

# Wizard of Oz Evaluation of Listening-oriented Dialogue Control using POMDP

Toyomi Meguro<sup>†</sup>, Yasuhiro Minami<sup>†</sup>, Ryuichiro Higashinaka<sup>‡</sup>, and Kohji Dohsaka<sup>†</sup>,

<sup>†</sup> NTT Communication Science Laboratories, NTT Corporation  
2-4 Hikaridai, Seika-cho, Soraku-gun, Kyoto 619-0237 Japan

<sup>‡</sup> NTT Cyber Space Laboratories, NTT Corporation  
1-1 Hikarinooka, Yokosuka, Kanagawa, 239-0847 Japan

{meguro.toyomi, minami.yasuhiro, higashinaka.ryuichiro, dohsaka.kohji}@lab.ntt.co.jp

**Abstract**—We have been working on dialogue control for listening agents. In our previous study [1], we proposed a dialogue control method that maximizes user satisfaction using partially observable Markov decision processes (POMDPs) and evaluated it by a dialogue simulation. We found that it significantly outperforms other stochastic dialogue control methods. However, this result does not necessarily mean that our method works as well in real dialogues with human users. Therefore, in this paper, we evaluate our dialogue control method by a Wizard of Oz (WoZ) experiment. The experimental results show that our POMDP-based method achieves significantly higher user satisfaction than other stochastic models, confirming the validity of our approach. This paper is the first to show the usefulness of POMDP-based dialogue control using human users when the target function is to maximize user satisfaction.

## I. INTRODUCTION

We have been working on listening-oriented dialogues (LoD) in which one conversational participant attentively listens to the other (See Fig. 1 for a typical LoD). Our aim is to build listening agents that can implement a human-like listening process so that users can satisfy their desire to speak and to be heard. As a first step for building such listening agents, we are focusing on a dialogue control module that chooses the best system dialogue act (i.e., meaning representation of an utterance) at every dialogue point using the dialogue history and the user dialogue act as input. We focus on this module because it is the most important one in a listening agent, since it decides the dialogue flow and because our analysis of LoDs showed that dialogue flow is what characterizes LoDs [2].

Since it is difficult to completely handcode a listening process, it is desirable that we train a dialogue control module from data. For this purpose, we have been adopting partially observable Markov decision processes (POMDPs), which can learn an optimal dialogue control policy from dialogue data. Our current objective is to train, by using POMDPs, a dialogue control policy that can satisfy users by attentively listening to them. Note that our target function is different from conventional POMDP-based work [3] that maximizes task-completion; little work has focused on maximizing user satisfaction because defining a target function is difficult. In our previous work [1], we introduced two kinds of rewards (user satisfaction and smoothness) for the target function of a POMDP, trained a dialogue control module, and verified its

Utterance	Dialogue act
S: Good evening.	GREETING
L: Good evening.	GREETING
S: Our dialogue topic is “food.”	GREETING
I had curry and rice for dinner.	S-DISC (sub: fact)
Do you like curry and rice?	QUESTION (sub: fact)
L: Nice to meet you.	GREETING
What kind of curry?	QUESTION (sub: fact)
S: Homemade.	S-DISC (sub: fact)
L: I love curry!	S-DISC (sub: fact)
S: No secret ingredients whatsoever.	S-DISC (sub: fact)
L: No secret ingredients!	REPEAT

Fig. 1. Excerpt of a typical listening-oriented dialogue. Dialogue acts corresponding to utterances are shown in parentheses (See Table 1 for their meanings). S-DISC stands for SELF-DISCLOSURE, S for speaker, and L for listener. The dialogue was translated from Japanese by the authors.

superiority to other stochastic methods by dialogue simulation; that is, we used a user simulator to create sets of simulated dialogues, which were then evaluated by human raters. Human raters were needed at the final evaluation stage because automatically estimating user satisfaction with precision remains difficult [4].

Although we successfully showed that our module performs well in simulations, this does not necessarily mean that it works as well in real dialogues with human users; when humans use a system, they greatly change their behavior based on system actions, which is not true in simulations. Simulations are also fraught with the fundamental difficulty of perfectly tracing human behavior, especially when there is less task-restriction as in LoDs. Since we are aiming at building listening agents that can talk to real users and satisfy them, it is crucial that we verify the validity of our dialogue control method with human users.

Therefore, in this paper, we evaluate our dialogue control method by a Wizard of Oz (WoZ) experiment. Since we only focus on the dialogue control performance in this work, we introduce WoZ experimenters for understanding user utterances and generating system utterances in natural language from system dialogue acts. Our contribution lies in being the first to examine the usefulness of POMDP-based dialogue control using real human users when the target function is to maximize user satisfaction.

## II. RELATED WORK

With increased attention on social dialogues and senior peer counseling, work continues to emerge on listening-oriented dialogues, ranging from the analysis of LoDs to system implementations.

Regarding analysis, in a WoZ experiment, Maatman et al. showed that virtual agents can provide users a sense of being heard using such gestures as nodding and head shaking [5]. Our previous work [2] analyzed the characteristics of LoDs, in which LoDs were compared with casual conversations, revealing that the two types of dialogues have significantly different flows and that listeners actively question with frequently inserted self-disclosures. The speaker utterances were mostly concerned with self-disclosure.

Shitaoka et al. [6] implemented listening agents and focused on their response generation components. Their system considers the confidence scores of speech recognition and changes the system response accordingly; it repeats the user utterance or makes an empathic utterance for high-confidence user utterances and provides backchanneling when the confidence is low. The system’s empathic utterances include “I’m happy” or “That’s too bad,” depending on whether a positive or negative expression is found in the user utterances. Their system’s response generation only uses the speech recognition confidence and the polarity of user utterances as cues to choose its actions. Currently, it does not consider the utterance content such as dialogue acts. For listening agents to achieve high smoothness, a switching mechanism between the “active listening mode,” in which the system is a listener, and the “topic presenting mode,” in which the system is a speaker, has been proposed [7], [8]. Here, the objective of their system is to maintain high user interest and to keep the system in an active listening mode.

All of these implementations use hand-crafted rules for dialogue control, which we considered needed automation and led to our previous work using POMDPs [1]. POMDPs were first introduced for robot action control to learn suitable movements for completing a certain task [9]. Only recently have they been actively applied to dialogue systems. Williams et al. [3] successfully used POMDPs for dialogue control in a ticket-buying domain to fix the departure and arrival places for tickets. Recent work on POMDPs has indicated that dialogue control modules can be trained in task-oriented dialogues, but it is not obvious whether they can be used for less task-oriented dialogues such as LoDs. We verified their effectiveness for LoDs by simulation, but further verification by experiments using human users is necessary.

## III. POMDP-BASED DIALOGUE CONTROL

Before the WoZ experiment, we made two enhancements to our proposed dialogue control that were necessary for natural interactions between the WoZ experimenters and the human users. We first describe our POMDP-based dialogue control and then our enhancements.

TABLE I

DEFINITIONS OF DIALOGUE ACTS WITH SAMPLE GENERATION TEMPLATES THAT WERE USED BY A WOZ EXPERIMENTER TO GENERATE UTTERANCES FOR DIALOGUE ACTS AND ADAPTED FOR THE “FOOD” TOPIC. EXCEPT FOR REPEAT, PARAPHRASE, AND OTHER, EACH DIALOGUE ACT HAS SEVERAL GENERATION TEMPLATES SHOWN IN BRACKETS. (TIME) CAN BE REPLACED WITH LUNCH, BREAKFAST, OR DINNER, AND (DAY) WITH TOMORROW, TODAY, ETC. (FOOD) IS FILLED WITH ARBITRARY ITEMS BY A WOZ EXPERIMENTER BASED ON THE DIALOGUE CONTEXT.

Dialogue act	Definition and sample generation templates
GREETING	Greeting and confirmation of a dialogue theme. e.g., [Hello.], [The dialogue theme is XX.]
INFORMATION	Delivery of objective information. [(place-name) is famous for (food).]
SELF-DISCLOSURE	Disclosure of preferences and feelings.
sub: fact	[I ate (food) for (time) (day).]
sub: experience	[I have eaten (food).]
sub: habit	[I always go out to eat.]
sub: preference (positive)	[I like (food).]
sub: preference (negative)	[ (food) is delicious.]
sub: preference (neutral)	[I don’t like (food).] [(food) is not delicious.]
sub: desire	[[food] is (adjective: either positive or negative).]
sub: plan	I want to eat (food) for (time) (day). I will eat (food) for (time) (day).
ACKNOWLEDGMENT	Encourages conversational partner to speak. [Well.] [Aha.]
QUESTION	Utterances that expect answers.
sub: information	[Please tell me about (food).]
sub: fact	[What did you eat for (time) (day)?]
sub: experience	[Have you eaten (food) before?]
sub: habit	[Do you usually go out to eat?]
sub: preference	[Do you like (food)?]
sub: desire	[What do you want to eat for (time) (day)?]
sub: plan	[What are you going to have for (time) (day)?]
SYMPATHY	Sympathetic utterances and praises. [Me, too]
NON-SYMPATHY	Negative utterances. [Me neither.]
CONFIRMATION	Confirm what conversation partner said. [Really?]
PROPOSAL	Encourage partner to act. [Please try (food).]
REPEAT	Repeat partner’s previous utterance.
PARAPHRASE	Paraphrase partner’s previous utterance.
APPROVAL	Bring up or show goodwill toward partner. [Absolutely!]
THANKS	Express thanks [Thank you.]
APOLOGY	Express regret [I’m sorry.]
FILLER	Filler between utterances. [Uh.], [Let me see.]
ADMIRATION	Express affection. [A-ha-ha.]
OTHER	Other utterances.

### A. Training a Dialogue Policy for LoDs

In a POMDP, we train a policy that uses the sum of two rewards: user satisfaction (i.e., how satisfied one is with the dialogue flow) and smoothness (i.e., how probable/natural the dialogue flow is) [1]. These two rewards are necessary because we found that if either is missing, unnatural or unsatisfactory dialogues result [10].

To obtain our POMDP structure, we first trained the dynamic Bayesian network (DBN) structure shown in Fig. 2 (left) from collected LoDs with their user satisfaction evaluation scores (obtained by questionnaires). Here,  $s_o$  and  $s_a$  are the dialogue and action states,  $o$  is a speaker observation,  $a$  is a listener action, and  $d$  is a random variable for an evaluation score. Then the DBN structure is converted into the POMDP structure (Fig.2, right), where  $r$  is the reward for taking action  $a$  in  $s_o$ , which is defined as the weighted sum of  $r1$  (user satisfaction) and  $r2$  (smoothness), where  $r1$  is the averaged evaluation score for  $(s_o, a)$  and  $r2$  is the reward for the

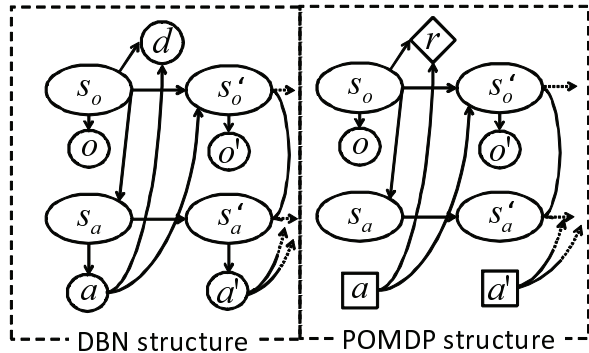


Fig. 2. DBN and POMDP structures

predictive probability of  $a$ .

We can calculate the average reward that will be gained in the future at time  $t$  by:

$$V_t = \sum_{\tau=0}^{\infty} \gamma^{\tau} \sum_s b_{\tau+t}((s_o, s_a)) r((s_o, s_a), a_{\tau+t}),$$

where  $b_t$  is the belief state (probability distribution over states) and  $\tau$  is a discount factor. Here,  $r((s_o, s_a), a_{\tau+t})$  is the weighted sum of the two rewards. Finally, a policy is learned by value iteration so that the action can be chosen that maximizes  $V_t$ .

#### B. Enhancement 1: Dealing with Multiple Dialogue Acts

In our POMDP structure (Fig. 2), a system dialogue act and a user observation have to occur alternately. However, as the example dialogue (Fig. 1) illustrates, humans do not necessarily produce a single dialogue act in one turn. Since the primary purpose of LoD is to let users speak freely and allow them to feel that they are being listened to, we should avoid limiting the length of user utterances. A system that only produces one dialogue act at all turns would also be unnatural.

Therefore, we modified our method to allow multiple dialogue acts per user/system turn by introducing a “skip” action. For training a policy that has “skip,” we first re-annotate the dialogue acts for each utterance in the collected LoDs. Since there can be one or more dialogue acts in one utterance, we insert “skip” between the dialogue acts, meaning that a “skip” operation was done by the conversational partner. Finally, using the skip-inserted data, we learn a policy, as described in Section III-A. By having “skip,” multiple dialogue acts are treated as follows:

- When the user input has multiple dialogue acts in one utterance, the system assumes that the system’s “skip” has been inserted between the dialogue acts and calculates the output and transition probabilities.
- When the system has the turn, the next user dialogue act is estimated by a user simulator. If the estimation result is “skip,” the system remembers that the user act was “skip” and calculates the next system action. At this step, the system still does not output its action. Only when the simulator estimate is not “skip” does this estimation process stop. The system outputs all dialogue

acts generated until the termination of the estimation. Here, the system basically chooses the first dialogue act that has the largest reward at each step. However, if the first dialogue act is “skip,” the system chooses the second one, because if the system chooses “skip” for the user’s estimated “skip”, this could result in uncooperative dialogue.

By these processes, the system can treat multiple dialogue acts.

#### C. Enhancement 2: Feedback from System Action

For a system to generate more appropriate system dialogue acts, we modified the states for system action  $s_a$  in Fig. 2. Here,  $s_a$  is a hidden state and the system has the distribution of  $s_a$  at each dialogue step. In our proposed model, each state of  $s_a$  has a one-to-one correspondence to each system dialogue act (See [1] for details, but we did this to facilitate the calculation of the predictive probability of  $s_a$ ).

In an on-going dialogue, when the system generates the best dialogue act, it can immediately be aware of the system’s dialogue act. With this information as feedback, the system can expect better distributions for the next step. To be more concrete, after the system’s generation of its dialogue acts,  $s_a$  can be set as non-distribution and the states can be clearly decided; that is, we set the state’s probability of the best dialogue act to one and the probabilities of all other states to zero.

## IV. EXPERIMENT

### A. Experimental Setup

We evaluated our system by a WoZ experiment. Here, our system means our proposed system in [1] with the enhancements explained in the previous section. In a dialogue of the experiment, one WoZ experimenter acted as a listener and one evaluator as the speaker. They chatted using a text-chat interface in separate rooms. The use of facial expressions was not allowed. Our system was compared to six other systems (described in Section IV-B). For each system including ours, a dialogue took place in the following steps:

- 1) The evaluator as a speaker inputs a natural language utterance in the chat interface to talk to the system. The evaluator does not know that a human is behind the system.
- 2) The WoZ experimenter sees the evaluator’s utterance and converts the natural language utterance into a dialogue act sequence based on the definition of dialogue acts in Table I. Then he/she inputs the converted dialogue acts as the speaker’s utterance to the system.
- 3) The system receives the dialogue act sequence and generates the system’s next dialogue act sequence.
- 4) The WoZ experimenter sees the dialogue act sequence and generates the system’s natural language utterance from the sequence. Since each dialogue act has several possible generation templates, as shown in Table I, the WoZ experimenter chooses an appropriate one matching the dialogue flow. Some templates have blanks to be

freely filled with well-suited words by the experimenter based on the context.

- 5) The WoZ experimenter inputs the generated utterance into the chat interface and waits for the next user utterance. (Return to Step 1).

We recruited 14 participants (seven males; seven females) as evaluators and two participants (one female; one male) as WoZ experimenters. The participants did not include the authors. The WoZ experimenters and the evaluators could not see each other and were strictly kept in different rooms during the experiments, including breaks between the dialogues. Each evaluator talked with two WoZ experimenters and talked seven times with each WoZ experimenter, each time WoZ experimenter using one of the seven different systems (the proposed system plus six systems for comparison). Therefore, each evaluator participated in 14 dialogues. Each dialogue lasted approximately 15 minutes. All dialogues were about “food,” and each had a more specific theme: “yesterday’s meal,” “favorite food,” “most impressive food,” “specialty food (If you do not have a specialty, your family, partner, friends’ are possible),” “food on trips,” “dining out (recommendations of restaurants, convenience stores, etc.),” “famous food in your hometown.” A theme was chosen randomly for each dialogue.

After each dialogue, the evaluators filled in questionnaires to rate their satisfaction level by answering the following question on a 7-point Likert scale: “Did you feel that you were listened to?” At the same time, the WoZ experimenters filled in similar questionnaires and answered, “Were the dialogue acts generated by the system good in terms of listening through the entire dialogue?”.

## B. Systems

In addition to our proposed system, we prepared six dialogue control systems for comparison. For training, the POMDP systems (Sections IV-B1 and IV-B2) and HMM (Section IV-B5) used the dialogue-act annotated LoDs between humans that we collected in our previous study (hereafter, **LoD corpus**) [1]. The LoD corpus was annotated using the dialogue act tag set in Table I. The corpus contains 1259 LoDs. In what follows, we describe the systems we used in the WoZ experiment one by one.

1) *Proposed POMDP*: This is our proposed system. The number of states and weight coefficients of the two rewards (user satisfaction and smoothness) is same as in our previous study [1].

2) *POMDP-only SA*: To clarify whether both user satisfaction and smoothness rewards in our method are effective, we prepared two systems, each of which used one of the rewards. This system’s structure and learning method are exactly the same as our proposed system, but the reward is only calculated using user satisfaction. We call it the POMDP system with only user satisfaction POMDP-only SA.

3) *POMDP-only Prob*: This system, which is the counterpart of POMDP-only SA, uses only the smoothness reward and is called POMDP-only Prob because it only uses the predictive probability of actions.

4) *Rule*: This system (hereafter, **Rule system**) uses dialogue control rules written by the authors. We conceived them based on our analysis of LoDs [2] and other counseling-related research. The system only sees the user’s last user dialogue act and generates up to four dialogue acts per turn (number of dialogue acts is determined randomly). The rules were verified to work well by a separate WoZ experiment [11].

5) *HMM*: This system (hereafter, **HMM system**) performs dialogue control using an HMM trained from LoDs. We have this system because an HMM is the most typical stochastic model in previous dialogue research [12]. It chooses the most probable dialogue act at each step and predicts whether it should continue its utterance or pass its turn to the user. If the most probable action at the next step is that of the user, the system stops generating dialogue acts. If the most probable dialogue act is that of the system, the system continues to generate dialogue acts by greedily choosing the most probable dialogue act at each step. The number of states for the HMM was 18.

6) *Human Dialogue Control*: When using this system, the WoZ experimenters choose the system dialogue acts (both the dialogue acts and their number) at their own discretion. We regard this system as an upper bound and call it the **Human system**.

7) *Random Dialogue Control*: This system randomly chooses the system dialogue acts. The number of dialogue acts is also randomly decided from one to four. We call this the **Random system**.

## C. Experimental Results

1) *Subjective Evaluation*: Figure 3 shows the averaged user satisfaction ratings for the seven systems and the result of the WoZ experimenters’ subjective evaluations. The user satisfaction rating of our proposed system is the same as the Rule system and superior to all of the following: the Random system, the HMM system, the POMDPs only Prob and the POMDPs only SA. Except for the Rule system, there are significant differences ( $p < 0.05$ ) by a statistical test (t-test) between our proposed system and the other systems.

Our method is superior to the other two POMDP systems with one reward, which means that both rewards are necessary. It also outperformed the HMM system, which is one representative of stochastic models, which means that the generation of dialogue acts based only on probability is not sufficient for dialogue control. The user satisfaction ratings of the Rule system and our proposed method are almost the same. However, since the rules took several months to write, our proposed method’s performance is surprising because it was only trained from data.

The Random system is superior to the HMM system and the two POMDP systems with one reward, probably because of the issue of the number of dialogue acts per utterance. In the LoD corpus (training data for the POMDP-based and the HMM systems), the average number of dialogue acts per utterance is 1.6, and consequently, the POMDP-based and HMM systems generated around one to two dialogue acts per utterance. This

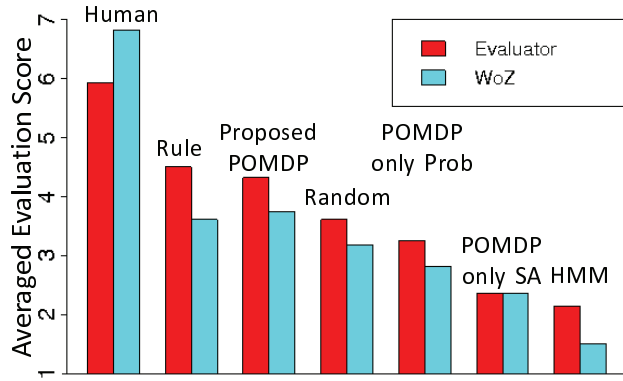


Fig. 3. Averaged user satisfaction ratings

number of dialogue acts was probably affected by the chat interface we used to collect the data; entering long (two or more) sentences was difficult. However, we used another interface with better stability for this WoZ experiment, and longer utterances were more easily entered. For example, in the Human system, the WoZ experimenters entered an average of 2.3 dialogue acts per utterance. In the Random system, they averaged 2.4 dialogue acts. Note that the Random system generates one to four dialogue acts per utterance. Since the number of dialogue acts for the Random system was closer to that of the Human system than others, it probably induced good subjective evaluations. Another possible reason is that the Random system’s generation was made more natural by the fact that it outputs more than two dialogue acts on average. Since the probability of generating self-disclosure or questions is almost half (Table I; there are 15 dialogue acts related to self-disclosure and questions based on their subcategories), each utterance of the Random system has at least one self-disclosure or question. This must have made it easier for the WoZ experimenters to recover from problems caused by random dialogue acts that do not fit the context. Even though the Random system benefited from such long dialogue acts, our proposed system outperformed it with a smaller number of dialogue acts (1.4), which indicates how accurately appropriate dialogue acts were selected by our method.

Figure 3 also shows the WoZ experimenters’ evaluation scores (See Section IV-A for the questionnaire item). Except for the Rule and Random systems, we found significant differences between our proposed system and other systems ( $p < 0.05$ ) by a statistical test (t-test). The orders of the systems of the user satisfaction ratings and the WoZ experimenters’ evaluation scores were almost the same except that the order of our proposed method and the Rule system was inverted; our proposed method outperformed the Rule system. In the WoZ experimenters’ evaluations, the system’s dialogue acts were perceived to be better than the Rule system, indicating the effectiveness of our proposed method. We also find the asymmetry of the results (proposed POMDP vs. Rule system) interesting; probably, the evaluators preferred rules because of its predictability whereas the experimenters liked some unpredictability to make conversations more engaging.




Fig. 4. Kullback-Leibler divergence of output distributions of listener’s dialogue acts when compared with LoDs used for LoD corpus and Human systems

2) *Objective Evaluation:* To further investigate the dialogue quality, for the dialogues collected by each system, we calculated the Kullback-Leibler (KL) divergence of the output distribution of the listener’s dialogue acts using the LoD corpus and the dialogues of the WoZ experimenters manipulating the Human system; the more closely the distribution resembles that of the distribution for the LoD corpus (human-human dialogues) or of the distribution of the WoZ experimenters using the Human system, the better the system should be.

Figure 4 shows the KL divergence results for each system. Low scores mean that the output distribution of the dialogue acts are similar. Our proposed method’s KL divergence for the Human system and the LoD corpus are smaller than the other systems. Therefore, we can confirm that it successfully generated dialogue acts that are more similar to human’s dialogue control than the others. The KL divergence between the Random system and the LoD corpus is small. This means that if we ignore the dialogue flow (NB. KL does not take into transition probabilities), humans generate dialogue acts rather evenly. However, the user satisfaction of the Random system is lower than the Human system, suggesting that dialogue flow is important for user satisfaction. The HMM system’s output distribution is far from that of the Human system and the LoD corpus, probably because it generated only highly frequent dialogue acts found in the LoD corpus.

To investigate the quality of the dialogue flow, we constructed language models using dialogue-act trigrams of the dialogues of each system and measured the perplexity when the dialogue act sequences of the Human system and those of the LoD corpus were input into each model. The results are shown in Fig. 5.

If the scores are low, the probability of the N-grams is similar between the Human system/LoD corpus and each system. The probability of the dialogue act sequences of our proposed system is as good as the Rule system and superior to the other systems, which is the same tendency we found for the user satisfaction ratings. This means that our proposed method and the Rule system generated dialogue act sequences similar to humans. The HMM system is just behind the Rule

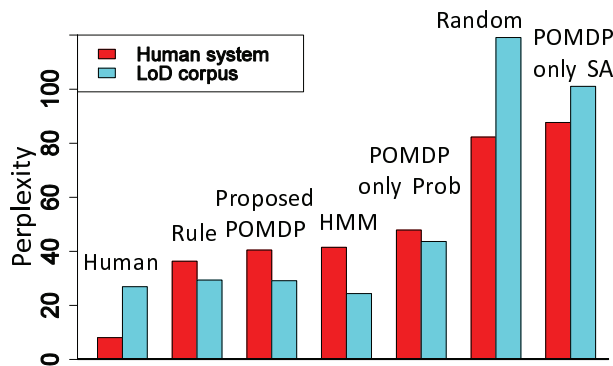


Fig. 5. Perplexity of each system when compared with the models of the dialogues of the Human system and the LoD corpus.

system and our proposed method, which is reasonable because it generates dialogue acts using the transition probabilities of dialogue acts in data.

#### D. Obtained Dialogues

Figures 6 and 7 show the dialogues with our proposed method (Proposed POMDP) and with the Human System. In Fig. 6, natural language utterances and their dialogue acts are shown in brackets. For the dialogue of Proposed POMDP, we also show the 5-best dialogue acts with their reward scores at each step. Since the proposed system generated timely self-disclosures and questions, the tendency of which is seen when the Human system is employed as in Fig. 7, we consider that our proposed method successfully learned human dialogue control from data.

### V. CONCLUSION AND FUTURE WORK

In this paper, we evaluated our POMDP-based dialogue control method by a Wizard of Oz (WoZ) experiment and found that our POMDP-based method achieves significantly higher user satisfaction than other stochastic models, confirming the validity of our approach. This paper is the first to show the usefulness of POMDP-based dialogue control using human users when the target function maximizes user satisfaction instead of task-completion.

Future work will implement our natural language understanding and generation components since currently, our proposed method can only treat dialogue acts. In addition, we want to consider other tasks than LoDs, in which we need to consider different rewards. Although we believe our approach is language-independent, we must also evaluate it in other languages such as English. In addition, we want to pursue methods to efficiently collect dialogue data for training.

#### REFERENCES

[1] T. Meguro, R. Higashinaka, Y. Minami, and K. Dohsaka, “Controlling listening-oriented dialogue using partially observable Markov decision processes,” in *Proc. COLING*, 2010, pp. 761–769.

[2] T. Meguro, R. Higashinaka, K. Dohsaka, Y. Minami, and H. Isozaki, “Analysis of listening-oriented dialogue for building listening agents,” in *Proc. SIGDIAL*, 2009, pp. 124–127.

Utterance and dialogue acts (5-best for system dialogue acts)	
L: Nice to talk to you. [GREETING:26.0 Q (sub:desire):24.9 APPROVAL:20.7 SKIP:20.2 S-DISC (sub:pref-):20.1]	
S: Nice to talk to you, too. [GREETING]	
L: Are you interested in trying green caviar? [QUESTION (sub: desire):24.0 SKIP:21.9 GREETING:21.1 APPROVAL:20.1 S-DISC (sub:pref-):20.1]	
S: Yes, I am. [SYMPATHY]	
I’ve never eaten it. [S-DISC (sub: experience)]	
L: I like to eat it in Okinawa. [SKIP:25.6 S-DISC (sub:pref+):21.0 CONFIRMATION:20.7 S-DISC (sub:desire):20.7 INFO:20.4]	
S: Let me talk about this dialogue’s theme: “food on trips.” [S-DISC (sub: fact)]	
I have some memories of Okinawa. [S-DISC (sub: pref+)]	
L: Oh, really? [SKIP:24.1 CONFIRMATION:20.5 S-DISC (sub: pref+):20.5 INFO:20.0 QUESTION (sub: fact):20.0]	
S: I ate rafute (braised pork belly) and goya champlo (sauteed vegetables) there. [S-DISC(sub: fact)]	

Fig. 6. Excerpt of dialogue between a WoZ experimenter (using our proposed method) and an evaluator. Theme is “food on trips.” INFO stands for INFORMATION and pref+ for positive preference. S for the evaluator; and L is the WoZ (listener).

Utterance	Dialogue act
L: Nice to talk to you.	GREETING
Hello.	GREETING
S: Hello.	GREETING
Nice to talk to you, too.	GREETING
Let’s talk about “Dining out.”	GREETING
L: Sure.	SYMPATHY
Let’s start.	GREETING
S: I rarely go out to eat.	S-DISC (sub: fact)
L: Really?	CONFIRMATION
Do you usually eat at home?	QUESTION (sub: habit)
S: Yes, I do.	SYMPATHY
I usually eat at home.	SYMPATHY
Do you frequently go out to eat?	QUESTION (sub: habit)

Fig. 7. Excerpt of dialogue between a WoZ experimenter (using Human system) and an evaluator. Theme is “dining out.”

[3] J. Williams and S. Young, “Partially observable Markov decision processes for spoken dialog systems,” *Computer Speech & Language*, vol. 21, no. 2, pp. 393–422, 2007.

[4] R. Higashinaka, Y. Minami, K. Dohsaka, and T. Meguro, “Issues in predicting user satisfaction transitions in dialogues: Individual differences, evaluation criteria, and prediction models,” in *Proc. IWSDS*, 2010, pp. 48–60.

[5] R. M. Maatman, J. Gratch, and S. Marsella, “Natural behavior of a listening agent,” *Lecture Notes in Computer Science*, vol. 3661, pp. 25–36, 2005.

[6] K. Shitaoka, R. Tokuhisa, T. Yoshimura, H. Hoshino, and N. Watanabe, “Active listening system for dialogue robot,” in *Proc. SIG-SLUD, JSAI*, vol. 58, 2010, pp. 61–66, (in Japanese).

[7] S. Yokoyama, D. Yamamoto, Y. Kobayashi, and M. Doi, “Development of dialogue interface for elderly people—switching the topic presenting mode and the attentive listening mode to keep chatting—,” in *IPSJ SIG Technical Report*, vol. 2010-SLP-80, no. 4, 2010, pp. 1–6, (in Japanese).

[8] Y. Kobayashi, D. Yamamoto, T. Koga, S. Yokoyama, and M. Doi, “Design targeting voice interface robot capable of active listening,” in *Proc. HRI*, 2010, pp. 161–162.

[9] J. Pineau, G. Gordon, and S. Thrun, “Point-based value iteration: An anytime algorithm for POMDPs,” in *Proc. IJCAI*, 2003, pp. 1025–1032.

[10] Y. Minami, R. Higashinaka, K. Dohsaka, T. Meguro, and E. Maeda, “Trigram dialogue control using POMDPs,” in *Proc. SLT*, 2010, pp. 336–341.

[11] T. Meguro, R. Higashinaka, Y. Minami, and K. Dohsaka, “Evaluation of listening-oriented dialogue control rules based on the analysis of HMMs,” in *Proc. Interspeech*, 2011, (to appear).

[12] K. Shirai, “Modeling of spoken dialogue with and without visual information,” in *Proc. ICSLP*, vol. 1, 1996, pp. 188–191.