## Probabilistic Graphical Models

# Infinite Feature Models: The Indian Buffet Process 



## Eric Xing

Lecture 21, April 2, 2014

Acknowledgement: slides first drafted by Sinead Williamson

## Limitations of a simple mixture model

- The Dirichlet distribution and the Dirichlet process are great if we want to cluster data into non-overlapping clusters.
- However, DP/Dirichlet mixture models cannot share features between clusters.
- In many applications, data points exhibit properties of multiple latent features
- Images contain multiple objects.
- Actors in social networks belong to multiple social groups.
- Movies contain aspects of multiple genres.


## Latent variable models

- Latent variable models allow each data point to exhibit multiple features, to varying degrees.
- Example: Factor analysis

$$
\mathbf{X}=\mathbf{W A}^{\top}+\varepsilon
$$

- Rows of $\mathbf{A}=$ latent features
- Rows of $\mathbf{W}=$ datapoint-specific weights for these features
- $\varepsilon=$ Gaussian noise.
- Example: Text Documents
- Each document represented by a mixture of features.


## Infinite latent feature models

- Problem: How to choose the number of features?
- Example: Factor analysis

$$
\mathbf{X}=\mathbf{W A}^{\top}+\varepsilon
$$

- Each column of $\mathbf{W}$ (and row of $\mathbf{A}$ ) corresponds to a feature.
- Question: Can we make the number of features unbounded a posteriori, as we did with the DP?
- Solution: allow infinitely many features a priori - ie let W (or A) have infinitely many columns (rows).
- Problem: We can't represent infinitely many features!
- Solution: make our infinitely large matrix sparse, and keep only the selected features

Griffiths and Ghaharamani, 2006

## The CRP: A distribution over indicator matrices

- Recall that the CRP gives us a distribution over partitions of our data.
- Which means that the CRP allows every data point to use one feature (table)

- We can use a similar scheme to represent a distribution over binary matrices recording "feature usage" across data, where each row corresponds to a data point, and each column to a feature
- And we want to encourage every data point to use a small subset of features sparsity


## The Indian Buffet Process (IBP)

- Another culinary experience: we describe a new unbounded multifeature model in terms of the following restaurant analogy.
- The first customer enters a restaurant with an infinitely large buffet
- He helps himself to Poisson $(\alpha)$ dishes.



## The Indian Buffet Process (IBP)

- Another culinary experience: we describe a new unbounded multifeature model in terms of the following restaurant analogy.
- The first customer enters a restaurant with an infinitely large buffet
- He helps himself to Poisson $(\alpha)$ dishes.
- The $n^{\text {th }}$ customer enters the restaurant
- He helps himself to each dish with probability $m_{k} / n$, where $m_{k}$ is the number of times dish $k$ was chosen
- He then tries Poisson $(\alpha / n)$ new dishes



## The Indian Buffet Process (IBP)

- Another culinary experience: we describe a new unbounded multifeature model in terms of the following restaurant analogy.
- The first customer enters a restaurant with an infinitely large buffet
- He helps himself to Poisson $(\alpha)$ dishes.
- The $n^{\text {th }}$ customer enters the restaurant
- He helps himself to each dish with probability $m_{k} / n$, where $m_{k}$ is the number of times dish $k$ was chosen
- He then tries Poisson $(\alpha / n)$ new dishes



## Example



## Data likelihood

- E.g.:

$$
\mathbf{X}=\mathbf{W A}^{\top}+\varepsilon
$$

- Rows of $\mathbf{A}=$ latent features (Gaussian)
- Rows of $\mathbf{W}=$ datapoint-specific weights for these features (Gaussian)
- $\varepsilon=$ Gaussian noise.

$$
\mathbf{W}=\mathbf{Z} \odot \mathbf{V}
$$

- Write
- $Z \sim \operatorname{IBP}(\alpha)$
- $\mathbf{V} \sim \mathcal{N}\left(0, \sigma_{v}{ }^{2}\right)$
- $\mathbf{A} \sim \mathcal{N}\left(0, \sigma_{\mathrm{A}}{ }^{2}\right)$


## This is equivalent to ...

- The infinite limit of a sparse, finite latent variable model:

$$
\begin{aligned}
\mathbf{X} & =\mathbf{W A}^{T}+\epsilon \\
\mathbf{W} & =\mathbf{Z} \odot \mathbf{V}
\end{aligned}
$$

for some sparse matrix $\mathbf{Z}$.

- Place a beta-Bernoulli prior on Z:

$$
\begin{aligned}
\pi_{k} & \sim \operatorname{Beta}\left(\frac{\alpha}{K}, 1\right), k=1, \ldots, K \\
z_{n k} & \sim \operatorname{Bernoulli}\left(\pi_{k}\right), n=1, \ldots, N .
\end{aligned}
$$

## Properties of the IBP

- "Rich get richer" property - "popular" dishes become more popular.
- The number of nonzero entries for each row is distributed according to Poisson $(\alpha)$ - due to exchangeability.
- Recall that if $x_{1} \sim \operatorname{Poisson}\left(\alpha_{1}\right)$ and $x_{2} \sim \operatorname{Poisson}\left(\alpha_{2}\right)$, then $\left(x_{1}+x_{2}\right) \sim$ Poisson $\left(\alpha_{1}+\alpha_{2}\right)$
- The number of nonzero entries for the whole matrix is distributed according to Poisson(Na).
- The number of non-empty columns is distributed according to Poisson $\left(\alpha H_{N}\right)$, where $H_{N}=\sum_{n=1}^{N} \frac{1}{n}$


## A two-parameter extension

- In the IBP, the parameter $\alpha$ governs both the number of nonempty columns and the number of features per data point.
- We might want to decouple these properties of our model.
- Reminder: We constructed the IBP as the limit of a finite betaBernoulli model where

$$
\begin{aligned}
\pi_{k} & \sim \operatorname{Beta}\left(\frac{\alpha}{K}, 1\right) \\
z_{n k} & \sim \operatorname{Bernoulli}\left(\pi_{k}\right)
\end{aligned}
$$

- We can modify this to incorporate an extra parameter:

$$
\begin{aligned}
\pi_{k} & \sim \operatorname{Beta}\left(\frac{\alpha \beta}{K}, \beta\right) \\
z_{n k} & \sim \operatorname{Bernoulli}\left(\pi_{k}\right)
\end{aligned}
$$

## A two-parameter extension

- Our restaurant scheme is now as follows:
- A customer enters a restaurant with an infinitely large buffet
- He helps himself to Poisson( $\alpha$ ) dishes.
- The $n^{\text {th }}$ customer enters the restaurant
- He helps himself to each dish with probability $m_{k} /(\beta+n-1)$
- He then tries Poisson( $\alpha \beta /(\beta+n-1))$ new dishes
- Note
- The number of features per data point is still marginally Poisson $(\alpha)$.
- The number of non-empty columns is now

$$
\operatorname{Poisson}\left(\alpha \sum_{n=1}^{N} \frac{\beta}{\beta+n-1}\right)
$$

- We recover the IBP when $\beta=1$.


## Two parameter IBP: examples

```
Im IBP
2
Prior sample from IBP with \(\alpha=10 \beta=1\)
```



Prior sample from IBP with $\alpha=10 \beta=5$


Image from Griffiths and Ghahramani, 2011

## Beta processes and the IBP

- Recall the relationship between the Dirichlet process and the Chinese restaurant process:
- The Dirichlet process is a prior on probability measures (distributions)
- We can use this probability measure as cluster weights in a clustering model cluster allocations are i.i.d. given this distribution.
- If we integrate out the weights, we get an exchangeable distribution over partitions of the data - the Chinese restaurant process.
- De Finetti's theorem tells us that, if a distribution $X_{1}, X_{2}, \ldots$ is exchangeable, there must exist a measure conditioned on which $X_{1}, X_{2}, \ldots$ are i.i.d.


## Beta processes and the IBP

- Recall the finite beta-Bernoulli model:

$$
\begin{aligned}
\pi_{k} & \sim \operatorname{Beta}\left(\frac{\alpha}{K}, 1\right) \\
z_{n k} & \sim \operatorname{Bernoulli}\left(\pi_{k}\right)
\end{aligned}
$$

- The $z_{n k}$ are i.i.d. given the $\pi_{k}$, but are exchangeable if we integrate out the $\pi_{k}$.
- The corresponding distribution for the IBP is the infinite limit of the beta random variables, as $K$ tends to infinity.
- This distribution over discrete measures is called the beta process.
- Samples from the beta process have infinitely many atoms with masses between 0 and 1 .

Thibaux and Jordan, 2007

## Posterior distribution of the beta process

- Question: Can we obtain the posterior distribution of the column probabilities in closed form?
- Answer: Yes!
- Recall that each atom of the beta process is the infinitesimal limit of a Beta( $\alpha / K, 1$ ) random variable.
- Our observation $m_{k}$ for that atom are a Binomial $\left(\pi_{k} N\right)$ random variable.
- We know the beta distribution is conjugate to the Binomial, so the posterior is the infinitesimal limit of a $\operatorname{Beta}\left(\alpha / K+m_{k}, N+1-m_{k}\right)$ random variable.


## A stick-breaking construction for the beta process

- We can construct the beta process using the following stickbreaking construction:
- Begin with a stick of unit length.
- For $k=1,2, \ldots$
- Sample a beta( $\alpha, 1$ ) random variable $\mu_{k}$.
- Break off a fraction $\mu_{k}$ of the stick. This is the $k^{\text {th }}$ atom size.
- Throw away what's left of the stick.
- Recurse on the part of the stick that you broke off

$$
\pi_{k}=\prod_{j=1}^{k} \mu_{j} \quad \mu_{j} \sim \operatorname{Beta}(\alpha, 1)
$$

- Note that, unlike the DP stick breaking construction, the atoms will not sum to one.


## Building latent feature models using the IBP

- We can use the IBP to build latent feature models with an unbounded number of features.
- Let each column of the IBP correspond to one of an infinite number of features.
- Each row of the IBP selects a finite subset of these features.
- The rich-get-richer property of the IBP ensures features are shared between data points.
- We must pick a likelihood model that determines what the features look like and how they are combined.


## 

- Problem with linear Gaussian model: Features are "all or nothing"
- Factor analysis: $\mathbf{X}=\mathbf{W A}^{\top}+\varepsilon$
- Rows of $\mathbf{A}=$ latent features (Gaussian)
- Rows of $\mathbf{W}=$ datapoint-specific weights for these features (Gaussian)
- $\varepsilon=$ Gaussian noise.

- Write $\mathbf{W}=\mathbf{Z} \odot \mathbf{V}$
- $\quad Z \sim \operatorname{IBP}(\alpha)$
- $\mathbf{V} \sim \mathcal{N}\left(0, \sigma_{\mathrm{v}}{ }^{2}\right)$
- $\mathbf{A} \sim \mathcal{N}\left(0, \sigma_{A}{ }^{2}\right)$



## A binary model for latent networks

- Motivation: Discovering latent causes for observed binary data
- Example:
- Data points = patients
- Observed features = presence/absence of symptoms
- Goal: Identify biologically plausible "latent causes" - eg illnesses.
- Idea:
- Each latent feature is associated with a set of symptoms
- The more features a patient has that are associated with a given symptom, the more likely that patient is to exhibit the symptom.

Wood et al, 2006

## A binary model for latent networks

- We can represent this in terms of a Noisy-OR model:

$$
\begin{aligned}
\mathbf{Z} & \sim \operatorname{IBP}(\alpha) \\
y_{d k} & \sim \operatorname{Bernoulli}(p) \\
p\left(x_{n d}=1 \mid \mathbf{Z}, \mathbf{Y}\right) & =1-(1-\lambda)^{\mathbf{Z}_{n} \mathbf{y}_{d}^{T}}(1-\epsilon)
\end{aligned}
$$

- Intuition:
- Each patient has a set of latent causes.
- For each sympton, we toss a coin with probability $\lambda$ for each latent cause that is "on" for that patient and associated with that feature, plus an extra coin with probability $\varepsilon$.
- If any of the coins land heads, we exhibit that feature.


## Inference in the IBP

- Recall inference methods for the DP:
- Gibbs sampler based on the exchangeable model.
- Gibbs sampler based on the underlying Dirichlet distribution
- Variational inference
- Particle filter.
- We can construct analogous samplers for the IBP


## Inference in the restaurant scheme

- Recall the exchangeability of the IBP means we can treat any data point as if it's our last.
- Let $K_{+}$be the total number of used features, excluding the current data point.
- Let $\Theta$ be the set of parameters associated with the likelihood - eg the Gaussian matrix A in the linear Gaussian model
- The prior probability of choosing one of these features is $m_{k} / N$
- The posterior probability is proportional to

$$
\begin{aligned}
& p\left(z_{n k}=1 \mid \mathbf{x}_{n}, \mathbf{Z}_{-n k}, \Theta\right) \propto m_{k} f\left(\mathbf{x}_{n} \mid z_{n k}=1, \mathbf{Z}_{-n k}, \Theta\right) \\
& p\left(z_{n k}=0 \mid \mathbf{x}_{n}, \mathbf{Z}_{-n k}, \Theta\right) \propto\left(N-m_{k}\right) f\left(\mathbf{x}_{n} \mid z_{n k}=0, \mathbf{Z}_{-n k}, \Theta\right)
\end{aligned}
$$

- In some cases we can integrate out $\Theta$, otherwise we must sample this.


## Inference in the restaurant scheme

- In addition, we must propose adding new features.
- Metropolis Hastings method:
- Let $K^{*}$ old be the number of features appearing only in the current data point.
- Propose $K_{\text {new }}^{*} \sim \operatorname{Poisson}(\alpha / N)$, and let $\mathbf{Z}^{*}$ be the matrix with $K_{\text {new }}^{*}$ features appearing only in the current data point.
- With probability

$$
\min \left(1, \frac{f\left(\mathbf{x}_{n} \mid \mathbf{Z}^{*}, \Theta\right)}{f\left(\mathbf{x}_{n} \mid \mathbf{Z}, \Theta\right)}\right)
$$

accept the proposed matrix.

## Inference in the stick-breaking construction

- We can also perform inference using the stick-breaking representation
- Sample $\mathbf{Z} \mid \boldsymbol{\pi}, \boldsymbol{O}$
- Sample m|Z
- The posterior for atoms for which $m_{k}>0$ is beta distributed.
- The atoms for which $m_{k}=0$ can be sampled using the stickbreaking proceedure.
- We can use a slice sampler to avoid representing all of the atoms, or using a fixed truncation level.

Teh et al, 2007

## Other distributions over infinite, exchangeable matrices

- Recall the beta-Bernoulli process construction of the IBP.
- We start with a beta process - an infinite sequence of values between 0 and 1 that are distributed as the infinitesimal limit of the beta distribution.
- We combine this with a Bernoulli process, to get a binary matrix.
- If we integrate out the beta process, we get an exchangeable distribution over binary matrices.
- Integration is straightforward due to the beta-Bernoulli conjugacy.
- Question: Can we construct other infinite matrices in this way?


## The infinite gamma-Poisson process

- The gamma process can be thought of as the infinitesimal limit of a sequence of gamma random variables.
- Alternatively,

$$
\begin{aligned}
\text { if } D & \sim \operatorname{DP}(\alpha, H) \\
\text { and } \gamma & \sim \operatorname{Gamma}(\alpha, 1) \\
\text { then } G & =\gamma D \sim \operatorname{GaP}(\alpha H)
\end{aligned}
$$

- The gamma distribution is conjugate to the Poisson distribution.


## The infinite gamma-Poisson process

- We can associate each atom $v_{k}$ of the gamma process with a column of a matrix (just like we did with the atoms of a beta process)
- We can generate entries for the matrix as $z_{n k} \sim \operatorname{Poisson}\left(v_{k}\right)$



IBP


infinite gamma-Poisson
Titsias, 2008

## The infinite gamma-Poisson process

- Predictive distribution for the $n^{\text {th }}$ row:
- For each existing feature, sample a count $z_{n k} \sim \operatorname{NegBinom}\left(m_{k}, n /(n+1)\right)$

| 4 | 2 | 4 | 7 | 0 | 0 | 0 | 0 | 0 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 5 | 0 | 2 | 9 | 4 | 1 | 0 | 0 | 0 |
| 3 | 2 | 1 | 6 | 2 | 1 | 0 | 0 | 0 |
| 7 | 1 | 3 | 6 | 3 | 0 | 0 | 0 | 0 |
|  |  |  |  |  |  |  |  |  |

## The infinite gamma-Poisson process

- Predictive distribution for the $n^{\text {th }}$ row:
- For each existing feature, sample a count $z_{n k} \sim \operatorname{NegBinom}\left(m_{k}, n /(n+1)\right)$

| 4 | 2 | 4 | 7 | 0 | 0 | 0 | 0 | 0 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 5 | 0 | 2 | 9 | 4 | 1 | 0 | 0 | 0 |
| 3 | 2 | 1 | 6 | 2 | 1 | 0 | 0 | 0 |
| 7 | 1 | 3 | 6 | 3 | 0 | 0 | 0 | 0 |
| 5 |  |  |  |  |  |  |  |  |

## The infinite gamma-Poisson process

- Predictive distribution for the $n^{\text {th }}$ row:
- For each existing feature, sample a count $z_{n k} \sim \operatorname{NegBinom}\left(m_{k}, n /(n+1)\right)$

| 4 | 2 | 4 | 7 | 0 | 0 | 0 | 0 | 0 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 5 | 0 | 2 | 9 | 4 | 1 | 0 | 0 | 0 |
| 3 | 2 | 1 | 6 | 2 | 1 | 0 | 0 | 0 |
| 7 | 1 | 3 | 6 | 3 | 0 | 0 | 0 | 0 |
| 5 | 0 |  |  |  |  |  |  |  |

## The infinite gamma-Poisson process

- Predictive distribution for the $n^{\text {th }}$ row:
- For each existing feature, sample a count $z_{n k} \sim \operatorname{NegBinom}\left(m_{k}, n /(n+1)\right)$

| 4 | 2 | 4 | 7 | 0 | 0 | 0 | 0 | 0 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 5 | 0 | 2 | 9 | 4 | 1 | 0 | 0 | 0 |
| 3 | 2 | 1 | 6 | 2 | 1 | 0 | 0 | 0 |
| 7 | 1 | 3 | 6 | 3 | 0 | 0 | 0 | 0 |
| 5 | 0 | 4 | 5 | 2 | 0 |  |  |  |

## The infinite gamma-Poisson process

- Predictive distribution for the $n^{\text {th }}$ row:
- For each existing feature, sample a count $z_{n k} \sim \operatorname{NegBinom}\left(m_{k}, n /(n+1)\right)$
- Sample $K_{n}^{*} \sim$ NegBinom $(\alpha, n /(n+1))$

| 4 | 2 | 4 | 7 | 0 | 0 | 0 | 0 | 0 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 5 | 0 | 2 | 9 | 4 | 1 | 0 | 0 | 0 |
| 3 | 2 | 1 | 6 | 2 | 1 | 0 | 0 | 0 |
| 7 | 1 | 3 | 6 | 3 | 0 | 0 | 0 | 0 |
| 5 | 0 | 4 | 5 | 2 | 0 |  |  |  |

## The infinite gamma-Poisson process

- Predictive distribution for the $n^{\text {th }}$ row:
- For each existing feature, sample a count $z_{n k} \sim \operatorname{NegBinom}\left(m_{k}, n /(n+1)\right)$.
- Sample $K^{*}{ }_{n} \sim \operatorname{NegBinom}(a, n /(n+1))$.
- Partition $K_{n}^{*}$ according to the CRP, and assign the resulting counts to new columns.

| 4 | 2 | 4 | 7 | 0 | 0 | 0 | 0 | 0 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 5 | 0 | 2 | 9 | 4 | I | 0 | 0 | 0 |
| 3 | 2 | I | 6 | 2 | I | 0 | 0 | 0 |
| 7 | I | 3 | 6 | 3 | 0 | 0 | 0 | 0 |
| 5 | 0 | 4 | 5 | 2 | 0 | 3 | 1 | 0 |

## Summary

- Infinite latent feature selection models
- IBP: generating random binary matrix
- Equivalence to beta-Bernoulli process
- Inference via MCMC
- Infinite latent feature weighting models
- The gamma-Poisson process


## Supplementary

- Proof of equivalence of IBP to the infinite limit of the betaBernoulli process


## A sparse, finite latent variable model

- If we integrate out the $\pi_{k}$, the marginal probability of a matrix $Z$ is:

$$
\begin{aligned}
& p(\mathbf{Z})=\prod_{k=1}^{K} \int\left(\prod_{n=1}^{N} p\left(z_{n k} \mid \pi_{k}\right)\right) p\left(\pi_{k}\right) d \pi_{k} \\
& =\prod_{k=1}^{K} \frac{B\left(m_{k}+\alpha / K, N-m_{k}+1\right)}{B(\alpha / K, 1)} \\
& =\prod_{k=1}^{K} \frac{\alpha}{K} \frac{\Gamma\left(m_{k}+\alpha / K\right) \Gamma\left(N-m_{k}+1\right)}{\Gamma(N+1+\alpha / K)} \\
& \text { where } m_{k}=\sum_{n=1}^{N} z_{n k}
\end{aligned}
$$

- This is exchangeable (doesn't depend on the order of the rows or columns


## An equivalence class of matrices

- We can naively take the infinite limit by taking $K$ to infinity
- Because all the columns are equal in expectation, as $K$ grows we are going to have more and more empty columns.
- We do not want to have to represent infinitely many empty columns!
- Define an equivalence class [Z] of matrices where the nonzero columns are all to the left of the empty columns.
- Let lof(.) be a function that maps binary matrices to leftordered binary matrices - matrices ordered by the binary number made by their rows.


## How big is the equivalence set?

- All matrices in the equivalence set [Z] are equiprobable (by exchangeability of the columns), so if we know the size of the equivalence set, we know its probability.
- Call the vector $\left(z_{1 k}, z_{2, k}, \ldots, z_{(n-1) k}\right)$ the history of feature $k$ at data point $n$ (a number represented in binary form).
- Let $K_{h}$ be the number of features possessing history $h$, and let $K_{+}$be the total number of features with non-zero history.
- The total number of lof-equivalent matrices in [Z] is

$$
\binom{K}{K_{0} \cdots K_{2^{N}-1}}=\frac{K!}{\prod_{n=0}^{2^{N}-1} K_{n}!}
$$

## Probability of an equivalence class of finite binary matrices.

- If we know the size of the equivalence class [Z], we can evaluate its probability:

$$
p([\mathbf{Z}])=\sum_{\mathbf{Z} \in[\mathbf{Z}]} p(\mathbf{Z})
$$

$$
\begin{aligned}
= & \frac{K!}{\prod_{n=0}^{2^{N}-1} K_{n}!} \prod_{k=1}^{K} \frac{\alpha}{K} \frac{\Gamma\left(m_{k}+\alpha / K\right) \Gamma\left(N-m_{k}+1\right)}{\Gamma(N+1+\alpha / K)} \\
= & \frac{\alpha^{K_{+}}}{\prod_{n=1}^{2^{N}-1} K_{n}!} \frac{K!}{K_{0}!K^{K_{+}}}\left(\frac{N!}{\prod_{j=1}^{N} j+\alpha / K}\right)^{K} \\
& \cdot \prod_{k=1}^{K_{+}} \frac{\left(N-m_{k}\right)!\prod_{j=1}^{m_{k}-1}(j+\alpha / K)}{N!}
\end{aligned}
$$

## Taking the infinite limit

- We are now ready to take the limit of this finite model as $K$ tends to infinity:

$$
\begin{aligned}
& \frac{\alpha^{K_{+}}}{\prod_{n=1}^{2^{N}-1} K_{n}!} \frac{K!}{K_{0}!K^{K_{+}}}\left(\frac{N!}{\prod_{j=1}^{N} j+\frac{\alpha}{K}}\right)^{K} \prod_{k=1}^{K_{+}} \frac{\left(N-m_{k}\right)!\prod_{j=1}^{m_{k}-1}\left(j+\frac{\alpha}{K}\right)}{N!} \\
& \downarrow K \rightarrow \infty \\
& \frac{\alpha^{K_{+}}}{\prod_{n=1}^{2^{N}-1} K_{n}!} \quad 1 \quad \exp \left\{-\alpha H_{N}\right\} \quad \prod_{k=1}^{K_{+}} \frac{\left(N-m_{k}\right)!\left(m_{k}-1\right)!}{N!}
\end{aligned}
$$

## Proof that the IBP is lof-equivalent to the infinite beta-Bernoulli model

- What is the probability of a matrix $\mathbf{Z}$ ?
- Let $K_{1}{ }^{(n)}$ be the number of new features in the $n^{\text {th }}$ row.

$$
\begin{aligned}
& p(\mathbf{Z})=\prod_{n=1} p\left(\mathbf{z}_{n} \mid \mathbf{z}_{1:(n-1)}\right) \\
= & \prod_{n=1}^{N} \operatorname{Poisson}\left(K_{1}^{(n)} \left\lvert\, \frac{\alpha}{n}\right.\right) \prod_{k=1}^{K_{+}}\left(\frac{\sum_{i=1}^{n-1} z_{i k}}{n}\right)^{z_{n k}}\left(\frac{n-\sum_{i=1}^{n-1} z_{i k}}{n}\right)^{1-z_{n k}} \\
= & \prod_{n=1}^{N}\left(\frac{\alpha}{n}\right)^{K_{1}^{(n)}} \frac{1}{K_{1}^{(n)}!} e^{-\alpha / n} \prod_{k=1}^{K_{+}}\left(\frac{\sum_{i=1}^{n-1} z_{i k}}{n}\right)^{z_{n k}}\left(\frac{n-\sum_{i=1}^{n-1} z_{i k}}{n}\right)^{1-z_{n k}} \\
= & \frac{\alpha^{K_{+}}}{\prod_{n=1}^{N} K_{1}^{(n)}!} \exp \left\{-\alpha H_{N}\right\} \prod_{k=1}^{K_{+}} \frac{\left.N-m_{k}\right)!\left(m_{k}-1\right)!}{N!}
\end{aligned}
$$

- If we include the cardinality of [Z], this is the same as before

