

Characterization of the *Freesound* Online Community

Frederic Font, Gerard Roma, Perfecto Herrera, Xavier Serra
Music Technology Group, Universitat Pompeu Fabra
Barcelona, Spain
first_name.last_name@upf.edu

Abstract—There are many online communities with member-generated and openly available multimedia content. Their success depends on having active contributing users and on producing useful content. With this criterion, the community of sound practitioners that has emerged in *Freesound* is a successful case of interest to be studied. But to understand it and support it further we need an appropriate analysis methodology. In this paper we propose some qualitative and quantitative approaches for its characterization, focusing on the analysis of organizational structure, shared goals, user interactions and vocabulary sharing. We think that the proposed approach can be applied to other online communities with similar characteristics.

I. INTRODUCTION

The characterization of online communities is becoming an important research topic for disciplines such as sociology, economics or computer science. But there are many types of communities and of platforms that support them, and it is not feasible to have a single approach for their characterization.

A number of studies have focused on defining and classifying the different types of online communities [1] [2] [3] and they show that researchers tend to categorize the communities according to the variables that are most relevant to their discipline. For example, sociologists tend to focus on the interaction structure, business researchers focus on applications aspects and information systems researchers put the emphasis on the platform perspective.

Freesound¹ is an online sound sample sharing site where a community of sound enthusiasts has emerged. That community has been instrumental in the successful creation of its large repository of sound recordings. Because of this natural evolution, we became interested in studying it, in understanding how knowledge is generated and shared within online communities [4] and in learning how to develop the platform to better support the community [5] [6].

By considering Freesound a community of practice, that is, a group of people who share a concern or a passion for something they do and learn how to do it better as they interact regularly, we have to find an appropriate general definition for the concept of online community. From the definition given in [3], and adapting it to our community of practice context, we will consider an online community as “a persistent group of users of an online social media platform, with shared goals, a specific organizational structure, strong interactions and a

common vocabulary.” From this definition we develop our approach to characterize Freesound. We first give a general description of Freesound and of the methodology used in our characterization. Then we present the analysis of the different community facets: the organizational structure, common goals, user interactions, and vocabulary sharing. We finish with our conclusions on how to use this characterization for improving the platform support of the community. The main contribution of this paper is the methodology we propose, which we think could be applied to characterize other online communities of similar characteristics.

II. OVERVIEW OF FREESOUND

Freesound was launched in April 2005 as an initiative of the Music Technology Group of the Universitat Pompeu Fabra. Its aim was to develop a website for exchanging Creative Commons audio material, supporting the researchers and artists working with audio samples.

After seven years of history, Freesound has become one of the most popular sites for sharing sound samples. It serves around 35,000 visits per day and has more than two million registered users accessing more than 140,000 uploaded sounds. Perhaps more importantly, there is a small number (compared to the total number of users) of highly engaged users that has been continuously contributing to the site, not only uploading sounds but also commenting, rating and discussing in the forums. These are the users that really conform the Freesound community and the ones that we are most interested in studying.

The Freesound platform has evolved over time. After several years of planning, a major software upgrade was done in September 2011 [7] but the basic architecture and functionalities have been maintained from the beginning. The core of the platform is a SQL database with the uploaded sound samples and the descriptions and tags that accompany them. The web front-end provides social interaction features, like a forum, sound comments, sound ratings, and private messaging. The search of sounds can be performed using text queries, content-based similarity (Query by Example) or by browsing tags or geotags. Sophisticated text queries are possible thanks to the tags and the text indexing of all the sound descriptions. A distributed architecture is used for processing the uploaded sounds, producing compressed previews and waveform/spectrogram images, as well as for obtaining the audio

¹www.freesound.org

features used in the search by content similarity. Registered users can download sounds in the original uploaded quality, while non-registered visitors can only listen to the sounds in streaming from the site.

The active contributors in the community have played an important role in the development of the platform, since quite a number of the software improvements have been the result of suggestions from the community or have been discussed within it. Two basic goals have driven most of the development of Freesound: quality and quantity of sounds. We wanted to gather as many user contributed sounds as possible, but not at any price. Sounds have to be of good quality and have to be well described in order to be easily found and reused.

The development of Freesound has been supported by the Music Technology Group through different research projects and with some donations from private companies. By emphasizing its relevance as a research platform and by involving the community of active users, we have been able to maintain Freesound with a very open model and using relatively few resources. The goal is to make it sustainable by continuing in that direction.

We have published some articles on various aspects of Freesound [8] [9] but we have never performed a general characterization of the community with which to identify its strengths and weakness, and with which to identify appropriate future directions for design and development.

III. METHODOLOGY

In this paper we characterize the community by using quantitative and qualitative analysis. Our main source of data has been the SQL database, studying all the data kept there since the creation of the platform in April 2005, until September 2011, time when a major software upgrade was done.

We have classified Freesound users according to the activities they carry on the site (e.g. downloading sounds, uploading sounds, posting in forums...) and analyzed how the number of users in these categories have been changing over time (IV-A). We have also conducted an online survey in the forums to gather qualitative data about the goals of the Freesound community (IV-B).

For the study of user interactions (Sec. IV-C) we have built two networks: one based on user conversations in the forums and the other on the comments that users post on other users sounds. We have then used network analysis techniques to calculate basic network properties that help us understand the nature of these networks and the user interactions they reflect.

In the last part of the community characterization (Sec. IV-D) we have focused on the analysis of sound annotations (i.e. tags that users have assigned to sounds) and the evolution in the usage of tags over time.

IV. COMMUNITY CHARACTERIZATION

In this section we present the results of the different characterization facets that arise from the community definition given in the introduction.

A. Organizational Structure

It is common in online communities to develop some sort of organization involving roles. When the community is united by common interests, these roles can result from the different levels of motivation of users and some practical needs related with the activity. We do not consider the non-registered users (or visitors) part of the community, since despite constituting a big part of the daily visits, these users do not contribute in any relevant way to the site. Within the community we have identified six well defined roles.

- **Regular users** are those that have registered to the site but have never contributed uploading sounds, posting in the forums, nor leaving comments or ratings. Since registering is mandatory for downloading original quality sounds, many users only register for this purpose, and so 89% of registered users belong to this group.
- **Active users** do not upload sounds but participate in forum conversations and leave comments and ratings. According to this definition, 11% of registered users are active users.
- **Contributors** are the ones that upload sounds. Most of these users are also active in the forums and leave comments and ratings to others users sounds. This group represents less than the 1% of the registered users.
- **Moderators** voluntarily participate in the moderation of newly uploaded sounds. They are all contributors and also very active in the community. Currently there are 14 moderators.
- **Developers** maintain the site at the software development level. Since the beginning of Freesound the number of active developers has oscillated between one and six and they are normally active in the community.
- **Community champions.** There is only one user in Freesound that has taken the role of *community champion* [10], being very active in the forums, in contact with users and was also the main software developer for a long time. He is most responsible for the development of a sense of community among Freesound users, therefore, a crucial element of the Freesound success.

The number of regular users started with an exponential growth and evolved into a linear growth during the last years, hitting almost two million users at September 2011. Active users and contributors have had a quite linear growth, reaching almost 260,000 active users and 5,500 contributors by the end of the studied period. In Fig. 1 we can see the ratios between number of users of these groups over time. The ratio between regular users and active users has remained quite constant, meaning that both groups grow every month in the same proportion. However, the ratio between regular users and contributors has been increasing (and also between active users and contributors, though less prominently), meaning that for every new user that becomes a contributor, we have more and more registered and active users every time.

Such a user structure with many more consumers than contributors has also been observed in many other online

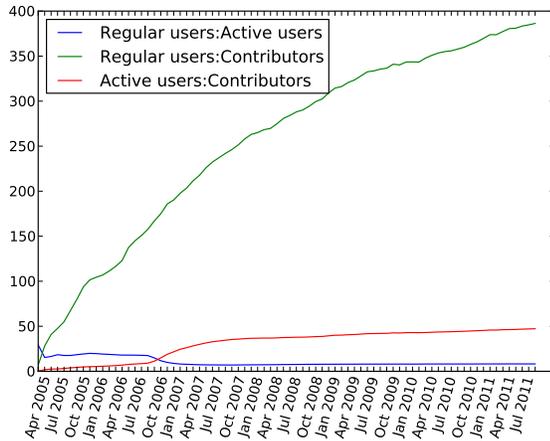


Fig. 1. Time evolution of the ratio between the accumulated number of regular users and active users, regular users and contributors, and between active users and contributors.

communities [11]. One aspect that generates this big difference is that many regular users only register to download some sounds and then never visit Freesound again. Actually 76.9% of regular users have downloaded less than 10 sounds. Another aspect is probably that the knowledge and gear needed for recording or generating sound samples of high quality are not as common as for other typical activities like photo sharing (while nowadays mobile phones can take photographs of acceptable quality, that is not usually the case for sound recordings). Therefore, contribution is almost totally restricted to sound enthusiasts. However, we believe that future designs of the platform should encourage more contributions.

B. Shared Goals

Common goals characterize task-oriented communities, that is to say, communities where members pursue some collective goals that benefit the whole community [3]. We have conducted a small online survey² in the Freesound forums, asking if users perceive the existence of shared goals in the Freesound community and, if so, which are these goals. We have received the responses of 86 participants, 50 of them agreeing with the existence of shared goals. Five users answered that there are no shared goals and the remaining 31 left the question unanswered or did not talk about shared goals. The shared goals that users described in their responses are quite diverse and it does not seem to be a strong common agreement. However, the most repeated goals could be summarized as “sharing sounds” (mentioned by 43.2% of those participants agreeing with the existence of shared goals), “building a big sound archive” (29.5%) and “helping each other by uploading useful sounds” (20.5%).

In another question of the survey, we asked users about the things that make Freesound different from other similar sites.

²Survey can be consulted online at <http://www.freesound.org/forum/freesound-project/32400/>.

In that case, 66.4% of users pointed either at the quantity, quality, diversity or “freeness” of accessible sounds. Other common answers are related with the user interface or the focus on sharing sound samples rather than music (24.1% of the responses). Therefore, although there does not seem to be a common explicit shared goal about creating a big, diverse and free database of quality sounds, it seems that users are implicitly contributing in that direction. This fact fits very well with the original design criteria of Freesound which focused on quantity, quality and royalty free sound samples.

C. User Interactions

Interaction between members is an important factor in real-world communities of practice. In online platforms such as forums or social media sites, it is common to find that users know little about each other. It is said that interactions in online platforms form weak ties in comparison with real-world social networks [12].

Nevertheless, interactions play an important role in online communities such as Freesound. The motivations for users to participate in these communities are typically related to social aspects such as reputation and reciprocity [13]. In our analysis, we focus on explicit interactions, i.e. actions where both parties are aware of each other. The main vehicles for such interactions in Freesound are the forum, and sound comments. The forum serves many purposes, such as promoting work done using Freesound sounds, discussing functionalities of the platform, or requesting specific samples. Comments on sounds, on the other hand, are usually showing support and appreciation. In both cases, participants can be divided between the user that starts the interaction by uploading a sound or creating a thread, and all the users that respond to the sound or the thread.

This form of interaction can be characterized using networks that link the initiator with all the respondents. A similar approach is used in [14] to analyze network-based algorithms for detecting prominent users in expertise (i.e. mostly question and answer) communities. Unlike these networks, in the Freesound forum one cannot assume that respondents have a high level of expertise. Links can be understood as a general measure of shared interest, but the specific purpose depends on the different functions as described. In the case of comments, directionality is always the same, and a directed network can be constructed where top sound uploaders will have large indegree (i.e. large number of incoming links corresponding to the users that have posted comments to their sounds).

Fig. 2a and 2b show the degree distributions for both networks. All of them are heavy tailed, as common in many real world networks, showing that a lot of the interactions are due to a few very active users, while a vast majority interacts occasionally with other users. Assortativity (degree correlation) is negative in both networks (-0.18 in the forum network, -0.2 in the comments network), which shows that users tend to interact with users of different degrees, and do not follow an “elitist” behavior. Nevertheless, a general perception is that the activity of the site is driven by a small group of

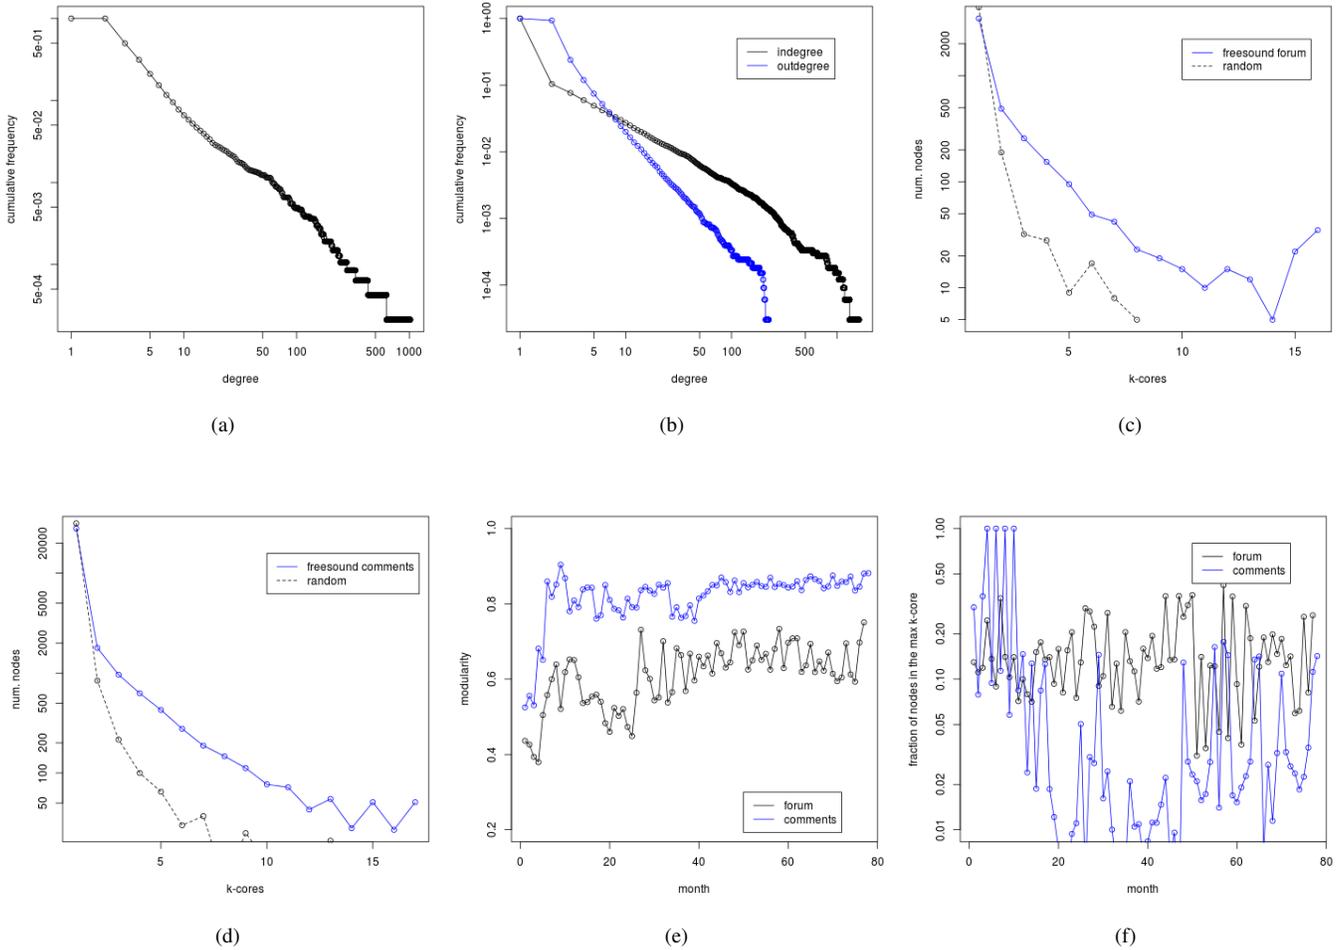


Fig. 2. Forum and comments network properties: (a) Degree distribution of the forum network; (b) Indegree and out degree distribution of the comments network; (c) “Coreness” of the forum network compared with an equivalent random graph; (d) “Coreness” of the comments network compared with an equivalent random graph; (e) Modularity evolution of forum and comments networks; and, (f) “Coreness” evolution of forum and comments networks.

core members (which is supported by the degree distributions). We analyzed this aspect using a k -cores algorithm [15]. K -cores are maximal subgraphs where all nodes have a minimum degree of k [16]. Comparison to a random graph with the exact same degree sequence (i.e. a random graph that has one node with the same degree for each node in the real graph, thus having the same exact degree distribution) shows a trend towards larger cores with higher degree (k) in the real networks than in the random ones (Fig. 2c and 2d). A similar result was found in a study analyzing blog networks [17].

In parallel with the core-periphery structure, both networks exhibit high modularity as measured using the “louvain” method [18] (0.49 for the forum network, 0.57 for the comments network), meaning that user interactions tend to form more densely connected sub groups related with different interests [8].

We analyzed the temporal evolution of modularity and “coreness”, measured as the portion of users that belong to the core with highest k , in networks formed taking into account

interactions within each month (Fig. 2e and 2f). Both networks show a trend towards increased modularity, while the fraction of core members oscillates around the average (0.1 in the forum network, 0.15 for the comments network). In general, our analysis supports the perception that interactions among users of Freesound are constantly lead by a core of very active users, but at the same they tend to form clearly identifiable sub groups, which suggests that further growth could be achieved by supporting sub communities, perhaps each one with a separate active core.

D. Vocabulary Sharing

An online community is also characterized by a common vocabulary that is shared among its members. This vocabulary helps in the understanding of messages being exchanged [3] and in the formation of a community identity [19]. In communities such as the one of Freesound, this vocabulary is very specialized and particular to the domain of discourse being treated (i.e. sound samples).

TABLE I
20 MOST FREQUENT TAGS IN FREESOUND.

#	Tag	Occ.	#	Tag	Occ.
1	field-recording	12556	11	percussion	4726
2	noise	8781	12	drone	4597
3	loop	7993	13	processed	4468
4	drum	7314	14	soundscape	4075
5	voice	7129	15	metal	4028
6	ambient	6719	16	pad	3829
7	electronic	6264	17	water	3698
8	synth	6170	18	vocal	3590
9	multisample	5472	19	glitch	3506
10	bass	4886	20	ambience	3498

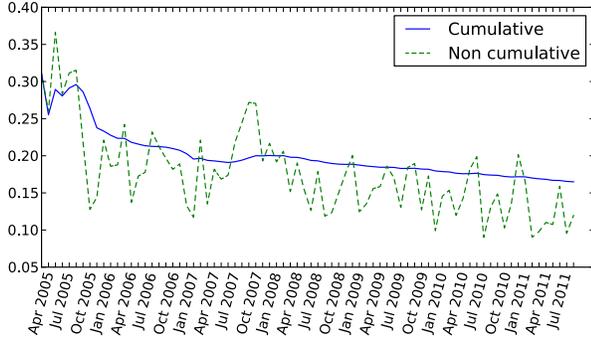
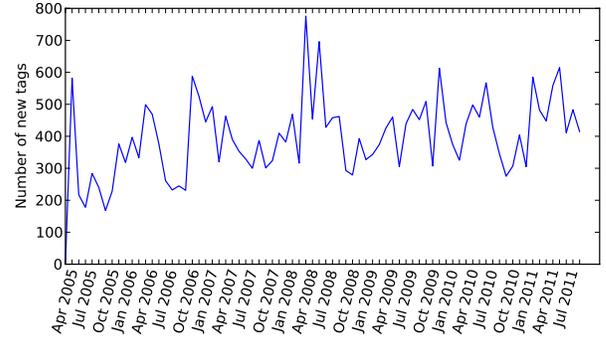


Fig. 3. Usage evolution of the 20 most frequent tags in Freesound.

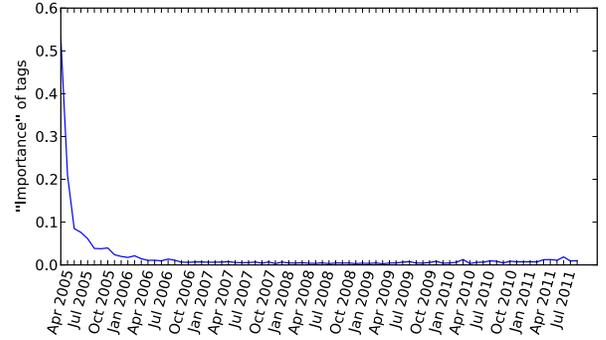
All sounds in Freesound have attached a number of tags which are added by the contributors at the time of uploading. These annotations (or descriptions) reflect the most sound-specific vocabulary used in Freesound. In order to characterize the vocabulary sharing in the Freesound community, we have focused on the analysis of these sound annotations.

Table I shows the 20 most frequent tags in Freesound ordered by their number of occurrences. This list reflects the big diversity of the sounds uploaded to Freesound, with tags referring to various kinds of sounds such as soundscape and outdoor recordings (field-recording, ambience, soundscape, water), music related (loop, drum, bass), voice recordings (voice, vocal) and sound fx (noise, glitch). Fig. 3 shows a measure of the degree of usage for these tags over time. For each sound in the Freesound collection, we have computed the percentage of its assigned tags that appear in the list of 20 most frequent tags. Then, we have averaged the results for all uploaded sounds *a*) during each single month (dashed curve in the plot), and *b*) for all uploaded sounds until each particular month (non-dashed curve). Both curves denote that there is a slow decreasing tendency of the usage of these tags, probably meaning that there is a certain increasing divergence in the usage of tags.

To get more insight in this hypothesis, we have looked at the number of new tags that are introduced every month (Fig. 4a). Despite the noisiness of the curve, we can see a slightly positive tendency which indicates that every month there are more and more newly created tags. This supports the idea that the top tags are becoming a bit less frequent due to the



(a)



(b)

Fig. 4. (a) Number of new tags introduced each month; and, (b) Average “importance” of the new tags introduced each month.

introduction of new tags.

Further analyzing this aspect, we have defined a very simple metric to measure tag “importance” which we calculate using the following equation:

$$importance(tag_i) = \frac{numberOfOccurrences(tag_i)}{longevity(tag_i)}$$

where *numberOfOccurrences* is the total count of annotations using *tag_i* and *longevity* is the number of days that have passed since the first appearance of *tag_i*. By normalizing with tag longevity we compensate the fact that oldest tags might have more occurrences just by being older. Fig. 4b shows the average “importance” of the newly introduced tags in each month. As shown in the figure, after the firsts months there have been no newly introduced tags that have achieved any relevant level of “importance”. What this suggests is that despite the fact that month after month the number of new introduced tags is increasing, these tags are not getting a wide adoption in the community. This evolution compared with the results shown in Fig. 3, supports the hypothesis of the increasing divergence in the vocabulary of sound descriptions, as top tags are losing importance over new tags but these are not achieving high occurrences.

This divergent tendency in the Freesound descriptions vocabulary is probably caused by different reasons. One reason could be that the growth in number of uploaded sounds

makes more difficult for a single user to be aware of all new contributions that are being done by other users. This reduces the visibility of newly introduced tags and after all diminishes the possibility that users get influenced by a common trend in other users' sound annotations. Another reason might be that increasing diversity in types of uploaded sounds requires the usage of new vocabulary that does not get high levels of "importance" because it is not used to label very popular types of sounds.

V. CONCLUSION

An online community like the one of Freesound is quite complex and difficult to describe. In this paper we have attempted to characterize what we believe are some of the fundamental aspects that define the community and with which we can learn how to improve the platform that supports it.

Freesound has been successful at creating a user community despite not being one of the main initial design goals. This community has been a key factor for the quantity and overall good quality of the user generated content. We observe that the core of the community shares the design goals. But the continuous increase in number of users and their activity makes it difficult to maintain the community spirit, mainly because of the large number of users, the growth of the diversity of interests (therefore the variety of uploaded sounds) and the increase of the general system complexity. To face this challenge, we propose some concrete ideas for further developments that should maintain the involvement of the community.

We believe that to maintain and promote the goals of the community they should be more explicit, and we should give more feedback to the users to reinforce the actions that contribute to the shared goals. For example, to reinforce quantity of uploaded sounds, we could place a ranking of contributors prominently in the front page, or reward users with a special community status.

The trend towards the diversification of interests and types of sounds (both suggested in the analysis of user interactions and vocabulary sharing), should be faced with the development of functionalities to support the creation of sub communities. That would allow the development of specialized contexts inside Freesound, facilitating communication between users with similar interests and the development of common specialized vocabularies.

ACKNOWLEDGMENT

We wish to acknowledge the Freesound development team and its whole user community. This research was partly funded 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.

REFERENCES

- [1] C. E. Porter, *A Typology of Virtual Communities: A Multi-Disciplinary Foundation for Future Research*. Journal of Computer-Mediated Communication, 10(1), 2006.
- [2] F. S. L. Lee, D. Vogel and M. Limayem, *Virtual community informatics: A review and research agenda*. Journal of Information Technology, 5(1), 47-61, 2003.
- [3] K. Stanoevska-Slabeva, *Toward a Community-Oriented Design of Internet Platforms*. International Journal of Electronic Commerce, 6(3), 71-95, 2002.
- [4] M. Kosonen, *Knowledge sharing in virtual communities*. PhD thesis, Lappeenranta University of Technology, 2008.
- [5] A. L. Blanchard and M.L. Markus, *Technology and Community Behavior in Online Environments*. Proceedings of the Third Communities and Technologies Conference, 323-350, 2007.
- [6] A. Iriberrri and G. Leroy, *A life-cycle perspective on online community success*. ACM Computing Surveys, 41(2), 1-29, 2009.
- [7] V. Akkermans et al., *Freesound 2: An Improved Platform for Sharing Audio Clips*. Late-breaking demo abstract in the International Society for Music Information Retrieval Conference, 2011.
- [8] G. Roma and P. Herrera, *Community structure in audio clip sharing*. Proceedings of the International Conference on Intelligent Networking and Collaborative Systems, 2010.
- [9] G. Roma et al., *Small world networks and creativity in audio clip sharing*. International Journal of Social Network Mining, 1(1), 112-127, 2012.
- [10] S. Mader, *Wikipatterns: a practical guide to improving productivity and collaboration in your organization*. Wiley Pub., 2008.
- [11] G. Beenen et al., *Using social psychology to motivate contributions to online communities*. Proceedings of the 2004 ACM conference on Computer Supported Cooperative Work, 2005.
- [12] B. Wellman, *An Electronic Group is Virtually a Social Network*. Social Science, 179-205, 1997.
- [13] P. Kollock, *The economies of online cooperation*. Communities in Cyberspace, 259-262, 1999.
- [14] J. Zhang, M.S.Ackerman and L. Adamic, *Expertise networks in online communities*. Proceedings of the 16th International Conference on World Wide Web, 2007.
- [15] V. Batagelj and M. Zaversnik, *An $O(m)$ Algorithm for Cores Decomposition of Networks*. Symposium A Quarterly Journal In Modern Foreign Literatures, 1-10, 2003.
- [16] S.B. Seidman, *Network structure and minimum degree*. Social Networks, 5, 269-287, 1983.
- [17] D. Obradovic and S. Baumann, *A Journey to the Core of the Blogosphere*. Proceedings of the International Conference on Advances in Social Network Analysis and Mining, 1-6, 2009.
- [18] V. D. Blondel et al., *Fast unfolding of communities in large networks*. Journal of Statistical Mechanics, 2008.
- [19] J. Zhang, M.S.Ackerman and L. Adamic, *Language And Adolescent Peer Groups*. Journal of Language and Social Psychology, 22(1), 112-118, 2003.