

Hybrid music recommender using content-based and social information

Paulo Esteban Chiliguano Torres

School of Electrical Engineering and Computer Science
Queen Mary University of London

September 1st, 2015

Outline

- 1 Motivation
 - Related work
- 2 Hybrid music recommendation
 - Design
 - Architecture
 - Item and user representation
- 3 Results
 - Music genre classifier
 - Hybrid recommender
- 4 Conclusions and future work

"Music doesn't have any special meaning; it depends what it's attached to." (Oliver Sacks 1933-2015)

Aim and Motivations

Design and implement a hybrid music recommender to mitigate the cold-start problem in a content-based recommendation strategy.

Design and implement a hybrid music recommender to mitigate the cold-start problem in a content-based recommendation strategy.

- Implement a convolutional deep neural network (CDNN) to obtain high-level representation of an audio file.

Design and implement a hybrid music recommender to mitigate the cold-start problem in a content-based recommendation strategy.

- Implement a convolutional deep neural network (CDNN) to obtain high-level representation of an audio file.
- Investigate Estimation of Distribution Algorithms (EDAs) to model user profiles in terms of probabilities of music genres preferences.

Hybrid music recommender (Yoshii et al. 2008)

- “bag of timbres” to represent acoustic features.
- Three-way aspect model: “unobserved” genre

Recommender Systems

Hybrid music recommender (Yoshii et al. 2008)

- “bag of timbres” to represent acoustic features.
- Three-way aspect model: “unobserved” genre

Deep content-based music recommendation (Oord et al. 2013)

- CDNN for latent vector representation
- Million Song Dataset

Recommender Systems

Hybrid music recommender (Yoshii et al. 2008)

- “bag of timbres” to represent acoustic features.
- Three-way aspect model: “unobserved” genre

Deep content-based music recommendation (Oord et al. 2013)

- CDNN for latent vector representation
- Million Song Dataset

Hybrid recommender based on EDA (Liang, T. et al. 2014)

- TF-IDF for item attributes
- Movielens dataset
- Permutation EDA

Hybrid music recommender design

Fundamental tasks:

- User modelling
- Information filtering

Required data:

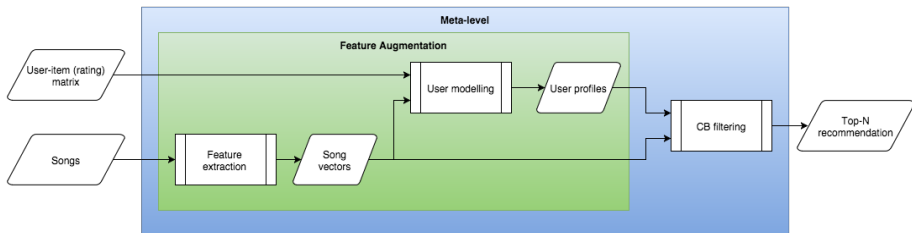
- User-item matrix: Taste profile dataset (53 users)
- Audio clips: 7digital UK catalogue (640 clips)

Song representation:

- 10-dimensional vector
- Probability to belong to a music genre

Hybrid music recommender approach

- Feature augmentation
- Meta-level



Probability of music genre

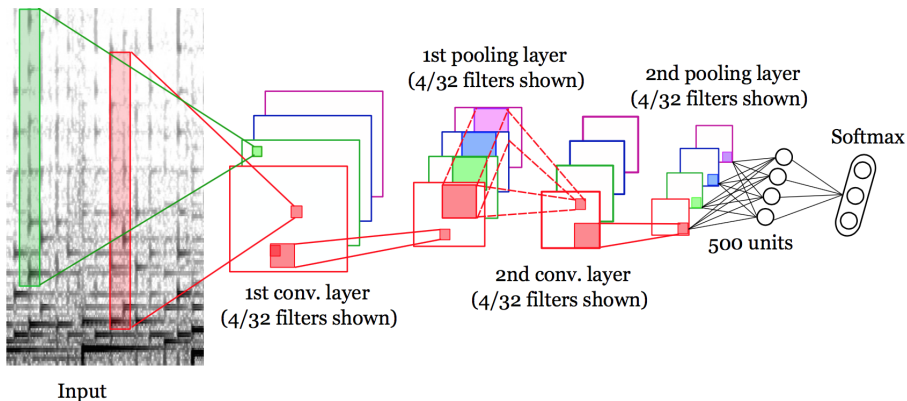


Figure: CDNN for music genre classification (Kereliuk et al. 2015)

Estimation of Distribution Algorithms (EDAs)

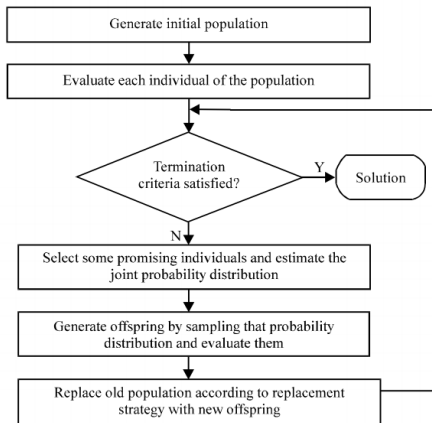


Figure: Flowchart for EDA (Ding et al. 2015)

With permutation EDA:

- 10 tags (GTZAN) equivalent to keywords
- 50 weights: evenly spaced over the interval $[0.1, \dots, 0.9]$

With continuous EDA:

- Each genre considered as a dimension
- Compute mean and covariance for each dimension along individuals
- Sample from normal distribution

Genre classification

Table: Genre classification results

Trial	Validation error (%)	Test error (%)	Iter.	Time elapsed (min.)
1	58.0	65.2	650	7.00
2	37.6	46.0	2150	13.07
3	39.6	46.0	700	7.54
4	35.6	36.8	550	6.01
5	36.4	40.0	250	5.47
6	40.4	44.8	150	5.41
7	32.4	40.4	800	8.64
8	36.0	38.8	250	5.42
9	34.0	38.8	850	9.14

Top - N recommendation

Table 5.2: Evaluation of recommender systems (N=5)

Recommender	Precision	Recall	F1	Accuracy
Content-based (baseline)	0.275 \pm 0.087	0.010 \pm 0.003	0.020 \pm 0.007	0.681 \pm 0.008
Hybrid (permutation EDA)	0.391 \pm 0.182	0.013 \pm 0.007	0.025 \pm 0.013	0.685 \pm 0.009
Hybrid (continuous UMDA)	0.318 \pm 0.142	0.011 \pm 0.005	0.021 \pm 0.011	0.683 \pm 0.009

Table 5.3: Evaluation of recommender systems (N=10)

Recommender	Precision	Recall	F1	Accuracy
Content-based (baseline)	0.301 \pm 0.059	0.022 \pm 0.007	0.041 \pm 0.012	0.678 \pm 0.007
Hybrid (permutation EDA)	0.370 \pm 0.073	0.024 \pm 0.007	0.045 \pm 0.013	0.682 \pm 0.009
Hybrid (continuous UMDA)	0.309 \pm 0.100	0.019 \pm 0.007	0.036 \pm 0.013	0.679 \pm 0.009

Table 5.4: Evaluation of recommender systems (N=20)

Recommender	Precision	Recall	F1	Accuracy
Content-based (baseline)	0.281 \pm 0.052	0.041 \pm 0.006	0.071 \pm 0.010	0.666 \pm 0.006
Hybrid (permutation EDA)	0.363 \pm 0.041	0.047 \pm 0.008	0.084 \pm 0.014	0.676 \pm 0.007
Hybrid (continuous UMDA)	0.302 \pm 0.067	0.039 \pm 0.011	0.070 \pm 0.019	0.671 \pm 0.010

Conclusions and future work

- CDNN produce similar results to long-established music genre classifiers
- Hybrid permutation EDA outperforms CB
- Investigate unsupervised deep learning
- Online evaluation

Questions?