

DATA JAM DAYS

Lausanne

November 24, 2017

FMA: A DATASET FOR MUSIC ANALYSIS

Michaël DEFFERRARD

École Polytechnique Fédérale de Lausanne (EPFL)

Joint work with Kirell BENZI (EPFL)
Pierre VANDERGHEYNST (EPFL)
Xavier BRESSON (NTU)



Fei-Fei Li

@drfeifei

Following



Datasets play crucial roles in advancing AI. **#ImageNet** helped to enable the latest deep learning advances.

	Datasets (first Available)
sch	Spoken Wall Street Journal articles and other texts (1992)
Kaplan	700,000 Grandmaster chess games, aka "The Extended Book" (1991)
in-English	1.0 trillion tokens from Google Web and News pages (collected in 2001)
frequency	8.6 million documents from Wikipedia, Wiktionary, Wikisource, and Project Gutenberg (distributed in 2010)
ification	ImageNet corpus of 1.5 million labeled images and 1,000 object categories (2010)
ation	Atari Learning Environment dataset of over 50 Atari games (2013)
ough	3 years

Datasets Over Algorithms

Content without method leads to fantasy; method without content to empty sophistry. — Johann Wolfgang von Goethe ("Maxims and Reflections", 1892) "Perhaps the most important news of ...

spacemachine.net

5:12 AM - 22 Apr 2016

103 Retweets 174 Likes



2



103



174



Motivations

Goal: open dataset for Music Information Retrieval (MIR)
mostly for Machine Learning (ML)

- ▶ Promote open evaluation and open data.
- ▶ Need both open benchmarking and standardized challenges.

There is a lot to learn from so much data!

The Free Music Archive

<https://freemusicarchive.org>



Free Music Archive



[Curators](#) | [Genres](#) | [Charts](#) | [About the FMA](#) | [Donate](#)

[Sign Up/Log In](#)

Rock

Genres > Rock



	Artist	Track	Album	Genre	Date Added
	half cocked	Magazines	Refractory	Rock, Punk, Lo-Fi	+ ↓
	Small Tall Order	Cameo Appearance	Perfect Situation	Pop, Folk, Indie-Rock	+ ↓
	Small Tall Order	Always Been Gone	Perfect Situation	Pop, Folk, Indie-Rock	+ ↓

Data: key numbers

- ▶ 106,574 tracks from 16,341 artists and 14,854 albums.
- ▶ 917 GiB and 343 days of Creative Commons-licensed audio.
- ▶ Arranged in a hierarchical taxonomy of 161 genres.
- ▶ 518 pre-computed features per track.


Data: metadata

100% track_id	100% title	93% number
2% information	14% language_code	100% license
4% composer	1% publisher	1% lyricist
98% genres	98% genres_all	98% genres_top
100% duration	100% bit_rate	100% interest
100% #listens	2% #comments	61% #favorites
100% date_created	6% date_recorded	22% tags
100% album_id	100% title	
94% type	96% #tracks	
76% information	16% engineer	18% producer
97% #listens	12% #comments	38% #favorites
97% date_created	64% date_released	18% tags
100% artist_id	100% name	25% members
38% bio	5% associated_labels	
43% website	2% wikipedia_page	
	5% related_projects	
37% location	23% longitude	23% latitude
11% #comments	48% #favorites	10% tags
99% date_created	8% active_year_begin	
	2% active_year_end	

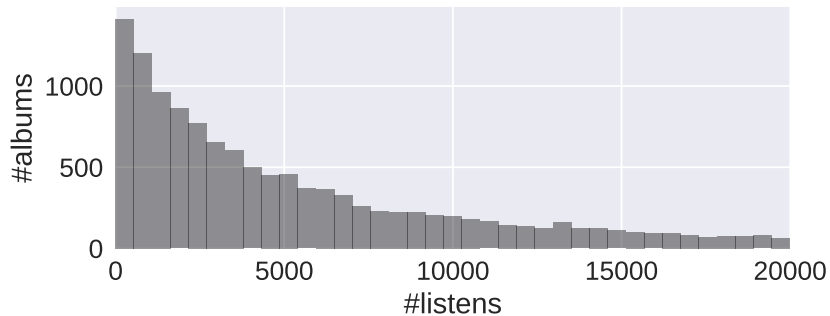
Data: metadata excerpt

track						album			
track_id	title	genres_all	genre_top	dur.	listens	title	listens	tags	
150073	Welcome to Asia	[2, 79]	International	81	683	Reprise	4091	[world music, dubtronica, fusion]	
140943	Sleepless Nights	[322, 5]	Classical	246	1777	Creative Commons Vol. 7	28900	[classical, alternate, soundtrack, piano, ...]	
64604	i dont want to die alone	[32, 38, 456]	Experimental	138	830	Summer Gut String	7408	[improvised, minimalist, noise, ...]	
23500	A Life In A Day	[236, 286, 15]	Electronic	264	1149	A Life in a Day	6691	[idm, candlestick, romanian, candle, ...]	
131150	Yeti-Bo-Betty	[25, 12, 85]	Rock	124	183	No Life After Crypts	3594	[richmond, fredericksburg, trash rock, ...]	

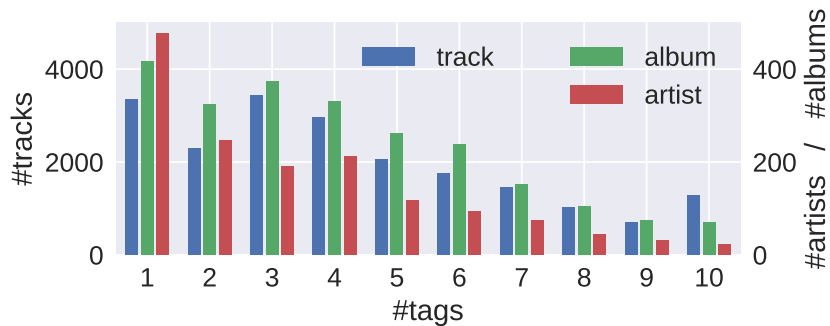
Data: genre distribution



Data: listening counts



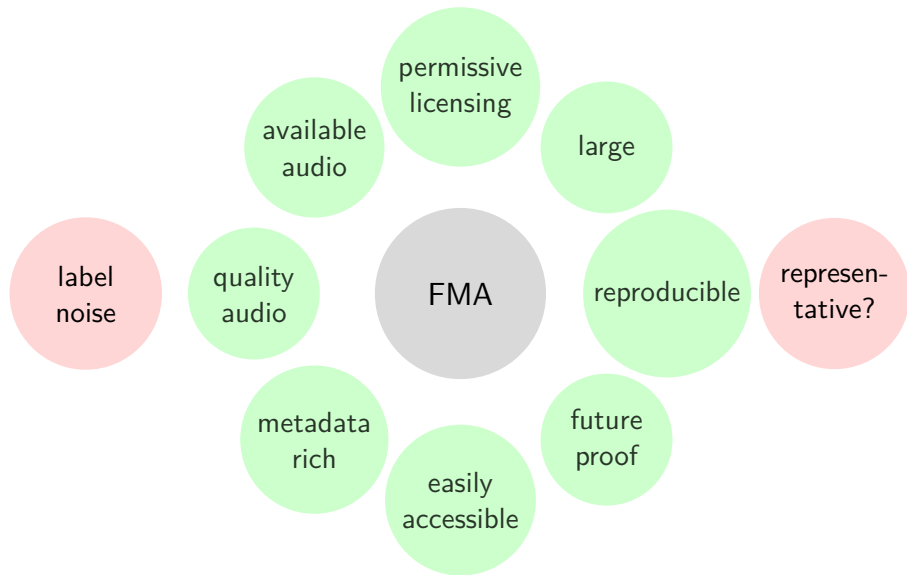
Data: tags



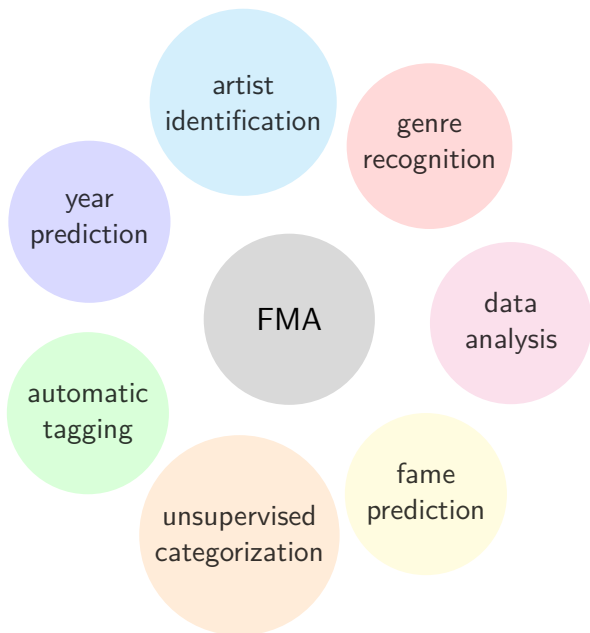
Data: subsets

dataset	clips	genres	length	size	
				[s]	[GiB] #days
small	8,000	8	30	7.4	2.8
medium	25,000	16	30	23	8.7
large	106,574	161	30	98	37
full	106,574	161	278	917	343

Qualities & Limitations



Usage



Usage example: genre classification on medium subset

feature set	dim.	LR	kNN	SVM	MLP
1 Chroma	84	44	44	48	49
2 Tonnetz	42	40	37	42	41
3 MFCC	140	58	55	61	53
4 Spec. centroid	7	42	45	46	48
5 Spec. bandwidth	7	41	45	44	45
6 Spec. contrast	49	51	50	54	53
7 Spec. rolloff	7	42	46	48	48
8 RMS energy	7	37	39	39	39
9 Zero-crossing rate	7	42	45	45	46
3 + 6	189	60	55	63	54
3 + 6 + 4	273	60	55	63	53
1 to 9	518	61	52	63	58

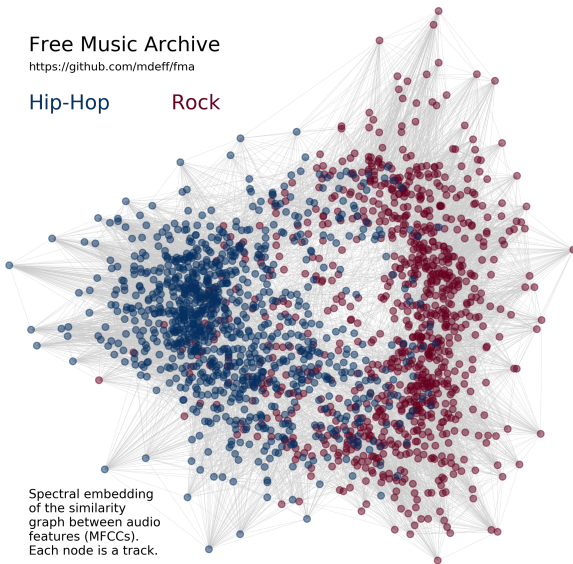
Usage example: clustering & visualization

Free Music Archive

<https://github.com/mdeff/fma>

Hip-Hop

Rock



Spectral embedding
of the similarity
graph between audio
features (MFCCs).
Each node is a track.

Online

<https://github.com/mdeff/fma>

Paper

- ▶ Dataset presentation & content analysis
- ▶ Proposed usage

Data

- ▶ 4 subsets from 10GB to 1TB (mp3 audio)
- ▶ Metadata (csv tables)

Code (Python)

- ▶ A guide to import and use the data.
- ▶ Data exploration. All the figures here and more!
- ▶ Baselines for some tasks.
- ▶ Collection and cleaning code. Reproducibility!

Slides <https://doi.org/10.5281/zenodo.1066119>

Paper Defferrard, Benzi, Vandergheynst, Bresson,
FMA: A Dataset For Music Analysis, ISMIR, 2017.
<https://arxiv.org/abs/1612.01840>

Code & Data <https://github.com/mdeff/fma>

Thanks Questions?