Ontology-based Identification of Music for Places

Marius Kaminskas[1], Ignacio Fernández-Tobías[2], Francesco Ricci[1], and Iván Cantador[2]

[1] Faculty of Computer Science
Free University of Bozen-Bolzano
39100 Bolzano, Italy
{mkaminskas, fricci}@unibz.it

[2] Escuela Politécnica Superior
Universidad Autónoma de Madrid
28049 Madrid, Spain
{i.fernandez, ivan.cantador}@uam.es

Abstract: Place is a notion closely linked with the wealth of human experience, and invested by values, attitudes, and cultural influences. In particular, many places are strongly linked to music, which contributes to shaping the perception and the meaning of a place. In this paper we propose a computational approach for identifying musicians and music suited for a place of interest (POI). We present a knowledge-based framework built upon the DBpedia ontology, and a graph-based algorithm that scores musicians with respect to their semantic relatedness to a POI and suggests the top scoring ones. We found that users appreciate and judge as valuable the musician suggestions generated by the proposed approach. Moreover, users perceived compositions of the suggested musicians as suited for the POIs.

Keywords: Meaning of a place, semantic networks, linked data, music information retrieval

1 Introduction

Place is a notion closely linked with the wealth of human experience, and invested by values, attitudes, and cultural influences. The humanistic school of geography, based on a perspective oriented towards understanding the dimensions of human experience within a physical environment, says that place is an "experienced space" (Tuan, 1977). From a philosophical perspective, Bachelard (1958) proposed a vision of space that takes into account the emotional dimension of one's experience of an environment. He claims that a specific space can trigger emotional responses according to the experiences that occurred within it, and memories associated with it. More recently, Augé (1995) suggested to stop considering the space as a mere shell, a container, or a location, and to start looking at it as a setting for action, experiences, and communication. Besides geography and philosophy, the analysis of destination image has become one of the most popular topics in the tourism literature (Pike, 2002).

Music is strictly connected to places: it is a cultural dimension, and human activity that contributes to give a meaning to a place. For instance, consider how important is music by Johann Strauss, or flamenco, for a place like Vienna or Seville respectively. There is no doubt that this music contributes to the image of these destinations, and we all deem this music as profoundly related to the places. But, finding musicians or music related to a given place is not a simple task; it requires knowledge of both domains, and it is clearly a difficult task to be solved automatically by an intelligent computer-based system (Gretzel, 2011).

Consider even a more specific place, a "place of interest" (POI), such as the State Opera of Vienna, Austria. It is one of the most famous opera houses dating back to the 19th century, and a prominent attraction for tourists visiting Vienna. A selection of well-fitting music for this place could consist of classical compositions by musicians who lived and worked in Vienna in that historical period. Another well suited selection could include operas by Austrian composers, which are frequently performed in this opera house, such as “Don Giovanni” by Wolfgang Amadeus Mozart. An even better match could be a composition by Arnold Schoenberg, who founded the so called Viennese School in that period, and revolutionized the classical music with the introduction of the atonal and dodecaphonic music. Such music selections, although rather straightforward to be done manually by a musicologist, can be a challenging task for a computer to perform. Performing this task automatically, for any given place of interest, requires a way to identify meaningful relations between POIs and music.

We note that finding music items that suit POIs can be exploited in a number of engaging information services. For instance, a mobile city guide providing an enhanced presentation of the place visited by a tourist, and playing music that is related to the place, i.e., music that is culturally or emotionally associated to the place (e.g. Mozart in Salzburg, or a Bach’s fugue in a Gothic Cathedral). Other examples include a car entertainment and navigation system that adapts music to the place the car is passing by, or a tourism website where the information on travel destinations is enhanced through a matching music accompaniment. Such information services can be used both to enhance the user’s travel experience, and to increase the sales of holiday destinations or music content.

The main challenge that one must face when addressing the above mentioned goal is related to the fact that POIs and musicians belong to two rather different domains, and there is no obvious way to match such heterogeneous items. But, with the advent of the Semantic Web, and specifically with its reference implementation in the Linked Data initiative (Bizer et al., 2009), new opportunities arise to face the above difficulties. In this paper we propose to exploit DBpedia (Auer et al., 2008) – the Linked Data version of Wikipedia – for building a framework in which focused semantic networks linking items belonging to some selected domains are automatically created. Over these networks we propose to use a graph-based spreading activation algorithm to rank and filter the items in the target domain (music) that are most related with certain given items in the source domain (POI).

We specifically address the following research questions. *RQ1:* how to automatically identify (using DBpedia) musicians semantically related to a given POI? *RQ2:* is the music of these musicians perceived by the users as well suited to the POI? By conducting user studies we found that users appreciate and judge as valuable the suggestions generated by the proposed approach. We also found that our approach is able to distinguish between relevant and non-relevant musicians, as those with a large or small number of semantic relations with the given POI. Moreover, users perceived compositions of the suggested musicians as suited for the POIs.

2 Related Work

Finding music that suits a POI can be viewed as a context-aware recommendation problem: the place being the context for consuming the recommendation (music) (Adomavicius et al., 2011). There have been several papers on context-aware music recommendation. Ankolekar and Sandholm (2011) presented a mobile audio application that plays audio content associated with a particular location with the goal of enhancing the sense of being in a place by creating its emotional atmosphere. Instead of establishing semantic relations between music and location, as we propose here, the presented approach relies on crowd-sourcing – users of the system are allowed to assign audio pieces (either music tracks or sound clips) to a specific location (represented by its geographical coordinates). Stupar and Michel (2011) described an approach to select music tracks for a given image. They rely on associations between the low-level music and image features mined from movies and their soundtracks.

Aiming to provide contextualized recommendations of music compositions when the user is visiting tourist attractions, Kaminskas and Ricci (2011) showed that emotional tags assigned to both types of items can be used to effectively select music content that fits a POI. From that result, and aiming to complement the tag-based recommendation model proposed there, the framework presented herein enables computing matching between places and musicians based on their semantic relatedness.

There have been some other attempts to establish semantic relations between items of different types. Loizou (2009) proposed to identify explicit semantic relations between items, and exploit such relations for cross-domain recommendations. Specifically, items were annotated and linked by concepts and properties extracted from Wikipedia. Then, with such relations, users and items were incorporated into a graph, upon which a probabilistic recommendation model was built. Passant (2010) developed dbrec – a system built upon DBpedia that computes semantic distances between concepts to recommend related music bands and solo artists.

As done by Loizou (2009), in this paper we exploit Wikipedia as a source of multi-domain information, but use Semantic Web technologies and Linked Data repositories to *automatically* build semantic networks interconnecting concepts from various domains – Architecture, Art, History, and Music. In previous work (Fernández-Tobías et al., 2011) we proposed a semantic-based framework that aims to extract and aggregate DBpedia concepts and relations between two different domains, but did not evaluate such approach. In this paper we extend that work by: a) developing a more efficient method to extract cross-domain information from DBpedia, b) finding richer semantic relations between Architecture and Music concepts, and c) conducting two user studies to evaluate the semantic relatedness between automatically matched POIs and musicians.

3 Knowledge-based Framework

Our recommendation framework is built upon an ontology-based knowledge representation model in the form of a graph/network of semantic entities (concepts) in different domains, and interlinked by semantic relations (properties). The entities can be roughly categorized as classes and instances. Classes are types or categories of concepts, such as ‘city’, while instances are particular members of classes, such as ‘Madrid.’ Moreover, semantic relations can link classes (e.g. a city ‘belongs to’ a country), instances (e.g. Madrid ‘is the capital of’ Spain), or both types of entities (e.g. Madrid 'is' a city). Links can express hierarchical relationships, e.g. 'subclass of' and 'instance of', or have an arbitrary meaning.

Our goal is to automatically find paths in the above semantic graph between an instance in a source domain (a POI) to instances in a target domain (musicians), and to select (recommend) some of the reached target instances according to certain criteria.

Our approach consists of three main stages. In the first stage we restrict the subspace of DBpedia by identifying classes belonging to the two domains of interest, and the relations existing between instances of such classes. We then build a network consisting of a directed acyclic graph (DAG) whose nodes represent the identified classes, and edges represent the selected relations. In this graph there is a target node that does not have out edges, and corresponds to the class whose instances will be recommended. This stage is detailed in Section 3.1. In a second stage we instantiate the built framework into an instance network for a particular source instance (e.g. 'Vienna State Opera') whose related musicians are sought. In practice, we query DBpedia to obtain instances related to the source instance according to the classes and relations defined in the above-mentioned graph. The retrieved instances are then incorporated into the network, and are used to query DBpedia for additional related instances. This process is repeated iteratively until all the identifiable instances of the selected classes are reached. This stage is described in Section 3.2. Finally, in the third stage we assign weights to the nodes (instances) and edges of the instance network that was created, and perform a graph-based ranking algorithm over the network to identify (and recommend) the target instances with the highest score values. This stage is explained in Section 3.3.

3.1 Class Network

As we mentioned above, the output of the first stage consists of a DAG that describes how the classes in two domains of interest – source and target – are linked by means of semantic relations available in the used knowledge repositories (DBpedia). Nodes in this graph are associated to classes, and edges are associated to relations between classes and/or instances. Moreover, there is a target node without exiting edges that corresponds to the class of the instances to recommend. The selection of relevant classes and relations is guided by experts of the domains of interest and knowledge repositories. The source and target domains are respectively characterized by POIs and musicians, so 'POI' and 'Musician' classes represent the starting and ending nodes in the class network of our framework. Analysing DBpedia, we identify three potential types of semantic paths from POIs to musicians:

- *Location paths*. A particular POI may be linked to musicians who were born, lived, or died in the city of the POI. For instance, Arnold Schoenberg *was born in* Vienna, which is the city where Vienna State Opera *is located*.
- *Time paths*. A POI may be linked to musicians who were born, lived, or died in the same time period (e.g. year, decade, century) when the POI was built or opened. For instance, Gustav Mahler *was born in* 1869, the same decade when Vienna State Opera *was built*.
- *Category paths*. A POI may have Architecture categories (e.g. architectural styles and eras, building types) that are related to Music categories (e.g. music genres and eras, musician types), through relations with History and Art categories. In

this way, musicians with matched Music categories could be linked to the input POI. For instance, Wolfgang Amadeus Mozart is a *classical music composer*, and classical compositions are played in *Opera houses*, which is the building type of the Vienna State Opera.

In total, the class network contains 14 classes (e.g., ‘Architectural style’, ‘Years in architecture’, ‘Music genre’, ‘Musical era’, etc.) and 11 different relations (e.g., ‘Construction start date’, ‘Place of birth’, etc.). Furthermore, in our framework we can assign relevance values for the considered semantic entities and relations, which may be used in the retrieval stage. These values may be assigned by the domain experts, or could come from the user’s profile. For instance, a domain expert may assign higher relevance to the class ‘City’ than to the class ‘Building type’, since the former can be considered more informative to link a particular POI with related musicians. Similarly, specific concepts like ‘Opera composers’ and ‘Classical music’ may receive high relevance if the user has a clear preference for them, hence producing personalized associations. Moreover, we consider the case in which relations also receive relevance values, e.g. to measure the strength with which ‘Art Deco’ architectural style and ‘Swing’ music genre (both emerged in the 1930s) are related with respect to other more/less related categories. For simplicity, in the framework implementation described here we set the above relevance values to 1 and leave for future work the investigation of strategies to establish distinct relevance values.

3.2 Instance Networks

By exploring the structured data repository (i.e., DBpedia) through the classes and relations established in the class network (Section 3.1), we could build new networks, each one linking a particular instance in the source domain with related instances in the target domain. **Fig. 1** shows part of a sub-network that links the POI ‘Vienna State Opera’ and the composer ‘Gustav Mahler.’ The full network of a POI is obtained by aggregating all the sub-networks linking the POI to semantically related musicians.

To test our approach, we have automatically extracted from DBPedia a dataset of 2860 POIs from 17 major city tourism destinations in Europe (Madrid, Berlin, Florence, Paris, etc.). On average, in this dataset, an instance network of a POI contains 684.5 musicians, i.e., this number of musicians can be retrieved per POI, by following the semantic paths established in section 3.1. We refer the reader to (Fernández-Tobías et al., 2011) for more details on the collected dataset.

An instance network has weights assigned to the relations between pairs of instances. These weights are computed from the relevance values of the relations linking the instance pairs and their classes. Specifically, a weight between two instances I, $I' \in \mathit{I}$ is given by a function $w: \mathit{I} \times \mathit{I} \rightarrow R_+$ that depends on the relevance values of the connections between the two instances, and between their classes C_I, $C_{I'} \in \mathit{C}$, that is, $w(I, I') = f(rel(I, I'), rel(C_I, C_{I'}))$, where $w(I, I')=0$ if there is no link from I to I’. We refer the reader to (Fernández-Tobías et al., 2011) for more details on the particular computation of the weights and relevance values.

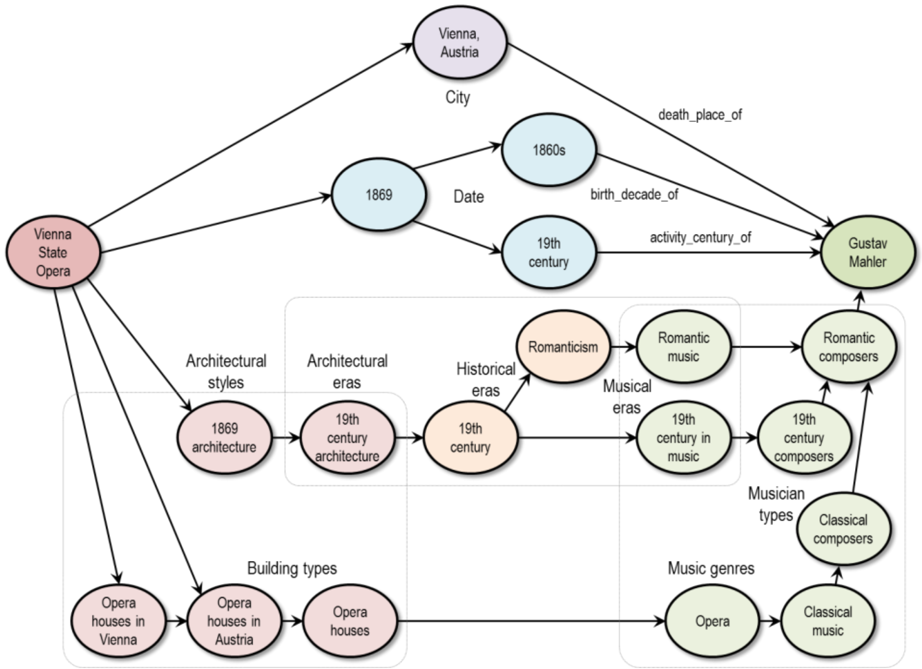

Fig. 1. Example of semantic sub-network linking the POI 'Vienna State Opera' with the composer 'Gustav Mahler'

3.3 Graph-based Ranking

In the final stage we run a graph-based ranking algorithm on the built instance networks. For each entity node in the network, the algorithm computes a relevance score to the source POI, by following a weight spreading strategy, and hence it provides a scoring of the target nodes (i.e., those belonging to the class 'Musician'). Then, the highest scored nodes are selected for recommendation. Thus, the score of an instance node *I* depends not only on the relevance values of the instance and its class, but also on the scores of instances that are connected to *I* along some paths in the semantic network.

Initially the score values of the instance nodes are set based on their relevance values, which, as explained in Section 3.1, may be related to domain characteristics and/or user preferences. Then, our spreading activation technique performs a single iteration for propagating instance scores in the instance network. The algorithm propagates the initial score of the source node through its weighted edges, updating the scores of its linked nodes. This is iteratively done for subsequent linked nodes until reaching the target nodes, whose scores cannot be further propagated because they do not have out edges (**Fig. 1**). We refer the reader to (Kaminskas et al., 2012) for more details on the networks building process and ranking algorithm.

4 Experiments

We conducted two user studies aimed at evaluating how users perceive and judge musician-to-POI recommendations provided by our approach. Since the users cannot perform a large number of judgements during an evaluation session, we limited the evaluation dataset to 25 POIs (from the full set of 2860 POIs described in section 3.2). On average, these POIs had instance networks consisting of 708.2 nodes (with 668 nodes representing musicians). For each POI, we obtained the top 5 ranked musicians computed by three different methods: the spreading activation algorithm that we designed, and two additional baseline algorithms, HITS and PageRank (Manning, 2008). Additionally, 5 musicians were randomly selected from the whole set of musicians belonging to the POI's instance network, as an additional baseline method with which to compare our approach. We also downloaded a representative music composition for each musician. So in total, we had 4 competing methods to find a musician (and its music composition) matching any given POI.

The first experiment was designed to evaluate how the users judge the semantic relations between POIs and retrieved musicians, while the second experiment was performed to evaluate if the users deem compositions performed by the retrieved musicians as relevant for the POIs.

4.1 Knowledge-Based Relations between POIs and Musicians

The graphical interface for the first user study had to be carefully designed since assessing the quality of the relatedness (matching) of items from different domains is not easy. Hence, we designed a tailored interface (**Fig. 2**) that may require some considerable user effort, but let collect important and interesting information about the user-perceived quality of different musician matches for POIs.

During each evaluation session a user was presented with a sequence of 10 pairs of POIs and musicians, where the musicians were obtained using one of the four evaluated matching methods. The information describing each POI and musician pair was presented in a structured way, according to the representation in the graph model – location, date, and category relations were clearly separated. The user was asked to carefully check the presented information, and assess whether the musician was actually related to the POI, and if yes to specify which parts of the structured musician information were contributing, and in which degree, to the match (right part of the system's interface, in **Fig. 2**). We aimed to understand which types of semantic paths, linking POIs to musicians (i.e., either location, or date, or category paths) contribute more to the matches, and were better appreciated by the users.

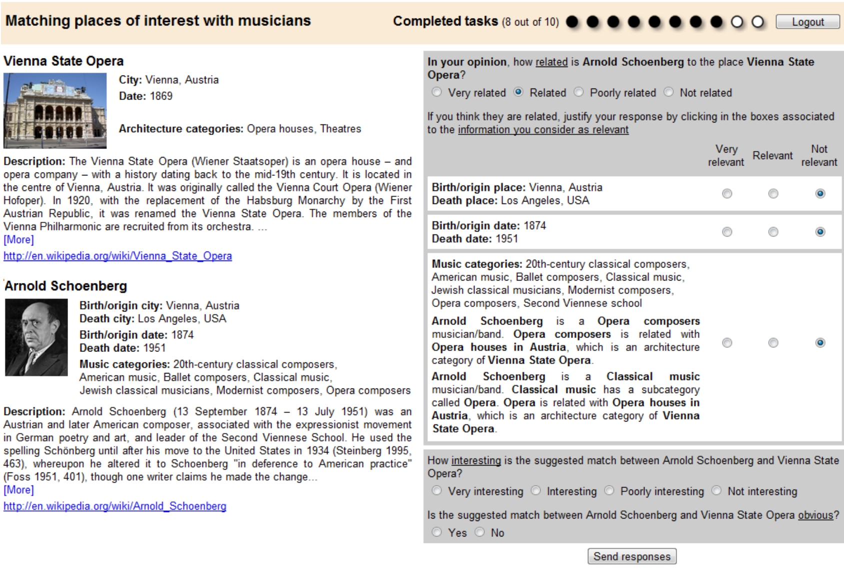

Fig. 2. Screenshot of the system used in assessing the knowledge-based relatedness of POIs and musicians

4.2 Matching Music to POIs Evaluation Study

For the second study we designed a simpler graphical interface (**Fig. 3**). During a single evaluation session a user was presented with a POI and a list of compositions of the musicians selected by each of the four evaluated methods. The order of the compositions was randomized, and the user was not aware of the algorithms that were used. The user was then asked to read the description of the POI, listen to the compositions, and select those that in her opinion well suited the POI.

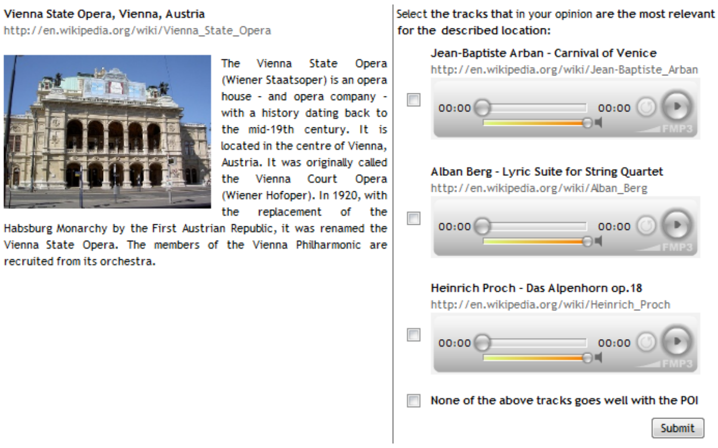

Fig. 3. Screenshot of the system used in the matching music to POIs evaluation study

Moreover, in this evaluation we asked the users to enter their music genre preferences prior to performing the evaluation. This was done to measure the influence of the users' music preferences on their decisions. The genre taxonomy was selected based on the musicians in our dataset, and included Classical, Pop, Medieval, Opera, Rock, Ambient, Folk, Hip Hop, Metal, and Electronic music genres.

5 Results

5.1 Relations between POIs and Musicians: Results

A total of 97 users participated in the study. They were PhD students and academic staff recruited via email, and covered an ample spectrum of ages and nationalities. They provided 1155 assessments for 356 distinct POI-musician pairs (note that a musician may match a POI for various scoring methods). Each of the 356 distinct pairs was assessed by at least 3 users. The Fleiss' Kappa correlation coefficient of the relatedness assessments per POI was 0.675, meaning a substantial agreement among users. **Table 1** clearly shows that the proposed spreading activation method outperforms the baselines, in the precision@K metric.

Table 1. Average precision values obtained for the top 1 to 5 ranked musicians for each POI. The values marked with * have differences statistically significant (Wilcoxon signed-rank test, $p<0.05$) with Spreading algorithm's

	P@1	P@2	P@3	P@4	P@5
Random	0.355*	0.391*	0.363*	0.435*	0.413*
HITS	0.688	0.706	0.711*	0.700*	0.694
PageRank	0.753	0.728	0.707*	0.660*	0.646*
Spreading	0.810	0.804	0.828	0.847	0.837

In order to understand which semantic information contributed, and in which degree, to the matches, **Fig. 4** shows the average numbers of semantic paths in the instance network between the input POIs and the retrieved musicians. The higher these numbers, the more semantic relations between POIs and musicians were found, and thus the richer the semantic relation between the retrieved musician and the source POI is. Based on the obtained results, we can highlight two aspects of the proposed spreading algorithm. First, differently from the other two baseline algorithms, it differentiates relevant and non-relevant musicians by finding a larger number of paths, between a POI and a relevant musician, compared with a non relevant one. Second, it uses in a balanced way all the different types of the considered relations.

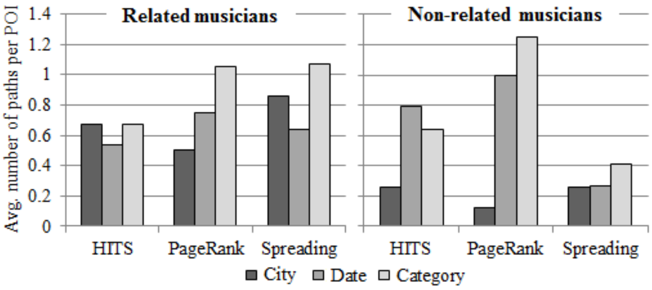

Fig. 4. Average number of semantic paths per POI

5.2 Matching Music to POIs Evaluation: Results

A total of 61 user participated in the second study. As in the first case, they were PhD students and academic staff recruited via email (some of them also participated in the previous study). 1125 evaluation sessions were performed (i.e., a POI shown to a user), and 1258 tracks were selected by the users as well-suited for a POI. **Fig. 5** shows the performance of the matching methods, computed as the ratio of the number of times a track produced by each method was considered as well-suited over the total number of evaluation sessions (1125). All methods performed significantly better than the random track selection ($p<0.01$ in a two-proportion z-test). Moreover, again, our weight spreading activation method outperformed the others ($p<0.01$).

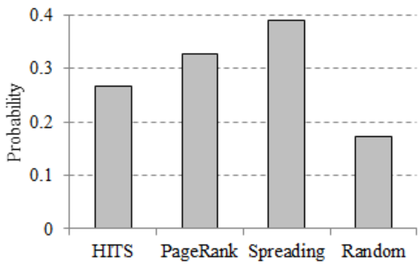

Fig. 5. Selection probability of the recommendation approaches

We also analysed the influence of the users' genre preferences on the music tracks that they selected as well-suited for the POIs. In the following, *TrkOK* represents the condition that a track was marked by a user as well-suited for a POI, $GTrk=g$ – the condition that the genre of a track is g, $g \in uPref$ – the condition that a user has included the genre g in her genre preferences, and $g \notin uPref$ – the condition that a user has not included the genre in her preferences.

As a baseline for the analysis we have computed, for each genre, the probability that a music track is of genre g, given that it was marked as suited for a POI – $P(GTrk=g \mid TrkOK)$ – as the ratio of the number of tracks of genre g selected as well-suited for the POIs over the total number of tracks selected as suited for POIs (1258). Then, to check the deviation from this baseline produced by the users' preferences, we measured $P(GTrk=g \mid TrkOK, g \in uPref)$ and $P(GTrk=g \mid TrkOK, g \notin uPref)$.

We have found that for the Classical, Medieval and Opera music genres the deviation from the baseline is significant (*p<0.01* in a two-proportion z-test). For Classical music, the probability that a user who likes this genre will mark a classical track as well-suited for a POI is higher, and in the opposite case – lower. For other music genres the deviation from the baseline probabilities is not significant.

We also measured the conditional probability for a track to be selected as suited for a POI, given that its genre is g – $P(TrkOK \mid GTrk=g)$ – as the ratio of the number of tracks of genre g selected as suited over the total number of times a track of genre g was displayed during the evaluation. The obtained results show that Medieval and Opera tracks are most often selected by the users (*0.49* probability), followed by Ambient (*0.36*) and Classical (*0.35*) music. Then, to check the deviation from this baseline produced by the users' preferences, we measured the probabilities $P(TrkOK \mid GTrk=g, g \in uPref)$ and $P(TrkOK \mid GTrk=g, g \notin uPref)$. From these probabilities the effect of user preferences is evident for Classical and Medieval music – the probability that a user will select a Classical/Medieval track as well suited for a POI is significantly ($p<0.05$) higher if the user likes these genres, compared to a user who does not.

We can thus confirm that for certain genres, in addition to the semantic matching between POIs and musician, there is a clear effect of user preferences on the decision for considering music to go well with a POI. Therefore, in the future we plan to take the users' music preferences into account when suggesting music suited for a POI.

6 Conclusions and Future Work

In this paper we have described an original approach for automatically identifying musicians semantically related to a given place of interest (*RQ1*). Two user studies showed that the musicians matched to a POI by the proposed approach are more likely to be considered as well suited to the POI by the user. Moreover, for the musicians that the users state as related to a POI, our approach can identify a high number of paths in the semantic network connecting them to the POI using all the proposed types of relations (i.e., city, date, and category). Conversely, these numbers are significantly lower for non-related musicians. This indicates that our approach, which is based on finding semantic paths between POIs and musicians, can be used effectively to define and evaluate the semantic relatedness between such instances. Finally, we have shown that users perceive music composed/performed by the recommended musicians as well-suited for the POIs (*RQ2*).

Future work will focus on identifying and exploiting other arbitrary semantic relations between POIs and musicians, e.g. direct relations such as 'Gustav Mahler *was the director of* Vienna State Opera', and complex non-directed relations such as 'Ana Belén (a famous Spanish singer) *composed a song* whose lyrics *are about* La Puerta de Alcalá (a well-known POI in Madrid, Spain).' We will search for these relations in different Linked Data repositories, and will use available tools such as RelFinder (Heim et al., 2009).

It is also important to better explore and evaluate strategies to initialize the importance weights of the relations and classes in the semantic network. In this context, we will explore other approaches to define and measure the semantic relatedness between concepts (Ponzetto & Strube, 2007), and will adapt or extend them for cases in which the concepts belong to different domains. We will also explore more complex constrained spreading activation techniques (Crestani, 1997) to propagate the weights through the semantic networks, by taking into account factors such as path lengths, node in/out-degrees, and weight propagation thresholds. Finally, we will adapt the proposed framework to incorporate user preferences into the music recommendation process

References

Adomavicius, G., Mobasher, B., Ricci, F., Tuzhilin, A. (2011). Context-aware recommender systems. *AI Magazine* 32(3), 67-80.

Ankolekar, A., Sandholm, T. (2011). Foxtrot: a soundtrack for where you are. In: *I*nteracting with Sound Workshop: Exploring Context-Aware, Local and Social Audio Applications, pp. 26-31.

Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak, R., Ives, Z. (2008). DBpedia: A nucleus for a web of open data. In: International Semantic Web Conference, ISWC 08,, pp. 722-735.

Augé, M. (1995). Non-Places. *Introduction to an Anthropology of Supermodernity*. Verso.

Bachelard, G. (1958). *The Poetics of Space*. Beacon Press.

Bizer, C., Heath, T., Berners-Lee, T. (2009). Linked data – the story so far. *International Journal on Semantic Web and Information Systems* 5(3), 1-22.

Crestani, F. (1997). Application of Spreading Activation Techniques in Information Retrieval. *Artificial Intelligence Review* 11, 453-482.

Fernández-Tobías, I., Cantador, I., Kaminskas, M., Ricci, F. (2011). A generic semantic-based framework for cross-domain recommendation. In: Proceedings of the 2nd International Workshop on Information Heterogeneity and Fusion in Recommender Systems., pp. 25-32.

Gretzel, U. (2011). Intelligent systems in tourism – a social science perspective. *Tourism Management* 38(3), 757-779.

Heim, P., Hellmann, S., Lehmann, J., Lohmann, S., Stegemann, T. (2009). RelFinder: Revealing relationships in RDF knowledge bases. In: SAMT'09, pp. 182-187.

Kaminskas, M., Ricci, F. (2011). Location-adapted music recommendation using tags. In: 19th International Conference on User Modeling, Adaptation and Personalization, pp. 183-194.

Kaminskas, M., Fernández-Tobías, I., Ricci, F., Cantador, I. (2012). Knowledge-based Music Retrieval for Places of Interest. In: Second International ACM Workshop on Music Information Retrieval with User-Centered and Multimodal Strategies. In press.

Loizou, A. (2009). How to recommend music to film buffs. Ph.D. thesis, University of Southampton.

Manning, C. (2008). Introduction to Information Retrieval. Cambridge University Press, Cambridge.

Passant, A. (2010). Dbrec – music recommendations using DBpedia. In: ISWC'10, pp. 209-224.

Pike, S. (2002). Destination image analysis: a review of 142 papers from 1973 to 2000. *Tourism Management* 23, 541-549.

Ponzetto, S., Strube, M. (2007). Knowledge derived from Wikipedia for computing semantic relatedness. *Journal of Artificial Intelligence Research* 30(1), 181-212.

Stupar, A., Michel, S. (2011). Picasso – to sing, you must close your eyes and draw. In: Proceeding of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 715-724.

Tuan, Y. F. (1977). Space and Place: The Perspective of Experience. University of Minnesota Press.