# Bayesian nonparametric models for bipartite graphs 

François Caron<br>Department of Statistics, Oxford<br>Statistics Colloquium, Harvard University<br>November 11, 2013



## Bipartite networks



- Scientists authoring papers
- Readers reading books
- Internet users posting messages on forums
- Customers buying items
- Objects sharing a set of features


## Bipartite networks



- Scientists authoring papers
- Readers reading books
- Internet users posting messages on forums
- Customers buying items
- Objects sharing a set of features


## Bipartite networks

## Readers

## Books



- Scientists authoring papers
- Readers reading books
- Internet users posting messages on forums
- Customers buying items
- Objects sharing a set of features


## Bipartite networks

## Readers

## Books



- Scientists authoring papers
- Readers reading books
- Internet users posting messages on forums
- Customers buying items
- Objects sharing a set of features


## Bipartite networks

## Readers

## Books



- Scientists authoring papers
- Readers reading books
- Internet users posting messages on forums
- Customers buying items
- Objects sharing a set of features


## Bipartite networks

Readers

## Books



- Scientists authoring papers
- Readers reading books
- Internet users posting messages on forums
- Customers buying items
- Objects sharing a set of features


## Bipartite networks

## Readers

Books


- Scientists authoring papers
- Readers reading books
- Internet users posting messages on forums
- Customers buying items
- Objects sharing a set of features


## Book-crossing community network

5000 readers, 36000 books, 50000 edges
Readers


Books

## Book-crossing community network

Degree distributions on log-log scale

(a) Readers

(b) Books

## Statistical network models

- Statistics literature
- Exponential random graph, stochastic block-models, Rasch models, etc
- Do not capture power-law behavior
- Inference do not scale well with the number of nodes
- Physics literature
- Preferential attachment
- Lacks interpretable parameters, non-exchangeability


## Bayesian nonparametrics

- Parameter of interest is infinite-dimensional
- Allows the complexity of the model to adapt to the data
- Dirichlet Process Mixtures: Clustering/density estimation with unknown number of modes
- Attractive power-law properties
- Language modeling, image segmentation


## BNP for networks

- Models with some latent structure (e.g. infinite relational model)
- Number of nodes is fixed and dimension of the latent structure unknown
- Here: Infinite number of nodes
- (stable) Beta-Bernoulli/Indian Buffet Process
- Can capture power-law degree distributions for books
- Poisson degree distribution for readers


## Bipartite networks

## Aims

- Bayesian nonparametric model for bipartite networks with a potentially infinite number of nodes of each type
- Each node is modelled using a positive rating parameter that represents its ability to connect to other nodes
- Captures power-law behavior
- Simple generative model for network growth
- Develop efficient computational procedure for posterior simulation.


## Hierarchical model

- Represent a bipartite network by a collection of atomic measures $Z_{i}$, $i=1,2, \ldots$ such that

$$
Z_{i}=\sum_{j=1}^{\infty} z_{i j} \delta_{\theta_{j}}
$$

- $z_{i j}=1$ if reader $i$ has read book $j, 0$ otherwise
- $\left\{\theta_{j}\right\}$ is the set of books
- Each book $j$ is assigned a positive "popularity" parameter $w_{j}$
- Each reader $i$ is assigned a positive "interest in reading" parameter $\gamma_{i}$
- The probability that reader $i$ reads book $j$ is

$$
P\left(z_{i j}=1 \mid \gamma_{i}, w_{j}\right)=1-\exp \left(-w_{j} \gamma_{i}\right)
$$

## Hierarchical model

- Represent a bipartite network by a collection of atomic measures $Z_{i}$, $i=1,2, \ldots$ such that

$$
Z_{i}=\sum_{j=1}^{\infty} z_{i j} \delta_{\theta_{j}}
$$

- $z_{i j}=1$ if reader $i$ has read book $j, 0$ otherwise
- $\left\{\theta_{j}\right\}$ is the set of books
- Each book $\boldsymbol{j}$ is assigned a positive "popularity" parameter $\boldsymbol{w}_{\boldsymbol{j}}$
- Each reader $i$ is assigned a positive
'interest in reading"
- The probability that reader $i$ reads book $j$ is

$$
\boldsymbol{P}\left(z_{i j}=1 \mid \gamma_{i}, w_{j}\right)=1-\exp \left(-w_{j} \gamma_{i}\right)
$$

## Hierarchical model

- Represent a bipartite network by a collection of atomic measures $Z_{i}$, $i=1,2, \ldots$ such that

$$
Z_{i}=\sum_{j=1}^{\infty} z_{i j} \delta_{\theta_{j}}
$$

- $z_{i j}=1$ if reader $i$ has read book $j, 0$ otherwise
- $\left\{\theta_{j}\right\}$ is the set of books
- Each book $\boldsymbol{j}$ is assigned a positive "popularity" parameter $\boldsymbol{w}_{\boldsymbol{j}}$
- Each reader $\boldsymbol{i}$ is assigned a positive "interest in reading" parameter $\gamma_{i}$
- The probability that reader $i$ reads book $j$ is

$$
P\left(z_{i j}=1 \mid \gamma_{i}, w_{j}\right)=1-\exp \left(-w_{j} \gamma_{i}\right)
$$

## Hierarchical model

- Represent a bipartite network by a collection of atomic measures $Z_{i}$, $i=1,2, \ldots$ such that

$$
Z_{i}=\sum_{j=1}^{\infty} z_{i j} \delta_{\theta_{j}}
$$

- $z_{i j}=1$ if reader $i$ has read book $j, 0$ otherwise
- $\left\{\theta_{j}\right\}$ is the set of books
- Each book $\boldsymbol{j}$ is assigned a positive "popularity" parameter $\boldsymbol{w}_{\boldsymbol{j}}$
- Each reader $\boldsymbol{i}$ is assigned a positive "interest in reading" parameter $\gamma_{i}$
- The probability that reader $i$ reads book $j$ is

$$
P\left(z_{i j}=1 \mid \gamma_{i}, w_{j}\right)=1-\exp \left(-w_{j} \gamma_{i}\right)
$$

## Data Augmentation

- Latent variable formulation
- Latent scores $s_{i j} \sim \operatorname{Gumbel}\left(\log \left(w_{j}\right), \mathbf{1}\right)$
- All books with a score above $-\log \left(\gamma_{i}\right)$ are retained, others are discarded



## Model for the book popularity parameters

- Random atomic measure

$$
G=\sum_{j=1}^{\infty} w_{j} \delta_{\theta_{j}}
$$

- Construction: two-dimensional Poisson process $N=\left\{w_{j}, \theta_{j}\right\}_{j=1, \ldots}$
- Completely Random Measure $\boldsymbol{G} \sim \operatorname{CRM}(\boldsymbol{\lambda}, \boldsymbol{h})$ characterized by a Lévy measure $\boldsymbol{\lambda}(\boldsymbol{w}) \boldsymbol{h}(\boldsymbol{\theta}) \boldsymbol{d} \boldsymbol{w} \boldsymbol{d} \boldsymbol{\theta}$

$$
\int_{0}^{\infty}\left(1-e^{-w}\right) \lambda(w) d w<\infty \quad \Rightarrow \text { finite total } \sum_{j=1}^{\infty} z_{i j}
$$

[Kingman, 1967, Regazzini et al., 2003, Lijoi and Prünster, 2010]

## Posterior characterization

- Observed bipartite network $Z_{1}, \ldots, Z_{n}$
- $\boldsymbol{n}$ readers and $\boldsymbol{K}$ books with degree at least one
- Cannot derive directly the conditional of $G$ given $Z_{1}, \ldots, Z_{n}$ nor the predictive of $Z_{n+1}$ given $Z_{1}, \ldots, Z_{n}$
- Let

$$
\boldsymbol{X}_{i}=\sum_{j=1}^{\infty} x_{i j} \delta_{\theta_{j}}
$$

where $\boldsymbol{x}_{i j}=\max \left(0, s_{i j}+\log \left(\gamma_{i}\right)\right) \geq 0$ are latent positive scores.



## Posterior Characterization

The conditional distribution of $\boldsymbol{G}$ given $\boldsymbol{X}_{1}, \ldots \boldsymbol{X}_{\boldsymbol{n}}$ can be expressed as

$$
G=G^{*}+\sum_{j=1}^{K} w_{j} \delta_{\theta_{j}}
$$

where $G^{*}$ and $\left(w_{j}\right)$ are mutually independent with

$$
G^{*} \sim \operatorname{CRM}\left(\lambda^{*}, h\right), \quad \lambda^{*}(w)=\lambda(w) \exp \left(-w \sum_{i=1}^{n} \gamma_{i}\right)
$$

and the masses are

$$
\boldsymbol{P}\left(w_{j} \mid \text { other }\right) \propto \lambda\left(w_{j}\right) w_{j}^{m_{j}} \exp \left(-w_{j} \sum_{i=1}^{n} \gamma_{i} e^{-x_{i j}}\right)
$$

Characterization related to that for normalized random measures [Prünster, 2002, James, 2002, James et al., 2009]

## Generative Process for network growth

Predictive distribution of $\boldsymbol{Z}_{n+1}$ given the latent process $\boldsymbol{X}_{1}, \ldots, \boldsymbol{X}_{n}$
Books

Reader 1

## Generative Process for network growth

Predictive distribution of $\boldsymbol{Z}_{n+1}$ given the latent process $\boldsymbol{X}_{1}, \ldots, \boldsymbol{X}_{n}$


## Generative Process for network growth

Predictive distribution of $\boldsymbol{Z}_{n+1}$ given the latent process $\boldsymbol{X}_{1}, \ldots, \boldsymbol{X}_{n}$
Books


## Generative Process for network growth

Predictive distribution of $\boldsymbol{Z}_{n+1}$ given the latent process $\boldsymbol{X}_{1}, \ldots, \boldsymbol{X}_{n}$
Books


Reader 2


## Generative Process for network growth

Predictive distribution of $\boldsymbol{Z}_{n+1}$ given the latent process $\boldsymbol{X}_{1}, \ldots, \boldsymbol{X}_{n}$
Books


## Generative Process for network growth

Predictive distribution of $\boldsymbol{Z}_{n+1}$ given the latent process $\boldsymbol{X}_{1}, \ldots, \boldsymbol{X}_{n}$
Books


## Generative Process for network growth

Predictive distribution of $\boldsymbol{Z}_{n+1}$ given the latent process $\boldsymbol{X}_{1}, \ldots, \boldsymbol{X}_{n}$
Books

| Reader 1 | 18 | 4 | 14 |  |  |  |  |  | $\cdots$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Reader 2 | 12 | 0 | 8 | 13 | 4 |  |  |  | $\cdots$ |



Generative Process for network growth
Predictive distribution of $\boldsymbol{Z}_{\boldsymbol{n + 1}}$ given the latent process $\boldsymbol{X}_{1}, \ldots, \boldsymbol{X}_{\boldsymbol{n}}$
Books

| Reader 1 |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Reader 2 |$\quad$| 18 | 4 | 14 |  |  |
| :---: | :---: | :---: | :---: | :---: |
|  |  |  |  | $\cdots$ |
| 12 | 0 | 8 | 13 | 4 |
|  |  |  | $\cdots$ |  |

Reader 3


Generative Process for network growth
Predictive distribution of $\boldsymbol{Z}_{\boldsymbol{n + 1}}$ given the latent process $\boldsymbol{X}_{1}, \ldots, \boldsymbol{X}_{\boldsymbol{n}}$
Books

| Reader 1 | 18 | 4 | 14 |  |  |  |  |  | $\cdots$ |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Reader 2 | 12 | 0 | 8 | 13 | 4 |  |  |  | $\cdots$ |
|  | Reader 3 |  |  |  |  |  |  |  |  |



Generative Process for network growth
Predictive distribution of $\boldsymbol{Z}_{\boldsymbol{n + 1}}$ given the latent process $\boldsymbol{X}_{1}, \ldots, \boldsymbol{X}_{\boldsymbol{n}}$
Books



Generative Process for network growth
Predictive distribution of $\boldsymbol{Z}_{\boldsymbol{n + 1}}$ given the latent process $\boldsymbol{X}_{1}, \ldots, \boldsymbol{X}_{\boldsymbol{n}}$
Books

| Reader 1 |  |  |  |  |  |  |  |  |  |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Reader 2 |  |  |  |  |  |  |  |  |  |
| Rea |  |  |  |  |  |  |  |  |  |
| Reader 3 | 12 | 4 | 14 |  |  |  |  |  | $\ldots$ |
|  | 16 | 0 | 8 | 13 | 4 |  |  |  | $\ldots$ |



## Prior Draws

Generalized Gamma process with $\lambda(w)=\frac{\alpha}{\Gamma(1-\sigma)} \boldsymbol{w}^{-\sigma-1} e^{-\tau w}, \tau=1, \gamma_{i}=2$.

(c) $\alpha=1, \sigma=0$

(f) $\alpha=2, \sigma=0.1$

(h) $\alpha=2, \sigma=0.9$

## Properties of the model

- Power-law behavior for the generalized gamma process with $\sigma>0$
- The total number of books read by $\boldsymbol{n}$ readers is $\boldsymbol{O}\left(\boldsymbol{n}^{\sigma}\right)$
- Asympt., the proportion of books read by $m$ readers is $O\left(m^{-1-\sigma}\right)$


- (stable) Beta-Bernoulli/Indian Buffet process as a special case when

$$
\lambda(w)=\frac{\alpha \Gamma(1+c)}{\Gamma(1-\sigma) \Gamma(\sigma+c)} \gamma\left(1-e^{-\gamma w}\right)^{-\sigma-1} e^{-\gamma w(c+\sigma)}
$$

## Bayesian Inference via Gibbs Sampling

## Readers

## Books



- Popularity parameters $\boldsymbol{w}_{j}$ of observed books.
- Sum $\boldsymbol{w}_{*}$ of popularity parameters of unobserved books.
- Latent scores $\boldsymbol{x}_{i j}$ associated to observed edges.
- Posterior distribution $P\left(\left\{w_{j}\right\}, w_{*},\left\{x_{i j}\right\} \mid Z_{1}, \ldots, Z_{n}\right)$

Gibbs sampler for the GGP

$$
\begin{aligned}
\boldsymbol{x}_{i j} \mid \text { rest } & \sim \text { Truncated Gumbel } \\
\boldsymbol{w}_{\boldsymbol{j}} \mid \text { rest } & \sim \text { Gamma } \\
\boldsymbol{w}_{*} \mid \text { rest } & \sim \text { Exponentially tilted stable }
\end{aligned}
$$

## Bayesian Inference via Gibbs Sampling

## Readers

Books


- Popularity parameters $w_{j}$ of observed books.
- Sum $\boldsymbol{w}_{*}$ of popularity parameters of unobserved books.
- Latent scores $\boldsymbol{x}_{i j}$ associated to observed edges.
- Posterior distribution $P\left(\left\{w_{j}\right\}, w_{*},\left\{x_{i j}\right\} \mid Z_{1}, \ldots, Z_{n}\right)$

Gibbs sampler for the GGP

$$
\begin{aligned}
\boldsymbol{x}_{i j} \mid \text { rest } & \sim \text { Truncated Gumbel } \\
\boldsymbol{w}_{\boldsymbol{j}} \mid \text { rest } & \sim \text { Gamma } \\
\boldsymbol{w}_{*} \mid \text { rest } & \sim \text { Exponentially tilted stable }
\end{aligned}
$$

## Bayesian Inference via Gibbs Sampling

## Readers

Books


- Popularity parameters $\boldsymbol{w}_{j}$ of observed books.
- Sum $\boldsymbol{w}_{*}$ of popularity parameters of unobserved books.
- Latent scores $\boldsymbol{x}_{i j}$ associated to observed edges.
- Posterior distribution $P\left(\left\{w_{j}\right\}, w_{*},\left\{x_{i j}\right\} \mid Z_{1}, \ldots, Z_{n}\right)$

Gibbs sampler for the GGP

$$
\begin{aligned}
\boldsymbol{x}_{i j} \mid \text { rest } & \sim \text { Truncated Gumbel } \\
\boldsymbol{w}_{\boldsymbol{j}} \mid \text { rest } & \sim \text { Gamma } \\
\boldsymbol{w}_{*} \mid \text { rest } & \sim \text { Exponentially tilted stable }
\end{aligned}
$$

## Bayesian Inference via Gibbs Sampling

## Readers

Books


- Popularity parameters $\boldsymbol{w}_{\boldsymbol{j}}$ of observed books.
- Sum $\boldsymbol{w}_{*}$ of popularity parameters of unobserved books.
- Latent scores $\boldsymbol{x}_{i j}$ associated to observed edges.
- Posterior distribution $P\left(\left\{w_{j}\right\}, w_{*},\left\{x_{i j}\right\} \mid Z_{1}, \ldots, Z_{n}\right)$

Gibbs sampler for the GGP

$$
\begin{aligned}
\boldsymbol{x}_{i j} \mid \text { rest } & \sim \text { Truncated Gumbel } \\
\boldsymbol{w}_{\boldsymbol{j}} \mid \text { rest } & \sim \text { Gamma } \\
\boldsymbol{w}_{*} \mid \text { rest } & \sim \text { Exponentially tilted stable }
\end{aligned}
$$

## Model for the "interest in reading" parameters

- Still Poisson degree distribution for readers
- Parametric: $\gamma_{i}$ are indep. and identically distributed from a gamma distribution
- Nonparametric: $\gamma_{i}$ are the points of a random atomic measure $\boldsymbol{\Gamma}$
- Gibbs sampler can be derived in the same way as for books


## Application

- Evaluate the fit of three models
- Stable Indian Buffet Process
- Proposed model where $G$ follows a Generalized Gamma process of unknown parameters $\left(\alpha_{w}, \sigma_{w}, \tau_{w}\right)$
- with shared and unknown $\gamma_{i}=\gamma$
- with nonparametric prior where $\Gamma$ follows a generalized gamma process of unknown parameters $\left(\alpha_{\gamma}, \tau_{\gamma}, \sigma_{\gamma}\right)$


## Application: IMDB Movie Actor network

280000 movies, 178000 actors, 341000 edges


Figure: Degree distributions for movies (a-d) and actors (e-h) for the IMDB movie-actor dataset with three different models. Data are represented by red plus fandnsamples from the model by blue crosses.

## Application: Book-crossing community network

5000 readers, 36000 books, 50000 edges


Figure: Degree distributions for readers (a-d) and books (e-h) for the book crossing dataset with three different models. Data are represented by red plus and Fsamples from the model by blue crosses.

## Application: Book-crossing community network

5000 readers, 36000 books, 50000 edges


Figure: Posterior distributions of the power-law parameters $\sigma_{\gamma}$ and $\sigma_{w}$

## Application

- Log-likelihood on test dataset

| Dataset | S-IBP | SG | GGP |
| :--- | :--- | :--- | :--- |
| Board | $\mathbf{9 . 8 2 ( 2 9 . 8 )}$ | $\mathbf{8 . 3 ( 3 0 . 8 )}$ | $-68.6(31.9)$ |
| Forum | -6.7 e 3 | -6.7 e 3 | $\mathbf{- 5 . 6 e \mathbf { 3 }}$ |
| Books | 83.1 | 214 | $\mathbf{4 . 4 e} \mathbf{4}$ |
| Citations | -3.7 e 4 | $-3.7 e 4$ | $\mathbf{- 3 . 4 e 4}$ |
| Movielens100k | -6.7 e 4 | -6.7 e 4 | $\mathbf{- 5 . 5 e 4}$ |
| IMDB | -1.5 e 5 | -1.5 e 5 | $\mathbf{- 1 . 1} \boldsymbol{e 5}$ |

## Summary

- Bayesian nonparametric model for bipartite networks with a potentially infinite number of nodes
- Captures power-law behavior
- Simple generative model for network growth
- Simple computational procedure for posterior simulation.
- Displays a good fit on a variety of social networks


## Future work

- BNP model for general (non-bipartite) networks

- BNP (dynamic) recommender systems
- Latent factorial models and dictionary learning


## Future work

- BNP model for general (non-bipartite) networks

- BNP (dynamic) recommender systems
- Latent factorial models and dictionary learning


## Future work

- BNP model for general (non-bipartite) networks

- BNP (dynamic) recommender systems
- Latent factorial models and dictionary learning


## Bibliography I

Brix, A. (1999).
Generalized gamma measures and shot-noise Cox processes.
Advances in Applied Probability, 31(4):929-953.
Devroye, L. (2009).
Random variate generation for exponentially and polynomially tilted stable distributions.
ACM Transactions on Modeling and Computer Simulation (TOMACS), 19(4):18.
Griffiths, T. and Ghahramani, Z. (2005).
Infinite latent feature models and the Indian buffet process.
In NIPS.
James, L., Lijoi, A., and Prünster, I. (2009).
Posterior analysis for normalized random measures with independent increments.
Scandinavian Journal of Statistics, 36(1):76-97.
James, L. F. (2002).
Poisson process partition calculus with applications to exchangeable models and bayesian nonparametrics.
arXiv preprint math/0205093.
Kingman, J. (1967).
Completely random measures.
Pacific Journal of Mathematics, 21(1):59-78.

## Bibliography II



Lijoi, A., Mena, R. H., and Prünster, I. (2007).
Controlling the reinforcement in bayesian non-parametric mixture models.
Journal of the Royal Statistical Society: Series B (Statistical Methodology),
69(4):715-740.
Lijoi, A. and Prünster, I. (2010).
Models beyond the Dirichlet process.
In N. L. Hjort, C. Holmes, P. M. S. G. W., editor, Bayesian Nonparametrics. Cambridge University Press.
盽 Prünster, I. (2002).
Random probability measures derived from increasing additive processes and their application to Bayesian statistics.
PhD thesis, University of Pavia.
Regazzini, E., Lijoi, A., and Prünster, I. (2003).
Distributional results for means of normalized random measures with independent increments.
The Annals of Statistics, 31(2):560-585.
宣
Teh, Y. and Görür, D. (2009).
Indian buffet processes with power-law behavior.
In Neural Information Processing Systems (NIPS'2009).

