
Hierarchical Compound Poisson Factorization

Mehmet E. Basbug

Princeton University, 35 Olden St., Princeton, NJ 08540 USA

MEHMETBASBUG@YAHOO.COM

Barbara E. Engelhardt

Princeton University, 35 Olden St., Princeton, NJ 08540 USA

BEE@PRINCETON.EDU

Abstract

Non-negative matrix factorization models based on a hierarchical Gamma-Poisson structure capture user and item behavior effectively in extremely sparse data sets, making them the ideal choice for collaborative filtering applications. Hierarchical Poisson factorization (HPF) in particular has proved successful for scalable recommendation systems with extreme sparsity. HPF, however, suffers from a tight coupling of sparsity model (absence of a rating) and response model (the value of the rating), which limits the expressiveness of the latter. Here, we introduce hierarchical compound Poisson factorization (HCPF) that has the favorable Gamma-Poisson structure and scalability of HPF to high-dimensional extremely sparse matrices. More importantly, HCPF decouples the sparsity model from the response model, allowing us to choose the most suitable distribution for the response. HCPF can capture binary, non-negative discrete, non-negative continuous, and zero-inflated continuous responses. We compare HCPF with HPF on nine discrete and three continuous data sets and conclude that HCPF captures the relationship between sparsity and response better than HPF.

1. Introduction

Matrix factorization has been a central subject in statistics since the introduction of principal component analysis (PCA) (Pearson, 1901). The goal is to embed data into a lower dimensional space with minimal loss of information. The dimensionality reduction aspect of matrix factorization has become increasingly important in exploratory data analysis as the dimensionality of data has exploded.

One alternative to PCA, non-negative matrix factorization (NMF), was first developed for factorizing matrices for face recognition (Lee & Seung, 1999). The idea behind NMF is that the contributions of each feature to a factor are non-negative. Although the motivation behind this choice has roots in cerebral representations of objects, non-negativeness has found great appeal in applications such as collaborative filtering (Gopalan et al., 2013), document classification (Xu et al., 2003), and signal processing (Févotte et al., 2009).

In collaborative filtering, the data are highly sparse user by item response matrices. For example, the *Netflix* movie rating data set includes 480K users, 17K movies and 100M ratings, meaning that 0.988 of the matrix entries are missing. In the *donors-choose* data set, the user (donor) response to an item (project) quantifies their monetary donation to that project; this matrix includes 1.3M donors, 525K projects and 2.66M donations, meaning that 0.999996 of the matrix entries are missing.

In the collaborative filtering literature, there are two school of thoughts on how to treat missing entries. The first one assumes that entries are missing at random; that is, we observe a uniformly sampled subset of the data (Marlin et al., 2012). Factorization is done on the premise that the response value provides all the information needed. The second method assumes that matrix entries are not missing at random, but instead there is a underlying mixture model: first, a coin is flipped to determine if an entry is missing. If the entry is missing, its value is set to zero; if it is not missing, the response is drawn from a specific distribution (Marlin & Zemel, 2009). In this framework, we postulate that absence of an entry carries information about the item and the user, and this information can be exploited to improve the overall quality of the factorization. The difficult part is in representing the connection between the absence of a response (sparsity model) and the numerical value of a response (response model) (Little & Rubin, 2014). We are concerned with the problem of sparse matrix factorization where the data are not missing at random.

Extensions to the NMF hint at a model that addresses this problem. Recent work showed that the NMF objective function is equivalent to a factorized Poisson likelihood (Cemgil, 2009). The authors proposed a Bayesian treatment of the Poisson model with Gamma conjugate priors on the latent factors, laying the foundation for hierarchical Poisson factorization (HPF) (Gopalan et al., 2013). The Gamma-Poisson structure is also used in earlier work for matrix factorization because of its favorable behavior (Canny, 2004; Ma et al., 2011). Long tailed Gamma priors were found to be powerful in capturing the underlying user and item behavior in collaborative filtering problems by applying strong shrinkage to the values near zero but allowing the non-zero responses to escape shrinkage (Polson & Scott, 2010). An extension of the Poisson factorization to non-discrete data using data augmentation has been considered (Févotte et al., 2009; Févotte & Idier, 2011). Along the similar lines, the connection between beta divergences and compound Poisson Gamma distribution is exploited to develop a non-negative matrix factorization model for sparse positive data (Simsekli et al., 2013). More recent work introduced a stochastic variational inference algorithm for scalable collaborative filtering using HPF (Gopalan et al., 2013).

HPF models each factor contribution to be drawn from a Poisson distribution with a long tail gamma prior. Thanks to the additive property of Poisson, the sum of these contributions are again a Poisson random variable which is used to model the response. For a collaborative filtering problem, HPF treats missing entries as true zero responses when applied to both missing and non-missing entries (Gopalan et al., 2013). When the overwhelming majority of the matrix is missing, the posterior estimates for the Poisson factors are close to zero. This mechanism has a profound impact on the response model: When the Poisson parameter of a zero truncated Poisson (ZTP) distribution approaches zero, the ZTP converges to a degenerate distribution at 1. In other words, if we condition on the fact that an entry is not missing, the HPF predicts that the response value is 1 with a very high probability. Since the response model and sparsity model are so tightly coupled, HPF is suitable only for sparse binary matrices. When using HPF on the full matrix (i.e., missing data and responses together), one might binarize the data to improve performance of the HPF (Gopalan et al., 2013). However, binarization ignores the impact of the response model on absence. For instance, a user is more likely to watch a movie that is similar to a movie that she gave a high rating to relative to one that she rated lower. This information is ignored in the HPF model.

In this paper, we introduce hierarchical compound Poisson factorization (HCPF), which has the same Gamma-Poisson structure as the HPF model and is equally computationally tractable. HCPF differs from the HPC in that it flexibility

decouples the sparsity model from the response model, allowing the HCPF to accurately model binary, non-negative discrete, non-negative continuous and zero-inflated continuous responses in the context of extreme sparsity. Unlike HPF, the ZTP distribution does not concentrate around 1, but instead converges to the response distribution that we choose. In other words, we effectively decouple the sparsity model and the response model. Decoupling does not imply independence, but instead the ability to capture the distributional characteristics of the response more accurately in the presence of extreme sparsity. HCPF still retains the useful property of HPF that the expected non-missing response value is related to the probability of non-absence, allowing the sparsity model to exploit information from the responses.

First we generalize HPF to handle non-discrete data. In Section 2, we introduce additive exponential dispersion models, a class of probability distributions. We show that any member of the additive exponential dispersion model family, including normal, gamma, inverse Gaussian, Poisson, binomial, negative binomial, and zero-truncated Poisson, can be used for the response model. In Section 3, we prove that a compound Poisson distribution converges to its element additive EDM distribution as sparsity increases.

Section 4 describes the generative model for hierarchical compound Poisson factorization (HCPF) and the mean field stochastic variational inference (SVI) algorithm for HCPF, which allows us to fit HCPF to data sets with millions of rows and columns quickly (Gopalan et al., 2013). In Section 5, we show the favorable behavior of HCPF as compared to HPF on twelve data sets of varying size including ratings data sets (*amazon*, *movielens*, *netflix* and *yelp*), social media activity data sets (*wordpress* and *tencent*), a web activity data set (*bestbuy*), a music data set (*echonest*), a biochemistry data set (*merck*), financial data sets (*donation* and *donorschoose*), and a genomics data set (*gevadis*).

2. Exponential Dispersion Models

Additive exponential dispersion models (EDMs) are a generalization of the natural exponential family where the nonzero dispersion parameter scales the log-partition function (Jorgensen, 1997). We first give a formal definition of additive EDM, and present seven useful members (Table 1).

Definition 1. A family of distributions $\mathcal{F}_\Psi = \{p_{(\Psi, \theta, \kappa)} \mid \theta \in \Theta = \text{dom}(\Psi) \subseteq \mathbb{R}, \kappa \in \mathbb{R}_{++}\}$ is called an additive exponential dispersion model if

$$p_{(\Psi, \theta, \kappa)}(x) = \exp(x\theta - \kappa\Psi(\theta))h(x, \kappa) \quad (1)$$

where θ is the natural parameter, κ is the dispersion parameter, $\Psi(\theta)$ is the base log-partition function, and $h(x, \kappa)$ is the base measure.

Table 1. Seven common additive exponential dispersion models. Normal, gamma, inverse Gaussian, Poisson, binomial, negative binomial, and zero truncated Poisson (ZTP) distributions written in additive EDM form with the variational distribution of the Poisson variable for the corresponding compound Poisson additive EDM. The gamma distribution is parametrized with shape (a) and rate (b).

DISTRIBUTION		θ	κ	$\Psi(\theta)$	$h(x, \kappa)$	$q(n_{ui} = n) \propto$
NORMAL	$N(\mu, \sigma^2)$	$\frac{\mu}{\sigma^2}$	σ^2	$\frac{\theta^2}{2}$	$\frac{1}{\sqrt{2\pi\kappa}} \exp(-\frac{x^2}{2\kappa})$	$\exp\left\{-\frac{n^2\mu^2 + y_{ui}^2}{2n\sigma^2}\right\} \frac{\Lambda_{ui}^n}{n!\sqrt{n}}$
GAMMA	$Ga(a, b)$	$-b$	a	$-\log(-\theta)$	$x^{\kappa-1}/\Gamma(\kappa)$	$\frac{(b^a y_{ui}^a \Lambda_{ui}^n)}{\Gamma(na)n!}$
INV. GAUSSIAN	$IG(\mu, \lambda)$	$\frac{-\lambda}{2\mu^2}$	$\sqrt{\lambda}$	$-\sqrt{-2\theta}$	$\frac{\kappa}{\sqrt{2\pi x^3}} \exp(-\frac{\kappa^2}{2x})$	$\exp\left\{n\frac{\lambda}{\mu} - \frac{n^2\lambda}{2y_{ui}}\right\} \frac{\Lambda_{ui}^n}{(n-1)!}$
POISSON	$Po(\lambda)$	$\log \lambda$	1	e^θ	$\frac{\kappa^x}{x!}$	$\exp\{-n\lambda\} \frac{n^{y_{ui}} \Lambda_{ui}^n}{n!}$
BINOMIAL	$Bi(r, p)$	$\log \frac{p}{1-p}$	r	$\log(1 + e^\theta)$	$\binom{\kappa}{x}$	$\frac{(nr)!(1-p)^{nr} \Lambda_{ui}^n}{n!(nr-y_{ui})!}$
NEG. BINOMIAL	$NB(r, p)$	$\log p$	r	$-\log(1 - e^\theta)$	$\binom{x+\kappa-1}{x}$	$\frac{(y_{ui}+nr-1)!(1-p)^{nr} \Lambda_{ui}^n}{n!(nr-1)!}$
ZTP	$ZTP(\lambda)$	$\log \lambda$	1	$\log(e^{e^\theta} - 1)$	$\frac{1}{x!} \sum_{j=0}^{\kappa} (-1)^j (\kappa - j)^x \binom{\kappa}{j}$	$\left(\frac{\Lambda_{ui}}{e^\lambda - 1}\right)^n \sum_{j=0}^n \frac{(-1)^j (n-j)^{y_{ui}}}{j!(n-j)!}$

The sum of additive EDMs with a shared natural parameter and base log partition function is again an additive EDM of the same type.

Theorem 1 (Jorgensen, 1997). *Let $X_1 \dots X_M$ be a sequence of additive EDMs such that $X_i \sim p_\Psi(x; \theta, \kappa_i)$, then $X_+ = X_1 + \dots + X_M \sim p_\Psi(x; \theta, \sum_i \kappa_i)$.*

In Table 1, we use a generalized definition of the zero truncated Poisson (ZTP) where the density of the sum of κ i.i.d. random variables from an ordinary ZTP distribution can be expressed with the same formula (Springael & Van Nieuwenhuysse, 2006). The last column in Table 1 shows the variational update for the Poisson parameter of the compound Poisson additive EDM, further discussed in Section 3.

3. Compound Poisson Distributions

We start with the general definition of a compound Poisson distribution and discuss compound Poisson additive EDM distributions. We then present the decoupling theorem and explain the implications of the theorem in detail.

Definition 2. *Let N be a Poisson distributed random variable with parameter Λ , X_1, \dots, X_N be i.i.d. random variables distributed with an element distribution $p_\Psi(x; \theta, \kappa)$. Then $X_+ \doteq X_1 + \dots + X_N \sim p_\Psi(x; \theta, \kappa, \Lambda)$ is called a compound Poisson random variable.*

In general, $p_\Psi(x; \theta, \kappa, \Lambda)$ does not have a closed form expression, but it is a well defined probability distribution. The conditional form, $X_+ | N$, is usually easier to manipulate, and we can calculate the marginal distribution of X_+ by integrating out N . When $N = 0$, X_+ has a degenerate distribution at zero. Furthermore, if the element distribu-

tion is an additive EDM, we have the following theorem:

Theorem 2. *Let $X \sim p_\Psi(x; \theta, \kappa)$ be an additive EDM and $X_+ \sim p_\Psi(x; \theta, \kappa, \Lambda)$ be the compound Poisson random variable with the element random variable X , then*

$$X_+ | N = n \sim p_\Psi(x; \theta, n\kappa) \quad (2)$$

$$X_+ | N = 0 \sim \delta_0. \quad (3)$$

Theorem 2 implies that the conditional distribution of a compound Poisson additive EDM is again an additive EDM. Hence, both the conditional and the marginal densities of X_+ can be calculated easily.

We make the following remark before the decoupling theorem. Later, we show that HPF is a degenerate form of our model using this remark.

Remark 1. *Compound Poisson random variable X_+ is an ordinary Poisson random variable with parameter Λ if and only if the element distribution is a degenerate distribution at 1.*

We now present the decoupling theorem. This theorem shows that the distribution of a zero truncated compound Poisson random variable converges to its element distribution as the Poisson parameter (Λ) goes to zero.

Theorem 3. *Let $X_{++} \doteq X_+ | X_+ \neq 0$ be a zero truncated compound Poisson random variable with element random variable X . If zero is not in the support of X , then*

$$Pr(X_+ = 0) = e^{-\Lambda} \quad (4)$$

$$E[X_{++}] = \frac{\Lambda}{1 - e^{-\Lambda}} E[X] \quad (5)$$

$$X_{++} \xrightarrow{D} X \text{ as } \Lambda \rightarrow 0. \quad (6)$$

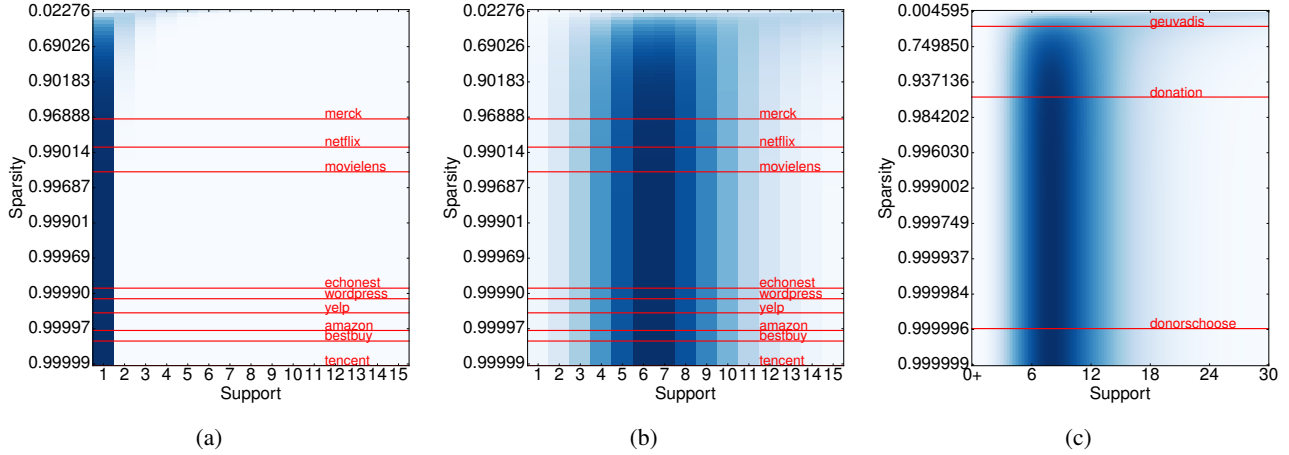


Figure 1. PDF of the zero truncated compound Poisson random variable X_{++} at various sparsity levels on log scale. The PDF is color coded, where darker colors correspond to greater density. The response distribution is a) a degenerate δ_1 , b) a zero truncated Poisson with $\lambda = 7$, c) a gamma distribution with $a = 5$, $b = 0.5$. Red vertical lines mark the sparsity levels of various data sets.

Proofs of Theorem 2, 3 and Remark 1 can be found in the Appendix.

Let X_+ be a compound Poisson variable with element random distribution $p_\Psi(x; \theta, \kappa)$ and let X_{++} be the zero truncated X_+ as in Theorem 3. We can study the probability density function (PDF) of X_{++} at various sparsity levels (Fig 1; Eq 4) with respect to the average sparsity levels of our 9 discrete data sets and 3 continuous data sets. Importantly, the element distribution is nearly identical across all levels of sparsity.

We will use X_+ to model an entry of a full sparse matrix, meaning that we are including both missing and non-missing entries. The zero truncated random variable, X_{++} , corresponds to the non-missing response. In Fig 1a, X_+ is an ordinary Poisson variable. As Remark 1 and Theorem 3 suggest, at levels of extreme sparsity (i.e., $> 90\%$ zeros), almost all of the probability mass of X_{++} concentrates at 1. That is, HPF predicts that, if an entry is not missing, then its value is 1 with a high probability. To get a more flexible response model, we might regularize X_+ with appropriate gamma priors; however, this approach would degrade the performance of the sparsity model.

On the other hand, when X_+ is a compound Poisson-ZTP random variable with $\lambda = 7$, extreme sparsity levels have virtually no effect on the distribution of the response (Fig 1b). Using the HCPF, we are free to choose any additive distribution for the response model. For discrete response data, we might opt for degenerate, Poisson, binomial, negative binomial or ZTP and for continuous response data, we might select gamma, inverse Gaussian or normal distribution. Furthermore, HCPF explicitly encodes a relationship between non-absence, $Pr(X_+ \neq 0)$, and the expected non-missing response value, $E[X_{++}]$ (Eq 4 and Eq 5), which is defined via the choice of response model.

HPF, as a degenerate HCPF model, defines this relationship as $X \sim \delta_1$ and $E[X_{++}] = 1$, which leads to the poor behavior outside of binary sparse matrices. Along with the flexibility of choosing the most natural element distribution, HCPF is capable of decoupling the sparsity and response models while still encoding a data-specific relationship between the sparsity model and the values of the non-zero responses in expectation.

4. Hierarchical Compound Poisson Factorization (HCPF)

Next, we describe the generative process for HCPF and the Gamma-Poisson structure. We explain the intuition behind the choices of the long-tailed Gamma priors. We then present the stochastic variational inference (SVI) algorithm for HCPF.

We can write the generative model of the HCPF with element distribution $p_\Psi(x; \theta, \kappa)$ with fixed hyperparameters θ and κ as follows, where C_U and C_I are the number of users and items, respectively:

- For each user $u = 1, \dots, C_U$
 1. Sample $r_u \sim Ga(\rho, \rho/\varrho)$
 2. For each component k , sample $s_{uk} \sim Ga(\eta, r_u)$
- For each item $i = 1, \dots, C_I$
 1. Sample $w_i \sim Ga(\omega, \omega/\varpi)$
 2. For each component k , sample $v_{ik} \sim Ga(\zeta, w_i)$
- For each user u and item i
 1. Sample count $n_{ui} \sim Po(\sum_k s_{uk}v_{ik})$
 2. Sample response $y_{ui} \sim p_\Psi(\theta, n_{ui}\kappa)$

Algorithm 1 SVI for HCPF

Initialize: Hyperparameters $\eta, \zeta, \rho, \varrho, \omega, \varpi, \theta, \kappa, \tau, \xi$ and parameters $t_u = \tau, t_i = \tau, b_u^r = \rho/\varrho, a_{uk}^s = \eta, b_{uk}^s = \varrho, b_i^w = \omega/\varpi, a_{ik}^v = \zeta, b_{ik}^v = \varpi$

Fix:

$$a_u^r = \rho + K\eta \quad a_i^w = \omega + K\zeta$$

repeat

Sample an observation y_{ui} uniformly from the data set
 Compute local variational parameters

$$\Lambda_{ui} = \sum_k \frac{a_{uk}^s a_{ik}^v}{b_{uk}^s b_{ik}^v}$$

$$q(n_{ui} = n) \propto \exp\{-\kappa n \Psi(\theta)\} h(y_{ui}, n\kappa) \frac{\Lambda_{ui}^n}{n!}$$

$$\varphi_{uik} \propto \exp\{\Psi(a_{uk}^s) - \log b_{uk}^s + \Psi(a_{ik}^v) - \log b_{ik}^v\}$$

Update global variational parameters

$$b_u^r = (1 - t_u^{-\xi})b_u^r + t_u^{-\xi} \left(\frac{\rho}{\varrho} + \sum_k \frac{a_{uk}^s}{b_{uk}^s} \right)$$

$$a_{uk}^s = (1 - t_u^{-\xi})a_{uk}^s + t_u^{-\xi} (\eta + C_I E[n_{ui}] \varphi_{uik})$$

$$b_{uk}^s = (1 - t_u^{-\xi})b_{uk}^s + t_u^{-\xi} \left(\frac{a_u^r}{b_u^r} + C_I \frac{a_{ik}^v}{b_{ik}^v} \right)$$

$$b_i^w = (1 - t_i^{-\xi})b_i^w + t_i^{-\xi} \left(\frac{\omega}{\varpi} + \sum_k \frac{a_{ik}^v}{b_{ik}^v} \right)$$

$$a_{ik}^v = (1 - t_i^{-\xi})a_{ik}^v + t_i^{-\xi} (\zeta + C_U E[n_{ui}] \varphi_{uik})$$

$$b_{ik}^v = (1 - t_i^{-\xi})b_{ik}^v + t_i^{-\xi} \left(\frac{a_i^w}{b_i^w} + C_U \frac{a_{uk}^s}{b_{uk}^s} \right)$$

Update learning rates

$$t_u = t_u + 1 \quad t_i = t_i + 1$$

(Optional) Update hyperparameters θ and κ

until Validation log likelihood converges

The mean field variational distribution for HCPF is given by

$$q(r_u | a_u^r, b_u^r) q(s_{uk} | a_{uk}^s, b_{uk}^s) q(w_i | a_i^w, b_i^w) q(v_{ik} | a_{ik}^v, b_{ik}^v) q(z_{ui} | \varphi_{ui}) q(n_{ui}).$$

The choice of long tail gamma priors has substantial implications for the response model in a collaborative filtering framework. The effect of the gamma prior on a particular user's responses is to effectively characterize her average response. Similarly, a gamma prior on a particular item models the average users' response for that item. The long

tail gamma prior assumption for users allows some users to have unusually high responses. For instance, in the donations data, we might expect to observe a few donors who make extraordinarily large donations to a few projects. This is not appropriate for movie ratings, since the maximum rating is 5, and a substantial number of non-missing ratings are fives. The long tail gamma prior for items allow a few items to receive unusually high average responses. This is a useful property of all of our data sets: we imagine that a few projects may attract particularly large donations, a few blogs may receive a lot of *likes*, or a few movies receive an unusually high average rating.

The choice of long tail gamma priors has different implications in terms of the sparsity model. The gamma prior on a particular user models how active she is, that is, how many items she has responses for. The long tail assumption on the sparsity model implies that there are unusually active users (e.g., cinephiles or frequent donors). The long tail assumption for items corresponds to very popular items with a large number of responses. Note that the movies with the most ratings are not necessarily the highest rated movies.

We leverage the fact that the contributions of Poisson factors can be written as a multinomial distribution (Cemgil, 2009). Using this, we can write out the stochastic variational inference algorithm for HCPF (Hoffman et al., 2013), where τ and ξ are the learning rate delay and learning rate power, and $\tau > 0$ and $0.5 < \xi < 1.0$ (Alg 1). Note that the HCPF is not a conjugate model due to $q(n_{ui})$. For other variational updates, we only need the statistics $E[n_{ui}]$. For that purpose, we calculate $q(n_{ui} = n)$ explicitly for $n = 1, \dots, N_{tr}$. The choice for the truncation value, N_{tr} , depends on θ, κ , and Λ_{ui} . We set N_{tr} using the expected range of Λ_{ui} and y_{ui} as well as fixed θ and κ .

HCPF reduces to HPF when we set $q(n_{ui}) = \delta_{y_{ui}}$. The specific form of $q(n_{ui})$ for different additive EDMs is given in Table 1.

5. Results

5.1. Data sets for collaborative filtering

We performed matrix factorization on 12 different data sets with different levels of sparsity, response characteristics, and sizes (Table 2). The rating data sets include *amazon* fine food ratings (McAuley & Leskovec, 2013), *movielens* (Harper & Konstan, 2015), *netflix* (Bell & Koren, 2007) and *yelp*, where the responses are a star rating from 1 to 5. The only exception is *movielens* where the maximum rating is 10. The social media data sets include *wordpress* and *tencent* (Niu et al., 2012), where the response is the number of likes, a non-negative integer value. Commercial data sets include *bestbuy*, where the response is the number of user visit to a product page. The biochem-

istry data sets include *merck* (Ma et al., 2015), which captures molecules (users) and chemical characteristics (items) where the response is the chemical activity. In *echonest* (Bertin-Mahieux et al., 2011), the response is the number of times a user listened a song. The donation data sets *donation* and *donorschoose* includes donors and projects, where the response is the total amount of a donation in US dollars. The genomics data set *geuvadis* includes genes and individuals, where the response is the gene expression level for a user of a gene (Lappalainen et al., 2013). Both *bestbuy* and *merck* are nearly binary matrices, meaning that the vast majority of the non-missing entries are one. On the other hand, *donation*, *donorschoose*, and *geuvadis* have a continuous response variable.

5.2. Experimental setup

We held out 20% and 1% of the non-missing entries for testing (\mathcal{Y}_{NM}^{test}) and validation, respectively. We also sampled an equal number of missing entries for testing (\mathcal{Y}_M^{test}) and validation. When calculating test and validation log likelihood, the log likelihood of the missing entries is adjusted to reflect the true sparsity ratio. Test log likelihood of the missing (\mathcal{L}_M) and non-missing entries (\mathcal{L}_{NM}) as well as the test log likelihood of a non-missing entry conditioned on that it is not missing (\mathcal{L}_{CNM}) are calculated as

$$\begin{aligned}\mathcal{L}_M &= \sum_{\mathcal{Y}_M^{test}} \log Po(n = 0 \mid \Lambda_{ui}) \\ \mathcal{L}_{NM} &= \sum_{\mathcal{Y}_{NM}^{test}} \log \sum_{n=0}^{N_{tr}} p_{\Psi}(y_{ui}^{test}; \theta, n\kappa) Po(n \mid \Lambda_{ui}) \\ \mathcal{L} &= \frac{0.2(\#total\ missing)}{|\mathcal{Y}_M^{test}|} \mathcal{L}_M + \mathcal{L}_{NM} \\ \mathcal{L}_{CNM} &= \sum_{\mathcal{Y}_{NM}^{test}} \log \sum_{n=1}^{N_{tr}} p_{\Psi}(y_{ui}^{test}; \theta, n\kappa) ZTP(n \mid \Lambda_{ui}).\end{aligned}$$

In HCPF, we fix $K = 160$, $\xi = 0.7$ and $\tau = 10,000$ after an empirical study on smaller data sets. To set hyperparameters θ and κ , we use the maximum likelihood estimates of the element distribution parameters on the non-missing entries. From the number of non-missing entries in the training data set, we inferred the sparsity level, effectively estimating $E[n_{ui}]$ empirically (note that we assume $E[n_{ui}]$ is the same for every user-item pair). We then used $E[n_{ui}]$ to set the factorization hyperparameters $\eta, \zeta, \rho, \varrho, \omega, \varpi$. To create heavy tails and uninformative gamma priors, we set $\varpi = \varrho = 0.1$ and $\omega = \rho = 0.01$. We then assumed that the contribution of each factor is equal, and set $\eta = \varrho \sqrt{E[n_{ui}]/K}$ and $\zeta = \varpi \sqrt{E[n_{ui}]/K}$. When training on the non-missing entries only, we simply assume that the sparsity level is very low (0.001), and from that we set the parameters as usual, and divide the maximum like-

Table 2. **Data set characteristics** Number of rows, columns, non-missing entries, and the ratio of missing entries to the total number of entries (sparsity) for the data sets we analyzed.

DATA SET	# ROWS	# COLS	SPARSITY	# NON-MISSING
AMAZON	256,059	74,258	0.999970	568,454
MOVIELENS	138,493	26,744	0.994600	20,000,263
NETFLIX	480,189	17,770	0.988224	100,483,024
YELP	552,339	77,079	0.999947	2,225,213
WORDPRESS	86,661	78,754	0.999915	581,508
TENCENT	1,358,842	878,708	0.999991	10,618,584
BESTBUY	1,268,702	69,858	0.999979	1,862,782
MERCK	152,935	10,883	0.970524	49,059,340
ECHONEST	1,019,318	384,546	0.999877	48,373,586
DONATION	394,266	82	0.965831	1,104,687
DCHOOSE	1,282,062	525,019	0.999996	2,661,820
GEUVDIS	9,358	462	0.462122	2,325,461

hood estimate of κ by $E[n_{ui}]$. Earlier work noted that HPF is not sensitive to hyperparameter settings within reason (Gopalan et al., 2013).

To identify the best response model for HCPF, we ran SVI with all seven additive EDM distributions in Table 1. We first fit all HCPF models by sampling from the full matrix. We calculated $\mathcal{L}, \mathcal{L}_M, \mathcal{L}_{NM}$ and \mathcal{L}_{CNM} . In Section 5.3, we compare HCPF and HPF in \mathcal{L} . In the second analysis, we only used the non-missing entries for training and calculated \mathcal{L}_{NM} . In Section 5.4, we compared \mathcal{L}_{CNM} of the first analysis to \mathcal{L}_{NM} of the second analysis.

5.3. Overall performance

In this analysis we quantify how well these models capture both the sparsity and response behavior in ultra sparse matrices. In a movie ratings data set, the question becomes ‘Can we predict if a user would rate a given movie and if she does what rating she would give?’. We report the test log likelihood of all twelve data sets. We fit HCPF with normal, gamma, and inverse Gaussian as element distributions for all the data sets. In discrete data sets, we additionally fit HPF and HCPF with Poisson, binomial, negative binomial, and zero truncated Poisson element distributions.

In all ratings data sets (*amazon*, *movielens*, *netflix*, and *yelp*), HCPF significantly outperforms HPF (Table 3). The relative performance difference is even more pronounced in sparser data sets (*amazon* and *yelp*). When we break down the test log likelihood into missing and non-missing parts, we see that, in sparser data sets, the relative performance of HPF for non-missing entries is much weaker than it is for less sparse data sets. This is expected, as response coupling in HPF is stronger in sparser data sets, forcing the response variables to zero. The opposite is true for HCPF: its performance improves with increasing data sparsity.

In social media activity data (*wordpress* and *tencent*) and the music data set (*echonest*), HCPF shows a significant improvement over HPF (Table 3). Unlike the ratings data sets, we have an unbounded response variable with an ex-

Hierarchical Compound Poisson Factorization

Table 3. Test log likelihood. Per-thousand-entry test log likelihood for HCPF and HPF trained on the full matrix for twelve data sets. Discrete HCPF models and HPF are not applicable to continuous data sets (N/A). Element distribution acronyms in Table 1.

	AMAZON	MOVIELENS	NETFLIX	YELP	WORDPRESS	TENCENT	ECHONEST	BESTBUY	MERCK	DONATION	DCHOOSE	GEUVADIS
HCPF-N	-0.388	-27.480	-50.461	-0.625	-0.909	-0.148	-1.502	-0.231	-129.225	-284.178	-0.100	-4424.288
HCPF-GA	-0.402	-28.546	-51.558	-0.648	-0.804	-0.129	-1.697	-0.221	-107.569	-260.704	-0.095	-2527.897
HCPF-IG	-0.401	-27.764	-51.152	-0.657	-0.876	-0.129	-1.648	-0.215	-93.989	-254.055	-0.095	-36327.644
HCPF-PO	-0.392	-27.056	-51.332	-0.636	-0.872	-0.183	-1.447	-0.298	-136.609	N/A	N/A	N/A
HCPF-BI	-0.387	-27.781	-50.152	-0.708	-0.930	-0.174	-1.871	-0.289	-116.886	N/A	N/A	N/A
HCPF-NB	-0.404	-28.886	-55.201	-0.701	-0.818	-0.143	-1.709	-0.342	-111.613	N/A	N/A	N/A
HCPF-ZTP	-0.389	-27.932	-52.698	-0.661	-0.850	-0.170	-1.846	-0.304	-113.856	N/A	N/A	N/A
HPF	-1.453	-76.182	-88.490	-2.026	-1.378	-0.583	-3.246	-0.289	-86.492	N/A	N/A	N/A

Table 4. Non-missing test log likelihood. Per non-missing entry test log likelihood of HCPF and HPF trained on the full matrix and on the non-missing entries only (marked as ‘F’ and ‘NM,’ respectively). When trained on the full matrix, the conditional non-missing test log likelihood is reported.

		AMAZON	MOVIELENS	NETFLIX	YELP	WORDPRESS	TENCENT	ECHONEST	BESTBUY	MERCK	DONATION	DCHOOSE	GEUVADIS
HCPF-N	F	-1.692	-2.290	-1.642	-1.734	-3.026	-4.391	-3.201	-0.872	-2.855	-4.985	-6.806	-6.841
	NM	-6.302	-2.268	-1.617	-3.573	-2.860	-3.933	-2.706	-16.754	-2.403	-4.881	-7.421	-7.173
HCPF-GA	F	-1.909	-2.417	-1.708	-1.876	-1.754	-2.500	-2.049	-0.835	-2.151	-4.338	-5.387	-3.341
	NM	-4.102	-2.310	-1.649	-2.776	-2.234	-2.761	-2.053	-17.265	-2.079	-4.213	-6.906	-4.129
HCPF-IG	F	-2.077	-2.363	-1.665	-1.999	-1.440	-2.040	-1.752	-0.782	-1.685	-4.132	-5.340	-66.228
	NM	-5.965	-2.690	-3.419	-5.589	-7.017	-7.470	-6.979	-8.247	-7.474	-7.135	-6.155	-10.121
HCPF-PO	F	-1.877	-2.255	-1.756	-1.856	-2.448	-7.960	-3.017	-1.001	-3.186	N/A	N/A	N/A
	NM	-7.206	-9.819	-5.608	-6.210	-5.234	-12.633	-6.145	-3.175	-6.673	N/A	N/A	N/A
HCPF-BI	F	-1.559	-2.208	-1.586	-1.704	-2.449	-6.911	-3.085	-0.694	-2.457	N/A	N/A	N/A
	NM	-2.113	-2.214	-1.626	-1.841	-2.649	-5.480	-2.448	-1.102	-1.993	N/A	N/A	N/A
HCPF-NB	F	-2.113	-2.510	-2.078	-2.067	-1.958	-3.441	-2.272	-1.349	-2.267	N/A	N/A	N/A
	NM	-3.772	-2.526	-2.082	-2.793	-2.344	-3.272	-2.178	-3.104	-2.085	N/A	N/A	N/A
HCPF-ZTP	F	-1.865	-2.356	-1.821	-1.832	-2.283	-6.924	-3.008	-0.010	-2.323	N/A	N/A	N/A
	NM	-6.853	-5.324	-5.541	-5.741	-5.406	-12.123	-5.952	-3.002	-6.653	N/A	N/A	N/A
HPF	F	-35.666	-9.703	-4.623	-27.025	-7.167	-63.794	-22.895	-0.017	-1.535	N/A	N/A	N/A
	NM	-1.868	-2.095	-1.672	-1.940	-2.522	-6.784	-2.640	-3.305	-1.563	N/A	N/A	N/A

ponentially decaying characteristic. At first HPF might seem to be a good model for such data; however, the sparsity level is so high that the non-missing point mass concentrates at 1. This is best seen in the comparison of HPF with *HCPF-ZTP*. Although the Poisson distribution seems to capture response characteristic effectively, the test performance degrades when zero is included as part of the response model (as in HPF).

In the *bestbuy* and *merck* data sets, where the response is near binary, HPF and HCPF performances are very similar. This confirms the observation that HPF is a sufficiently good model for sparse binary data sets. In the financial data sets (*donation* and *donors-choose*), we see that the gamma and inverse Gaussian are better distributional choices than the normal as the element distribution.

5.4. Response model evaluated using test log likelihood

In this section, we investigate which model captures the response most accurately. In a movie ratings data set, the

question becomes ‘Can we predict what rating a user would give to a movie given that we know she rated that movie?’. In Table 4, we report the conditional non-missing test log likelihood of models trained on the full matrix and test log likelihood of the models trained only on the non-missing entries.

First, we note that training HPF only on the non-missing entries results in a better response model than training HPF on the full matrix. The only exception is *bestbuy* where the conditional non-missing test log likelihood is near perfect. This is due to the near binary structure of the data set. When we know if an entry is not missing, then we are fairly confident that it has a value of 1.

Secondly, we investigate if modeling the missing entries explicitly helps the response model. We compare HPF trained on the non-missing entries to HCPF trained on the full matrix. HCPF-BI outperforms HPF in all ratings data sets except for *movielens*. A similar pattern can be seen in social media and music data sets where HCPF-IG seems to

Table 5. Test AUC Test AUC values for HCPF trained on the full matrix and HPF trained on the binarized full matrix.

		AMAZON	MOVIELENS	NETFLIX	YELP	WORDPRESS	TENCENT	ECHONEST	BESTBUY	MERCK	DONATION	DCHOOSE	GEUVADIS
HCPF-N	F	0.8020	0.9854	0.9720	0.8768	0.8991	0.9167	0.9111	0.8892	0.9872	0.8804	0.6002	0.7749
HCPF-GA	F	0.7986	0.9839	0.9723	0.8651	0.8943	0.9175	0.8912	0.8893	0.9880	0.8775	0.6543	0.7740
HCPF-IG	F	0.8017	0.9856	0.9722	0.8681	0.8953	0.9136	0.8892	0.8830	0.9880	0.8773	0.5967	0.7733
HCPF-PO	F	0.8034	0.9860	0.9721	0.8731	0.8989	0.9110	0.9124	0.8859	0.9884	N/A	N/A	N/A
HCPF-BI	F	0.8004	0.9848	0.9725	0.8452	0.8985	0.9164	0.8932	0.8878	0.9885	N/A	N/A	N/A
HCPF-NB	F	0.7989	0.9844	0.9729	0.8510	0.9002	0.9119	0.8908	0.8168	0.9877	N/A	N/A	N/A
HCPF-ZTP	F	0.8025	0.9849	0.9717	0.8640	0.8999	0.9176	0.8922	0.8440	0.9876	N/A	N/A	N/A
HPF	B	0.8021	0.9855	0.9718	0.8508	0.8971	0.9146	0.8901	0.8860	0.9879	0.8780	0.6552	0.7769

be the best model.

This phenomenon can be attributed to better identification of the relationship between the sparsity model and the response model. Although HPF trained on the full matrix can also capture this relationship, the high sparsity levels force HPF to fit near-zero Poisson parameters, hurting the prediction for response. In ratings data sets, HCPF captures the notion that more likely people are to consume an item, higher their responses are. In *movielens* and *netflix* for instance, we know that the most watched movies tend to have higher ratings. In social media data sets, the relation between the reach and popularity is captured. The more followers a blog has, the more content it is likely to produce, and the more reaction it will get. A similar correlation exists for the users as well. The more active users are also the more responsive ones. In summary, HCPF can capture the relationship between the non-missingness of an entry and the actual value of the entry if it is non-missing. Any NMF algorithm that makes the missing at random assumption would miss this relationship.

One might argue that perhaps HPF is not a good model because the underlying response distribution is not Poisson like. Comparing the rows marked ‘NM’, we observe that there is some truth to this argument. In *netflix* and *yelp* data sets, *HCPF-BI* outperforms HPF; however, we also see that training *HCPF-BI* on the full matrix is even better. A similar argument can be made for *HCPF-GA* in social media data sets (*wordpress* and *tencent*). The flexibility of HCPF is a key factor to identifying the underlying response distribution. The ability to model sparsity and response at the same time gives HCPF a further edge in modeling response.

5.5. Sparsity model evaluated using AUC

In a movie ratings data set or a purchasing data set, one important question is ‘Can we predict if a user would rate a given movie or buy a certain item?’. To understand the quality of our performance on this task, we evaluated the sparsity model separately by computing the area under the ROC curve (AUC), fitting all HCPF models by sampling from the full matrix and HPF to the binarized full matrix

(Table 5). To calculate the AUC, the true label is whether the data are missing (0) or non-missing (1), and we used the estimated probability of a non-missing entry, $Pr(X_+ \neq 0)$, as the model prediction. As discussed in Section 5.3, when we fit HPF to the full matrix, we compromise performance on sparsity and response. HCPF, on the other hand, enjoys the decoupling effect while preserving the relationship in expectation (see Eq 5). In 10 of the 12 data sets, we get an improvement over HPF (Table 5); this is somewhat surprising as the HPF is specialized to this task; this illustrates the benefit of coupling the sparsity and response models in expectation. Better modeling of the response would naturally lead to a better sparsity model.

6. Conclusion

In this paper, we first proved that a zero truncated compound Poisson distribution converges to its element distribution as sparsity increases. The implication of this theorem for HPF is that the non-missing response distribution concentrates at 1, which is not a good response model unless we have a binary data set. Inspired by the convergence theorem, we introduce HCPF. Similar to HPF, HCPF has the favorable Gamma-Poisson structure to model long-tailed user and item activity. Unlike HPF, HCPF is capable of modeling binary, non-negative discrete, non-negative continuous and zero-inflated continuous data. More importantly, HCPF decouples the sparsity and response models, allowing us to specify the most suitable distribution for the non-missing response entries. We show that this decoupling effect improves the test log likelihood dramatically when compared to HPF on high-dimensional, extremely sparse matrices. HCPF also shows superior performance to HPF trained exclusively on non-missing entries in terms of modeling response. Finally, we show that HCPF is a better sparsity model than HPF, despite HPF targeting this sparsity behavior.

For future directions, we will investigate the implications of the decoupling theorem in other Bayesian settings. We will also explore hierarchical latent structures for the element distribution.

Acknowledgements

BEE was funded by NIH R00 HG006265, NIH R01 MH101822, and a Sloan Faculty Fellowship. MB was funded in part by the Princeton Innovation J. Insley Blair Pyne Fund Award. We would like to thank Robert E. Schapire for valuable discussions.

References

- Bell, R. M. and Koren, Y. Lessons from the netflix prize challenge. *ACM SIGKDD Explorations Newsletter*, 2007.
- Bertin-Mahieux, T., Ellis, D. P. W., Whitman, B., and Lamere, P. The million song dataset. In *International Society for Music Information Retrieval Conference*, 2011.
- Canny, J. Gap: a factor model for discrete data. In *ACM SIGIR Conference on Research and Development in Information Retrieval*, 2004.
- Cemgil, A. T. Bayesian inference for nonnegative matrix factorisation models. *Computational Intelligence and Neuroscience*, 2009.
- Févotte, C. and Idier, J. Algorithms for nonnegative matrix factorization with the β -divergence. *Neural Computation*, 2011.
- Févotte, C., Bertin, N., and Durrieu, J. L. Nonnegative matrix factorization with the itakura-saito divergence: With application to music analysis. *Neural Computation*, 2009.
- Gopalan, P., Hofman, J. M., and Blei, D. M. Scalable recommendation with poisson factorization. *arXiv preprint arXiv:1311.1704*, 2013.
- Harper, F. M. and Konstan, J. A. The movielens datasets: History and context. *ACM Transactions on Interactive Intelligent Systems*, 2015.
- Hoffman, M. D., Blei, D. M., Wang, C., and Paisley, J. Stochastic variational inference. *The Journal of Machine Learning Research*, 2013.
- Jorgensen, B. *The theory of dispersion models*. CRC Press, 1997.
- Lappalainen, T., Sammeth, M., Friedlnder, M.R., ACt Hoen, P., Monlong, J., Rivas, M.A., Gonzlez-Porta, M., Kurbatova, N., Griebel, T., Ferreira, P.G., and Barann, M. Transcriptome and genome sequencing uncovers functional variation in humans. *Nature*, 2013.
- Lee, D. D. and Seung, H. S. Learning the parts of objects by non-negative matrix factorization. *Nature*, 1999.
- Little, R. J. A. and Rubin, D. B. *Statistical analysis with missing data*. John Wiley & Sons, 2014.
- Ma, H., Liu, C., King, I., and Lyu, M. R. Probabilistic factor models for web site recommendation. In *ACM SIGIR Conference on Research and Development in Information Retrieval*, 2011.
- Ma, J., Sheridan, R. P., Liaw, A., Dahl, G. E., and Svetnik, V. Deep neural nets as a method for quantitative structure-activity relationships. *Journal of Chemical Information and Modeling*, 2015.
- Marlin, B., Zemel, R. S., Roweis, S., and Slaney, M. Collaborative filtering and the missing at random assumption. *arXiv preprint arXiv:1206.5267*, 2012.
- Marlin, B. M. and Zemel, R. S. Collaborative prediction and ranking with non-random missing data. In *ACM Conference on Recommender Systems*, 2009.
- McAuley, J. J. and Leskovec, J. From amateurs to connoisseurs: modeling the evolution of user expertise through online reviews. In *International Conference on World Wide Web*, 2013.
- Niu, Y., Wang, Y., Sun, G., Yue, A., Dalessandro, B., Perlich, C., and Hamner, B. The tencent dataset and kddcup12. In *KDD-Cup Workshop*, 2012.
- Pearson, K. Liii. on lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 1901.
- Polson, N. G and Scott, J. G. Shrink globally, act locally: Sparse bayesian regularization and prediction. *Bayesian Statistics*, 2010.
- Simsekli, U., Cemgil, A. T., and Yilmaz, Y. K. Learning the beta-divergence in tweedie compound poisson matrix factorization models. In *International Conference on Machine Learning*, 2013.
- Springael, J. and Van Nieuwenhuysse, I. *On the sum of independent zero-truncated Poisson random variables*. University of Antwerp, Faculty of Applied Economics, 2006.
- Xu, W., Liu, X., and Gong, Y. Document clustering based on non-negative matrix factorization. In *ACM SIGIR Conference on Research and Development in Informaion Retrieval*, 2003.