Is Your Anchor Going Up or Down? Fast and Accurate Supervised Topic Models

Thang Nguyen, Jordan Boyd-Graber, Jeff Lund, Kevin Seppi, and Eric Ringger

daithang@umiacs.umd.edu, Jordan.Boyd.Graber@colorado.edu, {jefflund, kseppi}@byu.edu, ringger@cs.byu.edu

University of Maryland, College Park / University of Colorado, Boulder / Brigham Young University

Motivation

- Supervised topic models leverage latent document-level themes to capture nuanced sentiment, create sentiment-specific topics and improve sentiment prediction.
- Examples include Supervised LDA (Blei et al., 2007), Labelled LDA (Ramage et al., 2009), Med LDA (Zhu et al., 2009), etc.
- *The downside for Supervised LDA is that it is slow, which this work addresses.*

Contribution

- We create a supervised version of the anchor word algorithm (**ANCHOR**) (Arora et al., 2013).
- This supervised anchor word algorithm (**SUP ANCHOR**) is very fast because it inherits fast inference from the anchor algorithm.
- Experiments on three sentiment datasets show that **SUP ANCHOR** is comparable to Supervised LDA (**SLDA**) in terms of prediction accuracy.
- Anchor words learned by **SUP ANCHOR** provide great insight.

Anchor Algorithm (Arora et al., 2013)

- **Anchor Word Assumption**: each topic has at least one word that is uniquely indicative of the presence of the given topic. E.g. "wicket" for the cricket topic; "fifa" for the soccer topic.
- **Problem**: fast but no supervised information.

Supervised Anchor Word

$$\bar{Q} \equiv \begin{bmatrix} p(w_1|w_1)\dots & & \\ & \vdots & \\ & & p(w_j|w_i) \end{bmatrix}$$

$$S \equiv \begin{bmatrix} p(w_1|w_1)\dots & & p(y^{(l)}|w_1) \\ & \vdots & \vdots \\ & & p(w_j|w_i) \quad p(y^{(l)}|w_i) \end{bmatrix}$$

New column(s) encoding word-sentiment relationship

$$S_{i,(V+l)} \equiv \frac{\sum_d (\mathbb{1}[i \in d] \cdot \mathbb{1}[y_d = y^{(l)}])}{\sum_d \mathbb{1}[i \in d]}. \quad (1)$$

$$S_{i,\cdot} = \sum_{g_k \in \mathcal{G}} C_{i,k} S_{g_k,\cdot}. \quad (2)$$

- *We focus on learning sentiment-specific anchor words.* This will in turn lead to sentiment-specific topics.
- We explicitly represent the combination of words and discrete sentiment values by appending the conditional probability $Pr[sentiment|word]$ to word co-occurrence vector $\bar{Q}_{word}$.
- **Outcome**: we learn better topic representations and produce topics that are suitable for sentiment prediction.

Supervised Anchor Word (Cont.)

- **Anchor words form a convex hull that encloses all other words in the vocabulary.**
- **Adding sentiment related dimensions moves words UP or DOWN; forming sentiment-specific anchor words.**

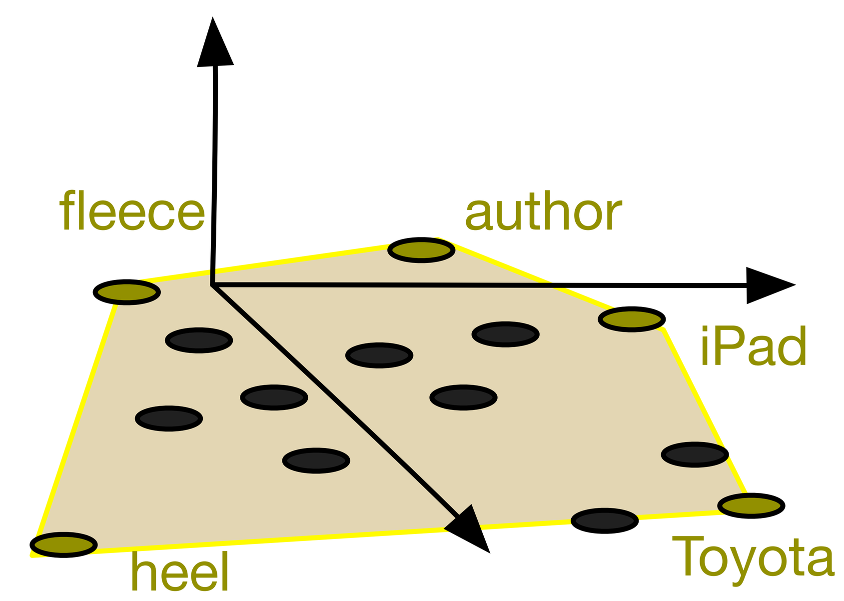

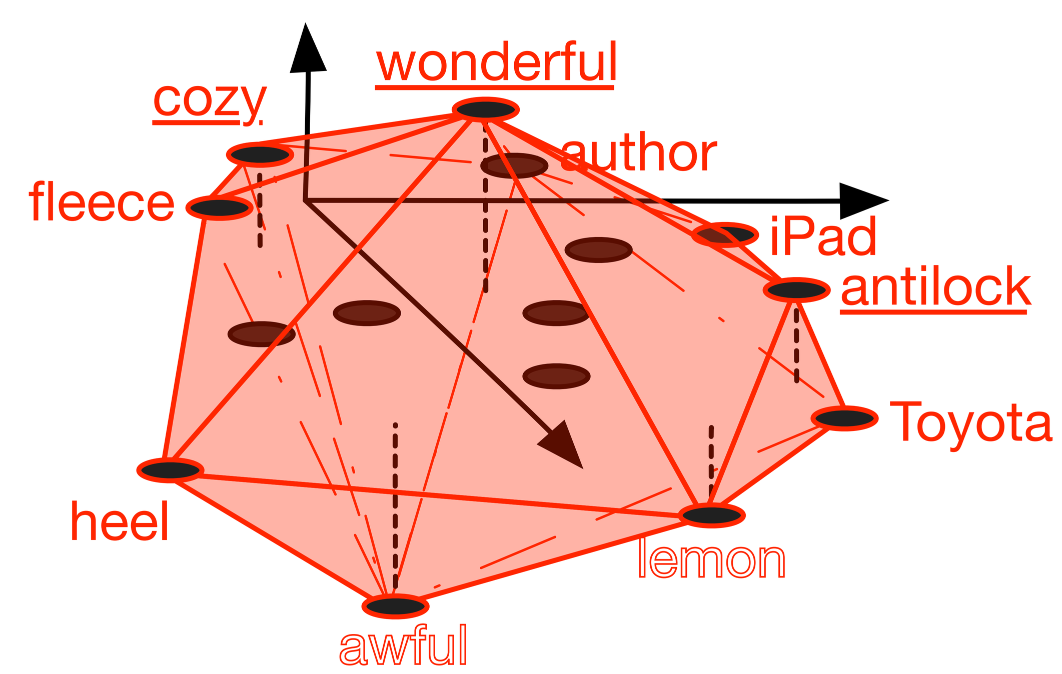

Figure : Adding an additional dimension to capture sentiment changes the convex hull: positive words appear above the original 2D plane (underlined) and negative words appear below (in outline).

Experiments

- **Goal**: Evaluate the new topics generated by the proposed model in a prediction task. We focus on binary classification in sentiment analysis datasets.
- Sentiment datasets.

Corpus	Train Docs	Test Docs	Tokens	Types	Positive Sentiment
AMAZON	13,300	3,314	1,031,659	2,662	52.2%
TRIPADVISOR	115,384	28,828	12,752,444	4,867	41.5%
YELP	13,955	3,482	1,142,555	2,585	27.7%

Table : Statistics for the datasets employed in the experiments.

Runtime Analysis

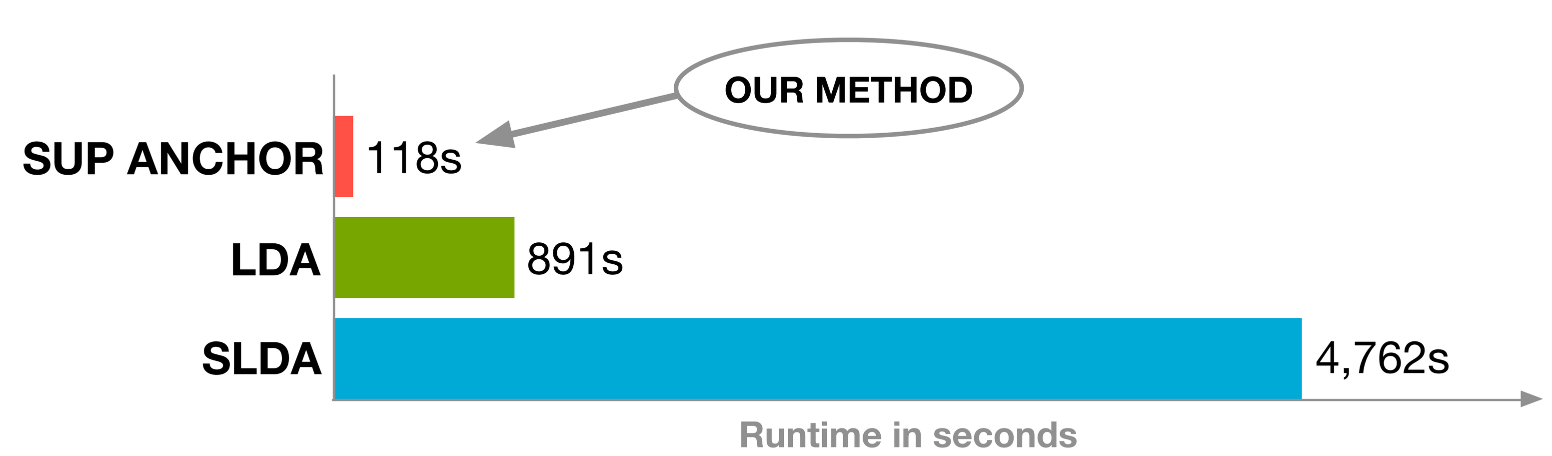

Figure : Total time for training and prediction on AMAZON dataset; **SUP ANCHOR** takes much less time than **SLDA**.

Prediction Accuracy & Topic Interpretability

- For prediction accuracy, **SUP ANCHOR** consistently outperforms all other models on all sentiment datasets.
- **SUP ANCHOR** has comparable topic interpretability (Lau et al., 2014, NPMI) to **ANCHOR**.
- Nguyen et al., 2014 show that with regularization anchor algorithms can produce comparable topics with those produced by **LDA**.

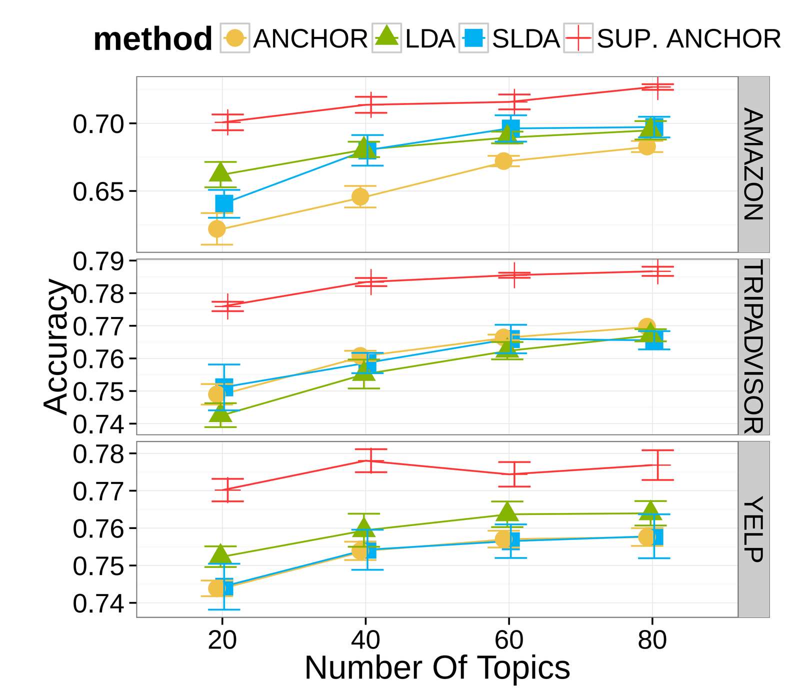

Figure : **SUP ANCHOR** outperforms **ANCHOR**, **LDA**, and **SLDA** on all three datasets.

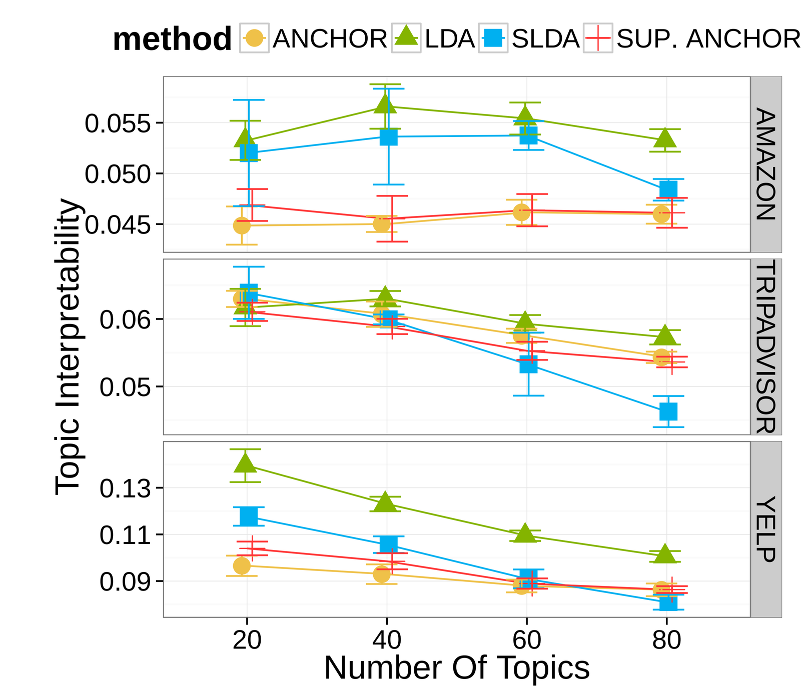

Figure : **SUP ANCHOR** and **ANCHOR** produce the same topic quality. **LDA** produces the best topics.

Anchor Words and Their Topics

- **SUP ANCHOR** produces anchor words around the same strong lexical cues that could discover better sentiment topics (e.g. positive reviews mentioning a *favorite* restaurant or negative reviews complaining about *long waits*).

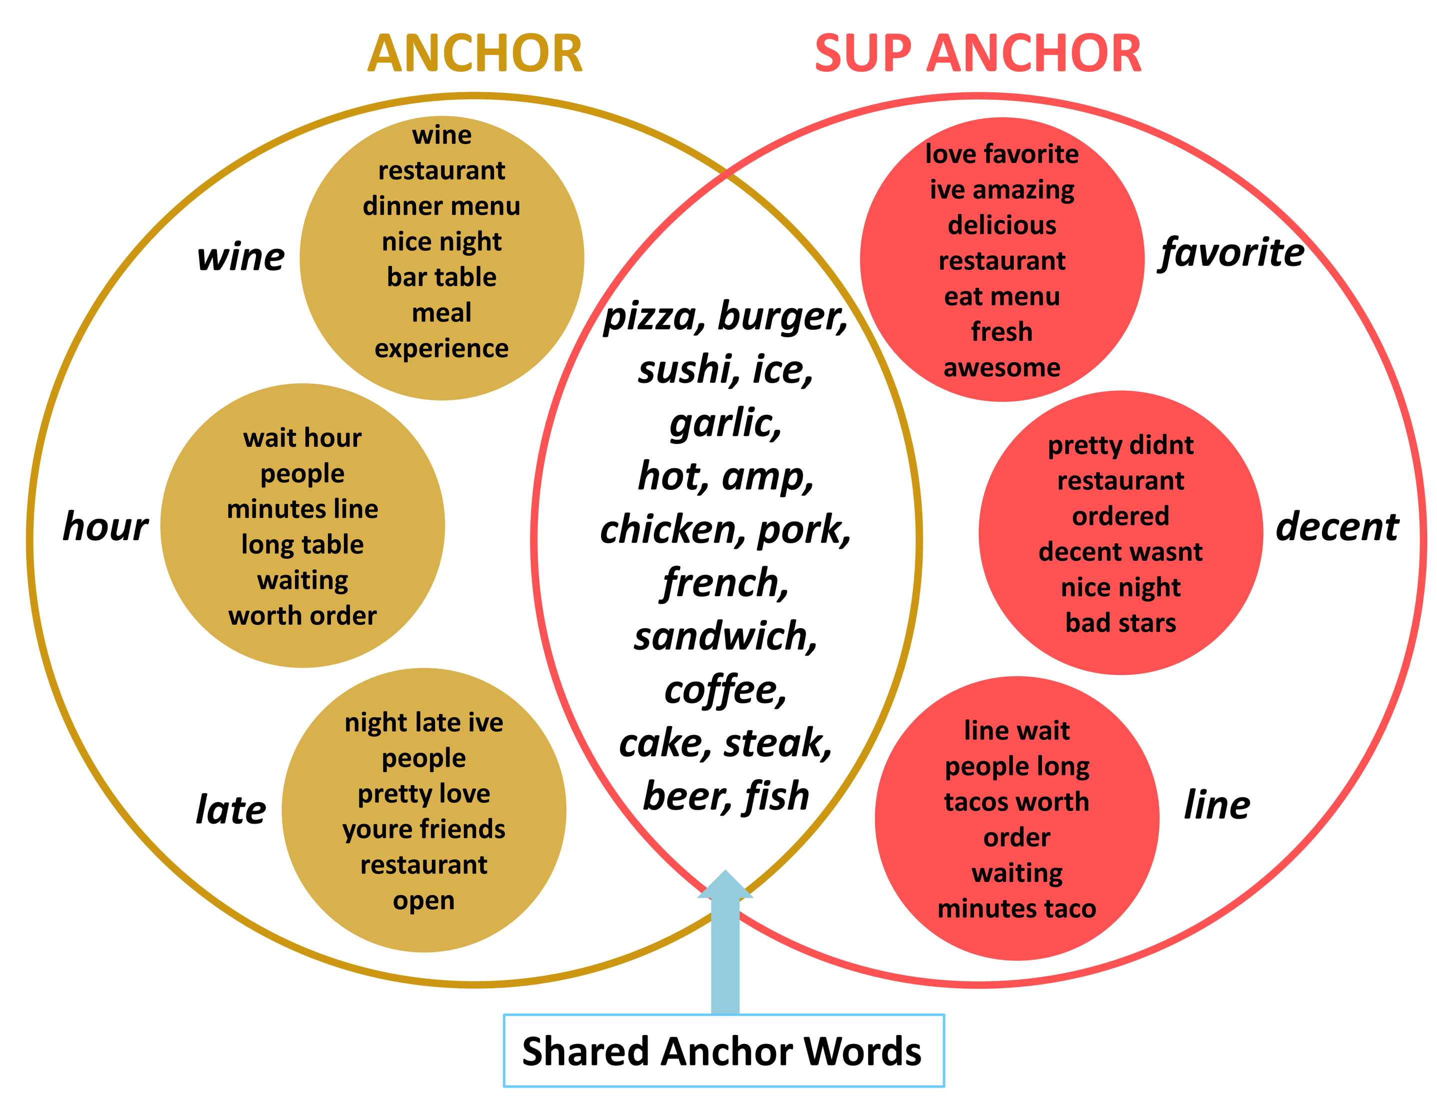

Figure : Topics generated for the YELP dataset: anchor words shared by both **ANCHOR** and **SUP ANCHOR** are listed. The distinct anchor words reflect positive ("favorite") and negative ("line") sentiment rather than less sentiment-specific qualities of restaurants (e.g., restaurants open "late").

Future Directions

- Explore other richer word representations (e.g. word2vec, more smoothing on added dimensions).
- Work on inferring document-topic distributions directly.
- Incorporate into Interactive Topic Modeling and Active Learning.

Email: daithang@umiacs.umd.edu, Jordan.Boyd.Graber@colorado.edu, {jefflund, kseppi}@byu.edu, ringger@cs.byu.edu