

Negative Binomial Process Count and Mixture Modeling

Mingyuan Zhou and Lawrence Carin, *Fellow, IEEE*

Abstract—The seemingly disjoint problems of count and mixture modeling are united under the negative binomial (NB) process. A gamma process is employed to model the rate measure of a Poisson process, whose normalization provides a random probability measure for mixture modeling and whose marginalization leads to an NB process for count modeling. A draw from the NB process consists of a Poisson distributed finite number of distinct atoms, each of which is associated with a logarithmic distributed number of data samples. We reveal relationships between various count- and mixture-modeling distributions and construct a Poisson-logarithmic bivariate distribution that connects the NB and Chinese restaurant table distributions. Fundamental properties of the models are developed, and we derive efficient Bayesian inference. It is shown that with augmentation and normalization, the NB process and gamma-NB process can be reduced to the Dirichlet process and hierarchical Dirichlet process, respectively. These relationships highlight theoretical, structural and computational advantages of the NB process. A variety of NB processes, including the beta-geometric, beta-NB, marked-beta-NB, marked-gamma-NB and zero-inflated-NB processes, with distinct sharing mechanisms, are also constructed. These models are applied to topic modeling, with connections made to existing algorithms under Poisson factor analysis. Example results show the importance of inferring both the NB dispersion and probability parameters.

Index Terms—Beta process, Chinese restaurant process, completely random measures, count modeling, Dirichlet process, gamma process, hierarchical Dirichlet process, mixed-membership modeling, mixture modeling, negative binomial process, normalized random measures, Poisson factor analysis, Poisson process, topic modeling.



1 INTRODUCTION

COUNT data appear in many settings, such as predicting the number of motor insurance claims [1], [2], analyzing infectious diseases [3] and modeling topics of document corpora [4]–[8]. There has been increasing interest in count modeling using the Poisson process, geometric process [9]–[13] and recently the negative binomial (NB) process [8], [14], [15]. It is shown in [8] and further demonstrated in [15] that the NB process, originally constructed for count analysis, can be naturally applied for mixture modeling of *grouped* data x_1, \dots, x_J , where each group $x_j = \{x_{ji}\}_{i=1, N_j}$. For example, in topic modeling (mixed-membership modeling), each document consists of a group of exchangeable words and each word is a group member that is assigned to a topic; the number of times a topic appears in a document is a latent count random variable that could be well modeled with an NB distribution [8], [15].

Mixture modeling, which infers random probability measures to assign data samples into clusters (mixture components), is a key research area of statistics and machine learning. Although the number of samples assigned to clusters are counts, mixture modeling is not typically considered as a count-modeling problem. It is often addressed under the Dirichlet-multinomial frame-

work, using the Dirichlet process [16]–[21] as the prior distribution. With the Dirichlet-multinomial conjugacy, the Dirichlet process mixture model enjoys tractability because the posterior of the random probability measure is still a Dirichlet process. Despite its popularity, the Dirichlet process is inflexible in that a single concentration parameter controls both the variability of the mass around the mean [21], [22] and the distribution of the number of distinct atoms [18], [23]. For mixture modeling of grouped data, the hierarchical Dirichlet process (HDP) [24] has been further proposed to share statistical strength between groups. The HDP inherits the same inflexibility of the Dirichlet process, and due to the non-conjugacy between Dirichlet processes, its inference has to be solved under alternative constructions, such as the Chinese restaurant franchise and stick-breaking representations [24]–[26]. To make the number of distinct atoms increase at a rate faster than that of the Dirichlet process, one may consider the Pitman-Yor process [27], [28] or the normalized generalized gamma process [23] that provide extra parameters to add flexibility.

To construct more expressive mixture models with tractable inference, in this paper we consider mixture modeling as a count-modeling problem. Directly modeling the counts assigned to mixture components as NB random variables, we perform joint count and mixture modeling via the NB process, using completely random measures [9], [22], [29], [30] that are easy to construct and amenable to posterior computation. By constructing a bivariate count distribution that connects the Poisson, logarithmic, NB and Chinese restaurant table distribu-

• M. Zhou is with the Department of Information, Risk, and Operations Management, McCombs School of Business, University of Texas at Austin, Austin, TX 78712. L. Carin is with the Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708.

tions, we develop data augmentation and marginalization techniques unique to the NB distribution, with which we augment an NB process into both the gamma-Poisson and compound Poisson representations, yielding unification of count and mixture modeling, derivation of fundamental model properties, as well as efficient Bayesian inference.

Under the NB process, we employ a gamma process to model the rate measure of a Poisson process. The normalization of the gamma process provides a random probability measure (not necessarily a Dirichlet process) for mixture modeling, and the marginalization of the gamma process leads to an NB process for count modeling. Since the gamma scale parameters appear as NB probability parameters when the gamma processes are marginalized out, they directly control count distributions on atoms and they could be conveniently inferred with the beta-NB conjugacy. For mixture modeling of grouped data, we construct hierarchical models by employing an NB process for each group and sharing their NB dispersion or probability measures across groups. Different parameterizations of the NB dispersion and probability parameters result in a wide variety of NB processes, which are connected to previously proposed nonparametric Bayesian mixture models. The proposed joint count and mixture modeling framework provides new opportunities for better data fitting, efficient inference and flexible model constructions.

1.1 Related Work

Parts of the work presented here first appeared in [2], [8], [15]. In this paper, we unify related materials scattered in these three conference papers and provide significant expansions. In particular, we construct a Poisson-logarithmic bivariate distribution that tightly connects the NB and Chinese restaurant table distributions, extending the Chinese restaurant process to describe the case that both the numbers of customers and tables are random variables, and we provide necessary conditions to recover the NB process and the gamma-NB process from the Dirichlet process and HDP, respectively.

We mention that a related beta-NB process has been independently investigated in [14]. Our constructions of a wide variety of NB processes, including the beta-NB processes in [8] and [14] as special cases, are built on our thorough investigation of the properties, relationships and inference of the NB and related stochastic processes. In particular, we show that the gamma-Poisson construction of the NB process is key to uniting count and mixture modeling, and there are two equivalent augmentations of the NB process that allow us to develop analytic conditional posteriors and predictive distributions. These insights are not provided in [14], and the NB dispersion parameters there are empirically set rather than inferred. More distinctions will be discussed along with specific models.

The remainder of the paper is organized as follows. We review some commonly used nonparametric Bayesian

priors in Section 2 and study the NB distribution in Section 3. We present the NB process in Section 4, the gamma-NB process in Section 5, and the NB process family in Section 6. We discuss NB process topic modeling in Section 7 and present example results in Section 8.

2 PRELIMINARIES

2.1 Completely Random Measures

Following [22], for any $\nu^+ \geq 0$ and any probability distribution $\pi(dp d\omega)$ on the product space $\mathbb{R} \times \Omega$, let $K^+ \sim \text{Pois}(\nu^+)$ and $(p_k, \omega_k) \stackrel{iid}{\sim} \pi(dp d\omega)$ for $k = 1, \dots, K^+$. Defining $\mathbf{1}_A(\omega_k)$ as being one if $\omega_k \in A$ and zero otherwise, the random measure $\mathcal{L}(A) \equiv \sum_{k=1}^{K^+} \mathbf{1}_A(\omega_k) p_k$ assigns independent infinitely divisible random variables $\mathcal{L}(A_i)$ to disjoint Borel sets $A_i \subset \Omega$, with characteristic functions

$$\mathbb{E}[e^{it\mathcal{L}(A)}] = \exp\left\{\int \int_{\mathbb{R} \times A} (e^{itp} - 1) \nu(dp d\omega)\right\}, \quad (1)$$

where $\nu(dp d\omega) \equiv \nu^+ \pi(dp d\omega)$. A random signed measure \mathcal{L} satisfying (1) is called a Lévy random measure. More generally, if the Lévy measure $\nu(dp d\omega)$ satisfies

$$\int \int_{\mathbb{R} \times S} \min\{1, |p|\} \nu(dp d\omega) < \infty \quad (2)$$

for each compact $S \subset \Omega$, the Lévy random measure \mathcal{L} is well defined, even if the Poisson intensity ν^+ is infinite. A nonnegative Lévy random measure \mathcal{L} satisfying (2) was called a completely random measure [9], [29], and it was introduced to machine learning in [31] and [30].

2.1.1 Poisson Process

Define a Poisson process $X \sim \text{PP}(G_0)$ on the product space $\mathbb{Z}_+ \times \Omega$, where $\mathbb{Z}_+ = \{0, 1, \dots\}$, with a finite and continuous base measure G_0 over Ω , such that $X(A) \sim \text{Pois}(G_0(A))$ for each subset $A \subset \Omega$. The Lévy measure of the Poisson process can be derived from (1) as $\nu(dud\omega) = \delta_1(du)G_0(d\omega)$, where $\delta_1(du)$ is a unit point mass at $u = 1$. If G_0 is discrete (atomic) as $G_0 = \sum_k \lambda_k \delta_{\omega_k}$, then the Poisson process definition is still valid that $X = \sum_k x_k \delta_{\omega_k}$, $x_k \sim \text{Pois}(\lambda_k)$. If G_0 is mixed discrete-continuous, then X is the sum of two independent contributions. Except where otherwise specified, below we consider the base measure to be finite and continuous.

2.1.2 Gamma Process

We define a gamma process [10], [22] $G \sim \text{GaP}(c, G_0)$ on the product space $\mathbb{R}_+ \times \Omega$, where $\mathbb{R}_+ = \{x : x \geq 0\}$, with scale parameter $1/c$ and base measure G_0 , such that $G(A) \sim \text{Gamma}(G_0(A), 1/c)$ for each subset $A \subset \Omega$, where $\text{Gamma}(\lambda; a, b) = \frac{1}{\Gamma(a)b^a} \lambda^{a-1} e^{-\lambda/b}$. The gamma process is a completely random measure, whose Lévy measure can be derived from (1) as $\nu(dr d\omega) = r^{-1} e^{-cr} dr G_0(d\omega)$. Since the Poisson intensity $\nu^+ = \nu(\mathbb{R}_+ \times \Omega) = \infty$ and $\int \int_{\mathbb{R}_+ \times \Omega} r \nu(dr d\omega)$ is finite, there are countably infinite atoms and a draw from the gamma process can be expressed as

$$G = \sum_{k=1}^{\infty} r_k \delta_{\omega_k}, \quad (r_k, \omega_k) \stackrel{iid}{\sim} \pi(dr d\omega),$$

where $\pi(dr d\omega)\nu^+ \equiv \nu(dr d\omega)$.

2.1.3 Beta Process

The beta process was defined by [32] for survival analysis with $\Omega = \mathbb{R}_+$. Thibaux and Jordan [31] modified the process by defining $B \sim \text{BP}(c, B_0)$ on the product space $[0, 1] \times \Omega$, with Lévy measure $\nu(dp d\omega) = cp^{-1}(1-p)^{c-1} dp B_0(d\omega)$, where $c > 0$ is a concentration parameter and B_0 is a base measure. Since the Poisson intensity $\nu^+ = \nu([0, 1] \times \Omega) = \infty$ and $\int \int_{[0, 1] \times \Omega} p\nu(dp d\omega)$ is finite, there are countably infinite atoms and a draw from the beta process can be expressed as

$$B = \sum_{k=1}^{\infty} p_k \delta_{\omega_k}, \quad (p_k, \omega_k) \stackrel{iid}{\sim} \pi(dp d\omega),$$

where $\pi(dp d\omega)\nu^+ \equiv \nu(dp d\omega)$.

2.2 Dirichlet and Chinese Restaurant Processes

2.2.1 Dirichlet Process

Denote $\tilde{G} = G/G(\Omega)$, where $G \sim \text{GaP}(c, G_0)$, then for any measurable disjoint partition A_1, \dots, A_Q of Ω , we have $[\tilde{G}(A_1), \dots, \tilde{G}(A_Q)] \sim \text{Dir}(\gamma_0 \tilde{G}_0(A_1), \dots, \gamma_0 \tilde{G}_0(A_Q))$, where $\gamma_0 = G_0(\Omega)$ and $\tilde{G}_0 = G_0/\gamma_0$. Therefore, with a space invariant scale parameter $1/c$, the normalized gamma process $\tilde{G} = G/G(\Omega)$ is a Dirichlet process [16], [33] with concentration parameter γ_0 and base probability measure \tilde{G}_0 , expressed as $\tilde{G} \sim \text{DP}(\gamma_0, \tilde{G}_0)$. Unlike the gamma process, the Dirichlet process is no longer a completely random measure as the random variables $\{\tilde{G}(A_q)\}$ for disjoint sets $\{A_q\}$ are negatively correlated.

A gamma process with a space invariant scale parameter can also be recovered from a Dirichlet process: if a gamma random variable $\alpha \sim \text{Gamma}(\gamma_0, 1/c)$ and a Dirichlet process $\tilde{G} \sim \text{DP}(\gamma_0, \tilde{G}_0)$ are independent with $\gamma_0 = G_0(\Omega)$ and $\tilde{G}_0 = G_0/\gamma_0$, then $G = \alpha\tilde{G}$ becomes a gamma process as $G \sim \text{GaP}(c, G_0)$.

2.2.2 Chinese Restaurant Process

In a Dirichlet process $\tilde{G} \sim \text{DP}(\gamma_0, \tilde{G}_0)$, we assume $X_i \sim \tilde{G}$; $\{X_i\}$ are independent given \tilde{G} and hence exchangeable. The predictive distribution of a new data sample X_{m+1} , conditioning on X_1, \dots, X_m , with \tilde{G} marginalized out, can be expressed as

$$X_{m+1}|X_1, \dots, X_m \sim \mathbb{E} \left[\tilde{G} \mid X_1, \dots, X_m \right] \\ = \sum_{k=1}^K \frac{n_k}{m+\gamma_0} \delta_{\omega_k} + \frac{\gamma_0}{m+\gamma_0} \tilde{G}_0, \quad (3)$$

where $\{\omega_k\}_{1,K}$ are distinct atoms in Ω observed in X_1, \dots, X_m and $n_k = \sum_{i=1}^m X_i(\omega_k)$ is the number of data samples associated with ω_k . The stochastic process described in (3) is known as the Pólya urn scheme [34] and also the Chinese restaurant process [24], [35], [36].

The number of nonempty tables l in a Chinese restaurant process, with concentration parameter γ_0 and m

customers, is a random variable generated as $l = \sum_{n=1}^m b_n$, $b_n \sim \text{Bernoulli}\left(\frac{\gamma_0}{n-1+\gamma_0}\right)$. This random variable is referred as the Chinese restaurant table (CRT) random variable $l \sim \text{CRT}(m, \gamma_0)$. As shown in [15], [17], [18], [24], it has probability mass function (PMF)

$$f_L(l|m, \gamma_0) = \frac{\Gamma(\gamma_0)}{\Gamma(m+\gamma_0)} |s(m, l)| \gamma_0^l, \quad l = 0, 1, \dots, m$$

where $s(m, l)$ are Stirling numbers of the first kind.

3 NEGATIVE BINOMIAL DISTRIBUTION

The Poisson distribution $m \sim \text{Pois}(\lambda)$ is commonly used to model count data, with PMF

$$f_M(m) = \frac{\lambda^m e^{-\lambda}}{m!}, \quad m \in \mathbb{Z}_+.$$

Its mean and variance are both equal to λ . Due to heterogeneity (difference between individuals) and contagion (dependence between the occurrence of events), count data are usually overdispersed in that the variance is greater than the mean, making the Poisson assumption restrictive. By placing a gamma prior with shape r and scale $\frac{p}{1-p}$ on λ as $m \sim \text{Pois}(\lambda)$, $\lambda \sim \text{Gamma}(r, \frac{p}{1-p})$ and marginalizing out λ , an NB distribution $m \sim \text{NB}(r, p)$ is obtained, with PMF

$$f_M(m|r, p) = \frac{\Gamma(r+m)}{m! \Gamma(r)} (1-p)^r p^m, \quad m \in \mathbb{Z}_+,$$

where r is the nonnegative dispersion parameter and p is the probability parameter. Thus the NB distribution is also known as the gamma-Poisson mixture distribution [37]. It has a mean $\mu = rp/(1-p)$ smaller than the variance $\sigma^2 = rp/(1-p)^2 = \mu + r^{-1}\mu^2$, with the variance-to-mean ratio (VMR) as $(1-p)^{-1}$ and the overdispersion level (ODL, the coefficient of the quadratic term in σ^2) as r^{-1} , and thus it is usually favored over the Poisson distribution for modeling overdispersed counts.

As shown in [38], $m \sim \text{NB}(r, p)$ can also be generated from a compound Poisson distribution as

$$m = \sum_{t=1}^l u_t, \quad u_t \stackrel{iid}{\sim} \text{Log}(p), \quad l \sim \text{Pois}(-r \ln(1-p)),$$

where $u \sim \text{Log}(p)$ corresponds to the logarithmic distribution [38], [39] with PMF $f_U(k) = -p^k/[k \ln(1-p)]$, $k = 1, 2, \dots$, and probability-generating function (PGF)

$$C_U(z) = \ln(1-pz)/\ln(1-p), \quad |z| < p^{-1}.$$

One may also show that $\lim_{r \rightarrow \infty} \text{NB}(r, \frac{\lambda}{\lambda+r}) = \text{Pois}(\lambda)$, and conditioning on $m > 0$, $m \sim \text{NB}(r, p)$ becomes $m \sim \text{Log}(p)$ as $r \rightarrow 0$.

The NB distribution has been widely investigated and applied to numerous scientific studies [40]–[43]. Although inference of the NB probability parameter p is straightforward with the beta-NB conjugacy, inference of the NB dispersion parameter r , whose conjugate prior is unknown, has long been a challenge. The maximum likelihood (ML) approach is commonly used to estimate r , however, it only provides a point estimate and does not allow incorporating prior information; moreover, the ML estimator of r often lacks robustness and may be severely

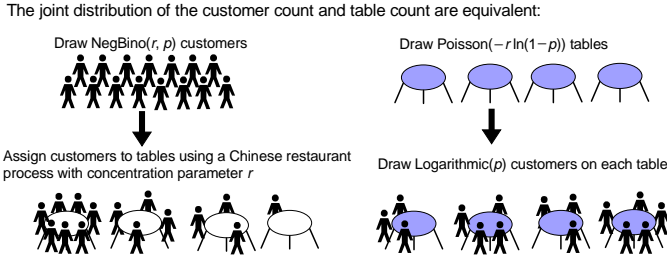


Fig. 1. The Poisson-logarithmic bivariate distribution models the total numbers of customers and tables as random variables. As shown in Theorem 1, it has two equivalent representations, which connect the Poisson, logarithmic, and negative binomial distributions and the Chinese restaurant process.

biased or even fail to converge, especially if the sample size is small [3], [44]–[48]. Bayesian approaches are able to model the uncertainty of estimation and incorporate prior information, however, the only available closed-form Bayesian inference for r relies on approximating the ratio of two gamma functions [49].

3.1 Poisson-Logarithmic Bivariate Distribution

Advancing previous research on the NB distribution in [2], [15], we construct a Poisson-logarithmic bivariate distribution that assists Bayesian inference of the NB distribution and unites various count distributions.

Theorem 1 (Poisson-logarithmic). *The Poisson-logarithmic (PoisLog) bivariate distribution with PMF*

$$f_{M,L}(m, l | r, p) = \frac{|s(m, l)| r^l}{m!} (1-p)^r p^m, \quad (4)$$

where $m \in \mathbb{Z}_+$ and $l = 0, 1, \dots, m$, can be expressed as a Chinese restaurant table (CRT) and negative binomial (NB) joint distribution and also a sum-logarithmic and Poisson joint distribution as

$$\begin{aligned} \text{PoisLog}(m, l; r, p) &= \text{CRT}(l; m, r) \text{NB}(m; r, p) \\ &= \text{SumLog}(m; l, p) \text{Pois}(l; -r \ln(1-p)), \end{aligned}$$

where $\text{SumLog}(m; l, p)$ denotes the sum-logarithmic distribution generated as $m = \sum_{t=1}^l u_t$, $u_t \stackrel{iid}{\sim} \text{Log}(p)$.

The proof of Theorem 1 is provided in Appendix A. As shown in Fig. 1, this bivariate distribution intuitively describes the joint distribution of two count random variables m and l under two equivalent circumstances:

- 1) There are $m \sim \text{NB}(r, p)$ customers seated at $l \sim \text{CRT}(m, r)$ tables;
- 2) There are $l \sim \text{Pois}(-r \ln(1-p))$ tables, each of which with $u_t \sim \text{Log}(p)$ customers, with $m = \sum_{t=1}^l u_t$ customers in total.

In a slight abuse of notation, but for added conciseness, in the following discussion we use $m \sim \sum_{t=1}^l \text{Log}(p)$ to denote $m = \sum_{t=1}^l u_t$, $u_t \stackrel{iid}{\sim} \text{Log}(p)$.

Corollary 2. *Let $m \sim \text{NB}(r, p)$, $r \sim \text{Gamma}(r_1, 1/c_1)$ represent a gamma-NB mixture distribution. It can be augmented as $m \sim \sum_{t=1}^l \text{Log}(p)$, $l \sim \text{Pois}(-r \ln(1-p))$, $r \sim \text{Gamma}(r_1, 1/c_1)$. Marginalizing out r leads to*

$$m \sim \sum_{t=1}^l \text{Log}(p), \quad l \sim \text{NB}(r_1, p'), \quad p' := \frac{-\ln(1-p)}{c_1 - \ln(1-p)},$$

where the latent count $l \sim \text{NB}(r_1, p')$ can be augmented as

$$l \sim \sum_{t'=1}^{l'} \text{Log}(p'), \quad l' \sim \text{Pois}(-r_1 \ln(1-p')),$$

which, using Theorem 1, is equivalent in distribution to

$$l' \sim \text{CRT}(l, r_1), \quad l \sim \text{NB}(r_1, p').$$

The connections between various distributions shown in Theorem 1 and Corollary 2 are key ingredients of this paper, which not only allow us to derive efficient inference, but also, as shown below, let us examine the posteriors to understand fundamental properties of various NB processes, clearly revealing connections to previous nonparametric Bayesian mixture models, including those based on the Dirichlet process, HDP and beta-NB processes.

4 JOINT COUNT AND MIXTURE MODELING

In this Section, we first show the connections between the Poisson and multinomial processes, and then we place a gamma process prior on the Poisson rate measure for joint count and mixture modeling. This construction can be reduced to the Dirichlet process and its restrictions for modeling grouped data are further discussed.

4.1 Poisson and Multinomial Processes

Corollary 3. *Let $X \sim \text{PP}(G)$ be a Poisson process defined on a completely random measure G such that $X(A) \sim \text{Pois}(G(A))$ for each subset $A \subset \Omega$. Define $Y \sim \text{MP}(Y(\Omega), \frac{G}{G(\Omega)})$ as a multinomial process, with total count $Y(\Omega) \sim \text{Pois}(G(\Omega))$ and random probability measure $\frac{G}{G(\Omega)}$, such that $(Y(A_1), \dots, Y(A_Q)) \sim \text{Mult}(Y(\Omega); \frac{G(A_1)}{G(\Omega)}, \dots, \frac{G(A_Q)}{G(\Omega)})$ for any disjoint partition $\{A_q\}_{1,Q}$ of Ω . According to Lemma 4.1 of [8], $X(A)$ and $Y(A)$ would have the same Poisson distribution for each $A \subset \Omega$, thus X and Y are equivalent in distribution.*

Using Corollary 3, we illustrate how the seemingly distinct problems of count and mixture modeling can be united under the Poisson process. For each $A \subset \Omega$, denote $X_j(A)$ as a count random variable describing the number of observations in x_j that reside within A . Given grouped data x_1, \dots, x_J , for any measurable disjoint partition A_1, \dots, A_Q of Ω , we aim to jointly model count random variables $\{X_j(A_q)\}$. A natural choice would be to define a Poisson process

$$X_j \sim \text{PP}(G)$$

with a shared completely random measure G on Ω , such that $X_j(A) \sim \text{Pois}(G(A))$ for each $A \subset \Omega$ and $G(\Omega) =$

$\sum_{q=1}^Q G(A_q)$. Following Corollary 3, with $\tilde{G} = G/G(\Omega)$, letting $X_j \sim \text{PP}(G)$ is equivalent to letting

$$X_j \sim \text{MP}(X_j(\Omega), \tilde{G}), \quad X_j(\Omega) \sim \text{Pois}(G(\Omega)).$$

Thus the Poisson process provides not only a way to generate independent counts from each A_q , but also a mechanism for mixture modeling, which allocates the $X_j(\Omega)$ observations into any measurable disjoint partition $\{A_q\}_{1,Q}$ of Ω , conditioning on the normalized random measure \tilde{G} .

4.2 Gamma-Poisson Process and Negative Binomial Process

To complete the Poisson process, it is natural to place a gamma process prior on the Poisson rate measure G as

$$X_j \sim \text{PP}(G), j = 1, \dots, J; \quad G \sim \text{GaP}(J(1-p)/p, G_0). \quad (5)$$

For a distinct atom ω_k , we have $n_{jk} \sim \text{Pois}(r_k)$, where $n_{jk} = X_j(\omega_k)$ and $r_k = G(\omega_k)$. Marginalizing out G of the gamma-Poisson process leads to an NB process

$$X = \sum_{j=1}^J X_j \sim \text{NBP}(G_0, p)$$

in which $X(A) \sim \text{NB}(G_0(A), p)$ for each $A \subset \Omega$.

Since $\mathbb{E}[e^{iuX(A)}] = \exp\{G_0(A)(\ln(1-p) - \ln(1-pe^{iu}))\} = \exp\{G_0(A) \sum_{m=1}^{\infty} (e^{ium} - 1) \frac{p^m}{m}\}$, the Lévy measure of the NB process can be derived from (1) as

$$\nu(dnd\omega) = \sum_{m=1}^{\infty} \frac{p^m}{m} \delta_m(dn) G_0(d\omega).$$

With $\nu^+ = \nu(\mathbb{Z}_+ \times \Omega) = -\gamma_0 \ln(1-p)$, a draw from the NB process consists of a finite number of distinct atoms almost surely and the number of samples on each of them follows a logarithmic distribution, expressed as

$$X = \sum_{k=1}^{K^+} n_k \delta_{\omega_k}, \quad K^+ \sim \text{Pois}(-\gamma_0 \ln(1-p)), \\ (n_k, \omega_k) \stackrel{iid}{\sim} \text{Log}(n_k; p) g_0(\omega_k), \quad k = 1, \dots, K^+, \quad (6)$$

where $g_0(d\omega) := G_0(d\omega)/\gamma_0$. Thus the NB probability parameter p plays a critical role in count and mixture modeling as it directly controls the prior distributions of the number of distinct atoms $K^+ \sim \text{Pois}(-\gamma_0 \ln(1-p))$, the number of samples at each of these atoms $n_k \sim \text{Log}(p)$, and the total number of samples $X(\Omega) \sim \text{NB}(\gamma_0, p)$. However, its value would become irrelevant if one directly works with the normalization of G , as commonly used in conventional mixture modeling.

Define $L \sim \text{CRTP}(X, G_0)$ as a CRT process that

$$L(A) = \sum_{\omega \in A} L(\omega), \quad L(\omega) \sim \text{CRT}(X(\omega), G_0(\omega))$$

for each $A \subset \Omega$. Under the Chinese restaurant process metaphor, $X(A)$ and $L(A)$ represent the customer count and table count, respectively, observed in each $A \subset \Omega$. A direct generalization of Theorem 1 leads to:

Corollary 4. *The NB process $X \sim \text{NBP}(G_0, p)$ augmented under its compound Poisson representation as*

$$X \sim \sum_{t=1}^L \text{Log}(p), \quad L \sim \text{PP}(-G_0 \ln(1-p))$$

is equivalent in distribution to

$$L \sim \text{CRTP}(X, G_0), \quad X \sim \text{NBP}(G_0, p).$$

4.3 Posterior Analysis and Predictive Distribution

Imposing a gamma prior $\text{Gamma}(e_0, 1/f_0)$ on γ_0 and a beta prior $\text{Beta}(a_0, 1/b_0)$ on p , using conjugacy, we have all conditional posteriors in closed-form as

$$(G|X, p, G_0) \sim \text{GaP}(J/p, G_0 + X) \\ (p|X, G_0) \sim \text{Beta}(a_0 + X(\Omega), b_0 + \gamma_0) \\ (L|X, G_0) \sim \text{CRTP}(X, G_0) \\ (\gamma_0|L, p) \sim \text{Gamma}\left(e_0 + L(\Omega), \frac{1}{f_0 - \ln(1-p)}\right). \quad (7)$$

If the base measure G_0 is finite and continuous, then $G_0(\omega) \rightarrow 0$ and we have $L(\omega) \sim \text{CRT}(X(\omega), G_0(\omega)) = \delta(X(\omega) > 0)$ and thus $L(\Omega) = \sum_{\omega \in \Omega} \delta(X(\omega) > 0)$, i.e., the number of nonempty tables $L(\Omega)$ is equal to K^+ , the number of distinct atoms. The gamma-Poisson process is also well defined with a discrete base measure $G_0 = \sum_{k=1}^K \frac{\gamma_0}{K} \delta_{\omega_k}$, for which we have $L = \sum_{k=1}^K l_k \delta_{\omega_k}$, $l_k \sim \text{CRT}(X(\omega_k), \gamma_0/K)$ and hence it is possible that $l_k > 1$ if $X(\omega_k) > 1$, which means $L(\Omega) \geq K^+$. As the data $\{x_{ji}\}_i$ are exchangeable within group j , conditioning on $X^{-ji} = X \setminus X_{ji}$ and G_0 , with G marginalized out, we have

$$X_{ji}|G_0, X^{-ji} \sim \frac{\mathbb{E}[G|G_0, p, X^{-ji}]}{\mathbb{E}[G(\Omega)|G_0, p, X^{-ji}]} \\ = \frac{G_0}{\gamma_0 + X(\Omega) - 1} + \frac{X^{-ji}}{\gamma_0 + X(\Omega) - 1}. \quad (8)$$

This prediction rule is similar to that of the Chinese restaurant process described in (3).

4.4 Relationship with the Dirichlet Process

Based on Corollary 3 on the multinomial process and Section 2.2.1 on the Dirichlet process, the gamma-Poisson process in (5) can be equivalently expressed as

$$X_j \sim \text{MP}(X_j(\Omega), \tilde{G}), \quad \tilde{G} \sim \text{DP}(\gamma_0, \tilde{G}_0) \\ X_j(\Omega) \sim \text{Pois}(\alpha), \quad \alpha \sim \text{Gamma}(\gamma_0, p/(J(1-p))), \quad (9)$$

where $G = \alpha \tilde{G}$ and $G_0 = \gamma_0 \tilde{G}_0$. Thus without modeling $X_j(\Omega)$ and $\alpha = G(\Omega)$ as random variables, the gamma-Poisson process becomes a Dirichlet process, which is widely used for mixture modeling [16], [18], [19], [21], [33]. Note that for the Dirichlet process, the inference of γ_0 relies on a data augmentation method of [18] when \tilde{G}_0 is continuous, and no rigorous inference for γ_0 is available when \tilde{G}_0 is discrete. Whereas for the proposed gamma-Poisson process augmented from the NB process, as shown in (7), regardless of whether the base measure G_0 is continuous or discrete, γ_0 has an analytic conditional gamma posterior, with conditional expectation $\mathbb{E}[\gamma_0|L, p] = \frac{e_0 + L(\Omega)}{f_0 - \ln(1-p)}$.

4.5 Restrictions of the Gamma-Poisson Process

The Poisson process has an equal-dispersion assumption for count modeling. For mixture modeling of grouped data, the gamma-Poisson (NB) process might be too restrictive in that, as shown in (9), it implies the same mixture proportions across groups, and as shown in (6),

it implies the same count distribution on each distinct atom. This motivates us to consider adding an additional layer into the gamma-Poisson process or using a different distribution other than the Poisson to model the counts for grouped data. As shown in Section 3, the NB distribution is an ideal candidate, not only because it allows overdispersion, but also because it can be augmented into either a gamma-Poisson or a compound Poisson representations and it can be used together with the CRT distribution to form a bivariate distribution that jointly models the counts of customers and tables.

5 JOINT COUNT AND MIXTURE MODELING OF GROUPED DATA

In this Section we couple the gamma process with the NB process to construct a gamma-NB process, which is well suited for modeling grouped data. We derive analytic conditional posteriors for this construction and show that it can be reduced to an HDP.

5.1 Gamma-Negative Binomial Process

For joint count and mixture modeling of grouped data, e.g., topic modeling where a document consists of a group of exchangeable words, we replace the Poisson processes in (5) with NB processes. Sharing the NB dispersion across groups while making the probability parameters be group dependent, we construct a gamma-NB process as

$$X_j \sim \text{NBP}(G, p_j), \quad G \sim \text{GaP}(c, G_0). \quad (10)$$

With $G \sim \text{GaP}(c, G_0)$ expressed as $G = \sum_{k=1}^{\infty} r_k \delta_{\omega_k}$, a draw from $\text{NBP}(G, p_j)$ can be expressed as $X_j = \sum_{k=1}^{\infty} n_{jk} \delta_{\omega_k}$, $n_{jk} \sim \text{NB}(r_k, p_j)$.

The gamma-NB process can be augmented as a gamma-gamma-Poisson process as

$$X_j \sim \text{PP}(\Theta_j), \quad \Theta_j \sim \text{GaP}\left(\frac{1-p_j}{p_j}, G\right), \quad G \sim \text{GaP}(c, G_0) \quad (11)$$

and with $\theta_{jk} = \Theta_j(\omega_k)$, we have $n_{jk} \sim \text{Pois}(\theta_{jk})$, $\theta_{jk} \sim \text{Gamma}(r_k, p_j/(1-p_j))$. This construction introduces gamma processes $\{\Theta_j\}$, whose normalization provide group-specific random probability measures $\{\Theta_j\}$ for mixture modeling. The gamma-NB process can also be augmented as

$$X_j \sim \sum_{t=1}^{L_j} \text{Log}(p_j), \quad L_j \sim \text{PP}(-G \ln(1-p_j)), \\ G \sim \text{GaP}(c, G_0), \quad (12)$$

which is equivalent in distribution to

$$L_j \sim \text{CRTP}(X_j, G), \quad X_j \sim \text{NBP}(G, p_j), \quad G \sim \text{GaP}(c, G_0) \quad (13)$$

according to Corollary 4. These three closely related constructions are graphically presented in Fig. 2.

With Corollaries 2 and 4, $p' := \frac{-\sum_j \ln(1-p_j)}{c-\sum_j \ln(1-p_j)}$ and $L := \sum_j L_j$, we further have two equivalent augmentations:

$$L \sim \sum_{t=1}^{L'} \text{Log}(p'), \quad L' \sim \text{PP}(-G_0 \ln(1-p')); \quad (14)$$

$$L' \sim \text{CRTP}(L, G_0), \quad L \sim \text{NBP}(G_0, p'). \quad (15)$$

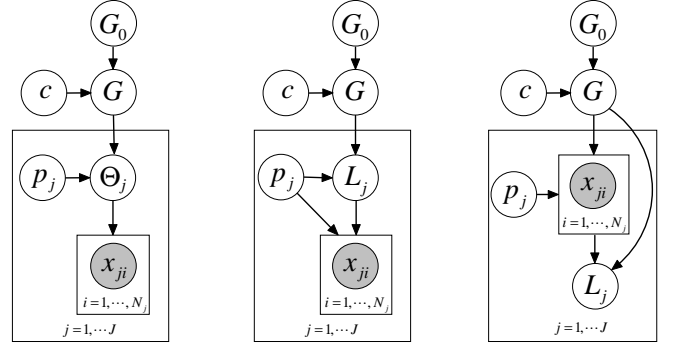


Fig. 2. Graphical models of the gamma-negative binomial process under the gamma-gamma-Poisson (left), gamma-compound Poisson (center), and gamma-negative binomial-Chinese restaurant table constructions (right). The center and right constructions are equivalent in distribution.

These augmentations allow us to derive a sequence of closed-form update equations, as described below.

5.2 Posterior Analysis and Predictive Distribution

With $p_j \sim \text{Beta}(a_0, b_0)$ and (10), we have

$$(p_j | -) \sim \text{Beta}(a_0 + X_j(\Omega), b_0 + G(\Omega)). \quad (16)$$

Using (13) and (15), we have

$$L_j | X_j, G \sim \text{CRTP}(X_j, G), \quad (17)$$

$$L' | L, G_0 \sim \text{CRTP}(L, G_0). \quad (18)$$

If G_0 is finite and continuous, we have $G_0(\omega) \rightarrow 0 \forall \omega \in \Omega$ and thus $L'(\Omega) | L, G_0 = \sum_{\omega \in \Omega} \delta(L(\omega) > 0) = \sum_{\omega \in \Omega} \delta(\sum_j X_j(\omega) > 0) = K^+$; if G_0 is discrete as $G_0 = \sum_{k=1}^K \frac{\gamma_0}{K} \delta_{\omega_k}$, then $L'(\omega_k) = \text{CRT}(L(\omega_k), \frac{\gamma_0}{K}) \geq 1$ if $\sum_j X_j(\omega_k) > 0$, thus $L'(\Omega) \geq K^+$. In either case, let $\gamma_0 = G_0(\Omega) \sim \text{Gamma}(e_0, 1/f_0)$, with the gamma-Poisson conjugacy on (14) and (12), we have

$$\gamma_0 | \{L'(\Omega), p'\} \sim \text{Gamma}\left(e_0 + L'(\Omega), \frac{1}{f_0 - \ln(1-p')}\right), \quad (19)$$

$$G | G_0, \{L_j, p_j\} \sim \text{GaP}\left(c - \sum_j \ln(1-p_j), G_0 + L\right). \quad (20)$$

Using the gamma-Poisson conjugacy on (11), we have

$$\Theta_j | G, X_j, p_j \sim \text{GaP}(1/p_j, G + X_j). \quad (21)$$

Since the data $\{x_{ji}\}_i$ are exchangeable within group j , conditioning on $X_j^{-i} = X_j \setminus X_{ji}$ and G , with Θ_j marginalized out, we have

$$X_{ji} | G, X_j^{-i} \sim \frac{\mathbb{E}[\Theta_j | G, X_j^{-i}]}{\mathbb{E}[\Theta_j(\Omega) | G, X_j^{-i}]} \\ = \frac{G}{G(\Omega) + X_j(\Omega) - 1} + \frac{X_j^{-i}}{G(\Omega) + X_j(\Omega) - 1}. \quad (22)$$

This prediction rule is similar to that of the Chinese restaurant franchise (CRF) [24].

5.3 Relationship with Hierarchical Dirichlet Process

With Corollary 3 and Section 2.2.1, we can equivalently express the gamma-gamma-Poisson process in (11) as

$$\begin{aligned} X_j &\sim \text{MP}(X_j(\Omega), \tilde{\Theta}_j), \quad \tilde{\Theta}_j \sim \text{DP}(\alpha, \tilde{G}), \\ X_j(\Omega) &\sim \text{Pois}(\theta_j), \quad \theta_j \sim \text{Gamma}(\alpha, p_j/(1-p_j)), \\ \alpha &\sim \text{Gamma}(\gamma_0, 1/c), \quad \tilde{G} \sim \text{DP}(\gamma_0, \tilde{G}_0), \end{aligned} \quad (23)$$

where $\Theta_j = \theta_j \tilde{\Theta}_j$, $G = \alpha \tilde{G}$ and $G_0 = \gamma_0 \tilde{G}_0$. Without modeling $X_j(\Omega)$ and θ_j as random variables, (23) becomes an HDP [24]. Thus the augmented and then normalized gamma-NB process leads to an HDP. However, we cannot return from the HDP to the gamma-NB process without modeling $X_j(\Omega)$ and θ_j as random variables. Theoretically, they are distinct in that the gamma-NB process is a completely random measure, assigning independent random variables into any disjoint Borel sets $\{A_q\}_{1,Q}$ in Ω , and the count $X_j(A)$ has the distribution as $X_j(A) \sim \text{NB}(G(A), p_j)$; by contrast, due to normalization, the HDP is not, and marginally

$$X_j(A) \sim \text{Beta-Binomial}(X_j(\Omega), \alpha \tilde{G}(A), \alpha(1 - \tilde{G}(A))).$$

Practically, the gamma-NB process can exploit Corollary 4 and the gamma-Poisson conjugacy to achieve analytic conditional posteriors. The inference of the HDP is a challenge and it is usually solved through alternative constructions such as the CRF and stick-breaking representations [24], [26]. In particular, both concentration parameters α and γ_0 are nontrivial to infer [24], [25] and they are often simply fixed [26]. One may apply the data augmentation method of [18] to sample α and γ_0 . However, if \tilde{G}_0 is discrete as $\tilde{G}_0 = \sum_{k=1}^K \frac{1}{K} \delta_{\omega_k}$, which is of practical value and becomes a continuous base measure as $K \rightarrow \infty$ [24], [25], [33], then using that method to sample γ_0 is only approximately correct, which may result in a biased estimate in practice, especially if K is not sufficiently large.

By contrast, in the gamma-NB process, the shared G can be analytically updated with (20) and $G(\Omega)$ plays the role of α in the HDP, which is readily available as

$$(G(\Omega)|-) \sim \text{Gamma}\left(\gamma_0 + L(\Omega), \frac{1}{c - \sum_j \ln(1-p_j)}\right) \quad (24)$$

and as in (19), regardless of whether the base measure is continuous, the total mass γ_0 has an analytic gamma posterior. Equation (24) also intuitively shows how the NB probability parameters $\{p_j\}$ govern the variations among $\{\tilde{\Theta}_j\}$ in the gamma-NB process. In the HDP, p_j is not explicitly modeled, and since its value appears irrelevant when taking the normalized constructions in (23), it is usually treated as a nuisance parameter and perceived as $p_j = 0.5$ when needed for interpretation.

Another related model is the DILN-HDP in [50], where group-specific Dirichlet processes are normalized from gamma processes, with the gamma scale parameters either fixed as $\frac{p_j}{1-p_j} = 1$ or learned with a log Gaussian process prior. Yet no analytic conditional posteriors are provided and Gibbs sampling is not considered as a

viable option. The main purpose of [50] is introducing correlations between mixture components. It would be interesting to compare the differences between learning the $\{p_j\}$ with beta priors and learning the gamma scale parameters with the log Gaussian process prior.

6 THE NEGATIVE BINOMIAL PROCESS FAMILY

The gamma-NB process shares the NB dispersion across groups while the NB probability parameters are group dependent. Since the NB distribution has two adjustable parameters, it is natural to wonder whether one can explore sharing the NB probability measure across groups, while making the NB dispersion parameters group specific or atom dependent. That kind of construction would be distinct from both the gamma-NB process and HDP in that Θ_j has space dependent scales, and thus its normalization $\tilde{\Theta}_j = \frac{\Theta_j}{\Theta_j(\Omega)}$, still a random probability measure, no longer follows a Dirichlet process.

It is natural to let the NB probability measure be drawn from the beta process [31], [32]. In fact, the first discovered member of the NB process family is a beta-NB process [8]. A main motivation of that construction is observing that the beta and Bernoulli distributions are conjugate and the beta-Bernoulli process is found to be quite useful for dictionary learning [51]–[54], whereas although the beta distribution is also conjugate to the NB distribution, there is apparent lack of exploitation of that relationship [8].

A beta-NB process [8], [14] is constructed by letting

$$X_j \sim \text{NBP}(r_j, B), \quad B \sim \text{BP}(c, B_0). \quad (25)$$

With $B \sim \text{BP}(c, B_0)$ expressed as $B = \sum_{k=1}^{\infty} p_k \delta_{\omega_k}$, a random draw from $\text{NBP}(r_j, B)$ can be expressed as

$$X_j = \sum_{k=1}^{\infty} n_{jk} \delta_{\omega_k}, \quad n_{jk} \sim \text{NB}(r_j, p_k). \quad (26)$$

Under this construction, the NB probability measure is shared and the NB dispersion parameters are group dependent. Note that if $\{r_j\}$ are fixed as one, then the beta-NB process reduces to the beta-geometric process, related to the one for count modeling discussed in [12]; if $\{r_j\}$ are empirically set to some other values, then the beta-NB process reduces to the one proposed in [14]. These simplifications are not considered in the paper, as they are often overly restrictive.

The asymptotic behavior of the beta-NB process with respect to the NB dispersion parameter is studied in [14]. Such analysis is not provided here as we infer NB dispersion parameters from the data, which usually do not have large values due to overdispersion. In [14], the beta-NB process is treated comparable to a gamma-Poisson process and is considered less flexible than the HDP, motivating the construction of a hierarchical-beta-NB process. By contrast, in this paper, with the beta-NB process augmented as a beta-gamma-Poisson process, one can draw group-specific Poisson rate measures for count modeling and then use their normalization to provide group-specific random probability measures

for mixture modeling; therefore, the beta-NB process, gamma-NB process and HDP are treated comparable to each other in hierarchical structures and are all considered suitable for mixed-membership modeling.

As in [8], we may also consider a marked-beta-NB process, with both the NB probability and dispersion measures shared, in which each point of the beta process is marked with an independent gamma random variable. Thus a draw from the marked-beta process becomes $(R, B) = \sum_{k=1}^{\infty} (r_k, p_k) \delta_{\omega_k}$, and a draw from the NB process $X_j \sim \text{NBP}(R, B)$ becomes

$$X_j = \sum_{k=1}^{\infty} n_{jk} \delta_{\omega_k}, \quad n_{jk} \sim \text{NB}(r_k, p_k). \quad (27)$$

With the beta-NB conjugacy, the posterior of B is tractable in both the beta-NB and marked-beta-NB processes [8], [14], [15]. Similar to the marked-beta-NB process, we may also consider a marked-gamma-NB process, where each point of the gamma process is marked with an independent beta random variable, whose performances is found to be similar.

If it is believed that there are excessive number of zeros, governed by a process other than the NB process, we may introduce a zero inflated NB process as $X_j \sim \text{NBP}(RZ_j, p_j)$, where $Z_j \sim \text{BeP}(B)$ is drawn from the Bernoulli process [31] and $(R, B) = \sum_{k=1}^{\infty} (r_k, \pi_k) \delta_{\omega_k}$ is drawn from a gamma marked-beta process, thus a draw from $\text{NBP}(RZ_j, p_j)$ can be expressed as $X_j = \sum_{k=1}^{\infty} n_{jk} \delta_{\omega_k}$, with

$$n_{jk} \sim \text{NB}(r_k b_{jk}, p_j), \quad b_{jk} = \text{Bernoulli}(\pi_k). \quad (28)$$

This construction can be linked to the focused topic model in [55] with appropriate normalization, with advantages that there is no need to fix $p_j = 0.5$ and the inference is fully tractable. The zero inflated construction can also be linked to models for real valued data using the Indian buffet process (IBP) or beta-Bernoulli process spike-and-slab prior [56]–[59]. Below we apply various NB processes for topic modeling and illustrate the differences between them.

7 NEGATIVE BINOMIAL PROCESS TOPIC MODELING AND POISSON FACTOR ANALYSIS

We consider topic modeling of a document corpus, a special case of mixture modeling of grouped data, where the words of the j th document x_{j1}, \dots, x_{jN_j} constitute a group \mathbf{x}_j (N_j words in document j), each word x_{ji} is an exchangeable group member indexed by v_{ji} in a vocabulary with V unique terms. Each word x_{ji} is drawn from a topic $\phi_{z_{ji}}$ as $x_{ji} \sim F(\phi_{z_{ji}})$, where $z_{ji} = 1, 2, \dots, \infty$ is the topic index and the likelihood $F(x_{ji}; \phi_k)$ is simply $\phi_{v_{ji}k}$, the probability of word x_{ji} under topic $\phi_k = (\phi_{1k}, \dots, \phi_{V_k})^T \in \mathbb{R}_+^V$, with $\sum_{v=1}^V \phi_{vk} = 1$. We refer to NB process mixture modeling of grouped words $\{\mathbf{x}_j\}_{1,J}$ as NB process topic modeling.

For the gamma-NB process described in Section 5, with the gamma process expressed as $G = \sum_{k=1}^{\infty} r_k \delta_{\phi_k}$, we can express the hierarchical model as

$$\begin{aligned} x_{ji} &\sim F(\phi_{z_{ji}}), \quad \phi_k \sim g_0(\phi_k), \quad N_j = \sum_{k=1}^{\infty} n_{jk}, \\ n_{jk} &\sim \text{Pois}(\theta_{jk}), \quad \theta_{jk} \sim \text{Gamma}(r_k, p_j/(1-p_j)) \end{aligned} \quad (29)$$

where $g_0(d\phi) = G_0(d\phi)/\gamma_0$. With $\theta_j = \Theta_j(\Omega) = \sum_{k=1}^{\infty} \theta_{jk}$, $\mathbf{n}_j = (n_{j1}, \dots, n_{j\infty})^T$ and $\boldsymbol{\theta}_j = (\theta_{j1}, \dots, \theta_{j\infty})^T$, using Corollary 3, we can equivalently express N_j and n_{jk} in (29) as

$$N_j \sim \text{Pois}(\theta_j), \quad \mathbf{n}_j \sim \text{Mult}(N_j; \boldsymbol{\theta}_j/\theta_j). \quad (30)$$

Since $\{x_{ji}\}_{i=1, N_j}$ are fully exchangeable, rather than drawing \mathbf{n}_j as in (30), we may equivalently draw it as

$$z_{ji} \sim \text{Discrete}(\boldsymbol{\theta}_j/\theta_j), \quad n_{jk} = \sum_{i=1}^{N_j} \delta(z_{ji} = k). \quad (31)$$

This provides further insights on uniting the seemingly distinct problems of count and mixture modeling.

Denote $n_{vjk} = \sum_{i=1}^{N_j} \delta(z_{ji} = k, v_{ji} = v)$, $n_{v \cdot k} = \sum_j n_{vjk}$ and $n_{\cdot k} = \sum_j n_{jk}$. For modeling convenience, we place Dirichlet priors on topics $\phi_k \sim \text{Dir}(\eta, \dots, \eta)$, then for the gamma-NB process topic model, we have

$$\Pr(z_{ji} = k | -) \propto \phi_{v_{ji}k} \theta_{jk}, \quad (32)$$

$$(\phi_k | -) \sim \text{Dir}(\eta + n_{1 \cdot k}, \dots, \eta + n_{V \cdot k}), \quad (33)$$

which would be the same for the other NB processes, since the gamma-NB process differs from them only on how the gamma priors of θ_{jk} and consequently the NB priors of n_{jk} are constituted. For example, marginalizing out θ_{jk} , we have $n_{jk} \sim \text{NB}(r_k, p_j)$ for the gamma-NB process, $n_{jk} \sim \text{NB}(r_j, p_k)$ for the beta-NB process, $n_{jk} \sim \text{NB}(r_k, p_k)$ for both the marked-beta-NB and marked-gamma-NB processes, and $n_{jk} \sim \text{NB}(r_k b_{jk}, p_j)$ for the zero-inflated-NB process.

7.1 Poisson Factor Analysis

Under the bag-of-words representation, without losing information, we can form $\{\mathbf{x}_j\}_{1,J}$ as a term-document count matrix $\mathbf{M} \in \mathbb{R}^{V \times J}$, where m_{vj} counts the number of times term v appears in document j . Given $K \leq \infty$ and \mathbf{M} , discrete latent variable models assume that each entry m_{vj} can be explained as a sum of smaller counts, each produced by one of the K hidden factors, or in the case of topic modeling, a hidden topic. We can factorize \mathbf{M} under the Poisson likelihood as

$$\mathbf{M} \sim \text{Pois}(\boldsymbol{\Phi} \boldsymbol{\Theta}),$$

where $\boldsymbol{\Phi} \in \mathbb{R}^{V \times K}$ is the factor loading matrix, each column of which is an atom encoding the relative importance of each term, and $\boldsymbol{\Theta} \in \mathbb{R}^{K \times J}$ is the factor score matrix, each column of which encodes the relative importance of each atom in a sample. This is called Poisson Factor Analysis (PFA) [8].

As in [8], [60], we may augment $m_{vj} \sim \text{Pois}(\sum_{k=1}^K \phi_{vk} \theta_{jk})$ as

$$m_{vj} = \sum_{k=1}^K n_{vjk}, \quad n_{vjk} \sim \text{Pois}(\phi_{vk} \theta_{jk}).$$

If $\sum_{v=1}^V \phi_{vk} = 1$, we have $n_{jk} \sim \text{Pois}(\theta_{jk})$, and with Corollary 3 and $\boldsymbol{\theta}_j = (\theta_{j1}, \dots, \theta_{jK})^T$, we

also have $(n_{vj1}, \dots, n_{vjK} | -) \sim \text{Mult}(m_{vj}; \frac{\phi_{v1}\theta_{j1}}{\sum_{k=1}^K \phi_{vk}\theta_{jk}}, \dots, \frac{\phi_{vK}\theta_{jK}}{\sum_{k=1}^K \phi_{vk}\theta_{jk}})$, $(n_{v\cdot 1}, \dots, n_{v\cdot K} | -) \sim \text{Mult}(n_{v\cdot}; \phi_k)$, and $(n_{j1}, \dots, n_{jK} | -) \sim \text{Mult}(N_j; \theta_j)$, which would lead to (32) under the assumption that the words $\{x_{ji}\}_i$ are exchangeable and (33) if $\phi_k \sim \text{Dir}(\eta, \dots, \eta)$. Thus topic modeling with the NB process can be considered as factorization of the term-document count matrix under the Poisson likelihood as $\mathbf{M} \sim \text{Pois}(\Phi\Theta)$.

PFA provides a unified framework to connect previously proposed discrete latent variable models, such as those in [5]–[7], [55], [61]. As discussed in detail in [8], these models mainly differ on how the priors of ϕ_{vk} and θ_{jk} are constituted and how the inferences are implemented. For example, nonnegative matrix factorization [61] with an objective function of minimizing the Kullback-Leibler (KL) divergence $D_{\text{KL}}(\mathbf{M} || \Phi\Theta)$ is equivalent to the ML estimation of Φ and Θ under PFA, and latent Dirichlet allocation (LDA) [5] is equivalent to a PFA with Dirichlet priors imposed on both ϕ_k and θ_j .

7.2 Negative Binomial Process Topic Modeling

From the point view of PFA, an NB process topic model factorizes the term-document count matrix under the constraints that each factor sums to one and the factor scores are gamma distributed random variables, and consequently, the number of words assigned to a topic (factor/atom) follows an NB distribution. Depending on how the NB distributions are parameterized, as shown in Table 1, we can construct a variety of NB process topic models, which can also be connected to a large number of previously proposed parametric and nonparametric topic models. For a deeper understanding on how the counts are modeled, we also show in Table 1 both the variance-to-mean ratio (VMR) and overdispersion level (ODL) implied by these settings. Eight differently constructed NB processes are considered:

- (i) The NB process described in Section 4 is used for topic modeling. It improves over the count-modeling gamma-Poisson process discussed in [11], [12] in that it unites mixture modeling and has closed-form conditional posteriors. Although this is a nonparametric model supporting an infinite number of topics, requiring $\{\theta_{jk}\}_{j=1, J} \equiv r_k$ may be too restrictive.
- (ii) Related to LDA [5] and Dir-PFA [8], the NB-LDA is also a parametric topic model that requires tuning the number of topics. It is constructed by replacing the topic weights of the Gamma-NB process in (29) as $\theta_{jk} \sim \text{Gamma}(r_j, p_j/(1-p_j))$. It uses document dependent r_j and p_j to learn the smoothing of the topic weights, and it lets $r_j \sim \text{Gamma}(\gamma_0, 1/c)$, $\gamma_0 \sim \text{Gamma}(e_0, 1/f_0)$ to share statistical strength between documents.
- (iii) Related to the HDP [24], the NB-HDP model is constructed by fixing $p_j/(1-p_j) \equiv 1$ (i.e., $p_j \equiv 0.5$) in (29). It is also comparable to the HDP in [50] that constructs group-specific Dirichlet processes with

normalized gamma processes, whose scale parameters are also set as one.

- (iv) The NB-FTM model is constructed by replacing the topic weights in (29) as $\theta_{jk} \sim \text{Gamma}(r_k b_{jk}, p_j/(1-p_j))$, with $p_j \equiv 0.5$ and b_{jk} drawn from a beta-Bernoulli process that is used to explicitly model zero counts. It is the same as the sparse-gamma-gamma-PFA ($S\gamma\Gamma$ -PFA) in [8] and is comparable to the focused topic model (FTM) [55], which is constructed from the IBP compound Dirichlet process. The Zero-Inflated-NB process improves over these approaches by allowing $\{p_j\}$ to be inferred, which generally yields better data fitting.
- (v) The Gamma-NB process, as shown in (10) and (29), explores sharing the NB dispersion measure across groups, and it improves over the NB-HDP by allowing the learning of $\{p_j\}$. As shown in (23), it reduces to the HDP in [24] without modeling $X_j(\Omega)$ and θ_j as random variables.
- (vi) The Beta-Geometric process is constructed by replacing the topic weights in (29) as $\theta_{jk} \sim \text{Gamma}(1, p_k/(1-p_k))$. It explores sharing the NB probability measure across groups, which is related to the one proposed for count modeling in [12]. It is restrictive that the NB dispersion parameters are fixed as one.
- (vii) The Beta-NB process is constructed by replacing the topic weights in (29) as $\theta_{jk} \sim \text{Gamma}(r_j, p_k/(1-p_k))$. It explores sharing the NB probability measure across groups, which improves over the Beta-Geometric process and the beta-NB process (BNBP) proposed in [14] by providing analytic conditional posteriors of $\{r_j\}$.
- (viii) The Marked-Beta-NB process constructed by replacing the topic weights in (29) as $\theta_{jk} \sim \text{Gamma}(r_k, p_k/(1-p_k))$. It is comparable to the BNBP proposed in [8], with the distinction that it provides analytic conditional posteriors of $\{r_k\}$.

7.3 Approximate and Exact Inference

Although all proposed NB process models have closed-form conditional posteriors, they contain countably infinite atoms that are infeasible to explicitly represent in practice. This infinite dimensional problem can be addressed by using a discrete base measure with K atoms, i.e., truncating the total number of atoms to be K , and then doing Bayesian inference via block Gibbs sampling [62]. This is a very general approach and is used in our experiments to make a fair comparison between a wide variety of models. Block gibbs sampling for the Gamma-NB process is described in Appendix B; block gibbs sampling for other NB processes and related algorithms in Table 1 can be similarly derived, as described in [15] and omitted here for brevity. The infinite dimensional problem can also be addressed by discarding the atoms with weights smaller than a small constant ϵ [22] or by modifying the Lévy measure to

TABLE 1

A variety of negative binomial processes are constructed with distinct sharing mechanisms, reflected with which parameters from r_k, r_j, p_k, p_j and π_k (b_{jk}) are inferred (indicated by a check-mark \checkmark), and the implied variance-mean-ratio (VMR) and overdispersion level (ODL) for counts $\{n_{jk}\}_{j,k}$. They are applied for topic modeling, a typical example of mixture modeling of grouped data. Related algorithms are shown in the last column.

Algorithms	θ_{jk}	r_k	r_j	p_k	p_j	π_k	VMR	ODL	Related Algorithms
NB	$\theta_{jk} \equiv r_k$	\checkmark					$(1-p)^{-1}$	r_k^{-1}	Gamma-Poisson [11], Gamma-Poisson [12]
NB-LDA	\checkmark		\checkmark		\checkmark		$(1-p_j)^{-1}$	r_j^{-1}	LDA [5], Dir-PFA [8]
NB-HDP	\checkmark	\checkmark			0.5		2	r_k^{-1}	HDP [24], DILN-HDP [50]
NB-FTM	\checkmark	\checkmark			0.5	\checkmark	2	$(r_k)^{-1} b_{jk}$	FTM [55], $S_\gamma\Gamma$ -PFA [8]
Beta-Geometric	\checkmark		1	\checkmark			$(1-p_k)^{-1}$	1	Beta-Geometric [12], BNB [8], BNB [14]
Beta-NB	\checkmark		\checkmark	\checkmark			$(1-p_k)^{-1}$	r_j^{-1}	BNB [8], BNB [14]
Gamma-NB	\checkmark	\checkmark			\checkmark		$(1-p_j)^{-1}$	r_k^{-1}	CRF-HDP [24], [25]
Marked-Beta-NB	\checkmark	\checkmark		\checkmark			$(1-p_k)^{-1}$	r_k^{-1}	BNB [8]

make its integral over the whole space be finite [8]. A sufficiently large (small) K (ϵ) usually provides a good approximation, however, there is an increasing risk of wasting computation as the truncation level gets larger.

To avoid truncation, the slice sampling scheme of [63] has been utilized for the Dirichlet process and normalized random measure based mixture models [64]–[66]. With auxiliary slice latent variables introduced to allow adaptive truncations in each MCMC interaction, the infinite dimensional problem is transformed into a finite one. This method has also been applied to the beta-Bernoulli process [67] and the beta-NB process [14]. It would be interesting to investigate slice sampling for the NB process based count and mixture models, which provide likelihoods that might be more amenable to posterior simulation since no normalization is imposed on the weights of the atoms. As slice sampling is not the focus of this paper, we leave it for future study.

Both the block Gibbs sampler and the slice sampler explicitly represent a finite set of atoms for posterior simulation, and algorithms based on these samplers are commonly referred as “conditional” methods [64], [68]. Another approach of solving the infinite dimensional problem is employing a collapsed inference scheme that marginalizes out the atoms and their weights [11], [18], [19], [21], [24], [69]. Algorithms based on the collapsed inference scheme are usually referred as “marginal” methods [64], [68]. A well-defined prediction rule is usually required to develop a collapsed Gibbs sampler, and the conjugacy between the likelihood and the prior distribution of atoms is desired to avoid numerical integrations. In topic modeling, a word is linked to a Dirichlet distributed atom with a multinomial likelihood, thus the atoms can be analytically marginalized out; since their weights can also be marginalized out as in (22), we may develop a collapsed Gibbs sampler for the gamma-NB process based topic models. As the collapsed inference scheme is not the focus of this paper and the prediction rules for other NB processes need further investigation, we leave them for future study.

8 EXAMPLE RESULTS AND DISCUSSIONS

Motivated by Table 1, we consider topic modeling using a variety of NB processes. We compare them with LDA [5], [70] and CRF-HDP [24], [25], in which the latent count n_{jk} is marginally distributed as

$$n_{jk} \sim \text{Beta-Binomial}(N_j, \alpha \tilde{r}_k, \alpha(1 - \tilde{r}_k))$$

with \tilde{r}_k fixed as $1/K$ in LDA and learned from the data in CRF-HDP. For fair comparison, they are all implemented with block Gibbs sampling using a discrete base measure with K atoms, and for the first fifty iterations, the Gamma-NB process with $r_k \equiv 50/K$ and $p_j \equiv 0.5$ is used for initialization. We set K large enough that only a subset of the K atoms would be used by the data. We consider 2500 Gibbs sampling iterations and collect the last 1500 samples.

We consider the Psychological Review¹ corpus, restricting the vocabulary to terms that occur in five or more documents. The corpus includes 1281 abstracts from 1967 to 2003, with $V = 2566$ and 71,279 total word counts. We randomly select 20%, 40%, 60% or 80% of the words from each document to learn a document dependent probability for each term v and calculate the per-word perplexity on the held-out words as

$$\text{Perplexity} = \exp\left(-\frac{1}{y_{..}} \sum_{j=1}^J \sum_{v=1}^V y_{jv} \log f_{jv}\right), \quad (34)$$

where $f_{jv} = \frac{\sum_{s=1}^S \sum_{k=1}^K \phi_{vk}^{(s)} \theta_{jk}^{(s)}}{\sum_{s=1}^S \sum_{v=1}^V \sum_{k=1}^K \phi_{vk}^{(s)} \theta_{jk}^{(s)}}$, y_{jv} is the number of words held out at term v in document j , $y_{..} = \sum_{j=1}^J \sum_{v=1}^V y_{jv}$, and $s = 1, \dots, S$ are the indices of collected samples. Note that the per-word perplexity is equal to V if $f_{jv} = \frac{1}{V}$, thus it should be no greater than V for a topic model that works appropriately. The final results are averaged over five random training/testing partitions. The performance measure is the same as the one used in [8] and similar to those in [26], [71], [72].

Note that the perplexity per held-out word is a fair metric to compare topic models. As analyzed in Section 7, NB process topic models can also be considered as factor analysis of the term-document count matrix under

1. http://psixp.ss.uci.edu/research/programs_data/toolbox.htm

the Poisson likelihood, with ϕ_k as the k th factor that sums to one and θ_{jk} as the factor score of the j th document on the k th factor, which can be further linked to other discrete latent variable models. If except for proportions $\bar{\theta}_j$ and \tilde{r} , the absolute values, e.g., θ_{jk} , r_k and p_k , are also of interest, then the NB process based count and mixture models would be more appropriate than the Dirichlet process based mixture models.

We show in Fig. 3 the NB dispersion and probability parameters learned by various NB process topic models listed in Table 1, revealing distinct sharing mechanisms and model properties. In Fig. 4 we compare the per-held-out-word prediction performance of various algorithms. We set the parameters as $c = 1$, $\eta = 0.05$ and $a_0 = b_0 = e_0 = f_0 = 0.01$. For LDA and NB-LDA, we search K for optimal performance. All the other NB process topic models are nonparametric Bayesian models that can automatically learn the number of active topics K^+ for a given corpus. For fair comparison, all the models considered are implemented with block Gibbs sampling, where $K = 400$ is set as an upper-bound.

When $\theta_{jk} \equiv r_k$ is used, as in the NB process, different documents are imposed to have the same topic weights, leading to the worst held-out-prediction performance.

With a symmetric Dirichlet prior $\text{Dir}(\alpha/K, \dots, \alpha/K)$ placed on the topic proportion for each document, the parametric LDA is found to be sensitive to both the number of topics K and the value of the concentration parameter α . We consider $\alpha = 50$, following the suggestion of the topic model toolbox provided for [70]; we also consider an optimized value as $\alpha = 2.5$, based on the results of the CRF-HDP on the same dataset. As shown in Fig. 4, when the number of training words is small, with optimized K and α , the parametric LDA can approach the performance of the nonparametric CRF-HDP; as the number of training words increases, the advantage of learning \tilde{r}_k in the CRF-HDP than fixing $\tilde{r}_k = 1/K$ in LDA becomes clearer. The concentration parameter α is important for both LDA and CRF-HDP since it controls the VMR of the count n_{jk} , which is equal to $(1 - \tilde{r}_k)(\alpha + N_j)/(\alpha + 1)$ based on (34). Thus fixing α may lead to significantly under- or over-estimated variations and then degraded performance, e.g., LDA with $\alpha = 50$ performs much worse than LDA-Optima- α , as shown in Fig. 4.

When (r_j, p_j) is used, as in NB-LDA, different documents are weakly coupled with $r_j \sim \text{Gamma}(\gamma_0, 1/c)$, and the modeling results in Fig. 3 show that a typical document in this corpus usually has a small r_j and a large p_j , thus a large overdispersion level (ODL) and a large variance-to-mean ratio (VMR), indicating highly overdispersed counts on its topic usage. NB-LDA is a parametric topic model that requires tuning the number of topics K . It improves over LDA in that it only has to tune K , whereas LDA has to tune both K and α . With an appropriate K , the parametric NB-LDA may outperform the nonparametric NB-HDP and NB-FTM as the training data percentage increases, showing that

even by learning both the NB parameters r_j and p_j in a document dependent manner, we may get better data fitting than using nonparametric models that fix the NB probability parameters.

When (r_j, p_k) is used to model the latent counts $\{n_{jk}\}_{j,k}$, as in the Beta-NB process, the transition between active and non-active topics is very sharp that p_k is either far from zero or almost zero, as shown in Fig. 3. That is because p_k controls the mean $\mathbb{E}[\sum_j n_{jk}] = p_k/(1-p_k) \sum_j r_j$ and the VMR $(1-p_k)^{-1}$ on topic k , thus a popular topic must also have large p_k and hence large overdispersion measured by the VMR; since the counts $\{n_{jk}\}_j$ are usually overdispersed, particularly true in this corpus, a small p_k indicating a small mean and small VMR is not favored and thus is rarely observed.

The Beta-Geometric process is a special case of the Beta-NB process that $r_j \equiv 1$, which is more than ten times larger than the values inferred by the Beta-NB process on this corpus, as shown in Fig. 3; therefore, to fit the mean $\mathbb{E}[\sum_j n_{jk}] = Jp_k/(1-p_k)$, it has to use a substantially underestimated p_k , leading to severely underestimated variations and thus degraded performance, as confirmed by comparing the curves of the Beta-Geometric and Beta-NB processes in Fig. 4.

When (r_k, p_j) is used, as in the Gamma-NB process, the transition is much smoother that r_k gradually decreases, as shown in Fig. 3. The reason is that r_k controls the mean $\mathbb{E}[\sum_j n_{jk}] = r_k \sum_j p_j/(1-p_j)$ and the ODL r_k^{-1} on topic k , thus popular topics must also have large r_k and hence small overdispersion measured by the ODL, and unpopular topics are modeled with small r_k and hence large overdispersion, allowing rarely and lightly used topics. Therefore, we can expect that (r_k, p_j) would allow more topics than (r_j, p_k) , as confirmed in Fig. 4 (a) that the Gamma-NB process learns 177 active topics, obviously more than the 107 ones of the Beta-NB process. With these analysis, we can conclude that the mean and the amount of overdispersion (measure by the VMR or ODL) for the usage of topic k is positively correlated under (r_j, p_k) and negatively correlated under (r_k, p_j) .

The NB-HDP is a special case of the Gamma-NB process that $p_j \equiv 0.5$. From a mixture modeling viewpoint, fixing $p_j = 0.5$ is a natural choice as p_j appears irrelevant after normalization. However, from a count modeling viewpoint, this would make a restrictive assumption that each count vector $\{n_{jk}\}_{k=1,K}$ has the same VMR of 2. It is also interesting to examine (24), which can be viewed as the concentration parameter α in the HDP, allowing the adjustment of p_j would allow a more flexible model assumption on the amount of variations between the topic proportions, and thus potentially better data fitting.

The CRF-HDP and Gamma-NB process have very similar performance on predicting held-out words, although they have distinct assumption on count modeling: n_{jk} is modeled as an NB distribution in the Gamma-NB process while it is modeled as a beta-binomial distribution in the CRF-HDP. The Gamma-NB process adjust both r_k and p_j to fit the NB distribution, whereas the CRF-HDP

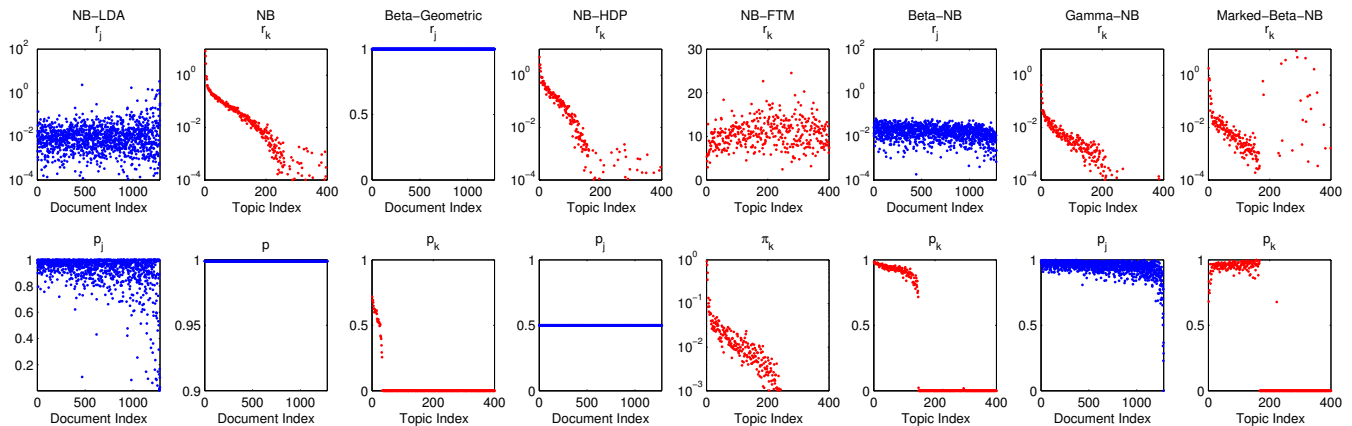


Fig. 3. Distinct sharing mechanisms and model properties are evident between various NB process topic models, by comparing their inferred NB dispersion parameters (r_k or r_j) and probability parameters (p_k or p_j). Note that the transition between active and non-active topics is very sharp when p_k is used and much smoother when r_k is used. Both the documents and topics are ordered in a decreasing order based on the associated number of words. These results are based on the last MCMC iteration, on the Psychological Review corpus with 80% of the words in each document used as training. The values along the vertical axis are shown in either linear or log scales for convenient visualization. Document-specific and topic-specific parameters are shown in blue and red colors, respectively.




Fig. 4. Comparison of per-word perplexity on held out words between various algorithms listed in Table 1 on the Psychological Review corpus. LDA-Optimal- α refers to an LDA algorithm whose topic proportion Dirichlet concentration parameter α is optimized based on the results of the CRF-HDP on the same dataset. (a) With 60% of the words in each document used for training, the performance varies as a function of K in both LDA and NB-LDA, which are parametric models, whereas the NB, Beta-Geometric, NB-HDP, NB-FTM, Beta-NB, CRF-HDP, Gamma-NB and Marked-Beta-NB all infer the number of active topics, which are 225, 28, 127, 201, 107, 161, 177 and 130, respectively, according to the last Gibbs sampling iteration. (b) Per-word perplexities of various algorithms as a function of the percentage of words in each document used for training. The results of LDA and NB-LDA are shown with the best settings of K under each training/testing partition. Nonparametric Bayesian algorithms listed in Table 1 are ranked in the legend from top to bottom according to their overall performance.

learns both α and \tilde{r}_k to fit the beta-binomial distribution. The concentration parameter α controls the VMR of the count n_{jk} as shown in (34), and we find through

experiments that prefixing its value may substantially degrade the performance of the CRF-HDP, thus this option is not considered in the paper and we exploit

the CRF metaphor to update α as in [24], [25].

When (r_k, π_k) is used, as in the NB-FTM model, our results in Fig. 3 show that we usually have a small π_k and a large r_k , indicating topic k is sparsely used across the documents but once it is used, the amount of variation on usage is small. This property might be helpful when there are excessive number of zeros that might not be well modeled by the NB process alone. In our experiments, the more direct approaches of using p_k or p_j generally yield better results, which might not be the case when excessive number of zeros could be better explained with the beta-Bernoulli processes, e.g., when the training words are scarce, the NB-FTM can approach the performance of the Marked-Beta-NB process.

When (r_k, p_k) is used, as in the Marked-Beta-NB process, more diverse combinations of mean and overdispersion would be allowed as both r_k and p_k are now responsible for the mean $\mathbb{E}[\sum_j n_{jk}] = Jr_k p_k / (1 - p_k)$. As observed in Fig. 3, there could be not only large mean and small overdispersion (large r_k and small p_k), indicating a popular topic frequently used by most of the documents, but also large mean and large overdispersion (small r_k and large p_k), indicating a topic heavily used in a relatively small percentage of documents. Thus (r_k, p_k) may combine the advantages of using only r_k or p_k to model topic k , as confirmed by the superior performance of the Marked-Beta-NB process.

9 CONCLUSIONS

We propose a variety of negative binomial (NB) processes for count modeling, which can be naturally applied for the seemingly disjoint problem of mixture modeling. The proposed NB processes are completely random measures, which assign independent random variables to disjoint Borel sets of the measure space, as opposed to Dirichlet processes, whose measures on disjoint Borel sets are negatively correlated. We reveal connections between various discrete distributions and discover unique data augmentation and marginalization methods for the NB process, with which we are able to unite count and mixture modeling, analyze fundamental model properties, and derive efficient Bayesian inference. We demonstrate that the NB process and the gamma-NB process can be recovered from the Dirichlet process and the HDP, respectively. We show in detail the theoretical, structural and computational advantages of the NB process. We examine the distinct sharing mechanisms and model properties of various NB processes, with connections made to existing discrete latent variable models under the Poisson factor analysis framework. Experimental results on topic modeling show the importance of modeling both the NB dispersion and probability parameters, which respectively govern the overdispersion level and variance-to-mean ratio for count modeling.

ACKNOWLEDGMENTS

The authors would like to thank the two anonymous reviewers and the editor for their constructive comments that help improve the manuscript.

REFERENCES

- [1] C. Dean, J. F. Lawless, and G. E. Willmot. A mixed Poisson-inverse-Gaussian regression model. *Canadian Journal of Stat.*, 1989.
- [2] M. Zhou, L. Li, D. Dunson, and L. Carin. Lognormal and gamma mixed negative binomial regression. In *ICML*, 2012.
- [3] J. O. Lloyd-Smith. Maximum likelihood estimation of the negative binomial dispersion parameter for highly overdispersed data, with applications to infectious diseases. *PLoS ONE*, 2007.
- [4] T. Hofmann. Probabilistic latent semantic analysis. In *UAI*, 1999.
- [5] D. Blei, A. Ng, and M. Jordan. Latent Dirichlet allocation. *J. Mach. Learn. Res.*, 2003.
- [6] J. Canny. Gap: a factor model for discrete data. In *SIGIR*, 2004.
- [7] W. Buntine and A. Jakulin. Discrete component analysis. In *Subspace, Latent Structure and Feature Selection Techniques*. Springer-Verlag, 2006.
- [8] M. Zhou, L. Hannah, D. Dunson, and L. Carin. Beta-negative binomial process and Poisson factor analysis. In *AISTATS*, 2012.
- [9] J. F. C. Kingman. *Poisson Processes*. Oxford University Press, 1993.
- [10] R. L. Wolpert and K. Ickstadt. Poisson/gamma random field models for spatial statistics. *Biometrika*, 1998.
- [11] M. K. Titsias. The infinite gamma-Poisson feature model. In *NIPS*, 2008.
- [12] R. J. Thibaux. *Nonparametric Bayesian Models for Machine Learning*. PhD thesis, UC Berkeley, 2008.
- [13] K. T. Miller. *Bayesian Nonparametric Latent Feature Models*. PhD thesis, UC Berkeley, 2011.
- [14] T. Broderick, L. Mackey, J. Paisley, and M. I. Jordan. Combinatorial clustering and the beta negative binomial process. *arXiv:1111.1802v3*, 2012.
- [15] M. Zhou and L. Carin. Augment-and-conquer negative binomial processes. In *NIPS*, 2012.
- [16] T. Ferguson. A Bayesian analysis of some nonparametric problems. *The Annals of Statistics*, 1973.
- [17] C. E. Antoniak. Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems. *Ann. Statist.*, 1974.
- [18] M. D. Escobar and M. West. Bayesian density estimation and inference using mixtures. *JASA*, 1995.
- [19] S. N. MacEachern and P. Müller. Estimating mixture of Dirichlet process models. *Journal of Computational and Graphical Statistics*.
- [20] R. M. Neal. Markov chain sampling methods for Dirichlet process mixture models. *JCGS*, 2000.
- [21] Y. W. Teh. Dirichlet processes. In *Encyclopedia of Machine Learning*. Springer, 2010.
- [22] R. L. Wolpert, M. A. Clyde, and C. Tu. Stochastic expansions using continuous dictionaries: Lévy Adaptive Regression Kernels. *Annals of Statistics*, 2011.
- [23] A. Lijoi, R. H. Mena, and I. Prünster. Controlling the reinforcement in Bayesian non-parametric mixture models. *Journal of the Royal Statistical Society: Series B*, 2007.
- [24] Y. W. Teh, M. I. Jordan, M. J. Beal, and D. M. Blei. Hierarchical Dirichlet processes. *JASA*, 2006.
- [25] E. Fox, E. Sudderth, M. Jordan, and A. Willsky. Developing a tempered HDP-HMM for systems with state persistence. *MIT LIDS, TR #2777*, 2007.
- [26] C. Wang, J. Paisley, and D. M. Blei. Online variational inference for the hierarchical Dirichlet process. In *AISTATS*, 2011.
- [27] J. Pitman and M. Yor. The two-parameter Poisson-Dirichlet distribution derived from a stable subordinator. *The Annals of Probability*, 1997.
- [28] H. Ishwaran and L. F. James. Gibbs sampling methods for stick-breaking priors. *JASA*, 2001.
- [29] J. F. C. Kingman. Completely random measures. *Pacific Journal of Mathematics*, 1967.
- [30] M. I. Jordan. Hierarchical models, nested models and completely random measures. In M.-H. Chen, D. Dey, P. Mueller, D. Sun, and K. Ye, editors, *Frontiers of Statistical Decision Making and Bayesian Analysis: in Honor of James O. Berger*. New York: Springer, 2010.
- [31] R. Thibaux and M. I. Jordan. Hierarchical beta processes and the Indian buffet process. In *AISTATS*, 2007.

- [32] N. L. Hjort. Nonparametric Bayes estimators based on beta processes in models for life history data. *Ann. Statist.*, 1990.
- [33] H. Ishwaran and M. Zarepour. Exact and approximate sum-representations for the Dirichlet process. *Can. J. Statist.*, 2002.
- [34] D. Blackwell and J. MacQueen. Ferguson distributions via Pólya urn schemes. *The Annals of Statistics*, 1973.
- [35] D. Aldous. Exchangeability and related topics. In *Ecole d'Ete de Probabilités de Saint-Flour XIII*, pages 1–198. Springer, 1983.
- [36] J. Pitman. *Combinatorial stochastic processes*. Lecture Notes in Mathematics. Springer-Verlag, 2006.
- [37] M. Greenwood and G. U. Yule. An inquiry into the nature of frequency distributions representative of multiple happenings with particular reference to the occurrence of multiple attacks of disease or of repeated accidents. *Journal of Royal Stat. Soc.*, 1920.
- [38] M. H. Quenouille. A relation between the logarithmic, Poisson, and negative binomial series. *Biometrics*, 1949.
- [39] N. L. Johnson, A. W. Kemp, and S. Kotz. *Univariate Discrete Distributions*. John Wiley & Sons, 2005.
- [40] C. I. Bliss and R. A. Fisher. Fitting the negative binomial distribution to biological data. *Biometrics*, 1953.
- [41] A. C. Cameron and P. K. Trivedi. *Regression Analysis of Count Data*. Cambridge, UK, 1998.
- [42] R. Winkelmann. *Econometric Analysis of Count Data*. Springer, Berlin, 5th edition, 2008.
- [43] M. D. Robinson and G. K. Smyth. Small-sample estimation of negative binomial dispersion, with applications to SAGE data. *Biostatistics*, 2008.
- [44] E. P. Pieters, C. E. Gates, J. H. Matis, and W. L. Sterling. Small sample comparison of different estimators of negative binomial parameters. *Biometrics*, 1977.
- [45] L. J. Willson, J. L. Folks, and J. H. Young. Multistage estimation compared with fixed-sample-size estimation of the negative binomial parameter k . *Biometrics*, 1984.
- [46] J. F. Lawless. Negative binomial and mixed Poisson regression. *Canadian Journal of Statistics*, 1987.
- [47] W. W. Piegorsch. Maximum likelihood estimation for the negative binomial dispersion parameter. *Biometrics*, 1990.
- [48] K. Saha and S. Paul. Bias-corrected maximum likelihood estimator of the negative binomial dispersion parameter. *Biometrics*, 2005.
- [49] E. T. Bradlow, B. G. S. Hardie, and P. S. Fader. Bayesian inference for the negative binomial distribution via polynomial expansions. *Journal of Computational and Graphical Statistics*, 2002.
- [50] J. Paisley, C. Wang, and D. Blei. The discrete infinite logistic normal distribution for mixed-membership modeling. In *AISTATS*, 2011.
- [51] M. Zhou, C. Wang, M. Chen, J. Paisley, D. Dunson, and L. Carin. Nonparametric Bayesian matrix completion. In *IEEE Sensor Array and Multichannel Signal Processing Workshop*, 2010.
- [52] M. Zhou, H. Yang, G. Sapiro, D. B. Dunson, and L. Carin. Dependent hierarchical beta process for image interpolation and denoising. In *AISTATS*, 2011.
- [53] L. Li, M. Zhou, G. Sapiro, and L. Carin. On the integration of topic modeling and dictionary learning. In *ICML*, 2011.
- [54] M. Zhou, H. Chen, J. Paisley, L. Ren, L. Li, Z. Xing, D. Dunson, G. Sapiro, and L. Carin. Nonparametric Bayesian dictionary learning for analysis of noisy and incomplete images. *IEEE TIP*, 2012.
- [55] S. Williamson, C. Wang, K. A. Heller, and D. M. Blei. The IBP compound Dirichlet process and its application to focused topic modeling. In *ICML*, 2010.
- [56] T. L. Griffiths and Z. Ghahramani. Infinite latent feature models and the Indian buffet process. In *NIPS*, 2005.
- [57] D. Knowles and Z. Ghahramani. Infinite sparse factor analysis and infinite independent components analysis. In *Independent Component Analysis and Signal Separation*, 2007.
- [58] J. Paisley and L. Carin. Nonparametric factor analysis with beta process priors. In *ICML*, 2009.
- [59] M. Zhou, H. Chen, J. Paisley, L. Ren, G. Sapiro, and L. Carin. Non-parametric Bayesian dictionary learning for sparse image representations. In *NIPS*, 2009.
- [60] D. B. Dunson and A. H. Herring. Bayesian latent variable models for mixed discrete outcomes. *Biostatistics*, 2005.
- [61] D. D. Lee and H. S. Seung. Algorithms for non-negative matrix factorization. In *NIPS*, 2001.
- [62] H. Ishwaran and M. Zarepour. Markov chain Monte Carlo in approximate Dirichlet and beta two-parameter process hierarchical models. *Biometrika*, 2000.
- [63] P. Damlén, J. Wakefield, and S. Walker. Gibbs sampling for Bayesian non-conjugate and hierarchical models by using auxiliary variables. *Journal of the Royal Statistical Society B*, 1999.
- [64] O. Papaspiliopoulos and G. O. Roberts. Retrospective Markov chain Monte Carlo methods for Dirichlet process hierarchical models. *Biometrika*, 2008.
- [65] S. G. Walker. Sampling the Dirichlet mixture model with slices. *Communications in Statistics Simulation and Computation*, 2007.
- [66] J. E. Griffin and S. G. Walker. Posterior simulation of normalized random measure mixtures. *JCGS*, 2011.
- [67] Y. W. Teh, D. Görür, and Z. Ghahramani. Stick-breaking construction for the Indian buffet process. In *AISTATS*, 2007.
- [68] M. Kalli, J. E. Griffin, and S. G. Walker. Slice sampling mixture models. *Statistics and computing*, 2011.
- [69] R. M. Neal. Markov chain sampling methods for Dirichlet process mixture models. *JCGS*, 2000.
- [70] T. L. Griffiths and M. Steyvers. Finding scientific topics. *PNAS*, 2004.
- [71] A. Asuncion, M. Welling, P. Smyth, and Y. W. Teh. On smoothing and inference for topic models. In *UAI*, 2009.
- [72] H. M. Wallach, I. Murray, R. Salakhutdinov, and D. Mimno. Evaluation methods for topic models. In *ICML*, 2009.

APPENDIX A PROOF OF THEOREM 1

With the PMFs of both the NB and CRT distributions, the PMF of the joint distribution of counts m and l is $f_{M,L}(m, l|r, p) = f_L(l|m, r)f_M(m|r, p) = \frac{\Gamma(r)|s(m,l)|r^l}{\Gamma(m+r)} \frac{\Gamma(r+m)(1-p)^r p^m}{m!\Gamma(r)} = \frac{|s(m,l)|r^l(1-p)^r p^m}{m!}$, which is the same as (4).

Since $m \sim \text{SumLog}(l, p)$ is the summation of l iid $\text{Log}(p)$ random variables, its PGF becomes $C_M(z) = C_U^l(z) = [\ln(1-pz)/\ln(1-p)]^l$, $|z| < p^{-1}$. With $[\ln(1+x)]^l = l! \sum_{n=1}^{\infty} s(n, l)x^n/n!$ and $|s(m, l)| = (-1)^{m-l}s(m, l)$ [39], its PMF can be expressed as

$$f_M(m|l, p) = \frac{C_M^{(m)}(0)}{m!} = \frac{p^m l! |s(m, l)|}{m! [-\ln(1-p)]^l}. \quad (35)$$

Letting $l \sim \text{Pois}(-r \ln(1-p))$, the PMF of the joint distribution of counts m and l is $f_{M,L}(m, l|r, p) = f_M(m|l, p)f_L(l|r, p) = \frac{p^m l! |s(m, l)|}{m! [-\ln(1-p)]^l} \frac{(-r \ln(1-p))^l e^{-r \ln(1-p)}}{l!} = \frac{|s(m, l)| r^l (1-p)^r p^m}{m!}$, which is the same as (4). \square

APPENDIX B BLOCK GIBBS SAMPLING FOR THE GAMMA-NEGATIVE BINOMIAL PROCESS

With $p_j \sim \text{Beta}(a_0, b_0)$, $\gamma_0 \sim \text{Gamma}(e_0, 1/f_0)$ and a discrete base measure as $G_0 = \sum_{k=1}^K \frac{\gamma_0}{K} \delta_{\omega_k}$, following Section 5.2, block Gibbs sampling for (29) proceeds as

$$\begin{aligned} \Pr(z_{ji} = k | -) &\propto F(x_{ji}; \omega_k) \theta_{jk}, \\ (l_{jk} | -) &\sim \text{CRT}(n_{jk}, r_k), \quad (l'_k | -) \sim \text{CRT}\left(\sum_j l_{jk}, \gamma_0/K\right), \\ (p_j | -) &\sim \text{Beta}(a_0 + N_j, b_0 + \sum_k r_k), \\ (\gamma_0 | -) &\sim \text{Gamma}\left(e_0 + \sum_k l'_k, \frac{1}{f_0 - \ln(1-p')}\right), \\ (r_k | -) &\sim \text{Gamma}\left(\gamma_0/K + \sum_j l_{jk}, \frac{1}{c - \sum_j \ln(1-p_j)}\right), \\ (\theta_{jk} | -) &\sim \text{Gamma}(r_k + n_{jk}, p_j), \\ p(\omega_k | -) &\propto g_0(\omega_k) \prod_{z_{ji}=k} F(x_{ji}; \omega_k), \end{aligned} \quad (36)$$

where $p' := \frac{-\sum_j \ln(1-p_j)}{c - \sum_j \ln(1-p_j)}$. Note that when $K \rightarrow \infty$, we have $(l'_k | -) = \delta(\sum_j n_{jk} > 0)$ and thus $\sum_k l'_k = K^+$.