Press ECCS to Doubt (Your Causal Graph)

Markos Markakis, Ziyu Zhang, Rana Shahout, Trinity Gao, Chunwei Liu, Ibrahim Sabek, Michael Cafarella

A Human-In-The-Loop Can Verify Causal Graphs Selectively

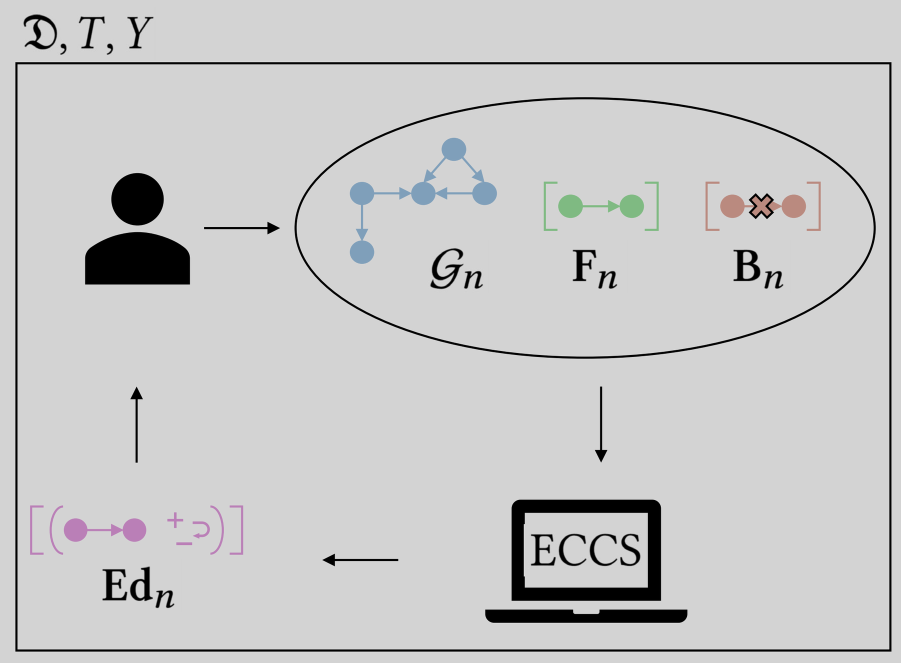

High-Quality Causal Graphs Are Hard to Obtain

- Causality has helped scientists across domains pose, discuss and test hypotheses.
- Pearl's framework [1] requires a causal graph for the problem.
- For large problem instances, hand-crafting a causal graph is prohibitive.
- But mining it automatically [2,3] from the data can also be error-and bias-prone.
- Comprehensive verification of the automated result is still needed!

Human Feedback Must Be Utilized Efficiently

- Our system **ECCS** helps with **E**xposing **C**ritical **C**ausal **S**tructures interactively.
- Users have a dataset D and a certain ATE question in mind: the effect of T on Y.
- In each Interaction Round:
  - Inputs: Causal graph, list of FIXED edges, list of BANNED edges
  - Outputs: Suggested causal graph edits with high impact on the ATE of interest.
- The user can revise the graph, fixed list and banned list and invoke ECCS again.

We Propose 3 Interactive Causal Graph Verification Strategies

Strategy 1: SingleEdit

1. Enumerate **graphs that are "one edge edit away"**:
   - For each edge in the starting graph, remove it if not FIXED.
   - For each edge in the starting graph, flip it if not FIXED and its inverse is not BANNED.
   - For each edge not in the starting graph, add it if not BANNED.
2. Find the one that **maximally affects the ATE of interest**.
3. Suggest the **corresponding single-edge edit** to the user.

Strategy 2: HeuristicEdit

1. **Create a search graph** where:
   - Nodes are causal graphs
   - Edges are causal edge operations
2. **Run A* search** [4] to find a frontier of graphs with maximally different ATE of interest.
3. Find **edits common to the paths to the "best" points** on the frontier.

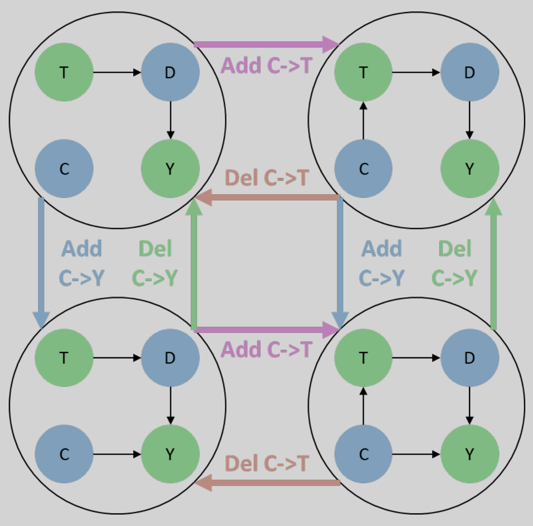

Strategy 3: AdjSetEdit

1. Derive an **adjustment set** implied by the starting causal graph.
2. **Toggle the inclusion status** of each variable in the adjustment set.
3. For each new adjustment set derived this way, **map it to a set of required graph edits.**
4. Return the set of edits that **maximizes the difference in the ATE of interest**.

Definition 1 (Back-Door Criterion). *A set $\mathbf{Z}$ satisfies the* ***backdoor criterion*** *relative to variables (T, Y) in a DAG $\mathcal{G}$ if:*

(i) $\mathbf{Z} \cap \mathbf{De}(T) = \emptyset$*; and*

(ii) $\mathbf{Z}$ *blocks every T-backdoor path between T and Y.*

In this case, $\mathbf{Z}$ is a sufficient adjustment set for correctly estimating ATE(T, Y). [1]

Our AdjSetEdit Strategy Offers Good Performance with Interactive Latency

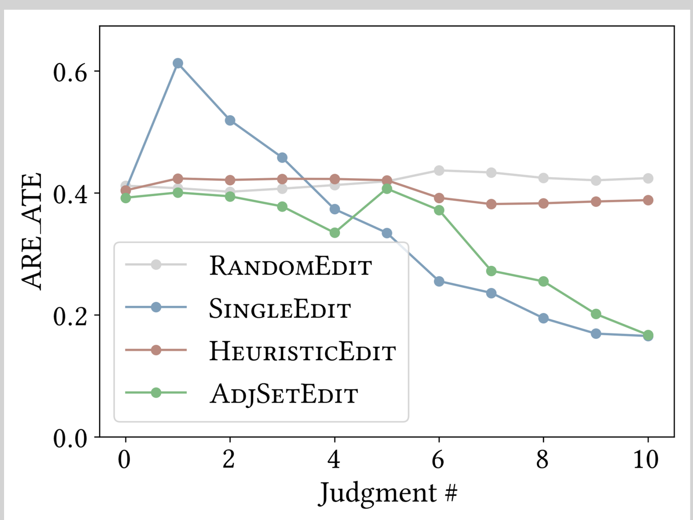

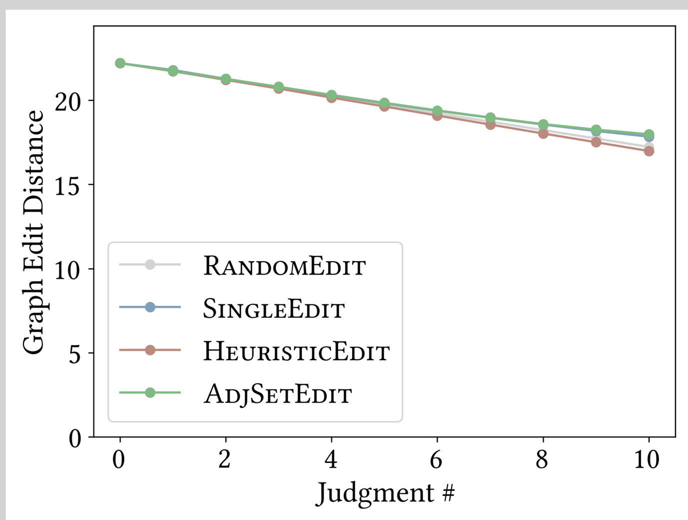

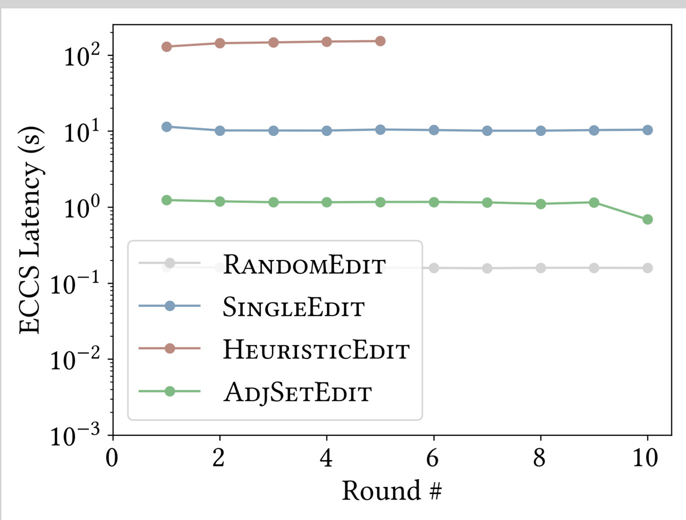

We Evaluate On Randomized Graphs

Ground Truth Graphs	GenDAG(10, 0.5, −10, 10, 0.001, 2)	10 runs
Datasets	GenData($\mathcal{G}$, 1000, −10, 10)	3 runs/graph
Starting Graphs	GenDAG(10, 0.5, −10, 10, 0.001, 2)	10 runs
Choices of T/Y	Combinations($\mathcal{G}.nodes$, 2)	$\binom{10}{2} = 45$
Strategies	RandomEdit SingleEdit HeuristicEdit AdjSetEdit	3 runs 1 run 1 run 1 run
Total	$\mathbf{10 \times 3 \times 10 \times 45 \times (1 + 1 + 1 + 3) = 81{,}000}$	

AdjSetEdit Emerges as Dominant

- **ARE_ATE:** Absolute Relative Error in the ATE of interest

$$ARE\_ATE = \left| \frac{ATE_{current} - ATE_{ground\_truth}}{ATE_{ground\_truth}} \right|$$

  - 0.17 after 10 judgements, -60.60% vs RandomEdit, similar to SingleEdit
- **Graph Edit Distance:** Closeness to ground truth graph after each judgment.
  - Similar to SingleEdit after 10 judgements, an average of 17.99.
- **Latency:** Time taken by ECCS to produce its suggestions in each round.
  - Each call takes 1.12s on average and can produce multiple suggested edits.

[1] Judea Pearl. 2009. Causality: Models, Reasoning and Inference. Cambridge University Press.
[2] Clark Glymour, Kun Zhang, and Peter Spirtes. 2019. Review of causal discovery methods based on graphical models. Frontiers in genetics 10 (2019), 524
[3] Peter Spirtes, Clark N Glymour, and Richard Scheines. 2000. Causation, prediction, and search. MIT press
[4] Peter E Hart, Nils J Nilsson, and Bertram Raphael. 1968. A formal basis for the heuristic determination of minimum cost paths. IEEE transactions on Systems Science and Cybernetics 4, 2 (1968), 100–107.