Modeling Influence and Innovation in Contemporary Western Music

Uri Shalit, Daphna Weinshall and Gal Chechik

ICML 2013

Case Study
Inference and Representation
DS-GA-1005, Fall 2015

What is the role of influence and innovation in art?

Las Meninas,
Velázquez 1656

Las Meninas,
Picasso 1957

What is the role of influence and innovation in art?

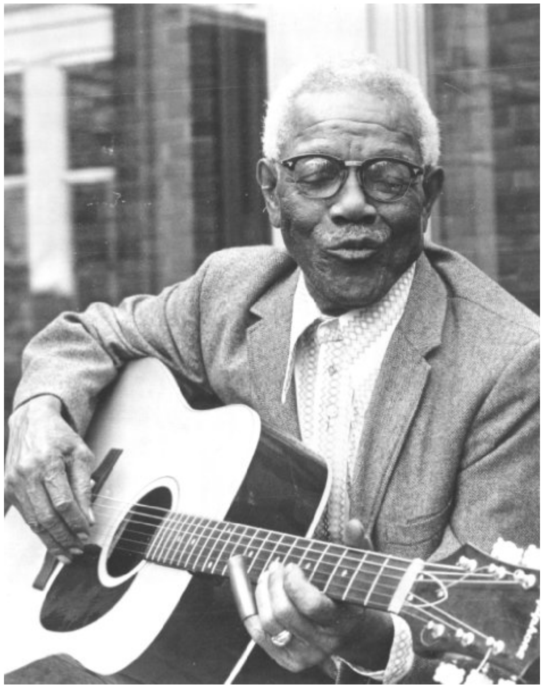

Furry Lewis
1927

Radiohead
1998

What is the role of influence and innovation in art?

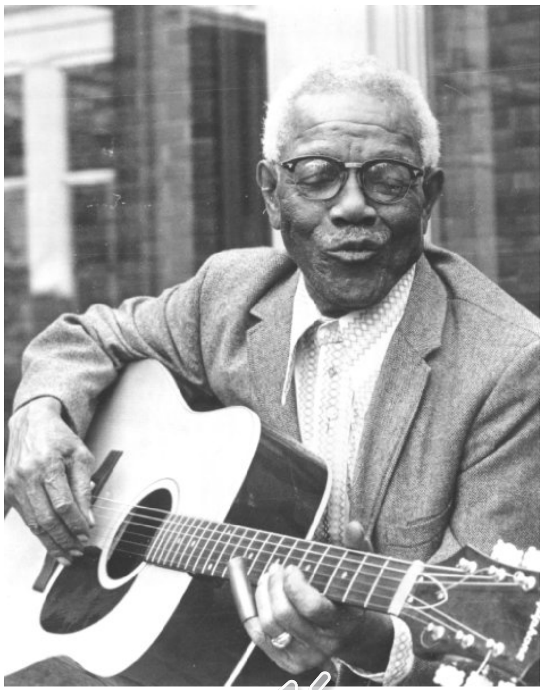

Furry Lewis
1927

Radiohead
1998

What is the role of influence and innovation in art?

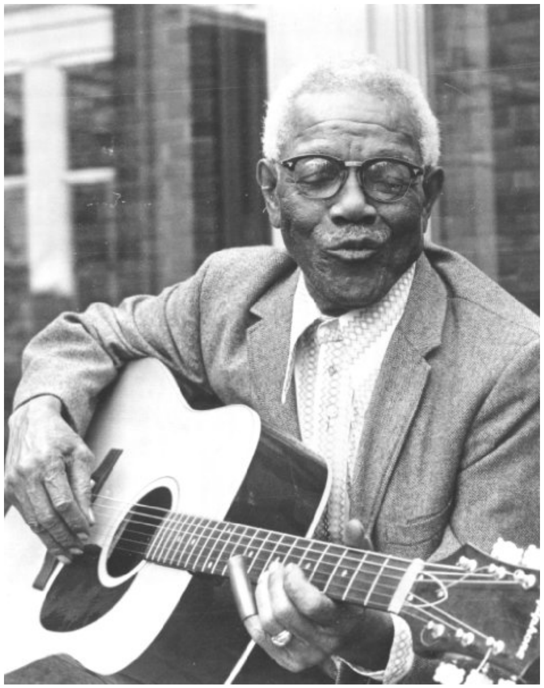

Furry Lewis
1927

Radiohead
1998

What is the role of influence and innovation in art?

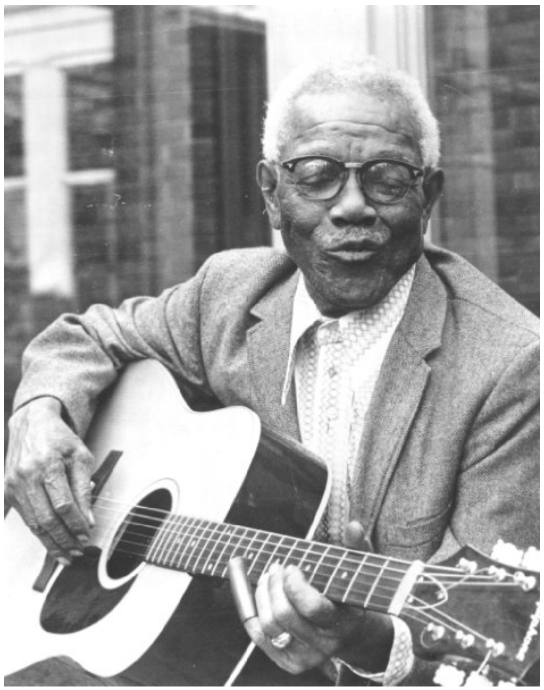

Furry Lewis
1927

Radiohead
1998

A (very!) hard perceptual problem

- How do people perceive music?
- Experts find individual artistic influence and innovation
- Difficult to grasp 10,000s of songs and diverse genres at once

Machine perception

- Process all the music in the world with a holistic view
- One step beyond object- and speech-recognition

Outline of our approach

The large scale data

25,000 songs, 9,000 artists, 70 years

The model

Unsupervised probabilistic topic models on acoustics

Evaluating perception

Properties of learned topics and influence

Analyzing innovation

Contrasting innovation with influence

Outline of our approach

The large scale data

25,000 songs, 9,000 artists, 70 years

The model

Unsupervised probabilistic topic models on acoustics

Evaluating perception

Properties of learned topics and influence

Analyzing innovation

Contrasting innovation with influence

The Million Songs Dataset

- Public dataset with 1,000,000 songs, each with detailed acoustic features
- Rich (but noisy) metadata: artist familiarity, genre tags

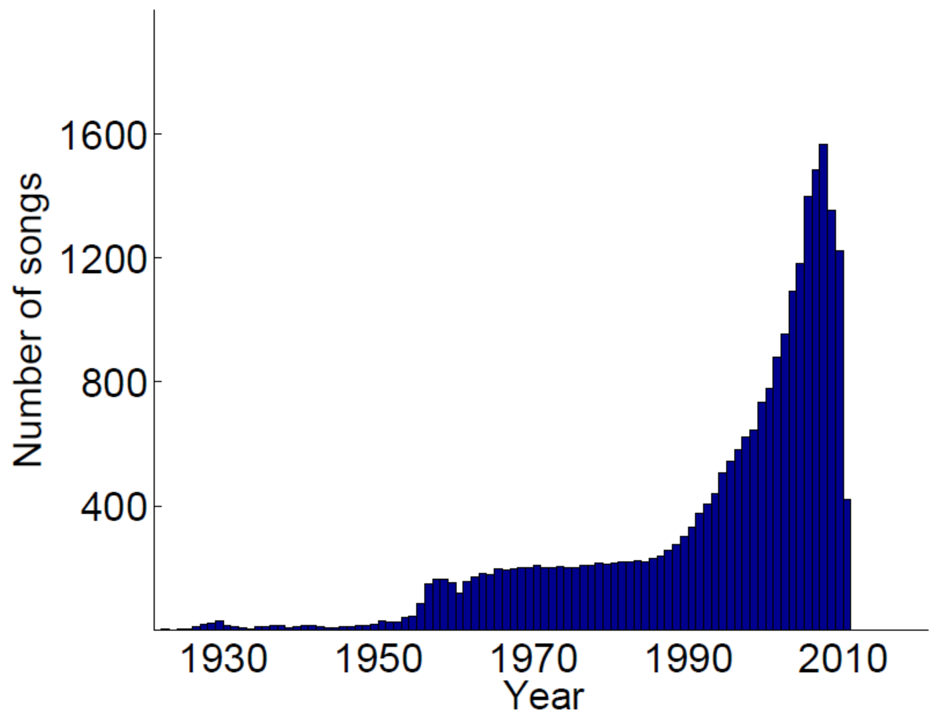

The Million Songs Dataset

NO BEATLES

Outline of our approach

The large scale data

25,000 songs, 9,000 artists, 70 years

The model

Unsupervised probabilistic topic models on acoustics

Evaluating perception

Properties of learned topics and influence

Analyzing innovation

Contrasting innovation with influence

Outline of our approach

The large scale data

25,000 songs, 9,000 artists, 70 years

The model

- Acoustic signal processing
- The probabilistic model – concept
- The probabilistic model - mathematically

Evaluating perception

Properties of learned topics and influence

Analyzing innovation

Contrasting innovation with influence

Acoustic Patterns

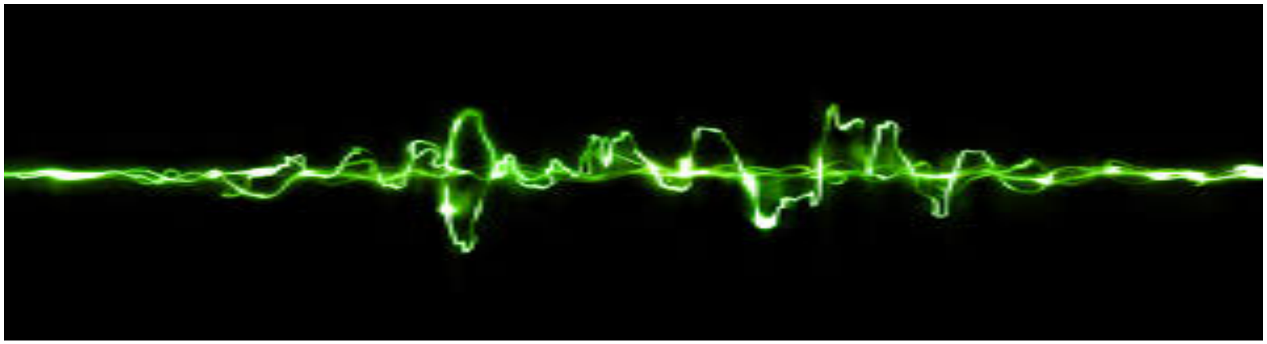

1 sec

Acoustic Patterns

X

[redacted]

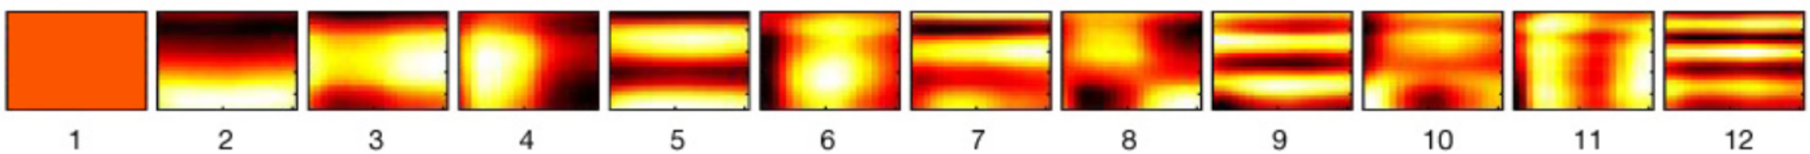

[redacted]

12 basis functions for timbre-descriptor

0.1	0.2	0.8	0.7	0.7	0.2	0.3	0.1	0.5	0.6	0.1	0.2

Acoustic Patterns

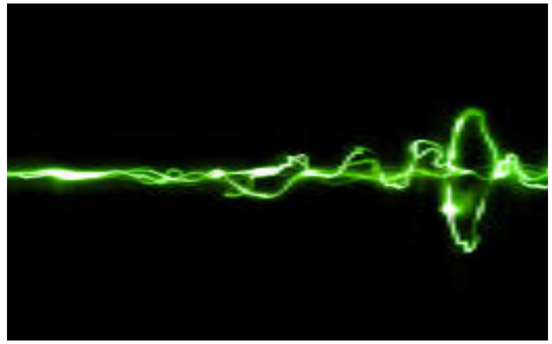

X

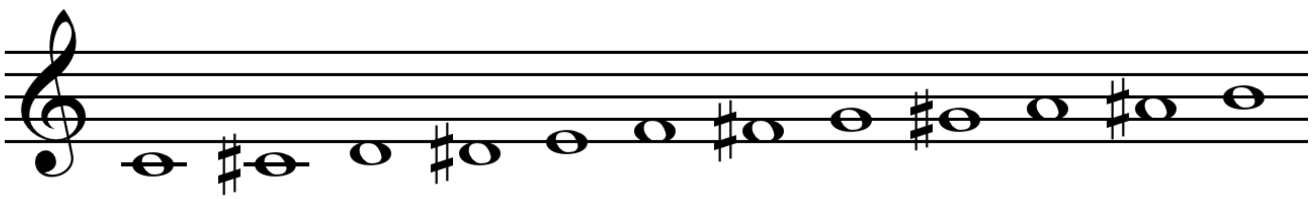

=

12 relative values for pitch-descriptor (tonal scale)

0.4	0.2	0.9	0.1	0.1	0.5	0.6	0.2	0.7	0.4	1.0	0.1

Acoustic Patterns

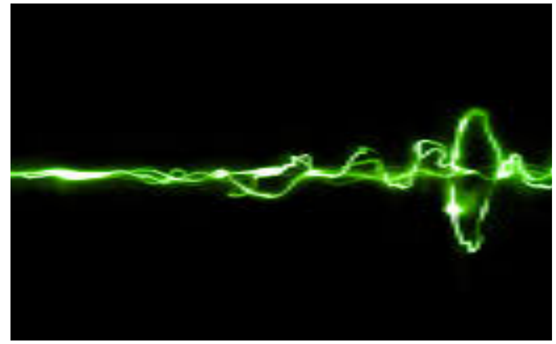

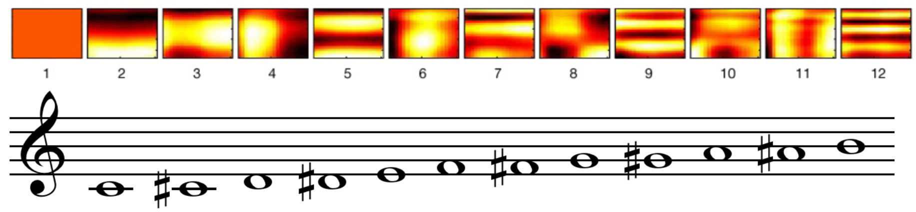

Cluster descriptors
into 5000 patterns

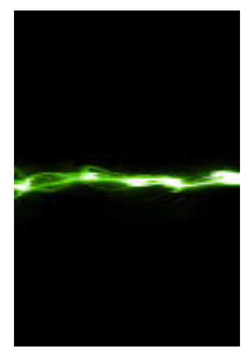

Outline of our approach

The large scale data

25,000 songs, 9,000 artists, 70 years

The model

- Acoustic signal processing
- The probabilistic model – concept
- The probabilistic model - mathematically

Evaluating perception

Properties of learned topics and influence

Analyzing innovation

Contrasting innovation with influence

Probabilistic Generative Model

Assume there is a complex probability distribution generating the **observed data**

Find the distribution parameters which maximize the likelihood of the observed data

Gerrish, S., & Blei, D. M. (2010). A language-based approach to measuring scholarly impact. Proceedings of the 27th International Conference on Machine Learning (ICML-10)

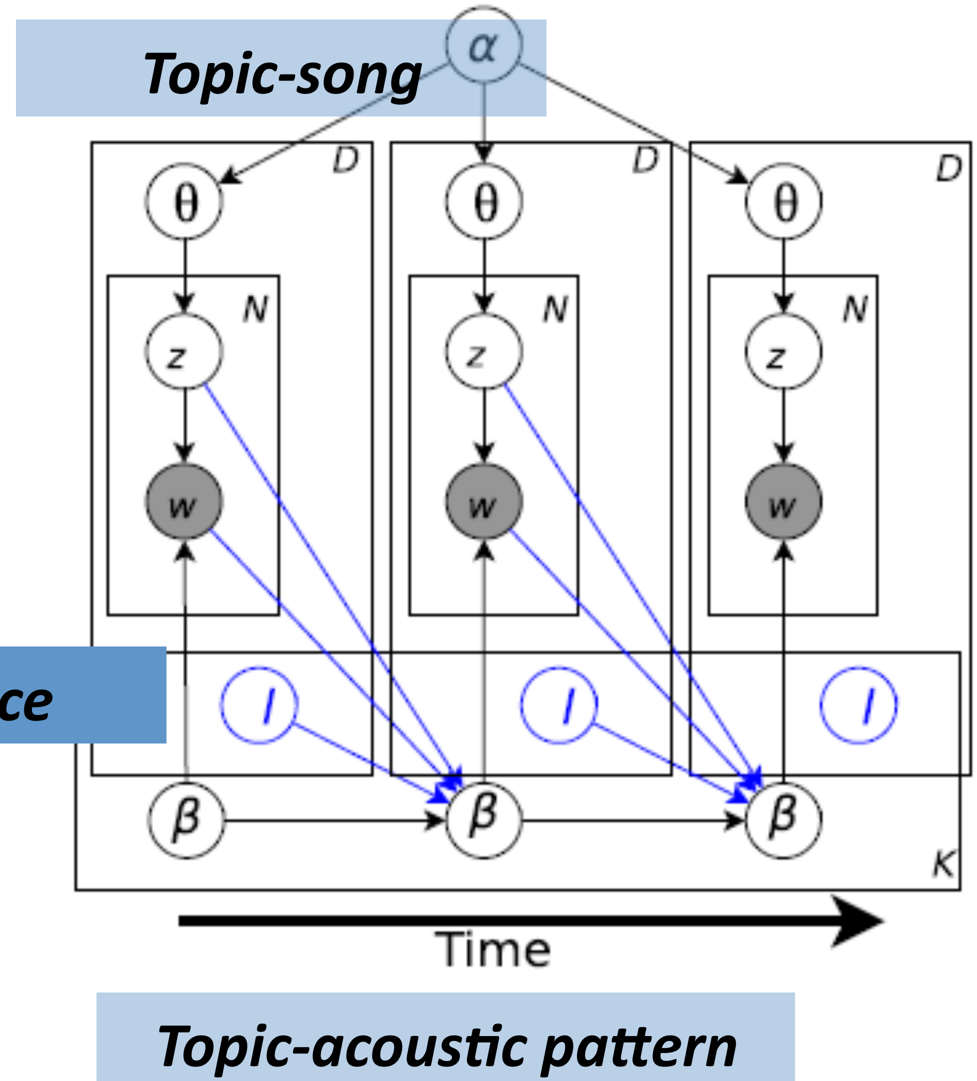

Genres → Topics

Humans understand music through **genres**

Automatically discover genre structure →**topics**

What defines a topic?

- Each topic is a distribution over *acoustic patterns*

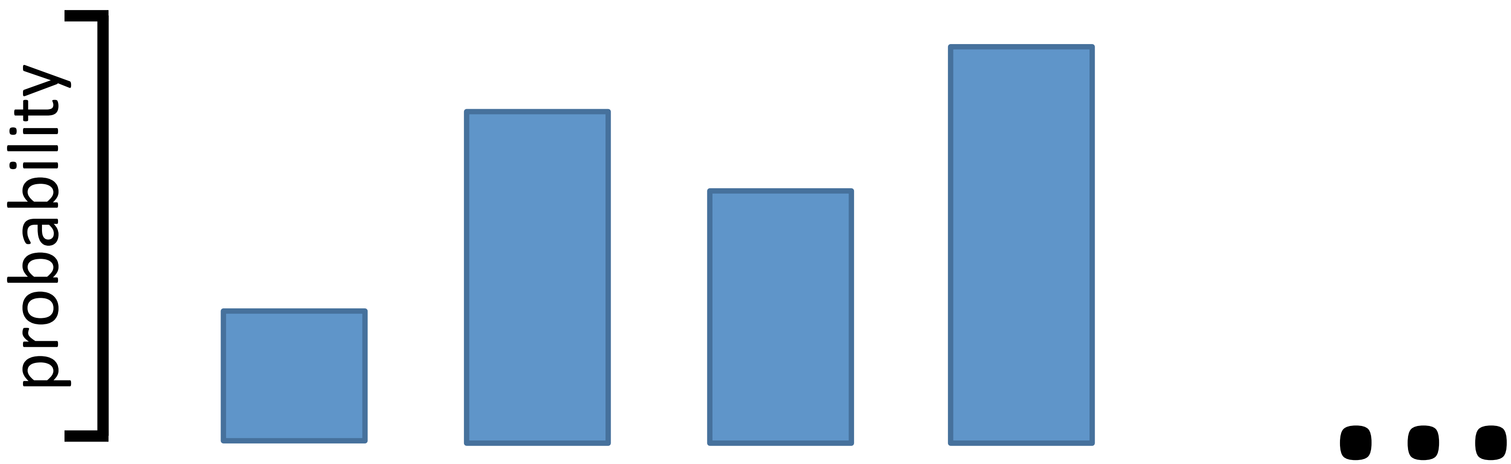

What defines a topic?

- Each topic is a distribution over *acoustic patterns*

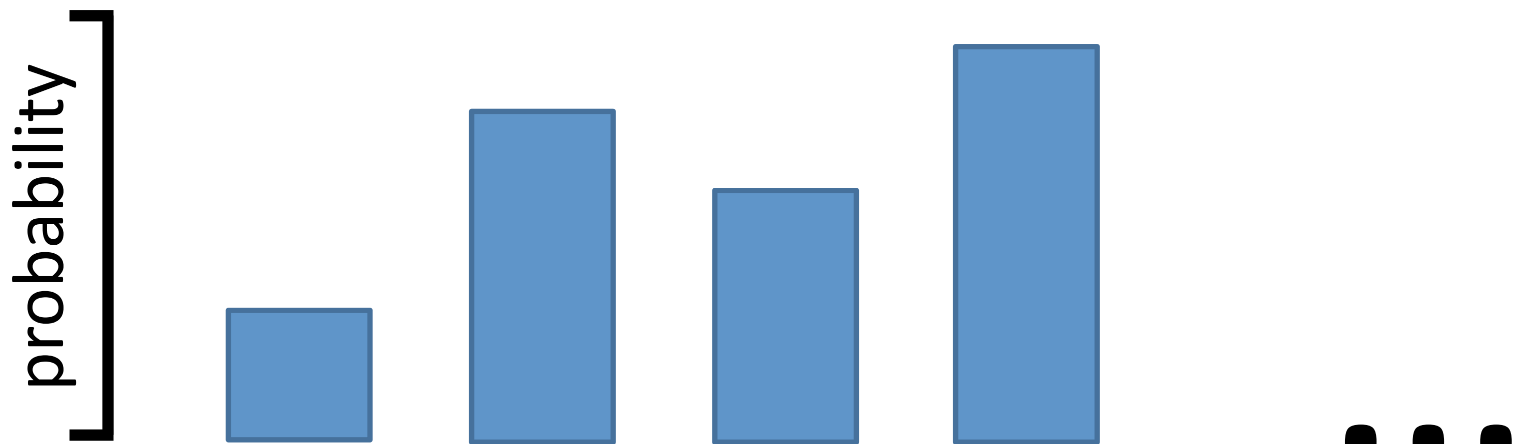

What defines a topic?

- Each topic is a distribution over *acoustic patterns*

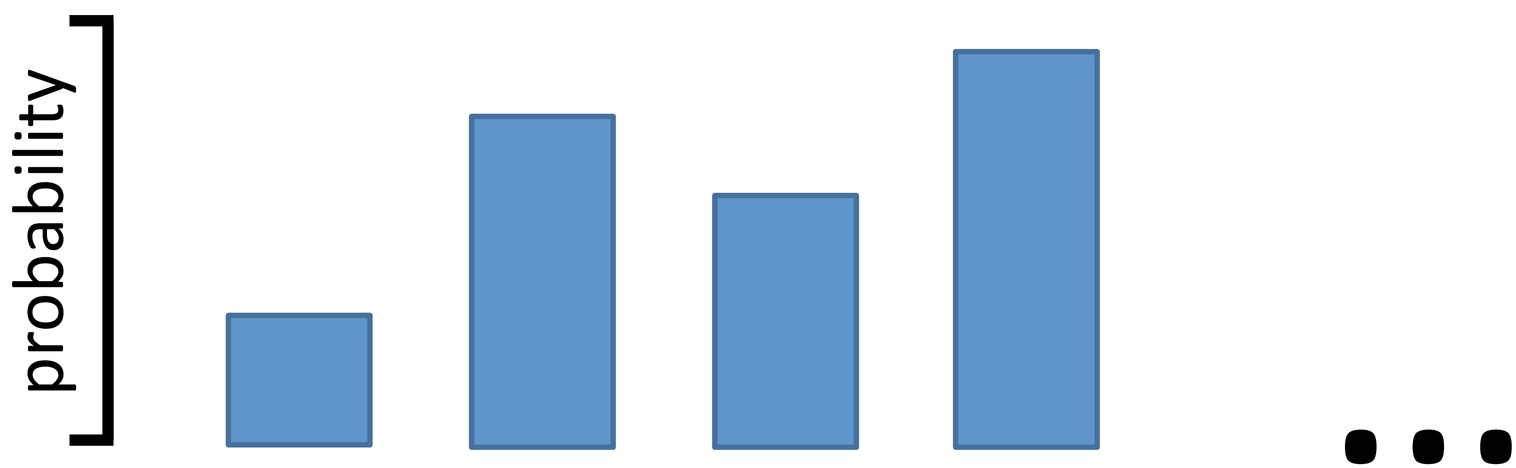

What defines a topic?

- Each topic is a distribution over *acoustic patterns*

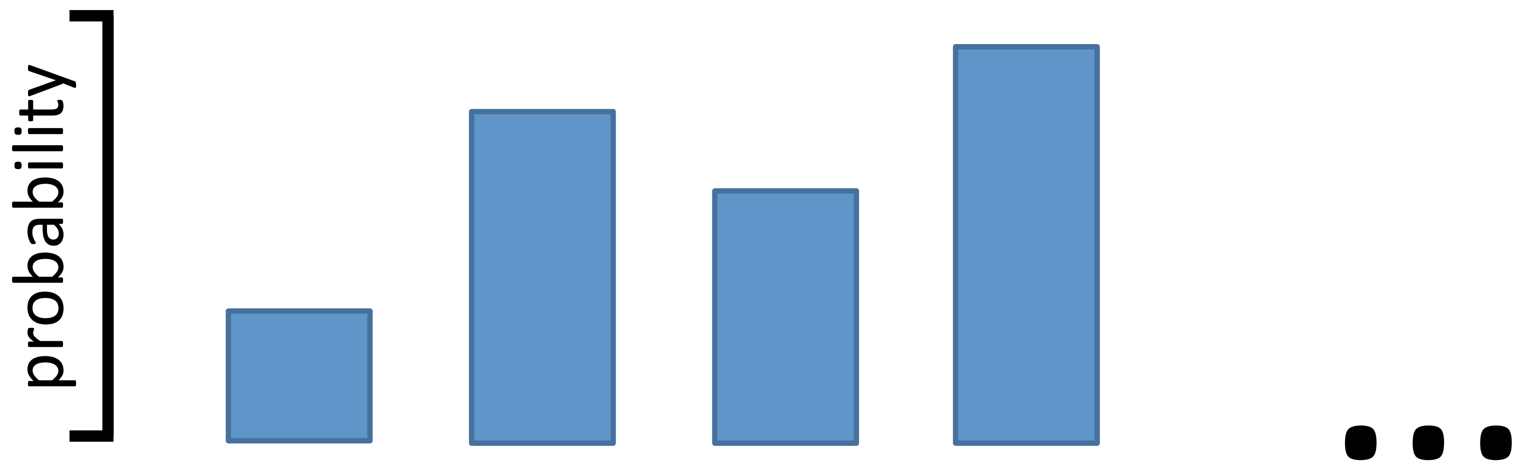

What defines a topic?

- Each topic is a distribution over *acoustic patterns*

Topic 1

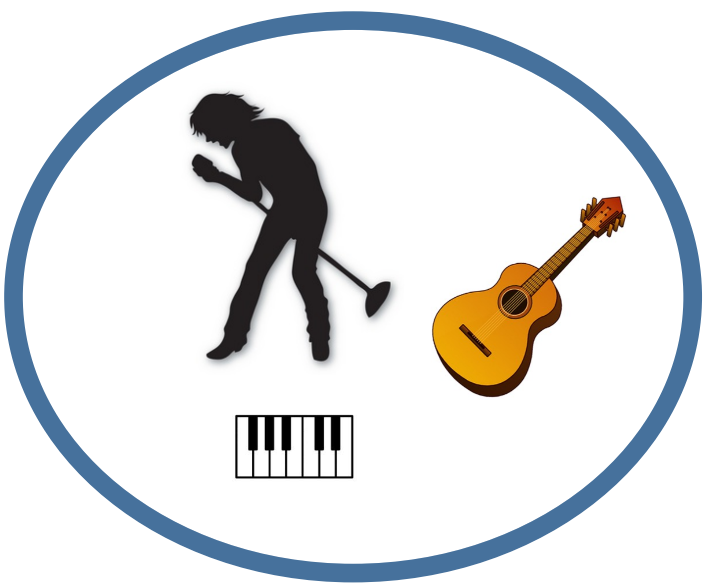

Topic 2

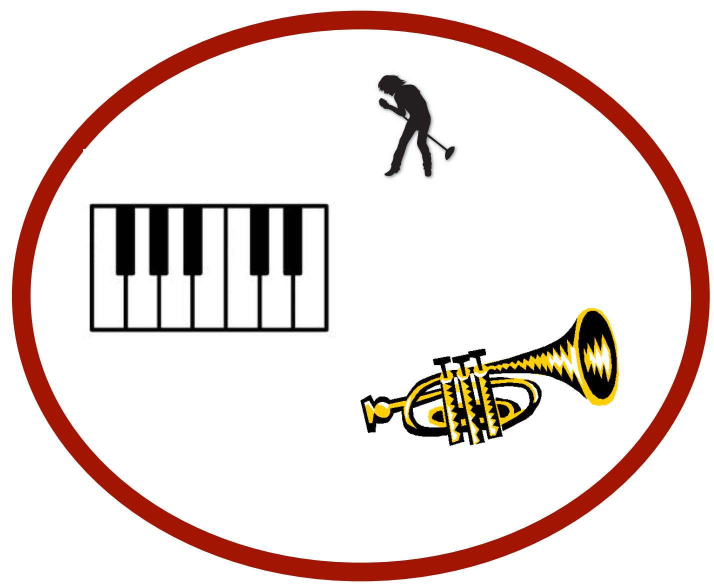

A song is a distribution over topics

- Model each song as a mixture of *topics*

A song is a distribution over topics

- Model each song as a mixture of *topics*
- For example Björk's 1995 song *Isobel*

$= 0.48 \times$ topic #3 $+ 0.38 \times$ topic #15
$+.14$ (18 other topics)

Topic 3 is about pop & classical
Topic 15 is about electronic & hip-hop

A song is a distribution over topics

- Model each song as a mixture of *topics*
- For example Björk's 1995 song *Isobel*

$= 0.48 \times$ topic #3 $+ 0.38 \times$ topic #15

$+.14$ (18 other topics)

Topic 3 is about pop & classical

Topic 15 is about electronic & hip-hop

A song is a distribution over topics

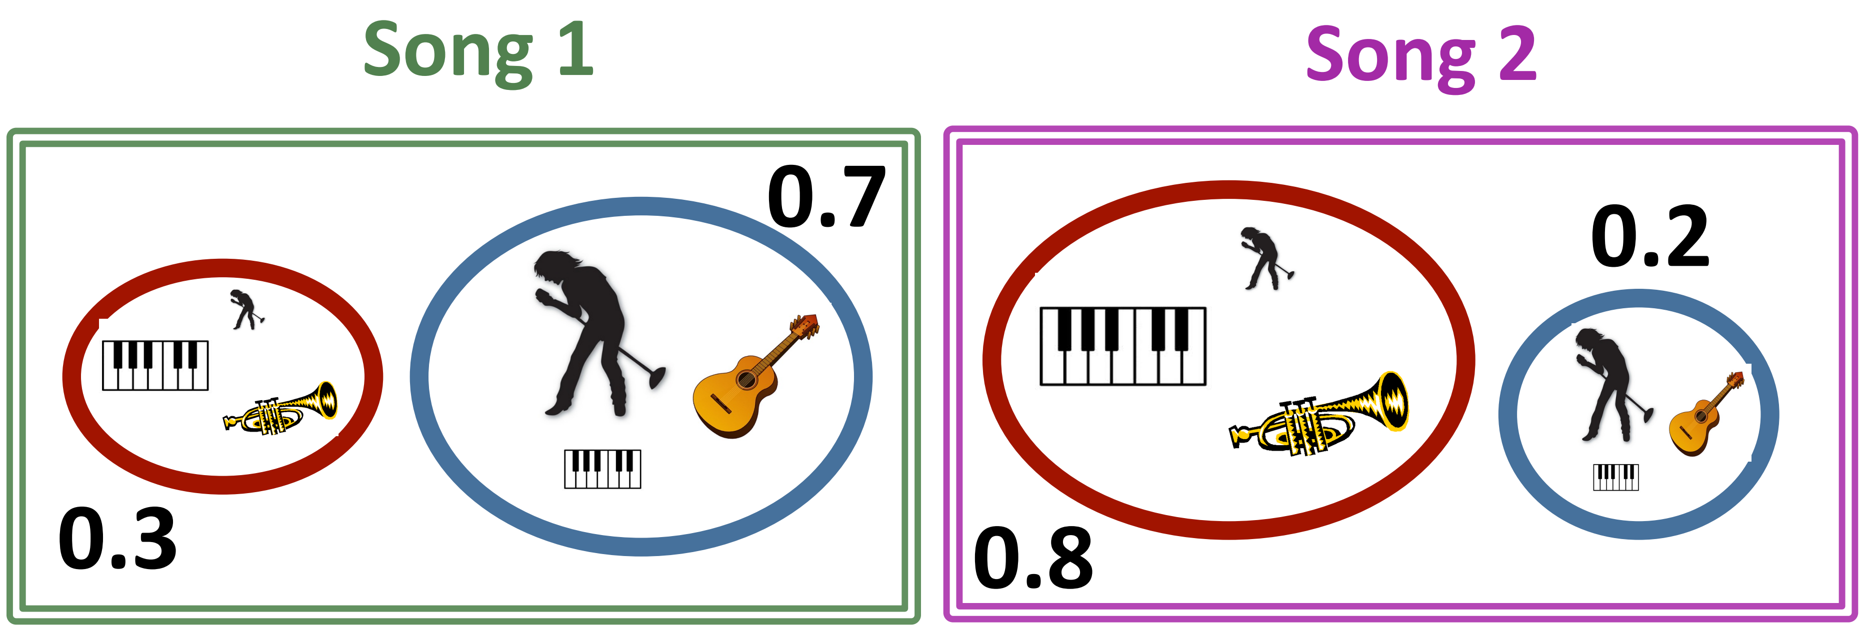

Topics evolve over time – topic #18

1956 Hound Dog
by Elvis Presley

1968 Big Sky
by The Kinks

1988 Michelle
by Guns'n'Roses

Songs influence the evolution of topics

In 1965, Bob Dylan switches from acoustic to electric guitar:

Songs influence the evolution of topics

In 1965, Bob Dylan switches from acoustic to electric guitar:

Songs influence the evolution of topics

In 1965, Bob Dylan switches from acoustic to electric guitar:

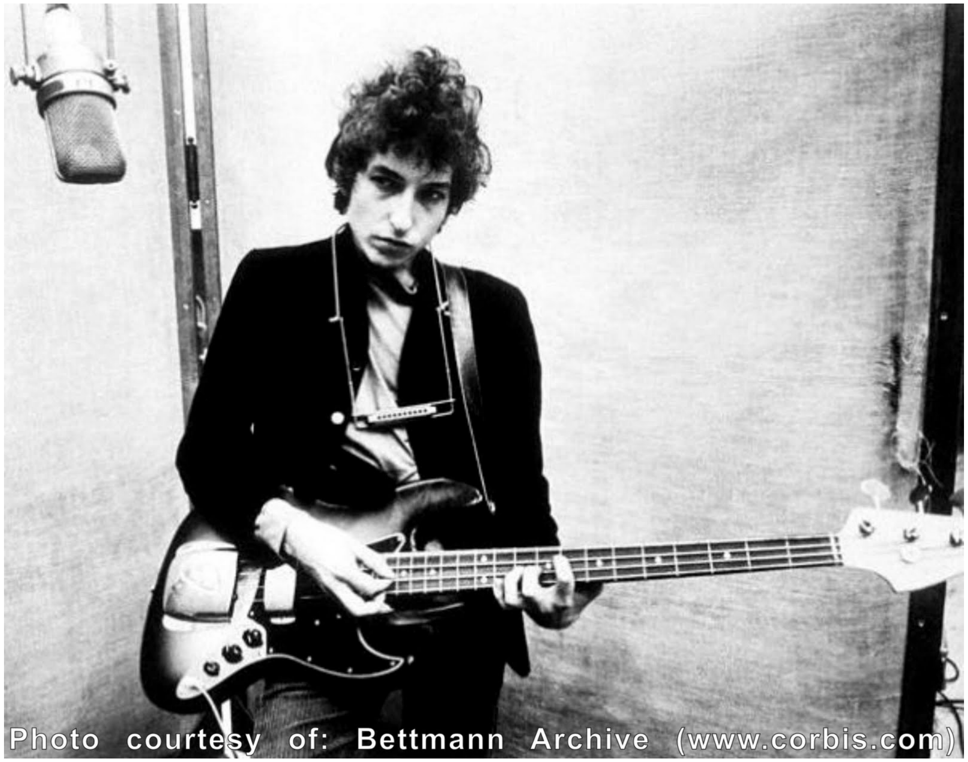

Songs influence the evolution of topics

In 1965, Bob Dylan switches from acoustic to electric guitar:

… and an entire topic goes with him, with artists such as The Velvet Underground and Jimmy Hendrix

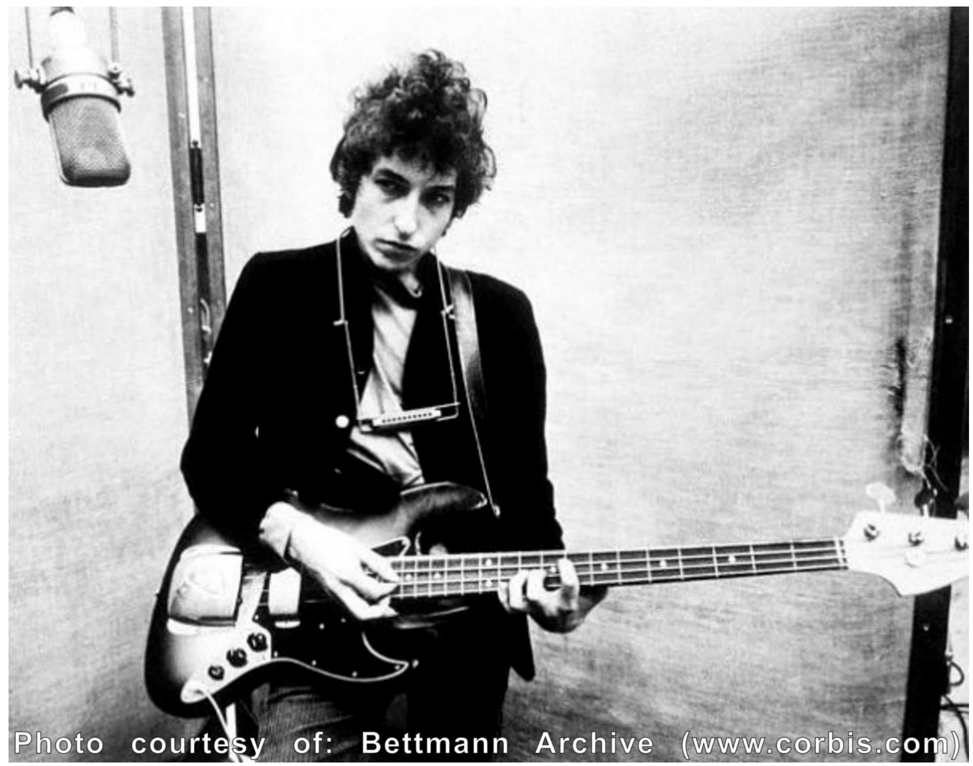

Each topic evolves in time, driven by the influence of songs

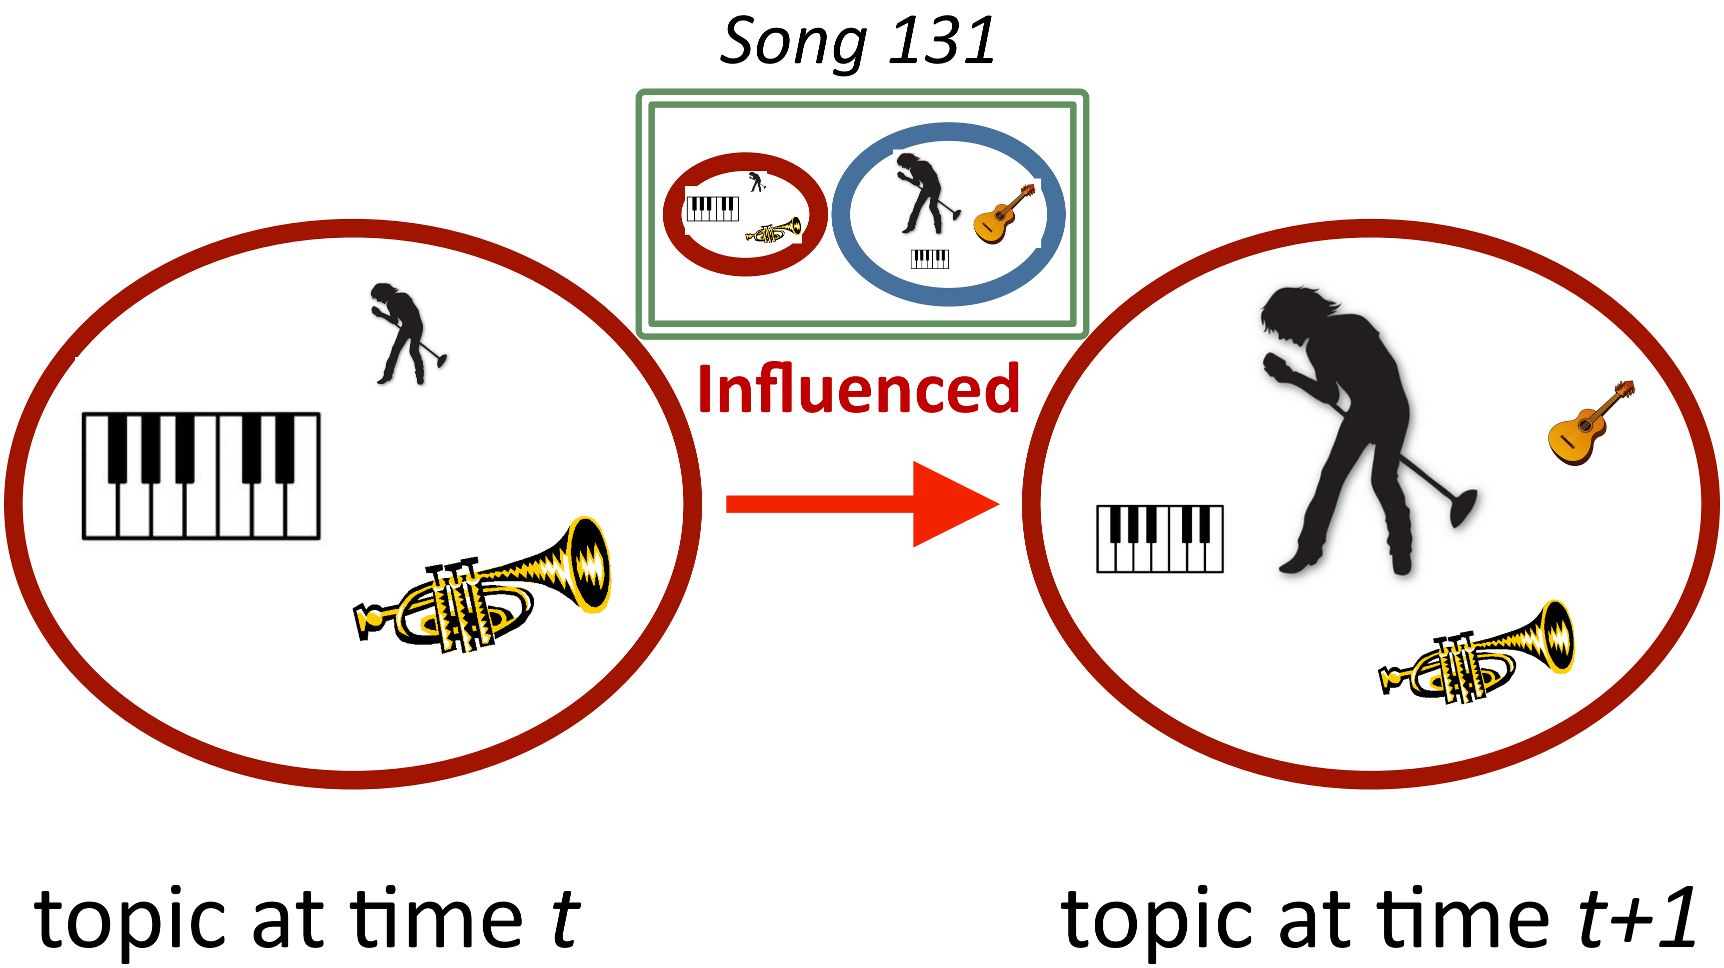

Probabilistic model overview

- Each topic is a distribution over acoustic patterns
- Each song is a distribution over topics
- Each topic evolves in time, driven by the influence scores of songs

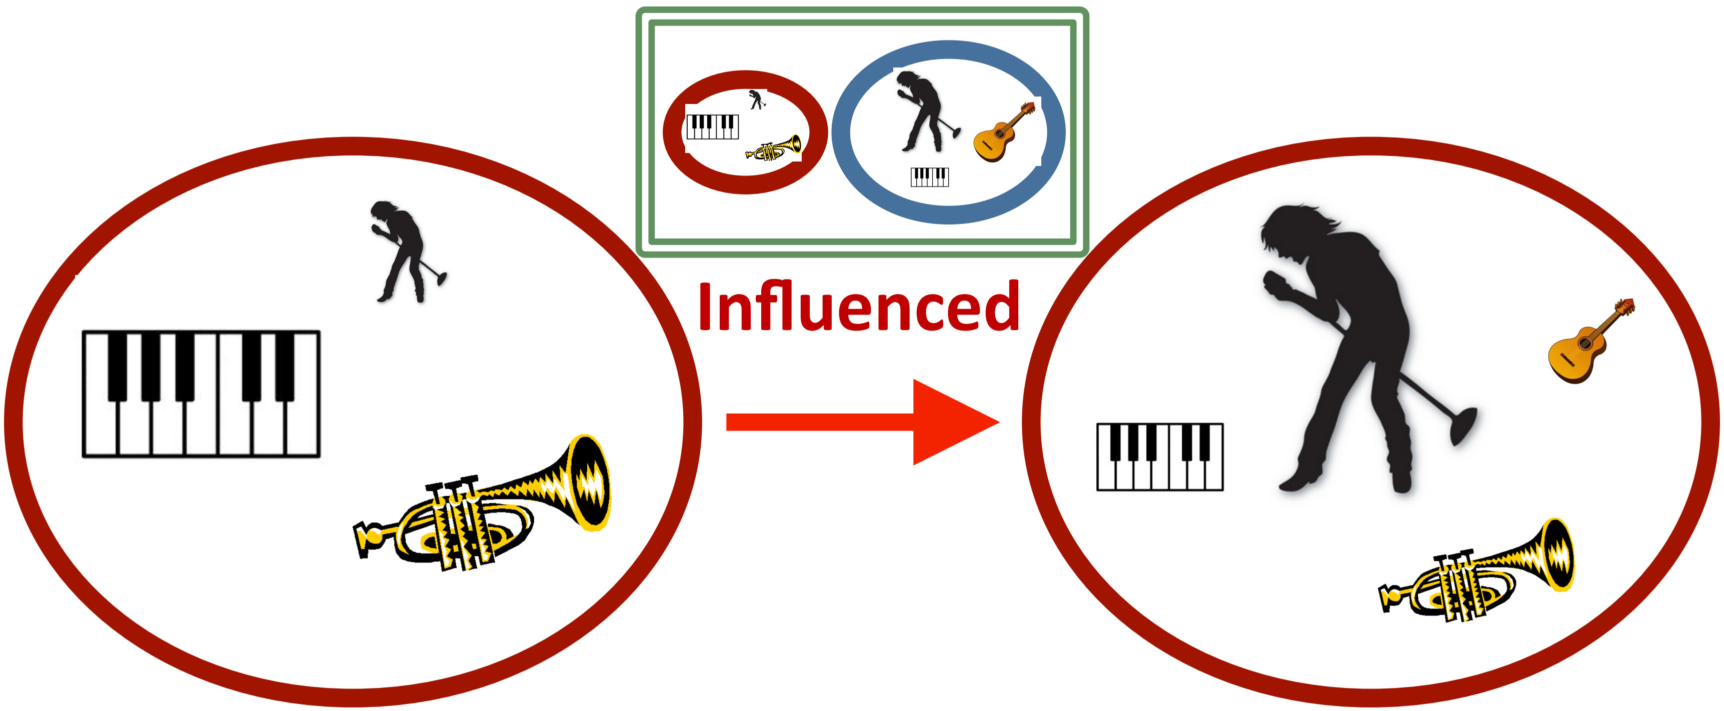

Outline of our approach

The large scale data

25,000 songs, 9,000 artists, 70 years

The model

- Acoustic signal processing
- The probabilistic model – concept
- The probabilistic model - mathematically

Evaluating perception

Properties of learned topics and influence

Analyzing innovation

Contrasting innovation with influence

Topic-Word Distribution

- For each year, each topic is a distribution over the 5033 acoustic patterns
- The distribution is parameterized by a vector $\beta_{k,t} \in \mathbf{R}^{5033}$
- The distribution is:

$$\Pr(word = i \mid \text{topic} = \text{k}, \text{time} = \text{t}) \propto \exp(\beta_{k,t}(i))$$

Song-Topic Distribution

- Each song *d* is assigned a multinomial distribution over the K topics $\theta_d \in \Delta^K$

Topic evolution – influence

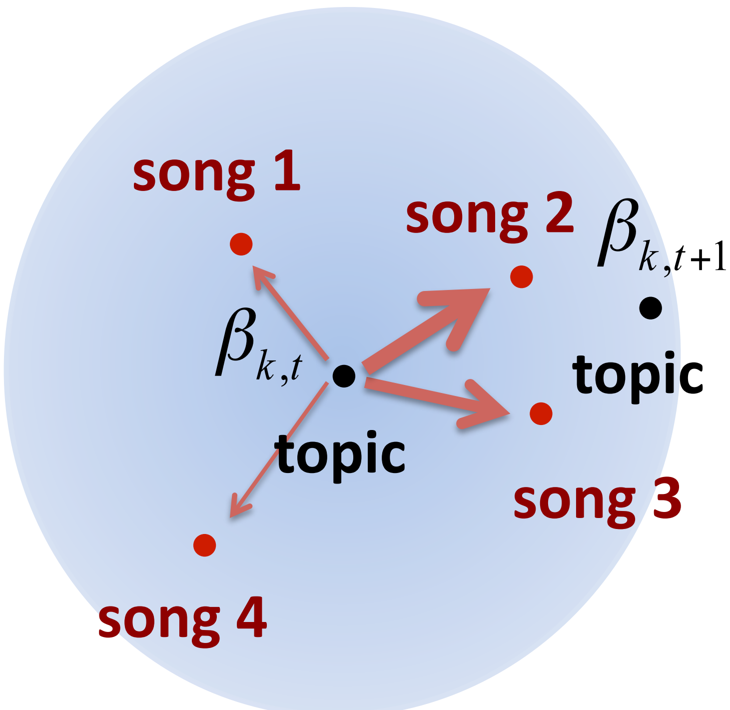

$$p\left(\beta_{k,t+1} \mid \beta_{k,t}\right) \sim N\left(\beta_{k,t} + \exp(-\beta_{k,t}) \cdot \gamma_{k,t}, \sigma^2\right)$$

$\gamma_{k,t}$ = weighted sum of influences

Probabilistic Generative Model
Gerrish & Blei (2010)

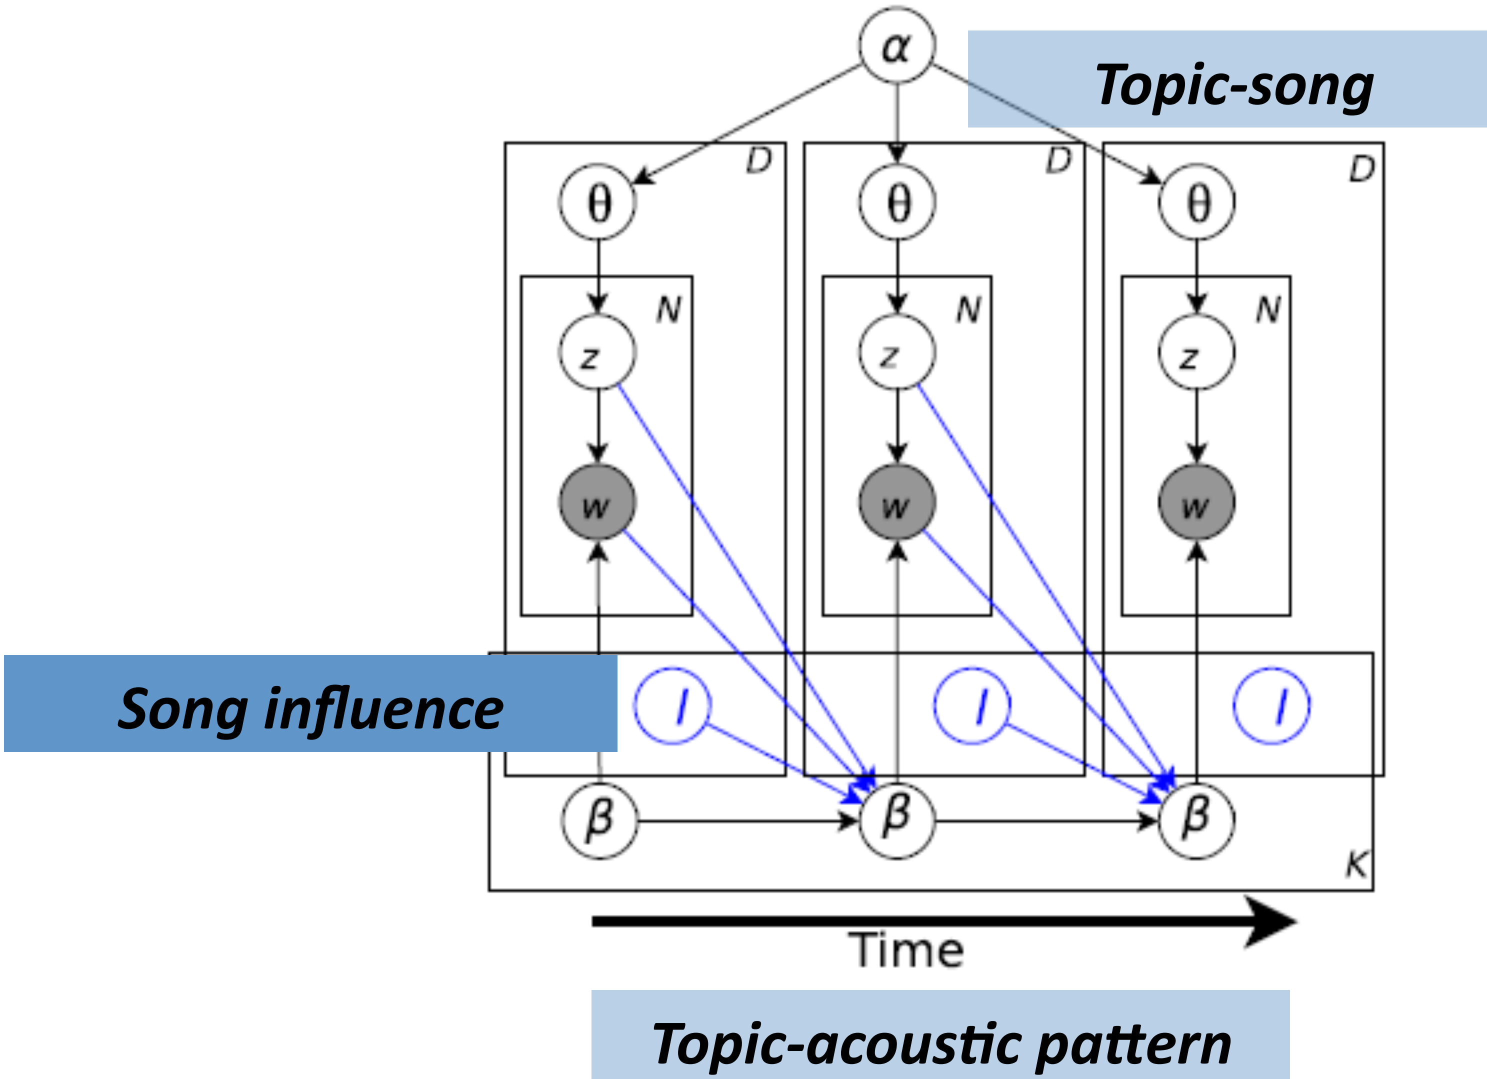

Probabilistic model - summary

- Each topic is a distribution over acoustic patterns:
  $\beta_{k,t}$ for all topics *k* and times *t*
- Each song is a distribution over topics:
  θ_{song} for all songs
- Each topic evolves in time, driven by the influence scores l_{song} for all songs

Optimization problem is intractable, solved approximately by introducing variational parameters and optimized using block-coordinate ascent. Code available at: https://github.com/Blei-Lab/dtm

maximize $pr(\text{corpus} \mid \beta, \theta, l)$

over β, θ, l

for all *topics*, *times* and *songs*

Outline of our approach

The large scale data

25,000 songs, 9,000 artists, 70 years

The model

Unsupervised probabilistic topic models on acoustics

Evaluating perception

Properties of learned topics and influence

Analyzing innovation

Contrasting innovation with influence

Results

Model's topics match known genres

(genre tags not used in training)

Results

Model's topics match known genres

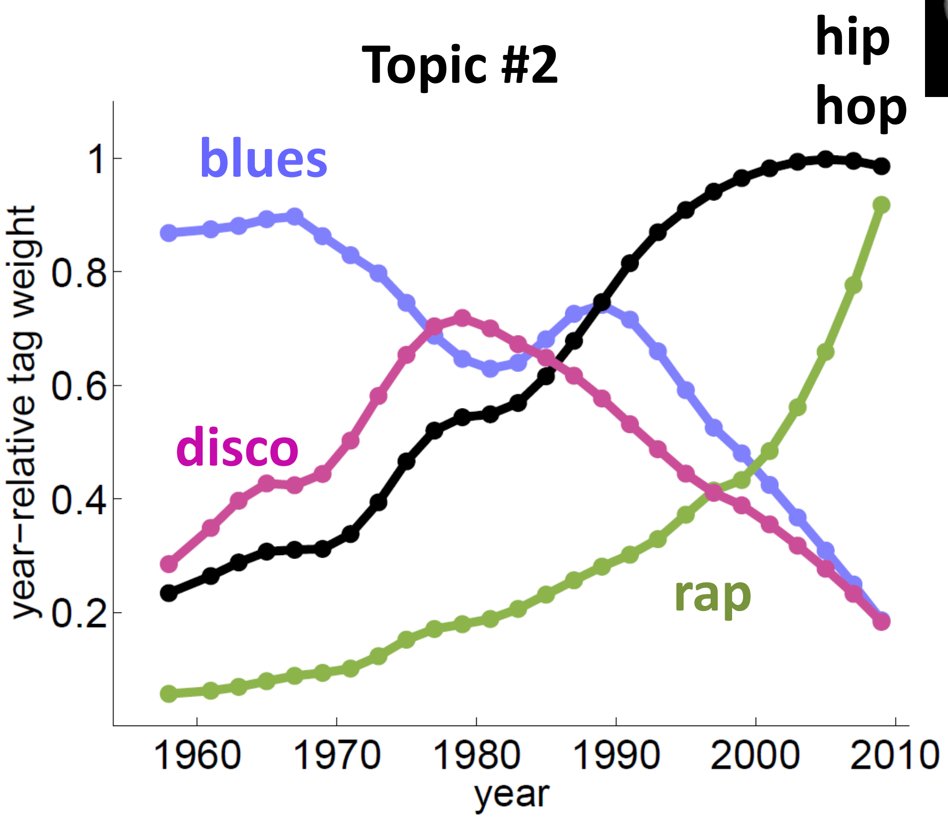

(genre tags not used in training)

Results

Model's topics match known genres

Jackson 5 – ABC (1970)

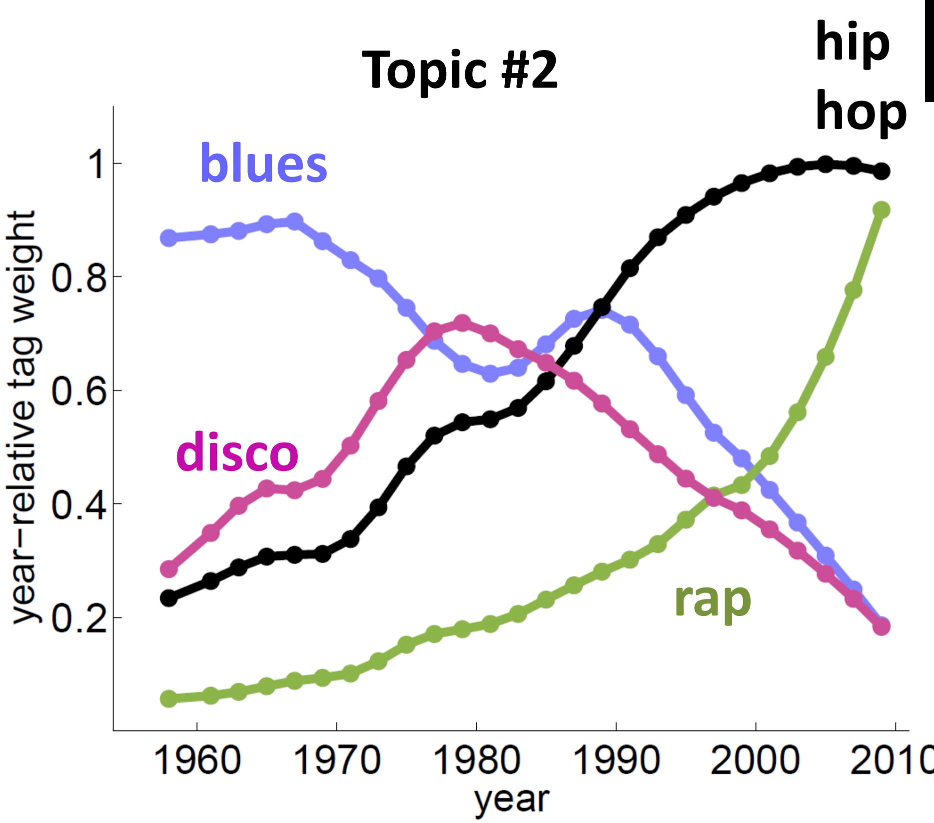

(genre tags not used in training)

Results

Model's topics match known genres

Public Enemy – Revolutionary Generation (1990)

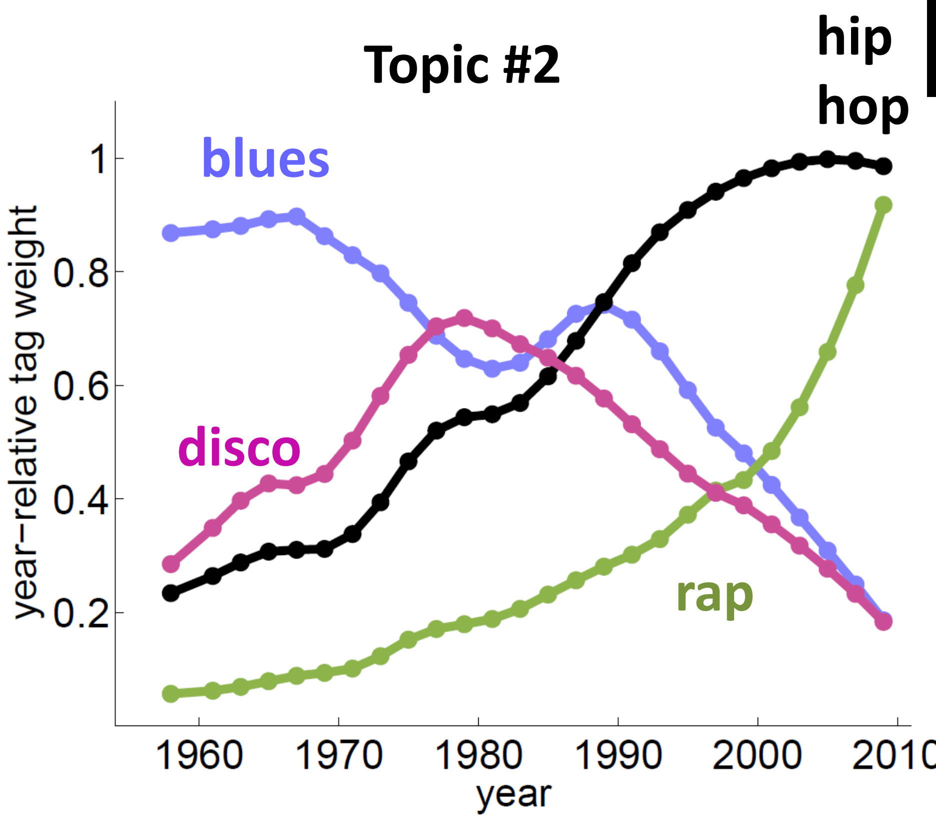

Jackson 5 – ABC (1970)

(genre tags not used in training)

Results

Model’s topics match known genres

Topic #13

The Who

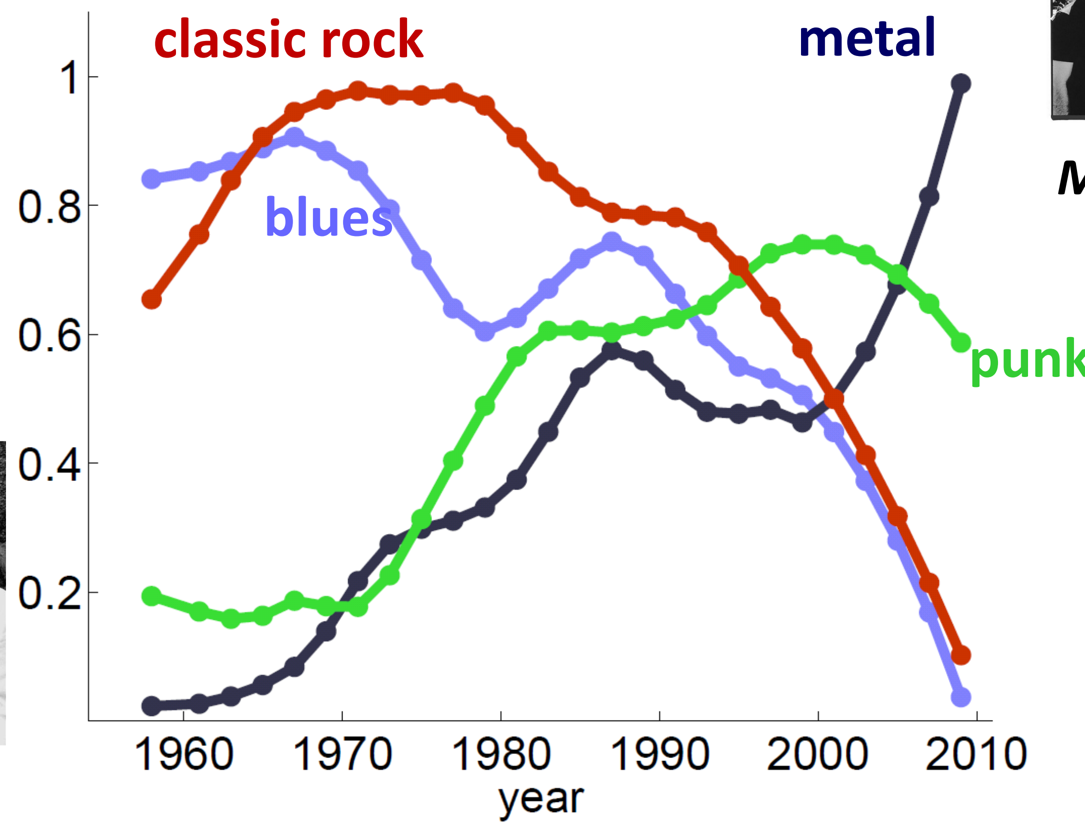

Metallica

(genre tags not used in training)

Results – validating influence measure

- Validating our influence measure against the influence graph of **allmusic.com**
- Figure shows $-\log_{10}$ p-value of Spearman correlation with allmusic.com influence measure

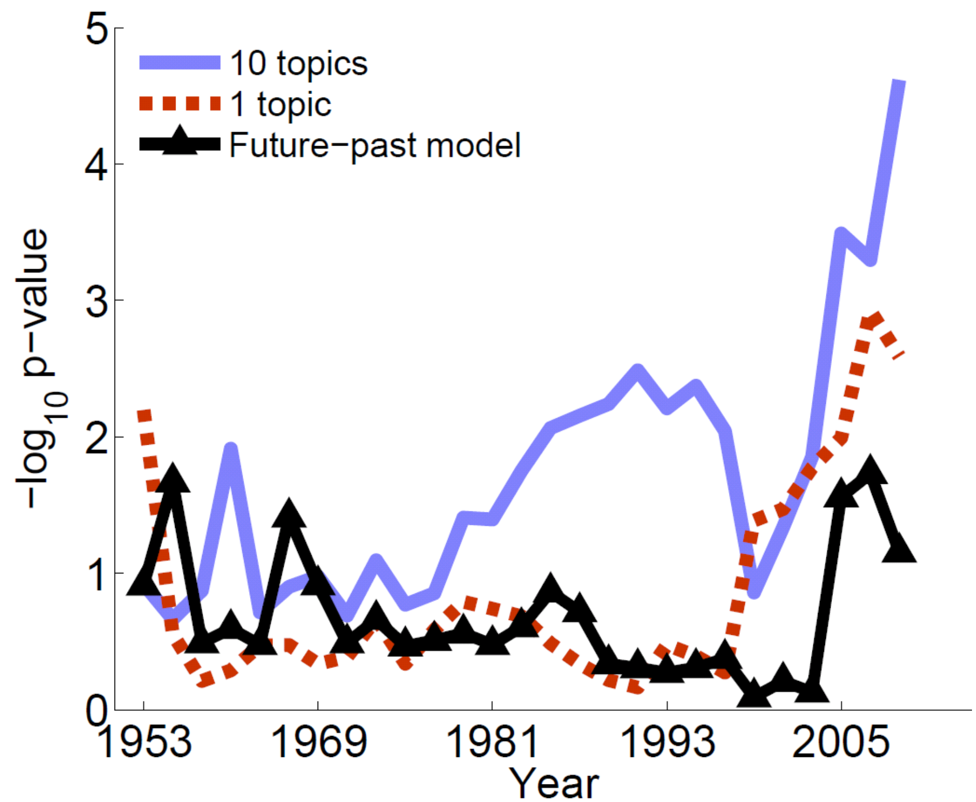

- Future-past model based on changes in musical-word frequencies between future and past

Results – validating influence measure

- Validating our influence measure against the influence graph of **allmusic.com**
- Figure shows $-\log_{10}$ p-value of Spearman correlation with allmusic.com influence measure

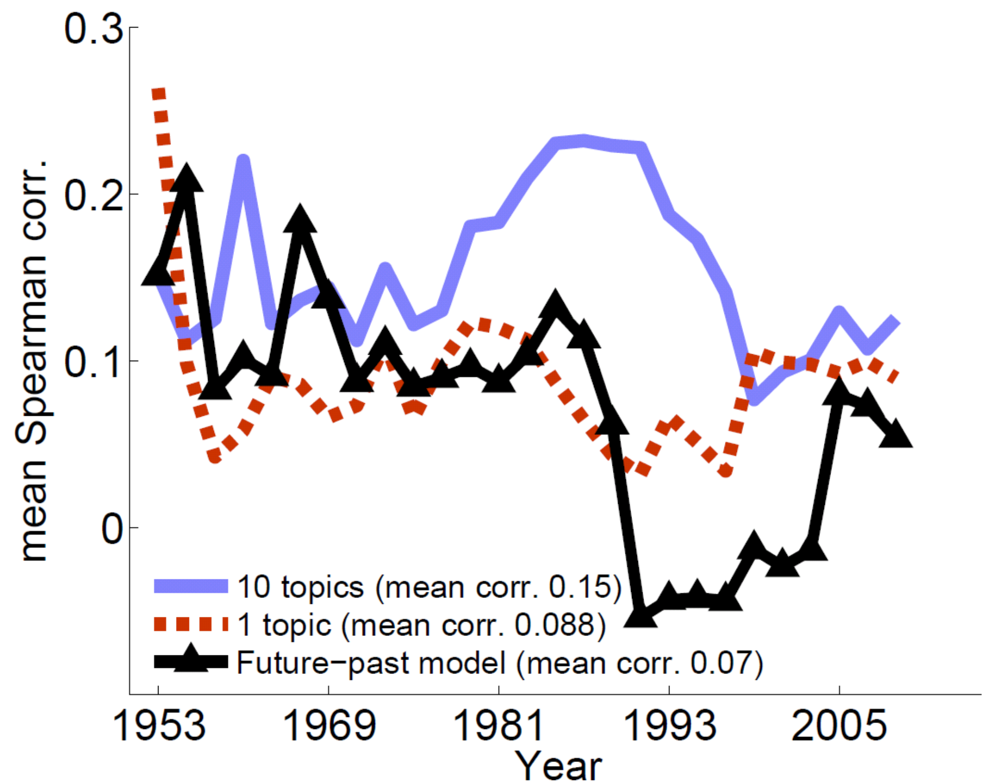

- Future-past model based on changes in musical-word frequencies between future and past

Results – examples of influential artists

- Many familiar artists: Bob Dylan, Rolling Stones, Bob Marley, Velvet Underground …
- But also lesser known artists:
  - *Model 500* "is widely credited as the originator of techno music"
  - *Killing Joke* "Finding modest commercial success, Killing Joke have influenced Nirvana, Metallica, Soundgarden…"
  - *Suicide* "Never widely popular amongst the general public, Suicide are highly influential… [many] sounds of the '80s and '90s gesture back to [Suicide]"

Outline of our approach

The large scale data

25,000 songs, 9,000 artists, 70 years

The model

Unsupervised probabilistic topic models on acoustics

Evaluating perception

Properties of learned topics and influence

Analyzing innovation

Contrasting innovation with influence

Musical innovation

How does being innovative relate to being influential?

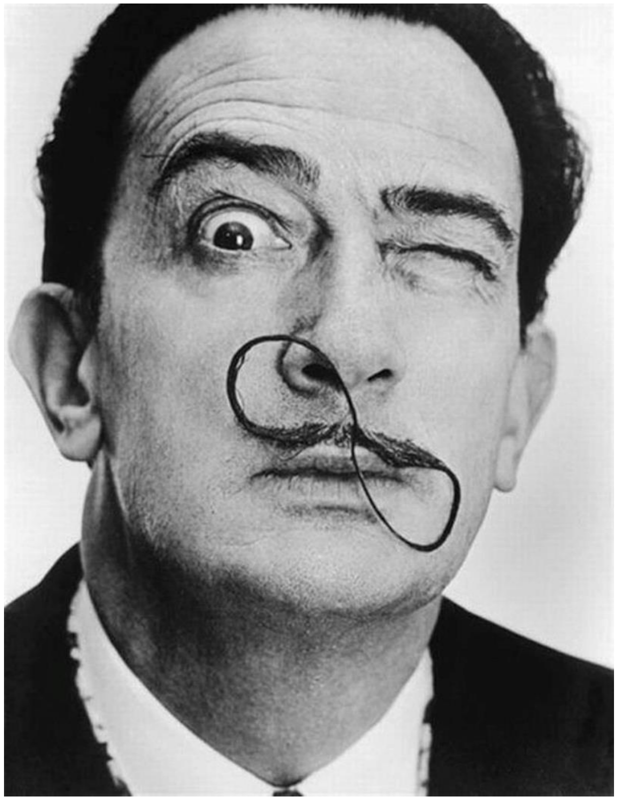

Musical innovation

- Probabilistic model: a way to model innovation
- Innovative songs will have **low likelihood** according to a model fitted only to earlier songs
- Low-likelihood songs:
  Described as innovative, experimental or unusual in the literature

Musical innovation vs. musical influence

- No monotonic correlation between the two (Spearman r=-0.019, p > 0.05)
- More complex relations seem to exist

Musical innovation vs. musical influence

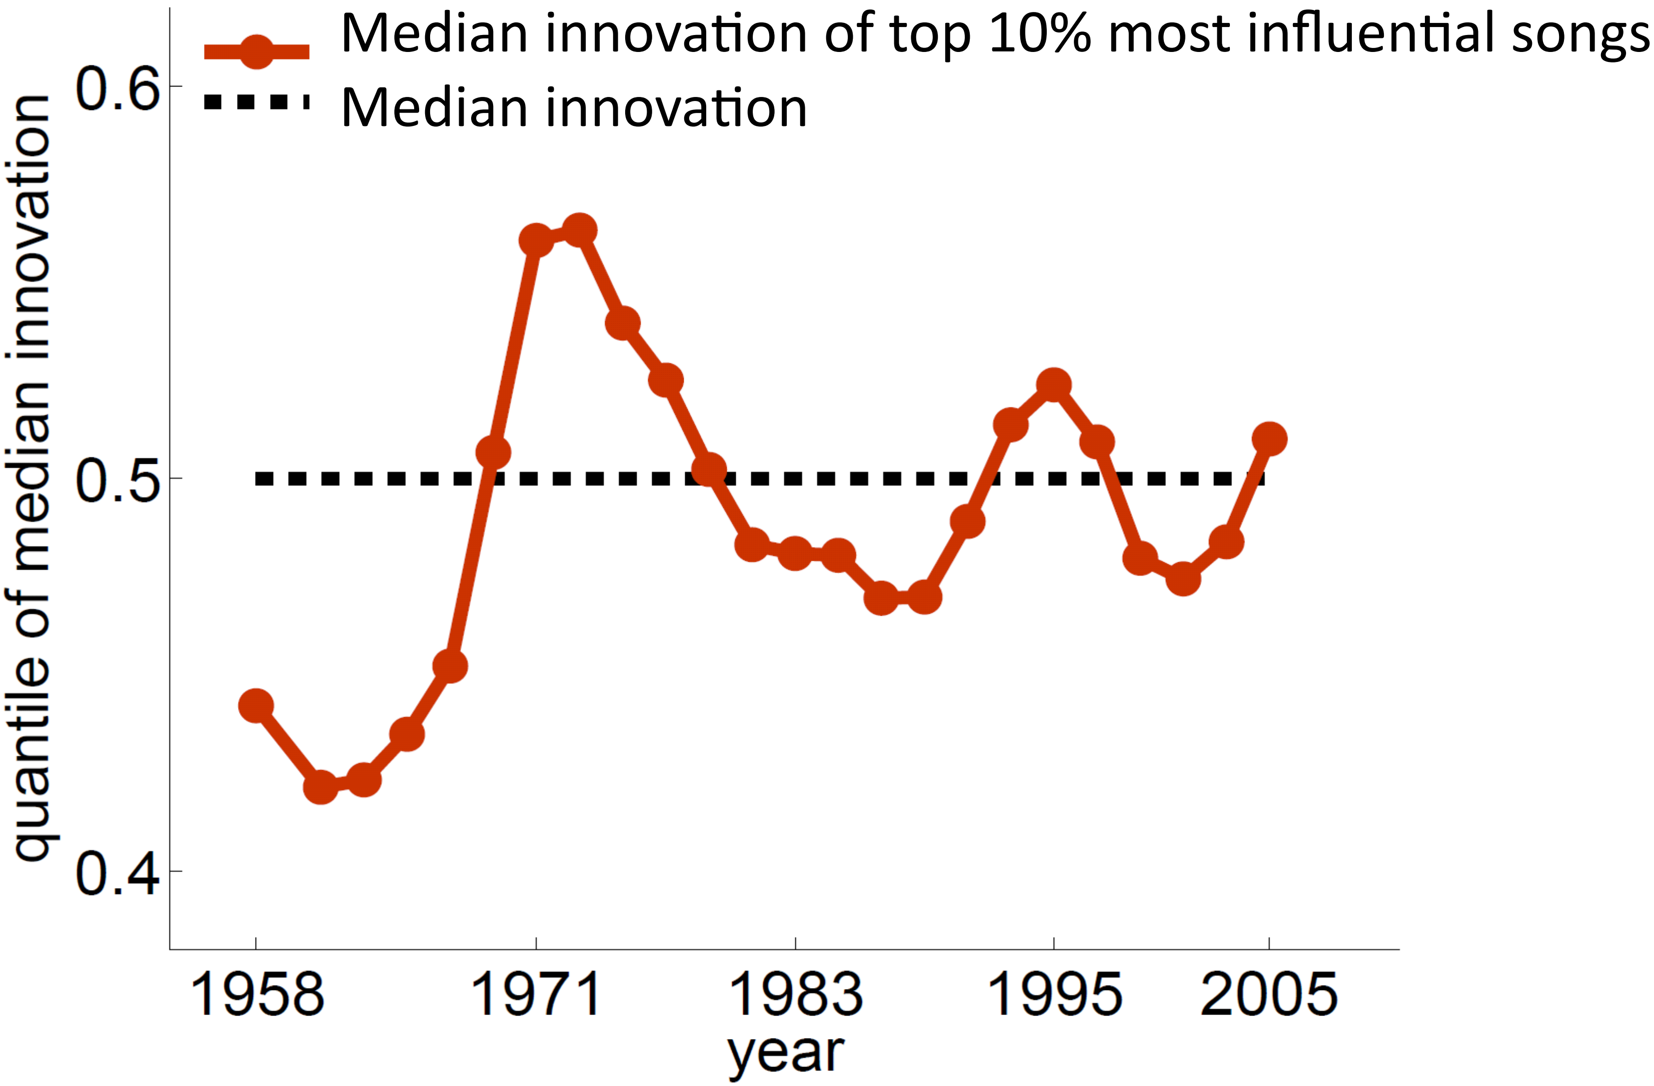

Musical innovation vs. musical influence

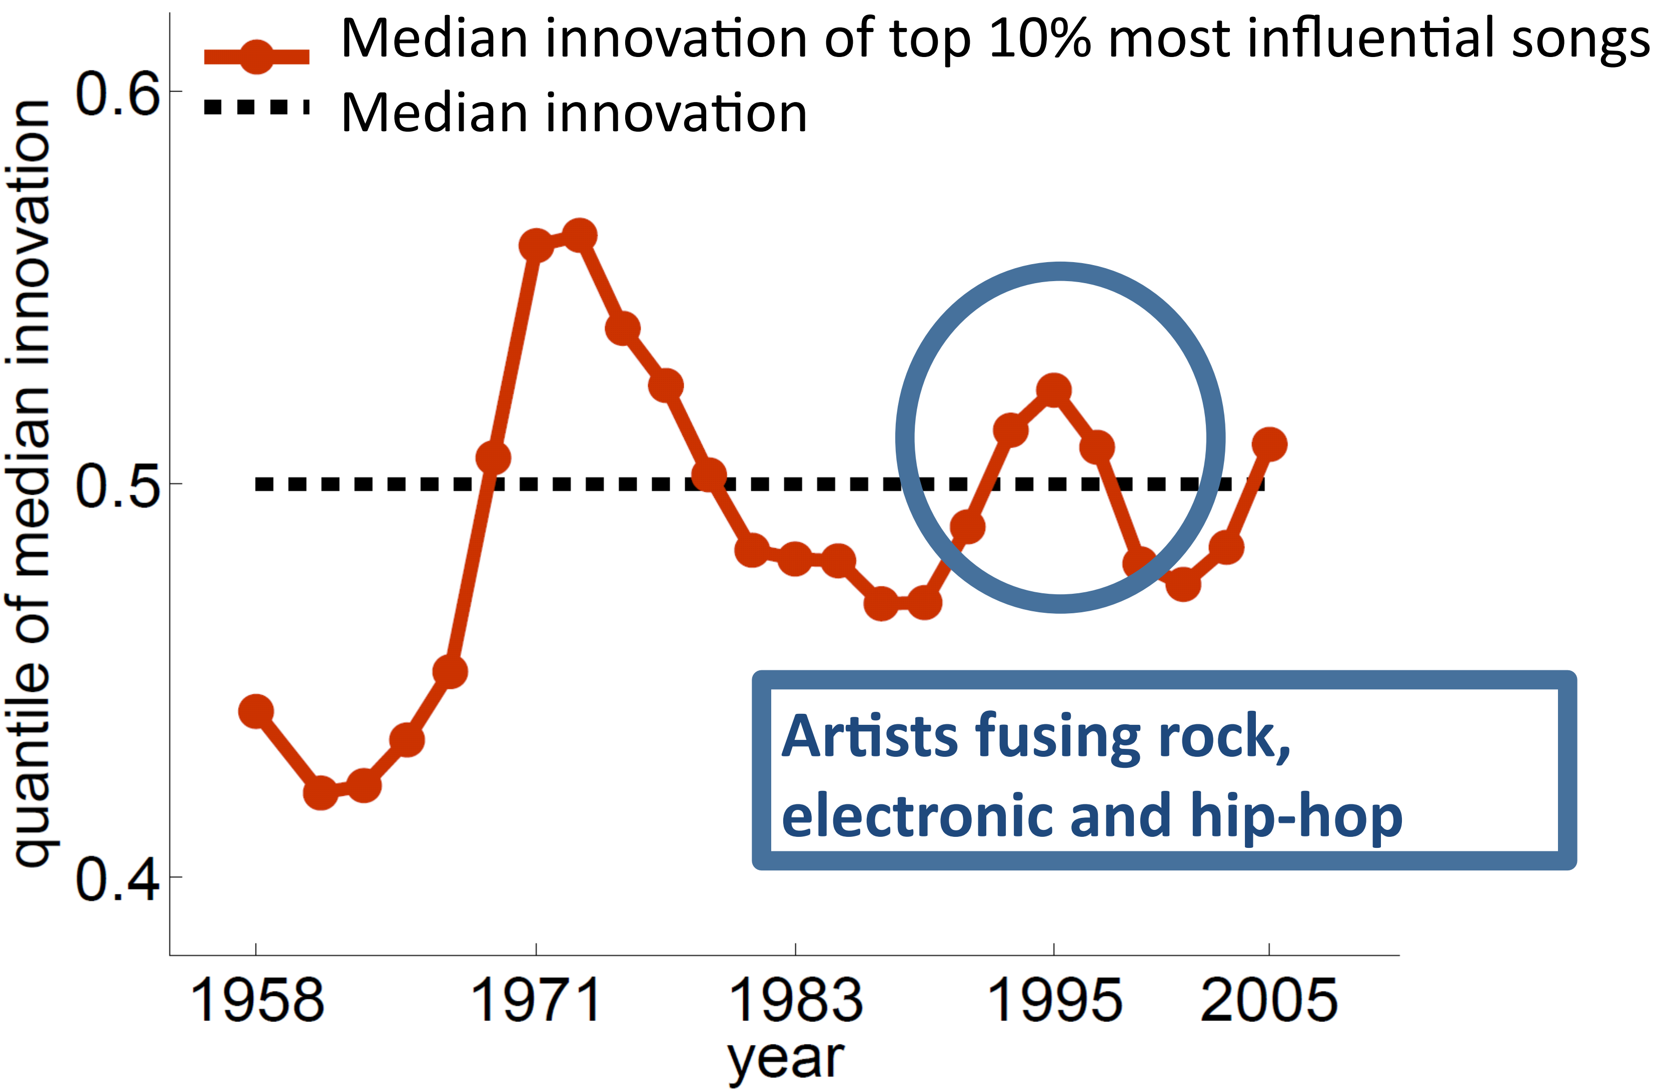

Musical innovation vs. musical influence

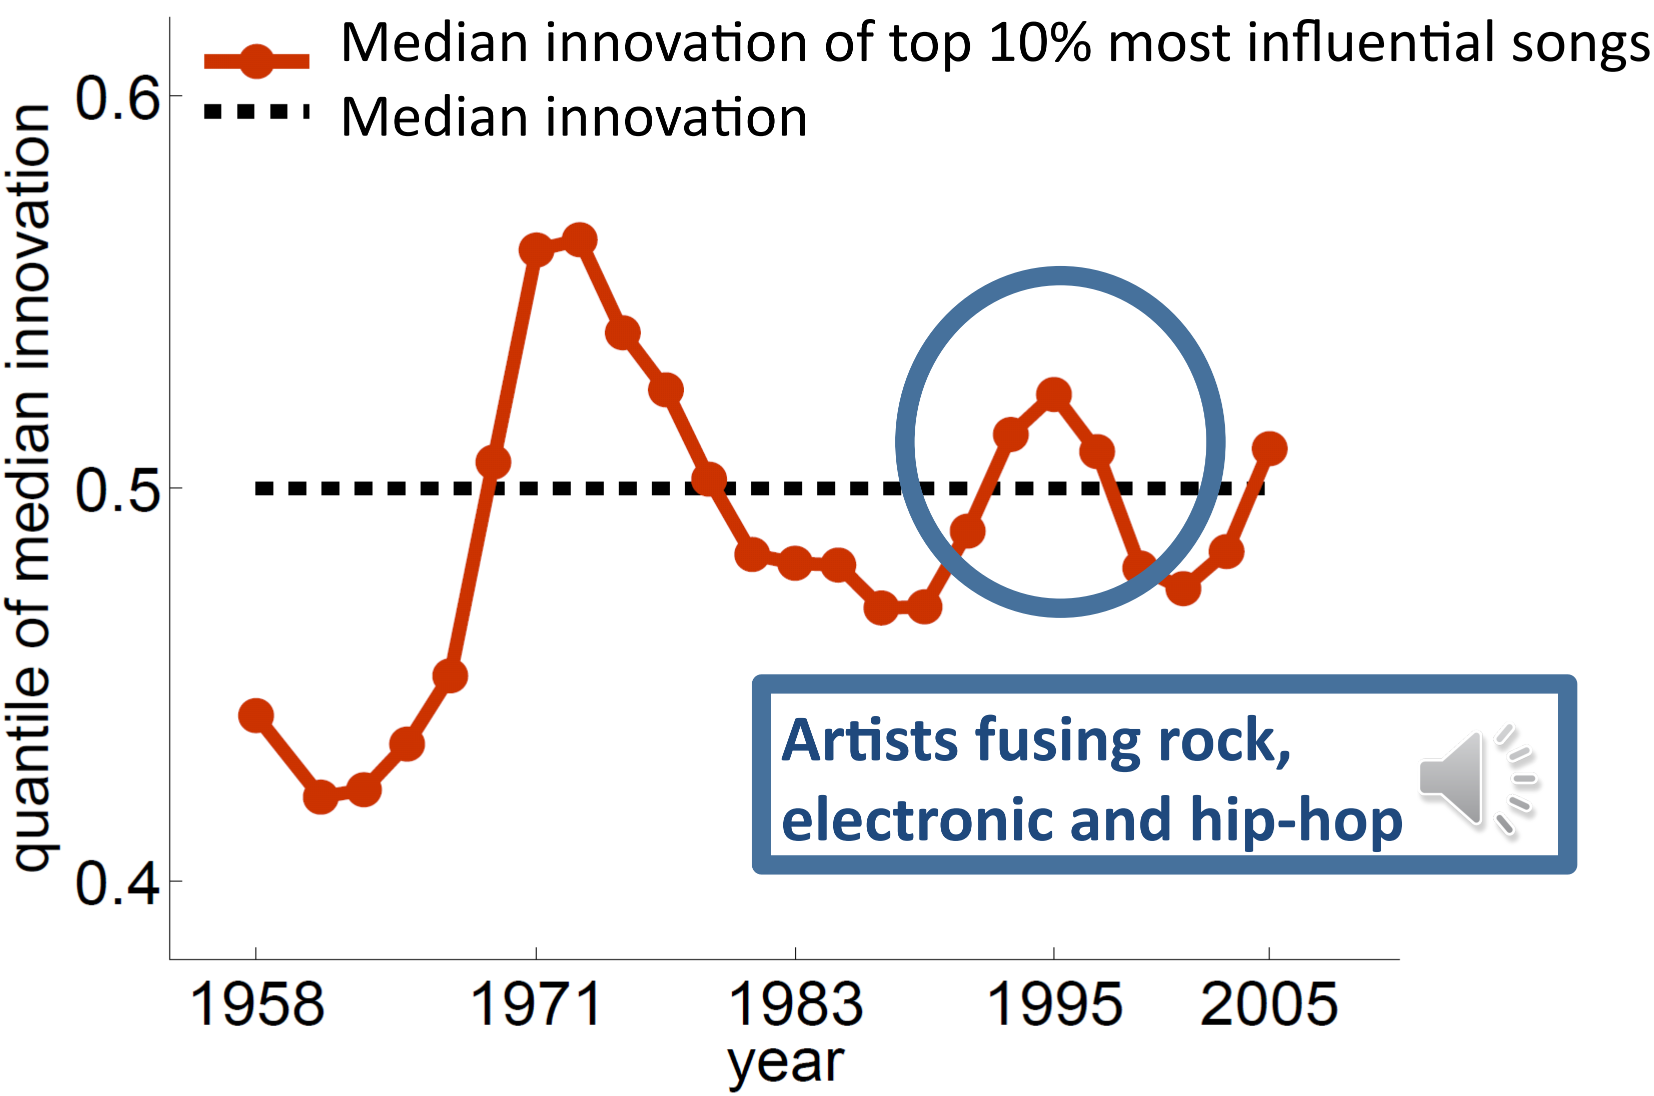

Conclusion

- First large-scale quantitative model of artistic influence
- Validated by human-curated influence measures
- Intriguing connections between innovation and influence