Elvin Isufi (TU Delft)

Relational Learning with Covariance Information

Abstract: Covariance information is commonly used in machine learning to reveal data interdependencies such as network topology inference (e.g., graphical lasso) and dimensionality reduction (e.g., Principal Component Analysis (PCA)). However, such information is often only the first step in a machine learning pipeline that is performed separately from the task. Because of finite-data estimation errors, we end up working with a sample covariance matrix that leads to uncertainties in its spectrum. For example, PCA is notoriously unstable to covariance estimation errors, i.e., small data perturbations might lead to large changes in principal directions. To address this, coVariance Neural Networks (VNNs) were introduced. These networks perform graph convolutions on the sample covariance matrix, an operation that, similarly to PCA, modulates the data principal components, but with enhanced representation power and greater stability against covariance estimation errors. However, in sparse, high-dimensional settings with limited data, covariance estimation is particularly difficult, which hinders VNNs' performance despite their stability. Sparse VNNs overcome this by using theoretically grounded covariance sparsification, which improves stability, reduces the impact of spurious correlations on performance and improves computation and memory efficiency. The success of VNNs motivates their extension to different settings. SpatioTemporal VNNs, for instance, process multivariate time series by applying graph convolutions on the online estimated covariance and temporal convolutions over time, achieving stability to estimation errors in both covariance and model parameters due to streaming data. Finally, VNNs' stability promotes fairness in datasets with poorly represented groups. Building on this, Fair VNNs leverage equitable covariance estimates and fairness penalties in the loss function to ensure a more balanced treatment of these groups.