



Repositorio Institucional de la Universidad Autónoma de Madrid

<https://repositorio.uam.es>

Esta es la **versión de autor** de la comunicación de congreso publicada en:
This is an **author produced version** of a paper published in:

Information Technology & Tourism 14.1 (2014): 73-95

DOI: <http://dx.doi.org/10.1007/s40558-014-0004-x>

Copyright: © Springer-Verlag Berlin Heidelberg 2014

El acceso a la versión del editor puede requerir la suscripción del recurso
Access to the published version may require subscription

Knowledge-based Identification of Music Suited for Places of Interest

Marius Kaminskas¹, Ignacio Fernández-Tobías², Francesco Ricci¹, and Iván Cantador²

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 related to music, which contributes to shaping the perception and meaning of a place. In this paper we propose a computational approach to identify musicians and music suited for a place of interest (POI) – which is based on a knowledge-based framework built upon the DBpedia ontology –, and a graph-based algorithm that scores musicians with respect to their semantic relatedness with a POI, and suggests the top scoring ones. Through empirical experiments we show that users appreciate and judge the musician recommendations generated by the proposed approach as valuable, and perceive compositions of the suggested musicians as suited for the POIs.

Keywords: meaning of a place, semantic networks, linked data, music recommendation

1 Introduction

Place is a notion closely linked with the wealth of human experience, and invested by values, attitudes and cultural influences (Ciolfi 2004). The humanistic school of geography, based on a perspective oriented towards understanding the dimensions of human experience within a physical environment, states that place is an “experienced space” (Tuan, 1977). From a philosophical perspective, Bachelard (1958) proposed a vision of space that takes the emotional dimension of one’s experience of an environment into account. He claimed that a specific space can trigger emotional responses according to the experiences that occurred within it, and to 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. Similarly, in tourism literature a destination, which is a target for a tourist travel, and comprises a set of places (of interest), cannot be reduced to its geographical boundaries (Werthner and Ricci, 2004). In fact what matters for a tourist (and for a Destination Management Organization, DMO) is the destination image, which is a mental construct based on a flood of information from promotional literature, opinion of others, and media (Echtner and Ritchie, 2003). In this context, we note that the analysis of the destination image has become one of the most popular topics in the tourism literature (Pike, 2002).

Therefore, these human-oriented perspectives of the notion of place go beyond the physical dimension, and connect to the user’s experience of a place. They thus motivate the analysis and introduction of new types of ubiquitous computing systems that could better interact with, and engage a situated user leveraging the information provided in a particular situation and context (Dey 2001). In particular, in this paper

we focus on technologies for identifying meaningful recommendations of musicians and music for situations where place plays an important role, e.g. the place where the user is located, or the place whose information and description the user is browsing.

Music, as any other artistic production, is often 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 the music by Johann Strauss is for a city like Vienna, flamenco music for Seville, or Beatles' songs to UK. There is no doubt that this music contributes to the image of such destinations, and we all deem this music as profoundly related to the places. In fact, many musical compositions have been motivated or inspired by specific places. Consider for instance the impressionistic compositions by Debussy, such as *La mer* or *Preludes*: they are all clearly dedicated to places (e.g. "La Puerta de Vino" and "La Terrasse Des Audiences Du Clair De Lune"). Moreover, a musician like Debussy is intimately linked to France, not only because of his nationality, but also due to his culture and role played in that country.

Finding musicians or music related to a given place, however, is not a simple task – it requires knowledge of both domains, and a methodology for establishing relations between the objects in the two domains, which is clearly a difficult problem to be solved automatically by an intelligent computer-based system (Gretzel, 2011). For example, consider a 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 that opera house, such as "Don Giovanni" by Wolfgang Amadeus Mozart. Another meaningful 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. Performing this task automatically, for any given POI – which may not be a music venue –, requires a way to identify meaningful relations between places and music.

We note that finding music items that suit POIs can be exploited in a number of engaging information services, which ultimately motivate this research. For instance, a mobile city guide may provide an enhanced presentation of the place visited by a tourist, and play music that is related, i.e., 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 to enhance the user's travel experience, to provide rich and engaging cultural information services, 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 and unique way to match items from such heterogeneous domains. 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 that are recognized by the users as semantically related to a given POI?
- *RQ2*: is the music of these musicians perceived by the users as well-suited for the POI?

We note that one could directly face the task of identifying music compositions well suited for a POI by matching these two types of items. We have addressed this problem in a previous work (Braunhofer et al., 2013). But in that research we did not exploit Linked Data and semantic relations (see the next Section on related work). In fact, matching a place directly to a music composition using semantic relations would require a very specific knowledge base of compositions, describing each single composition in terms of its semantic properties and relations, which is not available. For this reason, knowledge-based approaches for matching places to music must rely on an additional step of matching the places to composers.

By conducting two user studies we found that users judge the suggestions generated by the proposed approach as valuable. 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, in our experiments users perceived compositions of the suggested musicians as suited for the POIs.

The remainder of this paper is structured as follows. In Section 2 we review works in the areas related to our research problem, namely context-aware music recommendation and establishing item similarity across domains. In Section 3 we present our knowledge-based framework for identifying musicians suited for POIs. In Sections 4 and 5 we describe the evaluation of the framework, and present the obtained empirical results. In Section 6 we discuss the limitations of the proposed approach, and outline directions for future work. Finally, in Section 7 we provide conclusions of the conducted research.

2 Related Work

Finding music that suits a POI is a research problem that can be framed as context-aware recommendation, with the POI playing the role of the context, and the music being the content that should be adapted to or recommended for the specific situation determined by the context. Context-aware computing emerged in the mid 90's from the ubiquitous computing paradigm – the idea of omnipresent computing accessible to

the user at any time and in different conditions (Weiser, 1993). In computing applications, *context* is considered as any information that influences the interaction of the user with the system (Dey, 2000). Applications that can use such information to provide relevant information or services to the user are known as context-aware applications, and have become popular in the areas of mobile computing (Abowd et al., 1997), adaptive user interfaces (Brotherton et al., 1999), and recommender systems (Adomavicius et al., 2011).

For instance, in the area of recommender systems (Ricci et al., 2011), certain factors, such as time, location, and the purchasing purpose, identify the *context* in which recommendations are provided (Adomavicius et al., 2011). The situation defined by the contextual factors influences the information need of the user, and thus should be taken into account when generating recommendations (in addition to the more conventional knowledge of the user's short- and long-term preferences, needs and wants).

As noted above, identifying music that is relevant for a place can be viewed as a context-aware recommendation problem, where the place is the context for consuming recommendations (music). We stress that in this work we are not addressing the recommendation problem in its classical sense, i.e., personalizing the information content according to the user's preferences. Instead, we are proposing an approach to directly match the information items (music) with the user's context (a POI). Such approach may let exploit context information when generating personalized recommendations.

In the music recommendation field, there are several works that exploit context information (Kaminskas and Ricci, 2012). Ankolekar and Sandholm (2011) presented a mobile audio application that plays audio content associated with a particular location aiming to enhance 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, that 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 relied on associations between the low-level music and image features mined from movies and their soundtracks.

Braunhofer et al. (2013) presented an approach to select music fitting a POI by matching the emotional properties of music tracks and POIs. While emotions provide a natural way for humans to describe music (Zentner et al., 2008), places can trigger emotional responses in visitors (Bachelard, 1958). Through a series of user studies it was shown that a set of emotion tags can be utilized by users to describe both POIs and music tracks, and the commonality of emotions provides the base for establishing a degree of match between a place and a music composition. The tag-based representation lets match POIs and music tracks by comparing their tag profiles by means of standard vector similarity metrics.

Unlike the above work, in this paper we present an approach where the items may not be represented using the same set of features, and therefore standard similarity metrics may not be applicable. Instead, we are looking for general ways to establish meaningful semantic relations between the items across heterogeneous domains. This

is similar to the relatedness of concepts in WordNet developed by Pedersen et al. (2004), who implemented relatedness metrics for the lexical database WordNet that, unlike similarity metrics, can be applied across part of speech (POS) boundaries. In WordNet, nouns and verbs are organized into separate hierarchies. Consequently, while a similarity between two nouns can be computed using taxonomy-based metrics (Resnik, 1995; Lin, 1998; Seco et al., 2004), or graph-based metrics (Rada et al., 1989; Wu and Palmer, 1994; Jiang and Conrath, 1997), establishing how similar a noun and a verb are requires additional information sources (e.g. analyzing glossary definitions) that enable reasoning across POS domains.

Aiming to establish semantic relatedness metrics between concepts, more recently, the research focus has moved towards using new large-scale datasets with free link structure, such as Wikipedia, instead of using fixed hierarchical taxonomies and thesauri, such as WordNet. Hence, for instance, Milne and Witten (2008) proposed a path-based method, called Wikipedia Link-based Measure (WLM), which uses the inner hyperlinks between Wikipedia articles describing concepts. The semantic relatedness of two concepts (articles) is estimated averaging scores for their incoming and outgoing links. For the outgoing links, the relatedness is computed as the cosine similarity between the outgoing link vectors of the two articles. Thus, two articles are similar if they point to the same documents with relevant links. Differently to WLM, our approach – which exploits the structure of Wikipedia’s category graph – computes the semantic relatedness between two concepts by only considering the links to those concepts (articles) that belong to certain domains of interest, avoiding thus the noisy influence of scarcely related domains.

In the recommender systems field, 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. Narducci et al. (2013) used Wikipedia to model user preferences by representing each term in a user’s profile as a distribution over Wikipedia topics (articles).

As done by Loizou (2009) and Narducci et al. (2013), 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 a previous work (Kaminskas et al., 2013) we developed an efficient method to extract cross-domain information from DBpedia by finding semantic relations between architecture and music concepts, and conducted two user studies to evaluate the users’ subjective assessment of the semantic relatedness between automatically matched POIs and musicians. This paper extends our previous work by: a) providing a more articulated discussion of the motivations and the relevance of the proposed technique, b) providing more detailed descriptions of the matching techniques, c) performing a deeper analysis of the impact of users’ music preferences in the evaluation of music recommended for POIs.

3 Knowledge-based Framework

Our framework is composed of an ontology-based knowledge representation model in the form of a graph/network of semantic entities (concepts) belonging to different domains, and interlinked by semantic relations (properties). The entities can be categorized as classes and instances. Classes represent concepts, such as *city*, while instances represent particular members of the 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) and 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 domains of interest, and the relations existing between these 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 source node that does not have incoming edges, and corresponds to the class whose instances form the input recommendation context, i.e., the POIs. In this graph there is also a target node that does not have outgoing edges, and corresponds to the class whose instances will be recommended, i.e., the musicians. 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 POI instance (e.g. *Vienna State Opera*) whose related musicians are sought. In practice, we query DBpedia to obtain instances related with the source POI instance according to the classes and relations defined in the above-mentioned graph. The retrieved instances are then incorporated into the instance 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 a 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 scores. 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 – POIs and musicians – are linked by means of semantic relations (*paths*) available in the DBpedia knowledge repository.

As done in the majority of the Linked Data repositories, in DBpedia data is stored and retrieved in the form of triples, which are composed of *subject-property-object* elements, such as $\langle \textit{Arnold Schoenberg}, \textit{birth place}, \textit{Vienna} \rangle$, where a subject and an object belong to certain classes (e.g. *Person*, *City*), and the property denotes a relation

between the classes (e.g. a person *being born* in a city). We build the class network graph in such a way that nodes correspond to classes, and edges to relations between the classes. Moreover, in the graph there is a starting node without incoming edges that corresponds to the *POI* class, and a target node without outgoing edges that corresponds to the *Musician* class (i.e., the class of items to be recommended).

We define the concept of *path* in a class network as a sequence of connected nodes in the graph that connect the starting node (a *POI*) with the target node (a *Musician*). The number of such paths depends on the selection of relevant classes and relations to be included into the graph, which is guided by experts of the domains of interest and knowledge repositories. Since our goal was to identify direct and explicit semantic relations between the items, we aimed to identify the classes and properties that can provide meaningful explanations when a link between a *POI* and a musician is established and used. In fact, DBpedia provides a large number of potentially interesting properties that could be used in the class network (e.g. the *Person* class has more than 40 properties that denote its relations to other classes: *birth place*, *birth date*, *spouse*, *parent*, *university*, among others). However, we cannot include all possible relations in the graph since it may result in a semantic network with concepts belonging to domains not (closely) related to those of interest. Therefore, having carefully analyzed the available relations, we have identified the following paths for the graph:

- *City 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 the Vienna State Opera *is located*.
- *Date 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 the Vienna State Opera *was built*.
- *Art style (category) paths*. A *POI* may belong to certain architectural styles, eras and building types that are related to music genres and eras through relations with history and art periods. In this way, musicians with matched music styles 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 our approach the concepts of this type of paths are Wikipedia's categories¹, and the relations between such concepts correspond to a hierarchical property that we have called *subcategory_of*, which may relate categories from different domains.

Figure 1 shows a simplified view of the class network of our framework. In total, the defined class network contains 17 main classes and 16 distinct relations.

¹ Wikipedia's categorization, <http://en.wikipedia.org/wiki/Wikipedia:Category>

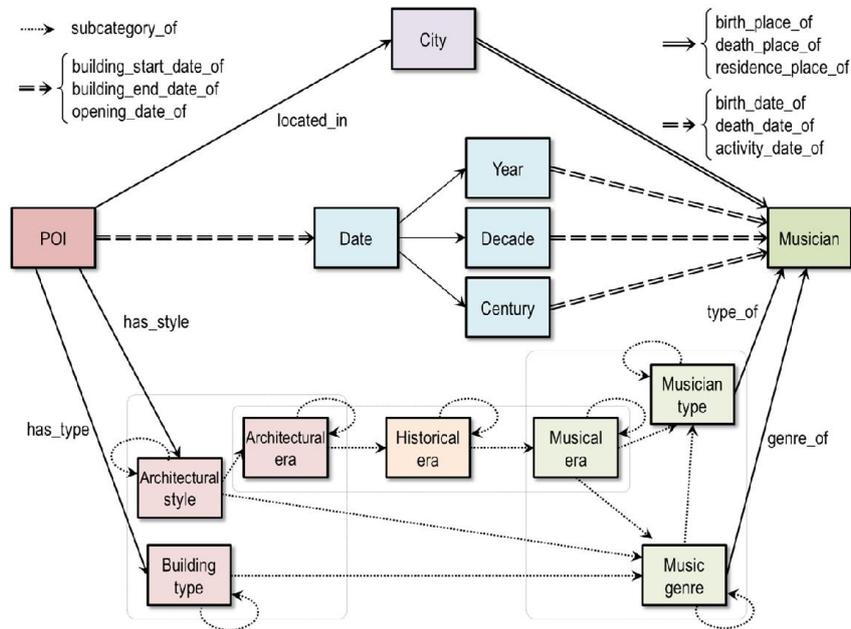


Figure 1. Simplified view of the class network used in the framework; some of the architecture, history and music concepts group various classes

The classes in the network are the following: *Visitor attraction (POI)*, *City*, *Year*, *Decade*, *Century*, *Architectural style*, *Building type*, *Architectural era*, *Historical era*, *Musical era*, *Musician type*, *Music genre*, *Musician*, *Composer*, *Singer*, and *Musical group*. As shown in the figure, the classes are linked using the following relations:

- *location city, birth place, origin place, residence place, and death place* for the city paths;
- *building start date, building end date, opening date, birth date, activity date, death date* for the date paths;
- *belonging to architectural style, having building type, belonging to music genre, being musician type, and having sub-category* for the art style paths.

Furthermore, in our framework we can assign relevance values to the considered classes and relations, which may be used in the ranking/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 important when linking a POI with related musicians. Or, the relation *residence place* may be assigned a higher relevance than the *death place* relation, since it is more important to consider the place where a musician lived during his/her career.

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 can build new networks, each one linking a particular *POI* instance in the source domain with related *Musician* instances in the target domain.

Figure 2 shows part of a sub-network that links the POI Vienna State Opera to 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.

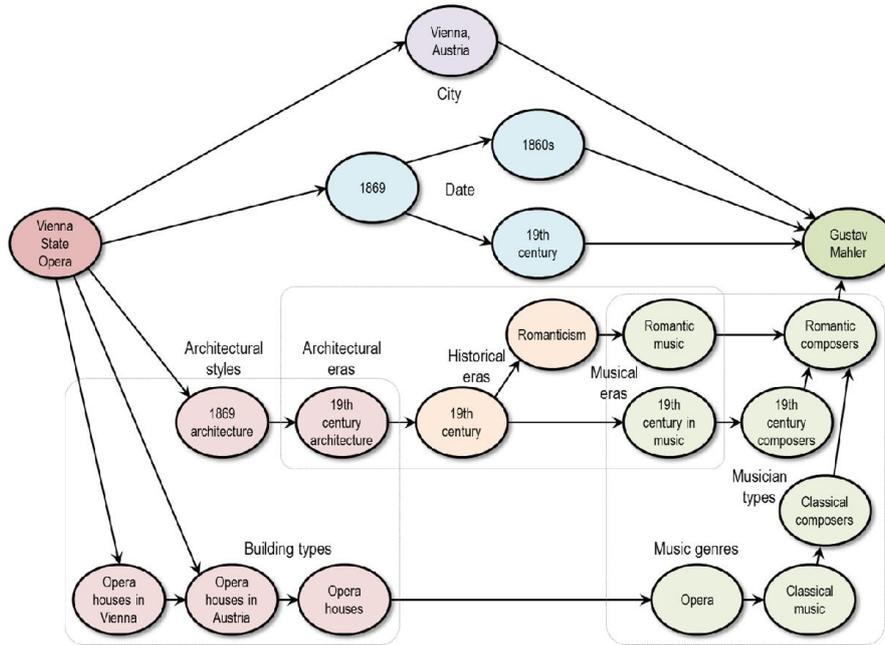


Figure 2. An example of semantic sub-network linking the POI Vienna State Opera with the composer Gustav Mahler

Similarly to the class network (see Section 3.1), on an instance level we can define relevance values for specific instances. For example, the *Classical music genre* may receive higher relevance than other music genres if the user has a clear preference for it, hence producing personalized associations. Moreover, relevance values can be assigned to relations between instances. For example, we can assign a higher relevance weight to the relation between *Art Deco* architectural style and *Swing* music genre (both emerged in the 1930s) than to relations between other less related music genres and architectural styles.

At this stage in our research, we manually assign the relevance values of classes $rel(C_i)$ and their relations $rel(C_i, C_j)$. In the future we intend to exploit and automatically set all the types of relevance values (i.e., the relevance of classes, instances, relations between classes, and relations between instances) for computing

personalized matching between POIs and musicians, and for fine tuning the semantic framework.

To test our approach, we have used DBpedia for extracting a dataset of 2860 POIs from 17 major city tourism destinations in Europe (Madrid, Berlin, Florence, Paris, etc.). On average, in this dataset the instance network of a POI contains 684.5 musicians; that is, 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 used dataset.

The relation between two instances may be assigned a weight $w(i, j)$, which can be computed from the relevance value of such a relation at the class and/or instance level. Specifically, a weight between two instances $i, j \in I$ is a function $w: I \times I \rightarrow \mathbb{R}_+$ that depends on the relevance values of the connections between the two instances, and the connections between their classes $C_i, C_j \in C$:

$$w(i, j) = f\left(\text{rel}(i, j), \text{rel}(C_i, C_j)\right)$$

where $w(i, j) = 0$ if there is no link from i to j .

We refer the reader to (Fernández-Tobías et al., 2011) for more details about the particular setting of the relevance values, and the computation of the weights. In that paper, we defined f as a linear combination of the relevance values $\text{rel}(i, j)$ and $\text{rel}(C_i, C_j)$, and proposed to set the relevance values of the relations based on the co-occurrence of entity pairs in a text corpus. In the experiments presented here we set the $\text{rel}(i, j)$ values to 1, so in the current implementation of our framework, the weight $w(i, j)$, and consequently the function f , depends solely on the relevance value of relation between their classes $\text{rel}(C_i, C_j)$. Moreover, we set the values of $\text{rel}(C_i, C_j)$ manually according to results in preliminary evaluations. Figure 3 shows examples of relevance values assigned to certain class relations.

3.3 Graph-based Ranking

In the final stage of our approach we run a graph-based ranking algorithm on the built instance networks. For each node in the instance network, the algorithm computes a relevance score from the source POI by following a score spreading strategy, and thus provides a scoring of the target nodes (i.e., those belonging to the class *Musician*). Then, nodes with the highest scores are selected for recommendation.

The score of an instance node j depends not only on the relevance values of the instance and its class, which are generic and not specific to a target POI, but also on the scores of the instances that connect the target POI to j along some paths in the semantic network.

Figure 3 shows an example of the score propagation (omitting the instance relevance values) in a part of the instance network created for the Vienna State Opera. 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 technique performs a single iteration for propagating

instance scores in the network. The algorithm propagates the initial score of the source node through its weighted edges, updating the scores of its linked nodes. The score of a node j with n incoming edges from the nodes i_1, \dots, i_n is computed as:

$$score(j) = rel(j) + \sum_n score(i_n) \times w(i_n, j)$$

The score propagation procedure is repeated for subsequent linked nodes until reaching the target nodes, whose scores cannot be further propagated because such nodes do not have outgoing edges. We refer the reader to (Kaminskas et al., 2012) for more details on the network building process and the ranking algorithm.

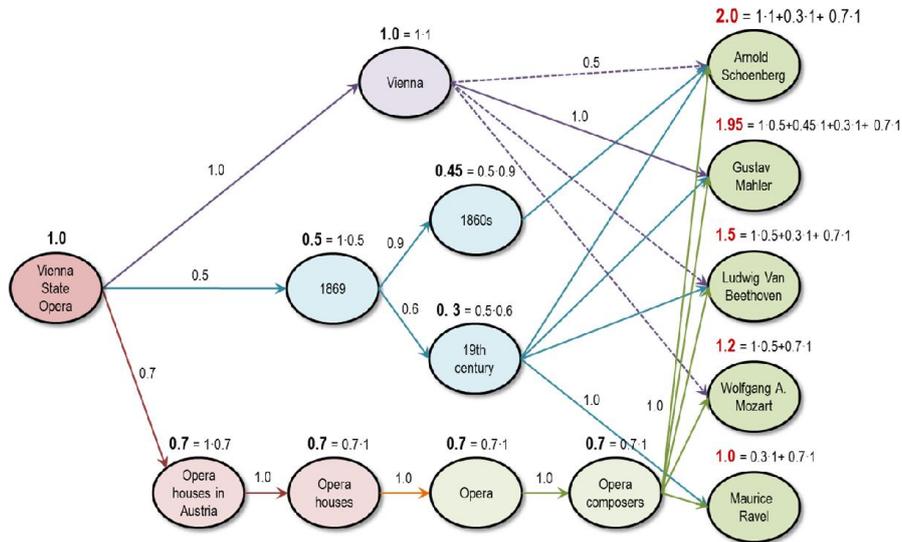


Figure 3. An example of score propagation in an instance sub-network of the Vienna State Opera; for the city paths, solid and dashed edges represent *birth_place_of* and *death_place_of* properties, respectively

4 Experiments

We conducted two user studies aimed to evaluate how users perceive and judge musician-to-POI recommendations provided by our approach. Since users cannot perform a large number of judgments during an evaluation session, we limited the evaluation dataset to 25 POIs (from the full set of 2860 POIs described in Section 3.2). These POIs have instance networks, which on average consist of 708.2 nodes (with 668 nodes representing musicians). For each POI, we obtained the top 5 ranked musicians computed by three different scoring methods: the spreading algorithm that we designed, and two additional algorithms, HITS and PageRank – which already being exploited successfully in many graph-based information retrieval tasks –, can be considered as state of the art simple, but effective baseline methods (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 music composition for each musician by taking the top-ranked result returned by the YouTube² search interface. This was aimed to ensure that the collected music compositions are well representative of the musicians. We note that more compositions for each musician should be considered in a more exhaustive study. In total, we had four competing methods to find a musician (and its music composition) matching any given POI.

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

4.1 Evaluating the 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 (Figure 4) that may require considerable user effort, but lets us 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 user was not aware of what method was used for identifying the musician recommended for a POI. The information describing each POI and musician pair was presented in a structured way, according to the representation in the class network – location, date and art style 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 description were contributing, and in which degree, to the match (right part of the system's interface, Figure 4). We aimed at understanding which types of semantic paths, linking POIs to musicians (i.e., either city, date or art style paths), contributed more to the matches and were better appreciated by the users.

² YouTube video-sharing website, <http://www.youtube.com>

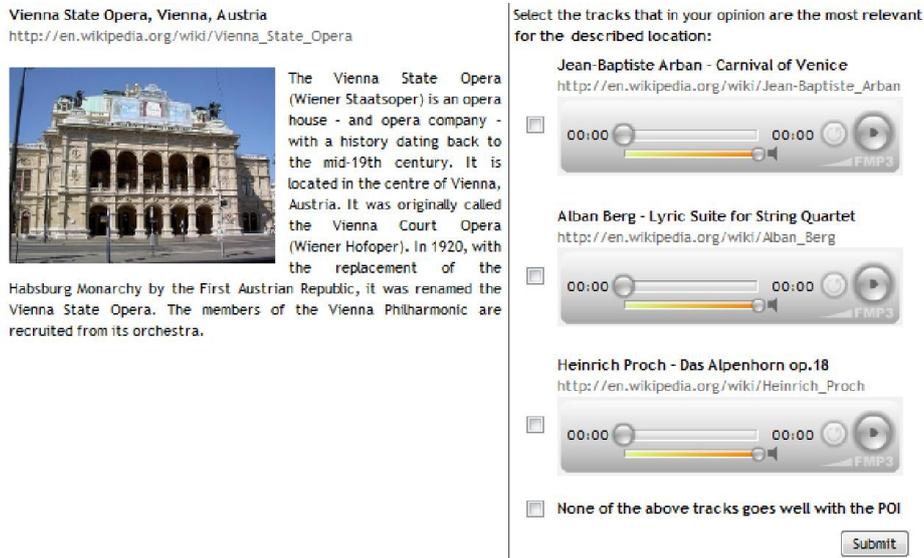


Figure 5. Screenshot of the system used in the matching music to POIs evaluation study

This application has been designed to find an appropriate soundtrack to the displayed Places of Interest (POIs).

Please mark the music genres that you like:

- Classical
- Pop
- Medieval
- Opera
- Rock
- Folk
- Hip Hop
- Electronic

Ok

Figure 6. Screenshot of the interface for collecting the users' genre preferences

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. The users provided 1155 assessments for 356 distinct POI-musician pairs (note that a musician may match the same POI using various ranking algorithms). 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.

In order to measure the performance of the different algorithms, we computed the precision for each algorithm, which is a standard performance metric in information retrieval and recommender systems. Formally, *precision at k* ($P@k$) is defined as the fraction of recommended items that the users judged as relevant among the top-k recommended results:

$$P@k = \frac{\#relevant\ items\ in\ top\ k}{k}$$

Table 1 shows *precision at k* results for all evaluated approaches. The results clearly show that the proposed spreading approach outperforms the baseline approaches.

Table 1. Average precision values obtained for the top 1 to 5 ranked musicians for each POI. The values marked with * are significantly different (Wilcoxon signed-rank test, $p < 0.05$) from the values obtained using the spreading algorithm

	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 recommendations produced by the different approaches, we analyzed the number of semantic paths in the instance network between an input POI and the musicians that each algorithm identified. The higher the number of these paths, 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.

Hence, we computed the average numbers of paths for each algorithm by averaging the path numbers separately for musicians judged as relevant/non-relevant by the users. Based on the obtained results (Figure 7), we observe that the spreading approach differentiates the relevant and non-relevant musicians by finding a larger number of paths between a POI and a relevant musician, compared to a non-relevant one. This indicates that the spreading approach agrees with the users' judgments about the relatedness of POIs and musicians better than the baseline approaches. Another difference worth noticing is that the relevant musicians have on average at least one category path. We also found that for the spreading algorithm the category paths were shorter. In this case, the connections through deeper categories in the architecture and music taxonomies led to longer, less informative paths that do not contribute to the overall relevance.

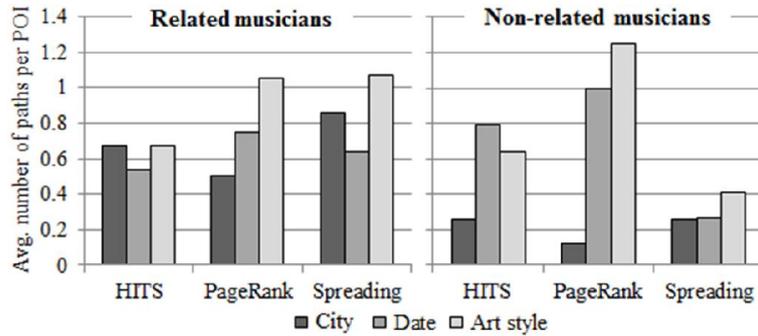


Figure 7. Average number of semantic paths per POI

As an illustrative example of the results reported in Figure 7, Table 2 shows the musicians suggested for the Vienna State Opera with top ranking positions by the three algorithms. These results show that, differently from the baseline approaches, the spreading algorithm tends to use all types of considered relations in a more balanced way.

Table 2. Top-ranked musicians for the Vienna State Opera (opened in 1869)

	Musician	City paths	Date paths	Art style paths
HITS	Jean-Baptiste Arban (French composer; 1825-1889)	0	3	2
PageRank	Carlotta Patti (Italian soprano; 1840-1889)	0	4	2
Spreading	Alban Berg (Austrian composer; 1885-1935)	2	1	4

As an illustrative example, Table 3 shows paths that link Alban Berg with Vienna State Opera. In the table paths are shown as ordered lists of semantic triples subject-property-object.

Table 3. Some paths linking Alban Berg and Vienna State Opera

Path type	Path triples
City	{(Alban_Berg, birth_place, Vienna), (Vienna, location, Vienna_State_Opera)}
City	{(Alban_Berg, death_place, Vienna), (Vienna, location, Vienna_State_Opera)}
Date	{(Alban_Berg, birth_date_century, 19th_century), (19th_century, opening_date_century, Vienna_State_Opera)}
Art style	{(Alban_Berg, music_genre, Opera_music), (Opera_music, super_category, Opera_houses), (Opera_houses, super_category, Opera_houses_in_Austria), (Opera_houses_in_Austria, building_type, Vienna_State_Opera)}
Art style	(Alban_Berg, music_genre, Classical_music), (Classical_music, super_category, Opera_music), (Opera_music, super_category, Opera_houses), (Opera_houses, super_category, Opera_houses_in_Austria), (Opera_houses_in_Austria, building_type, Vienna_State_Opera)}

In addition to analyzing the number of location, time and art style paths obtained for musicians stated as related and non-related with assessed POIs, in the evaluation (Figure 4) we also explicitly asked the users which semantic paths were relevant for each pair of related POI and musician. Figure 8 shows the average relevance values (in a [0, 3] scale) of the different types of paths for poorly related, related, and very related POI-musician pairs. It can be seen that having at least one city path and one art style path is necessary to consider a musician as related or very related with a POI. It also seems that having one or more date paths increases the users' perception of the semantic relatedness between POIs and musicians. Moreover, in general, the relevance of all the types of paths increases with the semantic relatedness degree.

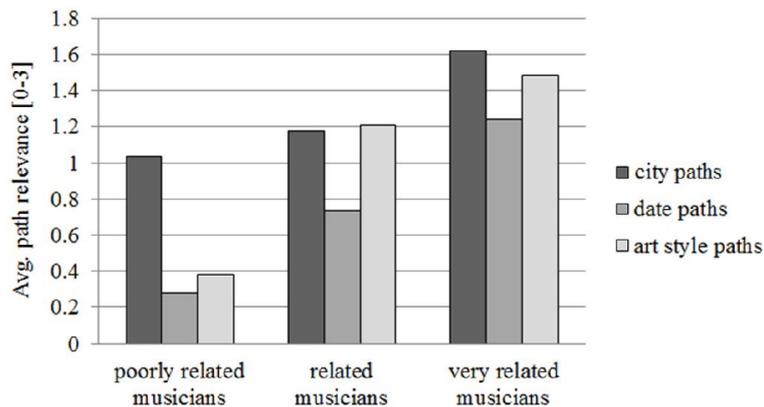


Figure 8. Average relevance (in a [0, 3] scale) of the different types of paths for musicians assessed by users as poor related, related and very related

5.2 Matching Music to POIs: Results

A total of 72 users participated in the second study. As in the first study, they were PhD students and academic staff recruited via email (some of them also participated in the previous study). 1388 evaluation sessions were performed (i.e., a POI shown to a user), and 1298 tracks were selected by the users as well-suited for a POI. Figure 9 shows the performance of the approaches, 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 (i.e., the number of times that a track was also offered to the user as a potential match). All approaches performed significantly better than the random track selection ($p < 0.01$ in a two-proportion z-test). Moreover, again, our spreading activation method outperformed the others ($p < 0.01$).

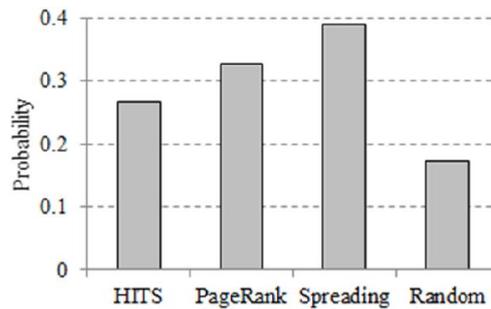


Figure 9. Selection probability of the music recommendation approaches

While in the described evaluation study users received music recommendations that matched the displayed POI, the recommendations were not adjusted to each user's music preferences. Therefore, we expected to observe a variance in the results depending on the users' music tastes. To confirm this, we analyzed 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, $GTrk=g$ is the condition that the genre of a track is g , $g \in uPref$ represents 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 computed, for each genre, the conditional probability for a track to be selected as suited for a POI, given that its genre is g – $P(TrkOK | GTrk=g)$ – 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 study.

As shown in Figure 10, the obtained results show that Medieval and Opera tracks are most often selected by the users (0.49 and 0.47 probabilities), followed by Classical (0.32) and Electronic (0.31) music. Interestingly, Medieval and Opera music were not the preferred genres among the study participants – these genres were liked by 19% and 27% of users respectively. Contrastingly, Rock and Pop music, which have low probabilities to be marked as suited for a POI, were liked by 66% and 63% of users respectively. Therefore, we can conclude that independently of their music preferences,

the users considered certain types of music as more suited for the POIs used in our evaluation.

To check the deviation from the baseline probabilities for each genre produced by the users' music preferences, we measured the probabilities $P(\text{TrkOK} \mid G\text{Trk}=g, g \in u\text{Pref})$ and $P(\text{TrkOK} \mid G\text{Trk}=g, g \notin u\text{Pref})$ by considering only the data provided by users who had/had not included genre g in their preferences. 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.

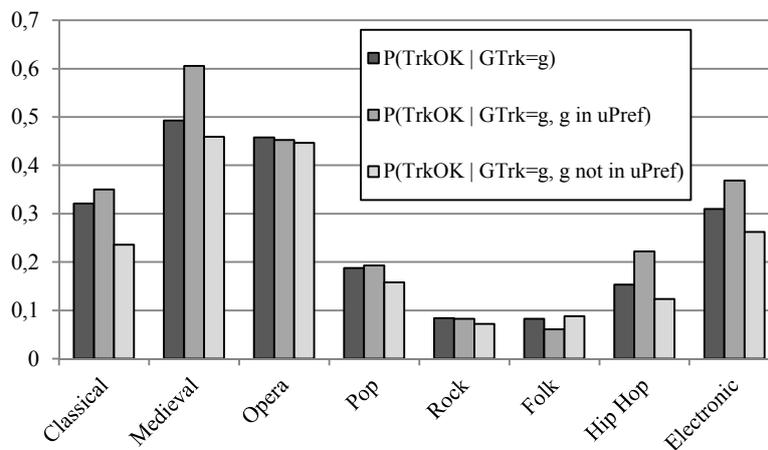


Figure 10. Probability for a track of certain genre to be selected as suited for a POI

We can thus confirm that for certain genres, in addition to the semantic matching between POIs and musicians, there is a clear effect of user preferences on the decision for considering music to go well with a POI: if a user likes a music genre that is generally suited for the POI (in our case, Classical or Medieval music), the probability that she will evaluate that music as suited will increase as well. But, if a music genre is not generally suited for the POIs (in our case, Rock or Pop music), then even if the user likes that music, it does not suffice to make this music type suited for a POI. Based on the results of this initial analysis, in the future we intend to take the users' music preferences into account when suggesting music suited for POIs.

6 Discussion

In this paper we have shown that knowledge-based identification of music for POIs is feasible. However, there are some important limitations that must be mentioned.

First of all, we must stress that, as in many other knowledge-based approaches (Jannach et al. 2010) the technology presented in this article requires a significant amount of expert knowledge and manual fine-tuning of the algorithms. For instance, acquiring musicological knowledge, identifying and measuring the relevance of the knowledge-based relations, eliminating cycles in the knowledge networks, which are all activities required for the implementation of the proposed approach, must be performed by a

domain expert. In the future we will investigate automatic methods to initialize the relevance values of the classes, instances and relations in the semantic networks. For instance, the relevance values of the classes may be set by a concept frequency-based heuristic that computes TF-IDF weights of concepts within Wikipedia's article corpus. The relevance of relations, on the other hand, may be set by a strategy that penalizes relations in long semantic paths between instances, which were found as less meaningful by the users in the conducted experiments.

Moreover, the presented implementation of the proposed framework fully relies on the data available in the DBpedia repository – in order for a musician to be recommended by our approach; its record has to be present in DBpedia. Also, the approach is currently limited to architectural POIs. We intend to address these limitations by exploiting additional Linked Data repositories that would let us use richer POI and music datasets. Furthermore, we intend to work 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 intend to search for these relations in different Linked Data repositories, such as the Europeana EU project's repository³, which contains data on European cultural heritage objects (images, sound recordings, texts). In addition, we intend to use the available tools for discovering semantic relations in the Linked Data cloud, e.g. RelFinder (Heim et al., 2009).

Another important limitation is the absence of personalization in the music selection process. Therefore, an important next step in this research is adapting the framework to incorporate user preferences into the music recommendation process. As mentioned in Section 3.1, this can be achieved by using the weights of relations between classes and instances in the semantic network. We intend to explore personalization approaches that translate user preferences into these weights. Moreover, we will explore other approaches to define and measure the semantic relatedness between concepts (Ponzetto and 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.

Furthermore, we note that there exist different ways of perceiving a match between a place and a music track. For instance, one can enjoy listening to a particular type of music when visiting a POI because the music “feels right” for the place, or one can appreciate a music composed by composer whose life and work are related with the POI. In a location-aware music recommender system, the two approaches can be combined to create an engaging music delivery service. Therefore, we are currently working on combining the ontology-based technique presented in this paper with a tag-based approach we developed in a previous work (Braunhofer et al., 2013). While the first results obtained in a web-based user study look promising, a real-world prototype of the combined recommendation approach remains to be implemented and evaluated in a user study.

³ Europeana's Linked Open Data repository, <http://pro.europeana.eu/linked-open-data>

Finally, we believe that the presented context-aware music recommendation solution may be adapted to other types of content, for instance, recommending music that fits movies, books or paintings.

7 Conclusions

In this paper we have described an original approach for automatically identifying musicians semantically related to a given place of interest (*RQ1*). The approach is based on a knowledge-based framework that is built upon the multi-domain DBpedia ontology, and consists of automatically created semantic networks linking items from POI and music domains. Over these networks a graph-based algorithm ranks and filters items in the target domain (music) with respect to their relatedness to an input item in a source domain (POI), providing thus a novel mechanism for location-aware music recommendation.

Two user studies showed that the musicians retrieved for POIs by the proposed approach are likely to be considered as well-suited for the POIs by the users. We conducted a first user study aimed to evaluate how users perceive and judge the musician recommendations provided by our approach. The evaluation results showed a good performance of our approach (over 80% precision for top-5 results in Table 1), which lets us claim that the approach is able to automatically identify musicians semantically related to a given place of interest. The results also showed that in our approach musicians evaluated as relevant tend to have a high numbers of paths in a POI's semantic network for all the proposed types of relations (i.e., location, date and art style), while these numbers are significantly lower for the musicians judged by users as non-relevant. This clearly implies 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.

Moreover, a second user study – aimed to evaluate how users perceive music recommendations of our approach for the given POIs – showed that users perceive music composed/performed by the recommended musicians as well-suited for the POIs, although for certain genres we found effects of user preferences on the decision in selecting music that goes well with a POI (*RQ2*).

The technology presented in this paper can be exploited in a number of engaging tourism information services, which ultimately motivate the conducted research. For instance, we have developed a mobile city guide providing an enhanced presentation of the place visited by a tourist that plays music that is related to the place, i.e., music emotionally associated to the place (e.g. a fast and happy tune in a lively place like a flea market) (Braunhofer et al. 2013). Other examples include a car entertainment and navigation system that adapts music to the place the car is passing by (Baltrunas et al. 2011), or a tourism website where the information on travel destinations is enhanced through a matching music accompaniment. Such information services can be used to enhance the user's travel experience, to provide rich and engaging cultural information services, and to increase the sales of holiday destinations or music content.

The problem of automatically identifying music for POIs has not been addressed before. Consequently, the solutions proposed in this paper have limitations and leave open research questions for further work. Nevertheless, the obtained results already demonstrate that matching music to POIs is feasible, and that the proposed approach could be used to create new and appealing music delivery services.

Acknowledgements

This work was supported by the Spanish Government (TIN2011-28538-C02) and the Regional Government of Madrid (S2009TIC-1542).

References

- Abowd, G., Atkeson, C., Hong, J., Long, S., Kooper, R., Pinkerton, M. (1997). Cyberguide: a mobile context-aware tour guide. *Wireless Networks* 3(5), pp. 421–433.
- Adomavicius, G., Mobasher, B., Ricci, F., Tuzhilin, A. (2011). Context-aware recommender systems. *AI Magazine* 32(3), pp. 67-80.
- Ankolekar, A., Sandholm, T. (2011). Foxtrot: a soundtrack for where you are. In: *Proceedings of the Interacting with Sound Workshop: Exploring Context-Aware, Local and Social Audio Applications (IwS'11)*, 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: *Proceedings of the 7th 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.
- Baltrunas, L., Kaminskas, M., Ludwig, B., Moling, O., Ricci, F., Aydin, A., Lüke, K. H., Schwaiger, R. (2011) InCarMusic: Context-Aware music recommendations in a car. In: *Proceedings of the 12th International Conference on E-Commerce and Web Technologies (EC-Web'11)*, pp. 89-100.
- Bizer, C., Heath, T., Berners-Lee, T. (2009). Linked data – the story so far. *International Journal on Semantic Web and Information Systems* 5(3), pp. 1-22.
- Braunhofer, M., Kaminskas, M., Ricci, F. (2013) Location-aware music recommendation. *International Journal of Music Information Retrieval* 2(1), pp. 31-44.
- Brotherton, J., Abowd, G., Truong, K. (1999). Supporting capture and access interfaces for informal and opportunistic meetings. Georgia Institute of Technology Technical Report.
- Ciolfi, L. (2004). Interaction design: Enhancing the user experience in interactive environments. PhD thesis, University of Limerick.
- Crestani, F. (1997). Application of spreading activation techniques in information retrieval. *Artificial Intelligence Review* 11, pp. 453-482.
- Dey, A. K. (2000). Providing architectural support for building context-aware applications. PhD thesis, Georgia Institute of Technology.
- Dey, A. K. (2001). Understanding and using context. *Personal and Ubiquitous Computing* 5(1), pp. 4-7.
- Echtner, C. M., Ritchie, J. R. B. (2003). The meaning and measurement of destination image. *The Journal of Tourism Studies* 14(1), pp. 37-48.
- 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 (HetRec'11)*, pp. 25-32.

- Gretzel, U. (2011). Intelligent systems in tourism – a social science perspective. *Tourism Management* 38(3), pp. 757-779.
- Heim, P., Hellmann, S., Lehmann, J., Lohmann, S., Stegemann, T. (2009). RelFinder: Revealing relationships in RDF knowledge bases. In: *Proceedings of the 4th International Conference on Semantic and Digital Media Technologies (SAMT'09)*, pp. 182-187.
- Jannach, D., Zanker, M., Felfernig, A., Friedrich, G. (2010) *Recommender systems: An introduction*. Cambridge University Press.
- Jiang, J. J., Conrath, D. W. (1997). Semantic similarity based on corpus statistics and lexical taxonomy. *CoRR, cmp-lg/970*.
- Kaminskas, M., Fernández-Tobías, I., Ricci, F., Cantador, I. (2012). Knowledge-based music retrieval for places of interest. In: *Proceedings of the 2nd International ACM Workshop on Music Information Retrieval with User-Centered and Multimodal Strategies (MIRUM'12)*, pp. 19-24.
- Kaminskas, M., Fernández-Tobías, I., Cantador, I., Ricci, F. (2013). Ontology-based identification of music for places. In: *Proceedings of the 13th International Conference on Information and Communication Technologies in Tourism (ENTER'13)*, pp. 436-447.
- Kaminskas, M., Ricci, F. (2012). Contextual music information retrieval and recommendation: State of the art and challenges. *Computer Science Review* 6, pp. 89-119.
- Lin, D. (1998). An information-theoretic definition of similarity. In: *Proceedings of the 15th International Conference on Machine Learning (ICML'98)*, pp. 296-304.
- Loizou, A. (2009). How to recommend music to film buffs. PhD thesis, University of Southampton.
- Manning, C. (2008). Introduction to information retrieval. Cambridge University Press.
- Milne, D., Witten, I. H. (2008). An effective, low-cost measure of semantic relatedness obtained from Wikipedia links. In: *Proceedings of the 1st AAAI Workshop on Wikipedia and Artificial Intelligence*.
- Narducci, F., Musto, C., Semeraro, G., Lops, P., de Gemmis, M. (2013). Leveraging encyclopedic knowledge for transparent and serendipitous user profiles. In: *Proceedings of the 21st International Conference on User Modelling, Adaptation, and Personalization (UMAP'13)*, pp. 350-352.
- Passant, A. (2010). Dbrec – music recommendations using DBpedia. In: *Proceedings of the 9th International Semantic Web Conference (ISWC'10)*, pp. 209-224.
- Pedersen T., Patwardhan S., Michelizzi J. (2004). WordNet:: Similarity: Measuring the relatedness of concepts. *Demonstration papers at the 2004 Human Language Technology conference / North American chapter of the Association for Computational Linguistics annual meeting (HLT-NAACL'04)*, pp. 38-41.
- Pike, S. (2002). Destination image analysis: A review of 142 papers from 1973 to 2000. *Tourism Management* 23, pp. 541-549.
- Ponzetto, S., Strube, M. (2007). Knowledge derived from Wikipedia for computing semantic relatedness. *Journal of Artificial Intelligence Research* 30(1), pp. 181-212.
- Rada, R., Mili, H., Bicknell, E., Blettner, M. (1989). Development and application of a metric on semantic nets. *IEEE Transactions on Systems, Man, and Cybernetics* 19(1), 17-30.
- Resnik, P. (1995). Using information content to evaluate semantic similarity in a taxonomy. In: *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI'95)*, pp. 448-453.
- Ricci, F., Rokach, L., Shapira, B., Kantor, P. B. (Eds.) (2011). *Recommender systems handbook*. Springer.
- Seco, N., Veale, T., Hayes, J. (2004). An intrinsic information content metric for semantic similarity in WordNet. In: *Proceedings of the 16th European Conference on Artificial Intelligence (ECAI'04)*, pp. 1089-1090.

- Stupar, A., Michel, S. (2011). Picasso – to sing, you must close your eyes and draw. In: *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'11)*, pp. 715-724.
- Tuan, Y. F. (1977). *Space and place: The perspective of experience*. University of Minnesota Press.
- Weiser, M. (1993). Some computer science issues in ubiquitous computing. *Communications of the ACM* 36(7) pp. 75–84.
- Werthner, H., Ricci, F. (2004) E-commerce and tourism. *Communications of the ACM* 47(12), pp. 101-105.
- Wu, Z., Palmer, M. S. (1994). Verb semantics and lexical selection. In: *Proceedings of the 32nd Annual Meeting of the Association for Computational Linguistics (ACL'94)*, pp. 133-138.
- Zentner, M., Grandjean D., Scherer, K. (2008). Emotions evoked by the sound of music: Characterization, classification, and measurement. *Emotion* 8(4), pp. 494-521.