

SGN-9206 Signal processing graduate seminar II

Music Information Retrieval

Anssi Klapuri (@tut.fi), Tampere University of Technology, Finland


Outline:

- Motivation for music retrieval
- User interfaces
- Music similarity
- Music classification
- Transcription and query by humming
- Music fingerprinting
- Indexing techniques
- Toolboxes
- Summary of seminar topics (choose yours!)

Introduction

- Music information retrieval (MIR) is currently an active research area
- See proceedings of ISMIR conference and annual MIREX evaluations
- Background for this growth:
We have huge amount of music available and most of that is easily and cheaply accessible over the Internet

Search Music



[Get alternatives for this song](#)

Now Playing

Favorites

My Playlists

Save Playlist

You can drag and drop songs here.

0 Songs

[Clear All](#) | [Save Playlist](#)

Artists



Pat Metheny
Jazz

1 Artists

Albums



Pat Metheny
Pat Metheny ...



Pat Metheny
Works I



Pat Metheny
Trio 99>00



Pat Metheny
Offramp



Pat Metheny
Secret Story



Pat Metheny
Bright Size ...



Pat Metheny
Question and ...



Pat Metheny
Still Life (...)

38 Albums

Next

Clips



Pat Metheny
Aprilwind



Pat Metheny
April Joy



Pat Metheny
Lone Jack



Pat Metheny
San Lorenzo



Pat Metheny
Phase Dance



Pat Metheny
Jaco



Pat Metheny
Yolanda, You Le...



Pat Metheny
Tears Inside

[Adult Album Alternative](#)
[British Invasion](#)
[Classic Rock](#)
[Garage Rock](#)

[Glam](#)
[Hard Rock](#)
[Jam Bands](#)

[Prog/Art Rock](#)
[Psychedelic](#)
[Rock & Roll](#)

[Rockabilly](#)
[Singer/Songwriter](#)
[Surf](#)



What began as a revolution of popular music in the 50s has become the most diverse style of modern music. Innovative singers like Chuck Berry, Elvis, and The Beatles ignited a culture that continues to shape the genre's massive landscape. With way too many to mention, some top **Rock** bands include the Rolling Stones, U2, and Green Day. -- [cutripgareg](#)

Related genres: [Alternative](#), [Electric Blues](#), [Metal](#), [Soft Rock](#), [Latin Rock](#)

[Talk about Rock](#)

Showing 1-15 of 639 *Rock* stations

1 | 2 | 3 | 4 | Next

#	Station Title / Description	Genre	Audio	TLH
1.	Modern Rock Classics Start the week with Live365's Best Rock Station! [Editor's pick] [Professional]	rock , modern rock , hard rock	CD 128k	★★★★☆
2.	Classic Hits Radio Start Your Work Week With The Best Mix of Classic Hits at 128K! [Professional]	classic rock , rock , adult album alternative	CD 128k	★★★★☆
3.	102.1 KPRI 102.1 KPRI - 102.1 KPRI FM - San Diego's Finest Rock [AM/FM] [Professional]	rock , blues , reggae	CD 96k ●	★★★☆☆
4.	ROCKLAND USA : POP ROCK ▶▶ CHANNEL 1 - NO LOG-INS REQUIRED - IF STATION IS FULL USE CHANNEL 2: live365.com/stations/rocklandusa2 - The Best At-Work Popular Top 40 Rock Hits. 2005 Broadcaster of the Year and Best Pop Station [Editor's pick] [Professional]	pop , rock , alternative	CD 64k ●	★★★★☆
5.	Rock-N-Radio.com If It Rocks We Play It! AC/DC to ZZ Top - Guns N' Roses Judas Priest Kiss Metallica Ozzy Osbourne Van Halen Def Leppard Bullet Boys [Professional]	rock , metal , hard rock	CD 64k ●	★★★★☆
6.	WUMB -- 56 kb/s WUMB Folk Radio -- Stereo [AM/FM] [Professional]	folk , classic rock , rock	FM 56k	★★★★☆
7.	Zappa Radio Frank Zappa: 'Music Is The Best' - That's right folks, you heard Right. Don't touch that dial! Zapparama & Randomonium too. Puts the Hrз on your nose. KILL UGLY RADIO!!!! [Official] [Professional]	rock , classical , rock&roll	CD 64k ●	★★★★★
8.	Zen Rock A Celebration of Rock Music featuring all the greatest hits from Guns N' Roses, Metallica, Red Hot Chili Peppers, U2, Queen, Scorpions, Poison, Pearl Jam, Ozzy Osbourne, Bon Jovi, Aerosmith, Creed, R.E.M., Raidohead, Nirvana, AC/DC and more!	rock , hard rock , classic rock	CD 64k ●	★★★★★
9.	New Rock - THE MACHINE 128k - Shinedown Alter Bridge Daughtry Seether Linkin Park Korn Velvet Revolver	rock , alternative , metal	CD 128k	★★★★☆
10.	ON THE ROCKS Clearly Classic! AC/DC, Aerosmith, Beatles, Boston, CCR, Clapton, Doors, Eagles, Elton John, Floyd, GNR, Hendrix, Journey, KISS, Led Zeppelin, Neil Young, Queen, Rush, Rolling Stones, Supertramp, Tom Petty, U2, Van Halen, Who, more!	classic rock , rock , hard rock	CD 64k ●	★★★★★
11.	aTTeNTioN sPaN raDio 1 Voted Live365's Best Jazz Station for 2004 & 2005. An eclectic mix of fusion, funk and rock. Over 650 artists & 7200 tracks! [Professional]	jazz , rock , fusion	CD 64k ●	★★★★☆
12.	Station6070S...60s&70s CLASSIC ROCK...60s&70s...All the treasures that defined an era including, The Beatles, Led Zeppelin, Pink Floyd, The Doors, The Moody Blues, CSNY, and other architects of the 60s and 70s consciousness....	classic rock , rock , oldies	CD 128k	★★★★★
13.	METAL SHOP Kick Ass 80's Metal	metal , rock , hard rock	CD 64k ●	★★★★★
14.	radio for musicheads B [AM/FM] [Professional]	rock , electric blues , singer/songwriter	CD 64k ●	★★★★★
15.	ROCKLAND USA : POP ROCK ▶▶ CHANNEL 2 - NO LOG-INS REQUIRED - IF STATION IS FULL USE CHANNEL 1: live365.com/stations/rocklandusa - The Best At-Work Popular Top 40 Rock Hits. 2005 Broadcaster of the Year and Best Pop Station [Professional]	pop , rock , alternative	CD 64k ●	★★★★★

Showing 1-15 of 639 *Rock* stations

1 | 2 | 3 | 4 | Next

[www.last.fm]

The screenshot displays the last.fm website interface. At the top, a red navigation bar contains the last.fm logo, menu items for Music, Videos, Listen, Events, and Charts, a search bar with a dropdown menu set to 'Music', and links for 'Log in | Sign up' and 'Help'. Below the navigation bar, the main content area is divided into three sections. On the left is a music player for 'Franz Ferdinand' with a progress bar and playback controls. In the center is a large image of the band Franz Ferdinand with the word 'Ulterior' overlaid at the bottom. On the right, there is a 'Share this Station' button, a 'Related stations' section with three radio-tagged links, and a promotional box for the Last.fm mobile app with an image of an iPhone.

last.fm Music Videos Listen Events Charts Music Search Log in | Sign up Help

Now playing Music like Muse

Share this Station

Related stations

- Music tagged rock
- Music tagged alternati...
- Music tagged alternati...

Get Last.fm on your iPhone or iPod Touch

Listen to free streaming radio and more

Buy

Franz Ferdinand

40'

00:14 03:08

From Franz Ferdinand

Ulterior



ARTISTS


CHANNELS

NEWS

ON TV





WIN

MY CONCERT.TV



VIDEO (6)

PHOTOS (5)

- 
APOCALYPTICA
 BETRAYAL
 METALMANIA 2005
ADD TO PLAYLIST +
- 
APOCALYPTICA
 INTRODUCTION
 INTERVIEW
 CONCERT.TV NEWS +
- 
APOCALYPTICA
 ON CELLOS AND METAL
 INTERVIEW
 CONCERT.TV NEWS +
- 
APOCALYPTICA
 ON 'WORLDS COLLIDE' AND SUCCESS
 INTERVIEW
 CONCERT.TV NEWS +

ARTIST NAME

SEARCH

CHANNELS



SELECT ▼

APOCALYPTICA

There aren't many rock bands on this planet that wouldn't like to be regarded as unique. Of course, in these times of nostalgia and endlessly recycled ideas, very few actually deserve that description. Finland's Apocalyptica deserve it more than most. Utterly original in every respect, these classically-trained merchants of genre-shattering metal sound quite unlike anything or anyone else in musical history.

Formed in the early '90s Apocalyptica chose an entirely individual path from day one. A few years earlier, while students at the Sibelius Academy in Helsinki, these young virtuosos joined a cello band that specialised in playing music by everyone from Bach to Hendrix.

MEMBER LOGIN

USERNAME



NOT A MEMBER?

LOGIN

Advertisement





Search bar containing 'pat metheny', dropdown menu set to 'Videos', 'Search' button, 'advanced' link, and 'Upload' button

"pat metheny" video results 1 - 20 of about 1,830

Filter tabs: Videos (selected), Channels. Sort by: Relevance. Uploaded: Anytime. Display: Grid view

Also try: pat metheny group, last train home, herbie hancock, minuano, pat matheny, パットメセニー, letter from home, don't know why

Only show partner videos



Pat Metheny Live In Switzerland

Awesome! ... Pat Metheny Live ...

Added: 1 year ago
From: Patrick88881234
Views: 328,442
★★★★★
07:46
More in Music



Pat Metheny - Don't Know Why

Pat metheny performing "Don't Know Why" by Norah Jones. ... pat metheny don't know why jazz guitar ...

Added: 1 year ago
From: clasguitar5
Views: 267,981
★★★★★
03:09
More in Music



pat metheny - last train home

youtube... more views mean more people turned onto your music... and more record buyers! ... last train home pat pmg metheny ...

Added: 1 year ago
From: pleasepatlea...
Views: 207,571
★★★★★
05:13
More in Music



Pat Metheny Group - Have you heard

We live here - Japan 1995 ... Pat Metheny Group Have you heard We live here jazz guitar ...

Added: 1 year ago
From: notelacreas
Views: 133,878
★★★★★
06:11
More in Music



PAT METHENY - LETTER FROM HOME

one of my favourite metheny tracks ... pat metheny group lyle mays letter from home more travels jazz guitar piano ...


Added: 1 year ago
From: 777i2i
Views: 191,742
★★★★★
02:22
More in Music

Introduction

Plenty of music – Consequences

- Traditional ways of finding music are no longer sufficient
 - we cannot browse through all the music we would potentially like
 - record companies and radio stations are no longer critical gatekeepers in music distribution
- Relying just on popularity statistics is not effective
 - music tastes are so different that averaging opinions does not produce precise information for an individual
 - "UK Singles Chart" etc. sales statistics work badly as a guide for the consumer

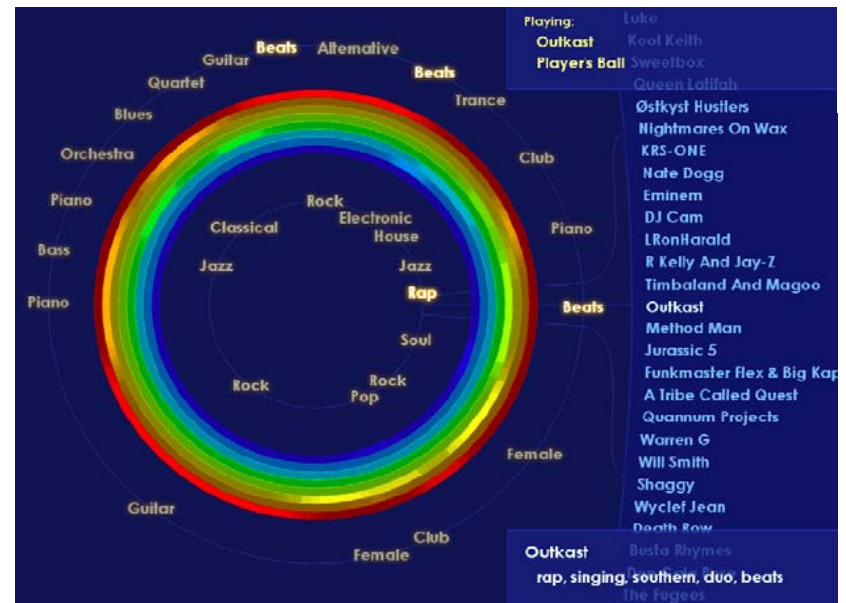
Music retrieval user interfaces



- Non-speech audio is difficult to describe with words
→ *expressing a music query is hard*

- Query mechanisms
 - "find similar artists"
 - query by example
 - browse by similarity (see [Figure](#))
 - query by humming or tapping
 - tempo
 - lyrics
 - music categories (genre, mood, tags)

- Tailor UI to match user abilities



Music rainbow [Pampalk-Goto-2006]

- Also, presenting (multiple) retrieval results is challenging
 - thumbnail extraction (chorus detection etc.)
 - spatialisation for simultaneous presentation [Stewart-dafx-2008]



Introduction

Modern ways of searching music

Two complementary approaches:

1. Collaborative filtering

- based on (users x items) matrix \approx "music likings metadata"
- recommend music by comparing *user profiles* and predicting likings for new pieces
- measure similarity of music pieces (acoustics, usage, etc.) based on *piece profiles*



2. Content-based retrieval

- either based on automatic signal analysis or collaborative tagging by users
- Old ways of discovering music are still relevant too (though ineffective)
 - talking to friends, relying on experts (e.g. listening to FM radio you like)


Introduction

Collaborative filtering vs. audio-based MIR

- Collaborative filtering (CF) has been very successful and is widely used but does not solve it all
 - CF does not allow separating the various dimensions of music similarity, but these are all mixed in the piece profile
 - CF alone is not able to deal with items that are new or do not have many listeners
- Audio-based MIR addresses the above problems
 - *enables “truly musical queries”* with specific musical criteria, such as requesting pieces with certain vocal characteristics or slow tempo
 - *can be employed even on media libraries* that do not have any audience of listeners
 - also, *enables musically interesting listening UIs* that encourage music understanding
- On the other hand, audio-based MIR alone cannot measure aspects like *quality, usage, or culture*
 - The two approaches are complementary
- Hybrid systems: for example novelty + relevance

Content-based retrieval system

- Overview of canonical content-based MIR system [Casey et al. 2008]



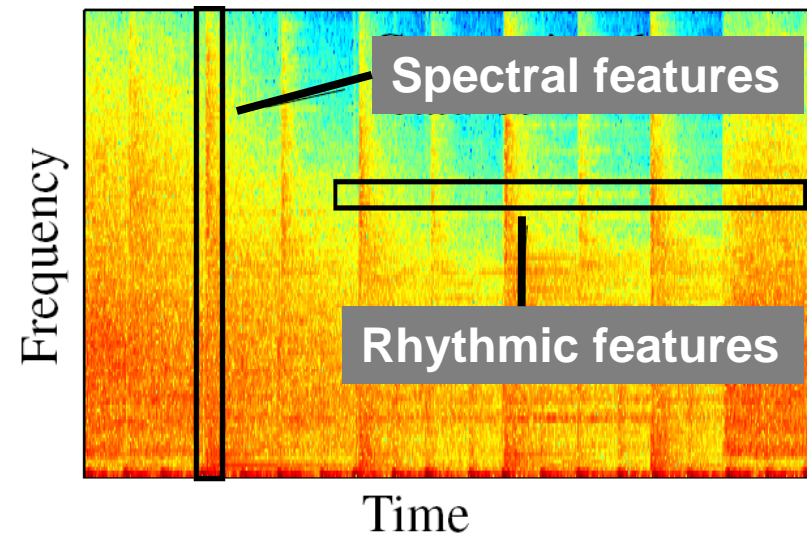
MIR Tasks

- Examples of MIR tasks and their "specificities" (high / middle / low) [Casey et al. 2008]
- On this course, we will study techniques for many of these tasks


Use Case	Specificity	Description
Music Identification	H	Identify a compact disk, provide metadata about an unknown track, mobile music information retrieval: e.g. <i>shazam.com</i>
Plagiarism detection	H	Identify mis-attribution of musical performances, mis-appropriation of music intellectual property.
Copyright monitoring	H	Monitor music broadcast for copyright infringement or royalty collection
Versions	H/M	Remixes, live vs. studio recordings, cover songs. Used for database normalization and near-duplicate results elimination
Melody	H/M	Find works containing a melodic fragment
Identical Work / Title	M	Retrieve performances of same opus number or song title
Performer	M	Find music by a specific artist
Sounds like	M	Find music that sounds like a given recording
Performance Alignment	M	Mapping one performance onto another independent of tempo and repetition structure
Composer	M	Find works by one composer
Recommendation	M/L	Find music that matches the user's personal profile
Mood	L	Find music using emotional concepts: <i>Joy, Energetic, Melancholy, Relaxing</i>
Style / Genre	L	Find music that belongs to a generic category: <i>Jazz, Funk, Female Vocal</i>
Instrument(s)	L	Find works with same instrumentation
Music-Speech	L	Radio broadcast segmentation, Music archives cataloguing

Music similarity

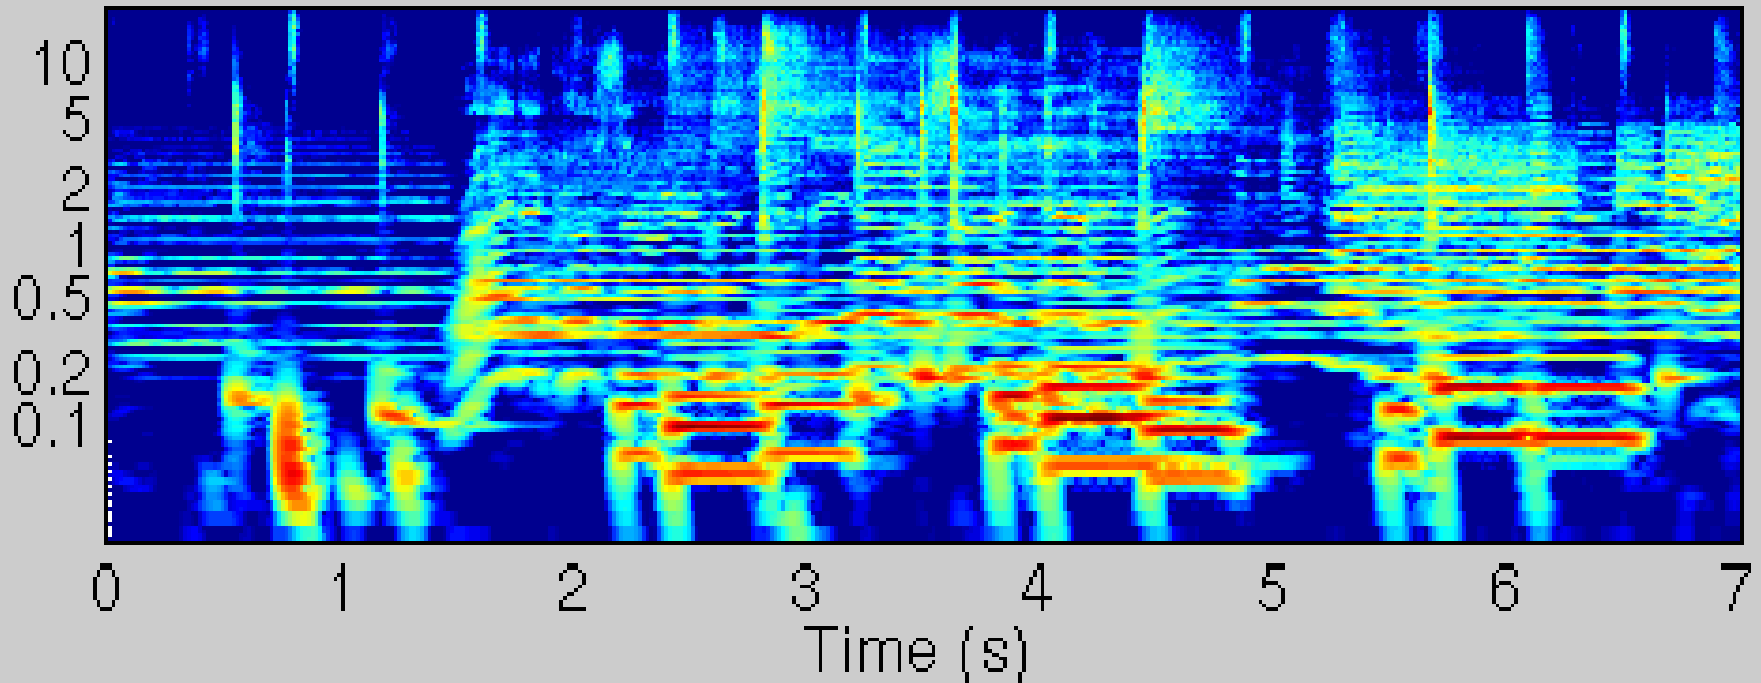
- Music similarity estimation enables query by example and browsing by artist similarity
- Requires defining a similarity measure or a distance metric btw signals
- Widely used acoustic features
 - Mel-frequency cepstral coefficients (MFCCs) → [timbre/instrumentation](#)
 - chroma [Bartsch-2001] → [harmonic content](#)
 - rhythmogram (or, fluctuation patterns) [Dixon-2003, Jensen-2007] → [rhythm](#)
 - different features model different (and nearly orthogonal) "perceptual dimensions" of audio
- For a more comprehensive list, see e.g. [Peeters-2004]




Amplitude



Frequency (kHz)





Features define "similarity"

- "Similarity" as such is not well-defined



- For example: Is *Bohemian rhapsody* by Queen more similar to



a) *Bohemian rhapsody* by London Symphony Orchestra, or



b) *A kind of magic* by Queen?

- The riddle is solved by choosing the acoustic features
 - chroma → a) is more similar (composition)
 - MFCCs → b) is more similar (instrumentation)
- User may wish to specify the features when doing query by example
- Narrowing down "a perfect piece" by using multiple examples
 - enabled when huge amounts of music is available

Similarity measures between audio clips

- Similarity between two audio signals is *typically* calculated based on the statistics of extracted features
- "Bag of features" approach: collapse all temporal structure in data
- Traditional distance measures utilize means and covariances of the features. For example *Mahalanobis distance* between clips f and g :

$$D_{Mah}(f, g) = (\mu_f - \mu_g)^T \Sigma^{-1} (\mu_f - \mu_g)$$

← μ_g mean of features in clip g
· Σ covariance of all features

Music similarity: recent evaluation results


- Results from MIREX 2007 evaluation [[Downie-ismir-2007](#)]
 - organised by IMIRSEL Group at University of Illinois

- Task
 - repeat query-by-example 100 times
 - return 5 closest from among 7000 songs
 - human listeners rated the relevance of the returned pieces

Audio Music Similarity		
Rank	Participant	Avg. Fine Score
1	Pohle & Schnitzer	0.568
2	Tzanetakis, G.	0.554
3	Barrington, Turnball, Torres & Lanskrict	0.541
4	Bastuck, C. (1)	0.539
5	Lidy, Rauber, Pertusa & Iñesta (1)	0.519
6	Mandel & Ellis	0.512
7	Lidy, Rauber, Pertusa & Iñesta (2)	0.491
8	Bastuck, C. (2)	0.446
9	Bastuck, C. (3)	0.439
10	Bosteels & Kerre (1)	0.412
11	Paradzinets & Chen	0.377
12	Bosteels & Kerre (2)*	0.178

pdf-based similarity measures

- Recently studied approach
- Idea: measure similarity by calculating distance between the probability density functions (pdfs) of features
 - more flexible and accurate than merely using the means and covariances of features
 - no need to go through the feature vectors (after the models have been trained)
- For example Euclidean distance, Kullback-Leibler (KL) divergence etc.



Using temporal sequences for similarity

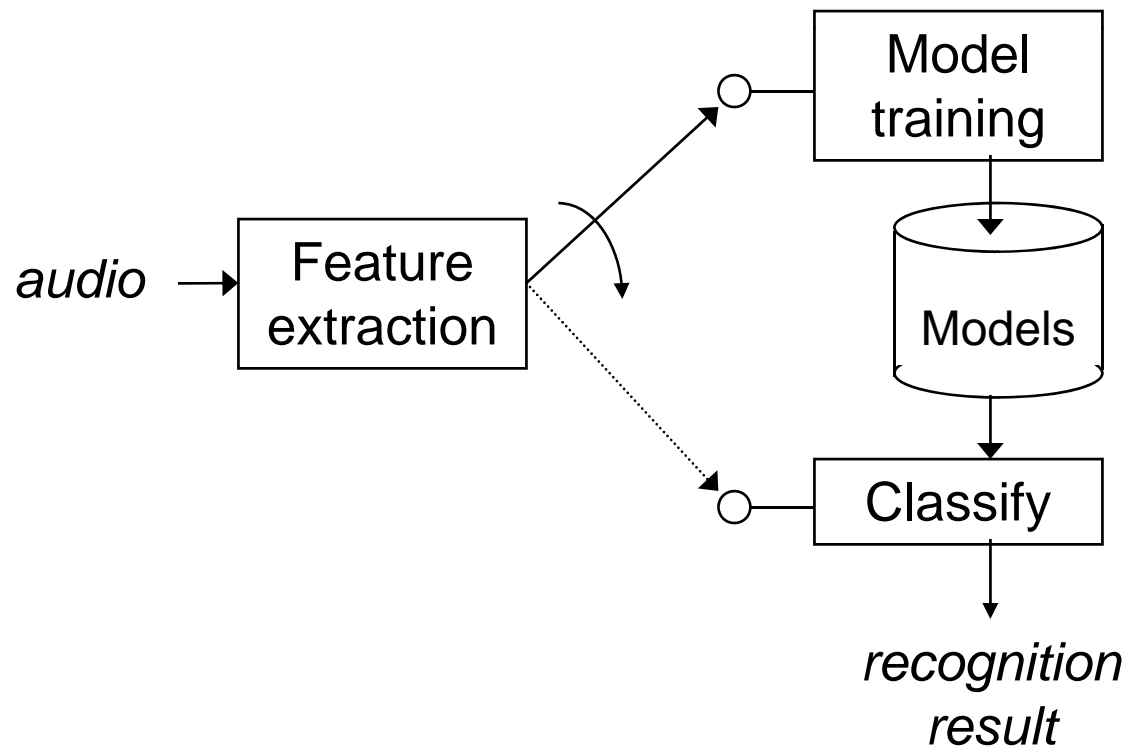
- Above methods collapse the time structure of feature sequences
- Using temporal sequences for similarity requires time-alignment
 - non-trivial: tempo differences, different numbers of sectional parts, etc.
- *Beat-synchronous feature extraction* reconciles for tempo differences
 - = track the beat of each song and extract one feature vector per inter-beat-interval
- Used previously in
 - cover song detection using beat-synchronised chroma features (e.g. [Ellis-2006])
 - analysis of the sectional form (verse, chorus,...) of a piece (e.g. [Paulus-dafx-2008])

MIREX 2007
[Downie-ismir-2007]

Audio Cover Song Identification		
Rank	Participant	Avg. Precision
1	Serrà & Gómez	0.521
2	Ellis & Cotton	0.330
3	Bello, J.	0.267
4	Jensen, Ellis, Christensen & Jensen	0.238
5	Lee, K. (1)	0.130
6	Lee, K. (2)	0.086
7	Kim & Perelstein	0.061
8	IMIRSEL	0.017

Music classification

- Music can be classified according to genre, mood, etc.
- Classical train/test *supervised classification* scenario:



Classification and identification tasks

- Classify music into categories
 - genre: rock, hip hop, jazz, classical,... (here 10)
 - mood: aggressive, passionate, humorous, cheerful,...(5)
 - artist identification (here 102 artists)
 - classical composer identification (here 11)
- Train-test setup
- Results from MIREX 2007 [[Downie-ismir-2007](#)]
 - organised by IMIRSEL Group at University of Illinois

Audio Artist Identification		
Rank	Participant	Avg. Raw Accuracy
1	IMIRSEL (svm)	48.14%
2	Mandel & Ellis (spec)	47.16%
3	Mandel & Ellis	40.46%
4	Lidy, Rauber, Pertusa & Iñesta	38.76%
5	Tzanetakis, G.	36.70%
6	IMIRSEL (knn)	35.29%
7	Lee, K.	9.71%

Audio Genre Classification		
Rank	Participant	Avg. Raw Accuracy
1	IMIRSEL (svm)	68.29%
2	Lidy, Rauber, Pertusa & Iñesta	66.71%
3	Mandel & Ellis	66.60%
4	Mandel & Ellis (spec)	65.50%
5	Tzanetakis, G.	65.34%
6	Guaus & Herrera	62.89%
7	IMIRSEL (knn)	54.87%

Audio Mood Classification		
Rank	Participant	Avg. Raw Accuracy
1	Tzanetakis, G.	61.50%
2	Laurier, C.	60.50%
3	Lidy, Rauber, Pertusa & Iñesta	59.67%
4	Mandel & Ellis	57.83%
5	Mandel & Ellis (spec)	55.83%
	IMIRSEL (svm)	55.83%
7	Lee, K. (1)	49.83%
8	IMIRSEL (knn)	47.17%
9	Lee, K. (2)	25.67%

Audio Classical Composer Identification		
Rank	Participant	Avg. Raw Accuracy
1	IMIRSEL (svm)	53.72%
2	Mandel & Ellis (spec)	52.02%
3	IMIRSEL (knn)	48.38%
4	Mandel & Ellis	47.84%
5	Lidy, Rauber, Pertusa & Iñesta	47.26%
6	Tzanetakis, G	44.59%
7	Lee, K. Klapuri - DAF08	29.09%

Transcription / separation approaches

- = Approaches where some musically meaningful part of the signal, such as the melody line, is extracted and analyzed

Music transcription



- Transcription of
 - melody [Goto; Paiva; Ellis-Poliner; Dressler; Ryyänen]
 - bass line [Goto; Hainsworth; Ryyänen]
 - drums [Paulus; Yoshii; FitzGerald]
 - chords [Sheh-Ellis; Bello-Pickens; Harte; Lee]
 - key/mode (e.g. [Gomez-phd])
 - tempo, meter (e.g. [Dixon; Klapuri])
 - instrument recognition
- Separation
 - vocals
 - drums
- MIREX 2007: polyphonic transcription task → [Downie-ismir-2007]

Multi F0 Note Tracking		
Rank	Participant	Avg. F-Measure
1	Ryyänen & Klapuri (2)	0.614
2	Vincent, Bertin & Badeau (4)	0.527
3	Poliner & Ellis (2)	0.485
4	Vincent, Bertin & Badeau (3)	0.453
5	Pertusa & Iñesta (2)	0.408
6	Egashira, Kameoka & Sagayama (4)	0.268
7	Egashira, Kameoka & Sagayama (3)	0.246
8	Pertusa & Iñesta (3)	0.219
9	Emiya, Badeau & David (2)	0.202
10	Cont, A. (4)	0.093
11	Cont, A. (3)	0.087

Music transcription – examples

- Example transcription [Ryynänen-cmj-2008]
 - melody, bass line and chords

input (RWC-P)



transcription





mixdown



The image displays a musical score for the example transcription. It consists of two systems of staves. The first system shows the input melody in the treble clef and the transcription of the bass line in the bass clef. Above the first staff, the chords Em, Am, C, and F are indicated. The second system shows the mixdown, with the melody in the treble clef and the bass line in the bass clef. Above the second staff, the chords Am, E, Am, C, and Dm are indicated. The notation includes various note values, rests, and accidentals.

- Drums [Paulus-ismir-2007]

- RWC-MDB-P-2001 No. 14 
- RWC-MDB-P-2001 No. 22 
- traditional way of writing down drum notation (not automatically generated):

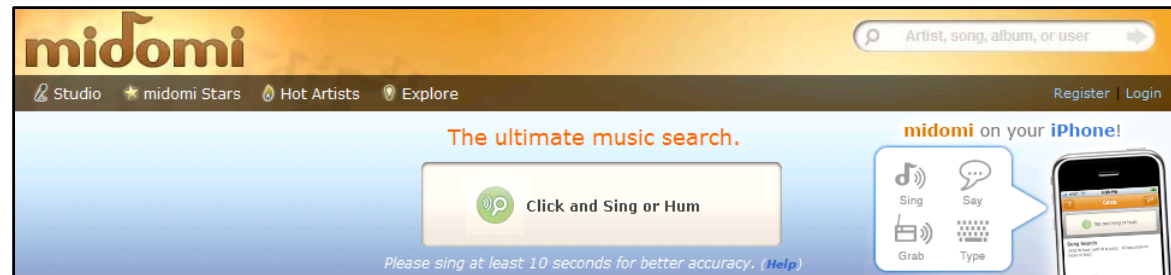
The image shows a traditional drum notation for RWC-MDB-P-2001 No. 14. It is written on a single staff with a 4/4 time signature. The notation uses 'x' marks to represent drum hits, with some marks grouped together to indicate specific drum patterns or accents.

Query by humming

- Research
 - QBH of MIDI data [Lemström-phd-00], [Meek-ismir-02], [Jang-00,06], **MIREX-2007**
 - QBH of audio only very recently [Duda-ismir-2007], [Ryynänen-icassp-2008]

- QBH services:

www.midomi.com



www.musicline.de



www.musipedia.org
(MIDI)















Query by humming at MIREX

- Results of MIREX-2008 evaluation [[Downie-ismir-2008](#)]

Query by Singing/Humming (Jang's Collection)		
Rank	Participant	MRR
1	Ryynänen & Klapuri	0.93
	Wang, L. (1)	0.93
3	Wu & Li (1)	0.92
4	Wu & Li (2)	0.91
5	Jang & Lee (3)	0.82
6	Jang & Lee (1)	0.81
7	Jang & Lee (4)	0.72
8	Jang & Lee (2)	0.69

Query by Singing/Humming (ThinkIT Collection)		
Rank	Participant	MRR
1	Wang, L. (3)	0.94
	Wu & Li (1)	0.94
3	Wang, L. (2)	0.92
	Wu & Li (2)	0.92
5	Ryynänen & Klapuri	0.89
6	Jang & Lee (4)	0.59
7	Jang & Lee (2)	0.53
8	Jang & Lee (1)	0.32
	Jang & Lee (3)	0.32

Query by humming of audio

- Example method [Ryynänen-icassp-2008]
 - preprocessing: extract melodies automatically from music pieces
 - transcribe the query
 - match by Euclidean distance between the two melodic contours (allow time scaling)
 - efficient indexing using locality sensitive hashing
- Demonstrations
 - Example A
 - query  retrieval results #1  #2  #3 
 - Example B
 - query  retrieval results #1  #2  #3 
 - Example C
 - query  retrieval results #1  #2  #3 

Music fingerprinting (identification)

- Acoustic fingerprint is a condensed digital summary that can be used to identify an audio sample
- Identify songs or records for example
- Shazam: an acoustic fingerprint-based service that allows for songs to be identified via cell phone. Figure below [Wang-2003]




Fig. 1A - Spectrogram





Fig. 1C - Combinatorial Hash Generation

Indexing techniques

- Locality sensitive hashing (LSH)
 - computationally efficient indexing technique to searching nearest neighbors
 - 1) in large databases and 2) in high-dimensional feature spaces [[Datar-2004](#)]
 - idea: project data points on random lines and subdivide the lines into hash buckets



Toolboxes

- See "Tools we use" page edited by Paul Lamere
<http://www.music-ir.org/evaluation/tools.html>

Some highlights

- Marsyas (Music analysis, retrieval, and synthesis for audio signals):
Open Source C++ tools. <http://marsyas.sourceforge.net/>
- CLAM (C++ library for audio and music): For research and application
development. <http://www.iaa.upf.es/mtg/clam/>
- Jyväskylä MIR Toolbox:
<https://www.jyu.fi/hum/laitokset/musiikki/en/research/coe/materials/mirtoolbox/mirtoolbox>
- M2K (Music to knowledge): Framework developed mainly for
evaluating MIR systems. <http://music-ir.org/evaluation/m2k/>

Seminar topics (choose yours!)

The order can be changed. There are 6+6 weeks in periods I-II.

1. Introduction to MIR (Anssi Klapuri) 8.9.2009
2. **User interfaces for MIR** →
3. **Collaborative filtering and Last.fm** → Annamaria Mesaros
4. **Acoustic features for content-based MIR** → Toni Heittola
5. **Music similarity and query by example** → Petri Hirvonen
6. Music classification (genre, mood, tag, etc.)
7. Query by humming
8. Tempo estimation and rhythm analysis
9. Cover song identification
10. **Music fingerprinting and Shazam system** → Katariina Mahkonen
11. Indexing by locality sensitive hashing (→ Anssi Klapuri?)
12. Vocals and lyrics analysis