A141 · Influence Causal Ordering: Scalable Causal Structure from Genome-Scale Perturbation Screens

Presenter: Ritwik Anand — Northeastern University

Authors: Ritwik Anand, Olga Vitek, Karen Sachs

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

Genome-scale perturbation screens, such as single-cell CRISPR-knockdown atlases and drug panels, now measure the effects of thousands of interventions, yet resolving their causal structure at this scale remains hard: combinatorial DAG search does not scale, pervasive biological feedback breaks acyclicity assumptions, and a dominant shared stress response swamps intervention-specific signal. We present Influence Causal Ordering (ICO), which converts a perturbation influence matrix into a hierarchical causal cluster graph, scaling to atlases spanning thousands of perturbations. In this graph, feedback loops and shared response programs are represented as clusters, and the clusters are arranged in causal order. ICO thus captures in one structure both the program-level responses shared across perturbations and each perturbation's specific effects, where typical causal methods don't scale nor accurately represent feedback cycles. Empirically, on single-cell knockdown atlases ICO recovers edge orientation up to the identifiability ceiling of the data and predicts the effects of held-out, never-perturbed knockdowns beyond a strong mean-response baseline; on dose-resolved drug phospho-proteomics, its causal cones predict held-out drug effects better than curated pathway priors. By representing cycles and programs as clusters and ordering them causally, ICO turns a large perturbation screen into a single interpretable causal map that accurately models both shared downstream responses and specific effects, and is capable of modeling cyclic feedback.