



NCS

Non-Clinical
Statistics
Conference

Padua, IT / 23 – 25 September 2026

Bayesian Nonparametrics Methods for Complex Biomedical Data




UNIVERSITÀ DI PADOVA

Dipartimento
di Scienze Statistiche

Francesco Denti
francesco.denti@unipd.it

The Bayesian framework

- I am a **Bayesian enthusiast!**
- Bayesian statistics provides a **coherent, unified probability framework** for data modeling and inference.
- **Intuitive:** It treats unknown quantities as **random variables** and learns new information from the data via conditional probabilities and **Bayes' rule**, updating our prior belief.
- Its popularity among statisticians exploded in the last few decades. However, wider scientific adoption has been slower.
- **Key hurdle:** Requires a paradigm shift away from familiar concepts like **p-values** and “**statistical significance**”.



ASA
News

AMERICAN STATISTICAL ASSOCIATION
Promoting the Practice and Profession of Statistics®

732 North Washington Street, Alexandria, VA 22314 • (703) 684-1221 • Toll Free: (888) 231-3473 • www.amstat.org • www.twitter.com/AmstatNews



AMERICAN STATISTICAL ASSOCIATION RELEASES STATEMENT ON
STATISTICAL SIGNIFICANCE AND P-VALUES

*Provides Principles to Improve the Conduct and Interpretation of Quantitative
Science*

March 7, 2016


Bayesian nonparametrics

- **Bayesian nonparametric methods** model complex objects: functions, random probability distributions, partitions...
- It can be useful to develop the **best statistical models** to answer compelling scientific questions.
 - “**Best**”: I want to provide a balance among:
 - Data modeling [adopt proper, flexible assumptions];
 - Performance [fitting, predictions];
 - Interpretability [clustering, avoid black-box solutions when possible];
 - Practical relevance [devise scalable, efficient algorithms].




- Deep learning uses **nonlinear, black-box methods** to flexibly estimate functions. **Bayesian nonparametric methods can provide a complementary role** when:
 - **Identifiability and interpretability** are of utmost importance;
 - Our **amount of data is limited** in terms of the number of observations needed to train and tune deep neural networks.

Pandarinath et al., 2022, Nature Comm.




RADICaL, Zhu et al., 2022, Nature Neurosc.



Gaussian Processes

- Flexible priors over a space of functions


Bivariate normal distribution




Gaussian Processes

- Flexible priors over a space of functions

Bivariate normal distribution




Gaussian process



Gaussian Processes

$$f \sim \text{GP}(m(x), K(x, x'))$$

- Flexible priors over a space of functions
- Fully specified by a mean function $m(x)$, often set to zero, and a covariance kernel $K(x, x')$
- The kernel is crucial to specify, since we can obtain very different functional behaviours



Gaussian Processes

$$f \sim \text{GP}(m(x), K(x, x'))$$


- Flexible priors over a space of functions
- Exact posterior inference requires operations on an $n \times n$ covariance matrix and scales as $\mathcal{O}(n^3)$ becoming computationally demanding for large datasets

$$y_i = f(x_i) + \varepsilon_i, \quad \varepsilon_i \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma^2),$$

$$f \mid \mathbf{y} \sim \text{GP}\left(m_{\text{post}}(x), K_{\text{post}}(x, x')\right),$$

$$m_{\text{post}}(x) = m(x) + \mathbf{k}_{x\mathbf{X}} \left(\mathbf{K}_{\mathbf{X}\mathbf{X}} + \sigma^2 \mathbf{I}_n\right)^{-1} (\mathbf{y} - \mathbf{m}_{\mathbf{X}})$$

$$K_{\text{post}}(x, x') = K(x, x') - \mathbf{k}_{x\mathbf{X}} \left(\mathbf{K}_{\mathbf{X}\mathbf{X}} + \sigma^2 \mathbf{I}_n\right)^{-1} \mathbf{k}_{\mathbf{X}x'}.$$




Mixture models

- Intuitively, think about mixture models as **recipes**
- We start from the **dish** and we want to **guess**:
 - The **ingredients**
 - Their **proportions**



Mixture models

- Intuitively, think about mixture models as **recipes**
- We start from the **data** and we want to **estimate**:
 - The **mixture components**
 - Their **mixture weights**



Bayesian Finite Mixture Model

$$\begin{aligned} X_i &\sim F \quad \forall i & F &= \sum_{k=1}^K \omega_k f(x | \theta_k) \\ \theta_k &\sim H \quad \forall k & \omega &\sim \text{Dirichlet}(\alpha) \end{aligned}$$

- flexible representation for complex distributions
- $f_k = f(x | \theta_k)$: **mixture kernel** used as building block
- θ_k is the parameter (location) characterizing the mixture component
- $\omega = (\omega_1, \dots, \omega_K)$ is the vector of **mixture weights**
- H is a distribution from which we sample the **mixture locations**


Bayesian Finite Mixture Model - Membership Labels

$$X_i \sim f_{z_i} \quad \forall i \qquad z_i \sim \text{Cat}_K(\omega)$$

$$\theta_k \sim H \quad \forall k \qquad \omega \sim \text{Dirichlet}(\alpha)$$


- To simplify the posterior computation we introduce, for each observation, the **membership label** z_i
- $z_i = k$: the i -th observation belongs to the k -th mixture component
- A partition is induced in the data: **model-based clustering**

Simple dataset




Mixture density

$$X_i \sim F \quad \forall i$$



Induced clustering

$$X_i \sim f_{z_i} \quad \forall i$$



The Dirichlet Process

Many characterizations for the same object

- Stochastic process frequently adopted as a prior for mixture models
- It is a **distribution over a space of distributions**
- Characterized by a base measure H and concentration parameter α
- Its realizations are probability distributions that are almost surely discrete:


$$G \sim DP(\alpha, H)$$

Sethuraman's Stick Breaking representation

$$G \sim DP(\alpha, H)$$

is equivalent to

$$G = \sum_{k=1}^{+\infty} \omega_k \delta_{\theta_k}(\cdot), \quad \omega \sim SB(\alpha), \quad \theta_k \sim H$$




Stick Breaking Procedure


- Sample a sequence of independent Beta proportions
- Combine them according to the formulas below to obtain the weights

$$\pi_k \sim \text{Beta}(1, \alpha) \quad \forall k$$

$$\omega_k = \pi_k \prod_{l < k} (1 - \pi_l)$$



Examples of DP realizations



Dirichlet Process Mixture Model

- The discreteness of its realizations make the DP difficult to use to model the data directly
- Instead, the DP is used as a mixing measure, convoluting a given parametric kernel

$$F(x) = \sum_{k=1}^{+\infty} \omega_k f(x | \theta_k),$$
$$\omega \sim SB(\alpha), \quad \theta_k \sim H.$$

Dirichlet Process Mixture Model



- The discreteness of its realizations make the DP difficult to use to model the data directly
- For every observation x_i we can introduce the latent membership indicators z_i to simplify posterior simulation:

$$F(x_i) = f\left(x_i \mid \theta_{z_i}\right), \quad z_i \sim \sum_{k=1}^{+\infty} \omega_k \delta_k(\cdot)$$

$$\omega \sim SB(\alpha), \quad \theta_k \sim H.$$

Dirichlet Process Mixture Model


- So we have a good model for
 - Density estimation
 - Clustering
- Many different extensions have been proposed:
 - we will focus on nested versions, where data are naturally grouped and
 - we will cluster both observations and groups




- **Pose and Poseidon:** Bayesian model for biclustering **large imaging data**
- MALDI-Mass Spectrometry Imaging mouse brain data
- **Biclustering and grouped data:** separate exchangeability
- **Image Segmentation:** accounting for spatial information with **HMRF**
- **Large dataset:** fast estimation via mean-field variational inference



Cecilia Balocchi
@ University of Edinburgh



Giulia Capitoli
@ University of Milan-Bicocca



Laura D'Angelo
@ University of Milan-Bicocca



Michele Guindani @ UCLA




Antonio Canale @ UniPD


- **SCDC**: a model for **simultaneous clustering** and **deconvolution** of calcium traces
- Calcium imaging data
 - **Spatially clustering** neurons via **PSB**
 - Deconvolution and **detection of spike trains** via latent **GPs** for temporal dependence

MALDI imaging

- **Matrix-Assisted Laser Desorption/Ionisation (MALDI)** - MSI tool extracts the distribution of **molecules** (such as **lipids**) in different locations of a biological sample
- It allows the detection of critical biological traits that would be overlooked with a simple visual morphological assessment
- A slide of a biological sample is divided into a grid of pixels of **50 microns each**
- For each pixel, the instrument acquires a **mass spectrum**, a **collection of molecular abundances** measured at different values of the mass-to-charge index **m/z** (also known as **signal**)




The motivating datasets



- It is obtained from a slice of the brain of a **healthy mouse**, to be used as a benchmark
- The atlas of the mouse brain is well known - this way, biologists can validate their preprocessing pipelines
- I will use this dataset as a tool to build the model


MALDI-MSI procedure

- We want to develop new statistical methods that take into account the information of these complex data




MALDI-MSI procedure

- We want to develop new statistical methods that take into account the information of these complex data
- The tissue is covered by a specific substance called *matrix*, which separates a specific molecule (here, **lipids**) from the rest and makes it “ready to be detected”



MALDI-MSI procedure




- We want to develop new statistical methods that take into account the information of these complex data



- The tissue is covered by a specific substance called *matrix*, which separates a specific molecule (here, **lipids**) from the rest and makes it “ready to be detected”
- A laser hits the tissue in a specific spot, and a mass spectrum is extracted (measuring signal abundance)
- Thus, the data can be organized into a large matrix. We have information for more than **4000 pixels** (i.e., different locations)
- Each pixel contains the abundance for approximately **80 signals** (i.e., different lipids)
- Thus, data can be organized in a **large matrix**



Visual summary

- The MALDI-MSI technique extracts the distributions of the abundance for different **lipid signals (rows)** in many **locations (pixels, columns)** given a biological sample




Visual summary

- For each lipid signal, we recover its expressions all over the brain slice
- There is spatial information to take into account!






Transformed Abundances -4 -3 -2 -1 0






Visual summary

- Within each pixel, we observe a **distribution of abundances**
- Recall that the dataset contains around 1000 pixels that we want to cluster




Visual summary

- Within each pixel, we observe an **abundance of particular signals**
- Recall that the dataset contains around 1000 pixels that we want to cluster



Biclustering algorithms


- We look for “two-dimensional” clustering
- **Biclustering** of the data: **simultaneously** group the pixels (columns) and, *within each column cluster*, group the different values of the signals (rows)
- The problem reminds us (or only me?) of the framework of **clustering grouped data**, for which lively literature has recently flourished (see, for example, Beraha et al. (2021), Balocchi et al. (2023), Camerlenghi et al. (2019), D’Angelo et al. (2023), Denti et al. (2023), Rebaudo et al. (2023), Rodriguez et al. (2011),...)
- All these models assume a **partially exchangeable** framework



Partial exchangeability

- **Partially exchangeable data:** the observations are exchangeable within groups and conditionally independent across groups
- For partially exchangeable data, a **matrix structure is not needed**, as the framework **disregards the identities of the rows**


Partial Exchangeability



Partial exchangeability

- **Partially exchangeable data:** the observations are exchangeable within groups and conditionally independent across groups
- For partially exchangeable data, a **matrix structure is not needed**, as the framework **disregards the identities of the rows**
 - It's not ideal for this application, as we want to take the meaning of the rows (lipids signals) into account!

Partial Exchangeability



Separate exchangeability

- There are cases in which one may want to draw inference about the rows' characteristics: we use **separate exchangeability** (Rebaudo et al., 2021)
- Separate exchangeability is an ideal framework for modeling **matrices where rows and columns have specific meanings**: once the **columns** are clustered together, the model does not disregard any potential information conveyed by the **rows**


Statistics > Methodology

[Submitted on 14 Dec 2021 (v1), last revised 20 Jun 2024 (this version, v2)]


Separate Exchangeability as Modeling Principle in Bayesian Nonparametrics

Giovanni Rebaudo, Qiaohui Lin, Peter Mueller

Partial Exchangeability




Separate Exchangeability



Separate Exchangeability through Nested Partitions

- Let $\mathbf{Y}$ indicate the MALDI-MSI abundance matrix, characterized by N rows (signals) and J columns (pixels)
- We need column-specific membership labels, denoted with $\mathbf{C} = (C_1, \dots, C_J)$, where $C_j \in \{1, \dots, K\}$ for $j = 1, \dots, J$. We will refer to the values in $\mathbf{C}$ as **distributional cluster** (DC) labels.

$$p(\mathbf{C} \mid \boldsymbol{\pi}) = \prod_{j=1}^J \pi_{C_j}$$




Separate Exchangeability through Nested Partitions

- Let $\mathbf{Y}$ indicate the MALDI-MSI abundance matrix, characterized by N rows (signals) and J columns (pixels)
- We need column-specific membership labels, denoted with $\mathbf{C} = (C_1, \dots, C_J)$, where $C_j \in \{1, \dots, K\}$ for $j = 1, \dots, J$. We will refer to the values in $\mathbf{C}$ as **distributional cluster** (DC) labels

$$p(\mathbf{C} \mid \boldsymbol{\pi}) = \prod_{j=1}^J \pi_{C_j}$$

- We consider a set of membership labels for the rows *within every possible column cluster*, defining $\mathbf{R}_k = (R_{1,k}, \dots, R_{N,k})$, with $R_{i,k} \in \{1, \dots, L\}$ for all $i = 1, \dots, N$ and $k = 1, \dots, K$.

$$p(\mathbf{R} \mid \boldsymbol{\omega}) = \prod_{i=1}^N \prod_{k=1}^K \omega_{R_{i,k},k}$$



Separate Exchangeability through Nested Partitions

- Let $\mathbf{Y}$ indicate the MALDI-MSI abundance matrix, characterized by N rows (signals) and J columns (pixels)
- We need column-specific membership labels, denoted with $\mathbf{C} = (C_1, \dots, C_J)$, where $C_j \in \{1, \dots, K\}$ for $j = 1, \dots, J$. We will refer to the values in $\mathbf{C}$ as **distributional cluster** (DC) labels.

$$p(\mathbf{C} \mid \boldsymbol{\pi}) = \prod_{j=1}^J \pi_{C_j}$$

- We consider a set of membership labels for the rows *within every possible column cluster*, defining $\mathbf{R}_k = (R_{1,k}, \dots, R_{N,k})$, with $R_{i,k} \in \{1, \dots, L\}$ for all $i = 1, \dots, N$ and $k = 1, \dots, K$.

$$p(\mathbf{R} \mid \boldsymbol{\omega}) = \prod_{i=1}^N \prod_{k=1}^K \omega_{R_{i,k},k}$$

- Finally, the data distribution is

$$p(\mathbf{Y} \mid \mathbf{R}, \mathbf{C}, \boldsymbol{\theta}) = \prod_{i=1}^N \prod_{j=1}^J f(y_{ij} \mid \theta_{R_{i,C_j}})$$

Statistics > Methodology

[Submitted on 14 Dec 2021 ([v1](#)), last revised 20 Jun 2024 (this version, v2)]

Separate Exchangeability as Modeling Principle in Bayesian Nonparametrics

[Giovanni Rebaudo](#), [Qiaohui Lin](#), [Peter Mueller](#)

Separate Exchangeability through Nested Partitions

- Let $\mathbf{Y}$ indicate the MALDI-MSI abundance matrix, characterized by N rows (signals) and J columns (pixels)
- We need column-specific membership labels, denoted with $\mathbf{C} = (C_1, \dots, C_J)$, where $C_j \in \{1, \dots, K\}$ for $j = 1, \dots, J$. We will refer to the values in $\mathbf{C}$ as **distributional cluster** (DC) labels.

$$p(\mathbf{C} \mid \boldsymbol{\pi}) = \prod_{j=1}^J \pi_{C_j}$$

- We consider a set of membership labels for the rows *within every possible column cluster*, defining $\mathbf{R}_k = (R_{1,k}, \dots, R_{N,k})$, with $R_{i,k} \in \{1, \dots, L\}$ for all $i = 1, \dots, N$ and $k = 1, \dots, K$.

$$p(\mathbf{R} \mid \boldsymbol{\omega}) = \prod_{i=1}^N \prod_{k=1}^K \omega_{R_{i,k},k}$$

- Finally, the data distribution is

$$p(\mathbf{Y} \mid \mathbf{R}, \mathbf{C}, \boldsymbol{\theta}) = \prod_{i=1}^N \prod_{j=1}^J f(y_{ij} \mid \theta_{R_i, C_j})$$

The column cluster is the second index of the row cluster! The row cluster is shared across all the elements of the column cluster!

Reminiscent of the biclustering model by Lee et al. (2013)

Presence of **Common Atoms**, as in Denti et al. (2023) and Chandra et al. (2023)

Statistics > Methodology

[Submitted on 14 Dec 2021 (v1), last revised 20 Jun 2024 (this version, v2)]

Separate Exchangeability as Modeling Principle in Bayesian Nonparametrics


Giovanni Rebaudo, Qiaohui Lin, Peter Mueller

Including spatial information in the model

- Include in the model the spatial correlation that may occur across pixels
- To do so, we use **hidden Markov random fields** (Besag, 1986) to specify a clustering distribution over the columns
- In particular, we use a **Potts model** with K potential clusters

$$p(C_j = k \mid C_{-j}) \propto \exp \left[\beta \sum_{q \sim \mathcal{N}_j} \mathbb{1}_{\{C_q = k\}} \right], \quad k = 1, \dots, K$$

- β is the inverse temperature parameter (crucial, but hard to estimate!) and $\mathcal{N}_j$ denotes the neighborhood of pixel j
- The larger the inverse temperature, the higher the probability of a column to be in the same cluster of its neighbors




Including spatial information in the model

- Include in the model the spatial correlation that may occur across pixels
- To do so, we use **hidden Markov random fields** (Besag, 1986) to specify a clustering distribution over the columns
- In particular, we use a **Potts model** with K potential clusters

$$p(C_j = k \mid C_{-j}) \propto \exp \left[\beta \sum_{q \sim \mathcal{N}_j} \mathbb{1}_{\{C_q = k\}} \right], \quad k = 1, \dots, K$$

- β is the inverse temperature parameter (crucial, but hard to estimate!) and $\mathcal{N}_j$ denotes the neighborhood of pixel j
- The larger the inverse temperature, the higher the probability of a column to be in the same cluster of its neighbors





Including spatial information in the model

- Include in the model the spatial correlation that may occur across pixels
- To do so, we use **hidden Markov random fields** (Besag, 1986) to specify a clustering distribution over the columns
- In particular, we use a **Potts model** with K potential clusters

$$p(C_j = k \mid C_{-j}) \propto \exp \left[\beta \sum_{q \sim \mathcal{N}_j} \mathbb{1}_{\{C_q = k\}} \right], \quad k = 1, \dots, K$$

- β is the inverse temperature parameter (crucial, but hard to estimate!) and $\mathcal{N}_j$ denotes the neighborhood of pixel j
- The larger the inverse temperature, the higher the probability of a column to be in the same cluster of its neighbors



Including spatial information in the model

- Include in the model the spatial correlation that may occur across pixels
- To do so, we use **hidden Markov random fields** (Besag, 1986) to specify a clustering distribution over the columns
- In particular, we use a **Potts model** with K potential clusters


$$p(C_j = k \mid C_{-j}) \propto \exp \left[\beta \sum_{q \sim \mathcal{N}_j} \mathbb{1}_{\{C_q = k\}} \right], \quad k = 1, \dots, K$$

- β is the inverse temperature parameter (crucial, but hard to estimate!) and $\mathcal{N}_j$ denotes the neighborhood of pixel j
- The larger the inverse temperature, the higher the probability of a column to be in the same cluster of its neighbors

Prob of green $\propto \exp [5\beta]$

Prob of blue $\propto \exp [\beta]$

Prob of yellow $\propto \exp [2\beta]$



Including spatial information in the model

- We can also extend the basic Potts model with the **BNP-HMRF prior** by Lü et al. (2021)
- Clever idea: adding as **first-order potentials** in the energy functions the element of a **log-stick breaking sequence**, combining two types of weights



$$p(\mathbf{C} \mid \boldsymbol{\pi}) = \frac{1}{\mathcal{K}(\boldsymbol{\pi}, \beta)} \exp \left[\sum_{j=1}^J \boxed{\log(\pi_{C_j})} + \beta \sum_{q \sim \mathcal{N}_j} \boxed{\mathbb{1}_{\{C_q = C_j\}}} \right]$$

First order potentials


Second order potentials

- **Key point:** the distribution above is valid without specifying the number of clusters. Thus, $\boldsymbol{\pi}$ can be specified non-parametrically as a Dirichlet Process or Pitman-Yor process
- Adding $\boldsymbol{\pi}$ allows to estimate the number of clusters nonparametrically
- We are also exploring the performance using sparse finite mixtures


Results

- We developed an efficient CAVI algorithm to perform fast mean-field variational inference
- The model runs in a matter of minutes - R packages are under development
- Really promising results:


K-means++









Biclustering model



Biclustering model
w/ HMRF



Comparing the content of two column clusters



Rinse and repeat! There is more than just lipids out there

- Denti et al. 2022, introduced a laboratory protocol that enables the **cleaning of the sample** from the matrix used after an initial MALDI analysis.
- This step is essential to removing residues from the initial matrix, which was applied to break specific bonds. Thus, the next analysis can take place.

Rinse and repeat! There is more than just lipids out there


- Denti et al. 2022, introduced a laboratory protocol that enables the **cleaning of the sample** from the matrix used after an initial MALDI analysis.
- This step is essential to removing residues from the initial matrix, which was applied to break specific bonds. Thus, the next analysis can take place.

**DISCLAIMER: I did not invent anything here,
the first author is simply a very smart homonymous**




Rinse and repeat! There is more than just lipids out there

- Denti et al. 2022, introduced a laboratory protocol that enables the **cleaning of the sample** from the matrix used after an initial MALDI analysis.
- This step is essential to removing residues from the initial matrix, which was applied to break specific bonds. Thus, the next analysis can take place.




Rinse and repeat! There is more than just lipids out there

- Once "washed," a new matrix can be applied to the sample, designed to break a different type of bond and extract other classes of molecules, such as peptides.
- The extraction order follows a biological rationale in which molecular bonds are broken: lipids, glycans, and finally, peptides.




Next step: Poseidon

- Our data are, in reality, even more complex: the MALDI-MSI extracts datasets for **three different molecules from the same tissue**: lipids, N-glycans, peptides
- They all **share the same set of pixels**, and potentially have a different number of signals
- We are developing **POSEIDON: P**Ott's model over **S**eparate **E**xchangeability **I**ntegrating **D**ifferent **O**mic's **N**ested data to perform image segmentation led by multiple sources




Datasets


Lipids




N-Glycans



Peptides








Lipids




N-Glycans

Peptides




Datasets



Shared
Column
Partition



Dataset-specific
Row Partitions



Next step: Poseidon

Dataset-specific Row Partitions

Shared Column Partition


$$p(\mathbf{Y} \mid \mathbf{R}, \mathbf{C}, \boldsymbol{\theta}) = \prod_{t=1}^T \prod_{i=1}^{N^{(t)}} \prod_{j=1}^J \prod_{l=1}^L \prod_{k=1}^K \left[\phi \left(y_{i,j}^{(t)} \mid \theta_l^{(t)} \right)^{\mathbb{1}_{\{r_{i,k}^{(t)}=l\}}} \cdot \mathbb{1}_{\{C_j=k\}} \right]$$

$$p(\mathbf{R} \mid \mathbf{C}, \boldsymbol{\theta}) = \prod_{t=1}^T \prod_{i=1}^{N^{(t)}} \prod_{l=1}^L \prod_{k=1}^K \left[\omega_{l,k}^{(t)} \right]^{\mathbb{1}_{\{r_{i,k}^{(t)}=l\}}}$$

$$p(\mathbf{C} \mid \boldsymbol{\pi}) = \frac{1}{\mathcal{K}_{\mathbf{C}}(\boldsymbol{\pi}, \beta)} \exp \left[\sum_{j=1}^J \log(\pi_{C_j}) + \beta \sum_{q \in \mathcal{N}_j} \mathbb{1}_{\{C_j=C_q\}} \right]$$



Some preliminary results

Shared column clustering



Some preliminary results

Shared column clustering





Some preliminary results

Shared column clustering


From glycans

From lipids


From peptides



Some preliminary results




Cool data




+

Smart collaborators




+

Nice results



Cool data




+

Smart collaborators



+

Nice results





Published paper



≠





Cool data



Smart collaborators



Nice results



Go to a conference in Greece, take a bus to the temple of Poseidon at Cape Sounion, and then ask for his favors and blessings



Published paper



Calcium imaging

A Bayesian nonparametric model for

Simultaneous **C**lustering and
Deconvolution of **C**alcium traces

Calcium imaging


A Bayesian nonparametric model for

Simultaneous **C**lustering and
Deconvolution of **C**alcium traces





Mapping Neuronal Activity in Real Time

- **Neuroimaging** uses advanced physiological mechanisms to **visually capture and measure brain function**.
 - Reveals how individuals and **groups of neurons respond** during tasks.
 - It helps decode how neurons work together to produce cognition, behavior, and perception
 - While the anatomical structure is well-known, the **functional properties** are still unknown.
 - Neurons operate at incredible speed - tracking their real-time activity is complex but crucial for breakthroughs in neuroscience.




Calcium imaging measurements

- **Miniaturized Microscopes:** Compact, high-resolution imaging systems
 - Allows real-time neural recording in freely moving subjects
- **Fluorescent Calcium Indicators:** Chemical markers that bind to calcium ions during neuronal activation
 - Enables precise measurement of the activity of an individual neuron




Calcium imaging measurements

- Physiological process behind calcium imaging:
 - **External Stimulus:** Triggers neuronal response
 - **Neuronal Activation:** Calcium floods the cell, causing a temporary increase in intracellular calcium concentration
 - **Return to Baseline:** When the neuron is at rest, calcium levels return to their normal state
- The fluorescent calcium traces can be used as **proxies of the activity** over time of individual neurons



- Example of (a **short** chunk of) a calcium trace for an individual neuron: it represents the level of calcium over time, the **spikes** in the observed concentration correspond to the **neuron's activation**.



- **Two-step procedure:**
 - **Deconvolution:** extraction of the spike train (i.e., the series of recorded firing times). The spike train is then regarded as the *true cleaned signal*.
 - **Analysis of the spike train:** definition of summary statistics, models, etc., to link the neuronal activity with the external conditions that induced it.

Deconvolution

- Many deconvolution methods are based on a **biophysical model** that relates the **observed fluorescence trace** to the **calcium dynamic** and the **spiking activity**.
- For $i = 1, \dots, n$ and $t = 1, \dots, T$, $y_{i,t}$ is the observed fluorescence of neuron i at time t , $c_{i,t}$ is the intracellular calcium concentration of neuron i at time t .
- A popular model* considers the observed fluorescence $y_{i,t}$ as a **noisy function of the intracellular calcium concentration** $c_{i,t}$.

$$y_{i,t} = b_i + c_{i,t} + \epsilon_{i,t}, \quad \epsilon_{i,t} \sim \mathcal{N}(0, \sigma^2)$$

- The calcium dynamic is then modeled using an **autoregressive process with jumps** at the neuron's **firing events**.


$$c_{i,t} = \gamma c_{i,t-1} + a_{i,t} \cdot s_{i,t} + \omega_{i,t}, \quad \omega_{i,t} \sim \mathcal{N}(0, \tau^2)$$

- Here, b_i denotes the baseline level, $\epsilon_{i,t}$ measurement error, γ the decay parameter.

The specification of the latent neuronal activity

$$c_{i,t} = \gamma c_{i,t-1} + a_{i,t} \cdot s_{i,t} + \omega_{i,t}, \quad \omega_{i,t} \sim \mathcal{N}(0, \tau^2)$$


- An autoregressive component models the calcium dynamics
- An additional component is considered for modeling the calcium spikes
 - We distinguish the **binary signal indicators** and the **amplitudes of the non-zero spikes**:
 - $s_{i,t} \in \{0, 1\}$ describes the absence/presence of activation for a neuron i at the time t
 - $a_{i,t} \in \mathbb{R}_+$ describes its magnitude




The specification of the latent neuronal activity

$$c_{i,t} = \gamma c_{i,t-1} + a_{i,t} \cdot s_{i,t} + \omega_{i,t}, \quad \omega_{i,t} \sim \mathcal{N}(0, \tau^2)$$

- An autoregressive component models the calcium dynamics
- An additional component is considered for modeling the calcium spikes
 - We distinguish the **binary signal indicators** and the **amplitudes of the non-zero spikes**:
 - $s_{i,t} \in \{0,1\}$ describes the absence/presence of activation for a neuron i at the time t
 - $a_{i,t} \in \mathbb{R}_+$ describes its magnitude



Spike trains




- The spike trains are the indicators of the neurons' activity over time
- They can be viewed as binary time series: for a neuron i

$$s_i = [s_{i,1}, \dots, s_{i,t}, \dots, s_{i,T}] \in \{0,1\}^T$$

- Neurons with **synchronized activity** should be characterized by **similarities in the spike trains**
- These series drive the identification of the clusters of co-activating neurons.


$s_{i,t}$



Spike trains

- Difficulties: $s_i = [s_{i,1}, \dots, s_{i,t}, \dots, s_{i,T}] \in \{0,1\}^T$
 - Often the firing events are **not strictly synchronous or overlapping in time**
 - Some **observed spikes might be due to measurement error**, or they might be random
 - Firing events are **not uniformly distributed over time** (an observed spike is produced by several consecutive activations): temporal dependence;
 - Not immediate how to cluster binary time series


$s_{i,t}$



Data augmentation: the spike probability

- We introduce an underlying continuous process that describes the evolution of the **spike probability** over time:
 - The spike probability is modeled by a **Probit transformation of a Gaussian process** $\tilde{s}_i(t)$.
- Conditionally on $\tilde{s}_i(t)$, the spikes are generated as independent Bernoulli r.v.

$$s_{i,t} \mid \tilde{s}_i(t) \sim \text{Bernoulli}(\Phi(\tilde{s}_i(t))) \quad \text{for } t = 1, \dots, T$$



Clustering spike trains

- **Idea:** similar evolutions of the spike probability will produce similar spike trains. We can cluster neurons through these latent series.

$$\tilde{s}_i \sim \sum_k \pi_k(\ell_i) \delta_{\tilde{s}_k^*}$$

- **Mixture weights:**
 - We would like to **inform the clustering** using the **anatomical location** ℓ_i
 - We use the **probit stick-breaking process** as prior distribution on the mixture probabilities:
 - More similar set of weights are induced on neighboring neurons

$$\pi_k(\ell_i) = \Phi(\alpha_k(\ell_i)) \prod_{r < k} (1 - \Phi(\alpha_r(\ell_i))) \quad \text{for } k = 1, 2, \dots$$

$$\left[\alpha_k(\ell_1), \alpha_k(\ell_2), \dots, \alpha_k(\ell_n) \right]^T \sim \mathbf{N}_n(\mathbf{0}, \Sigma).$$

Clustering spike trains

- **Idea:** similar evolutions of the spike probability will produce similar spike trains. We can cluster neurons through these latent series.

$$\tilde{s}_i \sim \sum_k \pi_k(\ell_i) \delta_{\tilde{s}_k^*}$$

- **Mixture atoms:**
 - Draws from a Gaussian process (indexed by time).
 - Its covariance function Ω describes the **temporal dependence between spikes**.

$$\tilde{s}_k^* \stackrel{iid}{\sim} \text{GP}(\mu, \Omega)$$

What about the amplitudes?


- We place an **additional DP** with a spike in zero on the distribution of the magnitudes, conditionally on the spike trains

$$a_{i,t} \mid s_{i,t}, P \sim (1 - s_{i,t}) \delta_0 + s_{i,t} P,$$
$$P \sim \text{DP}(\alpha, P_0)$$

- The model is then completed with the following priors (mostly dictated by conjugacy):


$$c_{i,0} \sim \text{N}(0, C_0), \quad b_i \sim \text{N}(b_0, B_0),$$
$$1/\sigma^2 \sim \text{Gamma}(h'_\sigma, h''_\sigma), \quad 1/\tau^2 \sim \text{Gamma}(h'_\tau, h''_\tau),$$
$$\gamma \sim \text{Beta}(h_{1\gamma}, h_{2\gamma})$$

A generative look at our model




A generative look at our model

The neurons are clustered according to a PSB, taking space into account




$$\tilde{s}_i \mid \ell_i, G_{\ell_i} \sim G_{\ell_i} \quad G_{\ell_i} = \sum_{k=1}^{\infty} \pi_k(\ell_i) \cdot \delta_{\tilde{s}_k^*}$$

$$\pi_k(\ell_i) = \Phi(\alpha_k(\ell_i)) \prod_{r < k} (1 - \Phi(\alpha_r(\ell_i)))$$



A generative look at our model





The neurons are clustered according to a PSB, taking space into account

$$\tilde{s}_i \mid \ell_i, G_{\ell_i} \sim G_{\ell_i}$$

$$G_{\ell_i} = \sum_{k=1}^{\infty} \pi_k(\ell_i) \cdot \delta_{\tilde{s}_k^*}$$

The corresponding atoms are Gaussian processes...



A generative look at our model



The corresponding atoms are Gaussian processes...

$$\tilde{s}_k^* \stackrel{iid}{\sim} \text{GP}(\mu, \Omega)$$


A generative look at our model



... that are transformed via a Probit transformation

$$\Phi(\tilde{s}_k^*(t))$$

A generative look at our model





These transformed processes model the probability of spiking over time


$$s_{i,t} \mid \tilde{s}_i(t) \sim \text{Bernoulli}(\Phi(\tilde{s}_i(t)))$$

The probability is the same across all the neurons. From here, we generate the **spike trains**, one for each neuron

A generative look at our model




A generative look at our model





A second DP models the magnitude of each spike

$$a_{i,t} \mid s_{i,t}, P \sim (1 - s_{i,t}) \delta_0 + s_{i,t} P,$$
$$P \sim \text{DP}(\alpha, P_0)$$




$$a^* \sim \bar{a} + \text{Gamma}(a, b)$$

A generative look at our model





A generative look at our model





Finally, given the spike trains and the magnitudes, we generate the calcium traces series with

$$y_{i,t} \mid b_i, \gamma, c_{i,t-1}, s_{i,t}, a_{i,t}, \sigma^2, \tau^2 \sim \text{N}(b_i + \gamma c_{i,t-1} + s_{i,t} \cdot a_{i,t}, \sigma^2 + \tau^2)$$

A generative look at our model




On real data




- We consider a short subset of the complete time series of the experiment
- The mouse is left to explore the circular arena
- Data: 229 neurons, more than 5000 time points that we divide into time windows according to the position of the mouse: center vs. outer ring

Results on a representative time window











- We group the neurons in 6 clusters
- Top: spike trains + amplitudes
- Right: time series and neurons' locations colored by cluster (cluster 1 with 122 inactive neurons not reported)



Cluster 1 2 3 4 5 6

Neuronal response to mouse position



- Postprocessing: analysis among time windows
- We select neurons frequently clustered together and try to understand their **spike patterns associated with the mouse position**

Discussion - Pose and Poseidon









- We introduced a model for biclustering matrix data, taking into account spatial information
- We developed a CAVI algorithm to obtain fast and accurate partitions + **easy to draw inference on nested partitions**
- We have extended the model by **integrating multiple datasets**
- **Open points**
 - CAVI is sensitive to initialization. How to pick the best starting configuration?
 - When to stop? Estimating β and π introduces approximation in the ELBO
 - Extend the framework to different patients, inducing an extra layer of clustering
 - Need for even more scalable computational solutions (stochastic VI?)

Discussion - SCDC

- We introduced a BNP model for flexibly modeling calcium imaging traces, performing simultaneous clustering of neurons and deconvolution.
 - Provide a one-step modeling approach to better quantify uncertainty
- **Open points**
 - C++ implementation of the MCMC algorithm. The time series are long, and the model could be slow. We focused on different subsets - suboptimal.
 - How can the Gaussian Processes be better approximated?
 - VI strategies for PSB?
 - Other modeling approaches: the firing of neurons suggests the use of Hawkes processes

BNP and GPU


- We are exploring the use of BNP mixture models applied to GLMs to capture genetic variation profiles in spatial transcriptomics
- Alex Cecchetto’ PhD thesis has a chapter devoted to scale up BNP applications! It seems to be working well, stay tuned!



Dataset	Dimensions	CPU time	GPU* time	Speed-up
Visium CRC	2,000 genes 2,511 spots	~12 minutes	~1.3 minutes	8.8x
Xenium CRC	384 genes 325,524 cells	~3.0 hours	~3.0 minutes	58.6x
CosMx CRC (Crowell et al. 2025)	2,000 genes 351,968 cells	~12.2 hours	~9.9 minutes	74x
Visium HD CRC	1,500 genes 545,553 bins	~13.5 hours	~12.2 minutes	66x



THANK YOU FOR YOUR ATTENTION!



UNIVERSITÀ DI PADOVA

Dipartimento
di Scienze Statistiche

Francesco Denti
francesco.denti@unipd.it
University of Padua

Common Atoms Model for nested data

- Let $(y_{1,j}, \dots, y_{n_j,j})$ indicate the scores of the songs of the j -th artist

$$y_{1,j}, \dots, y_{n_j,j} \mid p_j \stackrel{ind.}{\sim} p_j, \quad p_j(\cdot) = \int_{\Theta} f(\cdot \mid \theta) G_j(d\theta)$$

Mixture model for the energy distribution of **each artist**

Common Atoms Model for nested data

- Let $(y_{1,j}, \dots, y_{n_j,j})$ indicate the scores of the songs of the j -th artist

$$y_{1,j}, \dots, y_{n_j,j} \mid p_j \stackrel{ind.}{\sim} p_j,$$

$$p_j(\cdot) = \int_{\Theta} f(\cdot \mid \theta) G_j(d\theta)$$

$$G_1, \dots, G_J \mid Q \sim Q$$

$$Q = \sum_{k=1}^{\infty} \pi_k \delta_{G_k^*}$$

The group-specific mixing measures G_j are sampled from a common discrete distribution Q , where $\{\pi_k\}_{k \geq 1} \sim GEM(\alpha)$

Common Atoms Model for nested data

- Let $(y_{1,j}, \dots, y_{n_j,j})$ indicate the scores of the songs of the j -th artist

$$y_{1,j}, \dots, y_{n_j,j} \mid p_j \stackrel{\text{ind.}}{\sim} p_j,$$

$$p_j(\cdot) = \int_{\Theta} f(\cdot \mid \theta) G_j(d\theta)$$

$$G_1, \dots, G_J \mid Q \sim Q$$

$$Q = \sum_{k=1}^{\infty} \pi_k \delta_{G_k^*}$$

The group-specific mixing measures G_j are sampled from a common discrete distribution Q , where $\{\pi_k\}_{k \geq 1} \sim GEM(\alpha)$

We will refer to G_k^* 's as the **distributional atoms**

Q is discrete, so ties will occur among the mixing distributions: **distributional clustering**

Common Atoms Model for nested data

- Let $(y_{1,j}, \dots, y_{n_j,j})$ indicate the scores of the songs of the j -th artist

$$y_{1,j}, \dots, y_{n_j,j} \mid p_j \stackrel{ind.}{\sim} p_j,$$

$$p_j(\cdot) = \int_{\Theta} f(\cdot \mid \theta) G_j(d\theta)$$

$$G_1, \dots, G_J \mid Q \sim Q$$

$$Q = \sum_{k=1}^{\infty} \pi_k \delta_{G_k^*}$$

$$G_k^* = \sum_{l=1}^{\infty} \omega_{l,k} \delta_{\theta_l^*}$$

$$\theta_1^*, \theta_2^*, \dots \sim G_0$$

The distributional atoms are themselves discrete random measures, with

$$\{\omega_{l,k}\}_{l \geq 1} \sim GEM(\beta) \quad \forall k \geq 1$$

Common Atoms Model for nested data

- Let $(y_{1,j}, \dots, y_{n_j,j})$ indicate the scores of the songs of the j -th artist

$$y_{1,j}, \dots, y_{n_j,j} \mid p_j \stackrel{\text{ind.}}{\sim} p_j,$$

$$p_j(\cdot) = \int_{\Theta} f(\cdot \mid \theta) G_j(d\theta)$$

$$G_1, \dots, G_J \mid Q \sim Q$$

$$Q = \sum_{k=1}^{\infty} \pi_k \delta_{G_k^*}$$

$$G_k^* = \sum_{l=1}^{\infty} \omega_{l,k} \delta_{\theta_l^*}$$

$$\theta_1^*, \theta_2^*, \dots \sim G_0$$

The distributional atoms are themselves discrete random measures, with

$$\{\omega_{l,k}\}_{l \geq 1} \sim GEM(\beta) \quad \forall k \geq 1$$

Crucial is the presence of **common observational atoms** sampled from a base measure G_0

G_k^* is a discrete mixing distribution, so multiple observations will be assigned to the same atom θ_l^* : **observational clustering**

Common Atoms Model for nested data

- Let $(y_{1,j}, \dots, y_{n_j,j})$ indicate the gene expressions in the j -th pixel

$$y_{1,j}, \dots, y_{n_j,j} \mid p_j \stackrel{ind.}{\sim} p_j,$$

$$G_1, \dots, G_J \mid Q \sim Q$$

$$G_k^* = \sum_{l=1}^{\infty} \omega_{l,k} \delta_{\theta_l^*}$$

$$\{\pi_k\}_{k \geq 1} \sim GEM(\alpha)$$

$$p_j(\cdot) = \int_{\Theta} f(\cdot \mid \theta) G_j(d\theta)$$

$$Q = \sum_{k=1}^{\infty} \pi_k \delta_{G_k^*}$$

$$\theta_1^*, \theta_2^*, \dots \sim G_0$$

$$\{\omega_{l,k}\}_{l \geq 1} \sim GEM(\beta) \quad \forall k \geq 1$$

The CAM is one of the many proposals devised to address the degeneracy property of the nested DP (see, for example, Beraha et al. (2021), Balocchi et al. (2023), Camerlenghi et al. (2019), D'Angelo et al. (2023), Denti et al. (2023), Rebaudo et al. (2023), Rodriguez et al. (2011),...)

JOURNAL OF THE AMERICAN STATISTICAL ASSOCIATION
2021, VOL. 00, NO. 0, 1–12: Theory and Methods
<https://doi.org/10.1080/01621459.2021.1933499>

 Taylor & Francis
Taylor & Francis Group

 Check for updates

A Common Atoms Model for the Bayesian Nonparametric Analysis of Nested Data

Francesco Denti^{*a} , Federico Camerlenghi^{**b} , Michele Guindani^a , and Antonietta Mira^{c,d} 

^aDepartment of Statistics, University of California, Irvine, CA; ^bDepartment of Economics, Management and Statistics, University of Milano - Bicocca, Milan, Italy; ^cUniversità della Svizzera italiana, Lugano, Switzerland; ^dUniversity of Insubria, Como, Italy



CAM and correlation

- To measure the **dependence across random measures** generated by the process, one can look at $Corr(G_j(A), G_{j'}(A))$ with A generic Borel set
- In CAMs, the correlation does not depend on the specific set A , so it is used as a measure of **generic dependence** among the generated G_j 's, reflecting the **flexibility** of the process
- Thus we write $\rho_{j,j'} = Corr(G_j, G_{j'})$
- Turns out that, for CAM, this prior correlation is pretty high

$$\rho_{j,j'}^{(CAM)} = 1 - \frac{1}{1 + \alpha} \frac{\beta}{1 + 2\beta} \in (0.5, 1)$$



CAM and correlation - Toy example

- Such high correlation a priori also has an **impact a posteriori**, especially with density estimation
- Consider the following toy example, where the distributions of different subpopulations have **no shared atoms** and a few observations per group (e.g., $n_j = 10$)



CAM and correlation - Toy example

- Such high correlation a priori also has an **impact a posteriori**, especially with density estimation
- Consider the following toy example, where the distributions of different subpopulations have **no shared atoms** and a few observations per group (e.g., $n_j = 10$)
- The posterior density estimates contain spurious modes



An Orwellian problem



An Orwellian problem

Bayesian Analysis (2024)

TBA, Number TBA, pp. 1–34

A Finite-Infinite Shared Atoms Nested Model for the Bayesian Analysis of Large Grouped Data Sets

Laura D'Angelo^{*} and Francesco Denti[†]



Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

Econometrics and Statistics

journal homepage: www.sciencedirect.com/journal/econometrics-and-statistics




The generalized nested common atoms model

Francesco Denti^{a,*}, Laura D'Angelo^b

^a University of Padua, via Cesare Battisti 241, Padua, 35121, Italy

^b University of Milan-Bicocca, Piazza dell'Ateneo Nuovo 1, Milan, 20126, Italy



An Orwellian problem

- Weights arising from a SB construction have an implicit **decreasing stochastic ordering**

$$\mathbb{E} [\omega_{1,k}] \geq \mathbb{E} [\omega_{2,k}] \geq \mathbb{E} [\omega_{3,k}] \geq \mathbb{E} [\omega_{4,k}] \geq \dots \quad \forall k$$

- Therefore, the first atoms in the sequence have, in expectation, greater importance:

$$G_1^* = \sum_{l \geq 1} \omega_{l,1} \delta_{\theta_l^*} = \omega_{1,1} \delta_{\theta_1^*} + \omega_{2,1} \delta_{\theta_2^*} + \omega_{3,1} \delta_{\theta_3^*} + \dots$$

$$G_2^* = \sum_{l \geq 1} \omega_{l,2} \delta_{\theta_l^*} = \omega_{1,2} \delta_{\theta_1^*} + \omega_{2,2} \delta_{\theta_2^*} + \omega_{3,2} \delta_{\theta_3^*} + \dots$$

$$G_3^* = \sum_{l \geq 1} \omega_{l,3} \delta_{\theta_l^*} = \omega_{1,3} \delta_{\theta_1^*} + \omega_{2,3} \delta_{\theta_2^*} + \omega_{3,3} \delta_{\theta_3^*} + \dots$$

- And this holds across all the distributional atoms G_k^* 's
- This could imply forced similarities in the G_j 's, even when not motivated by the data

An Orwellian problem

- Weights arising from a SB construction have an implicit **decreasing stochastic ordering**

$$\mathbb{E} [\omega_{1,k}] \geq \mathbb{E} [\omega_{2,k}] \geq \mathbb{E} [\omega_{3,k}] \geq \mathbb{E} [\omega_{4,k}] \geq \dots \quad \forall k$$

- Therefore, the first atoms in the sequence have, in expectation, greater importance:

$$\begin{aligned} G_1^* &= \sum_{l \geq 1} \omega_{l,1} \delta_{\theta_l^*} = \omega_{1,1} \delta_{\theta_1^*} + \omega_{2,1} \delta_{\theta_2^*} + \omega_{3,1} \delta_{\theta_3^*} + \dots \\ G_2^* &= \sum_{l \geq 1} \omega_{l,2} \delta_{\theta_l^*} = \omega_{1,2} \delta_{\theta_1^*} + \omega_{2,2} \delta_{\theta_2^*} + \omega_{3,2} \delta_{\theta_3^*} + \dots \\ G_3^* &= \sum_{l \geq 1} \omega_{l,3} \delta_{\theta_l^*} = \omega_{1,3} \delta_{\theta_1^*} + \omega_{2,3} \delta_{\theta_2^*} + \omega_{3,3} \delta_{\theta_3^*} + \dots \end{aligned}$$

- And this holds across all the distributional atoms G_k^* 's
- This could imply forced similarities in the G_j 's, even when not motivated by the data

Nonparametric vs parametric solutions

Proposition 2.1 Let $G_j, G_{j'} \mid Q \sim Q$ and $Q \sim \text{geCAM}(\mathcal{L}(\boldsymbol{\pi}), \mathcal{L}(\boldsymbol{\omega}_k), H)$, with $j \neq j'$.
Then, the correlation between G_j and $G_{j'}$ is given by

Generic SB laws

$$\begin{aligned} \rho_{j,j'} &:= \text{Corr}(G_j, G_{j'}) \\ &= 1 - (1 - \mathbb{P}[G_j = G_{j'}]) \left(\frac{\mathbb{P}[y_{i,j} = y_{i',j'} \mid G_j \neq G_{j'}]}{\mathbb{P}[y_{i,j} = y_{i',j'} \mid G_j = G_{j'}]} \right) \\ &= 1 - (1 - q_1)(1 - q_2) \end{aligned} \quad (3)$$

where $q_1 = \mathbb{P}[G_j = G_{j'}] = \sum_{k \geq 1} \mathbb{E}[\pi_k^2]$ and

$$q_2 = 1 - \frac{\mathbb{P}[y_{i,j} = y_{i',j'} \mid G_j \neq G_{j'}]}{\mathbb{P}[y_{i,j} = y_{i',j'} \mid G_j = G_{j'}]} = \frac{\sum_{l \geq 1} \mathbb{E}[\omega_{l,k}]^2}{\sum_{l \geq 1} \mathbb{E}[\omega_{l,k}^2]}.$$

$(1 - q_2)$ is the **problematic term**, preventing the correlation from going below 0.5 - it stems only from the **observational weights**

Moreover, the correlation is always non-negative.

- We studied the evolution of the correlations for different specifications:
 - PY, 2PB, atom-skipping process (PAM) by Bi and Ji, 2023 - we call it the **skip-breaking process**, mostly for marketing reasons :)
- But... what if, instead, we **go finite**?

Finite-infinite Shared Atoms Nested model

- Let $(y_{1,j}, \dots, y_{n_j,j})$ indicate the scores of the songs of the j -th artist

$$y_{1,j}, \dots, y_{n_j,j} \mid p_j \stackrel{ind.}{\sim} p_j,$$

$$p_j(\cdot) = \int_{\Theta} f(\cdot \mid \theta) G_j(d\theta)$$

$$G_1, \dots, G_J \mid Q \sim Q$$

$$Q = \sum_{k=1}^{\infty} \pi_k \delta_{G_k^*}$$

$$G_k^* = \sum_{l=1}^{\boxed{L}} \omega_{l,k} \delta_{\theta_l^*}$$

$$\theta_1^*, \theta_2^*, \dots \sim G_0$$

$$\{\pi_k\}_{k \geq 1} \sim GEM(\alpha)$$


$$\boxed{\{\omega_{l,k}\}_{l=1}^L \sim Dirichlet_L(\beta)}$$

- We studied the behavior of this specification and adopted sparse finite mixtures
- Albeit it seems a simplification, it has an important impact on the model

- One can show that

$$\rho_{j,j'}^{(fiSAN)} = 1 - \frac{1}{1 + \alpha} \frac{L - 1}{(b + 1)L} \in (0,1)$$

- We called the **symmetric atoms**, “freed by the stochastic ordering” **shared** atoms, thus **SAN: Shared Atoms Nested** model
- This “simplification” does the trick, keeping the model tractable and interpretable



- One can show that

$$\rho_{j,j'}^{(fiSAN)} = 1 - \frac{1}{1 + \alpha} \frac{L - 1}{(b + 1)L} \in (0,1)$$

- We called the **symmetric atoms**, “freed by the stochastic ordering” **shared** atoms, thus **SAN: Share Atoms Nested** model
- This “simplification” does the trick, keeping the model tractable and interpretable
- Plus: this model specification lends itself to the derivation of a **mean-field variational inference algorithm** that
 - Allows us to scale the application to large datasets
 - Provides point estimates of parameters and partitions, avoiding the label-switching problem
- R packages available on CRAN! Currently working on **SANBA**

