How to Turn an MAS into a Graphical Causal Model

JAAMAS Track

H. Van Dyke Parunak
Parallax Advanced Research
Beavercreek, Ohio, USA
van.parunak@gmail.com

ABSTRACT

This paper shows that an appropriately configured multi-agent system (MAS) is formally equivalent to a graphical causal model (GCM, a broad category that includes many formalisms expressed as directed graphs), and offers benefits over other GCMs in modeling a social scenario. MASs often *use* GCMs to support their operation, but are not usually viewed as tools for *enhancing* their execution. We argue that the definition of a GCM should include its *update mechanism*, an often-overlooked component. We review a wide range of GCMs to validate this definition and point out limitations that they face when applied to the social and psychological dimensions of causality. Then we describe SCAMP (Social Causality using Agents with Multiple Perspectives), a causal language and multi-agent simulator that satisfies our definition and overcomes the limitations of other GCMs for social simulation.

KEYWORDS

Stigmergy, Causal Modeling, Agent-Based Modeling

ACM Reference Format:
H. Van Dyke Parunak. 2023. How to Turn an MAS into a Graphical Causal Model: JAAMAS Track. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 3 pages.

1 INTRODUCTION

This paper [2] makes two claims.

1. A stigmergic multi-agent system (MAS) with an appropriate environment has the same mathematical structure (defined in Section 2) as a graphical causal model (GCM).
2. Such an MAS has advantages over other GCMs for modeling social causality.

"Graphical causal models" (GCMs) include reasoning systems (e.g., Bayes nets, POMDPs, fuzzy cognitive maps, causal loop diagrams), based on directed graphs. These models have a common structure (Section 2). Multi-agent systems (MASs) often *use* GCMs, but are not usually viewed as *instances* of such models. Claim 1 asserts that a stigmergic MAS can share this structure in a way that other MASs do not. This insight allows the agent community to contribute in a new way to the causal modeling community.

Once we recognize an appropriately configured MAS as a GCM, we can compare it with other GCMs. Claim 2 arises from our experience with such an MAS, SCAMP (Social Causality using Agents with Multiple Perspectives), for a major experiment in social science as part of the DARPA Ground Truth program. SCAMP generated data from a synthetic society with known underlying causal structure [4]. Our society was inspired by civil strife in Syria, and agents belonged to distinctive *groups*, including the government, an armed opposition, radical insurgents, non-governmental organizations, and civilians. We captured the behaviors of these agents as they interacted, and teams of social scientists used this data to evaluate methods of causal discovery. Section 3 aligns SCAMP (described elsewhere [2, 4, 5]) with our definition to validate Claim 1, and shows how it supports Claim 2. We are not concerned with any specific context for a GCM, but rather highlight a potential application of agent-based reasoning to graphical representations of causality in general.

SCAMP implements Simon's Law [7]: the complexities of behavior, human as well as insect, can be explained as simple agent behaviors constrained by a complex environment. It uses stigmergic agents to explore a variety of graphical environments that encode psychological and social behavior. Two of these environments are directed graphs that satisfy our definition for the structural component of a GCM (Section 2): a Causal Event Graph (CEG) whose nodes are event types and whose edges show causal relations among them, and a hierarchical graph network (HGN) whose nodes are (sub)goals and whose edges show how they combine and relate to events. As we explain in more detail in Section 3,

- Every GCM includes not only a directed *graph*, but also a *process* that updates values on the nodes of the graph.
- If a graph is the environment in a stigmergic system, agents can update node values as they move over it.
- Such an architecture has advantages over other GCMs for modeling social causality.

2 DEFINITION OF A GCM

Humans often represent causality as a directed graph. Philosophers struggle to define causality [1] (in graph theoretic terms, the semantics of the directed edges). The formalisms we discuss sidestep this question. For example, Pearl refuses to define causality, instead treating "cause" as an undefined primitive, like "point" and "line" in Euclidean geometry [6, pp. 27, 48]. We adopt this position. If an approach presents causal information as a directed graph, we understand a directed edge as a causal claim, with a cause at the tail and an effect at the head, without quibbling over the precise nature of the edge's causality.

A graphical causal model (GCM) requires not only a digraph with values associated with the nodes, but also an *updating mechanism* that updates the values on nodes. This mechanism is a significant component of the causal semantics of the formalism. Formally, we

---

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

define a graphical causal model

$$C \equiv \langle N, E, V, F, U \rangle \tag{1}$$

where the components are

- a set of nodes N;
- a set of directed edges between nodes $E \subset N \times N$;
- a set of values V (possibly vectors) that nodes can carry;
- a function $F : N \rightarrow V$ ($F_k(n)$ indicates the kth element of the vector v associated with node n);
- an updating function $U : F \rightarrow F'$ that changes the values to which individual nodes evaluate.

U often uses information associated with the edges to propagate causal effects from one node to another. Formalisms intended for human inspection and not computation may lack an explicit V, F, and U and rely on informal, qualitative node updating, but do not evade this definition. The human users of such formalisms have an informal sense of the prominence of each node, and interpret the graph according to conventions that function as U.

In most methods with computational updating, U is *analytic*, and involves solving an equation. In an appropriately configured MAS, U, executed by agents, can be *algorithmic*, offering a much richer space of options.

Most GCMs focus on the directed graph $\mathbb{G} \equiv < N, E >$. In the Ground Truth program, the "causality" that the social science teams were challenged to recover consisted only of such a graph, without V, F, or U. Our interaction with the social science teams showed us that the semantics of a GCM involves U, and thus the values V assigned to nodes, as much as the structure $\mathbb{G}$, and configuring an MAS as a GCM greatly enhances the potential power of U.

The full paper validates this definition by aligning it with eleven different kinds of GCMs (and numerous subtypes) that have been proposed, including factor trees, causal loop diagrams, path diagrams, causal diagrams, influence nets, the causal influence models, influence diagrams, POMDPs, fuzzy cognitive maps, system dynamics models and their underlying ODEs, and stochastic Petri nets. This review identifies four desirable features in a causal model. No previous formalism satisfies all four.

(1) Does the formalism estimate the relative *probability* of different nodes and pathways? Decision-makers want to focus on the most likely outcomes, as well as those that are intrinsically most interesting.
(2) Does the formalism support *cycles and feedback*? Feedback loops are pervasive in real systems, and are critical for understanding stability, instability, and emergent behavior.
(3) Does the formalism model the quantitative passage of *time*? Users want to know not just that one thing is likely to happen after another, but how long it will take.
(4) Does the formalism represent *agency*, expressing who is responsible for the various causal influences? Certain dimensions of causality, such as considering the goals of different groups, can only be captured if we know who is doing what.

-

3 DEMONSTRATING THE CLAIMS

SCAMP demonstrates both claims from Section 1. SCAMP has the same four features in Equation 1 that we found in other formalisms, demonstrating the first claim:

- N (Nodes in directed graph): Event types in the CEG; Subgoals in the HGN;
- E (Directed edges among nodes): agency and influence edges in the CEG; *and*, *or*, *zip* edges and their inverses in the HGN.
- V, F (Values on nodes): On the CEG, wellbeing, urgency, presence features; the groups that are eligible to participate in (have agency for) an event type; the nominal duration of the event. On the HGN: satisfaction, urgency, frustration, tolerance.
- U (Update function): pheromone deposits on the CEG; update of urgency through the HGN.

Our second claim asserts that an appropriately configured MAS supports the four requirements addressed variously by other GCMs: probability, cycles, time, and agency. SCAMP supports all four requirements.

Probability is supported in two ways. 1) Roulette-based decisions model the non-determinism of human choice. The probability of these transitions is not static and defined exogenously, but emerges dynamically from psychologically realistic modeling primitives (agent preferences and event features) that vary over time. 2) The presence features on each event type are deposited by polyagent ghosts [3] as they plan paths for their avatars, and avatars follow the crest of the presence field for their group. Thus the presence features on event type nodes are, up to a normalizing constant, the probability that agents of each group will participate in that event type. While each avatar follows only a single path, we log the presence features over time, allowing us to recover the time-varying probability of alternative futures.

Cycles and feedback are possible because CEG nodes represent event types, not specific events: they have nominal durations but not start times. Time-anchored events emerge from SCAMP's execution, as agents begin and end their participation in event types. Thus an agent can meaningfully revisit a node by participating in different events of the same type.

Time is based on the duration feature of a CEG node $F_3(n)$ from which an agent samples the duration of its participation in event type n. Agents execute in order of their individual time, so the temporal order of events is respected.

Agency is supported in three ways. 1) An agent can only choose an event type n for which its group has agency (that is, the group is a member of $F_2(n)$). 2) Agents belong to groups within which preferences are similar. 3) Each agent's wellbeing preferences (and an overall wellbeing variable) vary with its experiences, so that different agents with different histories encountering the same environmental state may behave differently.

ACKNOWLEDGMENTS

The development of SCAMP was funded by the Defense Advanced Research Projects Agency (DARPA), under Cooperative Agreement HR00111820003. The content of this paper does not necessarily reflect the position or the policy of the US Government, and no official endorsement should be inferred.

REFERENCES

[1] Stephen Mumford and Rani Lill Anjum. 2013. *Causation: A Very Short Introduction.* Oxford University Press, Oxford, UK.

[2] H. Van Dyke Parunak. 2022. How to Turn an MAS into a Graphical Causal Model. *Journal of Autonomous Agents and Multi-Agent Systems* 36 (2022). https://doi.org/10.1007/s10458-022-09560-y
[3] H. Van Dyke Parunak and Sven Brueckner. 2006. Concurrent Modeling of Alternative Worlds with Polyagents. In *the Seventh International Workshop on Multi-Agent-Based Simulation (MABS06, at AAMAS06)*. Springer, 128–141.
[4] H. Van Dyke Parunak, Jason Greanya, Peggy McCarthy, Jonathan A. Morell, Srikanth Nadella, and Laura Sappelsa. 2021. SCAMP's Stigmergic Model of Social Conflict. *Computational and Mathematical Organization Theory* (2021). https://doi.org/10.1007/s10588-021-09347-8
[5] H. Van Dyke Parunak, Jonathan A. Morell, Laura Sappelsa, and Jason Greanya. 2020. *SCAMP User Manual*. Report. Parallax Advanced Research. https://www.abcresearch.org/abc/papers/SCAMPUserManual.zip
[6] Judea Pearl and Dana Mackenzie. 2018. *The Book of Why*. Basic Books, New York, NY.
[7] H.A. Simon. 1969. *The Sciences of the Artificial*. MIT Press, Cambridge, MA.