Using the Wizard of Oz Method to Train Persuasive Agents

Maiko Kawasoe, Tatsuya Narita, and Yasuhiko Kitamura

School of Science and Technology, Kwansei Gakuin University
2-1 Gakuen, Sanda-shi, Hyogo 669-1337, Japan
ykitamura@kwansei.ac.jp

Abstract. Persuasive conversational agents persuade users to change their attitudes or behaviors through conversation and are expected to be applied as virtual sales-clerks in e-shopping sites. Developing such an agent requires a conversation model that identifies the most appropriate responses to the user's inputs. To create such a model, we propose the approach of combining a learning agent with the Wizard of Oz method; in this approach, a person (called the Wizard) talks to the user pretending to be the agent. The agent learns from the conversations between the Wizard and the user and constructs its own conversation model. In this approach, the Wizard has to reply to most of the user's inputs at the beginning, but the burden gradually falls because the agent learns how to reply as the conversation model grows.

Every persuasive conversation has the goal of persuading the user and ends with success or failure. We introduce a goal-oriented conversation model that can represent the success probability of persuasion and a learning method to update the model depending on the success/failure of the persuasive conversation. We introduce a learning persuasive agent that implements the conversation model and the learning method and evaluate it in the situation wherein the agent persuades users to choose one type of digital camera over another. The agent could succeed in reducing the Wizard's inputs by 48%, and, more interestingly, succeeded in persuading 2 users without any help from the Wizard.

1 Introduction

Persuasive technology draws attention as a means to create interacting computing systems that can change people's attitudes and behaviors [1]. Conversational agents will play an important role in such systems. They can interact with users through conversation [2] and are expected to become virtual sales-clerks that persuade customers to Web shopping sites [3].

Developing a conversational agent requires a conversation model that represents how the agent responds to inputs from users. It is not easy to create a conversation model in which the agent interacts well with users and a large number of conversation rules must be created by experts. To reduce the burden, we integrate a learning agent and the Wizard of Oz method [4], in which a person called the Wizard talks with a user pretending to be the agent. The agent learns from the conversations between the Wizard and the users and constructs/refines a conversation model. At the beginning, the Wizard

M. Klusch, M. Pechoucek, and A. Polleres (Eds.): CIA 2008, LNAI 5180, pp. 177–190, 2008.

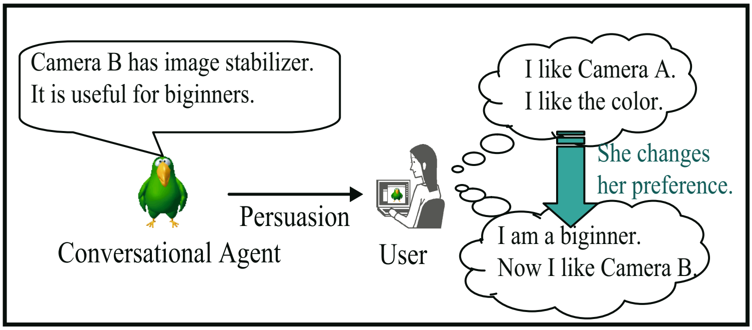

Fig. 1. Persuasion through conversation

has to input most of the replies, but gradually the agent learns to reply appropriately as the conversation model grows. When a reply made by the agent is not appropriate, the Wizard can correct it.

In this paper, we introduce a conversational agent that persuades users as shown in Fig. 1. The user initially prefers Camera A over Camera B, and the agent tries to persuade her to change her preference from A to B.

Every persuasive conversation has the goal of persuading a user and ends with success or failure. We introduce a goal-oriented conversation model that can represent the success probability of persuasion and a learning method to update the model depending on the success/failure of the persuasive conversation.

Section 2 of this paper addresses conversational agents and the Wizard of Oz method as the bases of persuasive conversational agents. In Section 3, we propose a goal-oriented conversation model and a learning method to update the model considering the success probability of persuasion. We then show a prototype system in Section 4 and evaluation results in Section 5. Finally, we conclude this paper with our future work in Section 6.

2 Persuasive Conversational Agents

2.1 Conversational Agents

Conversational agents interact with users though conversation to assist them in their information processing tasks such as information retrieval from the Web [3]. ALICE (Artificial Linguistic Internet Computer Entity) is representative of the conversational agents now available on the Web and is being used in a number of Web sites.[1]

The conversation model represents how an agent replies to inputs from users. There are two major approaches to constructing a conversation model. The first one is by describing scenarios or rules as is used in ALICE, PPP Persona [5,6], and so on. ALICE uses a language called AIML (Artificial Intelligence Markup Language), based on XML, to describe rules, each of which links a pattern, which represents an input from

[1] http://www.alicebot.org/

the user, to a template, which represents a reply from the agent. This approach forces us to write a large number of rules to make the agent reply fluently to various inputs from the user.

The second approach is to utilize a conversation corpus as is done in Command Talk [7]. In this approach, we need to establish a very large conversation corpus in advance to construct a conversation model. However, the agent cannot reply appropriately to an input if the input is not in the corpus.

2.2 Wizard of Oz Method

This paper takes the approach of integrating a learning agent and the Wizard of Oz method [8] as shown in Fig. 2. In the Wizard of Oz method, a person called the Wizard interacts with the user pretending to be the conversational agent. The Wizard can reply to input from the user when the agent cannot reply appropriately. The agent learns from the Wizard how to reply to an input by constructing a conversation model and can thereafter reply to the next instance of the same input. At the beginning, the Wizard has to reply to most of the inputs, but the burden of the Wizard falls because the agent learns to reply as the conversation model matures.

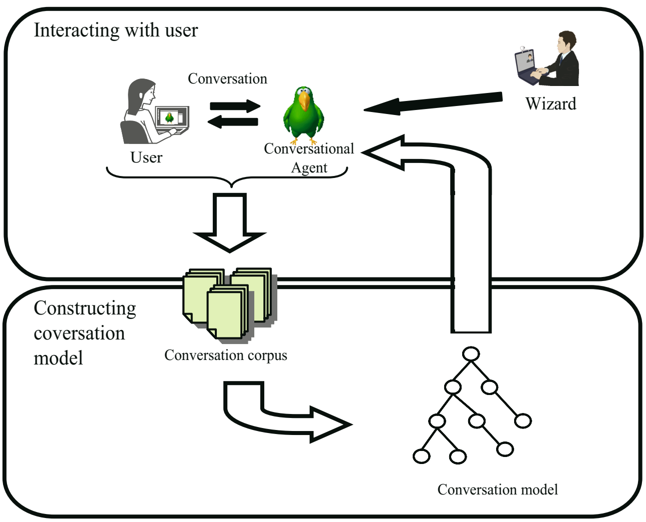

Fig. 2. Integrating a learning agent and the Wizard of Oz method

2.3 Persuasive Conversation

Persuasion is the action of changing people's attitudes and behaviors [1]. This paper considers the example of an agent that tries to persuade a user to change his/her preference from camera A to Camera B. If the user comes to prefer B, we define the persuasion as successful; otherwise, a failure.

Conventional conversational agents reply to an input from a user if the input matches a rule in the conversation model. When it matches multiple rules, one of them is selected. The selection process depends on the system and/or the applied domain. Persuasive agents, on the other hand, should select the rule that is more likely to lead to success. To this end, we propose a goal-oriented conversation model that considers the success probability of persuasion and a learning method to update the probability as derived from persuasive conversations between the Wizard and users.

3 Learning Persuasive Agents

To build persuasive agents that can learn, we need a goal-oriented conversation model and a learning method that can update the conversation model. Details of the model and the learning method are given below.

3.1 Goal-Oriented Conversation Model

The conversation model can be represented as a state transition tree where a statement is represented as a link to change a state from one to another as shown in Fig. 3. In this example, the agent tries to persuade a user to be a member of Kitamura laboratory. There are two types of states; user states, which represent the user talking and agent states, which represent the agent talking. They are interleaved on the conversation path. A conversation path represents the flow of conversation between the agent and one or more users and begins with the initial states and terminates with either success or failure. Each state is assigned a success probability score.

The agent decides how to respond to an input from the user following the conversation path held by the model. If the input matches a statement on a link to an agent state,

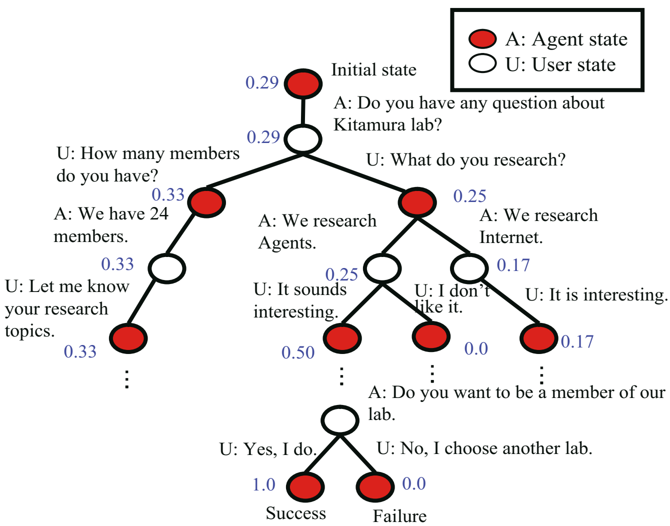

Fig. 3. Conversation model

it chooses a statement that links the agent state to the user state with greatest success probability.

For example in Fig. 3, the agent says "Do you have any question about Kitamura lab?" at the beginning. If the user asks "How many members do you have?" the agent replies "We have 24 members," following the stored conversation path. If the user asks "What do you research?" there are two reply candidates. The agent chooses the reply "We research Agents." because that link leads to a user state with higher success probability (0.25).

3.2 Updating Conversation Model

When an input from the user does not match any statement on the stored conversation path, the conversation path is branched and the success probability scores are updated depending on persuasion success/failure as shown in Fig. 4. If the persuasion succeeds (fails), 1.0 (0) is assigned to the terminal state. The success probability score of each state except terminal states in the conversation model is updated as below.

– Agent state s

$$Q(s) \leftarrow \max_{t \in succ(s)} Q(t)$$

– User state s

$$Q(s) \leftarrow \frac{1}{|succ(s)|} \sum_{t \in succ(s)} Q(t)$$

$succ(s)$ is a set of child states of s. At an agent state, the agent can choose what to say, so the success probability is set to be the maximum one among child user states. On the other hand, at a user node, the user chooses what to say, so the success probability is set to be the average one among child agent states. We here assume that the user takes a neutral attitude toward the agent. If we assume the user takes a negative attitude, the success probability should be the minimum one.

For example, when an agent says "We research Agents." using the conversation model shown in Fig. 3, if the user replies "What are Agents?" which is not contained in the model, a new conversation path is created by branching as shown in Fig. 4 following a persuasive conversation between the Wizard and the user. The persuasion succeeds, so 1.0 is attached to the terminal state of the branched path and each state on the conversation path is updated as mentioned above.

3.3 Reducing Redundancy in the Conversation Model

Because our conversation model is a tree, two conversation paths that virtually identical are treated as completely different if the first parts of the paths are different. Naive extension of the conversation model creates redundancy.

To reduce this redundancy, we transform the conversation model by using multiple phases as shown in Fig. 5. In this example, two conversation paths are different even though only the first parts are different as shown in Fig. 5 (a). We transform the model into one with two phases; greeting phase and persuasion phase, as shown in Fig. 5 (b) to reduce the redundancy.

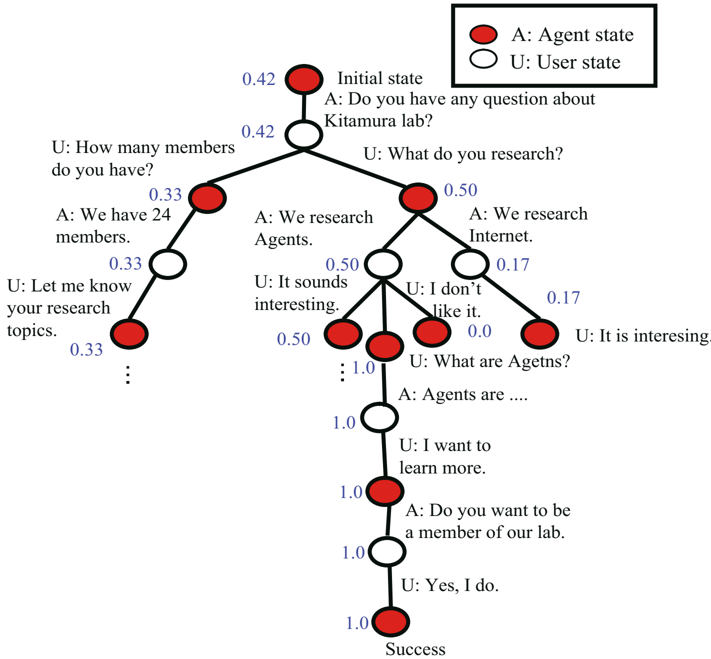

Fig. 4. Updated conversation model

4 Implementing a Persuasive Conversational Agent

We implemented a persuasive conversational agent as shown in Fig. 6 to chat with a user. Messages from the agent appear in the top panel and the user inputs messages to the agent in the message box at the bottom as shown in Fig. 7. When the agent receives a message from the user, it generates responses from the conversation model. There are two types of responses.

Context sensitive responses (CSR) are generated by the model by following a conversation path from the initial state. For example in Fig. 3, if the user inputs "What do you research?" in response to the message "Do you have any question about Kitamura lab?" from the agent in the initial state, the agent generates two context sensitive responses "We research Agents," and "We research Internet." Their success probabilities are 0.25 and 0.17, respectively.

Context free responses (CFR) are generated by the model by direct matching of the input from the user without following any conversation path. For example in Fig. 3, after the interchange of; "Do you have any question about Kitamura lab?", "How many members do you have?", and "We have 24 members," if the user asks "What do you research?" the agent generates two context free responses "We research Agents." and "We research Internet." In this case, because the responses do not follow any conversation path, success probabilities are not attached to the responses.

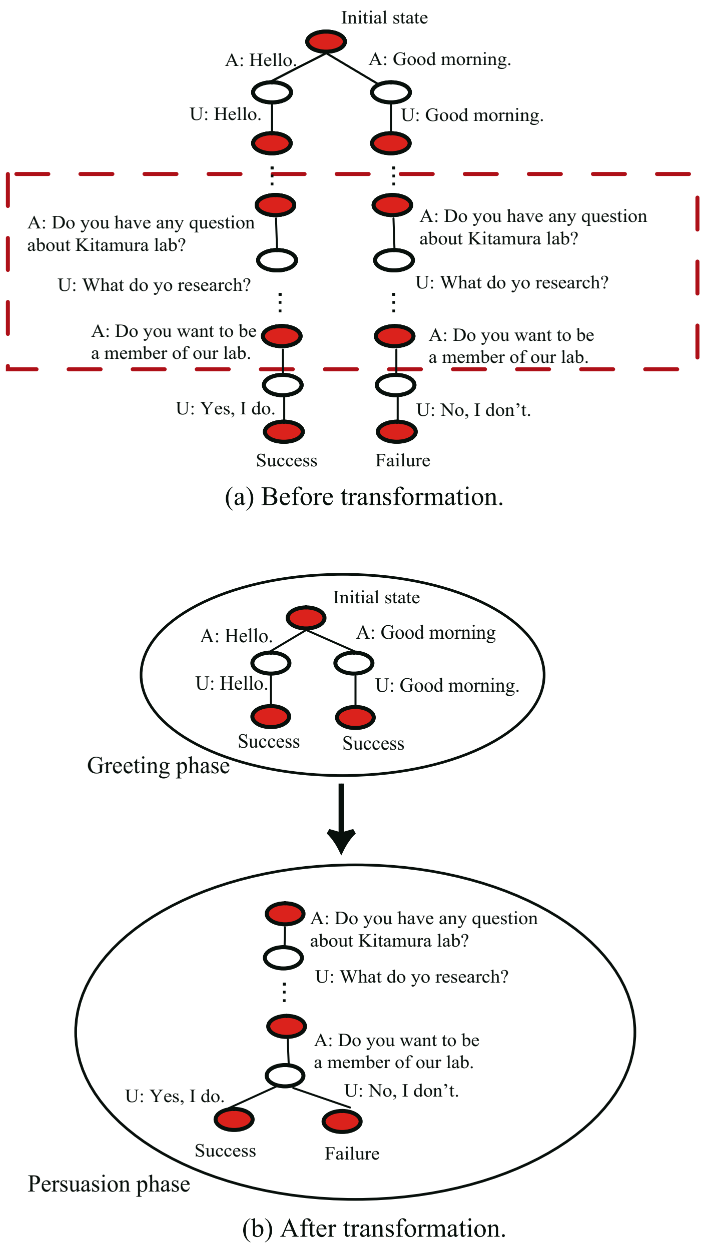

Fig. 5. Transforming conversation model

In the early stages of learning, the agent often fails to respond to the user if it uses only CSRs because the conversation model is small. By utilizing CFRs, the agent can generate more candidates.

Messages from the user appear on the top panel of the Wizard chat client and responses generated by the agent appear as in a pull-down menu as shown in Fig. 8. The Wizard can choose the most appropriate one from among them. If the Wizard does not like any response, he/she can input a new message directly into the message box.

When a persuasion succeeds/fails or a phase terminates, the Wizard notifies the state to the system through the pull-down menu in the left-bottom and a textbox to specify the next phase as shown in Fig. 9. A button on the right-bottom is used to show the conversation model, an example of which is shown in Fig. 10. When a link is clicked, the corresponding message appears in the window.

5 Evaluation

We evaluated our persuasive conversational agent from two viewpoints; (1) the input cost of Wizard when utilizing responses created by the agent, and (2) the persuasiveness of the conversation model constructed through conversations between users and the Wizard.

We performed two experiments in which the persuasive conversational agent tried to guide users to choose one of two digital cameras using the following procedure.

1. Each participant read the specifications of two digital cameras, A and B, as shown in Table 1. Camera A has better features than camera B, but the price of A is more than that of B.
2. The participant chooses one as his/her favorite from the first impression.
3. The agent, with help from the Wizard, tries to persuade the participant to choose the other one. It first asks why he/she chose the one selected, and then it tries to refute the reasoning. It asks him/her for situations in which he/she would use the camera and his/her taste, and recommends the camera that he/she did not initially choose.

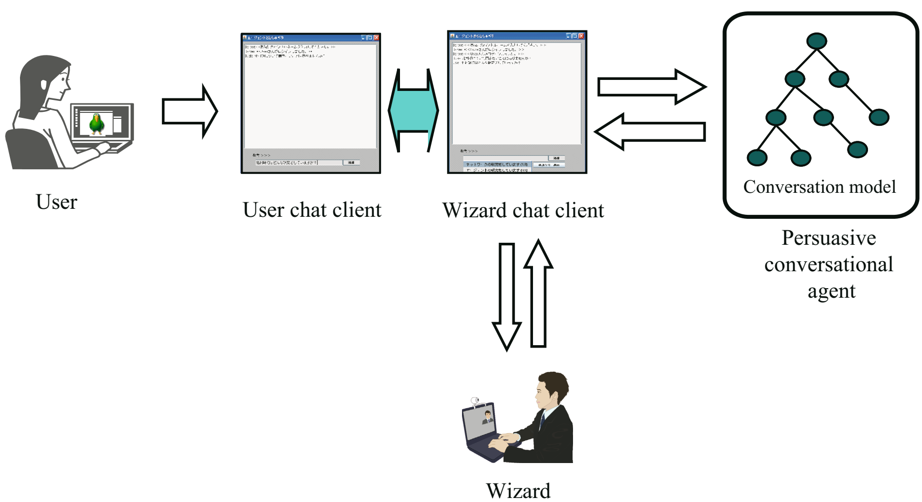

Fig. 6. System overview

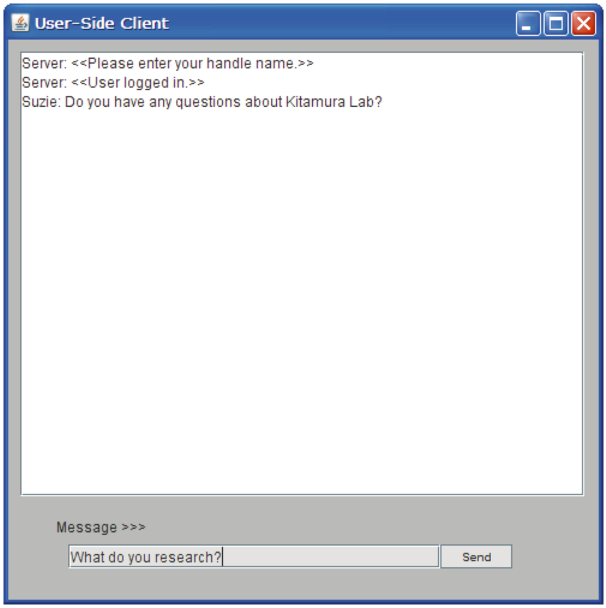

Fig. 7. User chat client

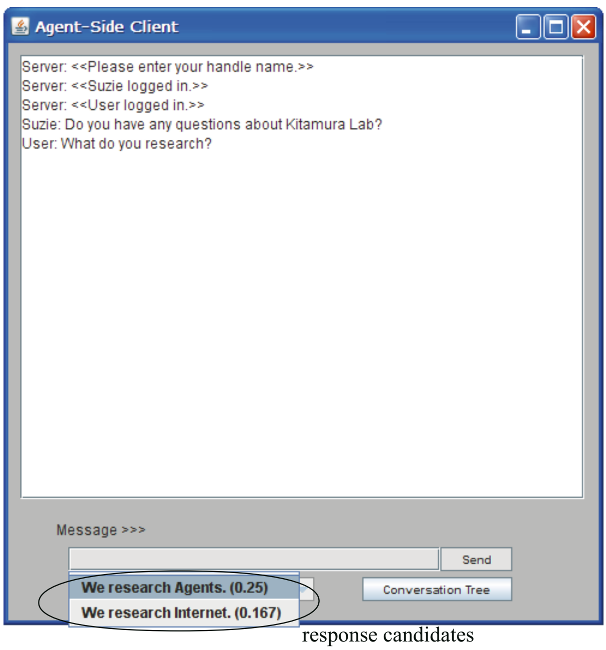

Fig. 8. Wizard chat client: choosing a response

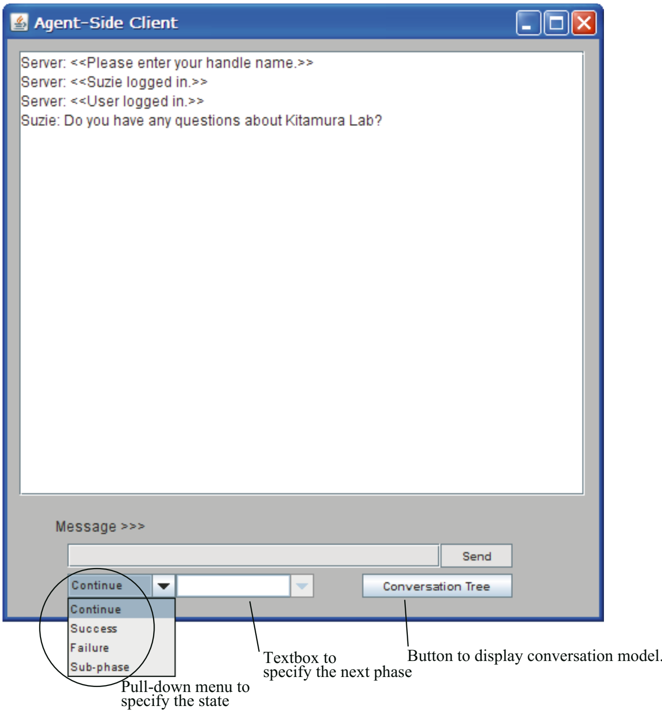

Fig. 9. Wizard chat client: specifying the state transition

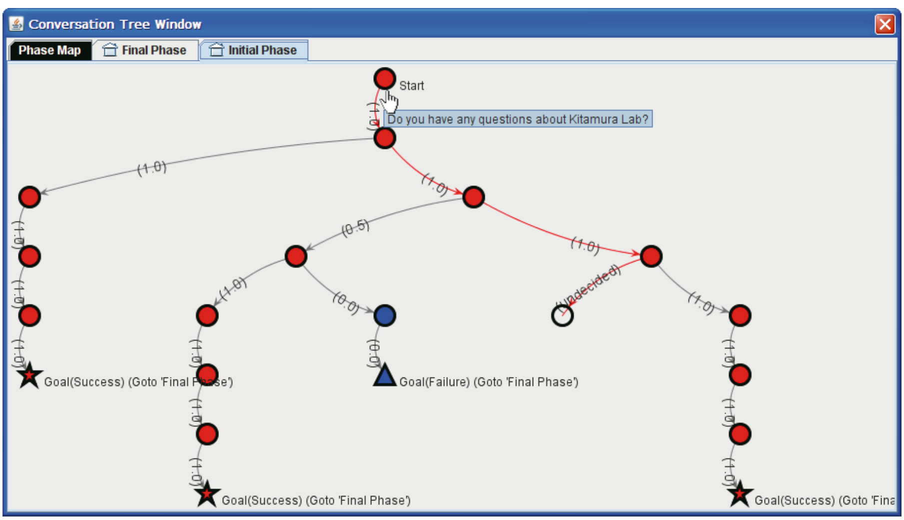

Fig. 10. Displaying a conversation model

Table 1. Specifications of digital cameras A and B

	A	B
Price	¥35,000	¥29,800
Resolution	10M pixels	7M pixels
Weight	154g	131g
Image stabilizer	Yes	No

Table 2. Experiment 1: Result of persuasion

Initial choice	Final choice	Number of participants	Success/Failure
A	A	22(54%)	Failure
	B	19(46%)	Success
B	A	6(30%)	Success
	B	13(70%)	Failure

4. The participant is then asked which camera he/she now prefers. The persuasion succeeds (fails) if he/she changes (does not change) his/her opinion.

5.1 Experiment 1: Input Cost of Wizard

In this experiment, we constructed a conversation model by collecting persuasive conversations between an agent and 60 university students (48 male, 12 female, the average age is 20.9) using the Wizard of Oz method. The results are shown in Table 2. In total, we succeeded in persuading 25 (42%) of the 60 participants.

Figure 11 shows the number of responses made by the agent to each participant. The responses selected by the Wizard are categorized into 4 groups. “CSR (best)” is the context sensitive response with the highest success probability generated by the agent. “CSR (2nd or worse)” covers the context sensitive responses that had 2nd or lower probability generated by the agent. “CFR” covers context free responses. “Wizard input” the responses input by the Wizard. At first, the Wizard has to input most of the responses, but gradually this number falls and the number of responses made by the agent increases. Overall, for the 60 persuasive conversations, the Wizard accepted 602 (48%) of the agent’s 1245 responses as appropriate, so this means that the input cost of the Wizard was reduced.

In a remarkable occurrence, the agent succeeded in persuading one participant (no.51) without any input from the Wizard.

5.2 Experiment 2: Persuasiveness of Conversation Model

To determine the persuasiveness of the conversation model described in the previous section, we performed another experiment to persuade 10 students. In this experiment, the agent was limited to one response at each turn. Only when the agent returned no

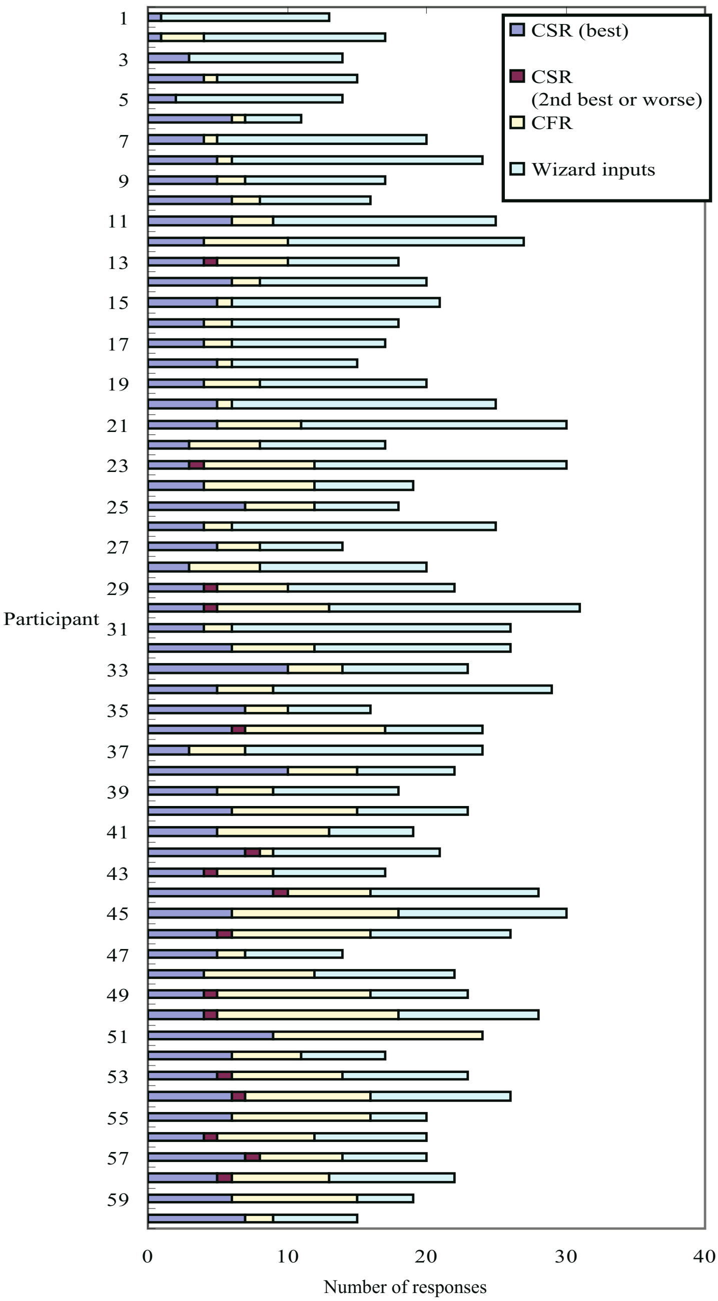

Fig. 11. Experiment 1: Categorized responses

Table 3. Experiment 2: Result of persuasion

Initial choice	Final choice	Number of participants	Success/Failure
A	A	4(57%)	Failure
	B	3(43%)	Success
B	A	0(0%)	Success
	B	3(100%)	Failure

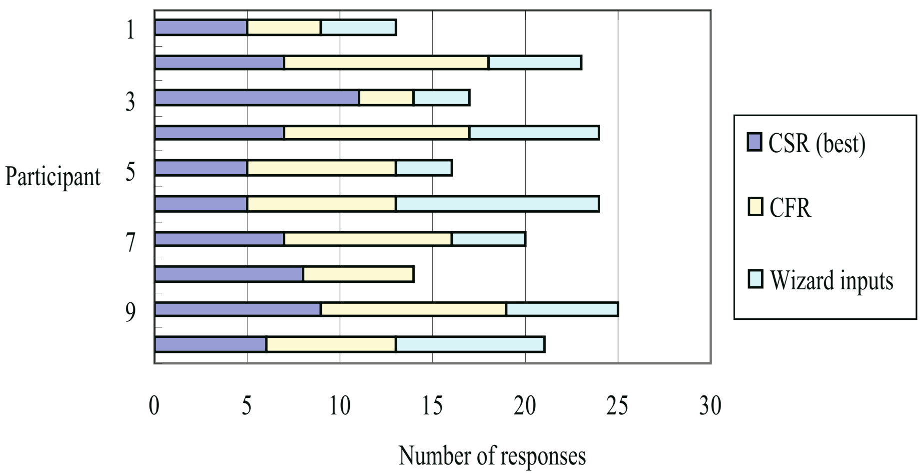

Fig. 12. Experiment 2: Categorized responses

Table 4. The number of generated states of each phase

Phase	Number Of states
Greeting	35
Initial choice	17
A: persuasion	1296
B: persuasion	491
	1839

response, the Wizard input a response as before. The result of this experiment is shown in Table 3. In total, the agent succeeded in persuading 3 (30%) out of the 10 participants.

The responses made by the agent are categorized in Fig. 12. The agent succeeded in persuading one participant (no.8) without any input from the Wizard.

The conversation model created from the two experiments consists of 1839 states as shown in Table 4.

6 Conclusion

Persuasive conversational agents are expected to be virtual shopping clerks on e-shopping sites. To create such agents, we need to create a conversation model that specifies how to reply to inputs from users. In this paper, we proposed an approach to create a conversation model by integrating a learning agent and the Wizard of Oz method. We evaluated the performance of the proposed persuasive conversational agent in a situation where it persuaded users to choose one digital camera over another.

In the 1st experiment with 60 subjects, we could reduce the number of Wizard inputs by 48%, and in the 2nd experiment, the agent that used the conversation model created in the first experiment succeeded in persuading one user (out of 10) without any input from the Wizard; another 2 subjects were persuaded with some assistance by the Wizard.

At present, our persuasive agent requires a lot of assistance from the Wizard to persuade human users. We need to work on to improve the ability of persuasion toward an agent that requires no assistance. In future work, we will improve the success ratio of persuasion. To this end, we need to collect more conversations to create a better conversation model that replies to a larger number of inputs from users. Further work is needed on reducing model redundancy by using natural language processing techniques to handle synonymous sentences.

Another future task is to increase the maintainability of the conversation model. At present, it is not easy to modify the conversation model. We need to develop a GUI for this and to visualize the persuasion strategies contained in the model.

Acknowledgment

This work is partly supported by the Grant-in-Aide for Scientific Research (No.17300050) from Japan Society for the Promotion of Science.

References

1. Fogg, B.J.: Persuasive Technology. Morgan Kaufmann, San Francisco (2003)
2. Cassell, J., et al.: Embodied Conversational Agents. MIT Press, Cambridge (2000)
3. Prendinger, H., Ishizuka, M. (eds.): Life-like Characters. Springer, Heidelberg (2004)
4. Fraser, N.M., Gilbert, G.N.: Simulating Speech Systems. Computer Speech and Language 5(1), 81–99 (1991)
5. Andre, E., Rist, T., Muller, J.: Integrating Reactive and Scripted Behaviours in a Life-Like Presentation Agent. In: Proceedings of the Second International Conference on Autonomous Agent, pp. 261–268 (1998)
6. Andre, E., Rist, T., Muller, J.: WebPersona: A Life-Like Presentation Agent for the World-Wide Web Knowledge-Based Systems 11(1), 25–36 (1998)
7. Stent, A., Dowding, J., Gawron, J.M., Bratt, E.O., Moore, R.: The CommandTalk spoken dialogue system. In: Proc. ACL 1999, pp. 183–190 (1999)
8. Okamoto, M., Yeonsoo, Y., Ishida, T.: Wizard of Oz Method for Learning Dialogue Agents, Cooperative Information Agents V, LNAI 2182. In: Klusch, M., Zambonelli, F. (eds.) CIA 2001. LNCS (LNAI), vol. 2182, pp. 20–25. Springer, Heidelberg (2001)