

An elementary derivation of the Chinese restaurant process from Sethuraman’s stick-breaking process

Jeffrey W. Miller

Harvard University, Department of Biostatistics

October 14, 2018

Abstract

The Chinese restaurant process (CRP) and the stick-breaking process are the two most commonly used representations of the Dirichlet process. However, the usual proof of the connection between them is indirect, relying on abstract properties of the Dirichlet process that are difficult for nonexperts to verify. This short note provides a direct proof that the stick-breaking process leads to the CRP, without using any measure theory. We also discuss how the stick-breaking representation arises naturally from the CRP.

1 Introduction

Sethuraman (1994) showed that the Dirichlet process has the following stick-breaking representation: if $\mathbf{v}_1, \mathbf{v}_2, \dots \stackrel{\text{iid}}{\sim} \text{Beta}(1, \alpha)$, $\boldsymbol{\pi}_k = \mathbf{v}_k \prod_{i=1}^{k-1} (1 - \mathbf{v}_i)$ for $k = 1, 2, \dots$, and $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \dots \stackrel{\text{iid}}{\sim} H$, then the random discrete measure

$$\mathbf{P} = \sum_{k=1}^{\infty} \boldsymbol{\pi}_k \delta_{\boldsymbol{\theta}_k} \tag{1}$$

is distributed according to the Dirichlet process $\text{DP}(\alpha, H)$ with concentration parameter α and base distribution H . This representation has been instrumental in the development of many nonparametric models (MacEachern, 1999, 2000; Hjort, 2000; Ishwaran and Zarepour, 2000; Ishwaran and James, 2001; Griffin and Steel, 2006; Dunson and Park, 2008; Chung and Dunson, 2009; Rodriguez and Dunson, 2011; Broderick et al., 2012), has facilitated the understanding of these models (Favaro et al., 2012; Teh et al., 2007; Thibaux and Jordan, 2007; Paisley et al., 2010), and is central to various inference algorithms (Ishwaran and James, 2001; Blei and Jordan, 2006; Papaspiliopoulos and Roberts, 2008; Walker et al., 2007; Kalli et al., 2011).

It is well-known that, as shown by Antoniak (1974), the Dirichlet process induces a distribution on partitions as follows: if $\mathbf{P} \sim \text{DP}(\alpha, H)$ where H is nonatomic (i.e., $H(\{\theta\}) =$

0 for any θ), $\mathbf{x}_1, \dots, \mathbf{x}_n | \mathbf{P} \stackrel{\text{iid}}{\sim} \mathbf{P}$, and \mathbf{C} is the partition of $\{1, \dots, n\}$ induced by $\mathbf{x}_1, \dots, \mathbf{x}_n$, then

$$\mathbb{P}(\mathbf{C} = C) = \frac{\alpha^{|\mathbf{C}|} \Gamma(\alpha)}{\Gamma(\alpha + n)} \prod_{c \in C} \Gamma(|c|). \quad (2)$$

The sequential sampling process corresponding to this partition distribution is known as the Chinese restaurant process (CRP), or Blackwell–MacQueen urn process.

The following key fact is a direct consequence of these two results (Sethuraman’s and Antoniak’s): if $\boldsymbol{\pi} = (\boldsymbol{\pi}_1, \boldsymbol{\pi}_2, \dots)$ is defined as above, $\mathbf{z}_1, \dots, \mathbf{z}_n | \boldsymbol{\pi} \stackrel{\text{iid}}{\sim} \boldsymbol{\pi}$, and \mathbf{C} is the partition induced by $\mathbf{z}_1, \dots, \mathbf{z}_n$, then the distribution of \mathbf{C} is given by Equation 2. This can be seen by noting that when H is nonatomic, the distribution of \mathbf{C} is the same as when it is induced by $\mathbf{x}_1, \dots, \mathbf{x}_n | \mathbf{P}$.

While this key fact follows directly from the results of Sethuraman and Antoniak, the proofs of their results are rather abstract and are not easy to verify, especially for those without expertise in measure theory. The purpose of this note is to provide a proof of this connection between the CRP and the stick-breaking representation using only elementary, non-measure-theoretic arguments. Our proof is completely self-contained and does not rely on any properties of the Dirichlet process or other theoretical results. Conversely, we also provide a sketch of how the CRP naturally leads to the stick-breaking representation.

In previous work, Broderick et al. (2013) used De Finetti’s theorem to provide an elegant derivation of the stick-breaking weights from the CRP. Also, Paisley (2010) showed by elementary calculations that if the base distribution H is a discrete distribution on $\{1, \dots, K\}$, and \mathbf{P} is defined by the stick-breaking process as in Equation 1, then $(\mathbf{P}(1), \dots, \mathbf{P}(K)) \sim \text{Dirichlet}(\alpha H(1), \dots, \alpha H(K))$; thus, despite the similar sounding title of the article by Paisley (2010), the result shown there is altogether different from what we show here.

2 Main result

We use $[n]$ to denote the set $\{1, \dots, n\}$, and \mathbb{N} to denote $\{1, 2, 3, \dots\}$. As is standard, we represent a partition of $[n]$ as a set $C = \{c_1, \dots, c_t\}$ of nonempty disjoint sets c_1, \dots, c_t such that $\bigcup_{i=1}^t c_i = [n]$. Thus, $t = |C|$ is the number of parts in the partition, and $|c|$ is the number of elements in a given part $c \in C$. We say that C is the partition of $[n]$ induced by $\mathbf{z}_1, \dots, \mathbf{z}_n$ if it has the property that for any $i, j \in [n]$, i and j belong to the same part $c \in C$ if and only if $z_i = z_j$. We use bold font to denote random variables.

Theorem 2.1. *Suppose*

$$\begin{aligned} \mathbf{v}_1, \mathbf{v}_2, \dots &\stackrel{\text{iid}}{\sim} \text{Beta}(1, \alpha) \\ \boldsymbol{\pi}_k &= \mathbf{v}_k \prod_{i=1}^{k-1} (1 - \mathbf{v}_i) \text{ for } k = 1, 2, \dots, \\ \mathbf{z}_1, \dots, \mathbf{z}_n | \boldsymbol{\pi} &= \boldsymbol{\pi} \stackrel{\text{iid}}{\sim} \boldsymbol{\pi}, \text{ that is, } \mathbb{P}(\mathbf{z}_i = k | \boldsymbol{\pi}) = \boldsymbol{\pi}_k, \end{aligned}$$

and \mathbf{C} is the partition of $[n]$ induced by $\mathbf{z}_1, \dots, \mathbf{z}_n$. Then

$$\mathbb{P}(\mathbf{C} = C) = \frac{\alpha^{|\mathbf{C}|} \Gamma(\alpha)}{\Gamma(n + \alpha)} \prod_{c \in C} \Gamma(|c|).$$

Our proof of the theorem relies on the following lemmas. Let us abbreviate $\mathbf{z} = (z_1, \dots, z_n)$. Given $\mathbf{z} \in \mathbb{N}^n$, let $C_{\mathbf{z}}$ denote the partition $[n]$ induced by \mathbf{z} .

Lemma 2.2. For any $\mathbf{z} \in \mathbb{N}^n$,

$$\mathbb{P}(\mathbf{z} = z) = \frac{\Gamma(\alpha)}{\Gamma(n + \alpha)} \left(\prod_{c \in C_{\mathbf{z}}} \Gamma(|c| + 1) \right) \left(\prod_{k=1}^m \frac{\alpha}{g_k + \alpha} \right)$$

where $m = \max\{z_1, \dots, z_n\}$ and $g_k = \#\{i : z_i \geq k\}$.

The proofs of the lemmas will be given in Section 3. We use $\mathbf{1}(\cdot)$ to denote the indicator function, that is, $\mathbf{1}(E) = 1$ if E is true, and $\mathbf{1}(E) = 0$ otherwise.

Lemma 2.3. For any partition C of $[n]$,

$$\frac{\alpha^{|\mathbf{C}|}}{\prod_{c \in C} |c|} = \sum_{\mathbf{z} \in \mathbb{N}^n} \mathbf{1}(C_{\mathbf{z}} = C) \prod_{k=1}^{m(\mathbf{z})} \frac{\alpha}{g_k(\mathbf{z}) + \alpha}$$

where $m(\mathbf{z}) = \max\{z_1, \dots, z_n\}$ and $g_k(\mathbf{z}) = \#\{i : z_i \geq k\}$.

Proof of Theorem 2.1.

$$\begin{aligned} \mathbb{P}(\mathbf{C} = C) &= \sum_{\mathbf{z} \in \mathbb{N}^n} \mathbb{P}(\mathbf{C} = C \mid \mathbf{z} = z) \mathbb{P}(\mathbf{z} = z) \\ &\stackrel{(a)}{=} \sum_{\mathbf{z} \in \mathbb{N}^n} \mathbf{1}(C_{\mathbf{z}} = C) \frac{\Gamma(\alpha)}{\Gamma(n + \alpha)} \left(\prod_{c \in C_{\mathbf{z}}} \Gamma(|c| + 1) \right) \left(\prod_{k=1}^{m(\mathbf{z})} \frac{\alpha}{g_k(\mathbf{z}) + \alpha} \right) \\ &= \frac{\Gamma(\alpha)}{\Gamma(n + \alpha)} \left(\prod_{c \in C} \Gamma(|c| + 1) \right) \sum_{\mathbf{z} \in \mathbb{N}^n} \mathbf{1}(C_{\mathbf{z}} = C) \left(\prod_{k=1}^{m(\mathbf{z})} \frac{\alpha}{g_k(\mathbf{z}) + \alpha} \right) \\ &\stackrel{(b)}{=} \frac{\Gamma(\alpha)}{\Gamma(n + \alpha)} \left(\prod_{c \in C} \Gamma(|c| + 1) \right) \frac{\alpha^{|\mathbf{C}|}}{\prod_{c \in C} |c|} \\ &\stackrel{(c)}{=} \frac{\Gamma(\alpha)}{\Gamma(n + \alpha)} \left(\prod_{c \in C} \Gamma(|c|) \right) \alpha^{|\mathbf{C}|} \end{aligned}$$

where (a) is by Lemma 2.2, (b) is by Lemma 2.3, and (c) is since $\Gamma(|c| + 1) = |c| \Gamma(|c|)$. \square

3 Proofs of lemmas

Proof of Lemma 2.2. Letting $e_k = \#\{i : z_i = k\}$, we have

$$\mathbb{P}(\mathbf{z} = z \mid \pi_1, \dots, \pi_m) = \prod_{i=1}^n \pi_{z_i} = \prod_{k=1}^m \pi_k^{e_k}$$

and thus

$$\mathbb{P}(\mathbf{z} = z \mid v_1, \dots, v_m) = \prod_{k=1}^m \left(v_k \prod_{i=1}^{k-1} (1 - v_i) \right)^{e_k} = \prod_{k=1}^m v_k^{e_k} (1 - v_k)^{f_k}$$

where $f_k = \#\{i : z_i > k\}$. Therefore,

$$\begin{aligned} \mathbb{P}(\mathbf{z} = z) &= \int \mathbb{P}(\mathbf{z} = z \mid v_1, \dots, v_m) p(v_1, \dots, v_m) dv_1 \cdots dv_m \\ &= \int \left(\prod_{k=1}^m v_k^{e_k} (1 - v_k)^{f_k} \right) p(v_1) \cdots p(v_m) dv_1 \cdots dv_m \\ &= \prod_{k=1}^m \int v_k^{e_k} (1 - v_k)^{f_k} p(v_k) dv_k \\ &\stackrel{(a)}{=} \prod_{k=1}^m \alpha B(e_k + 1, f_k + \alpha) \\ &= \prod_{k=1}^m \frac{\alpha \Gamma(e_k + 1) \Gamma(f_k + \alpha)}{\Gamma(e_k + f_k + \alpha + 1)} \\ &\stackrel{(b)}{=} \prod_{k=1}^m \frac{\alpha \Gamma(e_k + 1) \Gamma(g_{k+1} + \alpha)}{\Gamma(g_k + \alpha + 1)} \\ &\stackrel{(c)}{=} \left(\prod_{k=1}^m \Gamma(e_k + 1) \right) \left(\prod_{k=1}^m \frac{\alpha}{g_k + \alpha} \right) \left(\prod_{k=1}^m \frac{\Gamma(g_{k+1} + \alpha)}{\Gamma(g_k + \alpha)} \right) \\ &= \left(\prod_{c \in C_z} \Gamma(|c| + 1) \right) \left(\prod_{k=1}^m \frac{\alpha}{g_k + \alpha} \right) \frac{\Gamma(\alpha)}{\Gamma(n + \alpha)} \end{aligned}$$

where step (a) holds since

$$\int x^r (1-x)^s \text{Beta}(x|1, \alpha) dx = \frac{B(r+1, s+\alpha)}{B(1, \alpha)} = \alpha B(r+1, s+\alpha),$$

step (b) since $f_k = g_{k+1}$ and $g_k = e_k + f_k$, and step (c) since $\Gamma(x+1) = x\Gamma(x)$. \square

Let S_t denote the set of $t!$ permutations of $[t]$.

Lemma 3.1. For any $n_1, \dots, n_t \in \mathbb{N}$,

$$\sum_{\sigma \in S_t} \frac{1}{a_1(\sigma) \cdots a_t(\sigma)} = \frac{1}{n_1 \cdots n_t}$$

where $a_i(\sigma) = n_{\sigma_i} + n_{\sigma_{i+1}} + \cdots + n_{\sigma_t}$.

Proof. Consider an urn containing t balls of various sizes—specifically, suppose the balls are labeled $1, \dots, t$ and have sizes n_1, \dots, n_t . Consider the process of sampling without replacement t times from the urn, supposing that the probability of drawing any given ball is proportional to its size. This defines a distribution on permutations $\sigma \in S_t$ such that, letting $n = \sum_{i=1}^t n_i$,

$$\begin{aligned} p(\sigma_1) &= \frac{n_{\sigma_1}}{n} = \frac{n_{\sigma_1}}{a_1(\sigma)}, \\ p(\sigma_2|\sigma_1) &= \frac{n_{\sigma_2}}{n - n_{\sigma_1}} = \frac{n_{\sigma_2}}{a_2(\sigma)}, \\ p(\sigma_3|\sigma_1, \sigma_2) &= \frac{n_{\sigma_3}}{n - n_{\sigma_1} - n_{\sigma_2}} = \frac{n_{\sigma_3}}{a_3(\sigma)}, \end{aligned}$$

and so on. Therefore, since $n_{\sigma_1} \cdots n_{\sigma_t} = n_1 \cdots n_t$,

$$p(\sigma) = p(\sigma_1)p(\sigma_2|\sigma_1) \cdots p(\sigma_t|\sigma_1, \dots, \sigma_{t-1}) = \frac{n_1 \cdots n_t}{a_1(\sigma) \cdots a_t(\sigma)}. \quad (3)$$

Since $p(\sigma)$ is a distribution on S_t by construction, we have $\sum_{\sigma \in S_t} p(\sigma) = 1$; applying this to Equation 3 and dividing both sides by $n_1 \cdots n_t$ gives the result. \square

Proof of Lemma 2.3. Let $t = |C|$, and suppose c_1, \dots, c_t are the parts of C . For $\sigma \in S_t$, define $a_i(\sigma) = |c_{\sigma_i}| + \cdots + |c_{\sigma_t}|$. For any $z \in \mathbb{N}^n$ such that $C_z = C$, if $k_1 < \cdots < k_t$ are the distinct values taken on by z_1, \dots, z_n , then

$$\begin{aligned} \prod_{k=1}^{m(z)} \frac{\alpha}{g_k(z) + \alpha} &= \left(\frac{\alpha}{g_{k_1}(z) + \alpha} \right)^{k_1} \left(\frac{\alpha}{g_{k_2}(z) + \alpha} \right)^{k_2 - k_1} \cdots \left(\frac{\alpha}{g_{k_t}(z) + \alpha} \right)^{k_t - k_{t-1}} \\ &= \left(\frac{\alpha}{a_1(\sigma) + \alpha} \right)^{d_1} \left(\frac{\alpha}{a_2(\sigma) + \alpha} \right)^{d_2} \cdots \left(\frac{\alpha}{a_t(\sigma) + \alpha} \right)^{d_t} \end{aligned}$$

where $d_i = k_i - k_{i-1}$, with $k_0 = 0$, and σ is the permutation of $[t]$ such that $c_{\sigma_i} = \{j : z_j = k_i\}$. Note that the definition of $d = (d_1, \dots, d_t)$ and σ sets up a one-to-one correspondence (that is, a bijection) between $\{z \in \mathbb{N}^n : C_z = C\}$ and $\{(\sigma, d) : \sigma \in S_t, d \in \mathbb{N}^t\}$. Therefore,

$$\begin{aligned} \sum_{z \in \mathbb{N}^n} \mathbf{1}(C_z = C) \prod_{k=1}^{m(z)} \frac{\alpha}{g_k(z) + \alpha} &= \sum_{\sigma \in S_t} \sum_{d \in \mathbb{N}^t} \prod_{i=1}^t \left(\frac{\alpha}{a_i(\sigma) + \alpha} \right)^{d_i} \\ &= \sum_{\sigma \in S_t} \prod_{i=1}^t \sum_{d_i \in \mathbb{N}} \left(\frac{\alpha}{a_i(\sigma) + \alpha} \right)^{d_i} \\ &\stackrel{(a)}{=} \sum_{\sigma \in S_t} \prod_{i=1}^t \frac{\alpha}{a_i(\sigma)} \\ &\stackrel{(b)}{=} \frac{\alpha^t}{\prod_{i=1}^t |c_i|} = \frac{\alpha^t}{\prod_{c \in C} |c|} \end{aligned}$$

where step (a) follows from the geometric series, $\sum_{k=1}^{\infty} x^k = 1/(1-x) - 1$ for $x \in [0, 1)$, and step (b) is by Lemma 3.1. \square

4 Deriving the stick-breaking process from the CRP

We have provided an elementary derivation of the CRP from the stick-breaking process. What about going the other direction? Starting from the CRP, how might one arrive at the stick-breaking representation? Here, we sketch out how the stick-breaking process arises naturally from the CRP. This section should be viewed as a concise exposition of existing results; see Pitman (2006) for reference. In this section only, we appeal to Kingman's paintbox representation of exchangeable partitions, but otherwise our treatment is self-contained.

The CRP is a sequential allocation of customers $i = 1, 2, \dots$ to tables $k = 1, 2, \dots$ in which customer 1 sits at table 1, and each successive customer sits at a currently occupied table with probability proportional to the number of customers at that table, or sits at the next unoccupied table with probability proportional to α . The resulting random partition of customers by table is distributed according to Equation 2.

First, consider table 1. Let $\mathbf{y}_i = 1$ if customer i sits at table 1, and $\mathbf{y}_i = 0$ otherwise. Then $\mathbf{y}_1 = 1$, and $\mathbf{y}_2, \mathbf{y}_3, \dots$ is a two-color Pólya urn process in which $\mathbf{y}_i \mid \mathbf{y}_1, \dots, \mathbf{y}_{i-1} \sim \text{Bernoulli}(\sum_{j=1}^{i-1} \mathbf{y}_j / (\alpha + i - 1))$. Thus,

$$\begin{aligned} p(\mathbf{y}_1, \dots, \mathbf{y}_n) &= \frac{1 \cdot 2 \cdots (s_n - 1) \alpha (\alpha + 1) \cdots (\alpha + n - s_n - 1)}{(\alpha + 1)(\alpha + 2) \cdots (\alpha + n - 1)} \\ &= \frac{B(s_n, \alpha + n - s_n)}{B(1, \alpha)} = \int v^{s_n - 1} (1 - v)^{n - s_n} \text{Beta}(v \mid 1, \alpha) dv, \end{aligned}$$

where $s_n = \sum_{i=1}^n \mathbf{y}_i$. Therefore, the same distribution on $\mathbf{y}_1, \mathbf{y}_2, \dots$ can be generated by drawing $\mathbf{v} \sim \text{Beta}(1, \alpha)$, then setting $\mathbf{y}_1 = 1$ and drawing $\mathbf{y}_2, \mathbf{y}_3, \dots \mid \mathbf{v} \stackrel{\text{iid}}{\sim} \text{Bernoulli}(\mathbf{v})$. Letting $\mathbf{v}_1 = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbf{y}_i$ (the limit exists with probability 1), it follows that $\mathbf{v}_1 \sim \text{Beta}(1, \alpha)$. Note that \mathbf{v}_1 is the asymptotic proportion of customers at table 1.

Now, consider table k . Given the indices of the subsequence of customers that do not sit at tables $1, \dots, k - 1$, the customers in this subsequence sit at table k according to the same urn process as $\mathbf{y}_1, \mathbf{y}_2, \dots$ above, independently of the corresponding urn processes for $1, \dots, k - 1$. Thus, of the customers not at tables $1, \dots, k - 1$, the proportion at table k converges to a $\text{Beta}(1, \alpha)$ random variable, say \mathbf{v}_k , independent of $\mathbf{v}_1, \dots, \mathbf{v}_{k-1}$. Therefore, out of all customers, the proportion at table k converges to $\boldsymbol{\pi}_k := \mathbf{v}_k \prod_{j=1}^{k-1} (1 - \mathbf{v}_j)$ with probability 1. (The preceding urn-based derivation is adapted from Broderick et al., 2013.)

Although we have arrived at the stick-breaking process, our derivation is not yet complete because the partition distribution given $\boldsymbol{\pi}$ as defined above is different than the partition distribution induced by assignments $\mathbf{z}_1, \dots, \mathbf{z}_n \mid \boldsymbol{\pi} \stackrel{\text{iid}}{\sim} \boldsymbol{\pi}$. To establish that they are equivalent, marginally, we use Kingman's paintbox representation for exchangeable partitions.

By Kingman (1978), there exists a random sequence $\mathbf{q} = (\mathbf{q}_1, \mathbf{q}_2, \dots)$ with $\mathbf{q}_1 \geq \mathbf{q}_2 \geq \dots \geq 0$ and $\sum_{j=1}^{\infty} \mathbf{q}_j = 1$ (with probability 1) such that the random partition induced by $\mathbf{z}_1, \dots, \mathbf{z}_n \mid \mathbf{q} \stackrel{\text{iid}}{\sim} \mathbf{q}$ is marginally distributed according to Equation 2; for a concise proof, see Aldous (1985), Prop. 11.9. (Note that in general, $\sum_{j=1}^{\infty} \mathbf{q}_j < 1$ is possible, but not in this case because in the CRP, singleton clusters have probability 0, asymptotically.) Let $\mathbf{z}_1, \mathbf{z}_2, \dots \mid \mathbf{q} \stackrel{\text{iid}}{\sim} \mathbf{q}$ and $\hat{\mathbf{q}}_{nj} = \frac{1}{n} \sum_{i=1}^n \mathbb{1}(\mathbf{z}_i = j)$. With probability 1, for all $j \in \mathbb{N}$, $\hat{\mathbf{q}}_{nj} \rightarrow \mathbf{q}_j$ as $n \rightarrow \infty$, by the law of large numbers. Let $\boldsymbol{\sigma} = (\boldsymbol{\sigma}_1, \boldsymbol{\sigma}_2, \dots)$ be the permutation of \mathbb{N}

such that σ_k is the k th distinct value to appear in $\mathbf{z}_1, \mathbf{z}_2, \dots$; for reference, σ is called a size-biased permutation. In the CRP terminology, customer i is at table k when $\mathbf{z}_i = \sigma_k$, so the asymptotic proportion of customers at table k is \mathbf{q}_{σ_k} . Therefore, by the urn derivation above, $(\mathbf{q}_{\sigma_1}, \mathbf{q}_{\sigma_2}, \dots)$ is equal in distribution to $(\boldsymbol{\pi}_1, \boldsymbol{\pi}_2, \dots)$ where $\boldsymbol{\pi}_k = \mathbf{v}_k \prod_{j=1}^{k-1} (1 - \mathbf{v}_j)$ and $\mathbf{v}_1, \mathbf{v}_2, \dots \stackrel{\text{iid}}{\sim} \text{Beta}(1, \alpha)$. Finally, note that permuting the entries of \mathbf{q} (even via a random permutation that depends on \mathbf{q}) does not affect the marginal distribution of the partition induced by $\mathbf{z}_1, \dots, \mathbf{z}_n | \mathbf{q} \stackrel{\text{iid}}{\sim} \mathbf{q}$. This shows that the partition induced by $\mathbf{z}_1, \dots, \mathbf{z}_n | \boldsymbol{\pi} \stackrel{\text{iid}}{\sim} \boldsymbol{\pi}$ is marginally distributed according to Equation 2.

Acknowledgments

Thanks to David Dunson, Garritt Page, Tamara Broderick, and Steve MacEachern for helpful conversations.

References

- D. J. Aldous. Exchangeability and Related Topics. In *École d'Été de Probabilités de Saint-Flour XIII–1983*, pages 1–198. Springer, 1985.
- C. E. Antoniak. Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems. *The Annals of Statistics*, 2(6):1152–1174, 1974.
- D. M. Blei and M. I. Jordan. Variational inference for Dirichlet process mixtures. *Bayesian Analysis*, 1(1):121–143, 2006.
- T. Broderick, M. I. Jordan, and J. Pitman. Beta processes, stick-breaking and power laws. *Bayesian Analysis*, 7(2):439–476, 2012.
- T. Broderick, M. I. Jordan, and J. Pitman. Cluster and feature modeling from combinatorial stochastic processes. *Statistical Science*, 28(3):289–312, 2013.
- Y. Chung and D. B. Dunson. Nonparametric Bayes conditional distribution modeling with variable selection. *Journal of the American Statistical Association*, 104(488), 2009.
- D. B. Dunson and J.-H. Park. Kernel stick-breaking processes. *Biometrika*, 95(2):307–323, 2008.
- S. Favaro, A. Lijoi, and I. Pruenster. On the stick-breaking representation of normalized inverse Gaussian priors. *Biometrika*, 99(3):663–674, 2012.
- J. E. Griffin and M. J. Steel. Order-based dependent Dirichlet processes. *Journal of the American Statistical Association*, 101(473):179–194, 2006.
- N. L. Hjort. Bayesian analysis for a generalised Dirichlet process prior. *Technical Report, University of Oslo*, 2000.

- H. Ishwaran and L. F. James. Gibbs sampling methods for stick-breaking priors. *Journal of the American Statistical Association*, 96(453), 2001.
- H. Ishwaran and M. Zarepour. Markov chain Monte Carlo in approximate Dirichlet and beta two-parameter process hierarchical models. *Biometrika*, 87(2):371–390, 2000.
- M. Kalli, J. E. Griffin, and S. G. Walker. Slice sampling mixture models. *Statistics and Computing*, 21(1):93–105, 2011.
- J. F. C. Kingman. The representation of partition structures. *Journal of the London Mathematical Society*, 2(18):374–380, 1978.
- S. N. MacEachern. Dependent nonparametric processes. In *ASA Proceedings of the Section on Bayesian Statistical Science*, pages 50–55, 1999.
- S. N. MacEachern. Dependent Dirichlet processes. *Unpublished manuscript, Department of Statistics, The Ohio State University*, 2000.
- J. Paisley. A simple proof of the stick-breaking construction of the Dirichlet Process. Technical report, Princeton University, Department of Computer Science, August 2010.
- J. W. Paisley, A. K. Zaas, C. W. Woods, G. S. Ginsburg, and L. Carin. A stick-breaking construction of the beta process. In *Proceedings of the 27th International Conference on Machine Learning*, pages 847–854, 2010.
- O. Papaspiliopoulos and G. O. Roberts. Retrospective Markov chain Monte Carlo methods for Dirichlet process hierarchical models. *Biometrika*, 95(1):169–186, 2008.
- J. Pitman. *Combinatorial Stochastic Processes*. Springer–Verlag, Berlin, 2006.
- A. Rodriguez and D. B. Dunson. Nonparametric Bayesian models through probit stick-breaking processes. *Bayesian Analysis*, 6(1), 2011.
- J. Sethuraman. A constructive definition of Dirichlet priors. *Statistica Sinica*, 4:639–650, 1994.
- Y. W. Teh, D. Görür, and Z. Ghahramani. Stick-breaking construction for the Indian buffet process. In *International Conference on Artificial Intelligence and Statistics*, pages 556–563, 2007.
- R. Thibaux and M. I. Jordan. Hierarchical beta processes and the Indian buffet process. In *International Conference on Artificial Intelligence and Statistics*, pages 564–571, 2007.
- S. G. Walker, A. Lijoi, and I. Prünster. On rates of convergence for posterior distributions in infinite-dimensional models. *The Annals of Statistics*, 35(2):738–746, 2007.