

# Truncated Random Measures

TREVOR CAMPBELL<sup>1,2,\*</sup>, JONATHAN H. HUGGINS<sup>1,2,\*\*</sup>,  
JONATHAN P. HOW<sup>3,†</sup>, and TAMARA BRODERICK<sup>2,‡</sup>

<sup>1</sup>*First authorship is shared jointly by T. Campbell and J. H. Huggins.*

<sup>2</sup>*Computer Science and Artificial Intelligence Laboratory (CSAIL), Massachusetts Institute of Technology, Cambridge, MA, USA 02139*

*E-mail:* <sup>\*</sup>[tdjc@mit.edu](mailto:tdjc@mit.edu); <sup>\*\*</sup>[jhuggins@mit.edu](mailto:jhuggins@mit.edu); <sup>‡</sup>[tbroderick@csail.mit.edu](mailto:tbroderick@csail.mit.edu)

<sup>3</sup>*Laboratory for Information and Decision Systems (LIDS), Massachusetts Institute of Technology, Cambridge, MA, USA 02139*

*E-mail:* <sup>†</sup>[jhow@mit.edu](mailto:jhow@mit.edu)

Completely random measures (CRMs) and their normalizations are a rich source of Bayesian nonparametric priors. Examples include the beta, gamma, and Dirichlet processes. In this paper we detail two major classes of sequential CRM representations—*series representations* and *superposition representations*—within which we organize both novel and existing sequential representations that can be used for simulation and posterior inference. These two classes and their constituent representations subsume existing ones that have previously been developed in an ad hoc manner for specific processes. Since a complete infinite-dimensional CRM cannot be used explicitly for computation, sequential representations are often truncated for tractability. We provide truncation error analyses for each type of sequential representation, as well as their normalized versions, thereby generalizing and improving upon existing truncation error bounds in the literature. We analyze the computational complexity of the sequential representations, which in conjunction with our error bounds allows us to directly compare representations and discuss their relative efficiency. We include numerous applications of our theoretical results to commonly-used (normalized) CRMs, demonstrating that our results enable a straightforward representation and analysis of CRMs that has not previously been available in a Bayesian nonparametric context.

*Keywords:* completely random measure, normalized completely random measure, truncation, Bayesian nonparametrics, Poisson point process.

## 1. Introduction

In many data sets, we can view the data points as exhibiting a collection of underlying traits. For instance, each document in the *New York Times* might touch on a number of topics or themes, an individual’s genetic data might be a product of the populations to which their ancestors belonged, or a user’s activity on a social network might be dictated by their varied personal interests. When the traits are not directly observed, a common approach is to model each trait as having some frequency or rate in the broader population ([Airoldi et al., 2014](#)). The inferential goal is to learn these rates as well as

whether—and to what extent—each data point exhibits each trait. Since the traits are unknown a priori, their cardinality is also typically unknown.

As a data set grows larger, we can reasonably expect the number of traits to increase as well. In the cases above, for example, we expect to uncover more topics as we read more documents, more ancestral populations as we examine more individuals’ genetic data, and more unique interests as we observe more individuals on a social network. *Bayesian nonparametric* (BNP) priors provide a flexible, principled approach to creating models in which the number of exhibited traits is random, can grow without bound, and may be learned as part of the inferential procedure. By generating a countable infinity of potential traits—where any individual data point exhibits only finitely many—these models enable growth in the number of observed traits with the size of the data set.

In practice, however, it is impossible to store a countable infinity of random variables in memory or learn the distribution over a countable infinity of variables in finite time. Conjugate priors and likelihoods have been developed (Orbanz, 2010) that theoretically circumvent the infinite representation altogether and perform exact Bayesian posterior inference (Broderick et al., 2017). However, these priors and likelihoods are often just a single piece within a more complex generative model, and ultimately an approximate posterior inference scheme such as Markov Chain Monte Carlo (MCMC) or variational Bayes (VB) is required. These approximation schemes often necessitate a full and explicit representation of the latent variables.

One option is to approximate the infinite-dimensional prior with a related finite-dimensional prior: that is, to replace the infinite collection of random traits by a finite subset of “likely” traits. To do so, first enumerate the countable infinity of traits in the full model and write  $(\psi_k, \theta_k)$  for each paired trait  $\psi_k$  (e.g. a topic in a document) and its rate or frequency  $\theta_k$ . Then the discrete measure  $\Theta := \sum_{k=1}^{\infty} \theta_k \delta_{\psi_k}$  captures the traits/rates in a sequence indexed by  $k$ . The  $(\psi_k, \theta_k)$  pairs are random in the Bayesian model, so  $\Theta$  is a random measure. In many cases, the distribution of  $\Theta$  can be defined by specifying a sequence of simple, familiar distributions for the finite-dimensional  $\psi_k$  and  $\theta_k$ , known as a *sequential representation*. Given a sequential representation of  $\Theta$ , a natural way to choose a subset of traits is to keep the first  $K < \infty$  traits and discard the rest, resulting in an approximate measure  $\Theta_K$ . This approach is called *truncation*. Note that it is also possible to truncate by removing atoms with weights less than a specified threshold (Argiento et al., 2016; Muliere and Tardella, 1998), though this approach is not as easily incorporated in posterior inference algorithms.

Sequential representations have been shown to exist for *completely random measures* (CRMs) (Kingman, 1967; Ferguson and Klass, 1972), a large class of nonparametric priors that includes such popular models as the beta process (Hjort, 1990; Kim, 1999) and the gamma process (Ferguson and Klass, 1972; Kingman, 1975; Brix, 1999; Titsias, 2008; James, 2013). Numerous sequential representations of CRMs have been developed in the literature (Ferguson and Klass, 1972; Bondesson, 1982; Rosiński, 1990, 2001; James, 2014; Broderick et al., 2017). CRM priors are often paired with likelihood processes—such as the Bernoulli process (Thibaux and Jordan, 2007), negative binomial process (Zhou et al., 2012; Broderick et al., 2015), and Poisson likelihood process (Titsias, 2008). The likelihood process determines how much each trait is expressed by each data point. Se-

quential representations also exist for *normalized completely random measures* (NCRMs) (sometimes referred to as normalized random measures with independent increments) (Perman et al., 1992; Perman, 1993; James, 2002; Pitman, 2003; Regazzini et al., 2003; Lijoi and Prünster, 2010; James et al., 2009), which provide random distributions over traits, such as the Dirichlet process (Ferguson, 1973; Sethuraman, 1994). NCRMs are typically paired with a likelihood that assigns each data point to a single trait using the NCRM as a discrete distribution.

Since (N)CRMs have many possible sequential representations, a method is required for determining which to use for the application at hand and, once a representation is selected, for choosing a truncation level. Our main contributions enable the principled selection of both representation and truncation level using approximation error:

1. We provide a comprehensive characterization of the different types of sequential representations for (N)CRMs, filling in many gaps in the literature of sequential representations along the way. We classify these representations into two major groups: *series representations*, which are constructed by transforming a homogeneous Poisson point process; and *superposition representations*, which are the superposition of infinitely many Poisson point processes with finite rate measures. We also introduce two novel sequential representations for (N)CRMs.
2. We provide theoretical guarantees on the approximation error induced when truncating these sequential representations. We give the error as a function of the prior process, the likelihood process, and the level of truncation. While truncation error bounds for (N)CRMs have been studied previously, past work has focused on specific combinations of (N)CRM priors and likelihoods—in particular, the Dirichlet-multinomial (Sethuraman, 1994; Ishwaran and James, 2001; Ishwaran and Zarepour, 2002; Blei and Jordan, 2006), beta-Bernoulli (Paisley et al., 2012; Doshi-Velez et al., 2009), generalized beta-Bernoulli (Roy, 2014), and gamma-Poisson (Roychowdhury and Kulis, 2015) processes. In the current work, we give much more general results for bounding the truncation error.

Our results fill in large gaps in the analysis of truncation error, which is often measured in terms of the  $L^1$  (a.k.a. total variation) distance between the data distributions induced by the full and truncated priors. We provide the first analysis of truncation error for some sequential representations of the beta process with Bernoulli likelihood (Thibaux and Jordan, 2007), for the beta process with negative binomial likelihood (Zhou et al., 2012; Broderick et al., 2015), and for the normalization of the generalized gamma process (Brix, 1999), the  $\sigma$ -stable process, and the generalized inverse gamma (Lijoi et al., 2005; Lijoi and Prünster, 2010) with discrete likelihood. Moreover, even when truncation results already exist in the literature (Ishwaran and James, 2001; Doshi-Velez et al., 2009; Paisley et al., 2012; Roychowdhury and Kulis, 2015), we improve on those error bounds by a factor of two. The reduction arises from our use of the point process machinery of CRMs, circumventing the total variation bound used originally by Ishwaran and James (2001, 2002) upon which most modern truncation analyses are built. We obtain our truncation error guarantees by bounding the probability that data drawn from the full model will use a feature that is not available to the truncated model. Thinking in terms

of this probability provides a more intuitive interpretation of our bounds that can be communicated to practitioners and used to guide them in their choice of truncation level.

The remainder of this paper is organized as follows. In Section 2, we provide background material on CRMs and establish notation. In our first main theoretical section, Section 3, we describe seven different sequential CRM representations, including four series representations and three superposition representations, two of which are novel. Next, we provide a general theoretical analysis of the truncation error for series and superposition representations in Section 4. We provide analogous theory for the normalized versions of each representation in Section 5 via an infinite extension of the ‘‘Gumbel-max trick’’ (Gumbel, 1954; Maddison et al., 2014). We determine the complexity of simulating each representation in Section 6. In Section 7, we summarize our results (Table 1) and provide advice on how to select sequential representations in practice. Proofs for all results developed in this paper are provided in the appendices.

## 2. Background

### 2.1. CRMs and truncation

Consider a Poisson point process on  $\mathbb{R}_+ := [0, \infty)$  with rate measure  $\nu(d\theta)$  such that

$$\nu(\mathbb{R}_+) = \infty \quad \text{and} \quad \int \min(1, \theta) \nu(d\theta) < \infty. \quad (2.1)$$

Such a process generates a countable infinity of values  $(\theta_k)_{k=1}^\infty$ ,  $\theta_k \in \mathbb{R}_+$ , having an almost surely finite sum  $\sum_{k=1}^\infty \theta_k < \infty$ . In a BNP trait model, we interpret each  $\theta_k$  as the rate or frequency of the  $k$ -th trait. Typically, each  $\theta_k$  is paired with a parameter  $\psi_k$  associated with the  $k$ -th trait (e.g., a topic in a document or a shared interest on a social network). We assume throughout that  $\psi_k \in \Psi$  for some space  $\Psi$  and  $\psi_k \stackrel{\text{i.i.d.}}{\sim} G$  for some distribution  $G$ . Constructing a measure by placing mass  $\theta_k$  at atom location  $\psi_k$  results in a *completely random measure* (CRM) (Kingman, 1967). As shorthand, we will write  $\text{CRM}(\nu)$  for the completely random measure generated as just described:

$$\Theta := \sum_k \theta_k \delta_{\psi_k} \sim \text{CRM}(\nu).$$

The trait distribution  $G$  is left implicit in the notation as it has no effect on our results. Further, the possible fixed-location and deterministic components of a CRM (Kingman, 1967) are not considered here for brevity; these components can be added (assuming they are purely atomic) and the analysis modified without undue effort. The CRM prior on  $\Theta$  is typically combined with a likelihood that generates trait counts for each data point. Let  $h(\cdot | \theta)$  be a proper probability mass function on  $\mathbb{N} \cup \{0\}$  for all  $\theta$  in the support of  $\nu$  (though the present work may be easily extended to likelihoods with support in  $\mathbb{R}$ ).

Then a collection of conditionally independent observations  $X_{1:N} := \{X_n\}_{n=1}^N$  given  $\Theta$  are distributed according to the *likelihood process*  $\text{LP}(h, \Theta)$ , i.e.

$$X_n := \sum_k x_{nk} \delta_{\psi_k} \stackrel{\text{i.i.d.}}{\sim} \text{LP}(h, \Theta),$$

if  $x_{nk} \sim h(x | \theta_k)$  independently across  $k$  and i.i.d. across  $n$ . The desideratum that each  $X_n$  expresses a finite number of traits is encoded by the assumption that

$$\int (1 - h(0 | \theta)) \nu(d\theta) < \infty.$$

Since the trait counts are typically latent in a full generative model specification, define the observed data  $Y_n | X_n \stackrel{\text{indep}}{\sim} f(\cdot | X_n)$  for a conditional density  $f$  with respect to a measure  $\mu$  on some space. For instance, if the sequence  $(\theta_k)_{k=1}^\infty$  represents the topic rates in a document corpus,  $X_n$  might capture how many words in document  $n$  are generated from each topic and  $Y_n$  might be the observed collection of words for that document.

Since the sequence  $(\theta_k)_{k=1}^\infty$  is countably infinite, it may be difficult to simulate or perform posterior inference in this model. One approximation scheme is to define the *truncation*  $\Theta_K := \sum_{k=1}^K \theta_k \delta_{\psi_k}$ . Since it is finite, the truncation  $\Theta_K$  can be used for exact simulation or in posterior inference—but some error arises from not using the full CRM  $\Theta$ . To quantify this error, consider its propagation through the above Bayesian model. Define  $Z_{1:N}$  and  $W_{1:N}$  for  $\Theta_K$  analogous to the definitions of  $X_{1:N}$  and  $Y_{1:N}$  for  $\Theta$ :

$$Z_n | \Theta_K \stackrel{\text{i.i.d.}}{\sim} \text{LP}(h, \Theta_K), \quad W_n | Z_n \stackrel{\text{indep}}{\sim} f(\cdot | Z_n), \quad n = 1, \dots, N.$$

A standard approach to measuring the distance between  $\Theta$  and  $\Theta_K$  is to use the  $L^1$  metric between the marginal densities  $p_{N,\infty}$  and  $p_{N,K}$  (with respect to some measure  $\mu$ ) of the final observations  $Y_{1:N}$  and  $W_{1:N}$  (Ishwaran and James, 2001; Doshi-Velez et al., 2009; Paisley et al., 2012):

$$\frac{1}{2} \|p_{N,\infty} - p_{N,K}\|_1 := \frac{1}{2} \int_{y_{1:N}} |p_{N,\infty}(y_{1:N}) - p_{N,K}(y_{1:N})| \mu(dy_{1:N}).$$

All of our bounds on  $\frac{1}{2} \|p_{N,\infty} - p_{N,K}\|_1$  are also bounds on the probability that  $X_{1:N}$  contains a feature that is not in the truncation  $\Theta_K$  (cf. Sections 4 and 5). This interpretation may be easier to digest since it does not depend on the observation model  $f$  and is instead framed in terms of the underlying traits the practitioner is trying to estimate.

## 2.2. The gamma-Poisson process

To illustrate the practical application of the theoretical developments in this work, we provide a number of examples throughout involving the *gamma process* (Brix, 1999),

denoted  $\Gamma\text{P}(\gamma, \lambda, d)$ , with discount parameter  $d \in [0, 1)$ , scale parameter  $\lambda > 0$ , mass parameter  $\gamma > 0$ , and rate measure

$$\nu(d\theta) = \gamma \frac{\lambda^{1-d}}{\Gamma(1-d)} \theta^{-d-1} e^{-\lambda\theta} d\theta.$$

Setting  $d = 0$  yields the undiscounted gamma process (Ferguson and Klass, 1972; Kingman, 1975; Titsias, 2008). The gamma process is often paired with a Poisson likelihood,

$$h(x | \theta) = \frac{\theta^x}{x!} e^{-\theta}.$$

Throughout the present work, we use the rate parametrization of the gamma distribution (to match the gamma process parametrization), for which the density is given by

$$\text{Gam}(x; a, b) = \frac{b^a}{\Gamma(a)} x^{a-1} e^{-bx}.$$

Appendix A provides additional example applications of our theoretical results for two other CRMs: the beta process  $\text{BP}(\gamma, \alpha, d)$  (Teh and Görür, 2009; Broderick et al., 2012) with Bernoulli or negative binomial likelihood, and the beta prime process  $\text{BPP}(\gamma, \alpha, d)$  (Broderick et al., 2017) with odds-Bernoulli likelihood.

### 3. Sequential representations

Sequential representations are at the heart of the study of truncated CRMs. They provide an iterative method that can be terminated at any point to yield a finite approximation to the infinite process, where the choice of termination point determines the accuracy of the approximation. Thus, the natural first step in providing a coherent treatment of truncation analysis is to do the same for sequential representations. In past work, two major classes of sequential representation have been used: *series representations* of the form  $\sum_{k=1}^{\infty} \theta_k \delta_{\psi_k}$ , and *superposition representations* of the form  $\sum_{k=1}^{\infty} \sum_{i=1}^{C_k} \theta_{ki} \delta_{\psi_{ki}}$ , where each inner sum of  $C_k$  atoms is itself a CRM. This section examines four series representations (Ferguson and Klass, 1972; Bondesson, 1982; Rosiński, 1990, 2001) and three superposition representations (two of which are novel) (Broderick et al., 2012, 2017; James, 2014). We show how previously-developed sequential representations for specific CRMs fit into these seven general representations. Finally, we discuss a stochastic mapping procedure that is useful in obtaining new representations from the transformation of others. Proofs for the results in this section may be found in Appendix C.

#### 3.1. Series representations

Series representations arise from the transformation of a homogeneous Poisson point process (Rosiński, 2001). They tend to be somewhat difficult to analyze due to the

dependence between the atoms but also tend to produce very simple representations with small truncation error (cf. Sections 4 and 7). Throughout the paper we let  $\Gamma_k = \sum_{\ell=1}^k E_\ell$ ,  $E_\ell \stackrel{\text{i.i.d.}}{\sim} \text{Exp}(1)$ , be the ordered jumps of a unit-rate homogeneous Poisson process on  $\mathbb{R}_+$ , let  $\nu$  be a measure on  $\mathbb{R}_+$  satisfying the basic conditions in Eq. (2.1), and let  $\psi_k \stackrel{\text{i.i.d.}}{\sim} G$ .

**Inverse-Lévy (Ferguson and Klass, 1972)** Define  $\nu^\leftarrow(u) := \inf \{x : \nu([x, \infty)) \leq u\}$ , the inverse of the tail measuer  $\nu([x, \infty))$ . We say  $\Theta$  has an *inverse-Lévy* representation and write  $\Theta \leftarrow \text{IL-Rep}(\nu)$  if

$$\Theta = \sum_{k=1}^{\infty} \theta_k \delta_{\psi_k}, \quad \text{with} \quad \theta_k = \nu^\leftarrow(\Gamma_k).$$

Ferguson and Klass (1972) showed that  $\Theta \leftarrow \text{IL-Rep}(\nu)$  implies  $\Theta \sim \text{CRM}(\nu)$ . The inverse-Lévy representation is analogous to the inverse CDF method for generating an arbitrary random variable from a uniform random variable, with the homogenous Poisson process playing the role of the uniform random variable. It is also the optimal sequential representation in the sense that the sequence  $\theta_k$  that it generates is non-increasing. While an elegant and general approach, simulating the inverse-Lévy representation is difficult, as inverting the function  $\nu([x, \infty))$  is analytically intractable except in a few cases.

**Example 3.1** (Gamma process,  $\text{GP}(\gamma, \lambda, 0)$ ). We have  $\nu([x, \infty)) = \gamma \lambda E_1(\lambda x)$ , where  $E_1(x) := \int_x^\infty u^{-1} e^{-u} du$  is the exponential integral function (Abramowitz and Stegun, 1964). The inverse-Lévy representation for  $\text{GP}(\gamma, \lambda, 0)$  is thus

$$\Theta = \sum_{k=1}^{\infty} \lambda^{-1} E_1^{-1}(\gamma^{-1} \lambda^{-1} \Gamma_k) \delta_{\psi_k}.$$

Neither  $E_1$  nor its inverse can be computed in closed form, so one must resort to numerical approximations.

**Bondesson (Bondesson, 1982)** We say  $\Theta$  has a *Bondesson representation* and write  $\Theta \leftarrow \text{B-Rep}(c, g)$  if for  $c > 0$  and  $g$  a density on  $\mathbb{R}_+$ ,

$$\Theta = \sum_{k=1}^{\infty} \theta_k \delta_{\psi_k}, \quad \text{with} \quad \theta_k = V_k e^{-\Gamma_k/c}, \quad V_k \stackrel{\text{i.i.d.}}{\sim} g.$$

Theorem 3.1 shows that Bondesson representations can be constructed for a large, albeit restricted, class of CRM rate measures. We offer a novel proof of Theorem 3.1 in Appendix C using the induction strategy introduced by Banjevic et al. (2002). Similar proof ideas are also used to prove truncation error bounds for sequential representations in Section 4. We use a slight abuse of notation for brevity: if  $\nu(d\theta)$  is a measure on  $\mathbb{R}_+$  that is absolutely continuous with respect to Lebesgue measure, then  $\nu(\theta)$  is the density of  $\nu(d\theta)$  with respect to the Lebesgue measure.

**Theorem 3.1** (Bondesson representation (Bondesson, 1982)). *Let  $\nu(d\theta) = \nu(\theta)d\theta$  be a rate measure satisfying Eq. (2.1). If  $\theta\nu(\theta)$  is nonincreasing,  $\lim_{\theta \rightarrow \infty} \theta\nu(\theta) = 0$ , and  $c_\nu := \lim_{\theta \rightarrow 0} \theta\nu(\theta) < \infty$ , then  $g_\nu(v) := -c_\nu^{-1} \frac{d}{dv} [v\nu(v)]$  is a density on  $\mathbb{R}_+$  and*

$$\Theta \leftarrow \text{B-Rep}(c_\nu, g_\nu) \quad \text{implies} \quad \Theta \sim \text{CRM}(\nu).$$

**Example 3.2** (Bondesson representation for  $\text{GP}(\gamma, \lambda, 0)$ ). *The following representation for the gamma process with  $d = 0$  was described by Bondesson (1982) and Banjevic et al. (2002). Since  $\theta\nu(\theta) = \gamma\lambda e^{-\lambda\theta}$  is non-increasing and  $c_\nu = \lim_{\theta \rightarrow 0} \theta\nu(\theta) = \gamma\lambda$ , we obtain  $g_\nu(v) = \lambda e^{-\lambda v} = \text{Exp}(v; \lambda)$ . Thus, it follows from Theorem 3.1 that if  $\Theta \leftarrow \text{B-Rep}(\gamma\lambda, \text{Exp}(\lambda))$ , then  $\Theta \sim \text{GP}(\gamma, \lambda, 0)$ . The condition that  $\theta\nu(\theta)$  is non-increasing fails to hold if  $d > 0$ , so we cannot apply Theorem 3.1 to  $\text{GP}(\gamma, \lambda, d)$  when  $d > 0$ .*

**Thinning** (Rosiński, 1990) Using the nomenclature of Rosiński (2001), we say  $\Theta$  has a *thinning representation* and write  $\Theta \leftarrow \text{T-Rep}(\nu, g)$  if  $g$  is a probability measure on  $\mathbb{R}_+$  such that  $\nu$  is absolutely continuous with respect to  $g$ , i.e.  $\nu \ll g$ , and

$$\Theta = \sum_{k=1}^{\infty} \theta_k \delta_{\psi_k}, \quad \text{with} \quad \theta_k = V_k \mathbb{1} \left( \frac{d\nu}{dg}(V_k) \geq \Gamma_k \right), \quad V_k \stackrel{\text{i.i.d.}}{\sim} g.$$

Rosiński (1990) showed that  $\Theta \leftarrow \text{T-Rep}(\nu, g)$  implies  $\Theta \sim \text{CRM}(\nu)$ . Note that  $\Gamma_k \xrightarrow{\text{a.s.}} \infty$  as  $k \rightarrow \infty$ , so the probability that  $\frac{d\nu}{dg}(V_k) \geq \Gamma_k$  is decreasing in  $k$ . Thus, this representation generates atoms with  $\theta_k = 0$  (which have no effect and can be removed) increasingly frequently and becomes inefficient as  $k \rightarrow \infty$ .

**Example 3.3** (Thinning representation for  $\text{GP}(\gamma, \lambda, d)$ ). *If we let  $g = \text{Gam}(1 - d, \lambda)$ , then the thinning representation for  $\text{GP}(\gamma, \lambda, d)$  is*

$$\Theta = \sum_{k=1}^{\infty} V_k \mathbb{1}(V_k \Gamma_k \leq \gamma) \delta_{\psi_k}, \quad \text{with} \quad V_k \stackrel{\text{i.i.d.}}{\sim} \text{Gam}(1 - d, \lambda).$$

**Rejection** (Rosiński, 2001) Using the nomenclature of Rosiński (2001), we say  $\Theta$  has a *rejection representation* and write  $\Theta \leftarrow \text{R-Rep}(\nu, \mu)$  if  $\mu$  is a measure on  $\mathbb{R}_+$  satisfying Eq. (2.1) and  $\frac{d\nu}{d\mu} \leq 1$ , and

$$\Theta = \sum_{k=1}^{\infty} \theta_k \delta_{\psi_k}, \quad \text{with} \quad \theta_k = V_k \mathbb{1} \left( \frac{d\nu}{d\mu}(V_k) \geq U_k \right), \quad (V_k)_{k \in \mathbb{N}} \sim \text{PoisP}(\mu),$$

$$U_k \stackrel{\text{i.i.d.}}{\sim} \text{Unif}[0, 1].$$

Rosiński (2001) showed that  $\Theta \leftarrow \text{R-Rep}(\nu, \mu)$  implies  $\Theta \sim \text{CRM}(\nu)$ . This representation is very similar to the thinning representation, except that the sequence  $(V_k)_{k \in \mathbb{N}}$  is generated from a Poisson process on  $\mathbb{R}_+$  rather than i.i.d. This allows  $V_k \xrightarrow{\text{a.s.}} 0$  as  $k \rightarrow \infty$ ,

causing the frequency of generating ineffective atoms  $\theta_k = 0$  to decay as  $k \rightarrow \infty$ , assuming  $\mu$  is appropriately chosen such that  $\frac{d\nu}{d\mu}(\theta) \rightarrow 1$  as  $\theta \rightarrow 0$ . This representation can thus be constructed to be more efficient than the thinning representation. We can calculate the efficiency in terms of the expected number of rejections (that is, the number of  $\theta_k$  that are identically zero):

**Proposition 3.2.** *For  $\text{R-Rep}(\nu, \mu)$ , the expected number of rejections is*

$$\mathbb{E} \left[ \sum_{k=1}^{\infty} \mathbb{1}(\theta_k = 0) \right] = \int \left( 1 - \frac{d\nu}{d\mu}(x) \right) \mu(dx).$$

**Remark.** If  $\mu$  and  $\nu$  can be written as densities with respect to Lebesgue measure, then the integral in Proposition 3.2 can be rewritten as  $\int (\mu(x) - \nu(x)) dx$ .

**Example 3.4** (Rejection representation for  $\text{GP}(\gamma, \lambda, 0)$ ). *Following Rosiński (2001), consider  $\mu(d\theta) = \gamma\lambda\theta^{-1}(1+\lambda\theta)^{-1}d\theta$ . We call  $\text{CRM}(\mu)$  the Lomax process,  $\text{LomP}(\gamma, \lambda^{-1})$ , after the related Lomax distribution. We can use the inverse-Lévy method analytically with  $\mu$  since  $\mu^{\leftarrow}(u) = \frac{1}{\lambda(e^{(\gamma\lambda)^{-1}u} - 1)}$ . Thus, the rejection representation of  $\text{GP}(\gamma, \lambda, 0)$  is*

$$\Theta = \sum_{k=1}^{\infty} V_k \mathbb{1}(U_k \leq (1 + \lambda V_k) e^{-\lambda V_k}) \delta_{\psi_k}, \quad \text{with} \quad V_k = \frac{1}{\lambda(e^{(\gamma\lambda)^{-1}\Gamma_k} - 1)}, \quad U_k \stackrel{i.i.d.}{\sim} \text{Unif}[0, 1].$$

Unlike in the thinning construction given in Example 3.3, only a finite number of rates will be set to zero almost surely. In particular, the expected number of rejections is  $\gamma\lambda c_\gamma$ , where  $c_\gamma$  is the Euler-Mascheroni constant.

**Example 3.5** (Rejection representation for  $\text{GP}(\gamma, \lambda, d)$ ,  $d > 0$ ). *For the case of  $d > 0$ , we instead use  $\mu(d\theta) = \gamma \frac{\lambda^{1-d}}{\Gamma(1-d)} \theta^{-1-d} d\theta$ . We can again use the inverse-Lévy method analytically with  $\mu$  since  $\mu^{\leftarrow}(u) = (\gamma' u^{-1})^{1/d}$ , where  $\gamma' := \gamma \frac{\lambda^{1-d}}{d\Gamma(1-d)}$ . The rejection representation is then*

$$\Theta = \sum_{k=1}^{\infty} V_k \mathbb{1}(U_k \leq e^{-\lambda V_k}) \delta_{\psi_k}, \quad \text{with} \quad V_k = (\gamma' \Gamma_k^{-1})^{1/d}, \quad U_k \stackrel{i.i.d.}{\sim} \text{Unif}[0, 1].$$

The expected number of rejections is  $\gamma \frac{\lambda^{1-d}}{d}$ , so the representation is efficient for large  $d$ , but extremely inefficient when  $d$  is small.

## 3.2. Superposition representations

Superposition representations arise as an infinite sum of CRMs with finite rate measure. These tend to be easier to analyze than series representations as they decouple atoms between the summed CRMs, but can produce representations with larger truncation error (cf. Sections 4 and 7). Throughout, let  $\nu$  be a measure on  $\mathbb{R}_+$  satisfying the basic conditions in Eq. (2.1), and let  $\psi_k \stackrel{i.i.d.}{\sim} G$ .

**Decoupled Bondesson** We say  $\Theta$  has a *decoupled Bondesson representation* and write  $\Theta \leftarrow \text{DB-Rep}(c, g, \xi)$  if for  $c > 0, \xi > 0$ , and  $g$  a density on  $\mathbb{R}_+$ ,

$$\Theta = \sum_{k=1}^{\infty} \sum_{i=1}^{C_k} \theta_{ki} \delta_{\psi_{ki}}, \quad \text{with} \quad C_k \stackrel{\text{i.i.d.}}{\sim} \text{Pois}(c/\xi), \quad \theta_{ki} = V_{ki} e^{-T_{ki}}, \quad (3.1)$$

$$T_{ki} \stackrel{\text{indep}}{\sim} \text{Gam}(k, \xi), \quad V_{ki} \stackrel{\text{i.i.d.}}{\sim} g.$$

This is a novel superposition representation, though special cases are already known (Paisley et al., 2010; Roychowdhury and Kulis, 2015). Theorem 3.3 shows that the decoupled Bondesson representation applies to the same class of CRMs as the Bondesson representation from Section 3.1.

**Theorem 3.3** (Decoupled Bondesson representation). *Let  $\nu(d\theta) = \nu(\theta)d\theta$ ,  $c_\nu$ , and  $g_\nu$  be as specified in Theorem 3.1. Then for any fixed  $\xi > 0$ ,*

$$\Theta \leftarrow \text{DB-Rep}(c_\nu, g_\nu, \xi) \quad \text{implies} \quad \Theta \sim \text{CRM}(\nu).$$

The proof of Theorem 3.3 in Appendix C generalizes the arguments from Paisley et al. (2010) and Roychowdhury and Kulis (2015). The free parameter  $\xi$  controls the number of atoms generated for each outer sum index  $k$ ; its principled selection can be made by trading off computational complexity (cf. Section 6) and truncation error (cf. Section 4).

**Example 3.6** (Decoupled Bondesson representation for  $\text{GP}(\gamma, \lambda, 0)$ ). *Arguments paralleling those made in Example 3.2 show that the  $\text{GP}(\gamma, \lambda, 0)$  representation from Roychowdhury and Kulis (2015) follows directly from an application of Theorem 3.3: if  $\Theta \leftarrow \text{DB-Rep}(\gamma\lambda, \text{Exp}(\lambda), \xi)$ , then  $\Theta \sim \text{GP}(\gamma, \lambda, 0)$ . As in the Bondesson representation setting, Theorem 3.3 does not apply to  $\text{GP}(\gamma, \lambda, d)$  when  $d > 0$  because the condition that  $\theta\nu(\theta)$  is non-increasing fails to hold.*

**Size-biased** (Broderick et al., 2017; James, 2014) Let  $\pi(\theta) := h(0|\theta)$ . We say  $\Theta$  has a *size-biased representation* and write  $\Theta \leftarrow \text{SB-Rep}(\nu, h)$  if

$$\Theta = \sum_{k=1}^{\infty} \sum_{i=1}^{C_k} \theta_{ki} \delta_{\psi_{ki}}, \quad \text{with} \quad C_k \stackrel{\text{indep}}{\sim} \text{Pois}(\eta_k),$$

$$\theta_{ki} \stackrel{\text{indep}}{\sim} \frac{1}{\eta_k} \pi(\theta)^{k-1} (1 - \pi(\theta)) \nu(d\theta), \quad (3.2)$$

$$\eta_k := \int \pi(\theta)^{k-1} (1 - \pi(\theta)) \nu(d\theta).$$

Broderick et al. (2017) and James (2014) showed that  $\Theta \leftarrow \text{SB-Rep}(\nu, h)$  implies  $\Theta \sim \text{CRM}(\nu)$ . If the rate measure  $\nu$  and the likelihood  $h$  are selected to be a conjugate

exponential family then, noting that  $\sum_{x=1}^{\infty} h(x|\theta) = 1 - \pi(\theta)$ , the rate  $\theta_{ki}$  can be sampled from a mixture of exponential family distributions:

$$\theta_{ki} | z_{ki} \stackrel{\text{indep}}{\sim} \frac{1}{\eta_{kz_{ki}}} h(z_{ki}|\theta) \pi(\theta)^{k-1} \nu(d\theta), \quad z_{ki} \stackrel{\text{indep}}{\sim} \text{Categorical}((\eta_{kx}/\eta_k)_{x=1}^{\infty}),$$

$$\eta_{kx} := \int h(x|\theta) \pi(\theta)^{k-1} \nu(d\theta).$$

**Example 3.7** (Size-biased representation for  $\Gamma\text{P}(\gamma, \lambda, d)$ ). *For the Gamma process, values for  $\eta_{kx}$  and  $\eta_k$  can be found using integration by parts and the standard gamma distribution integral, while  $\theta_{ki} | z_{ki}$  is sampled from a gamma distribution by inspection:*

$$\eta_{kx} = \frac{\gamma \lambda^{1-d} \Gamma(x-d)}{x! \Gamma(1-d) (\lambda+k)^{x-d}}, \quad \eta_k = \begin{cases} \frac{\gamma \lambda^{1-d}}{d} ((\lambda+k)^d - (\lambda+k-1)^d) & d > 0 \\ \gamma \lambda (\log(\lambda+k) - \log(\lambda+k-1)) & d = 0 \end{cases},$$

$$\theta_{ki} | z_{ki} \stackrel{\text{indep}}{\sim} \text{Gam}(x-d, \lambda+k).$$

**Power-law** We say  $\Theta$  has a *power-law representation* and write  $\Theta \leftarrow \text{PL-Rep}(\gamma, \alpha, d, g)$  if for  $\gamma > 0$ ,  $0 \leq d < 1$ ,  $\alpha > -d$ , and  $g$  a density on  $\mathbb{R}_+$ ,

$$\Theta = \sum_{k=1}^{\infty} \sum_{i=1}^{C_k} \theta_{ki} \delta_{\psi_{\theta_{ki}}}, \quad \text{with } C_k \stackrel{\text{i.i.d.}}{\sim} \text{Pois}(\gamma), \quad \theta_{ki} = V_{ki} U_{kik} \prod_{j=1}^{k-1} (1 - U_{kij}), \quad (3.3)$$

$$V_{ki} \stackrel{\text{i.i.d.}}{\sim} g, \quad U_{kij} \stackrel{\text{indep}}{\sim} \text{Beta}(1-d, \alpha + jd).$$

This is a novel superposition representation, although it was previously developed in the special case of the beta process (where  $g(v) = \delta_1$ ) (Broderick et al., 2012). The name of this representation arises from the fact that it exhibits Types I and II power-law behavior (Broderick et al., 2012) under mild conditions when  $d > 0$ , as we show in Theorem C.1 in the appendix (note, however, that it will not exhibit power-law behavior when  $d = 0$ ). Theorem 3.5 below shows the conditions under which  $\Theta \leftarrow \text{PL-Rep}(\gamma, \alpha, d, g)$  implies  $\Theta \sim \text{CRM}(\nu)$ . Its proof in Appendix C relies on the notion of *stochastic mapping* (Lemma 3.4), a powerful technique for transforming one CRM into another. Note that in Lemma 3.4, the case where  $u$  is a deterministic function of  $\theta$  via the mapping  $u = \tau(\theta)$  may be recovered by setting  $\kappa(\theta, du) = \delta_{\tau(\theta)}$ .

**Lemma 3.4** (CRM stochastic mapping). *Let  $\Theta = \sum_{k=1}^{\infty} \theta_k \delta_{\psi_{\theta_k}} \sim \text{CRM}(\nu)$ . Then for any probability kernel  $\kappa(\theta, du)$ , we have  $\kappa(\Theta) \sim \text{CRM}(\nu_{\kappa})$ , where*

$$\kappa(\Theta) := \sum_{k=1}^{\infty} u_k \delta_{\psi_{u_k}}, \quad u_k | \theta_k \sim \kappa(\theta_k, \cdot), \quad \text{and} \quad \nu_{\kappa}(du) := \int \kappa(\theta, du) \nu(d\theta).$$

**Theorem 3.5** (Power-law representation). *Let  $\nu(d\theta) = \nu(\theta)d\theta$  be a rate measure satisfying Eq. (2.1), and let  $g_{\nu}$  be a density on  $\mathbb{R}_+$  such that*

$$\nu(u) = \int \theta^{-1} g_{\nu}(u\theta^{-1}) \nu_{\text{BP}}(d\theta),$$

where

$$\nu_{\text{BP}}(\text{d}\theta) = \gamma \frac{\Gamma(\alpha + 1)}{\Gamma(1 - d)\Gamma(\alpha + d)} \mathbb{1}[\theta \leq 1] \theta^{-1-d} (1 - \theta)^{\alpha+d-1} \text{d}\theta$$

is the rate measure for the beta process  $\text{BP}(\gamma, \alpha, d)$  from Eq. (A.1). Then

$$\Theta \leftarrow \text{PL-Rep}(\gamma, \alpha, d, g_\nu) \quad \text{implies} \quad \Theta \sim \text{CRM}(\nu).$$

**Example 3.8** (Power-law representation for  $\text{GP}(\gamma, \lambda, d)$ ). If we choose  $g_\nu = \text{Gam}(\lambda, \lambda)$ , then using the change of variable  $w = u(\theta^{-1} - 1)$ ,

$$\begin{aligned} & \int \theta^{-1} g_\nu(u\theta^{-1}) \nu_{\text{BP}}(\text{d}\theta) \\ &= \gamma \lambda \frac{\lambda^\lambda}{\Gamma(1-d)\Gamma(\lambda+d)} u^{\lambda-1} \int \theta^{-\lambda-d-1} e^{-\lambda u \theta^{-1}} (1-\theta)^{\lambda+d-1} \text{d}\theta \text{d}u \\ &= \gamma \lambda \frac{\lambda^\lambda}{\Gamma(1-d)\Gamma(\lambda+d)} u^{-1-d} e^{-\lambda u} \int w^{\lambda+d-1} e^{-\lambda w} \text{d}w \text{d}u \\ &= \gamma \frac{\lambda^{1-d}}{\Gamma(1-d)} u^{-1-d} e^{-\lambda u} \text{d}u. \end{aligned}$$

It follows immediately from Theorem 3.5 that if  $\Theta \leftarrow \text{PL-Rep}(\gamma, \lambda, d, \text{Gam}(\lambda, \lambda))$ , then  $\Theta \sim \text{GP}(\gamma, \lambda, d)$ . To the best knowledge of the authors, this power-law representation for the gamma process is novel.

## 4. Truncation analysis

Each of the sequential representations developed in Section 3 shares a common structural element—an outer infinite sum—which is responsible for generating a countably infinite number of atoms in the CRM. In this section, we terminate these outer sums at a finite *truncation level*  $K \in \mathbb{N}$ , resulting in a *truncated CRM*  $\Theta_K$  possessing a finite number of atoms. We develop upper bounds on the error induced by this truncation procedure. All of the truncated CRM error bounds in this section rely on Lemma 4.1, which is a tightening (by a factor of two) of the bound in Ishwaran and James (2001, 2002) (for its generalization to arbitrary discrete random measures, see Lemma D.1). Proposition 4.2 shows that the bound in Lemma 4.1 is tight without further assumptions on the data likelihood  $f$ .

**Lemma 4.1** (CRM protobound). *Let  $\Theta \sim \text{CRM}(\nu)$ . For any truncation  $\Theta_K$ , if*

$$\begin{aligned} X_n | \Theta &\stackrel{i.i.d.}{\sim} \text{LP}(h, \Theta), & Z_n | \Theta_K &\stackrel{i.i.d.}{\sim} \text{LP}(h, \Theta_K), \\ Y_n | X_n &\stackrel{\text{indep}}{\sim} f(\cdot | X_n), & W_n | Z_n &\stackrel{\text{indep}}{\sim} f(\cdot | Z_n), \end{aligned}$$

then, with  $p_{N,\infty}$  and  $p_{N,K}$  denoting the marginal densities of  $Y_{1:N}$  and  $W_{1:N}$ , respectively,

$$\frac{1}{2} \|p_{N,\infty} - p_{N,K}\|_1 \leq 1 - \mathbb{P}(\text{supp}(X_{1:N}) \subseteq \text{supp}(\Theta_K)),$$

**Proposition 4.2** (Protobound tightness). *If  $G$  is non-atomic and  $\Psi$  is Borel, then for any  $\delta > 0$ , there exists a likelihood  $f$  such that Lemma 4.1 is tight up to a factor of  $1 - \delta$ .*

The proof of all results in this section (including Lemma 4.1 and Proposition 4.2) can be found in Appendix D. All of the provided truncation results use the generative model in Lemma 4.1, and are summarized in Table 1 in Section 7. Throughout this section, for a given likelihood model  $h(x|\theta)$  we define  $\pi(\theta) := h(0|\theta)$  for notational brevity. The asymptotic behavior of truncation error bounds is specified with tilde notation:

$$a(K) \sim b(K), \quad K \rightarrow \infty \quad \iff \quad \lim_{K \rightarrow \infty} \frac{a(K)}{b(K)} = 1.$$

#### 4.1. Series representations

Each of the series representations can be viewed a functional of a standard Poisson point process and a sequence of i.i.d. random variables with some distribution  $g$  on  $\mathbb{R}_+$ . In particular, we may write each in the form

$$\Theta = \sum_{k=1}^{\infty} \theta_k \delta_{\psi_k}, \quad \text{with} \quad \theta_k = \tau(V_k, \Gamma_k), \quad V_k \stackrel{\text{i.i.d.}}{\sim} g, \quad (4.1)$$

where  $\Gamma_k$  are the jumps of a unit-rate homogeneous Poisson point process on  $\mathbb{R}_+$ , and  $\tau : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$  is a non-negative measurable function such that  $\lim_{u \rightarrow \infty} \tau(v, u) = 0$  for  $g$ -almost every  $v$ . The truncated CRM then takes the form

$$\Theta_K := \sum_{k=1}^K \theta_k \delta_{\psi_k}.$$

Theorem 4.3 provides a general truncation error bound for series representations of the form Eq. (4.1), specifies its range, and guarantees that the bound decays to 0 as  $K \rightarrow \infty$ .

**Theorem 4.3** (Series representation truncation error). *The error in approximating a series representation of  $\Theta$  with its truncation  $\Theta_K$  satisfies*

$$0 \leq \frac{1}{2} \|p_{N,\infty} - p_{N,K}\|_1 \leq 1 - e^{-B_{N,K}} \leq 1,$$

where

$$B_{N,K} := \int_0^\infty \left( 1 - \mathbb{E} \left[ \pi(\tau(V, u + G_K))^N \right] \right) du, \quad (4.2)$$

$G_0 := 0$ ,  $G_K \stackrel{\text{indep}}{\sim} \text{Gam}(K, 1)$  for  $K \geq 1$ , and  $V \stackrel{\text{indep}}{\sim} g$ . Furthermore,  $\forall N \in \mathbb{N}$ ,  $\lim_{K \rightarrow \infty} B_{N,K} = 0$ .

**Remark.** An alternate form of  $B_{N,K}$  that is sometimes easier to use in practice can be found by applying the standard geometric series formula to Eq. (4.2), which yields

$$B_{N,K} = \sum_{n=1}^N \int_0^\infty \mathbb{E} \left[ \pi(\tau(V, u + G_K))^{n-1} (1 - \pi(\tau(V, u + G_K))) \right] du.$$

A simplified upper bound on  $B_{N,K}$  can be derived by noting that  $\pi(\theta) \leq 1$ , so

$$B_{N,K} \leq N \int_0^\infty (1 - \mathbb{E}[\pi(\tau(V, u + G_K))]) du. \quad (4.3)$$

This bound usually gives the same asymptotics in  $K$  as Eq. (4.2).

The main task in using Theorem 4.3 to develop a truncation error bound for a series representation is evaluating the integrand in the definition of  $B_{N,K}$ . Thus, we next evaluate the integrand and provide expressions of the truncation error bound for the four series representations outlined in Section 3.1. Throughout the remainder of this section,  $G_K$  is defined as in Theorem 4.3,  $F_0 \equiv 1$ , and  $F_K$  is the CDF of  $G_K$ .

**Inverse-Lévy representation** For this representation we have

$$\tau(v, u) = \nu^{\leftarrow}(u) := \inf \{y : \nu([y, \infty)) \leq u\}.$$

To evaluate the bound in Eq. (4.3), we use the transformation of variables  $x = \nu^{\leftarrow}(u + G_K)$  and the fact that for  $a, b \geq 0$ ,  $\nu^{\leftarrow}(a) \geq b \iff a \leq \nu([b, \infty))$  to conclude that

$$B_{N,K} \leq N \int_0^\infty F_K(\nu[x, \infty))(1 - \pi(x)) \nu(dx). \quad (4.4)$$

Recent work on the inverse-Lévy representation has developed Monte Carlo estimates of the error of the truncated random measure moments for those  $\nu([x, \infty))$  with known inverse  $\nu^{\leftarrow}$  (Arbel and Prünster, 2017). In contrast, the result above provides an explicit bound on the  $L^1$  truncation error. Our bound does not require knowing  $\nu^{\leftarrow}$ , which is often the most challenging aspect of applying the inverse-Lévy representation.

**Example 4.1** (IL-Rep truncation for LomP( $\gamma, \lambda^{-1}$ ) with Poisson likelihood). *Recall from Example 3.4 that the Lomax process LomP( $\gamma, \lambda^{-1}$ ) is the CRM with rate measure  $\nu(d\theta) = \gamma\lambda\theta^{-1}(1 + \lambda\theta)^{-1}d\theta$ , so  $\nu[x, \infty) = \gamma\lambda \log\{1 + (\lambda x)^{-1}\}$ . Using Eq. (4.4), we have*

$$B_{N,K} \leq N\gamma\lambda \int_0^\infty F_K(\gamma\lambda \log\{1 + (\lambda x)^{-1}\})(1 - e^{-x})x^{-1}(1 + \lambda x)^{-1}dx.$$

Since  $F_K(t) \leq t^K/K! \leq (3t/K)^K$ , for any  $a > 0$  the integral is upper bounded by

$$\begin{aligned} & \int_0^a (1 + \lambda x)^{-1}dx + F_K\left(\gamma\lambda \log\left\{1 + \frac{1}{\lambda a}\right\}\right) \int_a^\infty x^{-1}(1 + \lambda x)^{-1}dx \\ & \leq a + F_K\left(\gamma\lambda \log\left\{1 + \frac{1}{\lambda a}\right\}\right) \log\left\{1 + \frac{1}{\lambda a}\right\} \\ & \leq \lambda^{-1}(e^b - 1)^{-1} + b(3\gamma\lambda b/K)^K \quad \text{where } b := \log\{1 + (\lambda a)^{-1}\}. \end{aligned} \quad (4.5)$$

Replacing  $(e^b - 1)^{-1}$  with the approximation  $e^{-b}$  and then setting the two terms in Eq. (4.5) equal, we obtain  $b = KW_0 \left( \left\{ 3\gamma\lambda^{\frac{K+2}{K+1}} (K+1)^{\frac{1}{K+1}} \right\}^{-1} \right)$ , where  $W_0$  is the product logarithm function, i.e.

$$W_0(y) = x \iff xe^x = y. \quad (4.6)$$

Thus, using the fact that  $e^{-t} \leq (e^t - 1)^{-1}$  and  $\lambda^{\frac{K+2}{K+1}} (K+1)^{\frac{1}{K+1}}$  reaches its maximum at  $K = \max(0, e\lambda^{-1} - 1)$ , we conclude that

$$\begin{aligned} B_{N,K} &\leq \frac{2N\gamma[1 + (3\gamma\lambda)^{-1}]}{\exp\left(KW_0\left(\{3\gamma\lambda \max(\lambda, e)\}^{-1}\right)\right) - 1} \\ &\sim 2N\gamma[1 + (3\gamma\lambda)^{-1}]e^{-KW_0(\{3\gamma\lambda \max(\lambda, e)\}^{-1})} \quad K \rightarrow \infty. \end{aligned}$$

**Bondesson representation** For this representation we have

$$\tau(v, u) = ve^{-u/c}, \quad g(dv) = -c^{-1} \frac{d}{dv} (v\nu(v)) dv.$$

Writing the expectation over  $V$  explicitly as an integral with measure  $g(v)dv$ , using the transformation of variables  $u = -c \log x/v$  (so  $x = ve^{-u/c}$ ), and given the definition of  $g(v) = -c^{-1} \frac{d}{dv} (v\nu(v))$  for the Bondesson representation, we have

$$B_{N,K} \leq N \int_0^\infty \left(1 - \mathbb{E} \left[ \pi \left( ve^{-G_K/c} \right) \right] \right) \nu(dv).$$

**Example 4.2** (Truncation of the Bondesson representation for  $\Gamma\text{P}(\gamma, \lambda, 0)$ ). Let  $\tilde{G}_K \stackrel{\mathcal{D}}{=} G_K/(\gamma\lambda)$ . Since  $\pi(\theta) = e^{-\theta}$  and  $c = \gamma\lambda$ , we have

$$\begin{aligned} \int_0^\infty \left(1 - \mathbb{E} \left[ \pi \left( ve^{-G_K/c} \right) \right] \right) \nu(dv) &= \gamma\lambda \mathbb{E} \left[ \int_0^\infty (1 - e^{-ve^{-\tilde{G}_K}}) v^{-1} e^{-\lambda v} dv \right] \\ &= \gamma\lambda \mathbb{E} \left[ \log(1 + e^{-\tilde{G}_K}/\lambda) \right] \\ &\leq \gamma \mathbb{E} \left[ e^{-\tilde{G}_K} \right] = \gamma \left( \frac{\gamma\lambda}{1 + \gamma\lambda} \right)^K. \end{aligned}$$

The second equality follows by using the power series for the exponential integral (Abramowitz and Stegun, 1964, Chapter 5). Thus,

$$B_{N,K} \leq N\gamma \left( \frac{\gamma\lambda}{1 + \gamma\lambda} \right)^K.$$

**Thinning representation** For this representation we have

$$\tau(v, u) = v \mathbb{1} \left[ \frac{d\nu}{dg}(v) \geq u \right], \quad g \text{ any distribution on } \mathbb{R}_+ \text{ s.t. } \nu \ll g.$$

Since  $\pi(0) = 1$  by Lemma B.3, we have that  $1 - \pi(v \mathbb{1}(A)) = (1 - \pi(v)) \mathbb{1}(A)$  for any event  $A$ . Using this fact, we have

$$B_{N,K} \leq N \int_0^\infty (1 - \pi(v)) \int_0^{\frac{d\nu}{dg}(v)} F_K(u) du g(dv). \quad (4.7)$$

Analytic bounds for the thinning representation of specific processes tend to be opaque and notationally cumbersome, so we simply compare its truncation error in Section 7 to the other representations by numerical approximation of Eq. (4.7).

**Rejection representation** Assume that we can use the inverse-Lévy representation to simulate  $\text{PoissP}(\mu)$ . Then for the rejection representation we have

$$\tau(v, u) = \mu^{\leftarrow}(u) \mathbb{1} \left[ \frac{d\nu}{d\mu}(\mu^{\leftarrow}(u)) \geq v \right], \quad g(dv) = \mathbb{1}[0 \leq v \leq 1]dv,$$

where  $\mu$  satisfies  $\nu \ll \mu$ ,  $\frac{d\nu}{d\mu} \leq 1$ , and  $\mu^{\leftarrow}(u) := \inf \{x : \mu([x, \infty)) \leq u\}$ . Using the same techniques as for the thinning and inverse-Lévy representations, we have that

$$B_{N,K} \leq N \int_0^\infty F_K(\mu[x, \infty))(1 - \pi(x)) \nu(dx). \quad (4.8)$$

**Example 4.3** (R-Rep truncation for  $\text{GP}(\gamma, \lambda, 0)$  with Poisson likelihood). *Using Eq. (4.8) and the fact that  $1 - e^{-x} \leq x$ , we have*

$$B_{N,K} \leq N\gamma\lambda \int_0^\infty F_K(\gamma\lambda \log\{1 + (\lambda x)^{-1}\}) e^{-\lambda x} dx. \quad (4.9)$$

Arguing as in Example 4.1, we see that the integral in Eq. (4.9) is upper bounded by

$$\begin{aligned} & \int_0^a e^{-\lambda x} dx + F_K \left( \gamma\lambda \log \left\{ 1 + \frac{1}{\lambda a} \right\} \right) \int_a^\infty e^{-\lambda x} dx \\ & \leq a + \lambda^{-1} F_K \left( \gamma\lambda \log \left\{ 1 + \frac{1}{\lambda a} \right\} \right) \\ & = \lambda^{-1} \left( (e^b - 1)^{-1} + (3\gamma\lambda b/K)^K \right), \end{aligned} \quad (4.10)$$

where  $b := \log\{1 + (\lambda a)^{-1}\}$ . Replacing  $(e^b - 1)^{-1}$  with the approximation  $e^{-b}$  and then setting the two terms in Eq. (4.10) equal to each other, we obtain  $b = KW_0(\{3\gamma\lambda\}^{-1})$  (where  $W_0$  is defined in Eq. (4.6)) and conclude that

$$B_{N,K} \leq \frac{2N\gamma}{e^{KW_0(\{3\gamma\lambda\}^{-1})} - 1} \sim 2N\gamma e^{-KW_0(\{3\gamma\lambda\}^{-1})} \quad K \rightarrow \infty.$$

**Example 4.4** (R-Rep truncation for  $\Gamma\text{P}(\gamma, \lambda, d)$  with Poisson likelihood,  $d > 0$ ). We have

$$B_{N,K} \leq N \frac{\gamma \lambda^{1-d}}{\Gamma(1-d)} \int_0^\infty F_K(\gamma' x^{-d}) (1 - e^{-x}) x^{-1-d} e^{-\lambda x} dx.$$

The integral can be upper bounded as

$$\begin{aligned} & \int_0^a x^{-d} dx + F_K(\gamma' a^{-d}) \int_a^\infty (1 + e^{-x}) x^{-1-d} e^{-\lambda x} dx \\ & \leq (1-d)^{-1} a^{1-d} + \Gamma(-d) (\lambda^d - (1+\lambda)^d) (3\gamma' K^{-1} a^{-d})^K. \end{aligned}$$

Setting the two terms equal and solving for  $a$ , we obtain

$$\begin{aligned} B_{N,K} & \leq 2N \frac{\gamma \lambda^{1-d}}{\Gamma(2-d)} [(1-d)\Gamma(-d)]^{\frac{d(1-d)}{d(1-d)+K}} \left[ \frac{3\gamma \lambda^{1-d}}{d\Gamma(1-d)K} \right]^{\frac{Kd(1-d)}{d(1-d)+K}} \\ & \sim 2N \frac{\gamma \lambda^{1-d}}{\Gamma(2-d)} \left[ \frac{3\gamma \lambda^{1-d}}{d\Gamma(1-d)} \right]^{d(1-d)} K^{-d(1-d)} \quad K \rightarrow \infty. \end{aligned}$$

## 4.2. Superposition representations

For superposition representations, the truncated CRM takes the form

$$\Theta_K := \sum_{k=1}^K \sum_{i=1}^{C_k} \theta_{ki} \delta_{\psi_{ki}}.$$

Let  $\Theta_K^+ := \Theta - \Theta_K$  denote the *tail measure*. By the superposition property of Poisson point processes (Kingman, 1993), the tail measure is itself a CRM with some rate measure  $\nu_K^+$  and is independent of  $\Theta_K$ :

$$\Theta_K^+ = \sum_{k=K+1}^\infty \sum_{i=1}^{C_k} \theta_{ki} \delta_{\psi_{ki}} \sim \text{CRM}(\nu_K^+), \quad \Theta_K^+ \perp\!\!\!\perp \Theta_K, \quad \Theta = \Theta_K + \Theta_K^+. \quad (4.11)$$

The following result provides a general truncation error bound for superposition representations, specifies its range, and guarantees that the bound decays to 0 as  $K \rightarrow \infty$ .

**Theorem 4.4** (Superposition representation truncation error). *The error in approximating a superposition representation of  $\Theta \sim \text{CRM}(\nu)$  with its truncation  $\Theta_K$  satisfies*

$$0 \leq \frac{1}{2} \|p_{N,\infty} - p_{N,K}\|_1 \leq 1 - e^{-B_{N,K}} \leq 1,$$

where

$$B_{N,K} := \int (1 - \pi(\theta)^N) \nu_K^+(d\theta). \quad (4.12)$$

Furthermore,  $\forall N \in \mathbb{N}$ ,  $\lim_{K \rightarrow \infty} B_{N,K} = 0$ .

**Remark.** As for series representations, an alternate form of  $B_{N,K}$  that is sometimes easier to use can be found by applying the standard geometric series formula to Eq. (4.12):

$$B_{N,K} = \sum_{n=1}^N \int \pi(\theta)^{n-1} (1 - \pi(\theta)) \nu_K^+(\mathrm{d}\theta).$$

A simplified upper bound on  $B_{N,K}$  can be derived by noting that  $\pi(\theta) \leq 1$ , so

$$B_{N,K} \leq N \int_0^\infty (1 - \pi(\theta)) \nu_K^+(\mathrm{d}\theta).$$

This bound usually gives the same asymptotics in  $K$  as Eq. (4.12).

The main task in using Theorem 4.4 to develop a truncation error bound for a superposition representation is determining its tail measure  $\nu_K^+$ . In the following, we provide the tail measure for the three superposition representations outlined in Section 3.2.

**Decoupled Bondesson representation** For each point process in the superposition, an average of  $c/\xi$  atoms are generated with independent weights of the form  $V e^{-T_k}$  where  $V \stackrel{\text{indep}}{\sim} g$  and  $T_k \stackrel{\text{indep}}{\sim} \text{Gam}(k, \xi)$ . Therefore, the tail measure is

$$\nu_K^+(\mathrm{d}\theta) = \frac{c}{\xi} \sum_{k=K+1}^{\infty} \tilde{g}_{k,\xi}(\theta) \mathrm{d}\theta,$$

where  $\tilde{g}_{k,\xi}$  is the density of  $V e^{-T_k}$ . The bound for the decoupled Bondesson representation can therefore be expressed as

$$B_{N,K} \leq N \frac{c}{\xi} \sum_{k=K+1}^{\infty} \mathbb{E} [1 - \pi(V e^{-T_k})].$$

**Example 4.5** (Decoupled Bondesson representation truncation for  $\Gamma\text{P}(\gamma, \lambda, 0)$ ). *Using the fact that  $1 - e^{-\theta} \leq \theta$ , we have*

$$\begin{aligned} B_{N,K} &= \frac{N\gamma\lambda}{\xi} \sum_{k=K+1}^{\infty} \mathbb{E}[1 - \pi(V_{k1} e^{-T_{k1}})] \leq \frac{N\gamma\lambda}{\xi} \sum_{k=K+1}^{\infty} \mathbb{E}[V_{k1} e^{-T_{k1}}] \\ &= \frac{N\gamma\lambda}{\xi} \sum_{k=K+1}^{\infty} \frac{1}{\lambda} \left( \frac{\xi}{1 + \xi} \right)^k = N\gamma \left( \frac{\xi}{1 + \xi} \right)^K, \end{aligned}$$

which is equivalent (up to a factor of 2) to the bound in Roychowdhury and Kulis (2015).

**Size-biased representation** The constructive derivation of the size-biased representation (Broderick et al., 2017, proof of Theorem 5.1) immediately yields

$$\nu_K^+(\mathrm{d}\theta) = \pi(\theta)^K \nu(\mathrm{d}\theta).$$

Therefore, the size-biased representation truncation error bound can be expressed using the formula for  $\eta_k$  from Eq. (3.2) as

$$B_{N,K} = \sum_{n=1}^N \int \pi(\theta)^{K+n-1} (1 - \pi(\theta)) \nu(\mathrm{d}\theta) = \sum_{n=1}^N \eta_{K+n}. \quad (4.13)$$

**Example 4.6** (Size-biased representation truncation for  $\Gamma\mathrm{P}(\gamma, \lambda, d)$ ). For  $d > 0$ , the standard gamma integral yields

$$\eta_k = \int \pi(\theta)^{k-1} (1 - \pi(\theta)) \nu(\mathrm{d}\theta) = \frac{\gamma \lambda^{1-d}}{d} ((\lambda + k)^d - (\lambda + k - 1)^d).$$

The sum from Eq. (4.13) is telescoping, so canceling terms,

$$B_{N,K} \leq \frac{\gamma \lambda^{1-d}}{d} ((\lambda + K + N)^d - (\lambda + K)^d) \sim \gamma N \lambda^{1-d} K^{d-1} \quad K \rightarrow \infty,$$

where the asymptotic result follows from Lemma B.4. To analyze the  $d = 0$  case, we use L'Hospital's rule to take the limit of the integral:

$$\lim_{d \rightarrow 0} \int \pi(\theta)^{k-1} (1 - \pi(\theta)) \nu(\mathrm{d}\theta) = \gamma \lambda (\log(\lambda + k) - \log(\lambda + k - 1)).$$

Canceling terms in the telescopic sum yields

$$B_{N,K} \leq \gamma \lambda (\log(\lambda + K + N) - \log(\lambda + K)) \sim \gamma \lambda N K^{-1} \quad K \rightarrow \infty,$$

where the asymptotic result follows from an application of Lemma B.4.

**Power-law representation** For each point process in the superposition, an average of  $\gamma$  atoms are generated with independent weights of the form  $VU_k \prod_{\ell=1}^{k-1} (1 - U_\ell)$ , where  $V \stackrel{\text{indep}}{\sim} g$  and  $U_\ell \stackrel{\text{indep}}{\sim} \text{Beta}(1 - d, \alpha + \ell d)$ . Therefore, the tail measure is

$$\nu_K^+(\mathrm{d}\theta) = \gamma \sum_{k=K+1}^{\infty} \tilde{g}_k(\theta) \mathrm{d}\theta,$$

where  $\tilde{g}_k$  is the density of the random variable  $VU_k \prod_{\ell=1}^{k-1} (1 - U_\ell)$ . The truncation error bound may be expressed as

$$B_{N,K} \leq N \gamma \sum_{k=K+1}^{\infty} \mathbb{E} \left[ 1 - \pi \left( VU_k \prod_{\ell=1}^{k-1} (1 - U_\ell) \right) \right].$$

**Example 4.7** (Power-law representation truncation for  $\text{GP}(\gamma, \lambda, d)$ ). Let  $\beta_k$  be a random variable with density  $\tilde{g}_k$  (with  $\lambda$  in the place of  $\alpha$ ). Using  $1 - e^{-\theta} \leq \theta$ , we have

$$\sum_{k=K+1}^{\infty} \mathbb{E}[1 - \pi(\beta_k)] \leq \sum_{k=K+1}^{\infty} \mathbb{E}[\beta_k] = \mathbb{E} \left[ \sum_{k=K+1}^{\infty} \beta_k \right] = \prod_{k=1}^K \frac{\lambda + kd}{\lambda + kd - d + 1},$$

where the final equality follows from [Ishwaran and James \(2001, Theorem 1\)](#). Thus,

$$B_{N,K} \leq \gamma N \prod_{k=1}^K \frac{\lambda + kd}{\lambda + kd - d + 1} \sim \gamma N \begin{cases} \left(\frac{\lambda}{\lambda+1}\right)^K & d = 0 \\ \frac{\Gamma(\frac{\lambda+1}{d})}{\Gamma(\frac{\lambda+d}{d})} K^{1-d^{-1}} & 0 < d < 1 \end{cases} \quad K \rightarrow \infty, \quad (4.14)$$

where the  $0 < d < 1$  case in [Eq. \(4.14\)](#) follows by [Lemma B.5](#) applied to

$$\prod_{k=1}^K \frac{\lambda + kd}{\lambda + kd - d + 1} = \frac{\Gamma((\lambda+1)/d) \Gamma(\lambda/d + K + 1)}{\Gamma((\lambda+d)/d) \Gamma(\lambda/d + K + d^{-1})}.$$

### 4.3. Stochastic mapping

We now show how truncation bounds developed elsewhere in this paper can be applied to CRM representations that have been transformed using [Lemma 3.4](#). For  $\Theta \sim \text{CRM}(\nu)$ , we denote its transformation by  $\tilde{\Theta} = \kappa(\Theta)$ . For any object defined with respect to  $\Theta$ , the corresponding object for  $\tilde{\Theta}$  is denoted with a tilde. For example, in place of  $N$  and  $X_{1:N}$  (for  $\Theta$ ), we use  $\tilde{N}$  and  $\tilde{X}_{1:\tilde{N}}$  (for  $\tilde{\Theta}$ ). We make  $B_{N,K}$  a function of  $\pi(\theta)$  in the notation of [Proposition 4.5](#); when one applies stochastic mapping to a CRM, one usually also wants to change the likelihood  $h(x|\theta)$ , and thus also changes  $\pi(\theta) = h(0|\theta)$ . The proof of [Proposition 4.5](#) may be found in [Appendix D](#).

**Proposition 4.5** (Truncation error under a stochastic mapping). *Consider a representation for  $\Theta \sim \text{CRM}(\nu)$  with truncation error bound  $B_{N,K}(\pi)$ . Then for any likelihood  $\tilde{h}(x|u)$ , if  $\tilde{\Theta}$  is a stochastic mapping of  $\Theta$  under the probability kernel  $\kappa(\theta, du)$ , its truncation error bound is  $B_{1,K}(\pi_{\kappa, \tilde{N}})$ , where  $\pi_{\kappa, \tilde{N}}(\theta) := \int \tilde{h}(0|u)^{\tilde{N}} \kappa(\theta, du)$ .*

### 4.4. Hyperpriors

In practice, prior distributions are often placed on the hyperparameters of the CRM rate measure (i.e.  $\gamma, \alpha, \lambda, d$ , etc.). We conclude our investigation of CRM truncation error by showing how bounds developed in this section can be modified to account for the use of hyperpriors. Note that we make the dependence of  $B_{N,K}$  on the hyperparameters  $\Phi$  explicit in the notation of [Proposition 4.6](#).

**Proposition 4.6** (CRM truncation error with a hyperprior). *Given hyperparameters  $\Phi$ , consider a representation for  $\Theta | \Phi \sim \text{CRM}(\nu)$ , and let  $B_{N,K}(\Phi)$  be given by Eq. (4.2) (for a series representation) or Eq. (4.12) (for a superposition representation). The error of approximating  $\Theta$  with its truncation  $\Theta_K$  satisfies*

$$0 \leq \frac{1}{2} \|p_{N,\infty} - p_{N,K}\|_1 \leq 1 - e^{-\mathbb{E}[B_{N,K}(\Phi)]} \leq 1.$$

**Example 4.8** (Decoupled Bondesson representation truncation for  $\Gamma\text{P}(\gamma, \lambda, 0)$ ). *A standard choice of hyperprior for the mass  $\gamma$  is a gamma distribution, i.e.  $\gamma \sim \text{Gam}(a, b)$ . Combining Proposition 4.6 and Example 4.5, we have that*

$$\mathbb{E}[B_{N,K}(\Phi)] \leq N \frac{a}{b} \left( \frac{\xi}{\xi + 1} \right)^K.$$

## 5. Normalized truncation analysis

In this section, we provide truncation error bounds for *normalized CRMs* (NCRMs). Examples include the Dirichlet process (Ferguson, 1973), the normalized gamma process (Brix, 1999; James, 2002; Lijoi and Prünster, 2003; Pitman, 2003; Lijoi et al., 2007; Lijoi and Prünster, 2010), and the normalized  $\sigma$ -stable process (Kingman, 1975; Lijoi and Prünster, 2010). Given a CRM  $\Theta$  on  $\Psi$ , we define the corresponding NCRM  $\Xi$  via  $\Xi(S) := \Theta(S)/\Theta(\Psi)$  for each measurable subset  $S \subseteq \Psi$ . Likewise, given a truncated CRM  $\Theta_K$ , we define its normalization  $\Xi_K$  via  $\Xi_K(S) := \Theta_K(S)/\Theta_K(\Psi)$ . Note that any simulation algorithm for  $\Theta_K$  can be used for  $\Xi_K$  by simply normalizing the result. This does not depend on the particular representation of the CRM, and thus applies equally to all the representations in Section 3.

The first step in the analysis of NCRM truncations is to define their approximation error in a manner similar to that of CRM truncations. Since  $\Xi$  and  $\Xi_K$  are both normalized, they are distributions on  $\Psi$ ; thus, observations  $X_{1:N}$  are generated i.i.d. from  $\Xi$ , and  $Z_{1:N}$  are generated i.i.d. from  $\Xi_K$ .  $Y_{1:N}$  and  $W_{1:N}$  have the same definition as for CRMs. As in the developments of Section 4, the theoretical results of this section rely on a general upper bound, provided by Lemma 5.1. Proposition 5.2 shows that the bound in Lemma 5.1 is tight without further assumptions on the data likelihood  $f$ .

**Lemma 5.1** (NCRM protobound). *Let  $\Theta \sim \text{CRM}(\nu)$ , and let its truncation be  $\Theta_K$ . Let their normalizations be  $\Xi$  and  $\Xi_K$  respectively. If*

$$\begin{aligned} X_n | \Xi &\stackrel{i.i.d.}{\sim} \Xi, & Z_n | \Xi_K &\stackrel{i.i.d.}{\sim} \Xi_K, \\ Y_n | X_n &\stackrel{indep}{\sim} f(\cdot | X_n), & W_n | Z_n &\stackrel{indep}{\sim} f(\cdot | Z_n), \end{aligned}$$

then

$$\frac{1}{2} \|p_{N,\infty} - p_{N,K}\|_1 \leq 1 - \mathbb{P}(X_{1:N} \subseteq \text{supp}(\Xi_K)),$$

where  $p_{N,\infty}$ ,  $p_{N,K}$  are the marginal densities of  $Y_{1:N}$  and  $W_{1:N}$ , respectively.

**Proposition 5.2** (Protobound tightness). *If  $G$  is non-atomic and  $\Psi$  is Borel, then for any  $\delta > 0$ , there exists a likelihood  $f$  such that Lemma 5.1 is tight up to a factor of  $1 - \delta$ .*

The analysis of CRMs in Section 4 relied heavily on the Poisson process structure of the rates in  $\Theta$  and  $X_{1:N}$ ; unfortunately, the rates in  $\Xi$  do not possess the same structure and thus lack many useful independence properties (the rates must sum to one). Likewise, sampling  $X_n$  for each  $n$  does not depend on the atoms of  $\Xi$  independently ( $X_n$  randomly selects a single atom based on their rates). Rather than using the basic definitions of the above random quantities to derive an error bound, we decouple the atoms of  $\Xi$  and  $X_{1:N}$  using a technique from extreme value theory. A *Gumbel* random variable  $T$  with location  $\mu \in \mathbb{R}$  and scale  $\sigma > 0$ , denoted  $T \sim \text{Gumbel}(\mu, \sigma)$ , is defined by the cumulative distribution function and corresponding density

$$\mathbb{P}(T \leq t) = e^{-e^{-\frac{t-\mu}{\sigma}}} \quad \text{and} \quad \frac{1}{\sigma} e^{-\left(\frac{t-\mu}{\sigma}\right) - e^{-\left(\frac{t-\mu}{\sigma}\right)}}.$$

An interesting property of the Gumbel distribution is that if one perturbs the log-probabilities of a finite discrete distribution by i.i.d.  $\text{Gumbel}(0, 1)$  random variables, the  $\arg \max$  of the resulting set is a sample from the discrete distribution (Gumbel, 1954; Maddison et al., 2014). This technique is invariant to normalization, as the  $\arg \max$  is invariant to the corresponding constant shift in the log-transformed space. For present purposes, we develop the infinite extension of this result:

**Lemma 5.3** (Infinite Gumbel-max sampling). *Let  $(p_i)_{i=1}^{\infty}$  be a collection of positive numbers such that  $\sum_i p_i < \infty$  and let  $\bar{p}_j := \frac{p_j}{\sum_i p_i}$ . If  $(T_i)_{i=1}^{\infty}$  are i.i.d.  $\text{Gumbel}(0, 1)$  random variables, then  $\arg \max_{i \in \mathbb{N}} T_i + \log p_i$  exists, is unique a.s., and has distribution*

$$\arg \max_{i \in \mathbb{N}} T_i + \log p_i \sim \text{Categorical} \left( (\bar{p}_j)_{j=1}^{\infty} \right).$$

The proof of this result, along with the others in this section, may be found in Appendix E. The utility of Lemma 5.3 is that it allows the construction of  $\Xi$  and  $X_{1:N}$  without the problematic coupling of the underlying CRM atoms due to normalization; rather than dealing directly with  $\Xi$ , we log-transform the rates of  $\Theta$ , perturb them by i.i.d.  $\text{Gumbel}(0, 1)$  random variables, and characterize the distribution of the maximum rate in this process. The combination of this distribution with Lemma 5.3 yields the key proof technique used to develop the truncation bounds in Theorems 5.4 and 5.5. The results presented in this section are summarized in Table 1 in Section 7.

## 5.1. Series representations

The following result provides a general truncation error bound for normalized series representations, specifies its range, and guarantees that it decays to 0 as  $K \rightarrow \infty$ . We again use the general series representation notation from Eq. (4.1), where  $g$  is a distribution on  $\mathbb{R}_+$ , and  $\tau : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$  is a measurable function such that  $\lim_{u \rightarrow \infty} \tau(v, u) = 0$  for  $g$ -almost every  $v$ .

**Theorem 5.4** (Normalized series representation truncation error bound). *The error of approximating a series representation of  $\Xi \sim \text{NCRM}(\nu)$  with its truncation  $\Xi_K$  satisfies*

$$0 \leq \frac{1}{2} \|p_{N,\infty} - p_{N,K}\|_1 \leq 1 - (1 - B_K)^N \leq 1,$$

where

$$B_K := \mathbb{E} \left[ \int_0^\infty J(\Gamma_K, t) \left( \int_0^1 J(\Gamma_K u, t) \, du \right)^{K-1} \left( -\frac{d}{dt} e^{\int_0^\infty (J(u+\Gamma_K, t)-1) du} \right) dt \right], \quad (5.1)$$

$$J(u, t) = \mathbb{E} \left[ e^{-t \cdot \tau(V, u)} \right], \quad V \sim g, \quad \text{and} \quad \Gamma_K \sim \text{Gam}(K, 1).$$

Furthermore,  $\lim_{K \rightarrow \infty} B_K = 0$ .

**Example 5.1** (Dirichlet process, DP( $\gamma$ ), B-Rep). *The Dirichlet process with concentration  $\gamma > 0$  is a normalized gamma process  $\text{NGP}(\gamma, 1, 0)$ . From Example 3.2 we have  $c_\nu = \gamma$  and  $g_\nu = \text{Exp}(1)$ , and from Section 4.1 we have  $\tau(v, u) = ve^{-u/c_\nu}$ . Therefore  $J$  and its antiderivative are*

$$J(u, t) = \mathbb{E} \left[ e^{-tV e^{-u/\gamma}} \right] = \left( 1 + t e^{-u/\gamma} \right)^{-1} \quad \text{and} \quad \int J(u, t) du = \gamma \log \left( e^{u/\gamma} + t \right).$$

Using the antiderivative to evaluate the integrals in the formula for  $B_K$ , writing the expectation over  $\Gamma_K \sim \text{Gam}(K, 1)$  explicitly, and making a change of variables we have

$$B_K = \frac{\gamma^{K+1}}{\Gamma(K)} \int_0^\infty \int_1^\infty \left( \log \left( \frac{s+t}{1+t} \right) \right)^{K-1} (s+t)^{-(\gamma+2)} \, ds \, dt = \left( \frac{\gamma}{1+\gamma} \right)^K,$$

where the last equality is found by multiplying and dividing the integrand by  $(1+t)^{-(\gamma+2)}$ , and making the change of variables from  $s$  to  $x = \log \frac{s+t}{1+t}$ . Therefore, the truncation error can be bounded by

$$\frac{1}{2} \|p_{N,\infty} - p_{N,K}\|_1 \leq 1 - \left( 1 - \left( \frac{\gamma}{\gamma+1} \right)^K \right)^N \sim N \left( \frac{\gamma}{\gamma+1} \right)^K \quad K \rightarrow \infty.$$

The bound in Example 5.2 has exponential decay, and reproduces earlier DP truncation error bound rates due to Ishwaran and James (2001) and Ishwaran and Zarepour (2002). However, the techniques used in past work do not generalize beyond the Dirichlet process, while those developed here apply to any NCRM.

## 5.2. Superposition representations

The following result provides a general truncation error bound for normalized superposition representations, specifies its range, and guarantees that it decays to 0 as  $K \rightarrow \infty$ . We once again rely on the property that the truncation  $\Theta_K$  and tail  $\Theta_K^\dagger$  are mutually independent CRMs, as expressed in Eq. (4.11), with the tail measure denoted  $\nu_K^\dagger$ .

**Theorem 5.5** (Truncation error bound for normalized superposition representations). *The error of approximating a superposition representation of  $\Xi \sim \text{NCRM}(\nu)$  with its truncation  $\Xi_K$  satisfies*

$$0 \leq \frac{1}{2} \|p_{N,\infty} - p_{N,K}\|_1 \leq 1 - (1 - B_K)^N \leq 1,$$

where

$$B_K := \int_0^\infty \left( \int \theta e^{-\theta t} \nu_K^+(\mathrm{d}\theta) \right) e^{\int (e^{-\theta t} - 1) \nu(\mathrm{d}\theta)} \mathrm{d}t. \quad (5.2)$$

Furthermore,  $\lim_{K \rightarrow \infty} B_K = 0$ .

This bound can be applied by using the tail measures derived earlier in Section 4.2.

**Example 5.2** (Dirichlet process, DP( $\gamma$ ), DB-Rep). *As in Example 5.1, we view the Dirichlet process with concentration  $\gamma > 0$  as a normalized gamma process NTP( $\gamma, 1, 0$ ). First, by Lemma B.8, the integral in the exponential is*

$$\exp \left( \int (e^{-t\theta} - 1) \nu(\mathrm{d}\theta) \right) = \exp \left( \gamma \int_0^\infty (e^{-t\theta} - 1) \theta^{-1} e^{-\theta} \mathrm{d}\theta \right) = (t+1)^{-\gamma}.$$

Example 3.2 shows  $c_\nu = \gamma$  and  $g_\nu(v) = e^{-v}$ , and Eq. (C.1) provides the tail measure  $\nu_K^+$  for the decoupled Bondesson representation,

$$\nu_K^+(\mathrm{d}\theta) = \frac{\gamma}{\xi} \sum_{k=K+1}^\infty \frac{\xi^k}{\Gamma(k)} \left( \int_0^1 (-\log x)^{k-1} x^{\xi-2} e^{-\theta x^{-1}} \mathrm{d}x \right) \mathrm{d}\theta.$$

Substituting this result, using Fubini's theorem to swap the order of integration and summation, evaluating the integral over  $\theta$ , and making the substitution  $x = e^{-s}$  yields

$$B_K = \frac{\gamma}{\xi} \sum_{k=K+1}^\infty \frac{\xi^k}{\Gamma(k)} \iint_{s,t \geq 0} \frac{s^{k-1} e^{-(\xi-1)s} (t+1)^{-\gamma}}{(e^s + t)^2} \mathrm{d}s \mathrm{d}t.$$

Noting that  $\forall s \geq 0, e^s \geq 1$ , we have for any  $a \in (0, 1] \cap (0, \gamma)$ ,

$$\begin{aligned} B_K &\leq \frac{\gamma}{\xi} \sum_{k=K+1}^\infty \frac{\xi^k}{\Gamma(k)} \int_0^\infty s^{k-1} e^{-(\xi+a)s} \mathrm{d}s \int_0^\infty (t+1)^{-(\gamma+1-a)} \mathrm{d}t \\ &= \frac{\gamma}{(\gamma-a)\xi} \sum_{k=K+1}^\infty \left( \frac{\xi}{\xi+a} \right)^k = \frac{\gamma}{a(\gamma-a)} \left( \frac{\xi}{\xi+a} \right)^K. \end{aligned}$$

Therefore, for any  $a \in (0, 1] \cap (0, \gamma)$ ,

$$\frac{1}{2} \|p_{N,\infty} - p_{N,K}\|_1 \leq 1 - \left( 1 - \frac{\gamma}{a(\gamma-a)} \left( \frac{\xi}{\xi+a} \right)^K \right)^N \sim \frac{N\gamma}{a(\gamma-a)} \left( \frac{\xi}{\xi+a} \right)^K \quad K \rightarrow \infty.$$

To find the tightest bound, one can minimize with respect to  $a$  given  $\gamma, \xi, K$ .

**Example 5.3** (Normalized gamma process,  $\text{NFP}(\gamma, \lambda, d)$ , SB-Rep). By Lemma B.8, the integral in the exponential is

$$\exp\left(\int (e^{-\theta t} - 1)\nu(d\theta)\right) = \begin{cases} \exp(-\gamma\lambda^{1-d}d^{-1}((t+\lambda)^d - \lambda^d)) & d > 0 \\ \left(\frac{\lambda}{t+\lambda}\right)^{\gamma\lambda} & d = 0, \end{cases} \quad (5.3)$$

and the standard gamma integral yields

$$\int \theta e^{-\theta t} \nu_K^+(d\theta) = \gamma \frac{\lambda^{1-d}}{\Gamma(1-d)} \int \theta^{-d} e^{-(K+t+\lambda)\theta} d\theta = \gamma \lambda^{1-d} (K+t+\lambda)^{d-1}. \quad (5.4)$$

When  $d > 0$ , multiplying the previous two displays and integrating over  $t \geq 0$  yields

$$B_K = \gamma \lambda^{1-d} e^{\gamma\lambda/d} \int_{\lambda}^{\infty} (K+t)^{d-1} e^{-\gamma\lambda^{1-d}t^d/d} dt \leq C_{\gamma,\lambda,d} (K+\lambda)^{d-1},$$

where we have used  $(K+t)^{d-1} \leq (K+\lambda)^{d-1}$  for  $t \geq \lambda$  and the change of variables  $u = \gamma\lambda^{1-d}d^{-1}t^d$  to find that  $C_{\gamma,\lambda,d} = e^{\sigma} \sigma^{1-d^{-1}} \lambda^{1-d} \Gamma(d^{-1}, \sigma)$ , where  $\sigma = \gamma\lambda d^{-1}$  and  $\Gamma(a, x) := \int_x^{\infty} \theta^{a-1} e^{-\theta} d\theta$  is the upper incomplete gamma function. Therefore,

$$\frac{1}{2} \|p_{N,\infty} - p_{N,K}\|_1 \leq 1 - (1 - C_{\gamma,\lambda,d} (K+\lambda)^{d-1})^N \sim N C_{\gamma,\lambda,d} K^{d-1} \quad K \rightarrow \infty.$$

When  $d = 0$ , multiplying Eqs. (5.3) and (5.4) and integrating over  $t \geq 0$  yields

$$\frac{B_K}{\gamma\lambda^{1+\gamma\lambda}} = \int_{\lambda}^{\infty} (K+t)^{-1} t^{-\gamma\lambda} dt \leq \begin{cases} (K+\lambda)^{-1} \left( \frac{\frac{1}{\gamma\lambda}(K+\lambda)^{1-\gamma\lambda} - \lambda^{1-\gamma\lambda}}{1-\gamma\lambda} \right) & \gamma\lambda \neq 1 \\ K^{-1} \log\left(\frac{K+\lambda}{\lambda}\right) & \gamma\lambda = 1, \end{cases}$$

where we obtain the bound for  $\gamma\lambda \neq 1$  by splitting the integral into the intervals  $[\lambda, K+\lambda]$  and  $[K+\lambda, \infty)$  and bounding each section separately, and we obtain the bound for  $\gamma\lambda = 1$  via the transformation  $u = t/(K+t)$ . Therefore, asymptotically

$$\frac{1}{2} \|p_{N,\infty} - p_{N,K}\|_1 \lesssim N \begin{cases} C_{\gamma,\lambda} K^{-\min(1,\gamma\lambda)} & \gamma\lambda \neq 1 \\ \lambda K^{-1} \log K & \gamma\lambda = 1 \end{cases} \quad K \rightarrow \infty,$$

where  $C_{\gamma,\lambda} := \max\left(\frac{\lambda^{\gamma\lambda}}{1-\gamma\lambda}, \frac{\gamma\lambda^2}{\gamma\lambda-1}\right)$ .

Truncation of the  $\text{NFP}(\gamma, \lambda, d)$  has been studied previously: [Argiento et al. \(2016\)](#) threshold the weights of the unnormalized CRM to be beyond a fixed level  $\epsilon > 0$  prior to normalization, and develop error bounds for that method of truncation. These results are not directly comparable to those of the present work due to the different methods of truncation (i.e. sequential representation termination versus weight thresholding).

### 5.3. Hyperpriors

As in the CRM case, we can place priors on the hyperparameters of the NCRM rate measure (i.e.  $\gamma, \alpha, \lambda, d$ , etc.). We conclude our investigation of NCRM truncation error by showing how bounds developed in this section can be modified to account for hyperpriors. Note that we make the dependence of  $B_K$  on the hyperparameters  $\Phi$  explicit in the notation of Proposition 5.6.

**Proposition 5.6** (NCRM truncation error with a hyperprior). *Given hyperparameters  $\Phi$ , consider a representation for  $\Theta | \Phi \sim \text{CRM}(\nu)$ , let  $\Xi | \Phi$  be its normalization, and let  $B_K(\Phi)$  be given by Eq. (5.1) (for a series representation) or Eq. (5.2) (for a superposition representation). The error of approximating  $\Xi$  with its truncation  $\Xi_K$  satisfies*

$$0 \leq \frac{1}{2} \|p_{N,\infty} - p_{N,K}\|_1 \leq 1 - (1 - \mathbb{E}[B_K(\Phi)])^N \leq 1.$$

**Example 5.4** (Dirichlet process, DP( $\gamma$ ), B-Rep). *If we place a Lomax prior on  $\gamma$ , i.e.  $\gamma \sim \text{LomP}(a, 1)$ , then combining Proposition 5.6 and Example 5.1 yields*

$$\frac{1}{2} \|p_{N,\infty} - p_{N,K}\|_1 \leq 1 - \left(1 - \frac{\Gamma(a+1)\Gamma(K+1)}{\Gamma(a+K+1)}\right)^N \sim N\Gamma(a+1)(K+1)^{-a} \quad K \rightarrow \infty.$$

## 6. Simulation and computational complexity

The sequential representations in Section 3 are each generated from a different finite sequence of distributions, resulting in a different expected computational cost for the same truncation level. Thus, the truncation level itself is not an appropriate parameter with which to compare the error bounds for different representations and we require a characterization of the computational cost. We investigate the mean complexity  $\mathbb{E}[R]$  of each representation, where  $R$  is the number of random variables sampled, as a function of the truncation level for each of the representations in Section 3.

We begin with the series representations. For each value of  $k = 1, \dots, K$ , each series representation generates a single trait  $\psi_k \sim G$  and a rate  $\theta_k$  composed of some transformation of random variables. Thus, all of the series representations in this work satisfy  $\mathbb{E}[R] = rK$  for some constant  $r$ : by inspection, the inverse-Lévy representation has  $r = 2$ , and all the remaining series representations have  $r = 3$ .

The superposition representations, on the other hand, generate a Poisson random variable to determine the number of atoms at each value of  $k = 1, \dots, K$ , and then generate those atoms. Therefore, the mean simulation complexity takes the form  $\mathbb{E}[R] = \sum_{k=1}^K 1 + r_k \mathbb{E}[C_k]$  for some constants  $r_k$  that might depend on the value of  $k$ . For the decoupled Bondesson representation,  $r_k = 3$  since each atom requires generating three values ( $\psi_{ki}$ ,  $V_{ki}$ , and  $T_{ki}$ ), and  $\mathbb{E}[C_k] = c/\xi$ , so  $\mathbb{E}[R] = \left(\frac{3c}{\xi} + 1\right) K$ . For the size-biased representation,  $r_k = 3$  since each atom requires generating three values ( $\psi_{ki}$ ,  $z_{ki}$ , and

**Table 1.** Asymptotic error bounds and simulation cost summary. Error bounds are presented up to a constant that varies between models. Be = Bernoulli, OBe = odds Bernoulli, Poi = Poisson.

Rep.	Random Measure	$h$	Asymptotic Error Bound	Complexity
IL	LomP( $\gamma, \lambda^{-1}$ )	Poi	$N e^{-KW_0(\{3\gamma\lambda \max(\lambda, e)\}^{-1})}$	$2K$
B	BP( $\gamma, \lambda \geq 1, 0$ )	Be	$N\gamma \left(\frac{\gamma\lambda}{\gamma\lambda+1}\right)^K$	$3K$
	$\Gamma P(\gamma, \lambda, 0)$	Poi		
	BPP( $\gamma, \lambda, 0$ )	OBe		
	DP( $\gamma$ )	—	$N \left(\frac{\gamma}{\gamma+1}\right)^K$	
T	—	—	See Eq. (4.7)	$3K$
R	BP( $\gamma, \lambda, 0$ ) $\Gamma P(\gamma, \lambda, 0)$ BPP( $\gamma, \lambda, 0$ )	Be Poi OBe	$N \begin{cases} e^{-KW_0(\{3\gamma\lambda\}^{-1})} & d = 0 \\ K^{-d(1-d)} & d > 0 \text{ (}\Gamma P\text{)} \\ K^{-1/d} & d > 0 \text{ (BP, BPP)} \end{cases}$	$3K$
DB	BP( $\gamma, \lambda \geq 1, 0$ )	Be	$N \left(\frac{\xi}{\xi+1}\right)^K$	$\left(\frac{3c}{\xi} + 1\right) K$
	$\Gamma P(\gamma, \lambda, 0)$	Poi		
	BPP( $\gamma, \lambda > 1, 0$ )	OBe		
	DP( $\gamma$ )	—	$\frac{N\gamma}{a(\gamma-a)} \left(\frac{\xi}{\xi+a}\right)^K, a \in (0, 1] \cap (0, \gamma)$	
SB	BP( $\gamma, \lambda, d$ )	Be	$NK^{d-1}$	$K$
	$\Gamma P(\gamma, \lambda, d)$	Poi		
	BPP( $\gamma, \lambda, d$ )	OBe		
	N $\Gamma P(\gamma, \lambda, d)$	—	$N \begin{cases} K^{-1} \log K & d = 0, \gamma\lambda = 1 \\ K^{-\min(1, \gamma\lambda)} & d = 0, \gamma\lambda \neq 1 \\ K^{d-1} & d > 0 \end{cases}$	
PL	BP( $\gamma, \lambda, d$ ) $\Gamma P(\gamma, \lambda, d)$ BPP( $\gamma, \lambda > 1, d$ )	Be Poi OBe	$N \begin{cases} \left(\frac{\lambda}{\lambda+1}\right)^K & d = 0 \text{ (BP, } \Gamma P\text{)} \\ 2^{-K} & d = 0 \text{ (BPP)} \\ K^{1-1/d} & d > 0 \end{cases}$	$\frac{\gamma}{2} K^2$

$\theta_{ki}$ ), and  $\mathbb{E}[C_k] = \eta_k$ , so  $\mathbb{E}[R] = K + 3 \sum_{k=1}^K \eta_k$ . Note that here  $\mathbb{E}[R] \sim K$ , for  $K \rightarrow \infty$  since  $\eta_k$  is a decreasing sequence. For the power-law representation,  $r_k = k + 2$ , since each atom requires generating  $\psi_{ki}$ ,  $V_{ki}$ , and  $k$  beta random variables, and therefore  $\mathbb{E}[R] = \left(1 + \frac{5\gamma}{2}\right) K + \frac{\gamma}{2} K^2$ .

## 7. Summary of results

Table 1 summarizes our truncation and simulation cost results as applied to the beta, (normalized) gamma, and beta prime processes. Results for the Bondesson representation of BP( $\gamma, 1, 0$ ) as well as the decoupled Bondesson representations of BP( $\gamma, \lambda, 0$ ) and  $\Gamma P(\gamma, \lambda, 0)$  were previously known, and are reproduced by our results. All other results in the table are novel to the best of the authors' knowledge. It is interesting to note that the bounds and expected costs within each of the representation classes often have the same form, aside from some constants. Across classes, however, they vary significantly, indicating that the chosen sequential representation of a process has more of an influence on the truncation error than the process itself.

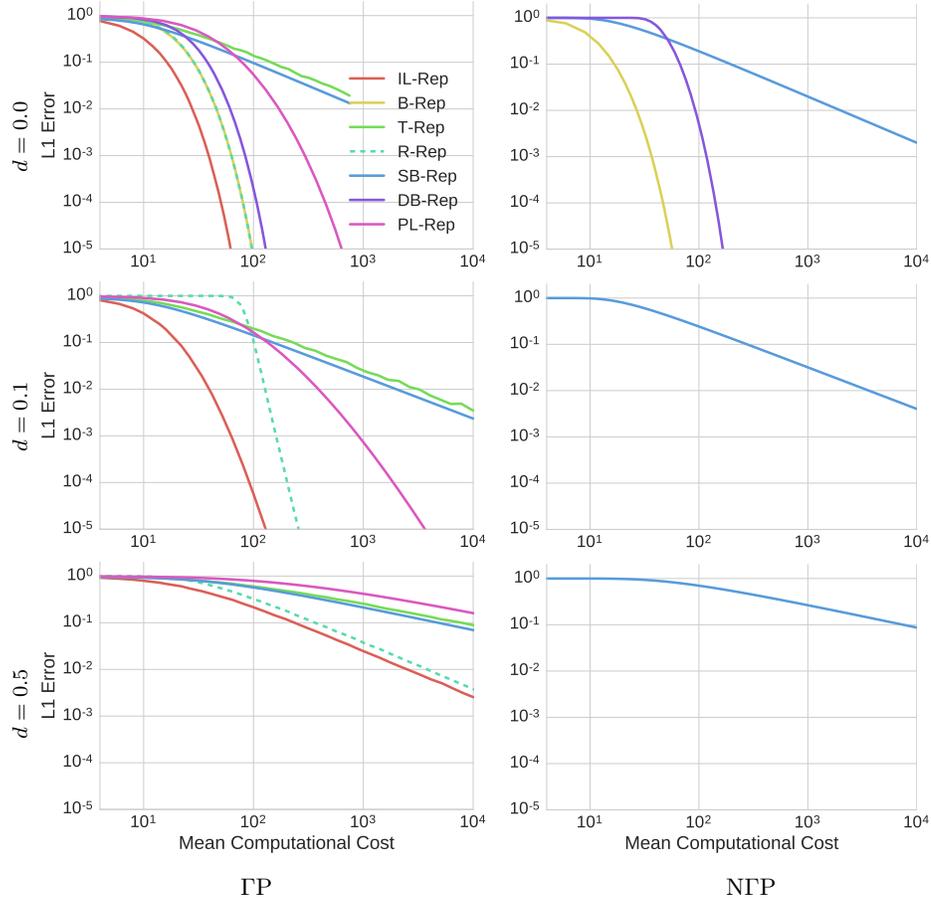


Figure 1: Truncation error bounds for representations of the (normalized) gamma-Poisson process, with  $\gamma = 1$ ,  $\lambda = 2$ , and  $\xi = \gamma\lambda$ . The left column is for the unnormalized process, while the right column is for the normalized process. Each row displays results for a different value of the discount parameter  $d \in \{0, 0.1, 0.5\}$ .

Fig. 1 shows a comparison of how the truncation error bounds vary with the expected computational cost  $\mathbb{E}[R]$  of simulation for the (normalized) gamma process and Poisson likelihood with  $N = 5$  observations. Results shown for the thinning, rejection, and inverse-Lévy representations are computed by Monte-Carlo approximation of the formula for  $B_{N,K}$  in Eq. (4.3), while all others use closed-form expressions from the examples in Sections 4 and 5. Note that the Bondesson and decoupled Bondesson representations do not exist when  $d > 0$ . Further, only those representations for which we provide closed-form bounds in the examples are shown for the normalized gamma process; we leave the numerical approximation of the results from Theorems 5.4 and 5.5 as an open problem. Similar figures for other processes (in particular, the beta-Bernoulli and beta prime-odds Bernoulli) are provided in Appendix A. Note that all bounds presented are improved by a factor of two versus comparable past results in the literature, due to the reliance on Lemmas 4.1 and 5.1 rather than the earlier bound found in Ishwaran and James (2001).

In Fig. 1, the top row shows results for the light-tailed process ( $\gamma = 1$ ,  $\lambda = 1$ ,  $d = 0$ , and  $\xi = c = \gamma\lambda$ ). All representations except for thinning and size-biased capture its exponential truncation error decay. This is due to the fact that the thinning representation generates increasingly many atoms of weight 0 as  $K \rightarrow \infty$ , and the expected number of atoms at each outer index for the size-biased representation decays as  $K \rightarrow \infty$ . The inverse-Lévy representation has the lowest truncation error as expected, as it is the only representation that generates a nonincreasing sequence of weights (and so must be the most efficient (Arbel and Prünster, 2017)). Based on this figure and those in Appendix A for other processes, it appears that the Bondesson representation typically provides the best tradeoff between simplicity and efficiency, and should be used whenever its conditions in Theorem 3.1 are satisfied. When the technical conditions are not satisfied, the rejection representation is a good alternative. If ease of theoretical analysis is a concern, the decoupled Bondesson representation provides comparable efficiency with the analytical simplicity of a superposition representation.

The bottom two rows of Fig. 1 show results for the heavy-tailed process ( $\gamma = 1$ ,  $\lambda = 2$ , and  $d \in \{0.1, 0.5\}$ ). The representation options are more limited, as the technical conditions of the Bondesson and decoupled Bondesson representations are not satisfied. Here the rejection representation is often the best choice due to its simplicity and competitive performance with the inverse-Lévy representation. However, one must take care to check its efficiency beforehand using Proposition 3.2 given a particular choice of  $\mu(d\theta)$ . For example, the choice of  $\mu(d\theta) \propto \theta^{-1-d}d\theta$  in the present work makes the rejection representation very inefficient when  $d \ll 1$  for both the gamma-Poisson (Fig. 1) and beta prime-odds Bernoulli (Fig. A.2) processes, but efficient for the beta-Bernoulli process (Fig. A.1). If no  $\mu(d\theta)$  yields reasonable results, the power-law representation is a good choice for  $d \ll 1$  as its truncation bound approaches the exponential decay of the light-tailed process. For larger  $d > 0$  the size-biased representation is a good alternative.

Based on the results in Fig. 1, it appears that there is no single dominant representation for all situations (provided the inverse-Lévy representation is intractable, as it most often is). However, as a guideline, the rejection and Bondesson representations tend to be good choices for light-tailed processes, while the rejection, size-biased, and power-law representations are good choices for heavy-tailed processes.

## 8. Discussion

We have investigated sequential representations, truncation error bounds, and simulation algorithms for (normalized) completely random measures. In past work, the development and analysis of these tools has occurred only on an ad hoc basis. The results in the present paper, in contrast, provide a comprehensive characterization and analysis of the different types of sequential (N)CRM representations available to the practitioner. However, there are a number of remaining open questions and limitations.

First, this work does not consider the influence of observed data: all analyses assume an *a priori* perspective, as truncation is typically performed before data are incorporated via posterior inference (e.g. in variational inference for the DP mixture (Blei and Jordan, 2006) and BP latent feature model (Doshi-Velez et al., 2009)). However, analysis of *a posteriori* truncation has been studied in past work as well (Ishwaran and James, 2001; Gelfand and Kottas, 2002; Ishwaran and Zarepour, 2002). In the language of CRMs, observations introduce a fixed-location component in the posterior process, while the unobserved traits are drawn from the (possibly normalized) ordinary component of a CRM (Ishwaran and Zarepour, 2002; Broderick et al., 2017). We anticipate that this property makes observations reasonably simple to include: the truncation tools provided in the present paper can be used directly on the unobserved ordinary component, while the fixed-location component may be treated exactly.

In addition, there are important open questions regarding the sequential representations developed in this work. It is unknown whether generalized versions of the Bondesson and decoupled Bondesson representations can be developed for larger classes of rate measures. The power-law representation does provide a partial answer in the decoupled Bondesson case. Regarding size-biased representations, one might expect that the use of conjugate exponential family CRMs (Broderick et al., 2017) would yield a closed-form expression for the truncation bound. In all of the cases provided in this paper, this was indeed the case; the integrals were evaluated exactly and a closed-form expression was found. However, we were unable to identify a general expression applicable to all conjugate exponential family CRMs. Based on the examples provided, we conjecture that such an expression exists. Finally, fundamental connections between some of the representations were left largely unexplored in this work. This is an open area of research, although progress has been made by connecting decoupled Bondesson and size-biased representations for (hierarchies of) generalized beta processes (Roy, 2014, Sec. 6.4).

A final remark is that one of the primary uses of sequential representations in past work has been in the development of posterior inference procedures (Paisley et al., 2010; Blei and Jordan, 2006; Doshi-Velez et al., 2009). The present work provides no guidance on which truncated representations are best paired with which inference methods. We leave this as an open direction for future research, which will require both theoretical and empirical investigation.

## Acknowledgments

The authors thank the anonymous reviewers for their thoughtful comments, which led to substantial improvements in both the presentation and content of the paper. All authors are supported by the Office of Naval Research under MURI grant N000141110688. J. Huggins is supported by the U.S. Government under FA9550-11-C-0028 and awarded by the DoD, Air Force Office of Scientific Research, National Defense Science and Engineering Graduate (NDSEG) Fellowship, 32 CFR 168a. T. Campbell and T. Broderick are supported by DARPA award FA8750-17-2-0019.

## References

- Abramowitz, M. and Stegun, I. (eds.) (1964). *Handbook of Mathematical Functions*. Dover Publications.
- Airoldi, E. M., Blei, D., Erosheva, E. A., and Fienberg, S. E. (2014). *Handbook of Mixed Membership Models and Their Applications*. CRC Press.
- Arbel, J. and Prünster, I. (2017). “A moment-matching Ferguson & Klass algorithm.” *Statistics and Computing*, 27(1): 3–17.
- Argiento, R., Bianchini, I., and Guglielmi, A. (2016). “A blocked Gibbs sampler for NGG-mixture models via a priori truncation.” *Statistics and Computing*, 26(3): 641–661.
- Banjevic, D., Ishwaran, H., and Zarepour, M. (2002). “A recursive method for functionals of Poisson processes.” *Bernoulli*, 8(3): 295–311.
- Blei, D. M. and Jordan, M. I. (2006). “Variational inference for Dirichlet process mixtures.” *Bayesian Analysis*, 1(1): 121–144.
- Bondesson, L. (1982). “On simulation from infinitely divisible distributions.” *Advances in Applied Probability*, 14: 855–869.
- Brix, A. (1999). “Generalized gamma measures and shot-noise Cox processes.” *Advances in Applied Probability*, 31: 929–953.
- Broderick, T., Jordan, M. I., and Pitman, J. (2012). “Beta processes, stick-breaking and power laws.” *Bayesian Analysis*, 7(2): 439–476.
- Broderick, T., Mackey, L., Paisley, J., and Jordan, M. I. (2015). “Combinatorial clustering and the beta negative binomial process.” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(2): 290–306.
- Broderick, T., Wilson, A. C., and Jordan, M. I. (2017). “Posteriors, conjugacy, and exponential families for completely random measures.” *Bernoulli*.
- Doshi-Velez, F., Miller, K. T., Van Gael, J., and Teh, Y. W. (2009). “Variational inference for the Indian buffet process.” In *International Conference on Artificial Intelligence and Statistics*.
- Ferguson, T. S. (1973). “A Bayesian analysis of some nonparametric problems.” *The Annals of Statistics*, 1(2): 209–230.
- Ferguson, T. S. and Klass, M. J. (1972). “A representation of independent increment

- processes without Gaussian components.” *The Annals of Mathematical Statistics*, 43(5): 1634–1643.
- Gelfand, A. and Kottas, A. (2002). “A computational approach for nonparametric Bayesian inference under Dirichlet process mixture models.” *Journal of Computational and Graphical Statistics*, 11(2): 289–305.
- Gumbel, E. J. (1954). *Statistical theory of extreme values and some practical applications: A series of lectures*. National Bureau of Standards Applied Mathematics Series 33. Washington, D.C.: U.S. Government Printing Office.
- Hjort, N. L. (1990). “Nonparametric Bayes estimators based on beta processes in models for life history data.” *The Annals of Statistics*, 18(3): 1259–1294.
- Ishwaran, H. and James, L. F. (2001). “Gibbs sampling methods for stick-breaking priors.” *Journal of the American Statistical Association*, 96(453): 161–173.
- (2002). “Approximate Dirichlet Process Computing in Finite Normal Mixtures: Smoothing and Prior Information.” *Journal of Computational and Graphical Statistics*, 11(3): 508–532.
- Ishwaran, H. and Zarepour, M. (2002). “Exact and approximate sum representations for the Dirichlet process.” *Canadian Journal of Statistics*, 30(2): 269–283.
- James, L. F. (2002). “Poisson Process Partition Calculus with applications to Exchangeable models and Bayesian Nonparametrics.” *arXiv preprint arXiv:0205093*.
- (2013). “Stick-breaking  $PG(\alpha, \zeta)$ -Generalized Gamma Processes.” *arXiv preprint arXiv:1308.6570*.
- (2014). “Poisson Latent Feature Calculus for Generalized Indian Buffet Processes.” *arXiv preprint arXiv:1411.2936v3*.
- James, L. F., Lijoi, A., and Prünster, I. (2009). “Posterior Analysis for Normalized Random Measures with Independent Increments.” *Scandinavian Journal of Statistics*, 36(1): 76–97.
- Kim, Y. (1999). “Nonparametric Bayesian estimators for counting processes.” *The Annals of Statistics*, 27(2): 562–588.
- Kingman, J. F. C. (1967). “Completely random measures.” *Pacific Journal of Mathematics*, 21(1): 59–78.
- (1975). “Random discrete distributions.” *Journal of the Royal Statistical Society B*, 37(1): 1–22.
- Kingman, J. F. C. (1993). *Poisson Processes*. Oxford Studies in Probability. Oxford University Press.
- Lijoi, A., Mena, R., and Prünster, I. (2005). “Bayesian nonparametric analysis for a generalized Dirichlet process prior.” *Statistical Inference for Stochastic Processes*, 8: 283–309.
- (2007). “Controlling the reinforcement in Bayesian nonparametric mixture models.” *Journal of the Royal Statistical Society B*, 69(4): 715–740.
- Lijoi, A. and Prünster, I. (2003). “On a normalized random measure with independent increments relevant to Bayesian nonparametric inference.” In *Proceedings of the 13th European Young Statisticians Meeting*, 123–124. Bernoulli Society.
- (2010). “Models beyond the Dirichlet process.” In Hjort, N. L., Holmes, C., Müller, P., and Walker, S. (eds.), *Bayesian Nonparametrics*, 80–136. Cambridge University

- Press.
- Maddison, C., Tarlow, D., and Minka, T. P. (2014). “A\* Sampling.” In *Advances in Neural Information Processing Systems*.
- Muliere, P. and Tardella, L. (1998). “Approximating distributions of random functionals of Ferguson–Dirichlet priors.” *Canadian Journal of Statistics*, 26(2): 283–297.
- Orbanz, P. (2010). “Conjugate projective limits.” *arXiv preprint arXiv:1012.0363*.
- Paisley, J. W., Blei, D. M., and Jordan, M. I. (2012). “Stick-breaking beta processes and the Poisson process.” *International Conference on Artificial Intelligence and Statistics*.
- Paisley, J. W., Zaas, A. K., Woods, C. W., Ginsburg, G. S., and Carin, L. (2010). “A stick-breaking construction of the beta process.” *International Conference on Machine Learning*.
- Perman, M. (1993). “Order statistics for jumps of normalised subordinators.” *Stochastic Processes and their Applications*, 46(2): 267–281.
- Perman, M., Pitman, J., and Yor, M. (1992). “Size-biased sampling of Poisson point processes and excursions.” *Probability Theory and Related Fields*, 92(1): 21–39.
- Pitman, J. (2003). “Poisson-Kingman partitions.” *Lecture Notes-Monograph Series*.
- Regazzini, E., Lijoi, A., and Prünster, I. (2003). “Distributional results for means of normalized random measures with independent increments.” *The Annals of Statistics*, 31(2): 560–585.
- Rosiński, J. (1990). “On series representations of infinitely divisible random vectors.” *Annals of Probability*, 18: 405–430.
- (2001). “Series representations of Lévy processes from the perspective of point processes.” In Barndorff-Nielsen, O., Resnick, S., and Mikosch, T. (eds.), *Lévy processes: theory and applications*, chapter VI, 401–415. Birkhäuser Boston.
- Roy, D. (2014). “The continuum-of-urns scheme, generalized beta and Indian buffer processes, and hierarchies thereof.” *arXiv preprint arXiv:1501.00208*.
- Roychowdhury, A. and Kulis, B. (2015). “Gamma Processes, Stick-Breaking, and Variational Inference.” In *International Conference on Artificial Intelligence and Statistics*.
- Sethuraman, J. (1994). “A constructive definition of Dirichlet priors.” *Statistica Sinica*, 4: 639–650.
- Teh, Y. W. and Görür, D. (2009). “Indian buffet processes with power-law behavior.” In *Advances in Neural Information Processing Systems*.
- Thibaux, R. and Jordan, M. I. (2007). “Hierarchical beta processes and the Indian buffet process.” In *International Conference on Artificial Intelligence and Statistics*.
- Titsias, M. (2008). “The infinite gamma-Poisson feature model.” In *Advances in Neural Information Processing Systems*.
- Zhou, M., Hannah, L., Dunson, D., and Carin, L. (2012). “Beta-negative binomial process and Poisson factor analysis.” In *Artificial Intelligence and Statistics*.