Reading Tea Leaves: How Humans Interpret Topic Models

Paper by Jonathan Chang, Jordan Boyd-Graber, Sean Gerrish, Chong Wang, David M. Blei

Presentation by M. Falis

Roadmap

- Intro to topic modelling
- Models
- LDA
- Measuring Performance
- MP and TLO
- Results
- Summary

Intro to Topic Modelling

- Motivation
- Topic = Distribution over words
- Document = Mixture of topics
- Problem definition

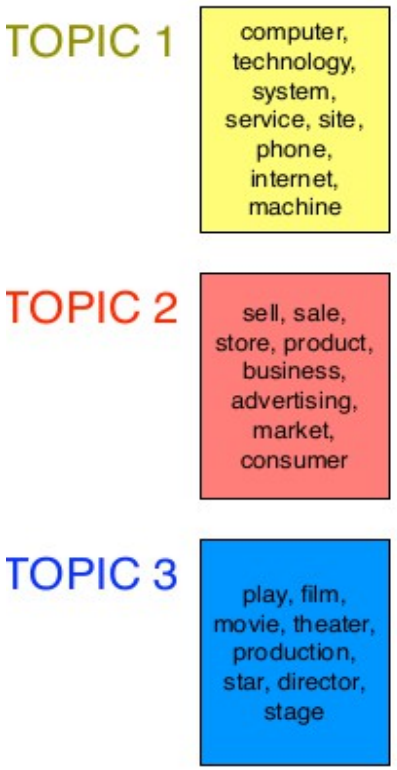

(a) Topics

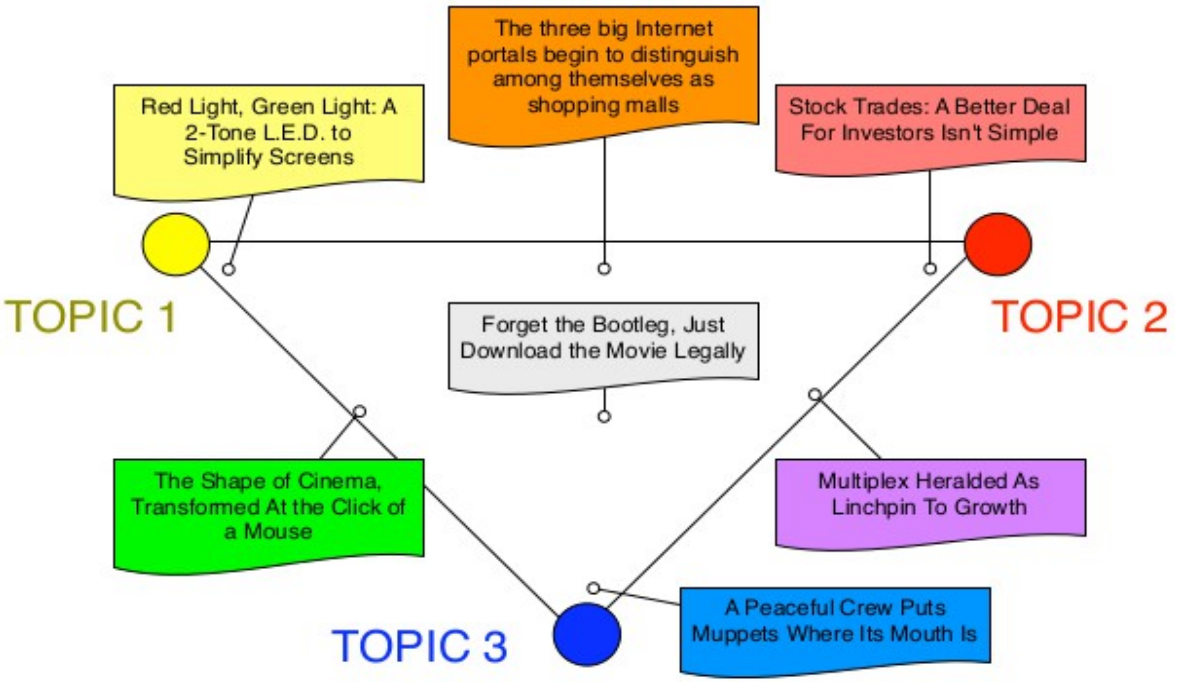

(b) Document Assignments to Topics

Models

- pLSI
- LDA
- CTM

$\alpha = 1$

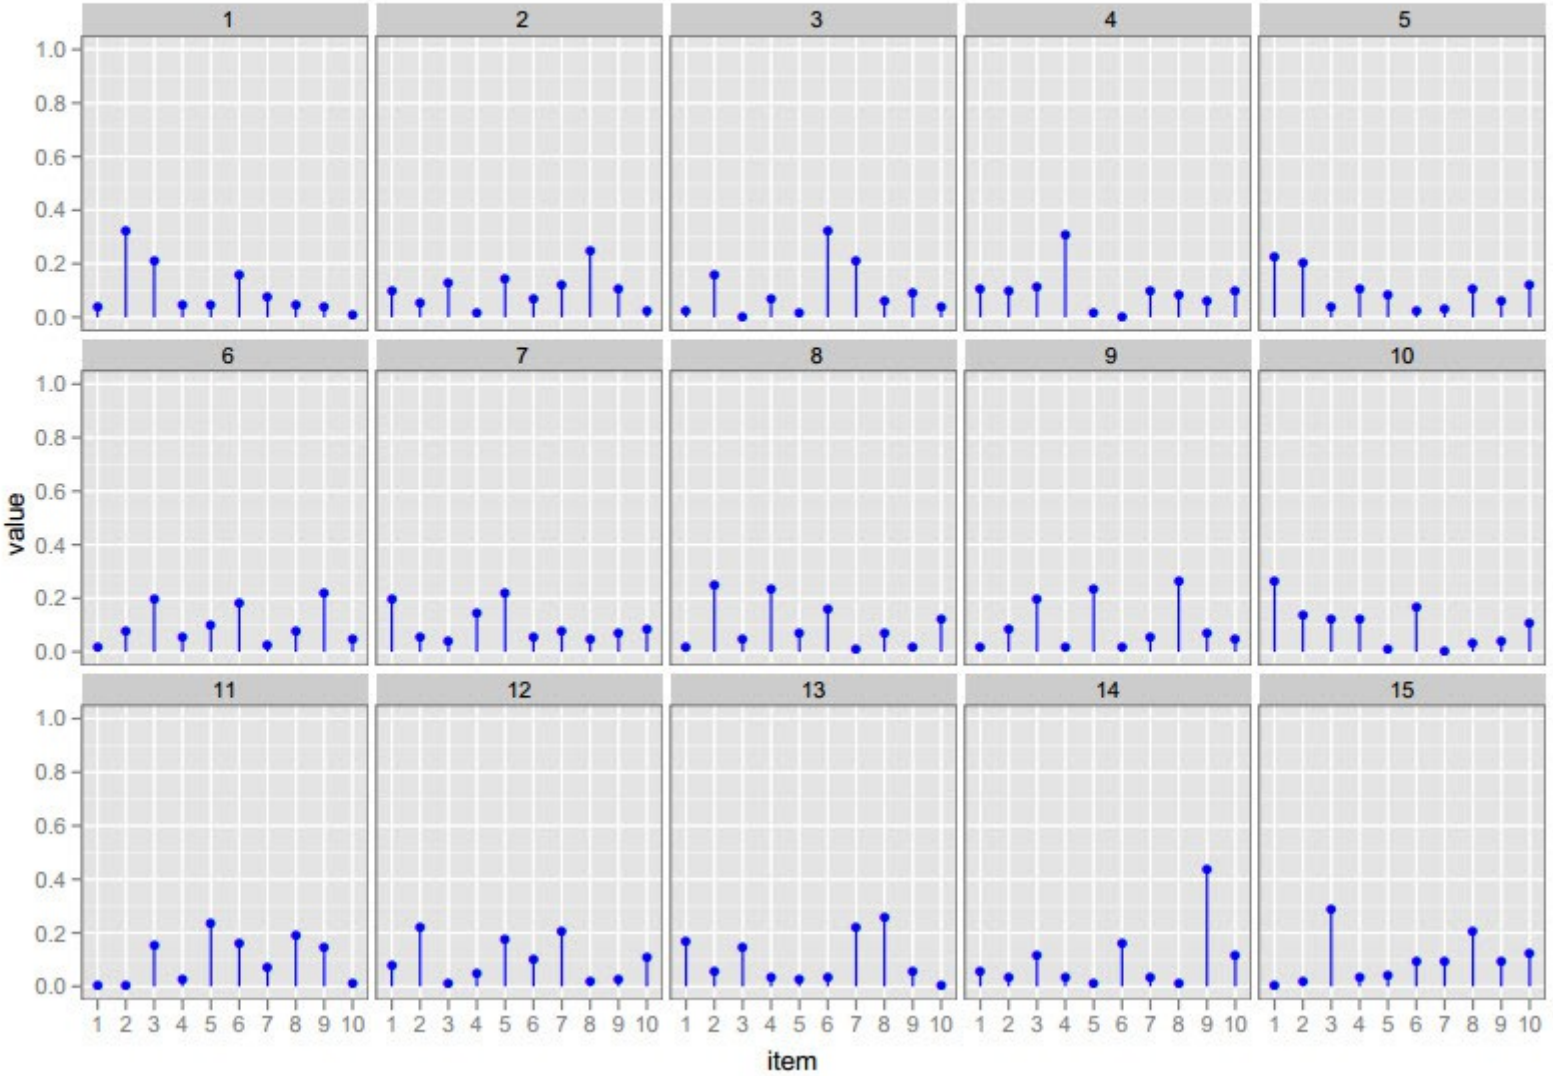
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
value
item
0.0
0.2
0.4
0.6
0.8
1.0
1 2 3 4 5 6 7 8 9 10

$\alpha = 100$

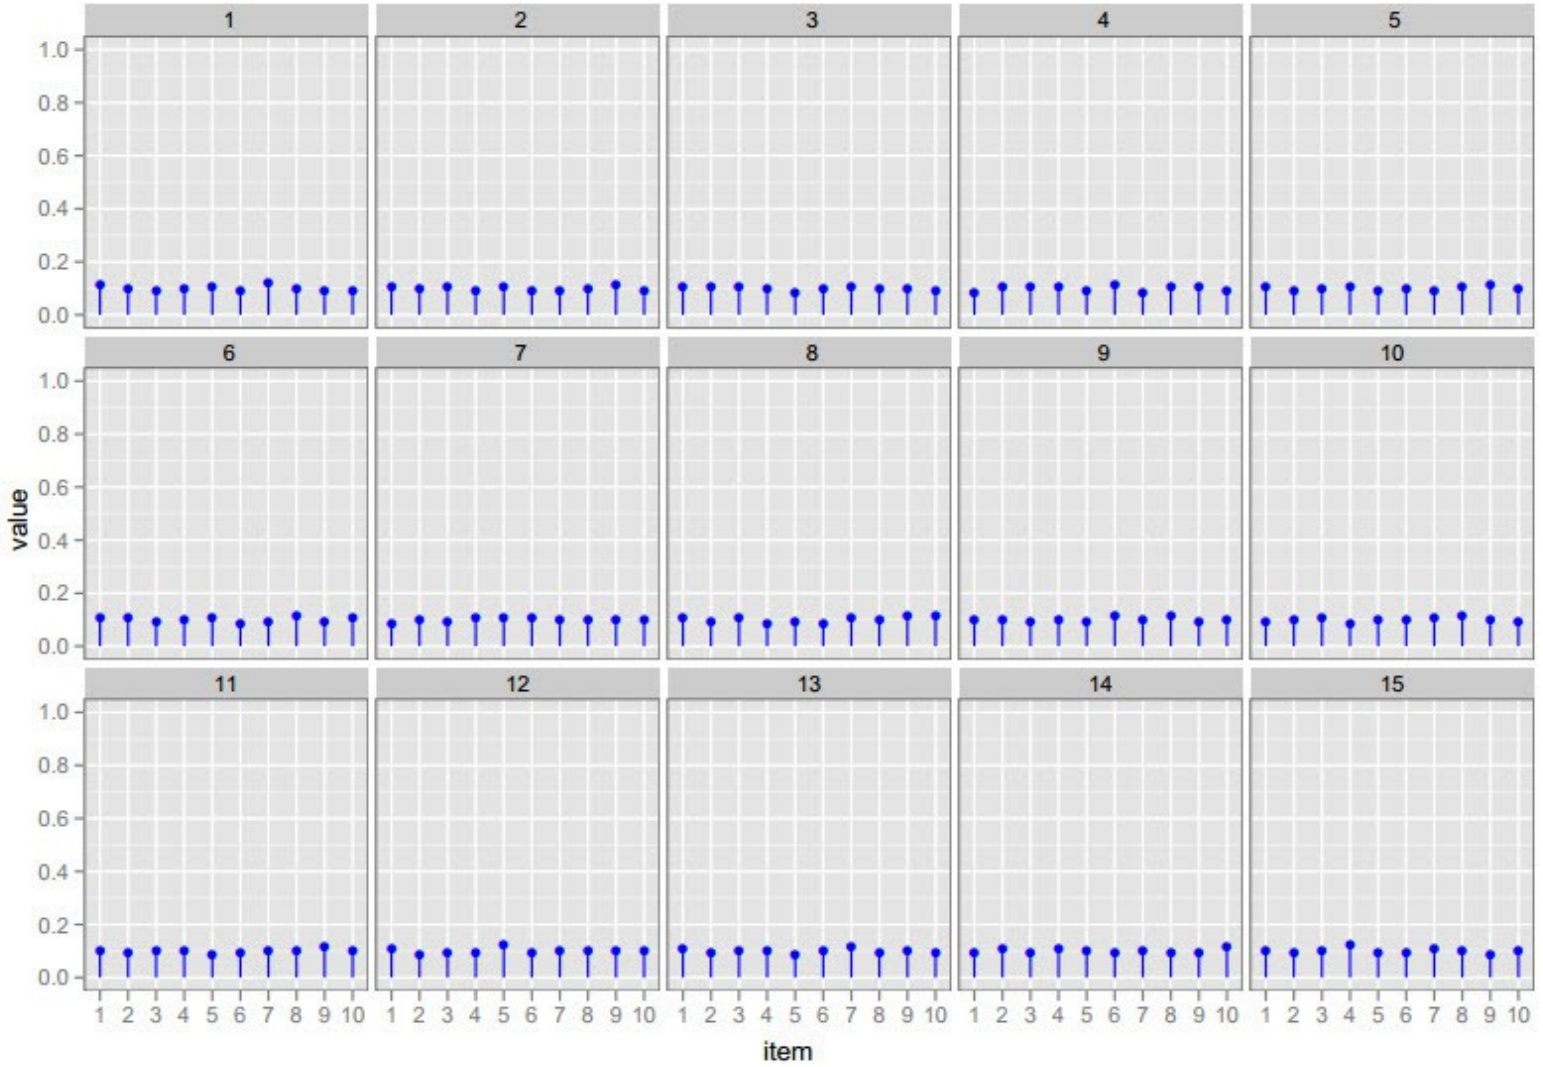
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
value
item
0.0
0.2
0.4
0.6
0.8
1.0
1 2 3 4 5 6 7 8 9 10

$\alpha = 0.1$

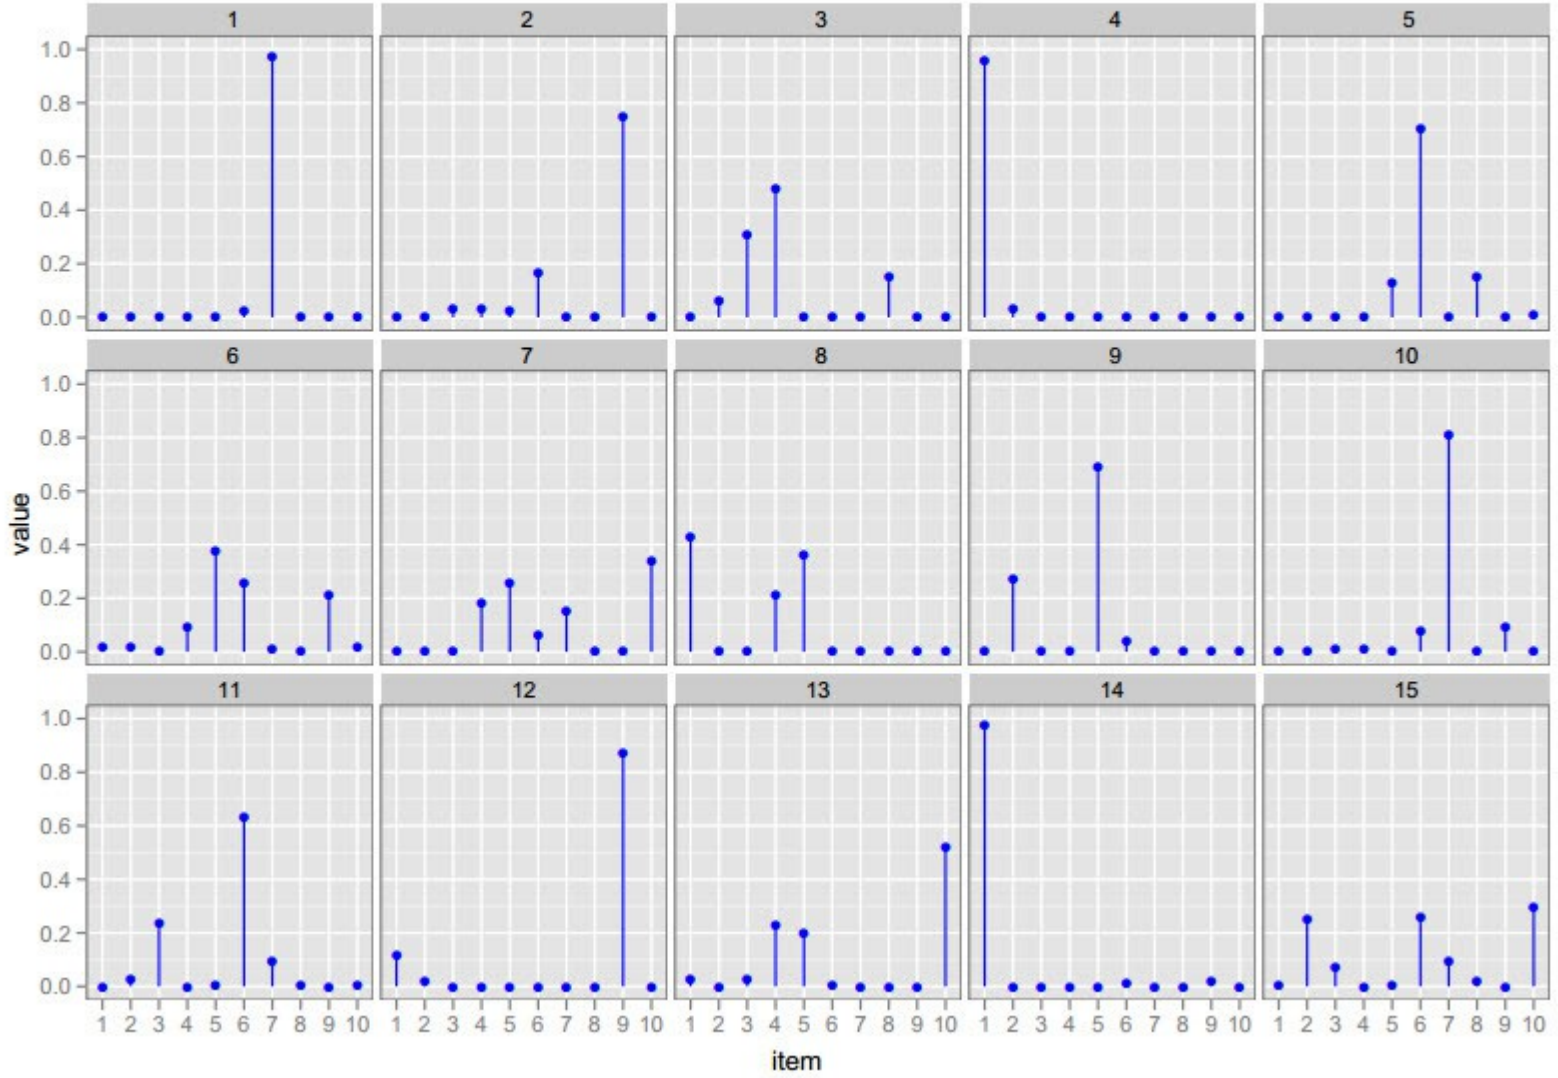
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
value
item
0.0
0.2
0.4
0.6
0.8
1.0
1 2 3 4 5 6 7 8 9 10

LDA
(fitting the model)

- Initial assignment: go through each document in the corpus and to each word assign a random topic from the set of topics T.
- For each document go through each word and for each topic calculate:
  - (a) the probability of topic given the document
  - (b) and the probability of the word given the topic
- reassign each word an new topic, chose the topic with probability (a)*(b)
- repeat until convergence

Measuring Performance

- External Tasks
- Held-out likelihood
- Humans?

$$\text{perplexity}(\text{test set } \boldsymbol{w}) = \exp\left\{-\frac{\mathcal{L}(\boldsymbol{w})}{\text{count of tokens}}\right\}$$

Corpus	Topics	LDA	CTM	pLSI
New York Times	50	**-7.3214 / 784.38**	-7.3335 / 788.58	-7.3384 / 796.43
	100	-7.2761 / 778.24	**-7.2647 / 762.16**	-7.2834 / 785.05
	150	-7.2477 / 777.32	-7.2467 / **755.55**	**-7.2382** / 770.36
Wikipedia	50	**-7.5257** / 961.86	-7.5332 / **936.58**	-7.5378 / 975.88
	100	-7.4629 / 935.53	**-7.4385 / 880.30**	-7.4748 / 951.78
	150	-7.4266 / 929.76	**-7.3872 / 852.46**	-7.4355 / 945.29

MP and TLO

- Word Intrusion and Model Precision
- Topic Intrusion and Topic Log Odds

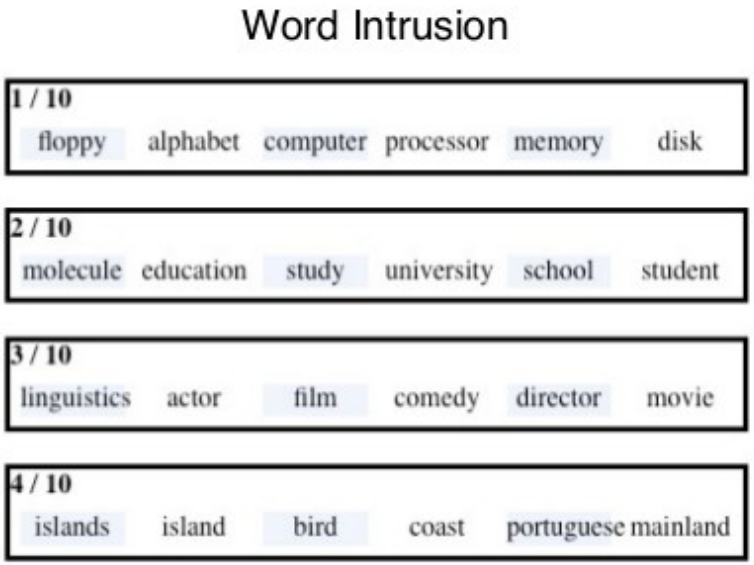

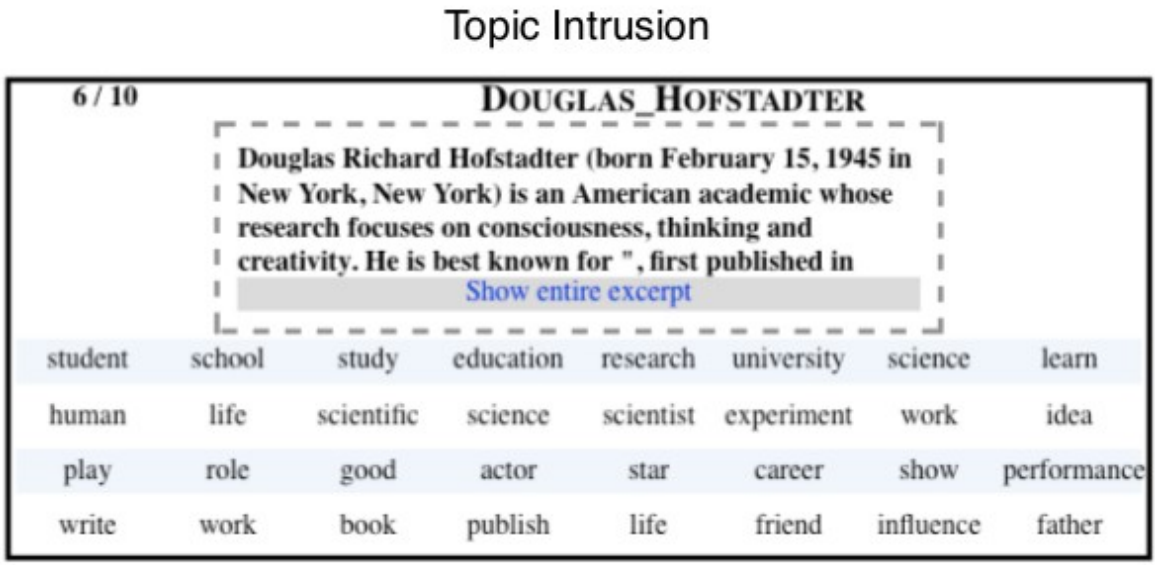

$$\mathrm{MP}_k^m = \textstyle\sum_s \mathbb{1}(i_{k,s}^m = \omega_k^m)/S.$$

$$\mathrm{TLO}_d^m = (\textstyle\sum_s \log \hat{\theta}_{d,j_{d,*}^m}^m - \log \hat{\theta}_{d,j_{d,s}^m}^m)/S.$$

Results

- Model Precision

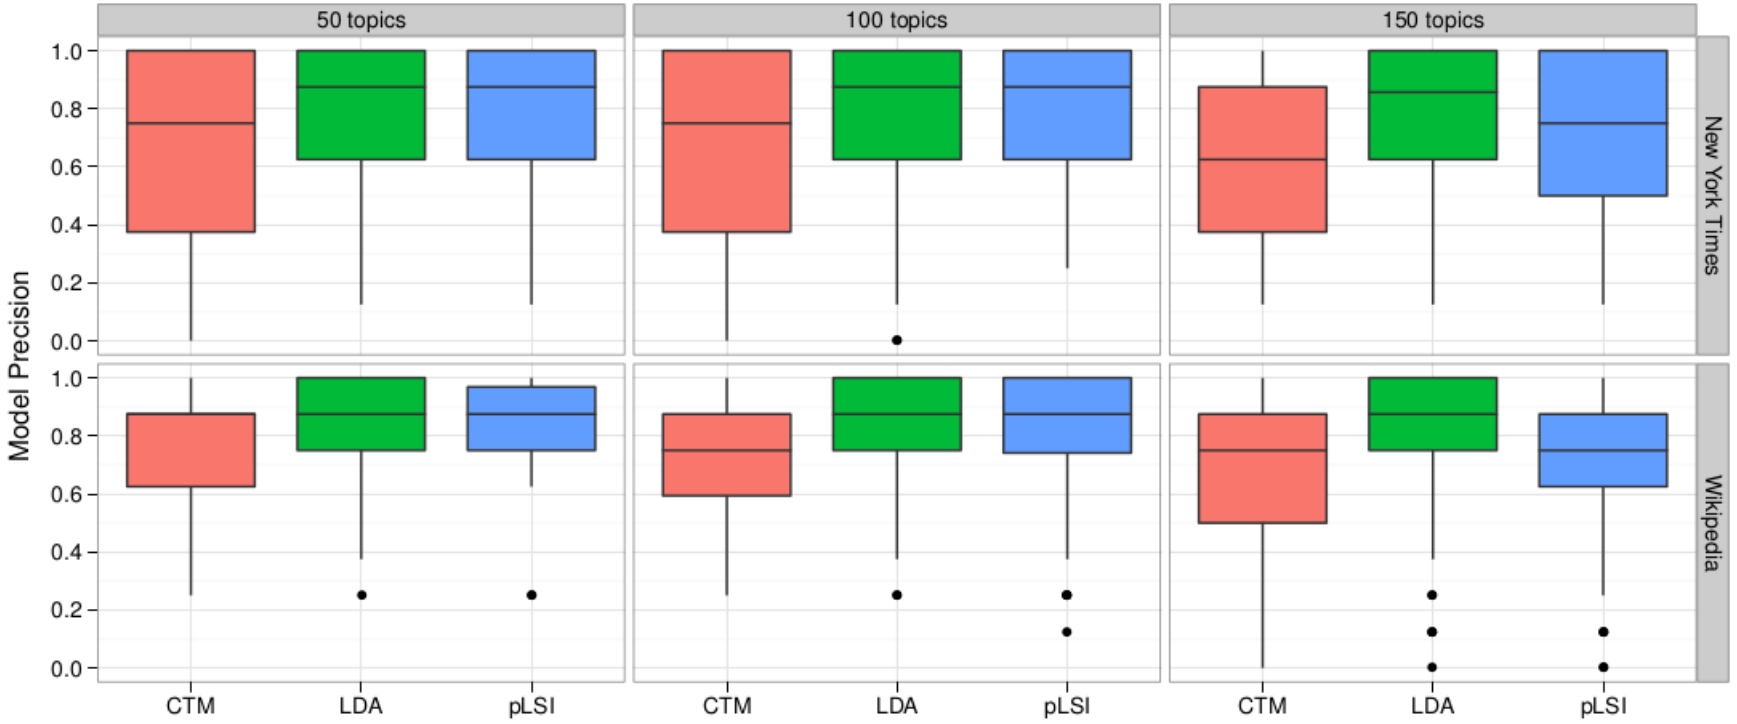

Results

- TLO

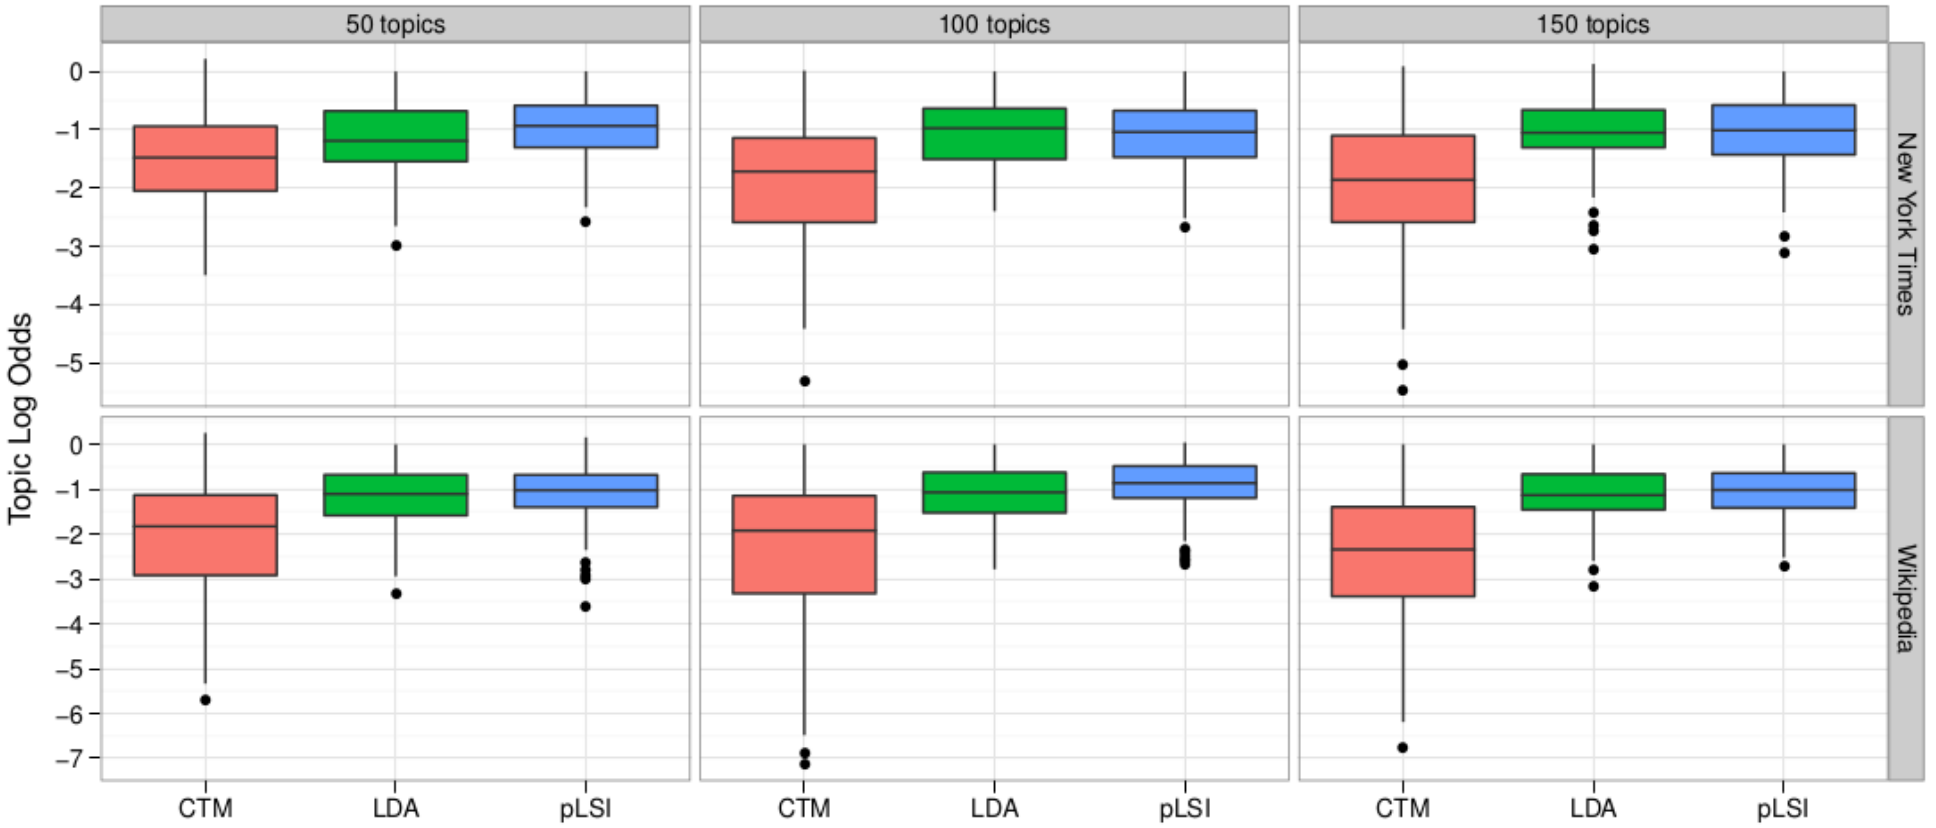
50 topics
100 topics
150 topics
New York Times
Wikipedia
Topic Log Odds
0
−1
−2
−3
−4
−5
−6
−7
CTM
LDA
pLSI

Results

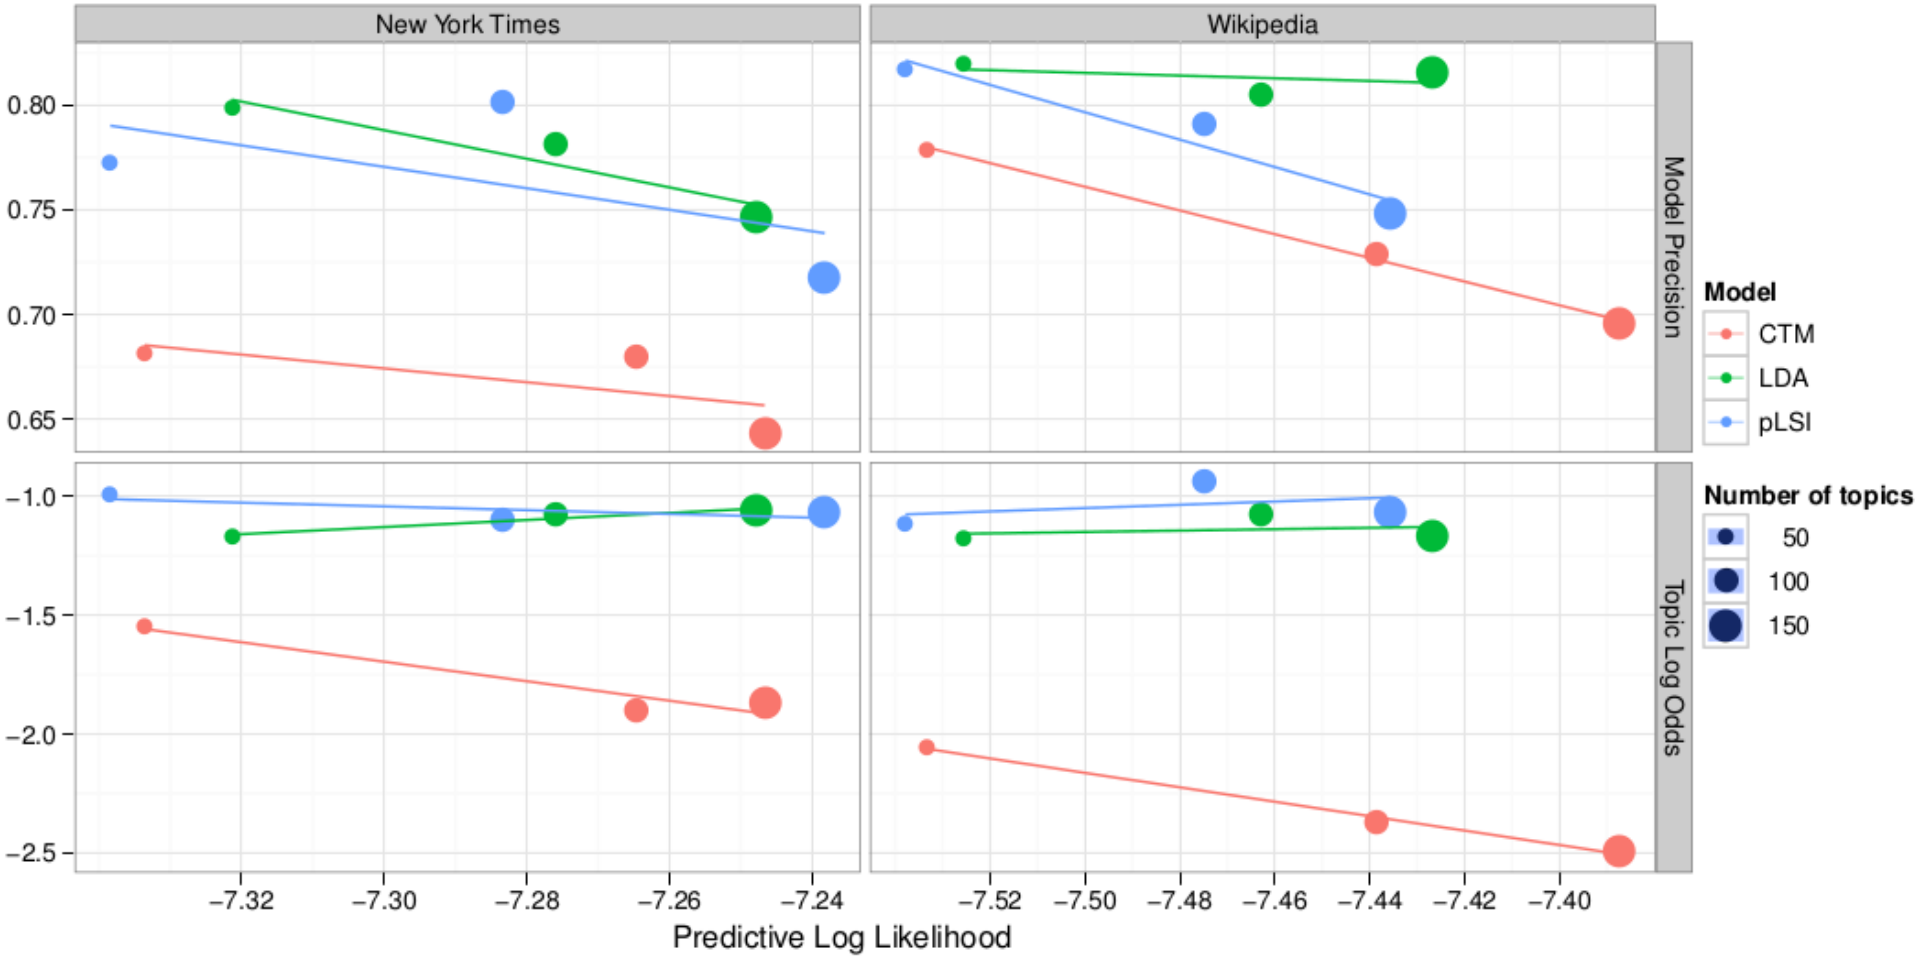
New York Times
Wikipedia
Model Precision
Topic Log Odds
0.80
0.75
0.70
0.65
−1.0
−1.5
−2.0
−2.5
−7.32
−7.30
−7.28
−7.26
−7.24
−7.52
−7.50
−7.48
−7.46
−7.44
−7.42
−7.40
Predictive Log Likelihood
Model
CTM
LDA
pLSI
Number of topics
50
100
150

Summary

- The use of metrics evaluating real world task performance
- Lower perplexity (higher likelihood) is not necessarily correlated to better coherence of topics

References

- http://www.umiacs.umd.edu/~jbg/docs/nips2009-rtl.pdf
- http://www.inf.ed.ac.uk/teaching/courses/nlu/lectures/nlu_l03-lda.pdf