



Published in final edited form as:

JMLR Workshop Conf Proc. 2016 June ; 48: 2445–2453.

Square Root Graphical Models: Multivariate Generalizations of Univariate Exponential Families that Permit Positive Dependencies

David I. Inouye, Pradeep Ravikumar, and Inderjit S. Dhillon

Dept. of Computer Science, University of Texas, Austin, TX 78712, USA

David I. Inouye: DINOUYE@CS.UTEXAS.EDU; Pradeep Ravikumar: PRADEEPR@CS.UTEXAS.EDU; Inderjit S. Dhillon: INDERJIT@CS.UTEXAS.EDU

Abstract

We develop Square Root Graphical Models (SQR), a novel class of parametric graphical models that provides multivariate generalizations of univariate exponential family distributions. Previous multivariate graphical models (Yang et al., 2015) did not allow positive dependencies for the exponential and Poisson generalizations. However, in many real-world datasets, variables clearly have positive dependencies. For example, the airport delay time in New York—modeled as an exponential distribution—is positively related to the delay time in Boston. With this motivation, we give an example of our model class derived from the univariate exponential distribution that allows for almost arbitrary positive and negative dependencies with only a mild condition on the parameter matrix—a condition akin to the positive definiteness of the Gaussian covariance matrix. Our Poisson generalization allows for both positive and negative dependencies without any constraints on the parameter values. We also develop parameter estimation methods using node-wise regressions with ℓ_1 regularization and likelihood approximation methods using sampling. Finally, we demonstrate our exponential generalization on a synthetic dataset and a real-world dataset of airport delay times.

1. Introduction

Gaussian, binary and discrete undirected graphical models—or Markov Random Fields (MRF)—have become popular for compactly modeling and studying the structural dependencies between high-dimensional continuous, binary and categorical data respectively (Friedman et al., 2008; Hsieh et al., 2014; Banerjee et al., 2008; Ravikumar et al., 2010; Jalali et al., 2010). However, real-world data does not often fit the assumption that variables come from Gaussian or discrete distributions. For example, word counts in documents are nonnegative integers with many zero values and hence are more appropriately modeled by the Poisson distribution. Yet, an independent Poisson distribution would be insufficient because words are often either positively or negatively related to other words—e.g. the words “machine” and “learning” would often co-occur together in ICML papers (positive dependency) whereas the words “deep” and “kernel” would rarely co-occur since they

usually refer to different topics (negative dependency). Thus, a Poisson-like model that allows for dependencies between words is desirable. As another example, the delay times at airports are nonnegative continuous values that are more closely modeled by an exponential distribution than a Gaussian distribution but an independent exponential distribution is insufficient because delays at different airports are often related (and sometimes causally related)—e.g. if a flight from Los Angeles, CA (LAX) to San Francisco, CA (SFO) is delayed then it is likely that the return flight of the same airplane will also be delayed. Other examples of non-Gaussian and non-discrete data include high-throughput gene sequencing count data, crime statistics, website visits, survival times, call times and delay times.

Though univariate distributions for these types of data have been studied quite extensively, multivariate generalizations have only been given limited attention. One basic approach to forming dependent multivariate distributions is to assume that the marginal distributions are exponentially distributed (Marshall & Olkin, 1967; Embrechts et al., 2003) or Poisson distributed (Karlis, 2003). This idea is related to copula-based models (Bickel et al., 2009) in which a probability distribution is decomposed into the univariate marginal distributions and a copula distribution on the unit hypercube that models the dependency between variables. However, the exponential model in (Marshall & Olkin, 1967; Embrechts et al., 2003) gives rise to a distribution that is composed of a continuous distribution and a singular distribution, which seems unusual and unlikely for general real-world situations. The multivariate Poisson distribution (Karlis, 2003) is based on the sum of independent Poisson variables and can only model *positive* dependencies. The copula versions of the multivariate Poisson distribution have significant issues related to non-identifiability because the Poisson distribution has a discrete domain (Genest & Neslehova, 2007). There has also been some recent work on semi-parametric graphical models (Liu et al., 2009) that use Gaussian copulas to relax the assumption of Gaussianity but these models are not parametric and only consider continuous real-valued data.

Another line of work assumes that the node conditional distributions—i.e. one variable given the values of all the other variables—are univariate exponential families¹ and determines under what conditions a joint distribution exists that is consistent with these node conditional distributions. Besag (1974) developed this multivariate distribution for pairwise dependencies, and Yang et al. (2015) extended this model to n-wise dependencies. Yang et al. (2015) also developed and analyzed an M-estimator based on ℓ_1 regularized node-wise regressions to recover the graphical model structure with high probability. Unfortunately, these models only allowed *negative* dependencies in the case of the exponential and Poisson distributions. Yang et al. (2013) proposed three modifications to the original Poisson model to allow positive dependencies but these modifications alter the Poisson base distribution or require the specification of unintuitive hyperparameters. Allen & Liu (2013) allowed positive dependencies by only requiring the Local Markov property rather than a consistent joint distribution that would have Global Markov properties. In a different approach, Inouye et al. (2015) altered the Poisson generalization by assuming the length of the vector is fixed or known similar to the multinomial distribution in which the number of trials is known.

¹See (Wainwright & Jordan, 2008) for an introduction to exponential families.

This allows a joint distribution that is decomposed into the marginal distribution of vector length and the distribution of the vector direction given the length. While the model in (Inouye et al., 2015) allowed for both positive and negative dependencies, the joint distribution needed to be modified by an ad hoc scalar weighting function to avoid very low likelihood values for vectors of long length—i.e. documents with many words.

Therefore, we develop a novel parametric generalization of univariate exponential family distributions with non-negative sufficient statistics—e.g. Gaussian, Poisson and exponential—that allows for *both positive and negative dependencies*. We call this novel class of multivariate distributions Square Root Graphical Models (SQR) because the square root function is fundamentally important as will be described in future sections. SQR models have a simple parametric form without needing to specify any hyperparameters and can be fit using ℓ_1 -regularized node-wise regressions similar to previous work (Yang et al., 2015). The independent model—e.g. independent Poisson or exponential—is merely a special case of this class unlike in (Yang et al., 2013). We show that the normalizability of the distribution puts little to no restriction on the values of the parameters, and thus SQR models give a very flexible multivariate generalization of well-known univariate distributions.

Notation

Let p and n denote the number of dimensions and number data instances respectively. We will generally use uppercase letters for matrices (e.g. Φ , X), boldface lowercase letters for vectors (i.e. \mathbf{x} , $\boldsymbol{\phi}$) and lowercase letters for scalar values (i.e. x , ϕ). Let \mathbb{R}_+ denote the set of nonnegative real numbers and \mathbb{Z}_+ denote the set of nonnegative integers.

2. Background

To motivate the form of our model class, we present a brief background on the graphical model class as in (Besag, 1974; Yang et al., 2015; 2013). Let $T(\mathbf{x})$ and $B(\mathbf{x})$ be the sufficient statistics and log base measure respectively of the base univariate exponential family and let $\mathcal{D} \subseteq \mathbb{R}_+^p$ be the domain of the random vector. We will denote $T(\mathbf{x}): \mathbb{R}^p \rightarrow \mathbb{R}^p$ to be the entry-wise application of the sufficient statistic function to each entry in the vector \mathbf{x} . With this notation, the previous class of graphical models can be defined as (Yang et al., 2015):

$$\Pr(\mathbf{x}|\boldsymbol{\theta}, \Phi) = \exp(\boldsymbol{\theta}^T T(\mathbf{x}) + T(\mathbf{x})^T \Phi T(\mathbf{x}) + \sum_{s=1}^p B(x_s) - A(\boldsymbol{\theta}, \Phi)) \quad (1)$$

$$A(\boldsymbol{\theta}, \Phi) = \int_{\mathcal{D}} \exp\left(\boldsymbol{\theta}^T T(\mathbf{x}) + T(\mathbf{x})^T \Phi T(\mathbf{x}) + \sum_{s=1}^p B(x_s)\right) d\mu(\mathbf{x}), \quad (2)$$

where $A(\boldsymbol{\theta}, \Phi)$ is the log partition function (i.e. log normalization constant) which is required for probability normalization, $\Phi \in \mathbb{R}^{p \times p}$ is symmetric with zeros along the diagonal and μ is either the standard Lebesgue measure or the counting measure depending on whether the domain \mathcal{D} is continuous or discrete. The only difference from a fully independent model is the quadratic interaction term $T(\mathbf{x})^T \Phi T(\mathbf{x})$ —i.e. $\mathcal{O}(T(\mathbf{x})^2)$ —which is

why the exponential and Poisson cases do not admit positive dependencies as will be described in later sections.

We will review the exponential instantiation of this previous graphical model in which the domain $\mathcal{D} \subseteq \mathbb{R}_+^p$, $T(x) = x$ and $B(x) = 0$. Suppose there is even one positive entry in Φ denoted ϕ_{st} . Then as $x \rightarrow \infty$ the positive quadratic term $x_s \phi_{st} x_t$ will dominate the linear term $\theta^T x$ and thus the log partition function will diverge (i.e. $A(\theta, \Phi) \rightarrow \infty$). Thus, $\Phi_{st} = 0$ is required for a consistent joint distribution. Similarly, in the case of the Poisson distribution where the domain $\mathcal{D} \in \mathbb{Z}_+^p$, $T(x) = x$ and $B(x) = -\ln(x!)$, suppose there is even one positive entry ϕ_{st} . The quadratic term $x_s \phi_{st} x_t$ will dominate the linear term *and* the log base measure term which is $O(x \ln(x))$; thus, $\Phi_{st} = 0$ is also required for the Poisson distribution.

In an attempt to allow positive dependencies for the Poisson distribution, Yang et al. (2013) developed three variants of the Poisson graphical model defined above. First, they developed a Truncated Poisson Graphical Model (TPGM) that kept the same parametric form but merely truncated the usual infinite domain to the finite domain $\mathcal{D}_{\text{TPGM}} = \{x \in \mathbb{Z}_+^p : x_s \leq R\}$. However, a user must a priori specify the truncation value R and thus TPGM is unnatural for normal count data that could be infinite. In addition, because of the quadratic term, even though the domain is finite, the quadratic term can dominate and push most of the mass near the boundary of the domain (Yang et al., 2013). The second proposal was to change the base measure from $\ln(x!)$ to x^2 . This proposal, however, gives the distribution Gaussian-like quadratic tails rather than the thicker tails of the Poisson distribution. Finally, the last proposal modified the sufficient statistic $T(x)$ to decrease from linear to constant as x increases. Similar to TPGM, this third proposal requires the a priori specification of two cutoff parameters (R_1, R_2) and behaves similarly to TPGM after the second cutoff point because the base measure of $-\ln(x!)$ will quickly make the probability approach 0 once the sufficient statistics become constant.

In a somewhat different direction, Inouye et al. (2015) proposed a variant called Fixed-Length Poisson MRF (LPMRF) that modifies the domain of the distribution assuming the length of the vector $L = \|x\|_1$ is fixed, i.e. $\mathcal{D} = \{x \in \mathbb{Z}_+^p : \|x\|_1 = L\}$. Because the domain is finite as in TPGM, the distribution is normalizable even with positive dependencies. However, as with TPGM, the quadratic term in the parametric form dominates the distribution if L is large, and thus Inouye et al. (2015) modify the distribution by introducing a weighting function that decreases the quadratic term as L increases. All of these variants of the Poisson graphical model attempt to deal with the quadratic interaction term in different ways but all of them significantly change the distribution/domain and often require the specification of new unintuitive hyperparameters to allow for positive dependencies. Also, according to the authors' best knowledge, no variants of the exponential graphical model have been proposed to allow for positive dependencies. Therefore, we propose a novel graphical model class that alleviates the problem with the quadratic interaction term and provides both exponential and Poisson graphical models that allow *positive and negative* dependencies.

3. Square Root Graphical Model

The amazingly simple yet helpful change from the previous graphical model class in Eqn. 1 is that we take the square root of the sufficient statistics in the interaction term. Essentially, this makes the interaction term linear in the sufficient statistics $\mathcal{O}(T(x))$ rather than quadratic $\mathcal{O}(T(x)^2)$ as in Eqn. 1. This change avoids the problem of the quadratic term overcoming the other terms while allowing both positive and negative dependencies. More formally, given any univariate exponential family with nonnegative sufficient statistics $T(x) \geq 0$, we can define the Square Root Graphical Model (SQR) class as follows:

$$Pr(\mathbf{x}|\boldsymbol{\theta}, \Phi) = \exp\left(\boldsymbol{\theta}^T \sqrt{T(\mathbf{x})} + \sqrt{T(\mathbf{x})}^T \Phi \sqrt{T(\mathbf{x})} + \sum_s B(x_s) - A(\Phi)\right) \quad (3)$$

$$A(\mathbf{x}|\boldsymbol{\theta}, \Phi) = \int_{\mathcal{D}} \exp\left(\boldsymbol{\theta}^T \sqrt{T(\mathbf{x})} + \sqrt{T(\mathbf{x})}^T \Phi \sqrt{T(\mathbf{x})} + \sum_s B(x_s)\right) d\mu(\mathbf{x}), \quad (4)$$

where $\sqrt{T(\mathbf{x})}$ is an entry-wise square root except when $T(x) = x^2$ in which case $\sqrt{T(x)} \equiv x$.² Figure 1 shows examples of the exponential and Poisson SQR distributions for no dependency, positive dependency and negative dependency. If $\boldsymbol{\theta} = 0$ and Φ is a diagonal matrix, then we recover an independent joint distribution so the SQR class of models can be seen as a direct relaxation of the independence assumption, similar to previous graphical models. In the next sections, we analyze some of the properties of SQR models including their conditional distributions.

3.1. SQR Conditional Distributions

We analyze two types of univariate conditional distributions of the SQR graphical models. The first is the standard *node* conditional distribution, i.e. the conditional distribution of one variable given the values for all other variables (see Fig. 2). The second is what we will call the *radial* conditional distribution in which the *unit direction* is fixed but the length of the vector is unknown (see Fig. 2). The node conditional distribution is helpful for parameter estimation as described more fully in Sec. 3.3. The radial conditional distribution is important for understanding the form of the SQR distribution as well as providing a means to succinctly prove that the normalization constant is finite (i.e. the distribution is valid) as described in Sec. 3.2.

Node Conditional Distribution—The probability distribution of one variable x_s given all other variables $\mathbf{x}_{-s} = [x_1, x_2, \dots, x_{s-1}, x_{s+1}, \dots, x_p]$ is as follows:

²This nuance is important for the Gaussian SQR in Sec. 4.

$$Pr(x_s | \mathbf{x}_{-s}, \boldsymbol{\theta}, \Phi) \propto \exp \left\{ \phi_{ss} T(x_s) + \left(\theta_s + 2\phi_{-s}^T \sqrt{T(\mathbf{x}_{-s})} \right) \sqrt{T(x_s) + B(x_s)} \right\},$$

where $\phi_{-s} \in \mathbb{R}^{p-1}$ is the s -th column of Φ with the s -th entry removed. This conditional distribution can be reformulated as a new two parameter exponential family:

$$Pr(x_s | \mathbf{x}_{-s}, \boldsymbol{\theta}, \Phi) = \exp \left(\eta_1 \tilde{T}_1(x_s) + \eta_2 \tilde{T}_2(x_s) + B(x_s) - A_{\text{node}}(\boldsymbol{\eta}) \right) \quad (5)$$

$$A_{\text{node}}(\boldsymbol{\eta}) = \int_{\mathcal{X}} \exp \left(\eta_1 \tilde{T}_1(x_s) + \eta_2 \tilde{T}_2(x_s) + B(x_s) \right) d\mu(x_s), \quad (6)$$

where $\eta_1 = \phi_{ss}$, $\eta_2 = \theta_s + 2\phi_{-s}^T \sqrt{T(\mathbf{x}_{-s})}$, $\tilde{T}_1(x) = T(x)$ and $\tilde{T}_2(x) = \sqrt{T(x)}$. Note that this reduces to the base exponential family only if $\eta_2 = 0$ unlike the model in Eqn. 1 which, by construction, has node conditionals in the base exponential family. Examples of node conditional distributions for the exponential and Poisson SQR can be seen in Fig. 3. While these node conditionals are different from the base exponential family and hence slightly more difficult to use for parameter estimation as described later in Sec. 3.3, the benefit of almost arbitrary positive and negative dependencies significantly outweighs the cost of using SQR over previous graphical models.

Radial Conditional Distribution—For simplicity, let us assume w.l.o.g. that $T(\mathbf{x}) = \mathbf{x}$.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. We call this the *radial* conditional distribution:

$$\begin{aligned} Pr(\mathbf{x} = z\mathbf{v} | \mathbf{v}, \boldsymbol{\theta}, \Phi) \\ &\propto \exp \left(\boldsymbol{\theta}^T \sqrt{z\mathbf{v}} + \sqrt{z\mathbf{v}}^T \Phi \sqrt{z\mathbf{v}} + \sum_s B(zv_s) \right) \\ &\propto \exp \left(\left(\boldsymbol{\theta}^T \sqrt{\mathbf{v}} \right) \sqrt{z} + \left(\sqrt{\mathbf{v}}^T \Phi \sqrt{\mathbf{v}} \right) z + \sum_s B(zv_s) \right). \end{aligned}$$

The radial conditional distribution can be rewritten as a univariate exponential family:

$$Pr(z | \mathbf{v}, \boldsymbol{\theta}, \Phi) = \exp \left(\underbrace{\bar{\eta}_1 z + \bar{\eta}_2 \sqrt{z}}_{O(z)} + \underbrace{\tilde{B}_{\mathbf{v}}(z)}_{O(B(z))} - A_{\text{rad}}(\bar{\boldsymbol{\eta}}) \right) \quad (7)$$

³If T is not linear than we can merely reparameterize the distribution so that this is the case.

$$A_{\text{rad}}(\bar{\boldsymbol{\eta}}) = \int_{\mathcal{D}} \exp \left(\underbrace{\bar{\eta}_1 z + \bar{\eta}_2 \sqrt{z}}_{O(z)} + \underbrace{\tilde{B}_{\mathbf{v}}(z)}_{O(B(z))} \right) d\mu(z), \quad (8)$$

where $\bar{\eta}_1 = \sqrt{\mathbf{v}^T} \Phi \sqrt{\mathbf{v}}$, $\bar{\eta}_2 = \boldsymbol{\theta}^T \sqrt{\mathbf{v}}$ and $\tilde{B}_{\mathbf{v}}(z) = \sum_s (z v_s)$. Note that if the log base measure of the base exponential family is zero $B(x) = 0$, then the radial conditional is the same as the node conditional distribution because the modified base measure is also zero $\tilde{B}_{\mathbf{v}}(z) = 0$. If both $\boldsymbol{\theta} = 0$ and $B(x) = 0$, this actually reduces to the base exponential family. For example, the exponential distribution has $B(x) = 0$, and thus if we set $\boldsymbol{\theta} = 0$, the radial conditional of an exponential SQR is merely the exponential distribution. Other examples with a log base measure of zero include the Beta distribution and the gamma distribution with a known shape. For distributions in which the log base measure is not zero, the distribution will deviate from the node conditional distribution based on the relative difference between $B(x)$ and $\tilde{B}_{\mathbf{v}}(x)$. However, the important point even for distributions with nonzero log base measures is that the terms in the exponent grow at the same rate as the base exponential family—i.e. $O(z) + O(B(z))$. This helps to ensure that the radial conditional distribution is normalizable even as $z \rightarrow \infty$ since the base exponential family was normalizable. As an example, the Poisson distribution has the log base measure $B(x) = -\ln(x!)$ and thus $\tilde{B}_{\mathbf{v}}(x)$ is $O(-x \ln x)$ whereas the other terms $\bar{\eta}_1 z + \bar{\eta}_2 \sqrt{z}$ are only $O(z)$. This provides the intuition of why the Poisson SQR radial distribution is normalizable as will be explained in Sec. 4.2.

3.2. Normalization

Normalization of the distribution was the reason for the negative-only parameter restrictions of the exponential and Poisson distributions in the previous graphical models (Besag, 1974; Yang et al., 2015) as defined in Eqn. 1. However, we show that in the case of SQR models, normalization is much simpler to achieve and generally puts little to no restriction on the value of the parameters—thus allowing both positive and negative dependencies. For our derivations, let $\mathcal{V} = \{\mathbf{v}: \|\mathbf{v}\|_1 = 1, \mathbf{v} \in \mathbb{R}_+^p\}$ be the set of unit vectors in the positive orthant. The SQR log partition function $A(\Phi)$ can be decomposed into nested integrals over the unit direction and the one dimensional integral over scaling, denoted z .

$$A(\boldsymbol{\theta}, \Phi) = \ln \int_{\mathcal{V}} \int_{\mathcal{L}(\mathbf{v})} \exp \left(\boldsymbol{\theta}^T \sqrt{z \mathbf{v}} + \sqrt{z \mathbf{v}^T} \Phi \sqrt{z \mathbf{v}} + \sum_s B(z v_s) \right) d\mu(z) d\mathbf{v} \quad (9)$$

$$= \ln \int_{\mathcal{V}} \int_{\mathcal{L}(\mathbf{v})} \exp(\bar{\eta}_1(\mathbf{v}) z + \bar{\eta}_2(\mathbf{v}) \sqrt{z} + \sum_s B(z v_s)) d\mu(z) d\mathbf{v}, \quad (10)$$

where $\mathcal{L}(\mathbf{v}) = \{z \in \mathbb{R}_+ : z \mathbf{v} \in \mathcal{D}\}$, and μ and \mathcal{D} are defined as in Eqn. 2. Because \mathcal{V} is bounded, we merely need that the radial conditional distribution is normalizable (i.e.

$A_{\text{rad}}(\bar{\boldsymbol{\eta}}) < \infty$ from Eqn. 8) for the joint distribution to be normalizable. As suggested in Sec. 3.1, the radial conditional distribution is similar to the base exponential family and thus likely only has similar restrictions on parameter values as the base exponential family. In Sec. 4, we give examples for the exponential SQR and Poisson SQR distributions showing that this condition can be achieved with little or no restriction on the parameter values.

3.3. Parameter Estimation

For estimating the parameters Φ and $\boldsymbol{\theta}$, we follow the basic approach of (Ravikumar et al., 2010; Yang et al., 2015; 2013) and fit p ℓ -regularized node-wise regressions using the node conditional distributions described in Sec. 3.1. Thus, given a data matrix $X \in \mathbb{R}^{p \times n}$ we attempt to optimize the following convex function:

$$\arg \min_{\Phi} -\frac{1}{n} \sum_s \sum_i (\eta_{1si} x_{si} + \eta_{2si} \sqrt{x_{si}} + B(x_{si}) - A_{\text{node}}(\eta_{1si}, \eta_{2si})) + \lambda \|\Phi\|_{1, \text{off}}, \quad (11)$$

where $\eta_{1si} = \phi_{s,s}$, $\eta_{2si} = \theta + 2\phi_{-s}^T \sqrt{\mathbf{T}(\mathbf{x}_{-si})}$, $\|\Phi\|_{1, \text{off}} = \sum_{s \neq t} |\phi_{st}|$ is the ℓ -norm on the off diagonal elements and λ is a regularization parameter. Note that this can be trivially parallelized into p independent sub problems which allows for significantly faster computation as in (Inouye et al., 2015). Unlike previous graphical models (Yang et al., 2015) that were known to have closed-form solutions to the node conditional log partition function, the main difficulty for SQR graphical models is that the node conditional log partition function $A_{\text{node}}(\boldsymbol{\eta})$ is not known to have a closed form in general.

For the particular case of exponential SQR models, there is a closed-form solution for A_{node} using the error function as will be seen in Sec. 4.1 on exponential SQR models. More generally, because A_{node} is merely a one dimensional summation or integral, standard numerical approximations such as Gaussian quadrature could be used. Similarly, the gradient of ∇A_{node} could be numerically approximated by:

$$\nabla A_{\text{node}} = \frac{1}{\varepsilon} \left[\left(\hat{A}(\eta_1 + \varepsilon, \eta_2) - \hat{A}(\eta_1, \eta_2) \right), \left(\hat{A}(\eta_1, \eta_2 + \varepsilon) - \hat{A}(\eta_1, \eta_2) \right) \right], \quad (12)$$

where ε is a small step such as 0.001. Notice that to compute the function value and the gradient, only three 1D numerical integrations are needed. Another significant speedup that could be explored in future work would be to use a Newton-like method as in (Hsieh et al., 2014; Inouye et al., 2015), which optimize a quadratic approximation around the current iterate. Because these Newton-like methods only need a small number of Newton iterations to converge, the number of numerical integrations could be reduced significantly compared to gradient descent which often require thousands of iterations to converge.

3.4. Likelihood Approximation

We use Annealed Importance Sampling (AIS) (Neal, 2001) similar to the sampling used in (Inouye et al., 2015) for likelihood approximation. In particular, we need to approximate the

SQR log partition function $A(\boldsymbol{\theta}, \Phi)$ as in Eqn. 4. First, we derive a slice sample for the node conditionals in which the bounds for the slice can be computed in closed form. Second, we use the slice sampler to develop a Gibbs sampler for SQR models. Finally, we derive an annealed importance sampler (Neal, 2001) using the Gibbs sampler as the intermediate sampler by linearly combining the off-diagonal part of the parameter matrix Φ_{off} with the diagonal part Φ_{diag} —i.e. $\tilde{\Phi} = \gamma\Phi_{\text{off}} + \Phi_{\text{diag}}$. We also modify $\tilde{\boldsymbol{\theta}} = \gamma\boldsymbol{\theta}$ similarly. For each successive distribution, we linearly change γ from 0 to 1. Thus, we start by sampling from the base exponential family independent distribution $\prod_{s=1}^p \text{Pr}(x|\eta_{1s} = \phi_{ss}, \eta_{2s} = 0)$ and slowly move towards the final SQR distribution $\text{Pr}(x|\boldsymbol{\theta}, \Phi)$. We maintain the sample weights as defined in (Neal, 2001) and from these weights, we can compute an approximation to the log partition function (Neal, 2001).

4. Examples from Various Exponential Families

We give several examples of SQR graphical models in the following sections (however, it should be noted that we have been developing a *class* of graphical models for *any* univariate exponential family with nonnegative sufficient statistics). The main analysis for each case is determining what conditions on the parameter matrix Φ allow the joint distribution to be normalized. As described in Sec. 3.2, for SQR models, this merely reduces to determining when the radial conditional distribution is normalizable. We analyze the exponential and Poisson cases in later sections but first we give examples of the discrete and Gaussian SQR graphical models.

The discrete SQR graphical model—including the binary Ising model—is equivalent to the standard discrete graphical model because the sufficient statistics are indicator functions $T_s(x) = \mathbf{I}(x = s)$, $\forall s \in \mathcal{p}$ and the square root of an indicator function is merely the indicator function. Thus, in the discrete case, the discrete graphical model in (Ravikumar et al., 2010; Yang et al., 2015) is equivalent to the discrete SQR graphical model. For the Gaussian distribution, we can use the nonnegative Gaussian sufficient statistic $T(x) = x^2$. Thus, the Gaussian SQR graphical model is merely $\text{Pr}(x|\Phi) \propto \exp(\boldsymbol{\theta}^T x + x^T \Phi x)$, which by inspection

is clearly the standard Gaussian distribution where $\boldsymbol{\theta} = \Sigma^{-1}\boldsymbol{\mu}$ and $\Phi = -\frac{1}{2}\Sigma^{-1}$ is required to be negative definite.⁴ Thus, the Gaussian graphical model can be seen as a special case of SQR graphical models.

4.1. Exponential SQR Graphical Model

We consider what are the required conditions on the parameters $\boldsymbol{\theta}$ and Φ for the exponential SQR graphical model. If $\bar{\eta}_1$ is positive, the log partition function will diverge because even the end point $\lim_{z \rightarrow \infty} \exp(\bar{\eta}_1 z) \rightarrow \infty$. On the other hand, if $\bar{\eta}_1$ is negative, then the radial conditional distribution is similar in form to the exponential distribution and thus the log partition function will be finite because the negative linear term $\bar{\eta}_1 z$ dominates in the exponent as $z \rightarrow \infty$.⁵ See appendix for proof. Thus, the basic condition on Φ is:

⁴This is by the slightly nuanced definition of the square root operator in Eqn. 3 and 4 such that $\sqrt{x^2} \equiv x$ rather than $|x|$.

$$\Phi_{\text{Exp}} \in \left\{ \Phi: \sqrt{\mathbf{v}}^T \Phi \sqrt{\mathbf{v}} < 0, \forall \mathbf{v} \in \mathcal{V} \right\}. \quad (13)$$

Note that this allows both positive and negative dependencies. A sufficient condition is that Φ be negative definite—as is the case for Gaussian graphical models. However, negative definiteness is far from necessary because we only need negativity of the interaction term for vectors in the positive orthant. It may even be possible for Φ to be positive definite but Eqn. 13 be satisfied; however, we have not explored this idea.

For fitting the SQR model, the node conditional log partition function $A_{\text{Exp}}(\boldsymbol{\eta})$ has a closed-form solution:

$$A_{\text{Exp}}(\boldsymbol{\eta}) = \ln \left(\frac{\sqrt{\pi} \eta_1 \exp\left(\frac{-\eta_2^2}{4\eta_1}\right) \left(1 - \operatorname{erf}\left(\frac{-\eta_2}{2\sqrt{-\eta_1}}\right)\right)}{-2(-\eta_1)^{\frac{3}{2}}} - \frac{1}{\eta_1} \right),$$

where $\operatorname{erf}(\cdot)$ is the error function. The erf function shows up because of an initial substitution of $u = \sqrt{x}$ to transform the exponent into a quadratic form. Note that $\eta_1 < 0$ by the condition on Φ_{Exp} in Eqn. 13 above. The derivatives of A_{Exp} can also be computed in closed form for use in the parameter estimation algorithm.

4.2. Poisson SQR Graphical Model

The normalization analysis for Poisson SQR graphical model is also relatively simple but requires a more careful analysis than the exponential SQR graphical model. Let us consider the form of the Poisson radial conditional:

$Pr_{\text{rad}}(z|\mathbf{v}) \propto \exp\left(\bar{\eta}_1 z + \bar{\eta}_2 \sqrt{z} - \sum_s \ln((zv_s)!) \right)$. Note that the domain of z , denoted $\mathcal{D}_z = \{z \in \mathbb{Z}_+ : z\mathbf{v} \in \mathbb{Z}_+^p\}$, is discrete. We can simplify analysis by taking a larger domain $\tilde{\mathcal{D}}_z = \{z \in \mathbb{Z}_+\}$ of all non-negative integers and changing the log factorial to the smooth gamma function, i.e. $\sum_s \ln((zv_s)!) \rightarrow \sum_s \ln(\Gamma(zv_s+1))$. Thus, the radial conditional log partition function is upper bounded by:

$$\sum_{z \in \mathbb{Z}_+} \exp \left(\underbrace{\bar{\eta}_1 z + \bar{\eta}_2 \sqrt{z}}_{O(z)} - \underbrace{\sum_s \ln(\Gamma(zv_s+1))}_{O(z \ln z)} \right) < \infty. \quad (14)$$

The basic intuition is that the exponent has a linear $O(z)$ term minus an $O(z \ln z)$ term, which will eventually overcome the linear term and hence the summation will converge. Note that we did not assume any restrictions on Φ except that all the entries are finite. Thus, for the

⁵On the edge case when $\bar{\eta}_1 = 0$, the log partition function will diverge if $\bar{\eta}_2 \geq 0$ and will converge if $\bar{\eta}_2 < 0$ by simple arguments. The normalizability condition when $\eta_2 = 0$ could slightly loosen the condition on Φ in Eqn. 13 but for simplicity we did not include this edge case.

Poisson distribution, Φ can have arbitrary positive and negative dependencies. A formal proof for Eqn. 14 is given in the appendix.

5. Experiments and Results

5.1. Synthetic Experiment

In order to show that our parameter estimation algorithm has the ability to find the correct dependencies, we develop a synthetic experiment on chain-like graphs. We construct Φ to be a k -dependent circular chain-like graph by first setting the diagonal of Φ to be 1. Then, we

add an edge between each node and its k neighbors with a value of $\frac{0.9}{2 * k}$, i.e. the s -th node is connected to the $(s + 1)$ -th, $(s + 2)$ -th, ..., $(s + k)$ -th nodes where the indices are modulo p (e.g. $k = 1$ is the standard chain graph). This ensures that Φ is negative definite by the Gershgorin disc theorem. We generate samples using Gibbs sampling with 1000 Gibbs iterations per sample and 10 slice samples for each node conditional sample. For this experiment, we set $p = 30$, $\lambda = 10^{-5}$, $k \in \{1, 2, 3, 4\}$, and $n \in \{100, 200, 400, 800, 1600\}$. We calculate the edge precision for the fitted model by computing the precision for the top kp edges—i.e. the number of true edges in the top kp estimated edges over the total number of true edges. The results in Fig. 4 demonstrate that our parameter estimation algorithm is able to easily find the edges for small k and is even able to identify the edges for large k , though the problem becomes more difficult when k is large (because there are more parameters, which are also smaller), and thus more samples are needed. With 1,600 samples, our parameter estimation algorithm is able to recover at least 95% of the edges even when $k = 4$.

5.2. Airport Delay Times Experiment

In order to demonstrate that the SQR graphical model class is more suitable for real-world data than the graphical models in (Yang et al., 2015) (which can only model *negative* dependencies), we fit an exponential SQR model to a dataset of airport delay times at the top 30 commercial USA airports—also known as Large Hub airports. We gathered flight data from the US Department of Transportation public “On-Time: On-Time Performance” database⁶ for the year 2014. We calculated the average delay time per day at each of the top 30 airports (excluding cancellations).

For our implementation, we set $\lambda \in \{0.05, 0.005, 0.0005\}$ and set a maximum of 5000 iterations for our proximal gradient descent algorithm. For approximating the log partition function using the AIS sampling defined in Sec. 3.4, we sampled 1000 AIS samples with 100 annealing distributions—i.e. γ took 100 values between 0 and 1—, 10 Gibbs steps per annealed distribution and 10 slice samples for every node conditional sampling. Generally, our algorithm with these parameter settings took roughly 35 seconds to train the model and about 25 seconds to compute the likelihood (i.e. AIS sampling) using MATLAB prototype code on the TACC Maverick cluster.⁷

⁶http://www.transtats.bts.gov/DL_SelectFields.asp?Table_ID=236&DB_Short_Name=On-Time

⁷<https://portal.tacc.utexas.edu/user-guides/maverick>

We computed the geometric mean of the relative log likelihood compared to the independent exponential model, i.e. $\exp((\mathcal{L}_{\text{SQR}} - \mathcal{L}_{\text{Ind}})/n)$, where \mathcal{L} is the log likelihood. These values can be seen in Fig. 4 (higher is better). Clearly, the exponential SQR model provides a major improvement in relative likelihood over the independent model suggesting that the delay times of airports are clearly related to one another. In Fig. 5, we visualize the non-zeros of Φ —which correspond to the edges in the graphical model—to show that our model is capturing intuitive positive dependencies.

First, it should be noted that all the dependencies are positive yet positive dependencies were not allowed by previous graphical models (Yang et al., 2015)! Second, as would be expected because of weather delays, the airports in the Chicago area seems to affect the delays of many other airports. Similarly, a weather effect seems to be evident for the airports near New York City. Third, as would be expected, some dependencies seem to be geographic in nature as seen by the west coast dependencies, Texas dependency (i.e. DFW-IAH), and east coast dependencies. Note that the geographic dependencies were found even though no location data was given to the algorithm. Fourth, the busiest airport in Atlanta, GA (ATL) is not strongly dependent on other airports. This seems reasonable because Atlanta rarely has snow and there are few major airports geographically close to Atlanta. These qualitative results suggest that the exponential SQR model is able to capture multiple interesting and intuitive dependencies.

6. Discussion

As full probability models, SQR graphical models could be used in any situation where a multivariate distribution is required. For example, SQR models could be used in Bayesian classification by modeling the probability of each class distribution instead of the classical Naive Bayes assumption of independence. As another example, SQR models could be used as the base distribution in mixtures or admixture composite distributions as in (Inouye et al., 2014; Inouye et al.)—similar to multivariate Gaussian mixture models. Another extension would be to consider mixed SQR graphical models in which the joint distribution has variables using different exponential families as base distributions as explored for previous graphical models in (Yang et al., 2014; Tansey et al., 2015).

7. Conclusion

We introduce a novel class of graphical models that creates multivariate generalizations for *any* univariate exponential family with nonnegative sufficient statistics—including Gaussian, discrete, exponential and Poisson distributions. We show that SQR graphical models generally have few restrictions on the parameters and thus can model *both positive and negative* dependencies unlike previous generalized graphical models as represented by (Yang et al., 2015). In particular, for the exponential SQR model, the parameter matrix Φ can have both positive and negative dependencies and is only constrained by a mild condition—akin to the positive-definiteness condition on Gaussian covariance matrices. For the Poisson distribution, there are no restrictions on the parameter values, and thus the Poisson SQR model allows for arbitrary positive and negative dependencies. We develop parameter estimation and likelihood approximation methods and demonstrate that the SQR model

indeed captures interesting and intuitive dependencies by modeling both synthetic datasets and a real-world dataset of airport delays. The general SQR class of distributions opens the way for graphical models to be effectively used with non-Gaussian and non-discrete data without the unintuitive restriction to negative dependencies.

Acknowledgments

This work was supported by NSF (DGE-1110007, IIS-1149803, IIS-1447574, IIS-1546459, DMS-1264033, CCF-1320746) and ARO (W911NF-12-1-0390).

References

- Allen GI, Liu Z. A local Poisson graphical model for inferring networks from sequencing data. *IEEE Trans on Nanobioscience*. 2013; 12(3):189–198.
- Banerjee O, El Ghaoui L, D'Aspremont A. Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data. *JMLR*. 2008; 9:485–516.
- Besag J. Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society. Series B (Methodological)*. 1974; 36(2):192–236.
- Bickel, P.; Diggle, P.; Fienberg, S.; Gather, U.; Olkin, I. *Copula Theory and Its Applications*. Springer; 2009.
- Embrechts, P.; Lindskog, F.; Mcneil, A. *Handbook of Heavy Tailed Distributions in Finance*. Elsevier; 2003. Modelling dependence with copulas and applications to risk management; p. 329-384.
- Friedman J, Hastie T, Tibshirani R. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*. 2008; 9(3):432–441. [PubMed: 18079126]
- Genest C, Neslehova J. A primer on copulas for count data. *ASTIN Bulletin*. 2007; 37(2):475–515.
- Hsieh CJ, Sustik MA, Dhillon IS, Ravikumar P. QUIC: Quadratic approximation for sparse inverse covariance estimation. *JMLR*. 2014; 15:2911–2947.
- Inouye DI, Ravikumar P, Dhillon IS. Capturing semantically meaningful word dependencies with an admixture of Poisson MRFs. *NIPS*. :3158–3166.
- Inouye DI, Ravikumar P, Dhillon IS. Admixture of Poisson MRFs: A topic model with word dependencies. *ICML*. 2014
- Inouye DI, Ravikumar P, Dhillon IS. Fixed-length Poisson MRF: Adding dependencies to the multinomial. *NIPS*. 2015:3195–3203.
- Jalali A, Ravikumar P, Vasuki V, Sanghavi S. On learning discrete graphical models using group-sparse regularization. *AISTATS*. 2010:378–387.
- Karlis D. An EM algorithm for multivariate Poisson distribution and related models. *Journal of Applied Statistics*. 2003; 30(1):63–77.
- Liu H, Lafferty J, Wasserman L. The nonpara-normal: Semiparametric estimation of high dimensional undirected graphs. *JMLR*. 2009; 10:2295–2328.
- Marshall AW, Olkin I. A multivariate exponential distribution. *Journal of the American Statistical Association*. 1967; 62(317):30–44.
- Neal RM. Annealed importance sampling. *Statistics and Computing*. 2001; 11(2):125–139.
- Ravikumar P, Wainwright M, Lafferty J. High-dimensional Ising model selection using L1-regularized logistic regression. *The Annals of Statistics*. 2010; 38(3):1287–1319.
- Tansey W, Padilla OHM, Suggala AS, Ravikumar P. Vector-space Markov random fields via exponential families. *ICML*. 2015
- Wainwright MJ, Jordan MI. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*. 2008; 1(12):1–305.
- Yang E, Ravikumar P, Allen GI, Liu Z. On Poisson graphical models. *NIPS*. 2013:1718–1726.
- Yang E, Baker Y, Ravikumar P, Allen GI, Liu Z. Mixed graphical models via exponential families. *AISTATS*. 2014; (2012):1042–1050.

Yang E, Ravikumar P, Allen GI, Liu Z. On graphical models via univariate exponential family distributions. *JMLR*. 2015; 16:3813–3847.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript




Figure 1.

These examples of 2D exponential SQR and Poisson SQR distributions with no dependency (i.e. independent), positive dependency and negative dependency show the amazing flexibility of the SQR model class that can intuitively model *positive and negative* dependencies while having a simple parametric form. The approximate 1D marginals are shown along the edges of the plots.

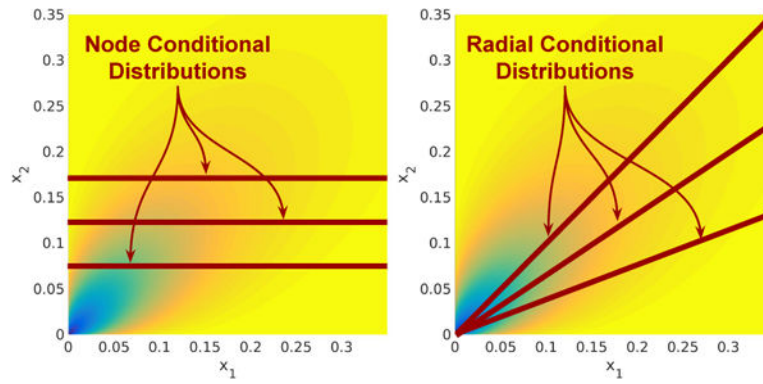


Figure 2. *Node* conditional distributions (left) are univariate probability distributions of one variable assuming the other variables are given while *radial* conditional distributions are univariate probability distributions of vector scaling assuming the vector direction is given. Both conditional distributions are helpful in understanding SQR graphical models.

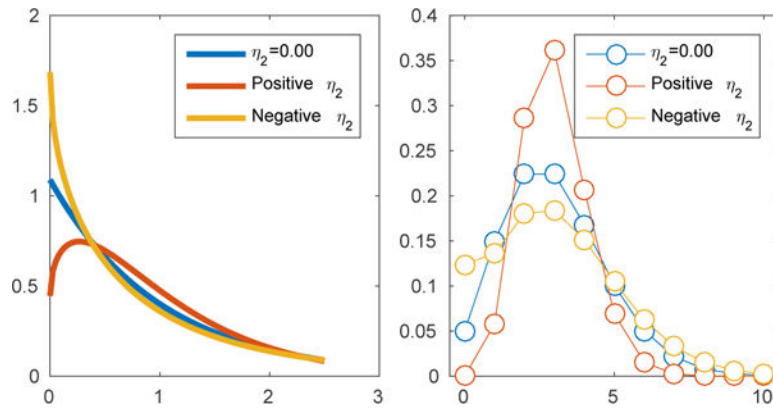


Figure 3. Examples of the node conditional distributions of exponential (left) and Poisson (right) SQR models for $\eta_2 = 0$, $\eta_2 > 0$ and $\eta_2 < 0$.




Figure 4. (Left) The fitted exponential SQR model improves significantly over the independent exponential model in terms of relative likelihood suggesting that a model with positive dependencies is more appropriate. (Right) The edge precision for the circular chain graph described in Sec. 5.1 demonstrate that our parameter estimation algorithm is able to effectively identify edges for small k , and if given enough samples, can also identify edges for larger k .

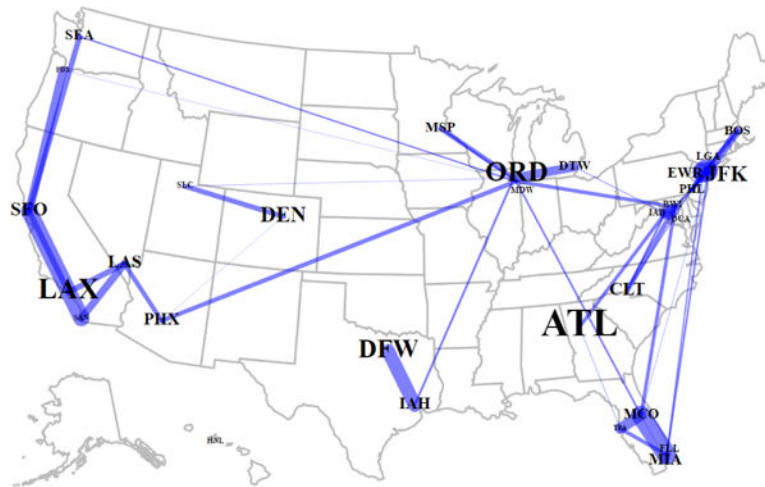


Figure 5. Visualizing the top 50 edges between airports shows that SQR models can capture interesting and intuitive *positive* dependencies even though previous exponential graphical models (Yang et al., 2015) were restricted to negative dependencies. The delays at the Chicago airports seem to affect other airports as would be expected because of Chicago weather delays. Other dependencies are likely related to weather or geography. (For this visualization, we set $\lambda = 0.0005$. Width of lines is proportional to the value of the edge weight, i.e. a non-zero in Φ , and the size of airport abbreviation is proportional to the average number of passengers.)