Causal Consistency Evaluation of Synthetic Features in Structural Causal Models

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 12 peer-reviewed papers addressing the following research question: Can the causal consistency of synthetic features generated by Structural Causal Models (SCMs) be quantitatively evaluated using causal benchmarks like FCI or PC algorithms, and how does this impact. 14 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 2.5/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Quantitative probing: Validating causal models using quantitative domain knowledge. Research question: Can the causal consistency of synthetic features generated by Structural Causal Models (SCMs) be quantitatively evaluated using causal benchmarks like FCI or PC algorithms, and how does this impact the robustness of TFMs against covariate shift?.

2 Methodology

Systematic literature search across multiple databases yielded 12 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 2.5/10.

3 Results

12 papers retrieved. 14 claims extracted; 0 independently verified. Quality review score: 2.5/10.

4 Limitations

This report is a machine-generated literature synthesis and does not constitute original research. Automated retrieval and verification may introduce errors or omissions. Review scores reflect automated assessment, not human peer review. Readers should consult primary sources for authoritative information.

5 Extracted Claims

Claim	Verified	Confidence
The aggregated results are based on 1378 runs.	×	0.03
In Figure 2, the x-coordinate indicates the hit rate of the run.	×	0.00
The upper left plot in Figure 2 shows the absolute difference between the estimated and the true value of the target eff	×	0.02
The upper right plot in Figure 2 shows the relative difference |τ-τ / τ| between the estimated value and the true value	×	0.01
All target effects in the study were chosen to be nontrivial to avoid division by zero.	×	0.03
The lower left plot in Figure 2 shows the structural hamming distance between the true and the discovered graph.	×	0.01
The structural hamming distance calculation includes edges present in only one graph as well as reversed edges.	×	0.03
The lower right plot in Figure 2 shows the absolute frequencies of the hit rates over all runs.	×	0.00
Almost all of the runs show a hit rate of at least 0.8.	×	0.00
Many of the randomly chosen treatment-outcome-pairs in the probes are not connected by a directed path in the true graph	×	0.01
Figure 3 replaces single data points for each hit rate with one data point representing the mean value of the plotted qu	×	0.02
The hit rate histogram in Figure 3 is unchanged from Figure 2 because the plots are derived from the same runs.	×	0.00
In Figure 3, the plots show a downward trend for regions containing a sufficient number of samples.	×	0.04
Runs with a higher hit rate performed better at estimating the target effect and recovering the causal graph.	×	0.03

References

- `http://arxiv.org/abs/2209.03013v1`

- http://arxiv.org/abs/2206.02435v2
- http://arxiv.org/abs/2412.12401v1