

Estimating Densities with Non-Parametric Exponential Families

Technical Report 12-02
Department of Statistics
Purdue University
West Lafayette, IN, USA 47907-2066

Lin Yuan*, Sergey Kirshner[†], Robert Givan*
{yuanl, skirshne, givan}@purdue.edu

* School of Electrical and Computer Engineering
[†] Department of Statistics
Purdue University, West Lafayette, IN 47907, USA

June 20, 2012

Abstract

We propose a novel approach for density estimation with exponential families for the case when the true density may not fall within the chosen family. Our approach augments the sufficient statistics with features designed to accumulate probability mass in the neighborhood of the observed points, resulting in a non-parametric model similar to kernel density estimators. We show that under mild conditions, the resulting model uses only the sufficient statistics if the density is within the chosen exponential family, and asymptotically, it approximates densities outside of the chosen exponential family. Using the proposed approach, we modify the exponential random graph model, commonly used for modeling small-size graph distributions, to address the well-known issue of model degeneracy.

1 Introduction

The problem of density estimation is ubiquitous in machine learning and statistics. A typical approach would assume a parametric family for the distribution from which the observed data is drawn and estimate the parameters by fitting them to the data. Among the parametric families, *exponential families* play a prominent role, as maximum likelihood estimation from complete data for exponential families is asymptotically unbiased, consistent, and efficient (van der Vaart, 1998, Chapter 5). Finding the maximum likelihood estimator (MLE) reduces to a convex optimization problem that requires knowing only the sufficient statistics from the data (Barndorff-Nielsen, 1978, Brown, 1986). When the functional form is not readily available, *non-parametric approaches* (e.g., kernel density estimation) provide a convenient alternative by allowing the number of parameters to grow with the available data. However, if the number of components in the data vectors is large relative to the number of the available data points, non-parametric approaches may suffer from the curse of dimensionality (Bellman, 1957) and overfit. From the standpoint of a bias-variance tradeoff, a parametric approach can be useful even if the number of observations is not large (as statistics of the data can often be estimated from relatively few samples) but could also be hindered by a misspecification bias if the true distribution falls outside the chosen parametric family. Conversely, non-parametric approaches can approximate any density with enough data, suffering a high variance when the data is limited or the dimensionality is even moderately large.

We propose a novel non-parametric density-estimation approach for exponential families that combines some of the strengths of parametric and non-parametric approaches. Our approach draws inspiration from kernel density estimators (KDEs), which approximate unknown densities by placing probability mass around the observations, and from the exponential families by imposing global constraints in matching the statistics. Exponential families are derived by maximizing the Shannon’s entropy of the estimated distribution subject to the constraint that the expected values of the chosen statistics (features) with respect to the empirical and the estimated distributions must match. Our proposed exponential family model imposes additional constraints requiring a small constant probability mass around each example point. We accomplish this by augmenting the set of given statistics (features) with kernel functions centered around the observations, so that the expected value for each of these functions represents the probability mass concentrated around an example point. The resulting exponential family model is non-parametric, as each data point has a parameter associated with it. The objective function for the parameter estimation is convex and contains an ℓ_1 -penalty term for each added parameter. These penalties encourage sparsity by potentially making many of the added parameters vanish. We show that if the true distribution is within the exponential family model with the chosen statistics, then as the sample size increases, all parameters associated with the added local features vanish and our approach converges to the true distribution. If the true distribution is *not* from the chosen exponential family, then, our approach provides a close approximation to the unknown density, comparable to KDEs.

Our work is in part motivated by a problem of learning distributions over graphs from examples of observed networks, typically from a *single* network. Such data arises in many domains, including social sciences, bioinformatics, and systems sciences. Among the approaches to this problem, one of the perhaps most-studied is the exponential random graph model (ERGM, or in the social network literature, p^* , e.g., Frank and Strauss, 1986, Holland and Leinhardt, 1981, Wasserman and Pattison, 1996). ERGMs use graph statistics as features to define an exponential family distribution over all possible graphs with a given number of nodes. Such models have the desirable property of learning a distribution that matches the observed graph statistics. the expected value for the chosen statistics in the observed population. However, ERGMs often suffer from issues of *degeneracy* (Handcock, 2003, Lunga and Kirshner, 2011, Rinaldo et al., 2009) manifested in placing most of the probability mass on unrealistic graphs (e.g., an empty or a complete graph), very dissimilar to the observed graph(s). As an illustration of our approach, we propose a modification to ERGMs which alleviates the above issue of degeneracy in moderate-sized graphs.

The main contributions of this paper are a novel framework for non-parametric estimation of densities with exponential family models that is applicable when the number of data points is relatively small, analysis of its convergence properties, and a modification of ERGMs that remedies one of the degeneracy issues. The paper is structured as follows. In Section 2, we briefly describe the exponential family models. In Section 3 we introduce the features we use to constrain the probability mass around the data points and derive a formulation for a new model from first principles. In Section 3.3, we derive some of the new model's properties, and then discuss the resulting parameter-estimation optimization problem and our approach to solving it in Section 3.4. In Section 4, we propose a new model for distributions over networks with a moderate number of nodes. We explore the properties of our estimator for 1-dimensional densities and for modeling network data via an empirical study in Section 5, and finally discuss our findings and outline possible future directions in Section 6.

2 Exponential Family

We briefly introduce the exponential family of distributions before describing our contribution, a non-parametric exponential family.

Suppose \mathbf{X} is a vector of random variables with support $\mathcal{X} \subseteq \mathbb{R}^m$. A distribution for \mathbf{X} belongs to the exponential family of distributions with sufficient statistics $\mathbf{t} : \mathcal{X} \rightarrow \mathcal{H} \subseteq \mathbb{R}^d$, if its probability density has a functional form:¹

$$f^E(\mathbf{x}|\boldsymbol{\lambda}) = \frac{1}{Z(\boldsymbol{\lambda})} q(\mathbf{x}) \exp \langle \boldsymbol{\lambda}, \mathbf{t}(\mathbf{x}) \rangle \quad \text{where} \quad (1)$$

$$Z(\boldsymbol{\lambda}) = \int_{\mathcal{X}} q(\mathbf{x}) \exp \langle \boldsymbol{\lambda}, \mathbf{t}(\mathbf{x}) \rangle d\mathbf{x} < \infty$$

is a partition function, $\boldsymbol{\lambda}$ is a vector of canonical parameters, $q : \mathcal{X} \rightarrow \mathbb{R}$ is a base measure, and $\langle \cdot, \cdot \rangle$ denotes the Euclidean inner product. We further assume that the exponential family is regular (i.e. the canonical parameter space is open). Assuming q is fixed, let $\mathcal{EF}_{\mathbf{t}}$ denote the set of all possible distributions of the form (1) with the set of sufficient statistics \mathbf{t} .

Given samples $\mathbf{x}^{1:n} \triangleq (\mathbf{x}^1, \dots, \mathbf{x}^n) \stackrel{i.i.d}{\sim} f$ where $f : \mathcal{X} \rightarrow \mathbb{R}$ is an unknown density with the same support as q . Let $\hat{f}_n : \mathcal{X} \rightarrow \mathbb{R}$ be the empirical distribution for $\mathbf{x}^{1:n}$, $\hat{f}_n(\mathbf{x}|\mathbf{x}^{1:n}) = \frac{1}{n} \sum_{i=1}^n \delta(\mathbf{x}^i)$ where $\delta(\mathbf{x})$ is a Dirac delta function. Exponential families can be obtained as a solution to the optimization problem of minimizing the relative entropy subject to matching the moment constraints of the empirical and the estimated distributions:

$$f_n^E(\mathbf{x}) = \arg \min_{f^E \in \mathcal{F}} KL(f^E \| q) \quad \text{subj to} \quad (2)$$

$$E_{f_n^E(\mathbf{x})}[\mathbf{t}(\mathbf{x})] = E_{\hat{f}_n(\mathbf{x}|\mathbf{x}^{1:n})}[\mathbf{t}(\mathbf{x})]. \quad (3)$$

A distribution $f_n^E(\mathbf{x}|\hat{\boldsymbol{\lambda}}_n) \in \mathcal{EF}_{\mathbf{t}}$ satisfying (3) can be found by maximizing the log-likelihood $l(\boldsymbol{\lambda}) = \langle \boldsymbol{\lambda}, \frac{1}{n} \sum_{i=1}^n \mathbf{t}(\mathbf{x}^i) \rangle - \log Z(\boldsymbol{\lambda})$, and provided $\mathbf{t}^* = E_{\hat{f}_n(\mathbf{x}|\mathbf{x}^{1:n})}[\mathbf{t}(\mathbf{x})] \in \text{rint}(\text{conv}(\mathcal{H}))$, a maximum likelihood estimate (MLE) $\hat{\boldsymbol{\lambda}}_n$ satisfying (3) exists (Wainwright and Jordan, 2008), and will be unique if $\mathcal{EF}_{\mathbf{t}}$ is minimal (Brown, 1986). If $f \in \mathcal{EF}_{\mathbf{t}}$, then $\hat{\boldsymbol{\lambda}}_n \xrightarrow{P} \boldsymbol{\lambda}$ (van der Vaart, 1998).

However, if the true distribution does not fall within the chosen exponential family, $f \notin \mathcal{EF}_{\mathbf{t}}$, the estimated model may provide a poor approximation to the true density. As will be illustrated in Section 4, for the case of discrete random vectors \mathbf{X} from the exponential family with a bounded support \mathcal{H} , finding

¹For notational convenience, we denote $\mathbf{X} = \mathbf{x}$ by \mathbf{x} .

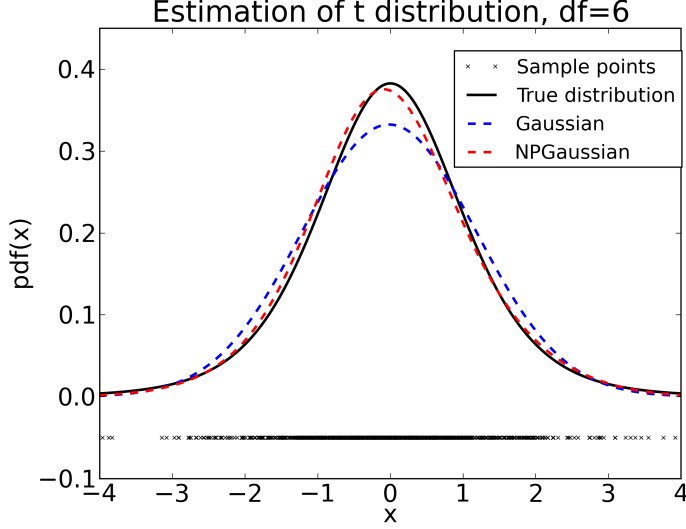


Figure 1: Density estimation from samples from a t -distribution. Black x's are samples; the black solid line is the true density, the blue dashed line is the fitted Gaussian density, and the red dashed line is the fitted non-parametric Gaussian density with non-parametric exponential family model with Gaussian kernel with width 1.5.

the MLE under the wrong modeling assumption $f \in \mathcal{EF}_t$ may assign very little probability mass to the observed samples $\mathbf{x}^{1:n}$.

3 Non-parametric Exponential Family

In this section, we propose a new family of distributions, a modification to the exponential family \mathcal{EF}_t . Our proposed approach modifies the set of features so that the estimated density (or a probability mass function for discrete vectors) places approximately the same amount of mass around each sample \mathbf{x}^i , $i = 1, \dots, n$ as the empirical distribution. This approach allows using exponential family models to approximate distributions outside of the exponential family (e.g., mixtures, heavy-tailed distributions). This approach can also be used to avoid degeneracy in cases where the set of features is poorly chosen (e.g., modeling of graphs with ERGMs).

3.1 Motivation

Suppose a set of samples from an unknown density “looks” Gaussian except perhaps for a few outliers in the tails (Figure 1). Should we fit a Gaussian? If not, should we use a non-parametric approach? Our approach combines both by using the exponential family with given features (e.g. $\mathbf{t}(\mathbf{x}) = (x, x^2)$ in the case of a univariate Gaussian) as a starting point and then adding features for each data point. It draws inspiration from KDEs (also known as Parzen windows, Parzen, 1962),

$$f_n^{\text{KDE}}(\mathbf{x}|\mathbf{x}^{1:n}) = \frac{1}{n} \sum_{i=1}^n K_{\mathbf{H}}(\mathbf{x}; \mathbf{x}^i) \text{ where}$$

$$K_{\mathbf{H}}(\mathbf{x}; \mathbf{x}^i) = |\mathbf{H}|^{-\frac{1}{2}} K\left(\mathbf{H}^{-\frac{1}{2}}(\mathbf{x} - \mathbf{x}^i)\right).$$

K is a univariate kernel function, a bounded probability density function on \mathbb{R} . $K_{\mathbf{H}}$ is a multivariate kernel function with a symmetric positive definite bandwidth matrix \mathbf{H} ; in this paper, we assume $\mathbf{H} = h^2 \mathbf{I}_d$ (assuming $\mathbf{x} \in \mathbb{R}^d$).

The uniform kernel is an indicator function on $[-\frac{1}{2}, \frac{1}{2}]$:

$$K_U(x) = 1 \text{ if } |x| \leq \frac{1}{2}, \text{ and } 0 \text{ otherwise,}$$

where a multi-dimensional version is a weighted indicator function for the hypercube centered at \mathbf{x}^i with each side equal to $\frac{h}{2}$. Most other kernels used with KDE are smooth approximations of K_U , e.g., Gaussian kernel $K_{\mathcal{N}}(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}$. KDE matches the mass around each data point (weighted according to the kernel) to that of the empirical distribution. Since the empirical distribution approaches the true distribution as n increases, the accuracy of KDE approximation improves with the increase in the number of data points and the decrease of bandwidth parameter h . The resulting representation however requires keeping all of the observations as parameters and requires exponentially many data points in the dimension d to approximate the underlying density well.

3.2 Formulation

Our approach preserves the mass around each data point by introducing additional moment constraints. Let $\mathcal{B} \subseteq \mathcal{X}$ be a region in the support of \mathbf{X} , and let $\mathcal{I}_{\mathcal{B}}(\mathbf{x}) = 1$ if $\mathbf{x} \in \mathcal{B}$, and $\mathcal{I}_{\mathcal{B}}(\mathbf{x}) = 0$ otherwise denote an indicator function for \mathcal{B} . Given metric space (\mathcal{X}, σ) , let $\mathcal{B}_i = \{\mathbf{x} \in \mathcal{X} : \sigma(\mathbf{x}^i, \mathbf{x}) \leq \varepsilon\}$ be an ε -neighborhood of \mathbf{x}_i . Then the probability mass for density f in the ε -neighborhood \mathcal{B}_i of \mathbf{x}_i is $P(\mathcal{B}_i) = E_f[\mathcal{I}_{\mathcal{B}_i}]$. We propose adding constraints to (3) which would approximately match the probability masses for \mathcal{B}_i ($i = 1, \dots, n$) between the empirical and the estimated distributions (\hat{f}_n and f_n respectively):

$$\left| E_{\hat{f}_n(\mathbf{x}|\mathbf{x}^{1:n})}[\mathcal{I}_{\mathcal{B}_i}(\mathbf{x})] - E_{f_n(\mathbf{x}|\mathbf{x}^{1:n})}[\mathcal{I}_{\mathcal{B}_i}(\mathbf{x})] \right| \leq \beta_i, \quad (4)$$

where $\beta_i \geq 0$ determine how closely the masses should match. Similar to KDEs, $\mathcal{I}_{\mathcal{B}_i}$ in (4) may be replaced with a multidimensional kernel $K_{\mathbf{H}}(\mathbf{x}^i; \mathbf{x})$, which assigns decaying importance of mass away from the center (e.g., a smoothed version of $\mathcal{I}_{\mathcal{B}_i}$). We will use $t_a^i \triangleq K_{\mathbf{H}}(\mathbf{x}^i; \mathbf{x})$, and use $\mathbf{t}_a(\mathbf{x}) \triangleq [t_a^1(\mathbf{x}), \dots, t_a^n(\mathbf{x})]$ to augment the statistics \mathbf{t} in estimating densities.² In addition to the canonical parameters $\boldsymbol{\lambda}$ for sufficient statistics, we add *augmented* parameters $\boldsymbol{\lambda}_a$ for the augmented statistics $\mathbf{t}_a(\mathbf{x})$.

Our proposed density approximation ($f_n^{NE}(\mathbf{x})$) is a solution to

$$\begin{aligned} f_n^{NE}(\mathbf{x}|\mathbf{x}^{1:n}) &= \arg \min_{f^{NE} \in \mathcal{F}} KL(f^{NE} \parallel q) \text{ subj to} \\ E_{f_n^{NE}(\mathbf{x}|\mathbf{x}^{1:n})}[\mathbf{t}(\mathbf{x})] &= E_{\hat{f}_n(\mathbf{x}|\mathbf{x}^{1:n})}[\mathbf{t}(\mathbf{x})], \\ \left| E_{f_n^{NE}}[t_a^i(\mathbf{x})] - E_{\hat{f}_n}[t_a^i(\mathbf{x})] \right| &\leq \beta_i, \quad i = 1, \dots, n. \end{aligned} \quad (5)$$

f_n^{NE} falls within the generalized MaxEnt framework (Dudik et al., 2007):

$$\begin{aligned} f(\mathbf{x}) &= \frac{1}{Z(\boldsymbol{\lambda}, \boldsymbol{\lambda}_a)} q(\mathbf{x}) \exp[\langle \boldsymbol{\lambda}, \mathbf{t}(\mathbf{x}) \rangle + \langle \boldsymbol{\lambda}_a, \mathbf{t}_a(\mathbf{x}) \rangle] \\ Z(\boldsymbol{\lambda}, \boldsymbol{\lambda}_a) &= \int_{\mathcal{X}} q(\mathbf{x}) \exp[\langle \boldsymbol{\lambda}, \mathbf{t}(\mathbf{x}) \rangle + \langle \boldsymbol{\lambda}_a, \mathbf{t}_a(\mathbf{x}) \rangle] d\mathbf{x}. \end{aligned} \quad (6)$$

²We omit h from t_a^i for the simplicity of notation. It is a tuning parameter that may be set globally for all $i = 1, \dots, n$.

Let $\mathbf{s}(\mathbf{x}) \triangleq (\mathbf{t}(\mathbf{x}), t_a(\mathbf{x}))$ and $\boldsymbol{\theta} \triangleq (\boldsymbol{\lambda}, \boldsymbol{\lambda}_a)$ be a combined set of statistics and parameters, respectively, for the augmented model. A specific set of parameter values for the distribution in (6) satisfying the constraints in (5) can be found by maximizing the penalized log-likelihood

$$l(\boldsymbol{\theta}) = \left\langle \boldsymbol{\theta}, \frac{1}{n} \sum_{i=1}^n \mathbf{s}(\mathbf{x}^i) \right\rangle - \ln Z(\boldsymbol{\theta}) - \sum_{i=1}^n \beta_i |\lambda_a^i|. \quad (7)$$

Note that the distribution f_n^{NE} satisfying the constraint in (5) will always exist since \hat{f}_n satisfies all of the above constraints.

We refer to the above class of models as *non-parametric exponential family models*, since the number of non-zero parameters $\boldsymbol{\theta} = (\boldsymbol{\lambda}, \boldsymbol{\lambda}_a)$ may increase with the number of available data points. We will denote this family by \mathcal{NEF}_s . Clearly $\mathcal{EF}_t \subseteq \mathcal{NEF}_s$ as all the augmented parameters can be set to 0. Let $\hat{\boldsymbol{\theta}}_n = (\tilde{\boldsymbol{\lambda}}_n, \tilde{\boldsymbol{\lambda}}_{a,n})$ be the MLE of $(\boldsymbol{\lambda}, \boldsymbol{\lambda}_a)$ for the case of n samples. The ℓ_1 -penalty in (7) is known to be sparsity-inducing (e.g., Bach et al., 2011), so in practice, many of $\tilde{\lambda}_{a,n}^i = 0$.

Note that the framework in (5) allows matching the moments of the estimated distribution to some predetermined vector $\mathbf{t}^* \in \text{rint}(\text{conv}(\mathcal{H}))$ instead of the empirical moments $E_{\hat{f}_n(\mathbf{x}|\mathbf{x}^{1:n})}[\mathbf{t}(\mathbf{x})]$, leading to the same functional form for the density (6), but with an additional linear term $\langle \boldsymbol{\lambda}, \mathbf{t}^* - \frac{1}{n} \sum_{i=1}^n \mathbf{t}(\mathbf{x}^i) \rangle$ in (7) (Dudik et al., 2007). Thus, our non-parametric density estimator is capable of satisfying global constraints (matching a set of provided moments, e.g., learning a distribution with a given covariance). We are not aware of other non-parametric methods with this capability.

3.3 Theoretical Properties

The proofs appear in the Appendix A.

Theorem 3.1. *Suppose a vector of random variables \mathbf{X} with support on \mathcal{X} has a density $f \in \mathcal{EF}_t$ with features $\mathbf{t} : \mathcal{X} \rightarrow \mathcal{H} \subseteq \mathbb{R}^d$ and a vector of canonical parameters $\boldsymbol{\lambda} \in \mathcal{C} \subseteq \mathbb{R}^d$. Suppose $\mathbf{x}^1, \dots, \mathbf{x}^n \triangleq \mathbf{x}^{1:n}$ is a sequence of i.i.d. random vectors drawn from f . Let $f_n^{NE}(\mathbf{x}|\hat{\boldsymbol{\theta}}_n, \mathbf{x}^{1:n}) \in \mathcal{NEF}_s$ be the MLE solution*

of (5), $\hat{\boldsymbol{\theta}}_n = (\tilde{\boldsymbol{\lambda}}_n, \tilde{\boldsymbol{\lambda}}_{a,n})$, with all $\beta_i = \beta > 0$, $i = 1, \dots, n$. Assuming

1. \mathcal{X} is compact,
2. \mathbf{t} is continuous,
3. \mathcal{EF}_t is a family of uniformly equicontinuous functions w.r.t \mathbf{x} ,
4. Kernel K has bounded variation and has a bandwidth parameter \mathbf{H} such that the series $\sum_{n=1}^{\infty} e^{-\gamma n|\mathbf{H}|}$ converges for every positive value of γ ,

then as $n \rightarrow \infty$, $\tilde{\lambda}_{a,n}^i \xrightarrow{p} 0, \forall i = 1, \dots, n$ and $\tilde{\boldsymbol{\lambda}}_n \xrightarrow{p} \boldsymbol{\lambda}$.

Intuitively, uniform convergence is required because we need the n additional constraints be satisfied with fixed threshold β as long as $n > N$. We use Assumption 2 to relate the pointwise convergence of the MLE $\boldsymbol{\lambda}$ in regular exponential families to pointwise convergence of $f_n^E(\mathbf{x}|\hat{\boldsymbol{\lambda}}_n)$. Assumptions 1 and 3 are further employed to convert the pointwise convergence to uniform convergence of f_n^E by considering \mathcal{X}, \mathcal{C} to be subsets of the original regular exponential family.⁴ Assumption 4 is the requirement used by Nadaraya (1965) for the uniform convergence of KDE satisfied by common kernels. Upon these uniform convergence results, the augmented constraints will be satisfied by the original MLE estimate $[\hat{\boldsymbol{\lambda}}_n, \mathbf{0}]$. Further, the nature

³ $f_n^{NE}(\mathbf{x}|\boldsymbol{\theta}, \mathbf{x}^{1:n})$'s functional form depends on both $\mathbf{x}^{1:n}$ and $\boldsymbol{\theta}$, for convenience, we sometimes omit $\boldsymbol{\theta}$ and $\mathbf{x}^{1:n}$ in notation when referring to f_n^{NE} .

⁴For example, Theorem 3.1 applies to the exponential family $\mathcal{N}(\mu, \sigma^2)$ with $\sigma \in [a, b], \forall b > a > 0$, but not if $\sigma^2 \in (0, \infty)$.

of $\hat{\lambda}_n$ being a maximum entropy solution guarantees that $[\hat{\lambda}_n, \mathbf{0}]$ is a maximum entropy solution under the additional constraints.

Theorem 3.1 shows that if the true distribution falls within the exponential family, then as sample size increases, the estimated density from the non-parametric exponential family will have vanishing reliance on the augmented parameters.

Theorem 3.2. *Given a probability density function $f(\mathbf{x}) : \mathcal{X} \rightarrow \mathbb{R}$, let $f_n^{NE}(\mathbf{x}|\hat{\theta}_n, \mathbf{x}^{1:n}) \in \mathcal{NEF}_s$ be a solution satisfying (5). If*

1. f is uniformly continuous on \mathcal{X} ,
2. $K_{\mathbf{H}}(\mathbf{x})$ is uniformly continuous on \mathcal{X} ,
3. $\sup_{\mathbf{x} \in \mathcal{X}} K_{\mathbf{H}}(\mathbf{x}) < \infty$,
4. $\lim_{\|\mathbf{x}\| \rightarrow \infty} K_{\mathbf{H}}(\mathbf{x}) \prod_{i=1}^m \mathbf{x}_i = 0$,
5. $\lim_{n \rightarrow \infty} |\mathbf{H}|^{\frac{1}{2}} = 0$,
6. $\lim_{n \rightarrow \infty} n |\mathbf{H}|^{\frac{1}{2}} = \infty$,

then $f_n^{NE}(\mathbf{x}|\hat{\theta}_n, \mathbf{x}^{1:n}) \xrightarrow{p} f(\mathbf{x})$ pointwise on \mathcal{X} .

Assumptions 3-6 are required for pointwise convergence of KDE (Parzen, 1962) at specific points $\mathbf{x}^{1:n}$. We then extend this pointwise convergence from $\mathbf{x}^{1:n}$ to \mathcal{X} , considering the probability of sampling a new $\mathbf{x} \in \mathcal{X}$ far away from existing $\mathbf{x}^{1:n}$ under the true density f . The monotone convergence theorem and the uniform continuity assumptions (1 and 2) lead to the pointwise convergence $f_n^{NE} \xrightarrow{p} f$ on \mathcal{X} .

Theorem 3.2 indicates the weak consistency of the non-parametric exponential family density estimator. Thus our proposed non-parametric approach can be used to approximate densities which are not from exponential families.

3.4 Estimating Parameters for Non-Parametric Exponential Families

Recently there have been a number of methods developed for optimization of convex non-smooth functions, some of them specifically aimed at log-linear problems such as (7) (e.g., Bach et al., 2011, Shalev-Shwartz and Tewari, 2011, Wu and Lange, 2008). We employed a coordinate descent algorithm similar to the SUM-MET algorithm of Dudik et al. (2007) (see Algorithm 1), primarily, due to its simplicity. Other possible approaches can be employed as well and may end up more efficient for this formulation.

The proposed algorithm iterates between optimizing canonical parameters λ (by setting $E_{\hat{f}_n}[t(\mathbf{x})] = E_{f_n^{NE}(\mathbf{x}|\theta^{(k)})}[t(\mathbf{x})]$) and sequentially optimizing the augmented parameters λ_a so that the Karush-Kuhn-Tucker conditions (e.g., Nocedal and Wright, 2006) are satisfied:

$$E_{\hat{f}_n}[t_a^i(\mathbf{x})] - E_{f_n^{NE}(\mathbf{x}|\theta)}[t_a^i(\mathbf{x})] \in \begin{cases} \{\beta_i\} & \lambda_a^i > 0, \\ \{-\beta_i\} & \lambda_a^i < 0, \\ (-\beta_i, \beta_i) & \beta_a^i = 0. \end{cases}$$

Algorithm 1 belies the inherent difficulty of: (1) calculation of the partial derivative $g_i^{(k)}$, and (2) an implicit search procedure to update $\lambda_a^{j,(k)}$, both involve calculating intractable integrals. If the support is low-dimensional and the mass is contained in a small volume, then the partition function (and thus the gradient) can be computed by numerical integration (quadrature). Alternatively, a common approach to MLE with an intractable partition function $Z(\theta)$ is Markov Chain Monte Carlo MLE (MCMC-MLE, Geyer and Thompson, 1992). For example, the time complexity at each iteration k is $O(Sn^2)$, where S is the

Algorithm 1 Non-Parametric Exponential Family Coordinate Descent

INPUT: Samples $\mathbf{x}^1, \dots, \mathbf{x}^n \in \mathbb{R}^d$, sufficient statistics $\mathbf{t} : \mathcal{X} \rightarrow \mathcal{H}$, augmented features $t_a^i : \mathcal{H} \rightarrow \mathbb{R}$, $i = 1, \dots, n$, ℓ_1 regularization parameters β

OUTPUT: MLE $\boldsymbol{\theta} = (\lambda^1, \dots, \lambda^d, \lambda_a^1, \dots, \lambda_a^n)$

Initialize $\boldsymbol{\theta}^{(0)}$

Compute the sufficient statistics $E_{\hat{f}_n(\mathbf{x}|\mathbf{x}^{1:n})}[\mathbf{t}(\mathbf{x})]$

repeat

iteration $k = k + 1$, $\boldsymbol{\theta}^{(k)} = \boldsymbol{\theta}^{(k-1)}$

for $i = 1, \dots, d$ **do**

$$g_i^{(k)} = E_{\hat{f}_n} [t^i(\mathbf{x})] - E_{f_n^{NE}(\mathbf{x}|\boldsymbol{\theta}^{(k)})} [t^i(\mathbf{x})]$$

Perform line search along $g_i^{(k)}$ to update $\lambda^{i,(k)}$

end for

for $j = 1 \dots n$ **do**

Solve two equations for λ_a^j ($\lambda_a^{j,-}$ and $\lambda_a^{j,+}$, respectively):

$$E_{f_n^{NE}(\mathbf{x}|\boldsymbol{\theta}^{(k)})} [t_a^j(\mathbf{x})] = E_{\hat{f}_n} [t_a^j(\mathbf{x})] - \beta_j$$

$$E_{f_n^{NE}(\mathbf{x}|\boldsymbol{\theta}^{(k)})} [t_a^j(\mathbf{x})] = E_{\hat{f}_n} [t_a^j(\mathbf{x})] + \beta_j$$

choose $\lambda_a^{j,(k)} = \lambda_a^{j,-}$ if $\lambda_a^{j,-} > 0$

choose $\lambda_a^{j,(k)} = \lambda_a^{j,+}$ if $\lambda_a^{j,+} < 0$

choose $\lambda_a^{j,(k)} = 0$ otherwise

end for

until convergence

return $\boldsymbol{\theta}^{(k)}$

number of Monte-Carlo samples we choose to use. However, we believe developments in optimization (Bach et al., 2011, Shalev-Shwartz and Tewari, 2011, e.g.) will help us find an efficient solution.

4 Application to Modeling of Graphs

In this section, we turn our attention to a problem of learning a distribution over $\mathcal{X} = \mathcal{G}_n$, a set of undirected graphs with n vertices and no self-loops, from a single observed instance $G^* \in \mathcal{G}_n$, an important branch in the analysis of social networks because of complicated relational structure (e.g. Goodreau, 2007). A commonly used approach to this problem which arises in the analysis of social networks is to estimate a distribution using exponential random graph models (ERGMs, e.g., Handcock, 2003, Robins et al., 2007a,b, Wasserman and Pattison, 1996, Wasserman and Robins, 2004). This approach however suffers from the model degeneracy, with estimated models placing probability mass on unrealistic graphs (e.g., complete or empty) and away from the observed instance. We propose a modification to ERGMs utilizing the non-parametric exponential family approach from Section 3 which alleviates the above issue of degeneracy.

4.1 Exponential Random Graph Models

An ERGM (or p^* model) is an exponential family model over \mathcal{G}_n which uses graph statistics as its features⁵. These features are typically motivated by the properties of the networks that are of interest to domain scientists (e.g., sociologists), and may include (among other local and global features) the number of edges

⁵sufficient statistics for the exponential family




Figure 2: Degenerate ERGM and Non-degenerate MPERG. The models are trained based on the observation $t_e(G^*) = 22, t_\Delta(G^*) = 29$. The orange \times is the observed statistics, and the red $+$ is the mode of the learned model. The color bar on the right from red to blue represents the probability mass changing from high to low.

$(t_e(G) = \sum \sum_{1 \leq i < j \leq n} e_{ij})$ and triangles $(t_\Delta(G) = \sum \sum \sum_{1 \leq i < j < k \leq n} e_{ij}e_{ik}e_{jk})$, where $e_{ij} = 1$ if there is edge between nodes i and j , and 0 otherwise. The probability mass for a graph $G \in \mathcal{G}_n$ is defined as

$$\begin{aligned}
 P(G|\hat{\lambda}) &= \frac{1}{Z(\lambda)} \exp \langle \lambda, t(G) \rangle, \\
 Z(\lambda) &= \sum_{G \in \mathcal{G}_n} \exp \langle \lambda, t(G) \rangle.
 \end{aligned} \tag{8}$$

The MLE $\hat{\lambda}$ makes mean statistics of the distribution match that of the observed graph: $E_{P(G|\lambda)} [t(G)] = t(G^*)$.

4.2 Degeneracy

Let $\mathcal{H} = \{t(G) : G \in \mathcal{G}_n\}$ be the set of all possible values for features. Even though in theory if the feature vector for the observed graph is in the relative interior of the convex hull $t(G^*) \in \text{rint}(\text{conv}(\mathcal{H}))$, MLE $\hat{\lambda}$ exists (and is unique if the set of features is linearly independent or minimal), in practice ERGMs often suffer from *degeneracy* (Handcock, 2003) manifested in one of the following ways: (1) MLE procedure does not converge due to numerical instabilities, and (2) MLE is found, but the resulting probability mass is placed mostly on unrealistic graphs (i.e., empty or complete graphs) and little mass is placed in the vicinity of the observed graph (around $t(G^*)$ in \mathcal{H} , c.f. Figure 2).

We focus on addressing the second type of degeneracy; for more information of the reasons of the first type of degeneracy see Handcock (2003), Rinaldo et al. (2009). Several attempts have been made to address the second type of degeneracy issue: Handcock et al. (2008) proposed to use domain knowledge specific feature sets in addition to edge and triangle features; Hunter and Handcock (2006), Hunter et al. (2008) used curved exponential families for ERGMs; Caimo and Friel (2010) suggested Bayesian ERGMs, and Jin and Liang (2012) devised an estimation procedure on stochastic approximation with varying truncation in parameter space. Lunga and Kirshner (2011) suggested the degeneracy issue for interior points may be due to the bounded support \mathcal{H} , and proposed spherical features for modifying the geometry of \mathcal{H} . In summary, there are two main approaches towards fixing degeneracy: 1) modifying the geometry (Handcock et al., 2008, Hunter et al., 2008, Lunga and Kirshner, 2011), and 2) limiting exploration in the canonical parameter space (Caimo and Friel, 2010, Jin and Liang, 2012). Our approach belongs in the first category.

4.3 Mass-Preserving ERGMs

To modify ERGMs, we solve the optimization problem in (5) with the uniform base measure $q(G)$ over possible graphs $G \in \mathcal{G}_n$. Let $t_a(G) = K_{\mathbf{H}}(\mathbf{t}(G^*); \mathbf{t}(G))$, a smoothed mass indicator in the neighborhood of the feature values for the observed graph. The solution is an exponential family probability mass function

$$f(G) = \frac{1}{Z(\boldsymbol{\lambda}, \lambda_a)} \exp[\langle \boldsymbol{\lambda}, \mathbf{t}(G) \rangle + \langle \lambda_a, t_a(G) \rangle]$$

$$Z(\boldsymbol{\lambda}, \lambda_a) = \sum_{G \in \mathcal{G}_n} \exp[\langle \boldsymbol{\lambda}, \mathbf{t}(G) \rangle + \langle \lambda_a, t_a(G) \rangle].$$

which we refer to as *mass-preserving* ERGM (MPERGM). The corresponding objective function

$$l(\boldsymbol{\lambda}, \lambda_a) = \langle \boldsymbol{\lambda}, \mathbf{t}(G^*) \rangle + \langle \lambda_a, t_a(G^*) \rangle - \ln Z(\boldsymbol{\lambda}, \lambda_a) - \beta |\lambda_a|$$

is concave.

There are several challenges with parameter estimation, most encountered before in ERGM fitting (e.g., Hunter et al., 2008). As in the continuous case, the gradient cannot be computed in closed form except for graphs of small size (\mathcal{G}_n for $n \leq 11$). We therefore apply MCMC-MLE approach of Hunter and Handcock (2006), computing $E_f[\mathbf{t}(G)]$ in Algorithm 1 as a sampled average $\frac{1}{S} \sum_{i=1}^S (\mathbf{t}(G^i))$ where $G^{1:S} \stackrel{i.i.d}{\sim} f(G|\boldsymbol{\lambda}, \lambda_a)$. There are, however, two complications with this approach. One, graph sampling from ERGMs is performed using Gibbs sampling and is computationally expensive. Therefore, graphs $G^{1:S}$ are re-sampled only once in several iterations, and reused for other iterations with weights equal to the posterior probabilities. Two, the resulting distribution over graphs can be multi-modal, and according to Hunter and Handcock (2006), Jin and Liang (2012), the sampler can get stuck around the closest mode leading to an incorrect estimate of the gradient. Instead of performing line search, we use the direction of the gradient with a predefined step-size.

5 Experimental Evaluation

5.1 Non-Parametric Exponential Family Density Estimation

We illustrate the behavior of the proposed non-parametric density estimator matching first and second order moment constraints (NPGaussian, i.e. $\mathbf{t}(x) = (x, x^2)$) in the univariate setting. Normal density (in \mathcal{EF} , $\mathcal{N}(0, 1)$), mixture of two normals (not in \mathcal{EF} , $\frac{1}{2}\mathcal{N}(-3, 1) + \frac{1}{2}\mathcal{N}(3, 1)$), and a t-distribution (not in \mathcal{EF} , $\text{df}=6$) are used for simulating i.i.d samples. We vary the sample size from 10 to 1000 for training and compute the out-of-sample likelihood with an evaluation set of 100000 samples for testing. We compared the performance of our non-parametric approach, the model from the true functional family, and another non-parametric approach (KDE). There are two sets of tuning parameters, bandwidth h and the box constraint parameter β , assumed to be the same for all $i = 1, \dots, n$. β was set according to a fixed schedule $\beta(n) = O(1/\sqrt{n})$. h (both for KDE and for our approach) was determined based on cross-validated log-likelihood.⁶ Gaussian kernel function is used for NPGaussian and for KDE. For estimating mixture distribution, the estimated NPGaussian model provides an approximation better than KDE, and perhaps not surprisingly, better than GMM when the training sample size is small (Figure 3(a)). For estimating normal density, the NPGaussian model quickly converges to the normal density as suggested by Theorem 3.1 (Figure 3(b)). We also consider the case when the true sufficient statistics are given to us (constrained NPGaussian, CNPG).

⁶Similar to KDE, the choice of kernel width h is important for obtaining good estimates. To test how the non-parametric exponential family is affected by the choice of h , we employed the same Gaussian kernel function to do density estimation with both KDE and non-parametric Gaussian. It appears that the best bandwidth are different.

The CNPG model shows improvement over NPGaussian for small n . However, as the training sample size increases, both CNPG and NPGaussian show similar performance as the moment constraints $\mathbf{t}(x)$ are more accurately approximated. We also experimented with $O(1/\log(n))$ regularization schedule for β s to estimate the mixed normal distribution. As n increase, the solution for NPGaussian is too sparse and gives a worse performance than KDE (Figure 3(a)). However, it also enjoys a sparse set of augmented parameters λ_a (Figure 3(d)), whereas with schedule $O(1/\sqrt{n})$, NPGaussian keeps adding non-zero λ_a s.

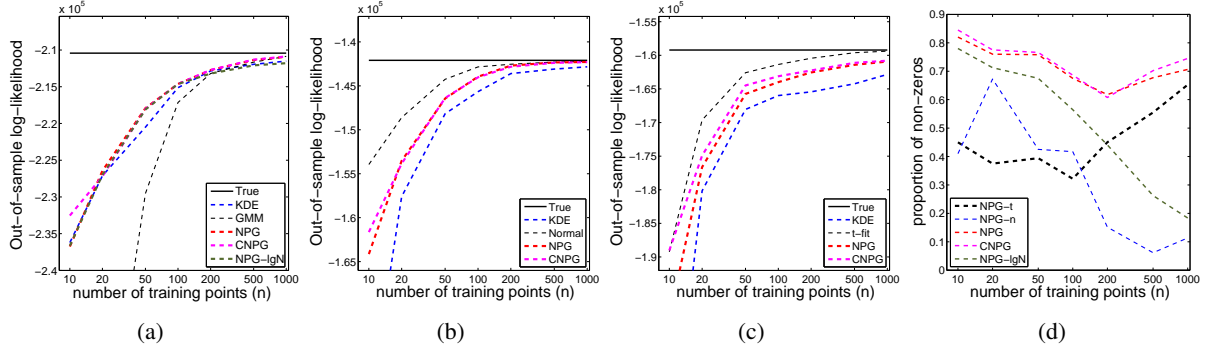


Figure 3: Estimating simple one dimensional densities. Results are averaged over 20 runs. The x axis is in log scale. (a) Mixed normal distribution (b) Normal distribution (c) t distribution (d) Number of non-zero λ_a s. In Figure(a,b,c), NPG: NPGaussian with $O(1/\sqrt{n})$ schedule, NPG-1gN: NPGaussian with $O(1/\log(n))$ schedule, CNPG: constrained NPGaussian with true global moment statistics. In (d), NPG-t: number of non-zeros for estimating t distribution with NPG, NPG-n: number of non-zeros for estimating normal distribution with NPG, NPG: number of non-zeros for estimating mixed normal distribution with NPG, CNPG: number of non-zeros for estimating mixed normal distribution with CNPG, NPG-1gN: number of non-zeros for estimating mixed normal distribution with NPG-1gN.

5.2 Modeling Graphs with MPERGMs

We evaluate the fit of the estimated models by comparing local statistics of the observed graph to that of the samples generated from the estimated distribution.⁷

We make use of three sets of local statistics commonly used as goodness-of-fit measures for ERGMs Hunter et al. (2008): the *degree distribution* (the proportion of nodes with exactly k neighbors), *edgewise shared partner distribution* (the proportion of edges joining nodes with exactly k neighbors in common), and the *minimum geodesic distance* (the proportion of connected node-pairs which has a minimum distance of k).

We consider the number of edges and triangles as sufficient statistics, $\mathbf{t}(G) = (t_e(G), t_\Delta(G))$. First, we consider the toy domain of graphs with 8 nodes, \mathcal{G}_8 . We enumerate all possible $K = 12346$ non-isomorphic graphs and resulting feature tuples, and compute probability mass entries π_1, \dots, π_K . We trained our MPERGM with a Gaussian kernel function with $h = 8, \beta = 0.2$. Figure 2 shows that MPERGM puts larger probability mass around G^* .

We also estimated MPERGMs for several social network data sets, ranging in the number of nodes from 16 to 1024, and with varying density of edges. Since the number of nodes n for these graphs are too large to enumerate \mathcal{G}_n , the graphs are drawn using Gibbs sampler, and the parameters for MPERGMs (and ERGMs, using the R package `ergm` (Hunter et al., 2008)) are estimated using MCMC-MLE. Then 100 samples were generated using the Markov Chain with learned parameters. For MPERGM, the Markov

⁷See Hunter et al. (2008) for a discussion on the evaluation of fit for social networks.

Table 1: Social network data sets. g8: The 8-node graph as in Figure 2(a); Do: The dolphins data set (Lusseau et al., 2003); Kp: The Kapferer data set (Hunter et al., 2008); F1: The Florentine Business data set (Hunter et al., 2008); Fa: The Faux.Mesa.High data set (Hunter et al., 2008); Ja: The Jazz data set (Gleiser and Danon, 2003); Ad: The AddHealth data set (Harris, 2008); Fb: The Facebook data set (Moreno and Neville, 2009); Em: The Email data set (Guimera et al., 2003).

	g8	Do	Kp	F1	Fa	Ja	Ad	Fb	Em
$ V $	8	62	39	16	206	198	803	1024	1133
$t_e(G^*)$	22	159	158	15	203	2742	1985	1012	5451
$t_\Delta(G^*)$	29	95	201	5	62	17899	649	116	5343
No. unique sampled graph	99	100	100	100	100	100	100	96	100
No. unique features	33	33	30	27	14	70	72	74	70
No. max hop of samples	22	306	263	39	341	1435	999	385	1275

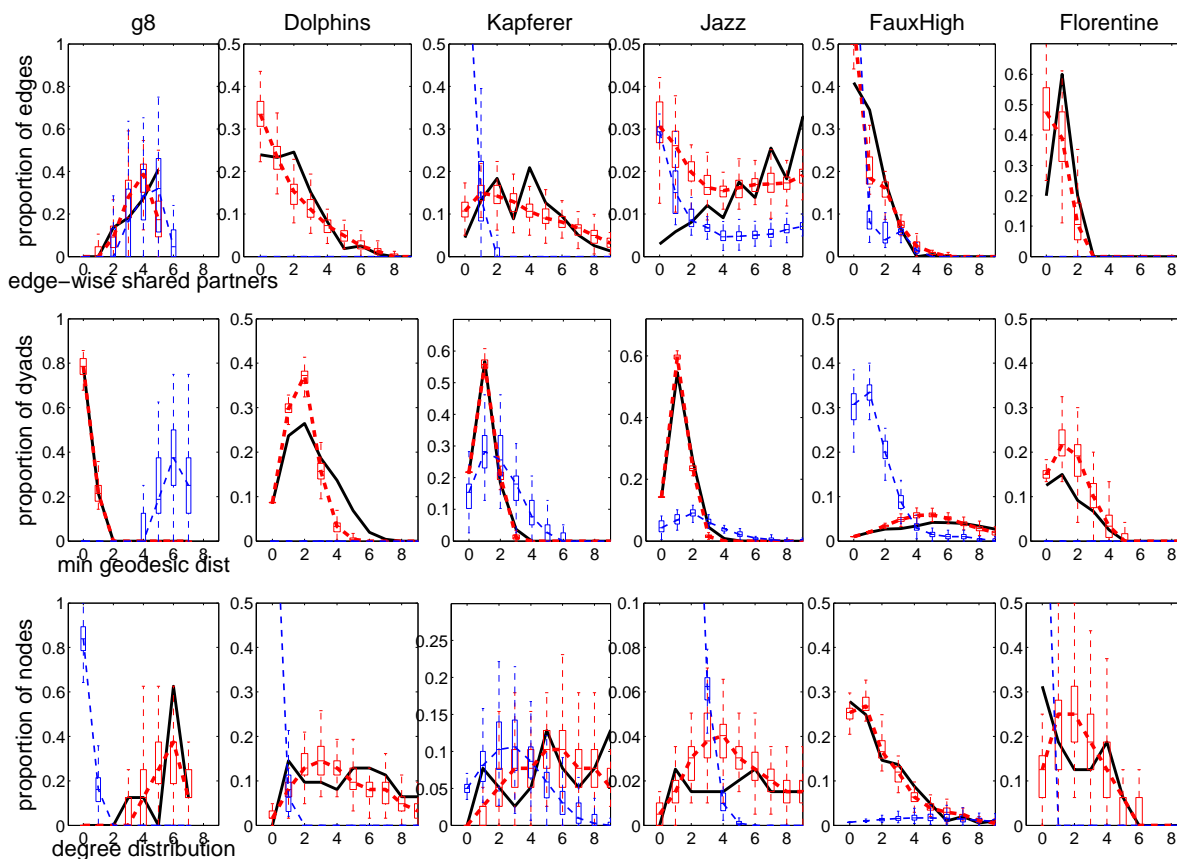


Figure 4: Goodness of fit for small graphs. Gaussian kernel functions are used for MPERGM. ERGM is shown in blue dashed lines, and MPERGM is shown in red dashed lines. Black lines are the statistics for G^* , being closer to black line means better fit.

chain was initialized with the example graph, whereas we are not sure what initial state was used by `ergm`. We then run the chain for a burn-in of 1000 iterations and then use 100 iterations between each draw. The 100 samples are then used to plot the graph statistics for goodness-of-fit test in Figure 4 and Figure 5. For each estimated model, the statistics in Figure 4 & 5 were generated from 100 sampled graphs obtained by running Gibbs with 1000 iterations for burn-in and 100 iterations between samples. We initialized our

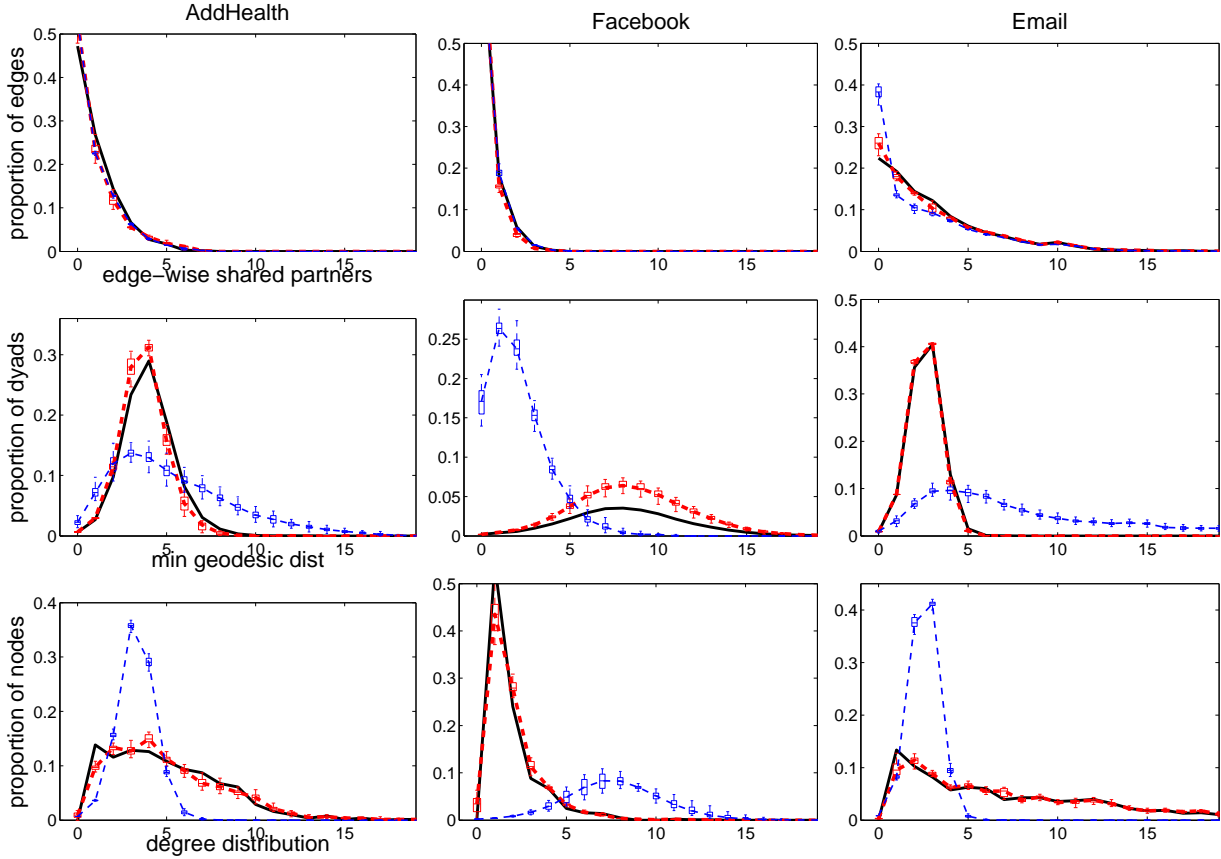


Figure 5: Goodness of fit for large graphs. Gaussian kernel functions are used for MPERGM. ERGM is shown in blue dashed lines, and MPERGM is shown in red dashed lines.

Markov chain with the example graph, whereas we are not sure what initial state did `ergm` use. We used a set of hand-tuned step-size and h for different data sets, and re-scaled the edge and triangle features by a factor of $\frac{1}{t_e(G^*)}$ and $\frac{1}{t_\Delta(G^*)}$. Empirically, we find $h \approx 8$ and a predefined step-size 10 works well for small graphs. For graphs with several thousands of nodes, the pre-defined step-size and h needs to be larger to guarantee reasonable variance and concentration of the model. In Figure 4, ERGM is degenerate for the Florentine and Dolphins dataset, because most sampled graphs have 0-degree nodes (third row), while MPERGM is able to generate samples scoring a similar set of graph statistics. In order to investigate the variance of the learned MPERGM, we count the number of samples that are different in structure (not counting isomorphism) or different in features (number of edges and triangles), while recording the maximum number of unique edge-flips needed to get from the initial state to the sampled graph (number of max hops). The results in Table 1 suggests that our sampler explores \mathcal{G}_n with a considerable range.

6 Discussion and Conclusions

We have proposed a non-parametric exponential family model that is capable of approximating distributions that do not fall within the exponential family empirically. As the data size increases, this model can approximate arbitrary (continuous) densities with tuning parameters controlling the sparsity. And if the true density falls within the exponential family characterized by the chosen features, the estimated non-parametric model converges to the parametric one. The proposed framework results in a non-parametric density estimator

which admits global constraints; if available, this information may require fewer data points to approximate the underlying density. The resulting MLE optimization problem is concave with ℓ_1 penalty term, but raises a computational challenge because the number of inequality constraints is proportional to the number of data points. We also adapted the approach to modify exponential random graph models for graphs to come up with an exponential family model averse to model degeneracy.

As future directions, we would like to investigate the rules for selecting bandwidth parameters, the acceleration of the optimization problem, and efficient sampling techniques for sampling from the non-parametric exponential family.

Acknowledgements

The authors thank Anthony Quas (via mathoverflow.net) and Herman Rubin for their help with proofs of the theoretical results. This research was supported by the NSF Award IIS-0916686.

References

- F. Bach, R. Jenatton, J. Mairal, and G. Obozinski. Convex optimization with sparsity-inducing norms. In S. Sra, S. Nowozin, and S. J. Wright., editors, *Optimization for Machine Learning*, chapter 2. MIT press, 2011.
- O. Barndorff-Nielsen. *Information and Exponential Families in Statistical Theory*. Wiley, New York, 1978.
- R. E. Bellman. *Dynamic Programming*. Princeton University Press, 1957.
- L. D. Brown. *Fundamentals of Statistical Exponential Families*, volume 9 of *Lecture Notes – Monograph Series*. Institute of Mathematical Statistics, Hayward, CA, 1986.
- A. Caimo and N. Friel. Bayesian inference for exponential random graph models. *ArXiv e-prints*, July 2010.
- M. Dudik, S.J. Phillips, and R.E. Schapire. Maximum entropy density estimation with generalized regularization and an application to species distribution modeling. *Journal of Machine Learning Research*, 8: 1217–1260, Jun 2007.
- O. Frank and D. Strauss. Markov graphs. *Journal of the American Statistical Association*, 81(395), September 1986.
- C. J. Geyer and E. A. Thompson. Constrained monte carlo maximum likelihood for dependent data. *Journal of the Royal Statistical Society*, 54(3):pp. 657–699, 1992.
- P. Gleiser and L. Danon. List of edges of the network of jazz musicians. *Adv. Complex Syst.* 6,565, 2003.
- S. Goodreau. Advances in exponential random graph (p) models applied to a large social network. *Social Networks*, 29(2):231–248, May 2007. ISSN 03788733.
- R. Guimera, L. Danon, A. Diaz-Guilera, F. Giralt, and A. Arenas. List of edges of the network of e-mail interchanges between members of the univeristy rovira i virgili (tarragona). *Physics Review E*, 68 (065103(R)), 2003.
- M. S. Handcock. Assessing degeneracy in statistical models of social networks. Technical Report 39, Center for Statistics and the Social Sciences, University of Washington, 2003.
- M. S. Handcock, D. R. Hunter, C. T. Butts, S. M. Goodreau, and M. Morris. statnet: Software tools for the representation, visualization, analysis and simulation of network data. *Journal of Statistical Software*, 24(3), 2008.
- K. Harris. The national longitudinal study of adolescent health (addhealth), waves i & ii, 1994–1996; wave iii, 2001–2002, 2008.
- P. W. Holland and S. Leinhardt. An exponential family of probability distributions for directed graphs (with discussion). *American Statistical Association*, 76(373), 1981.
- D. R. Hunter and M.S. Handcock. Inference in curved exponential family models for networks. *ASA, Journal of Computational and Graphical Statistics*, 15(2), 2006.
- D. R. Hunter, M. S. Handcock, C. T. Butts, S. M. Goodreau, and M. Morris. ergm: A package to fit, simulate and diagnose exponential-family models for networks. *Journal of Statistical Software*, 24(3), 2008.

- I. H. Jin and F. Liang. Fitting social network models using varying truncation stochastic approximation MCMC algorithm. *Journal of Computational and Graphical Statistics*, In Press, 2012.
- D. Lunga and S. Kirshner. Generating similar graphs from spherical features. In *Ninth Workshop on Mining and Learning with Graphs (MLG'11)*, San Diego, CA, August 2011.
- D. Lusseau, K. Schneider, O. J. Boisseau, P. Haase, E. Slooten, and Dawson S. M. The bottlenose dolphin community of Doubtful Sound features a large proportion of long-lasting associations. *Behavioral Ecology and Sociobiology*, 54, 2003.
- S. Moreno and J. Neville. An Investigation of the Distributional Characteristics of Generative Graph Models. In *Proceedings of the 1st Workshop on Information in Networks*, February 2009.
- É. A. Nadaraya. On Non-Parametric Estimates of Density Functions and Regression Curves. *Theory of Probability and its Applications*, 10:186–190, 1965.
- J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer Series in Operations Research. Springer, 2nd edition, 2006.
- E. Parzen. On Estimation of a Probability Density Function and Mode. *The Annals of Mathematical Statistics*, 33(3):1065–1076, 1962.
- A. Rinaldo, S. E. Fienberg, and Y. Zhou. On the geometry of discrete exponential families with application to exponential random graph models. *Electronic Journal of Statistics*, 3:446–484, 2009.
- G. Robins, P. Pattison, Y. Kalish, and D. Lusher. An introduction to exponential random graph (p^*) models for social networks. *Social Networks*, 29(2):173 – 191, 2007a.
- G. Robins, T. Snijders, P. Wang, M. Handcock, and P. Pattison. Recent developments in exponential random graph (p^*) models for social networks. *Social Networks*, 29(2):192 – 215, 2007b.
- H. L. Royden. *Real Analysis*. Prentice Hall, 3rd edition, 1988.
- S. Shalev-Shwartz and A. Tewari. Stochastic methods for l_1 regularized loss minimization. *Journal of Machine Learning Research*, 12:1865–1892, June 2011.
- A. W. van der Vaart. *Asymptotic Statistics*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 1998.
- M. J. Wainwright and M. I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1-2):1–305, 2008.
- S. Wasserman and P. Pattison. Logit models and logistic regression for social networks: An introduction to Markov graphs and p^* model. *Psychometria*, 61(3), September 1996.
- S. Wasserman and G. Robins. An introduction to random graphs, dependence graphs, and p^* . In P. J. Carrington, J. Scott, and S. Wasserman, editors, *Models and Methods in Social Network Analysis*. Cambridge University Press, 2004.
- T. Wu and K. Lange. Coordinate descent algorithms for lasso penalized regression. *The Annals of Applied Statistics*, 2(1):224–244, March 2008.

A Proofs

A.1 Theorem 3.1

To prove Theorem 3.1, let's take a closer look at the augmented constraints. First consider the augmented statistics $t_a^{i,*}$ for sample \mathbf{x}^i .

Lemma A.1. $E_{\hat{f}_n(\mathbf{x}|\mathbf{x}^{1:n})} [t_a^i(\mathbf{x})] = f_n^{KDE}(\mathbf{x}^i|\mathbf{x}^{1:n})$

Proof.

$$E_{\hat{f}_n(\mathbf{x}|\mathbf{x}^{1:n})} [t_a^i(\mathbf{x})] = \frac{1}{n} \sum_{j=1}^n K_{\mathbf{H}}(\mathbf{x}^i; \mathbf{x}^j) = f_n^{KDE}(\mathbf{x}^i|\mathbf{x}^{1:n}),$$

□

i.e., the augmented statistics are the KDE estimates for samples $\mathbf{x}^1 \dots \mathbf{x}^n$ respectively. Parzen (1962) proved the pointwise convergence of KDEs and Nadaraya (1965) proved the uniform convergence for KDEs under further assumptions. We include the first half of (Nadaraya, 1965, Theorem 1) below since it is essential for proving both Theorem 3.1 and Theorem 3.2.

Theorem A.2. Let $\bar{f}_n^{KDE}(\mathbf{x}) = E_{f(\mathbf{y}^{1:n})} [f_n^{KDE}(\mathbf{x}|\mathbf{y}^{1:n})]$ be the expected value of the KDE density given sample points $\mathbf{y}^{1:n}$ *i.i.d.* f .

Suppose $K_{\mathbf{H}}(\mathbf{x}) : \mathbf{x} \in \mathcal{X} \rightarrow \mathbb{R}$ is a function of bounded variation and $f(\mathbf{x})$ is a uniformly continuous density function, and the series $\sum_{n=1}^{\infty} e^{-\gamma n|\mathbf{H}|}$ converges for every positive value of γ .

Then $f_n^{KDE}(\mathbf{x}|\mathbf{x}^{1:n}) \xrightarrow{a.s.} \bar{f}_n^{KDE}(\mathbf{x})$ uniformly on \mathcal{X} .

That is, $\sup_{\mathbf{x} \in \mathcal{X}} |f_n^{KDE}(\mathbf{x}) - \bar{f}_n^{KDE}(\mathbf{x})| \xrightarrow{a.s.} 0$.

Remark 1. It is helpful to note that given $\mathbf{x}^1 \dots \mathbf{x}^n$ *i.i.d.* f ,

$$\begin{aligned} \bar{f}_n^{KDE}(\mathbf{x}^i) &= \int_{\mathbf{x}^1 \dots \mathbf{x}^n} \prod_{j=1}^n f(\mathbf{x}^j) \frac{1}{n} \sum_{j=1}^n K_{\mathbf{H}}(\mathbf{x}^i; \mathbf{x}^j) d\mathbf{x}^1 \dots d\mathbf{x}^n \\ &= \frac{1}{n} \sum_{j=1}^n \int_{\mathbf{x}^j \in \mathcal{X}} K_{\mathbf{H}}(\mathbf{x}^i; \mathbf{x}^j) f(\mathbf{x}^j) d\mathbf{x}^j = \int_{\mathbf{x} \in \mathcal{X}} K_{\mathbf{H}}(\mathbf{x}^i; \mathbf{x}) f(\mathbf{x}) d\mathbf{x} = E_f [t_a^i(\mathbf{x})]. \end{aligned}$$

Theorem 3.1 For a density $f(\mathbf{x}) \in \mathcal{EF}_t$, with domain space \mathcal{X} , feature space \mathcal{H} , and canonical parameters $\boldsymbol{\lambda} \in \mathcal{C} \in \mathbb{R}^d$. Suppose $\mathbf{x}^1, \dots, \mathbf{x}^n \triangleq \mathbf{x}^{1:n}$ is a sequence of independent and identical random vectors drawn from f . Let $f_n^{NE}(\mathbf{x}|\hat{\boldsymbol{\theta}}_n, \mathbf{x}^{1:n}) \in \mathcal{NEF}_s$ be the MLE solution of (5) $\hat{\boldsymbol{\theta}}_n = (\tilde{\boldsymbol{\lambda}}_n, \tilde{\boldsymbol{\lambda}}_{a,n})$, with all $\beta_i = \beta > 0, i = 1, \dots, n$. Assuming

1. \mathcal{X} is compact,
2. t is continuous,
3. \mathcal{EF}_t is a family of uniformly equicontinuous functions w.r.t \mathbf{x} ,
4. Kernel K has bounded variation and has a bandwidth parameter \mathbf{H} such that the series $\sum_{n=1}^{\infty} e^{-\gamma n|\mathbf{H}|}$ converges for every positive value of γ ,

then as $n \rightarrow \infty, \tilde{\boldsymbol{\lambda}}_{a,n} \xrightarrow{p} 0, \forall i = 1, \dots, n$ and $\tilde{\boldsymbol{\lambda}}_n \xrightarrow{p} \boldsymbol{\lambda}$.

Proof. Consider the estimated densities $f_n^E(\mathbf{x}|\hat{\boldsymbol{\lambda}}_n) \in \mathcal{EF}_t$ and $f_n^{NE}(\mathbf{x}|\hat{\boldsymbol{\theta}}_n, \mathbf{x}^{1:n}) \in \mathcal{NEF}_s$.

MLE for regular exponential families converges in probability to the true values of parameters, $\hat{\boldsymbol{\lambda}}_n \xrightarrow{P} \boldsymbol{\lambda}$. (e.g. Wainwright and Jordan, 2008, Chapter 6), (e.g. van der Vaart, 1998, Chapter 5.2)

Since $f(\mathbf{x}|\boldsymbol{\lambda})$ is continuous w.r.t $\boldsymbol{\lambda}$, given any $\mathbf{x} \in \mathcal{X}$, $f_n^E(\mathbf{x}|\hat{\boldsymbol{\lambda}}_n) \xrightarrow{P} f(\mathbf{x}|\boldsymbol{\lambda})$. Because \mathcal{EF}_t is an equicontinuous family, and \mathcal{X} is compact, according to (Royden, 1988, Problem 9.6), $f_n^E \xrightarrow{P} f$ uniformly on \mathcal{X} .

Thus given any $\delta, \beta > 0$, there exists N s.t. when $n > N$, $\forall \mathbf{x} \in \mathcal{X}$, $P(|f_n^E(\mathbf{x}|\hat{\boldsymbol{\lambda}}_n) - f(\mathbf{x}|\boldsymbol{\lambda})| > \frac{\beta}{2}) < \delta$. Then, we have

$$\begin{aligned} P\left(\left|E_{f_n^E(\mathbf{x}|\hat{\boldsymbol{\lambda}}_n)}[t_a^i(\mathbf{x})] - E_f[t_a^i(\mathbf{x})]\right| > \frac{\beta}{2}\right) &= P\left(\int_{\mathbf{x} \in \mathcal{X}} |f_n^E(\mathbf{x}|\hat{\boldsymbol{\lambda}}_n) - f(\mathbf{x}|\boldsymbol{\lambda})| t_a^i(\mathbf{x}) d\mathbf{x} > \frac{\beta}{2}\right) \\ &\leq P\left(\int_{\mathbf{x} \in \mathcal{X}} \sup_{\mathbf{x} \in \mathcal{X}} |f_n^E(\mathbf{x}|\hat{\boldsymbol{\lambda}}_n) - f(\mathbf{x}|\boldsymbol{\lambda})| t_a^i(\mathbf{x}) d\mathbf{x} > \frac{\beta}{2}\right) \\ &= P\left(\sup_{\mathbf{x} \in \mathcal{X}} |f_n^E(\mathbf{x}|\hat{\boldsymbol{\lambda}}_n) - f(\mathbf{x}|\boldsymbol{\lambda})| > \frac{\beta}{2}\right) \\ &< \delta. \end{aligned}$$

Thus for all $\mathbf{x}^i \in \mathcal{X}$, given any $\delta_1, \frac{\beta}{2} > 0$, there exist N_1 , such that when $n > N_1$,

$$P\left(\left|E_{f_n^E(\mathbf{x}|\hat{\boldsymbol{\lambda}}_n)}[t_a^i(\mathbf{x})] - E_f[t_a^i(\mathbf{x})]\right| > \frac{\beta}{2}\right) < \delta_1.$$

Meanwhile, since we have $E_{\hat{f}_n(\mathbf{x}|\mathbf{x}^{1:n})}[t_a^i(\mathbf{x})]$ for the box constraints in (5), we would like to have

$$E_{\hat{f}_n(\mathbf{x}|\mathbf{x}^{1:n})}[t_a^i(\mathbf{x})] \rightarrow E_f[t_a^i(\mathbf{x})]$$

for any $\mathbf{x}^i \in \mathcal{X}$ uniformly as well. This is satisfied under assumptions of Theorem A.2. Then we would have

$$E_{\hat{f}_n(\mathbf{x}|\mathbf{x}^{1:n})}[t_a^i(\mathbf{x})] \xrightarrow{a.s.} E_f[t_a^i(\mathbf{x})]$$

for all $i = 1 \dots n$. Because almost sure convergence implies convergence in probability, we have: for all $\mathbf{x}^i \in \mathcal{X}$, given any $\delta_2, \frac{\beta}{2} > 0$, there exists N_2 , s.t. when $n > N_2$,

$$P\left(\left|E_{\hat{f}_n(\mathbf{x}|\mathbf{x}^{1:n})}[t_a^i(\mathbf{x})] - E_f[t_a^i(\mathbf{x})]\right| > \frac{\beta}{2}\right) < \delta_2.$$

Thus using Triangle Inequality and Boole Inequality, when $n > N = \max(N_1, N_2)$,

$$P\left(\left|E_{\hat{f}_n(\mathbf{x}|\mathbf{x}^{1:n})}[t_a^i(\mathbf{x})] - E_{f_n^E(\mathbf{x}|\hat{\boldsymbol{\lambda}}_n)}[t_a^i(\mathbf{x})]\right| > \beta\right) \leq \delta_1 + \delta_2.$$

This means that the additional constraints will be matched by $[\hat{\boldsymbol{\lambda}}, \mathbf{0}]$ in probability. Because $\hat{\boldsymbol{\lambda}}_n$ is the solution to (3) without the additional constraints, thus it will have smaller KL divergence (larger entropy) than $\hat{\boldsymbol{\theta}}_n$. Therefore $[\hat{\boldsymbol{\lambda}}_n, \mathbf{0}]$ is bound to be the MLE solution to (5). Since $\hat{\boldsymbol{\lambda}}_n \xrightarrow{P} \boldsymbol{\lambda}$, we have $\tilde{\lambda}_{a,n}^i \xrightarrow{P} 0, \forall i = 1, \dots, n$ and $\tilde{\boldsymbol{\lambda}}_n \xrightarrow{P} \boldsymbol{\lambda}$. \square

Remark 2. It can be noted that a sufficient condition for \mathcal{EF}_t to be an equicontinuous exponential family is that \mathcal{X} and \mathcal{C} are both compact. Since in reality, we rarely deal with probability density functions with infinite density values, the conditions for Theorem 3.1, though look restrictive, do not constrain the application domain much. However, since for any regular exponential family the canonical parameter space is open (Brown, 1986, Theorem 3.6), this means Theorem 3.1 for non-parametric exponential family works only on a closed subset \mathcal{C} , of the original canonical parameter space.

A.2 Theorem 3.2

To prove the pointwise convergence result for our non-parametric exponential family, we rely on the convergence of KDEs and Triangle Inequalities. We start off by introducing several lemmas.

Lemma A.3. Let $\tilde{f}_n^{KDE}(\mathbf{x}) = E_{f_n^{NE}(\mathbf{y}^{1:n}|\hat{\boldsymbol{\theta}}_n, \mathbf{x}^{1:n})} [f_n^{KDE}(\mathbf{x}|\mathbf{y}^{1:n})]$ be the expected value of the KDE density given sample points $\mathbf{y}^{1:n} \stackrel{i.i.d.}{\sim} f_n^{NE}(\mathbf{x}|\hat{\boldsymbol{\theta}}_n, \mathbf{x}^{1:n})$. Then $\tilde{f}_n^{KDE}(\mathbf{x}^i) = E_{f_n^{NE}(\mathbf{x}|\hat{\boldsymbol{\theta}}_n, \mathbf{x}^{1:n})} [t_a^i(\mathbf{x})]$.

Proof. Similar as Remark 1,

$$\begin{aligned} \tilde{f}_n^{KDE}(\mathbf{x}^i) &= \int_{\mathbf{x}^1 \dots \mathbf{x}^n} \prod_{j=1}^n f_n^{NE}(\mathbf{x}^j) \frac{1}{n} \sum_{j=1}^n K_{\mathbf{H}}(\mathbf{x}^i; \mathbf{x}^j) d\mathbf{x}^1 \dots d\mathbf{x}^n \\ &= \frac{1}{n} \sum_{j=1}^n \int_{\mathbf{x}^j \in \mathcal{X}} K_{\mathbf{H}}(\mathbf{x}^i; \mathbf{x}^j) f_n^{NE}(\mathbf{x}^j) d\mathbf{x}^j = \int_{\mathbf{x} \in \mathcal{X}} K_{\mathbf{H}}(\mathbf{x}^i; \mathbf{x}) f_n^{NE}(\mathbf{x}) d\mathbf{x} \\ &= E_{f_n^{NE}(\mathbf{x})} [t_a^i(\mathbf{x})]. \end{aligned}$$

□

A pointwise convergence result for the expected KDE density can be found in Parzen (1962):

Theorem A.4. Suppose $f(\mathbf{x})$ is any probability density function which is continuous at point \mathbf{x}^0 . Let $\bar{f}_n^{KDE}(\mathbf{x}) = E_{f(\mathbf{y}^{1:n})} [f_n^{KDE}(\mathbf{x}|\mathbf{y}^{1:n})]$ be the expected value of the KDE density given sample points $\mathbf{y}^{1:n} \stackrel{i.i.d.}{\sim} f$.

Suppose $K_{\mathbf{H}}(\mathbf{x}) : \mathbf{x} \in \mathcal{X} \subset \mathbb{R}^m \rightarrow \mathbb{R}$ is a probability density function which satisfies:

$$\sup_{\mathbf{x} \in \mathcal{X}} K_{\mathbf{H}}(\mathbf{x}) < \infty, \quad (9)$$

$$\lim_{\|\mathbf{x}\| \rightarrow \infty} K_{\mathbf{H}}(\mathbf{x}) \prod_{i=1}^m \mathbf{x}_i = 0, \quad (10)$$

$$\lim_{n \rightarrow \infty} |\mathbf{H}|^{\frac{1}{2}} = 0. \quad (11)$$

Then $\lim_{n \rightarrow \infty} \bar{f}_n^{KDE}(\mathbf{x}^0) = f(\mathbf{x}^0)$.

That is, the expected KDE density at continuity point \mathbf{x}^0 converges to the sampling probability density function f at \mathbf{x}^0 .

Corollary A.5. Given a sample \mathbf{x}^i , $\lim_{n \rightarrow \infty} E_{f_n^{NE}(\mathbf{x})} [t_a^i(\mathbf{x})] = f_n^{NE}(\mathbf{x}^i)$ if $K_{\mathbf{H}}(\mathbf{x})$ is continuous and (9),(10),(11) are satisfied.

Proof. If $K_{\mathbf{H}}(\mathbf{x})$ is continuous, then f_n^{NE} satisfies Theorem A.4's conditions for f . Combining Lemma A.3, we have then have Corollary A.5. □

In addition, we have the mean-square convergence of the KDE density stated in Parzen (1962) as well:

Theorem A.6. Suppose $K_{\mathbf{H}}(\mathbf{x}) : \mathbf{x} \in \mathcal{X} \subset \mathbb{R}^m \rightarrow \mathbb{R}$ is a probability density function which in addition to (9),(10),(11), also satisfies:

$$\lim_{n \rightarrow \infty} n |\mathbf{H}|^{\frac{1}{2}} = \infty \quad (12)$$

and $f(\mathbf{x})$ is probability density function which is continuous at point \mathbf{x}^0 .

Then $\lim_{n \rightarrow \infty} E_{f(\mathbf{x}^{1:n})} \left[(f_n^{KDE}(\mathbf{x}^0 | \mathbf{x}^{1:n}) - f(\mathbf{x}^0))^2 \right] = 0$.

That is, the KDE density at continuity point \mathbf{x}^0 converges in mean square to the true density f at \mathbf{x}^0 .

Lemma A.7.

$$f_n^{NE}(\mathbf{x}^i) \xrightarrow{P} f(\mathbf{x}^i)$$

if

1. $K_{\mathbf{H}}(\mathbf{x})$ is continuous,
2. and (9),(10),(11),(12) are satisfied,
3. and $\lim_{n \rightarrow \infty} \beta_i = 0, \forall i = 1 \dots n$.

Proof. By the Triangle Inequality,

$$\begin{aligned} |f_n^{NE}(\mathbf{x}^i) - f(\mathbf{x}^i)| &\leq |f_n^{NE}(\mathbf{x}^i) - \tilde{f}_n^{KDE}(\mathbf{x}^i)| + |f_n^{KDE}(\mathbf{x}^i) - f(\mathbf{x}^i)| \\ &\quad + |\tilde{f}_n^{KDE}(\mathbf{x}^i) - f_n^{KDE}(\mathbf{x}^i | \mathbf{x}^{1:n})|. \end{aligned}$$

Fix $\xi, \zeta > 0$. Find N_1 using Corollary A.5, N_2 using Theorem A.6, N_3 by the schedule of β s.t.

$$\begin{aligned} P\left(|f_n^{NE}(\mathbf{x}^i) - \tilde{f}_n^{KDE}(\mathbf{x}^i)| > \xi/3\right) &< \zeta/3, \\ P\left(|f_n^{KDE}(\mathbf{x}^i) - f(\mathbf{x}^i)| > \xi/3\right) &< \zeta/3, \\ P\left(|\tilde{f}_n^{KDE}(\mathbf{x}^i) - f_n^{KDE}(\mathbf{x}^i | \mathbf{x}^{1:n})| > \xi/3\right) &< \zeta/3, \end{aligned}$$

for all $n \geq N_1, n \geq N_2$, and $n \geq N_3$, respectively. Set $N = \max\{N_1, N_2, N_3\}$. Then by Boole Inequality, when $n > N$,

$$\begin{aligned} P\left(|f_n^{NE}(\mathbf{x}^i) - f(\mathbf{x}^i)| > \xi\right) &\leq P\left(|f_n^{NE}(\mathbf{x}^i) - \tilde{f}_n^{KDE}(\mathbf{x}^i)| > \xi/3\right) \\ &\quad + P\left(|f_n^{KDE}(\mathbf{x}^i | \mathbf{x}^{1:n}) - f(\mathbf{x}^i)| > \xi/3\right) \\ &\quad + P\left(|\tilde{f}_n^{KDE}(\mathbf{x}^i) - f_n^{KDE}(\mathbf{x}^i | \mathbf{x}^{1:n})| > \xi/3\right) < \zeta. \end{aligned}$$

So $\forall \xi > 0, \lim_{n \rightarrow \infty} P\left(|f_n^{NE}(\mathbf{x}^i) - f(\mathbf{x}^i)| > \xi\right) = 0$. \square

Lemma A.8. Given any uniformly continuous probability density function $f : \mathcal{X} \rightarrow \mathbb{R}$, and n samples $\mathbf{x}^1 \dots \mathbf{x}^n \stackrel{i.i.d.}{\sim} f$. $\forall \xi > 0$, if we draw a new sample $\mathbf{x} \sim f$,

$$\lim_{n \rightarrow \infty} P\left(\left|f(\mathbf{x}) - f\left(\mathbf{x}^{\arg \min_{i=1..n} \|\mathbf{x} - \mathbf{x}^i\|}\right)\right| > \xi\right) = 0.$$

Proof. Since f is uniformly continuous, given any $\xi > 0$, there exists $\varepsilon > 0$, s.t. $\forall \mathbf{x}, \mathbf{y} \in \mathcal{X}$ and $\|\mathbf{x} - \mathbf{y}\| < \varepsilon$, we have $|f(\mathbf{x}) - f(\mathbf{y})| < \xi$. Therefore, let A be the event $\left|f(\mathbf{x}) - f\left(\mathbf{x}^{\arg \min_{i=1..n} \|\mathbf{x} - \mathbf{x}^i\|}\right)\right| > \xi$, let B be the event $\left\|\mathbf{x} - \mathbf{x}^{\arg \min_{i=1..n} \|\mathbf{x} - \mathbf{x}^i\|}\right\| > \varepsilon$ then $A \subset B$. Thus $P(A) < P(B)$, i.e.

$$P\left(\left|f(\mathbf{x}) - f\left(\mathbf{x}^{\arg \min_{i=1..n} \|\mathbf{x} - \mathbf{x}^i\|}\right)\right| > \xi\right) < P\left(\left\|\mathbf{x} - \mathbf{x}^{\arg \min_{i=1..n} \|\mathbf{x} - \mathbf{x}^i\|}\right\| > \varepsilon\right).$$

If we divide \mathcal{X} into countable number of hypercubes with each side length being ε , and index the hypercubes as $Q_1 \dots Q_k \dots$. Let $p_k = \int_{\mathbf{x} \in Q_k} f(\mathbf{x}) d\mathbf{x}$. Then the probability for event B to happen is that \mathbf{x} is the first sample to drop in $Q_k, \forall k \in \mathbb{Z}$. That is, $P_n(B) \triangleq \sum_{k \in \mathbb{Z}} p_k (1 - p_k)^n$. Because $p_k (1 - p_k)^n \rightarrow 0$, and is monotonically decreasing, using monotone convergence theorem, we have $\lim_{n \rightarrow \infty} P_n(B) = 0$.

Therefore,

$$\begin{aligned} \lim_{n \rightarrow \infty} P \left(\left| f(\mathbf{x}) - f \left(\mathbf{x}^{\arg \min_{i=1..n} \|\mathbf{x} - \mathbf{x}^i\|} \right) \right| > \xi \right) &< \lim_{n \rightarrow \infty} P \left(\left\| \mathbf{x} - \mathbf{x}^{\arg \min_{i=1..n} \|\mathbf{x} - \mathbf{x}^i\|} \right\| > \varepsilon \right) \\ &= 0. \end{aligned}$$

□

Now we are ready to prove Theorem 3.2 with Lemma A.7 and Lemma A.8.

Theorem 3.2 Given a probability density function $f(\mathbf{x}) : \mathcal{X} \rightarrow \mathbb{R}$, let $f_n^{NE}(\mathbf{x} | \hat{\theta}_n, \mathbf{x}^{1:n}) \in \mathcal{NEF}_s$ be a solution satisfying (5). If

1. f is uniformly continuous on \mathcal{X} ,
 2. $K_{\mathbf{H}}(\mathbf{x})$ is uniformly continuous on \mathcal{X} ,
 3. and (9),(10),(11),(12) holds,
 4. $\lim_{n \rightarrow \infty} \beta_i = 0, \forall i = 1 \dots n$,
- then $f_n^{NE}(\mathbf{x} | \hat{\theta}_n, \mathbf{x}^{1:n}) \xrightarrow{P} f(\mathbf{x})$ pointwise on \mathcal{X} .

Proof. By the Triangle Inequality, given any $\mathbf{x} \in \mathcal{X}$, and $f_n^{NE}(\mathbf{x})$ trained on n samples $\mathbf{x}^{1:n} \stackrel{i.i.d}{\sim} f$, let $i' \triangleq \arg \min_{i=1..n} \|\mathbf{x} - \mathbf{x}^i\|$.

$$\left| f_n^{NE}(\mathbf{x}) - f(\mathbf{x}) \right| \leq \left| f_n^{NE}(\mathbf{x}^{i'}) - f(\mathbf{x}^{i'}) \right| + \left| f_n^{NE}(\mathbf{x}) - f_n^{NE}(\mathbf{x}^{i'}) \right| + \left| f(\mathbf{x}) - f(\mathbf{x}^{i'}) \right|.$$

Fix $\xi, \zeta > 0$, use Lemma A.7 to find $N_1 \in \mathbb{N}$ and Lemma A.8 to find $N_2, N_3 \in \mathbb{N}$ s.t.

$$\begin{aligned} P \left(\left| f_n^{NE}(\mathbf{x}^{i'}) - f(\mathbf{x}^{i'}) \right| > \xi/3 \right) &< \zeta/3, \\ P \left(\left| f_n^{NE}(\mathbf{x}) - f_n^{NE}(\mathbf{x}^{i'}) \right| > \xi/3 \right) &< \zeta/3, \\ P \left(\left| f(\mathbf{x}) - f(\mathbf{x}^{i'}) \right| > \xi/3 \right) &< \zeta/3, \end{aligned}$$

for all $n \geq N_1, n \geq N_2$, and $n \geq N_3$, respectively. Set $N = \max \{N_1, N_2, N_3\}$. Then by Boole Inequality, when $n > N$,

$$\begin{aligned} P \left(\left| f_n^{NE}(\mathbf{x}) - f(\mathbf{x}) \right| > \xi \right) &\leq P \left(\left| f_n^{NE}(\mathbf{x}^{i'}) - f(\mathbf{x}^{i'}) \right| > \xi/3 \right) \\ &\quad + P \left(\left| f_n^{NE}(\mathbf{x}) - f_n^{NE}(\mathbf{x}^{i'}) \right| > \xi/3 \right) + P \left(\left| f(\mathbf{x}) - f(\mathbf{x}^{i'}) \right| > \xi/3 \right) \\ &< \zeta. \end{aligned}$$

So $\forall \xi > 0, \lim_{n \rightarrow \infty} P \left(\left| f_n^{NE}(\mathbf{x}) - f(\mathbf{x}) \right| > \xi \right) = 0$. □