A Culturally-Situated Agent to Support Intercultural Collaboration

Victoria Abou Khalil[1(✉)], Toru Ishida[1], Masayuki Otani[2], and Donghui Lin[1]

[1] Department of Social Informatics, Kyoto University, Kyoto, Japan
v.aboukhalil@gmail.com, ishida@i.kyoto-u.ac.jp
[2] Kinki University, Higashi-osaka, Japan
m-otani@i.kyoto-u.ac.jp

Abstract. While traveling, foreign visitors encounter new products that they need to understand. One solutionis by making Culturally Situated Associations (CSA) i.e. relating the products they encounter to products in their own culture. We propose the design of a system that provides tourists with CSA to help them understand foreign products. In order to provide tourists with CSA that they can understand, we must gather information about their culture, provide them with the CSA, and make sure they understand it. To deliver CSA to foreign visitors, two types of data are needed: data about the products, their associated properties and relationships, and data about the tourist cultural attributes such as country, region, language. The properties and relationships about countries, regions and products, can be extracted from open linked data on the web, and CSA can then be constructed. However, information about the tourist's cultural attributes and the knowledge they can relate to is unavailable. One way to tackle this problem would be to extract the tourist's cultural attributes that are needed in each situation through dialogue systems. In this case, a Culturally Situated Dialogue (CSD) must take place. To implement the dialogue, dialogue systems must follow a machine-learned dialogue strategy as previous work has shown that a machine-learned dialogue strategy outperform the handcrafted dialogue approach. We propose the design of a system that uses a reinforcement learning algorithm to learn CSD strategies that can support individual foreign tourists. Since no previous system providing CSA has been implemented, the system allows the creation of CSD strategies when no initial data or prototype exists. The method is used to generate 3 different agents that learn 3 different dialogue strategies.

Keywords: Automatic dialogue strategies · Reinforcement learning · Culturally situated associations · Wizard of Oz

1 Introduction

Japan, rich in both traditional culture and technical innovation, attracts people from all around the world and is a popular destination for tourists. Every year, tens of millions of visitors are walking in Zen gardens, shopping for strange gadgets, and experimenting with Japanese cuisine. However, the first complaint from foreign tourists is the paucity of foreign language services [1]. Some reported concerns with communication

T. Yoshino et al. (Eds.): CollabTech 2017, LNCS 10397, pp. 130–144, 2017.
DOI: 10.1007/978-3-319-63088-5_12

difficulties while shopping, particularly for food. This highlights a problem in intercultural collaboration. This area of research is becoming essential in a world that is losing its physical borders and in which people and cultures are more and more on the move and in contact [2].

To help tourists understand the food products they are about to buy, one possible solution is to display, in the tourist's language, a complete listing of the ingredients as well as a description of the food product. However, this kind of information might leave them with questions like: What does it taste like? What is the texture? How do we cook it?

In situations where providing a simple description of a product fails to deliver a complete understanding of the product, an efficient alternative is to relate the product to a similar product in the tourist's culture. This would mean offering Culturally Situated Associations (CSA) that allow foreign visitors to understand the usage, and taste of the food product they are inquiring about.

To be able to provide tourists with CSA, two types of data are needed: data about food products, their associated properties and relationships as well as data about the tourist's cultural attributes like country, region, and language. With today's available technologies and the prevalence of data on the web, we are able to offer solutions that use cultural associations to explain concepts and products. With the increased popularity and presence of open data on the web, we are able to query relationships and properties about products. Properties and relationships about countries, regions and food products can be found, and cultural associations can then be extracted. However, information about tourists' cultural attributes and the knowledge they can relate to is unavailable. One way to tackle this problem would be to extract the tourists' cultural attributes that are needed in each situation through dialogue systems.

A dialogue system that supports foreign tourists with CSA must deliver the associations and make sure that those associations were understood. The first requirement can be fulfilled by developing Culturally Situated Dialogue (CSD) strategies that support the realization of those objectives. However, when no initial observations or system exists, learning a dialogue strategy is a challenging task as developers or designers of the system may not be able to predict the most appropriate action to be taken by the system at each moment. Developers and designers would have to undertake the time consuming process of predicting what would be the most appropriate action in each situation. Moreover, a dialogue system is likely to need a considerable number of different utterances and previous work showed that automatic dialogue strategies outperform handcrafted dialogue strategies [3].

In conjunction with the demand for CSA and the challenge of automating CSD strategy generation where no initial system exists, reinforcement learning algorithms are needed to learn CSD strategies to support foreign tourists when no data or working prototype exists. To model the possible state spaces of the reinforcement learning algorithm, we first identified common dialogue patterns that take place between tourists and shop owners in Nishiki Market. Then, we extracted the attributes related to the tourists' culture as well as food properties that interest the tourists. By breaking down the extracted attributes into more fine grained attributes we created three attribute sets with different levels of granularity. Each of these three attribute sets was mapped into a different state space, resulting in the creation of three different agents.

2 Background

2.1 Culturally Situated Associations

A variety of intercultural communication models have been proposed by researchers. However, the most influential model is attributed to Byram because his approach provides holistic intercultural competence and has defined objectives and practical derivations [4].

Byram's model defines the five skills needed in order to accomplish successful intercultural communication: *intercultural attitudes, knowledge, interpreting and relating, discovery and interaction* as well as *critical cultural awareness* [5]. Two of those skills are necessary in the initial stages of becoming familiar with a new culture and are essential to understand foreign concepts or products [5]:

- *Discovery* or *knowledge*: knowledge about a social group and their products and practices in the foreign visitor's own country.
- *Interpreting and relating*: foreign visitors relate the information they get to information from their own culture.

Byram defines the skills of discovery as "the ability to recognize significant phenomena in a foreign environment and to elicit their meanings and connotations, and their relationship to other phenomena [4]". Those skills are of particular importance in contexts where the foreign visitor has very little information about the foreign culture and its related concepts or products, as it is the case of tourists in Japan. The skill of interpreting and relating consists of putting concepts or products from two or more cultures side by side and seeing how each might look from the other perspective [4]. However, in real life situations, interpreting and relating cannot be achieved in real time by tourists or shop clerks as CSA requires deep knowledge about the foreign culture. Automatic dialogue systems might be useful in this situation as they allow the identification of the tourist's culture and the retrieval of the needed association.

2.2 Linked Data

The term Linked Data was created in 2007 to describe a set of best practices for publishing and connecting structured data on the web. The data can be linked to form relationships and becomes more useful with the use of semantic queries. Linked Data allows the connection and query of data from different sources [6].

One of the main projects associated with the use of Linked Data has been the Linking Open Data project; it allows anyone to participate by publishing a dataset following the Linked Data recommendations and linking them with existing datasets. DBpedia is one of the biggest existing datasets. The DBpedia dataset contains data extracted from Wikipedia and consists of 3.4 million concepts described by 11 billion triples. As the information contained in DBpedia results from a crowd sourcing process and is extracted from unstructured and semi-structured information, there are still many problems with the dataset. The error rate in the DBpedia Dataset is 11.93%, which is considered moderate [7]. Previous studies explored the possibility of using Linked Data in combination with dialogue systems [8].

2.3 Automatic Dialogue Strategies

The recent literature shows a growing interest in the implementation and use of automatic dialogue systems. The development of such dialogue systems, and more particularly the development of dialogue strategies is challenging [9]. In order to achieve an application goal in an efficient way through a series of interactions with the user, dialogue strategies are needed. By quantifying the incremental achievements made as well as the efficiency, is it possible to describe the system as a stochastic model that can be used for learning those dialogue strategies. This method has many advantages including the possibility of automating the evaluation of the dialogue strategies as well as an automatic design and adaptation [10]. Moreover, this approach naturally utilizes large amounts of data [11].

Previous works on dialogue systems used reinforcement learning in order to learn Wizard of Oz' (WoZ) dialogue strategies of presenting information and replicating them. WoZ allows the learning of dialogue strategies when no initial system exists. The results showed that reinforcement learning combined with WoZ experiments allows the development of optimal strategies when no working prototype is available [12]. However, unlike standard dialogue systems that take into account user-related properties, the challenge in learning optimal CSD strategies consist of learning which information about the tourist's culture, if any, should be acquired and in which order.

3 Overview of the System

Figure 1 shows the system architecture. The WoZ experiment is used as no working prototype or initial CSD system is available. The tourist and the wizard communicate through Skype to allow the wizard to see the product the tourist is asking about. In order to provide the wizard with the optimal dialogue strategy, an agent is trained based on a reinforcement learning algorithm, and passes the optimal strategy to take at each step to the wizard. The wizard first reports its state of knowledge to the agent through a web interface (e.g.: I don't have any information yet). Once the agent receives the current state of knowledge of the system, it provides the wizard with the appropriate action to take (e.g.: Ask for the tourist's country). If the agent issues a query of the associated concept, the wizard retrieves the CSA from DBpedia. The dialogue, directed by the agent, and executed by the wizard is iterated until the CSA is provided to the tourist and understood. In practice, the dialogue would be translated into the tourist's language using Language Grid, a service-oriented collective intelligence that allows users to create language services from existing language resources [13].

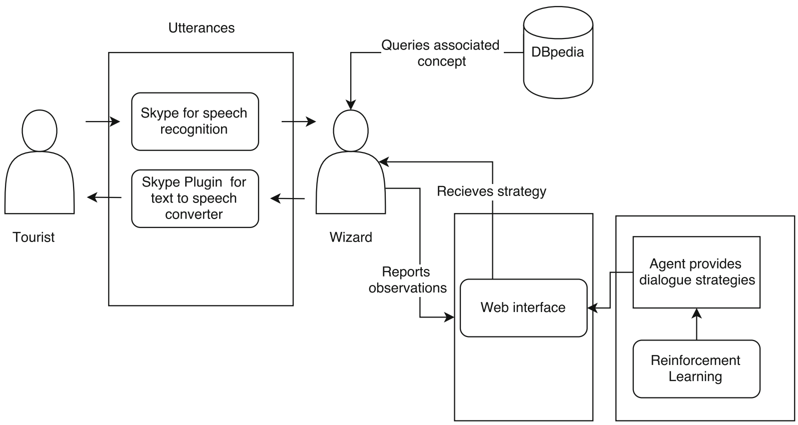

Fig. 1. System architecture

4 Extraction of Dialogue Patterns

In order to extract the necessary components needed to build the feature space of the reinforcement algorithm and create the automatic dialogue strategies, we first identify common natural dialogue patterns that should provide CSA to tourists.

To identify the possible dialogue patterns, we first conducted interviews with tourists in Nishiki Market, a traditional food market in Kyoto. We interviewed 15 tourists coming from western countries, chosen randomly during their visit to the market. The breakdown of genders was balanced and the participants were from Europe, New Zealand and U.S.A. The tourists were asked to list the questions that they would have wanted to ask if it was possible to get an answer. We received 34 questions from the participants. Similar questions were put together and the tourists' questions were categorized by question topic.

The questions of the tourists were classified into three categories shown in Table 1. The first category contains questions about the ingredients of a particular food. The second category covers questions about usage. The last category includes general questions about the composition and usage of the food.

Table 1. Categorization of questions asked by tourists by question topic

Category	Associated tourists' questions
Ingredients	What does it taste like?
	Is it suitable for vegetarians?
	Can I take it through customs?
Usage	How is it used?
	How do we eat it?
Ingredients and usage	What is this?
	What is the difference between X and Y?

Based on the questions provided by the tourists, we developed "typical" dialogues that could happen between the shop owners and the tourists. During those conversations, shop owners follow a CSD strategy to answer the questions of the tourists with CSA. We match each of the previous examples to a particular CSD pattern. To understand CSD, we define several terms as follows:

- *Target concept:* is the concept that needs to be explained.
- *Associated concept:* is used to explain a target concept. It is a concept that belongs to a different culture from the target concept.
- *Common attribute:* is an attribute or a property that belongs to both the target and the associated concepts.

Cultural attribute, such as a *location*, *language*, etc., is a common attribute which contributes to identifying a culture.

Using the previous terms, we classified culturally situated conversations into several CSD patterns, see the examples below.

Example conversation 1:

Tourist: "What is this and what does it taste like?"
Shop Owner: "It is Neri Goma. It is a paste made out of roasted sesame seeds. Where are you from?"
Tourist: "Iraq"
Shop Owner: "It is like Tahine."

Dialogue Pattern 1: Using cultural attribute as a pivot

Tourist: Question about the taste of the *target concept.*
Shop Owner: Question to identify the *cultural attributes* of the tourist.
Tourist: Tourist provides the *cultural attributes.*
Shop Owner: Finds the *associated concept* that possess *cultural attributes* that match the tourist's cultural attributes and *common attributes* related to the taste that are identical to the common attributes of the *target concept*.

Example conversation 2:

Tourist: "What is this? How do we use it?"
Shop Owner: "It is Neri Goma. It is a paste made out of roasted sesame seeds. Where are you from?"
Tourist: "Iraq."
Shop Owner: "It is like Tahine, but in Japan it is mainly used in sweets."

Dialogue Pattern 2: Comparative association

Tourist: Question about a *target concept.*
Shop Owner: Question to identify the *cultural attributes* of the tourist.
Tourist: Tourist provides the *cultural attributes.*
Shop Owner: Finds the *associated concept* that possess *cultural attributes that* match the tourist's cultural attributes and *common attributes* related to the taste that are identical to the common attributes of the *target concept*. If other common attributes differ

from the *target concept*'s *common attributes,* the differences are presented to the tourist.

Example conversation 3:

Tourist: "What is this?"
Shop Owner: "It is Udon, noodles made out of wheat and flour. They are usually served in a broth."
Tourist: "What is the difference from Soba?"
Shop Owner: "Udon is made out of wheat and Soba out of buckwheat Where are you from?"
Tourist: "Italy"
Shop Owner: "Udon is more like Spaghetti and Soba like Pizzoccheri"

Dialogue Pattern 3: Intra-Cultural Comparison

Tourist: Question about the difference between two *target concepts.*
Shop Owner: Question to identify the *cultural attributes* of the tourist.
Tourist: Tourist provides the *cultural attributes.*
Shop Owner: The difference between the two *target concepts* is identified by comparing all their common attributes. Based on the cultural attributes of the tourist, two *associated concepts* with the same difference to the common attributes are found.

Based on the previous dialogue patterns, we extract the components essential to conduct different types of CSD strategies:

- Target Concept
- Associated concept
- Cultural Attributes
- Common Attributes

5 Extraction of Culturally Situated Associations

In order to provide tourists with CSA about food products and answer their questions we need access to an adequate data source. DBpedia contains 3.4 million concepts including concepts about food and dishes, and also provides relationships between foods, countries, and usages.

5.1 Mapping the Common Attributes to DBpedia Properties

Based on the categorization of the questions of the foreign visitors and on the previous dialogue patterns, we map the questions asked by the tourists to the DBpedia properties that we need to answer those questions. Every set of questions could be answered by comparing the common attributes of a particular product from a foreign culture to the common attributes of another product from the tourist's culture. The DBpedia properties shown in Table 2 are the ones used to extract and compare the common attributes.

Table 2. Mapping of tourists questions to DBpedia properties needed to answer them

DBpedia properties	Tourists' questions
dbo:ingredients,	What does it taste like?
dbp:type,	Is it suitable for vegetarians?
dbp:similarDish	Can I take it through customs?
dbo:course, dbp:served,	How is it used?
dbp:similarDish	How do we eat it?
dbo:ingredients, dbp:type, dbp:similarDish, dbo:course, dbp:served, dbp:similarDish	What is this?
	What is the difference between X and Y?

5.2 Mapping the Cultural Attributes to DBpedia Properties

The cultural attributes allow the culture of the tourists to be identified in order to provide them with an associated concept from their own culture. As each country has different food, relating the target concept to an associated concept from the tourist's country is essential in realizing CSA. However, in several cases, particularly in big countries, different regions have different food products. Providing the tourists with an associated concept from their own region will either permit more precise information and/or allow them to understand the food better. We mapped the cultural attributes to existing DBpedia properties below:

Dbo:country: Country to which the associated concept belongs, usually country of the tourist
Dbp:region: Region to which the associated concept belongs, usually region of the tourist

5.3 Extraction of Culturally Situated Associations

We draw the knowledge representations to visualize the relationships, based on DBpedia's existing properties and relationships. The example below shows the knowledge representations for dialogue type 1. Figure 2 represents the products' similarity that answers the first type of conversation. In this case, the tourist from Italy asks 'What does Soba taste like?'. To answer this question, Soba is queried, as well as all the products that have the same ingredients and originate from Italy. The query gives us *Pizoccheri*. The red parts show the relationships and concepts that are taken into consideration to find the CSA for this particular question.

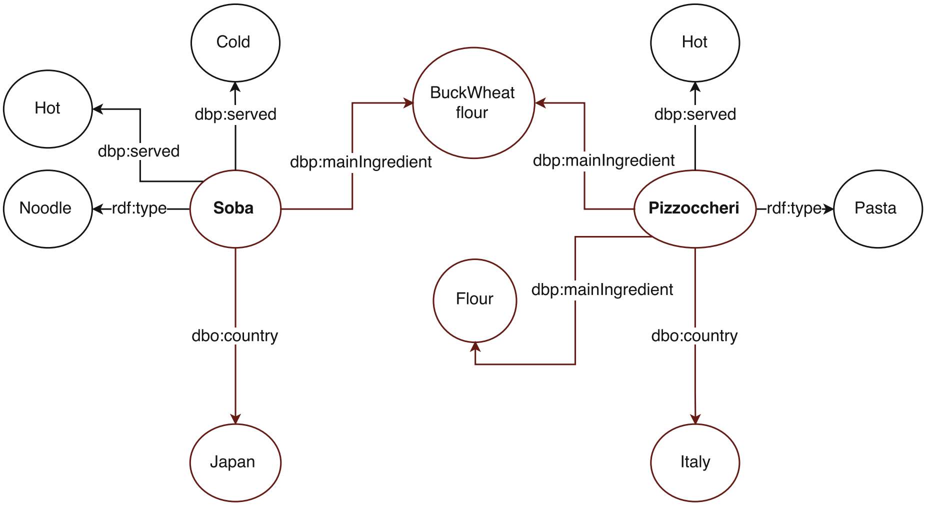

Fig. 2. Knowledge representation for conversation type 1 (Color figure online)

6 Reinforcement Learning Algorithm

6.1 Reinforcement Learning Algorithm

The Markov Decision process is a mathematical formalism that is used to implement the reinforcement learning algorithm. Reinforcement learning is the problem faced by an agent that must learn a behavior through trial-and-error interactions with a dynamic environment [14]. The main components of this formalism and their implementation are explained below. The algorithm was implemented using Python, following Nathan Epstein's implementation[1].

The State Space and Action Space

The state space represents the complete list of states that the system can be in. The action space is the all-inclusive list of actions that can be taken in the environment. The states and actions are usually set a priori.

The Transition Probabilities

The probabilities of transitioning between state *s* to state *s'* given action *a* taken are estimated from observed data. The estimated transition probability is computed as follows:

$$P(s, a, s') = \frac{Number\ times\ we\ transitioned\ from\ state s\ to\ state\ s'\ given\ action\ a}{Number\ of\ times\ we\ took\ action\ a\ in\ state\ s}$$

[1] https://github.com/NathanEpstein/reinforce/tree/664949173dfaabcc359f46c4f4c640fd577682b4.

In the case that action *a* was never taken in state *s*, we set the value of *P(s, a, s′)* to *1/number of states* in order to avoid the ratio of 0/0. In this situation we assume that the probability is equality distributed over all states. With an incremented number of observations, the state transition probability will be updated to become more precise.

The Reward

Our algorithm assumes that the reward is unknown. We can also compute the expected immediate reward in a specific state, as the average reward observed in state *s*. The goal is to find a policy that will produce the biggest possible cumulative reward.

Value Iteration and Policy

A policy is any function that maps states to actions. The value function for policy π is the expected sum of discounted rewards when we start in state s and take actions according to π. The value function of policy π is given by Bellman's equation.

$$V^{\pi}(s) = R(s) + \gamma \sum_{s' \in S} P_{s\pi(s)}(s')\, V^{\pi}(s')$$

Bellman's equation states that the expected sum of discounted rewards $V^{\pi}(s)$ consists of the sum of the immediate reward as the expected sum of future rewards. We define as well the optimal value function given by:

$$V^{*} = \max V^{\pi}(s)$$

$V^{\pi}(s)$ is the best expected sum of discounted rewards that can be reached using any policy. Based on the equations above, we will describe the algorithm that we use to calculate the value function and to get the best policy.

- For each state s, initialize $V(s) = 0$
- Repeat until convergence:
  - For each state, update:

$$V^{\pi}(s) = R(s) + \gamma * \max(a \in \mathrm{A}) \sum\nolimits_{s' \in S} P_{s\pi(s)}(s')\, V(s')$$

  - Policy in state s is the $a \in A$ that maximizes *V(s)*

In this algorithm we update the estimated value function based on Bellman's equation. For every state s, we calculate the new value of *V(s)*. After a certain number of iterations, the value is supposed to converge on *V*(s)*. We then choose the optimal policy that always maximizes *V(s)*.

The State Action Space

The states of the reinforcement learning algorithm amounts to all the states that the system (the wizard in our current system) possesses about internal and external resources that it is interacting with (e.g. input from the tourist, associated concepts). The action set of the dialogue system includes all possible actions it can accomplish. It includes the interactions with the user (e.g. asking the tourist for input, providing the tourist with output) as well as the interactions with other resources (e.g.: searching for the associated concepts). When the system's current state is s and action a is taken, the state changes to s'. For example, when the system is in an initial state and the wizard doesn't not have any information, the agent will ask the wizard to interact with the tourist and obtain specific information. The next state, s', will depend on whether the wizard obtained the information or not.

We identify the possible state spaces based on the components extracted from the dialogue patterns. The target concept is assumed to be known as the wizard would be interacting with the tourist and would be able to identify it. The cultural attributes are necessary in order to determine the culture of the tourist, and thus, in which culture the associated concepts should be found. Tourists usually have a question that is related to a particular common attribute (e.g.: usage, ingredients). The common attributes are necessary as they will be the basis of the comparison between the target concept and the associated concept. The action space is directly extracted from the state.

Based on the previously defined components, we create three levels of state spaces with different granularity in terms of the minimum number of observations to learn the dialogue strategies. In this work the minimum number of observations was calculated taking into consideration the case where every state is visited by every action. The three different agents are named: Novice agent, Intermediate agent and Advanced agent.

The Novice agent needs few observations to cover all the actions that could be taken from every state, however, it is expected to produce low quality dialogue strategies. The Intermediate agent needs more observations than the Novice agent to cover all the actions that could be taken from every state, but is expected to produce better quality dialogue strategies. The Advanced agent needs the most observations to cover all the actions that could be taken from every state, but is expected to produce the best dialogue strategies. Figure 3 plots the minimum number of observations versus the number of states.

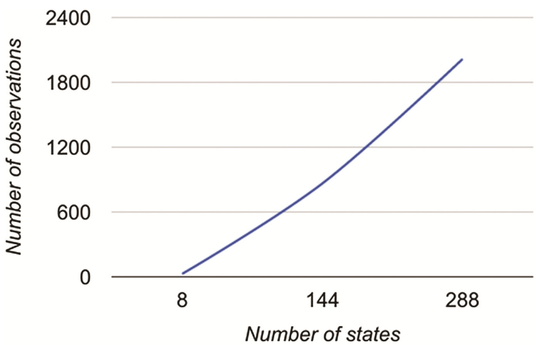

Fig. 3. Minimum number of observations needed versus number of states

6.2 The Novice Feature Space

The first level feature space produces the *Novice agent.* This State space includes only 3 entries that represent the mental state of the system, in other terms, the current state of the wizard.

- Doesn't Know the user's culture/Knows the user's culture
- Doesn't know the associated concept/Knows the associated concept
- Knows that the user doesn't understand the concept/Knows that the user understands the concept.

Every entry can take one of two values, giving us a total number of 8 states(includes 2 final states):

- Knows the user culture
- Knows the associated concept
- Knows that the user understands the concept

and

- Knows the associated concept
- Knows that the user understands the concept

For the first level feature space, the action space includes only three actions:

- Identify the user's culture
- Identify the associated concept
- Confirm that the user understood the concept

6.3 The Intermediate State Action Space

The second level State Action space produces the *Intermediate agent.* The second level state space is the result of breaking down the first level state space into more precise states of knowledge. It includes 6 entries that represent the mental state of the system.

- Doesn't know the user's country/Knows the user's country
- Doesn't know the user's region/Knows the user's region
- Doesn't know the common attributes/Knows the common attributes
- Doesn't know if there is an international associated concept/Knows that there is an international associated concept/Knows that there is not an international associated concept
- Doesn't know the cultural associated concept/Knows the cultural associated concept
- Doesn't know if the tourist understood the associated concept/Knows that the tourist understood the associated concept/Knows that the tourist didn't understand the associated concept

Every entry can take one of two values; with all permutations we get a total of 144 states, including 15 final states. To be in a final state, the agent should know the associated concept and should know that the user understood the associated concept. Moreover, the knowledge of the system should be consistent (E.g.: the system knows the cultural associated concept but doesn't know either of the cultural attributes, is not a final state). For the second state space, the action set includes seven actions:

- Identify the user's country
- Identify the user's region
- Identify the common attributes
- Identify if there is an international associated concept
- Identify if there is a cultural associated concept
- Confirm that the user understood the concept

6.4 The Advanced State Action Space

The third level State Action space produces the *Advanced agent.* The third level state space is the result of breaking down the second level state space in more precise states of knowledge. It includes 7 entries that represent the mental state of the system:

- Doesn't know the user's country/Knows the user's country
- Doesn't know the user's region/Knows the user's region
- Doesn't know the common attributes/Knows the common attributes
- Doesn't know if there is an international associated concept/Knows that there is an international associated concept/Knows that there is not an international associated concept
- Doesn't know the country associated concept/Knows the country associated concept
- Doesn't know the region associated concept/Knows the region associated concept
- Doesn't know if the tourist understood the associated concept/Knows that the tourist understood the associated concept/Knows that the tourist didn't understand the associated concept

Every entry can take one of two values; with all permutations we get a total number of 288 states, including 17 final states. To be in a final state, the agent should know the associated concept and should know that the user understood the associated concept. Moreover, the knowledge of the system should be consistent (e.g.: The system knows the cultural associated concept but doesn't know either of the cultural attributes is not a final state). For the third level state space, the action set include seven actions:

- Identify the user's country
- Identify the user's region
- Identify the common attributes
- Identify if there is an international associated concept
- Identify if there is a country associated concept
- Identify if there is a region associated concept
- Ask if the user understood the concept

7 Conclusion

This paper proposed a system that uses a reinforcement learning algorithm to learn Culturally Situated Dialogue strategies to support foreign tourists. Since no previous system has been implemented, the method allows the creation of dialogue strategies when no initial data or prototype exists.

As a first step, and in order to model the possible state spaces of the reinforcement learning algorithm, we identified common dialogue patterns that take place between tourists and shop owners in a market in Kyoto and extracted the attributes needed to implement Culturally Situated Dialogues. By breaking down the extracted attributes into more finely grained attributes we created three attribute sets with different levels of granularity. Each of these three attribute sets was mapped into a different state space, resulting in the creation of three different agents: The *Novice agent*, the *Intermediate agent* and the *Advanced agent*. Each agent needs a different minimum number observations and produces a different dialogue strategy. In order to provide the system with consistent data, we gathered open linked data concepts from DBpedia, after mapping the attributes with DBpedia properties.

Future work includes exploring the possibilities of automating the process of migrating to more complex agents depending on the available number of observations at each moment. This would allow the application of this technology to a variety of situations where Culturally Situated Associations are needed and no initial system or little observations exist.

Acknowledgments. This research was partially supported by a Grant-in-Aid for Scientific Research (A) (17H00759, 2017-2020) from Japan Society for the promotion of Science (JSPS), and the Leading Graduates Schools Program, 'Collaborative Graduate Program in Design', by the Ministry of Education, Culture, Sports, Science and Technology, Japan.

References

1. https://www.jnto.go.jp/eng/ttp/sta/PDF/E2016.pdf
2. Ishida, T.: Intercultural collaboration and support systems: a brief history. In: Baldoni, M., Chopra, Amit K., Son, T.C., Hirayama, K., Torroni, P. (eds.) PRIMA 2016. LNCS, vol. 9862, pp. 3–19. Springer, Cham (2016). doi:10.1007/978-3-319-44832-9_1
3. Scheffler, K., Young, S.: Automatic learning of dialogue strategy using dialogue simulation and reinforcement learning. In: Proceedings of the Second International Conference on Human Language Technology Research, pp. 12–19. Morgan Kaufmann Publishers Inc. (2002)
4. Yang, S., Chen, J.: Fostering foreign language learning through technology enhanced intercultural projects. Language Learn. Technol. **18**(1), 57–75 (2014)
5. Byram, M.: Teaching and assessing intercultural communicative competence, p. 34. Multilingual Matters, Clevedon (1997)
6. Berners-Lee, T.: Linked Data - Design Issues (2006). http://www.w3.org/DesignIssues/LinkedData.html. Accessed 23 July
7. Zaveri, A., et al.: User-driven quality evaluation of dbpedia. In: Proceedings of the 9th International Conference on Semantic Systems, pp. 97–104. ACM (2013)

8. Sonntag, D., Kiesel, M.: Linked data integration for semantic dialogue and backend access. In: AAAI Spring Symposium: Linked Data Meets Artificial Intelligence (2010)
9. Eckert, W., Levin, E., Pieraccini, R.: User modeling for spoken dialogue system evaluation. In: Proceedings of the IEEE Workshop on Automatic Speech Recognition and Understanding, pp. 80–87. IEEE (1997)
10. Levin, E., Pieraccini, R., Eckert, W.: Using Markov decision process for learning dialogue strategies. In: Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing, vol. 1, pp. 201–204. IEEE (1998)
11. Levin, E., Pieraccini, R., Eckert, W.: A stochastic model of human-machine interaction for learning dialog strategies. IEEE Trans. Speech Audio Process. **8**(1), 11–23 (2000)
12. Rieser, V., Lemon, O.: Learning Effective Multimodal Dialogue Strategies from Wizard-of-Oz Data: Bootstrapping and Evaluation, pp. 638–646. ACL (2008)
13. Ishida, T. (ed.): The Language Grid: Service-Oriented Collective Intelligence for Language Resource Interoperability. Springer, Heidelberg (2011)
14. Schmidhuber, J.: A general method for multi-agent reinforcement learning in unrestricted environments. In: Adaptation, Coevolution and Learning in Multiagent Systems, Papers from the 1996 AAAI Spring Symposium, pp. 84–87 (1996)