

# Natural Language Processing for Music Knowledge Discovery

Sergio Oramas<sup>1</sup>, Luis Espinosa-Anke<sup>2</sup>, Francisco Gómez<sup>3</sup>, and  
Xavier Serra<sup>1</sup>

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

<sup>2</sup>School of Computer Science and Informatics, Cardiff University

<sup>3</sup>Technical University of Madrid

*{sergio.oramas,xavier.serra}@upf.edu,  
espinosa-ankel@cardiff.ac.uk, fmartin@etsisi.upm.es*

## Abstract

Today, a massive amount of musical knowledge is stored in written form, with testimonies dated as far back as several centuries ago. In this work, we present different Natural Language Processing (NLP) approaches to harness the potential of these text collections for automatic music knowledge discovery, covering different phases in a prototypical NLP pipeline, namely corpus compilation, text-mining, information extraction, knowledge graph generation and sentiment analysis. Each of these approaches is presented alongside different use cases (i.e., flamenco, Renaissance and popular music) where large collections of documents are processed, and conclusions stemming from data-driven analyses are presented and discussed.

**Keywords** Musicology, Natural Language Processing, Information Extraction, Entity Linking, Sentiment Analysis

## 1 Introduction

One of the main tasks carried out in musicology is the development and validation of musical hypotheses. The seed that usually leads to most research involves looking for relevant information in written documents, which in general are organized as compilations, collections or anthologies. Today, it is unsurprising to find many of these collections stored in digitized machine-readable format, a scenario which has signified a great improvement on the way information is accessed. These digitized collections are mostly stored in digital libraries and managed by information systems where documents can be searched by textual keywords. This improvement has increased significantly the possibilities for

musicologists to access information. However, in these infrastructures the underlying semantics in the textual content of each document are not captured by search engines, which usually operate at an *exact text string matching* level, and therefore in the majority of cases do not take full advantage of the sophisticated processing tools that semantic search puts at our disposal. In this context, and in order to capture the subtle nuances in musical meaning and thus improving musicological research, we argue that it is not enough to put text corpora online and make them searchable. Indeed, there still remains the important and difficult task to transform text collections, from searchable repositories, into knowledge environments, in what can be seen as the next step in the evolution of digital libraries (Fast & Sedig, 2011).

This limitation coexists with an opportunity derived from the quick growth rate at which online content is generated. Today, specifically in the music domain, we have at our disposal vast amounts of knowledge, gathered for centuries by musicologists and music enthusiasts and made accessible by various agents. Most of this knowledge is encoded in artist biographies, reviews, facsimile editions, and other written media. The constant production of this music-related textual information results in large repositories of knowledge, which have great potential for musicological and philological studies. However, since most of it is recorded in *natural language*, processing and analyzing them effectively is a difficult task. We claim, however, that by leveraging *Natural Language Processing* (NLP) techniques, it is possible to unveil relevant information hidden in large domain-specific document collections, which would otherwise remain hidden.

Fortunately, targeting NLP techniques to text corpora in the music domain has been the main focus of several works so far (Oramas, Espinosa-Anke, Sordo, et al., 2016b; Tata & Di Eugenio, 2010; Oramas, Ostuni, et al., 2016; Sutcliffe et al., 2015; Knees & Schedl, 2011; Sordo et al., 2012; Fujinaga & Weiss, 2016). These and other contributions report experimental results of the application of intelligent text processing techniques to music-specific document collections. In addition, many of the upshots of these methods consist in large structured databases containing musical and musicological information, which can provide search engines with much richer and fine-grained information about musicians, their life and work, and even their relation with other musical entities (musicians, record labels, venues, and so forth). Conversely, information already structured in online knowledge repositories has also been exploited in the context of Computational Musicology. For example, a noteworthy example is provided in Crawford et al. (2014), where musicologists are provided with a means to create a linked and extensible knowledge structure over a collection of Early Music metadata and facsimile images. In Rose & Tuppen (2014), seven big datasets of musical and biographical metadata are aligned, showing how analysis and visualization of such data might transform musicological understanding. Despite these valuable contributions, scant musicological research has been carried out regarding the specific challenge of processing text collections.

We propose to specifically address the above challenge by presenting concrete methodologies aimed at exploring large musicological text corpora. With these methodologies we reconcile, on one the hand, intelligent text processing

techniques, and on the other, musical knowledge acquired both from structured and unstructured resources. First, we distill methods for gathering and combining information coming from different sources. The textual data used in our experiments comes in different flavors, namely (1) A knowledge base of flamenco music; (2) A corpus of biographies from artists of the Renaissance period; and (3) A dataset of music album reviews of diverse genres.

The underlying knowledge expressed in these corpora is thus extracted applying different NLP pipelines. First, shallow text-mining processing techniques are applied to understand main trends in the different schools of the Renaissance period. Second, Information Extraction (IE) techniques constitute the methodological basis for populating a novel fully automatic flamenco knowledge base, and to analogously study migratory tendencies and the role of different European capitals along the Renaissance period. Third, a methodology for the creation of a knowledge graph from a set of unstructured text documents is proposed and evaluated from different standpoints. We further show a direct application of this knowledge graph in automatically computing the ranked relevance of a given artist in the flamenco and Renaissance corpora. Finally, we present an approach for capturing the sentiment expressed in text. Using sentimental information as a starting point, we provide a diachronic study of music criticism via a quantitative analysis of the polarity associated to music album reviews gathered from Amazon<sup>1</sup>. Our analysis hints at a potential correlation between key cultural and geopolitical events and the language and evolving sentiment found in music reviews and, ultimately, opens exciting avenues for diachronic studies of music genres.

This paper is an extended version of two previous publications (Oramas, Gómez, et al., 2015; Oramas, Espinosa-Anke, Lawlor, & Others, 2016) , with the main novel contributions being the unification of approaches, the addition of more detailed results, and the introduction of an additional use case based on the study of the Renaissance Music period. The remainder of this paper is organized as follows. First, in Section 2, we describe the processes of gathering and combining information from different data sources, and apply it to the gathering of the three text corpora used throughout the paper: flamenco music, Renaissance artists, and albums reviews. Then, in Section 3, a text-mining approach based on word frequencies is described and applied to study the different music schools of the Renaissance period. Next, in Section 4, an information extraction pipeline for the extraction of biographical information is exposed and applied to populate a flamenco knowledge base and to study the Renaissance period. Later, in Section 5, an approach for the creation of knowledge graphs is presented and used to compute a relevance ranking of flamenco and Renaissance artists. In Section 6, an aspect-based sentiment analysis method is defined and applied to describe how sentiment associated with music reviews changes over time. Following this method, two experiments are performed, one aggregating sentiment scores by review publication year, and other by album publication year. Finally, we conclude with a discussion about our findings (Section 7).

---

<sup>1</sup><http://www.amazon.com>

## 2 Collecting text corpora

Gathering, structuring, and connecting data from different sources is a research problem in itself, where different and major challenges may arise. Although some existing repositories with music information, such as Wikipedia<sup>2</sup>, Oxford Music Online<sup>3</sup>, or MusicBrainz<sup>4</sup> are quite complete and accurate, there is still a vast amount of music information out there that is generally scattered across different sources on the Web. Selecting the sources and harvesting and combining data is a crucial step towards the creation of practical and meaningful music research corpora (Oramas, 2014)

In this work, three different datasets are built as testbeds of our knowledge extraction methodologies. First, we illustrate in detail a methodology for selecting and mixing data coming from different sources in the creation of a flamenco music knowledge base. Then, we apply some of the described approaches to collect a corpus of artist’s biographies about Renaissance artists and a collection of music album reviews and metadata. In what follows we describe the gathered corpora and the processes carried on for their compilation.

### 2.1 The flamenco corpus

In this section, we describe the methodology used for the creation of a knowledge base of flamenco music. To this end, a large amount of information is gathered from different data sources, and further combined by applying a process of pair-wise entity resolution.

#### 2.1.1 Flamenco music overview

Several musical traditions contributed to the genesis of flamenco music as we know it today. Among them, the influences of the Jews, Arabs, and Spanish folk music are recognizable, but indubitably the imprint of Andalusian Gypsies’ culture is deeply ingrained in flamenco music. The main components of flamenco music are: *cante* or singing, *toque* or guitar playing, and *baile* or dance. According to Gamboa (2005), flamenco music grew out of the singing tradition, as a melting process of all the traditions mentioned above, and therefore the role of the singer soon became dominant and fundamental. *Toque* is subordinated to *cante*, especially in more traditional settings, whereas *baile* enjoys more independence from voice.

In the flamenco jargon, styles are called *palos*. Criteria adopted to define flamenco *palos* are rhythmic patterns, chord progressions, lyrics, poetic structure, and geographical origin. In flamenco, geographical variation is important to classify *cantes* as often they are associated to a particular region where they were originated or where they are performed with gusto. Rhythm or *compás* is a unique feature of flamenco. Rhythmic patterns based on 12-beat cycles are

---

<sup>2</sup><http://wikipedia.org>

<sup>3</sup><http://www.oxfordmusiconline.com/>

<sup>4</sup><http://musicbrainz.org>

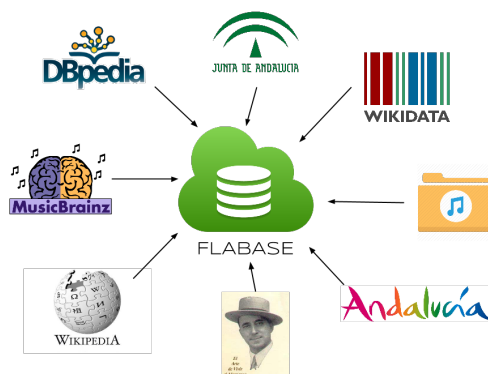


Figure 1: Selected data sources

mainly used. Those patterns can be classed as follows: binary patterns, such as *tangos* or *tientos*; ternary patterns, which are the most common ones, such as *fandangos* or *bulerías*; mixed patterns, where ternary and binary patterns alternate, such as *guajira*; free-form, where there is no a clear underlying rhythm, such as *tonás*. For further information on fundamental aspects of flamenco music, see the book Fernández (2004). For a comprehensive study of styles, musical forms and history of flamenco the reader is referred to the books of Blas Vega & Ríos Ruiz (1988), Navarro & Roperó (1995), and Gamboa (2005), and the references therein.

### 2.1.2 Data acquisition

Our aim is to gather an important amount of information about musical entities (e.g. artists, recordings), including textual descriptions and available metadata. A schema of the selected data sources is shown in Figure 1. We started by looking at Wikipedia. Each Wikipedia article may have a set of associated categories. Categories are intended to group together pages on similar subjects and are structured in a taxonomical way. To find Wikipedia articles related to flamenco music, we first looked for flamenco categories. The taxonomy of categories can be explored by querying DBpedia, a knowledge base with structured content extracted from Wikipedia. We queried the Spanish version of DBpedia<sup>5</sup> for categories related to flamenco. We obtained 17 different categories (e.g., *cantaors de flamenco*, *guitarristas de flamenco*).

We gathered all DBpedia resources related to at least one of these categories. We obtained a total number of 438 resources in Spanish, of which 281 were also in English. Each DBpedia resource is associated with a Wikipedia article. Text and HTML code were then extracted from Wikipedia articles in English and Spanish. Next, we classified the extracted articles according to the role of the biography subject (i.e., *cantaor*, guitarist, and *bailaor*). For this purpose, we

<sup>5</sup><http://es.dbpedia.org>

exploited classification information provided by DBpedia (DBpedia types and Wikipedia categories). At the end, from all gathered resources, we only kept those related to artists and *palos*, totaling 291 artists and 56 *palos*.

As the amount of information present in Wikipedia related to flamenco music is somewhat scarce, we decided to expand our knowledge base with information from two different websites. First, *Andalucia.org*, the touristic web from the Andalusia Government<sup>6</sup>. It contains 422 artist biographies in English and Spanish, and the description of 76 *palos* also in both languages. Second, a website called *El arte de vivir el flamenco*<sup>7</sup>, which includes 749 artist biographies among *cantaors*, *bailaors* and guitarists.

We used MusicBrainz to fill our knowledge base with information about flamenco album releases and recordings. For every artist mapped to MusicBrainz, all content related to releases and recordings was gathered. Thus, 814 releases and 9,942 recordings were collected.

The information gathered from MusicBrainz is a small part of the actual flamenco discography. Therefore, to complement it we used a flamenco recordings database gathered by Rafael Infante and available at CICA website<sup>8</sup> (Computing and Scientific Center of Andalusia). This database has information about releases from the early time of recordings until present time, counting 2,099 releases and 4,136 songs. For every song entry, a *cantaor* name is provided, and most of the times also guitarist and *palo*, which is an important piece of information to define flamenco recordings.

Finally, we supplied our knowledge base with information related to Andalusian towns and provinces. We gathered this information from the official database SIMA<sup>9</sup> (Multi-territorial System of Information of Andalusia).

### 2.1.3 Entity resolution

Entity resolution is the problem of extracting, matching and resolving entity mentions in structured and unstructured data (Getoor, 2012). There are several approaches to tackle the entity resolution problem. For the scope of this research, we selected a pair-wise classification approach based on string similarity between entity labels.

The first issue after gathering the data is to decide whether two entities from different sources are referring to the same one. Therefore, given two sets of entities  $A$  and  $B$ , the objective is to define an injective and non-surjective mapping function  $f$  between  $A$  and  $B$  that decides whether an entity  $a \in A$  is the same as an entity  $b \in B$ . To do that, a string similarity metric  $sim(a, b)$  based on the Ratcliff-Obershelp algorithm (Ratcliff & Metzener, 1988) has been applied. It measures the similarity between two entity labels and outputs a value between 0 and 1. We consider that  $a$  and  $b$  are the same entity if their similarity is bigger than a parameter  $\theta$ . If there are two entities  $b, c \in B$  that

<sup>6</sup><http://andalucia.org>

<sup>7</sup><http://www.elartedevivirelflamenco.com/>

<sup>8</sup><http://flun.cica.es/index.php/grabaciones>

<sup>9</sup><http://www.juntadeandalucia.es/institutodeestadisticaycartografia/sima>




Figure 2:  $F$ -measure for different values of  $\theta$

satisfy that  $sim(a,b) \geq \theta$  and  $sim(a,c) \geq \theta$ , we consider only the mapping with the highest score. To determine the value of  $\theta$ , we tested the method with several  $\theta$  values over an annotated dataset of entity pairs. To create this dataset, the 291 artists gathered from Wikipedia were manually mapped to the 422 artists gathered from Andalucia.org, obtaining a total amount of 120 pair matches. As it is shown in Figure 2 the best  $F$ -measure (0,97) was obtained with  $\theta = 0.9$ . Finally, we applied the described method with  $\theta = 0.9$  to all gathered entities from the three data sources. Thanks to the entity resolution process, we reduced the initial set of 1,462 artists and 132 *palos* to a set of 1,174 artists and 76 *palos*.

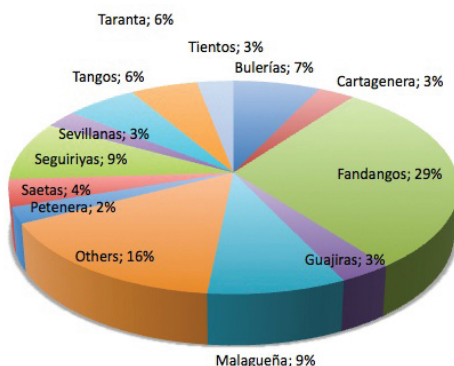
Once we had our artist entities resolved, we began to gather their related discography. First, we tried to find out the MusicBrainz ID of the gathered artists. Depending on the information about the entity, two different processes were applied. First, we leveraged mapping information between Wikipedia and MusicBrainz present in Wikidata<sup>10</sup> (Vrandečić & Krötzsch, 2014). Wikidata is a free linked database, which acts as a structured data storage of Wikipedia. For those artists without this mapping information, we queried the MusicBrainz API, and then applied our entity resolution method to the obtained results.

Finally, to integrate the discography database of CICA into our knowledge base, we applied the entity resolution method to the fields *cantaor*, guitarist and *palo* of each recording entry in the database. From the set of 202 *cantaores* and 157 guitarist names present in the recording entries of the database, a total number of 78 *cantaores* and 44 guitarists were mapped to our knowledge base. The number of mapped artists was low due to differences between the way of labeling an artist. An artist name may be written by using one or two of her surnames, or by using her nickname. In the case of *palos*, there were 162 different *palos* in the database, 54 of which were mapped with the 76 of our knowledge base. These 54 *palos* correspond to an 80% of *palo* assignments present in the recording entries.

#### 2.1.4 FlaBase

FlaBase (Flamenco Knowledge Base) is the acronym of the resulting knowledge base of flamenco music. It contains online editorial, biographical and musicolog-

<sup>10</sup><https://www.wikidata.org>

Figure 3: Songs by *palo*

ical information related to flamenco music. FlaBase is stored in JSON format, and it is freely available for download<sup>11</sup>. FlaBase contains information about 1,174 artists, 76 *palos* (flamenco genres), 2,913 albums, 14,078 tracks, and 771 Andalusian locations. In Figure 3 it is shown that the most representative *palos* in flamenco music are represented in our knowledge base, with a higher predominance of fandangos.

## 2.2 The Renaissance corpus

Renaissance is a period in history that starts around 1400 with the end of the medieval era, and closes around 1600, with the beginning of the Baroque period. Renaissance music refers to music written in Europe during this period.

In this work we experimented with two datasets of biographies about Renaissance composers, one gathered from Wikipedia, and another from The New Grove (Sadie, 2001), available through Oxford Music Online.

### 2.2.1 The Wikipedia corpus

In Wikipedia there is an important number of articles related to Renaissance music, most of them biographies of composers. For this research, we compiled the biographies of all composers that are linked in the Wikipedia page: List of Renaissance composers<sup>12</sup>. In this page, composers are classified by school. We collected biographies of composers from the five most representative schools: Spanish, German, English, Franco-Flemish, and Italian. A total number of 543 biographies were gathered. In addition to the biography texts, HTML links to other Wikipedia pages present in texts were also stored.

<sup>11</sup><http://mtg.upf.edu/download/datasets/flabase>

<sup>12</sup>[https://en.wikipedia.org/wiki/List\\_of\\_Renaissance\\_composers](https://en.wikipedia.org/wiki/List_of_Renaissance_composers)



### 2.2.2 The Grove corpus

The Grove Dictionary of Music and Musicians (Grove, 1878) is an encyclopedic dictionary, and one of the largest reference works in Western music. It was first published in four volumes in the last quarter of the XIX century by George Grove. In 1980 a new version called The New Grove (Sadie, 2001) was released with 20 volumes, where there are 22,500 articles and 16,500 biographies. The complete text of the second edition of The New Grove is available in machine-readable format on the online service Oxford Music Online as Grove Music Online. From this set of biographies, we gathered all of them classified as Early Renaissance and Late Renaissance. A total number of 1710 biographies were collected.

### 2.3 The albums reviews corpus

In this section, we put forward an integration procedure for enriching with music-related information a large dataset of Amazon customer reviews McAuley et al. (2015), with semantic metadata obtained from MusicBrainz. The initial dataset of Amazon customer reviews provides millions of review texts together with additional information such as overall rating (between 0 to 5), date of publication, or creator id. Each review is associated to a product and, for each product, additional metadata is also provided, namely, Amazon product id, list of similar products, price, sell rank, and genre categories. From this initial dataset, we selected the subset of products categorized as *CDs & Vinyls*, which also fulfill the following criteria. First, considering that the Amazon taxonomy of music genres contains 27 labels in the first hierarchy level, and about 500 in total, we obtain a music-relevant subset and select 16 of the 27 which really define a music style and discard for instance region categories (e.g., World Music) and other categories specifically non-related to a music style (e.g., Soundtrack, Miscellaneous, Special Interest), function-oriented categories (Karaoke, Holiday & Wedding), or categories whose albums might also be found under other categories (e.g., Opera & Classical Vocal, Broadway & Vocalists). We compiled albums belonging only to one of the 16 selected categories, i.e., no multi-label. Note that the original dataset contains not only reviews about CDs and Vinyls, but also about music DVDs and VHSs. Since these are not strictly speaking music audio products, we filter out those products also classified as "Movies & TV". Finally, since products classified as Classical and Pop are substantially more frequent in the original dataset, we compensate this unbalance by limiting the number of albums of any genre to 10,000. After this preprocessing, the dataset amounts to a total of 65,566 albums and 263,525 customer reviews. A breakdown of the number of albums per genre is provided in Table 1. The final dataset is called the Multimodal Album Reviews Dataset (MARD) and is freely available for download<sup>13</sup>.

Having performed genre filtering, we enrich the dataset by extracting artist names and record labels from the Amazon product page. We pivot over this

<sup>13</sup><https://www.upf.edu/web/mtg/mard>

<b>Genre</b>	<b>Amazon</b>	<b>MusicBrainz</b>
Alternative Rock	2,674	1,696
Reggae	509	260
Classical	10,000	2,197
R&B	2,114	2,950
Country	2,771	1,032
Jazz	6,890	2,990
Metal	1,785	1,294
Pop	10,000	4,422
New Age	2,656	638
Dance & Electronic	5,106	899
Rap & Hip-Hop	1,679	768
Latin Music	7,924	3,237
Rock	7,315	4,100
Gospel	900	274
Blues	1,158	448
Folk	2,085	848
<b>Total</b>	<b>66,566</b>	<b>28,053</b>

Table 1: Number of albums by genre with information from the different sources in the albums reviews dataset.

information to query the MusicBrainz search API to gather additional metadata such as release id, first release date, song titles and song ids. Mapping with MusicBrainz is performed using the same methodology described in Section 2.1.3, following a pair-wise entity resolution approach based on string similarity with a threshold value of  $\theta = 0.85$ . We successfully mapped 28,053 albums to MusicBrainz.

### 3 Text-mining

Text-mining is the process of deriving high-quality information from text. This high-quality information is typically derived through the devising of patterns and trends using statistical analysis over text. Many text-mining techniques are based on the analysis of frequencies of the words present in the set of studied documents. In what follows, we illustrate the potential of this technique with a simple application to the analysis of word frequencies in our corpus of Renaissance artist’s biographies.

#### 3.1 Renaissance music schools


The computational analysis of artist biographies may reveal interesting insights from the data that can be useful to musicologists. Using the Renaissance artist’s biographies gathered from Wikipedia (see Section 2.2.1), we applied a shallow analysis of the words used in the articles. We computed the frequencies of all

words present in the articles of every school. From the obtained frequencies we plot a word cloud for each school, where more frequent words are represented with bigger fonts. In Figure 4, the word clouds of the different schools are shown. We observe very clear insights from the images at first sight. We see, for instance, how madrigal is very important in the Italian, chanson in the French, and motet in the Franco-Flemish school. We also see the importance of the Church in the Spanish school, or the relevance of organ music in the German school. Although these observations may seem obvious to a musicologist, they are extracted directly from the data without human intervention. This approach can be applied to text corpora the researcher might not be familiar with, helping her in easily discovering some trends directly from the data.




(a) Spanish

(b) Italian



(c) English

(d) Franco-Flemish



(e) French

(f) German

Figure 4: Word clouds by school from Wikipedia biographies.

## 4 Information Extraction

Information extraction is the task of automatically extracting structured information from unstructured or semi-structured text sources. It is a widely studied topic within the NLP research community (Cowie & Lehnert, 1996). A major step towards understanding language is the extraction of meaningful terms (entities) from text as well as relationships between those entities. This statement involves two different tasks. First, the identification and categorization of entity mentions. This task is called named entity recognition (NER). However, when this task involves a latter step of disambiguation of entities against a knowledge base it is called named entity disambiguation or entity linking (EL). The second task consists of the identification of relevant semantic relations or attributes associated to these entities.

### 4.1 Entity linking

The advent of large knowledge repositories and collaborative resources has contributed to the emergence of entity linking, i.e., the task of discovering mentions of entities in text and link them to a suitable knowledge repository (Moro et al., 2014). It encompasses similar subtasks such as named entity disambiguation (Bunescu & Pasca, 2006), which is precisely linking mentions of entities to a knowledge base, or wikification (Mihalcea & Csomai, 2007), specifically using Wikipedia as knowledge base.

Entity linking is typically divided in two steps, namely, the identification of a text span from the text as an entity candidate, and the disambiguation of this entity with respect to a knowledge base. This disambiguation step can be directly applied to the surface form of the identified text span, or to the output of a NER system previously applied. The biggest difference here is that the NER system not only identifies the text span, but also provides a category that classify the identified candidate. We propose a method that employs a combination of both approaches, depending on the category of the entity. For NER, we used the Stanford NER system (Finkel et al., 2005), implemented in the library Stanford Core NLP<sup>14</sup> and trained on English and Spanish texts. For disambiguation we simply looked for exact string matches between entity labels in the knowledge base and identified text spans.

### 4.2 Studying the flamenco corpus

#### 4.2.1 Linking entities in the flamenco corpus

As the gathered flamenco texts are mostly written in Spanish, we needed an entity linking system that deals with Spanish texts. Although there are many entity linking tools available, state-of-the-art systems are well-tuned for English texts, but may not perform as well in languages other than English, and even less with music related texts (Oramas, Espinosa-Anke, Sordo, et al., 2016a). In

<sup>14</sup><http://nlp.stanford.edu/software/corenlp.shtml>

Approach	Precision	Recall	<i>F</i> -measure
1) no NER	0.829	<b>0.694</b>	0.756
2) NER to PERS & LOC	0.739	0.347	0.472
3) NER to LOC	<b>0.892</b>	0.674	<b>0.767</b>

Table 2: Precision, Recall and *F*-measure of entity linking approaches.

addition, we wanted to have a system that uses our own knowledge base for disambiguation. Therefore, we developed our own system, which is able to detect and disambiguate three categories of entities: Person, *Palo* and Location. Three different approaches for the selection of annotation candidates were defined by applying NER only on a subset of the categories of entities: only using text spans (no NER) for disambiguation; disambiguating Location and Person entities from the NER output, and *Palo* from text spans; and only disambiguating Location entities from the NER output, and Person and *Palo* directly from text spans.

To determine which approach performs better, three artist biographies were manually annotated, having a total number of 49 annotated entities. Results on the different approaches are shown in Table 2. We observe that applying NER to entities of the Person category worsens performance significantly, as recall suddenly decreases by half. After manually analyzing false negatives, we observed that this is caused because many artist names have definite articles between name and surname (e.g., *de*, *del*), and this is not recognized correctly by the NER system. In addition, many artists have a nickname that is not interpreted as a Person entity by the NER system. The best approach is the third one (NER to LOC), where NER output is used only for Locations, which is slightly better than the first one (no NER) in terms of precision. This is due to the fact that many artists have a town name as a surname or as part of his or her nickname. Therefore, applying entity linking directly to text spans is misclassifying Person entities as Location entities. Thus, by adding a previous step of NER to Location entities we have increased the overall performance, as it can be seen on the *F*-measure values.

#### 4.2.2 Extracting biographical data

While the created knowledge base of flamenco does already encode relevant culture and music-specific information, a notable portion of the data collected currently remains unexploited due to its unstructured nature. Consequently, to enhance the amount of structured data, a process of information extraction is carried out. We focus on extracting two specific pieces of information from the artist biographies: birth year and birth place, as they can be relevant for anthropological studies. We observed that this information is often in the first sentences of the biographies, and always near the word *nació* (Spanish translation of "was born"). Therefore, to extract this information, we look for this word in the first 250 characters of every biographical text. If it is found, we




Figure 5: FlaBase distributions.

apply our entity linking method to this piece of text. If a Location entity is found near the word "nació", we assume that this entity is the place of birth of the biography subject. In addition, by using regular expressions, we look for the presence of a year expression in the context of the Location entity. If it is found, we assume it as the year of birth. If more than one year is found, we select the one with the smaller value.

To evaluate our approach, we tested the extraction of birth places in all texts coming from the web [Andalucia.org](http://Andalucia.org) (442 artists). We manually annotated the province of provenance of these 442 artists for building ground truth data. After the application of the extraction process on the annotated test set, we obtained a precision value of 0.922 and a recall of 0.648. Therefore, we may argue that our method is extracting biographic information with high precision and quite reasonable recall. We finally applied the extraction process to all artist entities with biographical texts. Thus, 743 birth places and 879 birth years were extracted.

Using the information extracted, we computed the distribution of different items present in FlaBase. Data shown in Figures 5a and 5b was obtained thanks to the information extraction process applied. We can observe in Figure 5a that most flamenco artists are from the Andalusian provinces of Seville and Cadiz. Finally, in Figure 5b we observe a higher number of artists in the data were born from the 30's to the 80's of the 20th century.

### 4.3 Studying the Renaissance period

To study the Renaissance period, we applied a process of information extraction similar to the one described above for the flamenco corpus. Thus, we extracted biographical data from the artist biographies in the Grove corpus. We observed in this corpus that at the beginning of every biography there is a sentence between parentheses with information about the place and date of birth and death. Therefore, we automatically extracted this information using the same ad-hoc entity linking system and regular expressions used to extract information from the flamenco corpus. Using the extracted data, we first plotted the histograms of the distributions of birth and death dates (Figures 6a and 6b). As observed

City	Births	Deaths	Difference
Florence	25	26	+4%
Brescia	18	5	-72%
Parma	15	10	-33%
Nuremberg	15	17	+13%
Bologna	15	13	-13%

Table 3: Top cities by number of births, extracted from the Grove dataset.

City	Births	Deaths	Difference
Rome	9	59	+555%
London	7	39	+457%
Paris	6	32	+433%
Venice	11	29	+164%
Florence	25	26	+4%

Table 4: Top cities by number of deaths, extracted from the Grove dataset.

in Figures 6a and 6b, most Renaissance composers were born in the first half of the XVI century, and died at the beginning of the XVII century. This image gives as a simple overview of the activity in the period.




Figure 6: Distribution of birth and death dates.

Using the extracted places of birth and death, we also computed the difference between cities in number of births and deaths. We observe in Table 3 that Brescia and Parma are cities where many relevant composers were born, but few died. This perhaps implies a good educational environment in music, but less career opportunities for those composers. By contrast, we observe in Table 4 how big cities like Rome, London, Paris, or Venice are attractors of talent, with much larger number of deaths than births. Florence in contrast, typically considered as the cradle of the Renaissance, has a similar number of births and deaths.

City	Median year
Nuremberg	1563
Paris	1569
Venice	1576
Rome	1594
Florence	1597
London	1610

Table 5: Median of the distribution of deaths by city.

Finally, we computed the median of the distribution of death years by city of those with larger number of deaths. This data may be useful to observe when a city was in the middle of his success as an attractor of musical talent. In Table 5, we observe how the gravity center of Renaissance music moves from Nuremberg and Paris to Venice, Florence, and Rome, and finally to London. Again, this result may be very illustrative as a first impression of this musical period.

## 5 Knowledge graph construction

We assume that an entity mention inside an artist biography signals a semantic relation between the entity that constitutes the main theme of the biography (subject entity) and the mentioned entity. Based on this assumption, we build a semantic graph by applying the following steps. First, each artist in the corpus is added to the graph as a node. Second, entity linking is applied to artist’s biographical texts. For every linked entity identified in the biography, a new node is created in the graph (only if it was not previously created). Next, an edge is added, connecting the subject entity with the linked entity found in its biography. This way, a directed graph connecting the entities of the text corpus is obtained.

This graph may have multiple applications. It may be exploited to compute similarity measures between artists, as explored in Oramas, Sordo, et al. (2015), or it may provide a data structure well suited to the implementation of graphical navigational systems throughout the collection of documents, as explored in Oramas et al. (2014). In this work, we explore a different application: the measurement artist relevance.

### 5.1 Artists relevance

Entities identified in a text by an entity linking system may be seen as hyperlinks that connects one text to another. Thus, algorithms to measure the relevance of nodes in a network of hyperlinks can be applied to our semantic graph (Bellomi & Bonato, 2005). Hence, a knowledge graph constructed with the proposed methodology represents a network of hyperlinks that connect the different documents in the corpus. In order to measure artist relevance in our



constructed graph, we applied the PageRank (Brin & Page, 1998) and HITS (Kleinberg, 1999) algorithms. PageRank outputs a measure of relevance for each node, and HITS gives two different results: *authority* and *hubness*. We only take into consideration *authority* from HITS algorithm because it has been proven to be the most effective of both values as a metric of relevancy (Bellomi & Bonato, 2005).

### 5.1.1 Flamenco artists

Following the proposed methodology for the creation of a knowledge graph, we created a graph of flamenco artists after its application to the corpus of artist biographies gathered in FlaBase. We applied the entity linking system described in Section 4.1 and then constructed the graph. In this case, we also added other attributes present in FlaBase to the graph, such as the extracted attributes and the recordings associated to each artist. Once the graph was built, we applied the PageRank and HITS algorithms and built an ordered list with the top-10 entities of the different artist categories (*cantaor*, guitarist and *bailaor*) for each of the algorithms.

For evaluation purposes, we asked a reputed flamenco expert to build a list of top-10 artists for each category according to his knowledge and the available bibliography. The concept of artist relevance is somehow subjective and there is no unified or consensual criterion for flamenco experts about who the most relevant artists of all time are. Despite that, there is a high level of agreement among them on certain artists that should be on such a hypothetical list, based on their influence in the evolution of the genre. Thus, after consulting several documented sources and other flamenco experts, our expert provided us with this list of consensual top-10 artists by category and we considered it as ground truth.

We define precision as the number of identified artists in the resulting list that are also present in the ground truth list divided by the length of the list. We evaluated the output of the two algorithms by calculating precision over the entire list (top-10), and over the first five elements (top-5) (see Table 6). We can observe that PageRank results show the greatest agreement with the flamenco experts list. High values of precision, especially for the top-5 list, indicates that the information in the knowledge graph is highly complete and accurate, and the proposed methodology adequate to compute relevance of artists. In Table 7 the top-5 artists in each category obtained with the PageRank algorithm are shown. It is clear that this approach tend to favor ancient artists that have more probabilities of have been mentioned in other biographies. Therefore, we understand in this case artist relevance as a measure directly tied to the influence of the artist in the evolution of the genre.

### 5.1.2 Renaissance artists

We followed two different strategies for knowledge graph construction for the two datasets of Renaissance artist biographies. For the Wikipedia corpus, we

	Top-5	Top-10
PageRank	0.933	0.633
HITS Authority	0.6	0.4

Table 6: Precision values of artist relevance ranking.

<i>Cantaor</i>	Guitarist	<i>Bailaor</i>
Antonio Mairena	Paco de Lucía	Antonio Ruiz Soler
Manolo Caracol	Ramón Montoya	Rosario
La Niña de los Peines	Niño Ricardo	Antonio Gades
Antonio Chacón	Manolo Sanlúcar	Mario Maya
Camarón de la Isla	Sabicas	Carmen Amaya
Manuel Torre	Tomatito	Pilar López
José Mercé	Vicente Amigo	La Argentinita
Enrique Morente	Gerardo Núñez	Lola Flores
Pepe Marchena	Paco Cepero	Pastora Imperio
Manuel Vallejo	Pepe Habichuela	José Antonio

Table 7: PageRank Top-10 artists by category.

took advantage of the links already present in the Wikipedia pages instead of applying entity linking. We connected the biography main theme entity with all the entities linked in the biography text. There are many entities linked in the biographies that do not correspond to Renaissance composers (e.g., countries, events, kings). Therefore, we created a graph composed only of Renaissance composers and another with all the entities found in the biographies. Figure 7 shows the difference between the two graphs.

Following the same methodology described in Section 5.1, we computed the relevance ranking of the composers in the 2 graphs created from the Wikipedia corpus using the PageRank algorithm. We observe in Table 8 the most relevant composer of each school obtained from the 2 Wikipedia graphs, the one using only links between Renaissance composers (INTERNAL CONNECTIONS) and the one using links to any entity (ALL CONNECTIONS). From a musicological perspective, we observe that the results using only INTERNAL CONNECTIONS have

<b>School</b>	<b>Internal connections</b>	<b>All connections</b>
Spanish	Francisco Guerrero	Juan de la Encina
German	Hans Leo Hessler	Martin Luther
English	Thomas Morley	Henry VIII
Franco-Flemish	Josquin des Prez	Josquin des Prez
Italian	Palestrina	Monteverdi

Table 8: Relevance ranking of composers by school and graph creation approach using the Wikipedia dataset.

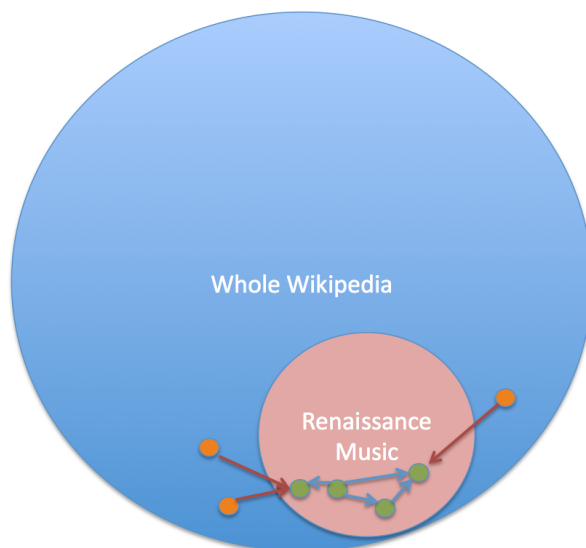


Figure 7: Knowledge graph construction approaches

Ranking	Internal connections	All connections
#1	Josquin des Prez	Henry VIII
#2	Palestrina	Martin Luther
#3	Orlande de Lassus	Henry V
#4	Adrian Willaert	Monteverdi

Table 9: Relevance ranking of all composers by graph creation approach using the Wikipedia dataset.

more sense than those obtained using ALL CONNECTIONS. For example, Henry VIII appears as the most prominent entity of the English school when using ALL CONNECTIONS. Henry VIII, in addition to the king of England, was a composer of the Renaissance era. However, his popularity is mainly due to his role as a king rather than as a composer. Using INTERNAL CONNECTIONS only we obtain Thomas Morley as the most prominent composer of the English school, who is really a cornerstone of this school. The same happens in the German school with Martin Luther, who is popular for other aspects different from music. In the Italian school we observe a slightly different situation. Claudio Monteverdi appears as the most prominent composer using ALL CONNECTIONS. He is actually one of the most prominent composers of the history of music, however, although he started his career in the Renaissance era, he is mostly considered as a Baroque composer. Palestrina, who was obtained using only INTERNAL CONNECTIONS, is also a very prominent composer in the history of music, but he is a prototypical composer of the Renaissance. We can infer from these results that the use of only inner connections helps the approach to obtain results that are

Ranking	Internal connections
#1	Palestrina
#2	Alessandro Damasceni Peretti di Montalto
#3	Petrarch
#4	Claudin de Sermisy
#5	Luca Marenzio
#6	Pierre Sandrin
#7	Jaques Arcadelt
#8	Jacob Obrecht

Table 10: Relevance ranking of all composers using the Grove dataset.

more musicologically meaningful. In Table 9 we observe the top-4 composers from both graphs independently of the music school. We notice here the same tendency in the results.

For the Grove corpus, we followed the same strategy described for the construction of the flamenco knowledge graph. We applied entity linking to the biographies, and then connected each biography subject with the entities mentioned in its biography. We employed a similar ad-hoc approach for entity linking as the one described in Section 4.1.

We computed the ranking list of the most relevant composers applying the PageRank algorithm over this graph. As shown in Table 10, composers obtained in this list are all very relevant musicians of the Renaissance period. However, there are other types of artists and relevant people in the list, similarly to the results obtained with the Wikipedia graph with ALL CONNECTIONS. This implies the need of a filtering process of the selected entities. This result confirms the findings shown in Section 5.1.1, which demonstrate the utility of the proposed approach to compute artist relevance ranking from unstructured texts by using entity linking.

## 6 Sentiment analysis

Sentiment analysis is the task to systematically identify, extract, quantify, and study affective states and subjective information in text. Among the different subtasks of sentiment analysis, we focus in this work on aspect-based sentiment analysis. This technique provides specific sentiment scores for different aspects present in the text, e.g. album cover, guitar, voice or lyrics. These scores represent how much the user likes or dislikes specific attributes expressed in text.

### 6.1 Aspect-based sentiment analysis

Following the work of Dong et al. (2013, 2014) we use a combination of shallow NLP, opinion mining, and sentiment analysis to extract opinionated features

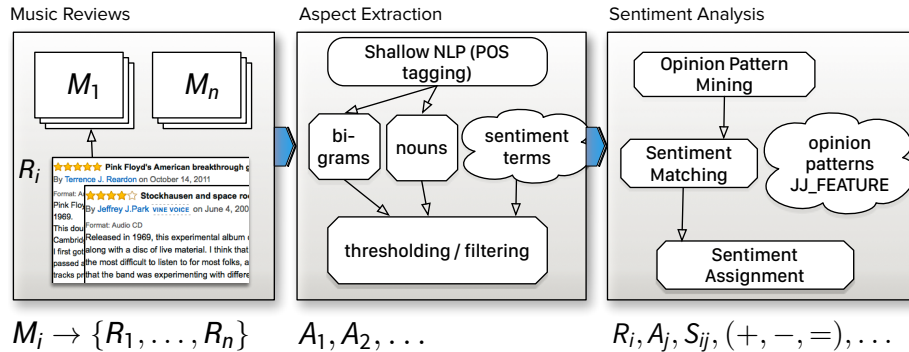


Figure 8: Overview of the opinion mining and sentiment analysis framework.

from reviews. For all reviews  $R_i$  of each album, we mine bi-grams and single-noun aspects (also called review features; see Hu & Liu (2004)). We consider bi-grams that conform to a noun followed by a noun (e.g., *chorus arrangement*) or an adjective followed by a noun (e.g., *original sound*), and excluded bi-grams whose adjective is a sentiment word (e.g., *excellent*, *terrible*). Separately, single-noun aspects are validated by eliminating nouns that are rarely associated with sentiment words in reviews, since such nouns are unlikely to refer to item aspects. We refer to each of these extracted aspects  $A_j$  as review aspects.

For a review aspect  $A_j$  we determine if there are any sentiment words in the sentence containing  $A_j$ . If not,  $A_j$  is marked neutral; otherwise, we identify the sentiment word  $w_{min}$  with the minimum word-distance to  $A_j$ . Next, we determine the part-of-speech tags for  $w_{min}$ ,  $A_i$  and any words that occur between  $w_{min}$  and  $A_i$ . We assign a sentiment score between -1 and 1 to  $A_j$  based on the sentiment of  $w_{min}$ , subject to whether the corresponding sentence contains any negation terms within 4 words of  $w_{min}$ . If there are no negation terms, then the sentiment assigned to  $A_j$  is that of the sentiment word in the sentiment lexicon; otherwise, this sentiment is reversed. Our sentiment lexicon is derived from SentiWordNet (Esuli & Sebastiani, 2006) and is not specifically tuned for music reviews. An overview of the process is shown in Figure 8. The end result of sentiment analysis is that we determine a sentiment score  $S_{ij}$  for each aspect  $A_j$  in review  $R_i$ . A sample annotated review is shown in Figure 9. Finally, the sentiment score of a review  $R_i$  is calculated as the average of the sentiment score  $S_{ij}$  of every aspect  $A_j$  in  $R_i$ .

## 6.2 Diachronic study of music criticism

We applied the proposed aspect-based sentiment analysis framework to the corpus of album customer reviews gathered from Amazon (see Section 2.3), obtaining specific sentiment scores for different aspects present in the text, e.g., album cover, guitar, voice or lyrics. In Figure 10 we observe that the sentiment scores follow a Gaussian distribution, with a median of 0.21, and remarkable picks at




Figure 9: A sentence from a sample review annotated with opinion and aspect pairs.




Figure 10: Distribution of sentiment scores

0 and 0.5.

In addition to the sentiment computed, this corpus includes music metadata such as genre, review rating, review publication date and album release date. We benefit from this substantial amount of information at our disposal for performing a diachronic analysis of music criticism. Specifically, we combine the metadata retrieved for each review with their associated sentiment information, and generate visualizations to help us investigate any potential trends in diachronic music appreciation and criticism. Based on this evidence, and since music evokes emotions through mechanisms that are not unique to music (Juslin & Västfjäll, 2008), we may go as far as using musical information as means for a better understanding of global affairs. Previous studies argue that national confidence may be expressed in any form of art, including music (Moïsi, 2010), and in fact, there is strong evidence suggesting that our emotional reactions to music have important and far-reaching implications for our beliefs, goals and actions, as members of social and cultural groups (Alcorta et al., 2008).

To investigate this matter, we carried out a study of the evolution of music criticism from two different temporal standpoints. Specifically, we consider when the review was written and, in addition, when the album was first published. We define the sentiment score of a review as the average score of all aspects in the review. Since we have sentiment information available for each review, we first computed an average sentiment score for each year of review

publication (between 2000 and 2014). In this way, we may detect any significant fluctuation in the evolution of affective language during the 21st century. Then, we also calculated an average sentiment score by year of album publication. The affective information is complemented with the averages of the Amazon rating scores.

In what follows, we show visualizations for sentiment scores and correlation with ratings given by Amazon users, according to these two different temporal dimensions. Although arriving to musicological conclusions is out of the scope of this paper, we provide *food for thought* and present the readers with hypotheses that may explain some of the facts revealed by these data-driven trends.

### 6.2.1 Evolution by review publication year

We applied sentiment and rating average calculations to the whole dataset, grouping album reviews by year of publication of the review. Figure 11a shows the average of the sentiment scores of all the reviews published in a specific year, whilst Figure 11b shows average review ratings per year. At first sight, we do not observe any correlation between the trends illustrated in the figures. However, the sentiment curve (Figure 11a) shows a remarkable peak in 2008, a slightly lower one in 2013, and a low between 2003 and 2007, and also between 2009 and 2012. Figure 11e shows the kernel density estimation of the distribution of reviews by year of the 16 genres. The shapes of these curves suggest that the 2008 peak in the sentiment score is not related to the number of reviews published that year. The peak persists if we construct the graphs with the average sentiment associated with the most repeated aspects in text (Figure 11d). It is not trivial to give a proper explanation of this variations on the average sentiment. We speculate that these curve fluctuations may suggest some influence of economical or geopolitical circumstances in the language used in the reviews, such as the 2008 election of Barack Obama as president of the US. As stated by the political scientist Dominique Moïsi in Moïsi (2010):

In November 2008, at least for a time, hope prevailed over fear. The wall of racial prejudice fell as surely as the wall of oppression had fallen in Berlin twenty years earlier [...] Yet the emotional dimension of this election and the sense of pride it created in many Americans must not be underestimated.

If we calculate the sentiment evolution curve for the different genres (see Figure 11c), we observe that 2008 constitutes an all-time-high for almost all genres. It is remarkable that genres traditionally related to more diverse communities such as Jazz and Latin Music experience such an increase, whilst other genres such as Country do not.

Another factor that might be related to the positiveness in use of language is the economical situation. After several years of continuous economic growth, in 2007 a global economic crisis started<sup>15</sup>, whose consequences were visible in the

<sup>15</sup><https://research.stlouisfed.org>

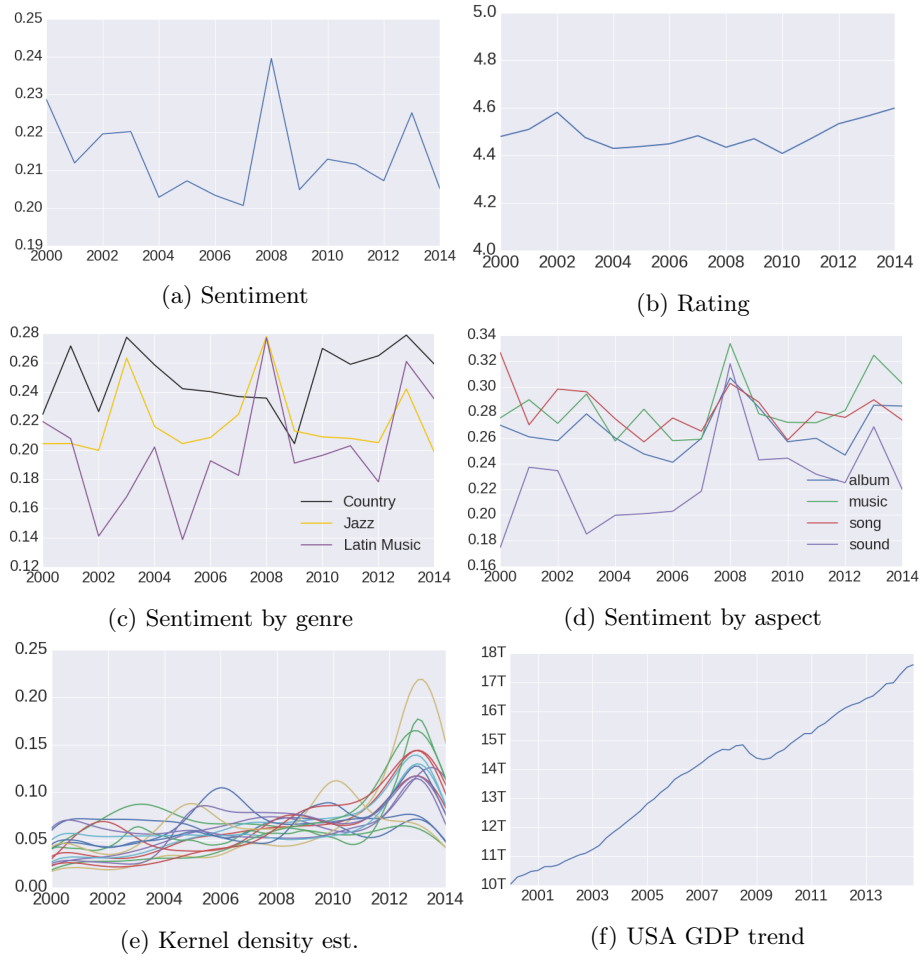


Figure 11: Sentiment (a, c, and d) and rating (b) averages by review publication year; Kernel density estimation of the distribution of reviews by year (e); GDP trend in USA from 2000 to 2014 (f)

society after 2008 (see Figure 11f). In any case, further study of the different implied variables is necessary to reinforce any of these hypotheses.

### 6.2.2 Evolution by album publication year

In this case, we study the evolution of the polarity of language by grouping reviews according to the album publication date. This date was gathered from MusicBrainz, meaning that this study is conducted on the 42,1% of the dataset that was successfully mapped. We compared again the evolution of the average sentiment polarity (Figure 12a) with the evolution of the average rating (Fig-



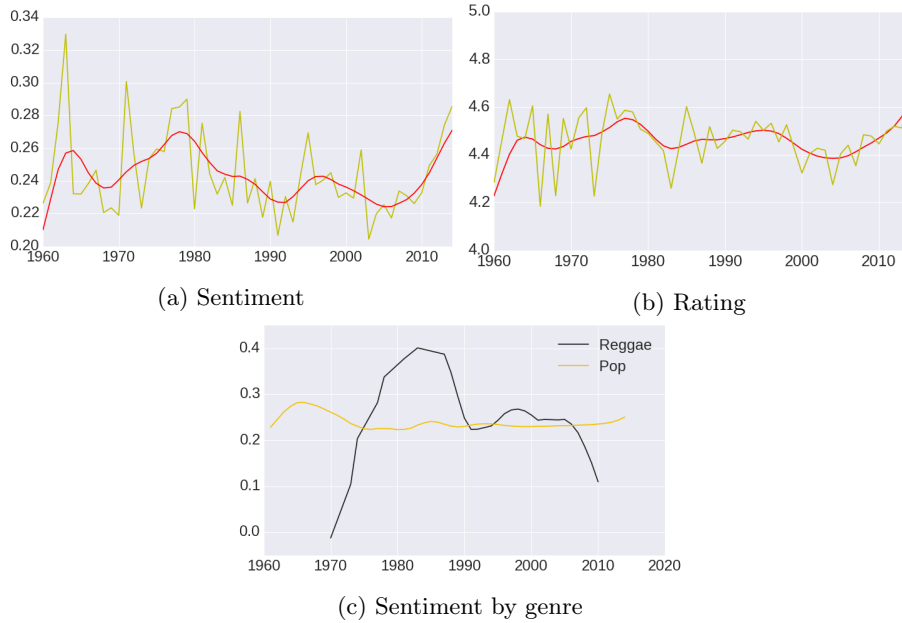


Figure 12: Sentiment (a), rating (b), and sentiment by genres (c) averages by album publication year.

ure 12b). Contrary to the results observed by review publication year, here we observe a strong correlation between ratings and sentiment polarity. To corroborate that, we computed first a smoothed version of the average graphs, by applying 1-D convolution (see line in red in Figures 12a and 12b). Then we computed Pearson’s correlation between smoothed curves, obtaining a correlation  $r = 0.75$ , and a p-value  $p \ll 0.001$ . This means that in fact there is a strong correlation between the polarity identified by the sentiment analysis framework in the review texts, and the rating scores provided by the users. This correlation reinforces the conclusions that may be drawn from the sentiment analysis data.

To further dig into the utility of this polarity measure for studying genre evolution, we also computed the smoothed curve of the average sentiment by genre, and illustrate it with two idiosyncratic genres, namely *Pop* and *Reggae* (see Figure 12c). We observe in the case of *Reggae* that there is a time period where reviews have a substantial use of a more positive language between the second half of the 70s and the first half of the 80s, an epoch which is often called the golden age of *Reggae* (Alleyne & Dunbar, 2012). This might be related to the publication of Bob Marley albums, one of the most influential artists in this genre, and the worldwide spread popularity of reggae music. In the case of *Pop*, we observe a more constant sentiment average. However, in the 60s and the beginning of 70s there are higher values, probably consequence by the release of albums by The Beatles and other iconic pop bands. These results show that the

use of sentiment analysis on music reviews over certain timelines may be useful to study genre evolution and identify influential events.

Finally, we observe a growing tendency for sentiment and rating average over albums published in the last years of the study. Are we experiencing a new golden age of music or is this due to an increase of the number of reviewers incentivized by 3rd party services (Gibbs, 2016)? Again, data-driven analysis has provided us with meaningful insights as well as novel hypotheses that open up vibrating avenues for further studies.

## 7 Conclusions

We have presented different methodologies to process large corpora of music-related documents from a musicological perspective, enabling data-driven analysis, which can be further used for complementing expert knowledge. These methodologies have been evaluated on three different use cases: Flamenco music (with the creation, population, and analysis of a flamenco knowledge base); the Renaissance period (by processing and studying a corpus of composers biographies); and the evolution of music criticism (with a cross-genre diachronic study using as target data online music reviews).

First, different challenges and techniques for gathering and combining large text corpora are presented and applied to compile the different datasets used throughout the paper. Then, a methodology for the analysis and visualization of word frequencies has been presented, showing some of the most important aspects of the different music schools in Renaissance Music. Next, an information extraction pipeline is presented where entities identified in text are connected to a knowledge base throughout an ad-hoc entity linking system. Then, a method for extracting biographical information from artist biographies is described, and applied to enrich a flamenco knowledge base and to study the Renaissance period. The analysis of extracted biographical data revealed migratory tendencies of composers and the oscillation of the gravitational center of the music activity among different European cities within the Renaissance period. A methodology to build knowledge graphs is described next and evaluated in the task of computing artist relevance rankings. Experimental results show high correlation between the automatically obtained ranking of artists and the opinion of flamenco experts. The application on the Renaissance corpus revealed also that the use of only inner connections helps the approach to obtain results that are more musicologically meaningful. Finally, an aspect-based sentiment analysis method is described. This methodology is then used to perform a diachronic study of the sentiment polarity expressed in customer reviews from two different standpoints. First, an analysis by year of review publication suggests that geopolitical events or macro-economical circumstances may influence the way people speak about music. Second, an analysis by year of album publication shows how sentiment analysis can be useful to study the evolution of music.

The main contribution of this work is a demonstration of the usefulness for musicological research of applying systematic linguistic processing techniques on

text collections *about music*. Although further work is necessary to elaborate on the hypotheses or claims that may be derived from purely data-driven analyses, these methodologies have shown their suitability in the quest of knowledge discovery from large amounts of documents, which may be highly useful for musicologists and humanities researchers in general. In fact, one of the strongest claims used in the Information Age is that *Big Data* can be used to reveal hidden patterns and meaningful variables hidden among unstructured information, and indeed in this work we provide a myriad of conclusions drawn from intelligent text processing, in the hopes that these may constitute the cornerstone of further musicological studies. Moreover, we envision that the combination of knowledge extracted from text with knowledge extracted from other data modalities (e.g., audio signals or music scores) would be a further step in the construction of real-world music understanding systems.

## 8 Acknowledgments

This work was partially funded by the Spanish Ministry of Economy and Competitiveness under the Maria de Maeztu Units of Excellence Programme (MDM-2015-0502) and by the COFLA2 research project (Proyectos de Excelencia de la Junta de Andalucía, FEDER P12-TIC-1362).

## References

- Alcorta, C. S., Sosis, R., & Finkel, D. (2008). Ritual harmony: Toward an evolutionary theory of music. *Behavioral and Brain Sciences*, 31(5), 576. doi: Doi10.1017/S0140525x08005311
- Alleyne, M. R., & Dunbar, S. (2012). *The Encyclopedia of Reggae: The Golden Age of Roots Reggae*. Sterling.
- Bellomi, F., & Bonato, R. (2005). Network analysis for wikipedia. In *Proceedings of wikimania 2005, the first international wikimedia conference*.
- Blas Vega, J., & Ríos Ruiz, M. (1988). *Diccionario enciclopédico ilustrado del flamenco*. Madrid: Cinterco.
- Brin, S., & Page, L. (1998). The Anatomy of a Large-Scale Hypertextual Web Search Engine. *Computer Networks*, 30, 107–117. Retrieved from <http://infolab.stanford.edu/~backrub/google.html>
- Bunescu, R., & Pasca, M. (2006). Using Encyclopedic Knowledge for Named Entity Disambiguation. In *Proceedings of the 11th conference of the european chapter of the association for computational linguistics (eacl-06)* (pp. 9–16). Retrieved from <http://www.cs.utexas.edu/users/ai-lab/?bunescu:eac106>

- Cowie, J., & Lehnert, W. (1996). Information extraction. *Communications of the ACM*, 39(1), 80–91.
- Crawford, T., Fields, B., Lewis, D., & Page, K. (2014). Explorations in linked data practice for early music corpora. *Digital Libraries 2014*.
- Dong, R., O'Mahony, M. P., & Smyth, B. (2014). Further Experiments in Opinionated Product Recommendation. In *Iccbr'14* (pp. 110–124). Retrieved from [http://dx.doi.org/10.1007/978-3-319-11209-1\\_{\\_}9](http://dx.doi.org/10.1007/978-3-319-11209-1_{_}9) doi: 10.1007/978-3-319-11209-1\_9
- Dong, R., Schaal, M., O'Mahony, M. P., & Smyth, B. (2013). Topic Extraction from Online Reviews for Classification and Recommendation. In (pp. 1310–1316).
- Esuli, A., & Sebastiani, F. (2006). Sentiwordnet: A publicly available lexical resource for opinion mining. In *Proceedings of Irec* (Vol. 6, pp. 417–422).
- Fast, K. V., & Sedig, K. (2011, November). Interaction and the epistemic potential of digital libraries. *International Journal on Digital Libraries*, 11(3), 169–207. Retrieved from <http://link.springer.com/10.1007/s00799-011-0066-8> doi: 10.1007/s00799-011-0066-8
- Fernández, L. (2004). *Teoría musical del flamenco*. Madrid: Acordes Concert.
- Finkel, J. R., Grenager, T., & Manning, C. (2005). Incorporating Non-local Information into Information Extraction Systems by Gibbs Sampling. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics* (pp. 363–370).
- Fujinaga, I., & Weiss, S. F. (2016). Digital prosopography for renaissance musicians: Discovery of social and professional networks. *NEH White Paper*. doi: <http://dx.doi.org/10.17613/M6FW85>
- Gamboa, J. M. (2005). *Una historia del flamenco*. Madrid: Espasa-Calpe.
- Getoor, L. (2012). Entity Resolution : Theory , Practice & Open Challenges. *Tutorial at AAAI-12*, pages 2018–2019.
- Gibbs, S. (2016, Sep). Incentivised reviews are warping amazon's product star ratings, report says. *The Guardian*. Retrieved from <https://www.theguardian.com/technology/2016/sep/20/incentivised-reviews-amazon-product-star-ratings>
- Grove, G. (1878). *A Dictionary of Music and Musicians*.
- Hu, M., & Liu, B. (2004). Mining Opinion Features in Customer Reviews. In *AAAI'04* (pp. 755–760). Retrieved from <http://dl.acm.org/citation.cfm?id=1597148.1597269>

- Juslin, P. N., & Västfjäll, D. (2008). Emotional responses to music: the need to consider underlying mechanisms. *The Behavioral and brain sciences*, 31(5), 559–621. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/18826699> doi: 10.1017/S0140525X08005293
- Kleinberg, J. M. (1999). Authoritative sources in a hyperlinked environment. *Journal of the ACM (JACM)*, 46, 604–632. doi: 10.1145/324133.324140
- Knees, P., & Schedl, M. (2011). Towards Semantic Music Information Extraction from the Web Using Rule Patterns and Supervised Learning. *Workshop on Music Recommendation and Discovery, colocated with ACM RecSys 2011*, 18. Retrieved from [http://ftp.informatik.rwth-aachen.de/Publications/CEUR-WS/Vol-793/womrad2011\\_complete.pdf#page=24](http://ftp.informatik.rwth-aachen.de/Publications/CEUR-WS/Vol-793/womrad2011_complete.pdf#page=24)
- McAuley, J., Pandey, R., & Leskovec, J. (2015). Inferring Networks of Substitutable and Complementary Products. In *Proceedings of the 21st acm sigkdd conference on knowledge discovery and data mining (kdd'15)* (p. 12). Retrieved from <http://arxiv.org/abs/1506.08839> doi: 10.1145/2783258.2783381
- Mihalcea, R., & Csomai, A. (2007). Wikify!: linking documents to encyclopedic knowledge. In *Proceedings of the sixteenth acm conference on conference on information and knowledge management* (pp. 233–242).
- Moisi, D. (2010). *The Geopolitics of Emotion: How Cultures of Fear, Humiliation, and Hope are Reshaping the World*. Anchor Books. Retrieved from <http://lccn.loc.gov/2008044798>
- Moro, A., Raganato, A., & Navigli, R. (2014). Entity Linking meets Word Sense Disambiguation: a Unified Approach. *Transactions of the Association for Computational Linguistics (TACL)*, 2, 231–244.
- Navarro, J. L., & Ropero, M. (1995). *Historia del flamenco*. Sevilla: Ed. Tartessos.
- Oramas, S. (2014). Harvesting and Structuring Social Data in Music Information Retrieval. *Extended Semantic Web Conference (ESWC). Lecture Notes in Computer Science*, 8465, 817–826.
- Oramas, S., Espinosa-Anke, L., Lawlor, A., & Others. (2016). Exploring Customer Reviews for Music Genre Classification and Evolutionary Studies. In *Proceedings of the 17th international society for music information retrieval conference (ismir 2016)*.
- Oramas, S., Espinosa-Anke, L., Sordo, M., Saggion, H., & Serra, X. (2016a). ELMD: An Automatically Generated Entity Linking Gold Standard Dataset in the Music Domain. In *In proceedings of the 10th international conference on language resources and evaluation, lrec 2016*.

- Oramas, S., Espinosa-Anke, L., Sordo, M., Saggion, H., & Serra, X. (2016b). Information extraction for knowledge base construction in the music domain. *Data and Knowledge Engineering*, 106, 70–83. Retrieved from <http://dx.doi.org/10.1016/j.datak.2016.06.001> doi: 10.1016/j.datak.2016.06.001
- Oramas, S., Gómez, F., Gómez, E., & Mora, J. (2015). FlaBase: Towards the Creation of a Flamenco Music Knowledge Base. In *Proceedings of the 16th international society for music information retrieval conference*.
- Oramas, S., Ostuni, V. C., Di Noia, T., Serra, X., & Di Sciascio, E. (2016). Sound and Music Recommendation with Knowledge Graphs. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 8(2), 21.
- Oramas, S., Sordo, M., Espinosa-Anke, L., & Serra, X. (2015). A Semantic-based Approach for Artist Similarity. In *Proceedings of 16th the international society for music information retrieval conference (ismir)*.
- Oramas, S., Sordo, M., & Serra, X. (2014). Automatic Creation of Knowledge Graphs from Digital Musical Document Libraries. In *Conference in interdisciplinary musicology*. Berlin.
- Ratcliff, J. W., & Metzener, D. (1988). Pattern matching: The Gestalt approach. *Dr. Dobb's Journal*, 13, 46–72.
- Rose, S., & Tuppen, S. (2014). Prospects for a Big Data History of Music. In *International workshop on digital libraries for musicology*.
- Sadie, S. (2001). *The New Grove Dictionary of Music and Musicians*, 2d ed. (Vol. 58). doi: 10.1353/not.2001.0195
- Sordo, M., Serrà, J., & Serra, X. (2012). A Method for Extracting Semantic Information from on-line Art Music Discussion Forums. In *2nd compmusic workshop* (pp. 55–60).
- Sutcliffe, R., Crawford, T., Fox, C., Root, D. L., & Hovy, E. (2015). Relating natural language text to musical passages. In *Proceedings of the 16th international society for music information retrieval conference*.
- Tata, S., & Di Eugenio, B. (2010). 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). Retrieved from <http://www.aclweb.org/anthology/P10-1140>
- Vrandečić, D., & Krötzsch, M. (2014). Wikidata: A Free Collaborative Knowledgebase. *Commun. ACM*, 57(10), 78–85. Retrieved from <http://doi.acm.org/10.1145/2629489> doi: 10.1145/2629489