

ARTIGO

Recebido em:
20/05/2014

Aceito em:
30/11/2014

Encontros Bibli: revista eletrônica de biblioteconomia e ciência da informação, v. 20, n. esp. 1, p. 33-54 Fev., 2015. ISSN 1518-2924. DOI: 10.5007/1518-2924.2015v20nesp1p33

Tagged at first listen: An examination of social tagging practices in a music recommender system

Etiquetagem em primeira audição: uma análise das práticas de etiquetagem social em um sistema de recomendação de música

Audrey LAPLANTE¹

Abstract

Social tagging has become a very common way to index different types of resources on the web. Less prevalent in music than in other domains, social tagging is nevertheless used in a popular recommender system, *Last.fm*. Although the number of publications on tagging and folksonomies has exploded in the last few years, music tagging is still not well studied. In this paper, we present a study of tagging practices of *Last.fm* users. We examine the social tagging of songs during the first three months after their release. Our analysis shows that the release of a song triggers a burst in tagging activity that lasts two weeks, after what it decreases sharply and then remains fairly constant for the next ten weeks. We also find that a majority of songs do not get tagged during the first week and that tagging was positively related to popularity. Finally, we find that tags that have been frequently applied to a given song are more likely to be genre related, shorter in length, and relatively objective than tags that have been applied only once.

Keywords: Social tagging; music indexing; music recommender systems

Resumo

A etiquetagem social (*social tagging*) tornou-se uma forma muito comum de indexar diferentes tipos de recursos na web. Menos predominante na música do que em outros domínios, a etiquetagem social é utilizada em um popular sistema de recomendação, *Last.fm*. Embora o número de publicações sobre atribuição de *tags* (etiquetas) e folksonomia tenha explodido nos últimos anos, a atribuição de *tags* à música permanece pouco estudada. Neste artigo, apresentamos um estudo das práticas de atribuição de *tags* dos usuários do *Last.fm*. Examinamos a etiquetagem social de músicas durante os primeiros três meses após sua liberação. Nossa análise mostra que a liberação de uma música desencadeia uma explosão nas atividades de atribuição de *tags* que dura duas semanas, depois essa atividade diminui de forma acentuada e, então, permanece razoavelmente constante nas próximas 10 semanas. Também verificamos que a maioria das músicas não recebem *tags* durante a primeira semana e que a atribuição de *tags* foi positivamente relacionada à popularidade. Finalmente, constatamos que as *tags* que são frequentemente aplicadas a determinada música são mais relacionadas ao gênero, são menores em extensão e relativamente mais objetivas do que *tags* que tenham sido aplicadas uma única vez.

Palavras-chave: Etiquetagem social; tag; indexação de música; sistemas de recomendação de música



v. 20, n. esp, 2015
p. 33-54
ISSN 1518-2924



Esta obra está licenciada sob uma [Licença Creative Commons](https://creativecommons.org/licenses/by/4.0/).

¹ Associate professor, École de bibliothéconomie et des sciences de l'information, Université de Montréal - audrey.laplante@umontreal.ca

1. INTRODUCTION

Indexing music for retrieval and discovery has become a much more complicated task in the last decade. Although indexing music albums with two or three subject headings in library catalogues or classifying albums by broad genres in music stores is often sufficient to allow users to find what they are looking for, these traditional indexing and classification methods cannot be maintained in current digital music libraries. Indeed, whereas brick-and-mortar music stores typically keep a few thousand unique albums on their shelves, online music retailers and on-demand streaming services give access to millions of tracks (for example, the music streaming service *Spotify*¹ has more than 20 million songs in its catalogue² and the online music retailer *iTunes Store*, more than 37 million songs³). Therefore, broad genres become almost useless since they are not discriminative enough. For instance, a search with “alternative” in *Amazon MP3* retrieves over 2 million songs, hence the need to index music with more specific terms, and at the song level rather than at the album level.

To reach that level of indexing and be able to provide personalized music recommendations to its users, *Pandora*⁴ has taken a rather traditional approach and relies on music analysts to index songs, with up to 450 attributes per song⁵. This approach is obviously time consuming and costly, which translates into a lower growth rate: *Pandora*'s catalogue currently contains just over one million tracks⁶, which is well below most music services. By comparison, another popular recommender service, *Last.fm*⁷, opted for a more scalable and affordable model. It relies on its users for indexing. *Last.fm*'s users are encouraged to assign free-form textual labels (i.e., social tags) to songs, albums, and artists. Although *Last.fm* does not make public the number of songs in its catalogue, it can be estimated to several millions⁸.

The terms “social tagging” and “collaborative tagging” refer to the process by which people assign tags or keywords, whereas the sum of all tags applied by the

¹<http://www.spotify.com>

²<http://press.spotify.com/us/information/>

³<http://www.apple.com/itunes/features/#store>

⁴<http://www.pandora.com>

⁵<http://www.pandora.com/about/mgp>

⁶<http://blog.pandora.com/press/pandora-company-overview.html>

⁷<http://www.last.fm>

⁸*Last.fm* does not reveal the exact number of tracks in its catalogue. However, according to CBS (the owner of *Last.fm*), it had more than 12 million tracks in 2013 (the page was taken down in the Fall of 2013 and can be viewed on Internet Archive at <https://web.archive.org/web/20130913010821/http://cbsimg.com/our-products/last-fm>).

user community of a given system is called a “folksonomy”(TRANT, 2009).Social tagging appeared as a component of the social web, in social bookmarking and photo-sharing services such as *Delicious* and *Flickr*. Although it was primarily meant to help users easily retrieve previously found items, it also provides additional access points for other users, for searching or exploring collections. This indexing method has attracted many criticisms. For instance, compared with indexing terms from controlled vocabulary lists, tags suffer from problems inherent to natural language (e.g., ambiguity due to polysemy, redundancy due to synonymy and spelling variants), which can result in poor precision and recall. Others, however, praised its flexibility and adaptability, qualities that could potentially represent significant benefits for music considering the blurred boundaries between music genres and the constant evolution of the vocabulary used to designate them.

Although collaborative tagging is fairly recent, there is already a large body of literature on this subject. The first tagging systems (i.e., *Delicious*¹and *Flickr*²) have received a lot of attention from researchers. Consequently, the social tagging of websites and images has been the subject of several studies. On the other had, research on the social tagging of music remains scarce: it has been examined in a few studies, but most of these studies are a few years old and are limited in scope. Since previous research shows that tagging behaviour varies depending on the type of resources being tagged (BISCHOFF et al., 2008), it seems important to focus on the tagging of music. In addition, the design of tagging systems continues to evolve, which might affect the way people tag resources in these systems.

In this paper, we examine collaborative tagging in *Last.fm*. Specifically, we analyze the tags applied to songs during the first three months after their release. We look at tag frequency at the song level and examine the nature of the tags that have been applied frequently compared with the tags that have not been reapplied by other users. We also look at the dynamics of tagging through time. We determine, among other things, how long it takes before a song gets its first tag and we examine the nature of these first tags. Finally, we look at the tag distribution by length.

2. RELATED RESEARCH

¹ <http://delicious.com>

² <http://www.flickr.com>

2.1 Social tagging of music

We find in the literature a certain number of articles related to social tags applied to music. Lamere (2008) was probably one of the first to write on the subject. He examined the 500 most popular tags in *Last.fm* and finds that 68 percent are genre related. The second most popular type of tags is what he calls “Locale,” which corresponds to tags related to the place of origin of the artist or group (e.g., “French,” “Seattle”). According to his sample, 12 percent of tags fall into this category. He also notices that the freedom users have when tagging results in concepts being often represented by several variants but insists that this allows for a more nuanced representation that should not be “corrected” by the system.

Bischoff et al.(2008) compare tagging practices in three systems that give access to different types of resources: *Last.fm* (music), *Delicious* (websites), and *Flickr* (images). They examine 300 tags from each system to determine how the nature of the resource affects the kinds of tags applied and their distribution. Instead of examining the most popular tags of each system, they opted for a stratified sample composed of the 100 most popular tags, 100 moderately popular tags (tags starting from 70 percent of probability density), and 100 not very popular tags (tags starting from 90 percent of probability density). They observe important differences in the tag distribution of the three systems. They note that the most prevalent tag category in *Delicious* and *Flickr* is “Topic”(it accounts for almost 60 percent of the tags in the former, and more than 45 percent in the latter¹), whereas “Type,” which includes musical genres, is the most important tag category in *Last.fm*(it accounts for about 50 percent of the tags). *Last.fm* is also the system in which there is the highest proportion of subjective/opinion tags, which the authors attribute to the role music tastes play in the formation of personal identity during young adulthood. Having a stratified sample of tags also allows them to compare the types of tags that are applied very frequently with those applied infrequently. They find that in *Last.fm*, although type/genre tags prevail in all layers of the sample, they are more prominent in the most often applied tags. This explains why their overall proportion of genre-related tags is slightly lower than what is reported by Lamere (2008), who limited his observations to the top tags. Bischoff *et al.*(2008) also notice

¹ Bischoff *et al.* only provides a 3D histogram, without the exact numbers. Therefore, the percentages given here are only approximations.

a sharp increase in the proportion of opinion tags in the long tail compared with the top tags.

Similarly, Zollers (2007) examines social tagging practices in *Last.fm*. She observes an even larger proportion of opinion tags than Bischoff et al. (2008). Her examination of the tags applied to 25 artists and 25 tracks (sampled from 50 randomly selected users) shows that 36 percent of artist-level tags and 67 percent of song-level tags are related to opinion expression.

Geleijnse, Schedl and Knees (2007) analyze tags applied to 1,995 artists in *Last.fm*. They note that a majority of tags (56 percent) had been applied to only one artist in their sample and that less than 2 percent had been applied to 50 artists or more. Levy and Sandler (2009) study social tagging in two music services, *Last.fm* and *MyStrands* (now defunct). They notice that tags are often very discursive: a third of the tags have at least three words, and more than 10 percent contain at least five words. They also observe that tags are unevenly distributed amongst artists, with a small proportion of artists receiving many tags while a large majority of artists receive none or very few tags on average per song. Indeed, their analysis shows that about a third of the 5,265 artists in their sample had never received a tag for any of their tracks and that for the artists who had tagged tracks, a third had no more than 5 unique tags on average per song. Only approximately one fifth of the artists in their sample had more than 30 distinct tags on average per song.

2.2 Social tag frequency distribution

Many tagging systems make part of their data available to the public through APIs. It is therefore sometimes possible to get tagging data (i.e., distinct tags and their frequency) for a given resource, a given user, or an entire database. Several researchers have taken these opportunities to study tag frequency distribution at different levels. Most researchers who examined the distribution of tag frequencies for a given resource (DELLSCHAFT; STAAB, 2008; GOLDBER; HUBERMAN, 2006; GUY; TONKIN, 2006; HALPIN; ROBU; SHEPHERD, 2007) find that the distribution looks like a power law curve, which means that there is usually a small number of tags that are extremely popular at the beginning of the distribution, followed by a sharp drop and a long tail of tags that have been applied very rarely. Levy and Sandler (2009) also detect the same type of distribution when examining tag frequencies at the database level in *Last.fm* and *MyStrands*. As pointed by Peters (2009), however, although this distribution is the most common, it is not the only one. She observes

that for some resources, the distribution of tag frequency rather follows an inverse logistic curve, which also has a long tail but differs at the beginning of the curve: it is characterized by a larger number of frequently applied tags (the 'long trunk') that shows a gently descending slope. Similarly, Kipp and Campbell (2006), who studied the tags assigned to 64 popular bookmarks in *Delicious*, observe a "much gentler slope" at the beginning of the curve before the sharp drop, a description that leads Peters(2009) to conclude that it resembles an inverse logistic distribution.

Different researchers have tried to explain why a few tags are very popular while a vast majority of tags are rarely reproduced. The most common explanation is that imitation plays a major role in the development of these types of distributions: users tend to reinforce tags previously applied by other users rather than adding new tags, a behaviour that would be exacerbated in systems where previously applied tags (usually the most popular tags for the resource being tagged) are recommended to the user when tagging. Dellschaft and Staab (2008)examine social tagging in *Delicious*, which offers such recommendations, and estimate that imitation would occur in as much as 60% to 90% of tag assignments. Golder and Huberman (2006), who also studied tagging in *Delicious*, suggest that users might imitate the choices of previous users because they perceive that selecting tags that have already been assigned by other users is a "safe choice" (a behaviour that would be supported by the theory of social proof), and because it is cognitively less demanding to select a tag that has been suggested by the system. However, the idea that tag suggestions might cause the power law distribution observed in tagging systems has been challenged by a tagging experiment conducted by Bollen and Halpin (2009). The researchers compare the tag frequency distribution in two systems: one that provides tag suggestions and one that does not. They find that a power law distribution emerges in both systems.

In addition to these studies, a few researchers have examined the dynamic aspects of tagging. Golder and Huberman (2006) find that the first tags applied to a website in *Delicious* usually has the highest median rank and tend to be more general in nature than the following tags. They also determine that after the first 100 bookmarks, the relative proportions of tags stabilize. Likewise, Kipp (2010), who focuses on a few very popular bookmarks in *Delicious* to analyze tagging patterns over a four-year period, observes consistency and convergence through time. However, her analysis also shows a slight increase in the proportion of tags that have only been applied once, which leads her to conclude that "in the aggregate

people tend to agree on certain tags which apply to bookmarked items while maintaining a substantial divergence of opinion in the area of personal opinion and differing terminology choice often referred to as the long tail.”

3. SOCIAL TAGGING IN *LAST.FM*

Although all tagging systems have in common to allow users to assign free text labels (tags) to resources to improve their retrieval and discovery, they also differ in several ways, in the nature of the resources to which these systems give access, but also in the way they are designed. Some of these design decisions may affect the way users tag resources. It is therefore important to explain how social tagging works in *Last.fm*. To do so, we use the taxonomy of tagging systems of Marlow et al. (2006), which was developed with the aim of providing researchers with common concepts and vocabulary to describe and analyze such systems. This taxonomy focuses on the system attributes that may affect tagging practices as well as on the incentives users have to tag resources in the system.

Starting with the tagging rights dimension, *Last.fm* is a free-for-all tagging system since it allows anyone to tag any music track, album, or artist. This contrasts with self-tagging systems in which users can only tag the resources they add to the system, such as in *Flickr*. The tags applied to a resource are always public: the most popular tags (up to 60 per resource) are displayed in a cloud, in alphabetic order, and the shade of the background color of each tag is used to reflect the relative frequency of its use (i.e., the most frequent tags are presented in a darker shade of blue) (see Figure 1). Users can click on any tag they come across and explore the resources to which it has been applied, or they can listen to the radio station that is associated with each tag or any combination of two tags. Tags can also be searched (one at a time) with the search engine provided by the system.

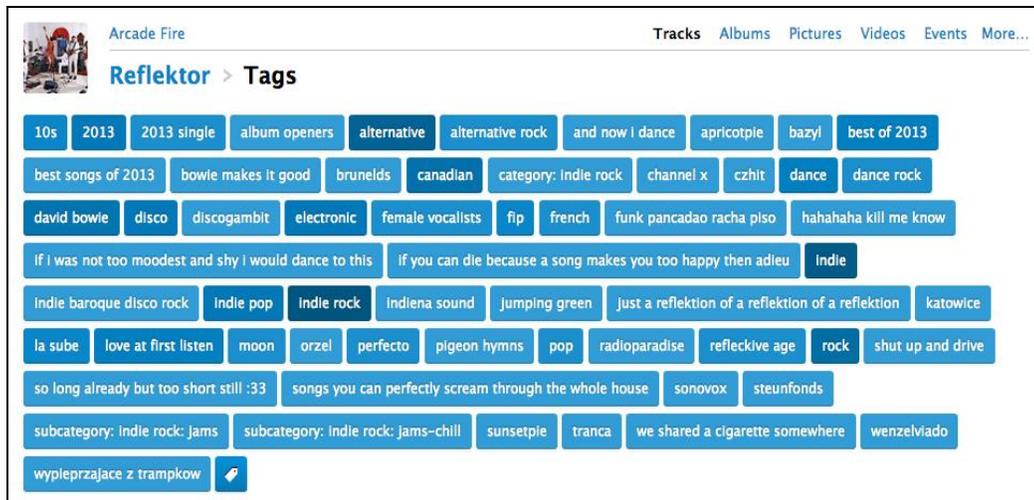


Figure 1. Tags for the song *Reflektor* by Arcade Fire in *Last.fm*

Last.fm is also what Marlow *et al.*(2006) call a “suggestive tagging system.” When a user wants to tag an item, a pop-up window is displayed in which suggestions of tags are presented to the user based on tags assigned by other users for that resource (“popular tags for this track”) and on tags assigned by the user to other resources (“your tags”). The user can either type the tags she wants to add in the text box or select tags amid the proposed tags. This type of support is meant to facilitate the process of tagging by suggesting tags that were deemed appropriate by other users, hence relieving users of having to come up with tag ideas. It is also expected to help consolidate tag usage for a given item (GOLDER; HUBERMAN, 2006; MARLOW *et al.*, 2006).

In terms of social connectivity, *Last.fm* provides networking tools that allow users to connect with each other in different ways. Users can have ‘friends’ and join groups. The system also automatically provides each user with a list of ‘neighbours’ based on their listening habits. This social structure may affect tagging. Groups of people could, for instance, develop a shared vocabulary or set of tags. Moreover, the list of tags a user has assigned appears in her personal profile, with the list of tracks associated with each.

In addition to the system design features included in the taxonomy of tagging system proposed by Marlow *et al.*, a few other attributes that could potentially influence the way people tag resources should be mentioned. *Last.fm* enables users to collaboratively describe tags in wiki pages, which could contribute to building a consensus on the meaning of certain tags. It could also allow people to share the

meaning they ascribe to highly personal tags or to new genres so that others could use them too.

In their taxonomy, Marlow et al.(2006) also include the incentives users have to tag resources in a particular system. Obviously, users' motivation is dependent on their personality and interests, but it is also linked to how the system is designed. In *Last.fm*, users are likely to be enticed to tag for future retrieval (or self-organization) since they can access their personal tags and the associated tracks. For example, a user could apply the tag "driving music" every time he hears a song he would like to listen to in his car. This would allow him to retrieve from his profile the complete list of songs to which he had applied this tag. Also, since tags are displayed in their profile, users might want to tag resources to mark their identity, or for "performance" as Zollers (2007) calls it. For instance, users might apply tags that convey their opinion about songs, albums, or artists, or tags that provide information about their relationship with the resource, such as "seen live." Finally, since tags are public and can be used for searching or browsing, users can also be motivated by altruistic considerations and engage in tagging with the objective of improving access for other users.

4. DATASET

Different methods were used to collect social tags in *Last.fm*. To be able to study the dynamics of social tagging, we wanted tagging data for individual songs for each day during the first three months after their release. These data were provided by the *Last.fm* team. Since *Last.fm* does not include the release date of a song in its database nor the date a song is added to the site, we needed to find another way to identify newly released songs. To do so, we consulted the "Reviews" section of *Billboard*, a trade journal of the music industry, during a two months period. This magazine was selected for its rather broad coverage (it includes about 12 album reviews per weekly issue) and provides the release date for each album. As we were interested in examining tags at the song level, we used *MusicBrainz*¹ to obtain the track lists of the albums. We retained the first four tracks of each album, unless the track length was less than 60 seconds (which often indicates a bridge or an introduction), for a total of 296 songs. The *Last.fm* team could provide tagging data for 288 songs. Songs that had been tagged more than seven days before the release

¹<http://musicbrainz.org>

of the album in which they appeared were removed from the sample. A verification of a few of these songs showed that they had been previously released as a single or, in some cases, the whole album had been made available on the Web for free for a week, before its official release. The final dataset was therefore composed of 184 songs, for which we had tagging data for each day of the three-month period following their release, for a total of 3,655 tag assignments and 1,814 distinct combinations tag/song. The analysis was also sometimes limited to the songs that had never been tagged before the official release (for instance to examine how long it takes before a song get tagged for the first time), which corresponded to 102 songs. Additional data about the popularity of each song the week after their release was found on the *Last.fm* site. Finally, the *Last.fm*API² was used to retrieve additional data, such as the most frequently applied tags and a list of the most popular songs.

5. FINDINGS

5.1 Distribution of social tags by type

Before examining the social tags of our dataset of 184 songs, we first extracted and examined the 250 most popular tags in *Last.fm* (as of August 8, 2014). These are the tags *Last.fm* users have assigned the most frequently to songs, albums, or artists. Each tag was manually classified by type. A few tags ($n = 21$) were placed in more than one category since they were multifaceted. For example, “britpop” was classified in both *Genre* and *Place of origin/language*, and “synthpop,” in *Genre* and *Instrumentation*. Table 1 shows the distribution of tags by type. Our results are very close to those obtained by Lamere (2008) who examined the top 500 tags in *Last.fm*. This suggests that tagging practices remain fairly stable across time. Our analysis shows that 64 percent of the tags are genre related, compared with 68 percent in Lamere (2008). In both studies, the second most popular category is *Place of origin/language* (which Lamere calls “Locale”), which accounts to 13 percent of the tags, compared with 12 percent in Lamere (2008). However, our analysis revealed a slightly higher proportion of tags for the next most popular categories, with 9 percent for *Instrumentation* tags (compared with 4 percent), 7 percent for *Opinion* tags (compared with 4 percent), and 6 percent for *Mood/Emotion* tags (compared with 5 percent).

²<http://www.lastfm.fr/api>

As revealed by previous research (BISCHOFF et al., 2008; HALPIN; ROBU; SHEPHERD, 2007), the distribution of tags by type varies along the power law curve. Therefore, with the aim of examining all tags and not only the most popular tags, we looked at the tags that had been assigned to the 184 songs in our sample. We randomly selected 181 tags (10 percent) from the 1,814 distinct combinations tag/song of the sample. The distribution is notably different from the one observed in the most popular tags (see Table 1 for details). The proportion of genre-related tags is lower at 40 percent. The analysis also shows a much higher proportion of tags in the *Opinion* (15 percent) and the *Mood/Emotion* (10 percent) categories. We also find a much higher number of tags (close to 9 percent) that did not fit in any of the main categories, most often because it was not possible to interpret what they meant (e.g., 332, AtS, MOONH8SUN). Finally, there is a significant proportion (more than 6 percent) of tags that consist in bibliographic information about a particular song, album, or artist (e.g., title of the song, artist name, music label), which, for obvious reasons, are not present in the most popular tags applied in *Last.fm*.

To examine how the distribution changes along the power law curve in *Last.fm*, we looked more closely at the tags that had been applied several times to a given song and those that had been applied only once. Among the 1,814 distinct combinations tag/song of our sample, 1,199 or 66 percent had been applied only once to a given track. We randomly selected 120 (10 percent) of these tags to study their distribution by type (see Table 1). Only slightly over a quarter of them (28 percent) are genre related. We also notice an important proportion of *Opinion* tags (19 percent) such as “best songs ever,” “magic,” or “i want to marry this song.” Nine percent refer to the mood or emotion of the song (“happy,” “sadness and the city”), and close to 7 percent fall under the *Self-reference* category, which includes tags that refer to personal experiences (e.g., “reminds me of you,” “my timeline”) or serve as a personal reminder (e.g., “download,” “albums i own”). Eight percent were classified under the *Other/Indecipherable* category, with tags such as “aitch,” or “my gang rotw.” The *Context/Usage* category accounts for 2.5 percent of the tags, including the very specific “songs i like to play whilst walking down the street at night.” It is also worth mentioning that 6 percent of the tags are in a language other than English.

The 42 tags that had been applied 10 times or more to a given track were also examined. The distribution by type was clearly different (see Table 1): 33 (79 percent) of them are genre related (e.g., “rock,” “alternative,” “indie rock”), while only 3 (7 percent) correspond to opinion expression (e.g., “love at first listen”). The

proportion of tags representing bibliographical information is slightly higher at 12 percent (compared with 8 percent), which could be explained by the fact that users rarely disagree about these tags. Moods or emotions, on the other hand, are completely absent in the set of popular tags. We can finally add that the genre-related tags that have been applied several times tend to be broad (e.g., “rock,” “alternative,” “indie”), whereas those that have only been applied once to a given track often represent more specific genres (e.g., “ambient neoclassical,” “alternative hip-hop”). This suggests that a consensus might form more readily around the broad genre(s) a song fits in. It could also indicate that users are not always familiar with more specific genres and are therefore less likely to use them to tag music. To sum up, the tags that had been applied only once to a given song—those that belong to the long tail—were generally more specific, idiosyncratic, and more likely to be in a language other than English, whereas the tags that had been applied frequently tended to be broader and more objective.

Type	Examples	250 most popular tags in <i>Last.fm</i>	Tags applied to the songs in the dataset (subsample: 181 tags or 10%)	Tags applied only once to a song in the dataset (subsample: 120 tags or 10%)	Tags applied 10 times or more to a song in the dataset ($n = 42$ tags)
Genre/Style	indie, ambient, alternative hip-hop	63.6%	39.8%	27.5%	81.0%
Place of origin/Language	french, american, Montreal	12.8%	2.2%	2.5%	0.0%
Instrumentation	female vocalist, bass, trumpet	8.8%	5.5%	5.0%	0.0%
Opinion	love at first listen, addictives, best of	6.8%	15.5%	19.2%	7.1%
Mood/Emotion	happy, dark, chill	6%	10.5%	9.2%	0.0%
Period	2014, 10s	2.8%	3.3%	5.0%	0.0%
Theme/Topic	love, peace, heartbreak	2.4%	2.8%	3.3%	0.0%
Usage context	train songs, night driving, staring out the window on a road trip	1.2%	2.2%	2.5%	0.0%
Self-reference	seen live, reminds me of you, download later	1.2%	3.9%	6.7%	0.0%
Bibliographic information	matador, linkin	0.0%	6.6%	8.3%	11.9%
Other/indecipherable		2.8%	8.8%	8.3%	0.0%
<i>Tags in a language other than English</i>		0.0%	3.9%	5.8%	0.0%

Table 1. Tag distribution by type

5.2 Tag frequency distribution at the song level

As mentioned in Section 2.2, several researchers observed that the distribution of social tag frequency at the resources or at the database level resembled a power law curve. Exceptions include Kipp and Campbell (2006), as well as Peters (2009), who found that the distribution for certain resources rather followed an inverse logistic distribution. To investigate this aspect in *Last.fm*, we looked at the tag frequency distribution of the 25 most popular songs of the site³ (as of August 8, 2014). The *Last.fm* API allows users to extract the top 100 tags that have been applied to each song. However, each song did not have that many tags. One song had as few as 25 distinct tags. Because of that and because the power law and the inverse logistic distributions differ primarily by the shape of the beginning of the curve (they both exhibit a long tail), we decided to focus on the distribution of the top 25 tags of each song. The *Last.fm* API does not provide the exact tag count but rather the relative weight of each tag for a given song; a weight of 100 being attributed to the most often applied tag. This facilitates the comparison of the tag frequency distribution across the 25 songs. It also allows us to compute the average weight of each of the first 25 positions of the tag frequency distribution to provide a global view for the 25 songs. Figure 2 presents this distribution, which very closely follows a power law curve. The examination of the individual distribution of the songs shows that all resembled a power law curve, although the initial drop was more pronounced in some than in others.

³ We chose to focus on the most popular songs since the number of tagging assignments for a given resource in *Last.fm* is much lower than in a site like *Delicious*, where it is common for a resource to be tagged hundreds if not thousands of times.

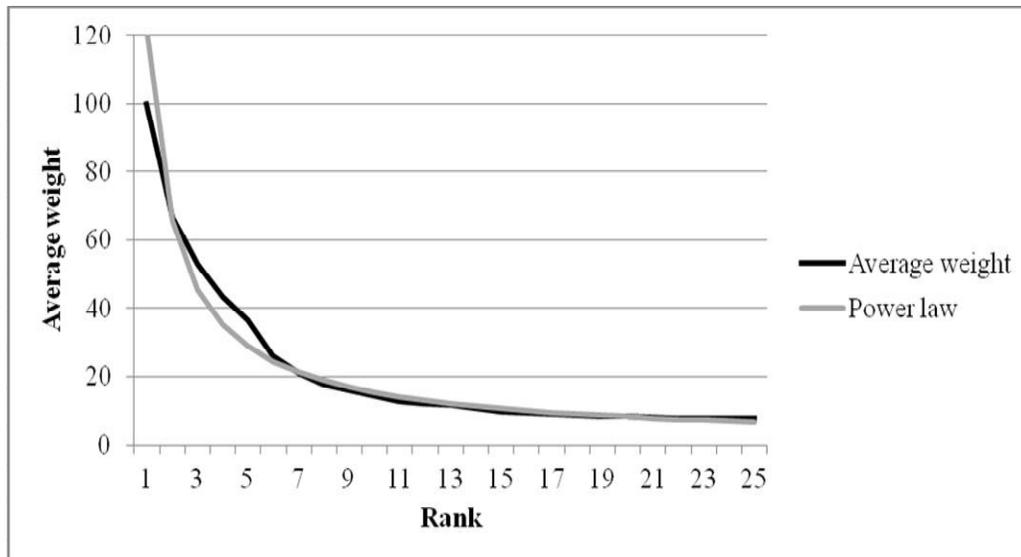


Figure 2. Distribution of the weight of the top 25 tags of the 25 most popular songs in *Last.fm*, compared with a maximum likelihood power law distribution.

5.3 Social tagging and popularity

We examined the tag frequency at the end of the third month after the release of each song. All songs of our sample had been tagged at least once. Each song had been tagged between 1 and 207 times, for an average of 19.9 tag assignments per song. The number of distinct tags varied between 1 and 85, for an average of 9.9 distinct tags per song. However, our analysis shows a clear asymmetry in the number of tags applied to tracks depending on their popularity. To estimate the popularity of a track, we used its *Last.fm* play count (i.e., the number of times a track has been played on *Last.fm* or with another media player with the application *Last.fmScrobbler*⁴) at the end of the first complete week after its release. Our analysis shows a strong positive correlation between the popularity of a track and the number of times it is tagged ($r = 0.73$; $p < .001$). The asymmetry is particularly apparent when we compare the most popular tracks with the rest of the sample. While the 10 most popular songs of our sample (i.e., those with a weekly play count superior to 15,000) had been tagged on average 140 times (median = 130), with 50 distinct tags, those who did not enjoy such popularity had been tagged a much smaller number of times. In fact, the number of tagging assignments drops abruptly after the 10 most popular tracks of our sample and continues to decrease with popularity. For instance, songs with a weekly play count under 1,000 had been tagged 4.9 times (median = 2.5), with an average of 3.56 distinct tags at the end of the third month after their release (see Figure 3 for a more detailed breakdown).

⁴*Last.fm Scrobbler* is an application that keeps track of the songs a user listens to on her computer and/or mobile device. The listening history of the user is sent to *Last.fm* to refine her taste profile, which is then used to make personalized recommendations.

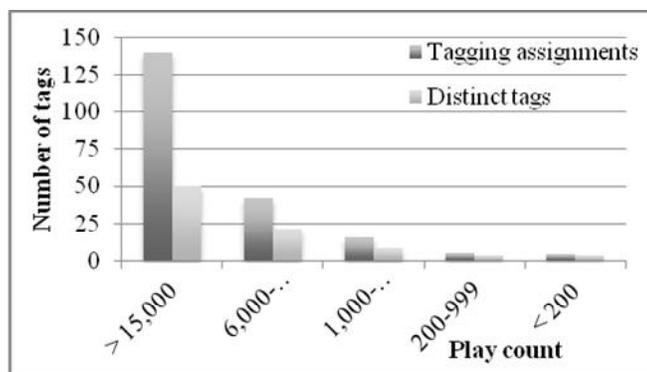


Figure 3. Number of tags applied to individual tracks by play count

5.4 Dynamics of tagging

Having tagging data for each day for the first three months after the release of the songs of our sample allowed us to examine how tagging patterns evolve over time. To do so, we focused on the 102 songs of our sample that had never been tagged before the release of the album in which they appear. On average, it took 16 days before a track received its first tag. However, like for tag frequency, we observe great variations depending on a song's popularity within the *Last.fm* community. While long tail songs (i.e., those with a weekly play count of less than 200) wait an average of 30 days before receiving their first tag, songs that are moderately to very popular (i.e., with a minimal weekly play count of 1,000) wait 11 days on average. We can add that only 49 percent of the songs in our sample had been tagged at least once at the end of the first week after their release. The percentage increases to 77 percent at the end of the first month. This shows that except for the very popular songs, most suffer at different levels from the so-called "cold start problem" that Levy and Sandler (2009) explains as follows: "tracks can be tagged only as listeners discover them, but untagged new tracks remain invisible within systems that depend on tags to give search results or recommendations".

It is also interesting to look at the first tags that are applied to a song. A manual classification of these tags shows that 55 percent are genre related. Twelve percent fall under the *Period* category (e.g., "2010," "10s"). It is impossible, however, to infer what these period tags mean to the users: it could refer to the year or decade the song during which it was released, or it could be applied by the users to remember the music they were listening to at a certain time in their life. Another 12 percent corresponds to opinion expression, and only 3 percent are self-reference tags (e.g., "download"). It seems reasonable to conclude from these data that *Last.fm*

users have a tendency to apply general tags first, which parallels what Golder and Huberman (2006) observe in *Delicious*.

Not surprisingly, our analysis also shows that the release of a song triggers a burst in tagging activity that lasts two weeks. After this peak, tagging assignments decreases sharply and then remains fairly constant for the next weeks (Figure 4 illustrates tagging activity for the whole sample). Furthermore, although for obvious reasons the number of distinct tags increases at a higher rate during the first two weeks, it continues to increase at a steady rate in the following weeks.

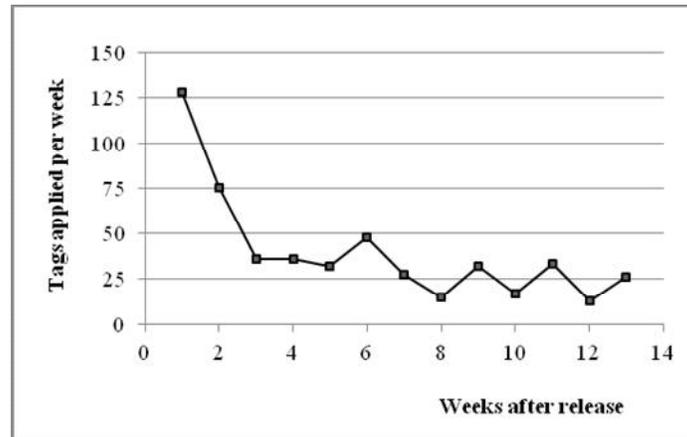


Figure 4. Tagging activity over a three-month period for the whole sample

5.5 Imitation and performance

As mentioned in Section 2.2, it has been demonstrated that people had a tendency to imitate the behaviour of others and select tags previously assigned rather than create new tags. Our analysis shows that *Last.fm* users, like those of other tagging systems, are influenced by others and often reapply existing tags for a given song, but perhaps not at the same rate than previous research suggests. Indeed, in our sample composed of 3,655 tag assignments, we had 1,813 distinct combination tag/song, which means that imitation occurred in 1,842 or 50.4 percent of the times. This is noticeably lower than the 60 to 90 percent estimate of Dellschaft and Staab (2008). This discrepancy could possibly be attributed to the system design and the incentives users have to tag in *Last.fm*. For instance, at the moment of tagging, the system not only suggests the most popular tags for the given resource but also tags the user has previously assigned to other resources, thus facilitating the use of tags for self-organization. It could also be attributed to the desire users might have to express their opinion considering the role music taste plays in the

formation of one's identity (DeNORA, 2000; FRITH, 1996). As mentioned in Section 3, *Last.fm* is also a social networking site: all users have a public personal profile, which includes the social tags they have applied to different resources in *Last.fm*. This explanation is supported by the findings of a study conducted by Zollers (2007) which shows that *Last.fm* users apply performance tags (which include self-reference tags) in a greater proportion than those of other systems. Zollers attributes this to the "performative aspect" associated with music.

The length of certain tags can also be seen as another indicator that there is a performance aspect to assigning social tags to resources in *Last.fm*, at least for some users. To our knowledge, the length in words of social tags have not been examined in other web sites, so it is not possible to affirm that the social tagging practices of *Last.fm* users are different from those of other tagging systems in that respect. It is nevertheless worth mentioning that, as observed by Levy and Sandler (2009), we also notice that *Last.fm* users often apply very lengthy tags. However, the proportion of lengthy tags in our sample differs from what Levy and Sandler observed: while 10 percent of the distinct tags of their sample are five words or more, the proportion in our sample is slightly lower at 8 percent. This difference could be attributed to the fact they their analysis is based on two music tagging systems. Another possibility would be that tagging practices have evolved during the last years. Figure 5 shows the detailed distribution by tag length for distinct combinations tag/song and for tag assignments. We notice important differences in these two distributions. Firstly, the proportion of short tags (tags that contain 1 or 2 words) is higher in the tag assignments distribution (89 percent) than in the distinct tags distribution (77 percent). This suggests that tags that are applied more frequently also tend to be shorter in length. Conversely, the proportion of tags that contain 5 words or more is much higher in the distinct tag distribution (8 percent) compared with the tag assignment distribution (3 percent). A closer look at the longer tags in our sample shows that these often consist in complete sentences. In terms of content, most refer to personal experiences or to the emotional/physical reaction of the user to the music (Table 2 shows the list of tags that contain 8 words or more in our sample). These tags demonstrate that tagging in *Last.fm* is much more than indexing music for retrieval or self-organization. For certain users, it is a way to express their tastes to the other members of this online community. This, again, recalls the close link that exists between music taste and identity. The effort they put into formulating

these tags and the poetic nature of many suggest the important role music plays in their life.

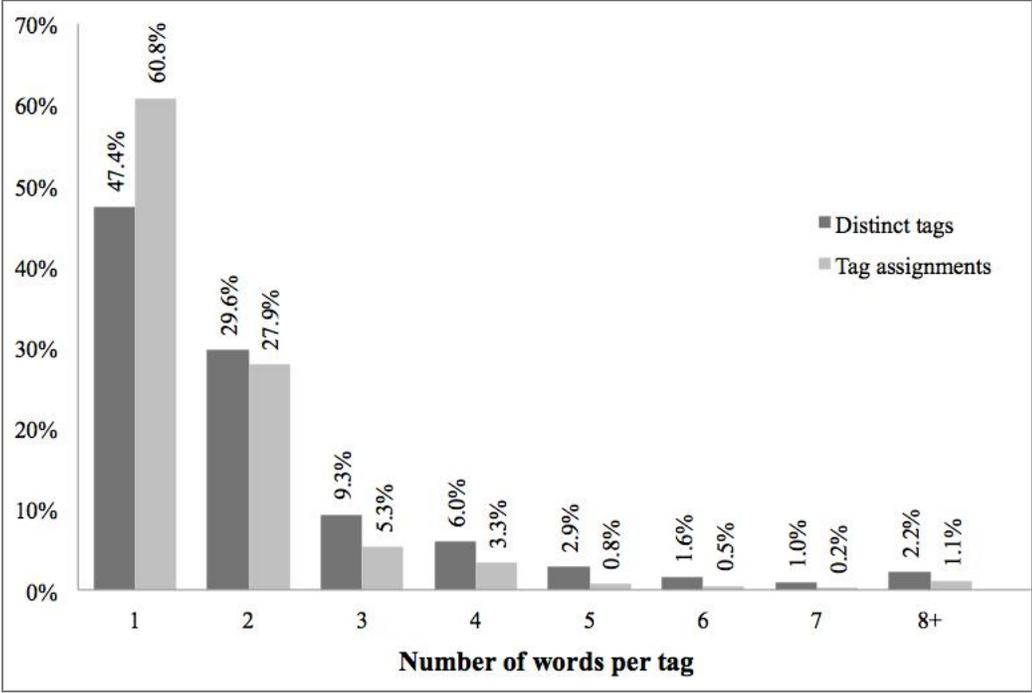


Figure 5. Distribution of tags by length

if you fall in love with me you should know these songs by heart
 music for when I feel like jumping off a very tall building
 songs i like to play whilst walking down the street at night
 songs that i will still be listening to in a billion years
 if this song were a pokemoni would catch it
 promised land hat society social aid and pleasure club fave
 songs that you wanna listen to over and over again
 ahh yes it made it on the record wooooooooooot
 i wish i could like this more than once
 Revisiting the infinite possibilities of life as a teenager
 total max point stand on your head and scream
 every aspect of this song is pretty epic
 glad to hear that bad religion still rocks
 Oh my god this song is fucking amazing
 rip out my fucking heart why dont you

songs I can actually listen to on repeat staring out the window on a road trip

Table 2. Tags containing 8 words or more

6. CONCLUSION

This study sheds light on some idiosyncrasies of collaborative tagging in *Last.fm*, which could be attributed to both the nature of the resource –music—and the particular design of this system. We find, for instance, that *Last.fm* users replicate the behaviour of other users less often than previous research on other tagging systems suggests. This study is also the first to examine the dynamics of the social tagging of music. We discover that, at the end of the first week after their release, a majority of songs have not been tagged and that a great proportion of the first tags to be applied to a song are genre-related and broad in nature. Finally, we observe a significant proportion of lengthy tags, sometimes containing as many as 14 words, many of which are poetic and very personal. This suggests that for some users, social tags are used as a means of expression of their relationship to music, and, to a certain extent, of their own identity.

REFERENCES

- BISCHOFF, Kerstin. et al. Can all tags be used for search? In: ACM CONFERENCE ON INFORMATION AND KNOWLEDGE MANAGEMENT, 17, 2008, Nova York.
Anais... Nova York: ACM Press, 2008. p. 193-202. Disponível em:
<http://www.csd.uoc.gr/~hy562/papers14/social-search/bischoff08.pdf> Acesso em: 12 nov. 2014.
- BOLLEN, Dirk; HALPIN, Harry. The role of tag suggestions in folksonomies. In: ACM CONFERENCE ON HYPERTEXT AND HYPERMEDIA, 20, 2009, Torino (Itália),
Anais... Nova York: ACM Press, 2009. p. 359-360. Disponível em:
<http://arxiv.org/pdf/0903.1788v1.pdf> Acesso em: 12 nov. 2014.
- DELLSCHAFT, Klaas; STAAB, Steffen. An epistemic dynamic model for tagging systems. In: ACM CONFERENCE ON HYPERTEXT AND HYPERMEDIA, 19, 2008, Pittsburgho (Estados Unidos). *Anais...* Nova York: ACM Press, 2008. p. 71-80.
- DeNORA, Tia. *Music in everyday life*. Cambridge; Nova York: Cambridge University Press, 2000.

- FRITH, Simon. Music and identity. In: HALL, S.; Du GAY, P. (Eds.). **Questions of cultural identity**. London; Thousand Oaks: Sage, 1996.p. 108-127.
- GELEIJNSE, Gijs; SCHEDL, Markus; KNEES, Peter. The quest for ground truth in musical artist tagging in the social web era. In: INTERNATIONAL CONFERENCE ON MUSIC INFORMATION RETRIEVAL, 8, 2007, Viena. **Anais...** p. 525-530.Disponível em: http://ismir2007.ismir.net/proceedings/ISMIR2007_p525_geleijnse.pdf Acesso em: 12 nov. 2014.
- GOLDER, S. A.; HUBERMAN, B. A. Usage patterns of collaborative tagging systems. *Journal of Information Science*, v. 32, n. 2, p. 198-208, 2006.
- GUY, Marieke; TONKIN, Emma. Folksonomies: tidying up tags?. **D-Lib Magazine**, v. 12, n.1, 2006. Disponível em: <http://www.dlib.org/dlib/january06/guy/01guy.html> Acesso em: 12 nov. 2014
- HALPIN, Harry; ROBU, Valentin; SHEPHERD, Hana. The complex dynamics of collaborative tagging. In: INTERNATIONAL CONFERENCE ON WORLD WIDE WEB, 16, 2007, Banff (Canadá). **Anais...** Nova York: ACM Press, 2007. p. 211-220. Disponível em: <http://www2007.org/papers/paper635.pdf> Acesso em: 12 nov. 2014.
- KIPP, Margaret E. I. Patterns and inconsistencies in collaborative tagging systems: an examination of tagging practices over a four year period. In: CANADIAN ASSOCIATION FOR INFORMATION SCIENCE ANNUAL CONFERENCE, 38, 2010, Montreal (Canada). **Anais...** Montreal (Canada): Concordia University, 2010
- KIPP, Margaret E. I.; CAMPBELL, D. Grant. Patterns and inconsistencies in collaborative tagging systems: an examination of tagging practices. **American Society for Information Science and Technology**, Maryland (Estados Unidos),v. 43, n. 1, p. 1-18, 2006.
- LAMERE, Paul. Social tagging and music information retrieval. **Journal of New Music Research**, v. 37, n. 2, p. 101-114, 2008.
- LEVY, Mark; SANDLER, Mark. (2009). Music information retrieval using social tags and audio. **IEEE Transactions on Multimedia**, v. 11, n. 3, p. 383-395, 2009.
- MARLOW, Cameron. et al. HT06, tagging paper, taxonomy, Flickr, academic article, to read. In: ACM CONFERENCE ON HYPERTEXT AND HYPERMEDIA, 17, 2006, Odense

(Dinamarca). *Anais...* Nova York: ACM Press, 2006. p. 31-40. Disponível em:
<http://www.danah.org/papers/Hypertext2006.pdf> Acesso em: 12 nov. 2014.

PETERS, Isabella. *Folksonomies: indexing and retrieval in Web 2.0*. Tradução de P. Becker. Berlim: De Gruyter/Saur, 2009.

TRANT, Jennifer. Studying social tagging and folksonomy: a review and framework. *Journal of Digital Information*, v. 10, n. 1, 2009. Disponível em:
<https://journals.tdl.org/jodi/index.php/jodi/article/view/269/278> Acesso em: 12 nov. 2014

ZOLLERS, Alla. Emerging motivations for tagging: expression, performance, and activism. In: INTERNATIONAL CONFERENCE ON WORLD WIDE WEB, 16, 2007, Banff (Canadá). *Anais...* Nova York: ACM Press, 2007. Disponível em:
http://www2007.org/workshops/paper_55.pdf Acesso em: 12 nov. 2014.