

A Fundamental Inconsistency Between Equilibrium Causal Discovery and Causal Reasoning Formalisms.

†DENVER DASH AND ‡MAREK DRUZDZEL

Decision Systems Laboratory^{†‡}, *Department of Physics and Astronomy*[†], *Department of Information Science and Telecommunications*[‡], and *Intelligent Systems Program*[‡],
University of Pittsburgh^{†‡}

Causal discovery, for the most part, is concerned with learning causal models in the form of directed acyclic graphs (DAGs) from equilibrium (as opposed to time series) data. Causal reasoning, by contrast, is concerned with using such causal DAGs to perform inferences. In particular, much work on causal reasoning has focused on the ability to predict the new probability distribution over a set of variables, V , given a causal graph $G = (V, E)$ and given the fact that some subset of variables $V' \subset V$ has been externally *manipulated* to some configuration. These types of *manipulation inferences* contrast with more common *diagnostic inferences*, in that the former may require the causal graph to be altered prior to performing probabilistic inference. Specifically, the ability to perform manipulation inferences is made possible by a critical postulate which we call the *Manipulation Postulate*. All formalisms for causal reasoning take the manipulation postulate as a fundamental starting point:

The Manipulation Postulate. *If $G = (V, E)$ is a causal graph and $V' \subset V$ is a subset of variables being manipulated, then the causal graph, G' , describing the manipulated system is such that $G' = (V, E')$, where $E' \subseteq E$ and E' differs from E by at most the set of arcs into V' .*

In other words, manipulating a variable can cause some of its incoming arcs to be removed from the causal graph, but can effect no other change in the causal graph. The Manipulation Theorem of Spirtes, *et. al.* (1992) proves that given the Manipulation Postulate and the Markov Condition, the probability distribution of the manipulated model can be calculated. Furthermore, the axiomatizations of causal reasoning of Galles and Pearl (1997) and of Halpern (1998) also take the Manipulation Postulate as a fundamental assumption.

The question that we pose in this paper is “Are these two lines of research (i.e., equilibrium causal discovery and manipulation reasoning) consistent?” Namely, what would happen if we took an equilibrium causal model (learned from data), and applied the manipulation formalisms to it? Are the resulting inferences guaranteed to be valid? We prove by explicit counterexample that such inferences are not guaranteed to be valid in the sense that conditional independencies in the manipulated model can differ from the conditional independencies in the learned model of the manipulated system. Symbolically, if M_S is a learned causal model of system S , and if we use the $\hat{\cdot}$ operator to denote manipulation, then we show that $\hat{M}_S \neq M_S$.

Our general strategy is as follows: We first present two extremely simple physical systems (an ideal gas trapped in a cylinder with a movable piston and a mass dangling from a damped spring), we show, based on physical laws what the “true” equilibrium causal graphs of these systems are. We further show that with an appropriate source of noise present in data taken from these systems, a constraint-based learning algorithm will learn the correct causal graphs. Finally, we show that the graph predicted by manipulation-type reasoning on these learned models will possess different conditional independence relations than the causal graph that would be learned from the true manipulated system. Furthermore, we will

show that under suitable manipulations, these systems will display dynamic instabilities; a phenomenon which is completely unaccounted for in any existing treatment of manipulation.

This inconsistency, i.e., the fact that a learned-then-manipulated causal model is not equal to the manipulated-then-learned model, is attributed to an inappropriate use of the Manipulation Postulate in manipulation formalisms. We review the work of Iwasaki and Simon (1994), which deals with representing causality in time-dependent systems based on structural equation models combined with differential equation systems. They show that physical systems possessing stable fixed points may possess multiple causal graphs depending on the time-scale being modeled.

We show that the Manipulation Postulate applied to Iwasaki–Simon-type graphs for our two paradoxical systems, *modeled on an infinitesimal time-scale* (graphs which we refer to as “differential causal graphs”), produce equilibrium causal graphs with the correct independence relations. Furthermore, we show how these differential causal models correctly predict the presence of instabilities under manipulations of the system. We conclude that the Manipulation Postulate, and thus all existing manipulation formalisms, are only guaranteed to be valid on differential causal models.

This work generates several new open questions and areas of research. These questions will be presented and possible solutions will be explored.

REFERENCES

- Galles, D. and Pearl, Judea 1991. Axioms of causal relevance. *Artificial Intelligence* 97(1–2):9–43.
- Halpern, Joseph Y. 1998. Axiomatizing causal reasoning. In *Proceedings of the Fourteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI–98)*, San Francisco, CA. Morgan Kaufmann Publishers. 202–210.
- Iwasaki, Yumi and Simon, Herbert A. 1994. Causality and model abstraction. *Artificial Intelligence* 67(1):143–194.
- Spirtes, Peter; Glymour, Clark; and Scheines, Richard 1993. *Causation, Prediction, and Search*. Springer Verlag, New York.

[School of Information Sciences, 740 SIS Building, 135 North Bellefield Avenue, Pittsburgh, PA 15260. ddash@sis.pitt.edu, marek@sis.pitt.edu]