ANALYSIS OF THE FOLKSONOMY OF FREESOUND

Frederic Font
Music Technology Gorup
Universitat Pompeu Fabra
Barcelona, Spain
frederic.font@upf.edu

Xavier Serra
Music Technology Gorup
Universitat Pompeu Fabra
Barcelona, Spain
xavier.serra@upf.edu

ABSTRACT

User generated content shared in online communities is often described using collaborative tagging systems where users assign labels to content resources. As a result, a folksonomy emerges that relates a number of tags with the resources they label and the users that have used them. In this paper we analyze the folksonomy of Freesound, an online audio clip sharing site which contains more than two million users and 150,000 user-contributed sound samples covering a wide variety of sounds. By following methodologies taken from similar studies, we compute some metrics that characterize the folksonomy both at the global level and at the tag level. In this manner, we are able to better understand the behavior of the folksonomy as a whole, and also obtain some indicators that can be used as metadata for describing tags themselves. We expect that such a methodology for characterizing folksonomies can be useful to support processes such as tag recommendation or automatic annotation of online resources.

1. INTRODUCTION

Web 2.0 has popularized the creation of online communities where users contribute huge amounts of content that is shared across the members and other visitors of the platform supporting the community. The content contributed by users can be of very different nature, from multimedia content such as music, sounds, photos and videos, to user reviews, hyperlinks and any other type of metadata in general. Many online communities have adopted the use of collaborative tagging as a way to describe the information resources and be able to organize and retrieve the content. These systems are of special importance in online communities where users share multimedia content such as sounds, music, photos or videos. In these cases, unless information items are described with some metadata, they can hardly be retrieved using standard text-based queries.

The idea behind collaborative tagging is that users freely associate different labels (tags) with online resources. Generally, users are not constrained by the use of any specific vocabulary or classification system, thus there is no explicit coordination between different users (no explicit agreement on the words to use). The collection of all the tags used by a community of users, along with all the assignments of these tags with the online resources, is called a *folksonomy*, and can be seen as a representation of the knowledge of the community. Figure 1 shows a general diagram of the idea of collaborative tagging.

In this paper we analyze the folksonomy of Freesound [1], an online audio clip sharing site which contains more than two million users and 140,000 user-contributed sound samples covering a wide variety of sounds, from field recordings and sound effects to drum loops and instrument samples. By following methodologies taken from other studies (specially from [2]), we analyze several general aspects that characterize the folksonomy as a whole. We also propose the use of some descriptors for characterizing tags themselves, including tag clustering in groups of domain-specific related concepts and semantic tag classification. Although in this paper we do not perform any further work than characterizing the folksonomy and its tags, our aim is that with that characterization we will be able to "smarten" the folksonomy and better support processes like tag recommendation, automatic tagging or cleaning the folksonomy of inherent tag inconsistencies.

The rest of the paper is organized as follows. In Sec. 2 we review the related work. In Sec. 3 we analyze several aspects of the Freesound folksonomy, both at the general level and at the tag level. In Sec. 4 we conclude the paper with a discussion about our findings and future work.

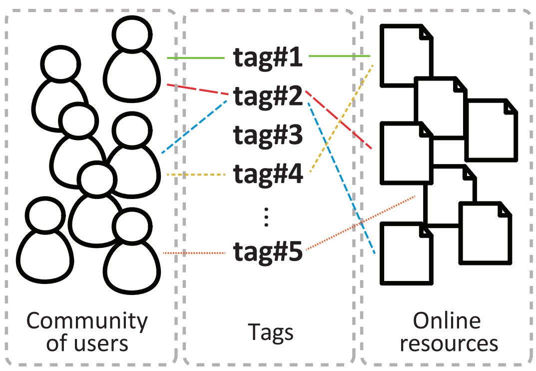

Figure 1. Collaborative tagging scheme. Each line links one user, one tag and one resource. That tripartite link is normally referred as tag "assignment" or "application".

2. RELATED WORK

There are some studies that characterize collaborative tagging systems [2–5]. These studies generally perform qualitative analysis of several collaborative tagging systems based on statistics regarding tag usage and the tagging vocabulary. In the present paper, we take the most relevant measures proposed in [2] and apply them to the folksonomy of Freesound.

Other studies look at the motivations that users have at the moment of tagging, and propose automatic tag classification methods to organize types of tags according to these motivations [6, 7]. We follow the methodology proposed in [7] to perform this step with our data.

Although there aren't many studies focused on the clusterization of tags of a folksonomy (aside from [8–10]), in general any graph-based or similarity-matrix-based clusterization method can be applied for that purpose [11–14]. Actually, in the literature this process is normally referred as "community detection" rather than clustering. However, in this paper we use the term "cluster" to avoid confusion with the concept of the community referring to a group of users of an online platform.

Most of the work done in the analysis of collaborative tagging systems takes as case studies well-known sites such as Delicious (bookmark sharing), CiteULike (scientific reference sharing) and Flickr (photo sharing). This work is, as far as we know, the first that uses tagging data coming from a large-scale audio clip sharing site.

3. ANALYSIS OF THE FREESOUND FOLKSONOMY

In Freesound users can upload sound samples and then describe them with as many tags as they feel appropriate. Since a software upgrade released in September 2011, a minimum of three tags was established for describing a sound. However, the average number of tags per sound has not significantly changed since then. For building the folksonomy we use in our experiments, we considered user annotations between April 2005 and May 2012. As opposite to other well studied collaborative tagging systems such as Delicious or CiteULike, Freesound has what is called a *narrow folksonomy* [15], meaning that sound annotations are shared among all users and therefore one single tag can only be assigned once to a particular sound (e.g. the tag `field-recording` cannot be added twice to the same sound).

The data we analyze comprises a total of 971,561 tag applications performed by 6,802 users to 143,188 resources (sounds), and involving 40,069 distinct tags. The average number of tags per resource is 6.79. Similar averages have been observed in well studied folksonomies with subsets of data coming from Flickr, Bibsonomy and Delicious, with 7.5, 3.66 and 5.63 tags per resource respectively [10]. Figure 2 shows the complementary cumulative distribution function of Freesound tag occurrences. Labels in the low part of the curve correspond to the most used tags. The curve (quite similar the one observed in the analysis of other folksonomies [6]) denotes that a relatively small

#	Tag	Occ.	#	Tag	Occ.
1	field-recording	14954	11	velocity	5468
2	drum	11967	12	bass	5369
3	multisample	11008	13	snare	5261
4	noise	9866	14	drone	4915
5	loop	9015	15	1-shot	4877
6	voice	8320	16	processed	4687
7	ambient	7707	17	soundscape	4619
8	electronic	6671	18	metal	4546
9	synth	6633	19	water	4355
10	percussion	5574	20	ambience	4240

Table 1. 20 most frequent tags in Freesound.

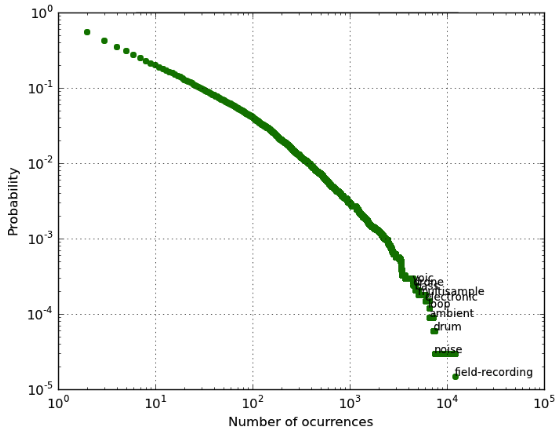

Figure 2. Complementary cumulative distribution function of Freesound tag occurrences.

group of the most used tags involves a big part of the total number of applications. Table 1 shows a relation of the 20 most used tags in Freesound.

The average number of tag applications per user is 142.83. Figure 3 shows a relation of the amount of users that have generated a particular number of tag applications. As it can be seen, the majority of tag applications have been performed by relatively few users. Again, this is a common behavior in other studied folksonomies.

We have computed the correlation between the number of uploaded sounds per user and the number of distinct tags per user (that is to say, the number tag applications per user involving distinct tags). The correlation is 0.51, indicating that the personal vocabulary of every user increases as the number of sounds he has uploaded also increases. This suggests that users feel with the need of using new tags as they upload new samples. A general intuition navigating in Freesound is that users tend to upload sounds of very different nature (except some users that are very specialized), and this might explain that correlation as new uploaded sounds require the use of distinct tags.

3.1 Tag growth

One important characteristic of a folksonomy is the growth of the total number of distinct tags (or vocabulary) that are being used. Figure 4 shows the number of new tags that are introduced every month in Freesound. As it can be seen, there is a slightly positive growing tendency and a sudden increase (approximately doubling the average) starting in September 2011. At that time, a major change in the soft-

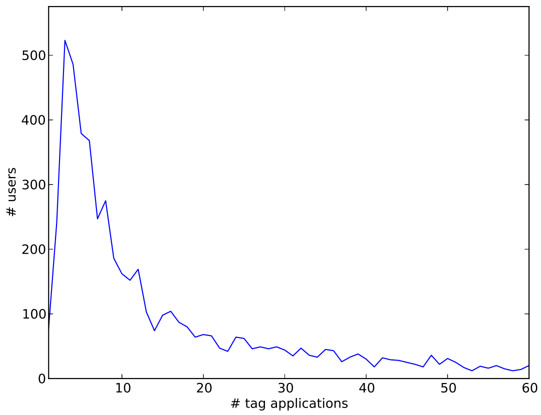

Figure 3. Number of users that have generated a particular number of tag applications.

ware was released with a completely redesigned interface that facilitates uploading and describing sounds. Therefore, this sudden increase in the number of new tags is probably due to this interface update and the increasing popularity that the site has gained since then. Figure 5 shows the cumulative number of new tags and new users per month. As we can see, both increase similarly with an almost perfect linear relationship (correlation is 0.99 if normalizing the two curves by the total number of tags and total number of users). This implies that as more new users are uploading and tagging sounds, more distinct tags are being created. Again, we can see that in September 2011 there is a sudden change in the growing rate of both users and tags.

Such a linear growth of the size of the tag vocabulary without any sign of stabilization or convergence suggests that users are tagging in an isolated fashion, without being really aware of the tags that other people is using to describe their sounds. Furthermore, it is easy to see that there are a lot of tags in the vocabulary which refer to the same concepts but use different string representations (synonymy). This has been observed to be a common problem in folksonomies [3].

3.2 Tag Reuse

The reuse is an important indicator of the collaborative aspects in a tagging system. A big degree of tag reuse means that users are sharing tags and therefore resource descriptions are coherent with respect to other resources. Just as a simple metric, we calculate the percentage of tag applications that correspond to previously used tags as follows:

$$p = 100\frac{M - N}{M},$$

where M is the total number of tag applications and N is the total number of distinct tags. For the Freesound folksonomy we obtain a percentage of 95.88%, which means that the vast majority of tag applications involve already used tags.

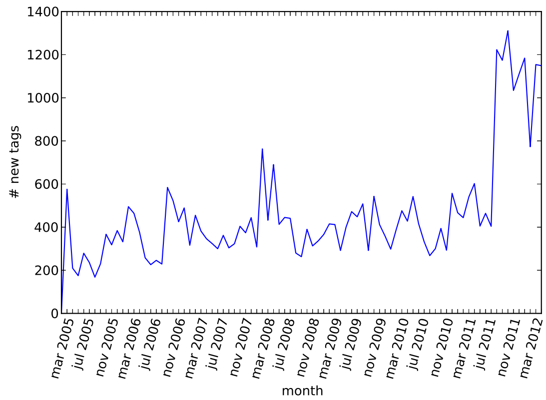

Figure 4. Number of new tags introduced every month.

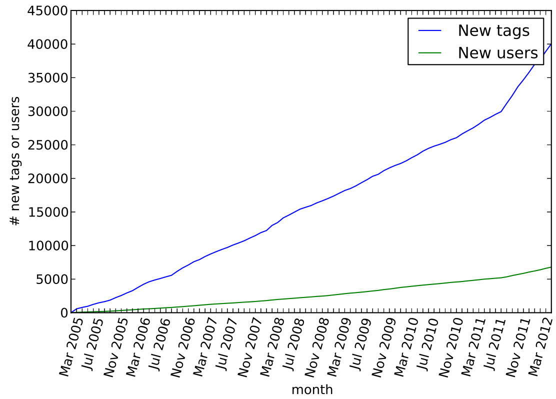

Figure 5. Cumulative number of new tags and new tagging users per month.

Another measure to characterize tag reuse [2, 16] can be computed as follows:

$$r = \frac{\sum U_t}{T},$$

where T is the total number of distinct tags (size of the vocabulary) and U_t is the total number of distinct users that have used tag t. In this way, a value of $r = 1$ indicates that tags are not shared among users across the folksonomy (there is no reuse across users). For the folksonomy we are analyzing we obtained a value of 4.84, which is higher that the values reported for CiteUlike and MovieLens in [2, 16] (1.59 and 1.76 respectively), but still is a very low value considering the upper limit of r which is U (the total number of users).

Figure 6 shows the relation between amount of tags and number of reuse occurrences per tag (that is to say, how many tags have been reused a particular number of times). It can be observed that only a few tags have been reused many times, and the majority have been reused less than 10 times ($\sim$80% of tags have been reused less than 10 times).

To get more insight in how are these tags reused, we looked at the amount of tag reuse from the particular vo-

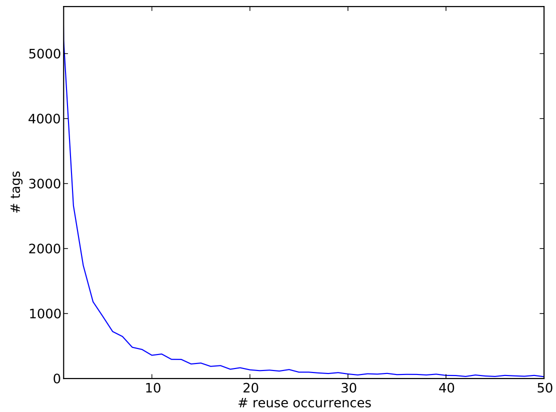

Figure 6. Number of tags that have been reused a particular number of times.

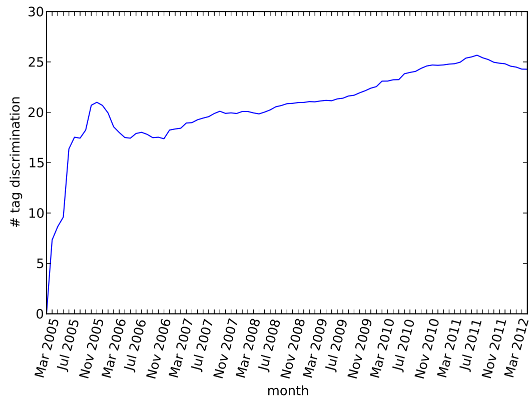

Figure 7. Evolution of tag discrimination for the Freesound folksonomy.

cabulary collections of each user. We computed the following equation:

$$k = \frac{\sum Tr_u}{U},$$

where U is the total number of users and Tr_u is the number of tags from the vocabulary of user u that have been reused (by u). This way we obtain an average of 11.41 tag reuse (more than twice than the average reported for CiteUlike in [2]), indicating a certain tendency that users have of reusing tags from their personal vocabulary.

Bringing all these results together suggests that although almost all tag applications involve reused tags, these tags that are being reused are only a small part of the whole vocabulary (only the most popular ones, tending to be quite generic as it can be seen in table 1). Although users tend to reuse a bit of their tags, they do not take tags from other users more than the few most popular. The less popular tags are probably much more specific and bring detail to the sound descriptions. Therefore, there is no vocabulary sharing among users and no agreement in how to describe the details. This might be expectable given that the tagging interface in Freesound does not reinforce the use of any particular tags nor the vocabulary sharing among users.

3.3 Tag discrimination

Tag discrimination can be understood as the ability of a tag to separate groups of online resources. A simple measure for tag discrimination can be calculated by averaging the number of distinct resources that have been labeled with each tag:

$$d = \frac{\sum R_t}{T},$$

where T is the total number of distinct tags and R_t is total number of distinct resources that have been tagged with tag t. In this way, a tag discrimination value of 1 indicates that all resources have been tagged with different tags, while a value of T means that all resources have been tagged with exactly the same tags. Applying this equation to the folksonomy of Freesound gives an average of 24.25 sounds per tag. That means that, in average, each tag discriminates 24.25 sounds from the rest. This is quite a low value considering that the folksonomy has a total of 143,188 distinct resources. However, it might not be surprising if we look at the tag occurrence distribution of Figure 2, where it is shown that the vast majority of the tags has been used only to label a few sounds.

In Figure 7 we have plotted the evolution of tag discrimination of the Freesound folksonomy. It is interesting to see that after a relatively constant growth during the first almost 6 years, it started getting lower since the previously commented software upgrade done in September 2011. This fact is probably due to the sudden raise in the creation of new tags that we observed in Figures 4 and 5.

We can also calculate a tag discrimination value for a particular tag (d_t) as the fraction of the number of resources tagged with a tag t with respect to the total number of resources. From an information theory point of view, the optimal value would be $d_t = R_t/R = 0.5$ (where R is the total number of resources and R_t is the total number of resources tagged with tag t). That would mean that a tag is able to separate half of the resources from the rest of the collection. If we have a close look to the five tags with more occurrences (listed in Table 1), we can observe that most discriminating tags are `field-recording` (0.103), `drum` (0.083), `multisample` (0.076), `noise` (0.068) and `loop` (0.062). These values are not too low considering the diversity of sounds present in the Freesound database (the most discriminating tag reported in [2] for CiteULike separates 0.0179 of the resources).

All these results suggest that although the majority of tags are only used a few times and that turns into a low tag discrimination average (there is a lot of diversity in the long tail of tags), the most used tags are quite useful to discriminate several regions of the database.

3.4 Tag semantic classification

In this section we follow the methodology proposed in [7] to semantically categorize the tags of the Freesound folksonomy with four categories. These categories indicate the type of information that tags tell us about resources, and

Category	Num.	Examples
Content	17,172	laugh, drum-beat, service-bell, folk-guitar, sitar, nice-music
Context	8,224	playground, mid-night, patagonia, studio-recording, barcelona
Subjective	3,962	oh-may-youre-so-beautiful, psychological, realistic, stressful
Organizational	830	i-love-calculus, sound-of-string, open-air-party, sonsdebarcelona-esther
None	17,772	AKGC1000s, mouvement, 60bpm, pasillo, archestra, grabaciones-de-campo

Table 2. Semantic categorization of Freesound tags.

#	Size	Tags of the cluster
1	777	field-recording, noise, ambient, soundscape, ambience, sound, atmosphere, birds, nature, ambiance, people, wind, talk, recording, car, city, street, engine, speak, woman
2	539	bass, guitar, techno, distortion, distorted, trance, drumloop, chord, bpm, free, delay, multi-sample, synthesis, lead, rock, dubstep, synthesized, dub, clean, hop
3	433	door, footsteps, open, walking, squeak, paper, scratch, household, scrape, floor, steps, walk, UPF-CS12, slide, creak, opening, closing, light, running, concrete
4	372	Synth, Water, Background, Effect, Soundscape, VST, Summer, Echo, Sub, Drum, Bass, Door, pull, Metal, Noise, Field-recording, Field-Recording, Click, FX, Ambient
5	348	kitchen, pop, fire, natural, snap, crack, crunch, crackle, aip09, up, bounce, ding, warm, blow, rubber, body, eating, mouth, bowl, balloon
6	341	train, announcement, station, heavy, bang, rumble, high, automated, road, clang, airport, jingle, rotterdam, stop, thump, ride, subway, passing, railway, steam
7	330	drum, loop, percussion, velocity, snare, 1-shot, metal, water, beat, sample, drums, hit, music, industrial, wood, hard, reverb, weird, dance, echo
8	324	synth, drone, fx, male, acoustic, effect, human, horror, electric, dark, sci-fi, bell, deep, house, synthesizer, computer, metallic, game, cinematic, sound-design
9	305	voices, barcelona, poznan, poland, freesound, image, japan, applause, h4n, seoul, korea, hall, clapping, performance, money, coin, ghent, japanese, desk, coins
10	274	electronic, electro, analog, digital, speech, english, radio, low, samples, beep, wave, tone, circuit, static, fm, plane, pulse, military, army, clip
11	265	click, synthetic, foley, switch, button, effects, soundeffect, strange, granular, press, abstract, dj, vintage, hi-tech, bleep, sounddesign, virus, sweep, ti, funk
12	299	multisample, pad, artificial, evolving, sax, strings, mezzoforte, violin, woodwind, jazz, zoom-h2n, saxophone, 120bpm, divine, non-vibrato, vst, chordophone, ppg, sampled-instruments, classical
13	288	buzz, animal, jungle, ice, south-spain, insects, snow, zoo, animals, tropical, france, waterfall, insect, cricket, exotic, fly, farm, horse, donana, rainforest

Table 3. Most popular tags of the biggest clusters that emerge using the standard modularity optimization technique. "Size" indicates the total number of tags of each cluster.

are: *i) content* (tags that describe the content of the sound such as instruments or sound sources that appear), *2) context* (tags that refer to the location of the recording or the action that generated the sound), *3) subjective* (related to subjective opinions of the users that tagged the resource) and *4) organizational* (tags useful for users personal organization).

To perform this categorization, we first map tags to YAGO [17] concepts. YAGO is an external semantic knowledge base that integrates information from Wikipedia and WordNet, therefore it "knows" about word meanings and relations, and also about world locations and other "facts". If a match is found, YAGO provides the possibility to navigate within semantic concepts of broader sense in a tree-structured fashion until a root category is reached. As proposed in [7], some of the concepts in the higher levels of the hierarchy can be assigned to the *content* and *context* categories (*e.g.* `physical_entity` is assigned to *content* and `location` is assigned to *context*). To maximize the possibility of a tag matching a YAGO concept, we perform a preprocessing step in which tags that are formed by a number of words separated by an hyphen (such as `field-recording`), are split apart and matched separately. The categorization resulting of each part of the tag is aggregated. Therefore, one single tag might be assigned to more than one category. On the other hand, if there are no matches found in the YAGO knowledge base, tags are analyzed using a natural language processing part of speech tagger to assign lexical categories such as "noun", "verb" or "adjective". These lexical categories are compared with a number of pre-defined patterns and if a pattern is matched the tag is assigned to the categories *organizational* or *subjective* (*e.g.* the pattern `[<adjective>]` corresponds to the category *subjective*). For a detailed explanation of the categorization process see [7].

Table 2 shows the number of tags that are categorized in each category along with some examples. As we can see, there are a lot more tags categorized under *content* or *context* than in *subjective* or *organizational*, meaning that they describe aspects of the sounds which are relevant for all users and not only suited to personal classification purposes or opinions. Nevertheless, a lot of tags remain uncategorized (they do not match with any YAGO concept nor with any lexical category pattern) and there are some errors and ambiguities in the categorization (as it can be observed looking at the examples). Some of these tags do not match

#	Size	Tags of the cluster
1	28	overtones, tabla, iran, zarb, hindustani, tambura, carnatic, middle, emotion, sitar, tanpura, bol, indian-classical, compmusic, tonic, raga, kanjira, harmonium, ganjira, eastern [8 more]
2	26	communication, bip, ham, bips, tuner, navigation, radio-static, receiver, telecommunication, interferences, vhf, ham-radio, sw, cb, fm-receiver, vhf-receiver, uhf, uhf-receiver, tv-tuner, cable-tuner [6 more]
3	20	distorted-guitar, guitar-chords, rhythm-guitar, strummed, ukulele, strumming, single-notes, 160bpm, power-chord, miscellaneous, lead-guitar, guitar-notes, uke, extras, drop-d, les-paul, 96khz, ukelele, 01, room-mic
4	20	pipe-organ, carousel, efteling, funfair, wurlitzer, street-organ, live-music, mechanical-music, 200a, e-piano, barrel-organ, parish-fair, annual, leisure, carrousel, parish-fair-organ, hurdy-gurdy, funfair-organ, historic-organ, merry-go-round
5	15	monk, tuva, yoga, undertone, mongolian, puja, tantric, umzie, tuvan, khumi, tantra, gyuto, yogic, kargyraa, sygyt
6	12	threatening, frightening, terrifying, phantom, frightful, shady, grisly, macabre, delusion, spectre, phantasm, imminent
7	7	deathmetal, guitar-riff, death-metal-riff, guitar-tapping, break-down, metal-riff, finger-tapping
8	7	development, blackjack, game-programmers, aplication, tool-kit, sound-set, game-developers
9	7	Step, Footstep, Run, Walk, Stairs, walkway, Hollow
10	5	percussion, bass, snare, beat, drums

Table 4. Most popular tags of the smallest clusters that emerge using the HGC technique.

due to typographical error, the use of words in other languages or for corresponding to too domain-specific concepts such as microphone models and brands. Therefore, these results must only be taken as an estimation, and further work should be needed to produce more accurate semantic categorizations.

3.5 Tag clusterization

The goal of this section is to analyze the Freesound folksonomy and extract clusters of semantically-related tags. For that purpose, we have used two different clustering techniques. Both techniques are based on a graph representation of the tags of a folksonomy, where nodes are tags and edges link similar tags. Similarity between tags is determined by comparing the number of times that two tags are used to label the same sound with their total number of occurrences. For computational complexity reasons, we only consider tags that have been used more than 10 times to build this graph (which are 7,628 of the total). Details on how this graph is extracted can be found in previous work of the authors of this paper [18].

The first clustering technique is a standard modularity optimization of the graph [12], which finds the node partitions that maximize local modularity (that is to say, groups of nodes with dense connections inside the group and sparse connections with nodes from other groups). This clustering technique does not allow node overlapping between clusters, meaning that each particular node can only belong to one cluster. The second clustering technique that we use (hybrid graph-based clusterization [10] or "HGC" for short) allows node overlapping between clusters. It is based on the selection of the most important nodes of the graph that will be the cores of each cluster. In a second step, these cores are expanded by adding similar nodes and, again, maximizing the modularity of the resulting clusters.

Using the standard modularity optimization techniques with the Freesound folksonomy results in the emergence of 59 clusters with an average of 129.29 tags per cluster. Table 3 shows an example of the most popular tags that appear in the biggest of these clusters. As it can be observed, these clusters seem to represent different types of sounds that can be found in the Freesound database at different levels of specificity, but tending to be quite general. For example, clusters 2, 7 and 12 include tags related to musical concepts, and clusters 1, 3, and 5 resemble ambient or field-recording concepts. At a more specific level, cluster 6 includes concepts of recordings done in "traveling" situations and cluster 13 resembles animal sounds.

When using the HGC clustering technique we obtain 561 clusters with an average of 158.44 tags per cluster. Although the average number of tags is quite similar to the standard modularity technique, HGC tends to produce much more smaller clusters (actually, 50% of the output clusters have less than 30 tags). We have seen that the degree of overlap between these clusters is very high, meaning that almost all tags belong to more than one cluster and some of them appear in many clusters. Actually, the biggest clusters tend to emerge more than one time with almost the same tags, meaning that similar tags might have been detected as important nodes and after the expansion step the resulting cluster is almost identical. We have observed that big clusters tend to be similar to the big clusters obtained with standard modularity technique (although a bit more generic). On the other side we find more interesting the emergence of the high number of small clusters, which seem to clearly reflect very specific groups of tags. Table 4 shows some examples of the emergence of small clusters obtained with HGC.

It is not in the scope of this paper to perform any formal evaluation of the clusters that emerge (we leave that to future work), but at first sight it is interesting to see how bigger clusters detected with the standard modularity optimization technique might be useful to form an idea of the different types of sounds that are uploaded in Freesound (at a very general level), and small clusters detected with HGC can, up to some extent, reveal groups of related tags belonging to several particular contexts.

4. CONCLUSIONS

The folksonomy analysis we have described in the previous sections can be useful to better understand how do users tag in Freesound and propose ways to improve the tagging system and thus the sound descriptions. We have observed that the folksonomy of Freesound is continuously growing and there are no signs of stabilization. One of the reasons for this continuous growth might be that new kinds of content are being uploaded that require new concepts to describe them. However, a probably more important reason is that the system does not promote tag reuse nor has any kind of "preferred" vocabulary to push forward. As a result, we find that the folksonomy is quite noisy, and reflects the typical problems also reported in other studies such as synonymy, polysemy and other kinds of inconsistencies. The noisiness of the folksonomy hardens the task of extracting structured information from the folksonomy like semantic classification or tag clusterization. However, we have shown that some techniques already produce interesting results.

A possible solution to help reduce the noisiness of the folksonomy of Freesound would be the inclusion of a tag recommendation system to aid users in the tagging process. Such a system has already been described in previous work of the authors [18], but could be enhanced by taking advantage of the analysis of the present paper. For example, candidate tags for a particular recommendation could be weighed by the tag discrimination values or the popularity of the tag according to the number of different users that use it. Furthermore, recommendations could be aware of the semantic category of the tags being recommended (e.g. recommending tags that belong to different semantic categories) and also take into account related tags according to automatically detected clusters.

Acknowledgments

This work is partially supported by the European Research Council under the European Unions Seventh Framework Program, as part of the CompMusic project (ERC grant agreement 267583), and by the Spanish Ministry of Science and Innovation under the BES-2010-037309 FPI grant for the TIN2009-14247-C02-01 DRIMS project.

5. REFERENCES

[1] V. Akkermans, F. Font, J. Funollet, and B. D. Jong, "Freesound 2: An Improved Platform for Sharing Audio Clips," in *Late-braking demo abstract of the Int. Soc. for Music Information Retrieval Conf.*, 2011.

[2] U. Farooq, T. G. Kannampallil, Y. Song, C. H. Ganoe, J. M. Carroll, and C. L. Giles, "Evaluating Tagging Behavior in Social Bookmarking Systems : Metrics and design heuristics," *Human-Computer Interaction*, vol. 1, pp. 351–360, 2007.

[3] H. Halpin and V. Robu, "The dynamics and semantics of collaborative tagging," in *Proc. of the 1st Semantic Authoring and Annotation Workshop*, 2006, pp. 1–21.

[4] C. Marlow, M. Naaman, M. Davis, and S. Hall, "HT06, Tagging Paper, Taxonomy, Flickr, Academic Article, To Read," in *Proc. of the 17th Conf. on Hypertext and Hypermedia*, 2006, pp. 31–39.

[5] S. a. Golder, "Usage patterns of collaborative tagging systems," *Journal of Information Science*, vol. 32, no. 2, pp. 198–208, Apr. 2006.

[6] K. Bischoff, C. S. Firan, W. Nejdl, and R. Paiu, "Can All Tags be Used for Search ? Categories and Subject Descriptors," in *Proc. of the 17th ACM Conf. on Information and Knowledge Management*, 2008, pp. 203–212.

[7] I. Cantador and I. Konstas, "Categorising social tags to improve folksonomy-based recommendations," *Web Semantics Science Services and Agents on the World Wide Web*, vol. 9, no. 1, pp. 1–15, Mar. 2011.

[8] C. Cattuto, A. Baldassarri, and V. Servedio, "Emergent Community Structure in Social Tagging Systems," *Advances in Complex Systems*, vol. 11, no. 4, pp. 1–13, 2008.

[9] A. Java, A. Joshi, and T. Finin, "Detecting Commmunities via Simultaneous Clustering of Graphs and Folksonomies," in *Proc. of the 10th Workshop on Web Mining and Web Usage Analysis*, 2008.

[10] S. Papadopoulos, Y. Kompatsiaris, and A. Vakali, "A Graph-Based Clustering Scheme for Identifying Related Tags in Folksonomies," in *Proc. of the 12th Int. Conf. on Data Warehousing and Knowledge Discovery*, 2010, pp. 65–76.

[11] B. J. Frey and D. Dueck, "Clustering by passing messages between data points." *Science*, vol. 315, no. 5814, pp. 972–6, Feb. 2007.

[12] V. D. Blondel, J. L. Guillaume, R. Lambiotte, and E. Lefebvre, "Fast unfolding of communities in large networks," *Journal of Statistical Mechanics: Theory and Experiment*, 2008.

[13] Y.-Y. Ahn, J. P. Bagrow, and S. Lehmann, "Link communities reveal multiscale complexity in networks." *Nature*, vol. 466, no. 7307, pp. 761–4, Aug. 2010.

[14] S. Ghosh, P. Kane, and N. Ganguly, "Identifying overlapping communities in folksonomies or tripartite hypergraphs," in *Proc. of the 20th Int. Conf. on World Wide Web*. New York, New York, USA: ACM Press, 2011, p. 39.

[15] T. Vander Wal, "Explaining and showing broad and narrow folksonomies," 2005. [Online]. Available: http://www.personalinfocloud.com/%202005/02/explaining%5C_and%5C_.html

[16] S. Sen, S. Lam, A. Rashid, and D. Cosley, "Tagging, communities, vocabulary, evolution," *Proc. of the 20th Conf. on Community Supported Cooperative Work*, pp. 181–190, 2006.

[17] F. M. Suchanek, G. Kasneci, and G. Weikum, "YAGO: A Core of Semantic Knowledge Unifying WordNet and Wikipedia," in *16th Int.World Wide Web Conf.*, 2007.

[18] F. Font, X. Serra, M. T. Gorup, and U. P. Fabra, "Folksonomy-based tag recommendation for online audio clip sharing," in *13th Int. Soc. for Music Information Retrieval Conf.*, 2012.