

Reading Tea Leaves: How Humans Interpret Topic Models

By Jonathan Chang, Jordan Boyd-Graber, (Chong Wang), et al.


NIPS 2009

Presented by Stephen Mayhew


Feb 2013

Motivation


- How to evaluate topic models?
- “Anecdotally”, “empirically”
- Intrinsic vs. extrinsic



SVM Document Classification on Reuters 21578



(a)



(b)

Human Metrics

1. Word intrusion
2. Topic intrusion

Crowdsourced approach using Amazon Mechanical Turk

Evaluating three different approaches: LDA, pLSI, CTM.

Word Intrusion

“Spot the intruder word”

Process:

1. Select a topic at random
2. Choose the 5 most probable words from the topic
3. Choose an improbable word from this topic (which is probable in another topic)
4. Shuffle
5. Present to subject

Word Intrusion

If the topic set is coherent, then the users will agree on the outlier.

If the topic set is incoherent, then the users will choose the outlier at random.

1 / 10
floppy alphabet computer processor memory disk

2 / 10
molecule education study university school student

3 / 10
linguistics actor film comedy director movie

4 / 10
islands island bird coast portuguese mainland

Topic Intrusion

“Spot the intruder topic”

Process:

1. Choose a document
2. Choose the three highest-prob. topics for this document
3. Choose one low-prob. topic for this document
4. Shuffle
5. Present to subject

Topic Intrusion

6 / 10

DOUGLAS_HOFSTADTER

Douglas Richard Hofstadter (born February 15, 1945 in New York, New York) is an American academic whose research focuses on consciousness, thinking and creativity. He is best known for "

[Show entire excerpt](#)

| | | | | | | | |
|---------|--------|------------|-----------|-----------|------------|-----------|-------------|
| student | school | study | education | research | university | science | learn |
| human | life | scientific | science | scientist | experiment | work | idea |
| play | role | good | actor | star | career | show | performance |
| write | work | book | publish | life | friend | influence | father |

Word Intrusion: how to measure it

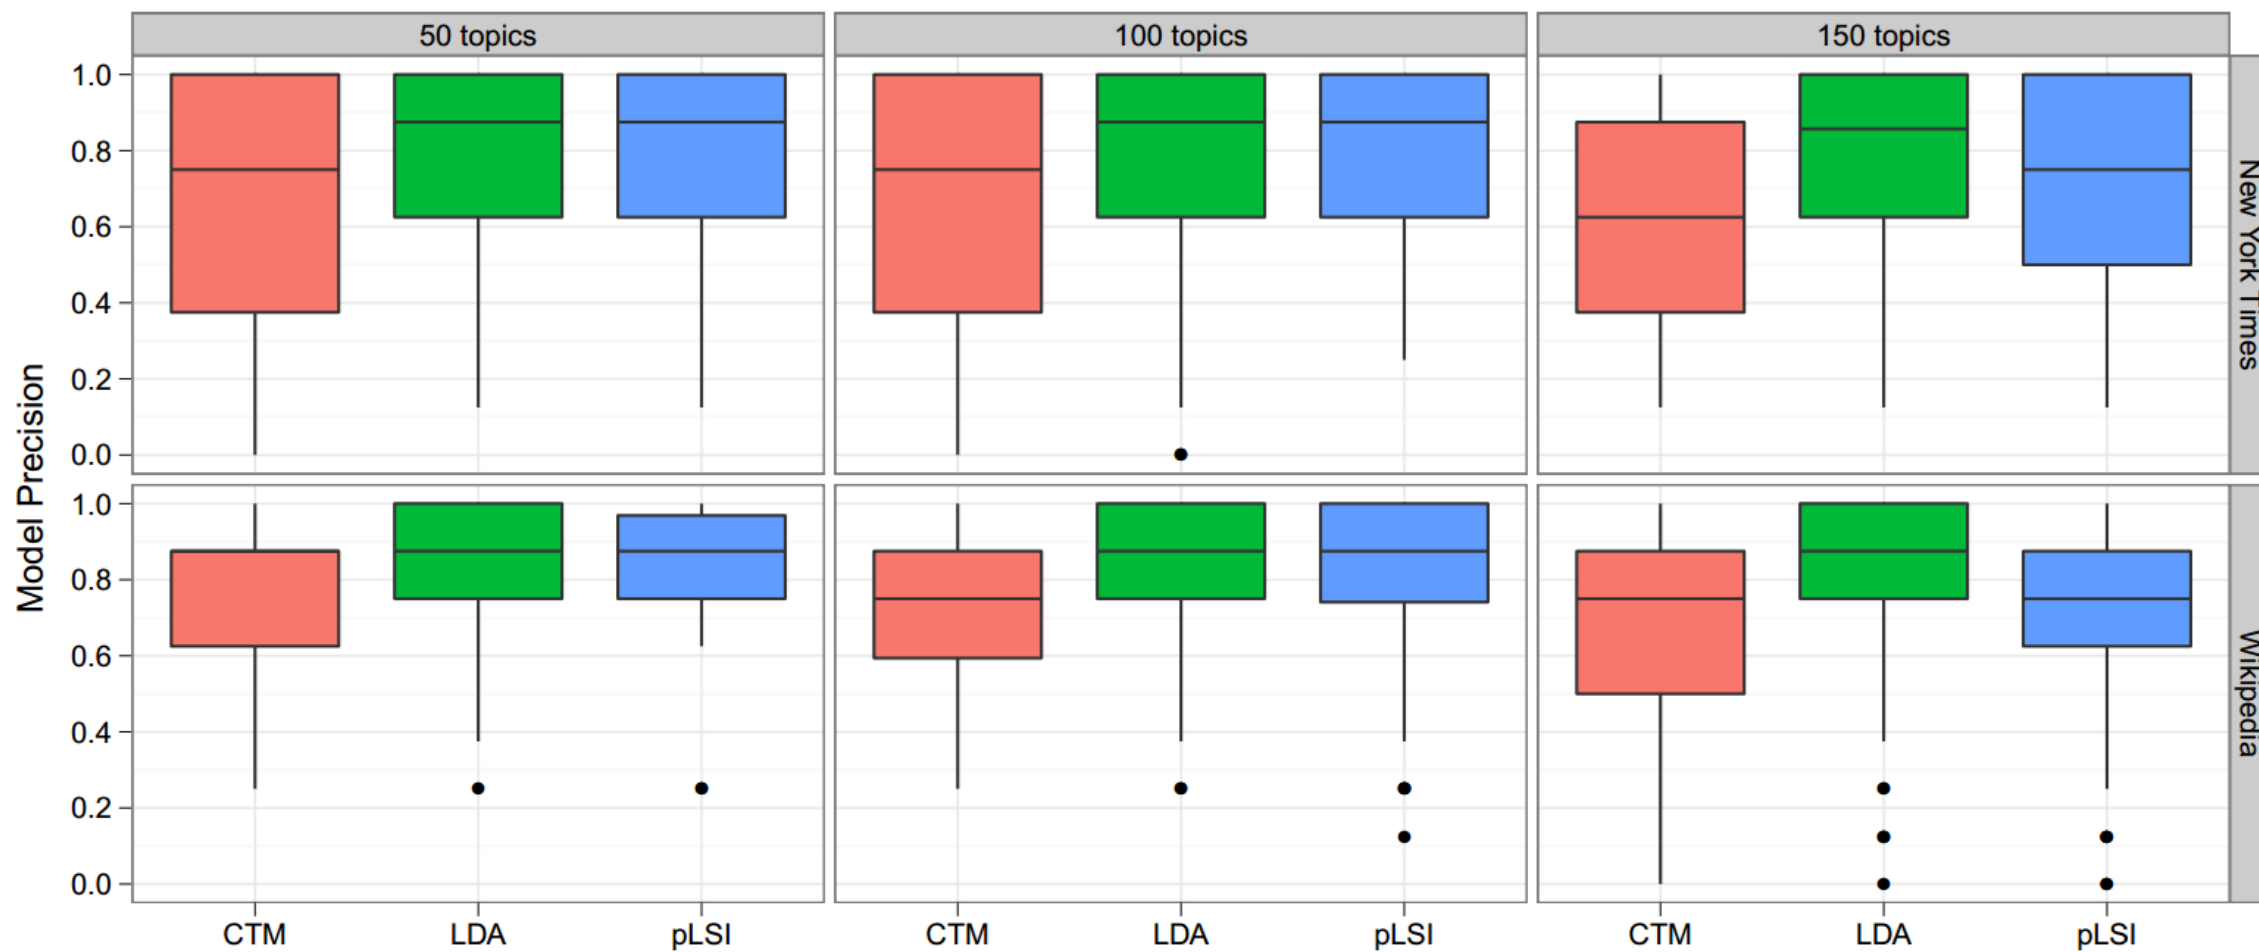
Model parameters:


$$\text{MP}_k^m = \sum_s \mathbb{1}(i_{k,s}^m = w_k^m) / S$$

Which is just a fancy way of saying:

$$\frac{\text{number of people correct}}{\text{total number of people}}$$

Word Intrusion





NYT corpus, 50 topic LDA model

Topic intrusion: how to measure It


Topic Log Odds (TLO):

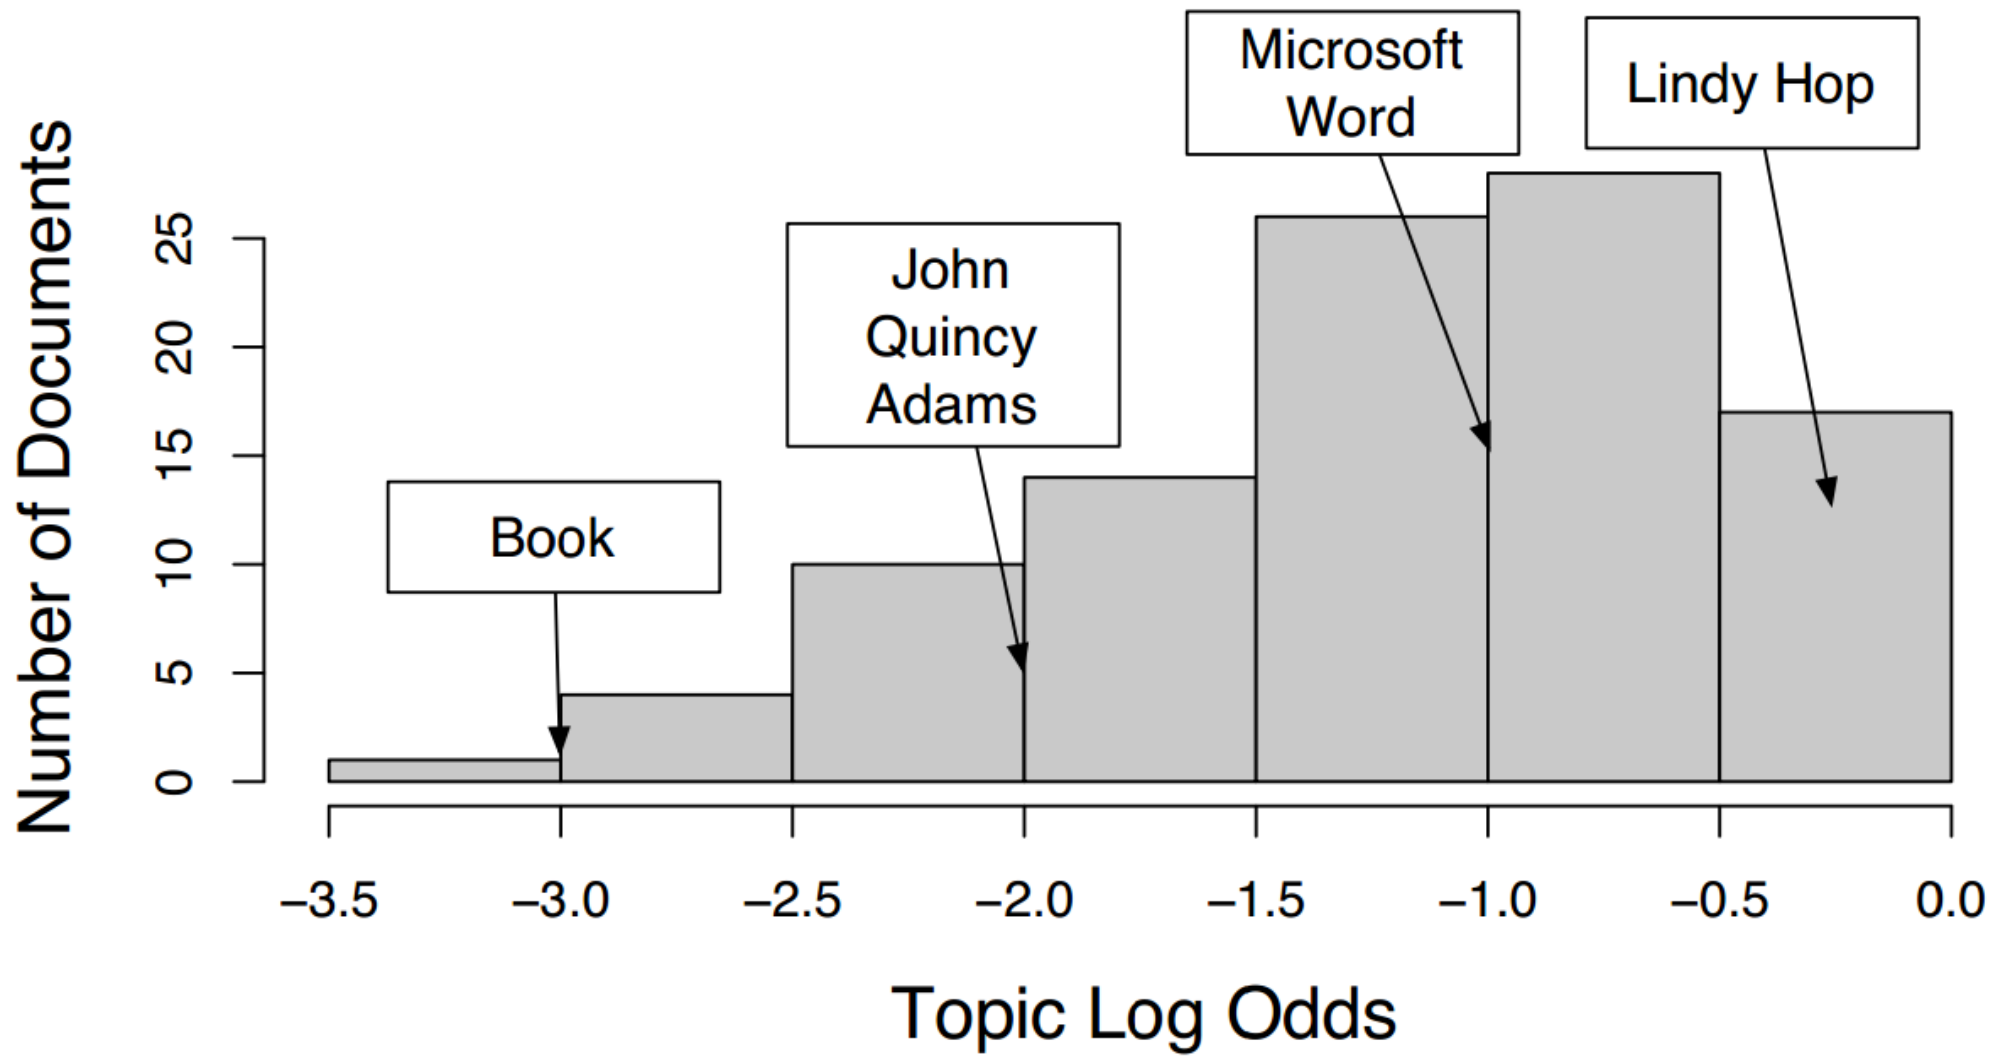
$$\text{TLO}_d^m = \left(\sum_s \log \hat{\theta}_{d,j_{d,*}^m}^m - \log \hat{\theta}_{d,j_{d,s}^m}^m \right) / S$$

Translation: normalized difference between probability mass of actual “intruder” and selected “intruder”.


Upper bound is 0, higher is better.

Topic Intrusion





Wikipedia corpus, 50 topic LDA model



Problems

Measures homogeneity (synonymy), not topic strength (coherence)

Example document: curling

Possible topic: broom, ice, Canada, rock, sheet, stone

Consider syntactic differences:

organization, physicality, proportions, red