

F#%@ That Noise: SoundCloud As (A-)Social Media?

Chris Hubbles
University of Washington
Information School
chubbles@uw.edu

David W. McDonald
University of Washington
Human-Centered Design &
Engineering
dwmc@uw.edu

Jin Ha Lee
University of Washington
Information School
jinhalee@uw.edu

ABSTRACT

SoundCloud is an audio uploading, streaming, and sharing website. The site allows listeners to leave public comments at specific time stamps within a graphical display of an audio file as it plays, and offers support for user communication through a publicly viewable reply function. We extracted user comments and conducted a qualitative content analysis in order to characterize and examine participation on the site. Commenting consisted primarily of short, positive comments toward the song or artist, and use of the reply function was minimal, though some artists made use of it to communicate with listeners. We speculate that the commenting function is used less as a means toward community-building or social interaction, and more for fostering or indicating shared mutual experiences. Comment design implications are also explored.

Keywords

SoundCloud, social media, music streaming, user behavior, communication

INTRODUCTION

Social platforms are compelling and interesting to users. Systems like Facebook, YouTube, Twitter, and Pinterest attract a great deal of attention and carry much of the mindshare of the social media landscape. Attempts to build social systems for audio, however, have been marked by failure after failure across multiple platforms. Music is a prominently social vehicle in modern society – the explosion of peer-to-peer services in the late 1990s and early 2000s, for instance, was driven partly by social sharing of music files – but online platforms have struggled to support social audio functionalities. Platforms for upload and download of music files (e.g., Napster, Kazaa, Grokster, Morpheus) were perpetually hobbled by legal issues; platforms for sharing personal collections (e.g., iTunes, Windows Media Player, Google Play Music) often have proprietary restrictions that make them functionally difficult to use socially. Furthermore, they rarely have interpersonal communication tools built into them. Sites for promoting, discussing, and sharing favorite musical artists (e.g., MySpace Music, PureVolume, Facebook, Reverb Nation) seem perpetually poised to get swallowed up by the ‘next big platform’; Facebook, perhaps the site *du jour* in

this arena, competes with upstarts such as BandCamp and Tumblr, which may themselves take Facebook’s place in a few years, just as Facebook did to MySpace not long ago. On-demand streaming services are sometimes built with social applications in mind, and services such as Spotify and (recently defunct) Grooveshark offer some opportunity for interpersonal sharing and discovery. However, they lack the open social discussion available on a site like YouTube, which is often thought of primarily as a vehicle for video uploading, but is used as a *de facto* audio streaming and sharing tool by many of its users (Lee et al., 2017).

One website for social audio, SoundCloud, exhibits several of these elements at once: user uploading, playlist building and sharing, artist promotion, and streaming. SoundCloud also includes interpersonal communication and public discussion tools, but neither artist nor listener behavior on SoundCloud has received much study, despite the site’s popularity and its unique combination of functionalities. Our study focuses on the listener side of the SoundCloud user base. We began with the research question: *What does the SoundCloud comment function provide to listeners?* In particular, how do SoundCloud listeners use the commenting function? Do listeners attempt to communicate with musical artists, and if so, how? Do listeners attempt to communicate with each other, and if so, how?

By addressing these questions through a qualitative study of user commenting, we seek to understand behavior in a social music system that offers novel avenues for personal expression and interaction. We collected and examined SoundCloud user comments using a modified grounded theory approach. We scraped comments from tracks on the site over the course of the entire year 2013 and developed a content codebook, qualitatively evaluating 5,608 comments using a consensus coding strategy. We found that commenting on SoundCloud tended toward short, positive comments, often forcefully and profanely stated, with sparse and light interpersonal communication, though some artists used the commenting function as a means of responding to fan outreach. Our analysis indicates that the site’s timestamped commenting function – an innovative tool indicative of SoundCloud’s latent capabilities as a platform for social interaction between music fans – goes largely unutilized for this purpose. Instead, user postings tend more toward the declaratory, functioning more like broadcasts than invitations to dialogue. We postulate that this creates, in aggregate, a sense of *shared mutual experience*, a virtual analog of spaces that are social even

when there is no interaction between attendees, such as the cinema or concert hall.

WHAT IS SOUNDCLLOUD?

SoundCloud is an audio social networking site that enables users to upload digital sound files. The files can then be played by listeners on the site, and can be embedded easily into other news sites, blogs, and social networking sites. The site was founded in Berlin in 2008, and soon became a hub for German electronic musicians to share tracks and ideas (Mac, 2013). By 2009, the site had launched internationally and was already challenging MySpace, the then-dominant platform, in the online music-sharing universe (Van Buskirk, 2009). SoundCloud grew rapidly, expanding from 10 million registered users in 2011 to 40 million by mid-2013 (Graham, 2013). As of August 2014, SoundCloud was reporting over 175 million monthly listeners (not all of whom make accounts) (Sisario, 2014), a figure the site has not since updated. Its creators aspirationally call it the ‘YouTube of audio’, and it is used as a distribution channel by major-label musicians, popular podcasters, and even U.S. government offices such as the White House (under Obama) and NASA. The site’s reach is worldwide, with significant user populations in North America, Europe, and Southeast Asian countries such as Indonesia, the Philippines, and Vietnam (Bonanos, 2014).

Registration for accounts is free, and up to three hours of audio may be uploaded with a free account; paid accounts have larger upload limits and increased customizability options. An uploaded track on SoundCloud displays as a two-dimensional graph of the sound file’s waveform, and a tracking line moves across the waveform during playback. If an uploader allows it (the default is to allow), users can leave time-stamped comments on tracks, which display publicly for a few moments, along with users’ avatars, as the song is played. Users may also post replies to existing comments, offering a conduit for communication. SoundCloud also supports uploader tags that function like a folksonomy, as they are not supported by a controlled vocabulary structure (Giannetti, 2013).

RELATED WORK

Studies of Social Media

SoundCloud is an instance of *social multimedia*, “an online source of multimedia resources that fosters an environment of significant individual participation and that promotes community curation, discussion and re-use of content” (Tian et al., 2010, p. 28, quoting Mor Naaman). Such platforms have recently attracted much academic attention, yet SoundCloud has been the subject of little empirical study, though it is occasionally mentioned by musicians in qualitative interviews (Baym, 2012; Hoare et al., 2014).

Content analysis of user comments is a common method of analyzing social media spaces such as Slashdot (Gomez 2008), blogs (Mishne & Glance, 2006), and Twitter (Honeycutt & Herring, 2009; Dann, 2010; Naaman et al.,

2010), but these sites differ starkly from SoundCloud in terms of content and purpose. User behavior on social multimedia site Tumblr has received recent study (Xu et al., 2014), but the site is known more for visual media than audio, and it farms much of its audio content out to SoundCloud and Spotify. Closer is YouTube, whose model of media uploads and user comments resembles SoundCloud’s. Benevenuto et al. (2009) examined video responses to other videos on YouTube and found evidence of significant opportunistic behavior and spam-like tendencies. De Choudhury et al. (2009)’s analysis of YouTube comments found that participation increases when conversations are interesting, when familiar people are speaking, and when engaging dialogue is observed. Madden et al. (2013) carried out a classification of YouTube comments, identifying ten top-level categories of behavior, though with little discussion of relative importance. Thelwall et al. (2012) conducted quantitative and sentiment analyses of YouTube comments, noting vigorous interaction over contentious issues such as politics and religion, and comparatively little response to positive comments in comparison to negative comments. A follow-up study by Tsou & Thelwall (2014) used a qualitative coding approach to compare user behavior on YouTube comment pages for TED talks with the TED website, though it did not address user communication. Shoham et al. (2013) qualitatively evaluated comments on a single YouTube video, finding more broadcasting of messages than interaction. Rotman et al. (2009) executed a grounded theory analysis of commenting on YouTube videos, finding that users tended to believe YouTube enabled them to hold a ‘sense of community’, even though this was at odds with the content of actual comments (which did not show evidence of fostering tight-knit groups, and which rarely hosted prolonged conversations).

We speculated that the real-time display of SoundCloud’s comment system might uniquely influence commenting behavior, since commenting at specific timepoints within a track allows for greater referential specificity. Moreover, the continuously unfolding comment layout results in a textual interface quite different from other platforms, which may affect commenting patterns. The design implications of SoundCloud’s model have not gone unnoticed; Guimaraes et al. (2012) proposed a tool for real-time display of comments on YouTube uploads, essentially adapting the SoundCloud comment structure for video, and video sites such as Viki (Russell, 2013) employ similar interfaces in practice. Empirical work is needed to understand how such design changes impact media use and users.

SoundCloud began attracting scholarly inquiry only recently. Several reports by Jourdanous (Jourdanous et al., 2014; Allington et al., 2015; Jourdanous et al., 2015) investigated SoundCloud qualitatively, through musician interviews, and quantitatively, by tracking site ‘following’ behavior. They noted clusters of social behavior (both via following of artists and commenting on tracks) revolving

around genre scenes and local geographies. Ishizaki et al. (2015) used a speech act analysis approach to qualitatively evaluate user commenting on SoundCloud and Last.FM. The analysis indicated that most SoundCloud comments were primarily short claims about, or reactions to, songs, with relatively little interpersonal interaction. In our study, we seek to extend Ishizaki's work by using an alternate qualitative coding strategy, which will help surface more about *what* commenters are talking about, rather than the linguistic function of their statements. Furthermore, by executing a more in-depth analysis of user replies, we hope to provide additional insight into the site's conversational reach, and evaluate the tentative conclusions reached by Ishizaki on SoundCloud as a social system.

Studies of Music Services

There is a substantial body of research on evaluation of online music services in the music information retrieval (MIR) domain, particularly within MIREX (Music Information Retrieval Evaluation eXchange), the annual evaluation campaign for music tools and services led by researchers at the University of Illinois. However, few studies evaluate performance, user interaction, or experience with *commercial* music services. The field of MIR is still relatively young, and has focused on testing algorithms for functions such as feature extraction, classification, and clustering, in order to assess their applicability to music data. Interest in evaluating and understanding user experiences of MIR systems has only surfaced recently, as evidenced by the MIREX User Experience Grand Challenge 2014. Scholarly evaluation of music services is also hindered by the fact that many commercial music services tend to be short-lived.

Commercial music services such as Last.fm, Apple iTunes Genius, Pandora, and Musicoverly have been reviewed and used in a few studies in MIR (Kaminskas & Ricci, 2012). Previous studies that make use of user data from commercial music services tend to either extract data (e.g., tags, user questions, playlist results) and use them for quantitative experiments, or to evaluate the performance of particular system features, rather than aiming to holistically understand how the service is being used. Representative examples of the former category include Hu & Downie (2007) and Laurier (2009), both of whom used social tags to understand music mood. Examples of the latter category include Barrington (2009) and Lamere (2011), who evaluated the performance of music recommender systems, including some commercial systems like Apple's iTunes Genius and Google's Instant Mix. In these studies, systems are evaluated through an experimental setup (e.g., the subject listens to a given seed song and evaluates the playlist) or via personal observation and opinion.

There are also a few studies aimed at improving general understanding of user needs and behaviors based on data from commercial music services. Bainbridge (2003) and Lee (2010), for instance, analyze users' music-related

queries from online Q&A services to understand different types of music information needs. Other studies like Lee & Waterman (2012) focus on collecting users' opinions about commercial music services to identify user requirements and desires. However, these studies are not focused on investigating users' interactions and experiences within a large-scale commercial music service, despite this being a common conduit for musical consumption in the digital age. Such a focus can uncover emergent and changing social behaviors, as well as offer insights on how to design services to support user needs effectively. Our qualitative study begins to fill this gap by exploring what kinds of user activities and interactions occur using the commenting function on a popular social music platform.

STUDY DESIGN

Pilot Study

We conducted a pilot study to explore strategies for collecting and analyzing data from SoundCloud. We settled on attempting to obtain a representative sample of popular songs, from which we would then draw user comments for analysis. SoundCloud does not provide a public ranking of its most popular tracks, but it hosts an application programming interface (API) which can be used to make a range of track-based requests. We piloted data collection and analysis by pulling tracks from SoundCloud's database through the API in February 2013 using the 'hotness' parameter. This returned random tracks with high play counts and favoriting counts that had been uploaded in the previous two weeks. We collected comments from these tracks and used them to develop initial content codes via a grounded theory approach. However, we could not use the same approach to draw a sample for the full study, because the 'hotness' parameter was deactivated by the site in April 2013 due to backend issues.¹

Our troubles with the 'hotness' parameter illustrate Karpf's Rule of Online Data (Karpf, 2012). SoundCloud is prone to rapid change in the behavior of its users (both listeners and artists) as well as in site architecture and server response, and the more visible it becomes commercially, the less representative and trustworthy its data will be for empirical analysis. Nevertheless, a certain measure of 'messy' data is necessary in order to study fast-moving new media settings such as SoundCloud, and we follow Karpf's embrace of transparency and 'kludginess' in data collection and analysis, in hopes of capturing insights regarding this valuable, but volatile, online environment.

Data Collection

For the full study, we wrote custom code to collect a wide sample of track, comment, and user data through the

¹ See Osman, P. Removing 'Hotness' Parameter. *SoundCloud Backstage Blog*, April 16, 2013. developers.soundcloud.com/blog/removing-hotness-param

SoundCloud API. We collected track metadata using the SoundCloud *search* API, which supported an unconstrained search of tracks that were present in the system as of a specified date. In response to that request, SoundCloud returned metadata for up to 8,000 tracks that had been uploaded to the system as of the date specified in the query. Exactly how the 8,000 items were selected in response to the query was not specified by the SoundCloud documentation, but, like the ‘hotness’ parameter, appears to be some function of popularity, related to the number of times the song has been played, the number of comments, the number of times it has been favorited, and how long the track has been in the system. As each track was collected, metadata from the track was stored in a database.

Comment data was collected using the SoundCloud *track* API, which provides resources associated with individual tracks. For each track in our database, we queried the track API and requested all comments available for the track. Comment metadata includes data about the individual who posted the comment, the date of the comment, and the millisecond location of the comment in the timecode of the track. We stored comment and user metadata in a database. A data collector, collecting tracks and comment data, was run repeatedly beginning in June 2013; we continued to run the collector through early 2014.

Our dataset is limited by the opacity of the SoundCloud API. For example, it is possible that a date search for tracks uploaded to SoundCloud as of a specific date could return tracks from any prior month. This is a function of how SoundCloud’s servers and code decide to respond to the request. We collected 240,774 tracks, some of which had been posted as early as 2007. 4,383,874 comments had been posted to these songs by 1,800,630 unique users. 43,568 of the tracks had no comments; 63,209 had 1-9 comments; 108,071 had 10-99 comments; 27,725 had 100-999 comments; and 530 had 1,000 or more comments.

Our dataset contained data for every day of 2013 and so our analysis focused on 2013. Figure 1 is a graph of the total number of tracks uploaded per day. Uploads per day is fairly consistent at first, generally ranging between 100 and 200 tracks per day, with a slight increasing trend at the start of the year. From the middle of June through the middle of July, there was a concerted spike in the number of uploaded tracks per day. The peak day was July 21, with 7,547 tracks uploaded. For subsequent months, the number of total uploads per day again ranged between 100 and 200 tracks per day, with a slight drop-off as the year ended.

Figure 2 is a graph of the total number of comments uploaded per day. Total comments per day showed a slightly different pattern, but also peaked in July, with 27,533 comments uploaded on July 19. While the number of track uploads seems to have returned at the end of the year to values similar to the beginning of the year, total comment uploads per day ended the year at a rate of nearly 5,000 comments per day more than early in the year.

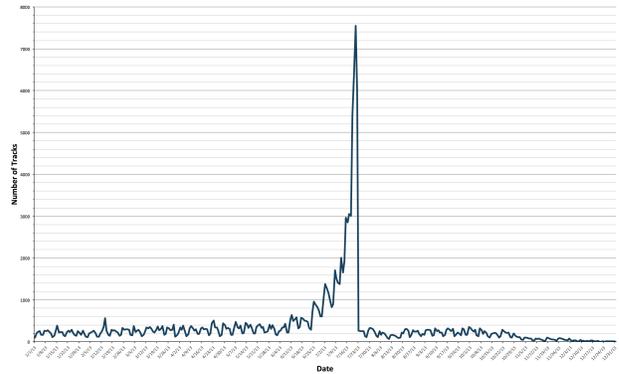


Figure 1. Total track uploads per day for 2013.

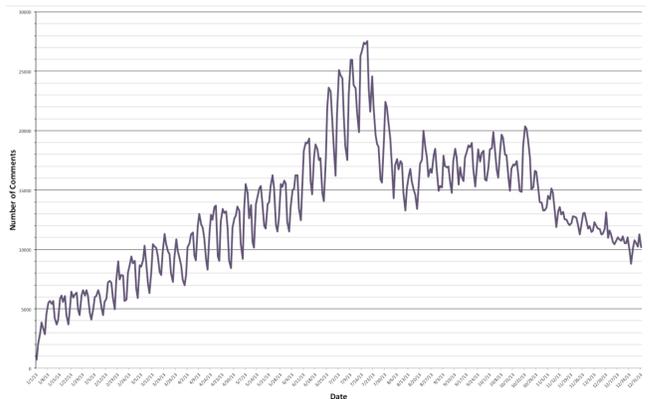


Figure 2. Total comment uploads per day for 2013.

Sampling

We decided to focus our analysis on the data from 2013, which narrowed the working dataset to 134,075 tracks. We chose to sample at the *track* level, rather than the comment level, because we believed it was important to understand comments in context. Our focus for this project was on SoundCloud’s *musical* content, so spoken-word tracks (podcasts, comedy, etc.) were removed from the dataset. While we found much commenting occurring on spoken-word tracks, we believe these would require a separate analysis with a *sui generis* qualitative coding scheme.

Tracks were stratified by duration into bins of less than ten minutes (primarily single songs and song previews), 10-25 minutes (mostly extended mixes and mini-album streams), and greater than 25 minutes (typically full album streams, concert recordings, and ‘megamixes’). Tracks with fewer than 10 comments were excluded from analysis due to anticipated paucity of social activity. To make coding manageable for an exploratory study, we restricted our sample to the 0-10 minute bin and to tracks with between 10 and 500 comments. Random track samples were taken from each calendar quarter, with slightly more samples taken from the third quarter because of the spike in activity.

Codebook Development

A qualitative codebook was drawn up based on patterns

observed in the pilot set. Each comment was manually evaluated and classed into categories. Iterative codebook refinements were made as analysis continued, following grounded theory methods (Corbin & Strauss, 2008). At least one code was assigned to each comment (including an ‘unintelligible’ code for comments having no interpretable content), and the number of codes assigned to a comment was not limited. The codebook underwent five iterations of sample coding. The three authors independently coded comments, then met and sought unity on the meaning and boundaries of each code, as well as the need for new codes. This resulted in 39 codes at the end of piloting. Four top-level code categories were recognized:

- *Mood*: whether the comment was positive, negative, or contained elements of both
- *Content*: what common attributes comments shared
- *Music*: essentially a subset of *Content*, but split out to cover specifically musical aspects of commenting
- *Reply*: whether the comment is a reply, and what type

Our comment analysis was conducted as a *consensus code*. Each comment set was coded independently by two coders, who then met and discussed disparities until as much common ground as possible was established. One project lead was assigned to look for codes the first two may have missed or misused; this researcher then engaged the first two in dialogue once more. The third researcher also acted as a tiebreaker when the first two coders could not initially come to consensus. We hoped that the use of three coders and the consensus strategy would help mitigate problems with comment ambiguity, short length, lack of context, and variety in slang and dialect present in the sample set.

RESULTS

Sample Composition

The sample set of SoundCloud tracks consisted of 64 different songs. The number of comments for a track ranged from 10 to 443, and track length ranged from 44 seconds to 6 minutes 44 seconds. Four songs were removed from the set because they had been deleted from SoundCloud, and two songs were removed because a preponderance of their comments were not in English. This left a workable set of 58 songs with 5,608 comments, all of which were consensus coded. Because our analytic strategy was so labor-intensive, it was not feasible to process a larger percentage of our dataset. Nevertheless, our analysis evaluates the same number of tracks as Ishizaki et al. (2015) (58 songs, 24,111 comments) and employs a more thorough coding method, which should offer new insights into the conversational world of SoundCloud’s listenership.

The set of songs retrieved was dominated by electronic music and hip-hop, which broadly mirrors the predominant tastes of the SoundCloud user base both historically and at current. In some cases, our codebook reflects the genre biases of our sample set. For instance, in pilot coding,

comments about ‘the drop’ (a sudden rhythmic shift, usually a slowdown, often coupled with heavy bass and a change in melodic material) were so common on electronic music tracks that we assigned the phenomenon its own code (Yadati et al., 2014). Commenting behavior likely differs in genres not well represented in our sample.² The type of track uploader also varied, which had the potential to affect commenting behavior. Some tracks were uploaded by the original artist, but others were uploaded by a remixer, record label, promotional company, blogger, or fan. Songs from both famous (e.g., Avicii, Busta Rhymes, Capital Cities, Black Violin) and relatively unknown artists were present in the sample; fan uploads were responsible for the presence of songs by several of the most famous artists.

Comment Contents

The first three code categories (Mood, Content, and Music) are discussed in this section. User comments are drawn from our sample set and are identified by a unique internal assignation. By far, the most common types of comments were simple statements praising the song or the artist. Examples include “Favorite song” (338113), “sick tune !!” (3772901), “fuck yeah!” (4483935), “shits hottt my gawd” (4483905), and “fuckingggggg amazing!” (3772879). The inclusion of cursing is by no means accidental; significant portions of the user base felt it necessary to express themselves using oaths, often seemingly in a battle to outdo each other in extremity. Occasionally, commenters posted overtly sexual or scatological comments such as “EAR SEX” (1808282) or “I just ruined a good pair of underwear.” (4483968), and at some points considerable cleverness came into play, just as on any men’s bathroom wall (Dundes, 1965). Furthermore, the frequency with which users explained their musical loves in terms of sexual gratification and drug references, such as “perfect for a nice doob. i love it keep it up guys” (437943), reflects a persistent cultural metaphor associating the emotional states experienced during these activities. ‘Sex, drugs, and rock-n-roll’ as a close-knit club of dionysian affairs is still alive and well, if this sample set is any indication.

Positive comments far outnumbered negative ones, confirming Jourdanous’s (2015) findings; over 92% of comments were judged by our coders to be generally positive in disposition. Even on the most-criticized track in our sample, a dance remix of Nirvana’s “Smells Like Teen Spirit”, 38 of 71 comments were coded positively. The general lack of negative comments suggests that users tend not to comment on things they do not like, but since much listening is self-selecting, it may be that users are not encountering music objectionable enough to comment on, or are clicking away instead of dwelling and responding.

² For a list of artists, song titles, genres, numbers of comments, and reply percentages for all tracks in the sample set, please contact the first author.

The use of caps lock, repeated letters, excess punctuation, and emoticons or ASCII art for emphasis and versatility of expression was pervasive.

Users would sometimes point to aspects of the music that drew them in or pushed them away. This could be as simple as “Dig the low climax” (706607) or “Sound quality sucks...” (1873962), and, in general, commenters were rather inarticulate about musical substance. Yet music comments could also get quite involved, with users providing detailed accounts of musical characteristics. Some users made observations about instrumentation, rhythm, vocals, stylistic peculiarities, genre designations, and song structures. Because of the timed commenting function, users were able to key their responses to the exact point in the song they wished to highlight. Comments such as “that snare just punched me in the face” (967477) and “I think megatron, just jizzed himself” (5734693) (the latter left in response to an abrupt, robotic interlude in an electronic music track) illustrate how this can allow some users to speak to precise musical moments with great effect. The most common music-related codes were in the realms of genre/style, lyrics (especially direct quotations), rhythm, production/remixing, and comments about artists.

Aside from spam and self-promotional comments, most comments pertained to the song, the artist, or something closely related, like genres, record labels, or similar artists. Many users divulged when, where, and how many times they listened to a song. Another recurrent topic was an artist’s concert appearances; some users noted instances in which they heard the artist play the track in a club or concert, or would express how they were looking forward to a future live event. Users would sometimes leave notice via comment if they had blogged about the song, posted it to a public playlist, or publicized the track in some other way, serving both promotional and informational purposes.

The different types of tracks posted could attract different types of responses. While many songs garnered comments requesting download links or inquiring about release details, these were more common on track previews, which left listeners with only part of the song to listen to. Likewise, electronic music frequently incorporates sampled recordings or interpolations of other musicians’ work, and it was common to see users commenting on, or asking about, sample sources on such tracks.

We also created a code that served as a ‘miscellaneous’ bin, to capture instances where comments recounted stories or personal anecdotes, experiences with the song, or generally unusual commenting behavior. Some examples include, on the simpler end, “This song made my summer” (642908) or “This track makes me dance while driving!!” (4373925), and on the more profuse and profane end, “This is proly the worst fucking song ive ever hear fuck tmills fight me you douchebag” (3843825). This reflects significant use of the commenting function to do more than *merely* make the simplest of telegraphic statements. There is no apparent

length limit to SoundCloud comments; we encountered comments with more than 900 characters (~175 words).

Reply Comments

Since reply functionality is built into the commenting system, we explored whether users take advantage of the reply function to communicate with one another, or with the artist or uploader of the song. For the 58 songs in the sample set, each comment was coded as a reply if it had been entered as a response to a previous comment using SoundCloud’s reply function. However, a simple count of this nature would be misleading. Some users reply to their own comments; these self-replies were removed. Occasionally, users repost the same text more than once as a reply; these were not removed. Some replies appear to be spam, and others were thematically unrelated to the original comment; these also were not removed. Additionally, many tracks had large numbers of replies to the first comment left at timepoint 0:00, sometimes in excess of commenting elsewhere in the song. This appears to have been the result of a defect in an earlier version of the commenting function. Unless context indicated genuine response to an earlier post, these were removed as replies. Finally, non-reply comments were combed for indicators of inter-user communication. Comments that mentioned other users in the form @username were classed as replies, as were comments that included contextual indicators of reply. The contribution of manually-classed replies was very small.

For each of the 58 songs, an adjusted reply count was figured according to these rules. 22 songs had no adjusted replies. While there seems to be a comment threshold over which at least one reply per track will be present, our sample set included tracks with as many as 77 comments without a single adjusted reply. 517 of the 5,608 comments in the sample set were found to be adjusted replies, a rate of 9.22%. However, the distribution of adjusted replies per track is strongly skewed; all but 9 of the 58 tracks had adjusted reply rates under 10%, and the median percentage of a track’s comments which were adjusted replies was only 2.36%. Most adjusted replies were single responses to comments, often approaching non sequitur, with no further discussion. Back-and-forth conversations are rare. There were 403 conversation threads between two or more parties in our sample set, with an average thread length of 2.21 comments per thread. Only four threads had lengths longer than four comments, though one additional conversation between a listener and an uploader was carried on across several short threads.

A small number of tracks, however, had larger numbers of replies. For all six of the tracks in our sample set with adjusted reply percentages greater than 20%, the track uploader had meticulously given responses to many – in some cases nearly all – of the comments that users had left on the track, often to express gratitude for listeners’ appreciation and support. In fact, the number of replies left by uploaders in our sample set approached the number left

by listeners; 231 replies were left by uploaders, while replies from listeners totaled 286. This indicates that fans sometimes used the service to reach musicians directly, and that some artists fostered that interaction via the reply function. Thirteen tracks had instances of musicians replying at least once to comments about their own music. However, we also looked for instances where listeners appeared to call out directly to the artist, to offer praise or blame, or to request something, such as more music, the identification of a sample, tour dates, changes to a song, or a link to a digital download of the track. Such comments were observed both on tracks artists uploaded to their own accounts and on tracks uploaded by labels, promoters, and blogs. They were posted to the songs of relative unknowns as well as those of stars, and few received a response. This seems to indicate that many users are speaking at artists without much expectation or hope of a reply.

DISCUSSION

The analysis we have conducted indicates that while the SoundCloud comment function is used extensively by listeners to respond to the music they are hearing, it is comparatively rare that they use it to respond to each other. Communication between listeners using the reply function is generally uncommon, and what communication does occur through the comment function is almost as likely to be artist-to-listener as it is to be listener-to-listener.

Our findings in part corroborate Rotman's work on YouTube comments, despite the fact that SoundCloud, in theory, allows for more focused conversations around particular moments in a musical work. Rotman et al. (2009) noted, "Comments do not generally create a prolonged discussion, and most are left unanswered" (p. 44), and that "Comments... create, at best, an interaction that culminates in 2-3 exchanges, and is seldom addressed by users other than the channel-owner and the original commentator" (p. 45). This is precisely what occurs in SoundCloud, at least among genres well represented in our sample set. In further explorations of YouTube sociability via user interviews, Rotman finds public commenting to be one small piece in a larger ecosystem of channels for sustaining the felt sense of community professed by some of its users (Rotman & Preece, 2010). It may be that SoundCloud's comment design is insufficiently attractive to draw users in for conversation, and interpersonal communication is being pushed out to the site's personal messaging system or to other venues such as Facebook (with which SoundCloud integrates well). Even if SoundCloud's commenting interface carries with it the appropriate conversational affordances (Norman, 1990; Reid & Reid, 2010) to support interaction, the mere presence of community-building tools does not mean that users will take them up in that way. The multiplicity of online communication spaces may make any single tool, even a well-designed one, minimally important to users for that explicit function (Baym, 2007).

The sociability of SoundCloud comments may be hampered

by real-time display design choices. Some tracks, especially short ones, have large numbers of comments such that, if all were displayed sequentially, the listener would see a blur of comments as the waveform progresses. SoundCloud now throttles the number of comments displayed to slightly less than one per second (in most browsers) where comments bunch together. The display appears to be randomized, and different comments may show if the song is replayed, though older comments do not appear to age out of display. On popular tracks, listeners may see only a small portion of comments posted to a song, which may discourage participation. Managing comment volume is a persistent challenge for streaming media services (Haimson & Tang, 2017), and changes in spatial positioning, size, and length of display may increase the likelihood of genuine response.

The initial impression given by our analysis seems to be of an environment akin to a rather asocial party; many people present, milling about, muttering to themselves, mostly ignoring each other. Yet an impressive mass of comments is being left here, hundreds in a matter of days for the site's more popular tunes. Who is their audience? Our findings lend support to Ishizaki et al. (2015)'s claim that user motivations for commenting relate more to expression than interaction. Consistent with Shoham et al. (2013), listeners may employ the comment system more as a medium for self-expressive broadcasting, directing speech at the whole world, but no one in particular. Similarly, one could conceptualize SoundCloud as a public space or platform for airing opinions about the music and attempting to influence others' experiences of it. On this view, SoundCloud's comment function acts like a town square soapbox; anyone may stand and shout, but no one is obligated to listen.

This raises a further question: *Is anybody actually reading the comments?* Listeners may well use SoundCloud merely as background music, hitting the play button without giving the comments a glance. SoundCloud's streaming function smoothly glides on to new tracks as the user browses, without displaying the new track's waveform or comments unless the user navigates to the new track's page. Those who leave comments are probably more likely to read comments themselves (though further empirical work would be needed to confirm this). Commenters post because they *believe* their words will be read, and so perhaps SoundCloud's comment function satisfies a more fundamental human desire to leave a mark of existence, which Gregory Benford, in homage to the ubiquitous mid-20th-century graffiti, called the 'Kilroy Was Here' impulse (Benford, 1999). The comment function may serve as a sort of sanctioned graffiti (Shoham et al., 2013), a wall to be filled with whatever comes to mind, regardless of how inconsequential or ignorant.

Nevertheless, there is probably more to such commenting from the user standpoint than merely flinging notice of one's existence into the ether. The comments alert listeners that there are other people listening in a more robust way

than play counts, favorite counts, or chronologically-arranged comment boxes can (as on YouTube or discussion threads). Listeners are asynchronous with respect to each other, but they see the comments synchronously as they play the song, which simulates a listening experience shared with other people. The notion of copresence, derived from Goffman (1963) and later refined in communications literature, refers to “the sense of being together with other people in a remote environment” (Zhao, 2003), and comes close to capturing this idea. Copresence has been connected to online multimedia commenting (Weisz et al., 2007), but it better describes systems that allow for synchronous interaction; a closer analogue would be live-streaming services with response functions, such as Twitch (Hamilton et al., 2014) or Facebook Live and Periscope (Haimson & Tang, 2017). Nevertheless, SoundCloud does afford a felt sense of having others around who are also appreciating (or belittling) the same entertainment one is consuming. In this way, it may serve as a stand-in for shared mutual experiences, like going to hear music at a concert or seeing a film in a theater. Despite having a large number of people gathered together in one place, there is often little direct conversation between attendees at these occurrences. We do not always think of such shared experiences as *social*, perhaps because they are not principally composed of interaction events. Yet these experiences are different than they would be if one were there by oneself, creating a sense of social presence (Ducheneaut et al., 2006) which helps explain their continuance as major entertainment draws.

While listener-to-listener interaction was scarce, the one peripheral, but significant, social activity we identified was occasional artist-listener communication via commenting, enriching Jourdanous et al. (2015)’s qualitative findings about the social value of SoundCloud for artists. If SoundCloud wishes to support artist-listener commenting more fully, it could change its comment notification structure. Uploaders may elect to be notified via e-mail when users comment on their tracks, but this is buried in the settings menu and is handled at the account level, rather than at the track level. If a notification box were offered at the time of upload, and a comment management interface included at the playlist and account levels (with e.g. radio buttons for each track, in addition to a blanket all-on/all-off option), more artists might begin to employ commenting as a vehicle for making personal connections with fans.

ISSUES AND CHALLENGES

SoundCloud’s interface has not remained static. A site redesign in late 2012 shrank the playback icons and comment display boxes, which led to complaints that the site was “slighting the worth of comments as a social feature” (Hill, 2012). Some early adopters felt alienated by this change; Peter Owen, a songwriter who used SoundCloud to share audio, claimed in 2013 that “Most musicians are finding that the number of plays are up and the number of comments has fallen....Those that haven’t

left are starting to use the site differently. Now they’re using it just as an advertising space...instead of the collaboration and creative communication that was there before” (Mac, 2013). Our sample set is composed of tracks uploaded after this redesign, so we cannot empirically evaluate Owen’s claim. In August 2014, the site began incorporating audio advertising (Sisario, 2014), and in 2016 it announced plans to launch a subscription streaming service modeled after sites like Spotify (Pierce, 2016). Commenting behavior may be strongly sensitive even to very small changes in site design, underscoring the importance of studying systems over time in order to provide a fuller picture of dynamic user conduct.

The site’s burgeoning popularity has also made it a magnet for stooge accounts and artificially inflated play counts; along with this has come the possibility of falsified comments posted to build a ‘buzz’ around an artist (Matthew, 2013). Since exposure and high numbers, even if faked or bought, bring in real fans, there is a strong incentive here, as on many online platforms, to spam. It remains possible that portions of our data beyond what we could identify are composed of social spam.

Lastly, the public nature of our dataset has eroded over time. Portions of the data we scraped, including data that was qualitatively coded, was subsequently removed from the site, often presumably because remixed, sampled, or fan-uploaded material was flagged for copyright infringement. SoundCloud now uses automated detection tools to remove allegedly infringing tracks, and, like YouTube, has granted Universal Music Group broad, direct power to remove material at Universal’s own discretion (Cushing, 2014). When tracks are removed, associated comments go down with them, meaning that future attempts to compile similar data will lack potentially valuable corpuses of user communication. As with most web content, the integrity of SoundCloud’s historical record receives little consideration from its designers, making our study a glimpse into a species of technological ephemera whose intellectual content disappears far faster than the paper, wax, and tape throwaways of prior generations.

CONCLUSION

Future work building on this study could include qualitative analysis of non-musical audio, which is a significant subset of SoundCloud’s holdings (Alcorn, 2014). Our dataset indicates much commenting is occurring on these tracks, which may differ substantially from what we analyzed. Other qualitative approaches might be employed to give a more well-rounded picture of SoundCloud’s conversational world; techniques such as sentiment analysis, linguistic inquiry and word count, or topic modeling could help evaluate larger and more varied genre sample sets.

Our analysis also excludes non-English music, and as we noted above, SoundCloud is truly a global platform, with users commenting in dozens of languages. Cross-cultural

comparisons of SoundCloud commenting behavior may yield a much richer picture of actual user behavior than can be presented in this limited study. Characterizing the user base demographically would also be of value. Surveys, interviews, or ethnographic work would help develop a more detailed picture of why and how users comment.

ACKNOWLEDGMENTS

Special thanks to Adrian DSa for help with SoundCloud's API, and to Alaina Bull, Rosanna Lui, Andy Magnusson, Maulik Mishra, Malavika Ravi, Bhargav Shah, Emily Spracklin, and Carolyn Tweedy for consensus coding.

REFERENCES

- Alcorn, S. (2014, Jan. 15). Why Audio Never Goes Viral. *Digg*. <http://digg.com/originals/why-audio-never-goes-viral>
- Allington, D., Dueck, B., & Jordanous, A. (2015). Networks of Value in Electronic Music: SoundCloud, London, and the Importance of Place. *Cultural Trends* 24, 211-222.
- Bainbridge, D., Cunningham, S. J., & Downie, J. S. (2003). How People Describe Their Music Information Needs: A Grounded Theory Analysis of Music Queries. In *Proc. ISMIR 2003*.
- Barrington, L., Oda, R., & Lanckriet, G. R. (2009). Smarter than Genius? Human Evaluation of Music Recommender Systems. In *Proc. ISMIR 2009*, 357-362.
- Baym, N. (2007). The New Shape of Online Community: The Example of Swedish Independent Music Fandom. *First Monday* 12:8.
- Baym, N. (2012). Fans or Friends: Seeing Social Media Audiences As Musicians Do. *Participations: Journal of Audience and Reception Studies* 9, 286-316.
- Benevenuto, F., Rodrigues, T., Almeida, V., Almeida, S., & Ross, K. (2009). Video Interactions in Online Video Social Networks. *ACM Trans. Multimedia Computing, Communications and Applications* 5, 1-25.
- Benford, G. (1999). *Deep Time: How Humanity Communicates Across Millennia*. New York: Avon.
- Bonanos, P. (2014, May 15). In Developing Asian Markets, Mobile and SoundCloud Dominate Listening. *Billboard*.
- Corbin, J. M., & Strauss, A. (2008). *Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory*. Los Angeles: Sage.
- Cushing, T. (2014, Jul. 2). SoundCloud Has Given Universal Music Group The Ability To Directly Move Content. *TechDirt*. <http://www.techdirt.com/articles/20140702/10252827760/soundcloud-has-given-universal-music-group-ability-to-directly-remove-content.shtml>
- Dann, S. (2010). Twitter Content Classification. *First Monday* 15:12.
- De Choudhury, M., Sundaram, H., John, A., & Seligmann, D.D. (2009). What Makes Conversations Interesting: Themes, Participants and Consequences of Conversations in Online Social Media. In *Proc. WWW 2009*, ACM Press, 331-340.
- Ducheneaut, N., Yee, N., Nickell, N., & Moore, R. (2006). "Alone Together?" Exploring the Social Dynamics of Massively Multiplayer Online Games. In *Proc CHI 2006*, ACM Press, 407-416.
- Dundes, A. (1965). Here I Sit: A Study of American Latrinalia. *Kroeber Anthropological Society Papers* 34, 91-105.
- Giannetti, F. (2014). Digital Media Reviews: SoundCloud. *MLA Notes* 70, 499-503.
- Goffman, E. (1963). *Behavior in Public Places*. New York: The Free Press.
- Gómez, V., Kaltenbrunner, A., & López, V. (2008). Statistical Analysis of the Social Network and Discussion Threads in Slashdot. In *Proc. WWW 2008*, ACM Press, 645-654.
- Graham, J. (2013, Jun. 17). Who's Listening to SoundCloud? 200 Million. *USA Today*.
- Guimarães, R.L., Cesar, P., & Bulterman, D.C. (2012). 'Let Me Comment on Your Video': Supporting Personalized End-User Comments within Third-Party Online Videos. In *Proc. WebMedia '12*, ACM Press, 253-260.
- Haimson, O., & Tang, J. (2017). What Makes Live Events Engaging on Facebook Live, Periscope, and Snapchat. In *Proc. CHI 2017*, ACM Press, 48-60.
- Hamilton, W., Garretson, O., & Kerne, A. (2014). Streaming on Twitch: Fostering Participatory Communities of Play within Live Mixed Media. In *Proc. CHI 2014*, ACM Press, 1315-1324.
- Hill, B. (2012, Dec. 7). Editorial: The New SoundCloud Is Optimized for Listeners, Not Uploaders. *Engadget*. <http://www.engadget.com/2012/12/07/editorial-soundcloud-optimized-for-listeners-not-artists/>
- Hoare, M., Benford, S., Jones, R., & Milic-Frayling, N. (2014). Coming in from the Margins: Amateur Musicians in the Online Age. In *Proc. CHI 2014*, ACM Press, 1295-1304.
- Honeycutt, C., & Herring, S. C. (2009). Beyond Microblogging: Conversation and Collaboration via Twitter. In *Proc. HICSS 2009*, 1-10.
- Hu, X. & Downie, J. S. (2007). Exploring Mood Metadata: Relationships with Genre, Artist and Usage Metadata. In *Proc. ISMIR 2007*, 67-72.
- Ishizaki, H., Herring, S. C., Hattori, G., & Takishima, Y. (2015). Understanding Behavior on Online Music Distribution Sites: A Discourse Approach. In *Proc. iConference 2015*, IDEALS.
- Jourdanous, A., Allington, D., & Dueck, B. (2014). Using Online Networks to Analyse the Value of Electronic

- Music. In *Proceedings of the 5th International Conference on Computational Creativity*.
- Jourdanous, A., Allington, D., & Dueck, B. (2015). Measuring Cultural Value Using Social Network Analysis: A Case Study on Valuing Electronic Musicians. In *Proceedings of the 6th International Conference on Computational Creativity*, 110-117.
- Kaminskas, M. & Ricci, F. (2012). Contextual Music Information Retrieval and Recommendation: State of the Art and Challenges. *Computer Science Review* 6, 89-119.
- Karpf, D. (2012). Social Science Research Methods in Internet Time. *Information, Communication, and Society* 15, 639-661.
- Lamere, P. (2011). How Good Is Google's Instant Mix? <http://musicmachinery.com/2011/05/14/how-good-is-googles-instant-mix/>
- Laurier, C., Sordo, M., Serrà, J., & Herrera, P. (2009). Music Mood Representations from Social Tags. In *Proc. ISMIR 2009*, 381-386.
- Lee, J. H. (2010). Analysis of User Needs and Information Features in Natural Language Queries Seeking Music Information. *Journal of the American Society for Information Science and Technology* 61, 1025-1045.
- Lee, J. H., & Waterman, N. M. (2012). Understanding User Requirements for Music Information Services. In *Proc. ISMIR 2012*, 253-258.
- Lee, J. H., Wishkoski, R., Aase, L., Meas, P., & Hubbles, C. (2017). Understanding Users of Cloud Music Services: Selection Factors, Management and Access Behavior, and Perceptions. *Journal of the American Society for Information Science and Technology* 68, 1186-1200.
- Mac, R. (2013, Apr. 26). SoundCloud Manages Criticism As It Moves to Become the YouTube of Audio. *Forbes*.
- Madden, A., Ruthven, I., & McMenemy, D. (2013). A Classification Scheme for Content Analyses of YouTube Videos. *Journal of Documentation* 69, 693-714.
- Matthew, T. (2013, Apr. 10). How to Become a SoundCloud Superstar, One Fake Fan at a Time. 5 *Magazine*. <http://www.5chicago.com/features/april2013/how-to-become-a-fake-soundcloud-superstar/>
- Mishne, G., & Glance, N. (2006). Leave a Reply: An Analysis of Weblog Comments. In *Proc. WWW 2006*, ACM Press.
- Naaman, M., Boase, J., & Lai, C-H. (2010). Is It Really About Me? Message Content in Social Awareness Streams. In *Proc. CSCW 2010*, ACM Press, 189-192.
- Norman, D. (1990). *The Design of Everyday Things*. New York: Doubleday.
- Pierce, D. (2016, Mar. 29). SoundCloud Go: An Audacious Answer to Spotify That's Dying to Stand Out. *Wired*.
- Reid, F., & Reid, D. (2010). The Expressive and Conversational Affordances of Mobile Messaging. *Behaviour & Information Technology*, 29, 3-22.
- Rotman, D., Golbeck, J., & Preece, J. (2009). The Community Is Where the Rapport Is – On Sense and Structure in the YouTube Community. In *Proc. C&T '09*, ACM Press, 41-49.
- Rotman, D., & Preece, J. (2010). The 'WeTube' in YouTube – Creating an Online Community through Video Sharing. *International Journal of Web Based Communities* 6, 317-333.
- Russell, J. (2013, May 9). Global Video Site Viki Is Bringing Social Back, As It Reintroduces Improved Real-Time Commenting System. *Thenextweb*. <http://thenextweb.com/media/2013/05/09/global-video-site-viki-is-bringing-back-social-rolls-out-improved-real-time-commenting-system>
- Shoham, M., Arora, A., & Al-Busaidi, A. (2013). Writing on the Wall: An Online "Community" of Patrons as Communication Network or Cyber-Graffiti? In *Proc. HICSS 2013*, 3951-3960.
- Sisario, B. (2014, Aug. 21). Popular and Free, SoundCloud Is Now Ready for Ads. *New York Times*.
- Thelwall, M., Sud, P., & Vis, F. (2012). Commenting on YouTube Videos: From Guatemalan Rock to El Big Bang. *Journal of the American Society for Information Science and Technology* 63, 616-629.
- Tian, Y., Srivastava, J., Huang, T., & Contractor, N. (2010). Social Multimedia Computing. *IEEE Computer* 43, 27-36.
- Tsou, A., Thelwall, M., Mongeon, P., & Sugimoto, C. A. (2014). Community of Curious Souls: An Analysis of Commenting Behavior on TED Talks Videos. *PLoS ONE* 9:4.
- Van Buskirk, E. (2009, Jul. 6). SoundCloud Threatens MySpace as Music Destination for Twitter Era. *Wired*.
- Weisz, J. D., Kiesler, S., Zhang, H., Ren, Y., Kraut, R.E., & Konstan, J.A. (2007). Watching Together: Integrating Text Chat with Video. In *Proc. CHI 2007*, ACM Press, 877-886.
- Xu, J., Compton, R., Lu, T., & Allen, D. (2014). Rolling through Tumblr: Characterizing Behavioral Patterns of the Microblogging Platform. In *Proc. WebSci '14*, ACM Press, 13-22.
- Yadati, K., Larson, M., Liem, C., & Hanjalic, A. (2014). Detecting Drops in Electronic Dance Music: Content Based Approaches to a Socially Significant Music Event. In *Proc. ISMIR 2014*, 143-148.
- Zhao, S. (2003). Toward a Taxonomy of Copresence. *Presence: Teleoperators & Virtual Environments* 12, 445-455.