A067 · Divide and Conquer: Scalable Partial Correlation Network Inference for High-Dimensional Omics Data

Presenter: Luke Berger — Boston University

Authors: Luke Berger, Stephano Monti

Session: Day 1 · Thu Oct 1, 2026 · 2:15–4:15 PM

Graphical models are widely used in systems biology to study the complex regulatory relationships underlying biological processes. The development of methods such as SILGGM has allowed for the inference of partial correlation networks, which improve interpretability by controlling for indirect interactions. Unfortunately, the application of these tools to high-throughput omics datasets is constrained by: 1) decreased stability and accuracy as the feature-to-sample (p/n) ratio increases, and 2) computational costs that scale rapidly with p.

These challenges can be mitigated by exploiting the fact that biological networks are typically sparse and modular in structure. SHINE exploits modular network structure to reduce graphical search complexity, but implements only Bayesian inference and is specifically tailored to hierarchical constraint learning. We developed a more generalized “Divide And Conquer” (DAC) approach that separates features into overlapping modules based on marginal correlations, performs partial correlation inference within modules, and constructs a network from the modular subgraphs. This allows for parallelization of subgraph inference and limits the dimensionality of each inference subproblem. We benchmarked DAC against a range of alternative algorithms and found that it substantially reduced runtime with a minimal loss in F1 score.

We created modularDAC: an R package implementing the DAC algorithm with customizable options for module detection and subgraph inference. By allowing for the rapid and accurate inference of partial correlation networks with thousands of features, modularDAC removes a bottleneck in systems biology workflows and enables network-based analysis of omics data at previously intractable scales.