Adaptive MIR – An Overview

Andreas Nürnberger and Sebastian Stober
Data & Knowledge Engineering Group, Faculty of Computer Science
Otto-von-Guericke-Universität Magdeburg, Germany

Email: andreas.nuernberger@ovgu.de, sebastian.stober@ovgu.de

Outline

- Day 1: Adaptation and Personalization: Concepts and Challenges
- **Day 2: Adaptive Music Retrieval: An Overview**
- Day 3: Adaptive Hierarchies: Constrained Clustering and Utility
- Day 4: Adaptive Music Similarity
- Day 5: User Interfaces and Gamification: Design and Evaluation

MIR Challenges [Downie’03]

music is multi-cultural

music information has many facets
and can be represented in multiple ways

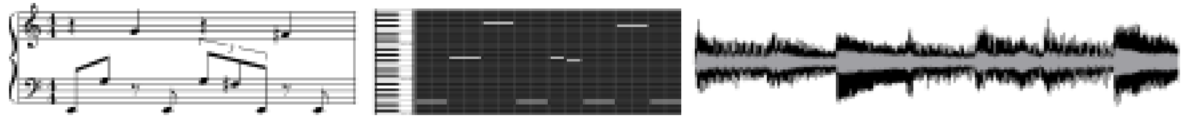

users of MIR systems are multi-disciplinary
and have varying information needs

music can be experienced in many ways
leading to different perceptions

Adaptive Systems – Definitions

- *behavior*:
  - (set-valued) input/output (I/O) function of a system
  - does not require knowledge about system internals
- *adaptable system*:
  - provides means to change its behavior
- *adaptation*:
  - change of internal system structure (invisible) and behavior (visible)
- *context*:
  - (operational) environment,
  - user context
  - data (i.e., input/output values and their characteristics)

Adaptive Systems

A system is (context) *adaptive* iff

1) it <u>behaves different</u> in different contexts given the same input [based on Broy et al. '09]

AND

2) the respective adaptation (i.e., the difference in behavior) is <u>goal-driven</u> in that it aims to optimize the system's behavior in the given context according to some pre-defined measure.

Adaptable → Adaptive System

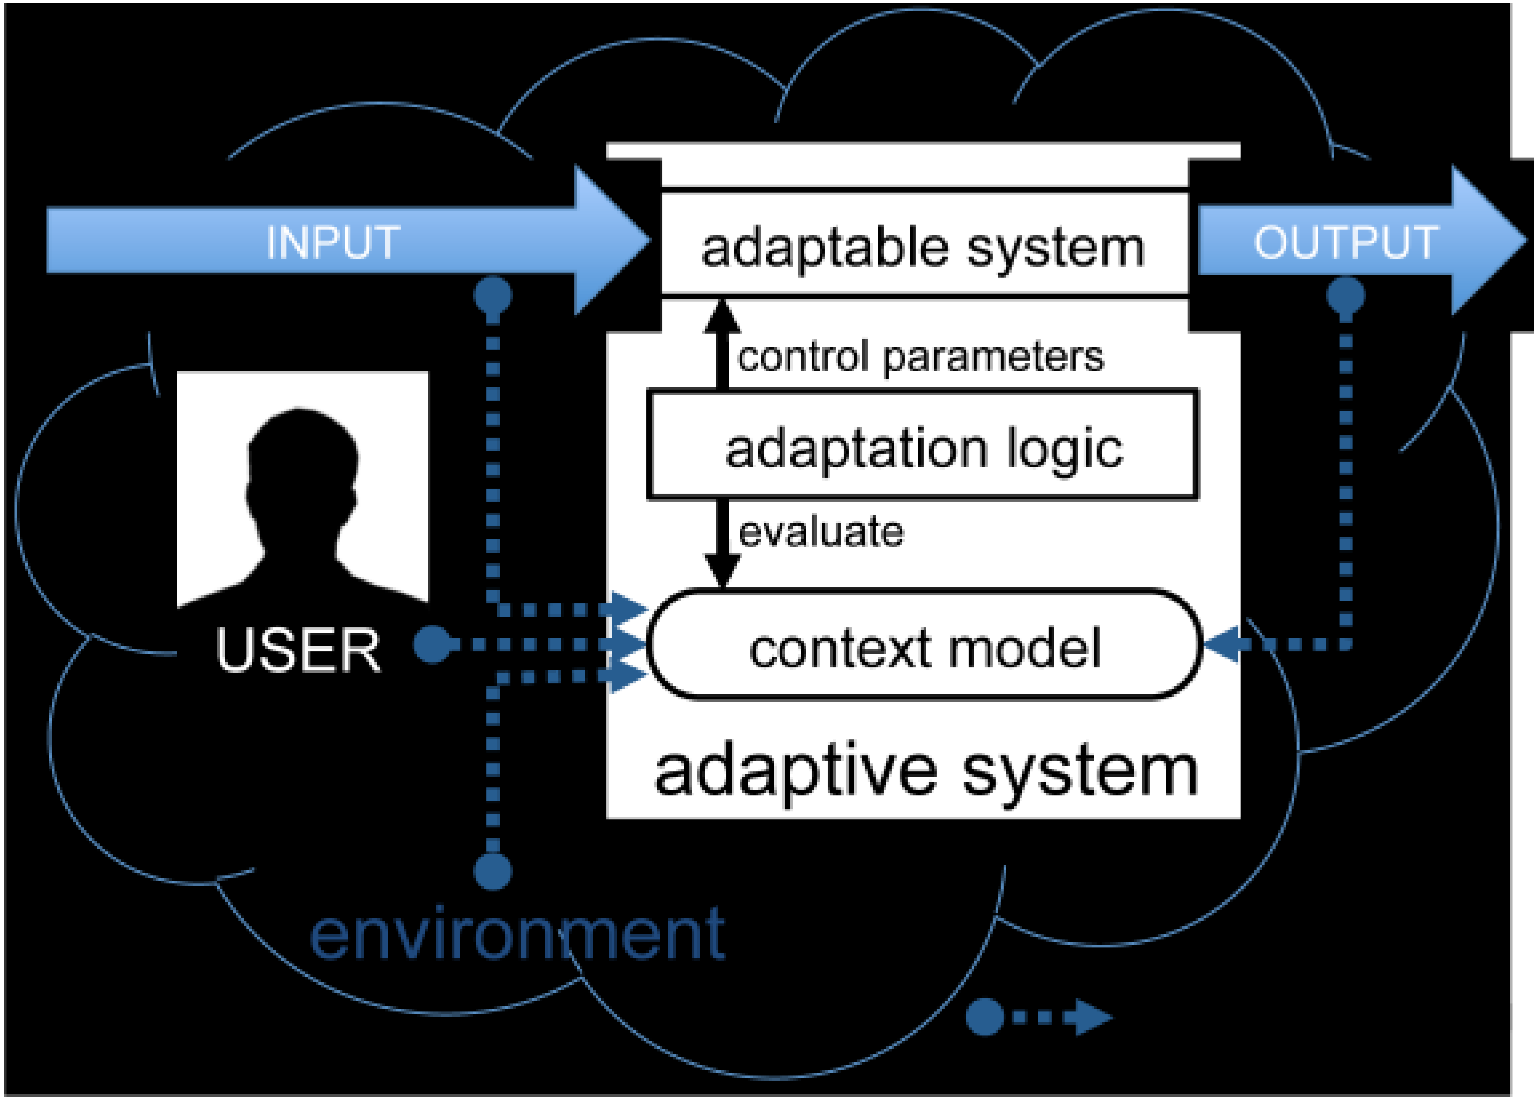

Task:

COLLECT EXAMPLES FOR ADAPTIVE APPROACHES IN MIR YOU KNOW OF!

A Typical MIR System

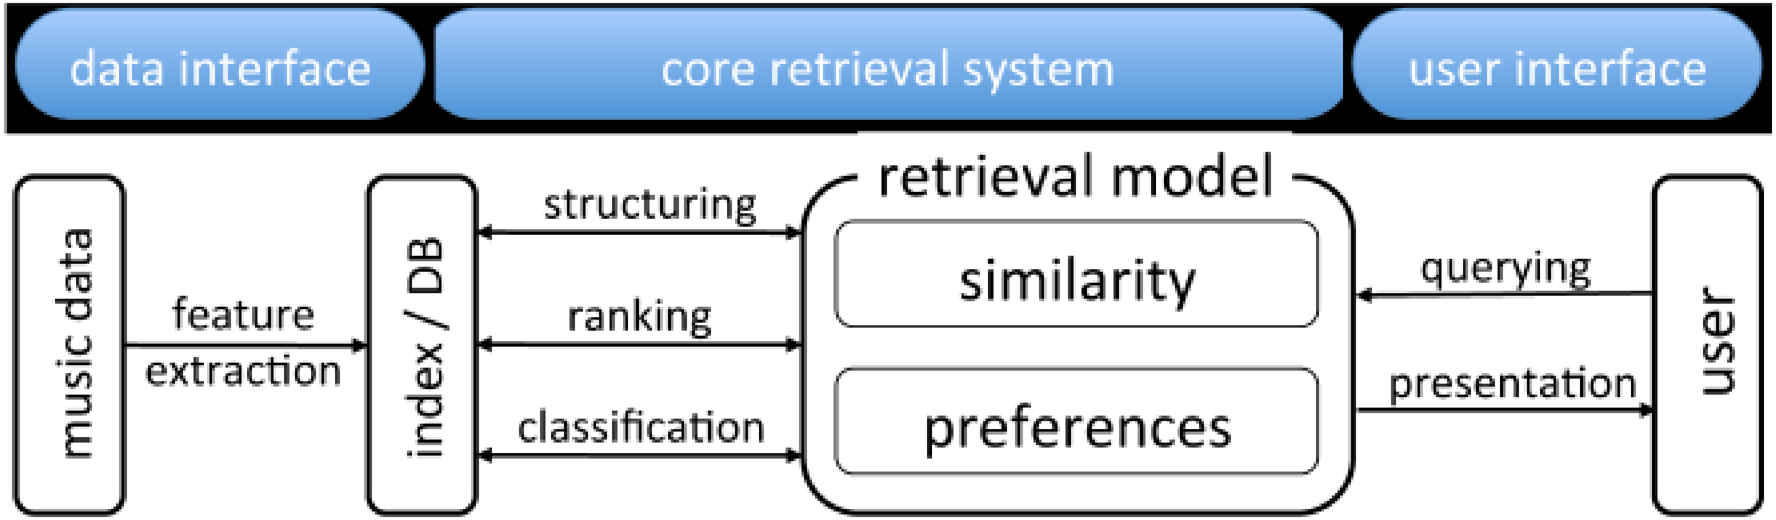

.... and typical adaptation goals:

adaptive component	adaptation goal
user-adaptive querying	better "understanding" of the user's information need
user-adaptive presentation	increase understandability
user-adaptive structuring	structures that reflect the user's individual way of structuring
data-adaptive structuring	respond to changes in the dataset
user-adaptive ranking	rank according to the user's understanding of relevance
data-adaptive ranking	increase diversity within the result list
user-adaptive classification	reflect the user's classification criteria
data-adaptive feature extraction	increase robustness and quality of the extracted features

ADAPTIVE FEATURE EXTRACTION

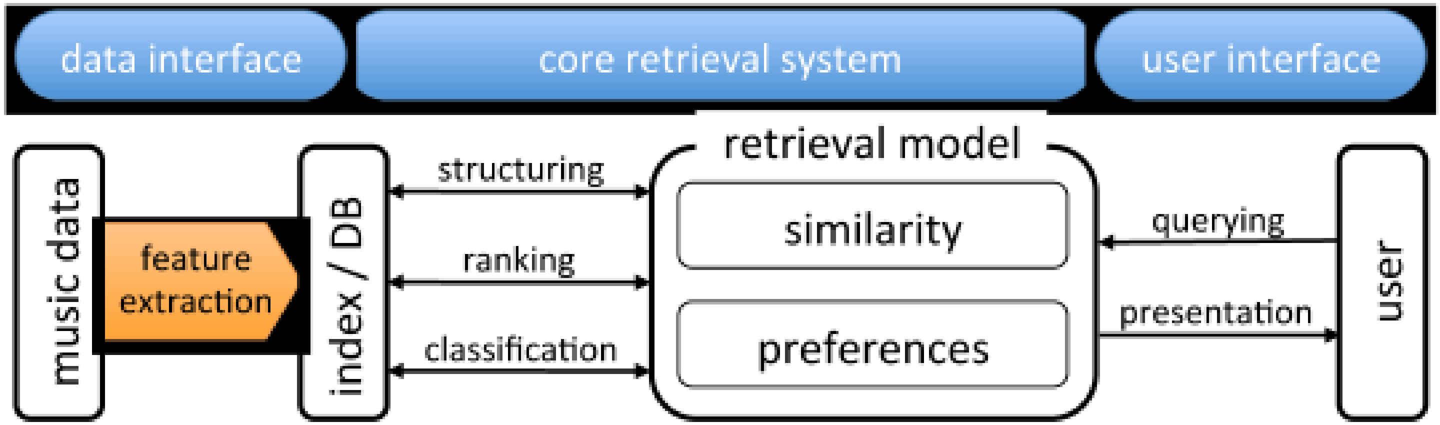

Adaptive Background Removal

- goal: extract melody voice (rough approximation)
- idea:
  - exploit standard noise removal techniques
  - use local properties of karaoke track (estimated) for noise profile

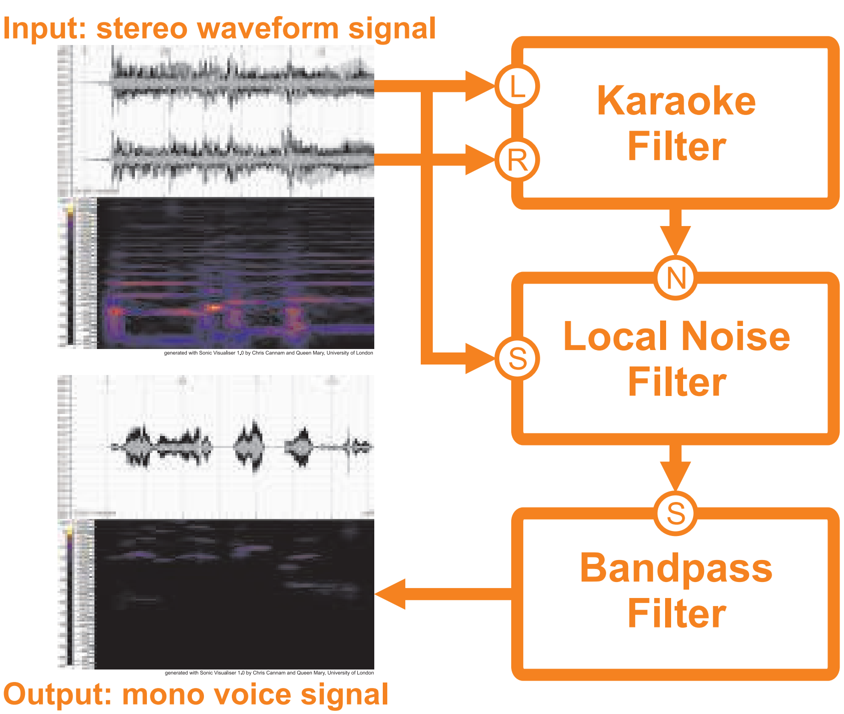

Idea: exploit spatial arrangement of instruments and voices in the mix

Karaoke Filter:
removes center pan (information contained in both channels) by inverting one channel and mixing it together with the other into a mono signal:

$$\text{output} = L - R$$

Requirements:
- stereo input signal
- lead voice (and possibly solo instruments) centered in the stereo mix
- instruments and backing vocals arranged out of center

Local Noise Filter:
derives a local (i.e. continuously updated) power spectrum of frequencies from a noise signal (N) which can then be removed from the signal (S)
(based on versions 1.34, Sep 23, 2006 and 1.39, Jul 27, 2007 of the NoiseRemoval effect by Dominic Mazzoni as part of Audacity)

Bandpass Filter (300-3000Hz):
keeps only frequency range of the input signal (S) that is relevant for human voice (lower bound is higher to filter out the bass guitar that might be in the center as well)

Adaptive Chord Recognition

- goal: correct chord misclassifications
- idea: exploit knowledge about (estimated) chord distribution

? n=16

... Am C C C C G Em Am Am Cm C G G A Am Am Am C ...

local distribution

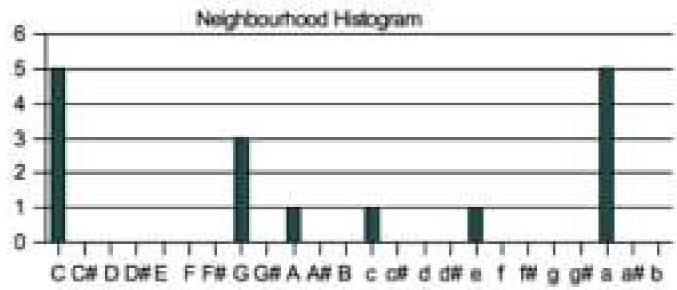

classifier output

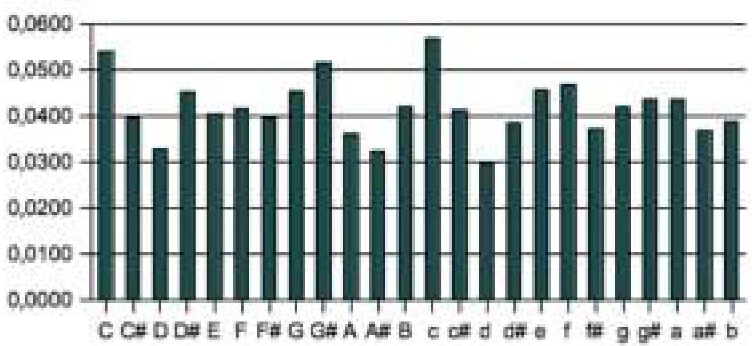

corrected output

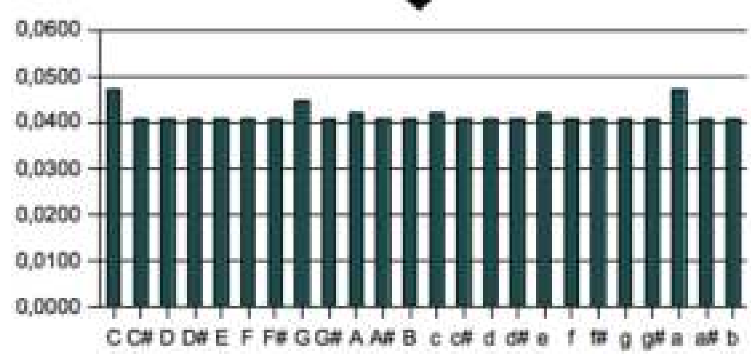

Adaptive Chord Recognition

- HMM-based approaches use Viterbi decoding
  - finds the optimal path in a sequence of states based on
    - similarity of the observed input to the state output probability distribution
    - transition probabilities contained in the transition matrix.
  - transition matrix can be adapted based on
    - general musical knowledge
    - the current music piece

- more adaptive approaches using context:
  - tuning
  - key
  - bass pitch
  - beat-synchronous chroma features

Peak Picking with Adaptive Thresholding

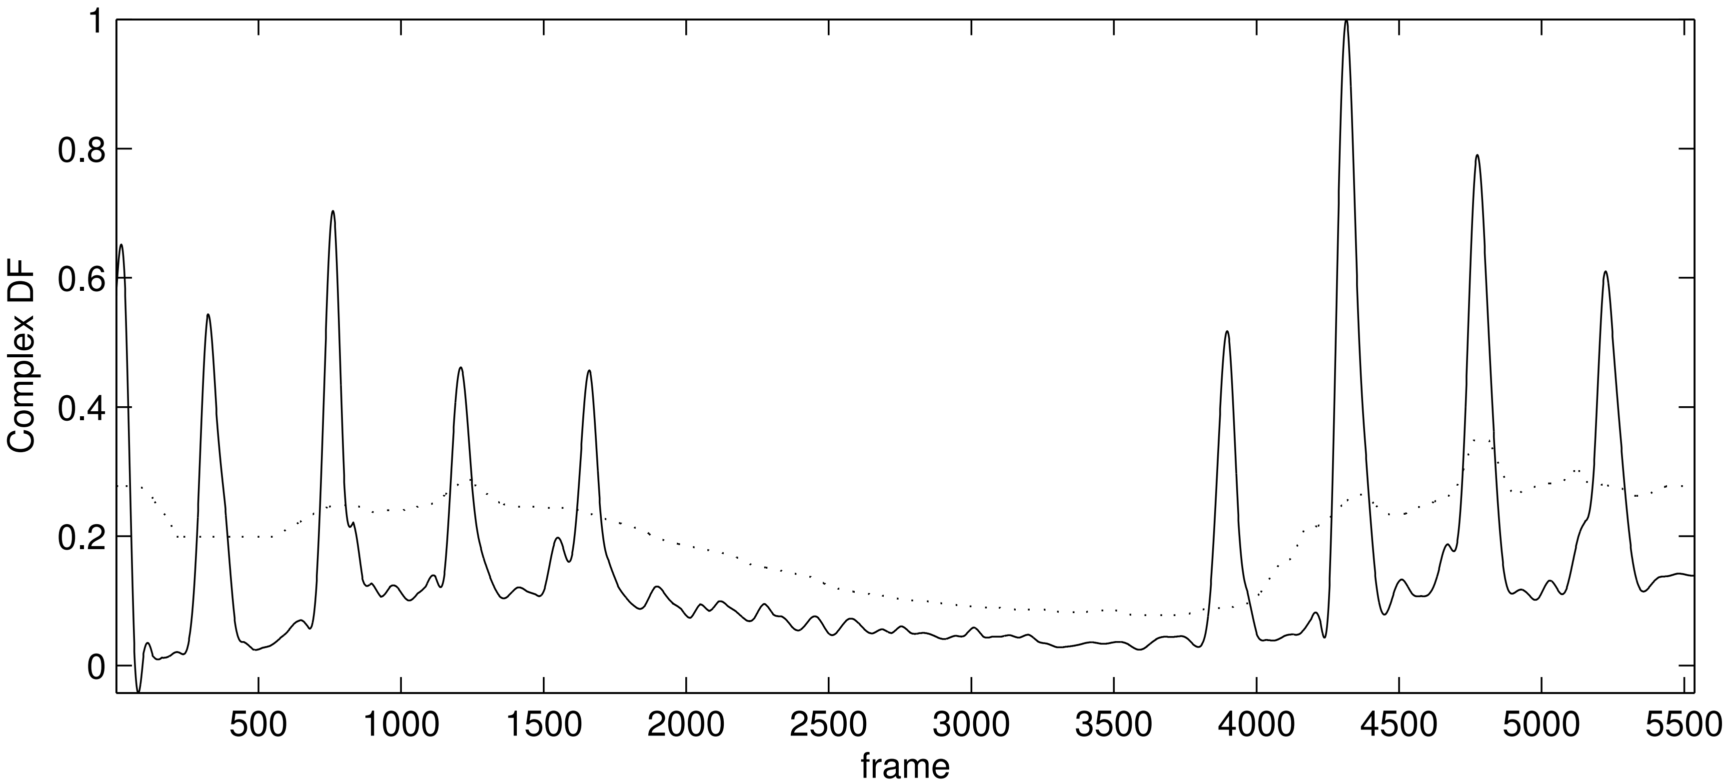

- moving median thresholding
  [Duxbury et al. 2003, Complex Domain Onset Detection for Musical Signals, DAFX]

- context: local value distribution

More Sophisticated Chord Context Models

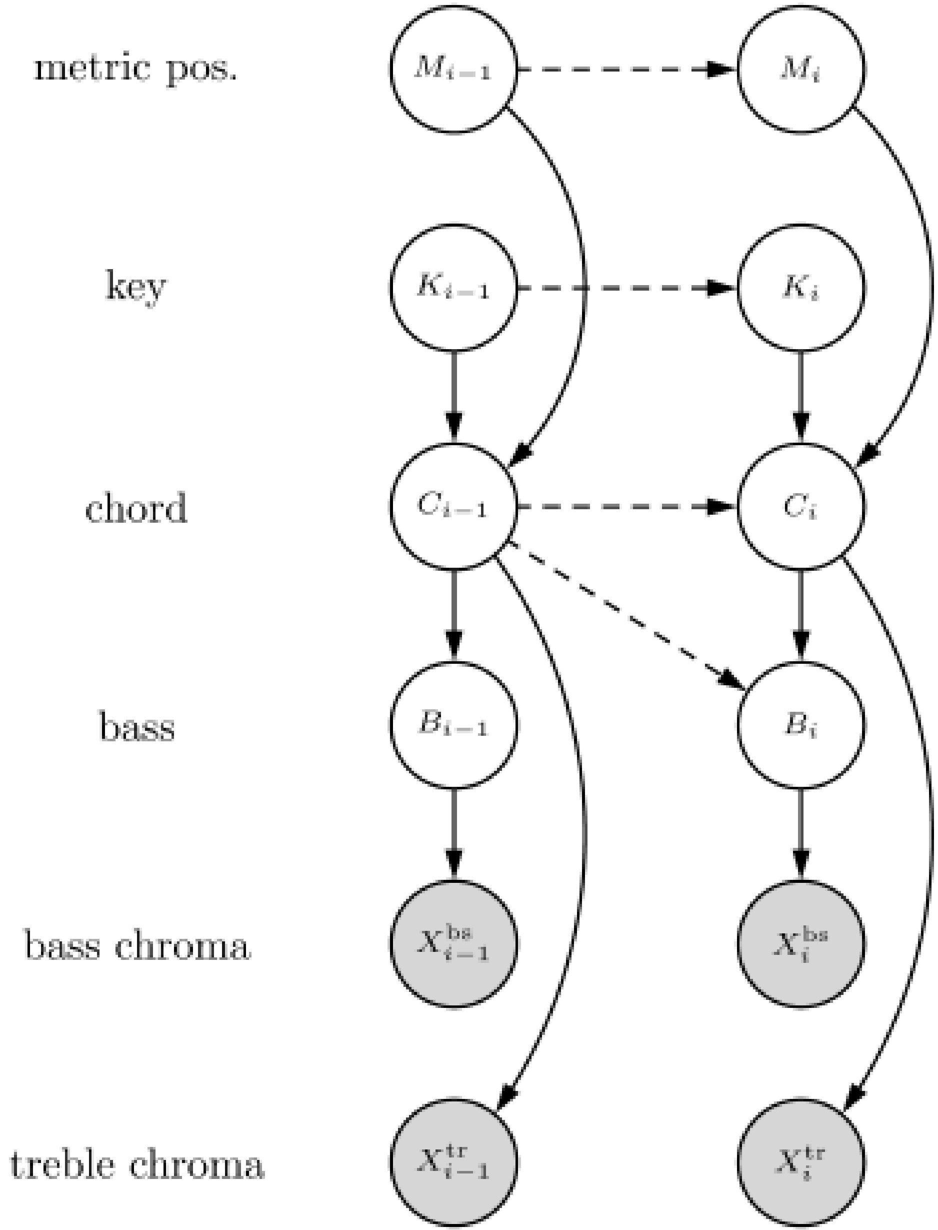

[Mauch & Dixon 2010]

Singer-Adaptive QBSH

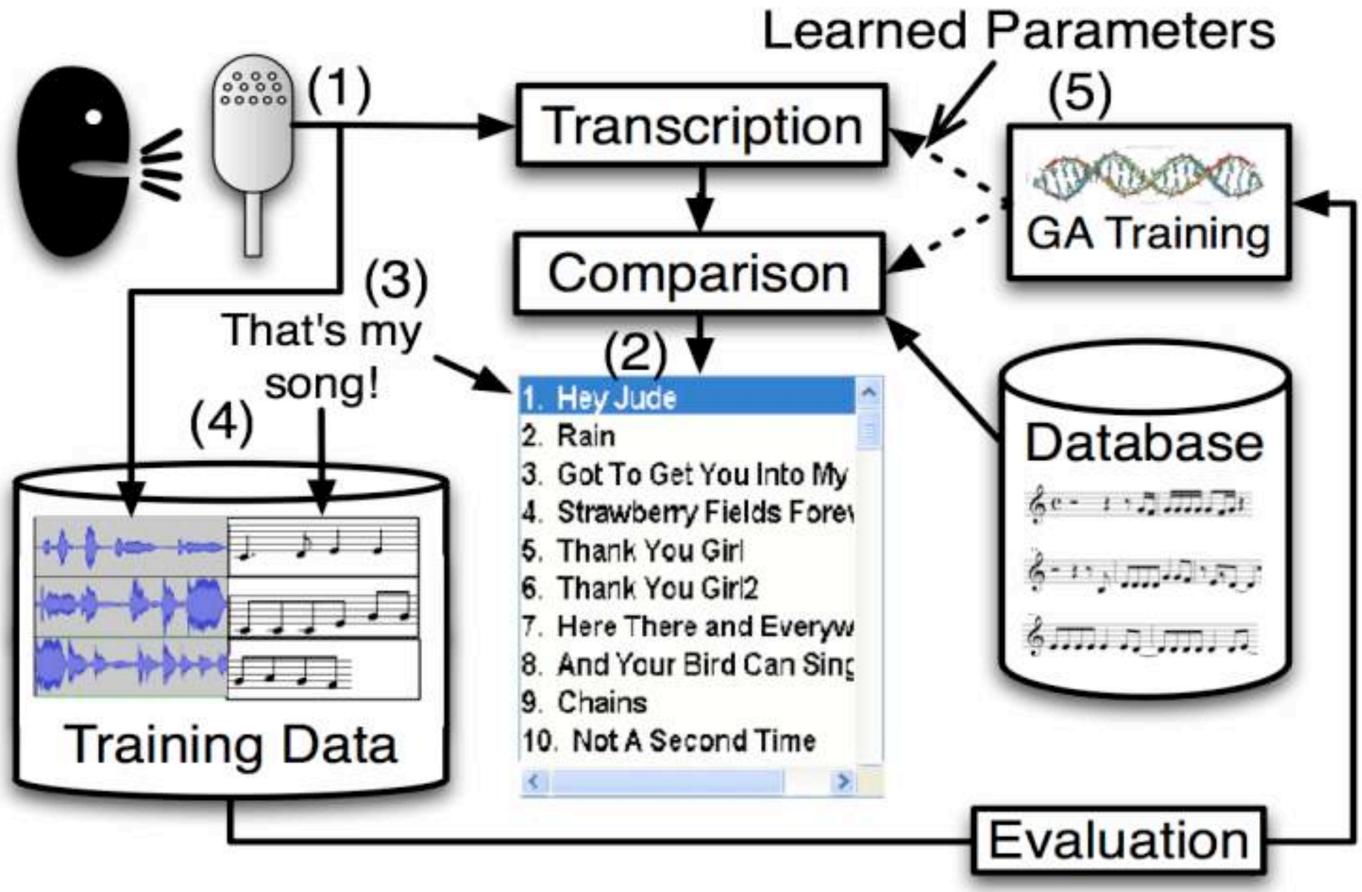

- trainable note segmentation system
- easily parametrized singer error model
  [Little et al. 2007, A Query By Humming System That Learns From Experience, ISMIR]

Adaptive Optical Music Recognition

(a) RISM 1528-2 (Attaignant, Paris, 1528)

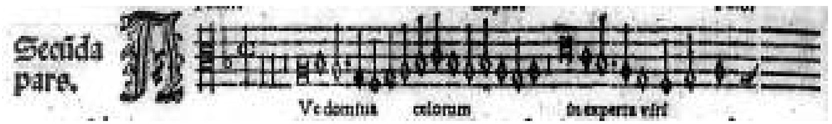

(b) RISM 1532-10 (Moderne, Lyon, 1532)

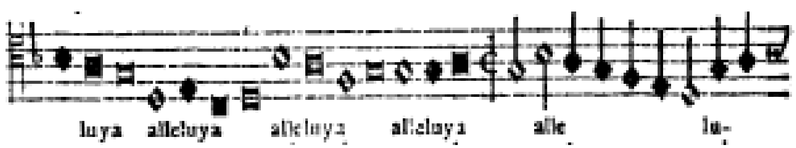

(c) RISM V-1421 (Figliuoli di Gardano, Venezia, 1572)

(d) RISM V-1433 (Basa, Roma, 1585)

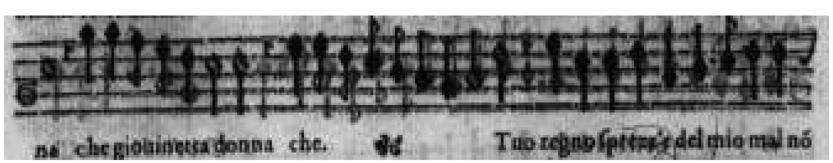

(e) RISM M-0582 (Le Roy & Ballard, Paris, 1598)

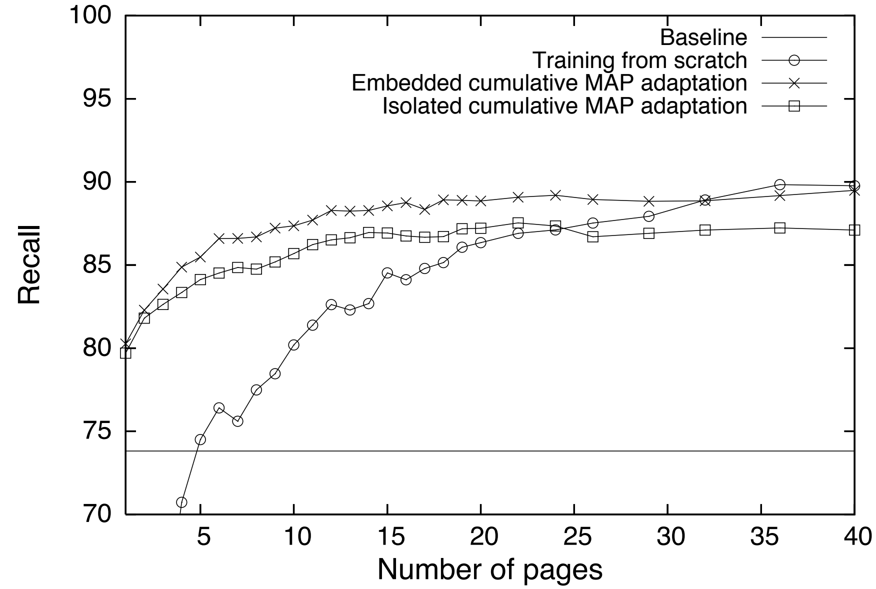

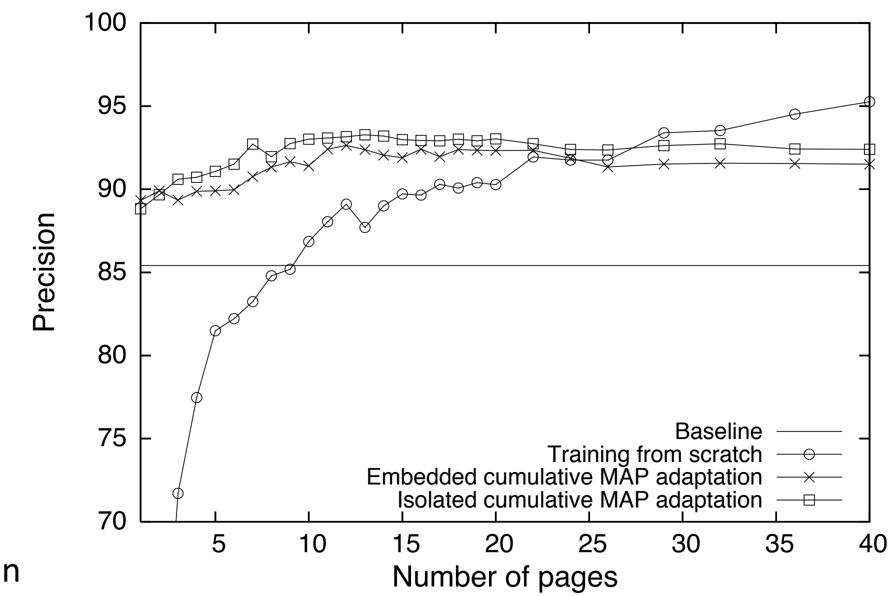

[Pugin et al. 2007, MAP adaptation to improve optical music recognition of early music documents using hidden markov models, ISMIR]

Some More Examples

- tempo estimation with adaptive window size
  - context: onsets
- pitch detection with adaptive strategy
  - context: monophonic vs. polyphonic signal
- F0-estimation with adaptive tone models
- key recognition with adaptive frequency weighting

ADAPTIVE DATA STRUCTURES

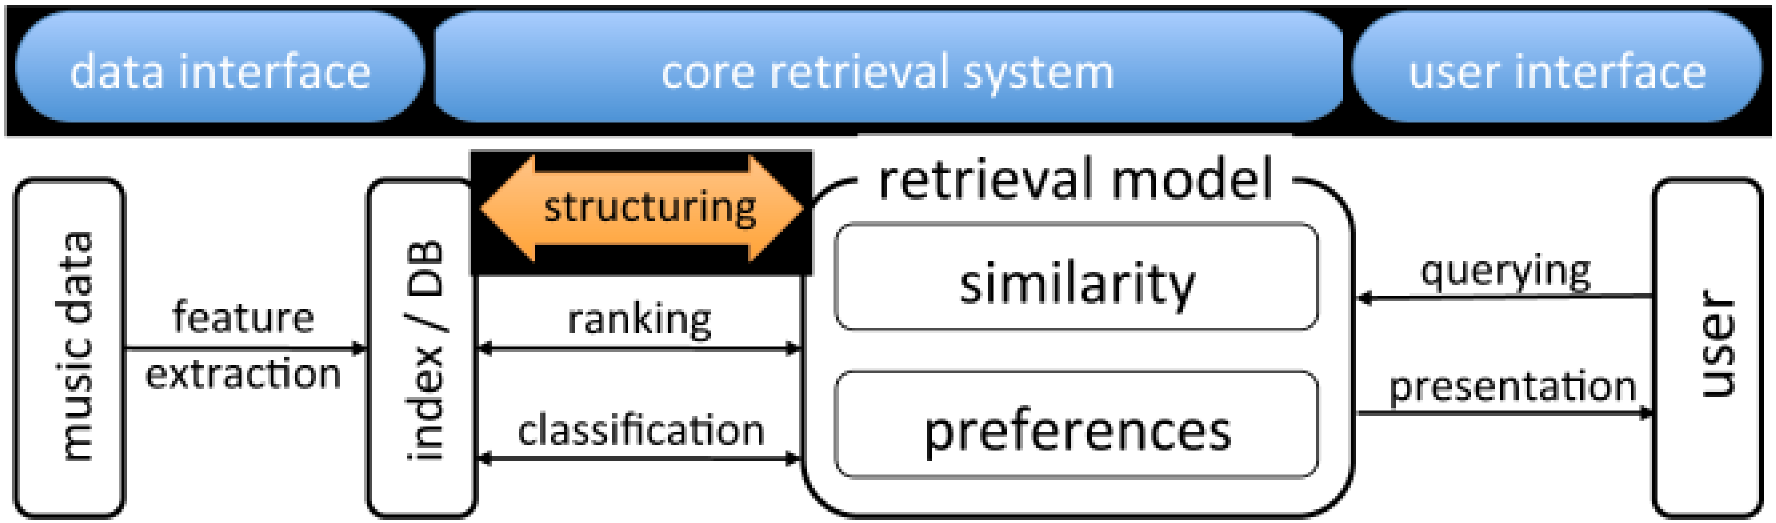

Self-Organizing Maps: Data-Adaptive Structures

Islands of Music
[Pampalk et al. 2003]

MusicMiner
[Mörchen et al. 2005]

BeatlesExplorer
[Stober et al. 2008]

- growing SOMs:

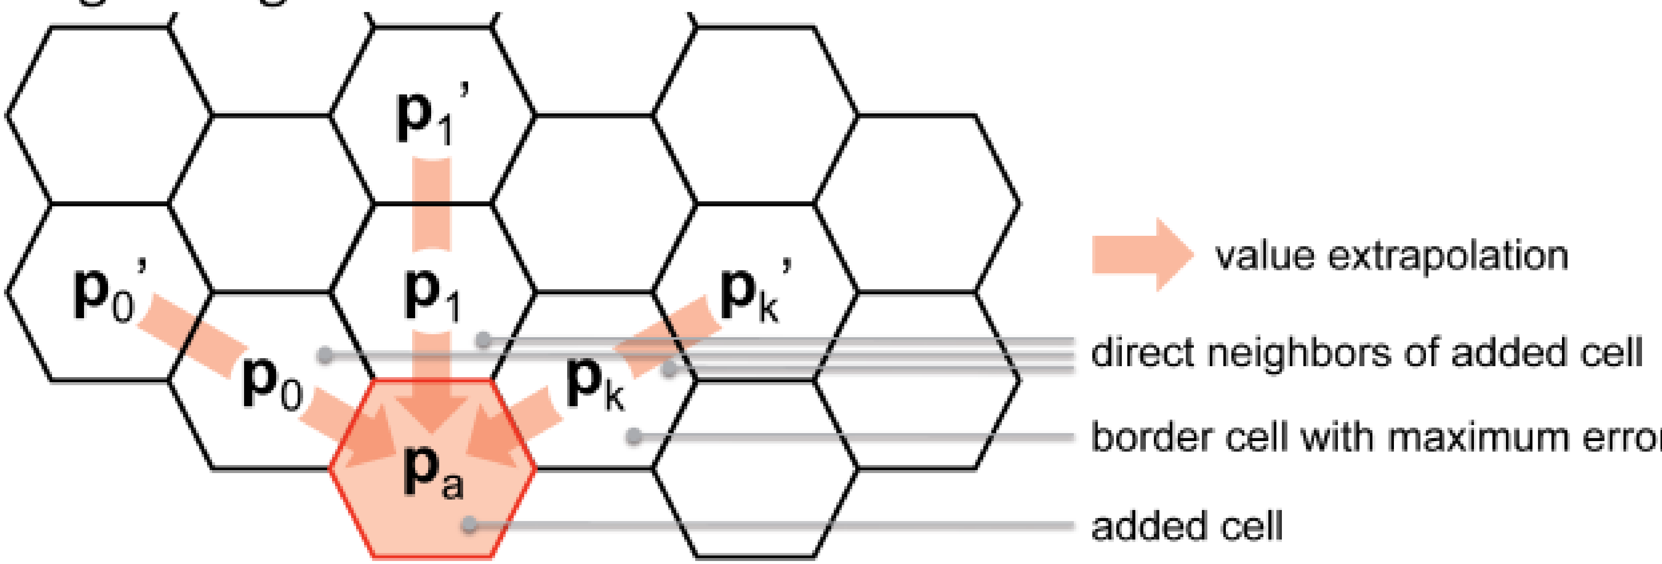

Adaptive Music Recommendation

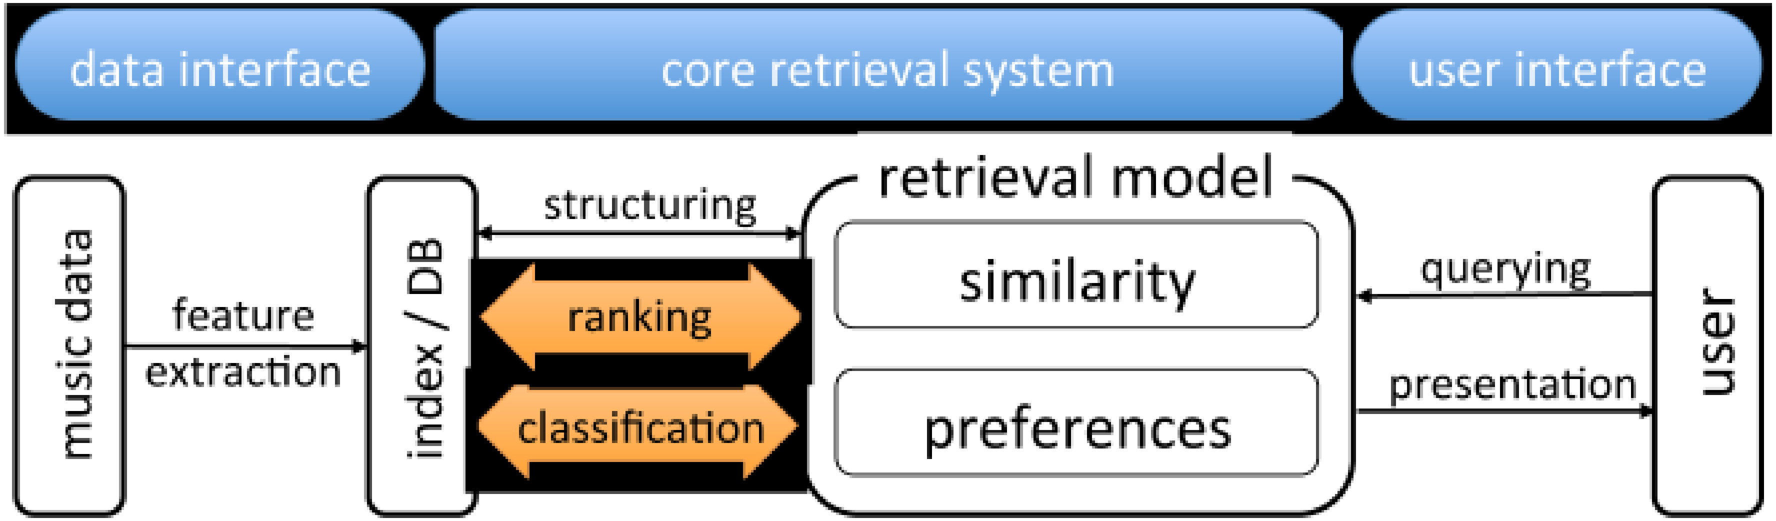

Select Songs for a Mobile Player [Pohle et al. 2008]

- given:
  - large music collection (on a desktop PC)
  - mobile player device with limited capacity (k songs)
- challenge:
  - Which songs to remove from and add to the mobile collection?
- "like value": $\mathtt{lv}_t = \dfrac{l_t}{l_t + s_t}$ (l = listened, s = skipped)
- removal strategies:
  a) skipped ≥ 1 (most removals)
  b) skipped > listened (most reluctant)
  c) lv < 0.75
  d) userbanned (most reliable)

Select Songs for a Mobile Player [Pohle et al. 2008]

- selection strategies:
  a) random
  b) in the order of their distance to the <u>initial</u> selection
  c) songs closest to any of the songs listened to during the last listening session
  d) songs closest to any of the last k tracks that have been accepted (listened to)
  e) select songs w.r.t. last k accepted songs (set A) and last k skipped songs (set S)
     - select songs more similar to A (sorted by distance)
     - continue adding songs sorted by ratio d_A/d_R
  - optional: weighting tracks in A and R by listened/skipped counts

- best (in simulation) weighted c/d/e

Adapting to the Listening Context

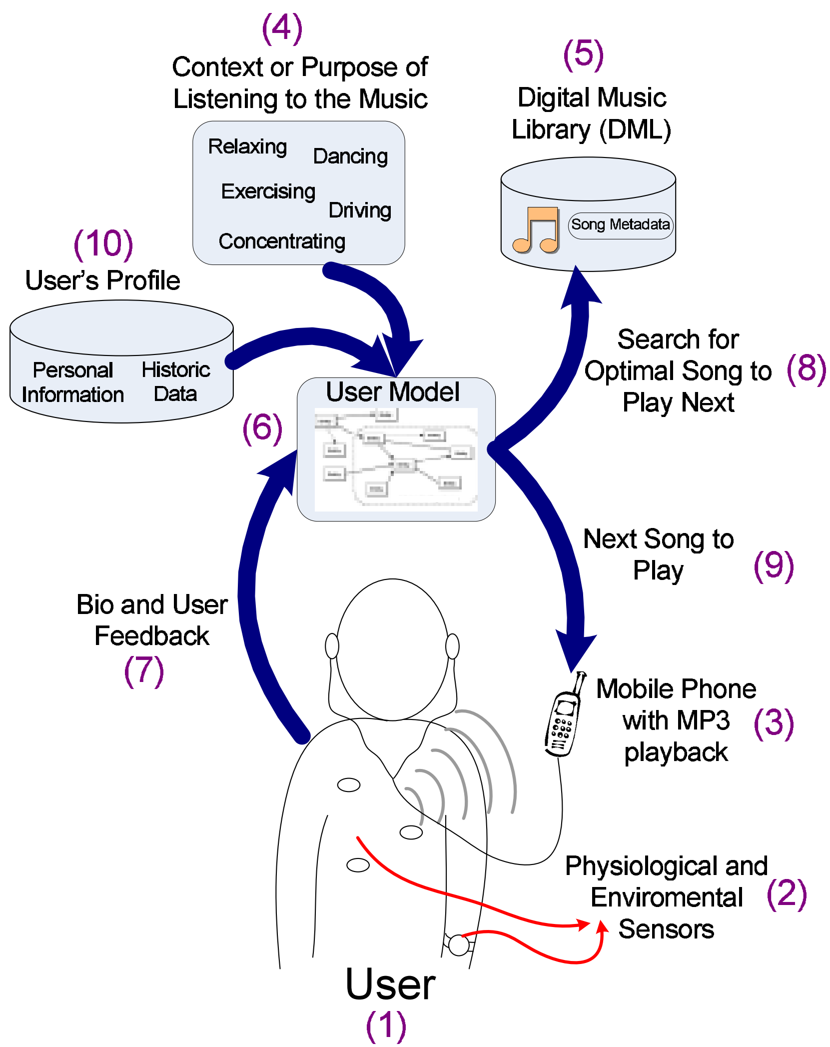

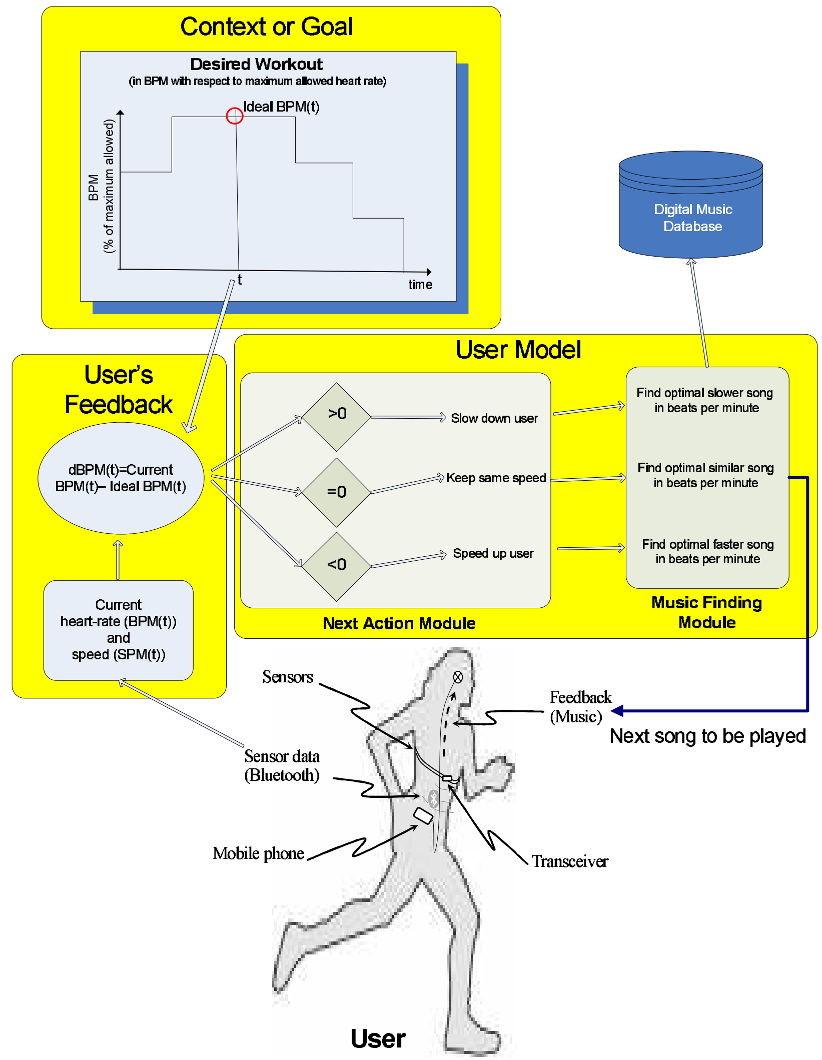

[Oliver & Kreger-Stickles 2006, PAPA: Physiology and Purpose-Aware Automatic Playlist Generation, ISMIR]

Commercial Applications

YAMAHA

日本サイトはこちら

BODiBEAT
BODY. MUSIC. IN SYNC.

PRODUCT / MUSIC IN EXERCISE / SONG DATABASE / REGISTRATION / SUPPORT / EVENTS / TESTIMONIALS

BUY NOW!

[http://www.yamaha.com/bodibeat/musicinexercise.asp]

BODiBEAT: MUSIC IN EXERCISE

Everyone knows that exercise is good for you, right? And through its Music & Wellness Institute, Yamaha conducts medical research that shows playing musical instruments, that is, creating your own music, is also good for you.

Well, how about listening to music while you exercise? It turns out that numerous scientific studies have shown what is probably intuitively obvious to you: Listening to music while you exercise decreases discomfort levels. In addition, it has been shown that listening to music with a tempo that matches the pace of your running reduces overall discomfort levels further still – The Yamaha BODiBEAT has been inspired by this knowledge. Your workout will be over before you know it when you use BODiBEAT.

Summary of Yamaha-sponsored research performed at the University of Tokyo on the Yamaha BODiBEAT.
The study clearly demonstrates that, for a constant running pace, the rating of perceived exertion (RPE) is minimized (and reported comfort levels are maximized) when subjects listen to music with a tempo that matches their pace. Read about it here: Verification of effects on reducing feelings of fatigue during exercise using the Yamaha BODiBEAT.

Other articles and research on the positive effects of music on exercise:

New York Times article that explores the motivational and fatigue-distracting effects of music during exercise. Specifically looks at the most popular motivational songs and beats per minute (BPM) of songs preferred by professional through casual runners. Read about it here: They're Playing My

Commercial Applications

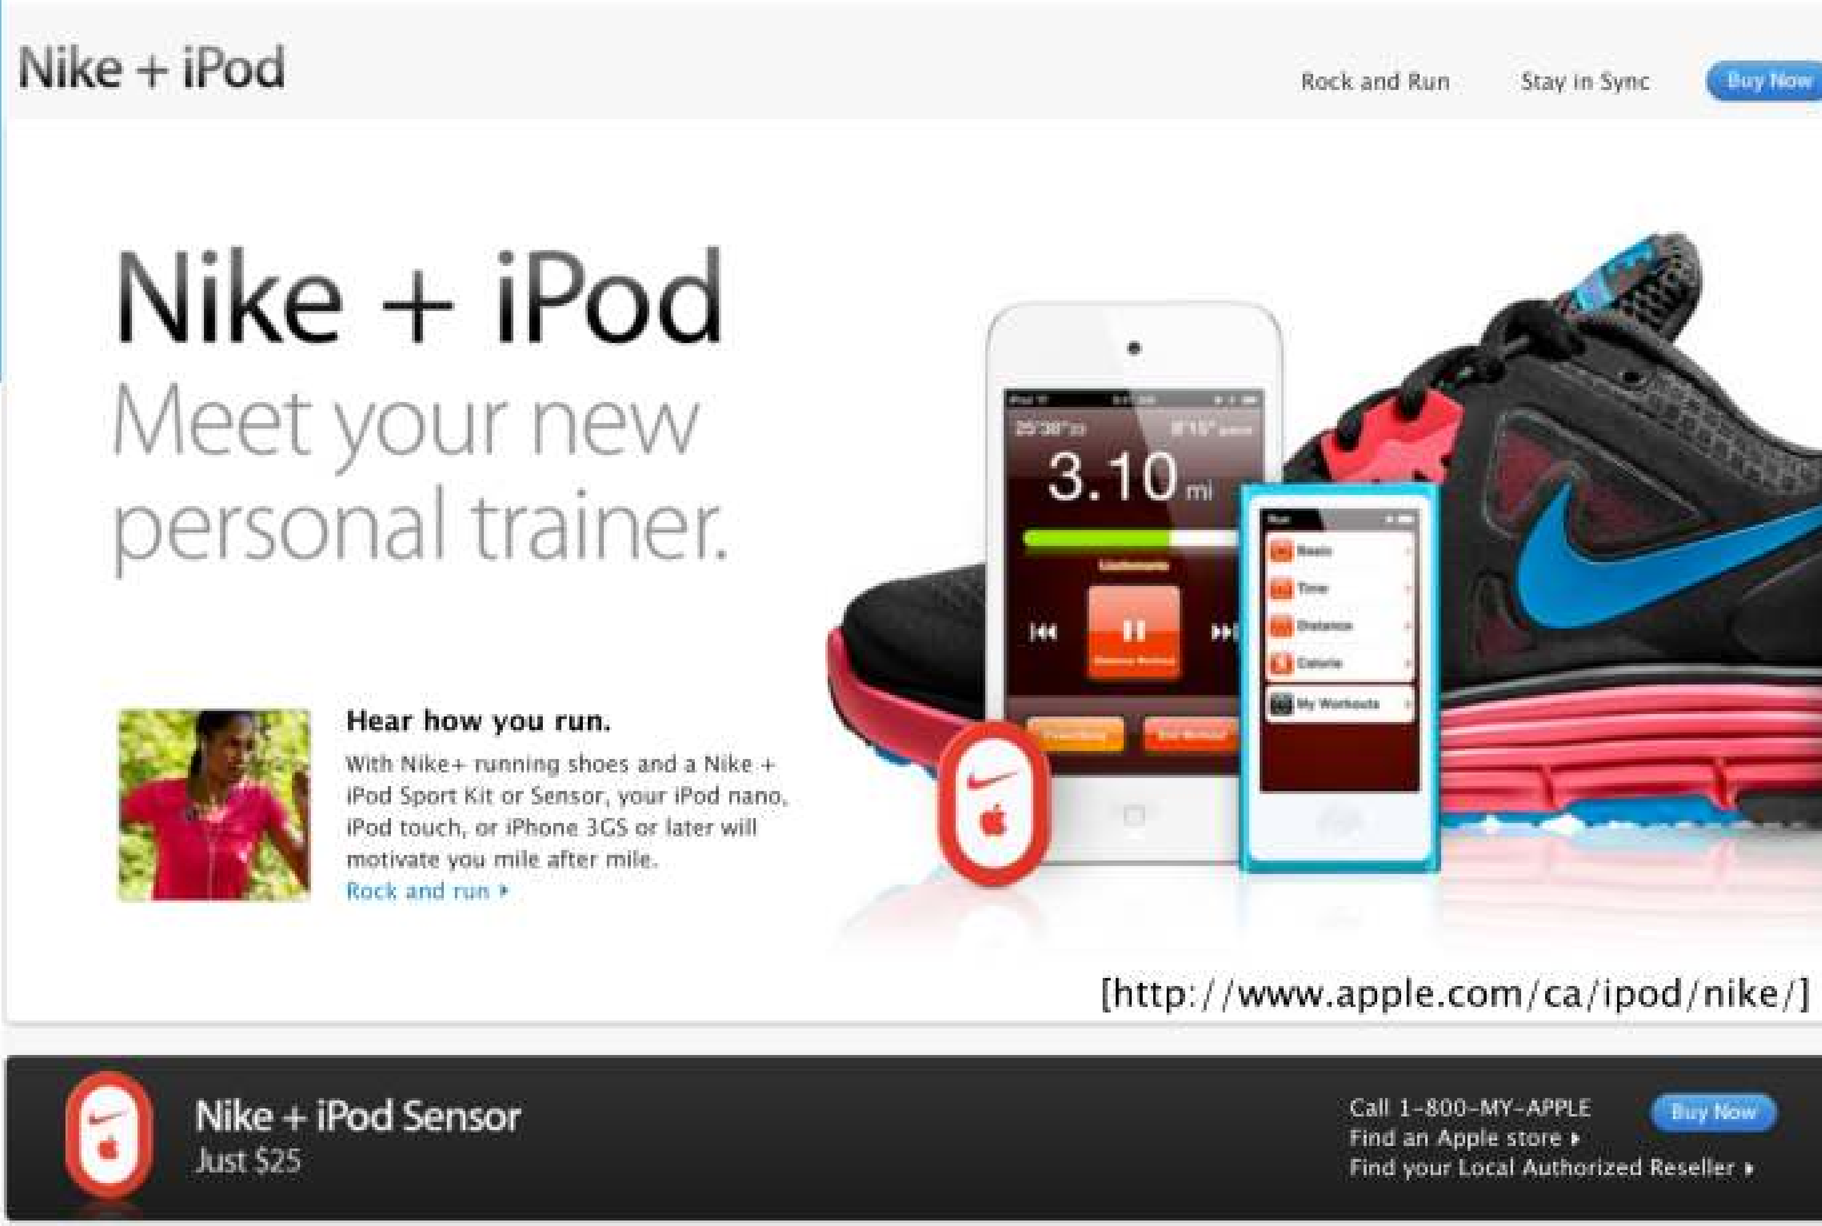

[http://www.apple.com/ca/ipod/nike/]

Thick of useless generic classifications into Rock/Pop/…?

HOW ABOUT USER-ADAPTIVE GENRES?

Why is usage context interesting?

- Jones et al. 2004:
  - One notable way that participants characterized music was by intended use - that is, based on the event or occasion at which they intended to listen to a particular set of music.
  - idiosyncratic genres (e.g. “driving music” “work music”)
- Lee & Downie 2004:
  - 41.9% of respondents said they would search or browse music information by “Associated usage.” […] This kind of extra-musical information is not traditionally incorporated in MIR systems.
- Hu et al. 2006:
  - […] conclude that the recommended usages specified by users reflect a meaningful source of user-generated metadata.

➡ exploit usage context for personalized structuring of music collections (possibly more meaningful than generic genres)

Related Work

- Pachet & Cazaly 2000:
  - Taxonomy of Music Genres
    - *Audience Location [...] describes the typical place where the music is usually listened to.*
    - *Danceability describes what dance type (if any), the music suggest, and can take various values such as “no” / ”rock” / ”salsa”, etc.*
- Govaerts et al. 2006:
  - Moody Tunes:
    - commercial application to select music for a desired atmosphere in hotels, restaurants and cafes
    - properties need to be assigned manually by experts and if necessary can only be adapted by hand

Related Work (cont.)

- Automatic context logging / guessing:
  - last.fm – audioscrobbler plug-in
    - tracks listening habits to build user profile and find users with similar taste
  - Lee & Lee 2006: M^3 music recommender system
    - season, month, day of the week, weather and temperature
  - Park et al. 2006: context-aware RS
    - weather (temp., humidity, current weather and forecast), time, ambient noise level, illuminance
  - Guan et al. 2006: RS for the "smart office"
    - location, time of day, people in the room, weather outside and the user's "stock portfolio" (???)
  - PAPA (Oliver and Kreger-Stickles 2006) and Yamaha BODiBEAT:
    - bio signals (e.g., pulse)

How can usage context be defined?

- context (Dey 2001):
  - any information that can be used to characterize the situation of a person, place or object of consideration
- four types of primary context (Dey & Abowd 2000):
  - location, identity, time and activity
- alternative way:

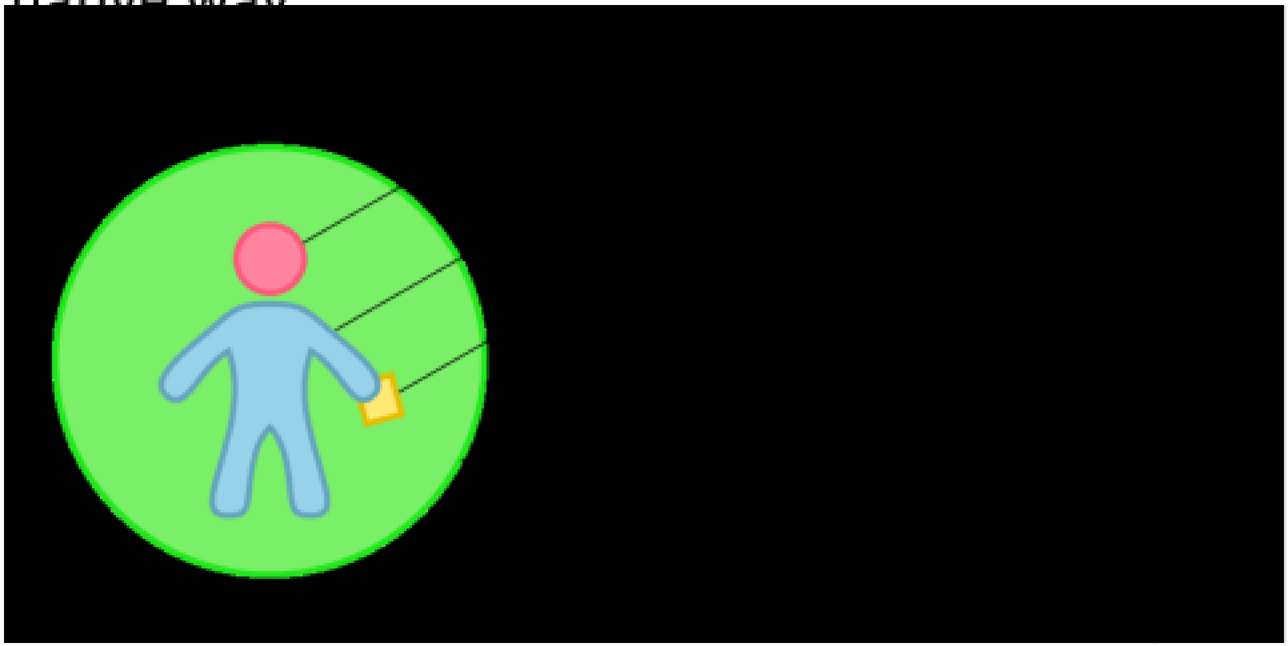

image © by Valentin Laube

Which data could be used as context?

- time
  - time of day, day of week, season, date
- location
  - device
  - GeoIP
  - GPS information (mobile devices)
    - infer speed and way of traveling (e.g. stationary, walking, running, cycling, driving, going by train, ...)
- activity / state (corr. with location, device, time)
  - access to calendar
  - active applications (e.g., browser, word processor, ...)
  - keyboard & mouse events per minute
  - people around (e.g., via bluetooth sensing)
  - ambient noise (level / classification)
  - illumination
  - acceleration / gyro sensor
  - bio signals
  - facial expression (web cam)

Prototype context logger

- for Foobar, Winamp and iTunes

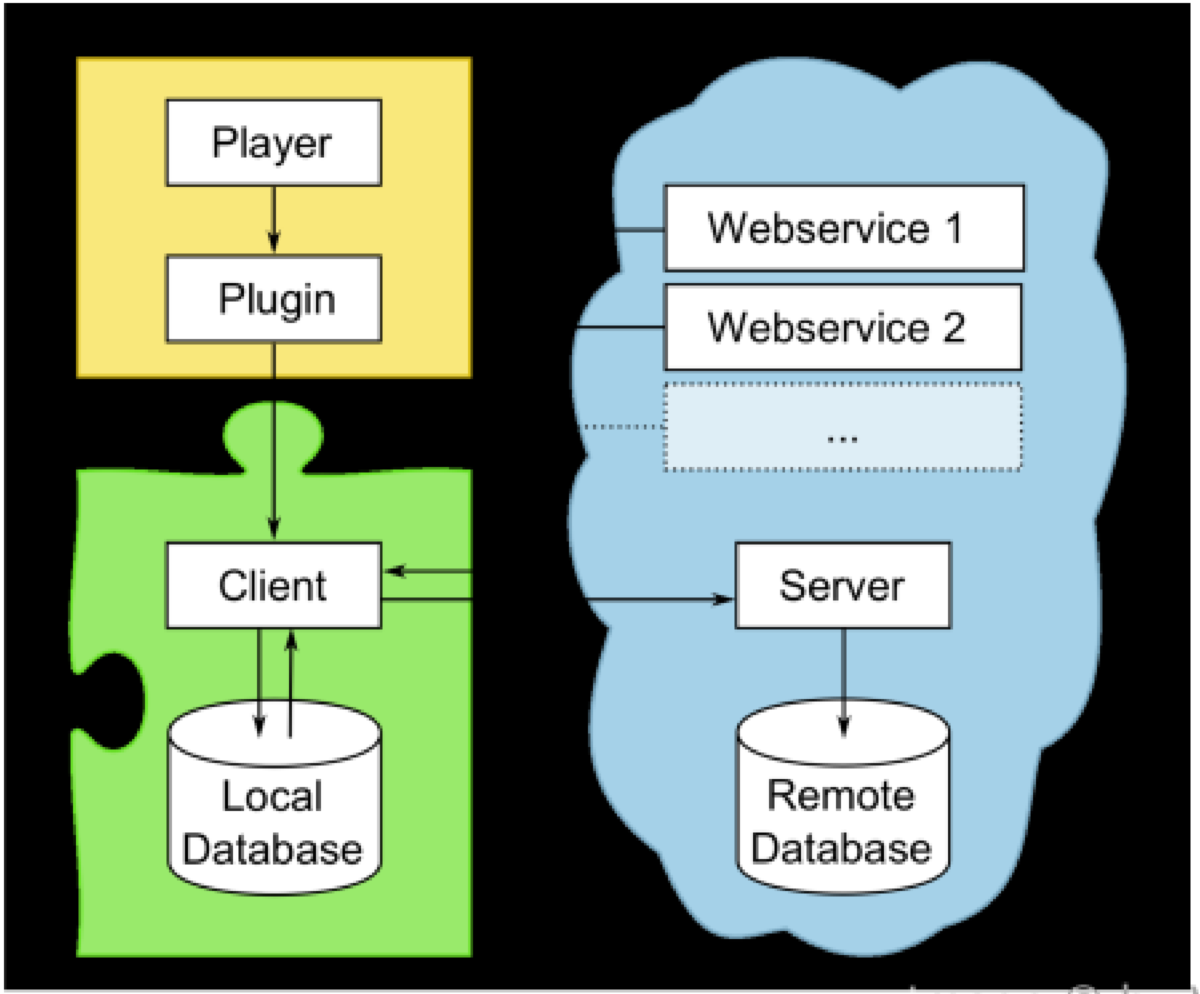

Facet-browser for context data

Listening Context Graphs (Concepts)

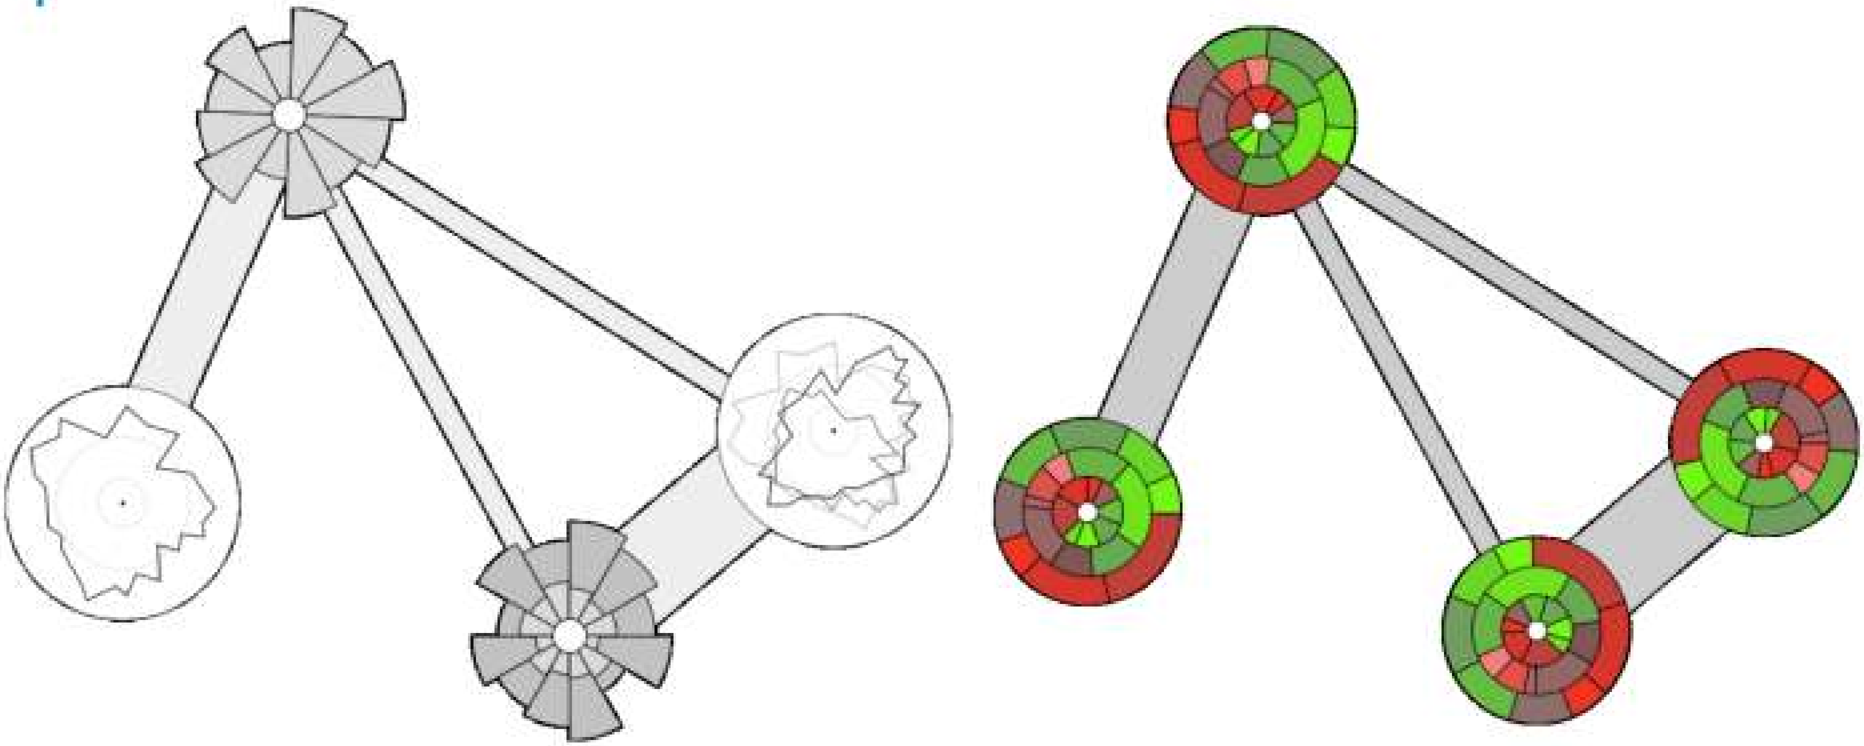

image © by Valentin Laube

Question

How much context is too much?

or

Where do we encounter privacy issues?

Survey Design & Context

- 8 questions
- 4 groups:
  1. demographic information
     - gender, age & country of residence
  2. general relation to music
  3. use of (web-) applications that collect, access and expose to some extend private data of their users
  4. acceptance of logging information about the listening context

- paper questionnaire at CeBIT 2009
  - in German only, 156 participants
- online survey from March till June 2009
  - in English and German, 305 participants

Acceptance of Listening Context Logging

- *Would you allow your music player (as software or as a self-contained device) to log the following information in order to enable it to learn personalized genres for sorting your music collection?*

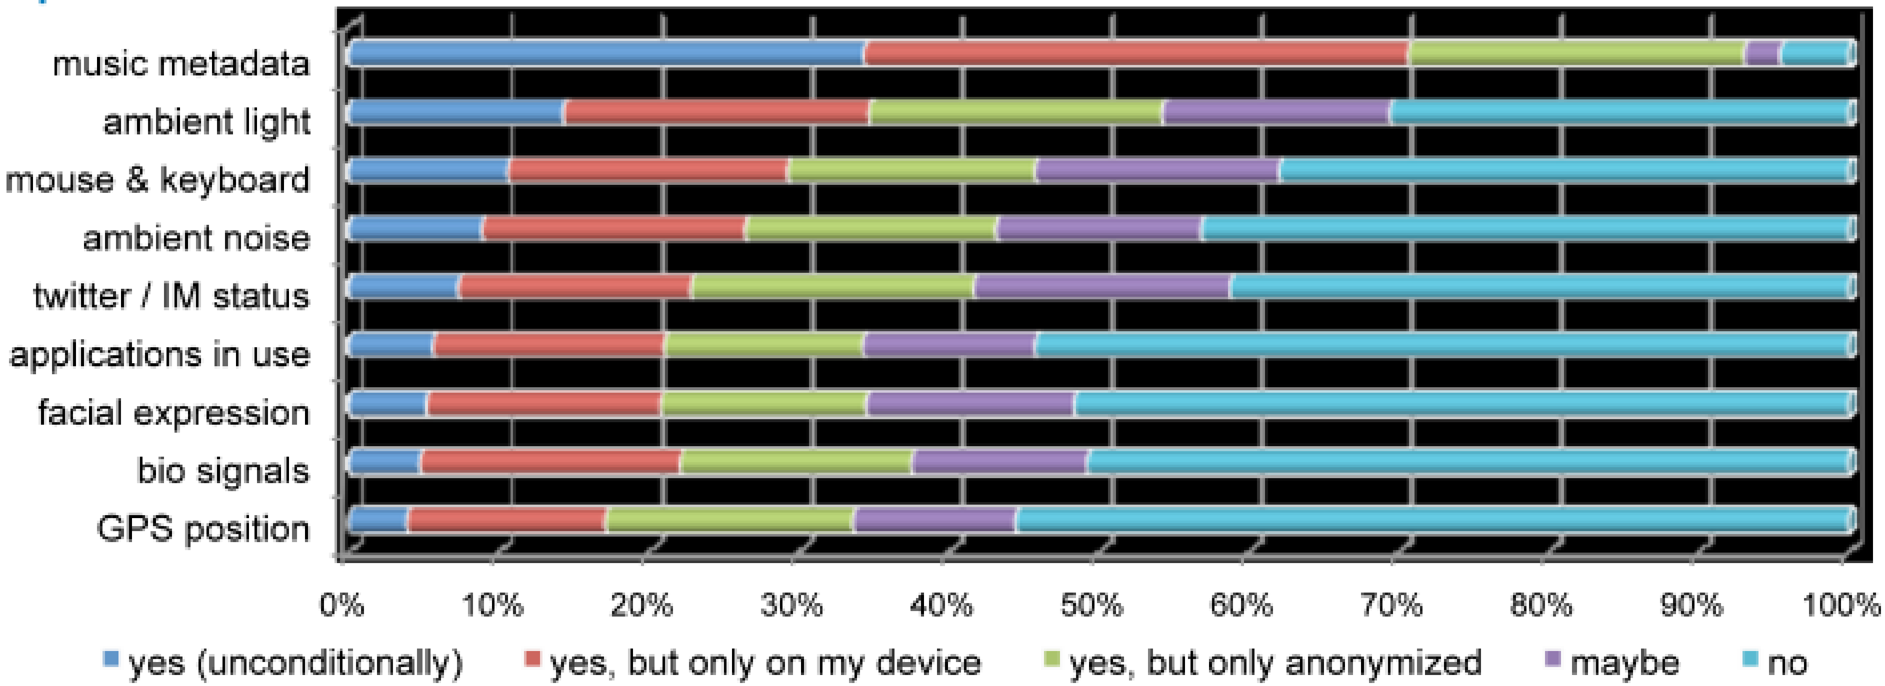

n=461

Reasons

- concerns about privacy
- against any kind of data collection
- no logging unless for visible benefit
  - doubts about relevance of collected information
- fear of being patronized by player's "intelligence"
- misuse of information for marketing purposes
- requirement of additional storage and processing power
  - increase of costs for hard- and software
- information leakage through hacking

Conclusions

- users must be:
  - fully informed about
    - when and
    - what information is logged,
    - where it is stored and
    - who has access to it
  - in full control of
    - the logging process
    - the adaptation of the MIR system

ADAPTIVE MUSIC SIMILARITY

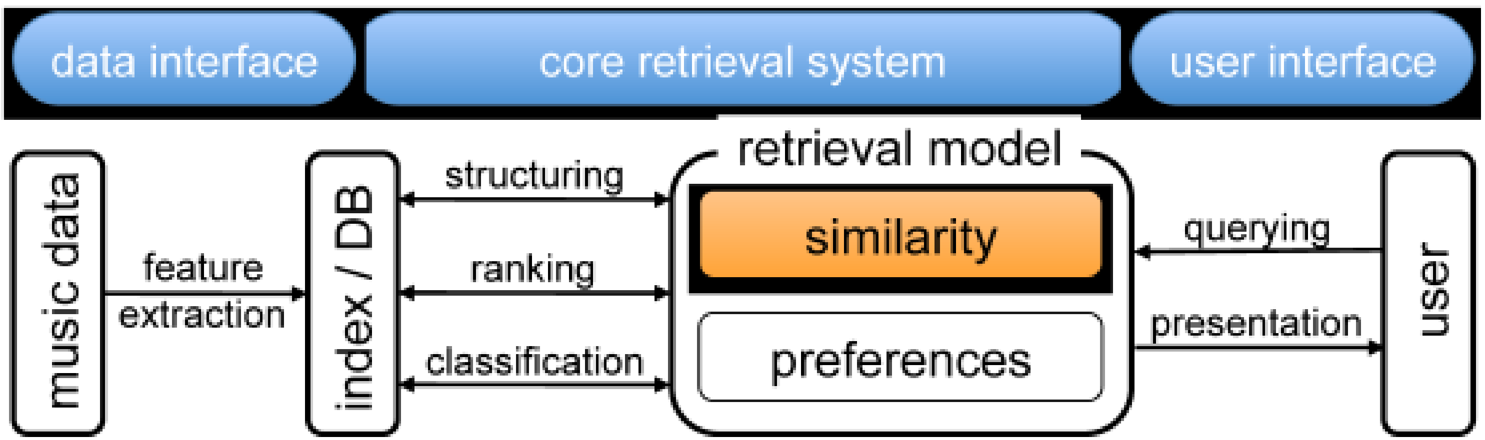

Listening Example

"His Master's Voice" (Francis Barraud)

- How should the software compare these songs?
  - melody, mood, timbre, lyrics, tempo, dynamics…
  - mode, instrumentation, key, harmonics, rhythm, meter …
  - ➡ music has many <u>facets</u> – How important is each one?

➡ How can the software learn how I compare songs?

Music Similarity Is Not A Static Concept!

- how music pieces are compared depends on context:
  - user
    - (musical) background/preferences
      - instruments played,
      - listening habits & history,
      - taste, …
  - retrieval task
    - e.g., cover song retrieval: harmonic progression/tonality and lyrics more important than timbre

Music Similarity

- key to many MIR applications & algorithms
  - query / ranking
  - recommendation (find more like…)
  - structuring (group similar)

- spectrum of music similarity specificity:

Music Signal Matching

Music Concept Matching

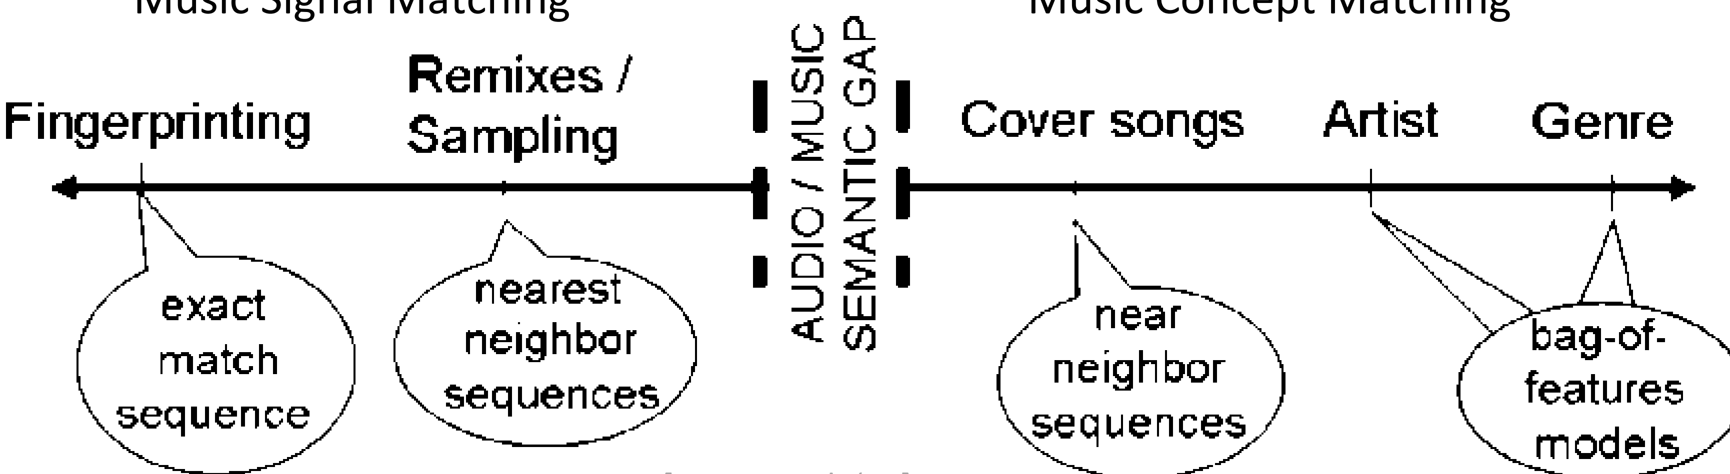

[Casey et al. '08]

Music Similarity Task/Specificity Spectrum

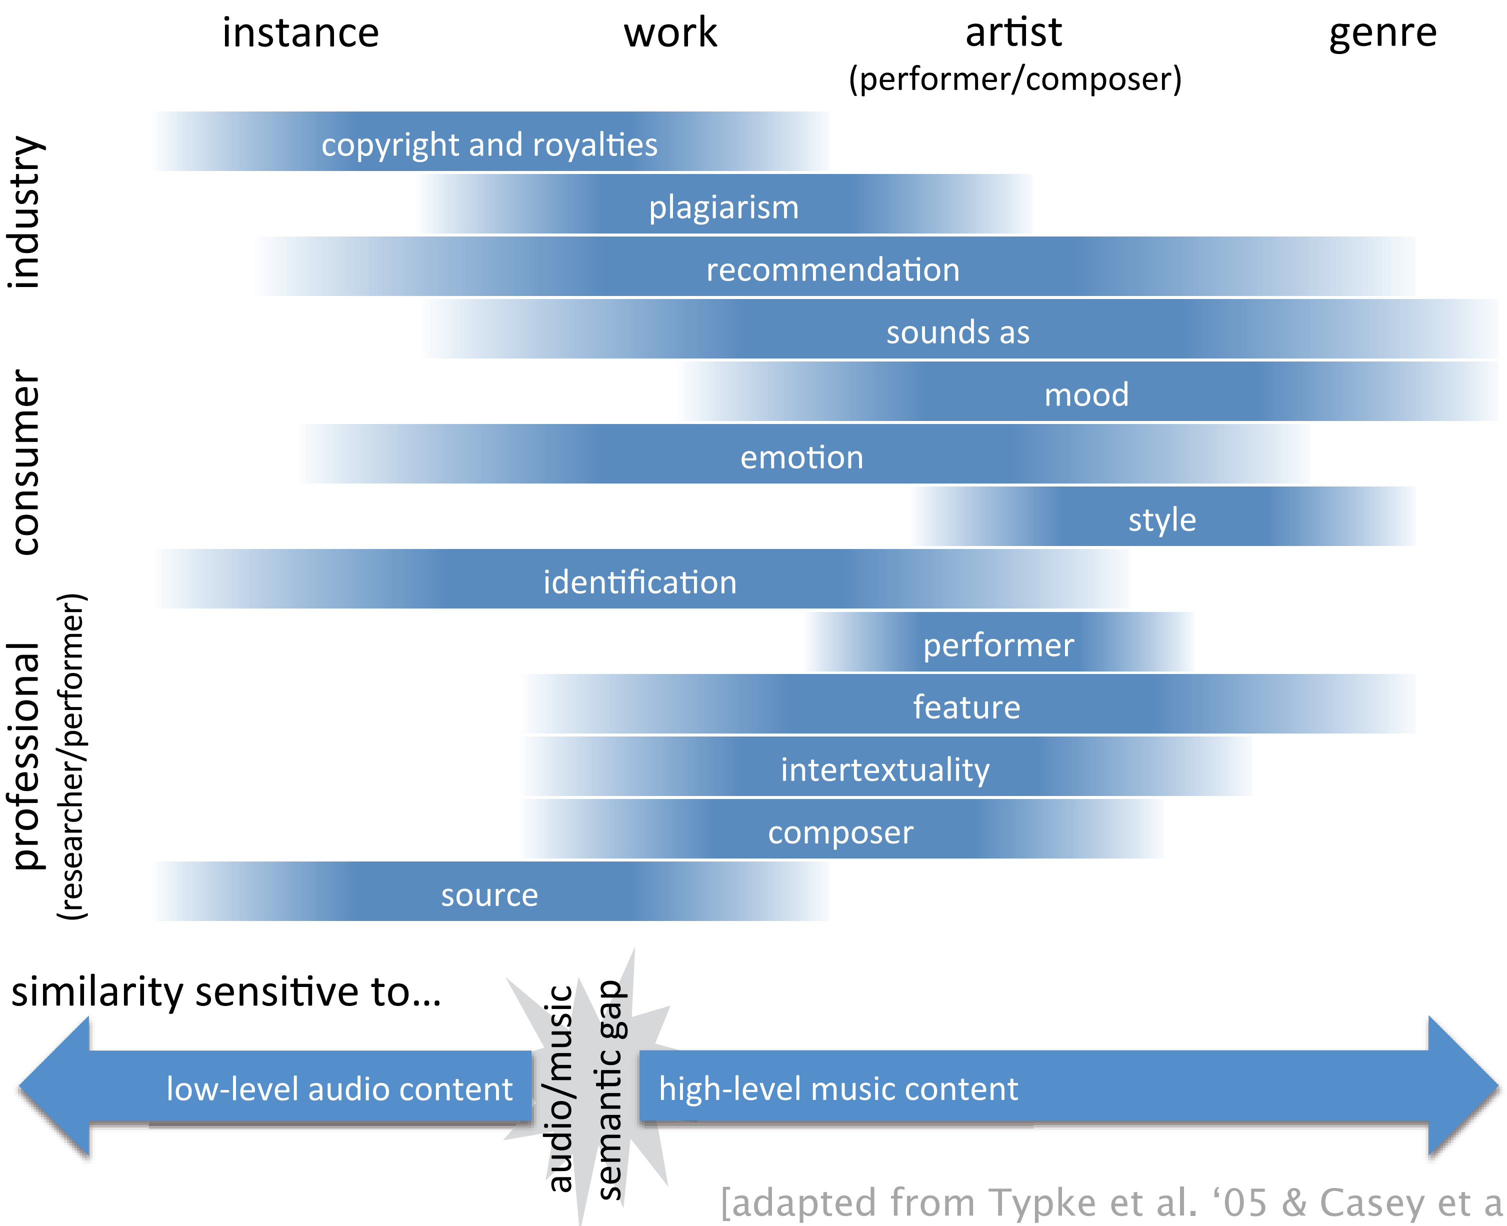

[adapted from Typke et al. '05 & Casey et al. '08]

Goal / Problem Formulation

Learn multi-facet music similarity measures
that reflect the user's information need and context!

Adaptable Model of Similarity

- objects of interest are described by various <u>features</u>
  - capture <u>different aspects</u> of similarity
  - may not be equally important for comparison

- distance facet
  = (set of) feature(s)
  \+ distance measure ←objective
    - non-negative: $d(a,b) \geq 0$ [and $d(a,b) = 0$ iff $a=b$]
    - symmetric: $d(a,b) = d(b,a)$
    - optionally: fulfills triangle inequality

➡ distance = weighted linear sum of facet distances
  - weights non-negative, constant weight sum
  - direct (manual) adaptation possible (simple & understandable model) ←subjective

Adaptable Systems

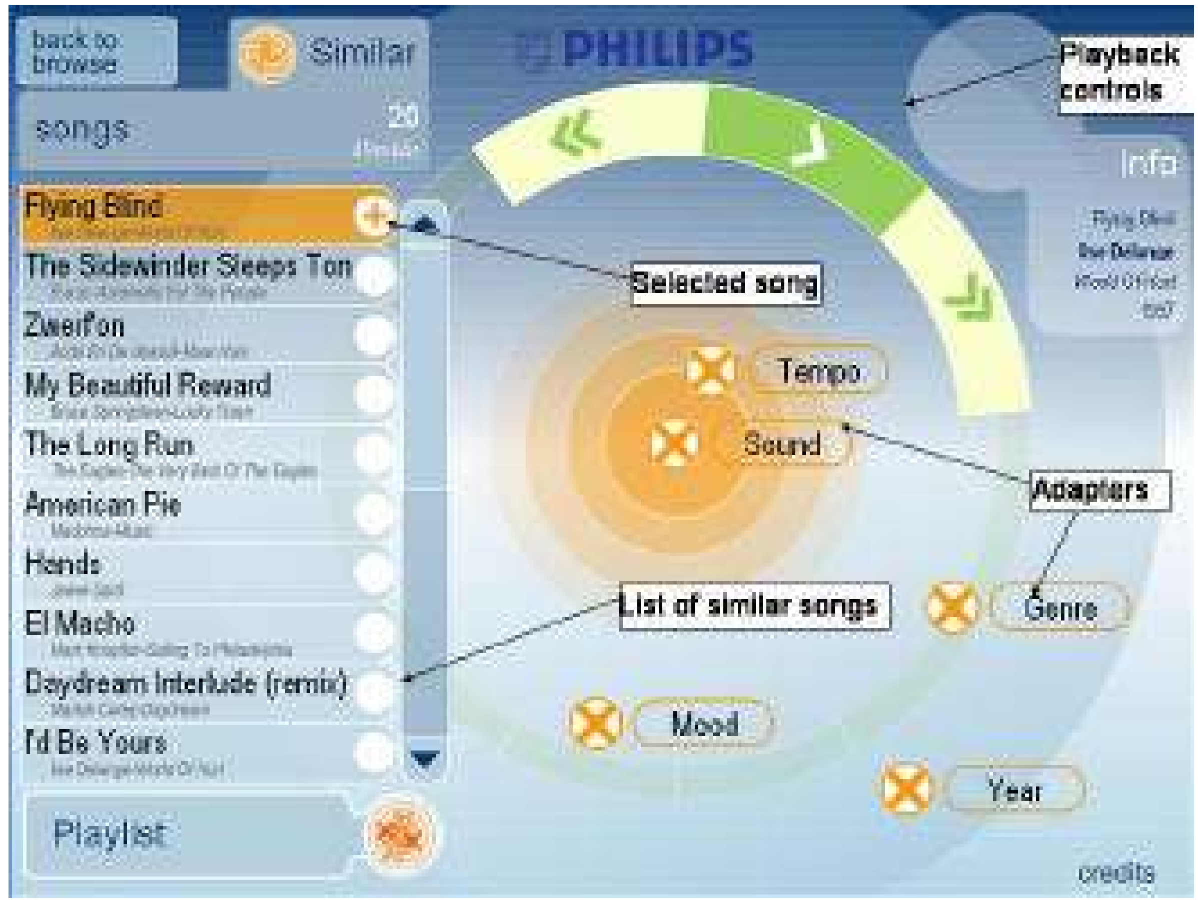

[Baumann & Halloran 2004, An ecological approach to multimodal subjective music similarity perception, CIM]

[Vignoli & Pauws 2005, A music retrieval system based on user driven similarity and its evaluation, ISMIR]

System Design

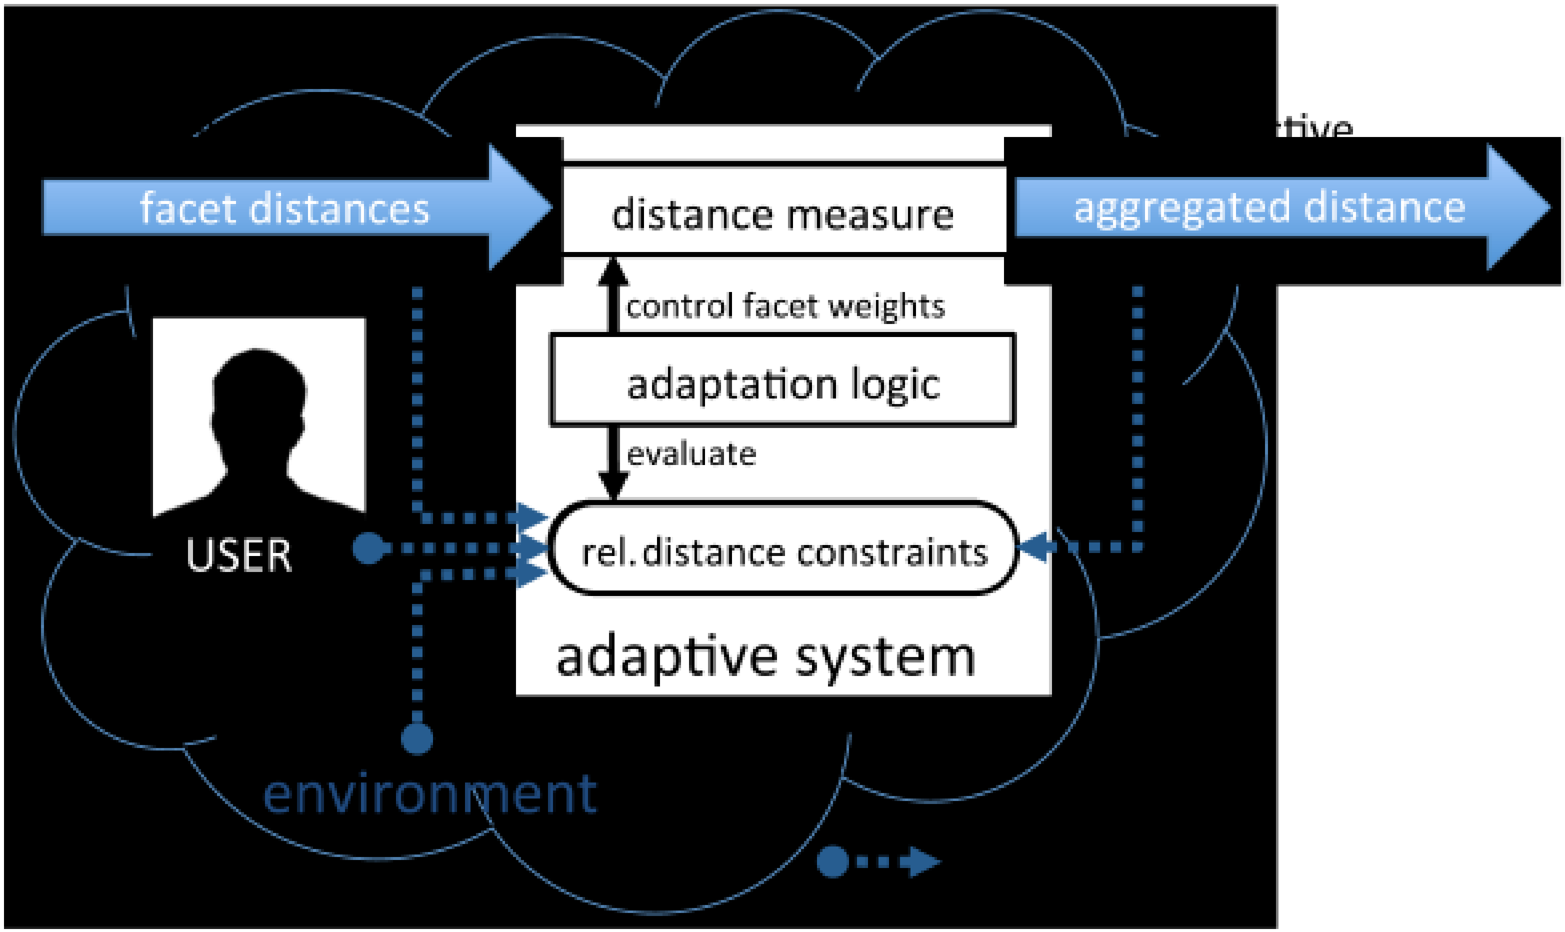

General Adaptation Approach

user actions

derive

relative distance constraints

as constraint

training examples

optimization problem:

- find (valid) weights that
  a) satisfy all constraints
  b) minimize error (#violations)

classification problem:

- learn <u>linear</u> classifier for
  \+ training examples *d(s,a)* < *d(s,b)*
  – training examples *d(s,a)* > *d(s,b)*
- weights are defined by the separating hyperplane
  [Cheng et al. '08]

Metric Learning as Classification Problem

relative distance constraints ⇨ linear classification problem

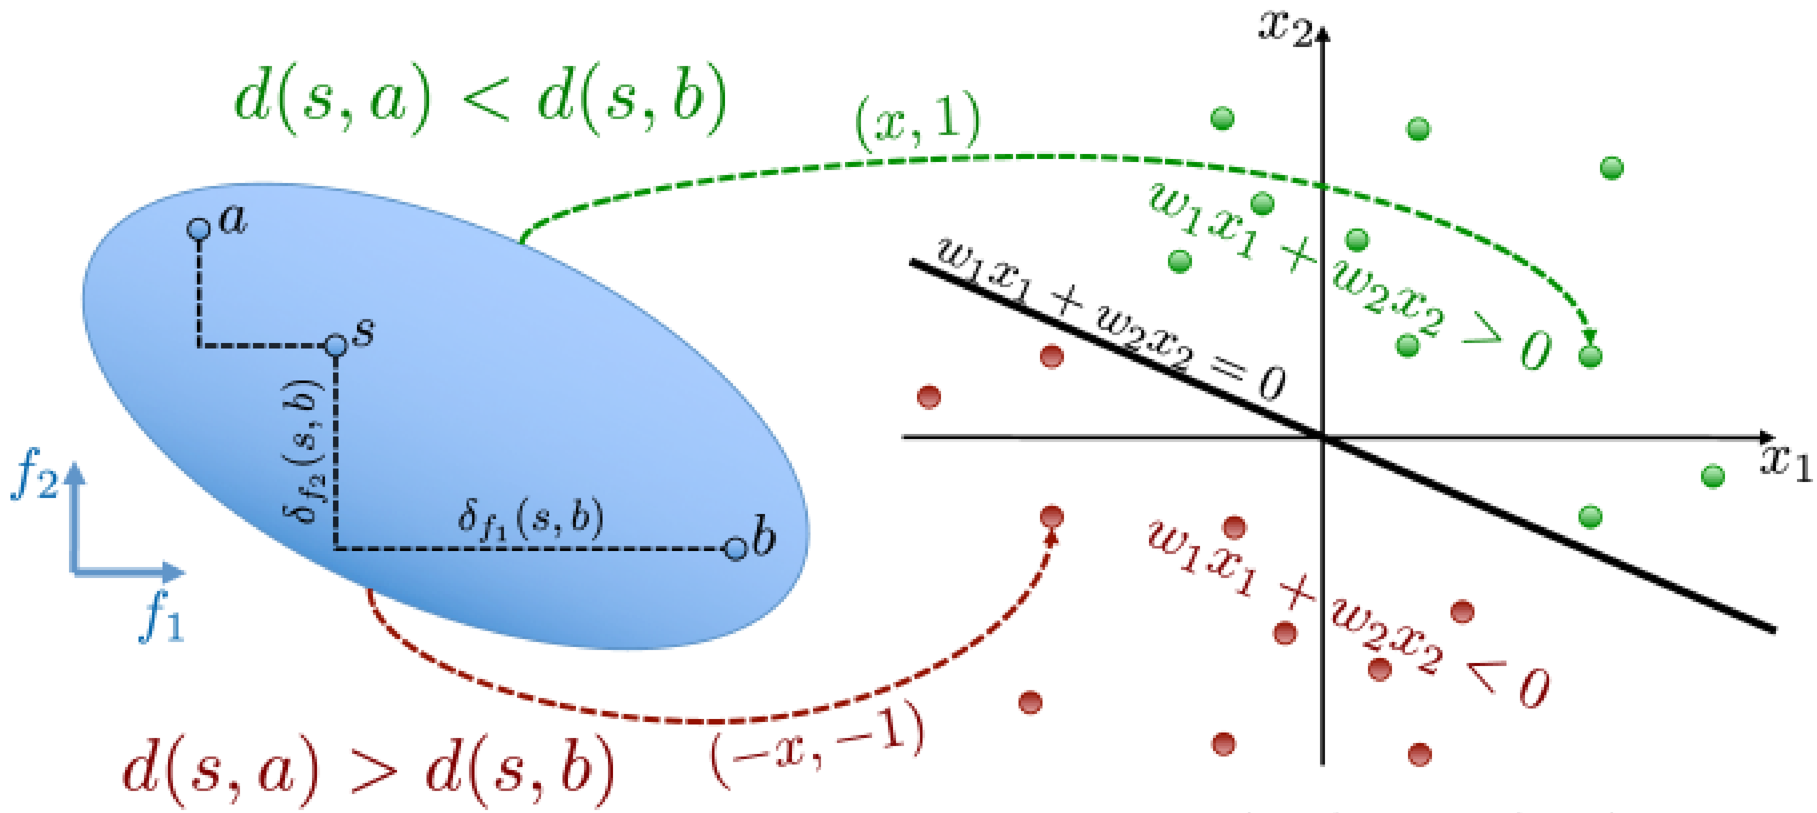

$$x_1 = \delta_{f_1}(s,b) - \delta_{f_1}(s,a)$$
$$x_2 = \delta_{f_2}(s,b) - \delta_{f_2}(s,a)$$

[Cheng et al. '08]

Facet Weight Adaptation Approaches

- Gradient Descent (optimization)
  - directly minimizes error (constraint violations)
  - problem: may get stuck in local minimum
- Quadratic Programming (optimization)
  - minimizes weight change subject to
    - hard weight bounds and
    - hard or soft distance constraints (additional slack variables)
  - continuity (no abrupt changes)
- Linear Support Vector Machine (classification)
  - maximizes margin (between + and – training examples)
  - favors "stable" solutions
  - problem: <u>soft</u> weight constraints may be violated (neg. weights)

Applications & Considered User Actions

- Liederenbank [ISMIR'09]
  - classifying Dutch folk songs
  - ➠ class annotations (by experts)

- BeatlesExplorer [AMR'08]
  - structuring the Beatles dataset
  - ➠ moving songs to other cells
  - ➠ correcting similarity rankings

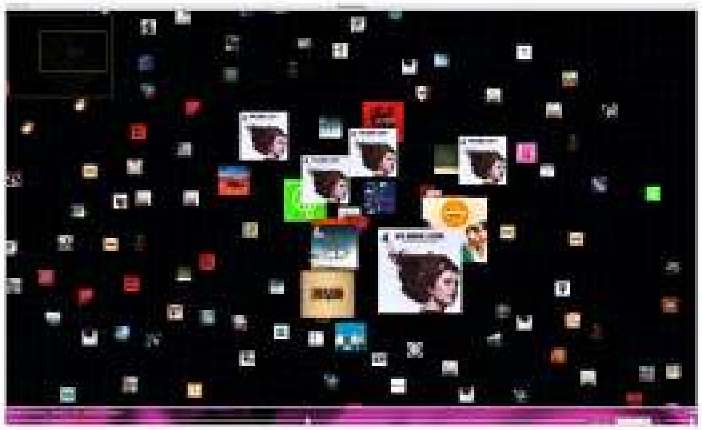

- MusicGalaxy [CMMR/SMC'10]
  - exploring media collections
  - ➠ tagging objects

More Adaptive Applications:

soniXplorer [Lübbers & Jarke 2009, Adaptive multimodal exploration of music collections, ISMIR]

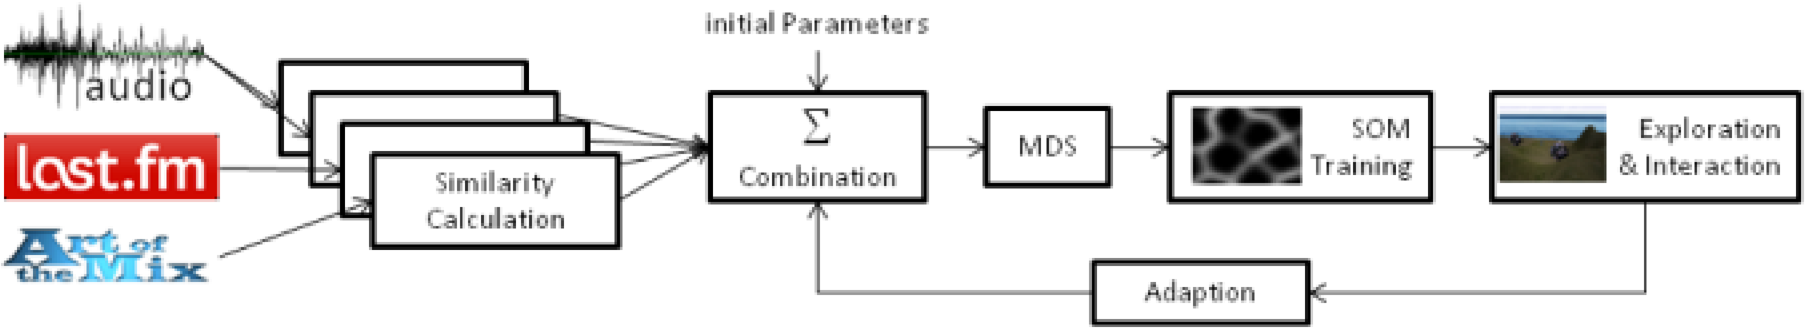

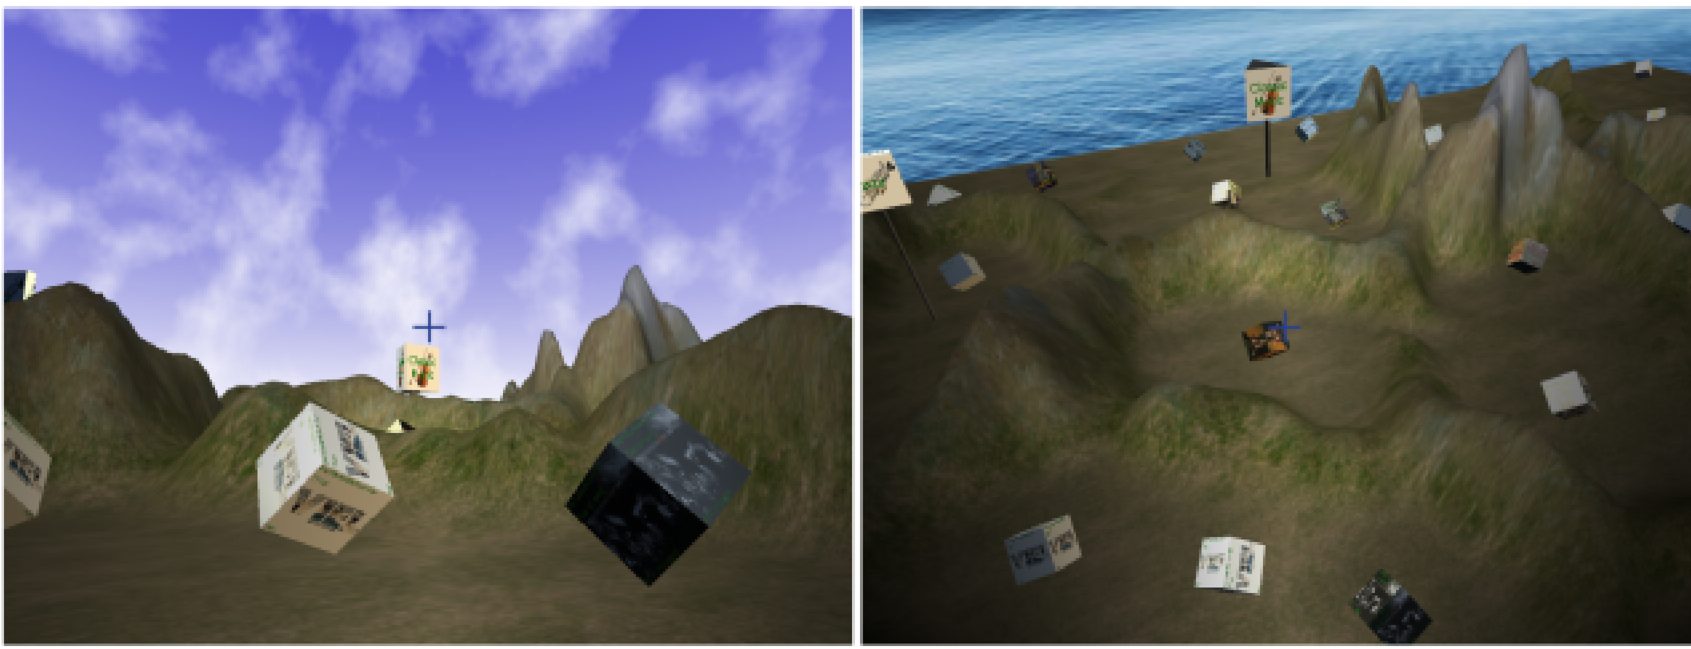

- users can change the terrain -> directly derive distance matrix

Focus-Adaptive Visualization

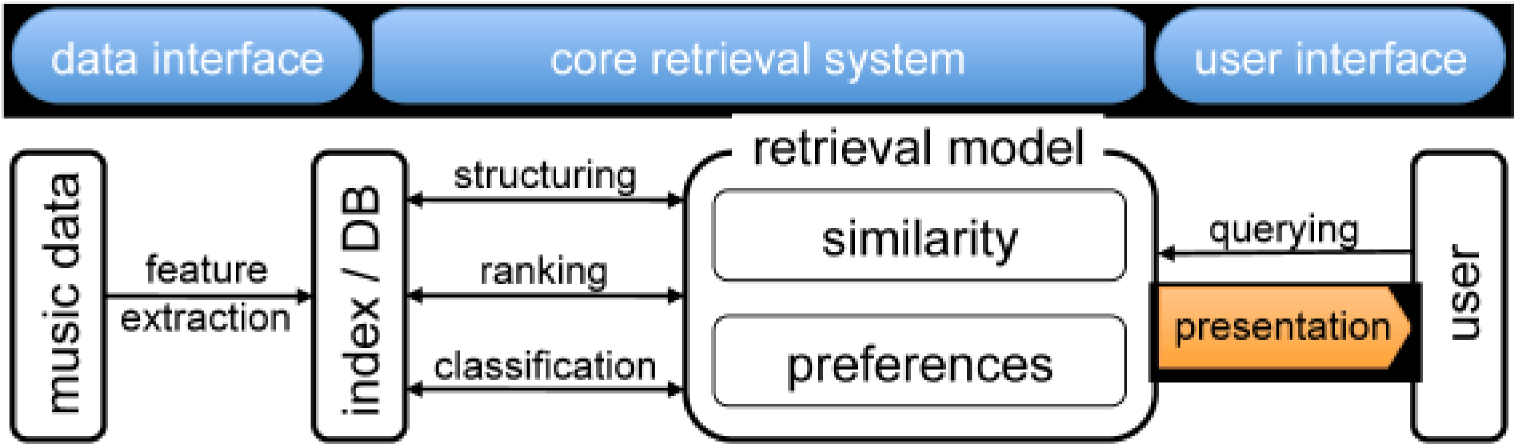

Motivation

- generate an overview of a music collection for exploration
- idea: use dimensionality reduction techniques

high-dimensional feature space

2D display

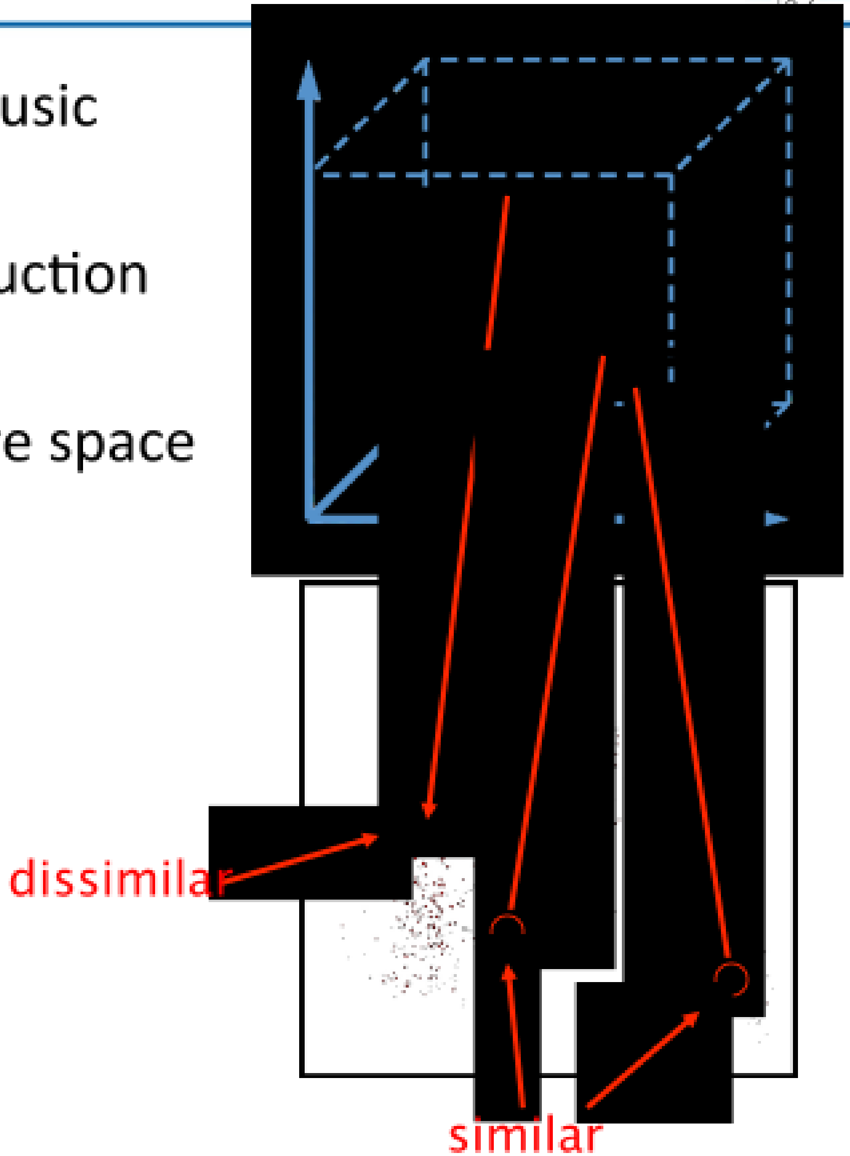

“Projection Errors” - They are everywhere!

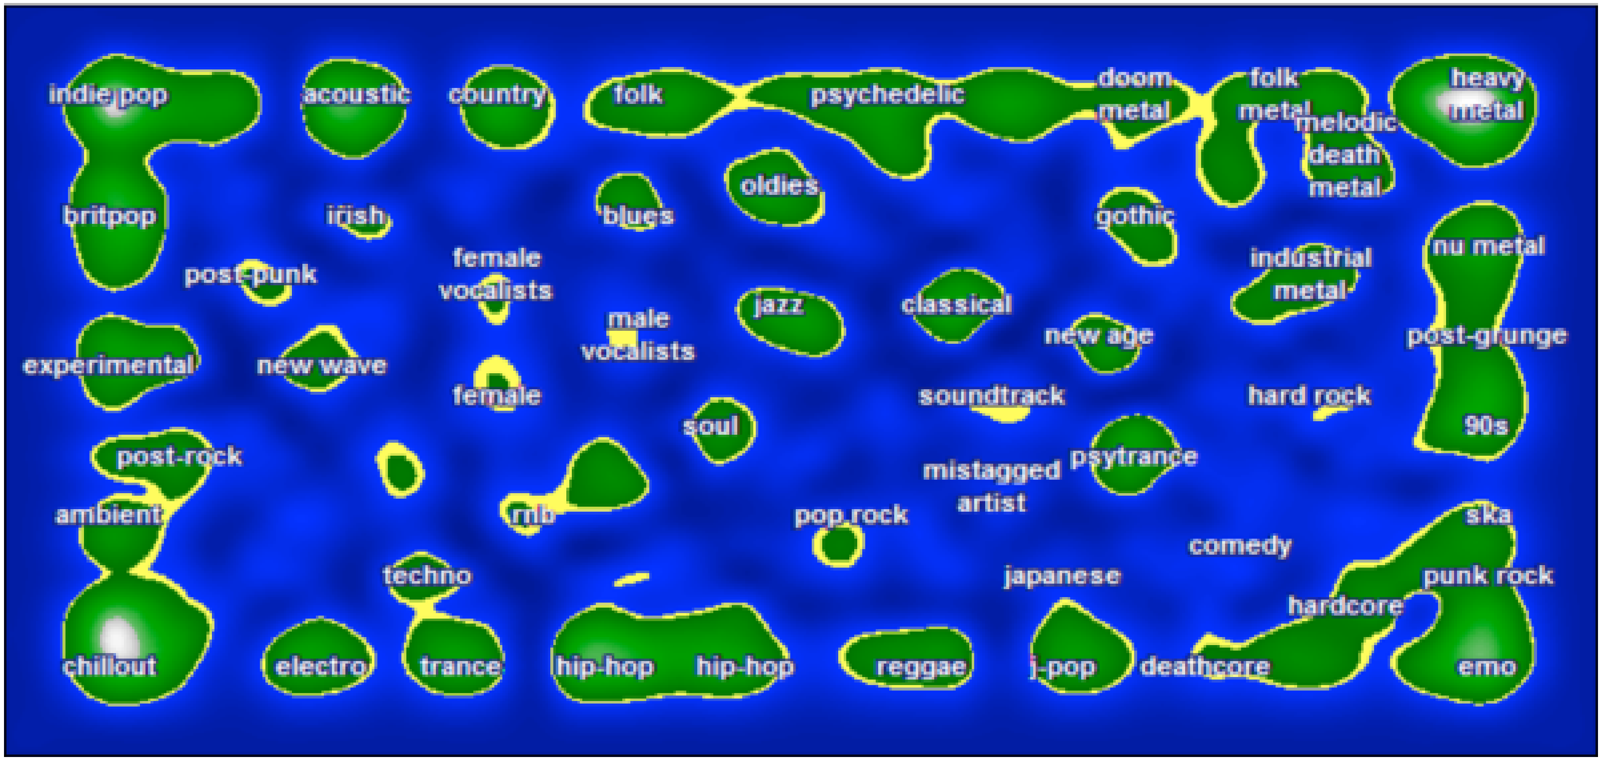

Islands of Music [Pampalk 2003]

"Note that it is not always the case that the first and second best matching units lay next to each other. (...) However, it is rather unlikely that the two units are separated completely on the map and mostly they will both be located in the same map area.“

“Projection Errors” - They are everywhere!

Islands
[Pamp

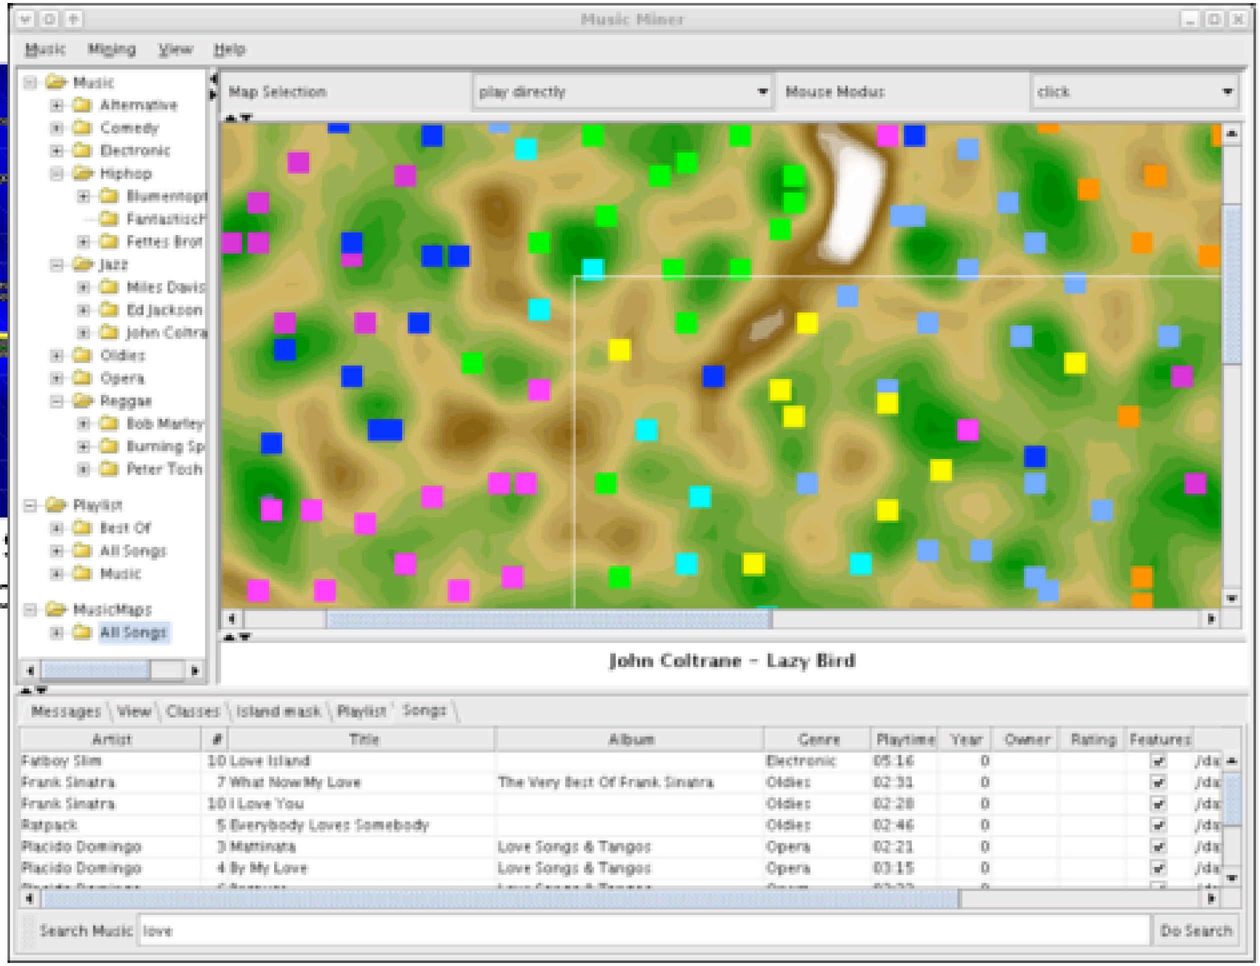

MusicMiner [Mörchen et al. 2005]

mountain ranges between dissimilar neighbors

“Projection Errors” - They are everywhere!

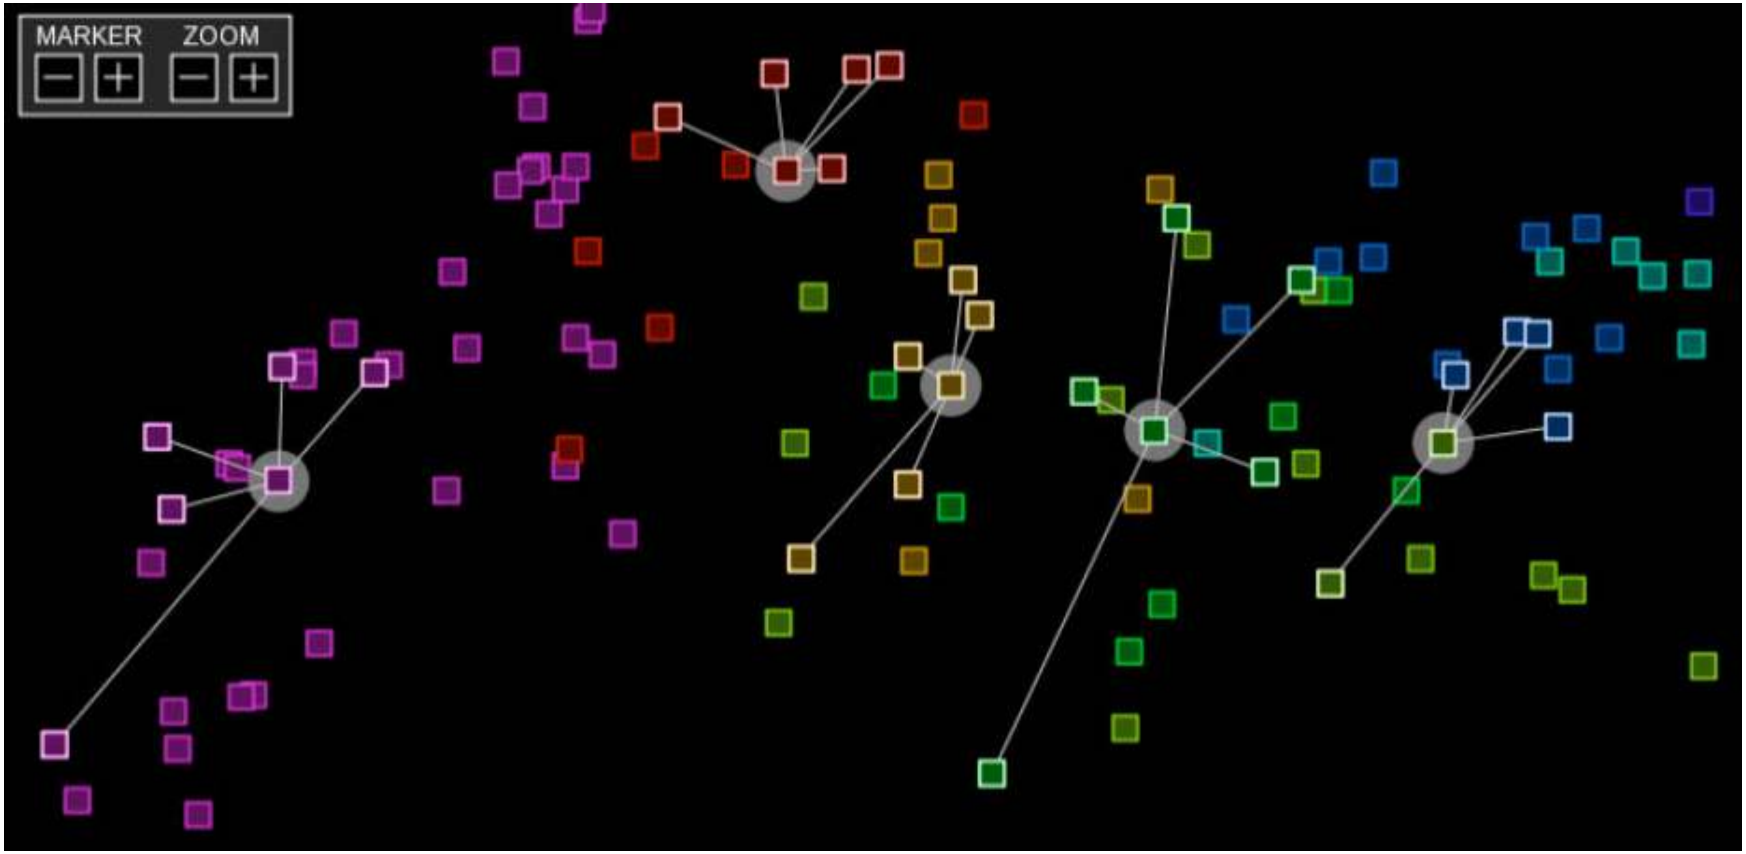

SoundBite for Songbird [Lloyd 2009]

edge connections to nearest neighbors

“Projection Errors” - They are everywhere!

Islands of Music
[Pampalk 2003]

MusicMiner
[Mörchen et al. 2005]

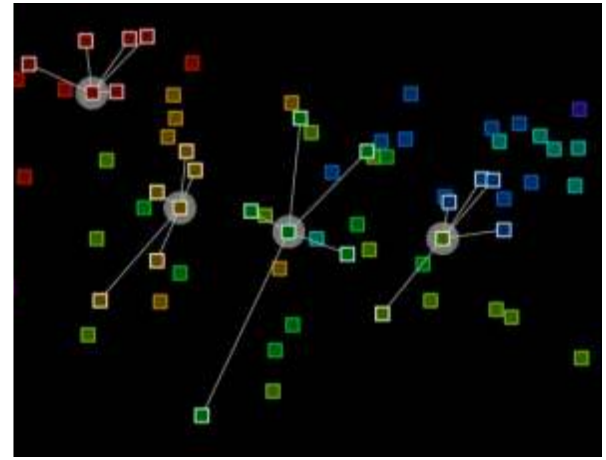

SoundBite for Songbird
[Lloyd 2009]

- Can’t we fix these “projection errors”?
  - No, we can’t (without causing damage elsewhere)
  - … but we can exploit them for exploration…

…using the wormhole metaphor

Focus-Adaptive SpringLens

<u>temporarily</u> fix / highlight

the neighborhood in <u>focus</u>

Focus-Adaptive SpringLens*

- multi-focus fish-eye distortion highlights nearest neighbors

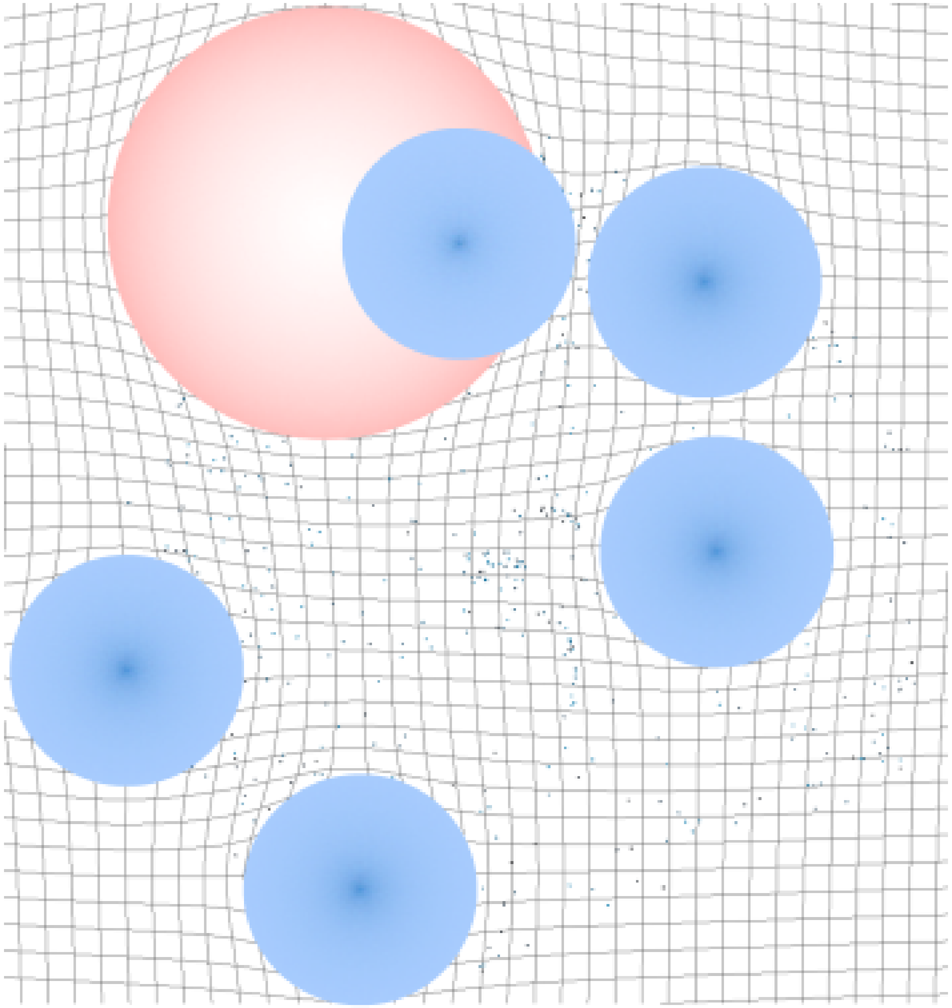

- primary lens
  - controlled by user
  - enlarges region of interest
  - more space for details
  - preserves context
- secondary lenses
  - data-driven
  - highlight nearest neighbors
  - show “wormholes”
  - neighbors come closer

*based on SpringLens non-linear distortion technique [Germer et al. ‘06]

System Design

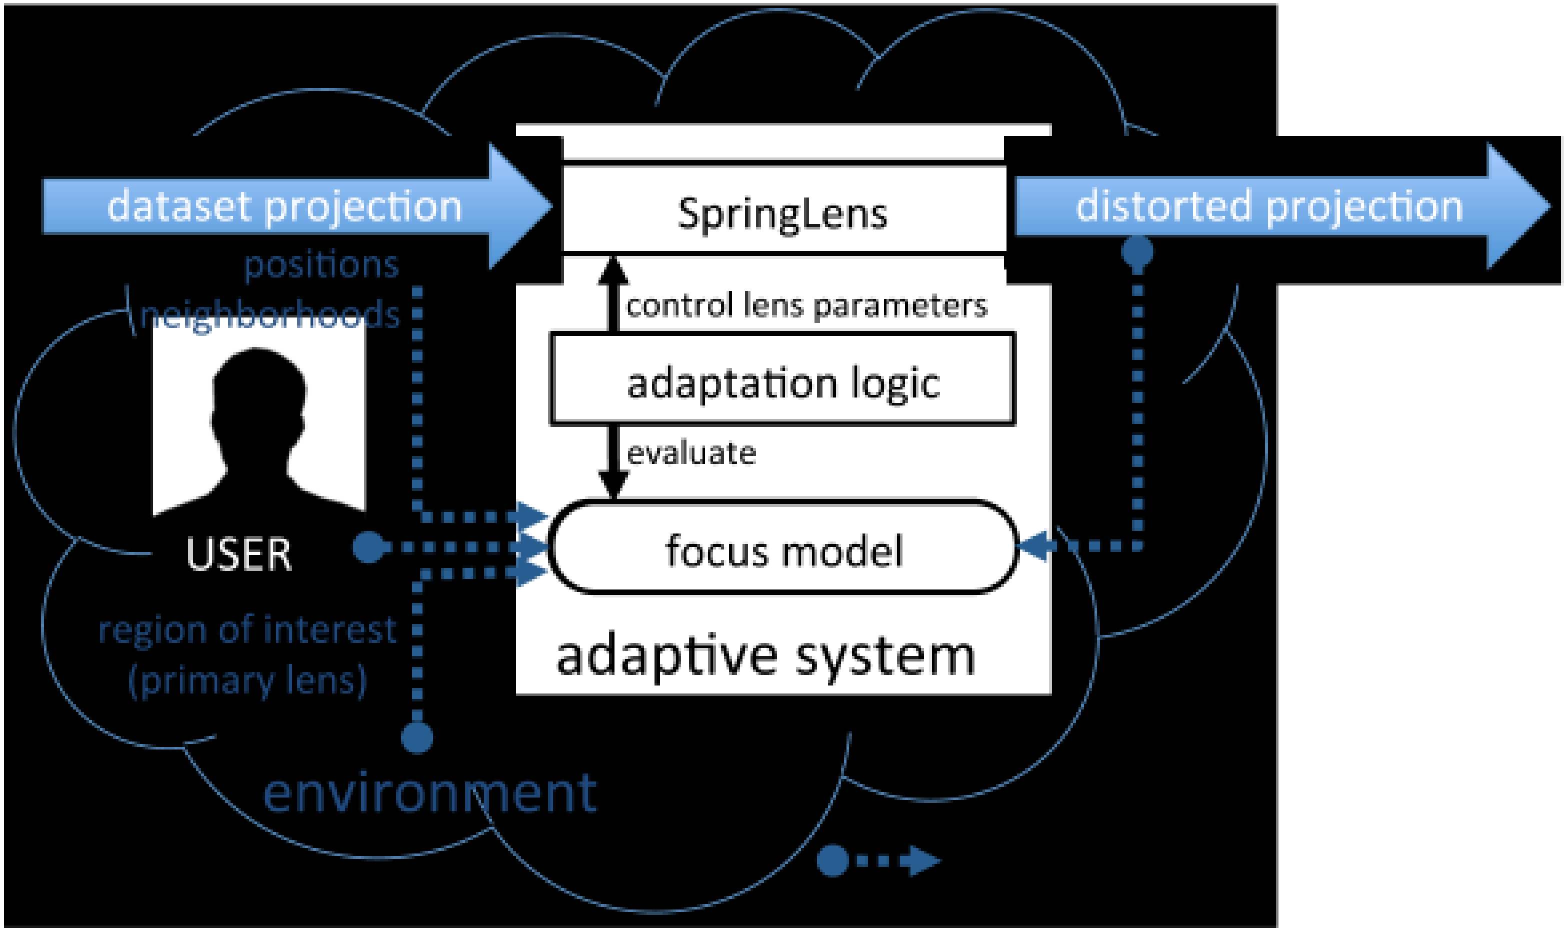

MusicGalaxy Prototype (Demo)

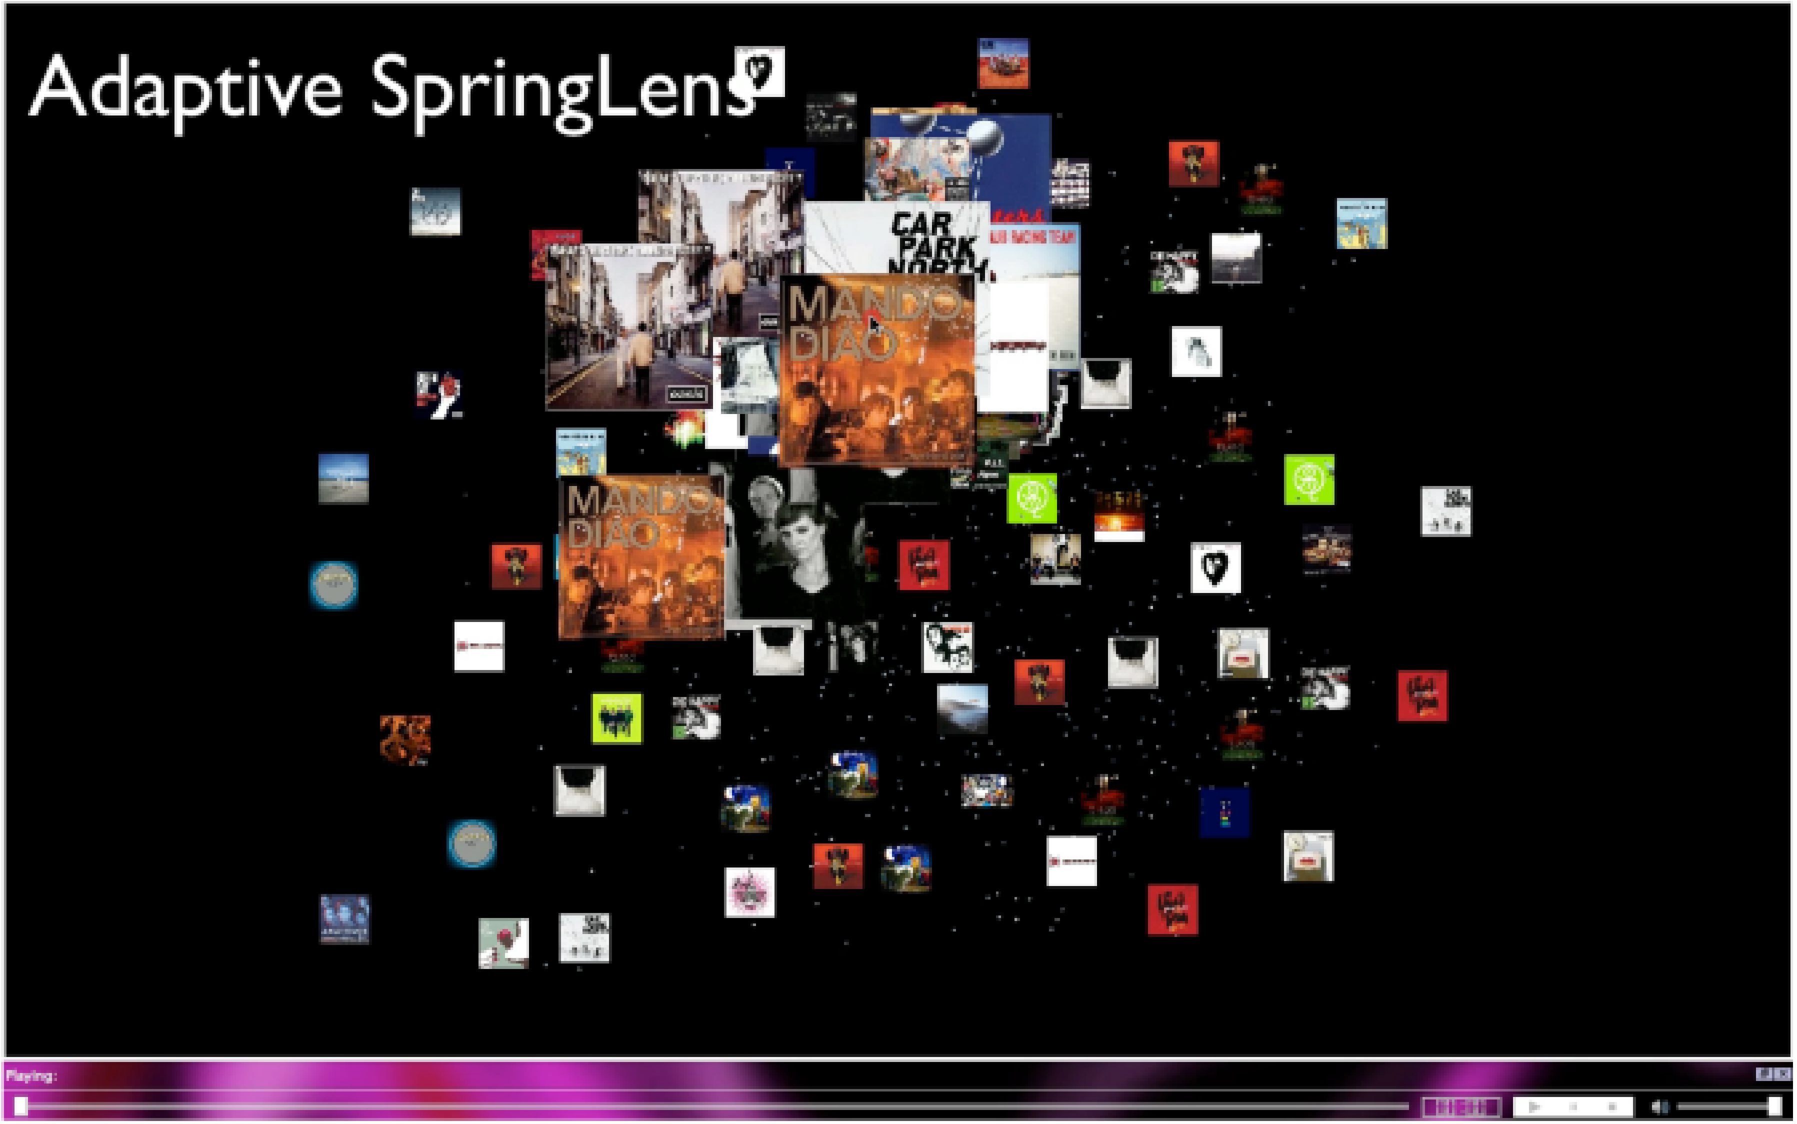

demo video at http://www.dke-research.de/aucoma

Gaze-Supported Interaction

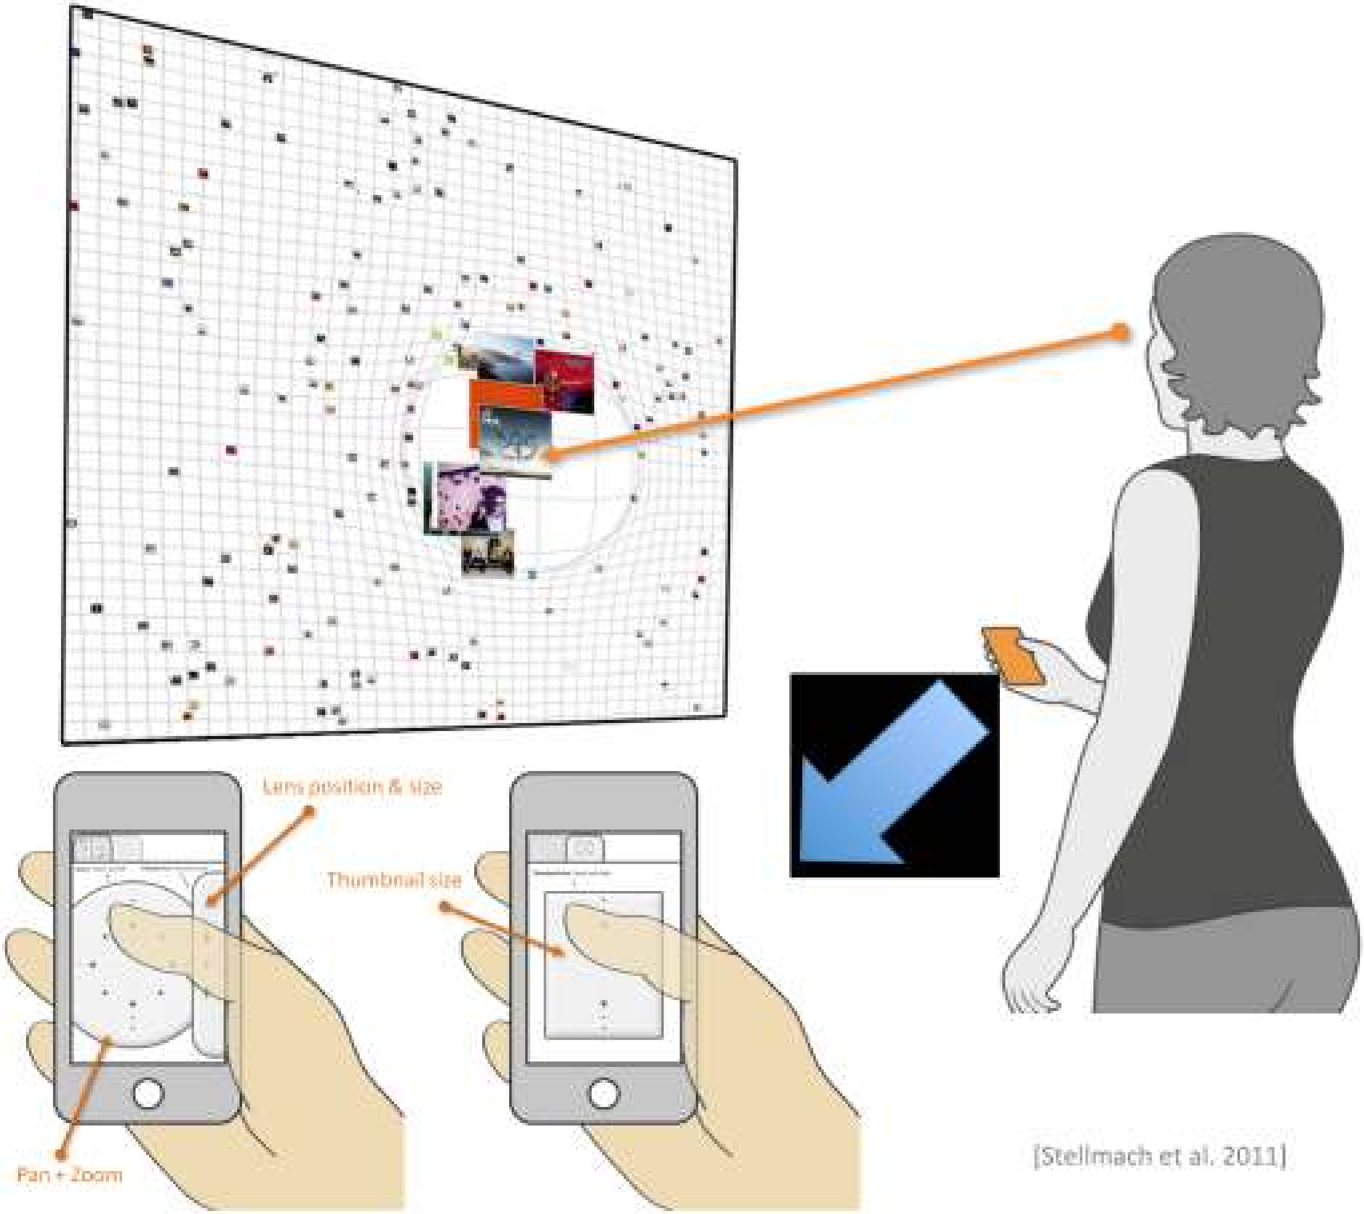

[Stellmach et al. 2011]

Complex Distortions

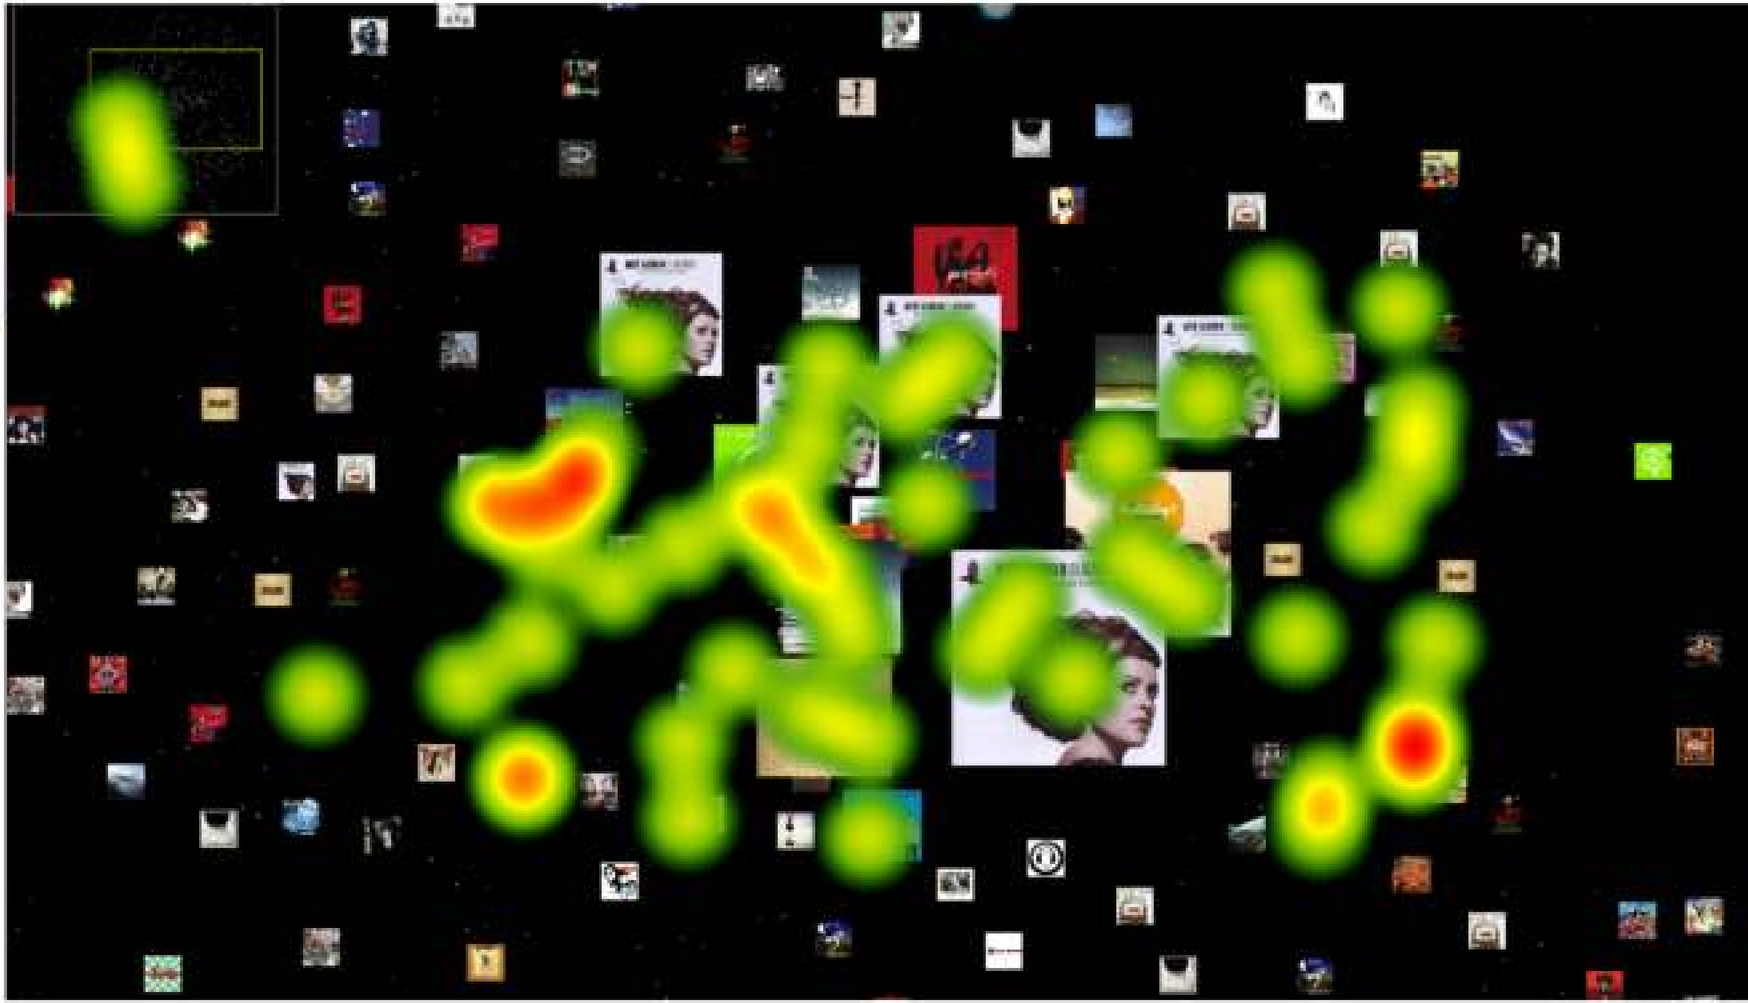

(gaze-based attention map)

WRAP-UP

Wrap-up

- there are many ways to enhance MIR systems through adaptive approaches

- improvements may come at a “cost”
  - system complexity
  - training effort
  - user acceptance

- important lesson learned:
  - let the user stay in control

    -> make controls for adaptable system accessible

    -> retain possibility to override or switch of adaptations
  - point out what is adapted and on which basis

Outlook

- tomorrow:
  - more on structuring
- Thursday:
  - more on similarity
- Friday:
  - more on interfaces and evaluation

THANK YOU FOR YOUR ATTENTION!

Further Background Reading

Stober, S. & Nürnberger, A.: "Adaptive Music Retrieval - A State of the Art." In: Multimedia Tools and Applications, August 2013, Volume 65, Issue 3, pp 467-49.

… or chapter 3 of my thesis available here: http://www.dke-research.de/aucona/thesis

Table 4: Overview of approaches sorted by adaptation technique.

technique (adaptation logic)	application area (section)	references	source for adaptation (context information)	target of adaptation	guidance of adaptation
BASIC DATA-DRIVEN TECHNIQUES USING DYNAMIC PARAMETERS					
dyn. thresholding	onset detection (3.3.1.1)	[39, 203, 252]	data distribution	threshold	internal objective function
	audio identification (3.3.1.1)	[16]			
	audio-score matching (3.3.1.1)	[162]			
	segmentation (3.3.1.1)	[144, 232]			
adaptive window size	tempo estimation (3.3.1.3)	[164]	inter-onset-intervals	window size	internal objective function
fuzzy freq. weighting	key recognition (3.3.1.4)	[40]	signal density distribution	pitch range weights	internal objective function
PROBABILISTIC (CLASSIFICATION) TECHNIQUES					
HMM / Viterbi	chord recognition (3.3.1.2)	[14, 186, 219]	data history	path - state	internal objective function
	key recognition (3.3.1.2)	[121, 170]			
	F0-estimation (3.3.1.7)	[69]			
MAP (EM)	chord recognition (3.3.1.2)	[14]	data distribution	model parameters	internal objective function
	optical music recognition (3.3.1.6)	[200]			explicit user feedback (corrections)
	F0-estimation (3.3.1.7)	[77]			internal objective function
naïve Bayes	rec. by listening-context (3.3.3.2)	[83]	location, time, people in the room, weather, stock portfolio	user mood	explicit user feedback
(fuzzy) Bayesian net	rec. by listening-context (3.3.3.2)	[187]	weather, time, ambient noise, illuminance	context (state)	explicit user preference
hypothesis search	chord recognition (3.3.1.2)	[233]	estimated bass pitch	chord probabilities	internal objective function
ARTIFICIAL NEURAL NETWORK TECHNIQUES					
RBFN	ranking (3.3.5.1)	[229]	track ratings	feature selection	explicit user feedback
SOM	collection structuring (3.3.4)	[54, 93, 105, 108, 123, 135, 163, 165, 180, 185, 230]	data	weights & thresholds	internal (amplify response w.r.t. input)
GHSOM	collection structuring (3.3.4)	[58, 178, 202]	data	structure (hierarchy)	internal objective function

Overview of Adaptive MIR Approaches (cont.)

technique (adaptation logic)	application area (section)	references	source for adaptation (context information)	target of adaptation	guidance of adaptation
OPTIMIZATION TECHNIQUES					
genetic algorithms	QBSH (3.3.1.8)	[129]	singer errors	singer model	explicit user feedback
simulated annealing	playlist generation (3.3.3)	[191, 192]	n/a	n/a	explicit constraints (adaptable!)
linear regression	collection structuring (3.3.5.2)	[135]	user interaction	similarity metric	implicit (derived distance matrix)
weight reinforcement	ranking (3.3.5.1)	[208]	alignment (pairings)	pairing type weights	explicit user feedback
whitening, LDA, RCA, NCA, LMNN	ranking (3.3.5.1)	[227]	artist/album/blog labels	similarity metric	explicit constraints
multi-kernel POE	ranking (3.3.5.1)	[152]	subjective judgments	similarity metric	explicit constraints
structural SVM	ranking (3.3.5.1)	[151]	collaborative filtering information	similarity metric	derived constraints
STRATEGY SELECTION TECHNIQUES					
rule-based	pitch detection (3.3.1.5)	[133]	harmonic descriptor (simple/polyphonic)	pitch estimator selection	internal objective function
(OTHER) CLASSIFICATION TECHNIQUES					
Rocchio	preference modeling (3.3.2)	[94]	track ratings	reference vectors	explicit user feedback
SVM	preference modeling (3.3.2)	[140]	track ratings	decision boundary	explicit user feedback
PA-L-EX	preference modeling (3.3.2)	[161]	track ratings	kernel	explicit user feedback
decision tree induction	playlist generation (3.3.3)	[190]	track ratings	decision tree	explicit user feedback
classifier fusion	playlist generation (3.3.3)	[248]	track ratings	fusion parameters	explicit user feedback
case-based reasoning	rec. by listening-context (3.3.3.2)	[118]	season, month, day of the week, weather and temperature	case selection	implicit (listening behavior)
heuristics	playlist generation (3.3.3)	[184, 199]	track skips	song selection/removal	implicit (skipping behavior)
	playlist generation (3.3.3)	[181]	track ratings	radio station profile	explicit user feedback

Some References

Broy, M. et al.: "Formalizing the notion of adaptive system behavior." In: *Proc. of the 2009 ACM Symposium on Applied Computing (SAC'09).* 2009, pp. 1029–1033

Germer, T. et al.: "SpringLens: Distributed Nonlinear Magnifications." In: *Eurographics 2006 - Short Papers.* 2006, pp. 123–126

Lloyd, S.: "Automatic Playlist Generation and Music Library Visualisation with Timbral Similarity Measures." MA thesis. Queen Mary University of London

Mauch, M. and Dixon, S.: "Approximate note transcription for the improved identification of difficult chords." In: ISMIR 2010, pages 135–140, 2010.

Mörchen, F. et al.: "Databionic Visualization of Music Collections According to Perceptual Distance." In: Proc. of ISMIR'05. 2005, pp. 396–403

Pampalk, E. et al.: "Exploring music collections by browsing different views." In: Proc. of ISMIR'03. 2003, pp. 201–208

Stellmach, S. et al.: "Designing Gaze-supported Multimodal Interactions for the Exploration of Large Image Collections." In: *Proc. of 1st Int. Conf. on Novel Gaze-Controlled Applications (NGCA'11).* 2011, pp. 1–8

(for more references, please refer to the MTAP paper)