



Augment-and-Conquer Negative Binomial Processes

Mingyuan Zhou

Lawrence Carin

Department of Electrical & Computer Engineering
Duke University, Durham, NC 27708, USA

Introduction

➤ We develop data augmentation methods unique to the negative binomial (NB) distribution.

➤ We perform joint count and mixture modeling under the NB process, using completely random measures (Poisson process, gamma process, beta process, NB process) that are simple to construct and amenable for posterior computation.

➤ We propose to augment-and-conquer the NB process: by “augmenting” a NB process into both the gamma-Poisson and compound Poisson representations, we “conquer” the unification of count and mixture modeling, the analysis of fundamental model properties, and the derivation of efficient Gibbs sampling.

➤ We show that the gamma-NB process can be reduced to the hierarchical Dirichlet process with normalization, highlighting its unique theoretical, structural and computational advantages.

➤ A variety of NB processes with distinct sharing mechanisms are constructed and applied to topic modeling, with connections to existing algorithms, showing the importance of inferring both the NB dispersion and probability parameters.

Poisson Process for Count & Mixture Modeling

➤ Mixture modeling infers probability random measures to assign data points into clusters (mixture components).

➤ Since the number of points assigned to clusters are counts, we consider mixture modeling as a count-modeling problem.

➤ The NB distribution, parameterized by a dispersion parameter and a probability parameter, is used to model overdispersed counts.

➤ The NB distribution is a gamma-Poisson mixture distribution (for hierarchical construction), and can also be represented as a compound Poisson distribution (for inference tractability).


➤ The Poisson process provides not only a way to generate independent counts from each partition of the space, but also a mechanism for mixture modeling, which allocates the observations into any measurable disjoint partition of space, conditioning on the total count and the normalized mean measure.

$$X_j(\Omega) = \sum_{q=1}^Q X_j(A_q), \quad X_j(A_q) \sim \text{Pois}(G(A_q));$$

$$X_j(\Omega) \sim \text{Poisson}(G(\Omega)), \quad [X_j(A_1), \dots, X_j(A_Q)] \sim \text{Mult}(X_j(\Omega); \tilde{G}(A_1), \dots, \tilde{G}(A_Q)).$$

Poisson-Logarithmic bivariate count distribution

The joint distribution of the customer count and table count are equivalent:



Augment-and-Conquer the Negative Binomial Distribution

Lemma 2.1. Denote $s(m, j)$ as Stirling numbers of the first kind [17]. Augment $m \sim \text{NB}(r, p)$ under the compound Poisson representation as $m \sim \sum_{l=1}^L \text{Log}(p)$, $l \sim \text{Pois}(-r \ln(1-p))$, then the conditional posterior of l has PMF

$$\Pr(l = j | m, r) = \frac{\Gamma(r)}{\Gamma(m+r)} |s(m, j)| r^j, \quad j = 0, 1, \dots, m.$$

Lemma 2.2. Let $m \sim \text{NB}(r, p)$, $r \sim \text{Gamma}(r_1, 1/c_1)$, denote $p' = \frac{-\ln(1-p)}{c_1 - \ln(1-p)}$, then m can also be generated from a compound distribution as

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

Gamma-Negative Binomial Process

Gamma-Negative Binomial Process

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

$$X_j \sim \text{PP}(\Lambda_j), \quad \Lambda_j \sim \text{GaP}((1-p_j)/p_j, G), \quad G \sim \text{GaP}(c, G_0)$$

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

$$L = \sum_j L_j \sim \sum_{l'=1}^{l'} \text{Log}(p'), \quad l' \sim \text{PP}(-G_0 \ln(1-p')), \quad p' = \frac{-\sum_j \ln(1-p_j)}{c - \sum_j \ln(1-p_j)}$$

Posterior Analysis

$$\Lambda_j | G, X_j, p_j \sim \text{GaP}(1/p_j, G + X_j) \quad (p_j | -) \sim \text{Beta}(a_0 + X_j(\Omega), b_0 + G(\Omega))$$

$$L_j | X_j, G \sim \text{CRTP}(X_j, G), \quad L' | L, G_0 \sim \text{CRTP}(L, G_0)$$

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

$$G | G_0, \{L_j, p_j\} \sim \text{GaP}(c - \sum_j \ln(1-p_j), G_0 + \sum_j L_j)$$

Predictive Distribution

$$X_{j+1} | G, X_j^{-1} \sim \frac{\mathbb{E}[\Lambda_j | G, X_j^{-1}]}{\mathbb{E}[\Lambda_j(\Omega) | G, X_j^{-1}]} = \frac{G}{G(\Omega) + X_j(\Omega) + 1} + \frac{X_j^{-1}}{G(\Omega) + X_j(\Omega) + 1}$$

Related to the Hierarchical Dirichlet Process

$$X_{j+1} \sim \tilde{\Lambda}_j, \quad \tilde{\Lambda}_j \sim \text{DP}(\alpha, \tilde{G}), \quad \alpha \sim \text{Gamma}(\gamma_0, 1/c), \quad \tilde{G} \sim \text{DP}(\gamma_0, \tilde{G}_0)$$

The NB Process Family and Related Algorithms

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 VMR and ODL for counts $\{n_{jk}\}_{j,k}$. They are applied for topic modeling of a document corpus, a typical example of mixture modeling of grouped data. Related algorithms are shown in the last column.

Algorithms	r_k	r_j	p_k	p_j	π_k	VMR	ODL	Related Algorithms
NB-LDA		\checkmark		\checkmark		$(1-p_j)^{-1}$	r_j^{-1}	LDA [32], Dir-PFA [5]
NB-HDP	\checkmark			0.5		2	r_k^{-1}	HDP[7], DILN-HDP [12]
NB-FTM	\checkmark			0.5	\checkmark	2	$(r_k)^{-1} b_{jk}$	FTM [27], S γ -PFA [5]
Beta-NB		\checkmark	\checkmark			$(1-p_k)^{-1}$	r_j^{-1}	BNBP [5], BNBP [6]
Gamma-NB	\checkmark			\checkmark		$(1-p_j)^{-1}$	r_k^{-1}	CRF-HDP [7, 24]
Marked-Beta-NB	\checkmark		\checkmark			$(1-p_k)^{-1}$	r_k^{-1}	BNBP [5]

$$\text{Beta-NB} \quad n_{jk} \sim \text{NB}(r_j, p_k) \quad \text{NB-FTM} \quad n_{jk} \sim \text{NB}(r_k b_{jk}, p_j)$$

$$\text{Gamma-NB} \quad n_{jk} \sim \text{NB}(r_k, p_j) \quad \text{Marked-Beta-NB} \quad n_{jk} \sim \text{NB}(r_k, p_k)$$

Topic Modeling (Mixture Modeling of Grouped Words)

Comparison of Per-Word Perplexities




Figure 1: Comparison of per-word perplexities on the held-out words between various algorithms. (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-HDP, NB-FTM, Beta-NB, CRF-HDP, Gamma-NB and Marked-Beta-NB all infer the number of active topics, which are 127, 201, 107, 161, 177 and 130, respectively, according to the last Gibbs sampling iteration. (b) Per-word perplexities of various models as a function of the percentage of words in each document used for training. The results of the LDA and NB-LDA are shown with the best settings of K under each training/testing partition.

Inferred NB Process Model Parameters




Figure 2: Distinct sharing mechanisms and model properties are evident between various NB processes, by comparing their inferred parameters. 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 number of words associated with each of them. These results are based on the last Gibbs sampling iteration. The values are shown in either linear or log scales for convenient visualization.