



Music recommendation and discovery at scale

Douglas Eck
Research Scientist, Google
Mountain View, CA, USA



Overview

- Research at Google
- What do users want from a music service?
- Music similarity using embeddings
 - Audio (content based)
 - User modeling (collaborative filtering)
- Structured data (Knowledge graph / Freebase)
- Casting recommendation as a search problem.

Research Collaboration with Google

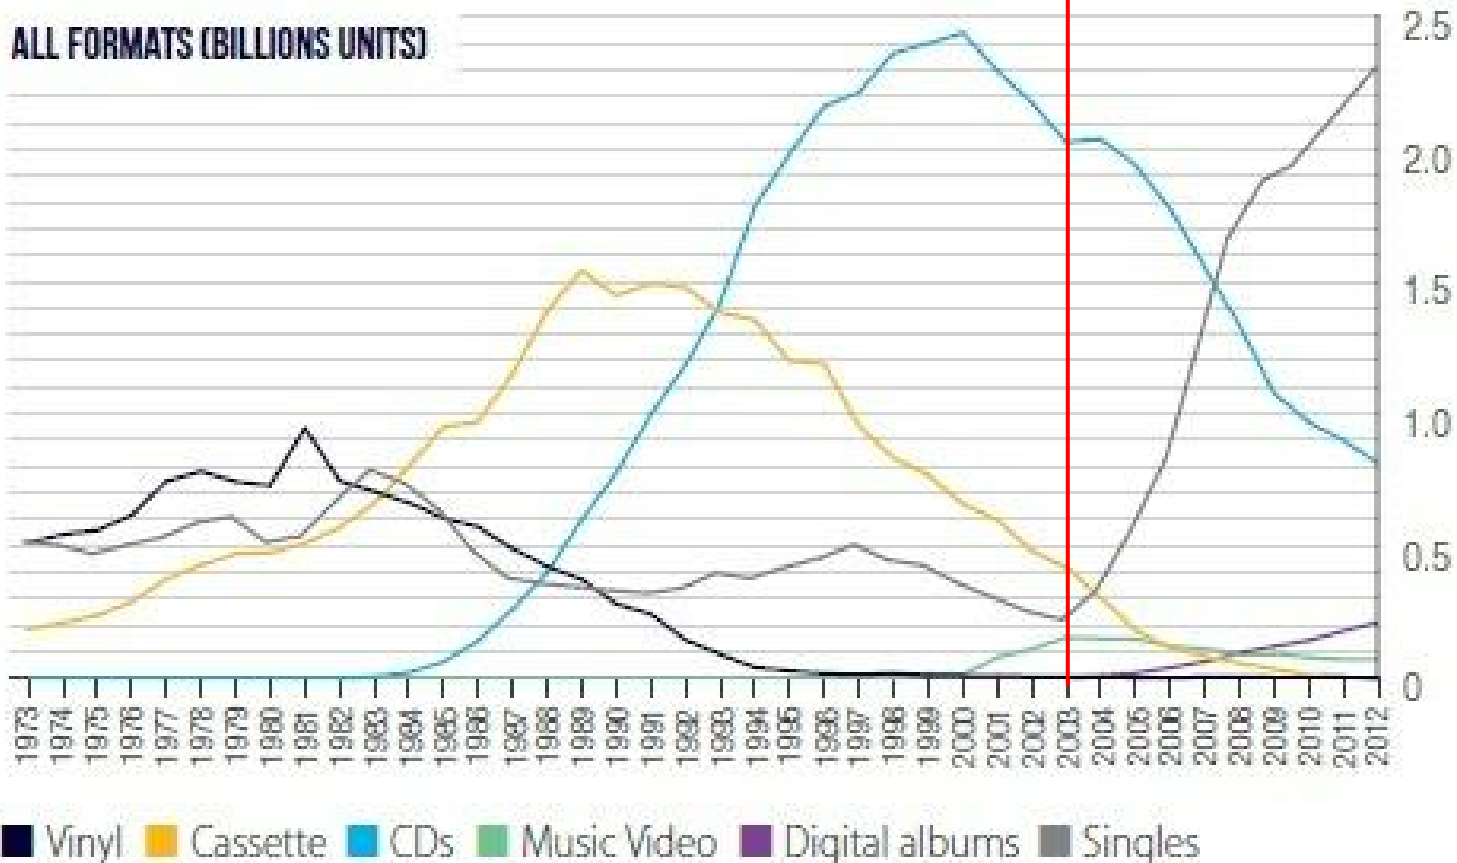
- “Hybrid Research Model”: blur the line between research and engineering; focus on high risk projects with big potential impact.
- Research areas (# publications)

Algorithms and Theory (368)	Information Retrieval and the Web (182)
AI and Machine Learning (386)	Machine Perception (226)
Data Management (92)	Machine Translation (938)
Data Mining (148)	Mobile Systems (47)
Distr. Systems and Parallel Computing (127)	Natural Language Processing (245)
Economics and Electronic Commerce (108)	Networking (109)
Education Innovation (20)	Security, Cryptography, and Privacy (176)
General Science (89)	Software Engineering (63)
Hardware and Architecture (47)	Software Systems (134)
HCI and Visualization (286)	Speech Processing (102)
- University collaborations
 - Internships for students and jobs for recent grads www.google.com/about/careers/students/
 - Google Research Awards for faculty research.google.com/university/relations/research_awards.html

Mp3 Revolution and iPod


2003: iTunes Store Opens

ALL FORMATS (BILLIONS UNITS)

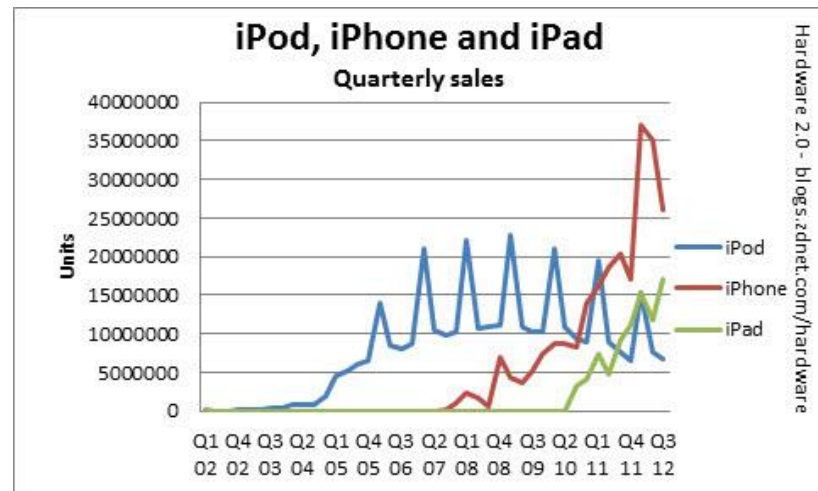


From iPod to Online Streaming From Desktop to Mobile

Pandora Listener Hours and Active Users



	FY3Q12	FY3Q11
Listener Hours (billions)	2.1	1.0
Active Users (millions)	40.0	24.0
Platform Hours		
Mobile	70%	52%
Non-mobile	30%	48%



Lean back (low friction)





Lean forward (high engagement)



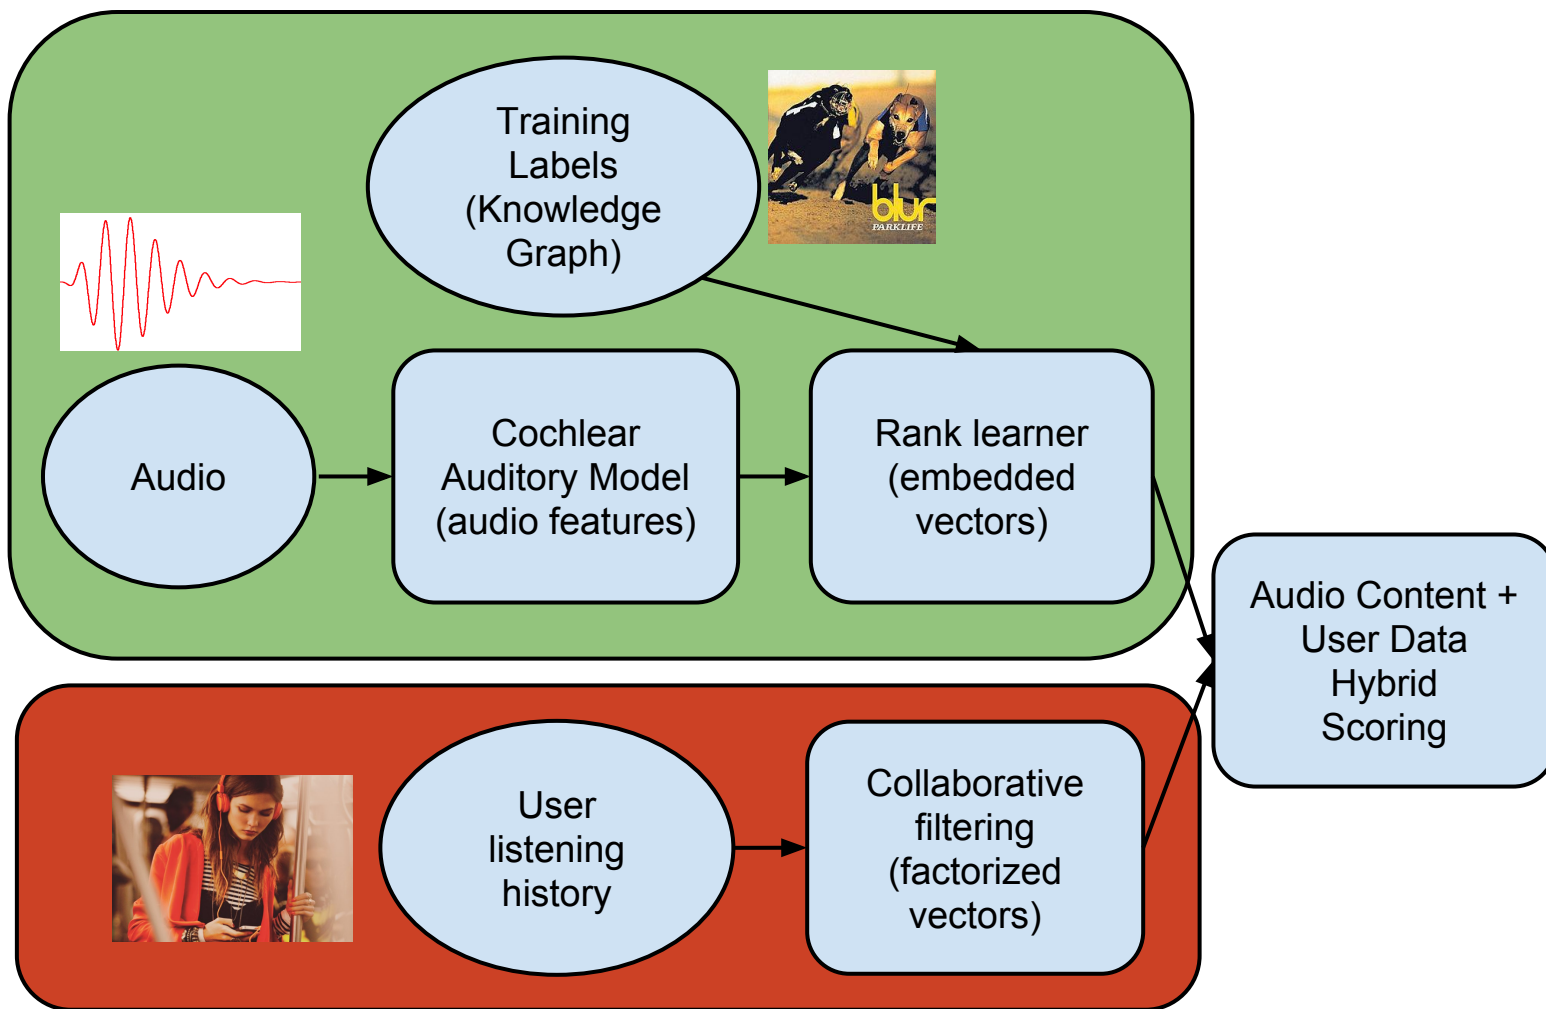


Users
Collaborative Filtering

Musicians
Knowledge and Connections




Music
Audio Signal Processing



Collaborative filtering

Jane likes the Thelonious Monk's album *Straight, No Chaser*"

Can we use this info to "filter" for other music Jane likes?



Customers Who Bought This Item Also Bought




Monk's Dream
 ▶ Thelonious Monk
 ★★★★★ (25)
 Audio CD
 \$8.30



Mingus Ah Um
 ▶ Charles Mingus
 ★★★★★ (60)
 Audio CD
 \$6.88



Criss Cross
 ▶ Thelonious Monk
 ★★★★★ (15)
 Audio CD
 \$8.24



Time Out
 ▶ Dave Brubeck
 ★★★★★ (277)
 Audio CD
 \$6.73



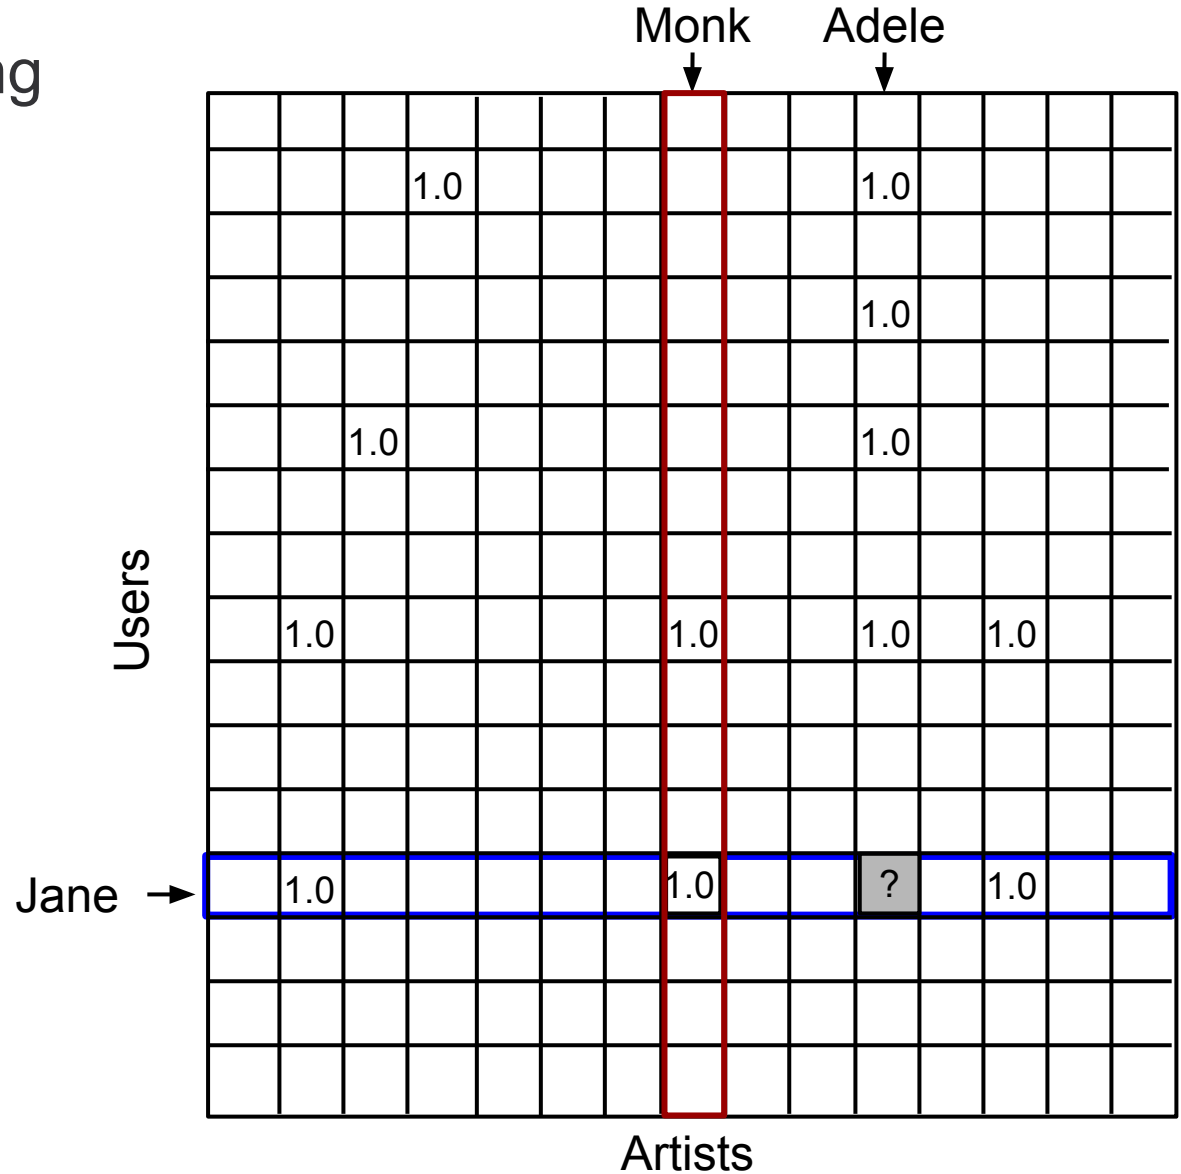
Giant Steps
 ▶ John Coltrane
 ★★★★★ (142)
 Audio CD
 \$7.99



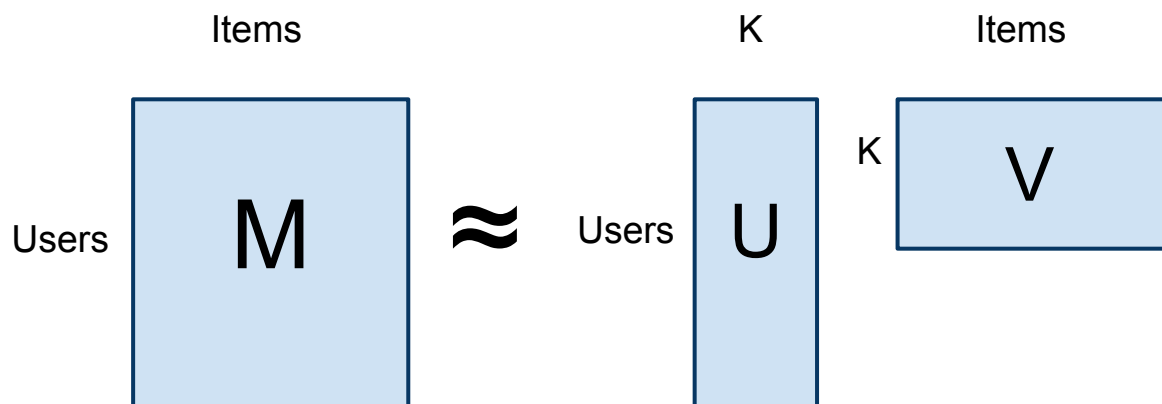
Essential Thelonious Monk
 ▶ Thelonious Monk
 ★★★★★ (11)
 Audio CD
 \$8.34

Collaborative filtering

Does Jane like Adele?



Matrix factorization for recommendation



Given a factorized ratings matrix $M \approx U \times V$:

- Take embedding vector u for a user, v for an item.
- Score by $d(u, v)$, usually dot product.

Intuition:

- $u \cdot v$ is the reconstructed matrix's guess if the user likes the item.

Embedding spaces: SVD vs WALS

SVD

Tracks

	1	0	0	0	1	1
	0	0	1	0	1	0
Users	0	1	0	0	0	1
	1	1	0	0	1	0

$$\min_{U,V} \sum_{i,j} (p_{i,j} - u_i \cdot v_j)^2$$

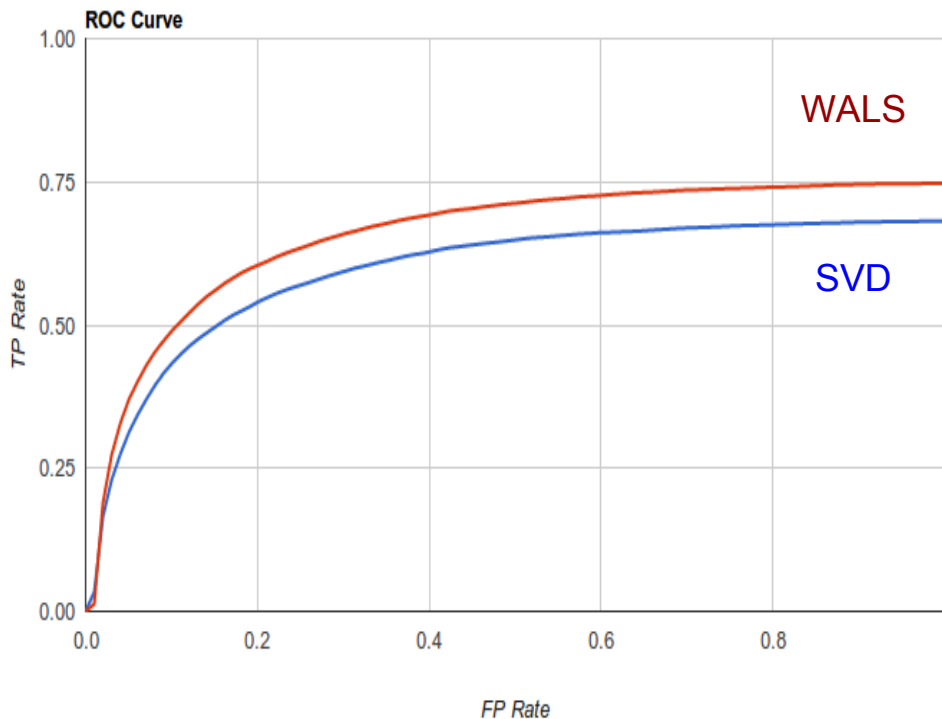
WALS

Tracks

	1	0	0	0	1	1
	0	0	1	0	1	0
Users	0	1	0	0	0	1
	1	1	0	0	1	0

$$\min_{U,V} \sum_{p_{i,j} \neq 0} c_{i,j} (p_{i,j} - u_i \cdot v_j)^2$$

WALS > SVD




WALS effects:

- Weakens popular songs' ability to link together less-popular ones.
- Helps combat correlated negatives: ABBA and AC/DC.

Recommendations for Coldplay

#	SVD	WALS
1	Coldplay	Coldplay
2	Faultline	Snow Patrol
3	Oasis	Keane
4	U2	The Fray
5	The Fray	U2
6	The Cranberries	Maroon 5
7	Snow Patrol	Oasis
8	Elton John	James Blunt
9	Creeper Lagoon	John Mayer
10	Bob Marley	Travis
11	Five For Fighting	Kings Of Leon
12	Kansas	David Gray
13	Frank Sinatra	Death Cab for Cutie
14	John Lennon	Damien Rice
15	Moby	Faultline
16	Journey	Aqualung
17	Morten Harket	R.E.M.
18	Nelly Furtado	Mat Kearney
19	Johnny Cash	Five For Fighting
20	Phil Collins	Foo Fighters



Similarity Functions For Recommenders

1. Dot product

$$u \cdot v$$

Tends to be biased toward popular items.

2. Cosine similarity

$$u \cdot v / \|u\| \cdot \|v\|$$

Tends to be biased toward niche items.

3. Limited inner product

$$u \cdot v / \|u\| \cdot \max(\|u\|, \|v\|)$$

Popularity of seed informs ranking.

Grateful Dead \Rightarrow Phish

Jerry Garcia Band \Rightarrow Phil Lesh

Arithmetic on embeddings: merging and steering

Efficient Estimation of Word Representations in Vector Space

Tomas Mikolov

Google Inc., Mountain View, CA
tmikolov@google.com

Greg Corrado


Google Inc., Mountain View, CA
gcorrado@google.com

Kai Chen


Google Inc., Mountain View, CA
kaichen@google.com

Jeffrey Dean

Google Inc., Mountain View, CA
jeff@google.com



CBOW



Skip-gram

Figure 1: New model architectures. The CBOW architecture predicts the current word based on the context, and the Skip-gram predicts surrounding words given the current word.

Relationship	Example 1
France - Paris	Italy: Rome
big - bigger	small: larger
Miami - Florida	Baltimore: Maryland
Einstein - scientist	Messi: midfielder
Sarkozy - France	Berlusconi: Italy
copper - Cu	zinc: Zn
Berlusconi - Silvio	Sarkozy: Nicolas
Microsoft - Windows	Google: Android
Microsoft - Ballmer	Google: Yahoo
Japan - sushi	Germany: bratwurst

Michael Jackson - Lady Gaga + Prince = ???

Radio from Michael Jackson Thriller

0		Michael Jackson - Thriller - 04 - Thriller
1		Beyoncé - I AM...SASHA FIERCE - 01 - Single Ladies (Put A Ring On It)
2		Lady Gaga - The Fame (US Version) - 02 - LoveGame
3		Justin Timberlake - Justified - 06 - Rock Your Body
4		Michael Jackson - Thriller - 08 - P.Y.T. (Pretty Young Thing)
5		Mariah Carey - The Emancipation of Mimi (U.S. Full Length CD) - 02 - We Belong Together
6		The Black Eyed Peas - Monkey Business - 01 - Pump It
7		The Jacksons - The Essential Michael Jackson - 08 - Blame It On The Boogie
8		Alicia Keys - As I Am - 04 - No One
9		Prince and the Revolution - Purple Rain - 06 - When Doves Cry (LP Version)
10		Michael Jackson - Thriller - 01 - Wanna Be Startin' Somethin'
11		Mariah Carey - E=MC ² - 02 - Touch My Body
12		Michael Jackson - Off the Wall - 01 - Don't Stop 'Til You Get Enough














Radio from Michael Jackson Thriller

Let's try to move it
towards 70s/80s
Michael Jackson.



0		Michael Jackson - Thriller - 04 - Thriller
1		Beyoncé - I AM...SASHA FIERCE - 01 - Single Ladies (Put A Ring On It)
2		Lady Gaga - The Fame (US Version) - 02 - LoveGame
3		Justin Timberlake - Justified - 06 - Rock Your Body
4		Michael Jackson - Thriller - 08 - P.Y.T. (Pretty Young Thing)
5		Mariah Carey - The Emancipation of Mimi (U.S. Full Length CD) - 02 - We Belong Together
6		The Black Eyed Peas - Monkey Business - 01 - Pump It
7		The Jacksons - The Essential Michael Jackson - 08 - Blame It On The Boogie
8		Alicia Keys - As I Am - 04 - No One
9		Prince and the Revolution - Purple Rain - 06 - When Doves Cry (LP Version)
10		Michael Jackson - Thriller - 01 - Wanna Be Startin' Somethin'
11		Mariah Carey - E=MC ² - 02 - Touch My Body
12		Michael Jackson - Off the Wall - 01 - Don't Stop 'Til You Get Enough

Michael Jackson
 + The Jacksons
 + Prince
 - Lady Gaga
 - Mariah Carey

0		Michael Jackson - Thriller - 04 - Thriller
1		Rihanna - Good Girl Gone Bad: Reloaded - 03 - Don't Stop The Music
2		Kool & The Gang - Celebrate! - 01 - Celebration
3		Michael Jackson - Thriller - 08 - P.Y.T. (Pretty Young Thing)
4		Prince - The Very Best Of Prince - 03 - Little Red Corvette (LP Version)
5		Stevie Wonder - Number 1's - 05 - Superstition
6		Michael Jackson - Off the Wall - 01 - Don't Stop 'Til You Get Enough
7		Stevie Wonder - Songs In The Key Of Life - 01 - Isn't She Lovely
8		Michael Jackson - Thriller - 01 - Wanna Be Startin' Somethin'
9		Rihanna - Good Girl Gone Bad: Reloaded - 13 - Disturbia
10		Jackson 5 - ABC - 01 - The Love You Save
11		The Black Eyed Peas - THE E.N.D. (THE ENERGY NEVER DIES) - 05 - I Gotta Feeling
12		Jackson 5 - ABC - 03 - ABC

Michael Jackson
 + The Jacksons
 + Prince
 - Lady Gaga
 - Mariah Carey














Add a couple more
 thumbs ups and
 downs.



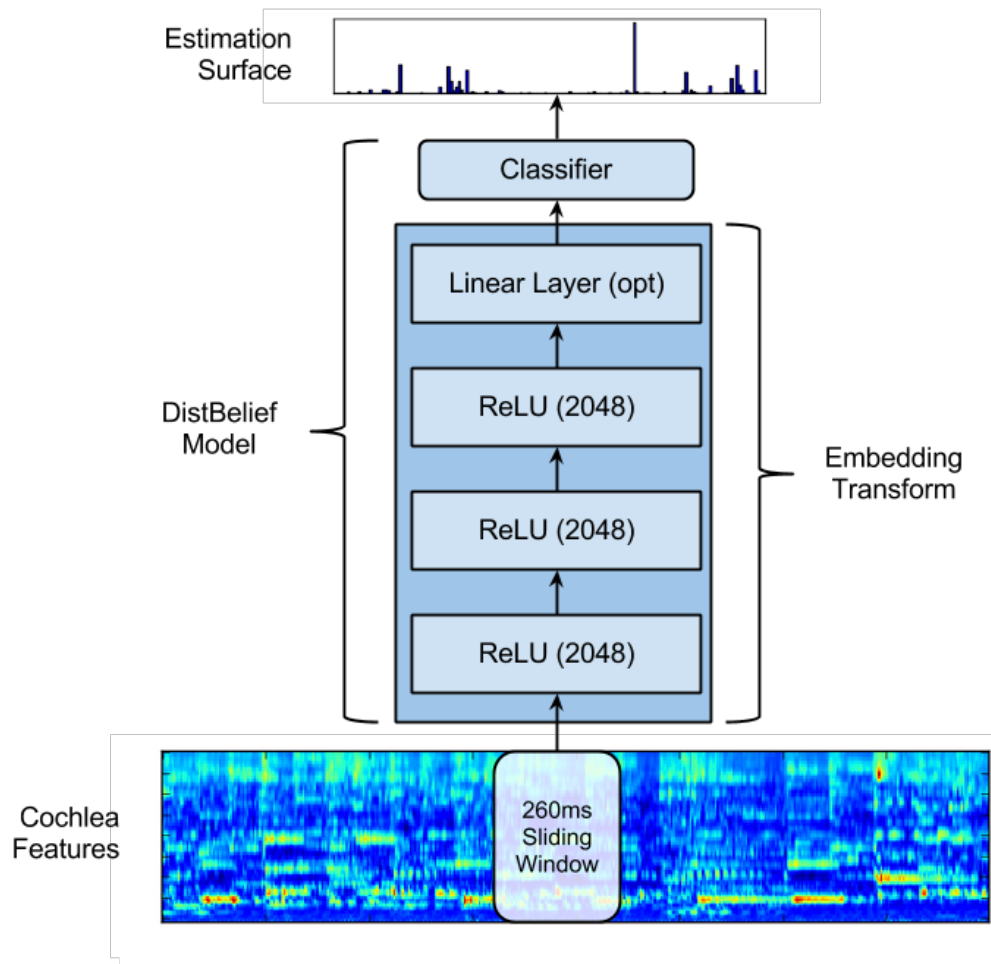
0		Michael Jackson - Thriller - 04 - Thriller
1		Rihanna - Good Girl Gone Bad: Reloaded - 03 - Don't Stop The Music
2		Kool & The Gang - Celebrate! - 01 - Celebration
3		Michael Jackson - Thriller - 08 - P.Y.T. (Pretty Young Thing)
4		Prince - The Very Best Of Prince - 03 - Little Red Corvette (LP Version)
5		Stevie Wonder - Number 1's - 05 - Superstition
6		Michael Jackson - Off the Wall - 01 - Don't Stop 'Til You Get Enough
7		Stevie Wonder - Songs In The Key Of Life - 01 - Isn't She Lovely
8		Michael Jackson - Thriller - 01 - Wanna Be Startin' Somethin'
9		Rihanna - Good Girl Gone Bad: Reloaded - 13 - Disturbia
10		Jackson 5 - ABC - 01 - The Love You Save
11		The Black Eyed Peas - THE E.N.D. (THE ENERGY NEVER DIES) - 05 - I Gotta Feeling
12		Jackson 5 - ABC - 03 - ABC

Composite Radio from Michael Jackson Thriller.

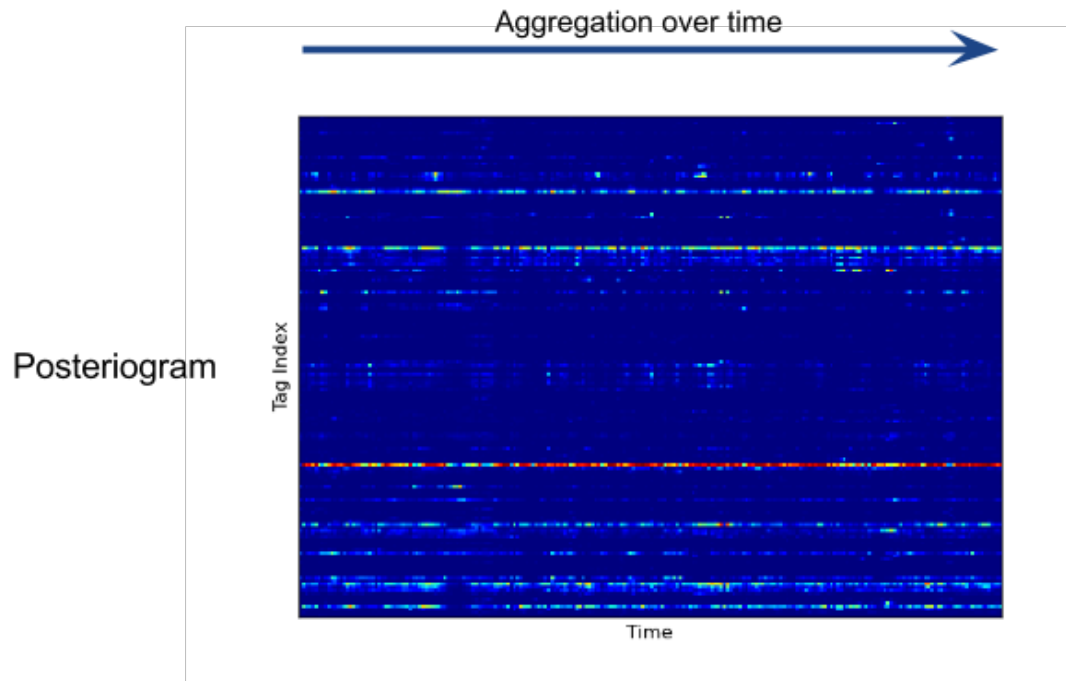

Michael Jackson
+ The Jacksons
+ Prince
+ The Jackson 5
- Lady Gaga
- Mariah Carey
- Black Eyed
Peas

0		Michael Jackson - Thriller - 04 - Thriller
1		The Jacksons - Victory - 08 - Body
2		PRINCE - Sign 'O' The Times - 15 - It's Gonna Be A Beautiful Night (Live LP Version)
3		Kool & The Gang - Celebrate! - 01 - Celebration
4		Earth, Wind & Fire - Greatest Hits - 08 - Serpentine Fire
5		Kool & The Gang - The Very Best Of Kool & The Gang - 13 - Joanna (Single Version)
6		The Jacksons - Triumph - 01 - Can You Feel It
7		Michael Jackson - Off the Wall - 05 - Off The Wall
8		Stevie Wonder - Song Review: A Greatest Hits Collection - 16 - Do I Do (1982 Musiquarium Version)
9		Michael Jackson - Off the Wall - 04 - Get On The Floor
10		Earth, Wind & Fire - Greatest Hits - 06 - Sing A Song
11		Commodores - The Ultimate Collection: The Commodores - 14 - Lady (You Bring Me Up)
12		Stevie Wonder - Number 1's - 15 - That Girl

Basic Deep Learning Architecture



Inference





Clusters in audio
feature space

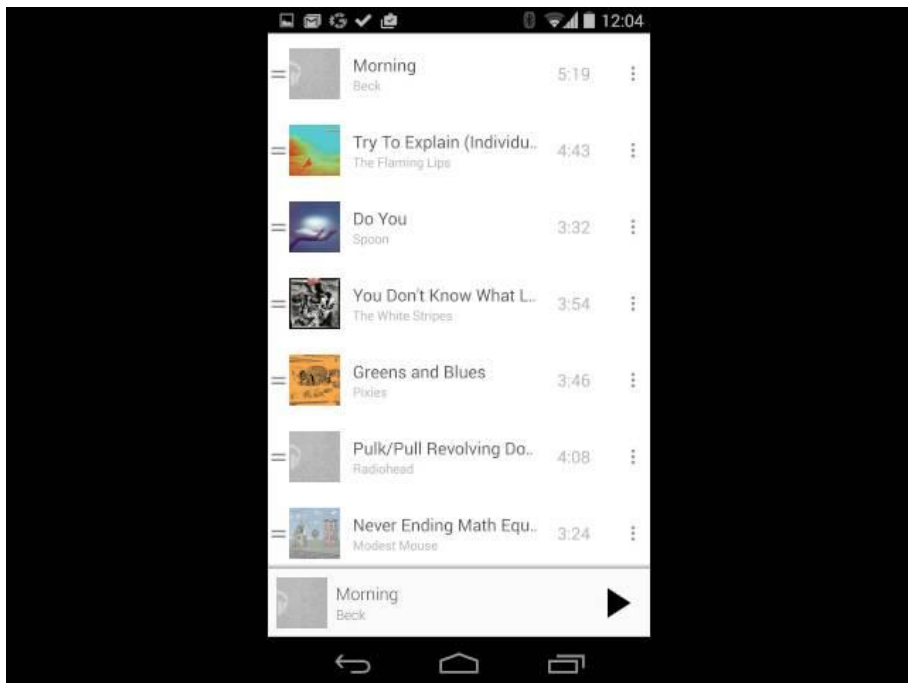
Learning
→



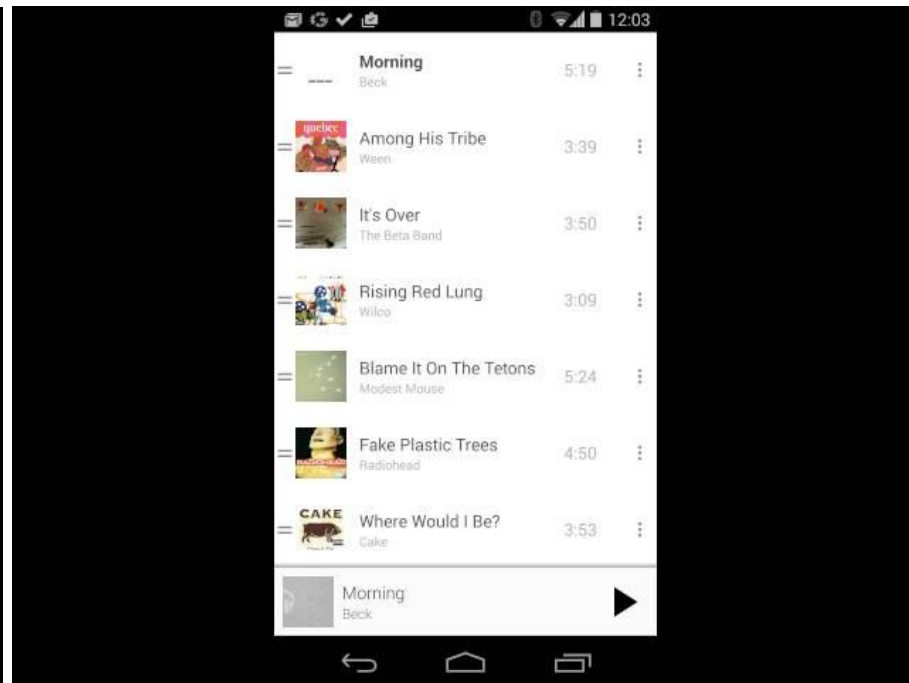
Semantic clusters
“Reggae”

Using audio in two-stage scoring

Collaborative filtering only.



Collaborative filtering + audio.



Knowledge and Connections



Gorillaz

Band

Gorillaz are an English musical and visual project created in 1998 by Damon Albarn and Jamie Hewlett. The project consists of Gorillaz itself and an extensive fictional universe depicting a "virtual band" of cartoon characters. [Wikipedia](#)

Members: [Damon Albarn](#), [Jamie Hewlett](#), [More](#)

Awards: Grammy Award for Best Pop Collaboration with Vocals, [More](#)






Origin: London, United Kingdom, Essex, United Kingdom

Record labels: Universal Music Group, EMI, Geffen Records, Virgin Records, Parlophone, Def Jam Recordings






Songs

Clint Eastwood	2001	Gorillaz
Feel Good Inc.	2005	Demon Days
Dirty Harry	2005	Demon Days
19/2000		
On Melancholy Hill	2010	Plastic Beach

Albums

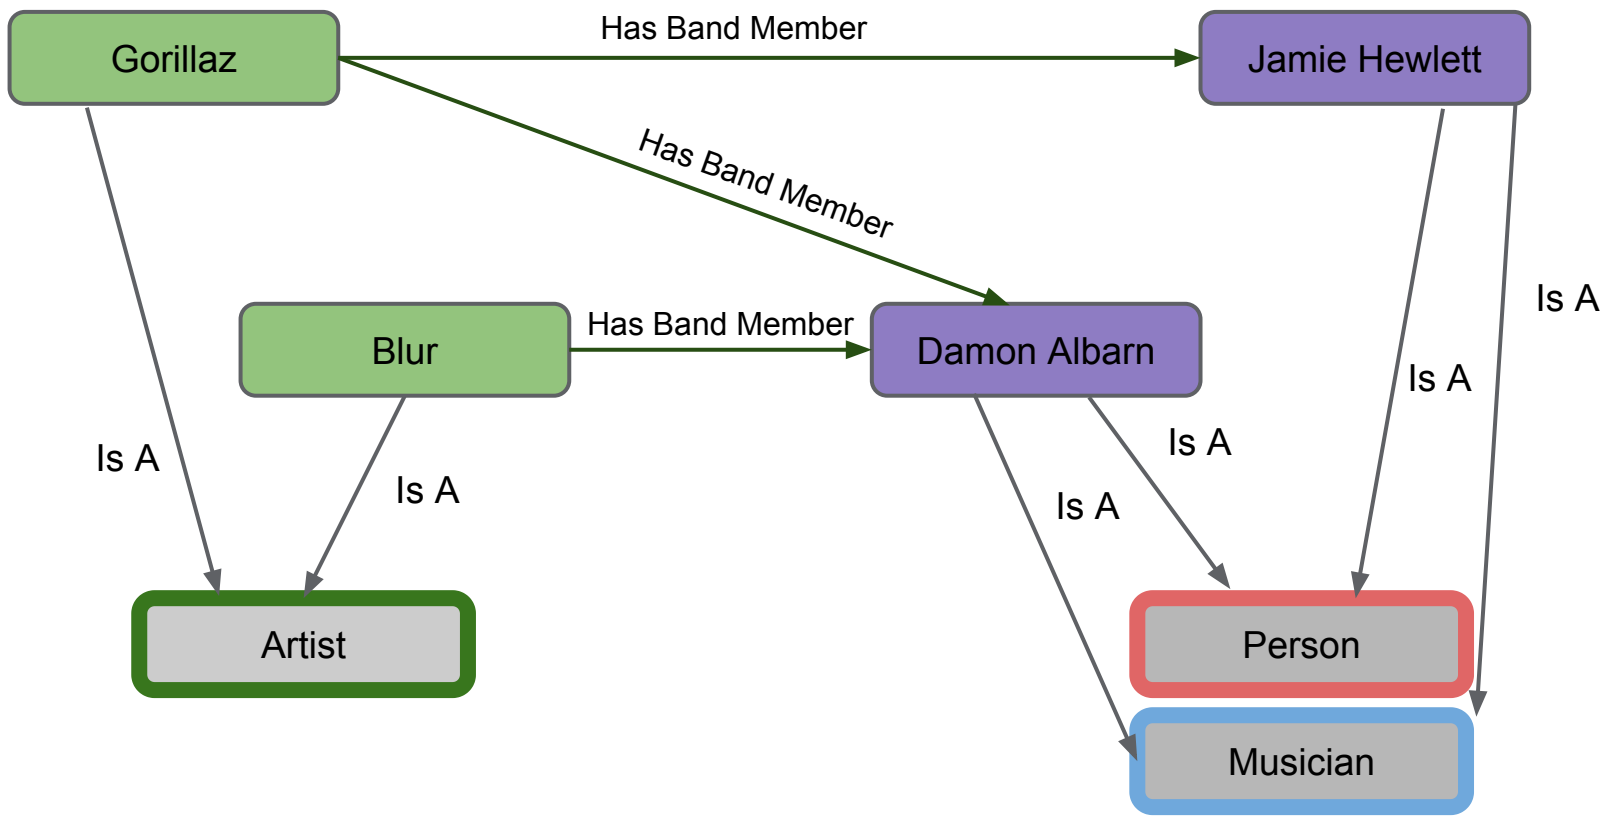
				
Plastic Beach	Demon Days	Gorillaz	The Fall	The Singles Collectio...
2010	2005	2001	2010	2011

People also search for

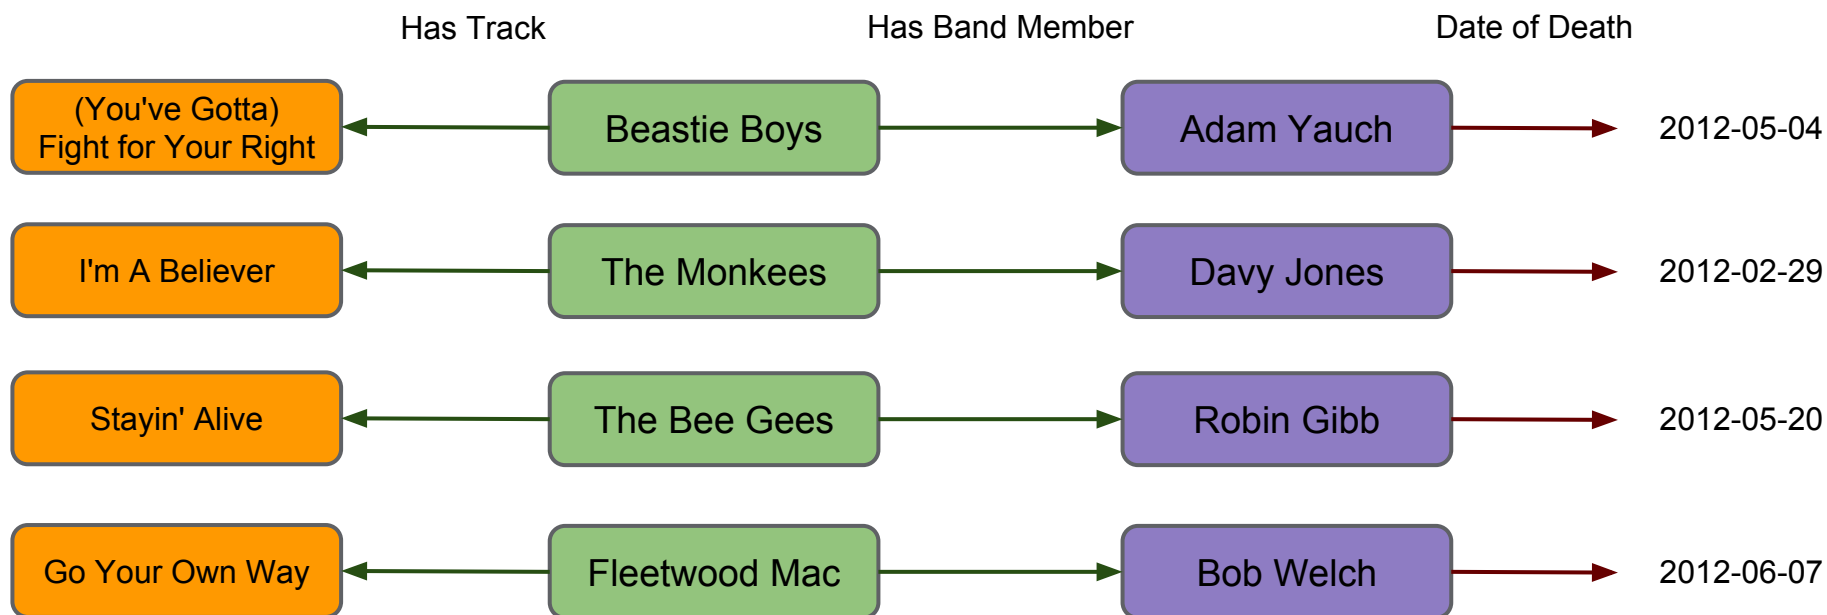
				
Noodle	Murdoc Niccals	2D	Blur	De La Soul

Knowledge Panel result from searching for "Gorillaz"

Freebase Music Schema



Playlist: Songs by Artists who had a member die in 2012



Freebase is a public dataset, so try it yourself!
<http://tinyurl.com/bj4bkdv>

[1] TRACK /id/t120
Pixies “Gigantic”
ROCK, ALT-INDIE

[2] TRACK /id/t250
John Coltrane, “Lazy
Bird” JAZZ, BEBOP

[3] ARTIST /id/a212
Gilberto Gil,
JAZZ, SAMBA

Inverted search index

ID	Term	Document
1	TRACK	1,2
2	JAZZ	2,3
3	John	1
4	/id/t250	2

- Map document terms onto documents.
- Fast **retrieval** of candidates for a query.
- Candidates are then ranked by **scoring** on popularity, relevance, etc.

Using a search index for music recommendation

- Index both text and embeddings
- Score based on
 - embeddings similarity
 - textual similarity
- Several strategies for limiting retrieval
 - Popularity tokens (top_track_by_artist)
 - Locality sensitive hashing tokens

```
TRACK [2]
Lazy Bird /id/250
Blue Train /id/B983
John Coltrane /id/A654
```

jazz, bebop, laid back

“John Coltrane was born...”

<CF embedding>

<Audio embedding>

Popularity: .012

top_track_by_artist

Bibliography

- Collaborative Filtering for Implicit Feedback Datasets. Y. Hu, [Y. Koren](#), C. Volinsky. ICDM 2008.
- WSABIE: scaling up to large vocabulary image annotation. [J. Weston](#), [S. Bengio](#), N. Usunier. IJCAI 2011.
- Efficient Estimation of Word Representations in Vector Space. T. Mikolov, K. Chen, [G. Corrado](#), [J. Dean](#). arXiv:1301.3781 2013.
- DeViSE: A Deep Visual-Semantic Embedding Model. [A. Frome](#), [G. Corrado](#), [J. Shlens](#), [S. Bengio](#), [J. Dean](#), [M. Ranzato](#), T. Mikolov NIPS (2013).
- WTA The Power of Comparative Reasoning, [J. Yagnik](#), D. Strelow, [D. Ross](#), R. Lin International Conference on Computer Vision, IEEE (2011).
- The Need for Music Information Retrieval with User-Centered and Multimodal Strategies. C. [C.S. Liem](#), M. Müller, [D. Eck](#), G. Tzanetakis, A. Hanjalic. MIRUM '11, ACM (2011).
- Temporal pooling and multiscale learning for automatic annotation and ranking of music audio [P. Hamel](#), S. Lemieux, Y. Bengio, [D. Eck](#). ISMIR (2011).
- Automatic generation of social tags for music recommendation. [D. Eck](#), P. Lamere, T. Bertin-Mahieux, S. Green. NIPS (2008).
- Aggregate features and AdaBoost for music classification. J. Bergstra, [N. Casagrande](#), [D. Erhan](#), [D. Eck](#), B Kégl. Machine learning 65 (2-3), 473-484, (2006).
- Sound Retrieval and Ranking Using Sparse Auditory Representations. [R. Lyon](#), M. Rehn, [S. Bengio](#), [T. Walters](#), [G. Chechik](#). Neural Computation, vol. 22, 2390-2416 (2010)..
- Building Musically-relevant Audio Features through Multiple Timescale Representations. [P. Hamel](#), Y. Bengio, [D. Eck](#) ISMIR, 553-558 (2013).
- Large Scale Distributed Deep Networks. [J. Dean](#), [G. Corrado](#), [R. Monga](#), K. Chen, M. Devin, Q. Le, M. Mao, M. Ranzato, [A. Senior](#), P. Tucker, [K. Yang](#), and A. Ng.

Conclusions

- Catalog access is now a commodity
- “Lean forward” vs “Lean back”
- Exploration vs exploitation

- Hybrid content / collaborative filtering approach
- Structured factual data important
- Search index provides unified approach
- User interface design remains a challenge

- Many open questions for neuroscience, cognitive psychology, music tech.

Thanks for your attention!

- Questions: deck@google.com

