

# Recognizing Musical Entities in User-generated Content

Lorenzo Porcaro<sup>1</sup> and Horacio Saggion<sup>2</sup>

<sup>1</sup> Music Technology Group, Universitat Pompeu Fabra

<sup>2</sup> TALN Natural Language Processing Group, Universitat Pompeu Fabra

**Abstract.** Recognizing Musical Entities is important for Music Information Retrieval (MIR) since it can improve the performance of several tasks such as music recommendation, genre classification or artist similarity. However, most entity recognition systems in the music domain have concentrated on formal texts (e.g. artists biographies, encyclopedic articles, etc.), ignoring rich and noisy user-generated content. In this work, we present a novel method to recognize musical entities in Twitter content generated by users following a classical music radio channel. Our approach takes advantage of both formal radio schedule and users tweets to improve entity recognition. We instantiate several machine learning algorithms to perform entity recognition combining task-specific and corpus-based features. We also show how to improve recognition results by jointly considering formal and user-generated content.

**Keywords:** Named Entity Recognition · Music Information Retrieval · User-generated Content.

## 1 Introduction

The increasing use of social media and microblogging services has broken new ground in the field of Information Extraction (IE) from user-generated content (UGC). Understanding the information contained in users' content has become one of the main goal for many applications, due to the uniqueness and the variety of this data [1]. However, the highly informal and noisy status of these sources makes it difficult to apply techniques proposed by the NLP community for dealing with formal and structured content [2].

In this work, we analyze a set of tweets related to a specific classical music radio channel, *BBC Radio 3*<sup>1</sup>, interested in detecting two types of musical named entities, *Contributor* and *Musical Work*.

The method proposed makes use of the information extracted from the radio schedule for creating links between users' tweets and tracks broadcasted. Thanks to this linking, we aim to detect when users refer to entities included into the schedule. Apart from that, we consider a series of linguistic features, partly taken from the NLP literature and partly specifically designed for this task, for building

---

<sup>1</sup> [www.twitter.com/BBCRadio3](http://www.twitter.com/BBCRadio3)

statistical models able to recognize the musical entities. To that aim, we perform several experiments with a supervised learning model, Support Vector Machine (SVM), and a recurrent neural network architecture, a bidirectional LSTM with a CRF layer (biLSTM-CRF).

The contributions in this work are summarized as follows:

- A method to recognize musical entities from user-generated content which combines contextual information (i.e. radio schedule) with Machine Learning models for improving the accuracy while recognizing the entities.
- The release of language resources such as an user-generated and bot-generated Twitter corpora manually annotated, usable for both MIR and NLP researches, and domain specific word embeddings.

The paper is structured as follows. In Section 2, we present a review of the previous works related to Named Entity Recognition, focusing on its application on UGC and MIR. Afterwards, in Section 3 it is presented the methodology of this work, describing the dataset and the method proposed. In Section 4, the results obtained are shown. Finally, in Section 5 conclusions are discussed.

## 2 Related Work

Named Entity Recognition (NER), or alternatively Named Entity Recognition and Classification (NERC), is the task of detecting entities in an input text and to assign them to a specific class. It starts to be defined in the early '80, and over the years several approaches have been proposed [3]. Early systems were based on handcrafted rule-based algorithms, while recently several contributions by Machine Learning scientists have helped in integrating probabilistic models into NER systems.

In particular, new developments in neural architectures have become an important resource for this task. Their main advantages are that they do not need language-specific knowledge resources [4], and they are robust to the noisy and short nature of social media messages [5]. Indeed, according to a performance analysis of several Named Entity Recognition and Linking systems presented in [6], it has been found that poor capitalization is one of the main issues when dealing with microblog content. Apart from that, typographic errors and the ubiquitous occurrence of out-of-vocabulary (OOV) words also cause drops in NER recall and precision, together with shortenings and slang, particularly pronounced in tweets.

**Table 1.** Examples of user-generated tweets.

1	No Schoenberg or Webern?? Beethoven is there but not his pno sonata op. 101??
2	Heard some of Opera 'Oberon' today... Weber... Only a little....
3	Cavalleria Rusticana...hm..from a Competition that very nearly didn't get entered!

Music Information Retrieval (MIR) is an interdisciplinary field which borrows tools of several disciplines, such as signal processing, musicology, machine learning, psychology and many others, for extracting knowledge from musical objects (be them audio, texts, etc.) [7]. In the last decade, several MIR tasks have benefited from NLP, such as sound and music recommendation [8], automatic summary of song review [9], artist similarity [10] and genre classification [11].

In the field of IE, a first approach for detecting musical named entities from raw text, based on Hidden Markov Models, has been proposed in [12]. In [13], the authors combine state-of-the-art Entity Linking (EL) systems to tackle the problem of detecting musical entities from raw texts. The method proposed relies on the *argumentum ad populum* intuition, so if two or more different EL systems perform the same prediction in linking a named entity mention, the more likely this prediction is to be correct. In detail, the off-the-shelf systems used are: DBpedia Spotlight [14], TagMe [15], Babelify [16]. Moreover, a first Musical Entity Linking, MEL<sup>1</sup> has been presented in [17] which combines different state-of-the-art NLP libraries and SimpleBrainz, an RDF knowledge base created from MusicBrainz<sup>2</sup> after a simplification process.

Furthermore, *Twitter* has also been at the center of many studies done by the MIR community. As example, for building a music recommender system [18] analyzes tweets containing keywords like *nowplaying* or *listeningto*. In [10], a similar dataset it is used for discovering cultural listening patterns. Publicly available *Twitter* corpora built for MIR investigations have been created, among others the *Million Musical Tweets* dataset<sup>3</sup> [19] and the *#nowplaying* dataset<sup>4</sup> [20].

### 3 Methodology

We propose a hybrid method which recognizes musical entities in UGC using both contextual and linguistic information. We focus on detecting two types of entities:

- *Contributor*: person who is related to a musical work (composer, performer, conductor, etc).
- *Musical Work*: musical composition or recording (symphony, concerto, overture, etc).

As case study, we have chosen to analyze tweets extracted from the channel of a classical music radio, *BBC Radio 3*. The choice to focus on classical music has been mostly motivated by the particular discrepancy between the informal language used in the social platform and the formal nomenclature of contributors and musical works. Indeed, users when referring to a musician or to a classical piece in a tweet, rarely use the full name of the person or of the work, as shown in Table 2.

---

<sup>1</sup> <http://mel.mtg.upf.edu>

<sup>2</sup> <https://musicbrainz.org>

<sup>3</sup> <http://www.cp.jku.at/datasets/MMTD>

<sup>4</sup> <http://dbis-nowplaying.uibk.ac.at>

**Table 2.** Example of entities annotated and corresponding formal forms, from the user-generated tweet (1) in Table 1.

<i>Informal form</i>	<i>Formal form</i>
Schoenberg Webern Beethoven pno sonata op. 101	Arnold Franz Walter Schoenberg Anton Friedrich Wilhelm Webern Ludwig Van Beethoven Piano Sonata No. 28 in A major, Op. 101

We extract information from the radio schedule for recreating the musical context to analyze user-generated tweets, detecting when they are referring to a specific work or contributor recently played. We manage to associate to every track broadcasted a list of entities, thanks to the tweets automatically posted by the *BBC Radio3 Music Bot*<sup>1</sup>, where it is described the track actually on air in the radio. In Table 3, examples of bot-generated tweets are shown.

**Table 3.** Examples of bot-generated tweets.

1	Now Playing Joaquín Rodrigo, Goran Listes - 3 Piezas españolas for guitar #joaquinrodrigo,#goranlistes
2	Now Playing Robert Schumann, Luka Mitev - Phantasiestcke, Op 73 #robertschumann,#lukamitev
3	Now Playing Pyotr Ilyich Tchaikovsky, MusicAeterna - Symphony No.6 in B minor #pyotrilyichtchaikovsky, #musicaeterna

Afterwards, we detect the entities on the user-generated content by means of two methods: on one side, we use the entities extracted from the radio schedule for generating candidates entities in the user-generated tweets, thanks to a matching algorithm based on time proximity and string similarity. On the other side, we create a statistical model capable of detecting entities directly from the UGC, aimed to model the informal language of the raw texts. In Figure 1, an overview of the system proposed is presented.

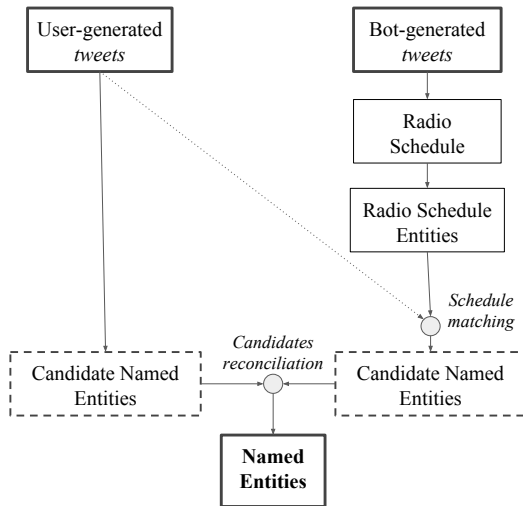
### 3.1 Dataset

In May 2018, we crawled Twitter using the Python library *Tweepy*<sup>2</sup>, creating two datasets on which *Contributor* and *Musical Work* entities have been manually annotated, using IOB tags.

The first set contains user-generated tweets related to the *BBC Radio 3* channel. It represents the source of user-generated content on which we aim to predict the named entities. We create it filtering the messages containing hashtags related to *BBC Radio 3*, such as *#BBCRadio3* or *#BBCR3*. We obtain a

<sup>1</sup> <https://twitter.com/BBCR3MusicBot>

<sup>2</sup> <http://www.tweepy.org>



**Fig. 1.** Overview of the NER system proposed

set of 2,225 unique user-generated tweets. The second set consists of the messages automatically generated by the *BBC Radio 3 Music Bot*. This set contains 5,093 automatically generated tweets, thanks to which we have recreated the schedule<sup>1</sup>.

In Table 4, the amount of tokens and relative entities annotated are reported for the two datasets. For evaluation purposes, both sets are split in a training part (80%) and two test sets (10% each one) randomly chosen. Within the user-generated corpora, entities annotated are only about 5% of the whole amount of tokens. In the case of the automatically generated tweets, the percentage is significantly greater and entities represent about the 50%.

**Table 4.** Tokens’ distributions within the two datasets: user-generated tweets (top) and bot-generated tweets (bottom)

	Training	TestA	TestB
User-generated tweets			
<i>Contributor</i>	1.069 (3,12%)	119 (2,96%)	127 (2,97%)
<i>Musical Work</i>	964 (2,81%)	118 (2,93%)	163 (3,81%)
Total tokens	34.247	4.016	4.275
Bot-generated tweets			
<i>Contributor</i>	15.162 (27,50%)	1.852 (22,93%)	1.879 (27,30%)
<i>Musical Work</i>	12.904 (23,40%)	1.625 (23,56%)	1.689 (24,48%)
Total tokens	55.122	6.897	6.881

<sup>1</sup> <https://github.com/LPorcaro/musicner>

### 3.2 NER system

According to the literature reviewed, state-of-the-art NER systems proposed by the NLP community are not tailored to detect musical entities in user-generated content. Consequently, our first objective has been to understand how to adapt existing systems for achieving significant results in this task.

**Table 5.** Example of musical named entities annotated

Beethoven	is	there	but	not	his	pno	sonata	op.	101
B-CONTR	O	O	O	O	O	B-WORK	I-WORK	I-WORK	I-WORK

In the following sections, we describe separately the features, the word embeddings and the models considered. All the resources used are publicly available<sup>1</sup>.

**Features’ description** We define a set of features for characterizing the text at the token level. We mix standard linguistic features, such as Part-Of-Speech (POS) and chunk tag, together with several gazetteers specifically built for classical music, and a series of features representing tokens’ left and right context. For extracting the POS and the chunk tag we use the Python library *twitter\_nlp*<sup>2</sup>, presented in [2].

In total, we define 26 features for describing each token: 1)POS tag; 2)Chunk tag; 3)Position of the token within the text, normalized between 0 and 1; 4)If the token starts with a capital letter; 5)If the token is a digit. Gazetteers: 6)Contributor first names; 7)Contributor last names; 8)Contributor types (“soprano”, “violinist”, etc.); 9)Classical work types (“symphony”, “overture”, etc.); 10)Musical instruments; 11)Opus forms (“op”, “opus”); 12)Work number forms (“no”, “number”); 13)Work keys (“C”, “D”, “E”, “F”, “G”, “A”, “B”, “flat”, “sharp”); 14)Work Modes (“major”, “minor”, “m”). Finally, we complete the tokens’ description including as token’s features the surface form, the POS and the chunk tag of the previous and the following two tokens (12 features).

**Word embedding** We consider two sets of GloVe word embeddings [21] for training the neural architecture, one pre-trained with 2B of tweets, publicly downloadable<sup>3</sup>, one trained with a corpora of 300K tweets collected during the 2014-2017 *BBC Proms Festivals* and disjoint from the data used in our experiments.

**Models** The first model considered for this task has been the John Platt’s sequential minimal optimization algorithm for training a support vector classifier [22], implemented in WEKA [23]. Indeed, in [24] results shown that SVM

<sup>1</sup> <https://github.com/LPorcaro/musicner>

<sup>2</sup> [https://github.com/aritter/twitter\\_nlp](https://github.com/aritter/twitter_nlp)

<sup>3</sup> <https://github.com/stanfordnlp/GloVe>

outperforms other machine learning models, such as Decision Trees and Naive Bayes, obtaining the best accuracy when detecting named entities from the user-generated tweets.

However, recent advances in Deep Learning techniques have shown that the NER task can benefit from the use of neural architectures, such as biLSTM-networks [4, 5]. We use the implementation<sup>1</sup> proposed in [25] for conducting three different experiments. In the first, we train the model using only the word embeddings as feature. In the second, together with the word embeddings we use the POS and chunk tag. In the third, all the features previously defined are included, in addition to the word embeddings. For every experiment, we use both the pre-trained embeddings and the ones that we created with our *Twitter* corpora. In section 4, results obtained from the several experiments are reported.

### 3.3 Schedule matching

The bot-generated tweets present a predefined structure and a formal language, which facilitates the entities detection. In this dataset, our goal is to assign to each track played on the radio, represented by a tweet, a list of entities extracted from the tweet raw text. For achieving that, we experiment with the algorithms and features presented previously, obtaining an high level of accuracy, as presented in section 4. The hypothesis considered is that when a radio listener posts a tweet, it is possible that she is referring to a track which has been played a relatively short time before. In this cases, we want to show that knowing the radio schedule can help improving the results when detecting entities.

Once assigned a list of entities to each track, we perform two types of matching. Firstly, within the tracks we identify the ones which have been played in a fixed range of time ( $t$ ) before and after the generation of the user’s tweet. Using the resulting tracks, we create a list of candidates entities on which performing string similarity. The score of the matching based on string similarity is computed as the ratio of the number of tokens in common between an entity and the input tweet, and the total number of token of the entity:

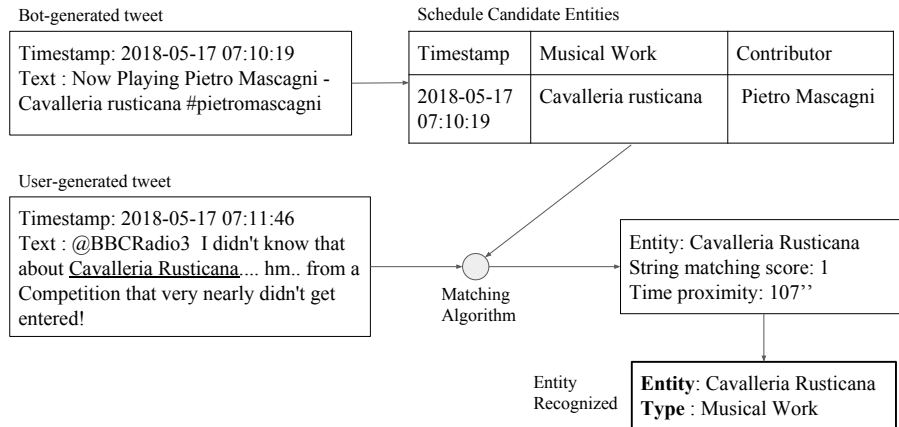
$$score_{string\ matching}(Entity) = \frac{\#(TokensEntity \cap input\ tweet)}{\#TokensEntity} \quad (1)$$

In order to exclude trivial matches, tokens within a list of stop words are not considered while performing string matching. The final score is a weighted combination of the string matching score and the time proximity of the track, aimed to enhance matches from tracks played closer to the time when the user is posting the tweet.

The performance of the algorithm depends, apart from the time proximity threshold  $t$ , also on other two thresholds related to the string matching, one for the *Musical Work* ( $w$ ) and one for the *Contributor* ( $c$ ) entities. It has been necessary for avoiding to include candidate entities matched against the schedule

<sup>1</sup> <https://github.com/UKPLab/emnlp2017-bilstm-cnn-crf>

with a low score, often source of false positives or negatives. Consequently, as last step *Contributor* and *Musical Work* candidates entities with respectively a string matching score lower than  $c$  and  $w$ , are filtered out. In Figure 2, an example of *Musical Work* entity recognized in an user-generated tweet using the schedule information is presented.



**Fig. 2.** Example of the workflow for recognizing entities in UGC using the information from the radio schedule

**Candidates Reconciliation** The entities recognized from the schedule matching are joined with the ones obtained directly from the statistical models. In the joined results, the criteria is to give priority to the entities recognized from the machine learning techniques. If they do not return any entities, the entities predicted by the schedule matching are considered. Our strategy is justified by the poorer results obtained by the NER based only on the schedule matching, compared to the other models used in the experiments, to be presented in the next section.

## 4 Results

The performances of the NER experiments are reported separately for three different parts of the system proposed.

Table 6 presents the comparison of the various methods while performing NER on the bot-generated corpora and the user-generated corpora. Results shown that, in the first case, in the training set the F1 score is always greater



**Table 6.** F1 score for *Contributor*(C) and *Musical Work*(MW) entities recognized from bot-generated tweets (top) and user-generated tweets (bottom)

Model	Features	GloVe vectors	Training		TestA		TestB	
<i>Bot-generated tweets</i>			<i>C</i>	<i>MW</i>	<i>C</i>	<i>MW</i>	<i>C</i>	<i>MW</i>
SVM	all	–	99.12	97.70	97.74	94.32	97.88	95.42
biLSTM-CRF	–	trained	98.95	97.07	98.06	92.99	98.33	95.59
		pre-trained	99.34	94.94	97.88	91.40	98.27	92.35
biLSTM-CRF	POS+chunk	trained	<b>99.94</b>	98.28	97.99	<b>94.68</b>	98.03	<b>95.97</b>
		pre-trained	99.69	97.23	98.12	93.30	98.49	93.61
biLSTM-CRF	all	trained	99.80	98.22	97.70	91.99	98.36	94.48
		pre-trained	99.90	<b>99.40</b>	<b>98.24</b>	90.46	<b>98.78</b>	94.23
<i>User-generated tweets</i>								
SVM	all	–	95.44	80.80	64.91	<b>33.48</b>	61.02	36.21
biLSTM-CRF	–	trained	79.09	51.51	60.00	26.66	67.02	31.48
		pre-trained	85.51	69.28	70.00	33.33	<b>71.26</b>	32.08
biLSTM-CRF	POS+chunk	trained	79.37	50.90	61.23	28.98	62.03	<b>40.00</b>
		pre-trained	73.51	37.28	<b>71.62</b>	25.00	63.74	25.53
biLSTM-CRF	all	trained	97.42	<b>88.92</b>	66.22	28.17	69.11	36.36
		pre-trained	<b>98.46</b>	87.35	68.79	23.68	70.41	29.51

than 97%, with a maximum of 99.65%. With both test sets performances decrease, varying between 94-97%. In the case of UGC, comparing the F1 score we can observe how performances significantly decrease. It can be considered a natural consequence of the complex nature of the users' informal language in comparison to the structured message created by the bot.

In Table 7, results of the schedule matching are reported. We can observe how the quality of the linking performed by the algorithm is correlated to the choice of the three thresholds. Indeed, the *Precision* score increase when the time threshold decrease, admitting less candidates as entities during the matching, and when the string similarity thresholds increase, accepting only candidates with an higher degree of similarity. The behaviour of the *Recall* score is inverted.

**Table 7.** *Precision* (P), *Recall* (R) and *F1* score for *Contributor* (C) and *Musical Work* (MW) of the schedule matching algorithm. *w* indicates the *Musical Work* string similarity threshold, *c* indicates the *Contributor* string similarity threshold and *t* indicates the time proximity threshold in seconds

		<i>t</i> =800			<i>t</i> =1000			<i>t</i> =1200		
		P	R	F1	P	R	F1	P	R	F1
w=0.33, c=0.33	<i>C</i>	72.49	16.49	26.87	69.86	17.57	28.08	68.66	<b>17.93</b>	<b>28.43</b>
	<i>MW</i>	26.42	4.78	8.10	26.05	5.29	8.79	23.66	5.29	8.65
w=0.33, c=0.5	<i>C</i>	<b>76.77</b>	14.32	24.14	74.10	15.64	25.83	73.89	16.00	26.30
	<i>MW</i>	27.1	4.95	8.37	26.67	5.46	<b>9.06</b>	24.24	<b>5.46</b>	8.91
w=0.5, c=0.5	<i>C</i>	<b>76.77</b>	14.32	24.14	74.71	15.64	25.87	73.89	16.00	26.30
	<i>MW</i>	<b>30.43</b>	4.78	8.26	30.30	5.12	8.76	27.52	5.12	8.63

**Table 8.** *Precision (P), Recall (R) and F1 score for Contributor (C) and Musical Work (MW) entities recognized from user-generated tweets using the biLSTM-CRF network together with the schedule matching. The thresholds used for the matching are  $t=1200$ ,  $w=0.5$ ,  $c=0.5$ .*

		Training			TestA			TestB		
		P	R	F1	P	R	F1	P	R	F1
biLSTM-CRF	<i>C</i>	<b>98.22</b>	96.64	<b>97.42</b>	69.01	63.64	66.22	<b>67.35</b>	70.97	<b>69.11</b>
	<i>MW</i>	<b>91.54</b>	86.44	<b>88.92</b>	<b>43.48</b>	20.83	28.17	<b>45.83</b>	30.14	36.36
biLSTM-CRF + Sch. Matcher	<i>C</i>	95.92	<b>97.81</b>	96.86	<b>74.19</b>	<b>71.88</b>	<b>73.02</b>	63.29	<b>74.63</b>	68.49
	<i>MW</i>	87.33	<b>87.03</b>	87.18	38.46	<b>22.73</b>	<b>28.57</b>	42.55	<b>32.26</b>	<b>36.70</b>

Finally, we test the impact of using the schedule matching together with a biLSTM-CRF network. In this experiment, we consider the network trained using all the features proposed, and the embeddings not pre-trained. Table 8 reports the results obtained. We can observe how generally the system benefits from the use of the schedule information. Especially in the testing part, where the neural network recognizes with less accuracy, the explicit information contained in the schedule can be exploited for identifying the entities at which users are referring while listening to the radio and posting the tweets.

## 5 Conclusion

We have presented in this work a novel method for detecting musical entities from user-generated content, modelling linguistic features with statistical models and extracting contextual information from a radio schedule. We analyzed tweets related to a classical music radio station, integrating its schedule to connect users’ messages to tracks broadcasted. We focus on the recognition of two kinds of entities related to the music field, *Contributor* and *Musical Work*.

According to the results obtained, we have seen a pronounced difference between the system performances when dealing with the *Contributor* instead of the *Musical Work* entities. Indeed, the former type of entity has been shown to be more easily detected in comparison to the latter, and we identify several reasons behind this fact. Firstly, *Contributor* entities are less prone to be shorten or modified, while due to their longness, *Musical Work* entities often represent only a part of the complete title of a musical piece. Furthermore, *Musical Work* titles are typically composed by more tokens, including common words which can be easily misclassified. The low performances obtained in the case of *Musical Work* entities can be a consequences of these observations. On the other hand, when referring to a *Contributor* users often use only the surname, but in most of the cases it is enough for the system to recognizing the entities.

From the experiments we have seen that generally the biLSTM-CRF architecture outperforms the SVM model. The benefit of using the whole set of features is evident in the training part, but while testing the inclusion of the features not always leads to better results. In addition, some of the features designed in

our experiments are tailored to the case of classical music, hence they might not be representative if applied to other fields. We do not exclude that our method can be adapted for detecting other kinds of entity, but it might be needed to redefine the features according to the case considered. Similarly, it has not been found a particular advantage of using the pre-trained embeddings instead of the one trained with our corpora. Furthermore, we verified the statistical significance of our experiment by using Wilcoxon Rank-Sum Test, obtaining that there have been not significant difference between the various model considered while testing.

The information extracted from the schedule also present several limitations. In fact, the hypothesis that a tweet is referring to a track broadcasted is not always verified. Even if it is common that radios listeners do comments about tracks played, or give suggestion to the radio host about what they would like to listen, it is also true that they might refer to a *Contributor* or *Musical Work* unrelated to the radio schedule.

**Acknowledgments** This work is partially supported by the European Commission under the TROMPA project (H2020 770376), and by the Spanish Ministry of Economy and Competitiveness under the Maria de Maeztu Units of Excellence Programme (MDM-2015-0502)

## References

1. Habib, M.B., Keulen, M.V: Information Extraction for Social Media. SWAIE@COLING (2014)
2. Ritter, A., Clark, S., Etzioni, O.: Named entity recognition in tweets: an experimental study. In: Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing, pp. 1524–1534 (2011)
3. Nadeau, D., Sekine, S.: A survey of named entity recognition and classification. *Linguisticae Investigationes*, **30**(1), 3–26 (2007)
4. Lample, G., Ballesteros, M., Subramanian, S., Kawakami, K., Dyer, C.: Neural Architectures for Named Entity Recognition. In: Proceedings of NAACL-HLT 2016, pp. 260–270 (2016)
5. Lin, B. Y., Xu, F. F., Luo, Z., Zhu, K. Q.: Multi-channel BiLSTM-CRF Model for Emerging Named Entity Recognition in Social Media. In: Proceedings of the 3rd Workshop on Noisy User-Generated Text, pp. 160–165 (2017)
6. Derczynski, L., Maynard, D., Rizzo, G., Van Erp, M., Gorrell, G., Troncy, R., Bontcheva, K.: Analysis of named entity recognition and linking for tweets. *Information Processing and Management*, **51**(2), pp. 32–49 (2015)
7. Müller, M.: *Fundamentals of Music Processing*. Springer (2015)
8. Oramas, S., Ostuni, V. C., Di Noia, T., Serra, X., Di Sciascio, E.: Sound and Music Recommendation with Knowledge Graphs. *ACM Transactions on Intelligent Systems and Technology*, **8**(2), pp. 1–21 (2015)
9. Tata, S., Di Eugenio, B.: Generating Fine-Grained Reviews of Songs from Album Reviews. In: Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics, pp. 1376–1385 (2010)

10. Schedl, M., Hauger, D.: Mining microblogs to infer music artist similarity and cultural listening patterns. In: Proceedings of the 21st International Conference on World Wide Web, pp. 877–886 (2012)
11. Oramas, S., Espinosa-anke, L., Lawlor, A., Serra, X., Saggion, H.: Exploring Customer Reviews for Music Genre Classification and Evolutionary Studies. In: Proceedings of the 17th International Society for Music Information Retrieval Conference, pp. 150–156 (2016)
12. Zhang, X., Liu, Z., Qiu, H., Fu, Y.: A hybrid approach for chinese named entity recognition in music domain. In: Proceedings of the 8th IEEE International Symposium on Dependable, Autonomic and Secure Computing, pp. 677–681 (2009)
13. Oramas, S., Espinosa-Anke, L., Sordo, M., Saggion, H., Serra, X.: ELMD: An Automatically Generated Entity Linking Gold Standard Dataset in the Music Domain. In: Proceedings of the Language Resources and Evaluation Conference, pp. 3312–3317 (2016)
14. Mendes, P. N., Jakob, M., Garca-silva, A., Bizer, C.: DBpedia Spotlight : Shedding Light on the Web of Documents. In: Proceedings of the 7th International Conference on Semantic Systems, pp. 1–8 (2011)
15. Ferragina, P., Scaiella, U.: Fast and Accurate Annotation of Short Texts with Wikipedia Pages. *IEEE Software*, **29**(1), pp. 70–75 (2012)
16. Moro, A., Raganato, A., Navigli, R.: Entity Linking meets Word Sense Disambiguation: a Unified Approach. *Transactions of the Association for Computational Linguistics*, **2**(0), pp. 231–244 (2014)
17. Oramas, S., Ferraro, A., Correya, A., Serra, X.: Mel: a Music Entity Linking System. In: Proceedings of the 18th International Society for Music Information Retrieval Conference (2017)
18. Zangerle, E., Gassler, W., Specht, G.: Exploiting Twitter’s Collective Knowledge for Music Recommendation. In: Proceedings of the WWW’12 Workshop on ‘Making Sense of Microposts’, pp. 14–17 (2012)
19. Hauger, D., Schedl, M., Koir, A., Tkalcic, M.: The Million Musical Tweets Dataset: What Can We Learn From Microblogs. In: Proceedings of the 14th International Society for Music Information Retrieval Conference (2013)
20. Zangerle, E., Pichl, M., Gassler, W., Specht, G.: #nowplaying Music Dataset: Extracting Listening Behavior from Twitter. In: Proceedings of the 1st International Workshop on Internet-Scale Multimedia Management, pp. 21–26 (2014)
21. Pennington, J., Socher, R., Manning, C.: Glove: Global Vectors for Word Representation. In: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, pp. 1532–1543 (2014)
22. Platt, J.: Fast Training of Support Vector Machines using Sequential Minimal Optimization. *Advances in Kernel Methods - Support Vector Learning* (1998)
23. Frank, E., Hall, M. A., Witten, I. H.: The WEKA Workbench. Online Appendix for “Data Mining: Practical Machine Learning Tools and Techniques”, Morgan Kaufmann, Fourth Edition (2016)
24. Porcaro, L.: Information Extraction from User-generated Content in the Classical Music Domain. Master thesis, Pompeu Fabra University, Barcelona, Spain (2018)
25. Reimers, N., Gurevych, I.: Reporting Score Distributions Makes a Difference: Performance Study of LSTM-networks for Sequence Tagging. In: Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, pp. 338–348 (2017)