

# Graphs in Financial Markets

Daniel P. Palomar

The Hong Kong University of Science and Technology (HKUST)

*Keynote Talk*

Graph Signal Processing Workshop (GSPW 2026)


Madrid, Spain

June 8–10, 2026

# Outline

- 1 Graphs
- 2 Graph Learning 101
- 3 Polynomial Graphical Lasso
- 4 Financial Data
- 5 Learning Heavy-Tailed Graphs
- 6 Learning Structured Graphs
- 7 Learning Time-Varying Graphs
- 8 Summary


# Graph Research Collaborators



Licheng Zhao



Sandeep Kumar



Jiaxi Ying



Vinícius Cardoso



Xiwen Wang



Amirhossein Javaheri



Arash Amini



Farokh Marvasti



Andrei Buciulea



Antonio G. Marques


# Outline

- 1 Graphs
- 2 Graph Learning 101
- 3 Polynomial Graphical Lasso
- 4 Financial Data
- 5 Learning Heavy-Tailed Graphs
- 6 Learning Structured Graphs
- 7 Learning Time-Varying Graphs
- 8 Summary


# Graphs Everywhere




Social network



Brain connectivity



Financial stocks




Foreign exchange

# Graphical Models

Graphical models are a way to represent and visualize the **relationships** between entities:

- **nodes** correspond to the entities (variables)
- **edges** encode the relationships between entities (dependencies between the variables)
- **edge weights  $W$**  encode the strength of the relationships



# Graph Matrices

The **adjacency matrix**  $\mathbf{W}$  directly characterizes a graph:  $W_{ij}$  contains the weight of the edge between nodes  $i$  and  $j$ .

The **degree matrix**  $\mathbf{D}$  is a diagonal matrix containing the node degrees (row sums of  $\mathbf{W}$ ):

$$\mathbf{D} = \text{Diag}(\mathbf{W}\mathbf{1}).$$

The **Laplacian**  $\mathbf{L} = \mathbf{D} - \mathbf{W}$  is the key matrix in graph signal processing:

- *Physical interpretation*: measures the **smoothness** of a graph signal  $\mathbf{x} = (x_1, \dots, x_p)$ :

$$\mathbf{x}^T \mathbf{L} \mathbf{x} = \frac{1}{2} \sum_{i,j} W_{ij} (x_i - x_j)^2$$


- *Mathematical properties*: symmetric and positive semidefinite ( $\mathbf{L} \succeq \mathbf{0}$ ); singular with  $\mathbf{L}\mathbf{1} = \mathbf{0}$

# Outline

- 1 Graphs
- 2 Graph Learning 101**
- 3 Polynomial Graphical Lasso
- 4 Financial Data
- 5 Learning Heavy-Tailed Graphs
- 6 Learning Structured Graphs
- 7 Learning Time-Varying Graphs
- 8 Summary


# Learning Graphs from Data

- Given a data matrix  $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_p] \in \mathbb{R}^{T \times p}$ , each column  $\mathbf{x}_i \in \mathbb{R}^T$  contains  $T$  observations from node  $i$  (and each row is one observation of the **graph signal**).
- The goal is to obtain a **graph representation** of the data.



# Graph Learning Illustration

Example of learning a graph from a dataset with points in  $\mathbb{R}^2$ :



Overview references: (Mateos et al. 2019) and (Dong et al. 2019).

# Learning Graphs from Smooth Signals

- Since  $\mathbf{x}^T \mathbf{L} \mathbf{x}$  measures the smoothness of a graph signal  $\mathbf{x}$ , the total smoothness over  $T$  observations collected along the rows of matrix  $\mathbf{X}$  is  $\text{Tr}(\mathbf{X} \mathbf{L} \mathbf{X}^T)$ .
- Learn the graph by finding the Laplacian that makes the signals smooth (Kalofolias 2016):

$$\begin{aligned} & \underset{\mathbf{L} \succeq \mathbf{0}}{\text{minimize}} && \text{Tr}(\mathbf{X} \mathbf{L} \mathbf{X}^T) - \alpha \mathbf{1}^T \log(\text{diag}(\mathbf{L})) + \frac{\beta}{2} \|\mathbf{L}\|_{F,\text{off}}^2 \\ & \text{subject to} && \mathbf{L} \mathbf{1} = \mathbf{0}, \quad L_{ij} = L_{ji} \leq 0, \quad \forall i \neq j. \end{aligned}$$

where  $\text{diag}(\mathbf{L})$  contains the node degrees, so  $-\log(\text{diag}(\mathbf{L}))$  penalizes zero degrees (no isolated nodes); and  $\|\mathbf{L}\|_{F,\text{off}}^2$  penalizes nodes with very few connections.

- The same formulation can be equivalently rewritten in terms of the adjacency matrix  $\mathbf{W}$  instead of  $\mathbf{L}$ .
- Other formulations: (Dong et al. 2015) and (Nie et al. 2016).

- **MLE of the covariance matrix**  $\Sigma$  from  $T$  observations (with sample covariance  $\mathbf{S}$ ):

$$\underset{\Sigma \succ \mathbf{0}}{\text{maximize}} \quad \log \det(\Sigma^{-1}) - \text{Tr}(\Sigma^{-1} \mathbf{S}).$$

- The **precision matrix**  $\Theta = \Sigma^{-1}$  encodes partial correlations:  $\Theta_{ij} = 0$  implies nodes  $i$  and  $j$  are conditionally independent given all others — this defines the graph edge structure.
- **Graphical LASSO (GLASSO)** adds an  $\ell_1$  penalty to promote sparsity in  $\Theta$  (Meinshausen and Bühlmann 2006):

$$\underset{\Theta \succ \mathbf{0}}{\text{maximize}} \quad \log \det(\Theta) - \text{Tr}(\Theta \mathbf{S}) - \rho \|\Theta\|_{1,\text{off}},$$

where  $\rho$  controls the sparsity level.

# Learning Graphs from a Gaussian Markov Random Field

- Modeling graph signals as a **Gaussian Markov Random Field (GMRF)**, the MLE of the graph Laplacian is (Lake and Tenenbaum 2010; Egilmez, Pavez, and Ortega 2017; Zhao et al. 2019):

$$\begin{aligned} & \underset{\mathbf{L} \succeq \mathbf{0}}{\text{maximize}} && \log \text{gdet}(\mathbf{L}) - \text{Tr}(\mathbf{S}\mathbf{L}) - \rho \|\mathbf{L}\|_{1,\text{off}} \\ & \text{subject to} && \mathbf{L}\mathbf{1} = \mathbf{0}, \quad L_{ij} = L_{ji} \leq 0, \quad \forall i \neq j, \end{aligned}$$

where  $\mathbf{S}$  is the sample covariance matrix.

- Surprisingly, the  $\ell_1$ -norm regularization  $\|\mathbf{L}\|_{1,\text{off}}$  does **not** promote sparsity — a counter-intuitive result (Ying, M. Cardoso, and Palomar 2020, 2021).
- Instead, one must use a **concave regularizer** to achieve sparse graphs.
- A concave regularizer leads to a nonconvex problem, but it can be solved iteratively via **majorization-minimization (MM)**, yielding a **reweighted  $\ell_1$ -norm** method (Kumar et al. 2020; Ying, M. Cardoso, and Palomar 2020).

## Trick: Laplacian Operator\*


- $\mathbf{L}$  contains  $p^2$  elements but the structural constraints reduce the effective degrees of freedom to  $p(p-1)/2$ . A more compact representation uses the vector  $\mathbf{w} \in \mathbb{R}_+^{p(p-1)/2}$  of edge weights (off-diagonal upper-triangular entries of  $\mathbf{L}$ ).
- The linear **Laplacian operator**  $\mathcal{L}(\mathbf{w})$  reconstructs the full Laplacian — e.g., for a 3-node graph:

$$\mathcal{L}(\mathbf{w}) = \begin{bmatrix} w_1 + w_2 & -w_1 & -w_2 \\ -w_1 & w_1 + w_3 & -w_3 \\ -w_2 & -w_3 & w_2 + w_3 \end{bmatrix}.$$


- $\mathbf{L} = \mathcal{L}(\mathbf{w})$  automatically satisfies all structural constraints  $\mathbf{L}\mathbf{1} = \mathbf{0}$ ,  $L_{ij} = L_{ji} \leq 0$ ,  $\forall i \neq j$ , so the optimization can be carried out directly over  $\mathbf{w} \geq \mathbf{0}$  (Kumar et al. 2019, 2020).

# Numerical Experiments: Effect of Sparsity Regularizer

S&P 500 stocks (3 sectors: Industrials, Consumer Staples, Energy; 2016–2019):



GMRF graph ( $\ell_1$ -norm)



GMRF graph (reweighted  $\ell_1$ -norm)

► The  $\ell_1$ -norm produces a dense graph — motivating the need for a concave (reweighted  $\ell_1$ ) regularizer.

# Outline

- 1 Graphs
- 2 Graph Learning 101
- 3 Polynomial Graphical Lasso**
- 4 Financial Data
- 5 Learning Heavy-Tailed Graphs
- 6 Learning Structured Graphs
- 7 Learning Time-Varying Graphs
- 8 Summary

# Graph Shift Operator and Graph Stationarity

- The **Graph Shift Operator (GSO)**  $\mathbf{S} \in \mathbb{R}^{p \times p}$  encodes the graph topology:  $S_{ij} \neq 0$  iff  $i = j$  or nodes  $i$  and  $j$  are connected by an edge. Common choices: adjacency matrix  $\mathbf{W}$  or Laplacian  $\mathbf{L}$ .
- A **graph filter** is a polynomial of the GSO:  $\mathbf{H} = \sum_{k=0}^K h_k \mathbf{S}^k$ . It operates locally — each node aggregates information from its  $K$ -hop neighborhood.
- **Physical interpretation:** since  $\mathbf{S}^k$  propagates a signal  $k$  hops through the graph,  $\mathbf{H}$  captures **multi-hop (delayed) interactions** among nodes — as opposed to an unstructured precision matrix  $\Theta$ , which only models instantaneous pairwise interactions.
- A graph signal  $\mathbf{x}$  is **stationary on  $\mathbf{S}$**  if its covariance  $\Sigma$  commutes with  $\mathbf{S}$ :  $\Sigma \mathbf{S} = \mathbf{S} \Sigma$ . Equivalently,  $\Sigma$  and  $\mathbf{S}$  share the same eigenvectors, i.e.,  $\Sigma$  is a polynomial of  $\mathbf{S}$ :

$$\Sigma = \mathbf{H}^2 = \left( \sum_{k=0}^K h_k \mathbf{S}^k \right)^2 = p(\mathbf{S}).$$

# The PGL Formulation

- **Recall:** the Graphical LASSO estimates the sparse precision matrix  $\Theta = \Sigma^{-1}$ :

$$\underset{\Theta \succ \mathbf{0}}{\text{maximize}} \quad \log \det(\Theta) - \text{Tr}(\Theta \hat{\Sigma}^{\text{SCM}}) - \rho \|\Theta\|_{1,\text{off}},$$

where  $\hat{\Sigma}^{\text{SCM}}$  is the sample covariance matrix; but ignores any underlying graph structure.

- **PGL** jointly estimates  $\Theta$  and the GSO  $\mathbf{S}$  under stationarity — the commutativity constraint  $\Theta \mathbf{S} = \mathbf{S} \Theta$  encodes that  $\Theta$  is a polynomial of  $\mathbf{S}$  (Buciulea et al. 2025):

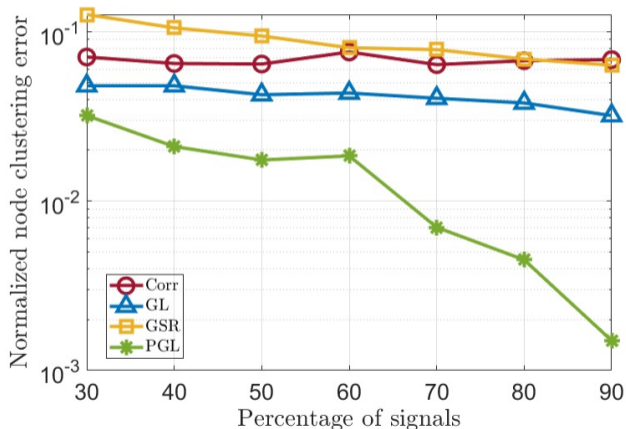
$$\begin{aligned} & \underset{\Theta \succ \mathbf{0}, \mathbf{S} \in \mathcal{S}}{\text{maximize}} && \log \det(\Theta) - \text{Tr}(\hat{\Sigma}^{\text{SCM}} \Theta) - \rho \|\mathbf{S}\|_0 \\ & \text{subject to} && \Theta \mathbf{S} = \mathbf{S} \Theta, \end{aligned}$$

where  $\mathcal{S}$  encodes structural constraints on  $\mathbf{S}$  (symmetric, non-negative off-diagonal, zero diagonal,  $\mathbf{S} \mathbf{1} \geq \mathbf{1}$ ).

- **Two challenges:** the  $\ell_0$  sparsity penalty and the commutativity constraint  $\Theta \mathbf{S} = \mathbf{S} \Theta$  are both non-convex.
- **Solution strategy:** relax  $\ell_0$  to a reweighted  $\ell_1$ , then apply **Majorization-Minimization (MM)** (Sun, Babu, and Palomar 2017) — alternating between two closed-form updates (Buciulea et al. 2025):
  - Update  $\Theta$  with  $\mathbf{S}$  fixed:  $\Theta^* = \mathbf{V} \text{diag}\left(\frac{1}{2}\left(-\hat{\lambda} + \sqrt{\hat{\lambda}^2 + 4}\right)\right) \mathbf{V}^T$ , where  $\hat{\lambda}$  are the eigenvalues of  $\hat{\Sigma}^{\text{SCM}}$  in the basis of  $\mathbf{S}$ .
  - Update  $\mathbf{S}$  with  $\Theta$  fixed: element-wise soft-thresholding in the eigenbasis of  $\Theta$ .
- **Guarantees:** converges to a stationary point; a fast variant (GST-fast) reduces computational complexity significantly.

# PGL Numerical Experiments: Stock Clustering

S&P 500 data (40 stocks, 4 sectors): clustering error vs. fraction of data used for graph learning (Buciulea et al. 2025):



► *PGL recovers the sector structure more accurately and with less data than GL, GSR, and correlation baselines.*

# Outline

- 1 Graphs
- 2 Graph Learning 101
- 3 Polynomial Graphical Lasso
- 4 Financial Data**
- 5 Learning Heavy-Tailed Graphs
- 6 Learning Structured Graphs
- 7 Learning Time-Varying Graphs
- 8 Summary

# Prices and Returns

- **Price:**  $p_t$  (discrete time periods: minutes to years)

- **Log-price:**  $y_t \triangleq \log p_t$

- **Linear return:**

$$r_t^{\text{lin}} \triangleq \frac{p_t - p_{t-1}}{p_{t-1}}$$

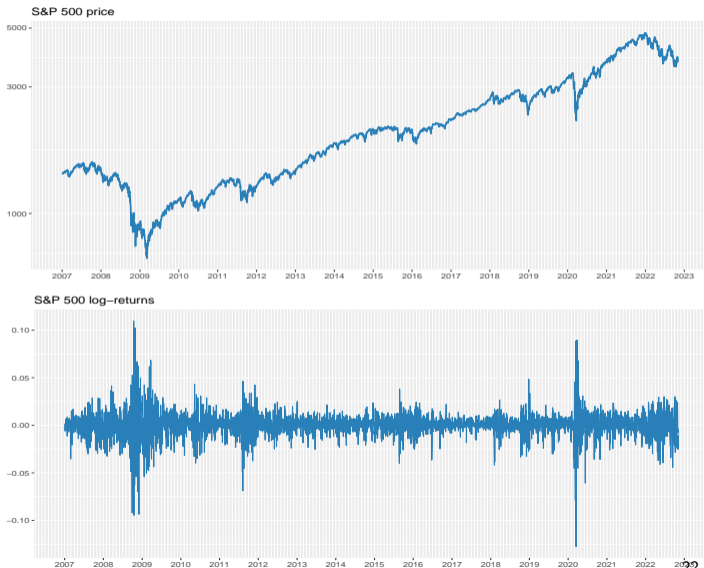
- **Log-return:**

$$r_t \triangleq y_t - y_{t-1} = \log\left(\frac{p_t}{p_{t-1}}\right)$$

- **Random walk model:**

$$y_t = \mu + y_{t-1} + \epsilon_t \Rightarrow r_t = \mu + \epsilon_t$$

(Ruppert 2010; Palomar 2025, ch.2)



Empirical regularities observed across instruments, markets, and time periods (Cont 2001):

## **Absence of Return Autocorrelations**

- Returns are largely unpredictable (efficient-market hypothesis, Fama 1970)
- But  $|r_t|$  and  $r_t^2$  are autocorrelated — nonlinear dependence persists

## **Volatility Clustering / Lack of Stationarity**

- Large moves tend to follow large moves; variance changes over time (Mandelbrot 1963)
- ⇒ *Addressed in Section 7: Learning Time-Varying Graphs*

## **Heavy Tails**

- Distributions exhibit much heavier tails than Gaussian; more frequent extreme events
- ⇒ *Addressed in Section 5: Learning Heavy-Tailed Graphs*

## **Correlation Structure**

- Assets co-move: pairwise correlations are predominantly positive and reveal market factor structure
- ⇒ *Addressed in Section 6: Learning Structured Graphs*

## Stylized Fact 1: Heavy Tails

**Gaussian distribution** (fully characterized by mean and variance):

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2\right)$$

**Skewness** (asymmetry) and **kurtosis** (tail thickness):

$$\gamma_1 = \frac{\mathbb{E}[(r-\mu)^3]}{\sigma^3}, \quad \kappa = \frac{\mathbb{E}[(r-\mu)^4]}{\sigma^4}$$

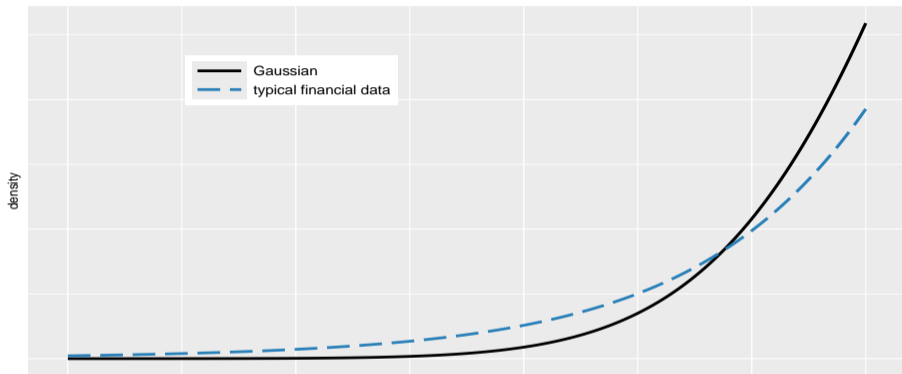
- Gaussian:  $\gamma_1 = 0$ ,  $\kappa = 3$  (excess kurtosis = 0)

**Heavy tails:** kurtosis  $> 3$  — tail decay slower than Gaussian

- More frequent extreme events than Gaussian predicts
- Financial returns exhibit heavier tails, especially on the left (large losses)

(Ruppert 2010; Palomar 2025, ch.2)

## Stylized Fact 1: Heavy Tails (cont.)




### Implications for finance:


- **Risk underestimation:** Gaussian models underestimate the probability of crashes and large losses
- **Portfolio optimization:** ignoring heavy tails leads to portfolios that are too aggressive

# Stylized Fact 1: Heavy Tails (S&P 500)

S&P 500 log-return histograms at increasing frequencies:



Q-Q plots vs. Gaussian (deviations from the line = heavy tails):




## Stylized Fact 2: Volatility Clustering

### Efficient-Market Hypothesis (EMH)

- Share prices reflect **all available information** at all times (Fama, 1970)
- Neither technical nor fundamental analysis can consistently generate excess returns
- In practice: financial returns behave approximately as **i.i.d.**.

### Linear temporal structure: ACF of returns


- ACF measures linear dependence between  $r_t$  and  $r_{t-k}$
- EMH predicts no significant autocorrelation in returns
- Empirically confirmed for S&P 500: ACF lags lie within noise bands



## Stylized Fact 2: Volatility Clustering (cont.)

### Nonlinear structure: the volatility envelope clusters


- Large changes tend to follow large changes (high volatility periods)
- Small changes tend to follow small changes (calm periods)
- Documented by Mandelbrot (1963) and Fama (1965)
- Modeled by ARCH/GARCH family of processes



## Stylized Fact 2: Volatility Clustering (cont.)

### ACF of $|r_t|$ : nonlinear dependence persists

- Unlike ACF of  $r_t$ , the ACF of  $|r_t|$  is significant at many lags
- Confirms that volatility — not returns — carries temporal structure



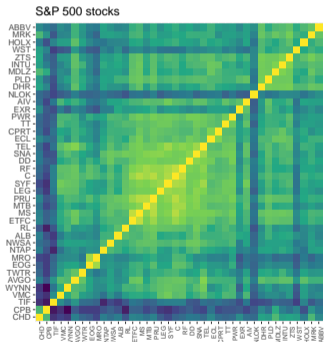
### Implications for finance:

- **Volatility is forecastable:** ARCH/GARCH models exploit this for risk management
- **Non-stationarity:** the statistics and distribution of financial returns change over time — addressed later when learning time-varying graphs

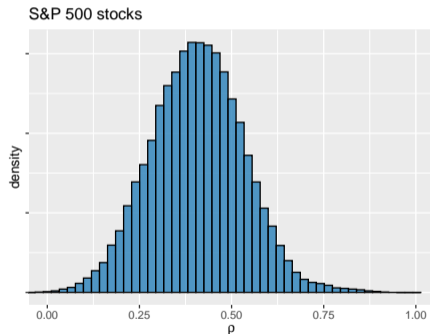
# Stylized Fact 3: Correlation Structure

**Assets are not independent — pairwise correlations are predominantly positive**

- Stocks co-move: when the market rises (falls), most stocks rise (fall) together
- Ignoring correlations leads to severe underestimation of portfolio risk



Correlation matrix of S&P 500 stocks




Distribution of pairwise correlations

(Cont 2001; Palomar 2025, ch.2)

## Stylized Fact 3: Correlation Structure (cont.)

### Factor model structure: one dominant eigenvalue

- Eigenvalues of the correlation matrix reveal the factor structure of the market
- One very large eigenvalue  $\gg$  all others: the **market factor** drives most of the variance
- Remaining smaller eigenvalues carry asset-specific (idiosyncratic) structure



**Implication:** after removing the dominant market factor, the residual correlations reflect genuine pairwise relationships between stocks — naturally modeled as a **graph** (Palomar 2025, ch.2)

# Implications for Graph Learning

The three stylized facts jointly motivate the models developed in this talk:

- **Heavy tails:** the Gaussian distribution is inadequate — a **heavy-tailed distribution model** is needed for robust graph learning (and perhaps even asymmetric)
- **Volatility clustering:** returns are not i.i.d. over time — a **time-varying graph** model is required to track the evolving market structure
- **Correlation structure:** pairwise asset relationships are non-trivial and structured (positive, low-rank, and with assets naturally **clustered** into groups of high correlation) — a **graph** is a natural and interpretable representation (connecting later to structured graph learning)

⇒ We need graph learning algorithms that handle **heavy tails**, **clustered assets**, and **time variation**

# Outline

- 1 Graphs
- 2 Graph Learning 101
- 3 Polynomial Graphical Lasso
- 4 Financial Data
- 5 Learning Heavy-Tailed Graphs**
- 6 Learning Structured Graphs
- 7 Learning Time-Varying Graphs
- 8 Summary

# Heavy-Tailed Graph Model

- Recall that the **Gaussian model**  $f(\mathbf{x}) \propto |\Sigma|^{-1/2} \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^\top \Sigma^{-1}(\mathbf{x} - \boldsymbol{\mu})\right)$  leads to the MLE problem for the graph Laplacian:

$$\begin{aligned} & \underset{\mathbf{L} \succeq \mathbf{0}}{\text{maximize}} && \log \text{gdet}(\mathbf{L}) - \text{Tr}(\mathbf{L}\mathbf{S}) \\ & \text{subject to} && \mathbf{L}\mathbf{1} = \mathbf{0}, \quad L_{ij} = L_{ji} \leq 0, \quad \forall i \neq j, \end{aligned}$$

where  $\mathbf{S} = \frac{1}{T} \sum_{t=1}^T \mathbf{x}^{(t)}(\mathbf{x}^{(t)})^\top$  is the sample covariance matrix.

- Student- $t$  model**  $f(\mathbf{x}) \propto |\Sigma|^{-1/2} \left(1 + \frac{1}{\nu}(\mathbf{x} - \boldsymbol{\mu})^\top \Sigma^{-1}(\mathbf{x} - \boldsymbol{\mu})\right)^{-(p+\nu)/2}$  with degrees of freedom  $\nu$  (heavier tails as  $\nu \downarrow$ ) leads to:

$$\begin{aligned} & \underset{\mathbf{L} \succeq \mathbf{0}}{\text{maximize}} && \log \text{gdet}(\mathbf{L}) - \frac{p + \nu}{T} \sum_{t=1}^T \log \left(1 + \frac{1}{\nu}(\mathbf{x}^{(t)})^\top \mathbf{L} \mathbf{x}^{(t)}\right) \\ & \text{subject to} && \mathbf{L}\mathbf{1} = \mathbf{0}, \quad L_{ij} = L_{ji} \leq 0, \quad \forall i \neq j. \end{aligned}$$

# Majorization-Minimization (MM) Algorithm

- The heavy-tailed formulation is **nonconvex**. The **MM framework** (Sun, Babu, and Palomar 2017) upper-bounds each log term:

$$\log\left(1 + \frac{1}{\nu}(\mathbf{x}^{(t)})^T \mathbf{L} \mathbf{x}^{(t)}\right) \leq \text{const} + \frac{(\mathbf{x}^{(t)})^T \mathbf{L} \mathbf{x}^{(t)}}{\nu + (\mathbf{x}^{(t)})^T \mathbf{L}_0 \mathbf{x}^{(t)}}.$$

- This yields a sequence of **Gaussianized** problems ( $k = 1, 2, \dots$ ):


$$\begin{aligned} & \underset{\mathbf{L} \succeq \mathbf{0}}{\text{maximize}} && \log \text{gdet}(\mathbf{L}) - \text{Tr}(\mathbf{L} \mathbf{S}^k) \\ & \text{subject to} && \mathbf{L} \mathbf{1} = \mathbf{0}, \quad L_{ij} = L_{ji} \leq 0, \quad \forall i \neq j, \end{aligned}$$

where  $\mathbf{S}^k = \frac{1}{T} \sum_{t=1}^T w_t^k \mathbf{x}^{(t)} (\mathbf{x}^{(t)})^T$  is a weighted sample covariance with weights (M. Cardoso, Ying, and Palomar 2021)


$$w_t^k = \frac{p + \nu}{\nu + (\mathbf{x}^{(t)})^T \mathbf{L}^k \mathbf{x}^{(t)}}.$$

# Numerical Experiments: Gaussian vs. Heavy-Tailed Graphs with Stocks


S&P 500 stocks (3 sectors: Industrials, Consumer Staples, Energy; 2016–2019):



GMRF graph ( $\ell_1$ -norm)



GMRF graph (reweighted  $\ell_1$ -norm)




Heavy-tailed MRF graph


► *Heavy-tailed graphs better capture the actual structure of financial data — fewer spurious edges, cleaner sector clusters.*

# Numerical Experiments: Gaussian vs. Heavy-Tailed Graphs with FX

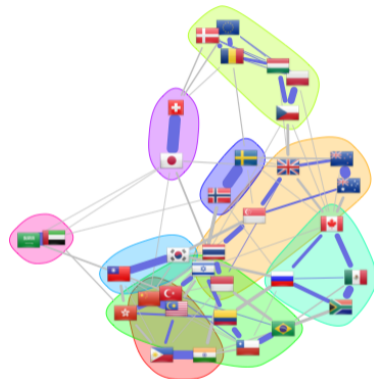
FX currencies (COVID period, 2020):



GMRF graph ( $\ell_1$ -norm)



GMRF graph (reweighted  $\ell_1$ -norm)




Heavy-tailed MRF graph

► *Heavy-tailed graphs highlight geographically close currency pairs (e.g., HKD/CNY, TWD/KRW, PLN/CZK) — invisible in Gaussian graphs.*


# Outline

- 1 Graphs
- 2 Graph Learning 101
- 3 Polynomial Graphical Lasso
- 4 Financial Data
- 5 Learning Heavy-Tailed Graphs
- 6 Learning Structured Graphs**
- 7 Learning Time-Varying Graphs
- 8 Summary


# Structured Graphs




Multi-component




Regular




Modular



Bipartite



Grid



Tree

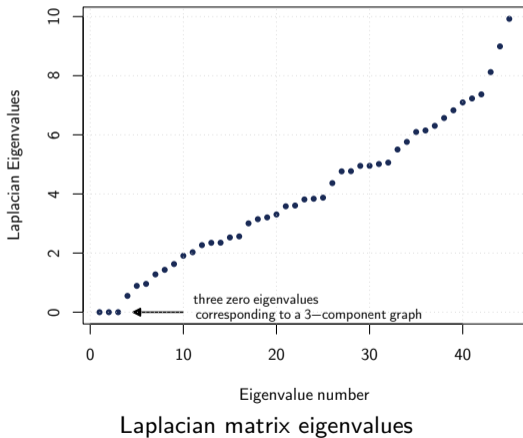

# Learning Structured Graphs via Eigenvalues

- Learning graphs with a **specific structure** is generally NP-hard; no universal algorithm exists.
- Simple structures (regular, grid): enforced by fixing elements of  $\mathbf{L}$  or controlling node degrees.
- Complex structures: characterized by **spectral properties** of  $\mathbf{L}$  or  $\mathbf{W}$  — enforced in the learning problem:
  - **$k$ -component graph:**  $\mathbf{L}$  has  $k$  zero eigenvalues (low-rank structure).
  - **Bipartite graph:**  $\mathbf{W}$  has eigenvalues symmetric around zero.

(Palomar 2025, ch. 5)

# Example of $k$ -Component Graph

3-component graph (3 clusters) along Laplacian matrix eigenvalues (3 zero eigenvalues):



# Learning $k$ -Component Graphs

- Enforce low-rank structure via **Ky Fan's theorem**:

$$\sum_{i=1}^k \lambda_i(\mathbf{L}) = \min_{\mathbf{F} \in \mathbb{R}^{p \times k}: \mathbf{F}^T \mathbf{F} = \mathbf{I}} \text{Tr}(\mathbf{F}^T \mathbf{L} \mathbf{F}),$$

so it suffices to add the penalty  $\text{Tr}(\mathbf{F}^T \mathbf{L} \mathbf{F})$  to the objective.

- Combining with the heavy-tailed MM framework yields a sequence of Gaussianized problems ( $k = 1, 2, \dots$ ) (Cardoso et al. 2021):


$$\begin{aligned} & \underset{\mathbf{L} \succeq \mathbf{0}, \mathbf{F}}{\text{maximize}} && \log \text{gdet}(\mathbf{L}) - \text{Tr}(\mathbf{L} \mathbf{S}^k) - \gamma \text{Tr}(\mathbf{F}^T \mathbf{L} \mathbf{F}) \\ & \text{subject to} && \mathbf{L} \mathbf{1} = \mathbf{0}, \quad L_{ij} = L_{ji} \leq 0, \quad \forall i \neq j, \\ & && \mathbf{F}^T \mathbf{F} = \mathbf{I}, \end{aligned}$$

where  $\mathbf{S}^k$  is the weighted sample covariance from the MM iteration.

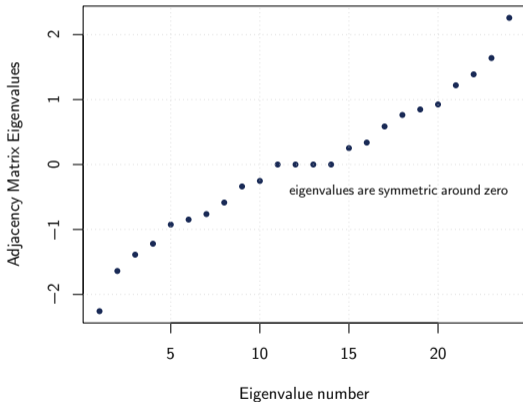
- Solved efficiently by alternating between updates of  $\mathbf{L}$  (with  $\mathbf{F}$  fixed) and  $\mathbf{F}$  (eigenvectors of  $\mathbf{L}$ ).

# Example of Bipartite Graph

Bipartite graph along adjacency matrix eigenvalues (symmetric around zero):



Graph



Adjacency matrix eigenvalues

# Learning Bipartite Graphs

- Two types of nodes with edges only between types — useful to learn connections between individual assets and sector indices.
- Enforce bipartite structure by constraining  $\mathbf{L}$  to have the block form (with  $\mathbf{B} \in \mathbb{R}_+^{r \times q}$  containing edge weights):

$$\mathbf{L} = \begin{bmatrix} \text{Diag}(\mathbf{B}\mathbf{1}) & -\mathbf{B} \\ -\mathbf{B}^\top & \text{Diag}(\mathbf{B}^\top\mathbf{1}) \end{bmatrix}.$$

- The heavy-tailed MM framework yields a sequence of Gaussianized problems ( $k = 1, 2, \dots$ ) (Cardoso et al. 2022):

$$\begin{aligned} & \underset{\mathbf{L} \succeq \mathbf{0}, \mathbf{B}}{\text{maximize}} && \log \text{gdet}(\mathbf{L}) - \text{Tr}(\mathbf{L}\mathbf{S}^k) \\ & \text{subject to} && \mathbf{L} = \begin{bmatrix} \text{Diag}(\mathbf{B}\mathbf{1}) & -\mathbf{B} \\ -\mathbf{B}^\top & \text{Diag}(\mathbf{B}^\top\mathbf{1}) \end{bmatrix}, \\ & && \mathbf{B} \geq \mathbf{0}. \end{aligned}$$

# Learning $k$ -Component Bipartite Heavy-Tailed Graphs

- Combine the two previous structures: a  $k$ -component graph with bipartite connectivity.
- This leads to a sequence of Gaussianized problems ( $k = 1, 2, \dots$ ) with both  $\mathbf{F}$  and  $\mathbf{B}$  as additional variables (M. Cardoso, Ying, and Palomar 2022):

$$\begin{aligned} & \underset{\mathbf{L} \geq \mathbf{0}, \mathbf{F}, \mathbf{B}}{\text{maximize}} && \log \text{gdet}(\mathbf{L}) - \text{Tr}(\mathbf{L}\mathbf{S}^k) - \gamma \text{Tr}(\mathbf{F}^T \mathbf{L} \mathbf{F}) \\ & \text{subject to} && \mathbf{L} = \begin{bmatrix} \text{Diag}(\mathbf{B}\mathbf{1}) & -\mathbf{B} \\ -\mathbf{B}^T & \text{Diag}(\mathbf{B}^T \mathbf{1}) \end{bmatrix}, \\ & && \mathbf{B} \geq \mathbf{0}, \quad \mathbf{F}^T \mathbf{F} = \mathbf{I}. \end{aligned}$$

# Learning Graphs with Arbitrary Structures

- Structural properties of many graph classes can be characterized by **spectral constraints**:
  - constraints on eigenvalues of the Laplacian matrix  $\mathbf{L}$
  - constraints on eigenvalues of the adjacency matrix  $\mathbf{W}$
- **General formulation** with spectral constraint (Kumar et al. 2019, 2020):


$$\begin{aligned} & \underset{\mathbf{L} \succeq \mathbf{0}}{\text{maximize}} && \log \text{gdet}(\mathbf{L}) - \text{Tr}(\mathbf{L}\mathbf{S}) - \rho \|\mathbf{L}\|_{0,\text{off}} \\ & \text{subject to} && \mathbf{L}\mathbf{1} = \mathbf{0}, \quad L_{ij} = L_{ji} \leq 0, \quad \forall i \neq j, \\ & && \lambda(\mathbf{L}) \in \mathcal{S}_\lambda, \end{aligned}$$

where  $\mathcal{S}_\lambda$  encodes the desired spectral structure (e.g.,  $k$  zero eigenvalues, symmetric eigenvalues).

- Nonconvex problem, but efficient algorithms exist based on MM and the Laplacian operator trick.

# Numerical Experiments: Gaussian vs. Heavy-Tailed $k$ -Component Graphs

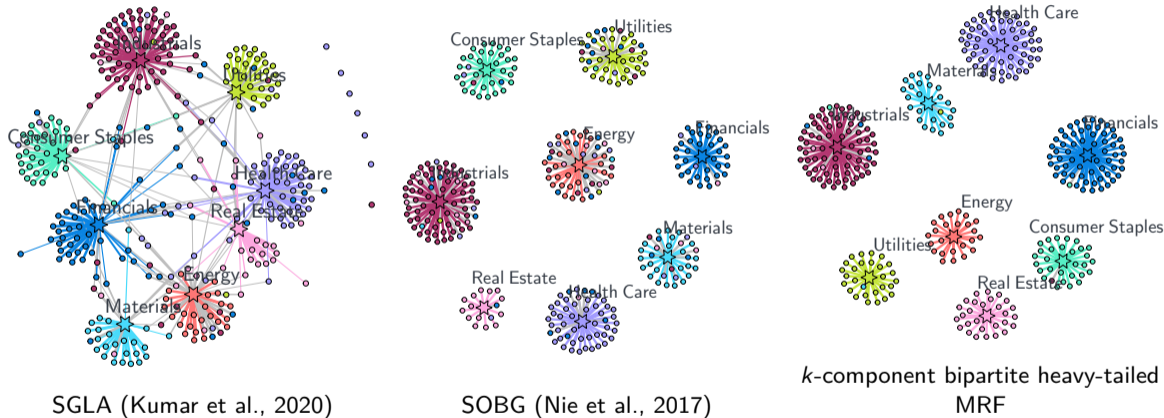
FX currencies (COVID period, 2020):



► *Heavy-tailed  $k$ -component graphs yield cleaner, more interpretable currency clusters than Gaussian counterparts.*

# Numerical Experiments: Comparison of $k$ -Component Bipartite Graphs

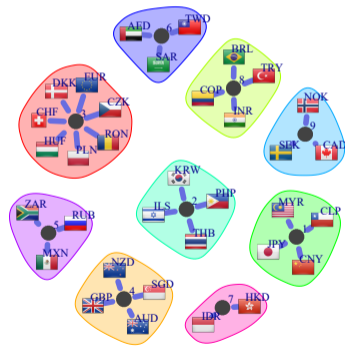
S&P 500 stocks (benchmarks: Kumar et al. 2020; Nie et al. 2017):




►  *$k$ -component bipartite heavy-tailed MRF graphs avoid isolated nodes and achieves better accuracy.*

# Numerical Experiments: $k$ -Component Bipartite Graphs with Unknown Indices

Applied to forex and crypto markets — unlike S&P 500, no predefined sector labels are available, so the  $k$  components and their cluster assignments are learned entirely from data (Javaheri and Palomar 2025).



Forex data



Crypto data

# Outline

- 1 Graphs
- 2 Graph Learning 101
- 3 Polynomial Graphical Lasso
- 4 Financial Data
- 5 Learning Heavy-Tailed Graphs
- 6 Learning Structured Graphs
- 7 Learning Time-Varying Graphs**
- 8 Summary

## Motivation: Graphs Change Over Time

- Financial markets are **non-stationary**: correlations between assets shift during crises, regime changes, and macroeconomic events.
- **Rolling-window approach** (re-estimate graph on a sliding window of data):
  - Introduces estimation delay proportional to window size
  - Ignores temporal continuity — consecutive graphs are estimated independently
  - Not robust to outliers or missing observations
- **Principled time-varying formulation**: learn a sequence of graphs  $\{\mathbf{L}_n\}$  jointly, coupling consecutive graphs through a temporal prior — capturing smooth evolution while adapting to regime shifts.

# Time-Varying Graph Model

- Data arrives in **frames**  $\mathcal{F}_n = \{t : t \in \text{frame } n\}$ ; each frame has its own graph  $\mathcal{L}(\mathbf{w}_n)$  with edge weights  $\mathbf{w}_n \in \mathbb{R}^{p(p-1)/2}$ .
- The standard approach is to penalize the difference  $\mathbf{w}_n - \mathbf{w}_{n-1}$  via the  $\ell_2$ -norm or  $\ell_1$ -norm.
- **Non-negative VAR(1) temporal evolution model**: edge weights evolve as

$$\mathbf{w}_n = (\mathbf{a} \odot \mathbf{w}_{n-1} + \boldsymbol{\varepsilon}_n)_+, \quad n = 1, 2, \dots$$

where  $\mathbf{a} \in \mathbb{R}_+^{p(p-1)/2}$  are VAR coefficients,  $\boldsymbol{\varepsilon}_n$  is zero-mean noise, and  $(\cdot)_+$  enforces non-negativity of edge weights.

- The **multiplicative factor**  $\mathbf{a}$  better captures proportional changes in edge strengths over time.

(Javaheri et al. 2025)

# Heavy Tails + Time-Varying: Combined Model

- **Laplacian operator** (Kumar et al. 2019, 2020): any non-negative edge weight vector  $\mathbf{w} \in \mathbb{R}_+^{p(p-1)/2}$  defines a valid Laplacian via the linear mapping  $\mathcal{L}(\mathbf{w})$ :  $\mathcal{L}(\mathbf{w})_{ij} = -w_{ij}$  for  $i \neq j$  and  $\mathcal{L}(\mathbf{w})_{ii} = \sum_{j \neq i} w_{ij}$ . The optimization is then carried out directly over  $\mathbf{w} \geq \mathbf{0}$ .
- **Student- $t$  likelihood** at each frame  $n$  combined with the VAR temporal prior yields (Javaheri et al. 2025):

$$\begin{aligned} \underset{\mathbf{w}_n \geq \mathbf{0}, \mathbf{a} \geq \mathbf{0}}{\text{maximize}} \quad & \log \text{gdet}(\mathcal{L}(\mathbf{w}_n)) - \frac{\nu + p}{T_n} \sum_{t \in \mathcal{F}_n} \log \left( 1 + \frac{\mathbf{x}_t^\top \mathcal{L}(\mathbf{w}_n) \mathbf{x}_t}{\nu} \right) \\ & - \alpha \|\mathbf{w}_n - \mathbf{a} \odot \hat{\mathbf{w}}_{n-1}\|_1 - \beta \|\mathbf{w}_n\|_0 - \gamma \|\mathbf{a}\|_1 \\ \text{subject to} \quad & \mathbf{w}_n \in \Omega_w, \end{aligned}$$

where  $\Omega_w$  encodes  $k$ -component structure and degree constraints.

# Handling Corrupted and Missing Data

- In practice, financial data contains **missing observations and outlier-corrupted measurements**:

$$\mathbf{y}_t = \mathbf{m}_t \odot (\mathbf{x}_t + \mathbf{n}_t), \quad \mathbf{n}_t \sim \mathcal{N}(\mathbf{0}, \sigma_n^2 \mathbf{I}),$$

where  $\mathbf{m}_t \in \{0, 1\}^P$  is a binary sampling mask ( $m_{ti} = 0$  means asset  $i$  is missing at time  $t$ ).

- The framework **jointly recovers the clean signal  $\mathbf{x}_t$  and learns the graph  $\mathbf{L}_{w_n}$**  — missing entries are imputed as a byproduct.
- The MM weights  $w_t^k \propto (\nu + \mathbf{x}_t^\top \mathbf{L}_{w_n} \mathbf{x}_t)^{-1}$  automatically **downweight outliers**: observations that are inconsistent with the graph structure receive lower weight.

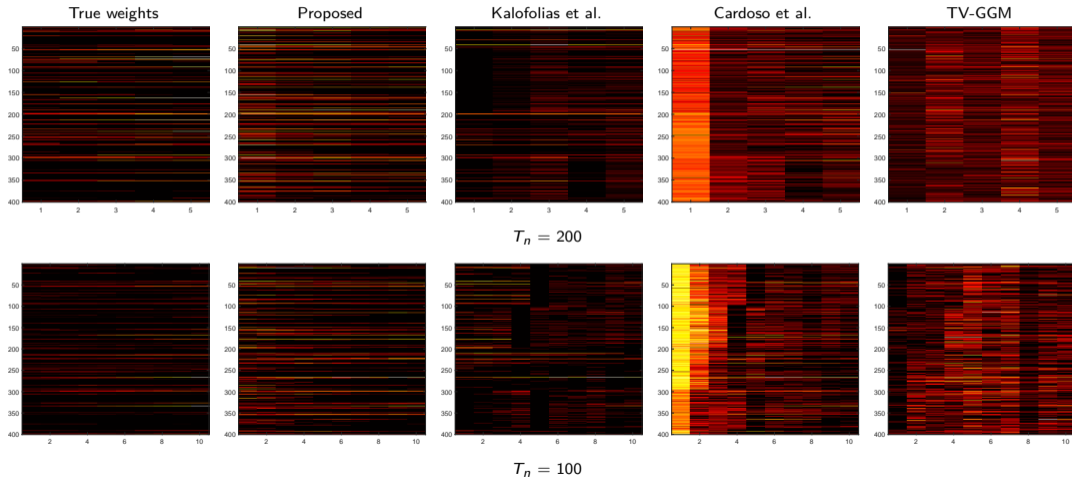
(Javaheri et al. 2025)

## Algorithm: MM + ADMM

- The combined problem is non-convex. An **MM** outer loop linearizes the Student- $t$  log terms, yielding a sequence of tractable subproblems solved via **ADMM**.
- Each iteration alternates five closed-form updates:
  - $\mathbf{L}_n$ -update: eigendecomposition enforcing  $k$ -component rank constraint
  - $\mathbf{w}_n$ -update: soft-thresholding enforcing sparsity ( $\ell_0$  relaxed to weighted  $\ell_1$ )
  - $\mathbf{X}_n$ -update: gradient step for signal recovery / missing-data imputation
  - $\mathbf{a}$ -update: soft-thresholding for VAR coefficient learning
  - Dual variable update: standard ADMM multiplier ascent
- **Complexity:**  $\mathcal{O}(p^3 + T_n p^2)$  per iteration; converges to a stationary point.
- The following slides show experiments on synthetic data and real S&P 500 market data (Javaheri et al. 2025).

# Numerical Experiments: Synthetic Data


Recovered edge weights over time (darker = stronger edge):




► *Proposed method recovers the ground-truth edge weight evolution most faithfully.*

# Numerical Experiments: $k$ -Component Graphs (S&P 500)


100 S&P 500 stocks partitioned into  $k = 8$  sectors (last frame):




CLR (ACC 0.56)




SGLA (ACC 0.27)



Fingraph (ACC 0.49)



Javaheri et al. (ACC 0.61)



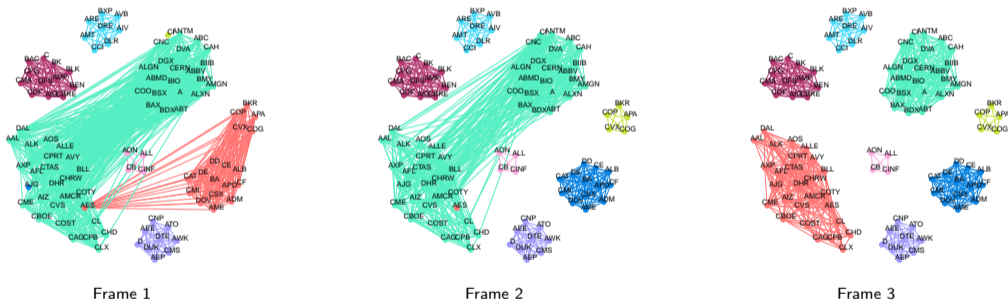
Proposed (ACC 0.69)

- Consumer Staples
- Energy
- Financials
- Health Care
- Industrials
- Materials
- Real Estate
- Utilities

► *Proposed method cleanly separates sectors with fewer misclassified stocks.*

# Numerical Experiments: Time Evolution of Graphs

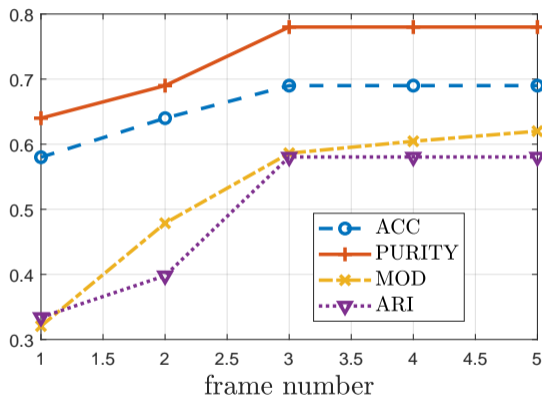
Graph structure evolves across three consecutive frames (S&P 500):



► *Cluster structure remains stable while edge strengths adapt to changing markets.*

# Numerical Experiments: Clustering Accuracy Over Time

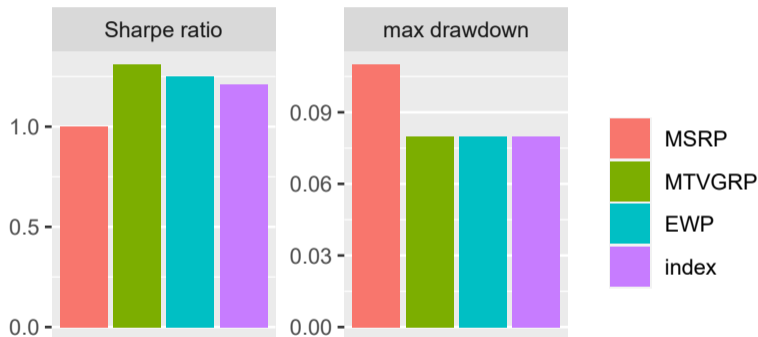
Clustering performance improves as more frames are observed (S&P 500,  $k = 8$ ):



► *Time-varying model accumulates information across frames, steadily improving accuracy.*

# Numerical Experiments: Portfolio Performance

Portfolio backtesting on S&P 500 (Sharpe ratio and maximum drawdown):



► *Time-varying graph portfolio (MTVGRP) achieves higher Sharpe ratio (1.31 vs. 1.00) and lower drawdown than benchmarks.*

# Outline

- 1 Graphs
- 2 Graph Learning 101
- 3 Polynomial Graphical Lasso
- 4 Financial Data
- 5 Learning Heavy-Tailed Graphs
- 6 Learning Structured Graphs
- 7 Learning Time-Varying Graphs
- 8 Summary**

- **Financial data is non-Gaussian and time-varying:** heavy tails, volatility clustering, and correlation structure all call for specialized graph models
- **Graph learning from smooth signals / GMRF** (Sections 2–3): a rich family of formulations — sparse, structured, and polynomial — for static graphs
- **Polynomial Graphical Lasso** (Section 3): exploits graph stationarity ( $\Sigma \mathbf{S} = \mathbf{S} \Sigma$ ) to jointly learn the precision matrix and the graph, capturing multi-hop interactions (Buciulea et al. 2025)
- **Heavy-tailed and structured graphs** (Sections 5–6): Student- $t$  model via MM yields robust estimation; spectral constraints enforce  $k$ -component and bipartite structure (M. Cardoso, Ying, and Palomar 2021, 2022; Kumar et al. 2020)
- **Time-varying graphs** (Section 7): non-negative VAR(1) temporal prior + Student- $t$  likelihood tracks market dynamics, detects crises, and improves portfolio Sharpe ratio (Javaheri et al. 2025)

# Thanks

For more information visit:

<https://portfoliooptimizationbook.com>

<https://www.danielpalomar.com>

<https://github.com/convexfi>

<https://www.youtube.com/danielpalomar>



- Buciulea, Andrei, Jiayi Ying, Antonio G. Marques, and Daniel P. Palomar. 2025. "Polynomial Graphical Lasso: Learning Edges from Gaussian Graph-Stationary Signals." *IEEE Transactions on Signal Processing* 73: 1153–67. <https://doi.org/10.1109/TSP.2025.3544376>.
- Cont, R. 2001. "Empirical Properties of Asset Returns: Stylised Facts and Statistical Issues." *Quantitative Finance* 1: 223–36.
- Dong, X., D. Thanou, P. Frossard, and P. Vandergheynst. 2015. "Laplacian Matrix Learning for Smooth Graph Signal Representation." In *Proc. IEEE Int. Conf. Acoustics, Speech, and Signal Processing (ICASSP)*. Brisbane, Australia.
- Dong, X., D. Thanou, M. Rabbat, and P. Frossard. 2019. "Learning Graphs from Data: A Signal Representation Perspective." *IEEE Signal Processing Magazine* 36 (3): 44–63.

- Egilmez, H. E., E. Pavez, and A. Ortega. 2017. “Graph Learning from Data Under Laplacian and Structural Constraints.” *IEEE Journal of Selected Topics in Signal Processing* 11 (6): 825–41.
- Javaheri, Amirhossein, and Daniel P. Palomar. 2025. “Clustering of Incomplete Data via a Bipartite Graph Structure.” In *Proc. European Signal Processing Conference (EUSIPCO)*. Palermo, Italy.
- Javaheri, Amirhossein, Jiayi Ying, Daniel P. Palomar, and Farokh Marvasti. 2025. “Time-Varying Graph Learning for Data with Heavy-Tailed Distribution.” *IEEE Transactions on Signal Processing* 73: 3044–60.  
<https://doi.org/10.1109/TSP.2025.3588173>.
- Kalofolias, V. 2016. “How to Learn a Graph from Smooth Signals.” In *Proc. Int. Conf. Artif. Intell. Statist.*, 920–29.

- Kumar, S., J. Ying, J. V. de M. Cardoso, and D. P. Palomar. 2019. “Structured Graph Learning via Laplacian Spectral Constraints.” In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*. Vancouver, Canada.
- . 2020. “A Unified Framework for Structured Graph Learning via Spectral Constraints.” *Journal of Machine Learning Research*, 1–60.
- Lake, B., and J. Tenenbaum. 2010. “Discovering Structure by Learning Sparse Graphs.” In *Proc. 33rd Annual Cognitive Science Conference*.
- M. Cardoso, J. V. de, J. Ying, and D. P. Palomar. 2021. “Graphical Models for Heavy-Tailed Markets.” In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*.
- . 2022. “Learning Bipartite Graphs: Heavy Tails and Multiple Components.” In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*. New Orleans, LA, USA.
- Mateos, G., S. Segarra, A. G. Marques, and A. Ribeiro. 2019. “Connecting the Dots.” *IEEE Signal Processing Magazine* 36 (3): 16–43.

- Meinshausen, N., and P. Bühlmann. 2006. “High-Dimensional Graphs and Variable Selection with the Lasso.” *The Annals of Statistics* 34 (3): 1436–62.
- Nie, F., X. Wang, M. Jordan, and H. Huang. 2016. “The Constrained Laplacian Rank Algorithm for Graph-Based Clustering.” In *Proc. AAAI Conf. On Artificial Intelligence*, 1969–76.
- Palomar, Daniel P. 2025. *Portfolio Optimization: Theory and Application*. Cambridge University Press.
- Ruppert, D. 2010. *Statistics and Data Analysis for Financial Engineering*. Springer.
- Sun, Y., P. Babu, and D. P. Palomar. 2017. “Majorization-Minimization Algorithms in Signal Processing, Communications, and Machine Learning.” *IEEE Trans. Signal Processing* 65 (3): 794–816.

- Ying, J., J. V. de M. Cardoso, and D. P. Palomar. 2020. “Nonconvex Sparse Graph Learning Under Laplacian Constrained Graphical Model.” In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*.
- . 2021. “Minimax Estimation of Laplacian Constrained Precision Matrices.” In *Proc. Int. Conf. Artificial Intelligence and Statistics (AISTATS)*.
- Zhao, L., Y. Wang, S. Kumar, and D. P. Palomar. 2019. “Optimization Algorithms for Graph Laplacian Estimation via ADMM and MM.” *IEEE Trans. Signal Processing* 67 (16): 4231–44.