

The Audio Commons initiative and the technologies for facilitating the reuse of open audio content

Xavier Serra

Music Technology Group

Universitat Pompeu Fabra, Barcelona

<http://audiocommons.org>

Outline

- Introduction
- AudioCommons project presentation
- The Audio Commons Ecosystem
- Relevant state of the art
- AudioCommons project tasks
- Conclusions



ABOUT THE AUDIO COMMONS INITIATIVE

The **Audio Commons Initiative** aims at *bringing Creative Commons audio content to the creative industries*. But what does this mean? We realise that significant amounts of user-generated audio content, such as **sound effects, field recordings, musical samples** and **music pieces** (among others), are uploaded to online repositories and made available under Creative Commons licenses. Furthermore, a constantly increasing amount of multimedia content, originally released with traditional copyright licenses, is becoming public domain as its copyright expires. However, we believe that the professional creative industries (e.g. videogames, film and music industries) are **not yet using much of all this content** in their media productions.

There are a number of reasons why such content is not yet extensively used in the professional sector. We think that a major one is the lack of a **shared culture** within the creative industries of open content and its potential use. But there are also **technical and practical issues** that do not facilitate this usage. Despite the amount of Creative Commons audio content available in online repositories such as [Jamendo](#) or [Freesound](#), other potentially useful content remains scattered around the web (if available at all) and typically not properly labeled with specific licenses or reachable through search engines. Also, the nature of this content, coming from a variety of sources and from authors with different levels of expertise, results in unstructured (or **not uniformly structured**) mass of resources, limiting its potential retrieval and reuse possibilities. Moreover, no tools are easily available to search and incorporate Creative Commons audio content in the **production workflows** of the creative industries.

Introduction: Motivation

- Creative Commons audio content has a huge potential for reuse which *is not being* exploited by the creative industries.
 - Limited understanding of CC licenses.
 - Content scattered.
 - Content not properly labeled, unstructured.
 - Lack of tools for seamless integration.

Project presentation: goals


- Promote publication of AC content and foster its reuse.
- Develop open technologies to support publication and reuse of AC content.
- Develop open technologies for the semantic annotation of AC content.
- Bootstrap the Audio Commons Ecosystem (ACE).
- Define standard procedures for joining the ACE.

Project presentation: consortium

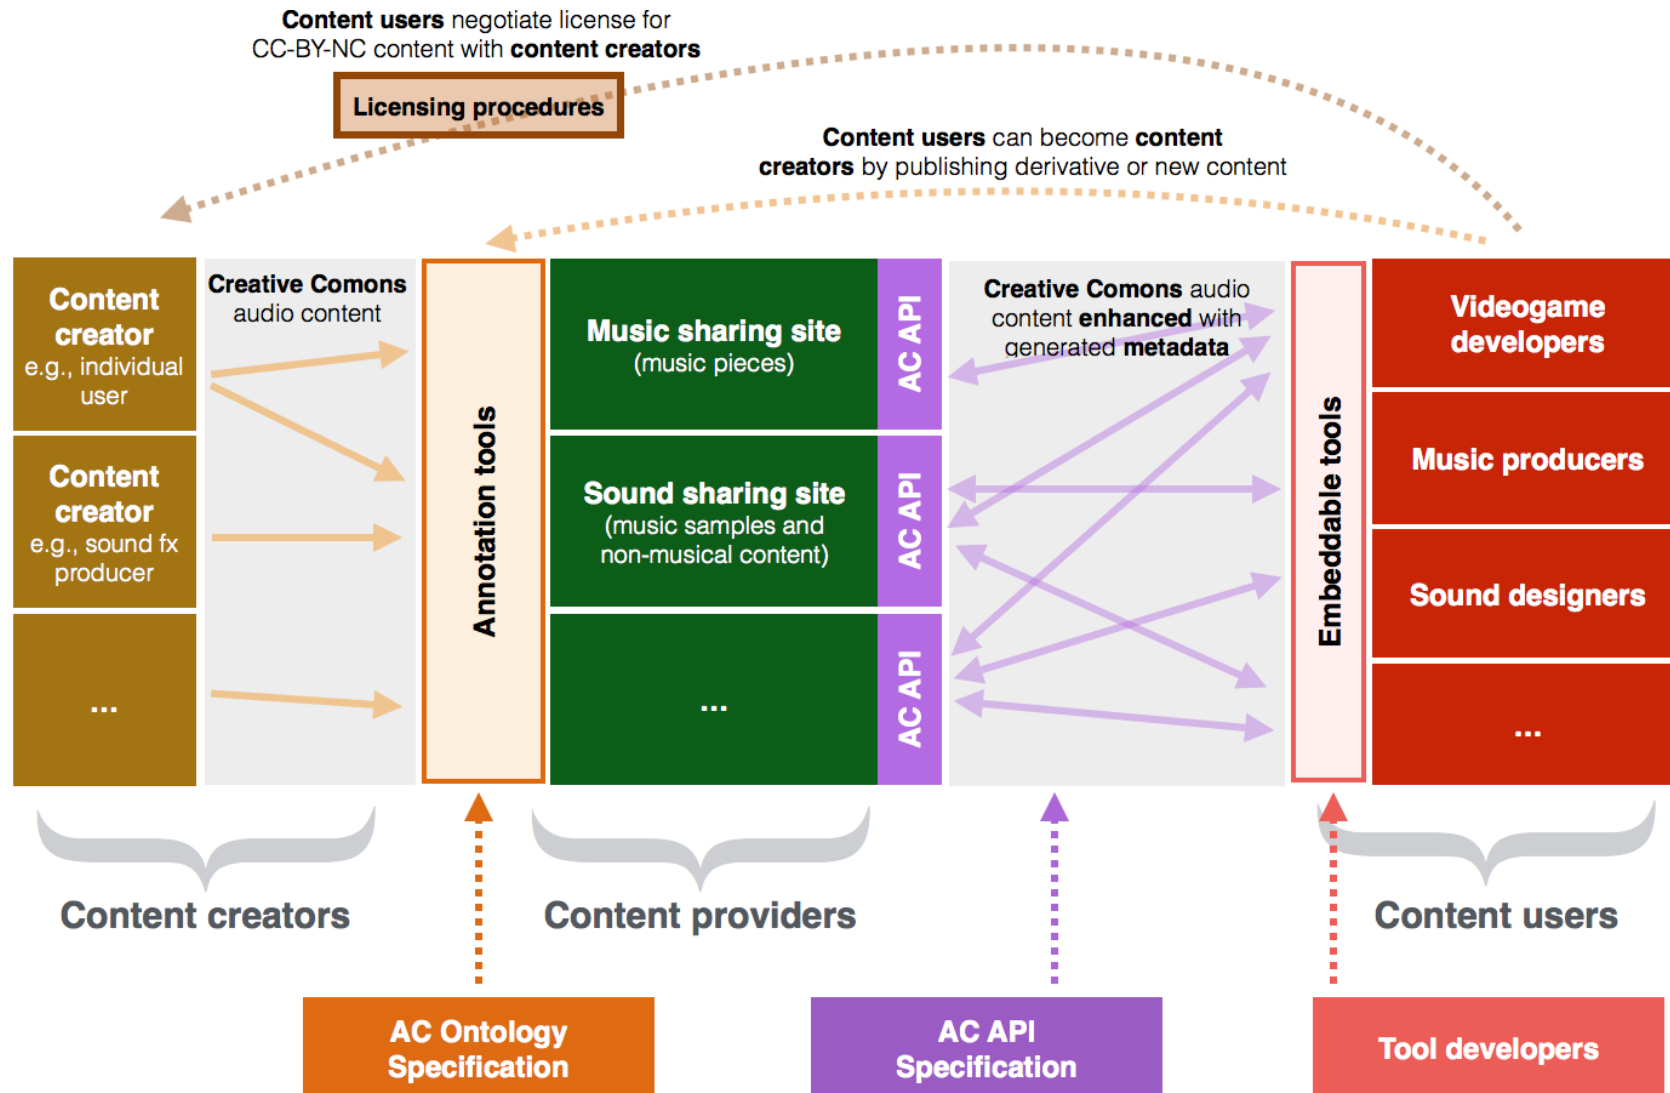
- Academic partners:



- Industrial partners:



The AudioCommons Ecosystem



AudioCommons initial content



300k sound samples



470k music pieces



Relevant state of the art

- Availability and access to CC audio content.
- Licensing procedures for CC audio resources.
- Semantic representation of sound and music.
- Semantic annotation of sound and music.
- Production tools supporting CC audio.

Availability and access to CC audio content

- Few CC audio content sites with APIs (e.g. freesound.org, jamendo.com, ...).
- Limited access due to lack of high quality and unified metadata.
- No unified access mechanism for APIs (APIs have different specifications).
- Inadequate content retrieval tools.
- CC audio content not frequently used in professional environments.

Freesound

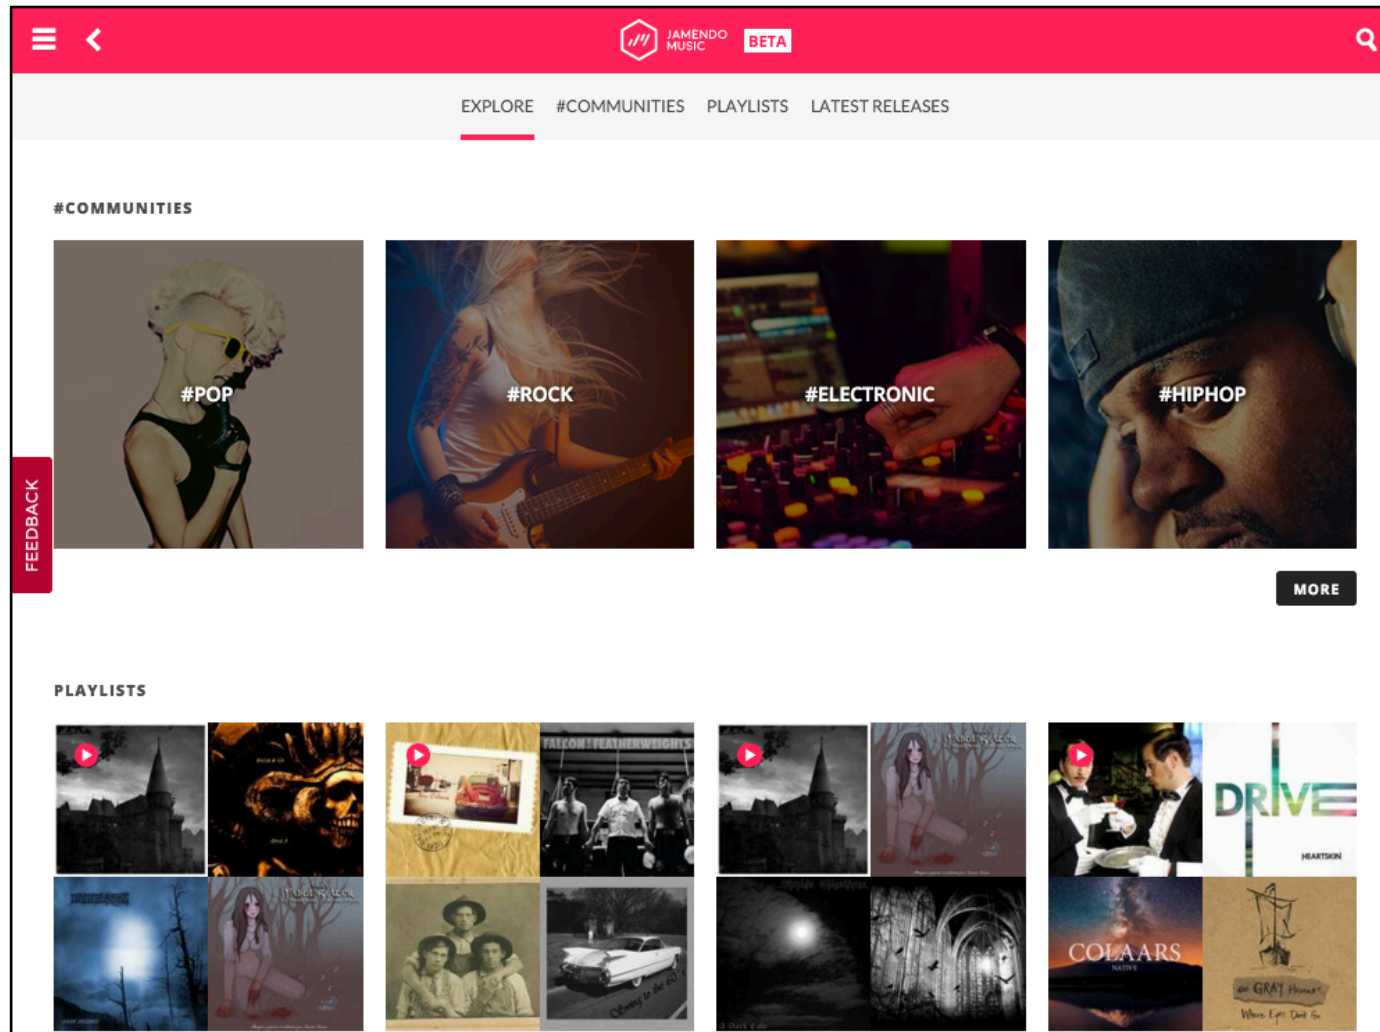
The screenshot shows the Freesound website interface. At the top, there is a navigation bar with the Freesound logo on the left and user options on the right: 'xserra Messages 0 Mod 2 Settings Log Out' and a yellow 'Upload Sounds' button. Below the navigation bar is a search bar with the placeholder text 'search sounds'. The main content area is divided into several sections:

- Random sound of the day:** A featured sound titled 'Gate made out of metal a...' by user 'zxx41'. It includes a waveform player, a description: 'A small square gate made out of metal and wire, being closed with a fairly strong force causing it to ...', and statistics: 'July 5th, 2014', '154 downloads', and '0 comments'. There are also icons for social media and a license.
- Freesound Blog:** A section with two entries:
 - Freesound Labs:** 'July 27th, 2015 frederic.font'. Text: 'Dear freesounders, we are happy to announce that we've just launched a new Freesound-side site called Freesound Labs. Freesound Labs is a directory of projects, hacks, apps, research and other initiatives that use content from Freesound or use the Freesound API. Our ... Continue reading → [Read Full Entry](#)'
 - Introducing the Modular Samples library:** 'July 7th, 2015 frederic.font'. Text: 'Hi everyone, We're glad to announce that we're in the process of uploading around 70GB of analog synthesizer samples to Freesound*, aka the Modular Samples library, consisting of 40,000 samples organised in 461 packs, recorded from 20 different analog synthesizers and ... Continue reading → [Read Full Entry](#)'
- Welcome to Freesound:** A yellow banner with the text: 'Freesound is a collaborative database of Creative Commons Licensed sounds. Browse, download and share sounds.'
- Support Us. Get your Freesound T-Shirt!:** A section featuring a black t-shirt with the Freesound logo. Below it are two buttons: 'Europe' and 'US Canada and Asia'. Text: 'Select your store:'
- Love Freesound?:** A section with a red heart icon, the text 'Only 2421€ to go!', and a 'Donate Now' button next to a progress bar.
- Active Forum Threads:** A section header for forum threads.

Freesound API

- Browse, search, and retrieve sounds and information about them.
- Find similar sounds to a given target (based on content analysis) and retrieve automatically extracted features from audio files.
- Perform advanced queries combining content analysis features with other metadata (tags, etc...).
- Upload, comment, rate and bookmark sounds.

Jamendo



Jamendo API

- 20 different read methods to access a catalog of half-a-million tracks.
- Some powerful features for music discovery like search and radios.
- OAuth2 based authentication, 5 write methods to manage user library.
- Website to monitor your app statistics.

CC audio repositories

| | Content type | Size | CC-licenses | License friendly | API | Access via API | Licensing tools |
|---------------------------|--|------|---|------------------|-----------|----------------|-----------------|
| Jamendo | Music pieces | 470k | All 6 CC variants + CC0 | Yes | Yes | Yes | Yes |
| Freesound | Music samples, sound effects, field-recordings | 290k | CC0, CC-BY, CC-BY-NC | Yes | Yes | Yes | No |
| CC-mixer | Music stems, music pieces | 38k | All 6 CC variants + CC0 | Yes | Yes (old) | No | Yes |
| Free Music Archive | Music pieces | 89k | All 6 CC variants + CC0 | Yes | Yes | No | No |
| Internet Archive | Music pieces, radios, live concerts | 2.6M | All 6 CC variants + CC0 + non-CC licenses | No | Yes | No | No |
| Eruopeana | | 574k | All 6 CC variants + CC0 + non-CC licenses | Yes | Yes | No | No |
| Looperman | Music stems, music pieces | 196k | No (looperman open license) | - | No | No | No |
| Soundcloud | Music pieces, audio books | ? | All 6 CC variants + non-CC licenses | No | Yes | No | No |
| Bandcamp | Music pieces | ? | CC-BY, CC-ND, CC-NC, CC-BY-NC-SA, CC-BY-NC-ND | No | No | No | Yes |
| Magnatune | Music pieces | ? | CC-BY-NC-SA | - | No | No | Yes |

Licensing procedures for CC resources

- Some license CC music pieces for commercial use (e.g. jamendo.com, magnatune.com, ...).
- Uncertainties about licensing for using CC content in commercial productions.
- Lack of standardized procedures for (re)licensing content with CC licenses.
- Uncertainties about business models based on CC audio content.

Creative Commons

The screenshot shows the Creative Commons website homepage. At the top is a green navigation bar with the Creative Commons logo and menu items: About, Licenses, Public Domain, Support CC, Projects, Blog, and News. Below the navigation bar is an orange banner with the text "Keep the internet creative, free and open." and a "Donate to Creative Commons" button. The main content area is divided into two columns. The left column contains sections for "History", "Founding", "Creative Commons licenses", and "Science". The right column contains a search bar, a "Get CC updates" subscription form, and a "Recent Posts" section with a list of five articles.

creative commons About ▾ Licenses ▾ Public Domain ▾ Support CC ▾ Projects ▾ Blog News ▾

Keep the internet creative, free and open. [Donate to Creative Commons](#)

History

Founding

Founded in 2001 with the generous support of the [Center for the Public Domain](#), CC is led by a [Board of Directors](#) comprised of thought leaders, education experts, technologists, legal scholars, investors, entrepreneurs and philanthropists.

Creative Commons licenses

In December 2002, Creative Commons released its first set of [copyright licenses](#) for free to the public. Creative Commons developed its licenses — [inspired in part](#) by the Free Software Foundation's GNU General Public License (GNU GPL) — alongside a Web application platform to help you license your works freely for certain uses, on certain conditions; or dedicate your works to the public domain.

In the years following the initial release, Creative Commons and its licenses have grown at an exponential rate around the world. The licenses have been further improved, and ported to over 50 [jurisdictions](#).

Science

Since 2005, Creative Commons has undertaken projects to build commons-based infrastructure for [science](#) through identifying and lowering unnecessary barriers to research, crafting policy guidelines and legal agreements, and developing technology to make research, data and materials easier to find and use.

Search the commons

Get CC updates

Recent Posts

- [Towards a Collaborative, Coordinated Strategy for OER Implementation](#)
- [Free Music Archive launches 2015 fundraising drive](#)
- [Trans-Pacific Partnership Would Harm User Rights and the Commons](#)
- [Message to our community about the Paris and Beirut attacks](#)
- [Creative Commons offers Bassel Khartabil position as Digital Cultural](#)

Some CC licenses



Public domain, no restrictions on use and redistribution, no attribution to source needed.



You can copy, modify and redistribute the sound, but you need to give an attribution to the original source.



You can copy, modify and redistribute the sound, but you need to give an attribution to the original source and you cannot use it commercially.

Jamendo Licensing

The screenshot shows the Jamendo Licensing website interface. At the top left is the Jamendo Licensing logo. To its right are navigation links for 'CATALOG' and 'IN-STORE RADIOS'. Further right is a 'SIGN IN' button and the language 'EN'. Below the navigation bar is a breadcrumb trail: 'Home / Search / Pricing / Credits / Subscription / Affiliation / Composition'. The main heading reads 'Over 200,000 royalty free tracks for any project' followed by 'Stock music library by thousands of independent artists, all rights included'. A search bar contains the placeholder text 'What are you looking for? (ex: happy, rock, nature...)' and a magnifying glass icon. Below the search bar is a pink button labeled 'go to advanced search'. The main content area features a grid of ten category tiles: 'CHRISTMAS' (with a Christmas tree image), 'ADVERTISING' (with a woman in a field), 'YOUTUBE' (with a grid of various images), 'CORPORATE' (with a meeting room), 'MOTIVATIONAL' (with hikers), 'HAPPY' (with people in the water), 'ROCK' (with a guitar), 'POP' (with a woman singing), and 'ELECTRONIC' (with a DJ set).

Freesound licenses

| License of sound of A | B wants to distribute the new sound under | Can B do this? |
|-----------------------|---|----------------|
| cc0 | cc0 | Yes |
| cc0 | by | Yes (*) |
| cc0 | by-nc | Yes (*) |
| by | cc0 | No |
| by | by | Yes (**) |
| by | by-nc | Yes (**) |
| by-nc | cc0 | No |
| by-nc | by | No |
| by-nc | by-nc | Yes (**) |

(*) If a third user C uses the sound from B, she must attribute to B.

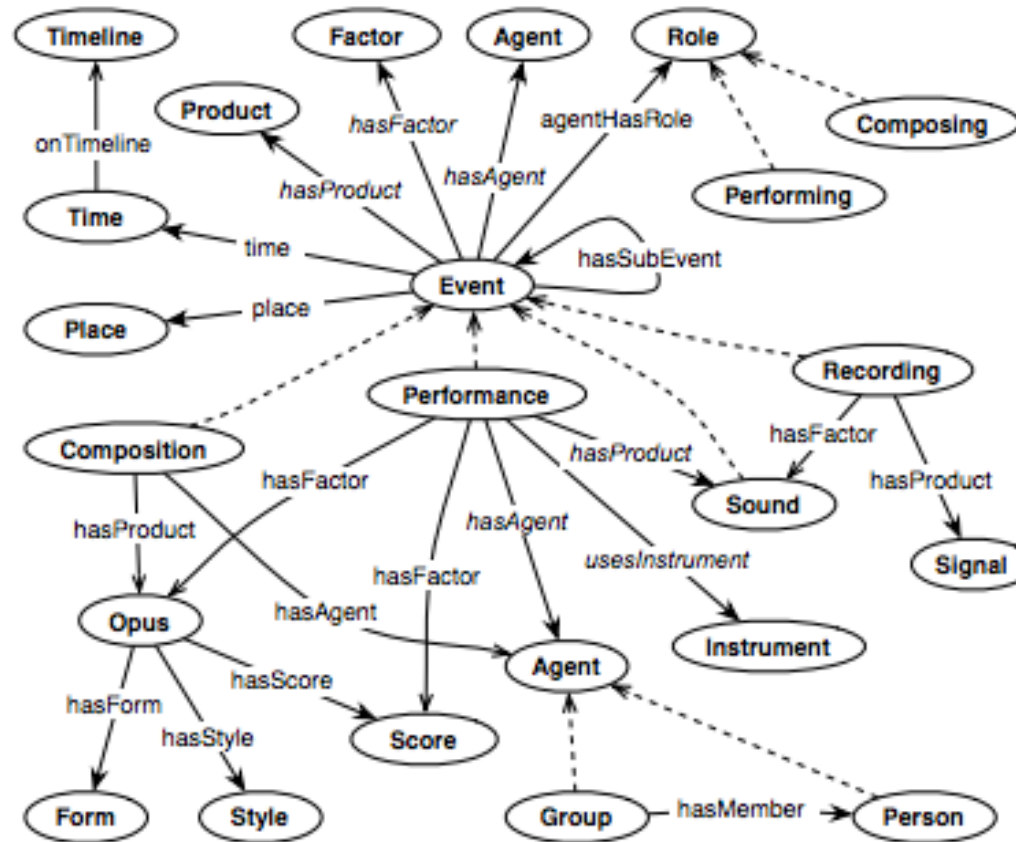
(**) B must attribute the sound to A. If a third user C uses the sound from B, she must attribute both A and B.

Semantic representation of sound and music

- Standardized generic metadata for media content (e.g. Dublin Core ontology, ...).
- A few sound and music models (e.g. Music Ontology, Europeana Data Model, MusicBrainz, ...)
- Some rights management ontologies in the media domain (e.g. MPEG-21 Media Value Chain Ontology, ...).
- Existing Semantic Web technologies (proposed by W3C) (e.g. Resource Description Framework, OWL Web Ontology Language, ...).

The Music Ontology

A vocabulary for publishing and linking music-related data on the Web.



MusicBrainz

The screenshot shows the MusicBrainz website homepage. At the top, there is a purple header with the MusicBrainz logo (a brain with musical notes) and the word "MusicBrainz" in large letters. To the right of the logo is a search bar with "search" text, a dropdown menu set to "Artist", and a "Search" button. Below the header is a navigation bar with links: "About", "Blog", "Products", "Search", "Documentation", "Contact Us", "English", "Log In", and "Create Account".

Welcome to MusicBrainz!

MusicBrainz is an open music encyclopedia that collects music metadata and makes it available to the public.

MusicBrainz aims to be:

1. **The ultimate source of music information** by allowing anyone to contribute and releasing the data under open licenses.
2. **The universal lingua franca for music** by providing a reliable and unambiguous form of music identification, enabling both people and machines to have meaningful conversations about music.

Like Wikipedia, MusicBrainz is maintained by a global community of users and we want everyone — including you — to participate and contribute.

[More Information](#) — [FAQs](#) — [Contact Us](#)

MusicBrainz is operated by the [MetaBrainz Foundation](#), a California based 501(c)(3) tax-exempt non-profit corporation dedicated to keeping MusicBrainz free and open source.

MusicBrainz Blog

Latest posts:

- [Recap of the MusicBrainz Summit 15](#)
- [Sandbox upgrade and downtime: Tuesday, 24 November, 2015 @ 17:00 UTC](#)
- [Server update, 2015-11-16](#)
- [Server update, 2015-11-02](#)
- [Server update, 2015-10-19](#)
- [Server update, 2015-10-05](#)

[Read more >>](#)

Tag Your Music

- [MusicBrainz Picard](#)
- [Magic MP3 Tagger](#)
- [Yate Music Tagger](#)

Quick Start

- [Beginners guide](#)
- [Editing introduction](#)
- [Style guidelines](#)
- [FAQs](#)

Community

Become a part of our **global community** and **start contributing!**

MusicBrainz Database

The majority of the data in the **MusicBrainz Database** is released into the **Public Domain** and can be downloaded and used **for free**.

Developers

Use our **XML web service** or **development libraries** to create your own MusicBrainz-enabled applications.

Recent Additions

THANGA MAGAN

OS IACRIVELS THE CLEVERS

COMMENT C'EST LOIN

LA VENGEANCE

SPONGE

MusicBrainz

December Season 1999
~ Release by T. M. Krishna (see all versions of this release, 1 available)

Overview | Disc IDs | Cover Art | Aliases | Tags | Details | Edit

Tracklist Display Credits at Bottom

| # | Title | Rating | Length |
|---|---|--------|--------|
| 1 | Viriboni violin: M A Sundaresan mridangam: K. Arun Prakash lead vocals: T. M. Krishna live recording of: Viriboni composer: Pachimiriyam Adiyappaiah | ★★★★★ | 8:01 |
| 2 | Rama Nee Vaddu violin: M A Sundaresan mridangam: K. Arun Prakash lead vocals: T. M. Krishna live recording of: Rama Nee Vaddu composer: Tyāgarāja (Indian composer) | ★★★★★ | 9:28 |
| 3 | Vazhi Maraittirukkude violin: M A Sundaresan mridangam: K. Arun Prakash lead vocals: T. M. Krishna live recording of: Vazhi Maraittirukkude composer: Gopalakrishna Bharathi | ★★★★★ | 7:25 |
| 4 | Ramanatham Bhajeham violin: M A Sundaresan mridangam: K. Arun Prakash lead vocals: T. M. Krishna live recording of: Ramanatham Bhajeham composer: Muthuswamy Dikshitar | ★★★★★ | 9:49 |

Release information
Format: CD
Length: 2:08:11

Additional details
Type: Album + Compilation
Status: Official
Language: [Multiple languages]
Script: Latin
Data Quality: Normal

Labels
■ Charsur Digital Workstation
CDW014

Release events
■ India
1999

Release group rating
★★★★★

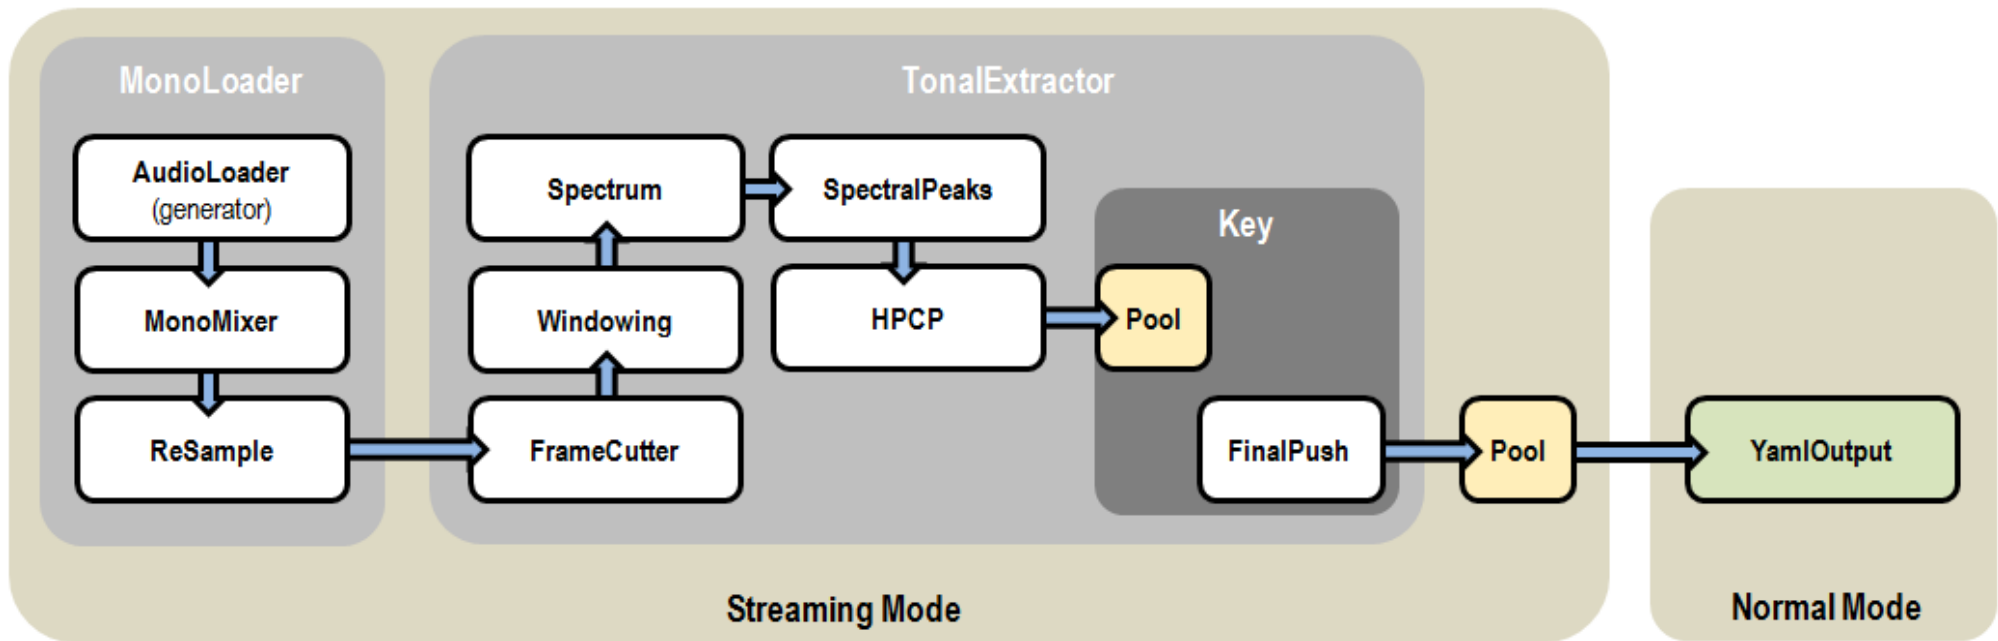
| Core Entities | |
|--------------------------------|--------------------|
| Artists: | 1,015,192 |
| Release Groups: | 1,239,983 |
| Releases: | 1,533,141 |
| Mediums: | 1,716,960 |
| Recordings: | 15,400,787 |
| Tracks: | 19,263,543 |
| Labels: | 100,508 |
| Works: | 644,963 |
| URLs: | 3,191,989 |
| Areas: | 114,635 |
| Places: | 13,547 |
| Series: | 3,805 |
| Instruments: | 756 |
| Events: | 8,421 |
| Other Entities | |
| Editors (valid / deleted): | 1,410,893 / 1,851 |
| Relationships: | 11,432,300 |
| CD Stubs (all time / current): | 331,472 / 283,353 |
| Tags (raw / aggregated): | 2,078,893 / 54,474 |
| Ratings (raw / aggregated): | 323,248 / 262,570 |
| Identifiers | |
| MBIDs: | 47,034,833 |
| ISRCs (all / unique): | 521,590 / 499,997 |
| ISWCs (all / unique): | 96,967 / 96,530 |
| Disc IDs: | 638,252 |
| Barcodes: | 514,326 |
| IPIs: | 22,088 |
| ISNIs: | 14,201 |

Semantic annotation of sound and music

- Existing audio features extraction tools (e.g. *essentia*, ...).
- Many algorithms for mid/high-level annotation of music signals: chord/key, tempo/rhythm/meter, pitch, genre/mood, structure, instrument,
- Satisfactory accuracies only achieved with algorithms trained for particular use-cases.

Essentia


Open-source C++ library for audio analysis and audio-based music information retrieval.




Essentia algorithms

- **Audio file I/O:** read and write audio files (wav, mp3, ogg, flac, etc.)
- **Standard signal processing blocks:** FFT, frame cutter, windowing, ...
- **Filters (FIR & IIR):** low/high/band pass, DC removal, equal loudness, ...
- **Statistical descriptors:** median, mean, variance, kurtosis, skewness, ...
- **Time-domain descriptors:** loudness, LARM, Leq, Vickers' loudness, zero-crossing-rate, log attack time , ...
- **Spectral descriptors:** Bark/Mel/ERB bands, MFCC, GFCC, LPC, spectral peaks, complexity, roll-off, contrast, HFC, inharmonicity, ...
- **Tonal descriptors:** Pitch salience function, predominant melody and pitch, HPCP (chroma), chords, key and scale, tuning frequency, ...
- **Rhythm descriptors:** beat detection, BPM, onset, rhythm transform, ...
- **Other high-level descriptors:** danceability, dynamic complexity, audio segmentation, SVM classifier, ...

AcousticBrainz

Contribute Downloads Goals API/Data FAQ Detailed Statistics xserra -



Welcome to AcousticBrainz!

The AcousticBrainz project aims to crowd source acoustic information for all music in the world and to make it available to the public. This acoustic information describes the acoustic characteristics of music and includes low-level spectral information and information for genres, moods, keys, scales and much more. The goal of AcousticBrainz is to provide music technology researchers and open source hackers with a massive database of information about music. We hope that this database will spur the development of new music technology research and allow music hackers to create new and interesting recommendation engines.

AcousticBrainz is a joint effort between **Music Technology Group** at **Universitat Pompeu Fabra** in Barcelona and the **MusicBrainz** project. AcousticBrainz was originally envisioned by Xavier Serra, the founder and head of the MTG. At the heart of this project lies the **Essentia toolkit** from the MTG – this open source toolkit enables the automatic analysis of music. The output from Essentia is collected by the AcousticBrainz project and made available to the public.

AcousticBrainz organizes the data on a recording basis, indexed by the **MusicBrainz ID** for **recordings**. If you know the MBID for a recording, you can easily fetch from AcousticBrainz. For details on how to do this, visit our **API documentation**.

All of the data contained in AcousticBrainz is licensed under the **CC0 license** (public domain).

Examples

If you're wondering what this collected data actually looks like, have a look at the last 5 recordings that have been submitted:

1. [Paradise Lost / Death Walks Behind You](#)
2. [Paradise Lost / Death Walks Behind You](#)
3. [Paradise Lost / Death Walks Behind You](#)
4. [Paradise Lost / Death Walks Behind You](#)
5. [Paradise Lost / Death Walks Behind You](#)


Data Statistics


| | Lossy | Lossless | Total |
|--------|-----------|-----------|------------------|
| Unique | 1,161,592 | 675,117 | 1,836,709 |
| All | 2,127,237 | 1,048,047 | 3,175,284 |

updated every 10 minutes
last updated: 2015-11-27 08:52
[more on the detailed statistics page](#)

Tweets


[Follow](#)

 **AcousticBrainz** @AcousticBrainz 30 Oct
Big database is big
[twitter.com/alsothings/sta...](#)
Expand


 **AcousticBrainz** @AcousticBrainz 30 Oct
The most frequently submitted song to AcousticBrainz is I'm Down, the B-Side to Help! by the Beatles
[acousticbrainz.org/ee898790-133f-...](#)
Weird, huh?
Expand

Tweet to @AcousticBrainz

[Blog](#) [Twitter](#) [GitHub](#) [Bug Tracker](#)



Essentia music extractor



[HOME](#) | [DOCUMENTATION](#) | [DOWNLOAD](#) | [APPLICATIONS](#) | [NEWS](#) | [ABOUT](#) |

Music extractor

Usage

streaming_extractor_music computes a large set of spectral, time-domain, rhythm, tonal and high-level descriptors. The frame-wise descriptors are **summarized** by their statistical distribution. This extractor is suited for batch computations on large music collections and is used within [AcousticBrainz project](#).

It is possible to customize the parameters of audio analysis, frame summarization, high-level classifier models, and output format, using a yaml profile file. For example, in the following profile, the extractor is set to analyze only the first 30 seconds of audio, output frame values as well as their statistical summarization, and apply two high-level models associated with the respective filepaths.

```
startTime: 0
endTime: 1e6
analysisSampleRate: 44100.0
outputFrames: 0
outputFormat: json
requireMbid: false
indent: 4

lowlevel:
  frameSize: 2048
  hopSize: 1024
  zeroPadding: 0
  windowType: blackmanharris62
  silentFrames: noise
  stats: ["mean", "var", "median", "min", "max", "dmean", "dmean2", "dvar", "dvar2"]

average_loudness:
  frameSize: 88200
  hopSize: 44100
  windowType: hann
  silentFrames: noise
```

Essentia Freesound extractor

Freesound API documentation » Analysis Documentation » previous | next | index

Analysis Descriptor Documentation

Contents

- Overview
- Terms of use
- Authentication
- Resources
- Analysis Documentation
 - Analysis Descriptor Documentation
 - Analysis Settings
 - Acronyms for the statistics
 - Metadata
 - Lowlevel Descriptors
 - Rhythm Descriptors
 - Tonal Descriptors
 - Sfx Descriptors
 - Analysis Data Example
 - Allowed descriptors for content based search
- APIv1 Documentation
- Client Libraries

Quick search

Enter search terms or a module, class or function name.

Contents

- Analysis Descriptor Documentation
 - Analysis Settings
 - Acronyms for the statistics
 - Metadata Descriptors
 - metadata.version
 - Lowlevel Descriptors
 - lowlevel.spectral_complexity
 - lowlevel.silence_rate_20dB
 - lowlevel.erb_bands
 - lowlevel.average_loudness
 - lowlevel.spectral_rms
 - lowlevel.spectral_kurtosis
 - lowlevel.barkbands_kurtosis
 - lowlevel.scvalleys
 - lowlevel.spectral_spread
 - lowlevel.pitch
 - lowlevel.dissonance
 - lowlevel.spectral_energyband_high
 - lowlevel.gfcc
 - lowlevel.spectral_flux
 - lowlevel.silence_rate_30dB
 - lowlevel.spectral_contrast
 - lowlevel.spectral_energyband_middle_high
 - lowlevel.barkbands_spread
 - lowlevel.spectral_centroid
 - lowlevel.pitch_salience
 - lowlevel.silence_rate_60dB
 - lowlevel.spectral_entropy

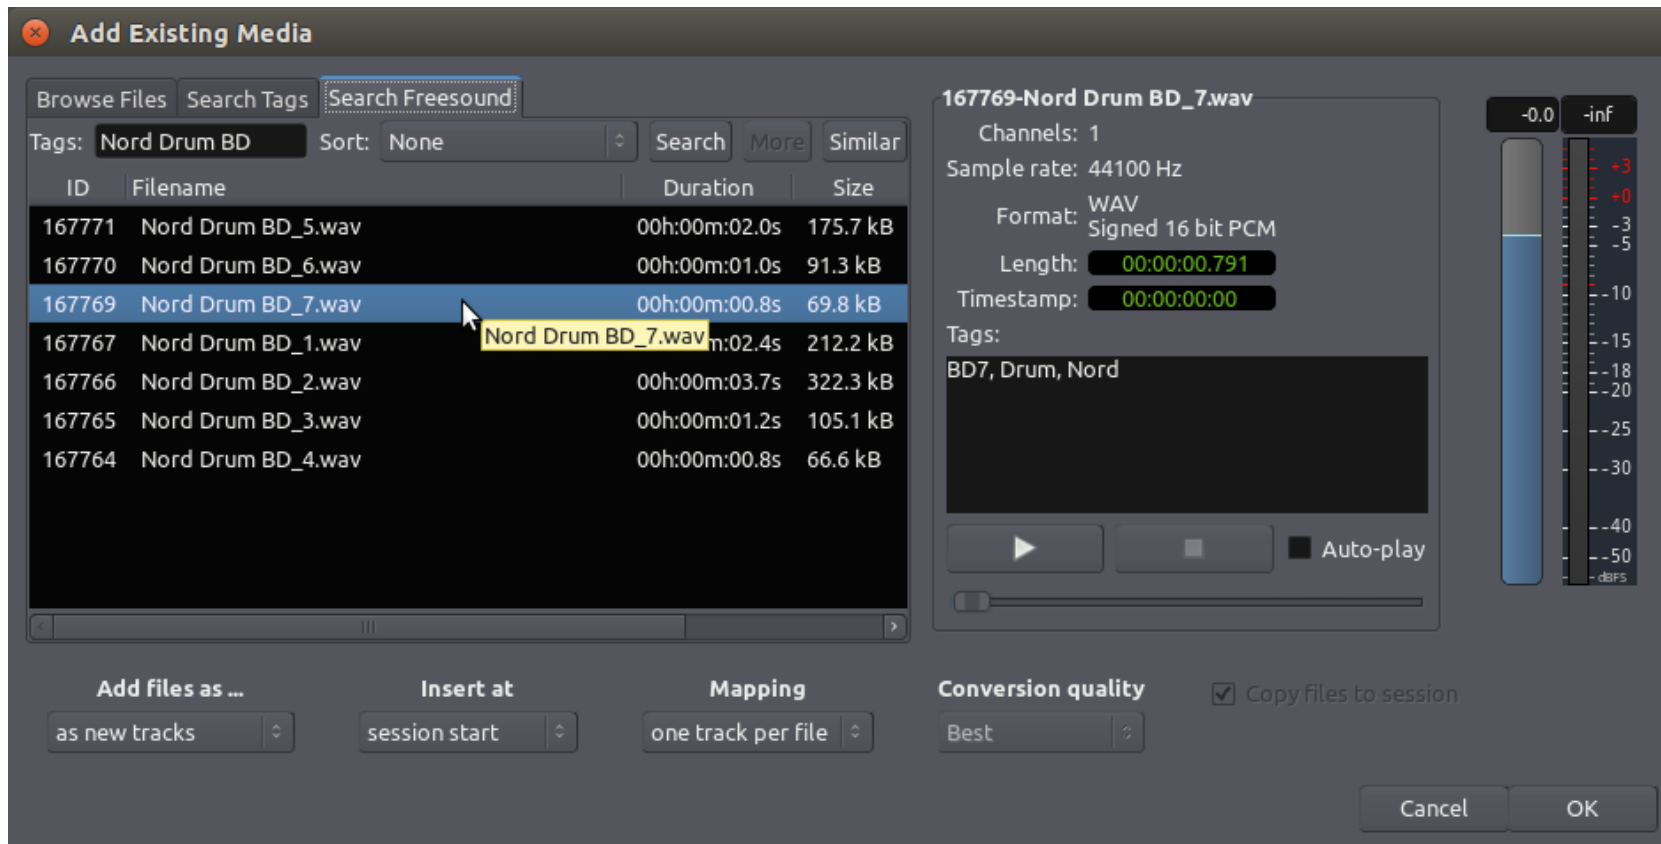
Audio annotation tools

| | Low-level features | Machine learning based features | Open license | Core language | Interface to other languages | Easy for large scale analysis | Last release |
|------------------------|--------------------|---------------------------------|--------------|---------------|------------------------------|-------------------------------|--------------|
| Essentia | Yes | Yes | Yes | C++ | Python, Matlab, Javascript | Yes | Nov 2015 |
| jMIR/jAudio | Yes | Yes | Yes | Java | | No | Oct 2013 |
| LibROSA | Yes | No | Yes | Python | | No | Oct 2015 |
| Aubio | Yes | No | Yes | C | Python | Yes | Aug 2015 |
| MIRtoolbox | Yes | Yes | Yes | Matlab | | No | Dec 2014 |
| Marsyas | Yes | Yes | Yes | C++ | | Yes | Feb 2015 |
| Sonic Annotator | Yes | Yes | Yes | C++ | Python | Yes | Nov 2015 |
| LibXtract | Yes | No | Yes | C | Python/Java | Yes | Jul 2014 |
| yaafe | Yes | No | Yes | C++ | Python/Matlab | Yes | Nov 2011 |

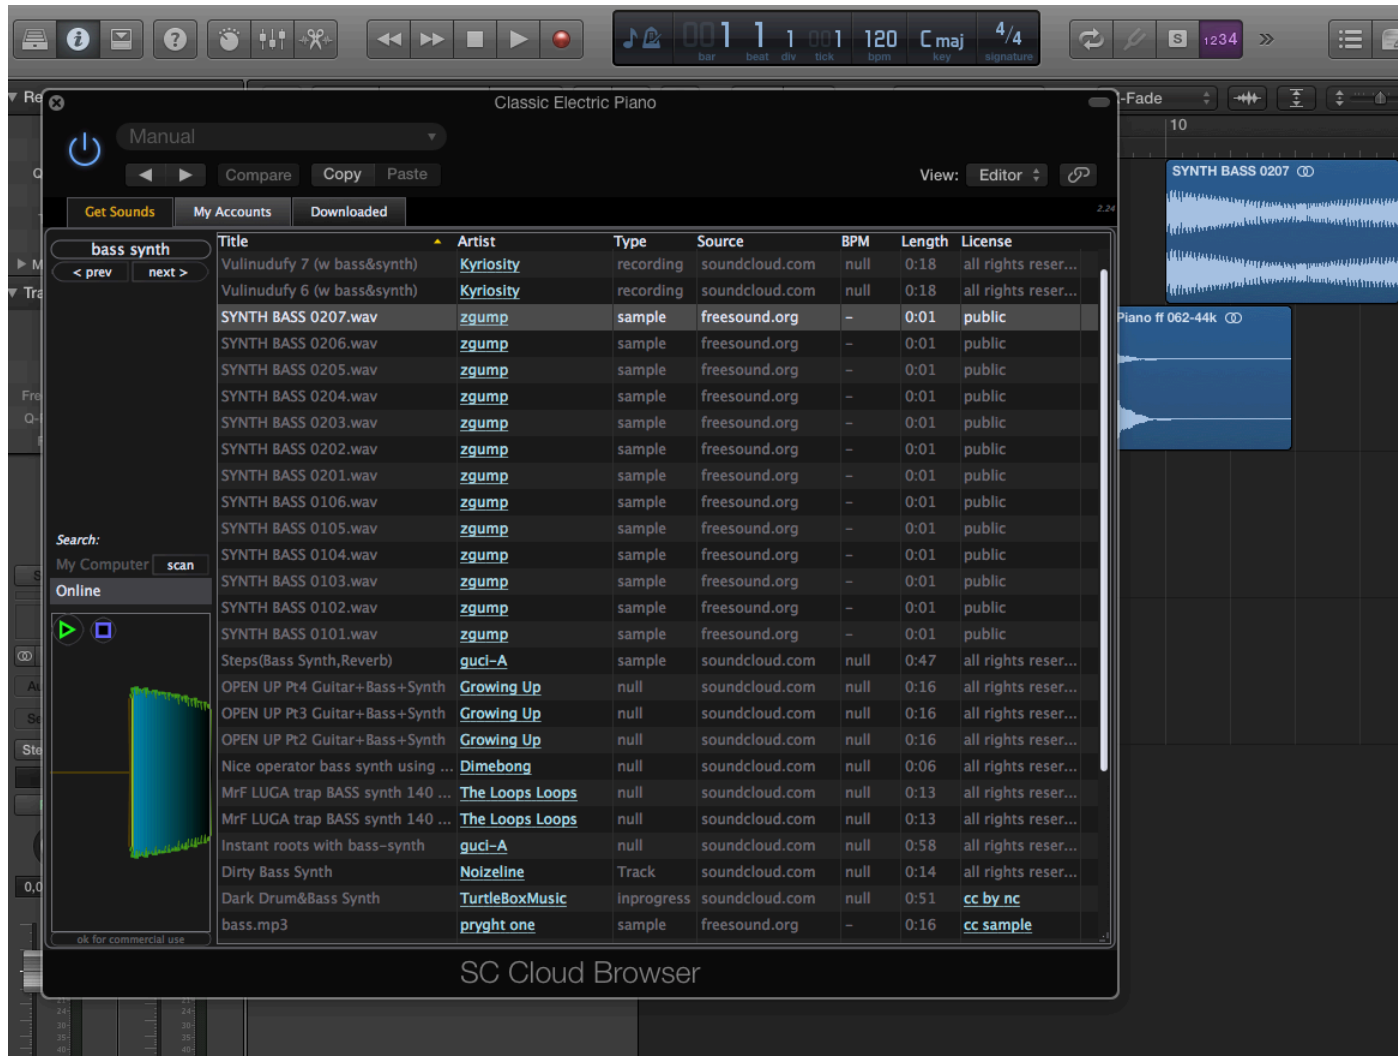
Production tools supporting CC audio

- Client libraries for some CC-ready APIs (e.g. for Freesound-API and Jamendo-API).
- Standard audio plugins architectures (e.g. VST, Audio Units, VAMP, ...).
- Open-source audio production tools easy to extent and supporting plugins (e.g. Ardour)
- No CC-ready audio production tools.

Ardour



Cloud browser (VST plugin)



Production tools

| | Retrieve | Upload | Repository |
|---------------------------|-----------------|---------------|----------------------|
| Ardour | Yes | No | Freesound |
| Cyberlink PowerDVD | Yes | No | Freesound |
| Logic Pro | No | Yes | Soundcloud |
| Cubase | No | Yes | Soundcloud |
| Pro-tools | No | Yes | Soundcloud |
| Soundly | Yes | No | Freesound |
| SoundCloud browser | Yes | Yes | Soundcloud/Freesound |

AC project tasks: production tools to be developed

- Music sampler
- Drum machine
- Loop navigator
- Music browser
- Sound FX browser
- Procedural audio SDK for videogames
- ...

AC project tasks: Ontology and API specification

- Define an ontology for the semantic annotation of musical and non-musical audio content (requirements from industry use cases)
- Design API specification for accessing content annotated using the ontology (+ implement it in Freesound)

AC project tasks: Description of musical sounds

- High-level description of music pieces (e.g. large-scale genre and mood classification, ...).
- Mid/High-level description of music samples (e.g. loop properties, instrument notes, ...).
- Manual/crowd-sourced annotation of datasets.

 AcousticBrainz

 freesound

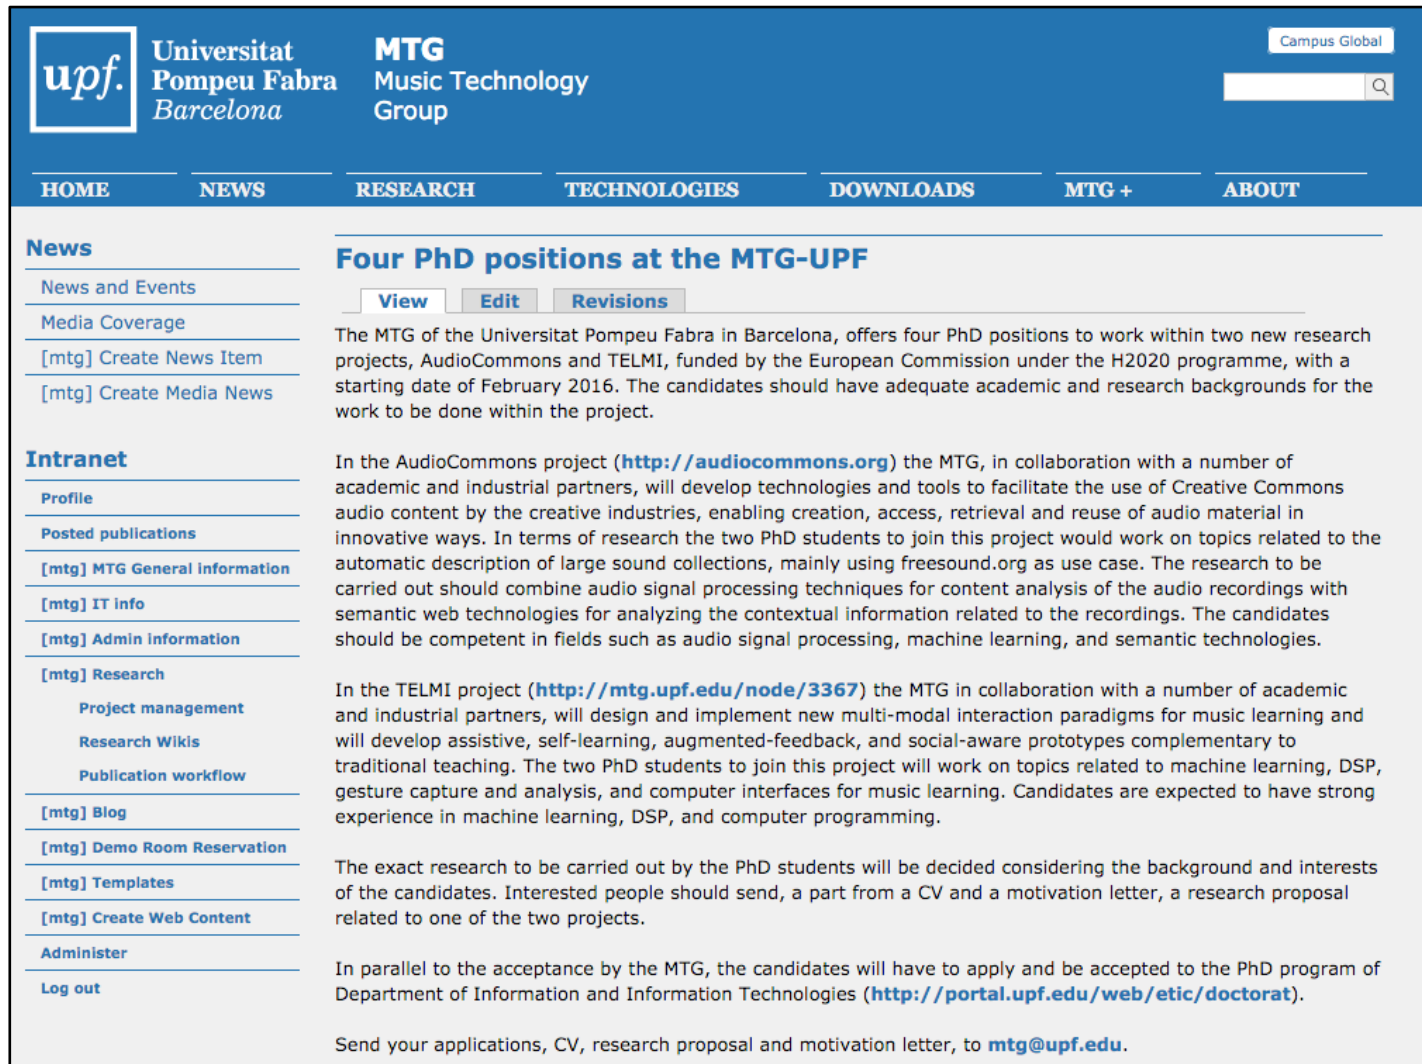
 SSENTIA

AC project tasks: Description of non-musical sounds

- High-level description of timbral properties for non-musical audio content (e.g. to query by “bright” or “dark” sounds).
- Manual/crowd-sourced annotation of datasets.



AC project: looking for people at MTG



The screenshot shows the website for the Music Technology Group (MTG) at the Universitat Pompeu Fabra (UPF) in Barcelona. The page features a blue header with the university and group logos, a search bar, and a navigation menu. The main content area displays a news article titled "Four PhD positions at the MTG-UPF". The article is structured with a title, a "View" button, and a "Revisions" button. The text describes two research projects: AudioCommons and TELMI, both funded by the European Commission. It details the requirements for candidates, including academic and research backgrounds, and provides information on how to apply, including the need for a CV, motivation letter, and research proposal. The application should be sent to mtg@upf.edu.

upf. Universitat Pompeu Fabra Barcelona

MTG Music Technology Group

Campus Global

HOME NEWS RESEARCH TECHNOLOGIES DOWNLOADS MTG + ABOUT

News

- News and Events
- Media Coverage
- [mtg] Create News Item
- [mtg] Create Media News

Intranet

- Profile
- Posted publications
- [mtg] MTG General information
- [mtg] IT info
- [mtg] Admin information
- [mtg] Research
 - Project management
 - Research Wikis
 - Publication workflow
- [mtg] Blog
- [mtg] Demo Room Reservation
- [mtg] Templates
- [mtg] Create Web Content
- Administer
- Log out

Four PhD positions at the MTG-UPF

[View](#) [Edit](#) [Revisions](#)

The MTG of the Universitat Pompeu Fabra in Barcelona, offers four PhD positions to work within two new research projects, AudioCommons and TELMI, funded by the European Commission under the H2020 programme, with a starting date of February 2016. The candidates should have adequate academic and research backgrounds for the work to be done within the project.

In the AudioCommons project (<http://audiocommons.org>) the MTG, in collaboration with a number of academic and industrial partners, will develop technologies and tools to facilitate the use of Creative Commons audio content by the creative industries, enabling creation, access, retrieval and reuse of audio material in innovative ways. In terms of research the two PhD students to join this project would work on topics related to the automatic description of large sound collections, mainly using freesound.org as use case. The research to be carried out should combine audio signal processing techniques for content analysis of the audio recordings with semantic web technologies for analyzing the contextual information related to the recordings. The candidates should be competent in fields such as audio signal processing, machine learning, and semantic technologies.

In the TELMI project (<http://mtg.upf.edu/node/3367>) the MTG in collaboration with a number of academic and industrial partners, will design and implement new multi-modal interaction paradigms for music learning and will develop assistive, self-learning, augmented-feedback, and social-aware prototypes complementary to traditional teaching. The two PhD students to join this project will work on topics related to machine learning, DSP, gesture capture and analysis, and computer interfaces for music learning. Candidates are expected to have strong experience in machine learning, DSP, and computer programming.

The exact research to be carried out by the PhD students will be decided considering the background and interests of the candidates. Interested people should send, a part from a CV and a motivation letter, a research proposal related to one of the two projects.

In parallel to the acceptance by the MTG, the candidates will have to apply and be accepted to the PhD program of Department of Information and Information Technologies (<http://portal.upf.edu/web/etic/doctorat>).

Send your applications, CV, research proposal and motivation letter, to mtg@upf.edu.

Conclusions: The AC initiative

- AudioCommons is a 3-year project (2016-2019).
- Think beyond the lifetime of the project.
- Engage new stakeholders, content providers, tool developers, ...

The Audio Commons initiative and the technologies for facilitating the reuse of open audio content

Xavier Serra

Music Technology Group

Universitat Pompeu Fabra, Barcelona

<http://audiocommons.org>
