

Regression from Dependent Observations

Constantinos Daskalakis *
EECS & CSAIL, MIT
costis@csail.mit.edu

Nishanth Dikkala *
EECS & CSAIL, MIT
nishanthd@csail.mit.edu

Ioannis Panageas †
ISTD & SUTD
ioannis@sutd.edu.sg

May 2, 2019

Abstract

The standard linear and logistic regression models assume that the response variables are independent, but share the same linear relationship to their corresponding vectors of covariates. The assumption that the response variables are independent is, however, too strong. In many applications, these responses are collected on nodes of a network, or some spatial or temporal domain, and are dependent. Examples abound in financial and meteorological applications, and dependencies naturally arise in social networks through peer effects. Regression with dependent responses has thus received a lot of attention in the Statistics and Economics literature, but there are no strong consistency results unless multiple independent samples of the vectors of dependent responses can be collected from these models. We present computationally and statistically efficient methods for linear and logistic regression models when the response variables are dependent on a network. Given one sample from a networked linear or logistic regression model and under mild assumptions, we prove strong consistency results for recovering the vector of coefficients and the strength of the dependencies, recovering the rates of standard regression under independent observations. We use projected gradient descent on the negative log-likelihood, or negative log-pseudolikelihood, and establish their strong convexity and consistency using concentration of measure for dependent random variables.

*Supported by NSF awards CCF-1617730 and IIS-1741137, a Simons Investigator Award, a Google Faculty Research Award, and an MIT-IBM Watson AI Lab research grant.

†Supported by SRG ISTD 2018 136. Part of this work was done while the authors were visiting Northwestern University for a spring program in Econometrics

1 Introduction

Linear and logistic regression are perhaps the two most prominent models in Statistics. In their most standard form, these models postulate that a collection of *response variables* y_1, \dots, y_n , which are scalar and binary respectively, are linearly related to a collection of *covariates* $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}^d$ through some *coefficient vector* θ , as follows:

- in vanilla *linear regression* it is assumed that:
 - for all $i \in \{1, \dots, n\}$: $y_i = \theta^\top \mathbf{x}_i + \epsilon_i$, where $\epsilon_i \sim \mathcal{N}(0, 1)$; and
 - y_1, \dots, y_n are independent.
- in vanilla *logistic regression* it is assumed that:
 - for all $i \in \{1, \dots, n\}$ and $\sigma_i \in \{\pm 1\}$: $\Pr[y_i = \sigma_i] = \frac{1}{1 + \exp(-2\theta^\top \mathbf{x}_i \sigma_i)}$; and
 - y_1, \dots, y_n are independent.

It is well-known that, given examples $(\mathbf{x}_i, y_i)_{i=1}^n$, where the y_i 's are sampled independently as specified above, the coefficient vector θ can be estimated to within ℓ_2 -error $O_d\left(\sqrt{\frac{1}{n}}\right)$ in both models, under mild assumptions about the smallest singular value of the matrix whose rows are $\mathbf{x}_1, \dots, \mathbf{x}_n$. In both cases, this can be achieved by solving the corresponding Maximum Likelihood Estimation (MLE) problem, which is concave. In fact, in linear regression, the optimum of the likelihood has a closed form, which is the familiar least-squares estimate.

The assumption that the response variables y_1, \dots, y_n are independent is, however, too strong. In many applications, these variables are observed on nodes of a network, or some spatial or temporal domain, and are *dependent*. Examples abound in financial and meteorological applications, and dependencies naturally arise in social networks through *peer effects*, whose study has recently exploded in topics as diverse as criminal activity (see e.g. [24]), welfare participation (see e.g. [2]), school achievement (see e.g. [36]), participation in retirement plans [18], and obesity (see e.g. [39, 11]). A prominent dataset where peer effects have been studied are data collected by the National Longitudinal Study of Adolescent Health, a.k.a. AddHealth study [26]. This was a major national study of students in grades 7-12, who were asked to name their friends—up to 10, so that friendship networks can be constructed, and answer hundreds of questions about their personal and school life, and it also recorded information such as the age, gender, race, socio-economic background, and health of the students. Estimating models that combine peer and individual effects to predict behavior in such settings has been challenging; see e.g. [31, 5].

1.1 Modeling Dependence

In this paper, we generalize the standard linear and logistic regression models to capture dependencies between the response variables, and show that if the dependencies are sufficiently weak, then both the coefficient vector θ and the strength of the dependencies among the response variables can be estimated to within error $O_d\left(\sqrt{\frac{1}{n}}\right)$. To define our models, we drop the assumption that the response variables y_1, \dots, y_n are independent, but maintain the form of the conditional distribution

that each response variable y_i takes, conditioning on a realization of the other response variables y_{-i} . In particular, for all i , conditioning on a realization of all other variables y_{-i} , the conditional distribution of y_i :

- (in our linear regression model) is a Gaussian of variance 1, as in standard linear regression, except that the mean of this Gaussian may depend on both $\theta^\top \mathbf{x}_i$ and in some restricted way the realizations y_j and the covariates \mathbf{x}_j , for $j \neq i$;
- (in our logistic regression model) takes value +1 with probability computed by the logistic function, as in standard logistic regression, except that the logistic function is evaluated at a point that may depend on both $\theta^\top \mathbf{x}_i$ and in some restricted way the realizations y_j and covariates \mathbf{x}_j , for $j \neq i$.

To capture network effects we parametrize the afore-described general models through a (known) interaction matrix $A \in \mathbb{R}^{n \times n}$ and an (unknown) strength of interactions $\beta \in \mathbb{R}$, as follows.

- In *linear regression* with (A, β) -dependent samples we assume that:

- $\epsilon = \mathbf{y} - X\theta$, with $\epsilon \sim \mathcal{N}(\mathbf{0}, (\beta A + D)^{-1})$.
- Or equivalently, for all i , conditioning on a realization of the response variables y_{-i} :

$$y_i = \theta^\top \mathbf{x}_i + \epsilon_i, \quad (1)$$

where $\epsilon_i \sim \mathcal{N}(\Sigma_i \Sigma_{ii}^{-1} \alpha_i, \frac{\det((\beta A + D)_{-i})}{\det(\beta A + D)} - \Sigma_i \Sigma_{ii}^{-1} \Sigma_i^\top)$, where Σ_i is the i -th row of $(\beta A + D)^{-1}$ by removing the coordinate (diagonal element) i -th, Σ_{ii} is $(\beta A + D)^{-1}$ by removing the i -th column and i -th row, $(\beta A + D)_{-i}$ is $\beta A + D$ by removing i -th row and column and finally column vector $\alpha_j = y_j - \theta^\top \mathbf{x}_j$ (this is the Schur complement for conditional multivariate Gaussians). Observe that $\Sigma_i \Sigma_{ii}^{-1} = -\frac{1}{D_{ii}} \beta A_i^1$ and hence the expectation becomes $-\frac{1}{D_{ii}} \sum_{j \neq i} \beta A_{ij} (y_j - \theta^\top \mathbf{x}_j)$ and moreover the variance becomes $\frac{1}{D_{ii}}$. By the transformation $\epsilon'_i = \epsilon_i + \frac{1}{D_{ii}} \sum_{j \neq i} \beta A_{ij} (y_j - \theta^\top \mathbf{x}_j)$ we get that

$$y_i = \theta^\top \mathbf{x}_i - \frac{1}{D_{ii}} \left[\sum_{j \neq i} \beta A_{ij} (y_j - \theta^\top \mathbf{x}_j) \right] + \epsilon'_i, \quad (2)$$

with $\epsilon'_i \sim \mathcal{N}\left(0, \frac{1}{D_{ii}}\right)$.

- Interpretation: The conditional expectation of y_i is additively perturbed from its expectation $\theta^\top \mathbf{x}_i$ by the weighted average, according to weights βA_{ij} , of how much the other responses are perturbed from their expectations in realization y_{-j} .
- Remark 2: The model proposed in Eq. (2) falls in the realm of auto-regressive models studied by Manski [31] and Bramoullé et al. [5], where it is shown that the model can be identified under conditions on the interaction matrix A . In contrast to our work, one of the conditions imposed on A is that it can be partitioned into many identical blocks (i.e. the weighted graph defined by A has many identical connected components). Thus the response variables cluster into multiple groups that are independently and

¹ A_i denotes the i -row of A by removing coordinate i , i.e., $n - 1$ vector

identically sampled, given the covariates. Instead we want to identify θ and β even when A corresponds to one strongly connected graph, and therefore there is no independence to be exploited.

- In *logistic regression* with (A, β) -dependent samples it is assumed that:

- For all i and $\sigma_i \in \{\pm 1\}$, conditioning on a realization of the response variables y_{-i} :

$$\Pr[y_i = \sigma_i] = \frac{1}{1 + \exp\left(-2\left(\theta^\top \mathbf{x}_i + \beta \sum_{j \neq i} A_{ij} y_j\right) \sigma_i\right)}. \quad (3)$$

- Interpretation: The probability that the conditional distribution of y_i assigns to $+1$ is determined by the logistic function applied to $2\left(\theta^\top \mathbf{x}_i + \beta \sum_{j \neq i} A_{ij} y_j\right)$ instead of $2\theta^\top \mathbf{x}_i$, i.e. it is increased by the weighted average, according to weights βA_{ij} , of the other responses in realization y_{-j} .
- Remark 3: It is easy to see that the joint distribution of random variables (y_1, \dots, y_n) , satisfying the requirements of Eq. (3), is an instance of the Ising model. See Eq. (6). In this Ising model each variable i has external field $\theta^\top \mathbf{x}_i$, and β controls the inverse temperature of the model. The Ising model was originally proposed to study phase transitions in spin systems [27], and has since found myriad applications in diverse research disciplines, including probability theory, Markov chain Monte Carlo, computer vision, theoretical computer science, social network analysis, game theory, and computational biology [29, 7, 21, 16, 22, 19, 35].

A particularly simple instance of our model arises when all covariates \mathbf{x}_i are single dimensional and identical. In this case, our model only has two free parameters, and this setting has been well-studied. [12] consider the consistency of maximum likelihood estimation in this setting. More recent work of Chatterjee [10], Bhattacharya and Mukherjee [4], and Ghosal and Mukherjee [23] has identified conditions on the interaction matrix A under which these parameters can be identified. Our work generalizes these works to the case of multi-dimensional covariates.

Now let us state our results for the above regression models with dependent response variables. We are given a set of observations $(\mathbf{x}_i, y_i)_{i=1}^n$ where the covariates \mathbf{x}_i are deterministic, and the response variables are assumed to have been sampled according to either of the models above, for a given interaction matrix A and an unknown scalar β and coefficient vector θ . Given our observations, we are interested in estimating β and θ . *It is important to stress that we only have **one sample** of the variables (y_1, \dots, y_n) .* In particular, we cannot redraw the response variables many times and derive statistical power from the independence of the samples. This is motivated by our application to network collected data, where we often have no access to independent snapshots of the responses at the nodes of the network. On a technical standpoint, estimating from a *single sample* distinguishes our work from other works in the literature of auto-regressive models and graphical models, and requires us to deal with the challenges of concentration of measure of functions of dependent random variables.

Our main results are stated as Theorems 3.1, for logistic regression, and 4.1, for linear regression. In both cases, the parameters β, θ can be estimated to within error $O_d\left(\sqrt{\frac{1}{n}}\right)$, the dependence of

the rate on n matching that of vanilla logistic regression and vanilla linear regression respectively. These results hold under the assumptions of Table 1. We note that the assumptions on θ , β , and the covariates are standard, even in the case of vanilla regression. Moreover, the bounds on the norm of A have been shown to be necessary for logistic regression by [4, 23]. And the minimum singular value condition for matrix AX is mild, and holds for various ensembles of A ; see e.g. Corollary 4.1 shown using Ky Fan inequalities [20].

Proof Overview: The estimation algorithms in both Theorem 3.1 and Theorem 4.1 are instances of Projected Gradient Descent (PGD). In the linear case (Theorem 4.1, PGD is applied to the negative log-likelihood of the observations (y_1, \dots, y_n) . However, the log-likelihood is not convex, so we perform a re-parametrization of the model, indeed an *overparametrization* of the model that renders it convex. Showing strong convexity of the re-parametrized negative log-likelihood requires some mild linear algebra. It has to be established that despite the overparametrization the optimum collapses to the right dimensionality, and can be used to recover the original parameters. A more complete overview of the approach is presented in the beginning of Section 4.

In the logistic case (Theorem 3.1), we do not run PGD on the negative log-likelihood but the negative log-*pseudolikelihood*. Pseudolikelihood is the product of the conditional probabilities of each response y_i , conditioning on all other responses y_{-i} . Pseudolikelihood is trivially convex, but we need to establish that its optimum is close to the true parameters and also that it is strongly convex. We show both properties via concentration results for functions of dependent random variables. To show that the maximum of the pseudolikelihood is close to the true parameters we use exchangeable pairs, adapting [9]. To show that it is strongly convex we show additional properties of A which are implied by our assumptions. Combining these with a new concentration inequality, we obtain the desired bound. A more complete overview of the approach is presented in Section 3.2.

Other Related Work: We have already reviewed the work that is most relevant to ours from the Economics, Probability Theory, and Statistics literature. Further discussion of the Econometrics and Statistics literature on the theory and applications of regression with dependent observations is discussed in [30]. There is another strand of literature studying generalization bounds that can be attained when learning from sequences of dependent observations; see e.g. [33, 32, 41, 38, 34, 1]. These works assume, however, that the sequence of observations is a *stationary process*, which does not hold in our models, and they impose strong mixing conditions on that sequence. Finally, we note that generalized linear regression, which accommodates dependencies among the response variables, cannot be applied directly to our linear regression setting to estimate θ , because the covariance matrix of our response variables depends on the parameter β , which is unknown and thus needs to be disentangled before bounding the error in the estimation of θ .

In the case of logistic regression, there has been a lot of work showing that under certain *high-temperature* conditions on the Ising model (which are similar to the assumptions we make in our paper), one can perform many statistical tasks such as learning, testing and sampling of Ising models efficiently [28, 14, 13, 15, 25, 17].

2 Preliminaries

We use bold letter such as \mathbf{x}, \mathbf{y} to denote vectors and capital letters X, Y, A, D to denote matrices. All vectors are assumed to be column vectors, i.e. $\text{dim} \times 1$. We will refer to A_{ij} as the $(i, j)^{\text{th}}$ entry of matrix A . We will use the following matrix norms. For a $n \times n$ matrix A ,

$$\|A\|_2 = \max_{\|x\|_2=1} \|Ax\|_2, \quad \|A\|_\infty = \max_{j \in [n]} \sum_{i=1}^n |A_{ij}|,$$

$$\|A\|_F = \sqrt{\sum_{i=1}^n \sum_{j=1}^n A_{ij}^2}.$$

When A is a symmetric matrix we have that $\|A\|_2 \leq \|A\|_\infty \leq \|A\|_F \leq \sqrt{n} \|A\|_2 \leq \sqrt{n} \|A\|_\infty$. We use λ to denote eigenvalues of a matrix and σ to denote singular values. λ_{\min} refers to the smallest eigenvalue and λ_{\max} to the largest, and similar notation is used for the singular values as well.

We will say an estimator $\hat{\theta}_n$ is consistent with a rate \sqrt{n} (or equivalently \sqrt{n} -consistent) with respect to the true parameter θ_0 if there exists an integer n_0 and a constant $C > 0$ such that for every $n > n_0$, with probability at least $1 - o(1)$,

$$\|\hat{\theta}_n - \theta_0\|_2 \leq \frac{C}{\sqrt{n}}$$

We utilize the following two well-known examples of graphical models to characterize dependencies in our logistic and linear regression models respectively.

1. **Ising Model:** Given an unweighted undirected graph $G(V, E)$ with adjacency matrix A and assignment $\sigma : V \rightarrow \{-1, +1\}^n$, an Ising model is the following probability distribution on the 2^n configurations of σ :

$$\Pr\{\mathbf{y} = \sigma\} = \frac{\exp(\sum_{v \in V} h_v \sigma_v + \beta \sigma^\top A \sigma)}{Z_G(\beta, \theta)} \quad (4)$$

where

$$Z(G) = \sum_{\tilde{\sigma}} \exp\left(\sum_{v \in V} h_v \tilde{\sigma}_v + \beta \tilde{\sigma}^\top A \tilde{\sigma}\right)$$

is the partition function of the system (or renormalization factor). Moreover the term $\sum_v h_v \sigma_v$ is called the external field. It can be observed that, without loss of generality, we can restrict the matrix A to have zeros on its diagonal.

2. **Gaussian Graphical Model:** Let $G = (V, E)$ be an undirected graph with $V = [n]$. A random vector $X \in \mathbb{R}^n$ is said to be distributed according to (undirected) Gaussian Graphical model with graph G if X has a multivariate Gaussian distribution $\mathcal{N}(\mu, \Sigma)$ with

$$(\Sigma^{-1})_{ij} = 0 \quad \forall (i, j) \notin E, \quad (5)$$

where the density function $f_{\mu, \Sigma}(\cdot)$ of $\mathcal{N}(\mu, \Sigma)$ is

$$f_{\mu, \Sigma}(x) = \frac{\exp\left(-\frac{1}{2}(x - \mu)^T \Sigma^{-1}(x - \mu)\right)}{(2\pi)^{n/2} \det(\Sigma)^{1/2}}$$

under the condition that Σ is positive semi-definite (Σ^{-1} is also known as the precision matrix).

2.1 Some Useful Lemmas from Literature

Weyl's inequalities are useful to understand how the spectra of symmetric matrices change under addition. We state them here for reference.

Lemma 2.1 (Weyl's Inequalities). *Let A , B and C be three $n \times n$ symmetric matrices with real entries such that $A = B + C$. Let $\lambda_1^A \geq \lambda_2^A \geq \dots \geq \lambda_n^A$, $\lambda_1^B \geq \lambda_2^B \geq \dots \geq \lambda_n^B$, $\lambda_1^C \geq \lambda_2^C \geq \dots \geq \lambda_n^C$ be their eigenvalues respectively. Then we have for all $i \in [n]$, $\lambda_i^B + \lambda_n^C \leq \lambda_i^A \leq \lambda_i^B + \lambda_1^C$.*

We will use the following concentration inequality which is standard in literature.

Theorem 2.1 ([40], Remark 5.40). *Assume that X is an $n \times d$ matrix whose rows X_i are independent sub-gaussian random vectors in R^d with second moment matrix Σ . Then for every $t \geq 0$, the following inequality holds with probability at least $1 - 2\exp(-ct^2)$,*

$$\left\| \frac{1}{n} X^T X - \Sigma \right\|_2 \leq \max(\delta, \delta^2)$$

with $\delta = C\sqrt{\frac{d}{n}} + \frac{t}{\sqrt{n}}$.

Remark 2.1. *By choosing t to be $\Theta(\sqrt{\ln n})$, it follows that with probability $1 - \frac{1}{\text{poly}(n)}$ we get that*

$$\left\| \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i^T - Q \right\|_2 \text{ is } O\left(\sqrt{\frac{\ln n}{n}}\right),$$

from which follows that $\lambda_{\min}(\frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i^T)$ is at least $\lambda_{\min}(Q) - O\left(\sqrt{\frac{\ln n}{n}}\right)$ with probability $1 - \frac{1}{\text{poly}(n)}$ (by Weyl's inequality).

Lemma 2.2 (Useful Inequalities on Singular Values). *The following inequalities hold:*

1. *Let W be a $n \times n$ matrix. It holds that $|\lambda_{\min}(W + W^T)| \leq 2\sigma_{\min}(W)$ (see [20]).*
2. *Let W, Z be matrices. It holds that $\sigma_{\min}(WZ) \leq \sigma_{\min}(W) \|Z\|_2$ (folklore).*
3. *Let W, Z be matrices, then $\|WZ\|_F^2 \leq \|W\|_2^2 \|Z\|_F^2$ (folklore).*

Lemma 2.3 (Expectation and Variance of a Quadratic form of a Gaussian Distribution). *Let $\mathbf{z} \sim \mathcal{N}(\mu, \Sigma)$ and we have the quadratic form $f(\mathbf{z}) := \mathbf{z}^T A \mathbf{z} + \mathbf{b}^T \mathbf{z} + c$. It holds that*

$$\mathbb{E}_z[f(\mathbf{z})] = \text{tr}(A\Sigma) + f(\mu), \mathbb{V}_z[f(\mathbf{z})] = 2\text{tr}(A\Sigma A\Sigma) + 4\mu^T A \Sigma A \mu + 4\mathbf{b}^T \Sigma A \mu + \mathbf{b}^T \Sigma \mathbf{b}.$$

Table 1 lists the assumptions under which our main theorems for logistic and linear regression hold.

Table 1: List of conditions under which our main consistency results (Theorems 3.1 and 4.1) hold.

Parameter	Logistic	Linear
θ	$(-\Theta, \Theta)^d$	$(-\Theta, \Theta)^d$
\mathbf{x}_i feature vectors with covariance matrix $Q = n^{-1}X^\top X^{-2}$	Support in $[-M, M]^d$ and $\lambda_{\max}(Q), \lambda_{\min}(Q)$ positive constants	No restriction in the support and $\lambda_{\max}(Q), \lambda_{\min}(Q)$ positive constants
D	Not Applicable	diagonal matrix with positive constant entries
A	symmetric, zero diagonal, $\ A\ _\infty \leq 1$ and $\ A\ _F^2 \geq cn$	symmetric, zero diagonal, $\ A\ _2 \leq 1$ and $\ A\ _F^2 \geq cn$
β	$(-B, B)$	$\lambda_{\min}((\beta A + D)^{-1}) > \rho_{\min}$, $\lambda_{\max}((\beta A + D)^{-1}) < \rho_{\max}$ and ρ_{\min}, ρ_{\max} positive constants for all $\beta \in (-B, B)$
$n^{-1}X^\top A^\top (I - DX(X^\top D^2 X)^{-1}X^\top D)AX$	No assumption	Minimum eigenvalue a positive constant ρ_{DAX}

3 Logistic Regression with Dependent Data

In this section we look at the problem of logistic regression with dependent data.

3.1 Our model

We are interested in a generalization of the Ising model on graph $G = (V, E)$ with $|V| = n$, where each vertex $i \in G$ has a feature vector $\mathbf{x}_i \in \mathbb{R}^d$. Moreover there is an unknown parameter $\theta \in \mathbb{R}^d$ and the corresponding probability distribution induces to the following:

$$\Pr\{\mathbf{y} = \sigma\} = \frac{\exp\left(\sum_{i=1}^n (\theta^\top \mathbf{x}_i) \sigma_i + \beta \sigma^\top A \sigma\right)}{Z(G)}, \quad (6)$$

where A is a symmetric matrix with zeros on the diagonal. Given one sample \mathbf{y} and the knowledge of the matrix A , we would like to infer β, θ .

We now study some conditions under which we can attain consistent estimates of the parameters of the model. Combined with some standard assumptions on the data-generating process of the feature vectors all our assumptions are listed in Table 1.

Theorem 3.1 (Logistic Regression with Dependent Samples). *Consider the model of (6). The Maximum Pseudo-Likelihood Estimate (MPLE) $(\hat{\theta}_{MPL}, \hat{\beta}_{MPL})$ is consistent with a rate of \sqrt{n} as long as (θ_0, β_0) and the features X satisfy the conditions of Column 2 in Table 1. Formally, for*

each constant $\delta > 0$ and n sufficiently large

$$\left\| (\hat{\theta}_{MPL}, \hat{\beta}_{MPL}) - (\theta_0, \beta_0) \right\|_2 \leq O_d \left(\sqrt{\frac{1}{n}} \right)$$

with probability $1 - \delta$. Moreover, we can compute a vector $(\tilde{\theta}, \tilde{\beta})$ with $\left\| (\hat{\theta}_{MPL}, \hat{\beta}_{MPL}) - (\tilde{\theta}, \tilde{\beta}) \right\|_2 \leq O_d \left(\sqrt{\frac{1}{n}} \right)$ in $O(\ln n)$ iterations of projected gradient descent (Algorithm in Section 5) where each iteration takes at most $O(dn)$ time, with probability $1 - \delta$.

Remark 3.1 (Necessity of an Upper Bound on $\|A\|_\infty$ and boundedness of β_0). *If $\|A\|_\infty$ scales with n then no consistent estimator might exist. This is because the peer effects through $\beta_0 A$ will dominate the outcome of the samples and will nullify the signal coming from $\theta_0^\top X$. Similarly one requires β_0 to be bounded as well to preserve some signal to enable recovery of θ_0 .*

Remark 3.2 (Necessity of the Lower Bound on $\|A\|_F$). *It was shown in [4] (Corollary 2.4 (b)) and [23] (Theorem 1.13) that when the condition $\|A\|_F^2 > cn$ is violated, we have specific examples where it is impossible to get consistent estimators for (θ_0, β_0) . The first instance is the Curie-Weiss model $CW(n, \beta, h)$ ($A_{ij} = \frac{1}{n}$ for all $i \neq j$). Note that $\|A\|_F^2 = O(1)$ in this case. The second instance is dense random graphs, i.e. $G(n, p)$ where p is a constant independent of n and A is chosen to be the adjacency matrix scaled down by the average degree of the graph, i.e. $A_{ij} = \frac{1}{(n-1)p} \mathbf{1}_{(i,j) \in E}$.*

Remark 3.3. *If the parameter β_0 is known, the condition that $\|A\|_F^2 \geq cn$ is not necessary for consistency of the MPL estimate $\hat{\theta}_{MPL}$. For instance, consider the independent case where $\beta_0 = 0$. Then, to recover θ , we do not need $\|A\|_F^2 \geq cn$.*

Remark 3.4. *Our approach achieves a $\sqrt{n/d}$ rate of consistency if $\|\mathbf{x}\|_2 \times \|\theta\|_2 = O(1)$.*

Example Instantiations of Theorem 3.1 Two example settings where the conditions required for Theorem 3.1 to hold are satisfied are

- A is the adjacency matrix of graphs with bounded degree d scaled down so that $\|A\|_2 \leq 1$.
- A is the adjacency matrix of a random d -regular graph.

3.2 Technical Overview

Estimation in Ising models is a well-studied problem which offers a lot of interesting technical challenges. A first approach one considers is maximum likelihood estimation. However the intractability of computing the partition function poses a serious obstacle for the MLE. Even if one could approximate the partition function, proving consistency of the MLE is a hard task. To circumvent these issues we take a maximum pseudo-likelihood approach. This was proposed by Julian Besag [3] and analyzed for inference problems on Ising models by Chatterjee [10] and others ([4],[23]). Given a sample of response variables \mathbf{y} let $f_i(\theta, \beta, \mathbf{y})$ denote the condition likelihood of observing y_i conditioned on everyone else. The pseudo-likelihood estimator of \mathbf{y} is

$$\left(\hat{\theta}_{MPL}, \hat{\beta}_{MPL} \right) = \operatorname{argmax}_{\theta, \beta} \prod_{i=1}^n f_i(\theta, \beta, \mathbf{y}). \quad (7)$$

This does away with the problematic partition function and retains concavity in the parameters θ, β . To show that the MPLE is consistent we need to show that its global optimum $(\hat{\theta}_{MPL}, \hat{\beta}_{MPL})$ is close in ℓ_2 distance to (θ_0, β_0) . We achieve this by showing two things hold simultaneously.

- The log pseudo-likelihood is strongly concave everywhere. This will tell us that the gradient of the log pseudo-likelihood quickly increases as we move away from $(\hat{\theta}_{MPL}, \hat{\beta}_{MPL})$ where it is 0.
- The norm of the gradient of the log pseudo-likelihood is small at when evaluated at (θ_0, β_0) hence implying proximity to the MPL estimates due to strong concavity.

We show that both these conditions are satisfied with high probability over the draw of our samples. Showing that the norm of the gradient is bounded involves obtaining variance bounds on two functions of the Ising model (Lemmas 3.2 and 3.3), and showing strong concavity amounts to showing a linear in n lower bound on a particular quadratic function (Lemma ??). Both these properties are challenging to prove because of the dependences between samples. To tackle the lack of independence, the proofs require a rich set of technical frameworks. In particular, to show the variance bounds we use the technique of exchangeable pairs developed by Chatterjee [8]. The boundedness of $\|A\|_\infty$ is necessary to have these concentration results. To show strong concavity of the log pseudolikelihood we first prove some properties of the matrix A together with an additional variance bound again shown via exchangeable pairs. The lower bound on $\|A\|_F$ is necessary to achieve strong concavity. Finally, we show in Section 5 that computing the MPLE can be achieved efficiently using projected gradient descent where after each step we project back into the space restricted by the conditions of Table 1. We describe each of these steps formally now.

3.3 Analyzing the Maximum Pseudolikelihood Estimator (MPLE)

We will treat terms not involving n as constants for the purposes of our analysis. We start by analyzing the maximum pseudo-likelihood estimator. Given the feature vector of the i^{th} sample \mathbf{x}_i , we denote by x_{ik} the k^{th} element of \mathbf{x}_i . Let $m_i(\mathbf{y}) := \sum_{j=1}^n A_{ij}y_j$ and let $\mathbb{B} = [-\Theta, \Theta]^d \times [-B, B]$ (the true parameters lie in the interior of \mathbb{B}). The pseudolikelihood for a specific sample \mathbf{y} is given by:

$$PL(\theta, \beta) := \prod_{i=1}^n \frac{\exp(\theta^\top \mathbf{x}_i y_i + \beta m_i(\mathbf{y}) y_i)}{\exp(\theta^\top \mathbf{x}_i + \beta m_i(\mathbf{y})) + \exp(-\theta^\top \mathbf{x}_i - \beta m_i(\mathbf{y}))}. \quad (8)$$

The normalized log pseudolikelihood for a specific sample \mathbf{y} is given by:

$$LPL(\theta, \beta) := \frac{1}{n} \log PL(\theta, \beta) = -\ln 2 + \frac{1}{n} \sum_{i=1}^n [y_i \beta m_i(\mathbf{y}) + y_i (\theta^\top \mathbf{x}_i) - \ln \cosh(\beta m_i(\mathbf{y}) + \theta^\top \mathbf{x}_i)]. \quad (9)$$

The first order conditions give:

$$\begin{aligned} \frac{\partial LPL(\hat{\theta}_{MPL}, \hat{\beta}_{MPL})}{\partial \beta} &= \frac{1}{n} \sum_{i=1}^n \left[y_i m_i(\mathbf{y}) - m_i(\mathbf{y}) \tanh(\hat{\beta}_{MPL} m_i(\mathbf{y}) + \hat{\theta}_{MPL}^\top \mathbf{x}_i) \right] = 0, \\ \frac{\partial LPL(\hat{\theta}_{MPL}, \hat{\beta}_{MPL})}{\partial \theta_k} &= \frac{1}{n} \sum_{i=1}^n \left[y_i x_{i,k} - x_{i,k} \tanh(\hat{\beta}_{MPL} m_i(\mathbf{y}) + \hat{\theta}_{MPL}^\top \mathbf{x}_i) \right] = 0. \end{aligned} \quad (10)$$

The Hessian $H_{(\theta, \beta)}$ is given by:

$$\begin{aligned} \frac{\partial^2 LPL(\theta, \beta)}{\partial \beta^2} &= -\frac{1}{n} \sum_{i=1}^n \frac{m_i^2(\mathbf{y})}{\cosh^2(\beta m_i(\mathbf{y}) + \theta^\top \mathbf{x}_i)}, \\ \frac{\partial^2 LPL(\theta, \beta)}{\partial \beta \partial \theta_k} &= -\frac{1}{n} \sum_{i=1}^n \frac{x_{i,k} m_i(\mathbf{y})}{\cosh^2(\beta m_i(\mathbf{y}) + \theta^\top \mathbf{x}_i)}, \\ \frac{\partial^2 LPL(\theta, \beta)}{\partial \theta_i \partial \theta_k} &= -\frac{1}{n} \sum_{i=1}^n \frac{x_{i,1} x_{i,k}}{\cosh^2(\beta m_i(\mathbf{y}) + \theta^\top \mathbf{x}_i)}. \end{aligned} \quad (11)$$

Writing the Hessian in a compact way we get

$$H_{(\theta, \beta)} = -\frac{1}{n} \sum_{i=1}^n \frac{1}{\cosh^2(\beta m_i(\mathbf{y}) + \theta^\top \mathbf{x}_i)} X_i X_i^\top$$

where $X_i = (\mathbf{x}_i, m_i(\mathbf{y}))^\top$. Thus $-H$ is a positive semidefinite matrix and LPL is concave. Moreover if $(\theta, \beta) \in \mathbb{B}$ it follows that

$$\frac{1}{\cosh^2(B+d \cdot M \cdot \Theta)} \cdot \left(\frac{1}{n} \sum_{i=1}^n X_i X_i^\top \right) \preceq -H_{(\theta, \beta)} \preceq \left(\frac{1}{n} \sum_{i=1}^n X_i X_i^\top \right). \quad (12)$$

Remark 3.5. Observe that $\|X_i\|_2^2 = \|x_i\|_2^2 + m_i^2(\mathbf{y}) \leq d\Theta^2 + 1$ (assuming that $\|A\|_\infty \leq 1$ trivially holds $|m_i(x)| \leq 1$). It is easy to see that $\lambda_{\max}(-H_{(\theta, \beta)}) \leq \sqrt{d\Theta^2 + 1}$ for all $(\theta, \beta) \in \mathbb{R}^{d+1}$, hence $-LPL$ is a $\sqrt{d\Theta^2 + 1}$ -smooth function, i.e. $\nabla -LPL$ is $\sqrt{d\Theta^2 + 1}$ -Lipschitz.

3.4 Consistency of the MPLE

Our argument for showing consistency of the MPLE uses Lemma 3.1.

Lemma 3.1. Let (θ_0, β_0) be the true parameter. We define $(\theta_t, \beta_t) = (1-t)(\theta_0, \beta_0) + t(\hat{\theta}_{MPL}, \hat{\beta}_{MPL})$ and let $\mathcal{D} \in [0, 1]$ be the largest value such that $(\theta_{\mathcal{D}}, \beta_{\mathcal{D}}) \in \mathbb{B}$ (if it does not intersect the boundary of \mathbb{B} , then $\mathcal{D} = 1$). Then,

$$\begin{aligned} \|\nabla LPL(\theta_0, \beta_0)\|_2 &\geq \mathcal{D} \min_{(\theta, \beta) \in \mathbb{B}} \lambda_{\min}(-H_{(\theta, \beta)}) \left\| (\theta_0 - \hat{\theta}_{MPL}, \beta_0 - \hat{\beta}_{MPL}) \right\|_2 \\ &= \min_{(\theta, \beta) \in \mathbb{B}} \lambda_{\min}(-H_{(\theta, \beta)}) \left\| (\theta_0 - \theta_{\mathcal{D}}, \beta_0 - \beta_{\mathcal{D}}) \right\|_2 \end{aligned}$$

Proof. We drop the subscript MPL from the estimates for brevity. We set

$$\begin{aligned} g(t) &:= (\theta_0 - \hat{\theta}, \beta_0 - \hat{\beta})^\top \nabla LPL(\theta_t, \beta_t), \\ g'(t) &= -(\theta_0 - \hat{\theta}, \beta_0 - \hat{\beta})^\top H_{(\theta_t, \beta_t)} (\theta_0 - \hat{\theta}, \beta_0 - \hat{\beta}). \end{aligned}$$

Observe that $\mathcal{D} = \frac{\|(\theta_{\mathcal{D}} - \theta_0, \beta_{\mathcal{D}} - \beta_0)\|_2}{\|(\hat{\theta} - \theta_0, \hat{\beta} - \beta_0)\|_2}$. Since H is negative semidefinite we have that $g'(t) \geq 0$ (*). It holds that

$$\begin{aligned} \left\| (\theta_0 - \hat{\theta}, \beta_0 - \hat{\beta}) \right\|_2 \cdot \|\nabla LPL(\theta_0, \beta_0)\|_2 &\geq |(\theta_0 - \hat{\theta}, \beta_0 - \hat{\beta})^\top \nabla LPL(\theta_0, \beta_0)| \\ &= |g(1) - g(0)| = \left| \int_0^1 g'(t) dt \right| \\ &\geq \left| \int_0^{\mathcal{D}} g'(t) dt \right| \text{ by (*)} \\ &\geq \mathcal{D} \min_{(\theta, \beta) \in \mathbb{B}} \lambda_{\min}(-H_{(\theta, \beta)}) \left\| (\theta_0 - \hat{\theta}, \beta_0 - \hat{\beta}) \right\|_2^2 \\ &= \min_{(\theta, \beta) \in \mathbb{B}} \lambda_{\min}(-H_{(\theta, \beta)}) \left\| (\theta_{\mathcal{D}} - \theta_0, \beta_{\mathcal{D}} - \beta_0) \right\|_2 \times \\ &\quad \times \left\| (\theta_0 - \hat{\theta}, \beta_0 - \hat{\beta}) \right\|_2 \end{aligned}$$

□

We apply Lemma 3.1 by showing

1. a concentration result for $\|\nabla LPL(\theta_0, \beta_0)\|_2^2$ around $1/n$ (Section 3.5) and
2. a (positive constant) lower bound for $\min_{(\theta, \beta) \in \mathbb{B}} \lambda_{\min}(-H_{(\theta, \beta)})$ (Section 3.6).

We combine the above with the observation that $\mathcal{D} \rightarrow 1$ as $n \rightarrow \infty$ (i.e., $\mathcal{D} \geq \frac{1}{2}$ for n sufficiently large). This is true because $\|(\theta_{\mathcal{D}} - \theta_0, \beta_{\mathcal{D}} - \beta_0)\|_2 \rightarrow 0$ as $n \rightarrow \infty$ (is of order $\frac{1}{\sqrt{n}}$ by showing the promised concentration result and the lower bound). Also note that any point on the boundary of \mathbb{B} has a fixed distance to (θ_0, β_0) since it lies in the interior. Hence $\|(\theta_{\mathcal{D}} - \theta_0, \beta_{\mathcal{D}} - \beta_0)\|_2 \rightarrow 0$ implies that $\mathcal{D} \rightarrow 1$. This gives the desired rate of consistency which we show in Section 3.7.

3.5 Variance Bounds using Exchangeable Pairs

In this Section we state the lemmata which are required to show that the norm of the gradient of the log pseudo-likelihood is bounded at the true parameters.

Lemma 3.2 (Variance Bound 1). *It holds that*

$$\mathbb{E}_{\theta_0, \beta_0} \left[\left(\sum_{i=1}^n y_i m_i(\mathbf{y}) - m_i(\mathbf{y}) \tanh(\beta_0 m_i(\mathbf{y}) + \theta_0^\top \mathbf{x}_i) \right)^2 \right] \leq (12 + 4B)n.$$

Proof. We use the powerful technique of exchangeable pairs as introduced by Chatterjee ([8]) and employed by Chatterjee and Dembo (see [9]). First it holds that $\frac{\partial m_i(\mathbf{y})}{\partial y_j} = A_{ij}$. Also note that since $\|A\|_\infty \leq 1$ it trivially follows that $|m_i(\mathbf{y})| \leq 1$ for all i and $\mathbf{y} \in \{-1, +1\}^n$. Set

$$Q(\mathbf{y}) := \sum_i (y_i - \tanh(\beta_0 m_i(\mathbf{y}) + \theta_0^\top \mathbf{x}_i)) m_i(\mathbf{y}), \quad (13)$$

hence we get

$$\frac{\partial Q(\mathbf{y})}{\partial y_j} = \sum_i \left(\mathbf{1}_{i=j} - \frac{\beta_0 A_{ij}}{\cosh^2(\beta_0 m_i(\mathbf{y}) + \theta_0^\top \mathbf{x}_i)} \right) m_i(\mathbf{y}) + \left(y_i - \tanh(\beta_0 m_i(\mathbf{y}) + \theta_0^\top \mathbf{x}_i) \right) \frac{\partial m_i(\mathbf{y})}{\partial y_j}. \quad (14)$$

We will bound the absolute value of each summand. First from above we can bound the second term as follows

$$\left| (y_i - \tanh(\beta_0 m_i(\mathbf{y}) + \theta_0^\top \mathbf{x}_i)) \frac{\partial m_i(\mathbf{y})}{\partial y_j} \right| \leq 2 |A_{ij}|. \quad (15)$$

Using the fact that $\frac{1}{\cosh^2(x)} \leq 1$ it also follows that

$$\left| \sum_i \left(\mathbf{1}_{i=j} - \frac{\beta_0 A_{ij}}{\cosh^2(\beta_0 m_i(\mathbf{y}) + \theta_0^\top \mathbf{x}_i)} \right) m_i(\mathbf{y}) \right| \leq |m_j(\mathbf{y})| + \sum_{i \neq j} \beta_0 |A_{ij} m_i(\mathbf{y})| \quad (16)$$

Using (15) and (16) it follows that $\left| \frac{\partial Q(\mathbf{y})}{\partial y_j} \right| \leq \sum_{i \neq j} |A_{ij}| (2 + \beta_0 |m_i(\mathbf{y})|) + |m_j(\mathbf{y})|$. Finally let $\mathbf{y}^j = (\mathbf{y}_{-j}, -1)$ and note that

$$|Q(\mathbf{y}) - Q(\mathbf{y}^j)| \leq 2 \cdot \left(\sum_{i \neq j} |A_{ij}| (2 + \beta_0 |m_i(\mathbf{y}_{-j}, w)|) + |m_j(\mathbf{y}_{-j}, w)| \right) \leq 2 \cdot \left(\sum_{i \neq j} |A_{ij}| (2+B) + 1 \right) \leq 6 + 2B, \quad (17)$$

where w is the argmax of $|m_i(\mathbf{y}_{-j}, w)|$ along the line with endpoints \mathbf{y}, \mathbf{y}^j (Taylor). In the last inequality we used that $\|A\|_\infty \leq 1$. We have all the ingredients to complete the proof. We first observe that

$$\sum_i \mathbb{E}_{\theta_0, \beta_0} [(y_i - \tanh(\beta_0 m_i(\mathbf{y}) + \theta_0^\top \mathbf{x}_i)) Q(\mathbf{y}^i) m_i(\mathbf{y})] = 0, \quad (18)$$

since

$$\begin{aligned} & \mathbb{E}_{\theta_0, \beta_0} [(y_i - \tanh(\beta_0 m_i(\mathbf{y}) + \theta_0^\top \mathbf{x}_i)) Q(\mathbf{y}^i) m_i(\mathbf{y})] = \\ & = \mathbb{E}_{\theta_0, \beta_0} [\mathbb{E}[(y_i - \tanh(\beta_0 m_i(\mathbf{y}) + \theta_0^\top \mathbf{x}_i)) Q(\mathbf{y}^i) m_i(\mathbf{y}) | \mathbf{y}_{-i}]] = 0. \end{aligned} \quad (19)$$

Therefore it follows

$$\begin{aligned} \mathbb{E}_{\theta_0, \beta_0} [Q^2(\mathbf{y})] &= \mathbb{E}_{\theta_0, \beta_0} \left[Q(\mathbf{y}) \cdot \left(\sum_i (y_i - \tanh(\beta_0 m_i(\mathbf{y}) + \theta_0^\top \mathbf{x}_i)) m_i(\mathbf{y}) \right) \right] \\ &= \mathbb{E}_{\theta_0, \beta_0} \left[\sum_i \left(Q(\mathbf{y}) (y_i - \tanh(\beta_0 m_i(\mathbf{y}) + \theta_0^\top \mathbf{x}_i)) m_i(\mathbf{y}) \right) \right] \\ &= \sum_i \mathbb{E}_{\theta_0, \beta_0} \left[(Q(\mathbf{y}) - Q(\mathbf{y}^i)) \cdot (y_i - \tanh(\beta_0 m_i(\mathbf{y}) + \theta_0^\top \mathbf{x}_i)) m_i(\mathbf{y}) \right] \\ &\leq \sum_i 2 \cdot (6 + 2B) = (12 + 4B)n. \end{aligned}$$

□

Lemma 3.3 (Variance Bound 2).

$$\mathbb{E}_{\theta_0, \beta_0} \left[\sum_{k=1}^d \left(\sum_{i=1}^n x_{i,k} y_i - x_{i,k} \tanh(\beta_0 m_i(\mathbf{y}) + \theta_0^\top \mathbf{x}_i) \right)^2 \right] \leq (4 + 4B) M^2 \cdot dn.$$

Proof. We use the powerful technique of exchangeable pairs as employed by Chatterjee and Dembo (see [9]). Note that since $\|A\|_\infty \leq 1$ it trivially follows that $|m_i(\mathbf{y})| \leq 1$ for all i and $\mathbf{y} \in \{-1, +1\}^n$. We fix a coordinate k and set

$$Q(\mathbf{y}) := \sum_i (y_i - \tanh(\beta_0 m_i(\mathbf{y}) + \theta_0^\top \mathbf{x}_i)) x_{i,k}, \quad (20)$$

hence we get $\frac{\partial Q(\mathbf{y})}{\partial y_j} = \sum_i \left(\mathbf{1}_{i=j} - \frac{\beta_0 A_{ij}}{\cosh^2(\beta_0 m_i(\mathbf{y}) + \theta_0^\top \mathbf{x}_i)} \right) x_{i,k}$. We will bound the term as follows

$$\left| \frac{\partial Q(\mathbf{y})}{\partial y_j} \right| \leq |x_{j,k}| + \sum_{i \neq j} \beta_0 |A_{ij} x_{i,k}|. \quad (21)$$

Finally let $\mathbf{y}^j = (\mathbf{y}_{-j}, -1)$ and note that

$$|Q(\mathbf{y}) - Q(\mathbf{y}^j)| \leq 2 \cdot \left(|x_{j,k}| + \sum_{i \neq j} \beta_0 |A_{ij} x_{i,k}| \right). \quad (22)$$

We have all the ingredients to complete the proof. We first observe that

$$\sum_i \mathbb{E}_{\theta_0, \beta_0} [(y_i - \tanh(\beta_0 m_i(\mathbf{y}) + \theta_0^\top \mathbf{x}_i)) Q(\mathbf{y}^i) x_{i,k}] = 0, \quad (23)$$

since

$$\begin{aligned} & \mathbb{E}_{\theta_0, \beta_0} [(y_i - \tanh(\beta_0 m_i(\mathbf{y}) + \theta_0^\top \mathbf{x}_i)) Q(\mathbf{y}^i) x_{i,k}] = \\ & = \mathbb{E}_{\theta_0, \beta_0} [\mathbb{E}[(y_i - \tanh(\beta_0 m_i(\mathbf{y}) + \theta_0^\top \mathbf{x}_i)) Q(\mathbf{y}^i) x_{i,k} | \mathbf{y}_{-i}]] = 0. \end{aligned} \quad (24)$$

Therefore it follows

$$\begin{aligned} \mathbb{E}_{\theta_0, \beta_0} [Q^2(\mathbf{y})] &= \mathbb{E}_{\theta_0, \beta_0} \left[Q(\mathbf{y}) \cdot \left(\sum_i (y_i - \tanh(\beta_0 m_i(\mathbf{y}) + \theta_0^\top \mathbf{x}_i)) x_{i,k} \right) \right] \\ &= \mathbb{E}_{\theta_0, \beta_0} \left[\sum_i \left(Q(\mathbf{y}) (y_i - \tanh(\beta_0 m_i(\mathbf{y}) + \theta_0^\top \mathbf{x}_i)) x_{i,k} \right) \right] \\ &= \sum_i \mathbb{E}_{\theta_0, \beta_0} \left[(Q(\mathbf{y}) - Q(\mathbf{y}^i)) \cdot (y_i - \tanh(\beta_0 m_i(\mathbf{y}) + \theta_0^\top \mathbf{x}_i)) x_{i,k} \right] \\ &\leq \sum_i 4 \cdot (x_{i,k}^2 + |x_{i,k}| \sum_{j \neq i} \beta_0 |x_{j,k}| A_{ij}) \\ &\leq 4 \sum_i |x_{i,k}|^2 + B |x_{i,k}| \max_j |y_{xj, k}|, \end{aligned}$$

and the claim follows by summing over all the coordinates. \square

3.6 Strong Concavity of Maximum Pseudolikelihood

In this Section, we set $F = I - X(X^\top X)^{-1} X^\top$ ³

Lemma 3.4 (Lower Bound on Smallest Eigenvalue of Hessian). *With probability $1 - o(1)$, $\lambda_{\min}(-H_{(\theta, \beta)}) \geq c$ for some constant $c > 0$ for all $(\theta, \beta) \in \mathbb{B}$.*

Proof. We have

$$-H := G := \begin{pmatrix} \frac{1}{n} X^\top X & \frac{1}{n} X^\top \mathbf{m} \\ \frac{1}{n} \mathbf{m}^\top X & \frac{1}{n} \|\mathbf{m}\|_2^2 \end{pmatrix}.$$

Recall our notation that $Q = \frac{1}{n} X^\top X$. By using the properties of Schur complement, we get that

$$\det(G - \lambda I) = \det(Q - \lambda I) \det\left(\frac{1}{n} \mathbf{m}^\top \left(I - \frac{1}{n} X (Q - \lambda I)^{-1} X^\top\right) \mathbf{m} - \lambda\right). \quad (25)$$

Therefore the minimum eigenvalue of G is at least a positive constant as long as the minimum eigenvalues of

$$Q \text{ and } \frac{1}{n} \mathbf{m}^\top \left(I - \frac{1}{n} X Q^{-1} X^\top\right) \mathbf{m}$$

are positive constants independent of n . Recall from our assumptions in Table 1, we have that $\lambda_{\min}(Q) \geq c_1$ where c_1 is a positive constant independent of n . Hence, it remains to show that

$$\lambda_{\min} \left(\frac{1}{n} \mathbf{m}^\top \left(I - \frac{1}{n} X Q^{-1} X^\top\right) \mathbf{m} \right) \geq c_2$$

for a positive constant c_2 with high probability. Our desired statement can be restated as

$$\|F \mathbf{m}\|_2^2 \geq c_2 n.$$

³ $I - F$ is called hat-matrix or projection matrix.

Note that F has the property that $F^2 = F$ (i.e. is idempotent) and hence all the eigenvalues of F are 0, 1 (since is of rank $n - d$, it has d eigenvalues zero and $n - d$ eigenvalues one). Moreover, from the sub-multiplicativity of the spectral norm, it holds that

$$\|FA\|_2 \leq \|F\|_2 \times \|A\|_2 \leq 1. \quad (26)$$

We also have that $\|FA\|_F^2$ is $\Omega(n)$. This is because $\sigma_i(FA) \geq \sigma_{d+i}(A)\sigma_{n-d-i+1}(F) = \sigma_{d+i}(A)$ for $1 \leq i \leq n - d$ ($\sigma_i(G)$ denotes the i -th largest eigenvalue of G). Since $\sigma_{\max}(A) \leq 1$ it follows that $\|FA\|_F^2 \geq \|A\|_F^2 - d$.

Below, we provide an important lemma that will be used to show that $\|F\mathbf{m}\|_2^2$ is $\Omega(n)$ with high probability.

Lemma 3.5. *Let W be an $n \times n$ matrix. Fix a pair of indices i, j . It holds that*

$$\mathbb{E}_{\theta_0, \beta_0} [(W_i \mathbf{y})^2 | \mathbf{y}_{-j}] \geq \frac{e^{-(B+d \cdot M \cdot \Theta)}}{2} W_{ij}^2,$$

where W_i is the i -th row of W .

Proof. For any realization of \mathbf{y} , consider the two summands $\left(\sum_{t \neq j} W_{it} y_t + W_{ij}\right)^2$ and $\left(\sum_{t \neq j} W_{it} y_t - W_{ij}\right)^2$. It is clear that if both $\sum_{t \neq j} W_{it} y_t, W_{ij}$ have same sign then the first term is at least W_{ij}^2 and if they have opposite sign then the second term is at least W_{ij}^2 . Additionally, at all times both terms are non-negative.

Moreover from above we have

$$\begin{aligned} \mathbb{E}_{\theta_0, \beta_0} [(W \mathbf{y})_i^2 | \mathbf{y}_{-j}] &\geq W_{ij}^2 \cdot \min(\Pr[y_j = +1 | \mathbf{y}_{-j}], \Pr[y_j = -1 | \mathbf{y}_{-j}]) \\ &\geq \frac{W_{ij}^2}{2} \exp\left(-\left|2\theta_0^\top \mathbf{x}_j + \beta_0 \sum_{t \neq j} A_{jt} y_j\right|\right) \\ &\geq \frac{W_{ij}^2}{2} \exp(-\Theta M d - B) \end{aligned}$$

where the inequality before the last holds because of (3) and last inequality by assumption. \square

Process to select indices: Let $h : [n] \rightarrow [n]$ a function to be defined later. We define the following iterative process. Assume that at time $t = 0$ we start with matrix $n \times n$, $W^1 = W$. At time step t we choose from W^t the row with maximum l2 norm (let i_t the index of that row, ties broken arbitrarily) and also let $j_t = \arg\max_j |W_{i_t, j}^t|$ (again ties broken arbitrarily). We set $h(i_t) = j_t$ and W^{t+1} is W^t by setting zeros the entries of i_t -th row and column j_t -th. We run the process above for n steps and define the **bijection** function h . Below we prove the following lemma.

Lemma 3.6. *Assume that $\|FA\|_F^2 \geq cn$ and $\|FA\|_\infty \leq 1$. We run the process described above on FA and get the function h . It holds that*

$$\sum_i |(FA)_{ih(i)}|^2 \geq \frac{c^4 n}{16}.$$

Proof. Let A^i be the i -th column of A . It is clear that $F \cdot A^i$ corresponds to the i -th column of FA . Since F has eigenvalues 0,1, we get that $\|FA^i\|_2^2 \leq \|A^i\|_2^2 \leq \|A^i\|_\infty \|A^i\|_1 \leq 1$ (a) since A is symmetric and hence $\|A\|_1 = \|A\|_\infty \leq 1$. Above we used the fact that for any n -dimensional vector \mathbf{u} , from Holder's inequality,

$$\|\mathbf{u}\|_2^2 \leq \|\mathbf{u}\|_\infty \|\mathbf{u}\|_1.$$

Also let $(FA)_i$ be the i -th row of FA . Since $\|FA\|_\infty \leq 1$ it holds that $\|(FA)_i\|_2^2 \leq \|FA\|_\infty \leq 1$ (b).

We run the process described in the previous paragraph on FA and let W^t be the matrix at time t (with $W^1 = FA$). Let i_1, \dots, i_n be the ordering of the indices of the rows the process chose. It is clear that at every step we remove a column and a row from the matrix, the frobenius norm is decreased by at most two (using facts (a) and (b)).

Hence by the definition of the process we have that

$$\|W_{i_t}^t\|_2^2 \geq \frac{\|A\|_F^2 - 2(t-1)}{n-t-1} \geq \frac{cn - 2t + 2}{n} = c - \frac{2(t-1)}{n}.$$

We set $T = \lfloor cn/2 \rfloor$ and we get

$$\sum_{t=1}^n \|W_{i_t}^t\|_2^2 \geq \sum_{t=1}^T \|W_{i_t}^t\|_2^2 \tag{27}$$

$$\geq cT - \frac{2(T-1)T}{2n} \geq cT - \frac{T^2}{n} \approx \frac{c^2 n}{4} \tag{28}$$

Therefore we have that (observe that $1 \geq \|(FA)_j\|_1 \geq \|W_j^t\|_1$ for all $t, j \in [n]$, i.e., the ℓ_1 norm of each row does not increase during the process and same is true for ℓ_∞)

$$\frac{c^2 n}{4} \leq \sum_{t=1}^n \|W_{i_t}^t\|_2^2 \leq \sum_{i=1}^n \|W_{i_t}^t\|_\infty \|W_{i_t}^t\|_1 \leq \sum_{i=1}^n \|W_{i_t}^t\|_\infty. \tag{29}$$

$$\implies \frac{c^4 n^2}{16} \leq \left(\sum_{t=1}^n \|W_{i_t}^t\|_2^2 \right)^2 \leq n \sum_{t=1}^n \|W_{i_t}^t\|_\infty^2 \tag{30}$$

$$\implies \sum_{t=1}^n \|W_{i_t}^t\|_\infty^2 \geq \frac{c^4 n}{16}. \tag{31}$$

Finally, it holds that $\sum_{t=1}^n \|(FA)_t\|_\infty^2 \geq \sum_{t=1}^n \|W_{i_t}^t\|_\infty^2$ and the claim follows. \square

Corollary 3.1.

$$\sum_i \mathbb{E}_{\theta_0, \beta_0} [(F\mathbf{m})_i^2 | \mathbf{y}_{-h(i)}] \geq Cn,$$

for some positive constant C .

Proof. It holds by Lemmas 3.5 and 3.6 by choosing for each i , index $h(i)$. \square

Lemma 3.7 (Bounding the variance). *It holds that*

$$\mathbb{E}_{\theta_0, \beta_0} \left[\left(\sum_{i=1}^n (F\mathbf{m})_i^2 - \sum_i \mathbb{E}_{\theta_0, \beta_0} [(Fm)_i^2 | \mathbf{y}_{-h(i)}] \right)^2 \right] \leq 48n + 16Bn.$$

Proof. For each i , we expand the term $\mathbb{E}_{\theta_0, \beta_0} [(F\mathbf{m})_i^2 | \mathbf{y}_{-h(i)}]$ and we get $\mathbb{E}_{\theta_0, \beta_0} [(F\mathbf{m})_i^2 | \mathbf{y}_{-h(i)}] = (\sum_{j \neq h(i)} (FA)_{ij} y_j)^2 + (FA)_{ih(i)}^2 + 2(\sum_{j \neq h(i)} (FA)_{ij} y_j) \tanh(\beta_0 m_{h(i)}(\mathbf{y}) + \theta_0^\top \mathbf{x}_{h(i)})$. We set $z_{it}(\mathbf{y}) = 2(\sum_{j \neq t} (FA)_{ij} y_j)$ and we get that the expectation we need to bound is equal to

$$\mathbb{E}_{\theta_0, \beta_0} \left[\left(\sum_i z_{ih(i)}(\mathbf{y}) y_{h(i)} - z_{ih(i)}(\mathbf{y}) \tanh(\beta_0 m_{h(i)}(\mathbf{y}) + \theta_0^\top \mathbf{x}_{h(i)}) \right)^2 \right].$$

First it holds that $\frac{\partial z_{it}}{y_j} = 2(FA)_{ij}$ and $\frac{\partial z_{it}}{y_t} = 0$. Also if $\|FA\|_\infty \leq 1$ it trivially holds that $|z_{it}| \leq 2$. We set

$$Q(\mathbf{y}) = \sum_i (y_{h(i)} - \tanh(\beta_0 m_{h(i)}(\mathbf{y}) + \theta_0^\top \mathbf{x}_{h(i)})) z_{ih(i)}(\mathbf{y}),$$

hence we get that

$$\frac{\partial Q(\mathbf{y})}{\partial y_j} = \sum_i \left(\mathbf{1}_{h(i)=j} - \frac{\beta_0 A_{h(i)j}}{\cosh^2(\beta_0 m_{h(i)}(\mathbf{y}) + \theta_0^\top \mathbf{x}_{h(i)})} \right) z_{ih(i)}(\mathbf{y}) \quad (32)$$

$$+ \left(y_{h(i)} - \tanh(\beta_0 m_{h(i)}(\mathbf{y}) + \theta_0^\top \mathbf{x}_{h(i)}) \right) \frac{\partial z_{ih(i)}(\mathbf{y})}{\partial y_j}. \quad (33)$$

We will bound the absolute value of each summand. First from above we can bound the second term as follows

$$\left| \left(y_{h(i)} - \tanh(\beta_0 m_{h(i)}(\mathbf{y}) + \theta_0^\top \mathbf{x}_{h(i)}) \right) \frac{\partial z_{ih(i)}(\mathbf{y})}{\partial y_j} \right| \leq 4 |(FA)_{ij}|. \quad (34)$$

Using the fact that $\frac{1}{\cosh^2(x)} \leq 1$ it also follows that

$$\left| \sum_i \left(\mathbf{1}_{h(i)=j} - \frac{\beta_0 A_{h(i)j}}{\cosh^2(\beta_0 m_{h(i)}(\mathbf{y}) + \theta_0^\top \mathbf{x}_{h(i)})} \right) z_{ih(i)}(\mathbf{y}) \right| \leq \sum_i \mathbf{1}_{h(i)=j} |z_{ih(i)}(\mathbf{y})| + \sum_{i \neq j} |\beta_0 A_{h(i)j} z_{ih(i)}(\mathbf{y})|, \quad (35)$$

which is at most $2 \sum_i \mathbf{1}_{h(i)=j} + 2B$.

Using (34) and (35) it follows that $\left| \frac{\partial Q(\mathbf{y})}{\partial y_j} \right| \leq 4 \sum_{i, h(i) \neq j} |(FA)_{ij}| + 2 \sum_i \mathbf{1}_{h(i)=j} + 2B$. We have all the ingredients to complete the proof. Let $\mathbf{y}^j = (\mathbf{y}_{-j}, -1)$. We first observe that

$$\sum_i \mathbb{E}_{\theta_0, \beta_0} [(y_{h(i)} - \tanh(\beta_0 m_{h(i)}(\mathbf{y}) + \theta_0^\top \mathbf{x}_{h(i)})) Q(\mathbf{y}^{\mathbf{h}(i)}) z_{ih(i)}(\mathbf{y})] = 0, \quad (36)$$

since

$$\begin{aligned} & \mathbb{E}_{\theta_0, \beta_0} [(y_{h(i)} - \tanh(\beta_0 m_{h(i)}(\mathbf{y}) + \theta_0^\top \mathbf{x}_{h(i)})) Q(\mathbf{y}^{\mathbf{h}(i)}) z_{ih(i)}(\mathbf{y})] = \\ & = \mathbb{E}_{\theta_0, \beta_0} [\mathbb{E}[(y_{h(i)} - \tanh(\beta_0 m_{h(i)}(\mathbf{y}) + \theta_0^\top \mathbf{x}_{h(i)})) Q(\mathbf{y}^{\mathbf{h}(i)}) z_{ih(i)}(\mathbf{y}) | \mathbf{y}_{-h(i)}]] = 0. \end{aligned} \quad (37)$$

Therefore it follows

$$\mathbb{E}_{\theta_0, \beta_0} [Q^2(\mathbf{y})] = \mathbb{E}_{\theta_0, \beta_0} \left[Q(\mathbf{y}) \cdot \left(\sum_i (y_{h(i)} - \tanh(\beta_0 m_{h(i)}(\mathbf{y}) + \theta_0^\top \mathbf{x}_{h(i)})) z_{ih(i)}(\mathbf{y}) \right) \right]$$

$$\begin{aligned}
&= \mathbb{E}_{\theta_0, \beta_0} \left[\sum_i \left(Q(\mathbf{y})(y_{h(i)} - \tanh(\beta_0 m_i(\mathbf{y}) + \theta_0^\top \mathbf{x}_i)) z_{ih(i)}(\mathbf{y}) \right) \right] \\
&= \sum_i \mathbb{E}_{\theta_0, \beta_0} \left[(Q(\mathbf{y}) - Q(\mathbf{y}^{h(i)})) \cdot (y_{h(i)} - \tanh(\beta_0 m_{h(i)}(\mathbf{y}) + \theta_0^\top \mathbf{x}_{h(i)})) z_{ih(i)}(\mathbf{y}) \right] \\
&= \sum_i \mathbb{E}_{\theta_0, \beta_0} \left[(Q(\mathbf{y}) - Q(\mathbf{y}^i)) \cdot (y_i - \tanh(\beta_0 m_i(\mathbf{y}) + \theta_0^\top \mathbf{x}_i)) z_{h^{-1}(i)i}(\mathbf{y}) \right] \\
&\leq 8 \sum_i \left[\sum_{t, h(t) \neq i} 4|(FA)_{ti}| + 2 \sum_t \mathbf{1}_{h(t)=i} + 2B \right] \\
&\leq 48n + 16Bn,
\end{aligned}$$

where we also used the fact that $\sum_i \sum_j (FA)_{ij} \leq \sum_i \|FA\|_\infty \leq n$. \square

Remark 3.6. *The proof of Lemma 3.6 depends on the assumption that $\|FA\|_\infty \leq c'_\infty$ for some positive constant c'_∞ . This assumption is not necessary for the proof to go through and below we will argue about removing this assumption.*

Lemma 3.8. *It holds that $\sum_{i,j} |F_{ij}| \leq n(d+1)$.*

Proof. Since F has rank $n-d$ and has eigenvalues 0,1 it holds that $F = I - \sum_{i=1}^d \mathbf{v}^i \mathbf{v}^{i\top}$ with $\|\mathbf{v}^i\|_2 = 1$. We choose vector \mathbf{x} so that $x_j = 1$ if $v_j^i \geq 0$ and $x_j = -1$ otherwise. It holds that $(\sum_j |v_j^i|)^2 = \mathbf{x}^\top \mathbf{v}^i \mathbf{v}^{i\top} \mathbf{x} \leq \|\mathbf{x}\|_2^2 = n$. Thus we get that $\sum_j \sum_t |v_j^i| \cdot |v_t^i| \leq n$. Therefore we get that

$$\sum_{i,j} |F_{i,j}| \leq n + \sum_{i=1}^d n = n(d+1).$$

\square

Corollary 3.2. *We can drop the assumption $\|FA\|_\infty$ is $O(1)$.*

Proof. From Lemma 3.8 we get that for every $\epsilon > 0$, there exists at most nd/ϵ rows of FA that have ℓ_1 norm larger than ϵ (by Markov's inequality), hence there exist $n(1-d/\epsilon)$ rows that have ℓ_1 norm at most ϵ . Moreover, since $\|FA\|_2 \leq 1$ (26) it holds that each row j of FA (denote by $(FA)_j$) has $\|(FA)_j\|_2^2 \leq 1$. The trick is to remove from matrix FA all the rows that have ℓ_1 norm larger than ϵ and we reduce the Frobenius norm squared of FA by nd/ϵ . By choosing ϵ to be $\frac{2nd}{\|FA\|_F^2}$, the resulting matrix (after removing the "overloaded" rows) has still Frobenius norm squared $\Omega(n)$ and moreover all the rows have ℓ_1 norm at most ϵ (which is $\Theta(1)$). The claim follows since the rows we are removing cannot increase the expression $\sum_i (FA\mathbf{y})_i^2$ (sum of squares). \square

Finally we are ready to complete the proof by showing that $\|F\mathbf{m}\|_2^2 \geq c_2 n$. \square

3.7 Completing the Proof

With the above results in hand, we can prove the main result of this Section, Theorem 3.1.

Proof of Theorem 3.1: First it holds that:

$$\mathbb{E}_{\theta_0, \beta_0} [\|\nabla LPL(\theta_0, \beta_0)\|_2^2] \quad (38)$$

$$= \sum_{k=1}^d \mathbb{E}_{\theta_0, \beta_0} \left[\frac{1}{n} \sum_{i=1}^n \left(x_{i,k} y_i - x_{i,k} \tanh(\beta_0 m_i(\mathbf{y}) + \theta_0^\top \mathbf{x}_i) \right) \right]^2 \quad (39)$$

$$+ \mathbb{E}_{\theta_0, \beta_0} \left[\frac{1}{n} \sum_{i=1}^n \left(y_i m_i(\mathbf{y}) - m_i(\mathbf{y}) \tanh(\beta_0 m_i(\mathbf{y}) + \theta_0^\top \mathbf{x}_i) \right) \right]^2. \quad (40)$$

From Lemmas 3.2 and 3.3 we have that

$$\mathbb{E}_{\theta_0, \beta_0} [\|\nabla LPL(\theta_0, \beta_0)\|_2^2] \leq \frac{c}{n} \quad (41)$$

for some constant c . By Markov's inequality, we get that

$$\Pr \left[\|\nabla LPL(\theta_0, \beta_0)\|_2^2 \leq \frac{c}{n\delta} \right] \geq 1 - \delta. \quad (42)$$

for any constant δ . Next, we have from Lemma 3.4 that, $\min_{(\theta, \beta) \in \mathbb{B}} \lambda_{\min}(-H_{(\theta, \beta)}) \geq C$ for some constant C independent of n . Plugging into Lemma 3.1, we get that

$$\|(\theta_{\mathcal{D}} - \theta_0, \beta_{\mathcal{D}} - \beta_0)\|_2 = \mathcal{D} \left\| (\hat{\theta}_{MPL} - \theta_0, \hat{\beta}_{MPL} - \beta_0) \right\|_2 \leq \frac{\|\nabla LPL(\theta_0, \beta_0)\|_2}{\min_{(\theta, \beta) \in \mathbb{B}} \lambda_{\min}(-H_{(\theta, \beta)})} \quad (43)$$

Now we have from the above that $\|(\theta_{\mathcal{D}} - \theta_0, \beta_{\mathcal{D}} - \beta_0)\|_2 \rightarrow 0$ as $n \rightarrow \infty$ (is of order $\frac{1}{\sqrt{n}}$). Also note that any point on the boundary of \mathbb{B} has a fixed distance to (θ_0, β_0) since it lies in the interior. Hence $\|(\theta_{\mathcal{D}} - \theta_0, \beta_{\mathcal{D}} - \beta_0)\|_2 \rightarrow 0$ implies that $\mathcal{D} \rightarrow 1$ which in turn implies that $\mathcal{D} \geq 1/2$ for sufficiently large n . Hence

$$(43) \implies \left\| (\hat{\theta}_{MPL} - \theta_0, \hat{\beta}_{MPL} - \beta_0) \right\|_2 \leq \frac{2 \|\nabla LPL(\theta_0, \beta_0)\|_2}{\min_{(\theta, \beta) \in \mathbb{B}} \lambda_{\min}(-H_{(\theta, \beta)})} \quad (44)$$

$$\leq O_d \left(\frac{1}{\sqrt{n}} \right) \quad (45)$$

with probability $\geq 1 - \delta$. □

4 Linear Regression with Dependent Observations

In this section we focus on linear regression under weakly dependent errors. As opposed to Logistic regression, in linear regression the log-likelihood is computationally tractable.

4.1 Our Model

We recall the model for dependent observations we consider. In our setting, we have that the errors $\epsilon_i = y_i - \theta x_i$ are distributed according to a Gaussian graphical model. Since each ϵ_i is zero mean, we have that $\mu = 0$ in our case. Also, similar to the logistic regression setting, we will assume that we have complete knowledge of the graph structure up to a scaling factor. That is, $(\Sigma)^{-1} = \beta A + D$ where the matrix A is a known symmetric matrix with $A_{ii} = 0$ and $D = [d_1 \dots d_n]^T$ is a known diagonal matrix with positive entries. Hence the probability distribution of the observations is:

$$\Pr[\mathbf{y} = \mathbf{a}] = \frac{\exp\left(-\frac{1}{2}(\mathbf{a} - \theta^\top \mathbf{x})^\top \Sigma^{-1}(\mathbf{a} - \theta^\top \mathbf{x})\right)}{(2\pi)^{n/2} \det(\Sigma)^{1/2}} \quad (46)$$

This section is devoted to showing that the Maximum Likelihood Estimator (MLE) under appropriate (over)re-parametrization - (the new parameter vector will be (θ, β, κ) where θ, β remain same after reparametrization and $\kappa = \beta \cdot \theta$) - is \sqrt{n} consistent in our linear regression model. We set $\mathbb{B} = [-\Theta, \Theta]^d \times [-B, B] \times [-\Theta B, \Theta B]^d$ (the set of which the parameters should be in the interior and B is defined in Table 1). Formally we prove the following theorem.

Theorem 4.1 (Main Linear). *Assume that*

1. Feature matrix X^4 with covariance matrix $Q = \frac{1}{n} X^\top X$ having $\lambda_{\min}(Q), \lambda_{\max}(Q)$ as positive constants.
2. $\|A\|_2$ is $\Theta(1)$, $\|A\|_F^2$ is $\Omega(n)$ and $\lambda_{\min}\left(\frac{1}{n} X^\top A^\top (I - DX(X^\top D^2 X)^{-1} X^\top D) AX\right)$ is $\Theta(1)$.

We show that the Maximum Log-likelihood Estimate (MLE) $(\hat{\theta}, \hat{\beta}, \hat{\kappa})$ is $O_d\left(\sqrt{\frac{1}{n}}\right)$ consistent as long as the true parameter vector $(\theta_0, \beta_0, \kappa_0) \in \mathbb{B}$ (in the interior of \mathbb{B}), i.e., for each $\delta > 0$ and n sufficiently large, $\left\|(\hat{\theta}, \hat{\beta}, \hat{\kappa}) - (\theta_0, \beta_0, \kappa_0)\right\|_2$ is $O_d\left(\sqrt{\frac{1}{n}}\right)$ with probability $1 - \delta$. Moreover, we can compute a vector $(\tilde{\theta}, \tilde{\beta})$ with $\left\|(\hat{\theta}, \hat{\beta}) - (\tilde{\theta}, \tilde{\beta})\right\|_2$ to be $O_d\left(\sqrt{\frac{1}{n}}\right)$ in $O(\ln n)$ iterations of projected gradient descent⁵ with probability $1 - o(1)$.

Corollary 4.1 (Application to Sherrington-Kirkpatrick (SK) model). *In the Sherrington-Kirkpatrick model [37], we have that $A_{ij} = \frac{g_{ij}}{\sqrt{n}}$ for $i < j$, where $g_{ij} \sim \mathcal{N}(0, 1)$ and $A_{ji} = A_{ij}, A_{ii} = 0$. From Lemma 4.6, it follows that A satisfies the assumptions of our main theorem, so we can infer β, θ (with a \sqrt{n} rate of consistency).*

Technically, we will reparametrize the log-likelihood function in such a way that the new parameter vector is not high-dimensional and the resulting log-likelihood becomes strongly convex. The reparametrization and the equations of log-likelihood, its gradient and Hessian can be found in Section 4.2. We will follow the same high level ideas as in the logistic regression. Under assumptions on $A, \beta, \theta, (\beta A + D)^{-1}, AX$ that are summarized in Table 1 we proceed as follows:

- We prove concentration results for the gradient of the reparametrized log-likelihood, see Lemmas 4.3, 4.1 and 4.2 in Section 4.3.1.
- We prove that the minimum eigenvalue of the negative Hessian of reparametrized log-likelihood is large enough, see Lemma 4.4 in Section 4.3.2.

Below we provide some important definitions.

⁴ \mathbf{x}_i can be subgaussian.

⁵Each iteration is polynomial time computable

4.2 Our Reparametrization and Log-likelihood

It is not hard to see that the negative log-likelihood is not convex with respect to the parameter vector $(\theta, \beta) \in \mathbb{R}^{d+1}$, for the linear regression model with dependent errors. Nevertheless, we can reparametrize the log-likelihood in such a way to make it convex. The classic way to do it sets $T := \Sigma^{-1}\nu = \Sigma^{-1}\mu$ (for a gaussian $\mathcal{N}(\mu, \Sigma)$). However, this creates a parameter vector of dimension $\Omega(n)$. It is crucial that after the reparametrization the dimensionality of the new parameter vector is $O(d)$ and not $\Omega(n)$ (for concentration purposes, see remark 4.1). Hence, we take a different route here.

We set $\kappa := \beta \cdot \theta \in \mathbb{R}^d$ and define our parametric vector to be $(\theta, \beta, \kappa) \in \mathbb{R}^{2d+1}$ (parameters θ, β remain the same and we introduce vector κ). The vector (θ, β, κ) is $(2d+1)$ -dimensional. Our reparameterization is an overparameterization which helps us achieve convexity of the negative log-likelihood function.

The negative log-likelihood is given by the following:

$$-LL = \frac{1}{2}(\mathbf{y} - X\theta)^\top (\beta A + D)(\mathbf{y} - X\theta) + \log \int_{\mathbb{R}^n} \exp\left(-\frac{1}{2}(\mathbf{z} - X\theta)^\top (\beta A + D)(\mathbf{z} - X\theta)\right) dz \quad (47)$$

$$= \frac{1}{2}\mathbf{y}^\top (\beta A + D)\mathbf{y} - \mathbf{y}^\top AX\kappa - \mathbf{y}^\top DX\theta + \log \int_{\mathbb{R}^n} \exp\left(-\frac{1}{2}\mathbf{z}^\top (\beta A + D)\mathbf{z} + \mathbf{z}^\top AX\kappa + \mathbf{z}^\top DX\theta\right) dz \quad (48)$$

The negative gradient of the log-likelihood is given below:

$$-\nabla_\theta LL(\theta, \beta, \kappa) = -\mathbf{y}^\top DX + \frac{\int_{\mathbb{R}^n} \mathbf{z}^\top DX \exp\left(-\frac{1}{2}\mathbf{z}^\top (\beta A + D)\mathbf{z} + \mathbf{z}^\top AX\kappa + \mathbf{z}^\top DX\theta\right) dz}{\int_{\mathbb{R}^n} \exp\left(-\frac{1}{2}\mathbf{z}^\top (\beta A + D)\mathbf{z} + \mathbf{z}^\top AX\kappa + \mathbf{z}^\top DX\theta\right) dz} \quad (49)$$

$$= -\mathbf{y}^\top DX + E_{\mathbf{z} \sim \mathcal{N}((\beta A + D)^{-1}(AX\kappa + DX\theta), (\beta A + D)^{-1})} \left[\mathbf{z}^\top \right] DX \quad (50)$$

$$-\nabla_\beta LL(\theta, \beta, \kappa) = \frac{1}{2}\mathbf{y}^\top A\mathbf{y} + \frac{\int_{\mathbb{R}^n} -\frac{1}{2}\mathbf{z}^\top A\mathbf{z} \exp\left(-\frac{1}{2}\mathbf{z}^\top (\beta A + D)\mathbf{z} + \mathbf{z}^\top AX\kappa + \mathbf{z}^\top DX\theta\right) dz}{\int_{\mathbb{R}^n} \exp\left(-\frac{1}{2}\mathbf{z}^\top (\beta A + D)\mathbf{z} + \mathbf{z}^\top AX\kappa + \mathbf{z}^\top DX\theta\right) dz} \quad (51)$$

$$= \frac{1}{2}\mathbf{y}^\top A\mathbf{y} - E_{\mathbf{z} \sim \mathcal{N}((\beta A + D)^{-1}(AX\kappa + DX\theta), (\beta A + D)^{-1})} \left[\frac{1}{2}\mathbf{z}^\top A\mathbf{z} \right] \quad (52)$$

$$-\nabla_\kappa LL(\theta, \beta, \kappa) = -\mathbf{y}^\top AX + \frac{\int_{\mathbb{R}^n} \mathbf{z}^\top AX \exp\left(-\frac{1}{2}\mathbf{z}^\top (\beta A + D)\mathbf{z} + \mathbf{z}^\top AX\kappa + \mathbf{z}^\top DX\theta\right) dz}{\int_{\mathbb{R}^n} \exp\left(-\frac{1}{2}\mathbf{z}^\top (\beta A + D)\mathbf{z} + \mathbf{z}^\top AX\kappa + \mathbf{z}^\top DX\theta\right) dz} \quad (53)$$

$$= -\mathbf{y}^\top AX + E_{\mathbf{z} \sim \mathcal{N}((\beta A + D)^{-1}(AX\kappa + DX\theta), (\beta A + D)^{-1})} \left[\mathbf{z}^\top \right] AX. \quad (54)$$

The negative hessian of the log-likelihood is given below (it is of size $(2d+1) \times (2d+1)$):

$$-H := -\nabla^2 LL = Cov_{\mathbf{z} \sim \mathcal{N}((\beta A + D)^{-1}(AX\kappa + DX\theta), (\beta A + D)^{-1})} \left[\begin{pmatrix} -\frac{1}{2}\mathbf{z}^\top A\mathbf{z} \\ X^\top D\mathbf{z} \\ X^\top A\mathbf{z} \end{pmatrix}, \begin{pmatrix} -\frac{1}{2}\mathbf{z}^\top A\mathbf{z} \\ X^\top D\mathbf{z} \\ X^\top A\mathbf{z} \end{pmatrix} \right]. \quad (55)$$

4.3 Consistency of Likelihood

Let $(\theta_0, \beta_0, \kappa_0)$ be the true parameter (observe that $\kappa_0 = \beta_0 \cdot \theta_0$). We define $(\theta_t, \beta_t, \kappa_t) = (1-t)(\theta_0, \beta_0, \kappa_0) + t(\hat{\theta}, \hat{\beta}, \hat{\kappa})$ where $(\hat{\theta}, \hat{\beta}, \hat{\kappa})$ satisfies the first order conditions for LL (i.e., $\nabla LL(\hat{\theta}, \hat{\beta}, \hat{\kappa}) =$

$\mathbf{0}$)⁶ and set

$$g(t) := (\theta_0 - \hat{\theta}, \beta_0 - \hat{\beta}, \kappa_0 - \hat{\kappa})^\top \nabla LL(\theta_t, \beta_t, \kappa_t), \quad g'(t) = -(\theta_0 - \hat{\theta}, \beta_0 - \hat{\beta}, \kappa_0 - \hat{\kappa})^\top H_{(\theta_t, \beta_t, \kappa_t)}(\theta_0 - \hat{\theta}, \beta_0 - \hat{\beta}, \kappa_0 - \hat{\kappa}).$$

Let $\mathcal{D} \in [0, 1]$ be such that $(\theta_{\mathcal{D}}, \beta_{\mathcal{D}}, \kappa_{\mathcal{D}})$ intersects the boundary of set \mathbb{B} (if it does not intersect the boundary of \mathbb{B} then $\mathcal{D} = 1$). Since H is negative semidefinite (from analysis in equation 55) we have that $g'(t) \geq 0$ (**). It holds that

$$\begin{aligned} \left\| (\theta_0 - \hat{\theta}, \beta_0 - \hat{\beta}, \kappa_0 - \hat{\kappa}) \right\|_2 \cdot \left\| \nabla LL(\theta_0, \beta_0, \kappa_0) \right\|_2 &\geq |(\theta_0 - \hat{\theta}, \beta_0 - \hat{\beta}, \kappa_0 - \hat{\kappa})^\top \nabla LL(\theta_0, \beta_0, \kappa_0)| \\ &= |g(1) - g(0)| \\ &= \left| \int_0^1 g'(t) dt \right| \\ &\geq \left| \int_0^{\mathcal{D}} g'(t) dt \right| \text{ by (**)} \\ &\geq \mathcal{D} \min_{(\theta, \beta, \kappa) \in \tilde{\mathbb{B}}} \lambda_{\min}(-H_{(\theta, \beta, \kappa)}) \left\| (\theta_0 - \hat{\theta}, \beta_0 - \hat{\beta}, \kappa_0 - \hat{\kappa}) \right\|_2^2. \end{aligned}$$

The aforementioned inequalities indicate that we need a concentration result for $\left\| \frac{1}{n} \nabla LL(\theta_0, \beta_0, \kappa_0) \right\|_2$ and a lower bound on the minimum eigenvalue of $-\frac{1}{n} H$ for consistency of the MLE. As in the logistic regression case, combining with the observation that $D \rightarrow 1^7$ as $n \rightarrow \infty$ (i.e., $D \geq \frac{1}{2}$ for n sufficiently large) we get the desired rate of consistency.

4.3.1 Concentration results

We have

$$\begin{aligned} \mathbb{E}_{\theta_0, \beta_0} \left[\left\| \nabla LL(\theta_0, \beta_0, \beta_0 \cdot \theta_0) \right\|_2^2 \right] &= \mathbb{E}_{\theta_0, \beta_0} \left[\left\| \nabla_{\theta} LL(\theta_0, \beta_0, \beta_0 \cdot \theta_0) \right\|_2^2 \right] + \\ &\mathbb{E}_{\theta_0, \beta_0} \left[\left| \nabla_{\beta} LL(\theta_0, \beta_0, \beta_0 \cdot \theta_0) \right|^2 \right] + \mathbb{E}_{\theta_0, \beta_0} \left[\left\| \nabla_{\kappa} LL(\theta_0, \beta_0, \beta_0 \cdot \theta_0) \right\|_2^2 \right] \end{aligned} \quad (56)$$

We prove below concentration results for each term separately.

Lemma 4.1 (Bounding the 1st term).

$$\mathbb{E}_{\theta_0, \beta_0} \left[\left\| \nabla_{\theta} LL(\theta_0, \beta_0, \kappa_0) \right\|_2^2 \right] = \left\| (\beta_0 A + D)^{-1/2} DX \right\|_F^2 \leq \left\| (\beta_0 A + D)^{-1/2} D \right\|_2^2 \|X\|_F^2.$$

Proof. Assume that $y \sim \mathcal{N}(X\theta_0, (\beta_0 A + D)^{-1})$, it follows that $-\nabla_{\theta} LL(\theta_0, \beta_0, \kappa_0) = -\mathbf{y}^\top DX - \mathbb{E}_{\theta_0, \beta_0}[-\mathbf{y}^\top DX]$. It is clear that the vector $\mathbf{w} = (\beta_0 A + D)^{1/2}(\mathbf{y} - X\theta_0) \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.

It follows that

$$-\mathbf{y}^\top DX - \mathbb{E}_{\theta_0, \beta_0}[-\mathbf{y}^\top DX] = -((\beta_0 A + D)^{-1/2} \mathbf{w} + X\theta_0)^\top DX + \mathbb{E}_{\theta_0, \beta_0}[(\beta_0 A + D)^{-1/2} \mathbf{w} + X\theta_0]^\top DX \quad (57)$$

$$= -\mathbf{w}^\top (\beta_0 A + D)^{-1/2} DX. \quad (58)$$

⁶Observe that is *not* necessarily true that $\hat{\kappa} = \hat{\beta}\hat{\theta}$.

⁷This is true because $\|(\theta_{\mathcal{D}} - \theta_0, \beta_{\mathcal{D}} - \beta_0, \kappa_{\mathcal{D}} - \kappa_0)\|_2 \rightarrow 0$ as $n \rightarrow \infty$ by showing the promised concentration result and the lower bound).

It holds that

$$\begin{aligned}
\mathbb{E}_{\theta_0, \beta_0} [\|\nabla_{\theta} LL(\theta_0, \beta_0, \kappa_0)\|_2^2] &= \mathbb{E}_{\theta_0, \beta_0} \left[\left\| \mathbf{w}^{\top} (\beta_0 A + D)^{-1/2} DX \right\|_2^2 \right] \\
&= \text{tr}((\beta_0 A + D)^{-1/2} DX X^{\top} D (\beta_0 A + D)^{-1/2}) \\
&= \left\| (\beta_0 A + D)^{-1/2} DX \right\|_F^2
\end{aligned}$$

□

Lemma 4.2 (Bounding the 3rd term). *Similarly to Lemma 4.1 we get*

$$\mathbb{E}_{\theta_0, \beta_0} \left[\|\nabla_{\kappa} LL(\theta_0, \beta_0, \kappa_0)\|_2^2 \right] = \left\| (\beta_0 A + D)^{-1/2} AX \right\|_F^2 \leq \left\| (\beta_0 A + D)^{-1/2} A \right\|_2^2 \|X\|_F^2.$$

Lemma 4.3 (Bounding the 2nd term). *It holds that*

$$\begin{aligned}
\mathbb{E}_{\theta_0, \beta_0, \kappa_0} [\|\nabla_{\beta} LL(\theta_0, \beta_0, \kappa_0)\|_2^2] &= \mathbb{V}_{\mathbf{z} \sim \mathcal{N}(X\theta_0, (\beta_0 A + D)^{-1})} [\mathbf{z}^{\top} A \mathbf{z}] \\
&\leq 2 \left\| (\beta_0 A + D)^{-1/2} A (\beta_0 A + D)^{-1/2} \right\|_F^2 + 4d\Theta^2 \left\| (\beta_0 A + D)^{-1/2} AX \right\|_2^2
\end{aligned} \tag{59}$$

Proof. We follow the calculations of Lemma 2.3. It holds that

$$\begin{aligned}
\mathbb{E}_{\theta_0, \beta_0, \kappa_0} [\|\nabla_{\beta} LL(\theta_0, \beta_0, \kappa_0)\|_2^2] &\leq 2\text{tr}((A(\beta_0 A + D)^{-1})^2) + 4d\Theta^2 \left\| (\beta_0 A + D)^{-1/2} AX \right\|_2^2 \\
&\leq 2\text{tr}((\beta_0 A + D)^{-1/2} A (\beta_0 A + D)^{-1/2} (\beta_0 A + D)^{-1/2} A (\beta_0 A + D)^{-1/2}) \\
&\quad + 4d\Theta^2 \left\| (\beta_0 A + D)^{-1/2} AX \right\|_2^2 \\
&= 2 \left\| (\beta_0 A + D)^{-1/2} A (\beta_0 A + D)^{-1/2} \right\|_F^2 + 4d\Theta^2 \left\| (\beta_0 A + D)^{-1/2} AX \right\|_2^2
\end{aligned}$$

□

Remark 4.1. *We note the dependence on the dimension d for the bound in Lemma 4.3. This indicates how crucial it is that the dimensionality of the parameter vector does not scale with n .*

4.3.2 Lower bound on the minimum eigenvalue

In this section we provide a lower bound on the minimum eigenvalue of the negative Hessian of the log-likelihood. We need this bound for strong concavity of the log-likelihood.

Lemma 4.4 (Bounding the minimum eigenvalue). *Let $z \sim \mathcal{N}((\beta A + D)^{-1}(AX\kappa + DX\theta), (\beta A + D)^{-1})$. There exists a constant C such that*

$$\lambda_{\min} \left(\text{Cov} \left[\begin{pmatrix} -\frac{1}{2} \mathbf{z}^{\top} A \mathbf{z} \\ X^{\top} D \mathbf{z} \\ X^{\top} A \mathbf{z} \end{pmatrix}, \begin{pmatrix} -\frac{1}{2} \mathbf{z}^{\top} A \mathbf{z} \\ X^{\top} D \mathbf{z} \\ X^{\top} A \mathbf{z} \end{pmatrix} \right] \right) \geq \frac{Cn}{d}. \tag{60}$$

Proof. We set $\mu = (\beta A + D)^{-1}(AX\kappa + DX\theta)$, $\Sigma = (\beta A + D)^{-1}$ and $\mathbf{w} = \Sigma^{-1/2}(\mathbf{z} - \mu)$ (i.e., $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$) and we consider the vector

$$\mathbf{h} := \begin{pmatrix} -\frac{1}{2} \mathbf{w}^{\top} \Sigma^{1/2} A \Sigma^{1/2} \mathbf{w} + \mathbf{w}^{\top} \Sigma^{1/2} A \mu - \mathbb{E}[-\frac{1}{2} \mathbf{w}^{\top} \Sigma^{1/2} A \Sigma^{1/2} \mathbf{w}] \\ X^{\top} D \Sigma^{1/2} \mathbf{w} \\ X^{\top} A \Sigma^{1/2} \mathbf{w} \end{pmatrix}.$$

Let $\mathbf{v} := (v_1, \mathbf{v}_2, \mathbf{v}_3)$ be a column vector where $v_1 \in \mathbb{R}$, $\mathbf{v}_2, \mathbf{v}_3 \in \mathbb{R}^d$ so that $\|\mathbf{v}\|_2 = 1$. It follows that $-\mathbf{v}^\top H \mathbf{v} = \mathbb{E}[(\mathbf{v}^\top \mathbf{h})^2]$ where $-H$ is the negative hessian computed in (55).

From Lemma 2.3 we get that

$$\mathbb{E}[(\mathbf{v}^\top \mathbf{h})^2] = v_1^2 \left(\frac{1}{2} \text{tr} \left((\Sigma^{1/2} A \Sigma^{1/2})^2 \right) \right) + \left\| v_1 \mu^\top A \Sigma^{1/2} + \mathbf{v}_2^\top X^\top D \Sigma^{1/2} + \mathbf{v}_3^\top X^\top A \Sigma^{1/2} \right\|_2^2.$$

If a positive constant mass at least $\sqrt{\frac{\epsilon}{d}}$ is put on v_1 then the above term is at least $\frac{\epsilon}{2d} \|\Sigma^{1/2} A \Sigma^{1/2}\|_F^2$ which is $\Theta(\frac{n}{d})$. If not then the above term is at least $\|\mathbf{v}_2^\top X^\top D \Sigma^{1/2} + \mathbf{v}_3^\top X^\top A \Sigma^{1/2}\|_2^2 - O(n\sqrt{\epsilon})$ with $\|\mathbf{v}_2\|_2^2 + \|\mathbf{v}_3\|_2^2 \geq 1 - \epsilon$. We will prove a $\Theta(n)$ lower bound on the term $\|\mathbf{v}_2^\top X^\top D \Sigma^{1/2} + \mathbf{v}_3^\top X^\top A \Sigma^{1/2}\|_2^2 \geq \|\mathbf{v}_2^\top X^\top D + \mathbf{v}_3^\top X^\top A\|_2^2 \sigma_{\min}(\Sigma)$.

It suffices to bound the minimum eigenvalue of the following matrix:

$$M := \begin{pmatrix} \frac{1}{n} X^\top D^2 X & \frac{1}{n} X^\top D A X \\ \frac{1}{n} X^\top A D X & \frac{1}{n} X^\top A^2 X \end{pmatrix}.$$

By the Schur complement, we get that $\det(M - \lambda I) = \det(\frac{1}{n} X^\top D^2 X - \lambda I) \det(\frac{1}{n} X^\top A^2 X - \frac{1}{n} X^\top A D X (\frac{1}{n} X^\top D^2 X - \lambda I)^{-1} \frac{1}{n} X^\top D A X - \lambda I)$. Therefore the minimum eigenvalue of M a positive constant if both matrices below are positive definite

$$\frac{1}{n} X^\top D^2 X \text{ and } \frac{1}{n} X^\top A^2 X - \frac{1}{n} X^\top A D X \left(\frac{1}{n} X^\top D^2 X \right)^{-1} \frac{1}{n} X^\top D A X.$$

The first matrix has clearly minimum eigenvalue a positive constant. The second matrix is equal to $\frac{1}{n} X^\top A (I - D X (X^\top D^2 X)^{-1} X D) A X$ which has minimum eigenvalue positive by assumption. \square

Remark 4.2 (Smoothness of Hessian). *If we want to find an upper bound on the eigenvalues of the negative Hessian by an easy argument using Lemma 2.3 follows that*

$$\begin{aligned} & \lambda_{\max}(-H_{\theta, \beta, \kappa}) & (61) \\ & \leq C \left\{ \left\| (\beta A + D)^{-1/2} A (\beta A + D)^{-1/2} \right\|_F^2 + d \sigma_{\max}^2((\beta A + D)^{-1/2} D X) + d \sigma_{\max}^2((\beta A + D)^{-1/2} A X) \right\} & (62) \end{aligned}$$

for all $(\theta, \beta, \kappa) \in \mathbb{B}$. From Lemma 4.5 we conclude that there exists a positive constant C_H such that $\lambda_{\max}(-\frac{1}{n} H_{\theta, \beta, \kappa}) \leq C_H$.

The following lemma indicates that the concentration results and the lower bound on the eigenvalues of the negative Hessian (for this section) are of the desirable order.

Lemma 4.5 (Bounding the norms). *The following claims hold:*

1. $\|X\|_F^2$ is $\Theta(n)$.
2. $\frac{1}{n} \|X\|_2^2, \frac{1}{n} \lambda_{\min}(X^\top X)$ are positive constants.
3. $\|\Sigma^{1/2} A \Sigma^{1/2}\|_F^2, \sigma_{\min}^2(\Sigma^{1/2} D X)$ and $\sigma_{\min}^2(\Sigma^{1/2} A X)$ are $\Theta(n)$.

Proof. For claim 1,2, it follows from Lemma 2.1 that $\frac{1}{n}X^\top X$ has minimum eigenvalue at least $\lambda_{\min}(Q) - o(1)$ and maximum eigenvalue at most $\lambda_{\max}(Q)$ with probability $1 - o(1)$. For claim 3 we have that

$$\left\| \Sigma^{1/2} A \Sigma^{1/2} \right\|_F^2 = \|\Sigma A\|_F^2 \geq \sigma_{\min}^2(\Sigma) \|A\|_F^2 \text{ which is } \Theta(n) \text{ by assumption on } A, \Sigma.$$

Moreover, $\sigma_{\min}^2(\Sigma^{1/2} D X) \geq \sigma_{\min}(\Sigma) \cdot \sigma_{\min}^2(D) \lambda_{\min}(X^\top X)$ which is $\Theta(n)$ with high probability (note that $\sigma_{\min}(\Sigma)$, $\sigma_{\min}^2(D)$ are positive constants). Similarly the proof goes for $\sigma_{\min}^2(\Sigma^{1/2} A X)$. \square

We are now ready to prove the main theorem:

Proof of Theorem 4.1. For any $\delta > 0$ and using Markov's inequality it follows from Lemmas 4.1, 4.3, 4.2 and Lemma 4.5 that

$$\Pr_{\theta_0, \beta_0} \left[\|\nabla LL(\theta_0, \beta_0, \beta_0 \cdot \theta_0)\|_2 \geq C_\delta \sqrt{dn} \right] \leq \delta$$

for some constant C_δ and $\lambda_{\min}(\nabla^2 LL) \geq \frac{Cn}{d}$ for some constant C . We conclude from the analysis in Section 4.3 that $\left\| (\theta_0, \beta_0) - \hat{\theta}, \hat{\beta} \right\|$ is $O\left(d\sqrt{\frac{d}{n}}\right)$ with probability at least $1 - \delta$. \square

We conclude by showing that the S-K model satisfies the assumptions we have made and hence Theorem 4.1 can be applied to it.

Lemma 4.6 ((SK) model satisfies the assumptions). *Let A be a $n \times n$ matrix such that $A_{ij} = \frac{g_{ij}}{\sqrt{n}}$ for $i < j$, where $g_{ij} \sim \mathcal{N}(0, 1)$ and $A_{ji} = A_{ij}$, $A_{ii} = 0$. Matrix A satisfies the assumptions of our main Theorem 4.1.*

Proof. We assume for simplicity of the calculations that $D = cI$ for some positive constant c .

Set $F = I - X(X^\top X)^{-1}X^\top$ (F is called a hat/projection matrix, it has the property that d eigenvalues are zero and the rest are one since $F^2 = F$). Let B be a matrix with i.i.d entries $\mathcal{N}(0, 1)$. It is clear that the matrix $W = \frac{1}{\sqrt{n}}B$ satisfies the following:

- $\left\| \frac{W+W^\top}{\sqrt{2}} \right\|_F^2$ is lower bounded by the sum of $\frac{n(n+1)}{2}$ i.i.d χ^2 variables with mean $\frac{1}{n}$ and hence $\left\| \frac{W+W^\top}{\sqrt{2}} \right\|_F^2$ is lower bounded from $\Theta(n)$ with high probability $1 - o(1)$. Moreover, $\left\| \frac{W+W^\top}{\sqrt{2}} \right\|_F^2 \leq 2\|W\|_F^2$ and is clear that $\|W\|_F^2$ is concentrated around n (n^2 i.i.d variables with mean $1/n$). Thus $\left\| \frac{W+W^\top}{\sqrt{2}} \right\|_F^2$ is concentrated around $\Theta(n)$. Same is true for $\left\| \frac{F(W+W^\top)}{\sqrt{2}} \right\|_F^2$
- $\left\| \frac{W+W^\top}{\sqrt{2}} \right\|_2, \|W\|_2$ are with high probability $\Theta(1)$ (it follows from semicircle law, see [40]). Moreover, the same is true for $\left\| \frac{F(W+W^\top)}{\sqrt{2}} \right\|_2, \|FW\|_2$.

Note that the reason behind the fact that multiplying by F does not change the claims above is because $\sigma_j(FW) \geq \sigma_{n-d-j+1}(F)\sigma_{d+j}(W) = \sigma_{d+j}(W)$ where σ_j denotes the j -th largest eigenvalue of the corresponding matrix and $n \gg d$.

We first show that FWX has singular values the eigenvalues Q plus $o(1)$ with high probability. First let $F = R^\top I_d R$ where R is a rotation matrix and I_d is the identity matrix by setting the last d rows to all zeros. It is clear that RW is also a matrix with i.i.d gaussians of mean zero and variance $1/n$ each. Condition on X (in case $X^\top X = Q$ then the analysis is simplified), it follows that the rows of $I_d W X$ are independent (except of the last d rows that are all zeros) and each row follows a gaussian $\mathcal{N}(\mathbf{0}, \frac{1}{n} X^\top X)$. Hence using Theorem 2.1 it follows that

$$\left\| \frac{1}{n-d} X^\top W^\top F^2 W X - \frac{1}{n-d} X^\top X \right\|_2 \text{ is } O\left(\sqrt{\frac{\ln(n-d)}{n-d}}\right)$$

with probability $1 - o(1)$. Finally since again by Theorem 2.1 we get that $\left\| \frac{1}{n} X^\top X - Q \right\|_2$ is $O\left(\sqrt{\frac{\ln n}{n}}\right)$ with probability $1 - o(1)$, using triangle inequality we conclude that $\left\| \frac{1}{n} X^\top W^\top F^2 W X - Q \right\|_2$ is $O\left(\sqrt{\frac{\ln n}{n}}\right)$. The claim follows by Weyl's inequality (Lemma 2.1).

Moreover, we prove that $\frac{1}{n} \left\| F(W + W^\top) X \right\|_2^2$ is $\Omega(1)$ with probability $1 - o(1)$. Let us assume without loss of generality that $\frac{1}{\sqrt{n}} \|X\|_2 < \frac{1}{2\|FW\|_2^2}$ (by appropriately rescaling X with a constant). It holds that

$$\begin{aligned} & \lambda_{\min} \left(\frac{1}{n} X^\top (W^\top + W) F (W + W^\top) X \right) \\ & \geq \lambda_{\min} \left(\frac{1}{n} X^\top (W^\top F W + W F W^\top) X \right) + \lambda_{\min} \left(\frac{1}{n} X^\top (W F W + W^\top F W^\top) X \right) \\ & \geq \lambda_{\min} \left(\frac{1}{n} X^\top (W^\top F W + W F W^\top) X \right) - \left| \lambda_{\min} \left(\frac{1}{n} X^\top (W F W + W^\top F W^\top) X \right) \right| \\ & \stackrel{\text{Lemma 2.2}}{\geq} \lambda_{\min} \left(\frac{1}{n} X^\top (W^\top F W + W F W^\top) X \right) - 2\sigma_{\min} \left(\frac{1}{n} X^\top W F W X \right). \end{aligned}$$

It is clear from the analysis above that the first term is with probability $1 - o(1)$ within error $O\left(\sqrt{\frac{\ln n}{n}}\right)$ from $2\lambda_{\min}(Q)$ (a). We analyze the other term and we get using Lemma 2.2

$$\begin{aligned} 2\sigma_{\min} \left(\frac{1}{n} X^\top W F W X \right) & \leq 2\sqrt{\lambda_{\min} \left(\frac{1}{n} X^\top X \right)} \left\| \frac{1}{\sqrt{n}} X \right\|_2 \|W F W\|_2 \\ & \leq 2(\lambda_{\min}(Q) + o(1)) \left\| \frac{1}{\sqrt{n}} X \right\|_2 \|F W\|_2^2 \\ & \leq \lambda_{\min}(Q) + o(1) \quad (b). \end{aligned}$$

Finally by combining (a), (b) it holds that

$$\lambda_{\min} \left(\frac{1}{n} X^\top (W^\top + W) F (W + W^\top) X \right) \geq \lambda_{\min}(Q) - o(1),$$

which is a positive constant. Hence we conclude that $\frac{1}{n} \left\| \frac{F(W+W^\top)}{\sqrt{2}} X \right\|_2^2$ is a positive constant.

We define the matrix A to be $A_{ii} = 0$ (zeros in the diagonal) and $A_{ij} = \frac{W_{ij}+W_{ji}}{\sqrt{2}}$ for $i \neq j$ (A is symmetric). It is clear that A captures the SK model. Moreover, it is easy to show that all the diagonal entries of $\frac{W+W^\top}{\sqrt{2}}$ are smaller than $O(\frac{\sqrt{\log n}}{\sqrt{n}})$ with probability $1 - o(1)$, hence it follows that $\left\|A - \frac{W+W^\top}{\sqrt{2}}\right\|_2$ is $o(1)$ with high probability.

Therefore $\sigma_{\min}(\frac{1}{n}X^\top A^\top FAX)$, $\|A\|_2$ are positive constants and $\|A\|_F^2$ is $\Theta(n)$, all the statements with probability $1 - o(1)$ and the assumptions on matrix A are satisfied for linear regression model. \square

5 Projected Gradient Descent Analysis

In this Section, we will present the projected gradient descent algorithms we use for our logistic and linear regression settings. We will use the following well known property of Projected Gradient Descent (Theorem 3.10 from [6]).

Theorem 5.1. *Let f be α -strongly convex and λ -smooth on compact set \mathcal{X} . Then projected gradient descent with stepsize $\eta = \frac{1}{\lambda}$ satisfies for $t \geq 0$*

$$\|\mathbf{x}_{t+1} - \mathbf{x}^*\|_2^2 \leq e^{-\frac{\alpha t}{\lambda}} \|\mathbf{x}_1 - \mathbf{x}^*\|_2^2. \quad (63)$$

Therefore, setting $R = \|\mathbf{x}_1 - \mathbf{x}^*\|_2$ and by choosing $t = \frac{2\lambda \ln \frac{R}{\epsilon}}{\alpha}$ it is guaranteed that $\|\mathbf{x}_{t+1} - \mathbf{x}^*\|_2 \leq \epsilon$.

5.0.1 Projected Gradient Descent for Logistic Regression

We consider the function $LPL(\theta, \beta)$ (log-pseudolikelihood as defined in Section 3) and we would like to approximate $(\hat{\theta}, \hat{\beta})$ within $\frac{1}{\sqrt{n}}$ in ℓ_2 distance. The stepsize in Theorem 5.1 should be $\eta = \frac{1}{\sqrt{d\Theta^2+1}}$

by Remark 3.5.

ALGORITHM 1: Projected Gradient Descent (Logistic)

Data: Vector sample \mathbf{y} , Magnetizations $m_i(\mathbf{y}) = \sum_j A_{ij}y_j$, Feature vectors \mathbf{x}_i

Result: Maximum Pseudolikelihood Estimate

```

1  $\beta^0 = 0, \theta^0 = \mathbf{0}, \text{normgrad} = +\infty, \eta = \frac{1}{\sqrt{d\Theta^2+1}};$ 
2  $t = 0;$ 
3 while  $\text{normgrad} > \frac{1}{\sqrt{n}}$  do
4    $\text{grad}_\theta = 0;$ 
5    $\text{grad}_\beta = -\frac{1}{n} \sum_{i=1}^n [y_i m_i(\mathbf{y}) - m_i(\mathbf{y}) \tanh(\beta^t m_i(\mathbf{y}) + \theta^t \top \mathbf{x}_i)];$ 
6   for  $k = 1; k \leq d; k++$  do
7      $\text{grad}_{\theta_k} = -\frac{1}{n} \sum_{i=1}^n [y_i x_{i,k} - x_{i,k} \tanh(\beta^t m_i(\mathbf{y}) + \theta^t \top \mathbf{x}_i)];$ 
8      $\text{grad}_\theta = \text{grad}_\theta + \text{grad}_{\theta_k}^2;$ 
9   end
10   $\text{normgrad} = \sqrt{\text{grad}_\beta^2 + \text{grad}_\theta};$ 
11   $\beta^{t+1} = \beta^t - \eta \text{grad}_\beta$  % update  $\beta^t$ ;
12  for  $k = 1; k \leq d; k++$  do
13     $\theta_k^{t+1} = \theta_k^t - \eta \text{grad}_{\theta_k}$  % update  $\theta_k^t$ ;
14  end
15   $t = t + 1;$ 
16  %  $\ell_2$  projection
17  if  $\beta^{t+1} < -B$  then
18     $\beta^{t+1} = -B;$ 
19  end
20  if  $\beta^{t+1} > B$  then
21     $\beta^{t+1} = B;$ 
22  end
23  for  $k = 1; k \leq d; k++$  do
24    if  $\theta_k^{t+1} < -\Theta$  then
25       $\theta_k^{t+1} = -\Theta;$ 
26    end
27    if  $\theta_k^{t+1} > \Theta$  then
28       $\theta_k^{t+1} = \Theta;$ 
29    end
30  end
31 end
32 return  $(\theta^t, \beta^t)$ 

```

5.0.2 Projected Gradient Descent for Linear Regression

We consider the function $LL(\theta, \beta, \kappa)$ (log-pseudolikelihood as defined in Section 4.2) and we would like to approximate $(\hat{\theta}, \hat{\beta}, \hat{\kappa})$ within $\frac{1}{\sqrt{n}}$ in ℓ_2 distance. The stepsize in Theorem 5.1 should be

$\eta = 1/C_H$ where C_H is the constant from Remark 4.2.

ALGORITHM 2: Projected Gradient Descent (Linear)

Data: Vector sample \mathbf{y} , Matrices A, D , Feature matrix X

Result: Maximum Likelihood Estimate

```

1  $\beta^0 = 0, \theta^0 = \mathbf{0}, \kappa^0 = \mathbf{0}, \text{normgrad} = +\infty;$ 
2  $t = 0;$ 
3 while  $\text{normgrad} > \frac{1}{\sqrt{n}}$  do
4    $\text{grad}_\theta = 0;$ 
5    $\text{grad}_\kappa = 0;$ 
6    $\text{grad}_\beta = \frac{1}{2}\mathbf{y}^\top A\mathbf{y} - \frac{1}{2}\text{tr}(A(\beta^t A + D)^{-1}) - \frac{1}{2}(AX\kappa^t + DX\theta^t)^\top ((\beta^t A + D)^{-1})A((\beta^t A + D)^{-1})(AX\kappa^t + DX\theta^t);$ 
7   for  $k = 1; k \leq d; k++$  do
8      $\text{grad}_{\theta_k} = -\sum_{i=1}^n y_i D_{ii} x_{ik} + \sum_{i=1}^n ((\beta^t A + D)^{-1}(AX\kappa^t + DX\theta^t))_i D_{ii} x_{ik};$ 
9      $\text{grad}_\theta = \text{grad}_\theta + \text{grad}_{\theta_k}^2;$ 
10     $\text{grad}_{\kappa_k} = -\sum_{i=1}^n \sum_{j=1}^n A_{ij} y_i x_{jk} + \sum_{i=1}^n \sum_{j=1}^n A_{ij} x_{jk} ((\beta^t A + D)^{-1}(AX\kappa^t + DX\theta^t))_i;$ 
11     $\text{grad}_\kappa = \text{grad}_\kappa + \text{grad}_{\kappa_k}^2;$ 
12  end
13   $\text{normgrad} = \sqrt{\text{grad}_\beta^2 + \text{grad}_\theta + \text{grad}_\kappa};$ 
14   $\beta^{t+1} = \beta^t - \eta \text{grad}_\beta$  % update  $\beta^t$ ;
15  for  $k = 1; k \leq d; k++$  do
16     $\theta_k^{t+1} = \theta_k^t - \eta \text{grad}_{\theta_k}$  % update  $\theta_k^t$ ;
17     $\kappa_k^{t+1} = \kappa_k^t - \eta \text{grad}_{\kappa_k}$  % update  $\kappa_k^t$ ;
18  end
19   $t = t + 1;$ 
20  %  $\ell_2$  projection
21  if  $\beta^{t+1} < -B$  then
22     $\beta^{t+1} = -B;$ 
23  end
24  if  $\beta^{t+1} > B$  then
25     $\beta^{t+1} = B;$ 
26  end
27  for  $k = 1; k \leq d; k++$  do
28    if  $\theta_k^{t+1} < -\Theta$  then
29       $\theta_k^{t+1} = -\Theta;$ 
30    end
31    if  $\theta_k^{t+1} > \Theta$  then
32       $\theta_k^{t+1} = \Theta;$ 
33    end
34    if  $\kappa_k^{t+1} < -B \cdot \Theta$  then
35       $\kappa_k^{t+1} = -B \cdot \Theta;$ 
36    end
37    if  $\kappa_k^{t+1} > B \cdot \Theta$  then
38       $\kappa_k^{t+1} = B \cdot \Theta;$ 
39    end
40  end
41 end
42 return  $(\theta^t, \beta^t)$ 

```

References

- [1] Alekh Agarwal and John C Duchi. The generalization ability of online algorithms for dependent data. *IEEE Transactions on Information Theory*, 59(1):573–587, 2013.
- [2] Marianne Bertrand, Erzo FP Luttmer, and Sendhil Mullainathan. Network effects and welfare cultures. *The Quarterly Journal of Economics*, 115(3):1019–1055, 2000.
- [3] Julian Besag. Statistical analysis of non-lattice data. *The statistician*, pages 179–195, 1975.
- [4] Bhaswar B Bhattacharya, Sumit Mukherjee, et al. Inference in ising models. *Bernoulli*, 24(1):493–525, 2018.
- [5] Yann Bramoullé, Habiba Djebbari, and Bernard Fortin. Identification of peer effects through social networks. *Journal of econometrics*, 150(1):41–55, 2009.
- [6] Sébastien Bubeck et al. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.
- [7] Sourav Chatterjee. *Concentration Inequalities with Exchangeable Pairs*. PhD thesis, Stanford University, June 2005.
- [8] Sourav Chatterjee. Concentration inequalities with exchangeable pairs (ph. d. thesis). *arXiv preprint math/0507526*, 2005.
- [9] Sourav Chatterjee and Amir Dembo. Nonlinear large deviations. *Advances in Mathematics*, 299:396–450, 2016.
- [10] Sourav Chatterjee et al. Estimation in spin glasses: A first step. *The Annals of Statistics*, 35(5):1931–1946, 2007.
- [11] Nicholas A Christakis and James H Fowler. Social contagion theory: examining dynamic social networks and human behavior. *Statistics in medicine*, 32(4):556–577, 2013.
- [12] Francis Comets and Basilis Gidas. Asymptotics of maximum likelihood estimators for the curie-weiss model. *The Annals of Statistics*, pages 557–578, 1991.
- [13] Constantinos Daskalakis, Nishanth Dikkala, and Siddhartha Jayanti. Hogwild!-gibbs can be panaccurate. In *Advances in Neural Information Processing Systems*, pages 32–41, 2018.
- [14] Constantinos Daskalakis, Nishanth Dikkala, and Gautam Kamath. Concentration of multi-linear functions of the ising model with applications to network data. In *Advances in Neural Information Processing Systems*, pages 12–23, 2017.
- [15] Constantinos Daskalakis, Nishanth Dikkala, and Gautam Kamath. Testing ising models. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1989–2007. Society for Industrial and Applied Mathematics, 2018.
- [16] Constantinos Daskalakis, Elchanan Mossel, and Sébastien Roch. Evolutionary trees and the Ising model on the Bethe lattice: A proof of Steel’s conjecture. *Probability Theory and Related Fields*, 149(1):149–189, 2011.

- [17] Christopher De Sa, Kunle Olukotun, and Christopher Ré. Ensuring rapid mixing and low bias for asynchronous gibbs sampling. In *JMLR workshop and conference proceedings*, volume 48, page 1567. NIH Public Access, 2016.
- [18] Esther Duflo and Emmanuel Saez. The role of information and social interactions in retirement plan decisions: Evidence from a randomized experiment. *The Quarterly journal of economics*, 118(3):815–842, 2003.
- [19] Glenn Ellison. Learning, local interaction, and coordination. *Econometrica*, 61(5):1047–1071, 1993.
- [20] K. Fan and A.J. Hoffman. Some metric inequalities in the space of matrices. *Proc. Amer. Math. Soc.*, pages 111–116, 1955.
- [21] Joseph Felsenstein. *Inferring Phylogenies*. Sinauer Associates Sunderland, 2004.
- [22] Stuart Geman and Christine Graffigne. Markov random field image models and their applications to computer vision. In *Proceedings of the International Congress of Mathematicians*, pages 1496–1517. American Mathematical Society, 1986.
- [23] Promit Ghosal and Sumit Mukherjee. Joint estimation of parameters in ising model. *arXiv preprint arXiv:1801.06570*, 2018.
- [24] Edward L Glaeser, Bruce Sacerdote, and Jose A Scheinkman. Crime and social interactions. *The Quarterly Journal of Economics*, 111(2):507–548, 1996.
- [25] Linus Hamilton, Frederic Koehler, and Ankur Moitra. Information theoretic properties of markov random fields, and their algorithmic applications. In *Advances in Neural Information Processing Systems*, pages 2463–2472, 2017.
- [26] Kathleen Mullan Harris, National Longitudinal Study of Adolescent Health, et al. Waves i & ii, 1994–1996; wave iii, 2001–2002; wave iv, 2007–2009 [machine-readable data file and documentation]. *Chapel Hill, NC: Carolina Population Center, University of North Carolina at Chapel Hill*, 10, 2009.
- [27] Ernst Ising. Beitrag zur theorie des ferromagnetismus. *Zeitschrift für Physik A Hadrons and Nuclei*, 31(1):253–258, 1925.
- [28] Adam Klivans and Raghu Meka. Learning graphical models using multiplicative weights. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 343–354. IEEE, 2017.
- [29] David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society, 2009.
- [30] Tianxi Li, Elizaveta Levina, and Ji Zhu. Prediction models for network-linked data. *arXiv preprint arXiv:1602.01192*, 2016.
- [31] Charles F Manski. Identification of endogenous social effects: The reflection problem. *The review of economic studies*, 60(3):531–542, 1993.

- [32] Ron Meir. Nonparametric time series prediction through adaptive model selection. *Machine learning*, 39(1):5–34, 2000.
- [33] Dharmendra S Modha and Elias Masry. Minimum complexity regression estimation with weakly dependent observations. *IEEE Transactions on Information Theory*, 42(6):2133–2145, 1996.
- [34] Mehryar Mohri and Afshin Rostamizadeh. Stability bounds for stationary φ -mixing and β -mixing processes. *Journal of Machine Learning Research*, 11(Feb):789–814, 2010.
- [35] Andrea Montanari and Amin Saberi. The spread of innovations in social networks. *Proceedings of the National Academy of Sciences*, 107(47):20196–20201, 2010.
- [36] Bruce Sacerdote. Peer effects with random assignment: Results for dartmouth roommates. *The Quarterly journal of economics*, 116(2):681–704, 2001.
- [37] D. Sherrington and S. Kirkpatrick. Solvable model of a spin-glass. *Phys. Rev. Lett.*, 35, 1975.
- [38] Ingo Steinwart and Andreas Christmann. Fast learning from non-iid observations. In *Advances in neural information processing systems*, pages 1768–1776, 2009.
- [39] Justin G Trogdon, James Nonnemaker, and Joanne Pais. Peer effects in adolescent overweight. *Journal of health economics*, 27(5):1388–1399, 2008.
- [40] Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.
- [41] Bin Zou, Luoqing Li, and Zongben Xu. The generalization performance of erm algorithm with strongly mixing observations. *Machine learning*, 75(3):275–295, 2009.