

# Correlated Mixed Membership Modeling for Somatic Mutations

Rahul Mehta

*Department of Bioengineering  
University of Illinois at Chicago  
Chicago, IL, USA  
mehta5@uic.edu*

Muge Karaman

*Department of Bioengineering  
Center for Magnetic Resonance Research  
University of Illinois at Chicago  
Chicago, IL, USA  
mkaraman@uic.edu*

**Abstract**—Recent studies of cancer somatic mutation profiles seek to identify mutations for targeted therapy in personalized medicine. Analysis of profiles, however, is not trivial, as each profile is heterogeneous and there are multiple confounding factors that influence the cause-and-effect relationships between cancer genes such as cancer (sub)type, biological processes, total number of mutations, and non-linear mutation interactions. Moreover, cancer is biologically redundant, i.e., distinct mutations can result in the alteration of similar biological processes, so it is important to identify all possible combinatorial sets of mutations for effective patient treatment. To model this phenomena, we propose the correlated zero-inflated negative binomial process to infer the inherent structure of somatic mutation profiles through latent representations. This stochastic process takes into account different, yet correlated, co-occurring mutations using profile-specific negative binomial dispersion parameters that are mixed with a correlated beta-Bernoulli process and a probability parameter to model profile heterogeneity. These model parameters are inferred by iterative optimization via amortized and stochastic variational inference using the Pan Cancer dataset from The Cancer Genomic Archive (TCGA). By examining the the latent space, we identify biologically relevant correlations between somatic mutations.

**Index Terms**—Bayesian non-parametric, Somatic Mutations, mixed-membership modeling

## I. INTRODUCTION

Discerning the relationships between somatic mutations in cancers is the foundation for targeted treatment and patient subtyping. Since somatic mutations in cancer genomes are often heterogeneous and sparse, where two patients with the same cancer may share only one mutation among thousands, models summarize the high-dimensional interactions into a simpler form. This requires a model that incorporates multiple confounding variables to determine relationships between somatic mutations. Based on current literature [1], [2] mutually exclusive and co-occurring mutations are influenced by non-linear relationships between gene mutation frequencies, biological processes, cancer (sub)types, total number of mutations in a tumor (TML), and positive selection for mutations. The combination of multiple confounding variables and the inherent sparsity of somatic mutation data poses a challenge to understand the underlying co-dependencies between mutations.

Statistical and computational models that try to discover relationships between somatic mutations often decompose a patient's mutation profile into a set of higher-level structures that

closely resemble known biological processes. This approach [3], [4] generally follows a random walk on an existing biological interaction network. This networks can be modeled as a graph,  $G = (V, E)$ , where each vertex,  $V$  is a gene, and the edge,  $E$  denotes the interaction among genes. The network is then modified into a weighted graph, with edge weights representing probability of interactions and vertex weights corresponding to the frequency of a mutation in a gene. A walk is then simulated by starting at a mutated gene and moving to another gene based on the probabilities of edge and vertex weights. The end result is a smaller subnetwork called a functional network that represents an altered biological process. While functional networks have been validated to discover some aberrant genes and pathways, they often result in false positives due to the inherent assumptions made.

The most common compendium of interaction networks widely used to generate functional networks is the Kyoto Encyclopedia of Genes and Genomes (KEGG) [5]. The KEGG interaction networks specify genetic pathways, which are complex graphical networks with directed and undirected edges connecting genes based on their physical and biochemical properties. The genetic pathways are then ascribed to specific biological processes. For example, the biological process of cell apoptosis (cell death) is controlled by two known genetic pathways compromising of a multitude of different genes. The networks within the KEGG database, however, are diverse and recapitulate a disease free patient. Functional networks therefore, assume the interaction networks are also cancer-relevant and disease-specific. As a result, functional networks are generalized to a common patient population and struggle to discriminate between different cancer types [6].

The second assumption is how functional networks take advantage of mutual exclusivity in somatic mutations. The process of mutual exclusivity in somatic mutations describes how mutations do not occur together if they are in the same genetic pathway [7]. In functional networks, accounting for mutual exclusivity corresponds to the frequency of a mutation, which is the weight of a vertex  $V_i$  in the graph. Theory, however suggests that there are multiple confounding factors that cause mutual exclusivity [1]. For example, the mutual exclusivity of mutations in the TP53 and MUC16 genes are better explained by cancer type and TML in colorectal cancer [1].

The last assumption is of preprocessing somatic mutation data to a limited number of mutations. Although preprocessing is done in numerous studies beyond genetics, there is no gold standard for somatic mutations. Furthermore, due to the high dimensionality of somatic mutation data and the limited number of samples, preprocessing will be biased towards frequently occurring mutations [8]. So while a model can identify novel co-occurring mutations within the functional networks, it may only reflect the model’s preference for a specific paradigm. For example, HotNet2 [3] removes samples with more than 400 somatic mutations, however, overdispersion of a gene and TML are both significant factors that influence the relationship between mutations.

From a machine learning perspective we can describe the problems faced by functional networks as overfitting. This is elucidated by the false positives produced from functional networks. Specifically, functional networks memorize the parameters of the interaction network instead of learning the parameters of the somatic mutation dataset. So while, functional networks may reproduce valid biological processes, they do not necessarily capture cancer-relevant relationships between mutations. For example, functional networks validated targeted treatment drugs erlotinib or gefitinib for mutations in the EGFR gene. Since the EGFR is omnipresent in genetic pathways and as a mutation, functional networks constrain a patient to be mainly influenced by the EGFR mutation. From the previous example, these treatments, however, are only temporarily effective and a relapse often occurs due to the presence of a co-occurring mutation in the same genetic pathway that had equal cancerous potential [9]. So, although functional networks correctly identified the EGFR mutation, it is the interplay between many mutations that influences cancer biology.

In this paper, we propose to exploit the inherent latent structure of a somatic mutation dataset with a generative probabilistic model. Instead of relying on interaction networks *a priori*, we use a prior based on a correlated random measure (CoRM). The CoRM enables the model to specify a notion of similarity on the possible latent distributions of the somatic mutations. In our case, we would like the latent distributions to correspond to two unique characteristics of the somatic mutation data: mutual exclusivity and cancer-related biological processes. Specifically, the CoRM assigns probabilities to particular configurations of latent distributions via a Zero Inflated Negative Binomial Process (ZINB) and enforces mutual exclusivity via a correlation structure through the conjunction of the Beta-Bernoulli Process and neural networks.

The main contributions of this work are as follows. We propose the Correlated Zero Inflated Binomial Process, CoZINB, a novel generative latent variable model with an implicit dependency structure and latent parameters that represent the sparsity of a somatic mutation dataset. Therefore, CoZINB avoids biases from frequency and interaction networks that confound functional network methods. The CoZINB is used to factorize a somatic mutation profile to infer positively and

negatively correlated latent factors<sup>1</sup> that consist of co-occurring mutations and represent biological processes. We create a computationally efficient inference scheme with the confluence of stochastic variational inference and amortized variational inference [10], [11], [12]. Our experiments on the pan cancer dataset from The Cancer Genomic Archive (TCGA) verify the benefits of our model to uncover co-occurring mutations while maintaining rules of mutual exclusivity and cancer specific processes.

## II. BACKGROUND

### A. Negative Binomial Distribution

The dispersion parameter  $r$  of the Negative Binomial (NB) distribution allows models self-excitation [13], an amenable property in biology, such that if a distinct cancer mutation occurs in a tumor, it is likely to occur again in the tumor. Another intuitive way to express the NB distribution is to model it as a Poisson-Gamma mixture distribution. The NB distribution is then a weighted mixture of Poisson distributions, where the rate parameter is an unknown gamma distribution. Thus, an NB distribution is analogous to an overdispersed Poisson distribution. Research in differential expression analysis [14], [15], regression, and clustering of single-cell data [16] has shown superior performance of NB distribution.

### B. Mixed Membership Modeling

---

**Algorithm 1** Latent Dirichlet Allocation a Mixed Membership Model

---

```

for Each cluster  $k : 1, \dots, K$  do
  Sample  $\phi_k \sim \text{Dirichlet}_M(\eta_1, \dots, \eta_m)$ 
end for
for Each sample  $j : 1, \dots, J$  do
  Sample  $\theta_j \sim \text{Dirichlet}_K(\alpha_1, \dots, \alpha_K)$ 
  for Each data point  $x_{j,m}$  in sample  $j$  do
    Sample Cluster  $z_{j,m} \sim \text{Discrete}_K(\theta_j)$ 
    Sample data point  $x_{j,m} \sim \text{Discrete}_M(\phi_{z_{j,m}})$ 
  end for
end for

```

---

Mixed-membership modeling is essentially a soft clustering problem. In contrast to hard clustering where a data point is assigned to one and only one cluster, each data point in a mixed-membership model is associated with a number of clusters. It is the interaction among these clusters that gives rise to an observed data point. In the context of cancer biology, mixed-membership models allow us to remove the restrictive assumption that mutations are mutually exclusive across clusters [17].

The generative process of LDA follows Algorithm 1. What specifically matters, is how choosing the priors: the factor score matrix,  $\theta$  and the factor loading matrix,  $\phi$  for different generative models. For example, if we change the priors in Algorithm

<sup>1</sup>There are a number of ways to describe sets based on the current nomenclature (clusters, factors, and topics). For the remainder of the paper we refer to the latent space of our model as factors.

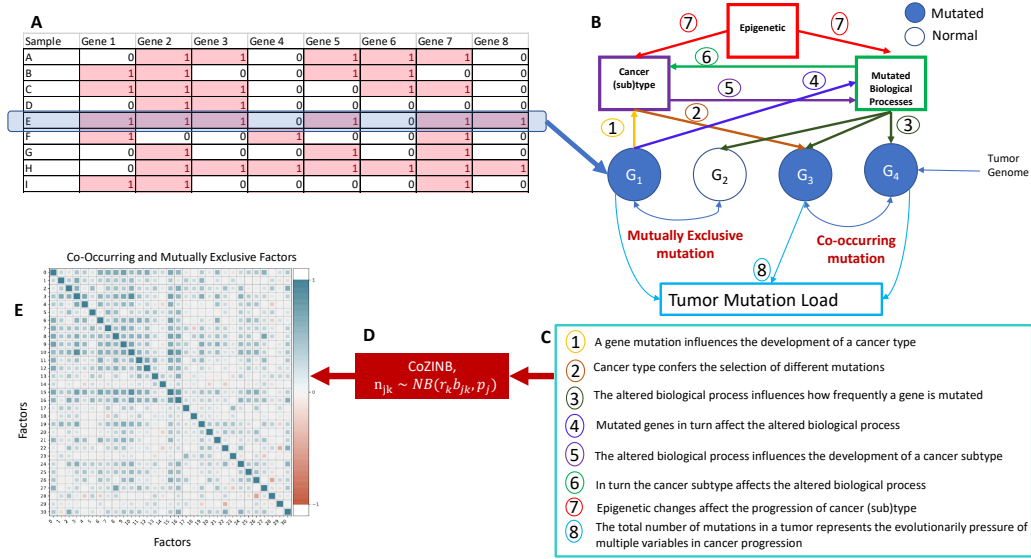


Fig. 1: CoZINB models multiple cofounding variables to determine co-occurring mutations and mutually exclusive mutations. (a) Is an example of a set of binarized somatic mutation profiles represented as a matrix, where 1 indicates that a gene is mutated (b-c) There is a complex relationship between cancer, biological processes, and the development of mutations. [1] (d) These relationships are modeled via a correlated Beta-Bernoulli Process and an independent Gamma process as seen in Equations 2 and 4, where  $n_{jk}$  is the number of times the  $k_{th}$  factor appears in sample  $j$ . (e) CoZINB learns a diverse latent structure of somatic mutation profiles, with positively and negatively correlated factors that contain a set of co-occurring mutations

1 to Gamma distributions and the likelihood to a Poisson distribution we can obtain Poisson Factor Analysis (PFA) [13]. Another extension are hierarchical mixed-membership models, which increase model capability by stacking stochastic processes such that algorithm infers the number possible clusters within a dataset [18].

Though powerful, mixed-membership models are limited as they do not explicitly model the correlations between the mixing proportions of any two clusters. Using science articles as an example, the LDA topic model cannot capture that the presence of a topic on cells is more positively correlated with a topic on cancer in comparison to a topic on astronomy. Recent innovations in joint modeling of correlation and mixture proportions has resulted in improved prediction and interpretability of the topics as seen in [19] and an hierarchical extensions in [10], [12], [11]. Continuing in the context of topic modeling, these methods inject a Gaussian covariance into the  $\theta$  such that the factor score matrix takes into account the correlation between topics. Our model follows this format, where the bottom stochastic process is based on correlated random measures.

### C. Latent Space as a Correlated Random Measure

Homogenous completely random measure (CRM) [20] is built upon a mean measure (levy measure) as  $v(da, dw) = H(da)v_w(dw)$  where  $H(da)$  is the base measure and  $v_w(dw)$  is the rate measure. As an example, the Gamma Process (GP) has a mean measure of  $v(da, dw) = H(da)e^{-cw}/wdw$ . Intuitively, we can think of the base measure as defining the

existence of factors and the rate measure assigning a weight to each factor. This definition allows us to have an infinite number of factors, but a finite mass, and is a foundation of many of the non-parametric Bayesian statistical models.

A correlated random measure (CoRM) is created by augmenting a CRM into a higher space to include locations such that the mean measure is now  $v(da, dw, dl) = H(da)v_w(dw)v_l(dl)$ , where  $v_l(dl)$  is a vector of locations in  $\mathbb{R}^d$ . We can then draw correlated weights  $x$  via a transformation function as  $x \sim T(\cdot|w, F(l))$ , given the uncorrelated weights,  $w$ , and a random kernel function on the locations  $F(l)$ . Based on Proposition 1 from [11] the transformed levy measure is now  $\tilde{v} = v(da, dw, dl)p(dx|F(l), w)$ . The transformation distribution, like a homogeneous CRM, is restricted to have finite mass to guarantee generation of useful statistical models.

## III. MODEL

### A. Correlated Zero Inflated Negative Binomial Process

The Correlated Zero-Inflated Negative Binomial Process (CoZINB) is a correlated hierarchical Bayesian model for learning the latent representations of somatic mutation profiles. We take advantage of neural networks for the random function  $F(l)$  to model non-linear correlations of somatic mutations. The transformation distribution based on the Beta-Bernoulli Process (BeBP) is then used to enforce sparsity of the globally shared latent factors. We describe this as a draw from the CoZINB,  $X_j \sim \text{NB}(RB_j, p)$  with a mean measure  $\tilde{v}$ :

$$H(da)v_l(dl)\alpha w^{-1}(1-w)^{w-1}p(dx|F(l), w)R_0(dr)dw$$

so  $B_j$  is a draw from the transformation distribution  $p(dx|F(l), w)$  and  $R$  is an independent  $\Gamma P$ . By deriving a latent representation with an unsupervised correlation structure, we use the inherent information shared across the somatic mutation profiles to ensure we capture mutual exclusivity between factors.

As shown in Figure 1, by using CoZINB we can model 1) the excessive amount of mutations that do not occur in a patient, 2) non-linear interactions of somatic mutations 3) implicit cancer subtype based on the factors 4) total number of mutations in a somatic mutation profile. We will show in our experiments the latent factors follow a pattern of co-occurrence that is also observed in cancer biology.

### B. Generative Structure

We represent our dataset as a bag of words commonly seen in text corpus,  $X \in \mathbb{N}^{J \times M}$ , where  $J$  and  $M$  are the number of patients and distinct number of mutations, respectively. As in many factor models we also introduce an additional latent variable,  $z_{jm}$ , indicating the factor assignment for mutation  $m$ . The mutation profile of each patient is realized as a mixture of latent factors ( $\phi_{z_{jm}}$ ) shared by all patients. Specifically, we model a distinct mutation count in a patient as:

$$n_{jmk} = \sum_{i=1}^{N_j} \delta(z_{ji} = k, m_{ji} = m) \quad (1)$$

$$n_{jk} \sim NB(r_k b_{jk}, p_j).$$

From a biological perspective  $n_{jk}$  is the number of times a  $k$  biological process occurs in patient  $j$ . The count  $n_{jmk}$  corresponds to the contribution of underlying biological process  $k$  to the count of mutation  $m$  in sample  $j$ . The total number of mutations in patient  $j$  is then  $N_j$ . The shape parameter  $r_k$  captures the popularity of the biological process  $k$  across all patients. The probability parameter  $p_j$  models the patient specific somatic mutation profile. Specifically,  $p_j$  accounts for heterogeneity among the patient population. To enable sparsity and correlations within the latent factors, we introduce Bernoulli latent variables  $b_{jk}$ . When  $b_{jk} = 1$ , the latent factor  $k$  is present in patient  $j$ , otherwise irrelevant. The correlations within the binary random variables are produced via a transformation of the BeBP as discussed in II-C. A transformed distribution follows [11] as

$$b_{jk} \sim \text{Bernoulli}(\sigma(\sigma^{-1}(\pi_k) + F(\cdot))) \quad (2)$$

$$\pi_k \sim \text{Beta}(\alpha/K, \alpha(1 - 1/K))$$

where  $\sigma$  is the sigmoid function. As a result, the proposed model assigns positive probability only to a subset of factors, based on the correlation structured created in  $F(\cdot)$ .

Inspired by Variational Autoencoders (VAE) [21] to model nonlinear correlations we use a deep neural network to create a kernel  $F(\cdot)$  is:

$$f(h_j, l_k) \sim N(u_f(h_j, l_k), \sigma_f^2(h_j, l_k)) \quad (3)$$

$$h_j \sim N(0, aI), l_k \sim N(0, bI).$$

where  $h_j$  is a patient specific vector as the output of an inference network of a VAE. The locations,  $l_k$ , are then concatenated with  $h_j$  as in input the decoder of a VAE to create the kernel. Unlike the standard VAE decoder which aims to recreate the input of a encoder, the decoder in our model supplies prior information to  $F(\cdot)$ . This paradigm of generating the kernel is similar to Empirical Bayes [22], where we use the structure of the data as a prior to generate  $F(\cdot)$ .

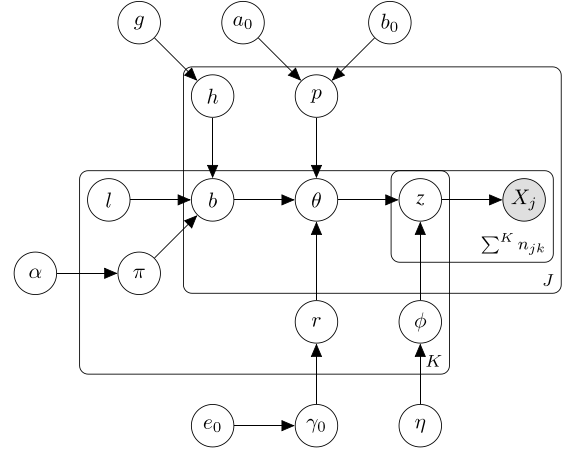


Fig. 2: Correlated Zero Inflated Negative Binomial Process Topic Model

Building upon the above equations, the full generative structure of the CoZINB topic model follows the paradigm of the Gamma-Poisson construction of a NB process:

$$x_{jm} \sim \text{Discrete}(\phi_{z_{jm}}), z_{jm,k} \sim \phi_{mk} \theta_{jk}$$

$$\phi_k \sim \text{Dirichlet}(\eta_0 \mathbf{1}_M), N_j = \sum_{k=1}^K n_{jk},$$

$$n_{jk} \sim \text{Poisson}(\theta_{jk}) \theta_{jk} \sim \text{Gamma}(r_k b_{jk}, p_j), \quad (4)$$

$$r_k \sim \text{Gamma}(\gamma_0, 1/\alpha), \gamma_0 \sim \text{Gamma}(e_0, 1/f_0),$$

$$p_j \sim \text{Beta}(a_0, b_0), L_{jk} \sim \text{CRT}(n_{jk}, r_k b_{jk}),$$

$$L'_k \sim \text{CRT}\left(\sum_{j=1}^J L_{jk}, \gamma_0\right)$$

and  $b_{jk}$  follows Equation 2. To aid in inference of  $r_k$ , we introduce a data augmentation latent variable,  $L_{jk}$  based on the Chinese Restaurant Table as in [13]. The full model plate is shown in Figure 2.

### C. Inference

Given the set of observed mutation profiles,  $X$  and their corresponding mutation counts  $m$ , our goal is to infer the factors  $z_{jm}$ , the factor score matrix,  $\theta_{jk}$ , the factor loading matrix,  $\phi_{mk}$ , and the factor locations  $l_k$ . For a biological analogue,  $\theta_{jk}$  is the proportion of the biological process  $k$  in sample  $j$ ,  $\phi_{mk}$  is the proportion of each mutation in factor  $k$ , and  $l_k$  is correlation between a set of co-occurring mutations. Denote  $n_{jk} = \sum_{m=1}^M \delta(z_{jm} = k)$ , where  $K$  is some finite, but large,

truncation level, we posit the fully factorize variational scheme:

$$q(\theta, \phi, \mathbf{p}i, Z, R, P, l, L, h) = \prod_{k=1}^K q(l_k)q(r_k)q(\pi_k) \\ q(\phi_k)q(\gamma_0) \prod_{j=1}^J \left[ q(h_j|X)q(p_j) \prod_{k=1}^K q(b_{jk})q(\theta_{jk})q(L_{jk}) \right. \\ \left. q(L'_k) \prod_{m=1}^{N_j} q(z_{jm}) \right] \quad (5)$$

The variational distribution for each latent variable is associated with it's own variational parameter(s) as follows

$$q(l_k) = \delta_{\tilde{l}_k}, q(\gamma_0) = \delta_{\tilde{\gamma}_0}, \\ q(\phi_k) = \text{Dirichlet}(\tilde{\eta}) \\ q(\pi_k) = \text{Beta}(\tau_{k1}, \tau_{k2}) \\ q(r_k) = \text{Gamma}(\tilde{r}_{k1}, \tilde{r}_{k2}) \\ q(p_j) = \text{Beta}(\tilde{a}_j, \tilde{b}_j) \\ q(b_{jk}) = \text{Bernoulli}(\nu_{jk}) \\ q(\theta_{jk}) = \text{Gamma}(\tilde{\theta}_{jk1}, \tilde{\theta}_{jk2}) \\ q(L_{jk}) = \delta_{\tilde{L}_{jk}} \\ q(L'_k) = \delta_{\tilde{L}'_k} \\ q(z_{jm}) = \text{Multinomial}(\psi_{jm})$$

We let  $q(h_j|X_j) = \delta_{g(X_j)}$  where  $g$  is the inference network of a VAE.

To update the variational parameters we can use stochastic variational inference (SVI) that assumes we subsample  $J_t \in J$  patients at every iteration  $t$  and optimize the noisy variational objective:

$$L_t = \mathbb{E}[\ln p(l, r, \phi, \gamma_0, \pi)] + \\ \frac{J}{J_t} \sum_{j \in J_t} \mathbb{E}[\ln h_j, \theta_j, z_j, p_j, b_j, L_j, X_j] + H[q(\phi)] + \\ \frac{J}{J_t} \sum_{j \in J_t} H[q(\theta_j, z_j, b_j, p_j)] \quad (6)$$

At each iteration  $t$ , we update the local variables,  $z, \theta, b$ , and,  $p$  using closed form equations, while the remaining variables, with the exception of  $\phi$  are updated via stochastic gradient descent via a decreasing step size [23]. The complete conditional updates are summarized in Table 1. <sup>2</sup>

#### IV. EXPERIMENTS

To show the importance of modeling sparsity we consider two latent variable models as comparisons:

**PRME** is a hierarchical mixed membership model based on the Dirichlet Process and the correlations are induced through neural networks with a Gaussian kernel, so  $\theta_{jk} \sim \text{Gamma}(r_k, f(h_j, l_k))$ . Where  $r_k$  is a stick-breaking Dirichlet Process instead of a Gamma Process as in CoZINB.

<sup>2</sup>Note that  $\mathbb{E}[\pi_{k2}] = E[\pi_{k1}](1 - p_j)^{\mathbb{E}[r_k]}$ .

**Prod-LDA** [24] is an extension of LDA to deep generative models. It replaces the Gaussian Prior in a VAE with a Dirichlet distributed prior to better learn the latent structure of text data.

To compare CoZINB's latent factors with functional networks, we use the HotNet2 model.

**Hotnet2** is used in a discriminate approach to compare functional networks with CoZINB. We train two distinct linear support vector machine (SVM) classifiers [25] with the subnetworks learned by Hotnet2 [3] and the factors from CoZINB as features. The SVMs are used to predict the cancer type of a patient given the features. Additionally we compare the correlation between functional networks and biological process and the correlation between factors from CoZINB and biological processes

Architecture encoder and decoder follow simple multilayer perceptrons (MLP) with dimensions shown in Table II. The hyperparameters for PRME and CoZINB are set as  $a = 1, b = 1, \alpha = 1, \eta = 0.2, a_0 = .001, b_0 = .001, e_0 = .001, f_0 = .001, M = 21332$ , and a truncation level  $K$  as 100 for all models. All gradient updates are done via Adam [26] with a learning rate of  $1e - 3$ .

We use the pan cancer dataset from the TCGA consisting of 10296 tumor samples ( $J$ ) and 21332 distinct mutations ( $M$ ). For training the models we select 70% of the dataset, setting 20% for validation of parameters, and 10% for testing. The only pre-processing we do is remove 'abParts,' which is not represented in any database.

#### A. Comparison Metrics

##### Per-heldout-word perplexity

$$\exp\left(\frac{1}{|X_{Test}|} \sum_{m \in X_{Test}} \ln p(X_{jm}|X_{j,Train})\right)$$

is used to measure the utilization of the latent dimensions of the two other latent variable models.

**Precision@1** measures the influence of the factors on determining the correct count of one distinct (@1) held-out mutation as TML increases. This is calculated by

$$\frac{1}{N} \sum_1^N \delta(X_{jm}^{\sim} = X_{jm}),$$

where  $X_{jm}^{\sim}$  is the predict count of mutation  $m$ .

**Specificity** is used to compare the discriminative power of the functional networks and the latent factors of CoZINB in predicting the correct cancer type Lung Adenocarcinoma (LUAD), Colon Cancer (COAD), Ovarian Cancer (READ), Stomach Cancer, and Breast Cancer (BRCA).

**Spearman-rho** is used to assess how well the factors produced from CoZINB match with biological processes.

**Qualitatively** we analyze the biological implications of the factors obtained by CoZINB for a subset of the TCGA dataset in LUAD and COAD. We chose these subsets as there are an established number of studies identifying co-occurring and mutual exclusive somatic mutations within these cancers.

Latent Variable	Update Type	Variational Update	Variational Parameter
$\phi_k$	Closed	$\eta + \sum_{J \in J_k} \sum_{n=1}^{N_j} \psi_{jm}(k) * \mathbf{1}(X_{jm} = \tilde{m})$	$\eta_{k\tilde{m}}$
$\pi_k$	Closed	$\alpha/K + \sum^J \nu_{jk}$	$\tau_{k1}$
$\pi_k$	Closed	$\alpha(1 - 1/K) + J - \sum^J \nu_{jk}$	$\tau_{k2}$
$r_k$	Closed	$\gamma_0 + \sum^J \mathbb{E}[L_{jk}]$	$r_{k1}$
$r_k$	Closed	$\alpha - \sum^J \mathbb{E}[b_{jk}] \ln 1 - p_j$	$r_{k2}$
$p_j$	Closed	$a_0 + N_j$	$\tilde{a}$
$p_j$	Closed	$b_0 + \sum^K \mathbb{E}[r_k] \mathbb{E}[b_{jk}]$	$\tilde{b}$
$b_{jk}$	Closed	$\mathbb{E}[\pi_{k2}] (\mathbb{E}[\pi_{k2}] + 1 - \mathbb{E}[\pi_{k1}])^{-1}$	$\nu_{jk}$
$\theta_{jk}$	Closed	$\mathbb{E}[r_k] \mathbb{E}[b_{jk}] + \sum^M \psi_{jm}(k)$	$\theta_{jk1}$
$\theta_{jk}$	Closed	$\mathbb{E}[p_j]$	$\theta_{jk2}$
$z_{jm}$	Closed	$\log \theta_{jk} + \log \phi_{km}$	$\psi_{jm}$
$\gamma_0$	Gradient Ascent	$\sum^K \nabla_{\gamma_0} \mathbb{E}[p(\gamma_0)] + \sum^J \sum^K \nabla_{\gamma_0} \mathbb{E}[p(L'_k   L_{jk}, \gamma_0)]$	$\delta_{\gamma_0}$
$l_k$	Gradient Ascent	$\sum^K \nabla_l \mathbb{E}[\ln p(l_k)] + \sum^J \sum^K \nabla_l \mathbb{E}[p(b_{jk}   \pi_k, h_n, l_k)]$	$\delta_{l_k}$
$h_n$	Gradient Ascent	$\sum^J \nabla_h \mathbb{E}[\ln p(h_j)] + \sum^J \sum^K \nabla_h \mathbb{E}[p(b_{jk}   \pi_k, h_n, l_k)]$	$\delta_{h_j}$
$L_{jk}$	Gradient Ascent	$\sum^J \sum^K \nabla_L \mathbb{E}[p(L_{jk}   n_{jk}, b_{jk}, r_k)]$	$\delta_{L_{jk}}$

TABLE I: A summary of the update information for the Variational Parameters.

## V. RESULTS

### A. Quantitative

	Perplexity@H	Encoder Sizes	Decoder Sizes
Prod-LDA	@0 1986.30		
	@100 1091.33	$[M \times 100]$	$[100 \times 100]$
	@200 1297.19	$[100 \times 100]$	$[100 \times 100]$
	@400 1691.76	$[100 \times 100]$	$[100 \times M]$
PRME	@0 1151.96	$[M \times 1000]$	$[40 \times 80]$
	@100 721.42	$[1000 \times 1000]$	$[80 \times 80]$
	@200 997.19	$[1000 \times 1000]$	$[80 \times 80]$
	@400 1091.76	$[1000 \times 40]$	$[80 \times 2]$
CoZINB	@0 779.69	$[M \times 1000]$	$[40 \times 80]$
	@100 721.77	$[1000 \times 1000]$	$[80 \times 80]$
	@200 699.35	$[1000 \times 1000]$	$[80 \times 80]$
	@400 709.34	$[1000 \times 40]$	$[80 \times 2]$

TABLE II: Perplexity of the held-out test Pan Cancer dataset and architecture of Prod-LDA, PRME, and CoZINB. The @H indicates if the models were trained with only the top frequently occurring mutations i.e., @100 is a training set with only top 100 frequently mutations, @0 = no mutations were held out

As Table II shows, CoZINB performs better than PRME and Prod-LDA when there are a excessive number of zeros in the dataset. Specifically, we show the importance of using a prior that can incorporate sparsity, by only using the top occurring mutations at levels of 100, 200, and 400 during. As we limit the amount of mutations the models need to learn, the perplexity of all models begin to improve. The sparsity of mutations, therefore poses a significant challenge in the learning paradigm. CoZINB achieved the best performance due to the use of selector variables  $b_{jk}$ . Specifically the selector variables assigns a finite number of latent factors for each sample and an explicit zero mass to the rest. Any random signals then will receive a zero probability, whereas the baseline models will assign some small probability to the latent factor resulting in a poor local optima.

In Figure 3 we observe the influence of the CoZINB shape and probability parameters. The PRME limits the distribution of the factor score matrix,  $\theta_{jk}$ , to the frequency of latent factors. In comparison, CoZINB models the count data as the frequency

of the latent factors in  $r_k$  and the probability of a count of somatic mutations in  $p_j$ . More precisely, the expected count of mutation  $m$  in sample  $j$  is proportional to  $\frac{p_j}{1-p_j}$  and the factor score matrix.




Fig. 3: Precision@1 for CoZINB is higher than PRME as it takes account into the sparsity of the dataset through the Negative-Binomial process

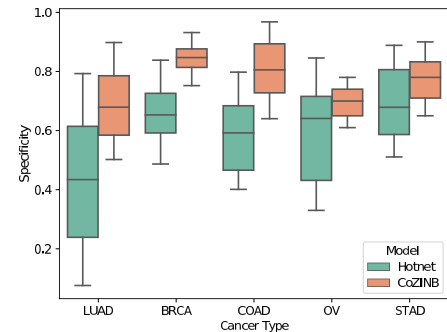


Fig. 4: Number of times a factor occurs for all the patients with Lung Adenocarcinoma (LUAD), Colon Cancer (COAD), OV (Ovarian Cancer), STAD (Stomach Cancer), and Breast Cancer (BRCA).




Fig. 5: Number of times a factor occurs for all the patients with Lung Adenocarcinoma (LUAD), Colon Cancer (COAD), Renal Cancer (READ), and Breast Cancer (BRCA).

Our assertion that functional networks via Hotnet2 are not specific to cancer types and overfit to the structure of interaction networks is shown in Figure 4. There is a significant variance in predicting cancer type when using the functional networks, especially for LUAD. The poor performance for HotNet2 in LUAD is likely due to the similarities of mutation profiles of LUAD and LUSC cancers, where LUSC generally has a higher TML. In contrast, CoZINB factors preserved cancer type information, especially in BRCA where it achieves an average specificity of 0.83. This suggests that interaction networks, while valuable, have much less influence in determining the relationship between mutations.

In Figure 5 we show the interplay of latent factors in LUAD, COAD, READ, and BRCA. Each bar represents how many times a latent factor occurs in a tissue. We can see there are distinct factors that only correspond to specific cancers. We use the spearman-rho coefficient to assess if the latent factors 11, 12, 39, 40, 47, and 52 from BRCA represent a biological process. We compare it against the spearman-rho coefficient between functional networks of BRCA and biological processes. Using the existing database provided by KEGG, CoZINB has a  $\rho$  of 0.67 for the latent factors of BRCA, while HotNet2 is significantly more correlated with a  $\rho$  of 0.78. This not unexpected since HotNet2 is built on the prior information from interaction networks.

A key insight from Figures 3, 4, and 5 is that we observe CoZINB's robust grouping of frequently occurring mutations into a few factors. The frequently occurring mutations incur a large penalty due to the ZINB process and the optimization procedure prevents the probability mass of frequently occurring mutations to spread across factors. Thereby, the remaining factors are more diverse and allow for better discriminative power between TML and cancer type. We see this again in the case studies below, where frequently occurring mutations are grouped into Factors 1-4.

### B. Lung Adenocarcinoma Case Study

A well studied mutual exclusive set of somatic mutations in LUAD are KRAS and EGFR. They occur in a significant fraction of LUAD patients, but are rarely if-ever observed together in the same tumor. Since these two genes are in the same pathway and activate similar downstream targets, it is

generally assumed that there is no selective pressure to favor cells with both mutations over cells with a mutation in one of them. This behavior is reproduced in CoZINB, where in the most common factors, KRAS and EGFR are at the opposite spectrums, with a mean distance of 7185.

The factors that are more unique to LUAD, 5 in our model are identified by factors 27 with top 5 mutations as: [DNAH5 PCLO ANK3 TTN TP53BP2] and 37 which include the mutations: [ USH2A TTN TSHZ2 DNAH3 MUC16]. To confirm this computationally, we use a simple Random Forest (RF) for classification of cancer type across all patients using the counts the top performing factors as features. With a significance of  $p < 0.05$  RF gives the ranks factors 27 and 37 as the most important features for LUAD, thereby, verifying our model can also implicitly determine patient cancer type.

Figure 6 shows the unique factors in every patient with LUAD as TML increases. Specifically LUAD with high TML will have some form of the background somatic mutations captured in factors 0 to 6, however as, TML begins to increase, occurrences of factors that contain mutually exclusive mutations will appear. Examining the top underlying somatic mutations of Factor 30 (MUC16, FLG, LRP1B, ZFH4 CSMD3, RYR) and Factor 1 (PIK3CA, PTEN, TP53, MUC16, MAP3K1.), we observe mutual exclusivity of the MUC15 and TP53 genes as discussed in section I.

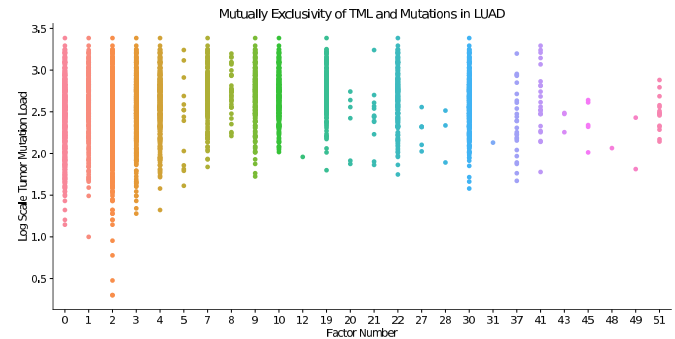


Fig. 6: A pattern of mutual exclusivity in LUAD. The X-axis represents a factor, while the Y-axis the total number of mutations in a tumor. As tumor mutation load increases, more unique factors appear, specifically ones that are mutually exclusive with certain somatic mutations.

### C. Colon Cancer Case Study

Figure 5 shows Factor 6 is unique to the 408 COAD tumor samples within the dataset, with top 5 mutations: [PMVK, BAZ2B, STARD3, SESTD1, KLHL4]. Using the same methodology as before, a RF classifier ranks factors 6 and 41 as the most significant ( $p < .05$ ) features for determining if a patient has Colon Cancer. Moreover, these specific mutations are also known to commonly occur together as passenger mutations [27].

COAD also has specific factors associated with increasing TML in Factors 8 and 21 ( $p < .01$ ) with mutations: [MUC16, DNAH5, TTN, PCLO, ANK3] and [BRCA2, CTNNAL1,

DDX52, NAH10, AK9]. Of note are the occurrences of MUC16, PCLO, and BRCA2, which are all hypermutated in colon cancers. Similar to LUAD we can see Factor 8 is mutually exclusive with factors that include the TP53 mutation with increasing TML.




Fig. 7: A pattern of mutual exclusivity in COAD. The X-axis represents a factor, while the Y-axis the total number of mutations in a tumor. As tumor mutation load increases, more unique factors appear, specifically ones that are mutually exclusive with certain somatic mutations.

## VI. DISCUSSION AND FUTURE WORK

The CoZINB, unlike existing methods for assessing co-occurring somatic mutations, is unsupervised and infers a latent structure in a sparse and complex dataset. A major concept shown in our model is that we can probabilistically capture mutual exclusivity between somatic mutations as a non-linear transformation of multiple latent variables. We also show that to correctly represent cancer biology and patterns of mutual exclusivity and co-occurrence in somatic mutations, a model needs to incorporate tumor mutational load, cancer type, and non-linear mutation correlations as confounding variables.

Although, we argued against the use of interaction networks, they might have a place in a semi-supervised framework. For example, it would be useful to understand the causal relationship of somatic mutations and unaltered genes, analogous to link-prediction in community network detection. This learning paradigm, however, is a complex combinatorial problem, considering the possible exponential search space of somatic mutation interactions.

## ACKNOWLEDGMENTS

We thank Aanon Zhang for discussing his work, Population Random Measure Embedding. We are also grateful for the support of NVIDIA Corporation with the donation of the Titan X Pascal GPU used for this research.

## REFERENCES

- [1] J. van de Haar, S. Canisius, K. Y. Michael, E. E. Voest, L. F. Wessels, and T. Ideker, "Identifying epistasis in cancer genomes: a delicate affair," *Cell*, vol. 177, no. 6, pp. 1375–1383, 2019.
- [2] B. Bauer, R. Siebert, and A. Traulsen, "Cancer initiation with epistatic interactions between driver and passenger mutations," *Journal of Theoretical Biology*, vol. 358, pp. 52–60, 2014.

- [3] M. D. Leiserson, F. Vandin, H.-T. Wu, J. R. Dobson, J. V. Eldridge, J. L. Thomas, A. Papoutsaki, Y. Kim, B. Niu, M. McLellan *et al.*, "Pan-cancer network analysis identifies combinations of rare somatic mutations across pathways and protein complexes," *Nature genetics*, vol. 47, no. 2, p. 106, 2015.
- [4] J. Zhang and S. Zhang, "The discovery of mutated driver pathways in cancer: Models and algorithms," *IEEE/ACM transactions on computational biology and bioinformatics*, vol. 15, no. 3, pp. 988–998, 2016.
- [5] M. Kanehisa and S. Goto, "Kegg: kyoto encyclopedia of genes and genomes," *Nucleic acids research*, vol. 28, no. 1, pp. 27–30, 2000.
- [6] L. Cowen, T. Ideker, B. J. Raphael, and R. Sharan, "Network propagation: a universal amplifier of genetic associations," *Nature Reviews Genetics*, vol. 18, no. 9, p. 551, 2017.
- [7] C.-H. Yeang, F. McCormick, and A. Levine, "Combinatorial patterns of somatic gene mutations in cancer," *The FASEB journal*, vol. 22, no. 8, pp. 2605–2622, 2008.
- [8] H. Yin and K. Gai, "An empirical study on preprocessing high-dimensional class-imbalanced data for classification," in *2015 IEEE 17th International Conference on High Performance Computing and Communications, 2015 IEEE 7th International Symposium on CyberSpace Safety and Security, and 2015 IEEE 12th International Conference on Embedded Software and Systems*, 2015, pp. 1314–1319.
- [9] F. Morgillo, C. M. Della Corte, M. Fasano, and F. Ciardiello, "Mechanisms of resistance to egfr-targeted drugs: lung cancer," *ESMO open*, vol. 1, no. 3, p. e000060, 2016.
- [10] A. Zhang and J. Paisley, "Random function priors for correlation modeling," *arXiv preprint arXiv:1905.03826*, 2019.
- [11] R. Ranganath and D. M. Blei, "Correlated random measures," *Journal of the American Statistical Association*, vol. 113, no. 521, pp. 417–430, 2018.
- [12] J. Paisley, C. Wang, and D. Blei, "The discrete infinite logistic normal distribution for mixed-membership modeling," in *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, 2011, pp. 74–82.
- [13] M. Zhou and L. Carin, "Negative binomial process count and mixture modeling," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 37, no. 2, pp. 307–320, 2013.
- [14] S. Anders and W. Huber, "Differential expression analysis for sequence count data," *Nature Precedings*, pp. 1–1, 2010.
- [15] S. Z. Dadaneh, X. Qian, and M. Zhou, "Bnp-seq: Bayesian nonparametric differential expression analysis of sequencing count data," *Journal of the American Statistical Association*, vol. 113, no. 521, pp. 81–94, 2018.
- [16] R. Lopez, J. Regier, M. B. Cole, M. I. Jordan, and N. Yosef, "Deep generative modeling for single-cell transcriptomics," *Nature methods*, vol. 15, no. 12, pp. 1053–1058, 2018.
- [17] B. Liu, L. Liu, A. Tsykin, G. J. Goodall, J. E. Green, M. Zhu, C. H. Kim, and J. Li, "Identifying functional mirna–mRNA regulatory modules with correspondence latent dirichlet allocation," *Bioinformatics*, vol. 26, no. 24, pp. 3105–3111, 2010.
- [18] Y. W. Teh and M. I. Jordan, "Hierarchical bayesian nonparametric models with applications," *Bayesian nonparametrics*, vol. 1, pp. 158–207, 2010.
- [19] J. D. Lafferty and D. M. Blei, "Correlated topic models," in *Advances in neural information processing systems*, 2006, pp. 147–154.
- [20] J. Kingman, "Completely random measures," *Pacific Journal of Mathematics*, vol. 21, no. 1, pp. 59–78, 1967.
- [21] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *arXiv preprint arXiv:1312.6114*, 2013.
- [22] B. P. Carlin and T. A. Louis, *Bayes and empirical Bayes methods for data analysis*. Chapman and Hall/CRC, 2010.
- [23] M. D. Hoffman, D. M. Blei, C. Wang, and J. Paisley, "Stochastic variational inference," *The Journal of Machine Learning Research*, vol. 14, no. 1, pp. 1303–1347, 2013.
- [24] A. Srivastava and C. Sutton, "Autoencoding variational inference for topic models," *arXiv preprint arXiv:1703.01488*, 2017.
- [25] M. A. Hearst, S. T. Dumais, E. Osuna, J. Platt, and B. Scholkopf, "Support vector machines," *IEEE Intelligent Systems and their applications*, vol. 13, no. 4, pp. 18–28, 1998.
- [26] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [27] E. Cerami, J. Gao, U. Dogrusoz, B. E. Gross, S. O. Sumer, B. A. Aksoy, A. Jacobsen, C. J. Byrne, M. L. Heuer, E. Larsson *et al.*, "The cbio cancer genomics portal: an open platform for exploring multidimensional cancer genomics data," 2012.