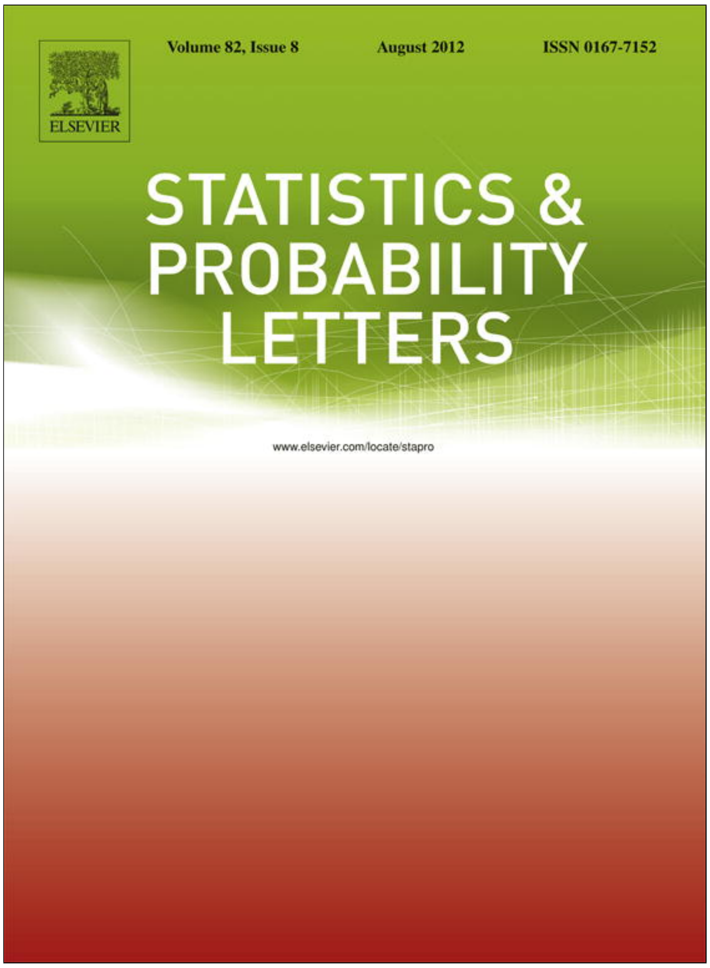
Volume 82, Issue 8
August 2012
ISSN 0167-7152
ELSEVIER
STATISTICS & PROBABILITY LETTERS
www.elsevier.com/locate/stapro

Statistics and Probability Letters 82 (2012) 1521–1529

Contents lists available at SciVerse ScienceDirect

Statistics and Probability Letters

journal homepage: www.elsevier.com/locate/stapro

Hierarchical reinforced urn processes

S. Fortini, S. Petrone *

Department of Decision Sciences, Bocconi University, Via Röntgen 1, 20136 Milano, Italy

ARTICLE INFO

Article history:
Received 21 November 2011
Received in revised form 12 April 2012
Accepted 13 April 2012
Available online 25 April 2012

Keywords:
Partial exchangeability
Mixtures of Markov chains
Hoppe's urn
Infinite hidden Markov models
Bayesian nonparametrics

ABSTRACT

We define a class of reinforced urn processes, based on Hoppe's urn scheme, that are Markov exchangeable, with a countable and possibly unknown state space. This construction extends the reinforced urn processes developed by Muliere et al. (2000) and widely used in Bayesian nonparametric inference and survival analysis. We also shed light on the connections with apparently unrelated constructions, recently proposed in the machine learning literature, such as the infinite hidden Markov model, offering a general framework for a deeper study of their theoretical properties.

1. Introduction

In this paper we define a class of reinforced urn schemes, which generate Markov exchangeable processes and have applications in Bayesian nonparametric inference. Our first aim is to extend the *reinforced urn processes* (RUPs) developed by Muliere et al. (2000) and widely used in Bayesian nonparametric survival analysis, to cover the case of infinite colors and unknown urn composition.

RUPs are informally defined as random walks on a space of Pólya urns. They rely on Diaconis and Freedman's (1980) results for Markov exchangeable sequences and their representation as mixtures of Markov chains, and can be regarded as a simple version of edge reinforced random walks (Coppersmith and Diaconis, 1987); the first results are contained in Pemantle (1988); see also Diaconis and Rolles (2006) for developments for reversible Markov chains, using undirected edges. A graph theoretic and an urn interpretation of the extreme points of Markov exchangeable measures, as a convex set, was given by Zaman (1984). RUPs have been fruitfully applied for Bayesian nonparametric inference in several areas, from survival analysis (Bulla et al., 2007b, 2009) and clinical trials (Bulla et al., 2007a) to credit risk analysis (Cirillo et al., 2010), and in the construction of dependent random measures (Paganoni and Secchi, 2004; Muliere et al., 2005; Trippa et al., 2011). The main property of RUPs is that they are partially exchangeable in the sense of Diaconis and Freedman (1980) or, following the terminology of Zaman (1984) and Zabell (1995), *Markov exchangeable*. Thus, when recurrent, a RUP can be represented as a mixture of Markov chains. The urn scheme characterizes the probability law of the process and therefore the mixing, or prior, distribution. Furthermore, it provides a generating algorithm that can be exploited for computations in Bayesian nonparametric inference.

However, one limitation is that RUPs assume a finite number of colors, which implies a rigid structure of zeros in the random transition matrix, and a known urn composition. Allowing more flexible transitions, and accounting for possible uncertainty on the states (colors) is in fact needed in many applications. We propose an extension of RUPs to the case of countably many colors and unknown initial urn composition. The basic step of our proposal is using Hoppe urns

* Corresponding author. Tel.: +39 0258365602.
E-mail addresses: sandra.fortini@unibocconi.it (S. Fortini), sonia.petrone@unibocconi.it (S. Petrone).

0167-7152/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
doi:10.1016/j.spl.2012.04.012

(Hoppe, 1984, 1987) rather than Pólya urns; more precisely, we suggest a slight variant of Hoppe urn, which generates Pólya sequences (Blackwell and MacQueen, 1973), and is therefore a natural extension of Pólya urns to the case of countably many colors. The proposed class of reinforced Hoppe urn processes includes RUPs as a special case, preserving their main property of being Markov exchangeable.

As a further generalization, we consider hierarchical RUPs, to account for uncertainty on the initial urn composition and on the state space. This is crucial in many applications, specifically to Bayesian inference for hidden Markov models (HMM). In fact, another underlying motivation of our work is to shed light on theoretical connections between RUPs and other, apparently unrelated, urn processes that have been recently proposed in the machine learning literature. In particular, we clarify the theoretical connections with the *infinite hidden Markov model* (iHMM; Beal et al., 2002), which has been developed for Bayesian inference in hidden Markov models, to allow for an unbounded number of states. Van Gael and Ghahramani (2011) discuss the equivalence between the iHMM urn process and the *hierarchical Dirichlet process* of Teh et al. (2006). We aim at giving a more complete picture, underlying theoretical relations with RUPs and Markov exchangeability and proving theoretical properties.

We remind basic results for Markov exchangeable sequences and RUPs in Section 2. Our generalization of RUPs is presented in Section 3, and extended to a hierarchial reinforced urn process in Section 4. Section 5 concludes the paper.

2. Brief review of reinforced urn processes

A RUP (Muliere et al., 2000) is defined by four elements: a countable state space I, a finite set of colors $E = \{c_1, \ldots, c_k\}$, and a law of motion $q : (I \times E) \to I$; finally, to each $x \in I$ it is associated an urn U_x, with known initial composition $\alpha(x) = (\alpha_x(c_1), \ldots, \alpha_x(c_k))$, where $\alpha_x(c) \geq 0$ is the number of balls of color c initially contained in urn U_x, and we let $\alpha_x = \sum_{j=1}^{k} \alpha_x(c_j)$. It is assumed that the law of motion q has the property that, for every $x, y \in I$, there is at most one color $c(x, y) \in E$ such that $q(x, c(x, y)) = y$.

Given these ingredients, a RUP is defined as follows. Fix $X_0 = x_0$ and go to urn U_{x_0}. Pick a ball from U_{x_0} and return it, along with another ball of the same color. If $c \in E$ is the color of the sampled ball, set $X_1 = q(x_0, c)$, and move to urn $U_{q(x_0,c)}$, as determined by the law of motion; and so on. Thus, balls are drawn from each urn according to a Pólya scheme, and one moves across urns according to the given law of motion. The process of colors (X_n) so defined is called RUP, with the four given elements.

The main property of RUPs is that they are Markov exchangeable. Let us briefly remind some basic results. Two sequences $x = (x_0, \ldots, x_n)$ and $y = (y_0, \ldots, y_n)$ in I^{n+1} are *equivalent*, $x \sim y$, if they start from the same state and have the same transitions counts. The sequence (X_n) is Markov exchangeable if $x \sim y$ implies $P(X_0 = x_0, \ldots, X_n = x_n) = P(X_0 = y_0, \ldots, X_n = y_n)$. The sequence (X_n) is *recurrent* if $P(X_n = X_0$ for infinitely many $n) = 1$.

Diaconis and Freedman (1980, Theorem 7) show that a recurrent sequence (X_n) is Markov exchangeable if and only if it is a mixture of Markov chains. That is, given the initial state x_0, there exists a unique probability measure $\mu(\cdot \mid x_0)$ on the space of transition matrices on I, such that

$$P(X_1 = x_1, \ldots, X_n = x_n \mid X_0 = x_0) = \int \prod_{i=1}^{n} \pi_{x_{i-1}}(x_i) d\mu(\pi \mid x_0),$$

where $\pi_i(j) := \pi_{i,j}$ (that is, π_i is the ith row of π, considered as a probability measure). In other words, there exists a random transition matrix Π such that, conditionally on Π and x_0, (X_n) is a Markov chain with transition matrix Π and initial state x_0. The prior distribution of Π is the probability measure μ in the above equation. The proof of the above result is based on the fact that Markov exchangeability and recurrence imply exchangeability of the sequence $(B_1, B_2, \ldots)$ of the successive x_0-blocks (a x_0-block for the sequence (X_n) is a finite sequence of states that begins at x_0 and contains no further x_0).

Muliere et al. (2000) compute the finite-dimensional laws of a RUP (X_n), and show that the process is Markov exchangeable. Therefore, a recurrent RUP (X_n) is a mixture of Markov chains. The reinforced urn scheme characterizes the probability law of the sequence (X_n) and therefore the prior μ. Muliere et al. (2000, Theorem 2.16) show that μ is such that the rows of Π are independent, and the xth row Π_x is a random probability measure on $(y_1 = q(x, c_1), \ldots, y_k = q(x, c_k))$, with probability masses $(\Pi_x(y_1), \ldots, \Pi_x(y_k))$ having a Dirichlet distribution with parameters $(\alpha_x(c_1), \ldots, \alpha_x(c_k))$.

An interesting example of RUP gives a characterization of the beta-Stacy process (Walker and Muliere, 1997), that has many applications in Bayesian nonparametric survival analysis. Suppose that: $I = \{0, 1, 2, \ldots\}$, the set of colors contains only two colors, white and black say, $E = \{w, b\}$, and the law of motion is such that $q(x, b) = x + 1$ and $q(x, w) = 0$, for all $x \in S$. From the previous results, when the resulting RUP is recurrent, it is a mixture of Markov chains. Furthermore, letting T_n be the length on the nth x_0-block, the sequence $(T_n, n \geq 1)$ is exchangeable. Muliere et al. (2000) show that its de Finetti measure is a beta-Stacy process on I with parameters $\{\alpha_j(w), \alpha_j(b), i, j \in I\}$. Thus T_n can be interpreted as the survival time for the n-individual and, assuming that individuals are exchangeable, this construction gives a characterization of the beta-Stacy process as a prior on the survival times. These results can be extended to characterize neutral to the right processes (Doksum, 1974).

Thus, recurrent RUPs provide a general class of mixtures of Markov chains, for which one can explicitly characterize the prior measure. However, a RUP has the restrictions that the initial urn composition must be known and the number of colors has to be finite. The latter assumptions implies that, in each step, the chain can only reach a finite number of states; in other words, each row of the transition matrix has at most k non-zero entries; and the states that are reachable in one step from x have to be fixed a priori. In the next section, we extend the construction to allow for a countable set of colors.

3. Generalization: reinforced Hoppe urns

As said, RUPs are informally defined as random walks on a space of Pólya urns. Informally as well, our idea for extending RUPs to infinite colors is to define a random walk on a space of Hoppe urns. More precisely, we consider the process of *colors* that can be associated to a Hoppe urn, which is a Pólya sequence (Blackwell and MacQueen, 1973), therefore a natural extension of the finite-color Pólya urn scheme.

3.1. Pólya sequences and colored Hoppe urns

Blackwell and MacQueen (1973) define a *Pólya sequence* with parameter $\alpha\, q(\cdot)$ as a sequence of random variables (X_n) satisfying the following predictive scheme: X_1 has distribution q and for any $n \geq 1$,

$$X_{n+1} \mid (X_1, \ldots, X_n) \sim \frac{\alpha}{\alpha+n} q + \frac{1}{\alpha+n} \sum_{i=1}^{n} \delta_{X_i}, \tag{1}$$

where δ_x denotes a measure degenerate on x. They proved that the sequence (X_n) is exchangeable, and its de Finetti measure is a Dirichlet process with parameter αq, $DP(\alpha q)$. The predictive rule (1) extends the one associated to the finite color Pólya urn sampling, therefore it is usually referred as Blackwell and McQueen's urn scheme. However, strictly speaking (1) define a Pólya *sequence*, not an *urn scheme*, since the latter would require physically meaningless urns with infinite colors. In our context, having a naturally interpretable urn scheme is crucial for defining the reinforced process, and in particular for the developments in Section 4. We obtain such urn representation as a colored version of the urn scheme proposed by Hoppe (1984).

Hoppe's urn is defined as follows. Consider sampling from an urn that initially contains $\alpha > 0$ black balls. At time n a ball is picked at random from the urn. If it is black, it is returned together with an additional ball of a previously unobserved color; if it is colored, it is returned together with an additional ball of the same color. Natural numbers are used to label the colors and they are chosen sequentially as the need arises. The sampling generates a process $(S_n,\ n \geq 1)$, where the random variable S_n is the label of the additional ball returned after the nth drawing. Initially there are only black balls, thus $S_1 = 1$; then $S_2 = 1$ or 2, $S_3 = 1, 2$ or 3, etc. For any $n \geq 1$, the random vector $(S_1, \ldots, S_n)$ defines a random partition ρ_n of $\{1, 2, \ldots, n\}$; Hoppe (1984) shows that the sequence (ρ_n) is Markov, with marginal distribution given by the celebrated Ewens sampling formula (Ewens, 1972).

Clearly, the sequence $(S_n,\ n \geq 1)$ generated by the Hoppe urn is not exchangeable. However, we can associate another process to the urn sampling, the process of *colors*, which is exchangeable. If one 'paints' the sequence (S_n), generating the colors at random from a diffuse color distribution q (i.e., as independent and identically distributed (i.i.d.) draws ξ_j from q, where $q(\{x\}) = 0$ for any x), then the resulting sequence of *colors* (X_n) has predictive rule (1), so it is a Pólya sequence with parameter $\alpha\, q$. The colored Hoppe's urn provides a natural way of decomposing the joint distribution of $(X_1, \ldots, X_n)$ in terms of the random partition, generated by $(S_1, \ldots, S_n)$, and the density of the distinct colors (see Antoniak, 1974). That is,

$$p(x_1, \ldots, x_n) = p(s_1, \ldots, s_n) \prod_j q(\xi_j) \tag{2}$$

where the ξ_j are the distinct values in $(x_1, \ldots, x_n)$ and the labels $(s_1, \ldots, s_n)$ identify the random partition generated by $(x_1, \ldots, x_n)$. In terms of the well known Chinese restaurant metaphor, the labels $(S_1, \ldots, S_n)$ generated by the Hoppe urn give the allocation of customers at tables, then tables are painted at random from the color distribution q.

The above scheme assumes that the color distribution is diffuse. To deal with a discrete color distribution, it is simpler to define the process of colors (X_n) more directly, through the following Hoppe-like urn scheme, that we call the *colored Hoppe urn*. As in Hoppe sampling, we draw from an urn that initially contains only α black balls. At time n, a ball is picked at random from the urn and if it is black, it is returned together with an additional ball of a color *drawn at random from a color distribution* q; if it is colored, it is returned together with an additional ball of the same color. We set X_n as the color of the additional ball returned in the urn. If the color distribution q is diffuse, the process (X_n) is the same as described above. However, we also allow a discrete color distribution, say $q = \sum_{j=1}^{k} q(a_j)\delta_{a_j}$, for $k \leq \infty$. This means that the set of colors is known a priori, and coincides with the support $\{a_1, a_2, \ldots, a_k\}$ of q. In this case, it is easy to show that $X_1 \sim q$ and for $n \geq 1$, $P(X_{n+1} = a_j \mid X_1 = x_1, \ldots, X_n = x_n) = (\alpha q(a_j) + \sum_{i=1}^{n} \delta_{x_i}(a_j))/(\alpha + n)$. Therefore, (X_n) is still a Pólya sequence, thus it is exchangeable, and its de Finetti measure is a $DP(\alpha q)$. Notice that, if $p \sim DP(\alpha q)$, then a.s. $p = \sum_{j=1}^{\infty} w_j \delta_{a_j}$, where the weights $(w_1, w_2, \ldots)$ define a random probability measure w on the positive integers, such that $w \sim DP(\alpha q^*)$, with $q^* = \sum_{j=1}^{k} q_0(a_j)\delta_j$.

For brevity, in the sequel, unless differently specified, we will refer to the colored Hoppe urns simply as Hoppe urns.

3.2. Generalized RUPs

We now define a stochastic process (X_n) through a reinforced colored Hoppe urn scheme. Let I denote the finite or countable state space (or color space), which, without loss of generality, we identify with the integers $I = \{1, \ldots, k\}$, where $k \leq \infty$. To each $i \in I$, we associate a Hoppe urn U_i, with α_i black balls and discrete color distribution $p_{0,i}$ on I (we will denote by the same symbol the probability measure and the corresponding probability mass function). Balls are extracted from each urn by Hoppe sampling, but we move across urns as follows. Fix $X_0 = x_0$ according to an initial distribution p_0, go to urn U_{x_0} and pick a ball from it. Since initially it only contains black balls, a color x_1 is sampled from p_{0,x_0} and a ball of color x_1 is added in the urn, together with the black ball. We set $X_1 = x_1$ and move to Hoppe urn U_{x_1}, and so on. We call the process of colors (X_n) so defined a generalized RUP with the defined elements, or a *reinforced Hoppe urn process* (Hoppe RUP), to underline its construction via colored Hoppe's urns.

If for each $i \in I$, the color distribution $p_{0,i}$ has a finite support, then the process (X_n) reduces to a RUP, that with no loss of generality is described by a set of colors that coincides with the state space I and a law of motion that, for each i, is given by $q(i, y) = y$ and is defined only on those colors y which are in the support of $p_{0,i}$.

Example. Suppose that $I = E = \{0, 1, 2, \ldots\}$ and that for each $i \in I$, the color distribution $p_{0,i}$ of urn U_i has positive masses only on $i + 1$ and $x_0 = 0$. The resulting process corresponds to a RUP where the set of colors contains only two colors, white and black say, $E = \{w, b\}$, and the law of motion is such that $q(i, b) = i + 1$ and $q(i, w) = 0$, for all $i \in I$. As discussed in Section 2, this urn scheme gives a characterization of the beta-Stacy process.

A Hoppe RUP maintains the main property of RUPs of being Markov exchangeable.

Proposition 1. *A reinforced Hoppe urn process is Markov exchangeable.*

Proof. For brevity, we will write $x = (x_0, \ldots, x_n)$, $p(x) := P(X_0 = x_0, \ldots, X_n = x_n)$ and $P(X_k = x_k \mid x_0, \ldots, x_{k-1}) := P(X_k = x_k \mid X_0 = x_0, \ldots, X_{k-1} = x_{k-1})$.

If two sequences x and y are equivalent, they have the same number of transitions $t(i, j)$ from state i to state j, for all i, j, and start with the same value. By construction, a transition from state i to state j in one step is not possible (we say not admissible) if color j is not contained in urn U_i, that is, $p_{0,i}(j) = 0$. If the sequence x contains a transition that is not admissible, then $p(x) = 0$. Since $x \sim y$, the same transition is also present in y, and being $p_{0,i}(j) = 0$, it is also not admissible; it follows that $p(y) = p(x) = 0$.

Now suppose that all the transitions in x, and therefore in y, are admissible. Then

$$P(X_1 = x_1, \ldots, X_n = x_n \mid x_0) = P(X_1 = x_1 \mid x_0)P(X_2 = x_2 \mid x_0, x_1) \cdots P(X_n = x_n \mid x_0, \ldots, x_{n-1}).$$

Let $x_1^*, \ldots, x_d^*$ $(d \leq n)$ denote the distinct values in the sequence $(x_0, \ldots, x_{n-1})$; in other words, $x_1^*, \ldots, x_d^*$ denote the urns visited along the sequence. Let $t_i = \sum_j t(i, j)$ be the number of draws from urn U_i in $(x_0, \ldots, x_n)$, and denote by $(x_{j,1}, \ldots, x_{j,t_j})$ the ordered successors of state x_j^* in $(x_0, \ldots, x_n)$, that is, the draws from urn $U_{x_j^*}$.

We can then reorder the factors $P(X_i = x_i \mid x_0, \ldots, x_{i-1})$ in the right hand side above, according to the value of x_{i-1}, obtaining

$$P(X_1 = x_1, \ldots, X_n = x_n \mid x_0) = \prod_{j=1}^{d} p_{0,x_j^*}(x_{j,1}) \frac{\alpha_{x_j^*} p_{0,x_j^*}(x_{j,2}) + \delta_{x_{j,1}}(x_{j,2})}{\alpha_{x_j^*} + 1} \cdots \frac{\alpha_{x_j^*} p_{0,x_j^*}(x_{j,t_j}) + \sum_{i=1}^{t_j - 1} \delta_{x_{j,i}}(x_{j,t_j})}{\alpha_{x_j^*} + t_j - 1}.$$

If $y \sim x$, it follows that $x_n = y_n$ and the set of distinct values is the same in $x = (x_0, \ldots, x_{n-1})$ and $y = (y_0, \ldots, y_{n-1})$, as well as the number of draws $t_{x_j^*}$ from urn $U_{x_j^*}$. Furthermore, the sequence of successors of x_j^* (draws from urn $U_{x_j^*}$) in x and y are the same, up to permutations. Since the above expression is invariant to permutations of the values $(x_{j,1}, \ldots, x_{j,t_j})$, it follows that $P(X_0 = x_0, \ldots, X_n = x_n) = P(X_0 = y_0, \ldots, X_n = y_n)$ and therefore the result. □

From Proposition 1 and the results in Diaconis and Freedman (1980), a recurrent Hoppe RUP is a mixture of Markov chains. In fact, as shown by Zabell (1995), and further developed by Fortini et al. (2002), mixtures of Markov chains can be also characterized in terms of exchangeability properties of the *successors states*. Let's associate to the process (X_n) its *successors matrix*, whose (i, m) element $X_{i,m}$ is the mth successor of state i. That is, $X_{i,m}$ is the value of the process immediately after the mth visit to state i. Informally, a process (X_n) is recurrent and Markov exchangeable if and only if the rows of the successors matrix are infinite, exchangeable sequences. In the reinforced urn scheme, the successors of i are the drawings from urn U_i; thus, being sampled according to an exchangeable scheme, Markov exchangeability is expected to hold, if recurrence properties guarantee that infinite draws are taken from each urn.

More formally, in order to properly define the successors matrix, avoiding having rows of finite length, let us introduce a 'dummy state', ∂, and, if a state is visited only a finite number of times m, let $X_{i,k}$ be equal to ∂ for $k > m$. Thus, the sequence of the successors of state i is well defined, as an infinite sequence of random variables with values in $I^* = I \cup \{\partial\}$. In our reinforced urn scheme, $X_{i,m}$ is the color obtained at the mth draw from urn i, or it is ∂ if urn i is visited less than m times; and

$(X_{i,m}, m \geq 1)$ represents the sequence of draws from urn i, in this extended sense. We say that the process (X_n) is strongly recurrent if, for *any* state $i \in I$, $P(X_n = i \text{ infinitely often} \mid i \text{ is visited}) = 1$, which implies that any state is either never visited or visited infinitely often a.s. In this case, the sequence of successors of any state i is either $(\partial, \partial, \ldots)$, if state i is never visited, or it is an infinite I-valued sequence, if state i is visited infinitely often. It can be shown that, if a process (X_n) is recurrent and Markov exchangeable, then it is also strongly recurrent (Fortini et al., 2002).

On this basis, it can be proved that a process (X_n) is recurrent and Markov exchangeable if and only if the successors matrix is partially exchangeable in the sense of de Finetti; that is, if and only if its distribution is invariant under permutations within rows (see Fortini et al., 2002, Theorem 1). Then, the process is a mixture of Markov chains, that is, there exists a stochastic transition matrix Π on I^*, such that, conditionally on Π, (X_n) is a Markov chain with transition matrix Π; furthermore, the prior distribution on Π is uniquely determined (provided the class of transition matrices is suitably defined; see Fortini et al., 2002). The marginal prior distribution of the ith row Π_i of Π is the de Finetti measure of the exchangeable sequence of successors of state i.

These results suggest that, for a Hoppe RUP (X_n), the prior on Π_i is a Dirichlet process, being the successors of i given by the sequence of draws from the Hoppe urn U_i. However, this holds if state i is visited infinitely often. Note that strong recurrence of (X_n) only guarantees that, if A_i is the event 'urn i is visited infinitely many times' and B_i = 'urn i is never visited', then $P(A_i \cup B_i) = 1$. In fact, the next lemma shows that an even stronger recurrence condition holds, namely the state space can be decomposed into accessible states that are visited infinitely often and non-accessible states that are never visited (a.s.).

For a sequence starting from x_0, let $I_{x_0}^{(0)} = \{x_0\}$ and $I_{x_0}^{(n)} = \{i \in I : p_{0,x}(i) > 0 \text{ for some } x \in I_{x_0}^{(n-1)}\}$, for $n \geq 1$. Since a transition from i to j is possible iff $p_{0,i}(j) > 0$, $I_{x_0}^{(n)}$ is the set of states that are accessible in n steps from x_0. Then, $I_{x_0} = \bigcup_{n=0}^{\infty} I_{x_0}^{(n)}$ is the set of accessible states from x_0. The next lemma shows that *each* state $i \in I_{x_0}$ is visited infinitely often a.s.

Lemma 1. *Suppose that (X_n) is a recurrent Hoppe RUP. For every $x_0 \in I$, $P(i \text{ is visited infinitely often} \mid X_0 = x_0) = 1$ if $i \in I_{x_0}$, while $P(i \text{ is never visited} \mid X_0 = x_0) = 1$ if $i \notin I_{x_0}$.*

Proof. Without loss of generality, we can fix $X_0 = x_0$. We show by induction on n that $P(A_i) = 1$ if $i \in I_{x_0}^{(n)}$, where A_i is the event 'urn i is visited infinitely many times', as defined above. It is true for $n = 0$ by hypothesis. Suppose now that for every $j \in I_{x_0}^{(n)}$, $P(X_n = j \text{ i.o.}) = 1$ and fix $i \in I_{x_0}^{(n+1)}$. Let $j_0 \in I_{x_0}^{(n)}$ be such that $p_{0,j_0}(i) > 0$. Since $X_n = j_0$ infinitely often (with probability one), urn U_{j_0} is visited infinitely often (with probability one). At the kth draw, the probability of a black ball is $\alpha/(\alpha + k - 1)$, independently of the previous draws. Hence the events $D_{j_0,k}$ = 'black ball at the kth draw from urn U_{j_0}' are independent, and $\sum_{k=1}^{\infty} P(D_{j_0,k}) = \sum_{k=1}^{\infty} \alpha/(\alpha + k - 1) = \infty$. By Borel–Cantelli second Lemma, black balls are drawn from urn U_{j_0} an infinite number of times. Let $Z_{j_0,k}$ be the color extracted from p_{0,j_0} when the kth black ball is observed. The random variables $Z_{j_0,k}$, $k \geq 1$, are i.i.d. and $P(Z_{j_0,k} = i) = p_{0,j_0}(i) > 0$. Hence again by Borel Cantelli Lemma, $P(Z_{j_0,k} = i \text{ i.o.}) = 1$. Since the $Z_{j_0,k}$ are some of the X_n's, $P(X_k = i \text{ i.o.}) = 1$. This completes the induction proof. It follows that for every $i \in I_{x_0}$, $P(A_i = 1)$.

The second part of the statement is obvious, since the states outside I_{x_0} are not accessible. □

Lemma 2. *For a recurrent Hoppe RUP starting at x_0, the sequences of successors $(X_{1,n}), (X_{2,n}), \ldots$ are independent. Furthermore, for every state i, the sequence $(X_{i,n})$ is exchangeable, with de Finetti measure $DP(\alpha p_{0,i})$ if $i \in I_{x_0}$, or degenerate on δ_∂ if $i \notin I_{x_0}$.*

Proof. Let A_i = 'urn i is visited infinitely many times'. By Lemma 1, given the starting value x_0, $P(A_i) = 1$ if $i \in I_{x_0}$. The conditional distribution of $(X_{i,n})$ given A_i and $(X_{j,n})$'s, $(j \neq i, n \geq 1)$, coincides with the probability law of a sequence of colors generated from a Hoppe urn, which are exchangeable, with de Finetti measure $DP(\alpha p_{0,i})$. Since $P(A_i) = 1$, the probability law of $(X_{i,n}, n \geq 1)$ coincides with its conditional law given A_i.

If $i \notin I_{x_0}$, then by Lemma 1 it is never visited (with probability one), therefore $X_{i,n} = \partial$ for every n, a.s. The sequence $(X_{i,n})$ is exchangeable, with de Finetti measure degenerate on ∂. □

From the above results we can characterize the mixing measure of a Hoppe RUP.

Proposition 2. *A recurrent Hoppe RUP (X_n) starting at x_0 is a mixture of Markov chains with state space I_{x_0}. That is, there exists a stochastic matrix Π on I_{x_0} such that, conditionally on Π, (X_n) is a Markov chain with transition matrix Π. The prior distribution of Π is such that the rows of Π are independent, with $\Pi_i \sim DP(\alpha_i p_{0,i})$.*

3.3. A Hoppe RUP for categorical variables

When the X_i are categorical variables, so that the colors are thought as unordered labels, it is natural to assume the same α and the same distribution of colors, say $p_0 = \sum_{j=1}^{k} p_0(i)\delta_i$, for each Hoppe urn. The process starts with a random draw x_0 from p_0, then moves to urn U_{x_0} and proceeds as in previous subsection. We prove that the resulting Hoppe RUP is recurrent, therefore the results in Proposition 2 apply.

Lemma 3. *For a Hoppe RUP (X_n), with common urn composition α and color distribution p_0, we have*
(a) $P(E)=1$, *where E is the event "infinitely many black balls are drawn from Hoppe urns".*
(b) *The draws from the color distribution along the Hoppe RUP are an infinite sequence of I-valued r.v.'s (ξ_n), i.i.d. from p_0.*
(c) *The process (X_n) is recurrent.*

Proof. (a) For every $m \geq 1$ let E_m be the event "black ball at the mth draw from Hoppe urns". We want to prove that

$$P(E)=P(\cap_{n=1}^{\infty}\cup_{m=n}^{\infty}E_m)=1.$$

It is sufficient to show that for any m, $\sum_{n=m+1}^{\infty}P(E_n \mid E_m^c\cap E_{m+1}^c\cap\cdots\cap E_{n-1}^c)=\infty$ (cfr. Billingsley, 1995, Problem 4.11). In our case,

$$\sum_{n=m+1}^{\infty}P(E_n \mid E_m^c\cap E_{m+1}^c\cap\cdots\cap E_{n-1}^c)\geq\sum_{n=m+1}^{\infty}\frac{\alpha}{\alpha+n-1}=\infty.$$

(b) By the above results, we can properly define an *infinite* sequence of I-valued r.v.'s (ξ_n), representing the draws from the color distribution along the Hoppe RUP. The ξ_n are then clearly i.i.d., with common distribution p_0.

(c) Without loss of generality, we can suppose that $p_0(x_0)>0$. By the second Borel–Cantelli Lemma, $P(\xi_n=x_0 \text{ i.o.})=1$. Then $P(X_n=x_0 \text{ i.o.})\geq P(\xi_n=x_0 \text{ i.o.})=1$. It follows that $P(X_n=x_0 \text{ i.o.} \mid X_0=x_0)=1$. □

In the sequel, with no real loss of generality, we assume that $p_0(i)>0$ for every $i\in I$. In this case $I_{x_0}=I$, whatever the initial state x_0. By Proposition 2, we have

Proposition 3. *A Hoppe RUP (X_n), with common urn composition α and color distribution p_0, is a mixture of recurrent Markov chains. More specifically, there exists a random transition matrix Π on I such that*

$(X_n)\mid\Pi$ is a Markov chain, with transition matrix Π and initial distribution p_0;
$(\Pi_i, i\in I)\overset{\text{i.i.d}}{\sim}DP(\alpha p_0)$.

4. Hierarchical RUP

In this section we extend the Hoppe RUP with common parameters α and p_0 to the case when the color distribution p_0 is not fixed a priori. We consider two cases: when the set of colors is known, and when the set of colors is unknown a priori.

4.1. Hierarchical Hoppe RUP, known colors

Consider the reinforced urn scheme described in Section 3.3, with color space $I=\{1,2,\ldots\}$, but suppose that the color distribution is not known. In this case, we can assume that the colors are drawn from an auxiliary *oracle urn*, that is a Hoppe urn with γ black balls and discrete color distribution $q=\sum_{i=1}^{\infty}q_i\delta_i$. We start with a draw from the oracle urn; being black, a color x_0 is generated from q and the black ball is returned in the oracle urn, together with an additional ball of color x_0. Then we move to the Hoppe urn U_{x_0}, pick a ball, and if it is black we interrogate the oracle urn. In this case, both the oracle urn and the urn U_{x_0} are reinforced with a ball of the observed color. We then proceed analogously and generate a sequence of colors $(X_n, n\geq 0)$, that we call hierarchical Hoppe RUP with urn compositions α and γ and color distribution q.

Let E be the event "infinitely many black balls are drawn from Hoppe urns". Reasoning as in the proof of Lemma 3, it can be proved that $P(E)=1$. Therefore, the sequence of random variables (ξ_i) describing the draws from the oracle urn is an infinite sequence of draws from a (colored) Hoppe urn with parameters γ and q; thus, it is exchangeable. Let p_0 be the random probability measure on I such that $\xi_i\mid p_0\overset{\text{i.i.d}}{\sim}p_0$; it comes from the urn scheme that $p_0\sim DP(\gamma q)$. Note that this implies that the support of p_0 coincides with the support I of q, so $p_0(i)>0$ a.s. for any $i\in I$.

It follows that, conditionally on p_0, the process (X_n) is a Hoppe RUP with urn composition α and color distribution p_0, as described in Section 3.3. Therefore, by the results obtained there, we have that, conditionally on p_0, (X_n) is a mixture of Markov chains, with a mixing distribution depending on p_0. Integrating out p_0 we get again a mixture of Markov chains. Thus, we have the following

Proposition 4. *The hierarchical Hoppe RUP (X_n), with urn compositions α and γ and discrete color distribution q on I, is a mixture of Markov chains with state space I. More specifically, there exist a random probability measure p_0 and a random transition matrix Π such that*

$(X_n)\mid\Pi, p_0$ is a Markov chain with state space I, transition matrix Π and initial distribution p_0;
$(\Pi_i,\ i\in I)\mid p_0\overset{\text{i.i.d}}{\sim}DP(\alpha p_0)$;
$p_0\sim DP(\gamma q)$.

4.2. Hierarchical RUP with unknown colors

In many applications, specifically in Bayesian inference for hidden Markov models, it is of interest to consider a Markov chain with unknown state space, that is, a finite or countable set $\{\xi_1^*,\xi_2^*,\ldots,\xi_k^*\}$, with $k\leq\infty$, where the ξ_j^*'s are unknown real numbers. The hierarchical Hoppe RUP defined in the previous subsection can be extended to cover this case, by sampling from an oracle urn with a diffuse color distribution on $\mathbb{R}$.

Suppose that, in the hierarchical Hoppe RUP described in Section 4.1, the colors of the oracle Hoppe urn are chosen from a *non-atomic* distribution q on $\mathbb{R}$. The process (X_n) is defined as above, but here, being q diffuse, a new color is generated a.s. when a black ball is drawn from the oracle urn, and a corresponding Hoppe urn is introduced; that is, urns are created as the need occurs. Note that each X_n takes values in $\mathbb{R}$.

It can be proved, analogously to Lemma 3, that the oracle urn is inquired infinitely many times a.s. Therefore, being sampled with the Hoppe scheme, the draws (ξ_n) from the oracle urn are exchangeable, and can be represented as $\xi_n \mid p_0 \overset{\text{i.i.d}}{\sim} p_0$ with $p_0 \sim DP(\gamma q)$. From the properties of the Dirichlet process, $p_0 = \sum_{j=1}^{\infty} w_j \delta_{\xi_j^*}$ a.s., where the atoms ξ_j^* are i.i.d. according to q, and the weights have a stick breaking, or $GEM(\gamma)$, distribution, independently on the ξ_j^*. Note that the atoms are thought as ordered in the order they would appear in a hypothetical sampling from p_0; see e.g. Pitman (1996).

It follows that, conditionally on $p_0 = \sum_{j=1}^{\infty} w_j \delta_{\xi_j^*}$, the process (X_n) is again a Hoppe RUP, as described in Section 3.3, with color distribution p_0. Therefore, by Proposition 3, conditionally on $p_0 = \sum_{j=1}^{\infty} w_j \delta_{\xi_j^*}$, the process (X_n) is a mixture of Markov chains, with state space $I(p_0) = \{\xi_1^*, \xi_2^*, \ldots\}$ specified as the support of p_0, and a random transition matrix whose rows are conditionally i.i.d. according to a $DP(\alpha p_0)$. More specifically, we have the following

Proposition 5. *The sequence (X_n) defined by the hierarchical Hoppe RUP, with urn compositions α and γ, and diffuse color distribution q, is a mixture of Markov chains. More specifically, there exist a random probability measure p_0 and a family $\Pi = (\Pi_{\xi,\xi'}, \xi, \xi' \in \mathbb{R})$ of random variables such that*

$p_0 \sim DP(\gamma q)$;
conditionally on p_0, the restriction $\Pi_{|I(p_0)\times I(p_0)}$ of Π on $I(p_0) \times I(p_0)$, where $I(p_0)$ is the support of p_0, is a random transition matrix whose rows are i.i.d. according to a $DP(\alpha p_0)$;
conditionally on p_0 and Π, the process (X_n) is a Markov chain with state space $I(p_0)$, initial distribution p_0 and transition matrix $\Pi_{|I(p_0)\times I(p_0)}$.

Note that (X_n) is a mixture of Markov chains with countable state space *conditionally on* p_0, where p_0 specifies both the state space and the distribution of the random transition matrix $\Pi_{|I(p_0)\times I(p_0)}$. Integrating the conditional distribution of $(X_n) \mid p_0$ with respect to the probability law of p_0, we have that the process (X_n) is a mixture of Markov chains, but notice that X_n takes values in $\mathbb{R}$. However it can be represented as a mixture of Markov chains, where the mixing distribution is such that the random transition matrix has exchangeable rows, whose joint probability law is a *hierarchical Dirichlet process* prior (Teh et al., 2006).

4.3. Hierarchical RUPs and infinite HMMs

The hierarchical Hoppe RUP described in Section 4.2 is strictly related to the urn scheme suggested by Beal et al. (2002) for Bayesian inference in hidden Markov models with an unbounded number of states, referred as the *infinite hidden Markov model* (iHMM). In fact, it can be regarded as a *colored* version of the iHMM urn process. Roughly speaking, as the colored Hoppe urn of Section 3.1 gives rise to an exchangeable, Pólya urn sequence, for which the Hoppe's urn describes the random partition, similarly our hierarchical reinforced urn scheme generates a Markov exchangeable sequence, in fact a mixture of Markov chains, for which the iHMM describes the color's allocation.

The iHMM is defined as in our hierarchical scheme, but it keeps track of the *labels* of the colors that are successively generated. More precisely, it defines two processes: $(S_i^{(o)},\ i \geq 1)$, that describes the labels of the colors drawn from the oracle urn, in the order they appear; and $(S_n,\ n \geq 0)$, which denotes the label of the additional ball returned in the urn at the nth draw. As before, we start with a draw from the oracle urn; since initially it contains only black balls, a new color with label 1 is generated, and we let $S_1^{(o)} = 1$ and $S_0 = 1$. Then, we create a Hoppe urn U_1 and pick a ball from it; being necessarily black, a new draw is done from the oracle urn, and if it gives a ball of color 1, we set $S_2^{(o)} = 1$ and $S_1 = 1$; if black, a new color with label 2 is generated from q, and we set $S_2^{(o)} = 2$ and $S_1 = 2$; and so on. Thus after n draws, given $S_1 = s_1, \ldots, S_n = s_n, M = m, S_1^{(o)} = s_1^{(o)}, \ldots, S_m^{(o)} = s_m^{(o)}$, where M denotes the random number of draws from the oracle urn, we generate S_{n+1} as follows.

- With probability $t_{i,j}/(\alpha + t_i)$, $S_{n+1} = j$, for $j = 1, \ldots, d_m$,
  where $t_{i,j}$ are the transitions from i to j in $(s_1, \ldots, s_n)$, i.e., the number of balls of label j extracted from urn U_i; $t_i = \sum_j t_{i,j}$ and $d_m = \max(S_1^{(o)}, \ldots, S_m^{(o)})$ is the number of colors generated from the oracle urn;
- with probability $\alpha/(\alpha + t_i)$, a black ball is sampled from U_i, thus a new draw is generated from the oracle urn:

$$S_{m+1}^{(o)} \mid M = m, \quad S_1^{(o)} = s_1^{(o)}, \ldots, S_m^{(o)} = s_m^{(0)} \sim \frac{\gamma}{\gamma + m}\delta_{d_m+1} + \sum_{j=1}^{d_m} \frac{m_j}{\gamma + m}\delta_j,$$

where m_j is the number of balls of color j extracted from the oracle urn. Then we let $S_{n+1} = S_{m+1}^{(o)}$.

Thus

$$S_{n+1} \mid S_1, \ldots, S_n = i, \quad M = m,\ S_1^{(o)}, \ldots, S_M^{(o)} \sim \sum_{j=1}^{d_m} \left(\frac{t_{i,j}}{\alpha + t_i} + \frac{\alpha}{\alpha + t_i} \frac{m_j}{\gamma + m} \right) \delta_j + \left(\frac{\alpha}{\alpha + t_i} \frac{\gamma}{\gamma + m} \right) \delta_{d_m+1}.$$

The process (S_n) is not Markov exchangeable. However, if we paint it with colors $\xi_j^* \overset{\text{i.i.d}}{\sim} q$, with q non-atomic, the resulting process of *colors* (X_n) is a hierarchical Hoppe RUP as defined in Section 4.2, and it is a mixture of Markov chains, for which the urn process characterizes the prior, as shown in Proposition 5.

The role of the labels S_n and $S_n^{(o)}$ is further clarified by a comparison with expression (2). In fact, similarly to (2), the joint distribution of $(X_1, \ldots, X_n)$ is obtained as

$$p(x_1, \ldots, x_n) = p(s_1, \ldots, s_n)\, p(\xi_1, \ldots, \xi_m);$$

but here, in turn, the distribution of colors $(\xi_1, \ldots, \xi_m)$ (the draws from the oracle urn) is expressed by

$$p(\xi_1, \ldots, \xi_m) = p(s_1^{(o)}, \ldots, s_m^{(o)}) \prod_{j=1}^{d_m} q(\xi_j^*).$$

Such joint distribution for the process of colors (X_n) is generated from the conditional distribution

$$X_{n+1} \mid S_1, \ldots, S_n = i, \qquad M = m,$$
$$S_1^{(o)}, \ldots, S_m^{(o)}, \xi_1^*, \ldots, \xi_{d_m}^* \sim \sum_{j=1}^{d_m} \left(\frac{t_{i,j}}{\alpha + t_i} + \frac{\alpha}{\alpha + t_i} \frac{m_j}{\gamma + m_j} \right) \delta_{\xi_j^*} + \frac{\alpha}{\alpha + t_i} \frac{\gamma}{\gamma + m}\, q.$$

Note that the predictive distribution of $X_{n+1} \mid X_1, \ldots, X_n$ is in fact more complex, requiring to average with respect to the conditional distribution of the labels' configurations. Efficient computational methods have been developed for applications in Bayesian inference; see e.g. Van Gael and Ghahramani (2011).

4.4. Related processes

Several extensions of iHMM have been recently proposed (Teh and Jordan, 2010 offer a wide review), and could be usefully reinterpreted in the framework of generalized RUPs. For example, the sticky iHMM (Fox et al., 2011 and references therein), developed to better capture state persistence in hidden Markov models, could be framed as a hierarchical Hoppe RUP where each Hoppe urn U_i initially contains α black balls and k balls of color i, so that infinite draws from U_i are exchangeable, with a $DP(\alpha + k, (\alpha q + k\delta_i)/(\alpha + k))$ de Finetti measure.

It is of interest to discuss comparisons between generalized RUPs and the Indian Buffet Process (IBP), recently proposed by Griffiths and Ghahramani (2006), and widely used in problems involving random binary matrices and infinite latent features. RUPs and IBP have in common the construction of the probability law of interest through a predictive scheme. However, generalized RUPs define Markov exchangeable sequences, while the Indian Buffet process is used to define an exchangeable probability law. Consider a process $(S_i,\ i \geq 1)$, where S_i is a binary sequence, $S_i = (S_{i,1}, S_{i,2}, \ldots)$ with $S_{i,j} \in \{0, 1\}$. S_i represents the choices of individual i among a countable vector of 'features'. A sample of size n gives a binary matrix with ith row $S_i,\ i = 1, \ldots, n$. Roughly speaking, the IBP is used to construct a probability law for the random binary matrix such that the rows are exchangeable, and the probability law of the sequence S_i reflects the idea that features are pure labels and sparse. No temporal dependence is assumed in the IBP, differently from the Markov exchangeability property of RUPs; there is however an analogy that is worth underline. As discussed in Section 3, the Hoppe's urn generates the labels that define random partitions, and once colored generates an exchangeable sequence; analogously, the infinite HMM keeps track of the labels sequences that, once colored, define the Markov exchangeable process described in Section 4.2. Analogously, the IBP generates a sequence (S_i) that is not exchangeable (to obtain exchangeability, one has to consider appropriate equivalence classes of binary matrices; see Griffiths and Ghahramani, 2006, 2011); however, a colored version of the IBP can directly generate an exchangeable sequence. Let us color the features with colors ξ_i i.i.d. from a diffuse color distribution q, and describe the choices of individual i as $\rho_i = \sum_{j=1}^{\infty} S_{i,j}\delta_{\xi_j}$. Then the sequence (ρ_i) is exchangeable.

One could associate a two-colors Pólya urn to each dish in the Indian buffet metaphor, creating the urn when the dish is chosen for the first time by some customer. Subsequent customers will decide whether taking a serving of the dish by sampling from the associated Pólya urn, having a serving if they pick a white ball. Notice that, a sort of walk along the urns is introduced, in some analogy with RUPs; however, moves from one urn to the next one, or enquiries of the 'oracle' Poisson distribution, do not depend on the results of the drawings, neither they are governed by a random mechanism. Thus, no temporal dependence is introduced, differently from what happens in generalized RUPs. However, an extension could be to associate a simple two-colors RUP to each dish. This envisaged construction may have connections with the Markov IBP (Van Gael et al., 2009), that we plan to investigate in future work.

5. Final remarks

We defined a class of reinforced urn processes that are Markov exchangeable, thus, when recurrent, can be represented as mixtures of Markov chains. We have discussed how the class of Hoppe RUPs and hierarchical Hoppe RUPs proposed here includes important urn processes in the Bayesian nonparametric literature, specifically RUPs and iHMMs, offering a general framework for their study. We believe that having clarified the relationship between apparently different urn schemes proposed in the statistical and machine learning literature may offer a clearer understanding of their theoretical properties, in the framework of Diaconis and Freedman (1980). In particular, the properties of the sequence of x_0-blocks that follow from Markov exchangeability and recurrence, could be explored more clearly for iHMMs and for more recent developments.

Acknowledgments

We acknowledge stimulating questions and comments by Steffen Ventz. This work has been partially supported by the Italian Ministry of University and Research, grant 2008MK3AFZ, and by Bocconi University research grants.

References

Antoniak, C.E., 1974. Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems. Annals of Statistics 2, 1152–1174.
Beal, M.J., Ghahramani, Z., Rasmussen, C.E., 2002. The infinite hidden Markov model. In: Machine Learning. MIT Press, pp. 239–245.
Billingsley, P., 1995. Probability and Measure. Wiley, New York.
Blackwell, D., MacQueen, J.B., 1973. Ferguson distributions via Pólya urn schemes. Annals of Statistics 1, 353–355.
Bulla, P., Mezzetti, M., Muliere, P., 2007a. An application of reinforced urn processes to determining maximum tolerated dose. Statistics and Probability Letters 77, 740–747.
Bulla, P., Muliere, P., Walker, S.G., 2007b. Bayesian nonparametric estimation of a bivariate survival function. Statistica Sinica 17, 427–444.
Bulla, P., Muliere, P., Walker, S., 2009. A Bayesian nonparametric estimator of a multivariate survival function. Journal of Statistical Planning and Inference 139, 3639–3648.
Cirillo, P., Hüsler, J., Muliere, P., 2010. A nonparametric urn-based approach to interacting failing systems with an application to credit risk modeling. International Journal of Theoretical and Applied Finance 13, 1223–1240.
Coppersmith, D., Diaconis, P., 1987. Random walk with reinforcement. Unpublished manuscript.
Diaconis, P., Freedman, D., 1980. De Finetti theorem for Markov chains. Annals of Probability 8, 115–130.
Diaconis, P., Rolles, S.W.W., 2006. Bayesian analysis for reversible Markov chains. Annals of Statistics 34, 1270–1292.
Doksum, K.A., 1974. Tailfree and neutral random probabilities and their posterior distributions. Annals of Probability 2, 183–201.
Ewens, W.J., 1972. The sampling theory of selectively neutral alleles. Theoretical Population Biology 3, 87–112.
Fortini, S., Ladelli, L., Petris, G., Regazzini, E., 2002. On mixtures of distributions of Markov chains. Stochastic Processes and their Applications 100, 147–165.
Fox, E.B., Sudderth, E.B., Jordan, M.I, Willsky, A.S., 2011. A sticky HDP–HMM with application to speaker diarization. Annals of Applied Statistics 5, 1020–1056.
Griffiths, T.L., Ghahramani, Z., 2006. Infinite latent features models and the Indian Buffet process. In: Advances in Neural Information Processing Systems, vol. 18. MIT Press, Cambridge, MA.
Griffiths, T.L., Ghahramani, Z., 2011. The Indian Buffet process: an introduction and review. Journal of Machine Learning Research 12, 1185–1224.
Hoppe, F.M., 1984. Pólya-like urns and the Ewens' sampling formula. Journal of Mathematical Biology 20, 91–94.
Hoppe, F.M., 1987. The sampling theory of neutral alleles and an urn model in population genetics. Journal of Mathematical Biology 25, 123–159.
Muliere, P., Secchi, P., Walker, S.G., 2000. Urn schemes and reinforced random walks. Stochastic Processes and their Applications 88, 59–78.
Muliere, P., Secchi, P., Walker, S., 2005. Partially exchangeable processes indexed by the vertices of a k-tree constructed via reinforcement. Stochastic Processes and their Applications 115, 661–667.
Paganoni, A.M., Secchi, P., 2004. Interacting reinforced-urn systems. Advances in Applied Probability 36, 791–804.
Pemantle, R., 1988. Random processes with reinforcement, Ph.D. Thesis, Department of Mathematics, Massachusetts Institute of Technology.
Pitman, J., 1996. Some developments of the Blackwell–MacQueen urn scheme. In: Ferguson, T.S., et al. (Eds.), Statistics, Probability and Game Theory. Papers in honor of David Blackwell. In: Lecture Notes-Monograph Series, vol. 30. Institute of Mathematical Statistics, Hayward, California, pp. 245–267.
Teh, Y.W., Jordan, M.I., 2010. Hierarchical Bayesian nonparametric models with applications. In: Hjort, N.L., Holmes, C., Muller, P., Walker, S.G. (Eds.), Bayesian Nonparametrics: Principles and Practice. Cambridge University Press, Cambridge, UK.
Teh, Y.W., Jordan, M.I., Beal, M.J., Blei, D.M., 2006. Hierarchical Dirichlet processes. Journal of the American Statistical Association 101, 1566–1581.
Trippa, L., Bulla, P., Petrone, S., 2011. Extended Bernstein prior via reinforced urn processes. Annals of the Institute of Statistical Mathematics 63, 481–496.
Van Gael, J., Ghahramani, Z., 2011. Non-parametric hidden Markov models. In: Barber, D., Cemgil, A.T., Chiappa, S. (Eds.), Bayesian Time Series Models. Cambridge University Press, Cambridge, UK, pp. 317–338.
Van Gael, J., Teh, Y.W., Ghahramani, Z., 2009. The infinite factorial hidden Markov model. Advances in Neural Information Processing Systems 21.
Walker, S., Muliere, P., 1997. Beta-Stacy processes and a generalization of the Pólya-urn scheme. Annals of Statistics 25, 1762–1780.
Zabell, S.L., 1995. Characterizing Markov exchangeable sequences. Journal of Theoretical Probability 8, 175–178.
Zaman, A., 1984. Urn models for Markov exchangeability. Annals of Probability 12, 223–229.