Music Recommendation Tutorial

Òscar Celma
Paul Lamere

ISMIR 2007
September 23, 2007

introduction

- Speaker Introductions
  - Òscar Celma – Music Technology Group
  - Paul Lamere – Sun Labs

- Goals for the tutorial

outline

- Introduction
- Formalization of the recommendation problem
- Recommendation algorithms
- Problems with recommenders
- Recommender examples
- Evaluation of recommenders
- Conclusions / Future

introduction:: what's the problem?

Photo by oh snap

- Today
  - iTunes: 6M tracks
  - iTunes: 3B Sales
  - P2P: 15B tracks
  - 53% buy music on line
- Tomorrow
  - **All** music will be on line
  - Billions of tracks
  - Millions more arriving every week
- Finding new, relevant music is hard!

introduction:: why music rec is important?

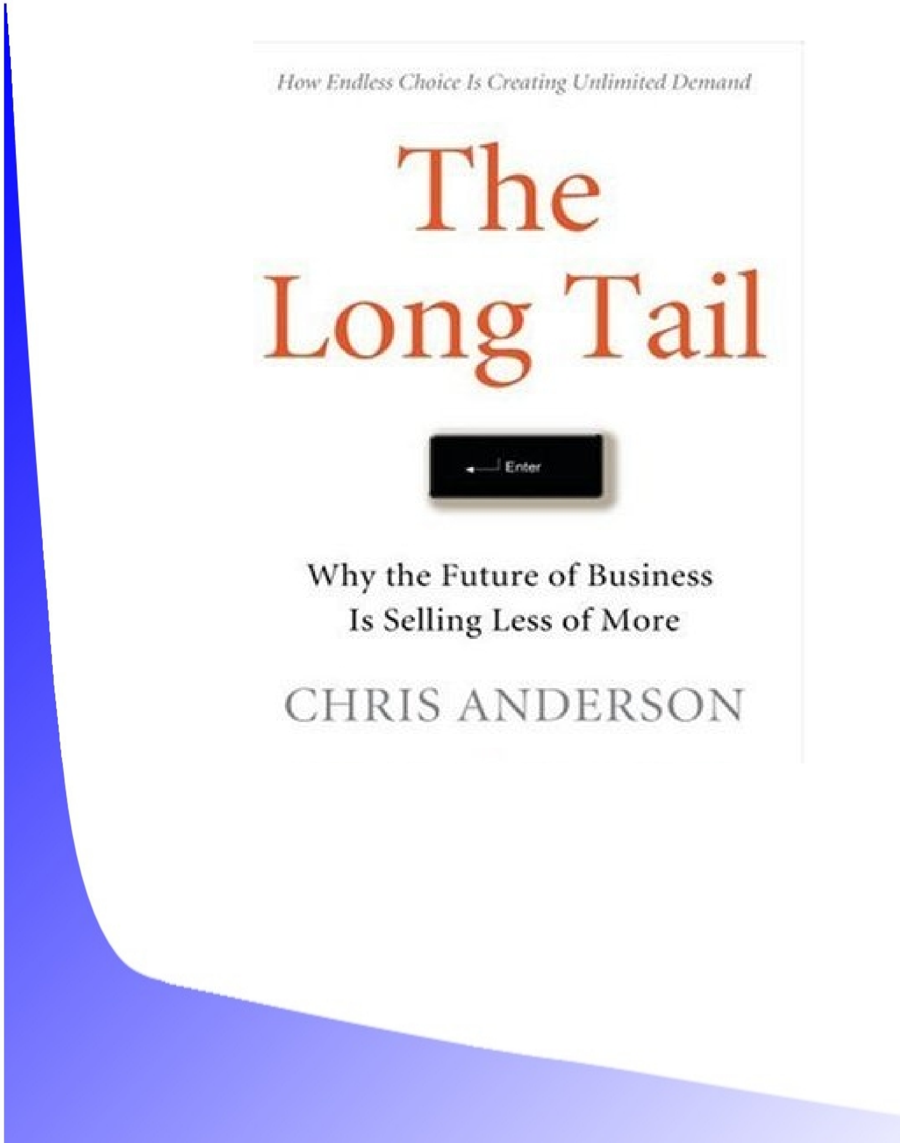

- **Long Tail Rules**
  - Make everything available
  - Help me find it
- **How do we find it?**
  - Experts
  - Friends
  - Content

introduction:: who is recommendation for?

- Different types of users
- Different types of recommendation

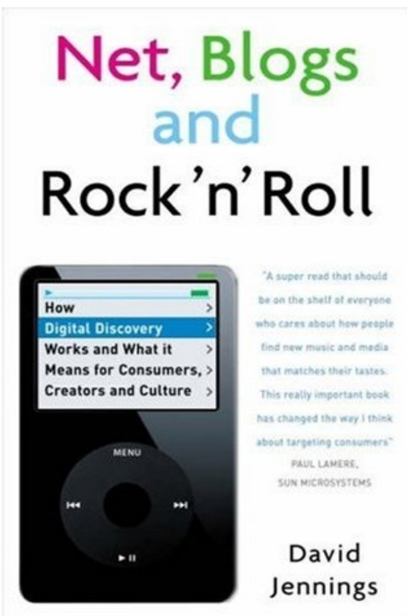

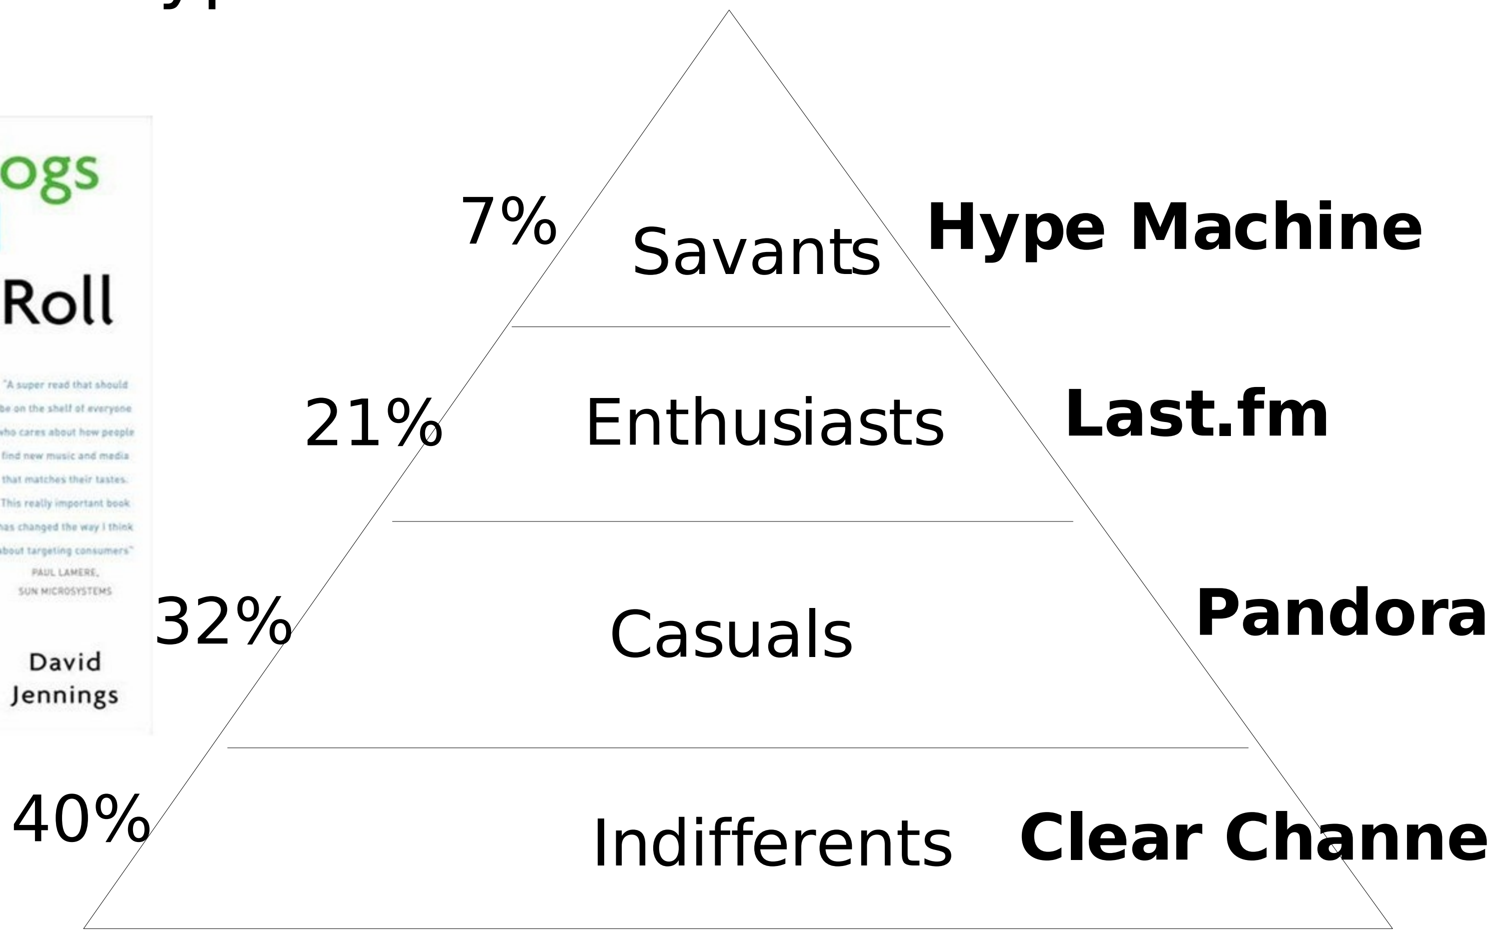

introduction:: savants

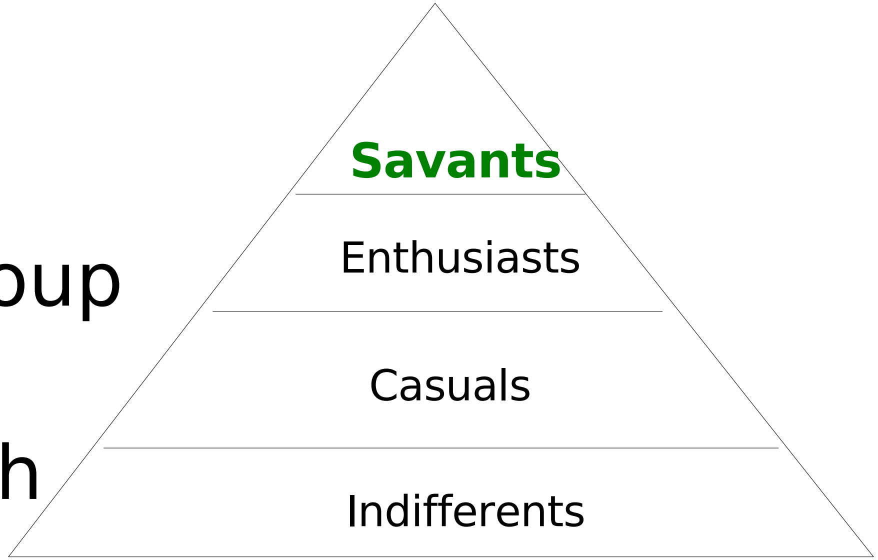

- 7% of the 16-45 age group
- Everything in life seems to be tied up with music
- Example identifying characteristics
  - Being knowledgeable about music is central to “who I am”
  - You reckon you could write better questions for the local bar’s music quiz

introduction:: enthusiasts

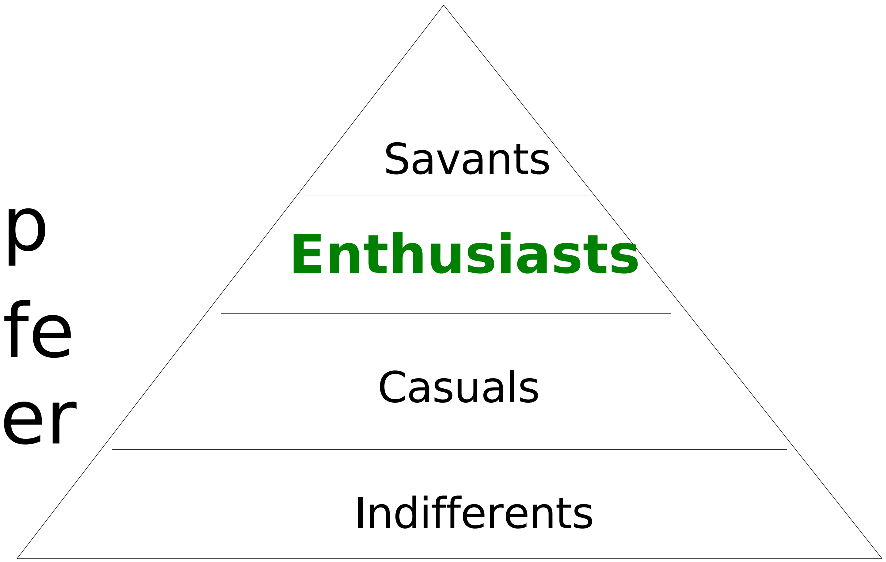

- 21% of 16-45 age group
- Music is a key part of life but is balanced by other interests
- Example identifying characteristics:
  - Believe that the iPod has made the world a better place
  - Get more of a kick from hearing a favorite song on CD than watching its video on television
  - Less “purist” in their musical tastes than savants

introduction:: casuals

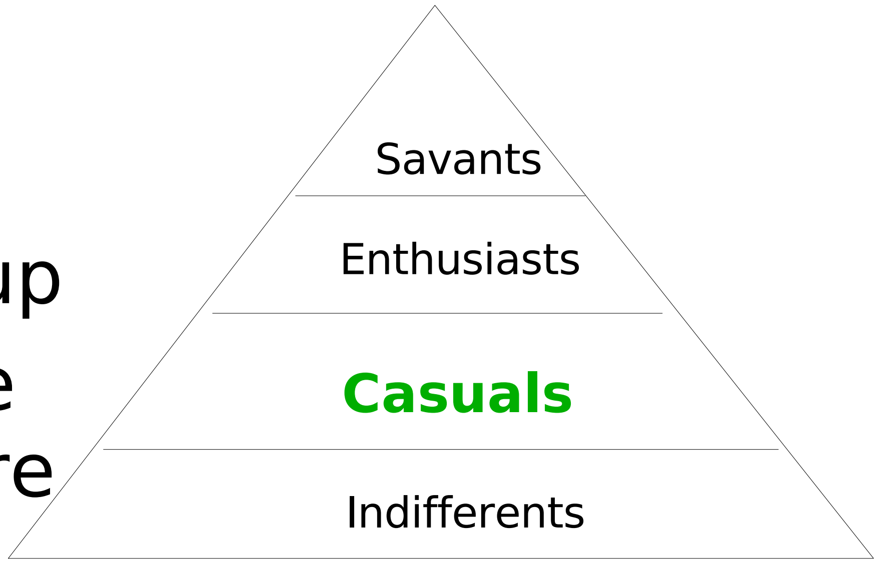

- 32% of 16-45 age group
- Music plays a welcome role, but other thing are far more important
- Example identifying characteristics:
  - Got into Coldplay about the same time that Gwyneth Paltrow did
  - Equally, or more, interested in the lifestyle and fashion trappings of the music world than the music itself

introduction:: indifferents

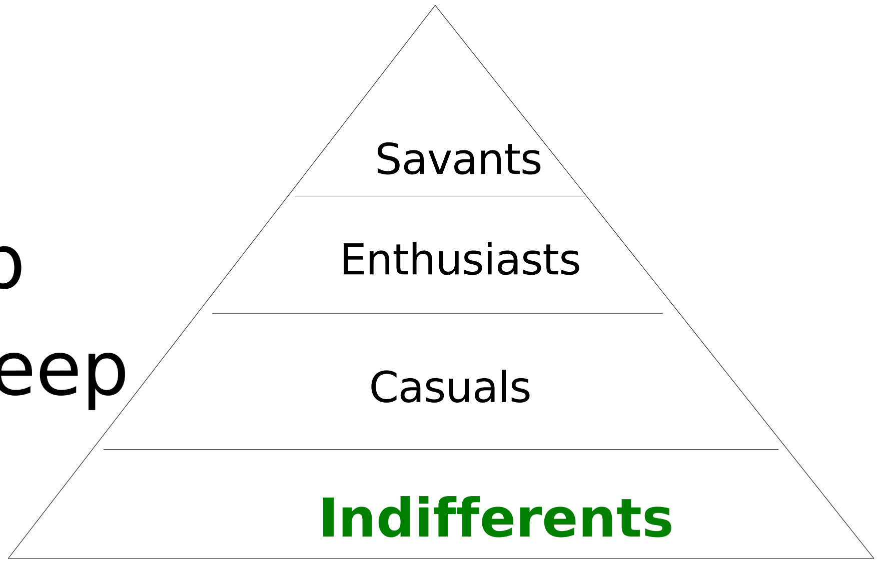

- 40% of 16-45 age group
- Would not lose much sleep if music ceased to exist

- Example identifying characteristics:
  - Most of the songs they hear at parties sound unfamiliar
  - Tend to listen to talk radio or sports rather than music

introduction:: music discovery in the small

- Personal music players:
  - No Experts to guide you
  - No Social network
  - Music discovery is ***random***
  - Shuffle Play doesn't scale
  - Results:
    - iPod **whiplash**
    - The music graveyard
- Study of 5,000 iPod users:
  - 80% of plays in 23% of songs
  - 64% of songs **never** played

introduction:: the value of recommendation

- **Netflix**:
  - 2/3 of movies rented were recommended
  - recommendation is “absolutely critical to retaining users”
- **Google News**:
  - Recommendations generate 38% more click-throughs
- **Amazon**:
  - claims 35% of product sales result from recommendations

introduction:: the value of recommendation

- **Greg Linden (Findory, Amazon):**
  - “recommendations generated a couple orders of magnitude more sales than just showing top sellers”
- **ChoiceStream survey**:
  - 28% would buy more music if they could find more that they liked

introduction:: sources of new music

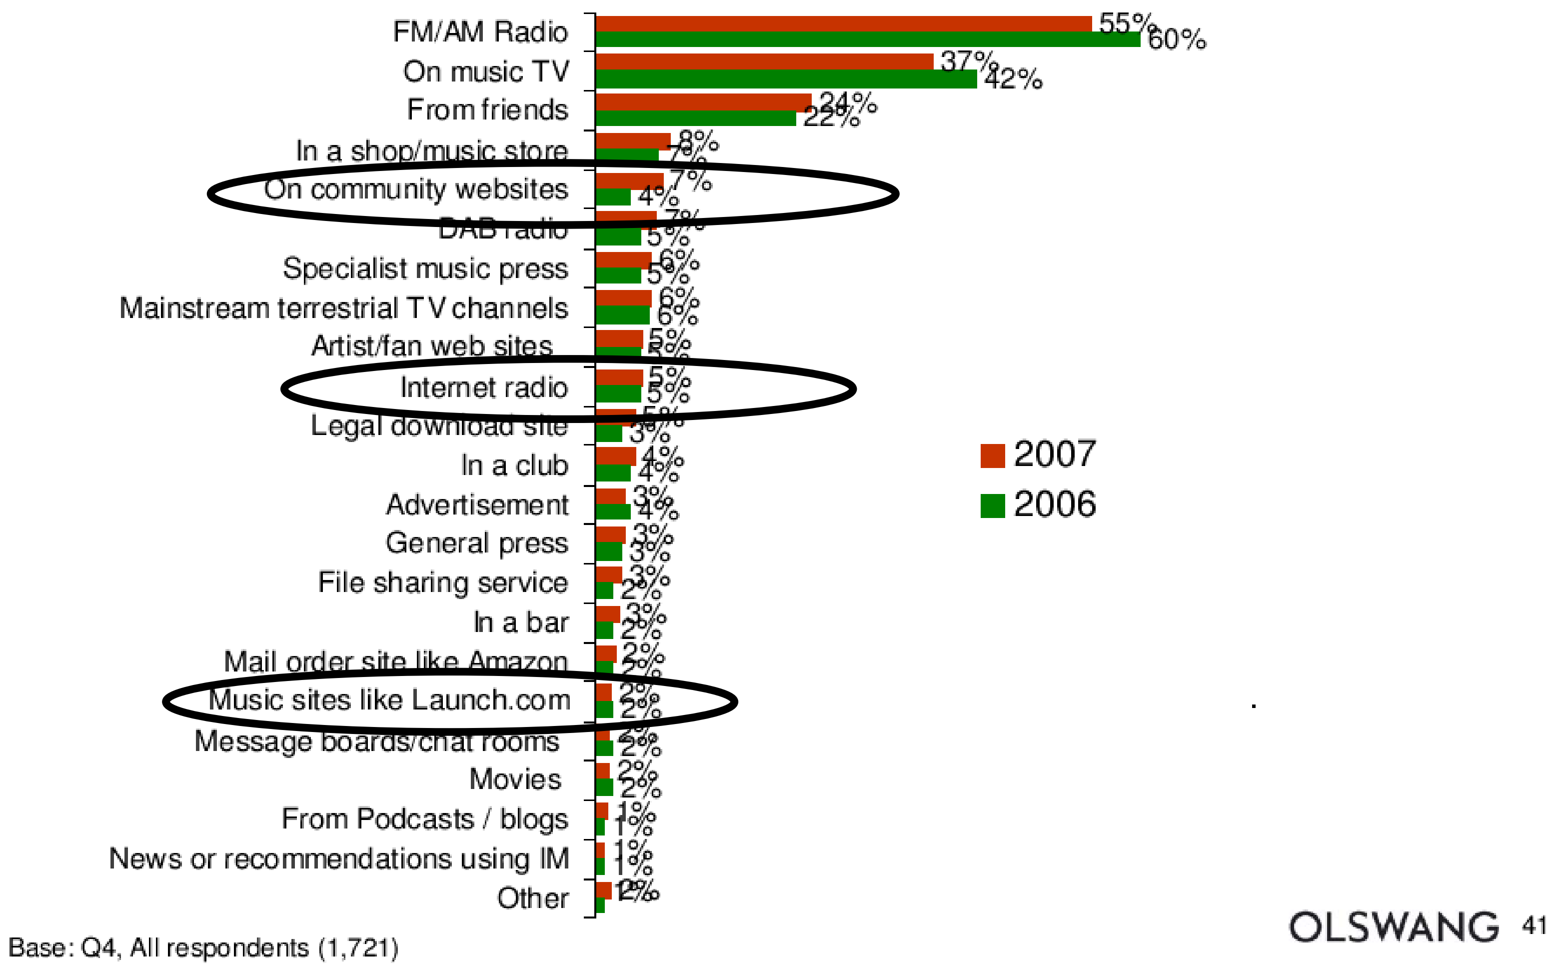

introduction:: sources of new music

- Hype Machine Survey

How do you Discover New Music?

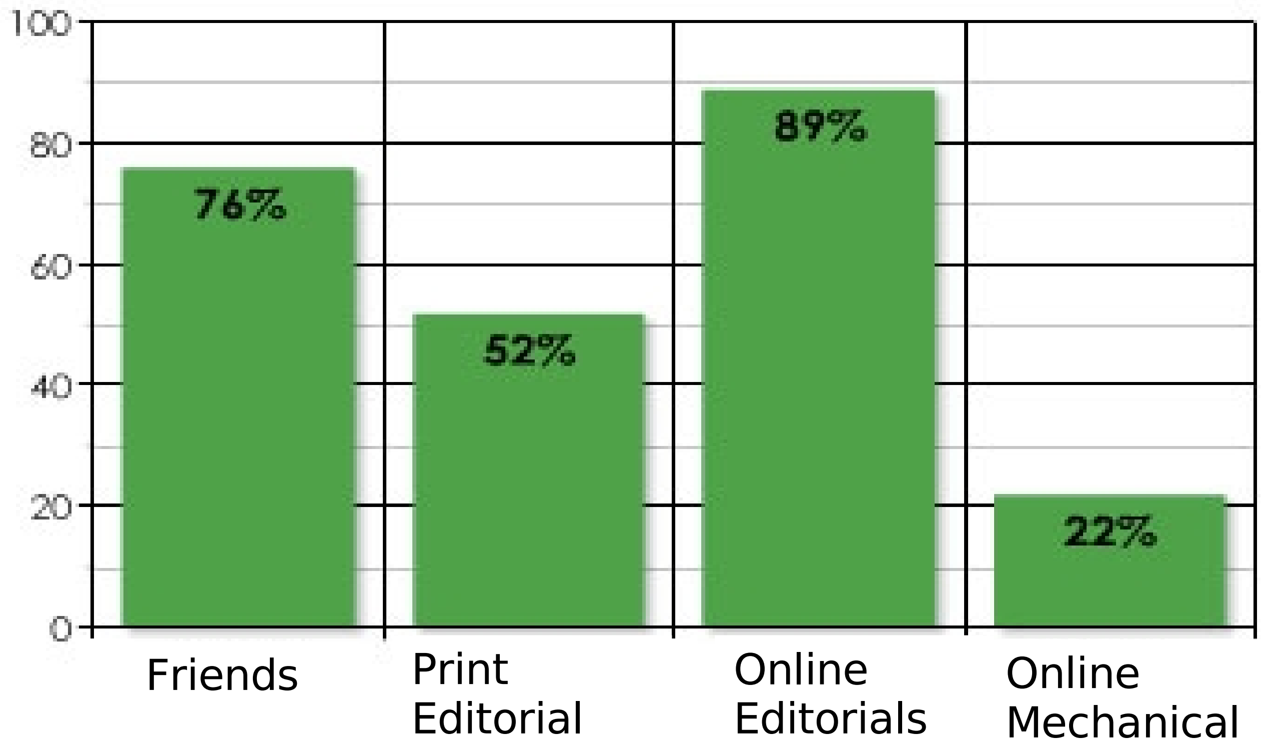

User Survey, July 2007

1430 responses

http://non-standard.net/blog/?p=85

introduction:: commercial interest

- 9Vibe
- Aggrega
- All Music Guide
- AOL Music
- **AudioBaba**
- Audiri
- bandBuzz
- SoundsLikeNow
- Bandwagon
- Buzzwire
- **BMAT**
- Earfeeder
- Goombah
- Grepr
- Guruger
- HooQs
- Idio
- iLike
- inDiscover
- iTunes
- LaunchCast
- **Last.fm**
- Mercora
- MOG
- MusicCodex
- MusicIP
- Musicovery
- Musicmobs
- Musio
- **MyStrands**
- **One Llama**
- **Owl Multimedia**
- Pandora
- QLoud
- RateYourMusic
- SeeqPod
- Slacker
- Soundflavor
- **Spotify**
- The Filter
- UpTo11.net
- ZuKool Music

introduction:: commercial interest

introduction:: recent investment

- Qloud – $1 million
- MOG - $1.4 million
- The Filter - $5 million
- Social.fm - $5 million
- Groove Mobile - $6 million
- Pandora – $6 million
- iLike – $13.3 million
- MyStrands - $25 million
- Slacker - $40 million
- Last.fm - $280 million

introduction:: summary

- Massive Increase in volume of online music
  - Huge shift from physical media to digital media
  - Huge drop in cost to produce new music
- Long Tail Economics
  - Make everything available
  - Help me find it
- Strong commercial interest
- Related Topics
  - Exploration
  - Discovery
  - Playlisting

outline

- Introduction
- **Formalization of the recommendation problem**
- Recommendation algorithms
- Problems with recommenders
- Recommender examples
- Evaluation of recommenders
- Conclusions / Future

formalization:: definition

- Definition of a recommender system
  - Recommender systems are a specific type of **information filtering** technique that attempt to present to the user information **items** (movies, *music*, books, news, web pages) the **user is interested in**. To do this the **user's profile** is compared to some reference characteristics.
  - from:
    - http://en.wikipedia.org/wiki/Recommendation_system

formalization:: definition

- Recommendation as a prediction problem
  - attempt to predict items that a user might be interested in
  - compute *similarity* between objects
    - user-user
    - item-item
  - form *predictions* based on the computed similarities

formalization:: use cases

- Use cases of a recommender system [Herlocker, 2004]
  - Find good items
    - provide a ranked list of items
    - expect some novel items
  - Find all good items
    - coverage
    - low false positive rate
  - Recommend sequence
    - an ordered sequence of items that is pleasing as a whole (i.e playlist generation)
  - Just browsing
  - Find credible recommender

formalization:: use cases

- Use cases of a recommender system [Herlocker, 2004]
  - Improve profile
    - important in recommenders that have a strong community component
  - Express self
    - communicate and interact with other users (messages, forums, weblogs, etc.)
  - Influence others
    - the most negative one
    - influence the community in viewing or purchasing a particular item (e.g labels trying to *promote* artists into the recommender)

formalization:: the whole picture

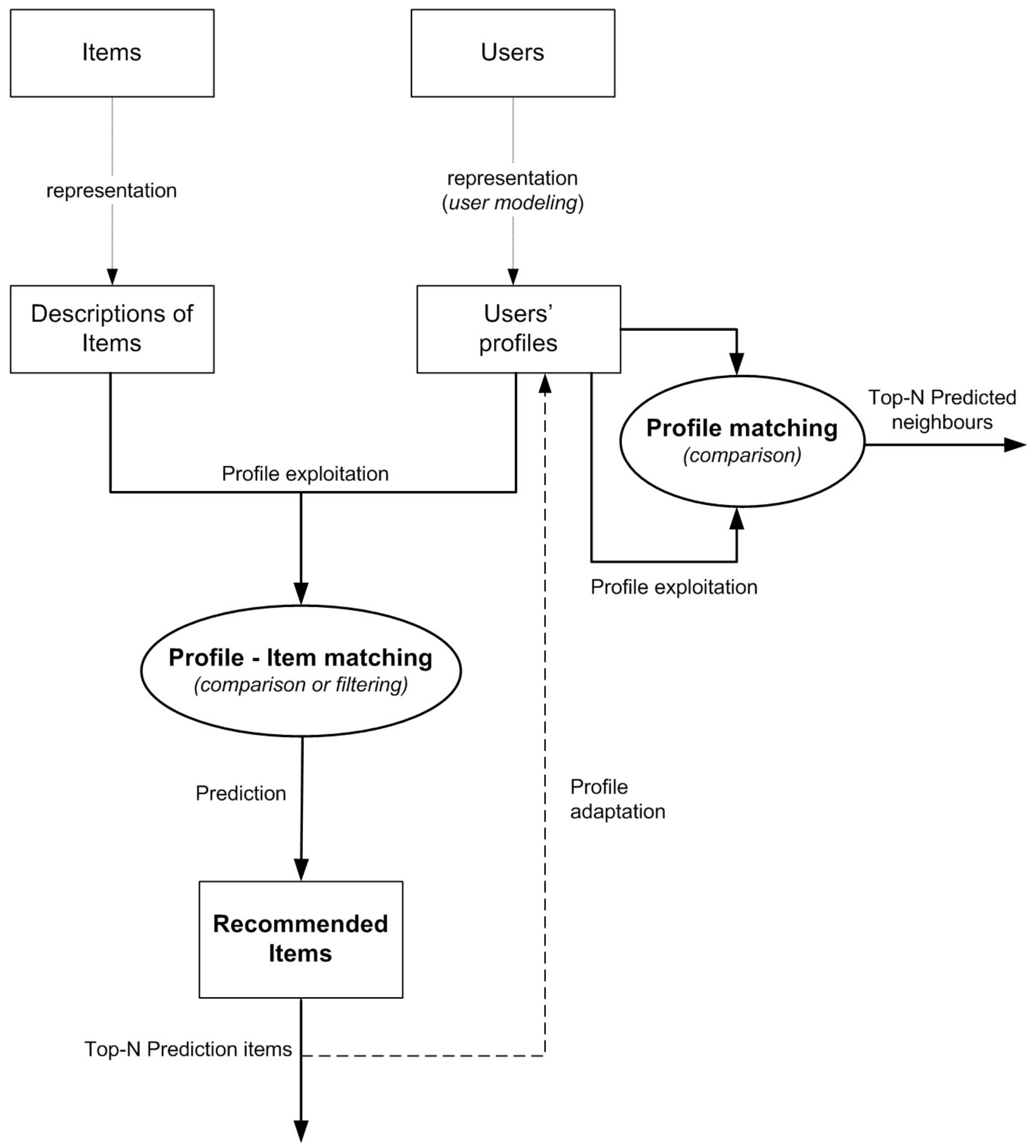

formalization:: describing users & items

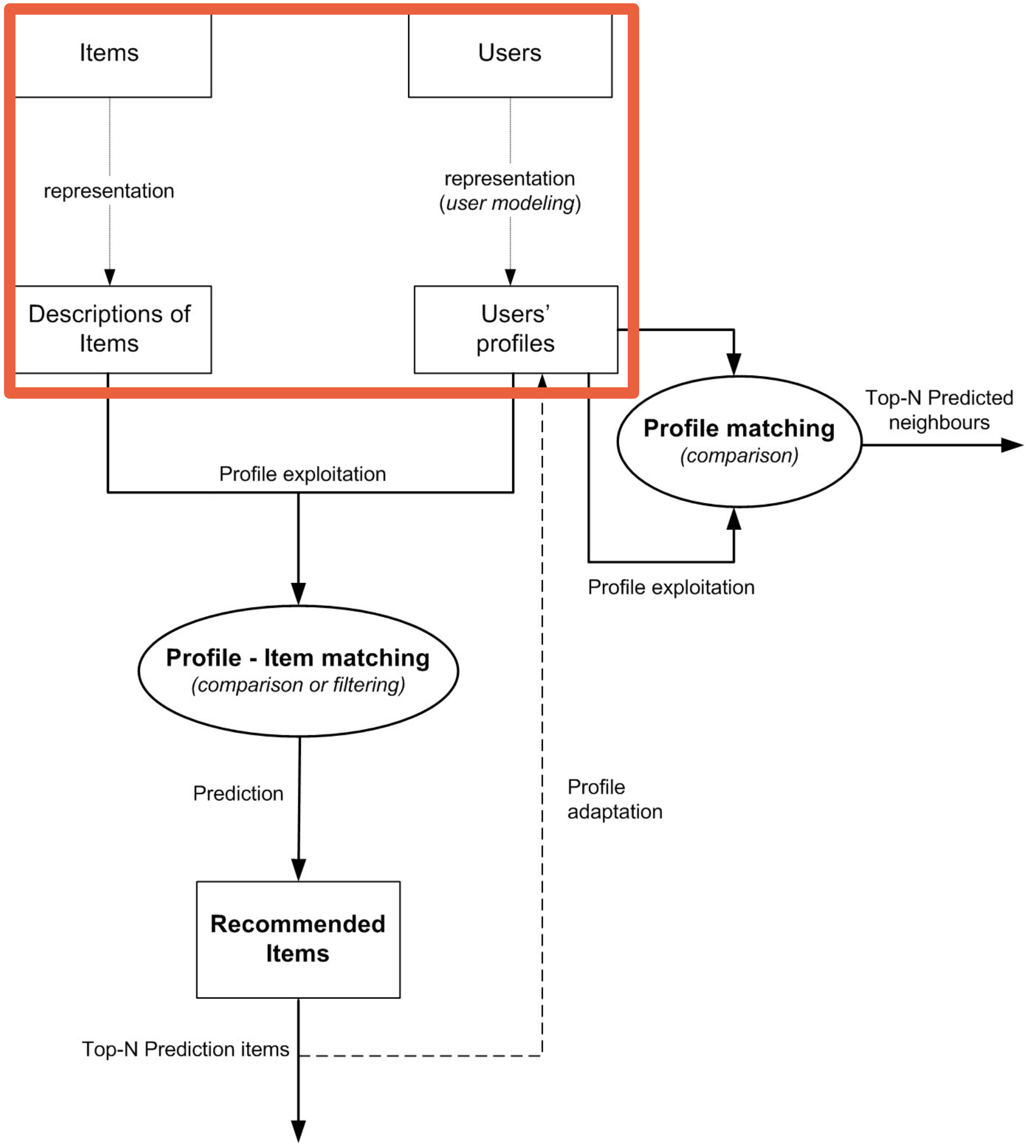

formalization:: describing users & items

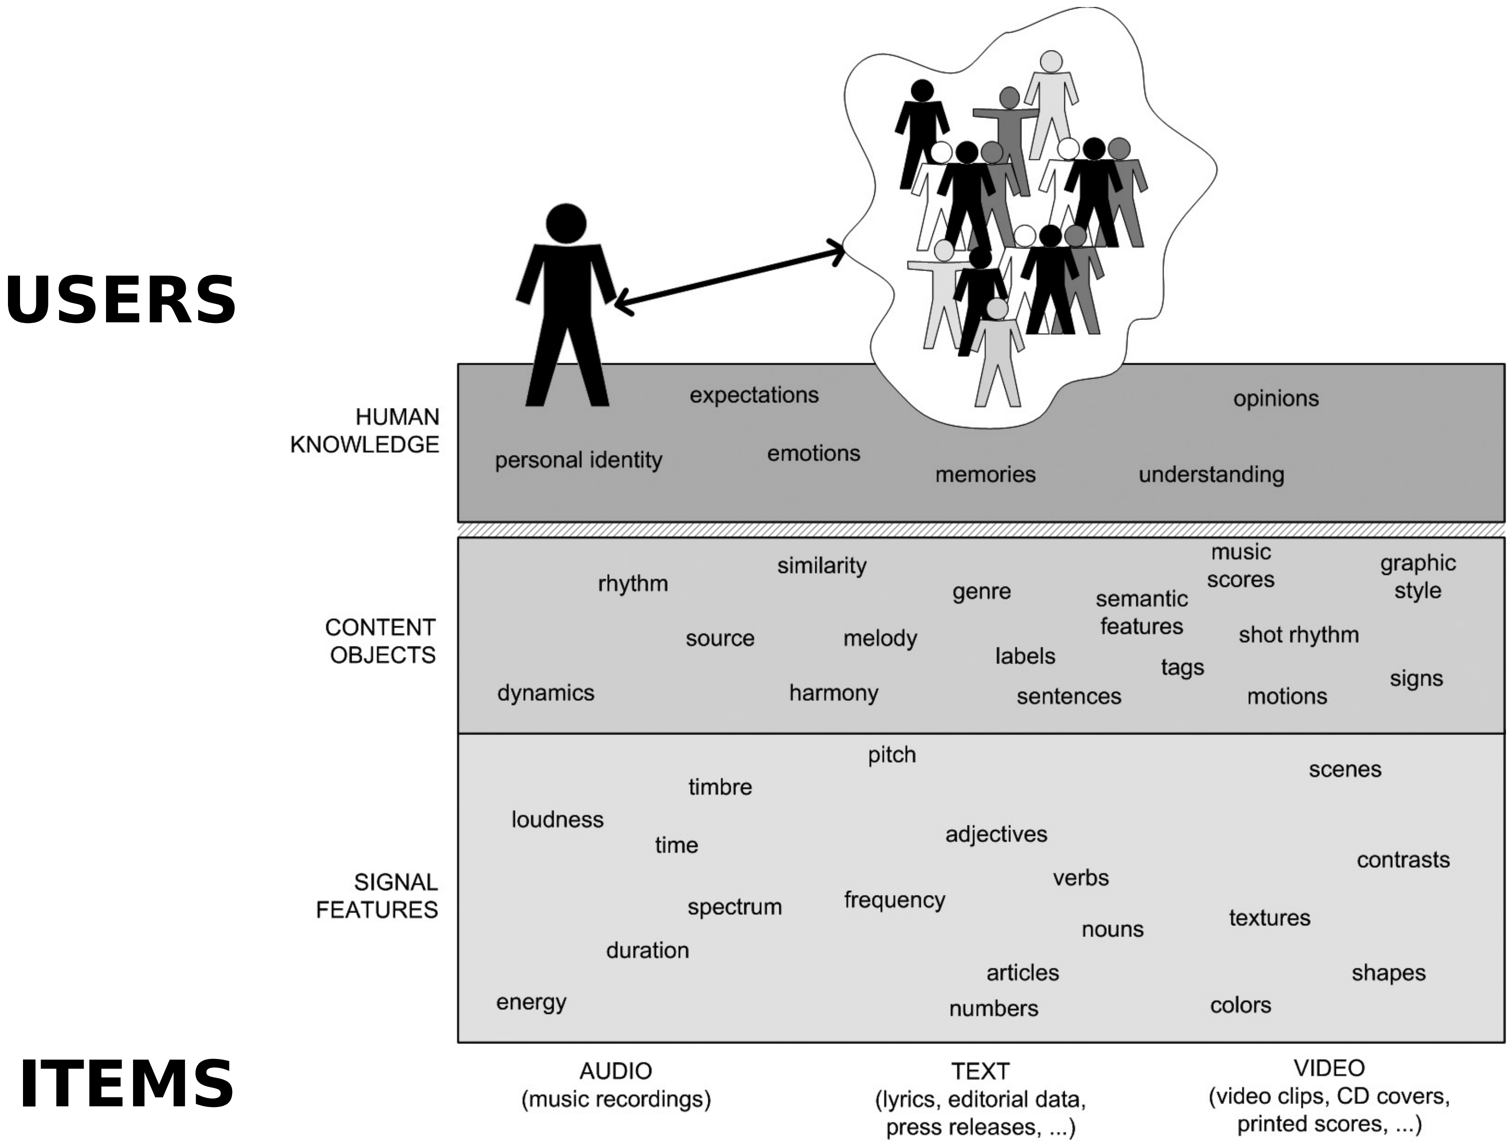

formalization:: describing users & items

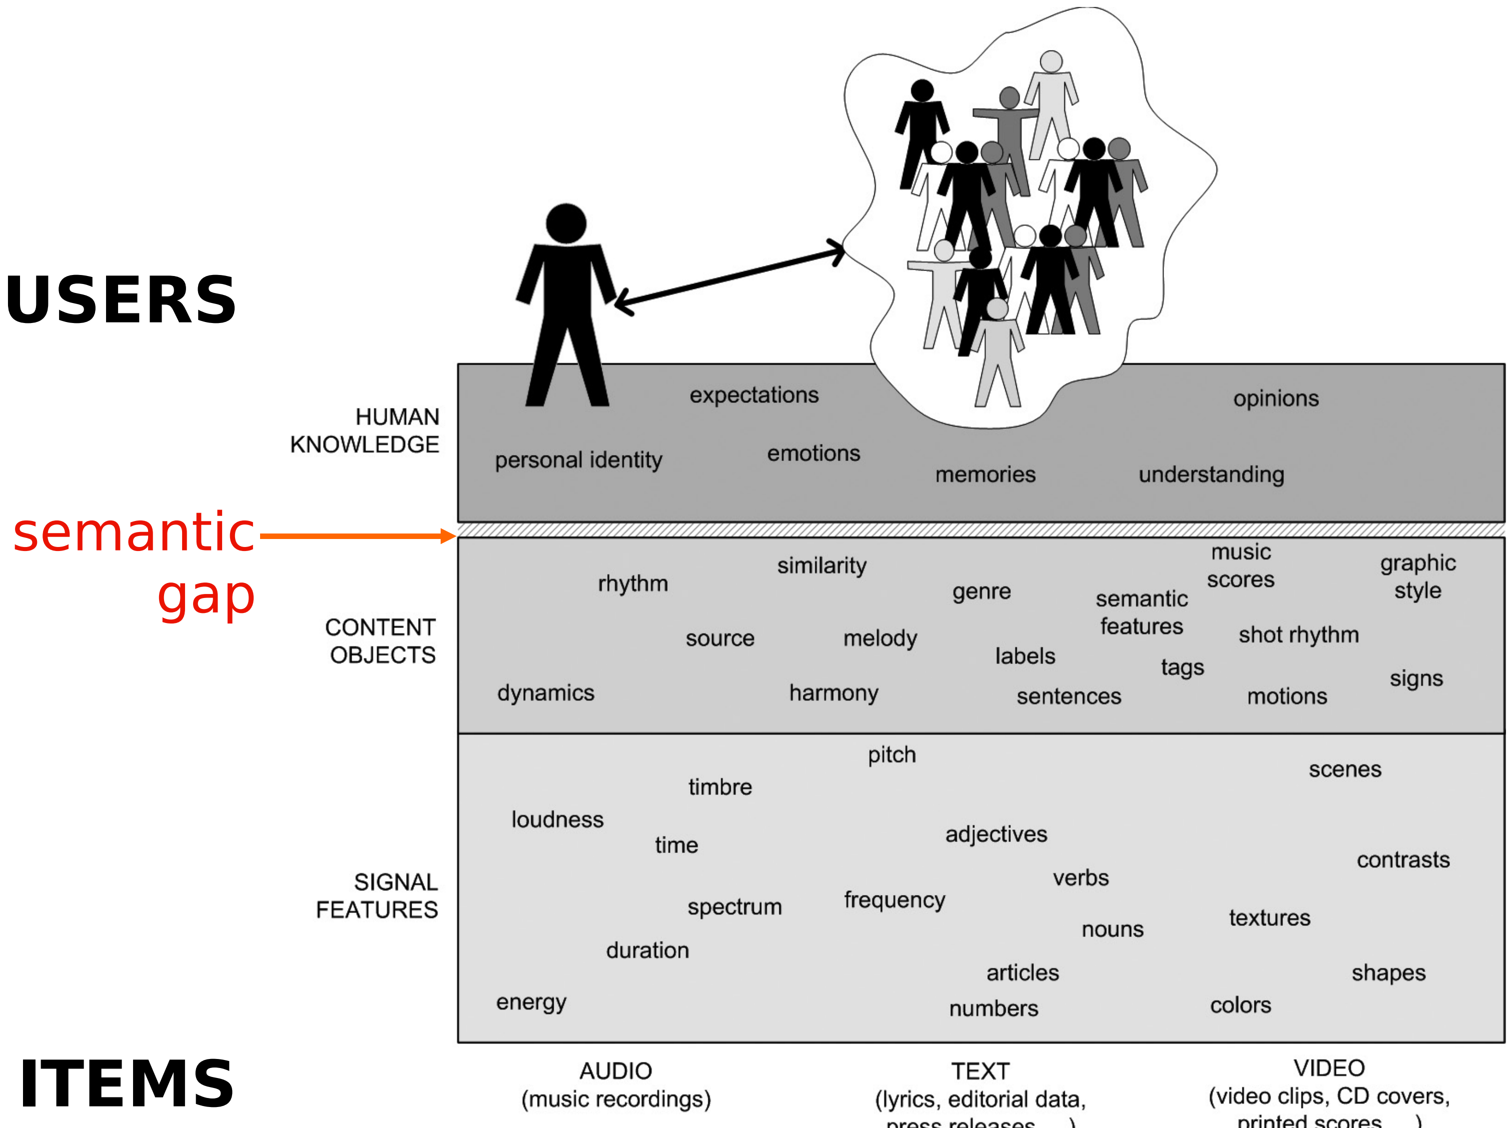

formalization:: describing users & items

USERS

semantic gap

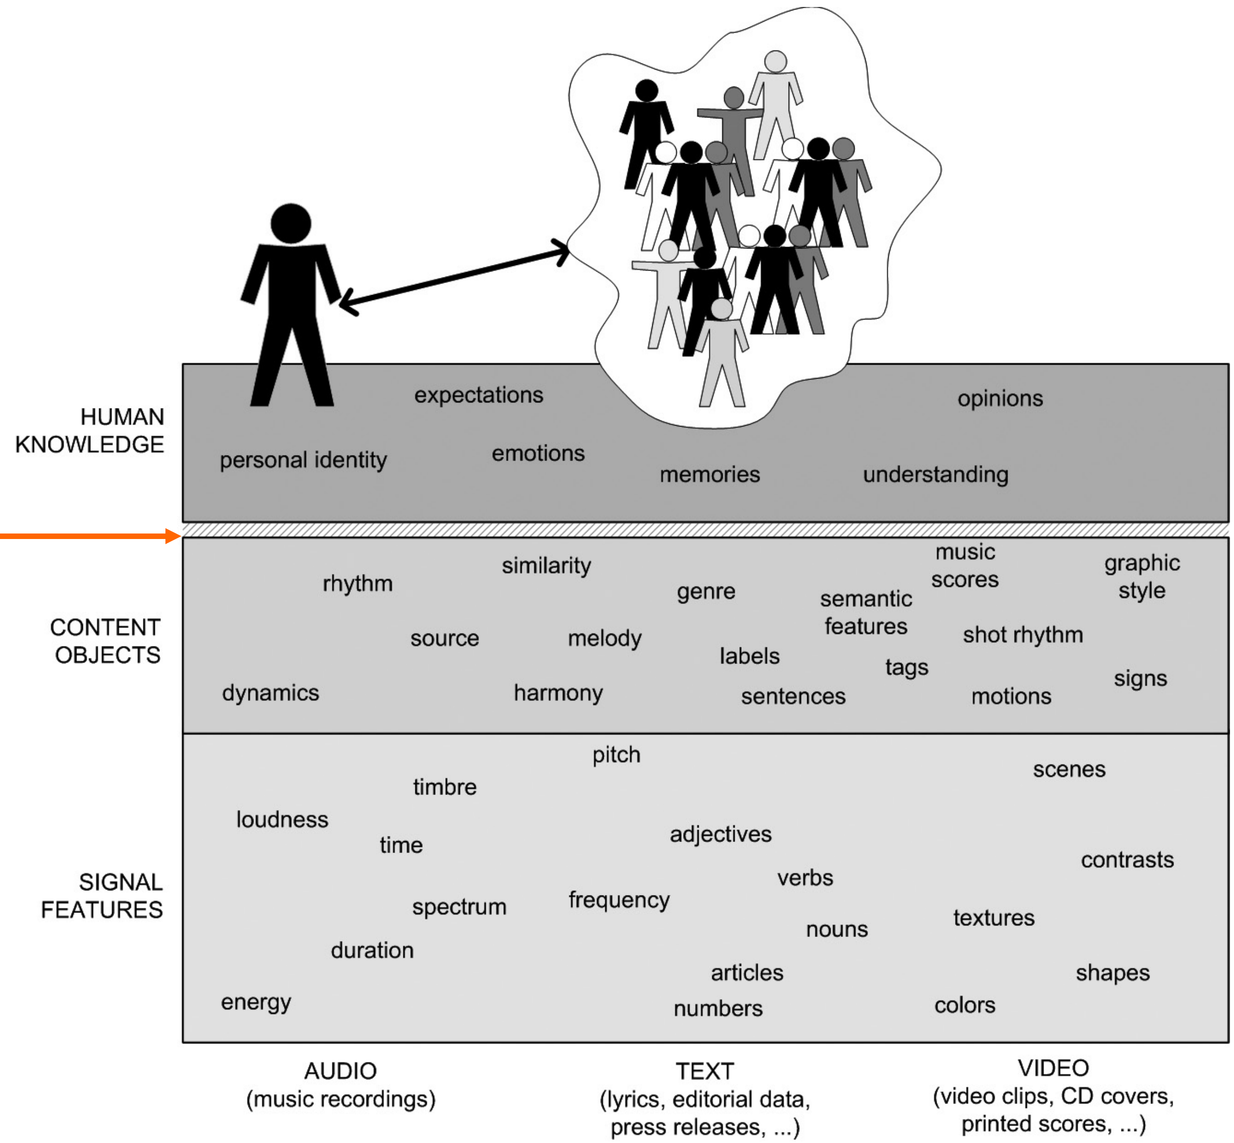

- get a better understanding of the music assets!

ITEMS

formalization:: describing users & items

USERS

- get a better understanding of the users preferences!

semantic gap

- get a better understanding of the music assets!

ITEMS

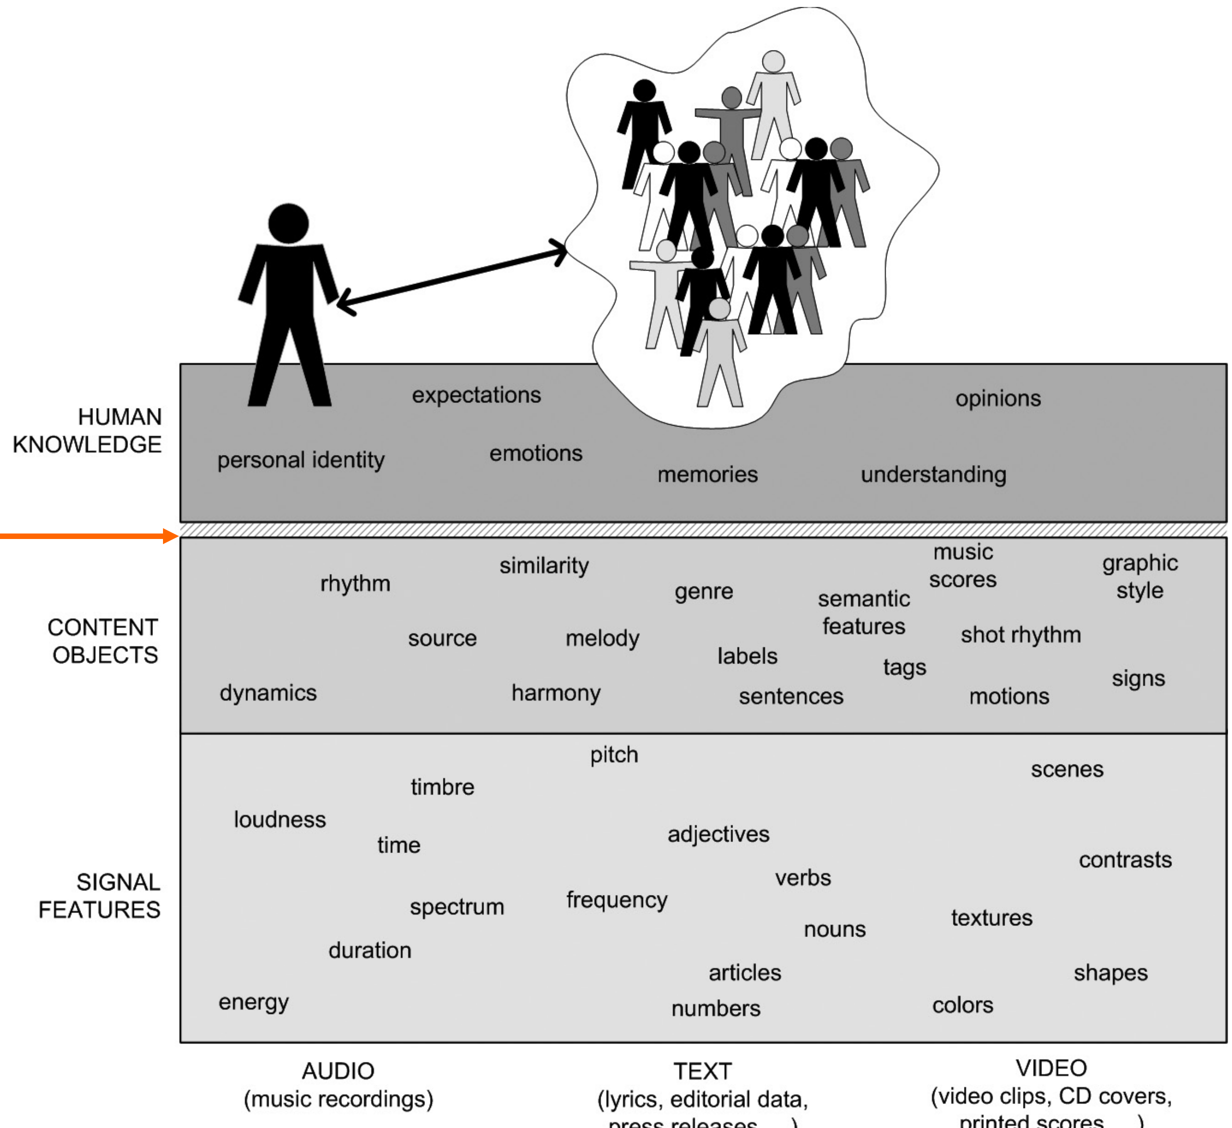

formalization:: describing users

USERS

- get a better understanding of the users preferences!

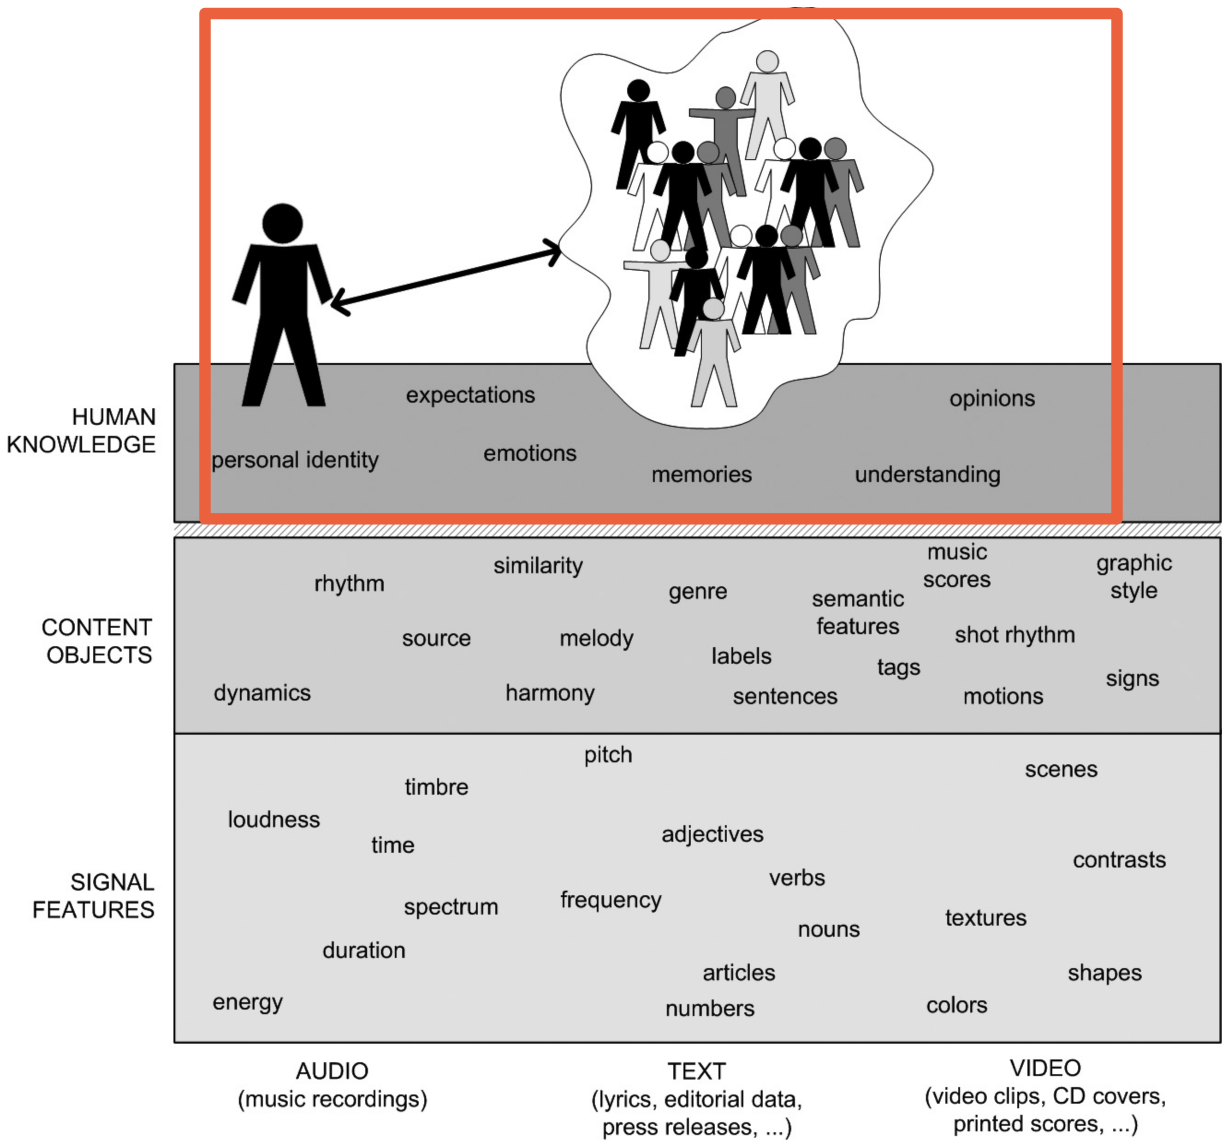

formalization:: describing users (me)

USERS

- get a better understanding of the users preferences!

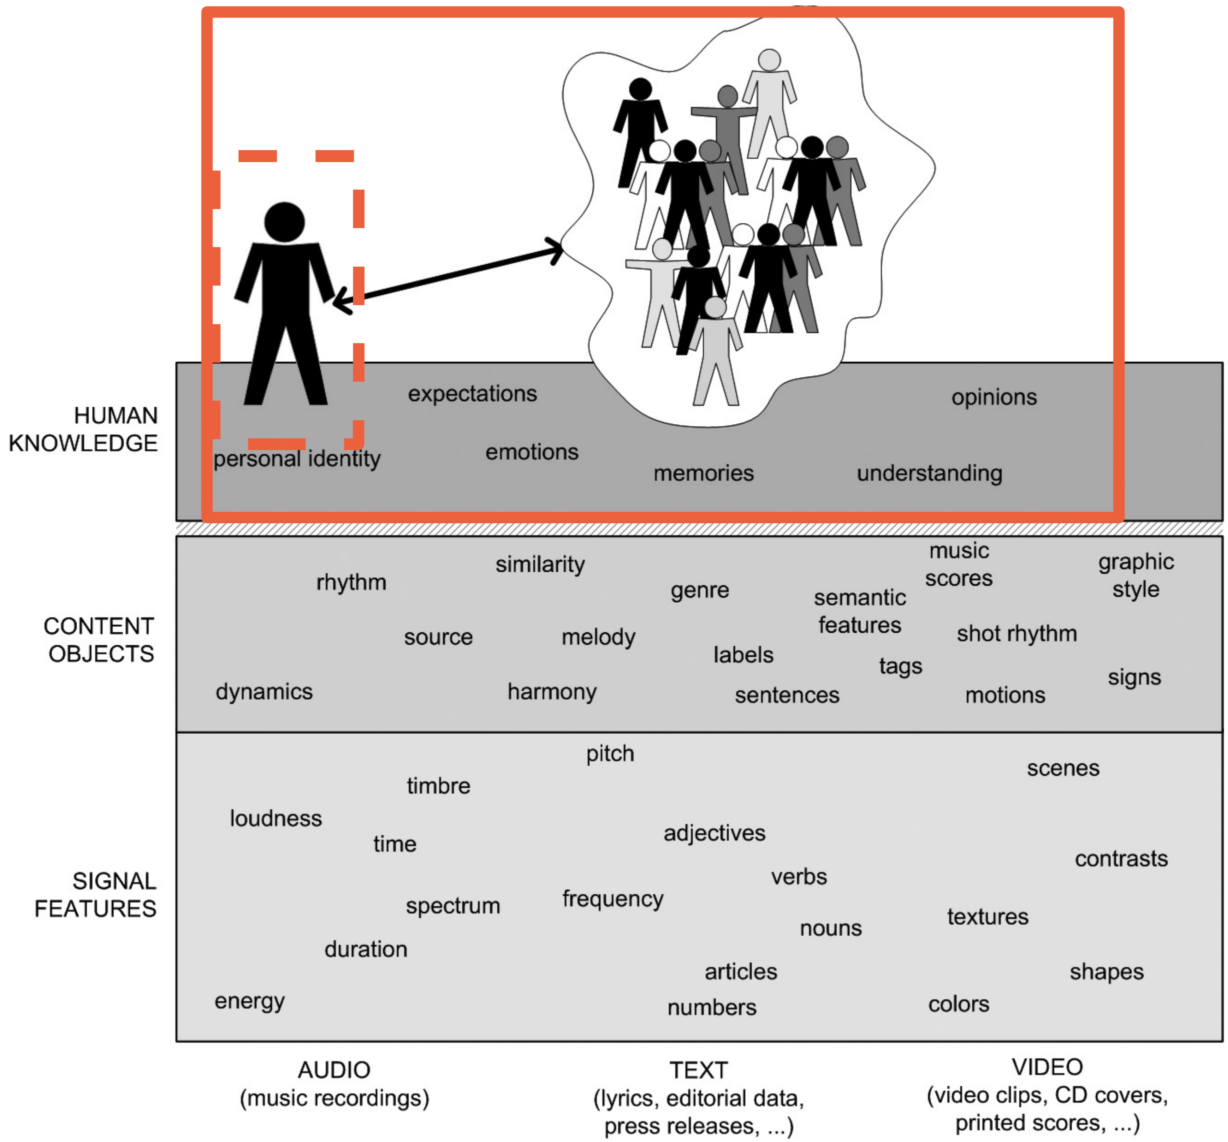

formalization:: describing users (me)

- me and myself (user profile) [Uitdenbogerd, 2002]
  - demographic
    - age, gender, languages, family status, income, education level, etc.
  - geographic
    - location
  - psychographic
    - general interests
    - hobbies
    - music preferences
    - ...

formalization:: describing users (me)

- ...me and myself (user profile)
  - music preferences
    - explicit
      - list of preferred / hated artists
      - list of preferred / hated songs
      - ratings / reviews / opinions (my blog)
      - (relevance feedback)
    - implicit
      - listening habits (play / stop / skip)
      - pages / blogs visited
      - ...

formalization:: describing users (me)

- …me and myself (user profile)
  - a note about **implicit** and **explicit** data
    - Implicit data like purchases may be noisy, but it also can be more accurate
    - “I love *cool* Jazz (especially Chet Baker), as well as J.S.Bach fugues”

formalization:: describing users (me)

- ...me and myself (user profile)
  - ❖ Yeah, yeah... cool jazz and Bach!

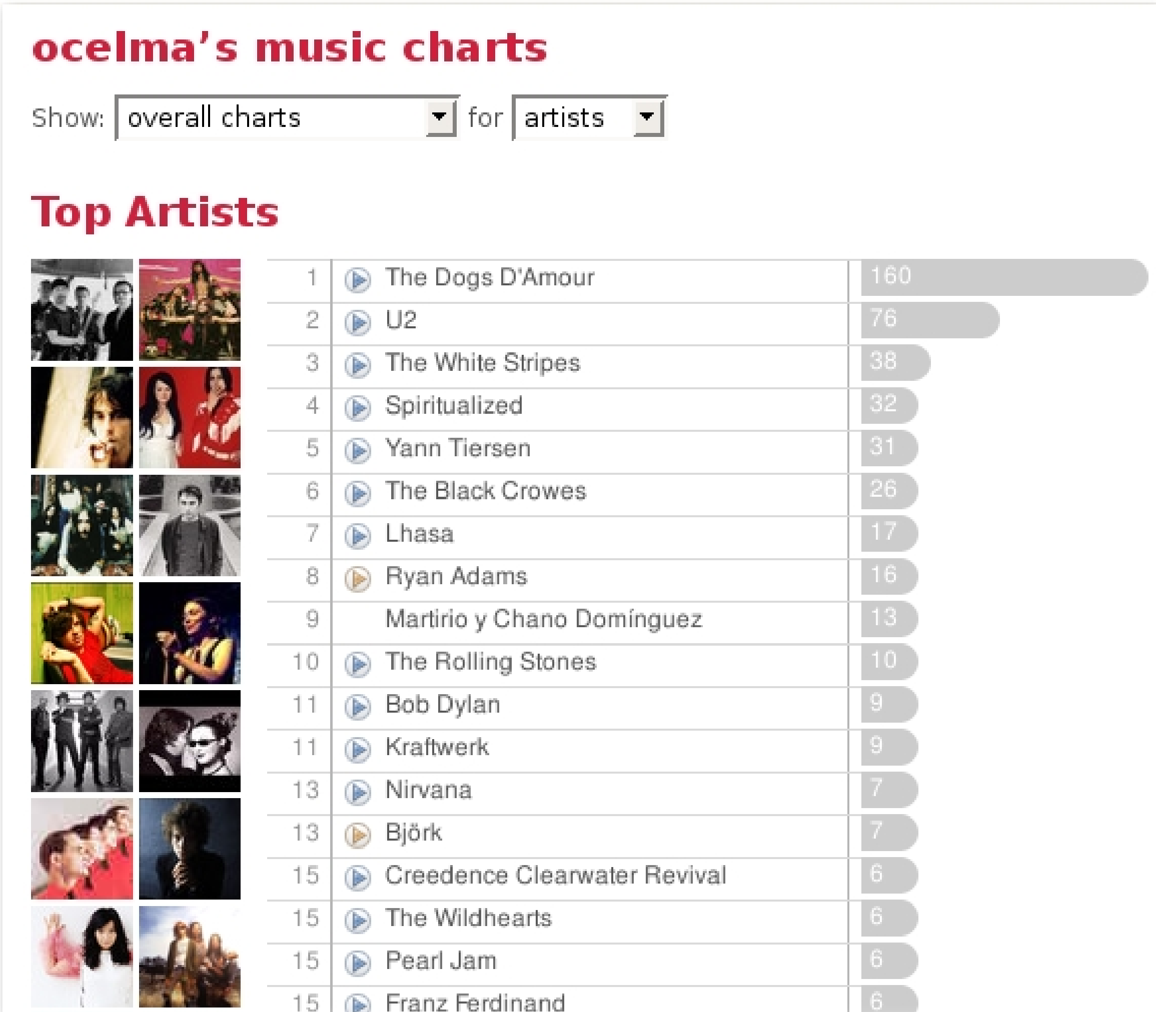

formalization:: describing users (me)

- ...me and myself (user profile)
  - Explicit data: People...
    - (1) usually **won't bother**,
    - (2) if they do bother, only provide **partial information** or even lie,
    - (3) even if they bother, tell the truth, and provide complete information, they usually **fail to update** their information over time."
  - From:
    - http://glinden.blogspot.com/2007/05/explicit-vs-implicit-data-for-news.html
    - http://glinden.blogspot.com/2004/05/interview-with-craig-silverstein-from.html

formalization:: describing users

USERS

- get a better understanding of the users preferences!

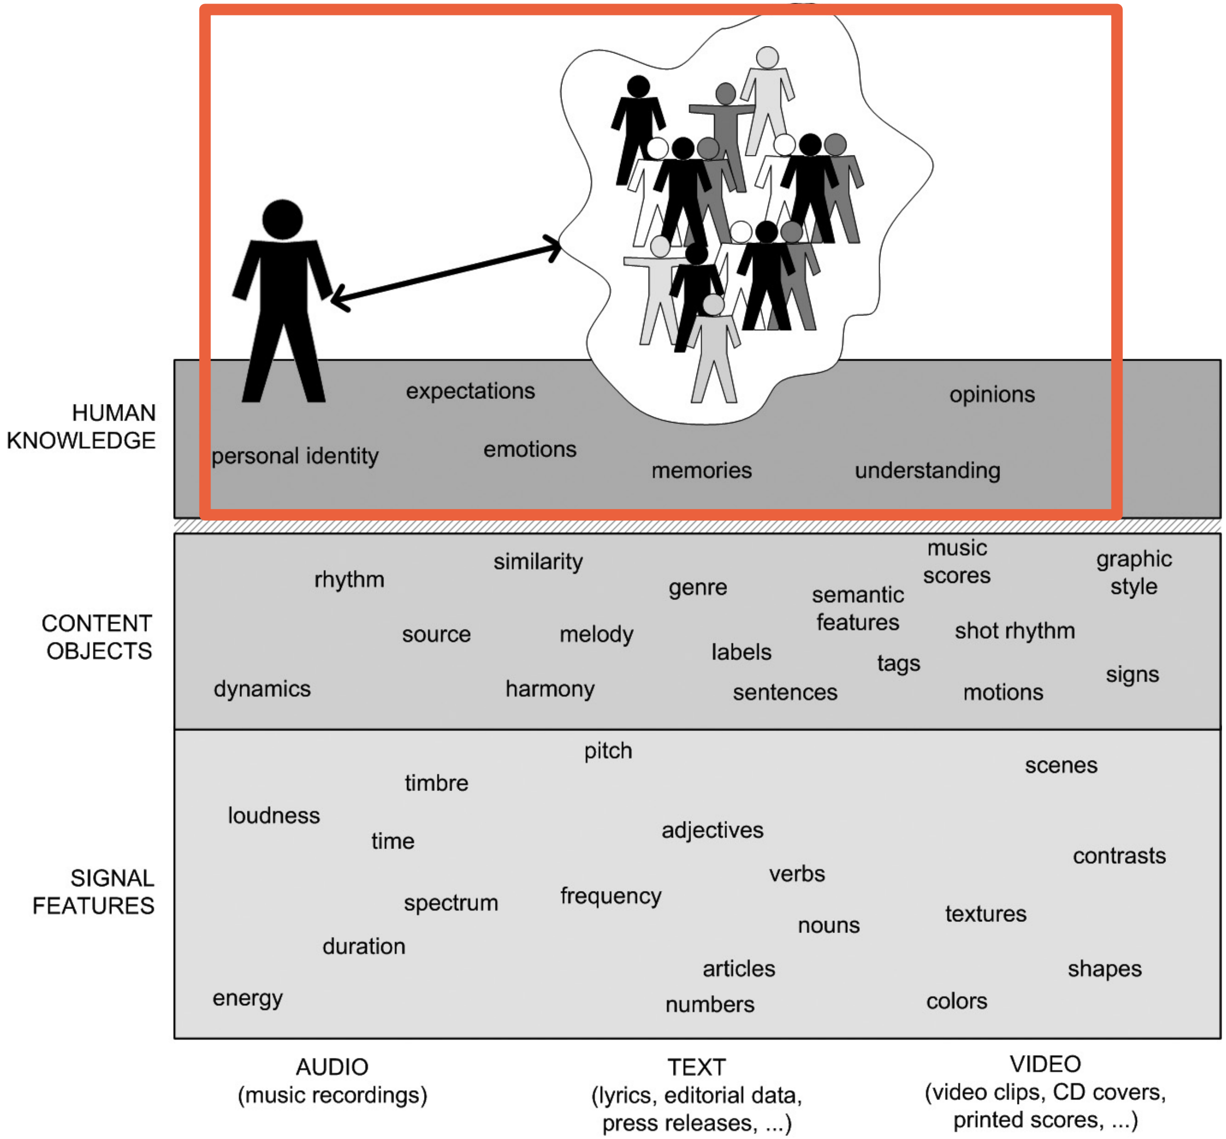

formalization:: describing users (us)

USERS

- get a better understanding of the users preferences!

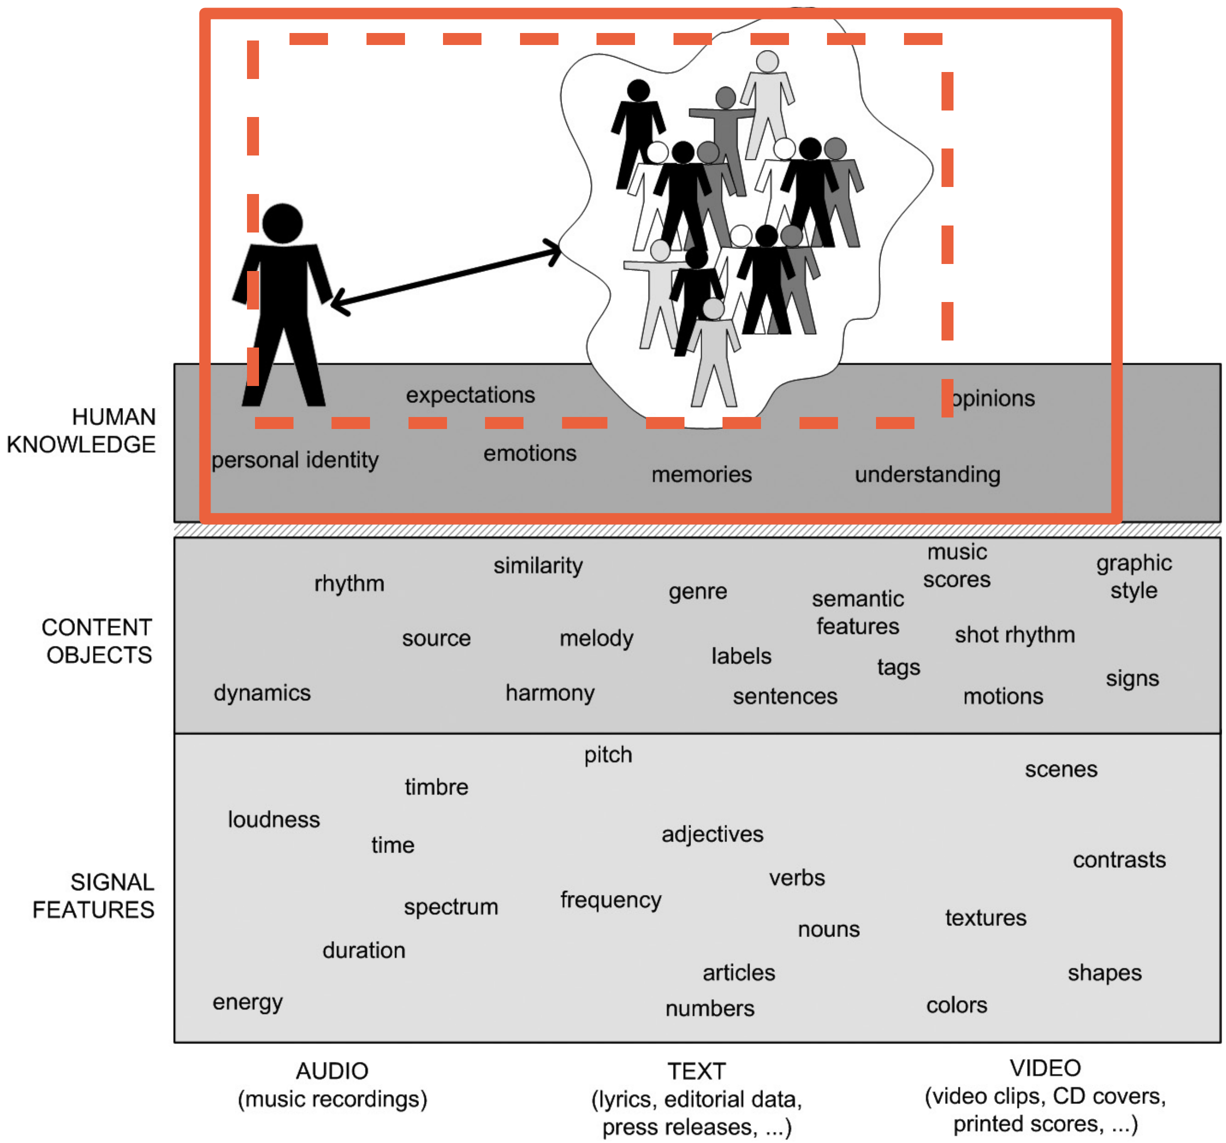

formalization:: describing users (us)

- me and the world (socializing) [Kazienko, 2006]
  - interaction with other users
  - relationships among users
    - duration
    - mutual watchings of blogs, artists pages, songs, etc.
    - common communications

formalization:: describing users (us)

- BlueTuna [Baumann, 2007]
  - a socializer: share your music tastes with people near by
  - meeting people who share the same music tastes
    - check with a mobile phone to see who in a close proximity has my tastes

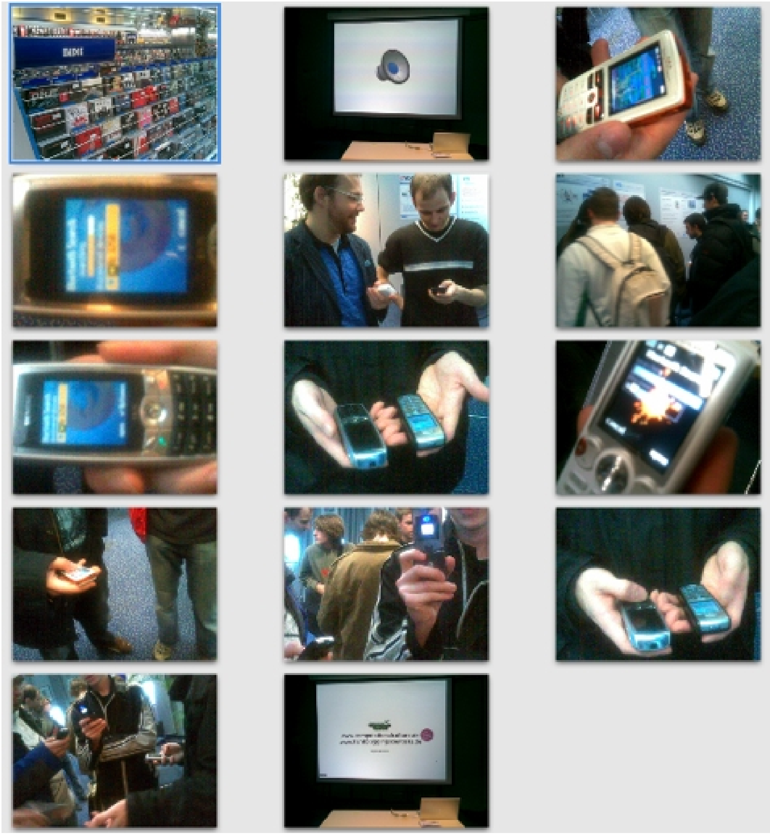

formalization:: describing users :: languages

- Some representations
  - User Modelling for Information Retrieval Language (UMIRL)
  - MPEG-7
  - Friend of a Friend (FOAF)
  - General User Model Ontology (GUMO)
- ...based on XML/XML-Schema or RDF/OWL

formalization:: describing users :: umirl

- User Modelling for Information Retrieval Language (UMIRL) [Chai, 2000]
  - demographic & geographic information
  - music background and music preferences
  - create definition of a perceptual feature, and its context (usage)
    - perceptual feature: "a **romantic** piece has a **slow tempo**, **lyrics** are related with *love*, and has a **soft intensity**"
    - usage: while **having** a special **dinner** with **girlfriend**
  - Languages
    - XML
    - No XML Schema (!)

formalization:: describing users :: umirl

- a complete example...

```
<user>
  <generalbackground>
    <name>Joan Blanc</name>
    <education>MS</education>
    <citizen>Catalan</citizen>
    <sex>male</sex>
  </generalbackground>
  <musicbackground>
    <education>none</education>
    <instrument>guitar</instrument>
    <instrument>vocal</instrument>
  </musicbackground>
```

formalization:: describing users :: umirl

- a complete example...

<**user**>

<**generalbackground**>

<name>Joan Blanc</name>

<education>MS</education>

<citizen>Catalan</citizen>

<sex>male</sex>

</generalbackground>

<**musicbackground**>

<education>none</education>

<instrument>guitar</instrument>

<instrument>vocal</instrument>

</musicbackground>

(continue...)

<musicpreferences>

<genre>blues</genre>

<genre>rock</genre>

<composer>Johann Sebastian Bach</composer>

<artist>The Dogs d'Amour</artist>

<sample>

<title>Two hearts beat as one</title>

<artist>U2</artist>

</sample>

</musicpreferences>

formalization:: describing users :: umirl

(continue....)

<**habit**>

<context>**Happy**

<tempo>very fast</tempo>

<genre>rock</genre>

</context>

<perceptualfeature>**Romantic**

<tempo>very slow</tempo>

<intensity>soft</intensity>

<lyrics>*love*</lyrics>

</perceptualfeature>

<context>**Dinner with fiance**

<perceptualfeature>**Romantic**</perceptualfeature>

</context>

</habit>

</user>

formalization:: describing users :: mpeg-7

- MPEG-7
  - “standard” for multimedia content description
  - Languages
    - XML
    - huge XML-Schema (!!!)
  - modeling user preferences
    - content filtering
    - searching and browsing
    - usage history
  - Example
    - “I like the album **To bring you my love**, from **P.J. Harvey**”

formalization:: describing users :: mpeg-7

```xml
<UserPreferences>
 <UserIdentifier protected="true">
  <Name xml:lang="ca">Joan Blanc</Name>
 </UserIdentifier>
 <FilteringAndSearchPreferences>
  <CreationPreferences>
   <Title preferenceValue="8">To bring you my love</Title>
   <Creator>
    <Role>
     <Name>Singer</Name>
    </Role>
    <Agent xsi:type="PersonType">
     <Name>
      <GivenName>Polly Jean</GivenName>
      <FamilyName>Harvey</FamilyName>
     </Name>
    </Agent>
   </Creator>
```

(continue...)

```xml
   <Keyword>dramatic</Keyword>
   <Keyword>fiery</Keyword>
   <DatePeriod>
    <TimePoint>1995-01-01</TimePoint>
   </DatePeriod>
  </CreationPreferences>
 </FilteringAndSearchPreferences>
</UserPreferences>
```

formalization:: describing users :: foaf

- Friend of a Friend (FOAF)
  - “a machine readable homepage”
  - semantic web flavour
    - add any available ontology
    - (in particular the Music Ontology)
  - Languages
    - OWL (the ontology)
    - RDF (the actual data)
  - Example
    - FOAFing the Music

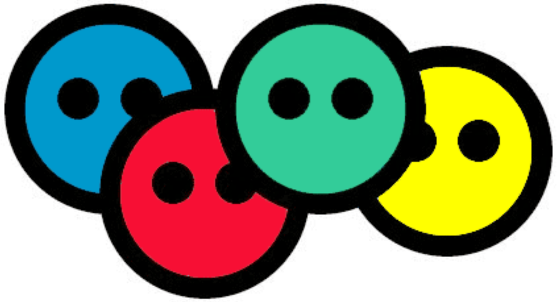

formalization:: describing users :: foaf

- Friend of a Friend (FOAF)
  - Linking Open Data

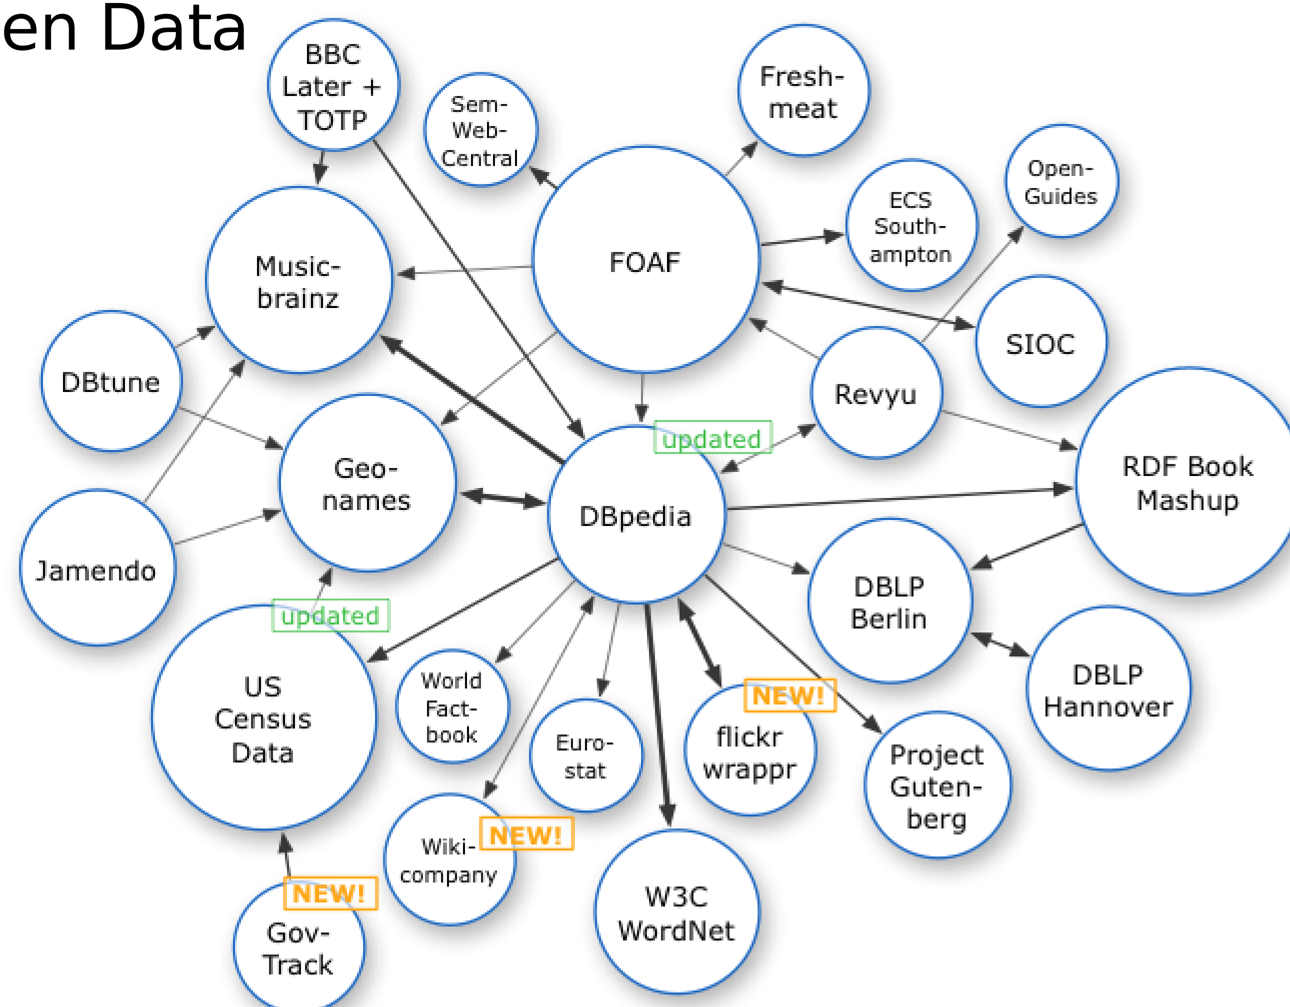

formalization:: describing users :: foafexample

```
<foaf:Person rdf:ID="me">
  <foaf:name>Oscar Celma</foaf:name>
  <foaf:nick>ocelma</foaf:nick>
  <foaf:gender>male</foaf:gender>
  <foaf:depiction rdf:resource="http://www.iua.upf.edu/mtg/img/faces/ocelma.png" />
  <foaf:homepage rdf:resource="http://www.iua.upf.edu/~ocelma"/>
  <foaf:workplaceHomepage rdf:resource="http://mtg.upf.edu"/>
  <foaf:mbox rdf:resource="mailto:oscar.celma@iua.upf.edu"/>
  <foaf:based_near geo:lat='41.385' geo:long='2.186' />
  <foaf:holdsAccount>
   <foaf:OnlineAccount>
    <foaf:accountName>ocelma</foaf:accountName>
    <foaf:accountServiceHomepage rdf:resource="http://last.fm"/>
   </foaf:OnlineAccount>
  </foaf:holdsAccount>
  <foaf:interest dc:title="Gretsch Guitars" rdf:resource="http://en.wikipedia.org/wiki/Gretsch" />
 </foaf:Person>
```

formalization:: describing users :: foaf

- Add explicit music interests and preferences into a FOAF profile
  - using the Music Ontology [Raimond, 2007]
  - Example
    - "I like the album **To bring you my love**, from **P.J. Harvey**"

formalization:: describing users :: foafexample

```
<foaf:interest>
  <mo:Record rdf:about="http://zitgist.com/music/record/24e5b7f5-14cd-4a65-b87f-
    91b5389a4e3a">
    <dc:title>To bring you my love</dc:title>
    <dcterms:created>1995-02-22T00:00:00Z</dcterms:created>
    <mo:releaseType rdf:resource="http://purl.org/ontology/mo/album"/>
    <mo:releaseStatus rdf:resource="http://purl.org/ontology/mo/official"/>
    <mo:discogs rdf:resource="http://www.discogs.com/release/379469"/>
    <foaf:img rdf:resource="http://g-ec2.images-amazon.com/images/I/21cwiSYVYkL.jpg"/>
    <foaf:made>
      <mo:MusicGroup rdf:about='http://zitgist.com/music/artist/e795e03d-b5d...fb308a2c'>
        <foaf:name>Polly Jean Harvey</foaf:name>
        <mo:discogs rdf:resource="http://www.discogs.com/artist/PJ+Harvey" />
        <mo:wikipedia rdf:resource="http://en.wikipedia.org/wiki/PJ_Harvey" />
      </mo:MusicGroup>
    </foaf:made>
  </mo:Record>
</foaf:interest>
```

formalization:: describing users :: gumo

- General User Model Ontology (GUMO) [Heckmann, 2007]
  - top level ontology
  - Aim
    - exchange of user profile data between adaptive systems
  - includes the Big Five personality traits
    - Neuroticism, Extraversion, Agreeableness, Conscientiousness, and Openness to Experience
  - Language
    - OWL

formalization:: describing users :: gumo

- main elements

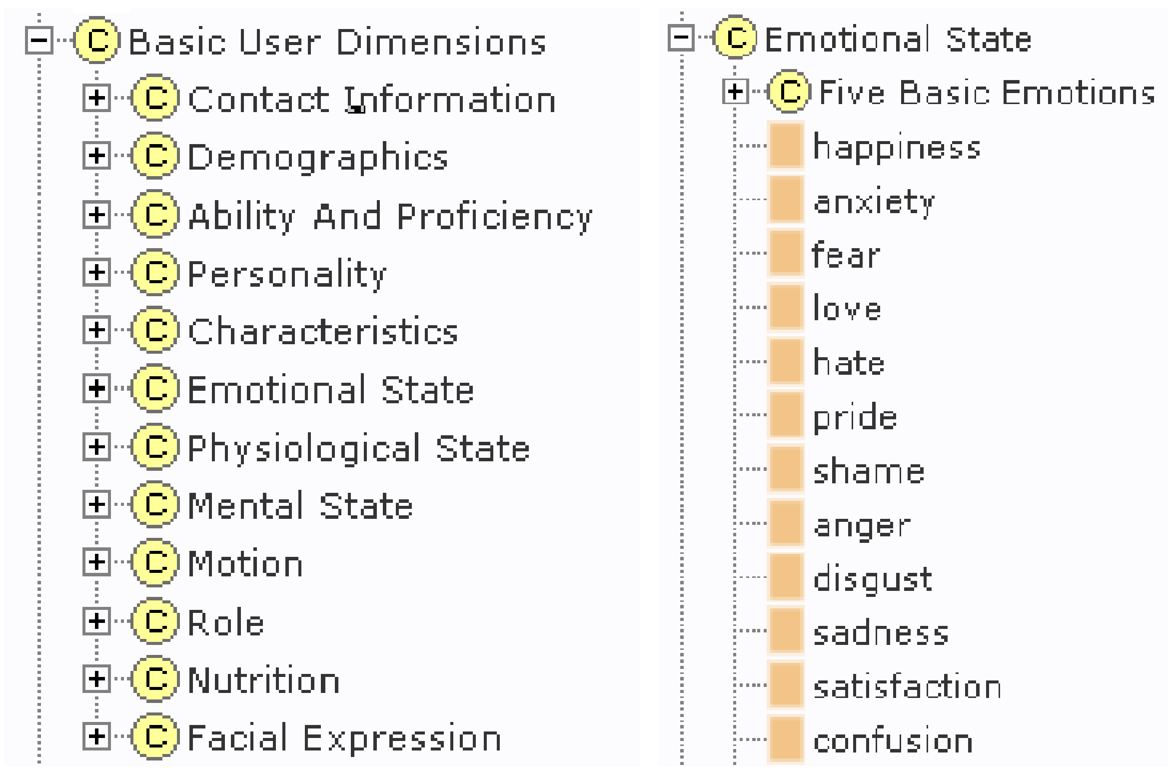

formalization:: describing users :: gumo

- main elements

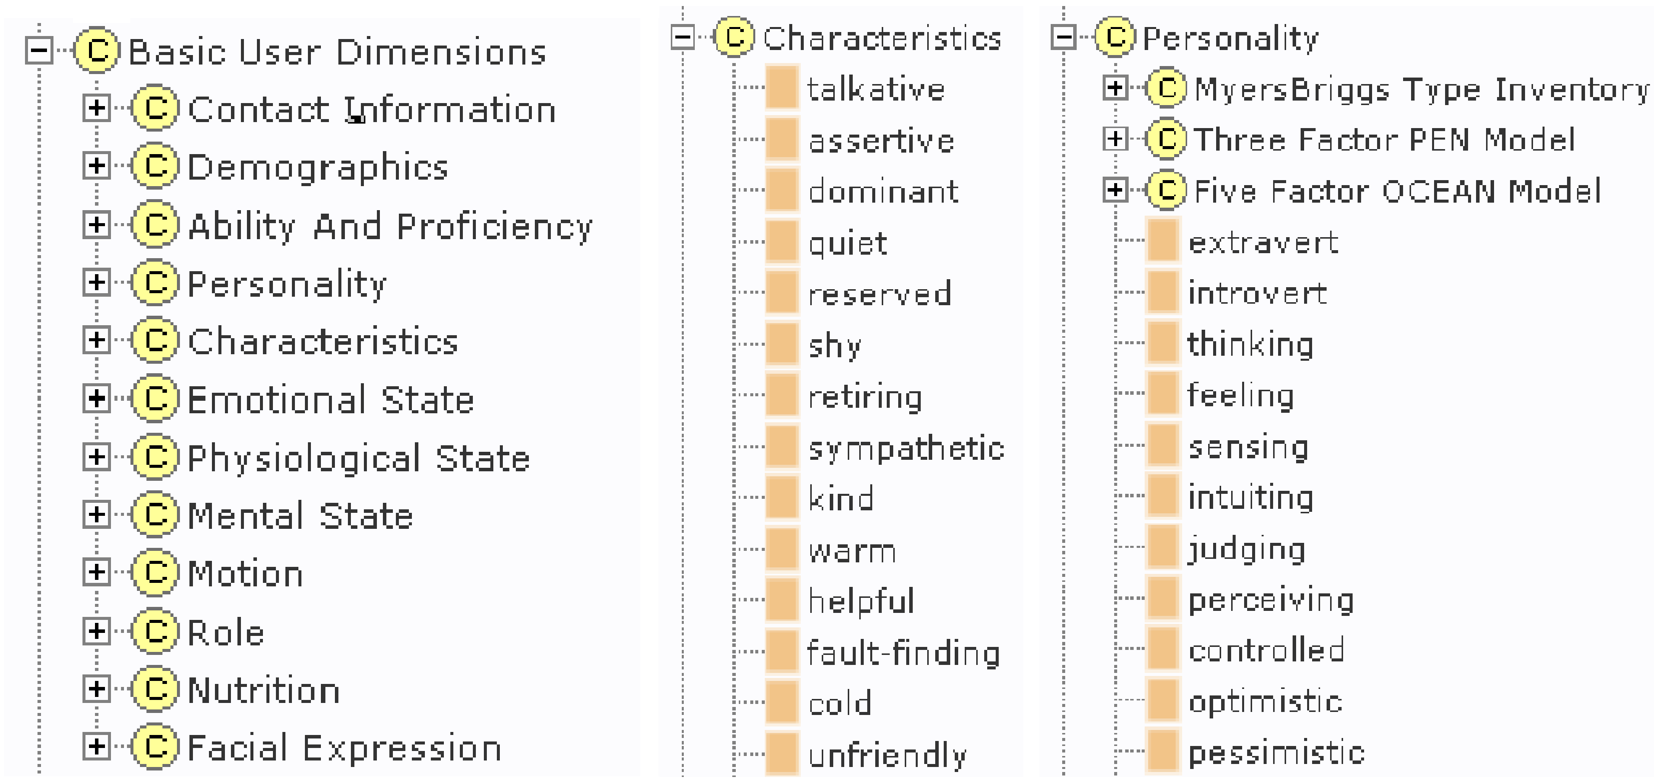

formalization:: describing users

- Complexity and expressiveness of the representations

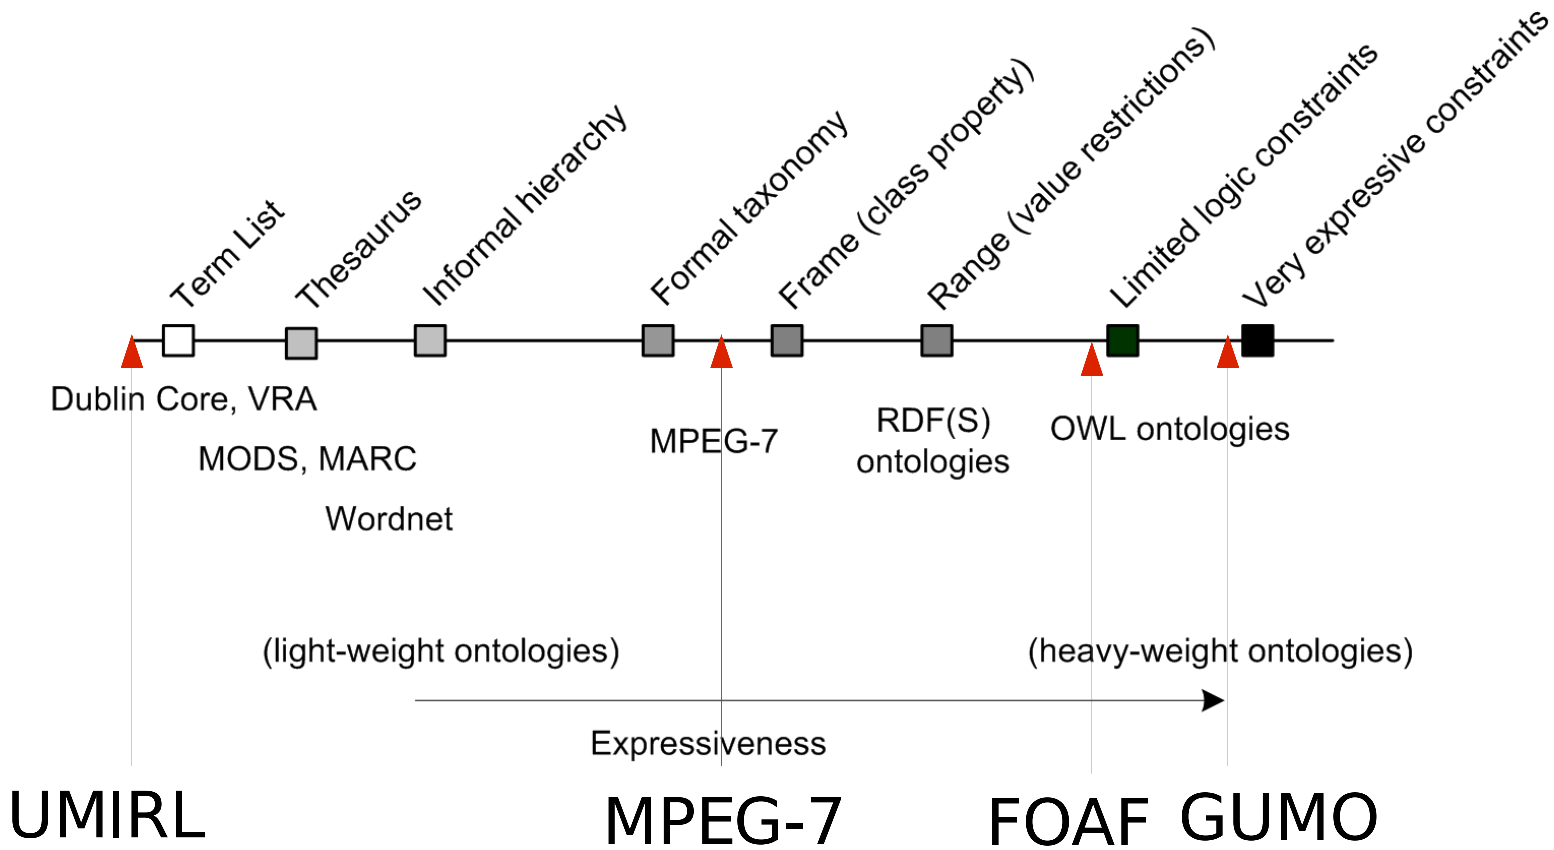

formalization:: describing users

- Issues
  - ❖ What about all the information that the user has on the "Web 2.0" (her blog, her *other* accounts, etc)? How to exploit them?
  - ❖ What about multiple-profiles?
    - me-at-work, me-at-home or me-on-weekends, etc.
  - ❖ Each system handles user information in their way
    - No interoperability between systems
    - Most information is not used

formalization:: describing items

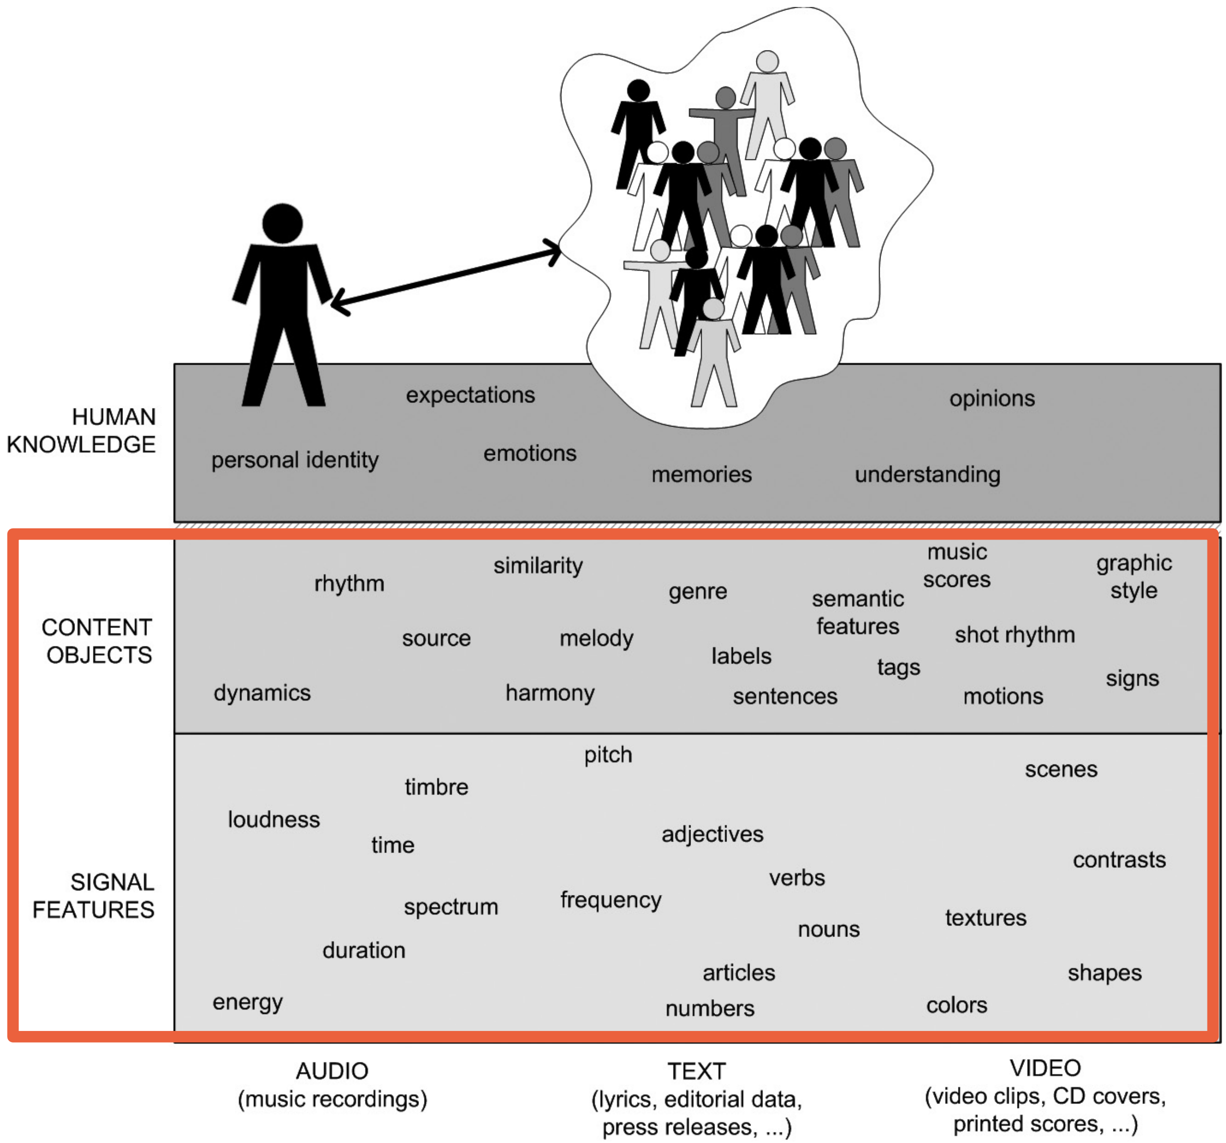

- get a better understanding of the music assets!

ITEMS

formalization:: describing items

- Text description – using the halo of text surrounding music
  - Expert-applied metadata
  - Web Mining
    - Reviews, Playlists, Lyrics
  - Tagging
    - last.fm, qloud
- Audio description
  - instrumentation / tonality / rhythm / timbre
    - Manual – Pandora, SoundFlavor
    - Automatic – Owl MM, MusicIP, One Llama, SITM, BMAT

formalization:: describing items :: text

- the text halo

Similarity based upon the text halo surrounding the music.

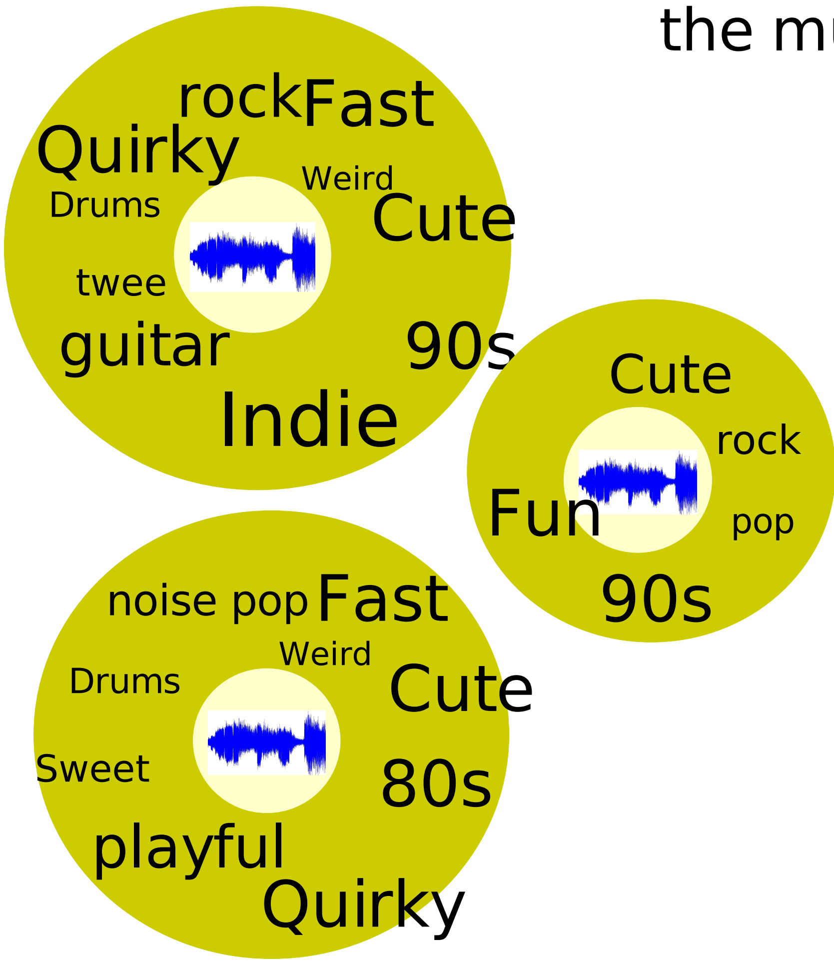

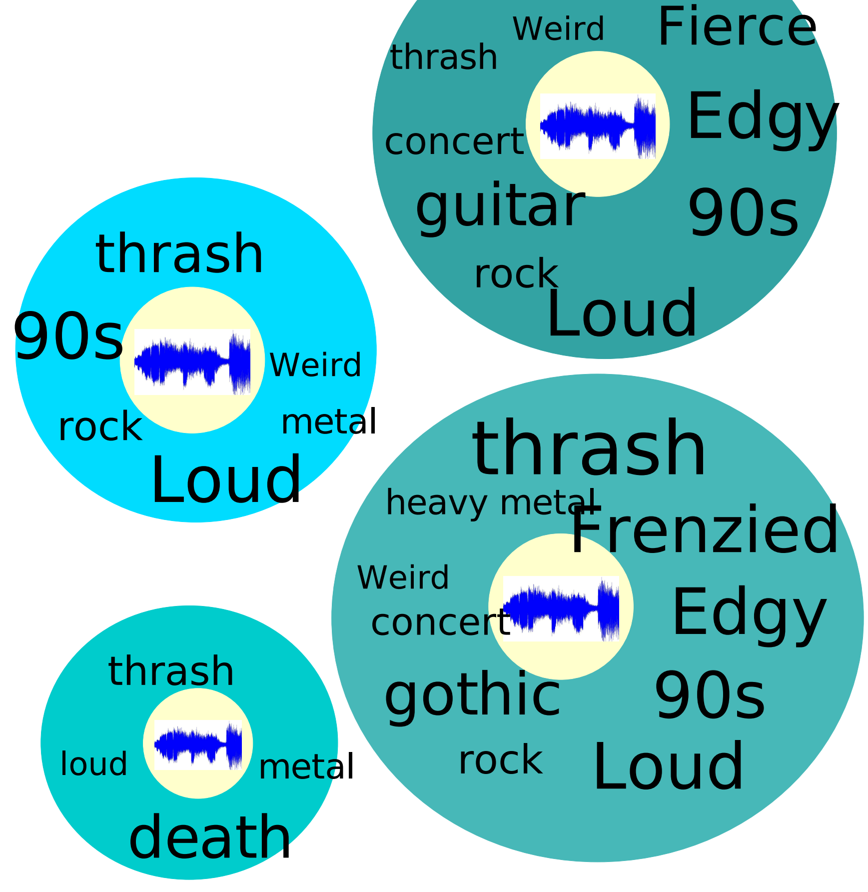

formalization:: describing items :: text

- determining similarity of text
  - Traditional text information retrieval technique: TF x IDF
    - TF – Term Frequency – a measure of the frequency of a term (word) within a document
    - IDF – Inverse Document Frequency
      - Measure of a term's rarity across the set of documents
      - Could be: 1 / document frequency
      - But more typically: log (n / df)

$$\mathrm{tf_i} = \frac{n_i}{\sum_k n_k}$$

$$\mathrm{idf_i} = \log \frac{|D|}{|\{d : t_i \ni d\}|}$$

$$\mathrm{tfidf} = \mathrm{tf} \cdot \mathrm{idf}$$

formalization:: describing items :: text

- determining similarity of text
  - ❖ cosine similarity

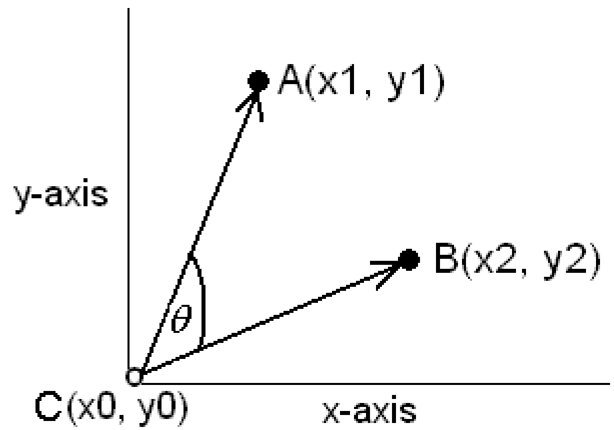

$$\mathrm{Sim}(A, B) = \text{cosine } \theta = \frac{A \bullet B}{|A||B|} = \frac{x1^*x2 + y1^*y2}{(x1^2 + y1^2)^{1/2} \; (x2^2 + y2^2)^{1/2}}$$

formalization:: describing items :: text

- Expert-applied metadata
  - **Deerhoof**:

allmusic

    - **Genre: rock**
    - **Styles:**
      - Indie Rock, Noise Pop, Noise-Rock, Post-Rock, Experimental
    - **Moods**:
      - Volatile, Freewheeling, Energetic, Whimsical, Playful, Rambunctious, Exuberant, Carefree, Irreverent, Springlike, Fun, Bittersweet, Cheerful, Cathartic, Innocent, Messy, Sweet, Precious, Naïve
    - **Similar Artists**: Persephone's Bees, Black Dice ...
    - **Influenced By:** Boredoms, Yoko Ono ...
    - **Followers:** The Mae Shi ...

formalization:: describing items :: text

Six Degrees of Black Sabbath

Starting Artist: Jeff Beck Ending Artist: Beck Find Path

I can find a path from **Jeff Beck** to **Beck** in 3 steps.

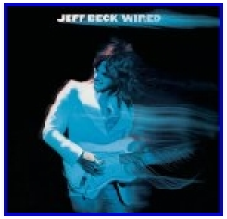

Jeff Beck who performed with
neighbors start end

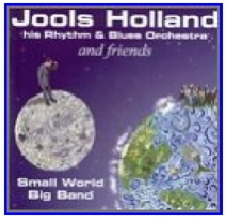

Jools Holland & His Rhythm & Blues Orchestra who performed with
bypass neighbors start end

Norah Jones who performed with
bypass neighbors start end

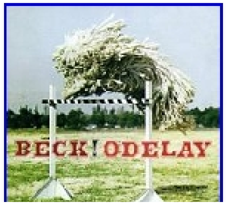

Beck
neighbors start end

Fun with metadata

formalization:: describing items :: text

- deerhoof biography

By turns cuddly and chaotic, San Francisco's Deerhoof mixes noise, sugary melodies, and an experimental spirit into sweetly challenging and utterly distinctive music. The group began as the brainchild of guitarist Rob Fisk and drummer/keyboardist Greg Saunier in 1994; early releases, such as 1995's 7"s Return of the Woods M'Lady and For Those of Us on Foot, had a more traditionally harsh, no wave-inspired sound, though they also included the quirky tendencies that dominated their later efforts.

formalization:: describing items :: text

- Web mining

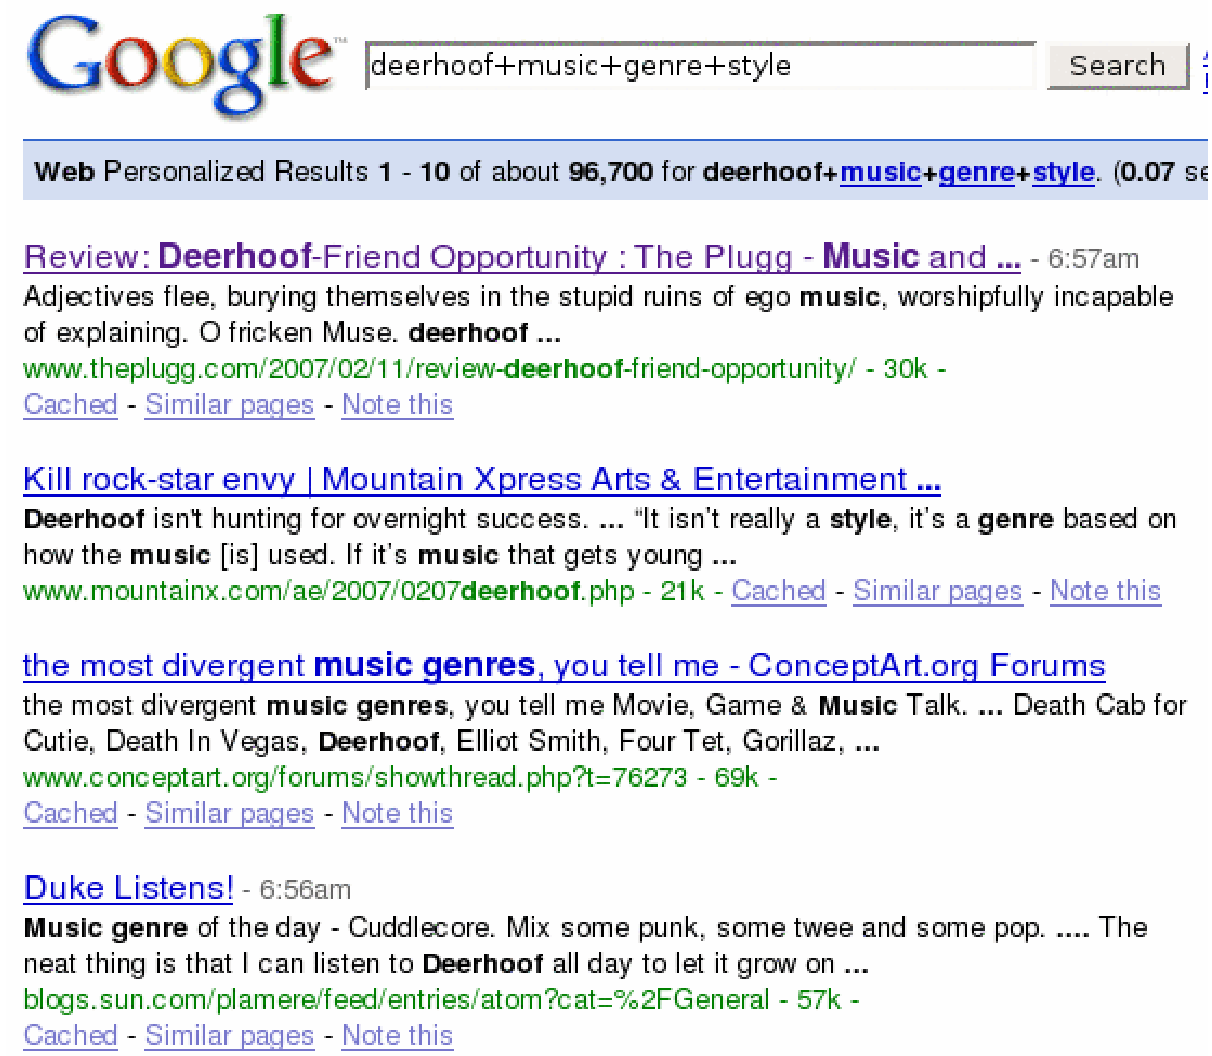

formalization:: describing items :: text

- Web mining
  - Music blogs

LINE OUT
THE STRANGER'S MUSIC BLOG

Deerhoof's new album, Friend Opportunity is **amazing**. I've never been a gung-ho **fan**, despite having numerous friends **rave** to me about how **awesome** these Bay Area **indie-rock** mainstays are. But this new full-length album strikes me immediately as their **finest to date**. **Not bad** for a group 12-years into its career. Its **radiant** yet skewed **beauty** and **surprising dynamics** set a towering example for how **indie** rock should sound and move in 2007. You can sense that they have intimate knowledge of **no wave**, **sunshine pop**, **astral jazz**, **Captain Beefheart**, **J pop**, **Raincoats**, **Polvo**, **Boredoms**, and many other exemplary touchstones. Yet they weave these styles and influences so adroitly that the resultant songs are instantly identifiable as only Deerhoof compositions.

formalization:: describing items :: text

- Web mining
  - ❖ Heavy metal terms

100	*sabbath	26	heavy	17	riff	12	butler
97	*pantera	26	ulrich	17	leaf	12	blackened
89	*metallica	26	vulgar	17	superjoint	12	bringin
72	*leppard	25	megadeth	17	maiden	12	purple
58	metal	25	pigs	17	armageddon	12	foolin
56	hetfield	24	halford	17	gillan	12	headless
55	hysteria	24	dio	17	ozzfest	12	intensity
53	ozzy	23	reinventing	17	leps	12	mob
52	iommi	23	lange	16	slayer	12	excitable
42	puppets	23	newsted	15	purify	12	ward
40	dimebag	21	leppards	15	judas	11	zeppelin
40	anselmo	21	adrenalize	15	hell	11	sandman
40	pyromania	21	mutt	15	fairies	11	demolition
40	paranoid	20	kirk	15	bands	11	sanitarium
39	osbourne	20	riffs	15	iron	11	*black
37	*def	20	s&m	14	band	11	appice
34	euphoria	20	trendkill	14	reload	11	jovi
32	geezer	20	snowblind	14	bassist	11	anger
29	vinnie	19	cowboys	14	slang	11	rocked
28	collen	18	darrell	13	wizard	10	drummer
28	hammett	18	screams	13	vivian	10	bass
27	bloody	18	bites	13	elektra	9	rocket
27	thrash	18	unforgiven	13	shreds	9	evil
27	phil	18	lars	13	aggression	9	loud
26	lep	17	trujillo	13	scar	9	hard

Knees; Pampalk ; Widmer *Artist Classification with Web-Based Data*

formalization:: describing items :: text

- Web mining
  - ❖ playlists: Mine music playlist sites for song and artist co-occurrence

Living Room	Basement Jaxx
Wonderland (The S-Mans Dark Tribe	The Psychedelic Walto
Lisbon Acid	AFX
Flashdance (Club Mix)	Deep Dish
Love And Imitation	Infusion
Do It Now (Extended Disco Version)	Dubtribe Sound Systen
Wir Sind Die Anderen (Boris Dl	2raumwohnung
Steppin Out	Kaskade
Black Water (Vocal Mix)	Octave One
Pornography	Client
Du What Ya Du (Trentemoller Mix)	Yoshimoto

project playlist beta

Art of the Mix

musicmobs

formalization:: describing items :: text

- Lyric similarity
  - ❖ Recommend music based upon similarity of the lyrics:
  - ❖ Seed song: Led Zeppelin's Gallows Pole
    - Peter, Paul & Mary:Hangman
    - Robert Plant: Hey Joe
    - Dust for life: The End
    - The Walkabouts: Hang Man
    - Smog: Hangman blue
    - Bay Laurel: We Lost
    - Samael: Worship him

formalization:: describing items :: text

- Lyric similarity

- **Led Zeppelin: Gallows Pole:** Hangman, hangman, hold it a little while, Think I see my friends coming, Riding a many mile.
- **Peter, Paul & Mary: Hangman**: Slack your rope hangman, slack it for a while think I see my father comin' ridin' many a mile Father have you brought me hope or have you paid my fee Or have you come to see me hangin' from the gallows tree?
- **The Walkabouts: Hang Man:** Hangman take these heads from me And swing 'em from your money tree Hear me laughing in my steps These heads are yours, they're yours to keep
- **Bay Laurel: We Lost:** Our hangman will wait until the end Our hangman will smile he knows you can There is no need to let me slip by Make me feel closer... my grave
- **Samael:Worship Him:** He is the fire of vengeance He is the blade of revenge He is the executioner's axe He is the hangman's rope

formalization:: describing items :: text :: tags

ISMIR – **I**sn't **S**ocial **M**usic **I**ncredibly **R**elevant?

- Social tags
  - Collaborative Categorization
  - 'Folksonomy'
  - Some examples:
    - Del.icio.us, Flickr, LibraryThing
    - Last.fm, Qloud, MusicMobs
  - Why do people tag?
    - Personal organization

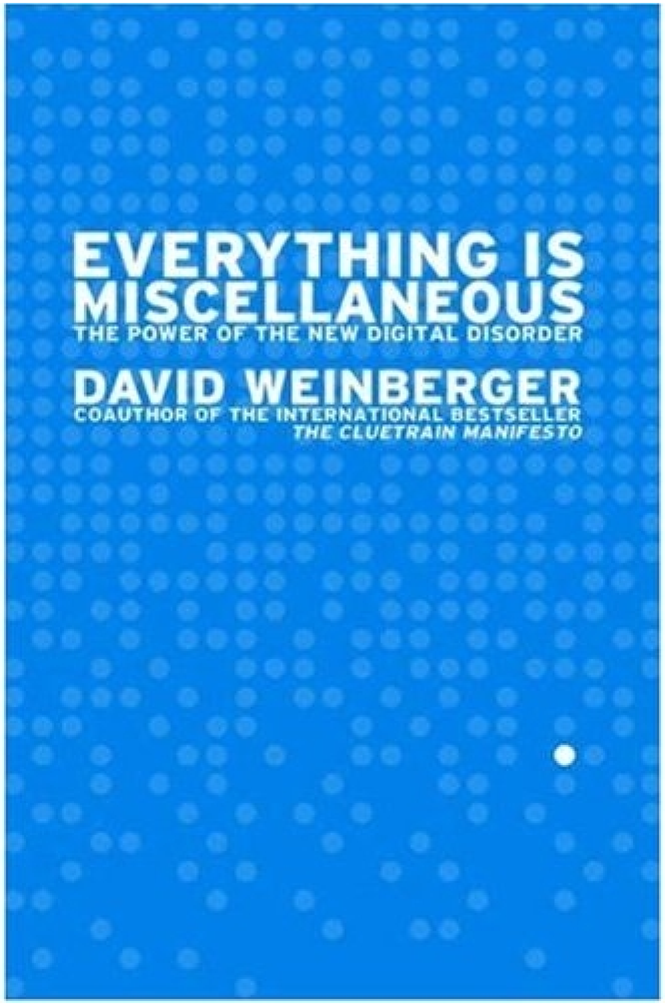

formalization:: describing items :: text :: tags

- Social tags

Browse by Tags

drums experimental instrumental punk sick drums

Deerhoof

Tell a friend about this artist

Popular Tags for This Artist

00s alternative alternative rock ambient american americana art punk art rock avant-garde california canadian classic rock downtempo drone electronic electronica energetic experimental experimental rock female vocalists folk fun funk fusion happy hip-hop indie indie pop indie rock industrial japanese jazz kill rock stars lo-fi math rock metal new wave noise noise pop noise rock noise-rock pop post rock post-punk post-rock power pop psychedelic rock punk rap rock san francisco seen live shoegaze singer-songwriter smooth soul stoner rock sweet trumpet weird

formalization:: describing items :: text :: tags

- Social tags
  - ❖ Tags – The Shins

Tag	Freq	Tag	Freq	Tag	Freq
Indie	2375	The Shins	190	Punk	49
Indie rock	1138	Favorites	138	Chill	45
Indie pop	841	Emo	113	Singer-songwriter	41
Alternative	653	Mellow	85	Garden State	39
Rock	512	Folk	85	Favorite	37
Seen Live	298	Alternative rock	83	Electronic	36
Pop	231	Acoustic	54	Love	35

Table 1: Top 21 tags applied to *The Shins*

formalization:: describing items :: text :: tags

- Social tags
  - Artist similarity based on tags for The Beatles

Top Tags	Distinctive Tags	Similar Artists via Tags
classic rock	The Beatles	John Lennon
rock	60s	Rolling Stones
pop	liverpool	Paul McCartney
british	british	The Kinks
60s	british psychedelia	The Who
oldies	oldies	Pink Floyd
psychedelic	britrock	Queen
alternative	psychedelic	The Police
indie	classic rock	Led Zeppelin
britpop	Rock and Roll	David Bowie

5628 unique tags have been applied to the Beatles

formalization:: describing items :: text :: tags

- Social tags
  - Tag similarity based on artists

Metal

- Metallica
- System of a down
- Iron Maiden
- Rammstein
- Slipknot
- In Flames
- Korn
- Pantera
- Judas Priest

Heavy Metal

- Iron Maiden
- Judas Priest
- Black Sabbath
- Manowar
- Motorhead
- Pantera
- Megadeth
- Ozzy Osbourne
- Dio

Pop

- Madonna
- The Beatles
- Black Eyed Peas
- Beach Boys
- Kelly Clarkson
- Michael Jackson
- Gwen Stefani
- Coldplay
- U2

formalization:: describing items :: text :: tags

- Social tags
  - ❖ Tag similarity based on artists

Metal

- ❖ Heavy Metal
- ❖ Death metal
- ❖ Hard Rock
- ❖ Thrash Metal
- ❖ Progressive Metal
- ❖ Rock
- ❖ Metalcore
- ❖ Seen live
- ❖ Melodic Death Metal
- ❖ Power Metal
- ❖ Gothic Metal

Pop

- ❖ Rock
- ❖ Alternative
- ❖ Female vocalists
- ❖ Indie
- ❖ Singer-Songwriter
- ❖ Classic Rock
- ❖ Favorites
- ❖ 80s
- ❖ Seen Live
- ❖ Dance
- ❖ Pop Rock

Classical

- ❖ Composers
- ❖ Clasica
- ❖ Eurite Music
- ❖ Baroque
- ❖ Classic
- ❖ Opera
- ❖ Instrumental
- ❖ Orchestral
- ❖ Piano
- ❖ Romantic
- ❖ Vivaldi

formalization:: describing items :: text :: tags

- Social tags
  - Tag clustering example

formalization:: describing items :: text :: tags

Social tags - Tag browsing: the world of metal

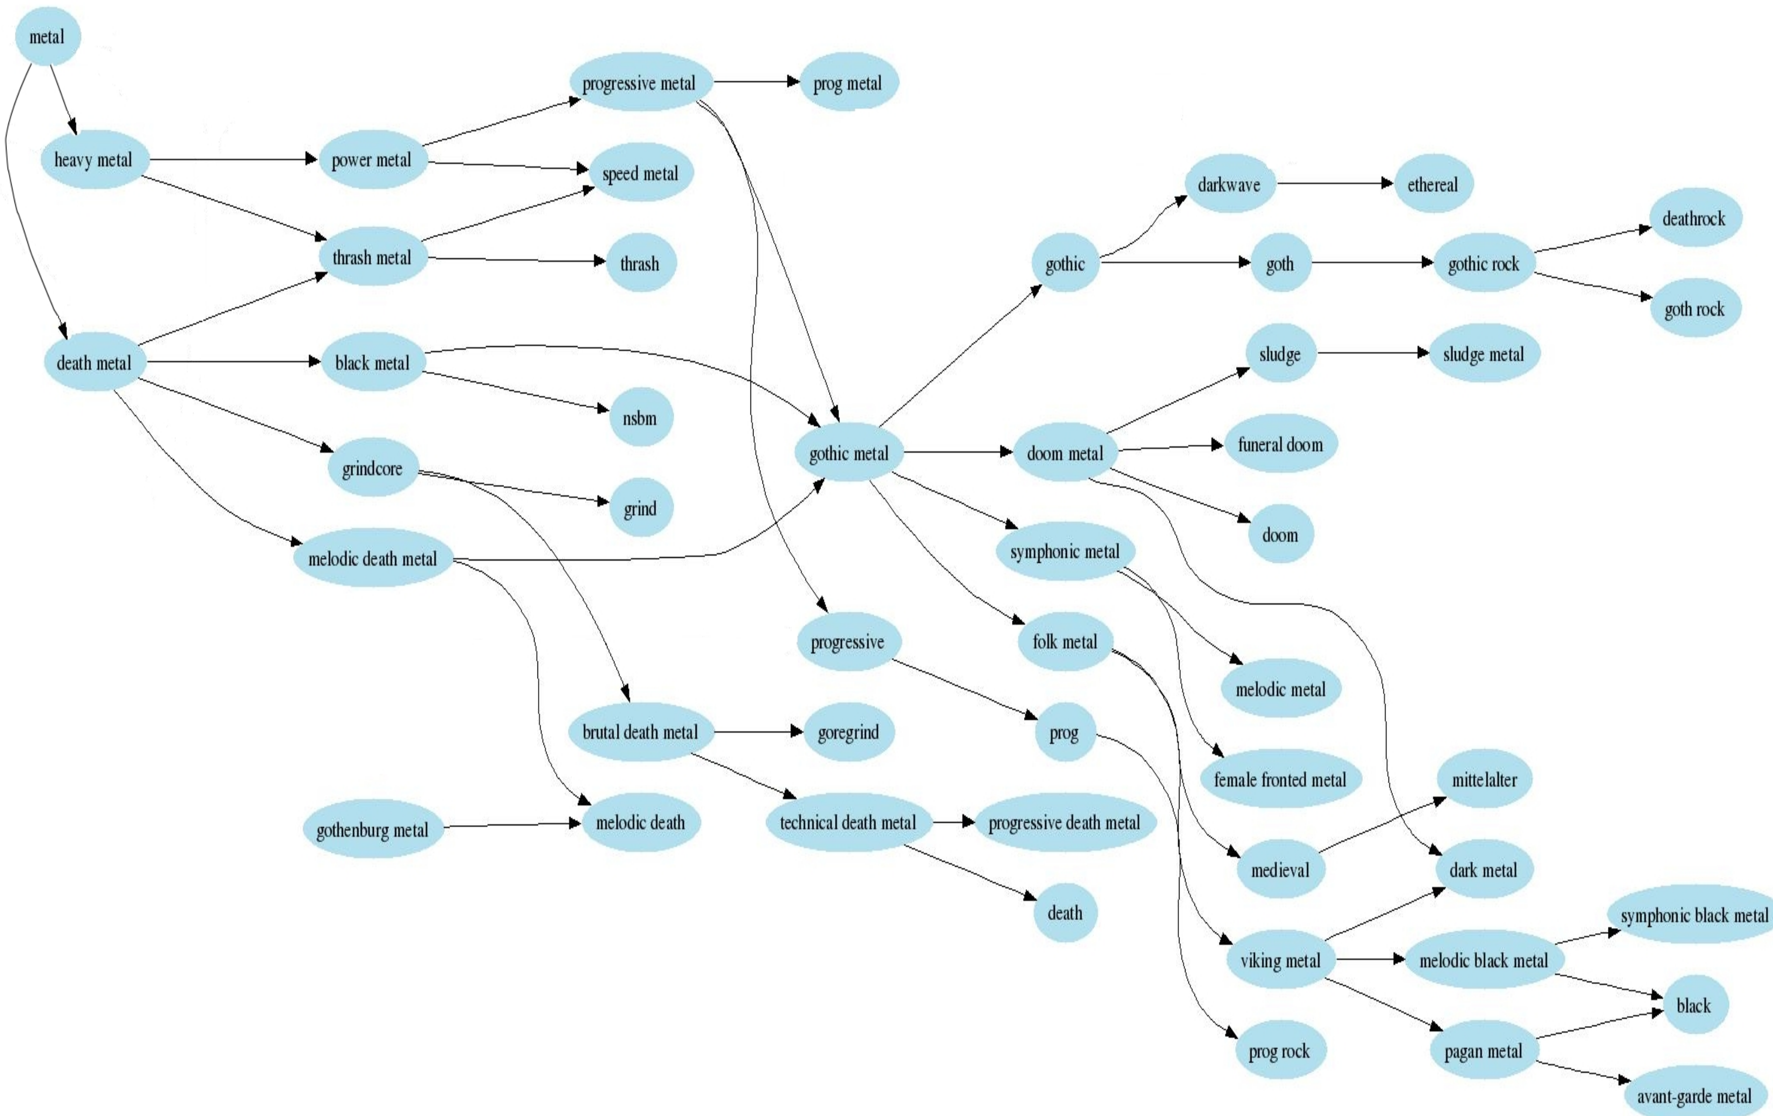

formalization:: describing items :: text :: tags

- Social tags -Faceted Tag browsing

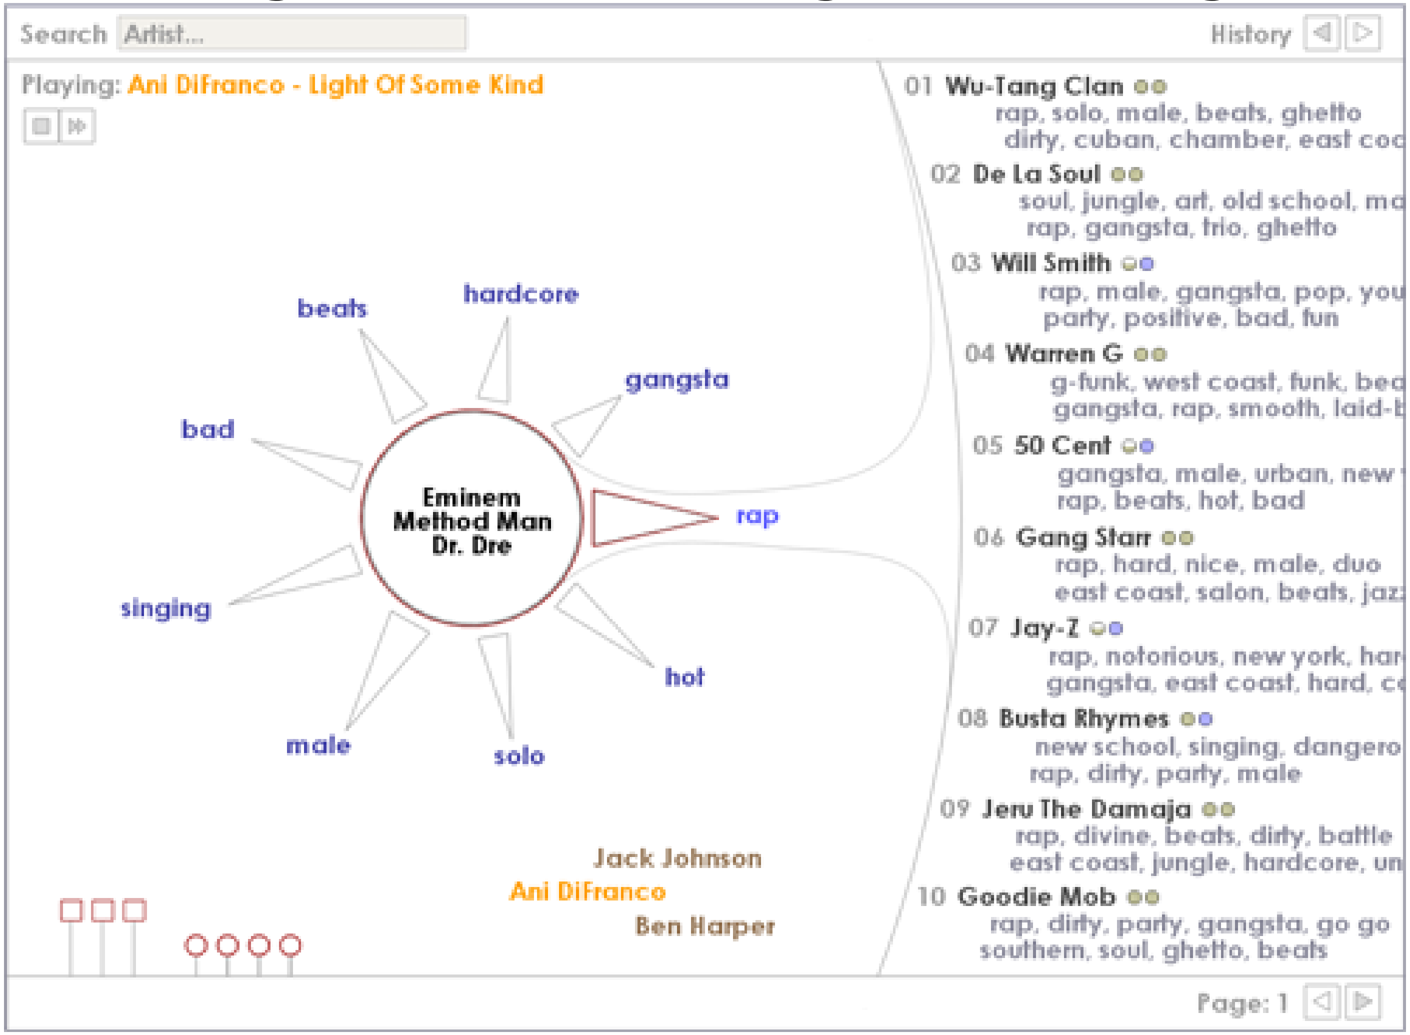

Elias Pampalk & Masataka Goto, "MusicSun: A New Approach to Artist Recommendation"

formalization:: describing items :: text :: tags

- Social tags
  - Distribution of Tags

Type	Freq	Examples
Genre	68%	Heavy metal, punk
Locale	12%	French, Seattle
Mood	5%	Chill, party
Opinion	4%	Love, favorite
Instrumentation	4%	Piano, female vocal
Style	3%	Political, humor
Misc	3%	Coldplay, composers
Personal	1%	Seen live, I own it

formalization:: describing items :: text :: tags

- Social tags: Issues
  - Polysemy
    - progressive
    - love
  - Synonymy
    - hip hop, hip-hop, hiphop, rap
  - Personal tags:
    - Seen live, I own it, Favorite
  - Noise
    - stuff a donut would like, woot, Lazy-eye

formalization:: describing items :: text :: tags

- Issues - Population bias: last.fm tags

All Music Genre	Rank	Volume
Rock	1	1
Electronica	5	0.36
Rap	9	0.21
Jazz	18	0.15
Classical	52	0.06
Blues	55	0.06
R&B	66	0.05
Country	68	0.04
World	121	0.02

Metal Tags	Rank
Metal	4
Death Metal	16
Black Metal	26
Metal Core	34
Power Metal	35
Thrash Metal	36
Progressive Metal	37
Melodic Death Metal	42
Gothic Metal	54
Doom Metal	59
Folk Metal	75
Nu Metal	83
Symphonic Metal	88
Industrial Metal	89
Viking Metal	103

Gothic Metal more popular than Country

formalization:: describing items :: text :: tags

- Social tags: Issues
  - Sparsity of data

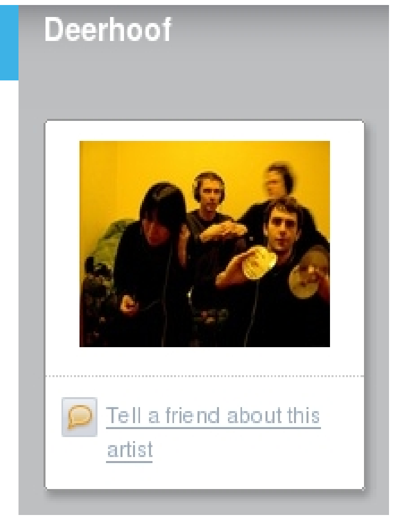

- Not enough tags for new bands
- Not enough tags at the track level
- MIR Techniques can help:
  - Learn to predict social tags
  - Apply social tags to new music
  - Related work at ISMIR2007:
    - Poster: *Autotagging Music Using Supervised Machine Learning - Douglas Eck, Thierry Bertin-Mahieux, Paul Lamere*
    - Short paper: *Annotating music collections: How content-based Similarity helps to propagate labels – Mohamed Sordo, Cyril Laurier, Oscar Celma*

formalization:: describing items :: text :: tags

- Social tags: Issues
  - Sources of tags

A Web-Based Game for Collecting Music Metadata
Michael I Mandel and Daniel P W Ellis

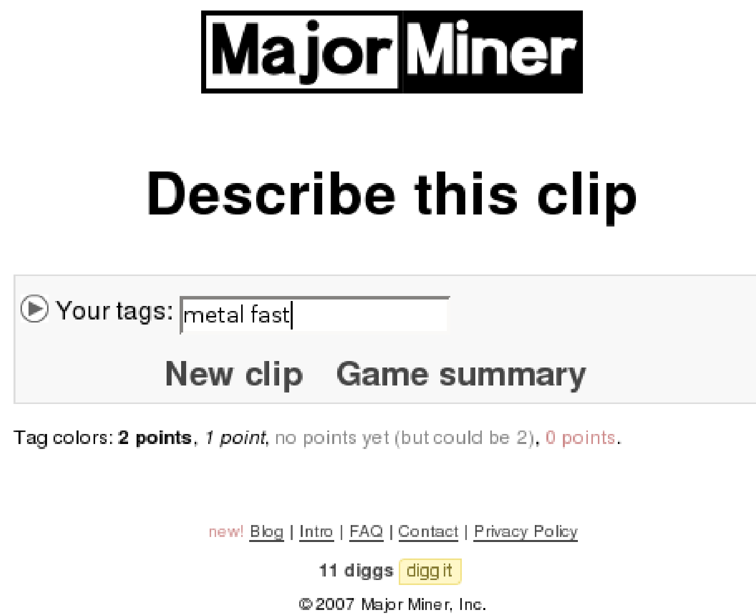

Avg 4.5 tags per clip

Identifying words that are musically meaningful
David Torres, Douglas Turnbull, Luke Barrington, and Gert Lanckriet

formalization:: describing items :: text :: tags

- Social tags: Issues
  - ❖ Hacking and Vandalism

Top Artists tagged “brutal death metal”

1	Paris Hilton	718
2	Nile	528
3	Cannibal Corpse	474
4	Suffocation	281
5	Aborted	259
6	Cryptopsy	241
7	Dying Fetus	181
8	Deicide	170
9	Devourment	166
10	Behemoth	142

CC by Metal Chris

formalization:: describing items :: text :: tags

- Social tags: Issues
  - Hacking: Paris Hilton – Raw tag counts
    - Brutal Death Metal 1145
    - atainwptiosb*: 508
    - Crap 290
    - Pop: 287
    - Officially Sh*t 248
    - Sh*t 143
    - Your ears will bleed: 140
    - emo 120
    - whore 103
    - in prison 98
    - female vocalist 80
    - whore untalented: 79
    - Best Singer in the World: 72
    - sexy 50
    - the worst thing ever to happen to music: 47
    - b*tch: 42
    - dance: 41
    - Guilty Pleasures: 40
    - Death Metal: 30
    - Female: 29
    - Slut: 29

**all things annoying in the world put together into one stupid b*tch*

formalization:: describing items :: text :: tags

- Social tags: Issues
  - Hacking: Dealing with vandals
    - Reduce influence of untrusted taggers
      - Does the tagger listen to the music they are tagging?
      - Does the tagger use the tags that they are applying?
      - Does anyone use the tags?
    - Tag Clustering
      - pop, female, sexy, guilty pleasure not often clustered with Brutal Death metal
      - Artists tagged with Brutal Death Metal are also tagged with:
        - *brutal, grind, death,*
        - *death metal, extreme metal,*
        - *gore metal, goregrind,grind, grindcore*
        - *tech technical death, technical death metal*

formalization:: describing items :: text :: tags

- Social tags: Issues
  - ❖ Hacking: last.fm strikes back!
  - ❖ Paris Hilton – Normalized Tag Counts

- Pop: 100
- Female Vocalists: 28
- Dance: 18
- American: 14
- Sexy: 13
- Brutal Death Metal: 11
- rnb: 8
- female vocalist: 8
- female: 7
- 00s: 6
- Guilty Pleasure: 6
- guilty pleasure: 6
- California: 5
- emo: 4
- Crap: 3
- Reggae: 3
- awful: 3
- party: 3
- underrated: 2
- Best Singer in the world: 2
- ataitwptiosb*: 2
- hot: 2

**all things annoying in the world put together into one stupid b*tch*

formalization:: describing items :: audio

- Manual annotation

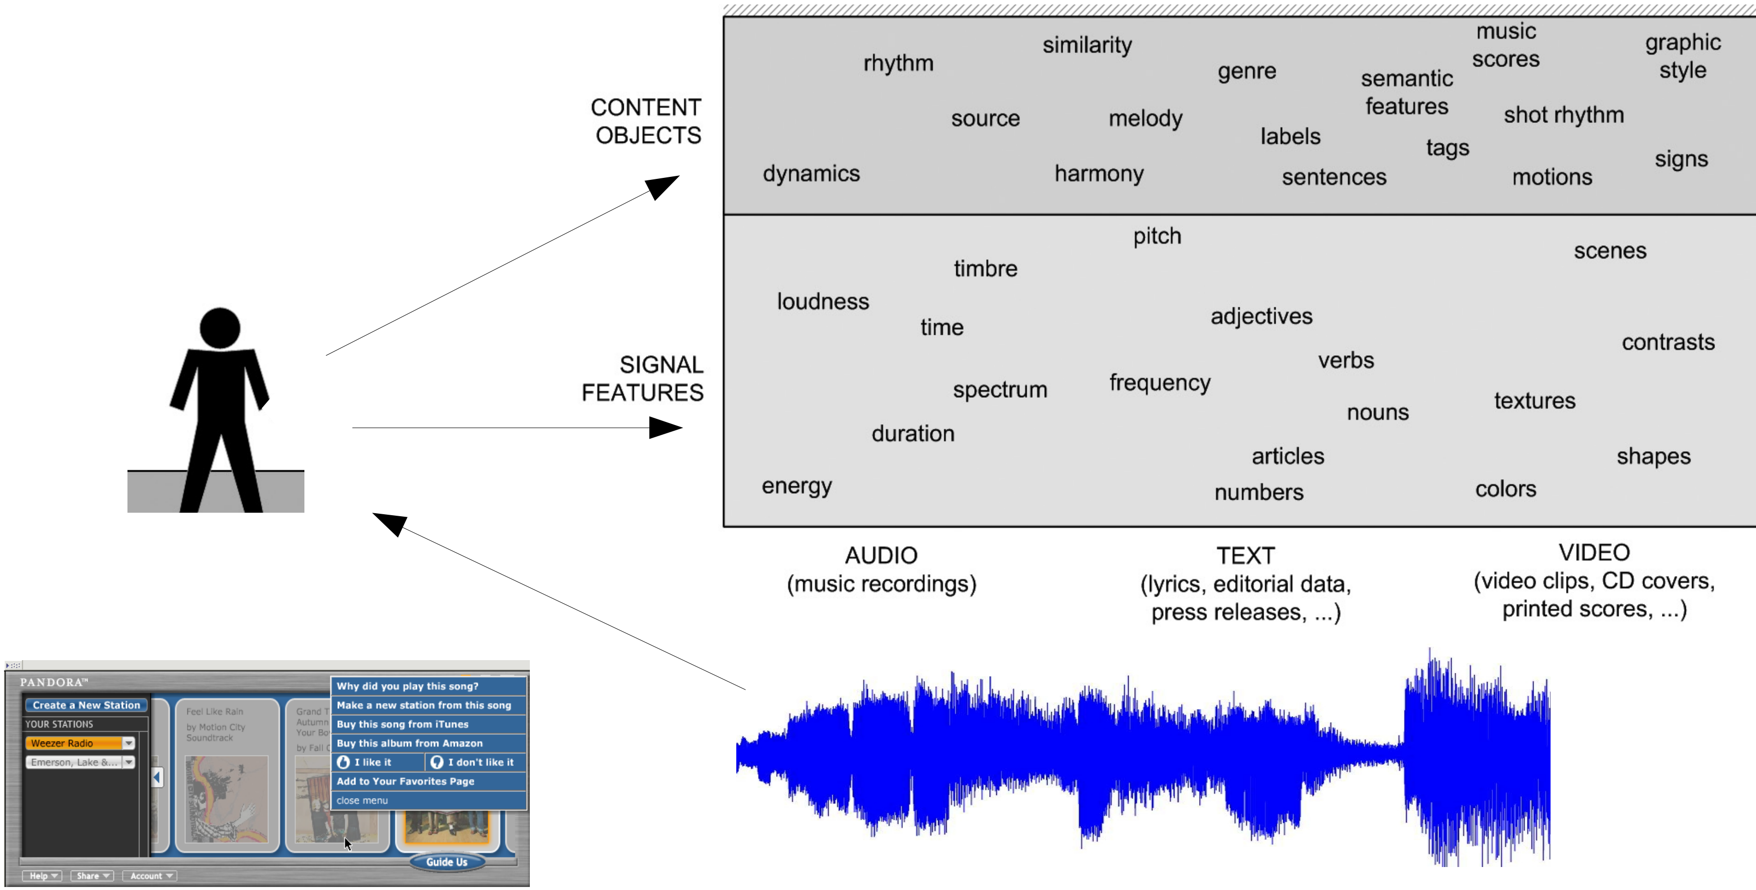

formalization:: describing items :: audio

- Manual annotation
  - Human-analysis of music:
    - Pandora
      - 50+ Musicians, 45 minutes per song
      - 400 Parameters per song, 500,000 song Catalog
    - SoundFlavor
      - '100s of parameters', 5 minutes per song
  - **But ... this doesn't scale to all music:**
    - Takes 5 years to analyze 1 year of new releases
    - Music editor becomes the Gatekeeper
    - Variability across 40 musician analysts
    - Ignore certain types of music (no 'Classical')
  - **Perhaps machines can do the job!**

formalization:: describing items :: audio

- Automatic annotation

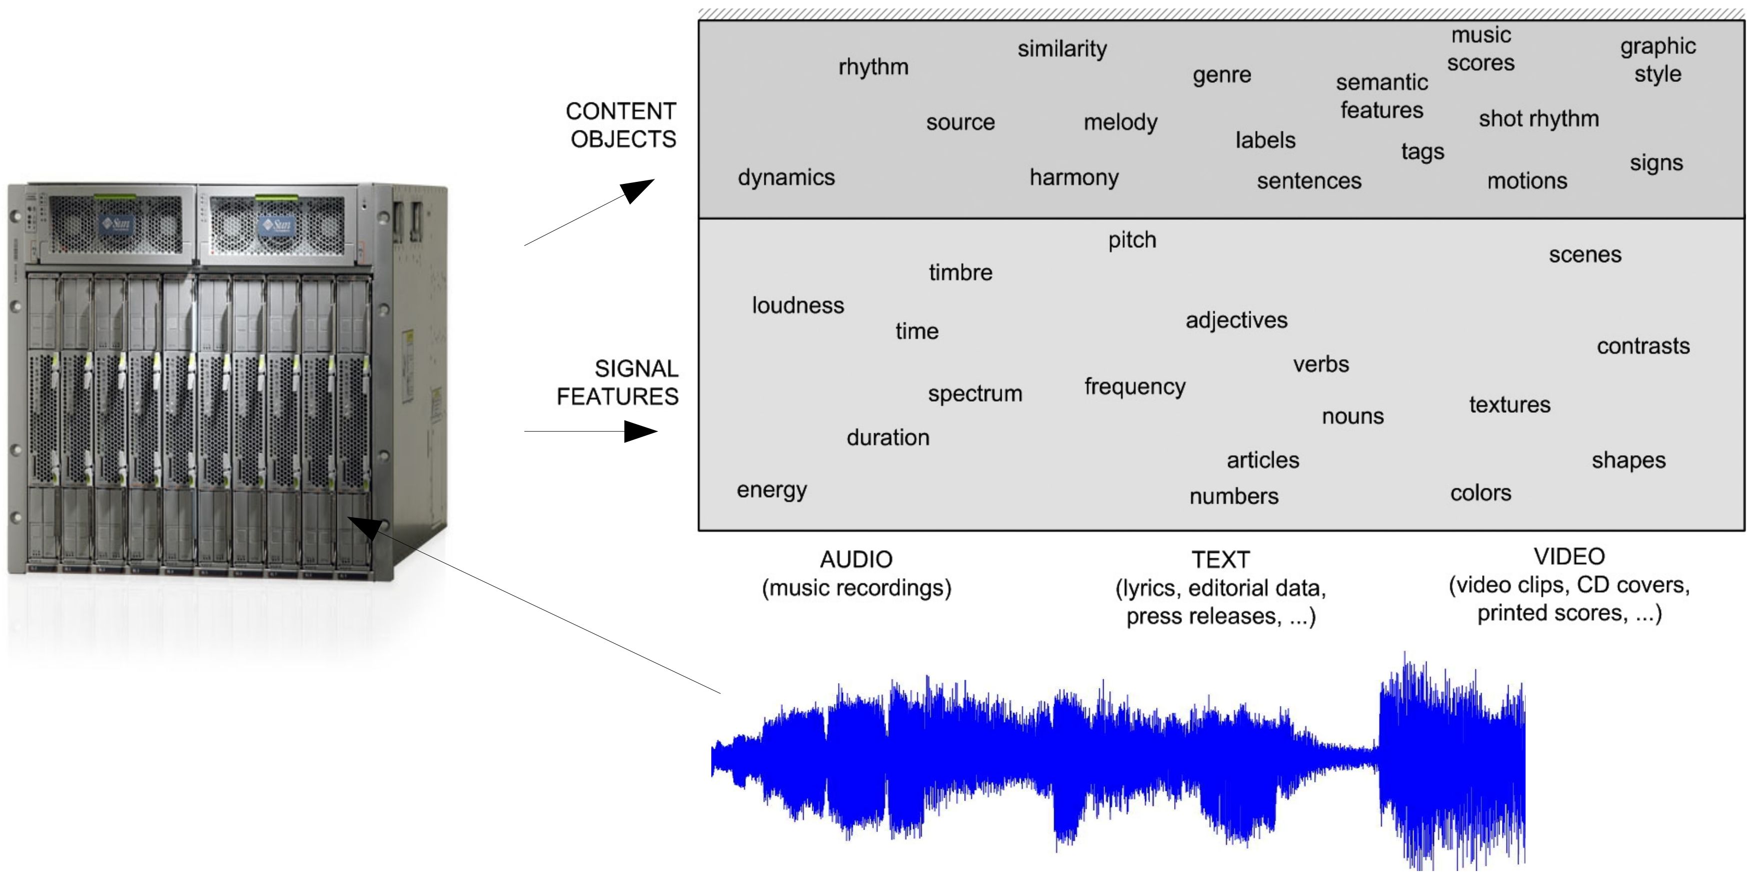

formalization:: describing items :: audio

- Automatic annotation
  - Can machines understand music?
    - **Music Similarity is hard:**
      - Hard to decide what is similar
      - Hard to evaluate
      - Start with something easier...
    - **Genre Classification**:
      - Manual : 72% (Perrot/Gjerdigen)
      - Automated (2002) 60% (Tzanetakis)
      - Automated (2005) 82% (Bergstra/Casagrande/Eck)
      - Automated (2007) 76%

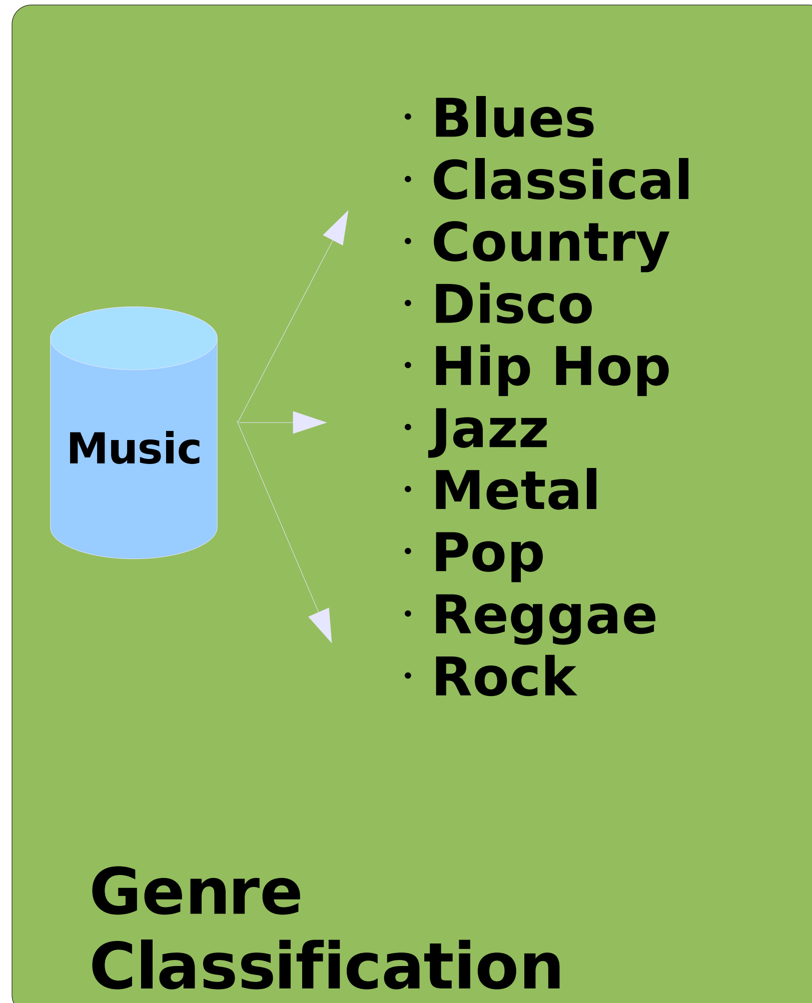

formalization:: describing items :: audio

- Automatic annotation
  - ❖ How does classification work?

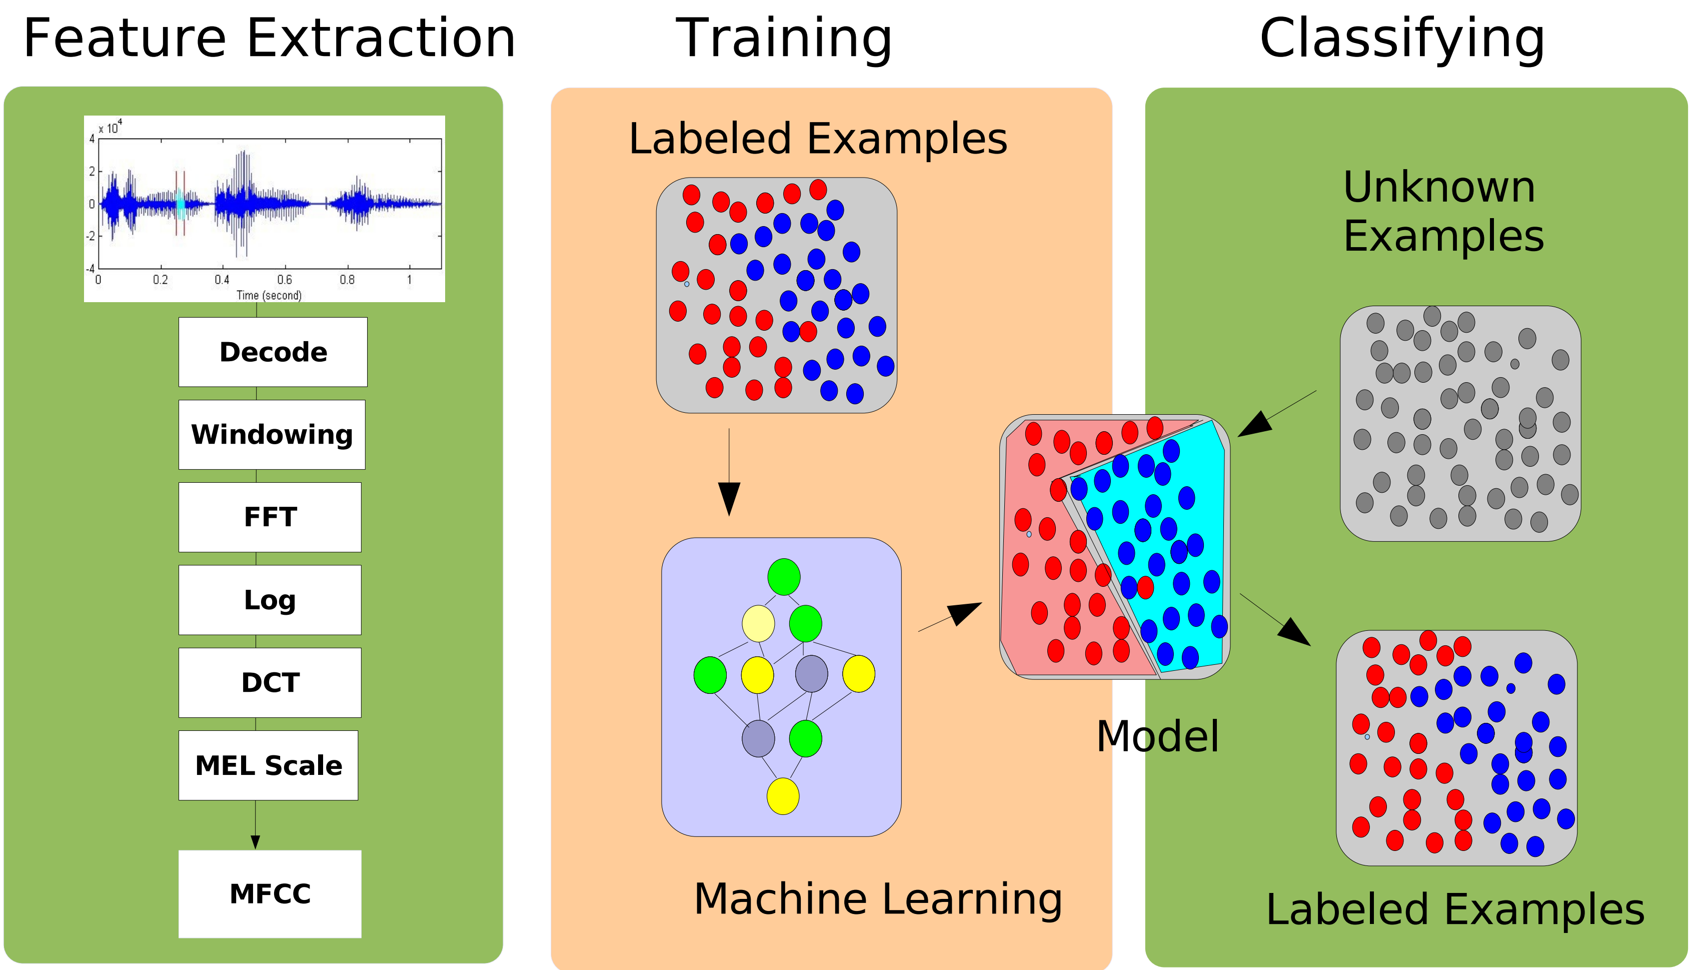

formalization:: describing items :: audio

- Automatic annotation
  - ❖ Feature extraction
    - Challenge: Too much audio data
    - Reduce audio to extract information about:
      - ❖ Pitch
      - ❖ Timbre
      - ❖ Rhythm

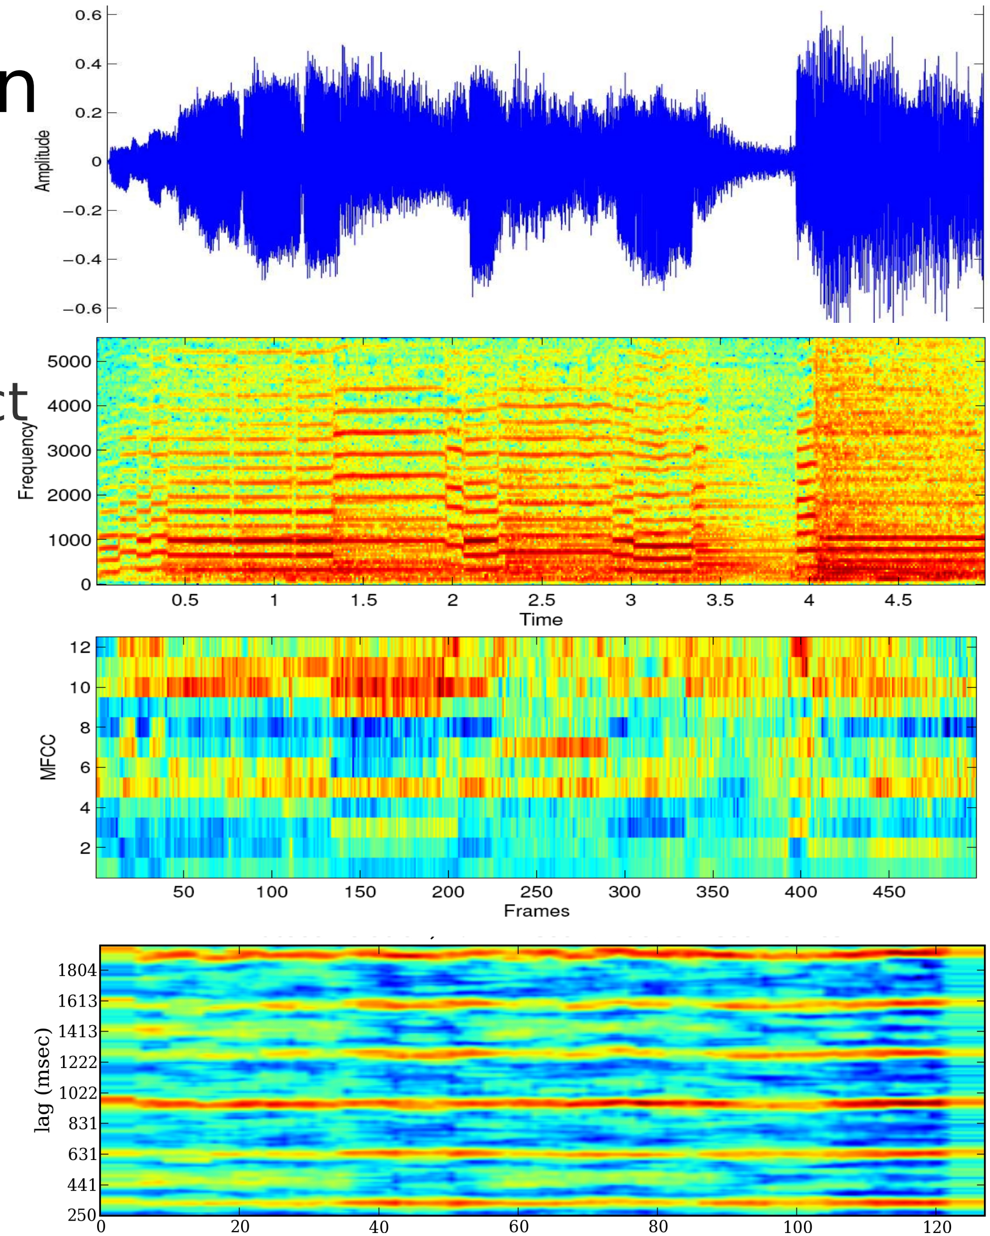

formalization:: describing items :: audio

- Automatic annotation
  - Feature extraction
    - MFCC
      - Used in speech recognition
      - Model human auditory response
      - Show rate of change in the different spectrum bands
      - Good for Timbre

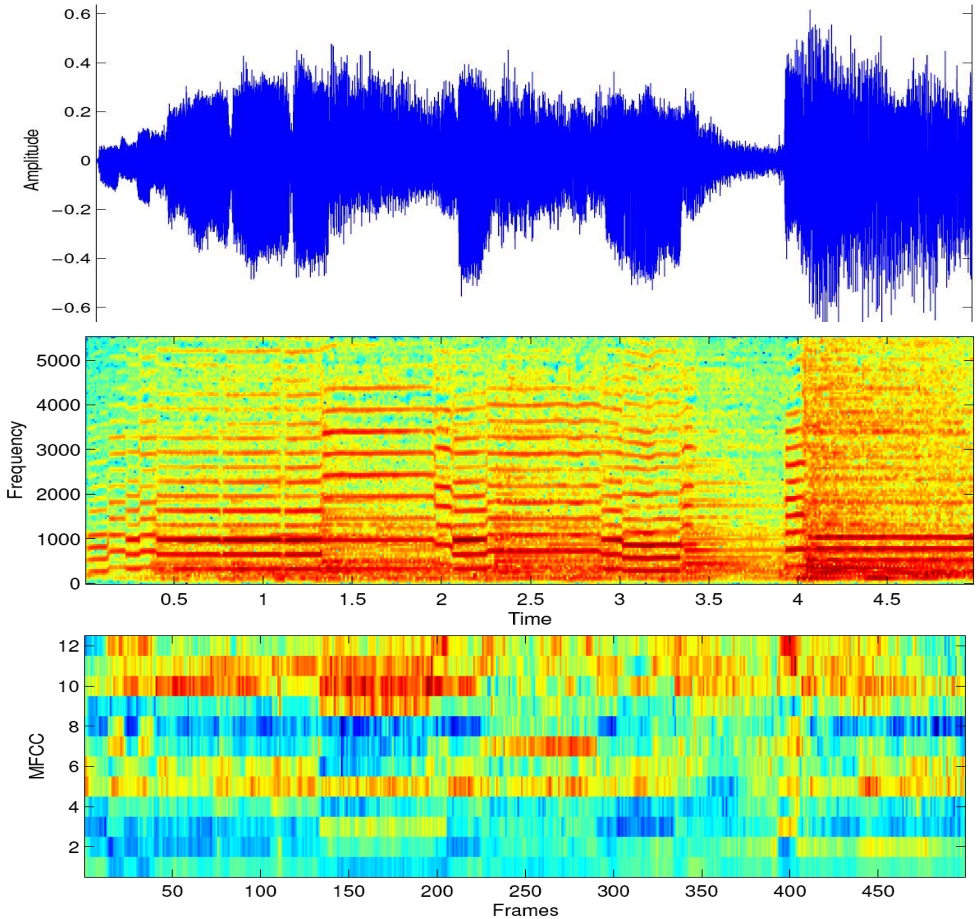

formalization:: describing items :: audio

- Automatic annotation
  - ❖ Feature extraction
    - Log Spectogram
      - ❖ Retains pitch info
    - Useful for:
      - ❖ Key identification
      - ❖ Mode identification
      - ❖ Mood classification
      - ❖ Style identification

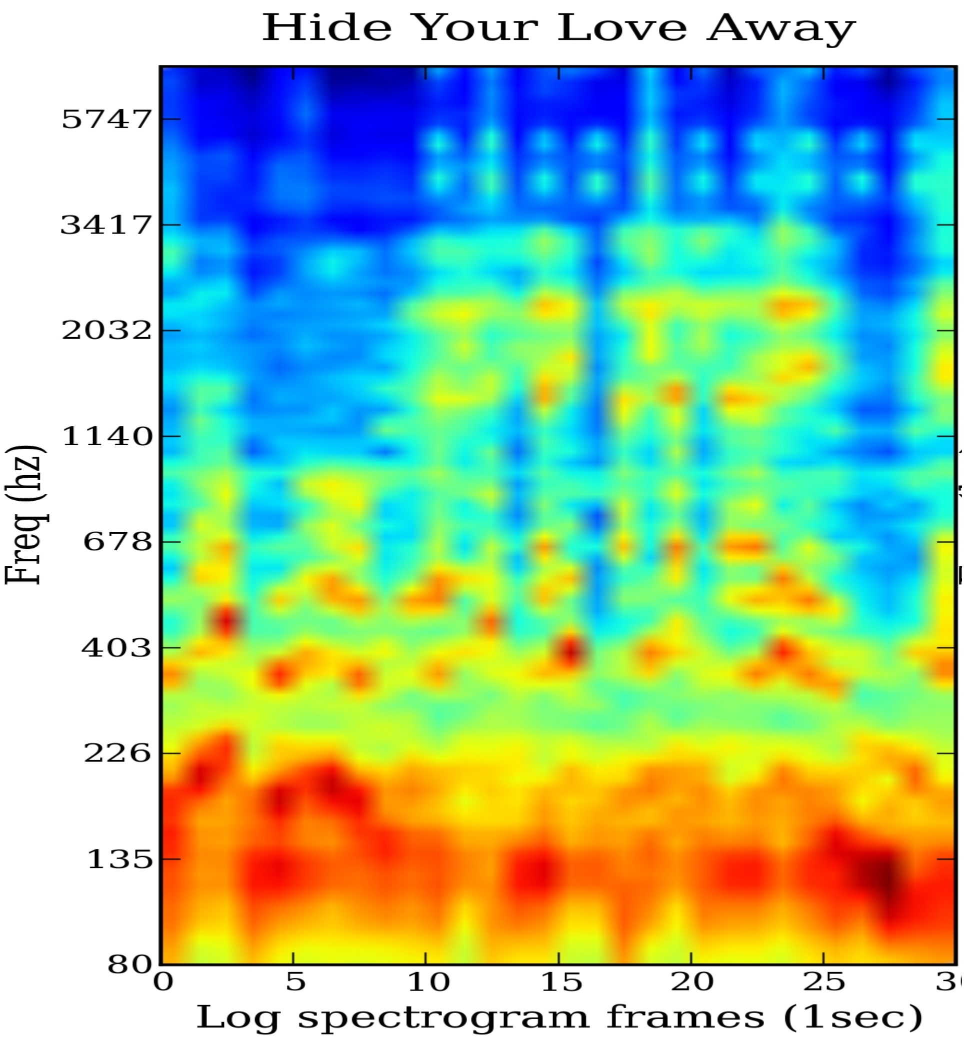

formalization:: describing items :: audio

- Automatic annotation
  - ❖ Feature extraction
    - Autocorrelation
      - ❖ Represents Timing information
    - Useful for:
      - ❖ Rhythm
      - ❖ Time signature
      - ❖ Tempo
      - ❖ Tempo drift

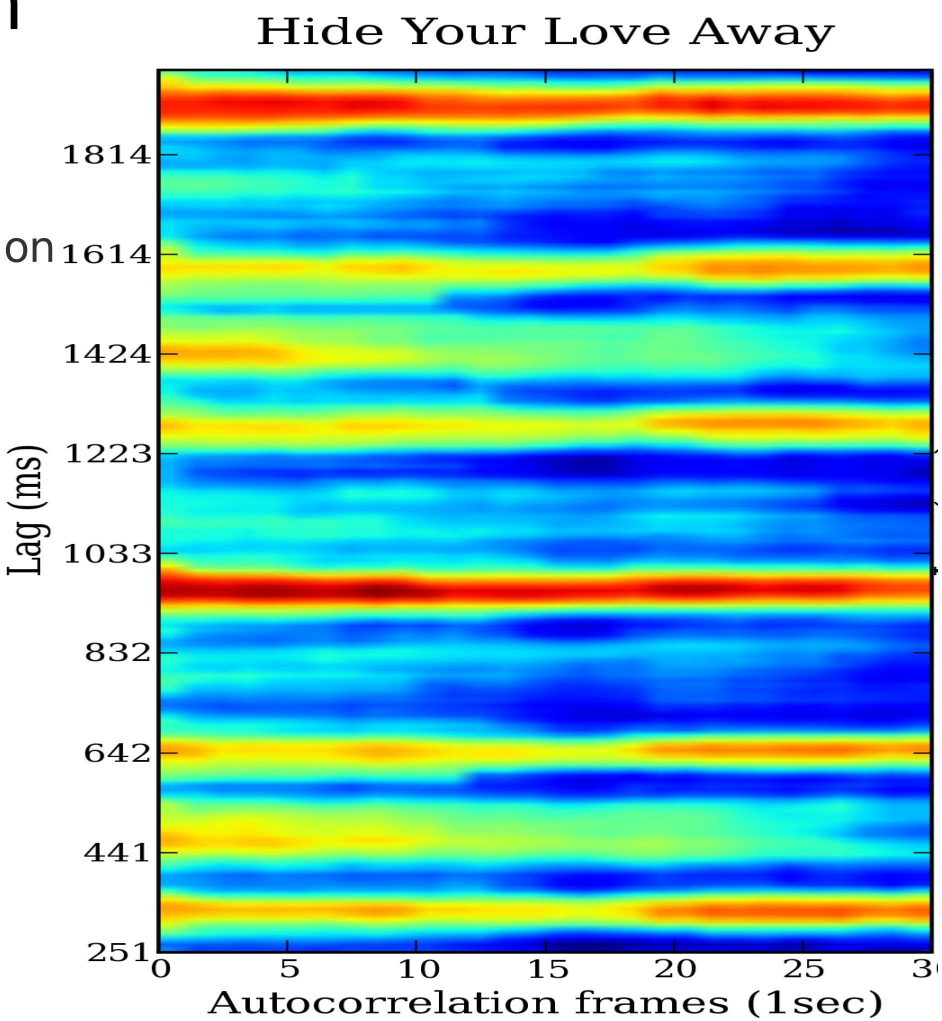

formalization:: describing items :: audio

- Automatic annotation
  - Feature extraction: Summary

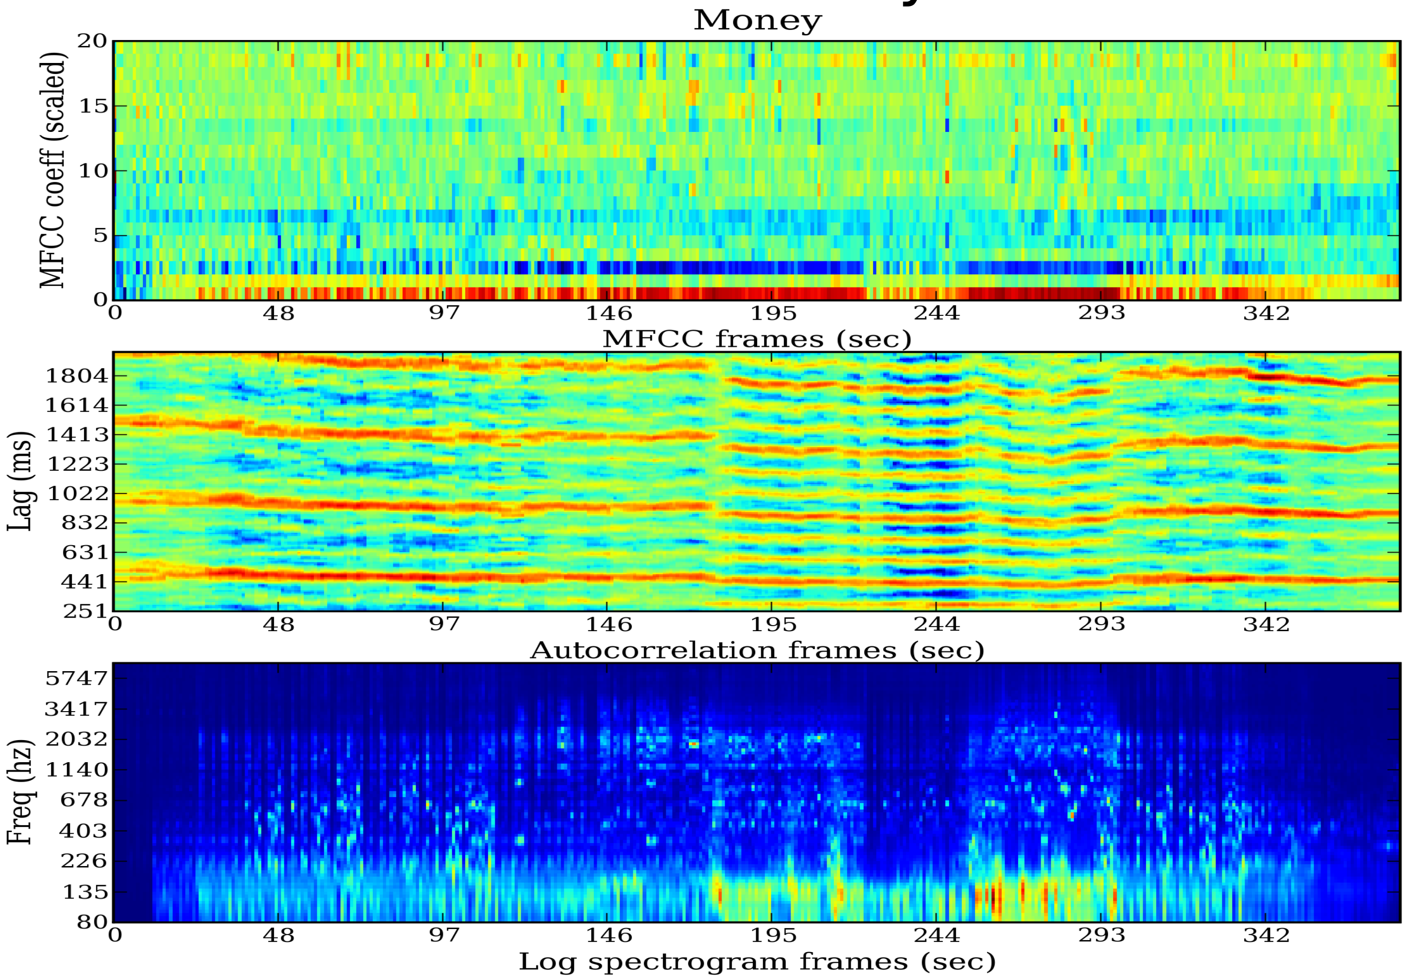

formalization:: describing items :: audio

- Automatic annotation
  - ❖ Machine learning
    - Statistical modeling
    - Exploit regularities in the data
    - Generalize to previously *unseen* examples
    - Predict without overfitting

Supervised learning example

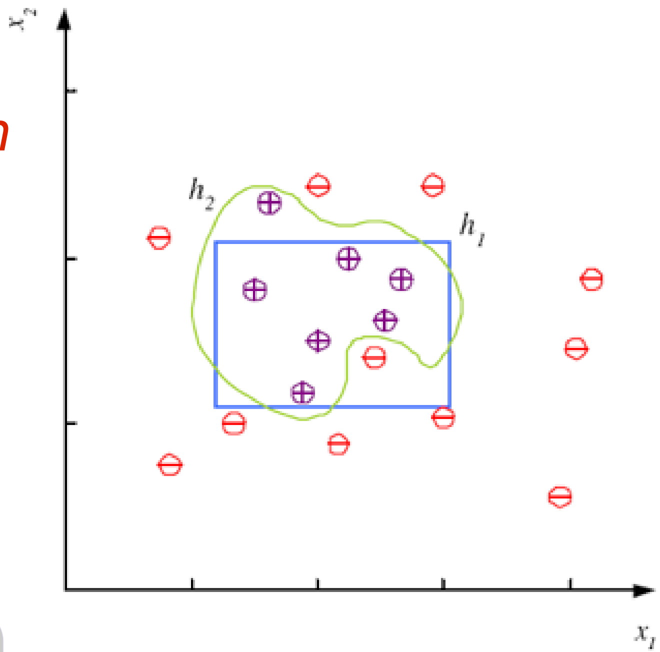

From Alpaydin (2004)

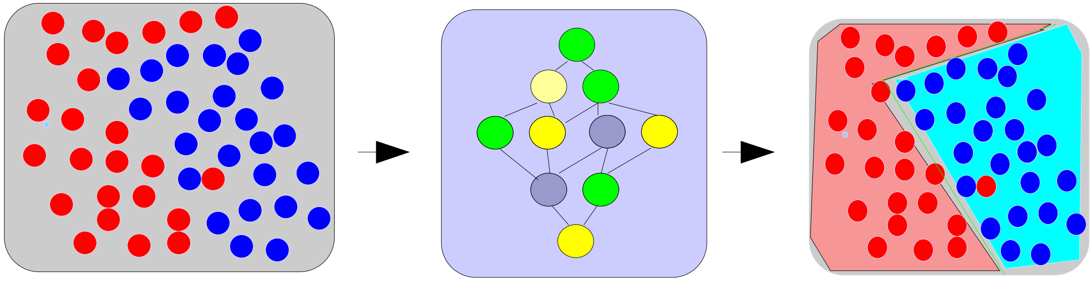

formalization:: describing items :: audio

- Automatic annotation
  - Multiple class machine classification

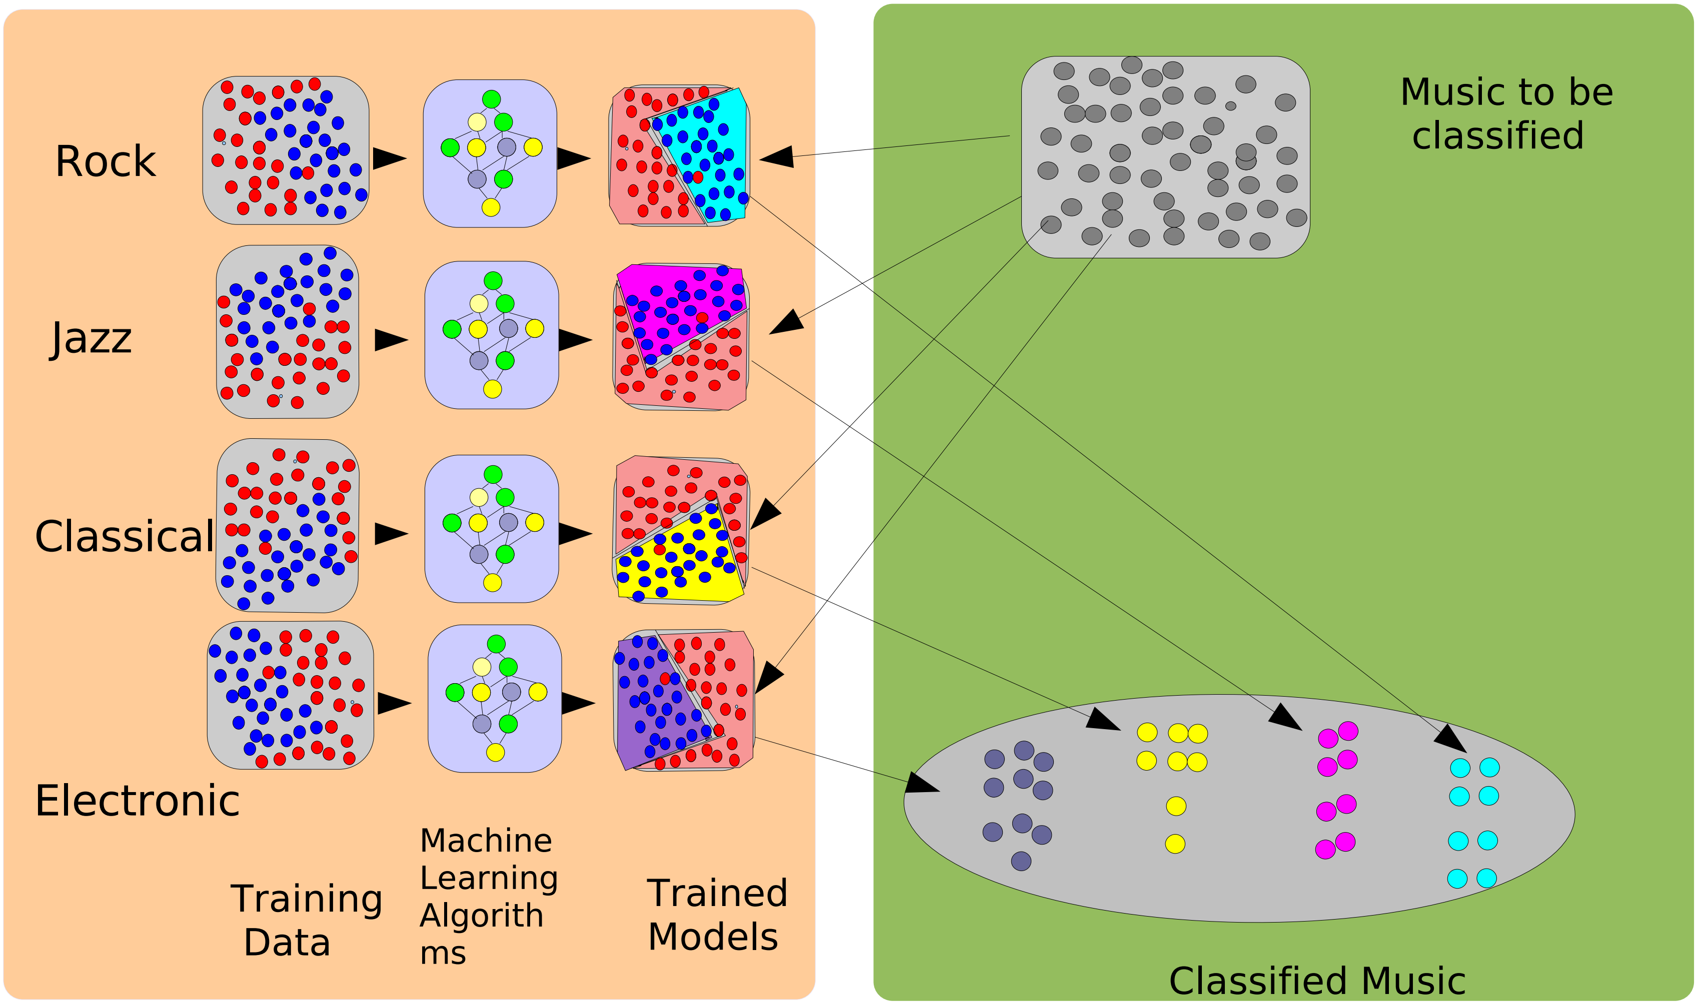

formalization:: describing items :: audio

- Automatic annotation
  - Similarity based on classification

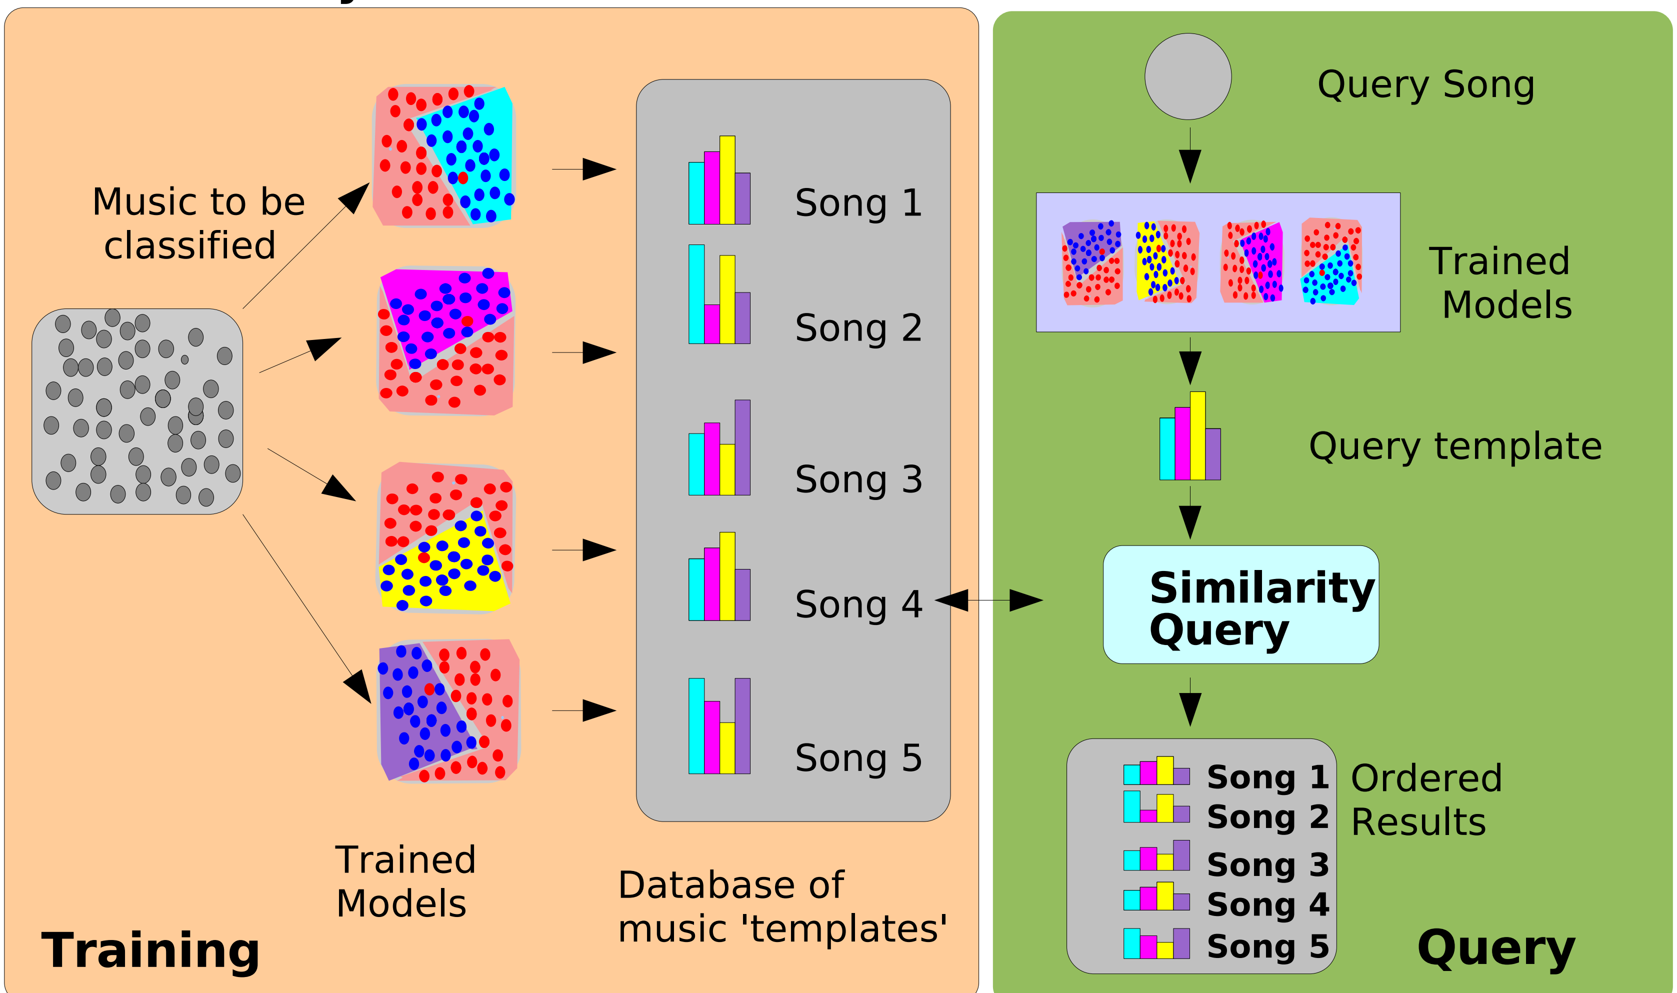

formalization:: describing items :: audio

- Automatic annotation
  - Content analysis: State of the art
    - Machines more accurate for simple tasks
    - Still early days for automated music similarity
    - **Time per million songs:**
      - Manual: with 100 people = 3 Years
      - Automatic: with 100 CPUs = 8 Hours
    - **Cost per million songs**
      - Manual: ~ $10,000,000
      - Automatic: ~ $1,000

outline

- Introduction
- Formalization of the recommendation problem
- **Recommendation algorithms**
- Problems with recommenders
- Recommender examples
- Evaluation of recommenders
- Conclusions / Future

music recommendation:: algorithms

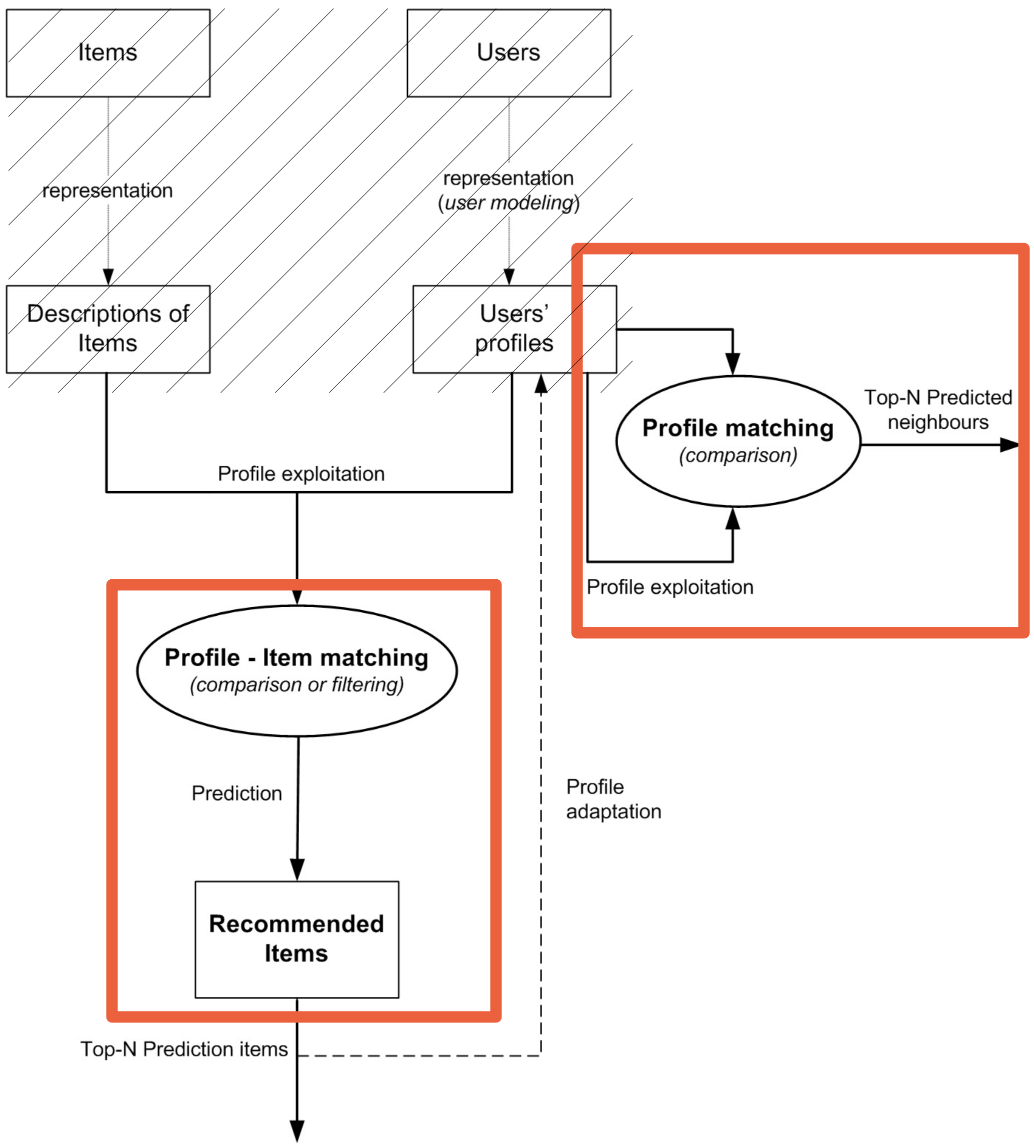

music recommendation:: algorithms

- elements
  - transactional dataset
  - user-item interaction
    - explicit data
      - rating
      - purchase
      - relevance feedback (e.g love & ban this song/artist)
      - etc.
    - implicit data
      - listen to (play / stop / skip)
      - time spent in a webpage
      - etc.

music recommendation:: algorithms

- general approaches
  - user-based
    - compute top-N neigbours for a given user
      - similarity measure
      - clustering
    - recommend items from the user's neigbourhood
  - item-based
    - compute item similarity
      - ratings / num. plays: linear regression, cosine, pearson correlation, adjusted cosine
      - content based: EMD, Manhattan distance, MFCC/GMM, etc.
    - recommend items to a user, based on her profile

music recommendation:: algorithms

- general approaches
  - model-based
    - create a model based on the user profile
      - probabilistic models (three way aspect model) [Yoshii, 2006], [Yoshii, 2007]
      - decision trees
      - neural networks
      - etc.
    - recommend items based on the user's model
      - Usually, recommendation seen as a classification problem

music recommendation:: algorithms

- Expert
- Demographic Filtering
- Collaborative Filtering
- Content-based Filtering
- Hybrid methods

music recommendation:: algorithms :: expert

- Expert
  - AMG editors
    - Genre
    - Styles
    - Moods
    - Themes
    - Similar artists

- eMusic in 2005-2006

  expands its editors staff to 120

  2nd licensed music download service, on a specialized market

emusic.

music recommendation:: algorithms :: expert

- Expert
  - AMG Tapestry, a playlist generator based on
    - Styles,
    - Moods,
    - Themes, and
    - tempo, dynamics, instrumentation, etc.

music recommendation:: algorithms :: expert

- Expert
  - rely on *experts* to recommend music
    - metadata by editors (Allmusic, eMusic, etc.)
    - expert reviews (pitchforkmedia, rollingstone, etc.)
    - mp3 blogs (hypem.com)
  - Pros
    - transparency of the recommendations
    - can differentiate between “good and bad” music, according to the expert
  - Cons
    - not personalized
    - limited coverage
    - no scaling

music recommendation:: algorithms :: demogr

- Demographic Filtering

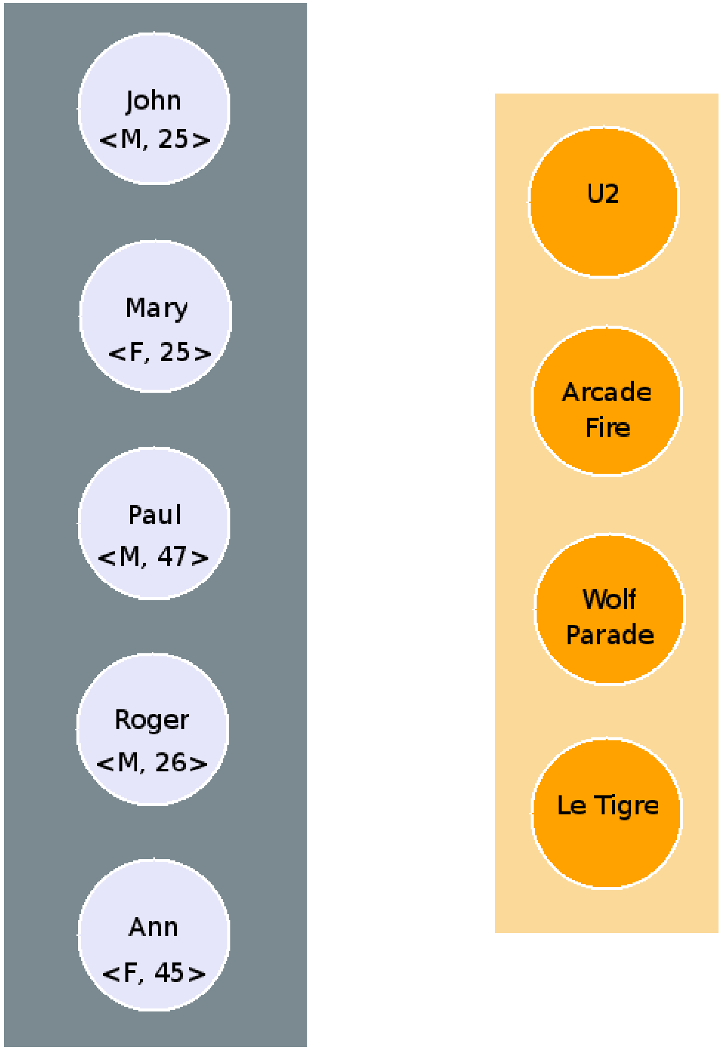

music recommendation:: algorithms :: demogr

- Demographic Filtering

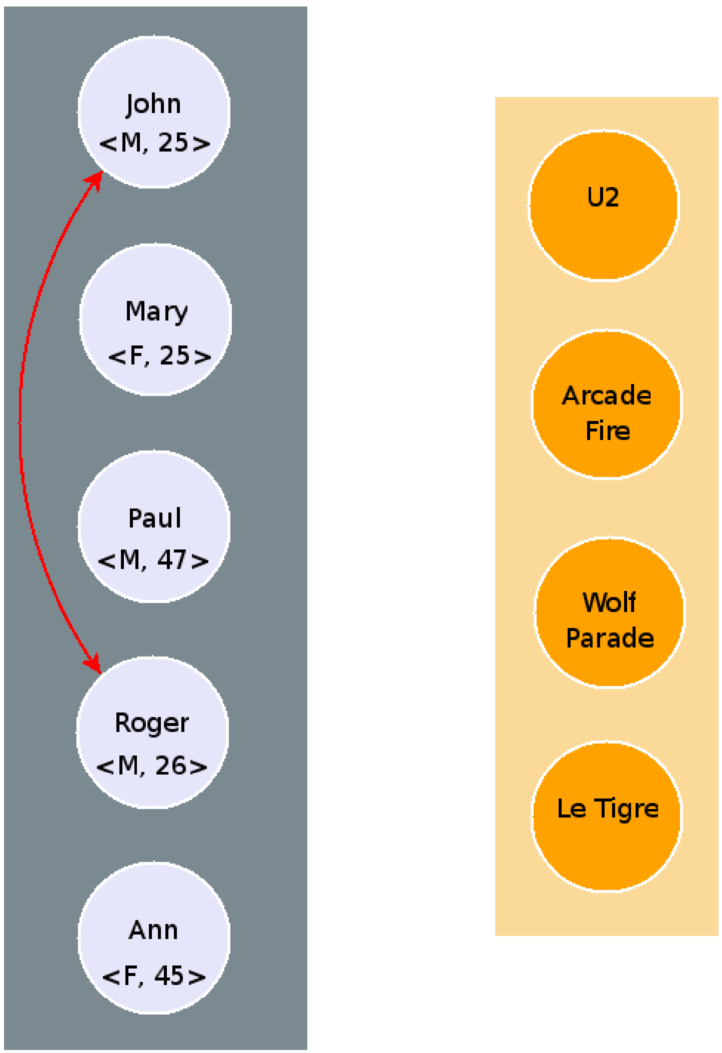

- **John** and **Roger** have similar profiles (gender, age)

music recommendation:: algorithms :: demogr

- Demographic Filtering

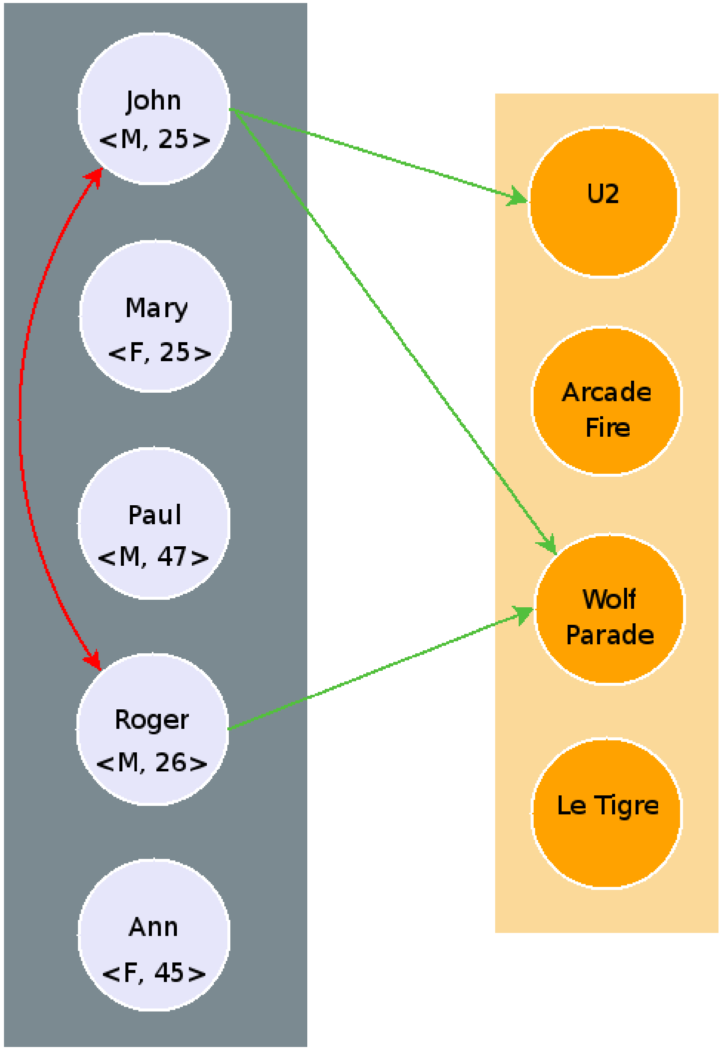

- Analyse **John** and **Roger** preferences for the artists

music recommendation:: algorithms :: demogr

- Demographic Filtering

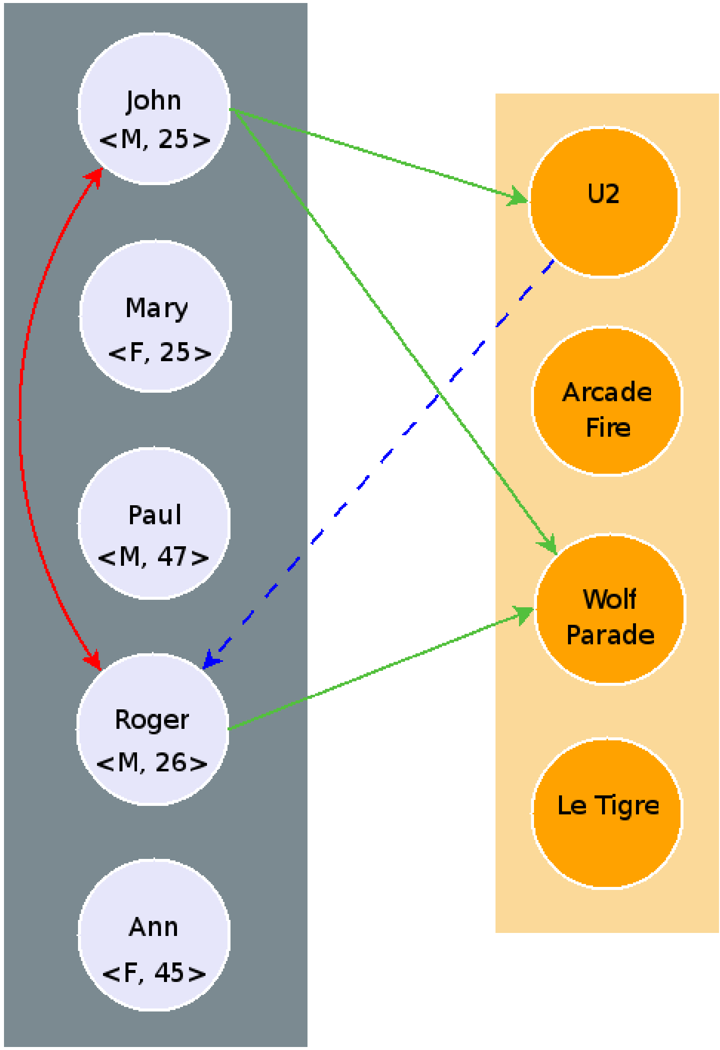

- Recommend **U2** to **Roger**

music recommendation:: algorithms :: demogr

- Demographic Filtering
  - Process
    - 1) find users with similar features
      - define similarity function among users
      - clustering based on the similarity distance
    - 2) recommend items preferred by similar users
      - prediction based on weighted average
  - Pros
    - avoids user cold-start problem (more later on...)
  - Cons
    - totally dependant on user's features (sometimes unknown / private / ...)
    - not personalized recommendations

music recommendation:: algorithms :: colfilter

- Collaborative Filtering
  - approaches
    - user-based
      - “recommend items from like-minded people”
    - item-based
      - Amazon example “people who buy this also bought that”
    - model-based
      - model the user behavior using bayesian network, clustering, association rules, neural networks, etc.

music recommendation:: algorithms :: colfilter

- Collaborative Filtering: User-based [Shardanand, 1995]

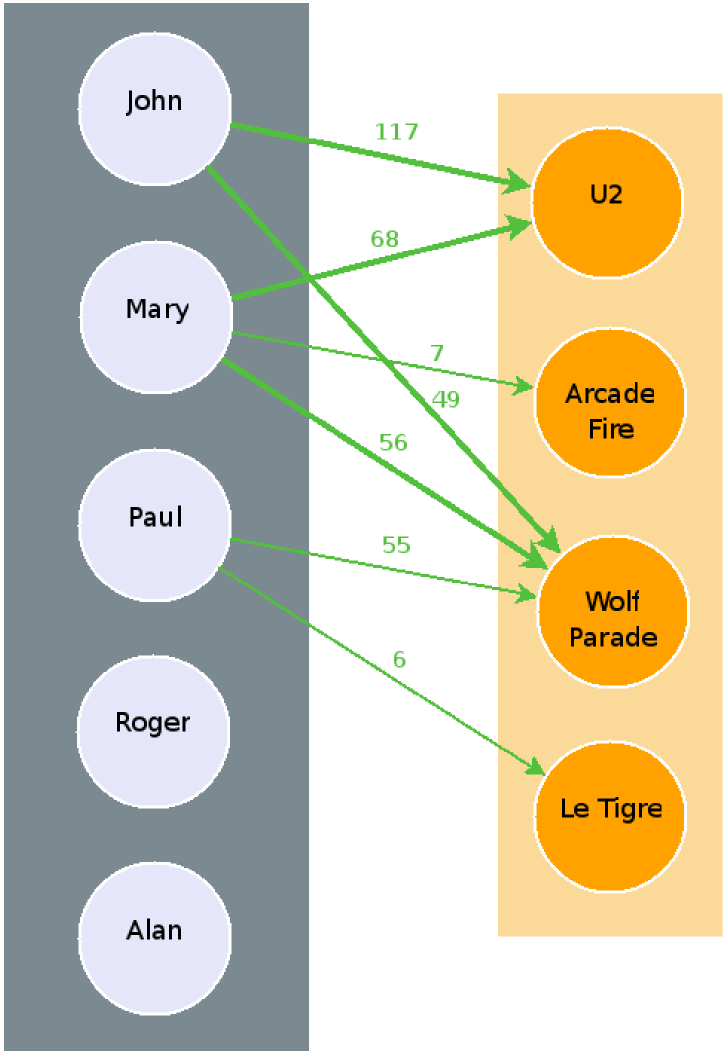

music recommendation:: algorithms :: colfilter

- Collaborative Filtering: User-based

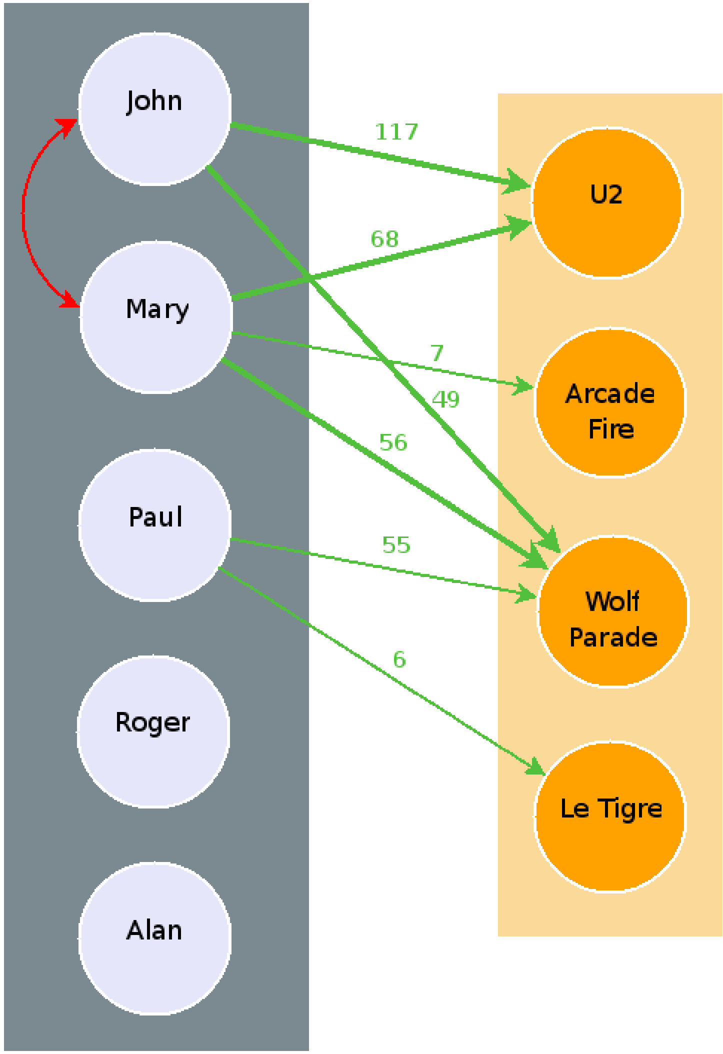

- **John** and **Mary** have similar listening habits

music recommendation:: algorithms :: colfilter

- Collaborative Filtering: User-based

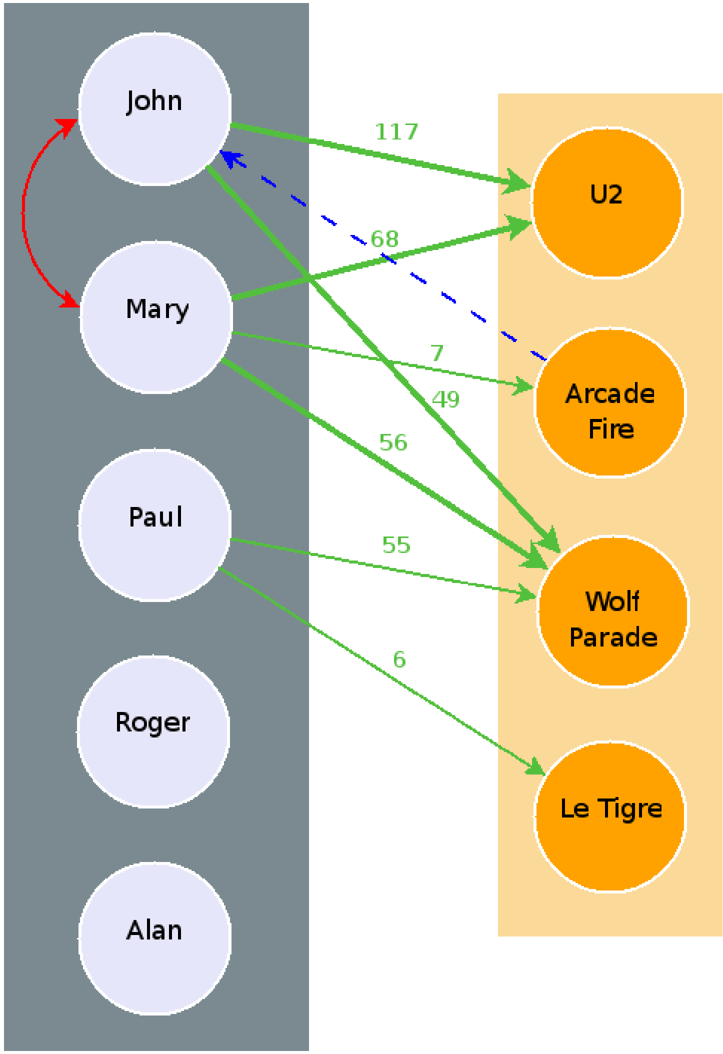

- Recommend **Arcade Fire** to **John**

music recommendation:: algorithms :: colfilter

- Collaborative Filtering: User-based
  - ❖ user-item matrix

	i_1	i_2			i_j		i_n
u_1					4		
u_2					Φ		
					4		
u_a					?		
					2		
					1		
u_m					Φ		

  - ❖ prediction (user U_a, item i_j): adjusted weighted sum

$$P_{a,j} = \bar{R}_a + \sum_{u \in Neighbours(u_a)} sim(u_a, u)(R_{u,j} - \bar{R}_u)$$

music recommendation:: algorithms :: colfilter

- Collaborative Filtering: Item-based [Sarwar, 2001]

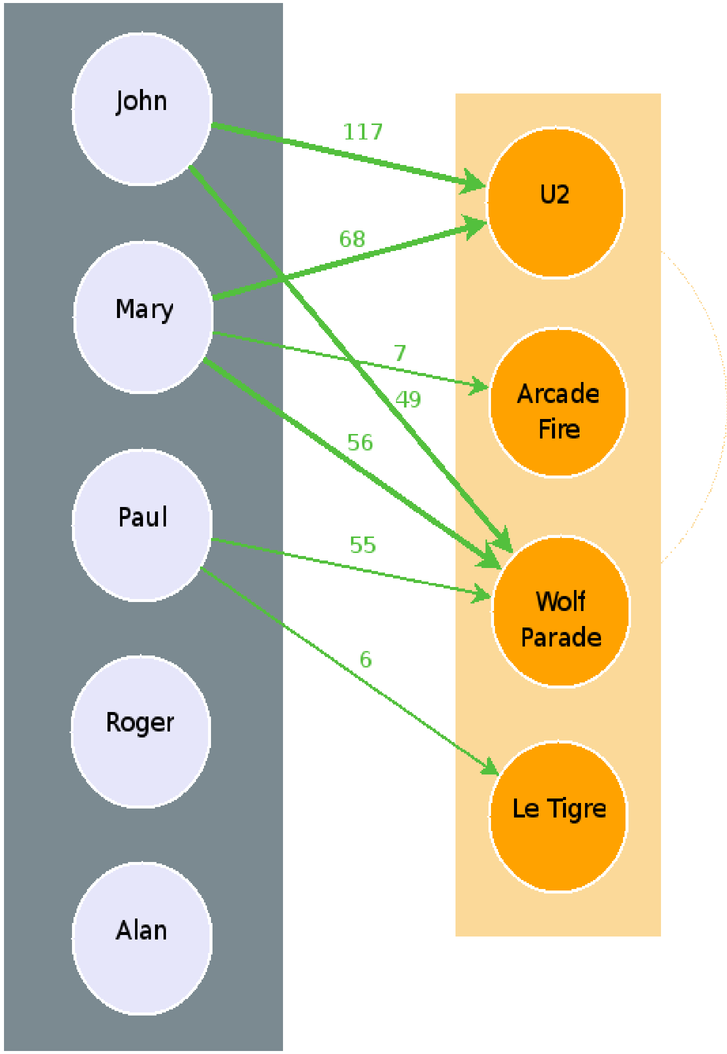

music recommendation:: algorithms :: colfilter

- Collaborative Filtering: Item-based

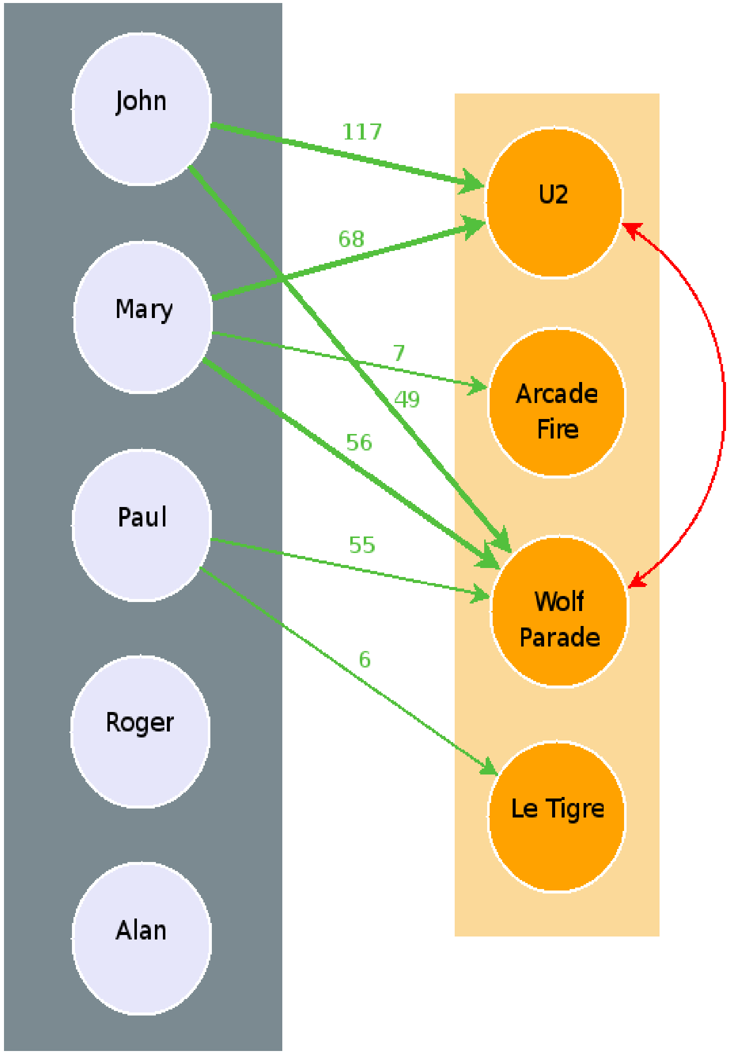

- **People** who listen to **U2** listen to **Wolf Parade**, too

music recommendation:: algorithms :: colfilter

- Collaborative Filtering: Item-based

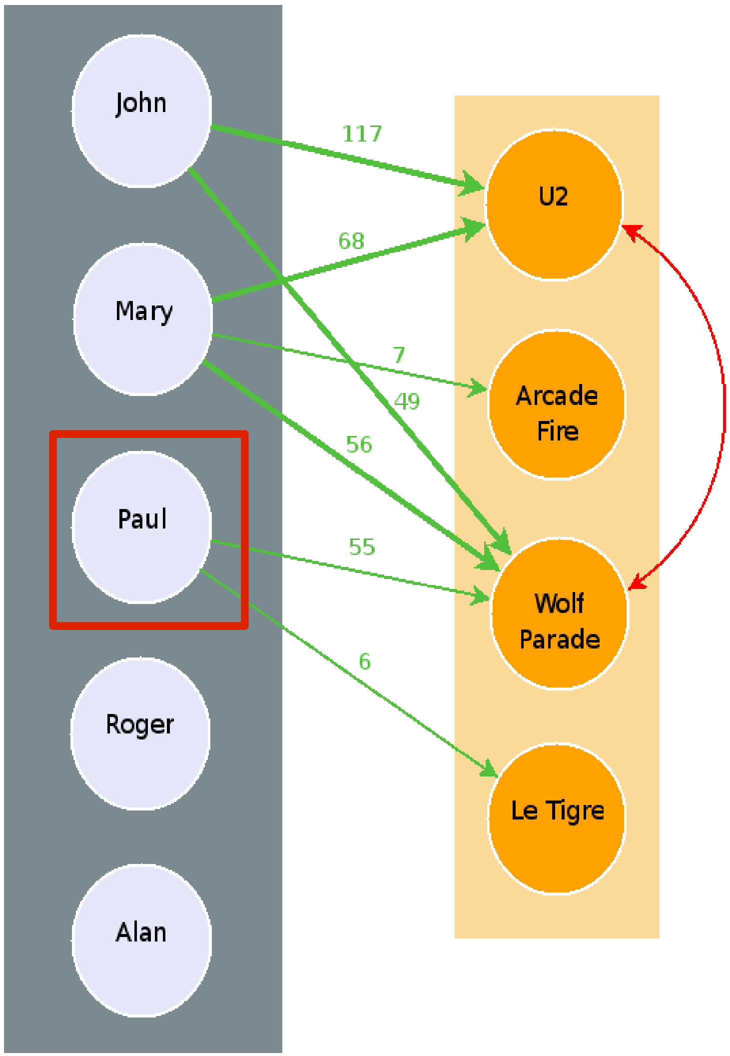

- **Paul** listens to **Wolf Parade**, then...

music recommendation:: algorithms :: colfilter

- Collaborative Filtering: Item-based

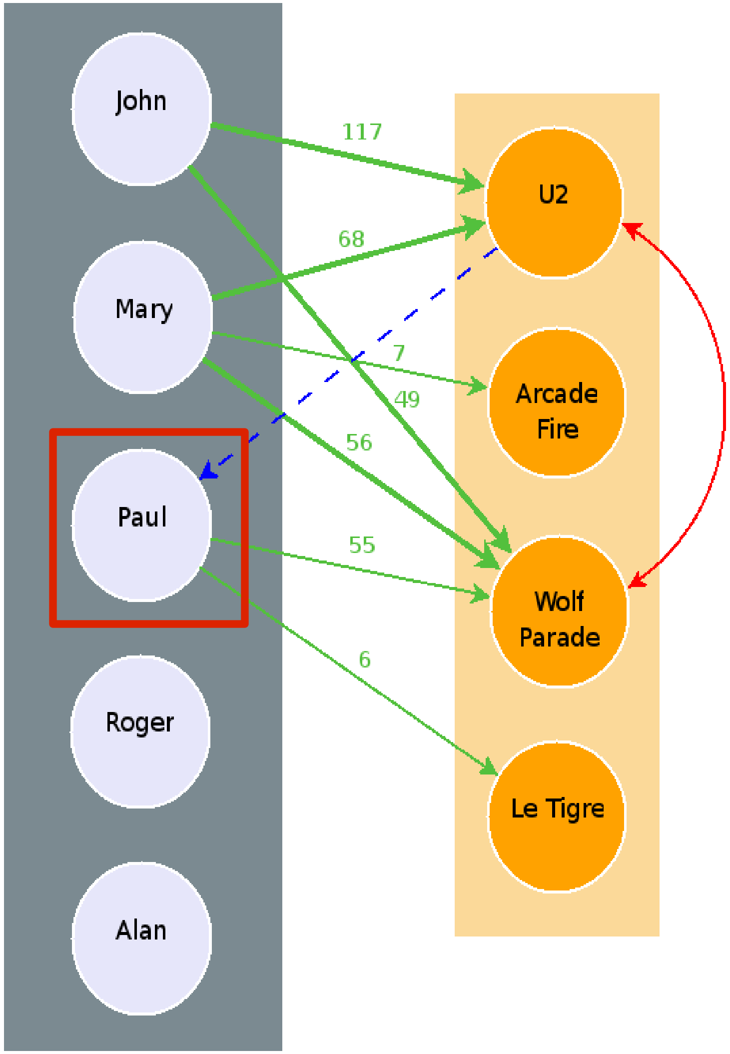

- Recommend **U2** to **Paul**

music recommendation:: algorithms :: colfilter

- Collaborative Filtering: Item-based
  - ❖ user-item matrix

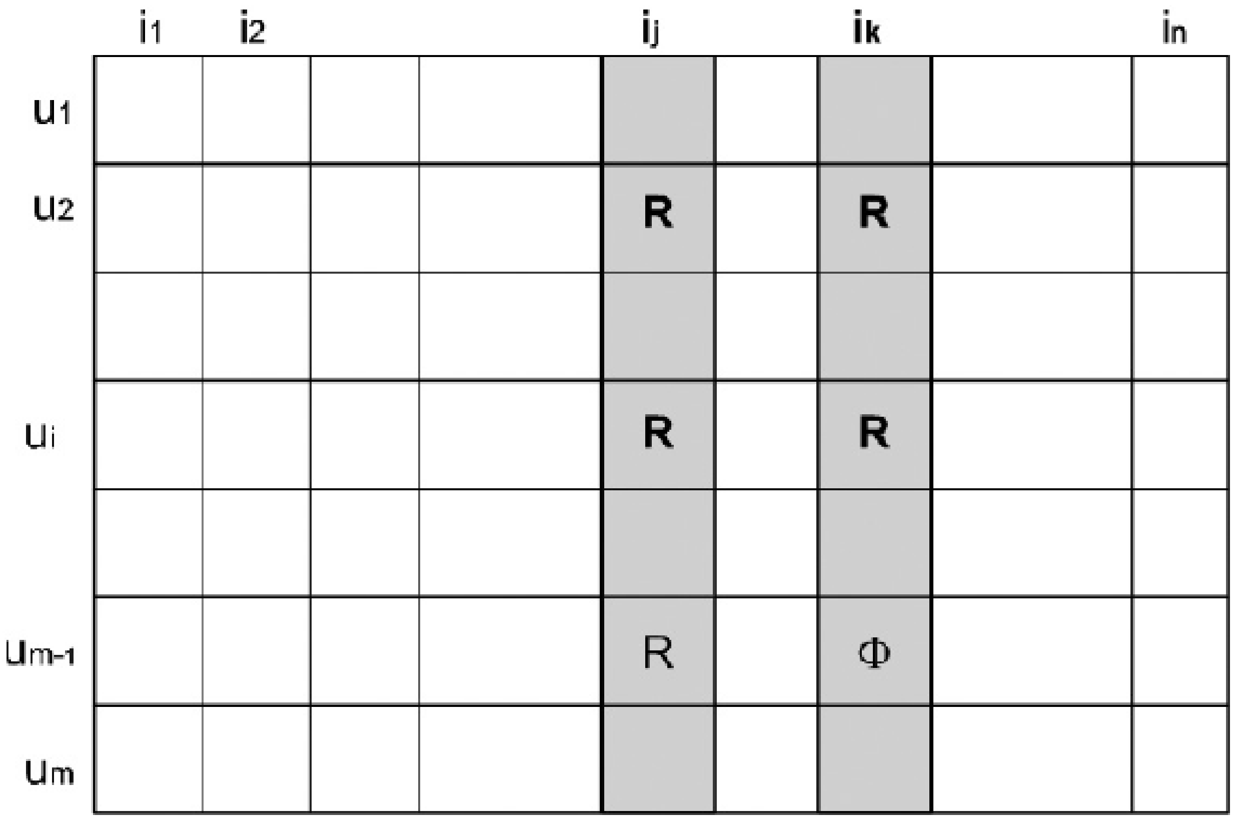

music recommendation:: algorithms :: colfilter

- Collaborative Filtering: Item-based
  - item similarity measures
    - cosine

$$sim(i,j) = cos(\vec{i},\vec{j}) = \frac{\vec{i}\cdot\vec{j}}{\|i\| * \|j\|} = \frac{\sum_{u\in U} R_{u,i}R_{u,j}}{\sqrt{\sum_{u\in U} R_{u,i}^2}\sqrt{\sum_{u\in U} R_{u,j}^2}}$$

    - adjusted cosine

$$sim(i,j) = \frac{\sum_{u\in U}(R_{u,i}-\bar{R}_u)(R_{u,j}-\bar{R}_u)}{\sqrt{\sum_{u\in U}(R_{u,i}-\bar{R}_u)^2}\sqrt{\sum_{u\in U}(R_{u,j}-\bar{R}_u)^2}}$$

    - pearson correlation

$$sim(i,j) = \frac{Cov(i,j)}{\sigma_i\sigma_j} = \frac{\sum_{u\in U}(R_{u,i}-\bar{R}_i)(R_{u,j}-\bar{R}_j)}{\sqrt{\sum_{u\in U}(R_{u,i}-\bar{R}_i)^2}\sqrt{\sum_{u\in U}(R_{u,j}-\bar{R}_j)^2}}$$

music recommendation:: algorithms :: colfilter

- Collaborative Filtering: Item-based
  - item similarity measures
    - cosine

$$sim(i,j) = cos(\vec{i},\vec{j}) = \frac{\vec{i}\cdot\vec{j}}{\|i\| * \|j\|} = \frac{\sum_{u\in U} R_{u,i}R_{u,j}}{\sqrt{\sum_{u\in U} R_{u,i}^2}\sqrt{\sum_{u\in U} R_{u,j}^2}}$$

    - adjusted cosine

$$sim(i,j) = \frac{\sum_{u\in U}(R_{u,i} - \bar{R}_u)(R_{u,j} - \bar{R}_u)}{\sqrt{\sum_{u\in U}(R_{u,i} - \bar{R}_u)^2}\sqrt{\sum_{u\in U}(R_{u,j} - \bar{R}_u)^2}}$$

    - pearson correlation

$$sim(i,j) = \frac{Cov(i,j)}{\sigma_i\sigma_j} = \frac{\sum_{u\in U}(R_{u,i} - \bar{R}_i)(R_{u,j} - \bar{R}_j)}{\sqrt{\sum_{u\in U}(R_{u,i} - \bar{R}_i)^2}\sqrt{\sum_{u\in U}(R_{u,j} - \bar{R}_j)^2}}$$

music recommendation:: algorithms :: colfilter

- Collaborative Filtering: Item-based
  - Prediction
    - user *u*, item *i*
    - Weighted sum

$$P_{u,i} = \frac{\sum_j s(i,j) R_{u,j}}{\sum_j s(i,j)}$$

music recommendation:: algorithms :: colfilter

- Collaborative Filtering
  - other approaches
    - Dimensionality reduction of the user-item matrix
      - SVD (LSA) [Hofmann, 2004]
      - Multidimensional Scaling [Platt, 2004]
    - graph based, with link prediction
      - consumer – product bipartite graphs [Huang, 2005], [Huang, 2007], [Mirza, 2003]

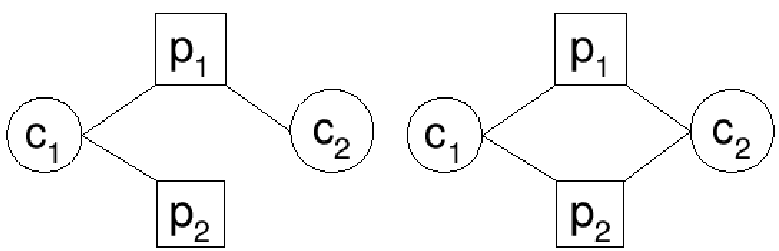

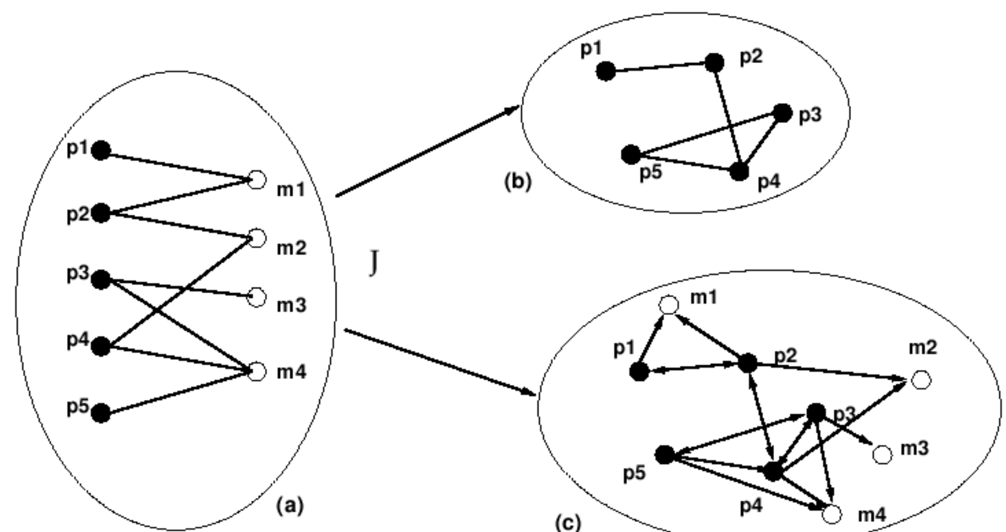

music recommendation:: algorithms :: content

- Content-based filtering
  - based on item similarity
  - usually at song level
  - similarity can be
    - content-based
    - editorial
    - tag-based
    - etc.

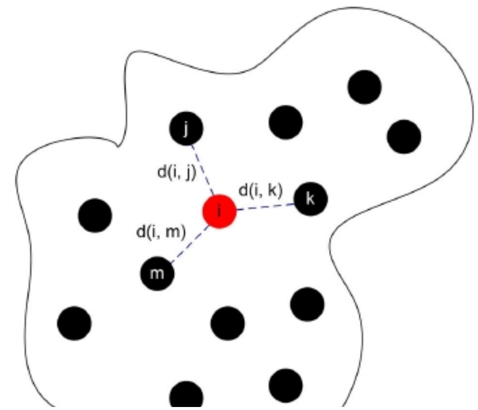

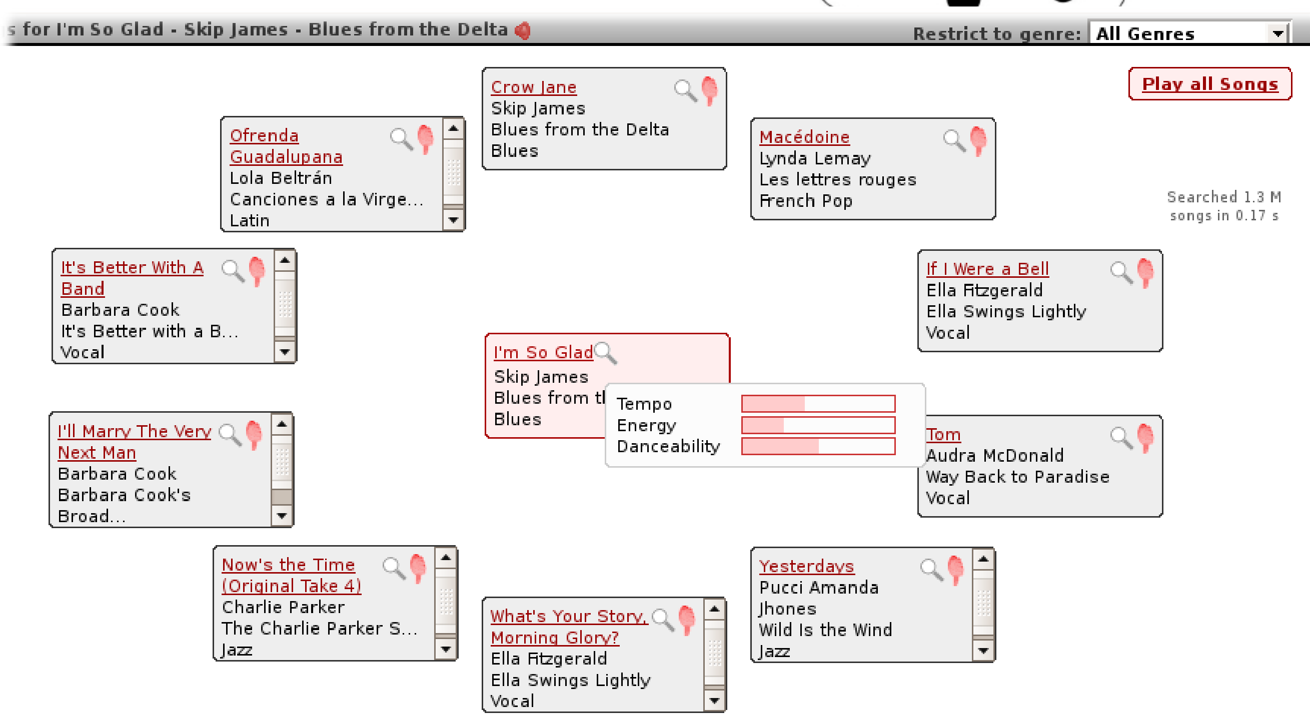

music recommendation:: algorithms :: hybrid

- Hybrid methods
  - combine previous approaches (mainly CF and CB)
    - combining both outputs (e.g linear combination)
    - cascade: CF -> CB, or CB -> CF
    - select the best method at anytime (CF | CB)
    - etc.

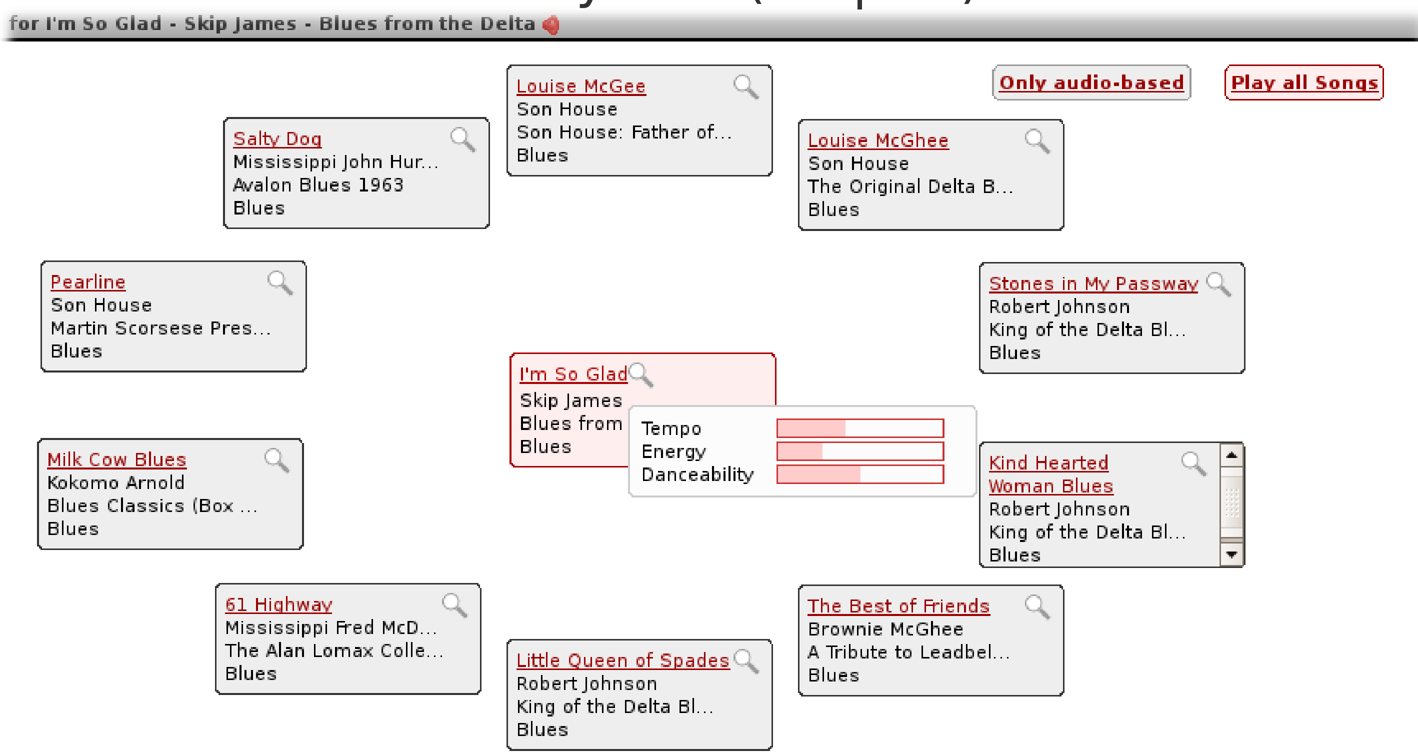

outline

- Introduction
- Formalization of the recommendation problem
- Recommendation algorithms
- **Problems with recommenders**
- Recommender examples
- Evaluation of recommenders
- Conclusions / Future

problems:: Social recommenders :: Cold Start

- Sparse Data can lead to poor recommendations
  - Postal Service - “Such Great Heights”
    - 2.4 million scrobbles
    - 1000s of tags
  - Mike Shupps's - “All Over Town”
    - 3 scrobbles
    - 0 Tags
- A problem for:
  - New artists/Tracks
  - New users
  - New recommenders

problems:: Social recommenders:: Cold Start

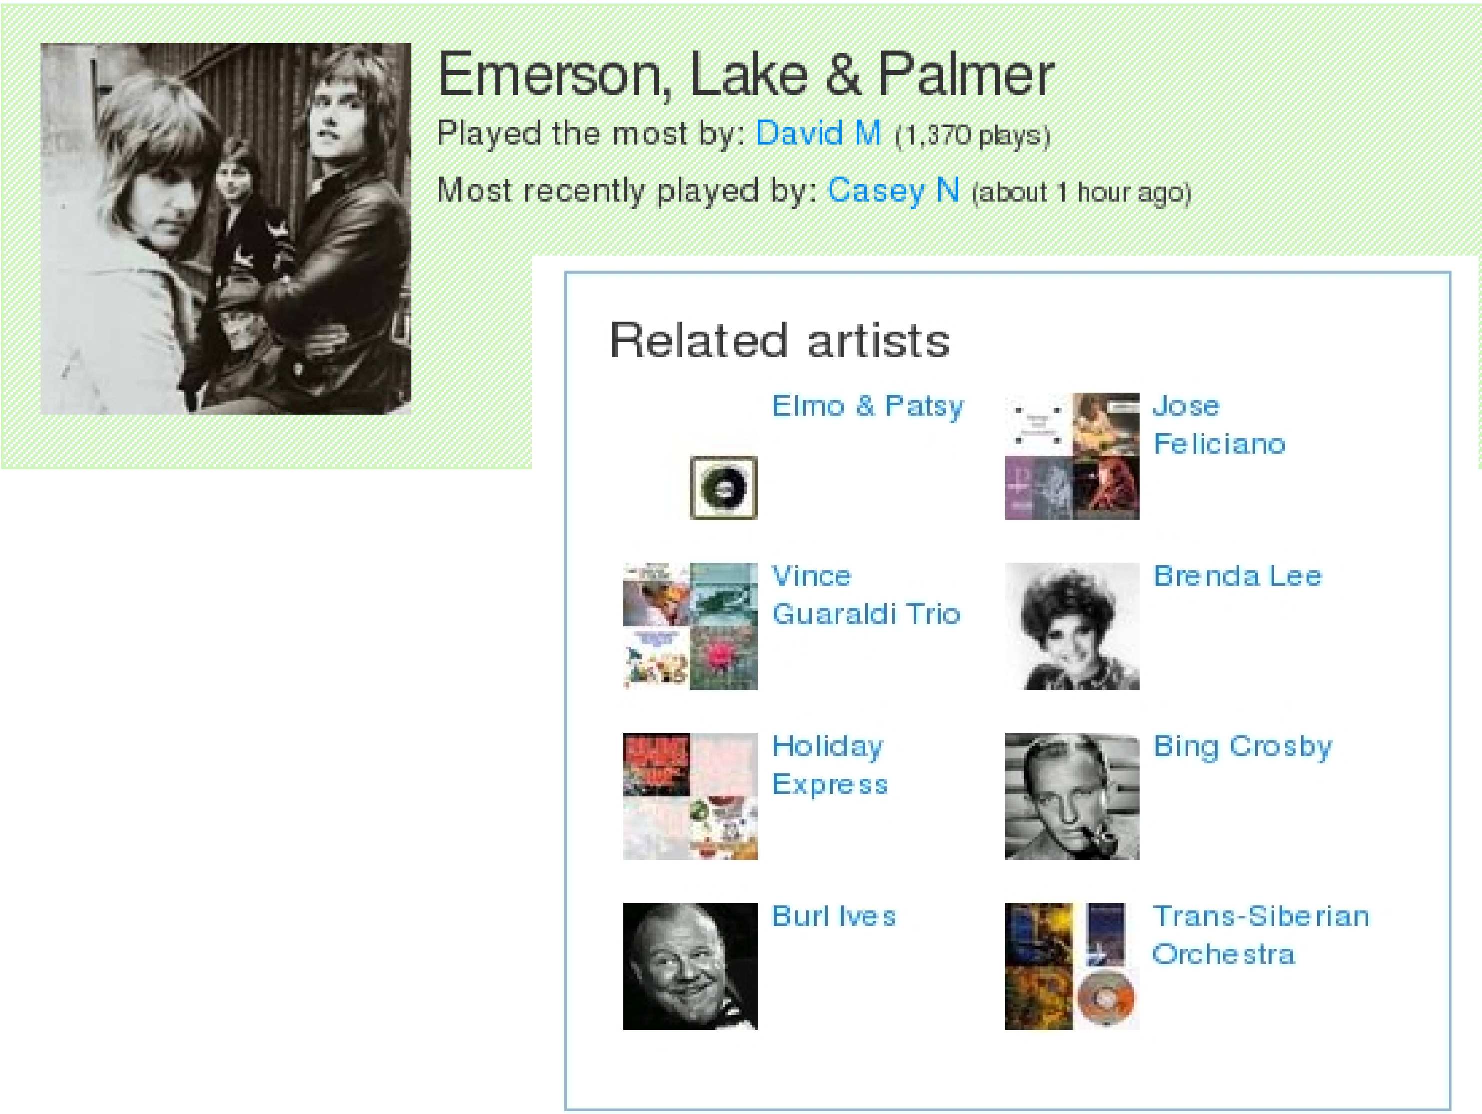

iLike Beta

problems:: Social recommenders:: Feedback Loops

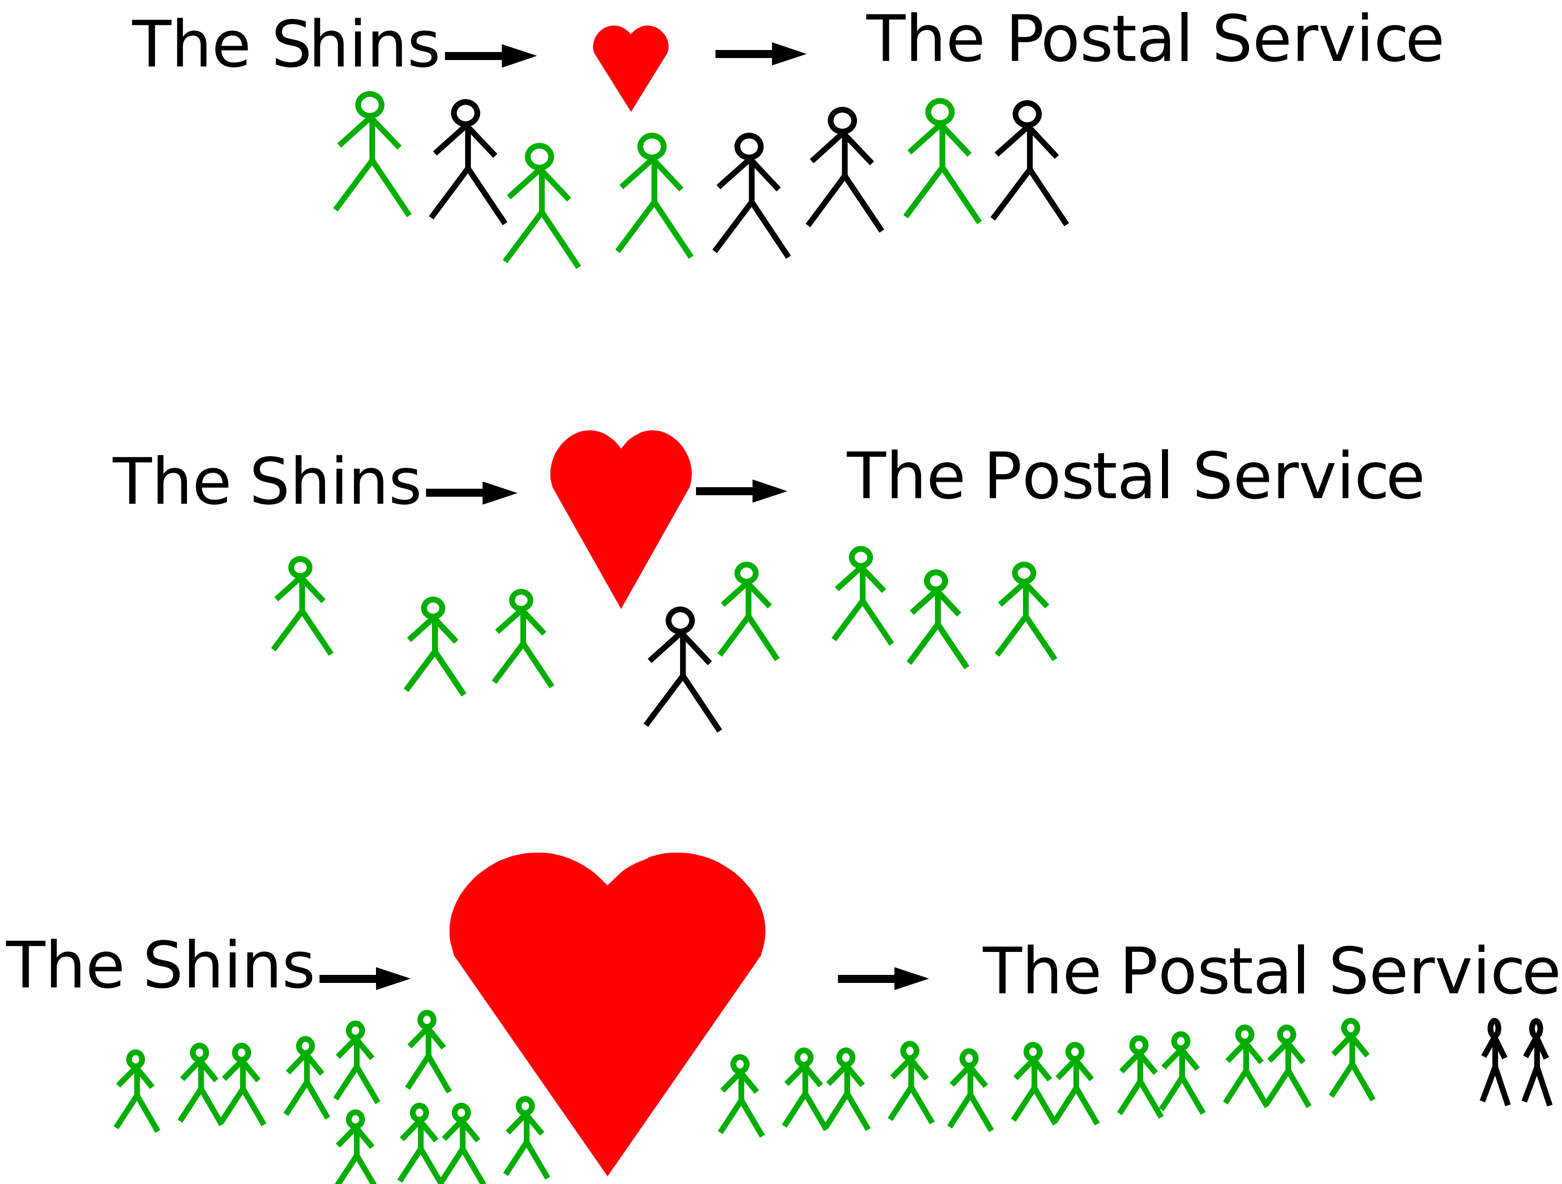

problems:: Social recommenders:: Popularity Bias

problems:: Social recommenders:: Scale

- Netflix:
  - 5 million customers
  - 50,000 items
  - 1.4B tatings
  - 2M ratings per day
  - 1B predictions per day
  - 2 days to retrain
- Amazon:
  - 30 million customers
  - 1 million items
- Yahoo! Music:
  - 25 million Users
  - 600 million minutes/month
  - 7 billion Ratings
  - 30 million user-customized radio stations
- Last.fm
  - 500 million 'scrobbles' per month
  - 20 million unique users
  - 100 million tracks
  - 2 million tags applied per month

problems:: Social recommenders

- Lack of transparency

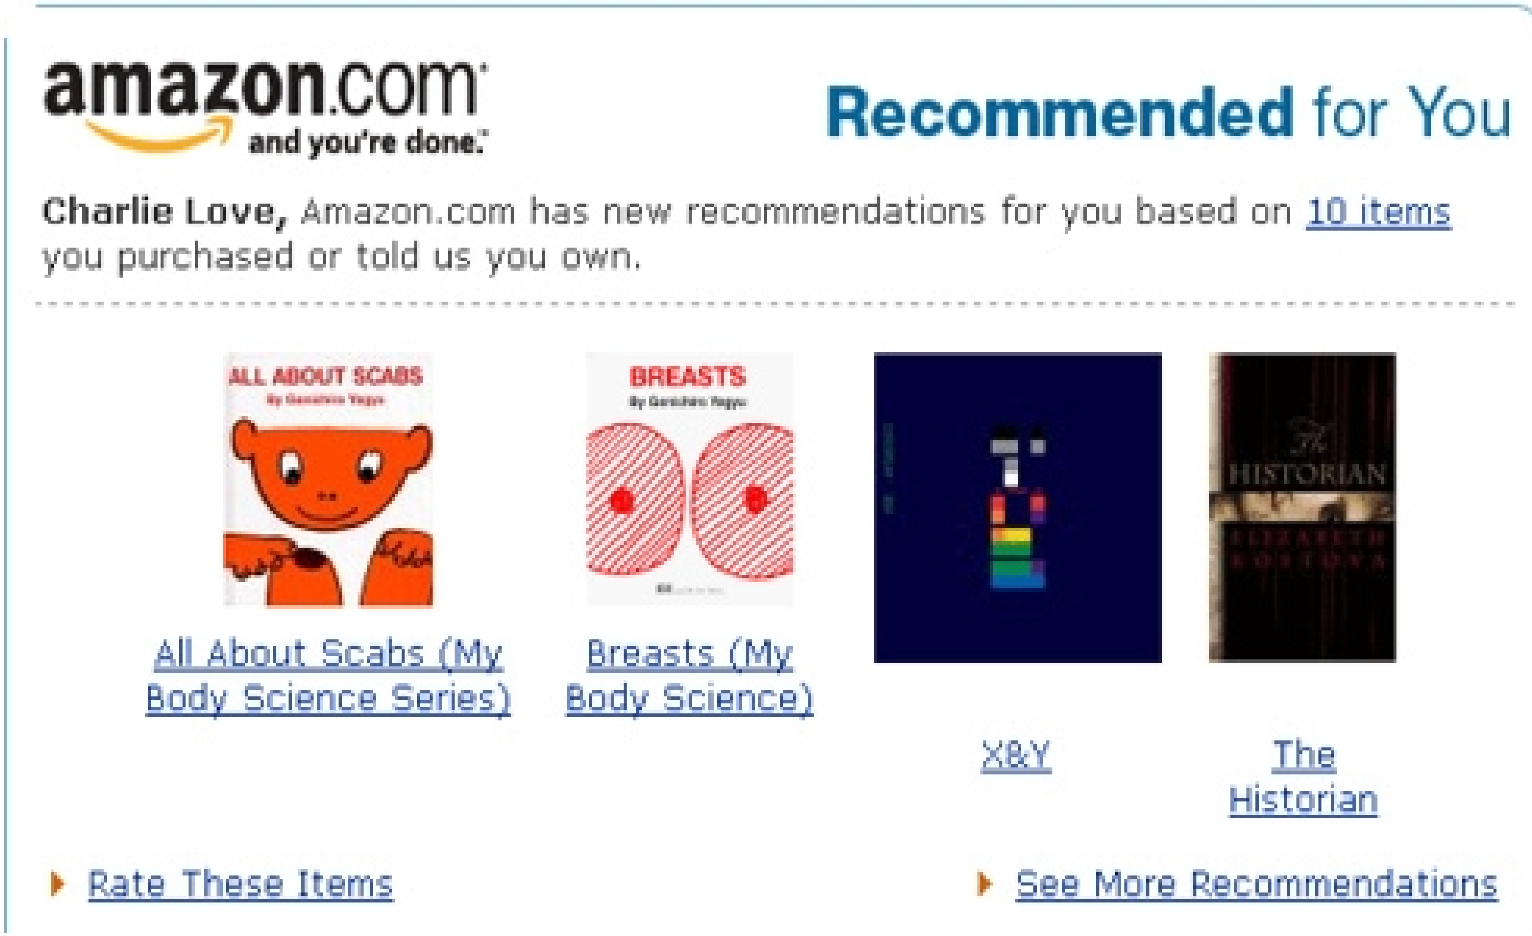

problems:: Social recommenders:: Early Rater Bias

- Early rater bias
  - Rich get Richer (Cumulative Advantage)
  - **Social Influence** as important as **quality**
  - Success of a song depends on the decisions of a few early-arriving individuals
  - The particular songs that became hits were different in different worlds

Is Justin Timberlake a Product of Cumulative Advantage?

Duncan Watts

problems:: Social recommenders

- Gray sheep
  - ❖ Common tastes mixed with uncommon tastes

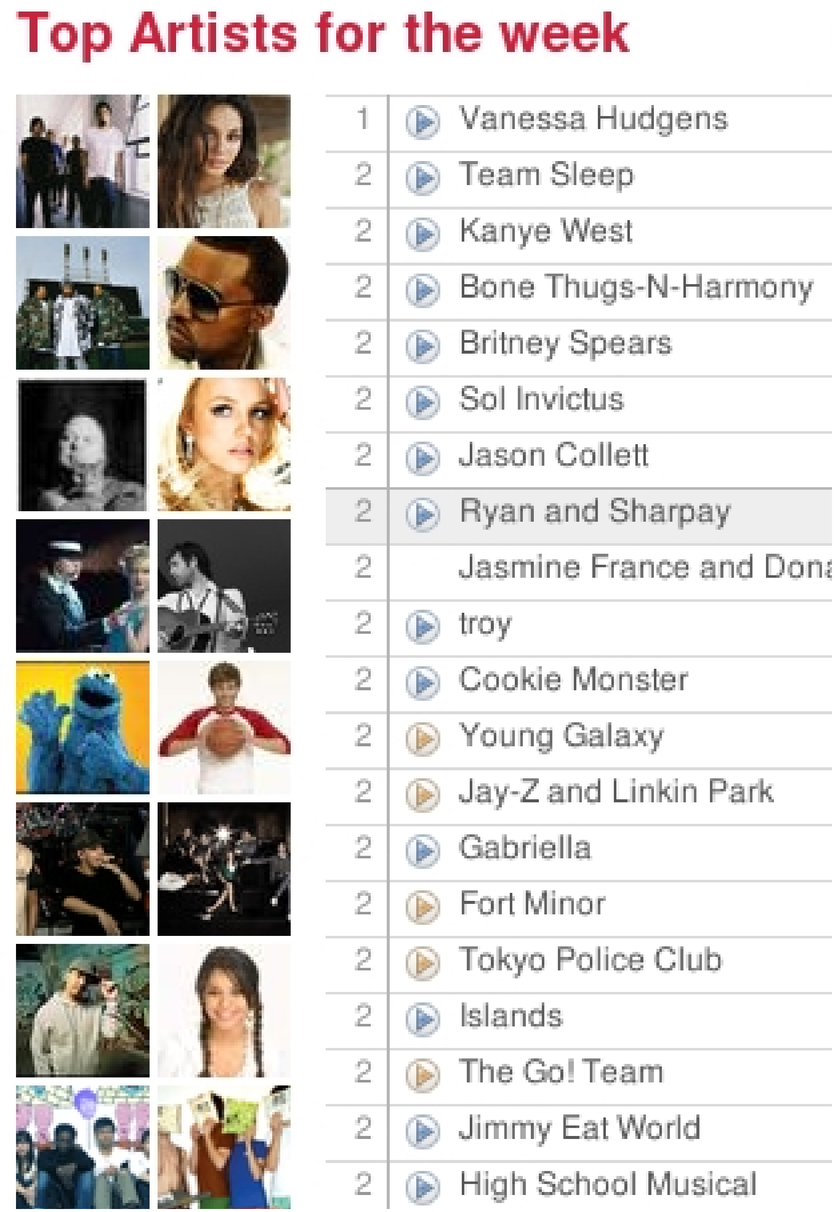

problems:: Social recommenders

- Hacking the recommender
  - ❖ Profile Injection Attack

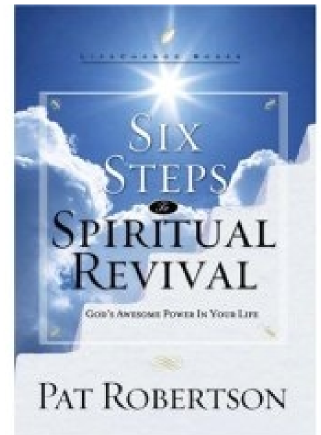

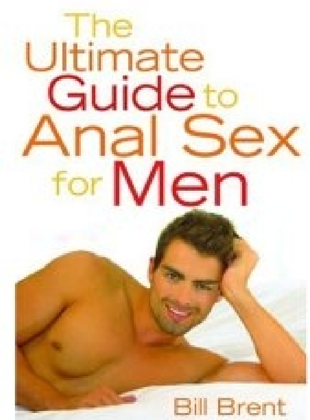

Courtesy of Bamshad Mobasher

problems:: Social recommenders

- Inertia / Aging

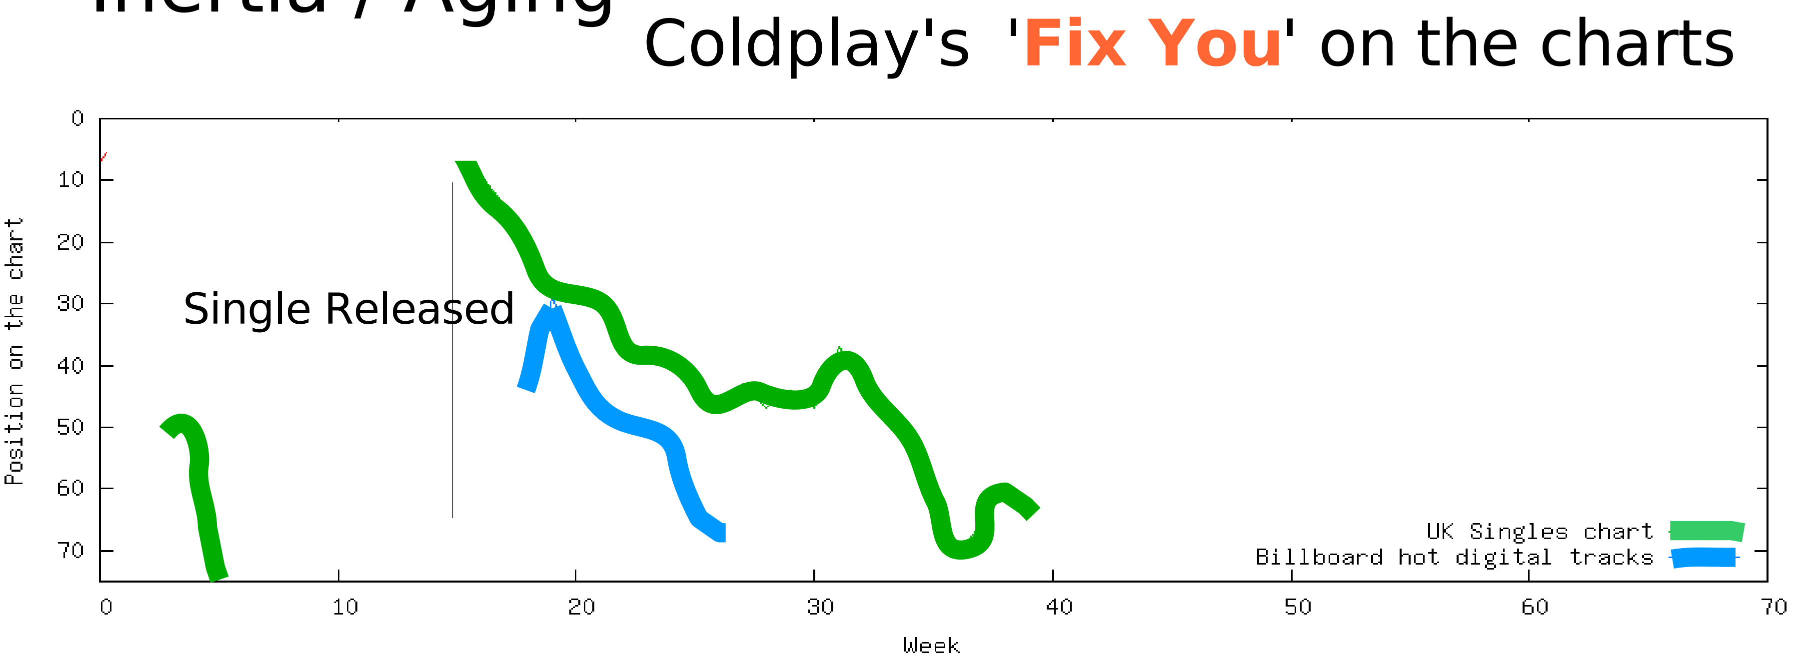

❖ Traditional sales charts have built-in decay

problems:: Social recommenders

- Inertia / Aging

Coldplay's '**Fix You**' on the charts

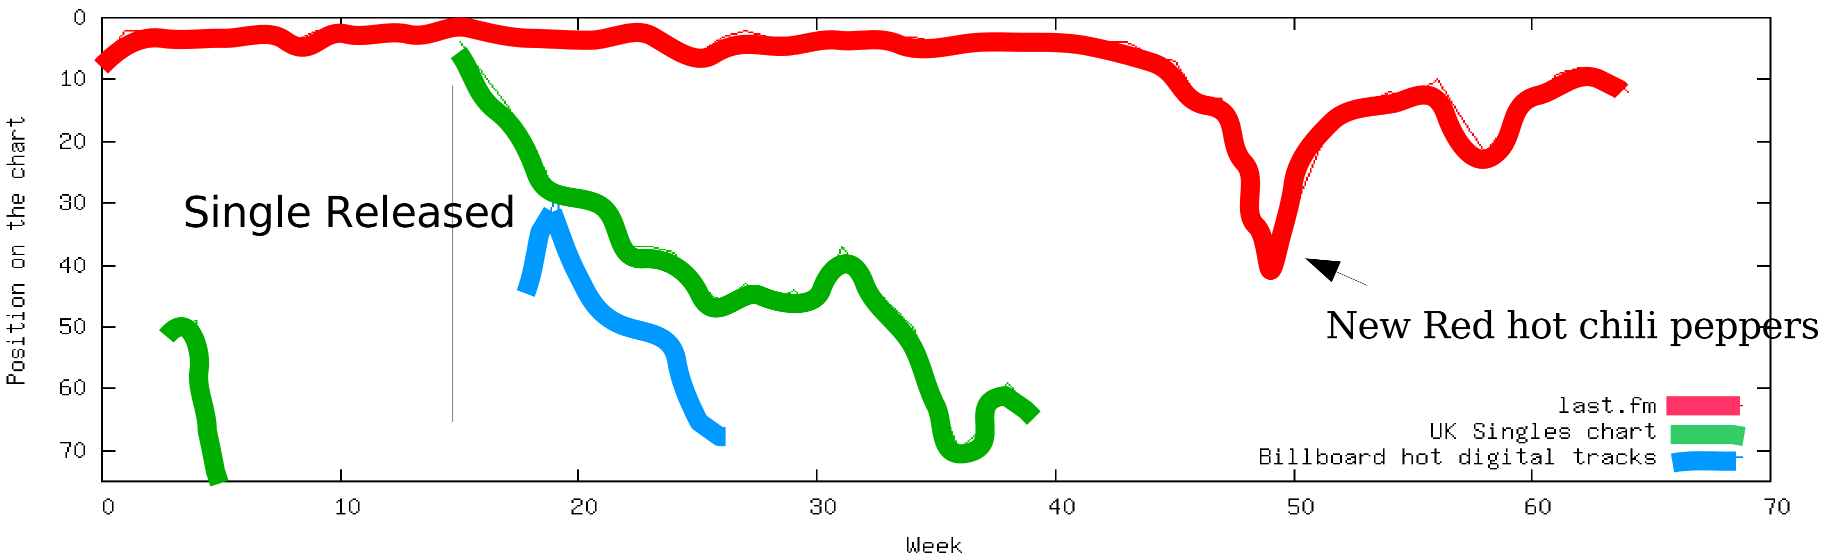

❖ Traditional charts have built-in decay

❖ New 'play' charts resist decay

problems:: Social recommenders

- Inertia / Aging

Coldplay's '**Fix You**' on the charts

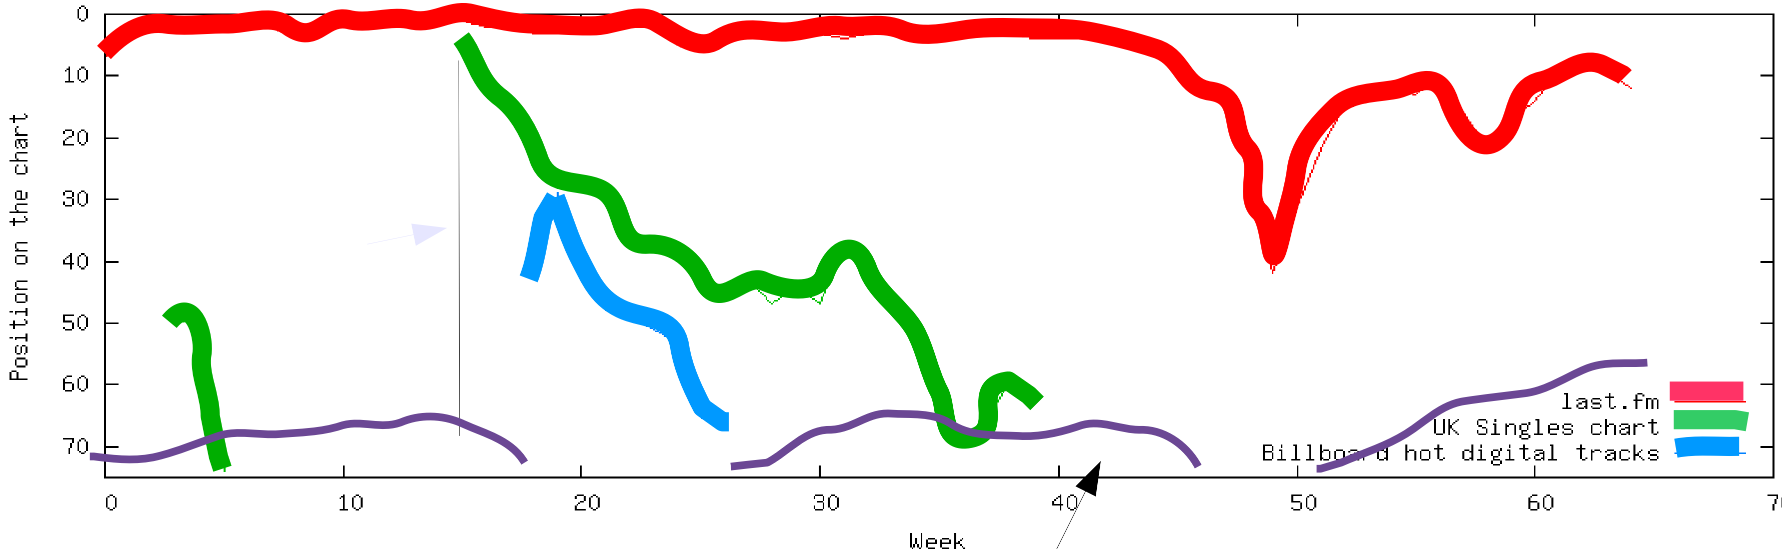

- #1 with a lead weight
  - ❖ "**Such Great Heights**" – 150 weeks in the top 10
  - ❖ "**Stairway to Heaven**" – #58 after 35 years

problems:: Social recommenders

- Inertia / Aging

Coldplay's '**Fix You**' on the charts

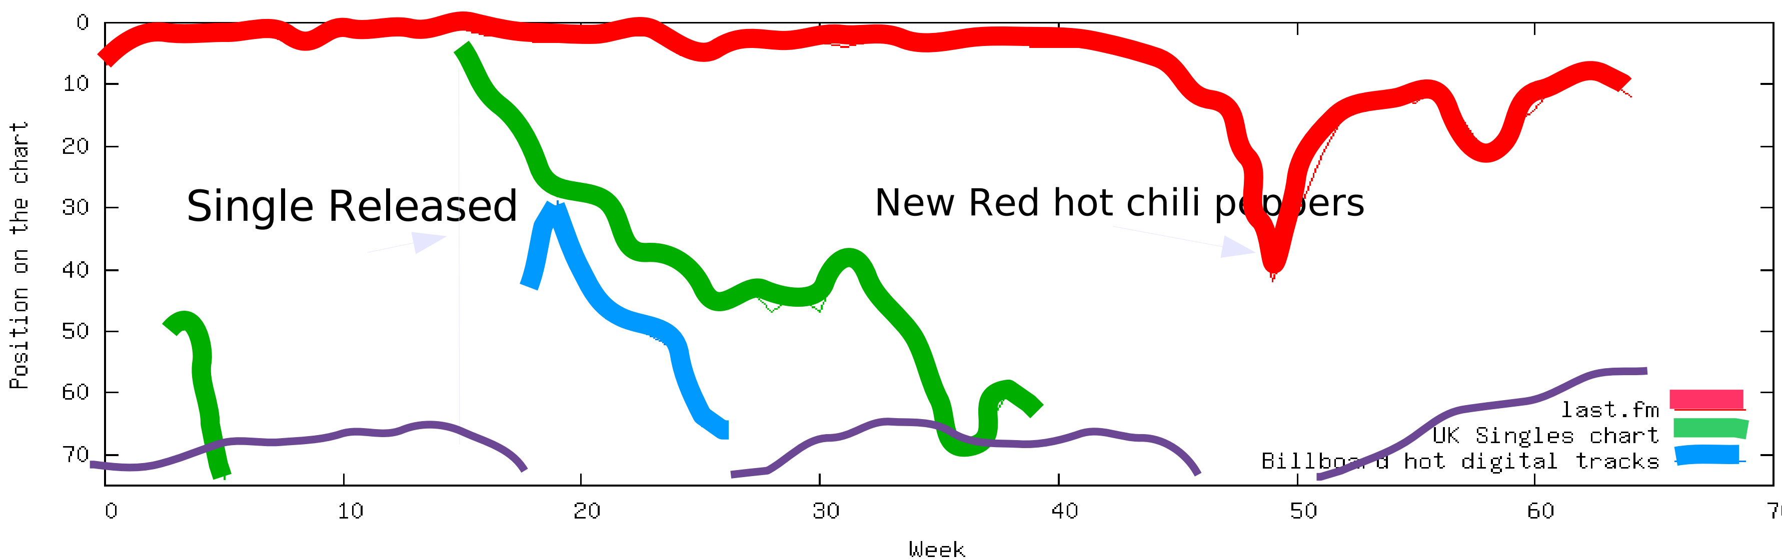

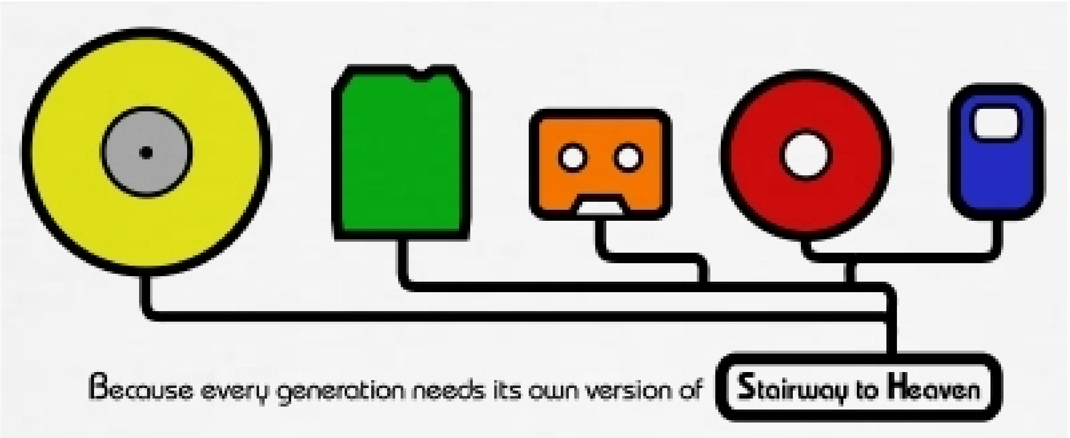

problems:: Social recommenders

- Novelty / Serendipity

Songs like "Hey Jude" by Elvis Presley

Discover | Share | Export to iTunes | Save | Advanced Controls | New Search

Preview	Play	Song	Artist
		Hey Jude	Elvis Presley
		Hey Jude	Arthur Fiedler
		Hey Jude	Luca Colombo
		Hey Jude	Espitia, J. Lennon/P. McCartney
		Hey Jude	Wilson Pickett
		Hey Jude	Chokocheeky
		Hey Jude	Dale Miller
		Hey Jude	Tiny Tim
		Mother Nature's Son	John Denver
		Strawberry Fields Forever	The Real Group

problems:: Social recommenders

- Trust
  - How can you trust a recommendation?
  - Payola
  - Amazon “Reviewer Scandal”
  - Pay-per-post
  - Pay-per-digg
- Coming soon:
  - Pay per Scrobble?

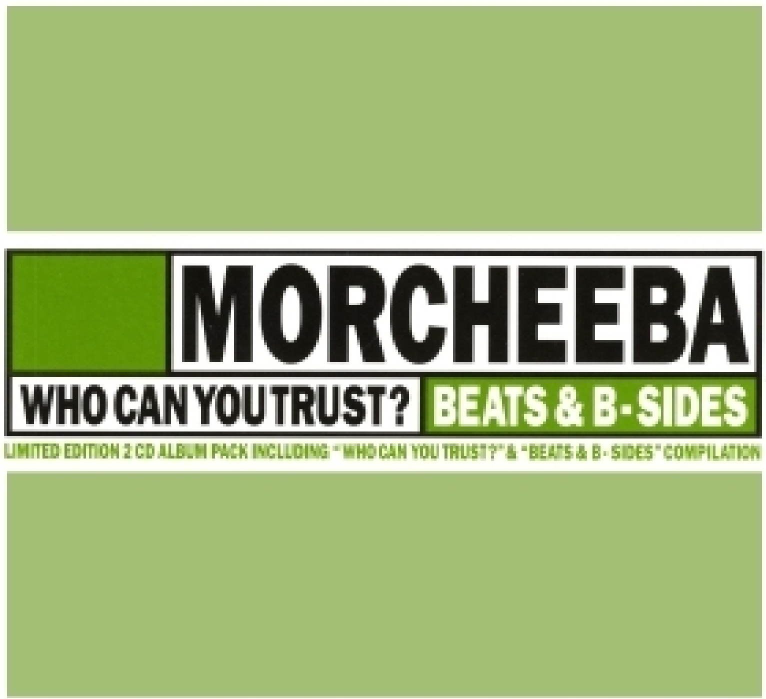

problems:: Social recommenders:: trust

- PayPerPost

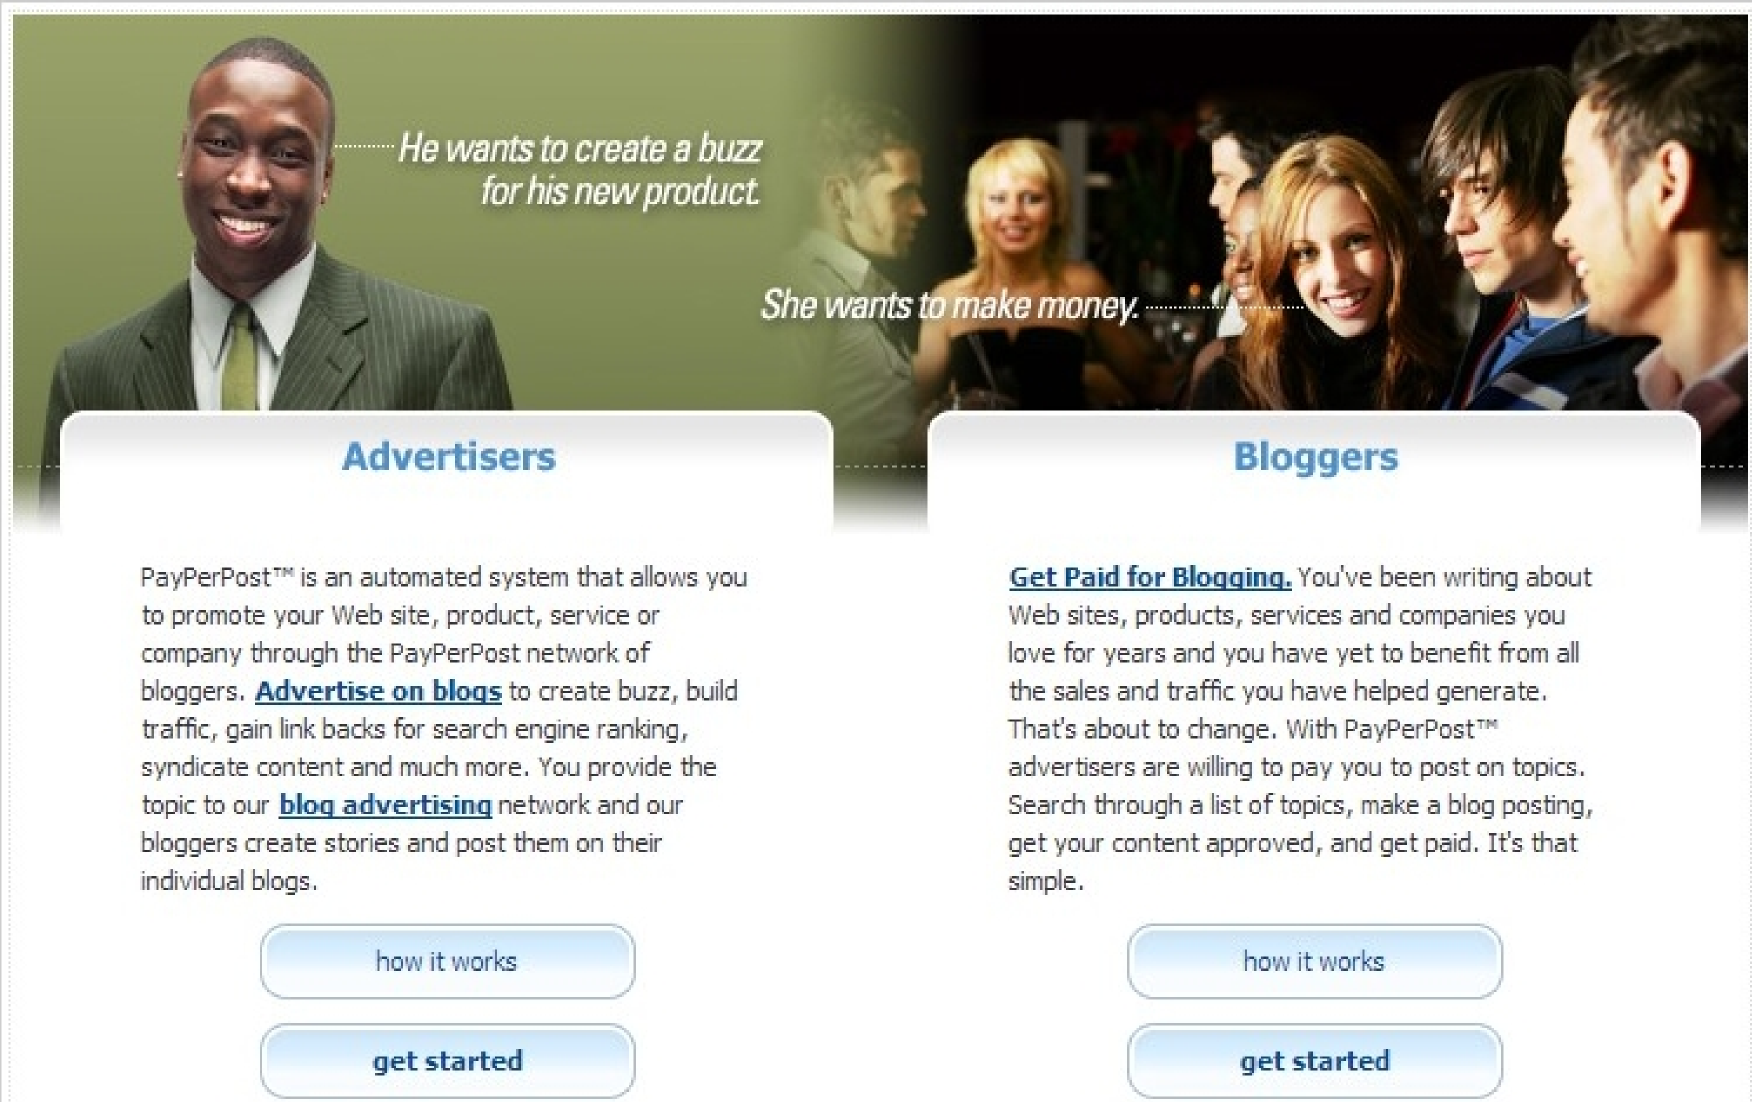

problems:: Social recommenders

- Gallery of bad recommendations

Are these really
Better Together?

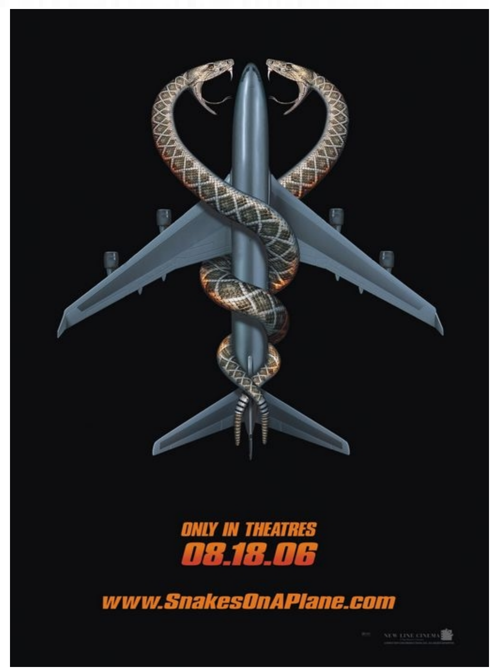

problems:: Social recommenders

- Strange connections

amazon.com

Rate this item

Durex Avanti Polyurethane Condoms, 6 Condoms
Durex
Price: $8.99
Add to Cart
Add to Wish List

...because you were interested in:

Netgear WG511 Wireless 802.11g PC Card
by Netgear
You purchased or rated this item

problems:: Social recommenders

Strange Connections

Panasonic ER411NC Nose and Ear Hair Groomer
Average Customer Review:
Release Date: May 16, 2002

Our Price: $14.88 Used & new from $14.88

Own It Not interested x| Saved
mmended because you rated The Chronicles of Narnia Boxed Set and more (edit)

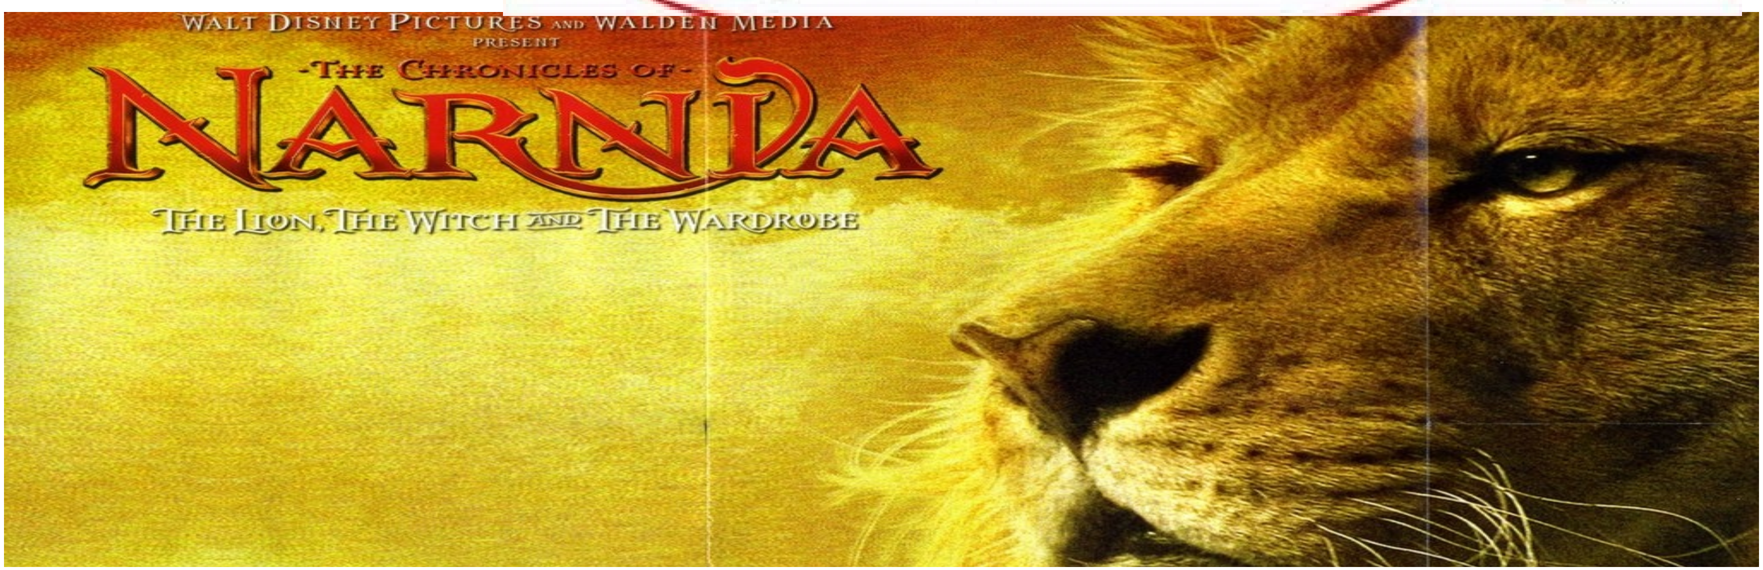

problems:: Social recommenders

If you like Gregorian Chants you might like Green Day

If you like Britney Spears you might like...

problems:: Content based recommenders

- Social effects are extremely important
- Early days:
  - ❖ Poor quality recommendations
  - ❖ Can't tell 'good' music from 'bad'
  - ❖ Can make mistakes no human would make:
    - harpsichord <-> distorted electric guitar
  - ❖ "Pandora isn't broken. The listener is."
  - ❖ DRM – listen only music

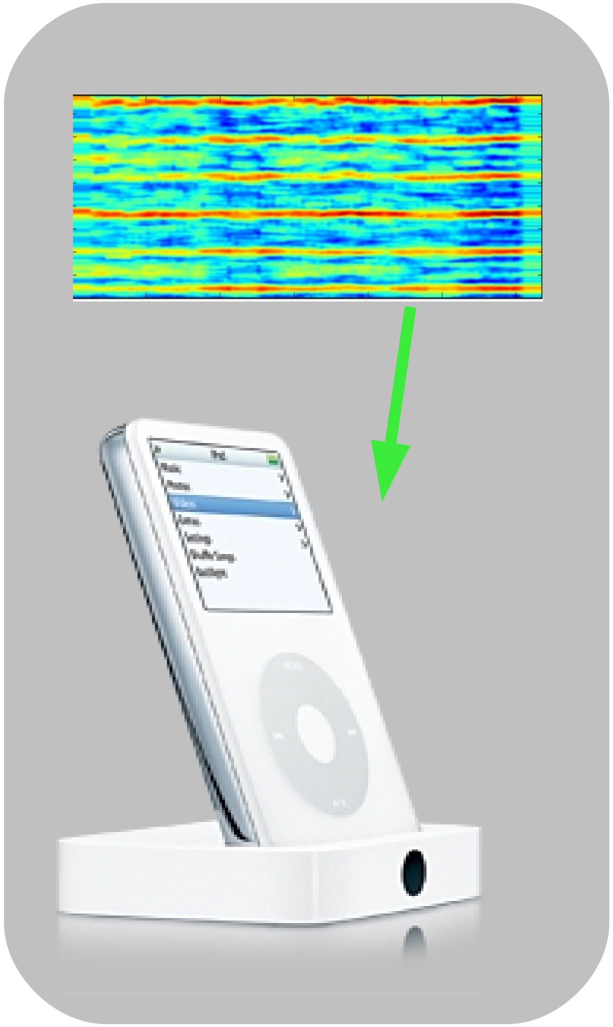

problems:: Content based recommenders

- Analysis is hard

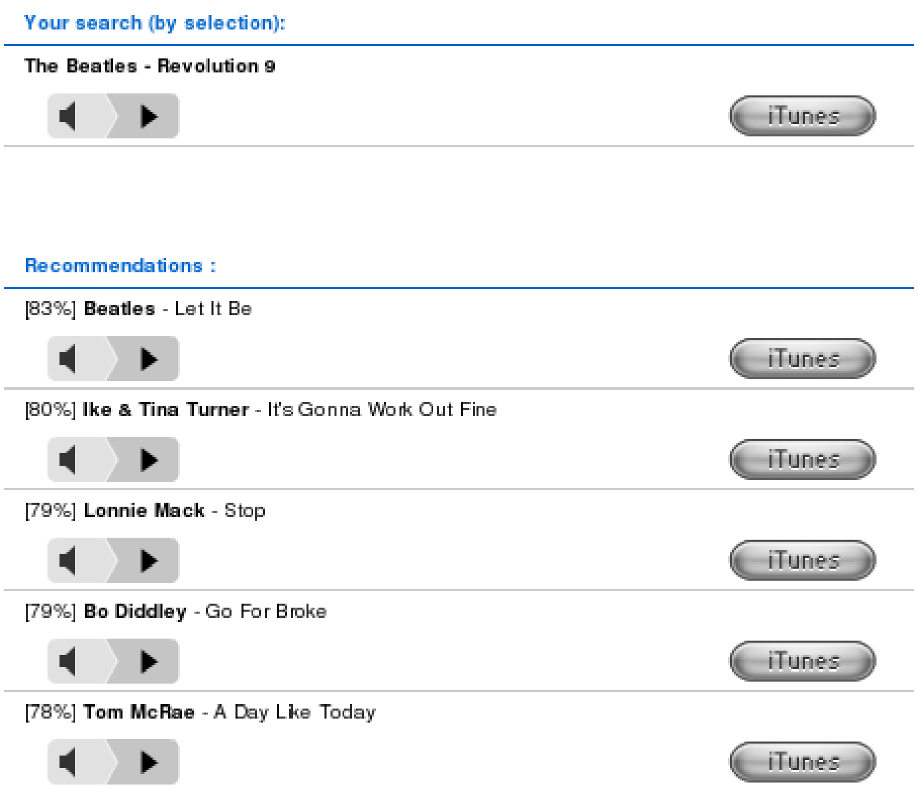

XXX"s technology analyses the music content to extract information about rhythm, tempo, timbre, instruments, vocals, and musical surface of a song.

In fact, they rely mostly (or entirely on metadata for recommendations).

outline

- Introduction
- Formalization of the recommendation problem
- Recommendation algorithms
- Problems with recommenders
- **Recommender examples**
- Evaluation of recommenders
- Conclusions / Future

examples:: Hard core music discovery

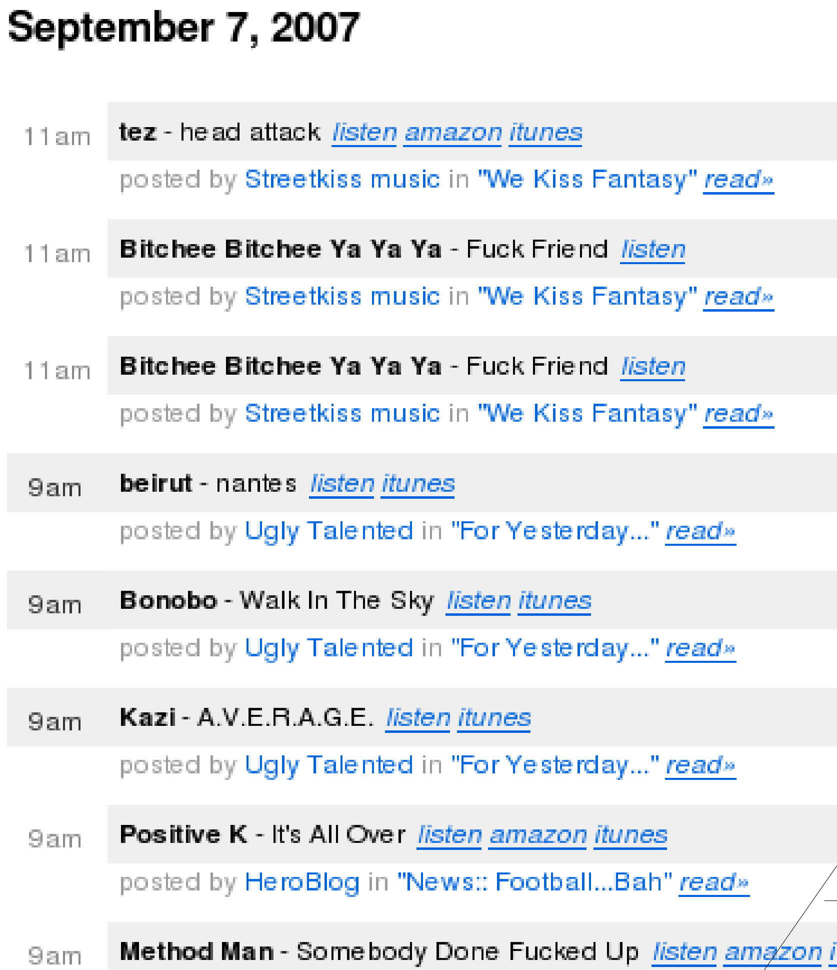

- No filter - High Risk listening
- Long tail - without the “help me find it”

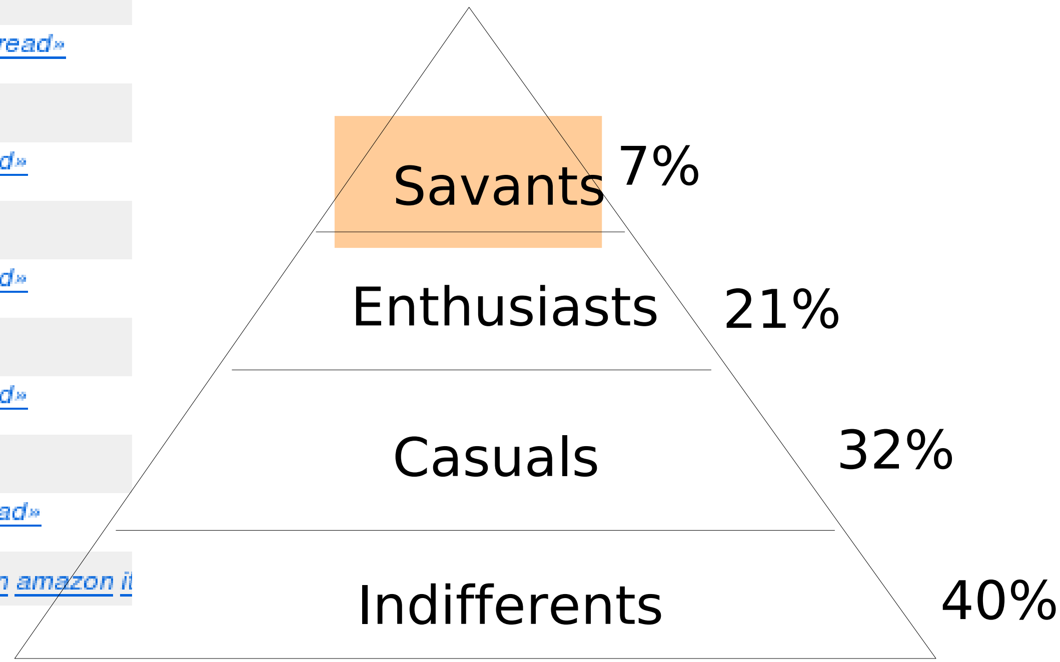

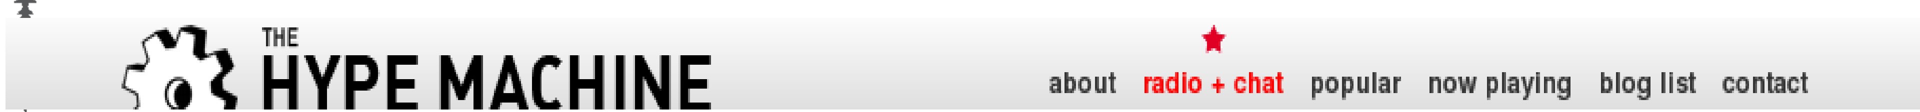

examples:: Hard core music discovery

- Primary discovery tool is “What's Hot” charts

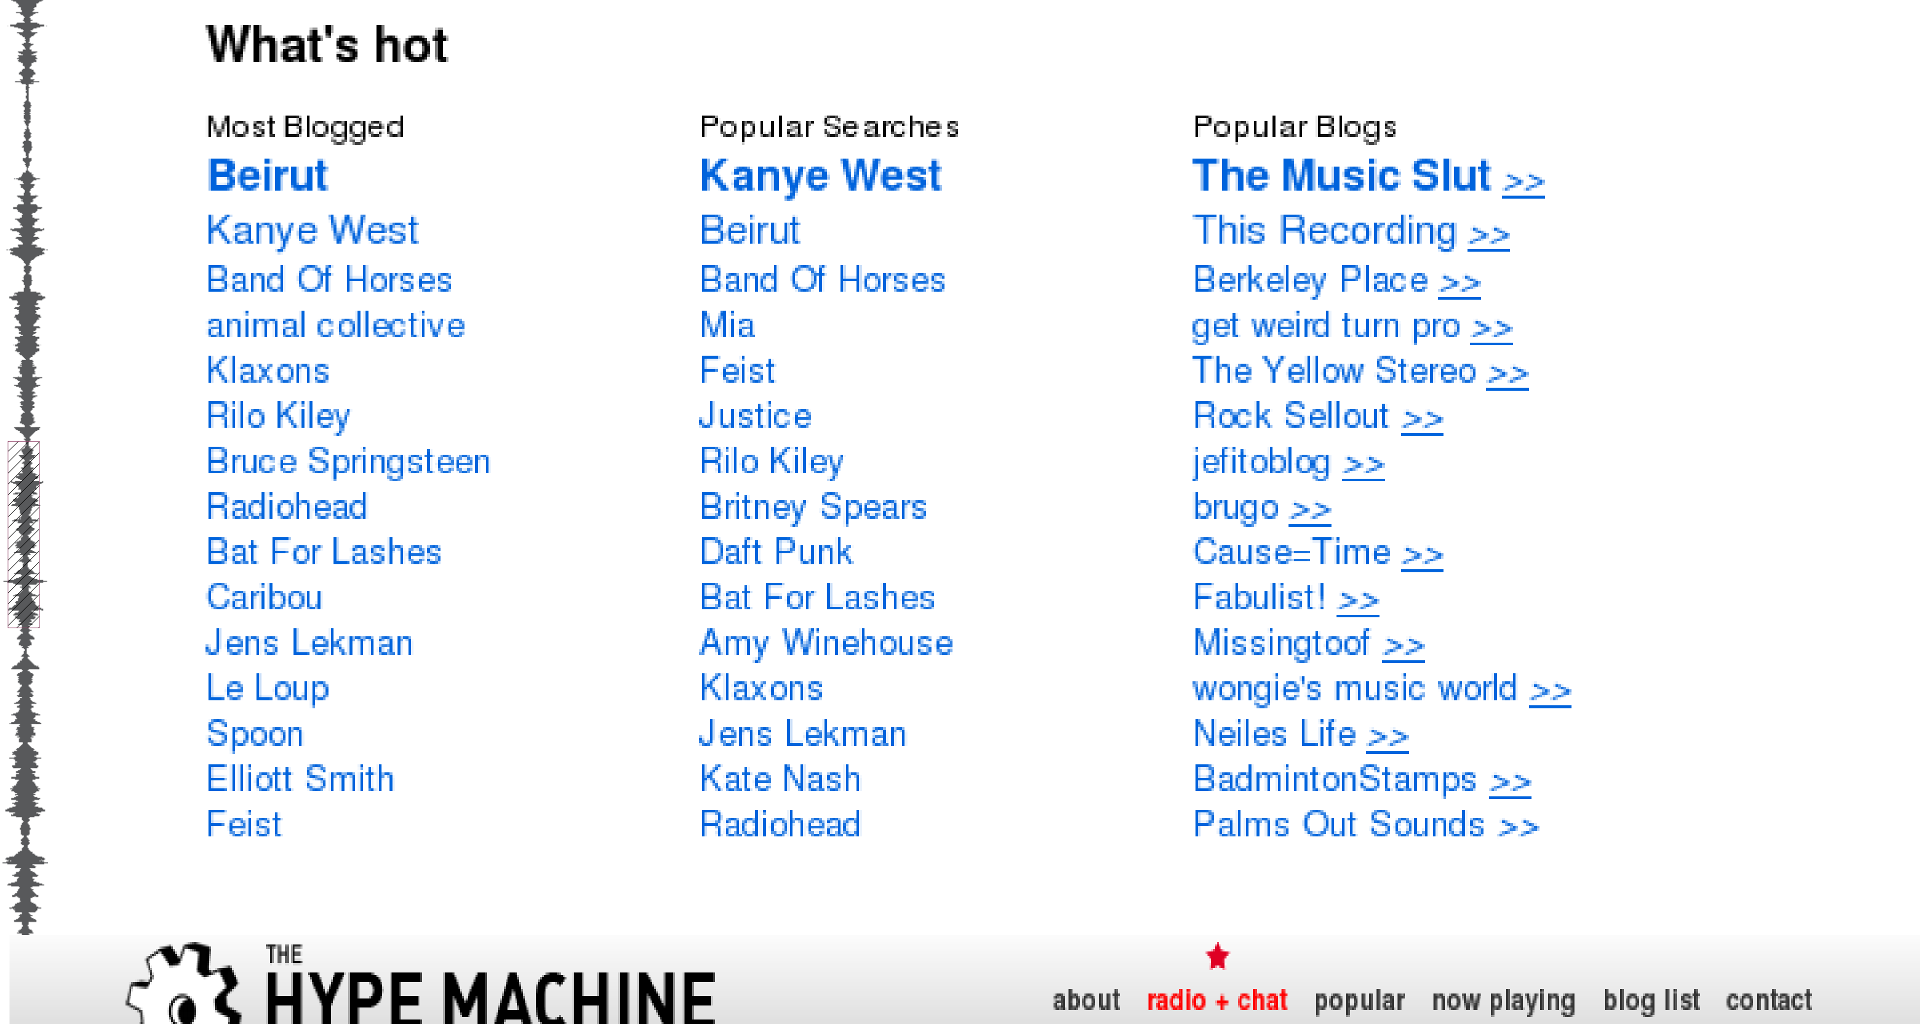

examples:: Social music

- Kitchen Sink Interface
- Focus is on music discovery
- Many tools for music discovery
  - recommendations, tags
  - friends, videos, charts ...

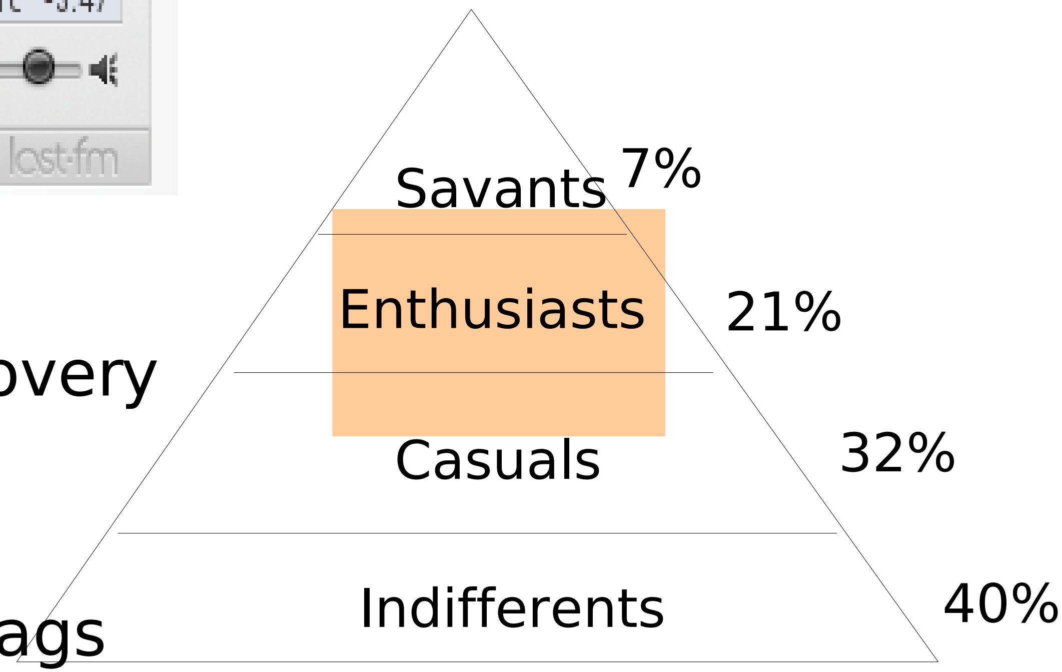

examples:: Social music

- Last.fm Features
  - ❖ Scrobbler
  - ❖ Artist, Album and Track tagging
  - ❖ Internet Radio
  - ❖ Recommendations
    - Similar artists,
    - Similar Listeners
  - ❖ Charts
    - Tracks, Artists, Albums, Videos, Movers, Tags
  - ❖ Videos, Events
  - ❖ Social Features
    - Friends, messages, shoutouts

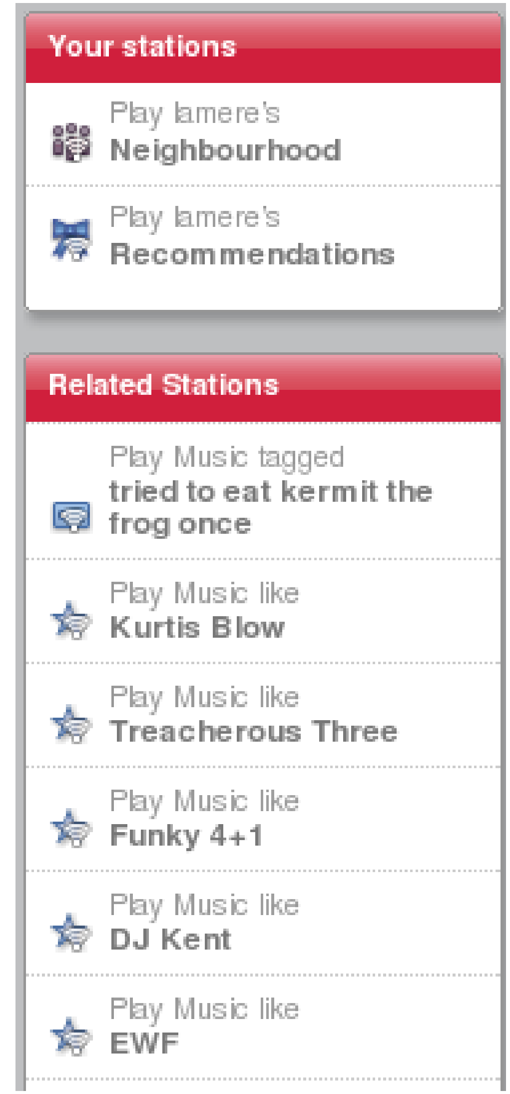

examples:: Social music

Last.fm recommendations

- Recommendations:
  - Primarily Collaborative Filtering
  - Item-Item (artist recommendations
  - User-User (Neighbors)
  - Could use:
    tags, audio, metadata
- Evaluating (rel. feedback)
  - Tracking Love/Ban behavior

examples:: Social music

Internet radio

- Neighbor radio
- Recommender Radio
- “My Radio”
- Tag Radio
- User Radio
- While listening:
  - Artist Bio
  - Related Stations
  - Top Listeners

examples:: Social music Last.fm Charts

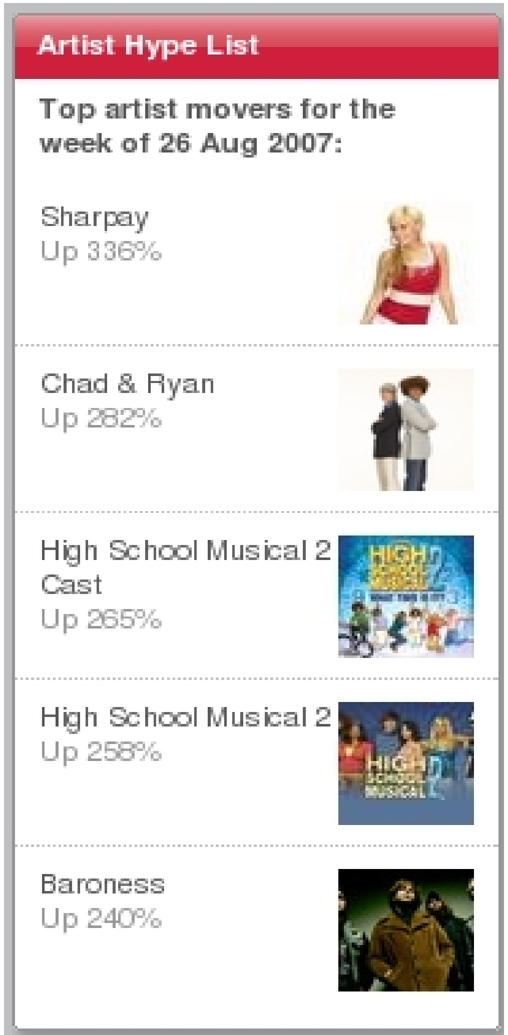

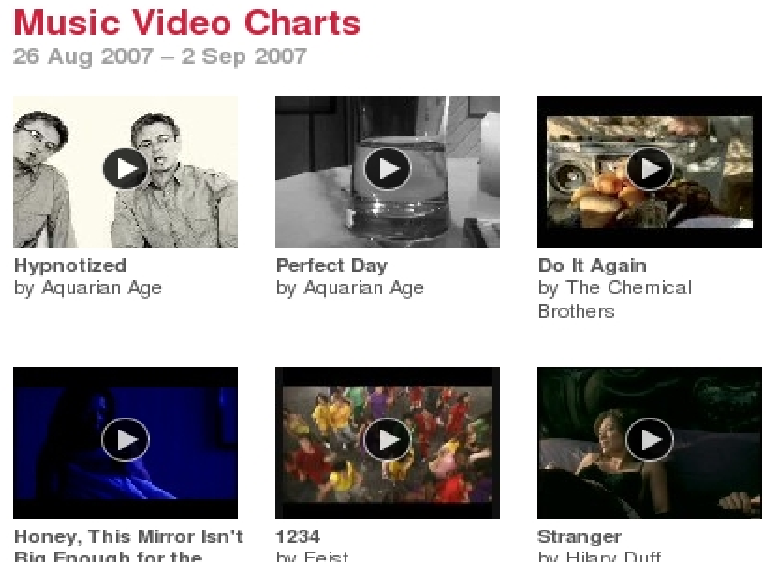

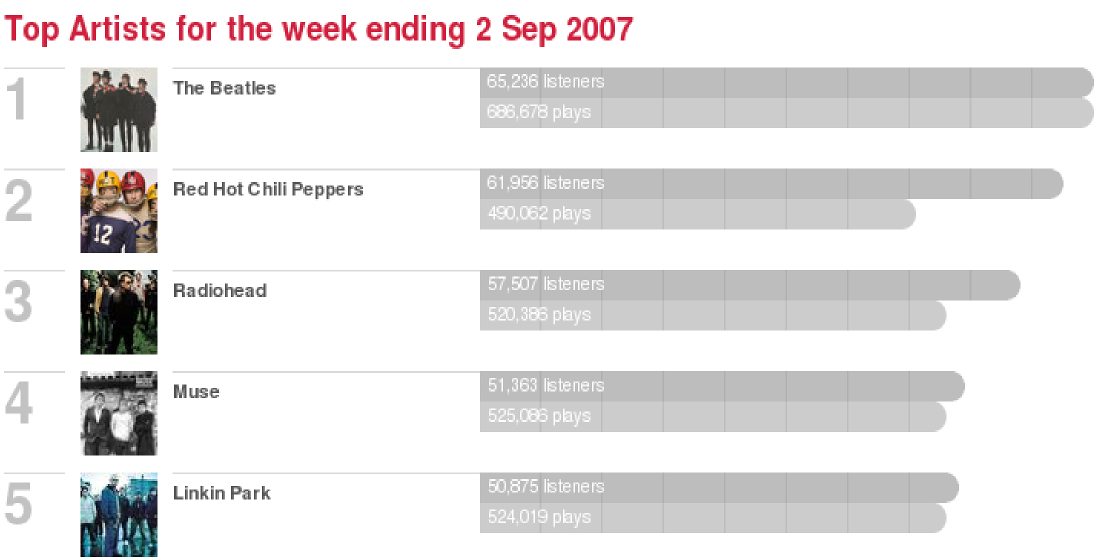

examples:: Social music

Last.fm - Scale

- 20 Million unique visitors per month
- 100 million unique tracks (including misspellings)
- 500 million 'scrobbles' per month
- 2 million tags applied per month
- Streamable tracks – 'millions'
- 100,000 independent artists
- 20,000 labels

examples:: Social music

Last.fm – web services

- Much of last.fm data is available via web services (Creative Commons License)
  - ❖ User Profile Data
  - ❖ Artist Data
  - ❖ Album Data
  - ❖ Track Data
  - ❖ Tag Data
  - ❖ Group Data
  - ❖ Forum Data
- http://www.audioscrobbler.net/data/webservices/

examples:: Social music

Last.fm – web services

http://ws.audioscrobbler.com/1.0/artist/Deerhoof/toptags.xml


```
<toptags artist="Deerhoof">
  <tag>
     <name>indie</name>
     <count>100</count>
     <url>http://www.last.fm/tag/indie</url>
  </tag>
  <tag>
    <name>experimental</name>
    <count>94</count>
    <url>http://www.last.fm/tag/experimental</url>
  </tag>
  <tag>
     <name>indie rock</name>
     <count>60</count>
     <url>http://www.last.fm/tag/indie%20rock</url>
  </tag>
  <tag> ... </tag>
</toptags>
```


Powered by Audioscrobbler™

examples:: Content-based:: Pandora

PANDORA™

- Low touch interface
- Focus is on music listening
- Transparent recommendations

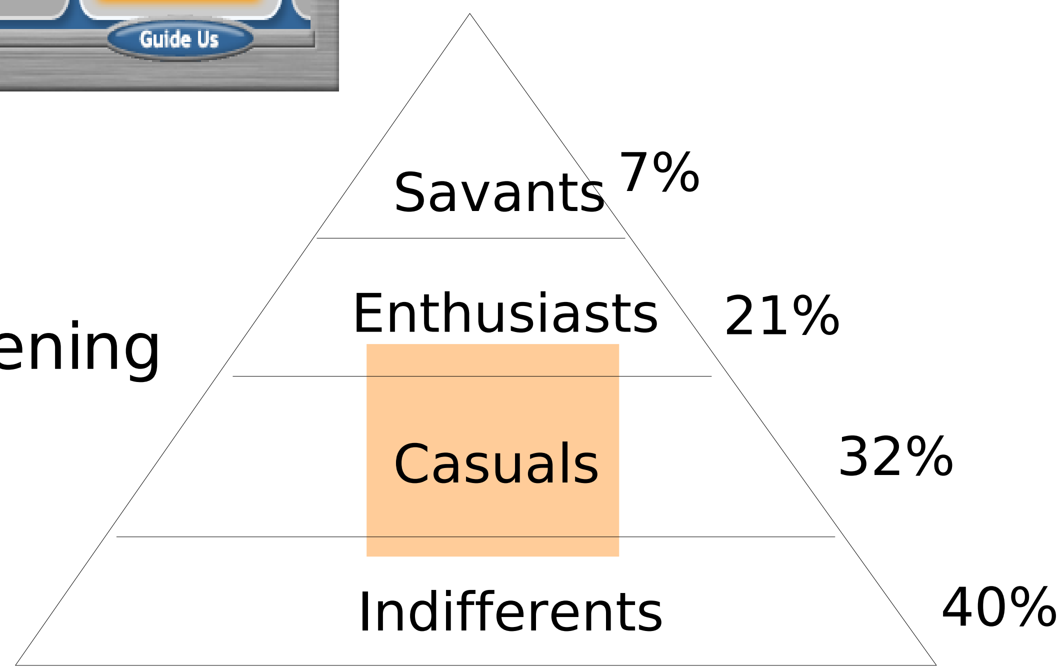

examples:: Content-based:: Pandora

- Transparency
- Technology losing its "cold"

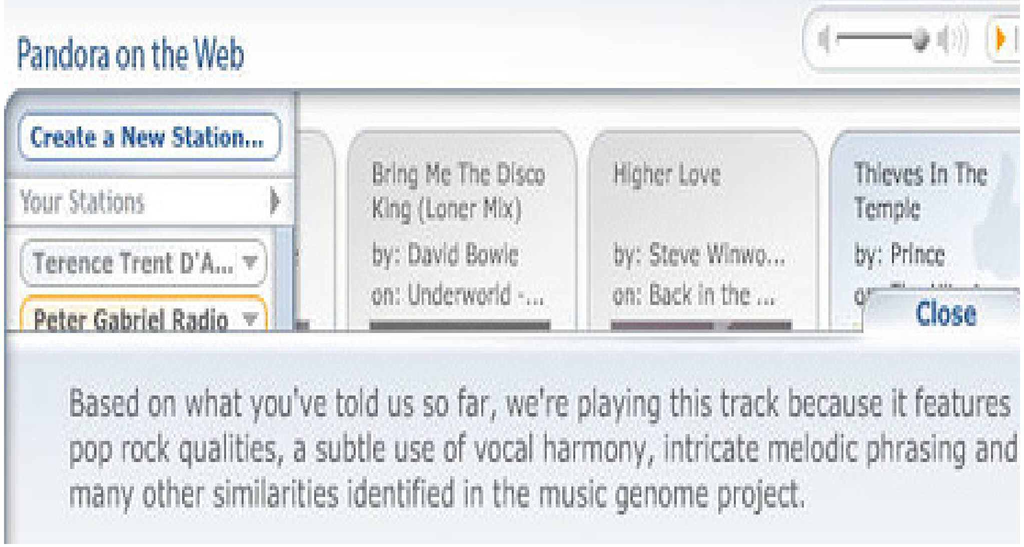

Pandora feels like a smart friend to me. This friend can articulate the reasons I love some of the things I love most (songs) better than I can, but only because I have told it what I like. This is one of my very favorite Prince songs and Pandora knows just why I like it so much. And I didn't know how to say it so well. Makes the technology seem very warm and reflective of my feelings and identity. It's an extension of the user, not a cold, isolating technology. I feel a part of Pandora some times. I'll bet they LOVE this.

http://flickr.com/photos/libraryman/1225285863/

examples:: Content-based:: Pandora

- Pandora: Scale
  - 35,000 artists
  - 500,000 songs
  - 15,000 songs analyzed per month
  - 8 Million registered users

examples:: Content-based:: Pandora

- Pandora: How does it work?
  - Curators select music to add to the catalog
  - Curators attach metadata (from AMG)
  - Music analysts characterize tracks across 400 features
  - Simple weighed Euclidean distance used to find similar songs
  - Playlists generated from seed songs/artists conforming to sets of rules

examples:: Content-based:: Pandora

- Enrolling a new CD
  - ❖ Phase 1: Curator
    - Curator Selects CD based on
      - Proof of audience
      - Curator judgment
      - “Does a song make listening better or worse?”
    - Curator rips CD, attaches metadata
  - ❖ Phase 2: Analysis
    - 160 Hours of training
    - Use double analysis to verify consistency of analysis
    - Analysis can be a collaborative process

examples:: Content-based:: Pandora

- Pandora analysis
  - ❖ Typically analyze 4 songs per album -
    - Chose songs that are representative of an artists career
  - ❖ Include music outliers
  - ❖ Data Entry – 400 attributes with a 10 point scale:
    - [0-1-2-3-4-5-6-7-8-9-10] – Back beat prominence
    - [0-1-2-3-4-5-6-7-8-9-10] – Electric Guitar Wah-Wah
    - [0-1-2-3-4-5-6-7-8-9-10] – light or breathy vocals
  - ❖ 400 Attributes are a trade secret

examples:: Content-based:: Pandora

- Pandora analysis
  - Dolly Parton – Stairway to heaven
    - country influences
    - bluegrass influences
    - folk influences
    - a subtle use of vocal harmony
    - mild rhythmic syncopation
    - acoustic sonority
    - demanding instrumental part writing
    - intricate melodic phrasing
    - thru composed melodic style
    - a clear focus on recording studio production
    - minor key tonality
    - melodic songwriting
    - a heavy twang in the vocals
    - acoustic rhythm guitars

Curiously, no Pandora attributes are given for Led Zeppelin's version

examples:: Content-based:: Pandora

- Pandora recommendation
  - 400 song parameters form euclidean space
  - Genre specific weights
    - Problems with cross-genre recommendations
  - Simple nearest neighbors for song selection – filtered for:
    - licensing compliance
    - mix of familiarity, variety, discovery
  - For artist similarity use specific songs – not an average of all songs.

examples:: Content-based:: Pandora

- Pandora goes social
  - Crowd understands things that the genome doesn't
  - CF used initially as a safety net
  - Started using 'thumbs-down' data to filter out songs
  - 'Thumbs-up data' correlates with 'familiarity'
  - Use 'familiarity' to select songs
  - Result: “Playlists are massively improved”

examples:: Content-based:: Pandora

- The New Pandora
  - Bigger is not always better
  - A Radio Station not a Recommender
    - Precision important, recall not.
  - All forms of song selection are good
  - New Hybrid approach:
    - Much happier listeners
    - Avoid some of the CF problems – coldstart and 'early rater' feedback loops
    - No longer strictly a Content-based recommender

examples:: Content-based:: Pandora

- The New Pandora
  - Pandora futures
    - Machine listening to add dimensions to their data
    - Social tags
  - Pandora issues
    - Adding new genres
    - Cross genre artists (the 'shakira' problem)
    - New music features – 'scratching'

examples:: Hybrid:: Mystrands

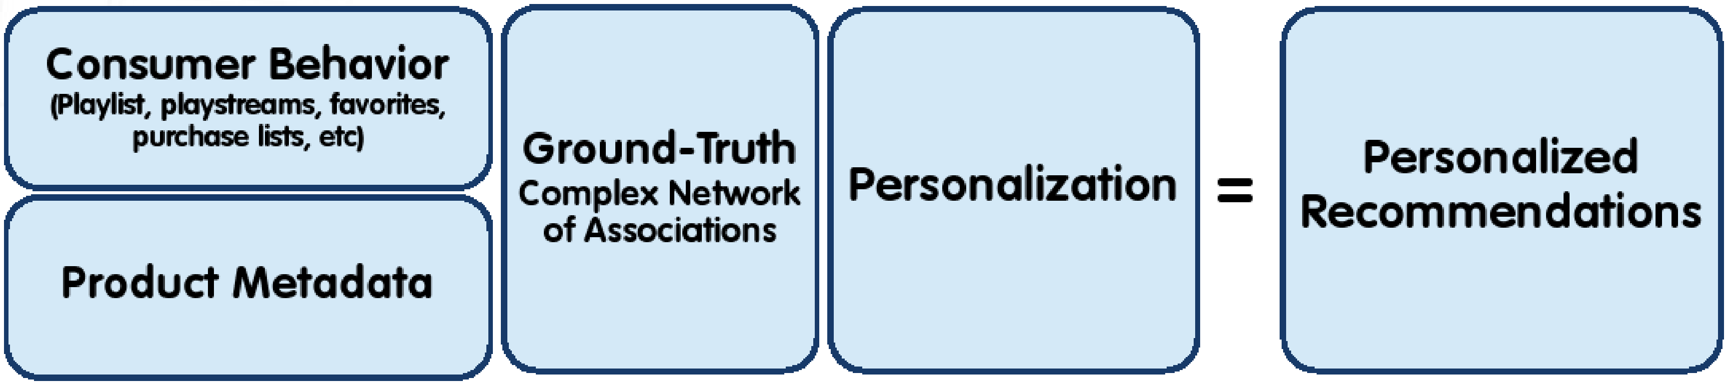

- Use
  - Precomputed top item-item correlations
  - Metadata – genre, year, popularity, tags, etc
  - User History – plays, skips, put on a playlist
  - User libraries

examples:: Hybrid:: One Llama

- Creating a Hybrid recommender
  - Content-based similarity:
    - Features = spectral, psycho-acoustic, MFCCs
    - Genetic algorithms for feature selection
    - Similarity metric – akin to VQ or hamming
    - Training similarity on production music
    - Fast similarity search – 200ms for 500K catalog
  - Social-based similarity:
    - Crawling web for playlists for co-occurrence similarity
- Still early days
- Strong ties to the MIR community (IMIRSEL)

ONE LLAMA BETA
LOGIN | REGISTER

examples:: Hybrid:: One Llama

- One Llama – playlist generator

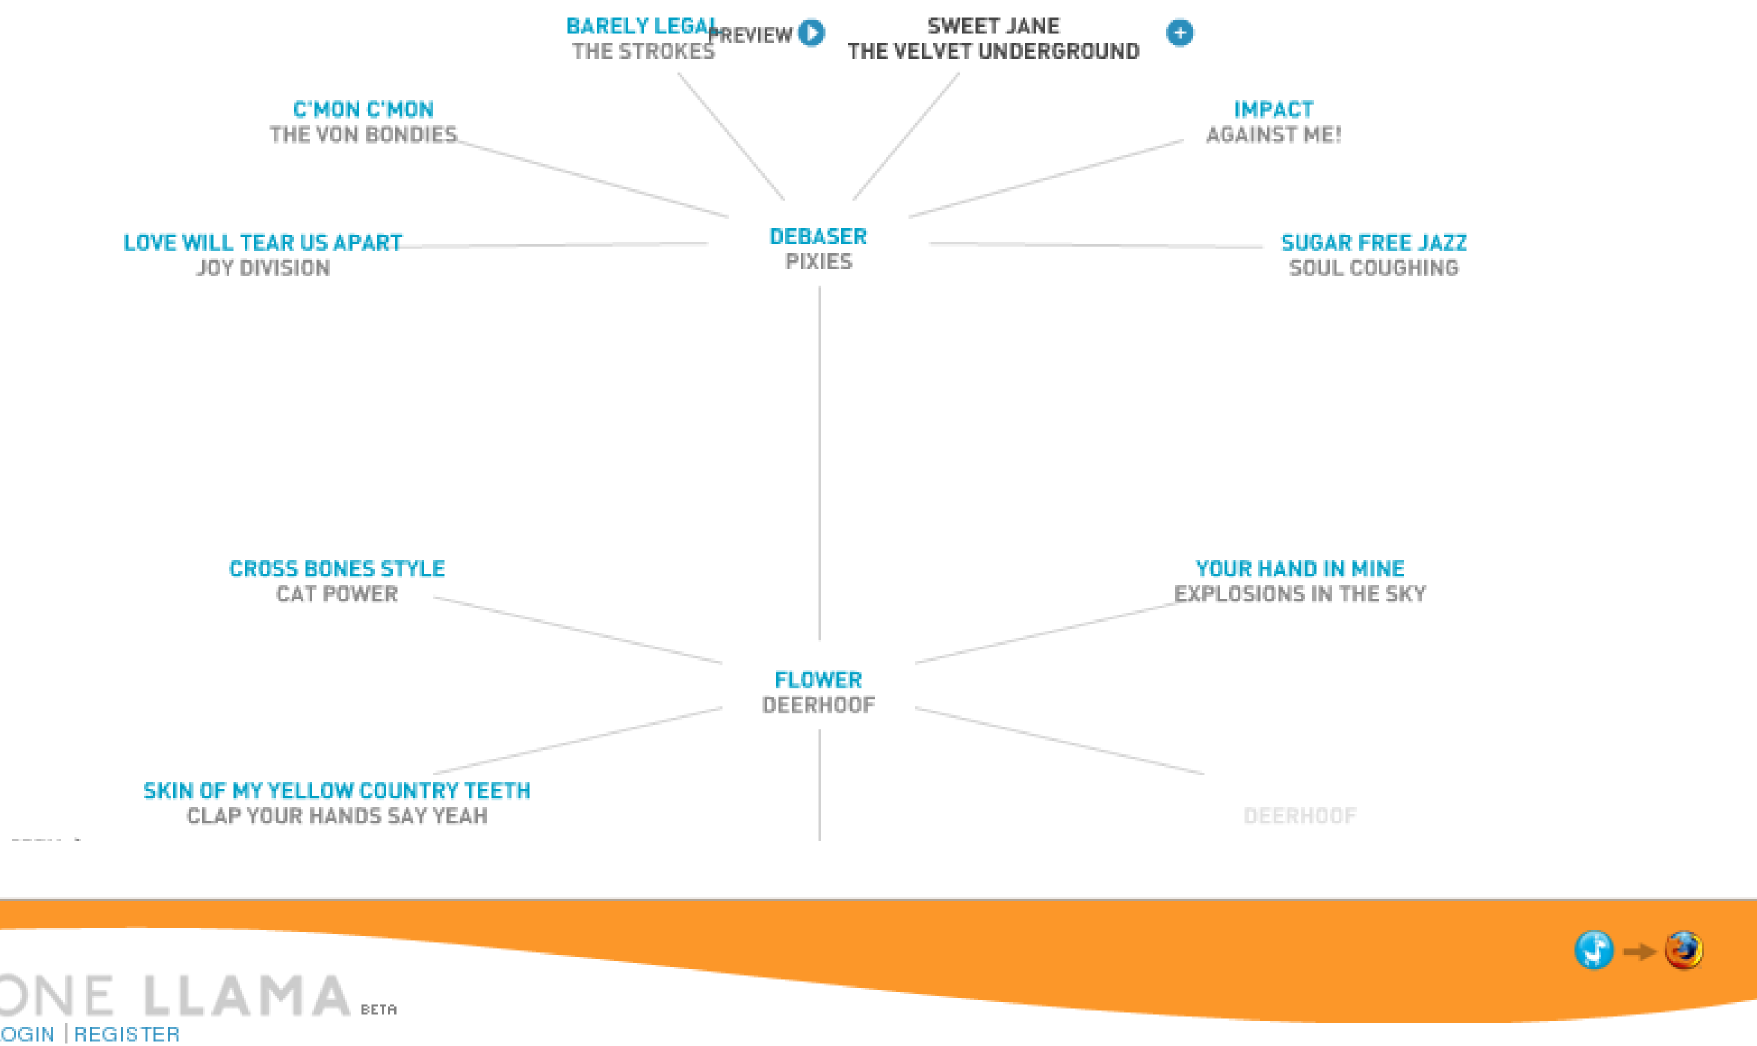

examples:: Hybrid:: BMAT

- spin-off of the MTG, started in 2006

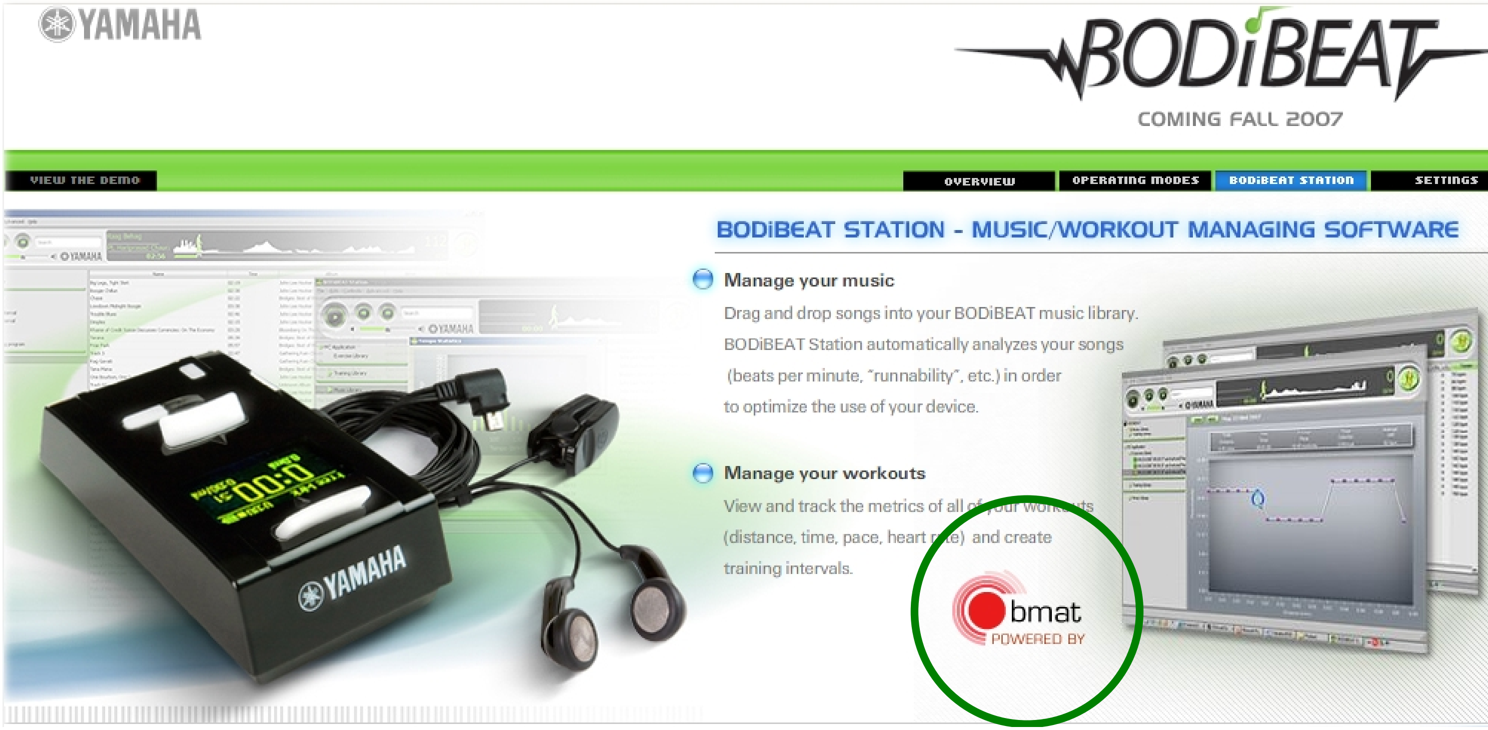

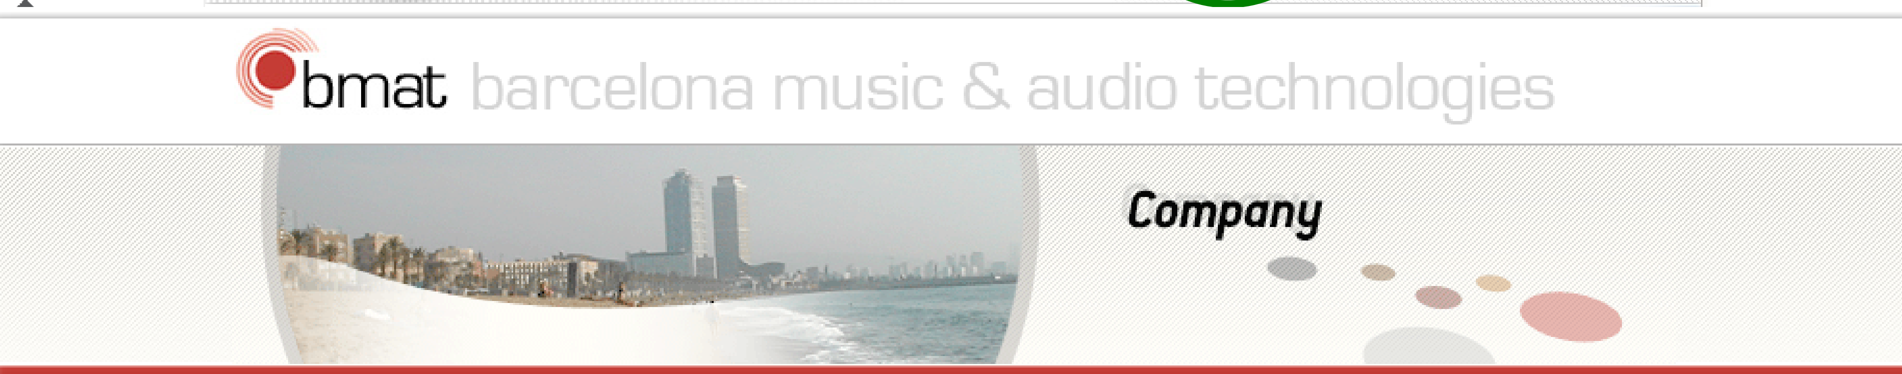

outline

- Introduction
- Formalization of the recommendation problem
- Recommendation algorithms
- Problems with recommenders
- Recommender examples
- **Evaluation of recommenders**
- Conclusions / Future

evaluation

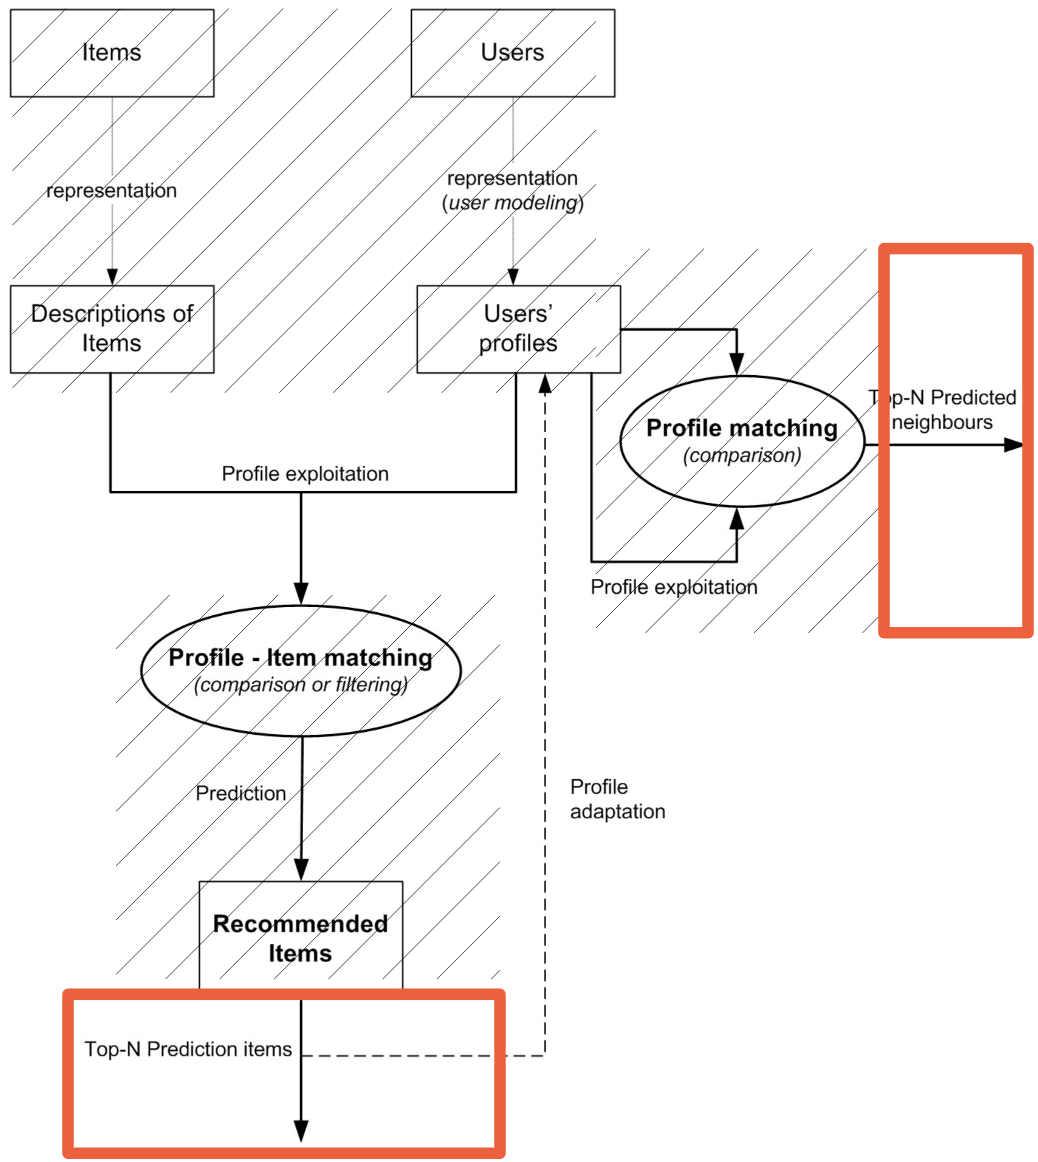
Items
Users
representation
representation
(user modeling)
Descriptions of Items
Users' profiles
Profile matching
(comparison)
Top-N Predicted neighbours
Profile exploitation
Profile exploitation
Profile - Item matching
(comparison or filtering)
Prediction
Profile adaptation
Recommended Items
Top-N Prediction items

evaluation

- Is it possible to create a *standard* ground truth?
- Need of a music dataset
  - such as Netflix for movies (ask last.fm :-)
  - split dataset (train / test set)?
- Goals
  1) measuring the quality of the items recommended to a given user
  2) measuring *how good* is the music collection
- Constraints
  1) do not recommend an item if the user has previously *purchased / listened to / etc.* that item

evaluation

- Outline
  - common metrics
  - new metrics to exploit the long tail (popularity)
  - complex network analysis
  - informal survey of different recommenders

evaluation:: common metrics

- Accuracy metrics
  - Statistical
    - measure deviation between prediction and the actual *rating*
    - Examples
      - Mean Absolute Error (MAE)
      - Root Mean Squared Error (RMSE) **Netflix Prize**
      - Correlation
  - Decision support
    - measure the selection of high-quality items from the whole list
    - Examples
      - (area of the) ROC curve
        - trade-off between True Positives and False Positives
      - Customer ROC curve
        - constrained to recommend the same number of items to each user

$$|E_{MAE}| = \frac{1}{N}\sum_{i=1}^{N}|x_i - \hat{x}_i|$$

$$\mathrm{MSE}(\hat{\theta}) = \mathrm{E}((\hat{\theta} - \theta)^2).$$

$$\mathrm{RMSE}(\hat{\theta}) = \sqrt{\mathrm{MSE}(\hat{\theta})}.$$

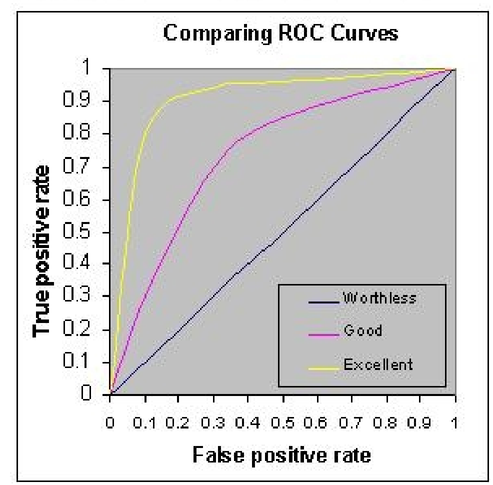

evaluation:: common metrics

- Relevance metric
  - Precision = TP / TP + FN
- Coverage metric
  - Recall = TP / TP + FP
- ...and both
  - F-measure
- Ranking quality of the recommendation list
  - Rank score
    - combine the hits and their position
  - if the training dataset contains ordered lists:
    - kendall tau
    - spearman rho

evaluation:: limitations

- Limitations of the current metrics
  - skewness
    - performed on test data that users *chose* to rate
  - do not take into account
    - usefulness
    - novelty / serendipity
    - *goodness* of the collection
      - analysis of the items' relationships

evaluation:: limitations

- Limitations of the current metrics
  - other components (user-dependent)

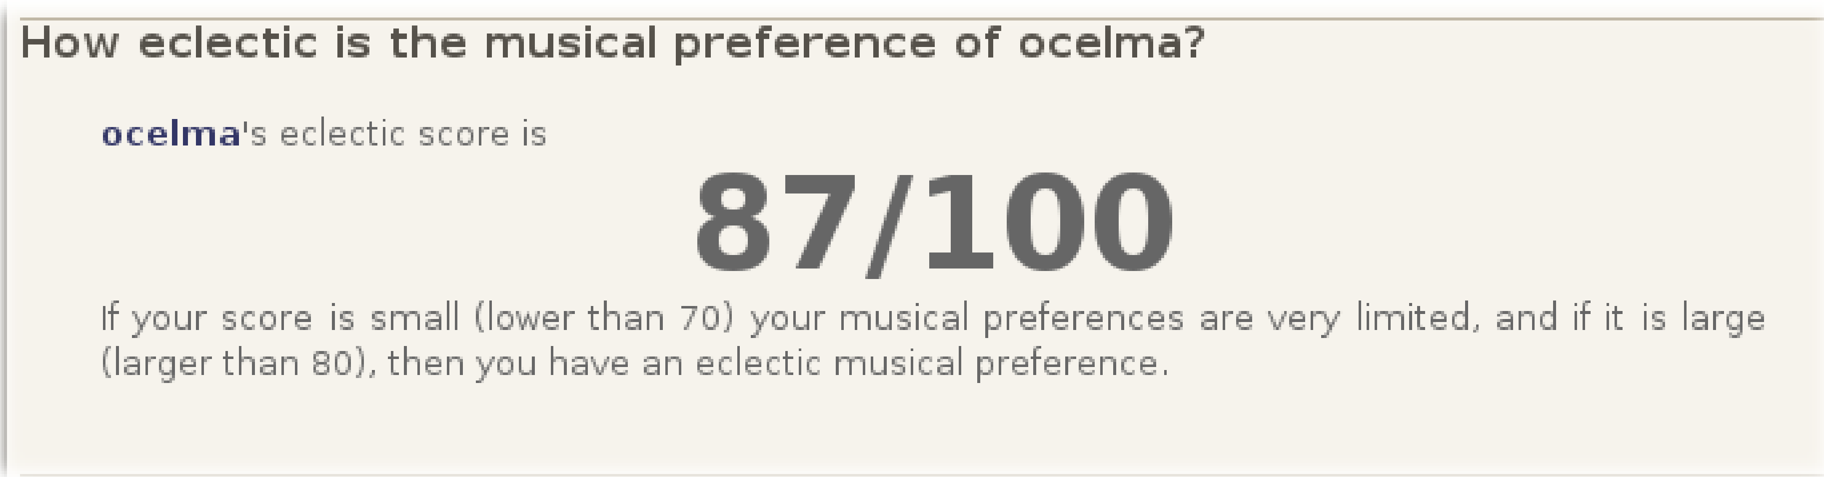

http://anthony.liekens.net/pub/scripts/last.fm/eclectic.php *http://mainstream.vincentahrend.com/*

Mainstream-O-Meter

Username: ocelma ok

Powered by Audioscrobbler

	Artist	Mainstreamness	Listeners	Weight
1.	The Dogs D'Amour	0.32 %	2,597	100 %
2.	U2	81.49 %	669,199	48 %
3.	The White Stripes	86.10 %	707,040	24 %
4.	Spiritualized	7.89 %	64,812	20 %
5.	Yann Tiersen	21.21 %	174,175	19 %
6.	The Black Crowes	9.55 %	78,438	16 %
7.	Lhasa	3.40 %	27,950	11 %
8.	Ryan Adams	22.22 %	182,505	10 %
9.	Martirio y Chano Domínguez	0.01 %	94	8 %
10.	The Rolling Stones	64.58 %	530,305	6 %
11.	Bob Dylan	59.02 %	484,687	6 %
12.	Kraftwerk	23.20 %	190,529	6 %
13.	Nirvana	85.65 %	703,330	4 %
14.	Björk	56.32 %	462,460	4 %
15.	Radiohead	103.00 %	845,815	4 %

30.82 % mainstream

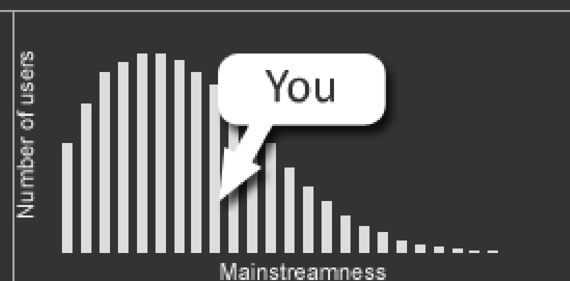

evaluation:: limitations

- New proposed metric: “novelty+relevance”
  - novelty (serendipity “Oh!”)
    - How? exploit the long-tail of the collection (popularity)
  - ...but still *relevant* to the user

evaluation:: novelty & relevance

- Dealing with novelty
  - Experiment with **last.fm** data
    - 249,753 artists
    - for each artist, get
      - total number of plays, and
      - similar artists (3,846,262 of relationships)

evaluation:: novelty & relevance

- **Last.fm** long-tail analysis (#artists plays)
  - data from July 2007

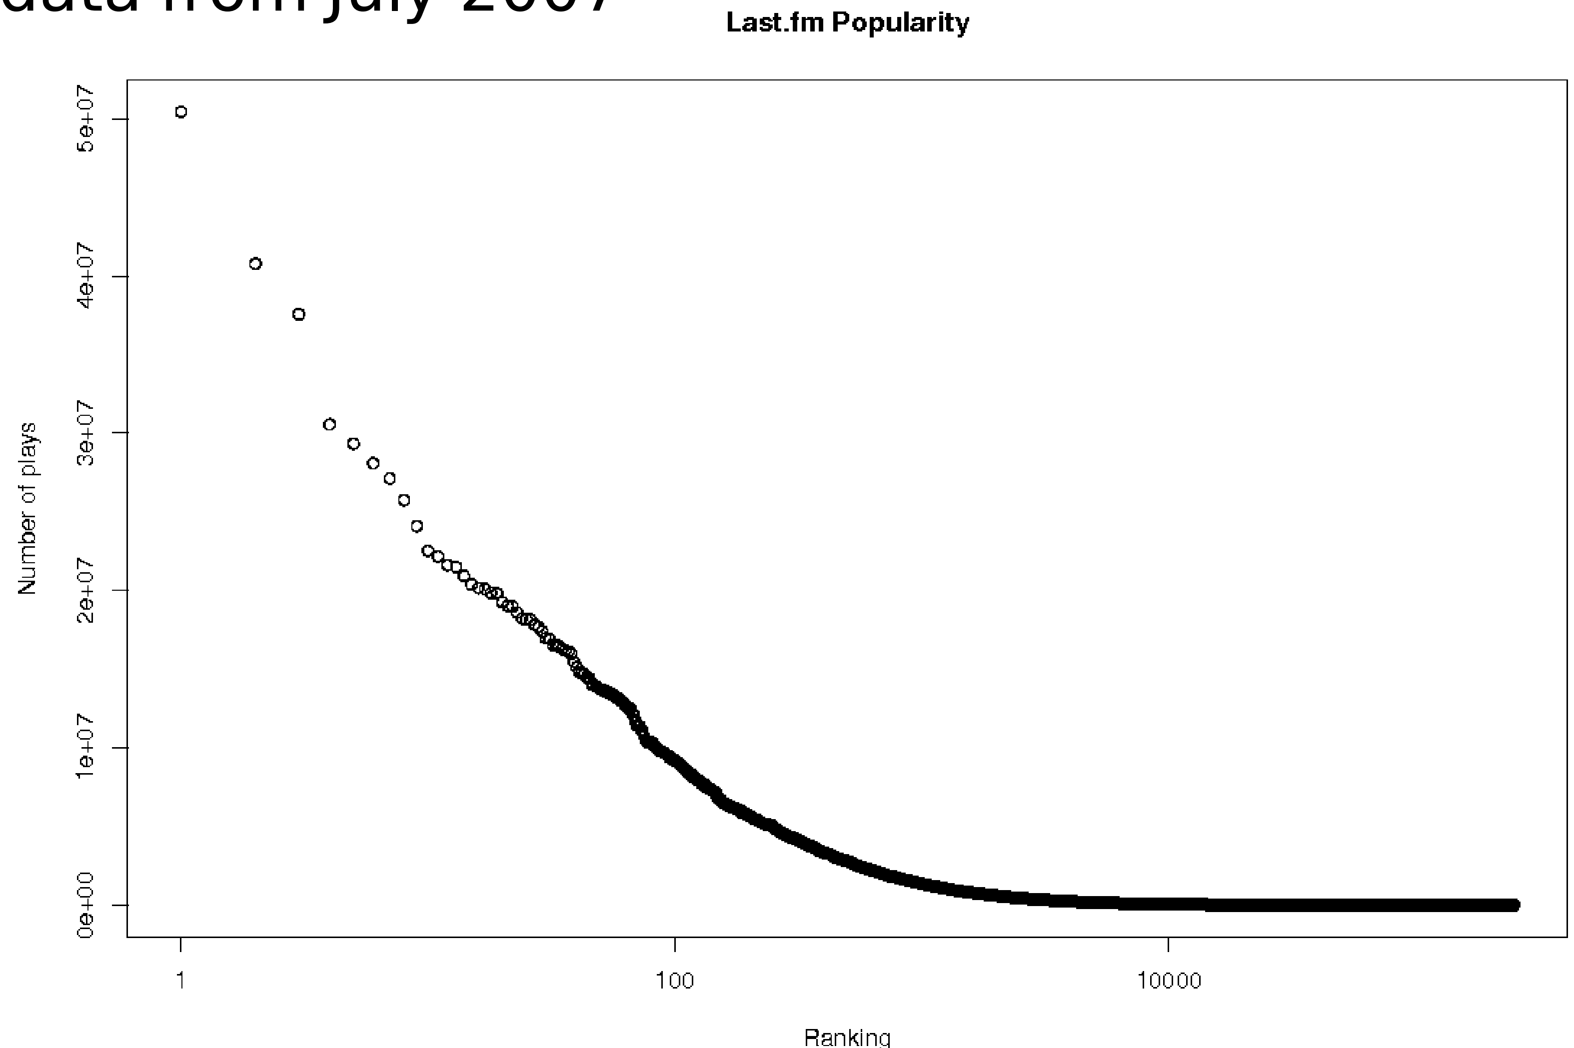

evaluation:: novelty & relevance

- **Last.fm** long-tail analysis (#artists plays)
  - 1st Example: explore the **long tail**, by means of content-based audio similarity

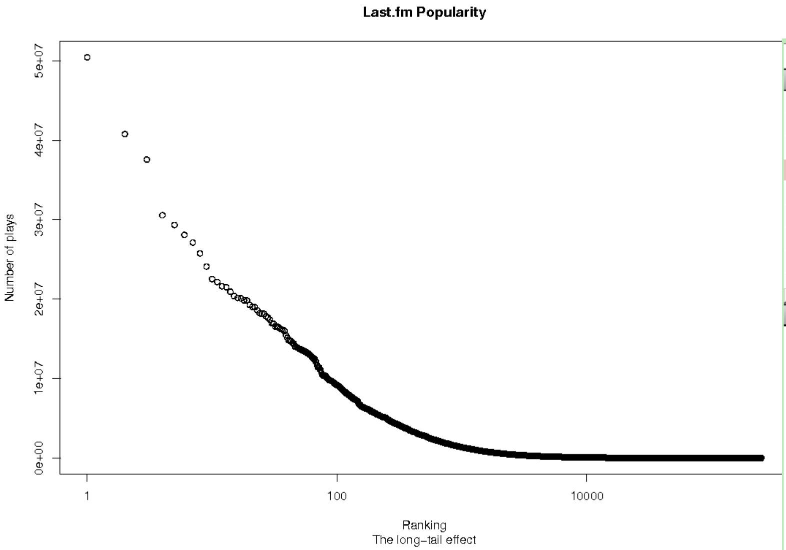

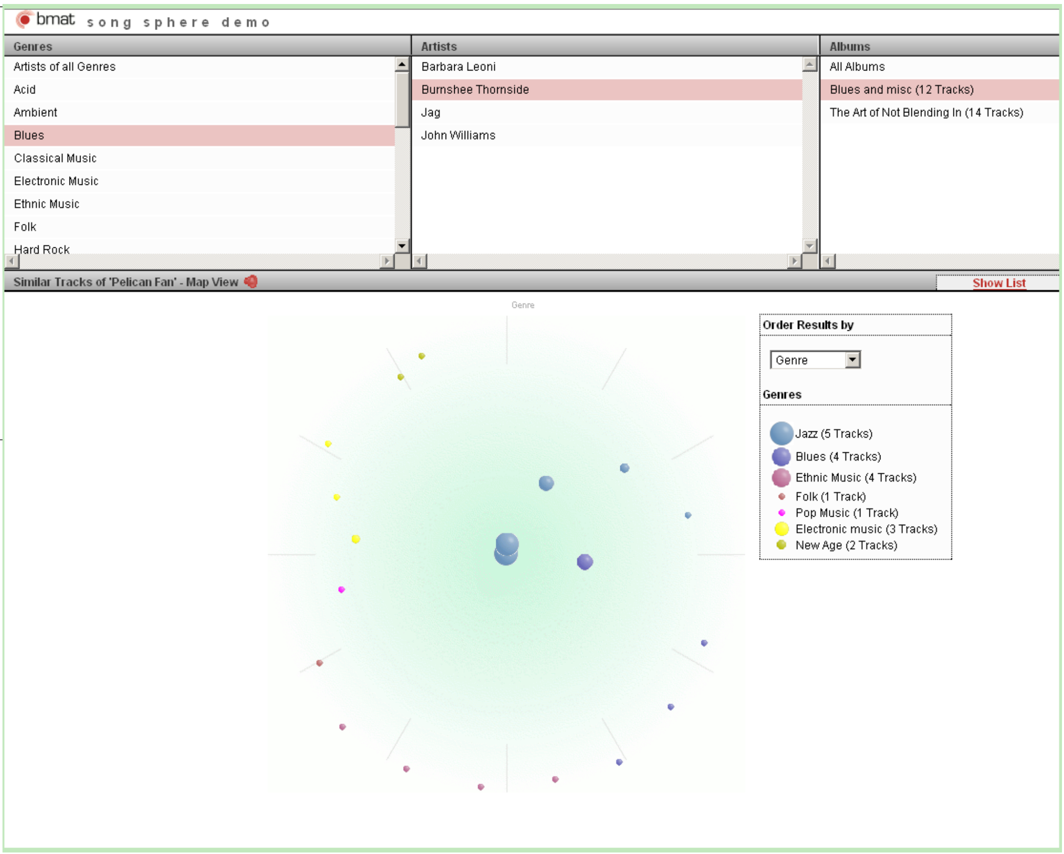

evaluation:: novelty & relevance

- Bruce Springsteen
  - total # plays in last.fm = 5,992,068
  - # plays for “Better days” (seed song) = 33,995
  - data from July 2007

evaluation:: novelty & relevance

- **Last.fm** long-tail analysis (#artists plays)
  - Bruce Springsteen

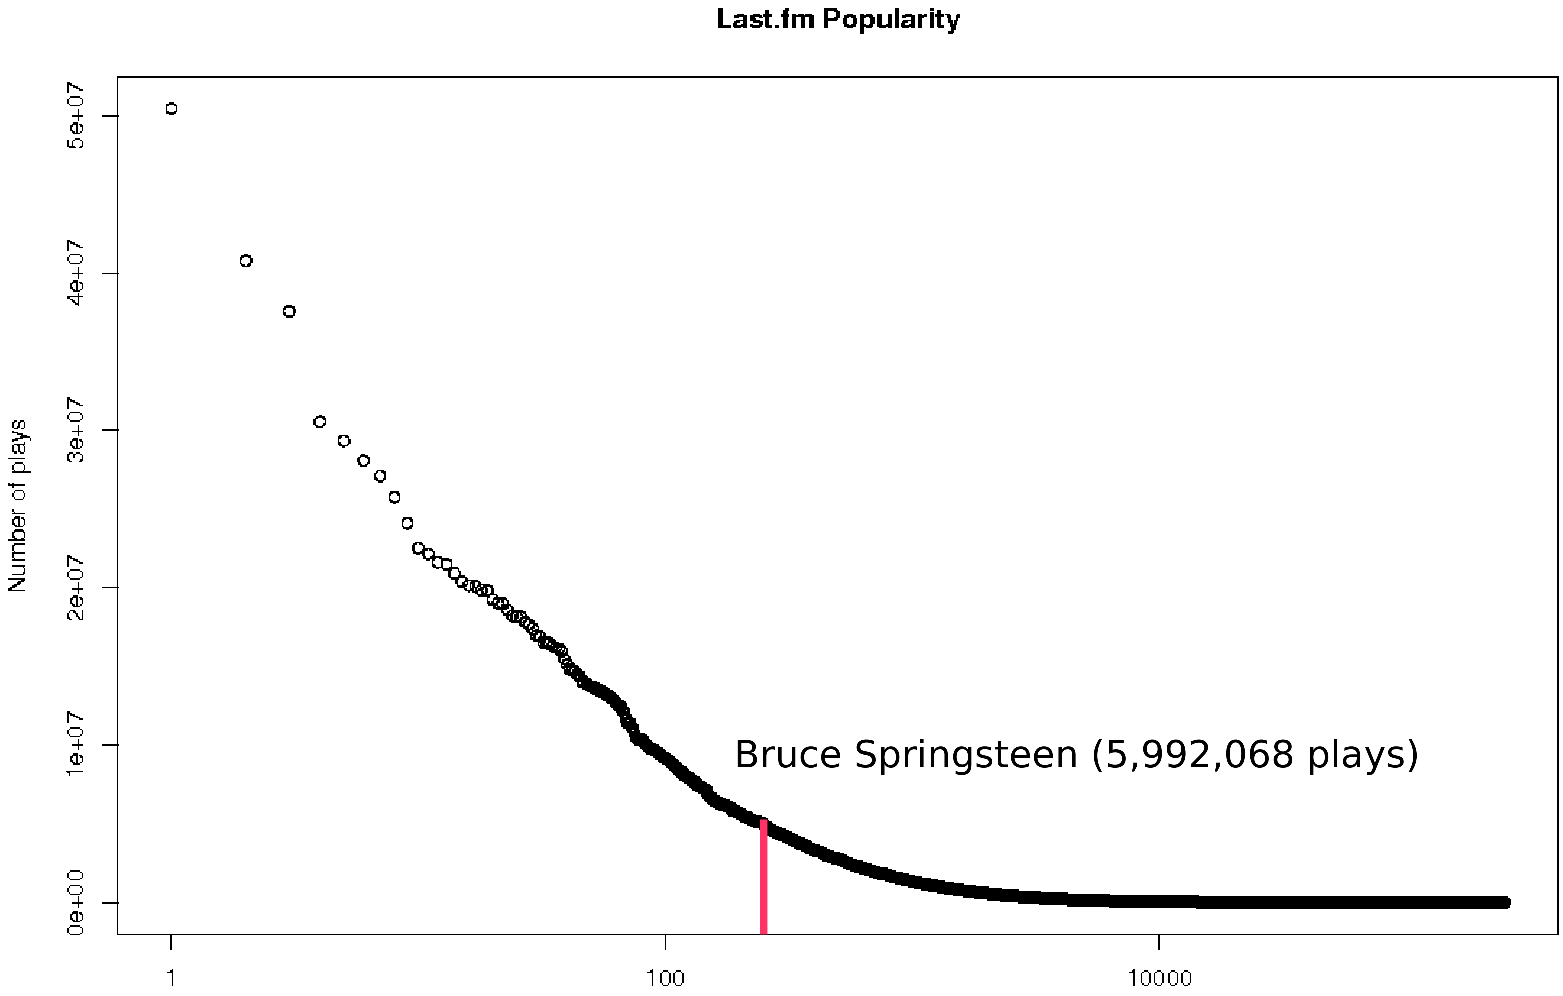

evaluation:: novelty & relevance

- The Rolling Stones
  - ❖ total # plays in last.fm = 11,239,824
  - ❖ # plays for “Mixed emotions” = 50,778
  - ❖ data from July 2007

evaluation:: novelty & relevance

- **Last.fm** long-tail analysis (#artists plays)
  - The Rolling Stones

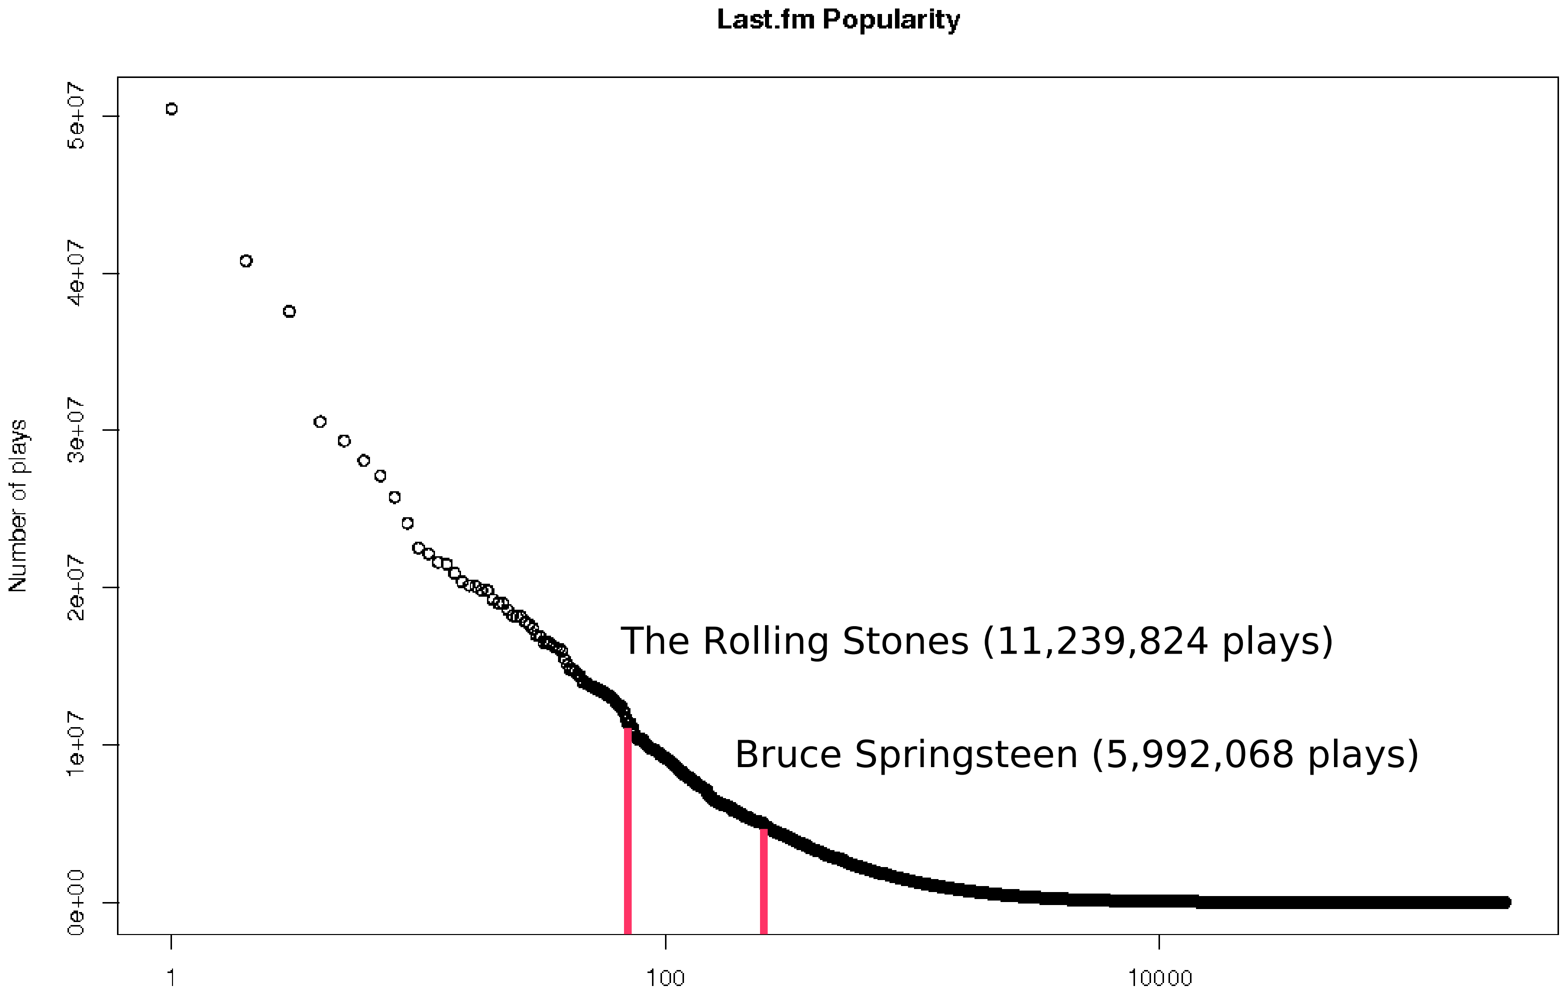

evaluation:: novelty & relevance

- Mike Shupp
  - total # plays in last.fm = 321
  - # plays for “Letter to Annete” = 0
  - data from July 2007

evaluation:: novelty & relevance

- **Last.fm** long-tail analysis (#artists plays)
  - ❖ Mike Shupp

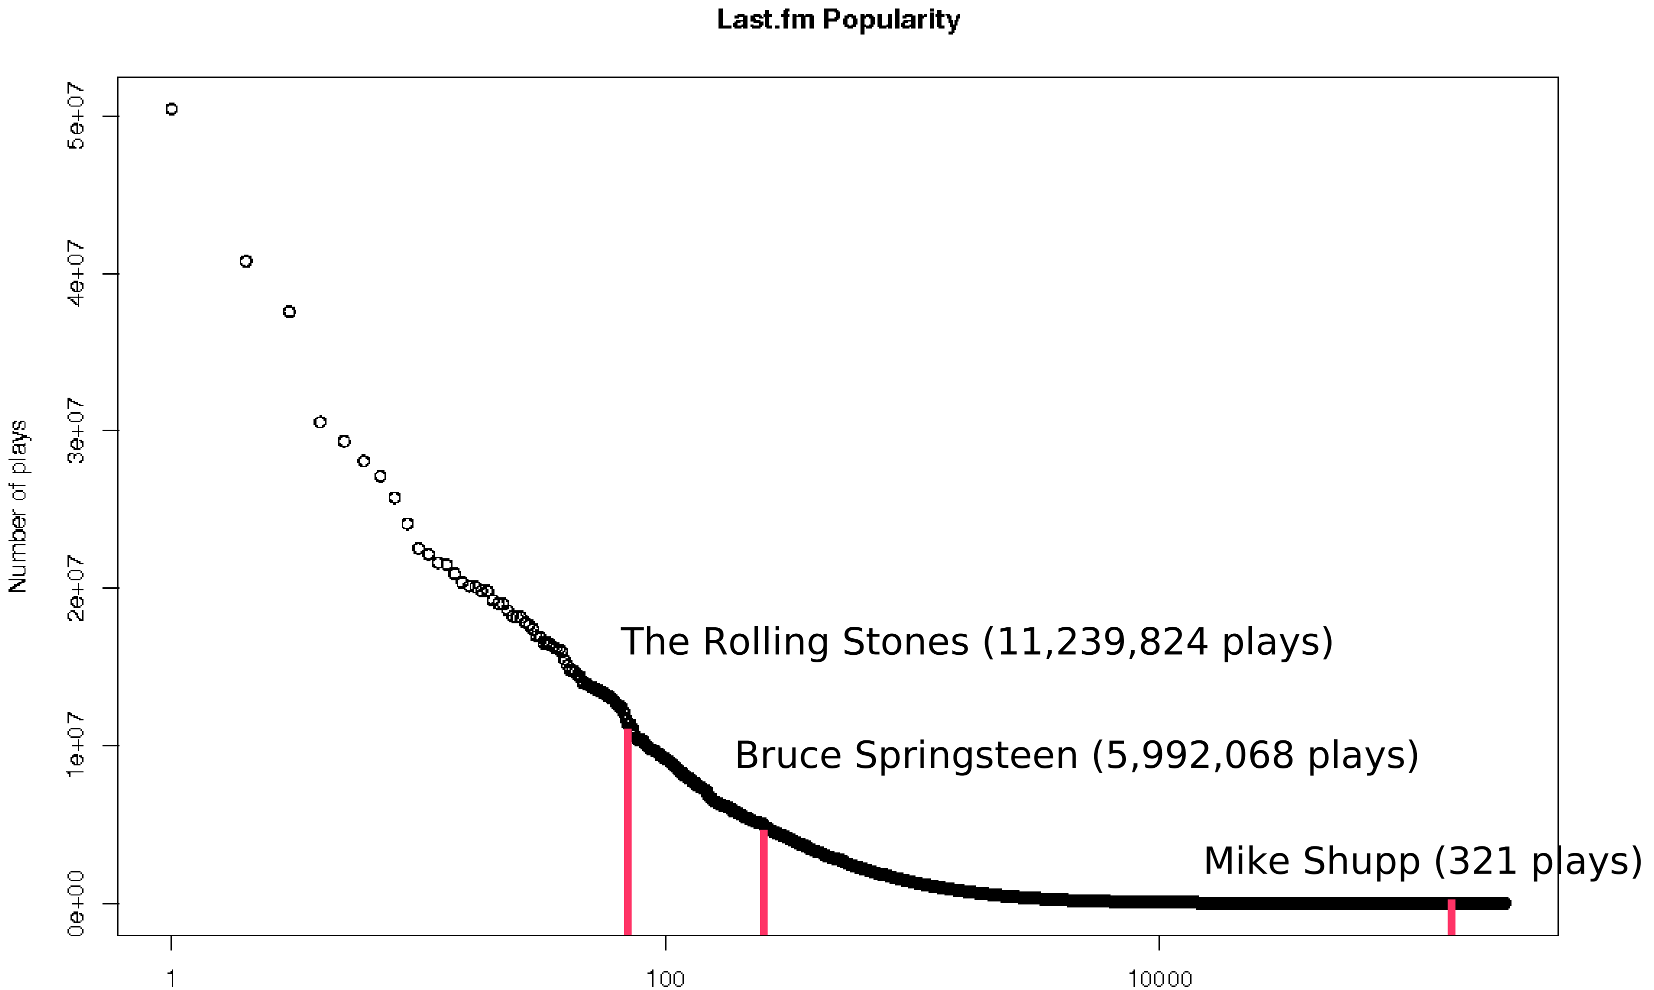

evaluation:: novelty & relevance

- Using CB similarity
  - Bruce Springsteen -> The Rolling Stones -> Mike Shupp

⇒ with collaborative filtering we would *never* reach Mike Shupp:

- Shortest path in the **last.fm** graph
  - Directed graph
    - Infinite! (in different graph components)
  - Undirected graph
    - Bruce Springsteen<->Steve Kilbey<->Mike Shupp

evaluation:: novelty & relevance

- CB democratizes the music, but who's voting?

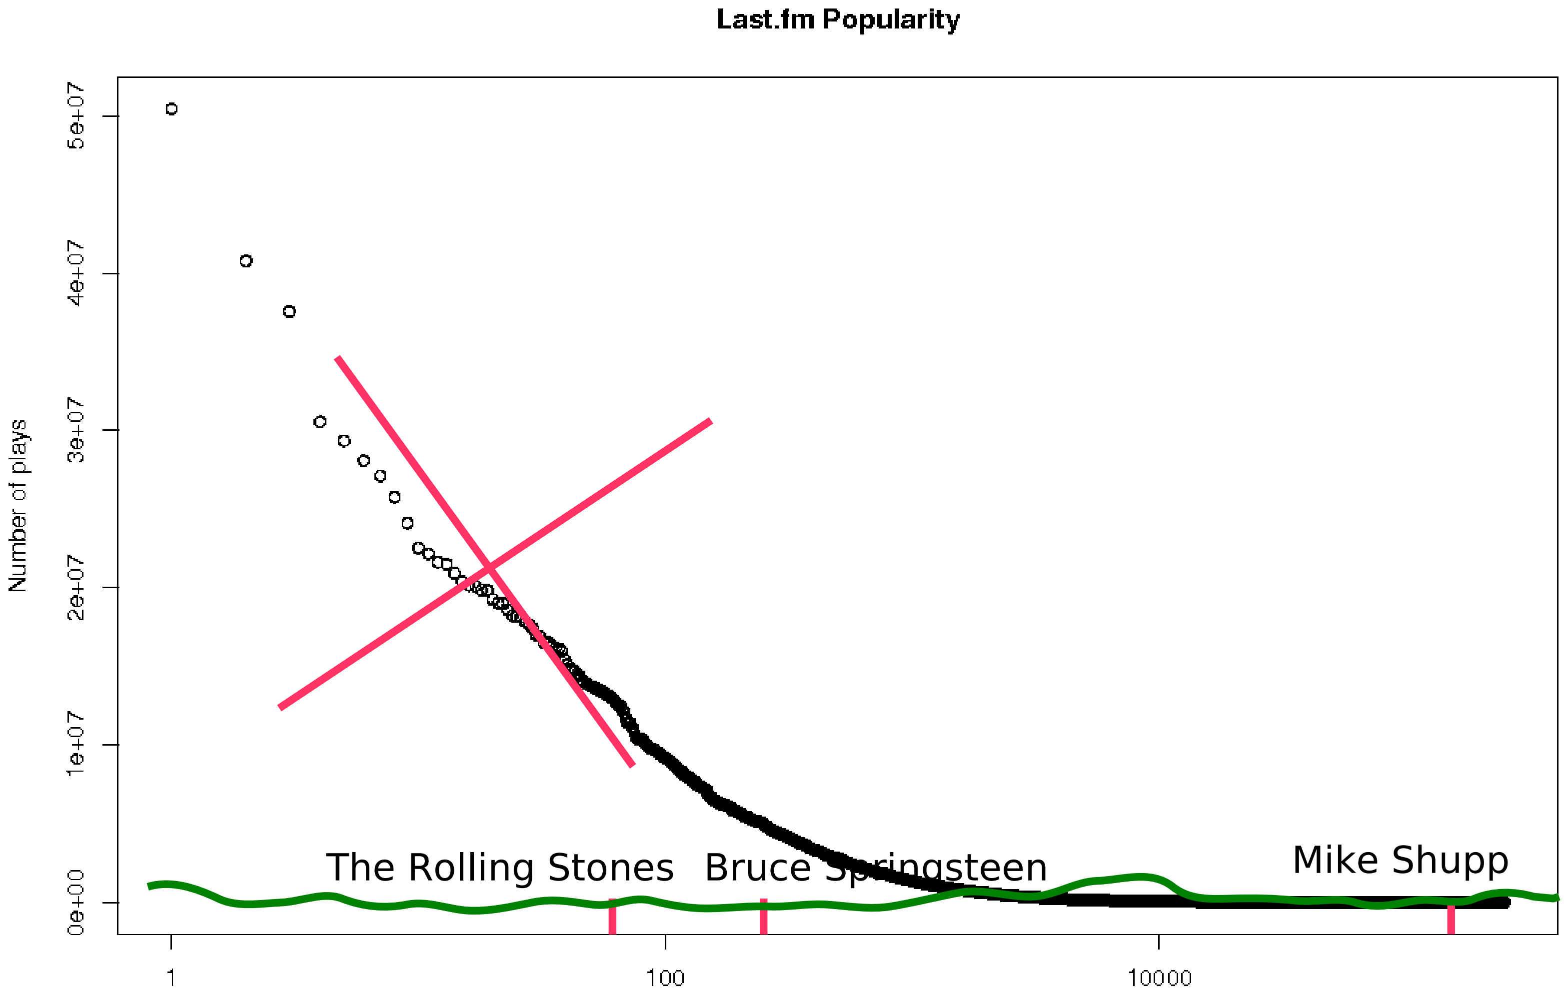

evaluation:: novelty & relevance

- And it seems that CB was not that wrong...
  - Mike Shupp review
    - “Letter to Annette”, “Right For You” and eight more (...). It's comforting to know that melodic rock and roll is still alive and kicking in the US (...) guitarist/songwriter Mike Shupp is deft with the straightahead country-inflected pop-rock that the likes of Paul Westerberg, **Bruce Springsteen**, Steve Forbert and Neil Young are renowned for. (...) -- *Kevin Mathews*

- Now, let's analyze the relationships between the long tail and the similar artists...

evaluation:: novelty & relevance

- **Last.fm** long-tail analysis (#artists plays)
  - 249,753 artists (data from July 2007)

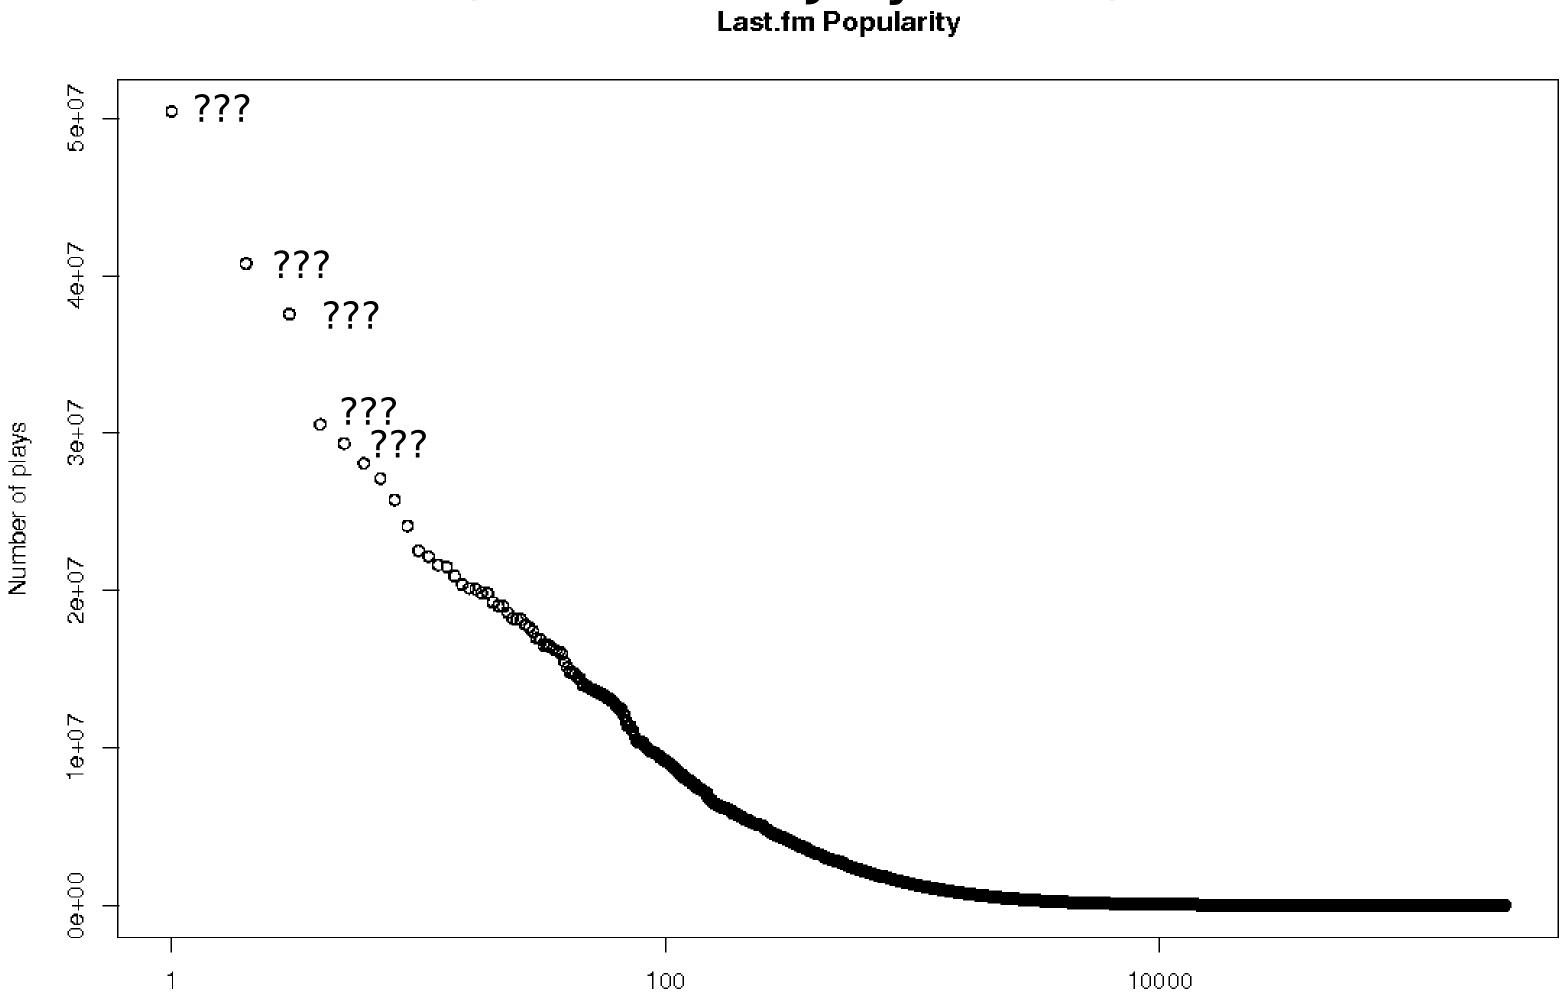

evaluation:: novelty & relevance

- **Last.fm** long-tail analysis (#artists plays)
  - 249,753 artists (data from July 2007)

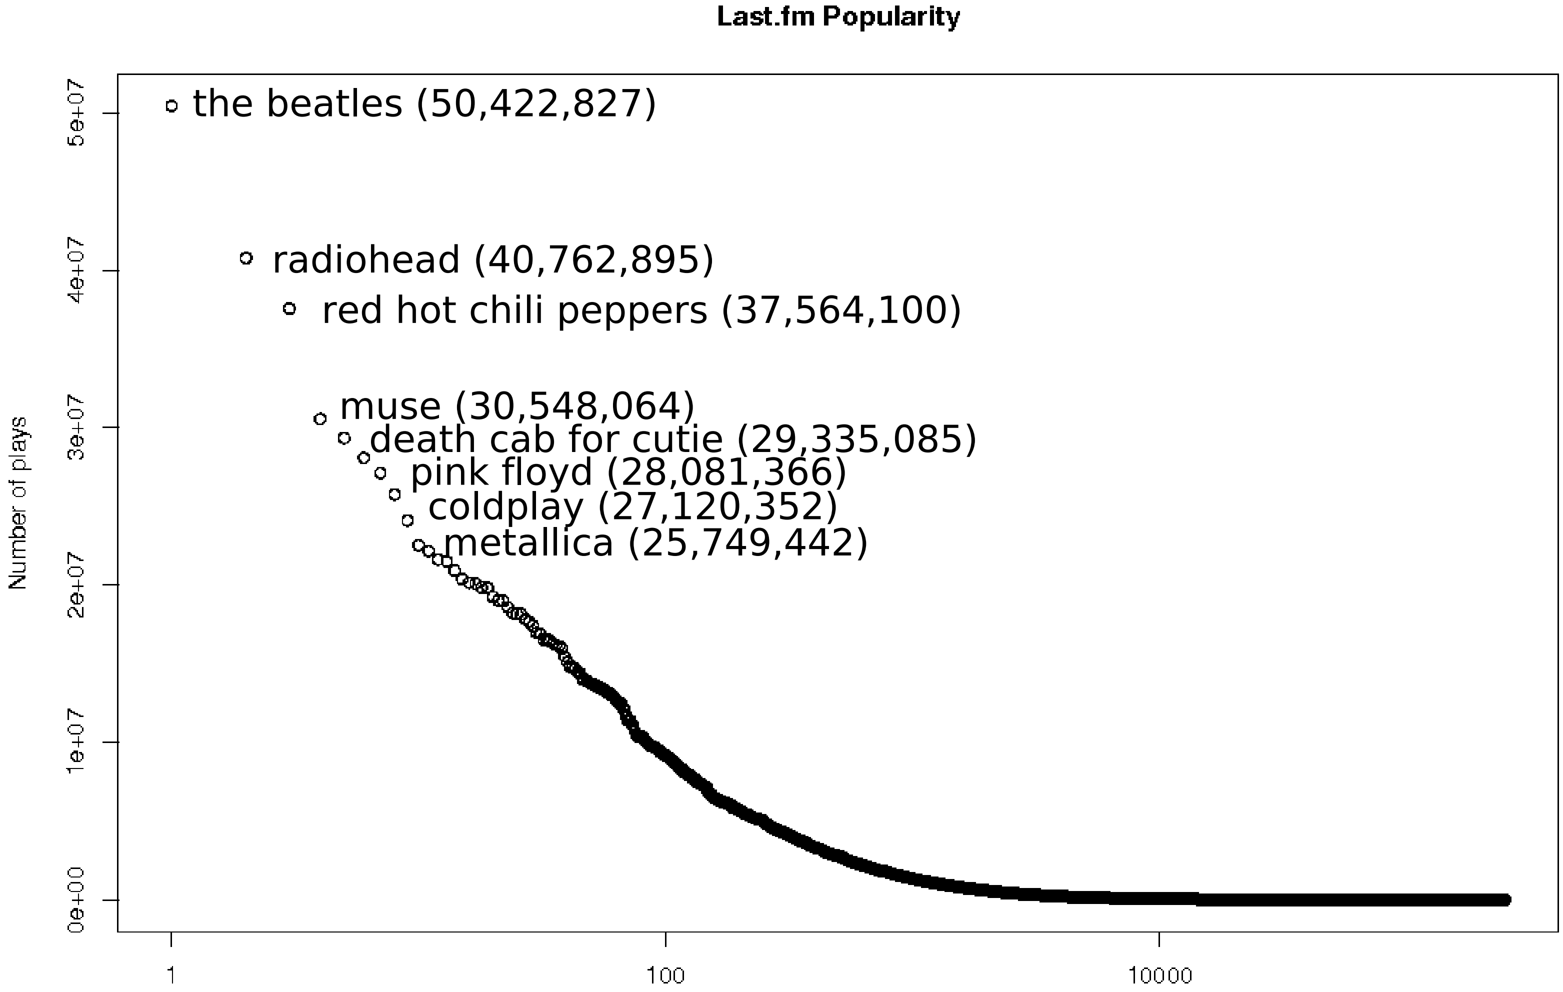

evaluation:: novelty & relevance

- **Last.fm** long-tail model [Kilkki, K., 2007]
  - cumulative percentage

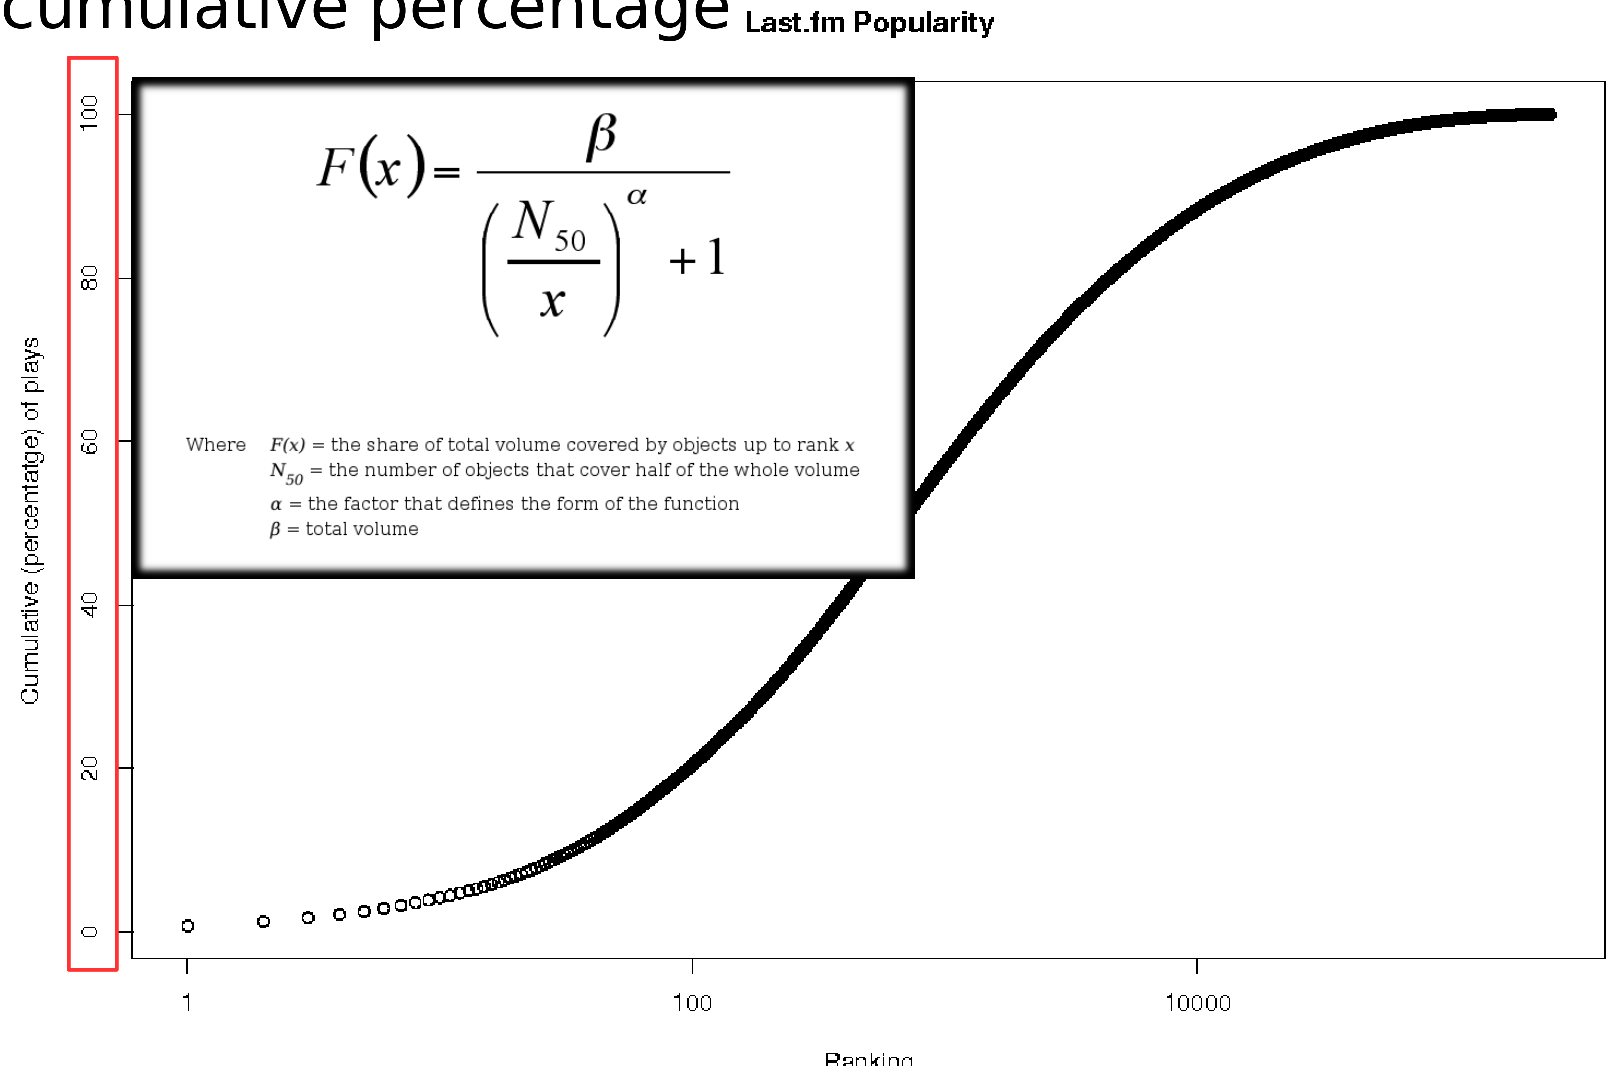

$$F(x) = \frac{\beta}{\left(\frac{N_{50}}{x}\right)^{\alpha} + 1}$$

Where $F(x)$ = the share of total volume covered by objects up to rank x
N_{50} = the number of objects that cover half of the whole volume
α = the factor that defines the form of the function
β = total volume

evaluation:: novelty & relevance

- **Last.fm** long-tail model

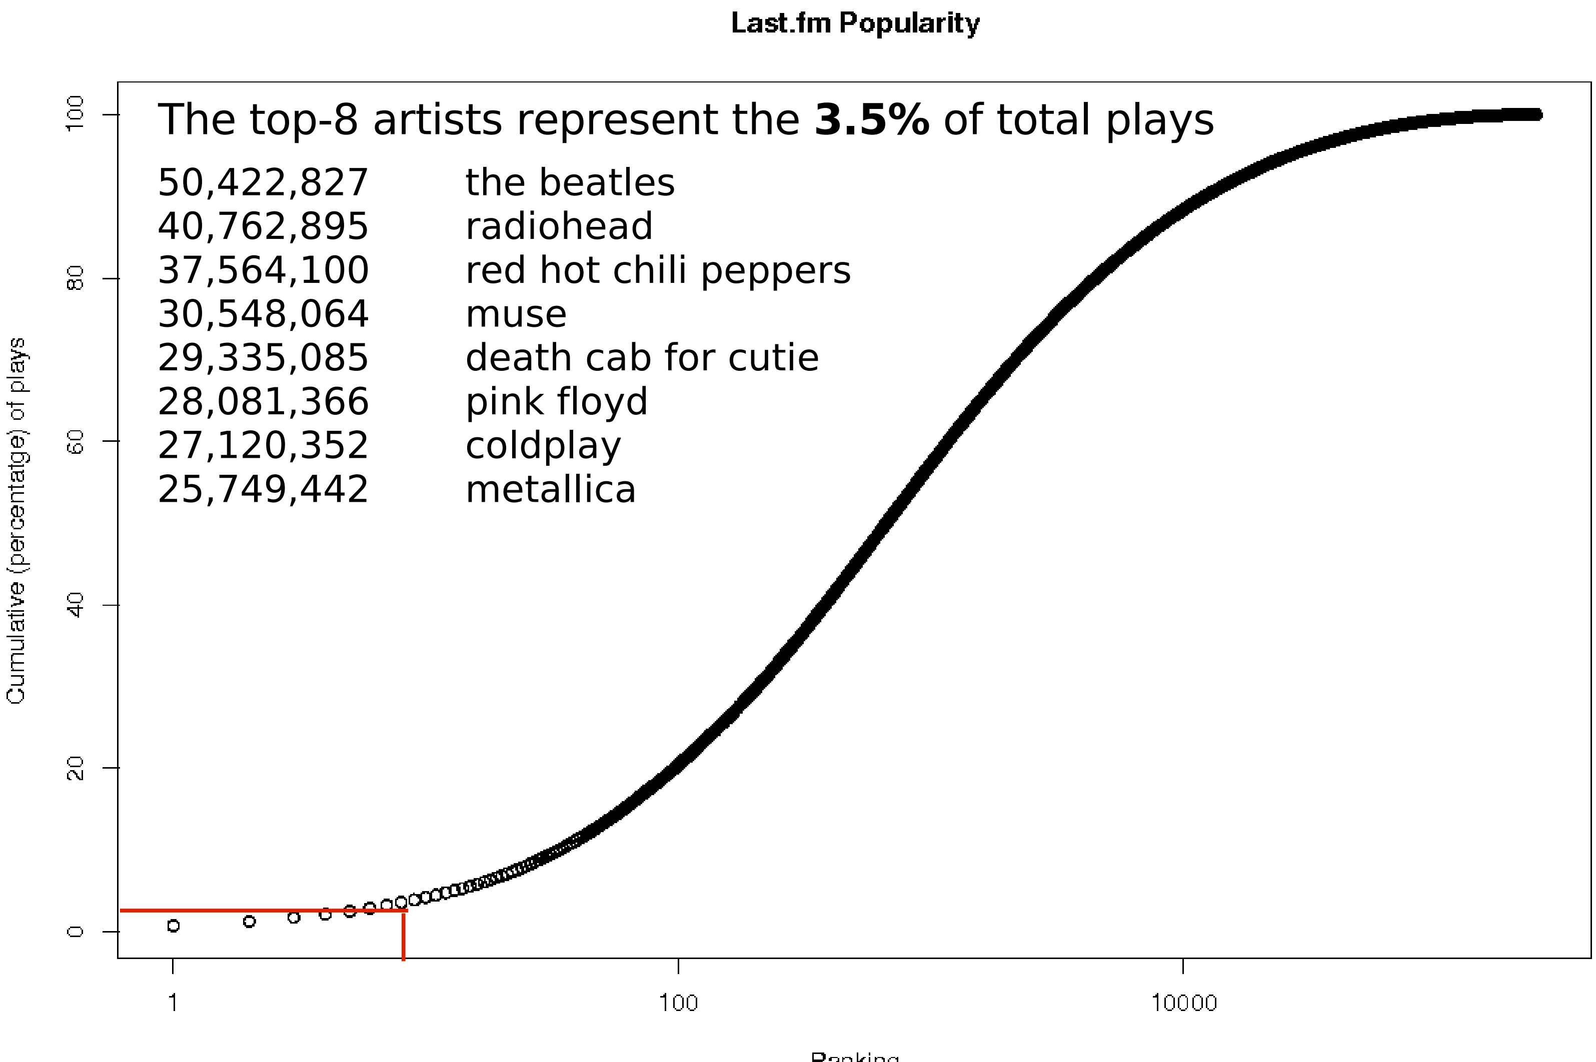

evaluation:: novelty & relevance

- **Last.fm** long-tail model

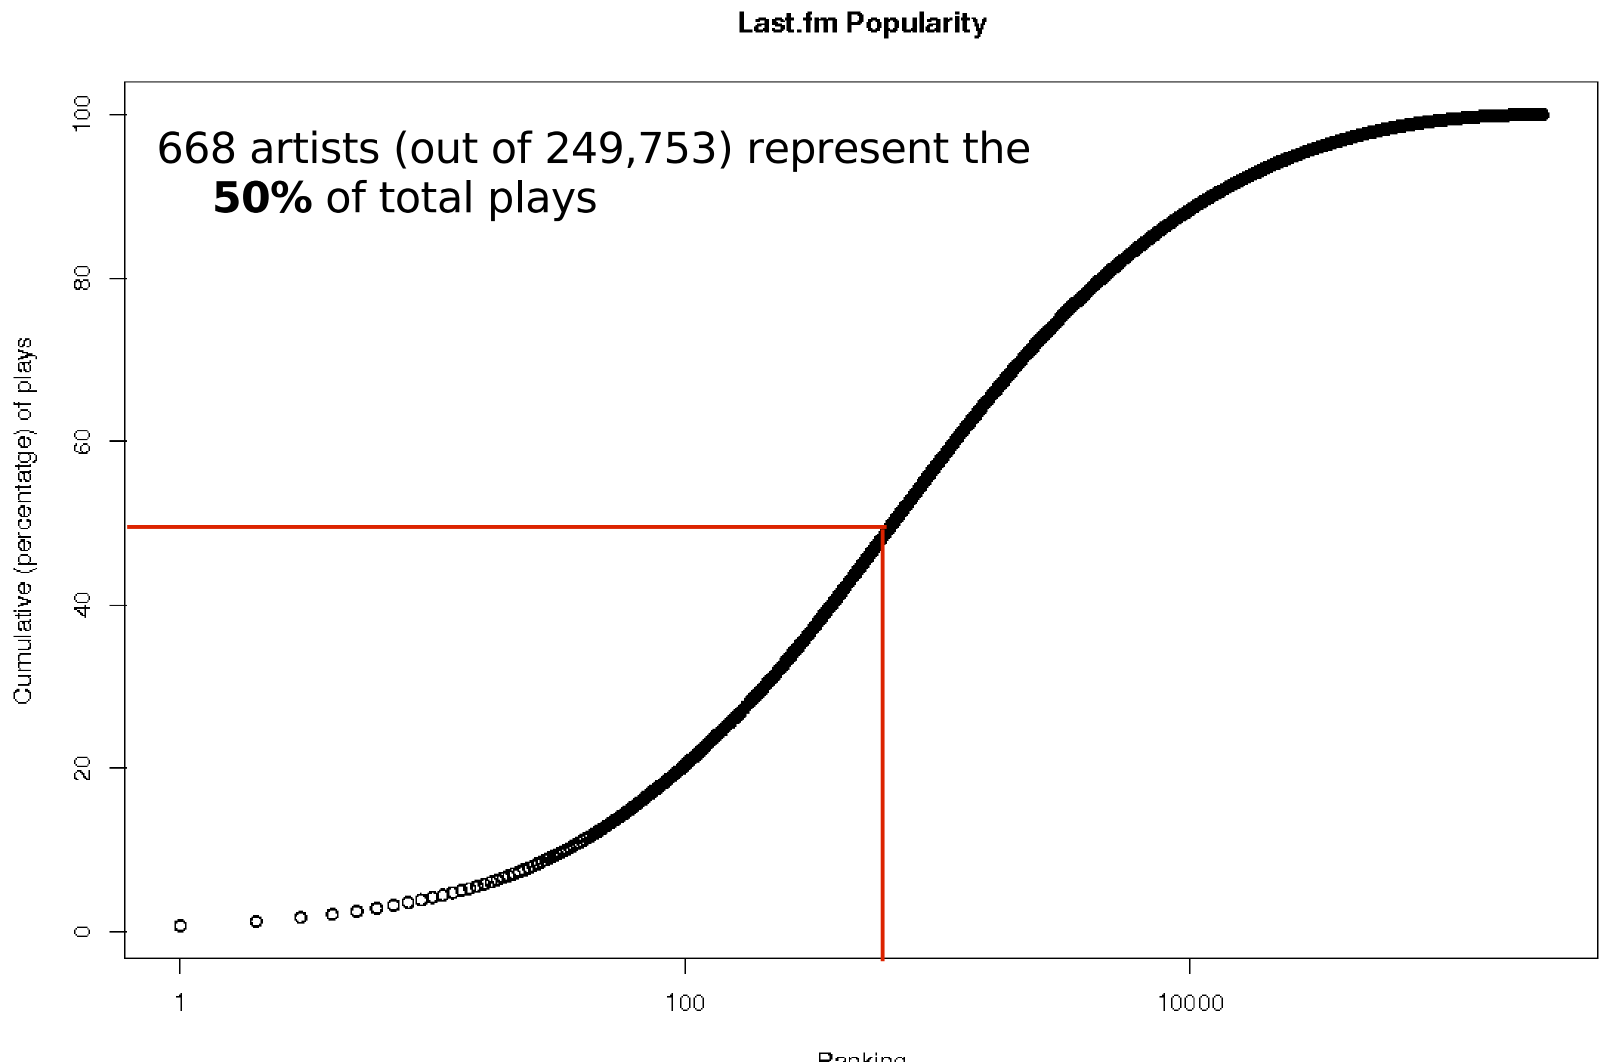

evaluation:: novelty & relevance

- **Last.fm** long-tail **& artist similarity**

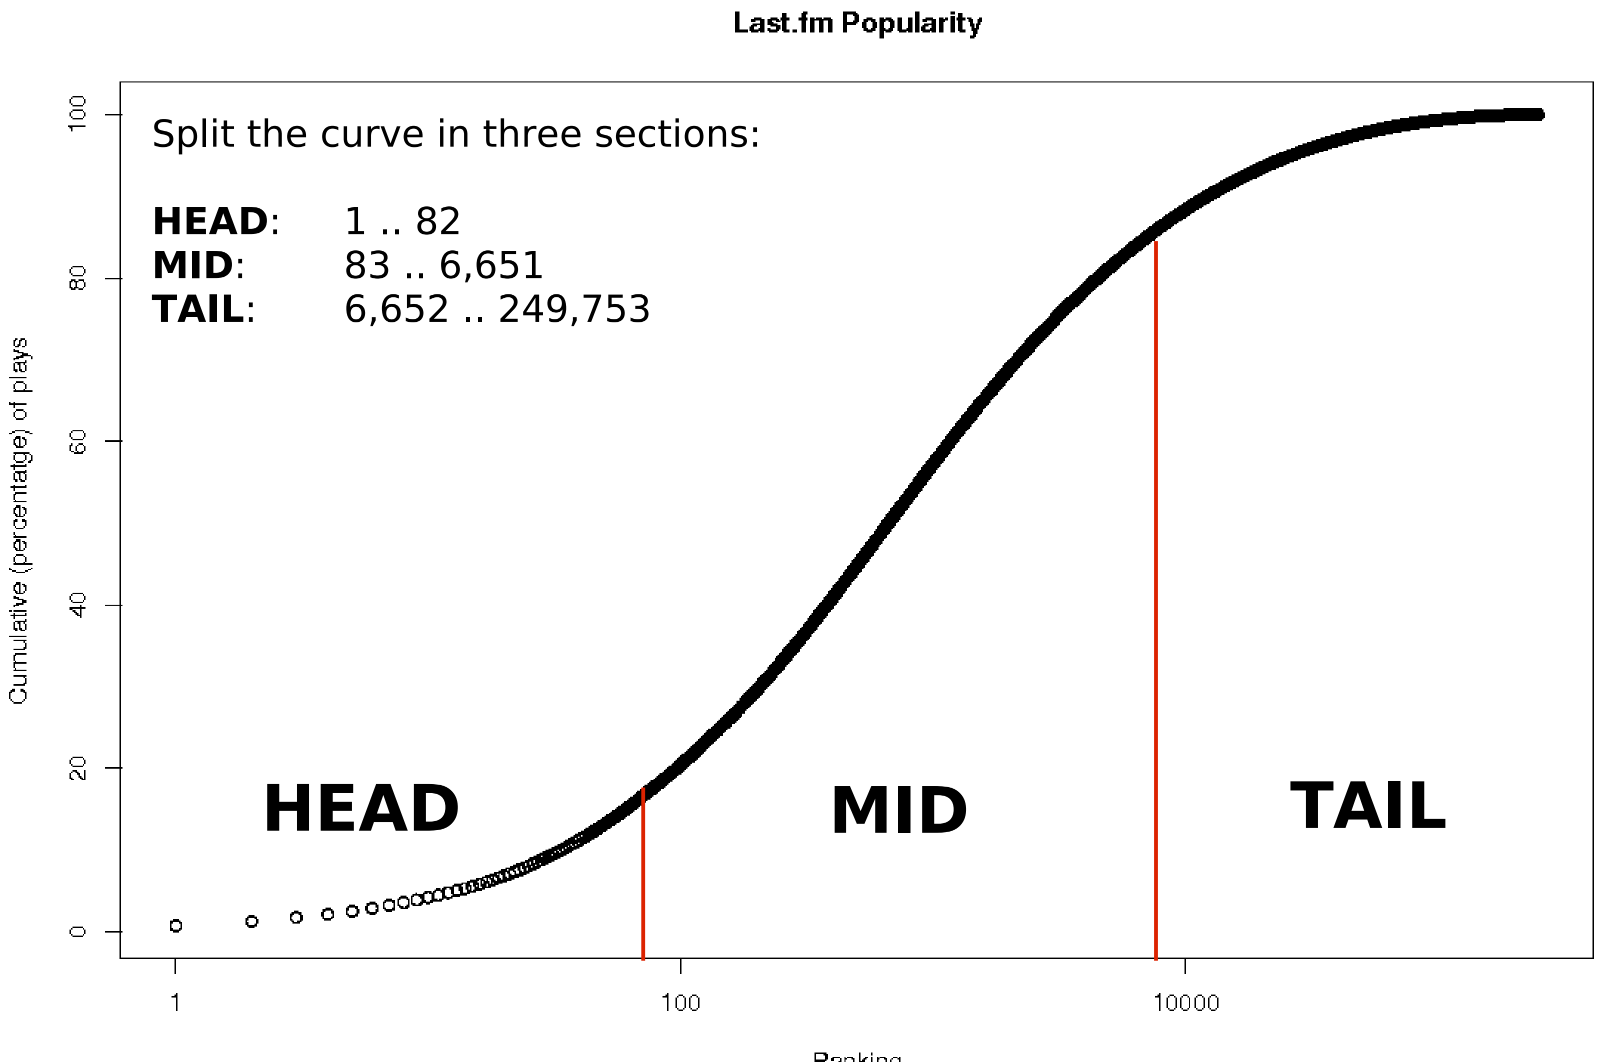

evaluation:: novelty & relevance

- **Last.fm** long-tail & artist similarity

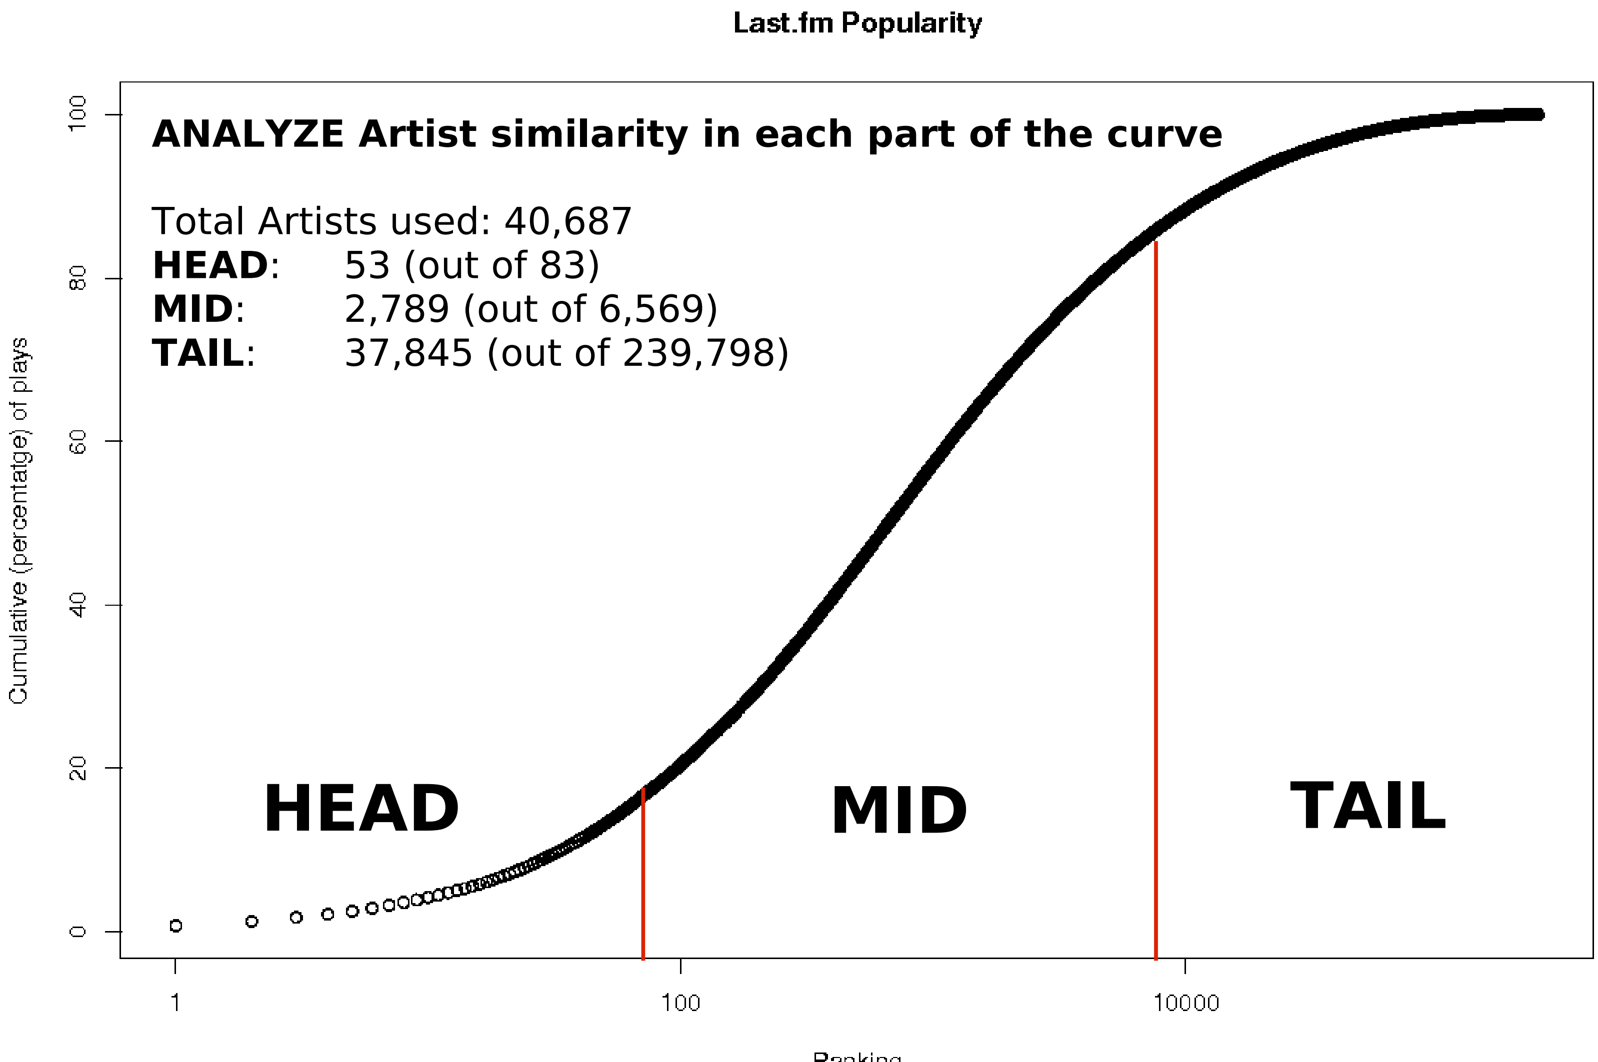

evaluation:: novelty & relevance

- **Last.fm** long-tail & artist similarity
  - The Beatles example

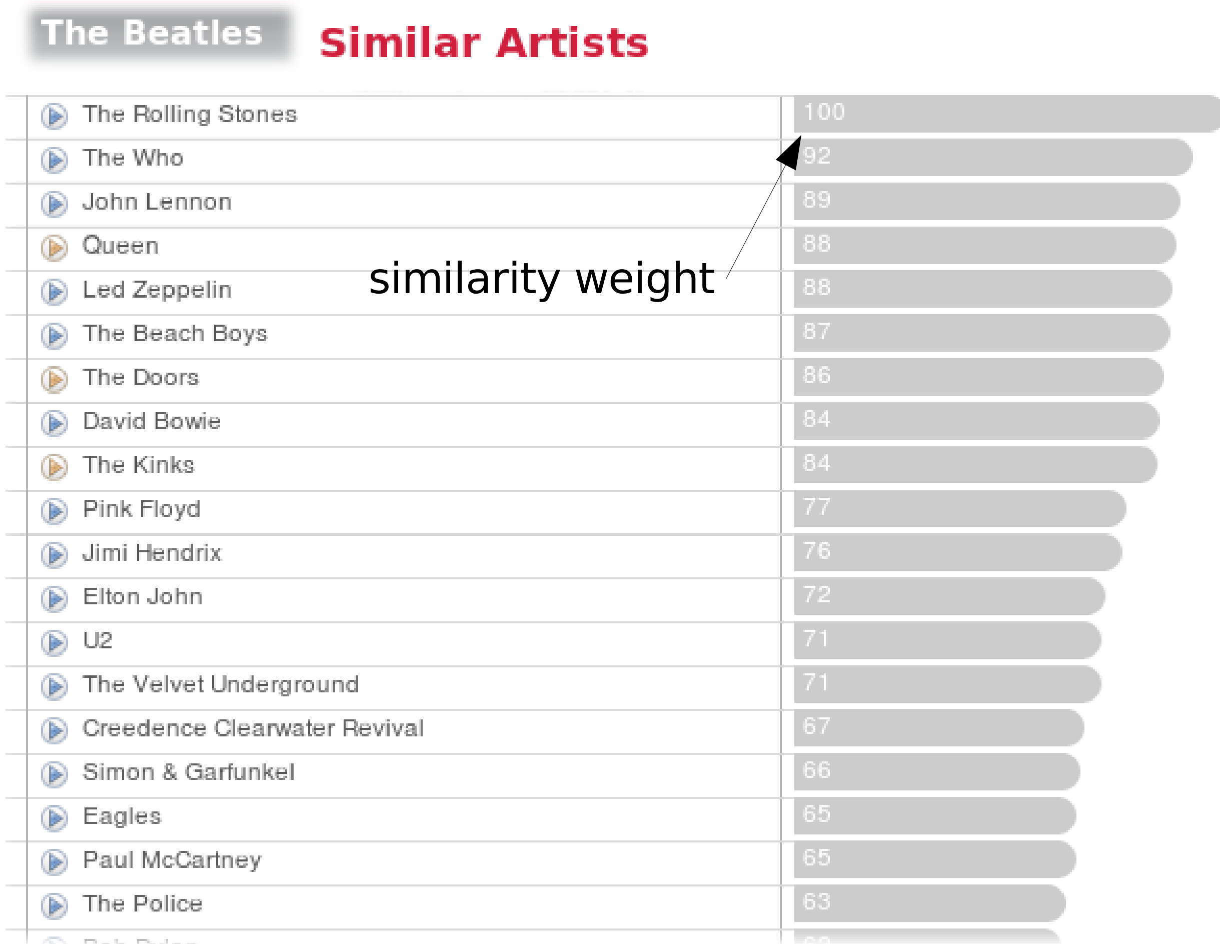

evaluation:: novelty & relevance

- **Last.fm** long-tail & artist similarity
  - The Beatles example

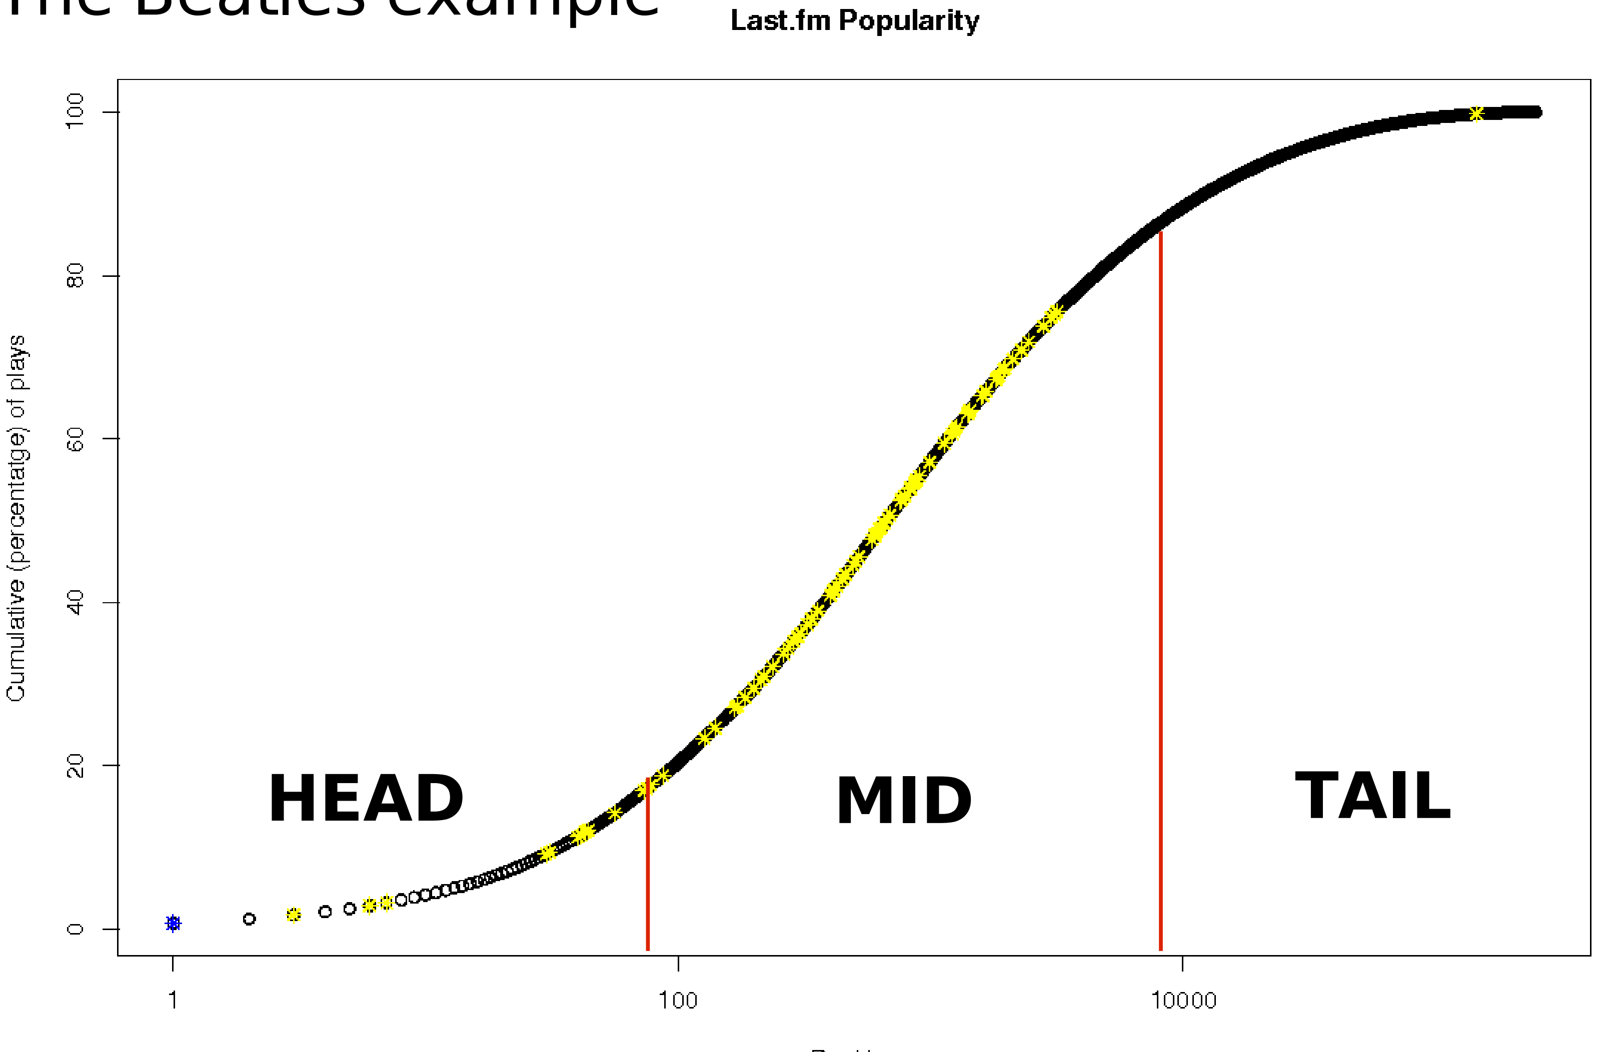

evaluation:: novelty & relevance

- **Last.fm** long-tail & artist similarity
  - Artists in the **HEAD**

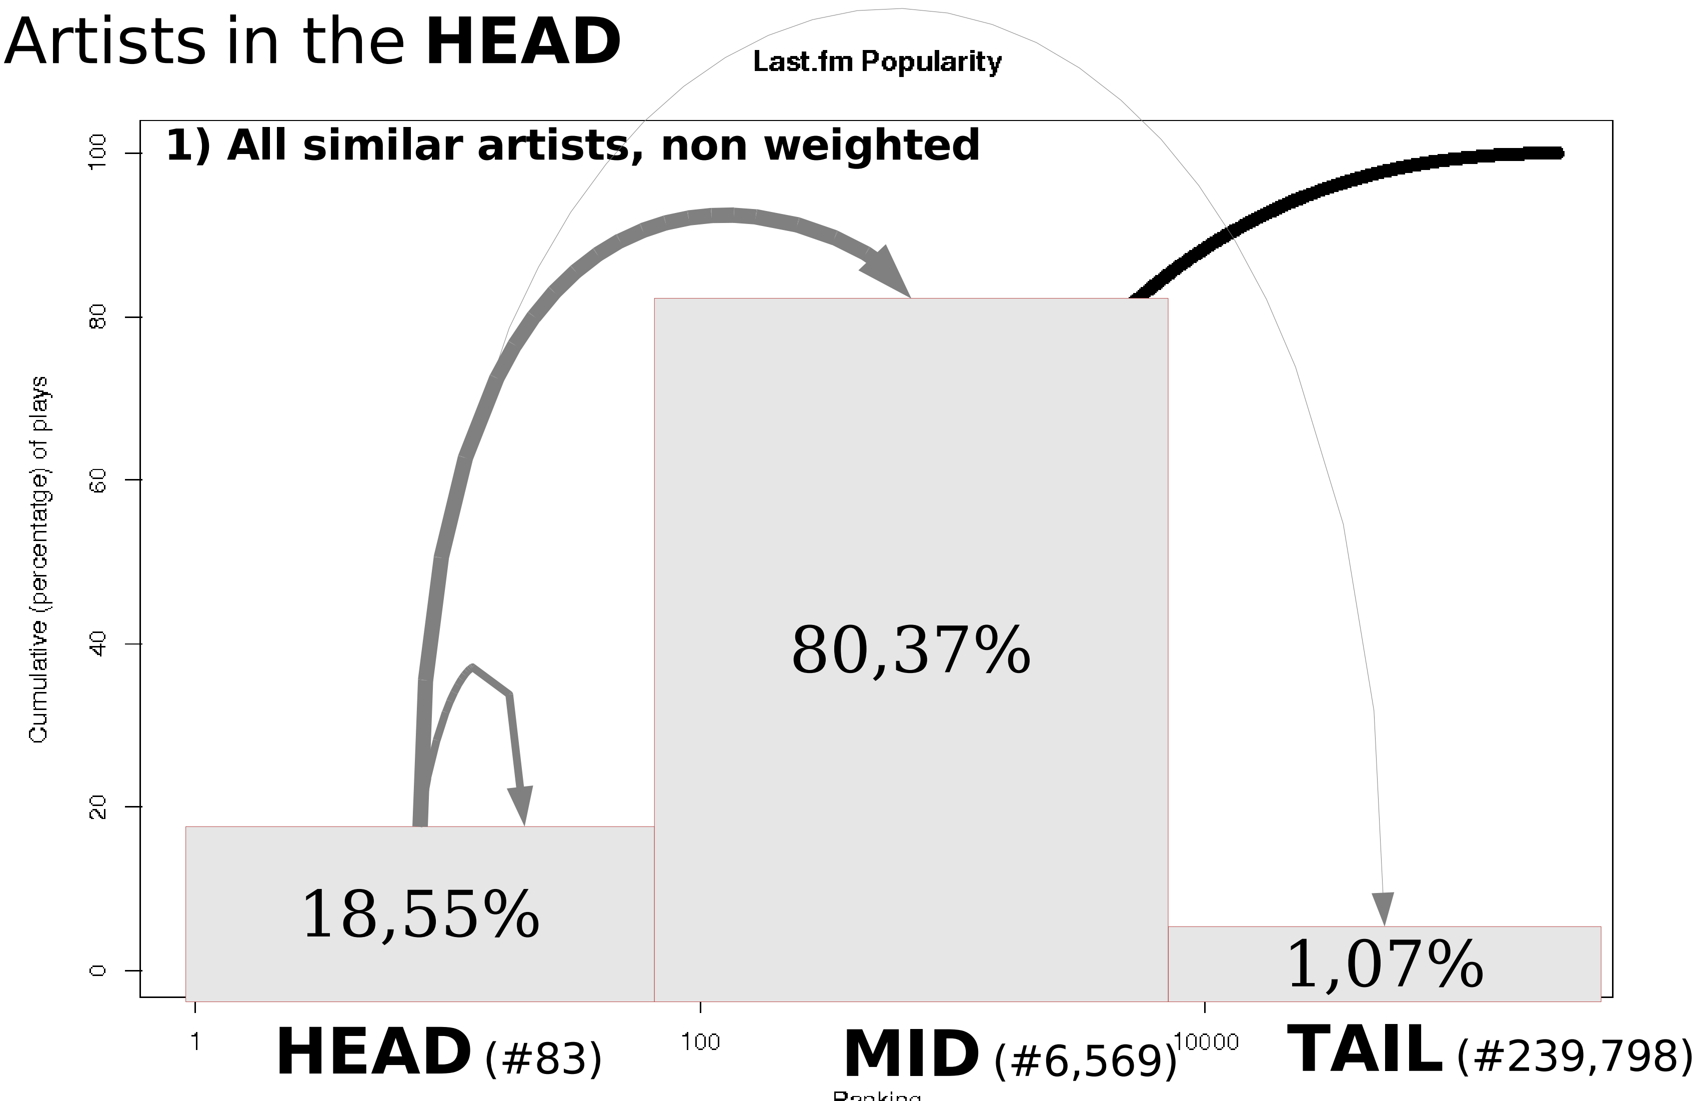

evaluation:: novelty & relevance

- **Last.fm** long-tail & artist similarity
  - Artists in the **HEAD**

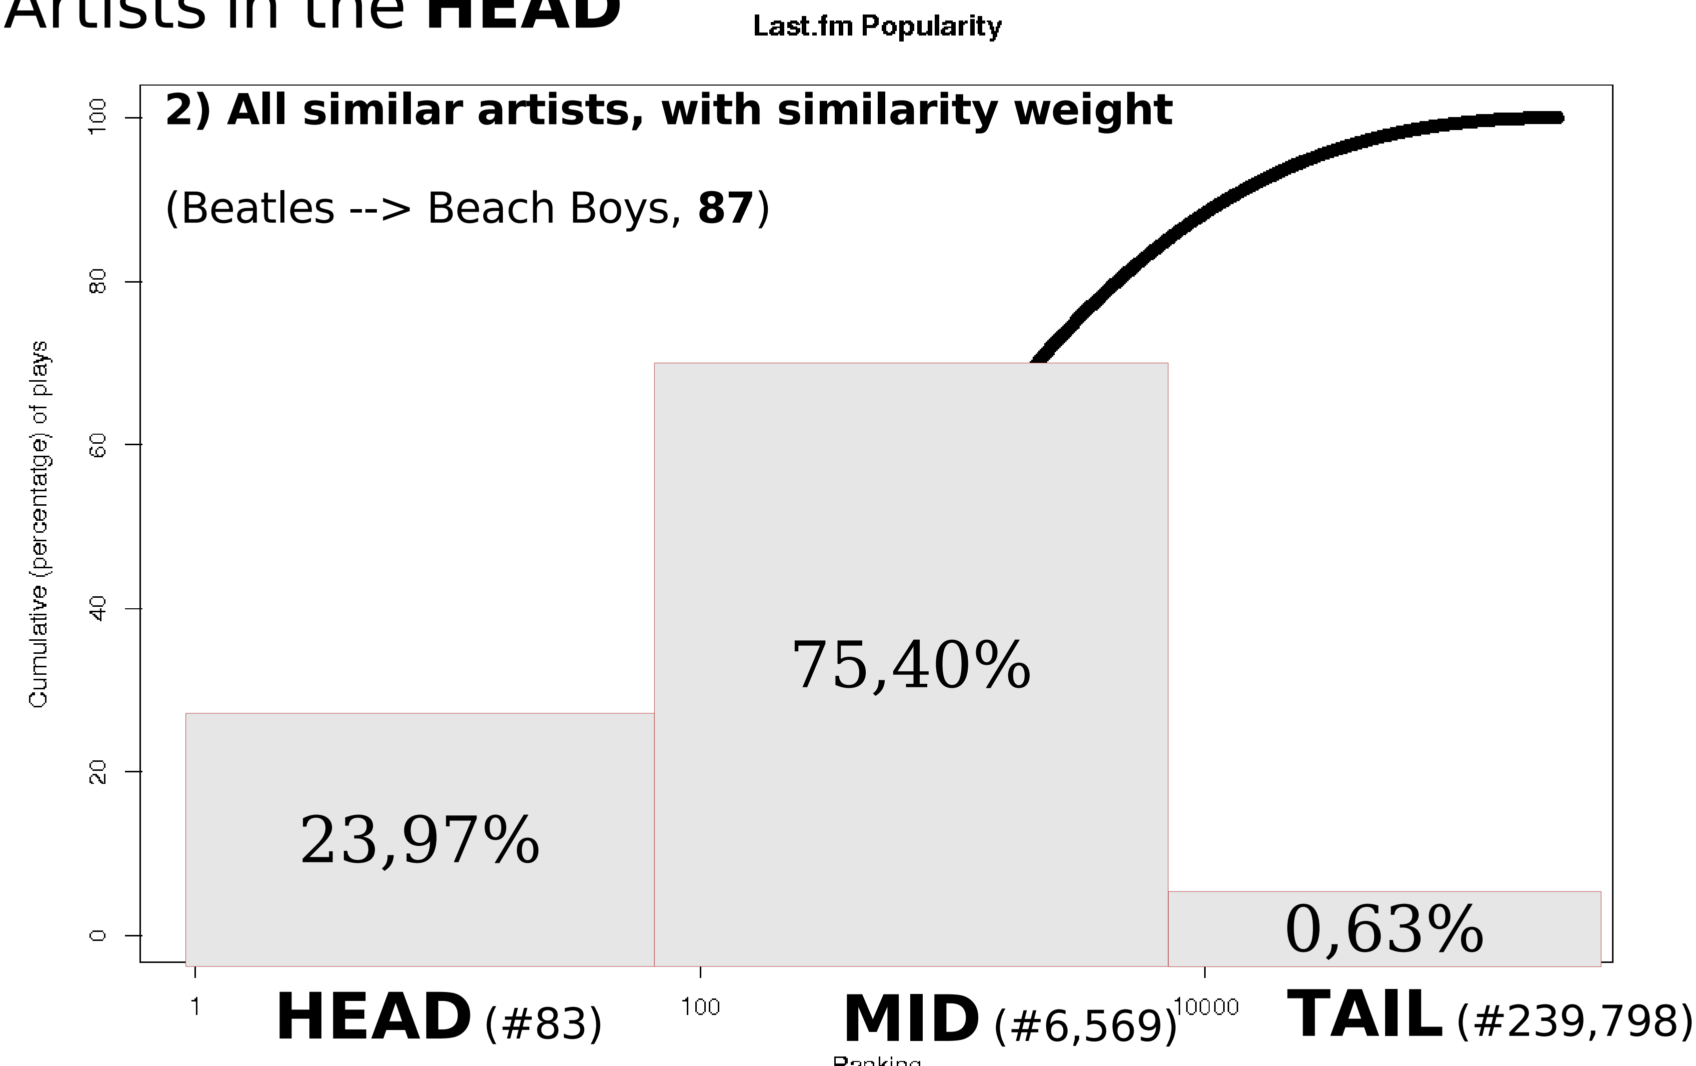

evaluation:: novelty & relevance

- **Last.fm** long-tail & artist similarity
  - ❖ Artists in the **HEAD**

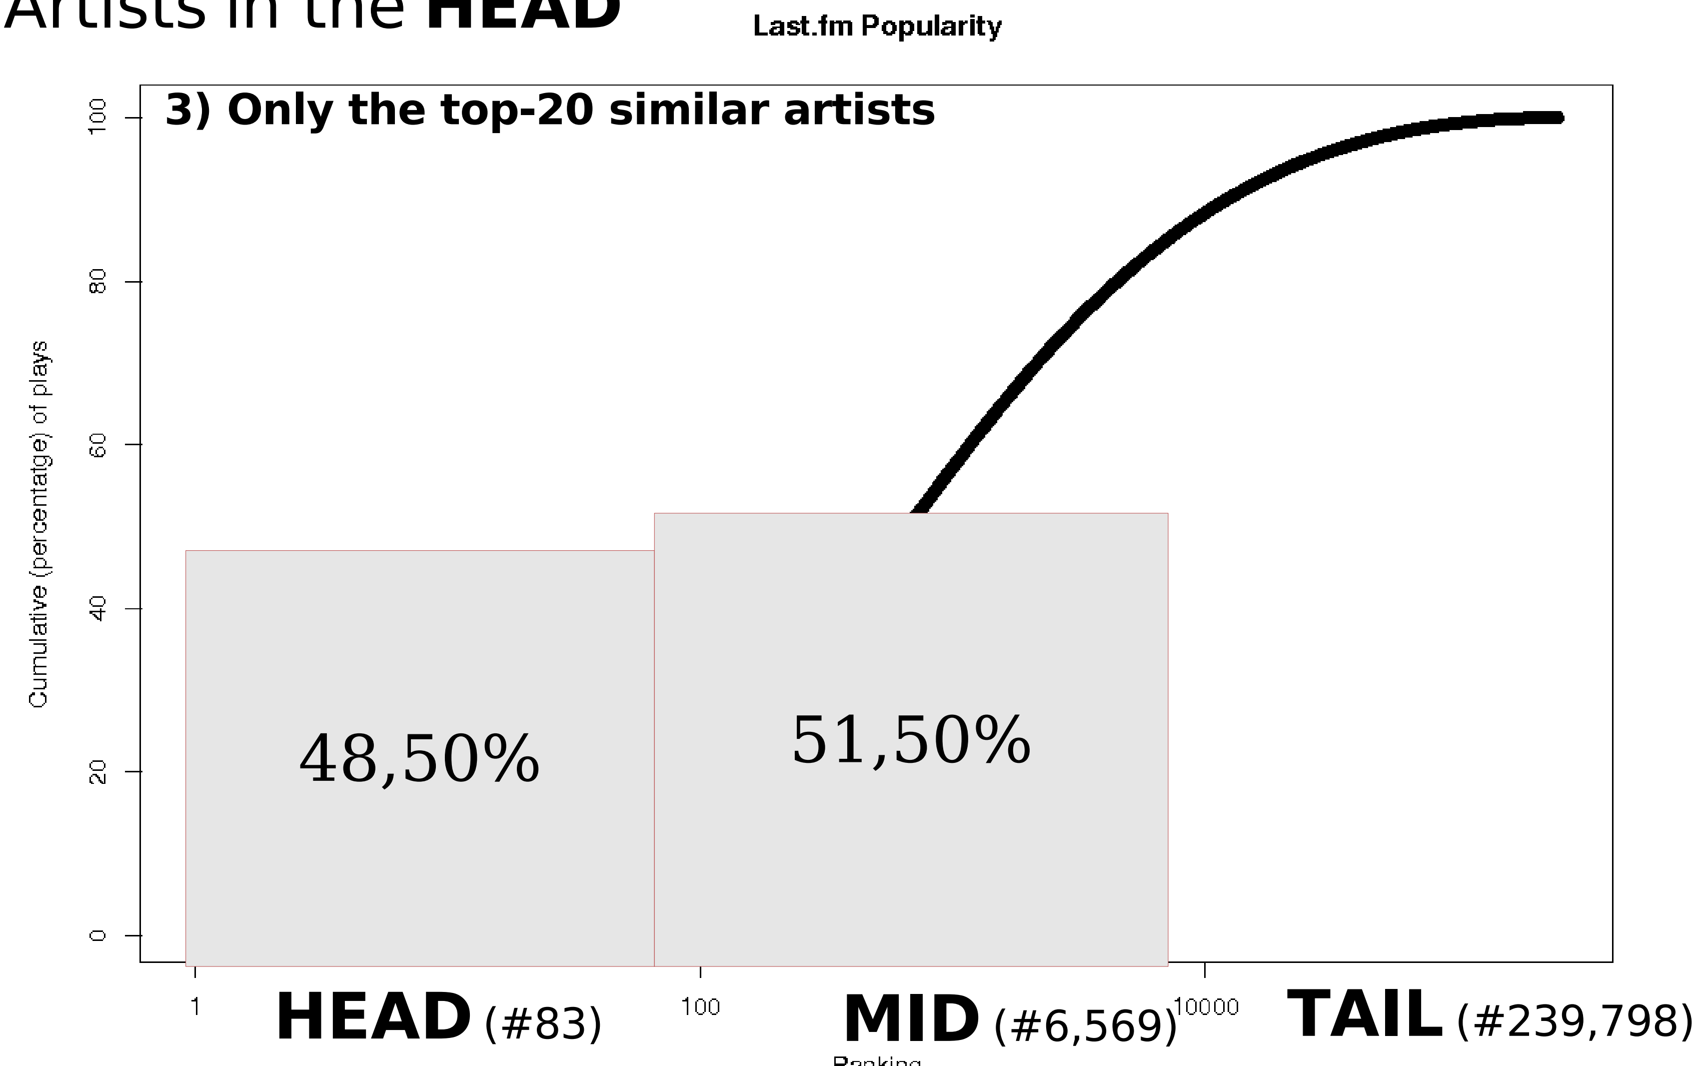

evaluation:: novelty & relevance

- **Last.fm** long-tail & artist similarity
  - ❖ Artists in the **MID**

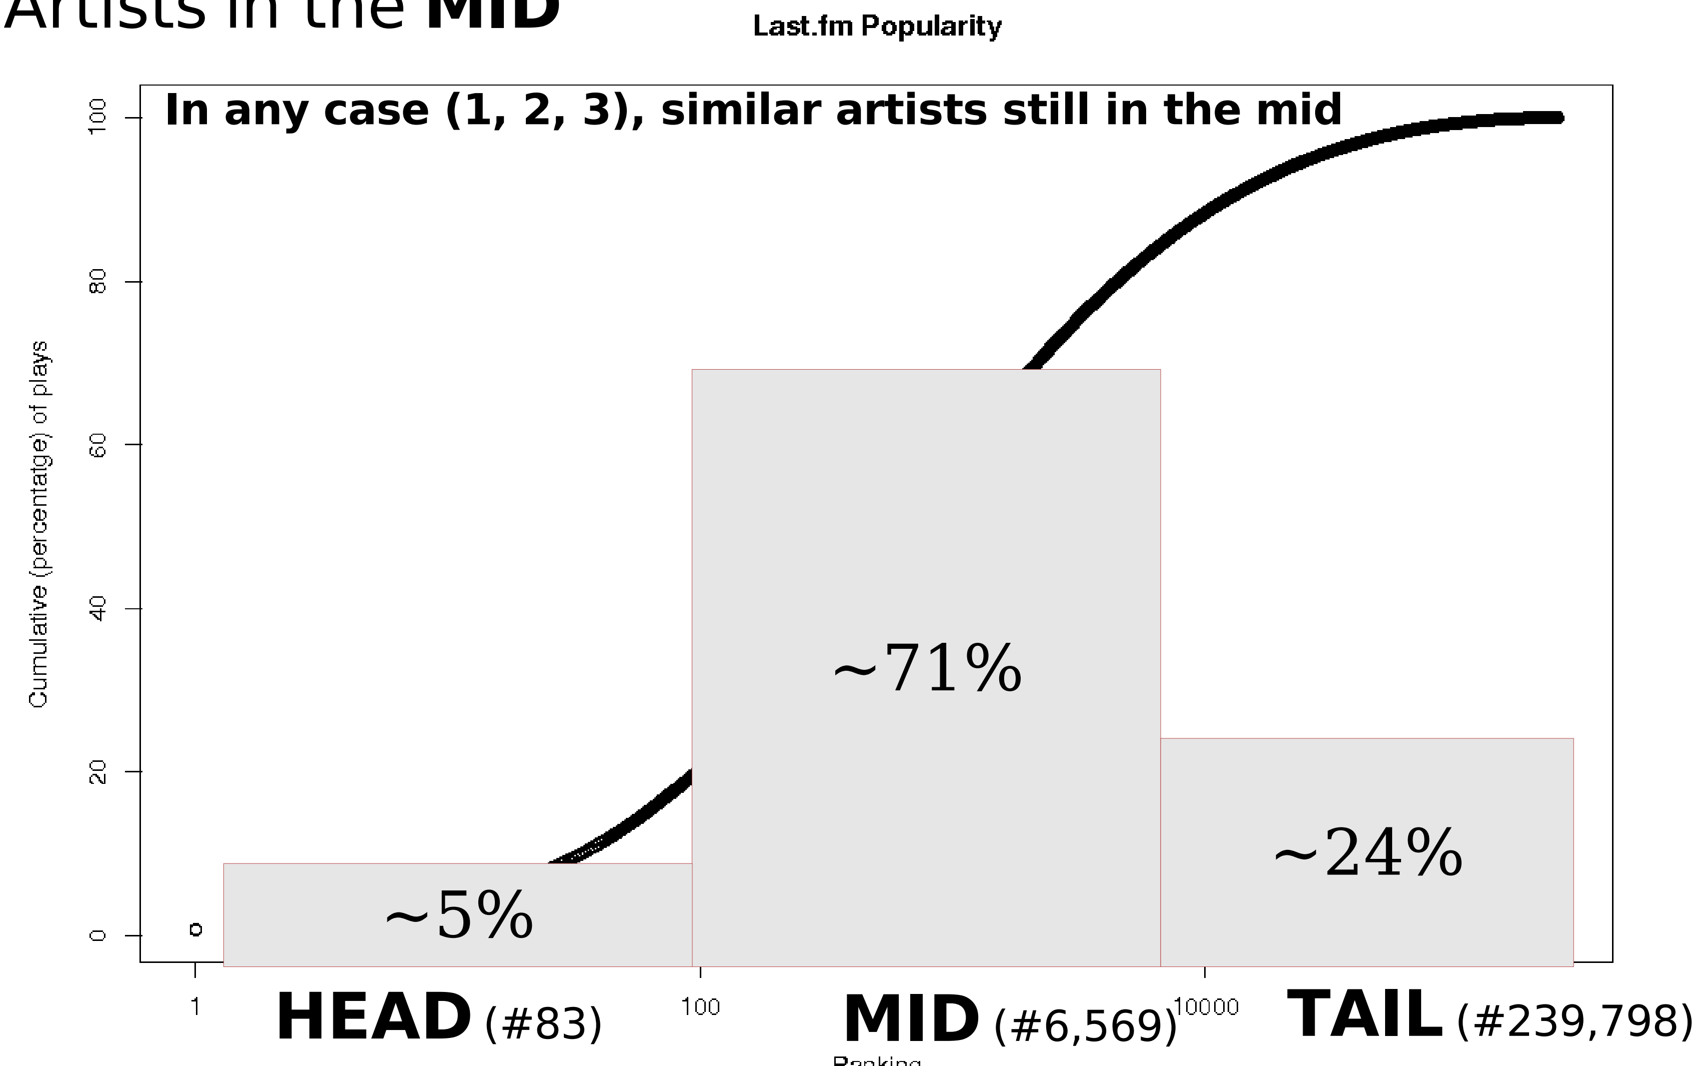

evaluation:: novelty & relevance

- **Last.fm** long-tail & artist similarity
  - Artists in the **TAIL**

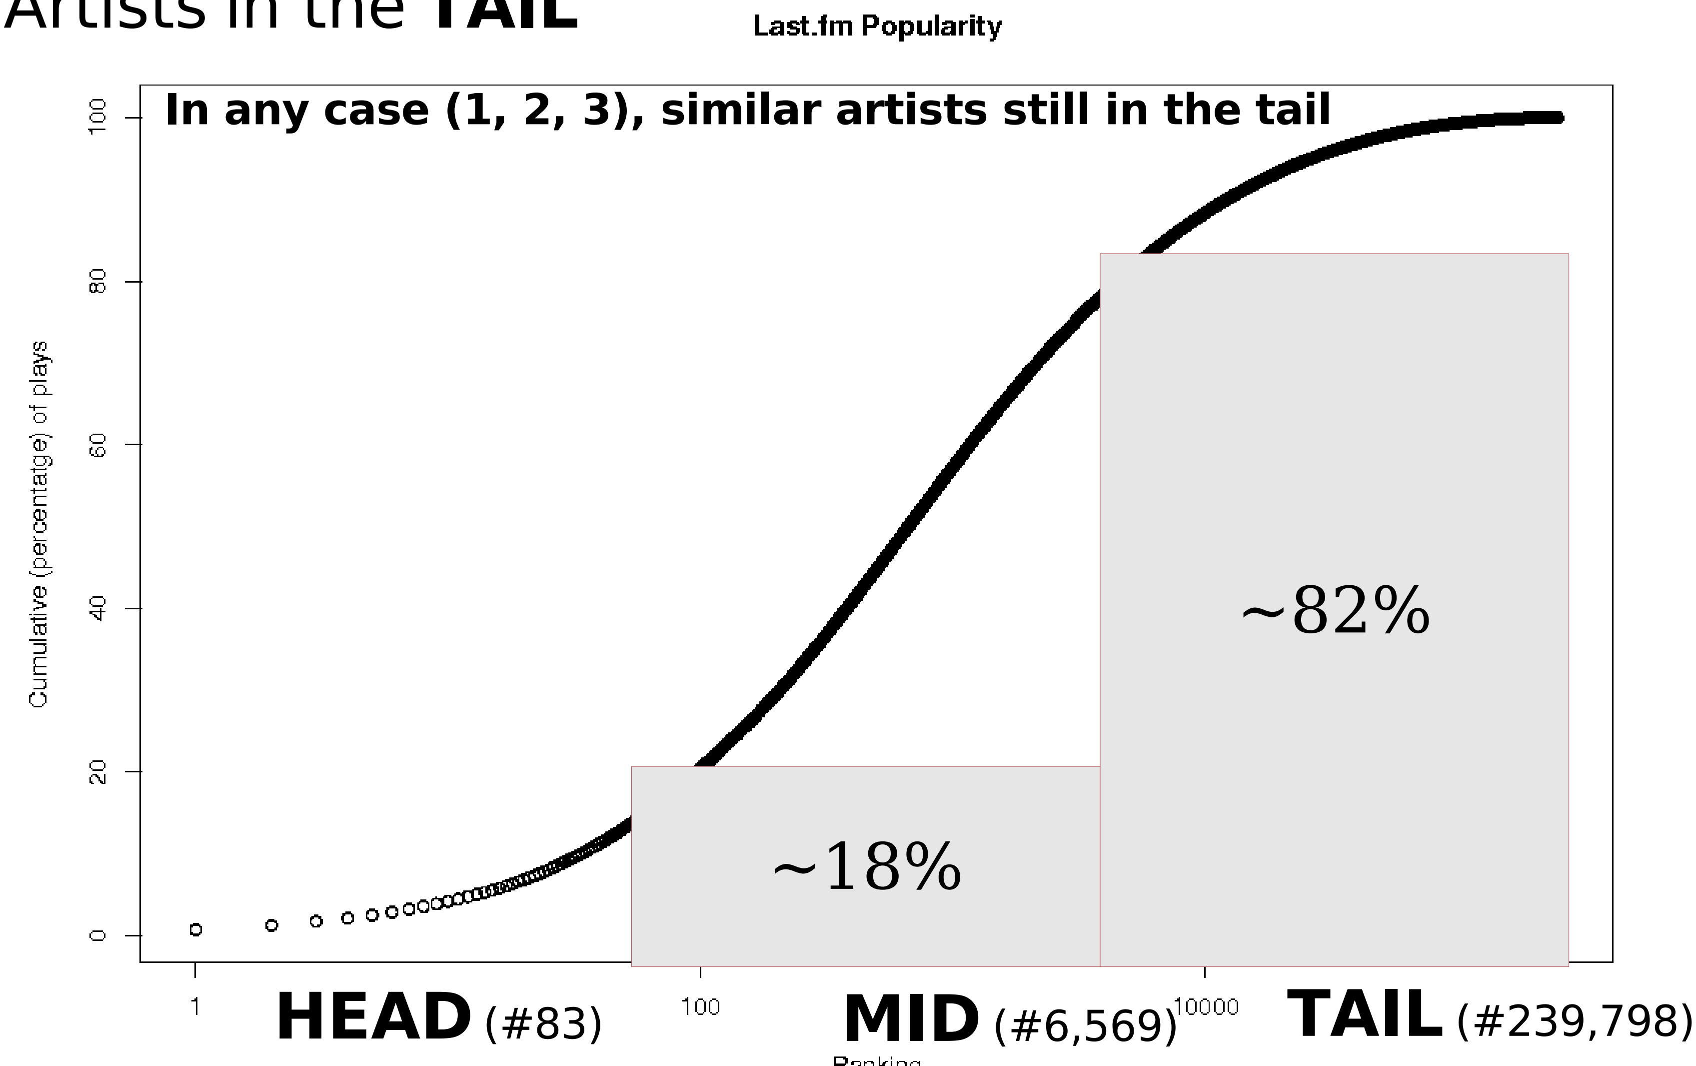

evaluation:: novelty & relevance

- **Last.fm** long-tail & artist similarity
  - implications on the navigation & discovery
    - From Bruce Springsteen to Mike Shupp, or the other way around?
  - implications on the recommendation algorithm itself
    - How to deal with the long-tail?

evaluation:: complex network analysis

- Complex network analysis
  - get hints about the *inherent* structure of the artist similarities
  - characterize the network
    - Small world effect? (“6 degrees of Black Sabbath”)
- Case Study 1: Last.fm artist similarity network
  - directed graph
  - 249,753 nodes
  - 3,846,262 edges
  - weighted graph (beatles --> the doors , *weight=86*)

evaluation:: complex network analysis

- **Last.fm** artist similarity network
  - Avg. degree, <k> = 15.4
  - Diameter = 20
  - Small world effect
    - Avg. shortest path, d = 6,24 (dr = 4,6)
    - Clustering Coefficient, C = 0.23 (Cr = 0,0053)

evaluation:: complex network analysis

- **Last.fm** artist similarity network
  - ❖ Cumulative indegree P(K>k)
    - **no** Power law distribution (log-log) => is not a scale-free network (“*highly connected hubs*”)

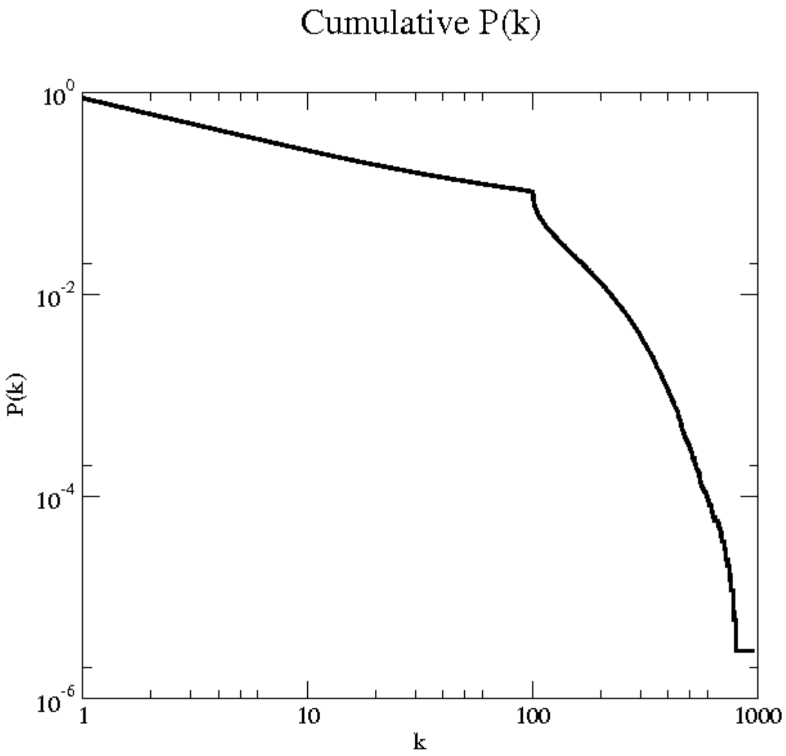

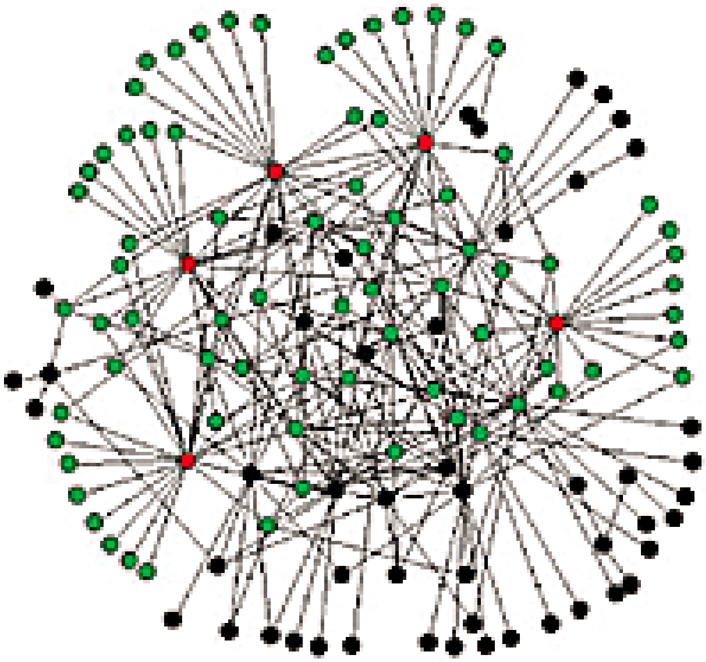

evaluation:: complex network analysis

- **Last.fm** artist similarity network
  - ❖ Cumulative indegree P(K>k)
    - follows an exponential decay (linear-log)

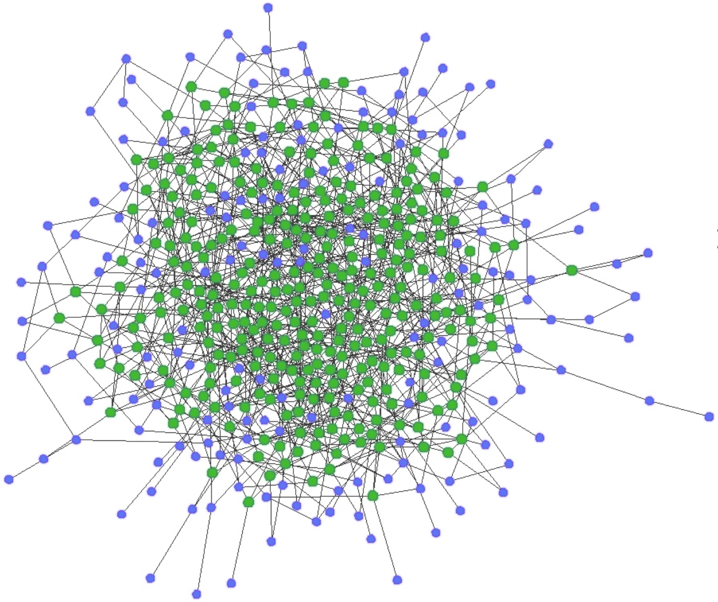

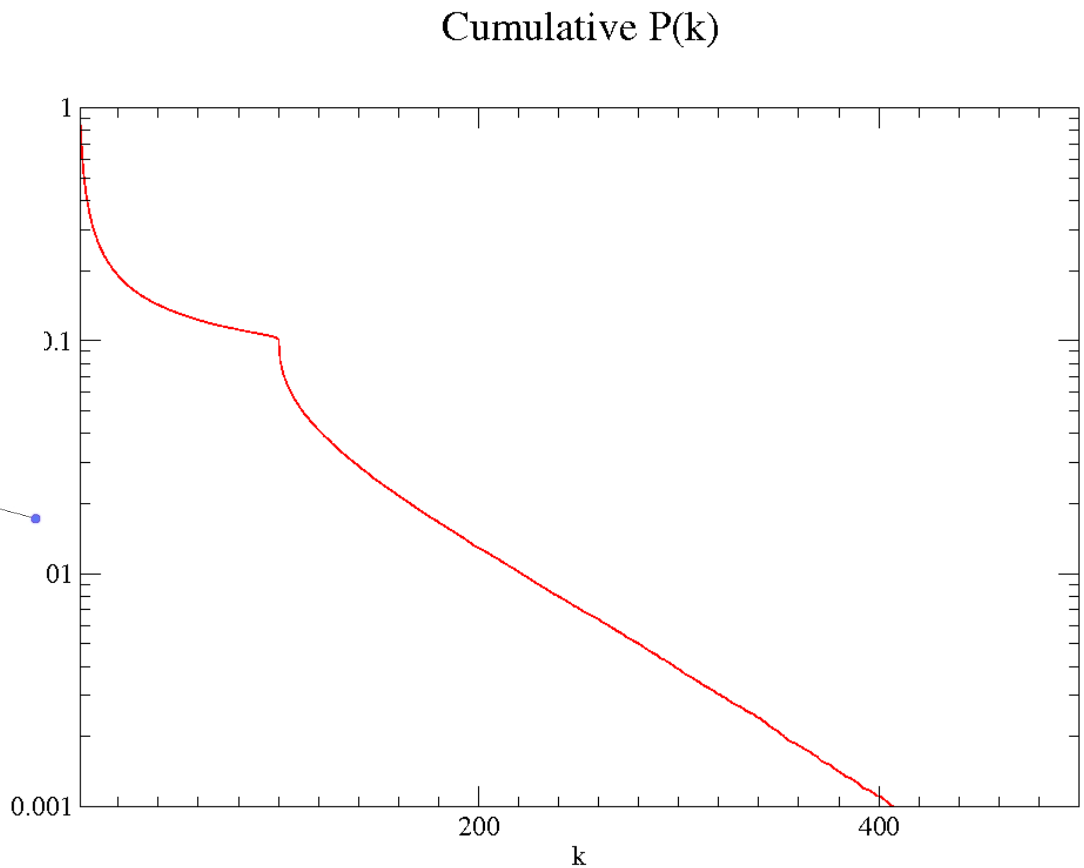

evaluation:: complex network analysis

- **Last.fm** long-tail & artist network
  - ❖ Top-100 artists incoming links

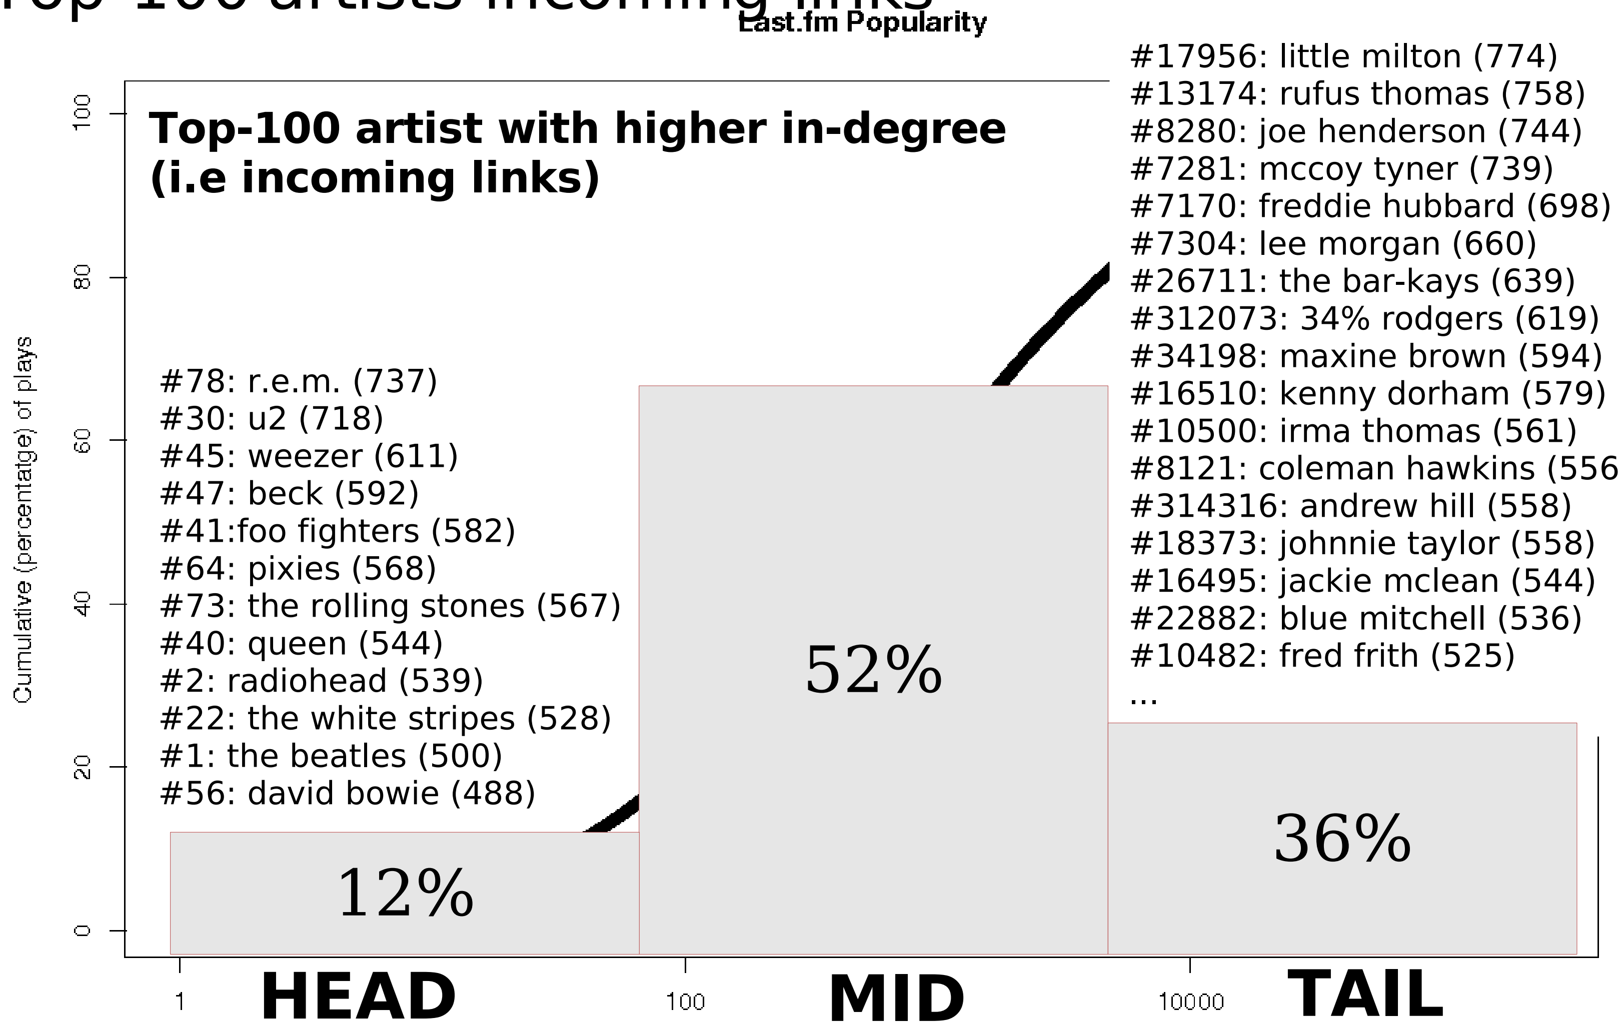

evaluation:: complex network analysis

- **Last.fm** artist similarity network
  - ❖ assortative mixing
    - degree correlation between adjacent nodes

in_degree(Robert Palmer)=430
=>
avg(in_degree(sim(Robert Palmer)))=342

in_degree(Ed Alton)=11
=>
avg(in_degree(sim(Ed Alton)))=18

Assortative mixing in Last.fm :: $K_{nn,in}(k_{in})$

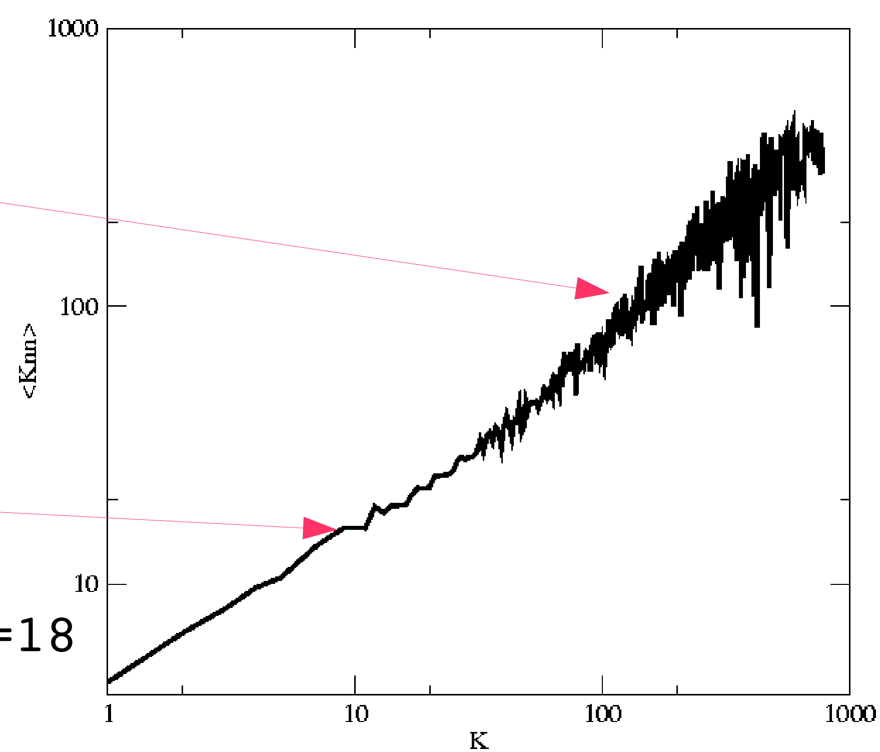

evaluation:: novelty & relevance

- **Last.fm** summary
  - Not (highly) influenced by the popularity effect
  - But...not exploiting the long-tail for discovery!

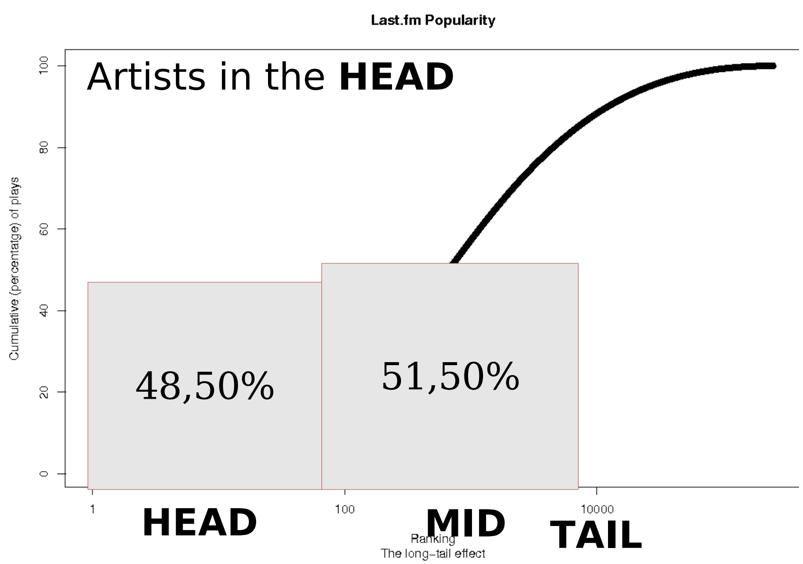

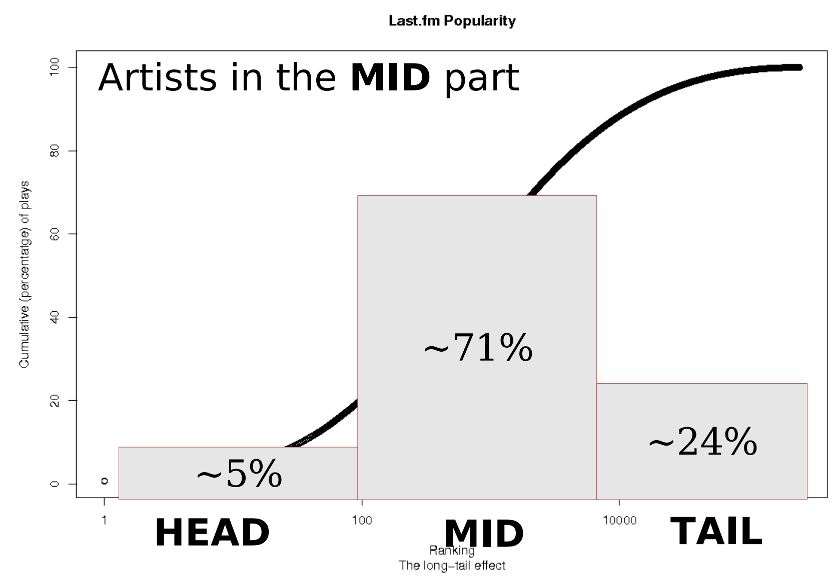

evaluation:: novelty & relevance

- **Last.fm** summary: exploit the long-tail
  - R.E.M "related" artists...

List #1

- U2
- Radiohead
- Coldplay
- Red Hot Chili Peppers
- The Smashing Pumpkins
- The White Stripes
- Foo Fighters
- Weezer
- Counting Crows
- Oasis
- Pixies
- ...

List #2

- Jesus Jones
- Primitive Radio Gods
- Love Spit Love
- Sprung Monkey
- Jeff Ament
- Flickerstic
- Lustre
- Loud Lucy
- The Primitives
- Mike Watt
- Weed
- ...

evaluation:: novelty & relevance

- **Last.fm** summary: exploit the long-tail
  - R.E.M “related” artists...
    - promoting the artists in the Long Tail
      - Jesus Jones, Sandy Rogers, (Primus is out!)

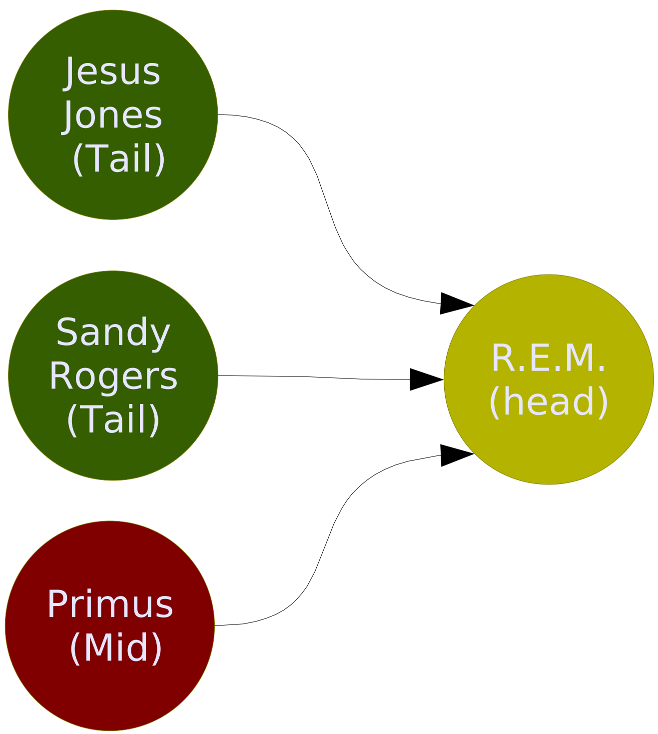

evaluation:: novelty & relevance

- **Last.fm** summary: exploit the long-tail
  - ❖ R.E.M "related" artists...
    - promoting the artists in the Long Tail
      - ❖ Jesus Jones, Sandy Rogers, (Primus is out!)
    - ...that are related, too, with R.E.M similar artists:
      - ❖ Jesus Jones

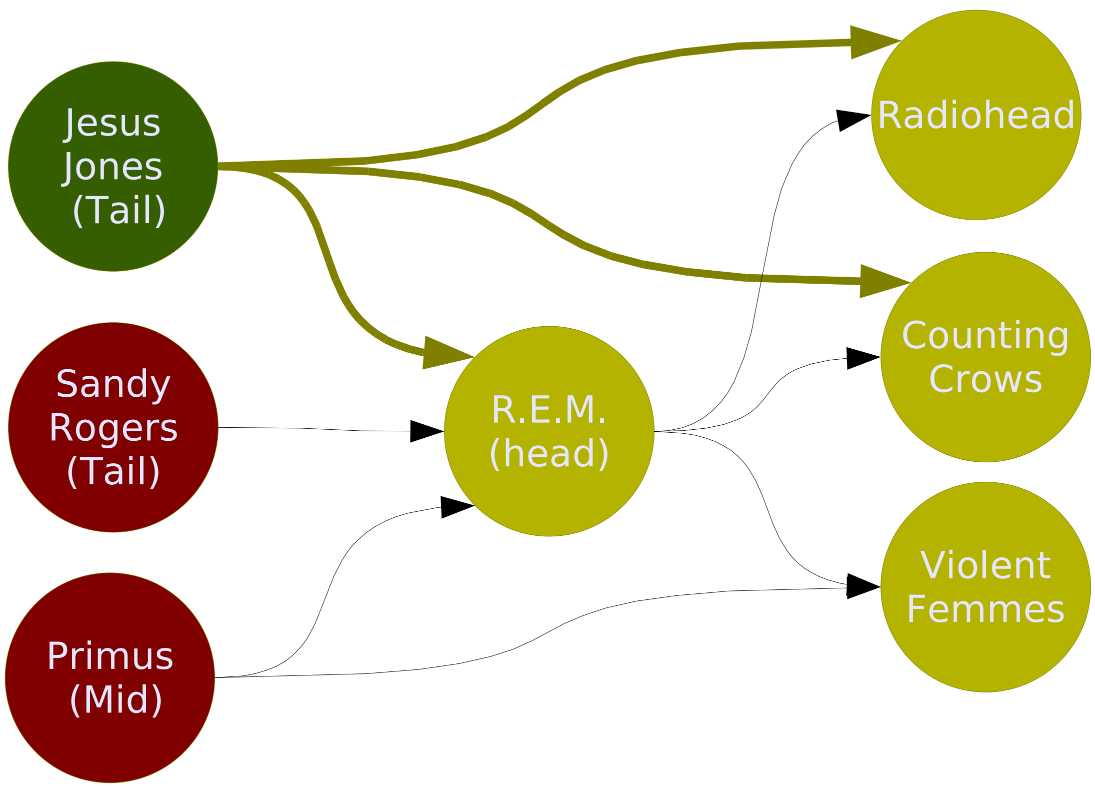

evaluation:: complex network analysis

- Case Study 2: CF vs. Expert recommenders [Cano, 2006]
  - Networks
    - Amazon and MSN: Collaborative filtering
    - AMG: Expert
    - Launch-Yahoo!: ???

	type	n	m	$\langle k \rangle$	C	C_r	d	d_r	γ_{in}	γ_{out}
MSN	directed	51,616	279,240	5.5	0.54	$1.0 \cdot 10^{-4}$	7.7	6.4	2.4 ± 0.01	-
Amazon	directed	23,566	158,866	13.4	0.14	$5.7 \cdot 10^{-4}$	4.2	3.9	2.3 ± 0.02	2.4 ± 0.04
AMG	directed	29,206	146,882	8.15	0.20	$2.8 \cdot 10^{-4}$	6.2	4.9	-	-
Yahoo	directed	16,302	511,539	62.8	0.38	$3.8 \cdot 10^{-3}$	2.7	2.3	-	-

  - All the networks present the Small World effect
    - low avg. shortest path
    - high clustering coefficient

evaluation:: complex network analysis

- Collaborative filtering networks
  - ❖ scale-free networks (power law distribution, log-log)

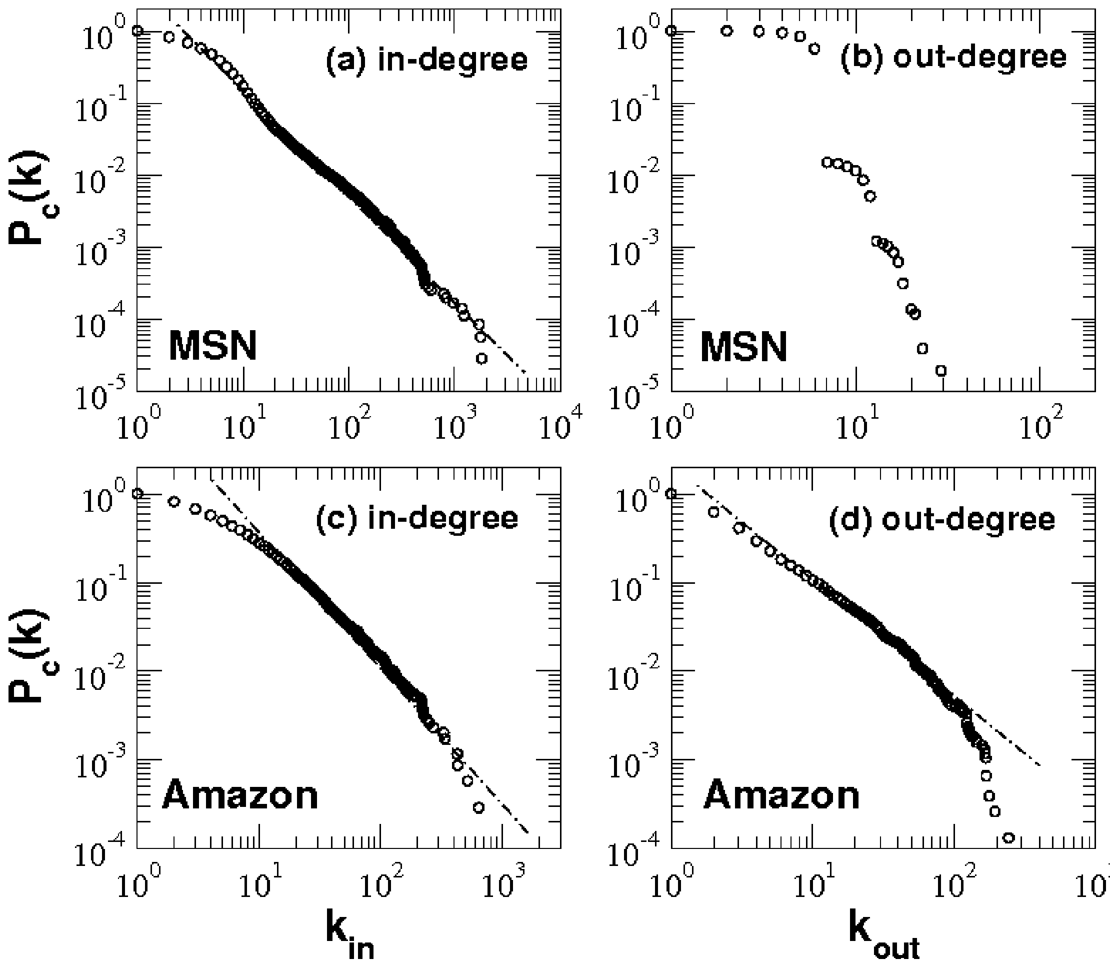

evaluation:: complex network analysis

- Expert network (AMG) ...and Yahoo!
  - ❖ exponential decay (linear-log)

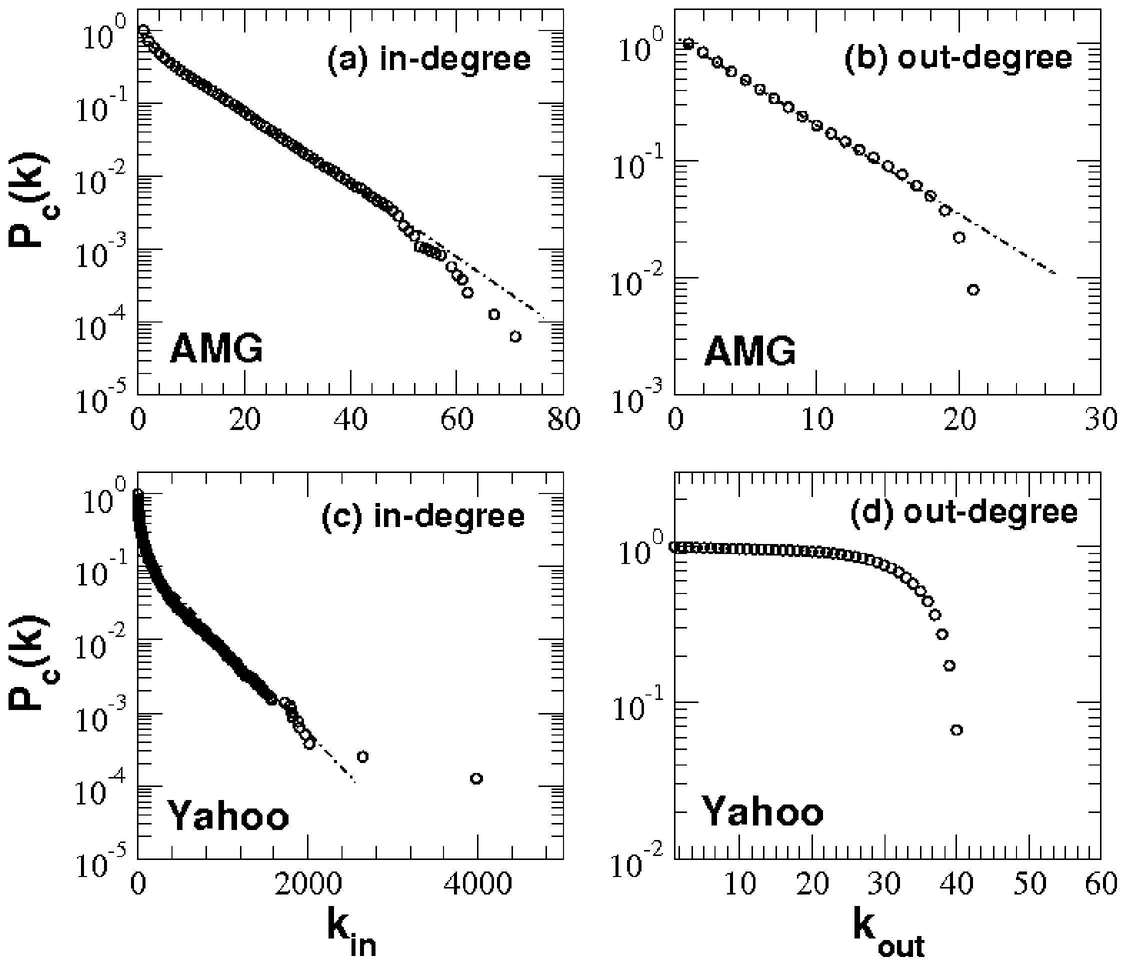

evaluation:: complex network analysis

- Summary
  - network structure clearly affects the recommendations, (as well as the navigation, and discovery)
    - `predecessors(R.E.M.)` could give more interesting information than `successors(R.E.M.)` !
  - Related work at ISMIR2007
    - Complex networks & clustering of users
      - Amélie Anglade, Marco Tiemann, Fabio Vignoli. '*Virtual Communities for Creating Shared Music Channels*'

evaluation:: informal survey

- Method:
  - Generate list of recommended artist based on seed artist
  - Compare lists to expert generated lists
  - Rank lists via survey
- Seed Artists:
  - The Beatles
  - Miles Davis
  - Emerson Lake and Palmer
  - Deerhoof
  - Arcade Fire
- Caveats

evaluation:: informal survey

- The recommenders
  - Commercial
    - All Music
    - iLike
    - last.fm
    - musicmobs
    - MusicMatch
    - MyStrands
    - Pandora
    - Unnamed Beta
    - Up To 11
  - Research
    - Sun tag-based
    - Sun user-based
  - Expert
    - Professional Critic
      - Brian
      - Chris
      - David
      - Dom
      - Joe

evaluation:: informal survey

- The Musical Turing Test

Which recommendation is from a human, which is from a machine?

Seed Artist: The Beatles

- Bob Dylan
- Beach Boys
- Billy Joel
- Rolling Stones
- Animals
- Aerosmith
- The Doors
- Simon & Garfunkel
- Crosby, Stills Nash & Young
- Paul Simon

- Chuck Berry
- Harry Nilsson
- XTC
- Marshall Crenshaw
- Super Furry Animals
- Badfinger
- The Raspberries
- The Flaming Lips
- Jason Faulkner
- Michael Penn

evaluation:: informal survey

- The Musical Turing Test

Which recommendation is from a human, which is from a machine?

Seed Artist: The Beatles

Machine: Up to 11

- Bob Dylan
- Beach Boys
- Billy Joel
- Rolling Stones
- Animals
- Aerosmith
- The Doors
- Simon & Garfunkel
- Crosby, Stills Nash & Young
- Paul Simon

Human

- Chuck Berry
- Harry Nilsson
- XTC
- Marshall Crenshaw
- Super Furry Animals
- Badfinger
- The Raspberries
- The Flaming Lips
- Jason Faulkner
- Michael Penn

evaluation:: informal survey

- The Musical Turing Test

Which recommendation is from a human, which is from a machine?

Seed Artist: Miles Davis

- John Coltrane
- Thelonious Monk
- Charlie Parker
- Herbie Hancock
- Chet Baker
- Bill Evans
- Charles Mingus
- Lee Morgan
- Sonny Rollins

- John Coltrane
- Ken Vandermark
- Talk Talk
- James Brown
- Ornette Coleman
- Norah Jones
- Dizzy Gillespie
- Duke Ellington
- Steely Dan
- Sea & Cake

evaluation:: informal survey

- The Musical Turing Test

Which recommendation is from a human, which is from a machine?

Seed Artist: Miles Davis

Machine: Sun Tags

- John Coltrane
- Thelonious Monk
- Charlie Parker
- Herbie Hancock
- Chet Baker
- Bill Evans
- Charles Mingus
- Lee Morgan
- Sonny Rollins

Human

- John Coltrane
- Ken Vandermark
- Talk Talk
- James Brown
- Ornette Coleman
- Norah Jones
- Dizzy Gillespie
- Duke Ellington
- Steely Dan
- Sea & Cake

evaluation:: informal survey

- The Musical Turing Test

Which recommendation is from a human, which is from a machine?

Seed Artist: Arcade Fire

- Interpol
- Bloc Party
- Modest Mouse
- The Shins
- Clap Your Hands Say Yeah
- Arctic Monkeys
- Editors
- The Strokes
- The Decemberists
- Kings of Leon

- Echo & the Bunnymen
- the Sound
- Comsat Angels
- The Church
- House of Love
- Stone Roses
- The Smiths
- Gene
- Interpol
- U2

evaluation:: informal survey

- The Musical Turing Test

Which recommendation is from a human, which is from a machine?

Seed Artist: Arcade Fire

Machine: last.fm

- Interpol
- Bloc Party
- Modest Mouse
- The Shins
- Clap Your Hands Say Yeah
- Arctic Monkeys
- Editors
- The Strokes
- The Decemberists
- Kings of Leon

Human

- Echo & the Bunnymen
- the Sound
- Comsat Angels
- The Church
- House of Love
- Stone Roses
- The Smiths
- Gene
- Interpol
- U2

evaluation:: informal survey

- Rankings

Agreement with Experts

System	Score
Sun tags	0.88
Sun users	0.82
last.fm	0.78
PC Dominique	0.78
PC Chris	0.74
Mystrands	0.70
Musicmatch	0.64
PC Brian	0.64
iLike	0.54
up to 11	0.53
All Music	0.47
PC Joe	0.46
PC David	0.38
MusicMobs	0.30
Pandora*	0.20
Unnamed beta	0.13

Agreement with Machines

System	Score
last.fm	3.20
Sun users	3.11
Musicmatch	2.66
Sun tags	2.60
Mystrands	2.58
iLike	2.49
PC Dom	1.92
PC Brian	1.80
up to 11	1.55
PC Chris	1.14
MusicMobs	0.88
Pandora*	0.88
PC Joe	0.76
All Music	0.69
PC David	0.66
Unnamed beta	0.27

Overall Agreement

System	Score
last.fm	3.98
Sun users	3.94
Sun tags	3.46
Musicmatch	3.30
Mystrands	3.30
iLike	3.02
PC Dom	2.70
PC Brian	2.44
up to 11	2.08
PC Chris	1.88
PC Joe	1.22
MusicMobs	1.18
All Music	1.16
Pandora*	1.08
PC David	1.04
Unnamed beta	0.39

evaluation:: informal survey

Music Recommendation Survey

Thanks for agreeing to participate in the music recommendation survey The survey is v
simple: you are asked to rate the quality of recommendations that are of the form "If yo
like XXX you might like YYY". The survey will take about 10 minutes per artist to
complete. You don't have to complete the survey for all artists. If you don't know anytl
about a particular seed artist, you can skip that artist.

The Survey

- If you like **The Beatles** you might like ...
- If you like **Miles Davis** you might like ...
- If you like **Emerson Lake and Palmer** you might like ...
- If you like **Deerhoof** you might like ...
- If you like **The Arcade Fire** you might like ...

Your answers will be used to evaluate and compare a set of music recommenders.
Send any questions or comments to Paul.Lamere@sun.com.

Privacy | Terms of Use | Trademarks | blog

200 Responses
>10,000 datapoints

Music Recommendation Survey

Instructions: Indicate how well the given artist answers the question:

If you like **The Beatles** you might like XXX?

Artist	Recommendation Rating				
The Turtles	Excellent:	Good:	Don't Know:	Fair:	Poor:
October Cherries	Excellent:	Good:	Don't Know:	Fair:	Poor:
Chills	Excellent:	Good:	Don't Know:	Fair:	Poor:
Counting Crows	Excellent:	Good:	Don't Know:	Fair:	Poor:
The Zombies	Excellent:	Good:	Don't Know:	Fair:	Poor:
Hoobastank	Excellent:	Good:	Don't Know:	Fair:	Poor:
War	Excellent:	Good:	Don't Know:	Fair:	Poor:
Aerosmith	Excellent:	Good:	Don't Know:	Fair:	Poor:
Dave Clark Five	Excellent:	Good:	Don't Know:	Fair:	Poor:
Michael Penn	Excellent:	Good:	Don't Know:	Fair:	Poor:
Donovan	Excellent:	Good:	Don't Know:	Fair:	Poor:

evaluation:: informal survey

- Results: systems ranked by survey

System	Average Rating	System	Novelty	System	Rel. Precision
Sun Tags	4.02	PC Joe	0.64	Sun Tags	0.49
Music Match	3.68	All Music	0.58	last.fm	-0.02
last.fm	3.50	PC David	0.57	All Music	-0.16
Sun Users	3.48	Pandora*	0.57	Sun Users	-0.38
Mystrands	3.26	Unnamed beta	0.48	PC Joe	-0.47
PC Brian	2.89	PC Chris	0.47	MusicMatch	-0.49
PC Dom	2.76	MusicMobs	0.47	PC Brian	-0.70
Up to 11	2.59	PC Dom	0.45	PC Dom	-1.01
All Music	2.06	iLike	0.38	Mystrands	-1.17
iLike	1.82	PC Brian	0.33	Up to 11	-1.18
PC Chris	1.64	Up to 11	0.32	PC David	-1.56
PC Joe	1.59	Sun Tags	0.31	Pandora*	-1.95
PC David	1.14	last.fm	0.28	PC Chris	-2.00
Pandora*	0.89	Sun Users	0.26	iLike	-3.31
Musicmobs	0.82	Musicmatch	0.24	MusicMobs	-4.32
Unnamed Beta	-2.39	Mystrands	0.23	Unnamed beta	-13.18

outline

- Introduction
- Formalization of the recommendation problem
- Recommendation algorithms
- Problems with recommenders
- Recommender examples
- Evaluation of recommenders
- **Conclusions / Future**

Conclusions / Future

- Coming soon: The Celestial Jukebox
  - All music in the world will be online
  - Millions of new tracks will be added every day
  - Lots of commercial interest
- Music tools will be essential
  - Exploration
  - Discovery
  - Recommendation
  - Organization
  - Playlisting

Conclusions / Future

- Current tools for finding music are inadequate
  - Like the web – pre-Google
  - Lots of problems:
    - scale, coldstart,
    - transparency
    - feedback loops
  - One size doesn't fit all
- Problems are opportunities for researchers

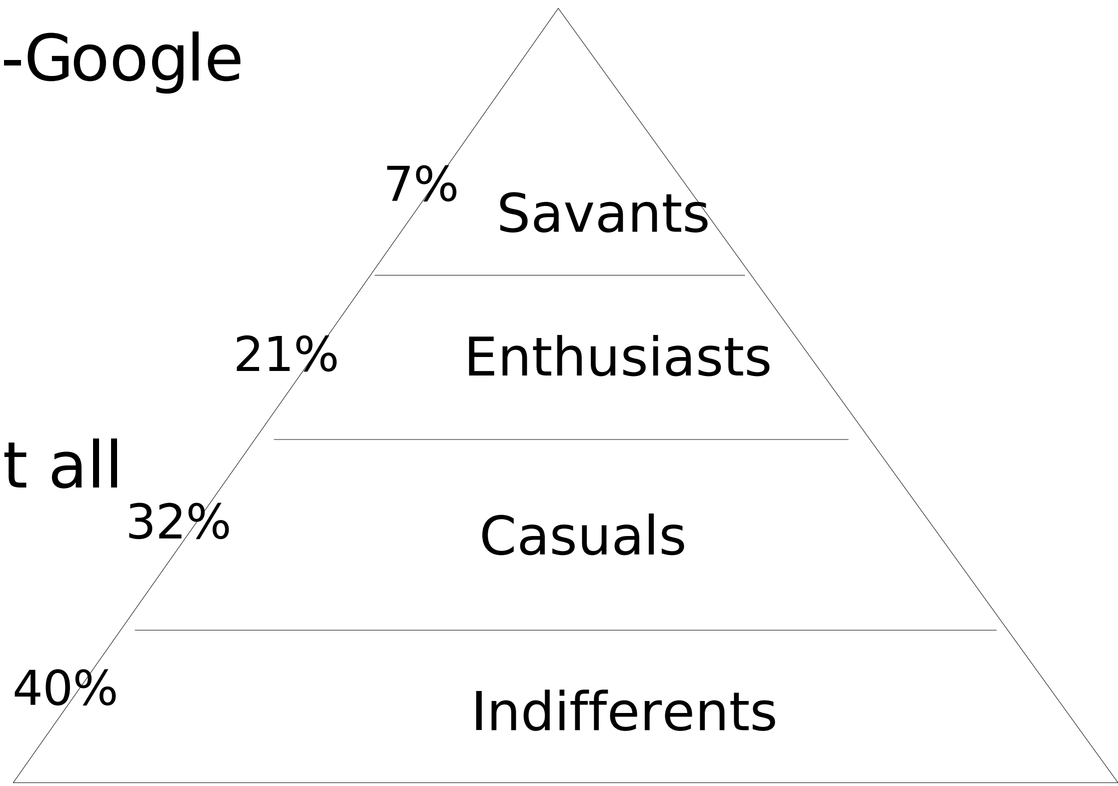

conclusions:: MIR Wishlist

- Big Problems
  - Coldstart
  - Feedback loops
  - Transparency
  - Scaling
  - Evaluation
    - Not just predicting ratings
    - Capture novelty / serendipity

conclusions:: MIR Wishlist

- Determining audio-based music similarity
- Extracting semantic descriptors from audio
- Recommendation for devices
- Recommendations for groups
- Combining different data sources
- Segmenting songs
- Creating intelligent playlists
- Creating user interfaces to browse music
- Learning from skipping behavior

conclusions:: MIR Wishlist

- Detecting cover songs
- Aligning lyrics with audio
- Separating audio sources
- Detecting and exploiting 'trendsetters'
- Extracting and aligning beat
- Supporting users in remixing/modifying their favorite songs
- Dealing with dirty, inconsistent metadata
  - As RJ from last.fm asks: “Just how many ways to write “Guns N’ Roses – Knockin’ on Heaven’s Door” are there?

And finally…

This is a very exciting time to be in a very fun field

Acknowledgments

- Stephan Baumann – DFKI
- Thierry Bertin-Mahieux - UdeM
- Jim Bennet – Netflix
- Norman Casagrande – last.fm
- Pedro Cano – BMAT
- Sean Columbo - LyricWiki
- Tom Conrad – Pandora
- Chris Dahlen – Pitchfork Media
- Justin Donaldson – Mystrands
- Douglas Eck – UdeM
- Dan Ellis - Columbia
- Bill Goldsmith – Radio Paradise
- Steve Green – Sun
- Brian Howe – Pitchfork Media
- David Jennings – NBRR Author
- Zac Johnson – All Music
- Richard Jones – last.fm
- Greg Linden – Amazon / Findory
- Alex Loscos – BMAT
- Dominique Leon – Pitchfork Media
- Francisco Martin – Mystrands
- Michael Mandel – Columbia
- Bamshad Mobasher – DePaul University
- Elias Pampalk – last.fm
- Jeremy Pickens – FXPAL
- David Raposa – Pitchfork Media
- Joe Tangari – Pitchfork Media
- David Tchang – One Llama
- Kris West – One Llama
- Ian Wilson – Zukool
- Mark Young – Itunes Registry

Music Recommendation Tutorial

Òscar Celma
Paul Lamere

ISMIR 2007
September 23, 2007

misc:: patents

- MusicIP
- Polyphonic HMI
- Philips
- MyStrands
- Microsoft

misc:: patents

- Music Recommendation system and method
  - Filing Date: 08/13/2004 – Publication: 02/17/2005
  - Authors (MusicIP, formerly Predixis)
    - Hicken, Wendell T. (La Verne, CA, US)
    - Holm, Frode (Santa Barbara, CA, US)
    - Clune, James Edmond III (Glendora, CA, US)
    - Campbell, Marc Elroy (Monrovia, CA, US)
  - Keywords
    - audio fingerprint, song similarity, playlist generation
  - United States Patent 20050038819

misc:: patents

- Method and system for music recommendation
  - Filing Date: 10/03/2003 – Publication: 06/10/2004
  - Authors (Polyphonic HMI)
    - Alcalde, Vicenc Gaitan (Castella del Valles, ES)
    - Ullod, Carlos Maria Lopez (Zaragoza, ES)
    - Bonet, Antonio Trias (Sant Cugat del Valles, ES)
    - Llopis, Antonio Trias (Sant Cugat del Valles, ES)
    - Marcos, Jesus Sanz (Barcelona, ES)
    - Ysern, Daniel Caldentey (Barcelona, ES)
    - Arkwright, Dominic (Barcelona, ES)
  - Keywords
    - song similarity (FFT, chunks, avg. values), vector similarity, user's taste vector, relevance feedback
  - United States Patent 20040107821

misc:: patents

- Sharing music essence in a recommendation system
  - ❖ Filing Date: 05/23/2006 – Publication: 11/23/2006
  - ❖ Author (MusicIP)
    - Hicken, Wendell T. (La Verne, CA, US)
  - ❖ Keywords
    - playlist characterization, playlist sharing, fill-in the gap, modify playlist
  - ❖ United States Patent 20060265349

misc:: patents

- Introducing new content items in a community-based recommendation system
  - ❖ Filing Date: 10/27/2003 - Publication: 04/20/2006
  - ❖ Authors (Philips)
    - Bodlaender, Maarten Peter (Eindhoven, NL)
    - Hollemans, Gerrit (Eindhoven, NL)
    - Vignoli, Fabio (Eindhoven, NL)
  - ❖ Keywords
    - user community, comparing user profiles, generating a recommended user set for the user
  - ❖ United States Patent 20060085818

misc:: patents

- Client-based generation of music playlists from a server-provided subset of music similarity vectors
  - ❖ Filing Date: 01/27/2005 - Publication: 05/25/2006
  - ❖ Authors (Microsoft)
    - Platt, John (Redmond, WA, US)
    - Renshaw, Erin (Kirkland, WA, US)
  - ❖ Keywords
    - music similarity, hybrid graph, MDS embedding, euclidean space
  - ❖ United States Patent 20060112082

misc:: patents

- MyStrands
  - http://labs.mystrands.com/patents.html
  - 16 pending patents (from 2004-2006)
  - Examples
    - “Personal music recommendation mapping applet overview”
    - “Sharing tags between individual user media libraries”
    - “User to user recommender”
    - “Freeing space for new media items on a mobile media playback device based on inferred user taste”
    - “Building and sharing a composite playlist from collective group tastes on multiple media playback devices”
    - ...and a long etc.