

Music Retrieval and Recommendation

Part IV: Listener-centric and Collaborative Music Similarity and Recommendation

Markus Schedl
Peter Knees

`{markus.schedl, peter.knees}@jku.at`

Department of Computational Perception
Johannes Kepler University (JKU)
Linz, Austria

Computational Factors Influencing Music Perception and Similarity

(Schedl et al., JIIS 2013)

Examples:

- mood
- activities
- social context
- spatio-temporal context
- physiological aspects



**user
context**

Examples:

- music preferences
- musical training
- musical experience
- demographics

user properties



**music
content**

Examples:

- rhythm
- timbre
- melody
- harmony
- loudness

Examples:

- semantic labels
- song lyrics
- album cover artwork
- artist's background
- music video clips

**music
context**



**music
perception
and similarity**



Department of
Computational
Perception

Computational Factors Influencing Music Perception and Similarity

(Schedl et al., JIIS 2013)

Examples:

- mood
- activities
- social context
- spatio-temporal context
- physiological aspects



**user
context**

**personalized/context-
aware methods:**
typically extend music
content or music context
with a user-category

Examples:

- music preferences
- musical training
- musical experience
- demographics

user properties




Examples:

- rhythm
- timbre
- melody
- harmony
- loudness



Examples:

- semantic labels
- song lyrics
- album cover artwork
- artist's background
- music video clips



Department of
Computational
Perception

Computational Factors Influencing Music Perception and Similarity

(Schedl et al., JIIS 2013)

Examples:

- mood
- activities
- social context
- spatio-temporal context
- physiological aspects



**user
context**

hybrid methods:
combine factors of at
least two categories

Examples:

- rhythm
- timbre
- melody
- harmony
- loudness



Examples:

- semantic labels
- song lyrics
- album cover artwork
- artist's background
- music video clips

**music
context**



Examples:

- music preferences
- musical training
- musical experience
- demographics

**user
properties**



Department of
Computational
Perception





Overview

- Categorization of listener-centric systems
- Data sources and features for user's interaction traces
- Co-occurrence approaches in collaborative data: microblogs (#nowplaying), playlists, P2P networks
- Other applications for #nowplaying data
- Context-aware music playlist adaptation
- Music recommendation for places of interest
- Music recommendation tailored to user characteristics

Categorization of Listener-centric Systems

- Personalized systems/methods
 - incorporate aspects of the *user properties*, i.e. static attributes
 - take into account music genre preference, music experience, age, etc.
- Context-aware systems/methods
 - incorporate aspects of the *user context*, i.e. dynamic aspects
 - **active user-awareness**: new user context is automatically incorporated into the system, adaptively changing its behavior
 - **passive user-awareness**: application presents the new context to the user for later retrieval/incorporation

Data sources for listener-related data



last.fm

Level of feature extraction for listener data



High-level

Examples:
activity, behavior, mood, social context



Mid-level

Examples: user feedback, user profiles, listening sessions,
devices in reach



Low-level

Examples: listening events, ratings, time, location, sensor data,
temperature, weather conditions, demographics

Categorization of User Features

- Implicit
 - sensors: GPS, heart rate, accelerometer, pressure, light intensity, environmental noise level (now available in abundance through smart phones)
 - derived features: location + time → weather
 - learned features (via ML): accelerometer, speed → user activity
- Explicit
 - via user involvement/feedback
 - e.g., mood, activity, listening events, item ratings, skipping behavior [Pampalk et al.; 2005]

Shared Folders in a P2P Network



Make use of meta-data transmitted as file names or ID3 tags in P2P network OpenNap (Whitman, Lawrence; 2002) (Ellis et al.; 2002)

Information gathered from users' shared folders (no file downloads!)

Similarities via artist co-occurrences in collections (cond. prob.)

Sparse co-occurrence matrix

Experiments on Gnutella network (Shavitt, Weinsberg; 2009):

- meta-data highly inconsistent
- can be used as song-based similarity measure and to estimate localized popularity/trends (matching IP addresses difficult!)

Playlist Co-Occurrences

Analysis of co-occurrences of artists and songs on radio station playlists and compilation CD databases (CDDDB) (Pachet et al.;2001)

$$sim_{pl_cooc}(A_i, A_j) = \frac{1}{2} \cdot \left[\frac{cooc(A_i, A_j)}{cooc(A_i, A_i)} + \frac{cooc(A_j, A_i)}{cooc(A_j, A_j)} \right]$$

Analysis of 29K playlists from “Art of the Mix” (Cano, Koppenberger;2004): artists similar if they co-occur in playlist (highly sparse)

Analysis of >1M playlists from “MusicStrands” (Baccigalupo et al.; 2008):

- distance in playlists taken into account $\beta_0 = 1, \beta_1 = 0.8, \beta_2 = 0.64$

$$dist_{pl_d}(A_i, A_j) = \sum_{h=0}^2 \beta_h \cdot [d_h(A_i, A_j) + d_h(A_j, A_i)]$$

- playlist prediction using case-based reasoning

Last.fm Playcounts



Interpret number of plays of a song/artist as a “rating”

Results in a user-track rating matrix

Use **co-occurrences** of items in listening histories to compute similarities (relative no. users that listen to item A and B)

Use standard **collaborative filtering** approaches to predict similarities (or to recommend unknown music)

e.g., (Resnick et al.; 1994)

Item-based: compare tracks by calculating similarity on vectors over all users


User-based: find similar users by comparing listening pattern vectors; use to find relevant/similar tracks yet unknown to user

#nowplaying Approaches: Basics

(Schedl, ECIR 2013)

Listening events extracted from microblogs (Twitter)

- Streaming API crawled from 2011–2015
- Hashtag- and rule-based filtering of stream to identify listening events (#nowplaying, #itunes, #np, ...)
- Multi-level, rule-based analysis (artists/songs) to find relevant tweets (MusicBrainz)
- Last.fm, Freebase, Allmusic, Yahoo! PlaceFinder to annotate tweets



```
{
  "id_str": "142338125895696385",
  "place": null,
  "text": "#NowPlaying Christmas Tree-  
Lady Gaga",
  "in_reply_to_user_id": null,
  "favorited": false,
  "geo": null,
  "retweet_count": 0,
  "in_reply_to_screen_name": null,
  "in_reply_to_status_id_str": null,
  "source": "web",
  "retweeted": false,
  "in_reply_to_user_id_str": null,
  "coordinates": null,
  "created_at": "Thu Dec 01 20:23:48 +0000 2011",
  "in_reply_to_status_id": null,
  "contributors": null,
  "user": {
    "id_str": "20209983",
    "profile_link_color": "2caba5",
    "screen_name": "tamse77",
    "follow_request_sent": null,
    "geo_enabled": false,
    "favourites_count": 26,
    "location": "Maryland",
    "following": null,
    "verified": false,
    "profile_background_color": "e80e0e",
    "show_all_inline_media": true,
    "profile_background_tile": true,
    "follower_count": 309,
    "profile_image_url": "http://va1.twimg.com/profile_images/1647613274V392960_10150559294659517_793614516_11700077_1689597400_n_normal.jpg",
    "description": "being awesome since 1990.",
    "is_translator": false,
    "profile_background_image_url_https": "https://va10.twimg.com/profile_background_images/359728130Vframes.gif",
    "friends_count": 148,
    "profile_sidebar_fill_color": "ffffff",
    "default_profile": false,
    "listed_count": 3,
    "time_zone": "Central Time (US & Canada)",
    "contributors_enabled": false,
    "created_at": "Fri Feb 06 01:51:10 +0000 2009",
    "profile_sidebar_border_color": "f5f8ff",
    "protected": false,
    "notifications": null,
    "profile_use_background_image": true,
    "name": "Katie",
    "default_profile_image": false,
    "statuses_count": 22172,
    "profile_text_color": "615d61",
    "url": null,
    "profile_image_url_https": "https://va10.twimg.com/profile_images/1647613274V392960_10150559294659517_793614516_11700077_1689597400_n_normal.jpg",
    "id": 20209983,
    "lang": "en",
    "profile_background_image_url": "http://va2.twimg.com/profile_background_images/359728130Vframes.gif",
    "utc_offset": -21600,
    "truncated": false,
    "id": 142338125895696385,
    "entities": {
      "hashtags": [
        {
          "text": "NowPlaying",
          "indices": [0, 11]
        }
      ],
      "urls": [],
      "user_mentions": []
    }
  }
}
```



last.fm
allmusic

Freebase

twitter-id user-id month weekday longitude latitude
country-id city-id artist-id track-id <tag-ids>

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

Most Active Countries

#nowplaying		#itunes	
Country	Tweets	Country	Tweets
Brazil	725,389	USA	78,460
USA	673,839	Japan	30,932
Japan	458,558	Mexico	23,047
Mexico	419,584	Brazil	16,390
Indonesia	284,082	UK	15,134
South Korea	251,132	Canada	11,266
China	183,178	South Korea	8,652
UK	128,744	Australia	5,119
Netherlands	121,134	China	4,492
Venezuela	110,336	Germany	3,157


Most Active Cities

#nowplaying		#itunes	
City	Tweets	City	Tweets
New York	126,952	New York	13,603
London	96,801	London	9,813
São Paulo	79,317	Los Angeles	9,030
Los Angeles	73,834	San Francisco	5,787
Amsterdam	66,021	San Jose	5,605
Guarulhos	58,453	Chicago	4,413
Osasco	57,512	Birmingham	3,869
São Bernardo	56,946	Toronto	3,363
Rotterdam	55,113	Hamilton	3,279
Mexico City	52,618	Baltimore	3,245

Most Frequently Listened Artists

#nowplaying		#itunes	
Artist	Tweets	Artist	Tweets
Paramore	9,066	The Beatles	939
Drake	7,697	Daft Punk	683
Katy Perry	6,998	Britney Spears	567
Bruno Mars	6,932	Adele	462
Lady Gaga	6,919	Coldplay	428
Coldplay	6,434	Bruno Mars	416
Eminem	6,352	Katy Perry	374
Rihanna	6,038	The Black Eyes Peas	373
Taylor Swift	5,844	Kanye West	367
Usher	5,445	Lady Gaga	358
Muse	5,383	Avril Lavigne	308
Justin Bieber	5,028	Arcade Fire	299
The Beatles	4,579	Radiohead	266
Michael Jackson	4,476	Kings of Leon	240
Linkin Park	4,285	Duran Duran	238
Oasis	4,190	Michael Jackson	229
Kanye West	4,013	Linkin Park	228
Chris Brown	3,943	Eminem	211
Avril Lavigne	3,780	Muse	209
Radiohead	3,756	The Black Keys	203





Computing Similarity

(Schedl et al., MM Syst., 2013)

Use (variant of) co-occurrence approach to compute similarity (cf. previous part)

Co-occurrence on **per-user** basis

popularity correction function

$$sim_{cooc}(a_i, a_j) = c(a_i, a_j) \cdot p(a_i, a_j)$$

co-occurrence function

$$c_1(a_i, a_j) = \frac{cooc_{a_i, a_j}}{occ_{a_i}}$$

$$c_2(a_i, a_j) = \frac{cooc_{a_i, a_j}}{\min(occ_{a_i}, occ_{a_j})}$$

$$c_3(a_i, a_j) = \frac{cooc_{a_i, a_j}}{\max(occ_{a_i}, occ_{a_j})}$$


$$c_4(a_i, a_j) = \frac{cooc_{a_i, a_j}}{\frac{1}{2} \cdot (occ_{a_i} + occ_{a_j})}$$

$$c_5(a_i, a_j) = \frac{cooc_{a_i, a_j}}{occ_{a_i} \cdot occ_{a_j}}$$

$$c_6(a_i, a_j) = \frac{cooc_{a_i, a_j}}{\sqrt{occ_{a_i} \cdot occ_{a_j}}}$$


$$p(a_i, a_j) = 1 - \frac{|occ_{a_i} - occ_{a_j}|}{\max_k occ_{a_k}}$$

Results



(Schedl et al., MM Syst., 2013)

Evaluation against Last.fm similar artists: <http://www.last.fm/api/show/artist.getSimilar>



$$sim_{cooc}(a_i, a_j) = c(a_i, a_j) \cdot p(a_i, a_j)$$

$$c_1(a_i, a_j) = \frac{cooc_{a_i, a_j}}{occ_{a_i}}$$

$$c_2(a_i, a_j) = \frac{cooc_{a_i, a_j}}{\min(occ_{a_i}, occ_{a_j})}$$


$$c_3(a_i, a_j) = \frac{cooc_{a_i, a_j}}{\max(occ_{a_i}, occ_{a_j})}$$

$$c_4(a_i, a_j) = \frac{cooc_{a_i, a_j}}{\frac{1}{2} \cdot (occ_{a_i} + occ_{a_j})}$$

$$c_5(a_i, a_j) = \frac{cooc_{a_i, a_j}}{occ_{a_i} \cdot occ_{a_j}}$$


$$c_6(a_i, a_j) = \frac{cooc_{a_i, a_j}}{\sqrt{occ_{a_i} \cdot occ_{a_j}}}$$

Results



(Schedl et al., MM Syst., 2013)

Evaluation against Last.fm similar artists: <http://www.last.fm/api/show/artist.getSimilar>



$$sim_{cooc}(a_i, a_j) = c(a_i, a_j) \cdot p(a_i, a_j)$$

$$c_1(a_i, a_j) = \frac{cooc_{a_i, a_j}}{occ_{a_i}}$$

$$c_2(a_i, a_j) = \frac{cooc_{a_i, a_j}}{\min(occ_{a_i}, occ_{a_j})}$$

$$c_3(a_i, a_j) = \frac{cooc_{a_i, a_j}}{\max(occ_{a_i}, occ_{a_j})}$$

$$c_4(a_i, a_j) = \frac{cooc_{a_i, a_j}}{\frac{1}{2} \cdot (occ_{a_i} + occ_{a_j})}$$

$$c_5(a_i, a_j) = \frac{cooc_{a_i, a_j}}{occ_{a_i} \cdot occ_{a_j}}$$

$$c_6(a_i, a_j) = \frac{cooc_{a_i, a_j}}{\sqrt{occ_{a_i} \cdot occ_{a_j}}}$$

Co-occurrence-based Approaches: Summary

	Microblogs	Playcounts	P2P nets	Playlists
Source	API	listening service	shared folders	radio, compilations, Web services
Community-based	yes	yes	yes	depends on source
Level	artists (tracks)	tracks (artists)	tracks	artists (tracks)
Specific Bias	community	popularity	community	low


Other Applications of #nowplaying Data

- Geospatial music taste analysis
- Music discovery in the microblogosphere
- Music trend analysis

Music taste analysis

Most mainstreamy countries


(Schedl, Hauger; 2012)



Music taste analysis

Least mainstreamy countries

(Schedl, Hauger; 2012)



Music discovery in the microblogosphere

[Hauger, Schedl; AMR 2012]

MusicTweetMap: <http://www.cp.jku.at/projects/MusicTweetMap>

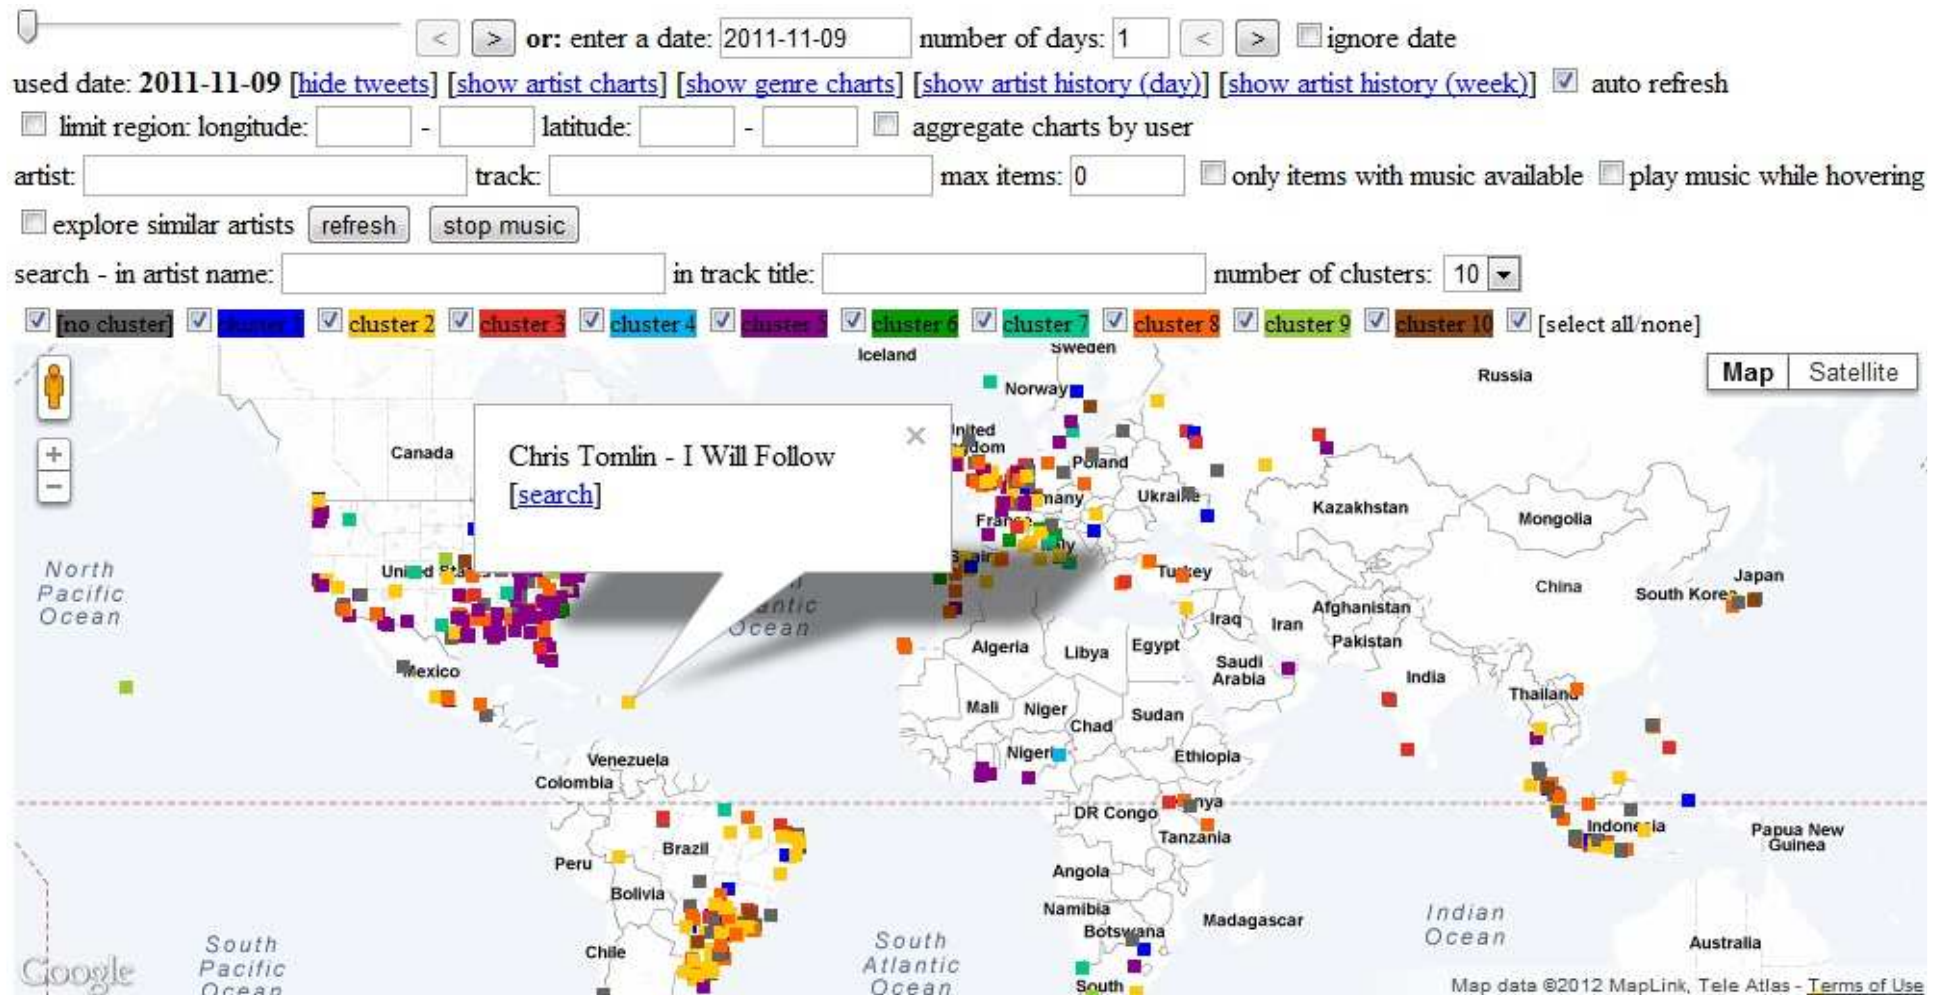
Exploring the “microblogosphere” of music from listening events
via...

- **time** (specific day or time range)
- **location** (country, state, city, longitude/latitude coordinates)
- **similar artists** (based on co-occurrences in tweets)
- **metadata**-based search (artist, track)
- **induced topics** with clustering based on Non-negative Matrix Factorization (NMF) on Last.fm tags → genres
- **30s snippets** where available
- artist **charts**, genre charts
- artist **play histories**



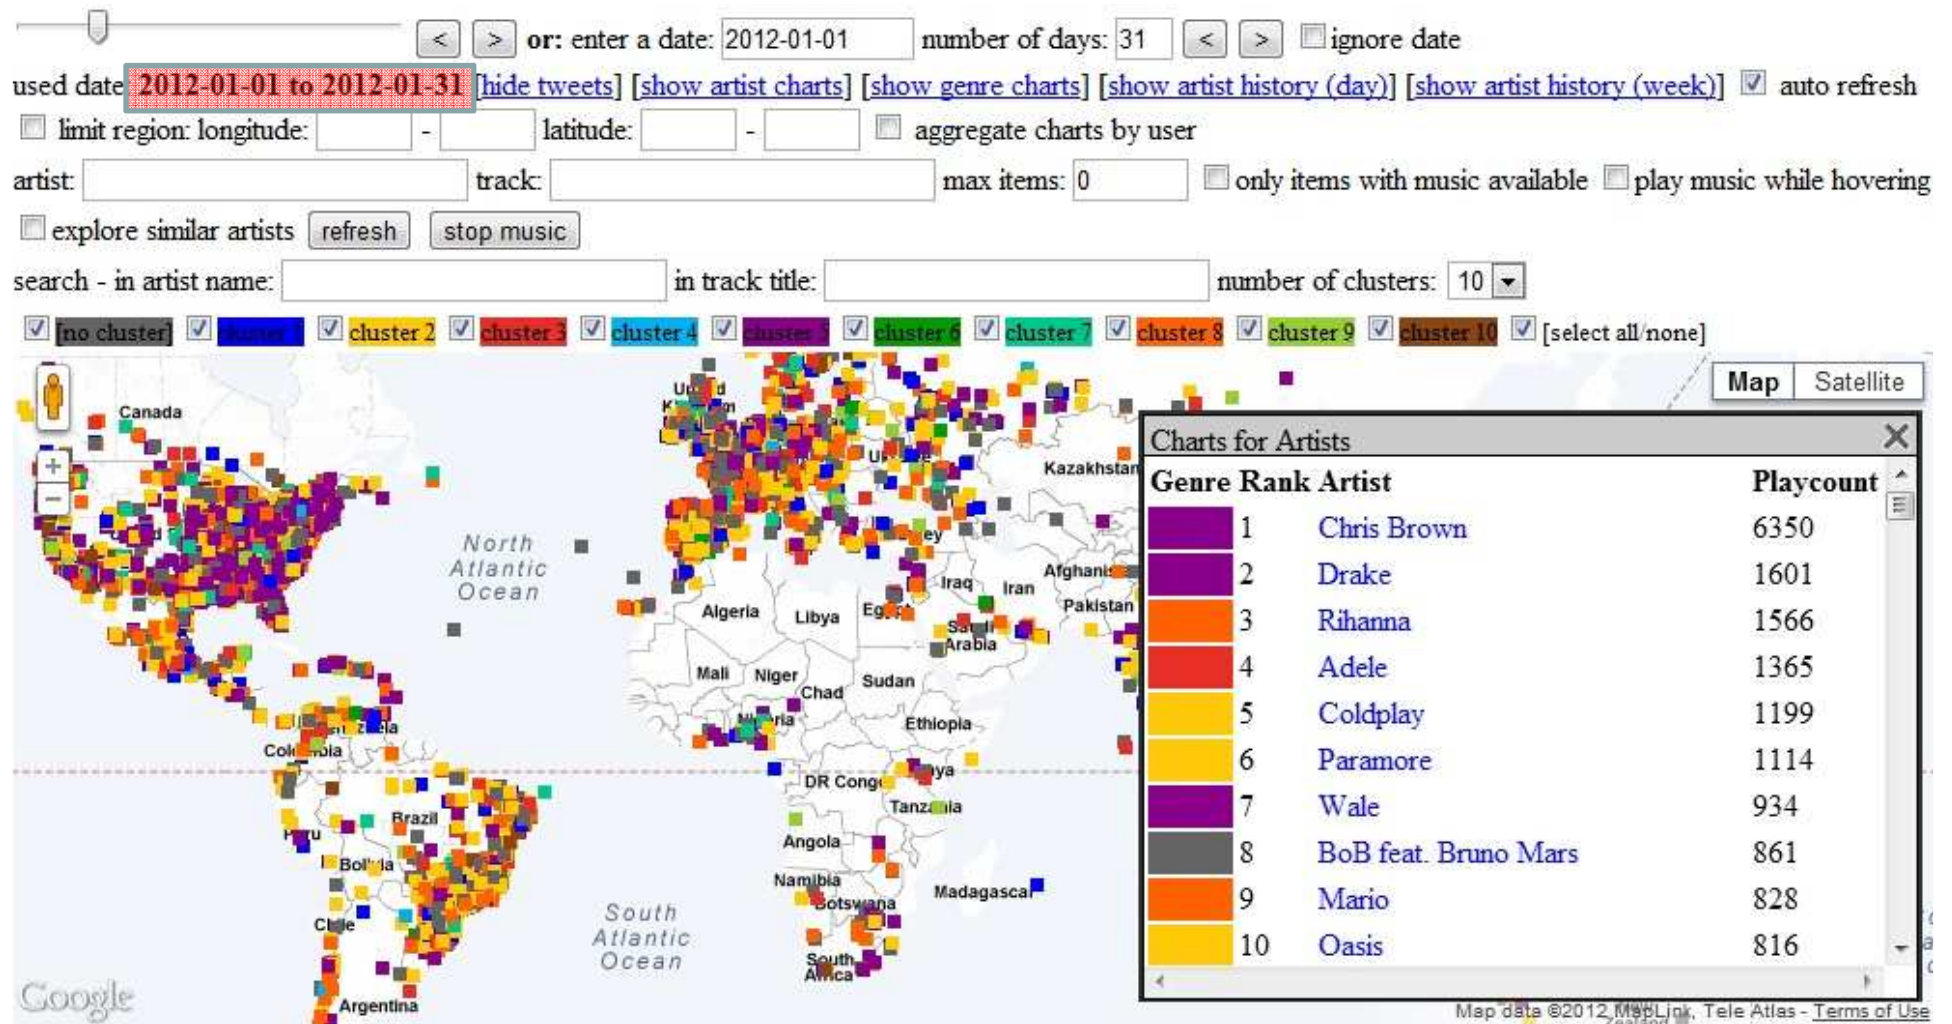
Music discovery in the microblogosphere

Visualization and browsing of geospatial music taste



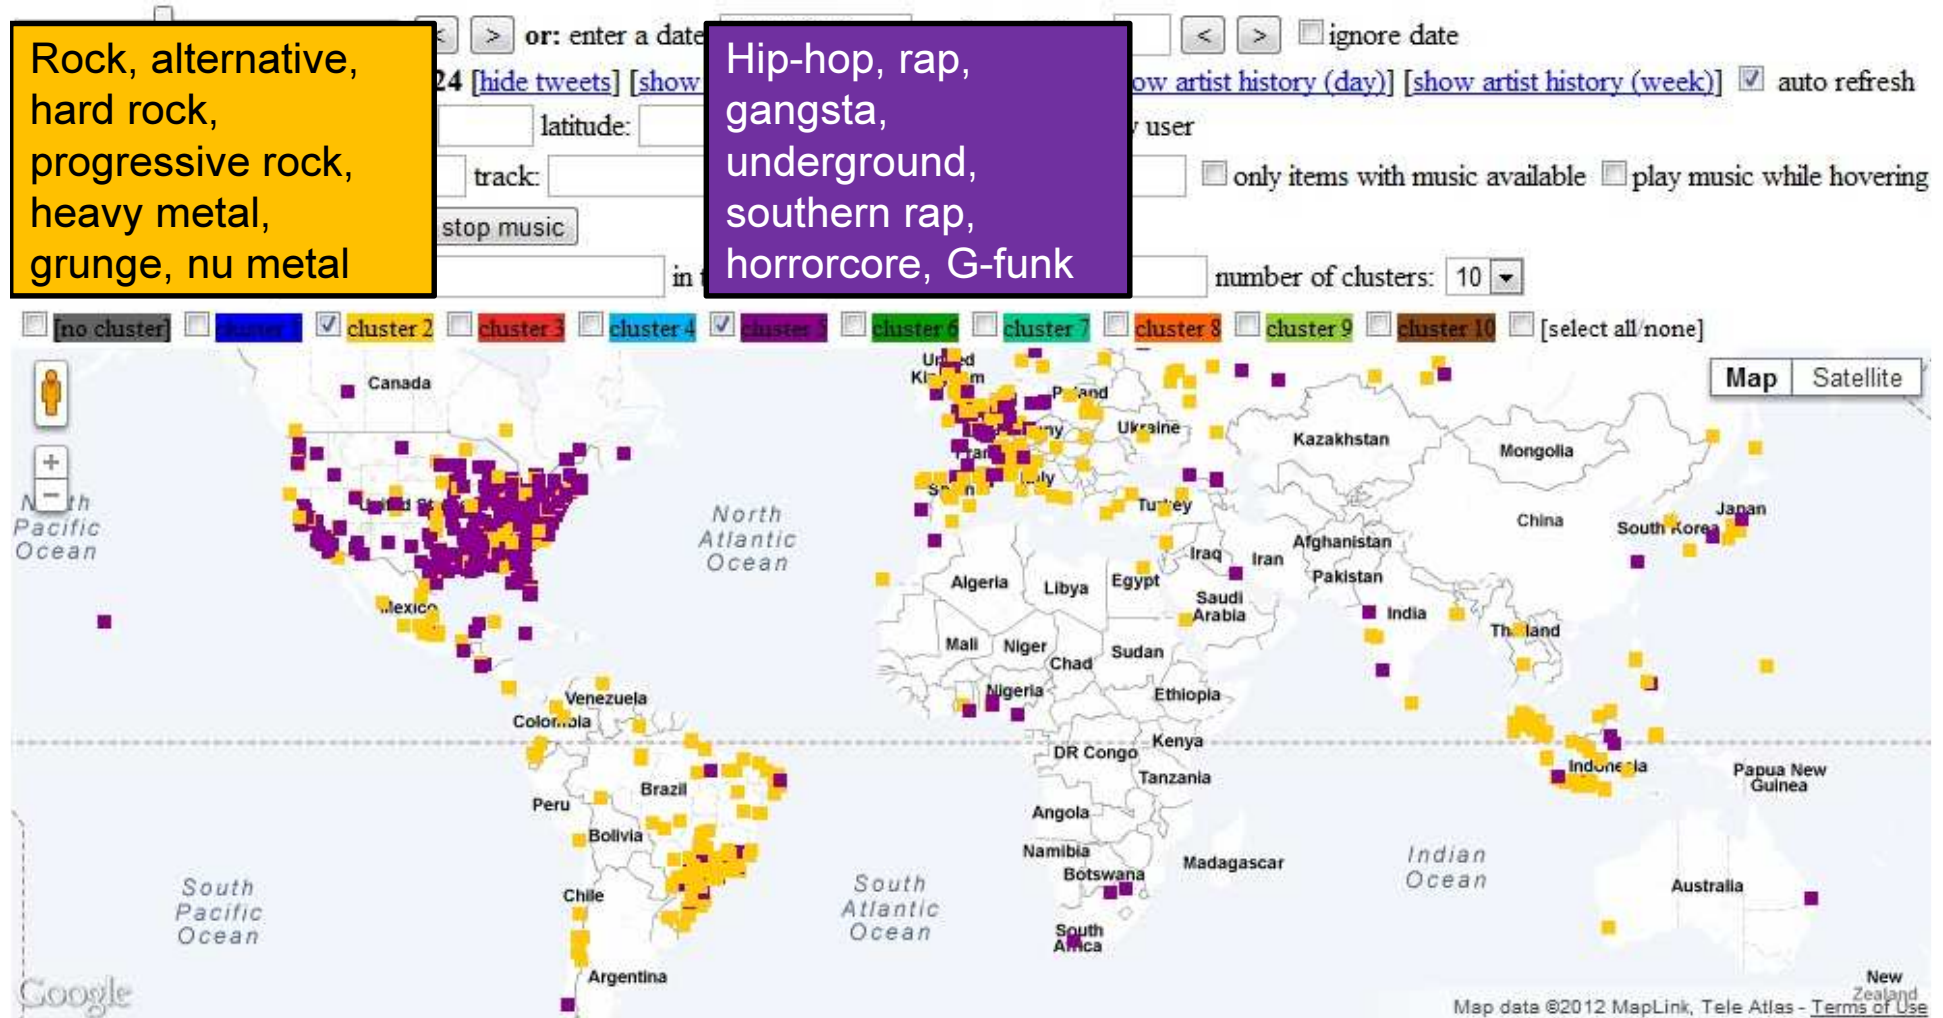
Music discovery in the microblogosphere

Browsing geospatial music taste by time



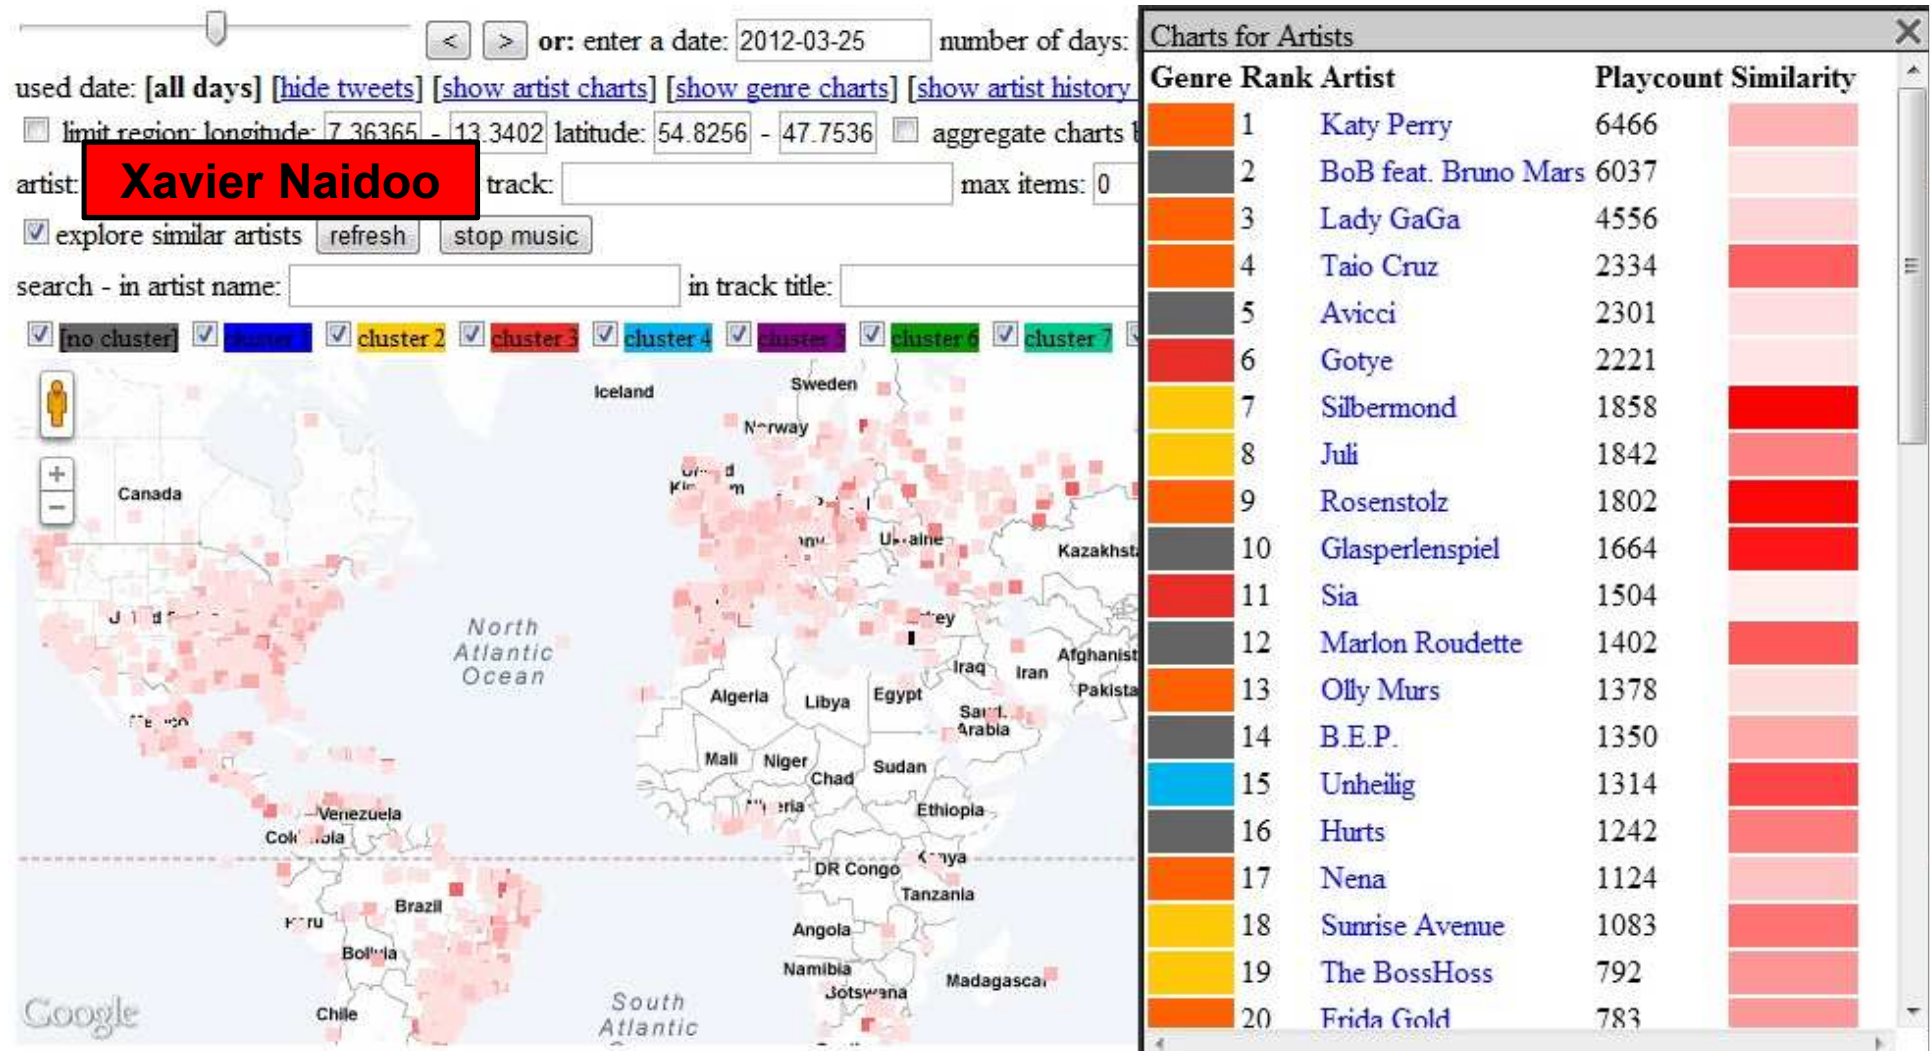
Music discovery in the microblogosphere

Geospatial music taste: “hip-hop” vs. “rock”




Music discovery in the microblogosphere

Exploring similar artists: Example “Xavier Naidoo”




Music trend analysis

Example: “The Beatles”



Music trend analysis

Example: “Madonna”




Context-aware Music Playlist Adaptation


(Schedl et al., ICMR 2014)

Mobile Music Genius (MMG):

<http://www.cp.jku.at/projects/MMG>

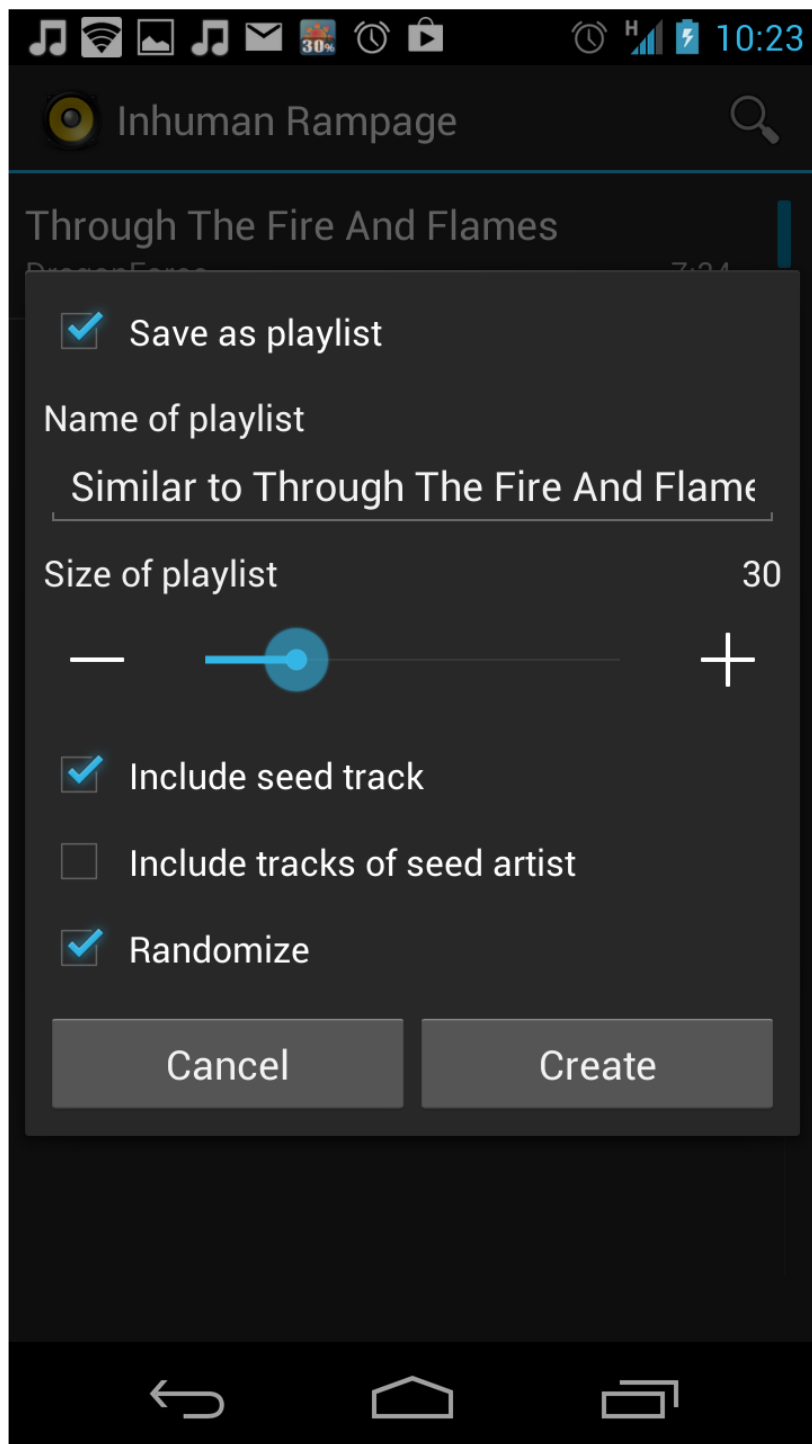
- music player for the Android platform
- collects user context and interaction data while playing
- adaptive system that learns user preferences from implicit feedback (player interaction: play, skip, duration played, playlists, ...)
- user modeling via relations:
user context – music preference (artist, track, genre)
- automatic playlist modification according to changes in listener's current context





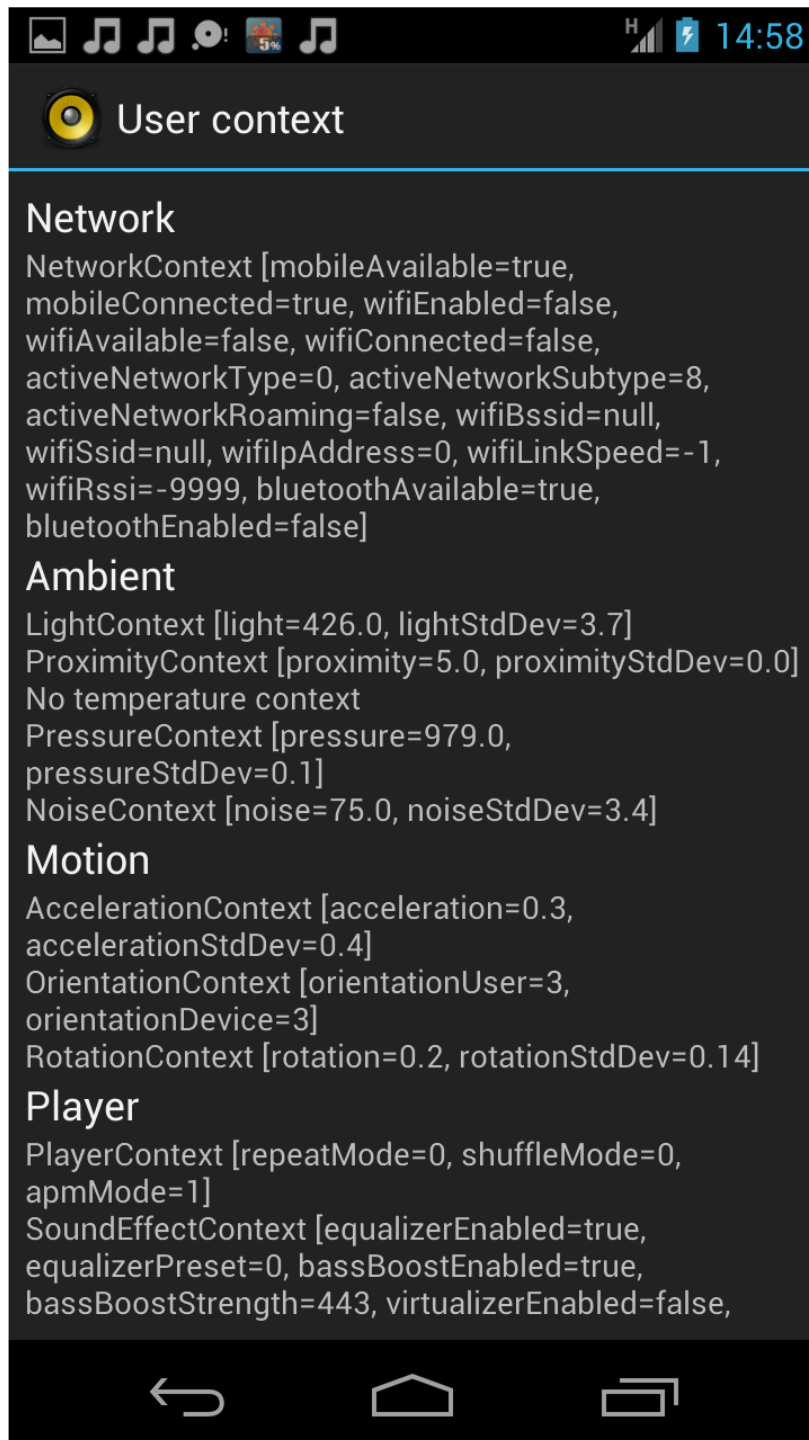
Mobile Music Genius

Music player in adaptive
playlist generation mode



Mobile Music Genius

Automatic playlist generation based on music context (features and similarity computed based on Last.fm tags)




Mobile Music Genius

Some user context
features gathered while
playing

Mobile Music Genius: Listener Features

<i>Time:</i>	day of week, hour of day
<i>Location:</i>	longitude/latitude, accuracy, speed, altitude
<i>Weather:</i>	wind direction, speed, clouds, temperature, dew point, humidity, air pressure
<i>Device:</i>	manufacturer, model, phone state, connectivity, storage, battery, various volume settings
<i>Phone:</i>	service state, roaming, signal strength, network type
<i>Task:</i>	recently used tasks/apps, screen on/off, docking mode
<i>Network:</i>	mobile network available/connected, active network, Bluetooth, WiFi
<i>Ambient:</i>	light, proximity, noise
<i>Motion:</i>	acceleration, user and device orientation
<i>Player:</i>	artist, album, track, track length, genre, playback position, playlist name, playlist type, player state (repeat, shuffle mode), audio output (headset plugged), sound effects
<i>Activity:</i>	mood and activity (direct user feedback)

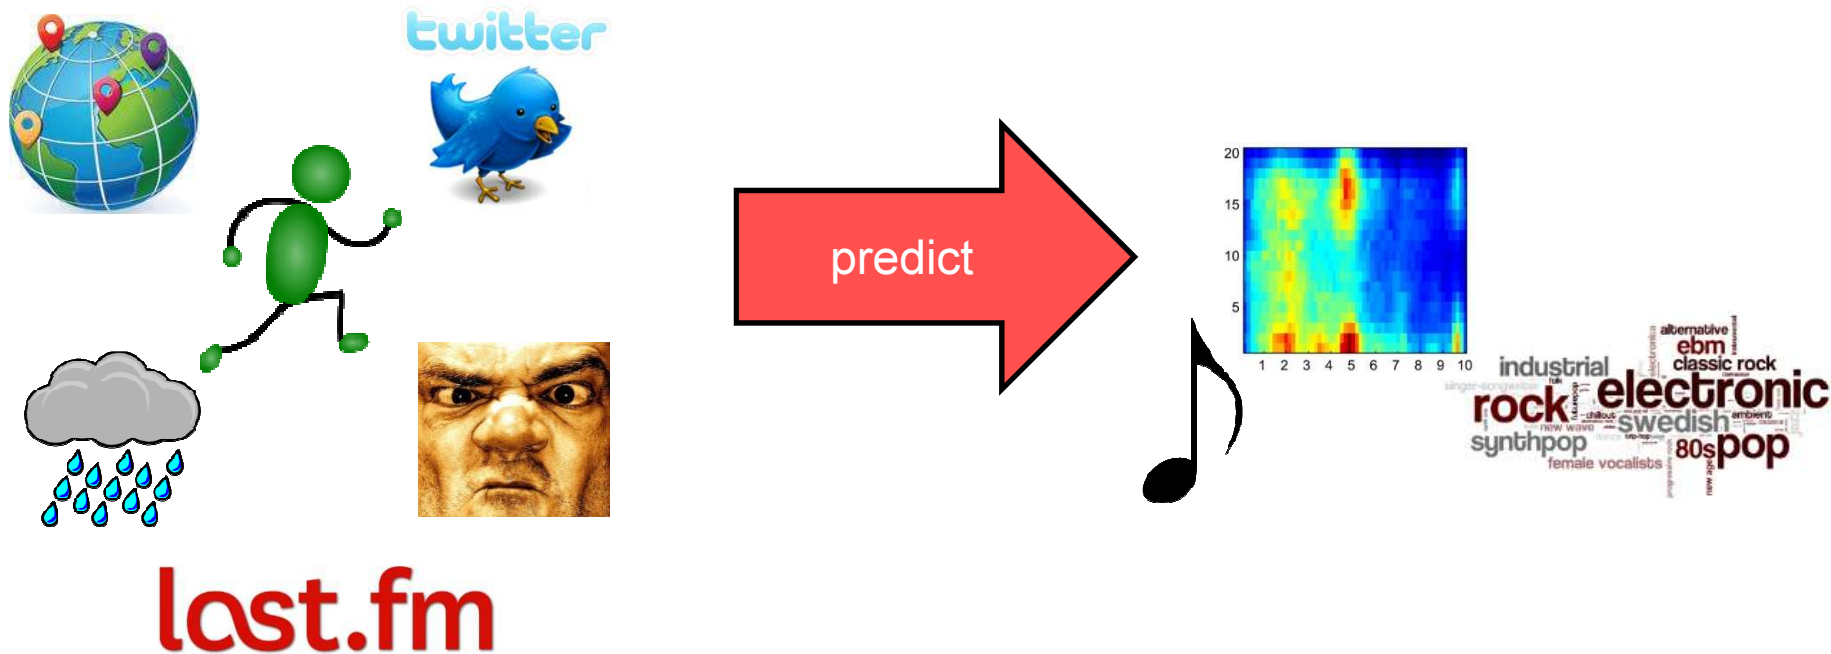


Context-aware Music Playlist Adaptation

(Schedl and Gillhofer, MMM 2015)

Evaluation via music preference prediction


Objective: To understand the complex interrelationship between user characteristics, context, and music taste.



Context-aware Music Playlist Adaptation

Music preference prediction: Data collection

- MMG used by students of JKU from January to July 2013
- collected 7,628 data points from 48 persons
- 4,149 unique tracks by 1,169 unique artists
- enriched by Last.fm: 24 genres and 70 moods



Property	Mean	users show quite diverse music taste
Artists per user	27.88	
Genres per user	5.14	
Moods per user	9.91	
Titles per user	89.16	

Context-aware Music Playlist Adaptation


Music preference prediction: Experimental setup

- 4 classifiers (+ majority voting “ZeroR” as baseline):
 - k-nearest neighbors (“IBk”)
 - decision tree (“J48”)
 - rule learner (“JRip”)
 - Random Forest
- 4 levels of prediction (artist, track, mood, genre)
- 10-fold cross validation
- different feature sets (time, location, weather, motion, ...)

Context-aware Music Playlist Adaptation

Music preference prediction: Results

Q1: Can we predict music taste from user context factors?




Predictive Accuracy for Title, Mood, Artist, Genre

Context-aware Music Playlist Adaptation

Music preference prediction: Results

Q1: Can we predict music taste from user context factors?



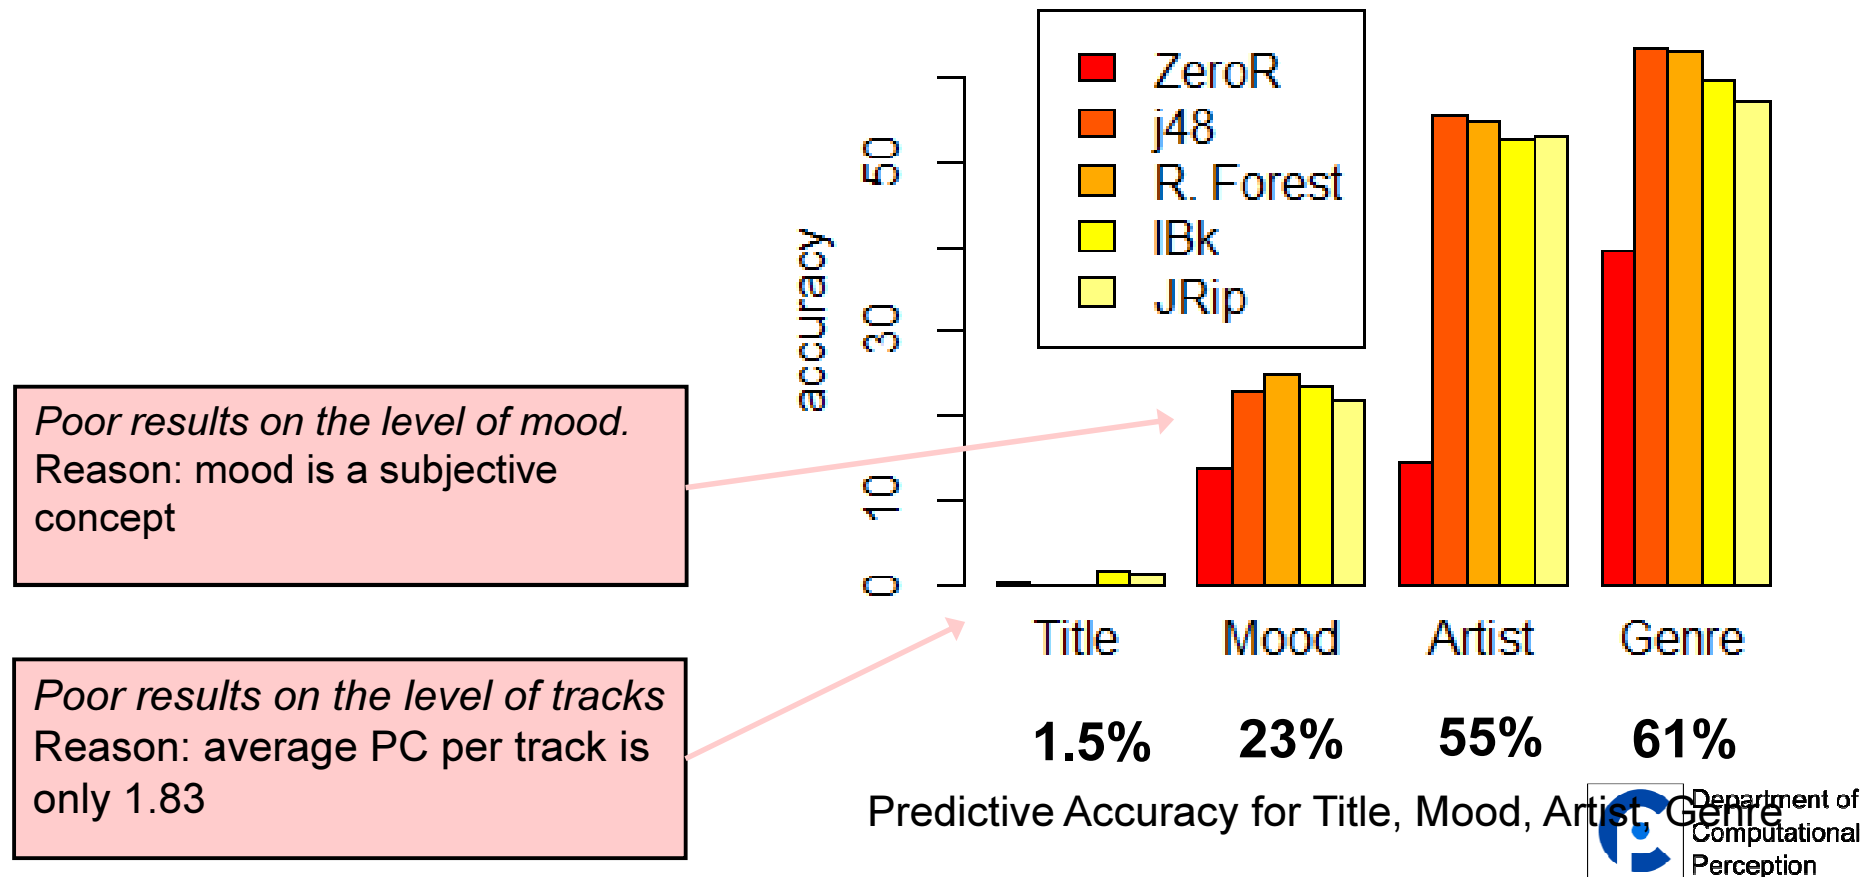
Poor results on the level of tracks
Reason: average PC per track is only 1.83

1.5% 23% 55% 61%
Predictive Accuracy for Title, Mood, Artist, Genre

Context-aware Music Playlist Adaptation

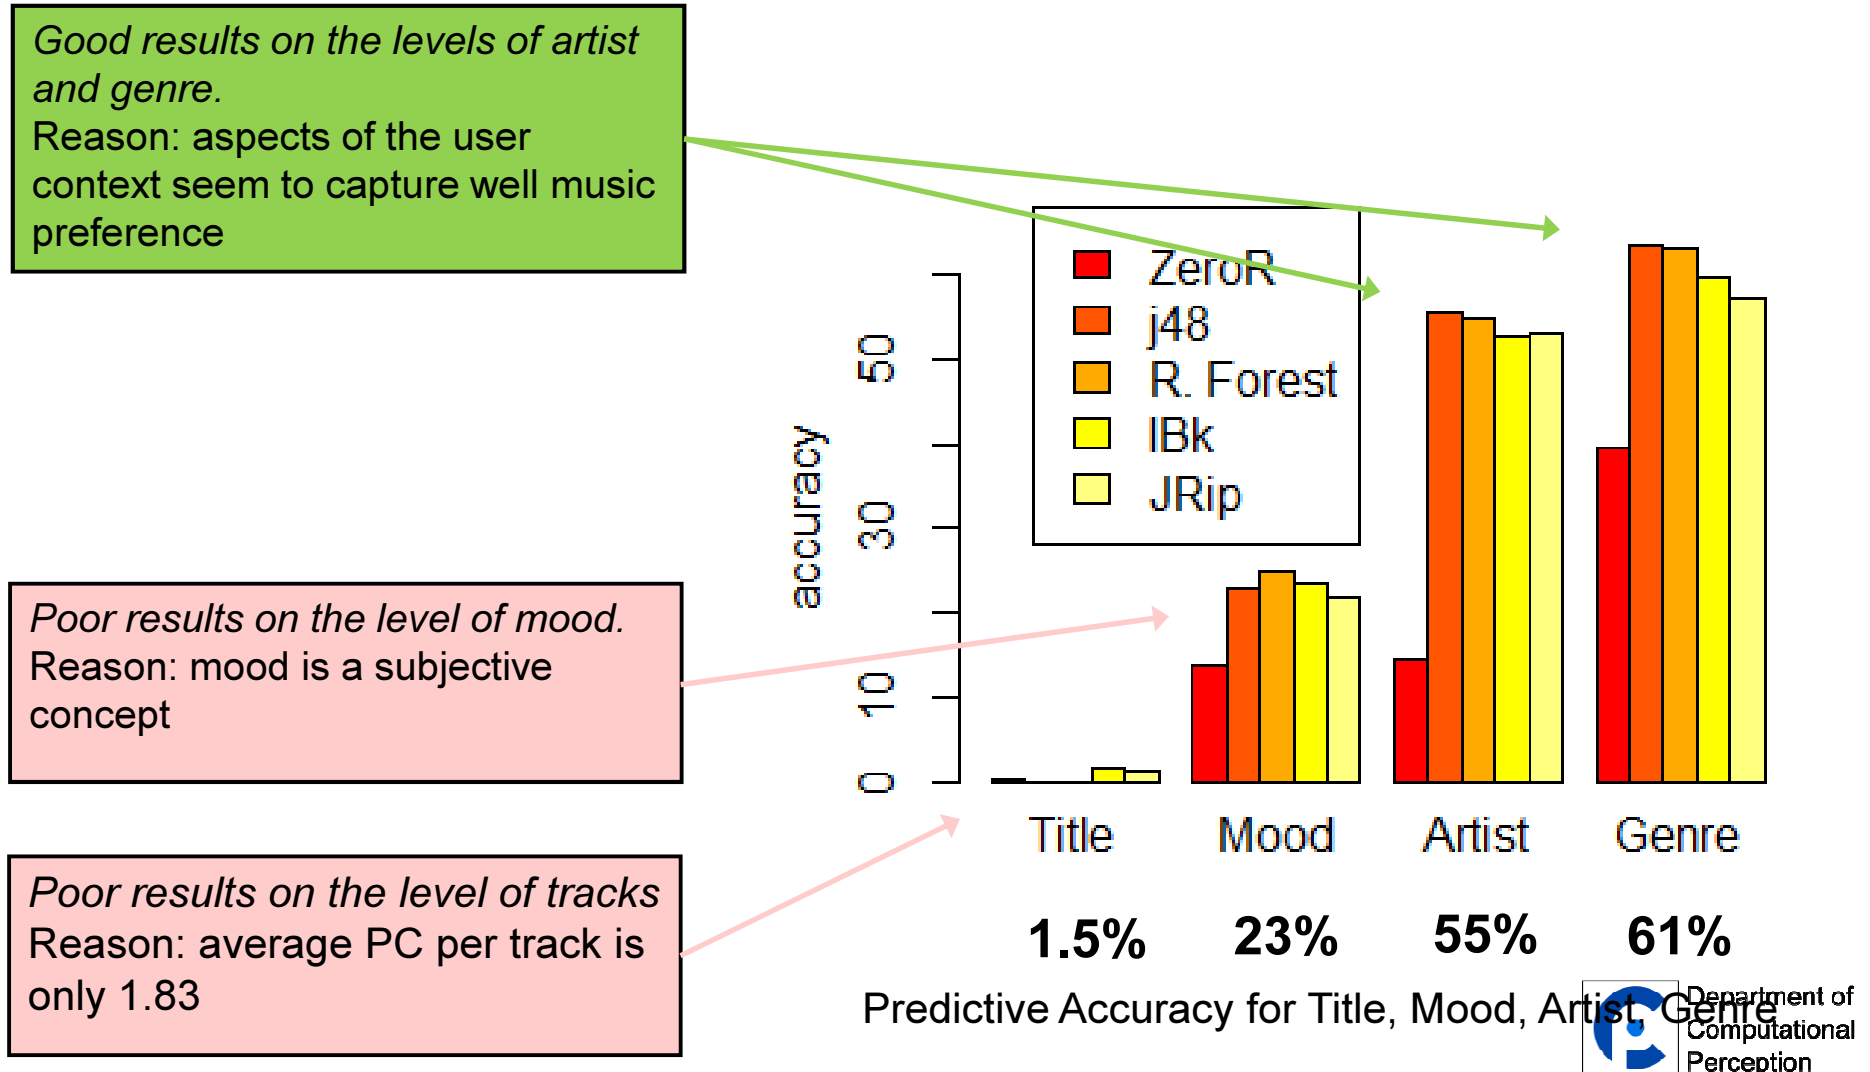
Music preference prediction: Results

Q1: Can we predict music taste from user context factors?



Context-aware Music Playlist Adaptation

Music preference prediction: Results




Context-aware Music Playlist Adaptation

Music preference prediction: Results

Q2: Which contextual factors are most promising for prediction?

Adjusting player settings (e.g., repeat mode) seems to be a reasonable indicator of song preference.



Relative importance of each feature group compared to the mean classification result (over all feature categories).

Context-aware Music Playlist Adaptation


Music preference prediction: Results

Q2: Which contextual factors are most promising for prediction?

Adjusting player settings (e.g., repeat mode) seems to be a reasonable indicator of song preference.

User-indicated activity is less important than apps running on the device.

Device features capture general music preference. Running apps, time, location, and weather capture dynamic adjustments.



Relative importance of each feature group compared to the mean classification result (over all feature categories).

Music Recommendation for Places of Interest

(Kaminskas et al.; RecSys 2013)

recommend music that is suited to a place of interest (POI) of the user (context-aware)

La Scala, Milan, Italy

http://en.wikipedia.org/wiki/La_Scala




La Scala is a world renowned opera house in Milan, Italy. The theatre was inaugurated on 3 August 1778 and was originally known as the New Royal-Ducal Theatre at La Scala. The premiere performance was Antonio Salieri's 'Europa riconosciuta'. Most of Italy's greatest operatic artists, and many of the finest singers from around the world, have appeared at La Scala during the past 200 years.

Session 1 out of 10: ◆ ◆ ◆ ◆ ◆ ◆ ◆ ◆ ◆ ◆

Listen to the tracks and select those that in your opinion are **suited** for the described place:

Reincidentes - Ay Dolores

<http://en.wikipedia.org/wiki/Reincidentes>



Vincenzo Pucitta - La Vestale, Opera seria 1st act

http://en.wikipedia.org/wiki/Vincenzo_Pucitta



The Shower Scene - This Is The Call Out

http://en.wikipedia.org/wiki/The_Shower_Scene



Duchess Maria Antonia of Bavaria - Pallid' ombra che d'intorno

http://en.wikipedia.org/wiki/Duchess_Maria_Antonia_of_Bavaria



Submit

Music Recommendation for Places of Interest

(Kaminskas et al.; RecSys 2013)

Approaches:


- *genre-based*: only play music belonging to the user's preferred genres (baseline)

Music Recommendation for Places of Interest

(Kaminskas et al.; RecSys 2013)

Approaches:

- *knowledge-based*: use the DBpedia knowledge base (relations between POIs and musicians)



As of September 2011

Music Recommendation for Places of Interest

(Kaminskas et al.; RecSys 2013)




Approaches:

- *tag-based*: user-assigned emotion tags describing images of POIs and music, Jaccard similarity between music-tag-vectors and POI-tag-vectors

Tag:

<input type="checkbox"/> Melancholic	<input type="checkbox"/> Bright
<input type="checkbox"/> Heavy	<input type="checkbox"/> Animated
<input checked="" type="checkbox"/> Tender	<input type="checkbox"/> Energetic
<input type="checkbox"/> Cold	<input type="checkbox"/> Spiritual
<input checked="" type="checkbox"/> Modern	<input checked="" type="checkbox"/> Serene
<input type="checkbox"/> Ancient	<input type="checkbox"/> Calm
<input type="checkbox"/> Affectionate	<input type="checkbox"/> Sad
<input checked="" type="checkbox"/> Dark	<input type="checkbox"/> Strong
<input checked="" type="checkbox"/> Lightweight	<input type="checkbox"/> Colorful
<input checked="" type="checkbox"/> Open	<input type="checkbox"/> Thrilling
<input type="checkbox"/> Warm	<input type="checkbox"/> Agitated
<input type="checkbox"/> Sentimental	<input type="checkbox"/> Bouncy

Fritz Kreisler - Liebesfreud
http://en.wikipedia.org/wiki/Fritz_Kreisler

00:08  00:31   FMP3

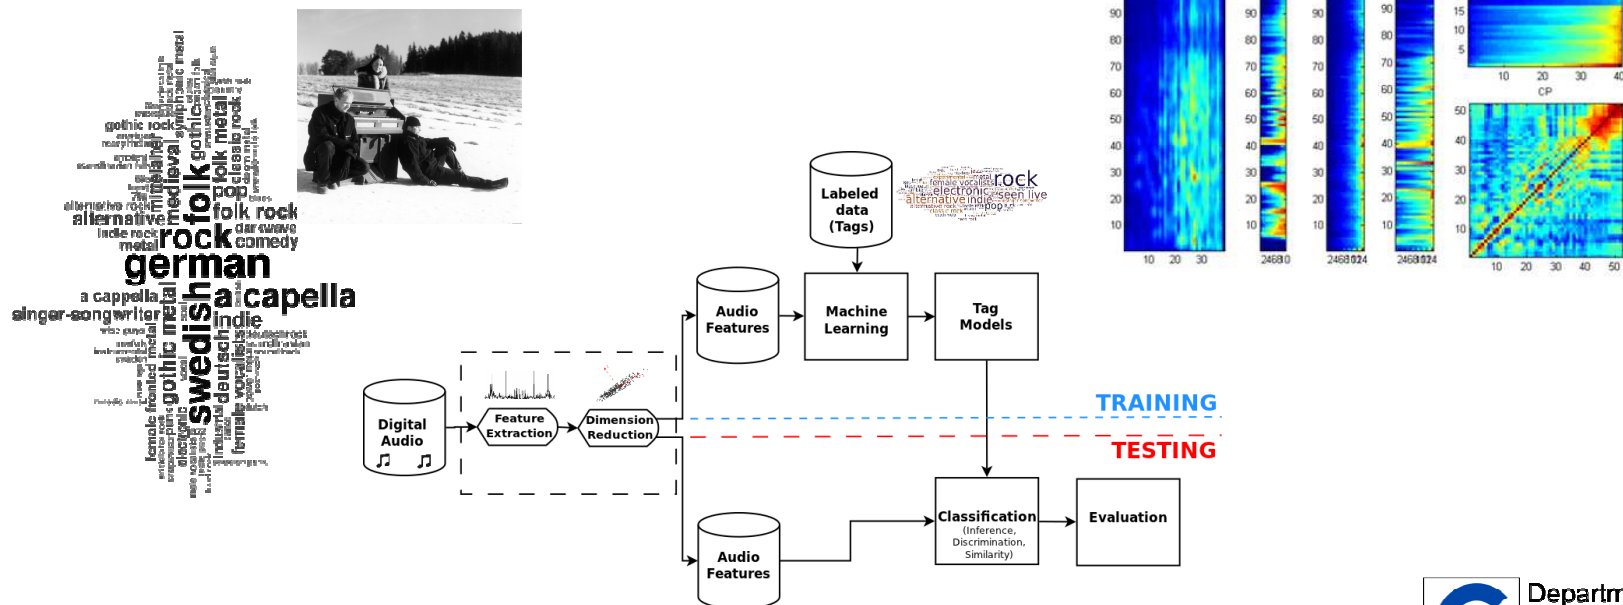
"Friedrich 'Fritz' Kreisler (February 2, 1875 – January 29, 1962) was an Austrian-born violinist and composer. One of the most famous violin masters of his or any other day, he was known for his sweet tone and expressive phrasing. Like many great violinists of his generation, he produced a characteristic sound which was immediately recognizable as his own. Although he derived in many respects from the Franco-Belgian school, his style is nonetheless reminiscent of the gemütlich (cozy) lifestyle of pre-war Vienna."

Music Recommendation for Places of Interest

(Kaminskas et al.; RecSys 2013)

Approaches:

- auto-tag-based*: use state-of-the-art music auto-tagger based on the Block-level Feature framework to automatically label music pieces; then again compute Jaccard similarity between music-tag-vectors and POI-tag-vectors



Music Recommendation for Places of Interest

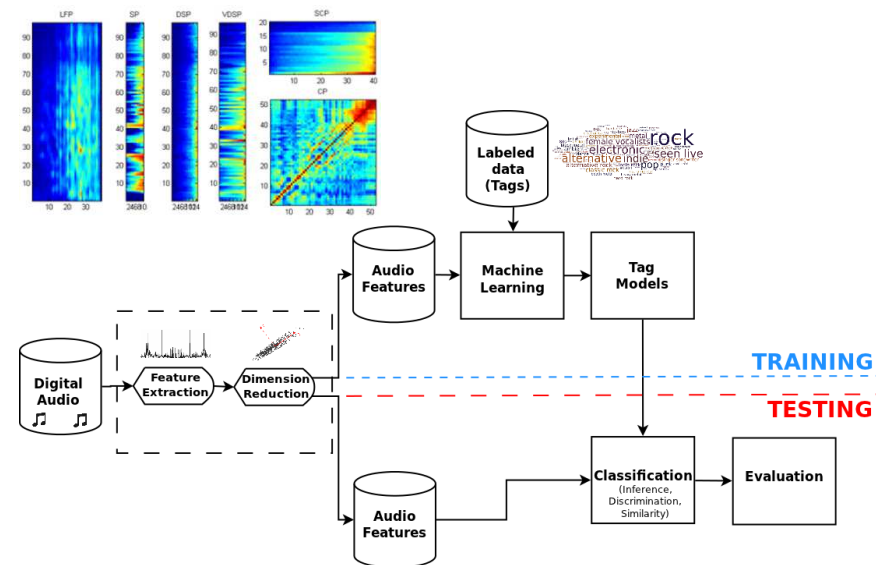
(Kaminskas et al.; RecSys 2013)

Approaches:

- *combined*: aggregate music recommendations w.r.t. ranks (Borda count) given by knowledge-based and auto-tag-based approaches



As of September 2011



Music Recommendation for Places of Interest

(Kaminskas et al.; RecSys 2013)

Approaches:

- *genre-based*: only play music belonging to the user's preferred genres (baseline)
- *knowledge-based*: using the DBpedia knowledge base (relations between POIs and musicians)
- *tag-based*: user-assigned emotion tags describing images of POIs and music, Jaccard similarity between music-tag-vectors and POI-tag-vectors
- *auto-tag-based*: using state-of-the-art music auto-tagger based on the Block-level Feature Framework to automatically label music pieces; then again use Jaccard similarity between music-tag-vectors and POI-tag-vectors
- *combined*: aggregate music recommendations w.r.t. ranks given by knowledge-based and auto-tag-based approaches

Music Recommendation for Places of Interest

(Kaminskas et al.; RecSys 2013)

Evaluation:

- user study via web interface (58 users, 564 sessions)

La Scala, Milan, Italy
http://en.wikipedia.org/wiki/La_Scala



La Scala is a world renowned opera house in Milan, Italy. The theatre was inaugurated on 3 August 1778 and was originally known as the New Royal-Ducal Theatre at La Scala. The premiere performance was Antonio Salieri's 'Europa riconosciuta'. Most of Italy's greatest operatic artists, and many of the finest singers from around the world, have appeared at La Scala during the past 200 years.

Session 1 out of 10: 

Listen to the tracks and select those that in your opinion are **suited** for the described place:

Reincidentes - Ay Dolores

<http://en.wikipedia.org/wiki/Reincidentes>



Vincenzo Pucitta - La Vestale, Opera seria 1st act

http://en.wikipedia.org/wiki/Vincenzo_Pucitta



The Shower Scene - This Is The Call Out

http://en.wikipedia.org/wiki/The_Shower_Scene



Duchess Maria Antonia of Bavaria - Pallid' ombra che d'intorno

http://en.wikipedia.org/wiki/Duchess_Maria_Antonia_of_Bavaria




Submit

Music Recommendation for Places of Interest

(Kaminskas et al.; RecSys 2013)

Evaluation:

- Performance measure: number of times a track produced by each approach was considered as well-suited in relation to total number of evaluation sessions, i.e. probability that a track marked as well-suited by a user was recommended by each approach




Music recommendation tailored to user characteristics

(Schedl and Hauger, SIGIR 2015)

(Schedl et al., ECIR 2015)

- Determine influence of user characteristics on music recommendation
- Select recommendation algorithm that performs best for a given user (category).



- Large-scale experiments on Last.fm datasets (100Ms listening events)

Music recommendation tailored to user characteristics

(Schedl and Hauger, SIGIR 2015)

(Schedl et al., ECIR 2015)



User aspects: age, gender, country, genre affinities, listening frequency, novelty/openness, diversity, mainstreaminess, time of day, ... personality?

Music recommendation tailored to user characteristics

(Schedl and Hauger, SIGIR 2015)

(Schedl et al., ECIR 2015)



Music recommendation algorithms:

PB	popularity-based recommender
CF	user-based collaborative filtering
CB (IB)	content-based (using term weights from tags)
LB (CULT)	location-based (CF enriched with location info)
RB	random baseline model: randomly picks users

Various algorithmic choices: normalization, score aggregation, fusion of combined methods, etc.

Aim: To identify music recommendation algorithms that perform superior for a given group of listeners, characterized by **diversity**, **mainstreamness**, and **novelty** in listening behavior.

User Characteristics

Diversity: How diverse is the user's music taste?

D_{PC} : avg. number of listening events per track
 D_{genres} : number of unique genre tags in user's listening history

Mainstreamness: How close to the mainstream is the user's listening history?

similarity between user's distribution of listening events (relative frequency per item) and global distribution
 M_{global} : considering user's entire listening history
 M_{avg_6m} : averaged over time windows of 6 months

Novelty: How open is the user to listen to new music?

percentage of items appearing for the first time in user's listening history, in a given time slot
 N_{avg_6m} : averaged over time windows of 6 months

Recommendation Models

Individual:

PB: popularity-based recommendation (recommends the most popular artists in the dataset)
 CF: user-based collaborative filtering
 IB: instance-based/content-based recommendation (based on Last.fm tags)
 LB: location-based recommendation (CF, considering only users in same country)
 RB: random baseline: recommendation based on randomly selected users

Hybrid:

PB+CF, PB+IB, PB+LB, CF+IB, PB+CF+IB

Aggregation functions:

mean, max

Normalization functions:

none, Gaussian, sum-to-1, max-to-1

Fusion functions:

mean, max, sum, multiply, Borda count

Experiments and Results:

- Last.fm users: ~200 million listening events, ~16,500 users, ~1.1 million artists
- 5-fold cross-validation for each user's listening history
- Precision, recall, and F-measure for varying numbers of recommended artists

Main Findings:

- Grouping users w.r.t. the defined categories of D, M, and N (low, medium, and high) yields better recommendations than using the entire user set.
- Hybrid methods involving PB tend to perform best for many user groups.
- IB alone performs poorly, but integrating IB diversifies recommendations (good for medium and high diversity users).
- Recommending music to low mainstreamness users is very hard.
- LB only works for users with high global mainstreamness and for low novelty users.

Results for user sets created according to level of global mainstreamness:

US_M_global_h			
Method	Precision	Recall	F-score
RB	3.78	8.35	4.86
PB	10.35	14.62	7.59
CF _{mean}	13.28	3.90	5.61
PB + LB _{mean}	8.72	13.20	9.66
US_M_global_m			
Method	Precision	Recall	F-score
RB	1.37	2.58	1.21
PB + IB _{mean}	2.00	26.00	3.58
PB + CF _{mean}	4.93	12.70	5.12
PB + CF _{mean} + IB _{mean}	8.25	0.86	1.48
US_M_global_l			
Method	Precision	Recall	F-score
RB	5.15	5.83	4.56
PB + LB _{mean}	1.74	10.67	2.25

Results on entire user set:

Method	Precision	Recall	F-score
RB	1.55	4.49	1.58
PB	4.82	11.83	4.26
CF _{mean}	5.41	9.26	4.42
CF _{max}	5.22	9.03	4.20
IB _{mean}	1.82	5.09	1.45
LB _{mean}	0.44	3.25	0.55
LB _{max}	2.59	4.74	2.22
PB + IB _{mean}	3.95	11.30	3.71
PB + CF _{mean}	6.35	10.00	5.14
PB + CF _{max}	6.34	9.70	4.96
PB + LB _{mean}	4.58	9.38	4.02
CF _{mean} + IB _{mean}	4.51	10.74	4.13
CF _{mean} + LB _{mean}	4.91	8.52	4.01
PB + CF _{mean} + IB _{mean}	4.99	13.85	4.75

Music recommendation tailored to user characteristics: Results

Come and see our poster on Tue, 16.30

SUMMARY

Music Information Retrieval is a great field

Various approaches to extract information from the audio signal

Various sources and approaches to extract contextual data and similarity information from the web, social media, and sensor data

Multimodal modeling and retrieval is important and allows for exciting applications

Next big challenges:

- considering cultural and multi-lingual aspects
- modeling user properties and context in a holistic way
- improve personalized and context-aware experience (serendipity)
- multimodal integration of complementary data sources
- new and better suited evaluation strategies