Twitter Opinion Topic Model:
Extracting Product Opinions from Tweets by Leveraging Hashtags and Sentiment Lexicon

Presenter: **Peter Christen**

Authors: Kar Wai Lim, Wray Buntine

6 November 2014

Aspect-based Opinion Aggregation

- Opinion Aggregation for reviews.
  - A process to collect reviews of products and services to analyze in aggregate.
- Aspect-based.
  - Groups reviews based on “aspects”.
  - Example:
    - Product types
      - Game consoles
      - Mobile phones
    - Product specs
      - Computer specs
      - Flight quality

Aspect	Examples
Game console	PS4, Xbox One, Wii U...
Mobile phone	iPhone, Samsung Note...
Computer spec	CPU, RAM, GPU...
Flight quality	Food, customer service...

Existing Method

- Independent Latent Dirichlet Allocation (ILDA).
  - Current state-of-the-art for aspect-based opinion aggregation (Moghaddam, 2012).
  - ILDA is a type of topic model.
  - Perform analysis on target-opinion pairs.
- Target-opinion pairs are extracted during preprocessing using Stanford dependency parser.
  - Examples:

Target	Opinion
iPhone	Awesome
Service	Good
Weather	Hot

ILDA

- Graphical model:

Note:
Shaded = observed
Unshaded = latent

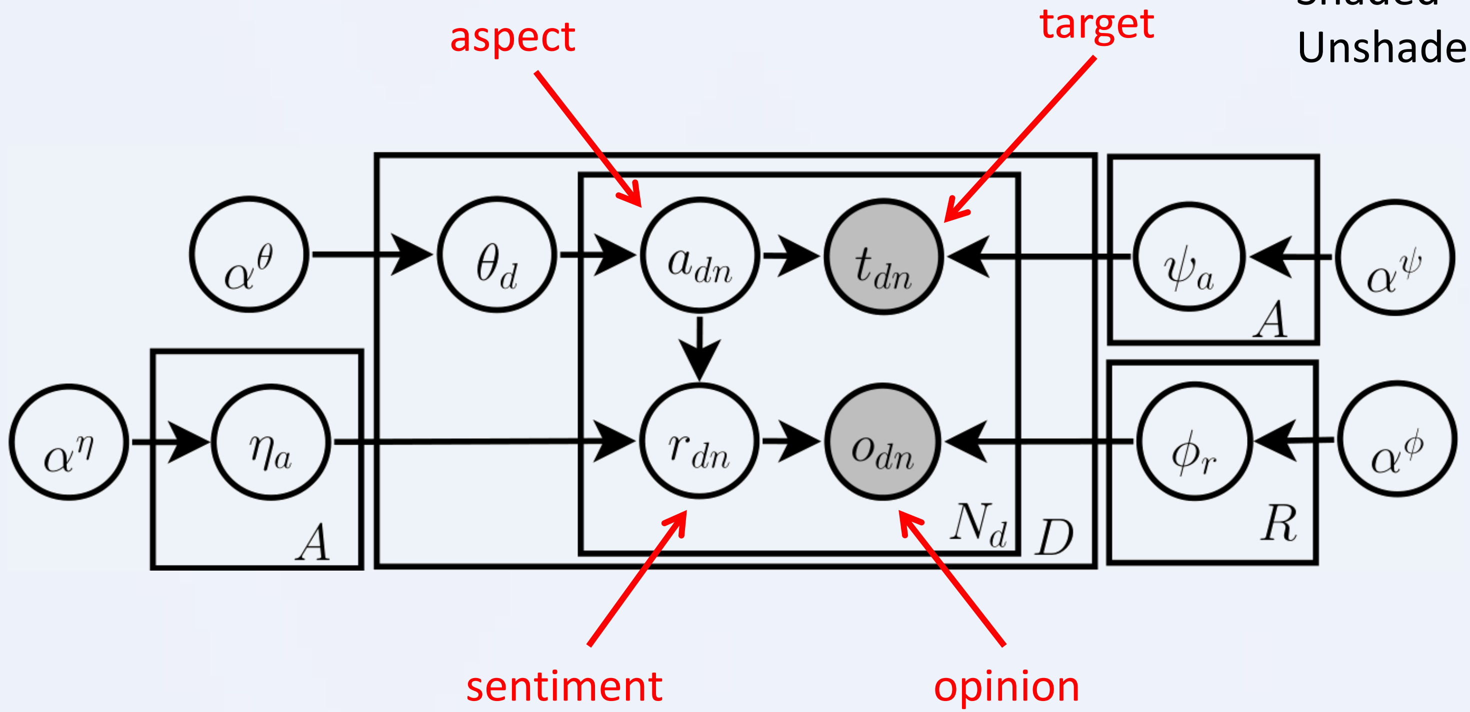

Aspects and sentiments are discrete labels learned by the model.

ILDA

- Graphical model:

Note:
Shaded = observed
Unshaded = latent

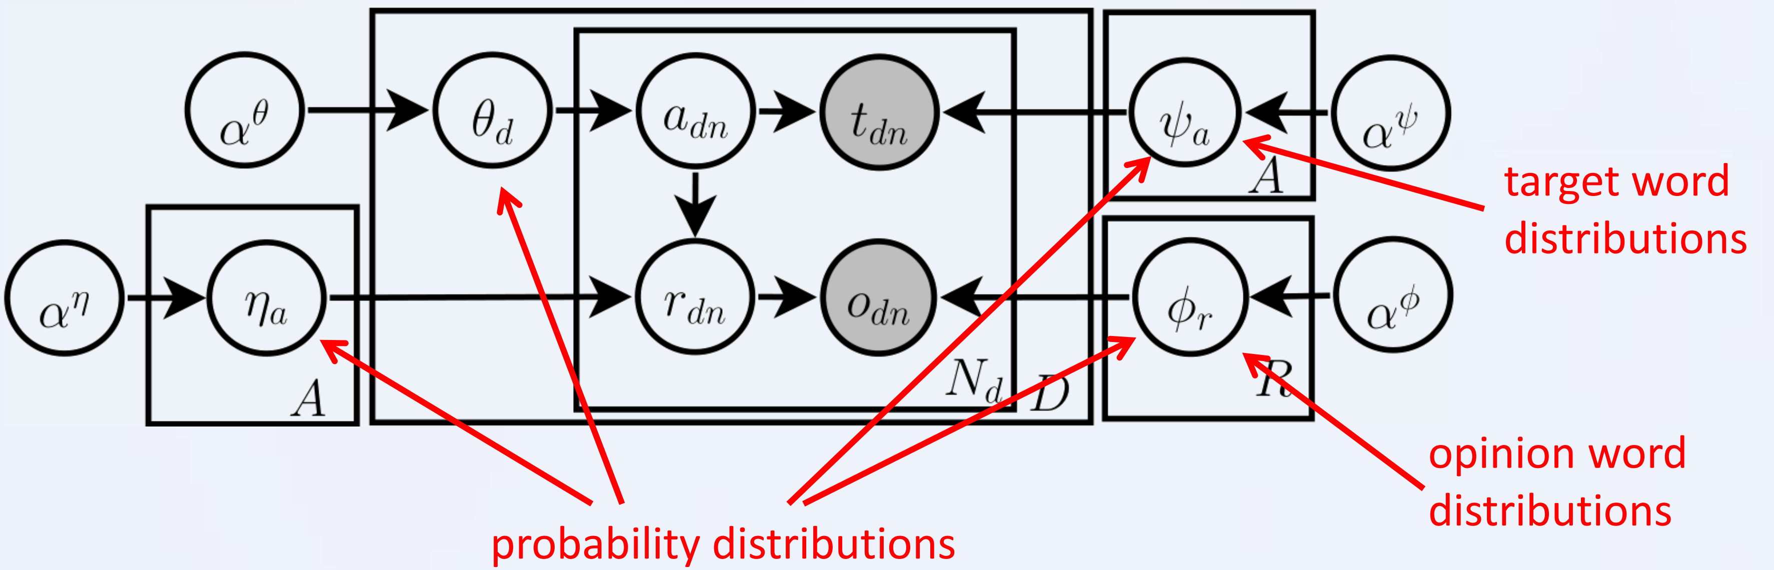

They capture the interaction between the variables, tell us about the corpus.

ILDA

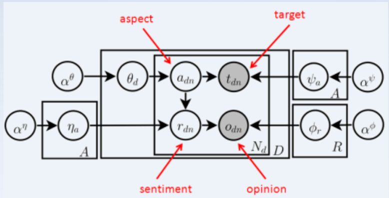

- What ILDA does:
  - Automatically groups target-opinion pairs into various aspects.
  - Learns the opinions corresponding to various sentiments.
- Limitation of ILDA:
  - Sentiment labels are arbitrary.
    - Need to manually inspect the associated opinions to know whether they are positive, neutral or negative.
  - Targets and opinions are related only via latent variables.
    - Interaction between targets and opinions are not considered.
    - The pair ‘friendly dumpling” is perfectly reasonable under ILDA.

* Picture stolen via Google search.

Opinion Aggregation on Tweets

- Why?
  - More opinions lying around, but less structured.
    - Easier to create than proper review.
  - Less targeted by fake review companies.
    - Tweets are usually written for friends and family.
- How?
  - Design Twitter Opinion Topic Model (TOTM) for Tweets.
  - Extension of ILDA but address its limitation.
  - Make use of emoticons (common in Tweets).
  - Use hashtags to aggregate Tweets.
  - Incorporate existing sentiment lexicons.

TOTM

- Graphical model:

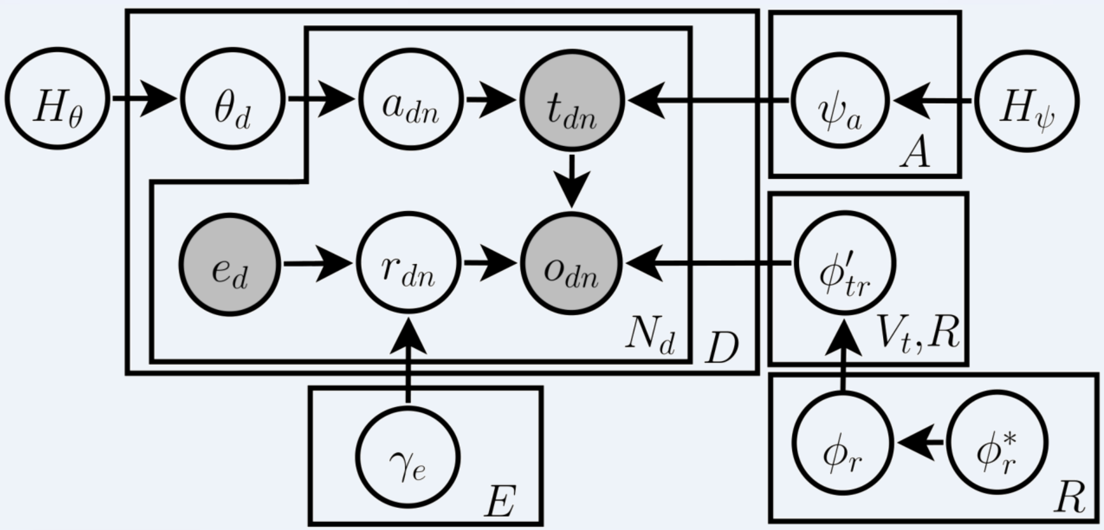

TOTM

- Graphical model:

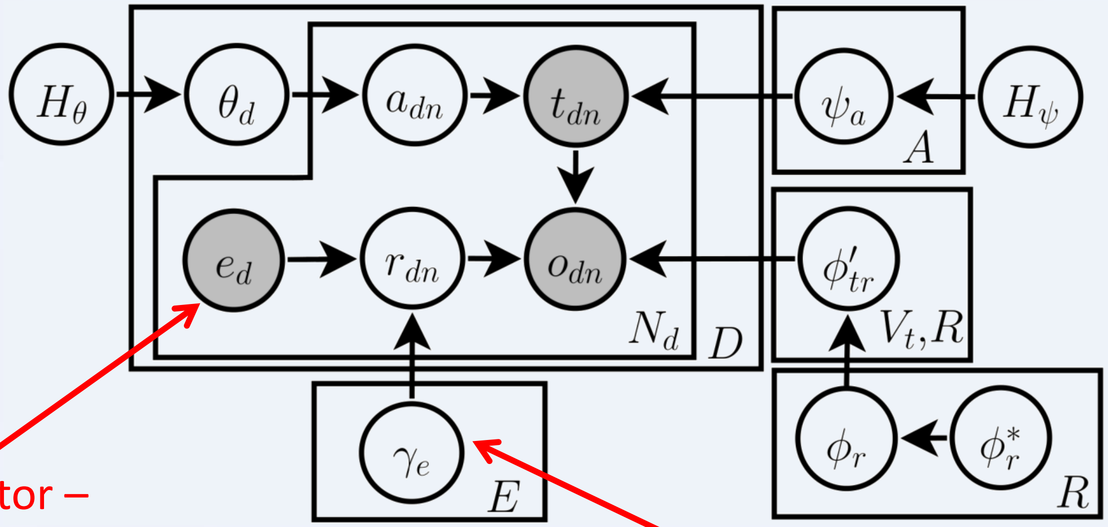

Emotion indicator – determined by seen emoticons or strong sentiment words (such as ‘happy’, ‘sad’).

Positive opinions tend to associate with positive emotions.

TOTM

- Graphical model:

Model target-opinion interaction directly.
This improves opinion prediction significantly.

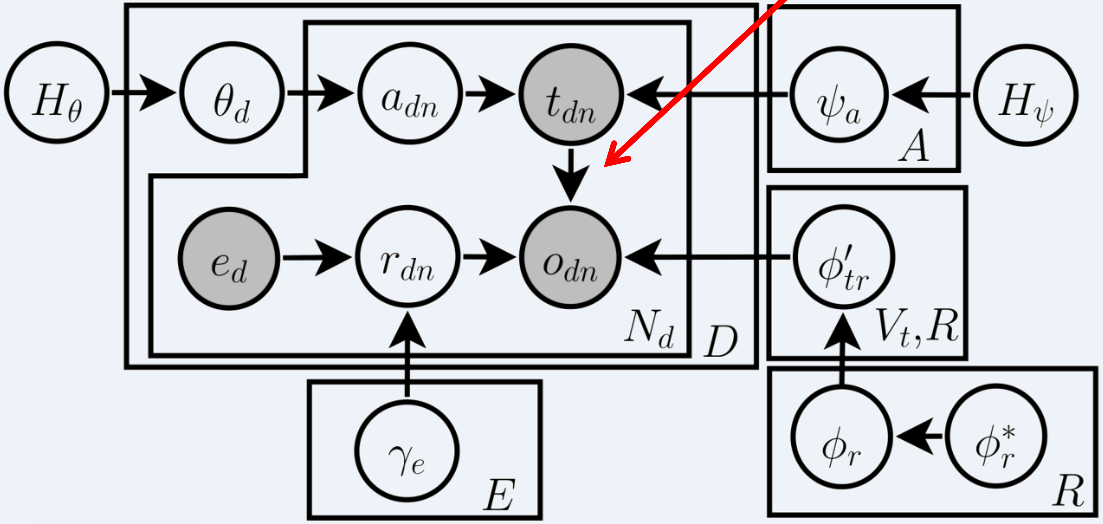

TOTM

- Graphical model:

target word distributions

opinion word distributions

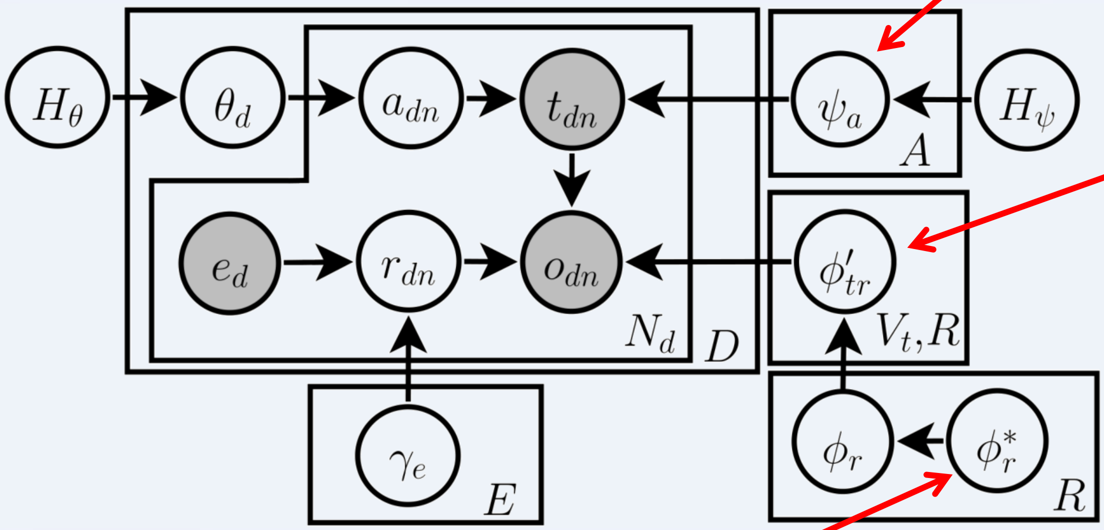

Incorporate sentiment lexicon as prior.

Incorporating Sentiment Lexicon

- Existing approach (He, 2012):
  - Rule-based system for topic models.
  - Modify the Dirichlet prior for opinion word distributions.
    - Note we have 3 opinion word distributions:
      - Positive-opinion distribution.
      - Neutral-opinion distribution.
      - Negative-opinion distribution.
    - The prior parameter is initialised as 0.33 for each opinion word of any sentiment-opinion distribution (uniform Dirichlet).
    - The prior parameter is then adjusted to 0.9 or 0.05 depending on the sentiment of a given opinion word (according to lexicon).

Incorporating Sentiment Lexicon

- Our approach:
  - Introduce a tunable parameter *b* to control the strength of sentiment prior.
  - The prior for the sentiment-opinion distribution is given by:

$$\phi^*_{rv} \propto (1+b)^{X_{rv}}$$

  - X_{rv} is the sentiment score of an opinion word determined from lexicon.
  - *b* is strictly positive, so positive X_{rv} enhances the prior while negative X_{rv} lowers the prior. (see details in paper)
  - Why exponential in the formula?
    - Ensures positivity of the priors.
    - Gives a simple learning algorithm for *b*.

Experiments

- Dataset:
  - Subset of Twitter 7 dataset (Yang & Leskovec, 2011).
    - 9 millions tweets on Electronic Products.
  - And 2 smaller corpus.
- Compare TOTM against
  - ILDA;
  - LDA-DP [Vanilla LDA but modify prior according to He (2012)].
- Evaluations:
  - Perplexity;
  - Sentiment prior evaluation;
  - Sentiment classification.

Perplexity

- Commonly used to evaluate topic models.
- Measure topic model's goodness of fit.
  - Negatively related to log likelihood so lower perplexity is better.

	Target	Opinion	Overall
LDA-DP	N/A	510.15 ± 0.08	N/A
ILDA	594.81 ± 13.61	519.84 ± 0.43	556.03 ± 6.22
TOTM	592.91 ± 13.86	$\mathbf{137.42} \pm 0.28$	$\mathbf{285.42} \pm 3.23$

Better fit for opinion words by modelling the target-opinion interaction directly.

Qualitative Analysis

- Inspect the top words from target word distributions.

Aspects (a)	Target Words (t)
Camera	camera, pictures, video camera, shots
Apple iPod	ipod, ipod touch, songs, song, music
Android phone	android, apps, app, phones, keyboard
Macbook	macbook, macbook pro, macbook air
Nintendo games	nintendo, games, game, gameboy

- Inspect the top words from opinion word distributions.

Target (t)	$+/-$	Opinions (o)
phone	$-$	**dead** damn stupid bad crazy
	$+$	**mobile smart** good great f***ing
battery life	$-$	terrible poor bad horrible **non-existence**
	$+$	good **long** great **7hr ultralong**
game	$-$	**addictive** stupid free **full addicting**
	$+$	great good awesome favorite **cat-and-mouse**
sausage	$-$	silly **argentinian cold** huge stupid
	$+$	**hot grilled** good **sweet** awesome

* Words in **bold** are more specific and can only describe certain targets.

Qualitative Analysis

- For comparison purpose, we can analyze hashtags that correspond to electronic companies such as #sony, #canon, #samsung…

Brands	Sentiment	Aspects / Targets' Opinions		
		Camera	Phone	Printer
Canon	−	*camera* → expensive small bad *lens* → prime cheap broken		*printer* → obscure violent digital *scanner* → cheap
	+	*camera* → great compact amazing *pictures* → great nice creative		*printer* → good great nice *scanner* → great fine
Sony	−	*camera* → big crappy defective *lens* → vertical cheap wide	*phone* → worst crappy shittest *battery life* → low	*printer* → stupid
	+	*photos* → great lovely amazing *camera* → good great nice	*phone* → great smart beautiful *reception* → perfect	
Samsung	−	*camera* → digital free crazy *shots* → quick wide	*phone* → stupid bad fake *battery life* → solid poor terrible	*scanner* → worst
	+	*camera* → gorgeous great cool *pics* → nice great perfect	*phone* → mobile great nice *service* → good sweet friendly	

Major Contributions

- Introduce TOTM for aspect-based opinion aggregation on Tweets.
  - Makes use of auxiliary information on Tweets.
- Novel way of incorporating sentiment prior information into topic models.
  - Simple to implement and allow automatic learning of the hyperparameter (b).

Please email Kar Wai (karwai.lim@gmail.com) if you have any questions, thank you.