

Compressive Bayesian non-negative matrix factorization for mutational signatures analysis

Alessandro Zito¹ and Jeffrey W. Miller¹

¹Department of Biostatistics, Harvard University, Boston, MA, 02115, U.S.A.

Abstract

Non-negative matrix factorization (NMF) is widely used in many applications for dimensionality reduction. Inferring an appropriate number of factors for NMF is a challenging problem, and several approaches based on information criteria or sparsity-inducing priors have been proposed. However, inference in these models is often complicated and computationally challenging. In this paper, we introduce a novel methodology for overfitted Bayesian NMF models using “compressive hyperpriors” that force unneeded factors down to negligible values while only imposing mild shrinkage on needed factors. The method is based on using simple semi-conjugate priors to facilitate inference, while setting the strength of the hyperprior in a data-dependent way to achieve this compressive property. We apply our method to mutational signatures analysis in cancer genomics, where we find that it outperforms state-of-the-art alternatives. In particular, we illustrate how our compressive hyperprior enables the use of biologically informed priors on the signatures, yielding significantly improved accuracy. We provide theoretical results establishing the compressive property, and we demonstrate the method in simulations and on real data from a breast cancer application.

1 INTRODUCTION

Non-negative matrix factorization (NMF) is a dimensionality reduction technique that decomposes a non-negative matrix into the product of two lower-dimensional non-negative matrices of a desired rank by minimizing a given loss function, such as the squared-error loss or Kullback–Leibler divergence (Lee and Seung, 2000). Several Bayesian NMF methods have been studied, using both parametric (Cemgil, 2009; Zhou, 2018) and nonparametric models (Hoffman et al., 2010; Gopalan et al., 2014; Ayed and Caron, 2021). See the general overview of Wang and Zhang (2013), and refer to Zhou and Carin (2015) for a detailed summary of Bayesian factorization methods for discrete outcomes.

In cancer genomics, non-negative matrix factorization has been successfully used to discover a wide range of mutational signatures corresponding to distinct processes, such as damaged DNA repair

mechanisms and environmental mutagens like tobacco smoking or metabolic byproducts (Alexandrov et al., 2013). These signatures, defined as vectors of the frequencies with which different types of point mutations occur, are inferred from mutation counts in whole-genome or whole-exome sequencing using NMF algorithms (Nik-Zainal et al., 2012, 2016; Alexandrov et al., 2013, 2020). Identifying these patterns in the DNA of cancer patients is a significant advance toward understanding the etiology of cancer (Koh et al., 2021), improving the effectiveness of precision therapies (Aguirre et al., 2018; Gulhan et al., 2019), and characterizing the evolution of the disease over time (Rubanova et al., 2020).

A number of NMF-based methods have been proposed for mutational signatures analysis; see Islam et al. (2022) and references therein. However, a difficult aspect of NMF is choosing an appropriate number of latent factors, which corresponds to the number of mutational signatures present in the data. Choosing too many factors can lead to the discovery of spurious signatures, while choosing too few factors can lead to incorrectly merging distinct signatures. Existing techniques for selecting the number of factors include the Bayesian information criterion (Rosales et al., 2016; Fischer et al., 2013), cross-validation (Lal et al., 2021), or even neural networks (Nebgen et al., 2021; Islam et al., 2022). The disadvantage of such approaches is that they require estimating a separate model for each choice of rank or regularization parameter, making them computationally intensive. Another popular approach is to use automatic relevance determination (Tan and Févotte, 2013; Kim et al., 2016), which is extremely fast but only provides point estimates without any uncertainty quantification. Bayesian nonparametric models with sparsity-inducing priors have been used to address this problem in Gaussian factor models (Bhattacharya and Dunson, 2011; Legramanti et al., 2020), however, these have not yet been applied to NMF, to our knowledge.

In this article, we introduce a novel Bayesian NMF method that yields accurate and reliable inference in a computationally simple way by using *compressive hyperpriors* to drive the weights of unneeded factors to zero. Specifically, we use a Poisson factorization model with semi-conjugate Dirichlet and gamma priors over the signatures and the loadings, respectively, and we induce sparsity by using a shrinkage hyperprior that strengthens with the amount of data in such a way that (a) unneeded factors are given negligible weight with probability tending to one, while (b) needed factors are given weights that are only mildly shrunk. This compressive property enables the method to select an appropriate number of latent factors in a continuous way, without the need to fit multiple models or discretely jump between models in Markov chain Monte Carlo samplers. Further, posterior inference can be carried out using a simple auxiliary variable Gibbs sampling algorithm as in Dunson and Herring (2005) and Cemgil (2009), making the method easy to implement. A crucial benefit of our use of semi-conjugate continuous priors and hyperpriors is that they yield a posterior surface that is easy to explore, in contrast with spike-and-slab constructions that could lead to stronger multimodality.

Additionally, the compressive hyperprior makes it straightforward to take advantage of prior information about the latent factors. This is especially useful for mutational signatures analysis, since the *Catalog Of Somatic Mutations In Cancer* (COSMIC) database provides a curated set of mutational signatures and their putative etiologies (Alexandrov et al., 2020). By using an informative prior centered at the COSMIC signatures, we find that the model obtains improved sensitivity to detect the presence of signatures and infers unambiguous matches to the original database.

The paper is organized as follows. Section 2 defines the model, introduces our compressive hyperprior, and provides a Gibbs sampler algorithm for posterior inference. In Section 3, we provide theoretical results establishing the compressive property of the model. Section 4 provides background on mutational signatures analysis. Section 5 contains a simulation study comparing to leading methods, and Section 6 presents an application to mutational signatures analysis using a benchmark breast cancer dataset. We conclude with a brief discussion in Section 7.

2 METHODOLOGY

2.1 POISSON NON-NEGATIVE MATRIX FACTORIZATION MODEL

We describe the model in the context of mutational signatures, our primary application of interest. Let X_{ij} represent the number of mutations for channel i in sample j , where $i = 1, \dots, I$ and $j = 1, \dots, J$, and let $X \in \mathbb{R}^{I \times J}$ denote the matrix with entries X_{ij} . Typically, the channels would consist of the 96 single-base substitution (SBS) types; see Sections 4 and S4 for details. Non-negative matrix factorization (NMF) consists of finding two non-negative matrices $R \in \mathbb{R}_+^{I \times K}$ and $\Theta \in \mathbb{R}_+^{K \times J}$ such that $X \approx R\Theta$, with the *rank* K typically chosen so that $K \leq \min\{I, J\}$. The k th column of R , denoted $r_k = (r_{1k}, \dots, r_{Ik})$, is referred to as the k th *mutational signature*. The k th row of Θ , denoted $\theta_k = (\theta_{k1}, \dots, \theta_{kJ})$, is the vector of weights representing the *loading* of signature k in each of the J samples, sometimes referred as the *signature activity*.

From a probabilistic perspective, it is natural to model the mutation counts as

$$X_{ij} \sim \text{Poisson}\left(\sum_{k=1}^K r_{ik}\theta_{kj}\right) \quad (1)$$

independently for $i = 1, \dots, I$ and $j = 1, \dots, J$, where $\text{Poisson}(\lambda)$ denotes the Poisson distribution with mean λ . In Section S4, we show that Equation (1) can be derived from first principles by modeling the occurrences of mutations as continuous-time Markov processes across the genome. Maximum likelihood estimation for R and Θ in Equation (1) is equivalent to minimizing the Kullback–Leibler divergence in the original formulation of Lee and Seung (2000); see Zhou and Carin (2015) for a discussion.

In mutational signatures analysis, it is common to impose the constraint that $\sum_{i=1}^I r_{ik} = 1$ for all $k = 1, \dots, K$. This avoids scaling ambiguities in both the signature vectors $r_k = (r_{1k}, \dots, r_{Ik})$ and their loadings $\theta_k = (\theta_{k1}, \dots, \theta_{kJ})$. Most methods do not impose such a constraint during inference, opting to enforce it as a post-processing step (Tan and Févotte, 2013; Drummond et al., 2023). However, we find that building $\sum_{i=1}^I r_{ik} = 1$ into our model has the additional benefits of simplifying the inference algorithm and enabling direct use of COSMIC signatures for constructing informative priors.

2.2 PRIOR

For the prior distribution on the signatures r_k and loadings θ_k , we take

$$r_k = (r_{1k}, \dots, r_{Ik}) \sim \text{Dirichlet}(\alpha, \dots, \alpha), \quad (2)$$

$$\theta_{k1}, \dots, \theta_{kJ} \mid \mu_k \sim \text{Gamma}(a, a/\mu_k), \quad (3)$$

$$\mu_k \sim \pi(\mu_k), \quad (4)$$

independently for $k = 1, \dots, K$, where $\alpha > 0$, $a > 0$, $\mu_k > 0$, and $\pi(\mu_k)$ is a hyperprior. The Dirichlet prior in Equation (2) automatically enforces the constraint that $\sum_{i=1}^I r_{ik} = 1$. Here, $\text{Gamma}(a, b)$ denotes the gamma distribution with mean a/b and variance a/b^2 . Thus, the prior mean of the loadings is $\mathbb{E}(\theta_{kj} \mid \mu_k) = \mu_k$, implying that μ_k controls the overall contribution of signature r_k to the factorization and, in turn, to the total number of mutations generated by process k .

We refer to μ_1, \dots, μ_K as *relevance weights*, following the usage of this type of prior structure in automatic relevance determination (ARD) for shrinking the weights of unneeded factors to near-zero values (Tan and Févotte, 2013). However, unlike Tan and Févotte (2013), we take a fully Bayesian approach, quantifying uncertainty in r_k and θ_k rather than just optimizing them. Also, unlike typical uses of ARD in Bayesian neural networks and Gaussian processes (Neal, 1996; Bishop, 2006), the marginal likelihood in the Poisson NMF model does not have a closed-form expression and thus is not amenable to direct optimization of μ_k . Nonetheless, it turns out that with a certain choice of data-dependent hyperprior on μ_k , we can obtain appealing computational properties similar to ARD, while performing inference with simple Gibbs sampling updates; we discuss this next.

2.3 COMPRESSIVE HYPERPRIOR

The hyperprior on the relevance weights in Equation (4) plays a crucial role in inferring the number of factors. Several approaches have been developed to provide sparsity in Gaussian factorization models, such as spike-and-slab priors that introduce exact zeros in the loadings (Carvalho et al., 2008; Ročková and George, 2016) or induce cumulative shrinkage to near-zero values for redundant factors (Legramanti et al., 2020; Frühwirth-Schnatter, 2023); also see Liu et al. (2019) for an extension to

NMF settings. However, in our model, we found that spike-and-slab priors over μ_k tend to make posterior inference difficult, likely due to the strong multimodal nature of the resulting posterior.

Instead, we propose a simpler alternative based on continuous shrinkage priors. Specifically, we let

$$\mu_k \sim \text{InvGamma}(aJ + 1, \varepsilon aJ) \quad (5)$$

independently for $k = 1, \dots, K$, where $\text{InvGamma}(a_0, b_0)$ denotes the inverse-gamma distribution with mean $b_0/(a_0 - 1)$ when $a_0 > 1$. Here, a is the shape parameter from the prior on θ_{kj} in Equation (3), and we set $\varepsilon > 0$ to be a small constant, such as $\varepsilon = 0.001$. Note that this makes μ_k small *a priori*, since $\mathbb{E}(\mu_k) = \varepsilon$. Further, the full conditional mean is

$$\mathbb{E}(\mu_k | -) = \frac{1}{2} \varepsilon + \frac{1}{2} \bar{\theta}_k, \quad (6)$$

where $\bar{\theta}_k = \frac{1}{J} \sum_{j=1}^J \theta_{kj}$, implying that the prior mean ε and the average loading $\bar{\theta}_k$ for signature k have equal influence on the posterior for μ_k . This *strength-matching* property of the hyperprior remains stable as the sample size J increases. It is important to note that this hyperprior should not be interpreted as a meaningful representation of prior uncertainty regarding μ_k . Instead, it is designed to yield a posterior with good properties in terms of computation and accuracy.

This deceptively simple choice of hyperprior has some key features. First, it favors sparse solutions, since $\mathbb{E}(\mu_k) = \varepsilon$. This makes it so that for any extra unneeded signatures, the hyperprior encourages the corresponding relevance weights μ_k to be small, on the order of ε . This leads to the compressive property of the model, which we analyze in more detail in Section 3.2.

Second, despite its growing strength with J , this hyperprior does not overly shrink the loadings θ_{kj} for factors that are needed to fit the data. To see why, suppose we instead used a fixed-strength hyperprior, say, $\mu_k \sim \text{InvGamma}(a_0, b_0)$. Then instead of $\varepsilon/2 + \bar{\theta}_k/2$, the full conditional mean of μ_k would be $(b_0 + aJ\bar{\theta}_k)/(a_0 + aJ)$, which is approximately $\bar{\theta}_k$ when J is sufficiently large. Consequently, when J is large and $\bar{\theta}_k \geq \varepsilon$, the prior mean of θ_{kj} differs by less than a multiple of 2 under the fixed-strength and strength-matching hyperpriors since $\mathbb{E}(\theta_{kj} | \mu_k) = \mu_k$, and this has only a small effect on the full conditional of θ_{kj} due to the overwhelming contribution from the likelihood.

Another important feature of the strength-matching hyperprior in Equation (5) is that small departures from the assumed Poisson NMF model do not strongly affect the number of factors used by the model. As J grows, a fixed-strength hyperprior on μ_k would be overwhelmed by the likelihood since the number of parameters θ_{kj} grows with J , leading to the inclusion of spurious extra signatures when the model is slightly misspecified. Meanwhile, the strength-matching drives out spurious extra signatures by maintaining a balance between the contribution from the loadings and from the hyperprior, as is evident from Equation (6).

Finally, the resulting posterior density is highly tractable due to the use of semi-conjugate distributions. This improves the performance of the sampling algorithm that we employ for inference, providing expeditious convergence to NMF solutions along with Bayesian uncertainty quantification.

2.4 POSTERIOR INFERENCE

Posterior inference for the hierarchical model defined by Equations (1) to (5) can be efficiently performed via Gibbs sampling. Since the sum of independent Poisson random variables is Poisson, we can equivalently write the hierarchical model as

$$\begin{aligned}
X_{ij} &= \sum_{k=1}^K Y_{ijk}, \\
Y_{ijk} \mid \mu_k, r_k, \theta_k &\sim \text{Poisson}(r_{ik}\theta_{kj}), \\
(r_{1k}, \dots, r_{Ik}) &\sim \text{Dirichlet}(\alpha, \dots, \alpha), \\
\theta_{k1}, \dots, \theta_{kJ} \mid \mu_k &\sim \text{Gamma}(a, a/\mu_k), \\
\mu_k &\sim \text{InvGamma}(aJ + 1, \varepsilon aJ).
\end{aligned} \tag{7}$$

Each auxiliary variable Y_{ijk} can be interpreted as the number of mutations due to signature k in channel i for sample j . Defining the vector $Y_{ij} = (Y_{ij1}, \dots, Y_{ijK})$, it turns out that $Y_{ij} \mid X_{ij}, R, \Theta$ follows a Multinomial($X_{ij}, (q_{ij1}, \dots, q_{ijK})$) distribution, where $q_{ijk} = r_{ik}\theta_{kj}/Q_{ij}$ and $Q_{ij} = \sum_{k=1}^K r_{ik}\theta_{kj}$. This auxiliary variable decomposition has been used in several previous methods (Dunson and Her-ring, 2005; Cemgil, 2009; Rosales et al., 2016; Zhou and Carin, 2015). The rest of the sampler relies on standard semi-conjugate updates, which are straightforward to derive.

GIBBS SAMPLER ALGORITHM

1. For $i = 1, \dots, I$ and $j = 1, \dots, J$, update the latent mutation counts by drawing

$$(Y_{ij} \mid -) \sim \text{Multinomial}(X_{ij}, (q_{ij1}, \dots, q_{ijK}))$$

where $q_{ijk} = r_{ik}\theta_{kj}/Q_{ij}$ and $Q_{ij} = \sum_{k=1}^K r_{ik}\theta_{kj}$.

2. For $k = 1, \dots, K$, update the signatures by drawing

$$(r_k \mid -) \sim \text{Dirichlet}\left(\alpha + \sum_{j=1}^J Y_{1jk}, \dots, \alpha + \sum_{j=1}^J Y_{Ijk}\right).$$

3. For $k = 1, \dots, K$ and $j = 1, \dots, J$, update the loadings by drawing

$$(\theta_{kj} \mid -) \sim \text{Gamma}\left(a + \sum_{i=1}^I Y_{ijk}, \frac{a}{\mu_k} + 1\right).$$

4. For $k = 1 \dots, K$, update the relevance weights by drawing

$$(\mu_k | -) \sim \text{InvGamma}\left(2aJ + 1, \varepsilon aJ + a \sum_{j=1}^J \theta_{kj}\right).$$

The model is symmetric with respect to the order of the factors, in the sense that the priors and likelihood are invariant to permutations of $k = 1, \dots, K$. While attractive from a modeling standpoint, this symmetry could potentially lead to label switching when running the Gibbs sampler, complicating the calculation of posterior expectations. However, we have not encountered label switching on either simulated or real data, so this has not been an issue in practice.

2.5 INFORMATIVE PRIORS BASED ON KNOWN SIGNATURES

A favorable aspect of mutational signatures analysis is the abundance of historical data on signatures across many cancer types. The COSMIC database contains a curated collection of signatures, annotated with associated cancer types and inferred etiologies (Alexandrov et al., 2020). It is natural to leverage such prior information as follows.

Suppose $s_k = (s_{1k}, \dots, s_{Ik})$, for $k = 1, \dots, K_{\text{pre}}$, are pre-defined mutational signatures known to occur in cancer. To allow for variation in signatures across studies, we let $\rho_k = (\rho_{1k}, \dots, \rho_{Ik})$ denote a study-specific version of s_k . We then generalize Equation (1) by modeling

$$X_{ij} \sim \text{Poisson}\left(\sum_{k=1}^{K_{\text{pre}}} \rho_{ik} \omega_{kj} + \sum_{k=1}^{K_{\text{new}}} r_{ik} \theta_{kj}\right) \quad (8)$$

independently, where r_{ik} and θ_{ik} are given the priors in Equations (2) and (3), respectively, and

$$\begin{aligned} \rho_k &\sim \text{Dirichlet}(\beta_k s_{1k}, \dots, \beta_k s_{1k}), \\ \omega_{kj} | \tau_k &\sim \text{Gamma}(b, b/\tau_k), \\ \tau_k &\sim \text{InvGamma}(bJ + 1, \varepsilon bJ). \end{aligned} \quad (9)$$

Thus, the prior on ρ_k is centered at s_k , with concentration parameter β_k . The loadings ω_{kj} and corresponding relevance weights τ_k are given the same prior and compressive hyperprior as θ_{kj} and μ_k , respectively, but with b in place of a .

The model in Equation (8) is reminiscent of the recovery-discovery model discussed in Grabski et al. (2023), when only a single study is taken into consideration. Notice that in such a framework, the prior rank for the lower-dimensional matrices, $K_{\text{pre}} + K_{\text{new}}$, is often greater than J . This is at odds with the classic approach to NMF, where the factorization rank is typically smaller than the rank of X . However, the compressive mechanism behind our priors in Equation (9) still ensures a

parsimonious representation in the posterior, such that only the active signatures have a nonnegligible relevance weight τ_k . Posterior inference can be performed using the same steps as in Section 2.4, with minor adjustments for handling ρ_k , ω_{kj} , and τ_k ; see Section S3 for details.

3 THEORY

In this section, we establish certain theoretical properties of the hierarchical Poisson NMF model in Section 2. In particular, we study the relationship between the latent mutation counts Y_{ijk} and the relevance weights μ_k , and we prove a concentration result for μ_k given the Y_{ijk} values.

3.1 CONDITIONAL DISTRIBUTION OF THE RELEVANCE WEIGHTS

We show that the distribution of $\mu_k | Y$ under the model in Equation (7), where Y is the tensor $(Y_{ijk}) \in \mathbb{R}^{I \times J \times K}$, has a closed-form expression in terms of confluent hypergeometric functions. We refer to the resulting family of distributions, which appears to be novel, as *inverse Kummer*.

Definition 1. The *inverse Kummer* distribution with parameters $\lambda > 0$, $\beta > 0$, $\delta > 0$, and $\gamma \in \mathbb{R}$ is a continuous distribution on $(0, \infty)$ with probability density function

$$\pi(\mu) = \frac{\mu^{-(\lambda-\gamma)-1} (1 + \mu/\delta)^{-\gamma} e^{-\beta/\mu}}{\delta^{\gamma-\lambda} \Gamma(\lambda) U(\lambda, \lambda + 1 - \gamma, \beta/\delta)}. \quad (10)$$

We write $\mu \sim \text{InvKummer}(\lambda, \beta, \gamma, \delta)$ to denote that μ has the density in Equation (10).

Here, $U(a, b, z)$ denotes the confluent hypergeometric function of the second kind,

$$U(a, b, z) = \frac{1}{\Gamma(a)} \int_0^\infty t^{a-1} (1+t)^{b-a-1} e^{-zt} dt,$$

with $z > 0$ (Abramowitz and Stegun, 1972). We call this an inverse Kummer distribution since if $\mu \sim \text{InvKummer}(\lambda, \beta, \gamma, \delta)$ then $1/\mu$ follows a Kummer distribution, which was introduced by Armero and Bayarri (1997) when studying a M/M/ ∞ queuing problem. The moments have closed-form expressions in terms of the hypergeometric function, following Equation (6.7) in Armero and Bayarri (1997): for $m < \lambda$, the m th moment of $\mu \sim \text{InvKummer}(\lambda, \beta, \gamma, \delta)$ is

$$\mathbb{E}(\mu^m) = \delta^m \frac{\Gamma(\lambda - m)}{\Gamma(\lambda)} \frac{U(\lambda - m, \lambda - m + 1 - \gamma, \beta/\delta)}{U(\lambda, \lambda + 1 - \gamma, \beta/\delta)}. \quad (11)$$

The inverse Kummer is a generalization of the inverse gamma distribution, since $\text{InvKummer}(\lambda, \beta, 0, \delta) = \text{InvGamma}(\lambda, \beta)$. To elucidate the relationship between the inverse Kummer and inverse gamma, in Section S2.1 we show that when $\lambda > 2$, the mean $\mathbb{E}(\mu)$ is monotonically increasing as a function of γ .

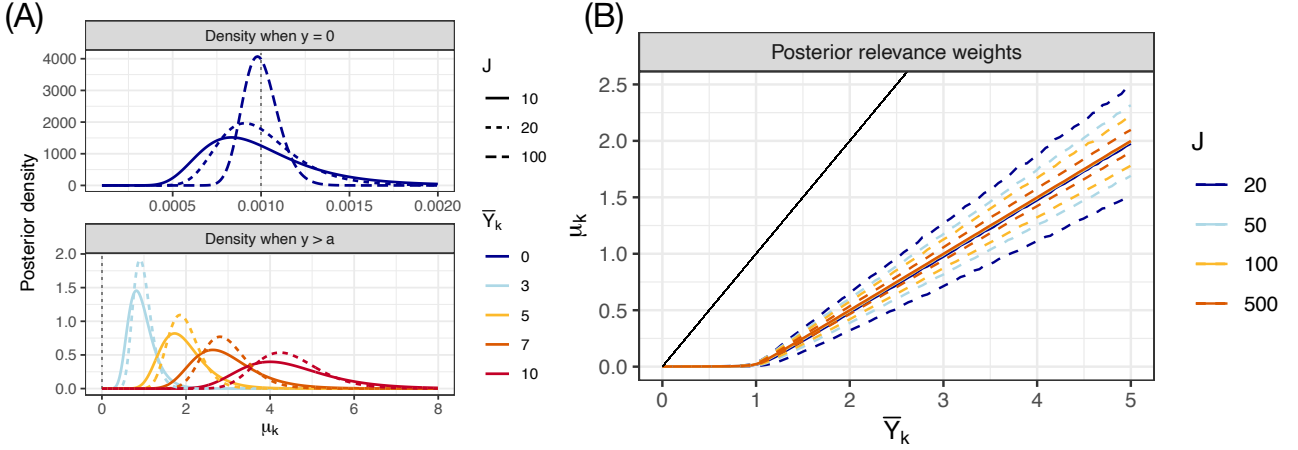


Figure 1: (A) Density of $\mu_k | Y$ as the average number of mutations \bar{Y}_k and the sample size J varies. Here, $a = 1$ and $\varepsilon = 0.001$. (B) Mean of $\mu_k | Y$ as a function of \bar{Y}_k , for varying values of J . Dashed lines indicate the 10th and the 90th percentiles of $\mu_k | Y$. The black line shows where $\mu_k = \bar{Y}_k$.

In the compressive NMF model in Equation (7), the inverse Kummer arises as the posterior distribution of the relevance weights given the latent counts, integrating out the signatures and loadings.

Theorem 1. Let $Y = (Y_{ijk}) \in \mathbb{R}^{I \times J \times K}$ denote the tensor of latent counts. Under the hierarchical model in Equation (7), it holds that

$$(\mu_k | Y) \sim \text{InvKummer}(2aJ + 1, \varepsilon aJ, J\bar{Y}_k + aJ, a),$$

where $\bar{Y}_k = \frac{1}{J} \sum_{i=1}^I \sum_{j=1}^J Y_{ijk}$ is the average number of mutations assigned to signature k .

See Section S1 for the proof. Figure 1(A) shows the density of $\mu_k | Y$ for various \bar{Y}_k and J values.

3.2 COMPRESSIVE PROPERTY

The essence of the compressive hyperprior is that for unneeded factors, the relevance weights are shrunk to $\approx \varepsilon$, while for needed factors they are only partly shrunk towards ε . Precisely this behavior is seen in Figure 1(B), which shows the relationship between \bar{Y}_k and the mean of $\mu_k | Y$, based on the result in Theorem 1. The average number of mutations due to the k th signature, \bar{Y}_k , is a minimal sufficient statistic for μ_k and plays a fundamental role in determining its posterior distribution. When $\bar{Y}_k \leq a$, we see that $\mu_k | Y$ becomes concentrated near ε . Meanwhile, when $\bar{Y}_k \geq a$, $\mathbb{E}(\mu_k | Y)$ grows approximately linearly as a function of \bar{Y}_k . Thus, given the latent counts Y , the model induces shrinkage by a fixed factor when \bar{Y}_k is large, while it has a sparsity inducing effect when \bar{Y}_k is small. Further, both Figure 1(A) and (B) illustrate that $\mu_k | Y$ concentrates as J increases.

In the following result, we make this rigorous by providing a concentration result for the posterior distribution of μ_k in Theorem 1 as $J \rightarrow \infty$.

Theorem 2. Consider the hierarchical model in Equation (7). If $\bar{Y}_k \rightarrow y \in [0, \infty)$ as $J \rightarrow \infty$, then for all $d > 0$,

$$\mathbb{P}(|\mu_k - \mu_*| \leq d \mid Y) \xrightarrow{J \rightarrow \infty} 1,$$

where $\mu_* = 2a\varepsilon / (\sqrt{(y - a + \varepsilon)^2 + 8a\varepsilon} - (y - a + \varepsilon))$.

In Section S2.2, we show that μ_* is monotone increasing as a function of y . Further, by a first-order Taylor approximation to the denominator of μ_* , when $\varepsilon \ll |y - a|$,

$$\mu_* \approx \begin{cases} \frac{y - a}{2} & \text{if } y > a \\ \frac{\varepsilon a(a - y)}{(a - y)^2 + (a + y)\varepsilon} & \text{if } 0 \leq y < a \end{cases} \quad (12)$$

This explains the shape of $\mathbb{E}(\mu_k \mid Y)$ seen in Figure 1; in particular, $\mu_* \approx \varepsilon a / (a + \varepsilon)$ as $y \approx 0$. The following corollary of Theorem 2 provides a criterion for thresholding the relevance weights.

Corollary 1. Under the assumptions of Theorem 2, if $\bar{Y}_k \rightarrow 0$ as $J \rightarrow \infty$, then for all $C > 1$,

$$\mathbb{P}(\mu_k > C\varepsilon \mid Y) \xrightarrow{J \rightarrow \infty} 0.$$

Hence, when signature k is not being used by the model, Corollary 1 shows that the posterior for μ_k concentrates on the interval $(0, C\varepsilon)$, for any $C > 1$. We refer to this as the *compressive property* of the model. This provides a natural criterion for selecting signatures for inclusion in the model, by using a threshold of $\mu_k > C\varepsilon$ to decide which signatures to keep and which to discard. Inspection of the proofs of Theorems 1 and 2 shows that they hold for any prior on signatures r_k such that $\sum_i r_{ik} = 1$. Consequently, the concentration results in Theorem 2 and Corollary 1 also hold for the relevance weights τ_k of the augmented model in Equation (9), which employs informative priors on r_k .

This compressive property parallels the automatic selection of the number of active components in overfitted mixture models described by [Rousseau and Mengersen \(2011\)](#), which has also been used in factor analysis (e.g., [Ferrari and Dunson, 2021](#)).

4 BACKGROUND ON MUTATIONAL SIGNATURES ANALYSIS

Cancer development in humans is connected to the accumulation of mutations in the DNA of somatic cells. When considering single-base substitutions, mutations are classified according to which of the four nucleotide bases was present before and after the mutation, on the strand containing the pyrimidine before the mutation occurred. Recalling that adenine (A) and guanine (G) are *purines* while cytosine (C) and thymine (T) are *pyrimidines* and that C always binds with G and T with A, there are six possible types of substitutions, namely, C>T, C>G, C>A, T>A, T>C and T>G. To

account for context-specific variability due to adjacent bases, mutations are further classified according to which bases (A, G, C, or T) occur on the 5' and the 3' sides on the strand containing the pre-substitution pyrimidine. This makes for a total of $6 \times 4 \times 4 = 96$ types of single-base substitutions, referred to as *mutational channels* (Alexandrov et al., 2013); see Section S4 for details.

It has been observed that many mutation-causing processes consistently produce each type of mutation at a particular rate: for instance, ultraviolet radiation has been observed to produce a large number of C>T substitutions in melanoma and glioma (Greenman et al., 2007). Due to this, the mutational processes acting on somatic cells can be characterized according to their *mutational signatures* (Nik-Zainal et al., 2012; Alexandrov et al., 2013), where the term “signature” refers to a vector containing the probability of each type of mutation under consideration.

A curated set of known signatures is maintained in the COSMIC database¹ (Alexandrov et al., 2020), which currently lists 86 single-base substitution (SBS) signatures. Many signatures can be attributed to a specific etiology that has been experimentally validated. For example, signatures SBS7a, b, c, and d are all linked to ultraviolet light exposure. Other signatures, such as SBS60, require further investigation to understand whether they arise from true biological processes or are due to technical artifacts. See Koh et al. (2021) for an overview.

The Poisson NMF model commonly used in mutational signatures analysis can be derived from first principles by modeling the occurrence of nucleotide substitutions at each base in the genome as a continuous-time Markov process (Section S4). Aggregating the resulting substitution counts across the entire genome, and modeling each mutational process as acting independently, it turns out that the counts are well approximated by the Poisson NMF model in Equation (1). This provides a compelling biological justification for the use of this model. We provide a detailed derivation in Section S4.

4.1 EXISTING METHODS

Several methods have been developed for mutational signature analysis; see the review by Islam et al. (2022). Currently, the most prominent method is SIGPROFILER (Alexandrov et al., 2020) and its successor, SIGPROFILEREXTRACTOR (Islam et al., 2022). SIGPROFILEREXTRACTOR uses a bootstrap-like procedure to resample the mutation count matrix from a Poisson model, and fits an NMF model to each resampled data matrix by minimizing the Poisson Kullback–Leibler divergence. The number of signatures is selected by applying the algorithm for a range of K values and using a neural network to choose K (Nebgen et al., 2021).

A leading Bayesian method is SIGNER (Rosales et al., 2016; Drummond et al., 2023) which is based on the Poisson NMF model in Equation (1), but employs independent hierarchical gamma priors over both the signatures and the loadings. Selection of K is performed using the Bayesian information

¹<https://cancer.sanger.ac.uk/signatures/>

criteria (BIC) after running a separate model for each K . Hence, both SIGPROFILEREXTRACTOR and SIGNER are particularly slow when the range of possible K values is moderate to large.

A faster alternative is offered by SIGNATUREANALYZER (Kim et al., 2016), which fits the Poisson NMF model in Equation (1) using a *maximum a posteriori* estimation algorithm (Tan and Févotte, 2013). Similarly to our approach, SIGNATUREANALYZER uses ARD with inverse-gamma hyperpriors on the relevance weights to determine the number of signatures. While the method is efficient and flexible in terms of the choice of the objective function (Kullback–Leibler or squared error) and prior (exponential or half-normal), it does not provide any uncertainty quantification.

A natural Bayesian approach to selecting the number of latent factors is to use spike-and-slab priors, where the relevance weight of each of the K signatures has some probability of being sampled from a spike close to zero. This is the approach taken in the elegant nonparametric factorization model proposed by Legramanti et al. (2020), which infers the number of factors using a cumulative shrinkage spike-and-slab process prior (CUSP). We adapt CUSP to the Poisson NMF model in Equation (1) by defining the loadings to be $\theta_{kj} = \mu_k \vartheta_{kj}$, where $\vartheta_{kj} \sim \text{Gamma}(a, a)$, and $\mu_k \sim (1 - \pi_k)\text{Gamma}(a_0, b_0) + \pi_k \delta_{\mu_\infty}$, with δ_x denoting the point mass at x . The spike location μ_∞ is typically chosen to be very small but nonzero, such as $\mu_\infty = 0.01$, while $\text{Gamma}(a_0, b_0)$ represents the slab. The spike probability π_k is constructed via a stick-breaking process, specifically, $\pi_k = \sum_{\ell=1}^k v_\ell \prod_{j=1}^{\ell-1} (1 - v_j)$ where $v_\ell \stackrel{\text{iid}}{\sim} \text{Beta}(1, \alpha)$ and $\alpha > 0$. Hence, this model assumes an infinite number of factors *a priori*, but the probability that μ_k comes from the spike—effectively removing that factor from the model—increases with k . Posterior inference is performed via an adaptive Metropolis algorithm; see Section S5.1 for details.

It is worth emphasizing that, unlike our model in Equation (8), none of these methods allow one to simultaneously use informative priors based on the COSMIC signatures and vague priors for *de novo* analysis. Some of these methods could potentially be adapted to Equation (8), but others cannot—for instance, the CUSP prior would result in an unnatural asymmetry across the known signatures.

5 SIMULATIONS

In this section, we conduct a simulation study to evaluate the performance of our compressive Poisson NMF method in terms of (a) detecting the true number of signatures that are active in the data and (b) accurately recovering the true signatures and their associated loadings.

5.1 SETUP OF THE SIMULATIONS

We simulate data and true parameters as follows. The mutation counts are generated as $X_{ij} \sim \text{NegBin}(1/\tau, 1/(1 + \tau\lambda_{ij}^0))$, where $\lambda_{ij}^0 = \sum_{k=1}^{K_{\text{pre}}^0} \rho_{ik}^0 \omega_{kj}^0 + \sum_{k=1}^{K_{\text{new}}^0} r_{ik}^0 \theta_{kj}^0$. We parametrize the negative binomial such that the mean and variance of X_{ij} are λ_{ij}^0 and $\lambda_{ij}^0(1 + \tau\lambda_{ij}^0)$, respectively, where $\tau > 0$

is a parameter controlling overdispersion. We set $K_{\text{pre}}^0 = 4$, and for $k = 1, \dots, 4$, we define $\rho_k^0 = (\rho_{1k}^0, \dots, \rho_{Ik}^0)$ to be COSMIC signatures SBS1, SBS2, SBS3, and SBS13, respectively. SBS1 is a sparse signature present in every cancer type, arising from the spontaneous deamination of 5-methylcytosine. SBS3 is a rather flat signature due to homologous recombination deficiencies in breast and pancreatic cancers. SBS2 and SBS13 are commonly occurring signatures associated with APOBEC activity. Meanwhile, we randomly generate $r_k^0 = (r_{1k}^0, \dots, r_{Ik}^0)$ as $r_k^0 \sim \text{Dirichlet}(0.25, \dots, 0.25)$, independently for $k = 1, \dots, K_{\text{new}}^0$. We generate loadings by setting $\omega_{kj}^0 = w_k \xi_{kj}$ where $w_k \sim \text{Gamma}(100, 1)$ and $\xi_{kj} \sim \text{Gamma}(0.5, 0.5)$ independently, and θ_{kj}^0 in the same way as ω_{kj}^0 .

For the overdispersion τ , we consider two settings: $\tau = 0$, in which case the negative binomial reduces to a Poisson (and thus, the assumed model is correct), and $\tau = 0.15$, resulting in mild misspecification. For the number of non-COSMIC signatures K_{new}^0 , we consider $K_{\text{new}}^0 = 2$ and $K_{\text{new}}^0 = 6$, so that the total number $K^0 = K_{\text{pre}}^0 + K_{\text{new}}^0$ is either 6 or 10. We consider a range of sample sizes $J \in \{50, 100, 200\}$. For each combination of τ , J , and K^0 , we generate 20 replicate sets of parameters and data matrices as described above. On each simulated data matrix, we run six methods:

- (i) CompNMF: our compressive NMF model in Equation (7) with $K = 20$ signatures and $\varepsilon = 0.001$,
- (ii) CompNMF+cosmic: our enhanced model in Equation (8) with $K_{\text{new}} = 15$ *de novo* signatures and the $K_{\text{pre}} = 67$ COSMIC v3.4 signatures that are not regarded as “possible sequencing artifacts”, and $\varepsilon = 0.001$,
- (iii) PoissonCUSP: the CUSP model with spike location $\mu_\infty = 0.01$, slab parameters $a_0 = b_0 = 1$, starting at $K = 20$ and using Algorithm 2 of Legramanti et al. (2020) to adaptively tune K ,
- (iv) signeR: the SIGNER model (Rosales et al., 2016; Drummond et al., 2023) with default parameters (`estimate_hyper = FALSE`) and with K ranging from 2 to 20,
- (v) SignatureAnalyzer: as implemented in the `sig_auto_extract` function of the `sigminer` package (Wang et al., 2020), with selection method set to L1KL and $K = 20$,
- (vi) SigProfiler: SIGPROFILEREXTRACTOR v1.1.23 (Islam et al., 2022) with random initialization and 5 replicates for each $K \in \{2, \dots, 20\}$, using the `sigprofiler_extract` wrapper in `sigminer`.

For methods (i), (ii), and (iii), we set $a = 1$ and $\alpha = 0.5$, and run the sampler for 5000 iterations, discarding the first 4000 as burn-in. In method (iii), adaptation of K was started after 500 iterations. In method (ii), we set the parameter β_k such that under the Dirichlet prior in Equation (9), the median cosine similarity between the prior mean and a sample from the prior is approximately 0.975. This makes β_k depend on the sparsity of the signature s_k . For example, we set $\beta_k = 17.29$ for the sparse signature SBS2, while $\beta_k = 1337.26$ for the rather flat SBS3.

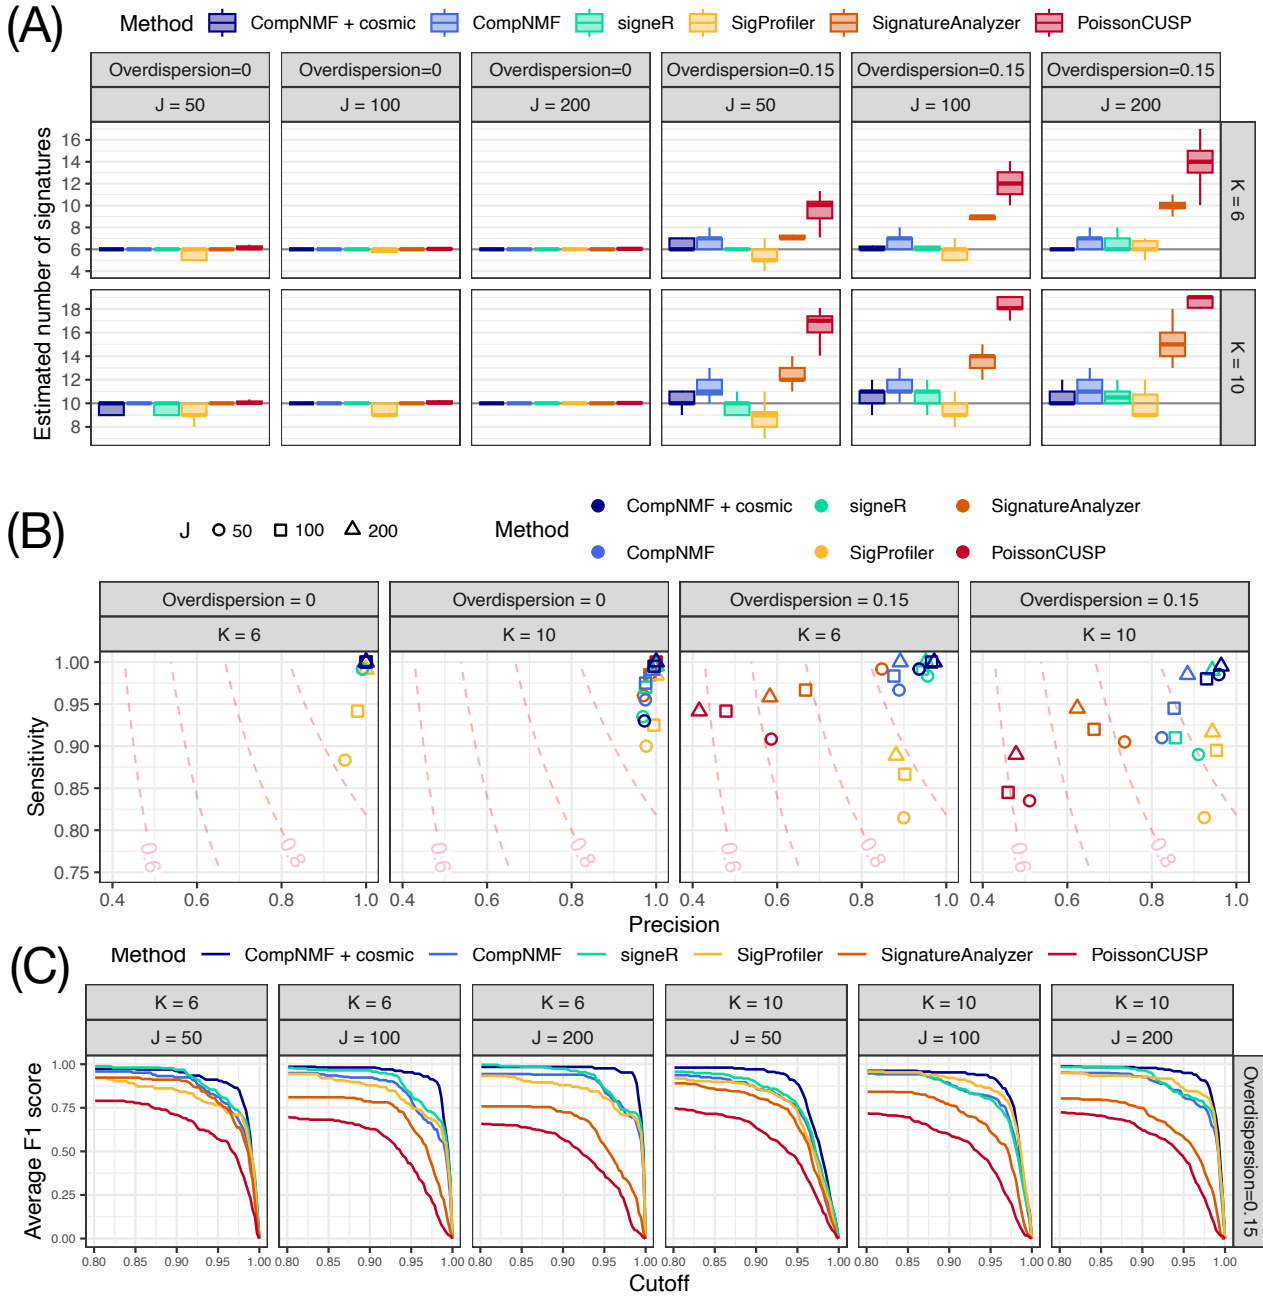


Figure 2: (A) Estimated number of signatures by each model, for varying choices of J , overdispersion, and true number of signatures, across 20 replicated data set in each scenario. The horizontal grey line indicates the true number ($K^0 = 6$ or $K^0 = 10$). (B) Average precision and sensitivity across 20 replicate data sets in each scenario, with a 0.9 cutoff for the cosine similarity. The dashed contour lines in the background indicate the F_1 score. (C) Average F_1 score as a function of the cosine similarity cutoff, across 20 replicates in each scenario, when the overdispersion is set to 0.15.

5.2 SIMULATION RESULTS

Figure 2 shows the main results of the simulation study. Figure 2(A) reports the estimated number of signatures for each method, for each combination of τ , J , and K^0 ; the boxplots summarize the distribution of estimated values across the 20 replicates. The estimated number of signatures K^* is defined as follows for each method: For (i) and (ii), K^* is the number of signatures for which the posterior mean of μ_k is greater than $5\varepsilon = 0.005$; for (iii), K^* is the number of signatures for which the posterior probability of being assigned to the spike is less than 0.05; and for (iv), (v), and (vi), K^* is the suggested solution returned by the corresponding package.

As expected, all methods accurately estimate K^0 when $\tau = 0$ (zero overdispersion), since the Poisson NMF model is correct in this case. Meanwhile, when $\tau = 0.15$ (mild overdispersion), we observe a noticeably different pattern of results. Our CompNMF+cosmic method, which leverages informative priors based on all 67 COSMIC signatures, correctly recovers K^0 more often than all other methods, both when $K^0 = 6$ and $K^0 = 10$. CompNMF, which does not rely on COSMIC, tends to slightly overestimate K^0 . This is likely due to the introduction of spurious signatures used by the model to accommodate the overdispersion. Interestingly, signeR works well even when the model is mildly misspecified, due to the tendency of the signeR to return a conservative estimate of the number of signatures, as discussed in Drummond et al. (2023). SigProfiler is even more conservative than signeR, frequently underestimating K^0 slightly in these simulations. In contrast, SignatureAnalyzer and PoissonCUSP strongly overestimate K^0 when there is overdispersion.

Figure 2(B) displays the precision and sensitivity, averaged over the 20 replicates, for each model in each setting. Here, we define *precision* as the proportion of estimated signatures that have a cosine similarity ≥ 0.9 with at least one of the ground truth signatures, and *sensitivity* is the proportion of ground truth signatures for which there is an estimated signature with cosine similarity ≥ 0.9 . Since this is a simulation study, the ground truth signatures ρ_k^0 and r_k^0 are known. The cutoff value of 0.9 was chosen following Islam et al. (2022), but we also vary it as described below. The contour lines in Figure 2(B) represent the F_1 score, defined as $2 \times \text{precision} \times \text{sensitivity} / (\text{precision} + \text{sensitivity})$.

Points in the top-right corner of the plot in Figure 2(B) indicate better performance. When the Poisson model is correct (no overdispersion), all of the methods perform well, except that SigProfiler has lower sensitivity for smaller sample sizes. CompNMF+cosmic consistently performs the best, with precision and sensitivity close to one in all of these settings. CompNMF and signeR also perform well, but with somewhat reduced precision and sensitivity when K^0 is larger and when there is overdispersion. SigProfiler has good precision but much lower sensitivity, particularly in the presence of overdispersion. Meanwhile, SignatureAnalyzer and PoissonCUSP struggle in the overdispersed scenarios, exhibiting severely degraded precision as well as low sensitivity.




Figure 3: Average computation time (in minutes) for each method. Times shown are the average across all combinations of overdispersion τ and number of signatures K^0 used in the simulation study.

To see the range of performance exhibited by each method as the cosine similarity cutoff varies, Figure 2(C) shows the average F_1 score versus the cutoff value. As before, we see that CompNMF+cosmic is a clear standout, generally exhibiting the best performance overall, and particularly excelling at cutoff values between 0.9 and 0.98. CompNMF, signeR, and SigProfiler are roughly comparable by this metric, whereas SignatureAnalyzer and especially PoissonCUSP suffer from significantly lower F_1 scores across all cutoff values.

Figure 3 shows the total computation time required by each method, as a function of the number of samples J . All computations were performed on an AMD Ryzen 3900-based dedicated server with 128GB of memory on Ubuntu 20.04, R version 4.3.1 linked to Intel MKL 2019.5-075. Calculations were split across 20 cores via the `foreach` package, allocating one dataset per core for each combination of τ , J , and K^0 , and running each method sequentially. Not surprisingly, SignatureAnalyzer is the fastest method by an order of magnitude, always taking under one minute to complete. PoissonCUSP is the second fastest method, because the number of signatures in the model varies adaptively within the Gibbs sampler, preventing unnecessary computations on inactive factors. CompNMF is slightly slower than PoissonCUSP since no adaptation is performed, but it is faster than signeR. SigProfiler was the slowest among all methods by a wide margin.

Thus, in terms of computation time, both CompNMF and CompNMF+cosmic are competitive with the next best-performing existing method, signeR. Furthermore, signeR is written in C++, whereas CompNMF and CompNMF+cosmic are written in R code, so they could be significantly faster if implemented in C++ or similar. Additionally, they could be sped up by updating only the non-compressed factors when running the Gibbs sampler, following the adaptive approach of PoissonCUSP.

The dominant performance of CompNMF+cosmic demonstrates the benefits of using our informative prior. Notably, since CompNMF+cosmic includes all $K_{\text{pre}} = 67$ COSMIC signatures—whereas only

$K_{\text{pre}}^0 = 4$ COSMIC signatures are used to generate the simulated data—this implies that 63 of these signatures are correctly compressed out of the model, illustrating the effectiveness of our compressive hyperprior. `signeR` performs impressively well in these simulations, tending to fall between `CompNMF` and `CompNMF+cosmic` in nearly all metrics. However, as mentioned above, the `CompNMF` methods would be significantly faster than `signeR` if implemented in the same language.

The poor performance of `SignatureAnalyzer` and `PoissonCUSP` in the misspecified settings appears to be due to overfitting. Indeed, the overestimation of K^0 seen in Figure 2(A), along with the fact that `SignatureAnalyzer` and `PoissonCUSP` yield the lowest root-mean-square error for the count matrices (see Section S5), is a clear sign of overfitting. In the case of `PoissonCUSP`, this is not surprising since the flexible nonparametric nature of the model allows it to closely fit the data, although unfortunately in this case it is fitting noise rather than signal.

6 APPLICATION: 21 BREAST CANCER DATA

In this section, we apply our compressive NMF method to the benchmark 21 breast cancer dataset considered by [Nik-Zainal et al. \(2012\)](#) in their landmark paper originating the study of mutational signatures in cancer. In particular, we aim to assess the effect of our compressive hyperprior for small sample sizes and evaluate whether our informative prior provides an advantage in practice.

The dataset is based on whole-genome sequencing of $J = 21$ patients, and consists of mutation counts for each of the $I = 96$ channels for each patient. We obtained the data from the `signeR` package ([Drummond et al., 2023](#)). The total mutation count is fairly homogeneous across patients, except for patient PD4120a, for whom a large number of mutations was detected (70,690) compared to the others (18,871 on average). This is attributable to the fact that PD4120a was sequenced more deeply at nearly 200x coverage compared to around 30x for the others.

We evaluate the performance of our compressive NMF method (with and without the informative COSMIC-based prior) compared to the same alternatives as in Section 5. Here, for both `CompNMF` and `CompNMF+cosmic`, we randomly initialize by sampling from the prior, and run the Gibbs sampler for 12,000 iterations, discarding the first 10,000 as burn-in. This is repeated four times, and we select the run yielding the highest average log-posterior. Details of the settings used in all of the other methods are reported in Section S6.

6.1 APPLICATION RESULTS

All of the methods fit the count matrix roughly equally well, with the exception of `SigProfiler`, which exhibited significantly higher error. Specifically, the root-mean-squared error (RMSE) between

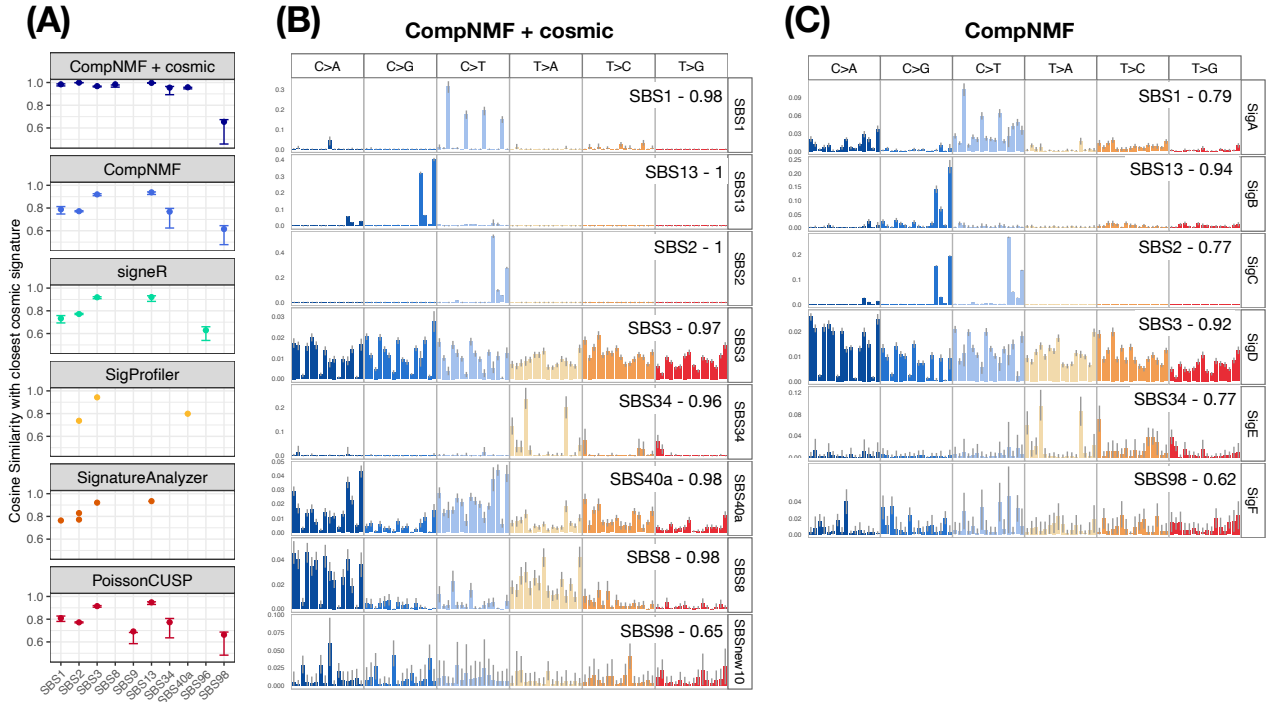


Figure 4: (A) Cosine similarity between the signatures inferred by each method and the best matching signature in COSMIC. Vertical error bars indicate 90% posterior credible intervals. (B) Signatures inferred by CompNMF+cosmic. (C) Signatures inferred by CompNMF. The vertical grey lines in panels (B) and (C) indicate univariate 90% posterior credible intervals for each mutational channel. The numbers on the top-right corners of each panel denote the cosine similarity with the best matching COSMIC signature.

the mutation count matrix X and $\hat{R}\hat{\Theta}$ was 9.51 for CompNMF, 9.57 for CompNMF+cosmic, 9.81 for signeR, 10.07 for PoissonCUSP, 10.08 for SignatureAnalyzer, and 37.08 for SigProfiler.

However, major differences were observed in the sets of estimated signatures. For each method, Figure 4(A) shows the cosine similarity between each estimated signature and the best matching COSMIC signature in terms of cosine similarity; see Section S6 for all estimated signatures. Overall, CompNMF+cosmic recovers the most signatures with the highest cosine similarities, which is not surprising since the COSMIC signatures are used to construct the informative prior used in CompNMF+cosmic. Nonetheless, this compellingly demonstrates that the informative prior enables more signal to be extracted from the data. For instance, CompNMF+cosmic finds evidence of the presence of SBS8 and SBS40a, which are not recovered by any other method except for SigProfiler, which finds a signature with moderate cosine similarity to SBS40a. SBS8 appears in other breast cancers and may be associated with homologous recombination deficiency. SBS40a is a rather flat signature that has unknown etiology but appears in every cancer type (Alexandrov et al., 2020).

CompNMF recovers all but two of the signatures found by CompNMF+cosmic, although with somewhat lower cosine similarities. PoissonCUSP yields similar results to CompNMF, but includes what appears to be a spurious match to SBS9, since it has low cosine similarity and SBS9 is not

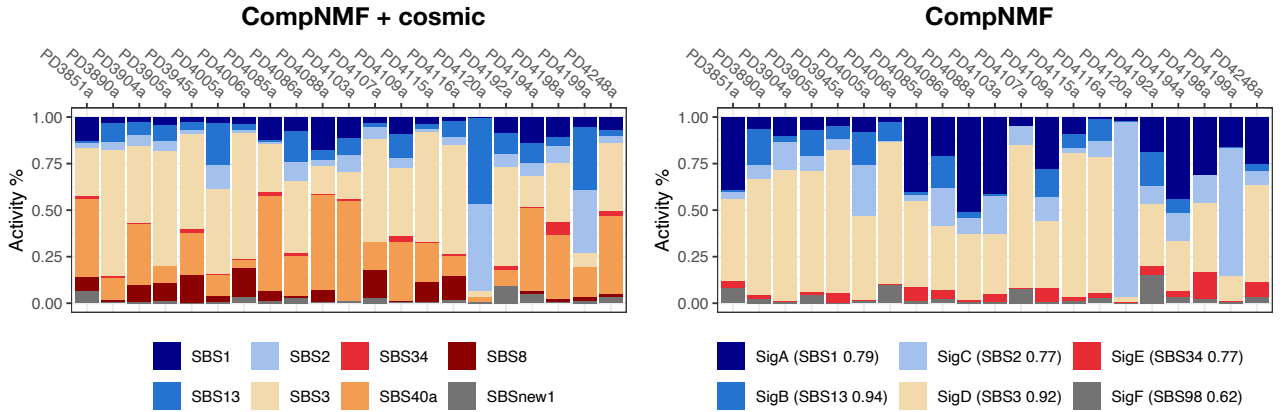


Figure 5: Average posterior loadings (%) by patient. Left: CompNMF+cosmic. Right: CompNMF.

found by any other method. Next, signerR is also similar to CompNMF, but does not recover SBS34, which does appear to be truly present since CompNMF+cosmic detects it in many samples. SignatureAnalyzer recovers four of the same signatures as CompNMF, with comparable cosine similarities; however, it misses SBS34 and produces two signatures that are both matched to SBS2, one of which is an amalgamation of SBS2 and SBS13. Lastly, SigProfiler recovers only three signatures. The strong performance of the CompNMF methods demonstrates the utility of our compressive hyperprior.

Figures 4(B) and (C) show the estimated signature vectors r_k for our two methods (CompNMF+cosmic and CompNMF); corresponding figures for all other methods are in Section S6. The informative prior enables CompNMF+cosmic to obtain clean signature estimates that are unambiguously matched to known COSMIC signatures, while still allowing for dataset-specific departures from the COSMIC signatures. In contrast, due to the small size of this dataset, all of the other methods produce “merged” signatures that are combinations of two or more COSMIC signatures. For instance, CompNMF+cosmic perfectly distinguishes SBS2 and SBS13, whereas these are merged into a single signature by all other methods; compare SBS2 and SBS13 in Figure 4(B) to CompNMF SigC in Figure 4(C) and the other methods’ results in Section S6. Both SBS2 and SBS13 are associated with APOBEC activity and appear frequently in breast cancers. Similarly, CompNMF+cosmic perfectly separates SBS1 and SBS40a, whereas these two are combined by all the other methods; see SigA in Figure 4(C) and see Section S6.

The vertical bars in Figure 4(A,B,C) indicate model-based uncertainty in the form of 90% credible intervals; there are no bars for SigProfiler and SignatureAnalyzer since they only provide point estimates. This uncertainty quantification is particularly useful for identifying signatures that may be spurious. For instance, for both CompNMF+cosmic and CompNMF, the estimated signature matched to SBS98 has high uncertainty in the cosine similarity and in the signature vector itself, indicating that one should be skeptical about the validity of the estimated signature and whether it should be matched to SBS98.

Figure 5 shows the loadings θ_{kj} estimated by CompNMF+cosmic and CompNMF; see Section S6 for all other methods. The CompNMF+cosmic loadings provide insight into the interpretation of the other methods’ results. For instance, in Figure 5, the loadings for SBS40a appear to have been split up and added to the loadings for SBS1 and SBS3 to make the CompNMF loadings for SigA and SigD, respectively. Indeed, looking at these signatures, it appears that SBS40a and SBS1 (Figure 4(B)) were agglomerated to make SigA (Figure 4(C)), and that SBS40a and SBS3 were agglomerated to make SigD. Similarly, CompNMF SigC appears to be a combination of SBS2 and SBS13. This is especially apparent in patient PD4120a, for whom almost all mutations are attributed to SigC by CompNMF, whereas CompNMF+cosmic splits the loading into equal contributions from SBS2 and SBS13.

7 DISCUSSION

This article introduces a novel Bayesian NMF method that obtains state-of-the-art performance for mutational signatures analysis, with lower computational burden. In particular, our compressive hyperprior provides a simple but effective technique for determining the subset of active factors. This enables the use of rich informative priors based on the COSMIC database of known signatures, significantly boosting the method’s precision and sensitivity for recovering true signatures. The informative prior also disambiguates the allocation of loadings to signatures, resulting in more accurate estimation of the contribution of each signature to each sample, as well as clarifying the results of other methods. Furthermore, the method provides posterior uncertainty quantification, which helps distinguish real from spurious signatures and can be used for downstream analyses.

There are several interesting directions for future work. First, one can envision several extensions of the model. Following [Grabski et al. \(2023\)](#), it would be interesting to jointly model multiple studies or multiple cancer types using a hierarchical model with study-specific or cancer type-specific parameters. Another useful extension of the model would be to include sample-specific covariates, which could be helpful in improving targeted therapies ([Aguirre et al., 2018](#)).

Additionally, the scope of applicability of the compressive hyperprior technique is potentially broader than Poisson NMF models, and might prove useful in other latent factorization models such as Gaussian factor models ([Bhattacharya and Dunson, 2011](#); [Legramanti et al., 2020](#)) or in user-item recommendation ([Gopalan et al., 2014](#)), especially when prior information on the factors is available.

Finally, while our method can handle mild misspecification in the form of small overdispersion, it is fundamentally based on the assumption that the counts are Poisson distributed – like all of the leading methods considered in our empirical results (signeR, SigProfiler, SignatureAnalyzer, PoissonCUSP). Consequently, larger departures from the assumed Poisson NMF model can be expected to negatively impact the performance of all of these methods. Of course, overdispersion can simply be handled by modeling the data as negative binomial rather than Poisson ([Lyu et al., 2020](#)). However, there are

many other plausible ways in which the model could be misspecified, and any parametric elaboration of the model will inevitably be misspecified in some way. Thus, an important area for future work is providing improved robustness to misspecification for mutational signatures analysis as well as non-negative matrix factorization more generally.

ACKNOWLEDGMENTS

We would like to thank Scott Carter for suggesting this research direction. We also thank Jonathan Huggins, Catherine Xue, and Giovanni Parmigiani for many helpful discussions. J.W.M. was supported by the Department of Defense grant W81XWH-18-1-0357 and the National Institutes of Health grant 5R01CA240299.

CODE AVAILABILITY

Code is publicly accessible via GitHub at <https://github.com/alessandrozito/CompressiveNMF>.

SUPPLEMENTARY MATERIAL

The supplementary material contains proofs of the theoretical results, the rationale for the Poisson NMF model likelihood, simulation details, and additional empirical results.

REFERENCES

- Abramowitz, M. and I. A. Stegun (1972). *Handbook of Mathematical Functions with formulas, graphs, and mathematical tables*. U.S. Dept. of Commerce, National Bureau of Standards.
- Aguirre, A. J., J. A. Nowak, N. D. Camarda, et al. (2018). Real-time genomic characterization of advanced pancreatic cancer to enable precision medicine. *Cancer Discovery* 8(9), 1096–1111.
- Alexandrov, L. B., J. Kim, N. J. Haradhvala, et al. (2020). The repertoire of mutational signatures in human cancer. *Nature* 578(7793), 94–101.
- Alexandrov, L. B., S. Nik-Zainal, D. C. Wedge, et al. (2013). Signatures of mutational processes in human cancer. *Nature* 500(7463), 415–421.
- Armero, C. and M. Bayarri (1997). A Bayesian analysis of a queueing system with unlimited service. *Journal of Statistical Planning and Inference* 58(2), 241–261.
- Ayed, F. and F. Caron (2021). Nonnegative Bayesian nonparametric factor models with completely random measures. *Statistics and Computing* 31(5), 63.

- Bhattacharya, A. and D. B. Dunson (2011). Sparse Bayesian infinite factor models. *Biometrika* 98(2), 291–306.
- Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. Springer.
- Carvalho, C. M., J. Chang, J. E. Lucas, J. R. Nevins, Q. Wang, and M. West (2008). High-dimensional sparse factor modeling: Applications in gene expression genomics. *Journal of the American Statistical Association* 103(484), 1438–1456.
- Cemgil, A. T. (2009). Bayesian inference for nonnegative matrix factorisation models. *Computational Intelligence and Neuroscience* 2009, 785152.
- Drummond, R. D., A. Defelicibus, M. Meyenberg, R. Valieris, E. Dias-Neto, R. A. Rosales, and I. T. da Silva (2023). Relating mutational signature exposures to clinical data in cancers via signER 2.0. *BMC Bioinformatics* 24(1), 439.
- Dunson, D. B. and A. H. Herring (2005). Bayesian latent variable models for mixed discrete outcomes. *Biostatistics* 6(1), 11–25.
- Ferrari, F. and D. B. Dunson (2021). Bayesian factor analysis for inference on interactions. *Journal of the American Statistical Association* 116(535), 1521–1532.
- Fischer, A., C. J. Illingworth, P. J. Campbell, and V. Mustonen (2013). EMu: probabilistic inference of mutational processes and their localization in the cancer genome. *Genome Biology* 14(4), R39.
- Frühwirth-Schnatter, S. (2023). Generalized cumulative shrinkage process priors with applications to sparse Bayesian factor analysis. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 381(2247), 20220148.
- Gopalan, P., F. J. Ruiz, R. Ranganath, and D. Blei (2014). Bayesian nonparametric Poisson factorization for recommendation systems. In S. Kaski and J. Corander (Eds.), *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*, Volume 33 of *Proceedings of Machine Learning Research*, pp. 275–283.
- Grabski, I., L. Trippa, and G. Parmigiani (2023). Bayesian multi-study non-negative matrix factorization for mutational signatures. *bioRxiv:10.1101/2023.03.28.534619*, 1–12.
- Greenman, C., P. Stephens, R. Smith, et al. (2007). Patterns of somatic mutation in human cancer genomes. *Nature* 446(7132), 153–158.
- Gulhan, D. C., J. J.-K. Lee, G. E. M. Melloni, I. Cortés-Ciriano, and P. J. Park (2019). Detecting the mutational signature of homologous recombination deficiency in clinical samples. *Nature Genetics* 51(5), 912–919.

- Hoffman, M. D., D. M. Blei, and P. R. Cook (2010). Bayesian nonparametric matrix factorization for recorded music. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, pp. 439–446.
- Islam, S. A., M. Díaz-Gay, Y. Wu, et al. (2022). Uncovering novel mutational signatures by de novo extraction with SigProfilerExtractor. *Cell Genomics* 2(11), 100179.
- Kim, J., K. W. Mouw, P. Polak, et al. (2016). Somatic ERCC2 mutations are associated with a distinct genomic signature in urothelial tumors. *Nature Genetics* 48(6), 600–606.
- Koh, G., A. Degasperi, X. Zou, S. Momen, and S. Nik-Zainal (2021). Mutational signatures: emerging concepts, caveats and clinical applications. *Nature Reviews Cancer* 21(10), 619–637.
- Lal, A., K. Liu, R. Tibshirani, A. Sidow, and D. Ramazzotti (2021). De novo mutational signature discovery in tumor genomes using SparseSignatures. *PLOS Computational Biology* 17(6), 1–24.
- Lee, D. and H. S. Seung (2000). Algorithms for non-negative matrix factorization. In T. Leen, T. Dietterich, and V. Tresp (Eds.), *Advances in Neural Information Processing Systems*, Volume 13. MIT Press.
- Legramanti, S., D. Durante, and D. B. Dunson (2020). Bayesian cumulative shrinkage for infinite factorizations. *Biometrika* 107(3), 745–752.
- Liu, Y., W. Dong, W. Song, and L. Zhang (2019). Bayesian nonnegative matrix factorization with a truncated spike-and-slab prior. In *2019 IEEE International Conference on Multimedia and Expo (ICME)*, pp. 1450–1455.
- Lyu, X., J. Garret, G. Rätsch, and K.-V. Lehmann (2020). Mutational signature learning with supervised negative binomial non-negative matrix factorization. *Bioinformatics* 36(Issue Supplement_1), i154–i160.
- Neal, R. M. (1996). *Bayesian Learning for Neural Networks*. Number 118. Springer. Lecture Notes in Statistics.
- Nebgen, B. T., R. Vangara, M. A. Hombrados-Herrera, S. Kuksova, and B. S. Alexandrov (2021). A neural network for determination of latent dimensionality in non-negative matrix factorization. *Machine Learning: Science and Technology* 2(2), 025012.
- Nik-Zainal, S., L. B. Alexandrov, D. C. Wedge, P. Van Loo, C. D. Greenman, K. Raine, D. Jones, J. Hinton, J. Marshall, L. A. Stebbings, A. Menzies, S. Martin, K. Leung, L. . Chen, and M. R. Stratton (2012). Mutational processes molding the genomes of 21 breast cancers. *Cell* 149(5), 979–993.

- Nik-Zainal, S., H. Davies, J. Staaf, et al. (2016). Landscape of somatic mutations in 560 breast cancer whole-genome sequences. *Nature* 534(7605), 47–54.
- Ročková, V. and E. I. George (2016). Fast Bayesian factor analysis via automatic rotations to sparsity. *Journal of the American Statistical Association* 111(516), 1608–1622.
- Rosales, R. A., R. D. Drummond, R. Valieris, E. Dias-Neto, and I. T. da Silva (2016). signeR: an empirical Bayesian approach to mutational signature discovery. *Bioinformatics* 33(1), 8–16.
- Rousseau, J. and K. Mengersen (2011). Asymptotic behaviour of the posterior distribution in overfitted mixture models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 73(5), 689–710.
- Rubanova, Y., R. Shi, C. F. Harrigan, R. Li, J. Wintersinger, N. Sahin, A. G. Deshwar, et al. (2020). Reconstructing evolutionary trajectories of mutation signature activities in cancer using TrackSig. *Nature Communications* 11(1), 731.
- Tan, V. Y. and C. Févotte (2013). Automatic relevance determination in nonnegative matrix factorization with the beta-divergence. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35(7), 1592–1605.
- Wang, S., Z. Tao, T. Wu, and X.-S. Liu (2020). Sigflow: an automated and comprehensive pipeline for cancer genome mutational signature analysis. *Bioinformatics* 37(11), 1590–1592.
- Wang, Y.-X. and Y.-J. Zhang (2013). Nonnegative matrix factorization: A comprehensive review. *IEEE Transactions on Knowledge and Data Engineering* 25(6), 1336–1353.
- Zhou, M. (2018). Nonparametric Bayesian negative binomial factor analysis. *Bayesian Analysis* 13(4), 1065–1093.
- Zhou, M. and L. Carin (2015). Negative binomial process count and mixture modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 37(2), 307–320.

Supplementary Material for “Compressive Bayesian non-negative matrix factorization for mutational signatures analysis”

Section S1 contains the proofs of the theoretical results. Section S2 provides additional results related to the inverse Kummer distribution and the marginal distribution of the latent counts. Section S3 gives a detailed description of the sampler for the NMF with informative priors. Section S4 provides additional background on mutational signatures, including a first-principles derivation of the model. Section S5 contains additional details of the simulations in Section 5, including an adaptation of the CUSP model of [Legramanti et al. \(2020\)](#) to the setting of Poisson non-negative matrix factorization. Finally, Section S6 provides additional results on the application.

S1 PROOFS

PROOF OF THEOREM 1

Proof. The proof proceeds by directly marginalizing r_k and θ_k from the joint distribution. For notational simplicity, we first handle the general case of $\mu_k \sim \text{InvGamma}(a_0, b_0)$, and then plug in the values of a_0 and b_0 for the compressive hyperprior. Under the model in Equation (7), the joint density of $Y = (Y_{ijk})$, $\Theta = (\theta_{kj})$, $R = (r_{ik})$, and $\mu = (\mu_k)$ is

$$\begin{aligned}
 \pi(Y, \Theta, R, \mu) \propto & \left\{ \prod_{i,j,k} e^{-r_{ik}\theta_{kj}} \frac{(r_{ik}\theta_{kj})^{Y_{ijk}}}{Y_{ijk}!} \right\} \times && \text{(Latent Poisson counts)} \\
 & \left\{ \prod_{i,k} r_{ik}^{\alpha-1} \right\} \times && \text{(Dirichlet prior on } r_k) \\
 & \left\{ \prod_{j,k} \left(\frac{a}{\mu_k} \right)^a \theta_{kj}^{a-1} e^{-a\theta_{kj}/\mu_k} \right\} \times && \text{(Gamma prior on } \theta_k) \\
 & \left\{ \prod_k \mu_k^{-a_0-1} e^{-b_0/\mu_k} \right\}, && \text{(InvGamma prior on } \mu_k)
 \end{aligned} \tag{S1}$$

dropping constants of proportionality. Since $\sum_i r_{ik} = 1$ for all k , we have $\prod_i e^{-r_{ik}\theta_{kj}} = e^{-\theta_{kj}}$. Thus, by Equation (S1), we have

$$\pi(Y, \Theta, R, \mu) = f(Y, R) \left\{ \prod_{j,k} \theta_{kj}^{\sum_i Y_{ijk} + a - 1} e^{-\theta_{kj} - a\theta_{kj}/\mu_k} \right\} \left\{ \prod_k \mu_k^{-Ja - a_0 - 1} e^{-b_0/\mu_k} \right\} \tag{S2}$$

where $f(Y, R)$ is a function that does not depend on Θ or μ . Hence,

$$\begin{aligned} \pi(\mu | Y) &\propto_{\mu} \pi(Y, \mu) = \int \int \pi(Y, \Theta, R, \mu) d\Theta dR \\ &= \left\{ \int f(Y, R) dR \right\} \left\{ \prod_k \mu_k^{-Ja-a_0-1} e^{-b_0/\mu_k} \right\} \left\{ \prod_{j,k} \int \theta_{kj}^{\sum_i Y_{ijk}+a-1} e^{-\theta_{kj}-a\theta_{kj}/\mu_k} d\theta_{kj} \right\} \\ &\propto_{\mu} \left\{ \prod_k \mu_k^{-Ja-a_0-1} e^{-b_0/\mu_k} \right\} \left\{ \prod_{j,k} \frac{\Gamma(a + \sum_i Y_{ijk})}{(1 + a/\mu_k)^{a+\sum_i Y_{ijk}}} \right\}. \end{aligned}$$

Since this factors over k into products that depend on μ_1, \dots, μ_K , respectively, it follows that

$$\begin{aligned} \pi(\mu_k | Y) &\propto \mu_k^{-Ja-a_0-1} e^{-b_0/\mu_k} \left(1 + \frac{a}{\mu_k}\right)^{-Ja-\sum_{i,j} Y_{ijk}} \\ &\propto \mu_k^{-(a_0-\sum_{i,j} Y_{ijk})-1} \left(\frac{\mu_k}{a} + 1\right)^{-Ja-\sum_{i,j} Y_{ijk}} e^{-b_0/\mu_k} \\ &\propto \text{InvKummer}\left(\mu_k \mid a_0 + Ja, b_0, Ja + \sum_{i,j} Y_{ijk}, a\right) \end{aligned}$$

by Definition 1. The proof is completed by letting $a_0 = aJ + 1$ and $b_0 = \varepsilon aJ$. \square

PROOF OF THEOREM 2

To establish Theorem 2, we prove the following concentration result for inverse Kummer distributions. Theorem 2 then follows immediately from Theorem 1 and Theorem 3.

Theorem 3. *Let $\varepsilon > 0$, $a > 0$, and $y_n \geq 0$ such that $y_n \rightarrow y$ for some $y \in [0, \infty)$ as $n \rightarrow \infty$. If $\mu_n \sim \text{InvKummer}(2an + 1, \varepsilon an, ny_n + an, a)$ then for any $d > 0$,*

$$\mathbb{P}(|\mu_n - \mu_*| \leq d) \xrightarrow{n \rightarrow \infty} 1$$

where $\mu_* = 2a\varepsilon / (\sqrt{(y-a+\varepsilon)^2 + 8a\varepsilon} - (y-a+\varepsilon))$.

Proof. We first prove concentration of $T_n = a/\mu_n$, and use this to show the concentration of μ_n . Applying this change of variables to Equation (10) and simplifying, we find that the density of T_n is

$$f_n(t) \propto t^{2an}(1+t)^{-ny_n-an} e^{-\varepsilon nt} = \exp(-ng_n(t)),$$

where $g_n(t)$ and its derivative are

$$g_n(t) = -2a \log(t) + (y_n + a) \log(1+t) + \varepsilon t \tag{S3}$$

$$g'_n(t) = -\frac{2a}{t} + \frac{y_n + a}{1+t} + \varepsilon. \tag{S4}$$

By Lemma 1, $g'_n(t)$ has a unique zero in $(0, \infty)$ at

$$t_n^* = \frac{\sqrt{(y_n - a + \varepsilon)^2 + 8a\varepsilon} - (y_n - a + \varepsilon)}{2\varepsilon}, \quad (\text{S5})$$

$g'_n(t)$ is strictly monotone increasing on $(0, t_n^*)$, and for all $t_0 > 0$, we have

$$\inf_{t \geq t_0} g'_n(t) \geq \min\{\varepsilon, g'_n(t_0)\}. \quad (\text{S6})$$

Further, $g'_n(t) < 0$ for all $t \in (0, t_n^*)$, and $g'_n(t) > 0$ for all $t > t_n^*$. Define $t^* = \lim_{n \rightarrow \infty} t_n^*$, noting that by continuity and the fact that $y_n \rightarrow y$ by assumption,

$$t^* = \frac{\sqrt{(y - a + \varepsilon)^2 + 8a\varepsilon} - (y - a + \varepsilon)}{2\varepsilon}. \quad (\text{S7})$$

Let $\delta \in (0, t_*)$ be arbitrary. Let $B_r(t^*) = \{t > 0 : |t - t^*| \leq r\}$ denote the ball of radius r around t^* in $(0, \infty)$, and let $B_r(t^*)^c = \{t > 0 : |t - t^*| > r\}$ denote its complement in $(0, \infty)$. Then

$$\mathbb{P}(|T_n - t^*| > \delta) = \frac{\int_{B_\delta(t^*)^c} e^{-ng_n(t)} dt}{\int_0^\infty e^{-ng_n(t)} dt} \leq \frac{\int_{B_\delta(t^*)^c} e^{-ng_n(t)} dt}{\int_{B_\delta(t^*)} e^{-ng_n(t)} dt}. \quad (\text{S8})$$

We seek a suitable upper bound for the numerator and a lower bound for the denominator to show that Equation (S8) converges to 0 as $n \rightarrow \infty$. Define $\alpha_n = \min\{g_n(t^* + \delta), g_n(t^* - \delta)\}$ and $\beta_n = \min\{\varepsilon, g'_n(t^* + \delta), -g'_n(t^* - \delta)\}$. To bound the numerator, we claim that for all n sufficiently large,

$$g_n(t) \geq \alpha_n + \beta_n(t - (t^* + \delta)) \text{ for } t > t^* + \delta, \quad (\text{S9})$$

$$g_n(t) \geq \alpha_n - \beta_n(t - (t^* - \delta)) \text{ for } 0 < t < t^* - \delta. \quad (\text{S10})$$

To see this, choose n large enough that $|t_n^* - t^*| < \delta$. Then by Equation (S6) with $t_0 = t^* + \delta$, we have $g'_n(t) \geq \min\{\varepsilon, g'_n(t^* + \delta)\} \geq \beta_n$ for all $t \geq t^* + \delta$. Hence, for $t \geq t^* + \delta$,

$$g_n(t) = g_n(t^* + \delta) + \int_{t^* + \delta}^t g'_n(\tau) d\tau \geq \alpha_n + \int_{t^* + \delta}^t \beta_n d\tau = \alpha_n + \beta_n(t - (t^* + \delta)),$$

proving Equation (S9). Meanwhile, since $g'_n(t)$ is monotone increasing on $(0, t_n^*)$ and $t^* - \delta \leq t_n^*$, we have $g'_n(t) \leq g'_n(t^* - \delta)$ for $0 < t < t^* - \delta$. Thus, for $0 < t < t^* - \delta$,

$$g_n(t) = g_n(t^* - \delta) - \int_t^{t^* - \delta} g'_n(\tau) d\tau \geq \alpha_n + \int_t^{t^* - \delta} \beta_n d\tau = \alpha_n - \beta_n(t - (t^* - \delta)),$$

proving Equation (S10). Hence, for the numerator in Equation (S8), for all n sufficiently large,

$$\begin{aligned}
\int_{B_\delta(t^*)^c} e^{-ng_n(t)} dt &= \int_0^{t^*-\delta} e^{-ng_n(t)} dt + \int_{t^*+\delta}^\infty e^{-ng_n(t)} dt \\
&\leq \int_0^{t^*-\delta} e^{-n(\alpha_n - \beta_n(t - (t^* - \delta)))} dt + \int_{t^*+\delta}^\infty e^{-n(\alpha_n + \beta_n(t - (t^* + \delta)))} dt \\
&\leq 2e^{-n\alpha_n} \int_0^\infty e^{-n\beta_n t} dt \\
&\leq \frac{2e^{-n\alpha_n}}{n\beta_n}.
\end{aligned} \tag{S11}$$

Next, we turn to finding a lower bound for the denominator of Equation (S8). Since $y_n \rightarrow y$, we have that $g_n(t) \rightarrow g(t)$ and $g'_n(t) \rightarrow g'(t)$ pointwise for all $t > 0$, where $g(t)$ is defined by replacing y_n with y in Equation (S3). Define $\alpha = \lim_{n \rightarrow \infty} \alpha_n$ and $\beta = \lim_{n \rightarrow \infty} \beta_n$. By Lemma 1 applied to $g'(t)$, we have $g'(t) < 0$ for all $t \in (0, t^*)$, and $g'(t) > 0$ for all $t > t^*$. Thus, $\beta > 0$ and $\alpha > g(t^*)$ since t^* is the unique minimizer of $g(t)$. Pick any γ such that $g(t^*) < \gamma < \alpha$. By the continuity of $g(t)$, we can choose $r \in (0, \delta)$ such that $\sup_{t \in B_r(t^*)} g(t) < \gamma$. Note that g_n converges uniformly to g on $B_\delta(t^*)$, since for all $t \in B_\delta(t^*)$,

$$|g_n(t) - g(t)| = |y_n - y| \log(1 + t) \leq |y_n - y| \log(1 + t^* + \delta).$$

Therefore, for all n sufficiently large, $\sup_{t \in B_r(t^*)} g_n(t) < \gamma$, and hence,

$$\int_{B_\delta(t^*)} e^{-ng_n(t)} dt \geq \int_{B_r(t^*)} e^{-ng_n(t)} dt \geq \int_{B_r(t^*)} e^{-n\gamma} dt = 2re^{-n\gamma}. \tag{S12}$$

Applying the bounds in Equations (S11) and (S12) to Equation (S8) yields

$$\mathbb{P}(|T_n - t^*| > \delta) \leq \frac{\int_{B_\delta(t^*)^c} e^{-ng_n(t)} dt}{\int_{B_\delta(t^*)} e^{-ng_n(t)} dt} \leq \frac{2e^{-n\alpha_n}/(n\beta_n)}{2re^{-n\gamma}} = \frac{e^{-n(\alpha_n - \gamma)}}{n\beta_n r} \xrightarrow{n \rightarrow \infty} 0$$

since $\alpha_n \rightarrow \alpha > \gamma$ and $\beta_n \rightarrow \beta > 0$. Since $\delta \in (0, t^*)$ was arbitrary, this implies $\mathbb{P}(|T_n - t^*| \leq \delta) \rightarrow 1$ for all $\delta > 0$.

Now, define $\mu_* = a/t^*$ and let $d \in (0, \mu_*)$. Then $|\mu_n - \mu_*| \leq d$ if and only if $\mu_* - d \leq \mu_n \leq \mu_* + d$, or equivalently, $a/(\mu_* + d) \leq a/\mu_n \leq a/(\mu_* - d)$. Choose $\delta > 0$ small enough that

$$\frac{a}{\mu_* + d} < \frac{a}{\mu_*} - \delta = t^* - \delta < t^* + \delta = \frac{a}{\mu_*} + \delta < \frac{a}{\mu_* - d}.$$

Then, recalling that $T_n = a/\mu_n$, we have $\mathbb{P}(|\mu_n - \mu_*| \leq d) \geq \mathbb{P}(|T_n - t^*| \leq \delta) \rightarrow 1$ as $n \rightarrow \infty$. \square

Lemma 1. *Let $a > 0$, $y \geq 0$, and $\varepsilon > 0$, and define $h(t) = -2a/t + (y + a)/(1 + t) + \varepsilon$ for $t > 0$. Then*

1. $h(t)$ has a unique zero on $(0, \infty)$ at $t^* = (\sqrt{(y-a+\varepsilon)^2 + 8a\varepsilon} - (y-a+\varepsilon))/(2\varepsilon)$,
2. $h(t) < 0$ for $0 < t < t^*$ and $h(t) > 0$ for $t > t^*$,
3. $h(t)$ is strictly monotone increasing on $(0, t^*)$, and
4. for all $t_0 > 0$, we have $\inf_{t \geq t_0} h(t) \geq \min\{\varepsilon, h(t_0)\}$.

Proof. To find the zeros, we set $h(t) = 0$ and solve. For $t > 0$, we have $h(t) = 0$ if and only if

$$0 = t(t+1)h(t) = -2a(t+1) + (y+a)t + \varepsilon t(t+1) = \varepsilon t^2 + (y-a+\varepsilon)t - 2a.$$

By the quadratic formula, this equation has zeros at $t^* = (-b \pm \sqrt{b^2 + 8a\varepsilon})/(2\varepsilon)$ where $b = y - a + \varepsilon$. Since $|-b| = \sqrt{b^2} < \sqrt{b^2 + 8a\varepsilon}$, there is exactly one positive solution t^* at

$$t^* = \frac{\sqrt{(y-a+\varepsilon)^2 + 8a\varepsilon} - (y-a+\varepsilon)}{2\varepsilon}.$$

Therefore, $h(t)$ has a unique zero on $(0, \infty)$ at t^* , proving part 1. Since $h(t) \rightarrow -\infty$ as $t \rightarrow 0$, and $h(t) \rightarrow \varepsilon$ as $t \rightarrow \infty$, the intermediate value theorem implies that $h(t) < 0$ for $0 < t < t^*$ and $h(t) > 0$ for $t > t^*$. This proves part 2.

For part 3, to show that $h(t)$ is strictly increasing on $(0, t^*)$, let $0 < t_1 < t_2 < t^*$. Define $\tilde{h}(t) = -2a/t + (y+a)/(1+t) + \tilde{\varepsilon}$ where $\tilde{\varepsilon} = \varepsilon - h(t_2)$, noting that $\tilde{\varepsilon} > 0$ since $h(t_2) < 0$ by part 2. Since $\tilde{h}(t_2) = 0$ by the definition of \tilde{h} , by applying part 1 to \tilde{h} we know t_2 is the unique zero of \tilde{h} . Therefore, by applying part 2 to $\tilde{h}(t)$, we know $\tilde{h}(t_1) < 0$ since $t_1 < t_2$. Since $\tilde{h}(t) = h(t) - h(t_2)$, this implies $h(t_1) < h(t_2)$ as desired.

To show part 4, we use proof by contradiction. Let $t_0 > 0$, and suppose there exists $t_1 > t_0$ such that $h(t_1) < \min\{\varepsilon, h(t_0)\}$. Define $\tilde{h}(t) = -2a/t + (y+a)/(1+t) + \tilde{\varepsilon}$ where $\tilde{\varepsilon} = \varepsilon - h(t_1)$, noting that $\tilde{\varepsilon} > 0$ since $h(t_1) < \varepsilon$ by assumption. Further, $\tilde{h}(t_1) = 0$ by the definition of \tilde{h} , and $\tilde{h}(t_0) > \tilde{h}(t_1)$ since $\tilde{h}(t_0) - \tilde{h}(t_1) = h(t_0) - h(t_1) > 0$ by assumption. Thus, by the intermediate value theorem, there exists $t' \in (0, t_0)$ such that $\tilde{h}(t') = \tilde{h}(t_1) = 0$, because $\tilde{h}(t) \rightarrow -\infty$ as $t \rightarrow 0$. But then $\tilde{h}(t)$ would have two zeros in $(0, \infty)$, which is a contradiction. Therefore, $h(t) \geq \min\{\varepsilon, h(t_0)\}$ for all $t > t_0$, which proves part 4. \square

PROOF OF COROLLARY 1

Proof. Suppose $\bar{Y}_k \rightarrow 0$. Then in Theorem 2, $y = 0$ and therefore the concentration point is $\mu^* = 2a\varepsilon/(\sqrt{(\varepsilon-a)^2 + 8a\varepsilon} - (\varepsilon-a))$. It holds that $\mu^* < \varepsilon$, since $(\varepsilon-a)^2 + 8a\varepsilon > (\varepsilon-a)^2 + 4a\varepsilon = (\varepsilon+a)^2$, and hence, $\sqrt{(\varepsilon-a)^2 + 8a\varepsilon} - (\varepsilon-a) > 2a$. Let $d = C\varepsilon - \mu^*$, noting that $d > 0$ since $C > 1$ and

$\mu^* < \varepsilon$. Then by Theorem 2,

$$\mathbb{P}(\mu_k > C\varepsilon \mid Y) \leq \mathbb{P}(|\mu_k - \mu^*| > d \mid Y) = 1 - \mathbb{P}(|\mu_k - \mu^*| \leq d \mid Y) \xrightarrow{J \rightarrow \infty} 0. \quad \square$$

S2 ADDITIONAL RESULTS

S2.1 FURTHER PROPERTIES OF THE INVERSE KUMMER

We study the relationship between the hyperparameter $\gamma \in \mathbb{R}$ of the inverse Kummer distribution (Definition 1), and the first moment of the distribution. In particular, the following proposition implies that the mean of an inverse Kummer with $\gamma > 0$ is larger than the mean of the corresponding inverse gamma distribution, when $\lambda > 2$.

Proposition 1. *Let $\mu \sim \text{InvKummer}(\lambda, \beta, \gamma, \delta)$. If $\lambda > 2$, then $\mathbb{E}(\mu)$ is monotone increasing as a function of γ , for $\gamma \in (0, \infty)$.*

Proof. Fix $\varepsilon \in (0, 1)$. Let $\mu_\gamma \sim \text{InvKummer}(\lambda, \beta, \gamma, \delta)$ and $\mu_{\gamma+\varepsilon} \sim \text{InvKummer}(\lambda, \beta, \gamma + \varepsilon, \delta)$, and define $f(t) = t^{\lambda-1}(1+t)^{-\gamma}e^{-\beta t/\delta}$. Our aim is to show that $\mathbb{E}(\mu_{\gamma+\varepsilon}) \geq \mathbb{E}(\mu_\gamma)$. By Equation (11), this is true if and only if

$$\frac{\int_0^\infty \frac{1}{t(1+t)^\varepsilon} f(t) dt}{\int_0^\infty \frac{1}{(1+t)^\varepsilon} f(t) dt} \geq \frac{\int_0^\infty \frac{1}{t} f(t) dt}{\int_0^\infty f(t) dt}. \quad (\text{S13})$$

Let X be the continuous random variable on $(0, \infty)$ with probability density function $p(x) = f(x)/\int_0^\infty f(t)dt$. Then, after multiplying and dividing the left-hand side by $\int_0^\infty f(t)dt$, Equation (S13) can be written in terms of expectations as

$$\frac{\mathbb{E}\left(\frac{1}{X(1+X)^\varepsilon}\right)}{\mathbb{E}\left(\frac{1}{(1+X)^\varepsilon}\right)} \geq \mathbb{E}\left(\frac{1}{X}\right). \quad (\text{S14})$$

or equivalently,

$$\text{Cov}\left(\frac{1}{X}, \frac{1}{(1+X)^\varepsilon}\right) \geq 0. \quad (\text{S15})$$

Define $g(x) = -1/x$ and $h(x) = -1/(1+x)^\varepsilon$ for $x \in (0, \infty)$. Observe that, by transformation of random variables, $1/X \sim \text{InvKummer}(\lambda, \beta, \gamma, \delta)$. Thus, by Equation (11), since $\lambda > 2$ by assumption,

$$\begin{aligned} \mathbb{E}|g(X)|^2 &= \mathbb{E}\left(\frac{1}{X^2}\right) < \infty, \\ \mathbb{E}|h(X)|^2 &= \mathbb{E}\left(\frac{1}{(1+X)^{2\varepsilon}}\right) \leq \mathbb{E}\left(\frac{1}{X^{2\varepsilon}}\right) < \infty. \end{aligned}$$

Therefore, since $g(x)$ and $h(x)$ are monotone increasing and have finite second moments, we have $\text{Cov}(g(X), h(X)) \geq 0$ by [Schmidt \(2003\)](#), which is equivalent to Equation (S15). This completes the proof. \square

S2.2 CHARACTERIZING THE CONCENTRATION POINT OF THE INVERSE KUMMER

We characterize the point at which the inverse Kummer concentrates, μ_* , in Theorem 2. Define

$$\mu_*(\varepsilon, y, a) = \frac{2a\varepsilon}{\sqrt{(y-a+\varepsilon)^2 + 8a\varepsilon} - (y-a+\varepsilon)} \quad (\text{S16})$$

for $\varepsilon > 0$, $y \geq 0$, and $a > 0$.

Proposition 2. *For all $\varepsilon > 0$ and $a > 0$, $\mu_*(\varepsilon, y, a)$ is monotone increasing as a function of y .*

Proof. Fix $\varepsilon > 0$ and define

$$g(y) = \sqrt{(y-a+\varepsilon)^2 + 8a\varepsilon},$$

so that $\mu_*(\varepsilon, y, a) = 2a\varepsilon/(g(y) - (y-a+\varepsilon))$. Then

$$\frac{\partial \mu_*}{\partial y} = \frac{-2a\varepsilon(g'(y) - 1)}{(g(y) - (y-a+\varepsilon))^2}.$$

Differentiating g , we find that

$$g'(y) = \frac{y-a+\varepsilon}{\sqrt{(y-a+\varepsilon)^2 + 8a\varepsilon}} < 1.$$

Therefore, $\partial \mu_*/\partial y > 0$, showing that μ_* is monotone increasing as a function of y . \square

Next, we derive Equation (12) using a first-order Taylor approximation. Fix $y \geq 0$ and $a > 0$, and define

$$h(\varepsilon) = \sqrt{(y-a+\varepsilon)^2 + 8a\varepsilon} \quad (\text{S17})$$

for $\varepsilon > 0$. Differentiating and simplifying, we find that

$$h'(\varepsilon) = \frac{y+\varepsilon+3a}{h(\varepsilon)}. \quad (\text{S18})$$

Thus, $h(0) = |y-a|$ and $h'(0) = (y+3a)/|y-a|$. Hence, a first-order Taylor approximation to h at $\varepsilon = 0$ yields

$$h(\varepsilon) \approx h(0) + h'(0)\varepsilon = |y-a| + \frac{y+3a}{|y-a|}\varepsilon \quad (\text{S19})$$

when ε is small relative to $|y - a|$. Plugging this into the definition of μ_* in Equation (S16), we obtain

$$\mu_*(\varepsilon, y, a) \approx \frac{2a\varepsilon}{|y - a| + \frac{y+3a}{|y-a|}\varepsilon - (y - a + \varepsilon)}. \quad (\text{S20})$$

When $y > a$, we have $|y - a| = y - a$, so in this case Equation (S20) becomes

$$\mu_*(\varepsilon, y, a) \approx \frac{y - a}{2}.$$

Meanwhile, when $0 \leq y < a$, we have $|y - a| = a - y$, so in this case

$$\mu_*(\varepsilon, y, a) \approx \frac{\varepsilon a(a - y)}{(a - y)^2 + (a + y)\varepsilon}$$

by collecting and rearranging terms. Therefore, the first-order Taylor approximation is

$$\mu_*(\varepsilon, y, a) \approx \begin{cases} \frac{y - a}{2} & \text{if } y > a \\ \frac{\varepsilon a(a - y)}{(a - y)^2 + (a + y)\varepsilon} & \text{if } 0 \leq y < a \end{cases} \quad (\text{S21})$$

as claimed in Equation (12).

S2.3 RELATIONSHIP BETWEEN THE RELEVANCE WEIGHTS AND THE LATENT COUNTS

We further characterize the relationship between the latent counts Y_{ijk} and relevance weights μ_k . In particular, we derive the distribution of $Y_{ijk} \mid \mu_k$, integrating out r_k and θ_k in Equation (7). We show that, appealingly, this distribution has a closed-form expression in terms of hypergeometric functions. First, the distribution of $Y_{ijk} \mid \mu_k, r_k$, integrating out θ_{kj} in Equation (7), is easily seen to be $Y_{ijk} \mid \mu_k, r_k \sim \text{NegBin}(a, a/(a + r_{ik}\mu_k))$, where the negative binomial is parametrized such that the mean and variance are $\mathbb{E}(Y_{ijk} \mid \mu_k, r_k) = \mu_k r_{ik}$ and $\text{Var}(Y_{ijk} \mid \mu_k, r_k) = \mu_k r_{ik}(a + \mu_k r_{ik})$. When both r_k and θ_k are integrated out, we obtain the following result, where ${}_2F_1(a, b; c; z) = \sum_{n=0}^{\infty} \frac{(a)_n (b)_n}{(c)_n} \frac{z^n}{n!}$ denotes the Gauss-hypergeometric function and $(a)_n = \Gamma(a + n)/\Gamma(a)$ is the ascending factorial ([Abramowitz and Stegun, 1972](#)).

Proposition 3. *The probability mass function of $Y_{ijk} \mid \mu_k$ under the model in Equation (7) is*

$$\mathbb{P}(Y_{ijk} = y \mid \mu_k) = \left(\frac{\mu_k}{a}\right)^y \frac{(a)_y (\alpha)_y}{y! (\alpha I)_y} {}_2F_1\left(y + a, y + \alpha, y + \alpha I, -\frac{\mu_k}{a}\right),$$

for $y \in \{0, 1, 2, \dots\}$. Furthermore, the mean of this distribution is $\mathbb{E}(Y_{ijk} \mid \mu_k) = \mu_k/I$.

Proof. Recall that $Y_{ijk} \mid \mu_k, r_k, \theta_k \sim \text{Poisson}(r_{ik}\theta_{kj})$ and $\theta_{kj} \mid \mu_k \sim \text{Gamma}(a, a/\mu_k)$. Also, $r_{ik} \sim \text{Beta}(\alpha, \alpha I - \alpha)$ by marginalization property of the Dirichlet distribution. Hence,

$$\begin{aligned}
\pi(Y_{ijk} = y \mid \mu_k) &= \int_0^1 \left\{ \int_0^\infty \mathbb{P}(Y_{ijk} = y \mid \mu_k, r_k, \theta_k) p(\theta_{kj} \mid \mu_k) d\theta_{kj} \right\} p(r_{ik}) dr_{ik} \\
&= \int_0^1 \left\{ \int_0^\infty (r_{ik}\theta_{kj})^y \frac{e^{-r_{ik}\theta_{kj}}}{y!} \frac{(a/\mu_k)^a}{\Gamma(a)} \theta_{kj}^{a-1} e^{-(a/\mu_k)\theta_{kj}} d\theta_{kj} \right\} \frac{\Gamma(\alpha I)}{\Gamma(\alpha)\Gamma(\alpha I - \alpha)} r_{ik}^{\alpha-1} (1 - r_{ik})^{\alpha I - \alpha - 1} dr_{ik} \\
&= \frac{(a/\mu_k)^a}{y! \Gamma(a)} \frac{\Gamma(\alpha I)}{\Gamma(\alpha)\Gamma(\alpha I - \alpha)} \int_0^1 \left\{ \int_0^\infty \theta_{kj}^{y+a-1} e^{-(r_{ik}+a/\mu_k)\theta_{kj}} d\theta_{kj} \right\} r_{ik}^{y+\alpha-1} (1 - r_{ik})^{\alpha I - \alpha - 1} dr_{ik} \\
&= \frac{(a/\mu_k)^a}{y! \Gamma(a)} \frac{\Gamma(\alpha I)}{\Gamma(\alpha)\Gamma(\alpha I - \alpha)} \int_0^1 \frac{\Gamma(y+a)}{(r_{ik} + a/\mu_k)^{y+a}} r_{ik}^{y+\alpha-1} (1 - r_{ik})^{\alpha I - \alpha - 1} dr_{ik} \\
&= \frac{(a)_y (\mu_k/a)^y}{y!} \frac{\Gamma(\alpha I)}{\Gamma(\alpha)\Gamma(\alpha I - \alpha)} \int_0^1 (t\mu_k/a + 1)^{-(y+a)} t^{y+\alpha-1} (1-t)^{(y+\alpha I) - (y+\alpha) - 1} dt \\
&= \left(\frac{\mu_k}{a}\right)^y \frac{(a)_y (\alpha)_y}{y! (\alpha I)_y} {}_2F_1(y+a, y+\alpha, y+\alpha I, -\mu_k/a),
\end{aligned}$$

where $(x)_n = \Gamma(x+n)/\Gamma(x)$ is the ascending factorial, and

$${}_2F_1(a, b, c, z) = \frac{\Gamma(c)}{\Gamma(b)\Gamma(c-b)} \int_0^1 t^{b-1} (1-t)^{c-b-1} (1-tz)^{-a} dt,$$

with $c > b > 0$ and $a, z \in \mathbb{R}$ is an alternative representation of the Gauss-hypergeometric function; see [Abramowitz and Stegun \(1972\)](#). \square

This result further explains the role of each μ_k in the mutational process: it directly controls the contribution of signature k in determining the number of mutations X_{ij} in channel i for patient j , for given values of hyperparameters a and α .

S3 GIBBS SAMPLER FOR NMF WITH INFORMATIVE PRIORS

In this section, we present the general Gibbs sampler for the model in Equation (8) with priors as in Equation (9). Each step follows from simple semi-conjugate prior updates, so we omit the derivations. Note that we set the value of β_k depending on the level of sparsity of each COSMIC signature. Specifically, for a given s_k , we calculate β_k by drawing 1000 samples $\rho_k \sim \text{Dirichlet}(\beta_k s_{1k}, \dots, \beta_k s_{Ik})$ for a range of plausible β_k values (from 10 to 5000, evenly spaced on a log scale), and we select the value for which the median cosine similarity between s_k and the sampled ρ_k vectors is closest to 0.975. This ensures that all signatures have approximately equal variance under the prior. All values are available at <https://github.com/alessandrozito/CompressiveNMF>.

Inference in the CompNMF+cosmic model is performed by iterating the following steps.

1. For $i = 1, \dots, I$ and $j = 1, \dots, J$, update the latent mutation counts by drawing

$$(Y_{ij} | -) \sim \text{Multinomial}(X_{ij}, (\tilde{q}_{ij1}, \dots, \tilde{q}_{ijK_{\text{pre}}}, q_{ij1}, \dots, q_{ijK_{\text{new}}}))$$

where $\tilde{q}_{ijk} = \rho_{ik}\omega_{kj}/Q_{ij}$ and $q_{ijk} = r_{ik}\theta_{kj}/Q_{ij}$, with $Q_{ij} = \sum_{k=1}^{K_{\text{pre}}} \rho_{ik}\omega_{kj} + \sum_{k=1}^{K_{\text{new}}} r_{ik}\theta_{kj}$.

2. For $k = 1, \dots, K_{\text{pre}}$, update the COSMIC signatures by drawing

$$(\rho_k | -) \sim \text{Dirichlet}\left(\beta_k s_{1k} + \sum_{j=1}^J Y_{1jk}, \dots, \beta_k s_{Ik} + \sum_{j=1}^J Y_{Ijk}\right).$$

3. For $k = K_{\text{new}} + 1, \dots, K_{\text{pre}} + K_{\text{new}}$, update the *de novo* signatures by drawing

$$(r_k | -) \sim \text{Dirichlet}\left(\alpha + \sum_{j=1}^J Y_{1jk}, \dots, \alpha + \sum_{j=1}^J Y_{Ijk}\right).$$

4. For $k = 1, \dots, K_{\text{pre}}$ and $j = 1, \dots, J$, update the loadings associated to the COSMIC signatures by drawing

$$(\omega_{kj} | -) \sim \text{Gamma}\left(b + \sum_{i=1}^I Y_{ijk}, \frac{b}{\tau_k} + 1\right).$$

5. For $k = K_{\text{new}} + 1, \dots, K_{\text{pre}} + K_{\text{new}}$ and $j = 1, \dots, J$, update the loadings associated to the *de novo* signatures by drawing

$$(\theta_{kj} | -) \sim \text{Gamma}\left(a + \sum_{i=1}^I Y_{ijk}, \frac{a}{\mu_k} + 1\right).$$

6. For $k = 1, \dots, K_{\text{pre}}$, update the relevance weights associated to the COSMIC signatures by drawing

$$(\tau_k | -) \sim \text{InvGamma}\left(2bJ + 1, \varepsilon bJ + b \sum_{j=1}^J \omega_{kj}\right).$$

7. For $k = K_{\text{new}} + 1, \dots, K_{\text{pre}} + K_{\text{new}}$, update the relevance weights associated to the *de novo* signatures by drawing

$$(\mu_k | -) \sim \text{InvGamma}\left(2aJ + 1, \varepsilon aJ + a \sum_{j=1}^J \theta_{kj}\right).$$

One important behavior we noticed is that, occasionally, the sampler either (i) morphs a novel signature into a COSMIC one even if that COSMIC signature has been specified in the prior, or (ii) morphs an existing COSMIC signature into another COSMIC one. This is due to the multi-modal nature

of the NMF model and can be influenced by initialization. While carefully eliciting each β_k as above does help, such incoherence sometimes can hold for relatively flat COSMIC signatures, such as SBS3, SBS5, or SBS40a,b. To solve the issue, we apply a label-switching step at 2/3 of the burn-in phase, where the signatures that have not been compressed out of the model are re-matched to the COSMIC signatures via the Hungarian algorithm. This does not invalidate the MCMC algorithm since it is only performed in the burn-in phase.

S4 RATIONALE FOR THE MODEL LIKELIHOOD

S4.1 TYPES OF BASE PAIR SUBSTITUTIONS

In DNA, there are four bases: cytosine (C), thymine (T), adenine (A), and guanine (G). Considering both strands of the double helix, cytosine always pairs with guanine, and thymine always pairs with adenine. Thus, if we distinguish one of the two strands of a given DNA molecule, there are four possible base pairs at each point: C-G, G-C, T-A, and A-T.

When considering base pair substitutions at a given point, the convention is to distinguish the strand containing the pyrimidine (C or T) before the substitution has been made. Recall that cytosine (C) and thymine (T) are *pyrimidines*, whereas adenine (A) and guanine (G) are *purines*. With this convention, there are six possible types of substitutions at any given point:

| | before | after | abbreviation |
|---|--------|-------|--------------|
| 1 | C-G | A-T | C>A |
| 2 | C-G | G-C | C>G |
| 3 | C-G | T-A | C>T |
| 4 | T-A | A-T | T>A |
| 5 | T-A | C-G | T>C |
| 6 | T-A | G-C | T>G |

Sometimes, these are abbreviated denoting only the pre-substitution pyrimidine and what it changes to, as seen above.

These six classes can be further divided by considering the trinucleotide context, that is, the bases directly adjacent to the base undergoing substitution. The convention is to label the context in terms of the bases (C, T, A, or G) on the 5' and 3' sides on the strand containing the pre-substitution pyrimidine. For instance, in a substitution C>A, the C may be flanked by a T on the 5' side and a G on the 3' side:

| | before | after | abbreviation |
|--|--------|-------|--------------|
| | TCG | TAG | T[C>A]G |
| | 5' 3' | 5' 3' | |

There are $4 \times 4 = 16$ different contexts for each of the original six substitution types. Therefore, there are $16 \times 6 = 96$ single-base substitution types when the trinucleotide context is taken into account. At each position in the genome, one of the two strands contains a pyrimidine C or T, flanked by bases on the 5' and 3' sides, say, X and Y, respectively: so the trinucleotide context is either XCY or XTY. Thus, each position in the genome can be in one of $2 \times 4 \times 4 = 32$ possible states. Since there are three possible single-base substitutions at every position, we arrive at a total of $32 \times 3 = 96$ *mutational channels*.

S4.2 CONTINUOUS-TIME MARKOV PROCESS FOR SUBSTITUTIONS

Focusing on one position ℓ in the genome, let us assume mutations at ℓ occur as a time-homogeneous continuous-time Markov process, holding the neighboring bases fixed. More precisely, when the current state is a , it remains a for an $\text{Exponential}(|\Lambda_{aa}|)$ amount of time and then transitions to $b \neq a$ with probability $\Lambda_{ab}/|\Lambda_{aa}|$, where Λ is a 32×32 matrix such that (i) $\Lambda_{ab} \geq 0$ for $a \neq b$, and (ii) $\sum_b \Lambda_{ab} = 0$. This is equivalent to saying that transitions from a to b occur with rate Λ_{ab} ; thus, Λ is called the transition rate matrix. See [Lawler \(2018\)](#) for background.

Let S_ℓ^t denote the state at locus ℓ at time t , and let S_ℓ^0 be the state at ℓ for the normal (germline) genome of the individual under consideration. Let $P_{ab}^t = \mathbb{P}(S_\ell^t = b \mid S_\ell^0 = a)$ be the probability that the state is b at time t given that the state is a at time 0. From the theory of continuous-time Markov processes, we have that

$$P^t = \exp(t\Lambda) = \sum_{k=0}^{\infty} \frac{(t\Lambda)^k}{k!}$$

where $\exp(\cdot)$ denotes the matrix exponential. Since the mutation rates Λ_{ab} are very small, it is reasonable to use a first-order Taylor approximation, $P^t \approx I + t\Lambda$.

S4.3 SUBSTITUTION COUNTS ARE APPROXIMATELY POISSON DISTRIBUTED

Let a_i and b_i denote the starting and ending states, respectively, for each of the substitution types $i = 1, \dots, 96$. Let $\lambda_i = \Lambda_{a_i b_i}$, and define $X_i^t = \#\{\ell : S_\ell^0 = a_i, S_\ell^t = b_i\}$, that is, X_i^t is the number of positions in the genome that undergo substitution i , starting at state a_i at time 0 and ending at state b_i at time t .

Now, consider all of the positions ℓ that are in state a at time 0, and to simplify the math, let us assume that (a) no two of these positions are adjacent, and (b) that substitutions occur independently across positions. Of the 32 states, only four of them can be reached from a : the state can remain at a , or one of three substitutions can occur. Suppose these three substitutions are $i = 1, 2, 3$, so that the starting states are $a_1 = a_2 = a_3 = a$ and the ending states are b_1, b_2, b_3 , respectively. Let $s^0 = (s_\ell^0 : \ell = 1, \dots, L)$ be a fixed vector of starting states for all positions ℓ , and let $n = \#\{\ell : s_\ell^0 = a\}$

be the number of positions starting in state a at time 0. Let $\lambda_0 = \lambda_1 + \lambda_2 + \lambda_3$ be the sum of the rates for substitution types 1, 2, 3. By the definition of P_{ab}^t and the assumption of independence across positions, letting $X_0^t = X_1^t + X_2^t + X_3^t$, the vector $(X_1^t, X_2^t, X_3^t, n - X_0^t)$ follows a multinomial distribution. Specifically, for non-negative integers x_1, x_2, x_3 such that $x_0 := x_1 + x_2 + x_3 \leq n$, we have

$$\begin{aligned} \mathbb{P}(X_{1:3}^t = x_{1:3} \mid S^0 = s^0) &= \frac{n!}{(n-x_0)!x_1!x_2!x_3!} (P_{aa}^t)^{n-x_0} \prod_{i=1}^3 (P_{a_i b_i}^t)^{x_i} \\ &\approx \frac{n!}{(n-x_0)!x_1!x_2!x_3!} (1-t\lambda_0)^{n-x_0} (t\lambda_1)^{x_1} (t\lambda_2)^{x_2} (t\lambda_3)^{x_3} \end{aligned} \quad (\text{S22})$$

by the first-order Taylor approximation, $P^t \approx I + t\Lambda$. Since the genome is large and mutation rates are small, it is natural to assume that n is large and $t\lambda_0 = O(1/n)$. Thus, letting $c = nt\lambda_0$ we have $(1-t\lambda_0)^n = (1-c/n)^n \approx e^{-c} = \exp(-nt\lambda_0)$ and $(1-t\lambda_0)^{-x_0} = (1-c/n)^{-x_0} \approx 1$ when $x_0 \ll n$, which is the case with high probability. Plugging these approximations into Equation (S22) yields

$$\begin{aligned} &\approx \frac{n!}{(n-x_0)!x_1!x_2!x_3!} \exp(-nt\lambda_0) (t\lambda_1)^{x_1} (t\lambda_2)^{x_2} (t\lambda_3)^{x_3} \\ &= \frac{n! n^{-x_0}}{(n-x_0)!} \prod_{i=1}^3 \exp(-nt\lambda_i) \frac{(nt\lambda_i)^{x_i}}{x_i!} \end{aligned}$$

since $x_0 = x_1 + x_2 + x_3$ and $\lambda_0 = \lambda_1 + \lambda_2 + \lambda_3$ by definition. By Stirling's approximation,

$$\frac{n! n^{-x_0}}{(n-x_0)!} \sim \frac{\sqrt{2\pi n} (n/e)^n n^{-x_0}}{\sqrt{2\pi(n-x_0)} ((n-x_0)/e)^{n-x_0}} = \sqrt{\frac{n}{n-x_0}} \frac{e^{-x_0} n^n}{(n-x_0)^n} \frac{(n-x_0)^{x_0}}{n^{x_0}} \rightarrow 1$$

as $n \rightarrow \infty$ with x_0 fixed, since $(1-x_0/n)^n \rightarrow e^{-x_0}$. Hence, we have

$$\mathbb{P}(X_{1:3}^t = x_{1:3} \mid S^0 = s^0) \approx \prod_{i=1}^3 \text{Poisson}(x_i \mid nt\lambda_i)$$

when n is large, $x_0 \ll n$, and $t\lambda_0 = O(1/n)$.

For each of the 32 distinct possible starting states a , the same approximation applies to the set of positions starting in state a . Modeling these 32 sets of positions independently, we have

$$\mathbb{P}(X_{1:96}^t = x_{1:96} \mid S^0 = s^0) \approx \prod_{i=1}^{96} \text{Poisson}(x_i \mid n_i t \lambda_i)$$

where $n_i = \#\{\ell : s_\ell^0 = a_i\}$. In other words, the counts of the 96 substitution types are approximately distributed as independent Poisson random variables with rates $n_i t \lambda_i$.

The preceding derivation ignores the fact that a substitution at one position changes the context of the two adjacent positions. However, since it is rare for single-base substitutions to occur at two adjacent positions, the effect of ignoring this should be negligible.

S4.4 MULTIPLE MUTATIONAL PROCESSES

Suppose x_{ij} is the number of mutations of substitution type i for subject j , for $i = 1, \dots, I$ and $j = 1, \dots, J$, where $I = 96$. The derivation above justifies modeling these mutation counts as

$$X_{ij} \sim \text{Poisson}(n_{ij}\lambda_{ij}t_j)$$

independently, where t_j is the age or exposure time of subject j , λ_{ij} is the mutation rate for substitution type i in subject j , and n_{ij} is the number of positions that are in state a_i in the normal genome of subject j , out of all positions that were measured. The positions measured may be a subset of the genome due to whole-exome/targeted sequencing or low sequencing depth, for example.

From birth, each subject is exposed to many mutational processes, such as environmental exposures, replication errors, defective DNA repair mechanisms, and so on. Each mutational process causes each substitution type to occur at a given rate, and the profile of rates across the 96 substitution types can be expected to vary depending on the mutational process. Since rates are additive in a continuous-time Markov process, it is natural to model the subject-specific mutation rates λ_{ij} as linear combinations of these mutational process rate profiles, with non-negative weights depending on the exposure of the subject to each process. Further, assuming the opportunity counts n_{ij} are constant (or nearly constant) across all subjects j , one can absorb n_{ij} into λ_{ij} , which changes the interpretation of λ_{ij} by reparametrizing it. This leads to using a representation of the form

$$n_{ij}\lambda_{ij}t_j = \sum_{k=1}^K r_{ik}\theta_{kj}$$

where the weight $\theta_{kj} \geq 0$ is the exposure of subject j to process k , and (r_{1k}, \dots, r_{Ik}) is the mutation rate profile for mutational process k , which is referred to as its mutational signature. Thus, we arrive at the Poisson non-negative matrix factorization model in Equation (1),

$$X_{ij} \sim \text{Poisson}\left(\sum_{k=1}^K r_{ik}\theta_{kj}\right).$$

A statistical issue with this representation is that there is a non-identifiability between the r_{ik} 's and θ_{kj} 's, since arbitrary multiplicative constants c_k can be moved between them. We deal with this

by normalizing the mutational signatures to sum to 1, that is, by enforcing the constraint $\sum_{i=1}^I r_{ik} = 1$ for all k .

S5 SIMULATION DETAILS AND ADDITIONAL RESULTS

S5.1 DETAILS OF POISSONCUSP

The CUSP model described in [Legramanti et al. \(2020\)](#) is a spike-and-slab shrinkage prior that enables automatic selection of the number of latent factors in Gaussian factorization models. We adapt it to the Poisson factorization model as follows. We begin by specifying the following prior structure for the signatures and the loadings:

$$(r_{1k}, \dots, r_{Ik}) \sim \text{Dirichlet}(\alpha, \dots, \alpha), \quad \theta_{kj} = \vartheta_{kj} \mu_k, \quad \vartheta_{kj} \sim \text{Gamma}(a, a).$$

Then, we let

$$\mu_k \sim (1 - \pi_k) \text{Gamma}(a_0, b_0) + \pi_k \delta_{\mu_\infty},$$

where δ_x denotes the point mass at x and π_k is the prior probability of sampling the spike $\mu_\infty = 0.01$, which is modeled as

$$\pi_k = \sum_{\ell=1}^k \phi_\ell, \quad \phi_\ell = v_\ell \prod_{m=1}^{\ell-1} (1 - v_m), \quad v_\ell \sim \text{Beta}(1, \alpha).$$

Hence, π_k increases as k increases. As a spike-and-slab prior, this enables automatic selection of the number of signatures. Each iteration of the Gibbs sampler for the basic PoissonCUSP model consists of the following steps:

1. For $i = 1, \dots, I$ and $j = 1, \dots, J$, sample auxiliary variables $Y_{ij} = (Y_{ij1}, \dots, Y_{ijK})$ according to $Y_{ij} \sim \text{Multinomial}(X_{ij}, (q_{ij1}, \dots, q_{ijK}))$, where $q_{ijk} = r_{ik} \theta_{kj} / \sum_{\kappa=1}^K r_{i\kappa} \theta_{\kappa,j}$.
2. Sample the individual loadings ϑ_{kj} from $\vartheta_{kj} \sim \text{Gamma}(a + \sum_{i=1}^I Y_{ijk}, a + \mu_k)$.
3. Sample the signatures $r_k = (r_{1k}, \dots, r_{Ik})$ from $r_k \sim \text{Dirichlet}(\alpha + \sum_{j=1}^J Y_{1jk}, \dots, \alpha + \sum_{j=1}^J Y_{Ijk})$.
4. For each $k = 1, \dots, K$, sample the categorical auxiliary variables Z_k as follows

$$\mathbb{P}(Z_k = \ell \mid -) = \begin{cases} \phi_\ell \mu_\infty^{\sum_{i,j} Y_{ijk}} \exp(-\mu_\infty \sum_j \vartheta_{kj}) & \text{if } 1 \leq \ell \leq k \\ \phi_\ell \frac{b_0^{a_0}}{\Gamma(a_0)} \frac{\Gamma(a_0 + \sum_{ij} Y_{ijk})}{(b_0 + \sum_j \vartheta_{kj})^{a_0 + \sum_{ij} Y_{ijk}}} & \text{if } k < \ell \leq K \end{cases}$$

where $\phi_\ell = \mathbb{P}(Z_k = \ell)$ is the prior probability of the auxiliary variables. These probabilities are derived by integrating out the parameters μ_k , relying on the conjugacy of the model.

5. For $\ell = 1, \dots, K - 1$, sample the sticks from

$$v_\ell \sim \text{Beta}\left(1 + \sum_{k=1}^K \mathbb{1}(Z_k = \ell), \alpha_\pi + \sum_{k=1}^K \mathbb{1}(Z_k > \ell)\right),$$

and fix $v_K = 1$.

6. Calculate ϕ_1, \dots, ϕ_K via the stick-breaking construction, namely $\phi_\ell = v_\ell \prod_{m=1}^{\ell-1} (1 - v_m)$.

7. For $k = 1, \dots, K$, if $Z_k \leq k$ then set $\mu_k = \delta_{\mu_\infty}$, otherwise sample

$$\mu_k \sim \text{Gamma}\left(a_0 + \sum_{i=1}^I \sum_{j=1}^J Y_{ijk}, b_0 + \sum_{j=1}^J \vartheta_{kj}\right).$$

Notice that here we assume a multiplicative structure for θ_{kj} , while our compressive hyperprior approach assumes a hierarchical one. In principle, one could specify $\theta_{kj} \sim \text{Gamma}(a, a/\mu_k)$ and $\mu_k \sim (1 - \pi_k)\text{InvGamma}(a_0, b_0) + \pi_k \delta_{\mu_\infty}$ to mimic our model. However, in practice, we found that this approach did not yield the desired shrinkage effect, since the Gibbs sampler preferred to sample from the prior slab rather than allocating the signatures to the spike. This happened both when $a_0 = 1$ and $b_0 = 1$ were fixed and when they were chosen according to our compressive model as $a_0 = aJ + 1$ and $b_0 = a\mu_\infty J$ with $\mu_\infty = 0.01$. This is likely due to a mixing issue and to the strong multimodal nature of the resulting posterior. Instead, we found that the sampler described above worked better.

Inference for the number of signatures K^* in PoissonCUSP can be performed using the same adaptive Metropolis sampler as in Algorithm 2 in [Legramanti et al. \(2020\)](#). This automatically tunes the number of columns in the factorization in a random manner, by eliminating the columns that in a given iteration fall within the spike (that is, the signatures for which $Z_k \leq k$) and potentially adding novel ones. Refer to [Legramanti et al. \(2020\)](#) for a description. We implement their algorithm with their choice of tuning hyperparameters.

S5.2 ADDITIONAL SIMULATION RESULTS

We now present additional results for the simulations in Section 5. In particular, we compare the models by assessing the root mean squared error (RMSE) for (i) the observed counts X , (ii) the true mean matrix $\Lambda^0 = (\lambda_{ij}^0)$, (iii) the true matrix of signatures R^0 , and (iv) the true loadings matrix Θ^0 .

First, we consider the RMSE for X and Λ^0 , calculated as $\text{RMSE}(X, \hat{\Lambda}) = (\sum_{ij} (X_{ij} - \hat{\lambda}_{ij})^2 / IJ)^{1/2}$ and $\text{RMSE}(\Lambda^0, \hat{\Lambda}) = (\sum_{ij} (\lambda_{ij}^0 - \hat{\lambda}_{ij})^2 / IJ)^{1/2}$, respectively, where $\hat{\lambda}_{ij} = \sum_{k=1}^{K^*} \hat{r}_{ik} \hat{\theta}_{kj}$ and K^* is the




Figure S1: RMSE between (top) $\hat{\Lambda}$ and the count matrix X , and (bottom) $\hat{\Lambda}$ and the true mean matrix Λ^0 , for each method across 20 replicates, when data are generated with overdispersion $\tau = 0$, i.e., the Poisson model.

estimated number of signatures for each method. Figures S1 and S2 display the comparison across all models and replicate data sets for overdispersion $\tau = 0$ and $\tau = 0.15$, respectively. In the correctly specified case ($\tau = 0$), no method performs overwhelmingly better than all the others. This is expected since all models correctly estimate the true number of signatures in this case; see Figure 2(A). One exception is SigProfiler, which performs noticeably worse than the others, particularly in terms of the RMSE for Λ^0 . The performance of SigProfiler does improve somewhat as the sample size increases, suggesting that the algorithm might struggle in small dimensions or on relatively low mutation counts.

The situation changes, however, in the overdispersed (negative binomial) case with $\tau = 0.15$, displayed in Figure S2. Here, SignatureAnalyzer and PoissonCUSP obtain the lowest $\text{RMSE}(X, \hat{\Lambda})$ across all values of J , however, their corresponding $\text{RMSE}(\Lambda^0, \hat{\Lambda})$ shows the opposite trend. This is a clear indication of overfitting, which is reinforced by the fact that both models overestimate the number of signatures; see Figure 2(A).

To enable a direct comparison of performance in terms of estimating the true signature and loadings matrices R^0 and Θ^0 , we first perform a matching step to maximize the total pairwise cosine similarity between the estimated \hat{R} and true R^0 using the Hungarian algorithm. If the true number of signatures and the estimated number differ, we also pad the smaller matrix with zeros for the non-matched signatures, which penalizes incorrect estimation of K . The rows of $\hat{\Theta}$ and Θ^0 are also permuted accordingly, and are also padded with zeros to make the dimensions the same. We then calculate $\text{RMSE}(R^0, \hat{R}) = (\sum_{ik} (r_{ik}^0 - \hat{r}_{ik})^2 / IK)^{1/2}$ and $\text{RMSE}(\Theta^0, \hat{\Theta}) = (\sum_{kj} (\theta_{kj}^0 - \hat{\theta}_{kj})^2 / KJ)^{1/2}$, where \hat{r}_{ik} and $\hat{\theta}_{kj}$ are the estimated signatures and loadings, obtained by averaging posterior samples in the case of the Bayesian models. The results for $\tau = 0$ and $\tau = 0.15$ are shown in Figures S3 and S4, respectively. The best performance is attained by CompNMF+cosmic. While this model uses information from

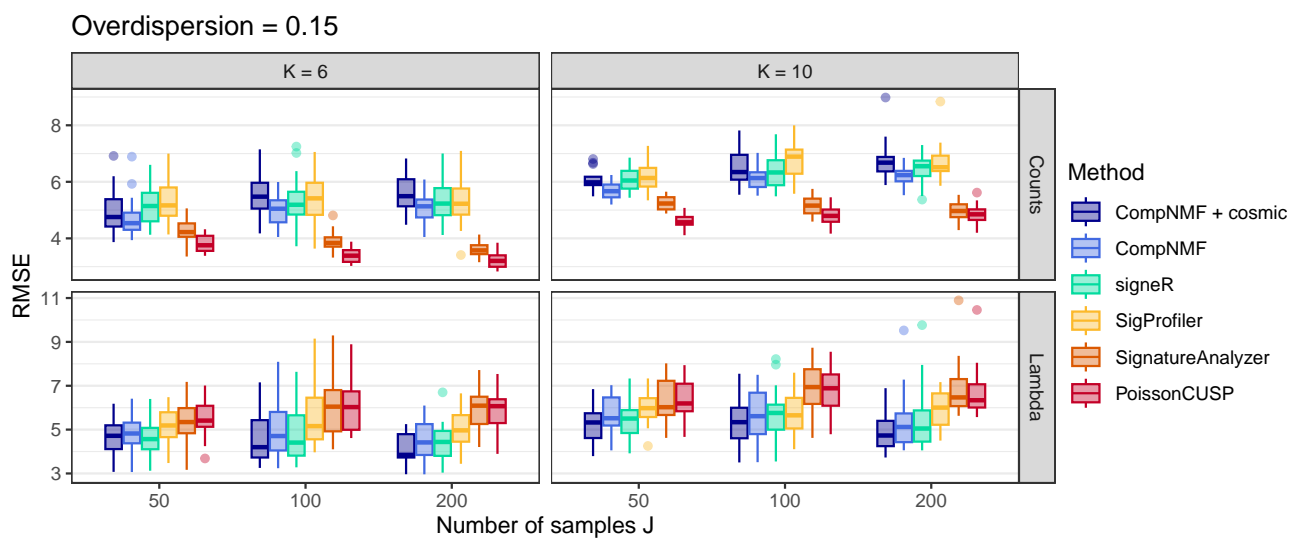


Figure S2: RMSE between (top) $\hat{\Lambda}$ and the count matrix X , and (bottom) $\hat{\Lambda}$ and the true mean matrix Λ^0 , for each method across 20 replicates, when data are generated with overdispersion be $\tau = 0.15$, i.e., negative binomial.

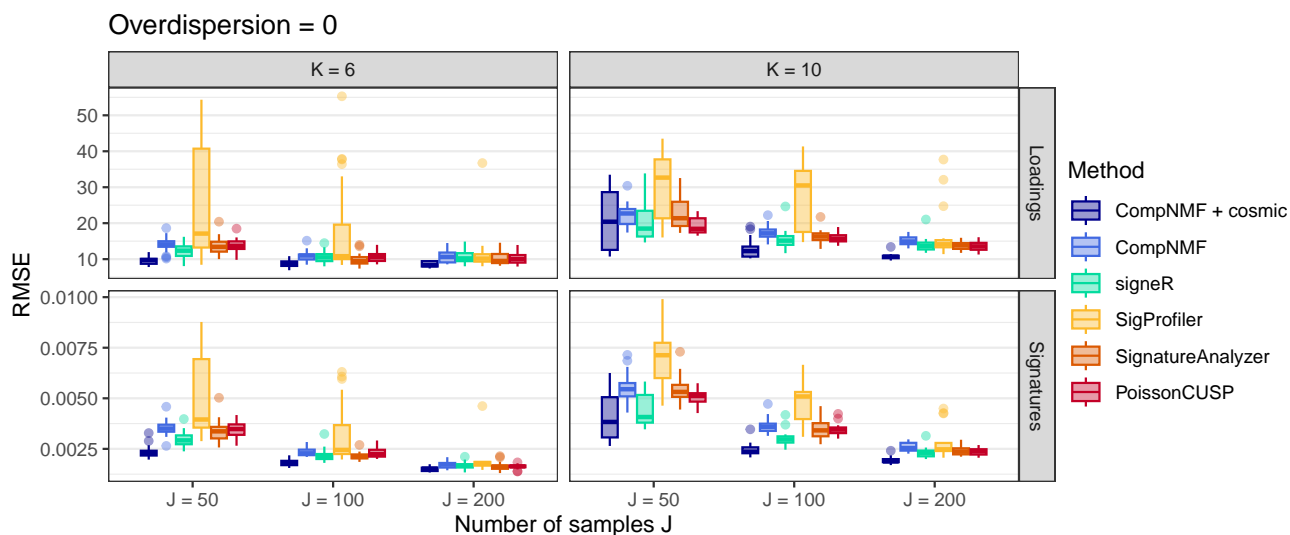


Figure S3: RMSE between (top) the estimated loadings $\hat{\Theta}$ and true loadings Θ^0 , and (bottom) the estimated signatures \hat{R} and true signatures R^0 , over 20 replicate datasets, when data are generated with overdispersion $\tau = 0$, i.e., the Poisson model.




Figure S4: RMSE between (top) the estimated loadings $\hat{\Theta}$ and true loadings Θ^0 , and (bottom) the estimated signatures \hat{R} and true signatures R^0 , over 20 replicate datasets, when data are generated with overdispersion $\tau = 0.15$, i.e., negative binomial.

the true COSMIC signatures in the prior, it correctly filters out the signatures that are not needed. Moreover, the improved estimation of the COSMIC signatures improves the estimation of the associated loadings. The performance of all the other models, with the exception of SigProfiler, is virtually identical when $\tau = 0$. Finally, when $\tau = 0.15$, PoissonCUSP and SignatureAnalyzer perform poorly due to overfitting, as before.

S6 ADDITIONAL RESULTS FROM THE 21 BREAST CANCER DATA

This section presents the complete sets of signatures inferred by each method on the 21 breast cancer data from Section 6. The models we run are the same as in Section 5, with only the following changes:

- (i) CompNMF: $K = 15$ signatures, $\varepsilon = 0.01$.
- (ii) CompNMF+cosmic: $K_{\text{new}} = 10$, *de novo* and $K_{\text{pre}} = 67$ COSMIC v3.4 signatures, $\varepsilon = 0.01$.
- (iii) PoissonCUSP: starting at $K = 15$,
- (iv) signeR: we set `estimate_hyper = TRUE` as in Rosales et al. (2016), and let K range from 2 to 15. Unlike Rosales et al. (2016), however, we do not include the opportunity counts in the model.
- (v) SignatureAnalyzer: $K = 15$, selection method set to L1W.L2H (same as Alexandrov et al. (2020)).
- (vi) SigProfiler: K ranging from 2 to 15.

Both compressive methods and PoissonCUSP are run for 12,000 iterations, discarding the first 10,000 as burn-in.

We obtain the following results. SigProfiler estimates only three signatures (matched to SBS2, SBS3, and SBS40a), and the cosine similarity is high only for SBS3; see Figure S5. This is likely a consequence of the small sample size, which can lead to the merging of two or more signatures due to insufficient signal to distinguish them. Interestingly, SignatureAnalyzer estimates five signatures (matched to SBS1, SBS2, SBS2, SBS3, and SBS13), but there is duplication since one of them appears to be a combination of SBS2 and SBS13, and ends up being matched to SBS2; see Figure S6.

We find that signeR infers five signatures (matched to SBS1, SBS2, SBS3, SBS13, and SBS96), similar to the results in the signeR paper (Rosales et al., 2016) but with slight differences since we do not account for the opportunity matrix in order to provide a consistent comparison between methods; see Figure S7. In the signeR results, SBS1 (which is sparse) appears to have been merged with a flatter signature and thus is retrieved with lower cosine similarity. The method also infers a signature that is matched to SBS96, but with low cosine similarity.

CompNMF and PoissonCUSP estimate a larger number of signatures (six and seven, respectively); see Figures S8 and S9. In both cases, SBS1, SBS2, SBS3, and SBS13 are inferred with cosine similarities comparable to the other methods. They also both infer SBS34 and SBS98, but these have larger credible intervals, indicating greater uncertainty. SBS98 also has particularly low cosine similarity in both cases, suggesting that it may be spurious. PoissonCUSP also estimates a signature matched to SBS9, but again with relatively high uncertainty and low cosine similarity. SBS9 has been found in other breast cancer types (Alexandrov et al., 2020), but its current hypothesized etiology (polymerase eta somatic hypermutation in lymphoid cells) has not been validated.

Finally, our CompNMF+cosmic model estimates eight signatures, all with cosine similarity near 1, except for SBS98; see Figure S10. The estimated signature matched to SBS98 has very high uncertainty and low cosine similarity, suggesting it is probably a spurious match. As in the simulations, the informative prior appears to provide significantly improved sensitivity to detect the presence of signatures, while still allowing for departures from the COSMIC signatures.

The loadings estimated by PoissonCUSP, SignatureAnalyzer, SigProfiler, and signeR are displayed in Figure S11 as percentages.

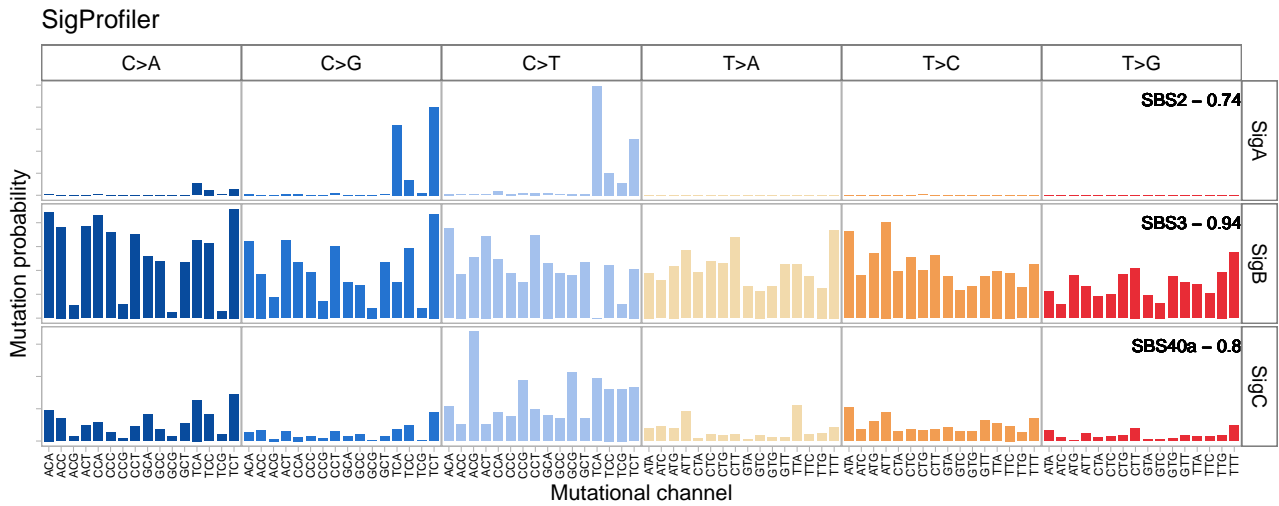


Figure S5: Mutational signatures inferred by SigProfiler on the 21 breast cancer dataset.




Figure S6: Mutational signatures inferred by SignatureAnalyzer on the 21 breast cancer dataset.




Figure S7: Mutational signatures inferred by signeR on the 21 breast cancer dataset.

PoissonCUSP




Figure S8: Mutational signatures inferred by PoissonCUSP on the 21 breast cancer dataset.

CompNMF




Figure S9: Mutational signatures inferred by CompNMF on the 21 breast cancer dataset.

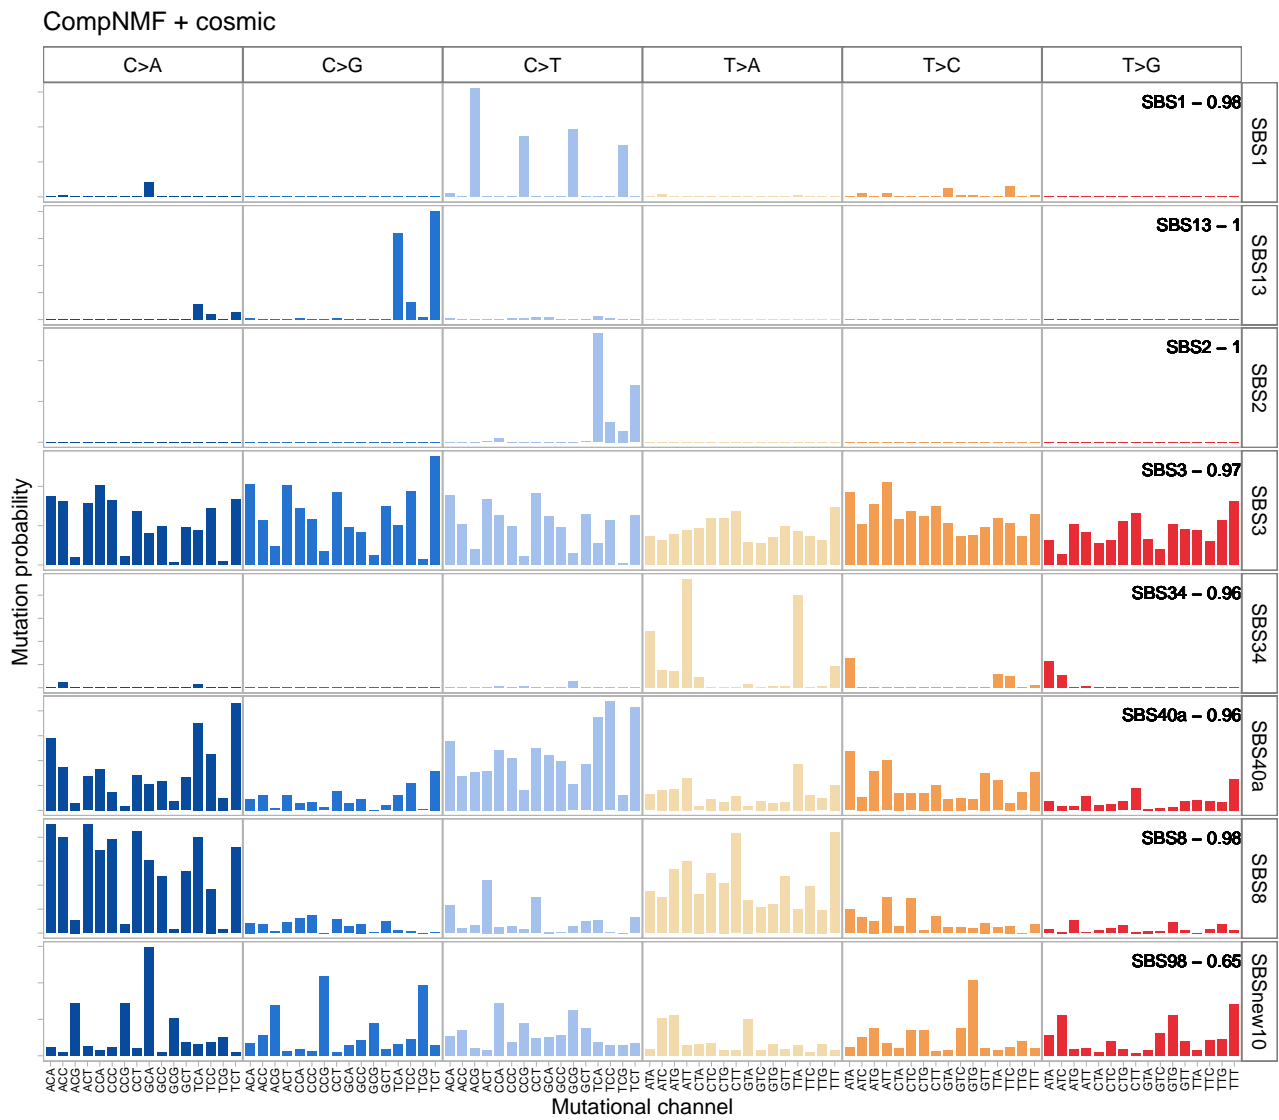


Figure S10: Mutational signatures inferred by CompNMF+cosmic on the 21 breast cancer dataset.

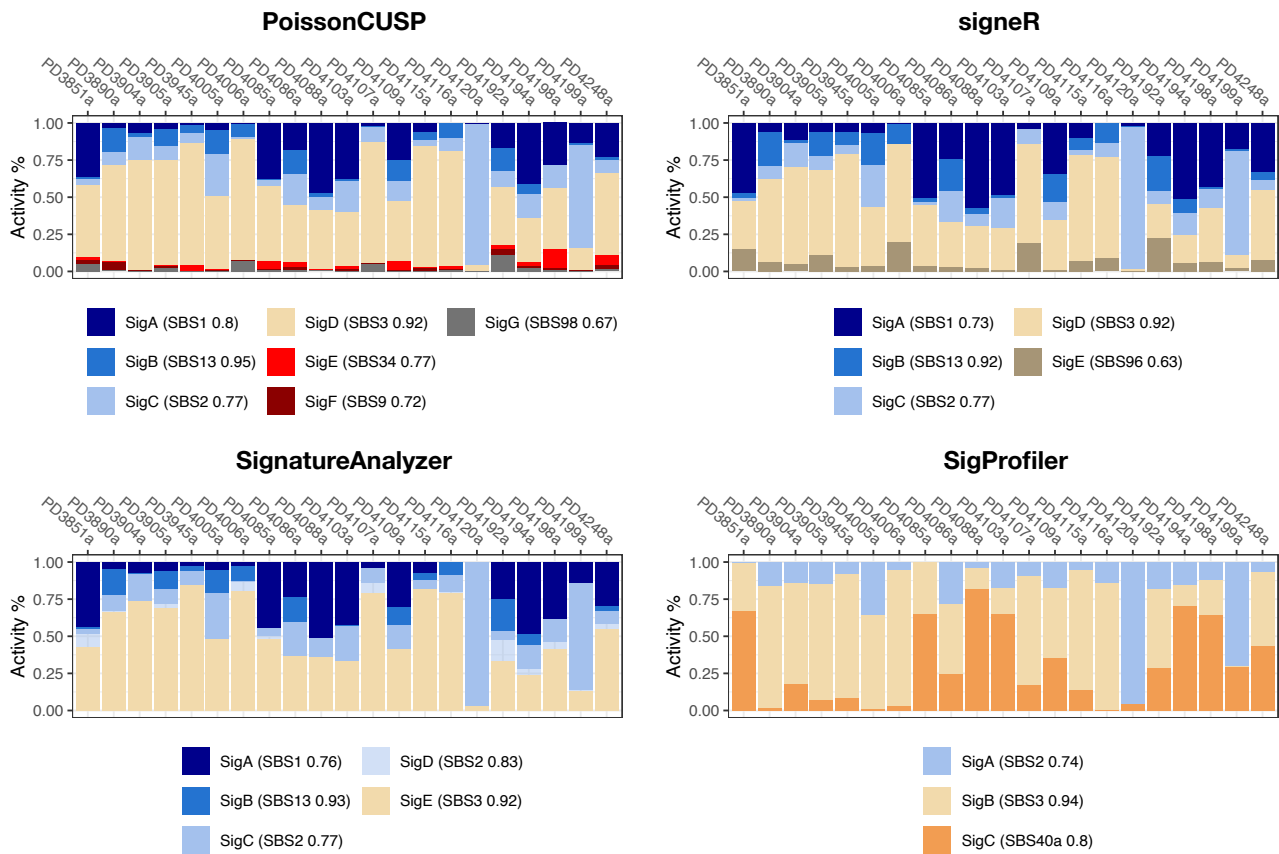


Figure S11: Loading of each signature (as a percentage of the total loading) for each patient, for PoissonCUSP, signeR, SignatureAnalyzer, and SigProfiler on the 21 breast cancer dataset.

REFERENCES

Abramowitz, M. and I. A. Stegun (1972). *Handbook of Mathematical Functions with formulas, graphs, and mathematical tables*. U.S. Dept. of Commerce, National Bureau of Standards.

Alexandrov, L. B., J. Kim, N. J. Haradhvala, et al. (2020). The repertoire of mutational signatures in human cancer. *Nature* 578(7793), 94–101.

Lawler, G. F. (2018). *Introduction to Stochastic Processes*. CRC Press.

Legramanti, S., D. Durante, and D. B. Dunson (2020). Bayesian cumulative shrinkage for infinite factorizations. *Biometrika* 107(3), 745–752.

Rosales, R. A., R. D. Drummond, R. Valieris, E. Dias-Neto, and I. T. da Silva (2016). signeR: an empirical Bayesian approach to mutational signature discovery. *Bioinformatics* 33(1), 8–16.

Schmidt, K. D. (2003). On the covariance of monotone functions of a random variable. Technical report. Inst. für Math. Stochastik.