

Extending Tagging Ontologies with Domain Specific Knowledge

Frederic Font¹, Sergio Oramas¹, György Fazekas², and Xavier Serra¹

¹Music Technology Group, Universitat Pompeu Fabra, Barcelona, Spain

²Centre for Digital Music, Queen Mary University of London, London, UK

{name.surname}@upf.edu, gyorgy.fazekas@eecs.qmul.ac.uk

Abstract. Currently proposed tagging ontologies are mostly focused on the definition of a common schema for representing the agents involved in a tagging process. In this paper we describe preliminary research around the idea of extending tagging ontologies by incorporating some domain specific class definitions and relations. We illustrate our idea with a particular use case where a tag recommendation system is driven by such an ontology. Besides our use case, we believe that such extended tagging ontologies can bring more meaningful structure into folksonomies and improve browsing and organisation functionalities of online platforms relying on tagging systems.

Keywords: Tagging ontology, Tag recommendation, Folksonomy, Freesound

1 Introduction

Tagging systems are extensively used in online sharing sites as a means of content browsing and organisation. In general, tagging systems allow users to annotate resources with free-form textual labels chosen by the users of the system. The resulting set of associations between tags, users and resources that arise in tagging systems is known as a folksonomy. Folksonomies suffer from a number of well-known issues including tag scarcity, ambiguities with synonymy and polysemy, typographical errors, the use of user-specific naming conventions, or even the use of different languages [1]. Despite these issues, folksonomies have succeeded in providing basic organisation and browsing functionalities to online sharing sites. However, their unstructured nature makes it difficult to allow more advanced capabilities such as hierarchical browsing or faceted searching.

In order to bring some structure to folksonomies, some studies have focused on the analysis of folksonomies to automatically derive structured or semi-structured representations of the knowledge of the domain, typically in the form of lightweight ontologies or hierarchical taxonomies [2–4]. However, these methods still tend to require significant amount of manual effort to provide meaningful representations. Some other studies have proposed modelling folksonomies and the tagging process using ontologies [5]. These ontologies are focused on defining a common schema for the agents involved in a tagging process. Current tagging ontologies may enhance interoperability between folksonomies, but do not generally provide ways of structuring a folksonomy with domain-specific knowledge. In this paper, we present

This work is partly supported by the European Research Council FP7-2007-2013 / ERC grant agreement 267583 (CompMusic).

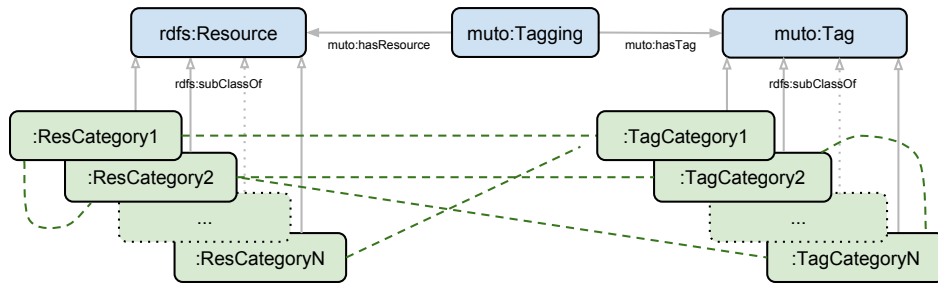


Fig. 1. Diagram of the extended parts of the tagging ontology.

some preliminary research on extending a tagging ontology by including the possibility to represent the semantics of a specific domain. The generic idea is presented and discussed in Sec. 2. In Sec. 3 we describe a practical application in a real-world tagging system where the tagging ontology is used to drive a tag recommendation system. Finally, in Sec. 4, we discuss about possible future directions.

2 Extending a tagging ontology

Our starting point for the extension of the tagging ontology is the Modular Unified Tagging Ontology (MUTO) [5]. In the core of the MUTO ontology, the `muto:Tagging` class is defined which supports several relations to indicate, among others, a resource that is tagged (`muto:hasResource` of type `rdfs:Resource`), the tag assigned to the resource (`muto:hasTag` of type `muto:Tag`), and the user that made the tag assignment (`muto:hasCreator` of type `sioc:UserAccount`).

We propose to extend the tagging ontology in two ways. First, we add a number of subclasses to the `muto:Tag` class which can be used instead of `muto:Tag` (right side of Fig. 1). These subclasses represent different tag categories (i.e. with a narrower scope than the generic `muto:Tag` class), similarly to the idea of `TagSet` introduced in the SCOT ontology [6], but in a semantic sense. A tag category represents a broad concept that groups a set of tags that share some semantic characteristics related to the specific domain. The same principle is applied to resources, and a number of `rdfs:Resource` subclasses are defined (left side of Fig. 1). Resource subclasses (or resource categories) are used to organise resources into groups with a narrower scope than the general `rdfs:Resource` class. The particular definition of tag and resource categories would depend on the particular application domain of the extended tagging ontology (an example is given below). Also, in the diagram of Fig. 1, both tag and resource subclasses are only shown as a flat hierarchy, but more complex class structures could be explored. Moreover, existing domain ontologies and taxonomies may be reused to extend the tagging ontology.

Second, we propose to extend the tagging ontology by adding object properties to model semantic relations among tag categories and resource categories (dashed lines in Fig. 1). These object properties are useful to, for example, model dependencies between categories of tags and resources. The specific meaning of these semantic relations would also depend on the particular application domain of the extended tagging ontology. In addition to semantic relations between tag and resource categories, and given that the `muto:Tag` class inherits from the Simple Knowledge Organization System (SKOS) [7] class `skos:Concept`, semantic relations between tag individuals can be also modelled [5].




Fig. 2. Screenshot of the interface of a prototype tag recommendation system driven by the extended tagging ontology.

3 Use case: tag recommendation in Freesound

We applied an extended tagging ontology as described above in a tag recommendation task in the context of Freesound, an online collaborative sound database with more than 200,000 uploaded sounds and 3,8 million registered users [8]. In previous work by the authors, a tag recommendation system was proposed which, given a set of input tags, is able to suggest other potentially relevant tags [9]. The system is based on the construction of five tag-tag similarity matrices tailored to five manually defined and rather generic audio categories (e.g. “Music”, “Effects”, etc.). The recommendation system uses a classifier to automatically predict one of these five categories depending on the input tags, and then uses the corresponding tag-tag similarity matrix for the recommendation process.

To improve that recommendation system, we used the extended tagging ontology to model the folksonomy and include some domain specific knowledge. On the one side, we extended the tagging ontology by adding 5 resource subclasses corresponding to the 5 sound categories mentioned above (e.g. `:EffectsSound`). Moreover, we defined 26 tag subclasses that are intended to group the tags in categories according to the type of information that they describe about sounds (i.e. grouped in audio properties). These include categories like “instrument”, “microphone”, “chord”, “material”, or “action” (e.g. `:InstrumentTag`). On the other side, we extended the ontology by defining several object properties that relate resource and tag categories. These object properties indicate that a particular tag category is relevant for one or more resource categories. For example, `:InstrumentTag` is relevant for `:MusicSound` audio category, and this is indicated with a `:hasInstrument` object property that relates instrument tag category with music resource category. Furthermore, we populated the extended ontology by manually classifying the 500 most used tags in Freesound into one of the 26 defined tag categories and added these tags as individuals (instances) of the corresponding tag category. This last step was necessary to bootstrap the tag recommendation system (see below).

Using this ontology we can extend the tag recommendation system in a way that, given the audio category detected by the classifier and the object properties that relate resource and tag categories, we can guide the annotation process by suggesting tag categories that are relevant for a particular sound. For example, for a sound belonging to the resource category `:MusicSound`, we can suggest tag categories like `:InstrumentTag` or `:TempoTag`, which are particularly relevant for musical sounds. Once tag categories are suggested, users can click on them and get a list of tag recommendations for every category. This list is obtained by computing the intersection of the tags provided by the aforementioned recommendation system (based on the tag-tag similarity matrix), with those that have been manually introduced in the ontology as tag instances of the selected tag category. See Fig. 2 for an screenshot of a prototype interface for that system.

4 Conclusions

In this paper we have shown some preliminary research on extending current tagging ontologies with structured knowledge specific to the domain of application of a tagging system. By incorporating domain specific knowledge in tagging ontologies, we expect to be able to bring some semantically-meaningful structure into folksonomies. We have illustrated the idea with a use case in the context of an audio clip sharing site where a tag recommendation system is driven by an extended tagging ontology. Formal evaluation of the ontology-driven tag recommendation system is planned for future work. Besides the described use case, we think that using extended tagging ontologies can improve other aspects of online platforms relying on tagging systems such as browsing and organisation functionalities. The main limitation for such improvements is the population of the ontology. In our use case, we use a manually populated ontology to bootstrap the recommender, but the tagging system could further populate the ontology by learning new “tag individuals-tag category” relations when users annotate new sounds. Furthermore, other knowledge extraction techniques could be used to automatically populate the ontology with information coming from other user-generated data (e.g. in our case could be sound comments or textual descriptions), and even from external data sources from linked open data.

References

1. H. Halpin, V. Robu, and H. Shepard, “The dynamics and semantics of collaborative tagging,” in *Proceedings of the 1st Semantic Authoring and Annotation Workshop*, pp. 1–21, 2006.
2. P. Mika, “Ontologies are us: A unified model of social networks and semantics,” *Web Semantics: Science, Services and Agents on the World Wide Web*, vol. 5, pp. 5–15, Mar. 2007.
3. P. Heymann and H. Garcia-Molina, “Collaborative Creation of Communal Hierarchical Taxonomies in Social Tagging Systems,” tech. rep., 2006.
4. F. Limpens, F. L. Gandon, and M. Buffa, “Linking folksonomies and ontologies for supporting knowledge sharing: a state of the art,” 2009.
5. S. Lohmann, P. Díaz, and I. Aedo, “MUTO: the modular unified tagging ontology,” *Proceedings of the 7th International Conference on Semantic Systems - I-Semantics '11*, pp. 95–104, 2011.
6. H. L. Kim, S. Scerri, J. G. Breslin, S. Decker, and H. G. Kim, “The State of the Art in Tag Ontologies : A Semantic Model for Tagging and Folksonomies,” pp. 128–137, 2008.
7. SKOS: Simple Knowledge Organization System. <http://www.w3.org/TR/skos-reference>.
8. F. Font, G. Roma, and X. Serra, “Freesound Technical Demo,” in *Proceedings of the 21st ACM Conference on Multimedia (ACM MM 13)*, pp. 411–412, 2013.
9. F. Font, J. Serrà, and X. Serra, “Class-based tag recommendation and user-based evaluation in online audio clip sharing,” *Journal on Knowledge Based Systems*, 2014, 10.1016/j.knosys.2014.06.003.