

THE ART OF MULTI-FACETED TAGGING

Interweaving Spatial Annotations, Categories, Meaningful URIs and Tags

Fabian Abel, Ricardo Kawase, Daniel Krause, Patrick Siehdnel and Nicole Ullmann
IVS – Semantic Web Group & L3S Research Center, Leibniz University Hannover, Appelstr. 4, 30167 Hannover, Germany

Keywords: Faceted tagging, Spatial tagging, Learning semantics.

Abstract: In this paper we present TagMe!, a tagging and exploration front-end for Flickr images, which enables users to attach tag assignments to a specific area within an image and to categorize tag assignments. We analyze the differences between tags and categories and show how both facets can be applied to learn semantic relations between concepts referenced by tags and categories. TagMe! automatically maps tags and categories to DBpedia URIs to clearly define the meaning. In our experiments we compare different strategies to realize such semantic mappings and show that already lightweight approaches map tags and categories with high precisions (86.85% and 93.77% respectively). We further discuss how multi-faceted tagging helps to improve the retrieval of folksonomy entities. The TagMe! system is currently available at <http://tagme.groupme.org>

1 INTRODUCTION

Tagging systems like Flickr or Delicious enable people to organize and search large item collections by utilizing the Web 2.0 phenomena: Users attach tags to resources and thereby create so-called tag assignments which are valuable metadata. However, imprecise or ambiguous tag assignments can decrease the performance of tagging systems regarding search and retrieval of relevant items.

For example a tag assignment, allotted to an image may only describe a small part of an image and hence cannot be used to derive the overall topic of the image correctly. Some tag assignments are valid for a user-specific point of view, e.g., a tourist would tag an image of a landmark in a different way than a geologist. And finally tag assignments suffer from ambiguity in natural languages.

For disambiguation, approaches like MOAT (Pasant and Laublet, 2008) exist, which support users to attach URIs describing the meaning of a tag to a particular tag assignment analogously to semantic tagging in Faviki¹. A more sophisticated approach, which exploits Wikipedia and WordNet² to detect the meaning of tags, is presented in (Marchetti et al., 2007).

In this paper, we extend the common folkson-

omy model by flexible, contextual tagging facets. We present the TagMe! system that introduces novel tagging facets: Tag assignments are enriched with a DBpedia URI (Auer et al., 2007) to disambiguate the meaning of a tag. So-called *area tags* enable users to tag a specific part of an image (spatial tagging). Furthermore, a *category* dimension is offered to categorize tag assignments.

In the evaluation we show that users appreciate the new tagging features. We present and examine different strategies to automatically map tags and categories to meaningful URIs. Further, we illustrate how the different context facets can be exploited to improve search and learn semantics among tags and categories. For example, we show that the introduced tagging facets are beneficial to identify similar tags and to learn semantic relations between semantic concepts referenced by the tags and DBpedia URIs.

The paper is structured as follows: In Section 2 we introduce the TagMe! system and outline how to integrate tagging facets in the user interface of a tagging system and explain how to extend traditional tagging models to offer additional tagging facets. The benefits of the additional contextual information are evaluated in Section 3. In Section 4 we discuss TagMe! with respect to related tagging systems. Finally, Section 5 summarizes the advantages of the multi-faceted tagging and gives an outlook on future work.

¹<http://faviki.com>

²<http://wordnet.princeton.edu/>

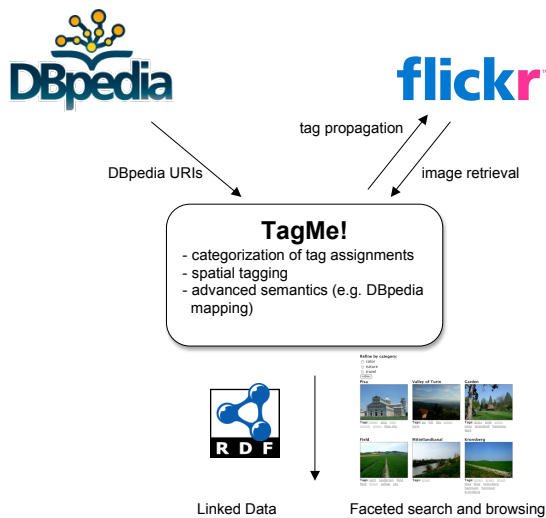


Figure 1: Conceptual architecture of TagMe!

2 TAGME! SYSTEM

TagMe! (Abel et al., 2009) is an online image tagging system where users can assign tags to pictures available in Flickr. Figure 1 outlines the conceptual architecture of TagMe!, which can basically be considered as an advanced tagging and search interface on top of Flickr. Users can directly import pictures from their own Flickr account or utilize the search interface to retrieve Flickr pictures. If users tag their own images in TagMe! then the tags are propagated to Flickr as well. Moreover, TagMe! maps DBpedia URIs to tag and category assignments by exploiting the DBpedia lookup service³ (cf. Section 3.3). Hence, all tags and categories have well-defined semantics so that applications, which operate on TagMe! data, can clearly understand the meaning of the tag and category assignments. The (meta-)data created in TagMe! is made available according to the principles of Linked Data (Berners-Lee, 2007) using the MOAT ontology⁴ and Tag ontology⁵ as primary schemata.

TagMe! extends the Flickr tagging functionality in two further facets, specifically *categories* and *area tags*. For each tag assignment the user can enter one or more categories that classify the annotation. While typing in a category, the users get auto-completion suggestions from the pre-existing categories of the user community (see bottom in Figure 2). TagMe! users can immediately benefit from the categories as TagMe! provides a faceted search interface that allows to refine tag-based search activities by category

³<http://lookup.dbpedia.org>

⁴<http://moat-project.org/ns>

⁵<http://www.holygoat.co.uk/projects/tags>



Figure 2: User tags an area within an image and categorizes the tag assignment with support of the TagMe! system.

(and vice versa). Additionally, users are enabled to perform *spatial tag assignments*, i.e. to attach a tag assignment to a specific area, which they can draw within the picture (see rectangle within the photo in Figure 2) similarly to *notes* in Flickr or annotations in LabelMe (Russell et al., 2008). When tagging, people usually only tag the main content of the picture, giving less or almost none importance to supplementary scenery images.

Area tags motivate the users to do so adding significant semantic value to each annotated image. Moreover, each spatial tag assignment has a globally unique URI and is therewith linkable, which allows users to share the link with others so that they can point their friends and other users directly to a specific part of an image. For example, if users follow the link of the spatial tag assignment “opera”⁶, shown in Figure 2 then they are directed to a page where the corresponding area is highlighted, which might be especially useful in situation where users discuss about specific things within a picture. While the area tags add an enjoyable visible feature for highlighting specific areas of an image and sharing the link to such areas with friends, we consider them as highly valu-

⁶<http://tagme.groupme.org/TagMe/resource/403/tas/1439>

able to improve search by detecting tag correlations or to enhance the identification of similar tags (see Section 3).

To express the introduced enhancements of the TagMe! tagging system in a formal way, current folksonomy models need to be extended.

2.1 Faceted Tagging

Formal models of a folksonomy are e.g. presented in (Halpin et al., 2007; Mika, 2005). They are based on bindings between users, tags, and resources. According to (Hotho et al., 2006) a folksonomy is defined as follows:

Definition 1 (Folksonomy). A *folksonomy* is a quadruple $\mathbb{F} := (U, T, R, Y)$, where:

- U, T, R , are finite sets of instances of *users*, *tags*, and *resources*, respectively, and
- Y defines a relation, the *tag assignment*, between these sets, that is, $Y \subseteq U \times T \times R$.

However, this simple folksonomy model is not sufficient to describe the tag assignments in more detail, i.e. assign context information to a tag assignment. To allow users to create these different facets of a tag assignment, we extend the given folksonomy:

Definition 2 (Faceted Folksonomy). A *faceted folksonomy* is a tuple $\mathbb{F} := (U, T, R, Y, C, Z)$, where:

- U, T, R, C are finite sets of instances of *users*, *tags*, *resources*, and *context-information* respectively,
- Y defines a relation, the *tag assignment* that is, $Y \subseteq U \times T \times R$ and
- Z defines a relation, the *context assignment* that is $Z \subseteq Y \times C$

In the TagMe! system, the *context information* can be a) an area, b) a DBpedia URI or c) a category. All context information are assigned to a tag assignment by a relation Z .

By utilizing the additional information, tag assignments become more connected to each other (see Figure 3). For example, two tags assigned to the same area within an image or having the same DBpedia concept can be considered as synonyms, while two tags that are assigned to different areas in an image are possibly not that strongly related to each other.

3 ANALYSIS AND BENEFITS OF TAGME!

An analysis of the TagMe! data set reveals that the users appreciate the multi-faceted tagging in TagMe!

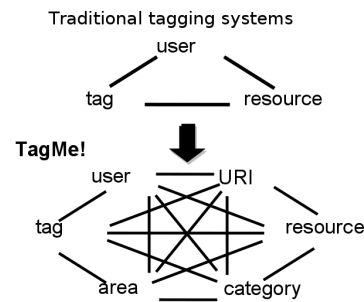


Figure 3: The Faceted Folksonomy in the TagMe! system.

as 874 of the 1295 tag assignments, which were performed within the three weeks after the launch of the system, were categorized and 645 times the users assigned a tag to a specific area within a picture. Given this initial data set, we analyzed the following questions.

- How are categories used in comparison to tags and what are the benefits of categorizing tag assignments?
- What are the benefits of assigning tags to specific areas within an image (spatial tag assignments)?
- How accurate can tags and categories be mapped to DBpedia URIs describing the meaning of the annotations?

3.1 Analysis of Category Usage and Benefits

Figure 4 shows the evolution of the number of distinct tags and categories: Although categories can be entered freely like tags, they grow much less than tags. Further, only 31 of the 79 distinct categories (e.g., “car” or “sea”) have also been used as tags, which means that users seem to use different kinds of concepts for categories and tags respectively.

The TagMe! system supports users in assigning categories by means of auto-completion (see Figure 2). During our evaluation we divided the users into two groups: 50% of the users (*group A*) got only those categories as suggestion, which they themselves used before, while the other 50% of the users (*group B*) got categories as suggestions, which were created by themselves or by another user within their group. This small difference in the functionality had a big impact on the alignment of the categories. The number of distinct categories in group A was growing 61.94% stronger than in group B. Hence, the vocabulary of the categories can be aligned much better if categories, which have been applied by other users, are provided as suggestions as well.

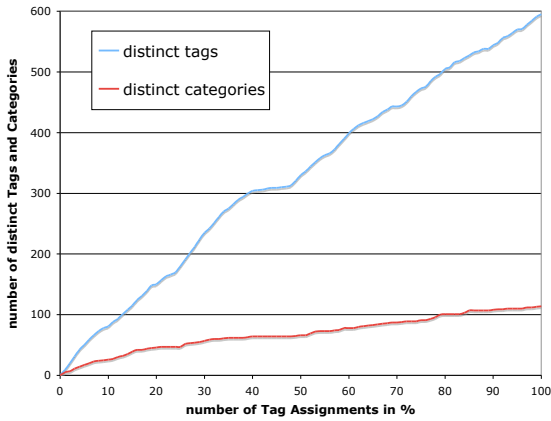


Figure 4: Growth of number of distinct tags in comparison to distinct categories.

Categories also enable to identify similar and related tags, which can, for example, be used for tag recommendations or query expansion. The identification of related tags is often based on tag co-occurrence analysis, e.g. (Sigurbjörnsson and van Zwol, 2008), i.e. two tags are related if they are often assigned to the same resource.

Table 1: Identifying tags related to “clouds”.

Rank	Tag-based	Category-based	Area-based
1	horse	sky	sky
2	sky	field	sun
3	tower	river	cloud
4	field	snow	cross
5	trees	water	sunset

Table 1 lists tags related to the tag “clouds”. Here, the *tag-based* co-occurrence strategy does not perform that well as it also ranks tags such as “horse” or “field” within as the top five most related tags. The *category-based* strategy promotes basically those tags to the top of the ranking that share the most categories with “clouds”. For example, “sky” and “clouds” share categories such as “nature” or “landscape”. In general, the category-based strategy for detecting related tags seems to work better. However, in the given example, it still ranks the rather unrelated tag “field” very high. In our experiments, the best results are produced by the *area-based* strategy, which refines the category-based approach: It ranks those tags higher that occur in spatial tag assignments, whose areas overlap with the areas of the given tag. As shown in Table 1, it also produces—in comparison to the other strategies—the most reasonable ranking of tags related to “clouds”. Four of the top five tags are apparently related (“cross” seems to be the only exception).

From our initial experiments on identifying similar tags, we draw the conclusion that tags, which

share the same category and are often assigned to similar areas within an image (cf. *area-based*), are closer related than tags that often co-occur at same resource. In our future work we will further investigate whether our conclusion holds, especially in larger datasets where categories might introduce noise as they increase the overall connectivity of the folksonomy graph (cf. Figure 3).

3.2 Analysis of Spatial Tag Assignments

Categories can be differentiated according to their usage. For example, some categories have never or very seldomly been used when a specific area of an image was tagged (e.g., “time”, “location”, or “art”) while others have been applied almost only for tagging a specific area (e.g., “people”, “animals”, or “things”).




Figure 5: Annotated areas.

The areas, can moreover be used to learn relations among categories and tags. Figure 5 shows (i) the areas that have been annotated whenever the categories “people” and “friends” have been used (the darker an area the more tags have been assigned to that area). As the areas that have been tagged in both categories strongly correlate and as category “people” was used more often than category “friends” one can deduce that “friends” is possibly a *sub-category* of “people” even if both categories would never co-occur at the same resource. Relations between tags can also deduced by analyzing the tagged areas. Figure 5 shows (ii) the areas that were tagged with “sky”, “clouds”,

“sun”, and “moon”⁷ and via the size and position of the area it is possible to learn that an entity *is part of* or *contained in* another entity (e.g., “sun, moon, and clouds are contained in sky”). The learned relations among tags and categories can moreover be used to learn and refine relations between URIs (ontology concepts) as TagMe! maps tags and categories to DBpedia URIs.

3.3 Mapping to DBpedia URIs

For realizing the feature of mapping tags and categories to DBpedia (Auer et al., 2007) URIs we compared the following two strategies.

DBpedia Lookup. The naive lookup strategy invokes the DBpedia lookup service with the tag/category that should be mapped to a URI as search query. DBpedia ranks the returned URIs similarly to PageRank (Bizer et al., 2009) and our naive mapping strategy simply assigns the top-ranked URI to the tag/category in order to define its meaning.

DBpedia Lookup + Feedback. The advanced mapping strategy is able to consider feedback while selecting an appropriate DBpedia URI. Whenever a tag/category is assigned, for which already a correctly validated DBpedia URI exists in the TagMe! database then that URI is selected. Otherwise the strategy falls back the naive DBpedia Lookup.

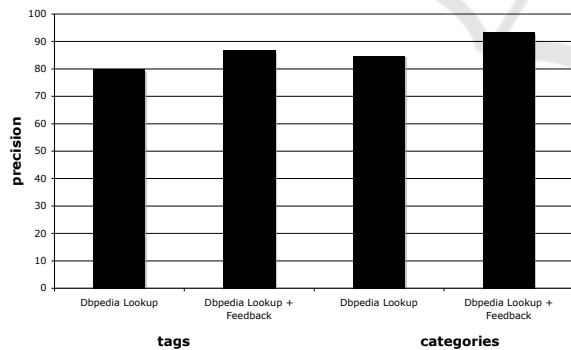


Figure 6: Precision of mapping tags and categories to DBpedia URI.

Figure 6 shows the accuracy of both strategies. The mappings of the naive approach result in a precision of 79.92% for mapping tags to DBpedia URIs and 84.94% for mapping categories. The consideration of feedback improves the precisions of the naive

⁷The visualizations are based on 25 (“sky”), 10 (“clouds”), 6 (“sun”), and 2 (“moon”) tag assignments respectively.

DBpedia Lookup clearly to 86.85% and 93.77% respectively, which corresponds to an improvement of 8.7% and 10.4%. As the mapping accuracy for categories is higher than the one for tags, it seems that the identification of meaningful URIs for categories is easier than for tags. In summary, the results of the DBpedia mapping are very encouraging. Moreover, the precision of the category mappings, which are determined by the strategy that incorporates feedback, will further improve, because—fostered by TagMe!’s category suggestion feature—the number of distinct categories seems to converge (cf. Figure 4). Further, the mapping strategies can be enhanced by also considering the context of the tag/category that should be mapped. For example, for mapping a tag assignment one could select the DBpedia URI, which best fits to the DBpedia URI of the category that is associated to the tag assignment. Implementation of such advanced mapping strategies is part of our future work.

The DBpedia mapping reduces the number of distinct tags and categories within TagMe! by 14.1% and 20.9% respectively, which has a positive impact on the recall when executing tag-based search. For example, while some users assigned the tag “car” to pictures showing cars other users chose “auto” to annotate other pictures that show cars. As both kinds of tag assignments are mapped to “<http://dbpedia.org/resource/Automobile>”, TagMe! can simply search via the DBpedia URI whenever users search via “car” or “auto” to increase recall of the tag-based search operations.

3.4 Synopsis

The two tagging facets, categories and areas, which are applied in TagMe! also have a positive impact on the retrieval of folksonomy entities such as searching for resources or receiving tag recommendations as those facets can be applied to detect correlations between the entities. For example, tag recommendations are usually based on tag co-occurrence, e.g. if different tags are often assigned to same resources then they can be considered as *tag pair* and whenever one of the tags occurs at some resource it is likely that the other tag is relevant for that resource as well. By exploiting the category facet, TagMe! can increase the number of such tag pairs by 367%. Further, the category dimension has potential to compute the similarity of two tags more precisely, e.g. in addition to the (relative) number of times two tags occur at same resources one can consider the (relative) number of times these tags have been used in the same category. The areas of tag assignments can be exploited similarly to refine the correlations between tags. The anal-

ysis of the size, position, and overlap of areas moreover promises to improve the quality of search and ranking.

The results of our analyses can be summarized as follows.

- The usage of categories differs from the usage of tags: Even for those users, who did not benefit from the category suggestions, the number of distinct categories is growing slower than the number of distinct tags.
- Categories are used to further describe and classify tag assignments, and allow the user to solve the problem of ambiguous tags.
- Categories and area tags enhance the connectivity of the folksonomy and provide big potential to improve search or recommender applications (e.g., categories in TagMe! increase the co-occurrence rate of tags 367%).
- For identifying related tags, tag assignments enriched with category and area facets are a more valuable source of information than traditional tag assignments: Tags, which share the same categories and are often assigned to similar areas within an image, are closer related than tags that simply co-occur at same resources.
- The spatial tag assignments can be used to learn typed relations among tags and categories such as *sub-category*, *sub-tag*, *part-of*, or *contained-in* relations. As tags and category assignments are mapped to meaningful URIs (ontological concepts), it is possible to propagate the learned relations to ontologies.
- A naive DBpedia lookup allows us to map tags and categories to ontological concepts (DBpedia URIs) with a high precision of 79.92% (tags) and 84.94% (categories).
- The consideration of feedback improves the accuracy of the mapping of tags and categories to 86.85% and 93.77% respectively.
- The DBpedia mapping reduces the number of distinct tags and categories and can therewith be used to improve recall of tag-based search.

Those findings motivate to exploit the different facets embedded in *faceted folksonomies* (cf. Definition 2.1) such as the TagMe! folksonomy. In our future work we will analyze the impact of those facets on search and ranking.

4 RELATED WORK

The analyses in the previous section revealed several technical advantages of the tagging facets available in the TagMe! system. In this section we compare the tagging and tag-based exploration features of TagMe! from the perspective of the end-users with other tagging systems: Flickr, Delicious, Faviki (Milicic, 2008) and LabelMe (Russell et al., 2008). Our comparison among the systems is partially based on the dimensions of the *tagging system design taxonomy* developed in (Marlow et al., 2006). For example, (i) “Tagging rights”, (ii) “Tagging support” and (iii) “Aggregation model” of those systems. These characteristics define respectively (i) who can tag, (ii) if the user gets assistance from the system during the tagging process and (iii) whether the system allows users to assign the same tag more than once to a particular resource (aggregation model = bag) or not (aggregation model = set).

We extend the tagging design taxonomy with the following additional dimensions related to tagging.

Semantic Tagging. We consider tagging as semantic tagging whenever the meaning of a tag is clearly defined, for example, by attaching a URI explaining the meaning of the tag (Passant and Laublet, 2008).

Spatial Tagging. The practice of annotating a specific piece of a resource, e.g., parts of an image or paragraphs in a text.

Tag Categorization. A method enabling users to categorize or classify the tags and tag assignments.

Further, we introduce two dimensions that characterize to which degree users can exploit the tags to retrieve resources within the system.

Tag-based Navigation. Not all systems that provide tagging functionality also allow their users to explore and browse content based on tags, e.g. initiating search by clicking on a tag.

Faceted Navigation. By faceted navigation, we mean the feature of filtering resources based on the different dimensions of a tag assignment, i.e. by user, tag, or resource, category, or area (cf. Folksonomy model, Section 2.1). For example, in Delicious people can navigate through bookmarks annotated with specific tags (tag dimension) by a specific user (user dimension).

Table 2 summarizes the characteristics of TagMe! and similar tagging systems according to the taxonomy explained above.

The social bookmarking system Faviki and TagMe! are the only systems listed in Table 2 that

Table 2: TagMe! system characteristics in comparison to other social tagging and annotating systems.

Dimension/System	Flickr	Delicious	Faviki	LabelMe	TagMe!
Semantic tagging	no	no	yes	no	yes
Spatial tagging	no	no	no	yes	yes
Tag categorization	no	tag bundles	no	no	tas categorization
Tagging support	viewable	suggested	suggested	viewable	suggested
Tagging rights	permission-based	free-for-all	free-for-all	free-for-all	free-for-all
Aggregation model	set	bag	bag	bag	bag
Tag-based navigation	yes	yes	yes	no	yes
Faceted navigation	yes (user, group)	yes (user)	yes (user)	no	yes (user, category)

allow for semantic tagging. Both systems primarily map tag assignments to DBpedia URIs (Bizer et al., 2009). Faviki requests the end-users to explicitly select the appropriate URIs while TagMe! is doing the mapping automatically. A fundamental restriction of Faviki is that only those tags, which correspond to a meaningful URI, can be assigned to a bookmark. Faviki supports users with a list of URI suggestions from which the users have to select one URI. Delicious and TagMe! provide tagging support by means of auto-completion. Flickr and LabelMe, which is an online annotation tool for images, do not provide tag suggestions but tags already assigned to a resource are *viewable* when adding new tags. In Flickr, users are not allowed to assign the same tag more than once to a particular resource (aggregation model = set) and moreover the owner of a picture has to grant others the permission to tag the picture (tagging rights: permission-based) which results in so-called *narrow folksonomies* (Vander Wal, 2005). In contrast, the other systems listed in Table 2 do not impose these restrictions which allows for *broad folksonomies*.

TagMe! provides two tagging features that are currently not sufficiently implemented in other systems: spatial tagging and tag categorization. Flickr and also MediaWiki⁸ platforms enable users to add notes or comments to specific areas within pictures. However, similarly to LabelMe, which allows users to attach keywords to arbitrarily formed shapes within an image, these systems do not provide means for tag-based navigation based on such spatial annotations, i.e. users cannot click on a spatial tag assignment to navigate to other resources that are related to the corresponding tag (and possibly to the area). TagMe! offers tag-based navigation, which is common in tagging systems such as Flickr and Delicious, also for spatial tag assignments. A further innovation of TagMe! is the tag categorization that is performed on the level of tag assignments (*tas categorization*) and can therewith be used to disambiguate the meaning of a particular tag assignment (cf. Section 3).

⁸<http://www.mediawiki.org>

Delicious, on the contrary, only supports grouping of tags in so-called *tag bundles*. These tag bundles enable users to organize tags but do not help them to disambiguate specific tag assignments. They are moreover seldomly used: Tonkin reports that approx. 10% of the Delicious users have more than five tag bundles (Tonkin, 2006).

The structure of folksonomies (see Section 2.1) can be exploited to navigate through the resource corpus of a tagging system with respect to different facets. For example, when clicking on a tag in Flickr to explore related pictures, users can filter the results to narrow down the results to pictures of a specific *user* or pictures that occur in a specific *group* of pictures. In addition to the feature of browsing resources in context of specific users—as possible in Flickr, Delicious, and Faviki—TagMe! allows such tag-based faceted navigation by applying the categories of tag assignments as filters.

5 CONCLUSIONS

In this paper we discussed multi-faceted tagging in the TagMe! system. TagMe! is a tagging and exploration interface for Flickr and enables users to (1) categorize their tag assignments and (2) attach tag assignments to a specific area within an image. Moreover, all tag assignments are mapped to DBpedia URIs that describe the meaning of the tag. Our analyses reveal that strategies, which exploit categories and spatial tag assignments, provide better results in detecting similar or related tags than naive tag-based co-occurrence strategies. Further, both facets can be exploited to automatically learn new relations among tags and categories (e.g., *contained-in* or *sub-tag*) and therewith also among the corresponding DBpedia URIs. Our feedback-based mapping strategy is able to map tag and category assignments with a precision of more than 85% and 90% respectively to the correct URIs. The DBpedia mapping itself has the potential to increase the precision and recall of search in tagging

systems as it solves the problem of ambiguous as well as synonymous tags. The new tagging facets give the users new means to navigate through images and further allow for advanced search and ranking algorithms.

In our future work we will examine whether it is possible to learn more fine-grained relations by connecting the semantic tags and categories in TagMe! with external domain ontologies. For example, if two objects within the same image are tagged with *person* or *friend* (spatial tagging) one could assume that there is a *foaf:knows* relation between both persons. Further, we will analyze the impact of spatial tagging on search and try to answer whether the size of a tagged area matter or whether the proximity of the tagged area is relevant to the midpoint of the picture. To explore these research questions on larger data sets, we would like to integrate the TagMe! tagging features into an other photo sharing platforms such as Arsmeteo (<http://www.arsmeteo.org>).

REFERENCES

- Abel, F., Kawase, R., Krause, D., and Siehdnel, P. (2009). Multi-faceted Tagging in TagMe! In *8th International Semantic Web Conference (ISWC2009)*.
- Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak, R., and Ives, Z. (2007). DBpedia: A Nucleus for a Web of Open Data. In et al., A., editor, *The Semantic Web, 6th International Semantic Web Conference (ISWC), 2nd Asian Semantic Web Conference (ASWC)*, pages 715–728.
- Berners-Lee, T. (2007). Linked Data - design issues. Technical report, W3C. <http://www.w3.org/DesignIssues/LinkedData.html>.
- Bizer, C., Lehmann, J., Kobilarov, G., Auer, S., Becker, C., Cyganiak, R., and Hellmann, S. (2009). Dbpedia - a crystallization point for the web of data. *Web Semantics: Science, Services and Agents on the World Wide Web*.
- Halpin, H., Robu, V., and Shepherd, H. (2007). The Complex Dynamics of Collaborative Tagging. In *Proc. of 16th Int. World Wide Web Conference (WWW '07)*, pages 211–220, New York, NY, USA. ACM Press.
- Hotho, A., Jäschke, R., Schmitz, C., and Stumme, G. (2006). BibSonomy: A Social Bookmark and Publication Sharing System. In *Proc. First Conceptual Structures Tool Interoperability Workshop*, pages 87–102, Aalborg.
- Marchetti, A., Tesconi, M., Ronzano, F., Rosella, M., and Minutoli, S. (2007). SemKey: A Semantic Collaborative Tagging System. In *Workshop on Tagging and Metadata for Social Information Organization at WWW '07, May 8-12, 2007, Banff, Canada*.
- Marlow, C., Naaman, M., Boyd, D., and Davis, M. (2006). HT06, tagging paper, taxonomy, flickr, academic article, to read. In *Proc. of the 17th Conf. on Hypertext and Hypermedia*, pages 31–40. ACM Press.
- Vander Wal, T. (2005). Explaining and showing broad and narrow folksonomies. http://www.personalinfocloud.com/2005/02/explaining_and_.html.
- Mika, P. (2005). Ontologies Are Us: A unified model of social networks and semantics. In *Proc. Int. Semantic Web Conference (ISWC 2005)*, pages 522–536.
- Milicic, V. (2008). Case study: Semantic tags. W3C Semantic Web Case Studies and Use Cases. <http://www.w3.org/2001/sw/sweo/public/UseCases/Faviki/>.
- Passant, A. and Laublet, P. (2008). Meaning Of A Tag: A collaborative approach to bridge the gap between tagging and Linked Data. In *Proceedings of the WWW 2008 Workshop Linked Data on the Web (LDOW2008), Beijing, China*.
- Russell, B. C., Torralba, A. B., Murphy, K. P., and Freeman, W. T. (2008). LabelMe: A Database and Web-based tool for Image Annotation. *International Journal of Computer Vision*, 77(1-3):157–173.
- Sigurbjörnsson, B. and van Zwol, R. (2008). Flickr tag recommendation based on collective knowledge. In *Proc. of 17th Int. World Wide Web Conference (WWW '08)*, pages 327–336. ACM Press.
- Tonkin, E. (2006). Searching the long tail: Hidden structure in social tagging. *Proceedings of the 17th SIG Classification Research Workshop*.