The State of Music at LabROSA

Dan Ellis
Laboratory for Recognition and Organization of Speech and Audio
Dept. Electrical Eng., Columbia Univ., NY USA

dpwe@ee.columbia.edu http://labrosa.ee.columbia.edu/

1. The State of LabROSA
2. Music Projects
3. The Big Picture

Lab ROSA
Laboratory for the Recognition and Organization of Speech and Audio

Columbia University
In the City of New York

1. The State of LabROSA

- LabROSA Music Publications per year

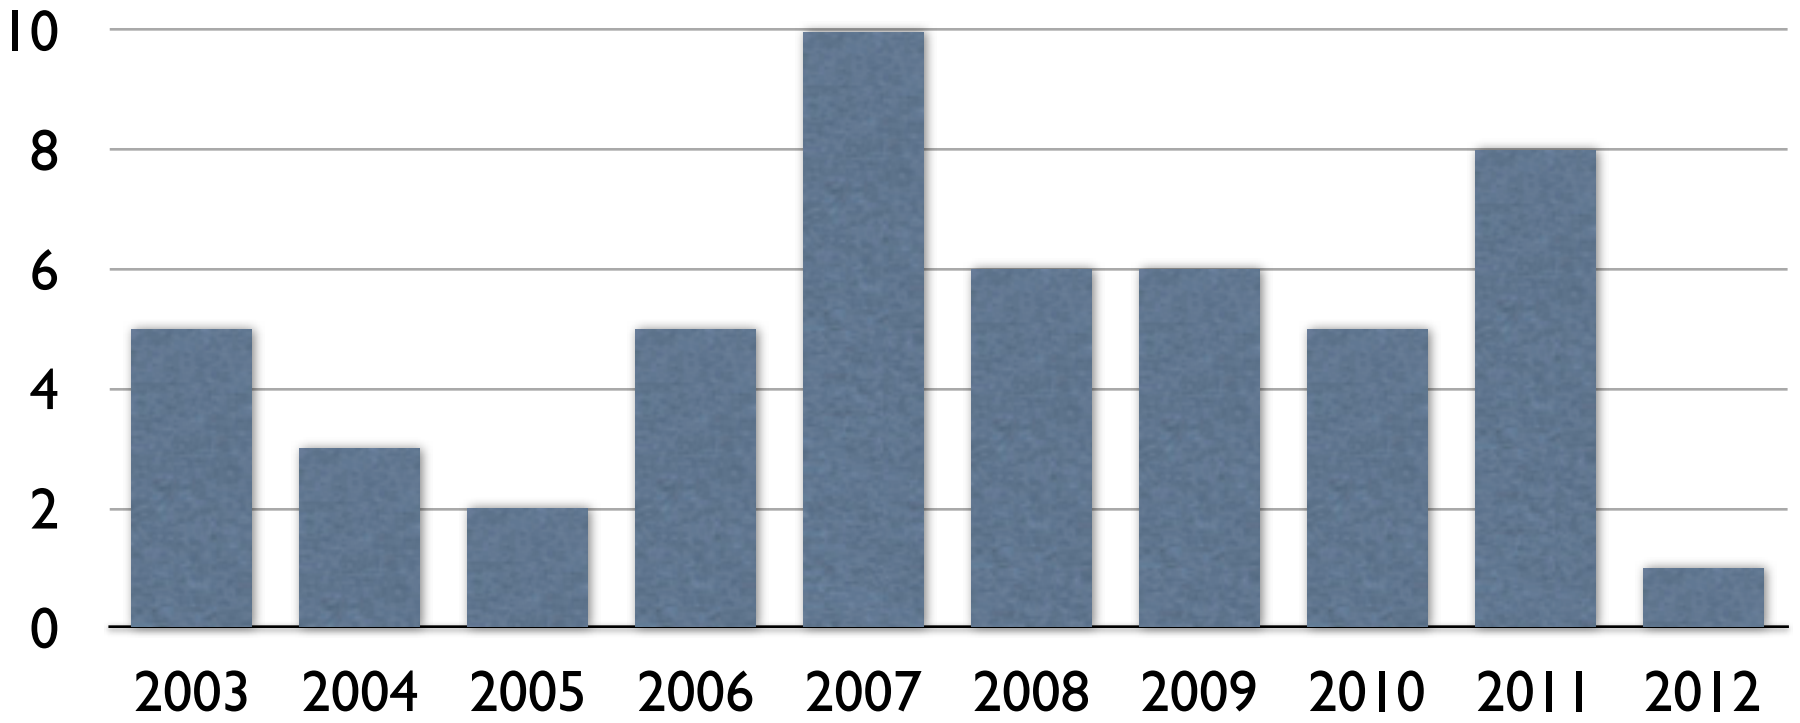

But!

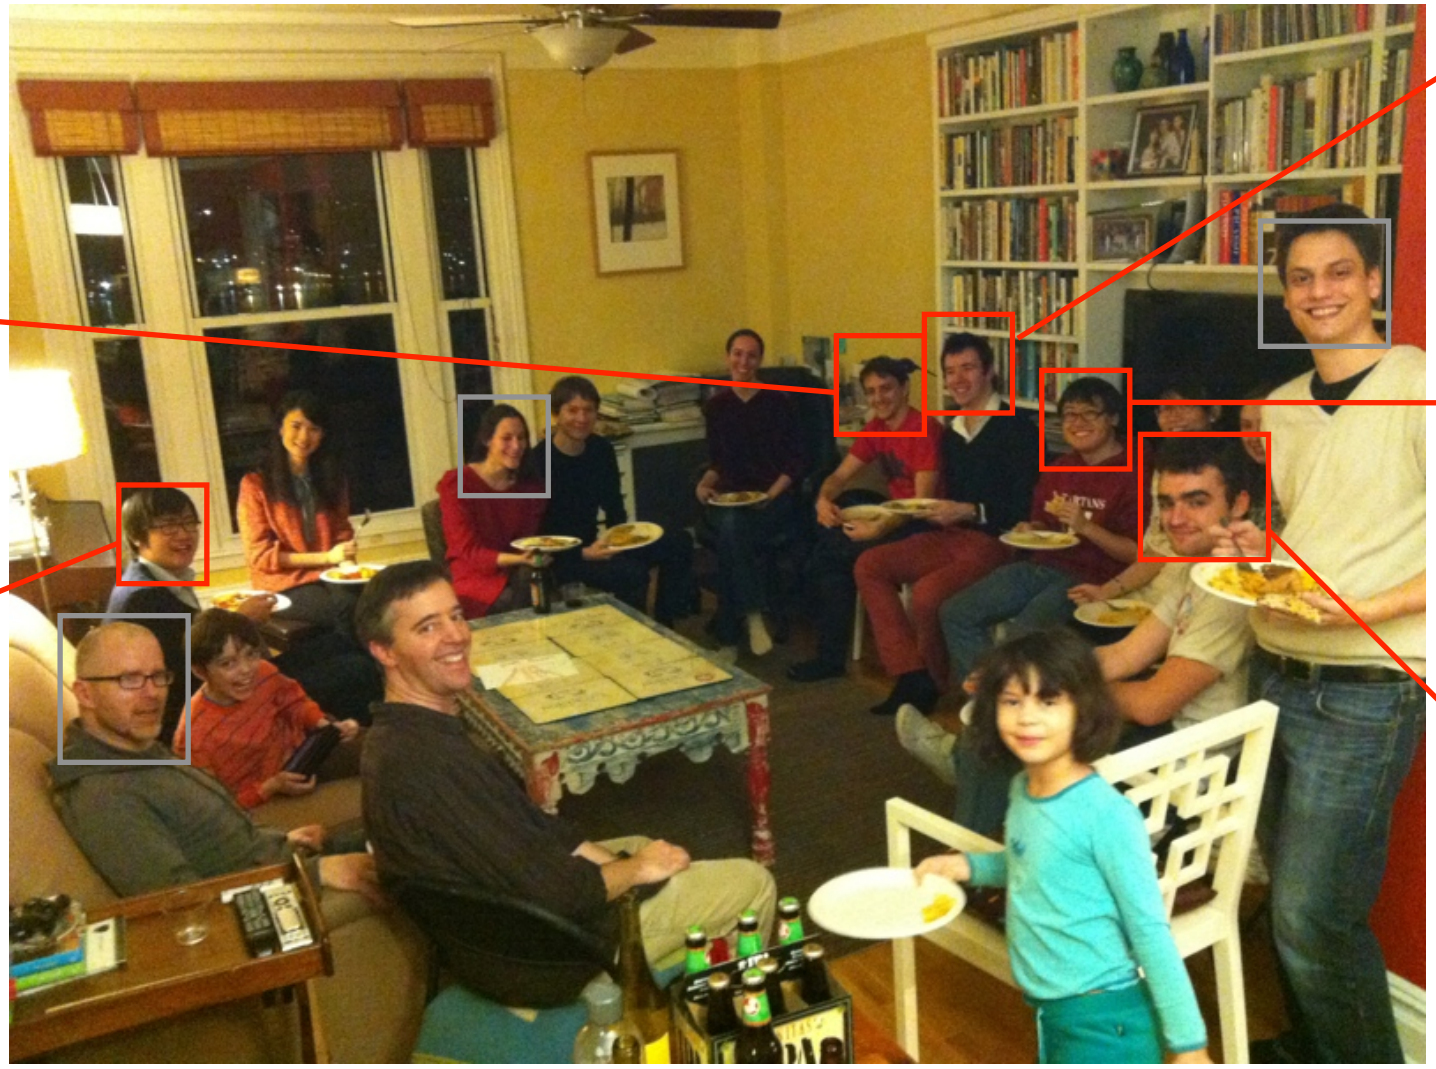
Matt McVicar
Brian McFee
Dawen Liang
Zhuo Chen
Colin Raffel

Current Projects at LabROSA

Zhuo Chen

- Parametric source-note models

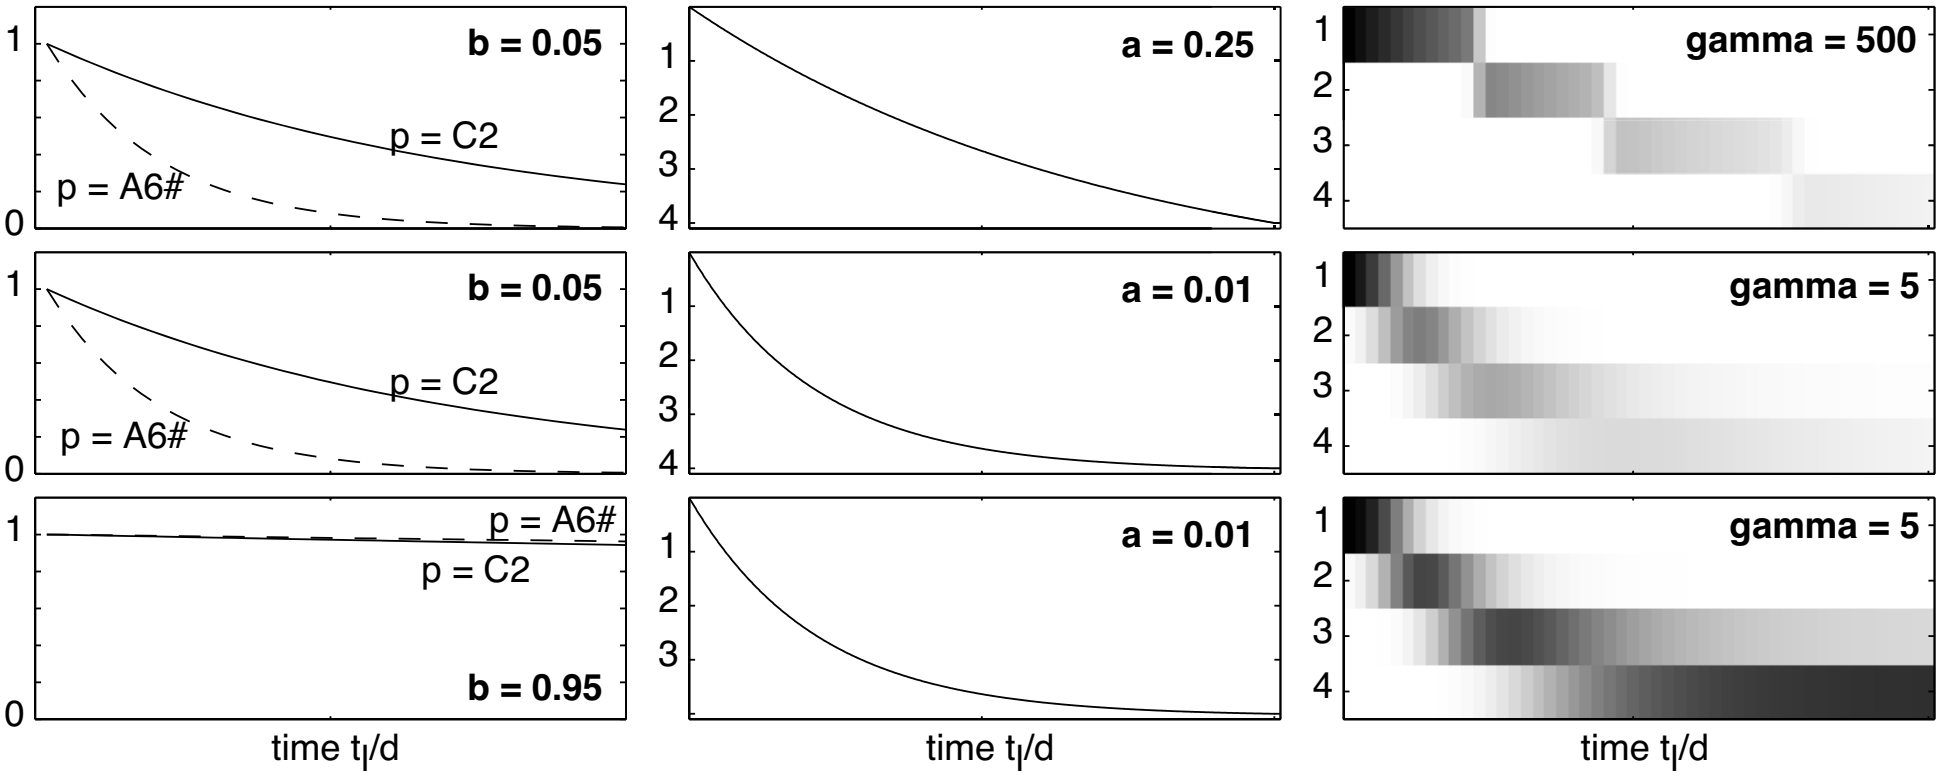

Large-Scale Cover Recognition

Bertin-Mahieux

- 2D Fourier Transform Magnitude (2DFTM)
  - fixed-size feature to capture "essence" of chromagram:

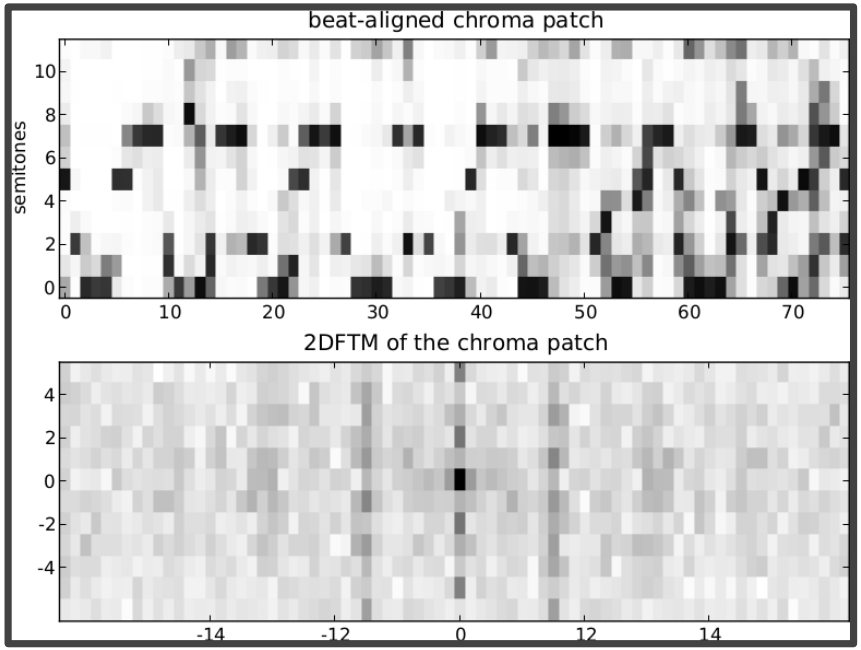

- First results on finding covers in 1M songs

	Average rank	meanAP
random	500,000	0.000
jumpcodes 2	308,369	0.002
2DFTM (50 PC)	**137,117**	**0.020**

Million Song Dataset

Bertin-Mahieux McFee

- **Many Facets**
  - Echo Nest audio features + metadata
  - Echo Nest "taste profile" user-song-listen count
  - Second Hand Song covers
  - musiXmatch lyric BoW
  - last.fm tags

- **MSD Challenge**
  - Kaggle.com competition, 153 teams
  - winning entry used simple Collaborative Filtering

Jazz Discography Project

- **How can MIR help organize jazz collections?**
  - our tools are quite genre-specific
  - e.g. beat tracker is fine for pop, useless for Jazz

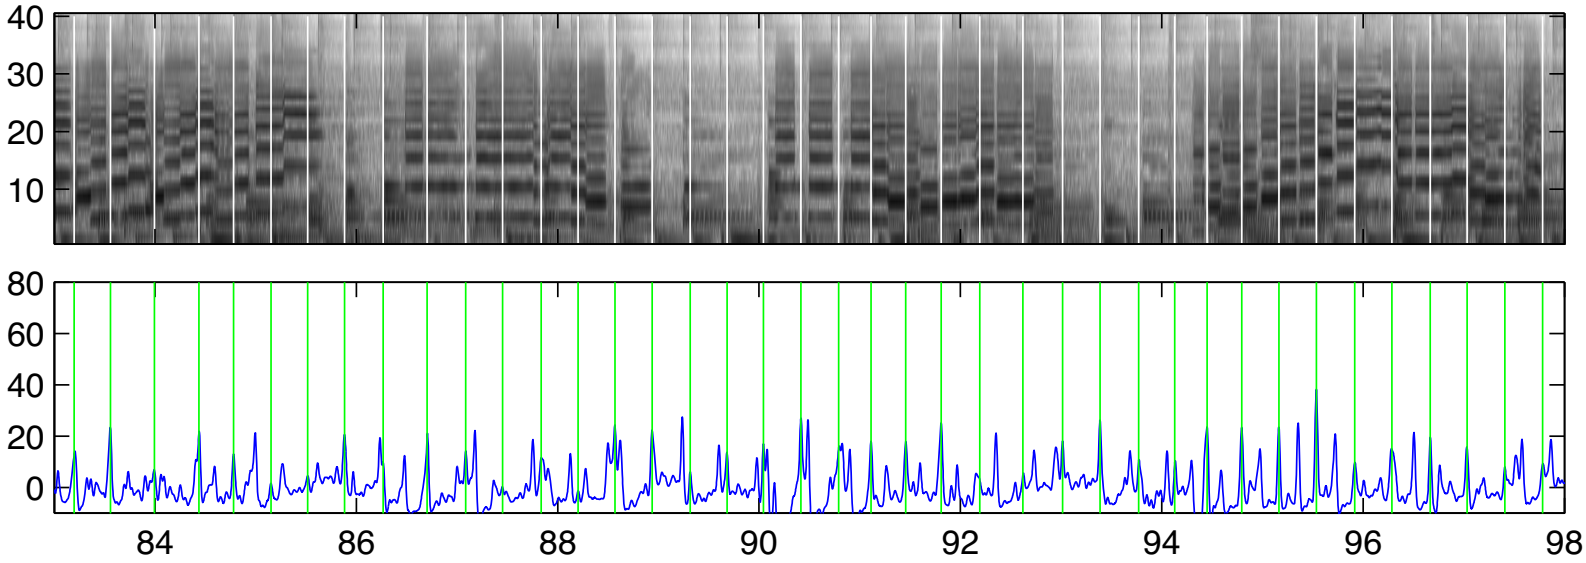

Local Tagging

- MFCC-statistics classifiers on 5 sec windows trained from MajorMiner data

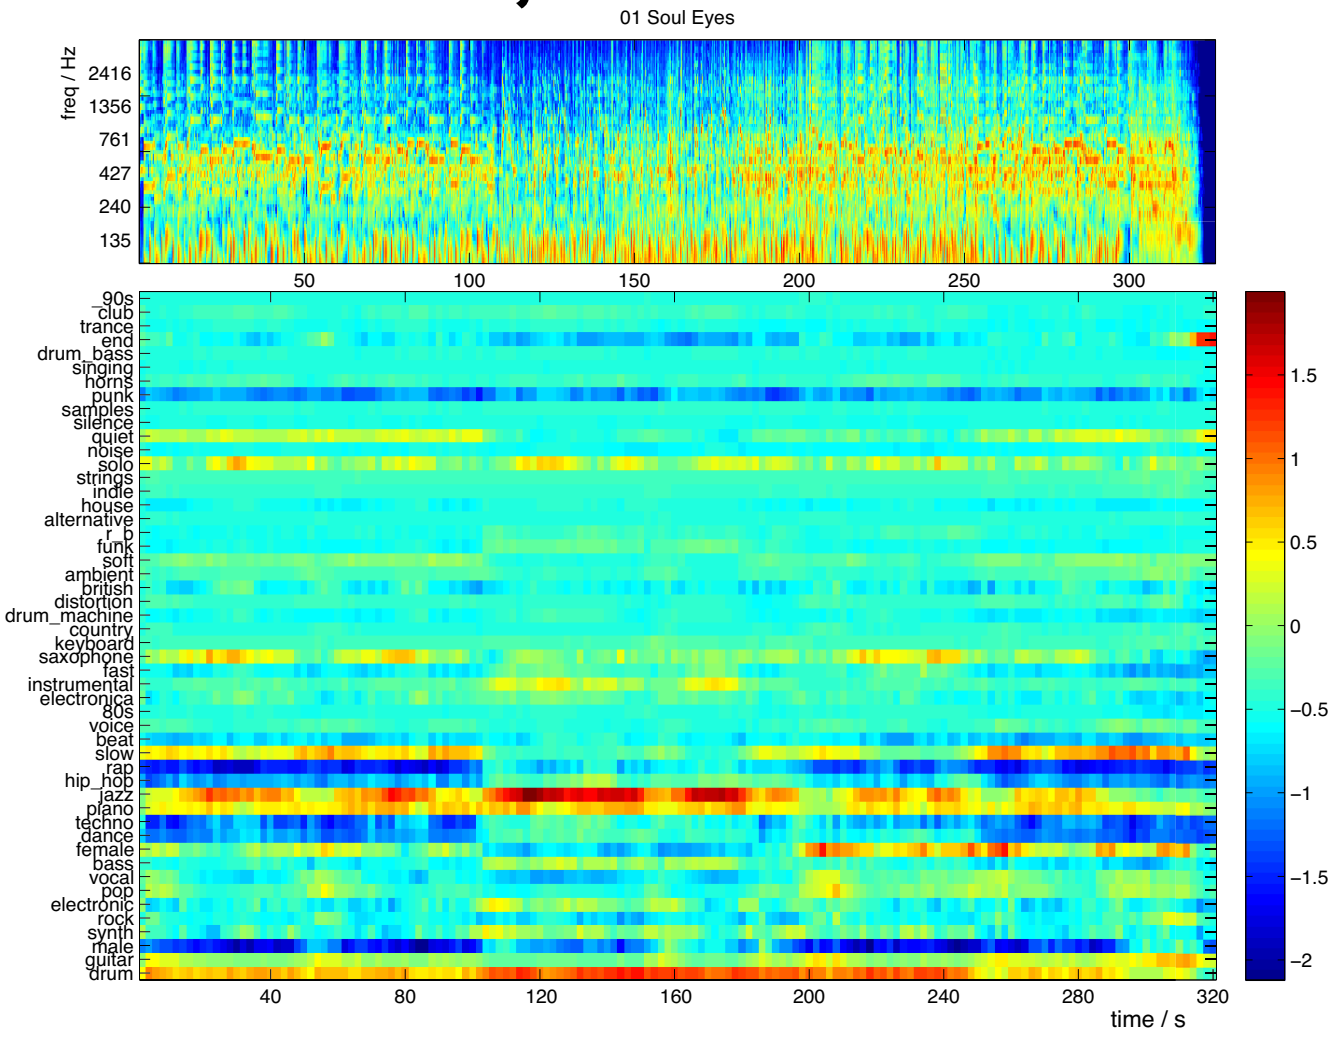

Onset Correlation

Brian McFee

- "Ahead of" or "behind" the beat?

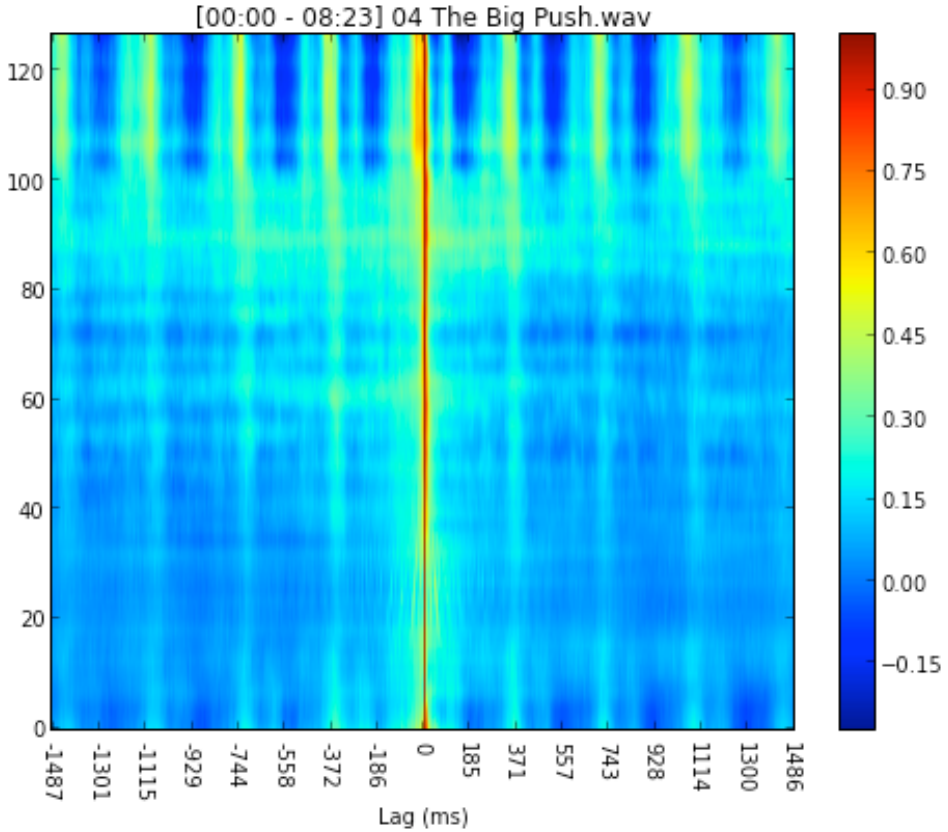

Tony Williams

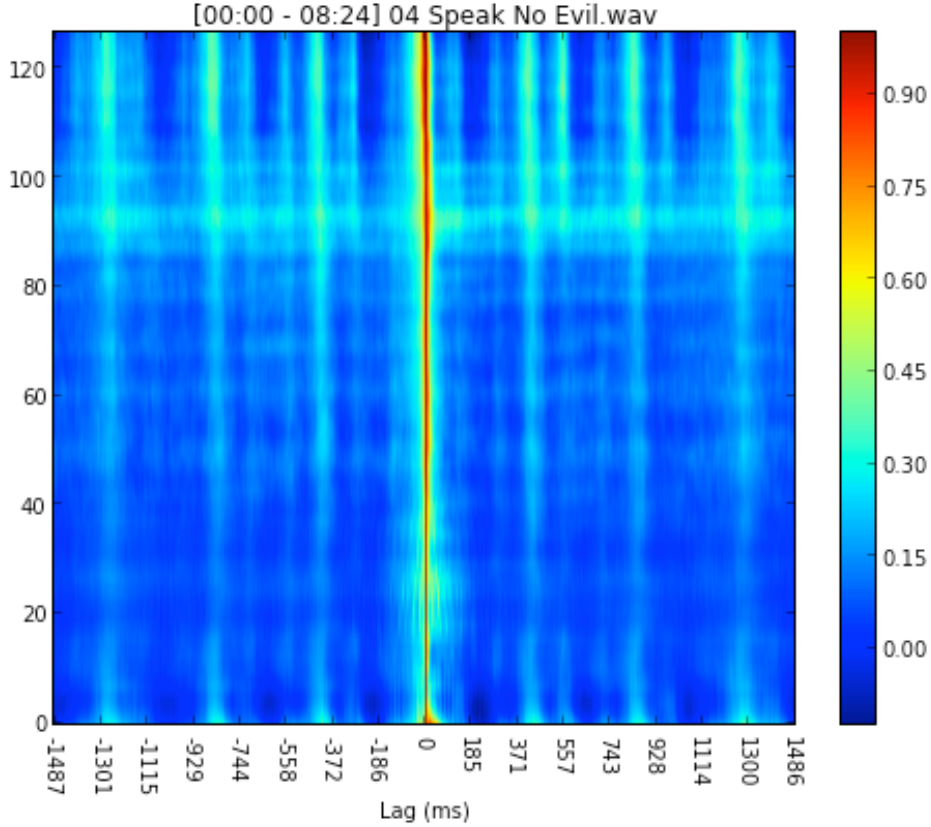

Elvin Jones

The Big Picture

- **After The Echo Nest solved music recommendation, what's left for MIR?**
  - decomposing audio into events
  - nuance
  - musical structure & motifs
  - good vs. bad

Music Research Hackathon?

- Like a hack day, except
  - no cash prizes, corporate sponsors, flash factor
  - aimed at academic researchers & similar
  - deeper problems more than pretty demos
- @Columbia, spring/summer 2013
- See Colin's breakout