facial processing service[2]. We use two outputs of the facial processing service for annotating our training dialogue corpus: user's gender and smile, along with a confidence value. According to the gender attribute, a brief intervention SDS can adapt its behavior because there exist different thresholds for males and females (e.g. recommended drinking limits). The history of smiles, on the other hand, can give important information about the user's experience with the system (e.g. engagement and enjoyment levels can be inferred). These two outputs do not currently have any effect on our system's behavior; they are used for annotation to create an exploratory data set in the current version of the system. Nonverbal communication is also important in delivering health interventions but the focus of this paper is verbal aspect of the interaction.

[1] Currently, Microsoft Speech API (SAPI).

[2] Currently, Sky Biometry `http://www.skybiometry.com/`

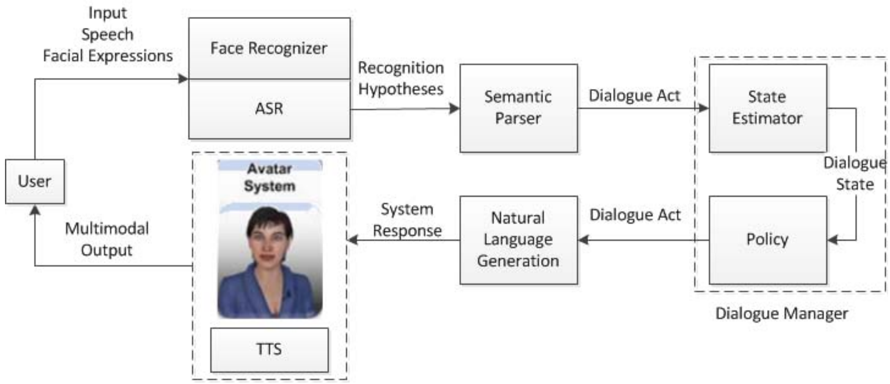


**Fig. 1.** Multimodal Dialogue System Architecture

*Automatic Speech Recognition (ASR)* hypotheses are parsed by the *Semantic Parser*[3] using context-free grammars and converted to dialogue acts. In addition to the parsing, we used a named-entity recognizer for behavioral health [40] to tag relevant entities such as alcoholic beverages. Semantic dialogue act output is passed to the *State Estimator*. The state estimator updates the *Dialogue State*, a random dialogue action is selected from the corresponding *Policy* table (state action mappings). The *Natural Language Generation* module uses a matching template for the dialogue act: if it is a question, it directly passes it to the *Text-To-Speech (TTS)* engine; if it is a confirmation, it fills the necessary parts in the template and passes it to the TTS engine. The TTS engine generates phonemes and the *Avatar System* automatically performs lip synchronization[4].

## 4 Approach

Compared to previous applications of RL-based dialogue systems, our brief intervention domain has several challenges: 1) the dialogue length and complexity; 2) the lack of a baseline system and of a dialogue corpus. The first challenge makes modeling harder and the dialogue size creates very large state-spaces which may cause data sparsity problems. The second challenge prevents us to optimize our dialogue policies and to evaluate optimized policies.

As we have discussed, we represent our problem with MDP framework which can be characterized by a tuple (S,A,T,R), where:

- S is a finite set of states
- A is a finite set of actions
- T is a state-transition function such that $T(s', a, s) = p(s' \mid s,a)$ which describes the probability of performing action **a** in state **s'** will lead to state **s**
- R(s, s') is a local reward function, and the objective of the SDS is to maximize the gained reward.

[3] Currently, Phoenix Parser [39].
[4] Currently, Haptek$^{TM}$.

### 4.1 Brief Intervention for Alcohol-related Health Problems

According to the clinician's guide for conducting brief interventions from the National Institute on Alcohol Abuse and Alcoholism (NIAAA) [20], a brief intervention can be delivered in three sequential steps:

- Step 1: Asking About Alcohol Use
- Step 2: Assessing for Alcohol Use Disorders
  - Assessment of Abuse
  - Assessment of Dependence
- Step 3: Advising and Assisting according to degree of alcohol problem
  - At-risk drinkers
  - Drinkers with alcohol use disorder

To develop our dialogue content, we follow the brief intervention guide for alcohol prepared by NIAAA [6]. The goal of our dialogue system is to deliver alcohol screening and brief interventions based on the guide. We explain the first steps in detail in the following three sections (step 3 in less details for lack of space).

Brief intervention dialogue for alcohol problems can be modeled as slot-filling dialogue. However, the number of slots are larger than the applications discussed in the tourist information domain. Moreover, the number of slots that are needed to be filled by the system is not fixed. Another aspect which differs in brief intervention dialogue is that the strategy that the system needs to follow is not always constant, and needs to adapt according to inputs that the system receives from the user: there can be *no* fixed dialogue plan. According to the NIAAA guide, we identified the number of slots we need by minimizing the complexity of dialogue: the number of the slots for at-risk alcohol users is 11, and it is 9 for users with alcohol use disorder. The dialogue needs to branch according to user inputs. For the users who do not have harmful drinking patterns, interaction may end earlier. Therefore, the dialogue strategy needs to be adapted according to user's pattern of drinking.

For each dialogue state, the system has two options for dialogue action selection based on *initiative type*. One type is *user initiative* dialogue action which are usually open-ended questions. The second type is the *system initiative* dialogue action which are closed questions. According to the initiative type, the system uses different grammar types for automatic speech recognition (ASR). If the question type is user initiative (open-ended question), the system uses non-restrictive grammar. If the question type is the system imitative, the system uses restrictive grammar which only recognizes particular entities mentioned in the question (e.g. number of alcoholic beverages consumed).

### 4.2 Step 1: Asking about Alcohol Use

The system starts interaction by greeting and asking permission to talk about user's drinking. After receiving consent of a user, it asks single question about alcohol use (e.g. "Do you sometimes drink beer, wine, or other alcoholic beverages?"). If the client's answer is no, there is no need to continue to screening. If the client's answer is yes, the system asks about the amount of alcohol the client consumes to find out if the client is an at-risk drinker (e.g. "How many times in the past year have you had 5 or more drinks in a day?").

If a client is not an at-risk drinker, the system may finalize interaction by advising to maintain or lower drinking limits according to the situation and offer re-screening annually. If a client is an at-risk drinker, to get the complete picture of drinking, the system asks two more questions to query the drinking pattern of the client. We have demonstrated the sample dialogue for Step 1 in Table 1. The example dialogue actions to query pattern of drinking are performed in S4 and S5 dialogue turns. Since the questions asked in S4 and S5 are open-ended, the type of dialogue action is user initiative.

**Table 1.** Sample Dialogued during Step 1: Asking About Alcohol Use

| | |
|---|---|
| S1: | Hi, I am [:::], Do you mind, if we talk about your drinking? |
| C1: | No, it is okay! |
| S2: | Do you sometimes drink, wine, or other alcoholic beverages? |
| C2: | Yes, I drink sometimes! |
| S3: | How many times in the past year have you had 5 or more drinks in a day? |
| C3: | I think at least once a week, I had around 5 drinks or more a day. |
| S4: | How frequently do you have an alcoholic beverage? |
| C4: | I think at least 3 days in a week. |
| S5: | On a typical drinking day, how many drinks you have? |
| C5: | I think 4 or 5 whiskeys. |
| S6: | Thanks for the information you have provided about your drinking. Next I will give you feedback about some important effects of your drinking. |

In Step 1, there can be maximum 4 slots if the user is an at-risk drinker (see Table 1). The system continues to Step 2 only if a user is an at-risk drinker. Since the dialogue is branching, we have represented each distinct step or sub-step with a separate MDP (see Fig. 2). We have elicited a state-space for each MDP separately which greatly reduced state-space. We represented dialogue states in Step 1 with 5 features: **1)** Greet (G) indicates whether or not the system greeted the user and asked for permission to talk about client's drinking; **2)** Question (Q) indicates which question is being queried in the current state; **3)** Confidence (C) indicates confidence level of the speech recognizer (low, medium and high are represented by 0,1,2 respectively; confidence values 3 and 4 stand for confirmed and non-confirmed, respectively); **4)** Value (V) indicates, is the value was obtained or not for the current question; **5)** Grammar(Gram) indicates the type of grammar (restrictive or non-restrictive) used by the ASR.

For example, dialogue state *11210* indicates that the system greeted the user (G=1), the first question is queried (Q=1), the ASR confidence level is high (C=2), the type of grammar is restrictive (Gram=0). The Confidence (C), Value

(V) and Grammar (Gram) features are also used in state representations in Step 2 and Step 3. Since the Greet (G) and Question (Q) are not relevant to represent the state of the dialogue, if the system is performing Step 2 or Step 3, it is possible not to use them in order to reduce state-space. We used the same approach to reduce state spaces in each of the separate MDPs. In each step we only used relevant state features. The compact state representation helps to avoid the data-sparsity problem with limited number of training dialogues.

**Table 2.** Generating the dialogue for Step 1 shown in Table 1

| States | | | | | Actions | Turn |
|---|---|---|---|---|---|---|
| **G** | **Q** | **C** | **V** | **Gram** | | |
| 0 | 0 | 0 | 0 | 0 | GreetS | S1 |
| 1 | 0 | 2 | 1 | 0 | NoConf | - |
| 1 | 1 | 0 | 0 | 0 | AskQ1S | S2 |
| 1 | 1 | 2 | 1 | 0 | NoConf | - |
| 1 | 2 | 0 | 0 | 0 | AskQ2S | S3 |
| 1 | 2 | 2 | 1 | 0 | NoConf | - |
| 1 | 3 | 0 | 0 | 0 | AskQ3U | S4 |
| 1 | 3 | 2 | 1 | 1 | NoConf | - |
| 1 | 4 | 0 | 0 | 0 | AskQ4U | S5 |
| 1 | 4 | 2 | 1 | 1 | NoConf | - |
| 1 | 5 | 0 | 0 | 0 | InformTrans1S | S6 |

As shown in Table 2, **dialogue actions** represent 2 types of actions for asking each question during the first time, according to the type of initiative (user or system). For each question, there are 2 types of actions to re-ask the question with user and system initiative types which are performed in the dialogue states when the system did not receive answer (i.e. Value feature of the state equals to 0 indicates that the answer was not obtained). There is also explicit confirmation action for each question to verify that the input was received, the system may also select not to confirm. If the system selects not to confirm action (i.e. showed as NoConf), it updates the dialogue state as input is received and continues with randomly selecting a dialogue action in the updated dialogue state.

A **dialogue policy** is a mapping between dialogue states and available dialogue actions in each state. In our training system, the dialogue action are randomly selected. This approach will create exploratory dialogue corpus for optimizing dialogue strategies with RL.

### 4.3 Step 2: Assessing for Alcohol Use Disorders

In Step 2, the system aims to determine whether or not there is a maladaptive pattern of alcohol use that is causing clinically significant impairment or distress. In this step, the system queries with 4 questions whether a client has alcohol abuse (e.g. risk of bodily harm, relationship trouble) and alcohol dependence (e.g. kept drinking despite problems, not been able to stick to drinking limits) problem. If a patient does not meet the criteria for alcohol abuse or dependence, the patient is still at-risk for developing alcohol related problems. If a patient has an alcohol use disorder (dependence or abuse), the next step (Step 3) will be different than at-risk drinkers.

Querying abuse and dependence are represented by two separate MDPs for this step. The dialogue state is represented by different features in addition to common features (C, V, Gram) for all states as discussed in Step 1. For *abuse*, we used two specific features. Question (Q) indicates which question is being queried. Since there are 4 questions (slots) for querying abuse, Q can take 1,2,3,4. The second feature specific to abuse is boolean feature Abuse (A). Since it is enough to elicit one abuse indicator, this feature is binary (0 or 1). Since it is enough to elicit one indicator, the system continues to the next step as soon as it elicits one abuse indicator. For *dependence*, a dialogue state is represented by 2 specific and 3 common features. The first specific feature is Question (range 1-7, since there are 7 questions for dependence). The second specific feature is Dependence (D) (range 0-3 which shows the number of indicators elicited: system may not elicit any dependence indicator and 3 dependence indicators is enough to elicit).

### 4.4 Step 3: Advising and Assisting According to Degree of Alcohol Problem

In Step 3, if the client is at-risk, the system states its conclusion according to the guideline [20] and recommends to the user to cut down his/her drinking. Then it tries to asses readiness to change based on readiness ruler approach (e.g. "On a scale of 1 to 10, how important is it for you to make a change?"). If the client is not ready to change, the system restates its concern for client's health, encourages reflection by asking positive versus negatives of drinking and reaffirms its willingness to help when the client is ready. If a client is ready to change, the system sets a goal (e.g. "How could I assist you in getting to a 7?"), agrees on a change plan and provides educational materials (e.g. pamphlets). In Step 3, for the clients who has alcohol abuse or dependence problems, the system states its conclusion, negotiates a drinking behavior goal and refers to an addiction specialist.

### 4.5 Modeling World with Interconnected MDPs

To address the data sparsity problem, we aimed at minimizing the number of system states used. Since the BI dialogue requires many dialogue turns between the system and a client, the number of available dialogue strategies is very large, and can make learning optimal policies infeasible with limited number of training data. To alleviate this problem, we used separate MDPs for each phase.

We represent each step or phase of the BI with one MDP with local goals and reward functions. This divided the problem into 5 interconnected MDPs (see Figure 2) but, in any interaction with the system, we use a maximum 4 MDPs,

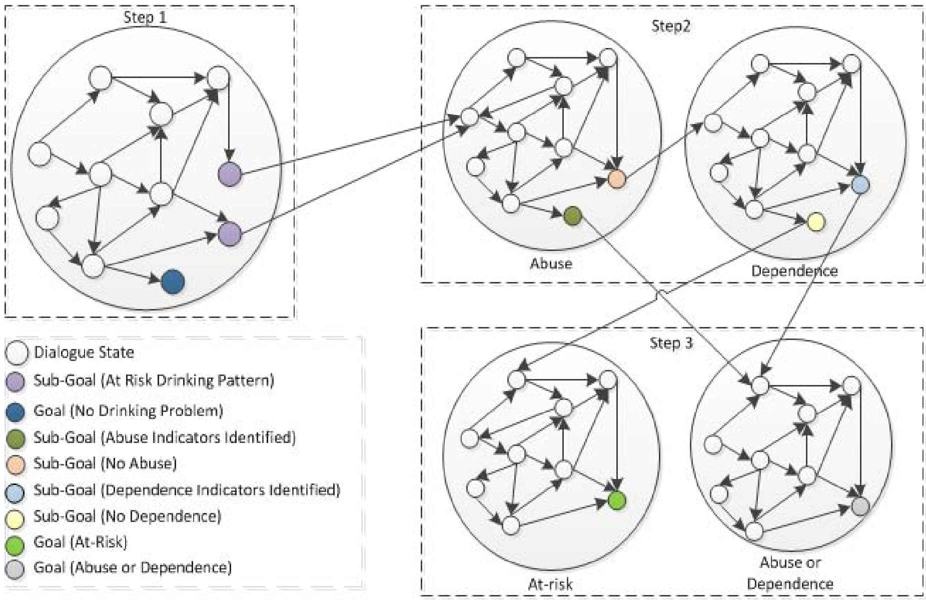


**Fig. 2.** Representation Of World Model With MDPs

i.e. 1) Step 1; 2) Abuse; 3) Dependence; and 4) one MDP from Step 3 based on Abuse or Dependence problem. This approach also reduced the number of required state features for each step, thus reducing the number of states required. For example for Step 1, there are 45 states with 2 action choices, which results in $2^{45}$ possible exploratory dialogue policies.

In *Step 1*, the reward can be awarded based on reaching one of the sub-goal states (i.e. there can be many sub-goals and goal states in each MDP) showed in Figure 2. Goal states, in the MDP representing Step 1, represent completion of the interaction, which means that a client does not have alcohol problem. Sub-goals represent identification of at-risk drinking patterns. If the system reaches one of the sub-goal states in Step 1, it is rewarded with a local reward function and the state is transited to the abuse assessment MDP in Step 2. The transitions between MDPs do not require to have stochastic transition model, thus they are deterministic.

Since there are two phases in *Step 2*, one for querying alcohol abuse and one for querying alcohol dependence, we represent Step 2 with two distinct MDPs (as shown in Figure 2), which greatly reduces number of exploratory policies without compromising fine-grained distinctions between dialogue strategies. Because the two phases are independent from each other, representing each phase with a separate MDP is appropriate. For each of the MDPs, the rewards are awarded based on reaching one of the local sub-goal states. The reward can be awarded in the stage of assessing alcohol abuse, as soon as eliciting one indicator of alcohol abuse, or completing the assessment with 4 questions without eliciting any indicator. The reward in the dependence stage is awarded based on reaching one of the sub-goal states in the MDP. To reach the goal state for dependence, the system needs to identify 3 indicators of the dependence or finish asking all of the questions without eliciting any indicator.

There are two separate MDPs for representing different phases in *Step 3*. One for representing the model for “At-risk” drinkers who does not have alcohol use disorder problems (i.e abuse and dependence). The reward is awarded upon reaching the goal state which is end of the intervention. For the client’s with abuse and dependence problems, the model is represented by MDP which is labeled with ”Abuse or Dependence”. The reward is awarded in the same way as for at-risk drinkers, although the dialogue actions are different.

In conclusion, the system is modeled with 5 MDPs. In each MDP, there are goals and/or sub-goals. Sub-goals represent that the system completed a step but that the interaction is not completed yet. Therefore each sub-goal deterministically transits dialogue state to start state of a successor MDP. At the same time, it awards the agent with the local reward. Local rewards shows how good is a dialogue policy selection for performed dialogue strategy. With this approach, learning the optimal dialogue strategy for an entire dialogue is reduced to learning optimal dialogue strategy for the each MDPs. Finally each goal represents that the interaction is completed and that there is no need to transition to another MDP. As discussed before this approach alleviates the data sparsity problem.

## 5 Components of the Corpus

We plan on creating a corpus from anonymized real user interactions with our training system, it will be used later for learning approximately optimal dialogue strategies. We want to make this as exploratory as possible by rich annotation. We are annotating each dialogue with several objective and subjective performance metrics. We annotate each interaction with best hypothesis of the ASR, ASR confidence scores, n-best hypothesis of ASR (with confidence), system prompts, dialogue acts from NLU, named entities, filled/confirmed slots, dialog context (speech act history), rewards and reward history, dialogue length, number of errors and confirmations.

The corpus represents the dialogues in hierarchical XML structure. Each interaction contains a sequence of turns which includes the system and client utterances, dialogue context (e.g. named entities, filled slots) and rewards. We also annotate subjective reward signals elicited from the user upon completion of the interaction by asking a few questions about ease of use, future intention to use, perceived task completion. We also annotate each dialogue with gender and smile labels with confidence value (for reasons described earlier).

## 6 Conclusion

In this paper, we demonstrated our approach to model relatively long and branching brief intervention dialogue with MDPs. We build an initial training system which can deliver basic unoptimized functionality. Using this training system, we will collect dialogue corpus to help solve optimization problems with RL. One of the largest obstacle building a system for a new domain is the lack of annotated data for training a model. In this project, we addressed the infrastructure needed to collect annotated data.

**Acknowledgment.** This material is based in part upon work supported by the National Science Foundation under Grant Nos. CNS-0821345, CNS-1126619, HRD-0833093, IIP-0829576, CNS-1057661, IIS-1052625, CNS-0959985, OISE-1157372, IIP-1237818, IIP-1215201, IIP-1230661, IIP-1026265, IIP-1058606, IIS-1213026, OISE-0730065, CCF-0938045, CNS-0747038, CNS-1018262, CCF-0937964.

## References

1. Mokdad, A.H., Marks, J.S., Stroup, D.F., Gerberding, J.L.: Actual causes of death in the united states, 2000. JAMA 291(10), 1238–1245 (2000)
2. Bouchery, E.E., Harwood, H.J., Sacks, J.J., Simon, C.J., Brewer, R.D.: Economic costs of excessive alcohol consumption in the U.S., 2006. American Journal of Preventive Medicine 41(5), 516–524 (2011)
3. Moyer, A., Finney, J.W., Swearingen, C.E., Vergun, P.: Brief interventions for alcohol problems: a meta-analytic review of controlled investigations in treatment-seeking and non-treatment-seeking populations. Addiction 97(3), 279–292 (2002)
4. Miller, W.R., Wilbourne, P.L.: Mesa grande: a methodological analysis of clinical trials of treatments for alcohol use disorders. Addiction 97(3), 265–277 (2002)
5. Babor, T.F., Grant, M.: Programme on substance abuse: project on identification and management of alcohol-related problems. Report on phase ii, an randomized clinical trial of brief interventions in primary health care (1992)
6. NIAAA: NIAAA Alcohol Alert No. 66: Brief interventions (2006)
7. Riper, H., Spek, V., Boon, B., Conijn, B., Kramer, J., Martin-Abello, K., Smit, F.: Effectiveness of e-self-help interventions for curbing adult problem drinking: A meta-analysis. J. Med. Internet Res. 13(2), e42 (2011)
8. White, A., Kavanagh, D., Stallman, H., Klein, B., Kay-Lambkin, F., Proudfoot, J., Drennan, J., Connor, J., Baker, A., Hines, E., Young, R.: Online alcohol interventions: A systematic review. J. Med. Internet Res. 12(5), e62 (2010)
9. Portnoy, D.B., Scott-Sheldon, L.A.J., Johnson, B.T., Carey, M.P.: Computer-delivered interventions for health promotion and behavioral risk reduction: a meta-analysis of 75 randomized controlled trials, 1988-2007. Preventive Medicine 47(1), 3–16 (2008)
10. Hester, R.K., Squires, D.D., Delaney, H.D.: The Drinker's Check-up: 12-month outcomes of a controlled clinical trial of a stand-alone software program for problem drinkers. Journal of Substance Abuse Treatment 28(2), 159–169 (2005)
11. Mauriello, L.M., Gökbayrak, N.S., Marter, D.F.V., Paiva, A.L., Prochaska, J.M.: An Internet-Based Computer-Tailored Intervention to Promote Responsible Drinking: Findings from a Pilot Test with Employed Adults. Alcoholism Treatment 30(1), 1–15 (2011)
12. Lisetti, C., Yasavur, U., Leon, C.D., Amini, R., Rishe, N., Visser, U.: Building an On-demand Avatar-based Health Intervention for Behavior Change. FLAIRS (2012)

13. Cassell, J., Nakano, Y.I., Bickmore, T.W., Sidner, C.L., Rich, C.: Non-verbal cues for discourse structure. In: Proceedings of the 39th Annual Meeting on Association for Computational Linguistics, ACL 2001, pp. 114–123 (2001)
14. Moridis, C.N., Economides, A.A.: Affective Learning: Empathetic Agents with Emotional Facial and Tone of Voice Expressions. IEEE Transactions on Affective Computing 3(3), 260–272 (2012)
15. Lisetti, C., Yasavur, U., Leon, C.D., Amini, R.: Building an On-Demand Avatar-Based Health Intervention for Behavior Change. In: Proceedings of the Twenty-Fifth International Florida Artificial Intelligence Research Society Conference (2012)
16. Schulman, D., Bickmore, T.W., Sidner, C.L.: An Intelligent Conversational Agent for Promoting Long-Term Health Behavior Change using Motivational Interviewing. In: Association for the Advancement of Artificial Intelligence (AAAI) Spring Symposium Series, pp. 61–64. Association for the Advancement of Artificial Intelligence (2011), http://www.aaai.org
17. Morbini, F., Forbell, E., DeVault, D., Sagae, K., Traum, D.R., Rizzo, A.A.: A mixed-initiative conversational dialogue system for healthcare. In: Proceedings of the 13th Annual Meeting of the Special Interest Group on Discourse and Dialogue, SIGDIAL 2012, pp. 137–139. Association for Computational Linguistics, Stroudsburg (2012)
18. Wang, N., Gratch, J.: Don't just stare at me. In: CHI, Atlanta, GA, USA, pp. 1241–1249 (2010)
19. Kang, S.H., Gratch, J., Wang, N., Watt, J.: Does the contingency of agents' nonverbal feedback affect users' social anxiety? In: Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS, vol. 1, pp. 120–127. International Foundation for Autonomous Agents and Multiagent Systems (2008)
20. NIAAA: Helping Patients Who Drink Too Much, A Clinician's Guide (2007)
21. Bickmore, T.W., Schulman, D., Sidner, C.L.: A reusable framework for health counseling dialogue systems based on a behavioral medicine ontology. Journal of Biomedical Informatics 44(2), 183–197 (2011)
22. Ferguson, G., Quinn, J., Horwitz, C., Swift, M., Allen, J., Galescu, L.: Towards a Personal Health Management Assistant. Journal of Biomedical Informatics 43(5 suppl.), S13–S16 (2010)
23. Bickmore, T.W., Puskar, K., Schlenk, E.A., Pfeifer, L.M., Sereika, S.M.: Maintaining reality: Relational agents for antipsychotic medication adherence. Interacting with Computers 22(4), 276–288 (2010)
24. Ferguson, G., Allen, J., et al.: Trips: An integrated intelligent problem-solving assistant. In: Proceedings of the National Conference on Artificial Intelligence, pp. 567–573. John Wiley & Sons Ltd. (1998)
25. Bohus, D., Rudnicky, A.I.: The ravenclaw dialog management framework: Architecture and systems. Computer Speech & Language 23(3), 332–361 (2009)
26. Traum, D., Larsson, S.: The Information State Approach to Dialogue Management, pp. 325–353. Kluwer Academic Publishers (2003)
27. Bos, J., Klein, E., Lemon, O., Oka, T.: Dipper: Description and formalisation of an information-state update dialogue system architecture. In: 4th SIGdial Workshop on Discourse and Dialogue, pp. 115–124 (2003)
28. Sutton, R.S., Barto, A.G.: Reinforcement learning: An introduction, vol. 1. Cambridge Univ. Press (1998)

29. Lemon, O., Pietquin, O., et al.: Machine learning for spoken dialogue systems. In: Proceedings of the European Conference on Speech Communication and Technologies (Interspeech 2007), pp. 2685–2688 (2007)
30. Singh, S., Litman, D., Kearns, M., Walker, M.: Optimizing dialogue management with reinforcement learning: Experiments with the NJFun system. Journal of Artificial Intelligence Research 16, 105–133 (2002)
31. Young, S., Gašić, M., Thomson, B., Williams, J.: Pomdp-based statistical spoken dialog systems: A review (2013)
32. Papangelis, A., Kouroupas, N., Karkaletsis, V., Makedon, F.: An adaptive dialogue system with online dialogue policy learning. In: Maglogiannis, I., Plagianakos, V., Vlahavas, I. (eds.) SETN 2012. LNCS, vol. 7297, pp. 323–330. Springer, Heidelberg (2012)
33. Frampton, M., Lemon, O.: Recent research advances in reinforcement learning in spoken dialogue systems. Knowledge Engineering Review 24(4), 375–408 (2009)
34. Young, S., Gašić, M., Keizer, S., Mairesse, F., Schatzmann, J., Thomson, B., Yu, K.: The hidden information state model: A practical framework for pomdp-based spoken dialogue management. Computer Speech & Language 24(2), 150–174 (2010)
35. Chi, M., VanLehn, K., Litman, D.: Do micro-level tutorial decisions matter: Applying reinforcement learning to induce pedagogical tutorial tactics. In: Aleven, V., Kay, J., Mostow, J. (eds.) ITS 2010, Part I. LNCS, vol. 6094, pp. 224–234. Springer, Heidelberg (2010)
36. Walker, M.A., Passonneau, R., Boland, J.E.: Quantitative and qualitative evaluation of DARPA communicator spoken dialogue systems. In: Proceedings of the 39th Annual Meeting on Association for Computational Linguistics, ACL 2001, pp. 515–522. Association for Computational Linguistics, Stroudsburg (2001)
37. Georgila, K., Lemon, O., Henderson, J.: Automatic annotation of communicator dialogue data for learning dialogue strategies and user simulations. In: Ninth Workshop on the Semantics and Pragmatics of Dialogue (SEMDIAL: DIALOR). Citeseer (2005)
38. Frampton, M., Lemon, O.: Learning more effective dialogue strategies using limited dialogue move features. In: Proceedings of the 21st International Conference on Computational Linguistics and the 44th Annual Meeting of the Association for Computational Linguistics, pp. 185–192. Association for Computational Linguistics (2006)
39. Ward, W.: Understanding spontaneous speech: the Phoenix system. In: Proceedings of the 1991 International Conference on Acoustics, Speech, and Signal Processing, ICASSP 1991, pp. 365–367. IEEE Computer Society, Washington, DC (1991)
40. Yasavur, U., Amini, R., Lisetti, C., Rishe, N.: Ontology-based named entity recognizer for behavioral health. In: The Twenty-Sixth International FLAIRS Conference (2013)