

A review of multivariate distributions for count data derived from the Poisson distribution

David I. Inouye,¹ Eunho Yang,² Genevera I. Allen^{3,4} and Pradeep Ravikumar^{5*}

The Poisson distribution has been widely studied and used for modeling univariate count-valued data. However, multivariate generalizations of the Poisson distribution that permit dependencies have been far less popular. Yet, real-world, high-dimensional, count-valued data found in word counts, genomics, and crime statistics, for example, exhibit rich dependencies and motivate the need for multivariate distributions that can appropriately model this data. We review multivariate distributions derived from the univariate Poisson, categorizing these models into three main classes: (1) where the marginal distributions are Poisson, (2) where the joint distribution is a mixture of independent multivariate Poisson distributions, and (3) where the node-conditional distributions are derived from the Poisson. We discuss the development of multiple instances of these classes and compare the models in terms of interpretability and theory. Then, we empirically compare multiple models from each class on three real-world datasets that have varying data characteristics from different domains, namely traffic accident data, biological next generation sequencing data, and text data. These empirical experiments develop intuition about the comparative advantages and disadvantages of each class of multivariate distribution that was derived from the Poisson. Finally, we suggest new research directions as explored in the subsequent Discussion section.

© 2017 Wiley Periodicals, Inc.

How to cite this article:

WIREs Comput Stat 2017, 9:e1398. doi: 10.1002/wics.1398

Keywords: Poisson, Multivariate, Graphical models, Copulas, High dimensional

Additional Supporting Information may be found in the online version of this article.

*Correspondence to: pradeep@cs.cmu.edu

¹Department of Computer Science, The University of Texas at Austin, Austin, TX, USA

²School of Computing, Korea Advanced Institute of Science and Technology, Daejeon, Republic of Korea

³Department of Statistics, Rice University, Houston, TX, USA

⁴Department of Pediatrics-Neurology, Baylor College of Medicine, Houston, TX, USA

⁵School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA

Conflict of interest: The authors have declared no conflicts of interest for this article.

INTRODUCTION

Multivariate count-valued data has become increasingly prevalent in modern big data settings. Variables in such data are rarely independent and instead exhibit complex positive and negative dependencies. We highlight three examples of multivariate count-valued data that exhibit rich dependencies: text analysis, genomics, and crime statistics. In text analysis, a standard way to represent documents is to merely count the number of occurrences of each word in the vocabulary and create a word-count vector for each document. This representation is often known as the bag-of-words representation, in which the word order and syntax are ignored. The vocabulary size—i.e., the number of variables

in the data—is usually much greater than 1000 unique words, and thus, a high-dimensional multivariate distribution is required. Also, words are clearly not independent. For example, if the word ‘Poisson’ appears in a document, then the word ‘probability’ is *more likely* to also appear, signifying a positive dependency. Similarly, if the word ‘art’ appears, then the word ‘probability’ is less likely to also appear, signifying a negative dependency. In genomics, RNA-sequencing technologies are used to measure gene and isoform expression levels. These technologies yield counts of reads mapped back to DNA locations, which even after normalization yield non-negative data that is highly skewed with many exact zeros. These genomics data are both high dimensional, with the number of genes measuring in the tens of thousands, and strongly dependent as genes work together in pathways and complex systems to produce particular phenotypes. In crime analysis, counts of crimes in different counties are clearly multidimensional, with dependencies between crime counts. For example, the counts of crime in adjacent counties are likely to be correlated with one another, indicating a positive dependency. While positive dependencies are probably more prevalent in crime statistics, negative dependencies might be very interesting. For example, a negative dependency between adjacent counties may suggest that a criminal gang has moved from one county to the other.

These examples motivate the need for a high-dimensional, count-valued distribution that permits rich dependencies between variables. In general, a good class of probabilistic models is a fundamental building block for many tasks in data analysis. Estimating such models from data could help answer exploratory questions such as: Which genomic pathways are altered in a disease, e.g., by analyzing genomic networks? Or which county seems to have the strongest effect, with respect to crime, on other counties? A probabilistic model could also be used in Bayesian classification to determine questions such as: Does this Twitter post display positive or negative sentiment about a particular product (fitting one model on positive posts and one model on negative posts)?

The classical model for a count-valued random variable is the univariate Poisson distribution, whose probability mass function for $x \in \{0, 1, 2, \dots\}$ is:

$$\mathbb{P}_{\text{Pois}}(x|\lambda) = \lambda^x \exp(-\lambda)/x!, \quad (1)$$

where λ is the standard mean parameter for the Poisson distribution. A trivial extension of this to a multivariate distribution would be to assume independence between variables and take the product of node-wise univariate Poisson distributions, but such a model would be ill suited for many examples of multivariate count-valued data that require rich dependence structures. We review multivariate probability models that are derived from the univariate Poisson distribution and permit nontrivial dependencies between variables. We categorize these models into three main classes based on their primary modeling assumption. The first class assumes that the univariate marginal distributions are derived from the Poisson. The second class is derived as a mixture of independent multivariate Poisson distributions. The third class assumes that the univariate conditional distributions are derived from the Poisson distribution—this last class of models can also be studied in the context of probabilistic graphical models. An illustration of each of these three main model classes can be seen in Figure 1. While these models might have been classified by primary application area or performance on a particular task, a classification based on modeling assumptions helps emphasize the core abstractions for each model class. In addition, this categorization may help practitioners from different disciplines learn from the models that have worked well in different areas. We discuss multiple instances of these classes in the later sections and highlight the strengths and weaknesses of each class. We then provide a short discussion on the differences between classes in terms of interpretability and theory. Using two different empirical measures, we empirically compare multiple models from each class on three real-world datasets that have varying data characteristics from different domains, namely traffic accident data, biological next generation sequencing data, and text data. These experiments develop intuition about the comparative advantages and disadvantages of the models and suggest new research directions as explored in the subsequent Discussion section.

Notation

\mathbb{R} denotes the set of real numbers, \mathbb{R}_+ denotes the *non-negative* real numbers, and \mathbb{R}_{++} denotes the *positive* real numbers. Similarly, \mathbb{Z} denotes the set of integers. Matrices are denoted as capital letters (e.g., X, Φ), vectors are denoted as boldface lower-case letters (e.g., \mathbf{x}, ϕ), and scalar values are nonbold lowercase letters (e.g., x, ϕ).




FIGURE 1 | (Left) The first class of Poisson generalizations is based on the assumption that the univariate marginals are derived from the Poisson. (Middle) The second class is based on the idea of mixing independent multivariate Poissons into a joint multivariate distribution. (Right) The third class is based on the assumption that the univariate conditional distributions are derived from the Poisson.

MARGINAL POISSON GENERALIZATIONS

The models in this section generalize the univariate Poisson to a multivariate distribution with the property that the marginal distributions of each variable are Poisson. This is analogous to the fact that the marginal distributions of the multivariate Gaussian distribution are univariate Gaussian distributions and thus seems like a natural constraint when extending the univariate Poisson to the multivariate case. Several historical attempts at achieving this marginal property have incidentally developed the same class of models, with different derivations.^{1–4} This marginal Poisson property can also be achieved via the more general framework of copulas.^{5–7}

Multivariate Poisson Distribution

The formulation of the multivariate Poisson^a distribution goes back to M'Kendrick² where the author used differential equations to derive the bivariate Poisson process. An equivalent but more readable interpretation to arrive at the bivariate Poisson distribution would be to use the summation of independent Poisson variables as follows¹: let y_1 , y_2 , and z be univariate Poisson variables with parameters λ_1 , λ_2 , and λ_0 , respectively. Then, by setting $x_1 = y_1 + z$ and $x_2 = y_2 + z$, (x_1, x_2) follows the bivariate Poisson distribution, and its joint probability mass is defined as:

$$\mathbb{P}_{\text{BiPoi}}(x_1, x_2 | \lambda_1, \lambda_2, \lambda_0) = \exp(-\lambda_1 - \lambda_2 - \lambda_0) \times \frac{\lambda_1^{x_1} \lambda_2^{x_2}}{x_1! 2!} \sum_{z=0}^{\min(x_1, x_2)} \binom{x_1}{z} \binom{x_2}{z} z! \left(\frac{\lambda_0}{\lambda_1 \lambda_2} \right)^z. \quad (2)$$

As the sum of independent Poissons is also Poisson (whose parameter is the sum of those of two components), the marginal distribution of x_1 (similarly x_2) is still a Poisson with the rate of $\lambda_1 + \lambda_0$. It can be easily seen that the covariance of x_1 and x_2 is λ_0 , and as a result, the correlation coefficient is somewhere between 0 and $\min\left\{\frac{\sqrt{\lambda_1 + \lambda_0}}{\sqrt{\lambda_2 + \lambda_0}}, \frac{\sqrt{\lambda_2 + \lambda_0}}{\sqrt{\lambda_1 + \lambda_0}}\right\}$.⁸ Independently, Wicksell⁴ derived the bivariate Poisson as the limit of a bivariate binomial distribution. Campbell¹ shows that the models in M'Kendrick² and Wicksell⁴ can identically be derived from the sums of three independent Poisson variables.

This approach to directly extend the Poisson distribution can be generalized further to handle the multivariate case $x \in \mathbb{Z}_+^d$, in which each variable x_i is the sum of individual Poisson y_i and the common Poisson x_0 as before. The joint probability for a multivariate Poisson is developed in Teicher³ and further considered by other works^{9–12}:

$$\mathbb{P}_{\text{MulPoi}}(x; \lambda) = \exp\left(-\sum_{i=1}^d \lambda_i\right) \left(\prod_{i=1}^d \frac{\lambda_i^{x_i}}{x_i!}\right) \times \sum_{z=0}^{\min_i x_i} \left(\prod_{i=1}^d \binom{x_i}{z}\right) z! \left(\frac{\lambda_0}{\prod_{i=1}^d \lambda_i}\right)^z. \quad (3)$$

Several studies have shown that this formulation of the multivariate Poisson can also be derived as a limiting distribution of a multivariate binomial distribution when the success probabilities are small and the

number of trials is large.^{13–15} As in the bivariate case, the marginal distribution of x_i is Poisson with parameter $\lambda_i + \lambda_0$. As λ_0 controls the covariance between *all* variables, an extremely limited set of correlations between variables is permitted.

Mahamunulu¹⁶ first proposed a more general extension of the multivariate Poisson distribution that permits a full covariance structure. This distribution has been studied further by many studies.^{13,17–20} While the form of this general multivariate Poisson distribution is too complicated to spell out for $d > 3$, its distribution can be specified by a multivariate reduction scheme. Specifically, let y_i for $i = 1, \dots, (2^d - 1)$ be independently Poisson distributed with parameter λ_i . Now, define $\mathbf{A} = [A_1, A_2, \dots, A_d]$ where A_i is a $d \times \binom{d}{i}$ matrix consisting of ones and zeros, where each column of A_i has exactly i ones with no duplicate columns. Hence, A_1 is the $d \times d$ identity matrix, and A_d is a column vector of all ones. Then, $\mathbf{x} = \mathbf{A}\mathbf{y}$ is a d -dimensional multivariate Poisson-distributed random vector with a full covariance structure. Note that the simpler multivariate Poisson distribution with constant covariance in Eq. 0.3 is a special case of this general form, where $\mathbf{A} = [A_1, A_d]$.

The multivariate Poisson distribution has not been widely used for real data applications. This is likely due to two major limitations of this distribution. First, the multivariate Poisson distribution only permits *positive* dependencies; this can easily be seen because the distribution arises as the sum of independent Poisson random variables, and hence, covariances are governed by the positive rate parameters λ_i . The assumption of positive dependencies is likely unrealistic for most real count-valued data examples. Second, the computation of probabilities and inference of parameters is especially cumbersome for the multivariate Poisson distribution; these are only computationally tractable for small d and hence not readily applicable in high-dimensional settings. Kano and Kawamura¹⁷ proposed multivariate recursion schemes for computing probabilities, but these schemes are only stable and computationally feasible for small d , thus complicating likelihood-based inference procedures. Karlis¹⁸ more recently proposed a latent variable-based EM algorithm for parameter inference of the general multivariate Poisson distribution. This approach treats every pairwise interaction as a latent variable and conducts inference over both the observed and hidden parameters. While this method is more tractable than recursion schemes, it still requires inference over $\binom{d}{2}$ latent variables and

is hence not feasible in high-dimensional settings. Overall, the multivariate Poisson distribution introduced above is appealing in that its marginal distributions are Poisson; yet, there are many modeling drawbacks, including severe restriction on the types of dependencies permitted (e.g., only positive relationships), a complicated and intractable form in high dimensions, and challenging inference procedures.

Copula Approaches

A much more general way to construct valid multivariate Poisson distributions with Poisson marginals is by pairing a *copula* distribution with Poisson marginal distributions. For continuous multivariate distributions, the use of copula distributions is founded on the celebrated Sklar's theorem: any continuous joint distribution can be decomposed into a copula and the marginal distributions, and conversely, any combination of a copula and marginal distributions gives a valid continuous joint distribution.²¹ The key advantage of such models for continuous distributions is that copulas fully specify the dependence structure, hence separating the modeling of marginal distributions from the modeling of dependencies. While copula distributions paired with continuous marginal distributions enjoy wide popularity (see e.g., Ref 22 in finance applications), copula models paired with discrete marginal distributions, such as the Poisson, are more challenging both for theoretical and computational reasons.^{23–25} However, several simplifications and recent advances have attempted to overcome these challenges.^{24–26}

Copula Definition and Examples

A copula is defined by a joint cumulative distribution function (CDF), $C(\mathbf{u}) : [0, 1]^d \rightarrow [0, 1]$, with uniform marginal distributions. As a concrete example, the Gaussian copula (see left subfigure of Figure 2 for an example) is derived from the multivariate normal distribution and is one of the most popular multivariate copulas because of its flexibility in the multidimensional case; the Gaussian copula is defined simply as:

$$C_R^{\text{Gauss}}(u_1, u_2, \dots, u_d) = H_R(H^{-1}(u_1), \dots, H^{-1}(u_d)),$$

where $H^{-1}(\cdot)$ denotes the standard normal inverse CDF, and $H_R(\cdot)$ denotes the joint CDF of a $\mathcal{N}(0, R)$ random vector, where R is a correlation matrix. A similar multivariate copula can be derived from the multivariate Student's t distribution if extreme values are important to model.²⁷




FIGURE 2 | A copula distribution (left)—which is defined over the unit hypercube and has uniform marginal distributions—paired with univariate Poisson marginal distributions for each variable (middle) defines a valid discrete joint distribution with Poisson marginals (right).

The Archimedean copulas are another family of copulas that have a *single* parameter that defines the global dependence between all variables.²⁸ One property of Archimedean copulas is that they admit an explicit form, unlike the Gaussian copula. Unfortunately, the Archimedean copulas do not directly allow for a rich dependence structure like the Gaussian because they only have one dependence parameter rather than a parameter for each pair of variables.

Pair copula constructions (PCCs)²⁹ for copulas, or vine copulas, allow combinations of different bivariate copulas to form a joint multivariate copula. PCCs define multivariate copulas that have an expressive dependency structure, like the Gaussian copula, but may also model asymmetric or tail dependencies available in Archimedean and *t* copulas. Pair copulas only use univariate CDFs, conditional CDFs, and bivariate copulas to construct a multivariate copula distribution and hence can use combinations of the Archimedean copulas described previously. The multivariate distributions can be factorized in a variety of ways using bivariate copulas to flexibly model dependencies. *Vines*, or graphical tree-like structures, denote the possible factorizations that are feasible for PCCs.³⁰

Copula Models for Discrete Data

As per Sklar's theorem, any copula distribution can be combined with marginal distribution CDFs, $\{F_i(x_i)\}_{i=1}^d$, to create a joint distribution:

$$G(x_1, x_2, \dots, x_d | \theta, F_1, \dots, F_d) = C_\theta(u_1 = F_1(x_1), \dots, u_d = F_d(x_d)).$$

If sampling from the given copula is possible, this form admits simple direct sampling from the joint

distribution (defined by the CDF $G(\cdot)$) by first sampling from the copula $\mathbf{u} \sim \text{Copula}(\theta)$ and then transforming \mathbf{u} to the target space using the inverse CDFs of the marginal distributions: $\mathbf{x} = [F_1^{-1}(u_1), \dots, F_d^{-1}(u_d)]$.

A valid multivariate discrete joint distribution can be derived by pairing a copula distribution with Poisson marginal distributions. For example, a valid joint CDF with Poisson marginals is given by

$$G(x_1, x_2, \dots, x_d | \theta) = C_\theta(F_1(x_1 | \lambda_1), \dots, F_d(x_d | \lambda_d)),$$

where $F_i(x_i | \lambda_i)$ is the Poisson CDF with mean parameter λ_i , and θ denotes the copula parameters. If we pair a Gaussian copula with Poisson marginal distributions, we create a valid joint distribution that has been widely used for generating samples of multivariate count data^{7,31,32}; an example of the Gaussian copula paired with Poisson marginals to form a discrete joint distribution can be seen in Figure 2.

Nikolouloupoulos⁵ present an excellent survey of copulas to be paired with discrete marginals by defining several desired properties of a copula (quoted from Ref 5):

1. Wide range of dependence, allowing both positive and negative dependence.
2. Flexible dependence, meaning that the number of bivariate marginals is (approximately) equal to the number of dependence parameters.
3. Computationally feasible CDF for likelihood estimation.
4. Closure property under marginalization, meaning that lower-order marginals belong to the same parametric family.

5. No joint constraints for the dependence parameters, meaning that the use of covariate functions for the dependence parameters is straightforward.

Each copula model satisfies some of these properties but not all of them. For example, Gaussian copulas satisfy properties (1), (2), and (4) but not (3) because the Gaussian CDF is not known in closed form nor (5) because the correlation parameter matrix is constrained to be in the set of positive definite matrices. Nikoloulopoulos⁵ recommends Gaussian copulas for general models and vine copulas if modeling dependence in the tails or asymmetry is needed.

Theoretical Properties of Copulas Derived from Discrete Distributions

From a theoretical perspective, a multivariate discrete distribution can be viewed as a continuous copula distribution paired with discrete marginals, but the derived copula distributions are not unique and, hence, are unidentifiable.²³ Note that this is in contrast to continuous multivariate distributions where the derived copulas are uniquely defined.³³ Because of this nonuniqueness property, Genest and Nešlehová²³ caution against performing inference on and interpreting dependencies of copulas derived from discrete distributions. A further consequence of nonuniqueness is that when copula distributions are paired with discrete marginal distributions, the copulas no longer fully specify the dependence structure, as with continuous marginals.²³ In other words, the dependencies of the joint distribution will depend in part on which marginal distributions are employed. In practice, this often means that the range of dependencies permitted with certain copula and discrete marginal distribution pairs is much more limited than the copula distribution would otherwise model. However, several studies have suggested that this nonuniqueness property does not have major practical ramifications.^{5,34}

We discuss a few common approaches used for the estimation of continuous copulas with discrete marginals.

Continuous Extension for Parameter Estimation

For the estimation of continuous copulas from data, a two-stage procedure, called Inference Function for Marginals (IFM)³⁵ is commonly used in which the marginal distributions are estimated first and then used to map the data onto the unit hypercube using the CDFs of the inferred marginal distributions. While this is straightforward for continuous

marginals, this procedure is less obvious for discrete marginal distributions when using a continuous copula. One idea is to use the continuous extension (CE) of integer variables to the continuous domain³⁶ by forming a new ‘jitter’ continuous random variable \tilde{x} :

$$\tilde{x} = x + (u - 1),$$

where u is a random variable defined on the unit interval. It is clear that this new random variable is continuous, and $\tilde{x} \leq x$. An obvious choice for the distribution of u is the uniform distribution. With this idea, inference can be performed using a surrogate likelihood by randomly projecting each discrete data point into the continuous domain and averaging over the random projections, as performed in Refs 37,38. Madsen³⁹ and Madsen and Fang⁴⁰ use the CE idea as well but generate *multiple* jittered samples $\{\tilde{x}^{(1)}, \tilde{x}^{(1)}, \dots, \tilde{x}^{(m)}\}$ for each original observation x to estimate the discrete likelihood rather than merely generating one jittered sample \tilde{x} for each original observation x as in Refs 37,38. Nikoloulopoulos²⁴ finds that CE-based methods significantly underestimate the correlation structure because the CE jitter transform operates independently for each variable instead of considering the correlation structure between the variables.

Distributional Transform for Parameter Estimation

In a somewhat different direction, Rüschendorf²⁶ proposed the use of a generalization of the CDF distribution function $F(\cdot)$ for the case with discrete variables, which they term a *distributional transform* (DT), denoted by $\tilde{F}(\cdot)$:

$$\tilde{F}(x, v) \equiv F(x) + v\mathbb{P}(x) = \mathbb{P}(X < x) + v\mathbb{P}(X = x),$$

where $v \sim \text{Uniform}(0, 1)$. Note that in the continuous case, $\mathbb{P}(X = x) = 0$, and thus, this reduces to the standard CDF for continuous distributions. One way of thinking of this modified CDF is that the random variable v adds a random jump when there are discontinuities in the original CDF. If the distribution is discrete (or more generally, if there are discontinuities in the original CDF), this transformation enables the simple proof of a theorem akin to Sklar’s theorem for discrete distributions.²⁶

Kazianka⁴¹ and Kazianka and Pilz⁴² propose using the DT from Ref 26 to develop a simple and intuitive approximation for the likelihood.

Essentially, they simply take the expected jump value of $\mathbb{E}(v) = 0.5$ (where $v \sim \text{Uniform}(0, 1)$) and thus transform the discrete data to the continuous domain through the following:

$$u_i \equiv F_i(x_i - 1) + 0.5\mathbb{P}(x_i) = 0.5(F_i(x_i - 1) + F_i(x_i)),$$

which can be seen as simply taking the average of the CDF values at $x_i - 1$ and x_i . Then, they use a continuous copula such as the Gaussian copula. Note that this is much simpler to compute than the simulated likelihood (SL) method in Ref 24 or the CE methods in Refs 37–40, which require averaging over many different random initializations.

SL for Parameter Estimation

Finally, Nikolouloupoulos²⁴ proposes a method to directly approximate the maximum likelihood estimate (MLE) by estimating a discretized Gaussian copula. Essentially, unlike the CE and DT methods, which attempt to transform discrete variables to continuous variables, the MLE for a Gaussian copula with discrete marginal distributions F_1, F_2, \dots, F_d can be formulated as estimating multivariate normal rectangular probabilities:

$$\mathbb{P}(x|\gamma, R) = \int_{\phi^{-1}[F_1(x_1 - 1|\gamma_1)]}^{\phi^{-1}[F_1(x_1|\gamma_1)]} \dots \int_{\phi^{-1}[F_d(x_d - 1|\gamma_d)]}^{\phi^{-1}[F_d(x_d|\gamma_d)]} \Phi_R(z_1, \dots, z_d) dz_1 \dots dz_d, \quad (4)$$

where γ is the marginal distribution parameters, $\phi^{-1}(\cdot)$ is the univariate standard normal inverse CDF, and $\Phi_R(\cdot \dots \cdot)$ is the multivariate normal density with correlation matrix R . Nikolouloupoulos²⁴ proposes to approximate the multivariate normal rectangular probabilities via fast simulation algorithms discussed in Ref 43. Because this method directly approximates the MLE via simulated algorithms, this method is called SL. Nikolouloupoulos²⁵ compares the DT and SL methods for small sample sizes and find that the DT method tends to overestimate the correlation structure. However, because of the computational simplicity, Nikolouloupoulos²⁵ provides some heuristics of when the DT method might work well compared to the more accurate but more computationally expensive SL method.

Vine Copulas for Discrete Distributions

Panagiotelis et al.⁴⁴ provide conditions under which a multivariate discrete distribution can be decomposed as a vine PCC copula paired with discrete marginals. In addition, Panagiotelis et al.⁴⁴ show that the

likelihood computation for vine PCCs with discrete marginals is quadratic as opposed to exponential, as would be the case for general multivariate copulas such as the Gaussian copula with discrete marginals. However, computation in truly high-dimensional settings remains a challenge as $2d(d - 1)$ bivariate copula evaluations are required to calculate the probability mass function (PMF) or likelihood of a d -variate PCC using the algorithm proposed by Panagiotelis et al.⁴⁴. These bivariate copula evaluations, however, can be coupled with some of the previously discussed computational techniques, such as CE, DT, and SL, for further computational improvements. Finally, while vine PCCs offer a very flexible modeling approach, this comes with the added challenge of selecting the vine construction and bivariate copulas,⁴⁵ which has not been well studied for discrete distributions. Overall, Nikolouloupoulos⁵ recommends using vine PCCs for the complex modeling of discrete data with tail dependencies and asymmetric dependencies.

Summary of Marginal Poisson Generalizations

We have reviewed the historical development of the multivariate Poisson, which has Poisson marginals, and then reviewed many of the recent developments of using the much more general copula framework to derive Poisson generalizations with Poisson marginals. The original multivariate Poisson models based on latent Poisson variables are limited to positive dependencies and require computationally expensive algorithms to fit. However, the estimation of copula distributions paired with Poisson marginals—although theoretically has some caveats—can be performed efficiently in practice. Simple approximations, such as the expectation under the DT, can provide nearly trivial transformations that move the discrete variables to the continuous domain in which all the tools of continuous copulas can be exploited. More complex transformations, such as the SL method²⁴ can be used if the sample size is small, or high accuracy is needed.

POISSON MIXTURE GENERALIZATIONS

Instead of directly extending univariate Poissons to the multivariate case, a separate line of work proposes to indirectly extend the Poisson based on the mixture of independent Poissons. Mixture models are often considered to provide more flexibility by allowing the parameter to vary according to a mixing

distribution. One important property of mixture models is that they can model *overdispersion*. Overdispersion occurs when the variance of the data is larger than the mean of the data—unlike in a Poisson distribution in which the mean and variance are equal. One way of quantifying dispersion is the dispersion index:

$$\delta = \frac{\sigma^2}{\mu}. \quad (5)$$

If $\delta > 1$, then the distribution is overdispersed, whereas if $\delta < 1$, then the distribution is underdispersed. In real-world data, as will be seen in the experimental section, overdispersion is more common than underdispersion. Mixture models also enable dependencies between the variables, as will be described in the following paragraphs.

Suppose we are modeling a univariate random variable x with a density of $f(x|\theta)$. Rather than assuming θ is fixed, we let θ itself be a random variable following some *mixing* distribution. More formally, a general *mixture* distribution can be defined as⁴⁶:

$$\mathbb{P}(x|g(\cdot)) = \int_{\Theta} f(x|\theta)g(\theta)d\theta, \quad (6)$$

where the parameter θ is assumed to come from the mixing distribution $g(\theta)$, and Θ is the domain of θ .

For the Poisson case, let $\lambda \in \mathbb{R}_{++}^d$ be a d -dimensional vector whose i -th element λ_i is the parameter of the Poisson distribution for x_i . Now, given some mixing distribution $g(\lambda)$, the family of Poisson mixture distributions is defined as

$$\mathbb{P}_{\text{MixedPoi}}(x) = \int_{\mathbb{R}_{++}^d} g(\lambda) \prod_{i=1}^d \mathbb{P}_{\text{Poi}}(x_i|\lambda_i) d\lambda, \quad (7)$$

where the domain of the joint distribution is any count-valued assignment (i.e., $x_i \in \mathbb{Z}_+, \forall i$). While the probability density function (Eq. (7)) has a complicated form involving a multidimensional integral (a complex, high-dimensional integral when d is large), the mean and variance are known to be expressed succinctly as

$$\mathbb{E}(x) = \mathbb{E}(\lambda), \quad (8)$$

$$\text{Var}(x) = \mathbb{E}(\lambda) + \text{Var}(\lambda). \quad (9)$$

Note that Eq. (9) implies that the variance of a mixture is always larger than the variance of a single distribution. The higher-order moments of x are also

easily represented by those of λ . Besides the moments, other interesting properties (convolutions, identifiability, etc.) of Poisson mixture distributions are extensively reviewed and studied in Ref 46.

One key benefit of Poisson mixtures is that they permit both positive as well as negative dependencies simply by properly defining $g(\lambda)$. The intuition behind these dependencies can be more clearly understood when we consider the sample generation process. Suppose we have the distribution $g(\lambda)$ in two dimensions (i.e., $d = 2$) with a strong positive dependency between λ_1 and λ_2 . Then, given a sample (λ_1, λ_2) from $g(\lambda)$, x_1 and x_2 are likely to also be positively correlated.

In an early application of the model, Arbous and Kerrich⁴⁷ constrain the Poisson parameters as the different scales of the common gamma variable λ : for $i = 1, \dots, d$, the time interval t_i is given, and λ_i is set to $t_i\lambda$. Hence, $g(\lambda)$ is a univariate gamma distribution specified by $\lambda \in \mathbb{R}_{++}$, which only allows a simple dependency structure. As another early attempt, Steyn⁴⁸ choose the multivariate normal distribution for the mixing distribution $g(\lambda)$ to provide more flexibility on the correlation structure. However, the normal distribution poses problems because λ must reside in \mathbb{R}_{++} , while the normal distribution is defined on \mathbb{R} .

One of the most popular choices for $g(\lambda)$ is the log-normal distribution because of its rich covariance structure and natural positivity constraint^b:

$$\mathcal{N}_{\log}(\lambda|\mu, \Sigma) = \frac{1}{\prod_{i=1}^d \lambda_i \sqrt{(2\pi)^d |\Sigma|}} \exp\left(-\frac{1}{2}(\log\lambda - \mu)^T \Sigma^{-1}(\log\lambda - \mu)\right). \quad (10)$$

The log-normal distribution above is parameterized by μ and Σ , which are the mean and the covariance of $(\log\lambda_1, \log\lambda_2, \dots, \log\lambda_d)$, respectively. Setting the random variable x_i to follow the Poisson distribution with parameter λ_i , we have the multivariate Poisson log-normal distribution⁴⁹ from Eq. (7):

$$\mathbb{P}_{\text{PoiLogN}}(x|\mu, \Sigma) = \int_{\mathbb{R}_{++}^d} \mathcal{N}_{\log}(\lambda|\mu, \Sigma) \prod_{i=1}^d \mathbb{P}_{\text{Poi}}(x_i|\lambda_i) d\lambda. \quad (11)$$

While the joint distribution (Eq. (11)) does not have a closed-form expression, and hence, as d increases, it becomes computationally cumbersome to work with, its moments are available in closed form as a special case of Eq. (9):

$$\begin{aligned}\alpha_i &\equiv \mathbb{E}(x_i) = \exp\left(\mu_i + \frac{1}{2}\sigma_{ii}\right), \\ \text{Var}(x_i) &= \alpha_i + \alpha_i^2(\exp(\sigma_{ii}) - 1), \\ \text{Cov}(x_i, x_j) &= \alpha_i \alpha_j (\exp(\sigma_{ij}) - 1).\end{aligned}\quad (12)$$

The correlation and the degree of overdispersion (defined as the variance divided by the mean) of the marginal distributions are strictly coupled by α and σ . Also, possible Spearman's ρ correlation values for this distribution are limited if the mean value α_i is small. To briefly explore this phenomenon, we simulated a two-dimensional Poisson log-normal model with log-normal parameters $\mu = [0, 0]$ and

$$\Sigma = 2\log(\gamma) \begin{bmatrix} 1 & \pm 0.999 \\ \pm 0.999 & 1 \end{bmatrix},$$

which corresponds to a mean vector of $[\gamma, \gamma]$ per Eq. (12) and the strongest positive and negative correlation possible between the two variables. We simulated one million samples from this distribution and found that when fixing $\gamma = 2$, Spearman's ρ values are between -0.53 and 0.58 . When fixing $\gamma = 10$, Spearman's ρ values are between -0.73 and 0.81 . Thus, for small mean values, the log-normal mixture is limited in modeling strong dependencies, but for large mean values, the log-normal mixture can model stronger dependencies. Besides the examples provided here, various Poisson mixture models from different mixing distributions are available, although limited in the applied statistical literature due to their complexities. See Ref 46 and the references therein for more examples of Poisson mixtures. Karlis and Xekalaki⁴⁶ also provide the general properties of mixtures as well as the specific properties of Poisson mixtures, such as moments, convolutions, and the posterior.

While this review focuses on modeling multivariate, count-valued responses without any extra information, several extensions of multivariate Poisson log-normal models have been proposed to provide more general correlation structures when covariates are available.^{50–55} These works formulate the mean parameter of log-normal mixing distribution, $\log \mu_i$, as a linear model on given covariates in the Bayesian framework.

In order to alleviate the computational burden of using log-normal distributions as an infinite mixing density as above, Karlis and Meligkotsidou⁵⁶ proposed an EM-type estimation for a finite mixture of $k > 1$ Poisson distributions, which still preserves similar properties, such as both positive and negative

dependencies as well as closed form moments. While Karlis and Meligkotsidou⁵⁶ consider mixing multivariate Poissons with positive dependencies, the simplified form—where the component distributions are independent Poisson distributions—is much simpler to implement using an expectation-maximization (EM) algorithm. This simple finite mixture distribution can be viewed as a middle ground between a single Poisson and a nonparametric estimation method where a Poisson is located at every training point—i.e., the number of mixtures is equal to the number of training data points ($k = n$).

The gamma distribution is another common mixing distribution for the Poisson because it is the conjugate distribution for the Poisson mean parameter λ . For the univariate case, if the mixing distribution is gamma, then the resulting univariate distribution is the well-known negative binomial distribution. The negative binomial distribution can handle overdispersion in count-valued data when the variance is larger than the mean. Unlike the Poisson log-normal mixture, the univariate gamma-Poisson mixture density—i.e., the negative binomial density—is known in closed form:

$$\mathbb{P}(x|r, p) = \frac{\Gamma(r+x)}{\Gamma(r)\Gamma(x+1)} p^r (1-p)^x.$$

As $r \rightarrow \infty$, the negative binomial distribution approaches the Poisson distribution. Thus, this can be seen as a generalization of the Poisson distribution. Note that the variance of this distribution is always larger than the Poisson distribution with the same mean value.

In a similar vein of placing a gamma prior distribution on the mean parameter λ , a log-gamma prior distribution can be placed on the log mean parameter $\theta = \log(\lambda)$. Bradley et al.⁵⁷ recently leveraged the log-gamma conjugacy to the Poisson log-mean parameter θ by introducing the Poisson log-gamma hierarchical mixture distribution. They particularly discuss the multivariate log-gamma distribution that can have a flexible dependency structure similar to the multivariate log-normal distribution and illustrate some modeling advantages over the log-normal mixture model.

Summary of Mixture Model Generalizations

Overall, mixture models are particularly helpful if there is overdispersion in the data, which is often the case for real-world data as seen in the experiments section, while also allowing for variable dependencies to be modeled implicitly through the mixing

distribution. If the data exhibits overdispersion, then the log-normal or log-gamma distributions⁵⁷ give somewhat flexible dependency structures. The principal caveat with the complex mixture of Poisson distributions is computational; exact inference of the parameters is typically computationally difficult due to the presence of latent mixing variables. However, simpler models such as the finite mixture using simple EM may provide good results in practice (see Comparison section).

CONDITIONAL POISSON GENERALIZATIONS

While the multivariate Poisson formulation in Eq. (3), as well as the distribution formed by pairing a copula with Poisson marginals, assumes that univariate *marginal distributions* are derived from the Poisson, a different line of work generalizes the univariate Poisson by assuming that the univariate *node-conditional distributions* are derived from the Poisson.^{58–63} Like the assumption of Poisson marginals in previous sections, this conditional Poisson assumption seems a different yet natural extension of the univariate Poisson distribution. The multivariate Gaussian can be seen to satisfy such a conditional property as the node-conditional distributions of a multivariate Gaussian are univariate Gaussian. One benefit of these conditional models is that they can be seen as undirected graphical models or Markov Random Fields, and they have a simple parametric form. In addition, estimating these models generally reduces to estimating simple node-wise regressions, and some of these estimators have theoretical guarantees on estimating the global graphical model structure even under high-dimensional sampling regimes, where the number of variables (d) is potentially even larger than the number of samples (n).

Background on Exponential Family Distributions

We briefly describe exponential family distributions and graphical models, which form the basis for the conditional Poisson models. Many commonly used distributions fall into this family, including Gaussian, Bernoulli, exponential, gamma, and Poisson, among others. The exponential family is specified by a vector of sufficient statistics, denoted by $T(\mathbf{x}) \equiv [T_1(\mathbf{x}), T_2(\mathbf{x}), \dots, T_m(\mathbf{x})]$, the log base measure $B(\mathbf{x})$, and the domain of the random variable \mathcal{D} . With this notation, the generic exponential family is defined as:

$$\mathbb{P}_{\text{ExpFam}}(\mathbf{x}|\eta) = \exp\left(\sum_{i=1}^m \eta_i T_i(\mathbf{x}) + B(\mathbf{x}) - A(\eta)\right)$$

$$A(\eta) = \log \int_{\mathcal{D}} \exp\left(\sum_{i=1}^m \eta_i T_i(\mathbf{x}) + B(\mathbf{x})\right) d\mu(\mathbf{x}),$$

where η is called the *natural or canonical parameters* of the distribution, μ is the Lebesgue or counting measure depending on whether \mathcal{D} is continuous or discrete, respectively, and $A(\eta)$ is called the *log partition function* or *log normalization constant* because it normalizes the distribution over the domain \mathcal{D} . Note that the sufficient statistics $\{T_i(\mathbf{x})\}_{i=1}^m$ can be any arbitrary function of \mathbf{x} ; for example, $T_i(\mathbf{x}) = x_1 x_2$ could be used to model an interaction between x_1 and x_2 . The log partition function $A(\eta)$ will be a key quantity when discussing the following models: $A(\eta)$ must be finite for the distribution to be valid so that the realizable domain of parameters is given by $\{\eta \in \mathcal{D} : A(\eta) < \infty\}$. Thus, for instance, if the realizable domain only allows positive or negative interaction terms, then the set of allowed dependencies would be severely restricted.

Let us now consider the exponential family form of the univariate Poisson:

$$\begin{aligned} \mathbb{P}_{\text{Pois}}(\mathbf{x}|\lambda) &= \lambda^x \exp(-\lambda)/x! \\ &= \exp(\log(\lambda^x) - \log(x!) - \lambda) \\ &= \exp\left(\underbrace{\log(\lambda)}_{\eta} \underbrace{x}_{T(x)} + \underbrace{(-\log(x!))}_{B(x)} - \lambda\right), \end{aligned}$$

and therefore

$$\mathbb{P}_{\text{Pois}}(\mathbf{x}|\eta) = \exp(\eta x - \log(x!) - \exp(\eta)), \quad (13)$$

where $\eta \equiv \log(\lambda)$ is the natural parameter of the Poisson, $T(x) = x$ is the Poisson sufficient statistic, $-\log(x!)$ is the Poisson log base measure, and $A(\eta) = \exp(\eta)$ is the Poisson log partition function. Note that for the general exponential family distribution, the log partition function may not have a closed form.

Background on Graphical Models

The graphical model over \mathbf{x} , given some graph \mathcal{G} —a set of nodes and edges—is a set of distributions on \mathbf{x} that satisfy the Markov independence assumptions with respect to \mathcal{G} .⁶⁴ In particular, an undirected graphical model provides a compact way to represent *conditional* independence among random variables—the *Markov* properties of the graph. Conditional independence relaxes the notion of full independence by defining which variables are independent *given* that the other variables are fixed or known.

More formally, let \mathcal{V} be a set of d nodes corresponding to the d random variables in \mathbf{x} , let \mathcal{E} be the set of undirected edges connecting nodes in \mathcal{V} , and let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be the corresponding undirected graph. By the Hammersley-Clifford theorem,⁶⁵ any such distribution has the following form:

$$\mathbb{P}(\mathbf{x}|\eta) = \exp\left(\sum_{C \in \mathcal{C}} \eta_C T_C(\mathbf{x}_C) - A(\eta)\right) \quad (14)$$

where \mathcal{C} is a set of cliques (fully-connected sub-graphs) of \mathcal{G} , and $T_C(\mathbf{x}_C)$ are the clique-wise sufficient statistics. For example, if $C = \{1, 2, 3\} \in \mathcal{C}$, then there would be a term $\eta_{1,2,3} T_{1,2,3}(x_1, x_2, x_3)$ that involves the first, second, and third random variables in \mathbf{x} . Hence, a graphical model can be understood as an exponential family distribution with the form given in Eq. (14). An important special case—which will be the focus in this paper—is a pairwise graphical model, where \mathcal{C} consists of merely \mathcal{V} and \mathcal{E} , i.e., $|\mathcal{C}| = \{1, 2\}, \forall C \in \mathcal{C}$, so that we have

$$\mathbb{P}(\mathbf{x}|\eta) = \exp\left(\sum_{i \in \mathcal{V}} \eta_i T_i(x_i) + \sum_{(i,j) \in \mathcal{E}} \eta_{ij} T_{ij}(x_i, x_j) - A(\eta)\right).$$

As graphical models provide direct interpretations on the Markov independence assumptions, for the Poisson-based graphical models in this section, we can easily investigate the conditional independence relationships between random variables rather than marginal correlations.

As an example, we will consider the Gaussian graphical model formulation of the standard multivariate normal distribution (for simplicity, we will assume the mean vector is zero, i.e., $\mu = 0$):

Standard Form \Leftrightarrow Graphical Model Form

$$\Sigma = -\frac{1}{2}\Theta^{-1} \Leftrightarrow \Theta = -\frac{1}{2}\Sigma^{-1} \quad (15)$$

$$\begin{aligned} \mathbb{P}(\mathbf{x}|\Sigma) &\propto \exp\left(-\frac{1}{2}\mathbf{x}^T \Sigma^{-1} \mathbf{x}\right) \\ &\Leftrightarrow \mathbb{P}(\mathbf{x}|\Theta) \propto \exp(\mathbf{x}^T \Theta \mathbf{x}) \\ &= \exp\left(\sum_i \theta_{ii} x_i^2 + \sum_{i \neq j} \theta_{ij} x_i x_j\right). \end{aligned} \quad (16)$$

Note how Eq. (16) is related to Eq. (14) by setting $\eta_i = \theta_{ii}$, $\eta_{ij} = \theta_{ij}$, $T_i(x_i) = x_i^2$, $T_{ij}(x_i, x_j) = x_i x_j$, and $\mathcal{E} = \{(i, j) : i \neq j, \theta_{ij} \neq 0\}$ —i.e., the edges in the graph correspond to the non-zeros in Θ . In addition, this

example shows that the marginal moments—i.e., the covariance matrix Σ —are quite different from the graphical model parameters—i.e., the negative of the inverse covariance matrix $\Theta = -\frac{1}{2}\Sigma^{-1}$. In general, for graphical models, such as the Poisson graphical models (PGMs) defined in the next section, the transformation from the covariance to the graphical model parameter (Eq. (15)) is not known in closed form; in fact, this transformation is often very difficult to compute for non-Gaussian models.⁶⁶ For more information about graphical models and exponential families, see Refs 66,67.

Poisson Graphical Model

The first to consider multivariate extensions constructed by assuming that conditional distributions are univariate exponential family distributions, such as and including the Poisson distribution, was Besag⁵⁸. In particular, suppose all node-conditional distributions—the conditional distribution of a node conditioned on the rest of the nodes—are univariate Poisson. Then, there is a unique joint distribution consistent with these node-conditional distributions under some conditions, and moreover, this joint distribution is a graphical model distribution that factors according to a graph specified by the node-conditional distributions. In fact, this approach can be uniformly applicable for any exponential family beyond the Poisson distribution and can be extended to more general graphical model settings^{61,63} beyond the pairwise setting in Ref 58. The particular instance of the univariate Poisson as the exponential family underlying the node-conditional distributions is called a PGM.^c

Specifically, suppose that for every $i \in \{1, \dots, d\}$, the node-conditional distribution is specified by univariate Poisson distribution in exponential family form as specified in Eq. (13):

$$\mathbb{P}(x_i|\mathbf{x}_{-i}) = \exp\{\psi(\mathbf{x}_{-i})x_i - \log(x_i!) - \exp(\psi(\mathbf{x}_{-i}))\}, \quad (17)$$

where \mathbf{x}_{-i} is the set of all x_j except x_i , and the function $\psi(\mathbf{x}_{-i})$ is *any* function that depends on the rest of all the random variables except x_i . Furthermore, suppose that the corresponding joint distribution on \mathbf{x} factors according to the set of cliques \mathcal{C} of a graph \mathcal{G} . Yang et al.⁶³ then show that such a joint distribution consistent with the above node-conditional distributions exists and, moreover, necessarily has the form

$$\mathbb{P}(\mathbf{x}|\eta) = \exp\left\{\sum_{C \in \mathcal{C}} \eta_C \prod_{i \in C} x_i - \sum_{i=1}^d \log(x_i!) - A(\eta)\right\}, \quad (18)$$

where the function $A(\eta)$ is the log-partition function on all parameters $\eta = \{\eta_C\}_{C \in \mathcal{C}}$. The pairwise PGM, as a special case, is defined as follows:

$$\mathbb{P}_{\text{PGM}}(\mathbf{x}|\eta) = \exp \left\{ \sum_{i=1}^d \eta_i x_i + \sum_{(i,j) \in \mathcal{E}} \eta_{ij} x_i x_j - \sum_{i=1}^d \log(x_i!) - A_{\text{PGM}}(\eta) \right\}, \quad (19)$$

where \mathcal{E} is the set of edges of the graphical model, and $\eta = \{\eta_1, \eta_2, \dots, \eta_d\} \cup \{\eta_{ij}, \forall (i, j) \in \mathcal{E}\}$. For notational simplicity and development of extensions to PGM, we will gather the node parameters η_i into a vector $\theta = [\eta_1, \eta_2, \dots, \eta_d] \in \mathbb{R}^d$ and gather the edge parameters into a symmetric matrix $\Phi \in \mathbb{R}^{d \times d}$ such that $\phi_{ij} = \phi_{ji} = \eta_{ij}/2$, $\forall (i, j) \in \mathcal{E}$ and $\phi_{ij} = 0, \forall (i, j) \notin \mathcal{E}$. Note that for PGM, Φ has zeros along the diagonal. With this notation, the pairwise PGM can be equivalently represented in a compact vectorized form as:

$$\mathbb{P}_{\text{PGM}}(\mathbf{x}|\theta, \Phi) = \exp \left\{ \theta^T \mathbf{x} + \mathbf{x}^T \Phi \mathbf{x} - \sum_{i=1}^d \log(x_i!) - A_{\text{PGM}}(\theta, \Phi) \right\}, \quad (20)$$

Parameter estimation in a PGM is naturally suggested by its construction: all of the PGM parameters in Eq. (20) can be estimated by considering the node-conditional distributions for each node separately and solving an ℓ_1 -regularized Poisson regression for each variable. In contrast to the previous approaches in the sections above, this parameter estimation approach is not only simple but is also guaranteed to be consistent even under high-dimensional sampling regimes and under some other mild conditions, including a sparse graph structural assumption (see Yang et al.^{61,63} for more details on the analysis). As in Poisson log-normal models, the parameters of PGM can be made to depend on covariates to allow for more flexible correlations.⁶⁸

In spite of its simple parameter estimation method, the major drawback with this vanilla PGM distribution is that it only permits *negative* conditional dependencies between variables:

Proposition 1. (Ref 58) *Consider the PGM distribution in Eq. (20). Then, for any parameters θ and Φ , $A_{\text{PGM}}(\theta, \Phi) < +\infty$ only if the pairwise parameters are nonpositive: $\phi_{ij} \leq 0$, $\forall (i, j) \in \mathcal{E}$.*

Intuitively, if any entry in Φ , say Φ_{ij} , is positive, the term $\Phi_{ij} x_i x_j$ in Eq. (20) would grow

quadratically, whereas the log base measure terms $-\log(x_i!) - \log(x_j!)$ only decrease as $O(x_i \log x_i + x_j \log x_j)$, so $A(\theta, \Phi) \rightarrow \infty$ as $x_i, x_j \rightarrow \infty$. Thus, even though the PGM is a natural extension of the univariate Poisson distribution (from the node-conditional viewpoint), it entails a highly restrictive parameter space, with severely limited applicability. Thus, multiple PGM extensions attempt to relax this negativity restriction to permit positive dependencies, as described next.

Extensions of PGMs

To circumvent the severe limitations of the PGM distribution, which particularly only permits negative conditional dependencies, several extensions to PGM that permit a richer dependence structure have been proposed.

Truncated PGM

Because the negativity constraint is due in part to the infinite domain of count variables, a natural solution would be to truncate the domain of variables. It was Kaiser and Cressie⁶⁹ who first introduced an approach to truncate the Poisson distribution in the context of graphical models. Their idea was simply to use a Winsorized Poisson distribution for node-conditional distributions: x is a Winsorized Poisson if $z = \mathbb{I}(z' < R)z' + \mathbb{I}(z' \geq R)R$, where z' is Poisson, $\mathbb{I}(\cdot)$ is an indicator function, and R is a fixed positive constant denoting the truncation level. However, Yang et al.⁶² showed that Winsorized node-conditional distributions actually do *not* lead to a consistent joint distribution.

As an alternative way of truncation, Yang et al.⁶² instead keep the same parametric form as PGM (Eq. (20)) but merely truncate the domain to non-negative integers less than or equal to R —i.e., $\mathcal{D}_{\text{TPGM}} = \{0, 1, \dots, R\}$ —so that the joint distribution takes the form⁶³:

$$\mathbb{P}_{\text{TPGM}}(\mathbf{x}) = \exp \left\{ \theta^T \mathbf{x} + \mathbf{x}^T \Phi \mathbf{x} - \sum_i \log(x_i!) - A_{\text{TPGM}}(\theta, \Phi) \right\}. \quad (21)$$

As Yang et al.⁶² show, the node-conditional distributions of this graphical model distribution belong to an exponential family that is Poisson-like but with the domain bounded by R . Thus, the key difference from the vanilla PGM (Eq. (20)) is that the domain is finite, and hence, the log partition function $A_{\text{TPGM}}(\cdot)$ only involves a finite number of summations. Thus,

no restrictions are imposed on the parameters for the normalizability of the distribution.

Yang et al.⁶² discuss several major drawbacks to TPGM. First, the domain needs to be bounded a priori, so R should ideally be set larger than any unseen observation. Second, the effective range of parameter space for a nondegenerate distribution is still limited: as the truncation value R increases, the effective values of pairwise parameters become increasingly negative or close to zero; otherwise, the distribution can be degenerate, placing most of its probability mass at 0 or R .

Quadratic PGM and Sublinear PGM

Yang et al.⁶² also investigate the possibility of PGMs that (1) allow both positive and negative dependencies as well as (2) allow the domain to range over all non-negative integers. As described previously, a key reason for the negative constraint on the pairwise parameters ϕ_{ij} in Eq. (20) is that the log base measure $\sum_i \log(x_i!)$ scales more slowly than the quadratic pairwise term $\mathbf{x}^T \Phi \mathbf{x}$ where $\mathbf{x} \in \mathbb{Z}_+^d$. Yang et al.⁶² thus propose two possible solutions: increase the base measure or decrease the quadratic pairwise term.

First, if we modify the base measure of Poisson distribution with ‘Gaussian-esque’ quadratic functions (note that for the linear sufficient statistics with positive dependencies, the base measures should be quadratic at the very least⁶²), then the joint distribution, which they call a quadratic PGM, is normalizable while allowing both positive and negative dependencies⁶²:

$$\mathbb{P}_{\text{QPGM}}(\mathbf{x}) = \exp\{\theta^T \mathbf{x} + \mathbf{x}^T \Phi \mathbf{x} - A_{\text{QPGM}}(\theta, \Phi)\}. \quad (22)$$

Essentially, QPGM has the same form as the Gaussian distribution but where its domain is the set of non-negative integers. The key differences from PGM are that Φ can have negative values along the diagonal, and the Poisson base measure $\sum_i -\log(x_i!)$ is replaced by the quadratic term $\sum_i \phi_{ii} x_i^2$. Note that a sufficient condition for the distribution to be normalizable is given by:

$$\mathbf{x}^T \Phi \mathbf{x} < -c \|\mathbf{x}\|_2^2 \quad \forall \mathbf{x} \in \mathbb{Z}_+^d, \quad (23)$$

for some constant $c > 0$, which in turn can be satisfied if Φ is negative definite. One significant drawback of QPGM is that the tail is Gaussian-esque and thin rather than Poisson-esque and thicker as in PGM.

Another possible modification is to use sublinear sufficient statistics in order to preserve the Poisson base measure and possibly heavier tails.

Consider the following univariate distribution over count-valued variables:

$$\mathbb{P}(z) \propto \exp\{\theta T(z; R_0, R) - \log z!\}, \quad (24)$$

which has the same base measure $\log z!$ as the Poisson but with the following sublinear sufficient statistics:

$$T(z; R_0, R) = \begin{cases} z & \text{if } z \leq R_0 \\ -\frac{1}{2(R-R_0)} z^2 + \frac{R}{R-R_0} z - \frac{R_0^2}{2(R-R_0)} & \text{if } R_0 < z \leq R \\ \frac{R+R_0}{2} & \text{if } z \geq R. \end{cases} \quad (25)$$

For values of x upto R_0 , $T(x)$ increases linearly, while after R_0 , its slope decreases linearly, and finally after R , $T(x)$ becomes constant. The joint graphical model, which they call a sublinear PGM (SPGM), specified by the node-conditional distributions belonging to the family in Eq. (24), has the following form:

$$\mathbb{P}_{\text{SPGM}}(\mathbf{x}) = \exp\left\{\theta^T T(\mathbf{x}) + T(\mathbf{x})^T \Phi T(\mathbf{x}) - \sum_i \log(x_i!) - A_{\text{SPGM}}(\theta, \Phi | R_0, R)\right\}, \quad (26)$$

where

$$A_{\text{SPGM}}(\theta, \Phi | R_0, R) = \log \sum_{\mathbf{x} \in \mathbb{Z}_+^d} \exp\{\theta^T T(\mathbf{x}) + T(\mathbf{x})^T \Phi T(\mathbf{x}) - \sum_i \log(x_i!)\}, \quad (27)$$

and $T(\mathbf{x})$ is the entry-wise application of the function in Eq. (25). SPGM is always normalizable for $\phi_{ij} \in \mathbb{R} \quad \forall i \neq j$.⁶²

The main difficulty in estimating the PGM variants above with an infinite domain is the lack of closed-form expressions for the log partition function, even only for the node-conditional distributions that are needed for parameter estimation. Yang et al.⁶² propose an approximate estimation procedure that uses the univariate Poisson and Gaussian log-partition functions as upper bounds for the node-conditional log-partition functions for the QPGM and SPGM models, respectively.

Poisson Square Root Graphical Model

In the similar vein as SPGM in the earlier section, Inouye et al.⁶⁰ consider the use of exponential

families with square root-sufficient statistics. While they consider general graphical model families, their PGM variant called the square root graphical model (SQR) can be written as:

$$\mathbb{P}_{\text{SQR}}(\mathbf{x}|\theta) = \exp\left\{\theta^T \sqrt{\mathbf{x}} + \sqrt{\mathbf{x}}^T \Phi \sqrt{\mathbf{x}} - \sum_i \log(x_i!) - A_{\text{SQR}}(\theta, \Phi)\right\}, \quad (28)$$

where ϕ_{ii} can be non-zero in contrast to the zero diagonal of the parameter matrix in Eq. (20). As with PGM, when there are no edges (i.e., $\phi_{ij} = 0 \forall i \neq j$) and $\theta = 0$, this reduces to the independent Poisson model. The node conditionals of this distribution have the form:

$$\mathbb{P}(x_i|\mathbf{x}_{-i}) \propto \exp\left\{\phi_{ii}x_i + \left(\theta_i + 2\phi_{i,-i}^T \sqrt{\mathbf{x}_{-i}}\right)\sqrt{x_i} - \log(x_i!)\right\}, \quad (29)$$

where $\phi_{i,-i}$ is the i -th column of Φ with the i -th entry removed. This can be rewritten in the form of a two-parameter exponential family:

$$\mathbb{P}(x_i|\eta_1, \eta_2) = \exp\{\eta_1 x_i + \eta_2 \sqrt{x_i} - \log(x_i!) - A(\eta_1, \eta_2)\}, \quad (30)$$

where $\eta_1 = \phi_{ii}$, $\eta_2 = \theta_i + 2\phi_{i,-i}^T \sqrt{\mathbf{x}_{-i}}$, and $A(\eta_1, \eta_2)$ is the log partition function. Note that a key difference with the PGM variants in the previous section is that the diagonal of Φ_{SQR} can be non-zero, whereas the diagonal of Φ_{PGM} must be zero. Because the interaction term $\sqrt{\mathbf{x}}^T \Phi \sqrt{\mathbf{x}}$ is asymptotically linear rather than quadratic, the Poisson SQR graphical

model does not suffer from the degenerate distributions of TPGM as well as the fixed-length PGM (FLPGM) discussed in the next section while still allowing both positive and negative dependencies.

To show that SQR graphical models can easily be normalized, Inouye et al.⁶⁰ first define *radial*-conditional distributions. The *radial*-conditional distribution assumes that the *unit direction* is fixed, but the length of the vector is unknown. The difference between the standard 1D *node*-conditional distributions and the 1D *radial*-conditional distributions is illustrated in Figure 3. Suppose we condition on the unit direction $\mathbf{v} = \frac{\mathbf{x}}{\|\mathbf{x}\|_1}$ of the sufficient statistics, but the scaling of this unit direction $z = \|\mathbf{x}\|_1$ is unknown. With this notation, Inouye et al.⁶⁰ define the *radial*-conditional distribution as:

$$\begin{aligned} \mathbb{P}(\mathbf{x} = z\mathbf{v}|\mathbf{v}, \theta, \Phi) &\propto \exp\left\{\theta^T \sqrt{z\mathbf{v}} + \sqrt{z\mathbf{v}}^T \Phi \sqrt{z\mathbf{v}} - \sum_i \log((zv_i)!)\right\} \\ &\propto \exp\left\{(\theta^T \mathbf{v})\sqrt{z} + (\sqrt{\mathbf{v}}^T \Phi \sqrt{\mathbf{v}})z - \sum_i \log((zv_i)!)\right\}. \end{aligned}$$

Similar to the node-conditional distribution, the radial-conditional distribution can be rewritten as a two-parameter exponential family:

$$\mathbb{P}(z|\mathbf{v}, \theta, \Phi) = \exp\left(\underbrace{\bar{\eta}_1 z + \bar{\eta}_2 \sqrt{z}}_{O(z)} + \underbrace{\tilde{B}_{\mathbf{v}}(z)}_{O(-z \log(z))} - A_{\text{rad}}(\bar{\eta}_1, \bar{\eta}_2)\right), \quad (31)$$




FIGURE 3 | Node-conditional distributions (left) are univariate probability distributions of one variable conditioned on the other variables, while *radial*-conditional distributions are univariate probability distributions of the vector scaling conditioned on the vector direction. Both conditional distributions are helpful in understanding square root (SQR) graphical models. (Illustration from Ref 60)

where $\bar{\eta}_1 = \sqrt{\mathbf{v}^T \Phi \mathbf{v}}$, $\bar{\eta}_2 = \theta^T \mathbf{v}$, and $\tilde{B}_{\mathbf{v}}(z) = -\sum_{i=1}^d \log((zv_i!))$. The only difference between this exponential family and the node-conditional distribution is the different base measure—i.e., $-\sum_{i=1}^d \log((zv_i!)) \neq -\log(z!)$. However, note that the log base measure is still $O(-z \log(z))$, and thus, the log base measure will overcome the linear term as $z \rightarrow \infty$. Therefore, the *radial*-conditional distribution is normalizable for *any* $\bar{\eta}_1, \bar{\eta}_2 \in \mathbb{R}$.

With the radial-conditional distributions notation, Inouye et al.⁶⁰ show that the log partition function for Poisson SQR graphical models is finite by separating the summation into a nested radial direction and scalar summation. Let $\mathcal{V} = \{\mathbf{v} : \|\mathbf{v}\|_1 = 1, \mathbf{v} \in \mathbb{R}^d\}$ be the set of unit vectors in the positive orthant. The SQR log partition function $A_{\text{SQR}}(\theta, \Phi)$ can be decomposed into nested summation over the unit direction and the one-dimensional radial conditional:

$$A_{\text{SQR}}(\theta, \Phi) = \log \int \sum_{\mathbf{v} \in \mathcal{V}} \sum_{z \in \hat{\mathbb{Z}}} \exp\{\bar{\eta}_1(\mathbf{v}|\Phi)z + \bar{\eta}_2(\mathbf{v}|\theta)\sqrt{z}\} \sqrt{z} - \sum_i \log(zv_i!) \} d\mathbf{v}, \quad (32)$$

where $\bar{\eta}_1(\mathbf{v}|\Phi)$ and $\bar{\eta}_2(\mathbf{v}|\theta)$ are the radial-conditional parameters as defined above, and $\hat{\mathbb{Z}} = \{z : z\mathbf{v} \in \mathbb{Z}_+^d\}$. Note that $\hat{\mathbb{Z}} \subset \mathbb{Z}$, and thus, the inner summation can be replaced by the radial-conditional log partition function. Therefore, because \mathcal{V} is a bounded set, and the radial-conditional log partition function is finite for *any* $\bar{\eta}_1(\mathbf{v}|\theta)$ and $\bar{\eta}_2(\mathbf{v}|\Phi)$, $A_{\text{SQR}} < \infty$ and the Poisson SQR joint distribution is normalizable.

The main drawback of the Poisson SQR is that for parameter estimation, the log partition function $A(\eta_1, \eta_2)$ of the node conditionals in Eq. (30) is not known in closed form in general. Inouye et al.⁶⁰ provide a closed-form estimate for the exponential SQR, but a closed-form solution for the Poisson SQR model seems unlikely to exist. Inouye et al.⁶⁰ suggest numerically approximating $A(\eta_1, \eta_2)$ as it only requires a one-dimensional summation.

Local PGM

Inspired by the neighborhood selection technique of Meinshausen and Bühlmann⁷⁰, Allen and Liu^{71,72} propose to estimate the network structure of count-valued data by fitting a series of ℓ_1 -regularized Poisson regressions to estimate the node-neighborhoods. Such an estimation method may yield interesting network estimates, but as Allen and Liu⁷² note, these estimates do not correspond to a consistent joint

density. Instead, the underlying model is defined in terms of a series of local models where each variable is conditionally Poisson given its node-neighbors; this approach is thus termed the local PGM (LPGM). Note that LPGM does not impose any restrictions on the parameter space or types of dependencies; if the parameter space of each local model was constrained to be non-positive, then the LPGM reduces to the vanilla PGM as previously discussed. Hence, the LPGM is less interesting as a candidate multivariate model for count-valued data, but many may still find its simple and interpretable network estimates appealing. Recently, several studies have proposed to adopt this estimation strategy for alternative network types.^{73,74}

Fixed-Length Poisson MRFs

In a somewhat different direction, Inouye et al.⁵⁹ propose a distribution that has the same parametric form as the original PGM but allows positive dependencies by decomposing the joint distribution into two distributions. The first distribution is the marginal distribution over the length of the vector denoted $\mathbb{P}(L)$ —i.e., the distribution of the ℓ_1 -norm of the vector or the total sum of counts. The second distribution, the FLPGM, is the conditional distribution of PGM *given* the fact that the vector length L is known or fixed, denoted $\mathbb{P}_{\text{FLPGM}}(\mathbf{x} | \|\mathbf{x}\|_1 = L)$. Note that this allows the marginal distribution on length and the distribution given the length to be specified independently.^d The restriction to negative dependencies is removed because the second distribution, given the vector length $\mathbb{P}_{\text{FLPGM}}(\mathbf{x} | \|\mathbf{x}\|_1 = L)$, has a finite domain $\mathcal{D}_{\text{FLPGM}} = \{\mathbf{x} : \mathbf{x} \in \mathbb{Z}_+^d, \|\mathbf{x}\|_1 = L\}$ and is thus trivially normalizable—similar to the normalizability of the finite-domain TPGM. More formally,⁵⁹ defined the FLPGM as:

$$\mathbb{P}(\mathbf{x} | \theta, \Phi, \lambda) = \mathbb{P}(L | \lambda) \mathbb{P}_{\text{FLPGM}}(\mathbf{x} | \|\mathbf{x}\|_1 = L, \theta, \Phi), \quad (33)$$

$$\mathbb{P}_{\text{FLPGM}}(\mathbf{x} | \|\mathbf{x}\|_1 = L, \theta, \Phi) = \exp\{\theta^T \mathbf{x} + \mathbf{x}^T \Phi \mathbf{x} - \sum_i \log(x_i!) - A_L(\theta, \Phi)\}, \quad (34)$$

where λ is the parameter for the marginal length distribution, which could be Poisson, negative binomial, or any other distribution on non-negative integers. In addition, FLPGM could be used as a replacement for the multinomial distribution because it has the same domain as the multinomial and actually reduces to the multinomial if there are no dependencies. Earlier, Altham and Hankin⁷⁵ developed an identical model by generalizing an earlier bivariate generalization of

the binomial.⁷⁶ However, in Refs 75,76, the model assumed that L was constant over all samples, whereas Ref 59 allowed for L to vary for each sample according to some distributions $\mathbb{P}(L)$.

One significant drawback is that FLPGM is not amenable to the simple node-wise parameter estimation method of the previous PGM models. Nonetheless, in Inouye et al.⁵⁹, the parameters are heuristically estimated with Poisson regressions similar to PGM, although the theoretical properties of this heuristic estimate are unknown. Another drawback is that while FLPGM allows for positive dependencies, the distribution can yet yield a degenerate distribution for large values of L —similar to the problem of TPGM where the mass is concentrated near 0 or R . Thus, Inouye et al.⁵⁹ introduce a decreasing weighting function $\omega(L)$ that scales the interaction term:

$$\begin{aligned} \mathbb{P}_{\text{FLPGM}}(\mathbf{x} \mid \|\mathbf{x}\|_1 = L, \theta, \Phi, \omega(\cdot)) \\ = \exp \left\{ \theta^T \mathbf{x} + \omega(L) \mathbf{x}^T \Phi \mathbf{x} - \sum_i \log(x_i!) - A_L(\theta, \omega(L) \Phi) \right\}. \end{aligned} \quad (35)$$

While the log likelihood is not available in tractable form, Inouye et al.⁵⁹ approximate the log likelihood using annealed importance sampling,⁷⁷ which might be applicable to the extensions covered previously as well.

Summary of Conditional Poisson Generalizations

The conditional Poisson models benefit from the rich literature in exponential families and undirected graphical models, or Markov Random Fields. In addition, the conditional Poisson models have a simple parametric form. The historical PGM—or the auto-Poisson model⁵⁸—only allowed negative dependencies between variables. Multiple extensions have sought to overcome this severe limitation by altering the PGM so that the log partition function is finite even with positive dependencies. One major drawback to the graphical model approach is that computing the likelihood requires the approximation of the joint log partition function $A(\theta, \Phi)$; a related problem is that the distribution moments and marginals are not known in closed form. Despite these drawbacks, parameter estimation using composite likelihood methods via ℓ_1 -penalized node-wise regressions (in which the joint likelihood is not computed) has solid theoretical properties under certain conditions.

MODEL COMPARISON

We compare models by first discussing two structural aspects of the models: (a) interpretability and (b) the relative stringency and ease of verifying theoretical assumptions and guarantees. We then present and discuss an empirical comparison of the models on three real-world datasets.

Comparison of Model Interpretation

Marginal models can be interpreted as weakly decoupling the univariate marginal distributions from the dependency structure between the variables. However, in the discrete case, specifically for distributions based on pairing copulas with Poisson marginals, the dependency structure estimation is also dependent on the marginal estimation, unlike for copulas paired with continuous marginals.²³ *Conditional models* or graphical models, on the other hand, can be interpreted as specifying generative models for each variable given the variable's neighborhood (i.e., the conditional distribution). In addition, dependencies in graphical models can be visualized and interpreted via networks. Here, each variable is a node, and the weighted edges in the network structure depict the pair-wise conditional dependencies between variables. The simple network depiction for graphical models may enable domain experts to interpret complex dependency structures more easily compared to other models. Overall, marginal models may be preferred if modeling the statistics of the data, particularly the marginal statistics over individual variables, is of primary importance, while conditional models may be preferred if the prediction of some variables, given others, is of primary importance. *Mixture models* may be more or less difficult to interpret depending on whether there is an application-specific interpretation of the latent mixing variable. For example, a finite mixture of two Poisson distributions may model the crime statistics of a city that contains downtown and suburban areas. On the other hand, a finite mixture of Poisson distributions or a log-normal Poisson mixture when modeling crash severity counts (as seen in the empirical comparison section) seems more difficult to interpret; even if the model empirically fits the data well, the hidden mixture variable might not have an obvious application-specific interpretation.

Comparison of Theoretical Considerations

The estimation of *marginal models* from data has various theoretical problems, as evidenced by the analysis of copulas paired with discrete marginals in

Ref 23. The extent to which these theoretical problems cause any significant practical issues remains unclear. In particular, the estimators of the marginal distributions themselves typically have easily checked assumptions as the empirical marginal distributions can be inspected directly. On the other hand, the estimation of *conditional models* is both computationally tractable and comes with strong theoretical guarantees even under high-dimensional regimes where $n < d$.⁶³ However, the assumptions under which the guarantees of the estimators hold are difficult to check in practice and could cause problems if they are violated (e.g., outliers caused by unobserved factors). The estimation of *mixture models* tends to have limited theoretical guarantees. In particular, finite Poisson mixture models have very weak assumptions on the underlying distribution—eventually becoming a nonparametric distribution if $k = O(n)$ —but the estimation problems in theory are extremely difficult, with very few theoretical guarantees for practical estimators. Yet, as seen in the next section, empirically estimating a finite mixture model using EM iterations performs well in practice.

Empirical Comparison

In this section, we seek to empirically compare models from the three classes presented to assess how well they fit real-world count data.

Comparison Experimental Setup

We empirically compare models on selected datasets from three diverse domains which, have different data characteristics in terms of their mean count values and dispersion indices (Eq. 0.5), as can be seen in Table 1. The crash severity dataset is a small accident dataset from Ref 78 with three different count variables corresponding to crash severity classes: ‘Property-only’, ‘Possible Injury’, and ‘Injury’. The crash severity data exhibit high count values and

high overdispersion. We retrieve raw next generation sequencing data for breast cancer (BRCA) using the software TCGA2STAT⁷⁹ and compute a simple log-count transformation of the raw counts: $\lfloor \log(x+1) \rfloor$, a common preprocessing technique for RNA-Seq data. The BRCA data exhibit medium counts and medium overdispersion. We collect the word count vectors from the Classic3 text corpus, which contains abstracts from aerospace engineering, medical, and information sciences journals.^e The Classic3 dataset exhibits low counts—including many zeros—and medium overdispersion. In the Supporting information, we also give results for a crime statistics dataset and the 20 Newsgroup dataset, but they have similar characteristics and perform similarly to the BRCA and Classic3 datasets, respectively; thus, we omit them for simplicity. We select variables (e.g., for $d = 10$ or $d = 100$) by sorting the variables by mean count value or sorting by variance in the case of the BRCA dataset as highly variable genes are of more interest in biology.

In order to understand how each model might perform under varying data characteristics, we consider the following two questions: (1) How well does the model (i.e., the joint distribution) fit the underlying data distribution? and (2) How well does the model capture the dependency structure between variables? To help answer these questions, we evaluate the empirical fit of models using two metrics, which only require samples from the model. The first metric is based on a statistic called maximum mean discrepancy (MMD),⁸⁰ which estimates the maximum moment difference over all possible moments. The empirical MMD can be approximated as follows from two sets of samples $X \in \mathbb{R}^{n_1 \times d}$ and $Y \in \mathbb{R}^{n_2 \times d}$:

$$\widehat{\text{MMD}}(\mathcal{G}, X, Y) = \sup_{f \in \mathcal{G}} \frac{1}{n_1} \sum_{i=1}^{n_1} f(x_i) - \frac{1}{n_2} \sum_{j=1}^{n_2} f(y_j), \quad (36)$$

TABLE 1 | Dataset Statistics

Dataset	(Per Variable \Rightarrow)		Means			Dispersion Indices			Spearman's ρ		
	d	n	Min	Med	Max	Min	Med	Max	Min	Med	Max
Crash Severity	3	275	3.4	3.8	9.7	6	9.3	16	0.61	0.73	0.79
BRCA	10	878	3.2	5	7.7	1.5	2.2	3.8	−0.2	0.25	0.95
	100	878	1.1	4	9	0.63	1.7	4.6	−0.5	0.08	0.95
	1000	878	0.51	3.5	11	0.26	1	4.6	−0.64	0.06	0.97
Classic3	10	3893	0.26	0.33	0.51	1.4	3.4	3.8	−0.17	0.12	0.82
	100	3893	0.09	0.14	0.51	1.1	2.1	8.3	−0.17	0.02	0.82
	1000	3893	0.02	0.03	0.51	0.98	1.7	8.5	−0.17	−0	0.82

where \mathcal{G} is the union of the RKHS spaces based on the Gaussian kernel using 21 σ values log-spaced between 0.01 and 100. In our experiments, we estimate the MMD between the pairwise marginals of model samples and the pairwise marginals of the original observations:

$$D_{st}^{\text{MMD}} = \begin{cases} \widehat{\text{MMD}}\left(\mathcal{G}, [\mathbf{x}^{(s)}], [\hat{\mathbf{x}}^{(s)}]\right), & s = t \\ \widehat{\text{MMD}}\left(\mathcal{G}, [\mathbf{x}^{(s)}, \mathbf{x}^{(t)}], [\hat{\mathbf{x}}^{(s)}, \hat{\mathbf{x}}^{(t)}]\right), & \text{otherwise} \end{cases} \quad (37)$$

where $\mathbf{x}^{(s)}$ is the vector of data for the s -th variable of the true data, and $\hat{\mathbf{x}}^{(s)}$ is the vector of data for the n -th variable of samples from the estimated model—i.e., $\mathbf{x}^{(s)}$ are observations from the true underlying distribution, and $\hat{\mathbf{x}}^{(s)}$ are samples from the estimated model distribution. In our experiments, we use the fast approximation code for MMD from⁸¹ with 2^6 number of basis vectors for the FastMMD approximation algorithm. The second metric merely computes the absolute difference between pairwise Spearman's ρ values of model samples and Spearman's ρ values of the original observations:

$$D_{st}^{\rho} = |\rho(\mathbf{x}^{(s)}, \mathbf{x}^{(t)}) - \rho(\hat{\mathbf{x}}^{(s)}, \hat{\mathbf{x}}^{(t)})|, \quad \forall s, t. \quad (38)$$

The MMD metric is of more general interest because it evaluates whether the models actually fit the empirical data distribution, while the Spearman metric may be more interesting for practitioners who primarily care about the dependency structure, such as biologists who specifically want to study gene dependencies rather than gene distributions.

We empirically compare the model fits on these real-world datasets for several types of models from the three general classes presented. As a baseline, we estimate an independent Poisson model ('Ind Poisson'). We include Gaussian copulas and vine copulas, both paired with Poisson marginals ('Copula Poisson' and 'Vine Poisson') to represent the marginal model class. We estimate the copula-based models via the two-stage IFM method³⁵ via the DT.²⁶ For the mixture class, we include both a simple finite mixture of independent Poissons ('Mixture Poiss') and a log-normal mixture of Poissons ('Log-Normal'). The finite mixture was estimated using a simple EM algorithm; the log-normal mixture model was estimated via MCMC sampling using the code from Ref 55. For the conditional model class, we estimate the simple PGM ('PGM'), which only allows negative dependencies, and three variants that allow for positive dependencies: the truncated PGM ('Truncated PGM'), the

FLPGM with a Poisson distribution on the vector length $L = \|\mathbf{x}\|_1$ ('FLPGM Poisson'), and the Poisson square root graphical model ('Poisson SQR'). Using composite likelihood methods of penalized ℓ_1 node-wise regressions, we estimate these models via code from Refs 63,82,83, and the XMRF^f R package. After parameter estimation, we generate 1000 samples for each method using different types of sampling for each of the model classes.

To avoid overfitting to the data, we employ threefold cross-validation and report the average over the three folds. Because the conditional models (PGM, TPGM, FLPGM, and Poisson SQR) can be significantly different depending on the regularization parameter—i.e., the weight for the ℓ_1 regularization term in the objective function for these models—we select the regularization parameter of these models by computing the metrics on a tuning split of the training data. For the mixture model, we similarly tune the number of components k by testing $k = \{10, 20, 30, \dots, 100\}$. For the very high-dimensional datasets where $d = 1000$, we use a regularization parameter near the tuning parameters found when $d = 100$ and fix $k = 50$ in order to avoid the extra computation of selecting a parameter. More sampling and implementation details for each model are available in the Supporting information.

Empirical Comparison Results

The full results for both the MMD and Spearman's ρ metrics for the crash severity, breast cancer RNA-Seq, and Classic3 text datasets can be seen in Figures 4, 5, and 6, respectively. The low-dimensional results ($d \leq 10$) give evidence across all the datasets that three models outperform the others in their classes^g:

The Gaussian copula paired with Poisson marginals model ('Copula Poisson') for the marginal model class, the mixture of Poissons distribution ('Mixture Poiss') for the mixture model class, and the Poisson SQR distribution ('Poisson SQR') for the conditional model class. Thus, we only include these representative models along with an independent Poisson baseline in the high-dimensional experiments when $d > 10$. We discuss the results for specific data characteristics as represented by each dataset.^h

For the crash severity dataset with high counts and high overdispersion (Figure 4), mixture models (i.e., 'Log-Normal' and 'Mixture Poiss') perform the best as expected as they can model overdispersion well. However, if dependency structure is the only object of interest, the Gaussian copula paired with Poisson marginals ('Copula Poisson') performs well. For the BRCA dataset with medium counts and medium overdispersion (Figure 5), we note similar

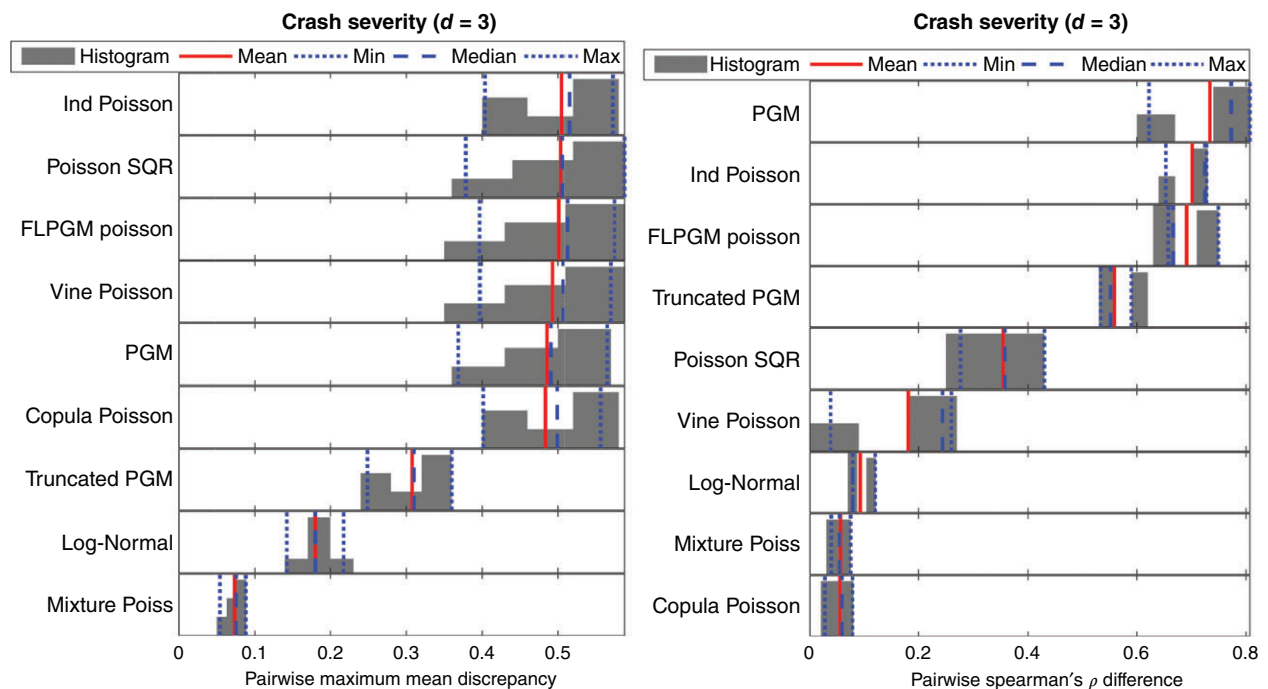


FIGURE 4 | Crash severity dataset (high counts and high overdispersion): maximum mean discrepancy (left) and Spearman ρ 's difference (right). As expected, for high overdispersion, mixture models ('Log-Normal' and 'Mixture Poiss') seem to perform the best.

trends with two notable exceptions: (1) The Poisson SQR model actually performs reasonably in low dimensions, suggesting that it can model moderate overdispersion, and (2) the high-dimensional ($d \geq 100$) Spearman's ρ difference results show that the Gaussian copula paired with Poisson marginals ('Copula Poisson') performs significantly better than the mixture model; this result suggests that copulas paired with Poisson marginals are likely better for modeling dependencies than mixture models. Finally, for the Classic3 dataset with low counts and medium overdispersion (Figure 6), the Poisson SQR model seems to perform well in this low-count setting, especially in low dimensions, unlike in previous data settings. While the simple independent mixture of Poisson distributions still performs well, the Poisson log-normal mixture distribution ('Log-Normal') performs quite poorly in this setting with small counts and many zeros. This poor performance of the Poisson log-normal mixture is somewhat surprising as the dispersion indices are almost all greater than one, as seen in Table 1. The differing results between low counts and medium counts with similar overdispersion demonstrate the importance of considering both the overdispersion and the mean count values when characterizing a dataset.

In summary, we note several overall trends. Mixture models are important for overdispersion when counts are medium or high. The Gaussian copula with

Poisson marginals joint distribution can estimate dependency structure (per the Spearman metric) for a wide range of data characteristics even when the distribution does not fit the underlying data (per the MMD metric). The Poisson SQR model performs well for low count values with many zeros (i.e., sparse data) and may be able to handle moderate overdispersion.

DISCUSSION

While this review analyzes each model class separately, it would be quite interesting to consider combinations or synergies between the model classes. Because negative binomial distributions can be viewed as a gamma-Poisson mixture model, one simple idea is to consider pairing a copula with negative binomial marginals or developing a negative binomial SQR graphical model. As another example, we could form a finite mixture of copula-based or graphical model-based models. This might combine the strengths of a mixture model in handling multiple modes and overdispersion with the strengths of the copula-based models and graphical models, which can explicitly model dependencies.

We may also consider how one type of model informs the other. For example, by the generalized Sklar's theorem,²⁶ each conditional PGM actually induces a copula—just as the Gaussian graphical




FIGURE 5 | BRCA RNA-Seq dataset (medium counts and medium overdispersion): maximum mean discrepancy (MMD) (top) and Spearman ρ 's difference (bottom) with different number of variables: 10 (left), 100 (middle), 1000 (right). While mixtures ('Log-Normal' and 'Mixture Poiss') perform well in terms of MMD, the Gaussian copula paired with Poisson marginals ('Copula Poisson') can model dependency structure well as evidenced by the Spearman metric.

model induces the Gaussian copula. Studying the copulas induced by graphical models seems to be a relatively unexplored area. On the other hand, it may be useful to consider fitting a Gaussian copula paired with discrete marginals using the theoretically grounded techniques from graphical models for sparse dependency structure estimation, especially for the small sample regimes in which $d > n$; this has been studied for the case of continuous marginals in Ref 84. Overall, bringing together and comparing these diverse paradigms for probability models opens up the door for many combinations and synergies.

CONCLUSION

We have reviewed three main approaches to constructing multivariate distributions derived from the Poisson using three different assumptions: (1) the

marginal distributions are derived from the Poisson, (2) the joint distribution is a mixture of independent Poisson distributions, and (3) the node-conditional distributions are derived from the Poisson. The first class based on Poisson marginals, and particularly the general approach of pairing copulas with Poisson marginals, provides an elegant way to partially decouple the marginals from the dependency structure and gives strong empirical results despite some theoretical issues related to nonuniqueness. While advanced methods to estimate the joint distribution of copulas paired with discrete marginals, such as SL²⁵ or vine copula constructions, provide more accurate or more flexible copula models, respectively, our empirical results suggest that a simple Gaussian copula paired with Poisson marginals with the trivial DT can perform quite well in practice. The second class based on mixture models can be particularly helpful for handling overdispersion that often occurs

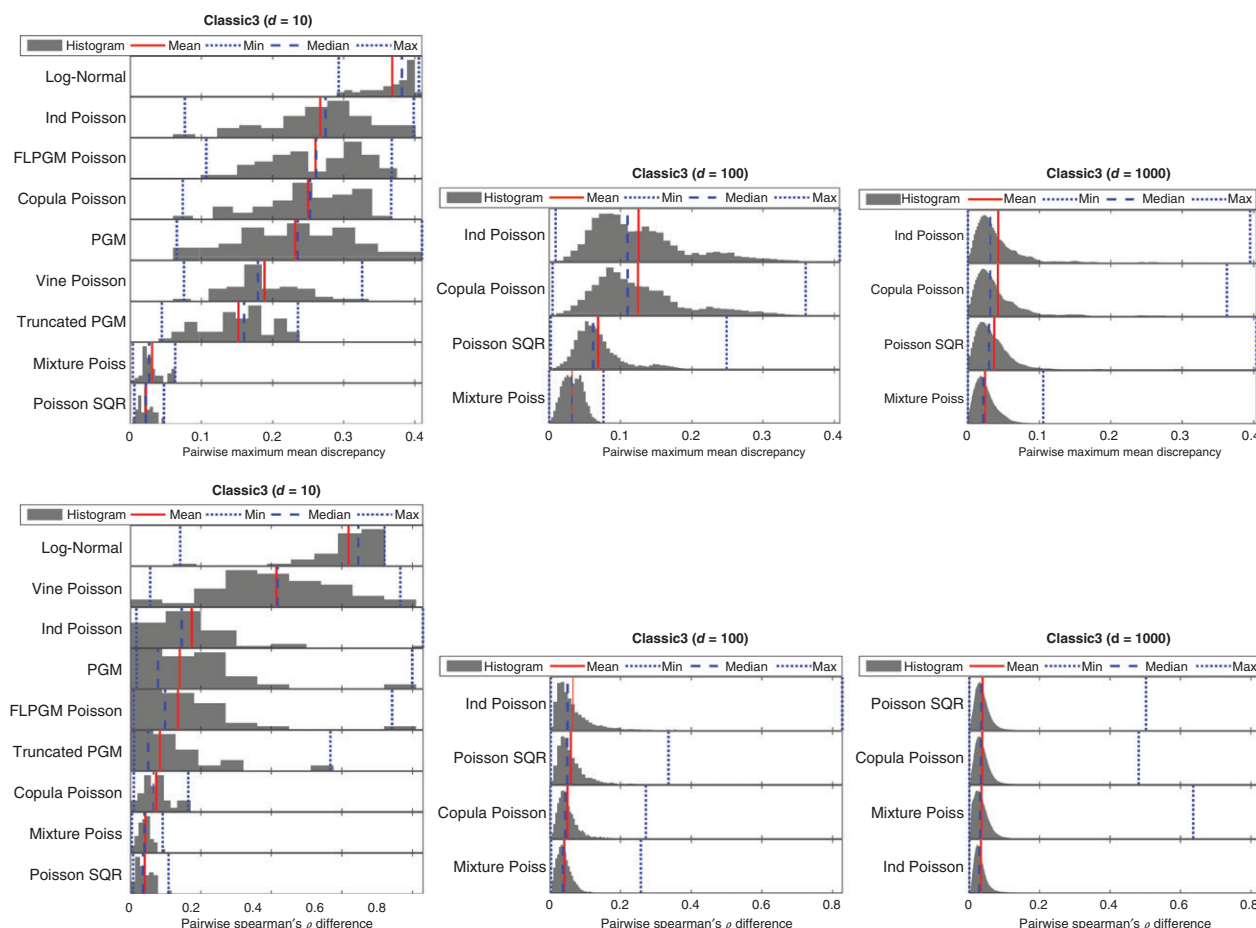


FIGURE 6 | Classic3 text dataset (low counts and medium overdispersion): maximum mean discrepancy (top) and Spearman ρ 's difference (bottom) with different number of variables: 10 (left), 100 (middle), 1000 (right). The Poisson SQR model performs better on this low count dataset than in previous settings.

in real count data with the log-normal–Poisson mixture and a finite mixture of independent Poisson distributions being prime examples. In addition, mixture models have closed-form moments and, in the case of a finite mixture, closed-form likelihood calculations—something not generally true for the other classes. The third class based on Poisson conditionals can be represented as graphical models, thus providing both compact and visually appealing representations of joint distributions. Conditional models benefit from strong theoretical guarantees about model recovery given certain modeling assumptions. However, checking conditional modeling assumptions may be impossible and may not always be satisfied for real-world count data. From our empirical experiments, we found that (1) mixture models are important for overdispersion when counts are medium or high, (2) the Gaussian copula with Poisson marginals joint distribution can estimate dependency structure for a wide range of data

characteristics even when the distribution does not fit the underlying data, and (3) Poisson SQR models perform well for low count values with many zeros (i.e., sparse data) and can handle moderate overdispersion. Overall, in practice, we would recommend comparing the three best-performing methods from each class, namely the Gaussian copula model paired with Poisson marginals, the finite mixture of independent Poisson distributions, and the Poisson SQR model. This initial comparison will likely highlight some interesting properties of a given dataset and suggest which class to pursue in more detail.

This review has highlighted several key strengths and weaknesses of the main approaches to constructing multivariate Poisson distributions. Yet, there remain many open questions. For example, what are the marginal distributions of the PGMs that are defined in terms of their conditional distributions? Or conversely, what are the conditional distributions of the copula models that are defined in

terms of their marginal distributions? Can novel models be created at the intersection of these model classes that could combine the strengths of different classes as suggested in the Discussion section? Could certain model classes be developed in an application area that has been largely dominated by another model class? For example, graphical models are well known in the machine learning literature, while copula models are well known in the financial modeling literature. Overall, multivariate Poisson models are poised to increase in popularity given the wide potential applications for real-world, high-dimensional, count-valued data in text analysis, genomics, spatial statistics, economics, and epidemiology.

NOTES

^a The label ‘multivariate Poisson’ was introduced in the statistics community to refer to the particular model introduced in this section, but other generalizations could also be considered multivariate Poisson distributions.

^b This is because if $y \in \mathbb{R} \sim \text{Normal}$, then $\exp(y) \in \mathbb{R}_+ \sim \text{LogNormal}$.

^c Besag⁵⁸ originally named these Poisson auto models, focusing on pairwise graphical models, but here, we consider the general graphical model setting.

^d If the marginal distribution on the length is set to be the same as the marginal distribution on length for the PGM—i.e., if $\mathbb{P}(L) = \sum_{\mathbf{x}: \|\mathbf{x}\|_1 = L} \mathbb{P}_{\text{PGM}}(\mathbf{x})$, then the PGM distribution is recovered.

^e http://ir.dcs.gla.ac.uk/resources/test_collections/

^f <https://cran.r-project.org/web/packages/XMRP/index.html>

^g For the crash-severity dataset, the truncated Poisson graphical model (‘Truncated PGM’) outperforms the Poisson SQR model under the pairwise MMD metric. After inspection, however, we realized that the Truncated PGM model performed better merely because outlier values were truncated to the 99th percentile as described in the supplementary material. This reduced the overfitting of outlier values caused by the crash-severity dataset’s high overdispersion.

^h These basic trends are also corroborated by the two datasets in the supplementary material.

ⁱ In the discrete case, the dependency structure cannot be perfectly decoupled from the marginal distributions, unlike in the continuous case where the dependency structure and marginals can be perfectly decoupled.

ACKNOWLEDGMENTS

D.I. and P.R. acknowledge the support of ARO via W911NF-12-1-0390, NSF via IIS-1149803, IIS-1447574, DMS-1264033, and NIH via R01 GM117594-01 as part of the Joint DMS/NIGMS Initiative to Support Research at the Interface of the Biological and Mathematical Sciences. E.Y. acknowledges the support of MSIP/IITP via ICT R&D program 2016-0-00563. G.A. acknowledges support from NSF DMS-1264058 and NSF DMS-1554821.

REFERENCES

1. Campbell J. The Poisson correlation function. *Proc Edinburgh Math Soc* 1934, 4:18–26.
2. M’Kendrick A. Applications of mathematics to medical problems. *Proc Edinburgh Math Soc* 1925, 44:98–130.
3. Teicher H. On the multivariate Poisson distribution. *Scand Actuar J* 1954, 1954:1–9.
4. Wicksell SD. Some theorems in the theory of probability, with special reference to their importance in the theory of homographic correlation. *Svenska Aktuarietidskrift*. 1916:165–213.
5. Nikoloulopoulos AK. Copula-based models for multivariate discrete response data. In: Jaworski P, Durante F, Härdle W, eds. *Copulae in Mathematical and Quantitative Finance*. Berlin Heidelberg: Springer-Verlag; 2013.
6. Nikoloulopoulos AK, Karlis D. Modeling multivariate count data using copulas. *Commun Stat Simul Comput* 2009, 39:172–187.
7. Xue-Kun Song P. Multivariate dispersion models generated from gaussian copula. *Scand J Stat* 2000, 27:305–320.
8. Holgate P. Estimation for the bivariate Poisson distribution. *Biometrika* 1964, 51:241–287.
9. Dwass M, Teicher H. On infinitely divisible random vectors. *Ann Math Stat* 1957, 28:461–470.
10. Kawamura K. The structure of multivariate Poisson distribution. *Kodai Math J* 1979, 2:337–345.
11. Srivastava R, Srivastava A. On a characterization of Poisson distribution. *J Appl Probab* 1970, 7:497–501.
12. Wang Y. Characterizations of certain multivariate distributions. *Mathematical Proceedings of the Cambridge Philosophical Society* 1974, 75:219–234.

13. Johnson NL, Kotz S, Balakrishnan N. *Discrete Multivariate Distributions*, vol. 165. New York: John Wiley & Sons; 1997.
14. Krishnamoorthy A. Multivariate binomial and Poisson distributions. *Sankhyā: Indian J Stat* 1951, 11:117–124.
15. Krummenauer F. Limit theorems for multivariate discrete distributions. *Metrika* 1998, 47:47–69.
16. Mahamunulu D. A note on regression in the multivariate Poisson distribution. *J Am Stat Assoc* 1967, 62:251–258.
17. Kano K, Kawamura K. On recurrence relations for the probability function of multivariate generalized Poisson distribution. *Commun Stat Theory Methods* 1991, 20:165–178.
18. Karlis D. An em algorithm for multivariate Poisson distribution and related models. *J Appl Stat* 2003, 30:63–77.
19. Loukas S, Kemp C. On computer sampling from trivariate and multivariate discrete distributions: multivariate discrete distributions. *J Stat Comput Simul* 1983, 17:113–123.
20. Tsiamyrtzis P, Karlis D. Strategies for efficient computation of multivariate Poisson probabilities. *Commun Stat Simul Comput* 2004, 33:271–292.
21. Sklar A. Fonctions de répartition à n dimensions et leurs marges. *Publ Inst Statist Univ Paris* 1959, 8:229–231.
22. Cherubini U, Luciano E, Vecchiato W. *Copula Methods in Finance*. Hoboken, NJ: John Wiley & Sons; 2004.
23. Genest C, Nešlehová J. A primer on copulas for count data. *Astin Bull* 2007, 37:475–515.
24. Nikoloulopoulos AK. On the estimation of normal copula discrete regression models using the continuous extension and simulated likelihood. *J Stat Plann Inference* 2013, 143:1923–1937.
25. Nikoloulopoulos AK. Efficient estimation of high-dimensional multivariate normal copula models with discrete spatial responses. *Stoch Environ Res Risk Assess* 2016, 30:493–505.
26. Rüschendorf L. Copulas, Sklar's theorem, and distributional transform. In: *Mathematical Risk Analysis*. Berlin, Heidelberg: Springer-Verlag; 2013, 3–34 (Chapter 1).
27. Demarta S, McNeil AJ. The t copula and related copulas. *Int Stat Rev* 2005, 73:111–129.
28. Trivedi PK, Zimmer DM. Copula Modeling: An Introduction for Practitioners. Foundations and Trends® in Econometrics, vol. 1. Boston - Delft: now publishers; 2007, 1–111.
29. Aas K, Czado C, Frigessi A, Bakken H. Pair-copula constructions of multiple dependence. *Insur: Math Econ* 2009, 44:182–198.
30. Bedford T, Cooke RM. Vines: a new graphical model for dependent random variables. *Ann Stat* 2002, 30:1031–1068.
31. Cook RJ, Lawless JF, Lee K-A. A copula-based mixed Poisson model for bivariate recurrent events under event-dependent censoring. *Stat Med* 2010, 29:694–707.
32. Yahav I, Shmueli G. On generating multivariate Poisson data in management science applications. *Appl Stoch Models Bus Ind* 2012, 28:91–102.
33. Sklar A. Random variables, joint distribution functions, and copulas. *Kybernetika* 1973, 9:449–460.
34. Karlis D. Models for multivariate count time series. In: Davis RA, Holan SH, Lund R, Ravishanker N, eds. *Handbook of Discrete-Valued Time Series*. Boca Raton, FL: CRC Press; 2016, 407–424(Chapter 19).
35. Joe H, Xu JJ. The estimation method of inference functions for margins for multivariate models. Technical Report 166, The University of British Columbia, Vancouver, Canada, 1996.
36. Denuit M, Lambert P. Constraints on concordance measures in bivariate discrete data. *J Multivariate Anal* 2005, 93:40–57.
37. Heinen A, Rengifo E. Multivariate autoregressive modeling of time series count data using copulas. *J Empir Finance* 2007, 14:564–583.
38. Heinen A, Rengifo E. Multivariate reduced rank regression in nonGaussian contexts, using copulas. *Comput Stat Data Anal* 2008, 52:2931–2944.
39. Madsen L. Maximum likelihood estimation of regression parameters with spatially dependent discrete data. *J Agric Biol Environ Stat* 2009, 14:375–391.
40. Madsen L, Fang Y. Joint regression analysis for discrete longitudinal data. *Biometrics* 2011, 67: 1171–1175.
41. Kazianka H. Approximate copula-based estimation and prediction of discrete spatial data. *Stoch Environ Res Risk Assess* 2013, 27:2015–2026.
42. Kazianka H, Pilz J. Copula-based geostatistical modeling of continuous and discrete data including covariates. *Stoch Environ Res Risk Assess* 2010, 24: 661–673.
43. Genz A, Bretz F. *Computation of Multivariate Normal and t Probabilities*, vol. 195. Dordrecht Heidelberg London New York: Springer; 2009.
44. Panagiotelis A, Czado C, Joe H. Pair copula constructions for multivariate discrete data. *J Am Stat Assoc* 2012, 107:1063–1072.
45. Czado C, Brechmann EC, Gruber L. Selection of vine copulas. In: *Copulae in Mathematical and Quantitative Finance*. Berlin Heidelberg: Springer; 2013, 17–37.
46. Karlis D, Xekalaki E. Mixed Poisson distributions. *Int Stat Rev* 2005, 73:35–58.

47. Arbous AG, Kerrich J. Accident statistics and the concept of accidentproneness. *Biometrics* 1951, 7: 340–432.
48. Steyn H. On the multivariate Poisson normal distribution. *J Am Stat Assoc* 1976, 71:233–236.
49. Aitchison J, Ho C. The multivariate Poisson-log normal distribution. *Biometrika* 1989, 76:643–653.
50. Agüero-Valverde J, Jovanis PP. Bayesian multivariate Poisson lognormal models for crash severity modeling and site ranking. *Transp Res Rec: J Transp Res Board* 2009, 2136:82–91.
51. Chib S, Winkelmann R. Markov chain Monte Carlo analysis of correlated count data. *J Bus Econ Stat* 2001, 19:428–435.
52. El-Basyouny K, Sayed T. Collision prediction models using multivariate Poisson-lognormal regression. *Accid Anal Prev* 2009, 41:820–828.
53. Ma J, Kockelman KM, Damien P. A multivariate Poisson-lognormal regression model for prediction of crash counts by severity, using Bayesian methods. *Accid Anal Prev* 2008, 40:964–975.
54. Park E, Lord D. Multivariate Poisson-lognormal models for jointly modeling crash frequency by severity. *Transp Res Rec: J Transp Res Board* 2007, 2019:1–6.
55. Zhan X, Abdul Aziz HM, Ukkusuri SV. An efficient parallel sampling technique for multivariate Poisson-lognormal model: Analysis with two crash count datasets. *Anal Methods Accid Res* 2015, 8:45–60.
56. Karlis D, Meligkotsidou L. Finite mixtures of multivariate Poisson distributions with application. *J Stat Plann Inference* 2007, 137:1942–1960.
57. Bradley JR, Holan SH, Wikle CK. Computationally efficient distribution theory for Bayesian inference of high-dimensional dependent count-valued data. arXiv preprint arXiv:1512.07273, 2015.
58. Besag J. Spatial interaction and the statistical analysis of lattice systems. *J R Stat Soc Ser B Stat Methodol* 1974, 36:192–236.
59. Inouye DI, Ravikumar P, Dhillon IS. Fixed-length Poisson MRF: adding dependencies to the multinomial. In: *Neural Information Processing Systems*, 28, 2015.
60. Inouye DI, Ravikumar P, Dhillon IS. Square root graphical models: multivariate generalizations of univariate exponential families that permit positive dependencies. In: *International Conference on Machine Learning*, 2016.
61. Yang E, Ravikumar P, Allen G, Liu Z. Graphical models via generalized linear models. In: *Neural Information Processing Systems*, 25, 2012.
62. Yang E, Ravikumar P, Allen G, Liu Z. On Poisson graphical models. In: *Neural Information Processing Systems*, 26, 2013.
63. Yang E, Ravikumar P, Allen G, Liu Z. Graphical models via univariate exponential family distribution. *J Mach Learn Res* 2015, 16:3813–3847.
64. Lauritzen SL. *Graphical Models*. Oxford: Oxford University Press; 1996.
65. Clifford P. Markov random fields in statistics. In: *Disorder in Physical Systems*. Oxford: Oxford Science Publications; 1990.
66. Wainwright MJ, Jordan MI. Graphical models, exponential families, and variational inference. *Found Trends R Mach Learn* 2008, 1:1–305.
67. Koller D, Friedman N. *Probabilistic Graphical Models: Principles and Techniques*. Cambridge, MA: MIT Press; 2009.
68. Yang E, Ravikumar P, Allen G, Liu Z. Conditional random fields via univariate exponential families. In: *Neural Information Processing Systems*, 26, 2013.
69. Kaiser MS, Cressie N. Modeling Poisson variables with positive spatial dependence. *Stat Probab Lett* 1997, 35:423–432.
70. Meinshausen N, Bühlmann P. High-dimensional graphs and variable selection with the lasso. *Ann Stat* 2006, 34:1436–1462.
71. Allen GI, Liu Z. A log-linear graphical model for inferring genetic networks from high-throughput sequencing data. In: *2012 I.E. International Conference on Bioinformatics and Biomedicine*, IEEE, 2012, 1–6.
72. Allen GI, Liu Z. A local Poisson graphical model for inferring networks from sequencing data. *IEEE Trans NanoBiosci* 2013, 12:189–198.
73. Hadji F, Molina A, Natarajan S, Kersting K. Poisson dependency networks: gradient boosted models for multivariate count data. *Mach Learn* 2015, 100:477–507.
74. Han SW, Zhong H. Estimation of sparse directed acyclic graphs for multivariate counts data. *Biometrics* 2016, 72:791–803.
75. Altham PME, Hankin RKS. Multivariate generalizations of the multiplicative binomial distribution: introducing the mm package. *J Stat Softw* 2012, 46:1–23. doi:10.18637/jss.v046.i12.
76. Altham PME. Two generalizations of the binomial distribution. *J R Stat Soc Ser C Appl Stat* 1978, 27:162–167.
77. Neal RM. Annealed importance sampling. *Stat Comput* 2001, 11:125–139.
78. Milton JC, Shankar VN, Mannering FL. Highway accident severities and the mixed logit model: an exploratory empirical analysis. *Accid Anal Prev* 2008, 40:260–266.
79. Wan Y-W, Allen GI, Liu Z. TCGA2STAT: simple TCGA data access for integrated statistical analysis in R. *Bioinformatics* 2016, 32:952–954.
80. Gretton A. A kernel two-sample test. *J Mach Learn Res* 2012, 13:723–773.
81. Zhao J, Meng D. FastMMD: ensemble of circular discrepancy for efficient two-sample test. *Neural Comput* 2015, 27:1345–1372.

82. Inouye DI, Ravikumar P, Dhillon IS. Admixtures of Poisson MRFs: a topic model with word dependencies. In: *International Conference on Machine Learning*, vol 31, 2014.
83. Inouye DI, Ravikumar P, Dhillon IS. Generalized root models: beyond pairwise graphical models for univariate exponential families. arXiv preprint arXiv:1606.00813, 2016.
84. H. Liu, F. Han, M. Yuan, J. Lafferty, and L. Wasserman. High dimensional semiparametric Gaussian copula graphical models. *Ann Stat*, 40:34, 2012. 10.1214/12-AOS1037.