

Tutorial presented to
ISMIR 2006: Victoria BC
October 8, 2006

User Interfaces for Music Information Retrieval

David Gerhard
Department of Computer Science
Department of Music
University of Regina
david.gerhard@uregina.ca

Overview

- Introduction
- Motivation
- Principles of User Interface Design
- MIR interfaces
 - Activities
 - Presenting information
 - Designing interaction
 - User goals and activities
- Evaluating interactive MIR

ISMIR 2006: Victoria, BC
October 8, 2006

2

Introduction

- David Gerhard
 - david.gerhard@uregina.ca
 - www2.cs.uregina.ca/~gerhard
- aRMADILo lab
 - <http://armadilo.cs.uregina.ca/>



ISMIR 2006: Victoria, BC
October 8, 2006

3

Introduction

- Who are you?
 - Computer Scientist?
 - Musician?
 - Musicologist?
 - Student?
- What is your knowledge or interest in
 - MIR?
 - Human-Computer Interaction?

ISMIR 2006: Victoria, BC
October 8, 2006

4

Motivation

- Have you ever been frustrated by a computer program?
 - When?
 - Why?

ISMIR 2006: Victoria, BC
October 8, 2006

5

Motivation

- Consider a computer scientist or engineer interested in MIR
 - What do they want to do?
 - What barriers exist to prevent them from doing it?
- Consider a musician or musicologist interested in MIR
 - What do they want to do?
 - What barriers exist to prevent them from doing it?

ISMIR 2006: Victoria, BC
October 8, 2006

6

The Tools we Use

- <http://www.music-ir.org/evaluation/tools.html>
- How many are easily usable by a non-CS type person?
- How many are easily usable by a CS-type person?
- What is the learning curve required to get into one of these tools?
 - What is the guarantee that the tool will do what you want it to do?

ISMIR 2006: Victoria, BC
October 8, 2006

7


The Tools we Use

- Most assume computer science knowledge
 - familiar with C++, Java or Matlab
 - Access to algorithms like SVM, Bayesian networks, neural nets, MFCC, Auditory models, HMMs, SOMs...
 - What if musicologists don't know what these are?
 - Must a CS person be fluent in C, C++, Java and matlab to use these? Must the platform be the same as the developer platform?

ISMIR 2006: Victoria, BC
October 8, 2006

8

“Just Good Enough” psychology



ISMIR 2006: Victoria, BC
October 8, 2006

9

“Just Good Enough”

- Find a gap and build an implementation.
 - Numerous underlying assumptions
- Distribute the solution
 - Implementation of a fix to *your* problem
 - Making it work with other systems? Not your problem
- Macintosh users face this problem regularly
 - Someone has built it, but it only runs on Windows
 - or commonly in the MIR business: Linux

ISMIR 2006: Victoria, BC
October 8, 2006

10

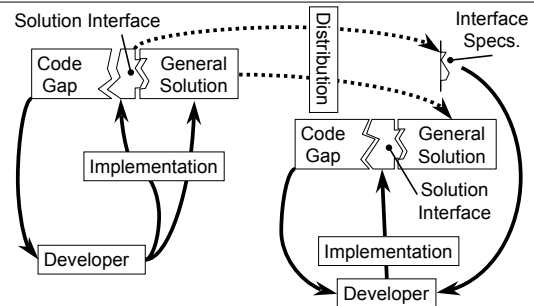
“Just Good Enough” usability

- Developer builds a system
 - “I’ll add the interface later”
 - Builds a basic interface for testing
 - Becomes *familiar* with the basic interface
 - Basic interface begins to seem intuitive
 - System is distributed without interface improvements
 - Leads to:
 - Poor design from good intentions

ISMIR 2006: Victoria, BC
October 8, 2006

11

Intentional interface design



ISMIR 2006: Victoria, BC
October 8, 2006

12

Intentional interface design

- Developer builds general solution + interface
 - Or adheres to a standardized interface
 - Works for human interface or code libraries
 - e.g. standard features, feature ranges, M2K
- New developer acquires the solution and the interface specs
 - Much easier than trying to fit to poorly specified or non-generalized code or interactions

ISMIR 2006: Victoria, BC
October 8, 2006

13

Interface design motivation

- Interaction between communities
 - Computing community is good at building software tools
 - Musicology community is good at thinking about music
 - Each community can benefit from the other
 - Sometimes a musicologist will learn some programming, or a computing professional will learn some musicology
 - Can lead to *more* difficult-to-use implementations:
 - Additional assumptions about user knowledge.

ISMIR 2006: Victoria, BC
October 8, 2006

14

Interface design motivation

- the point, for developers:
 - It is hard to judge initial usability based on your personal experience
- the point, for users:
 - Things don't have to be this hard to use
- the point, from both perspectives
 - we need to work together

ISMIR 2006: Victoria, BC
October 8, 2006

15

Principles of User Interface Design

- Note: these are general principles
- Three ways to characterize the relationship between humans and computers:
 - Manual: humans do all the work
 - Automatic: computers do all the work
 - Augmented: computers support humans doing the work
- Can an interface be *humane*?
 - If it is responsive to human needs and considerate of human frailties

ISMIR 2006: Victoria, BC
October 8, 2006

16

Principles of User Interface Design

- In a *humane* interface:
- People bring: holistic pattern matching, creativity, initiative, exception handling, ability to learn from experience, ability with ill-defined problems, good motor skills, judgment, sense of ethics, ability to apply social context, ability to fail gracefully, flexibility and adaptability
- Computers bring: precision and repeatability, fast and accurate calculations, reliable memory, tirelessness, objectivity, patience, physical robustness

ISMIR 2006: Victoria, BC
October 8, 2006

17

Principles of User Interface Design

- What else are people good at that computers are not good at?
- What else are computers good at that people are not good at?
- Is there anything that both people and computers are good at?
- Is there anything that both people and computers are not good at?
- MIR applications: what are people good at? What are computers good at?

ISMIR 2006: Victoria, BC
October 8, 2006

18

Principles of User Interface Design

- Visibility and Affordance:
 - a teapot handle is visible and it affords grasping in order to pour tea
 - A door handle may be visible but it may not be clear whether to push or pull
 - A door may afford entry to a room but it may not be clear how to operate it
 - “how does the user know that an action is possible through an interface”

ISMIR 2006: Victoria, BC
October 8, 2006

19

Principles of User Interface Design

- Evaluation of interfaces: 3 scales:
- A good interface is:
 - Effective:
 - functionality exists in the software
 - Comprehensible:
 - user can determine how to access functionality
 - Satisfying:
 - the software ceases to get in the way of the user

ISMIR 2006: Victoria, BC
October 8, 2006

20

Principles of User Interface Design

- Utility versus Usability:
 - Software with lots of functionality has high *utility*, but it may be impossible to access it
 - Software that is easy to use has high *usability*, but it may lack functionality
 - Engelbart (designer of the mouse):
 - “it’s easier to learn to ride a tricycle, but you can do more with a bicycle”
 - Not impossible to have high utility *and* high usability.

ISMIR 2006: Victoria, BC
October 8, 2006

21

Principles of User Interface Design

- Utility and Usability in New Music Instruments
- David Wessel:
 - “a low entry fee with no ceiling on virtuosity”
 - A kazoo is easy to use, but you can only go so far
 - A violin takes years of training, but there’s always room to get better
 - No entry fee: CD player, record player
 - Press play for beautiful music
 - But people find ways to add virtuosity
 - Turntableism

ISMIR 2006: Victoria, BC
October 8, 2006

22

Principles of User Interface Design


- operator error = system failure
 - You can’t attribute a problem to “human error”
 - Operators are part of the system
- What is an error, in MIR?
- novice/expert dichotomy is problematic
 - beginners need simplicity, clarity of function, and visibility
 - experts need aptness to task, modelessness, and monotony (repeatable ways to accomplish a task)
 - People might be at either stage with respect to any one feature

ISMIR 2006: Victoria, BC
October 8, 2006

23

Principles of User Interface Design

- Gulfs between computer and user



ISMIR 2006: Victoria, BC
October 8, 2006

24

Classic examples of user interface failure

(December, 1972)

Normally, a green light in the airliner's cockpit signals that landing gear is down. When the indicator failed to light, the pilot decided to circle at 2,000 feet and the autopilot was engaged. All the crew tried to change the bulb but they could not get it out. During these efforts, the autopilot became disengaged. Soon, an automatic warning sounded to indicate they were 250 feet below their assigned altitude. A yellow indicator also lit up. These warnings weren't noticed by the crew.

ISMR 2006: Victoria, BC
October 8, 2006

25

Classic examples of user interface failure

A little later while still struggling with the bulb, the copilot noticed that the altimeter indicated 150 feet -- alarmingly low. He asked the pilot, "we're still at two thousand, right?" The pilot responded, "hey, what's happening here?" As the pilot spoke, a low altitude warning horn went off. But amidst all the warnings, the crew was so sure that they were still at 2,000 feet that 8 seconds later, they crashed into the everglades

ISMR 2006: Victoria, BC
October 8, 2006

26

Classic examples of user interface failure

Lessons:

- If the computer behaves unexpectedly while you are using an interface, you become less likely to see hints, help messages, or other user aids as you become increasingly agitated about the problem.
- The interface has to work, whatever the user's state of absorption.

ISMR 2006: Victoria, BC
October 8, 2006

27

Classic examples of user interface failure

The Therac-25 was used in radiation therapy in the mid-1980's. One of its innovations was software based control and monitoring. Safety mechanisms and interlocks were removed because any software errors could be caught in testing and the software would avoid risks due to wear of the mechanical components. It was a paradigm case of bad software engineering. There were many critical design flaws, no documentation, and no component testing.

ISMR 2006: Victoria, BC
October 8, 2006

28

Classic examples of user interface failure

An operator controlled the machine from a VT100 terminal where he or she would enter parameters, including energy modes. The screen update could not handle rapid typing (from corrections, say) and when a mismatch was found, a "Malfunction 54" message was displayed. Its wording allowed operators to believe too low a dose was delivered, and so they gave another treatment. 6 patients died from massive overdoses.

ISMR 2006: Victoria, BC
October 8, 2006

29

Classic examples of user interface failure

Lessons:

- Developers must understand the user's task (in this case, the likelihood of typos and rapid correction for treatment plans)
- Good feedback is essential (the displays were poorly coordinated with program state, incomprehensible error messages and inadequate documentation)
- Test, test, test (neither basic functionality nor system use was fully tested)

ISMR 2006: Victoria, BC
October 8, 2006

30

Modern examples of user interface failure

- TV/DVD/Set-Top Box remote control
- This interface is *modal*: you must select TV, VCR, or Set-Top before issuing a command to any of these devices (usually buttons at the top of the remote)
- Can be made easier to use by programming the remote to always use the audio box for volume
 - Programming a remote control is harder than programming a VCR.

ISMIR 2006: Victoria, BC
October 8, 2006

31

User interface design

- Consider three separate steps
 - Activities
 - Information
 - Interaction

ISMIR 2006: Victoria, BC
October 8, 2006

32

Activities

- What the software will enable the user to do
- Different from *Techniques*: How the software will perform these activities
 - How the system evaluates similarity or extracts information
 - If the user is a programmer, then the technique is the activity
 - if the user is not a programmer, the technique is secondary, and often invisible

ISMIR 2006: Victoria, BC
October 8, 2006

33

Activity design

- Activities must be
 - Effective
 - Comprehensible
 - Satisfying
- “Just good enough” design often accomplishes only the first requirement

ISMIR 2006: Victoria, BC
October 8, 2006

34

Activity design

- Software that is effective, but not comprehensible or satisfying may exceed a user's “Threshold of indignation”
 - so fed up that they give up and walk away
 - such software is non-functional

ISMIR 2006: Victoria, BC
October 8, 2006

35

Activities in music information retrieval

- Music Query
 - by humming
 - by rhythm
 - by lyric
 - by metadata
 - by score
 - by symbolic data
 - by theme
 - by ...

ISMIR 2006: Victoria, BC
October 8, 2006

36

Activities in music information retrieval

- Playlist generation
 - from library
 - from new music
- Rhythm recognition
 - for beat matching
 - for metadata indexing
 - for playlist generation
- Score analysis

ISMIR 2006: Victoria, BC
October 8, 2006

37

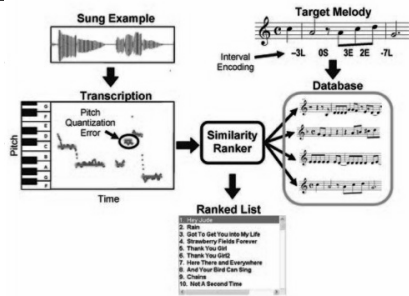
Activities in music information retrieval

- New Music Suggestion
 - by existing library
 - by existing playlist
 - by metadata
 - by genre
- What activities are you interested in or working on?

ISMIR 2006: Victoria, BC
October 8, 2006

38


MIR activities: Query by humming (Birmingham, Communications ACM)



ISMIR 2006: Victoria, BC
October 8, 2006

39

MIR activities: Score alignment (Dannenberg, Communications ACM)



ISMIR 2006: Victoria, BC
October 8, 2006

40

Expand activities into scenarios

- Not just what can be done
- Imagine a specific user doing a specific task
- Give that user context, opinions
- Like a mini case study
- Allows you to "try out" interactions and aspects of the interface

ISMIR 2006: Victoria, BC
October 8, 2006

41

Cognitive models

- Designers and users have different models
 - Designer's models are implementation-based
 - Systematic, logical, comprehensive
 - What's under the hood
 - What's possible, what's expected
 - User's models are experience-based
 - What happens when we do this
 - This looks like something else I've seen
 - early user models may be incomplete or incorrect
 - constantly updated through experience

ISMIR 2006: Victoria, BC
October 8, 2006

42

Metaphors

- Help users develop mental models
 - MP3 player that looks like a tape deck
 - sound editing software that looks like a rack of components
- Not only show what's possible, but may imply what is perceived to be possible
 - can't jump from track to track with a tape deck
 - Our computer player might have that ability
 - Tape deck metaphor may inadvertently alter perceived functionality

ISMIR 2006: Victoria, BC
October 8, 2006

43

Metaphor: "ponds"




Figure 3: To the right is an empty shoal, representing an on-going query using the keyword 'Vivaldi'. To the left is a shoal of creatures representing a finished query.

44

Metaphor: "ponds"

- interesting and novel metaphor: shoals, creatures,
 - but what do they mean?
 - what activities, interaction and information are available from a shoal?
- metaphor for the sake of metaphor
- looks nice, is interesting, but the main advantage of metaphor is *inheritance of familiarity*
- The curse of the "shiny" factor

ISMIR 2006: Victoria, BC
October 8, 2006

45

Metaphor breakdown



Why metaphor breakdown?

- Intimidating for people who don't know about using a rack of components
- Frustrating for people who do, because a knob on a rack is tuned with the fingers, not a mouse
- Physical interfaces can be connected to the software
 - disconnect between action and feedback
- Ahh, but it's shiny!

ISMIR 2006: Victoria, BC
October 8, 2006

47

Automation

- Some tasks should be automated
 - indexing a music library
- Some tasks should not be automated
 - *deciding* to index a music library
 - even if it is necessary for proper operation
 - Inform the user that it could take a while
 - let the user decide
- Default values are automated choices
 - invisible unless on the main interface

ISMIR 2006: Victoria, BC
October 8, 2006

48


Information Design

- How to present information to the user
- What information the user knows
 - Explicit knowledge
 - obvious, written down knowledge
 - Tacit knowledge
 - implied, often subconscious knowledge
 - “everyone knows that”

ISMIR 2006: Victoria, BC
October 8, 2006

49

Gestalt: Perceptual organization

- German: shape, form *gə -'shtält (-'shtölt)*
- "a unified symbolic configuration having properties that cannot be derived from its parts."
 - Visually: a group of elements make a shape
 - Sonically: a group of notes make a melody
- Acoustic example... 
- cohesive units are perceived as *figure* against *background*

ISMIR 2006: Victoria, BC
October 8, 2006

50

Gestalt Principles

- Similarity: elements that share characteristics tend to be perceived as a group
- Proximity: elements near each other tend to be perceived as a group
- Continuity: perception favours smooth contours
- Closure: elements tend to be perceived as complete closed figures




ISMIR 2006: Victoria, BC
October 8, 2006

51

Gestalt Principles

- Area: smaller elements are seen as foreground, larger elements as background
 - Symmetry: symmetrical elements tend to be seen as part of the same figure
- [][][]][][][]
- Illusions occur when two or more gestalt principles compete:
 - Symmetry, Area, background.
 - also, familiarity plays a role



ISMIR 2006: Victoria, BC
October 8, 2006

52

Gestalt principles in interfaces

- *proximity*: words on a menu bar, columns in a tabular display
- *similarity*: toolbar icons
- *closure*: overlapping windows, menus etc.
- *area*: icons, pop-up menus
- *symmetry*: scrollbars
- *continuity*: paragraphs, lists
 - Many can be used together to reinforce the perceptual grouping (redundant coding)
 - radio buttons: shape, colour, proximity, continuity
 - data visualization

ISMIR 2006: Victoria, BC
October 8, 2006

53

Organization tradeoff

- Presenting information about all current options helps people understand what is possible and supports flexible interaction
- BUT
- Every piece of information or control that is presented increases the complexity of the visual display

ISMIR 2006: Victoria, BC
October 8, 2006

54

Perceptual Organization Overload

- Each visual cue (colour, shape etc.) adds complexity and clutter to a display
- Each added feature is less valuable in guiding perception
 - too many colour categories are hard to distinguish.
- “When everything is emphasized, nothing is emphasized.” - Edward Tufte, 1997

ISMIR 2006: Victoria, BC
October 8, 2006

55

Perceptual Design example

#	OPAC Name Headings Search	Title Long	De
[1.1]	Tufte, Edward R., 1942- Library Location: MAIN	Envisioning information / Edward R. Tufte. Call Number: P 93.5 T84 1990	Status: Renewed
[1.2]	Tufte, Edward R., 1942- Library Location: MAIN	Political control of the economy / Edward R. Tufte. Call Number: HB 73 T83	Status: Not Charged
[1.3]	Tufte, Edward R., 1942- Library Location: MAIN	Quantitative analysis of social problems. Edited by Edward R. Tufte. Call Number: H 61 T9	Status: Not Charged
[1.4]	Tufte, Edward R., 1942- Library Location: MAIN	Size and democracy [by] Robert A. Dahl and Edward R. Tufte. Call Number: JC 364 D33	Status: Not Charged
[1.5]	Tufte, Edward R., 1942- Library Location: MAIN	Visual display of quantitative information / Edward R. Tufte. Call Number: QA 276.3 T83 1983	Status: Renewed
[1.6]	Tufte, Edward R., 1942- Library Location: SIFC	Visual display of quantitative information / Edward R. Tufte. Call Number: QA 276.3 T83 1983	Status: Not Charged
[1.7]	Tufte, Edward R., 1942- Library Location: MAIN	Visual explanations : images and quantities, evidence and narrative / Edward R. Tufte. Call Number: P 93.5 T846 1997	Status: Not Charged

ISMIR 2006: Victoria, BC
October 8, 2006

56

Perceptual Design example

- What are the problems?
 - lots and lots of bounding boxes and separators
 - lots of use of colour, italics etc (but these have a purpose)
 - unnecessary things: brackets around numbers [1], boxes at the end of rows, undefined acronyms, jargon (title long)
- What are the good things about it?
 - horizontal and vertical groupings, familiarity (blue underlined = link), redundant coding (red+italics)

ISMIR 2006: Victoria, BC
October 8, 2006

57

The screenshot shows a search interface with a query: "What is the status of the trade deficit with Japan?". Below the query is a table of search results with columns: RANK, PROGRAM, DATE, STORY, SCORE, LENGTH, and HITS. The top result is from NPR All Things Considered, dated 05/21, with a score of 15.83 and 27.65 hits. Below the table is an overview section titled "OVERVIEW - NPR All Things Considered: 05/21" showing a bar chart of the deficit status for Japan and trade. Below that is a transcript section titled "TRANSCRIPTS - NPR All Things Considered: 05/21" containing text from the broadcast.

58

Whittaker et al. retrieval from speech archives

- Good use of gestalt
 - Colour matching reinforces link between information
 - Fonts and background colour used to differentiate text labels
- Be careful with colour
 - More people are colourblind than you might think
 - Not a bad idea to have a second correlating emphasis
 - Colour plus font, for example

ISMIR 2006: Victoria, BC
October 8, 2006

59

presenting lists of music results

Top » Bands and Artists » H

Subcategories:


- (hed) Planet Earth (6)
- H2O (2)
- H2SO4 (3)
- Hackberry Ramblers (7)
- Hackett, Steve (6)
- Hadden, Martin (0)
- Hafez, Abdel Halim (3)
- Haffner Trio, The (3)
- Hager, Regan (0)
- Hagar, Sammy (13)
- Hagen, Nina (13)
- Haggard, Merle (28)
- Haggard, Steve (2)
- Hahn, Hilary (11)
- Haines, Luke (2)
- Haircut 100 (3)
- Haklin (10)
- Haley, Bill and the Comets (5)
- Hemming, Fujiiko (2)
- Henderson, Ainslie (7)
- Henderson, Joe (12)
- Henderson, Murray (1)
- Hendrix, Jimi (77)
- Henley, Don (9)
- Hennessy, Carly (3)
- Henry, Joe (3)
- Henske, Judy (3)
- Hensley, Ken (10)
- Hensley, Tom (1)
- Henson, Jim (11)
- Heptones, The (4)
- Heretic's Fork (2)
- Heritage (0)
- Hermans Hermits (11)
- Herman, Ian (0)
- Herndon, Ty (11)

ISMIR 2006: Victoria, BC
October 8, 2006

60

List of musical results


- page goes on forever. Alphabetical is not a good choice
- if everything has a bullet, bullets aren't informative
- Subcategories:
 - One element is not a list



ISMR 2006: Victoria, BC
October 8, 2006

61

Presentation of tags by popularity



ISMR 2006: Victoria, BC
October 8, 2006

62

Presentation of tags by popularity

- Popular with social computing networks
 - Flickr
 - Blogspot
 - Del.icio.us
- Perception is fast and clear
 - Bigger font represents more of ... whatever the system is trying to display

ISMR 2006: Victoria, BC
October 8, 2006

63

Graphic Elements: familiar?

Be careful using icons - they may not mean what you think they should mean.



ISMR 2006: Victoria, BC
October 8, 2006

64

Graphical element implies meaning

- What does yellow mean? green?
 - no legend given
 - no apparent difference in interaction




ISMR 2006: Victoria, BC
October 8, 2006

65

Visual consistency

- If things look the same, users will assume they have the same functionality
- one interface unit per song
- loading indicator
- options
- preview
 - Icon used implies "play"



ISMR 2006: Victoria, BC
October 8, 2006

66

Visual consistency

- Problem: same interface for an album



- can't play this thing. Why?
 - albums aren't playable?
 - this particular album isn't available?
- turns out albums aren't playable
 - I dismissed 10 albums before I discovered that I had actually listened to a song on one in the same session

ISMIR 2006: Victoria, BC
October 8, 2006

67

A note on Modality

- This interface tries to get rid of modes
 - album interaction appears same as song interaction
 - Usually, modelessness is good
 - a command performs the same function, regardless of state of the system
 - means the command can be executed from anywhere
 - Here, a mode would be useful
 - album browsing is *different* that song browsing
 - or remove albums from consideration

ISMIR 2006: Victoria, BC
October 8, 2006

68

Affordances: What can you do?



ISMIR 2006: Victoria, BC
October 8, 2006

69

Affordances

- What can you do?
 - to look at a different song, do you roll over or click?
 - clicking centers the search on a new song
 - roll-over looks at the song without selecting it
 - Back button
 - shows that there is an “undo” function available

ISMIR 2006: Victoria, BC
October 8, 2006

70

Affordance Breakdown

- Doors which you don't know whether to push or pull
- Visual elements that look like buttons but are not
- Visual elements that are meant to be pushed but don't look like buttons
- Underlined text in a web page that is not a link
- Links in a web page are not underlined

ISMIR 2006: Victoria, BC
October 8, 2006

71

Lack of affordance:

%>_

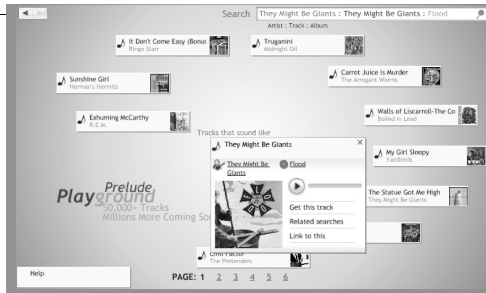


- This time, not enough information. No metaphor, no idea what the thing might be able to do
- “the tyranny of the blank page”

ISMIR 2006: Victoria, BC
October 8, 2006

72

What information is present?



ISMR 2006: Victoria, BC
October 8, 2006

73

What information is present?

- Similarity of songs
 - Is closer similar?
 - Is there a reason they are on a spiral?

ISMR 2006: Victoria, BC
October 8, 2006

74


Information Visualization

- visual features used to code data attributes
 - 2- or 3-D spatial displays
 - temporal displays
 - Colour, shape, texture, positioning, material, connectedness
 - rely heavily on gestalt perception
 - Show similarity across multiple features

ISMR 2006: Victoria, BC
October 8, 2006

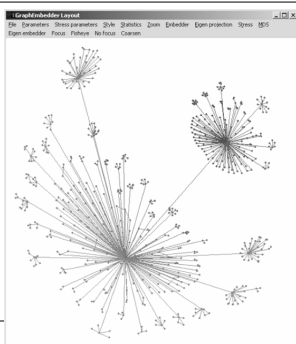
75

Information Visualization Example: Pacific Northwest National Lab



76

Information Visualization Example: ATT Labs



77

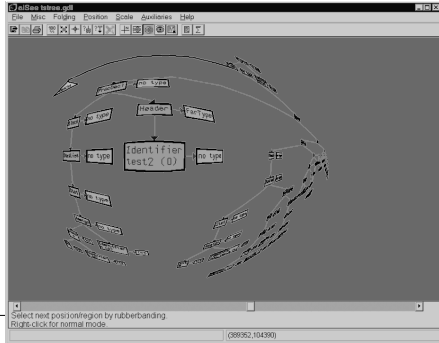
Focus-plus-Context Dynamic Displays

- Details of the area in focus, plus an overview of the areas out of focus
- Selecting a new area shifts the focus to that location
- Examples:
 - Fish-eye visualization
 - Hyperbolic visualization
 - Localized zooming, perhaps?

ISMR 2006: Victoria, BC
October 8, 2006

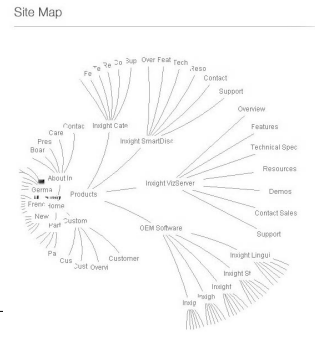
78

Fish-Eye Visualization: <http://www.aisee.com/>




79

Hyperbolic Visualization: <http://www.inxight.com/>



80

Music Similarity Visualization



- Piotr D. Adamczyk, ACM MM 04

ISMIR 2006: Victoria, BC
October 8, 2006

81

Adamczyk

- from "allmusic.com" database
- compared 2d and 3d models
- both were preferred over text models
 - 2d not preferred over 3d or vice versa

ISMIR 2006: Victoria, BC
October 8, 2006

82

Tonality Visualization (Krumhansl, ACM CIE 05)



Fig. 4. The map of musical tones used in Lougnot-Higgins and Soodman's [1971] key-finding algorithm.




Fig. 5. Two helical representations of musical pitch proposed by Shepard [1982], the single helix (left) based on pitch proximity on the chroma circle and octave equivalence, and the double helix (right) based on two whole tone scales where the projection onto the horizontal is the circle of fifths.

Interactive Visualization of Serial Periodic Data (Carlis, UIST 98)



ISMIR 2006: Victoria, BC
October 8, 2006

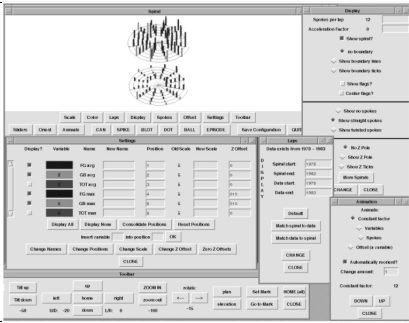
84

- Beautiful elements
- Structural explosion



ISMR 2006: Victoria, BC October 8, 2006 Figure 3. A spiral display of monthly consumption percentages for all 112 foods during the period 1980 - 1989.

But look at the interface...




ISMR 2006: Victoria, BC October 8, 2006 86

Interface to Periodic data visualizer

- Buttons everywhere
- All buttons look the same
- No form, no organization
- Palette windows overlapping
- Generally not great. But really nice visualization method

ISMR 2006: Victoria, BC October 8, 2006 87

Music Key structure visualization



- Chew and Francois, ACM MIE 05


ISMR 2006: Victoria, BC October 8, 2006 88

Chew and Francois

- Complex, multi-level information visualization
- Dynamic
 - visual continuity leads to better gestalt perception
 - things that move together are connected
- From

ISMR 2006: Victoria, BC October 8, 2006 89

Self similarity display



Foote, ACM MM99

ISMR 2006: Victoria, BC October 8, 2006 90

Foote

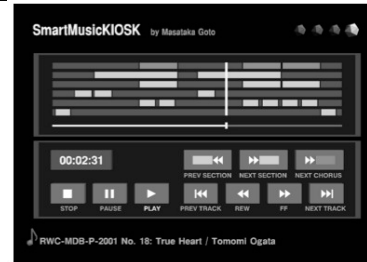
- self-similarity matrix in one song
- shows rhythm, choruses etc

- Similarity matrix between songs
 - for lyric alignment, score alignment, detection of variations or covers

ISMIR 2006: Victoria, BC
October 8, 2006

91

Self-similarity interface



Goto
UIST 03

[video]

ISMIR 2006: Victoria, BC
October 8, 2006

92

Goto

- presents tracks of extracted similarity
 - jump from track to track, from chorus to chorus
 - makes use of a “piano roll” metaphor
 - alternative to score or other symbolic presentations
 - Simplistic, but effective in interactive contexts

ISMIR 2006: Victoria, BC
October 8, 2006

93

Music animation machine

- visualization of symbolic music data
- Piano roll visualization
 - bars are note-ons and note-offs
 - doesn't show velocity(in these versions)
 - colour can be voice, pitch, etc

ISMIR 2006: Victoria, BC
October 8, 2006

94


Music animation machine



ISMIR 2006: Victoria, BC
October 8, 2006

95


Music animation machine



ISMIR 2006: Victoria, BC
October 8, 2006

96

Score annotation visualization



- Sapp CIE 2005

ISMIR 2006: Victoria, BC
October 8, 2006

97

Sapp


- Visual hierarchical key analysis
- similar songs will have similar patterns
- Extract overarching key information
 - considerable built-in musicological knowledge

ISMIR 2006: Victoria, BC
October 8, 2006

98

Music similarity visualization

- Pampalk
- “islands of music”



ISMIR 2006: Victoria, BC
October 8, 2006

100

Pampalk

- Self-organizing map
- Metaphor of a terrain
 - Natural, since it is for exploration of a musical “space”
 - Movement on a horizontal implies similar music
 - Movement up is toward the prototype for a class
- Implications
 - Distance between points is consistent through some measure

ISMIR 2006: Victoria, BC
October 8, 2006

100

Music visualization:

- the ubiquitous spectrum
- anything seem funny to you?
 - it’s a little too broadband for my taste...



ISMIR 2006: Victoria, BC
October 8, 2006

101

the ubiquitous spectrum

- music players should have spectrum displays
 - why? because they’re pretty
- spectrum displays are not cheap
 - so let’s just fake it!
- The good: it’s feedback, telling the user that music should be coming out of the speakers
- the bad: people interested in music are likely to know better.

ISMIR 2006: Victoria, BC
October 8, 2006

102

Feedback

- Indicates that input is being processed
 - e.g. typed text appearing on the screen
 - lets users track progress; adjust behaviour
- Constant and complete feedback is an idealization (system resources etc.)
- How much / how accurate is appropriate?
- e.g. dragging a window
 - show contents of window or just outline?

ISMIR 2006: Victoria, BC
October 8, 2006

103

Usability Testing

- Representative users interact with system prototypes
 - study their behaviour, subjective reactions
- can test all usability aspects:
 - what users will *expect*
 - how users will pursue their *goals*
 - how users will respond to *feedback*
 - what the users subjective *reactions* will be

ISMIR 2006: Victoria, BC
October 8, 2006

104

Usability Testing on a *Working System*

- brings a sense of *realism* to the task
- User behaviour is relatively natural and unhindered
 - all the aspects of the user interface are complete and functional *for the given task*
- must wait until we are well into the development of the system before we can start this type of testing

ISMIR 2006: Victoria, BC
October 8, 2006

105

Rapid Prototype for Usability Testing

- Build a realistic simulation with rapid prototyping tools
- The prototype is temporary and will eventually be replaced by the real system
- There is the risk that if the prototype exhibits enough system functionality, the team (or management) may believe that this is the final system

ISMIR 2006: Victoria, BC
October 8, 2006

106

Working Partial System

- An executable prototype that includes a subset of the intended functionality
- Horizontal Prototype
 - all intended functionality, only at top level
 - for studying high-level goals, action plans
- Vertical Prototype
 - implement a task in full detail
 - for studying usability of complex task

ISMIR 2006: Victoria, BC
October 8, 2006

107

Some Online services

- <http://www.songtapper.com/>
 - Tap out the rhythm to find a song
- <http://search.singingfish.com/>
 - text-based search
- <http://www.musicrobot.com/>
 - text-based search
- <http://hype.non-standard.net/>
 - The Hype machine: collects publicly posted mp3s from blogs
- <http://www.last.fm/explore/>
 - Scrobbling

ISMIR 2006: Victoria, BC
October 8, 2006

108

Online services

- musicbrainz
- musicip
- amg lasso
 - <http://www.allmediaguide.com/lasso/>
 - for fingerprinting existing music
 - Similar to Shazam

ISMIR 2006: Victoria, BC
October 8, 2006

109

<http://www.soundjunction.org/>

- tagline: listen. explore. discover. create
- mostly music education, some composition games
- Intriguing “Journey Mode” for web exploration

ISMIR 2006: Victoria, BC
October 8, 2006

110

Now what do I do?



113

www.mystrands.com



ISMIR 2006: Victoria, BC
October 8, 2006

112

www.mystrands.com

- requires a user account
- requires downloading of a client
- requires indexing of user's music library
- Can look at popular music from the rest of the online community without an account
- Similar to other existing library indexers and metadata retrieval systems

ISMIR 2006: Victoria, BC
October 8, 2006

www.mystrands.com

- “Music Strands”
 - Playlist generation, music suggestions
 - metadata-based
 - if your songs are improperly labeled or tagged, no luck
 - I have a few “unnamed” mp3s in my library
 - no suggestions, not unusual songs
 - Builds database from your music library
 - Connects to existing database which cross-references

ISMIR 2006: Victoria, BC
October 8, 2006

114

library indexing systems: issues

- Entry fee: client system
 - time to download
 - risk of malicious/malignant software
- Entry fee: user account
 - risk of providing personal information
- Entry fee: library indexing
 - risk of exposing not-quite-legal music ownership
 - risk of suggestions on novelty music
 - unless listening frequency is encoded

ISMIR 2006: Victoria, BC
October 8, 2006

115

library indexing systems: issues

- Personal Libraries
 - Obscure or unusual music not recognized
 - Incorrectly tagged music not recognized

ISMIR 2006: Victoria, BC
October 8, 2006

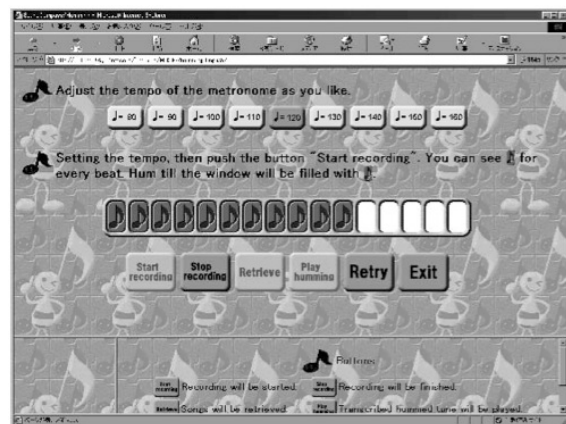
116

To Finish: Some not-so-good interfaces

- Sorry if these are yours...
- I'll keep them anonymous
- See if you can spot the problems

ISMIR 2006: Victoria, BC
October 8, 2006

117

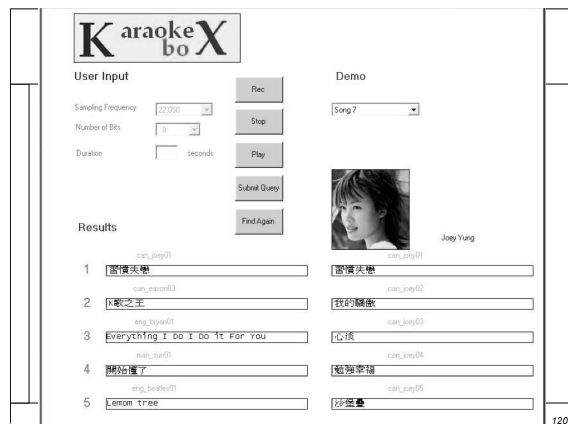


Problems

- Background is distracting
- Musical notes indicate beats recorded
 - Are they really eighth notes? Unlikely
 - Crossover metaphor
 - By implying music notation, designer implies meaning
- Text instructions aren't great
 - Some say a good user interface doesn't need text to suggest the affordances
- Collection of buttons for tempo - slider would be better
- Big buttons below beat display seem to be all one group, but do three different things

ISMIR 2006: Victoria, BC
October 8, 2006

119



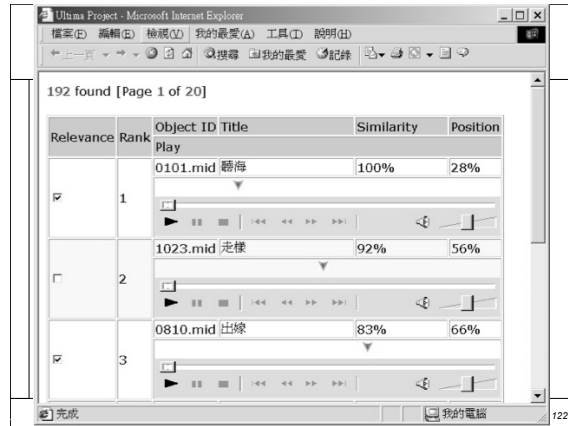
120

Problems

- Who is that person? Why is she there?
- What's the difference between "user input" and "demo"
- What's the difference between "submit query" and "find again"
- Sound input settings shouldn't be right on the main interface
 - Adds unnecessary clutter, system should handle this

ISMIR 2006: Victoria, BC
October 8, 2006

121



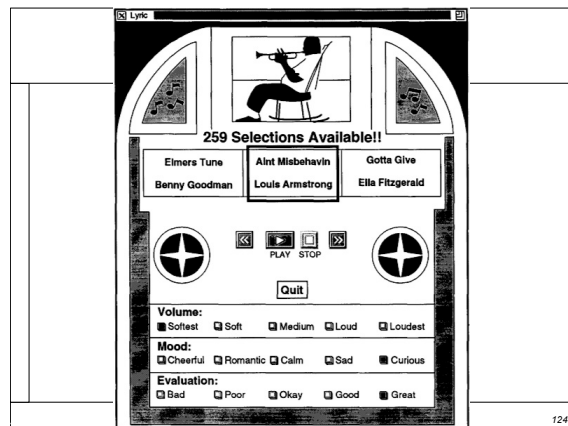
122

Problems

- Again, separator overload
- Similarity 100%? Really? Was that the query that was sent?
- What is "Position?"
- 192 results is overwhelming
- Player in each is nice,
 - But would you ever play them all at once?
- Red triangle means what?

ISMIR 2006: Victoria, BC
October 8, 2006

123



124

Problems

- Quit button in the middle of the interface.
 - Most users are used to the top right corner of the window
- 3 sets of radio buttons
 - One set is discrete and unrelated (like radio stations - good metaphor for radio buttons)
 - two of them are continuous scales
 - most radios use a dial for volume
- What does this "evaluation" mean
 - do you like the song? do you like the suggestion?
 - a song you like in a bad context could be poor, as could a song you don't like in the right context
- Gotta love double exclamation marks!!

ISMIR 2006: Victoria, BC
October 8, 2006

125

References

- [1] P. D. Adamczyk. Seeing sounds: exploring musical social networks. In MULTIMEDIA '04: Proceedings of the 12th annual ACM international conference on Multimedia, pages 512–515, New York, NY, USA, 2004. ACM Press.
- [2] M. Allen, J. Gluck, K. MacLean, and E. Tang. An initial usability assessment for symbolic haptic rendering of music parameters. In ICMI '05: Proceedings of the 7th international conference on Multimodal interfaces, pages 244–251, New York, NY, USA, 2005. ACM Press.
- [3] J. L. Alty and D. I. Rigas. Communicating graphical information to blind users using music: the role of context. In CHI '98: Proceedings of the SIGCHI conference on Human factors in computing systems, pages 574–581, New York, NY, USA, 1998. ACM Press/Addison-Wesley Publishing Co.
- [4] L. Bass and B. E. John. Linking usability to software architecture patterns through general scenarios. The Journal of Systems and Software, 66:187–197, 2003.
- [5] T. Beamish, K. Maclean, and S. Fels. Manipulating music: multimodal interaction for djs. In CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems, pages 327–334, New York, NY, USA, 2004. ACM Press.
- [6] P. Bellini, P. Nesi, and M. B. Spina. Cooperative visual manipulation of music notation. ACM Trans. Comput.-Hum. Interact., 9(3):194–237, 2002.
- [7] F. Bentley, C. Metcalf, and G. Harboe. Personal vs. commercial content: the similarities between consumer use of photos and music. In CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems, pages 667–676, New York, NY, USA, 2006. ACM Press.

ISMIR 2006: Victoria, BC
October 8, 2006

126

[8] W. Birmingham, R. Dannenberg, and B. Pardo. Query by humming with the vocalsearch system. *Commun. ACM*, 49(8):49-52, 2006.

[9] J. O. Borchers. Worldbeat: designing a baton-based interface for an interactive music exhibit. In CHI '97: Proceedings of the SIGCHI conference on Human factors in computing systems, pages 131-138, New York, NY, USA, 1997. ACM Press.

[10] D. Byrd. Music-notation searching and digital libraries. In JCDL '01: Proceedings of the 1st ACM/IEEE-CS joint conference on Digital libraries, pages 239-246, New York, NY, USA, 2001. ACM Press.

[11] R. Cai, L. Lu, and H.-J. Zhang. Using structure patterns of temporal and spectral feature in audio similarity measure. In MULTIMEDIA '03: Proceedings of the eleventh ACM international conference on Multimedia, pages 219-222, New York, NY, USA, 2003. ACM Press.

[12] P. Cano, M. Koppenberger, and N. Wack. Content-based music audio recommendation. In MULTIMEDIA '05: Proceedings of the 13th annual ACM international conference on Multimedia, pages 211-212, New York, NY, USA, 2005. ACM Press.

[13] J. V. Carlis and J. A. Konstan. Interactive visualization of serial periodic data. In UIST '98: Proceedings of the 11th annual ACM symposium on User interface software and technology, pages 29-38, New York, NY, USA, 1998. ACM Press.

[14] S. Carroll. Deepwave: music visualization. *SIGGRAPH Comput. Graph.*, 34(2):47, 2000.

[15] B. P. Chellis and A. D. N. Edwards. Wiesel: a computer based system for providing non-visual access to music notation. *SIGGRAPH Comput. Phys. Handicap.*, (69):1-12, 2000.

ISMIR 2006: Victoria, BC
October 8, 2006 127

[16] E. Chew. Foreword to special issue on music visualization. *Comput. Entertain.*, 3(4):1-3, 2005.

[17] E. Chew and A. R. J. Francois. Interactive multi-scale visualizations of tonal evolution in musa.rt opus 2. *Comput. Entertain.*, 3(4):1-16, 2005.

[18] S. J. Cunningham, N. Reeves, and M. Britland. An ethnographic study of music information seeking: implications for the design of a music digital library. In JCDL '03: Proceedings of the 3rd ACM/IEEE-CS joint conference on Digital libraries, pages 5-16, Washington, DC, USA, 2003. IEEE Computer Society.

[19] R. B. Dannenberg and C. Raphael. Music score alignment and computer accompaniment. *Commun. ACM*, 49(8):38-43, 2006.

[20] D. Daveco, D. Cakmakov, and V. Cabukovski. A multimedia cognitive-based information retrieval system. In CSC '91: Proceedings of the 19th annual conference on Computer Science, pages 430-438, New York, NY, USA, 1991. ACM Press.

[21] J. S. Downie and J. Futrel. Terascala music mining. In SC '05: Proceedings of the 2005 ACM/IEEE conference on Supercomputing, page 71, Washington, DC, USA, 2005. IEEE Computer Society.

[22] S. Downie and M. Nelson. Evaluation of a simple and effective music information retrieval method. In SIGIR '00: Proceedings of the 23rd annual international ACM SIGIR conference on Research and development in information retrieval, pages 73-80, New York, NY, USA, 2000. ACM Press.

[23] J. W. Dunn, D. Byrd, M. Ntess, J. Riley, and R. Scherle. Variations2: retrieving and using music in an academic setting. *Commun. ACM*, 49(8):53-58, 2006.

ISMIR 2006: Victoria, BC
October 8, 2006 128

[24] J. W. Dunn and E. J. Isaacson. Variations2: a digital music library system. In JCDL '02: Proceedings of the 2nd ACM/IEEE-CS joint conference on Digital libraries, pages 374-374, New York, NY, USA, 2002. ACM Press.

[25] J. Dunne, L. Lapal, M. Flury, M. Shabb, T. Warner, K. Hammond, and L. Birnbaum. mpmel: music recommendation and exploration. In IUI '02: Proceedings of the 7th international conference on Inteligent user interfaces, pages 235-235, New York, NY, USA, 2002. ACM Press.

[26] J. Foote. Visualizing music and audio using self-similarity. In MULTIMEDIA '99: Proceedings of the seventh ACM international conference on Multimedia (Part 1), pages 77-80, New York, NY, USA, 1999. ACM Press.

[27] M. Goto. Smartmusiciosk: music listening station with chorus-search function. In UIST '03: Proceedings of the 16th annual ACM symposium on User interface software and technology, pages 31-40, New York, NY, USA, 2003. ACM Press.

[28] M. Hatala, L. Kalantari, R. Wakkary, and K. Newby. Ontology and rule based retrieval of sound objects in augmented audio reality system for museum visitors. In SAC '04: Proceedings of the 2004 ACM symposium on Applied computing, pages 1045-1050, New York, NY, USA, 2004. ACM Press. 1

[29] G. Haus and E. Pollastri. A multimodal framework for music inputs (poster session). In MULTIMEDIA '00: Proceedings of the eighth ACM international conference on Multimedia, pages 362-364, New York, NY, USA, 2000. ACM Press.

[30] A. Hazan. Towards automatic transcription of expressive oral percussive performances. In IUI '05: Proceedings of the 10th international conference on Inteligent user interfaces, pages 296-298, New York, NY, USA, 2005. ACM Press.

ISMIR 2006: Victoria, BC
October 8, 2006 129

[31] R. Hiraga and N. Matsuda. Visualization of music performance as an aid to listener's comprehension. In AVI '04: Proceedings of the working conference on Advanced visual interfaces, pages 103-106, New York, NY, USA, 2004. ACM Press.

[32] R. Hiraga, R. Mizaki, and I. Fujishiro. Performance visualization: a new challenge to music through visualization. In MULTIMEDIA '02: Proceedings of the tenth ACM international conference on Multimedia, pages 239-242, New York, NY, USA, 2002. ACM Press.

[33] K. Hirata, Y. Hara, N. Shibata, and F. Hirabayashi. Media-based navigation for hypermedia systems. In HYPERTEXT '93: Proceedings of the fifth ACM conference on Hypertext, pages 159-173, New York, NY, USA, 1993. ACM Press.

[34] K. Hoashi, K. Matsumoto, and N. Inoue. Personalization of user profiles for content-based music retrieval based on relevance feedback. In MULTIMEDIA '03: Proceedings of the eleventh ACM international conference on Multimedia, pages 110-119, New York, NY, USA, 2003. ACM Press.

[35] H. Hochheiser. Browsers with changing parts: a catalog explorer for philip glass' website. In DIS '00: Proceedings of the conference on Designing interactive systems, pages 105-115, New York, NY, USA, 2000. ACM Press.

[36] J.-L. Hsu, A. L. P. Chen, H.-C. Chen, and N.-H. Liu. The effectiveness study of various music information retrieval approaches. In CIKM '02: Proceedings of the eleventh international conference on Information and knowledge management, pages 422-429, New York, NY, USA, 2002. ACM Press.

[37] M. Jones and S. Jones. The music is the message. *interactions*, 13(4):24-27, 2006.

ISMIR 2006: Victoria, BC
October 8, 2006 130

[38] B. Jung, H. Hwang, S. Lee, G. J. Kim, and H. Kim. Incorporating co-presence in distributed virtual music environment. In VRST '00: Proceedings of the ACM symposium on Virtual reality software and technology, pages 206-211, New York, NY, USA, 2000. ACM Press.

[39] J. Kominek and R. Kazman. Accessing multimedia through concept clustering. In CHI '97: Proceedings of the SIGCHI conference on Human factors in computing systems, pages 19-26, New York, NY, USA, 1997. ACM Press.

[40] N. Kosugi, Y. Nishihara, T. Sakata, M. Yamamoto, and K. Kushima. A practical query-by-humming system for a large music database. In MULTIMEDIA '00: Proceedings of the eighth ACM international conference on Multimedia, pages 333-342, New York, NY, USA, 2000. ACM Press.

[41] N. Kosugi, Y. Sakurai, and M. Morimoto. Soundcompass: a practical query-by-humming system; normalization of scalable and shiftable time-series data and effective subsequence generation. In SIGMOD '04: Proceedings of the 2004 ACM SIGMOD international conference on Management of data, pages 881-886, New York, NY, USA, 2004. ACM Press.

[42] C. L. Krumhansl. The geometry of musical structure: a brief introduction and history. *Comput. Entertain.*, 3(4):1-14, 2005.

[43] M. S. Lew, N. Sebe, C. Djeraba, and R. Jain. Content-based multimedia information retrieval: State of the art and challenges. *ACM Trans. Multimedia Comput. Commun. Appl.*, 2(1):1-19, 2006.

[44] H. Lieberman and J. Espinosa. A goal-oriented interface to consumer electronics using planning and commonsense reasoning. In IUI '06: Proceedings of the 11th international conference on Inteligent user interfaces, pages 226-233, New York, NY, USA, 2006. ACM Press.

ISMIR 2006: Victoria, BC
October 8, 2006 131

[45] S. Loeb. Architecting personalized delivery of multimedia information. *Commun. ACM*, 35(12):39-47, 1992.

[46] R. J. McNab, L. A. Smith, I. H. Witten, C. L. Henderson, and S. J. Cunningham. Towards the digital music library: tune retrieval from acoustic input. In DL '96: Proceedings of the first ACM international conference on Digital libraries, pages 11-18, New York, NY, USA, 1996. ACM Press.

[47] M. Malucchi and N. Orso. Musical information retrieval using melodic surface. In DL '99: Proceedings of the fourth ACM conference on Digital libraries, pages 152-160, New York, NY, USA, 1999. ACM Press.

[48] T. MIURA and I. SHIOYA. Similarity among melodies for music information retrieval. In CIKM '03: Proceedings of the twelfth international conference on Information and knowledge management, pages 61-68, New York, NY, USA, 2003. ACM Press.

[49] M. Ntess and J. Dunn. Variations2: improving music findability in a digital library through work-centric metadata. In JCDL '04: Proceedings of the 4th ACM/IEEE-CS joint conference on Digital libraries, pages 422-422, New York, NY, USA, 2004. ACM Press.

[50] K. O'Hara, M. Lipson, M. Jansen, A. Unger, H. Jeffries, and P. Macer. Jukola: democratic music choice in a public space. In DIS '04: Proceedings of the 2004 conference on Designing interactive systems, pages 145-154, New York, NY, USA, 2004. ACM Press.

[51] E. Pampalk, A. Rauber, and D. Merkl. Content-based organization and visualization of music archives. In MULTIMEDIA '02: Proceedings of the tenth ACM international conference on Multimedia, pages 570-579, New York, NY, USA, 2002. ACM Press.

ISMIR 2006: Victoria, BC
October 8, 2006 132

[52] B. Pardo. Introduction. *Commun. ACM*, 49(8):28–31, 2006.

[53] L. Prechelt and R. Typke. An interface for melody input. *ACM Trans. Comput.-Hum. Interact.*, 8(2):133–149, 2001.

[54] J. Raskin. *The Humane Interface*. Addison-Wesley, 2000.

[55] M. B. Rossion and J. M. Carroll. *Usability Engineering: Scenario-Based Development of Human-Computer Interaction*. Morgan-Kaufmann, 2002.

[56] C. S. Sapp. Visual hierarchical key analysis. *Comput. Entertain.*, 3(4):1–19, 2005.

[57] S. M. Smith and G. N. Williams. A visualization of music. In *VIS '97: Proceedings of the 8th conference on Visualization '97*, pages 499–501, Los Alamitos, CA, USA, 1997. IEEE Computer Society Press, 2.

[58] O. Sjöberg, H. A. Wallberg, J. Sjöberg, J. Humble, L. E. Fahlgren, A. Bullock, and J. Lundberg. Information exploration using the pond. In *CVI '02: Proceedings of the 4th international conference on Collaborative virtual environments*, pages 72–79, New York, NY, USA, 2002. ACM Press.

[59] D. Tao, H. Liu, and X. Tang. K-box: a query-by-singing based music retrieval system. In *MULTIMEDIA '04: Proceedings of the 12th annual ACM international conference on Multimedia*, pages 464–467, New York, NY, USA, 2004. ACM Press.

[60] P. Toivainen. Visualization of tonal content with self-organizing maps and self-similarity matrices. *Comput. Entertain.*, 3(4):1–10, 2005.

ISMIR 2006: Victoria, BC
October 8, 2006 133

[61] E. Unal, S. S. Narayanan, H. H. Shih, E. Chew, and C. C. J. Kuo. Creating data resources for designing user-centric frontends for query by humming systems. In *MIR '03: Proceedings of the 5th ACM SIGMM international workshop on Multimedia information retrieval*, pages 116–121, New York, NY, USA, 2003. ACM Press.

[62] A. van Breemen and C. Bartneck. An emotional interface for a music gathering application. In *IUI '03: Proceedings of the 8th international conference on Intelligent user interfaces*, pages 307–309, New York, NY, USA, 2003. ACM Press.

[63] A. Voids, R. E. Grinter, N. Ducheneaut, W. K. Edwards, and M. W. Newman. Listening in: practices surrounding itunes music sharing. In *CHI '05: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 191–200, New York, NY, USA, 2005. ACM Press.

[64] A. Wang. The shazam music recognition service. *Commun. ACM*, 49(8):44–48, 2006.

[65] Y. Wang, M.-Y. Kan, T. L. Nwe, A. Shenoy, and J. Yin. Lyrically: automatic synchronization of acoustic musical signals and textual lyrics. In *MULTIMEDIA '04: Proceedings of the 12th annual ACM international conference on Multimedia*, pages 212–219, New York, NY, USA, 2004. ACM Press.

[66] J. Watanabe, H. Nii, Y. Hashimoto, and M. Inami. Visual resonator: interface for interactive cocktail party phenomenon. In *CHI '06: CHI '06 extended abstracts on Human factors in computing systems*, pages 1505–1510, New York, NY, USA, 2006. ACM Press.

ISMIR 2006: Victoria, BC
October 8, 2006 134

[67] M. Weideman and W. Kritzing. Concept mapping vs. web page hyperlinks as an information retrieval interface: preferences of postgraduate culturally diverse learners. In *SAICSIT '03: Proceedings of the 2003 annual research conference of the South African institute of computer scientists and information technologists on Enablement through technology*, pages 69–82, Republic of South Africa, 2003. South African Institute for Computer Scientists and Information Technologists.

[68] G. Weinberg and S. Driscoll. Robot-human interaction with an anthropomorphic percussionist. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 1229–1232, New York, NY, USA, 2006. ACM Press.

[69] S. Whittaker, J. Hirschberg, J. Choi, D. Hindle, F. Pereira, and A. Singhal. Scan: designing and evaluating user interfaces to support retrieval from speech archives. In *SIGIR '99: Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval*, pages 26–33, New York, NY, USA, 1999. ACM Press.

[70] T. Winograd, editor. *Bringing Design to Software*. Addison-Wesley, 1996.

ISMIR 2006: Victoria, BC
October 8, 2006 135

Tutorial presented to
ISMIR 2006: Victoria BC
October 8, 2006

User Interfaces for Music Information Retrieval

David Gerhard
Department of Computer Science
Department of Music
University of Regina
david.gerhard@uregina.ca