
Poisson–Gamma Dynamical Systems

Aaron Schein

College of Information and Computer Sciences
University of Massachusetts Amherst
Amherst, MA 01003
aschein@cs.umass.edu

Mingyuan Zhou

McCombs School of Business
The University of Texas at Austin
Austin, TX 78712
mingyuan.zhou@mcombs.utexas.edu

Hanna Wallach

Microsoft Research New York
641 Avenue of the Americas
New York, NY 10011
hanna@dirichlet.net

Abstract

We introduce a new dynamical system for sequentially observed multivariate count data. This model is based on the gamma–Poisson construction—a natural choice for count data—and relies on a novel Bayesian nonparametric prior that ties and shrinks the model parameters, thus avoiding overfitting. We present an efficient MCMC inference algorithm that advances recent work on augmentation schemes for inference in negative binomial models. Finally, we demonstrate the model’s inductive bias using a variety of real-world data sets, showing that it exhibits superior predictive performance over other models and infers highly interpretable latent structure.

1 Introduction

Sequentially observed count vectors $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(T)}$ are the main object of study in many real-world applications, including text analysis, social network analysis, and recommender systems. Count data pose unique statistical and computational challenges when they are high-dimensional, sparse, and overdispersed, as is often the case in real-world applications. For example, when tracking counts of user interactions in a social network, only a tiny fraction of possible edges are ever active, exhibiting bursty periods of activity when they are. Models of such data should exploit this sparsity in order to scale to high dimensions and be robust to overdispersed temporal patterns. In addition to these characteristics, sequentially observed multivariate count data often exhibit complex dependencies within and across time steps. For example, scientific papers about one topic may encourage researchers to write papers about another related topic in the following year. Models of such data should therefore capture the topic structure of individual documents as well as the excitatory relationships between topics.

The linear dynamical system (LDS) is a widely used model for sequentially observed data, with many well-developed inference techniques based on the Kalman filter [1, 2]. The LDS assumes that each sequentially observed V -dimensional vector $\mathbf{r}^{(t)}$ is real valued and Gaussian distributed: $\mathbf{r}^{(t)} \sim \mathcal{N}(\Phi \boldsymbol{\theta}^{(t)}, \Sigma)$, where $\boldsymbol{\theta}^{(t)} \in \mathbb{R}^K$ is a latent state, with K components, that is linked to the observed space via $\Phi \in \mathbb{R}^{V \times K}$. The LDS derives its expressive power from the way it assumes that the latent states evolve: $\boldsymbol{\theta}^{(t)} \sim \mathcal{N}(\Pi \boldsymbol{\theta}^{(t-1)}, \Delta)$, where $\Pi \in \mathbb{R}^{K \times K}$ is a transition matrix that captures between-component dependencies across time steps. Although the LDS can be linked to non-real observations via the extended Kalman filter [3], it cannot efficiently model real-world count data because inference is $\mathcal{O}((K + V)^3)$ and thus scales poorly with the dimensionality of the data [2].

Many previous approaches to modeling sequentially observed count data rely on the generalized linear modeling framework [4] to link the observations to a latent Gaussian space—e.g., via the Poisson–lognormal link [5]. Researchers have used this construction to factorize sequentially observed count matrices under a Poisson likelihood, while modeling the temporal structure using well-studied Gaussian techniques [6, 7]. Most of these previous approaches assume a simple Gaussian state-space model—i.e., $\theta^{(t)} \sim \mathcal{N}(\theta^{(t-1)}, \Delta)$ —that lacks the expressive transition structure of the LDS; one notable exception is the Poisson linear dynamical system [8]. In practice, these approaches exhibit prohibitive computational complexity in high dimensions, and the Gaussian assumption may fail to accommodate the burstiness often inherent to real-world count data [9].

We present the *Poisson–gamma dynamical system (PGDS)*—a new dynamical system, based on the gamma–Poisson construction, that supports the expressive transition structure of the LDS. This model naturally handles overdispersed data. We introduce a new Bayesian nonparametric prior to automatically infer the model’s rank. We develop an elegant and efficient algorithm for inferring the parameters of the transition structure that advances recent work on augmentation schemes for inference in negative binomial models [10] and scales with the number of non-zero counts, thus exploiting the sparsity inherent to real-world count data. We examine the way in which the dynamical gamma–Poisson construction propagates information and derive the model’s steady state, which involves the Lambert W function [11]. Finally, we use the PGDS to analyze a diverse range of real-world data sets, showing that it exhibits excellent predictive performance on smoothing and forecasting tasks and infers interpretable latent structure, an example of which is depicted in figure 1.

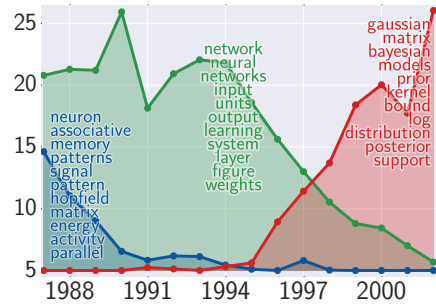


Figure 1: The time-step factors for three components inferred by the PGDS from a corpus of NIPS papers. Each component is associated with a feature factor for each word type in the corpus; we list the words with the largest factors. The inferred structure tells a familiar story about the rise and fall of certain subfields of machine learning.

2 Poisson–Gamma Dynamical Systems

We can represent a data set of V -dimensional sequentially observed count vectors $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(T)}$ as a $V \times T$ count matrix Y . The PGDS models a single count $y_v^{(t)} \in \{0, 1, \dots\}$ in this matrix as follows:

$$y_v^{(t)} \sim \text{Pois}(\delta^{(t)} \sum_{k=1}^K \phi_{vk} \theta_k^{(t)}) \text{ and } \theta_k^{(t)} \sim \text{Gam}(\tau_0 \sum_{k_2=1}^K \pi_{kk_2} \theta_{k_2}^{(t-1)}, \tau_0), \quad (1)$$

where the latent factors ϕ_{vk} and $\theta_k^{(t)}$ are both positive, and represent the strength of feature v in component k and the strength of component k at time step t , respectively. The scaling factor $\delta^{(t)}$ captures the scale of the counts at time step t , and therefore obviates the need to rescale the data as a preprocessing step. We refer to the PGDS as *stationary* if $\delta^{(t)} = \delta$ for $t = 1, \dots, T$. We can view the feature factors as a $V \times K$ matrix Φ and the time-step factors as a $T \times K$ matrix Θ . Because we can also collectively view the scaling factors and time-step factors as a $T \times K$ matrix Ψ , where element $\psi_{tk} = \delta^{(t)} \theta_k^{(t)}$, the PGDS is a form of Poisson matrix factorization: $Y \sim \text{Pois}(\Phi \Psi^T)$ [12, 13, 14, 15].

The PGDS is characterized by its expressive transition structure, which assumes that each time-step factor $\theta_k^{(t)}$ is drawn from a gamma distribution, whose shape parameter is a linear combination of the K factors at the previous time step. The latent transition weights $\pi_{11}, \dots, \pi_{k_1 k_2}, \dots, \pi_{KK}$, which we can view as a $K \times K$ transition matrix Π , capture the excitatory relationships between components. The vector $\theta^{(t)} = (\theta_1^{(t)}, \dots, \theta_K^{(t)})$ has an expected value of $\mathbb{E}[\theta^{(t)} | \theta^{(t-1)}, \Pi] = \Pi \theta^{(t-1)}$ and is therefore analogous to a latent state in the the LDS. The concentration parameter τ_0 determines the variance of $\theta^{(t)}$ —specifically, $\text{Var}(\theta^{(t)} | \theta^{(t-1)}, \Pi) = (\Pi \theta^{(t-1)}) \tau_0^{-1}$ —without affecting its expected value.

To model the strength of each component, we introduce K component weights $\nu = (\nu_1, \dots, \nu_K)$ and place a shrinkage prior over them. We assume that the time-step factors and transition weights for component k are tied to its component weight ν_k . Specifically, we define the following structure:

$$\theta_k^{(1)} \sim \text{Gam}(\tau_0 \nu_k, \tau_0) \text{ and } \pi_k \sim \text{Dir}(\nu_1 \nu_k, \dots, \xi \nu_k \dots, \nu_K \nu_k) \text{ and } \nu_k \sim \text{Gam}(\frac{\gamma_0}{K}, \beta), \quad (2)$$

where $\boldsymbol{\pi}_k = (\pi_{1k}, \dots, \pi_{Kk})$ is the k^{th} column of Π . Because $\sum_{k_1=1}^K \pi_{k_1 k} = 1$, we can interpret $\pi_{k_1 k}$ as the probability of transitioning from component k to component k_1 . (We note that interpreting Π as a stochastic transition matrix relates the PGDS to the discrete hidden Markov model.) For a fixed value of γ_0 , increasing K will encourage many of the component weights to be small. A small value of ν_k will shrink $\theta_k^{(1)}$, as well as the transition weights in the k^{th} row of Π . Small values of the transition weights in the k^{th} row of Π therefore prevent component k from being excited by the other components and by itself. Specifically, because the shape parameter for the gamma prior over $\theta_k^{(t)}$ involves a linear combination of $\theta^{(t-1)}$ and the transition weights in the k^{th} row of Π , small transition weights will result in a small shape parameter, shrinking $\theta_k^{(t)}$. Thus, the component weights play a critical role in the PGDS by enabling it to automatically turn off any unneeded capacity and avoid overfitting.

Finally, we place Dirichlet priors over the feature factors and draw the other parameters from a non-informative gamma prior: $\boldsymbol{\phi}_k = (\phi_{1k}, \dots, \phi_{V_k}) \sim \text{Dir}(\eta_0, \dots, \eta_0)$ and $\delta^{(t)}, \xi, \beta \sim \text{Gam}(\epsilon_0, \epsilon_0)$. The PGDS therefore has four positive hyperparameters to be set by the user: τ_0, γ_0, η_0 , and ϵ_0 .

Bayesian nonparametric interpretation: As $K \rightarrow \infty$, the component weights and their corresponding feature factor vectors constitute a draw $G = \sum_{k=1}^{\infty} \nu_k \mathbb{1}_{\phi_k}$ from a gamma process $\text{GamP}(G_0, \beta)$, where β is a scale parameter and G_0 is a finite and continuous base measure over a complete separable metric space Ω [16]. Models based on the gamma process have an inherent shrinkage mechanism because the number of atoms with weights greater than $\varepsilon > 0$ follows a Poisson distribution with a finite mean—specifically, $\text{Pois}(\gamma_0 \int_{\varepsilon}^{\infty} d\nu \nu^{-1} \exp(-\beta \nu))$, where $\gamma_0 = G_0(\Omega)$ is the total mass under the base measure. This interpretation enables us to view the priors over Π and Θ as novel stochastic processes, which we call the *column-normalized relational gamma process* and the *recurrent gamma process*, respectively. We provide the definitions of these processes in the supplementary material.

Non-count observations: The PGDS can also model non-count data by linking the observed vectors to latent counts. A binary observation $b_v^{(t)}$ can be linked to a latent Poisson count $y_v^{(t)}$ via the Bernoulli–Poisson distribution: $b_v^{(t)} = \mathbb{1}(y_v^{(t)} \geq 1)$ and $y_v^{(t)} \sim \text{Pois}(\delta^{(t)} \sum_{k=1}^K \phi_{vk} \theta_k^{(t)})$ [17]. Similarly, a real-valued observation $r_v^{(t)}$ can be linked to a latent Poisson count $y_v^{(t)}$ via the Poisson randomized gamma distribution [18]. Finally, Basbug and Engelhardt [19] recently showed that many types of non-count matrices can be linked to a latent count matrix via the compound Poisson distribution [20].

3 MCMC Inference

MCMC inference for the PGDS consists of drawing samples of the model parameters from their joint posterior distribution given an observed count matrix Y and the model hyperparameters $\tau_0, \gamma_0, \eta_0, \epsilon_0$. In this section, we present a Gibbs sampling algorithm for drawing these samples. At a high level, our approach is similar to that used to develop Gibbs sampling algorithms for several other related models [10, 21, 22, 17]; however, we extend this approach to handle the unique properties of the PGDS. The main technical challenge is sampling Θ from its conditional posterior, which does not have a closed form. We address this challenge by introducing a set of auxiliary variables. Under this augmented version of the model, marginalizing over Θ becomes tractable and its conditional posterior has a closed form. Moreover, by introducing these auxiliary variables and marginalizing over Θ , we obtain an alternative model specification that we can subsequently exploit to obtain closed-form conditional posteriors for $\Pi, \boldsymbol{\nu}$, and ξ . We marginalize over Θ by performing a “backward filtering” pass, starting with $\boldsymbol{\theta}^{(T)}$. We repeatedly exploit the following three definitions in order to do this.

Definition 1: If $y = \sum_{n=1}^N y_n$, where $y_n \sim \text{Pois}(\theta_n)$ are independent Poisson-distributed random variables, then $(y_1, \dots, y_N) \sim \text{Mult}(y, (\frac{\theta_1}{\sum_{n=1}^N \theta_n}, \dots, \frac{\theta_N}{\sum_{n=1}^N \theta_n}))$ and $y \sim \text{Pois}(\sum_{n=1}^N \theta_n)$ [23, 24].

Definition 2: If $y \sim \text{Pois}(c\theta)$, where c is a constant, and $\theta \sim \text{Gam}(a, b)$, then $y \sim \text{NB}(a, \frac{c}{b+c})$ is a negative binomial-distributed random variable. We can equivalently parameterize it as $y \sim \text{NB}(a, g(\zeta))$, where $g(z) = 1 - \exp(-z)$ is the Bernoulli–Poisson link [17] and $\zeta = \ln(1 + \frac{c}{b})$.

Definition 3: If $y \sim \text{NB}(a, g(\zeta))$ and $l \sim \text{CRT}(y, a)$ is a Chinese restaurant table-distributed random variable, then y and l are equivalently jointly distributed as $y \sim \text{SumLog}(l, g(\zeta))$ and $l \sim \text{Pois}(a\zeta)$ [21]. The sum logarithmic distribution is further defined as the sum of l independent and identically logarithmic-distributed random variables—i.e., $y = \sum_{i=1}^l x_i$ and $x_i \sim \text{Log}(g(\zeta))$.

Marginalizing over Θ : We first note that we can re-express the Poisson likelihood in equation 1 in terms of latent subcounts [13]: $y_v^{(t)} = y_{v\cdot}^{(t)} = \sum_{k=1}^K y_{vk}^{(t)}$ and $y_{vk}^{(t)} \sim \text{Pois}(\delta^{(t)} \phi_{vk} \theta_k^{(t)})$. We then define $y_{\cdot k}^{(t)} = \sum_{v=1}^V y_{vk}^{(t)}$. Via definition 1, we obtain $y_{\cdot k}^{(t)} \sim \text{Pois}(\delta^{(t)} \theta_k^{(t)})$ because $\sum_{v=1}^V \phi_{vk} = 1$.

We start with $\theta_k^{(T)}$ because none of the other time-step factors depend on it in their priors. Via definition 2, we can immediately marginalize over $\theta_k^{(T)}$ to obtain the following equation:

$$y_{\cdot k}^{(T)} \sim \text{NB}(\tau_0 \sum_{k_2=1}^K \pi_{kk_2} \theta_{k_2}^{(T-1)}, g(\zeta^{(T)})), \text{ where } \zeta^{(T)} = \ln(1 + \frac{\delta^{(T)}}{\tau_0}). \quad (3)$$

Next, we marginalize over $\theta_k^{(T-1)}$. To do this, we introduce an auxiliary variable: $l_k^{(T)} \sim \text{CRT}(y_{\cdot k}^{(T)}, \tau_0 \sum_{k_2=1}^K \pi_{kk_2} \theta_{k_2}^{(T-1)})$. We can then re-express the joint distribution over $y_{\cdot k}^{(T)}$ and $l_k^{(T)}$ as

$$y_{\cdot k}^{(T)} \sim \text{SumLog}(l_k^{(T)}, g(\zeta^{(T)})) \text{ and } l_k^{(T)} \sim \text{Pois}(\zeta^{(T)} \tau_0 \sum_{k_2=1}^K \pi_{kk_2} \theta_{k_2}^{(T-1)}). \quad (4)$$

We are still unable to marginalize over $\theta_k^{(T-1)}$ because it appears in a sum in the parameter of the Poisson distribution over $l_k^{(T)}$; however, via definition 1, we can re-express this distribution as

$$l_k^{(T)} = l_{k\cdot}^{(T)} = \sum_{k_2=1}^K l_{kk_2}^{(T)} \text{ and } l_{kk_2}^{(T)} \sim \text{Pois}(\zeta^{(T)} \tau_0 \pi_{kk_2} \theta_{k_2}^{(T-1)}). \quad (5)$$

We then define $l_{\cdot k}^{(T)} = \sum_{k_1=1}^K l_{k_1 k}^{(T)}$. Again via definition 1, we can express the distribution over $l_{\cdot k}^{(T)}$ as $l_{\cdot k}^{(T)} \sim \text{Pois}(\zeta^{(T)} \tau_0 \theta_k^{(T-1)})$. We note that this expression does not depend on the transition weights because $\sum_{k_1=1}^K \pi_{k_1 k} = 1$. We also note that definition 1 implies that $(l_{1k}^{(T)}, \dots, l_{Kk}^{(T)}) \sim \text{Mult}(l_{\cdot k}^{(T)}, (\pi_{1k}, \dots, \pi_{Kk}))$. Next, we introduce $m_k^{(T-1)} = y_{\cdot k}^{(T-1)} + l_{\cdot k}^{(T)}$, which summarizes all of the information about the data at time steps $T-1$ and T via $y_{\cdot k}^{(T-1)}$ and $l_{\cdot k}^{(T)}$, respectively. Because $y_{\cdot k}^{(T-1)}$ and $l_{\cdot k}^{(T)}$ are both Poisson distributed, we can use definition 1 to obtain

$$m_k^{(T-1)} \sim \text{Pois}(\theta_k^{(T-1)} (\delta^{(T-1)} + \zeta^{(T)} \tau_0)). \quad (6)$$

Combining this likelihood with the gamma prior in equation 1, we can marginalize over $\theta_k^{(T-1)}$:

$$m_k^{(T-1)} \sim \text{NB}(\tau_0 \sum_{k_2=1}^K \pi_{kk_2} \theta_{k_2}^{(T-2)}, g(\zeta^{(T-1)})), \text{ where } \zeta^{(T-1)} = \ln(1 + \frac{\delta^{(T-1)}}{\tau_0} + \zeta^{(T)}). \quad (7)$$

We then introduce $l_k^{(T-1)} \sim \text{CRT}(m_k^{(T-1)}, \tau_0 \sum_{k_2=1}^K \pi_{kk_2} \theta_{k_2}^{(T-2)})$ and re-express the joint distribution over $l_k^{(T-1)}$ and $m_k^{(T-1)}$ as the product of a Poisson and a sum logarithmic distribution, similar to equation 4. This then allows us to marginalize over $\theta_k^{(T-2)}$ to obtain a negative binomial distribution. We can repeat the same process all the way back to $t=1$, where marginalizing over $\theta_k^{(1)}$ yields $m_k^{(1)} \sim \text{NB}(\tau_0 \nu_k, g(\zeta^{(1)}))$. We note that just as $m_k^{(t)}$ summarizes all of the information about the data at time steps t, \dots, T , $\zeta^{(t)} = \ln(1 + \frac{\delta^{(t)}}{\tau_0} + \zeta^{(t+1)})$ summarizes all of the information about $\delta^{(t)}, \dots, \delta^{(T)}$.

As we mentioned previously, introducing these auxiliary variables and marginalizing over Θ also enables us to define an alternative model specification that we can exploit to obtain closed-form conditional posteriors for Π , ν , and ξ . We provide part of its generative process in figure 2. We define $m_k^{(T)} = y_{\cdot k}^{(T)} + l_{\cdot k}^{(T+1)}$, where $l_{\cdot k}^{(T+1)} = 0$, and $\zeta^{(T+1)} = 0$ so that we can present the alternative model specification concisely.

$$\begin{aligned} l_{\cdot k}^{(1)} &\sim \text{Pois}(\zeta^{(1)} \tau_0 \nu_k) \\ (l_{1k}^{(t)}, \dots, l_{Kk}^{(t)}) &\sim \text{Mult}(l_{\cdot k}^{(t)}, (\pi_{1k}, \dots, \pi_{Kk})) \text{ for } t > 1 \\ l_{\cdot k}^{(t)} &= \sum_{k_2=1}^K l_{kk_2}^{(t)} \text{ for } t > 1 \\ m_k^{(t)} &\sim \text{SumLog}(l_{\cdot k}^{(t)}, g(\zeta^{(t)})) \\ (y_{\cdot k}^{(t)}, l_{\cdot k}^{(t+1)}) &\sim \text{Bin}(m_k^{(t)}, (\frac{\delta^{(t)}}{\delta^{(t)} + \zeta^{(t+1)} \tau_0}, \frac{\zeta^{(t+1)} \tau_0}{\delta^{(t)} + \zeta^{(t+1)} \tau_0})) \\ (y_{1k}^{(t)}, \dots, y_{V_k}^{(t)}) &\sim \text{Mult}(y_{\cdot k}^{(t)}, (\phi_{1k}, \dots, \phi_{V_k})) \end{aligned}$$

Figure 2: Alternative model specification.

Steady state: We draw particular attention to the backward pass $\zeta^{(t)} = \ln(1 + \frac{\delta^{(t)}}{\tau_0} + \zeta^{(t+1)})$ that propagates information about $\delta^{(t)}, \dots, \delta^{(T)}$ as we marginalize over Θ . In the case of the stationary PGDS—i.e., $\delta^{(t)} = \delta$ —the backward pass has a fixed point that we define in the following proposition.

Proposition 1: The backward pass has a fixed point of $\zeta^* = -\mathbb{W}_{-1}(-\exp(-1 - \frac{\delta}{\tau_0})) - 1 - \frac{\delta}{\tau_0}$.

The function $\mathbb{W}_{-1}(\cdot)$ is the lower real part of the Lambert W function [11]. We prove this proposition in the supplementary material. During inference, we perform the $\mathcal{O}(T)$ backward pass repeatedly. The existence of a fixed point means that we can assume the stationary PGDS is in its steady state and replace the backward pass with an $\mathcal{O}(1)$ computation¹ of the fixed point ζ^* . To make this assumption, we must also assume that $l_{\cdot,k}^{(T+1)} \sim \text{Pois}(\zeta^* \tau_0 \theta_k^{(T)})$ instead of $l_{\cdot,k}^{(T+1)} = 0$. We note that an analogous steady-state approximation exists for the LDS and is routinely exploited to reduce computation [25].

Gibbs sampling algorithm: Given Y and the hyperparameters, Gibbs sampling involves resampling each auxiliary variable or model parameter from its conditional posterior. Our algorithm involves a “backward filtering” pass and a “forward sampling” pass, which together form a “backward filtering–forward sampling” algorithm. We use $-\setminus \Theta^{(\geq t)}$ to denote everything excluding $\theta^{(t)}, \dots, \theta^{(T)}$.

Sampling the auxiliary variables: This step is the “backward filtering” pass. For the stationary PGDS in its steady state, we first compute ζ^* and draw $(l_{\cdot,k}^{(T+1)} | -) \sim \text{Pois}(\zeta^* \tau_0 \theta_k^{(T)})$. For the other variants of the model, we set $l_{\cdot,k}^{(T+1)} = \zeta^{(T+1)} = 0$. Then, working backward from $t = T, \dots, 2$, we draw

$$(l_{k\cdot}^{(t)} | - \setminus \Theta^{(\geq t)}) \sim \text{CRT}(y_{\cdot,k}^{(t)} + l_{\cdot,k}^{(t+1)}, \tau_0 \sum_{k_2=1}^K \pi_{kk_2} \theta_{k_2}^{(t-1)}) \text{ and} \quad (8)$$

$$(l_{k_1}^{(t)}, \dots, l_{k_K}^{(t)} | - \setminus \Theta^{(\geq t)}) \sim \text{Mult}(l_{k\cdot}^{(t)}, (\frac{\pi_{k_1} \theta_1^{(t-1)}}{\sum_{k_2=1}^K \pi_{kk_2} \theta_{k_2}^{(t-1)}}, \dots, \frac{\pi_{k_K} \theta_K^{(t-1)}}{\sum_{k_2=1}^K \pi_{kk_2} \theta_{k_2}^{(t-1)}})). \quad (9)$$

After using equations 8 and 9 for all $k = 1, \dots, K$, we then set $l_{\cdot,k}^{(t)} = \sum_{k_1=1}^K l_{k_1 k}^{(t)}$. For the non-steady-state variants, we also set $\zeta^{(t)} = \ln(1 + \frac{\delta^{(t)}}{\tau_0} + \zeta^{(t+1)})$; for the steady-state variant, we set $\zeta^{(t)} = \zeta^*$.

Sampling Θ : We sample Θ from its conditional posterior by performing a “forward sampling” pass, starting with $\theta^{(1)}$. Conditioned on the values of $l_{\cdot,k}^{(2)}, \dots, l_{\cdot,k}^{(T+1)}$ and $\zeta^{(2)}, \dots, \zeta^{(T+1)}$ obtained via the “backward filtering” pass, we sample forward from $t = 1, \dots, T$, using the following equations:

$$(\theta_k^{(1)} | - \setminus \Theta) \sim \text{Gam}(y_{\cdot,k}^{(1)} + l_{\cdot,k}^{(2)} + \tau_0 \nu_k, \tau_0 + \delta^{(1)} + \zeta^{(2)} \tau_0) \text{ and} \quad (10)$$

$$(\theta_k^{(t)} | - \setminus \Theta^{(\geq t)}) \sim \text{Gam}(y_{\cdot,k}^{(t)} + l_{\cdot,k}^{(t+1)} + \tau_0 \sum_{k_2=1}^K \pi_{kk_2} \theta_{k_2}^{(t-1)}, \tau_0 + \delta^{(t)} + \zeta^{(t+1)} \tau_0). \quad (11)$$

Sampling Π : The alternative model specification, with Θ marginalized out, assumes that $(l_{1k}^{(t)}, \dots, l_{Kk}^{(t)}) \sim \text{Mult}(l_{\cdot,k}^{(t)}, (\pi_{1k}, \dots, \pi_{Kk}))$. Therefore, via Dirichlet–multinomial conjugacy,

$$(\pi_k | - \setminus \Theta) \sim \text{Dir}(\nu_1 \nu_k + \sum_{t=1}^T l_{1k}^{(t)}, \dots, \xi \nu_k + \sum_{t=1}^T l_{kk}^{(t)}, \dots, \nu_K \nu_k + \sum_{t=1}^T l_{Kk}^{(t)}). \quad (12)$$

Sampling ν and ξ : We use the alternative model specification to obtain closed-form conditional posteriors for ν_k and ξ . First, we marginalize over π_k to obtain a Dirichlet–multinomial distribution. When augmented with a beta-distributed auxiliary variable, the Dirichlet–multinomial distribution is proportional to the negative binomial distribution [26]. We draw such an auxiliary variable, which we use, along with negative binomial augmentation schemes, to derive closed-form conditional posteriors for ν_k and ξ . We provide these posteriors, along with their derivations, in the supplementary material.

We also provide the conditional posteriors for the remaining model parameters— $\Phi, \delta^{(1)}, \dots, \delta^{(T)}$, and β —which we obtain via Dirichlet–multinomial, gamma–Poisson, and gamma–gamma conjugacy.

4 Experiments

In this section, we compare the predictive performance of the PGDS to that of the LDS and that of gamma process dynamic Poisson factor analysis (GP-DPFA) [22]. GP-DPFA models a single count in Y as $y_v^{(t)} \sim \text{Pois}(\sum_{k=1}^K \lambda_k \phi_{vk} \theta_k^{(t)})$, where each component’s time-step factors evolve as a simple gamma Markov chain, independently of those belonging to the other components: $\theta_k^{(t)} \sim \text{Gam}(\theta_k^{(t-1)}, c^{(t)})$. We consider the stationary variants of all three models.² We used five data sets, and tested each model on two time-series prediction tasks: smoothing—i.e., predicting $y_v^{(t)}$ given

¹Several software packages contain fast implementations of the Lambert W function.

²We used the `pykalman` Python library for the LDS and implemented GP-DPFA ourselves.

$y_v^{(1)}, \dots, y_v^{(t-1)}, y_v^{(t+1)}, \dots, y_v^{(T)}$ —and forecasting—i.e., predicting $y_v^{(T+s)}$ given $y_v^{(1)}, \dots, y_v^{(T)}$ for some $s \in \{1, 2, \dots\}$ [27]. We provide brief descriptions of the data sets below before reporting results.

Global Database of Events, Language, and Tone (GDEL T): GDEL T is an international relations data set consisting of country-to-country interaction events of the form “country i took action a toward country j at time t ,” extracted from news corpora. We created five count matrices, one for each year from 2001 through 2005. We treated directed pairs of countries $i \rightarrow j$ as features and counted the number of events for each pair during each day. We discarded all pairs with fewer than twenty-five total events, leaving $T = 365$, around $V \approx 9,000$, and three to six million events for each matrix.

Integrated Crisis Early Warning System (ICEWS): ICEWS is another international relations event data set extracted from news corpora. It is more highly curated than GDEL T and contains fewer events. We therefore treated undirected pairs of countries $i \leftrightarrow j$ as features. We created three count matrices, one for 2001–2003, one for 2004–2006, and one for 2007–2009. We counted the number of events for each pair during each three-day time step, and again discarded all pairs with fewer than twenty-five total events, leaving $T = 365$, around $V \approx 3,000$, and 1.3 to 1.5 million events for each matrix.

State-of-the-Union transcripts (SOTU): The SOTU corpus contains the text of the annual SOTU speech transcripts from 1790 through 2014. We created a single count matrix with one column per year. After discarding stopwords, we were left with $T = 225$, $V = 7,518$, and 656,949 tokens.

DBLP conference abstracts (DBLP): DBLP is a database of computer science research papers. We used the subset of this corpus that Acharya et al. used to evaluate GP-DPFA [22]. This subset corresponds to a count matrix with $T = 14$ columns, $V = 1,771$ unique word types, and 13,431 tokens.

NIPS corpus (NIPS): The NIPS corpus contains the text of every NIPS conference paper from 1987 to 2003. We created a single count matrix with one column per year. We treated unique word types as features and discarded all stopwords, leaving $T = 17$, $V = 9,836$, and 3.1 million tokens.

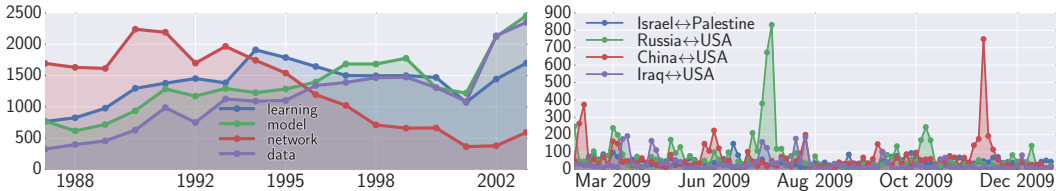


Figure 3: $y_v^{(t)}$ over time for the top four features in the NIPS (left) and ICEWS (right) data sets.

Experimental design: For each matrix, we created four masks indicating some randomly selected subset of columns to treat as held-out data. For the event count matrices, we held out six (non-contiguous) time steps between $t = 2$ and $t = T - 3$ to test the models’ smoothing performance, as well as the last two time steps to test their forecasting performance. The other matrices have fewer time steps. For the SOTU matrix, we therefore held out five time steps between $t = 2$ and $t = T - 2$, as well as $t = T$. For the NIPS and DBLP matrices, which contain substantially fewer time steps than the SOTU matrix, we held out three time steps between $t = 2$ and $t = T - 2$, as well as $t = T$.

For each matrix, mask, and model combination, we ran inference four times.³ For the PGDS and GP-DPFA, we performed 6,000 Gibbs sampling iterations, imputing the missing counts from the “smoothing” columns at the same time as sampling the model parameters. We then discarded the first 4,000 samples and retained every hundredth sample thereafter. We used each of these samples to predict the missing counts from the “forecasting” columns. We then averaged the predictions over the samples. For the LDS, we ran EM to learn the model parameters. Then, given these parameter values, we used the Kalman filter and smoother [1] to predict the held-out data. In practice, for all five data sets, V was too large for us to run inference for the LDS, which is $\mathcal{O}((K + V)^3)$ [2], using all V features. We therefore report results from two independent sets of experiments: one comparing all three models using only the top $V = 1,000$ features for each data set, and one comparing the PGDS to just GP-DPFA using all the features. The first set of experiments is generous to the LDS because the Poisson distribution is well approximated by the Gaussian distribution when its mean is large.

³For the PGDS and GP-DPFA we used $K = 100$. For the PGDS, we set $\tau_0 = 1, \gamma_0 = 50, \eta_0 = \epsilon_0 = 0.1$. We set the hyperparameters of GP-DPFA to the values used by Acharya et al. [22]. For the LDS, we used the default hyperparameters for `pykalman`, and report results for the best-performing value of $K \in \{5, 10, 25, 50\}$.

Table 1: Results for the smoothing (“S”) and forecasting (“F”) tasks. For both error measures, lower values are better. We also report the number of time steps T and the burstiness \hat{B} of each data set.

	T	\hat{B}	Task	Mean Relative Error (MRE)			Mean Absolute Error (MAE)		
				PGDS	GP-DPFA	LDS	PGDS	GP-DPFA	LDS
GDELT	365	1.27	S	2.335 ± 0.19	2.951 ± 0.32	3.493 ± 0.53	9.366 ± 2.19	9.278 ± 2.01	10.098 ± 2.39
			F	2.173 ± 0.41	2.207 ± 0.42	2.397 ± 0.29	7.002 ± 1.43	7.095 ± 1.67	7.047 ± 1.25
ICEWS	365	1.10	S	0.808 ± 0.11	0.877 ± 0.12	1.023 ± 0.15	2.867 ± 0.56	2.872 ± 0.56	3.104 ± 0.60
			F	0.743 ± 0.17	0.792 ± 0.17	0.937 ± 0.31	1.788 ± 0.47	1.894 ± 0.50	1.973 ± 0.62
SOTU	225	1.45	S	0.233 ± 0.01	0.238 ± 0.01	0.260 ± 0.01	0.408 ± 0.01	0.414 ± 0.01	0.448 ± 0.00
			F	0.171 ± 0.00	0.173 ± 0.00	0.225 ± 0.01	0.323 ± 0.00	0.314 ± 0.00	0.370 ± 0.00
DBLP	14	1.64	S	0.417 ± 0.03	0.422 ± 0.05	0.405 ± 0.05	0.771 ± 0.03	0.782 ± 0.06	0.831 ± 0.01
			F	0.322 ± 0.00	0.323 ± 0.00	0.369 ± 0.06	0.747 ± 0.01	0.715 ± 0.00	0.943 ± 0.07
NIPS	17	0.33	S	0.415 ± 0.07	0.392 ± 0.07	1.609 ± 0.43	29.940 ± 2.95	28.138 ± 3.08	108.378 ± 15.44
			F	0.343 ± 0.01	0.312 ± 0.00	0.642 ± 0.14	62.839 ± 0.37	52.963 ± 0.52	95.495 ± 10.52

Results: We used two error measures—mean relative error (MRE) and mean absolute error (MAE)—to compute the models’ smoothing and forecasting scores for each matrix and mask combination. We then averaged these scores over the masks. For the data sets with multiple matrices, we also averaged the scores over the matrices. The two error measures differ as follows: MRE accommodates the scale of the data, while MAE does not. This is because relative error—which we define as $\frac{|y_v^{(t)} - \hat{y}_v^{(t)}|}{1 + y_v^{(t)}}$,

where $y_v^{(t)}$ is the true count and $\hat{y}_v^{(t)}$ is the prediction—divides the absolute error by the true count and thus penalizes overpredictions more harshly than underpredictions. MRE is therefore an especially natural choice for data sets that are bursty—i.e., data sets that exhibit short periods of activity that far exceed their mean. Models that are robust to these kinds of overdispersed temporal patterns are less likely to make overpredictions following a burst, and are therefore rewarded accordingly by MRE.

In table 1, we report the MRE and MAE scores for the experiments using the top $V = 1,000$ features. We also report the average burstiness of each data set. We define the burstiness of feature v in matrix Y to be $\hat{B}_v = \frac{1}{T-1} \sum_{t=1}^{T-1} \frac{|y_v^{(t+1)} - y_v^{(t)}|}{\hat{\mu}_v}$, where $\hat{\mu}_v = \frac{1}{T} \sum_{t=1}^T y_v^{(t)}$. For each data set, we calculated the burstiness of each feature in each matrix, and then averaged these values to obtain an average burstiness score \hat{B} . The PGDS outperformed the LDS and GP-DPFA on seven of the ten prediction tasks when we used MRE to measure the models’ performance; when we used MAE, the PGDS outperformed the other models on five of the tasks. In the supplementary material, we also report the results for the experiments comparing the PGDS to GP-DPFA using all the features. The superiority of the PGDS over GP-DPFA is even more pronounced in these results. We hypothesize that the difference between these models is related to the burstiness of the data. For both error measures, the only data set for which GP-DPFA outperformed the PGDS on both tasks was the NIPS data set. This data set has a substantially lower average burstiness score than the other data sets. We provide visual evidence in figure 3, where we display $y_v^{(t)}$ over time for the top four features in the NIPS and ICEWS data sets. For the former, the features evolve smoothly; for the latter, they exhibit bursts of activity.

Exploratory analysis: We also explored the latent structure inferred by the PGDS. Because its parameters are positive, they are easy to interpret. In figure 1, we depict three components inferred from the NIPS data set. By examining the time-step factors and feature factors for these components, we see that they capture the decline of research on neural networks between 1987 and 2003, as well as the rise of Bayesian methods in machine learning. These patterns match our prior knowledge.

In figure 4, we depict the three components with the largest component weights inferred by the PGDS from the 2003 GDELT matrix. The top component is in blue, the second is in green, and the third is in red. For each component, we also list the sixteen features (directed pairs of countries) with the largest feature factors. The top component (blue) is most active in March and April, 2003. Its features involve USA, Iraq (IRQ), Great Britain (GBR), Turkey (TUR), and Iran (IRN), among others. This component corresponds to the 2003 invasion of Iraq. The second component (green) exhibits a noticeable increase in activity immediately after April, 2003. Its top features involve Israel (ISR), Palestine (PSE), USA, and Afghanistan (AFG). The third component exhibits a large burst of activity

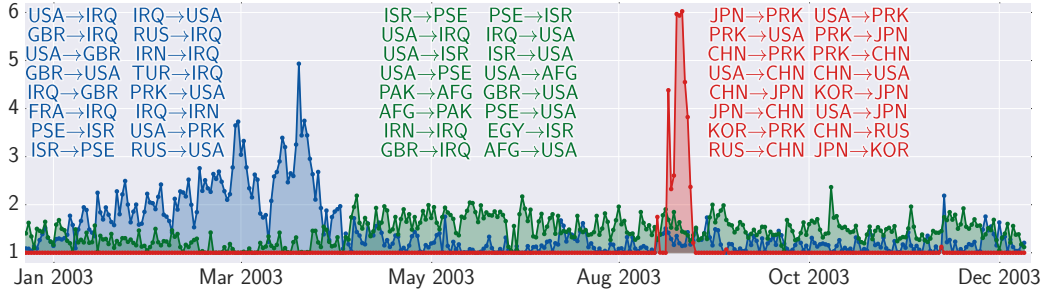


Figure 4: The time-step factors for the top three components inferred by the PGDS from the 2003 GDELT matrix. The top component is in blue, the second is in green, and the third is in red. For each component, we also list the features (directed pairs of countries) with the largest feature factors.

in August, 2003, but is otherwise inactive. Its top features involve North Korea (PRK), South Korea (KOR), Japan (JPN), China (CHN), Russia (RUS), and USA. This component corresponds to the six-party talks—a series of negotiations between these six countries for the purpose of dismantling North Korea’s nuclear program. The first round of talks occurred during August 27–29, 2003.

In figure 5, we also show the component weights for the top ten components, along with the corresponding subset of the transition matrix II. There are two components with weights greater than one: the components that are depicted in blue and green in figure 4. The transition weights in the corresponding rows of II are also large, meaning that other components are likely to transition to them. As we mentioned previously, the GDELT data set was extracted from news corpora. Therefore, patterns in the data primarily reflect patterns in media coverage of international affairs. We therefore interpret the latent structure inferred by the PGDS in the following way: in 2003, the media briefly covered various major events, including the six-party talks, before quickly returning to a backdrop of the ongoing Iraq war and Israeli–Palestinian relations. By inferring the kind of transition structure depicted in figure 5, the PGDS is able to model persistent, long-term temporal patterns while accommodating the burstiness often inherent to real-world count data. This ability is what enables the PGDS to achieve superior predictive performance over the LDS and GP-DPFA.

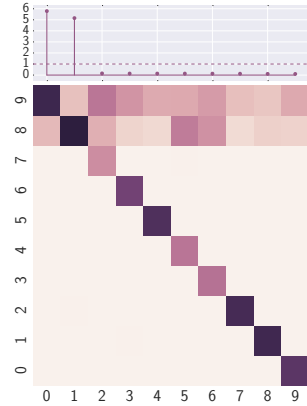


Figure 5: The latent transition structure inferred by the PGDS from the 2003 GDELT matrix. Top: The component weights for the top ten components, in decreasing order from left to right; two of the weights are greater than one. Bottom: The transition weights in the corresponding subset of the transition matrix. This structure means that all components are likely to transition to the top two components.

5 Summary

We introduced the Poisson–gamma dynamical system (PGDS)—a new Bayesian nonparametric model for sequentially observed multivariate count data. This model supports the expressive transition structure of the linear dynamical system, and naturally handles overdispersed data. We presented a novel MCMC inference algorithm that remains efficient for high-dimensional data sets, advancing recent work on augmentation schemes for inference in negative binomial models. Finally, we used the PGDS to analyze five real-world data sets, demonstrating that it exhibits superior smoothing and forecasting performance over two baseline models and infers highly interpretable latent structure.

Acknowledgments

We thank David Belanger, Roy Adams, Kostis Gourgoulis, Ben Marlin, Dan Sheldon, and Tim Vieira for many helpful conversations. This work was supported in part by the UMass Amherst CIIR and in part by NSF grants SBE-0965436 and IIS-1320219. Any opinions, findings, conclusions, or recommendations are those of the authors and do not necessarily reflect those of the sponsors.

References

- [1] R. E. Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35–45, 1960.
- [2] Z. Ghahramani and S. T. Roweis. Learning nonlinear dynamical systems using an EM algorithm. In *Advances in Neural Information Processing Systems*, pages 431–437, 1998.
- [3] S. S. Haykin. *Kalman Filtering and Neural Networks*. 2001.
- [4] P. McCullagh and J. A. Nelder. *Generalized linear models*. 1989.
- [5] M. G. Bulmer. On fitting the Poisson lognormal distribution to species-abundance data. *Biometrics*, pages 101–110, 1974.
- [6] D. M. Blei and J. D. Lafferty. Dynamic topic models. In *Proceedings of the 23rd International Conference on Machine Learning*, pages 113–120, 2006.
- [7] L. Charlin, R. Ranganath, J. McInerney, and D. M. Blei. Dynamic Poisson factorization. In *Proceedings of the 9th ACM Conference on Recommender Systems*, pages 155–162, 2015.
- [8] J. H. Macke, L. Buesing, J. P. Cunningham, B. M. Yu, K. V. Krishna, and M. Sahani. Empirical models of spiking in neural populations. In *Advances in Neural Information Processing Systems*, pages 1350–1358, 2011.
- [9] J. Kleinberg. Bursty and hierarchical structure in streams. *Data Mining and Knowledge Discovery*, 7(4):373–397, 2003.
- [10] M. Zhou and L. Carin. Augment-and-conquer negative binomial processes. In *Advances in Neural Information Processing Systems*, pages 2546–2554, 2012.
- [11] R. Corless, G. Gonnet, D. E. G. Hare, D. J. Jeffrey, and D. E. Knuth. On the LambertW function. *Advances in Computational Mathematics*, 5(1):329–359, 1996.
- [12] J. Canny. GaP: A factor model for discrete data. In *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 122–129, 2004.
- [13] A. T. Cemgil. Bayesian inference for nonnegative matrix factorisation models. *Computational Intelligence and Neuroscience*, 2009.
- [14] M. Zhou, L. Hannah, D. Dunson, and L. Carin. Beta-negative binomial process and Poisson factor analysis. In *Proceedings of the 15th International Conference on Artificial Intelligence and Statistics*, 2012.
- [15] P. Gopalan, J. Hofman, and D. Blei. Scalable recommendation with Poisson factorization. In *Proceedings of the 31st Conference on Uncertainty in Artificial Intelligence*, 2015.
- [16] T. S. Ferguson. A Bayesian analysis of some nonparametric problems. *The Annals of Statistics*, 1(2):209–230, 1973.
- [17] M. Zhou. Infinite edge partition models for overlapping community detection and link prediction. In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics*, pages 1135–1143, 2015.
- [18] M. Zhou, Y. Cong, and B. Chen. Augmentable gamma belief networks. *Journal of Machine Learning Research*, 17(163):1–44, 2016.
- [19] M. E. Basbug and B. Engelhardt. Hierarchical compound Poisson factorization. In *Proceedings of the 33rd International Conference on Machine Learning*, 2016.
- [20] R. M. Adelson. Compound Poisson distributions. *OR*, 17(1):73–75, 1966.
- [21] M. Zhou and L. Carin. Negative binomial process count and mixture modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(2):307–320, 2015.
- [22] A. Acharya, J. Ghosh, and M. Zhou. Nonparametric Bayesian factor analysis for dynamic count matrices. *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics*, 2015.
- [23] J. F. C. Kingman. *Poisson Processes*. Oxford University Press, 1972.
- [24] D. B. Dunson and A. H. Herring. Bayesian latent variable models for mixed discrete outcomes. *Biostatistics*, 6(1):11–25, 2005.
- [25] W. J. Rugh. *Linear System Theory*. Pearson, 1995.
- [26] M. Zhou. Nonparametric Bayesian negative binomial factor analysis. arXiv:1604.07464.
- [27] J. Durbin and S. J. Koopman. *Time Series Analysis by State Space Methods*. Oxford University Press, 2012.