Carnegie Mellon

Music Understanding
and the
Future of Music Performance

Roger B. Dannenberg
Professor of Computer Science, Art, and Music
Carnegie Mellon University

Carnegie Mellon

Why Computers and Music?

- Music in every human society!
- Computing can make music:
  - More Fun
  - More Available
  - Higher Quality
  - More Personal

Carnegie Mellon

My Background

- Always interested in math and music and making things
- Discovered synthesizers in high school
- Discovered computers about the same time
- Discovered computer music in college
- Research motivated by musical experience:
  - Computer accompaniment
  - Expressive programming languages for music
  - Audacity
  - … current work

Carnegie Mellon

Overview

- Introduction

- How Is Computation Used in Music Today?
- New Capabilities:
  What Can Computers Do Tomorrow?
- What Will Music Be Like in the Future?

Carnegie Mellon

How Is Computation Used in Music Today?

Indabamusic.com

making-music.blogspot.com

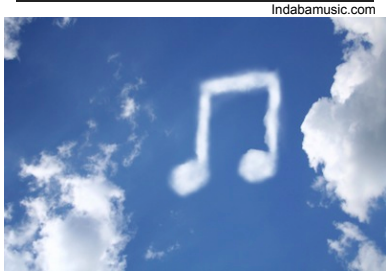

http://venturebeat.com/

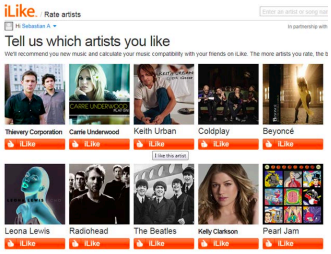

Carnegie Mellon

Music Computation Today

Production: digital recording, editing, mixing

- Nearly all music production today...
  - Records audio to (digital) disk
  - Edit/manipulate audio digitally
    - Equalization
    - Reverberation
  - Convert to media:
    - CD
    - MP3
    - Etc.

protools.com

Carnegie Mellon

Music Computation Today

Musical Instruments: synthesizers and controllers

Synthesizer (Solaris)

Drum Machine (Yamaha)

Linnstrument (Roger Linn)

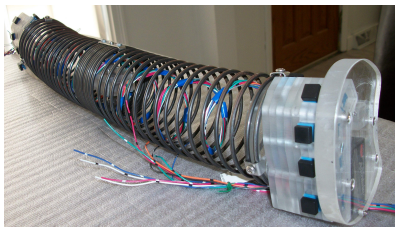

Sonic Spring (Tomas Henriques)

Carnegie Mellon

Music Computation Today

Distribution: compression, storage, networks

Napster

Apple iPod

Apple iTunes

Amazon Cloud Player

Carnegie Mellon

Music Computation Today

Search, recommendation, music fingerprinting

Google Music China

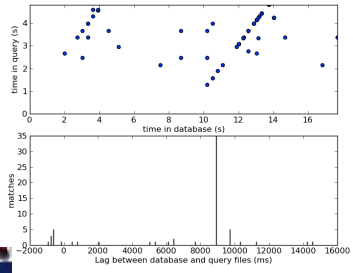

Music Fingerprinting

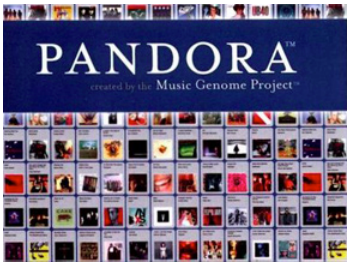

Pandora
Music Recommendation

Carnegie Mellon

Overview

- Computer Music Introduction
- How Is Computation Used in Music Today?

- New Capabilities:
  What Can Computers Do Tomorrow?
- What Will Music Be Like in the Future?

Carnegie Mellon

New Capabilities: What Can Computers Do Tomorrow?

- Computer accompaniment
- Style classification
- Score alignment
- Onset detection
- Sound synthesis

Carnegie Mellon

Accompaniment Video

Carnegie Mellon

Computer Accompaniment

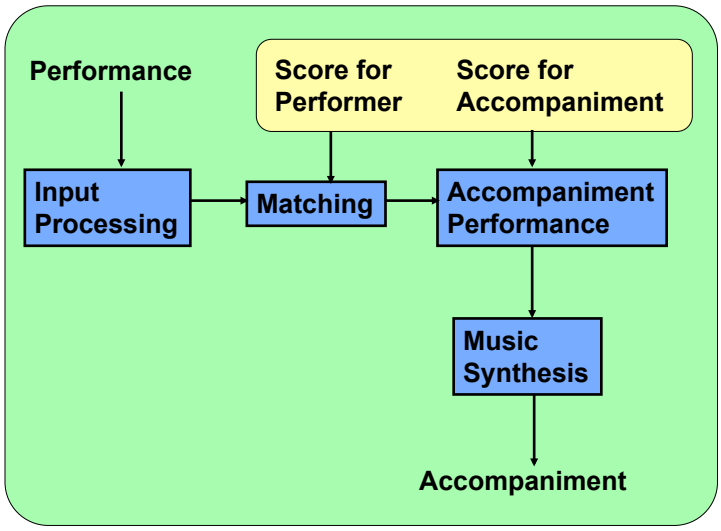

13

Carnegie Mellon

Computer Accompaniment

Performance →

Score →

	A	B	A				
A	1	1					
B	1	2	2				
B	1	2	2				
A	1	2	3				
C		2	3				
B			3				
G							

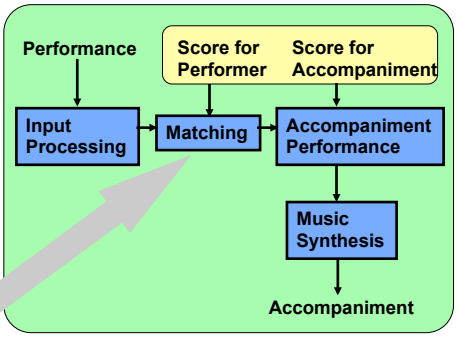

Dynamic Programming, plus ...

On-line, column-by-column evaluation

Windowing for real-time evaluation

Heuristics for best-yet matching

Penalty for skipping notes

14

Carnegie Mellon

Computer Accompaniment

Rule-based system:

E.g. If matcher is confident and accompaniment is ahead < 0.1s, stop until synchronized.

If matcher is confident and accompaniment is behind <0.5s, speed up until synchronized.

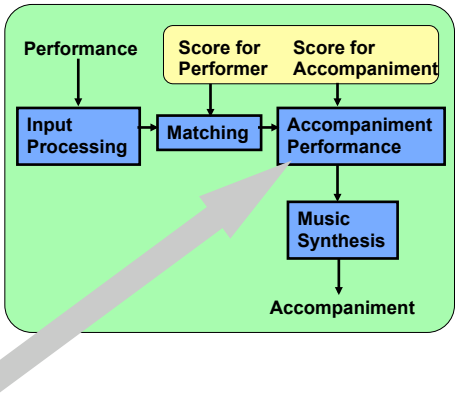

Carnegie Mellon

Vocal Accompaniment

- Lorin Grubb's Ph.D. (CMU CSD)
- Machine learning used to:
  - Learns what kinds of tempo variation are likely
  - Characterize sensors
    - When is a notated G sensed as a G#?
- Machine learning necessary for good performance

Carnegie Mellon

Vocal Accompaniment

17

Carnegie Mellon

How It Works

- Score position modeled as a probability density function
- Bayesian update rule: $P(s|o) \propto P(o|s)P(s)$
- P(o|s) is e.g. "probability of observing pitch G if the score says play an A." Simple statistics on labeled training data.
- Prior P(s) by fast *convolution* with a log normal (describes tempo and tempo variation)

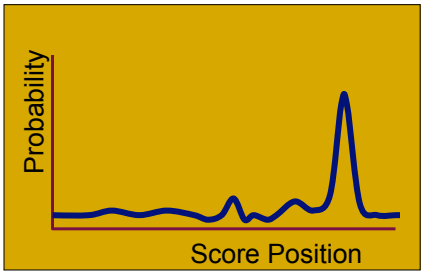

18

Carnegie Mellon

Commercial Implementation

rtsp://qt.partner-streaming.com/makemusic/wm_03_1.mov rtsp://qt.partner-streaming.com/makemusic/wm_04_1.mov

Carnegie Mellon

Style Classification: Listening to Jazz Styles

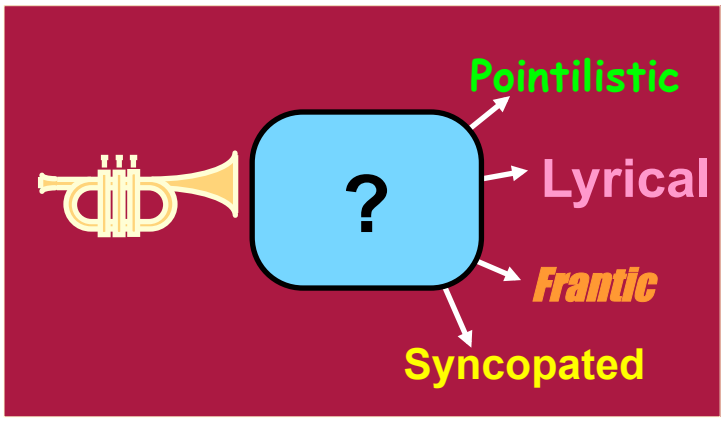

Carnegie Mellon

Jazz Style Recognition

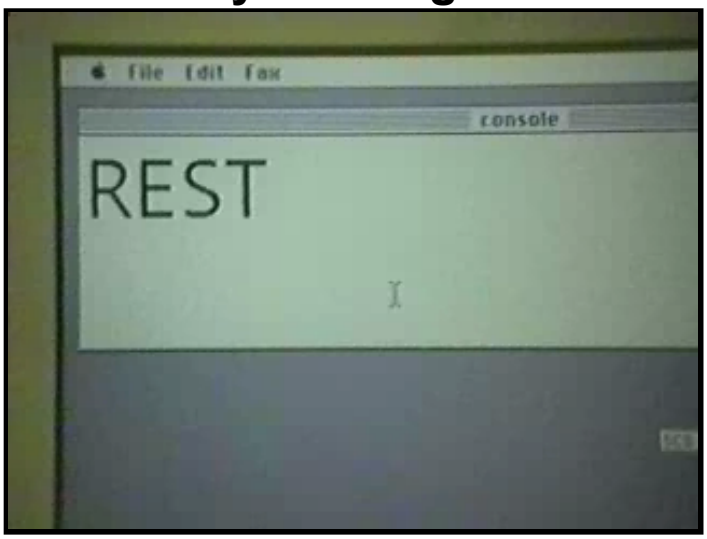

Carnegie Mellon

Techniques

- Extract features from audio:
  - Note density
  - Mean & Std. Dev. of pitch range
  - Mean & Std. Dev. of pitch intervals
  - Silence vs. Sounding ("duty factor")
  - ... and many more
- Features over 5-second windows
- Standard Classifiers (Naive Bayes, Linear, Neural Net)

Carnegie Mellon

Polyphonic Audio-to-Score Alignment

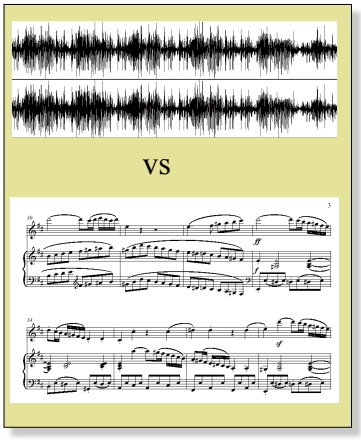

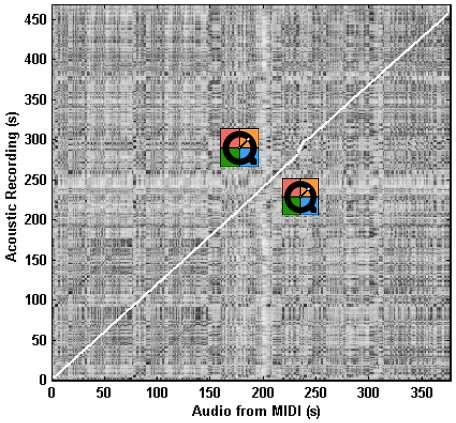

Carnegie Mellon

Audacity Editor with Automatic Audio-to-MIDI Alignment

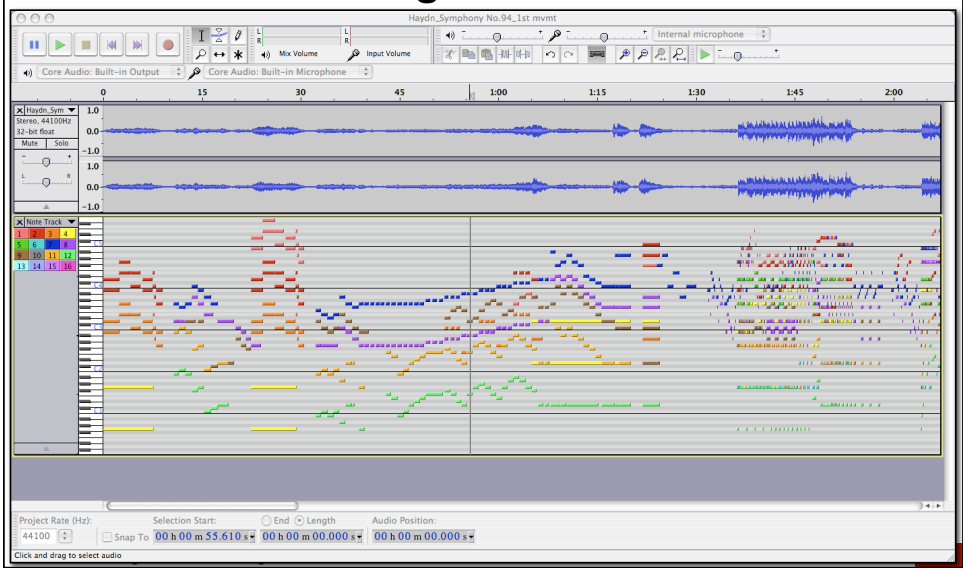

Carnegie Mellon

Finding Note Onsets
(How to segment music audio into notes.)

- Not all attacks are clean
- Slurs do not have obvious (or fast) transitions
- We can use score alignment to get a rough idea of where the notes are (~1/10 second)
- Then, machine learning can create programs that do an even better job (bootstrap learning).

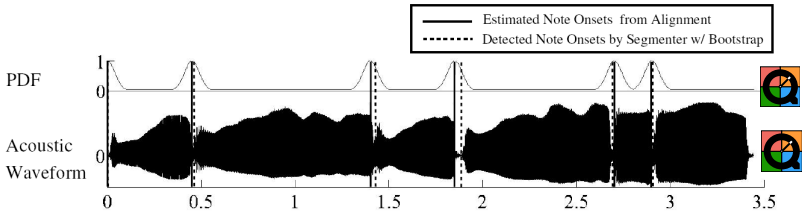

Carnegie Mellon

Expressive Performance

Carnegie Mellon

Phrase-based Synthesis

Note-by-Note Synthesis

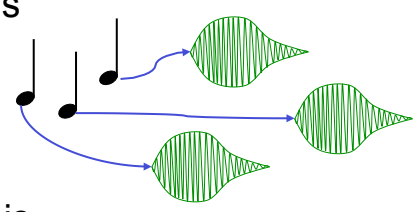

Phrase-based Synthesis

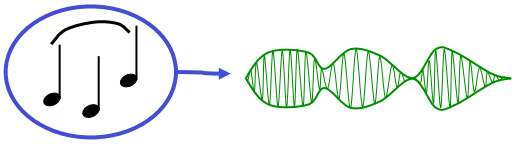

Carnegie Mellon

Example Envelopes

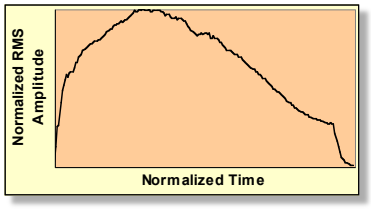

Tongued Note

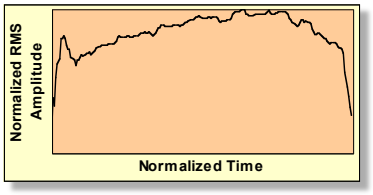

Slurred Note

Carnegie Mellon

Synthesis Examples

- Good trumpet sounds, mechanically performed: 
- Same sounds, but performed with AI-based model of trumpet performance: 
- Another example: 
- Trumpet example from Ning Hu's thesis: 
- Bassoon example from Ning Hu's thesis: 

Carnegie Mellon

Overview

- Computer Music Introduction
- How Is Computation Used in Music Today?
- New Capabilities:
  What Can Computers Do Tomorrow?

- What Will Music Be Like in the Future?

Carnegie Mellon

Human Computer Music Performance

- The most advanced computer music research is applied to esoteric art music.
  - There is a widespread practice of interactive computer (art) music
  - … but relatively little sophistication in popular music
- OPPORTUNITY
  - State-of-the-art computer music systems for popular music performance
  - *Autonomous Intelligent Machine Musicians*

 31

Carnegie Mellon

Example

- Suppose you want to get together and play music ... BUT, you're missing a *bass* player.

credit: Green Day

 32

Carnegie Mellon

What Research Is Needed?

- Synchronization
  - Signal processing
  - Machine learning
  - Human interface
- Sketchy notation
  - Representation issues
- Improvisation
  - Models of style
- Sound Production
  - Phrase-based synthesis?
- Modularity/Systems issues
  - Real-time systems
  - Software architecture
- Interaction
  - HCI

Carnegie Mellon

Is There a Market? What's the Impact?

- $8B annual US music sales
  - ***Excluding*** recordings, education, performances
- 5 million musical instruments per year
- Performance revenue is on the order of $10B
- Recording revenue is similar; order of $10B
- Approximately 1/2 of all US households have a practicing musician
- ... so *very roughly* $10+B and 100M people!

Carnegie Mellon

Rock Prodigy

- Guitar Hero for Real Guitars
- Game design, content, animation, etc. by others
- (Play Video)
- Unsolicited comment: "The best part about it is polyphonic pitch detection"

Carnegie Mellon

An Example

Carnegie Mellon

Online, collaborative development of creative content is already here…

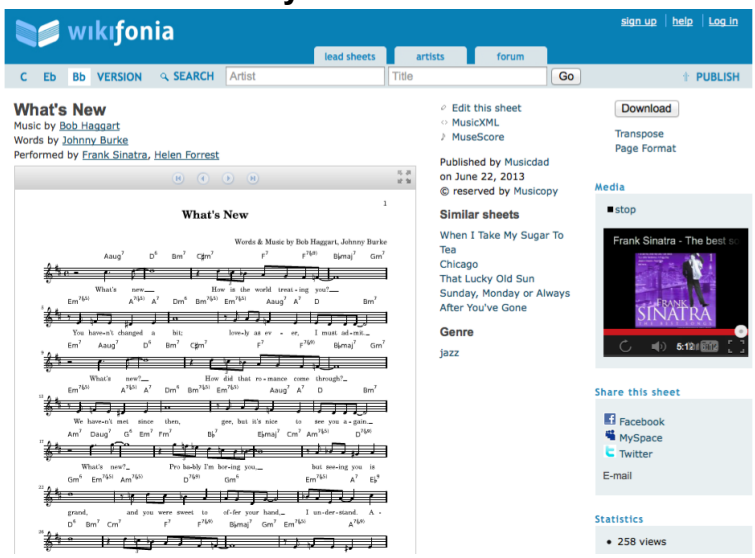

Carnegie Mellon

What Will People Do With HCMP?

- Practice with virtual bands.
- Create their own arrangements.
- Post machine-readable music online, share.
- Blend conventional performance with algorithmic composition, new sounds, new music.
- Robot performers.
- Eventually ... new art forms
- Think of the electric guitar, drum machine in music, camera in visual art, ...

Carnegie Mellon

Another Example

Carnegie Mellon

Conclusion

- Music Understanding and Human Computer Music Performance will *enrich musical experiences* for millions of people, including both amateurs and professionals.
- If we build computers that can perform popular music interactively with intelligence, great music will be made. *That is the future of music performance*.