Towards Learning Personalized Semantic Relevance Paths In Dialogue Systems

Panos Alexopoulos
iSOCO
Av. del Partenon, 16-18, 1-7, 28042, Madrid, Spain
Email: palexopoulos@isoco.com

Manolis Wallace
Department of Computer Science and Technology
University of Peloponnese
End of Karaiskaki St., 22100, Tripolis, Greece
Email: wallace@uop.gr

***Abstract*—When interacting with information access systems, users have typically distinct interests, preferences and goals that may significantly influence the way they judge the relevance of these systems' outputs. User modeling aims at capturing such interests and preferences in the form of personalized user profiles which could be then used by the systems to provide tailored information services to their users. In this context, we present in this paper a novel ontology-based user profile for personalized information access which, different from the majority of existing approaches, captures user preferences in the form of semantic relevance paths within the ontology graph. More importantly, we present an automated approach for learning and maintaining such profiles in dialogue systems by taking advantage of the dialogue-based interaction between the system and the user and eliciting from the latter important feedback regarding the relevance of the provided information.**

I. INTRODUCTION

User modeling and personalization in information retrieval is a major research area that aims at enhancing the retrieval process by complementing explicit user requests with implicit user interests, preferences and goals. The latter are typically used to better interpret a given user request and return results that are more accurate, more complete and, consequently, more likely to satisfy his/her information need [1] [2] [3].

The main issues to be addressed in personalization approaches are three, namely the representation, the acquisition and the exploitation of user profiles. Many user profile representation approaches utilize bags of terms and keywords as a way to represent user preferences [1]. More sophisticated approaches use more formal knowledge structures, such as taxonomies or ontologies for that task and relate users to ontological entities (concepts or individuals) which they are important for them in the given application scenario [2] [4] [5] [3] [6]. This kind of user profile representation has more advantages than the term-based one as i) it is more accurate and ii) the reasoning capabilities of the ontology may be exploited during the retrieval process.

As an enhancement to the above, a relatively few number of approaches propose to include in the user profile not only ontological concepts but also relations [4] [7]. This is an important addition for two reasons. First, it is often the case that a user is interested in a particular entity only when the latter is somehow related to specific other entities. For example, a user might not be generally interested in historical personalities but only to those that are related to the specific historical event for which he looks information. In such a case, user preference representation solely by means of entities is not enough as, in practice, the user is not really interested in historical events or historical personalities but practically in the relation between them. That means that he wants each time he makes a query about an event and only then to get as results also personalities that are related to this event.

The second reason why relations need to be included in user profiles is that many approaches that utilize entity-based user profiles assume that the user is also interested in all the entities that are semantically related to those included in his/her profile. However, it is often the case that not all of the ontology's relations are (equally) important for the user. For instance, in the above history-related example, the user might be interested in the relation between events and personalities but not in the one between events and locations. Again, an entity-based user profile is inadequate for capturing this type of preference.

Nevertheless, the way ontological relations are typically considered in the relevant approaches has the following limitations:

1) Most approaches relate users to single ontological relations according to some interest score. However, the same relation might link multiple different concepts to each other, yet not having the same importance for all of them. For example, the relation between a person and the place he/she was born in might be interesting to a user only when the person is an actor. This practically means that merely including in the user profile the relation *"isBornIn"* (which relates any type of person to his/her birthplace) will fail to accurately reflect the user's preference.
2) The user may be interested to particular combinations (or chains) of relations without necessarily being interested in each of them separately. For example, a user might be interested in persons who where born in a location where an event took place. This means that the relations *"isBornInLocation"* and *"tookPlaceInLocation"* need to be included in the user's profile,

yet not individually but as an ordered pair. This is something that current approaches do not support.

3) The same relation might have a different interest score in a different situation context. For example, if a thesis writing student searches in the university's electronic library catalogue for information about some historical event (e.g. the battle of Waterloo), he/she will probably be also interested in books about relevant entities (e.g. a biography of Napoleon). On the other hand, if the same student is in an electronic bookstore seeking to buy a book about the same battle, it is more likely that he/she will prefer a book about the battle rather than its participants.

Our goal in this paper is to address the above limitations in the context of (limited-domain) dialogue systems that are used for interactive information search by human users. Such systems comprise an increasingly popular interaction paradigm that aims to facilitate effective human-machine communication by means of natural conversations and are found in a large variety of applications, including intelligent assistants, call centers, navigation systems and human-robot interaction [8]. Personalization and user modeling are key research topics in this area and a variety of relevant approaches have been proposed [9] [10]. Yet, to the best of our knowledge, most of these approaches follow the entity-based user modeling paradigm described above, thus not taking full advantage of the relational information that a domain ontology may offer.

In this context, our proposed contribution in this paper is twofold. On the one hand we propose a novel type of ontology-based user profile which captures user preferences in the form of **semantic relevance paths**. Informally, a semantic relevance path can be defined as a path connecting two concepts in the ontology graph by means of a sequence of (an arbitrary number of) intermediate relations and concepts. As we will show in the rest of the paper, relating users to such paths (instead of merely to entities or relations) results into more comprehensive user profiles and enables systems to provide better responses to user requests.

The second part of our contribution is an automated approach for learning and maintaining such profiles in dialogue systems by taking advantage of the dialogue-based interaction between the system and the user. The proposed method involves eliciting from the users important feedback regarding the relevance of the provided information and using it to learn and maintain their profiles.

The structure of the rest of the paper is as follows. In the next section we describe related work in the area of personalized information access both within the context of dialogue systems and outside of it. In section III we describe in detail our proposed framework and its components and in section IV we summarize the key aspects of our work and we discuss the potential directions it could take in the future.

II. Related Work

Information search personalization approaches are typically analyzed in two dimensions: i) the way they model and represent user preferences, interests and goals and ii) the process they use to acquire and populate these user models.

A. User Model Representation

In the first dimension, a variety of user models have been proposed in the literature. The majority of these models have the form of vectors of weighted terms [1] [11], vectors of weighted ontological concepts and entities [5] [3] and/or instances of predefined ontologies [12] [2] [6] [13]. For example, in [5] a profile is represented as a vector of weights denoting the degree to which a concept is related to the user's current activities (tasks, goals, short-term needs). Similarly, in [2] the user profile is represented as a subgraph of the ODP ontology[1], consisting of the concepts the user has used in a specific search session.

Inclusion of ontological relations in the user profile is found in [4] and [7]. In the first work, the notion of a user ontology is defined as a specialization of a domain ontology in which both concepts and relations are assigned a specific value for indicating a user's interests. In the second approach, an ontology design pattern based on fuzzy description logics is employed in order to represent the relative importance of ontological relations in different contexts.

Compared to the above, the main differentiating characteristic of our proposed personalization framework, in terms of user model representation, is the assignment of user interest weights to semantic relevance paths rather than merely concepts or relations.

B. User Model Acquisition

As defining user models in a manual fashion is a non-tractable task, personalization approaches typically define, along with their user models, some process to acquire and maintain these models in an automated way. In most of the cases, this happens by considering previous interactions between the user and the system and analyzing relevant data [1] [5] [2] [6].

For example, in [5] a user profile is built as a cumulative combination of the concepts involved in successive user requests, in such a way that the importance of concepts fades away with time. In [3] the users usage history is the result of a number of different user actions (querying, browsing within results, feedback provision etc.) and it is used to mine from its containing data concept-based user preferences by means of clustering algorithms. Usage history is also utilized in [2] and [6] with the first work applying a spreading activation algorithm to build and update the user model and the second doing the same thing by means of a linear combination formula.

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

The user model acquisition approach we propose in this paper is different from the above in that it utilizes user-system dialogues to elicit preference-related user information. We use the term "elicit" to denote the fact that the system in our approach does not merely extract concepts or relations from the dialogue history but it also uses particular dialogue strategies to influence the contents of this history so as to be more related to the user preferences. Similar interactive user model elicitation and acquisition approaches are found in a number of works [14] [15] [16], though the type of user model that we seek to elicit in this paper, has not been considered, to the best of our knowledge, in any of these.

III. Proposed Framework

To enable the automatic acquisition of personalized semantic relevance paths in a dialogue system we work as follows. First, we create and initialize one or more sets of relevance paths. Second, we use these patterns within the dialogue-based system and we collect from the users related feedback through an appropriate dialogue strategy. Third, we use the feedback to update the patterns and learn new ones. To that end, our proposed framework comprises the following components:

- **A Semantic Relevance Path Model** which defines semantic relevance paths and relates them to users and contexts.
- **A Dialogue-Based Feedback Elicitation Mechanism** which elicits related feedback from the users.
- **A Path Adaptation Mechanism** which uses the gathered feedback to adapt the existing paths and discover new ones for each user.

In the subsequent sections we elaborate on each of these components.

A. Semantic Relevance Path Model

1) Model Representation: For the purposes of this paper we define an ontology as a tuple $O = \{C, R, I, i_C, i_R\}$ where:

- C is a set of concepts.
- I is a set of instances.
- R is a set of binary relations that may link pairs of concept instances.
- i_C is a concept instantiation function $C \rightarrow I$.
- i_R is a relation instantiation function $R \rightarrow I \times I$.

We also define a semantic relevance path as an ordered set $p = \{c_1, r_1, ..., r_l, c_{l+1}\}$ where $c_i \in C$ is a concept, $r_i \in R$ is a relation and l is the path length. The meaning of such a path is that if a query asks for information relevant to the concept c_1 then information relevant to concept c_l should be returned as well.

As already suggested in the introduction, different users may assign different degrees of importance to different paths in different contexts. Therefore we model the set of user relevance path preferences as a function $Pref : U \times Ctx \times P \rightarrow [0, 1]$, where U is the set of all users, Ctx is the set of all possible contexts and P is the set of all possible semantic relevance paths. If $u \in U$, $ctx \in Ctx$ and $p \in P$ then $Pref(u, ctx, p)$ is the degree to which the semantic relevance path p should be used by the system for query expansion purposes for the user u and in the context ctx. Based on this formalization, our goal is practically to learn the function $Pref$ for all users and contexts.

2) Model Initialization: Model initialization involves the definition of an initial $Pref$ function that is to be used by the dialogue system in the beginning of its interaction with the users. What this function might look like depends on the application domain and on any a priori knowledge that is available or can be derived regarding expected user relevance paths in this domain. In most cases, we expect to be able to define for all the users some initial paths for some initial contexts. Then, after the system is deployed and used, these paths should be (automatically) tailored to each individual user and new ones should be discovered.

B. Feedback Elicitation Mechanism

1) Required Feedback: The feedback we want to elicit from the user consists practically of semantic relevance path assertions, namely statements expressing how much does the user consider a given relevance path to be useful, important or applicable for his/her purposes. For example, the statement *"Yes, I am very interested in events that led to this battle"* suggests that the path relating battles to events that led to them is important for the user. On the other hand, the statement *"No, I don't want the key personalities of this event"* suggests that the system should not utilize the path(s) between events and related persons when answering the user's request.

More formally, we model the relevant feedback assertions as a mapping function $A : U \times Ctx \times P \rightarrow L$, where L is the set of linguistic expressions that the users may use to express the importance of the path for them. This function practically returns all the linguistic expressions that a given user has uttered for a given path $p \in P$ in a given context. Table I shows an example of such a function for the library assistant scenario.

2) Dialogue Game for Feedback Elicitation: To derive the above assertion sets from the users we define a dialogue game, namely a set of rules and constraints that determines in general the way the system engages into dialogues with the users and help it decide how to respond to specific user utterances. In our scenario, a dialogue typically starts with the user making a query to the system involving some ontology entity (e.g. *"I would like books about the Battle of Marathon"*). If the related domain/scenario has known contexts, then the one meant in the dialogue might be already contained in the request (e.g. *"I would like books about*

Table I
EXAMPLE DESIRED FEEDBACK FOR THE LIBRARY ASSISTANT SCENARIO

User	Context	Path	Importance
John	Thesis Writing	{HistoricalEvent, *hasParticipant*, HistoricalPerson }	High
John	Thesis Writing	{HistoricalEvent, *isPartOf*, HistoricalEvent }	Low
John	Thesis Writing	{Location, *hasEvent*, HistoricalEvent }	High
John	Thesis Writing	{Location, *isNearTo*, Location, *hasEvent*, HistoricalEvent }	Low

the Battle of Marathon, I'm writing my thesis"). If it's not contained, then it can be explicitly enquired by the system (e.g. *"What is the reason you want this information for?"*) with the user possibly suggesting a previously unknown context which is added to the relevance model.

Given a query the system has two options; either it will use its initialized relevance model to provide a response to the user or, before that, it will proactively ask the user what other related entities he/she might be also interested in (e.g. *"Are you also interested in personalities that participated in this war?"*). This latter option enables the system to acquire information about relevance paths that are not defined in the initialized relevance model.

In any case, since it's possible that the user disagrees with the system's relevance model (e.g. that historical events should not be expanded to their locations) it is important to capture this disagreement and use it to adapt the model. To do that, the system provides, along with its response, an explanation of why this result was retrieved, practically exposing the path utilized for the relevance assessment. (e.g. *"I suggest book X, it is a biography of Napoleon who led the French to the Battle of Waterloo"*). This explanation is even more important when multiple different paths have been used for the production of the response.

Given an explanation, the user will either i) accept the response and proceed with the dialogue or ii) reject the response by also providing an explanatory assertion (e.g. *"no, I don't want books about event participants, I want to stay focused"*) or iii) reject the response without any explanation.

In case of an explanatory rejection (where the explanation is related to the relevance path used) and if a context is already provided, the assertion is recorded and then used to give a new response to the user. If no particular context has been given, the system enquires about the possible existence of one. If the user provides a context, the assertion is stored with it, otherwise the default context is used. On the other hand, if no explanation is given for the rejection, the system can only provide an alternative response.

Given this analysis, the dialogue game for relevance feedback elicitation from the system's users is shown in figure 1. According to this, the dialogue acts and their associated generation rules are the following:

- **Request-Information**
  - **Preconditions:** No specific preconditions
  - **Meaning:** The user seeks information from the system.
  - **Response:** Provide-Answer or Enquire-Context.
- **Enquire-Path**
  - **Preconditions:** The user has made a request involving an ontology entity.
  - **Meaning:** The system asks the user whether he/she would also be interested in other entities, connected to the requested one through specific paths.
  - **Response:** Provide-Path-Feedback.
- **Provide-Path-Feedback**
  - **Preconditions:** The system has enquired the potential importance of a relevance path
  - **Meaning:** The user provides feedback on the suggested path.
  - **Response:** Provide-Answer.
- **Enquire-Context**
  - **Preconditions:** The user has made a request for which there are known contexts but he/she has not indicated one. Alternatively, the user has rejected a response without having provided a context as an explanation.
  - **Meaning:** The system asks the user whether his request has a particular context.
  - **Response:** Provide-Context.
- **Provide-Context**
  - **Preconditions:** The system has enquired a context
  - **Meaning:** The user provides a context or states the lack of one
  - **Response:** Provide-Answer.
- **Provide-Answer**
  - **Preconditions:** The user has previously uttered one of the following locutions: i) Request-Information, ii) Reject-Answer, iii) Provide-Context iv) Provide-Path-Feedback.
  - **Meaning:** The system provides an answer to the user's request, based on the information gathered from the previous locutions, as well as an explanation for it.
  - **Response:** Accept-Answer, Reject-Answer, or Quit-Dialogue.
- **Accept-Answer**
  - **Preconditions:** The system has provided an answer

Figure 1. Dialogue Game for Feedback Elicitation

- **Meaning:** The user accepts the answer.
- **Response:** -

- **Quit-Dialogue**
  - **Preconditions:** The system has provided an answer
  - **Meaning:** The user quits the dialogue, probably because the system seems unable to answer his/her request.
  - **Response:** -
- **Reject-Answer**
  - **Preconditions:** The system has provided an answer
  - **Meaning:** The user rejects the answer, optionally providing an explanation.
  - **Response:** Provide-Answer.

3) Dialogue Example: As an example of the above dialogue game consider the following dialogue between a university student (U) and the electronic assistant of the university' library:

- U: Hi, I am looking for books on the Battle of Marathon for my thesis.
- S: Books about this battle are B1 and B2. There are also books B3 and B4 that are about the Persian Wars and you will probably find useful as well as B5 and B6 about events E1 and E2 that according to historians have led to this battle.
- U: Well show me the books about Persian Wars but not about the other events, I want to stay focused.
- S: Ok, but for thesis writing you need also information about the key personalities of this event.
- U: Yes you are right, show me these books too, but leave them last.

In this dialogue the student seeks to find books that are relevant to the Battle of Marathon. In the context she is into, namely the thesis writing, it is expected (according to the system's initialized user model) that she will be also interested in books describing greater events that contain the event she is looking for (Persian Wars in this example). Similarly it is expected to be interested in events that have caused the event she is looking for as well as in personalities that have participated in this event. Therefore, the system provides the relevant content along with corresponding explanations in order to elicit positive or negative reactions from the student. The student does verifies some and dismisses some other suggestions and their rationale, thus providing the system with the feedback it requires to update the student's user model.

C. Model Adaptation

The update of the relevance model is performed using the feedback collected by the users through the dialogue game of the previous section. In particular, for every user $u \in U$, context $ctx \in Ctx$ and path $p \in P$ we derive the set of linguistic expressions $\{l_1, l_2, ..., l_n\} = A(u, ctx, p)$ contained in the feedback.

Then, by considering the set of all possible linguistic expressions $\{l_0, l_1, ..., l_h\}$, ordered in an ascending way, we derive a new feedback function $A'(u, ctx, c, i) = \{s_1, s_2, ..., s_n\}$ where s_j is 0 if $l_j = l_0$, 1 if $l_j = l_1$, ..., h if $s_j = l_h$. Using that new function we derive the new importance degree of the p as follows:

$$Pref(u, ctx, p) = \frac{\sum_{s_j \in A'(u,ctx,p)} s_j}{h * |A'(u, ctx, p)|} \quad (1)$$

IV. Conclusions and Future Work

In this paper we proposed a novel ontology-based user personalization model in which user preferences are captured in the form of semantic relevance paths within the ontology graph. Moreover, we described a framework for (semi-) automatic acquisition of such models that exploits the dialogical interaction paradigm in order to elicit relevance related feedback from the users and use it to learn a relevance model for each of them.

Our immediate future work involves the finalization of the framework and its evaluation in two phases. In the first phase we will test and evaluate our framework's ability to learn semantic relevance paths by means of a dialogue simulation process. This kind of evaluation, in which a simulated user interacts with a dialogue system, has been widely used by researchers [17] [18] as a way to avoid the problems and difficulties related to using humans to evaluate such systems (training and employing human evaluators is expensive, qualified human users are not always immediately available, etc.), at least in the early stages of development.

In the second phase, we will evaluate the framework with real users (most likely through some crowdsourcing platform), focusing on assessing the ability of the dialogue game to elicit the required feedback from the users. In this context, we shall implement and evaluate different dialogue strategies for feedback elicitation (e.g. more aggressive ones) so as to see which works best for our purposes.

Acknowledgment

This work was supported by the European Commission under contract FP7-287615 PARLANCE.

References

[1] N. Matthijs and F. Radlinski, "Personalizing web search using long term browsing history," in *Proceedings of the fourth ACM international conference on Web search and data mining*, ser. WSDM '11, 2011, pp. 25–34.

[2] M. Daoud, L. Tamine-Lechani, M. Boughanem, and B. Chebaro, "A session based personalized search using an ontological user profile," in *Proceedings of the 2009 ACM symposium on Applied Computing*, ser. SAC '09, 2009, pp. 1732–1736.

[3] P. Mylonas, D. Vallet, P. Castells, M. FernÁndez, and Y. Avrithis, "Personalized information retrieval based on context and ontological knowledge," *Knowl. Eng. Rev.*, vol. 23, no. 1, pp. 73–100, Mar. 2008.

[4] X. Jiang and A.-H. Tan, "Learning and inferencing in user ontology for personalized semantic web search," *Inf. Sci.*, vol. 179, no. 16, pp. 2794–2808, Jul. 2009.

[5] D. Vallet, P. Castells, M. Fernandez, P. Mylonas, and Y. Avrithis, "Personalized content retrieval in context using ontological knowledge," *IEEE Trans. Cir. and Sys. for Video Technol.*, vol. 17, no. 3, pp. 336–346, Mar. 2007.

[6] A. Sieg and B. Mobasher, "learning ontology-based user profiles semantic approach to personalized web search."

[7] F. Bobillo, M. Delgado, and J. Gómez-Romero, "Representation of Context-Dependent Relevance Relations with Fuzzy Ontologies," in *ESWC'08 Workshop on Advancing Reasoning on the Web: Scalability and Commonsense (ARea2008)*, 2008.

[8] M. McTear, "Spoken dialogue applications: Research directions and commercial deployment," in *Spoken Dialogue Technology*. Springer London, 2004, pp. 19–43.

[9] C. A. Thompson, M. H. Göker, and P. Langley, "A personalized system for conversational recommendations," *J. Artif. Int. Res.*, vol. 21, no. 1, pp. 393–428, Mar. 2004.

[10] D. Griol, J. M. Molina, and Z. Callejas, "Providing personalized internet services by means of context-aware spoken dialogue systems." *JAISE*, vol. 5, no. 1, pp. 23–45, 2013.

[11] J. Teevan, S. T. Dumais, and E. Horvitz, "Personalizing search via automated analysis of interests and activities," in *Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval*, ser. SIGIR '05, 2005, pp. 449–456.

[12] M. Daoud, L. Tamine-Lechani, and M. Boughanem, "Learning user interests for a session-based personalized search," in *Proceedings of the second international symposium on Information interaction in context*, ser. IIiX '08, 2008, pp. 57–64.

[13] P. A. Chirita, W. Nejdl, R. Paiu, and C. Kohlschütter, "Using odp metadata to personalize search," in *Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval*, ser. SIGIR '05, 2005, pp. 178–185.

[14] R. Denaux, L. Aroyo, and V. Dimitrova, "OWL-OLM: Interactive Ontology-based Elicitation of User Models." Proceedings of the Workshop on Personalisation in the Semantic Web, User Modelling, PerSWeb05, pp. 34–46.

[15] J.-H. Hong and S.-B. Cho, "Domain-adaptive conversational agent with two-stage dialogue management," in *Proceedings of the 17th Australian joint conference on Advances in Artificial Intelligence*, ser. AI'04. Berlin, Heidelberg: Springer-Verlag, 2004, pp. 1147–1153.

[16] J. Szymanski and W. Duch, "Semantic memory knowledge acquisition through active dialogues." in *IJCNN*. IEEE, 2007, pp. 536–541.

[17] H. Ai and F. Weng, "User simulation as testing for spoken dialog systems," in *Proceedings of the 9th SIGdial Workshop on Discourse and Dialogue*, ser. SIGdial '08, 2008, pp. 164–171.

[18] S. Jung, C. Lee, K. Kim, M. Jeong, and G. G. Lee, "Data-driven user simulation for automated evaluation of spoken dialog systems," *Comput. Speech Lang.*, vol. 23, no. 4, pp. 479–509, Oct. 2009.