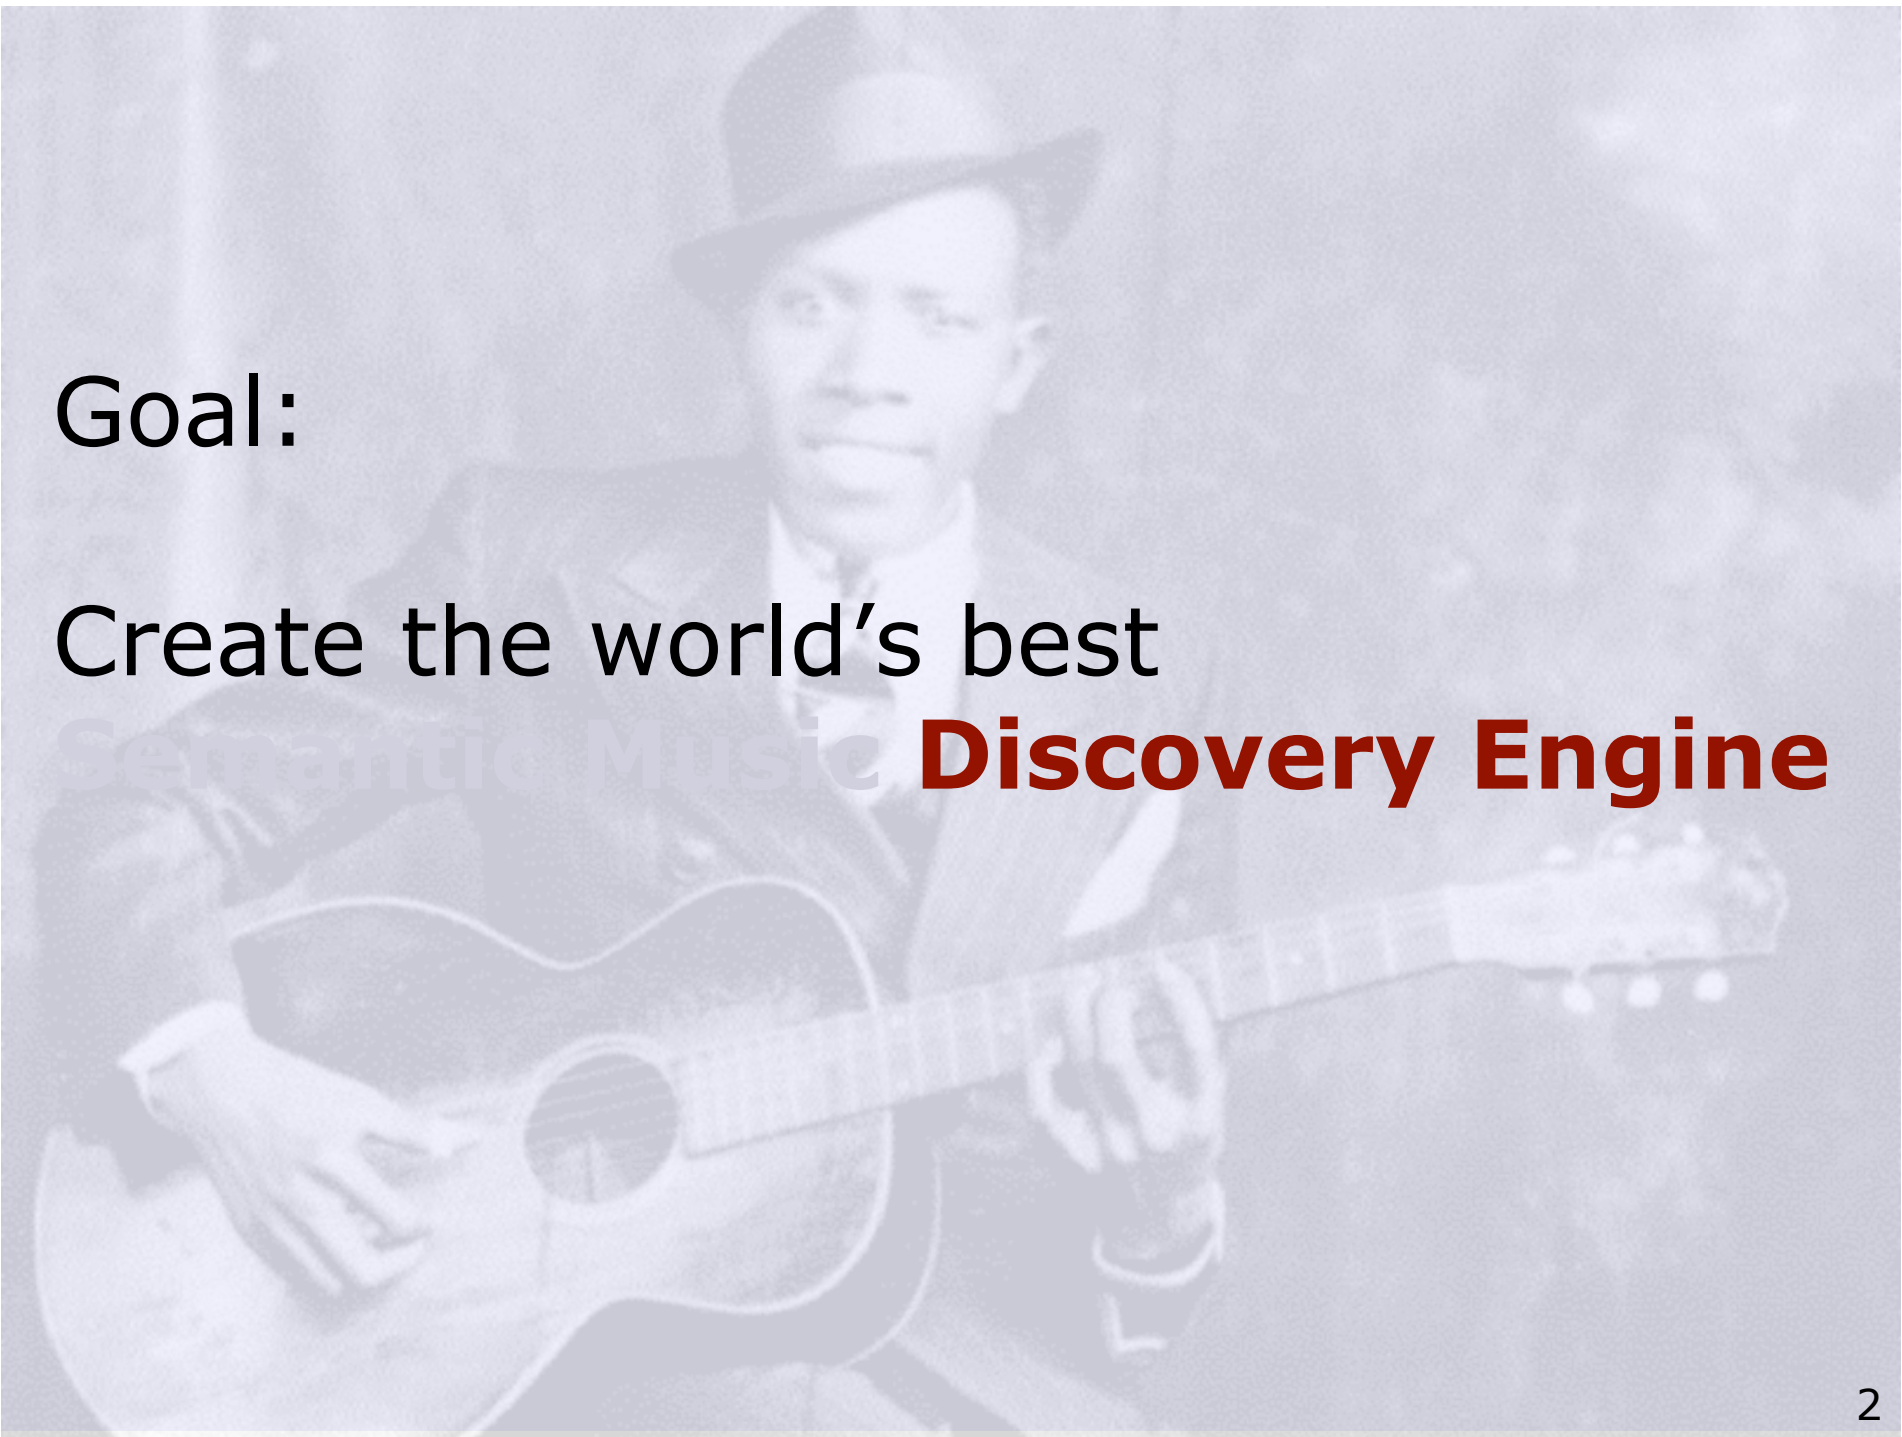




Goal:

Create the world's best
Semantic Music Discovery Engine




Goal:

Create the world's best
Semantic Music Discovery Engine



Goal:

Create the world's best
Semantic Music Discovery Engine



Goal:

Create the world's best
Semantic Music Discovery Engine



Meerkat



I Made A Lovers Prayer by Gillian Welch

Blue Diamond Mines by The Johnson Mountain Boys

Whiskey Lullaby by Alison Krauss

Tramp On The Street by Dian & The Greenbriar Boys

Count Me In by Deana Carter

'Whiskey Lullaby' Tags:

folk influences

country women

chicks

country

easy listening

female-vocalists

Station Tags:

Add Tag

old time country ✕

bluegrass ✕

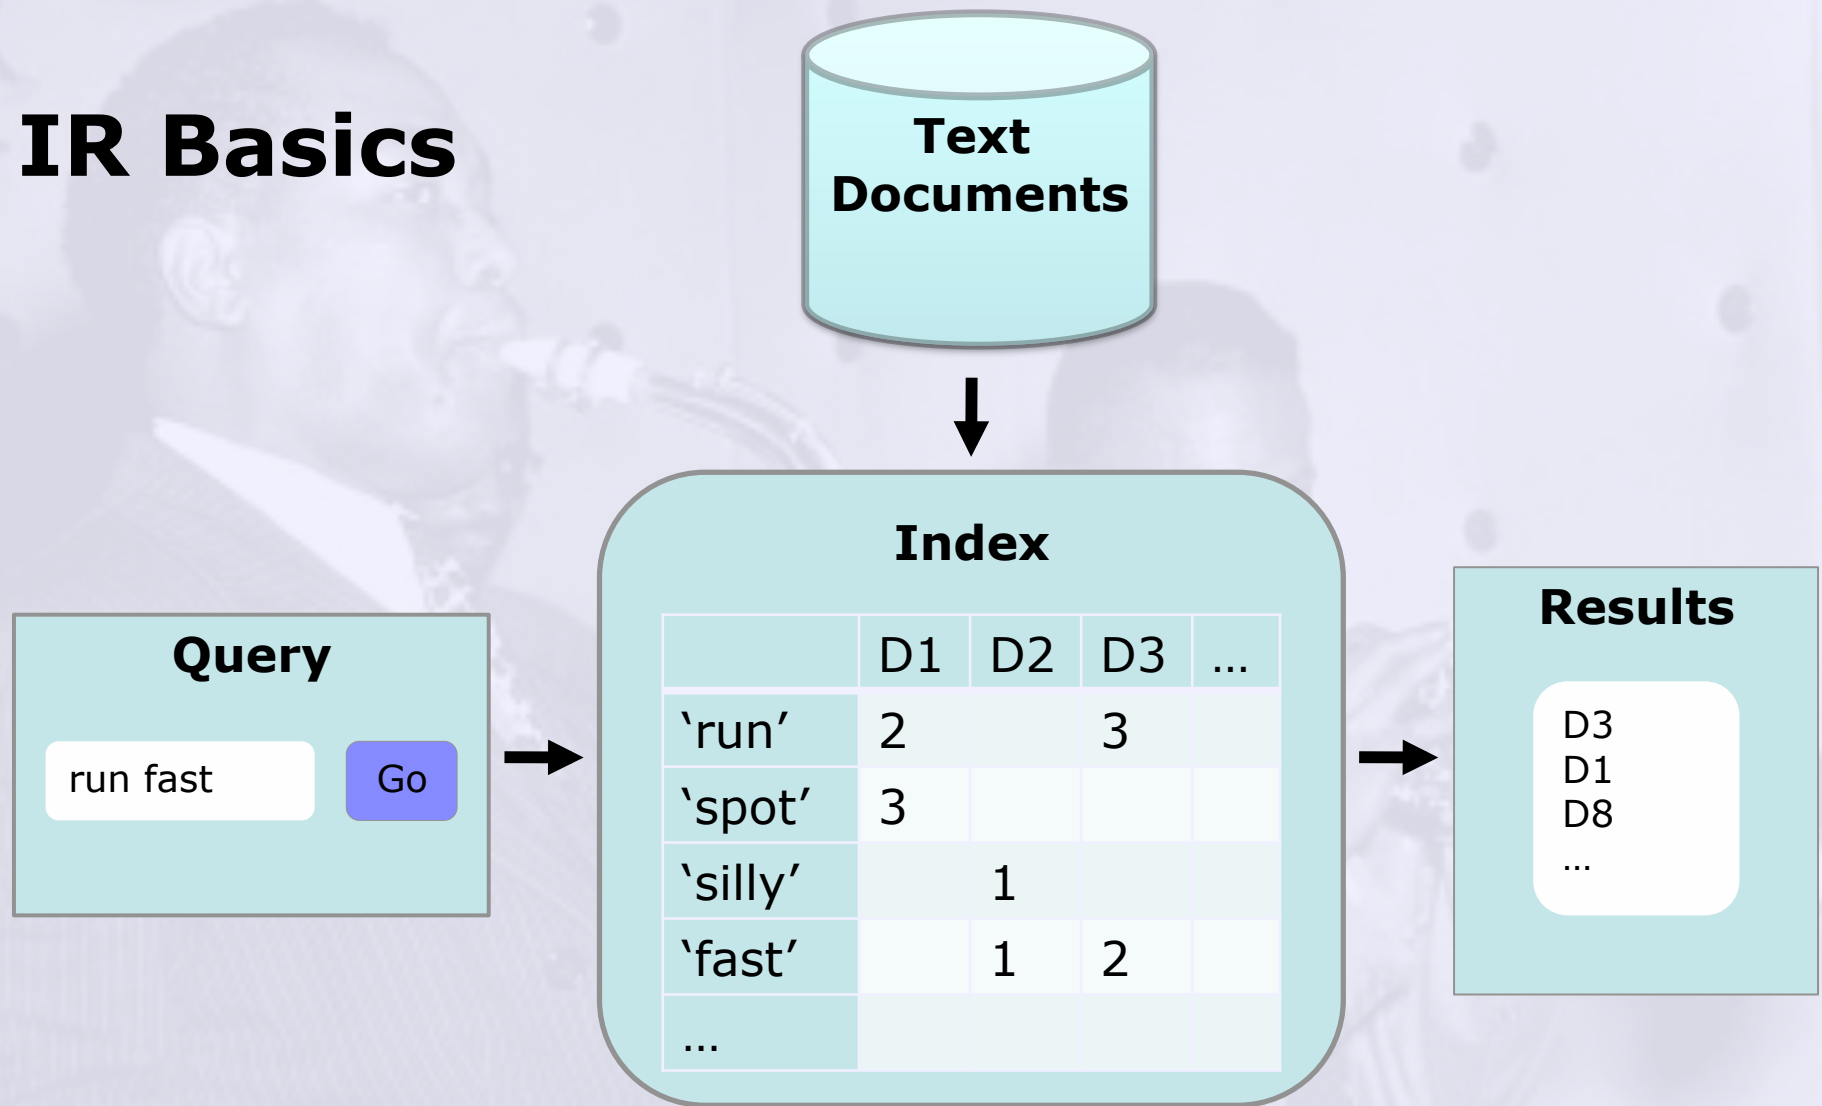
female vocals ✕



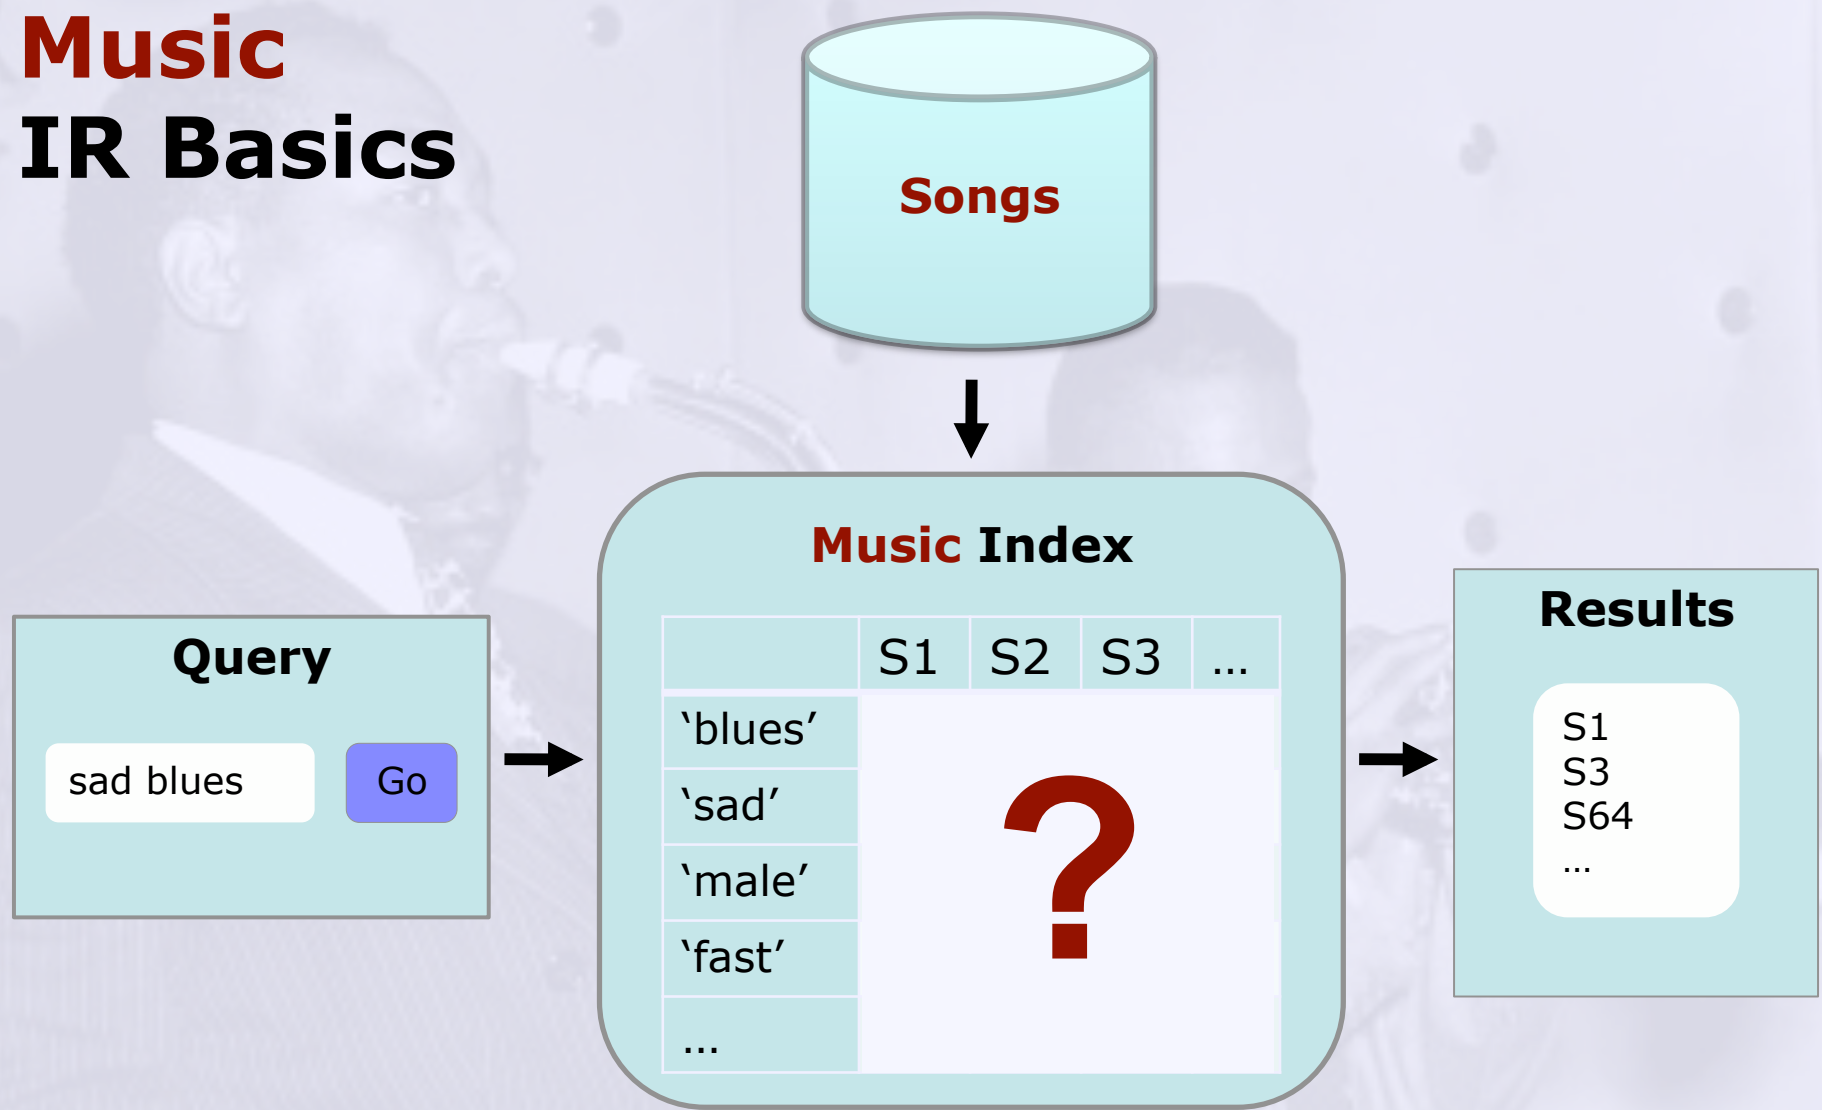
Problem:

Music
Annotation
is Hard

IR Basics



Music IR Basics



Four Sources of Music Information



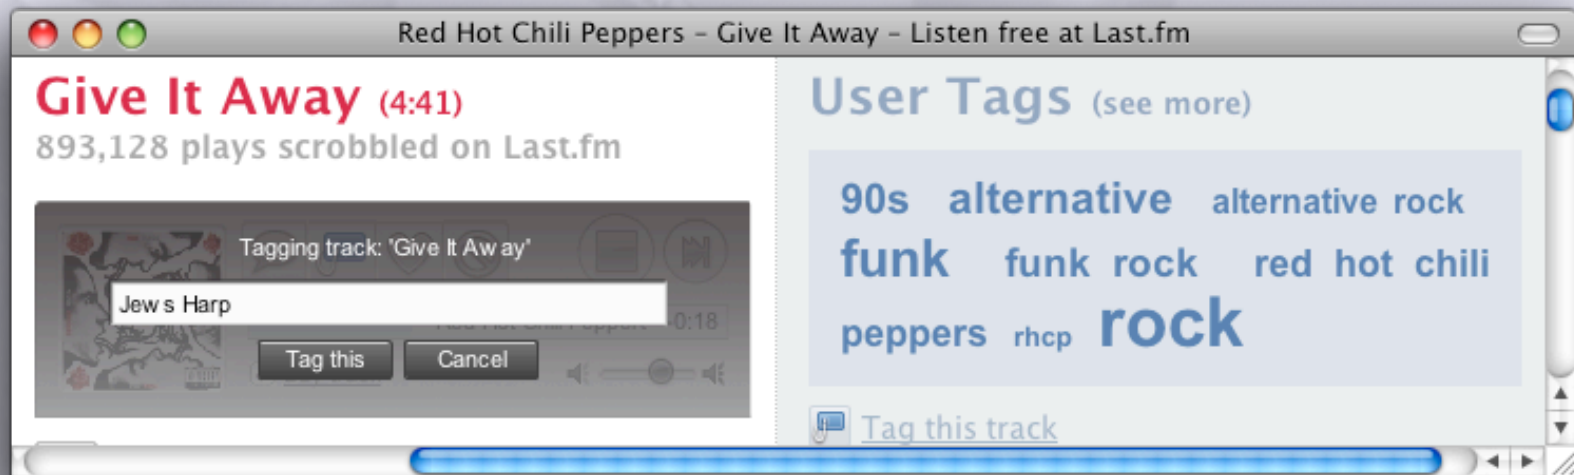
1. Human Survey

CAL500 [SIGIR 07]

- 174 tags
- 55 undergrads, 120 hours
- 500 songs annotated by 3+ people



2. Social Tagging




3. Mining Web Documents [Knees 07]

The screenshot shows a web browser window with the address bar displaying `http://www.allmusic.com/cg/amg.dll?p=amg&sq`. The page title is "Give It Away" by Red Hot Chili Peppers. Below the title, there are two sections: "Composed By" (Flea/John Frusciante/Anthony Kiedis/Chad Smith) and "Other Links" (All Performers that have performed this Title). The main content is a "Song Review" by Steve Huey. The review text is as follows:

The first single off the Red Hot Chili Peppers' quadruple-platinum smash *Blood Sugar Sex Magik*, "Give It Away" didn't achieve the massive **pop** success of its follow-up, "Under the Bridge," barely peaking in the Top 75, but it did become one of the band's most instantly recognizable songs, thanks mostly to Anthony Kiedis' nonsensical **raps**. Flea's **jumping, sliding, popping bass line**, and Pete Weiss' **jew's-harp** boinging away in the background; plus, MTV jumped all over the visually distinctive video, which featured the scantily clad bandmembers cavorting in the desert wearing silver body paint. Kiedis' lyrics were a free-associative mixture of positive vibes, tributes to musical heroes, and free love, and their literal meaning was often as difficult to understand as Kiedis' **nasal, staccato enunciation**. But that **distinctive vocal style** helped make the most comprehensible lines even catchier and more memorable, greatly enhancing the song's appeal. John Frusciante's **guitar** should not be underappreciated either, his **noisy, scratchy funk-rock** work adding depth and texture to the powerhouse rhythm section of Flea and Chad Smith. Frusciante also adds the song's two most unpredictable change-ups: a sudden contrast to Kiedis' hyperactivity in the form of a **languid solo** pre-recorded and dubbed backwards over the rhythm track, and a **hard-rocking riff** which is not introduced until the song's outro and bears a more than suspicious resemblance to Black Sabbath's "Sweet Leaf." **Funky** and **hard-rocking, horny and cheerfully loopy** all at the same time, "Give It Away" stands as one of the Chili Peppers' best singles, and a landmark single in relation to popular acceptance of **funk-metal**.

4. Content-based AutoTagging [SIGIR 07]

Learn a probabilistic model that captures a relationship between **audio content** and **tags**.



4. Content-Based AutoTagging

Supervised Multi-class Labeling model [SIGIR 07]

- One Gaussian Mixture Model (GMM) per tag
- Mixture Hierarchies EM Algorithm

Notes:

- Developed for image annotation [Carneiro 06]
- Scalable and Parallelizable
- Top system on 2008 MIREX Autotagging Task

Comparing Data Sources

Groundtruth

- **CAL500** - {72 tags} x {500 songs} binary labels

Approaches

1. **Social Tags** - Last.fm
2. **Web Autotags** - Site-specific relevance scoring
3. **Audio Autotags** - SML model w/ MFCCs

Comparing Data Sources

For each 3 approach:

For each 72 tag:

1. Rank songs
2. Calculate **Area under the ROC curve** (AUC)
 - 0.5 random ranking (Bad)
 - 1.0 perfect ranking (Good)

Calculate mean AUC

Constructing a Music Index

Approach	AUC	Best for # Tags
Social Tags	0.623	9
Text Mining Web Documents	0.625	12
Content-Based Autotagging	0.731	51
Single Source Oracle	0.756	

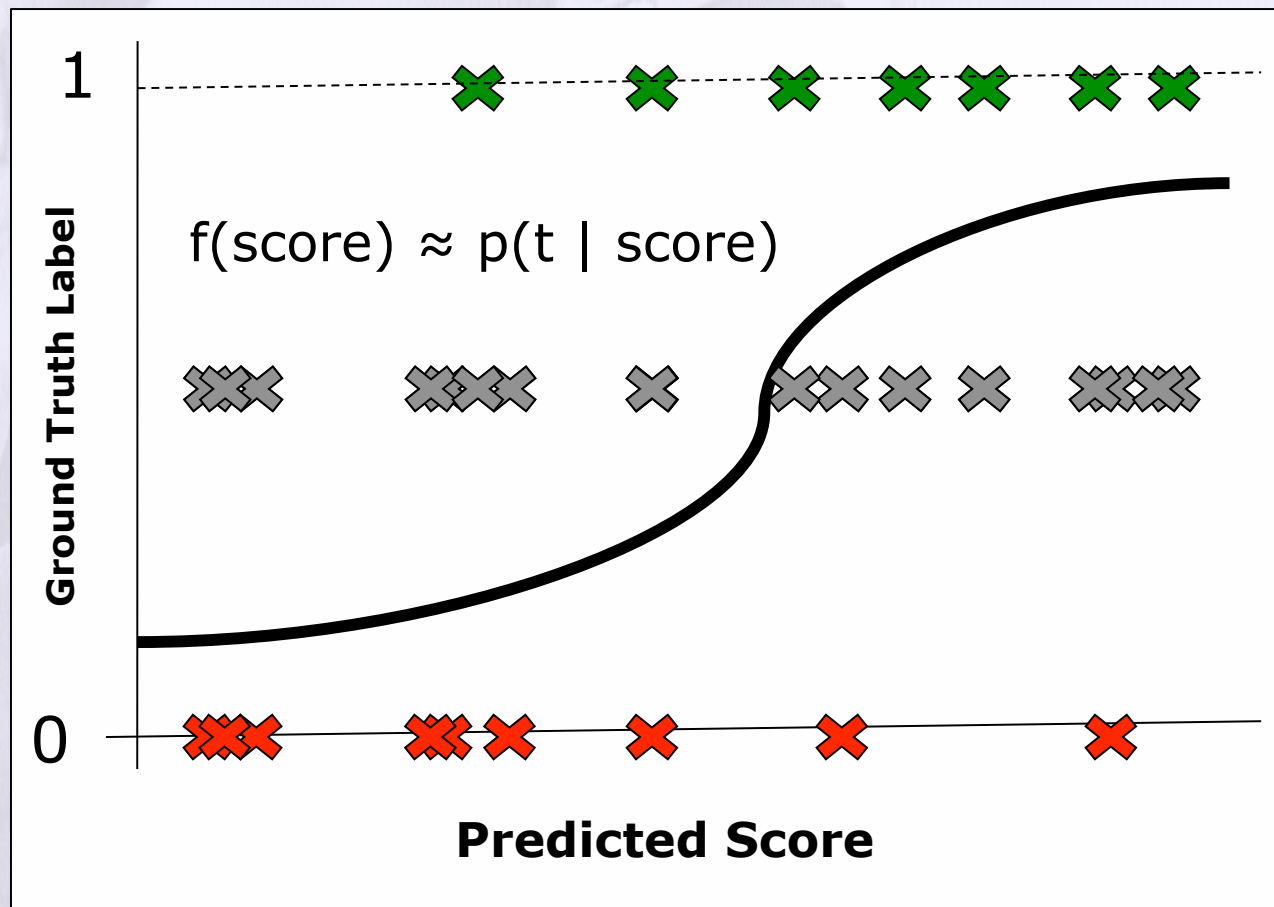
Combining Data Sources

Combining Data Sources

Approaches


- 1. Calibrated Score Averaging** - [Zadrozny 02]
- 2. RankBoost** - [Freund 03]
- 3. Kernel Combination SVM**- [Lanckriet 04]

1. Calibrated Score Averaging




1. Calibrated Score Averaging


Autotag
Score



Web
Document
Score



Social Tag
Score



$$p(t|\mathbf{s}) = \prod_i p(t|s_i)$$

2. RankBoost [Freund 03]

- Greedy iterative algorithm:
 - pick **weak ranker** & assigns **weight** based on loss function
- **Strong ranker** is a weighted voted by weak rankers

2. RankBoost

Weak Ranker = (Source, Threshold, Weight)

AutoTags		
id	Score	Label
8	0.82	1
2	0.77	0
13	0.76	1
7	0.69	1
3	0.68	0
...		

Diagram illustrating the AutoTags weak ranker. The table shows scores and labels for various items. A green horizontal line is drawn at the score of 0.69, separating items with Label 1 (above) from Label 0 (below). A green arrow labeled w_1 points up from the line, and a black arrow labeled 0 points down from the line. Another green arrow labeled w_3 points up from the line, and a black arrow labeled 0 points down from the line.

Web Docs		
id	Score	Label
7	19	1
14	18	0
9	16	1
11	16	0
21	15	1
...		

Diagram illustrating the Web Docs weak ranker. The table shows scores and labels for various items. A blue horizontal line is drawn at the score of 15, separating items with Label 1 (above) from Label 0 (below). A blue arrow labeled w_4 points up from the line, and a black arrow labeled 0 points down from the line.

Social Tags		
id	Score	Label
13	100	1
4	95	1
17	95	0
11	89	0
8	82	1
...		

Diagram illustrating the Social Tags weak ranker. The table shows scores and labels for various items. An orange horizontal line is drawn at the score of 95, separating items with Label 1 (above) from Label 0 (below). An orange arrow labeled w_2 points up from the line, and a black arrow labeled 0 points down from the line.

2. RankBoost

Strong Ranker: $f(\mathbf{s})$

Test Song $\mathbf{s} = [s_{\text{audio}}, s_{\text{web}}, s_{\text{social}}]$

$$f(\mathbf{s}) = w_1 I(s_{\text{audio}} > t_1) + w_2 I(s_{\text{social}} > t_2) + \dots$$

3. Kernel Combo SVM

1. Compute **kernel matrix** for each data source.

- 3 Song X Song similarity matrices

$$W_1 \boxed{K_{\text{audio}}} + W_2 \boxed{K_{\text{web}}} + W_3 \boxed{K_{\text{social}}} = \boxed{K_{\text{combo}}}$$

2. Learn an optimal **linear combination of the kernel matrices** using convex optimization

- Produces single kernel matrix for SVM

3. Rank songs based on score from SVM

- **positive distance** from **separating hyperplane**.

Summary

Approach	AUC
Single Source Oracle	0.756
Kernel Combination	0.756
RankBoost	0.760
Calibrated Score Averaging	0.763

Summary

Data Sources provide complementary information

- Sparse **Social Context**
- Dense **Audio Content**

Calibrated Score Averaging & RankBoost work well

- Good performance, easy to implement, easy to tune, quick to train & evaluate

Chapter 4

The Future (Man)

Research Challenges

1. Combine Music Information Sources

- Games, Social Networks, Popularity Info

2. Improving **Autotagging**

- Discriminative Approaches [Mandel 08, Eck 07]

3. **Personalization**

- Demographic Groups
- Psychographic Groups
- Individual Preference
- Emotional states of Individual

What's on tap

1. Big new data set

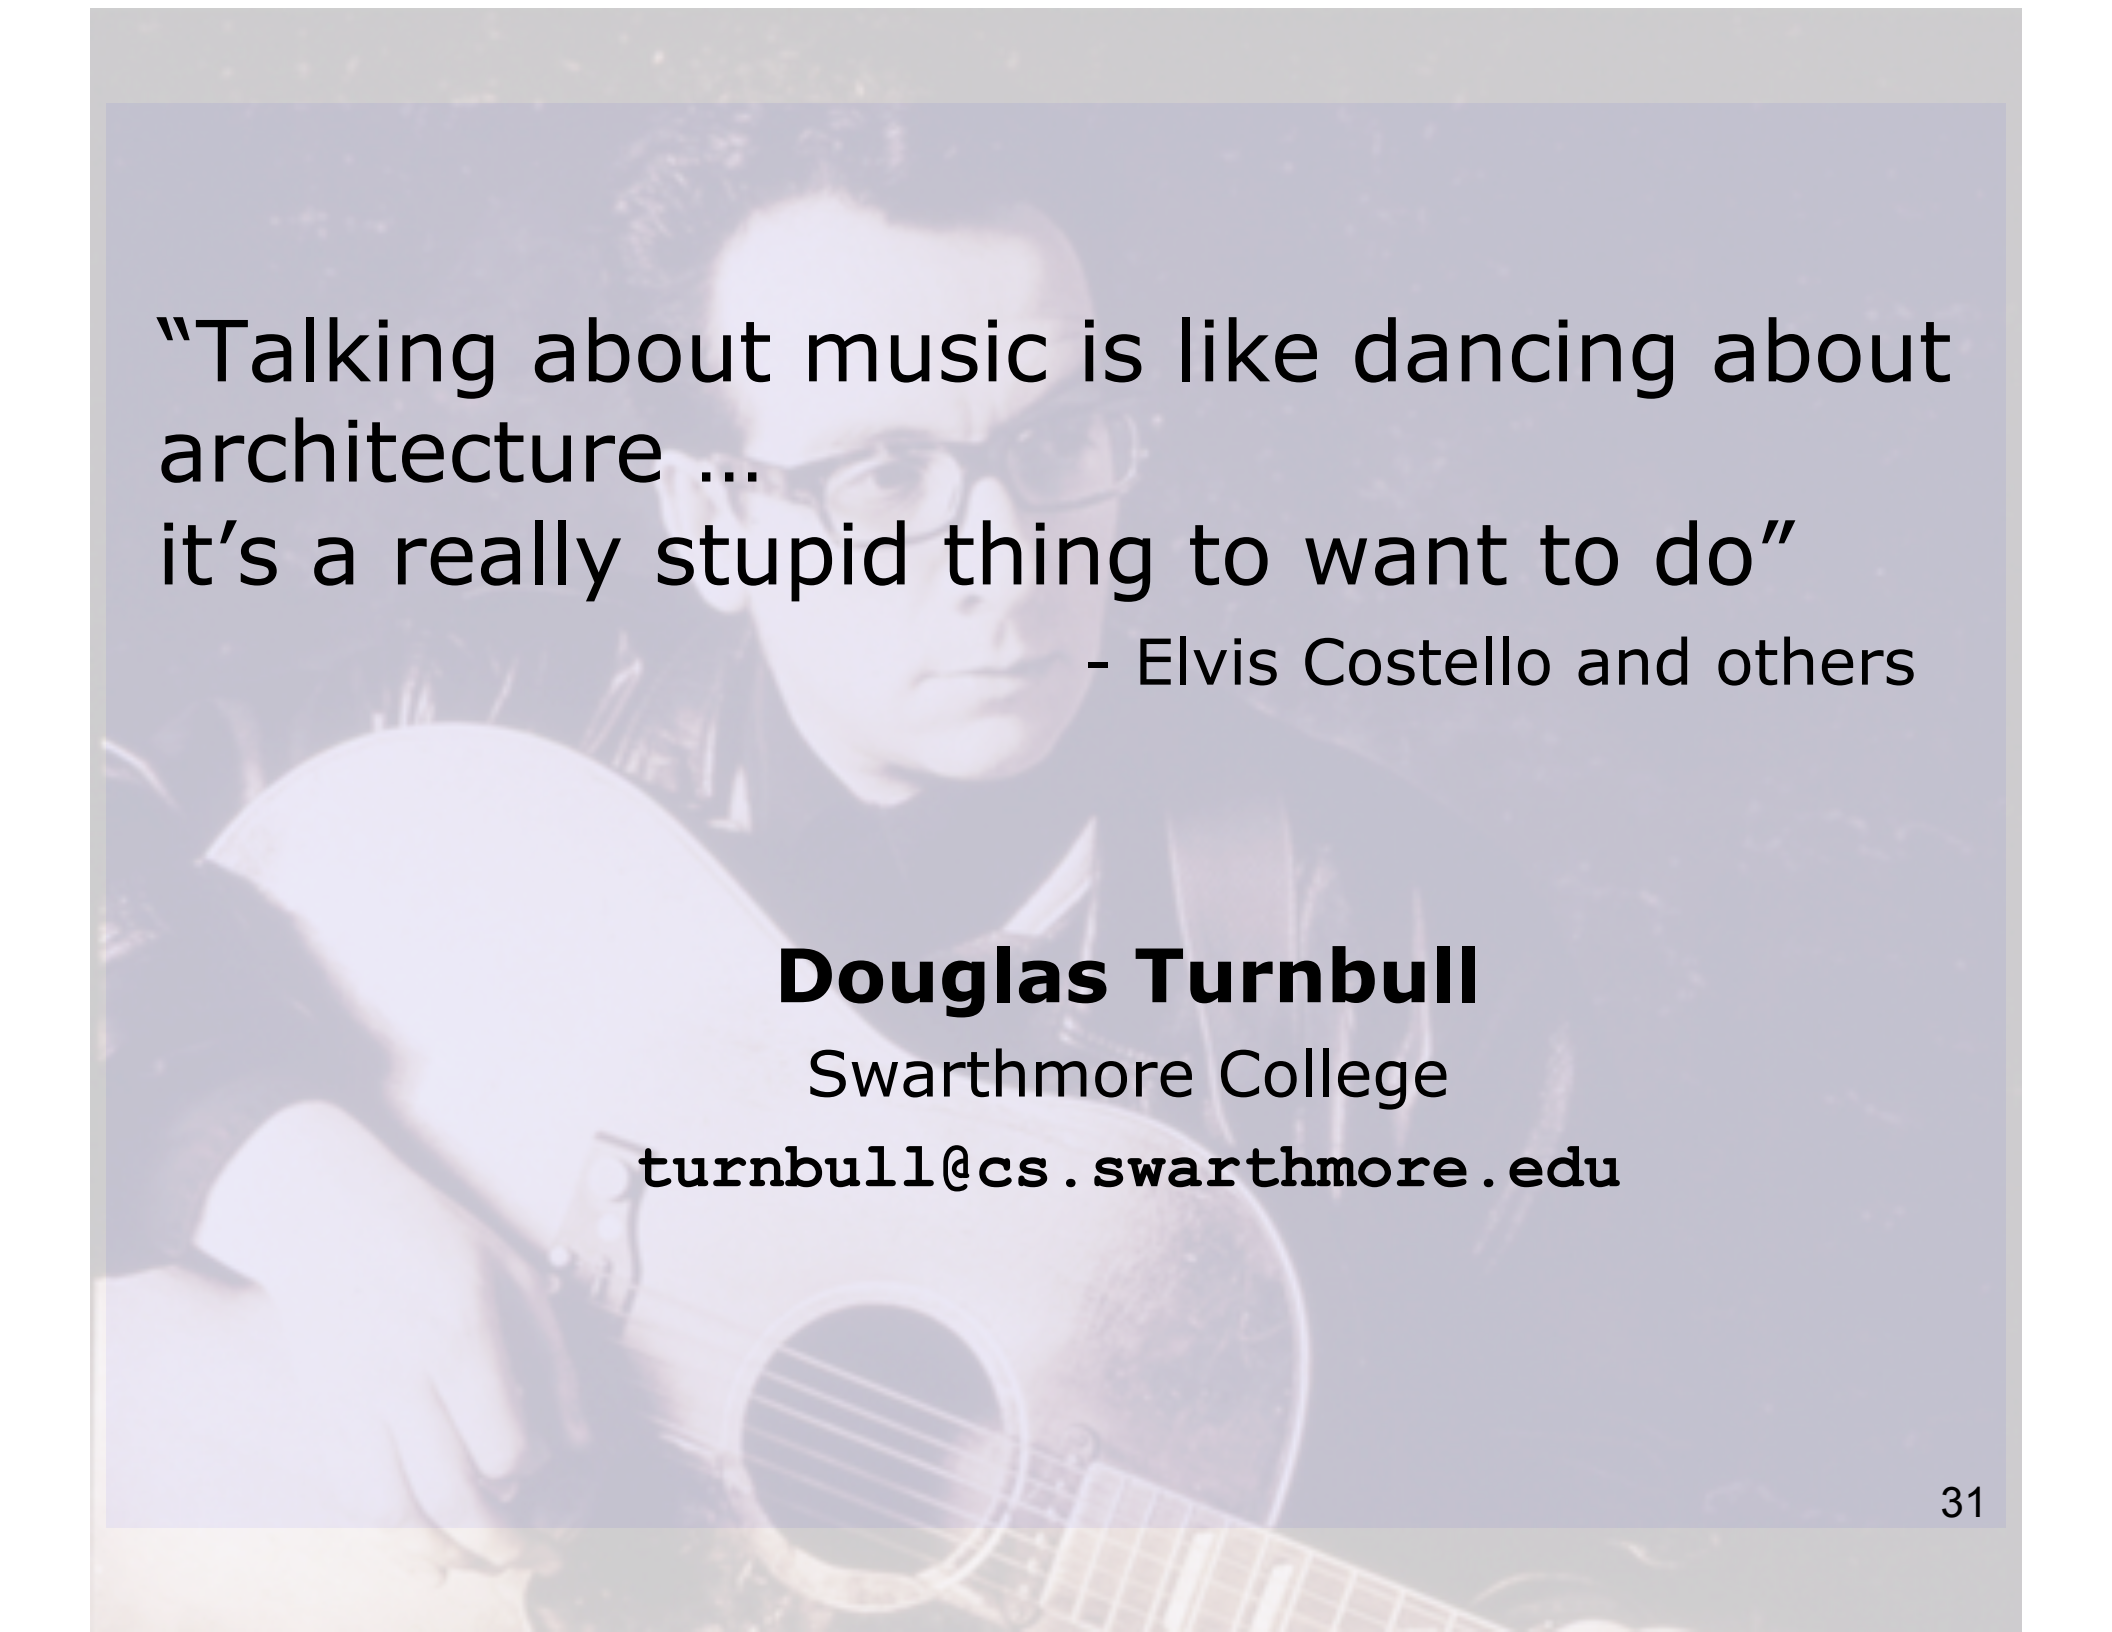
- 10,000 songs
- Acoustic, Genre, Social Tags

2. Herd-it Game [HCOMP 09]

- <http://apps.facebook.com/herd-it>
- Analyzing Game Data

3. New Everything

- Autotagging Approaches
- Content-Context Approach
- User Interfaces



“Talking about music is like dancing about
architecture ...
it’s a really stupid thing to want to do”

- Elvis Costello and others

Douglas Turnbull

Swarthmore College

`turnbull@cs.swarthmore.edu`

References

Social Context-Audio Content [Tech, SIGIR 09, ISMIR 08]

Autotagging [IEEE TASLP 08, SIGIR 07]

Related:

Music Annotation Games [HCOMP 09, ISMIR 07a]

Query-by-Semantic-Similarity [ICASSP 07, MIREX 07]

Tag Vocabulary Selection with Sparse CCA [ISMIR 07b]

Supervised Music Boundary Detection [ISMIR 07c]

Collaborative Filtering & Tag Propagation [Tech]