

Statistical models of music-listening sessions in social media

Presenter: Elena Zheleva

Joint work with

John Guiver, Eduarda Mendes Rodrigues, Natasa Milic-Frayling

WWW, April 29, 2010



Microsoft
Research

U. PORTO

Online social media

facebook

flickr*

You Tube

last.fm

twitter


BibSonomy ::

dogster
FOR THE LOVE OF DOGS

orkut^{beta}

bebo

Linked in®



myspace.com
a place for friends..



Online **music** communities



Zune Social



← Search

collection social
music videos pictures podcasts

artists genres albums SO

Song	Song artist	Album	Genre
Amanda	Aisha Duo	Quiet Songs	Jazz
Amazing Grace	Various Artists	The Bagpipes & Drums of	National Mus
Amen UK - Passion (Herbst In	Paul Van Dyk	Perspective (A Collection	Electronic
Amon Hen	Howard Shore	The Lord of the Rings: Th	Soundtrack
Amor de Loca Juventud	Buena Vista Social Club	Buena Vista Social Club	Latin
Amy Hit the Atmosphere	Counting Crows	This Desert Life	Rock
The Ananda Project - Cascades	Various Artists	Clubbed Out - Chilled Be	General Hous
Anastasia (Original)	Anastasia	Anastasia (Single)	Electronic



Buena Vista Social Club

0:09 -3:14

zune products music video social support

search MossyRoc intro (2) sign out

my social my profile

WELCOME MOSSYROC!

my friends MOST RECENT

- Cesarisok (21,715 plays)

my friend's recently played songs

- Get This Money by B-Legit
- You Give Good Love by Whitney Houston
- Sweetheart by Mariah Carey
- Never Too Late by Three Days Grace
- Shadow Of The Day by Linkin Park

my friend's most played albums

- Apologize by Timbaland
- Bleed It Out by Linkin Park
- Low by Flo Rida
- Night (Geneues Remix) by Benga & Coki
- Crimewave by Crystal Castles

my friend's most played artists

what's new

- LouMarvin has left a comment on [BocDinger's page](#) 1/29/2008 10:04AM
- Capn Salty has become a **Bronz Power Listener** for [Radio Head](#) 1/29/2008 10:04AM
- Chadro has added [YoYoMama](#) as a friend 1/29/2008 10:04AM
- Doctorkds has become a **Gold Power Listener** for [Never Too Late](#) by [Three Days Grace](#) 1/29/2008 10:04AM
- Shemp Dizzle has left a comment on [Chadro's page](#) 1/29/2008 10:04AM
- EnglishEmperor added [With or Without You](#) to his favorites 1/29/2008 10:04AM
- StanleyVinegar has sent [With or Without You](#) to Doctorkds 1/29/2008 10:04AM
- mamasan212 has changed her **Zune Background** 1/29/2008 10:04AM

Zune Social

▶ Song-listening instances

The screenshot displays the Zune Social application interface. At the top, there is a search bar, navigation links for 'collection social', 'music videos pictures podcasts', and 'artists genres albums songs playlists'. A 'Sign in' button is also visible. The main content area shows a list of songs with columns for Song, Song artist, Album, and Genre. The song 'Amor de Loca Juventud' by Buena Vista Social Club is highlighted. The playback controls at the bottom show the song is currently playing at 0:09 of a 3:14 duration. Two red circles highlight the 'Genre' column and the heart icon in the right-hand column of the song list.

Song	Song artist	Album	Genre
Amanda	Aisha Duo	Quiet Songs	Jazz
Amazing Grace	Various Artists	The Bagpipes & Drums of	National Music
Amen UK - Passion (Herbst In	Paul Van Dyk	Perspective (A Collection	Electronic
Amon Hen	Howard Shore	The Lord of the Rings: Th	Soundtrack
Amor de Loca Juventud	Buena Vista Social Club	Buena Vista Social Club	Latin
Amy Hit the Atmosphere	Counting Crows	This Desert Life	Rock
The Ananda Project - Cascades	Various Artists	Clubbed Out - Chilled Be	General House
Anastasia (Original)	Anastasia	Anastasia (Single)	Electronic

0:09 3:14

VOLUME 22

Genre taxonomy – 2-level

Rock

Hip Hop

R&B / Soul

Pop

Electronic / Dance

Latin

Reggae / Dancehall

World

Country

Classical

Jazz

Blues / Folk

Comedy / Spoken

Word

Christian / Gospel

Soundtracks

Kids'

More

Classic Rock

Indie / Modern Rock

Metal

New Wave

Punk / Ska

Rock & Roll

Baroque

Chamber

Choral

Classical

Classical Guitar

Crossover

Early

Opera

Operettas

Religious

Renaissance

Romantic



Motivation

- ▶ Variety of ways of **organizing media content** and **summarizing user preferences**
 - ▶ Genre taxonomy is one way
 - ▶ Does not necessarily **reflect user song-listening patterns**
- ▶ Goal: Cluster media according to the listening patterns of users
 - ▶ **Approximate the generative mechanisms** which give rise to the listening patterns
- ▶ Propose: two models for organizing media
 - ▶ *Taste model* – capture **latent taste**, aggregates patterns across users
 - ▶ *Session model* – capture **latent mood**, patterns across users and sessions





Applications

- ▶ **Playlist completion**
 - ▶ create an automatic playlist based on a few seed songs given by the user – e.g. “I am in the mood for songs like this.”
- ▶ **Offer deals on songs**
 - ▶ Target people with a certain taste
- ▶ **Recommendation systems**



Models for clustering media

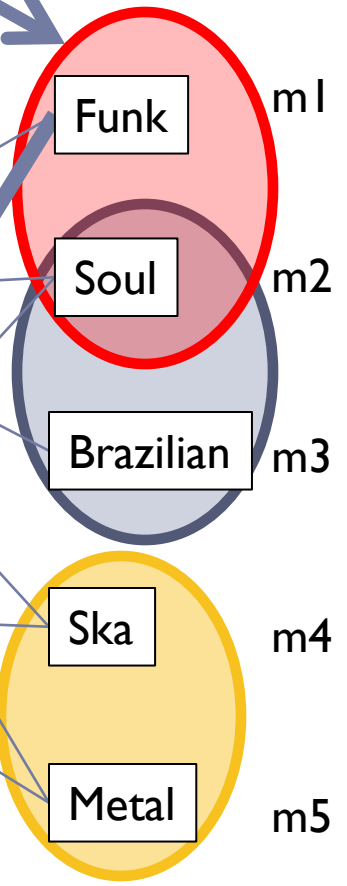
Media clusters

► Collaborative filtering

User	Media item	Timestamp
u1	m1	07-11-09 08:12:06
u1	m2	07-11-09 08:15:42
u2	m5	07-11-09 08:15:42
u1	m3	07-11-09 08:15:42
u2	m4	07-11-09 08:15:42
u2	m1	07-11-09 08:15:42
u2	m1	07-11-09 08:15:42
u2	m2	07-11-09 08:15:42
u2	m1	07-11-09 08:15:42
u1	m4	07-15-09 07:56:43
u1	m5	07-15-09 07:56:43

What are the clusters of media items that *emerge* from the user listening patterns?

Media clusters



Genre – observed descriptor for each media item



Topic models and social media

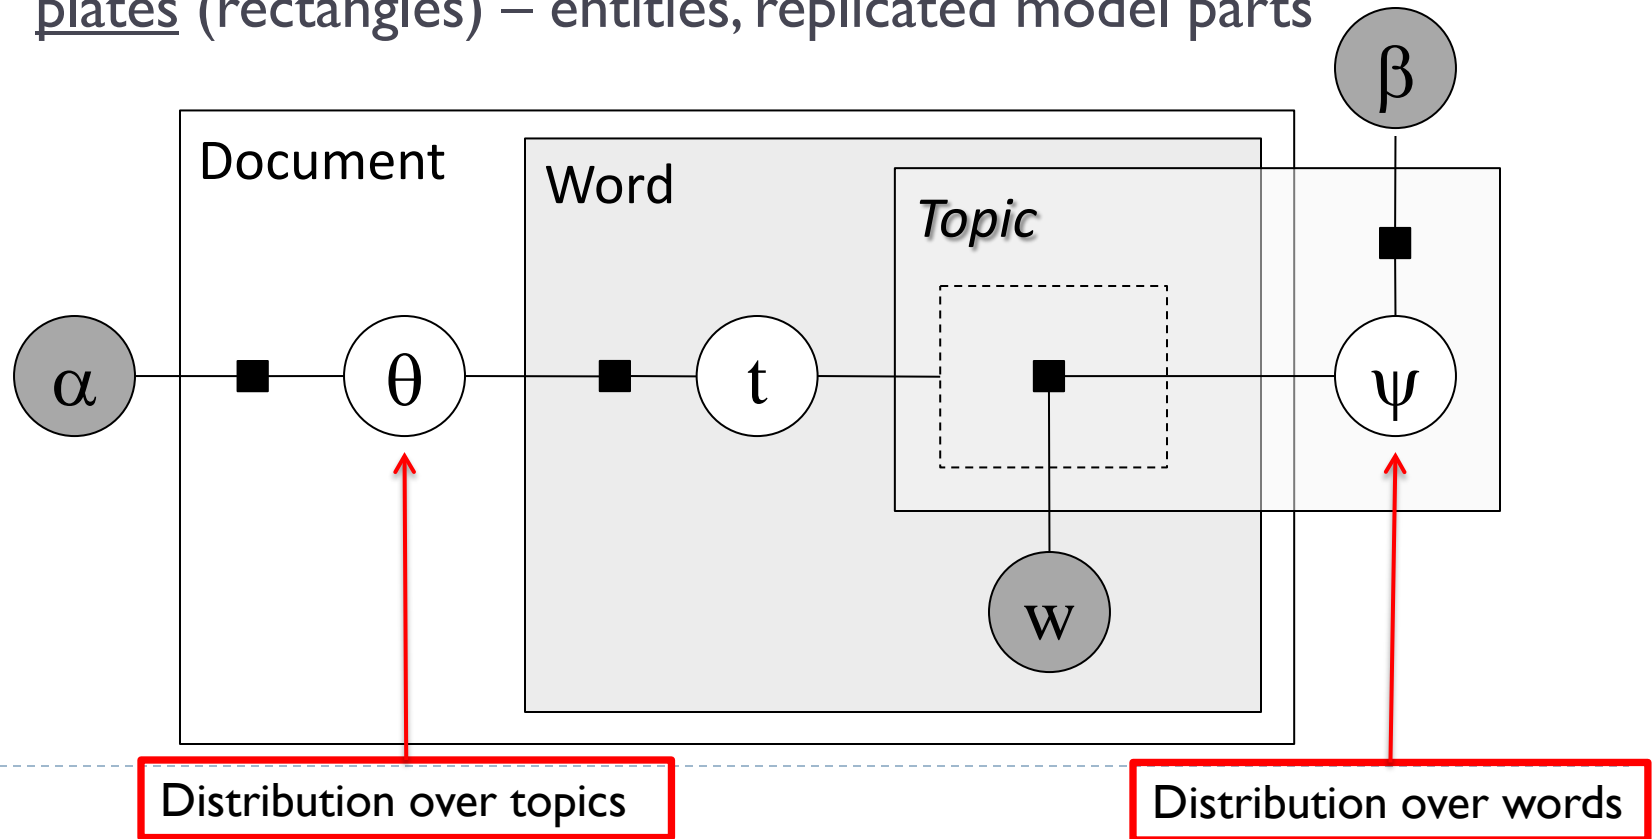
- ▶ We are interested in *latent media clusters*
 - ▶ Derived from aggregate patterns of media listening
 - ▶ Soft clusters – distributions over media items
- ▶ LDA model (Blei, Ng, Jordan, JMLR 2003)
 - ▶ Hierarchical probabilistic model
 - ▶ Discovers *latent topics* in documents using word co-occurrences
 - ▶ Learn *distribution over words* for each topic
 - ▶ Learn *distribution over topics* for each document



LDA model

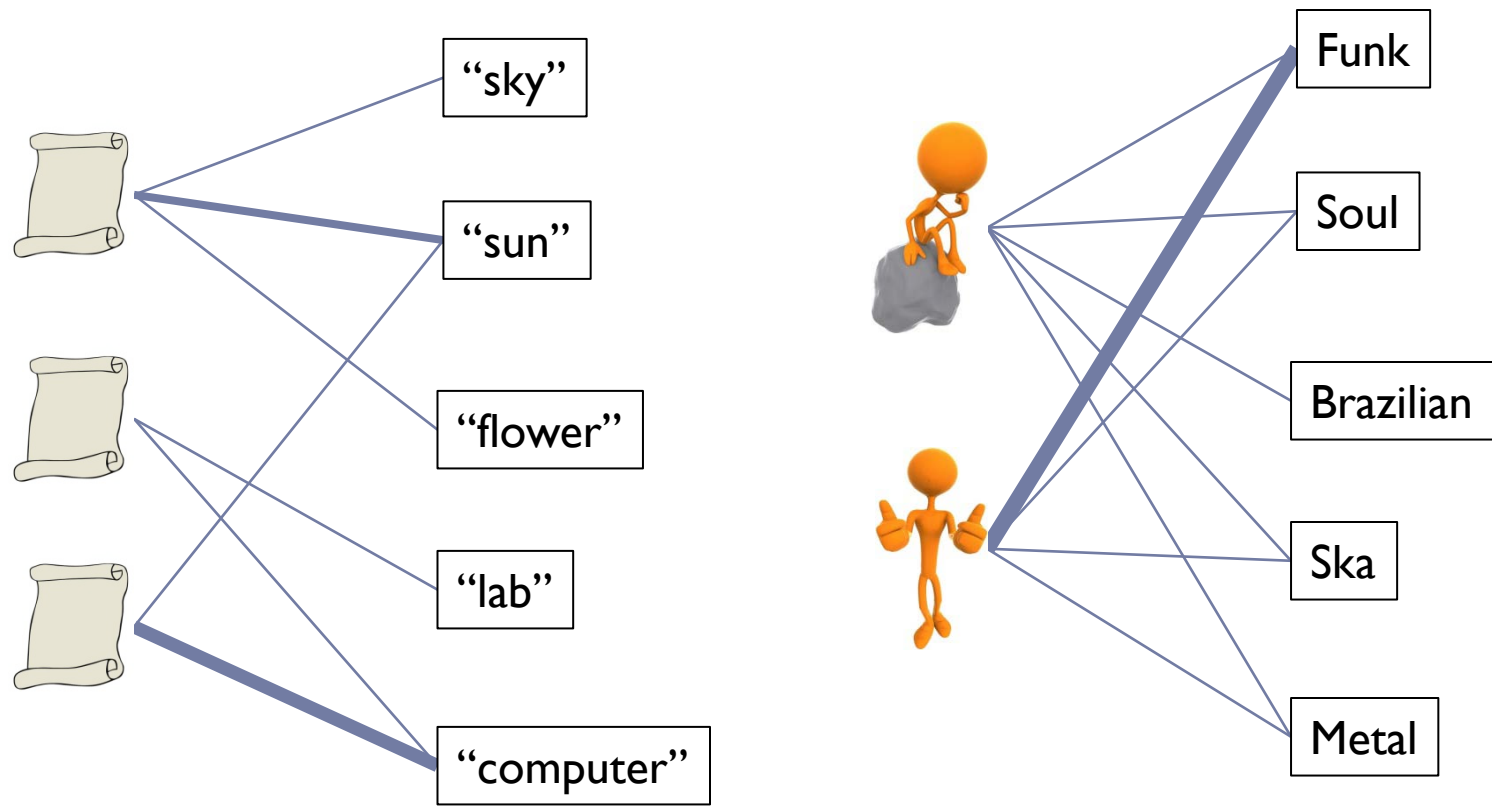
▶ Factor graph representation

- ▶ random variables (circles) – observed and unobserved
- ▶ factors (squares) – functions of random variables
- ▶ plates (rectangles) – entities, replicated model parts



Topic models and social media

- ▶ Documents and words vs. Users and media items
- ▶ Discover latent clusters



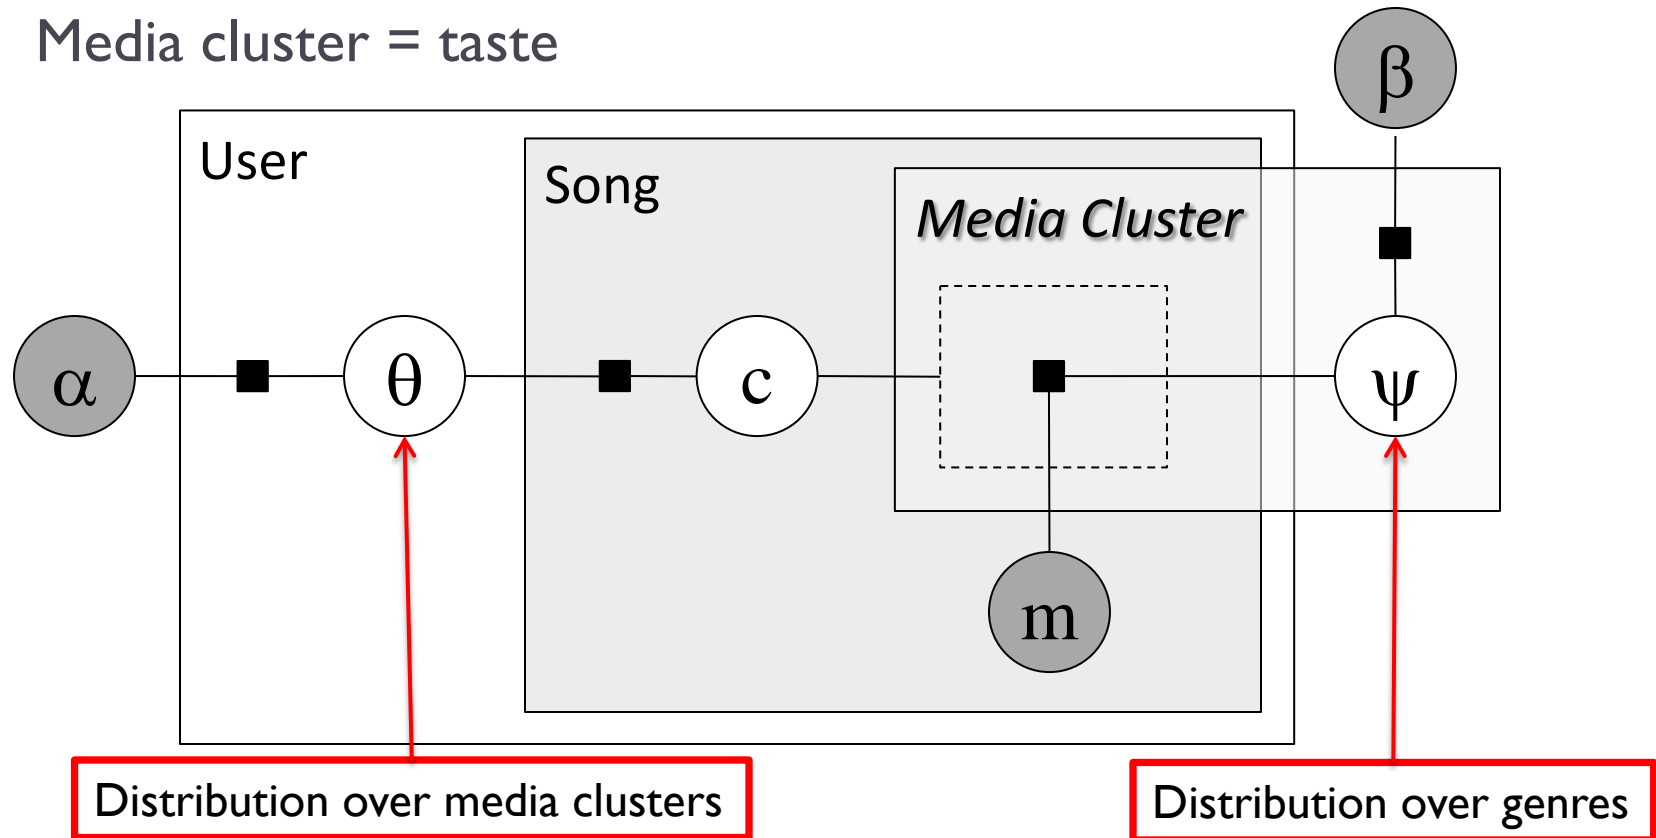
Topic models and social media

- ▶ A number of ways in which LDA can be adapted and applied
 - ▶ We chose two
 - ▶ Taste model
 - ▶ Session model
- ▶ Cluster the media items and learn
 - ▶ Distribution of media items (genre) for each media cluster
 - ▶ Distribution of media clusters (taste, mood) for each user



Taste model

- ▶ Direct application of LDA model
 - ▶ Media cluster = taste

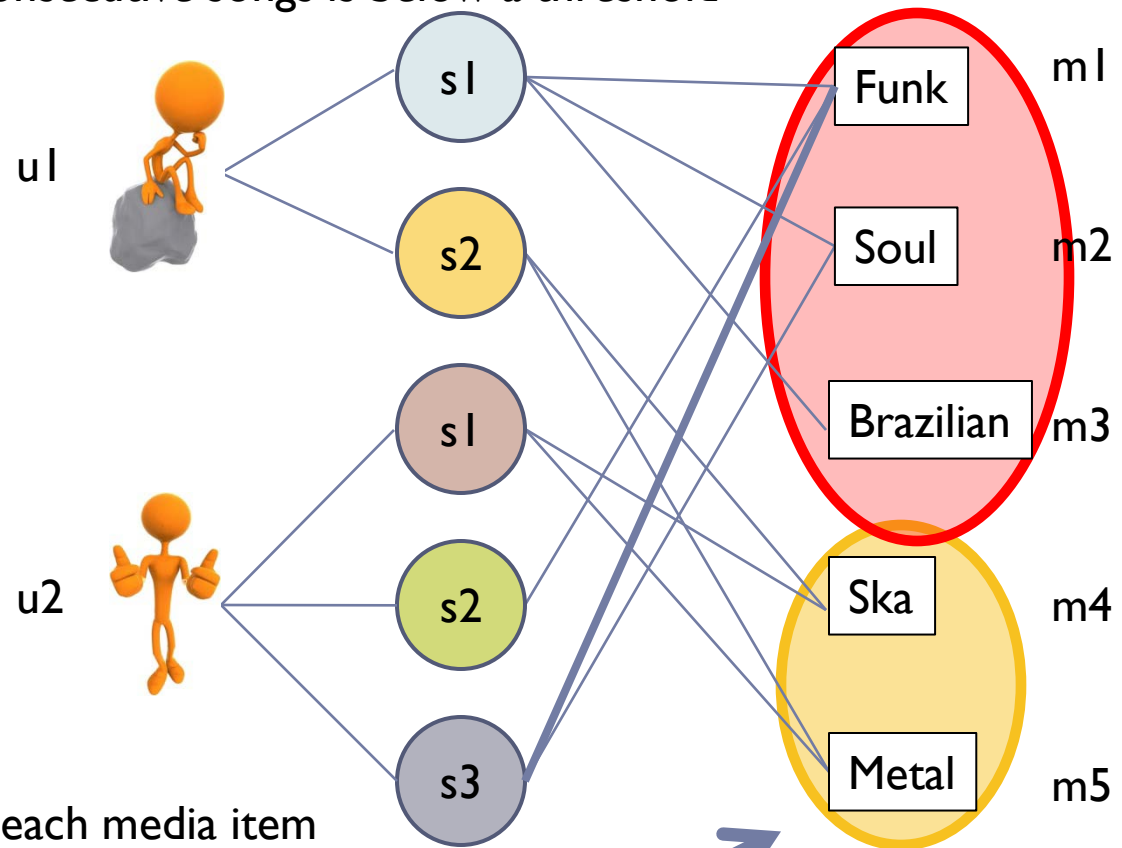


$$P(\theta, \psi, c, m; \alpha, \beta) = \prod_{u \in U} P(\theta_u; \alpha) \prod_{m \in M_u} P(c_{u,m} | \theta_u) P(m | \psi_c) \prod_{c \in C} P(\psi_c; \beta)$$

Media clusters

- ▶ **Session** - a contiguous sequence of songs listened to by a user
 - ▶ Time interval between consecutive songs is below a threshold

User	Media item	Timestamp
u1	m1	07-11-09 08:12:06
u1	m2	07-11-09 08:15:42
u2	m5	07-11-09 08:16:41
u1	m3	07-11-09 08:20:17
u2	m4	07-11-09 08:20:33
u2	m1	07-11-09 18:56:43
u2	m1	07-12-09 23:54:45
u2	m2	07-12-09 23:58:21
u2	m1	07-13-09 00:01:12
u1	m4	07-15-09 07:54:04
u1	m5	07-15-09 07:56:43

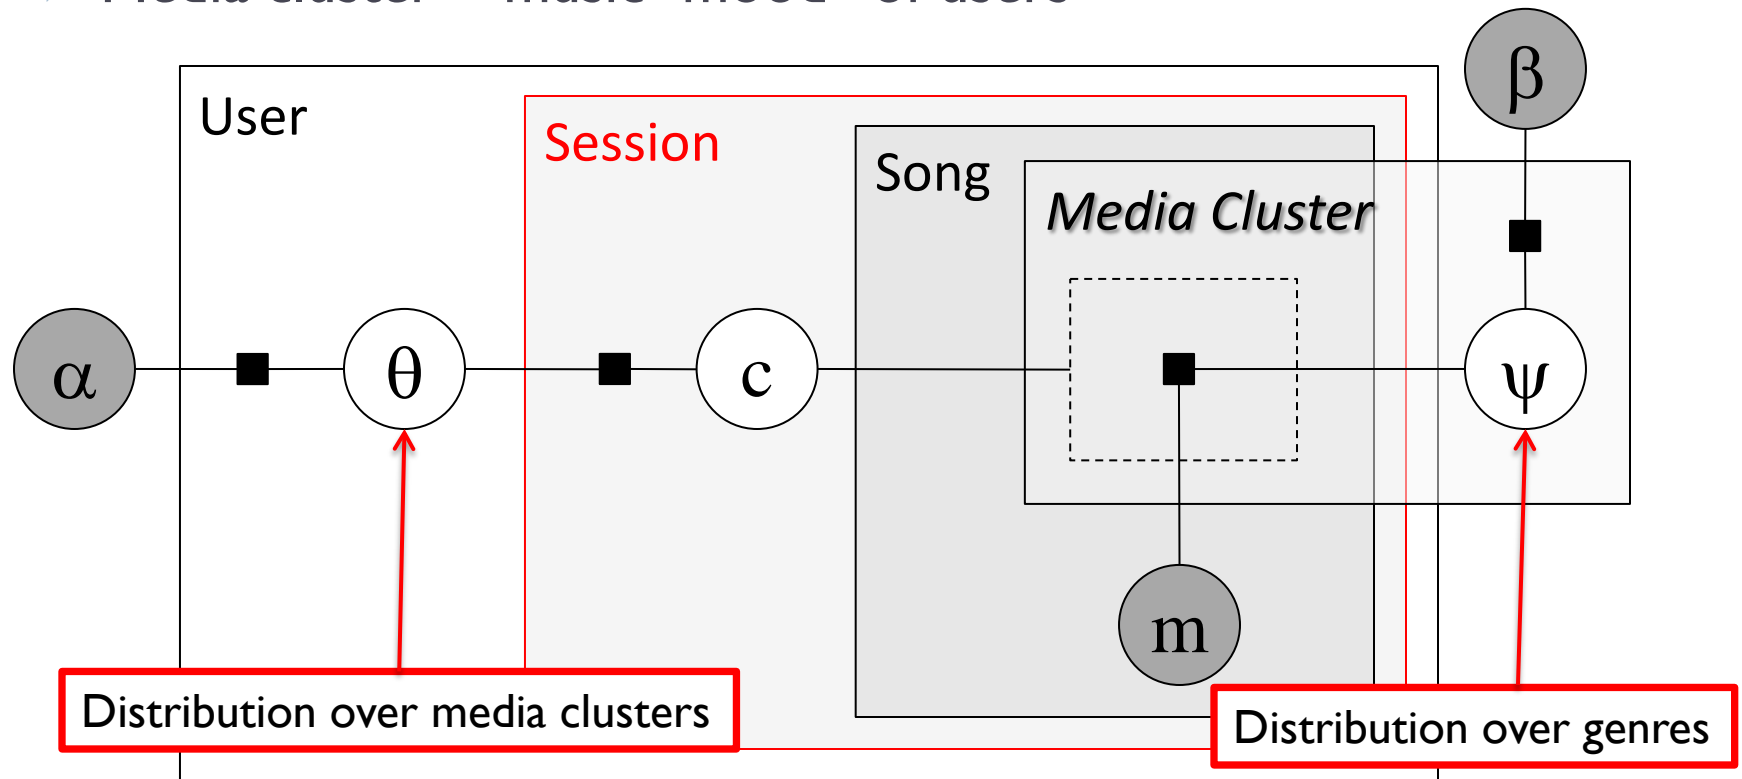


Genre – observed descriptor for each media item

How to cluster media items?

Session model

- ▶ Extend taste model to use temporal session data
 - ▶ Media cluster = music “mood” of users



$$P(\theta, \psi, c, m; \alpha, \beta) = \prod_{u \in U} P(\theta_u; \alpha) \prod_{s \in S_u} P(c_{u,s} | \theta_u) \prod_{m \in M_{u,s}} P(m | \psi(c_{u,s})) \prod_{c \in C} P(\psi_c; \beta)$$



Evaluation



Zune Social data set

- ▶ **14 weeks of data**
- ▶ **2,014 users**
 - ▶ New users – complete history
- ▶ **478,056 genre instances**
 - ▶ Only of positively rated songs
 - ▶ 84 possible genres (sub-genres in the taxonomy)
- ▶ **19,782 sessions**
 - ▶ Gap between sessions – chosen from peak in distribution (30 mins)



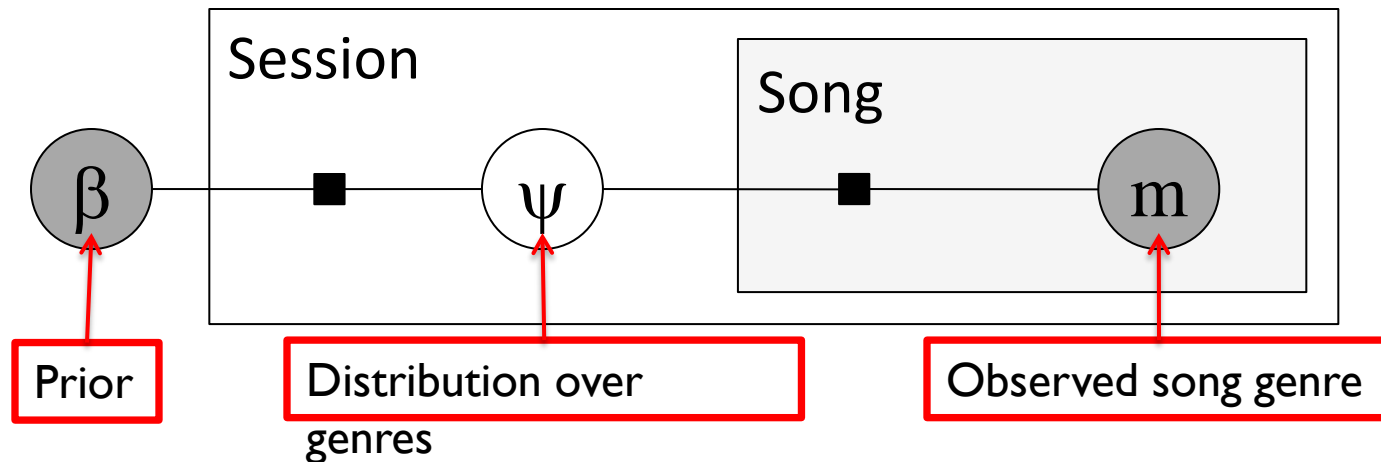
Experimental setup

- ▶ Implemented in Infer.NET
 - ▶ Bayesian inference using variational message passing
- ▶ Train models on 2 months of data
 - ▶ Learn posterior distributions (for ψ and θ)
- ▶ For every test session in the rest of the data
 - ▶ Use learned posteriors as priors
 - ▶ Learn distribution over clusters using 5 or 10 seed songs
 - ▶ Report on *perplexity score* of the other songs in the session
 - ▶ Standard entropy-based measure for topic models
 - ▶ Lower perplexity -> better predictive power



Experimental setup

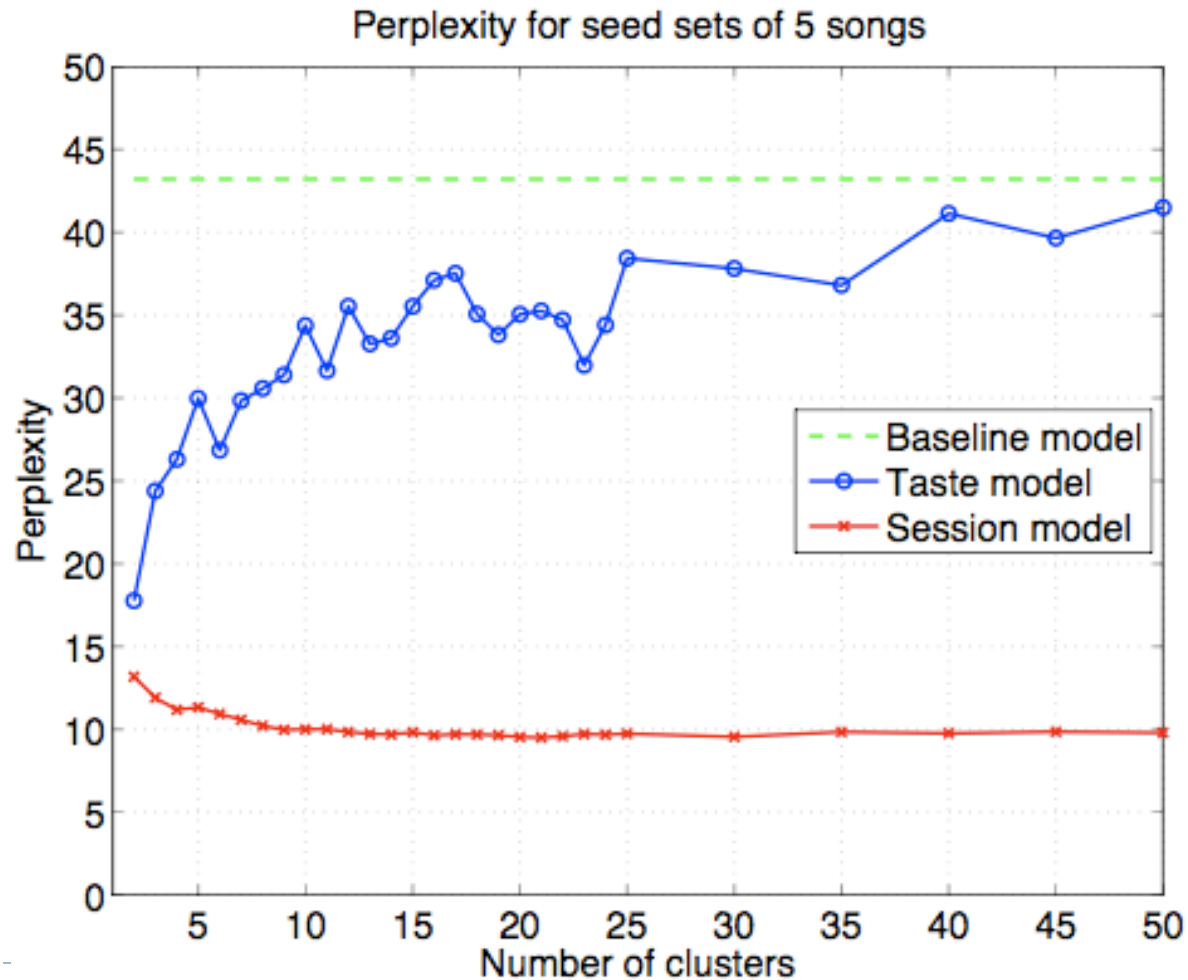
- ▶ Test 3 models:
 - ▶ **taste model**
 - ▶ **session model**
 - ▶ simple **baseline model**
 - ▶ Assumes that there are *no* latent media clusters
 - ▶ Each session defines its own media cluster



$$P(\psi, m; \beta) = \prod_{s \in S} P(\psi_s; \beta) \prod_{m \in M_s} P(m | \psi_s)$$


Predictive power

- ▶ Lower perplexity -> better model

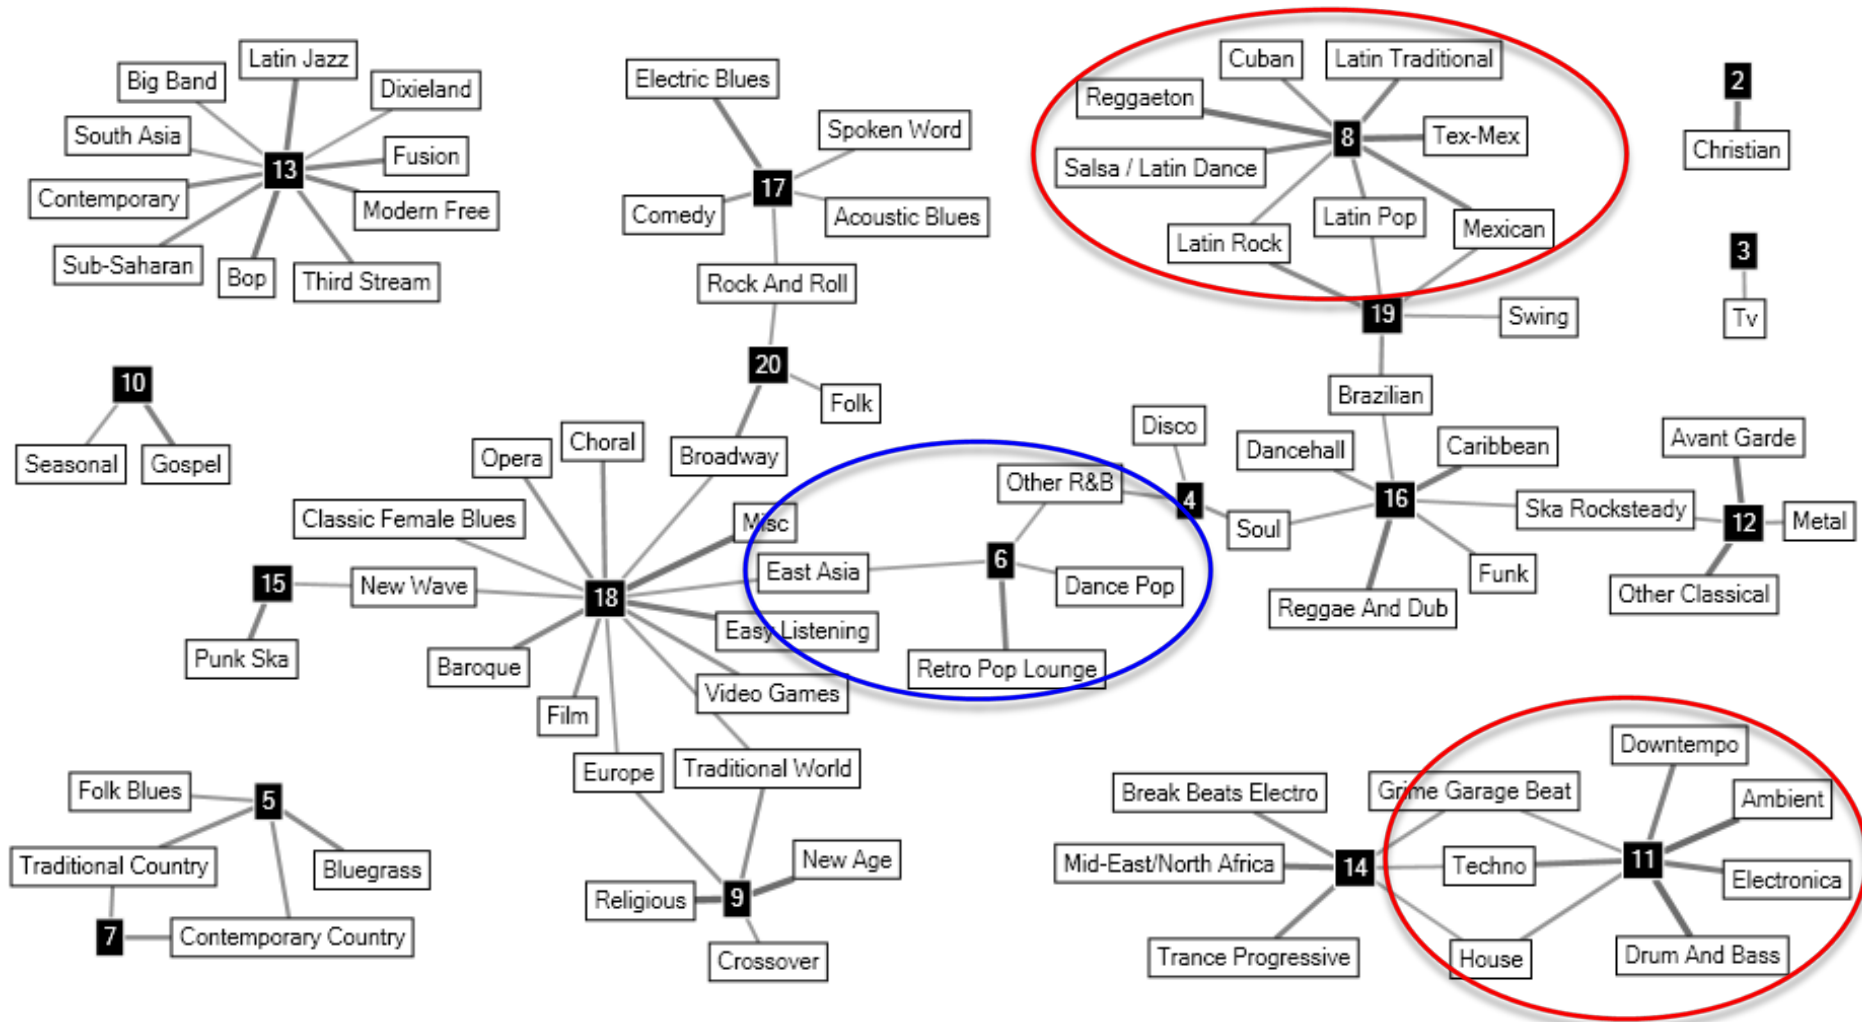


Predictive power

- ▶ Lower perplexity -> better model




Discovered clusters – session model



Training time

- ▶ Using variational message passing





Conclusions

- ▶ There are many ways to adapt *topic models* to social media
 - ▶ We explored two models for an online music community
 - ▶ To capture the latent **music taste** and **mood** of a user
- ▶ Having a richer, mood model (with **sessions**) paid off
 - ▶ Better predictive power for playlist generation
 - ▶ Faster training time
 - ▶ Useful in this setting, expect it to be in others





Applications and future work

- ▶ **Playlist completion** based on 'mood'
 - ▶ dynamically adapting recommendations at session level
- ▶ **Personalization of recommendations**
 - ▶ choosing the right level – individual vs. community
- ▶ **New radios**
 - ▶ serving the emerging media clusters rather than genre separation
- ▶ **Interpretations: macro level -> micro level**
 - ▶ distributions to **human experience**



Statistical models of music-listening sessions in social media

QUESTIONS?

Elena Zheleva
elena@cs.umd.edu