## Gregor Heinrich

CTO, vsonix GmbH, Darmstadt
http://vsonix.com
http://arbylon.net

Research Seminar, Institute for Statistics and Mathematics Vienna University of Economics and Business
http://www.wu.ac.at/statmath/resseminar
Vienna, 1 June 2012

## Topic models



- Probabilistic representations of grouped discrete data
- Illustrative for text: words grouped in documents
- Latent topics (a.k.a. concepts, components) = cluster semantically related words (Landauer and Dumais 1997; Griffiths et al. 2007)
- Language $=$ semantic meaning (topics) + noise
$\rightarrow$ Reduce vocabulary problem by discovery of semantic relations
$\rightarrow$ Reduce sparsity problem by dimensionality reduction $\leftrightarrow$ discrete principal components analysis (Buntine and Jakulin 2005)
- Topic models - motivation and review
- Networks of mixed membership (NoMMs)
- Inference - a Gibbs "meta-sampler"
- NoMM typology and design
- Application to tag-enhanced expertise finding
- Conclusions and outlook

Bayesian networks: Dirichlet-generated multinomials


## Bayesian networks:

- Graphical modelling of joint probability distributions
- Node: random variable
- Edge: conditional probability distribution
- Plate: repeated i.i.d. samples


## Example document-topic distributions

Document $m=42$ (column): Traditional machine learning relies on the availability of a large amount of data to train a model, which is then applied to test data in the same feature space. However, labeled data are often scarce and expensive to obtain...
Strongest topics: $k=\{25,21,48, \ldots\}$

transposed view: rows = topics, columns = documents
Figure: Excerpt from document-topic matrix $\vartheta(M=200, K=50)$.

## Latent Dirichlet Allocation



Draw word from term distribution of topic 2, "learning"

## Example topic-term distributions

Topic $k=21$ (row): data word feature label data scarce obtain. Topic $k=25$ (row): machine learning train model test feature space.. Topic $k=48$ (row): computing support grant project system method..


Figure: Excerpt from topic-term matrix $\varphi(V=200, K=50)$.

## Example: Text mining for semantic clusters

| Topic label | Most likely terms according to $\varphi_{k, t}=p$ (wordltopic) |
| :--- | :--- |
| Politische Parteien | CDU Partei Kohl Aufklärung Schäuble Zeitung Union Krise Wahrheit Affäre Christ- <br> demokraten Glaubwürdigkeit Konsequenzen |
| Bundesliga | FC SC München Borussia SV VfL Kickers SpVgg Uhr Köln Bochum Freiburg VfB <br> Eintracht Bayern Hamburger Bayern+München |
| Polizei / Unfall | Polizei verletzt schwer Auto Unfall Fahrer Angaben schwer+verletzt Menschen Wa- <br> gen Verletzungen Lawine Mann vier Meter Straße |
| Tschetschenien | Rebellen russischen Grosny russische Tschetschenien Truppen Kaukasus Moskau <br> Angaben Interfax tschetschenischen Agentur |
| Politik / Hessen | FDP Koch Hessen CDU Koalition Gerhardt Wagner Liberalen hessischen Wester-- <br> welle Wolfgang Roland+Koch Wolfgang+Gerhardt |
| Wetter | Grad Temperaturen Regen Schnee Süden Norden Sonne Wetter Wolken Deutsch- <br> land zwischen Nacht Wetterdienst Wind |
| Politik / Kroatien | Parlament Partei Stimmen Mehrheit Wahlen Wahl Opposition Kroatien Präsident <br> Parlamentswahlen Mesic Abstimmung HDZ |
| Die Grünen | Grünen Parteitag Atomausstieg Tritin Grüne Partei Trennung Mandat Ausstieg Amt <br> Roestel Jahren Müller Radcke Koalition |
| Russische Politik | Russland Putin Moskau russischen russische Jelzin Wladimir Tschetschenien Rus- <br> slands Wladimir+Putin Kreml Boris Präsidenten |
| Polizei / Schulen | Polizei Schulen Schüler Täter Polizisten Schule Tat Lehrer erschossen Beamten <br> Mann Polizist Beamte verletzt Waffe |

Bigram LDA topics, 18400 German news messages, Jan. 2000 (Heinrich et al. 2005)
Gregor Heinrich
A generic approach to topic models

## Typical derivation method (Is it really that complex?)


(e) Expert-tag-topic model (ETT) (Heinrich 2011)


```
    M
```



Topic models: Example structures

(a) Latent Dirichlet allocation (LDA)
(a) Latent Dirichlet allocation (LDA)

(c) Pachinko allocation model (PAM4)

(b) Author-topic model (ATM)

(d) Hierarchical PAM (hPAM)
(Blei et al. 2003; Rosen-Zvi et al. 2004; Li and McCallum 2006; Li et al. 2007)

## Topic models - bottom line

- Expanding research field with practical relevance
- No existing analysis as generic model class
$\rightarrow$ Conjecture:
- Important properties generic across models
- Simplifications in the derivation of model properties, inference algorithms and design methods
- Topic models - motivation and review
- Networks of mixed membership (NoMMs)
- Inference - a Gibbs "meta-sampler"
- NoMM typology and design
- Application to tag-enhanced expertise finding
- Conclusions and outlook

Generic topic models - "NoMMs"


- Generic characteristics of topic models:
- Levels with multinomial components, generated from Dirichlet
- Coupling via values of discrete variables
$\rightarrow$ "Network of mixed membership" (NoMM): directed graph
- Compact, domain-specific alternative to Bayesian network
- Node: Sample from mixture component, selected via incoming edges

- Terminal node: observation
- Edge: Propagation of discrete values to children


## Topic models in NoMM representation

$$
\begin{aligned}
& \left.\left.\bigcirc \xrightarrow[{[M}]]{m} \xrightarrow[\vec{\vartheta}_{m} \mid \alpha]{z_{m, n}=k} \xrightarrow[{[K}]\right]{\vec{\beta}_{k} \mid \eta} \xrightarrow[{[V}]\right]{w_{m, n}=t} \\
& \text { (a) Latent Dirichlet allocation (LDA) } \\
& \left.\left.\bigcirc \xrightarrow[{[M}]]{m} \xrightarrow[\vec{a}_{m}]{x_{m, n}=x} \xrightarrow[{[A}]\right]{\vec{\vartheta}_{x} \mid \alpha} \xrightarrow[{[K}]\right]{z_{m, n}=k} \xrightarrow[\vec{\varphi}_{k} \mid \beta]{\substack{w_{m, n}=t}} \xrightarrow{w_{[ }} \\
& \text {(b) Author-topic model (ATM) }
\end{aligned}
$$

> (c) Pachinko allocation model (PAM4)
> (d) Hierarchical pachinko allocation model (hPAM)

## Example NoMM generative process: PAM4



## Bayesian inference problem

- Bayesian inference: "Reverse generative process"
- Estimate (distributions over) parameters $\Theta$ and latent variables ("topics") $H$ given observations $V$ and hyperparameters $A$
$\rightarrow$ Find posterior distribution $p(H, \Theta \mid V, A) \rightarrow$ exponential complexity

- Topic models - motivation and review
- Networks of mixed membership (NoMMs)
- Inference - a Gibbs "meta-sampler"
- NoMM typology and design
- Application to tag-enhanced expertise finding
- Conclusions and outlook


## Collapsed Gibbs sampling



- Collapsed Gibbs sampling: stochastic EM / MCMC:
- NoMMs: parameters $\Theta$ correlated with $H \rightarrow$ integrated out
- For each data token $i$ : Sample latent variables $H_{i}=\left(y_{i}, z_{i}, \ldots\right)$, given all other data, latent $H_{\neg i}$ and visible $V$ :

$$
\begin{equation*}
H_{i} \sim p\left(H_{i} \mid H_{-i}, V, A\right) . \tag{1}
\end{equation*}
$$

- Stationary state: full conditional simulates posterior
- Faster absolute convergence for NoMMs than, e.g., variational Bayes (Heinrich and Goesele 2009)


## Collapsed Gibbs full conditionals



- NoMM full conditionals can be generically derived (Heinrich 2009)
- Typical case leads to weights with straight-forward factor structure:

$$
\begin{equation*}
p\left(H_{i} \mid H_{\neg i}, V, A\right) \propto \prod_{\ell}\left[\frac{n_{k, t}^{\neg i}+\alpha}{n_{k}^{\neg i}+T \alpha}\right]^{[\ell]} \tag{2}
\end{equation*}
$$

- $n_{k, t}=$ count of co-occurrences between input and output values of a NoMM level $\ell$
- More generally: $p\left(H_{i} \mid \cdot\right) \propto \prod_{\ell}[q(k, t)]^{[\ell]}$ with $t=$ set of values/edges


## Implementation: Gibbs "meta-sampler"



- Code generator for topic models in Java and C
- Separation of knowledge domains: topic model applications vs. machine learning vs. computing architecture


## $q$-functions and Pólya urn



Figure: Pólya urn and multinomial parameters.

$$
\begin{aligned}
q(k, t) \triangleq \frac{\mathrm{B}\left(\vec{n}_{k}+\alpha\right)}{\mathrm{B}\left(\vec{n}_{k}^{\urcorner i}+\alpha\right)} & \stackrel{|t|=1}{=} \frac{n_{k, t}^{\neg t_{i}}+\alpha}{n_{k}^{\neg t_{i}}+T \alpha}=\text { smoothed ratio of co-occurrence counts } \\
& \stackrel{t=\{u, v\}}{=} \frac{n_{k, u}^{\neg u_{i}}+\alpha}{n_{k}^{\neg u_{i}}+T \alpha} \cdot \frac{n_{k, v}^{\neg v_{i}}+\alpha+\delta(u-v)}{n_{k}^{\neg v_{i}}+T \alpha+1} \triangleq q(k, u \oplus v)
\end{aligned}
$$

$$
\ldots
$$

Example NoMM script and generated kernel: hPAM2


## Example document-topic distributions



Figure: Excerpt from document-topic matrix $\vartheta(M=200, K=50)$.

## Overview

- Topic models - motivation and review
- Networks of mixed membership (NoMMs)
- Inference - a Gibbs "meta-sampler"
- NoMM typology and design
- Application to tag-enhanced expertise finding
- Conclusions and outlook


## Fast sampling: hybrid acceleration methods

Serial:

- Exploit saliency of few weights, e.g., generalising (Porteous et al. 2008): compute only few weights on average + estimate normalisation term
- Complex data structures, especially for larger models


## Parallel:

- Distribute local parameters (document-specific etc.)
- Need to sync global parameters: different methods, e.g., generalising (Newman et al. 2009)
- Occupancy: balance communication and computation (architecture-spec.
Independence assumption:
- Reduce complexity: $\prod_{\ell} T^{\ell} \gg \sum_{\ell} T^{\ell}$

| method | model | parameters | speedup (it | onverge) |
| :---: | :---: | :---: | :---: | :---: |
| P4 | LDA | $K=100$ | 6.3 |  |
| S×P4 | LDA | $K=500$ | 30.2 |  |
| I | PAM4 | $K, L=40,40$ | 21.8 | 7.4 |
| P4×I | PAM4 | $K, L=40,40$ | 78.7 | 24.1 |
| S×P4×I | PAM4 | $K, L=40,40$ | 163.2 | 49. |
| S×P4×I | PAM4 | $K, L=20,100$ | 143.6 | 43. |

$\qquad$


## $q$-functions and Pólya urn revisited



Figure: Pólya urn and multinomial parameters.
$q(k, t) \triangleq \frac{\mathrm{B}\left(\vec{n}_{k}+\alpha\right)}{\mathrm{B}\left(\vec{n}_{k}^{\neg i}+\alpha\right)} \stackrel{|t|=1}{=} \frac{n_{k, t}^{\neg t_{i}}+\alpha}{n_{k}^{\neg t_{i}}+T \alpha}=$ smoothed ratio of co-occurrence counts

$$
\stackrel{t=\{u, v\}}{=} \frac{n_{k, u}^{\neg u_{i}}+\alpha}{n_{k}^{\neg u_{i}}+T \alpha} \cdot \frac{n_{k, v}^{\neg v_{i}}+\alpha+\delta(u-v)}{n_{k}^{\neg v_{i}}+T \alpha+1} \triangleq q(k, u \oplus v)
$$

## NoMM sub-structure typology

N1. Dirichlet-multinomial parameters E2. Autonomous edges

$$
q(a, z) q(z, b) \quad q(a, x \oplus y) q(x, b) q(y, c) \quad q(a, x) q(b, y) q(k, c), k=f(x, y)
$$

N2. Observed parameters E3. Coupled edges C3. Interleaved indices

Gibbs full conditional assembled via:

$$
\begin{equation*}
p\left(H_{i} \mid \cdot\right) \propto \prod_{\ell}[q(k, t)]^{\ell} \tag{3}
\end{equation*}
$$

Towards a design process


Figure: NoMM design process.


## Overview

- Topic models - motivation and review
- Networks of mixed membership (NoMMs)
- Inference - a Gibbs "meta-sampler"
- NoMM typology and design
- Application to tag-enhanced expertise finding
- Conclusions and outlook

$$
\begin{aligned}
& \vartheta_{a, z}^{\mathrm{c}} q(z, b) \quad q(a, z) q(z, b) q(z, c) \quad \approx q\left(a, z^{1}\right) q\left(b, z^{2}\right) q\left(z^{1}, c \oplus \tilde{c}\right) q\left(z^{2}, \tilde{c} \oplus c\right)
\end{aligned}
$$

## Define tasks + metrics; set up terminals

- Retrieval of experts $a$ for term queries $\vec{w}$ and tag queries $\vec{c}$ : query likelihood model: $p(\vec{w} \mid a)$ and $p(\vec{c} \mid a) \rightarrow$ measure retrieval precision
- Topic quality $\rightarrow$ measure coherence score
- Baseline: Author-topic model ATM (Rosen-Zvi et al. 2004), LDA (Blei et al. 2003)


Figure: Model design: Terminals

解 scientific community of Neural Information Processing Systems (NIPS) conference

Tags: probabilistic methods, variational inference, learning algorithms


## Modelling assumptions

## Compose model




$p(\ldots \mid \vec{a}, \vec{w}, \vec{c}) \propto \ldots$

Starting from terminals
(a) Expertise of authors weighted by the portion of authorship $a_{m, a}$.
(b) Expertise semantics expressed by topics $z$. Each author has a single field of expertise (topic distribution).
(c) Tag semantics expressed by topics $y$. Tag topics $y$ could be $\equiv z$.


$$
p(x, \ldots \mid \cdot) \propto a_{m, x} q(x, \ldots) \ldots
$$

## Up-stream evidence $\vec{a}_{m}$

$\rightarrow$ observed parameter node samples word author $x$


$$
p(x, z, \ldots \mid \cdot) \propto a_{m, x} q(x, z) \ldots
$$

Each author only one field of expertise (topic distribution) $\rightarrow q$-term $q(x, z)$ assigns topics to sampled author $x$ (cf. ATM)

## Compose model


$p(x, z \mid \cdot) \propto a_{m, x} q(x, z) q(z, w) q(z, c)$
Incorporate tags via $q(z, c)$ conditioned on the same topic
$\rightarrow$ Problem: How to determine tag $c_{m, n}$ for word?

## Compose model



$$
p(x, z, y \mid \cdot) \propto a_{m, x} q(x, z \oplus y) q(z, w) q(y, c)
$$

$\rightarrow$ Incorporate tag topics $y_{m, j}$ on separate sequence ( $m, j$ )
$\rightarrow$ Tag boosting: adjust tag influence via tag sequence length $J_{m}$

## ETT1 model



Assembled $q$-terms:

$$
\begin{equation*}
p(x, z, y \mid \cdot) \propto a_{m, x} q(x, z \oplus y) q(z, w) q(y, c) \tag{4}
\end{equation*}
$$

Easy expansion to standard Gibbs full conditionals:

$$
\begin{align*}
& p\left(x_{m, n}=x, z_{m, n}=z \mid \cdot\right) \propto a_{m, x} \cdot \frac{n_{x, z}^{\neg\{x, z\}_{m, n}}+\alpha}{n_{x}^{\neg}\{x,\}_{m, n}+K \alpha} \cdot \frac{n_{z, w} \frac{\neg z_{m, n}}{}+\beta}{n_{z}^{\neg z_{m, n}}+V \beta}  \tag{5}\\
& p\left(x_{m, j}=x, y_{m, j}=y \mid \cdot\right) \propto a_{m, x} \cdot \frac{n_{x, y}^{\tau\{x, y\}_{m, j}}+\alpha}{n_{y}^{\neg\{x, y\}_{m, j}}+K \alpha} \cdot \frac{n_{y, c}^{\neg y_{m, j}}+\gamma}{n_{y}^{\neg y_{m, j}}+C \gamma} \tag{6}
\end{align*}
$$

Retrieval via query likelihood model:

$$
\begin{equation*}
p(\vec{w} \mid a)=\prod_{w \in \vec{w}} \sum_{z} \vartheta_{a, z} \varphi_{z, w} \quad p(\vec{c} \mid a)=\prod_{c \in \vec{c}} \sum_{y} \vartheta_{a, y} \psi_{y, c} \tag{7}
\end{equation*}
$$

## ETT1 model



Assembled $q$-terms:

$$
\begin{equation*}
p(x, z, y \mid \cdot) \propto a_{m, x} q(x, z \oplus y) q(z, w) q(y, c) \tag{4}
\end{equation*}
$$

Easy expansion to standard Gibbs full conditionals:

$$
\begin{align*}
& p\left(x_{m, n}=x, z_{m, n}=z \mid \cdot\right) \propto a_{m, x} \cdot \frac{n_{x, z}^{\neg\{x, z\}_{m, n}}+\alpha}{n_{x}^{\neg\{x, z\}_{m, n}}+K \alpha} \cdot \frac{n_{z, w}^{\neg z_{m, n}}+\beta}{n_{z}^{\neg z_{m, n}}+V \beta}  \tag{5}\\
& p\left(x_{m, j}=x, y_{m, j}=y \mid \cdot\right) \propto a_{m, x} \cdot \frac{n_{x, y}^{\neg\{x, y\}_{m, j}}+\alpha}{n_{y}^{\neg\{x, y\}_{m, j}}+K \alpha} \cdot \frac{n_{y, c} \neg y_{m, j}}{n_{y} y_{m, j}+\gamma}  \tag{6}\\
& n_{y}+C \gamma
\end{align*}
$$

Retrieval via query likelihood model:

$$
\begin{equation*}
p(\vec{w} \mid a)=\prod_{w \in \vec{w}} \sum_{z} \vartheta_{a, z} \varphi_{z, w} \quad p(\vec{c} \mid a)=\prod_{c \in \vec{c}} \sum_{y} \vartheta_{a, y} \psi_{y, c} \tag{7}
\end{equation*}
$$

Gregor Heinrich A generic approach to topic models ..... 37/50
Typical derivation method (Is it really that complex?)
(E.10)
In

$$
=\int
$$

$$
\iint_{\eta_{n}}^{\|}
$$

## Model evaluation

Truncated average precision


Figure: ETT1 example query in community browser.
Gregor Heinrich A generic approach to topic models

## Retrieval and clustering results


(a) Term queries

(b) Tag queries

- Term retrieval improved by tag influence during training time
- Mutual information between a-priori tag clusterings $p(c \mid a)$ and topic clusterings $p(z \mid a)$ : ETT1 $\geq 1.002$ vs. ATM $=0.865$.
- Semi-supervised features: find relevant items with missing tags
- Tag strength: bias towards strong tags in combinations


## Conclusions

- Networks of mixed membership:

Domain-specific compact representation

- Inference:
- Generic Gibbs sampling: $q$-functions as central quantity in model behaviour
- Gibbs meta-sampler: simplify implementation
- Hybrid acceleration methods
- Alternatives: variational Bayes (Heinrich and Goesele 2009), collapsed VB
- Typology and design method:
- Model structure types: literature + novel
- Building blocks for design with predictable properties
- Application:
- Expert-tag-topic model demonstrates design

- Application to tag-enhanced expertise finding
- Conclusions and outlook


## Towards an R-based Gibbs meta-sampler

- R environment becoming popular for topic models, e.g.:
- topicmodels package implementing general and various special cases (Grün and Hornik 2011), based on text mining package tm
- lda package with LDA, supervised, relational topic models (Blei et al. 2003; Blei and McAuliffe 2007; Chang and Blei 2009)
- Vision: Use Gibbs meta-sampler as front-end to create R-based high-performance code $\leftrightarrow$ use $R$ as experimental front-end

- Extend to non-parametric distributions, e.g., based on DPpackage (Jara et al. 2012):
- NoMMs as polymorphism of parametric and non-parametric models (with different Bayesian networks)


## References I

## References II

## References

Bei, D. and J. McAuliffe (2007)
In Advances in Neural Information Processing Systems.
Blei, D., A. Ng, and M. Jordan (2003, January).
Latent Dirichlet allocation.
Journal of Machine Learning Research 3, 993-1022.
Buntine, W. and A. Jakulin (2005).
Discrete principal components analysis.
In Proc. ECML.
Chang, J. and D. M. Blei (2009).
Relational topic models for document networks.
n AISTATS.
Chang, J., J. Boyd-Graber, S. Gerrish, C. Wang, and D. Blei (2009)
Reading tea leaves: How humans interpret topic models.
In Proc. Neural Information Processing Systems (NIPS).
Griffiths, T. L., J. B. Tenenbaum, and M. Steyvers (2007)
Topics in semantic representation
Psychological Review 114(2), 211-244


Jara, A., T. Hanson, F. A. Quintana, P. Mueller, and G. L. Rosner (2012, Feb.) DPpackage: Bayesian nonparametric modeling in R.
software documentation
andauer, T. K. and S. T. Dumais (1997)
Solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge.
Psych. Rev. 104(2), 211-240
Cognitive view on LSA.
Li, W., D. Blei, and A. McCallum (2007).
Mixtures of hierarchical topics with pachinko allocation.
In International Conference on Machine Learning.
Li, W. and A. McCallum (2006).
Pachinko allocation: DAG-structured mixture models of topic correlations
n ICML '06: Proceedings of the 23rd international conference on Machine learning, New York, NY, USA, pp. 577-584. ACM.

Mimno, D., H. M. Wallach, E. Talley, M. Leenders, and A. McCallum (2011, July)
Optimizing semantic coherence in topic models.
In Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing, Edinburgh, UK, pp. 262272.

Grün, B. and K. Hornik (2011)
opicmodels: An R package for fitting topic models.
Journal of Statistical Software 43(13)
Heinrich, G. (2009).
A generic approach to topic models.
In Proc. European Conf. on Mach. Learn. / Principles and Pract. of Know. Discov. in Databases (ECML/PKDD), Part 1, pp. 517-532
Heinrich, G. (2011)
Typology of mixed-membership models: Towards a design method
In Proc. European Conf. on Mach. Learn. / Principles and Pract. of Know. Discov. in Databases (ECML/PKDD).
Heinrich, G. and M. Goesele (2009).
Variational Bayes for generic topic models
In Proc. 32nd Annual German Conference on Artificial Intelligence (KI2009),
Heinrich, G., J. Kindermann, C. Lauth, G. Paaß, and J. Sanchez-Monzon (2005)
nvestigating word correlation at different scopes - a latent concept approach
In Workshop Lexical Ontology Learning at Int. Conf. Mach. Learning.

Newman, D., A. Asuncion, P. Smyth, and M. Welling (2009, August) Distributed algorithms for topic models.
JMLR 10, 1801-1828.
Porteous, I., D. Newman, A. Ihler, A. Asuncion, P. Smyth, and M. Welling (2008). Fast collapsed Gibbs sampling for latent Dirichlet allocation.
In KDD '08: Proceeding of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining, New York, NY, USA, pp. 569-577. ACM.

Rosen-Zvi, M., T. Griffiths, M. Steyvers, and P. Smyth (2004).
The author-topic model for authors and documents.
In Proc. 20th Conference on Uncertainty in Artificial Intelligence (UAI)

