Music Understanding and the Future of Music

Roger B. Dannenberg

Professor of Computer Science, Art, and Music

Carnegie Mellon University

Why Computers and Music?

- Music in every human society!
- Computers are everywhere!
- Music is technological!
- Computing can make music:
  - More Fun
  - More Available
  - Higher Quality
  - More Personal

My Background

- Always interested in math and music and making things
- Trumpet player since age 11
- Discovered synthesizers in high school
- Discovered computers about the same time
- Discovered computer music in college
- Musical Acoustics with Art Benade
- Research motivated by musical experience:
  - Computers as performers
  - Expressive programming languages for music
  - Audacity Audio Editor (co-creator)

Overview

- Introduction

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

How Is Computation Used in Music Today?

Indabamusic.com

making-music.blogspot.com

http://venturebeat.com/

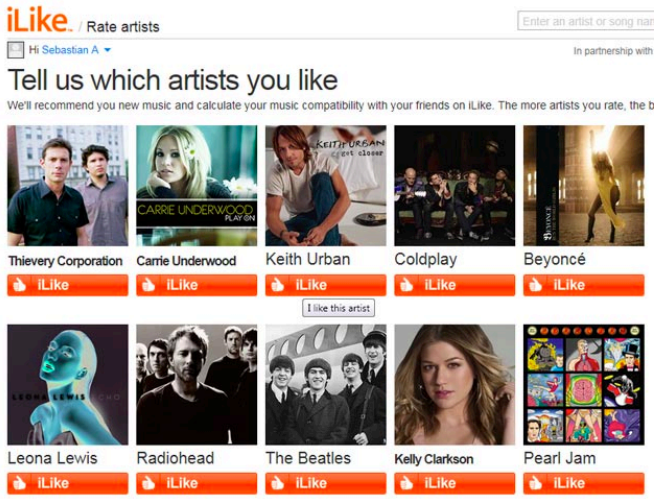

Music Computation Today

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

Music Computation Today

Musical Instruments: synthesizers and controllers

Synthesizer (Solaris)

Drum Machine (Yamaha)

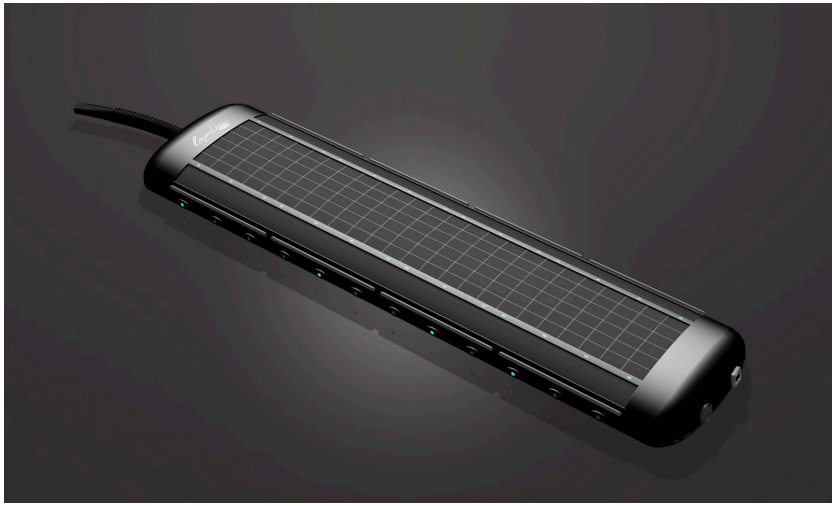

Linnstrument (Roger Linn)

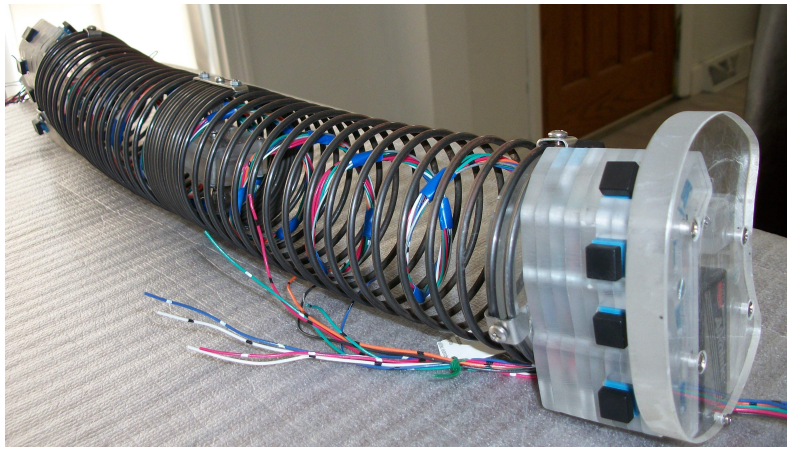

Sonic Spring (Tomas Henriques)

Music Computation Today

Distribution: compression, storage, networks

Napster

Apple iPod

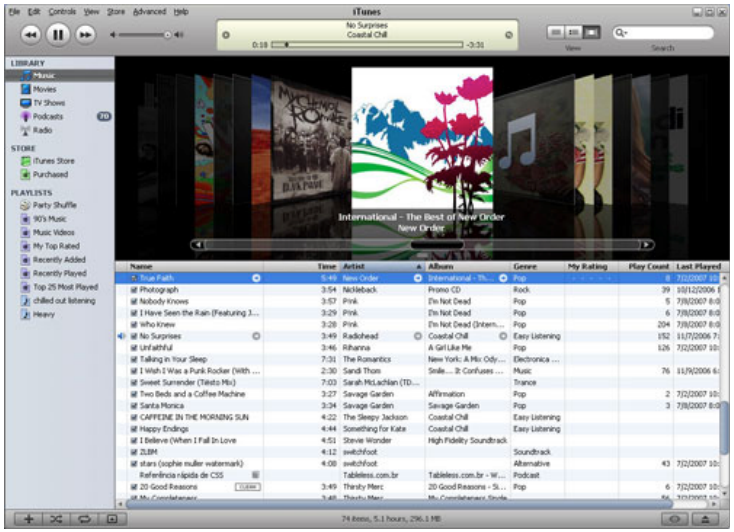

Apple iTunes

Music Computation Today

Search, recommendation, music fingerprinting

Google Music China

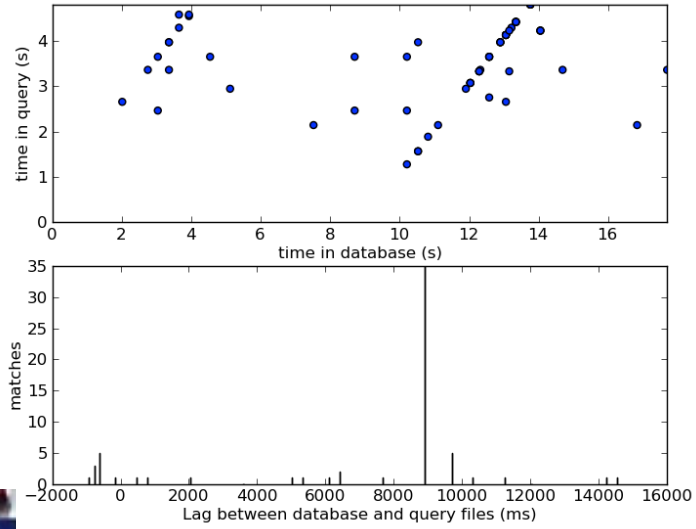

Music Fingerprinting

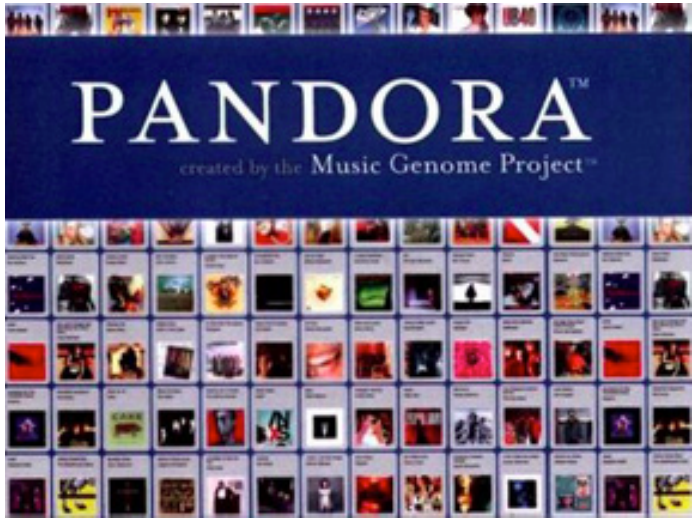

Pandora
Music Recommendation

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?

New Capabilities: What Can Computers Do Tomorrow?

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

Accompaniment Video

Computer Accompaniment

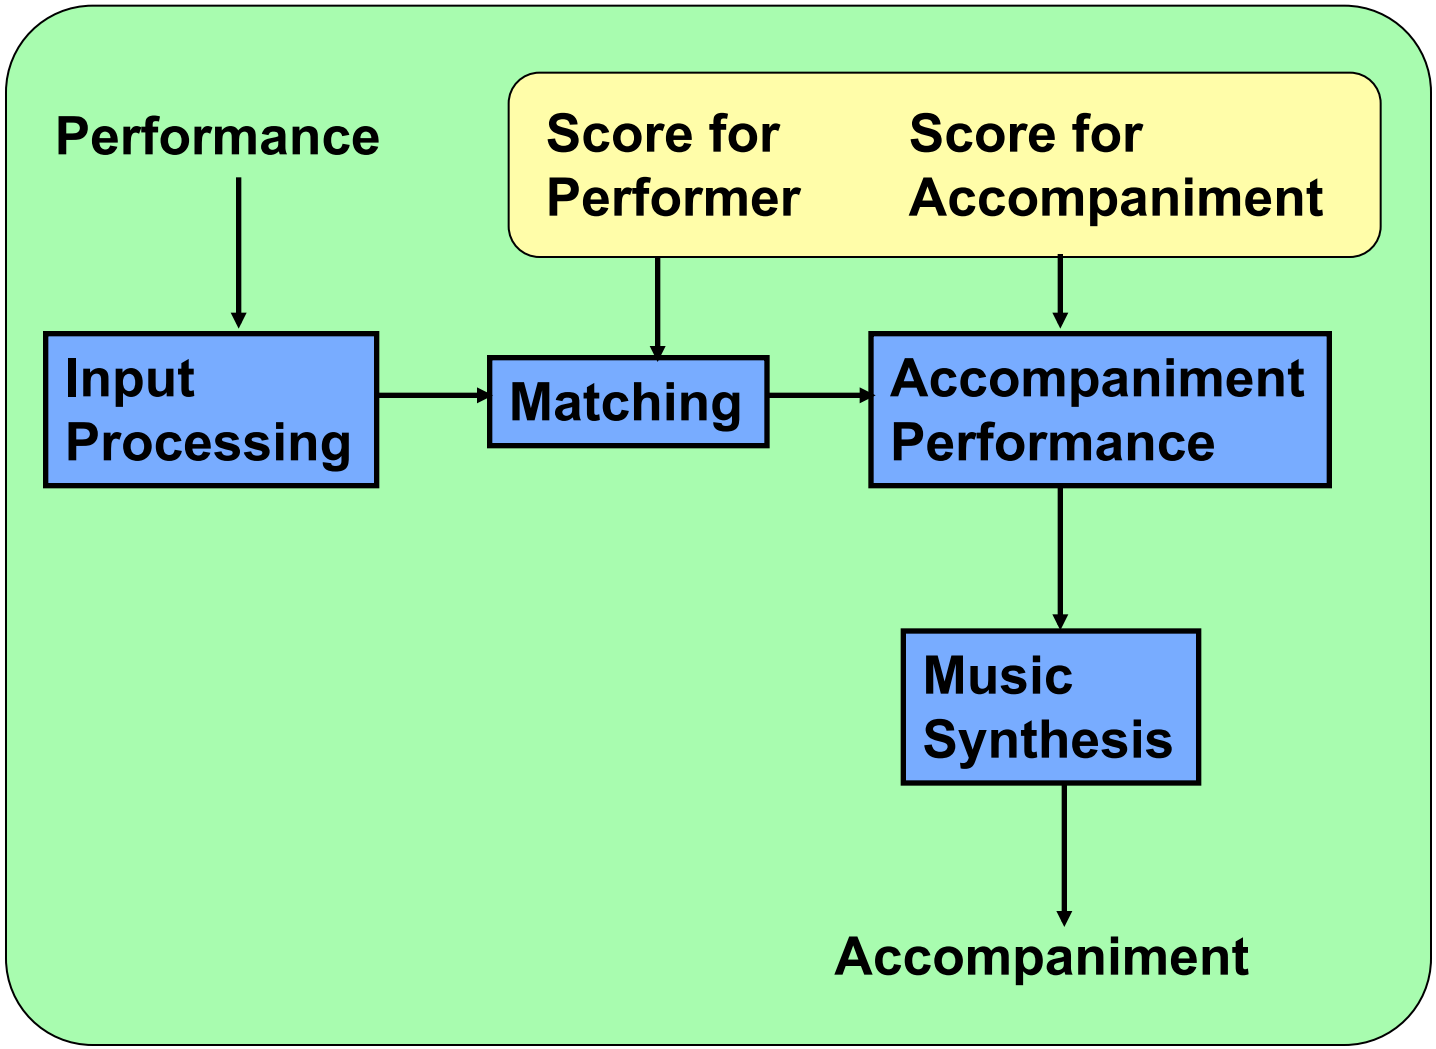

Computer Accompaniment

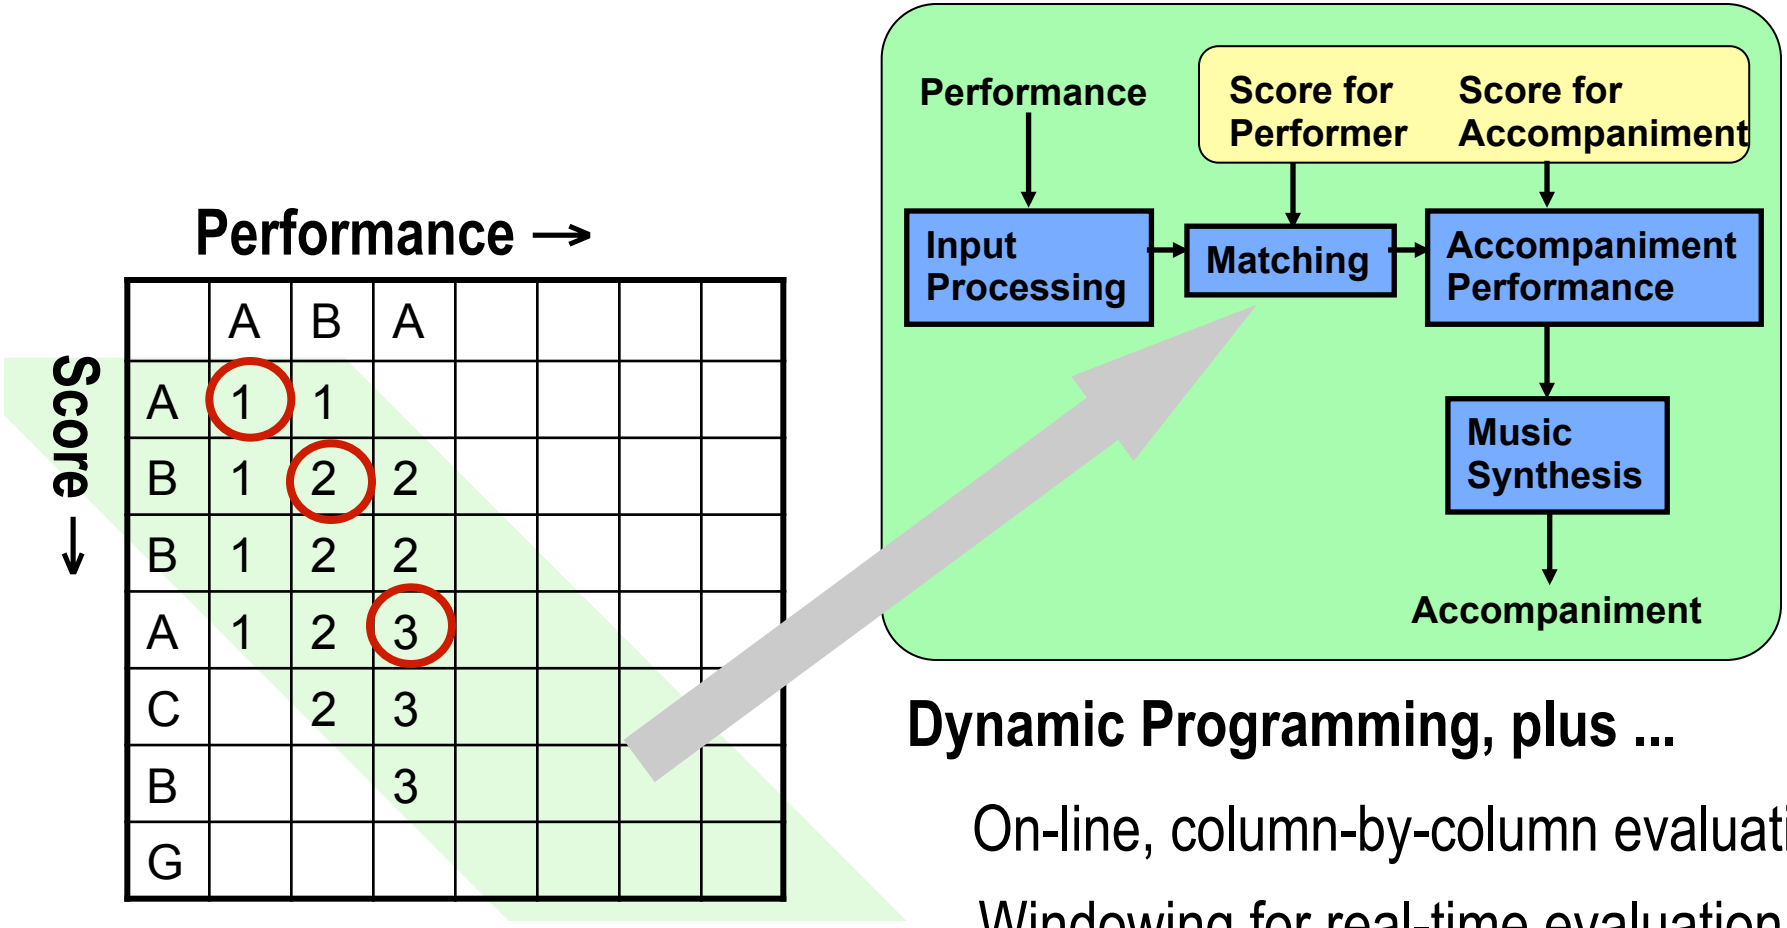

Dynamic Programming, plus ...

On-line, column-by-column evaluation

Windowing for real-time evaluation

Heuristics for best-yet matching

Penalty for skipping notes

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.

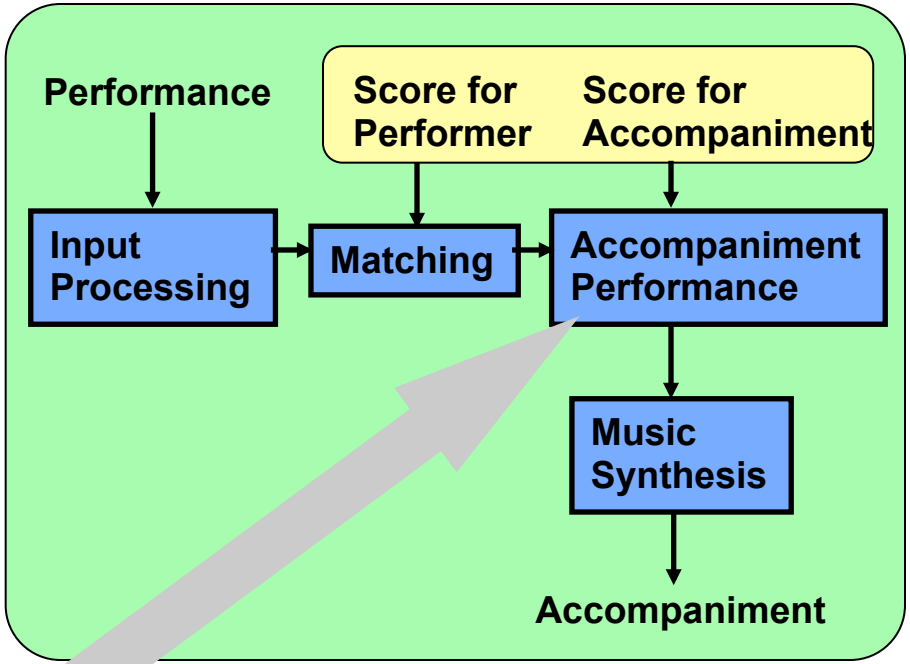

Vocal Accompaniment

- 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

Vocal Accompaniment

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)

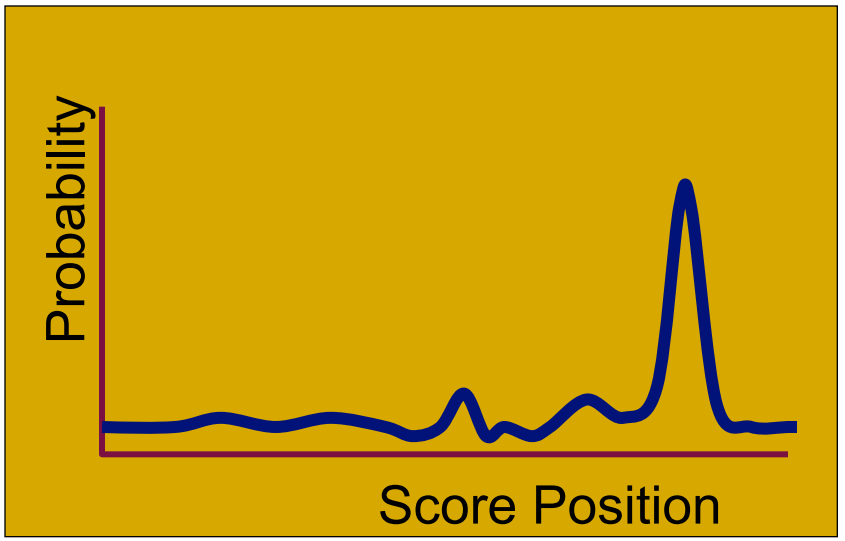

Commercial Implementation

rtsp://qt.partner-streaming.com/makemusic/wm_03_l.mov rtsp://qt.partner-streaming.com/makemusic/wm_04_l.mov

Style Classification: Listening to Jazz Styles

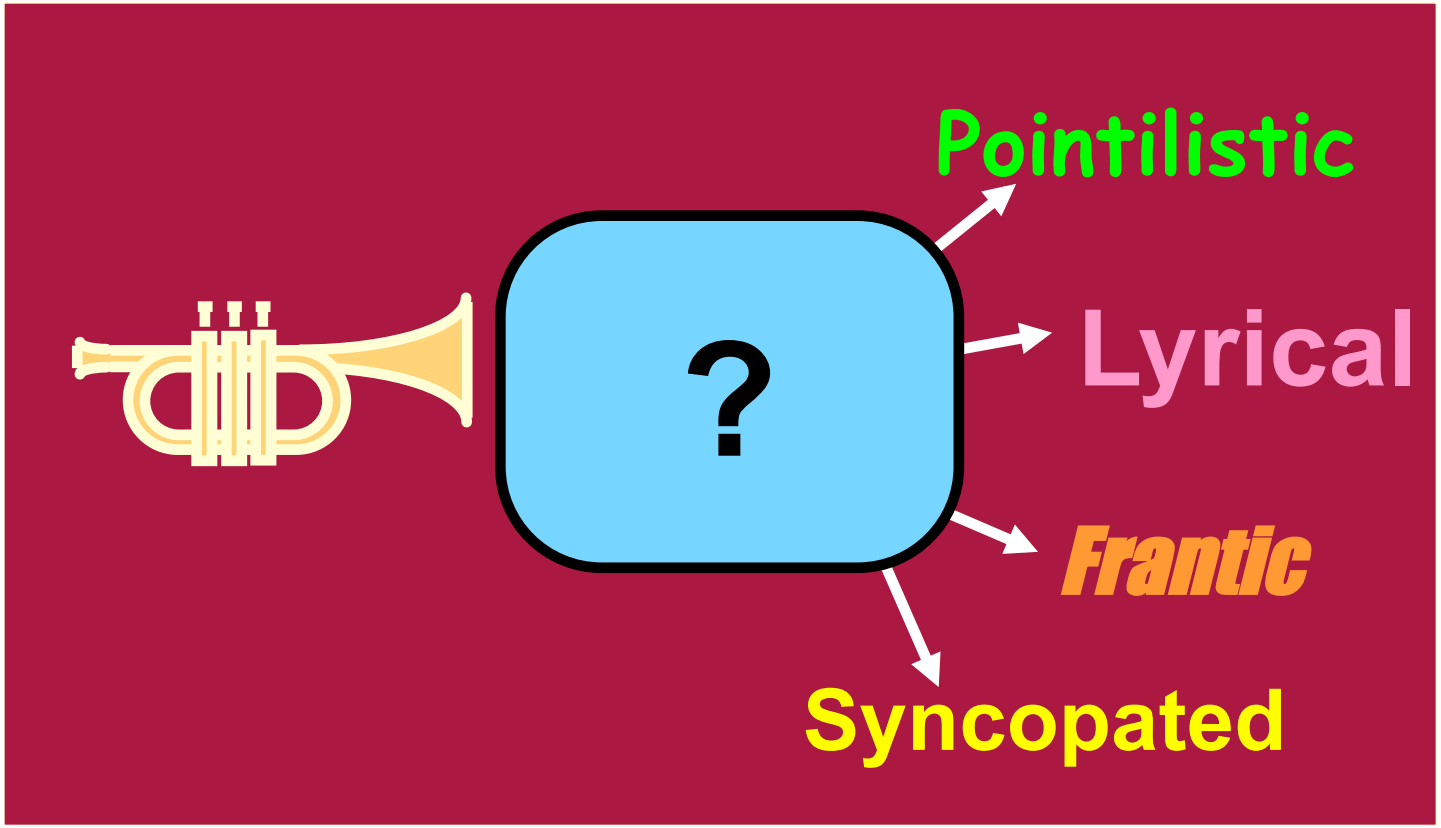

Jazz Style Recognition

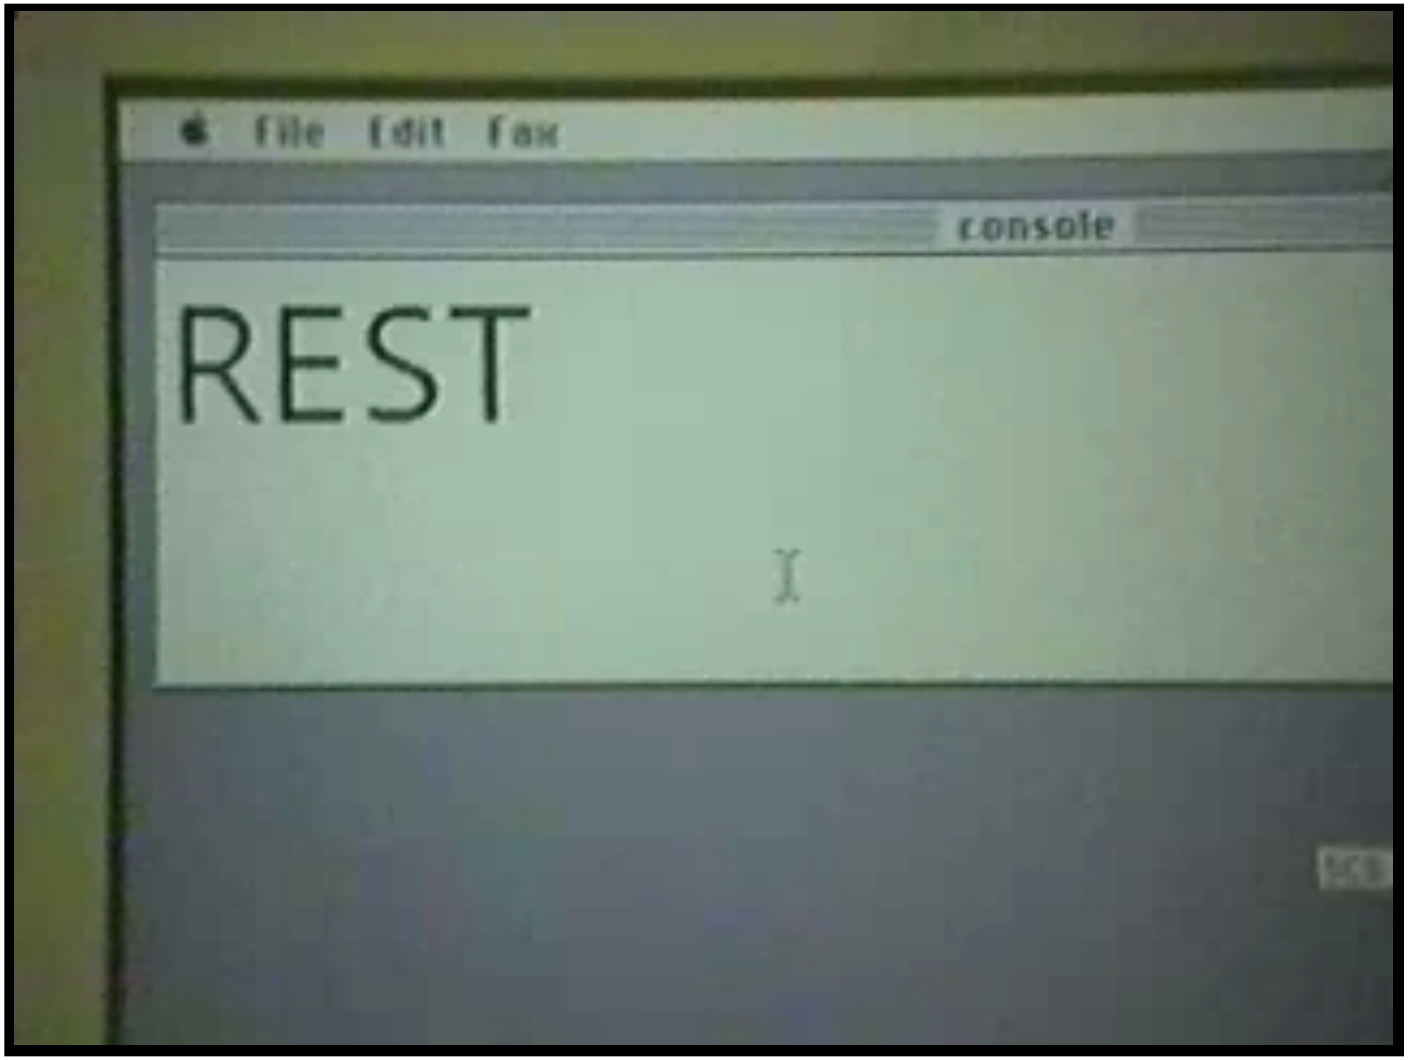

Techniques

- 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)

Polyphonic Audio-to-Score Alignment

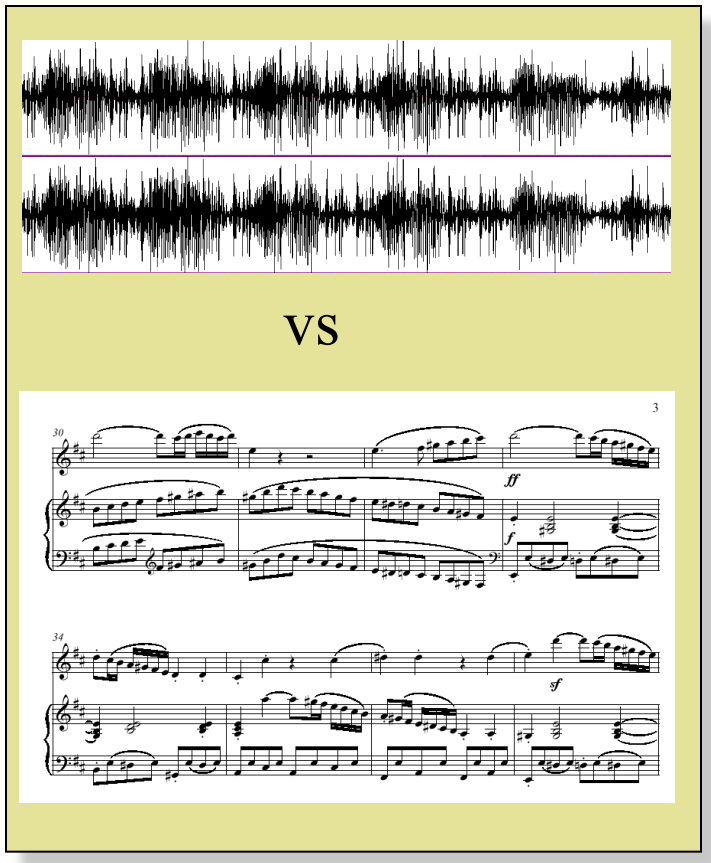

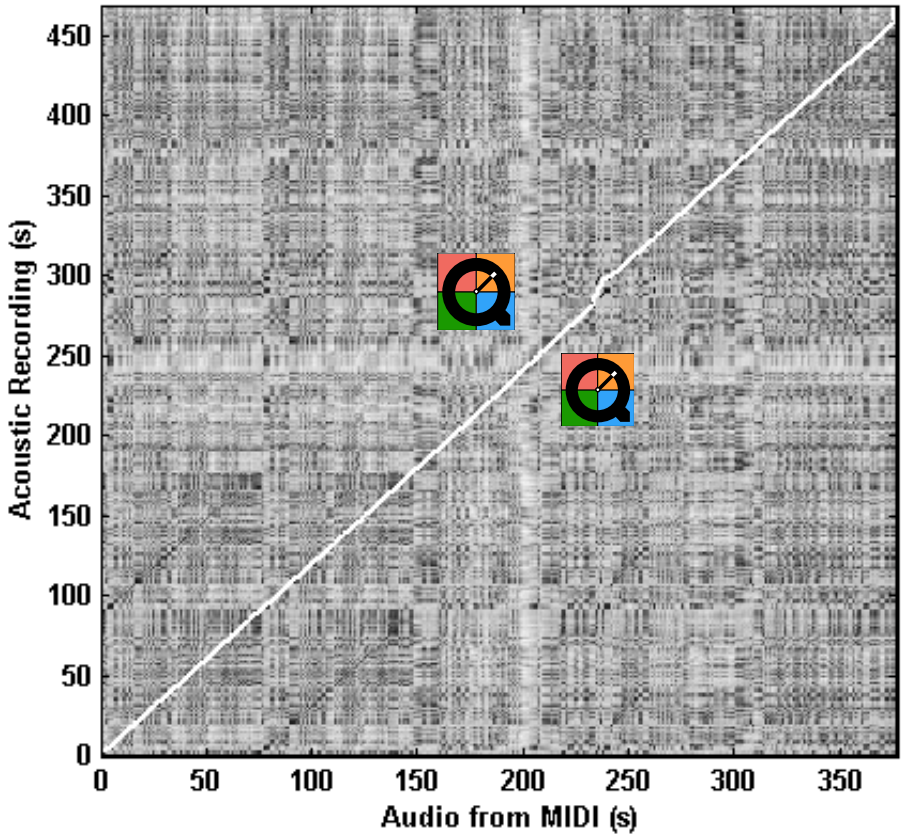

Audacity Editor with Automatic Audio-to-MIDI Alignment

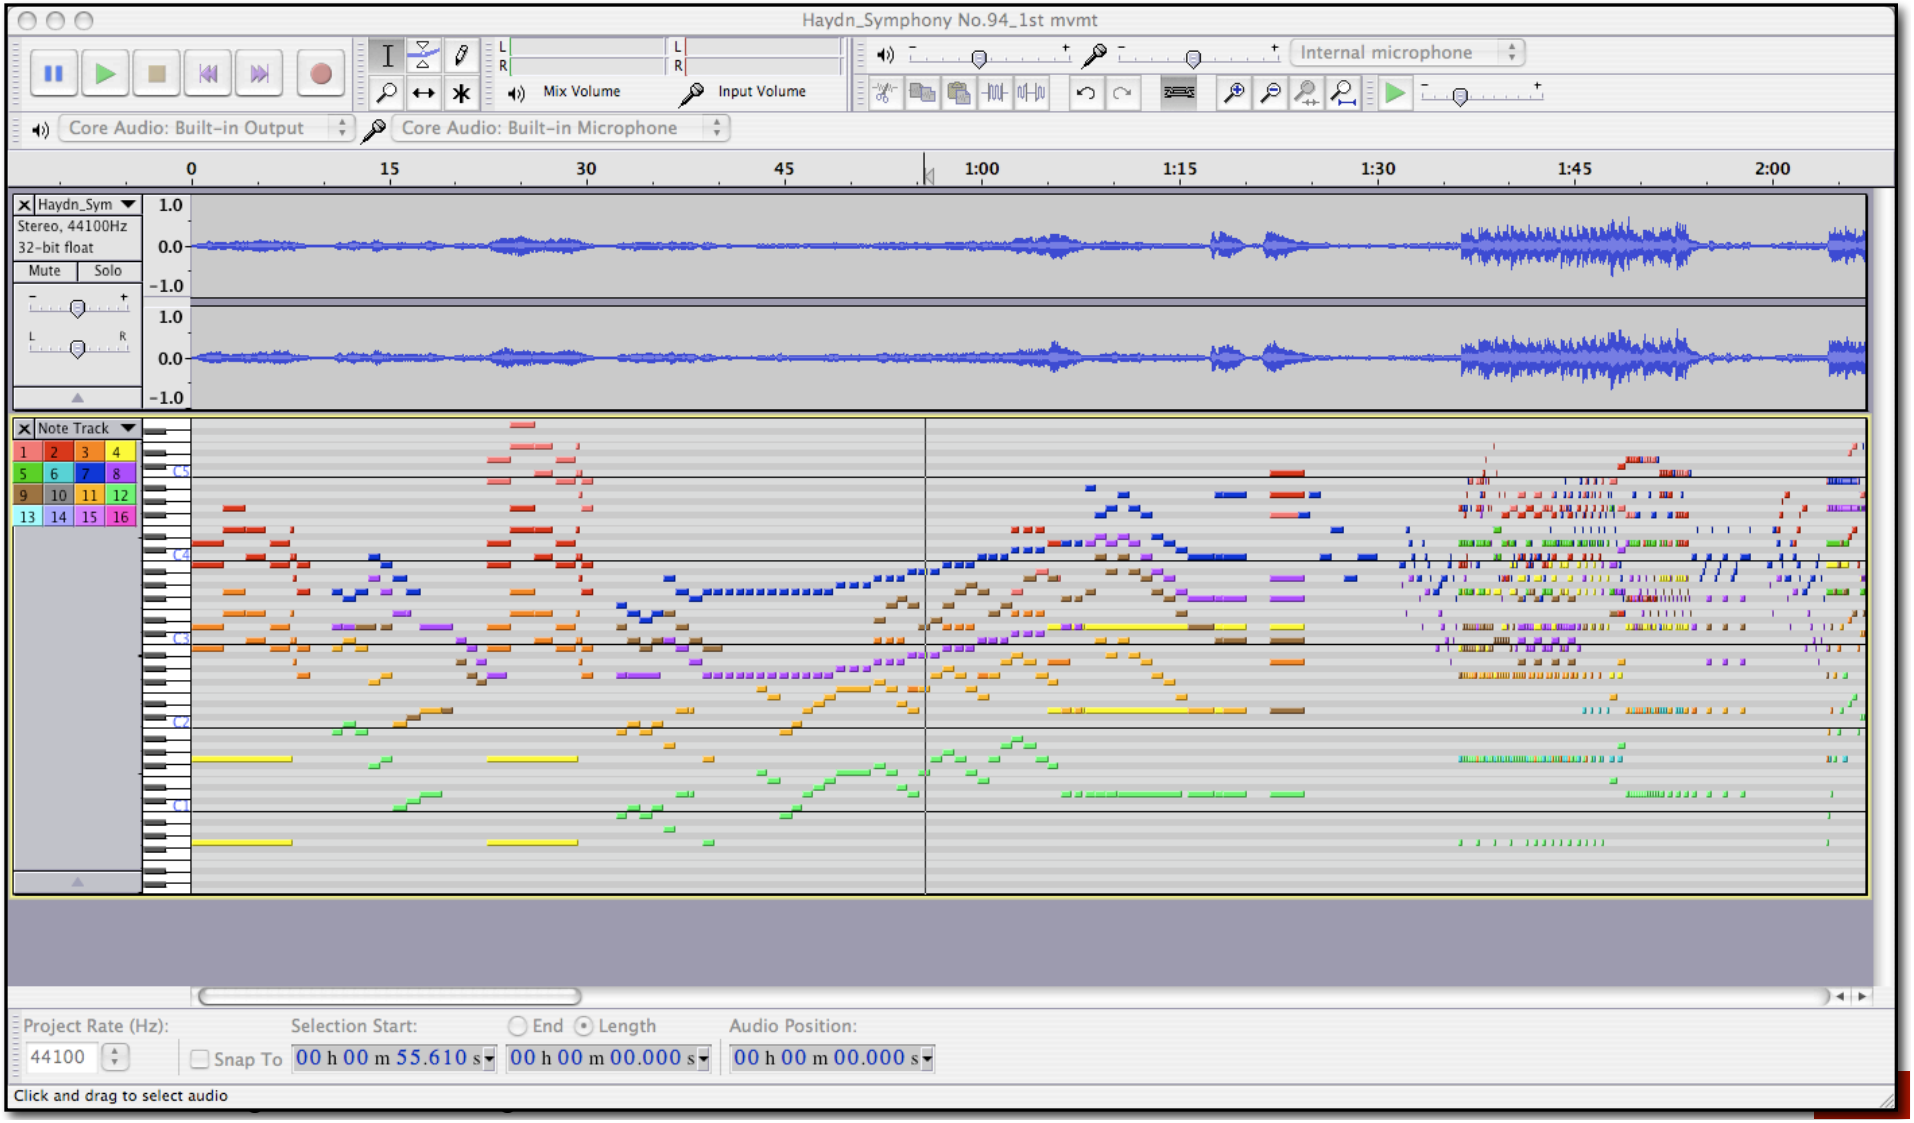

Intelligent Audio Editor

- This excerpt is included in the audio examples:

- Before:  After: 

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).

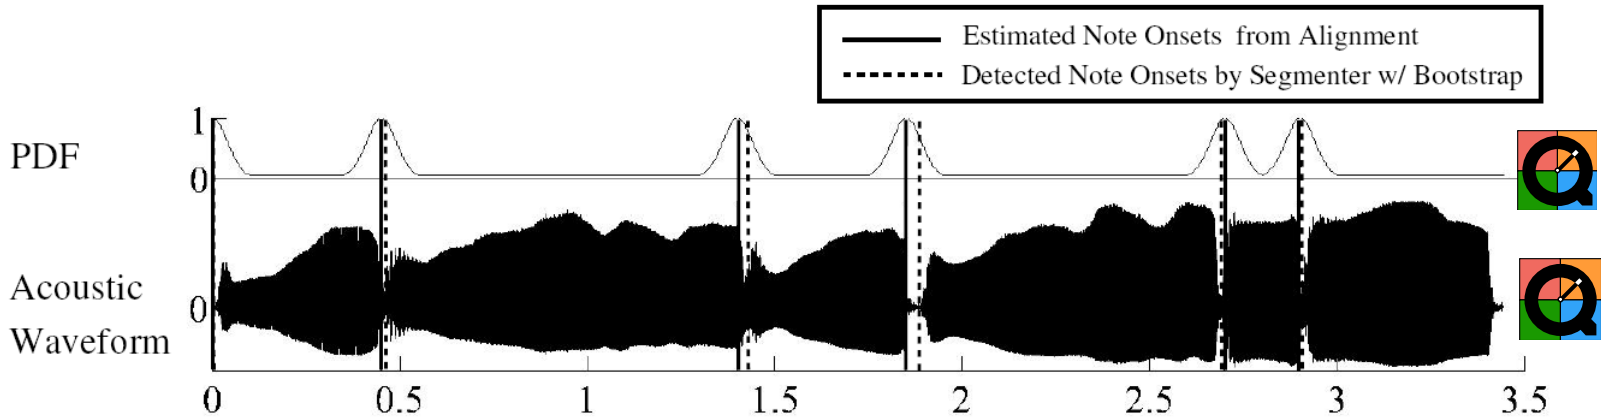

Expressive Performance

Phrase-based Synthesis

Note-by-Note Synthesis

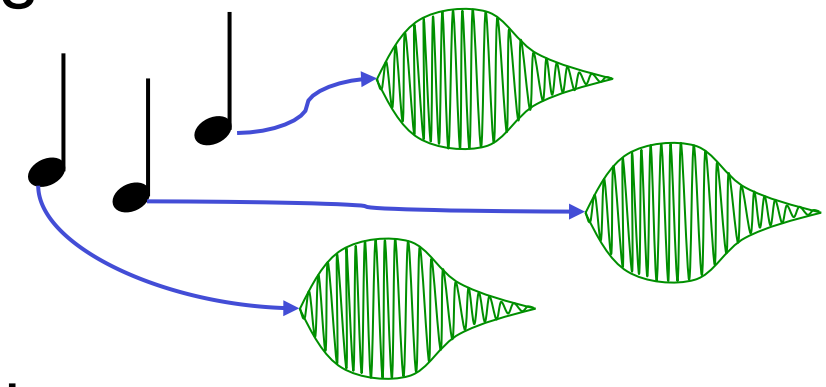

Phrase-based Synthesis

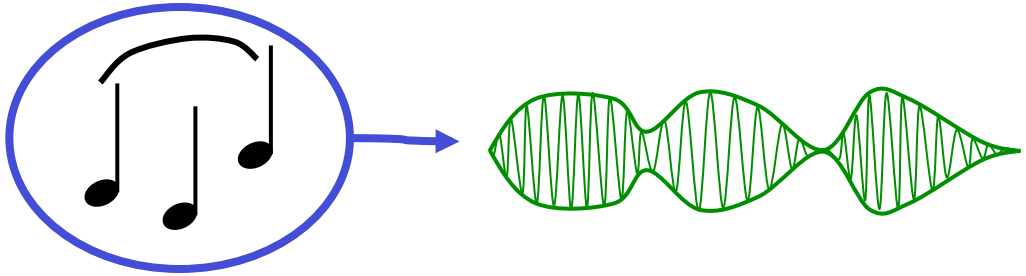

Example Envelopes

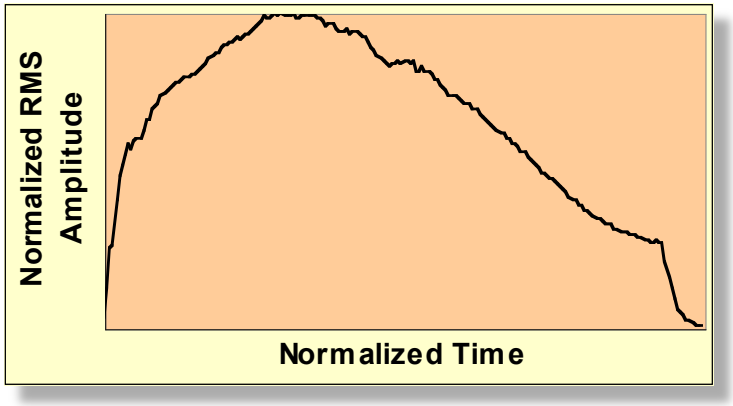

Tongued Note

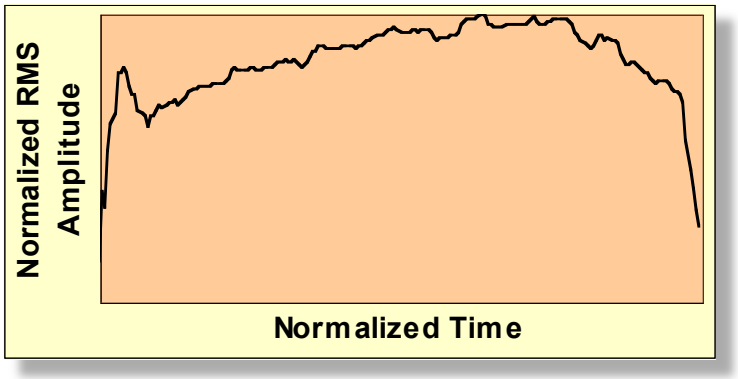

Slurred Note

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: 

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?

Human Computer Music Performance

- OPPORTUNITY
  - State-of-the-art computer music systems for popular music performance
  - *Autonomous Intelligent Machine Musicians*

Example

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

bass

credit: Green Day

What Research Is Needed?

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

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!

Rock Prodigy

- Guitar Hero for Real Guitars
- Game design, content, animation, etc. by others
- (Play Video)

Rock Prodigy

- Unsolicited comment: "The best part about it is polyphonic pitch detection"

An Example

Arrangements: Flexibly Adapting Music Data for Live Performance

Roger B. Dannenberg
Andrew Russell

Carnegie Mellon University

1

Another Application: Internet Drum Circle

- Latency is key:

OK

Shakers with 0.1s delay

- Can computers
  - Play drums?
  - Lead humans to keep it interesting 24x7?
  - Help keep the beat steady?

Online, collaborative development of creative content is already here…

wikifonia

sign up | help | Log in

lead sheets | artists | forum

C Eb Bb VERSION SEARCH Artist Title Go PUBLISH

What's New

Music by Bob Haggart
Words by Johnny Burke
Performed by Frank Sinatra, Helen Forrest

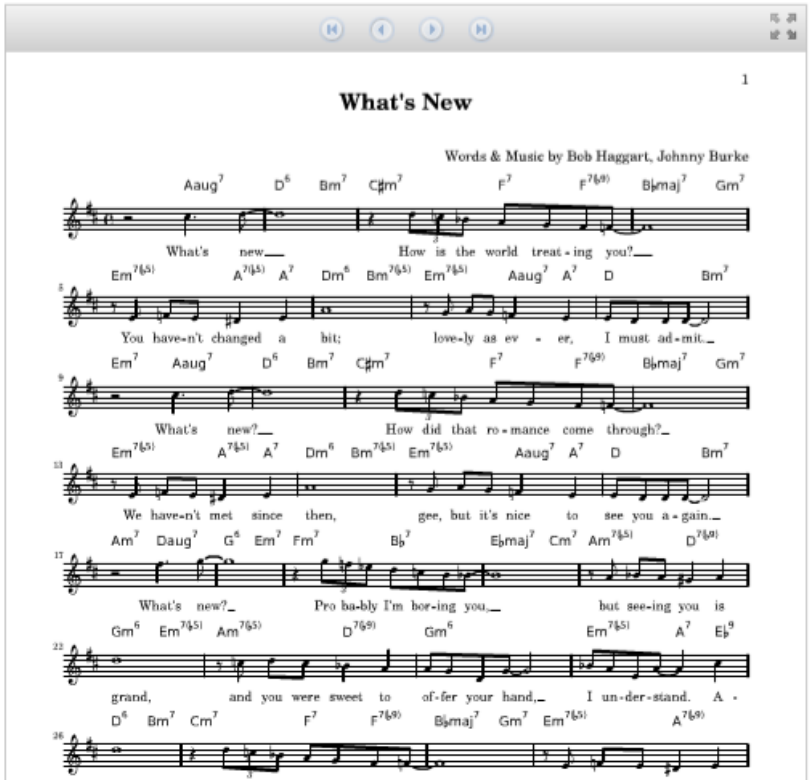

Edit this sheet
MusicXML
MuseScore

Published by Musicdad
on June 22, 2013
© reserved by Musicopy

Similar sheets

When I Take My Sugar To Tea
Chicago
That Lucky Old Sun
Sunday, Monday or Always
After You've Gone

Genre

jazz

Download

Transpose
Page Format

Media

■ stop

Share this sheet

Facebook
MySpace
Twitter

E-mail

Statistics

• 258 views

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, ...

Conclusion

- Automating 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 use understanding and intelligence to perform popular music, *great music will be made*.
  That is the future of music performance.