News Navigation System based on Proactive Dialogue Strategy

Koichiro Yoshino and Tatsuya Kawahara

Abstract This paper addresses the concept of information navigation and the system that navigates news articles updated day-by-day. In the information navigation, the system has a back-end knowledge base and users can access information through a natural interaction. It is composed of several modules that interact with users in different manners. Both the system and the user can take an initiative of dialogue depending on the specification of the user interest. The system allows ambiguous user queries and proactively presents information related to the user interest by tracking the user focus. An experimental result shows that the proposed system based on POMDP and user focus tracking can interact with users effectively by selecting the most appropriate dialogue modules.

1 Introduction

Studies on spoken dialogue systems now enter a new stage. A large number of spoken dialogue systems have been investigated and many systems are now deployed in the real world, most typically as smart phone applications, which interact with a diversity of users. However, a large majority of current applications is based on a specific task description which includes a definite task goal and necessary slots, such as place and date, for the task completion [1, 2]. Users are required to follow these concepts and they need to be aware of the clear task goal according to the system's capability. On the other hand, keyword search systems and question answering systems with a speech interface are also developed for smart-phone applications. Such

Koichiro Yoshino

Academic Center for Computing and Media Studies, Kyoto University, Sakyo, Kyoto, 606-8501
e-mail: yoshino@ar.media.kyoto-u.ac.jp

Tatsuya Kawahara
Academic Center for Computing and Media Studies, Kyoto University, Sakyo, Kyoto, 606-8501
e-mail: kawahara@i.kyoto-u.ac.jp

systems can provide answers to a variety of queries from users, but these systems do not conduct dialogue which involves an interaction with users, as they do not incorporate the domain knowledge and dialogue histories [3, 4]. Moreover, these systems work well only for simple keyword queries and factoid questions, but it is hard to deal with ambiguous user queries or non-factoid questions. These systems assume a clear goal of the user, a unique destination of a dialogue, and the aim of the system is to reach the goal as soon as possible.

However, when users ask something beyond the system's capability of the goal-oriented systems, current systems usually reply "I can't answer the question", or turns to the Web search and returns the retrieval list in the display. This kind of dialogue is not a natural interaction since people want to converse with them besides simple commands. A user-friendly conversational system should not reply with "I can't answer the question" even if the system cannot find the result exactly matching the user query [5]. Instead, it should present relevant information according to the user's intention and preference by using domain knowledge and dialogue management that considers the dialogue history. There are several studies towards this direction [6, 7, 8]. This kind of system is realized by information navigation that is addressed in this paper.

2 Task of Information Navigation

In human-human dialogue, people usually have topics they plan to talk about, and they progress the dialogue in accordance with the topics [9]. Dialogue participants have a role of speaker and listener, and they converse with each other by changing their role of speaker and listener. The proposed system realizes information navigation by taking a role of the speaker who provides information to the user.

An example is shown in **Figure 1**. First, the speaker offers a new topic and probes the interest of the listener. If the listener shows interest, the speaker describes details of the topic. If the listener asks a specific question, the speaker answers it. On the other hand, if the listener is not interested in the topic, the speaker avoids the details of that topic and changes the topic.

The task of information navigation is designed as non-goal-oriented dialogue according to the above-described manner. The aim of dialogue is to fulfill information demand of the user through an interaction. When the user demands are not always clear, the information navigation system clarifies the user through interactions. The system presents relevant information even if the user request is not necessarily clear and there is no exactly matching result to the user query. Moreover, the system can occasionally present potentially useful information without any explicit request by following the dialogue context.

The task design of information navigation is defined as a selection of information navigation modules. The initiative of dialogue comes and goes between the system and the user because it depends on the specification of the user demand. If the user has a clear demand, the user can ask a specific question that matches to his demand.
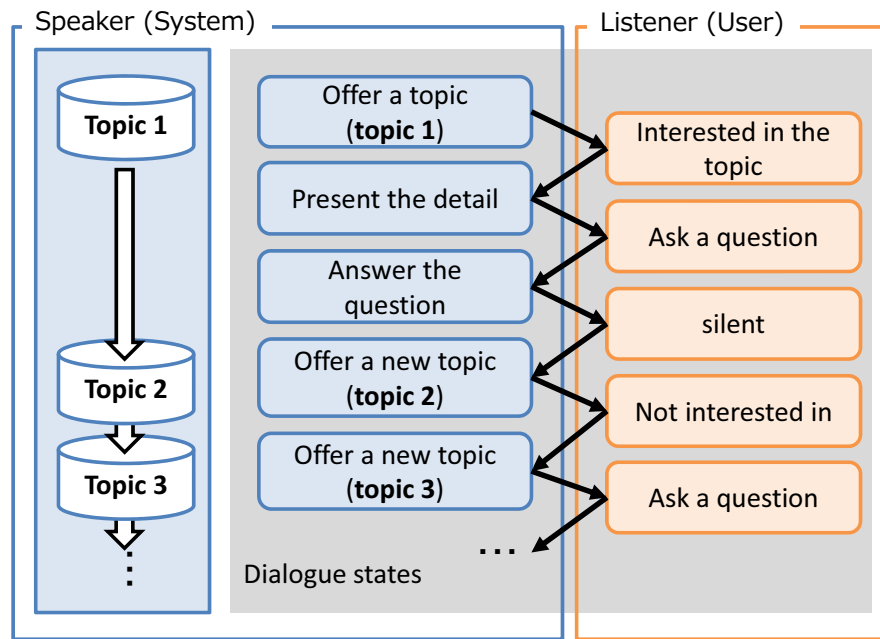
Fig. 1 *An example of information navigation in human-human conversation.*

When the user demand is not clear, the system takes an initiative to clarify the user demand by showing candidates that is related to the ambiguous query of the user. This function is achieved by modules that refer to the domain knowledge, the user intention and the user focus. Here, we define the user focus as "the main piece of information of interest to the user."

In information navigation, the system presents topics that it can talk about, describes the detail of the current topic, or presents topics related to the dialogue history when the system has an initiative. In contrast, the system answers the question of the user, replies to the information demand of the user, or receives a request of changing the topic. The functions of the system modules depends on the kind of information navigation. An example of information navigation modules is shown in **Figure 2**.

3 News Navigation System

We develop a news navigation system that realizes the information navigation described above.
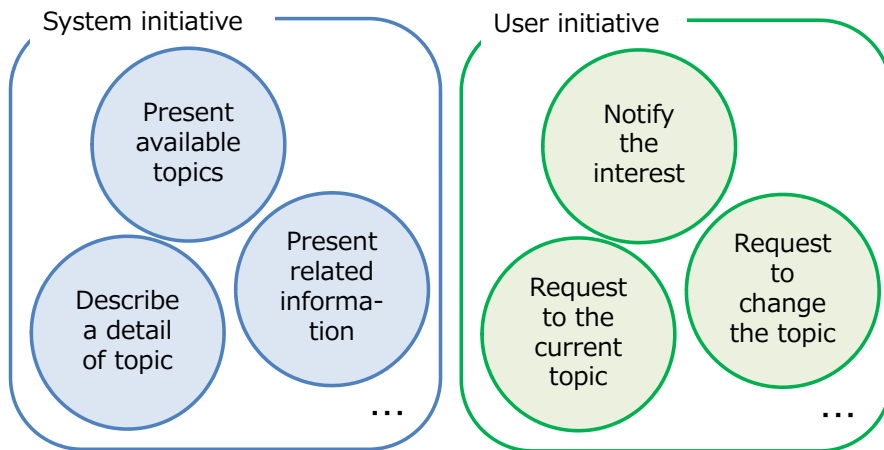
Fig. 2 *An example of information navigation modules.*

3.1 Task of News Navigation

The news navigation system assumes a large number of news articles in raw text as a back-end knowledge source. The knowledge source is limited to the news articles, but the articles are updated day by day. The system navigates this dynamic content by parsing the articles and extracting information from the huge back-end knowledge source. Moreover, it uses a tag of the domain in the news articles to extract the domain knowledge from the text source.

The news navigation system is designed based on the dialogue structure of information navigation depicted in Figure 1. The system gives a briefing on what happened on the day that is written in the articles, and the user can retrieve information through an interaction according to his interests and queries.

3.2 System Modules

An overview of the proposed system is illustrated in **Figure 3**. The system has seven modules, each of which implements a different dialogue acts. Each module takes as input a recognized user utterance, an analyzed predicate-argument (P-A) structure, and the detected user focus.

The system begins a dialogue with the "topic presentation (TP)" module, which presents a new topic selected from news articles. It chooses the next module based on the user's response. In this work, it is assumed that each news article corresponds to a single topic, and the system presents a headline of the news in the TP module. If the user shows interest (positive response) in the topic without any specific questions, the system selects the "story telling (ST)" module to give details of the news.
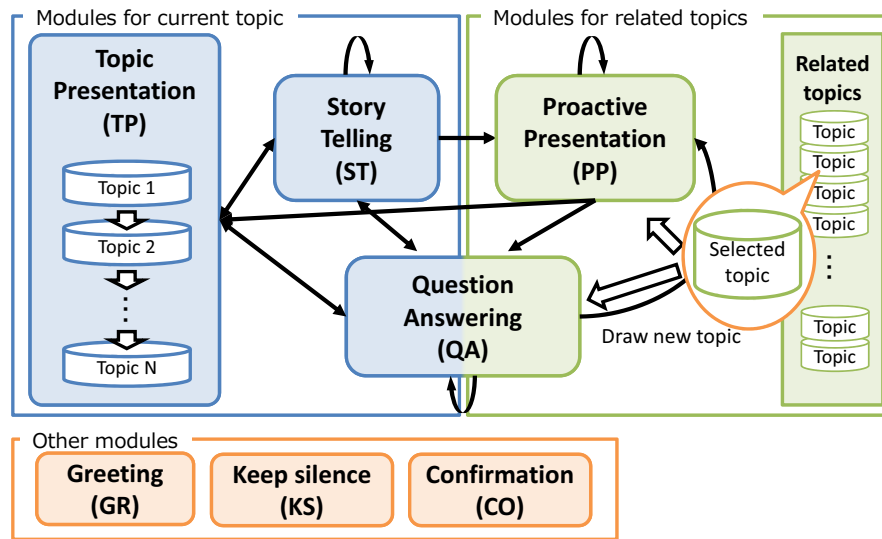
Fig. 3 *An overview of the information navigation system.*

In the ST module, the system provides a summary of the news article by using lead sentences. The system can also provide related topics with the "proactive presentation (PP)" module. This module is invoked by the system's initiative; this module is not invoked by any user request. If the user asks a specific question regarding the topic, the system switches to the "question answering (QA)" module to answer the question. This module deals with questions on the presented topic and related topics.

The modules of PP and QA are based on a dialogue framework which uses the similarity of the P-A structure between user queries and news articles, and retrieves or recommends the appropriate sentence from the news articles. This method searches for appropriate information from automatically parsed documents by referring to domain knowledge that is automatically extracted from a domain corpus [10].

Transitions between the modules are allowed as shown in Figure 3. The modules "greeting (GR)", "keep silence (KS)" and "confirmation (CO)" are also prepared. The GR module generates fixed greeting patterns by using regular expression matching. The CO module makes a confirmation if the system does not have certainty about the user query. In terms of dialogue flow, these modules can be called at any time.

The proposed scheme enables the system to answer not only clear requests, but also ambiguous requests that do not have any specified goal. The system can respond with flexible matching between the user query and the back-end knowledge source by using the statistical learning result of the semantic P-A structure [10]. As a result, the system has a capability to answer not only factoid questions, but also non-factoid

questions such as "How was today's Ichiro?" or "How do you feel about the all-star game?". By responding to these questions with some specified news such as "Ichiro hit a home-run" or "28 members are selected for the all-star game", the user can know the outline of the news that he may be interested in, and some more specific questions are invoked.

The dialogue is generated based on the news articles in the knowledge source texts. All modules of the system are automatically trained from the knowledge source, and they are easily portable to different domains.

3.3 Dialogue Control of the Proposed System

The proposed system is controlled by the dialogue management based on partially observable Markov decision process (POMDP) and conducts information navigation by selecting the most appropriate dialogue module to respond the user [11]. Markov decision processes (MDPs) and partially observable Markov decision processes (POMDPs) are the most successful and now widely used to model and train dialogue managers [12, 13, 14, 15, 16]. These approaches allow us to consider all possible future actions of a dialogue system, and thus to obtain a new optimal dialogue strategy which could not be anticipated in conventional hand-crafted dialogue systems.

The conventional scheme for goal-oriented systems assumes that the task and dialogue goal are clearly defined and readily encoded in the reinforcement learning (RL) reward function. This is not true in casual conversation or information navigation addressed in this work.

Some previous work has tackled with this problem. Pan et al. [7] designed a spoken document retrieval system whose goal is user's information need satisfaction, and defined rewards by using the structure of the target document set. This is possible only for well-defined document search problems. The strategy requires a structure of the document set and definition of user demand satisfaction. Shibata et al. [17] developed a conversational chatting system. It asks users to make evaluation at the end of each dialogue session to define rewards for reinforcement learning. Meguro et al. [18] proposed a listening dialogue system. In their work, levels of satisfaction were annotated in the logs of dialogue sessions to train a discriminative model. These approaches require costly input from users or developers, who provide evaluation and supervision labels. In the proposed dialogue management, a framework in which reward is defined for the quality of system actions and also for encouraging long interactions is explored, in contrast to the previous approaches. Moreover, user focus is tracked to make appropriate actions, which are more rewarded.

As described in Section 3.2, the task of information navigation is decoded as a module selection of seven dialogue modules: topic presentation (TP), story telling (ST), question answering (QA), proactive presentation (PP), greeting (GR), keep silence (KS), and confirmation (CO). The dialogue manager selects a module (ac-

tion decision) based on an input of a user intention. A user intention is encoded as a request to the system; the user intention has six classes and each intention has a corresponding system action.

- *TP*: request to the TP module.
- *ST*: request to the ST module.
- *QA*: request to the QA module.
- *GR*: greeting to the GR module.
- *NR*: silence longer than a threshold.
- *II*: irrelevant input due to ASR errors or noise.

Logistic regression (LR) based dialogue act tagging [19] is adopted for the user intention analysis. The existence of the user focus in the utterance is also detected by a discriminative model based on conditional random field (CRF). The system tracks the user focus to select an appropriate action module according to the user interest. The probabilities of the user intention analysis and the user focus detection are used as inputs of belief update of POMDP.

The POMDP updates its belief of the user intention by the recurrence formula,

$$b^{t+1}_{s'_j, f'_m} = \underbrace{P(o^{t+1}_s, o^{t+1}_f | s'_j, f'_m)}_{\textbf{Obs}} \sum_i \sum_l \underbrace{P(s'_j, f'_m | s_i, f_l, \hat{a}_k)}_{\textbf{Trans}} b^t_{s_i, f_l}. \tag{1}$$

Here, t is a time step and $b^t_{s_i, f_l}$ is a belief of the user intention s_i and the user focus f_l. o_s and o_f are observation results of the user intention and the user focus, and $\hat{a}_k$ is the optimal system action selected by the optimal policy function of the POMDP. The POMDP is trained by Q-learning and grid-based value iteration using a user simulator that is constructed from the annotated dialogue data [11] of news navigation.

Simplified reward for the end of each turn is defined in **Table 1** to constrain the module selection as an expected behavior. In Table 1, + is a positive reward given to appropriate actions, 0 to acceptable actions, and - is a negative reward to inappropriate actions. Here, pairs of a state and its apparently corresponding action, *TP* and TP, *ST* and ST, *QA* and QA, *GR* and GR, and *II* and KS, have positive rewards.

Other positve rewards are defined for the following reasons. If a user asks a question (*QA*) without a focus (e.g. "What happened on the game?"), the system can continue by story telling (ST). If the system cannot find an answer, it can present relevant information (PP). When the user says nothing (*NR*), the system action should be determined by considering the user focus; present a new topic if the user is not interested in the current topic (f=0), or present an article related to the dialogue history (f=1). Keeping silence (KS) is a safe action to the user silence (*NR*), thus, its reward is 0. However, we give 1 frustration point if the system selects KS in this case because the strategy conflicts with the concept of information navigation. Confirmation (CO) is a safe action to every user input, but it also frustrates the user. Thus, the reward of CO is defined as 0 for every intention, but 2 frustration points are given to the system. If the system selects an inappropriate action (action

Table 1 Rewards in each turn.

state s	focus f	action a						
		TP	ST	QA	PP	GR	KS	CO
TP	0	+	-	-	-	-	-	0
	1							
ST	0	-	+	-	0	-	-	0
	1							
QA	0	-	+	+	-	-	-	0
	1		-		+			
GR	0	-	-	-	-	+	-	0
	1							
NR	0	+	-	-	-	-	0	0
	1	-			+			
II	0	-	-	-	-	-	+	0
	1							

of $r = -10$), 2 frustration points are given to the system. If the frustration points accumulate more than 10, a large penalty is given to the system and the dialogue is terminated. A large positive reward is given if 20 turns are passed to reward a long continued dialogue.

4 Experimental Evaluation

For evaluation of the system, 626 utterances (12 users, 24 dialogues; 2 dialogues with each user) were collected with the proposed dialogue system.

For comparison, we also constructed a rule-based system (=Rule) and a POMDP-based system that does not track the user focus (=POMDP w.o. focus). We evaluated the system performance by the accuracy of action selection. The gold-standard is annotated by two annotators. The agreement for the user states was 0.958 and Cohen's kappa was 0.932. The agreement for the system actions was 0.944 and Cohen's kappa was 0.915. We reprioritized the first annotator who is familiar with the task if the annotation was not agreed.

A breakdown is shown in **Table 2**. The table shows precision (P), recall (R) and F-measure (F) of each intention tag. Here, the results of TP, ST, QA and PP are presented because the number of KS and GR was very small (#GR=2, #KS=4), and CO was not labeled as a correct action. The proposed method outperformed the compared systems for all actions. The proposed method improved the accuracy for topic presentation (TP) and proactive presentation (PP) especially when the user intention was no request (*NR*). The POMDP without the user focus always selected the keep silence (KS) module if the user said nothing (*NR*).

The proposed method also made more effective confirmations (CO) when the SLU result was not correct. It made confirmations (CO) 18 times, and 15 times of them was done when the SLU result was incorrect (15/18=83.3%). The POMDP

Table 2 Performance of action selection (precision, recall and F-measure).

tag	Rule			POMDP w.o. focus			POMDP proposed		
	P	R	F	P	R	F	P	R	F
TP	0.884	0.822	0.852	0.917	0.764	0.834	0.959	0.803	**0.874**
ST	1.000	0.022	0.043	0.900	0.500	0.643	0.910	0.789	**0.845**
QA	0.678	0.993	0.806	0.797	0.962	0.872	0.843	0.945	**0.891**
PP	0.929	0.342	0.500	0.000	0.000	0.000	0.854	0.921	**0.886**

without the user focus made only two confirmations, when the detected user intention was correct (0/2=0.0%).

The proposed method made 35 proactive presentations (PP), and 17 times of them (17/35=48.6%) invoked new user questions. This result demonstrates that the proposed system encouraged interactions in news navigation.

4.1 Discussion of Trained Policy

An example dialogue is shown in **Figure 4**. In the example, the system selects appropriate actions even if the observation likelihood is low. At the 4th turn of Dialogue 1 in this example, the system with the user focus responds with an action of proactive presentation a=PP, but the system without the user focus responds with an action of topic presentation a=TP. At the 2nd turn of Dialogue 2, the user asks a question without a focus. The confidence of s=QA is lowered by the belief update, and the system selects the story telling module a=ST. These examples show that the trained policy reflects the design of information navigation proposed in this paper. It is better to make a proactive presentation when the user is interested in the topic.

5 Conclusions

We proposed a new concept of spoken dialogue system that navigates news information. The task of information navigation is along with the direction from conventional task-oriented dialogue systems to general non-task-oriented dialogue systems. In information navigation, users are not forced to accommodate the task goal of the system. Instead, the user can make ambiguous queries. The system provides information that the user wants to know by probing and clarifying the potential demands of the user. The proposed news navigation system can converse with users in a user-friendly manner. It does not respond "I can't answer the question", or turns to the Web search even if it cannot find exact information.

The proposed system is controlled by the dialogue management based on the POMDP that tracks user focus. The system presents relevant information according to the user's interest by tracking the user focus. A POMDP framework is extended

Dialogue 1

Sys: Hello.	
Usr: Hello. Tell me today's news.	$s=TP$, $P(o\|s)=0.8$, $b(s=TP)=0.800$, $f=0$
Sys: Hanshin won the game by the game winning hit of Toritani.	$a=TP$, $b(s=TP\ f=0)=0.8$
Usr: Did Nohmi come into pitch?	$s=QA$, $P(o\|s)=0.9$, $b(s=QA)=0.991$, $f=1$
Sys: Nohmi was the winning pitcher, by shutting out for seven innings.	$a=QA$, $b(s=QA, f=1)=1.0$
Usr: Did he get strikeout?	$s=QA$, $P(o\|s)=0.9$, $b(s=QA)=0.994$, $f=1$
Sys: Nohmi struck out Abe on a slider with two out two-on jam.	$a=QA$, $b(s=QA, f=1)=1.0$
Usr: (pause)	$s=NR$, $P(o\|s)=1.0$, $b(s=NR)=1.00$, $f=1$
Sys: By the way, Nohmi struck out Sakamoto on a 147km/h four-seam in August 29th.	$a=PP$, $b(s=NR, f=1)=1.0$
Usr: Is four-seam his best pitch?	$s=QA$, $P(o\|s)=0.9$, $b(s=QA)=0.995$, $f=1$
Sys: He struck out Murata with split, his best pitch.	$a=QA$, $b(s=QA, f=1)=0.9$
Usr: OK, any other news?	$s=TP$, $P(o\|s)=0.8$, $b(s=TP)=0.984$, $f=0$
Sys: Hiroshima stopped losing streak.	$a=TP$, $b(s=TP, f=0)=1.0$
Usr: Tell me more.	$s=ST$, $P(o\|s)=0.9$, $b(s=ST)=0.871$, $f=0$
Sys: Hiroshima Carp stopped a losing streak at 3 games.	$a=ST$, $b(s=ST, f=0)=0.9$
...	

Dialogue 2

Sys: Hello.	
Usr: Hello. Tell me today's news?	$s=TP$, $P(o\|s)=0.8$, $b(s=TP)=0.800$, $f=0$
Sys: Hanshin won the game by the game winning hit of Toritani.	$a=TP$, $b(s=TP\ f=0)=0.8$
Usr: What happened?	$s=QA$, $P(o\|s)=0.8$, $b(s=QA)=0.532$, $f=0$
Sys: Hanshin-tigers won second straight against Yomiuri-giants by the game winning hit of team-leader, Toritani.	$a=ST$, $b(s=QA, f=0)=0.5$
Usr: When did he hit the game winning hit?	$s=QA$, $P(o\|s)=1.0$, $b(s=QA)=0.999$, $f=1$
...	

Fig. 4 A dialogue example.

to track the user focus to select the appropriate action module. In the experimental evaluations, the system responded with the appropriate action modules according to the user intention and the focus.

References

1. Lichan Hong, Shigeru Muraki, Arie Kaufman, Dirk Bartz, and Taosong He. Virtual voyage: Interactive navigation in the human colon. In *Proceedings of Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH 1997, pages 27–34, 1997.
2. Deborah A. Dahl, Madeleine Bates, Michael Brown, William Fisher, Kate Hunicke-Smith, David Pallett, Christine Pao, Alexander Rudnicky, and Elizabeth Shriberg. Expanding the scope of the ATIS task: the ATIS-3 corpus. In *Proceedings of the Workshop on Human Language Technology*, HLT 1994, pages 43–48, 1994.

3. Julian Kupiec. Murax: A robust linguistic approach for question answering using an on-line encyclopedia. In *Proceedings of Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR 1993, pages 181–190, 1993.
4. Robin D Burke, Kristian J Hammond, Vladimir Kulyukin, Steven L Lytinen, Noriko Tomuro, and Scott Schoenberg. Question answering from frequently asked question files: Experiences with the faq finder system. *AI Magazine*, 18(2):57–66, 1997.
5. Tatsuya Kawahara. New perspectives on spoken language understanding: Does machine need to fully understand speech? In *IEEE Workshop on Automatic Speech Recognition & Understanding*, ASRU 2009, pages 46–50, 2009.
6. Teruhisa Misu and Tatsuya Kawahara. Bayes risk-based dialogue management for document retrieval system with speech interface. *Speech Communication*, 52(1):61–71, 2010.
7. Yi-Cheng Pan, Hung yi Lee, and Lin shan Lee. Interactive spoken document retrieval with suggested key terms ranked by a Markov decision process. *IEEE Transactions on Audio, Speech, and Language Processing*, 20(2):632–645, 2012.
8. Larry Heck, Dilek Hakkani-Tur, Madhu Chinthakunta, Gokhan Tur, Rukmini Iyer, Partha Parthasarathy, Lisa Stifelman, Elizabeth Shriberg, and Ashley Fidler. Multimodal conversational search and browse. In *IEEE Workshop on Speech, Language and Audio in Multimedia*, August 2013.
9. Emanuel A. Schegloff and Harvey Sacks. Opening up closings. *Semiotica*, 8(4):289–327, 1973.
10. Koichiro Yoshino, Shinsuke Mori, and Tatsuya Kawahara. Spoken dialogue system based on information extraction using similarity of predicate argument structures. In *Proceedings of Annual SIGdial Meeting on Discourse and Dialogue*, SIGDIAL 2011, pages 59–66, Portland, Oregon, June 2011.
11. Koichiro Yoshino and Tatsuya Kawahara. Information navigation system based on POMDP that tracks user focus. In *Proceedings of Annual SIGdial Meeting on Discourse and Dialogue*, SIGDIAL 2014, pages 32–40, Philadelphia, Pennsylvania, June 2014.
12. Nicholas Roy, Joelle Pineau, and Sebastian Thrun. Spoken dialogue management using probabilistic reasoning. In *Proceedings of Annual Meeting on Association for Computational Linguistics*, ACL 2000, pages 93–100, 2000.
13. Esther Levin, Roberto Pieraccini, and Wieland Eckert. A stochastic model of human-machine interaction for learning dialog strategies. *IEEE Transactions on Speech and Audio Processing*, 8(1):11–23, 2000.
14. Jason D. Williams and Steve Young. Partially observable Markov decision processes for spoken dialog systems. *Computer Speech & Language*, 21(2):393–422, 2007.
15. Steve Young, Milica Gašić, Simon Keizer, François Mairesse, Jost Schatzmann, Blaise Thomson, and Kai Yu. The hidden information state model: A practical framework for POMDP-based spoken dialogue management. *Computer Speech & Language*, 24(2):150–174, 2010.
16. Koichiro Yoshino, Shinji Watanabe, Jonathan Le Roux, and John R. Hershey. Statistical dialogue management using intention dependency graph. In *Proceedings of International Joint Conference on Natural Language Processing*, IJCNLP 2013, pages 962–966, Nagoya, Japan, October 2013.
17. Tomohide Shibata, Yusuke Egashira, and Sadao Kurohashi. Chat-like conversational system based on selection of reply generating module with reinforcement learning. In *Proceedings of International Workshop Series on Spoken Dialog Systems*, IWSDS 2014, pages 124–129, 2014.
18. Toyomi Meguro, Ryuichiro Higashinaka, Yasuhiro Minami, and Kohji Dohsaka. Controlling listening-oriented dialogue using partially observable markov decision processes. In *Proceedings of International Conference on Computational Linguistics*, COLING 2010, pages 761–769, 2010.
19. Gokhan Tur, Umit Guz, and Dilek Hakkani-Tur. Model adaptation for dialog act tagging. In *Proceedings of IEEE workshop on Spoken Language Technology*, IWSLT 2006, pages 94–97, 2006.