Unsupervised Prediction of Citation Influences
Dietz, Bickel, Scheffer (2007)

Alex Klibisz, alex.klibisz.com, UTK STAT645

October 20, 2016

Motivation

Researchers need a bird's-eye visualization of a research area.

- Overview of ideas.
- Important publications.
- Indicates which publications significantly impact one another.
- Complements in-depth publication graphs.

Example Results

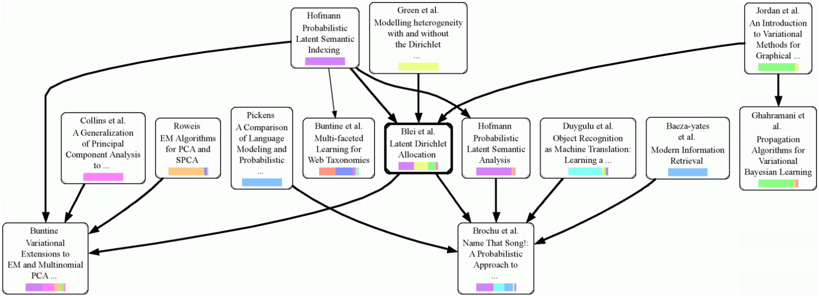

Figure 3. The filtered citation graph contains only edges which represent a significant influence.

Example Results

Table 3. Words in the abstract of the research paper "Latent Dirichlet Allocation" are assigned to citations. The probabilities in parentheses indicate $p(w, c|d, \cdot)$.

Cited Title	Associated Words	γ
Probabilistic Latent Semantic Indexing	text(0.04), latent(0.04), modeling(0.02), model(0.02), indexing(0.01), semantic(0.01), document(0.01), collections(0.01)	0.49
Modelling heterogeneity with and without the Dirichlet process	dirichlet(0.02), mixture(0.02), allocation(0.01), context(0.01), variable(0.0135), bayes(0.01), continuous(0.01), improves(0.01), model(0.01), proportions(0.01)	0.25
Introduction to Variational Methods for Graphical Methods	variational(0.01), inference(0.01), algorithms(0.01), including(0.01), each(0.01), we(0.01), via(0.01)	0.22

Problem Statement

Given

1. Universe of publications (full text or abstracts)
2. Citation graph (publications are nodes, directed edges indicate citing).

Find

1. Weights of citations that correlate to ground-truth impact:

- $\gamma_d(c)$: *impact of cited publication c on citing publication d*

Evaluate

- Ground truth is not available; results compared to expert opinion.

Steps

1. Models
   1.1 Two extensions of LDA: LDA-JS, LDA-post.
   1.2 Copycat Model.
   1.3 Citation Influence Model.
2. Evaluation
   2.1 Narrative evaluation on LDA paper.
   2.2 Predictive performance against expert-labeled influences.
   2.3 Topic differences for duplicated publications.

Related Work

- Bibliometric measures such as co-coupling as a similarity measure in digital library projects.
- Graph-based analyses such as community detection, node ranking according to authorities and hubs, link prediction.
- How paper networks evolve over time.
- Identifying latent communities via HITS or stochastic blockmodels.
- Unsupervised learning of hidden topics from text publications via pLSA and LDA.
- Community analysis via pHITS and pLSA.

To our knowledge, no one has included text and links into a probabilistic model to infer topical influences of citations.

Estimating the Influence of Citations with LDA

Two Assumptions

1. Publications with strong impact are directly cited.
2. Citing publication's topics not influenced by cited publications' topics.

Strength of Influence Heuristics

Strength of influence is not an integral part of the model , but has to be determined in a later step using a heuristic measure.

LDA-JS Model

Heuristic

- Measure compatibility between topic distributions of citing and cited publications.
- Similar topic distribution $\rightarrow$ strong influence.

Weight function

- Based on Jensen-Shannon Divergence:

$$\gamma_d(c) = \exp(-D_{JS}(\theta_d \| \theta_c)) \, , c \in L(d) \tag{1}$$

with $D_{JS}(\theta_d \| \theta_c) =$

$$\frac{1}{2} D_{\mathrm{KL}} \left(\theta_d \middle\| \frac{\theta_d + \theta_c}{2} \right) + \frac{1}{2} D_{\mathrm{KL}} \left(\theta_c \middle\| \frac{\theta_d + \theta_c}{2} \right)$$

LDA-post Model

Heuristic

- Measure $p(c|d)$, probability of a citation given a publication.
- Assumes posterior of a cited publication given a topic $p(c|t) \propto p(t|c)$.

Weight function

$$\gamma_d(c) = p(c|d) = \sum_t p(t,c|d) = \sum_t p(t|d) \cdot p(c|t) \quad (2)$$

LDA plate diagram

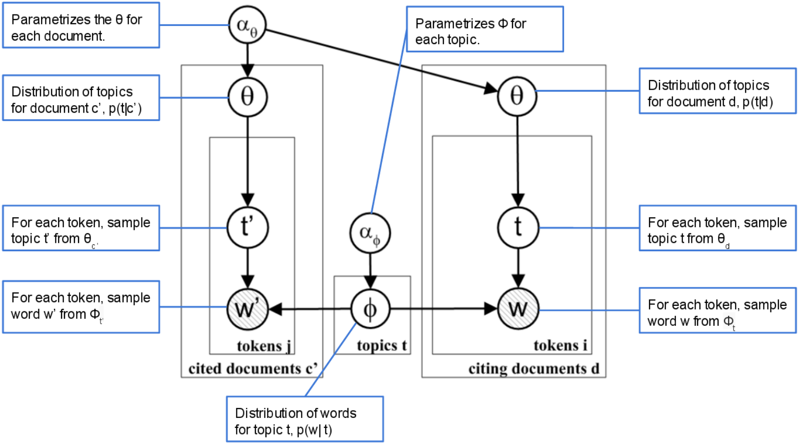

Copycat Model

Intuition

- Attribute every word in a citing publication to a topic from one of the cited publications.

Requires Bipartite Citation Graph

1. **D** nodes have outgoing links (citing).
2. **C** nodes have incoming links (cited).
   - Nodes that both cite and get cited are duplicated.

Mutual Influence of Citing Publications

- Allows associations between fields.
  - e.g. Gibbs sampling in both physics and ML.
- Creates noise, doesn't model *innovation* (all words taken from a cited publication).

Copycat Model Plate Diagram

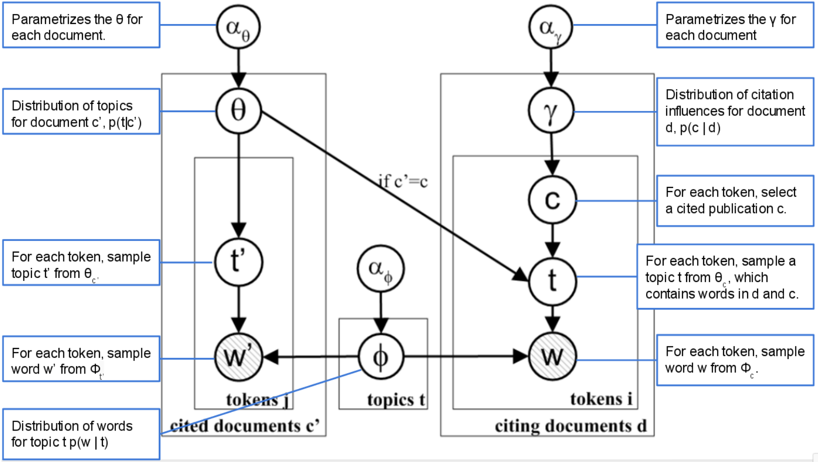

Citation Influence Model

Intuition

1. Flip an unfair coin **s** from distribution λ parameterized by α_λ.
   - If $\mathbf{s} = 0$, draw topic from a cited document's topic mixture $\theta_{c_{d,i}}$.
   - If $\mathbf{s} = 1$, draw topic from innovation topic mixture ψ_d.
2. Draw words from the selected topic.

Properties

- λ is an estimate for how well a publication fits its citations.
- $\lambda \cdot \gamma$ gives the absolute strength of influence, useful for visualizing influence.

Citation Influence Generative Process

1 • for all topics $t \in [1:T]$ do

1.1 • draw the word distribution for each latent topic $\phi_t = p(w|t) \sim dirichlet(\vec{\alpha}_\phi)$

2 • for all cited documents $c' \in C$ do

2.1 • draw a topic mixture $\theta_{c'} = p(t'|c') \sim dirichlet(\vec{\alpha}_\theta)$

2.2 • for all tokens j do

2.2.1 • draw a topic $t'_{c',j} \sim \theta_{c'}$ from the topic mixture

2.2.2 • draw a word $w_{c',j} \sim \phi_{t'_{c',j}}$ from the topic specific word distribution

3 • for all citing documents $d \in D$ do

3.1 • draw a citation mixture $\gamma_d = p(c|d)|_{L(d)} \sim dirichlet(\vec{\alpha}_\gamma)^1$ restricted to the publications c cited by this publication d

3.2 • draw an innovation topic mixture $\psi_d = p(t|d) \sim dirichlet(\vec{\alpha}_\psi)$

3.3 • draw the proportion between tokens associated with citations and those associated with the innovation topic mixture $\lambda_d = p(s = 0|d) \sim beta(\alpha_{\lambda_\theta}, \alpha_{\lambda_\psi})$

3.4 • for all tokens i do

3.4.1 • toss a coin $s_{d,i} \sim bernoulli(\lambda_d)$

3.4.2 • if $s_{d,i} = 0$

3.4.2.1 • draw a cited document $c_{d,i} \sim multi(\gamma_d)$

3.4.2.2 • draw a topic $t_{d,i} \sim multi(\theta_{c_{d,i}})$ from the cited document's topic mixture

3.4.3 • else $(s_{d,i} = 1)$

3.4.3.1 • draw the topic $t_{d,i} \sim multi(\psi_d)$ from the innovation topic mixture

3.4.4 • draw a word $w_{d,i} \sim multi(\phi_{t_{d,i}})$ from the topic specific word distribution

Citation Influence Generative Process Plate Diagram

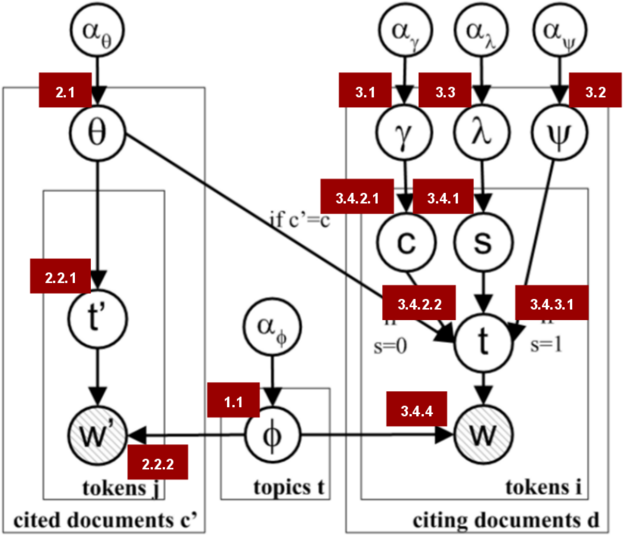

Citation Influence Plate Diagram

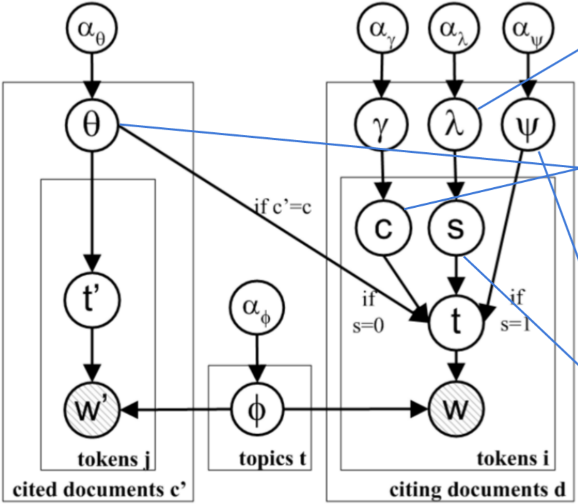

Coin parameter λ determines coin-flip distribution. λ learned by model given assymetric beta prior α_λ

For each token, draw a topic from cited document c's topic distribution, θ_c

For each token, draw a topic from topic mixture ψ_d, which models document d's own innovations.

Flipping unfair coin s determines if t is taken from θ_c or ψ_d.

Citation-influence Gibbs Sampling

Learn the model via Gibbs Sampling

- Iteratively updates each latent variable given fixed remaining variables.
- Update equations computed in constant time using count caches.
  - e.g. $C_{d,c,s}(1,2,0)$ holds the number of tokens in document 1 that are assigned to citation 2 with coin result $s = 0$.

Update equations

$$p(c_i | \vec{c}_{\neg i}, d_i, s_i = 0, t_i \cdot) \qquad (3)$$

$$p(s_i = 0 | \vec{s}_{\neg i}, d_i, c_i, t_i, \cdot) \qquad (4)$$

$$p(s_i = 1 | \vec{s}_{\neg i}, d_i, t_i \cdot) \qquad (5)$$

$$p(t_i | \vec{t}_{\neg i}, w_i, s_i = 0, c_i \cdot) \qquad (6)$$

$$p(t_i | \vec{t}_{\neg i}, w_i, d_i, s_i = 1, c_i \cdot) \qquad (7)$$

Experiments

Data

- Original LDA paper (Blei et al., 2003)
- Subset of CiteSeer

Evaluations

1. Narrative evaluation of original LDA paper
2. Prediction performance
3. Duplication of publications

Narrative Evaluation

Goal

- Check quality on a known topic and popular paper.

Method

- Consider LDA paper plus two levels of cited and citing papers.
- Fixed hyperparameters:
  - $\alpha_\phi = 0.01$, $\alpha_\theta = \alpha_\psi = 0.1$, $\alpha_{\lambda_\theta} = 3.0$, $\alpha_{\lambda_\psi} = 0.1$, $\alpha_\gamma = 1.0$
  - $T = 30$
- Only include edges with influence weight $\gamma_d(c) > 0.05$.

Narrative Evaluation

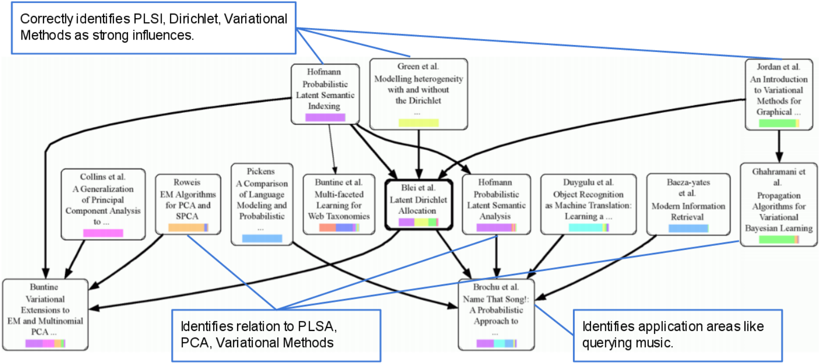

Figure 3. The filtered citation graph contains only edges which represent a significant influence.

Predictive Performance Evaluation

Goal

- Compare influence weights to expert opinions.

Method

- Include six models: 1) Citation Influence, 2) Copycat, 3) LDA-JS, 4) LDA-post, 5) PageRank of cited nodes, 6) Cosine similarity of TF-IDF vectors.
- Run models for $T = 10, 15, 30, 50$ with hyperparmeters:
  - Citation influence model: $\alpha_\phi = 0.01$, $\alpha_\theta = \alpha_\psi = 0.1$, $\alpha_{\lambda_\theta} = 3.0$, $\alpha_{\lambda_\psi} = 0.1$, $\alpha_\gamma = 1.0$
  - Copycat model: $\alpha_\phi = 0.01$, $\alpha_\theta = 0.1$, $\alpha_\gamma = 1.0$
  - LDA-JS: $\alpha_\phi = 0.01$, $\alpha_\theta = 0.1$
  - LDA-post: $\alpha_\phi = 0.01$, $\alpha_\theta = 0.1$
- Three experts label 22 seed publications and their citations - total 132 abstracts - using Likert scale.
- Predictive performance represented as Area under ROC Curve (area $= 1 \rightarrow$ perfect match).

Predictive Performance Evaluation

Results

- Citation Influence significantly better than LDA-post.
- Citation Influence has no significant improvement over Copycat.
- Copycat has no significant improvement over LDA-post.
- LDA-JS slightly below LDA-post
- LDA degenerates at $T = 30, 50$
- Copycat is significantly better than LDA-post at $T = 30, 50$
- TF-IDF and PageRank can't predict strength of influence.

Predictive Performance Evaluation

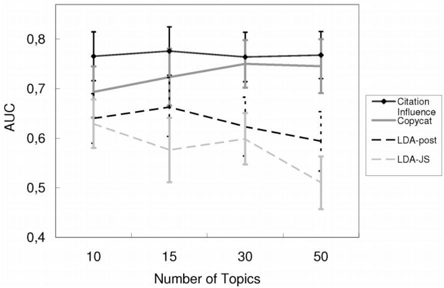

Interpretation

- Little difference between citation-influence and copycat models might indicate:
  1. Papers contained little innovation.
  2. Human judges over-attribute innovations to cited papers.

Duplicated Publications Evaluation

Goal

- Citation Influence model holds cited and citing versions of same publication independently.
- Does the model assign a similar mixture to the cited and citing instances?

Method

- Compare topic mixtures via Jensen-Shannon divergence.

Results

- Mean divergence for duplicated $= 0.07$.
- Mean divergence otherwise $= 0.69$.

Summary

Contributions

1. Copycat and citation influence models to model influence of citations in a collection of publications.
2. Practical technique for transforming data to visualize publication influence.

Questions, Critique

1. Evaluation with three experts on 132 abstracts is subjective and might lack rigor.
2. A very simple baseline might be to simply parse text and rank influence by the number of times citations (e.g. `[1]`, `[2]`, etc.) occur.