
Large-Scale Optimization Algorithms for Sparse Conditional Gaussian Graphical Models

Calvin McCarter
Machine Learning Department
Carnegie Mellon University
calvinm@cmu.edu

Seyoung Kim
Computational Biology Department
Carnegie Mellon University
sssykim@cs.cmu.edu

Abstract

This paper addresses the problem of scalable optimization for l_1 -regularized conditional Gaussian graphical models. Conditional Gaussian graphical models generalize the well-known Gaussian graphical models to conditional distributions to model the output network influenced by conditioning input variables. While highly scalable optimization methods exist for sparse Gaussian graphical model estimation, state-of-the-art methods for conditional Gaussian graphical models are not efficient enough and more importantly, fail due to memory constraints for very large problems. In this paper, we propose a new optimization procedure based on a Newton method that efficiently iterates over two sub-problems, leading to drastic improvement in computation time compared to the previous methods. We then extend our method to scale to large problems under memory constraints, using block coordinate descent to limit memory usage while achieving fast convergence. Using synthetic and genomic data, we show that our methods can solve problems with millions of variables and tens of billions of parameters to high accuracy on a single machine.

1 INTRODUCTION

Sparse Gaussian graphical models (GGMs) [2] have been extremely popular as a tool for learning a network structure over a large number of continuous variables in many different application domains including neuroscience [7] and biology [2]. A sparse GGM can be estimated as a sparse in-

verse covariance matrix by minimizing the convex function of l_1 -regularized negative log-likelihood. Highly scalable learning algorithms such as graphical lasso [2], QUIC [3], and BigQUIC [4] have been proposed to learn the model.

In this paper, we address the problem of scaling up the optimization of sparse conditional GGM (CGGM), a model closely related to sparse GGM, to very large problem sizes without requiring excessive time or memory. Sparse CGGMs have been introduced as a discriminative extension of sparse GGMs to model a sparse network over outputs conditional on input variables [8, 10]. CGGMs can be viewed as a Gaussian analogue of conditional random field [6], while GGMs are a Gaussian analogue of Markov random field. A sparse CGGM can be estimated by minimizing a convex function of l_1 -regularized negative log-likelihood. This optimization problem is closely related to that for sparse GGMs because CGGMs also model the network over outputs. However, the presence of the additional parameters in CGGMs for the functional mapping from inputs to outputs makes the optimization significantly more complex than in sparse GGMs.

Several different approaches have been previously proposed to estimate sparse CGGMs, including OWL-QN [8], accelerated proximal gradient method [11], and Newton coordinate descent algorithm [10]. In particular, the Newton coordinate descent algorithm extends the QUIC algorithm [3] for sparse GGM estimation to the case of CGGMs, and has been shown to have superior computational speed and convergence. This approach finds in each iteration a descent direction by minimizing a quadratic approximation of the original negative log-likelihood function along with l_1 regularization. Then, the parameter estimate is updated with this descent direction and a step size found by line search.

Although the Newton coordinate descent method [10] is state-of-the-art for its scalability and fast convergence, it is still not efficient enough to be applied to many real-world problems even with tens of thousands of variables. More importantly, it suffers from a large space requirement, because for very high-dimensional problems, several large

Appearing in Proceedings of the 19th International Conference on Artificial Intelligence and Statistics (AISTATS) 2016, Cadiz, Spain. JMLR: W&CP volume 51. Copyright 2016 by the authors.

dense matrices need to be precomputed and stored during optimization. For a CGGM with p inputs and q outputs, the algorithm requires storing several $p \times p$ and $q \times q$ dense matrices, which cannot fit in memory for large p and q .

We propose new algorithms for learning l_1 -regularized CGGMs that significantly improve the computation time of the previous Newton coordinate descent algorithm and also remove the large memory requirement. We first propose an optimization method, called an alternating Newton coordinate descent algorithm, for improving computation time. Our algorithm is based on the key observation that the computation simplifies drastically, if we alternately optimize the two sets of parameters for output network and for mapping inputs to outputs, instead of updating all parameters at once as in the previous approach. The previous approach updated all parameters simultaneously by forming a second-order approximation of the objective on all parameters, which requires an expensive computation of the large Hessian matrix of size $(p + q) \times (p + q)$ in each iteration. Our approach of alternate optimization forms a second-order approximation only on the network parameters, which requires the Hessian of size $q \times q$, as the other set of parameters can be updated easily using a simple coordinate descent.

In order to overcome the constraint on the space requirement, we then extend our algorithm to an alternating Newton block coordinate descent method that can be applied to problems of unbounded size on a machine with limited memory. Instead of recomputing each element of the large matrices on demand, we divide the parameters into blocks for block-wise updates such that the results of computation can be reused within each block. Block-wise parameter updates were previously used in BigQUIC [4] for learning a sparse GGM, where the block sparsity pattern of the network parameters was leveraged to overcome the space limitations. We propose an approach for block-wise update of the output network parameters in CGGMs that extends their idea. We then propose a new block-wise update strategy for the parameters for mapping inputs to outputs. In our experiments, we show that we can solve problems with a million inputs and hundreds of thousands of outputs on a single machine.

The rest of the paper is organized as follows. In Section 2, we provide a brief review of sparse CGGMs and the current state-of-the-art Newton coordinate descent algorithm [10] for learning the models. In Section 3, we propose an alternating Newton coordinate descent algorithm that significantly reduces computation time compared to the previous method. In Section 4, we further extend our algorithm to perform block-wise updates in order to scale up to very large problems on a machine with bounded memory. In Section 5, we demonstrate our proposed algorithms on synthetic and real-world genomic data.

2 BACKGROUND

2.1 The Conditional Gaussian Graphical Model

A CGGM [8, 10] models the conditional probability density of $\mathbf{x} \in \mathbb{R}^p$ given $\mathbf{y} \in \mathbb{R}^q$ as follows:

$$p(\mathbf{y}|\mathbf{x}; \mathbf{\Lambda}, \mathbf{\Theta}) = \exp\{-\mathbf{y}^T \mathbf{\Lambda} \mathbf{y} - 2\mathbf{x}^T \mathbf{\Theta} \mathbf{y}\} / Z(\mathbf{x}),$$

where $\mathbf{\Lambda}$ is a $q \times q$ matrix for modeling the network over \mathbf{y} and $\mathbf{\Theta}$ is a $p \times q$ matrix for modeling the mapping between the input variables \mathbf{x} and output variables \mathbf{y} . The normalization constant is given as $Z(\mathbf{x}) = (2\pi)^{q/2} |\mathbf{\Lambda}|^{-1} \exp(\mathbf{x}^T \mathbf{\Theta} \mathbf{\Lambda}^{-1} \mathbf{\Theta}^T \mathbf{x})$. Inference in a CGGM gives $p(\mathbf{y}|\mathbf{x}) = \mathcal{N}(\mathbf{B}^T \mathbf{x}, \mathbf{\Lambda}^{-1})$, where $\mathbf{B} = -\mathbf{\Theta} \mathbf{\Lambda}^{-1}$, showing the connection to multivariate linear regression.

Given a mean-centered dataset of $\mathbf{X} \in \mathbb{R}^{n \times p}$ and $\mathbf{Y} \in \mathbb{R}^{n \times q}$ for n samples, and their covariance matrices $\mathbf{S}_{\mathbf{xx}} = \frac{1}{n} \mathbf{X}^T \mathbf{X}$, $\mathbf{S}_{\mathbf{xy}} = \frac{1}{n} \mathbf{X}^T \mathbf{Y}$, $\mathbf{S}_{\mathbf{yy}} = \frac{1}{n} \mathbf{Y}^T \mathbf{Y}$, a sparse estimate of CGGM parameters can be obtained by minimizing l_1 -regularized negative log-likelihood:

$$\min_{\mathbf{\Lambda} > 0, \mathbf{\Theta}} f(\mathbf{\Lambda}, \mathbf{\Theta}) = g(\mathbf{\Lambda}, \mathbf{\Theta}) + h(\mathbf{\Lambda}, \mathbf{\Theta}), \quad (1)$$

where $g(\mathbf{\Lambda}, \mathbf{\Theta}) = -\log |\mathbf{\Lambda}| + \text{tr}(\mathbf{S}_{\mathbf{yy}} \mathbf{\Lambda} + 2\mathbf{S}_{\mathbf{xy}}^T \mathbf{\Theta} + \mathbf{\Lambda}^{-1} \mathbf{\Theta}^T \mathbf{S}_{\mathbf{xx}} \mathbf{\Theta})$ for the smooth negative log-likelihood and $h(\mathbf{\Lambda}, \mathbf{\Theta}) = \lambda_{\mathbf{\Lambda}} \|\mathbf{\Lambda}\|_1 + \lambda_{\mathbf{\Theta}} \|\mathbf{\Theta}\|_1$ for the non-smooth elementwise l_1 penalty. $\lambda_{\mathbf{\Lambda}}, \lambda_{\mathbf{\Theta}} > 0$ are regularization parameters. As observed in [8, 10, 11], this objective is convex.

2.2 Optimization

The current state-of-the-art method for solving Eq. (1) for l_1 -regularized CGGM is the Newton coordinate descent algorithm [10] that extends QUIC [3] for l_1 -regularized GGM estimation. In each iteration, this algorithm found a generalized Newton descent direction by forming a second-order approximation of the smooth part of the objective and minimizing this along with the l_1 penalty. Given this Newton direction, the parameter estimates were updated with a step size found by line search using Armijo's rule [1].

In each iteration, the Newton coordinate descent algorithm found the Newton direction as follows:

$$\mathbf{D}_{\mathbf{\Lambda}}, \mathbf{D}_{\mathbf{\Theta}} = \underset{\Delta_{\mathbf{\Lambda}}, \Delta_{\mathbf{\Theta}}}{\text{argmin}} \bar{g}_{\mathbf{\Lambda}, \mathbf{\Theta}}(\Delta_{\mathbf{\Lambda}}, \Delta_{\mathbf{\Theta}}) + h(\mathbf{\Lambda} + \Delta_{\mathbf{\Lambda}}, \mathbf{\Theta} + \Delta_{\mathbf{\Theta}}), \quad (2)$$

where $\bar{g}_{\mathbf{\Lambda}, \mathbf{\Theta}}$ is the second-order approximation of g given by Taylor expansion:

$$\begin{aligned} \bar{g}_{\mathbf{\Lambda}, \mathbf{\Theta}}(\Delta_{\mathbf{\Lambda}}, \Delta_{\mathbf{\Theta}}) &= \text{vec}(\nabla g(\mathbf{\Lambda}, \mathbf{\Theta}))^T \text{vec}([\Delta_{\mathbf{\Lambda}} \ \Delta_{\mathbf{\Theta}}]) \\ &\quad + \frac{1}{2} \text{vec}([\Delta_{\mathbf{\Lambda}} \ \Delta_{\mathbf{\Theta}}])^T \nabla^2 g(\mathbf{\Lambda}, \mathbf{\Theta}) \text{vec}([\Delta_{\mathbf{\Lambda}} \ \Delta_{\mathbf{\Theta}}]). \end{aligned}$$

The gradient and Hessian matrices above are given as:

$$\begin{aligned}\nabla g(\Lambda, \Theta) &= [\nabla_{\Lambda} g(\Lambda, \Theta) \quad \nabla_{\Theta} g(\Lambda, \Theta)] \\ &= [\mathbf{S}_{\mathbf{y}\mathbf{y}} - \Sigma - \Psi \quad 2\mathbf{S}_{\mathbf{x}\mathbf{y}} + 2\Gamma] \quad (3) \\ \nabla^2 g(\Lambda, \Theta) &= \begin{bmatrix} \nabla_{\Lambda}^2 g(\Lambda, \Theta) & \nabla_{\Lambda} \nabla_{\Theta} g(\Lambda, \Theta) \\ \nabla_{\Lambda} \nabla_{\Theta} g(\Lambda, \Theta)^T & \nabla_{\Theta}^2 g(\Lambda, \Theta) \end{bmatrix} \\ &= \begin{bmatrix} \Sigma \otimes (\Sigma + 2\Psi) & -2\Sigma \otimes \Gamma^T \\ -2\Sigma \otimes \Gamma & 2\Sigma \otimes \mathbf{S}_{\mathbf{x}\mathbf{x}} \end{bmatrix}, \quad (4)\end{aligned}$$

where $\Sigma = \Lambda^{-1}$, $\Psi = \Sigma \Theta^T \mathbf{S}_{\mathbf{x}\mathbf{x}} \Theta \Sigma$, and $\Gamma = \mathbf{S}_{\mathbf{x}\mathbf{x}} \Theta \Sigma$. Given the Newton direction in Eq. (2), the parameters can be updated as $\Lambda \leftarrow \Lambda + \alpha \mathbf{D}_{\Lambda}$ and $\Theta \leftarrow \Theta + \alpha \mathbf{D}_{\Theta}$, where step size $0 < \alpha \leq 1$ ensures sufficient decrease in Eq. (1) and positive definiteness of Λ .

The *Lasso* problem [9] in Eq. (2) was solved using coordinate descent. Despite the efficiency of coordinate descent for *Lasso*, applying coordinate updates repeatedly to all $q^2 + pq$ variables in Λ and Θ is costly. So, the updates were restricted to an active set of variables given as:

$$\begin{aligned}\mathcal{S}_{\Lambda} &= \{(\Delta_{\Lambda})_{ij} : |(\nabla_{\Lambda} g(\Lambda, \Theta))_{ij}| > \lambda_{\Lambda} \vee \Lambda_{ij} \neq 0\} \\ \mathcal{S}_{\Theta} &= \{(\Delta_{\Theta})_{ij} : |(\nabla_{\Theta} g(\Lambda, \Theta))_{ij}| > \lambda_{\Theta} \vee \Theta_{ij} \neq 0\}.\end{aligned}$$

Because the active set sizes $m_{\Lambda} = |\mathcal{S}_{\Lambda}|$, $m_{\Theta} = |\mathcal{S}_{\Theta}|$ approach the number of non-zero entries in the sparse solutions for Λ^* and Θ^* over iterations, this strategy yields a substantial speedup.

To further improve the efficiency of coordinate descent, intermediate results were stored for the large matrix products that need to be computed repeatedly. At the beginning of the optimization for Eq. (2), $\mathbf{U} := \Delta_{\Lambda} \Sigma$ and $\mathbf{V} := \Delta_{\Theta} \Sigma$ were computed and stored. Then, after a coordinate descent update to $(\Delta_{\Lambda})_{ij}$, row i and j of \mathbf{U} were updated. Similarly, after an update to $(\Delta_{\Theta})_{ij}$, row i of \mathbf{V} was updated.

2.3 Computational Complexity and Scalability

Although the Newton coordinate descent method is computationally more efficient than other previous approaches, it does not scale to problems even with tens of thousands of variables. The main computational cost of the algorithm comes from computing the large $(p+q) \times (p+q)$ Hessian matrix in Eq. (4) in each application of Eq. (2) to find the Newton direction. At the beginning of the optimization in Eq. (2), large dense matrices Σ , Ψ , and Γ , for computing the gradient and Hessian in Eqs. (3) and (4), are pre-computed and reused throughout the coordinate descent iterations. Initializing $\Sigma = \Lambda^{-1}$ via Cholesky decomposition costs up to $O(q^3)$ time, although in practice, sparse Cholesky decomposition exploits sparsity to invert Λ in much less than $O(q^3)$. Computing $\Psi = \frac{1}{n} \mathbf{R}^T \mathbf{R}$, where $\mathbf{R} = \mathbf{X} \Theta \Sigma$, requires $O(nm_{\Theta} + nq^2)$ time, and computing Γ costs $O(npq + nq^2)$. After the initializations, the cost of coordinate descent update per each active variable $(\Delta_{\Lambda})_{ij}$ and $(\Delta_{\Theta})_{ij}$ is $O(p+q)$. During the coordinate descent for

solving Eq. (2), the entire $(p+q) \times (p+q)$ Hessian matrix in Eq. (4) needs to be evaluated, whereas for the gradient in Eq. (3) only those entries corresponding to the parameters in active sets are evaluated.

A more serious obstacle to scaling up to problems with large p and q is the space required to store dense matrices Σ (size $q \times q$), Ψ (size $q \times q$), and Γ (size $p \times q$). In our experiments on a machine with 104 Gb RAM, the Newton coordinate descent method exhausted memory when $p+q$ exceeded 80,000.

In the next section, we propose a modification of the Newton coordinate descent algorithm that significantly improves the computation time. Then, we introduce block-wise update strategies to our algorithm to remove the memory constraint.

3 ALTERNATING NEWTON COORDINATE DESCENT

In this section, we introduce our alternating Newton coordinate descent algorithm for learning an l_1 -regularized CGGM that significantly reduces computation time compared to the previous method. Instead of performing Newton descent for all parameters Λ and Θ simultaneously, our approach alternately updates Λ and Θ , optimizing Eq. (1) over Λ given Θ and vice versa until convergence.

Our approach is based on the key observation that with Λ fixed, the problem of solving Eq. (1) over Θ becomes simply minimizing a quadratic function with l_1 regularization. Thus, it can be solved efficiently using a coordinate descent method, without the need to form a second-order approximation or to perform line search. On the other hand, optimizing Eq. (1) for Λ given Θ still requires forming a quadratic approximation to find a generalized Newton direction and performing line search to find the step size. However, this computation involves only $q \times q$ Hessian matrix and is significantly simpler than performing the same type of computation on both Λ and Θ jointly as in the previous approach.

3.1 Coordinate Descent Optimization for Λ

Given fixed Θ , the problem of minimizing the objective in Eq. (1) with respect to Λ becomes

$$\operatorname{argmin}_{\Lambda \succ 0} g_{\Theta}(\Lambda) + \lambda_{\Lambda} \|\Lambda\|_1,$$

where $g_{\Theta}(\Lambda) = -\log |\Lambda| + \operatorname{tr}(\mathbf{S}_{\mathbf{y}\mathbf{y}} \Lambda + \Lambda^{-1} \Theta^T \mathbf{S}_{\mathbf{x}\mathbf{x}} \Theta)$. In order to solve this, we first find a generalized Newton direction that minimizes the l_1 -regularized quadratic approximation of $g_{\Theta}(\Lambda)$:

$$\mathbf{D}_{\Lambda} = \operatorname{argmin}_{\Delta_{\Lambda}} \bar{g}_{\Lambda, \Theta}(\Delta_{\Lambda}) + \lambda_{\Lambda} \|\Lambda + \Delta_{\Lambda}\|_1, \quad (5)$$

Algorithm 1: Alternating Newton Coordinate Descent

input : Inputs $X \in \mathbb{R}^{n \times p}$ and $Y \in \mathbb{R}^{n \times q}$; regularization parameters $\lambda_\Lambda, \lambda_\Theta$
output: Parameters Λ, Θ
 Initialize $\Theta \leftarrow 0, \Lambda \leftarrow I_q$
for $t = 0, 1, \dots$ **do**
 Determine active sets $\mathcal{S}_\Lambda, \mathcal{S}_\Theta$
 Solve via coordinate descent:
 $D_\Lambda = \underset{\Delta_\Lambda}{\operatorname{argmin}} \bar{g}_{\Lambda, \Theta}(\Lambda + \Delta_\Lambda, \Theta) + h(\Lambda + \Delta_\Lambda, \Theta)$
 Update $\Lambda^+ = \Lambda + \alpha D_\Lambda$, where step size α is found with line search
 Solve via coordinate descent:
 $\Theta^+ = \underset{\Theta}{\operatorname{argmin}} g_\Lambda(\Theta) + \lambda_\Theta \|\Theta\|_1$

where $\bar{g}_{\Lambda, \Theta}(\Delta_\Lambda)$ is obtained from a second-order Taylor expansion and is given as

$$\bar{g}_{\Lambda, \Theta}(\Delta_\Lambda) = \operatorname{vec}(\nabla_\Lambda g(\Lambda, \Theta))^T \operatorname{vec}(\Delta_\Lambda) + \frac{1}{2} \operatorname{vec}(\Delta_\Lambda)^T \nabla_\Lambda^2 g(\Lambda, \Theta) \operatorname{vec}(\Delta_\Lambda).$$

The $\nabla_\Lambda g(\Lambda, \Theta)$ and $\nabla_\Lambda^2 g(\Lambda, \Theta)$ above are components of the gradient and Hessian matrices corresponding to Λ in Eqs. (3) and (4). We solve the *Lasso* problem in Eq. (5) via coordinate descent. Similar to the Newton coordinate descent method, we maintain $\mathbf{U} := \Delta_\Lambda \Sigma$ to reuse intermediate results of the large matrix-matrix product. Given the Newton direction for Λ , we update $\Lambda \leftarrow \Lambda + \alpha \Delta_\Lambda$, where α is obtained by line search.

Restricting the generalized Newton descent to Λ simplifies the computation significantly for coordinate descent updates, compared to the previous approach [10] that applies it to both Λ and Θ jointly. Our updates only involve $\nabla_\Lambda g(\Lambda, \Theta)$ and $\nabla_\Lambda^2 g(\Lambda, \Theta)$, and no longer involve $\nabla_\Theta g(\Lambda, \Theta)$ and $\nabla_\Lambda \nabla_\Theta g(\Lambda, \Theta)$, eliminating the need to compute the large $p \times q$ dense matrix Γ in $O(npq + nq^2)$ time. Our approach also reduces the computational cost for the coordinate descent update of each element of Δ_Λ from $O(p + q)$ to $O(q)$.

3.2 Coordinate Descent Optimization for Θ

With Λ fixed, the optimization problem in Eq. (1) with respect to Θ becomes

$$\underset{\Theta}{\operatorname{argmin}} g_\Lambda(\Theta) + \lambda_\Theta \|\Theta\|_1, \quad (6)$$

where $g_\Lambda(\Theta) = \operatorname{tr}(2\mathbf{S}_{xy}^T \Theta + \Lambda^{-1} \Theta^T \mathbf{S}_{xx} \Theta)$. Since $g_\Lambda(\Theta)$ is a quadratic function itself, there is no need to form its second-order Taylor expansion or to determine a step size via line search. Instead, we solve Eq. (6) directly with coordinate descent method, storing and maintaining $\mathbf{V} := \Theta \Sigma$. Our approach reduces the computation time for

updating Θ compared to the corresponding computation in the previous algorithm [10]. We avoid computing the large $p \times q$ matrix Γ , which had dominated overall computation time with $O(npq)$. Our approach also eliminates the need for line search for updating Θ . Finally, it reduces the cost for each coordinate descent update in Θ to $O(p)$, compared to $O(p + q)$ for the corresponding computation for Δ_Θ in the previous method.

Our approach is summarized in Algorithm 1. We provide the details of the coordinate descent update equations in Appendix. In practice, we approximately solve Eqs. (5) and (6) by using a warm-start for Λ and Θ with the results of the previous iteration and making a single pass over the active set. This ensures decrease in the objective in Eq. (1) and reduces the overall computation time in practice.

4 ALTERNATING NEWTON BLOCK COORDINATE DESCENT

The alternating Newton coordinate descent algorithm in the previous section improves the computation time of the previous state-of-the-art method, but is still limited by the space required to store large matrices during coordinate descent computation. Solving Eq. (5) for updating Λ requires precomputing and storing $q \times q$ matrices, Σ and Ψ , whereas solving Eq. (6) for updating Θ requires Σ and a $p \times p$ matrix \mathbf{S}_{xx} . A naive approach to reduce the memory footprint would be to recompute portions of these matrices on demand for each coordinate update, which would be very expensive.

In this section we describe how our algorithm in the previous section can be combined with block coordinate descent to scale up the optimization to very large problems on a machine with limited memory. During coordinate descent optimization, we update blocks of Λ and Θ so that within each block, the computation of the large matrices can be cached and re-used, where these blocks are determined automatically by exploiting the sparse structure. For Λ , we extend the block coordinate descent approach in BigQUIC [4] developed for GGMs to take into account the conditioning variables in CGGMs. For Θ , we describe a new approach for block coordinate descent update. Our algorithm can, in principle, be applied to problems of any size on a machine with limited memory.

4.1 Blockwise Optimization for Λ

4.1.1 Block Coordinate Descent Method

A coordinate-descent update of $(\Delta_\Lambda)_{ij}$ requires the i th and j th columns of Σ and Ψ . If these columns are in memory, they can be reused. Otherwise, it is a cache miss and we should compute them on demand. Σ_i for the i th column of Σ can be obtained by solving linear system $\Lambda \Sigma_i = \mathbf{e}_i$

Algorithm 2: Alternating Newton Block Coordinate Descent

input : Inputs $\mathbf{X} \in \mathbb{R}^{n \times p}$ and outputs $\mathbf{Y} \in \mathbb{R}^{n \times q}$;
 regularization parameters $\lambda_\Lambda, \lambda_\Theta$
output: Parameters Λ, Θ
 Initialize $\Theta \leftarrow 0, \Lambda \leftarrow I_q$
for $t = 0, 1, \dots$ **do**
 Determine active sets $\mathcal{S}_\Lambda, \mathcal{S}_\Theta$
 Partition columns of Λ into k_Λ blocks
 \triangleright Minimize over Λ
 Initialize $\Delta_\Lambda \leftarrow 0$
 for $z = 1$ to k_Λ **do**
 Compute $\Sigma_{C_z}, \mathbf{U}_{C_z}$, and Ψ_{C_z}
 for $r = 1$ to k_Λ **do**
 if $z \neq r$ **then**
 Identify columns $B_{zr} \subset C_r$ with active
 elements in Λ
 Compute $\Sigma_{B_{zr}}, \mathbf{U}_{B_{zr}}$, and $\Psi_{B_{zr}}$
 Update all active $(\Delta_\Lambda)_{ij}$ in (C_z, C_r)
 Update $\Lambda^+ \leftarrow \Lambda + \alpha \Delta_\Lambda$, where step size α is found
 with line search
 Partition columns of Θ into k_Θ blocks
 \triangleright Minimize over Θ
 for $r = 1$ to k_Θ **do**
 Compute Σ_{C_r} , and initialize $\mathbf{V} \leftarrow \Theta \Sigma_{C_r}$
 for row $i \in \{1, \dots, p\}$ if $I_\phi(\mathcal{S}(i, C_r))$ **do**
 Compute $(\mathbf{S}_{xx})_{ij}$ for non-empty rows j in V_{C_r}
 Update all active Θ_{ij} in (i, C_r)

with conjugate gradient method in $O(m_\Lambda K)$ time, where K is the number of conjugate gradient iterations. Then, Ψ_i can be obtained from $\frac{1}{n} \mathbf{R}^T \mathbf{R}_i$ in $O(nq)$ time, where $\mathbf{R} = \mathbf{X} \Theta \Sigma$.

In order to reduce cache misses, we perform block coordinate descent, where within each block, the columns of Σ are cached and re-used. Suppose we partition $\mathcal{N} = \{1, \dots, q\}$ into k_Λ blocks, $C_1, \dots, C_{k_\Lambda}$. We apply this partitioning to the rows and columns of Δ_Λ to obtain $k_\Lambda \times k_\Lambda$ blocks. We perform coordinate-descent updates in each block, updating all elements in the active set within that block. Let \mathbf{A}_{C_r} denote a q by $|C_r|$ matrix containing columns of \mathbf{A} that corresponds to the subset C_r . In order to perform coordinate-descent updates on (C_z, C_r) block of Δ_Λ , we need $\Sigma_{C_z}, \Sigma_{C_r}, \Psi_{C_z}$, and Ψ_{C_r} . Thus, we pick the smallest possible k_Λ such that we can store $2q/k_\Lambda$ columns of Σ and $2q/k_\Lambda$ columns of Ψ in memory. When updating the variables within block (C_z, C_r) of Δ_Λ , there are no cache misses once $\Sigma_{C_z}, \Sigma_{C_r}, \Psi_{C_z}$, and Ψ_{C_r} are computed and stored. After updating each $(\Delta_\Lambda)_{ij}$ to $(\Delta_\Lambda)_{ij} + \mu$, we maintain \mathbf{U}_{C_z} and \mathbf{U}_{C_r} by $\mathbf{U}_{it} \leftarrow \mathbf{U}_{it} + \mu \Sigma_{jt}, \mathbf{U}_{jt} \leftarrow \mathbf{U}_{jt} + \mu \Sigma_{it}, \forall t \in \{C_z \cup C_r\}$.




Figure 1: Schematic of block coordinate descent for Λ . The Λ of size $q = 9$ is updated for each of the k_Λ^2 blocks in turn with $k_\Lambda = 3$. Filled elements denote the parameters in the active set. The green arrows denote rows of Σ and Ψ that are computed once and reused while sweeping through a row of blocks. The red arrows denote cache misses and the corresponding columns of Σ and Ψ need to be recomputed.




Figure 2: Schematic of block coordinate descent for Θ . The Θ of size $p = 8, q = 6$ is updated for each of the $p \times k_\Theta$ blocks with $k_\Theta = 2$. Filled elements denote the parameters in the active set. Green arrows denote columns of Σ that are computed once and reused while sweeping through the column of p blocks. The red arrows denote cache misses for $(\mathbf{S}_{xx})_i$.

To go through all blocks, we update blocks $(C_z, C_1), \dots, (C_z, C_{k_\Lambda})$ for each $z \in \{1, \dots, k_\Lambda\}$. Since all of these blocks share Σ_{C_z} and Ψ_{C_z} , we pre-compute and store them in memory. When updating an off-diagonal block $(C_z, C_r), z \neq r$, we compute Σ_{C_r} and Ψ_{C_r} . In the worst case, overall Σ and Ψ will be computed k_Λ times.

4.1.2 Reducing Computational Cost Using Graph Clustering

In typical real-world problems, the graph structure of Λ will exhibit clustering, with an approximately block diagonal structure. We exploit this structure by choosing a partition $\{C_1, \dots, C_{k_\Lambda}\}$ that reduces cache misses. Within diagonal blocks (C_z, C_z) 's, once Σ_{C_z} and Ψ_{C_z} are computed, there are no cache misses. For off-diagonal blocks (C_z, C_r) 's, $r \neq z$, we have a cache miss only if some vari-

ables in $\{(\Delta_{\Lambda})_{ij} | i \in C_z, j \in C_r\}$ lie in the active set. We thus minimize the active set in off-diagonal blocks via clustering, following the strategy for sparse GGM estimation in [4]. In the best case, if all parameters in the active set appear in the diagonal blocks, Σ and Ψ are computed only once with no cache misses. We use the METIS [5] graph clustering library. Our method for updating Λ is illustrated in Figure 1.

4.2 Blockwise Optimization for Θ

4.2.1 Block Coordinate Descent Method

The coordinate descent update of Θ_{ij} requires $(S_{xx})_i$ and Σ_j to compute $(S_{xx})_i^T \mathbf{V}_j$, where $\mathbf{V}_j = \Theta \Sigma_j$. If $(S_{xx})_i$ and Σ_j are not already in the memory, it is a cache miss. Computing $(S_{xx})_i$ takes $O(np)$, which is expensive if we have many cache misses.

We propose a block coordinate descent approach for solving Eq. (6) that groups these computations to reduce cache misses. Given a partition of $\{1, \dots, q\}$ into k_{Θ} subsets, $C_1, \dots, C_{k_{\Theta}}$, we divide Θ into $p \times k_{\Theta}$ blocks, where each block comprises a portion of a row of Θ . We denote each block (i, C_r) , where $i \in \{1, \dots, p\}$. Since updating block (i, C_r) requires $(S_{xx})_i$ and Σ_{C_r} , we pick the smallest possible k_{Θ} such that we can store q/k_{Θ} columns of Σ in memory. While performing coordinate descent updates within block (i, C_r) of Θ , there are no cache misses, once $(S_{xx})_i$ and Σ_{C_r} are in memory. After updating each Θ_{ij} to $\Theta_{ij} + \mu$, we update \mathbf{V}_{C_r} by $\mathbf{V}_{it} \leftarrow \mathbf{V}_{it} + \mu \Sigma_{jt}, \forall t \in C_r$.

In order to sweep through all blocks, each time we select a $q \in \{1, \dots, k_{\Theta}\}$ and update blocks $(1, C_r), \dots, (p, C_r)$. Since all of these p blocks with the same C_r share the computation of Σ_{C_r} , we compute and store Σ_{C_r} in memory. Within each block, the computation of $(S_{xx})_i$ is shared, so we pre-compute and store it in memory, before updating this block. The full matrix of Σ will be computed once while sweeping through the full Θ , whereas in the worst case S_{xx} will be computed k_{Θ} times.

4.2.2 Reducing Computational Cost Using Row-wise Sparsity

We further reduce cache misses for $(S_{xx})_i$ by strategically selecting partition $C_1, \dots, C_{k_{\Theta}}$, based on the observation that if the active set is empty in block (i, C_r) , we can skip this block and forgo computing $(S_{xx})_i$. We therefore choose a partition where the active set variables are clustered into as few blocks as possible. Formally, we want to minimize $\sum_{i,r} I_{\phi}(\mathcal{S}_{(i,C_r)})$, where $I_{\phi}(\mathcal{S}_{(i,C_r)})$ is an indicator function that outputs 1 if the active set $\mathcal{S}_{(i,C_r)}$ within block (i, C_r) is not empty. We therefore perform graph clustering over the graph $G = (V, E)$ defined from the active set in Θ , where $V = \{1, \dots, q\}$ with one node for each column of Θ , and $E = \{(j, k) | \exists i : \Theta_{ij} \in \mathcal{S}_{\Theta}, \Theta_{ik} \in \mathcal{S}_{\Theta}\}$,

connecting two nodes j and k with an edge if both Θ_{ij} and Θ_{ik} are in the active set. This edge set corresponds to the non-zero elements of $\Theta^T \Theta$, so the graph can be computed quickly in $O(m_{\Theta}q)$.

We also exploit row-wise sparsity in Θ to reduce the cost of each cache miss. Every empty row in Θ corresponds to an empty row in $\mathbf{V} = \Theta \Sigma$. Because we only need elements in $(S_{xx})_i$ for the dot product $(S_{xx})_i^T \mathbf{V}_j$, we skip computing the k th element of $(S_{xx})_i$ if the k th row of Θ is all zeros. Our strategy for updating Θ is illustrated in Figure 2.

Our method is summarized in Algorithm 2. See Appendix for analysis of the computational cost.

4.3 Parallelization

The most expensive computations in our algorithm are embarrassingly parallelizable, allowing for further speedups on machines with multiple cores. Throughout the algorithm, we parallelize matrix and vector multiplications. In addition, for block-wise Λ updates, we compute multiple columns of Σ_{C_z} and Ψ_{C_z} as well as multiple columns of Σ_{C_r} and Ψ_{C_r} for multiple cache misses in parallel, running multiple conjugate gradient methods in parallel. For block-wise Θ updates, we compute multiple columns of Σ in parallel before sweeping through blocks and perform a parallel computation within each cache miss, computing elements within each $(S_{xx})_i$ in parallel.

5 EXPERIMENTS

We compare the performance of our methods with the existing state-of-the-art Newton coordinate descent algorithm, using synthetic and real-world genomic datasets. All methods were implemented in C++ with parameters represented in sparse matrix format. All experiments were run on 2.6GHz Intel Xeon E5 machines with 8 cores and 104 Gb RAM, running Linux OS. We run the Newton coordinate descent and alternating Newton coordinate descent algorithms as a single thread job on a single core. For our alternating Newton block coordinate descent method, we run it on a single core and with parallelization on 8 cores.

5.1 Synthetic Data Experiments

We compare the different methods on two sets of synthetic datasets, one for chain graphs and another for random graphs with clustering for Λ , generated as follows. For chain graphs, the true sparse parameters Λ is set with $\Lambda_{i,i-1} = 1$ and $\Lambda_{i,i} = 2.25$ and the true Θ is set with $\Theta_{i,i} = 1$. We perform one set of chain graph experiments with $p = q$, and another set with an additional q irrelevant features unconnected to any outputs, so that $p = 2q$. For random graphs with clustering, following the procedure in [4] for generating a GGM, we set the true Λ to a graph with




Figure 3: Comparison of scalability on chain graphs. We vary p and q , where (a) $p = q$ and (b) $p = 2q$. The Newton coordinate descent and alternating Newton coordinate descent methods could not be run beyond the problem sizes shown due to memory constraint. (c) Convergence when $q = 20,000$ and $p = 40,000$.




Figure 4: Comparison of scalability on random graphs with clustering. (a) Varying p with q fixed at 10,000. (b) Varying q with p fixed at 40,000. (c) Active set size versus time with $p = 20,000$ and $q = 10,000$.

clusters of nodes of size 250 and with 90% of edges connecting randomly-selected nodes within clusters. We set the number of edges so that the average degree of each node is 10, with edge weights set to 1. We then set the diagonal values so that Λ is positive definite. To set the sparse patterns for Θ , we randomly select $100\sqrt{p}$ input variables as having edges to at least one output and distribute total $10q$ edges among those selected inputs to influence randomly selected outputs. We set the edge weights of Θ to 1.

Then, we draw samples from the CGGM defined by these true Λ and Θ . We generate datasets with $n = 100$ samples for the chain graphs and $n = 200$ samples for random graphs with clustering. We choose λ_Λ and λ_Θ so that the number of edges in the estimated Λ and Θ is close to ground truth. Following the strategy used in GGM estimation [4], we use the minimum-norm subgradient of the objective as our stopping criterion: $\|\text{grad}^S(\Lambda_t, \Theta_t)\|_1 < 0.01(\|\Lambda\|_1 + \|\Theta\|_1)$.

We compare the scalability of the different methods on chain graphs of different sizes. We show the computation times for datasets with $p = q$ in Figure 3(a) and for datasets with $p = 2q$ with q additional irrelevant features in Figure 3(b). For large problems, computation times are not shown for Newton coordinate descent and alternating Newton coordinate descent methods because they could not complete the optimization with limited memory. In addition, for large problems, alternating block coordinate descent was terminated after 60 hours of computation. We provide results

on varying the sample size n in Appendix. In Figure 3(c), using the dataset with $p = 40,000$ and $q = 20,000$, we plot the suboptimality in the objective $f - f^*$ over time, where f^* is obtained by running alternating Newton coordinate descent algorithm to numerical precision. Our new methods converge substantially faster than the previous approach, regardless of desired accuracy level. We notice that as expected from the convexity of the optimization problem, all algorithms converge to the global optimum and find nearly identical parameter estimates.

In Figure 4, we compare scalability of different methods for random graphs with clustering. In Figure 4(a), we vary p , while setting q to 10,000. In Figure 4(b), we vary q , fixing p to 40,000. Similar to the results from chain graph, for larger problems, Newton coordinate descent and alternating Newton coordinate descent methods ran out of memory and alternating block coordinate descent was terminated after 60 hours. For all problem sizes, our alternating Newton coordinate descent algorithm significantly reduces the computation time of the previous method, the Newton coordinate descent algorithm. This gap in the computation time increases for larger problems. In Figure 4(c), we compare the convergence in sparsity pattern for the different methods as measured by the active set size, for $p = 20,000$ and $q = 10,000$. All our methods recover the optimal sparsity pattern much more rapidly than the previous approach.

Figures 3 and 4 show that our alternating Newton block coordinate descent can run on much larger problems than

Table 1: Computation time in hours on genomic dataset. ‘*’ indicates running out of memory.

p	q	$\ \Lambda^*\ _0$	$\ \Theta^*\ _0$	Newton CD	Alternating Newton CD	Alternating Newton BCD
34,249	3,268	34,914	28,848	22.0	0.51	0.24
34,249	10,256	86,090	103,767	> 50	2.4	2.3
442,440	3,268	26,232	30,482	*	*	11




Figure 5: Speedup with parallelization for alternating Newton block coordinate descent.

any other methods, while those methods without block coordinate descent run out of memory. For example, in Figure 4(a) alternating Newton block coordinate descent could handle problems with one million inputs, while without block-wise optimization it ran out of memory when $p > 100,000$. We also notice that on a single core, the alternating Newton block coordinate descent is slightly slower than the same method without block-wise optimization because of the need to recompute Σ and S_{xx} . However, it is still substantially faster than the previous method.

Finally, we evaluate the parallelization scheme for our alternating Newton block coordinate descent method on multi-core machines. Given a dataset generated from cluster graph with $p = 40,000$ and $q = 20,000$, in Figure 5, we show the folds of speedup for different numbers of cores with respect to a single core. We obtained about 7 times speedup on a 8-core machine with 104Gb RAM, and about 12 times speedup on a 16-core machine with 28Gb RAM. In general, we observe greater speedup on larger problems and for random graphs, because such problems tend to have more cache misses that can be computed in parallel.

5.2 Genomic Data Analysis

We compare the different methods on a genomic dataset. The dataset consists of genotypes for 442,440 single nucleotide polymorphisms (SNPs) and 10,256 gene expression levels for 171 individuals with asthma, after removing genes with variance < 0.01 . We fit a sparse CGGM using SNPs as inputs and expressions as outputs to model a gene network influenced by SNPs. We also compared the methods on a smaller dataset of 34,249 SNPs from chromosome 1 and expression levels for 3,268 genes with variance > 0.1 . As typically sparse model structures are of interests in this type of analysis, we chose regularization parameters so that the number of non-zero entries in each of Θ and




Figure 6: Convergence results on genomic dataset. (a) Suboptimality and (b) active set size over time.

Λ at convergence was approximately 10 times the number of genes. The alternating Newton block coordinate descent was run on a 8-core machine with parallelization.

The computation time of different methods are provided in Table 1. On the largest problem, the previous approach could not run due to memory constraint, whereas our block coordinate descent converged in around 11 hours. We also compare the convergence of the different methods on the dataset with 34,249 SNPs and 3,268 gene expressions in Figure 6, and find that our methods provide vastly superior convergence than the previous method.

6 CONCLUSION

In this paper, we addressed the problem of large-scale optimization for sparse CGGMs. We proposed a new optimization procedure, called alternating Newton coordinate descent, that reduces computation time by alternately optimizing for the two sets of parameters Λ and Θ . Further, we extended this with block-wise optimization so that it can run on any machine with limited memory.

Acknowledgements

This material is based upon work supported by an NSF CAREER Award No. MCB-1149885, Sloan Research Fellowship, and Okawa Foundation Research Grant.

References

- [1] L. Armijo et al. Minimization of functions having lipschitz continuous first partial derivatives. *Pacific Journal of mathematics*, 16(1):1–3, 1966.
- [2] J. Friedman, T. Hastie, and R. Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, 2008.
- [3] C.-J. Hsieh, I. S. Dhillon, P. K. Ravikumar, and M. A. Sustik. Sparse inverse covariance matrix estimation using quadratic approximation. In *Advances in Neural Information Processing Systems 24*, pages 2330–2338. 2011.
- [4] C.-J. Hsieh, M. A. Sustik, I. S. Dhillon, P. K. Ravikumar, and R. Poldrack. BIG & QUIC: Sparse inverse covariance estimation for a million variables. In *Advances in Neural Information Processing Systems 26*, pages 3165–3173. 2013.
- [5] G. Karypis and V. Kumar. Metis-unstructured graph partitioning and sparse matrix ordering system, version 2.0. 1995.
- [6] A. Lafferty, J. McCallum and F. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the 18th International Conference on Machine Learning*, 2001.
- [7] B. Ng, R. Abugharbieh, G. Varoquaux, J. B. Poline, and B. Thirion. Connectivity-informed fMRI activation detection. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2011*, pages 285–292. Springer, 2011.
- [8] K. Sohn and S. Kim. Joint estimation of structured sparsity and output structure in multiple-output regression via inverse-covariance regularization. In *Proceedings of the 15th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 16. JMLR W&CP, 2012.
- [9] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of Royal Statistical Society, Series B*, 58(1):267–288, 1996.
- [10] M. Wytock and J. Kolter. Sparse Gaussian conditional random fields: algorithms, theory, and application to energy forecasting. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28. JMLR W&CP, 2013.
- [11] X.-T. Yuan and T. Zhang. Partial Gaussian graphical model estimation. *IEEE Transactions on Information Theory*, 60(3):1673–1687, 2014.