

CSC475 Music Information Retrieval

Introduction, overview and MIR tasks

George Tzanetakis

University of Victoria

2014

Table of Contents I



- 1 Course Logistics
- 2 My Background and MIR
- 3 Syllabus and Schedule
- 4 History
- 5 MIR Tasks

Prerequisites

- Basic high school math
- Probability and Statistics
- Linear Algebra
- Computer programming
- Basic music theory

The associated textbook tries to be self-contained and has chapters covering the fundamentals required for understanding the published MIR literature.

Learning outcomes

- Basic knowledge of DSP
- Basic knowledge of Machine Learning (ML)
- Basic knowledge of Music Theory
- Familiarity with the basic tasks that have been explored in MIR research and the algorithms used to solve them
- **Being able to read and understand the majority of published literature in ISMIR**
- Experience with designing and developing MIR algorithms and systems

- 3 lectures/week (MWR 14:30-15:30 PST)
- 2 hours of associated homework expected for each lecture
- Total weekly commitment approximately 9-10 hours
- 4 assignments each worth 10% of final grade and done individually.
- Assignments typically will combine some reading and understanding of MIR literature as well as some programming of MIR algorithms
- 1 final group project (2-3 students per group)

External participants

- Course is open to external participants
- No credit or grading provided for externals
- What is provided:
 - Access to video/audio recording of lectures
 - Answers to questions
 - Assignment model solutions and grading rubric
- Two main points of contact for both internal and external participants:
 - Google+ community
 - Textbook/course webpage
 - Also Hangouts on Air + Google Moderator for questions

Notes for this course are in the process of evolving into a textbook. It is under heavy construction and I hope to use this offering of the course as motivation to get it closer to completion. Please be patient about the process.

Current draft on webpage - will be frequently updated and dated. Also there is much more material for the remaining chapters that exists in other documents, papers, etc that I have that I will be editing and transferring as the term progresses.

Table of Contents I



- 1 Course Logistics
- 2 My Background and MIR
- 3 Syllabus and Schedule
- 4 History
- 5 MIR Tasks

- Main focus of my research has been Music Information Retrieval (MIR)
- Involved from the early days of the field
- Have published papers in almost every ISMIR conference
- Organized ISMIR in 2006
- Tutorials on MIR in several conferences

Education and Academic Work Experience



- 1997 BSc in Computer Science (CS), University of Crete, Greece
- 1999 MA in CS, Princeton University, USA
- 2002 PhD in CS, Princeton University, USA
- 2003 PostDoc in CS, Carnegie Mellon University, USA
- 2004 Assistant Professor in CS, Univ. of Victoria, Canada
- 2010 Associate Professor in CS, Univ. of Victoria, Canada
- 2010 Canada Research Chair (Tier II) in Computer Analysis of Audio and Music
- Music theory, saxophone and piano performance, composition, improvisation both in conservatory and academic settings

Inherently inter-disciplinary and cross-disciplinary work.
Connecting theme: making computers better understand music to create more effective interactions with musicians and listeners. Audio analysis is challenging due to large volume of data - did big data before it became fashionable.

- Music Information Retrieval
- Digital Signal Processing
- Machine Learning
- Human-Computer Interaction
- Software Engineering
- Artificial Intelligence
- Multimedia
- Robotics
- Visualization
- Programming Languages

Inherently inter-disciplinary and cross-disciplinary work.
Connecting theme: making computers better understand music to create more effective interactions with musicians and listeners. Audio analysis is challenging due to large volume of data - did big data before it became fashionable.

- Music Information Retrieval
- Digital Signal Processing
- Machine Learning
- Human-Computer Interaction
- Software Engineering
- Artificial Intelligence
- Multimedia
- Robotics
- Visualization
- Programming Languages

Work Experience beyond Academia

Many internships in research labs throughout studies. Several consulting jobs while in academia. A few representative examples:

- Moodlogic Inc (2000). Designed and developed one of the earliest audio fingerprinting systems (patented) - 100000 users matching to 1.5 million songs
- Teligence Inc (2005). Automatic male/female voice discrimination for voice messages used in popular phone dating sites - processing of 20000+ recordings per day.

- Music Analysis, Retrieval and Synthesis for Audio Signals
- Open source in C++ with Python Bindings
- Started by me in 1999 - core team approximately 4-5 developers
- Approximately 400 downloads per month
- Many projects in industry and academia
- State-of-the-art performance while frequently orders of magnitude faster than other systems

Visiting Scientist at Google Research (6 months)



Things I worked on (of course as part of larger teams):

- Cover Song Detection (applied to every uploaded YouTube video).
- 100 hours of video are uploaded to YouTube every minute
- Content ID scans over 250 years of video every day - 15 million references
- Audio Fingerprinting (part of Android Jelly Bean)
- Named inventor on 6 pending US patents related to audio matching and fingerprinting

Table of Contents I



- 1 Course Logistics
- 2 My Background and MIR
- 3 Syllabus and Schedule**
- 4 History
- 5 MIR Tasks

Schedule

The detailed lecture *TENTATIVE* schedule can be found in the G+ community as well as the webpage. Basically the first few weeks we will cover some fundamental concepts and after that we will examine how they can be used to develop algorithms and systems for several MIR tasks.

Table of Contents I



- 1 Course Logistics
- 2 My Background and MIR
- 3 Syllabus and Schedule
- 4 History
- 5 MIR Tasks

History of MIR before computers

How did a listener encounter a new piece of music throughout history ?

- Live performance
- Music Notation
- Physical recording
- Radio

Brief History of computer MIR

- Pre-history (< 2000): scattered papers in various communities. Symbolic processing mostly in digital libraries and information retrieval venues and audio processing (less explored) mostly in acoustics and DSP venues.
- The birth 2000: first International symposium on Music Information Retrieval (ISMIR) with funding from NSF Digital Libraries II initiative organized by J. Stephen Downie, Time Crawford and Don Byrd. First contact between the symbolic and the audio side.
- 2000-2006 Rapid growth
- 2006-2014 Slower growth and steady state

Table of Contents I



- 1 Course Logistics
- 2 My Background and MIR
- 3 Syllabus and Schedule
- 4 History
- 5 MIR Tasks

Conceptual MIR dimensions I

- Stages
 - Representation/Hearing
 - Analysis/Learning
 - Interaction/Action
- Specificity
 - Audio fingerprinting
 - Common score performance
 - Cover song detection
 - Artist identification
 - Genre classification
 - Recommendation ?

Conceptual MIR dimensions II

Data sources:

- Audio
- Track metadata
- Score
- Lyrics
- Reviews
- Ratings
- Download patterns
- Micro-blogging

- Similarity retrieval, playlists, recommendation
- Classification and clustering
- Tag annotation
- Rhythm, melody, chords
- Music transcription and source separation
- Query by humming
- Symbolic MIR
- Segmentation, structure, alignment
- Watermarking, fingerprinting and cover song detection