

Exercises

Portfolio Optimization: Theory and Application Chapter 3 – Financial Data: I.I.D. Modeling

Daniel P. Palomar (2025). *Portfolio Optimization: Theory and Application*.
Cambridge University Press.

portfoliooptimizationbook.com

Exercise 3.1: Unbiasedness and consistency of sample mean estimator

Consider a univariate Gaussian-distributed i.i.d. time series with mean 0.01 and variance 1, $x_t \sim \mathcal{N}(0.01, 1)$, $t = 1, \dots, T$.

- Generate data for $T = 10$ and compute the sample mean. Repeat the experiment multiple times and plot the histogram of the estimated mean value. Confirm that the expected value of the histogram coincides with the true mean value.
- Now repeat the experiment with $T = 20$ observations and compare the histograms (also compute the standard deviation of each histogram).
- Finally, repeat the experiment multiple times, for different numbers of observations $T = 10, 20, \dots, 100$, and plot the mean squared error of the estimation as a function of T .

Exercise 3.2: Bias of sample covariance matrix

Suppose we have T i.i.d. N -dimensional observations $\mathbf{x}_1, \dots, \mathbf{x}_T$ distributed as $\mathbf{x}_t \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$.

- Derive the following expected value based on the true $\boldsymbol{\mu}$:

$$\mathbb{E} \left[\sum_{t=1}^T (\mathbf{x}_t - \boldsymbol{\mu})(\mathbf{x}_t - \boldsymbol{\mu})^\top \right].$$

- Derive the following expected value based now on the sample mean $\hat{\boldsymbol{\mu}} = \frac{1}{T} \sum_{t=1}^T \mathbf{x}_t$:

$$\mathbb{E} \left[\sum_{t=1}^T (\mathbf{x}_t - \hat{\boldsymbol{\mu}})(\mathbf{x}_t - \hat{\boldsymbol{\mu}})^\top \right].$$

- c. Discuss the appropriate normalization factor, $1/(T - 1)$ or $1/T$, to be used in the expression of the sample covariance matrix.

Exercise 3.3: Location estimators

Consider a two-dimensional ($N = 2$) Gaussian-distributed i.i.d. time series with zero mean and identity covariance matrix, $\mathbf{x}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $t = 1, \dots, T$.

- Generate data for $T = 20$ and estimate the mean vector $\boldsymbol{\mu}$ via the sample mean, the median, and the spatial median. Visualize the results in a scatter plot.
- Repeat the experiment multiple times, for different numbers of observations $T = 10, 20, \dots, 100$, and plot the mean squared error as a function of T .

Exercise 3.4: Location estimators with outliers

Consider a two-dimensional ($N = 2$) Gaussian-distributed i.i.d. time series with zero mean and identity covariance matrix, $\mathbf{x}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $t = 1, \dots, T$.

- Generate data for $T = 20$ and estimate the mean vector $\boldsymbol{\mu}$ via the sample mean, the median, and the spatial median. Visualize the results in a scatter plot. Repeat the experiment multiple times and compute the mean squared error of the estimators.
- Then, add some small percentage of outliers in the observations, for example, distributed as $\mathbf{x}_t \sim \mathcal{N}(0.1 \times \mathbf{1}, \mathbf{I})$, and compute again the mean squared error of the estimators.
- Finally, repeat the experiment multiple times and plot the estimation error as a function of the percentage of outliers. Observe the robustness of the three estimators against outliers and discuss.

Exercise 3.5: Derivation of sample mean as location estimator

Given the observations \mathbf{x}_t , $t = 1, \dots, T$, the sample mean can be derived as the solution to the following optimization problem:

$$\underset{\boldsymbol{\mu}}{\text{minimize}} \quad \sum_{t=1}^T \|\mathbf{x}_t - \boldsymbol{\mu}\|_2^2.$$

- Is this problem convex? What class of optimization problem is it?
- Derive the solution in closed form by setting the gradient with respect to $\boldsymbol{\mu}$ to zero.

Exercise 3.6: Computation of spatial median as location estimator

Given the observations \mathbf{x}_t , $t = 1, \dots, T$, the spatial median can be derived as the solution to the

following optimization problem:

$$\underset{\boldsymbol{\mu}}{\text{minimize}} \quad \sum_{t=1}^T \|\mathbf{x}_t - \boldsymbol{\mu}\|_2.$$

- Is this problem convex? What class of optimization problem is it?
- Can a closed-form solution be obtained as in the case of the sample mean?
- Develop an iterative algorithm to compute the spatial median by solving a sequence of weighted sample means. Hint: find a majorizer of the ℓ_2 -norm in the form of a squared ℓ_2 -norm and then employ the majorization–minimization framework.

Exercise 3.7: ML estimation of covariance matrix

Consider an N -dimensional i.i.d. time series with zero mean and identity covariance matrix, $\mathbf{x}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $t = 1, \dots, T$.

- Generate Gaussian data for $N = 10$ and $T = 50$ and estimate the covariance matrix $\boldsymbol{\Sigma}$ via the Gaussian ML estimator and the heavy-tailed ML estimator. Run the experiment multiple times and compute the mean squared error of the estimators.
- Now repeat the whole experiment but generating instead heavy-tailed data (e.g., following a t distribution) with the same mean and covariance matrix. Observe the robustness of the two estimators against heavy tails and discuss.

Exercise 3.8: Derivation of Gaussian ML estimators

Given T N -dimensional observations $\mathbf{x}_1, \dots, \mathbf{x}_T$, the Gaussian ML estimation for $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ is formulated as

$$\underset{\boldsymbol{\mu}, \boldsymbol{\Sigma}}{\text{minimize}} \quad \log \det(\boldsymbol{\Sigma}) + \frac{1}{T} \sum_{t=1}^T (\mathbf{x}_t - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1} (\mathbf{x}_t - \boldsymbol{\mu}).$$

Derive the estimators by setting the gradient of the objective function with respect to $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}^{-1}$ to zero.

Exercise 3.9: Derivation of heavy-tailed ML estimators

Given T N -dimensional observations $\mathbf{x}_1, \dots, \mathbf{x}_T$, the heavy-tailed ML estimation (under the t distribution with degrees of freedom ν) for $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ is formulated as

$$\underset{\boldsymbol{\mu}, \boldsymbol{\Sigma}}{\text{minimize}} \quad \log \det(\boldsymbol{\Sigma}) + \frac{\nu + N}{T} \sum_{t=1}^T \log \left(1 + \frac{1}{\nu} (\mathbf{x}_t - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1} (\mathbf{x}_t - \boldsymbol{\mu}) \right).$$

Derive the fixed-point equations characterizing the estimators by setting the gradient of the objective function with respect to $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}^{-1}$ to zero.

Exercise 3.10: Shrinkage James–Stein estimator for the sample mean

Consider a Gaussian-distributed i.i.d. N -dimensional time series with zero mean and identity covariance matrix, $\mathbf{x}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $t = 1, \dots, T$.

- Generate data for $N = 10$ and $T = 20$, and estimate the mean vector with the sample mean and with the shrinkage James–Stein estimator.
- Run the experiment multiple times and compute the mean squared error of the estimators.
- Finally, repeat the experiment multiple times, for different numbers of observations $T = 10, 20, \dots, 100$, and plot the mean squared error as a function of T .

Exercise 3.11: Shrinkage sample covariance matrix estimator

Consider a Gaussian-distributed i.i.d. N -dimensional time series with zero mean and identity covariance matrix, $\mathbf{x}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $t = 1, \dots, T$.

- Generate data for $N = 10$ and $T = 20$, and estimate the covariance matrix with the sample covariance matrix and with the shrinkage Ledoit–Wolf estimator.
- Run the experiment multiple times and compute the mean squared error of the estimators.
- Finally, repeat the experiment multiple times, for different numbers of observations $T = 10, 20, \dots, 100$, and plot the mean squared error as a function of T .

Exercise 3.12: Factor model estimator

Consider a Gaussian-distributed i.i.d. N -dimensional time series with zero mean and covariance matrix with a single-factor structure $\Sigma = \beta\beta^\top + \mathbf{I}$ (e.g., with $\beta = \mathbf{1}$), $\mathbf{x}_t \sim \mathcal{N}(\mathbf{0}, \Sigma)$, $t = 1, \dots, T$.

- Generate data for $N = 10$ and $T = 20$, and estimate the covariance matrix with the sample covariance matrix and with the single-factor model structure (e.g., with PCA).
- Run the experiment multiple times and compute the mean squared error of the estimators.
- Finally, repeat the experiment multiple times, for different numbers of observations $T = 10, 20, \dots, 100$, and plot the mean squared error as a function of T .