Understanding Engagement and Sentiment in MOOCs using Probabilistic Soft Logic (PSL)

Lise Getoor

UC Santa Cruz

@lgetoor

    

Arti Ramesh
UMD
=> Suny Binghamton

Dan Goldwasser
UMD
=> Purdue

Bert Huang
UMD
=> VTech

Hal Daume III
UMD

Sabina Tomkins
UC Santa Cruz

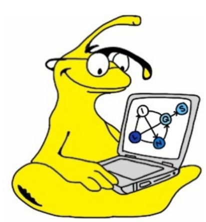

Machine Learning for Education NIPS Workshop
December 10, 2016

MOOC Domains

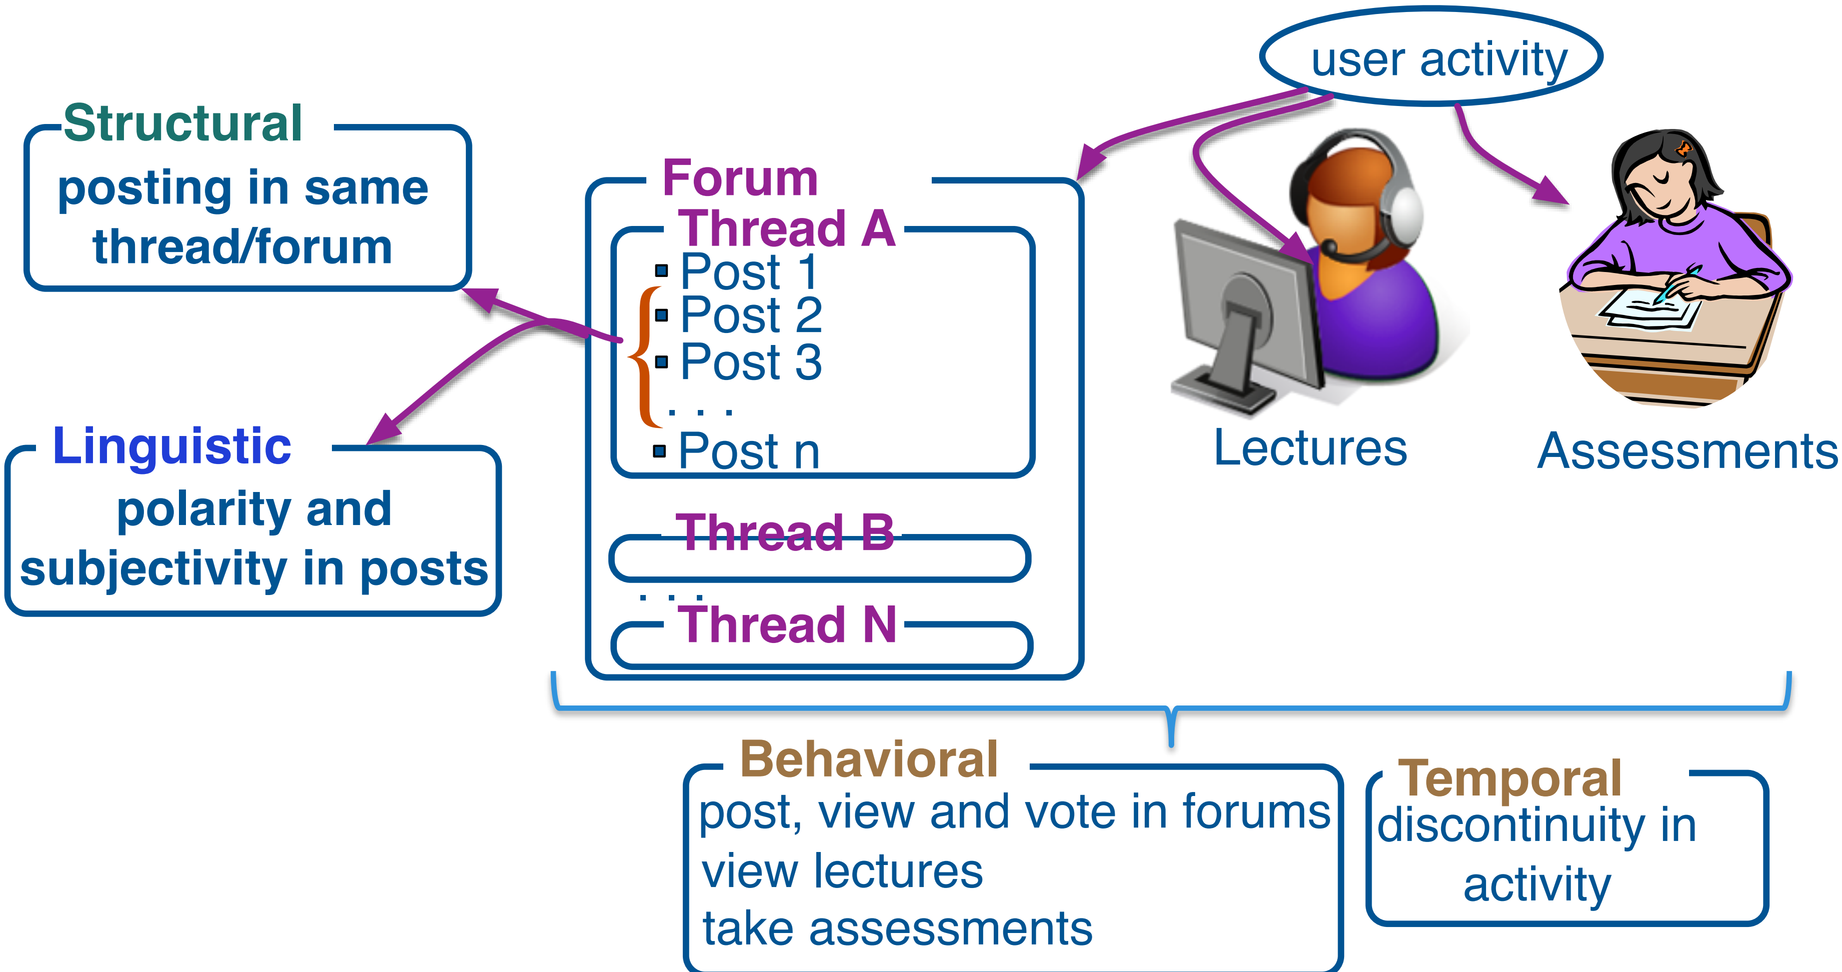

Characterized by:

- Rich Socio-Behavioral Data
- Rich Outcomes Space
- Opportunity (and Need for!!) Interventions

Overview

- Data
  - 14 different UMD courses spanning varied disciplines:
    - Business, Sciences, History, Computer Science, Philosophy, and Religion
    - Includes most popular courses, Innovative Ideas and Android, which have run successfully for several iterations and attract a large number of participants each offering
  - Courses have on average 100,000 students, 10,000 posts
  - On average, # latent + target variables ≈ 400,000
  - Largest course has students ~230,000 and 50,000 posts
  - Temporal data for 34 repeated offerings of Business course and 15 offerings of CS course
- Latent Variable Models
  - Engagement, sentiment, topics, fine-grained course aspects
  - Validated on outcome, completion and aspect prediction
  - Used for exploratory and descriptive analysis

References: Ramesh, PhD Thesis 2016, Ramesh et al., ACL 2015, Ramesh et al., AAAI 2014, Ramesh et al., L@S 2014, Ramesh et al., BEA 2014.

Modeling Approach

http://psl.umiacs.umd.edu

Probabilistic Soft Logic (PSL)

A probabilistic programming language for collective probabilistic inference problems

- Predicate = relationship or property
- Atom = **(continuous)** random variable
- Rule = capture dependency or **constraint**
- Set = define **aggregates**

PSL Program = Weighted Rules + Input DB

Reference: *Hinge-Loss Markov Random Fields and Probabilistic Soft Logic, Stephen H. Bach, Matthias Broecheler, Bert Huang, Lise Getoor, arXiv 2015*

PSL Foundations

- PSL makes large-scale reasoning scalable by mapping logical rules to convex functions and defines a ***hinge-loss Markov Random field***:

$$P(\mathbf{Y} \mid \mathbf{X}) = \frac{1}{Z} \exp \left[- \sum_{j=1}^{m} w_j \max\{\ell_j(\mathbf{Y}, \mathbf{X}), 0\}^{p_j} \right]$$

- Three principles justify this mapping [Bach et al., AIStats 15]:
  - LP programs for MAX SAT with approximation guarantees [Goemans & Williamson 94]
  - Pseudomarginal LP relaxations of Boolean Markov random fields [Wainwright et al. 02]
  - Łukasiewicz logic, a logic for reasoning about continuous values [Klir & Yuan 95]

PSL Summary in a Slide

Probabilistic Soft Logic

psl.umiacs.umd.edu

- PSL is a probabilistic programming language that supports declarative features, collective reasoning and lifted models
- MAP Inference in PSL translates into convex optimization problem -> **inference is really fast.** Inference further enhanced with **state-of-the-art optimization and distributed processing paradigms** such as ADMM & GraphLab -> **inference even faster**
- **Outperforms discrete MRFs** in speed and often accuracy
- **Learning** methods for **rule weights** & **latent variables**
- Good fit for many **structured prediction** problems in NLP, computer vision, social computing, information integration, knowledge construction, and more
- PSL is **open-source**, code, data, tutorials available online

Application Domains

- Computational Biology & Health Informatics
  - Drug-target prediction
  - Drug interaction prediction
- Computational Social Science
  - Social trust prediction
  - Latent Group Modeling in Twitter
  - Learner engagement in MOOCs
  - Inferring bias in political discourse
  - Psychological modeling on online social networks
- Computer Vision
  - Low-level image reconstruction
  - Activity recognition in Video
- Information Integration & Extraction
  - Entity resolution
  - Knowledge graph identification
  - Ontology alignment & schema mapping

Application Domains

- Computational Biology & Health Informatics
  - Drug-target prediction
  - Drug interaction prediction
- Computational Social Science
  - Social trust prediction
  - Latent Group Modeling in Twitter
  - **Learner engagement in MOOCs**
  - Inferring bias in political discourse
  - Psychological modeling on online social networks
- Computer Vision
  - Low-level image reconstruction
  - Activity recognition in Video
- Information Integration & Extraction
  - Entity resolution
  - Knowledge graph identification
  - Ontology alignment & schema mapping

Latent-Variable Models

Student Engagement

Ramesh et al., AAAI 2014

Student Engagement in MOOCs

Engagement in MOOCs different from classrooms

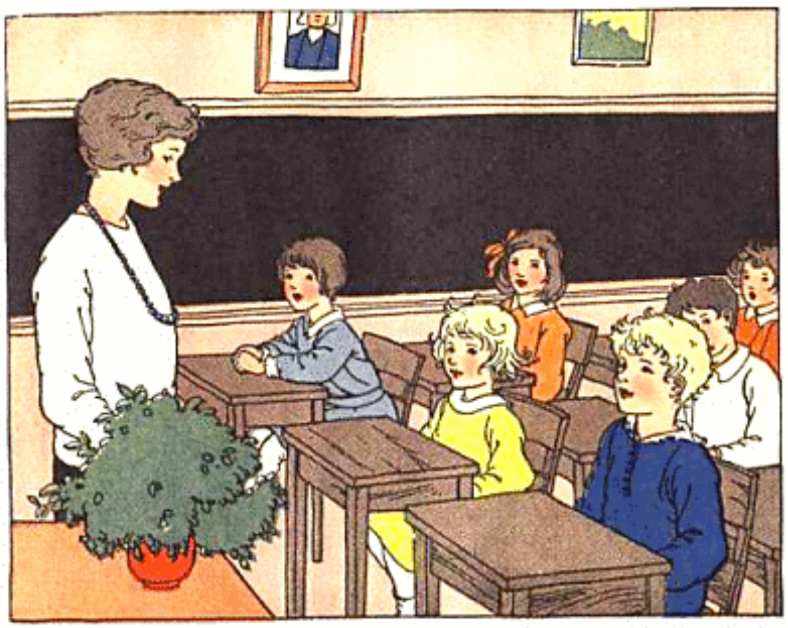

How do students engage with the MOOC?

Problem: Modeling Student Engagement

Large number of registrants

Low completion rate

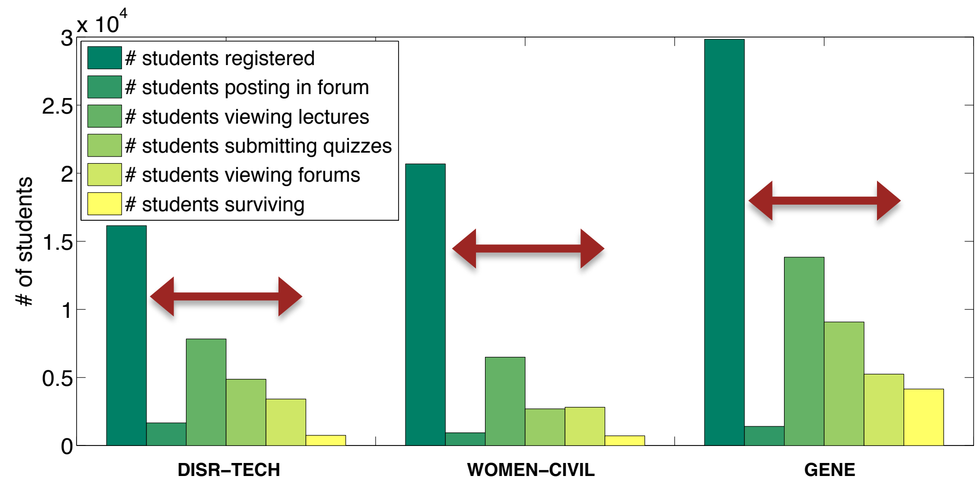

Problem: How to model student engagement in MOOCs?

Model engagement using

- **Online behavior**
- **Linguistic Analysis of forum posts**
- **Structural attributes from forum interactions**
- **Temporal attributes modeling continuous behavior**

Prediction Framework

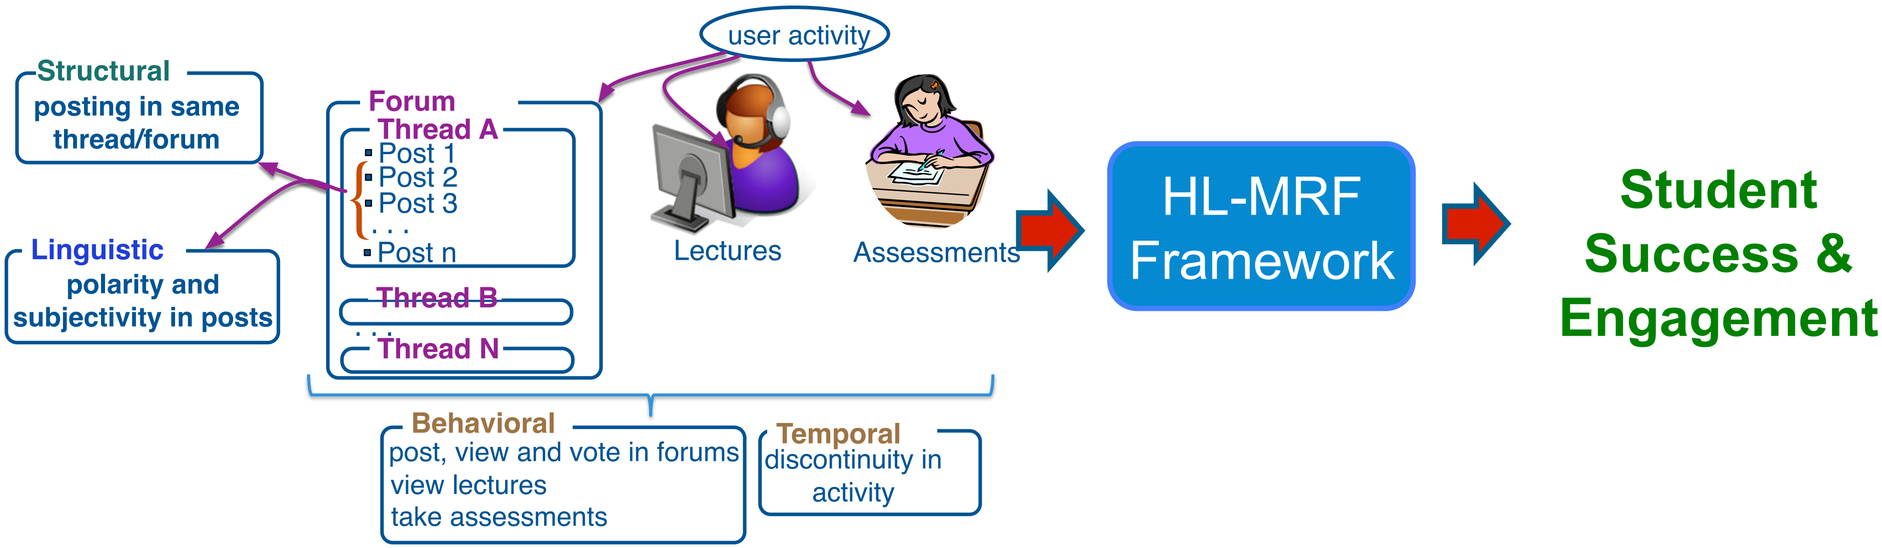

Three forms of engagement: **active engagement, passive engagement, and disengagement**

Two measures of student success: **course performance** and **course completion**

Two models: **direct** and **latent**

HL-MRF Course Success Prediction Models

- **Direct**
  - Dependencies among observed variables predict **course success**

PSL-DIRECT

"John posts positive sentiment posts indicates course completion"

Observed Features — implies — Target Variable

$$\boxed{\text{POSTS}(``John", P) \wedge \text{POSITIVE}(P)} \Rightarrow \boxed{\text{COURSE-COMPLETION}(``John")}$$

"Mary votes on positive sentiment indicates course completion"

Observed Features — implies — Target Variable

$$\boxed{\text{VOTES}(``Mary", P) \wedge \text{POSITIVE}(P)} \Rightarrow \boxed{\text{COURSE-COMPLETION}(``Mary")}$$

PSL Course Success Prediction Models

- **PSL-LATENT**
  - Latent variable capturing engagement types
    - Active Engagement, Passive Engagement, Disengagement
  - Dependencies between **observed behavior** and **course-success** through latent engagement type

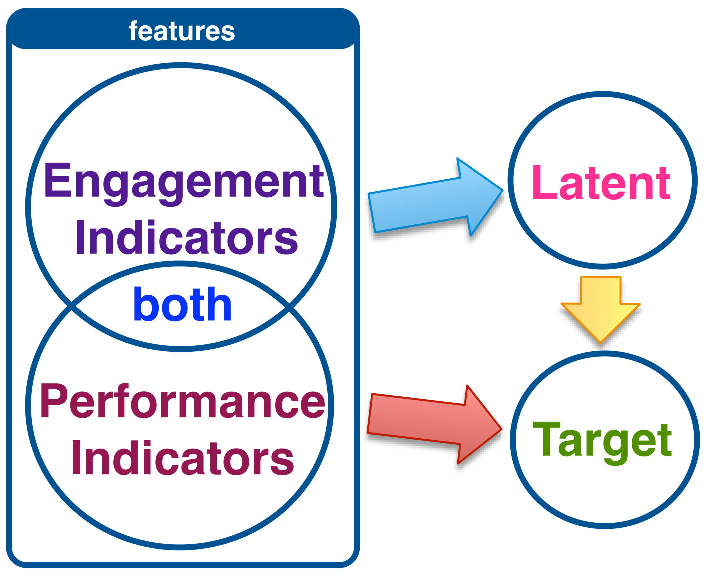

PSL-LATENT

"John posts in forums indicates his active engagement in class"

Observed Features — implies — Latent Engagement Variables

$$\boxed{\text{POSTS}(\text{"}John\text{"}, P) \wedge \text{POSITIVE}(P)} \Rightarrow \boxed{\text{ENGAGEMENT-ACTIVE}(\text{"}John\text{"})}$$

"Mary views lectures, views forum posts indicates her passive engagement in class"

Observed Features — implies

$$\boxed{\text{VIEW-LECTURE}(\text{"}Mary\text{"}, L) \wedge \text{VIEW-FORUM}(\text{"}Mary\text{"}, P)} \Rightarrow \boxed{\text{ENGAGEMENT-PASSIVE}(\text{"}Mary\text{"})}$$

Latent Engagement Variable

Experimental Results: Performance

Course	Model	certificate	~certificate	AUC-ROC
DISRUPTIVE TECHNOLOGIES (BUSINESS)	LECTURE-RANK	0.630	0.421	0.512
	DIRECT	0.739	0.546	0.667
	LATENT	**0.759**	**0.575**	**0.692**
WOMEN & CIVIL RIGHTS (HISTORY)	LECTURE-RANK	0.263	0.761	0.503
	DIRECT	0.794	0.881	0.862
	LATENT	**0.922**	**0.950**	**0.948**
GENE & HUMAN CONDITION (SCIENCE)	LECTURE-RANK	0.503	0.482	0.476
	DIRECT	0.814	0.755	0.817
	LATENT	**0.943**	**0.879**	**0.931**

Performance: predicting whether a student earns a certificate

Area under Curve (AUC) scores for predicting **certificate** and **~certificate**

LATENT PSL Model performs better at predicting **performance**

Experimental Results: Course Completion

Course	Model	course-completion	dropout	AUC-ROC
DISRUPTIVE TECHNOLOGIES (BUSINESS)	LECTURE-RANK	0.333	**0.998**	0.957
	DIRECT	0.393	0.997	0.936
	LATENT	**0.546**	**0.998**	**0.969**
WOMEN & CIVIL RIGHTS (HISTORY)	LECTURE-RANK	0.508	0.995	0.946
	DIRECT	0.565	0.995	0.940
	LATENT	**0.816**	**0.998**	**0.983**
GENE & HUMAN CONDITION (SCIENCE)	LECTURE-RANK	0.688	0.984	0.938
	DIRECT	0.757	**0.985**	0.939
	LATENT	**0.818**	**0.985**	**0.944**

LATENT Model performs better at predicting **completion** and **dropout**

Completion harder to predict due to **high percentage of dropouts**

Early Prediction

Course	Model	start	mid	start-mid
DISRUPTIVE TECHNOLOGIES (BUSINESS)	LECTURE-RANK	0.204	0.280	0.269
	DIRECT	0.304	0.400	0.372
	LATENT	**0.417**	**0.454**	**0.451**
WOMEN & CIVIL RIGHTS (HISTORY)	LECTURE-RANK	0.538	0.518	0.533
	DIRECT	0.593	0.647	0.596
	LATENT	**0.674**	**0.722**	**0.699**
GENE & HUMAN CONDITION (SCIENCE)	LECTURE-RANK	0.552	0.648	0.650
	DIRECT	0.647	0.755	0.692
	LATENT	**0.705**	**0.755**	**0.778**

Early prediction helps identify students engaged with the MOOC

LATENT outperforms **DIRECT** and **LECTURE-RANK** in **early prediction**

Topics & Sentiment

Ramesh et al., 9th ACL Workshop on Innovative Use of NLP for Building Educational Applications (BEA), 2014

Example MOOC Posts

MOOC Post
Our history keeps silence about **violence** in families, **contempt** and **unfair** treatment of women.
I am from New York as well! Really **love** the city!
What I **love** about the lesson is how music influences social change. So many **great** songs can **motivate** change in people and society.
The video is very **jumpy**. I **hated** the experience.
Will everyone get a certificate or only people in the signature track?
Will subtitles be made available for the lectures for this week? I **liked** the transcripts from last week.
When is quiz 4 due?

Sentiment in MOOC Posts

MOOC Post	Sentiment
Our history keeps silence about **violence** in families, **contempt** and **unfair** treatment of women.	Negative
I am from New York as well! Really **love** the city!	Positive
What I **love** about the lesson is how music influences social change. So many **great** songs can **motivate** change in people and society.	Positive
The video is very **jumpy**. I **hated** the experience.	Negative
Will everyone get a certificate or only people in the signature track?	Neutral
Will subtitles be made available for the lectures for this week? I **liked** the transcripts from last week.	Positive
When is quiz 4 due?	Neutral

http://psl.umiacs.umd.edu

Sentiment in MOOC Posts

MOOC Post	Sentiment
Our history keeps silence about **violence** in families, **contempt** and **unfair** treatment of women.	Negative
I am from New York as well! Really **love** the city!	Positive
What I **love** about the lesson is how music influences social change. So many **great** songs can **motivate** change in people and society.	Positive
The video is very **jumpy**. I **hated** the experience.	Negative
Will everyone get a certificate or only people in the signature track?	Neutral
Will subtitles be made available for the lectures for this week? I **liked** the transcripts from last week.	Positive
When is quiz 4 due?	Neutral

Academic

Sentiment in MOOC Posts

MOOC Post	Sentiment
Our history keeps silence about **violence** in families, **contempt** and **unfair** treatment of women.	Negative
I am from New York as well! Really **love** the city!	Positive
What I **love** about the lesson is how music influences social change. So many **great** songs can **motivate** change in people and society.	Positive
The video is very **jumpy**. I **hated** the experience.	Negative
Will everyone get a certificate or only people in the signature track?	Neutral
Will subtitles be made available for the lectures for this week? I **liked** the transcripts from last week.	Positive
When is quiz 4 due?	Neutral

Academic

Logistics

MOOC Forum Categories

- Discussions about course material (**academic**)
- Meta-level discussions about course: logistics and feedback (**logistics**)
- Other general discussions: Introductions, Study Groups (**social**)

Topic Modeling for MOOC Forums

- Unsupervised methods are useful
  - Many classes, less commonality between classes
  - LDA, variants of LDA are good choices
- Sample LDA Topics
  - newspaper, paper, model, business, print, **course, assignments, grade**
  - **time, grading,** different**, course, class**, major, **submit**, product, like
  - companies**, interesting, class, thanks**, print, far, wonder, article
- Words in **LOGISTICS** get mixed up with **ACADEMIC** words
- **SeededLDA for MOOCs**
  - Many classes but a **common** set of seed words
  - Course-specific seed words from syllabus

Jagarlamudi et al. 2010

SeededLDA for MOOC Forums

- Seed topics with possible words from logistics and social posts

LOGISTICS: problem, issue, lecture, assignment, question
SOCIAL: introduction, study, group, coursera, learning

- Seed academic topics with words from syllabus

Course - Disruptive Technologies: disruptive, technology, innovation, survival,

Course - Women and Civil Rights: women, civil, rights, movement, political, …

Course – Gene and the Human Data: gene, sequence, disease, immunity, …

- *k* non-seeded topics to capture other topics

Adding Topics as Features

DIRECT +

Behavioral

Structural

Linguistic

- Sentiment
- **Topic of posts**

Temporal

→ Course Success

Topic as features in our success prediction models

Encoding Topics in PSL Rules

- LOGISTICS posts with negative sentiment implies dropping out

posts(John, P) ∧ **topic(P, LOGISTICS)** ∧ negative(P) → ¬success(John)

- SOCIAL posts indicate dropping out

posts(Mary, P) ∧ **topic(P, SOCIAL)** → ¬success(Mary)

- ACADEMIC posts that receive positive feedback indicate survival

posts(Lily, P) ∧ **topic(P, ACADEMIC)** ∧ upvote(P) → success(Lily)

- Students posting in similar topics have similar survival tendencies

posts(John, P1) ∧posts(Mary, P2) ∧ **topic(P1, ACADEMIC)** ∧ **topic(P2, ACADEMIC)** ∧ success(John) → success(Mary)

Experimental Results

COURSE	MODEL	AUC-PR (POS)	AUC-PR (NEG)	AUC-ROC
DISR	DIRECT	0.76	0.62	0.68
	DIRECT+TOPIC	**0.79**	**0.64**	**0.71**
WOMEN	DIRECT	0.65	0.89	0.82
	DIRECT+TOPIC	**0.67**	**0.90**	**0.83**
GENE	DIRECT	0.87	0.78	0.86
	DIRECT+TOPIC	**0.89**	**0.79**	**0.88**

- 3 Courses: Disruptive Technologies, Women & Civil Rights, and Gene & Human Condition
- Including topics improves AUC-PR and AUC-ROC for predicting student survival

Negative Sentiment in ACADEMIC VS LOGISTICS

Negative sentiment in ACADEMIC post (survived student)

- Our history keeps silence about **violence** in families, **contempt** and **unfair** treatment of women.

Negative sentiment in LOGISTICS post (dropped out student)

- The videos are totally **unsynchronized**. It's really **confusing** to hear the instructor talking about something while the image is telling a whole different story.

ACADEMIC posts with negative sentiment often indicate student engagement, hence student completion

Instructor Intervention in LOGISTICS posts

Unanswered **LOGISTICS** posts (dropped out student)

- There are some **mistakes** on quiz 2. Questions 3, 5, and 15 mark you **wrong** for answers that are correct.

Answered **LOGISTICS** posts (survived student)

- Lecture slides for the videos (week 5) **don't open** (at all)
- Hopefully the Terrell reading and the Lecture PowerPoints now open for you. Thanks for reporting this.

Answering logistics questions leads to increased student satisfaction, hence survival

Finding LOGISTICS posts and the problems mentioned automatically can help instructor intervention and avoid dropout

Fine-grained Aspect-Sentiment Models

Ramesh et al., ACL 2015

Sentiment in MOOC Posts

MOOC Post	Sentiment
Our history keeps silence about **violence** in families, **contempt** and **unfair** treatment of women.	Negative
I am from New York as well! Really **love** the city!	Positive
What I **love** about the lesson is how music influences social change. So many **great** songs can **motivate** change in people and society.	Positive
The video is very **jumpy**. I **hated** the experience.	Negative
Will everyone get a certificate or only people in the signature track?	Neutral
Will subtitles be made available for the lectures for this week? I **liked** the transcripts from last week.	Positive
When is quiz 4 due?	Neutral

Logistics

Aspect-Sentiment in MOOC Posts

MOOC Post	Sentiment	Aspect
Our history keeps silence about **violence** in families, **contempt** and **unfair** treatment of women.	Negative	Lecture-Content
I am from New York as well! Really **love** the city!	Positive	Social
What I **love** about the lesson is how music influences social change. So many **great** songs can **motivate** change in people and society.	Positive	Lecture-Content
The video is very **jumpy**. I **hated** the experience.	Negative	Lecture-Video
Will everyone get a certificate or only people in the signature track?	Neutral	Certificate
Will subtitles be made available for the lectures for this week? I **liked** the transcripts from last week.	Positive	Lecture-Subtitles
When is quiz 4 due?	Neutral	Quiz-Deadlines

http://psl.umiacs.umd.edu

Aspect Hierarchy

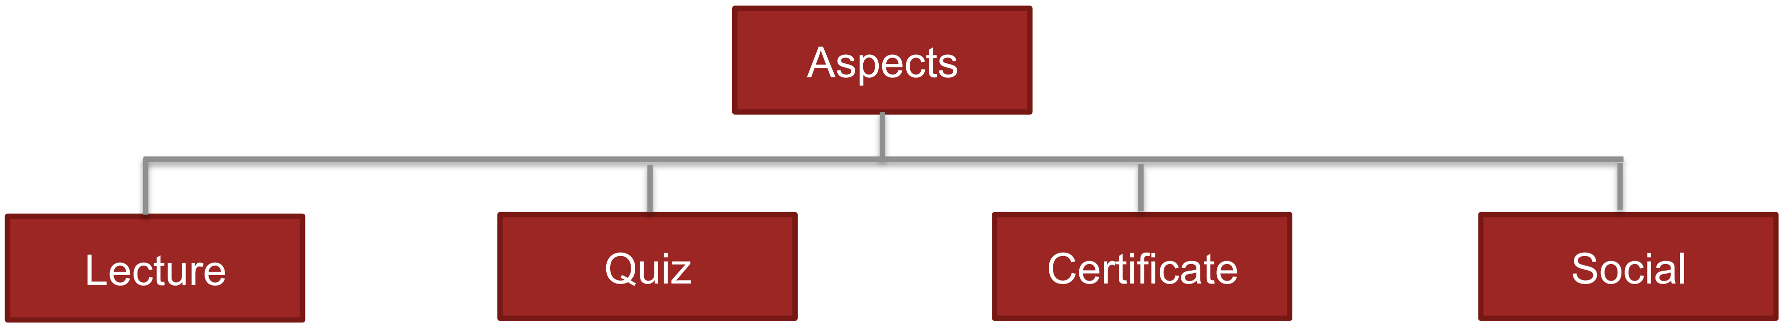

Aspect Hierarchy

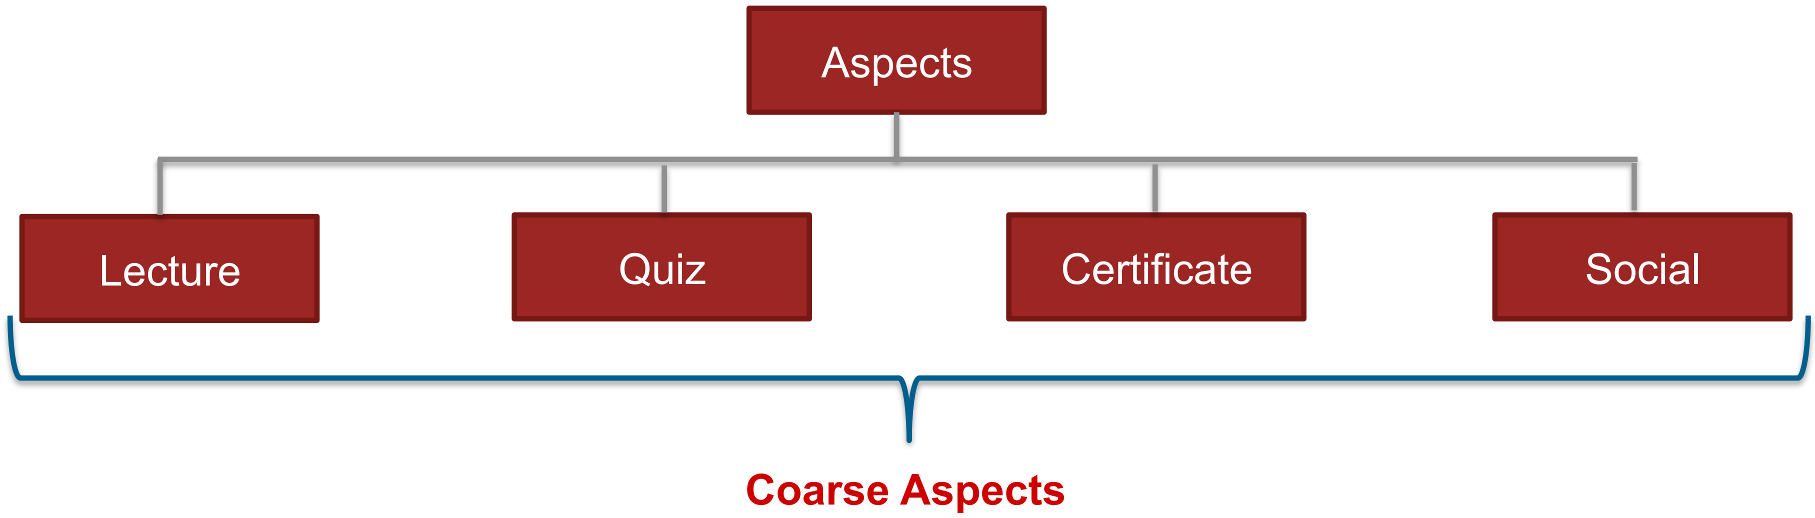

http://psl.umiacs.umd.edu

Aspect Hierarchy

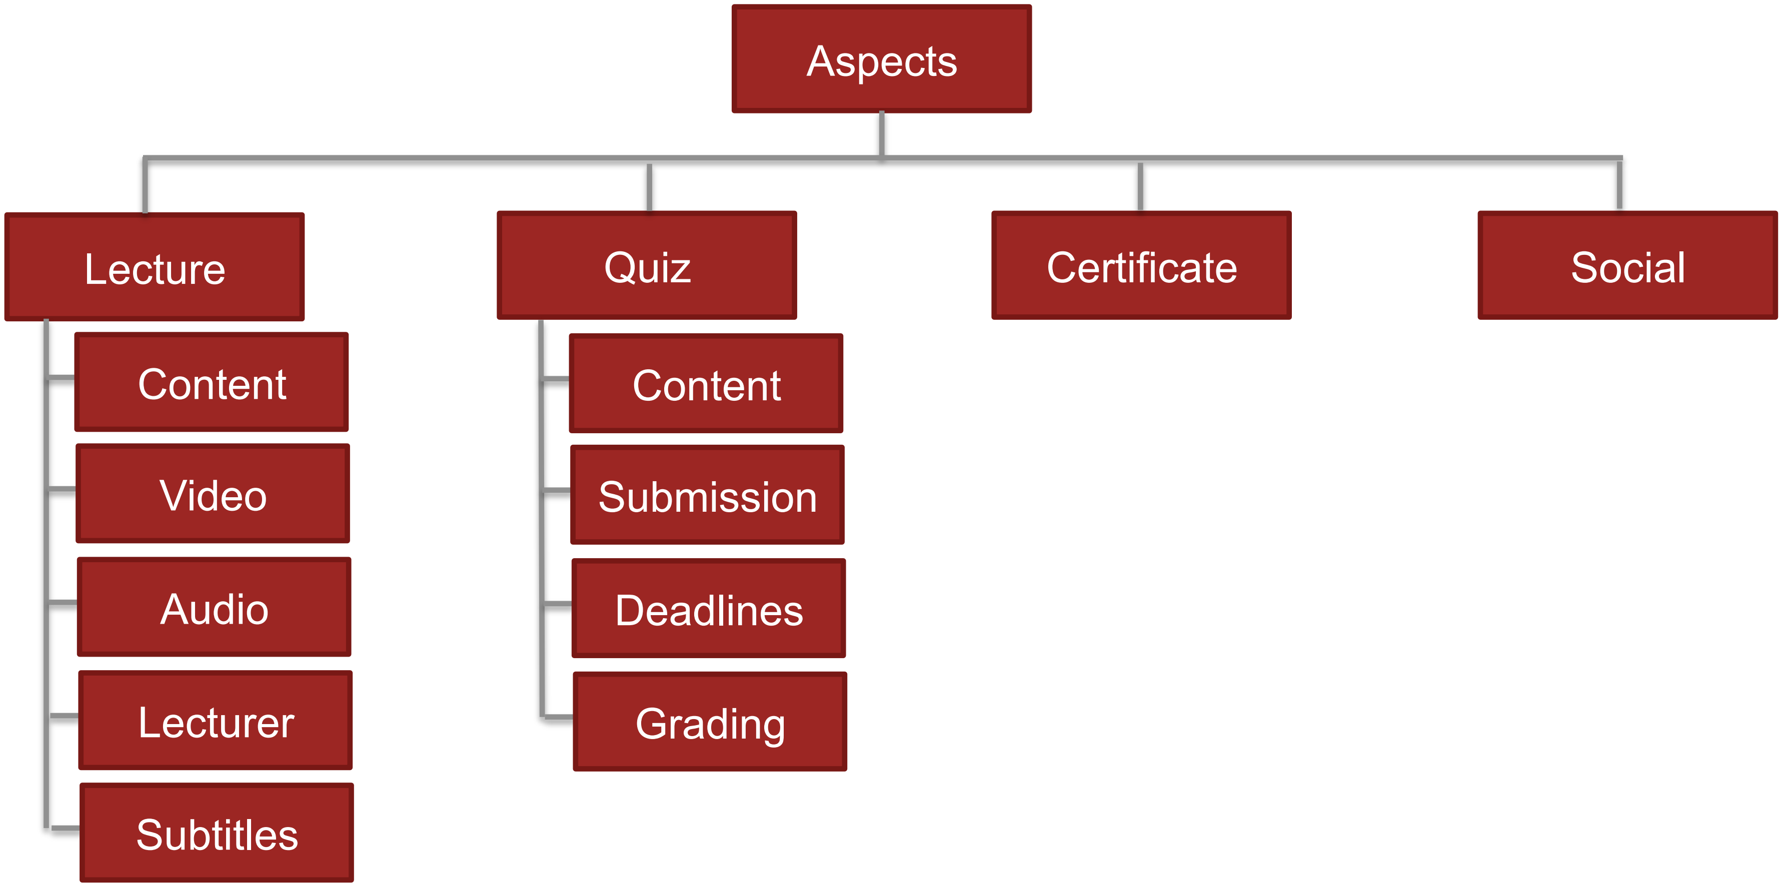

Aspect Hierarchy

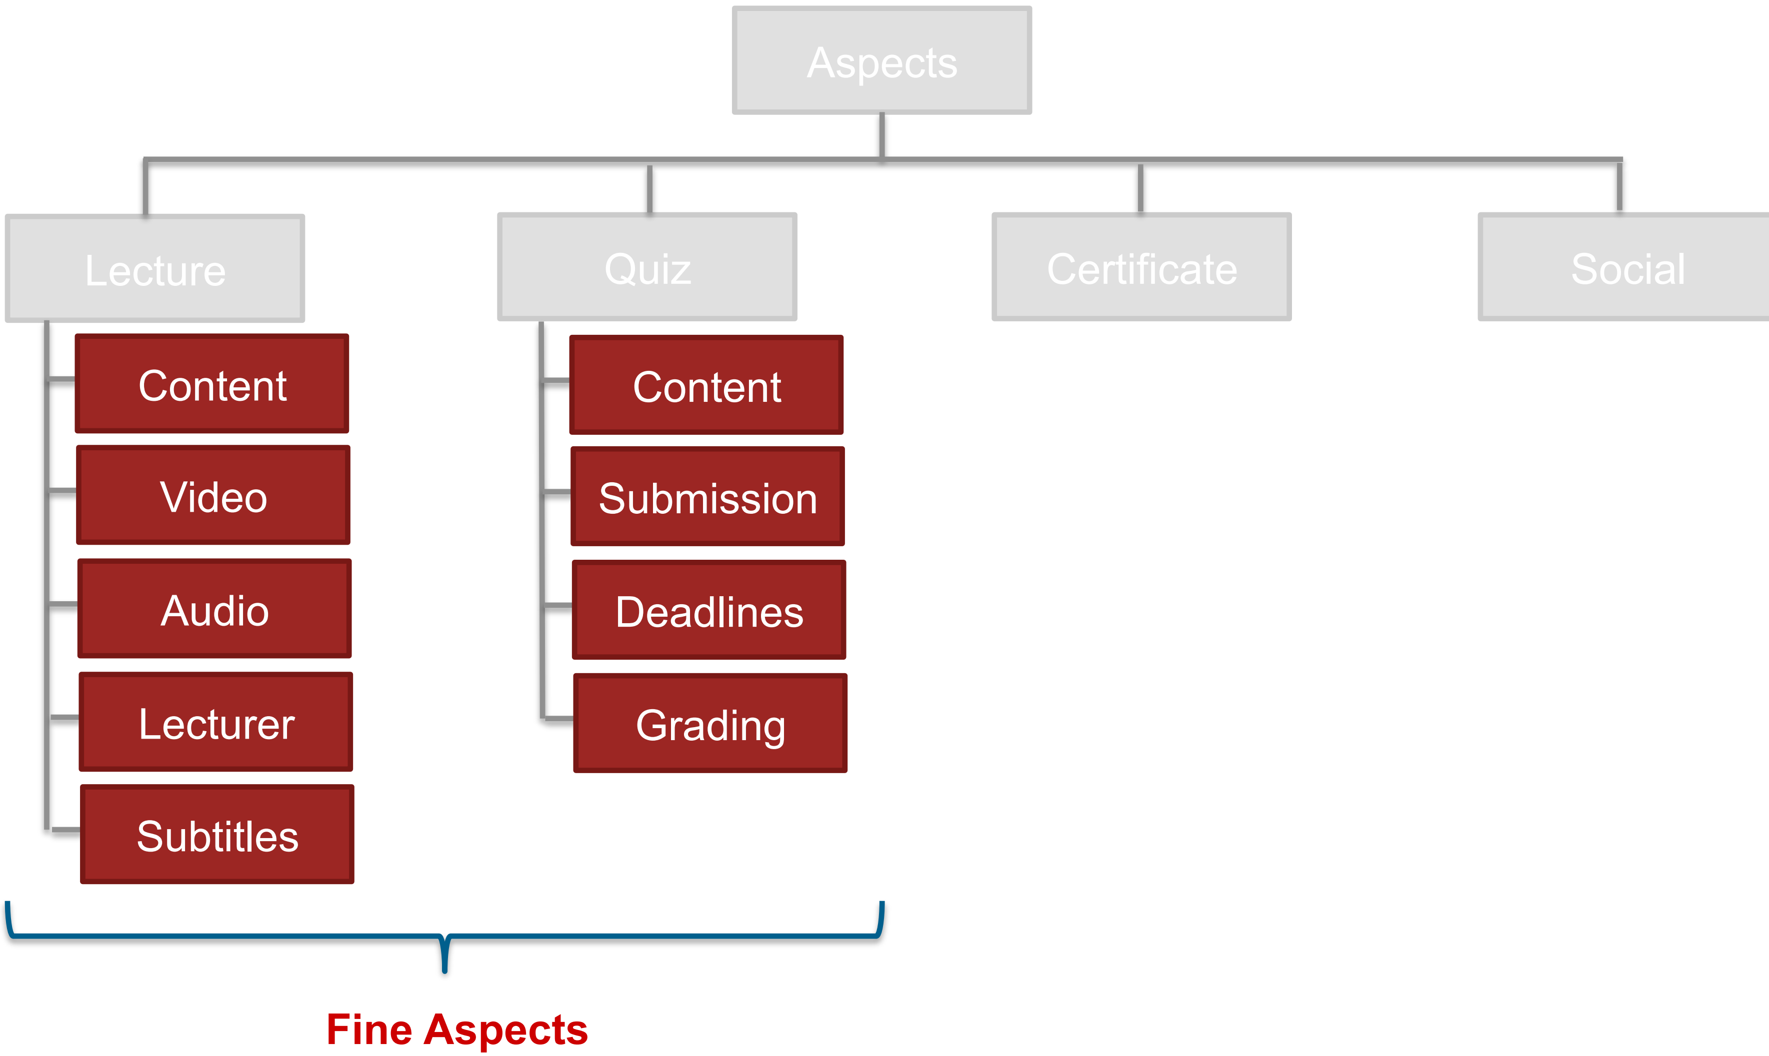

Predicting Fine-grained Aspects: Challenges

- Labeled data hard to obtain
  - **5-10%** posts contain problems
  - Unsupervised/weakly supervised approaches desirable
  - Privacy concerns around data sharing
- Aspects differ across courses
  - Careful mining of posts required to come up with exhaustive aspect categories
  - System not fine-tuned to one course, but can adapt across courses

Aspect Sentiment Prediction Models

- **Seeded LDA**
  - 3 seeded LDA models for coarse aspect, fine aspect and sentiment
- **PSL-Joint**
  - Combining features
    - Multiple seeded LDA features
  - Encoding dependencies
    - Hierarchical dependence between aspects
    - Dependence between aspect and sentiment

Empirical Evaluation

- Twelve courses across multiple disciplines
  - Disruptive Technologies (*Business*)
  - Gene and Human Condition (*Science*)
  - Women and the Civil Rights Movement (*History*)
  - Developing Innovative Ideas for New Companies (*Business*)
  - Android Applications (*Computer Science*)
  - Making Better Group Decisions (Philosophy)
  - Religious Tolerance in Religious Society (History)
- Dataset contains
  - Posts sampled from 12 courses
  - Crowdsourced labels for
    sentiment, coarse aspect, and fine aspect (evaluation)

Evaluation: Coarse Aspect and Sentiment

F1 scores for SeededLDA and PSL-Joint for **coarse** aspects

Model	Lecture	Quiz	Certificate	Social
SeededLDA	0.632	0.657	0.459	0.654
PSL-Joint	0.630	**0.706**	**0.621**	**0.659**

F1 scores for SeededLDA and PSL-Joint for **sentiment**

Model	Positive	Negative	Neutral
SeededLDA	0.182	0.517	0.356
PSL-Joint	**0.189**	**0.615**	**0.434**

Evaluation: Fine Aspect

Fine-grained aspects under coarse aspect **lecture**

Model	Content	Video	Audio	Lecturer	Subtitles
SeededLDA	0.08	0.240	**0.684**	0.06	0.397
PSL-Joint	**0.410**	**0.485**	0.582	**0.323**	**0.461**

Fine-grained aspects under coarse aspect **quiz**

Model	Content	Submission	Deadlines	Grading
SeededLDA	0.011	**0.437**	0.214	0.514
PSL-Joint	**0.36**	0.416	**0.611**	**0.550**

Interpreting PSL-Joint Predictions

"There is a typo or other mistake in the **assignment** instructions (e.g., essential information omitted)."

SeededLDA Prediction: **Lecture-content**
PSL-Joint Prediction: **Quiz-content**

"Thanks for the suggestion about downloading the video and referring to the subtitles. **The audio is barely audible, even when the volume is set to 100%**"

SeededLDA Prediction: **Lecture-subtitles**
PSL-Joint Prediction: **Lecture-audio**

Topic Evolution Models

Data

- Business Course: 34 iterations

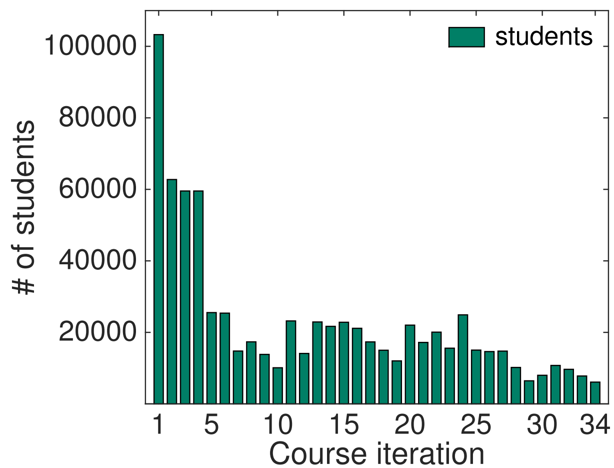

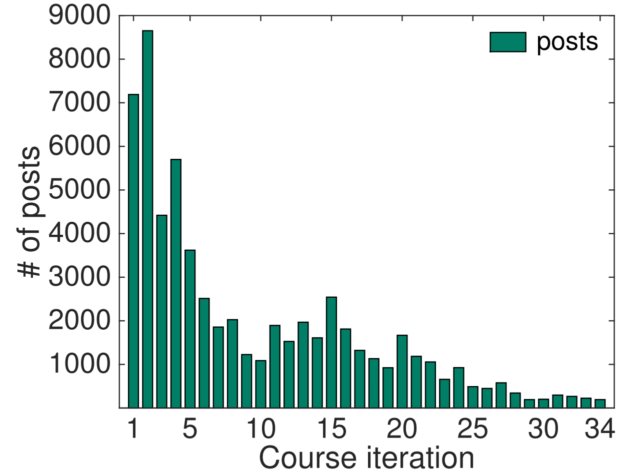

- CS Course: 15 offerings, 9 total iterations

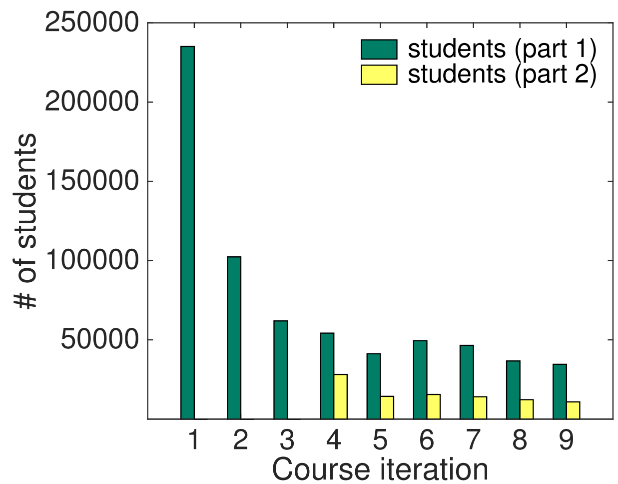

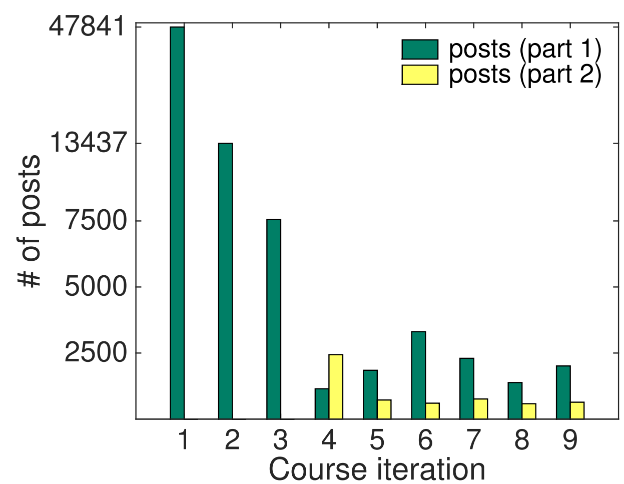

Business Course: Primary Use of Forums

Evolution of social, issue and technical topics across iterations

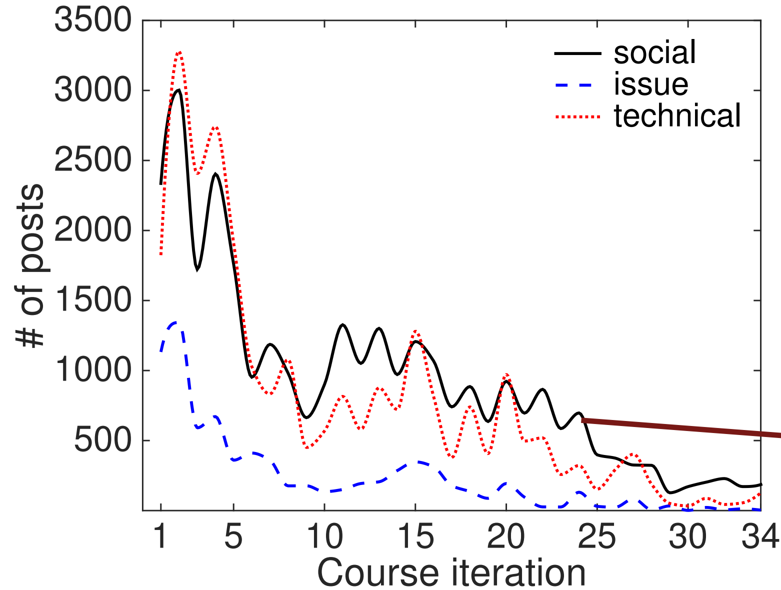

Social posts dominate the forums across iterations

Business Course: Primary Use of Forums

Evolution of social, issue and technical topics across iterations

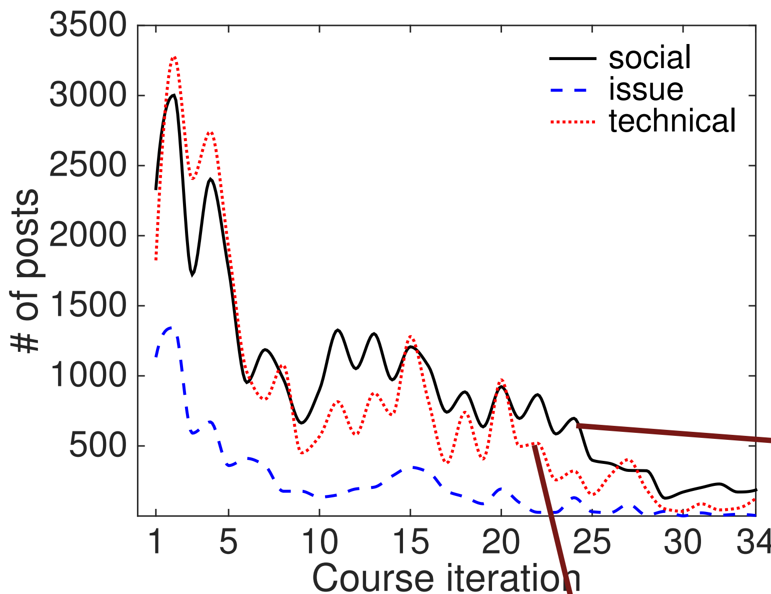

Social posts dominate the forums across iterations

Followed by technical posts

Business Course: Primary Use of Forums

Evolution of social, issue and technical topics across iterations

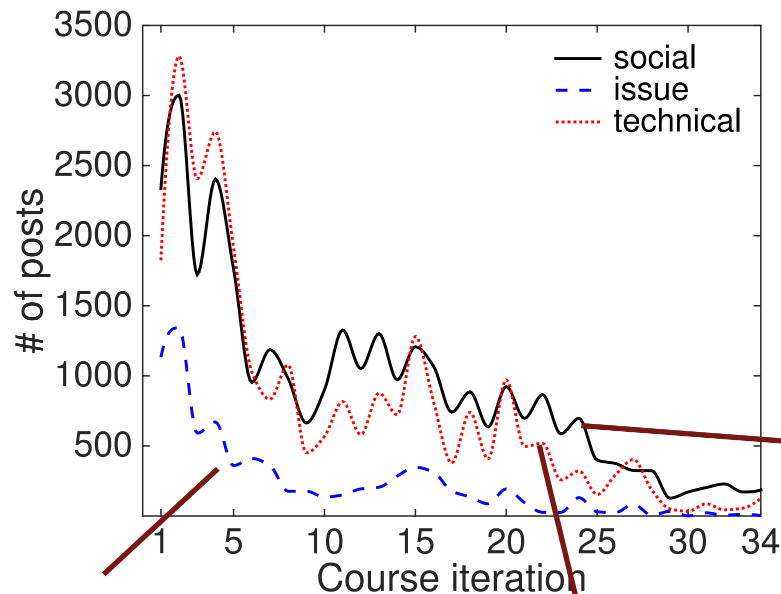

Issues decline significantly, falling to negligible numbers after 20 iterations

Followed by technical posts

Social posts dominate the forums across iterations

Forums mostly used for **socializing** and **course-related discussions** in later iterations

CS Course: Primary Use of Forums

Evolution of social, issue and technical topics across iterations

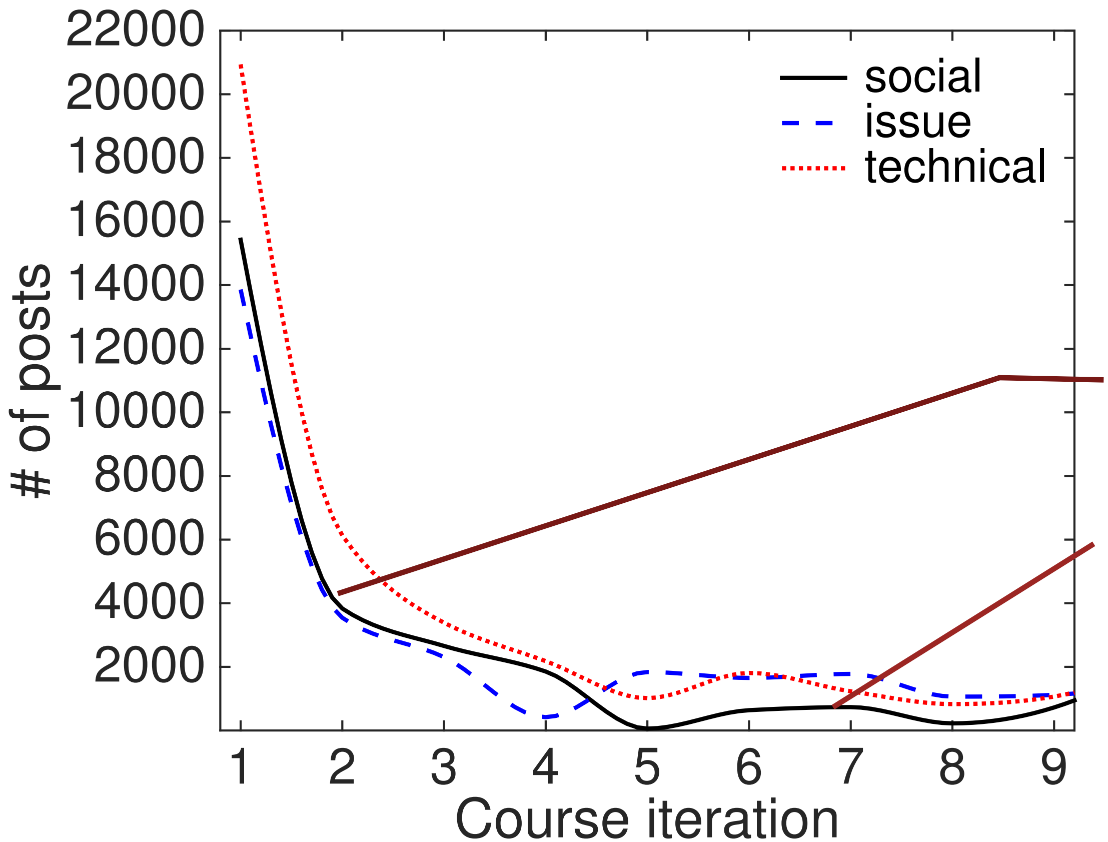

Social posts are more prominent in the early iterations, but fall lower than issue and technical posts toward the end

CS Course: Primary Use of Forums

Evolution of social, issue and technical topics across iterations

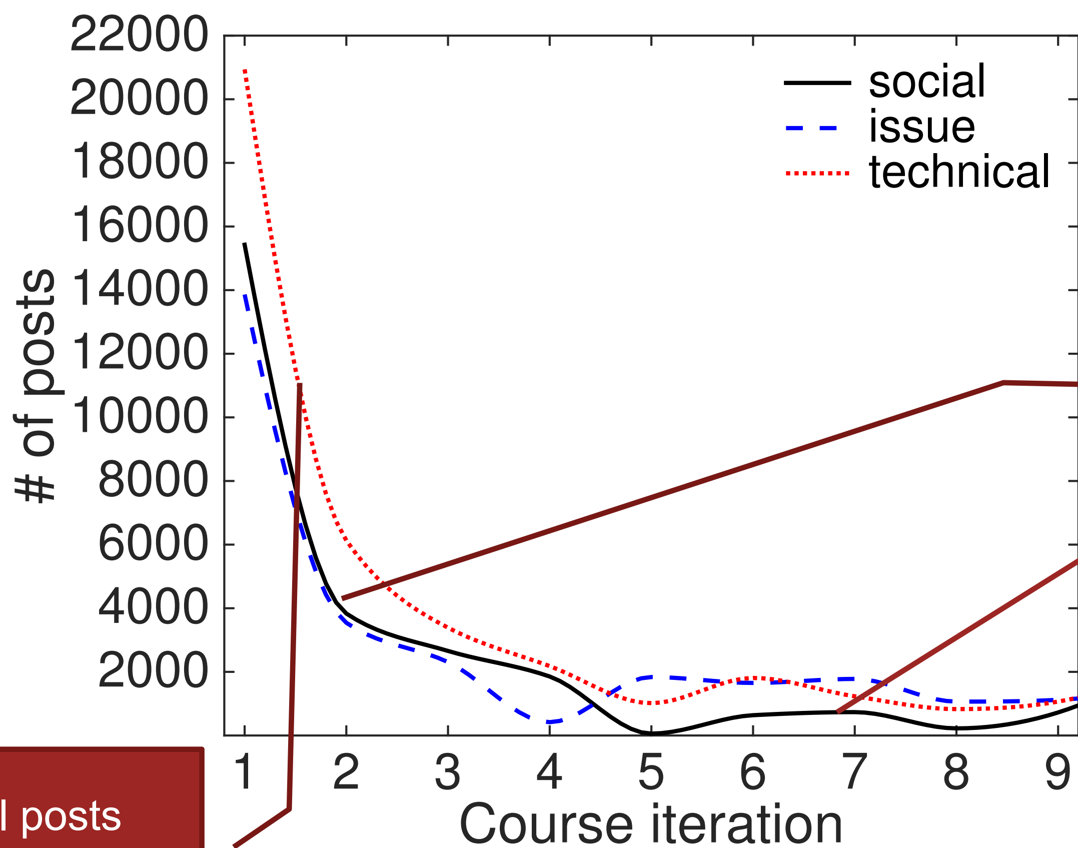

Social posts are more prominent in the early iterations, but fall lower than issue and technical posts toward the end

Technical posts dominate the course across all iterations

CS Course: Primary Use of Forums

Evolution of social, issue and technical topics across iterations

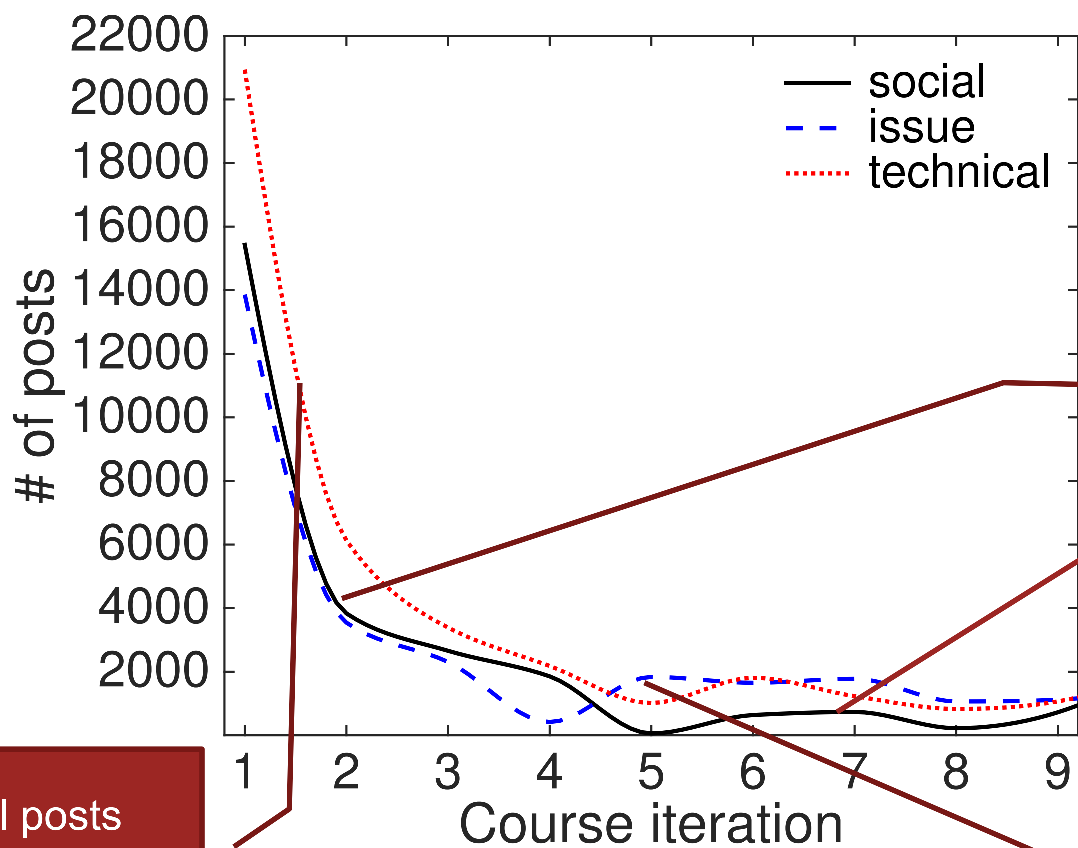

Social posts are more prominent in the early iterations, but fall lower than issue and technical posts toward the end

Technical posts dominate the course across all iterations

Issue posts more significant in later iterations

Effectiveness of Interventions: Coaching

High-school MOOCs

- Performance of MOOC students
- Effectiveness of coaching & forums
- Predictive behaviors for post-test performance (AP exam)

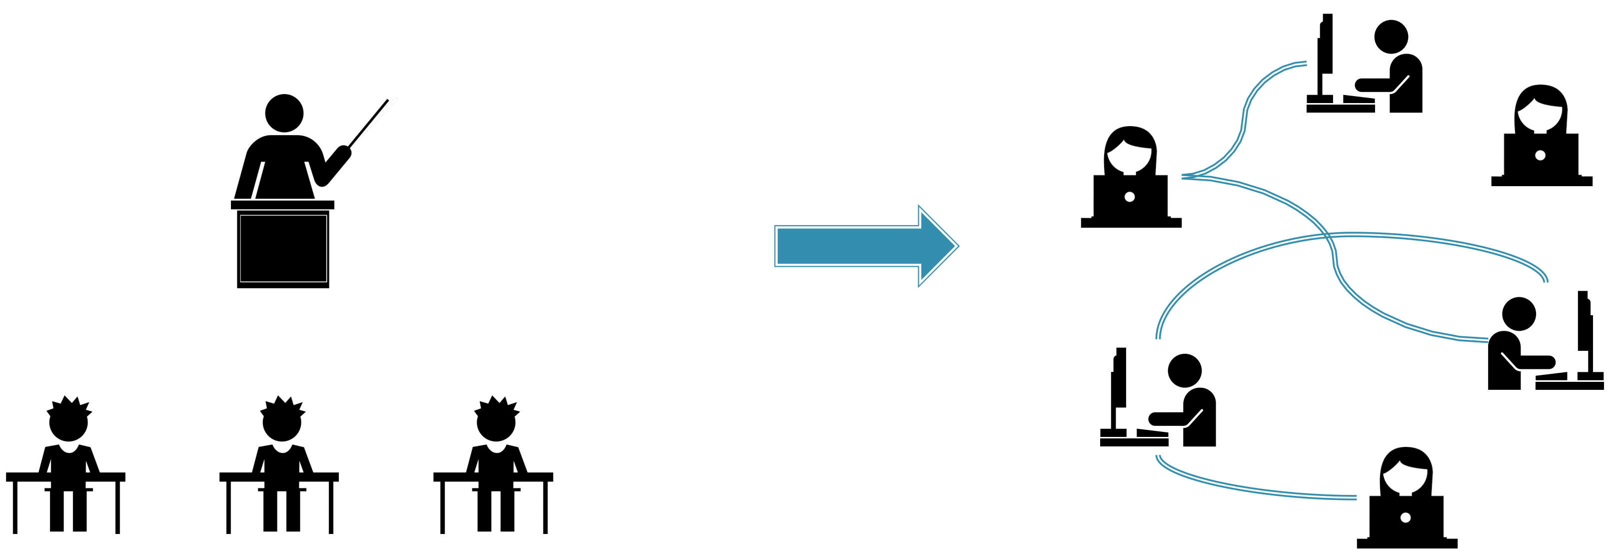

Tomkins et al., EDM 2016

Discussion & Summary

Socio-Behavioral Modeling in Education

- Requires rich models that capture content, behavior and outcomes
- Latent variables important
  - Engagement, sentiment, topics
  - Outcomes to validate latent modeling
    - Course completion, performance, response
- Opportunities
  - mix data-driven course-specific modeling with pedagogical knowledge-driven modeling
  - Direct interventions

Thank You!

Probabilistic
Soft
Logic

psl.umiacs.umd.edu