Statistical Methods for Integration and Analysis of Online Opinionated Text Data

ChengXiang ("Cheng") Zhai

Department of Computer Science
University of Illinois at Urbana-Champaign
http://www.cs.uiuc.edu/homes/czhai

Joint work with Yue Lu, Qiaozhu Mei, Kavita Ganesan, Hongning Wang, and others

Online opinions cover all kinds of topics

Topics:
People
Events
Products
Services, …

Sources:
Blogs
Microblogs
Forums
Reviews ,…

45M reviews **53M** blogs **1307M** posts **65M** msgs/day **115M** users **10M** groups …

tripadvisor

blogScope PREVIEW

twitter

YAHOO! GROUPS

Great opportunities for many applications

Opinionated Text Data

Decision Making & Analytics

"Which cell phone should I **buy**?"

"What are the winning features of **iPhone** over **blackberry**?"

"How do people like this new **drug**?"

"How is Obama's health care **policy** received?"

"Which presidential candidate should I **vote** for?"

...

However, it's not easy to for users to make use of the online opinions

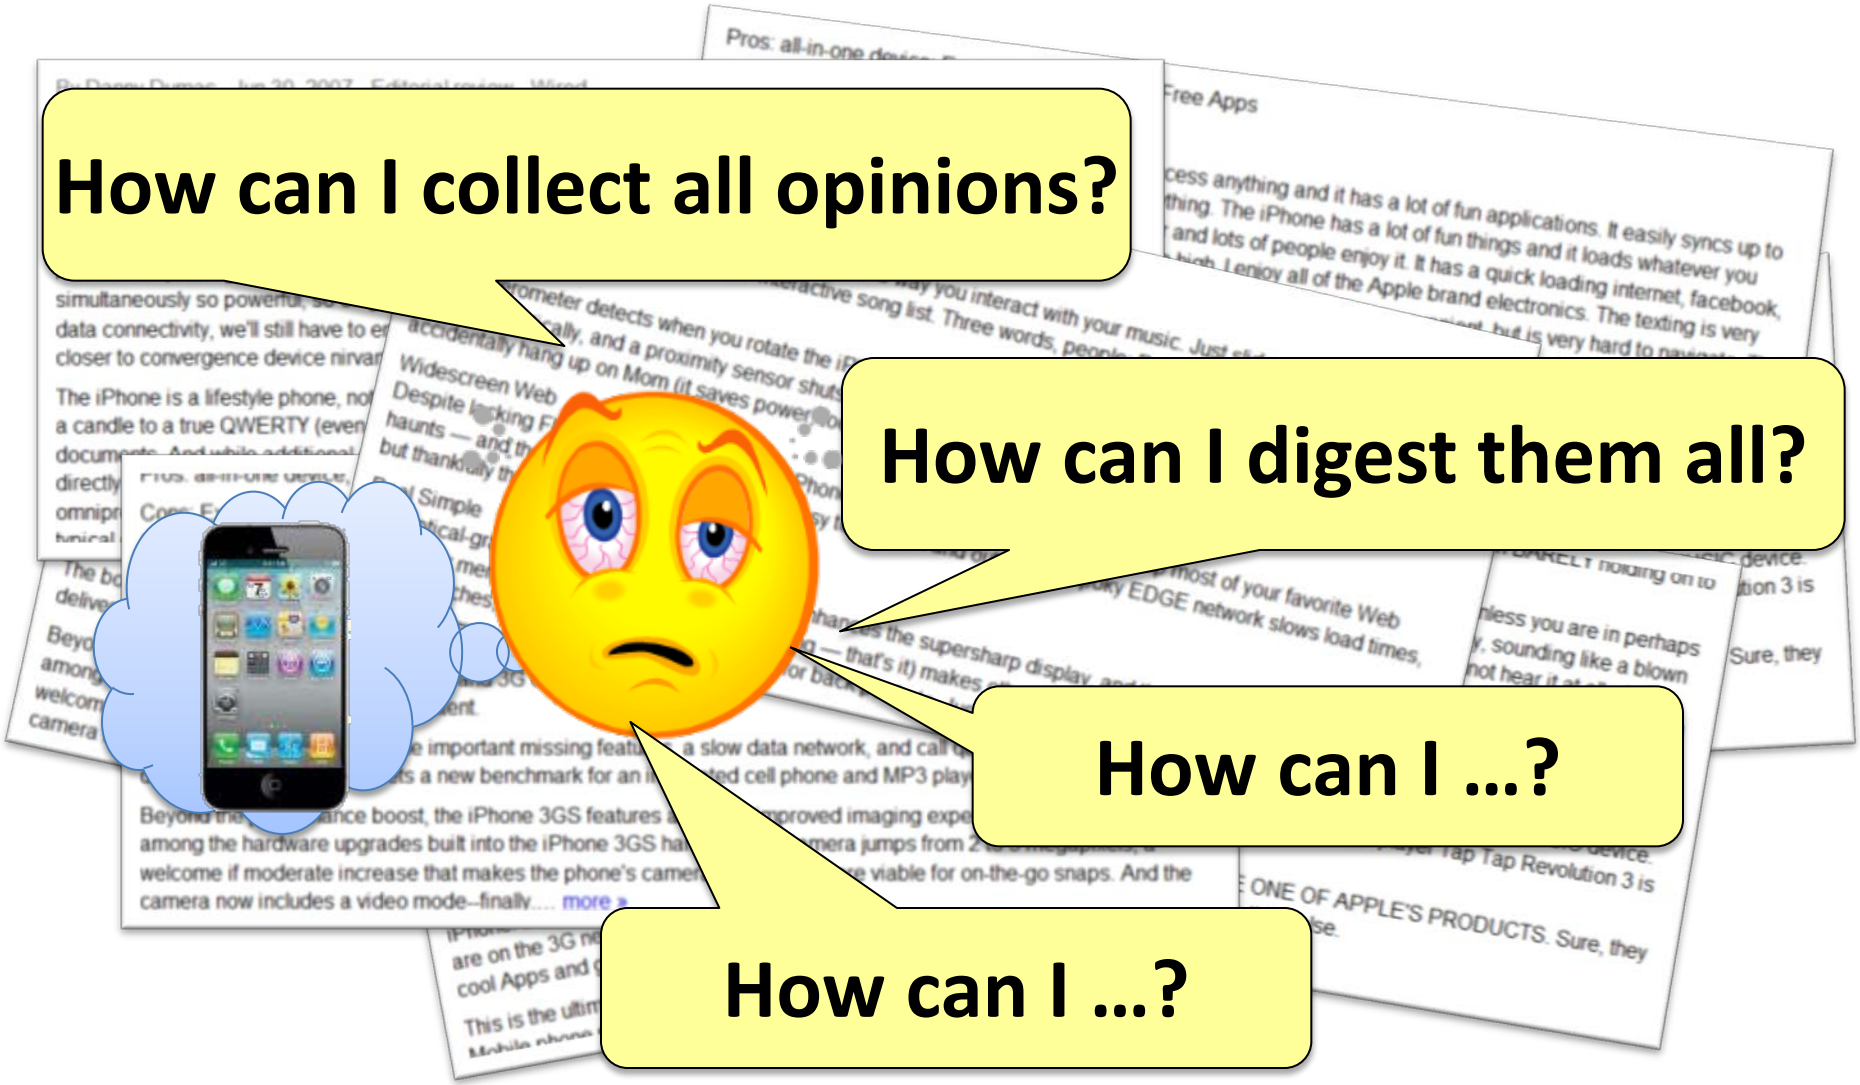

Research Questions

- How can we integrate scattered opinions?
- How can we summarize opinionated text articles?
- How can we analyze online opinions to discover patterns and understand consumer preferences?
- How can we do all these in a general way with no or minimum human effort?
  - **Must work for all topics**
  - **Must work for different natural languages**

Solutions: Knowledge-Lean Statistical Methods (Statistical Language Models)

Lots of related work (usually not as general):

Bing Liu, Sentiment Analysis and Opinion Mining, Morgan & Claypool Publishers, 2012

Rest of the talk: general methods for

1. Opinion Integration

2. Opinion Summarization

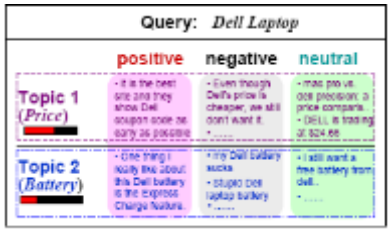

3. Opinion Analysis

Decision Making & Analytics

"Which cell phone should I **buy**?"

"What are the winning features of **iPhone** over **blackberry**?"

"How do people like this new **drug**?"

"How is Obama's health care **policy** received?"

"Which presidential candidate should I **vote** for?"

...

Outline

1. Opinion Integration

2. Opinion Summarization

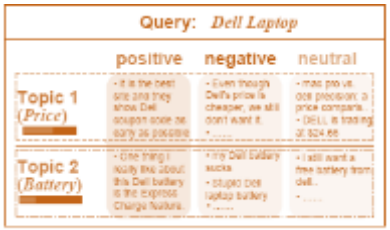

3. Opinion Analysis

Decision Making & Analytics

"Which cell phone should I **buy**?"

"What are the winning features of **iPhone** over **blackberry**?"

"How do people like this new **drug**?"

"How is Obama's health care **policy** received?"

"Which presidential candidate should I **vote** for?"

...

How to digest all scattered opinions?

Need tools to automatically integrate all scattered opinions

Observation: two kinds of opinions

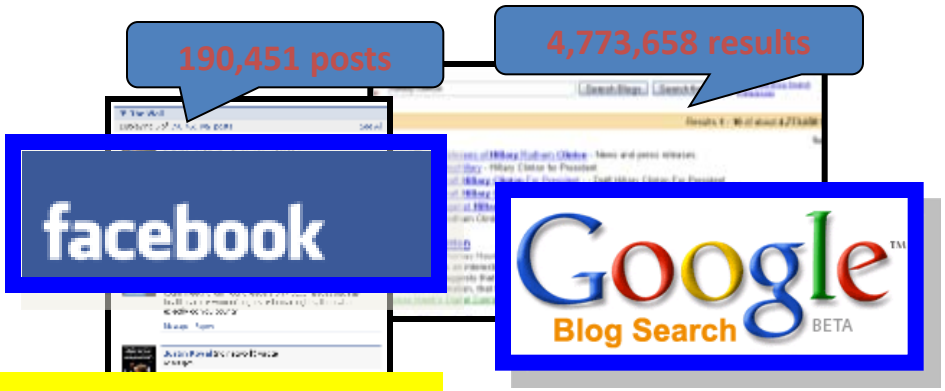

Can we combine them?

Expert opinions	Ordinary opinions
•CNET editor's review •Wikipedia article	•Forum discussions •Blog articles
•Well-structured •Easy to access	•Represent the majority •Up to date
•Maybe biased •Outdated soon	•Hard to access •fragmented

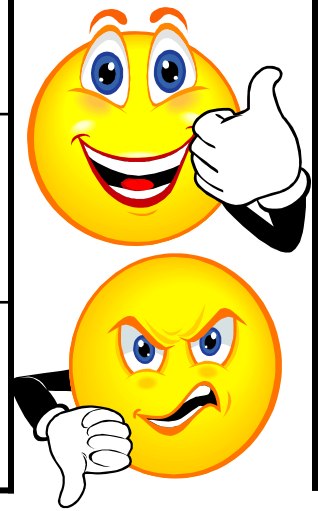

Opinion Integration Strategy 1

[Lu & Zhai WWW 08]

Align scattered opinions with well-structured expert reviews

Yue Lu, ChengXiang Zhai. Opinion Integration Through Semi-supervised Topic Modeling, *Proceedings of the World Wide Conference 2008* (**WWW'08**), pages 121-130.

Review-Based Opinion Integration

Input

Topic: iPod

Expert review with aspects

DesignBatteryPrice..

Text collection of ordinary opinions, e.g. Weblogs

Output

Review Aspects

	Similar opinions	Supplementary opinions
Design	cute… tiny…	..thicker..
Battery	last many hrs	die out soon
Price	could afford it	still expensive

Extra Aspects

iTunes	… easy to use…
warranty	…better to extend..

Integrated Summary

Solution is based on probabilistic latent semantic analyis (PLSA) [Hofmann 99]

Topic model
= unigram language model
= multinomial distribution

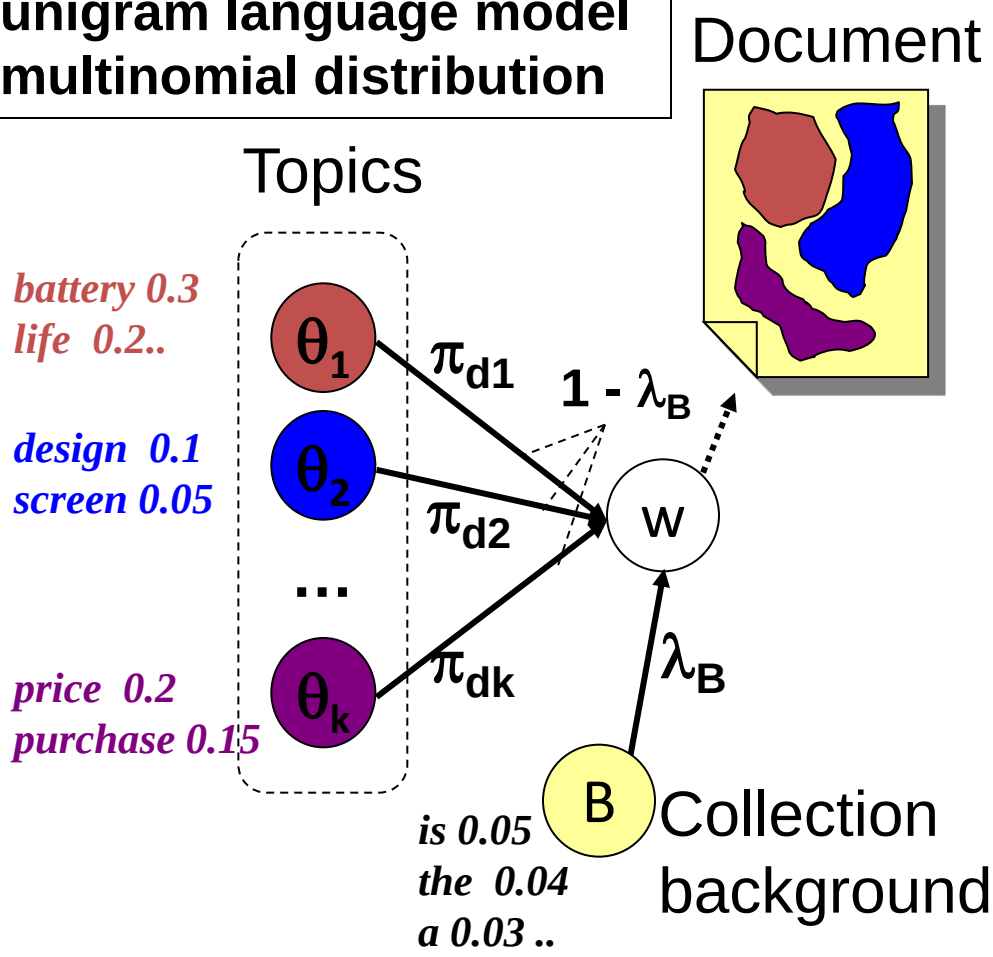

Generate a word in a document

$$p_d(w) = \lambda_B p(w|\theta_B) + (1-\lambda_B)\sum_{j=1}^{k}[\pi_{d,j}p(w|\theta_j)]$$

Basic PLSA: Estimation

Generate a word in a document

$$p_d(w) = \lambda_B p(w|\theta_B) + (1-\lambda_B)\sum_{j=1}^{k}[\pi_{d,j} p(w|\theta_j)]$$

Log-likelihood of the collection

Count of word in the document

$$\log p(\mathcal{C}_O|\Lambda) = \sum_{d\in\mathcal{C}_O}\sum_{w\in V}\{c(w,d)\times \log p_d(w)\}$$

$p(w|\theta_j)$ $\pi_{d,j}$

- **Parameters estimated with Maximum Likelihood Estimator (MLE) through an EM algorithm**

$$\hat{\Lambda} = \arg\max_{\Lambda} \log p(\mathcal{C}_O|\Lambda)$$

Semi-supervised Probabilistic Latent Semantic Analysis

Cast review aspects as conjugate Dirichlet priors

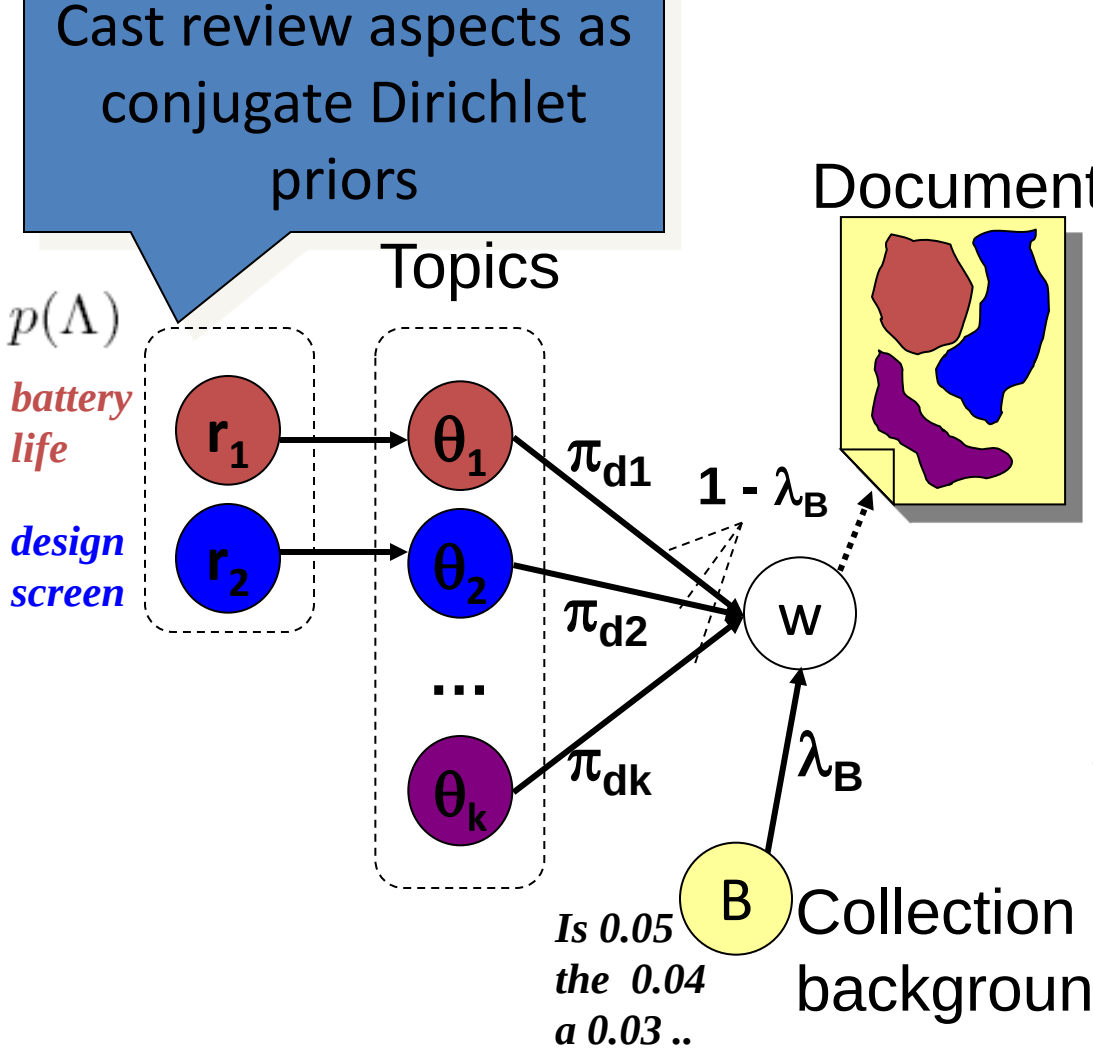

Document

Maximum Likelihood Estimation (MLE)

$$\hat{\Lambda} = \arg\max_{\Lambda} \log p(\mathcal{C}_O|\Lambda)$$

Maximum A Posterior (MAP) Estimation

$$\hat{\Lambda} = \arg\max_{\Lambda} \log p(\mathcal{C}_O|\Lambda)\, p(\Lambda)$$

Results: Product (iPhone)

- **Opinion Integration with review aspects**

Review article	Similar opinions	Supplementary opinions
You can make emergency calls, but you can't use any other functions…	N/A	… methods for **unlocking** the iP[illegible]ed on the Ir[illegible] few weeks, although they involve tinkering with the iPhone hardware…
rated battery life of **8 hours talk time**, **24 hours of music** playback, **7 hours of video** playback, and 6 [illegible]ernet use.	Up to **8 Hours of Talk Time**, 6 Hours of Internet Use, **7 Hours of Video** Playback or **24 Hours of Audio** Playback	Playing relatively high bitrate VGA H.264 videos, our iPhone lasted almost exactly **9 freaking hours** of continuous playback with cell and WiFi on (but Bluetoo[illegible]

Activation

Confirm the opinions from the review

Unlock/hack iPhone

Battery

Additional info under real usage

Results: Product (iPhone)

- **Opinions on extra aspects**

support	Supplementary opinions on extra aspects
15	You may have heard of **iASign** ... an iP... t allows you to **activate** your phone wit... e iTunes rigamarole.
13	**Cisco** has owned the **trademark** on the name "**iPhone**" since 2000, when it acquired InfoG... hich originally registered the name...
13	With the imminent availabilit... hone, a look at 10 things current smartphones like the **Nokia N95** have been able ... the **iPhone can't currently match**...

Another way to activate iPhone

iPhone trademark originally owned by Cisco

A better choice for smart phones?

As a result of integration…

What matters most to people?

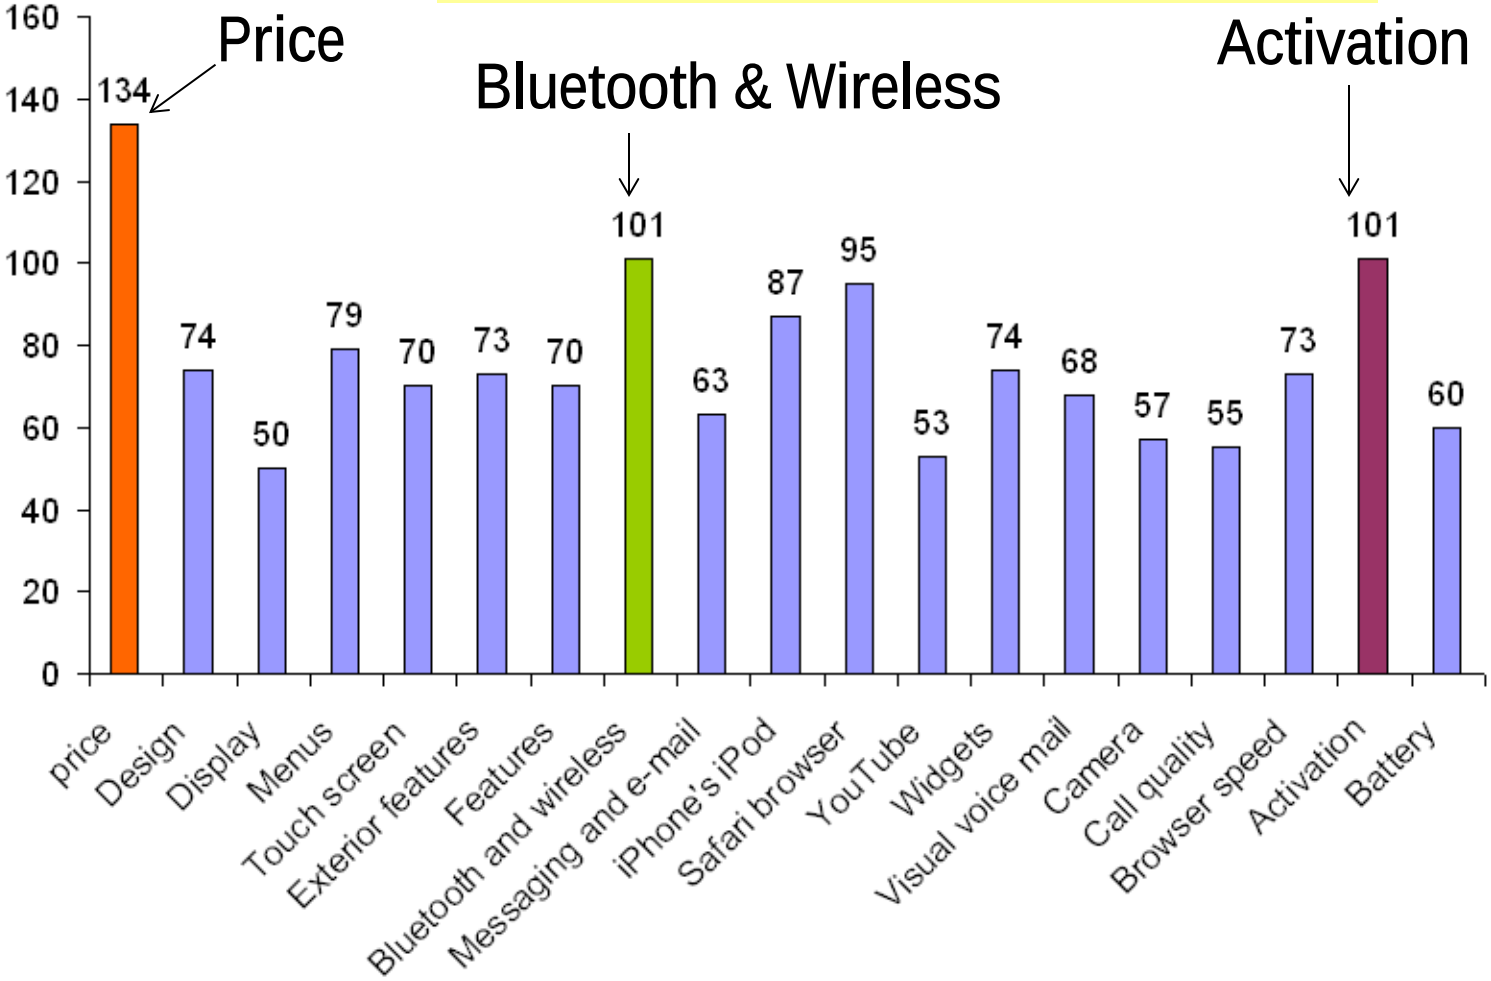

What if we don't have expert reviews?

How can we organize scattered opinions?

Exploit online ontology!

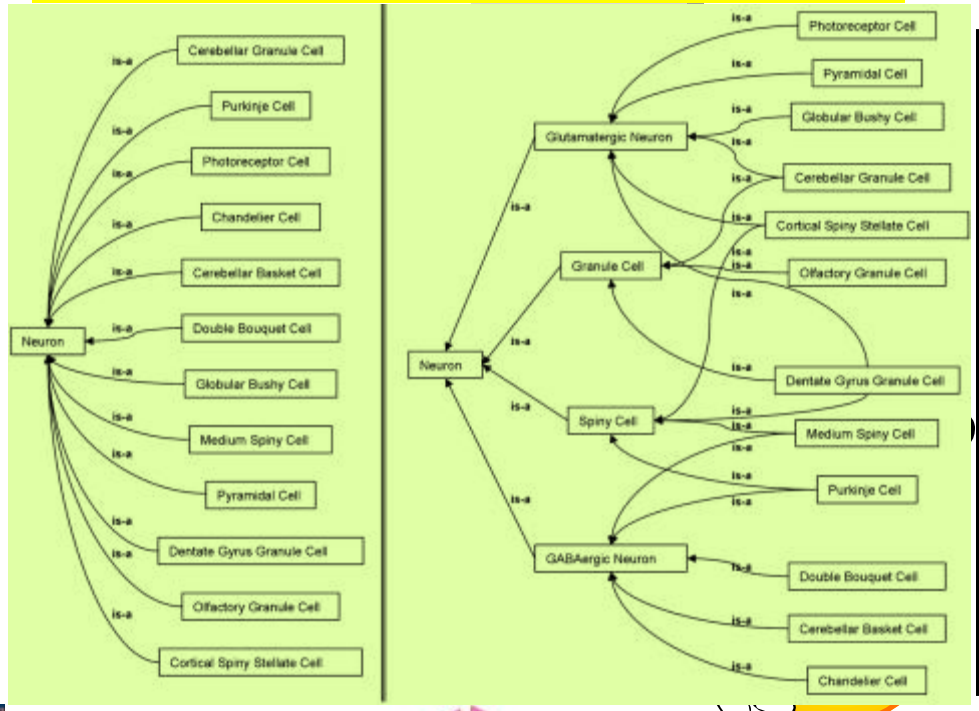

Ordinary opinions
•Forum discussions •Blog articles
•Represent the majority •Up to date
•Hard to access •fragmented

Opinion Integration Strategy 2

[Lu et al. COLING 10]

Organize scattered opinions using an ontology

Yue Lu, Huizhong Duan, Hongning Wang and ChengXiang Zhai. Exploiting Structured Ontology to Organize Scattered Online Opinions, *Proceedings of COLING 2010 (**COLING 10**)*, pages 734-742.

Sample Ontology: Freebase

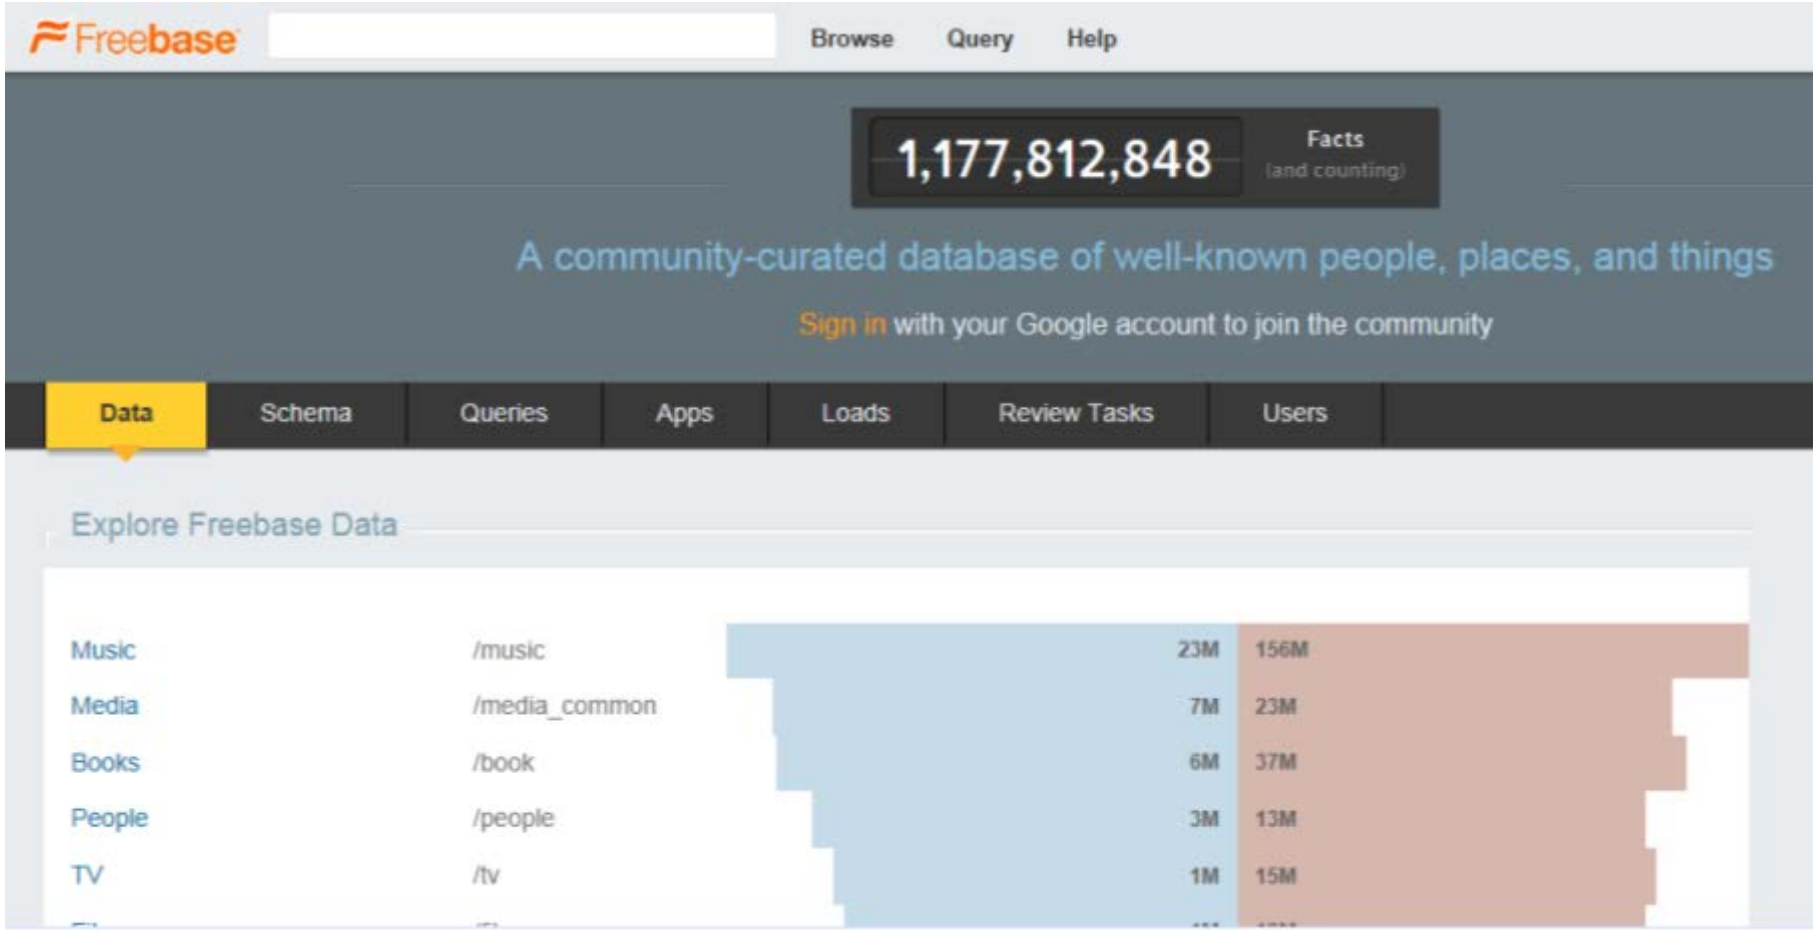

Ontology-Based Opinion Integration

Two key tasks: 1. Aspect Selection. 2. Aspect Ordering

Topic = "Abraham Lincoln"
(Exists in ontology)

Aspects from Ontology
(more than 50)

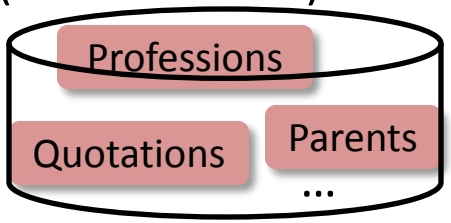

Online Opinion Sentences

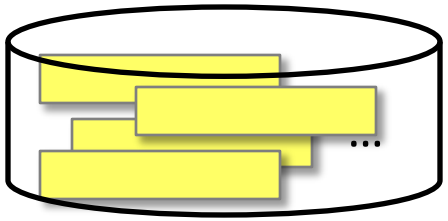

Subset of Aspects
Ordered to optimize readability

Matching Opinions

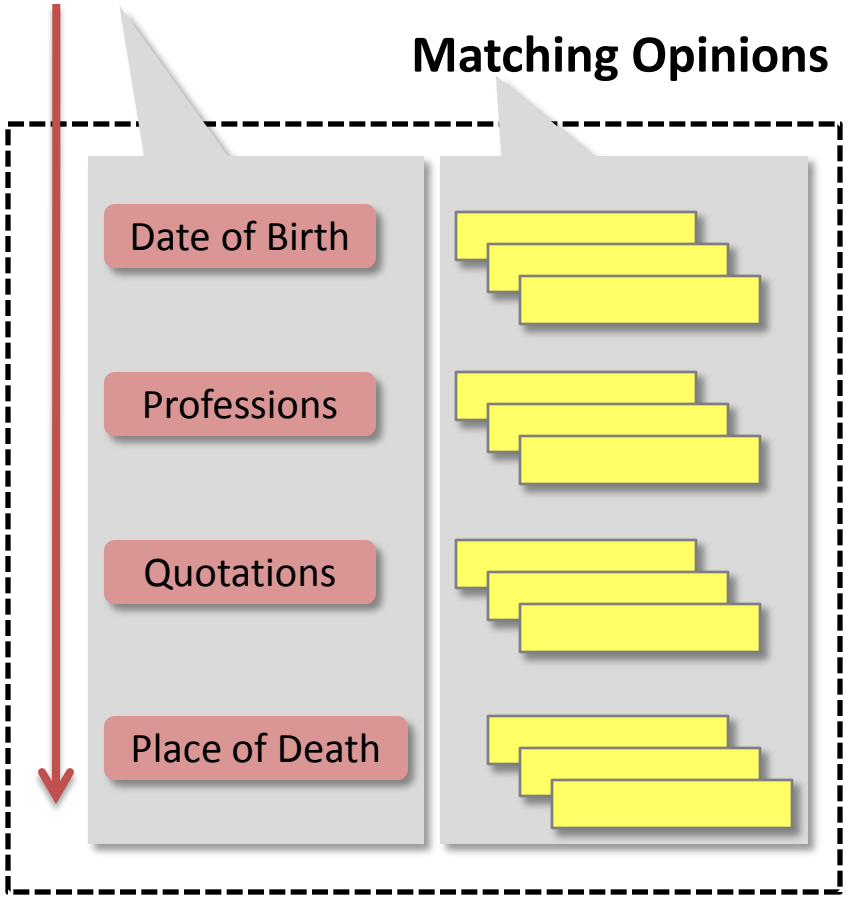

1. Aspect Selection: Conditional Entropy-based Method

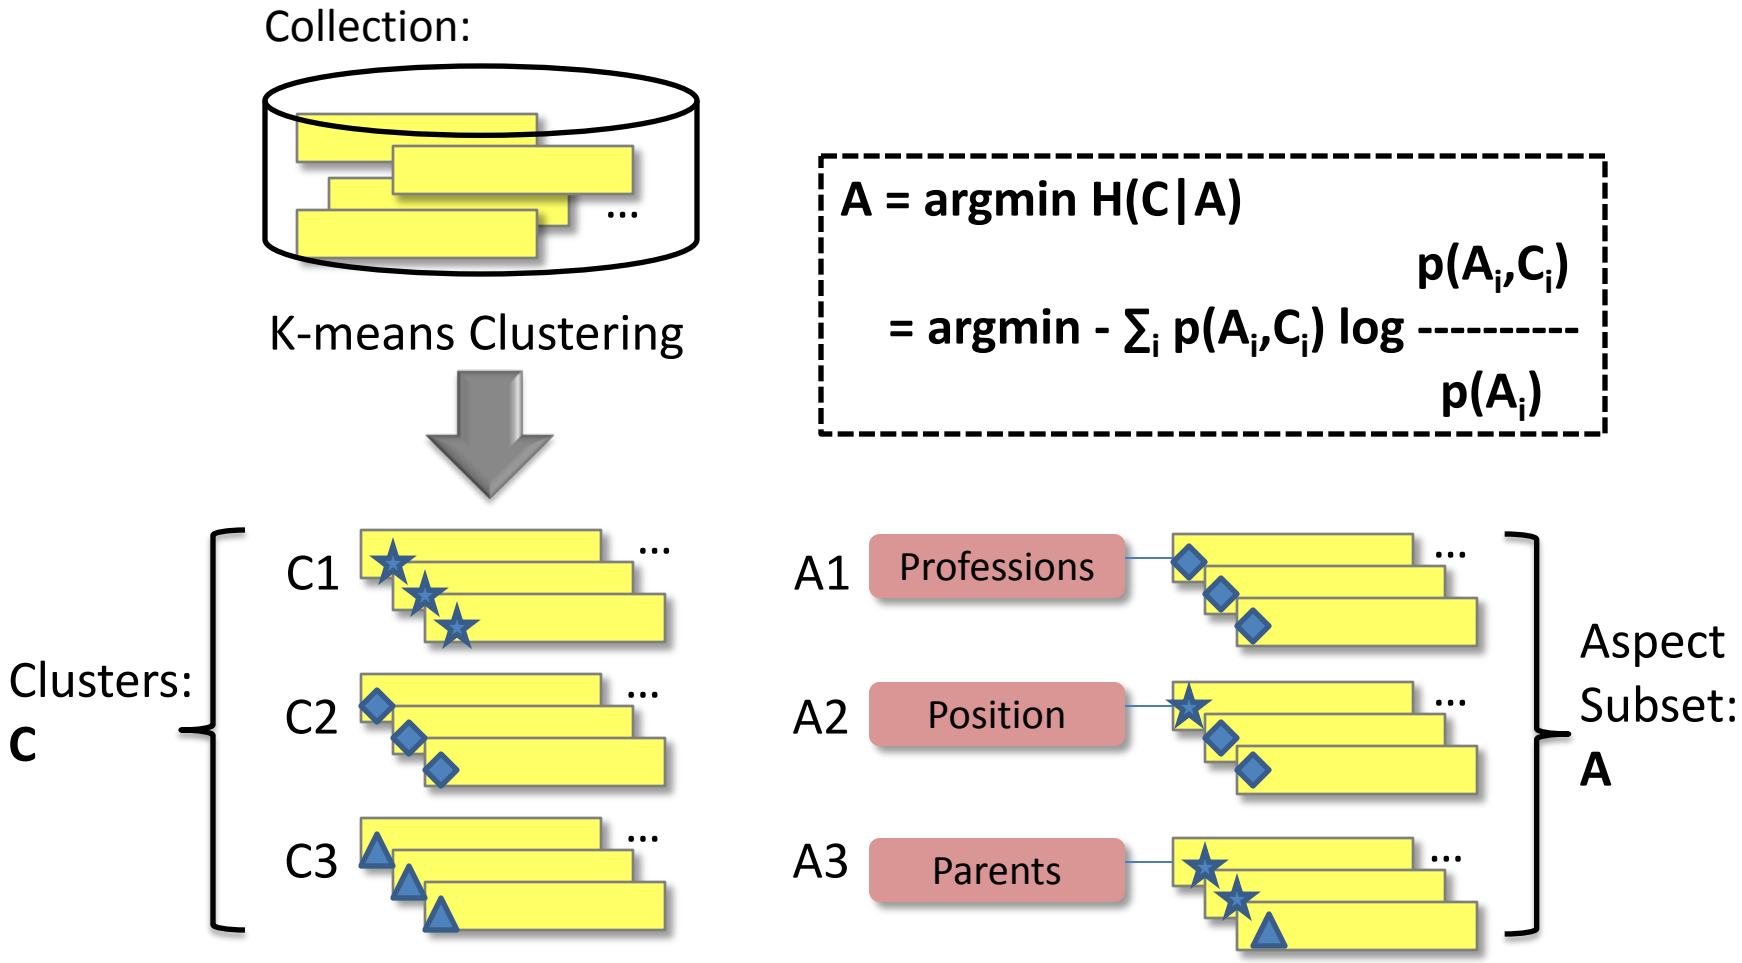

2. Aspect Ordering: Coherence Order

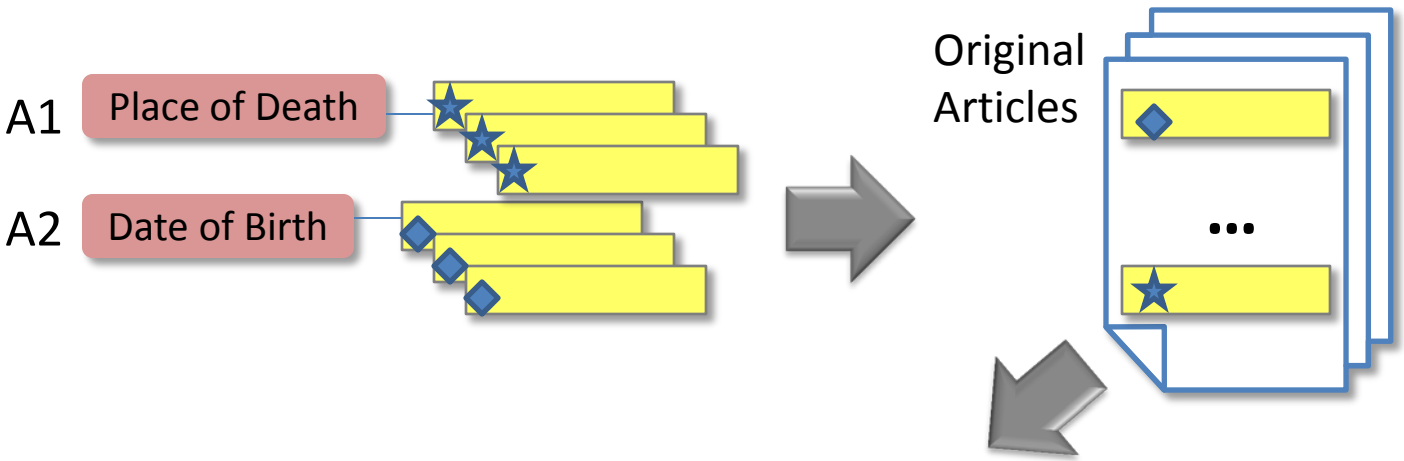

Coherence(A1, A2) ← #(★ is before ◆)

Coherence(A2, A1) ← #(◆ is before ★)

So, Coherence(A2, A1) > Coherence (A1, A2)

$$\Pi(A) = \text{argmax} \sum_{A_i \text{ before } A_j} \text{Coherence}(A_i, A_j)$$

Sample Results:Sony Cybershot DSC-W200

Freebase Aspects	sup	Representative Opinion Sentences
Format: Compact	13	Quality pictures in a compact package. …amazing is that this is such a small and compact unit but packs so much power
Supported Storage Types: Memory Stick Duo	11	This camera can use Memory Stick Pro Duo up to 8 GB Using a universal storage card and cable (c'mon Sony)
Sensor type: CCD	10	I think the larger ccd makes a difference. but remember this is a small CCD in a compact point-and-shoot.
Digital zoom: 2X	47	once the digital :smart" zoom kicks in you get another 3x of zoom. I would like a higher optical zoom, the W200 does a great digital zoom translation...

More opinion integration results are available at:

http://sifaka.cs.uiuc.edu/~yuelu2/opinionintegration/

Outline

1. Opinion Integration

2. Opinion Summarization

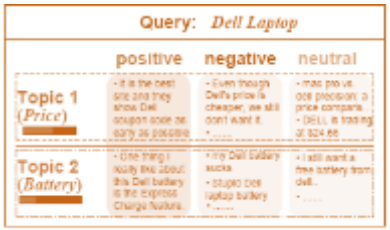

3. Opinion Analysis

Decision Making & Analytics

"Which cell phone should I **buy**?"

"What are the winning features of **iPhone** over **blackberry**?"

"How do people like this new **drug**?"

"How is Obama's health care **policy** received?"

"Which presidential candidate should I **vote** for?"

...

Need for opinion summarization

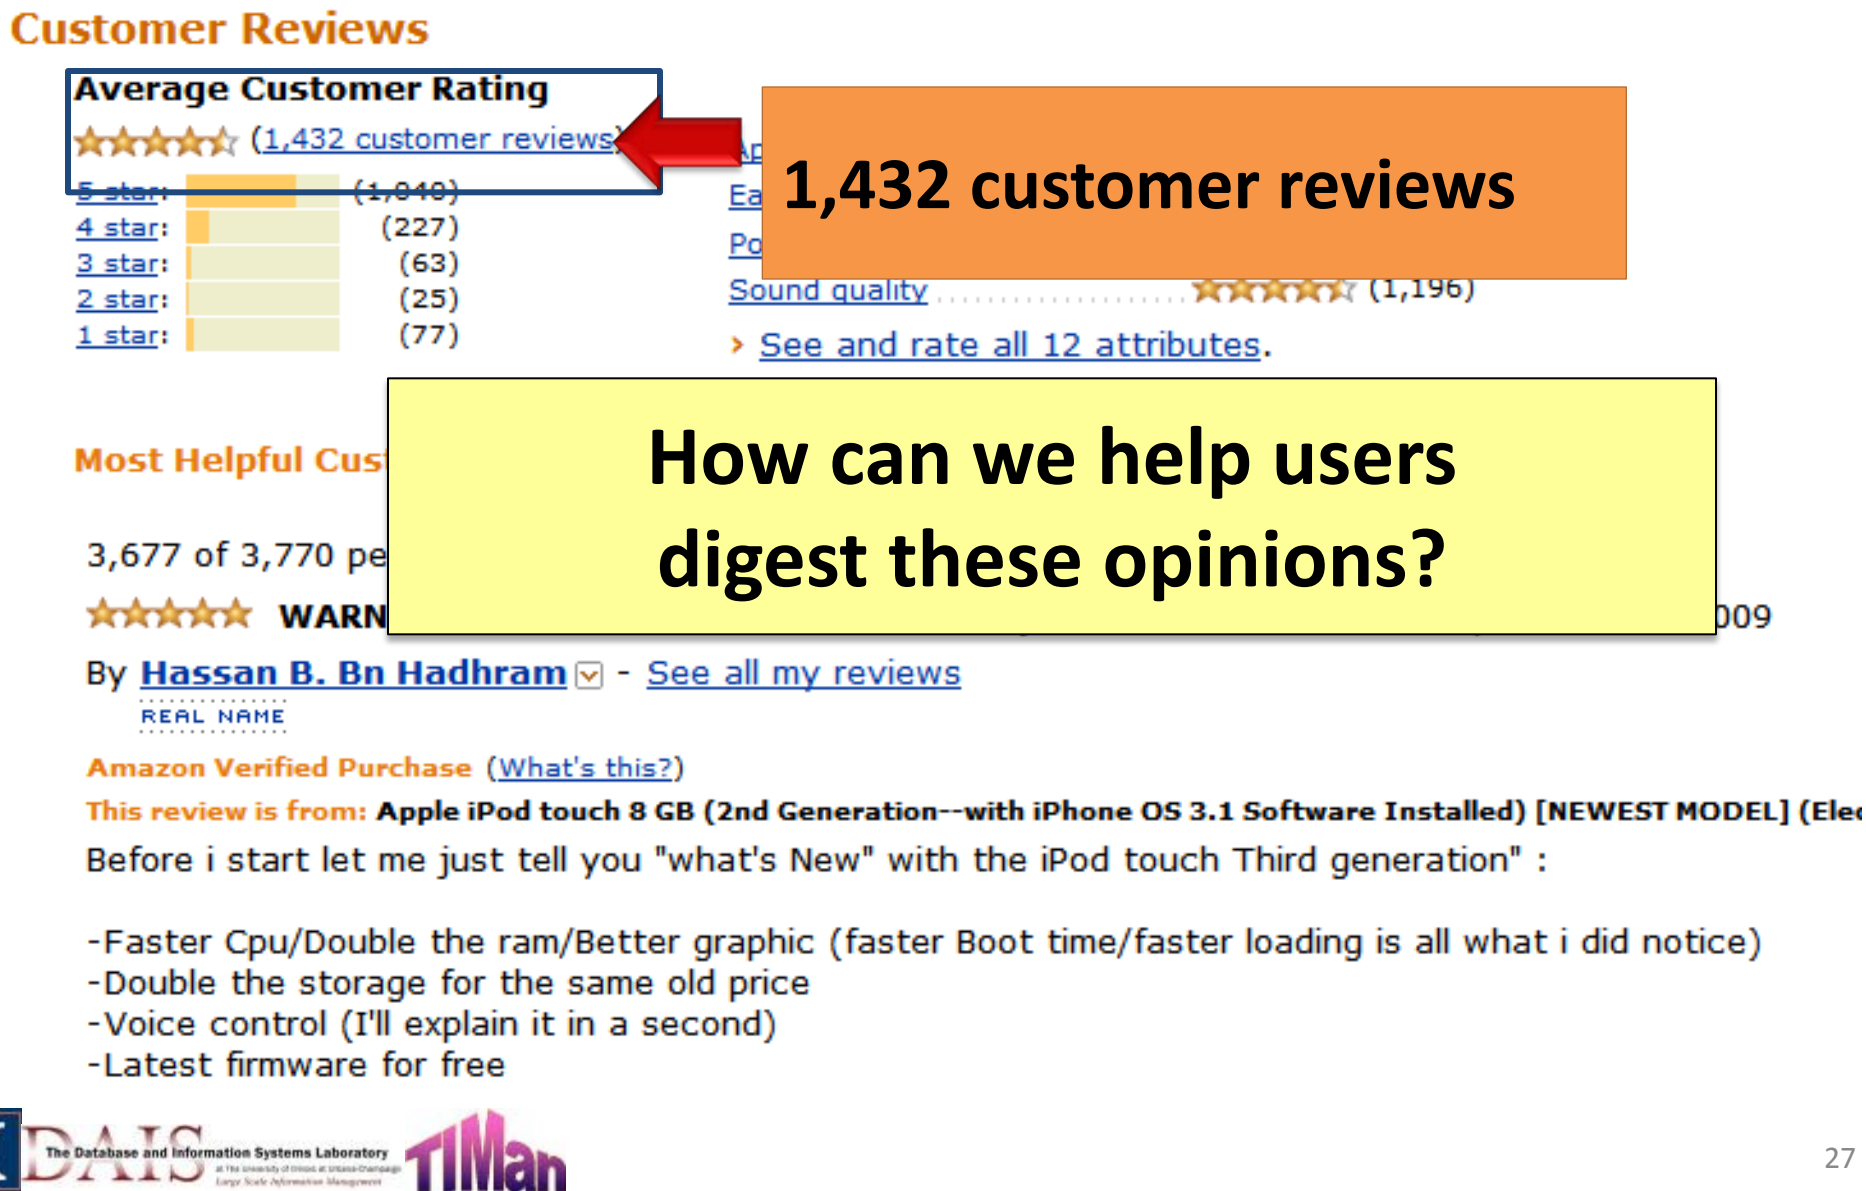

Nice to have….

Topic-sentiment summary

Query: *Dell Laptop*

	positive	negative	neutral
Topic 1 (*Price*)	• It is the best site and they show Dell coupon code as early as possible	• Even though Dell's price is cheaper, we still don't want it. • ……	• mac pro vs. dell precision: a price comparis.. • DELL is trading at $24.66
Topic 2 (*Battery*)	• One thing I really like about this Dell battery is the Express Charge feature.	• my Dell battery sucks • Stupid Dell laptop battery • ……	• I still want a free battery from dell.. • ……

Topic-sentiment dynamics (Topic = *Price*)

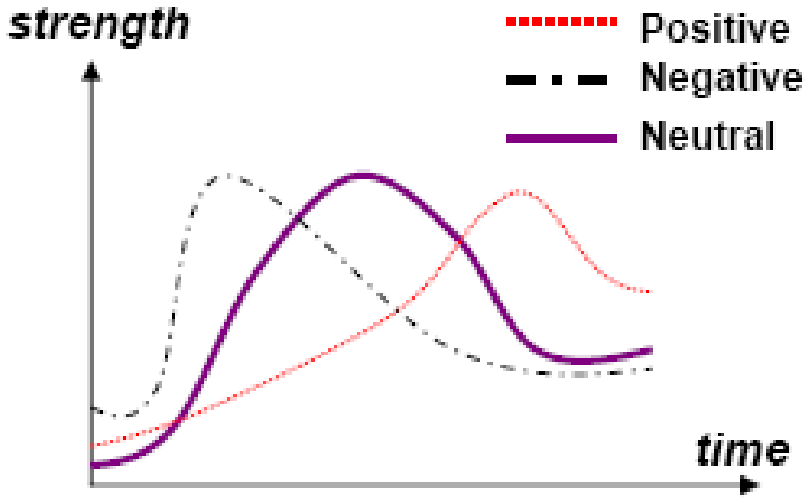

Can we do this in a general way?

Opinion Summarization 1:

[Mei et al. WWW 07]

Multi-Aspect Topic Sentiment Summarization

Qiaozhu Mei, Xu Ling, Matthew Wondra, Hang Su, ChengXiang Zhai, Topic Sentiment Mixture: Modeling Facets and Opinions in Weblogs, *Proceedings of the World Wide Conference 2007* (**WWW'07**), pages 171-180

A Topic-Sentiment Mixture Model

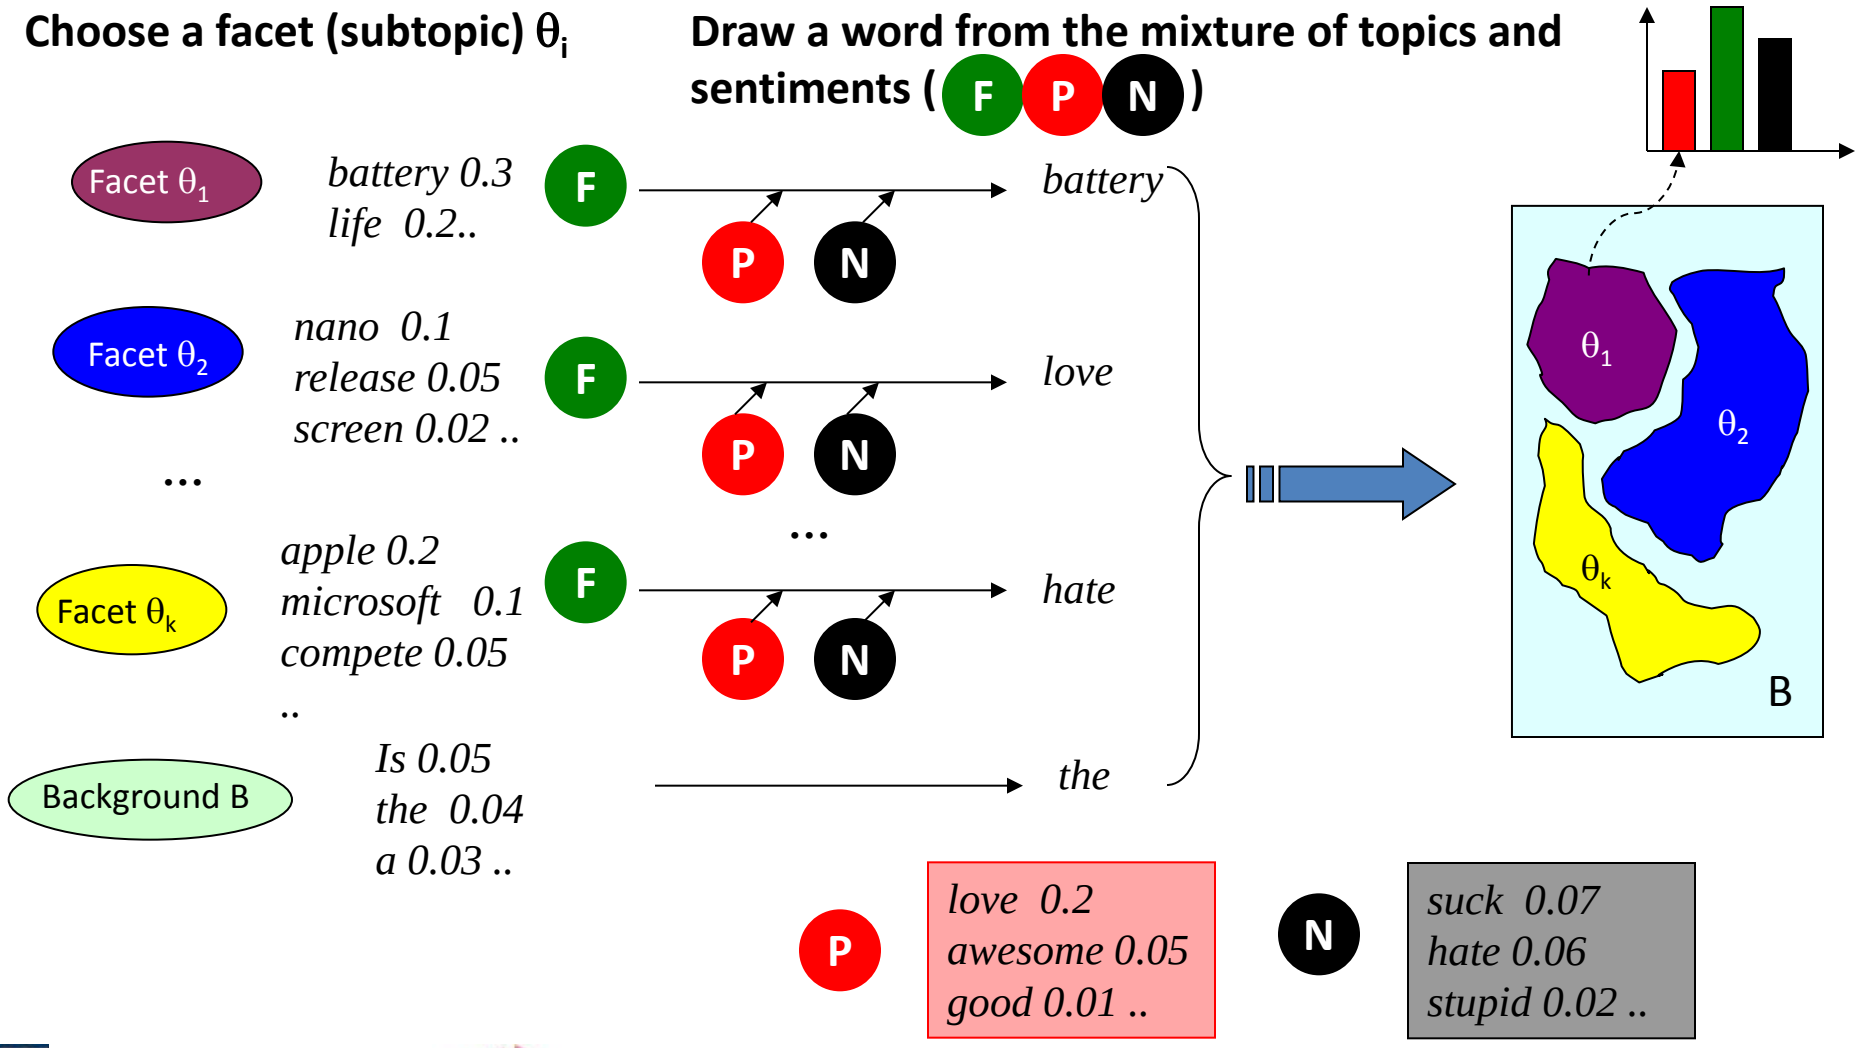

The Likelihood Function

$$\log p(C) = \sum_{d \in C} \sum_{w \in V} c(w,d) \log[\lambda_B p(w|B) + (1-\lambda_B) \sum_{j=1}^{k} \pi_{dj} \times (\delta_{j,d,F} p(w|\theta_j) + \delta_{j,d,P} p(w|\theta_P) + \delta_{j,d,N} p(w|\theta_N))]$$

Count of word w in document d

Generating w using the neutral topic model

Generating w using the background model

Generating w using the positive sentiment model

Choosing a faceted opinion

Generating w using the negative sentiment model

Two Modes for Parameter Estimation

- **Training Mode: Learn the sentiment model**

$$\log(C) = \sum_{d \in C} \sum_{w \in V} c(w,d) \log[\lambda_B p(w|B) + (1-\lambda_B) \sum_{j=1}^{k} \pi_{dj}$$
$$\times (\delta_{j,d,F}\, p(w|\theta_j) + \delta_{j,d,P}\, p(w|\theta_P) + \delta_{j,d,N}\, p(w|\theta_N))]$$

Fixed for each d

One of them is zero for d

- **Testing Mode: Extract the Topic models**

$$\log(C) = \sum_{d \in C} \sum_{w \in V} c(w,d) \log[\lambda_B p(w|B) + (1-\lambda_B) \sum_{j=1}^{k} \pi_{dj}$$
$$\times (\delta_{j,d,F}\, p(w|\theta_j) + \delta_{j,d,P}\, p(w|\theta_P) + \delta_{j,d,N}\, p(w|\theta_N))]$$

Feed strong prior on sentiment models

EM algorithm can be used for estimation

The Database and Information Systems Laboratory at The University of Illinois at Urbana-Champaign
DAIS
TIMan

Results: General Sentiment Models

- Sentiment models trained from diversified topic mixture v.s. single topics

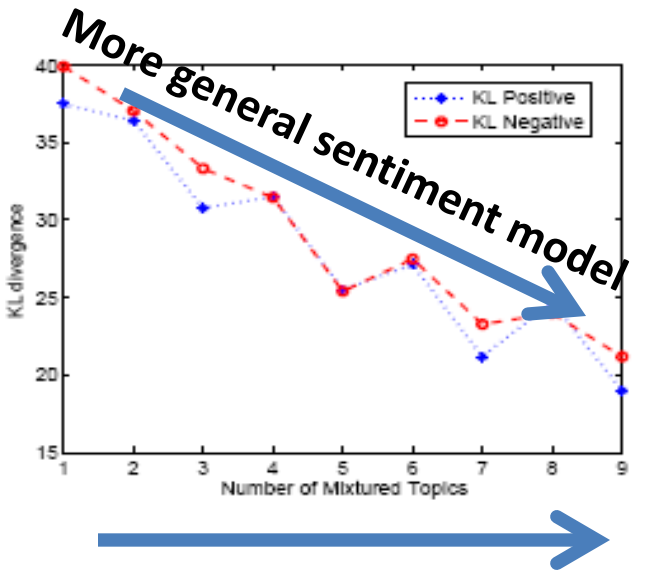

Pos-Mix	Neg-Mix	Pos-Cities	Neg-Cities
love	suck	beautiful	hate
awesome	hate	love	suck
good	stupid	awesome	people
miss	ass	amaze	traffic
amaze	fuck	live	drive
pretty	horrible	good	fuck
job	shitty	night	stink
god	crappy	nice	move
yeah	terrible	time	weather
bless	people	air	city
excellent	evil	greatest	transport

Multi-Faceted Sentiment Summary (query="Da Vinci Code")

	Neutral	Positive	Negative
Facet 1: Movie	... Ron Howards selection of Tom Hanks to play Robert Langdon.	Tom Hanks stars in the movie,who can be mad at that?	But the movie might get delayed, and even killed off if he loses.
	Directed by: Ron Howard Writing credits: Akiva Goldsman ...	Tom Hanks, who is my favorite movie star act the leading role.	protesting ... will lose your faith by ... watching the movie.
	After watching the movie I went online and some research on ...	Anybody is interested in it?	... so sick of people making such a big deal about a FICTION **book** and **movie**.
Facet 2: Book	I remembered when i first read the book, I finished the book in two days.	Awesome book.	... so sick of people making such a big deal about a FICTION **book** and **movie**.
	I'm reading "Da Vinci Code" now. …	So still a good book to past time.	This controversy book cause lots conflict in west society.

Separate Theme Sentiment Dynamics

“book”

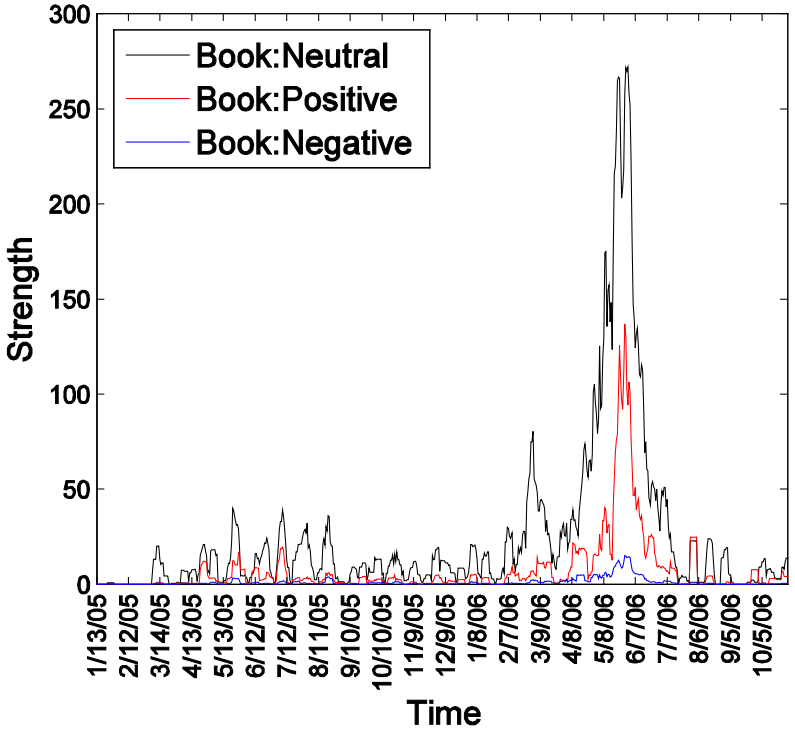

“religious beliefs”

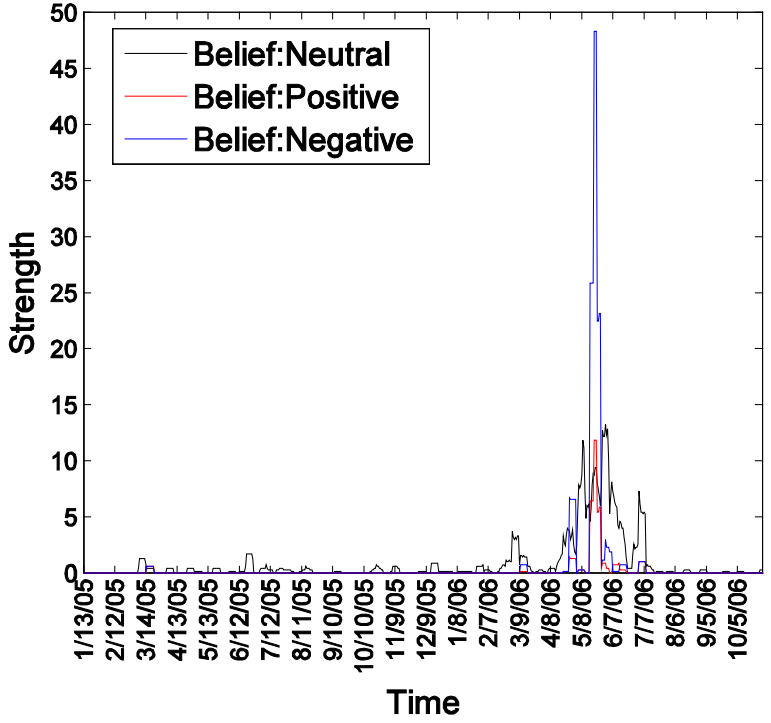

Can we make the summary more concise?

	Neutral	Positive	Negative
Facet 1: Movie	... Ron Howards selection of Tom Hanks to play Robert Langdon.	Tom Hanks stars in the movie,who can be mad at that?	But the movie might get delayed, and even killed off if he loses.
	Directed by: Ron Howard Writing credits: Akiva Goldsman ...	Tom Hanks, who is my favorite movie star act the leading role.	protesting ... will lose your faith by ... watching the movie.
	research on ...		FICTION **book** and **movie**.
Facet 2: Book	I remembered when i first read the book, I finished the book in two days.	Awesome book.	... so sick of people making such a big deal about a FICTION **book** and **movie**.
	I'm reading "Da Vinci Code" now. …	So still a good book to past time.	This controversy book cause lots conflict in west society.

What if the user is using a smart phone?

Opinion Summarization 2:

[Ganesan et al. WWW 12]

"Micro" Opinion Summarization

Kavita Ganesan, Chengxiang Zhai and Evelyne Viegas, Micropinion Generation: An Unsupervised Approach to Generating Ultra-Concise Summaries of Opinions, *Proceedings of the World Wide Conference 2012* (**WWW'12**), pages 869-878, 2012.

Micro Opinion Summarization

- **Generate a set of non-redundant phrases:**
  - Summarizing key opinions in text
  - Short (2-5 words)
  - Readable

Micropinions

Micropinion summary for a restaurant:

> **"Good service"**
> **"Delicious soup dishes"**

- **Emphasize (1) ultra-concise nature of phrases; (2) abstractive summarization**

"Room is large"
"Room is clean"

"large clean room"

A general unsupervised approach

- **Main idea:**
  - use existing words in original text to compose meaningful summaries
  - leverage Web-scale n-gram language model to assess meaningfulness

- **Emphasis on 3 desirable properties of a summary:**
  - Compactness
    - **summaries should use as few words as possible**
  - Representativeness
    - **summaries should reflect major opinions in text**
  - Readability
    - **summaries should be fairly well formed**

Optimization Framework to capture *compactness, representativeness & readability*

$$M = \arg\max_{\{m1 \ldots mk\}} \sum_{i=1}^{k} S_{rep}(m_i) + S_{read}(m_i)$$

Micropinion Summary, M

2.3	very clean rooms
2.1	friendly service
1.8	dirty lobby and pool
1.3	nice and polite staff

subject to

$$\sum_{i=1}^{k} |m_i| \leq \sigma_{ss}$$ **Size of summary**

$$S_{rep}(m_i) \geq \sigma_{rep}$$

$$S_{read}(m_i) \geq \sigma_{read}$$ **Minimum rep. & readability**

$$sim(m_i, m_j) \leq \sigma_{sim} \forall i, j \in [1, k]$$ **Redundancy**

Representativeness scoring: Srep(mi)

- **2 properties of a highly representative phrase:**
  - Words should be strongly associated in text
  - Words should be sufficiently frequent in text
- **Captured by modified pointwise mutual information**

$$pmi'(w_i, w_j) = \log_2 \frac{p(w_i, w_j) \times c(w_i, w_j)}{p(w_i) \times p(w_j)}$$

Add frequency of occurrence within a window

$$pmi_{local}(w_i) = [\frac{1}{2C} \sum_{j=i-C}^{i+C} pmi'(w_i, w_j)]$$

$$S_{rep}(w_1..w_n) = \frac{1}{n} \sum_{i=1}^{n} pmi_{local}(w_i)$$

Readability scoring, Sread(mi)

- **Phrases are constructed from seed words, thus we can have new phrases not in original text**
- **Readability scoring based on N-gram language model (normalized probabilities of phrases)**
  - **Intuition:** A phrase is more readable if it occurs more frequently on the web

$$S_{read}(w_k \ldots w_n) = \frac{1}{K} \log_2 \prod_{k=q}^{n} p(w_k \mid w_{k-q+1} \ldots w_{k-1})$$

Ungrammatical		Grammatical	
"sucks life battery"	**-4.51**	*"battery life sucks"*	**-2.93**
"life battery is poor"	**-3.66**	*"battery life is poor"*	**-2.37**

Overview of summarization algorithm

Input

Text to be summarized

Unigrams

….
very
nice
place
clean
problem
dirty
room …

Step 1: Shortlist high freq unigrams **(count > median)**

Seed Bigrams

very + nice
very + clean
very + dirty
clean + place
clean + room
dirty + place …

Srep > σ_{rep}

Step 2: Form seed bigrams by pairing unigrams. Shortlist by $S_{rep.}$ **($S_{rep} > \sigma_{rep}$)**

Overview of summarization algorithm

Higher order n-grams

Candidates	+	Seed Bi-grams	=	New Candidates
very clean	+	clean rooms clean bed	=	very clean rooms very clean bed
very dirty	+	dirty room dirty pool	=	very dirty room very dirty pool
very nice	+	nice place nice room	=	very nice place very nice room

$S_{rep} < \sigma_{rep}$; $S_{read} < \sigma_{read}$

Summary

0.9 very clean rooms
0.8 friendly service
0.7 dirty lobby and pool
0.5 nice and polite staff
.....
.....

Sorted Candidates

Step 3: Generate higher order n-grams.

- Concatenate existing candidates + seed bigrams
- Prune non-promising candidates (S_{rep} & S_{read})
- Eliminate redundancies (sim(mi,mj))
- Repeat process on shortlisted candidates (until no possbility of expansion)

Step 4: Final summary.

Sort by objective function value. Add phrases until $|M| < \sigma_{ss}$

Performance comparisons (reviews of 330 products)

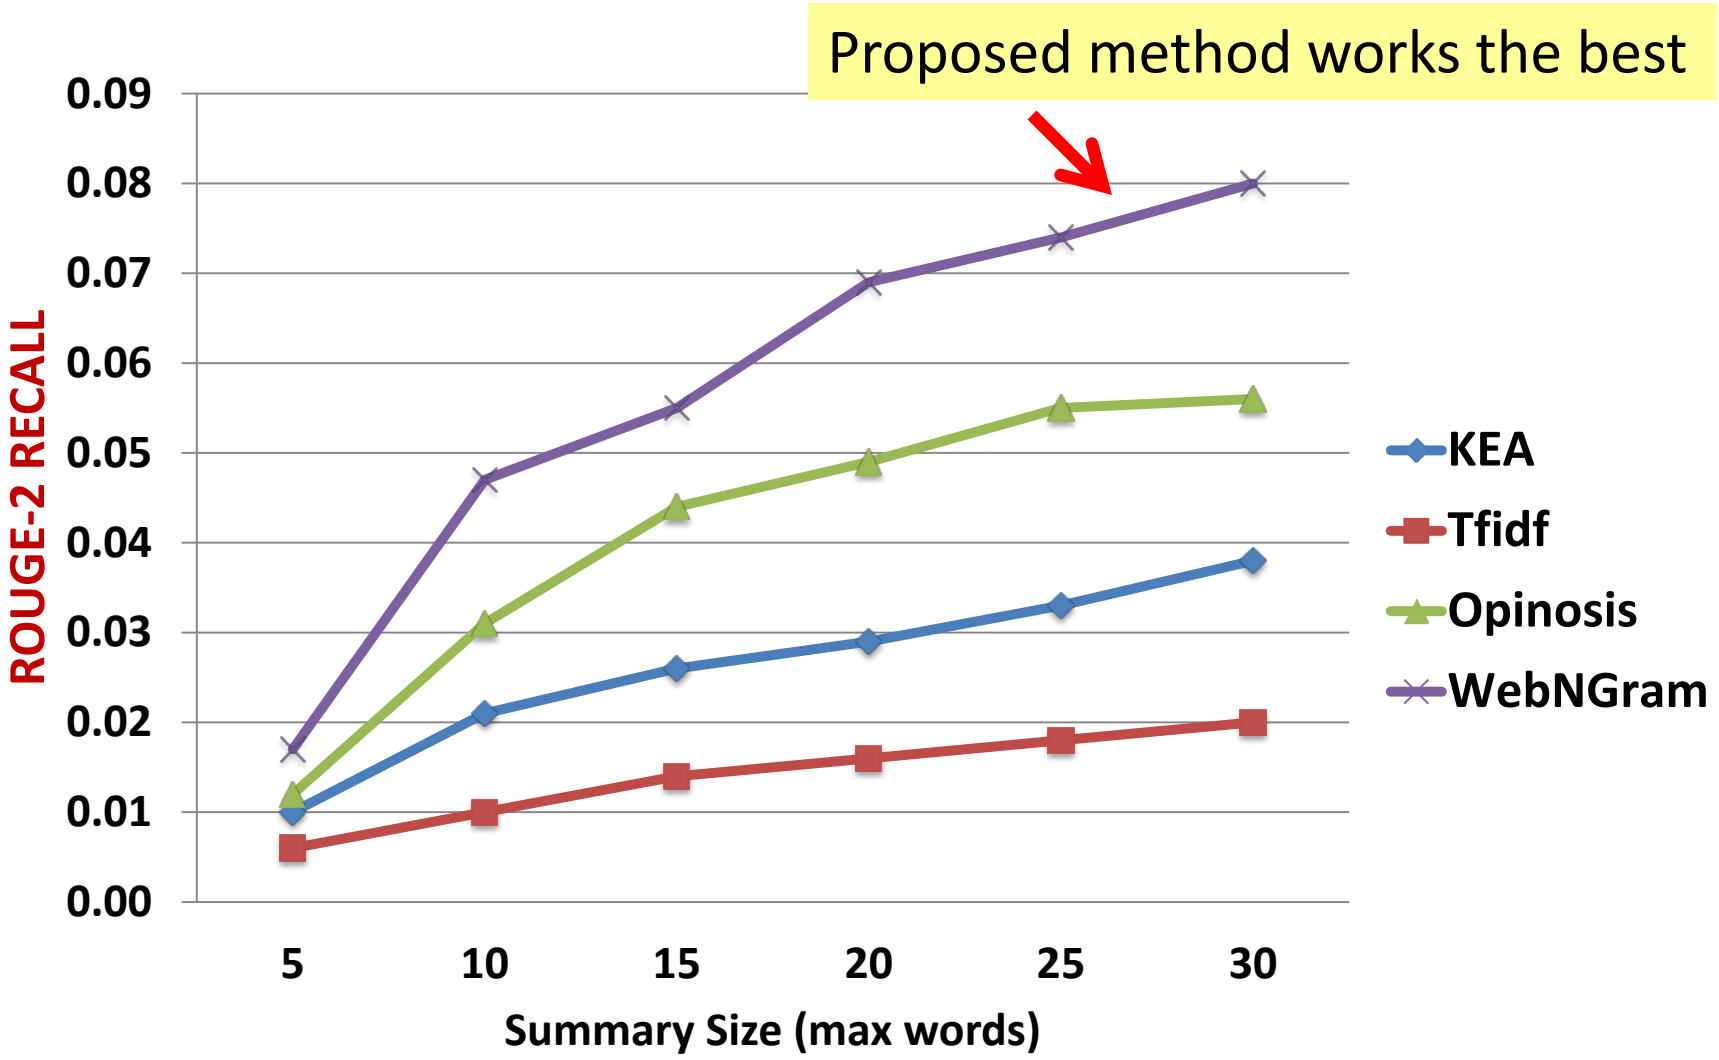

The program can generate meaningful novel phrases

Example:

Unseen N-Gram (Acer AL2216 Monitor)

"wide screen lcd monitor is bright"

readability : -1.88

representativeness: 4.25

"…plus the monitor is very bright…"

"…it is a wide screen, great color, great quality…"

"…this lcd monitor is quite bright and clear…"

Related snippets in original text

A Sample Summary

Canon Powershot SX120 IS

Easy to use
Good picture quality
Crisp and clear
Good video quality

Useful for pushing opinions to devices where the **screen is small**

E-reader/ Tablet

Smart Phones

Cell Phones

Outline

1. Opinion Integration

2. Opinion Summarization

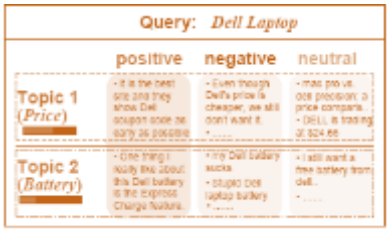

3. Opinion Analysis

Decision Making & Analytics

"Which cell phone should I **buy**?"

"What are the winning features of **iPhone** over **blackberry**?"

"How do people like this new **drug**?"

"How is Obama's health care **policy** received?"

"Which presidential candidate should I **vote** for?"

...

Motivation

Hotel Palomar Chicago: Traveler Reviews

"Great location+spacious room =happy traveler"

Save Review

leos_10 3 contributions
Boston

Jul 11, 2010 | Trip type: Couples NEW

Stayed for a weekend in July. Walked everywhere, enjoyed the comfy bed and quiet hallways. more

"terrific service and gorgeous facility"

Save Review

ahickling 1 contribution
Greensboro, North Carolina

Jul 7, 2010 | Trip type: Family NEW

I stayed at the Palomar with my young daughter for three nights June 17-20, 2010 and absolutely loved the hotel. The room was one of the nicest I've ever stayed in (My daughter loved the Fuji jetted tub so much that she wanted to take 2 baths a day!) in terms of decor, design, and size. (It compared favorably to... more

How to infer aspect ratings?

My ratings for this hotel

- Value
- Rooms
- Location
- Cleanliness
- Service
- Sleep Quality

My ratings for this hotel

- Value
- Rooms
- Location
- Cleanliness
- Service
- Sleep Quality

How to infer aspect weights?

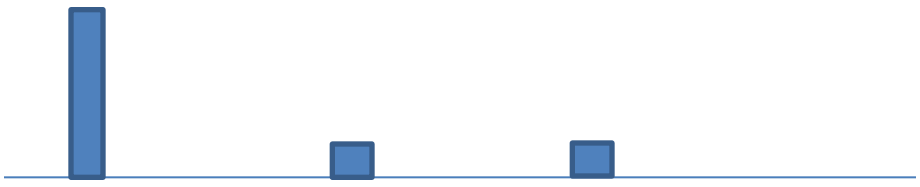

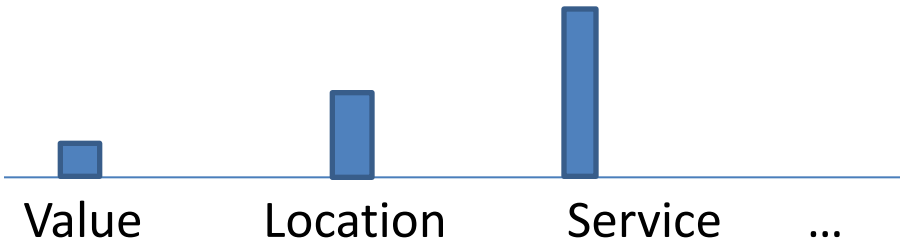

Opinion Analysis:

[Wang et al. KDD 2010] & [Wang et al. KDD 2011]

Latent Aspect Rating Analysis

Hongning Wang, Yue Lu, ChengXiang Zhai. Latent Aspect Rating Analysis on Review Text Data: A Rating Regression Approach, *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (**KDD'10**), pages 115-124, 2010.

Hongning Wang, Yue Lu, ChengXiang Zhai, Latent Aspect Rating Analysis without Aspect Keyword Supervision, *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (**KDD'11**), 2011, pages 618-626.

Latent Aspect Rating Analysis

- **Given a set of review articles about a topic with overall ratings**
- **Output**
  - Major aspects commented on in the reviews
  - Ratings on each aspect
  - Relative weights placed on different aspects by reviewers
- **Many applications**
  - Opinion-based entity ranking
  - Aspect-level opinion summarization
  - Reviewer preference analysis
  - Personalized recommendation of products
  - …

Solving LARA in two stages: Aspect Segmentation + Rating Regression

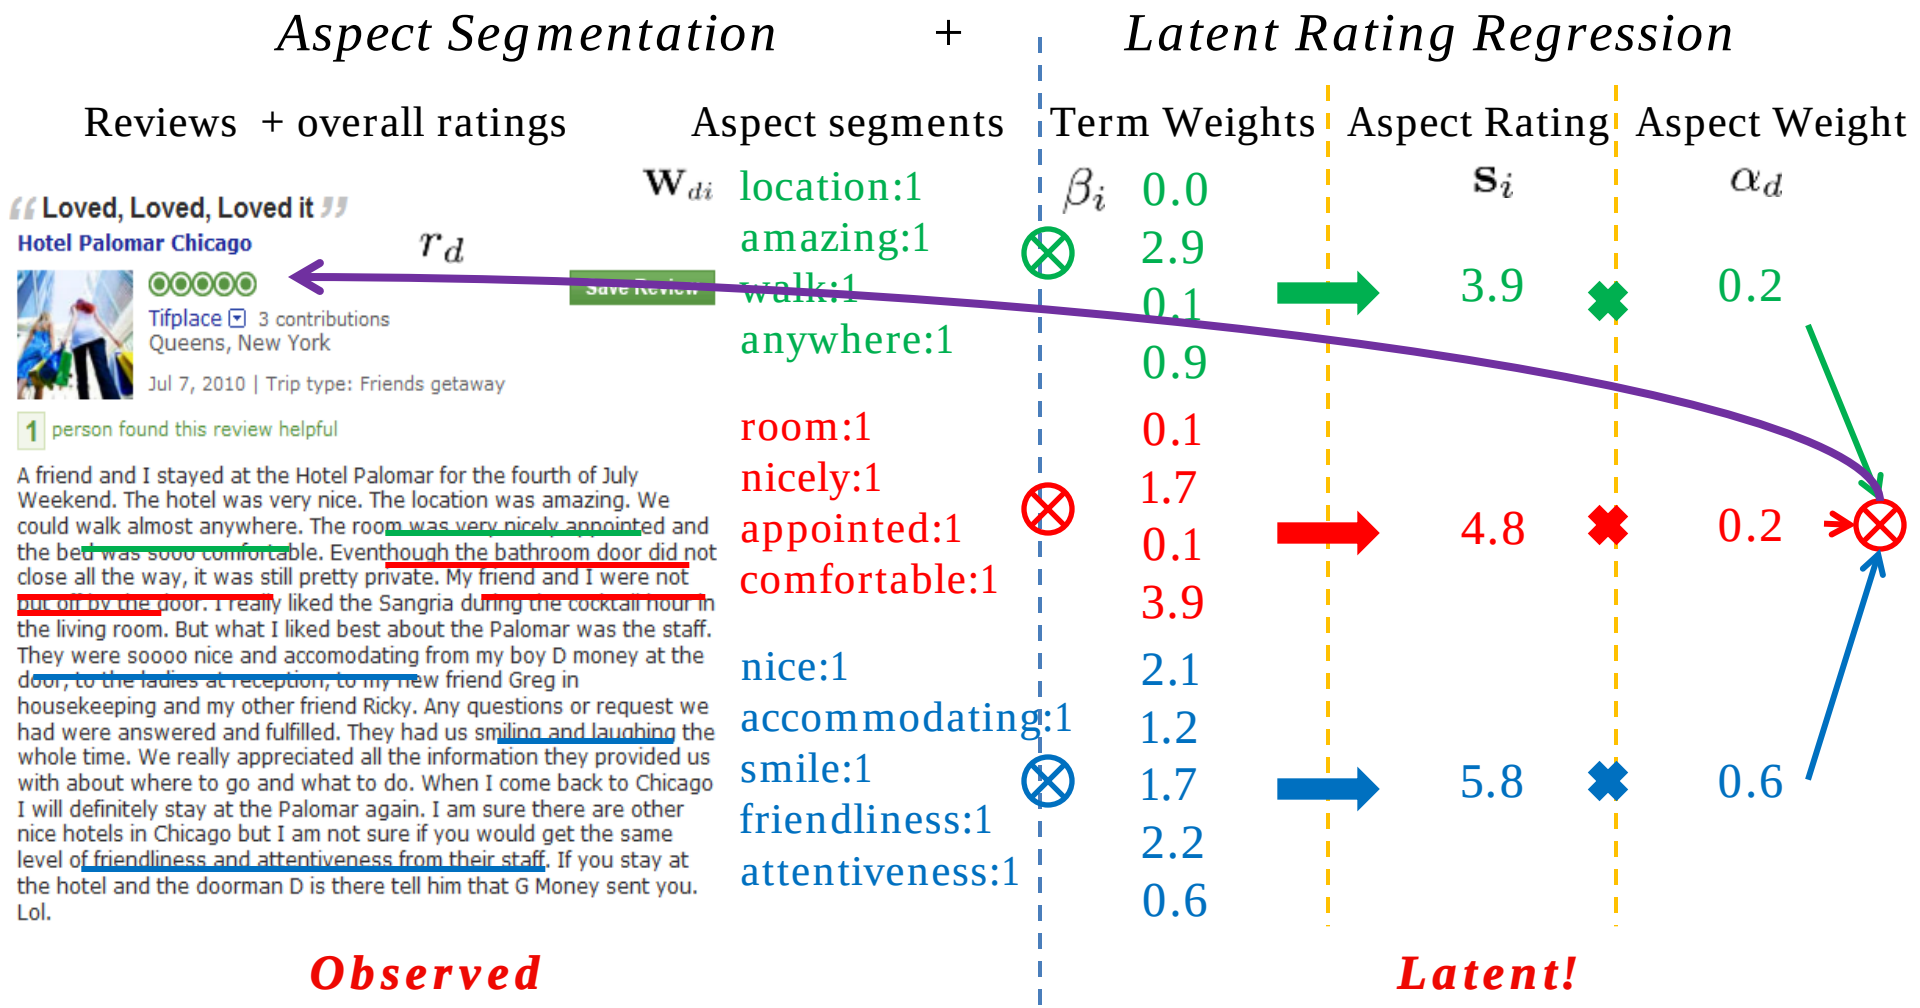

Latent Rating Regression

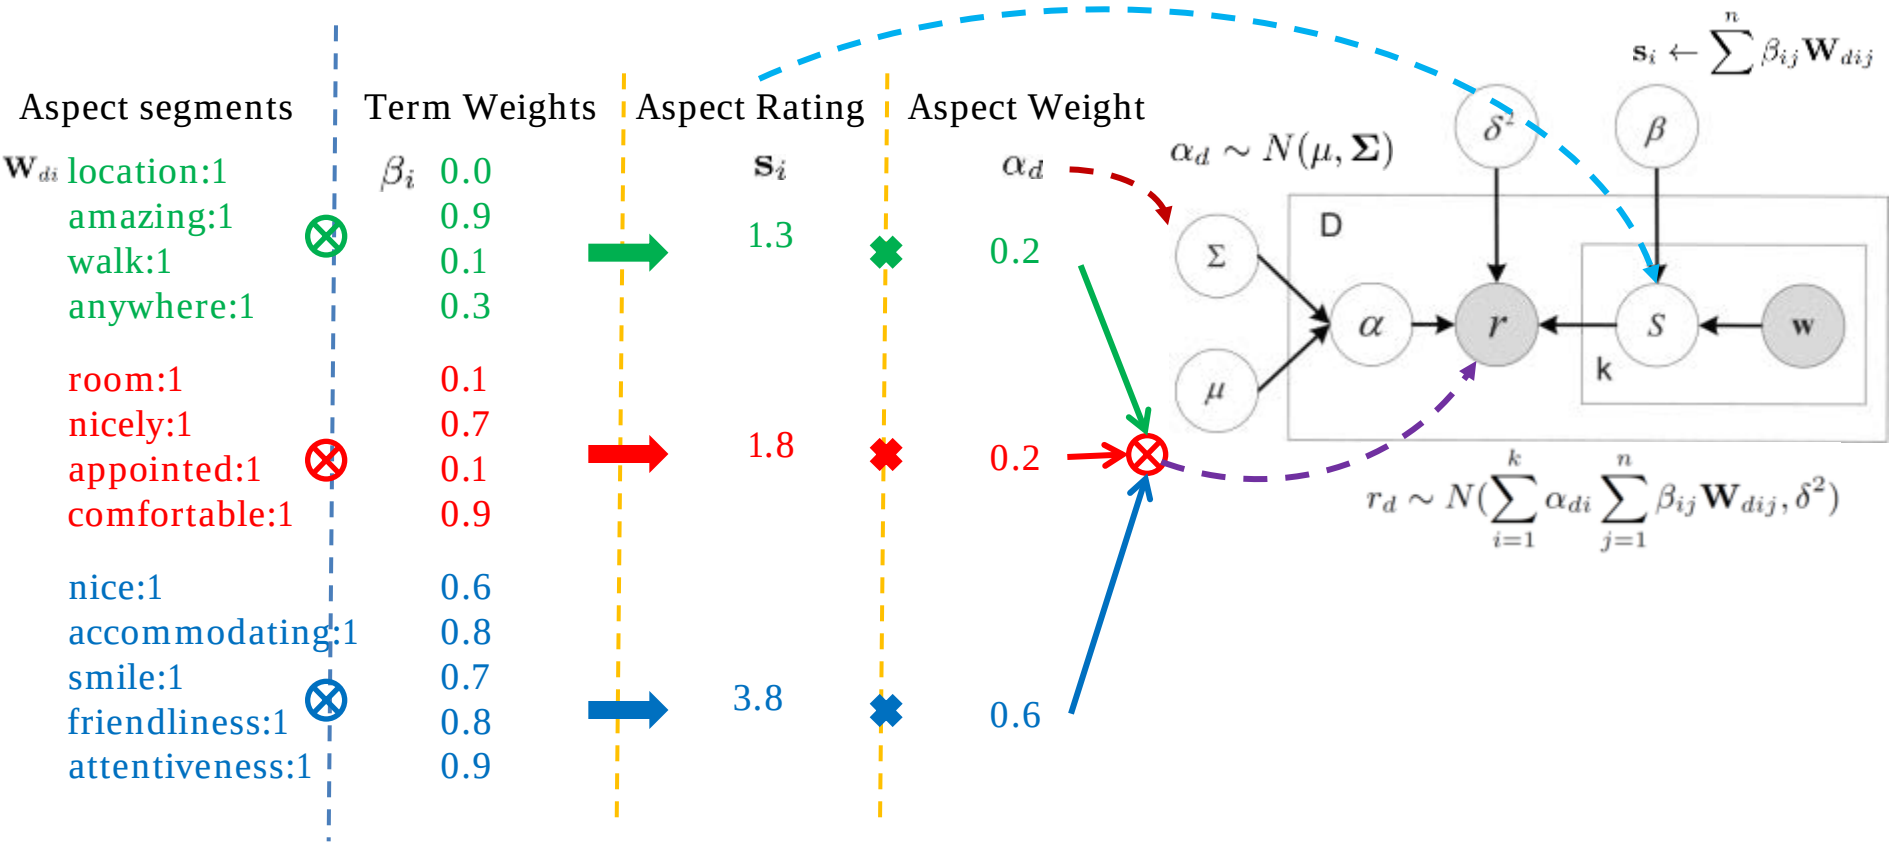

Conditional likelihood

$$P(r|d) = P(r_d|\mu, \mathbf{\Sigma}, \delta^2, \beta, \mathbf{W}_d)$$

$$= \int p(\alpha_d|\mu, \mathbf{\Sigma}) p(r_d|\sum_{i=1}^{k} \alpha_{di} \sum_{j=1}^{n} \beta_{dij} \mathbf{W}_{dij}, \delta^2) \mathrm{d}\alpha_d$$

A Unified Generative Model for LARA

Entity

Aspects

Review

Aspect Rating

Aspect Weight

Location

location
amazing
walk
anywhere

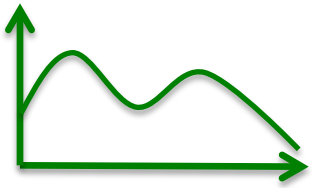

Room

room
dirty
appointed
smelly

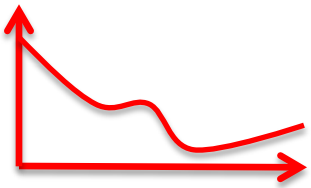

Service

terrible
front-desk
smile
unhelpful

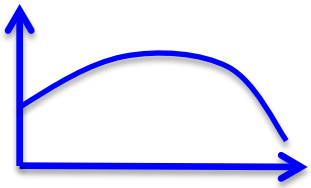

"Spend your money elsewhere"

Reviewed September 19, 2010

Excellent location in walking distance to Tiananmen Square and shopping streets. That's the best part of this hotel! The rooms are getting really old. Bathroom was nasty. The fixtures were falling off, lots of cracks and everything looked dirty. I don't think it worth the price. Service was the most disappointing part, especially the door men. this is not how you treat guests, this is not hospitality.

0.86

0.04

0.10

Latent Aspect Rating Analysis Model

- **Unified framework**

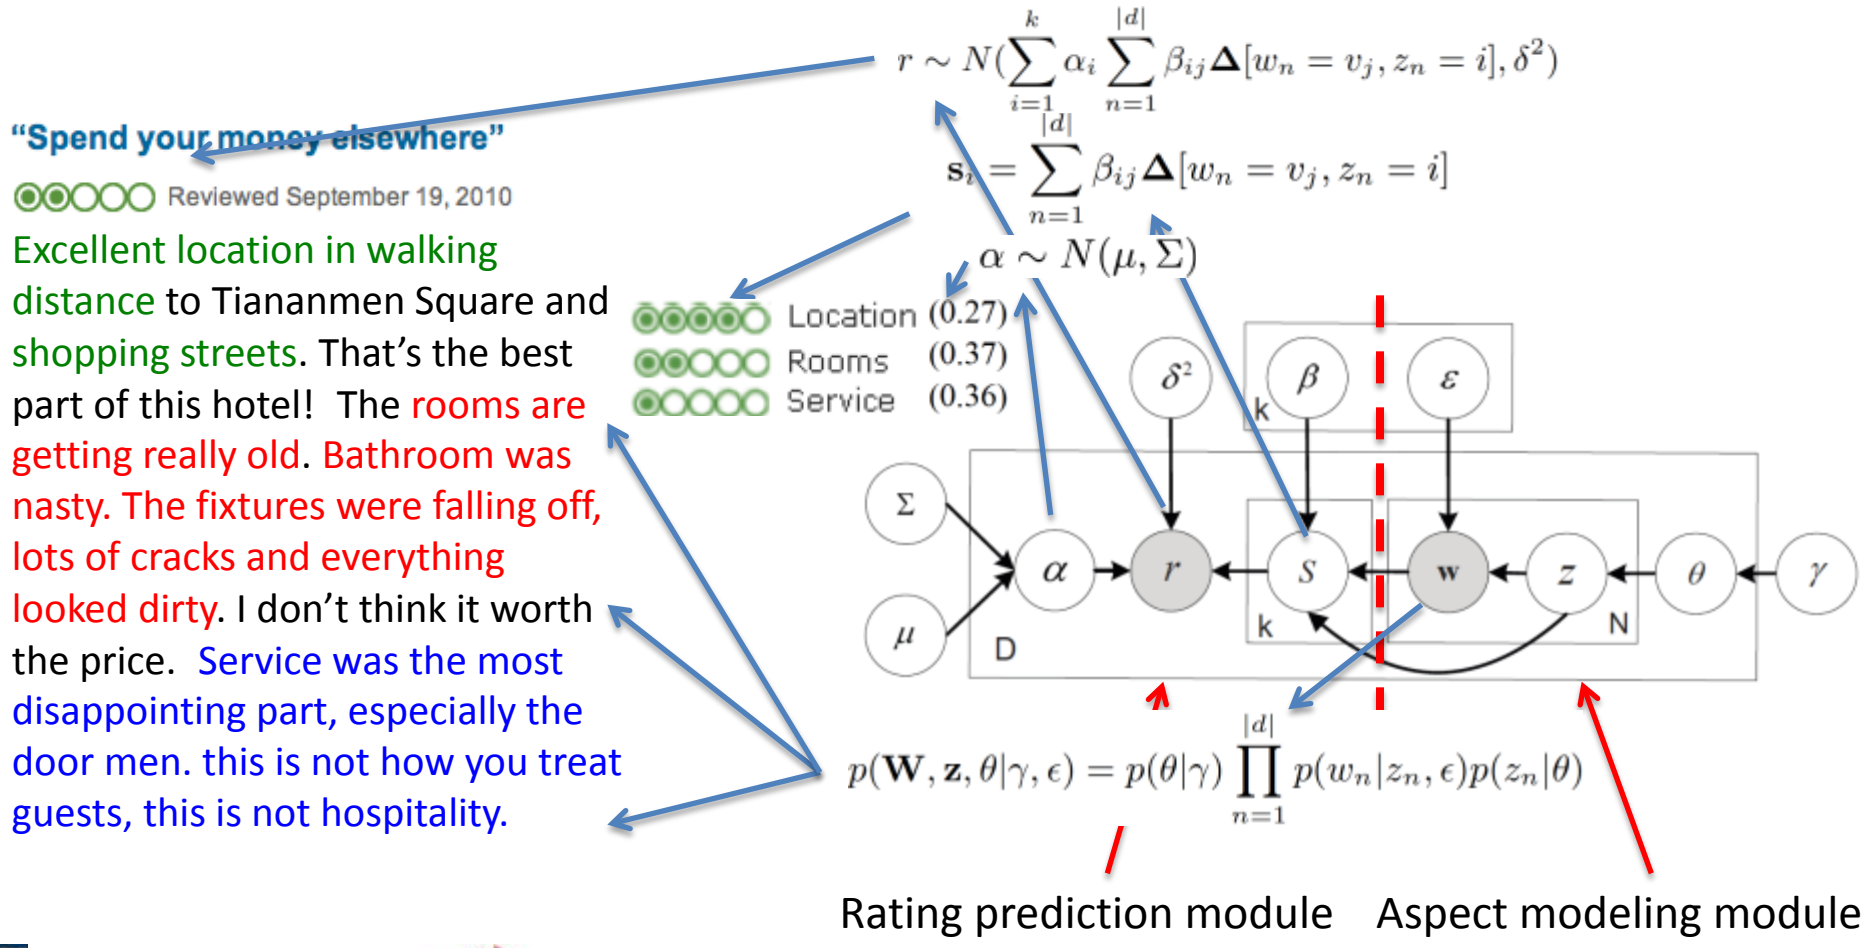

Sample Result 1: Rating Decomposition

- **Hotels with the same overall rating but different aspect ratings**

(All 5 Stars hotels, ground-truth in parenthesis.)

Hotel	*Value*	*Room*	*Location*	*Cleanliness*
Grand Mirage Resort	4.2(4.7)	3.8(3.1)	4.0(4.2)	4.1(4.2)
Gold Coast Hotel	4.3(4.0)	3.9(3.3)	3.7(3.1)	4.2(4.7)
Eurostars Grand Marina Hotel	3.7(3.8)	4.4(3.8)	4.1(4.9)	4.5(4.8)

- **Reveal detailed opinions at the aspect level**

Sample Result 2: Comparison of reviewers

- **Reviewer-level Hotel A**
  - Different reviewers' ratinç
  - Reveal differences in opi

Reviewer	*Value*	
Mr.Saturday	3.7(4.0)	
Salsrug	5.0(5.0)	

"Good Price for what we got"

Riu Palace Punta Cana

Save Review

salsrug 13 contributions
Marylander

Oct 27, 2008 | Trip type: Family

1 person found this review helpful

We stayed for six days, five nights. Overall, we had a very good time. The was pretty good and the staff was very friendly. They definitely do not skimp on the free alcohol. The room was a little smelly, which we had read on trip advisor so we bought a candle with us - no problem. They only thing I had an issue with was the little bugs. They were like gnats or fleas but they weren't either. I had some candy and popcorn which we brought from the States to munch on. I left it out on the table and within 40 minutes, the bag was infested. DO NOT keep any open food in your room. Also we ended up having to wash all of our clothes (clean and dirty) and airing out our luggage when we got home because we could still smell the room on them. For the price we paid, we really did have an excellent time besides those small things. The pool was awesome and the beach was spectacular. Out of the nearby resorts that we saw, Riu Palace Punta Cana was the best (it was also the nicest out of the other Riu's on Punta Cana). We went on the 1/2 day Outback Safari and had a great time. We got coffee and souvenirs cheapers than other places and the hotel. General - not good or bad just things that we noticed - There were a lot of topless sunbathers. The crowd is middle aged (35 - 55) so we were on the younger side and the majority of the people were European or Brazilian. It helps to know some spanish but it's not a necessity.

Liked — The beach was excellent.

Disliked — Room smell and little bugs.

My ratings for this hotel

Value
Rooms
Location
Cleanliness
Check in / front desk
Service
Business service (e.g., internet access)

Sample Result 3:Aspect-Specific Sentiment Lexicon

Value	*Rooms*	*Location*	*Cleanliness*
resort 22.80	view 28.05	restaurant 24.47	clean 55.35
value 19.64	comfortable 23.15	walk 18.89	smell 14.38
excellent 19.54	modern 15.82	bus 14.32	linen 14.25
worth 19.20	quiet 15.37	beach 14.11	maintain 13.51
bad -24.09	*carpet -9.88*	*wall -11.70*	*smelly -0.53*
money -11.02	*smell -8.83*	*bad -5.40*	*urine -0.43*
terrible -10.01	*dirty -7.85*	*road -2.90*	*filthy -0.42*
overprice -9.06	*stain -5.85*	*website -1.67*	*dingy -0.38*

Uncover sentimental information directly from the data

Sample Result 4: Validating preference weights

- **Analysis of hotels preferred by different types of reviewers**

City	AvgPrice	Group	Val/Loc	Val/Rm	Val/Ser
Amsterdam	241.6	top-10	190.7	214.9	221.1
		bot-10	270.8	333.9	236.2
Barcelona	280.8	top-10	270.2	196.9	263.4
		bot-10	330.7	266.0	203.0
San Francisco	261.3	top-10	214.5	249.0	225.3
		bot-10	321.1	311.1	311.4
Florence	272.1	top-10	269.4	248.9	220.3
		bot-10	298.9	293.4	292.6

  - Reviewers emphasizing the ‘value’ aspect more would prefer cheaper hotels

Application 1: Rated Aspect Summarization

Aspect	*Summary*	*Rating*
Value	Truly unique character and a great location at a reasonable price Hotel Max was an excellent choice for our recent three night stay in Seattle.	3.1
	Overall not a negative experience, however considering that the hotel industry is very much in the impressing business there was a lot of room for improvement.	1.7
Location	The location, a short walk to downtown and Pike Place market, made the hotel a good choice.	3.7
	When you visit a big metropolitan city, be prepared to hear a little traffic outside!	1.2
Business Service	You can pay for wireless by the day or use the complimentary Internet in the business center behind the lobby though.	2.7
	My only complaint is the daily charge for internet access when you can pretty much connect to wireless on the streets anymore.	0.9

(Hotel Max in Seattle)

Application 2: Discover consumer preferences

- **Amazon reviews: no guidance**

Table 2: Topical Aspects Learned on MP3 Reviews

Low Overall Ratings			High Overall Ratings		
unit	jack	service	files	player	vision
usb	headphone	charge	format	music	video
battery	warranty	problem	included	download	player
charger	replacement	support	easy	headphones	quality
reset	problem	hours	convert	button	great
time	player	months	mp3	set	product
hours	back	weeks	videos	hours	sound
work	months	back	file	buds	radio
thing	buy	customer	wall	volume	accessory
wall	amazon	time	hours	ear	fm

battery life accessory service file format volume video

Application 3: User Rating Behavior Analysis

	Expensive Hotel		*Cheap Hotel*	
	5 Stars	*3 Stars*	*5 Stars*	*1 Star*
Value	0.134	0.148	**0.171**	0.093
Room	0.098	0.162	0.126	0.121
Location	0.171	0.074	0.161	0.082
Cleanliness	0.081	**0.163**	0.116	**0.294**
Service	**0.251**	0.101	0.101	0.049

People like expensive hotels because of good service

People like cheap hotels because of good value

Application 4: Personalized Ranking of Entities

Query: 0.9 value
0.1 others

Non-Personalized

Personalized

Table 10: Personalized Hotel Ranking

Hotel	Overall Rating	Price	Location
Majestic Colonial	5.0	339	Punta Cana
Agua Resort	5.0	753	Punta Cana
Majestic Elegance	5.0	537	Punta Cana
Grand Palladium	5.0	277	Punta Cana
Iberostar	5.0	157	Punta Cana
Elan Hotel Modern	5.0	216	Los Angeles
Marriott San Juan Resort	4.0	354	San Juan
Punta Cana Club	5.0	409	Punta Cana
Comfort Inn	5.0	155	Boston
Hotel Commonwealth	4.5	313	Boston

Summary

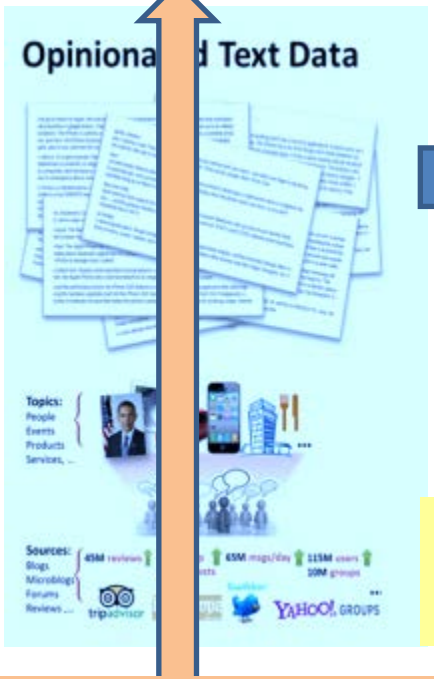

Rapidly growing opinionated text data open up many applications

1. Opinion Integration

- Leverage expert reviews [WWW 08]
- Leverage ontology [COLING 10]

2. Opinion Summarization

- Aspect sentiment summary [WWW 07]
- Micro opinion summary [WWW 12]

3. Opinion Analysis

- Two-stage rating analysis [KDD 10]
- Unified rating analysis [KDD 11]

Decision Making & Analytics

"Which cell phone should I buy?"

Users face significant challenges in collecting and digesting opinions

"How is Obama's health care policy received?"

"Which presidential candidate should I vote for?"

...

applicable to any topic, any natural language, with no/minimum human effort

Open Questions

- **How can we combine all these methods in a general unified decision-support system?**
  - What are the basic common functions required by all applications?
  - How do we support users to interact with the system?
- **How far can we go with such pure statistical approaches?**
  - How can we maximize the benefit of unsupervised learning? Continuous learning from the Web?
  - How can we combine unsupervised learning naturally with supervised learning through user interactions?
- **How can we incorporate linguistic resources & knowledge?**
  - How can we build a sentiment analyzer to take advantage all the resources available today?
  - Can we automatically construct sophisticated features for sentiment analysis? Deep learning?

Demo: FindILike System

1. Opinion Integration

2. Opinion Summarization

www.findilike.com

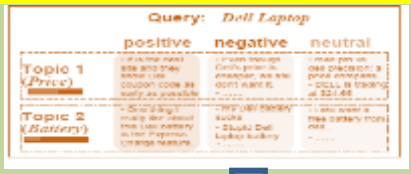

3. Opinion Analysis

Decision Making & Analytics

"Which cell phone should I buy?"

res of

v drug?"

e policy

"Which presidential candidate should I vote for?"

...

Findilike: Opinion-Based Decision-Support

www.findilike.com

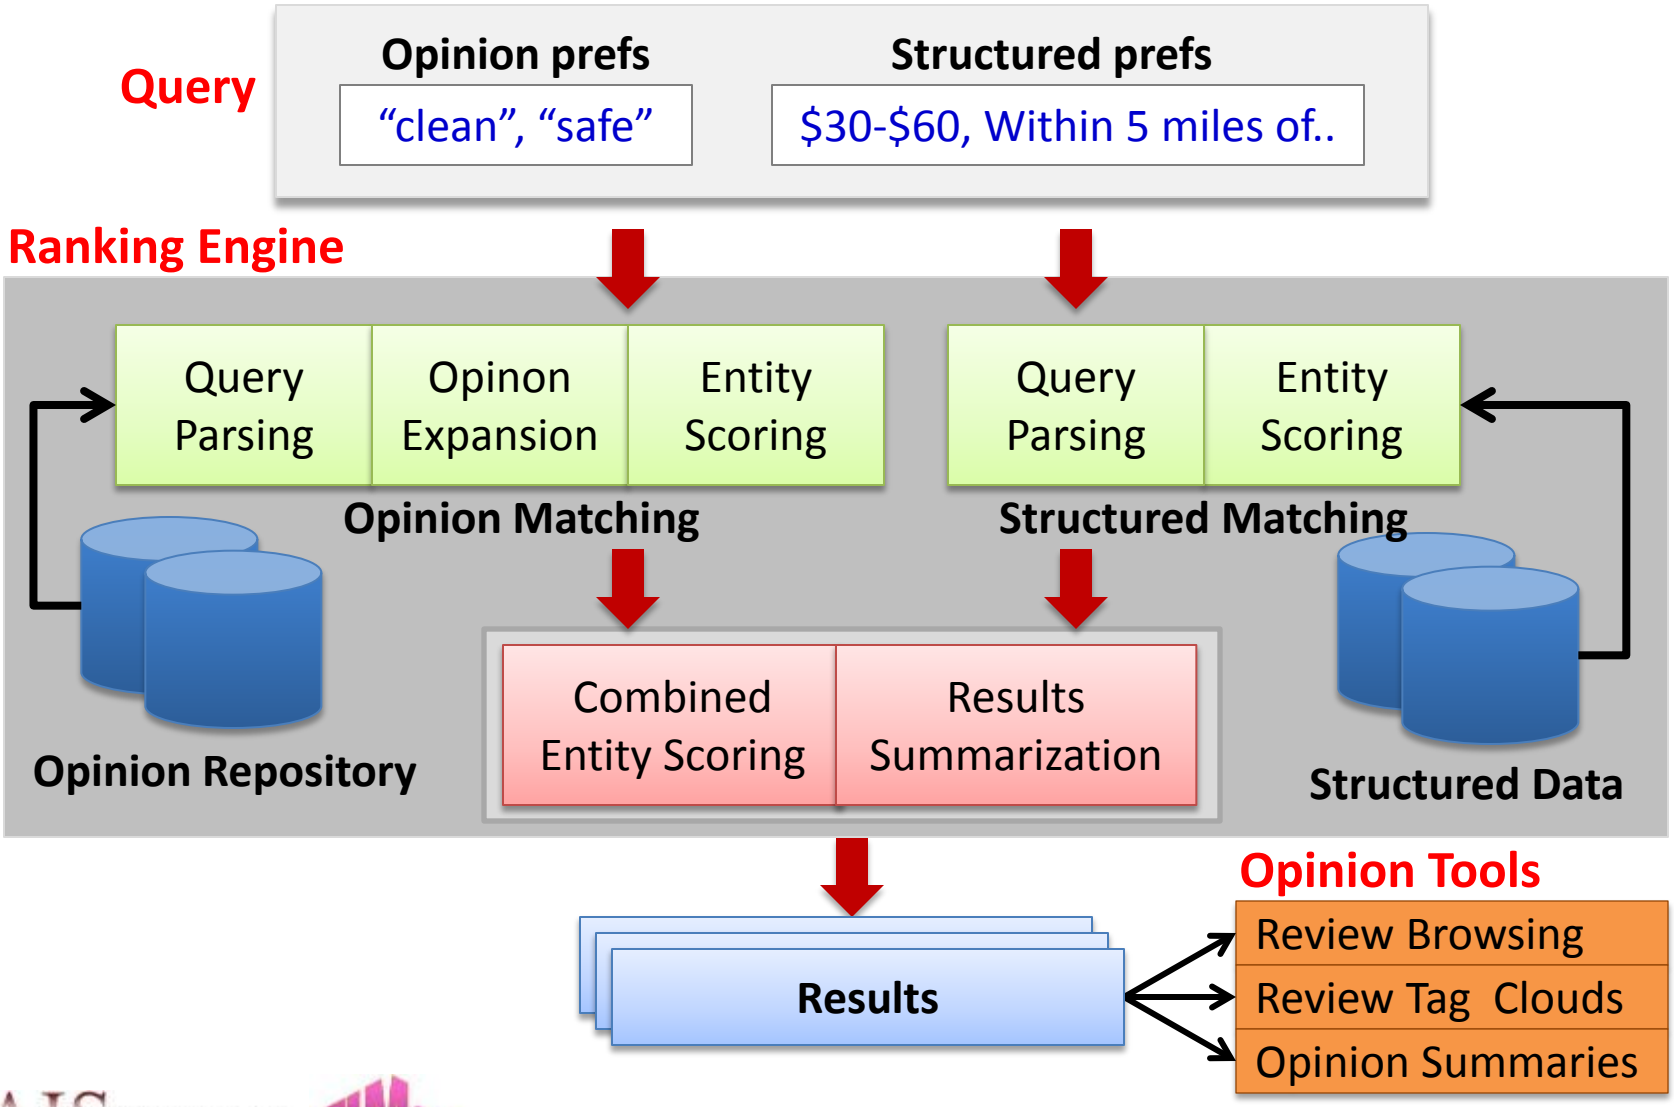

Opinion-Based Entity Ranking

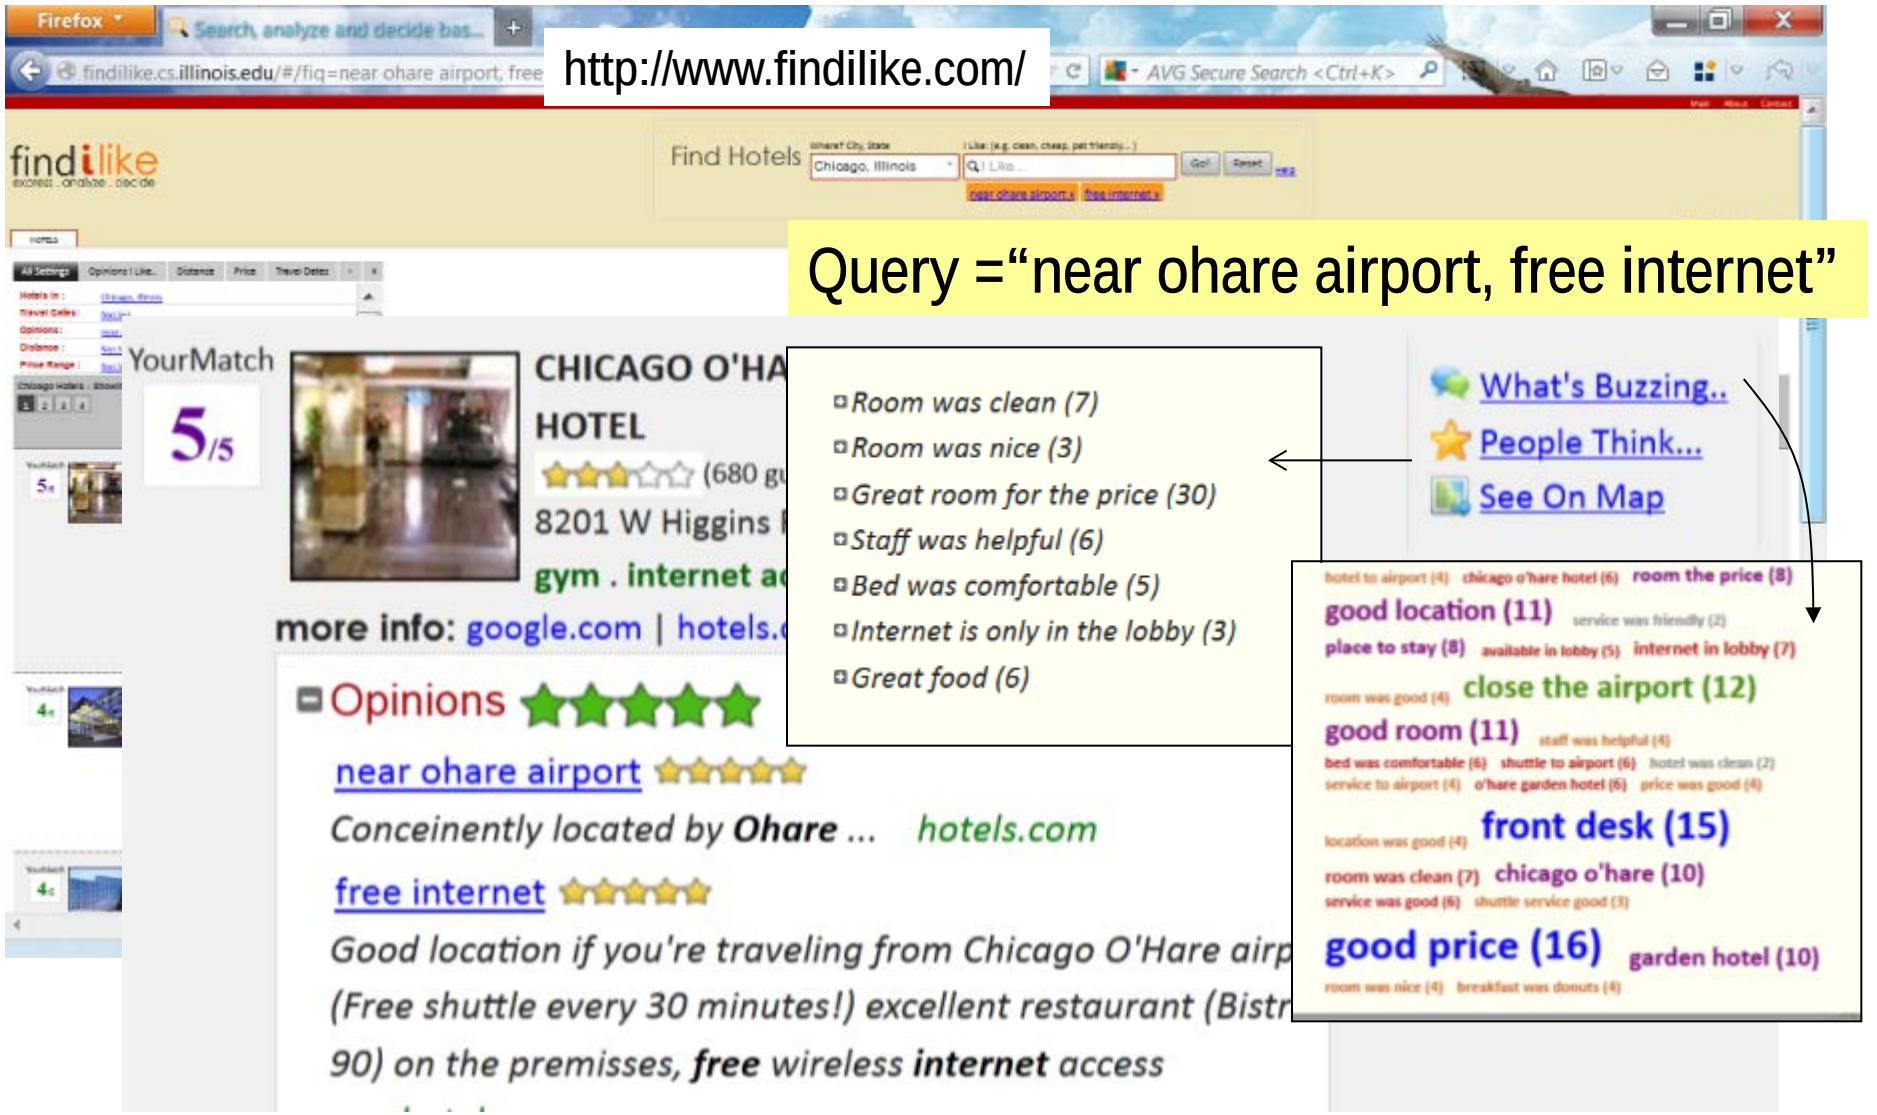

Map Review

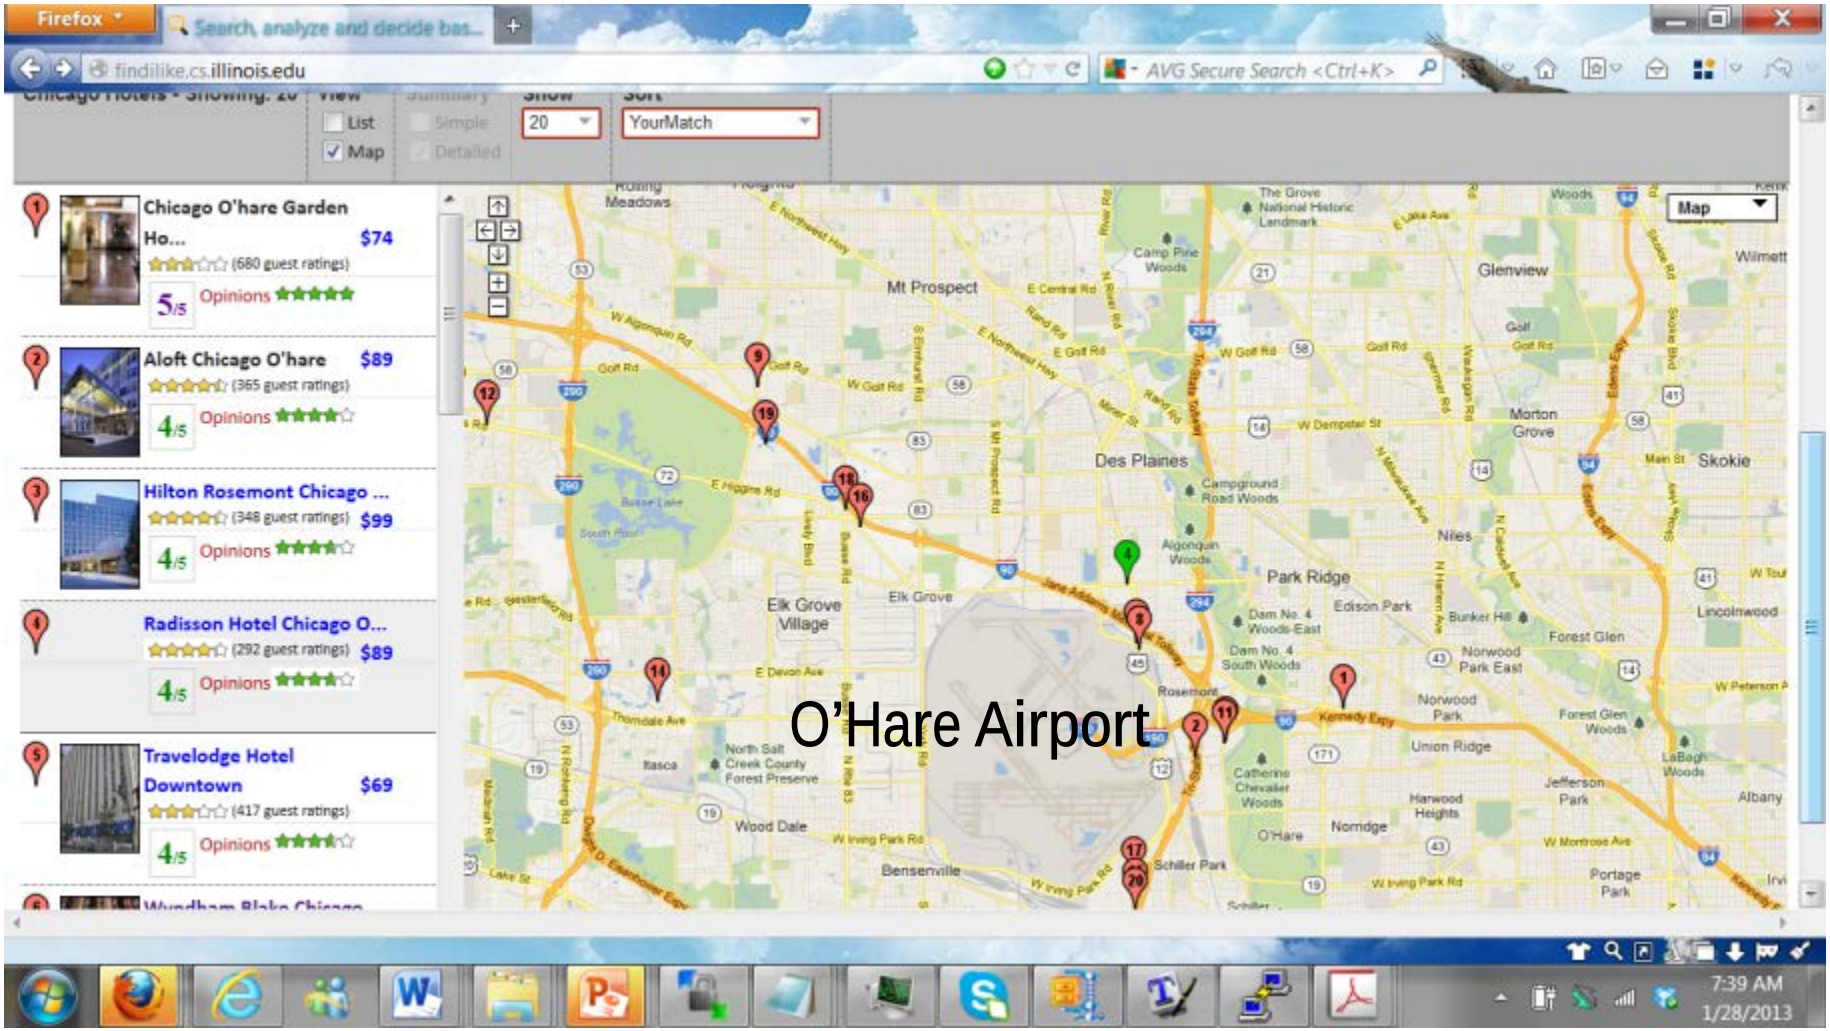
Firefox
findilike.cs.illinois.edu
AVG Secure Search <Ctrl+K>
List
Map
Simple
Detailed
20
YourMatch
Chicago O'hare Garden Ho... $74
(680 guest ratings)
5/5 Opinions
Aloft Chicago O'hare $89
(365 guest ratings)
4/5 Opinions
Hilton Rosemont Chicago ... $99
(348 guest ratings)
4/5 Opinions
Radisson Hotel Chicago O... $89
(292 guest ratings)
4/5 Opinions
Travelodge Hotel Downtown $69
(417 guest ratings)
4/5 Opinions
Map
O'Hare Airport
Mt Prospect
Des Plaines
Elk Grove Village
Park Ridge
Glenview
Wood Dale
Bensenville
Itasca
Norridge
O'Hare
7:39 AM
1/28/2013

Acknowledgments

- Collaborators: Yue Lu, Qiaozhu Mei, Kavita Ganesan, Hongning Wang, and many others
- Funding

Thank You!

Questions/Comments?

More information can be found at http://timan.cs.uiuc.edu/