Technical University of Munich

Department of Mathematics

Undirected Structures in Graphical Continuous Lyapunov Models

Master's Thesis

von

Patrick Ruang Schmidhalter

Supervisor:	Prof. Dr. Mathias Drton
Advisor:	Philipp Dettling
Submission Date:	17.04.2023

I hereby declare that this thesis is my own work and that no other sources have been used except those clearly indicated and referenced.

P. Schmidwalter

Munich, 17.04.2023

Acknowledgements

I would like to thank Professor Mathias Drton for giving me the opportunity to write my thesis under his supervision. I am very grateful for the many insightful discussions we had, and the suggestions he provided throughout the course of this thesis. His advice and feedback were a great guidance.
I would also like to extend my gratitude to my advisor Philipp Dettling for his great help, time, and patience. His expertise and comments enabled me to gain a deeper understanding of the material.
I am deeply thankful to my parents for their support during my years of study and the many opportunities they provided me. Last but not least, I would like to thank my brothers, friends, and girlfriend for all their support and encouragement.

German Abstract

Graphische Lyapunov Modelle sind ein neuartiger Ansatz, um korrelierte multivariate Daten statistisch zu modellieren. Dabei nehmen wir an, dass jede Stichprobe eine Momentaufnahme des multivariaten Ornstein-Uhlenbeck Prozess im Gleichgewicht ist. Letzterer besitzt eine eindeutige Kovarianzmatrix, falls diese eine Lösung der stetigen Lyapunov Gleichung ist, welche von einer stabilen Driftmatrix parametrisiert wird.
Jede dünnbesetzte Driftmatrix bestimmt die Adjazenz eines ungerichteten Graphen. Wir verwenden das group lasso, um eine dünnbesetzte ungefähre Lösung für die stetige Lyapunov Gleichung zu bestimmen für eine gegebene Kovarianzmatrix. Darüberhinaus ermitteln wir hinreichende Bedingungen um die Adjazenz eines Graphen korrekt zu bestimmen zu können. Dazu führen wir sogenannte duale Normen ein und adaptieren die primal-dual witness technique für das group lasso. Eine der hinreichenden Bedingungen, die wir aufzeigen, ist eine sogenannte group irrepresentability condition, welche wir genauer untersuchen.
Im zweiten Teil dieser Arbeit widmen wir uns algebraischen Fragestellungen bezüglich graphischen Lyapunov Modellen. Wir untersuchen dabei die identifiability und Kovarianzäquivalenz. Für letztere beweisen wir, dass Bäume in niedrigen Dimensionen immer unterschiedliche statistische Modelle induzieren.
Im letzten Teil führen wir eine Reihe von numerischen Experimenten durch, um die Performance von dem lasso und group lasso für gerichtete und ungerichtete Graphenschätzung zu vergleichen.

English Abstract

Graphical continuous Lyapunov models present a novel approach to statistically model correlated multivariate data. In this setting, every observation is treated as a one-time cross-sectional snapshot of the multivariate Ornstein-Uhlenbeck process in equilibrium. The preceding process has a unique covariance matrix obtained as a solution to the continuous Lyapunov equation specified by a stable drift matrix.
The sparsity pattern of the drift matrix is related to the support of an undirected graph. We propose the group lasso as a regularization approach to model selection, i.e., we find a sparse approximate solution to the continuous Lyapunov equation for a given covariance matrix. Moreover, we derive sufficient conditions for consistent correct recovery by introducing the notion of dual norms and adapting the primal-dual witness method for the group lasso. One of the key assumptions will be a group irrepresentability condition, which we will further investigate.
In the second part of the thesis, we review several algebraic questions related to graphical continuous Lyapunov models. We start by providing a theory for the identifiability problem. In addition, we study covariance equivalence for undirected graphs, where we derive that low-dimensional trees will always induce different statistical models.
Lastly, we conduct a series of numerical experiments comparing the performance of the lasso and group lasso for both directed and undirected structure recovery.

Contents

1 Introduction

Graphical models allow us to better understand and interpret complex dependencies in multivariate data. Moreover, they enable us to establish cause-effect relationships or, in other words causality, (Pearl (2009); Peters et al. (2017); Spirtes et al. (2000)). In particular, this is done using structural equation modeling (SEM), where we consider each variable X to be a function of parent variables and independent noise ϵ, i.e.,

$$X = \Lambda^{\mathsf{T}} X + \epsilon$$

with Λ being a matrix of unknown parameters. Assuming $\mathrm{Var}(\epsilon) = \Omega$ the covariance matrix Σ of X satisfies

$$(I - \Lambda)^{\mathsf{T}} \Sigma (I - \Lambda) = \Omega. \tag{1.1}$$

For an acyclic graph, SEMs provide a parametrization of the observables from a directed acyclic model with possible latent variables (Richardson and Spirtes (2002)). These so-called directed acyclic graphs (DAG) can be interpreted, and their parameters can be easily estimated due to statistically favorable density factorization properties (Maathuis et al. (2019)).

By contrast, if we consider the cyclic case, i.e., we allow the graph to contain directed cycles, the statistical model becomes more involved. In this setting, cycles represent feedback loops (Bongers et al. (2021)). While we can still define the statistical model by solving structural equations, a number of issues arise. For instance, the presence of directed cycles impedes density factorizations, thus complicating the computation of maximum likelihood estimates (Drton et al. (2019)). Furthermore appropriate model selection is more difficult for cyclic graphs (Améndola et al. (2020); Richardson (1996)). Consequently, the interpretation of these models is not as straightforward as in the acyclic case. Fisher (1970) proposes an interpretation based on data that are time averages. Mooji et al. (2013) and Bongers and Mooji (2018) relate the equilibrium states of differential equations to structural equations.

For some continuous-time stochastic processes, the equilibrium covariance matrix does not have the simple graphical representation provided by equation (1.1) given above. Rather we require a parametrization corresponding to the graphical representation of the dynamics of the process. Consider the p-dimensional Ornstein-Uhlenbeck process defined as the solution to the stochastic differential equation

$$dX_t = M(X_t - a)dt + DdW_t, \tag{1.2}$$

where $M, D \in \mathbb{R}^{p\times p}$ are non-singular parameter matrices, $a \in \mathbb{R}^p$ and W_t is a standard Brownian motion in $\mathbb{R}^p$. Assuming M is stable, i.e., all real parts of the eigenvalues are strictly negative, X_t has a stationary distribution that is multivariate normal with mean vector a and positive definite covariance matrix Σ (Fitch (2019), Theorem 2). Σ is defined to be the unique matrix that solves the continuous Lyapunov equation

$$M\Sigma + \Sigma M^{\top} + C = 0 \tag{1.3}$$

with $C := DD^{\top}$. Fitch (2019) and Varando and Hansen (2020) now assume the data to be generated by a Ornstein-Uhlenbeck process in equilibrium. Namely, we have a sample $X_1, \ldots, X_n \in \mathbb{R}^p$ from the aforementioned multivariate process, where X_i represents a single cross-sectional observation of the i-th process in equilibrium. Thus X_i is multivariate normal with expectation a and covariance Σ. The focal point of graphical continuous Lyapunov models is the drift matrix M as it quantifies temporal cause-effect relations among the coordinates of the Ornstein-Uhlenbeck process X_t.

Estimating the support or sparsity pattern of the drift matrix M corresponds to estimating the mixed graph associated with each graphical continuous Lyapunov model. Dettling et al. (2022) proposes the lasso as a structure recovery method. They establish sufficient conditions for correct recovery. The crucial assumption here is the irrepresentability condition.

For this thesis, we take a step back and restrict ourselves to estimating the undirected structure of a mixed graph, i.e., the skeleton. We apply the group lasso to the continuous Lyapunov equation and develop conditions for consistent support recovery. We review existing literature for consistent support recovery of the group lasso and develop our own approach based on dual norms and the primal-dual witness method. Specifically, we introduce the group irrepresentability condition as a sufficient condition and study its properties.

Besides estimating the underlying graph, we will also study several algebraic questions concerning the graphical continuous Lyapunov model. Specifically, whether the graphical continuous Lyapunov model is uniquely determined by the multivariate distribution of the observations. This question is referred to as identifiability, and we will provide a review of existing results on this topic (Dettling et al. (2022)). Secondly, we examine whether two different undirected graphs may induce the same graphical continuous Lyapunov model, i.e., we study the notion of covariance equivalence for undirected graphs.

Structure of the thesis. In Chapter 2, we review the Kronecker product and related theory necessary to formulate graphical continuous Lyapunov models. Chapter 3 deals with undirected structure estimation via the group lasso and develops sufficient conditions for consistent sparsity pattern estimation. In Chapter 4, we discuss different notions of identifiability and covariance equivalence for undirected graphs. Lastly, in Chapter 5, we perform a series of numerical studies to gain more insight into the behavior of the lasso and group lasso for structure estimation.

Notation

- Let $p \in \mathbb{N}$. Then $[p] = \{1, \dots, p\}$.
- $\mathbb{C}^{p \times q}$ and $\mathbb{R}^{p \times q}$ denotes the set of $p \times q$ matrices in $\mathbb{C}$ or $\mathbb{R}$, respectively.
- For a $p \times q$ matrix A, we write $A_k = A_{\cdot k}$ to select k-th column of A and $A_{l \cdot}$ to select the l-th row of A.
- Let A be a a matrix, then $\sigma(A)$ denotes the set of all eigenvalues of A.
- For $v \in \mathbb{R}^p$ and $b \in [1, \infty]$ the ℓ_b-norm of v is $\|v\|_b = (\sum_{i=1}^p |v_i|^b)^{1/b}$ with $\|v\|_\infty = \max_{1 \leqslant i \leqslant p} |v_i|$.
  Let $A \in \mathbb{R}^{p \times q}$. The associated operator norm is denoted by $|||A|||_b = \max_{\|x\|_b = 1} \|Ax\|_b$. $\|A\|_F = (\sum_{i=1}^p \sum_{j=1}^q |A_{ij}|^2)^{1/2}$ denotes the Frobenius norm.
- We use the abbreviations LHS and RHS to refer to the left-hand side or right-hand side of an equation or inequality.

2 Graphical Continuous Lyapunov Models

2.1 Kronecker product and linear matrix equations

The Kronecker product is essential for the formulation and study of graphical continuous Lyapunov models (GCLM). In the following, we will provide interesting properties and key facts about the Kronecker product that will allow us to derive important results with regard to GCLMs. The main reference for these results is Horn and Johnson (1991, Chapter 4).

Definition 2.1 (*Kronecker product*)**.** The Kronecker product of $A = (a_{ij}) \in \mathbb{C}^{m\times n}$ and $B = (b_{ij}) \in \mathbb{C}^{p\times q}$ is denoted by $A \otimes B$ and is defined to be the block matrix

$$A \otimes B := \begin{pmatrix} a_{11}B & \cdots & a_{1n}B \\ \vdots & \ddots & \vdots \\ a_{m1}B & \cdots & a_{mn}B \end{pmatrix} \in \mathbb{C}^{mp\times nq}.$$

Note that $A\otimes B \neq B\otimes A$ in general. A few basic properties of the Kronecker product include:

$$\begin{aligned}
&(\alpha A) \otimes B = A \otimes (\alpha B) \text{ for all } \alpha \in \mathbb{C} \\
&(A + B) \otimes C = A \otimes C + B \otimes C \text{ for all } A, B \in \mathbb{C}^{m\times n} \text{ and } C \in \mathbb{C}^{p\times q} \\
&A \otimes (B + C) = A \otimes C + B \otimes C \text{ for all } A \in \mathbb{C}^{m\times n} \text{ and } B, C \in \mathbb{C}^{p\times q} \\
&(A \otimes B) \otimes C = A \otimes (B \otimes C) \text{ for all } C \in \mathbb{C}^{r\times s} \\
&(A \otimes B)^{\mathsf{T}} = A^{\mathsf{T}} \otimes B^{\mathsf{T}} \\
&(A \otimes B)^* = A^* \otimes B^*
\end{aligned}$$

Another useful fact about the Kronecker product is the mixed-product property, which combines ordinary matrix multiplication and the Kronecker product.

Lemma 2.2. (*mixed-product property*) Let $A \in \mathbb{C}^{m\times n}, B \in \mathbb{C}^{p\times q}, C \in \mathbb{C}^{n\times k}$ and $D \in \mathbb{C}^{q\times r}$. Then

$$(A \otimes B)(C \otimes D) = AC \otimes BD.$$

Proof. We define $A = (a_{il})$ and $C = (c_{lj})$, then $A \otimes B = (a_{il}B)$ and $C \otimes D = (c_{lj}D)$. We obtain the following expression for the i, j block of $(A \otimes B)(C \otimes D)$ by writing the matrix multiplication explicitly as

$$\sum_{l=1}^{n} a_{il} B c_{lj} D = \left(\sum_{l=1}^{n} a_{il} c_{lj} \right) BD = (AC)_{ij} BD.$$

Note that this is exactly the i, j block of $AC \otimes BD$, which concludes the proof. □

The Kronecker product allows for convenient representations of linear matrix equations. For this purpose, we introduce the following definition.

Definition 2.3. For a matrix $A = (a_{ij}) \in \mathbb{C}^{m \times n}$ we associate the vector $\text{vec}(A) \in \mathbb{C}^{mn}$ defined by

$$\text{vec}(A) := (a_{11}, \ldots, a_{m1}, a_{12}, \ldots, a_{m2}, \ldots, a_{1n}, \ldots, a_{mn})^{\mathsf{T}}.$$

We call $\text{vec}(A)$ the *vectorization* of A.

The next lemma provides the connection between linear matrix equations and the Kronecker product.

Lemma 2.4. Let $A \in \mathbb{C}^{m \times n}, B \in \mathbb{C}^{p \times q}$ and $X \in \mathbb{C}^{n \times p}$. Then the following equation holds true

$$\text{vec}(AXB) = (B^{\mathsf{T}} \otimes A)\text{vec}(X).$$

Proof. For $k = 1, \ldots, q$ we get

$$\begin{aligned} (AXB)_k = A(XB)_k = AXB_k &= (b_{1k}A, b_{2k}A, \cdots, b_{pk}A)\text{vec}(X) \\ &= (B_k^{\mathsf{T}} \otimes A)\text{vec}(X). \end{aligned}$$

Thus,

$$\text{vec}(AXB) = \begin{pmatrix} B_1^{\mathsf{T}} \otimes A \\ \vdots \\ B_q^{\mathsf{T}} \otimes A \end{pmatrix} \text{vec}(X).$$

Note that the above Kronecker products are just $B^{\mathsf{T}} \otimes A$ since the transpose of a column of B is a row of B^{T}. Finally, we obtain

$$\text{vec}(AXB) = (B^{\mathsf{T}} \otimes A)\text{vec}(X).$$

□

Example 2.5. Consider the continuous Lyapunov equation with M, Σ symmetric, $C \in \mathbb{R}^{p \times p}$ and

$$M\Sigma + \Sigma M^{\mathsf{T}} + C = 0.$$

This can be equivalently expressed as

$$\text{vec}(M\Sigma) + \text{vec}(\Sigma M^{\mathsf{T}}) + \text{vec}(C) = 0.$$

Applying Lemma 2.4 to the first two summands and solving for Σ, i.e. $X = \Sigma$, yields

$$((M \otimes I_p) + (I_p \otimes M))\mathrm{vec}(\Sigma) + \mathrm{vec}(C) = 0. \tag{2.1}$$

Alternatively, we could also solve for M

$$(\Sigma \otimes I_p)\mathrm{vec}(M) + (I_p \otimes \Sigma)\mathrm{vec}(M^\top) + \mathrm{vec}(C) = 0. \tag{2.2}$$

For a more compact representation of equation (2.2) in the Example given above, it would be useful to represent $\mathrm{vec}(M^\top)$ in terms of $\mathrm{vec}(M)$. The next theorem provides exactly this.

Theorem 2.6. *There exists a unique matrix $K^{(m,n)} \in \mathbb{C}^{m\times n}$ such that*

$$vec(X^\top) = K^{(m,n)} vec(X) \text{ for all } X \in \mathbb{C}^{m\times n}.$$

$K^{(m,n)}$ is only dependent on the dimensions m and n and is given by

$$K^{(m,n)} := \sum_{i=1}^{m}\sum_{j=1}^{n} E_{ij} \otimes E_{ij}^\top = (E_{ij})_{\substack{i=1,\dots,m\\ j=1,\dots,n}} \tag{2.3}$$

with $E_{ij} \in \mathbb{C}^{m\times n}$ having entry 1 at position (i,j) and all other entries being 0. Moreover, $K^{(m,n)}$ is a permutation matrix with $K^{(m,n)} = (K^{(n,m)})^\top = (K^{(n,m)})^{-1}$.

Proof. see proof of Theorem 4.3.8 in Horn and Johnson (1991). □

We also include the following corollary.

Corollary 2.7. Let $K^{(p,m)} \in \mathbb{C}^{pm\times pm}$ and $K^{(n,q)} \in \mathbb{C}^{nq\times nq}$ be permutation matrices. Then

$$B \otimes A = K^{(p,m)}(A \otimes B)K^{(n,q)}. \tag{2.4}$$

Proof. see proof of Corollary 4.3.10 in Horn and Johnson (1991). □

Example 2.5 (continued)**.** (2.2) can now be written as

$$((\Sigma \otimes I_p) + (I_p \otimes \Sigma)K^{(m,n)})\mathrm{vec}(M) + \mathrm{vec}(C) = 0. \tag{2.5}$$

The matrices we obtain in Example 2.5 are examples of a more general structure, so-called *Kronecker sums*, which have the following form $(A \otimes I_n) + (I_m \otimes B)$ with $A \in \mathbb{C}^{m\times m}$ and $B \in \mathbb{C}^{n\times n}$. The following two theorems will allow us to compute the eigenvalues of Kronecker sums and discuss the solvability of linear matrix equations.

Theorem 2.8. *Let $A \in \mathbb{C}^{m\times m}$ and $B \in \mathbb{C}^{n\times n}$ with eigenvalues $\lambda \in \sigma(A) = \{\lambda_1, \dots, \lambda_m\}$ and $\mu \in \sigma(B) = \{\mu_1, \dots, \mu_n\}$, respectively. Then $\lambda+\mu$ is an eigenvalue of $(A\otimes I_n)+(I_m\otimes B)$. In particular, $\sigma((A \otimes I_n) + (I_m \otimes B)) = \{\lambda_i + \mu_j : i = 1, \dots, m \text{ and } j = 1, \dots, n\}$.*

Proof. Using Schur's decomposition, we can find two unitary matrices $U \in \mathbb{C}^{m\times m}$ and $V \in \mathbb{C}^{n\times n}$ such that $U^*AU = \Delta_A$ and $V^*BV = \Delta_B$ are upper triangular. In particular, note that $W = V \otimes U \in \mathbb{C}^{mn\times mn}$ is unitary, which can be easily verified using the mixed-induct property.
Subsequent application of the mixed-induct property yields

$$\begin{aligned} W^*(A\otimes I_n)W &= (V^*\otimes U^*)(A\otimes I_n)(V\otimes U) = (V^*A\otimes U^*)(V\otimes U) = V^*AV\otimes U^*U \\ &= \Delta_A \otimes I_n = \begin{pmatrix} \lambda_1 I_n & & * \\ & \ddots & \\ 0 & & \lambda_m I_n \end{pmatrix}. \end{aligned}$$

For the last equality, recall that the eigenvalues of A are the diagonal entries of Δ_A.
Analogously, we obtain

$$W^*(I_m \otimes B)W = I_m \otimes \Delta_B = \begin{pmatrix} \Delta_B & & 0 \\ & \ddots & \\ 0 & & \Delta_B \end{pmatrix}$$

(m copies of Δ_B on the diagonal).
Consequently, we can write the Kronecker sum as

$$W^*[(A\otimes I_n) + (I_m \otimes B)]W = (\Delta_A \otimes I_n) + (I_m \otimes \Delta_B)$$

The RHS of the above equation forms an upper triangular matrix with the eigenvalues of the Kronecker sum on its diagonal. Closer inspection reveals that each diagonal entry of Δ_A is paired with all diagonal entries of Δ_B. Thus, noting that similar matrices have the same eigenvalues, we can infer that

$$\sigma((A\otimes I_n) + (I_m \otimes B)) = \{\lambda_i + \mu_j : i = 1, \dots, m \text{ and } j = 1, \dots, n\}.$$

□

Theorem 2.9. *Let $A \in \mathbb{C}^{m\times m}$ and $B \in \mathbb{C}^{n\times n}$. The equation $AX + XB = C$ has a unique solution $X \in \mathbb{C}^{m\times n}$ for each $C \in \mathbb{C}^{m\times n}$ if and only if $\sigma(A) \cap \sigma(-B) = \varnothing$.*

Proof. Lemma 2.4 states that

$$AX + XB = C$$

is equivalent to

$$((A^\mathsf{T} \otimes I_n) + (I_m \otimes B))\mathrm{vec}(X) = \mathrm{vec}(C)$$

Note that A^T has the same eigenvalues as A and thus according to Theorem 2.8 $(A^\mathsf{T} \otimes I_n) + (I_m \otimes B)$ has a zero eigenvalue if and only if $\sigma(A) \cap \sigma(-B) \neq \varnothing$. Now recall that the determinant of a matrix can be expressed as the product over all its eigenvalues, hence $(A^\mathsf{T} \otimes I_n) + (I_m \otimes B)$ is non-singular if and only if $\sigma(A) \cap \sigma(-B) = \varnothing$. □

$\sigma(A) \cap \sigma(-B) = \varnothing$ is equivalent to saying that the sum of any two eigenvalues from A and B is nonzero.

Example 2.5 (continued)**.** The eigenvalues of $(M \otimes I_p) + (I_p \otimes M)$ are sums of pairs of eigenvalues of M. The solution to (2.1) is thus unique if and only if the sum of any two eigenvalues of M is nonzero.

2.2 Definition of graphical continuous Lypanunov models

We will study models of covariance matrices Σ given as solutions to the Lyapunov equation (1.3). As we have seen in Example 2.5, this can be equivalently expressed as the linear equation

$$((M \otimes I_p) + (I_p \otimes M))\mathrm{vec}(\Sigma) + \mathrm{vec}(C) = 0. \tag{2.6}$$

In order for (2.6) to have a unique solution, we require the sum of any two eigenvalues of M to be nonzero. We denote the unique solution by $\Sigma(M, C)$.

We introduce a few more definitions to describe the solutions of (1.3). Firstly, let $\mathrm{Mat}_0(p)$ be the set of all $p \times p$ matrices that do not have eigenvalues summing to zero, i.e.

$$\mathrm{Mat}_0(p) := \{X \in \mathbb{R}^{p \times p} : \sigma(X) \cap \sigma(-X) = \varnothing\}.$$

Furthermore, we denote the set of all symmetric matrices by

$$\mathrm{Sym}(p) := \{X \in \mathbb{R}^{p \times p} : X^\top = X\}$$

and the set of all stable matrices, that is, all matrices whose eigenvalues have strictly negative parts, is denoted by

$$\mathrm{Stab}(p) := \{X \in \mathbb{R}^{p \times p} : \mathrm{Re}(\sigma(X)) < 0\}.$$

Note that $\mathrm{Stab}(p) \subseteq \mathrm{Mat}_0(p)$. Lastly, the cone of all positive definite matrices is given by

$$\mathrm{PD}(p) := \{X \in \mathbb{R}^{p \times p} : v^\top X v > 0 \text{ for all } v \in \mathbb{R}^p\}.$$

We can associate the sparsity patterns of M and C to a mixed graph.

Definition 2.10 (*mixed graph*)**.** A mixed graph $G = (V, E)$ is a graph with vertex set V and edge set E containing directed as well as bidirected edges. In particular, we allow for self-loops and multiple edges between two nodes.

A pair of matrices $(M, C) \in \mathrm{Mat}_0(p) \times \mathrm{Sym}(p)$ are called *compatible* with a mixed graph G if $M_{ji} \neq 0$ implies $i \to j \in E$ and $C_{ij} \neq 0$ implies $i \leftrightarrow j \in E$. We denote the set of G-compatible matrix pairs by $\Xi_G \subseteq \mathrm{Mat}_0(p) \times \mathrm{Sym}(p)$ and $\Theta_G = \Xi_G \cap (\mathrm{Stab}(p) \times \mathrm{PD}(p))$.

Example 2.11. Consider the following pair of matrices

$$M = \begin{pmatrix} -1 & 1 & 0 & 0 & 0 \\ -1 & 0 & 0.2 & 0 & 0 \\ 0 & 0 & -1 & -0.5 & 0 \\ 0 & 0 & 0 & -1 & 1 \\ 0 & 0 & 1 & 0 & -1 \end{pmatrix} \qquad \text{and} \qquad C = I_5.$$

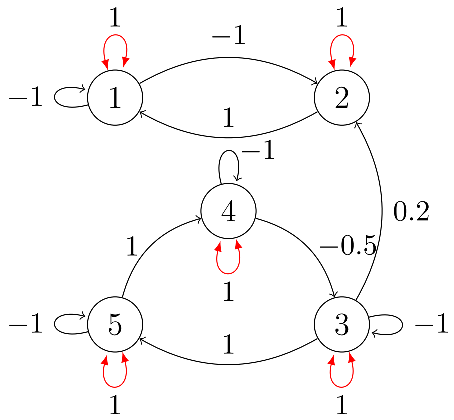

Figure 2.1: Compatible graph $\mathcal{G}$ with edge weights for (M, C) given in Example 2.11, where the black edges are specified by M and the red edges by C.

Note that M has eigenvalues $-1.79, -0.60 \pm 0.69i$ and $-0.50 \pm 0.87i$, i.e., all real parts are strictly negative, and thus M is stable. (M, C) are compatible with the mixed graph G as shown in Figure 2.1. Thus $(M, C) \in \Theta_{\mathcal{G}}$.

For a mixed graph G the map $(M, C) \mapsto \Sigma(M, C)$ is well defined on Ξ_G with image $\Sigma(\Xi_G)$ in $\mathrm{Sym}(p)$. Moreover, the image of Θ_G denoted by $\Sigma(\Theta_G)$ is in $\mathrm{PD}(p)$, which follows from Theorem 2.12.

Theorem 2.12. *Let $(M, C) \in \Theta_G$, then*

$$\Sigma(M, C) = \lim_{s\to\infty} \Sigma(s) = \int_0^\infty e^{uM} C e^{uM^T} du, \tag{2.7}$$

where $\Sigma(s) := \int_0^s e^{uM} C e^{uM^T} du$. Moreover, $\Sigma(M, C) \in \mathrm{PD}(p)$.

Proof. Note that by Theorem 2.9, the stability of M guarantees that the Lyapunov equation (1.3) has a unique solution. Moreover, stability also implies that the improper integral in (2.7) is convergent. We can prove the equality by verifying that the integral on the RHS of (2.7) is indeed a correct solution to the Lyapunov equation.

$$\begin{aligned} M \lim_{s\to\infty} \Sigma(s) + \lim_{s\to\infty} \Sigma(s) M^\mathsf{T} &= \int_0^\infty M e^{uM} C e^{uM^\mathsf{T}} + e^{uM} C e^{uM^\mathsf{T}} M^\mathsf{T} du \\ &= \int_0^\infty \frac{d}{du} e^{uM} C e^{uM^\mathsf{T}} du = -C \end{aligned}$$

Let $x \in \mathbb{R}^p \backslash \{0\}$, then we obtain the following expression using 2.7

$$x^\mathsf{T} \Sigma(B, C) x = \int_0^\infty x^\mathsf{T} e^{uM} C e^{uM^T} x du = \int_0^\infty (e^{uM^\mathsf{T}} x)^\mathsf{T} C (e^{uM^\mathsf{T}} x) du > 0.$$

For the last inequality, we used that C is positive-definite. $\square$

We can now state the definition for graphical continuous Lyapunov models (GCLM).

Definition 2.13 (*GCLM*)**.** Let G be a mixed graph, then the *graphical continuous Lyapunov model* of G is the set of covariance matrices

$$\mathcal{M}_G := \Sigma(\Theta_G) = \{\Sigma \in \mathrm{PD}(p) : M\Sigma + \Sigma M^{\mathsf{T}} + C = 0 \text{ with } (M, C) \in \Theta_G\} \subseteq \mathrm{PD}(p).$$

The *extended* GCLM is defined as $\mathcal{M}^e_G := \Sigma(\Xi_G)$

When we refer to the undirected structure of a GCLM, we mean the skeleton of a compatible graph G for a given GCLM $\mathcal{M}_G$.

Definition 2.14 (*skeleton*)**.** The *skeleton* of a mixed graph $\mathcal{G} = (V, E)$ is the undirected graph $\mathcal{G}^{skel} = (V, E^{skel})$ with $v - w \in E^{skel}$ if and only if v and w are adjacent in $\mathcal{G}$.

Example 2.15. For the graph given in Example 2.11 we obtain the skeleton given in Figure 2.2. Note that we omitted edge weights and self-loops as they are not of primary

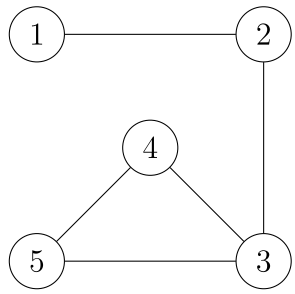

Figure 2.2: Skeleton of the compatible graph G given in Example 2.11

interest when discussing the undirected structure of a GCLM.

We call a pair of matrices $(M, C) \in \mathrm{Mat}_0(p) \times \mathrm{Sym}(p)$ *compatible* with an undirected graph G^{skel} if $M_{ij} \neq 0$ or $M_{ji} \neq 0$ implies $i - j \in E^{skel}$, and $C_{ij} \neq 0$ or $C_{ji} \neq 0$ implies $i - j \in E^{skel}$. We denote the set of G^{skel}-compatible matrix pairs by $\Xi_{G^{skel}} \subseteq \mathrm{Mat}_0(p) \times \mathrm{Sym}(p)$ and $\Theta_{G^{skel}} = \Xi_{G^{skel}} \cap (\mathrm{Stab}(p) \times \mathrm{PD}(p))$. Analogously to directed graphs, we can define undirected graphical continuous Lyapunov models as follows.

Definition 2.16 (*undirected GCLM*)**.** Let G^{skel} be an undirected graph, then the *undirected graphical continuous Lyapunov model* of G^{skel} is the set of covariance matrices

$$\mathcal{M}_{G^{skel}} := \Sigma(\Theta_{G^{skel}}) = \{\Sigma \in \mathrm{PD}(p) : M\Sigma + \Sigma M^{\mathsf{T}} + C = 0 \text{ with } (M, C) \in \Theta_{G^{skel}}\} \subseteq \mathrm{PD}(p).$$

The *extended* undirected GCLM is defined as $\mathcal{M}^e_{G^{skel}} := \Sigma(\Xi_{G^{skel}})$.

3 Undirected Structure Estimation

3.1 Group Lasso

Our method of choice to estimate the undirected structure of GCLMs will be the group lasso. It was first introduced by Yuan and Lin (2006) with the intent to select groups of variables in regression problems in which covariates admit a natural group structure. There exists a vast amount of literature on the group lasso. Our introduction to it will mostly follow Hastie et al. (2015).

Consider the linear regression problem with response $y \in \mathbb{R}^p$, design matrix $X \in \mathbb{R}^{p \times p}$ and regression coefficients $\beta \in \mathbb{R}^p$. Moreover, let the parameter β have a natural grouping with m total groups, i.e.,, $\beta = (\beta_{11}, \ldots, \beta_{1p_1}, \beta_{21}, \ldots, \beta_{2p_2}, \ldots, \beta_{m1}, \ldots, \beta_{mp_m})$ with $\sum_{j=1}^m p_j = m$. Alternatively, we write $\beta = (\beta_{G_1}, \ldots, \beta_{G_m})$ where each $\beta_{G_j} := (\beta_{j1}, \ldots, \beta_{jp_j}) \in \mathbb{R}^{p_j}$ for every $j \in [m]$.

Definition 3.1 (*group lasso*)**.** The *group lasso* estimator solves the following convex optimization problem

$$\min_{\beta \in \mathbb{R}^p} \frac{1}{2} \|y - X\beta\|_2^2 + \lambda \sum_{j=1}^m \sqrt{p_j} \left\| \beta_{G_j} \right\|_2 , \tag{3.1}$$

where $\lambda \geqslant 0$ is a tuning parameter that can be chosen freely.

Note that if we set each group size $p_j = 1$, the groups are all singletons, and the above optimization problem (3.1) reduces to the ordinary lasso estimator. Depending on the choice of λ, the parameter estimate $\hat{\beta}_{G_j}$ for β_{G_j} will either be equal to zero, or all its elements will be nonzero. We refer to the second summand in (3.1) as the penalty term.

Remark 3.2. The choice of the factor in front of $\left\| \beta_{G_j} \right\|_2$ is somewhat subjective. Here we set it to $\sqrt{p_j}$, thus weighting the penalty for each group j according to their size. Other possible choices include setting it to 1 or $\left\| X_{G_j} \right\|_F$.

We connect the problem of estimating undirected structures in GCLM to the group lasso by introducing the *group Lyapunov lasso.* Given an i.i.d. sample of centered observations $X_1, \ldots, X_n \in \mathbb{R}^p$ from the Ornstein-Uhlenbeck process in equilibrium, the group Lyapunov lasso estimator finds the optimal solution to the following problem

$$\min_{M \in \mathbb{R}^{p \times p}} \frac{1}{2} \left\| M\hat{\Sigma} + \hat{\Sigma}M^\mathsf{T} + C \right\|_F^2 + \lambda \sum_{j=1}^m \sqrt{p_j} \left\| \mathrm{vec}(M)_{G_j} \right\|_2 , \tag{3.2}$$

where $\lambda \geqslant 0$ is a tuning parameter and

$$\hat{\Sigma} = \frac{1}{n} \sum_{i=1}^{n} X_i X_i^{\mathsf{T}} \tag{3.3}$$

is the sample covariance matrix. $\mathrm{vec}(M)_{G_j}$ denotes the vector containing all entries of M that belong to group j. The exact structure of G_j will be specified later.

To see how (3.2) relates to the group lasso, we transform the objective function. As we have seen in Example 2.5 equation (2.5), the term in the Frobenius norm can be rewritten as the linear equation

$$((\hat{\Sigma} \otimes I_p) + (I_p \otimes \hat{\Sigma})K^{(p,p)})\mathrm{vec}(M) + \mathrm{vec}(C).$$

Setting

$$A(\Sigma) := ((\Sigma \otimes I_p) + (I_p \otimes \Sigma)K^{(p,p)}) \in \mathbb{R}^{p^2 \times p^2} \tag{3.4}$$

and using the fact that for a general matrix X

$$\|X\|_F = \|\mathrm{vec}(X)\|_2$$

we obtain an alternative version of (3.2)

$$\min_{M \in \mathbb{R}^{p \times p}} \frac{1}{2} \left\| A(\hat{\Sigma})\mathrm{vec}(M) + \mathrm{vec}(C) \right\|_2^2 + \lambda \sum_{j=1}^{m} \sqrt{p_j} \left\| \mathrm{vec}(M)_{G_j} \right\|_2 . \tag{3.5}$$

This now more closely resembles the standard definition of the group lasso, where $A(\hat{\Sigma}) \in \mathbb{R}^{p^2 \times p^2}$ takes on the role of the design matrix in the linear regression setting, $\mathrm{vec}(M) \in \mathbb{R}^{p^2}$ represent the regression coefficients and $-\mathrm{vec}(C) \in \mathbb{R}^{p^2}$ is the response vector.

Remark 3.3. There is a key difference between the group Lyapunov lasso and the group lasso. While the latter has a random component in the form of an additive error ϵ, i.e.,

$$Y = X\beta + \epsilon,$$

the group Lyapunov lasso introduces randomness through the design matrix $A(\hat{\Sigma})$, which is based on the sample covariance matrix $\hat{\Sigma}$.

We index the rows and columns of $A(\Sigma)$ by the pairs $(i, j) \in \{1, \ldots, p\}^2$. The design matrix $A(\Sigma)$ has an interesting structure, as the next example shows.

Example 3.4. For $p = 3$ $A(\Sigma)$ is a 9×9 matrix and has the following form

$$\begin{array}{c} \\ (1,1) \\ (1,2) \\ (1,3) \\ \mathit{(2,1)} \\ (2,2) \\ (2,3) \\ \mathit{(3,1)} \\ \mathit{(3,2)} \\ (3,3) \end{array} \begin{array}{c} \begin{array}{ccccccccc} (1,1) & (1,2) & (1,3) & (2,1) & (2,2) & (2,3) & (3,1) & (3,2) & (3,3) \end{array} \\ \begin{pmatrix} 2\Sigma_{11} & 0 & 0 & 2\Sigma_{12} & 0 & 0 & 2\Sigma_{13} & 0 & 0 \\ \Sigma_{21} & \Sigma_{11} & 0 & \Sigma_{22} & \Sigma_{12} & 0 & \Sigma_{23} & \Sigma_{13} & 0 \\ \Sigma_{31} & 0 & \Sigma_{11} & \Sigma_{23} & 0 & \Sigma_{12} & \Sigma_{33} & 0 & \Sigma_{13} \\ \Sigma_{21} & \Sigma_{11} & 0 & \Sigma_{22} & \Sigma_{12} & 0 & \Sigma_{23} & \Sigma_{13} & 0 \\ 0 & 2\Sigma_{21} & 0 & 0 & 2\Sigma_{22} & 0 & 0 & 2\Sigma_{23} & 0 \\ 0 & \Sigma_{31} & \Sigma_{21} & 0 & \Sigma_{23} & \Sigma_{22} & 0 & \Sigma_{33} & \Sigma_{23} \\ \Sigma_{31} & 0 & \Sigma_{11} & \Sigma_{23} & 0 & \Sigma_{12} & \Sigma_{33} & 0 & \Sigma_{13} \\ 0 & \Sigma_{31} & \Sigma_{21} & 0 & \Sigma_{23} & \Sigma_{22} & 0 & \Sigma_{33} & \Sigma_{23} \\ 0 & 0 & 2\Sigma_{31} & 0 & 0 & 2\Sigma_{23} & 0 & 0 & 2\Sigma_{33} \end{pmatrix} \end{array}.$$

Rows with an italicized index correspond to the strictly upper triangular entries in the Lyapunov equation defined in (1.3).

Note that it holds in general that $A(\Sigma)$ has two copies of each row corresponding to an off-diagonal entry in the Lyapunov equation. This is due to the symmetry of the Lyapunov equation.

We define the Gram matrix

$$\Gamma(\Sigma) := A(\Sigma)^{\mathsf{T}} A(\Sigma) \in \mathbb{R}^{p^2 \times p^2} \tag{3.6}$$

and

$$g(\Sigma) := -A(\Sigma)\mathrm{vec}(C) \in \mathbb{R}^{p^2}.$$

(3.5) can then be reformulated as

$$\min_{M \in \mathbb{R}^{p \times p}} \frac{1}{2} \mathrm{vec}(M)^{\mathsf{T}} \Gamma(\hat{\Sigma}) \mathrm{vec}(M) - g(\hat{\Sigma})^{\mathsf{T}} \mathrm{vec}(M) + \lambda \sum_{j=1}^{m} \sqrt{p_j} \left\| \mathrm{vec}(M)_{G_j} \right\|_2.$$

Observe that we omitted a term that is constant in M as it has no influence on the minimization problem.

Recall from Definition 2.14 that the undirected structure of the GCLM is based on the adjacency of two nodes. Consequently, given two nodes i and j, there exists an undirected edge connecting them if $m_{ij} > 0$ or $m_{ji} > 0$. Thus it is only natural to consider the $\binom{p}{2}$ groups each of length 2 defined as

$$G_{\mathrm{offdiag}} := \{\{(i,j),(j,i)\} : i \neq j,\ i,j \in [p]\}. \tag{3.7}$$

This leaves us with the entries on the diagonal of M, which we will consider as p groups each of length 1, i.e.,

$$G_{\mathrm{diag}} := \{\{(i,i)\} : i \in [p]\}. \tag{3.8}$$

We propose two versions of the group Lyapunov lasso. The first one penalizes all entries of M, whereas the second only penalizes the off-diagonal entries.

Definition 3.5 (*group Lyapunov lasso including the diagonal*)**.** The *group Lyapunov lasso* estimator, including the diagonal, finds the optimal solution to the following problem

$$\min_{M \in \mathbb{R}^{p \times p}} \frac{1}{2} \left\| A(\hat{\Sigma})\mathrm{vec}(M) + \mathrm{vec}(C) \right\|_2^2 + \lambda \sum_{j=1}^{\binom{p}{2}+p} \sqrt{p_j} \left\| \mathrm{vec}(M)_{G_j} \right\|_2 , \tag{3.9}$$

where $G_j \in G_{\mathrm{offdiag}} \cup G_{\mathrm{diag}}$ for $j \in [\binom{p}{2} + p]$.

We can elect to not penalize the entries on the diagonal of M as they describe the edge $j \to j$ for all $j \in [p]$, which is not directly a part of the undirected structure. Hence, we obtain the second version of the *group Lyapunov lasso*.

Definition 3.6 (*group Lyapunov lasso excluding the diagonal*)**.** The *group Lyapunov lasso* excluding the diagonal estimator finds the optimal solution to the following problem

$$\min_{M \in \mathbb{R}^{p \times p}} \frac{1}{2} \left\| A(\hat{\Sigma})\mathrm{vec}(M) + \mathrm{vec}(C) \right\|_2^2 + \lambda \sum_{j=1}^{\binom{p}{2}} \sqrt{2} \left\| \mathrm{vec}(M)_{G_j} \right\|_2 , \tag{3.10}$$

where $G_j \in G_{\mathrm{offdiag}}$ for $j \in [\binom{p}{2}]$.

3.2 Review of support recovery conditions for the group lasso

When it comes to available literature for conditions that guarantee consistent support recovery, there exist a few papers whose main ideas we aim to present in the following. The results are slightly modified to match the group lasso setting defined in (3.1) and thus allow for better comparison. Specifically, we omit the specification of a bias term and fix the weights of our penalties to be the group size squared $\sqrt{p_j}$ for all $j \in [m]$.

Definition 3.7 (*support*)**.** We define the *support* S of a vector $\beta \in \mathbb{R}^p$ as

$$S = S(\beta) := \{j \in [m] : \beta_{G_j} \neq 0\}.$$

We use the convention for vectors that $\beta_S := \{\beta_{G_j}\}_{j \in S}$ and for matrices X_S is formed by the columns of X with index in $\bigcup_{j \in S} G_j$. Moreover, let $\hat{\beta}$ be the solution to (3.1) with associated support $\hat{S} = S(\hat{\beta})$ and β^* be the true signal with support $S = S(\beta^*)$.

Remark 3.8. We will use the terms correct support recovery, correct model selection and correct sparsity pattern synonymously.

One of the first papers to study consistent model selection for the group lasso was Bach (2008). It proposes sufficient and necessary conditions for correct model selection under the following assumptions. Let $Y \in \mathbb{R}$ be a response from covariates $X \in \mathbb{R}^p$, both satisfying

1. $\mathbb{E}[\|X\|_2^4] < \infty$ and $\mathbb{E}[\|Y\|_2^4] < \infty$,
2. $\Sigma_{XX} := \mathbb{E}[XX^\top] - \mathbb{E}[X]\mathbb{E}[X]^\top$ is invertible,
3. $\mathbb{E}[(Y - (\beta^*)^\top X)^2|X]$ is almost surely greater than $\sigma^2_{min} > 0$ with $\beta^* \in \mathbb{R}^p$ denoting the minimizer of $\mathbb{E}[(Y - (\beta^*)^\top X)^2]$.

Observe that the last assumption does not require $\mathbb{E}[Y|X]$ to be an affine function of X and the conditional variance to be constant, as is commonplace for most results derived for consistency in linear supervised learning.

Leading us to the consistency conditions

$$\max_{i \in S^c} \frac{1}{\sqrt{p_i}} \left\| (\Sigma_{XX})_{G_i S}((\Sigma_{XX})_{SS})^{-1} \mathrm{diag}\left(\sqrt{p_j} / \left\|\beta^*_{G_j}\right\|_2\right) \beta^*_S \right\|_2 < 1, \tag{3.11}$$

$$\max_{i \in S^c} \frac{1}{\sqrt{p_i}} \left\| (\Sigma_{XX})_{G_i S}((\Sigma_{XX})_{SS})^{-1} \mathrm{diag}\left(\sqrt{p_j} / \left\|\beta^*_{G_j}\right\|_2\right) \beta^*_S \right\|_2 \leqslant 1, \tag{3.12}$$

where $\mathrm{diag}\left(\sqrt{p_j} / \left\|\beta^*_{G_j}\right\|_2\right)$ is a block-diagonal matrix (with block sizes p_j) in which each diagonal block is equal to $\frac{\sqrt{p_j}}{\left\|\beta^*_{G_j}\right\|_2} I_{p_j}$. We refer to (3.11) as the *strong condition* and (3.12) is called the *weak condition*. The strong condition is sufficient for consistent support recovery, as the next theorem shows.

Theorem 3.9. *Under the assumptions 1. - 3. stated earlier, if condition* (3.11) *is fulfilled, then for any sequence of regularization parameters λ_n with $\lambda_n \to 0$ and $\lambda_n\sqrt{n} \to +\infty$, the group lasso estimate $\hat{\beta}$ (solution to* (3.1)*) will converge in probability to β^* and $\mathbb{P}(\hat{S} = S) \to 1$.*

On the other hand, if there exists a consistent solution along a regularization sequence, the weak condition (3.12) must be met.

Theorem 3.10. *Under the previously stated assumptions, if there exists a regularization path λ_n such that both $\hat{\beta}$ converges to β^* and $\hat{S}$ converges to S in probability, then* (3.12) *holds true.*

Theorem 3.9 implies that if there is low correlation between the predictors in S and the predictors in S^c, the group lasso will be consistent. By contrast, Theorem 3.10 states that we can not hope for a consistent solution if the weak condition is not met. In analogy to the theory for model consistency developed for the lasso by Zhao and Yu (2006), (3.11) and (3.12) are called *irrepresentability conditions.* Moreover, note that for the lasso Yuan and Lin (2007) proved that the strong condition is both necessary and sufficient.

The two previous theorems provide results on a particular type of consistency, namely consistency in both the norm $\left\|\hat{\beta} - \beta^*\right\|$ and the sparsity pattern simultaneously. If we are only interested in sparsistency, i.e., consistent estimation of the support, we have the following result.

Theorem 3.11. *Under the assumptions stated earlier, if $\lambda_n \to \lambda_0 > 0$, the group lasso estimate $\hat{\beta}$ is sparsity-consistent if and only if the solution of*

$$\min_{\beta \in \mathbb{R}^p} \frac{1}{2}(\beta - \beta^*)^\top \Sigma_{XX}(\beta - \beta^*) + \lambda_0 \sum_{j=1}^{m} \sqrt{p_j} \left\| \beta_{G_j} \right\|_2$$

has the correct sparsity pattern.

Consequently, we may have consistent estimation of the support but inconsistent estimation of the parameters, even though the weak condition is not satisfied. Bach (2008) also showed this fact in a simulation study, where 10000 covariance matrices were sampled based on a certain procedure. The study revealed that even in cases where the weak condition was violated the group lasso was able to consistently estimate the correct sparsity pattern for roughly 40% of the samples.

Although the assumptions in Bach (2008) accommodate for quite general settings, they are not directly applicable to the group Lyapunov lasso. In particular, Σ_{XX}, which would be equal to $\Gamma^* = A(\Sigma^*)^\top A(\Sigma^*)$ in our setting, is not invertible. This is due to the identical rows that appear in the design matrix $A(\Sigma^*)$ (cf. also Example 3.4).

For a more special setting, namely the linear model

$$Y = X\beta^* + \epsilon \quad \text{ with } X \in \mathbb{R}^{n \times p}, \beta^* \in \mathbb{R}^p \text{ and } \epsilon \sim \mathcal{N}(0, \sigma^2 I_n),$$

Nardi and Rinaldo (2008) proposes the following conditions for sparsistency. First, they make the simplifying assumption that

$$\frac{1}{n} X_{G_j}^\top X_{G_j} = I_{p_j} \quad \text{ for all } j \in [m], \tag{3.13}$$

which can be achieved by utilizing the Gram-Schmidt orthogonalization procedure. Moreover, they assume that

1. The smallest eigenvalue of $\frac{1}{n}(X_S^\top X_S)$ is bounded from below by a constant $C_{min} > 0$.
2. For $\alpha := \min_{j \in S} \left\| \beta^*_{G_j} \right\|_\infty$ and $s := \sum_{j \in S} p_j$

$$\frac{1}{\alpha}\left(\sqrt{\frac{\log s}{n}} + \sqrt{s}\lambda \max_{j \in S} \sqrt{p_j} \right) \to 0.$$

3. For some $0 < \epsilon < 1$ and every $j \in S^c$

$$\left\| X_{G_j}^\top X_S (X_S^\top X_S)^{-1} \right\|_2 \leqslant \frac{1 - \epsilon}{\sqrt{\sum_{j \in S} p_j}}. \tag{3.14}$$

4.
$$\frac{\log(p-s)}{n\lambda^2} \to 0.$$

They derive the following theorem for sparsistency.

Theorem 3.12. *Let assumptions* 1. − 4. *hold. Then for $\mathcal{O}$, which is the event that there exists a group lasso solution $\hat{\beta}$ with $\left\|\hat{\beta}_{G_j}\right\|_2 > 0$ for all $j \in S$ and $\hat{\beta}_{G_j} = 0$ for all $j \in S^c$, it holds that*
$$\mathbb{P}(\mathcal{O}) \to 1 \textit{ for } n \to \infty.$$

The irrepresentability condition stated in (3.14) is more restrictive than the one proposed by Bach (2008) in (3.11). However, this is because Nardi and Rinaldo (2008) considers a more general asymptotic scenario. They study the so-called *double-asymptotic scenario*, where the parameter space spanned by the columns of the $n \times p_j$ submatrices X_{G_j} with $j \in [m]$ is allowed to change with n. In particular, they allow for $m \to \infty$ and dimensions of the groups p_j to change with n. Specifically, they include situations where $s >> n$, meaning s grows faster than n. By contrast, Bach (2008) assumes the parameter space is fixed.

Lounici et al. (2011) established slightly different sufficient conditions for correct support recovery for the linear model. Let X_{G_j} be the $n \times p_j$ sub-matrix of X obtained by taking the columns of X indexed in G_j. Assume there exists some integer $s \geqslant 1$ and a constant $\alpha > 0$ such that

1. For any $j \in [m]$ and $t \in [p_j]$ it holds that $(X_{G_j}^\top X_{G_j}/n)_{t,t} = \phi$ and
$$\max_{1 \leqslant t,t' \leqslant p_j, t \neq t'} |(X_{G_j}^\top X_{G_j}/n)_{t,t'}| \leqslant \frac{\min_{i \in [m]} \sqrt{p_i}\phi}{14\alpha \max_{i \in [m]} \sqrt{p_i} s} \frac{1}{p_j}. \tag{3.15}$$

2. For any $j \neq j' \in [m]$ it holds that
$$\max_{1 \leqslant t \leqslant \min\{p_j, p_{j'}\}} |(X_{G_j}^\top X_{G_{j'}}/n)_{t,t}| \leqslant \frac{\min_{i \in [m]} \sqrt{p_i}\phi}{14\alpha \max_{i \in [m]} \sqrt{p_i} s} \tag{3.16}$$
and
$$\max_{1 \leqslant t \leqslant p_j, 1 \leqslant t' \leqslant p_{j'}, t \neq t'} |(X_{G_j}^\top X_{G_{j'}}/n)_{t,t'}| \leqslant \frac{\min_{i \in [m]} \sqrt{p_i}\phi}{14\alpha \max_{i \in [m]} \sqrt{p_i} s} \frac{1}{\sqrt{p_j p_{j'}}}. \tag{3.17}$$

Note that $(X_{G_j}^\top X_{G_j}/n)_{t,t} = \phi$ can be easily achieved by normalizing the design matrix X. Compared to the orthogonality assumption (3.13) made in Nardi and Rinaldo (2008), this allows for more general design matrices. Assumptions 1. and 2. take on the form of a *mutual coherence condition* as they restrict the maximum absolute value of the cross-correlations between the columns of X. Given the above two assumptions, we can state the following theorem.

Theorem 3.13. *Let $m \geqslant 2$,$n \geqslant 1$, $|S| \leqslant s$ and choose*

$$\lambda \geqslant \max_{j\in[m]} \frac{2\sigma}{\sqrt{np_j}} \sqrt{tr(X_{G_j}^{\mathsf{T}} X_{G_j}/n) + 2 \left\| X_{G_j}^{\mathsf{T}} X_{G_j}/n \right\|_2 (2q \log m + \sqrt{p_j q \log m}}, \tag{3.18}$$

where q is a positive parameter. Moreover, assume the above-stated assumptions are satisfied. In particular, let (3.15), (3.16) *and* (3.17) *hold with the same s. Define*

$$c := \left(\frac{3}{2} + \frac{16}{7(\alpha - 1)} \right).$$

Then with probability at least $1 - 2m^{1-q}$, it holds that for any solution $\hat{\beta}$ of the group lasso

$$\max_{j\in[m]} \left\| (\hat{\beta} - \beta^*)_{G_j} \right\|_2 \leqslant \frac{c}{\phi} \max_{i\in[m]} \sqrt{p_i}.$$

If additionally,

$$\min_{j\in S} \left\| (\beta^*)_{G_j} \right\|_2 > \frac{2c}{\phi} \max_{i\in[m]} \sqrt{p_i}, \tag{3.19}$$

then with the same probability, the set of indices

$$\hat{S} = \left\{ j : \left\| \hat{\beta}_{G_j} \right\|_2 > \frac{c}{\phi} \max_{i\in[m]} \sqrt{p_i} \right\}$$

will correctly estimate the true sparsity pattern S, i.e., $\hat{S} = S$.

(3.19) is referred to as the *beta min condition*. It states that $(\beta^*)_{G_j}$ cannot be arbitrarily close to zero if $j \in S$. In other words, the norms should be at least somewhat larger than the noise level. For the group lasso this implies that at least one component in $(\beta^*)_{G_j}$ has to be sufficiently large since we aim to select entire groups and not individual components.

Note the theory developed in Nardi and Rinaldo (2008) and Lounici et al. (2011) again cannot be applied to the group Lyapunov lasso as the underlying regression problem does not contain an additive error. For more details, we refer to Remark 3.3.

For some more specialized variants of the group lasso, there also exists some sufficient conditions for correct support recovery. Wei and Huang (2010) studied the adaptive group lasso, which iteratively uses the group lasso twice. Specifically, it uses the group lasso to compute an initial estimate $\hat{\beta}$ and reduce the dimension of the problem. Afterward, the weights of the penalty are set to $1/\left\| \hat{\beta}_{G_k} \right\|_2$, and the group lasso algorithm is run for a second time. The crucial assumption for correct model selection is that the initial estimator is consistent at zero. The latter is true for the group lasso given that the *sparse Riesz condition* (SRC)

$$c_* \leqslant \frac{\|X_A v\|_2^2}{n \|v\|_2^2} \leqslant c^* \quad \text{for all } A \text{ with } q^* = |A| \text{ and } v \in \mathbb{R}^{\sum_{k\in A} p_k}$$

is satisfied by the submatrix X_A with rank q^* and $0 < c_* < c^* < \infty$. In words, the SRC restricts the range of eigenvalues of the submatrix.

The group square root lasso is obtained by taking the square of the prediction error in the group lasso, while the penalty remains the same. Bunea et al. (2014) showed that assuming an irrepresentability condition

$$\max_{v:\left\|v_{G_k}\right\|_2 \leqslant \sqrt{p_k}} \max_{j \in S^c} \frac{\left\|X_{G_j}^\intercal X_S (X_S^\intercal X_S)^{-1} v\right\|_2}{\sqrt{p_j}} < \eta \quad \text{with } 0 < \eta < 1$$

and a beta min condition, the group square root lasso can consistently select the correct sparsity pattern.
The advantage of the group square root lasso is that the optimal tuning parameter λ can be selected independently of the variance of the errors σ. For the group lasso, this is not the case, as can be seen from the choice of λ in (3.18).

In summary, to establish sufficient conditions for correct support recovery we need a condition on the design matrix X. This is usually in the form of an irrepresentability condition or mutual coherence condition. Secondly, we require some restrictions on the tuning parameter λ. Particularly, we either specify the speed of convergence or assume a lower bound for λ. Lastly, we may need a beta min condition to distinguish the true signal from the noise.
Surveying the available literature, we were not able to find appropriate theory for the group Lyapunov lasso. To prove sufficient conditions for the latter we will proceed as follows. We will adapt the primal-dual witness construction developed by Wainwright (2009) to the group lasso setting. To achieve this it is useful to introduce the concept of dual norms.

3.3 Dual norm

Recall that the result of consistent support recovery for the lasso

$$\min_{\beta \in \mathbb{R}^p} \frac{1}{2} \|y - X\beta\|_2^2 + \lambda \|\beta\|_1$$

is dependent on the irrepresentability condition formulated in terms of the ℓ_∞-norm

$$\left\|X_{S^c}^\intercal X_S (X_S^\intercal X_S)^{-1} \text{sign}(\beta^*)\right\|_\infty < 1.$$

This is by no means coincidental. In fact, the ℓ_∞-norm is the dual norm to the ℓ_1-norm in the lasso penalty. We will show that the group lasso penalty is a norm and derive its dual norm, which will be the basis for a group irrepresentability condition. The theory presented here is based on van de Geer (2016) and Miccheli et al. (2010).

Let Ω be a norm on $\mathbb{R}^p$, we can define the *dual norm* Ω_* of Ω as follows.

Definition 3.14 (*dual norm*).

$$\Omega_*(\omega) := \max_{\Omega(\beta)\leqslant 1} |w^\mathsf{T}\beta|,\ \omega \in \mathbb{R}^p.$$

We can immediately verify that the *dual norm inequality* holds by Definition 3.14

$$|w^\mathsf{T}\beta| \leqslant \Omega_*(\omega)\Omega(\beta). \tag{3.20}$$

Example 3.15. Consider the ℓ_1-norm $\|\cdot\|_1$ whose dual norm is the ℓ_∞-norm $\|\cdot\|_\infty$. To see that we write

$$\max_{\|\beta\|_1\leqslant 1} |w^\mathsf{T}\beta| \leqslant \max_{\|\beta\|_1\leqslant 1} \sum_{i=1}^p |\omega_i\beta_i| \leqslant \max_{i\in[p]} |\omega_i| \max_{\|\beta\|_1\leqslant 1} \|\beta\|_1 \leqslant \|\omega\|_\infty ,$$

where we used the triangle inequality for the first inequality. Note the maximum on the RHS is attained by choosing β as the appropriate canonical basis vector.

Definition 3.16 (($\Omega-$)*allowed*). Let $S \subseteq [p]$ and Ω^{-S} be a norm defined on $\mathbb{R}^{p-|S|}$. We say S is ($\Omega-$)*allowed* if

$$\Omega(\beta) \geqslant \Omega(\beta_S) + \Omega^{-S}(\beta_{-S}) \quad \text{for all} \beta \in \mathbb{R}^p. \tag{3.21}$$

We call Ω weakly decomposable for the set S.

Observe also that by the triangle inequality, it holds in general that

$$\Omega(\beta) \leqslant \Omega(\beta_S) + \Omega(\beta_{-S}).$$

Thus for allowed sets, we have the reverse implication of the above inequality, albeit that $\Omega(\beta_{-S})$ is replaced by a different norm.

Example 3.17. For the ℓ_1-norm choosing Ω^{-S} as the ℓ_1-norm again yields that any subset $S \subseteq [p]$ is allowed. In fact, (3.21) then holds with equality, i.e.,

$$\|\beta\|_1 = \|\beta_S\|_1 + \|\beta_{-S}\|_1 .$$

The group lasso penalty is an example of a more general class of norms that satisfy the weak decomposability property (cf. Definition 3.16).

Definition 3.18. Let $\mathcal{A}$ be a convex cone in $[0,\infty)^p$. The norm Ω generated by $\mathcal{A}$ is defined as

$$\Omega(\beta) := \min_{a\in\mathcal{A}} \frac{1}{2}\sum_{j=1}^p \left[\frac{|\beta_j|^2}{a_j} + a_j\right],\ \beta \in \mathbb{R}^p. \tag{3.22}$$

We use the convention $0/0 = 0$. Observe that if $\beta_j \neq 0$, one is forced to choose an $a_j \neq 0$ in (3.22).

Lemma 3.19. The function Ω given in Definition 3.18 above is a norm.

Proof. Ω is clearly non-negative by definition and can only be zero when $\beta \equiv 0$. Moreover for $\lambda \geqslant 0$ we have that

$$\Omega(\lambda\beta) = \min_{a\in\mathcal{A}} \frac{1}{2}\sum_{j=1}^{p}\left[\frac{\lambda^2|\beta_j|^2}{a_j} + a_j\right] = \lambda \min_{a\in\mathcal{A}} \frac{1}{2}\sum_{j=1}^{p}\left[\frac{|\beta_j|^2}{\frac{a_j}{\lambda}} + \frac{a_j}{\lambda}\right] = \lambda\Omega(\beta),$$

thus proving homogeneity. For the last equality, we used that $\mathcal{A}$ is a cone, i.e., $\frac{a}{\lambda} \in \mathcal{A}$. To see that the triangle inequality holds note $\frac{1}{2}\sum_{j=1}^{p}\left[\frac{|\beta_j|^2}{a_j} + a_j\right]$ is convex as it is a sum of the convex functions $(a, b) \mapsto \frac{b^2}{a}$ and $a \mapsto a$. Hence, when $\mathcal{A}$ is convex, Ω is convex since it is the minimum of convex functions over $\mathcal{A}$. Assume $\alpha, \beta \in \mathbb{R}^p$, then applying homogeneity and convexity of Ω we obtain that

$$\Omega(\alpha + \beta) = 2\Omega\left(\frac{\alpha + \beta}{2}\right) \leqslant \Omega(\alpha) + \Omega(\beta).$$

□

Assume $\{G_j : j \in [m]\}$ forms a partition of the index set $[p]$ into m groups with $|G_j| = p_j$ and

$$\mathcal{A} = \{\alpha \in [0, \infty)^p : \alpha_i = \theta_l,\ i \in G_j,\ l \in [m],\ \theta_l > 0\},$$

the set of all non-negative vectors which are constant within groups. Thus using the arithmetic-geometric mean inequality, i.e. $\frac{a+b}{2} \geqslant \sqrt{ab}$

$$\Omega(\beta) = \min_{a\in\mathcal{A}} \frac{1}{2}\sum_{j=1}^{p}\left[\frac{|\beta_j|^2}{a_j} + a_j\right] = \min_{a\in\mathcal{A}} \frac{1}{2}\sum_{j=1}^{m}\left[\frac{\left\|\beta_{G_j}\right\|_2^2}{a_j} + p_j a_j\right] \geqslant \sum_{j=1}^{m}\sqrt{p_j}\left\|\beta_{G_j}\right\|_2$$

Hence, we have proven that the group lasso penalty is a norm of the form specified in Definition 3.18.

Definition 3.20 (*group norm*)**.** For $\beta \in \mathbb{R}^p$ we define the *group norm* as

$$\|\beta\|_g := \sum_{j=1}^{m}\sqrt{p_j}\left\|\beta_{G_j}\right\|_2.$$

Note that any union of groups G_j form an allowed set, and for any allowed set S

$$\Omega^{-S}(\beta_{-S}) = \|\beta_{-S}\|_g$$

and

$$\|\beta\|_g = \|\beta_S\|_g + \|\beta_{-S}\|_g.$$

To derive the corresponding dual norm, we simply use Definition 3.14 and the Cauchy-Schwarz inequality.

$$\begin{aligned}\Omega_*(\omega) = \max_{\|\beta\|_g \leqslant 1} |w^\top \beta| &\leqslant \max_{\|\beta\|_g \leqslant 1} \sum_{j=1}^m \|w_{G_j}\|_2 \|\beta_{G_j}\|_2 \\ &\leqslant \max_{j \in [m]} \frac{\|w_{G_j}\|_2}{\sqrt{p_j}} \max_{\|\beta\|_g \leqslant 1} \sum_{j=1}^m \sqrt{p_j} \|\beta_{G_j}\|_2 \\ &\leqslant \max_{j \in [m]} \frac{\|w_{G_j}\|_2}{\sqrt{p_j}}.\end{aligned}$$

Definition 3.21 (*dual group norm*)**.** For vectors $\omega \in \mathbb{R}^p$ the *dual group norm* is given as

$$\|\omega\|_* := \max_{j \in [m]} \frac{\|\omega_{G_j}\|_2}{\sqrt{p_j}}.$$

If $M \in \mathbb{R}^{p \times p}$, we define the dual group matrix norm as

$$|||M|||_* := \max_{\|\omega\|_* = 1} \|M\omega\|_* = \max_{\omega \neq 0} \frac{\|M\omega\|_*}{\|\omega\|_*}.$$

Remark 3.22. If we choose the partition $G_j = \{j\}$ for all $j \in [p]$, we obtain the ℓ_1-norm for the group norm and the ℓ_∞-norm for the dual group norm. On the contrary, if we set the partition to be $[p]$, we obtain a scaled version of the ℓ_2-norm for the group norm and the dual group norm, respectively.

For two norms Ω and Ω' on the Euclidean space, we say Ω is a *stronger norm* than Ω' if

$$\Omega \geqslant \Omega' \Leftrightarrow \Omega(\beta) \geqslant \Omega'(\beta) \text{ for all } \beta \in \mathbb{R}^p$$

In particular, this implies for the associated dual norms $\Omega'_* \geqslant \Omega_*$ since $\{\beta : \Omega(\beta) \leqslant 1\} \subseteq \{\beta : \Omega'(\beta) \leqslant 1\}$ and thus

$$\Omega'_*(\omega) = \max_{\Omega'(\beta) \leqslant 1} |w^\top \beta| \geqslant \max_{\Omega(\beta) \leqslant 1} |w^\top \beta| = \Omega_*(\omega).$$

Recall that for $\beta \in \mathbb{R}^p$ the $\ell_1 - \ell_2$-inequality $\|\beta\|_1 \leqslant \sqrt{p}\,\|\beta\|_2$ holds. As an immediate consequence, we have that $\|\cdot\|_g$ is stronger than $\|\cdot\|_1$. Indeed, for all norms of the form given in Definition 3.18, it holds true that they are stronger than the ℓ_1-norm (van de Geer (2016), Lemma 6.9) and hence their dual norms are weaker. When specifying oracle inequalities, weaker dual norms can result in more relaxed bounds on the tuning parameter λ (van de Geer (2016), Section 6.7).

We finish this section by providing several bounds for the dual group matrix norm in terms of the associated matrix norm $|||\cdot|||_2$ and $|||\cdot|||_\infty$.

Lemma 3.23. Let $M \in \mathbb{R}^{p\times p}$, then

$$|||M|||_* \leqslant \max_{i\in[m]} \sum_{j=1}^{m} \sqrt{\frac{p_j}{p_i}} \,|||M_{G_iG_j}|||_2 \,,$$

where $M_{G_iG_j}$ is the submatrix formed by selecting the rows of M with index in G_i and columns of M with index in G_j.

Proof.

$$\begin{aligned}|||M|||_* &= \max_{\|\omega\|_*=1} |||M\omega|||_* = \max_{\|\omega\|_*=1} \max_{i\in[m]} \frac{\|M_{G_i}\omega\|_2}{\sqrt{p_i}} = \max_{i\in[m]} \max_{\|\omega\|_*=1} \frac{\sum_{j=1}^m \left\|M_{G_iG_j}\omega_{G_j}\right\|_2}{\sqrt{p_i}} \\ &\leqslant \max_{i\in[m]} \max_{\|\omega\|_*=1} \sum_{j=1}^{m} \frac{|||M_{G_iG_j}|||_2 \left\|\omega_{G_j}\right\|_2}{\sqrt{p_i}} \leqslant \max_{i\in[m]} \sum_{j=1}^{m} \sqrt{\frac{p_j}{p_i}} \,|||M_{G_iG_j}|||_2\end{aligned}$$

For the last inequality, we used that $\frac{\left\|\omega_{G_j}\right\|_2}{\sqrt{p_j}} \leqslant \|\omega\|_* = 1$ for all $j \in [m]$ by Definition 3.21. ☐

Lemma 3.24. Let $M \in \mathbb{R}^{p\times q}$ and S an allowed set of size q, then

$$|||M|||_* \leqslant \frac{\sqrt{q}}{\min_{j\in[m]} \sqrt{p_j}} \,|||M|||_2 \,.$$

Proof. Observe that $\|\omega\|_* \leqslant \frac{1}{\min_{j\in[m]} \sqrt{p_j}} \|\omega\|_2$, thus

$$|||M|||_* \leqslant \max_{\|\omega\|_*=1} \frac{1}{\min_{j\in[m]} \sqrt{p_j}} \|M\omega\|_2$$

Let m_q be the number of groups in S, then $\|\omega\|_* = \max_{j\in[m_q]} \frac{\left\|\omega_{G_j}\right\|_2}{\sqrt{p_j}} = 1$ implies

$$\|\omega\|_2 = \sqrt{\sum_{j=1}^{m_q} p_j \left(\frac{\left\|\omega_{G_j}\right\|_2}{\sqrt{p_j}}\right)^2} \leqslant \sqrt{\sum_{j=1}^{m_q} p_j} = \sqrt{q},$$

which yields that

$$|||M|||_* \leqslant \frac{\sqrt{q}}{\min_{j\in[m]} \sqrt{p_j}} \,|||M|||_2 \,.$$

☐

Lemma 3.25. Let $M \in \mathbb{R}^{p\times p}$, then

$$|||M|||_* \leqslant \max_{j\in[m]} \sqrt{p_j} \,|||M|||_\infty \,.$$

Proof. Recall that the group norm is stronger than the ℓ_1-norm, which implies that the dual group norm $\|\cdot\|_*$ is weaker than the ℓ_∞-norm. Therefore,

$$\begin{aligned}|||M|||_* &= \max_{\|\omega\|_*=1} \|M\omega\|_* \leqslant \max_{\|\omega\|_*=1} \|M\omega\|_\infty \\ &\leqslant |||M|||_\infty \max_{\|\omega\|_*=1} \|\omega\|_\infty \leqslant |||M|||_\infty \max_{j\in[m]} \sqrt{p_j}.\end{aligned}$$

☐

3.4 Primal-dual witness construction for group lasso

A key technique to prove consistency results for the group Lyapunov lasso will be the primal-dual witness construction, which was introduced by Wainwright (2009) to prove consistent support recovery of the lasso. We start by establishing general optimality conditions for the group lasso. Then we will adapt the primal-dual witness construction and provide a lemma showing that this method is viable.

Definition 3.26 (*subgradient*, *subdifferential*)**.** Given a convex function $f : \mathbb{R}^p \to \mathbb{R}$ a vector $z \in \mathbb{R}^p$ is said to be a *subgradient* of f at β if

$$f(\beta') \geqslant f(\beta) + \langle z, \beta' - \beta \rangle \text{ for all } \beta' \in \mathbb{R}^p. \tag{3.23}$$

We denote the set of all subgradients of f at β - called *subdifferential* - by $\partial f(\beta)$.

Example 3.27. We consider the group lasso penalty $\|\beta\|_g$ and calculate the subgradient to be Dz, where $D \in \mathbb{R}^{p\times p}$ is a block-diagonal matrix with m blocks $\sqrt{p_j} I_{p_j} \in \mathbb{R}^{p_j \times p_j}$ with $j \in [m]$ and $z \in \partial \|\beta\|_g$ is defined as

$$z_{G_j} = \begin{cases} \frac{\beta_{G_j}}{\|\beta_{G_j}\|_2} & \text{if } \beta_{G_j} \neq 0, \\ \in \{x \in \mathbb{R}^{p_j} : \|x\|_2 \leqslant 1\} & \text{if } \beta_{G_j} = 0. \end{cases} \tag{3.24}$$

Considering the case $\beta_{G_j} \neq 0$, we used that the subdifferential reduces to a single vector for differentiable functions, namely the derivative $\nabla \left\|\beta_{G_j}\right\|_2$. Clearly, the euclidean norm is not differentiable at 0, hence we have to show that for $\beta_{G_j} = 0$ the subgradient condition (3.23) holds, i.e.,

$$\left\|\beta'_{G_j}\right\|_2 \geqslant \|0\|_2 + \langle z_{G_j}, \beta'_{G_j} - 0 \rangle.$$

Assuming $\left\|z_{G_j}\right\|_2 \leqslant 1$, we have with Cauchy-Schwarz that

$$|\langle z_{G_j}, \beta'_{G_j} \rangle| \leqslant \left\|z_{G_j}\right\|_2 \left\|\beta'_{G_j}\right\|_2 \leqslant \left\|\beta'_{G_j}\right\|_2 .$$

We can now state a necessary and sufficient optimality condition for the group lasso, derived from the Karush-Kuhn-Tucker (KKT) conditions for constrained convex optimization problems. In particular, the solution $\hat{\beta}$ of the group lasso satisfies the following subgradient equations

$$X^\top X \hat{\beta} - X^\top y + \lambda D \hat{z} = 0, \tag{3.25}$$

where $\hat{z}$ is defined as in (3.24).

Let S be the support defined in Definition 3.7. We use the convention for vectors that $\beta_S := \{\beta_{G_j}\}_{j \in S}$ and for matrices X_S is formed by the columns of X with index in $\bigcup_{j\in S} G_j$.

The *primal-dual witness* (PDW) method is based on an explicit construction of a pair of vectors $(\hat{\beta}, \hat{z})$ that (when the procedure succeeds) are primal and dual optimal solutions for the group lasso, and act as witnesses for the correct recovery of the support S. We construct $(\hat{\beta}, \hat{z})$ according to the following steps:

1. Obtain $\hat{\beta}_S$ by solving the restricted group lasso problem,

$$\hat{\beta}_S = \arg\min_{\beta_S} \left\{ \frac{1}{2} \left\| y - X_S \beta_S \right\|_2^2 + \lambda \left\| \beta_S \right\|_g \right\}.$$

   We set $\hat{\beta}_{S^c} = 0$.

2. Choose $\hat{z}_S$ as an element of the subdifferential of the group norm evaluated at $\hat{\beta}_S$, i.e., $\hat{z}_S \in \partial \left\| \hat{\beta}_S \right\|_g$.

3. Assuming $y = X\beta^* + \epsilon$, solve the zero subgradient equation (3.25) for $\hat{z}_{S^c}$ and check whether or not the dual feasibility condition $\left\| \hat{z}_{S^c} \right\|_* < 1$ is satisfied.

Lemma 3.28. If the primal-dual witness construction succeeds with $\left\| \hat{z}_{S^c} \right\|_* < 1$, $(\hat{\beta}_S, 0)$ is an optimal solution for the group lasso. In particular, this implies $\hat{S} \subseteq S$, where $\hat{S} = S(\hat{\beta})$.

Proof. By construction of the PDW-method $\hat{\beta} = (\hat{\beta}_S, 0)$ is an optimal solution to (3.1) with associated subgradient vector $\hat{z} \in \mathbb{R}^p$ satisfying $\left\| \hat{z}_{S^C} \right\|_* < 1$.
Assume $\tilde{\beta}$ to be another optimal solution of the group lasso. Thus denoting $F(\beta) := \frac{1}{2} \left\| y - X\beta \right\|_2^2$, we have

$$F(\hat{\beta}) + \lambda \sum_{j=1}^{m} \sqrt{p_j} \left\| \hat{\beta}_{G_j} \right\|_2 = F(\tilde{\beta}) + \lambda \sum_{j=1}^{m} \sqrt{p_j} \left\| \tilde{\beta}_{G_j} \right\|_2. \tag{3.26}$$

Let $D \in \mathbb{R}^{p \times p}$ be a block-diagonal matrix with m blocks $\sqrt{p_j} I_{p_j} \in \mathbb{R}^{p_j \times p_j}$ with $j \in [m]$. Then by definition of the subgradient vector $\hat{z}$ and $\hat{\beta}$

$$\langle D\hat{z}, \hat{\beta} \rangle = \sum_{j \in S} \sqrt{p_j} \frac{\hat{\beta}_{G_j}^\intercal \hat{\beta}_{G_j}}{\left\| \hat{\beta}_{G_j} \right\|_2} = \sum_{j \in S} \sqrt{p_j} \left\| \hat{\beta}_{G_j} \right\|_2 = \sum_{j=1}^{m} \sqrt{p_j} \left\| \hat{\beta}_{G_j} \right\|_2.$$

We can now express (3.26) as

$$\begin{aligned} F(\hat{\beta}) + \lambda \langle D\hat{z}, \hat{\beta} \rangle &= F(\tilde{\beta}) + \lambda \sum_{j=1}^{m} \sqrt{p_j} \left\| \tilde{\beta}_{G_j} \right\|_2 \\ F(\hat{\beta}) - \lambda \langle D\hat{z}, \tilde{\beta} - \hat{\beta} \rangle &= F(\tilde{\beta}) + \lambda \sum_{j=1}^{m} \sqrt{p_j} \left\| \tilde{\beta}_{G_j} \right\|_2 - \lambda \langle D\hat{z}, \tilde{\beta} \rangle. \end{aligned}$$

Using the subgradient equations required for optimality, namely $\lambda D\hat{z} = -\nabla F(\hat{\beta})$ yields

$$F(\hat{\beta}) + \langle \nabla F(\hat{\beta}), \tilde{\beta} - \hat{\beta} \rangle - F(\tilde{\beta}) = \lambda \sum_{j=1}^{m} \sqrt{p_j} \left\| \tilde{\beta}_{G_j} \right\|_2 - \lambda \langle D\hat{z}, \tilde{\beta} \rangle. \tag{3.27}$$

Note that by convexity of F, the LHS of equation (3.27) is negative and hence

$$\left\|\tilde{\beta}\right\|_g = \sum_{j=1}^m \sqrt{p_j} \left\|\tilde{\beta}_{G_j}\right\|_2 \leqslant \langle D\hat{z}, \tilde{\beta}\rangle \leqslant \|D\hat{z}\|_* \left\|\tilde{\beta}\right\|_g = \left\|\tilde{\beta}\right\|_g .$$

For the last inequality, we used the dual norm inequality (3.20), while for the last equality, note that $\|D\hat{z}\|_* = 1$ by definition of $\hat{z}$.
Finally, we obtain

$$\begin{aligned}\left\|\tilde{\beta}\right\|_g = \langle D\hat{z}, \tilde{\beta}\rangle = \sum_{j=1}^m \sqrt{p_j}\hat{z}_{G_j}^\mathsf{T}\tilde{\beta}_{G_j} \\ \leqslant \sum_{j\in S} \sqrt{p_j} \left\|\hat{\beta}_{G_j}\right\|_2 + \sum_{j\in S^c} \sqrt{p_j} \left\|\hat{z}_{G_j}\right\|_2 \left\|\hat{\beta}_{G_j}\right\|_2 .\end{aligned}$$

From this we can infer that $\tilde{\beta}_{G_j} = 0$ for all $j \in S^c$ because $\|\hat{z}_{S^c}\|_* < 1$. □

3.5 Consistent support recovery of group Lyapunov lasso

The goal of this section is to establish sufficient conditions under which the group Lyapunov lasso correctly recovers the undirected structure of a GCLM. In other words, we are interested in finding conditions that guarantee the correct identification of the support of the population drift matrix $S(M^*)$.

Definition 3.29. We define the support of a drift matrix M as

$$S = S(M) := \{j : \mathrm{vec}(M)_{G_j} \neq 0,\ j \in [m]\}.$$

The population drift matrix M^* is a parameter of the continuous Lyapunov equation (1.3) and $\hat{M}$ denotes the solution of the group lasso problem defined in (3.9). For ease of notation, let $S \equiv S(M^*)$ be the support of the population drift matrix M^* and $\hat{S} \equiv S(\hat{M})$ the support of the estimate $\hat{M}$.

Example 3.30. Let M be a matrix defined as

$$M = \begin{pmatrix} -2 & 0 & 0 \\ 0.1 & -3 & 0.1 \\ 0 & 0 & -4 \end{pmatrix}$$

and assume we use the partition $G_{\text{offdiag}} \cup G_{\text{diag}}$ defined in (3.7) and (3.8), i.e.

$$G_{\text{offdiag}} = \{\{(1,2),(2,1)\},\{(1,3),(3,1)\},\{(2,3),(3,2)\}\} = \{G_2, G_3, G_5\}$$

and

$$G_{\text{diag}} = \{\{(1,1)\},\{(2,2)\},\{(3,3)\}\} = \{G_1, G_4, G_6\}$$

then the support is

$$S(M) = \{1, 2, 4, 5, 6\}.$$

We introduce some more relevant notations.

Definition 3.31. For $A \in \mathbb{R}^{p\times p}$, the matrix $A_S = A_{\cdot S}$ is obtained from A by selecting the columns specified by $\bigcup_{j\in S} G_j$. We order the columns by $(i,j) < (k,l)$ if $i < k$ and $(i,j) < (k,l)$ if $i = k$ and $j < l$. Analogously, we can define the matrices $A_{S\cdot}$ and A_{SS}. By contrast, the matrix A_{-S} is defined by removing the columns specified in S.

Furthermore, let Σ^* and $\hat{\Sigma}$ be the associated covariance matrices with M^* and $\hat{M}$, obtained by solving the continuous Lyapunov equation (2.1) and using the formula for the sample covariance matrix (3.3), respectively. Then $\Gamma^* = \Gamma(\Sigma^*)$,$\hat{\Gamma} = \Gamma(\hat{\Sigma})$, $g^* = g(\Sigma^*)$ and $\hat{g} = g(\hat{\Sigma})$. Moreover, we define the differences $\Delta_\Gamma = \hat{\Gamma} - \Gamma^*$ and $\Delta_g = \hat{g} - g^*$. The constants

$$c_{\Gamma^*} = \left|\!\left|\!\left|(\Gamma^*_{SS})^{-1}\right|\!\right|\!\right|_* \quad \text{and} \quad c_{M^*} = \|\mathrm{vec}(M^*)\|_*$$

appear in the statement of our theorem as well.

For the proof of the theorem, we adapt the ideas presented in Lin et al. (2016) and Dettling et al. (2022). In particular, we use the specific PDW construction presented in both papers.

Theorem 3.32. *Let $M^* \in Stab_p$ be the true signal with support set S. Assume Γ^*_{SS} is invertible and the group-irrepresentable condition*

$$\left|\!\left|\!\left|\Gamma^*_{S^cS}(\Gamma^*_{SS})^{-1}\right|\!\right|\!\right|_* \leqslant \sqrt{d}(1-\alpha) \tag{3.28}$$

holds with parameter $\alpha \in [0,1)$, where d is the constant group size on S^c. Furthermore, assume that $\hat{\Gamma}$ is a matrix, such that

$$\left|\!\left|\!\left|(\Delta_\Gamma)_{\cdot S}\right|\!\right|\!\right|_* < \epsilon_1 \leqslant \frac{\alpha}{\left(3+\frac{3}{\sqrt{d}}\right)c_{\Gamma^*}} \quad \textit{and} \quad \|\Delta_g\|_* < \epsilon_2.$$

For

$$\lambda > \left(1 + \frac{1}{\sqrt{d}} - \alpha\right)\frac{3}{\alpha}\max\{\epsilon_1 c_{M^*}, \epsilon_2\},$$

$\hat{M}$ has its support $\hat{S}$ included in the true support S, i.e., $\hat{S} \subseteq S$ and $\hat{M}$ satisfies

$$\left\|vec(\hat{M}) - vec(M^*)\right\|_* \leqslant \left(\frac{1+\frac{1}{\sqrt{d}}}{1+\frac{1}{\sqrt{d}}-\alpha}\right)c_{\Gamma^*}\lambda.$$

In addition, if

$$\min_{j\in S}\left\|M^*_{G_j}\right\|_* > \left(\frac{1+\frac{1}{\sqrt{d}}}{1+\frac{1}{\sqrt{d}}-\alpha}\right)c_{\Gamma^*}\lambda,$$

then $\hat{S} = S$.

Proof. We will use the PDW method to prove the theorem. First, note that our optimization problem (3.9) is convex since $\hat{\Gamma}$ is positive semi-definite by construction. Recall that the KKT conditions are then necessary and sufficient conditions for a solution to be optimal.
The KKT conditions are given by the following sub-gradient equations

$$\hat{\Gamma}\mathrm{vec}(\hat{M}) - \hat{g} + \lambda D\hat{z} = 0, \tag{3.29}$$

where $\hat{z} \in \partial \left\|\mathrm{vec}(\hat{M})\right\|_g$ is a sub-gradient vector with entries

$$\hat{z}_{G_j} = \begin{cases} = \dfrac{\mathrm{vec}(\hat{M})_{G_j}}{\|\mathrm{vec}(\hat{M})_{G_j}\|_2} & j \in S, \\ \in \{x \in \mathbb{R}^{p_j} : \|x\|_2 \leqslant 1\} & j \in S^c. \end{cases} \tag{3.30}$$

Recall that by Theorem 2.9 choosing $M^* \in \mathrm{Stab}(p)$ and $C \in \mathrm{PD}(p)$ yields a unique positive-definite Σ^* determined by the continuous Lyapunov equation, i.e.,

$$\Gamma^*\mathrm{vec}(M^*) - g^* = 0. \tag{3.31}$$

The goal of the PDW technique is to establish a primal-dual pair $(\hat{M}, \hat{z})$, which solves (3.29) and (3.31) with the correct support structure. As we have seen earlier, this consists of three steps.

1. Solve the restricted optimization problem

$$\mathrm{vec}(\tilde{M}) = \underset{\mathrm{vec}(M)_{S^c=0}}{\arg\min} \frac{1}{2}\mathrm{vec}(M)^\mathsf{T}\hat{\Gamma}\mathrm{vec}(M) - \hat{g}^\mathsf{T}\mathrm{vec}(M) + \lambda\|\mathrm{vec}(M)\|_g$$

.

2. Set $\tilde{z}_S = \left(\frac{\mathrm{vec}(\tilde{M})_{G_j}}{\|\mathrm{vec}(\tilde{M})_{G_j}\|_2}\right)_{j\in S}$. Then $\tilde{z}_S \in ||\mathrm{vec}(\tilde{M})||_g$.

We define $\Delta_\Gamma := \hat{\Gamma} - \Gamma^*$, $\tilde{\Delta}_M := \mathrm{vec}(\tilde{M}) - \mathrm{vec}(M^*)$ and $\Delta_g := \hat{g} - g^*$. Using (3.29) and (3.31) we obtain

$$\Gamma^*\tilde{\Delta}_M + \Delta_\Gamma\mathrm{vec}(\tilde{M}) - \Delta_g + \lambda D\tilde{z} = 0. \tag{3.32}$$

We can rewrite the above equation (3.32) in block-matrix form

$$\begin{pmatrix} \Gamma^*_{SS} & \Gamma^*_{SS^c} \\ \Gamma^*_{S^cS} & \Gamma^*_{S^cS^c} \end{pmatrix} \begin{pmatrix} (\tilde{\Delta}_M)_S \\ 0 \end{pmatrix} + \begin{pmatrix} (\Delta_\Gamma)_{SS} & (\Delta_\Gamma)_{SS^c} \\ (\Delta_\Gamma)_{S^cS} & (\Delta_\Gamma)_{S^cS^c} \end{pmatrix} \begin{pmatrix} \mathrm{vec}(\tilde{M})_S \\ 0 \end{pmatrix} - \begin{pmatrix} (\Delta_g)_S \\ (\Delta_g)_{S^c} \end{pmatrix} + \lambda \begin{pmatrix} D_S\tilde{z}_S \\ D_{S^c}\tilde{z}_{S^c} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}.$$

This can then be solved for

$$\tilde{z}_{S^c} = \frac{1}{\lambda}D_{S^c}\left[-\Gamma^*_{S^cS}(\tilde{\Delta}_M)_S - (\Delta_\Gamma)_{S^cS}\mathrm{vec}(\tilde{M})_S + (\Delta_g)_{S^c}\right]$$

and

$$(\tilde{\Delta}_M)_S = (\Gamma^*_{SS})^{-1}\left[-(\Delta_\Gamma)_{SS}\text{vec}(\tilde{M})_S + (\Delta_g)_S - \lambda D_S \tilde{z}_S\right].$$

Substituting $(\tilde{\Delta}_M)_S$ in $\tilde{z}_{S^c}$ yields our third and final step. Note this step differs from the general PDW construction for the group lasso, as our problem does not have an additive error.

3. We have

$$\begin{aligned}\tilde{z}_{S^c} = \frac{1}{\lambda} D^{-1}_{S^c}\Big[\Gamma^*_{S^cS}\left(\Gamma^*_{SS}\right)^{-1}\Big((\Delta_\Gamma)_{SS}\text{vec}(\tilde{M})_S - (\Delta_g)_S\Big) \\ -(\Delta_\Gamma)_{S^cS}\text{vec}(\tilde{M})_S + (\Delta_g)_{S^c}\Big] \\ + D^{-1}_{S^c}\Gamma^*_{S^cS}\left(\Gamma^*_{SS}\right)^{-1} D_S\tilde{z}_S.\end{aligned} \tag{3.33}$$

According to Lemma 3.28 we need to show that $\|\tilde{z}_{S^c}\|_* < 1$ to prove $\hat{S} \subseteq S$. Using the triangle inequality we obtain the following bound

$$\begin{aligned}\|\tilde{z}_{S^c}\|_* \leqslant & \frac{1}{\lambda}\left\|\!\left|D^{-1}_{S^c}\right|\!\right\|_*\Big[\left\|\!\left|\Gamma^*_{S^cS}\left(\Gamma^*_{SS}\right)^{-1}\right|\!\right\|_*\Big(\left\|(\Delta_\Gamma)_{SS}\text{vec}(\tilde{M})_S\right\|_* + \|(\Delta_g)_S\|_*\Big) \\ & + \left\|(\Delta_\Gamma)_{S^cS}\text{vec}(\tilde{M})_S\right\|_* + \|(\Delta_g)_{S^c}\|_*\Big] \\ & + \left\|\!\left|D^{-1}_{S^c}\right|\!\right\|_*\left\|\!\left|\Gamma^*_{S^cS}\left(\Gamma^*_{SS}\right)^{-1}\right|\!\right\|_*\|D_S\tilde{z}_S\|_*.\end{aligned}$$

Looking at the individual dual norm terms, we can further simply the above expression. First note that on S^c all groups are of size d, thus

$$\left\|\!\left|D^{-1}_{S^c}\right|\!\right\|_* = \max_{\|\omega\|_*=1}\left\|\frac{1}{\sqrt{d}}\omega\right\|_* = \frac{1}{\sqrt{d}}.$$

Furthermore, we observe that

$$\|D_S\tilde{z}_S\|_* = \max_{j\in S}\sqrt{p_j}\frac{\|\tilde{z}_{G_j}\|_2}{\sqrt{p_j}} = \max_{j\in S}\frac{\left\|\text{vec}(\tilde{M}_{G_j})\right\|_2}{\left\|\text{vec}(\tilde{M}_{G_j})\right\|_2} = 1.$$

By assumption we have $\left\|\!\left|\Gamma^*_{S^cS}\left(\Gamma^*_{SS}\right)^{-1}\right|\!\right\|_* \leqslant \sqrt{d}(1-\alpha)$. Therefore,

$$\begin{aligned}\|\tilde{z}_{S^c}\|_* \leqslant & \frac{1}{\lambda}\Big[(1-\alpha)\Big(\left\|(\Delta_\Gamma)_{SS}\text{vec}(\tilde{M})_S\right\|_* + \|(\Delta_g)_S\|_*\Big) \\ & + \frac{1}{\sqrt{d}}\Big(\left\|(\Delta_\Gamma)_{S^cS}\text{vec}(\tilde{M})_S\right\|_* + \|(\Delta_g)_{S^c}\|_*\Big)\Big] \\ & + (1-\alpha).\end{aligned}$$

Definition 3.21 of the dual group norm yields that

$$\max\left\{\left\|(\Delta_\Gamma)_{SS}\text{vec}(\tilde{M})_S\right\|_*, \left\|(\Delta_\Gamma)_{S^cS}\text{vec}(\tilde{M})_S\right\|_*\right\} \leqslant \left\|(\Delta_\Gamma)_{\cdot S}\text{vec}(\tilde{M})_S\right\|_*$$

and $\max\{\|(\Delta_g)_S\|_*, \|(\Delta_g)_{S^c}\|_*\} \leqslant \|\Delta_g\|_*$. Hence,

$$\begin{aligned}
\|\tilde{z}_{S^c}\|_* &\leqslant \frac{1+\frac{1}{\sqrt{d}}-\alpha}{\lambda}\left[\left\|(\Delta_\Gamma)_{\cdot S}\mathrm{vec}(\tilde{M})_S\right\|_* + \|\Delta_g\|_*\right] + (1-\alpha) \\
&= \frac{1+\frac{1}{\sqrt{d}}-\alpha}{\lambda}\left[\left\|(\Delta_\Gamma)_{\cdot S}\left(\mathrm{vec}(M^*)_S + (\tilde{\Delta}_M)_S\right)\right\|_* + \|\Delta_g\|_*\right] + (1-\alpha) \\
&\leqslant \frac{1+\frac{1}{\sqrt{d}}-\alpha}{\lambda}\left[|||(\Delta_\Gamma)_{\cdot S}|||_*\,\|\mathrm{vec}(M^*)_S\|_* + |||(\Delta_\Gamma)_{\cdot S}|||_*\left\|(\tilde{\Delta}_M)_S\right\|_* + \|\Delta_g\|_*\right] + (1-\alpha).
\end{aligned} \tag{3.34}$$

We can bound the first summand by

$$\frac{1+\frac{1}{\sqrt{d}}-\alpha}{\lambda}|||(\Delta_\Gamma)_{\cdot S}|||_*\,\|\mathrm{vec}(M^*)_S\|_* \leqslant \frac{1+\frac{1}{\sqrt{d}}-\alpha}{\lambda}\epsilon_1 c_{M^*} < \frac{\alpha}{3}, \tag{3.35}$$

where the last equality follows from our choice of λ.
By assumption $\|\Delta_g\|_* < \epsilon_2$, thus we deduce the following bound for the third summand

$$\frac{1+\frac{1}{\sqrt{d}}-\alpha}{\lambda}\|\Delta_g\|_* < \frac{1+\frac{1}{\sqrt{d}}-\alpha}{\lambda}\epsilon_2 < \frac{\alpha}{3}. \tag{3.36}$$

It remains to bound the second summand or, to be more precise, $\left\|(\tilde{\Delta}_M)_S\right\|_*$. Analogously to before, we can combine (3.29) and (3.31) (this time adding the null sum $\hat{\Gamma}\mathrm{vec}(M^*) - \hat{\Gamma}\mathrm{vec}(M^*)$) to obtain

$$\hat{\Gamma}\tilde{\Delta}_M + \Delta_\Gamma \mathrm{vec}(M_*) - \Delta_g + \lambda D\tilde{z} = 0.$$

Note that

$$\tilde{\Delta}_M = \begin{pmatrix}(\tilde{\Delta}_M)_S \\ 0\end{pmatrix} \quad \text{and} \quad \mathrm{vec}(M^*) = \begin{pmatrix}\mathrm{vec}(M^*)_S \\ 0\end{pmatrix}.$$

Thus solving for $(\tilde{\Delta}_M)_S$ we have

$$(\tilde{\Delta}_M)_S = (\hat{\Gamma}_{SS})^{-1}\left(-(\Delta_\Gamma)_{SS}\mathrm{vec}(M^*)_S + (\Delta_g)_S - \lambda D_S\tilde{z}_S\right).$$

Consequently,

$$\begin{aligned}
\left\|(\tilde{\Delta}_M)_S\right\|_* &\leqslant \left|\left|\left|(\hat{\Gamma}_{SS})^{-1}\right|\right|\right|_*\left(|||(\Delta_\Gamma)_{SS}|||\,\|\mathrm{vec}(M^*)_S\| + \|(\Delta_g)_S\|_* + \lambda\|D_S\tilde{z}_S\|_*\right) \\
&\leqslant \left|\left|\left|(\hat{\Gamma}_{SS})^{-1}\right|\right|\right|_*(\epsilon_1 c_{M^*} + \epsilon_2 + \lambda).
\end{aligned}$$

Additionally, using our lower bounds for λ gives us

$$\begin{aligned}
\left\|(\tilde{\Delta}_M)_S\right\|_* &\leqslant \left|\left|\left|(\hat{\Gamma}_{SS})^{-1}\right|\right|\right|_*\left(\frac{1}{1+\frac{1}{\sqrt{d}}-\alpha}\frac{2\alpha}{3}\lambda + \lambda\right) \\
&= \left|\left|\left|(\hat{\Gamma}_{SS})^{-1}\right|\right|\right|_*\left(\frac{1+\frac{1}{\sqrt{d}}-\frac{\alpha}{3}}{1+\frac{1}{\sqrt{d}}-\alpha}\lambda\right).
\end{aligned} \tag{3.37}$$

The next step now is to find an upper bound for $\left\|\left|(\hat{\Gamma}_{SS})^{-1}\right|\right\|_*$. For this purpose observe that by assumption

$$|||(\Delta_\Gamma)_{SS}|||_* \leqslant |||(\Delta_\Gamma)_{\cdot S}|||_* < \epsilon_1 \leqslant \frac{\alpha}{c_{\Gamma^*}(3+\frac{3}{\sqrt{d}})} < \frac{1}{c_{\Gamma^*}}$$

therefore implying

$$|||(\Gamma^*_{SS})^{-1}(\Delta_\Gamma)_{SS}|||_* \leqslant |||(\Gamma^*_{SS})^{-1}|||_* \, |||(\Delta_\Gamma)_{SS}|||_* < 1. \tag{3.38}$$

We aim to bound the dual group norm of the inverse of the estimated Gram matrix $\left|\left|\left|(\hat{\Gamma}_{SS})^{-1}\right|\right|\right|_*$ by $|||(\Gamma^*_{SS})^{-1}|||_*$. Consider the difference

$$\begin{aligned}
(\hat{\Gamma}_{SS})^{-1} - (\Gamma^*_{SS})^{-1} &= \left[\Gamma^*_{SS} + (\Delta_\Gamma)_{SS}\right]^{-1} - (\Gamma^*_{SS})^{-1} \\
&= \left[\Gamma^*_{SS}(I + (\Gamma^*_{SS})^{-1}(\Delta_\Gamma)_{SS})\right]^{-1} - (\Gamma^*_{SS})^{-1} \\
&= \left\{\left[I + (\Gamma^*_{SS})^{-1}(\Delta_\Gamma)_{SS}\right]^{-1} - I\right\} (\Gamma^*_{SS})^{-1}.
\end{aligned}$$

(3.38) allows us to express the above difference as the Neumann series (Horn and Johnson (2013), Corollary 5.6.16)

$$(\hat{\Gamma}_{SS})^{-1} - (\Gamma^*_{SS})^{-1} = \left[-(\Gamma^*_{SS})^{-1}(\Delta_\Gamma)_{SS} \sum_{k=0}^{\infty} \left(-(\Gamma^*_{SS})^{-1}(\Delta_\Gamma)_{SS}\right)^k\right] (\Gamma^*_{SS})^{-1}.$$

This yields for the dual group norm difference

$$\begin{aligned}
\left|\left|\left|(\hat{\Gamma}_{SS})^{-1} - (\Gamma^*_{SS})^{-1}\right|\right|\right|_* &\leqslant |||(\Gamma^*_{SS})^{-1}(\Delta_\Gamma)_{SS}|||_* \sum_{k=0}^{\infty} \left(|||(\Gamma^*_{SS})^{-1}(\Delta_\Gamma)_{SS}|||_*\right)^k |||(\Gamma^*_{SS})^{-1}|||_* \\
&= \frac{|||(\Gamma^*_{SS})^{-1}(\Delta_\Gamma)_{SS}|||_* \, |||(\Gamma^*_{SS})^{-1}|||_*}{1 - |||(\Gamma^*_{SS})^{-1}(\Delta_\Gamma)_{SS}|||_*}.
\end{aligned}$$

For the equality, we used that $\sum_{k=0}^{\infty} \left(|||(\Gamma^*_{SS})^{-1}(\Delta_\Gamma)_{SS}|||_*\right)^k$ is a convergent geometric series due to (3.38). Then

$$\begin{aligned}
\left|\left|\left|(\hat{\Gamma}_{SS})^{-1}\right|\right|\right|_* &\leqslant |||((\Gamma^*_{SS})^{-1}|||_* + \left|\left|\left|(\hat{\Gamma}_{SS})^{-1} - (\Gamma^*_{SS})^{-1}\right|\right|\right|_* \\
&\leqslant \frac{|||((\Gamma^*_{SS})^{-1}|||_*}{1 - |||(\Gamma^*_{SS})^{-1}(\Delta_\Gamma)_{SS}|||_*} \leqslant \frac{|||((\Gamma^*_{SS})^{-1}|||_*}{1 - |||(\Gamma^*_{SS})^{-1}|||_* \, |||(\Delta_\Gamma)_{SS}|||_*} \\
&\leqslant \frac{c_{\Gamma^*}}{1 - c_{\Gamma^*}\epsilon_1} \leqslant \frac{c_{\Gamma^*}}{1 - \frac{\alpha}{3+\frac{3}{\sqrt{d}}}}.
\end{aligned} \tag{3.39}$$

We can substitute (3.39) into (3.37) to find the new bound

$$\left\|\tilde{\Delta}_M\right\|_* = \left\|(\tilde{\Delta}_M)_S\right\|_* \leqslant \frac{c_{\Gamma^*}}{1 - \frac{\alpha}{3+\frac{3}{\sqrt{d}}}} \left(\frac{1 + \frac{1}{\sqrt{d}} - \frac{\alpha}{3}}{1 + \frac{1}{\sqrt{d}} - \alpha}\lambda\right) = \left(\frac{1 + \frac{1}{\sqrt{d}}}{1 + \frac{1}{\sqrt{d}} - \alpha}\right) c_{\Gamma^*}\lambda,$$

which also proves the second claim of the theorem.
Now it is straightforward to see that

$$\frac{1+\frac{1}{\sqrt{d}}-\alpha}{\lambda} |||(\Delta_\Gamma)_{\cdot S}|||_* \left\|(\tilde{\Delta}_M)_S\right\|_* < \frac{\alpha}{3}. \tag{3.40}$$

Inserting (3.35), (3.36) and (3.40) into (3.34) yields

$$\|\tilde{z}_{S^c}\| < \frac{\alpha}{3} + \frac{\alpha}{3} + \frac{\alpha}{3} + (1-\alpha) = 1,$$

thus establishing strict dual feasibility.

Lastly, to prove the last claim of the theorem for any $j \in S$, the reverse triangle inequality implies

$$\left\|\hat{M}_{G_j}\right\|_* \geqslant \left\|M^*_{G_j}\right\|_* - \left\|\hat{M}_{G_j} - M^*_{G_j}\right\|_* \geqslant \min_{j\in S}\left\|M^*_{G_j}\right\|_* - \left\|\mathrm{vec}(\hat{M}) - \mathrm{vec}(M^*)\right\|_* > 0.$$

An immediate consequence is that $\hat{S} = S$. □

Remark 3.33. Dettling et al. (2022) investigated irrepresentability conditions for the direct Lyapunov lasso, which solves the following optimization problem

$$\min_{M\in\mathbb{R}^{p\times p}} \frac{1}{2}\left\|A(\hat{\Sigma})\mathrm{vec}(M) + \mathrm{vec}(C)\right\|_2^2 + \lambda\left\|\mathrm{vec}(M)\right\|_1.$$

They obtained the irrepresentability condition

$$|||\Gamma^*_{S^cS}\left(\Gamma^*_{SS}\right)^{-1}|||_\infty \leqslant (1-\alpha).$$

Note Theorem 3.32 contains the direct Lyapunov lasso as a special case since for groups that of size $p_j \equiv 1$, the group lasso is equal to the lasso. Moreover, recall from Remark 3.22 for a partition containing only singletons the group norm equals the ℓ_1-norm and the associated dual group norm becomes the ℓ_∞-norm. Consequently, the group irrepresentability condition is then identical to the one given above. Furthermore, the bounds on Σ and M match the ones given in Dettling et al. (2022) (Theorem 3.2).

Remark 3.34. One can slightly relax the group irrepresentability condition if we elect by not penalizing the diagonal elements of M in the group lasso problem (cf. Definition 3.6). While the proof remains mostly the same, there are some notable differences. Firstly, observe that due to the KKT conditions every optimal dual variable $\hat{z}$ has to satisfy

$$\hat{z}_{G_j} = \begin{cases} = \dfrac{\mathrm{vec}(\hat{M})_{G_j}}{\|\mathrm{vec}(\hat{M})_{G_j}\|_2} & j\in S \text{ and } G_j \in G_{\mathrm{offdiag}} \\ = 0 & j \in S \text{ and } G_j \in G_{\mathrm{diag}} \\ \in \{x\in\mathbb{R}^2 : \|x\|_2 \leqslant 1\} & j \in S^c \end{cases}.$$

The 0 entry in $\hat{z}$ comes from taking the partial derivative of the penalty wrt. m_{ii} (with m_{ii} not appearing in the penalty for all $i \in [p]$). We can leverage this particular structure of $\hat{z}$ to find a more efficient bound for

$$\begin{aligned}\left\| D_{S^c}^{-1} \Gamma^*_{S^c S} \left(\Gamma^*_{SS}\right)^{-1} D_S \tilde{z}_S \right\|_* &\leqslant \left\|\!\left| D_{S^c}^{-1} \right|\!\right\|_* \left\| \Gamma^*_{S^c S} \left(\Gamma^*_{SS}\right)^{-1} D_S \tilde{z}_S \right\|_* \\ &= \frac{1}{\sqrt{d}} \left\| \Gamma^*_{S^c S} \left((\Gamma^*_{SS})^{-1} \right)_{-S_{\text{diag}}} (D_S \tilde{z}_S)_{-S_{\text{diag}}} \right\|_* \\ &\leqslant \frac{1}{\sqrt{d}} \left\|\!\left| \Gamma^*_{S^c S} \left((\Gamma^*_{SS})^{-1} \right)_{-S_{\text{diag}}} \right|\!\right\|_*\end{aligned}$$

where $S_{\text{diag}} := \{j \in S : G_j \in G_{\text{diag}}\}$.
The new group-irrepresentability condition therefore is

$$\left\|\!\left| \Gamma^*_{S^c S} \left((\Gamma^*_{SS})^{-1} \right)_{-S_{\text{diag}}} \right|\!\right\|_* \leqslant \sqrt{d}(1-\alpha),$$

which is smaller than or equal to the group-irrepresentability condition given in (3.28) by definition of $|\!|\!|\cdot|\!|\!|_*$.

Theorem 3.32 provides a deterministic result on estimation error and support recovery under some conditions on Δ_Γ and Δ_g. Assuming our data is sub-Gaussian, we can provide more specific bounds on the preceding quantities. We start by introducing a series of lemmas, derived in Dettling et al. (2022) with the goal of bounding Δ_Γ in terms of $\Delta_\Sigma = \hat{\Sigma} - \Sigma^*$. The key idea here is to leverage the special structure of the Gram matrix Γ (cf. Lemma 3.35) to obtain a sharp bound and thus a good sample size requirement.

Lemma 3.35. For a given covariance matrix $\Sigma \in \mathbb{R}^p$ the Gram matrix $\Gamma(\Sigma) \in \mathbb{R}^{p^2 \times p^2}$ is equal to

$$\Gamma(\Sigma) = A(\Sigma)^\mathsf{T} A(\Sigma) = 2(\Sigma^2 \otimes I_p) + (\Sigma \otimes \Sigma) K^{(p,p)} + K^{(p,p)} (\Sigma \otimes \Sigma).$$

Proof. Recall that for two matrices A and B, $(A \otimes B)^\mathsf{T} = (A^\mathsf{T} \otimes B^\mathsf{T})$ and $K^{(p,p)}$ is symmetric by definition (cf. (2.3)).

$$A(\Sigma)^\mathsf{T} A(\Sigma) = \left[(\Sigma \otimes I_p) + K^{(p,p)}(I_p \otimes \Sigma)\right] \left[(\Sigma \otimes I_p) + (I_p \otimes \Sigma) K^{(p,p)}\right].$$

The mixed-product property implies

$$= \Sigma^2 \otimes I_p + (\Sigma \otimes \Sigma) K^{(p,p)} + K^{(p,p)}(\Sigma \otimes \Sigma) + K^{(p,p)}(I_p \otimes \Sigma^2) K^{(p,p)}.$$

In addition, Corollary 2.7 yields

$$= 2(\Sigma^2 \otimes I_p) + (\Sigma \otimes \Sigma) K^{(p,p)} + K^{(p,p)}(\Sigma \otimes \Sigma).$$

□

From the above Lemma we can deduce that $\Gamma(\Sigma) = \Gamma_1(\Sigma) + \Gamma_2(\Sigma)$ with

$$\Gamma_1(\Sigma) := 2(\Sigma^2 \otimes I_p) \quad \text{and} \quad \Gamma_2(\Sigma) := (\Sigma \otimes \Sigma)K^{(p,p)} + K^{(p,p)}(\Sigma \otimes \Sigma).$$

The following two Lemmas provide bounds on the spectral norm of $\Gamma_1(\Sigma)$ and $\Gamma_2(\Sigma)$ in terms of the spectral norm of Σ. We define the constant $c_{\Sigma*} := |||\Sigma^*|||_2$.

Lemma 3.36.

$$\left|\left|\left|\Gamma_1(\hat{\Sigma}) - \Gamma_1(\Sigma^*)\right|\right|\right|_2 \leqslant 2\,|||\Delta_\Sigma|||_2^2 + 4c_{\Sigma*}\,|||\Delta_\Sigma|||_2\,.$$

Proof. The Kronecker product is distributive, hence

$$\left|\left|\left|\Gamma_1(\hat{\Sigma}) - \Gamma_1(\Sigma^*)\right|\right|\right|_2 = 2\left|\left|\left|(\hat{\Sigma}^2 - (\Sigma^*)^2) \otimes I_p)\right|\right|\right|_2\,.$$

From Theorem 4.2.15 in Horn and Johnson (1991) we can infer that $|||A \otimes B|||_2 = |||A|||_2\,|||B|||_2$, therefore

$$= 2\left|\left|\left|\hat{\Sigma}^2 - (\Sigma^*)^2)\right|\right|\right|_2 \leqslant 2\,|||\Delta_\Sigma^2|||_2 + 2\,|||\Delta_\Sigma \Sigma^*|||_2 + 2\,|||\Sigma^* \Delta_\Sigma|||_2\,.$$

Note that for a symmetric matrix, the spectral norm is equal to the absolute maximal eigenvalue. Moreover, the eigenvalues of a squared matrix are the squared eigenvalues of the original matrix. Thus, we find as claimed that

$$\leqslant 2\,|||\Delta_\Sigma|||_2^2 + 4c_{\Sigma*}\,|||\Delta_\Sigma|||_2\,.$$

□

Lemma 3.37.

$$\left|\left|\left|\Gamma_2(\hat{\Sigma}) - \Gamma_2(\Sigma^*)\right|\right|\right|_2 \leqslant 2\,|||\Delta_\Sigma|||_2^2 + 4c_{\Sigma*}\,|||\Delta_\Sigma|||_2\,.$$

Proof. Observe that by (2.3) the commutation matrix $K^{(p,p)}$ is an orthonormal matrix, implying $|||K^{(p,p)}|||_2 = 1$. Additionally, it holds that for orthonormal matrices Q and a general matrix A, $|||QA|||_2 = |||AQ|||_2 = |||A|||_2$. Thus,

$$\left|\left|\left|K^{(p,p)}(\hat{\Sigma} \otimes \hat{\Sigma} - \Sigma^* \otimes \Sigma^*)\right|\right|\right|_2 = \left|\left|\left|(\hat{\Sigma} \otimes \hat{\Sigma} - \Sigma^* \otimes \Sigma^*)K^{(p,p)}\right|\right|\right|_2 = \left|\left|\left|(\hat{\Sigma} \otimes \hat{\Sigma} - \Sigma^* \otimes \Sigma^*)\right|\right|\right|_2\,.$$

Applying the above equation and the same tricks as in the proof of Lemma 3.36 yields

$$\begin{aligned}
\left|\left|\left|\Gamma_2(\hat{\Sigma}) - \Gamma_2(\Sigma^*)\right|\right|\right|_2 &\leqslant 2\left|\left|\left|\hat{\Sigma} \otimes \hat{\Sigma} - \Sigma^* \otimes \Sigma^*\right|\right|\right|_2 \\
&= 2\,|||(\Delta_\Sigma + \Sigma^*) \otimes (\Delta_\Sigma + \Sigma^*) - \Sigma^* \otimes \Sigma^*|||_2 \\
&= 2\,|||\Delta_\Sigma \otimes \Delta_\Sigma + \Delta_\Sigma \otimes \Sigma^* + \Sigma^* \otimes \Delta_\Sigma + \Sigma^* \otimes \Sigma^* - \Sigma^* \otimes \Sigma^*|||_2 \\
&\leqslant 2\,|||\Delta_\Sigma \otimes \Delta_\Sigma|||_2 + 2\,|||\Delta_\Sigma \otimes \Sigma^*|||_2 + 2\,|||\Sigma^* \otimes \Delta_\Sigma|||_2 \\
&= 2\,|||\Delta_\Sigma|||_2^2 + 4c_{\Sigma*}\,|||\Delta_\Sigma|||_2\,.
\end{aligned}$$

□

Let $A \in \mathbb{R}^{p\times s}$ with $s := \sum_{j\in S} p_j$, then $|||A|||_* \leqslant \frac{\sqrt{s}}{\min_{j\in[m]}\sqrt{p_j}} |||A|||_2$ by Lemma 3.24. It follows that

$$|||(\Delta_\Gamma)_{\cdot S}|||_* \leqslant \sqrt{s}\, |||(\Delta_\Gamma)_{\cdot S}|||_2\,, \tag{3.41}$$

and by Lemma 3.36 and Lemma 3.37

$$\leqslant \sqrt{s}(4\, |||\Delta_\Sigma|||_2^2 + 8c_{\Sigma^*}\, |||\Delta_\Sigma|||_2). \tag{3.42}$$

We now shift our focus to bounding $\|\Delta_g\|_*$.

Lemma 3.38. Let $c_C := \|\mathrm{vec}(C)\|_2$, then

$$\|\Delta_g\|_* \leqslant 2c_C\, \|\Delta_\Sigma\|_2\,.$$

Proof. From Lemma 3.24 we can infer that

$$\begin{aligned}\|\Delta_g\|_* &\leqslant \|\Delta_g\|_2\\ &\leqslant c_C \left|\left|\left|\hat{\Sigma}\otimes I_p + (I_p\otimes\hat{\Sigma})K^{(p,p)} - \Sigma^*\otimes I_p - (I_p\otimes\Sigma^*)K^{(p,p)}\right|\right|\right|_2.\end{aligned}$$

Using the distributivity of the Kronecker product and orthonormality of $K^{(p,p)}$ we have

$$\leqslant c_C\left(\left|\left|\left|(\hat{\Sigma}-\Sigma^*)\otimes I_p\right|\right|\right|_2 + \left|\left|\left|I_p\otimes(\hat{\Sigma}-\Sigma^*)\right|\right|\right|_2\right).$$

As mentioned earlier $|||A\otimes B|||_2 = |||A|||_2\, |||B|||_2$. Therefore,

$$= 2c_C\, |||\Delta_\Sigma|||_2\,,$$

which completes the proof. □

The bounds derived in (3.42) and Lemma 3.38 both depend on the spectral norm of Δ_Σ. If we choose the latter to be sufficiently small, we can ensure that our assumptions for $|||(\Delta_\Gamma)_{\cdot S}|||_*$ and $\|\Delta_g\|_*$ in Theorem 3.32 are met.

Lemma 3.39. If we choose sample size n to be large enough such that

$$|||\Delta_\Sigma|||_2 = \left|\left|\left|\hat{\Sigma}-\Sigma^*\right|\right|\right|_2 < \min\left\{\frac{\epsilon_1}{\sqrt{s}(4+8c_{\Sigma^*})}, \frac{\epsilon_2}{2c_C}, 1\right\}.$$

Then

$$|||(\Delta_\Gamma)_{\cdot S}|||_* < \epsilon_1 \quad\text{and}\quad \|\Delta_g\|_* < \epsilon_2.$$

Proof. The result is an immediate consequence of (3.42), where $|||\Delta_\Sigma|||_2^2 \leqslant |||\Delta_\Sigma|||_2$, and Lemma 3.38. □

The following theorem provides a probabilistic upper bound on $|||\Delta_\Sigma|||_2$ (concentration inequality) under the assumption that $\{X_i\}_{i=1}^n$ are sub-Gaussian. We start by recalling the definition of sub-Gaussian random vectors.

Definition 3.40 (*sub-Gaussian random vector*)**.** Let $X \in \mathbb{R}^p$ be a random vector with mean zero. It is said to be *sub-Gaussian* with parameter at most σ if for each fixed $v \in \mathbb{K}^{p-1}$

$$\mathbb{E}[e^{\lambda\langle v,X\rangle}] \leqslant e^{\frac{\lambda^2\sigma^2}{2}} \quad \text{for all } \lambda \in \mathbb{R}, \tag{3.43}$$

where $\mathbb{K}^{p-1}$ denotes the closed unit ball in $\mathbb{R}^p$.

Example 3.41. A typical example of a sub-Gaussian random vector is the Gaussian random vector $X \sim \mathcal{N}(0, \Sigma)$. To see this, observe that

$$\langle v, X\rangle = v^\mathsf{T} X \sim \mathcal{N}(0, v^\mathsf{T}\Sigma v). \tag{3.44}$$

The moment generating function of $\langle v, X\rangle$ equals

$$\mathbb{E}[e^{\lambda\langle v,X\rangle}] = e^{(\lambda^2 v^\mathsf{T}\Sigma v)/2} \quad \text{for all } \lambda \in \mathbb{R}. \tag{3.45}$$

Lastly, $v^\mathsf{T}\Sigma v \leqslant |||\Sigma|||_2$ for all $v \in \mathbb{K}^{p-1}$ proves that X is indeed sub-Gaussian with parameter at most $\sigma^2 = |||\Sigma|||_2$.

Theorem 3.42. *Let $\{X_i\}_{i=1}^n$ be sub-Gaussian random vectors with parameter σ. Then there exists universal constants $\{c_j\}_{j=1}^3$ such that the sample covariance matrix $\hat{\Sigma}$ defined in* (3.3) *satisfies*

$$\mathbb{P}\left(\frac{\left|\left|\left|\hat{\Sigma} - \Sigma^*\right|\right|\right|_2}{\sigma^2} \geqslant c_1\left(\sqrt{\frac{p}{n}} + \frac{p}{n}\right) + \delta\right) \leqslant c_2 \exp(-c_3 n \min(\delta, \delta^2)) \quad \textit{for all } \delta \geqslant 0. \tag{3.46}$$

Proof. see proof of Theorem 6.5 in Wainwright (2019). $\square$

We can rewrite Theorem 3.42 in a more convenient form.

Corollary 3.43. Suppose $\{c_j\}_{j=1}^3$ are the universal constants from Theorem 3.42. We assume w.l.o.g. that $c_1 > 1$. Moreover, let $\{X_i\}_{i=1}^n \sim \mathcal{N}(0, \Sigma^*)$ i.i.d.. Then for any $\epsilon / c_{\Sigma^*} \in (4c_1\sqrt{p/n}, 2)$, we have

$$\mathbb{P}\left(\left|\left|\left|\hat{\Sigma} - \Sigma^*\right|\right|\right|_2 \geqslant \epsilon\right) \leqslant c_2 \exp\left(-\frac{c_3}{4c_{\Sigma^*}^2} n\epsilon^2\right).$$

Proof. Let $\delta := \frac{\epsilon}{2c_{\Sigma^*}}$. By assumption $\frac{p}{n} < \frac{\epsilon^2}{16c_1^2 c_{\Sigma^*}^2}$, thus

$$c_{\Sigma^*}\left[c_1\left(\sqrt{\frac{p}{n}} + \frac{p}{n}\right) + \delta\right] \leqslant \frac{\epsilon}{4} + \frac{\epsilon^2}{16c_1 c_{\Sigma^*}} + \frac{\epsilon}{2} < \frac{\epsilon}{4} + \frac{\epsilon}{8} + \frac{\epsilon}{2} < \epsilon,$$

where the second to last inequality follows from $c_1 > 1$ and $\epsilon / c_{\Sigma^*} < 2$.
The above inequality immediately yields

$$\mathbb{P}\left(\left|\left|\left|\hat{\Sigma} - \Sigma^*\right|\right|\right|_2 \geqslant \epsilon\right) \leqslant \mathbb{P}\left(\left|\left|\left|\hat{\Sigma} - \Sigma^*\right|\right|\right|_2 \geqslant c_{\Sigma^*}\left[c_1\left(\sqrt{\frac{p}{n}} + \frac{p}{n}\right) + \delta\right]\right).$$

We showed in Example 3.41 that a multivariate normal distribution is sub-Gaussian with parameter at most $\sigma^2 = c_{\Sigma^*}$. Furthermore, our choice of ϵ implies $\delta < 1$ and therefore $\delta^2 < \delta$. Hence, from Theorem 3.42 we can infer that

$$\leqslant c_2 \exp(-c_3 n \delta^2) = c_2 \exp\left(-\frac{c_3}{4c_{\Sigma^*}^2} n\epsilon^2\right).$$

□

Lastly, for sub-Gaussian data, we can now state a more explicit version of Theorem 3.32 with specific error bounds.

Theorem 3.44. *Assume the data is drawn from a p-dimensional Ornstein-Uhlenbeck process in equilibrium. Let $M^* \in \mathrm{Stab}(p)$ be the corresponding true signal with support set S and $C \in \mathrm{PD}(p)$. Moreover, assume the Gram matrix Γ^*_{SS} is invertible and the group irrepresentability condition (3.28) holds for $\alpha \in [0,1)$. For $s := \sum_{j\in S} p_j$, we define*

$$c_{\Sigma^*} := \left\|\Sigma^*\right\|_2, \quad c_C := \|\mathrm{vec}(C)\|, \quad c_* := \frac{\left(3+\frac{3}{\sqrt{d}}\right) c_{\Gamma^*}}{\alpha},$$
$$\hat{c} := \max\{\sqrt{s}(4+8c_{\Sigma^*}), 2c_C\}, \quad \tilde{c} := \frac{4c_{\Sigma^*}^2}{c_3 s}\hat{c}^2$$

with $\{c_j\}_{j=1}^3$ the universal constants from Theorem 3.42 and w.l.o.g $c_1 > 1$. If we choose $\tau_1 > (4c_1^2 c_3 \hat{c} p)/\log p$, $n > \tilde{c} s \log p^{\tau_1} \max\{1/(4c_{\Sigma^*}^2 \hat{c}), 1/\hat{c}, c_*^2\}$ and the regularization parameter

$$\lambda > \left(1+\frac{1}{\sqrt{d}}-\alpha\right)\frac{3}{\alpha}\sqrt{\frac{\tilde{c} s \log p^{\tau_1}}{n}} \max\{c_{M^*}, 1\},$$

then the following statements hold with probability at least $1 - c_2 p^{-\tau_1}$.

(i) The estimate $\hat{M}$ has its support $\hat{S}$ included in the true support S, i.e. $\hat{S} \subseteq S$ and satisfies

$$\left\|\mathit{vec}(\hat{M}) - \mathit{vec}(M^*)\right\|_* \leqslant \left(\frac{1+\frac{1}{\sqrt{d}}}{1+\frac{1}{\sqrt{d}}-\alpha}\right) c_{\Gamma^*}\lambda.$$

(ii) In addition, if

$$\min_{j\in S}\left\|M^*_{G_j}\right\|_* > \left(\frac{1+\frac{1}{\sqrt{d}}}{1+\frac{1}{\sqrt{d}}-\alpha}\right) c_{\Gamma^*}\lambda,$$

then $\hat{S} = S$.

Proof. Set $\epsilon = \sqrt{\tilde{c} s \log p^{\tau_1}/n}$. Under the assumption for τ_1, we have

$$\min\left\{\frac{\epsilon}{\sqrt{s}(4+8c_{\Sigma^*})}, \frac{\epsilon}{2c_C}\right\} = \frac{\epsilon}{\max\{\sqrt{s}(4+8c_{\Sigma^*}), 2c_C\}} = \sqrt{\tau_1}\frac{\sqrt{\tilde{c} s \log p/n}}{\hat{c}} > 4c_1 c_{\Sigma^*}\sqrt{\frac{p}{n}}.$$

On the other hand, $n > \tilde{c} s \log p^{\tau_1} \max\{1/(4c_{\Sigma*}^2 \hat{c}), 1/\hat{c}\}$ yields

$$\min\left\{\frac{\epsilon}{\sqrt{s}(4+8c_{\Sigma*})}, \frac{\epsilon}{2c_C}\right\} < \min\{2c_{\Sigma*}, 1\}. \tag{3.47}$$

Recall that the Ornstein-Uhlenbeck process in equilibrium has a stationary distribution that is multivariate normal with covariance matrix Σ^*. Hence, together with the inequalities above, we can apply Corollary 3.43 and obtain that

$$\begin{aligned}\mathbb{P}\left(\left|\!\left|\!\left|\hat{\Sigma} - \Sigma^*\right|\!\right|\!\right|_2 < \min\left\{\frac{\epsilon}{\sqrt{s}(4+8c_{\Sigma*})}, \frac{\epsilon}{2c_C}\right\}\right) &> 1 - c_2 \exp\left(-\frac{c_3}{4c_{\Sigma*}^2} n \min\left\{\frac{\epsilon}{\sqrt{s}(4+8c_{\Sigma*})}, \frac{\epsilon}{2c_C}\right\}^2\right) \\ &= 1 - c_2 p^{-\tau_1}.\end{aligned}$$

Therefore, by Lemma 3.39 and (3.47), we have that $|\!|\!|(\Delta_\Gamma)_{\cdot S}|\!|\!|_* < \epsilon$ and $\|\Delta_g\|_* < \epsilon$. Moreover, note that $n > \tilde{c} s \log p^{\tau_1} c_*^2$ implies

$$\epsilon < c_* = \frac{\left(3 + \frac{3}{\sqrt{d}}\right) c_{\Gamma*}}{\alpha}.$$

In conclusion, we have shown that we can attain the error bounds assumed in Theorem 3.32, and thus the remaining statements immediately follow from the preceding theorem. □

Closer inspection of Theorem 3.44 reveals that the group Lyapunov lasso requires a sample size of $n = \Omega(s \log p)$ to recover the support of M^* correctly, where s is the number of non-zero entries in M^*. This result is identical to the sample size requirement for the lasso in the Lyapunov setting (Dettling et al. (2022), Corollary 3.2). If the true signal is relatively sparse, s will be much smaller than the number of unknown parameters p^2. Consequently, Theorem 3.44 then states that the group Lyapunov lasso can also be employed in a high-dimensional setting.
However, compared to other undirected structure recovery methods it is less powerful for high-dimensional data. Consider, for example, the graphical lasso for multivariate normal data. Ravikumar et al. (2011) showed that consistent support recovery requires a sample size of the order $s^2 \log p$, where s is the number of non-zero entries in the rows of the true precision matrix. This is due to the precision matrix being simply the inverse of the covariance matrix. There is a straightforward one-to-one relation, whereas the covariance matrix and the drift matrix M^* for GCLMs don't exhibit a similarly simple connection (Dettling et al. (2022)).

3.6 Analysis of the irrepresentability condition

The irrepresentability condition (3.28) is a crucial assumption for Theorem 3.32. Unlike in standard lasso or group lasso regression, there exists no intuitive interpretation of the preceding quantity for GCLMs. For both classical regression settings given a design matrix X, the irrepresentability condition has the following form

$$\left\|\left\|(X^{\mathsf{T}}X)_{S^cS}(X^{\mathsf{T}}X)_{SS}^{-1}\right\|\right\| < 1,$$

where $\|\|\cdot\|\|$ is either the ℓ_∞-norm $\|\|\cdot\|\|_\infty$ or the dual group norm $\|\|\cdot\|\|_*$, respectively. Here the interpretation is that the predictors in the active set S are not strongly correlated with the variables in the inactive set S^c. For the Lyapunov model, the design matrix is replaced by $A(\Sigma)$, which is not only dependent on the predictors but also on the signal M itself. Recall, for example, the fixed structure of $A(\Sigma)$ given in Example 3.4, whose individual entries are given as the solution to the continuous Lyapunov equation (2.1). For the entirety of this section, we will assume that the volatility matrix C in the continuous Lyapunov equation is a multiple of the identity matrix. To be more precise, we will assume that $C = 2I_p$. The results presented here would also hold for other diagonal matrices C.

3.6.1 Irrepresentability condition for the direct Lyapunov lasso

Before we turn to the analysis of the irrepresentability condition for the group Lyapunov lasso, we will investigate the irrepresentability conditon for the direct Lyapunov lasso. The theory for this subsection was developed in Dettling et al. (2022).

First of, we define the following irrepresentability constant.

Definition 3.45 (*irrepresentability constant*)**.** For $M^* \in \mathrm{Stab}(p)$, the true signal with support set S, the *irrepresentability constant* is defined as

$$\rho(M^*) := \left\|\left\|\Gamma^*_{S^cS}\left(\Gamma^*_{SS}\right)^{-1}\right\|\right\|_\infty .$$

Note the definition for the support S slightly changes if we work in the lasso setting. In particular, as the lasso is equal to the group lasso if all considered groups are singletons, we have

$$S = S(M^*) = \{(i,j) : M^*_{ij} \neq 0\}.$$

The appropriate irrepresentability condition can then be obtained from Theorem 3.32 by setting $p_j = 1$ for all $j \in [m]$ and adjusting the norms (cf. Remark 3.33). Hence, we deduce that

$$\rho(M^*) \leqslant (1-\alpha) \quad \text{for } \alpha > 0.$$

The key idea to gain more insight to the behaviour of $\rho(M^*)$ is the following fact from standard lasso regression. Let X be a design matrix from a linear regression problem. Assuming the Gram matrix $X^{\mathsf{T}}X$ is diagonal, then $(X^{\mathsf{T}}X)_{S^cS} = 0$, and thus $\left|\!\left|\!\left|(X^\top X)_{S^cS}(X^\top X)^{-1}_{SS}\right|\!\right|\!\right|_\infty = 0$. The irrepresentability condition is trivially satisfied. Moreover, in a neighborhood around the diagonal matrix $X^\top X$ the irrepresentability condition will also hold true. We can adapt this idea for GCLMs by considering drift matrices M that yield a diagonal Gram matrix. This leads us to the following definition.

Definition 3.46 (*local irrepresentability constant*)**.** Let $G = (V, E)$ be a graph with associated support set $S_G := \{(i, j) : i \to j \in E\}$. For a diagonal matrix $M^0 \in \mathrm{Stab}(p)$ the *local irrepresentability constant* is defined as

$$\tilde{\rho}_G(M^0) := \left|\!\left|\!\left|\Gamma^0_{S_G^c S_G}(\Gamma^0_{S_G S_G})^{-1}\right|\!\right|\!\right|_\infty$$

with $\Gamma^0 := \Gamma(\Sigma^0)$ and Σ^0 obtained from solving the continous Lyapunov equation (1.3) with M^0.

The local irrepresentability constant is not well-defined for non-simple graphs.

Definition 3.47. A graph $G = (V, E)$ is *simple* if it contains no 2-cycle. In other words, there exist no two nodes $i, j \in V$ with $i \neq j$ such that both $i \to j$ and $j \to i$ are in E.

Lemma 3.48. Let $G = (V, E)$ be a directed graph with a 2-cycle and $M^0 \in \mathrm{Stab}(p)$ be an arbitrary diagonal matrix. Then $\tilde{\rho}_G(M^0)$ is not well-defined.

Proof. We start with a general fact that holds for $A(\Sigma)$. Let $i, j \in V$ be two nodes with $i \to j \in E$ and $j \to i \in E$. The columns of the design matrix $A(\Sigma)_{\cdot(i,j)}$ and $A(\Sigma)_{\cdot(j,i)}$ (recall the way we index columns of $A(\Sigma)$, cf. Example 3.4), will have diagonal entries of Σ in the same two positions. Moreover, the remaining entries will either be off-diagonal entries of Σ or 0.
Note that when we M^0 is diagonal, Σ^0 will, in turn, also be diagonal. In particular, this implies for the columns $A(\Sigma^0)_{\cdot(i,j)}$ and $A(\Sigma^0)_{\cdot(j,i)}$ that they are identical. Hence, $A(\Sigma^0)$ contains linearly dependent columns and $\Gamma^0_{S_G S_G} = A(\Sigma^0)^\top A(\Sigma^0)$ cannot be inverted, i.e., $\tilde{\rho}_G(M^0)$ is not well-defined. □

However, simply choosing M^0 to be a diagonal matrix will not always result in the Gram matrix Γ^0 being diagonal as the next example shows.

Example 3.49. Consider the 3-chain graph $\mathcal{G} = ([3], \{1 \to 2, 2 \to 3\})$ in Figure 3.1 and the diagonal matrix

$$M^0 = \mathrm{diag}(-d_1, -d_2, -d_3) \quad \text{with } d_i > 0 \text{ for } i \in [3].$$

M^0 is obviously stable, thus specifying $C = 2I_3$ yields the covariance matrix

$$\Sigma^0 = \mathrm{diag}(1/d_1, 1/d_2, 1/d_3)$$

as the unique solution to the Lyapunov equation in (1.3).
Moreover,

$$A(\Sigma^0) = \begin{pmatrix} 2/d_1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1/d_1 & 0 & 1/d_2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1/d_1 & 0 & 0 & 0 & 1/d_3 & 0 & 0 \\ 0 & 1/d_1 & 0 & 1/d_2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 2/d_2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1/d_2 & 0 & 1/d_3 & 0 \\ 0 & 0 & 1/d_1 & 0 & 0 & 0 & 1/d_3 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1/d_2 & 0 & 1/d_3 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 2/d_3 \end{pmatrix}$$

and

$$\Gamma^0 = \begin{pmatrix} 4/d_1^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 2/d_1^2 & 0 & 2/d_1 d_2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 2/d_1^2 & 0 & 0 & 0 & 2/d_1 d_3 & 0 & 0 \\ 0 & 2/d_1 d_2 & 0 & 2/d_2^2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 4/d_2^2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 2/d_2^2 & 0 & 2/d_2 d_3 & 0 \\ 0 & 0 & 2/d_1 d_3 & 0 & 0 & 0 & 2/d_3^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 2/d_2 d_3 & 0 & 2/d_3^2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 4/d_3^2 \end{pmatrix}.$$

Clearly, Γ^0 is not diagonal, and we cannot immediately conclude that the local irrepresentability condition is satisfied as in the classical lasso setting.

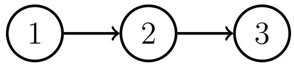

Figure 3.1: Directed graph on 3 nodes, where we omitted to draw the self-loops induced by M and C

We require one more step. Namely, we can establish a sufficient and necessary condition for the local irrepresentability constant to be met by ordering the diagonal entries in a certain way.

Theorem 3.50. *Let $G = ([p], E)$ be a simple directed graph. For any diagonal matrix $M^0 \in Stab(p)$ with diagonal entries $-d_i < 0$ for $i \in [p]$, $\tilde{\rho}(M^0) < 1$ if and only if $d_i < d_j$ for every edge $j \to i \in E$.*

Proof. Solving the Lyapunov equation (1.3) with M^0 and $C = 2I_p$, we obtain that

$$\Sigma^0 = \mathrm{diag}(1/d_1, \ldots, 1/d_p).$$

Let $i, j, k, l \in [p]$. Since Σ^0 is diagonal, $A(\Sigma^0) = (\Sigma^0 \otimes I_p) + (I_p \otimes \Sigma^0)K^{(p,p)}$ is symmetric as well. The entries of $A(\Sigma^0)$ are given as

$$A(\Sigma^0)_{(i,j),(k,l)} = \begin{cases} 2/d_k & \text{if } i = j = k = l, \\ 1/d_k & \text{if } i = k,\ j = l \text{ and } k \neq l, \\ 1/d_k & \text{if } i = l,\ j = k \text{ and } k \neq l, \\ 0 & \text{otherwise.} \end{cases}$$

Furthermore, recall that the entries of the Gram matrix Γ^0 are given by the inner products of the columns of $A(\Sigma^0)$, hence

$$\Gamma^0_{(i,j),(k,l)} = \begin{cases} 4/d_k^2 & \text{if } i = j = k = l, \\ 2/d_k^2 & \text{if } i = k,\ j = l \text{ and } k \neq l, \\ 2/(d_k d_l) & \text{if } i = l,\ j = k \text{ and } k \neq l, \\ 0 & \text{otherwise.} \end{cases}$$

Observe that Γ^0 has an off-diagonal entry only when the row index is (i, j) and the column index is (j, i) with $i \neq j$. By assumption we only consider simple graphs, therefore $\Gamma^0_{S_G S_G}$ is diagonal with entries

$$(\Gamma^0_{S_G S_G})_{(k,l),(k,l)} = \begin{cases} 4/d_k^2 & \text{if } k = l, \\ 2/d_k^2 & \text{if } k \neq l, \end{cases}$$

with $k \to l \in E$. The second submatrix appearing in the local irrepresentability constant $\Gamma^0_{S_G^c S_G}$ has exactly one non-zero entry in each column as well. Assuming that $k \to l \in E$, then this entry equals

$$(\Gamma^0_{S_G^c S_G})_{(l,k),(k,l)} = 2/(d_k d_l).$$

Combining the above equations, we get that

$$(\Gamma^0_{S_G^c S_G}(\Gamma^0_{S_G S_G})^{-1})_{(i,j),(k,l)} = \begin{cases} d_k/d_l & \text{if } (i,j) = (l,k) \text{ and } (k,l) \in S_G,\ (l,k) \in S_G^c, \\ 0 & \text{otherwise.} \end{cases}$$

$|||\cdot|||_\infty$ is defined as the maximum absolute row sum, thus $\tilde{\rho}(M^0) < 1$ if and only if $d_i/d_j < 1$ for all pairs $(i, j) \in S_G$, or equivalently, $j \to i \in E$. □

To illustrate the claims made in the proof above for a graph with 3 nodes, we refer to Example 3.49 above and Example 3.51 below.

Example 3.51. Consider the same setting as in Example 3.49. We have

$$S_G = \{(1,1), (2,1), (2,2), (3,2), (3,3)\}$$

and

$$S_G^c = \{(1,2), (1,3), (3,1), (2,3)\}.$$

Consequently,

$$(\Gamma^0_{S_G S_G})^{-1} = \operatorname{diag}(d_1^2/4, d_2^2/2, d_2^2/4, d_3^2/2, d_3^2/4),$$

$$\Gamma^0_{S_G^c S_G} = \begin{pmatrix} 0 & 2/d_1 d_2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2/d_2 d_3 & 0 \end{pmatrix},$$

and

$$\Gamma^0_{S_G^c S_G}(\Gamma^0_{S_G S_G})^{-1} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & d_2/d_1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & d_3/d_2 & 0 \end{pmatrix}.$$

From this we can immediately see that $\left|\!\left|\!\left| \Gamma^0_{S_G^c S_G}(\Gamma^0_{S_G S_G})^{-1} \right|\!\right|\!\right|_\infty = \max\{d_2/d_1, d_3/d_2\} < 1$ if and only if $d_2/d_1 < 1$ and $d_3/d_2 < 1$. Noting that $1 \to 2 \in E$ and $2 \to 3 \in E$, this is exactly what Theorem 3.50 requires.

Another difficulty in analyzing the local irrepresentability constant arises when we consider graphs containing directed cycles.

Corollary 3.52. Let $G = ([p], E)$ be a directed graph containing a directed cycle that is not a 2-cycle. For any diagonal matrix $M^0 \in \operatorname{Stab}(p)$, it holds that $\tilde{\rho}(M^0) \geqslant 1$.

Proof. Wlog. assume that we have an n-cycle $(1 \to n \to (n-1) \to \cdots \to 2 \to 1)$ with $n \leqslant p$. Similar to before denote the diagonal entries of M^0 by $-d_1, \ldots, -d_p$. Then for the local irrepresentability condition to be met Theorem 3.50 assumes

$$\frac{d_1}{d_2} < 1,\ \frac{d_2}{d_3} < 1, \ldots,\ \frac{d_{n-1}}{d_n} < 1,\ \frac{d_n}{d_1} < 1.$$

Yet this implies

$$\frac{d_n}{d_1} > \frac{d_n}{d_1}\frac{d_{n-1}}{d_n} = \frac{d_{n-1}}{d_1} > \cdots > \frac{d_2}{d_1}\frac{d_1}{d_2} = 1,$$

and thus by Theorem 3.50 again $\tilde{\rho}(M^0) \geqslant 1$. □

Hence, for the local irrepresentability constant to be smaller than 1, we should only consider directed graphs that are acylic.

Definition 3.53 (*DAG*)**.** A *directed acyclic graph* (DAG) is a directed graph $G = (V, E)$ that contains no cycles. Meaning for any node $i \in V$, there exists no sequence of edges in E such that there is a directed walk $i \to \cdots \to i$, excluding self-loops.

The next theorem shows how we can extend the local irrepresentability condition to the regular irrepresentability condition. Note we say that the irrepresentability condition for a support S holds *uniformly* over a set $U \subseteq \operatorname{Stab}(p)$, if there exists $\alpha > 0$ such that $\rho(M^*) \leqslant 1 - \alpha$ for all $M^* \in U$ with support $S = S(M^*)$.

Theorem 3.54. *Let $G = ([p], E)$ be a DAG and $M^0 = diag(-d_1, \dots, -d_p) \in Stab(p)$. Then the irrepresentability condition for S_G holds uniformly over a neighborhood of M^0 if and only if*

$$d_i < d_j \quad \text{for every edge } j \to i \in E.$$

Proof. By Theorem 3.50 $\tilde{\rho}(M^0) < 1$ if and only if $d_i < d_j$ for every edge $j \to i \in E$. We define the matrix function $M(e)$ to be the matrix with diagonal entries equal to M^0 and off-diagonal entries $e_1, \dots, e_{|E|-p} \in \mathbb{R}$. Moreover, let $D_{M(e)} \subseteq \mathbb{R}^{|E|-p}$ be the domain of $M(e)$ such that it is stable and $S(M(e)) = S_G$. The natural matrix function $\Gamma(e) := \Gamma(\Sigma(e))$ obtained by solving the Lyapunov equation (1.3) with $M(e)$ for $\Sigma(e)$ and plugging the latter into the formula for the design matrix $A(\Sigma(e))$ (cf. (3.4)) and computing the Gram matrix (cf. (3.6)). We can now define the function

$$\Phi : D_{M(e)} \to \mathbb{R}$$
$$e \to \begin{cases} \left|\!\left|\!\left| \Gamma_{S_G^c S_G}(e)(\Gamma_{S_G S_G}(e))^{-1} \right|\!\right|\!\right|_\infty & \text{for } e \neq 0 \\ \left|\!\left|\!\left| \Gamma^0_{S_G^c S_G}(\Gamma^0_{S_G S_G})^{-1} \right|\!\right|\!\right|_\infty & \text{for } e = 0 \end{cases},$$

which is continuous because Φ is a rational function. It is also well defined since $\Gamma_{S_G S_G}(e)$ is invertible for all DAGs. To see this we refer to the proof of Theorem 4.10, which states that $\det(A(\Sigma(e))_{S_G, S_G}) \neq 0$ and thus $\det(\Gamma_{S_G S_G}(e)) = \det(A(\Sigma(e))_{S_G, S_G})^2 \neq 0$.

As M^0 is chosen such that $\Phi(0) = \tilde{\rho}(M^0) < 1$ and Φ is continuous, we can find a small open ball O_M around $e = 0$ such that $\Phi(e) < 1$ for all $e \in O_M$. In summary, we have found a neighborhood U around M^0 such that $\rho(M^*) < 1$ for all $M^* \in U$, which proves the theorem. □

Remark 3.55. In other words, Theorem 3.54 states for any given DAG G, we can find an associated drift matrix M^* such that the irrepresentability condition is met. We start by ordering the diagonal elements of M^* to satisfy condition (3.54) in the following way. Recall that for each DAG we can establish a topological ordering, i.e., we can enumerate the nodes such that for any two nodes $i, j \in V$ with $i \neq j$

$$i \to j \text{ implies } i < j.$$

Consequently, reordering the diagonal elements such that $d_i < d_j$ for $j \to i \in E$, then satisfies (3.54). The final consists of suitably choosing the off-diagonal entries of M^*.

3.6.2 Irrepresentability condition for the group Lyapunov lasso

Throughout this section we assume $[p^2]$ to be partitioned according to $G_{\text{offdiag}} \cup G_{\text{diag}}$. There are a several challenges when it comes to analyzing the group irrepresentability condition

$$\left|\!\left|\!\left| \Gamma^*_{S^c S}(\Gamma^*_{SS})^{-1} \right|\!\right|\!\right|_* \leq \sqrt{2}(1 - \alpha) \quad \text{with } \alpha > 0$$

for the group Lyapunov lasso. We refer to $|\!|\!| \Gamma^*_{S^c S}(\Gamma^*_{SS})^{-1} |\!|\!|_*$ as the group irrepresentability constant.

Firstly, we start by noting that the group irrepresentability constant cannot be readily computed in closed form. Hence we provide an upper bound on the preceding quantity based on Lemma 3.23

$$\begin{aligned}\left\|\!\left|\Gamma^*_{S^cS}(\Gamma^*_{SS})^{-1}\right|\!\right\|_* &\leqslant \max_{i\in S^c}\sum_{j\in S}\sqrt{\frac{p_j}{p_i}}\left\|\!\left|(\Gamma^*_{S^cS}(\Gamma^*_{SS})^{-1})_{G_iG_j}\right|\!\right\|_2 \\ &= \max_{i\in S^c}\sum_{j\in S}\sqrt{\frac{p_j}{p_i}}\left\|\!\left|\sum_{k\in S}\Gamma^*_{G_iG_k}((\Gamma^*_{SS})^{-1})_{G_kG_j}\right|\!\right\|_2 . \end{aligned} \tag{3.48}$$

Bach (2008) suggests that the above bound (3.48) may be improved by solving a semidefinite programming problem. Recall that by definition of the dual group norm,

$$\left\|\!\left|\Gamma^*_{S^cS}(\Gamma^*_{SS})^{-1}\right|\!\right\|_* = \max_{i\in S^c}\left\|\!\left|\Gamma^*_{G_iS}(\Gamma^*_{SS})^{-1}\right|\!\right\|_* .$$

Lemma 3.56. The quantity $\left\|\!\left|\Gamma^*_{G_iS}(\Gamma^*_{SS})^{-1}\right|\!\right\|_*^2$ can be bounded from above by

$$\max_{X\succeq 0,\operatorname{tr}(X_{G_jG_j})=1, j\in S}\operatorname{tr}\left(X\left(D_S(\Gamma^*_{SS})^{-1}\Gamma^*_{SG_i}\Gamma^*_{G_iS}(\Gamma^*_{SS})^{-1}D_S\right)\right), \tag{3.49}$$

where X is a matrix defined by blocks following the block structure of Γ^*_{SS}. The blocks are defined as $\Gamma^*_{G_iG_j}$ with $i, j \in S$.

Proof. Let $X = uu^\top$ with $\|u_{G_j}\|_2 = 1$ for all $j \in S$. Then using basic properties of the trace,

$$\operatorname{tr}(X_{G_jG_j}) = \operatorname{tr}(u_{G_j}u_{G_j}^\top) = \operatorname{tr}(u_{G_j}^\top u_{G_j}) = \|u_{G_j}\|_2^2 = 1.$$

Moreover, by the symmetry of the Gram matrix Γ^*_{SS} we have that

$$\begin{aligned}\operatorname{tr}\left(uu^\top\left(D_S(\Gamma^*_{SS})^{-1}\Gamma^*_{SG_i}\Gamma^*_{G_iS}(\Gamma^*_{SS})^{-1}D_S\right)\right) &= \operatorname{tr}\left(u^\top D_S(\Gamma^*_{SS})^{-1}\Gamma^*_{SG_i}\Gamma^*_{G_iS}(\Gamma^*_{SS})^{-1}D_Su\right) \\ &= \left\|\Gamma^*_{G_iS}(\Gamma^*_{SS})^{-1}D_Su\right\|_2^2 .\end{aligned}$$

The claim then follows from

$$\left\|\!\left|\Gamma^*_{G_iS}(\Gamma^*_{SS})^{-1}\right|\!\right\|_*^2 = \max_{\|u_{G_j}\|_2=1, j\in S}\left\|\Gamma^*_{G_iS}(\Gamma^*_{SS})^{-1}D_Su\right\|_2^2 ,$$

since the RHS of the above equation is upper-bounded by (3.49). □

Secondly, the irrepresentability condition is only well-defined if the inverse of Γ^*_{SS} exists. The next lemma provides a necessary condition on the size $|s|$ of the support S for Γ^*_{SS} to be invertible.

Lemma 3.57. If $s = |\bigcup_{j\in S} G_j| > \frac{p(p+1)}{2}$, Γ^*_{SS} is not invertible.

Proof. We use the fact that for a matrix A, $\operatorname{rank}(A^\top A) = \operatorname{rank}(A)$. Thus

$$\operatorname{rank}(\Gamma^*_{SS}) = \operatorname{rank}((A(\Sigma^*)_{\cdot S})^\top A(\Sigma^*)_{\cdot S}) = \operatorname{rank}(A(\Sigma^*)_{\cdot S}) \tag{3.50}$$

$$\leqslant \operatorname{rank}(A(\Sigma^*)) \leqslant \frac{p(p+1)}{2}. \tag{3.51}$$

The last inequality follows from the symmetry of the Lyapunov equation. $A(\Sigma)$ contains two copies of each row corresponding to an off-diagonal entry in the Lyapunov equation. Consequently, $A(\Sigma)$ can have at most $\frac{p(p+1)}{2}$ linearly independent rows. Choosing $|\bigcup_{j\in S} G_j| > \frac{p(p+1)}{2}$ results in Γ^*_{SS} not having full rank, i.e., it is not invertible. □

Lemma 3.57 restricts the number and type of undirected structures for which an irrepresentability condition is well-defined. In Table 3.1, we illustrate how the number of vertices p relates to the maximal number of undirected edges ($\lfloor \frac{p(p-1)}{4} \rfloor$, assuming all diagonal entries of M are nonzero) for which the irrepresentability condition is still well-defined. Observe, for instance, that a graph containing only 2 vertices can not be investigated according to the framework proposed in Theorem 3.32.

p	2	3	4	5	6	10	20
maximal number of undirected edges	0	1	3	5	7	22	95

Table 3.1: Maximal number of undirected edges of $\mathcal{G}^{skel} = ([p], E)$ for which the irrepresentability condition is well-defined under the assumption that the diagonal entries of M are all nonzero

As we have discussed earlier for the direct Lyapunov lasso, the interpretation of the irrepresentability condition for GCLM is somewhat involved. In particular, we resort to analyzing a diagonal signal M^0 first to infer more general statements for sparse non-diagonal signals. However, this approach does not apply to the group Lyapunov lasso. The main issue is that the local irrepresentability constant is not well-defined since $\Gamma^0_{S_G S_G}$ is not invertible for any choice of the support set $S_G \neq \{(i,i) : i \in [p]\}$.
Assume we include one undirected edge $l - k$ with $l \neq k$, that is a single group $\{(l,k),(k,l)\}$ that corresponds to an off-diagonal entry $M_{lk} > 0$ or $M_{kl} > 0$ in M. Then $S_G = \{(i,i) : i \in [p]\} \cup \{(l,k),(k,l)\}$. Hence, by Lemma 3.48 $\Gamma^0_{S_G S_G}$ is not invertible as we have included a de facto 2-cycle. As every group defined by an off-diagonal signal entry is considered a 2-cycle in the viewpoint of Lemma 3.48, we cannot define the local irrepresentability condition for the group Lyapunov lasso.
Establishing a different starting point than a diagonal signal M^0 for a more in-depth analysis of the group irrepresentability condition proved to be quite difficult. The crux lies in describing $(\Gamma^0_{S_G S_G})^{-1}$ for non-diagonal M.

In general, it is rather difficult to find examples that meet the group irrepresentability condition or, to be more precise, yield that an upper bound for it that is strictly smaller than $\sqrt{2}$. One potential explanation is the presence of strongly correlated columns in the submatrix of the design matrix $A(\Sigma^*)_{\cdot S}$. For the regression setting, this is also referred to as multicollinearity. Collinearity causes the matrix $\Gamma^*_{SS} = (A(\Sigma^*)_{\cdot S})^\top A(\Sigma^*)_{\cdot S}$ to be ill-conditioned, i.e. close to being singular (Draper and Smith (1998)). Consequently, $(\Gamma^*_{SS})^{-1}$ will have very large entries, thus making it difficult to meet the group irrepresentability condition. We illustrate this problem with the following example.

Example 3.58. Consider the line graph $G = ([4], \{1 \to 2, 2 \to 3, 3 \to 4\})$ with associated drift matrix

$$M^* = \begin{pmatrix} -2 & 0 & 0 & 0 \\ 0.1 & -3 & 0 & 0 \\ 0 & 0.1 & -4 & 0 \\ 0 & 0 & 0.1 & -5 \end{pmatrix} \quad \text{and} \quad C = 2I_4.$$

We have support $S = \{1, 2, 5, 6, 8, 9, 10\}$ corresponding to the groups

$$\bigcup_{j \in S} G_j = \{\{(1,1)\}, \{(1,2), (2,1)\}, \{(2,2)\}, \{(2,3), (3,2)\}, \{(3,3)\}, \{(3,4), (4,3)\}, \{(4,4)\}\}.$$

Recall that the items in G_j represent the indices of the corresponding entries in M^*. The correlation matrix for $A(\Sigma^*)_{\cdot S}$ reveals an interesting structure as it equals (rounded up to 3 significant digits)

$$\begin{array}{cc}
 & \begin{array}{cccccccccc} (1,1) & (1,2) & (2,1) & (2,2) & (2,3) & (3,2) & (3,3) & (3,4) & (4,3) & (4,4) \end{array} \\
\begin{array}{c} (1,1) \\ (1,2) \\ (2,1) \\ (2,2) \\ (2,3) \\ (3,2) \\ (3,3) \\ (3,4) \\ (4,3) \\ (4,4) \end{array} &
\begin{pmatrix}
1 & -0.086 & -0.042 & -0.071 & -0.105 & -0.103 & -0.07 & -0.103 & -0.101 & -0.069 \\
-0.086 & 1 & 0.998 & -0.05 & -0.153 & -0.149 & -0.103 & -0.151 & -0.148 & -0.101 \\
-0.042 & 0.998 & 1 & -0.084 & -0.157 & -0.154 & -0.106 & -0.155 & -0.152 & -0.104 \\
-0.071 & -0.05 & -0.084 & 1 & -0.096 & -0.065 & -0.072 & -0.106 & -0.104 & -0.071 \\
-0.105 & -0.153 & -0.157 & -0.096 & 1 & 0.999 & -0.068 & -0.155 & -0.151 & -0.104 \\
-0.103 & -0.149 & -0.154 & -0.065 & 0.999 & 1 & -0.089 & -0.153 & -0.15 & -0.103 \\
-0.07 & -0.103 & -0.106 & -0.072 & -0.068 & -0.089 & 1 & -0.096 & -0.072 & -0.07 \\
-0.103 & -0.151 & -0.155 & -0.106 & -0.155 & -0.153 & -0.096 & 1 & 0.999 & -0.074 \\
-0.101 & -0.148 & -0.152 & -0.104 & -0.151 & -0.15 & -0.072 & 0.999 & 1 & -0.09 \\
-0.069 & -0.101 & -0.104 & -0.071 & -0.104 & -0.103 & -0.07 & -0.074 & -0.09 & 1
\end{pmatrix}
\end{array}.$$

Observe that we have high correlation $\operatorname{cor}(A(\Sigma^*)_{\cdot(i,i+1)}, (A(\Sigma^*)_{\cdot(i+1,i)}) \geqslant 0.998$ for $i \in [3]$. These are exactly the columns of $A(\Sigma^*)_{\cdot S}$ that belong to the same group (G_2, G_6, G_9). Hence, columns within the same group belonging to the support exhibit strong collinearity. Consequently, Γ^*_{SS} is very near to singularity, i.e. $\det(\Gamma^*_{SS}) \approx 7.6 * 10^{-24}$. Moreover, computing the upperbound (3.48), we obtain $|||\Gamma^*_{S^cS}(\Gamma^*_{SS})^{-1}|||_* < 652550.75$. Clearly, the group irrepresentability condition cannot be guaranteed by this bound.

By contrast, the lasso irrepresentability condition is met for this choice of M^*. We compute the irrepresentability constant to be $|||\Gamma^*_{S^cS}(\Gamma^*_{SS})^{-1}|||_\infty \approx 0.82 < 1$. For the lasso, the definition of the support S changes and specifically excludes the correlated columns that are present for the group lasso. Thus, the entries of $(\Gamma^*_{SS})^{-1}$ will be significantly smaller due to lack of collinearity and the irrepresentability condition is more likely to be met.

It is worth pointing out that the group lasso condition is only sufficient and not necessary. Consider, for example, the drift matrices M given in Figure 5.2. In simulations, the group Lyapunov lasso was able to correctly recover the support of M for $n \to \infty$ (cf. Figure 5.1).

There exists one scenario where meeting the lasso irrepresentability condition directly implies that the group irrepresentability conditon is met as the next lemma shows.

Lemma 3.59. Let G be a graph were every edge is a 2-cycle, then the lasso irrepresentability condition $\vert\vert\vert\Gamma^*_{S^cS}(\Gamma^*_{SS})^{-1}\vert\vert\vert_\infty \leqslant (1-\alpha)$ implies the group lasso irrepresentability condition $\vert\vert\vert\Gamma^*_{S^cS}(\Gamma^*_{SS})^{-1}\vert\vert\vert_* \leqslant \sqrt{d}(1-\alpha)$.

Proof. If every edge is a 2-cycle, the support defined for the lasso and the group lasso coincide. Moreover, by Lemma 3.25 we obtain

$$\left\vert\left\vert\left\vert\Gamma^*_{S^cS}(\Gamma^*_{SS})^{-1}\right\vert\right\vert\right\vert_* \leqslant \sqrt{d}\left\vert\left\vert\left\vert\Gamma^*_{S^cS}(\Gamma^*_{SS})^{-1}\right\vert\right\vert\right\vert_\infty \leqslant \sqrt{d}(1-\alpha). \tag{3.52}$$

□

Lastly, we want to provide an example for Remark 3.34, where we show that we can obtain a more relaxed group irrepresentability condition, that is,

$$\left\vert\left\vert\left\vert\Gamma^*_{S^cS}\left((\Gamma^*_{SS})^{-1}\right)_{-S_{\mathrm{diag}}}\right\vert\right\vert\right\vert_* \leqslant \sqrt{2}(1-\alpha) \quad \text{with } \alpha > 0$$

for the group Lyapunov lasso not penalizing the diagonal entries of M.

Example 3.60. Choose the drift matrix

$$M^* = \begin{pmatrix} -0.468074 & -0.839606 & -0.976979 & -0.856221 \\ -0.020352 & -0.59448 & 0 & 0 \\ 0.118385 & 0 & -0.941955 & 0 \\ -0.162255 & 0 & 0 & -0.479817 \end{pmatrix} \quad \text{and} \quad C = \begin{pmatrix} 2 & 2 & 0 & 0 \\ 2 & 2 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 2 \end{pmatrix}.$$

Then the relaxed group irrepresentability condition is satisfied since $\left\vert\left\vert\left\vert\Gamma^*_{S^cS}\left((\Gamma^*_{SS})^{-1}\right)_{-S_{\mathrm{diag}}}\right\vert\right\vert\right\vert_* \leqslant 1.354331 < \sqrt{2}$.

On the other hand, using the same method to compute an upper-bound (cf. (3.48)) for the regular group irrepresentability constant, we obtain $\vert\vert\vert\Gamma^*_{S^cS}(\Gamma^*_{SS})^{-1}\vert\vert\vert_* \leqslant 53584.53$. Extensive computations were required to derive the above example for M^*. Out of roughly 10^8 generated drift matrices with matching sparsity pattern to M^*, the given example was the only drift matrix where we could derive an upper bound smaller than $\sqrt{2}$.

4 Algebraic Results for GCLMs

In this section, we aim to address a few algebraic questions for graphical continuous Lyapunov models. First of all, we provide an overview of the theory developed in Dettling et al. (2022) to discuss the notion of identifiability in GCLMs. Secondly, we investigate the problem of covariance equivalence for undirected structures in GCLMs. Recall the definition for the continuous Lyapunov equation

$$M\Sigma + \Sigma M^{\mathsf{T}} + C = 0, \tag{4.1}$$

where $M, C \in \mathbb{R}^{p\times p}$ are parameters and $\Sigma \in \mathbb{R}^{p\times p}$ is the covariance matrix for random multivariate observations in $\mathbb{R}^p$. For our investigations, we need to slightly refine the definition of a GCLM (cf. Definition 2.13).

Definition 4.1 (*GCLM given* C)**.** For a mixed graph $G = ([p], E)$ and $C \in \mathrm{PD}(p)$ we define the *graphical continuous Lyapunov model given* C as the set of covariance matrices

$$\mathcal{M}_{G,C} := \{\Sigma \in \mathrm{PD}(p) : M\Sigma + \Sigma M^{\mathsf{T}} + C = 0 \text{ with } M \in \mathrm{Stab}(E)\},$$

where $\mathrm{Stab}(E)$ represents the set of stable matrices $M = (M_{ij}) \in \mathrm{Stab}(p)$ with $M_{ji} = 0$ whenever $i \to j \notin E$.

4.1 Identifiability in GCLMs

A key question when discussing graphical models is the notion of parameter identification. Specifically, given a mixed graph G, are the effects of interest identifiable, i.e., uniquely determined by the multivariate distribution of the observations? For structural equations, Drton (2016) provides a good overview of the corresponding theory.

Adapted for GCLMs identifiability, thus asks if the continuous Lyapunov equation with a given covariance matrix $\Sigma \in \mathcal{M}_{G,C}$ is solvable for more than one choice of a matrix $M \in \mathrm{Stab}(E)$. Hence we investigate the injectivity of the parametrization map

$$\begin{aligned} \phi_{G,C} : \mathrm{Stab}(E) &\to \mathrm{PD}(p) \\ M &\mapsto \Sigma(M, C), \end{aligned}$$

with $\Sigma(M, C)$ being the unique matrix Σ obtained by solving the continuous Lyapunov equation with the stable matrix M and positive-definite matrix C.

Without any constraints on $\mathrm{Stab}(E)$, the function $\phi_{G,C}$ is not injective. To see this, recall that the continuous Lyapunov equation is a symmetric matrix equation hence providing $p(p+1)/2$ constraints. On the other hand, M is a $p \times p$ matrix that may not be symmetric. Consequently, we need some constraints on the sparsity pattern of M, or else M is never uniquely determined by Σ.

4.1.1 Notions of identifiability

To introduce the different notions of identifiability it is helpful to define the concept of a fiber.

Definition 4.2 (*fiber*)**.** Let $\mathcal{M}_{G,C}$ be the GCLM for a given $C \in \mathrm{PD}(p)$, associated with a mixed graph $G = ([p], E)$. Then the *fiber* of a matrix $M_0 \in \mathrm{Stab}(E)$ is defined as

$$\mathcal{F}_{G,C}(M_0) := \{M \in \mathrm{Stab}(E) : \phi_{G,C}(M) = \phi_{G,C}(M_0)\}.$$

In words, the fiber consists of all drift matrices $M \in \mathrm{Stab}(E)$, whose Lyapunov equation is solved by a given covariance matrix Σ for a fixed $C \in \mathrm{PD}(p)$.

We distinguish three types of identifiability.

Definition 4.3. Let $\mathcal{M}_{G,C}$ be the GCLM for a given $C \in \mathrm{PD}(p)$, associated with a mixed graph $G = ([p], E)$. Then the model $\mathcal{M}_{G,C}$ is

(i) *globally identifiable* if $\mathcal{F}_{G,C}(M_0) = \{M_0\}$ for all $M_0 \in \mathrm{Stab}(E)$;

(ii) *generically identifiable* if $\mathcal{F}_{G,C}(M_0) = \{M_0\}$ for almost all $M_0 \in \mathrm{Stab}(E)$, i.e., the matrices with $\mathcal{F}_{G,C}(M_0) \neq \{M_0\}$ form a Lebesgue null set in $\mathrm{Stab}(E)$

(iii) *non-identifiable* if $|\mathcal{F}_{G,C}(M_0)| = \infty$ for all $M_0 \in \mathrm{Stab}(E)$

To illustrate these definitions consider the following examples.

Example 4.4. Consider G to be the directed 3-cycle displayed in Figure 4.1. It can be shown that solving the Lyapunov equation for M is equivalent to solving the linear equation

$$A_r(\Sigma)\mathrm{vec}(M) = -\mathrm{vech}(C),$$

where $\mathrm{vech}(C)$ represents the half-vectorization of $C \in \mathrm{PD}(p)$ (cf. (4.5)) and the $p(p+1)/2 \times p^2$-matrix $A_r(\Sigma)$ (cf. Definition 4.6) is equal to

$$\begin{array}{c} \\ (1,1) \\ (1,2) \\ (1,3) \\ (2,2) \\ (2,3) \\ (3,3) \end{array}
\begin{array}{c}
\begin{array}{ccccccccc} 1\to1 & 1\to2 & 1\to3 & 2\to1 & 2\to2 & 2\to3 & 3\to1 & 3\to2 & 3\to3 \end{array} \\
\begin{pmatrix}
2\Sigma_{11} & 0 & 0 & 2\Sigma_{12} & 0 & 0 & 2\Sigma_{13} & 0 & 0 \\
\Sigma_{12} & \Sigma_{11} & 0 & \Sigma_{22} & \Sigma_{12} & 0 & \Sigma_{23} & \Sigma_{13} & 0 \\
\Sigma_{13} & 0 & \Sigma_{11} & \Sigma_{23} & 0 & \Sigma_{12} & \Sigma_{33} & 0 & \Sigma_{13} \\
0 & 2\Sigma_{12} & 0 & 0 & 2\Sigma_{22} & 0 & 0 & 2\Sigma_{23} & 0 \\
0 & \Sigma_{13} & \Sigma_{12} & 0 & \Sigma_{23} & \Sigma_{22} & 0 & \Sigma_{33} & \Sigma_{23} \\
0 & 0 & 2\Sigma_{13} & 0 & 0 & 2\Sigma_{23} & 0 & 0 & 2\Sigma_{33}
\end{pmatrix}
\end{array}.$$

Columns indexed by $i \to j$ belong to the entry M_{ji} of the drift matrix M. Observe that for the directed 3-cycle $M_{12} = M_{23} = M_{31} = 0$, hence the unique solvability of the equation system in (4.4) is solely determined by the invertibility of the submatrix

$$A_r(\Sigma)_{\cdot,E} = \begin{array}{c} \\ (1,1) \\ (1,2) \\ (1,3) \\ (2,2) \\ (2,3) \\ (3,3) \end{array} \begin{array}{c} \begin{array}{cccccc} 1\to 1 & 1\to 2 & 2\to 2 & 2\to 3 & 3\to 1 & 3\to 3 \end{array} \\ \begin{pmatrix} 2\Sigma_{11} & 0 & 0 & 0 & 2\Sigma_{13} & 0 \\ \Sigma_{12} & \Sigma_{11} & \Sigma_{12} & 0 & \Sigma_{23} & 0 \\ \Sigma_{13} & 0 & 0 & \Sigma_{12} & \Sigma_{33} & \Sigma_{13} \\ 0 & 2\Sigma_{12} & 2\Sigma_{22} & 0 & 0 & 0 \\ 0 & \Sigma_{13} & \Sigma_{23} & \Sigma_{22} & 0 & \Sigma_{23} \\ 0 & 0 & 0 & 2\Sigma_{23} & 0 & 2\Sigma_{33} \end{pmatrix} \end{array}.$$

To show invertibility we examine the determinant of $A_r(\Sigma)_{\cdot,E}$, which factorizes as

$$\det(A_r(\Sigma)_{\cdot,E}) = 2^3 \cdot \det(\Sigma) \cdot (\Sigma_{11}\Sigma_{22}\Sigma_{33} - \Sigma_{12}\Sigma_{13}\Sigma_{23}). \tag{4.2}$$

Note that by assumption Σ is positive definite and thus $\det(\Sigma) > 0$. Furthermore, by the Cauchy-Schwarz inequality $\det(\Sigma_{ij,ij}) = \Sigma_{ii}\Sigma_{jj} - \Sigma_{ij}^2 > 0$ for $i \neq j$, directly implying $\Sigma_{11}^2\Sigma_{22}^2\Sigma_{33}^2 = \Sigma_{11}\Sigma_{22}\Sigma_{11}\Sigma_{33}\Sigma_{22}\Sigma_{33} > \Sigma_{12}^2\Sigma_{13}^2\Sigma_{23}^2$. Therefore, both factors in (4.2) are positive. In summary, for every $\Sigma \in \mathcal{M}_{G,C}$ there exists a unique matrix M such that $\Sigma = \phi_{G,C}(M)$. Hence, $\mathcal{F}_{G,C}(M_0) = \{M_0\}$ for all $M_0 \in \mathrm{Stab}(E)$, i.e., $\mathcal{M}_{G,C}$ is globally identifiable for the 3-directed cycle.

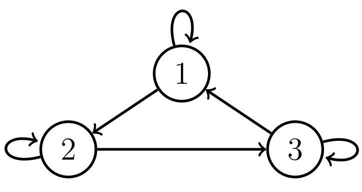

Figure 4.1: The directed 3-cycle.

Example 4.5. For the 2-cycle $G = (\{1,2\}, 1 \to 2, 2 \to 1$, $\phi_{G,C}$ maps the 4-dimensional parameter space $\mathrm{Stab}(E)$ to the 3-dimensional PD(2)-cone. The fiber is then specified by an undetermined linear system with 3 equations in 4 unknowns. We can therefore conclude that regardless of the choice of C, $\mathcal{M}_{G,C}$ is non-identifiable.

We can formulate equivalent rank conditions for identifiability. As shown in Example 2.5 the continuous Lyapunov equation (4.1) can be vectorized and thus transformed into the linear equation

$$((\Sigma \otimes I_p) + (I_p \otimes \Sigma)K^{(m,n)})\mathrm{vec}(M) = -\mathrm{vec}(C). \tag{4.3}$$

Observe that by the symmetry of the Lyapunov equation, the matrix on the LHS of the above equation (4.3) has redundant rows (cf. also Example 3.4). We define a row-restricted version of the preceding matrix as follows.

Definition 4.6. Let $\Sigma \in \mathbb{R}^{p\times p}$ be a symmetric matrix. The $p(p+1) \times p^2$ matrix $A_r(\Sigma)$ is obtained by selecting the rows of

$$A(\Sigma) := (\Sigma \otimes I_p) + (I_p \otimes \Sigma)K^{(m,n)} \tag{4.4}$$

indexed by pairs (k,l) with $k \leqslant l$.

We use the same indexing method introduced for $A(\Sigma)$ in Section 3.1. To better illustrate how each entry of the drift matrix M_{ij} relates to $A_r(\Sigma)$, we change the column index from (i,j) to $j \to i$. Moreover, we define the half-vectorization of the symmetric matrix C to be

$$\mathrm{vech}(C) := (C_{kl} : k \leqslant l). \tag{4.5}$$

Therefore, we can rewrite (4.3) as

$$A_r(\Sigma)\mathrm{vec}(M) = -\mathrm{vech}(C) \tag{4.6}$$

Note that for a given GCLM, the drift matrix M has non-zero entries only for pairs (j,i), where the associated graph G has an edge $i \to j$. Consequently, the solvability of (4.6) can be determined by a subset of the columns of the coefficient matrix $A_r(\Sigma)$.

Lemma 4.7. Let $G = ([p], E)$ be a mixed graph and $C \in \mathrm{PD}(p)$. We denote the submatrix of $A_r(\Sigma)$ restricted to the columns with indexes in E by $A_r(\Sigma)_{\cdot,E}$. Then the GCLM $\mathcal{M}_{G,C}$ is

(i) globally identifiable if and only if $A_r(\Sigma)_{\cdot,E}$ has full column rank $|E|$ for all $\Sigma \in \mathcal{M}_{G,C}$;

(ii) generically identifiable if and only if there exists a matrix $\Sigma \in \mathcal{M}_{G,C}$ such that $A_r(\Sigma)_{\cdot,E}$ has full column rank $|E|$;

(iii) non-identifiable if $\mathcal{M}_{G,C}$ is not generically identifiable.

Proof. Choose $M_0 \in \mathrm{Stab}(E)$ and let Σ_0 be the associated covariance matrix. Let $\mathrm{vec}(M)_E$ be the subvector of $\mathrm{vec}(M)$ obtained by selecting the entries (j,i) with $i \to j \in E$, then by the definition of a fiber

$$\mathcal{F}_{G,C}(M_0) = \{M \in \mathrm{Stab}(E) : A_r(\Sigma_0)_{\cdot,E}\mathrm{vec}(M)_E = -\mathrm{vech}(C)\}$$

For $\mathcal{M}_{G,C}$ to be globally identifiable $\mathcal{F}_{G,C}(M_0) = \{M_0\}$. Hence, this is satisfied if and only if $A_r(\Sigma_0)_{\cdot,E}$ has full column rank, which proves (i).

To prove the second claim, first observe that $A_r(\Sigma)_{\cdot,E}$ has full column rank if and only if the vector of all maximal minors of $A_r(\Sigma_0)_{\cdot,E}$ is non-zero.
From the continuous Lyapunov equation (4.3) it is clear that $\phi_{G,C}$ is a rational map. Moreover, the function mapping $M \in \mathrm{Stab}(E)$ to the maximal minors of $A_r(\phi_{G,C}(M)))_{\cdot,E}$ will then be rational as well. Recall that a rational map is non-zero outside a Lebesgue measure zero set if and only if we can find a single point where it is non-zero. Thus,

if we can find a $\Sigma \in \mathcal{M}_{G,C}$, where $A_r(\Sigma)_{\cdot,E}$ has full column rank, we can infer generic identifiability of $\mathcal{M}_{G,C}$.

Finally, to show the last claim (iii), note that $\mathcal{M}_{G,C}$ is not generically identifiable if the column rank of $A_r(\Sigma_0)_{\cdot,E}$ is strictly smaller than $|E|$ for all $\Sigma_0 = \phi_{G,C}(M_0) \in \mathcal{M}_{G,C}(M_0)$. Consequently, the fiber $\mathcal{F}_{G,C}(M_0) \subseteq \mathrm{Stab}(E)$ will then be an affine subspace with dimension larger than 1. Therefore, $|\mathcal{F}_{G,C}(M_0)| = \infty$ for all $M_0 \in \mathrm{Stab}(E)$ and $\mathcal{M}_{G,C}$ is non-identifiable. □

4.1.2 Identifiability for simple graphs

For an entire class of graphs, namely directed acyclic graphs (DAG), we can prove that they are globally identifiable. To motivate our choice of graphs, consider the following lemma.

Lemma 4.8. Let $\mathcal{M}_{G,C}$ be a globally identifiable GCLM with underlying mixed graph $G = ([p], E)$ and $C \in \mathrm{PD}(p)$. Consider $E' \subseteq E$ a subset of the edges specified by G. Then the model $\mathcal{M}_{H,C}$ specified by the subgraph $H = ([p], E')$ is globally identifiable.

Proof. Clearly, $\mathrm{Stab}(E') \subseteq \mathrm{Stab}(E)$. Hence, for any matrix $M_0 \in \mathrm{Stab}(E')$, it holds that the fiber

$$\mathcal{F}_{H,C}(M_0) \subseteq \mathcal{F}_{G,C}(M_0) = \{M_0\}$$

since $\mathcal{M}_{G,C}$ is assumed to be globally identifiable. □

An immediate consequence of Lemma 4.8 and Example 4.5 is that the graph of a globally identifiable GCLM cannot contain any 2-cycles. We call graphs that don't include any 2-cycles simple. One example of simple graphs are DAGs. Moreover, by Lemma 4.8 we only need to consider complete DAGs to prove global identifiability for all DAGs. Recall the definition for complete graphs.

Definition 4.9 (*complete*)**.** A simple graph $G = ([p], E)$ is *complete* if there exists an edge between every pair of distinct nodes.

Assuming the simple graph contains all self-loops $i \to i,\ i \in [p]$, it is complete if and only if $|E| = p(p+1)/2$. Since vertex relabelling will not affect identifiability, we may assume w.l.o.g. that the DAG adheres to a topological ordering. In summary, it suffices to consider the single complete DAG G^* with edge set $E^* = \{i \to j : i \geqslant j,\ i,j \in [p]\}$.

Theorem 4.10. *Let $G = ([p], E)$ be a DAG. Then the GCLM $\mathcal{M}_{G,C}$ is globally identifiable for every $C \in \mathrm{PD}(p)$.*

Proof. As explained above, we may restrict ourselves to the complete and topologically ordered DAG $G^* = (V, E^*)$. By Lemma 4.7 global identifiability is equivalent to $\det(A_r(\Sigma)_{\cdot,E^*}) \neq 0$ for all $\Sigma \in \mathcal{M}_{G^*,C}$. The key idea of the proof is that the coefficient matrix $A(\Sigma)$ exhibits a particular block structure, thus simplifying the calculation of the determinant.

Consider the partition of the edge set $E^* = \bigcup_{i=1}^{p} E_i^*$ with $E_i^* := \{j \to i : j \geqslant i,\ j \in [p]\}$. Furthermore, we analogously partition the row index set of $A_r(\Sigma)$ into the disjoint union of the sets $R_k := \{(k,l) : l \geqslant k\}$, $k \in [p]$. Recall that the entries of $A(\Sigma)$ have the following form

$$A_r(\Sigma)_{(k,l),i\to j} = \begin{cases} 0, & \text{if } \ j \neq k, l; \\ \Sigma_{li}, & \text{if } \ j = k,\ k \neq l; \\ \Sigma_{ki}, & \text{if } \ j = l,\ l \neq k; \\ 2\Sigma_{ji}, & \text{if } \ j = k = l. \end{cases} \tag{4.7}$$

From (4.7) we can infer that the submatrix

$$A_r(\Sigma)_{R_k, E_i^*} = 0 \quad \text{if } k > i.$$

Thus, $A(\Sigma)$ can be rearranged in a block upper-triangular form, and

$$\det(A_r(\Sigma)_{\cdot, E^*}) = \prod_{i=1}^{p} \det(A_r(\Sigma)_{R_i, E_i^*}).$$

Moreover, observe that by (4.7) $A_r(\Sigma)_{R_i, E_i^*}$ is equal to the principal submatrix $P(\Sigma)_{\geqslant i} := \Sigma_{\{i,\dots,p\},\{i,\dots,p\}}$ with the first row (indexed by i) being multiplied by 2. Applying Sylvester's criterion for $\Sigma \in \mathrm{PD}(p)$ implies that all principal minors are positive and therefore

$$|\det(A_r(\Sigma)_{\cdot, E^*})| = 2^p \prod_{i=1}^{p} \det(P(\Sigma)_{\geqslant i}) > 0 \qquad \text{for all } \Sigma \in \mathrm{PD}(p).$$

Specifically, for all $\Sigma \in \mathcal{M}_{G^*, C} \subseteq \mathrm{PD}(p)$ we have $\det(A_r(\Sigma)_{\cdot, E^*}) \neq 0$. □

In fact, we can show that global identifiability holds for simple graphs in general, i.e., we can also include simple cyclic graphs. To prove this, first, recall the following fact from Barnett and Storey (1967) that holds for the drift matrices M of the continuous Lyapunov equation.

Lemma 4.11. Let $\Sigma, C \in \mathrm{PD}(p)$ be given. The continuous Lyapunov equation (4.1) is solved by a matrix $M \in \mathbb{R}^{p\times p}$ if and only if there exists a skew-symmetric matrix $K \in \mathbb{R}^{p\times p}$ ($K^\mathsf{T} = -K$) such that

$$M = \left(K - \frac{1}{2}C\right)\Sigma^{-1}.$$

Proof. Let M be a matrix that solves the continuous Lyapunov equation (4.1) for a given Σ and C. Since Σ and C are symmetric we can rewrite (4.1) as

$$(M\Sigma)^\mathsf{T} + \frac{1}{2}C^\mathsf{T} = -M\Sigma - \frac{1}{2}C.$$

The above equation yields that $K := M\Sigma + \frac{1}{2}C$ is a skew-symmetric matrix. Thus,

$$M = \left(K - \frac{1}{2}C\right)\Sigma^{-1},$$

which proves the lemma. □

Theorem 4.12. *Let $G = ([p], E)$ be a simple directed graph. Then the GCLM $\mathcal{M}_{G,C}$ is globally identifiable for all $C \in \mathrm{PD}(p)$.*

Proof. Consider $M_1, M_2 \in \mathrm{Stab}(E)$ that both solve the Lyapunov equation (4.1) for the same $\Sigma \in \mathcal{M}_{G,C}$. Applying Lemma 4.11 yields that there exists two skew-symmetric matrices K_1 and K_2 with $M_1 = (K_1 - \frac{1}{2}C)\Sigma^{-1}$ and $M_2 = (K_2 - \frac{1}{2}C)\Sigma^{-1}$. We define the difference

$$M := M_1 - M_2 = \left(K_1 - \frac{1}{2}C\right)\Sigma^{-1} - \left(K_2 - \frac{1}{2}C\right)\Sigma^{-1} = (K_1 - K_2)\Sigma^{-1}.$$

Note that the difference $K := K_1 - K_2$ is again skew-symmetric. Hence, M is again the product of a skew-symmetric matrix M and the positive-definite matrix Σ^{-1}.

The square M^2 is equal to

$$M^2 = K\Sigma^{-1}K\Sigma^{-1}.$$

Since Σ is positive-definite, the square root $\Sigma^{\frac{1}{2}}$ and its inverse $\Sigma^{-\frac{1}{2}}$ both exist. Thus, M^2 is similar to

$$\Sigma^{-\frac{1}{2}}M^2\Sigma^{\frac{1}{2}} = \Sigma^{-\frac{1}{2}}K\Sigma^{-1}K\Sigma^{-\frac{1}{2}}.$$

Moreover, recalling that K is skew-symmetric yields

$$\Sigma^{-\frac{1}{2}}K\Sigma^{-1}K\Sigma^{-\frac{1}{2}} = -(\Sigma^{-\frac{1}{2}}K)\Sigma^{-1}(\Sigma^{-\frac{1}{2}}K)^{\mathsf{T}}.$$

Therefore, M^2 is similar to a symmetric and negative semi-definite matrix, hence shares the same eigenvalues as the preceding matrix. Consequently, M^2 has eigenvalues that are non-positive and $\mathrm{tr}(M^2) \leqslant 0$ as it is the sum over all eigenvalues counted with multiplicity.

G is assumed to be simple, thus $M_{ij} \neq 0$ directly implies $M_{ji} = 0$ for all pairs of indices $i \neq j$. In particular, the diagonal of M^2 is then only specified by the squared diagonal elements of M, i.e. $(M^2)_{ii} = (M_{ii})^2$. Furthermore,

$$0 \leqslant \sum_{i=1}^{p}(M_{ii})^2 = \mathrm{tr}(M^2) \leqslant 0.$$

Let $\lambda_1, \ldots, \lambda_p \in \mathbb{C}$ be the eigenvalues of M. Then the eigenvalues of M^2 are $\lambda_1^2, \ldots, \lambda_p^2$. Recall from above that all the eigenvalues of M^2 are non-positive, i.e. $\lambda_1^2 \leqslant 0, \ldots, \lambda_p^2 \leqslant 0$. Combining this with the fact that

$$0 = \mathrm{tr}(M^2) = \sum_{i=1}^{p}\lambda_i^2 \leqslant 0,$$

leads us to the conclusion that $\lambda_i^2 = 0$ for all $i \in [p]$ or equivalently $\lambda_i = 0$ for all $i \in [p]$.

Note that $M = K\Sigma^{-1}$ is similar to $\tilde{M} = \Sigma^{-\frac{1}{2}}K\Sigma^{-1}\Sigma^{\frac{1}{2}}$, which is a skew-symmetric matrix as

$$\tilde{M}^{\mathsf{T}} = (\Sigma^{-\frac{1}{2}}K\Sigma^{-\frac{1}{2}})^{\mathsf{T}} = \Sigma^{-\frac{1}{2}}K^{\mathsf{T}}\Sigma^{-\frac{1}{2}} = -\Sigma^{-\frac{1}{2}}K\Sigma^{-\frac{1}{2}} = -\tilde{M}.$$

All skew-symmetric matrices are diagonalizable over $\mathbb{C}$, implying that M is similar to the zero matrix. However, then $M = 0$ and thus $M_1 = M_2$, which proves that the Lyapunov equation has a unique solution if the underlying graph is simple. □

Remark 4.13. If we restrict $C \in \mathrm{PD}(p)$ to be diagonal, G being simple can be shown to be a sufficient and necessary condition for global identifiability (Dettling et al. (2022), Theorem 31 (ii)).

4.1.3 Identifiability for non-simple graphs

For graphs that are non-simple, making general statements wrt. identifiability is much more difficult. Although one can always check the rank conditions from Lemma 4.7 to determine identifiability, finding a general class of graphs that will guarantee identifiability proves to be much harder. However, we can at least provide a necessary combinatorial condition that has to be satisfied if $\mathcal{M}_{G,C}$ is generically identifiable. First, note the following fact for non-identifiability.

Lemma 4.14. Let $G = ([p], E)$ be a directed graph and $C \in \mathrm{PD}(p)$. If $|E| > \dim(\mathcal{M}_{G,C})$, then the GCLM $\mathcal{M}_{G,C}$ will be non-identifiable. In other words, all graphs with more than $p(p+1)/2$ edges give non-identifiable models.

Proof. The set of sparse stable matrices $\mathrm{Stab}(E)$ is semi-algebraic by the Hurwitz criterion (Horn and Johnson (1991), Theorem 2.3.3). Moreover, $\dim(\mathrm{Stab}(E)) = |E| > \dim(\mathcal{M}_{G,C})$ implies that the rational map $\phi_{G,C}$ is generically infinite-to-one (Barber et al. (2022), Lemma 2.5). Hence, all fibers are infinite, and $\mathcal{M}_{G,C}$ is non-identifiable. □

We can refine the above criterion by observing that the existence of no treks between two nodes implies that the corresponding covariance entry is 0. We define a trek as follows.

Definition 4.15 (*trek*)**.** A *trek* τ is a sequence of edges of the form

$$i \leftarrow \cdots \leftarrow i_1 \leftarrow k \leftrightarrow l \rightarrow j_1 \rightarrow \cdots \rightarrow j,$$

where k and l are connected by a bidirected edge. To the left and right of k and l we have directed paths $i \leftarrow \cdots \leftarrow i_1 \leftarrow k$ and $l \rightarrow j_1 \rightarrow \cdots \rightarrow j$ of length $n(\tau)$ and $m(\tau)$, respectively. In particular, we allow for $n(\tau) = 0$ and $n(\tau) = 0$, thus we consider directed paths as well as single nodes to be treks.

Recall that by Theorem 2.12 for a given M and C the covariance matrix $\Sigma \in \mathcal{M}_{G,C}$ can be expressed as the following limit

$$\Sigma = \lim_{s\to\infty} \Sigma(s) = \lim_{s\to\infty} \int_0^s e^{uM} C e^{uM^T} du.$$

Furthermore, Varando and Hansen (2020) also proved that $\Sigma(s)$ can be expressed in terms of the treks as follows.

Theorem 4.16. *Let $M \in \mathrm{Stab}(E)$ and $C \in \mathrm{PD}(p)$. We define*

$$\kappa(s,\tau) := \frac{s^{n(\tau)+m(\tau)+1}}{(n(\tau)+m(\tau)+1)n(\tau)!m(\tau)!}$$

for any trek τ and $s \in \mathbb{R}$. Moreover, let

$$\omega(M,C,\tau) := C_{kl} \prod_{g\to h\in\tau} M_{hg}.$$

Then

$$\Sigma(s)_{ij} = \sum_{\tau\in\mathcal{T}(i,j)} \kappa(s,\tau)\omega(M,C,\tau),$$

where $\mathcal{T}(i,j)$ represents the set of all treks from i to j.

Proof. By the series expansion of the matrix exponential e^{uM}, we have

$$\begin{aligned}\Sigma(s)_{ij} &= \int_0^s (e^{uM} C e^{uM^T})_{ij} du = \int_0^s \sum_{k,l=1}^p (e^{uM})_{ik} C_{kl} (e^{uM})_{lj} du \\ &= \int_0^s \sum_{n=0}^\infty \sum_{m=0}^\infty \sum_{k,l=1}^p \frac{t^n t^m}{n!m!} (M^n)_{ik} C_{kl} (M^n)_{lj} du.\end{aligned}$$

Using Fubini's theorem, we can exchange integration and summation. Therefore,

$$\begin{aligned}&= \sum_{n=0}^\infty \sum_{m=0}^\infty \sum_{k,l=1}^p \frac{s^{n+m+1}}{(n+m+1)n!m!} (M^n)_{ik} C_{kl} (M^n)_{lj} \\ &= \sum_{\tau\in\mathcal{T}(i,j)} \kappa(s,\tau)\omega(M,C,\tau)\end{aligned}$$

$\square$

From Theorem 2.12 and Theorem 4.16 it is straightforward to deduce the next corollary.

Corollary 4.17. Let $G = ([p], E)$ be a mixed graph and $C \in \mathrm{PD}(p)$. Moreover, assume there exists no treks between the nodes i and j in G. Then $\Sigma_{ij} = 0$ for all $\Sigma \in \mathcal{M}_{G,C}$.

Observe that if there exists no treks between certain nodes, the dimension of $\mathcal{M}_{G,C}$ is upper bounded by

$$\dim(\mathcal{M}_{G,C}) \leqslant \frac{p(p+1)}{2} - |\{\{i,j\} : i,j \in [p] \text{ with no trek between them}\}|.$$

In combination with Lemma 4.14 we can now establish the following necessary condition for generic identifiability.

Corollary 4.18. Let $G = ([p], E)$ be a mixed graph and $C \in \mathrm{PD}(p)$. Assume $\mathcal{M}_{G,C}$ is generically identifiable. Then

$$|E| \leqslant \frac{p(p+1)}{2} - |\{\{i,j\} : i,j \in [p] \text{ with no trek between them}\}|.$$

4.2 Covariance equivalence for undirected GCLMs

The undirected structure of a GCLM is defined as the skeleton G^{skel} of the underlying mixed graph G for a GCLM (cf. Definition 2.14). Similar as for the directed case, we adjust our definition for an undirected GCLM (cf. Definition 2.16).

Definition 4.19 (*undirected GCLM given* C)**.** For an undirected graph $G^{skel} = ([p], E^{skel})$ and $C \in \mathrm{PD}(p)$ we define the *undirected graphical continuous Lyapunov model given* C as the set of covariance matrices

$$\mathcal{M}_{G^{skel},C} := \{\Sigma \in \mathrm{PD}(p) : M\Sigma + \Sigma M^{\mathsf{T}} + C = 0 \text{ with } M \in \mathrm{Stab}(E_{skel})\},$$

where $\mathrm{Stab}(E^{skel})$ represents the set of stable matrices $M = (M_{ij}) \in \mathrm{Stab}(p)$ with $M_{ij} = M_{ji} = 0$ whenever $i - j \notin E^{skel}$.

We are now interested in the question of whether different graphs may induce the same model. Assume we have two different undirected graphs G_1^{skel} and G_2^{skel}, which represent different scientific hypotheses. If $\mathcal{M}_{G_1^{skel},C} = \mathcal{M}_{G_2^{skel},C}$, then we cannot differentiate between the associated hypotheses based on the data alone. This leads us to the following definition.

Definition 4.20 (*covariance equivalence*)**.** Let $G_1^{skel} = ([p], E_1^{skel})$ and $G_2^{skel} = ([p], E_2^{skel})$ be undirected graphs and $C \in \mathrm{PD}(p)$. If the associated GCLMs satisfy the following equality

$$\mathcal{M}_{G_1^{skel},C} = \mathcal{M}_{G_2^{skel},C},$$

we call G_1 and G_2 *covariance equivalent.*

Equivalently, we could look at the images of the rational maps $\phi_{G_1^{skel},C}$ and $\phi_{G_2^{skel},C}$ defined as

$$\begin{aligned}\phi_{G_i^{skel},C} : \mathrm{Stab}(E_i^{skel}) &\to \mathrm{PD}(p)\\ M &\mapsto \Sigma(M,C), \quad \text{for } i = 1,2.\end{aligned}$$

Then G_1 and G_2 are covariance equivalent if

$$\operatorname{img}(\phi_{G_1^{skel},C}) = \operatorname{img}(\phi_{G_2^{skel},C}).$$

We will restrict our investigations to connected graphs.

Definition 4.21 (*connected*)**.** An undirected graph $G = (V, E)$ is said to be *connected* if every pair of nodes in G is connected. Meaning there exists a path between every pair of nodes. An undirected graph that is not connected is called *disconnected*.

Note for disconnected graphs we can infer the presence of zero entries in the covariance matrix Σ by Corollary 4.17. In particular, for two disconnected nodes i and j it is not possible to find a trek connecting them in a corresponding directed graph, and thus $\Sigma_{ij} = \Sigma_{ji} = 0$. Hence, any potential covariance equivalent graph needs to have a matching sparsity pattern for Σ. For connected graphs, we cannot make a similar statement as the covariance matrix Σ does not contain zero entries in general.

To investigate covariance equivalence for connected graphs, we perform a simulation study in `Mathematica`. The simulation setup is as follows.

1. Generate a random diagonal matrix $C \in \mathrm{PD}(p)$, where the diagonal entries C_{ii} for $i \in [p]$ are independently drawn from a uniform distribution on $[100]$. Moreover, generate all possible symbolic drift matrices M^{sym} each belonging to a connected graph G.

2. We populate the off-diagonal non-zero entries of the symbolic drift matrices with values independently drawn from a uniform distribution on $\{-100, 99, \dots, -1\} \cup \{1, 2, \dots, 100\}$, i.e., for $i \neq j$

$$M_{ij}^{sample} = \begin{cases} u_{ij} \sim \mathrm{Uniform}([100]) & \text{if } M_{ij}^{sym} \text{ is not zero,} \\ 0 & \text{if } M_{ij}^{sym} = 0. \end{cases}$$

   The diagonal entries of M are chosen to be

$$M_{ii}^{sample} = -\sum_{i \neq j} |M_{ij}^{sample}| - |u_{ii}| \quad \text{for } i \in [p],$$

   where $u_{ii} \sim \mathrm{Uniform}([100])$.

3. We solve the following continuous Lyapunov equation

$$M^{sample}\Sigma + \Sigma(M^{sample})^{\mathsf{T}} + C = 0$$

   for Σ.

4. Next, we solve the linear matrix equation with given Σ and C

$$M^{sym}\Sigma + \Sigma(M^{sym})^{\mathsf{T}} + C = 0 \tag{4.8}$$

   for all possible M^{sym} generated in the first step.

5. Steps 2. - 4. are repeated n = 4 times in total for each graph. Lastly, we tally the results for each graph G, i.e., we note whether a symbolic drift matrix M^{symb} belonging to a graph G_1^{skel} was able to solve the continuous Lyapunov equation with a given $\Sigma \in \mathcal{M}_{G_2^{skel},C}$ associated to a graph G_2.

The first observation we make is that for low-dimensional graphs, all connected graphs are potentially covariance equivalent. We call two graphs G_1^{skel} and G_2^{skel} potentially covariance equivalent if there exists a covariance matrix $\Sigma \in \mathrm{PD}(p)$ such that $\Sigma \in \mathcal{M}_{G_1^{skel},C}$ and $\Sigma \in \mathcal{M}_{G_2^{skel},C}$. We stress the fact potential covariance equivalence is merely an indication for covariance equivalence as for the latter, it is required that $\mathcal{M}_{G_1^{skel},C} = \mathcal{M}_{G_2^{skel},C}$.

nodes p	# of connected graphs	# of potential equivalent graphs by edge count ({# of edges, # of graphs})
2	1	$\{1,1\}$
3	4	$\{2,4\}$, $\{3,4\}$
4	38	$\{3,38\}$, $\{4,38\}$, $\{5,38\}$, $\{6,38\}$

Table 4.1: Simulation results for undirected graphs with $p = 2,3,4$ nodes. The last column displays the # of graphs that can solve (4.8) for a Σ provided by a graph with a certain # of edges.

From Table 4.1 we can deduce that for $p = 2,3,4$, all connected graphs are potentially covariance equivalent since the number of potential equivalent graphs for each edge count equals the total number of connected graphs. There is strong evidence that for $p = 2,3,4$, all connected graphs are, in fact, covariance equivalent as the results above hold true even if we increase our number of repetitions n. The following lemma and subsequent corollary prove that they are indeed covariance equivalent.

Lemma 4.22. Let $G_1^{skel} = ([p], E_1^{skel})$ and $G_2^{skel} = ([p], E_2^{skel})$ be two undirected graphs and $C \in \mathrm{PD}(p)$ a diagonal matrix. If

$$\min\{|E_1^{skel}|, |E_2^{skel}|\} \geqslant \frac{p(p-1)}{4},$$

G_1 and G_2 are covariance equivalent.

Proof. Note that the dimension of the undirected GCLM $\mathcal{M}_{G^{skel},C}$ is given by

$$\dim(\mathcal{M}_{G^{skel},C}) = \min\left\{p + 2|E^{skel}|, \frac{p(p+1)}{2}\right\}.$$

Assuming $\min\{|E_1^{skel}|, |E_2^{skel}|\} \geqslant p(p-1)/4$ then yields

$$\dim(\mathcal{M}_{G_1^{skel},C}) = \dim(\mathcal{M}_{G_2^{skel},C}) = \frac{p(p+1)}{2}.$$

Recall that by definition $\mathcal{M}_{G_1^{skel},C}, \mathcal{M}_{G_2^{skel},C} \subseteq \mathrm{PD}(p)$. Therefore, both undirected GCLMs cover the entire cone of positive definite matrices, i.e.

$$\mathcal{M}_{G_1^{skel},C} = \mathrm{PD}(p) = \mathcal{M}_{G_2^{skel},C},$$

which proves the lemma. □

It suffices to look at one of the minimally connected undirected graph $G^{skel} = ([p], E^{skel})$, i.e., G^{skel} is connected, and there is no edge that can be removed while still leaving the graph connected. In this case $|E^{skel}| = p - 1$. Observe that

$$|E^{skel}| = p - 1 \geqslant \frac{p(p-1)}{4} \quad \text{for } p = 2, 3, 4.$$

Hence, we can directly infer from Lemma 4.22 this next corollary.

Corollary 4.23. Let $p \leqslant 4$. Consider two undirected connected graphs $G_1^{skel} = ([p], E_1^{skel})$ and $G_2^{skel} = ([p], E_2^{skel})$ and $C \in \mathrm{PD}(p)$ a diagonal matrix. Then G_1 and G_2 are covariance equivalent.

We now consider connected graphs with $p = 5$ nodes. Moreover, we introduce the following definition.

Definition 4.24 (*tree*)**.** A *tree* is an undirected graph $G = ([p], E)$ in which any two vertices are connected by exactly one path. Equivalently, G is a tree if and only if G is connected and G has $p - 1$ edges.

nodes p	# of connected graphs (# of trees)	# of potential equivalent graphs by edge count ({# of edges, # of graphs})
5	728 (125)	$\{4, 604\}$, $\{i, 603\}$ for $5 \leqslant i \leqslant 10$

Table 4.2: Simulation results for undirected graphs with $p = 5$ nodes. The last column displays the # of graphs that can solve (4.8) for a Σ provided by a graph with a certain # of edges.

According to Table 4.2, connected graphs with more $i \leqslant 5$ edges have exactly 603 potentially covariance equivalent graphs. Note that for graphs with $p = 5$ nodes, we have exactly 125 trees. The preceding $603 = 728 - 125$ graphs are all non-trees. This result is consistent with Lemma 4.22, which states that all 603 non-trees are covariance equivalent since the number of edges is equal to or larger than $p(p-1)/4$. Moreover, we observe that the trees are not covariance equivalent to the non-trees. This can also be explained theoretically since $\dim(M_{G^{skel},C}) = 13 < 15 = \dim(\mathrm{PD}(5))$ if G^{skel} is a tree and the corresponding GCLM for non-trees are equal to the entire PD-cone.
For the second key insight from Table 4.2, note that all trees have 604 potentially covariance equivalent graphs. A closer inspection reveals that of the 604 graphs, 1 is the original tree itself, whereas the remaining 603 graphs are non-trees. Consequently, for $p = 5$ all trees are not covariance equivalent. This fact holds for higher-dimensional graphs as well.

Theorem 4.25. *Let $p = 5, 6$. All trees corresponding to an undirected GCLM are not covariance equivalent.*

Proof. We can prove this fact computationally by running the simulation setup described earlier, which will net a counterexample to covariance equivalence for each possible tree combination. □

The simulation results for trees with 6 nodes are displayed in Table 4.3.
We conjecture that the statement from Theorem 4.25 will hold for all $p \geqslant 5$. However, we are not able to provide a proof for this fact yet. In particular, we note that to run the simulation setup for higher-dimensional trees, considerable computational effort is required. The number of trees for a graph with p labeled nodes equals p^{p-2} (Gross and Yellen (2004)). Our simulation setup then has to solve $4p^{2(p-2)}$ Lyapunov equations in total to check the validity of the above theorem. For example, for $p = 8$, we would need to check 262144 possible trees against each other.

nodes p	# of trees	# of potential equivalent trees
6	1296	1

Table 4.3: Simulation results for trees with $p = 6$ nodes. The last column shows the number of potential equivalent graphs for each tree

5 Numerical Studies

5.1 Structure recovery methods and tuning

In the following numerical studies, we will use the direct Lyapunov lasso and the group Lyapunov lasso. For the lasso, we use the implementation provided by the R package `glmnet` and the group lasso is computed using the R package `gglasso` (Yang and Zou (2015)). While there exist several other packages for these methods, both choices provide some of the fastest computation times, which are necessary for the high-dimensional setting we operate in.
For both methods, we use a decreasing regularization parameter sequence of length 100 with slight differences:

$$\begin{aligned} \text{lasso } &: \lambda_{100} = \lambda_{max,l} > \cdots > \lambda_1 = \frac{\lambda_{max,l}}{10^4} > 0 \\ \text{group lasso } &: \lambda_{100} = \lambda_{max,gl} > \cdots > \lambda_1 = \frac{\lambda_{max,gl}}{10^3} > 0. \end{aligned}$$

$\lambda_{max,l}$ is chosen to be the smallest penalization parameter such that $\hat{M}$ is diagonal, whereas $\lambda_{max,gl}$ is set to the smallest penalization parameter for which all entries of $\hat{M}$ are zero. Moreover, we apply individual weights to the regularization parameters. For the lasso, we set all weights to be 1 for the off-diagonal elements of M and 0 for the diagonal entries of M. The weights for the group lasso are chosen to be the squared group length, i.e., $\sqrt{2}$ for the off-diagonal elements of M, and 0 or 1 for the diagonal elements of M. Furthermore, we set the intercept in both algorithms to 0.

5.2 Performance comparison for fixed GCLMs

We start by comparing the performance of the lasso and the group lasso for correct undirected graph recovery of fixed graphs. For the group lasso, we consider two variants. The first one penalizes the diagonal elements of the signal M (cf. Definition 3.5), whereas the second does not (cf. Definition 3.6).

5.2.1 Simulation setup

We will consider the following graphs with 10 nodes ($p = 10$) each as displayed in Figure 5.1.

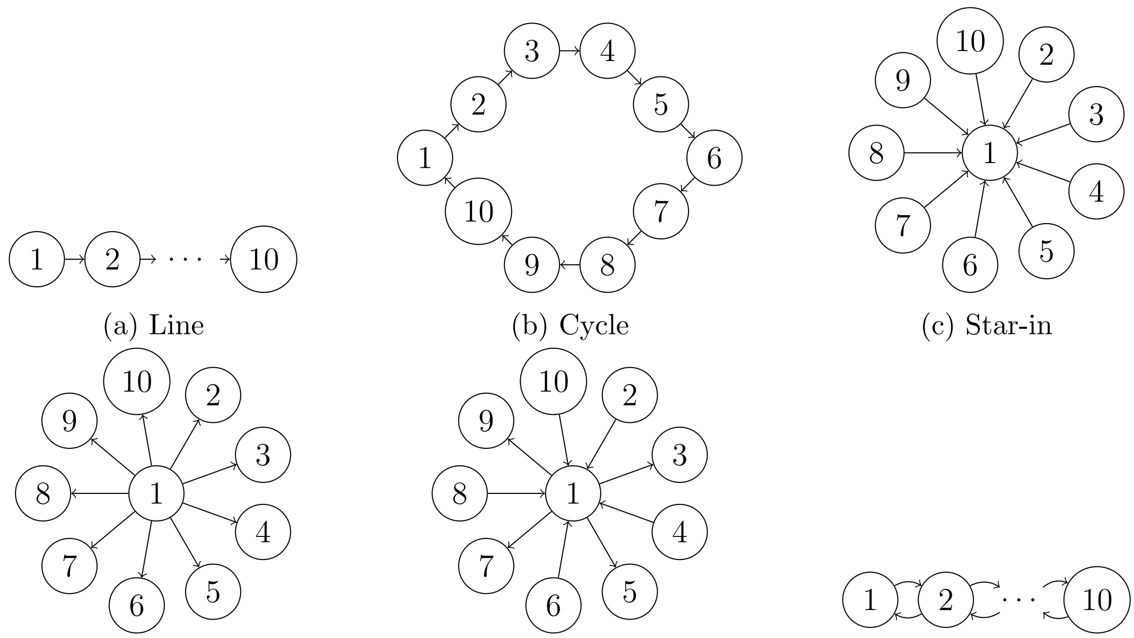

(a) Line (b) Cycle (c) Star-in

(d) Star-out (e) Star-mixed (f) 2-cycles

Figure 5.1: Overview of the considered graphs

For each of these graphs we fix the diagonal elements of the corresponding matrix to be $(-2, -3, \dots, -11)$ and the nonzero off-diagonal entries to be 0.1. An overview of the matrices M is given in Figure 5.2.

We set $C = 2I_{10}$. Choosing C to be diagonal results in the undirected graph being only determined by M as C introduces no additional edges. Based on the pair (M, C) we solve the Lyapunov equation for Σ, which is then used to generate $n = 10^2, 10^3, 10^4, 10^5, \infty$ samples from a multivariate Gaussian $\mathcal{N}(0, \Sigma)$. For each sample size, we compute a sample covariance matrix $\hat{\Sigma}$. Note for $n = \infty$ we just set $\hat{\Sigma} = \Sigma$. We repeat the preceding process 100 times per sample size.

5.2.2 Results

Since we initialise both the lasso and the group lasso with a lambda sequence of length 100, we obtain solution paths of drift matrices $\hat{M}$. Specifically, for each given regularization parameter λ, we compute a corresponding $\hat{M}$. We consider the following performance measures computed by considering the undirected graph recovery as a classification problem over the $\frac{p(p-1)}{2}$ possible edges. The ground truth is represented by the edges specified by the non-zero entries of M.

- *True positives* (TP): The number of undirected edges that are specified by $\hat{M}$ and M, i.e.,
  $TP = |\{(i,j) : \hat{M}_{ij} \neq 0 \text{ or } \hat{M}_{ji} \neq 0,\ i < j\} \cap \{(i,j) : M_{ij} \neq 0 \text{ or } M_{ji} \neq 0,\ i < j\}|$.
- *False positives* (FP): The number of undirected edges that are specified by $\hat{M}$ but

$$\begin{pmatrix} -2 & 0 & 0 & \cdots & 0 \\ 0.1 & -3 & 0 & \cdots & 0 \\ 0 & 0.1 & -4 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & 0.1 & -11 \end{pmatrix}$$

(a) Line

$$\begin{pmatrix} -2 & 0 & \cdots & 0 & 0.1 \\ 0.1 & -3 & 0 & \cdots & 0 \\ 0 & 0.1 & -4 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & 0.1 & -11 \end{pmatrix}$$

(b) Cycle

$$\begin{pmatrix} -2 & 0.1 & 0.1 & \cdots & 0.1 \\ 0 & -3 & 0 & \cdots & 0 \\ 0 & 0 & -4 & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & 0 \\ 0 & 0 & \cdots & 0 & -11 \end{pmatrix}$$

(c) Star-in

$$\begin{pmatrix} -2 & 0 & 0 & \cdots & 0 \\ 0.1 & -3 & 0 & \cdots & 0 \\ 0.1 & 0 & -4 & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & 0 \\ 0.1 & 0 & \cdots & 0 & -11 \end{pmatrix}$$

(d) Star-out

$$\begin{pmatrix} -2 & — & v & — & 0 \\ 0.1 & -3 & 0 & \cdots & 0 \\ | & 0.1 & -4 & \ddots & \vdots \\ v & \ddots & \ddots & \ddots & 0 \\ | & \cdots & 0 & 0.1 & -11 \end{pmatrix}$$

(e) Star-mixed

$$\begin{pmatrix} -2 & 0.1 & 0 & \cdots & 0 \\ 0.1 & -3 & 0.1 & \ddots & \vdots \\ 0 & 0.1 & -4 & \ddots & 0 \\ \vdots & \ddots & \ddots & \ddots & 0.1 \\ 0 & \cdots & 0 & 0.1 & -11 \end{pmatrix}$$

(f) 2-cycles

Figure 5.2: Overview of corresponding matrices M. The vector v in Star-mixed is defined as $v := (0, 0.1, 0, 0.1, 0, 0.1, 0, 0.1)$.

not M, i.e.,
$FP = |\{(i,j) : \hat{M}_{ij} \neq 0 \text{ or } \hat{M}_{ji} \neq 0,\ i < j\} \cap \{(i,j) : M_{ij} = 0 \text{ or } M_{ji} = 0,\ i < j\}|$.

- *True negatives* (TN): The number of absent undirected edges that are specified by $\hat{M}$ and M, i.e.,
  $TN = |\{(i,j) : \hat{M}_{ij} = 0 \text{ or } \hat{M}_{ji} = 0,\ i < j\} \cap \{(i,j) : M_{ij} = 0 \text{ or } M_{ji} = 0,\ i < j\}|$.
- *False negatives* (FN): The number of absent undirected edges that are specified by $\hat{M}$ but not M, i.e.,
  $FN = |\{(i,j) : \hat{M}_{ij} = 0 \text{ or } \hat{M}_{ji} = 0,\ i < j\} \cap \{(i,j) : M_{ij} \neq 0 \text{ or } M_{ji} \neq 0,\ i < j\}|$.
- *Path-wise maximum accuracy* ($maxacc$): The maximal accuracy achieved for undirected edge recovery along a solution path
  $maxacc = \max_\lambda \frac{TP(\lambda)+TN(\lambda)}{p(p-1)/2}$, where $TP(\lambda)$ and $TN(\lambda)$ refer to the TP and TN for an estimate $\hat{M}$ specified by their respective λ (one entry of the solution path).

Note we omitted the diagonal elements of $\hat{M}$ and M in our definition of the various metrics since they do not represent undirected edges in our interpretation of the drift matrix and its corresponding undirected graph.

The results can be seen in Figure 5.3. As one would expect, the accuracy for all algorithms improves with increasing sample size. For finite sample sizes both variants of the group lasso are slightly better than the lasso across all considered graphs. Moreover, when the sample size is infinite, all structure recovery methods achieve perfect support selection for all graphs but the line and the cycle. However, only the group lasso with no penalties on the diagonal of M was able to correctly estimate the support for all graphs.

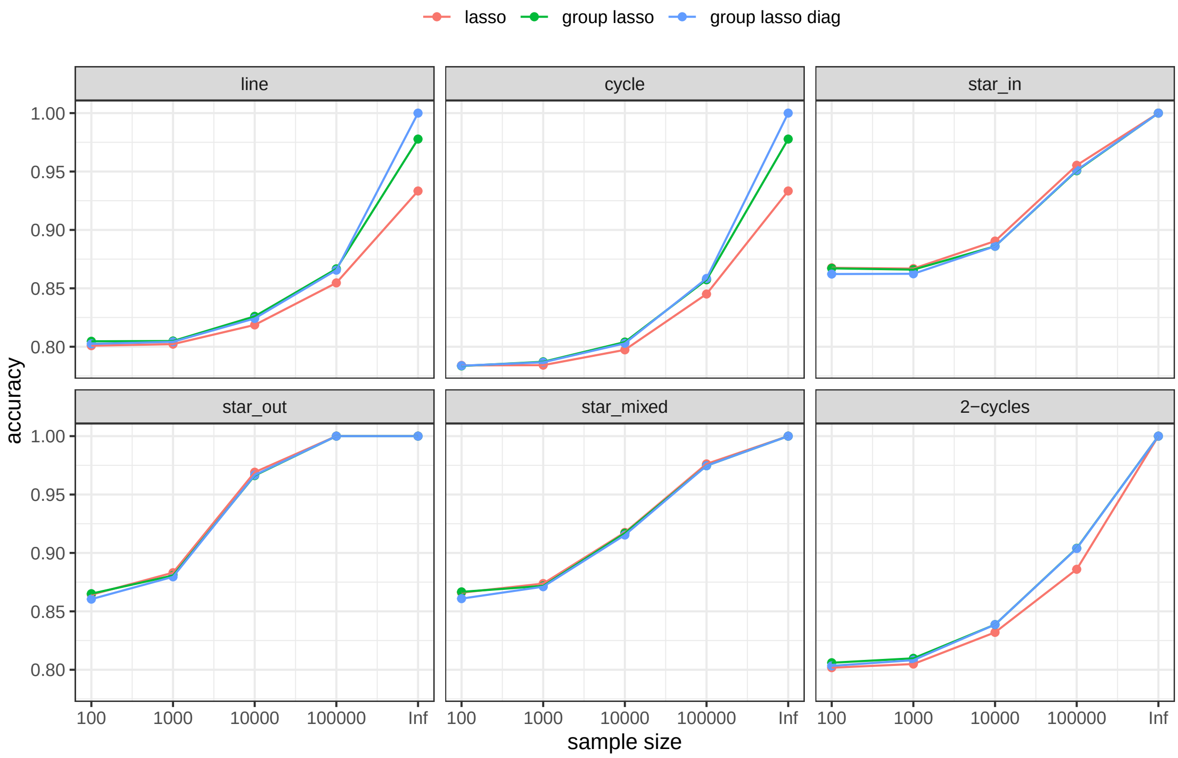

Figure 5.3: Structure recovery results for selected graphs. Accuracy given for different sample sizes and algorithms (color). The label group lasso represents the group lasso penalizing the diagonal entries, whereas group lasso diag represents the second variant not penalizing the diagonal entries.

5.2.3 Permuting the diagonal

This experiment is motivated by the fact that the ordering of the diagonal elements of M has an impact on the successful directed structure recovery via the direct Lyapunov lasso. For instance, Theorem 3.54 states that around a neighborhood of a diagonal signal, where $M_{ii} < M_{jj}$ holds for every $j \to i \in E$, we can find a DAG that will satisfy the irrepresentability condition.
To explore the effect of the ordering of diagonal elements on undirected structure recovery, we permute the diagonal elements of M. For a random subset (of size 1000) of all permutations (10! in total), we evaluate the accuracy of the estimated $\hat{M}$. We restrict ourselves to the $n = \infty$ case, i.e., we use the population covariance matrix Σ.
According to Figure 5.4, all methods achieve 100% accuracy for the 2-cycle graph regardless of the permutation. The group lasso that does not penalize the diagonal elements

of M is able to correctly recover the true sparsity pattern for all graphs. On the other hand, there exist permutations that will make it impossible for the first variant of the group lasso and the lasso to correctly estimate the undirected structure of a GCLM. However, the former achieves a higher accuracy on average across all considered graphs. In particular, the number of cases where the group lasso achieves 100% accuracy is significantly larger compared to the lasso.

In summary, we conclude that the group lasso not penalizing the diagonal entries outperforms the standard group lasso and the lasso for all considered graphs. While the performance is relatively comparable for smaller sample sizes, for $n \to \infty$, it was the only method able to achieve 100% accuracy consistently. Moreover, not penalizing the diagonal entries for the group lasso seems to make the correct support recovery invariant to the ordering of the diagonal elements.

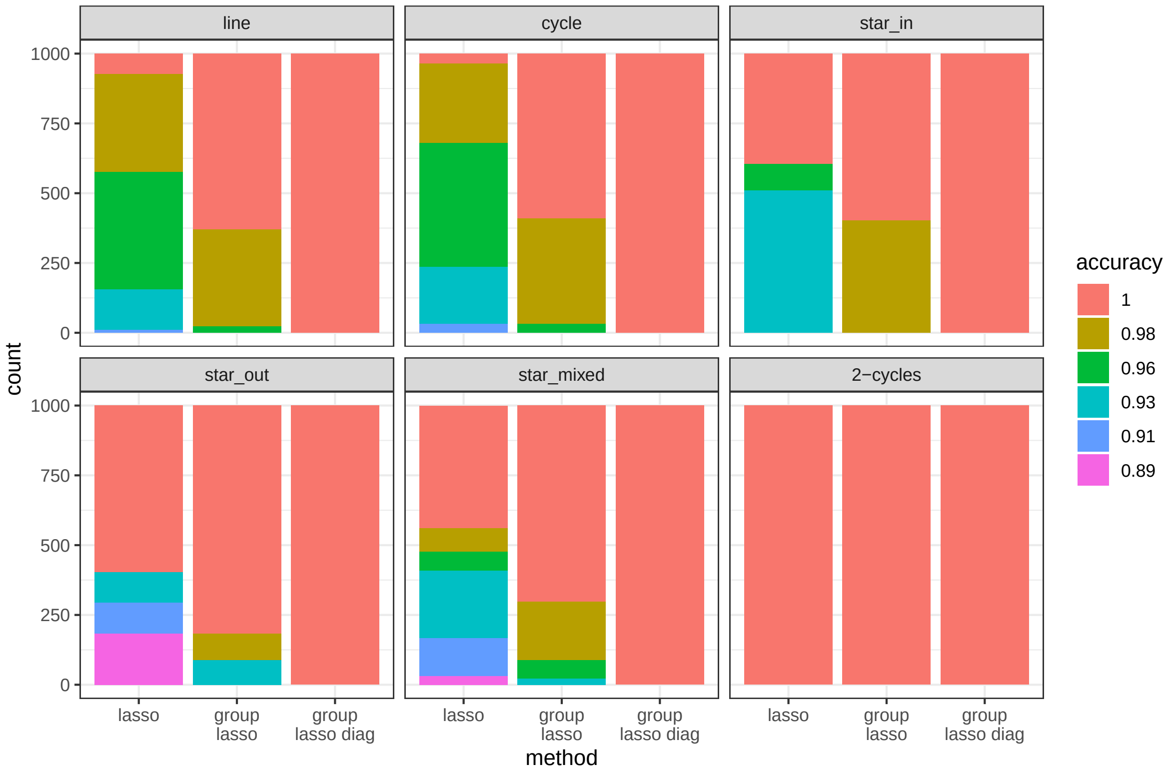

Figure 5.4: Accuracy for structure recovery under different permutations of the diagonal elements for selected graphs. The label group lasso represents the group lasso penalizing the diagonal entries, whereas group lasso diag represents the second variant not penalizing the diagonal entries.

5.3 Performance comparison for random GCLMs

We run a simulation study to compare the performance in terms of recovering the directed and undirected part of a random GCLMs for different methods, namely the lasso and the group lasso. By random GCLMs, we mean GCLMs associated with randomly generated graphs. Note we use the variant of the group lasso, which does not penalize the diagonal entries (Definition 3.6).

5.3.1 Simulation setup

The simulation setup is similar to the experiment conducted in Varando and Hansen (2020). We explore the performance of our structure recovery methods for different models of size $p = 10, ..., 50$ and edge probability $d = \frac{k}{p}$ with $k \in \{1, 2, 3, 4\}$. For each pair (p, k), we create 100 drift matrices M as mentioned below.

1. Generate a matrix M with entries
$$M_{ij} = \begin{cases} \omega_{ij}|u_{ij}| & \text{for } i \neq j, \\ -\sum_{k:k\neq i} |M_{ik}| - |u_{ii}| & \text{for } i = j, \end{cases}$$
where $\omega_{ij} \sim \text{Bernoulli}(d)$ and $u_{ij} \sim \text{Uniform}([0.1, 1])$. We set C to be a diagonal matrix with entries $C_{ii} \sim \text{Uniform}([0, 1])$.
2. For each (M, C) we solve the continuous Lyapunov equation (1.3) for Σ.
3. Generate $n = 1000$ observations $x_i, \ldots, x_n$ from a multivariate Gaussian $\mathcal{N}(0, \Sigma)$.
4. Finally, we obtain our design matrix $A(\hat{\Sigma})$ by computing the covariance matrix $\hat{\Sigma}$ based on our observations and plugging it into the respective equation (3.4).

Step 1. produces stable Metzler matrices M (Briat (2017), Lemma 2.8) that will guarantee unique solutions to the Lyapunov equation (cf. Theorem 2.9). The drift matrices M correspond to mixed graphs $G = ([p], E)$ with self-loops at every node and directed edges independently generated with uniform probability d.

5.3.2 Results

We introduce a few additional metrics to assess the performance for random GCLMs.

- *Path-wise maximum F1 score* (*maxf*1): $maxf1 = \max_\lambda \frac{2TP(\lambda)}{2TP(\lambda)+FP(\lambda)+FN(\lambda)}$.
- *Area under ROC curve* (*auroc*): The ROC curve displays the true positive rate on the y-axis and the false positive rate on the x-axis for each value of the regularization parameter λ.
- *Area under precision-recall curve* (*auprc*): The PR curve displays the precision on the y-axis and the recall on the x-axis for each value of the regularization parameter λ.

We also calculate the same metrics for the recovery of the directed graph. They can be computed analogously to the undirected case, with the only difference being the definitions for TP, FP, TN and FN, which are now computed by considering the directed graph recovery as a classification problem over the $p(p-1)$ possible edges. We obtain the directed graph from the undirected graph estimated by the group lasso by translating every undirected edge into the two possible directed edges.

Figure 5.5 displays the results of our experiment averaged over 100 repetitions and different edge probabilities d. For increasingly sparser models, i.e., increasing model size p, the maximal accuracy maxacc and auroc improve for all methods and classification problems. Accounting for the class imbalance due to the sparse models, the curves for aupr and maxf1 reveal a different trend. Both metrics are actually decreasing with increasing model size p.
According to all four metrics, the lasso is superior to the group lasso when considering randomly drawn sparse graphs. This is true for both directed and undirected structure recovery. Moreover, the evaluations for lasso and group lasso are highly similar in the sense that the evaluation curves for all metrics run almost parallel for all dimensions p.
Comparing the performance of undirected and directed structure recovery against each

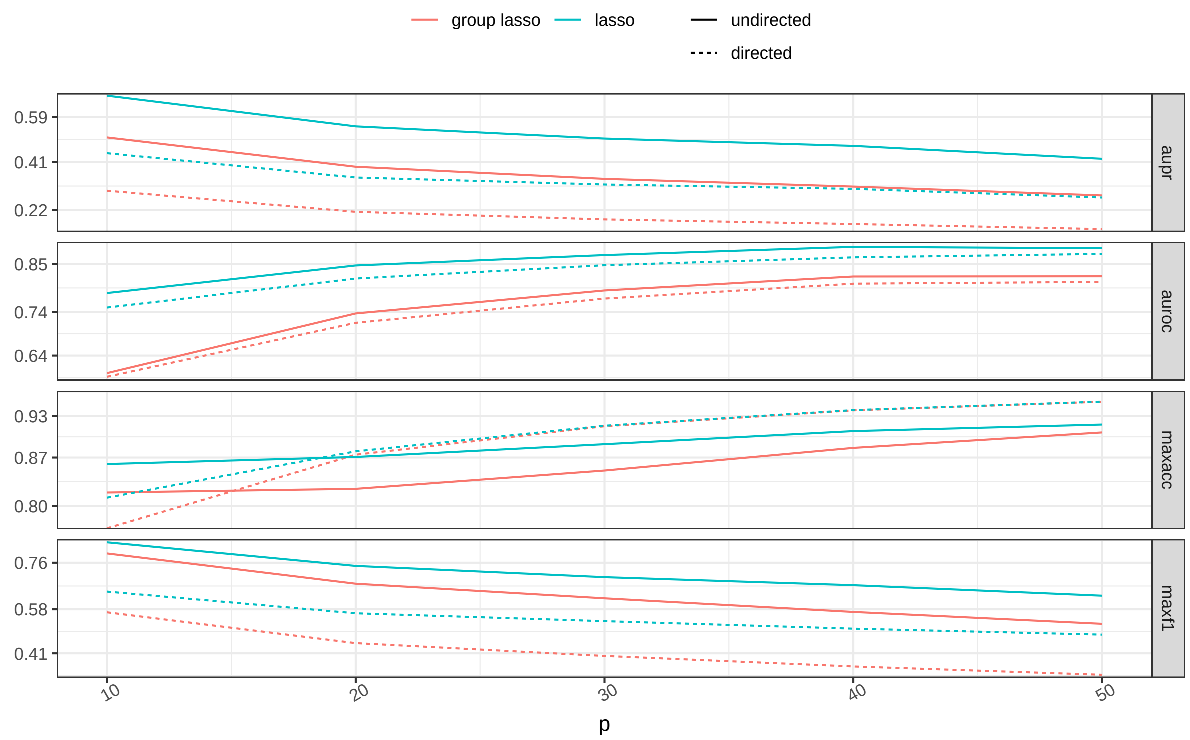

Figure 5.5: Structure recovery simulation results for undirected and directed graphs (line type). Average evaluation metrics (rows) as a function of the model size for different algorithms (colors)

other has to be done with care. Recall that they both correspond to classification problems with a different number of total samples $p(p-1)/2$ vs. $p(p-1)$. Consequently, for example misclassifying one undirected edge for an undirected graph will always lead to a lower accuracy compared to misclassifying one directed edge for a directed graph. This also explains why maxacc for directed structure recovery is higher than for undirected structure recovery. However, considering the remaining metrics, undirected structure recovery achieves a better performance.

One possible scenario where one could hope that the group lasso would outperform the lasso in terms of correct support recovery for undirected graphs is the case where we only consider 2-cycle graphs, i.e., directed graphs where every edge corresponds to a 2-cycle. Lemma 3.59, for instance, states that the group irrepresentability condition is always satisfied when the irrepresentability condition for the lasso is met. Recall that irrepresentability is a sufficient condition for correct support recovery.
We employ the same simulation setup as before (cf. Subsection 5.3.1) with a slight alteration in the first step. Namely,

1. For $i \leqslant j$ generate the entries of M as

$$M_{ij} = \begin{cases} \omega_{ij}|u_{ij}| & \text{for } i \neq j, \\ -\sum_{k:k\neq i}|M_{ik}| - |u_{ii}| & \text{for } i = j, \end{cases} \quad \text{and} \quad M_{ji} = |u_{ji}| \text{ if } M_{ij} \neq 0 \text{ and } i \neq j.$$

Using this procedure, we generate drift matrices associated with random 2-cycle graphs. However, as we can see in Figure 5.6 the group lasso is not able to outperform the lasso in terms of better support recovery for undirected graphs.

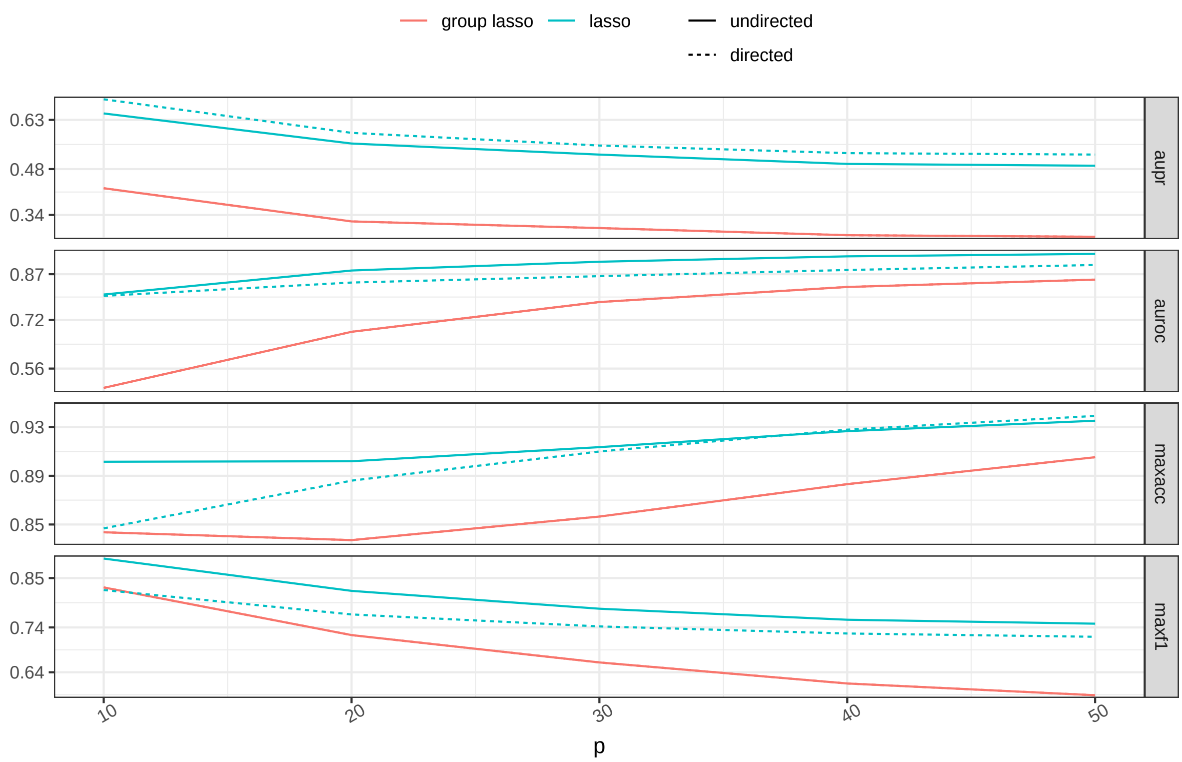

Figure 5.6: Structure recovery simulation results for undirected and directed 2-cycle graphs (line type). Average evaluation metrics (rows) as a function of the model size for different algorithms (colors)

6 Discussion

In this thesis, we proposed the group lasso as an undirected structure recovery method for GCLMs. Specifically, we showed how to find sparse solutions to the continuous Lyapunov equation using the group Lyapunov lasso. Moreover, we adapted the primal-dual witness method for the group lasso by introducing the concept of dual norms and its corresponding theory. Applying this technique, we were able to prove a deterministic result for consistent model selection for GCLMs. Restricting ourselves to Gaussian data, we can provide a bound on sample complexity for the group Lyapunov lasso by carefully investigating the Gram matrix Γ_{SS}.
Furthermore, we identified the group irrepresentability condition as a crucial sufficient assumption for consistent model selection. While it is possible to formulate conditions under which the lasso irrepresentability condition holds for DAGs, a similar analysis for the group irrepresentability condition proved to be quite difficult. The existence of no closed-form solutions for the group irrepresentability condition and restricted invertibility of the Gram matrix further complicated the analysis. For a particular type of graph, namely graphs where every edge corresponds to a 2-cycle, we showed that the lasso irrepresentability condition directly implies the group irrepresentability condition.

In the second part of the thesis, we studied algebraic questions related to GCLMs. We provided an overview for identifiability in GCLMs. The main result here is that every simple graph is globally identifiable. In addition, we examined covariance equivalence for undirected graphs. We derived a sufficient condition based on the number of edges of two connected graphs for them to be covariance equivalent. In particular, we showed that connected graphs with up to 4 nodes are always covariance equivalent. Moreover, we were able to computationally prove that trees defined for $p = 5, 6$ nodes are never covariance equivalent. However, it remains an open question to prove this fact for higher dimensions.

Lastly, we compared the performance of the direct Lyapunov lasso and the group Lyapunov lasso for directed and undirected structure recovery. For a choice of fixed graphs, the group Lyapunov has favorable properties for sample size $n \to \infty$. Namely, it seems to be invariant to the ordering of the diagonal elements of the drift matrix M. However, for randomly drawn sparse graphs, the direct Lyapunov lasso was able to outperform the group Lyapunov lasso for both directed and undirected structure recovery.

Bibliography

C. Améndola, P. Dettling, M. Drton, F. Onori, J. Wu, Structure learning for cyclic linear causal models, In *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence* (2020) pp. 999-1008.

F.R. Bach, Consistency of the Group Lasso and Multiple Kernel Learning, *Journal of Machine Learning Research* **9** (2008) pp. 1179–1225.

R.F. Barber, M. Drton, N. Sturma, and L. Weihs, Half-trek criterion for identifiability of latent variable models, arXiv, abs/2201.04457 (2022).

S. Barnett, C. Storey, Analysis and synthesis of stability matrices, *J. Differential Equations* **3** (1967) pp. 414–422.

S. Bongers, J.M. Mooji, From random differential equations to structural causal models: the stochastic case, arXiv:1803.08784 (2018).

S. Bongers, P. Forré, J. Peters, J. Mooij, Foundations of structural causal models with cycles and latent variables, arXiv: 1611.06221 (2021).

C. Briat, Sign properties of Metzler matrices with applications, *Linear Algebra and its Applications* **515** (2017) pp. 53-86.

F. Bunea, J. Lederer, Y. She, The Group Square-Root Lasso: Theoretical Properties and Fast Algorithms, *IEEE Transactions on Information Theory* **60:2** (2014) pp. 1313–1325.

P. Dettling, M. Drton, M. Kolar, On the lasso for graphical continuous Lyapunov models, arXiv:2208.13572 (2022).

P. Dettling, R. Homs, C. Amendola, M. Drton, N.R. Hansen, Identifiability in continuous Lyapunov models, arXiv:2209.03835 (2022).

N.R. Draper, H. Smith, Applied Regression Analysis, Wiley Series in Probability and Statistics (2014) pp. 369–386.

M. Drton, C. Fox, Y.S. Wang, Computation of maximum likelihood estimates in cyclic structural equation models, *Ann. Statist.* **47:2** (2019) pp. 663-690.

M. Drton, Algebraic Problems in Structural Equation Modeling, arXiv:1612.05994 (2016).

F.M. Fisher, A correspondence principle for simultaneous equation models, *Econometrica* **38:1** (1970) pp. 73-92.

K.E. Fitch, Learning directed graphical models from Gaussian data, CoRR, abs/1906.08050 (2019).

J. L. Gross and J. Yellen, Handbook of Graph Theory, CRC Press (2004) p. 524.

T. Hastie, R. Tibshirani, M. Wainwright, Statistical learning with sparsity, CRC Press (2015).

R.A. Horn, C.R. Johnson, Topics in Matrix Analysis, Cambridge University Press (1991).

R.A. Horn, C.R. Johnson, Matrix Analysis, Cambridge University Press (2013).

L. Lin, M. Drton, and A. Shojaie, Estimation of high-dimensional graphical models using regularized score matching, *Electron. J. Stat.* **10:1** (2016) pp. 806-854.

K. Lounici, M. Pontil, S. van de Geer, A.B. Tsybakov, Oracle inequalities and optimal inference under group sparsity, *Annals of Statistics* **39:4** (2011) pp. 2164-2204.

M. Maathuis, M. Drton, S. Lauritzen, M. Wainwright, Handbook of graphical models, CRC Press (2019).

C.A. Miccheli, J. Morales, M. Pontil, A Family of Penalty Functions for Structured Sparsity, In *Advances in Neural Information Processing Systems, NIPS 2010* volume 23 (2010) pp. 1612-1623.

S.W. Mogensen, D. Malinsky, N. R. Hansen, Causal learning for partially observed stochastic dynamical systems, *In Proceedings of the UAI* (2018).

J.M. Mooji, D. Janzing, B. Schölkopf, From ordinary differential equations to structural causal models: The deterministic case, In *In Proceedings of the 29th Conference on Uncertainty in Artificial Intelligence* (2013) pp. 440–448.

Y. Nardi, A. Rinaldo, On the asymptotic properties of the group lasso estimator for linear models, *Electronic Journal of Statistics* **2** (2008) pp. 605-633.

J. Pearl, Causality, Cambridge University Press (2009).

J. Peters, D. Janzing, B. Schölkopf, Elements of causal inference, MIT Press (2017).

P. Ravikumar, M.J. Wainwright, G. Raskutti, and B. Yu, High-dimensional covariance estimation by minimizing ℓ_1-penalized log-determinant divergence, *Electron. J. Stat.* **5** (2011) pp.935-980.

T.Richardson, P.Spirtes, Ancestral graph Markov models, *Ann. Statist.* **30:4** (2002) pp. 962-1030.

T. Richardson, A discovery algorithm for directed cyclic graphs, In *Uncertainty in artificial intelligence* (1996) pp. 454-461.

P. Spirtes, C. Glymour, R. Scheines, Causation, prediction, and search, MIT Press (2000).

S. van de Geer, Estimation and Testing Under Sparsity, Springer (2016).

G. Varando, N.R. Hansen, Graphical continouous Lyapunov models, In *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence* (2020) pp. 989-998.

M.J. Wainwright, Sharp thresholds for high-dimensional and noisy sparsity recovery using $\ell 1$-constrained quadratic programming (lasso), *IEEE Transactions on Information Theory* **55:5** (2009) pp.2183–2202.

M.J. Wainwright, High-Dimensional Statistics, Cambridge University Press (2019).

F. Wei, J. Huang, Consistent group selection in high-dimensional linear regression, *Bernoulli* **16:4** (2010) pp. 1369-1384.

Y. Yang, H. Zou, A Fast Unified Algorithm for Computing Group-Lasso Penalized Learning Problems, *Statistics and Computing* **25(6)** (2015) pp. 1129-1141.

M. Yuan, Y. Lin, Model selection and estimation in regression with grouped variables, *Journal of The Royal Statistical Society Series B* **68:1** (2006) pp. 49–67.

M.Yuan, Y. Lin, On the non-negative garrotte estimator, *Journal of The Royal Statistical Society Series B* **69:2** (2007) pp. 143-161.

P. Zhao, B. Yu, On model selection consistency of Lasso, *Journal of Machine learning research* **7** (2006) pp. 2541-2563.