The Generation of User Interest Profiles from Semantic Quiz Games

Magnus Knuth, Nadine Ludwig, Lina Wolf, and Harald Sack

Hasso-Plattner-Institut für Softwaresystemtechnik GmbH,
Prof.-Dr.-Helmert-Str. 2–3, 14482 Potsdam, Germany
{magnus.knuth,nadine.ludwig,
lina.wolf,harald.sack}@hpi.uni-potsdam.de
http://www.hpi.uni-potsdam.de

Abstract. Personalized web search and recommendation systems aim to provide results that match the users' personal interests and lead to a more satisfactory and effective information access. Building user profiles that reflect a large spectrum from continuous (long-term) to specific (short-term) interests is an essential task when developing personalized web applications. In this paper we present a method to generate user interest profiles without direct user interaction generated out of data sourced from quiz games played by the user. Both utilized games, *WhoKnows?* and *RISQ!* , have originally been developed as serious games with the intention to rank facts in knowledge representations as well as to find inconsistencies in a given knowledge base.

Key words: user profiles, interests, DBpedia, serious games

1 Introduction

The anyhow enormous number of information objects on a diverse range of topics on the world wide web (WWW) is continually growing. Search engines can be regarded as signposts in this information universe mandatory for any information access in the WWW today. But even with the help of search engines, the overwhelming amount of information resources being delivered as search results often simply overloads the user.

One possible way out of this information overflow is considered in progressing personalization. Ideally, the users should only be provided with information that fit to their personal information needs. Search engines apply various technologies to provide personalized search results, as e. g., user history, bookmarks, community behaviour as well as click-through rate or stickiness to a specific web page. In general, the main approaches for personalization are the reranking and filtering of search results and the development of personal recommendation systems. But, to achieve results according to the user's interests and information needs, not only statistical usage information but also explicit descriptions of the user's interests are necessary that are able to reflect a large spectrum of user interests in an interoperable way. Besides personalized search results user interest profiles

can also be used to determine the user's expertise and know-how. These profiles can be applied in social networking as well as in finding experts for specific topic areas.

In this paper we present a knowledge-based approach for the creation of user interest profiles by mining and aggregating logfile data from quiz games that shift the conventional user interrogation to an entertaining setting.

There are various approaches to collect data for building user profiles, some of them are summarized in Sec. 2. In contrast to most of the existing approaches, the use of quiz games provides valuable and sufficient reliable information about how firm a user's knowledge is on given topic areas. In Sec. 3 the two games, *WhoKnows?* and *RISQ!* are introduced that have been utilized for this research. Sec. 4 summarizes the applied ontologies and category systems used to represent the user interests, while Sec.5 explains our approach in detail. In Sec. 6 first results are shown and possible future applications are pointed out. Sec. 7 provides a brief evaluation and discussion of the achieved results and Sec.8 concludes the paper with an outlook of ongoing and future work.

2   Related Work

Personalised services demand the aggregation of user profiles that represent the interests of users adequately to fulfil their mission. Such services can be personalised and adaptive web applications, such as e. g. *Persona* [1], *PResTo!* [2] or *OBIWAN* [3], as well as recommendation systems like *Quickstep* [4] for scientific papers that apply interest profiles.

Observing user behaviour is a common way to create user interest profiles, as e. g., by web usage mining or relevance feedback, c. f. [2]. These behaviour-based approaches rely on machine learning or clustering algorithms and suffer from a cold start problem, i. e. initially there are no valid recommendations for the user that can be suggested to request feedback. Another approach to obtain user interest profiles is knowledge-based profiling that employs questionaires and interviews to acquire the users' interests, which often appears intrusive or disturbing to the user.

Though there are vocabularies to model user interests only few user profiling systems apply Semantic Web technologies. The widely used *FOAF vocabulary* [5] allows to represent interests by linking documents with the `interest` property. Based on that, the *E-foaf:interest Vocabulary*[1] provides the possibility to specify more detailed statements about interests like the period of time and the value of interest. The *Cognitive Characteristics Ontology*[2] likewise allows to specify skills, expertise and interests having a weight and time relation. The *Recommendation Ontology*[3] provides a vocabulary for describing recommendations that can be ranked and addressed to groups or agents.

---

[1] http://wiki.larkc.eu/e-foaf:interest, released January 2010
[2] http://smiy.sourceforge.net/cco/spec/cognitivecharacteristics.html
[3] http://smiy.sourceforge.net/rec/spec/recommendationontology.html

For Quickstep [4] a research topic ontology containing 32 classes has been purified from the *Open Directory Project*[4] to model user interests in the computer science domain. Research papers are classified according to classes within the topic ontology by a k-Nearest-Neighbour classifier based on the documents' term vectors, which uses a manually created training set. The user interest profiles are computed from the classifications of recently accessed research papers. A 7 % to 15 % higher topic acceptance is observed using a topic hierarchy compared to a flat topic list and a small improvement in recommendation accuracy.

Serious games have been utilized before in the area of semantic web. The games *Guess What?!* [6] and the *Virtual Pet Game* [7] are used for ontology building while the quiz game *SpotTheLink* [8] tries to align concepts from the DBpedia to the PROTON upper ontology. Serious games have also been developed to annotate images including the *ESP Game* [9], and *Phetch* [10]. Most of these games can only be played in competitive multiplayer mode, in contrast to the games utilized in this paper, where a single player can play alone. Therefore, correctly and wrongly identified questions can only be identified assuming the statements within the applied ontology are correct. The purpose of these quiz games is the ranking of properties in an existing ontology.

3 Utilized Games

We have analyzed the log files of two games, namely *WhoKnows?* [11] and *RISQ!* [12] that have been developed for relevance ranking of facts in DBpedia in order to get information about entities that are known to single players. The data is anonymized, but we managed to reassign the identity of 14 persons from our research institute that can be used for evaluation.

3.1 *WhoKnows?*

Subject	Property	Object
Chile	**language**	**Spanish language**
Iraq	language	Arabic language
Brazil	language	Portuguese language
Italy	language	Italian language

Fig. 1. Screenshot and triples used to generate a One-To-One question.

---

[4] http://dmoz.org/

WhoKnows? is based on the principle to present questions to the user that have been generated out of facts stemming from DBpedia RDF triples. The game has been designed to evaluate the ranking heuristics proposed in [13]. These heuristics are based on the RDF graph structure and use statistical arguments to rank RDF properties according to their relevance. Fig. 1 shows the sample question *'Spanish language is the language of ...?'* with the correct answer *'Chile'*. The question originates from the RDF triple

```
dbp:Chile dbpprop:language dbp:Spanish_language .
```

and is composed by turning the order upside down:

```
Object is the property of: subject1, subject2, subject3...
```

Fig. 1 also shows the RDF triples for the remaining choices. In addition, false answers 'Iraq', 'Brazil', and 'Italy' are randomly selected from other triples meeting the requirement that the RDF triples' subjects belong to the same or a similar category and are not related to the object used in the question.

To add variety and to increase the user's motivation, the game is designed with different game variants: *One-To-One:* only one answer is correct, *One-To-N:* one or more answers are correct, and the *Hangman* game asks to fill the correct answer in a cloze. While playing the game, the variants are used alternately. After selecting the answer, the user immediately receives feedback about the correctness of her choice. *WhoKnows?* is described in more detail in [11].

Within the log files used for this study 5,889 rounds have been played. Approximately half of the rounds have been played in Facebook[5] by 83 players, the remaining have used the anonymous standalone version[6], whereas 11 players have given an answer in 100 to 300 rounds.

3.2 RISQ!

RISQ! has been developed as a serious game to rank the facts about renowned persons in DBpedia. The game can be played in the social network Facebook[7] as well as standalone[8]. The flow of *RISQ!* is similar to the famous TV-show *Jeopardy!*. Questions are presented to the contestants in four different topics and three different price categories. In contrast to the original *Jeopardy!* game less topics and price levels are used in order to decrease the number of questions and increase the game speed.

In each question a clue is presented to the player that points to the solution. The clues are constucted by using an RDF triple from DBpedia, replacing the property by a template to form a valid sentence, and replacing the solution by a category it belongs to. An example for such a hint is *'This New York State Senator was nominee of United States presidential election, 1940.'*, whose solution would be *'Franklin D. Roosevelt'*. Since automatically constructed hints
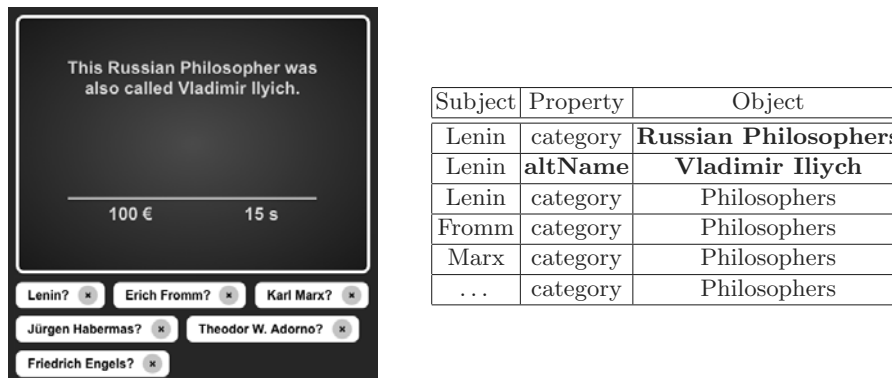
---

[5]

[6] http://tinyurl.com/whoknowsgame

[7] http://tinyurl.com/facebook-risq

[8] http://tinyurl.com/risqgamefb

Fig. 2. Screenshot and triples used to generate an active clue.

are not helpful at times and the game aims at finding the most helpful properties to identify a person, the contestants can buy additional clues with the game money.

In the orginal TV-show three contestants are playing against each other, whereas in *RISQ!* the game can be played only in single user mode so far. We introduced a timeout to prevent people from looking up the correct solution and log all player actions for later analysis.

RISQ! logged 117 unique players of which 14 has been identified as being members of our research institute. The users have played an average of 197 questions. In total 23,093 questions have been logged.

4 Entity Hierarchies

To classify the users' interests we refer to multiple entity classifications. We assume that if a user knows facts about several entities of a certain category, she is interested therein, as e. g., a player frequently answers questions about individual german soccer players correctly, it can be said she is interested in this domain, represented by Wikipedia category `GermanFootballers`, and further generalized in the `Football` domain.

The entities utilized in the games originate from DBpedia and are therefore organized in multiple hierarchical category systems, namely the DBpedia Ontology [14] and YAGO [15], and also linked to Wikipedia categories by `dc:subject` and the Freebase type system [16] via `owl:sameAs`. Each hierarchy constitutes one layer in a huge directed acyclic graph with leaf nodes containing the entities.

The DBpedia Ontology [14] is a high-level ontology, which has been manually created and contains 272 classes. Its subsumption hierarchy is comparatively shallow having a maximum depth of 6. The entities' types base on mappings of infoboxes within Wikipedia article pages to the DBpedia ontology classes.

YAGO is an automatically generated ontology based on Wikipedia and Word-Net [15]. The class system is extracted from Wikipedia categories and WordNet hyponym relations, it embraces 149,162 classes.

On Wikipedia, categories are used to organize the articles, enabling users to find and navigate related articles. According to the guidelines each article should be placed in at least one category and all of the most specific ones it logically belongs to. The category system embraces more than 450,000 categories and forms a poly-hierarchie.

Freebase [16] is a collaborative database of structured knowledge having a rather lightweight type system that consists of conceptual 'topics', that are grouped in 'types', i. e. the fixed depth is 2. The type hierarchy is not determinated, hence types can be mixed independently as needed (e. g. to assign a certain property) and are created by users. Each entity is at least of type `common/topic`.

5 Method

Answering a question in a quiz game demands that the player has knowledge about that certain entity. We assume that frequently known entities are in the users scope of interest, hence the categories, whose member entities are known frequently are assumed to form a topic relevant to the user. The level of 'proven' knowledge about an entity e is represented by a numeric value, which is calculated from the number of correctly and wrongly answered questions. As shown in (1) this value gets weighted by the ratio of given facts, respectively RDF statements in the knowledge base, that have been answered. We employed a square root for the weight to reduce the impact of subjects having a great many of statements. The rated interest in e is calculated by

$$int_{e,u} = \left(\frac{factsplayed_{e,u}}{facts_e} \right)^{1/2} \cdot \frac{correct_{e,u} - wrong_{e,u}}{correct_{e,u} + wrong_{e,u}}, (-1 \le int_{e,u} \le 1). \quad (1)$$

The users' interest in a certain category c is determined as the mean value of interests in the entities played within this category, which gets weighted by the ratio of played entities.

$$int_{c,u} = \left(\frac{entitiesplayed_{c,u}}{|\{c/e \in c\}|} \right)^{1/2} \cdot \frac{\sum_{entitiesplayed_{c,u}} int_e}{entitiesplayed_{c,u}}, (-1 \le int_{c,u} \le 1). \quad (2)$$

These ratings can be specified for a single user ($int_{x,u}$) by applying only answers of user u, as well as for the whole group of users ($int_{x,\forall}$), applying all answers, which gives us an indicator for the general knowledge. Since both games are embedded in a social network, whose members considerably differ in characteristical attributes like age, gender, origin and social background, one can suppose an adequate diversity of interests within the group, who plays the games. To distinguish special user interests from general knowledge we derive the users' performance for an entity or within a category as shown in (3). By subtracting the general knowledge of the peer group, the personal impact becomes

recognizable.

$$perf_{x,u} = \frac{int_{x,u} - int_{x,\forall}}{2}, (-1 \leq perf_{x,u} \leq 1). \tag{3}$$

A positive performance value indicates a special interest of the user in x and the higher this value, the more distinctive is the knowledge of the user compared to the general knowledge within the group.

The interests are ranked according to the users interest and performance. These outcomes can be modeled using the *Cognitive Characteristics Ontology* and integrated in the users FOAF profile, exemplary we described a main interest of Magnus here:

```
ex:magnus a foaf:Person ;
    foaf:name "Magnus Knuth" ;
    cco:interest <http://dbpedia.org/resource/Greece> ;
    cco:habit [
        a cco:CognitiveCharacteristic ;
        cco:topic <http://dbpedia.org/resource/Greece> ;
        cco:characteristic cco:interest ;
        wo:weight [
            a wo:Weight ;
            wo:weight_value 0.22 ;
            wo:scale ex:AScale
            ]
        ] .
ex:AScale a wo:Scale ;
    wo:min_weight -1.0 ;
    wo:max_weight 1.0 ;
    wo:step_size 0.01 .
```

Having computed the interest for a player in each category, we also have tried to deduce the interest in entities have never been played. Therefore the value of interest in an entity e is determined as the mean value of the categories it is a member of.

$$dint_{e,u} = \frac{\sum_{c/e \in c} rating_{c,u}}{|\{c/e \in c\}|}, (-1 \leq rating_{e,u} \leq 1). \tag{4}$$

The interest value of entities that never have been played is computed by the membership categories that serve here as a common feature. Entities that are located in branches of the hierarchy that not have been played will earn the interest of the parent categories. As an extension of this it would be possible to use further common features or relationships that could be retrieved from other properties than `rdfs:typeOf` and `dc:subjectOf`.

6 Results

Table 1 shows the top and least ranked entities for the authors of this paper. Though the complete rankings could not be strictly validated, the general rank-

ings of entities appeared to the individual players, with few exceptions, entirely acceptable.

Table 1. Entity rankings

Rank	Player A	Player B
1	Gwyneth Paltrow (0.28)	Country music (0.25)
2	Ludwig van Beethoven (0.26)	Major League Baseball (0.18)
3	Sylvester Stallone (0.24)	India (0.17)
...	...	...
n-1	Duke Ellington (-0.20)	Ohio (-0.03)
n	Richard Wagner (-0.24)	Michigan (-0.10)

Rank	Player C	Player D
1	Jack Nicholson (0.37)	Pop music (0.27)
2	Sophia Loren (0.32)	Rudyard Kipling (0.24)
3	Robert De Niro (0.30)	Royal Navy (0.24)
...	...	...
n-1	Franz Marc (-0.26)	Nigeria (-0.03)
n	Gustav Mahler (-0.29)	Wolfgang Pauli (-0.14)

One advantage of semantic search is also a disadvantage: the user, who is searching for information by intitially entering a keyword needs to decide for an entity to resolve disambiguities originating from homonymous terms. To support this task we can rank the resources according to his personal interests. As e. g., looking for the programming language 'Python' this can support someone fimiliar to information science or other programming languages. In Table 2 a comparison of the rankings made for a computer scientist and for the average user is shown. Although none of these entities have ever been played within one of the games, a tendency can be observed. Unfortunately, 'Pythonydae' and 'Monty Python' lag behind, though they seem of interest in this context.

Table 2. Rankings of different entities for the term 'Python'

Rank	computer scientist	average player
1	Python (programming language)	Python (Efteling)
2	Python (mythology)	Python (roller coaster)
3	Python (Efteling)	Python (programming language)
4	Python (film)	Python (missile)
5	Python II	Python (mythology)
6	Python of Byzantium	Colt Python
7	Python Automobile	Armstrong Siddeley Python
...	...	...

7 Evaluation

To evaluate the achieved interest profiles for the identified 14 persons, they have been asked to select ten categories from each hierarchy system and to order them according to their personal interests. We received the asked for orderings back from 12 participants. At a first glance some users' self-assessment correponds quite well with the computed ranking, though for others there have been extensive differences. To compare the players' ordering with the computed list, the longest common subsequences have been computed. Therefore, we put the computed rankings in the players' stated order and extract the *longest increasing subsequence* of that permutation with a *patience sort algorithm* [17]. The longest common subsequences have an average length of 5.5, that is more than the half was ordered correctly. There was no hierarchy that performed preferably better.

It is sometimes surprising how categories are ranked, but considering the underlying entities reveals the relationships to the known entities. In subsequent interviews we could figure out, that one reason for some players' striking differences was their intensional interpretation of the category names which deviated from the extension of the category, which is used for computation. Someone might not at all be interested in the *'Rectors of the University of Edinburgh'* but still know facts about Sir Winston Churchill or Mr Gordon Brown, who are members of this Wikipedia category.

8 Summary and Outlook on Future Work

The data gathered from the games allowed us to derive players' interests in certain entities and categories. The mapping there has been straight-forward without any detours like natural language processing or machine learning.

Given the increasing importance of social semantic web it can be valuable to publish user interests within FOAF profiles automatically or give recommendations about which topics to use.

The main problem of the application is that of incomplete coverage, since the applied datasets in both games have been filtered in advance, the data of *Who-Knows?* was filtered for entities having maximum divergence in their statements while *RISQ!* comprises solely data related to persons. This partiality excludes entire domains from being reasonably rated and must be dissolved to achieve more reliable results. Therefore, we plan to extend the entity base for both games. Nonetheless, it seems reasonable to observe further kinds of relatedness than the entities categorization.

The user profiles derived with this approach reflect rather long-term interests, in combination with specific short-term interest, that e. g. originate from log file analysis, they can be purposed for personalization of semantic search, which is one objective in our project Yovisto[9], an academic video search engine.

---

[9] http://yovisto.com/

9 Acknowledgement

References

1. Tanudjaja, F., Mui, L.: Persona: A contextualized and personalized web search (2002)
2. Keenoy, K., Levene, M.: Personalisation of web search. In Mobasher, B., Anand, S., eds.: Intelligent Techniques for Web Personalization. Volume 3169 of Lecture Notes in Computer Science. Springer Berlin / Heidelberg (2005) 201–228
3. Pretschner, A., Gauch, S.: Ontology based personalized search. In: ICTAI. (1999) 391–398
4. Middleton, S., Shadbolt, N., De Roure, D.: Ontological User Profiling in Recommender Systems. In: ACM Transactions on Information Systems. Volume 22. (2004) 54–88
5. Brickley, D., Miller, L.: The Friend Of A Friend (FOAF) vocabulary specification (November 2007) `http://xmlns.com/foaf/spec/`.
6. Markotschi, T., Völker, J.: Guess What?! Human Intelligence for Mining Linked Data. Proceedings of the Workshop on Knowledge Injection into and Extraction from Linked Data (KIELD) at the International Conference on Knowledge Engineering and Knowledge Management (EKAW) (2010)
7. Kuo, Y.l., Lee, J.C., Chiang, K.y., Wang, R., Shen, E., Chan, C.w., Hsu, J.Y.j.: Community-based game design: experiments on social games for commonsense data collections. In: Proceedings of the ACM SIGKDD Workshop on Human Computation. HCOMP '09, New York, NY, USA, ACM (2009) 15–22
8. Thaler, S., Simperl, E., Siorpaes, K.: Spotthelink: playful alignment of ontologies. In: Proceedings of the 2011 ACM Symposium on Applied Computing. SAC '11, New York, NY, USA, ACM (2011) 1711–1712
9. von Ahn, L., Dabbish, L.: Labeling images with a computer game. In: Proceedings of the 2004 Conference on Human Factors in Computing Systems, CHI 2004, Vienna, Austria, April 24 - 29, 2004. (2004) 319–326
10. von Ahn, L., Ginosar, S., Kedia, M., Liu, R., Blum, M.: Improving accessibility of the web with a computer game. In Grinter, R.E., Rodden, T., Aoki, P.M., Cutrell, E., Jeffries, R., Olson, G.M., eds.: CHI, ACM (2006) 79–82
11. Waitelonis, J., Ludwig, N., Knuth, M., Sack, H.: WhoKnows? - Evaluating Linked Data Heuristics with a Quiz that Cleans Up DBpedia. International Journal of Interactive Technology and Smart Education (ITSE) **8**(3) (2011)
12. Wolf, L., Knuth, M., Osterhoff, J., Sack, H.: Risq! renowned individuals semantic quiz a jeopardy like quiz game for ranking facts. In: Proceedings of 7th Int. Conf. on Semantic Systems I-SEMANTICS 2011, Graz, Austria, ACM (September 2011)
13. Waitelonis, J., Sack, H.: Towards Exploratory Video Search Using Linked Data. In: Proc. of the 2009 11th IEEE Int. Symp. on Multimedia (ISM2009), Washington, DC, USA (December 2009)

14. Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak, R., Ives:, Z.: DBpedia: a nucleus for a web of open data. In: The Semantic Web: 6th International Semantic Web Conference, 2nd Asian Semantic Web Conference, ISWC 2007 + ASWC 2007, Busan, Korea. (2008) 722–735
15. Suchanek, F.M., Kasneci, G., Weikum, G.: Yago: A Core of Semantic Knowledge. In: 16th international World Wide Web conference (WWW 2007), New York, NY, USA, ACM Press (2007)
16. Bollacker, K., Evans, C., Paritosh, P., Sturge, T., Taylor, J.: Freebase: a collaboratively created graph database for structuring human knowledge. In: SIGMOD '08: Proceedings of the 2008 ACM SIGMOD international conference on Management of data, New York, NY, USA, ACM (2008) 1247–1250
17. Aldous, D., Diaconis, P.: Longest increasing subsequences: from patience sorting to the baik-deift-johansson theorem. Bulletin of The American Mathematical Society **36** (1999) 413–433