Autonomous Agents and Multi-Agent Systems (2022) 36:31
https://doi.org/10.1007/s10458-022-09560-y

How to turn an MAS into a graphical causal model

H. Van Dyke Parunak[1]

Accepted: 16 April 2022

Abstract
This paper proposes 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 Social Causality using Agents with Multiple Perspectives (SCAMP), 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 · Social simulation

1 Introduction

This paper makes two claims.

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

"Graphical causal models" include reasoning systems, such as Bayes nets, POMDPs, fuzzy cognitive maps, Petri Nets, and causal loop diagrams, based on directed graphs. MASs often *use* GCMs, but are not usually viewed as *instances* of such models. We offer a definition of such models in Sect. 2, supported by an extensive review in Sect. 3 and Appendix 1.

Claim 1 asserts that a stigmergic MAS can satisfy this definition. This claim arises from our experience with such an MAS, Social Causality using Agents with Multiple Perspectives (SCAMP), for a major experiment in social science. Section 4 describes SCAMP, aligning it with our definition to validate Claim 1.

✉ H. Van Dyke Parunak
van.parunak@gmail.com

[1] Parallax Advanced Research, Beavercreek, OH, USA

Once we recognize an appropriately configured MAS as a GCM, we can compare it with other GCMs to substantiate Claim 2. This claim allows the agent community to contribute in a new way to the causal modeling community. Section 7.1 discusses this claim, as well as summarizing the evidence for Claim 1.

We[1] developed SCAMP as part of the DARPA Ground Truth program to generate data from synthetic societies with known underlying causal structure (as discussed in [50]). Teams of social scientists then used this data to evaluate methods of causal discovery.

SCAMP exploits a wide range of techniques that the author and previous colleagues[2] developed in stigmergic MASs, in which agents coordinate their activity by making and sensing changes in a shared environment. The environments in these systems included not only spatial lattices [49], where stigmergy is widely used in robotics, but also hierarchical task networks [46] and directed graphs of events [48]. The last of these satisfies our definition for the structural component of a GCM (Sect. 2). In developing this model we contrasted it with more conventional GCMs such as Bayes networks, but did not recognize the deeper points of similarity. Those insights, developed in this paper, became clear only as we integrated these techniques into a production-strength social simulator and presented its data to social scientists in response to causally-oriented questions in the DARPA Ground Truth program. Briefly, and as we explain in more detail in Section 4:

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

Section 5 presents some experimental results from the deployment of a full-scale model in the Ground Truth program. Section 6 reviews previous work at the intersection of GCMs and MASs. Section 7 summarizes the evidence for our two claims, discusses some implications of this work, and outlines future directions. Appendix 1 documents other GCMs in detail, and Appendix 2 summarizes the variables used throughout the paper.

2 Definition of a GCM

Humans often represent causality as a directed graph. Philosophers struggle to define causality [37] (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 [54, 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. (However, we will try to justify the causal intuition behind such edges informally.)

[1] In addition to the author, the SCAMP team included Mike Cox, Jason Greanya, Peggy McCarthy, Jonny Morell, Sri Nadella, and Laura Sappelsa, with consulting input from Kathleen Carley.

[2] Major collaborators include, alphabetically by last name, Rafael Alonso, Ted Belding, Rob Bisson, Sven Brueckner, Mike Cox, Keith Decker, Liz Downs, Jason Greanya, Rainer Hilscher, Hua Li, Bob Matthews, Scott Page, Rich Rohwer, Mike Samples, Laura Sappelsa, John Sauter, Bob Savit, Peter Weinstein, and Andrew Yinger.

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

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

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 that nodes can carry;
- a node description function $F : N \rightarrow V$ that gives the values associated with each node;
- an update mechanism $U : F \rightarrow F'$ that changes the node description function.

U often uses information associated with the edges to propagate causal effect 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. To emphasize this generality, we characterize U as an update *mechanism*, even though it implements a functional.

Most GCMs focus on the directed graph $\mathbb{G} \equiv \langle N, E \rangle$. 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.

3 A review of graphical causal models

We motivate our approach to modeling causality with agents by summarizing several formalisms that exploit the causal intuition of a directed graph (Table 1), including SCAMP to facilitate comparison. Columns 2, 3, and 5 describe the N, V, and U components of Eq. 1 in each formalism, validating our characterization of a GCM and the framework against which our first claim will be demonstrated. The last four columns reflect four requirements of a causal model for social scenarios that are, partially or completely, unmet by previous models. Our second claim, the relative advantage of an MAS-based GCM, is based on the observation that SCAMP, unlike previous GCMs, supports all four of these requirements. These four requirements are:

- *Column 6*: 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 serious.
- *Column 7*: Does the formalism support *cycles and feedback*? Feedback loops are pervasive in real systems, and are critical for understanding stability, instability, and emergent behavior.

Table 1 Common causal formalisms, grouped according to the sections in Appendix 1.

1 Name	2 Nodes *N*	3 Values *V*	4 Edge Info *E*	5 Update *U*	6 Probability?	7 Cycles?	8 Time?	9 Agency?
Factor Tree Analysis	Informal	Informal	Informal	Informal	No	No	No	No
Causal Loop Diagrams	Variables	Informal	Sign, delay	Informal	No	Yes	Qualitative	No
Path Diagrams	Variables	$\mathbb{R}$	Correlation coefficient	Multiplication of path coefficients	Paths	No	No	No
Causal Diagram	Statements	[0,1] (belief)	Conditional probabilities	Multiplication	Yes	No	No	No
Influence Net	Events	[0,1] (occurrence)	Weighted influences	Multiplication	Yes	No	No	No
Causal Influence Model	Booleans, ordinals, categories	[0,1]	Weighted influences	Multiplication	Yes	No	No	No
Influence Diagram	Decisions; Chance	Sets	Conditional probabilities	Multiplication	Yes	Some	Some	See discussion
POMDPs	S(tates); A(ctions); O(bservations)	Sets	Conditional probabilities	Dynamic Programming	Yes	Yes	See discussion	See discussion
Fuzzy Cognitive Map	Concepts as fuzzy sets	$[-1,1]$	Weight	Multiplication; Thresholding	See discussion	Yes	No	No
System Dynamics, ODEs	Variables	$\mathbb{R}$	Equation	Integration	Yes	Yes	Yes	No
Stochastic Petri Net	Variables; Transitions	$\mathbb{R}+;[0,1]$	Sequence	Simulation	Yes	Yes	Yes	No
SCAMP	Event types	$[-1,1]^n$	Relations	Simulation	Yes	Yes	Yes	Yes

– *Column 8*: 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.
– *Column 9*: 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.

Capturing agency is a distinctive benefit that agents can bring to causal modeling. Strong support for agency includes accounting for two contrasting features of causal actors. First, they form related *groups* with similar behaviors. Second, their individual histories give each agent a *distinctive state* that might lead two actors in the same situation to behave differently. Lack of the latter capability can lead to unrealistic results in equation-based models [71], warning of the consequences of ignoring it in GCMs as well.

Existing formalisms vary in their support for these features. In particular, support for agency is very weak, yet this feature is critical to modeling human behavior involving different interacting groups. SCAMP's agent-based update mechanism supports all four features.

Appendix 1 provides references, details, and examples for each of these formalisms.

4 The SCAMP formalism

We summarize SCAMP's stigmergic architecture (Sect. 4.1), then present selected elements of SCAMP to substantiate our two claims. Full details are available in an ODD protocol [18] describing the architecture [43], and in the SCAMP user manual [51]. The current implementation is in Java on the Repast platform [3].

The "MP" in SCAMP means "multiple perspectives." SCAMP currently supports four perspectives, two of which are discussed in this paper. For the others, and for further details on the two that we do discuss, see [42, 43, 45, 50].

– The *event* perspective is the causal event graph or CEG (Sect. 4.3). This graph has two kinds of internal edges: agency edges (Sect. 4.4) and influence edges (Sect. 4.7).
– The *goal* perspective (Sect. 4.5) consists a hierarchical goal network (HGN) for each group in the model, and is linked to the CEG.
– The *geospatial* perspective supports event types that require agents to move spatially, and allows agents to interact based on their spatial proximity.
– The *social* perspective allows agents to construct dynamic networks of relations with one another based on their co-participation in events or their encounters in geospace, and to modify their preferences based on those acquaintances. It also supports the dynamic addition of agents to the population, or their removal from it.

The CEG is obligatory. The other three perspectives can be included or omitted, independently of one another. The CEG alone is sufficient to demonstrate the correspondence between SCAMP and a GCM, but because of the importance of goal-based reasoning in agent-based modeling in general, this paper discusses the goal perspective as well. We describe the CEG and HGN selectively, focusing on those details necessary to support our claims.

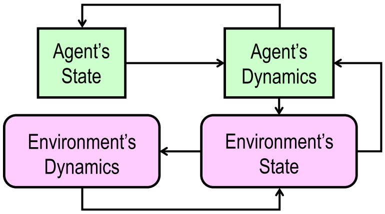

Fig. 1 Basic stigmergic schema (Color figure online)

4.1 SCAMP's stigmergic architecture

SCAMP is based on an agent architecture known as "stigmergy." Grassé [16] coined this word in 1959 from the Greek *stigma* (sign) and *ergon* (action) to describe insect *actions* that are mediated by *signs* in the environment (Fig. 1). An agent's state and local environment determine its actions, and are modified by its actions. The environment modifies its own state, and agents interact by sensing signs left by other agents.

The parade example of stigmergy is the use of pheromones, chemical markers that many social insects deposit in the environment (for example, to construct paths between the nest and food sources) [6]. In this example, following Fig. 1,

- The *agent's state* includes its hunger level and whether or not it is currently holding food.
- The *agent's dynamics*, based on its current state and the pheromone concentrations that it senses in its environment, are to choose the type and strength of pheromones to deposit, and to deposit them at its current location in the environment.
- The *environment's state* is pheromone concentration as a function of location and time.
- The *environment's dynamics* evaporates pheromones over time and propagates them through space.

Stigmergy usually models animal behavior without reference to anthropomorphic cognitive states [6], and has not classically been applied to realistic social behavior for humans. SCAMP uses stigmergy to generate realistic human behavior for testing social science methods, as described in Sect. 1. Section 7 summarizes some of the benefits that stigmergy brings to this application.

Stigmergy represents the agent's behavior externally to the agent. SCAMP's behavioral representations are directed graphs, the set of nodes N and set of directed edges E in Eq. 1 (Sect. 4.3). U for SCAMP (the update mechanism) has three components: the movement of SCAMP agents over their environment (Sect. 4.4), influences between elements of N that are not directly accessible from one another (Sect. 4.7), and computation of the urgency of alternative behaviors from a group's goals (Sect. 4.5). Sections 4.5 and 4.6 discusses how SCAMP's HGNs and agents, respectively, contribute to U. Thus U is an algorithm that is executed, rather than an equation that is evaluated.

4.2 Groups and feature space

Social scenarios usually involve a set of *groups* G with which actors can be affiliated. In our model of civil conflict inspired by Syria, G includes the Government, Armed Opposition, Relief Organizations, People, Violent Extremists, and the Military. Each group has a unique index, returned by $GpIndex{:}G \rightarrow [0..|G|-1]$. Each agent belongs to one group, but can have weighted affiliations with others.

We overload *GpIndex* in two ways.

1. Every agent $a \in A$ belongs to one of these groups, so we allow $GpIndex : A \rightarrow [0..|G|-1]$ to map each agent to the index of its group.
2. Each group has an HGN Γ made up of goals $\gamma \in \Gamma$, so we allow $GpIndex : \Gamma \rightarrow [0..|G|-1]$ to map any goal in an HGN to the index of that HGN's group.

The critical element in the definition of a group is a vector of preferences $P \in [-1, 1]^k$, initially set by the modeler. k is the dimension of the system's *feature space*, so-called because this space also defines features of each node in N, as discussed in Sect. 4.3. Elements of the preference vector that are greater than zero describe the degree to which an agent is attracted to nodes with the corresponding feature, while negative elements describe degree of repulsion. The semantics of different elements in feature space depend on the model, but in the SCAMP model for the Ground Truth program, they have the following semantics:

- The first three elements, the *wellbeing* preferences, record an agent's concern for its physical, emotional, and economic wellbeing.
- The next $|G|$ elements, the *urgency* preferences, record the degree to which an agent wants to advance or oppose the goals of each group.
- The final $|G|$ elements, the *presence* preferences, record the agent's attraction to or repulsion from other agents belonging to each group.

Thus the SCAMP feature space in Ground Truth has dimension $k = 3 + 2|G|$.

The preferences of individual agents are initialized by sampling around the group's preference vector, and they change over time as described in Sect. 4.6. Thus the state of an individual agent depends on its individual history, and may differ from the state of other agents in its group.

4.3 Causal event graph

Stigmergy requires a shared environment in which agents interact. In SCAMP, the heart of this environment is a Causal Event Graph (CEG) inspired by narrative graphs. A narrative environment respects evidence that a fundamental construct underlying human cognition is the narrative [11, 30], a sequence of events. Such graphs are common in intelligence analysis [22], cyber security planning [69], discrete event simulation [67], analysis of social disagreement [68], computer games [35], and the study of natural-language texts [62], among other applications. These formalisms share the following features with the SCAMP CEG:

1. Nodes (members of N in Eq. 1) are *event types*, not variables as in most other causal formalisms.
2. A directed edge between two nodes indicates the *narrative coherence* of moving from one type of event to the next. That is, it would make sense in a narrative of the scenario for an agent who had just participated in the source event to participate next in the target event. Section 4.4 formalizes and encodes this relation.

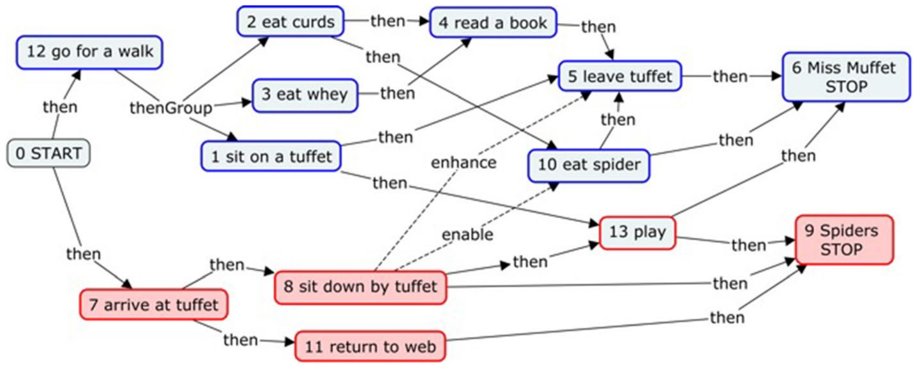

Fig. 2 Causal event graph for Miss Muffet (Color figure online)

3. Any trajectory through the graph represents a plausible narrative.
4. The graph summarizes many possible narratives.

We use the term "event type" rather than "event" because agents can visit a single node repeatedly. An event instance (or event) is defined by an event type node and a time period during which that node has continuous participation by one or more agents. Each instance of an event type begins when an agent chooses to participate in an event type in which no other agents are currently participating, and ends when no agents are participating in the event type. Where there is no danger of confusion, we sometimes overload "event" to describe a node in the CEG.

The CEG we developed in SCAMP's original context has 459 event types and has the potential to generate on the order of 10^{11} possible trajectories (Sect. 5.2). Section 5 summarizes our experience with this full model. For clarity, Fig. 2 presents a much simpler CEG representing a classic children's rhyme.

Little Miss Muffet sat on a Tuffet
Eating her curds and whey.
Along came a spider, and sat down beside her,
And frightened Miss Muffet away.

To illustrate the ability of a CEG to capture multiple possible narratives, this CEG supports not only the original poem, but also some variants.

- The original poem concludes,

  Along came a spider and sat down beside her And frightened Miss Muffet away

- At about age 10, boys discover they can get an entertaining response from girls by modifying the last line to "And she ate that too." (Event type 10, *eat spider*, results in the removal of the spider, using mechanisms in the social perspective not discussed in this paper.)
- Pacifists might prefer a third conclusion, "And they began to play."

Figure 2 generates all three of these variants, and others besides.

According to Eq. 1, each node $n \in N$ is mapped to a value in V by F. The modeler sets initial values, though some can change as SCAMP executes. For clarity we factor $F(n)$ into three parts.

- $F_1 \in [-1, 1]^k$, where $k = 3 + 2|G|$, is a vector in feature space (Sect. 4.2) describing the attractiveness of the node to agents.
- It is not meaningful for agents of every group to participate in every event type. $F_2 \in 2^G$ records the groups whose agents have *agency* for the node. In our example (Fig. 2), Miss Muffet has agency for the blue (upper) nodes, the spiders have agency for the red (lower) nodes, and both groups have agency for node 13.
- F_3 is the parameter of an exponential distribution (corresponding to inter-arrival times in a Poisson distribution) that an agent samples to learn how long its participation lasts. Thus durations can vary not only between event types, but also between agents on the same event type.

For Ground Truth and Miss Muffet, the semantics of F_1 are defined in the same space as preference vectors (Sect. 4.2), as follows:

- The first three elements (wellbeing features) record the impact of the event type on a participating agent's physical, emotional, and economic wellbeing. An agent with a high preference for one of these elements will seek events with a high value for the corresponding feature.
- The next $|G|$ elements (urgency features) record how urgent execution of the event type is to satisfying the HGN of each group, based on the current state of the HGN (Sect. 4.5). An agent with a high preference for advancing a group's goals will seek events with a high urgency feature for the group. Depending on its preferences, an agent can seek to advance the goals of groups other than its own, or pursue actions contrary to its own group's goals.
- The final $|G|$ elements (presence features) record recent degree of participation of agents belonging to each group in the event type, generated as described in Sect. 4.6. These features are the most direct parallel to insect pheromones. An agent that desires to encounter agents of a given group will seek events with a high presence feature for that group.

We use $F_1[wellbeing], F_1[urgency]$, and $F_1[presence]$ to designate these subvectors, and further index them to select a single element. In the present version of SCAMP, agency F_2, the event duration parameter F_3, and $F_1[wellbeing]$ are defined by the modeler, and U changes only $F_1[urgency]$ (via goal edges, Sect. 4.5) and $F_1[presence]$ (via the direct activity of agents, Sect. 4.6. But one can envision processes that modulate $F_1[wellbeing]$ as well, based on external data.

Preferences of agents and *features* of nodes in the CEG represent the time-dependent states of individual agents and individual event types, respectively.

The edges E in the CEG are of two types: agency edges and influence edges. We explain agency edges next. Influence edges will make more sense after we discuss agents in Sect. 4.6.

4.4 Agency edges

The solid edges in Fig. 2 are *agency edges*. An agent on the source of an edge includes the edge's destination among its options for its next choice of participation. For example, if Miss Muffet is currently sitting on her Tuffet (event type 1), she can subsequently choose either to leave the Tuffet (event type 5) or play with the spider (event type 13). The agency edges labeled "then" connect a single current event to a single option for the next event.

An agency edge extends from m to n, $m, n \in N$, on two conditions:

1. $F_2(m) \cap F_2(n) \neq \emptyset$, that is, one or more groups have agency for both events.
2. The edge is *narratively coherent*, that is, it would make sense in a narrative for an agent who has just participated in m to participate next in n.

The *thenGroup* multiedge specifies event types in which agents participate concurrently, if they enter the nodes through the thenGroup edge. For example, thenGroup would allow an agent to participate concurrently in event types *walk* and *chew gum*. At first glance, such a capability violates the principle that an agent can only participate in one event at a time. Here is how thenGroup works in SCAMP.

On parsing a thenGroup edge, SCAMP constructs an event group, a single node in the CEG that acts as a container for the set of nodes to which the thenGroup edge leads. (Each of these nodes may also have its own incoming edges, if it makes sense for some agents to participate in it alone.) The event group has the following features and constraints.

- F_3 for the event group, its duration, is the maximum of the durations returned by the contained nodes, since an agent needs to complete all contained events before moving on.
- F_1 for the event group, its feature vector, is the mean of the feature vectors for the contained nodes.
- Some event types require a participating agent to move through geospace from its current location to event-dependent destination. A physical agent cannot concurrently move to different destinations, so at most one geospatial event type is allowed in an event group.
- The outgoing agency edges of the event group are the combined outgoing agency edges of all contained events, and an agent participating in the event group considers all of these agency edges in choosing its next event.

Agency edges are the foundation of SCAMP's update mechanism U, because an agent can augment the presence features of event types in the CEG only by visiting those nodes. These presence features drive other components of U via influence and goal edges.

The agency edges in Fig. 2 do not in themselves fully satisfy our causal intuitions. Event type 2 *eat curds* does not in itself cause either 4 *read a book* or 10 *eat spider*. However, the combination of these edges with an agent on event type 2 who is motivated by its preferences to choose one of these successors *is* causal. As we asserted in Eq. 1 and showed in Sect. 3, in all causal formalisms, the update mechanism U is an essential part of the causal semantics, and SCAMP is no exception.

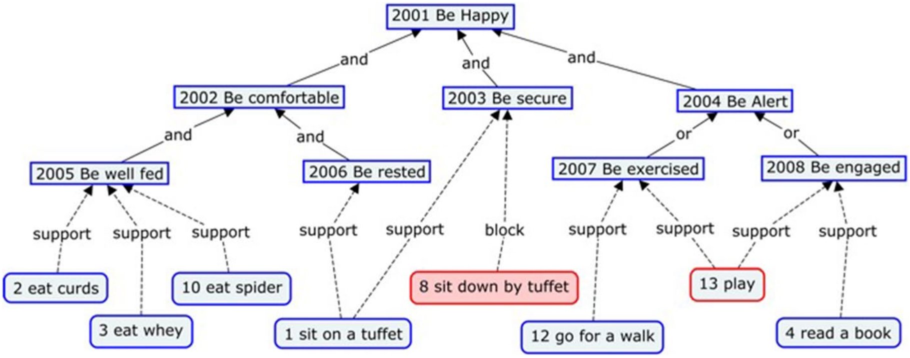

Fig. 3 Hierarchical goal network for Miss Muffet (Color figure online)

4.5 Hierarchical goal network

SCAMP supports a hierarchical goal network (HGN) [70] for each group $g \in G$, describing the high-level goal for the group and its decomposition into subgoals. In this discussion, for conciseness, the term "goal" can indicate either the root goal or a subgoal. A goal is a state of the world that the members of the group wish to achieve, in contrast to the events or actions that describe the event types in the CEG. For example, states do not have agency or duration, as events do. (The same features distinguish hierarchical *goal* networks from hierarchical *task* networks, whose nodes are events.) The *satisfaction* of a goal is an estimate of how well the goal is being achieved in the current state of the world, what TÆMS [23] calls *quality*. A goal's *urgency* is an estimate of how critical the goal currently is to achieving the overall (root) goal.

HGNs satisfy Eq. 1, but the components N, E, F, V, U function differently than they do in the CEG, because agents do not move over the HGN as they do over the CEG and geospatial lattice.

The nodes N of an HGN for group g are a set of goals $\gamma \in \Gamma(g)$. The HGN is acyclic and rooted. γ_0, the root, is the highest-level goal, to which all the others are subgoals. $\Gamma_{leaves}(g)$ is the set of leaf goals, or leaves (goals whose subgoals are not defined). From this point on, we understand Γ to refer to the goals for a single group and omit the reference to g.

E contains two kinds of edges, *combining* and *zip*, whose function is discussed below under U. *Combining* edges connect goals to each other, while *zip* edges combine leaf goals to events in the CEG.

Figure 3 shows an HGN for Miss Muffet and its relation to the CEG. The rounded rectangles across the bottom are event nodes in the CEG. For clarity, we suppress agency and influence edges. The squared rectangles are goals in the HGN, culminating in the root goal at the top. Solid arrows between goals are *combining* edges from subgoals to their higher-level goals. *Or* edges indicate that any subset of subgoals can satisfy the higher-level goal. *And* indicates that all of the subgoals are required to satisfy the higher goal. Dashed arrows are *zip* edges connecting specific event types in the CEG to leaves in the HGN.

For clarity, Fig. 3 shows only edges directed toward the root. For each of these edges there is also an *inverse* edge (of types *invAnd*, *invOr*, *invSupport*, and *invBlock*) between the same two nodes, in the reverse direction. Informally, U propagates satisfaction, generated by activity on events zipped to the leaf nodes, to the root along the edges in the diagram, and propagates urgency from the root back to the CEG nodes along the inverse edges. The HGN thus functions as a multiedge among CEG nodes that are zipped to it. It translates presence features on these nodes into urgency features on the same nodes, allowing agents to modulate their preference-based selection along agency edges in light of the relevance of the alternatives to the goals of each group $g \in G$.

The values V returned by F from goals in the HGN have two components, which we subscript to distinguish them from F in the CEG.

- $F_4 : \Gamma \rightarrow [0, 1]$ is a goal's *satisfaction*.
- $F_5 : \Gamma \rightarrow [0, 1]$ is a goal's *urgency*.

For HGNs, the fourth element in Eq. 1, U, has two parts, one that updates F_4 (satisfaction), and the other that updates F_5 (urgency). Both satisfaction and urgency update through bounded addition and subtraction, constraining their values to [0, 1] on each goal. This mechanism is inspired by quality propagation in TÆMS [23], as extended in our earlier work on rTÆMS [46].

To update F_4, each $\gamma \in \Gamma_{leaves}$ consults the CEG nodes that are zipped to it along zip edges. It adds $\sum F_1(n)[presence]$ for each CEG node n that has a *support* zip edge to it, and subtracts $\sum F_1(n)[presence]$ for each n that has a *block* zip edge to it. In both cases the summation is over the group-specific elements of F_1. Informally, the higher the recent agent participation (by all groups) on supporting CEG nodes and the less recent participation on blocking nodes, the higher the satisfaction of the subgoal. Satisfaction propagates upward through *or* edges as the maximum of the satisfaction levels of the subgoals, and through *and* edges as the minimum.

Once the root γ_0 knows its satisfaction, it updates its urgency, $F_5(\gamma_0) = 1 - F_4(\gamma_0)$, which it propagates along the inverse edges. A goal passes its urgency directly to subgoals to which it has an *invAnd* combining edge. Subgoals joined via *invOr* subtract their own *satisfaction* from their parent's *urgency*. Finally, each leaf goal γ adds its urgency $F_5(\gamma)$ to $F_1(n)[urgency][GpIndex(\gamma)]$ for each *invSupport* zip it originates, and subtracts its urgency from $F_1(n)[urgency][GpIndex(\gamma)]$ for each *invBlock* zip it originates, thus contributing to U for F_1.

An event type for which one group has agency can change the satisfaction of goals of other groups, and also respond to the urgency levels in other HGNs, if it is zipped to subgoals in those HGNs: the HGN in Fig. 3 is for Miss Muffet, but is blocked by spider event type 8. As a result, agents can modulate their decisions by the desire to advance or hinder the goals not only of their own group, but of other groups as well.

4.6 Agents

With these elements in mind, we can trace the operation of SCAMP's agents.

Each agent $a \in A$ has a preference vector $P : A \rightarrow [-1, 1]^k, k = 3 + 2|G|$ over feature space, sampled around the preference vector of its group, with weighted contributions from other groups with which it is affiliated. Let $Succ(n)$ be the set of CEG nodes at the destinations of agency edges originating at node n. An agent a currently on node n computes the dot product between the agent's preference vector and the node's feature vector for each $m \in Succ(n)$ for which $a \in F_2(m)$ (that is, each successor for which a has agency). It exponentiates each value (to make it positive), and normalizes the set to form a roulette wheel that it spins. Before normalization, each value is raised to a power that controls the determinism of the roulette. A determinism of 0 sets all values to 1 and results in a completely random choice, while determinism greater than 1 increases the chance that the largest value will be chosen. This process yields a probabilistic transition, as in POMDPs, but the probabilities emerge from psychologically realistic (event type) features and (agent) preferences, which modelers can define much more naturally than transition probabilities.

SCAMP represents each actor by a polyagent [47], a single *avatar* that continuously deploys a swarm of *ghosts* to explore the future. This construct is an important part of SCAMP's time model. Each avatar maintains its current *domain time*, which it updates each time it ends its participation in an event (in other words, when it advances to the next event node in the CEG). When it leaves its current node $n \in N$, it computes the length of time it has spend on that node. If the current event node is a geospatial node, the duration of the event is the length of time the avatar took to move to the geospatial goal, as outlined in [43]. Otherwise the avatar samples an exponential distribution with parameter $F_3(n)$ and adds it to its current domain time to obtain its new domain time. Avatars begin execution ordered by their domain time, so that they do not run ahead of each other in time. Their ghosts travel into the future a configurable distance to explore the value of each decision alternative.

An avatar's execution has two phases.

First, it sends out several successive *shifts* of ghosts, each with multiple ghosts. The number of ghosts per shift and the number of shifts are parameters of the avatar's group.

All ghosts carry the avatar's preference vector, but because of roulette selection, they do not necessarily follow the same trajectory. Thus ghosts within a shift sample alternative futures for the avatar. The number of ghosts per shift controls the *breadth* of the avatar's reasoning about the future, that is, number of alternative futures that it considers. Each ghost a contributes to U by augmenting the presence features of the event types in which it participates ($F_1[presence][GpIndex(a)]$), a digital analog of pheromones in social insects. The strength of the deposit depends on the quality of the ghost's individual path (the sum of the dot products of its preference vector and the feature vectors of the nodes actually visited).

$F_1[presence]$ is augmented only by ghosts, not avatars. The first shift of ghosts sees no presence features. Ghosts in later shifts respond to presence features deposited by earlier shifts. The number of shifts thus controls the *recursive depth* of the avatar's reasoning about the future.

Second, after all shifts are complete, the avatar a chooses among the successors of its current node n ($Succ(n)$) by applying roulette selection to its group's presence features $F_1[presence][GpIndex(a)]$ on $Succ(n)$. Thus the agent's decision is probabilistic, not deterministic.

When the agent leaves one CEG node and moves to another, it updates not only its time to reflect the duration of the completed event, but also its wellbeing preferences and overall wellbeing, based on the wellbeing features of that event. If a given wellbeing feature is positive, the agent's wellbeing increases, but its preference for that feature decreases, reflecting satiation, while if a feature is negative, the agent's wellbeing decreases and its preference increases. Thus agents' decisions (and the U that they implement) are not Markovian in event space, but reflect their past history and experiences.

Like pheromones, presence features evaporate exponentially over time. For example, consider a stationary ghost that augments its group's presence feature at its current location by d each time step, while the feature evaporates each step by the factor $e \in (0, 1)$. The most recent deposit contributes d, the one from the previous time step de, the one from two time steps back de^2, and so forth. Thus the total presence feature after k steps is

$d + de + de^2 + \cdots + de^{k-1} = d \sum_{i=0}^{k-1} e^i$, which is just the geometric series, and asymptotes to $d/(1-e)$.

4.7 Influence edges

Influence edges (dashed edges in Fig. 2) capture causal influences among event types between which agents do not move directly. For example, a spider sitting by the Tuffet (node 8) may influence Miss Muffet to leave her Tuffet (node 5), but the spider does not have agency for node 5 and cannot participate in it.

An influence edge from CEG node m to n adjusts the segments in any roulette wheel that includes n, based on $\sum F_1(m)[presence]$, the total recent participation by all groups on m.

SCAMP supports four kinds of influence edges. The hard influences *prevent* and *enable* probabilistically exclude or include a successor's segment in the roulette, weighted by $\sum F_1(m)[presence]$. The soft influences *enhance* and *inhibit* adjust the size of successor's segment, based on $\sum F_1(m)[presence]$.

5 Experimental results

For clarity of exposition, this paper has focused on a toy example causal model. In the Ground Truth program, SCAMP modeled and simulated a much more complex causal scenario, described in detail in [50]. To allow this paper to stand alone, this section repeats some of the experimental details from that paper that concern the event and goal perspectives. The full paper also gives results from the geospatial and social perspectives. In evaluating these results, it is important to keep in mind that the purpose of simulators in the Ground Truth program was to generate *realistic* data, not to replicate a real situation. Thus our work so far has not included fitting the model to data and generating testable predictions, though we are currently extending methods that we previously applied for this purpose in geospatial [66] and emotional [40] models.

5.1 Scenario

The model we built in SCAMP, Conflict World, reflects multipartisan conflict in an imaginary country (Tharum) inspired by (but not a detailed replica of) Syria. Every agent is belongs to one of the groups involved in the conflict, here listed with their initial population:

1. The *government* (GO) is authoritarian, bent on retaining its own control of the situation, and willing to oppress its people to keep them in line (16 agents).
2. In some configurations, a distinct *military* (MIL) is initially aligned with the government, but can diverge (6).
3. The *armed opposition* (AO), inspired by the Syrian opposition, is a movement within the country that seeks to reform or replace the government with democratic institutions (11).
4. The *violent extremists* (VE), inspired by ISIS, are an ideologically driven foreign faction that seeks to include the local territory in a larger religious state (16).

5. *Relief agencies* (RA) seek to provide humanitarian relief for civilians, largely in the form of refugee camps both within and just outside of Tharum (6).
6. *People* (PEO) are pro-opposition, anti-government civilians, trying to live their lives (11).

The small initial population (66 agents) may seem unrealistic. However, each avatar decision comes from a swarm of 24 ghost agents, so the actual exploration of alternatives involves nearly 1600 agents. In addition, the social dynamics perspective [50] allows new agents to join the simulation as it runs (for example, an influx of foreign fighters). In one run, the simulation peaked at 850 avatars (20,400 ghosts) and ended with 548 avatars (over 13,000 ghosts).

5.2 Event types and the causal event graph

The model of Tharum supports 459 event types, including:

– large numbers of people move to urban areas
– public demands democratic reforms
– government security forces arrest minority leader
– military refuses to carry out government's orders
– government and opposition leaders commence official talks
– neighbors leave
– people arm themselves
– relief agencies identify an increase in unplanned need
– military bombs opposition-controlled neighborhoods
– funders of relief agencies lose interest in conflict
– protesters share political news on social media
– quality of life at internally displaced people camp improves
– government and armed opposition forces cease negotiations
– head of state/government calls for end to violence
– transitional government invites election monitors

An important part of intelligence analysis (or scenario modeling) is anticipating possible patterns of behavior that the actors of interest might exhibit. Analysts commonly describe scenarios in terms of possible narratives, making the CEG a natural representation for capturing complex social situations, and we originally developed the CEG in support of intelligence analysis [65]. One benefit of this representation is that it amplifies the creativity of analysts by combining narratives that they explicitly formulate to yield a huge number of other narratives that are consistent with these. Ground Truth is not concerned with analytic creativity, but this same amplification means that the CEG can generate an incredibly large number of different behavioral trajectories as data for the research teams.

A simple example illustrates this amplification. An analyst might consider possible narratives $A \rightarrow B \rightarrow C$ and $D \rightarrow B \rightarrow F$, offering agents a total of two possible histories. But if we merge the narratives on event type B, the number of possible trajectories doubles (adding $A \rightarrow B \rightarrow F$ and $D \rightarrow B \rightarrow C$), without defining any additional events.

The amount of combinatorial amplification depends on the length of the analyst's individual narratives and the number of common event types among them, but we can calibrate

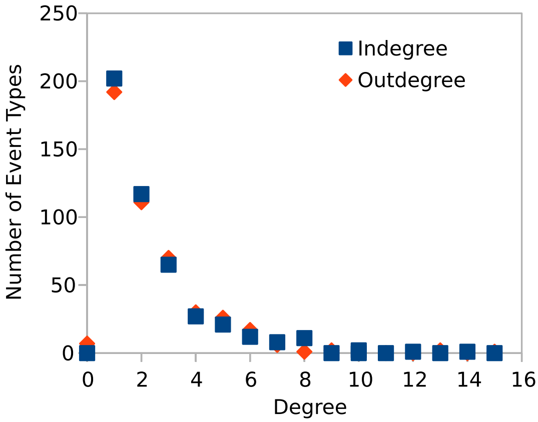

Fig. 4 Degree distribution in challenge 3 CEG (Color figure online)

our intuitions. If we add START and STOP nodes, the CEG is an irregular directed lattice.[3] Consider the paths between diagonally opposite corners in a square directed lattice of side k. Such a lattice contains $(1 + k)^2$ nodes. Except for the START and STOP nodes, this is the number of event types that the analyst must define. The average node degree in a square lattice asymptotically approaches 4, and by symmetry in-degree = out-degree = 2. Each path from START to STOP has $2k$ steps. To go from the lower-left corner to the upper-right one, the agent must choose k out of its $2k$ steps to go up, and go right on the other half, so the number of simple paths is given by the central binomial coefficients, ${}_{2k}C_k$ [8].

Our CEG for Ground Truth has 459 nodes, or 461 with the addition of START and STOP. 461 is not an exact square, but we can build intuition with a lattice of 441 nodes (k = 20), which generates more than 1.3E11 possible trajectories, each of length $2k = 40$. The analyst needs to conceptualize only enough narratives to generate the desired number of event types, with enough overlap to link them into a lattice. For a square lattice with indegree = outdegree = 2, each event type should appear on average in two narratives. Thus 22 narratives of length 40 ($2 \times 441/40$) covering 441 distinct event types suffice to yield a 441 node lattice, far fewer than the 1.3E11 trajectories such a lattice contains.

Our model of Tharum has an average degree of 4.7. The event indegree and outdegree distributions are highly skewed (Fig. 4), and nearly identical. The events with no outgoing edges are STOP nodes for each group, and one event (START) has no incoming edges. This skewing reduces the generative power of the CEG, but even so the number of possible paths greatly exceeds those envisioned by the modeler. SCAMP's swarming ghosts develop a probability field over this massive space, sampling it for data generation (in Ground Truth) or for intelligence analysis.

Most event types are restricted to agents of one or a few groups, according to F_2. Thus the narrative space is partitioned into smaller subgraphs for each group. However, agents can move from one subgraph to another, if they are affiliated with both groups.

Figure 5 shows the CEG for the Conflict World. Colors reflect the agency of the various events. This CEG generates a huge number of alternative narratives, giving a very rich *event space* within which agents move.

[3] SCAMP allows cycles, so the CEG no longer defines a partial order over events, but this discussion excludes such cycles. Their presence strengthens our argument by adding to the space of possible narratives.

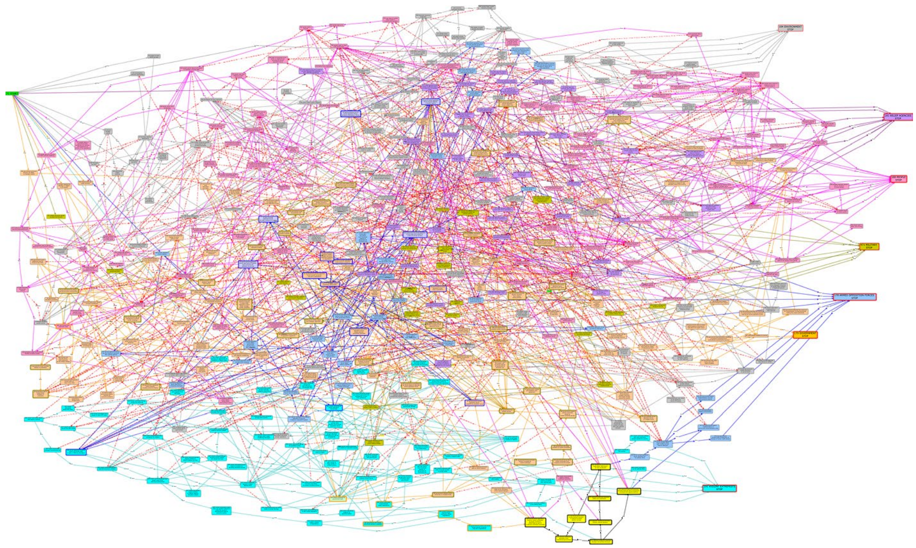

Fig. 5 Causal event graph for Tharum Conflict World (Color figure online)

Table 2 Example event trajectory for an unaffiliated military agent

Domain time	Event number	Event name
1	E521	Military perceives threat to peace and security
17	E536	Military implements operational plans
28	E351	Government sets up checkpoints at official border crossings
143	E24	Government and armed opposition forces increase fighting in border regions
209	E281	Government and armed opposition forces wage fierce battle for control of critical territory
409	E281	Agent dies

The START node for all groups is at the upper left of Fig. 5, and has edges to all of the group subgraphs. Each group has its own STOP node, at the right. When an agent reaches a STOP node, or some other node that (because of *prevent* influence edges) offers no next choices, it returns to START, but in most cases will not retrace its previous path, since its own state and the state of the event nodes will have changed. Each time the overall participation on a CEG node goes from zero to non-zero, a new instance of that event type begins, and ends when the node's participation next drops to zero.

The main data about events available to the social science research teams is an agent history reporting, for each agent, the event in which it is participating at successive times. For example, Table 2 shows the trajectory for an unaffiliated military agent, that is, the history of events in which it participates. Each agent maintains its history internally, but after each agent movement, we also log the move in a file for subsequent analysis.

The similarity among agents in the same group can be detected by comparing their trajectories. One way to do this is with the "string edit" or "Levenshtein" distance [34], which is the minimum number of changes (additions, deletions, or replacements of one element

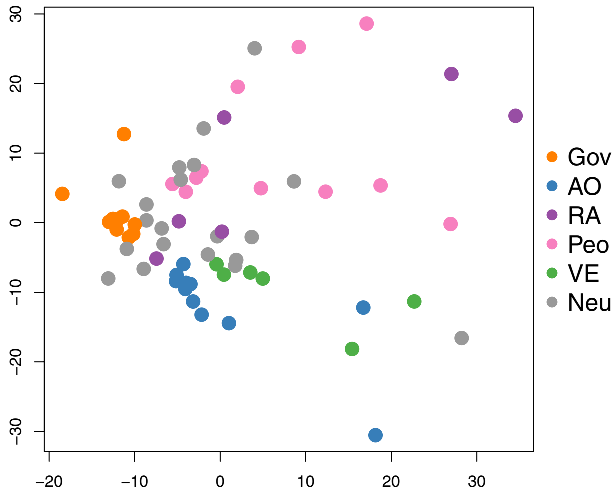

Fig. 6 Trajectory similarities (Color figure online)

by another) needed to change one list of elements into the other. Figure 6 shows the result of plotting such similarities from a run, using multidimensional scaling (R's isoMDS).

The plot clearly reflects the impact of a per-group parameter that defines how much each agent's preferences vary from the group baseline. The preferences of Gov agents are sampled ± 0.2 around the values in the group baseline, AO, VE, and RA vary ± 0.3, and Peo vary ± 0.4. (In this particular run, there was no separate Military; Neu(tral) agents have randomly assigned preferences.) Groups with low sampling variation generate similar trajectories, while those with higher variation generate more diverse trajectories. Affiliations of individual agents with multiple groups also diversify their trajectories.

The dynamics of SCAMP, driven by agent decisions, drastically restrict the state space generated by linked narratives. All nodes in Fig. 5 are reachable from the START node, but when we run SCAMP, the actors visit only 252 of the nodes and traverse only 461 of the edges. A lattice of 252 nodes contains over 10^8 trajectories, but in fact the 850 actors, belonging to 6 groups with different goals and preferences, generate only 164 distinct trajectories in all. This huge reduction is an emergent property of the system's causal dynamics. Analysts could not anticipate it by inspection. Nor is this a brittle result of a single run, because SCAMP samples stochastically and generates a probability distribution over possible futures.

5.3 Goals

The exogenous and presence features on event nodes support tactical decisions by agents, based on their preferences for these features. Hierarchical goal networks (HGNs) for each group modulate the urgency features, supporting strategic decisions. The HGNs for the six groups have a total of 122 goals and subgoals. Of these, 77 subgoals are leaves, zipped to 177 event types.

Figure 7 shows an example of how satisfaction levels at the root of each HGN vary over time in a version of the model without a distinct Military group. The Armed Opposition rapidly gains satisfaction, but then becomes more frustrated as Government satisfaction increases. Relief Agencies take longer to achieve their goals, but as long as they do,

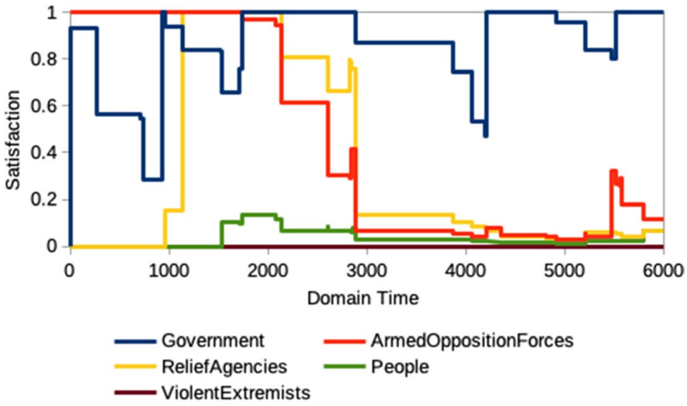

Fig. 7 Satisfaction levels in an early experiment (Color figure online)

the People achieve some satisfaction. However, as the satisfaction of Armed Opposition decreases, so does that of Relief Agencies, and then that of People.

The paper on which this section is based [50] includes details on the geospatial and social perspectives, and experimental results from those perspectives, as well.

6 Previous work combining MASs and GCMs

Every MAS embodies causality, because it models how the environment changes the agent's state and how the agent seeks to modify its environment. Thus every MAS may be described by an external causal model, but not every MAS may be said to reason over a causal model while updating values on its nodes. Examples of MASs that do satisfy this criterion do not provide functionality beyond what a non-agent form of the model would support.

The literature reflects two kinds of relations among agents and GCMs. In some cases, the entire GCM is internal to each agent. In others, different parts of the GCM are assigned to different agents. Both of these approaches differ from SCAMP. Unlike the first, our agents do not contain the GCM. Rather, the GCM contains the agents. Unlike the second, our agents represent domain actors, not segments of the GCM, and move over the GCM in updating it. Table 3 summarizes some previous examples, and includes SCAMP for ease of comparison.

For our purposes, the most distinctive feature of each approach is the relation between the agent and the GCM.

Multiply-sectioned Bayes networks (MSBNs) [79] assign homogeneous subnetworks of a Bayes network to different agents, to limit the amount of inference that is needed when the network is updated with new data. In their original form, agents in MSBNs have no natural alignment with actors in the domain, though extensions to MSBNs support domain agents [78].

Agents often have internal GCMs to support part or all of their reasoning. A GCM such as a Bayes network or its generalization as an influence diagram may be the means by which an agent interprets its environment (for example, maintaining models of other agents) as one part of its cognitive processing [1]. In a more advanced version of this

Table 3 Examples of previous relevant work

Technique	Agent represents	Agent-GCM Relation	Communications	Motivation
Multiply-Sectioned Bayes Networks	Homogeneous subnetwork	Agents partition GCM	Shared variables	Efficiency through localized updating
POMDP	Domain actor	Agents contain GCMs	Outside model	Reasoning about other agents' reasoning
Agent-Encapsulated Bayes Network	Domain actor	Agents contain GCMs	Shared variables	Efficient communications
Communicative Multiagent Team Decision Problem	Domain actor	Agents contain GCMs	Generic	Abstract model for complexity and optimality analysis
SCAMP	Domain actor	GCM contains agents	Via effect of actions	GCM accessible to non-programmer domain experts

approach, each agent's entire behavior is driven by an explicit causal model, for example, a POMDP [57, 58]. In these systems, each agent is indeed updating a GCM, using the U appropriate to the formalism, but the focus is on how the causal model supports the agent, not on how the agent can enhance causal reasoning. Neither embedding the causal model in an agent nor having multiple agents addresses the limitations intrinsic to the specific GCM used.

Agent-encapsulated Bayes networks identify specific input and output variables in each agent's internal network, unifying the inter-agent communication model with the internal structure of the network.

Initial efforts to distribute processing of POMDPs across agents without communications [5] led to the discouraging result that the complexity of processing for two or more agents is complete in non-determininstic exponential time (NEXP). Allowing communications yields the communicative multiagent team decision problem (COM-MTDP) [59], which makes the problem tractable. But as with other multi-agent approaches discussed so far, the GCM is internal to the agents. The research emphasis in these programs is in managing information sharing among components of the network, not in advancing technical support for causal reasoning.

We know of no prior work comparable to the approach described here, in which

– agents representing domain actors
– move dynamically over a GCM to update it,
– fully supporting all four of the requirements summarized in the Introduction (probabilistic analysis of alternative pathways, cycles and feedback, quantitative time, and agency associated with different causal factors).

SCAMP was constructed for a specific purpose (generating realistic social data). But it leads us to the novel insight that an MAS, as a computational mechanism, may have as its primary purpose the analysis of a GCM, extending the GCM's expressiveness compared with current causal models.

7 Discussion

Five issues deserve discussion.

1. How does SCAMP demonstrate our two main claims?
2. How can a stigmergic system, inspired by biological agents with low levels of cognition, generate data consistent with human social and psychological behavior?
3. What are the benefits of a stigmergic approach?
4. What limitations does it impose?
5. To what further research topics do our insights lead?

7.1 Demonstrating the claims

We began with two claims:

1. A stigmergic MAS with an appropriate stigmergic environment is a GCMs.
2. Such an MAS has advantages over other GCMs for modeling causality in social systems.

Table 4 SCAMP satisfies Eq. 1

Component in Eq. 1	Function	SCAMP (Sect. 4)
N	Nodes in directed graph	Event types in CEG; Subgoals in HGN
E	Directed edges among nodes	CEG: agency and influence edges; HGN: *and*, *or*, *zip* edges and their inverses
V, F	Values on nodes	CEG: F_1 (wellbeing, urgency, presence features), F_2 (agency), F_3 (duration parameter); HGN: F_4 (satisfaction), F_5 (urgency)
U	Update mechanism	CEG: modifies F_1[*presence*] by ghost deposits; HGN: modifies F_1[*urgency*] along inverse zip edges; F_4 along zip and combining edges; and F_5 along inverse combining edges

SCAMP demonstrates both of these claims.

An extensive survey of GCMs (Sect. 3 and Appendix 1) shows that they all satisfy the four features in Eq. 1. SCAMP has these same four features, as summarized in Table 4. This alignment demonstrates the first claim.

In Sect. 1, we avoided a formal definition of causality, simply declaring that a directed edge in a causal graph implies causality. Still, it is helpful to consider informally how each type of edge in a SCAMP model reflects our intuitions about causality.

The SCAMP edges that correspond most closely to edges in traditional causal modes are *influence* edges in the CEG. These edges directly assert that participation of agents in one type of event directly modulates the likelihood that another type of event will attract participation.

The CEG's *agency* edges record which next types of events make sense for an agent currently on a given event. The simple structure of these edges is not causal: the fact that one event follows another does not mean that the first caused the other. But we have repeatedly emphasized that causality in any GCM includes not only its structure, expressed in N and E, but also its update mechanism U. SCAMP's update mechanism in the CEG is driven by decisions that agents make about which next events best satisfy their immediate preferences and their strategic goals. In modeling social scenarios, human choice is a critical component of causality, and the agency edges show how the environment constrains and guides that choice.

The edge types in the HGN record the causal relation between events that happen in the world and the level of satisfaction that agents feel, which in turn provides strategic direction to their behavior. The normal edges (*zip*, *and*, and *or*) adjust the agents' satisfaction based on what is happening in the world, while the inverse edges (*invZip*, *invAnd*, and *invOr*) modulate the attractiveness of event types in the CEG as agents consider their next steps.

Our second claim asserts that an appropriately configured MAS supports the four requirements where other GCMs fall short: probability, cycles, time, and agency. SCAMP supports these 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 from psychologically realistic modeling primities (agent preferences and event features) that vary over time. (2) The presence features on each event type are deposited by polyagent ghosts 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 have nominal durations but not start times. Time-anchored events are an emergent feature of the system's operation, as agents begin and end their participation in event types. Thus an agent can meaningfully revisit an event node multiple times.

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 start and end times of each event are well defined.

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.

7.2 Stigmergy and human cognition

It may seem counterintuitive to use stigmergy, a mechanism inspired by social insects, to generate cognitively realistic behaviors. The most direct approach to such constraints would seem to be an agent model such as Belief-Desire-Intention (BDI) [61], or one based on Bayesian formalisms believed to reflect human cognition. However, these approaches embed cognitive behavior in the agent code. We want professional analysts, subject-matter experts with no computer programming experience, to generate our causal ground truth.

Our solution, justified in more detail in [42], lies in "Simon's Law" [72]:

> An ant, viewed as a behaving system, is quite simple. The apparent complexity of its behavior over time is largely a reflection of the complexity of the environment in which it finds itself.

Simon extends this principle to human behavior:

> Human beings, viewed as behaving systems, are quite simple. The apparent complexity of our behavior over time is largely a reflection of the complexity of the environment in which we find ourselves.

In spite of stigmergy's simplicity, Simon's Law enables SCAMP to capture psychologically and socially realistic dynamics.

- The preference-feature method for selecting options models deliberate choice.
- The use of a roulette reflects modern insights into non-deterministic decision-making [7].
- Strategic (goal-driven) decisions are a recognized feature of human rationality [70].
- So is use of mental simulation (modeled by polyagents) to look ahead in time [30].
- In SCAMP's social perspective (not discussed in this paper), agents adjust their preferences as they interact with other agents encountered on event types or geospatial tiles [13].
- The foundation of the CEG in studies of narrative reflects the centrality of narrative as a mental representation [11].

- Indirect constraints on agents via influence edges, HGNs viewed as multiedges across event types, and (in the full model) interactions in geospace provide realistic bounds on the rationality of individual agents [73].

[42] discusses these features in more detail.

7.3 Benefits of stigmergic processing

While SCAMP's stigmergic architecture is unusual compared with other architectures commonly used in agent-based modeling, it has several benefits.

- SCAMP captures numerous features of modern psychological theory (Sect. 7.2).
- SCAMP's causal rules are human-readable graphs that analysts can readily understand and personally construct and modify. Other formalisms may use graphs to document the internal structure of the agent code, but the risk remains of a mismatch between the understanding of the analyst and that of the software engineer. The isomorphism between SCAMP and a GCM is particularly clear because SCAMP's stigmergic mechanisms allow it to reason over a graph with clearly defined N and E.
- SCAMP's different perspectives are integrated through their effect on the scalar components of the feature vectors of CEG nodes and the preference vectors of the agents. This interface among different perspectives is minimal, compared with the representations often required by conventional agents to connect (say) group goals, social connections, and agent decisions with one another. Integrating perspectives by modifying vectors of scalars makes implementation of additional perspectives straightforward.
- With conventional agents, estimating distributions over alternative possible outcomes requires repeated runs. SCAMP's polyagent technology generates a distribution of possible outcomes (the ghost pheromone field) with a single run. In our current configuration, we use this distribution to generate a single trajectory, but the pheromone field could be analyzed to develop (for example) uncertainty bounds around the results of a single run.

7.4 Limitations of stigmergic processing

Researchers familiar with other agent technologies that might be used in social modeling (such as BDI architectures [61]) may find SCAMP limiting in two ways: the informal nature of the events and goals that it uses to define a scenario, and the lack of direct interagent messaging.

Given the rich tradition of knowledge representation and symbolic reasoning in classical AI, the natural-language labels on CEG nodes and HGN goals in SCAMP seem naïve and undisciplined. In fact, the program does not understand these labels. Decisions are based on

- the structure of agency and influence edges in the CEG,
- the pattern of *and* and *or* edges in the HGNs,
- the *zip* edges that connect the CEG and the HGNs together,
- and the preferences and features assigned to groups and events, respectively.

Viewed by themselves, these details are semantically impoverished. But they do not operate by themselves. One design objective of SCAMP was to enable domain experts who are not programmers to construct and modify complex social models [45], and in the course of the Ground Truth program, such experts successfully constructed and meaningfully interpreted multiple models. The labels they assigned to events and goals guided their selection of edges, preferences, and features, and we repeatedly observed how discussions of the natural-language labels influenced analyst decisions about model connectivity and parameters. The human modeler is the link that relates SCAMP's formal structure and its semantics. The alternative presented by a semantically rich representation raises the knowledge acquisition barrier *between* a software engineer and an analyst, resulting in a model that is opaque to the user.

One disadvantage of a purely stigmergic model is that agents cannot exchange messages directly with one another. But stigmergy is not an all-or-nothing design decision for agents. Humans are the parade example of explicit rational reasoning and symbolic communication, and yet numerous human behaviors are most naturally viewed as stigmergic [41]. SCAMP itself violates strict stigmergy in the social perspective, discussed in other publications [50]: if two agents participate in the same event type or visit the same geospatial tile concurrently, they exchange their preference vectors and use them (according to individual parameters) to modify their own preferences, modeling the influence of our associates on our own attitudes. Though highly stereotyped, this extension demonstrates that a stigmergic code base can be extended to provide more conventional features, such as a rich interagent language.

7.5 Directions for future work

Our insights in this paper can be usefully extended in several ways.

- In our current system, exogenous components of the feature vector ($F_1[wellbeing]$) are static, defined by the modeler. But they would be a natural place to couple SCAMP's directed graph to external sensors or other reasoning modules, extending it from an off-line planning and modeling tool to a real-time control system.
- The stigmergic process of making and sensing quantitative changes provides a simple interface for adding new perspectives. For example, a computer network perspective could model a cyber attacker traversing a network, and a resource perspective could model events that produce and consume resources.
- Analytic versions of U (in conventional GCMs) and algorithmic versions (in an MAS) are not mutually exclusive. The representational benefits of agent-based over equation-based models are well documented [52], but these benefits come at the cost of increased computational complexity and slower execution. It would be interesting to explore hybrid mechanisms for U that combine both analytic and agent-based mechanisms. For example, how could stigmergic agents enhance the update process on directed graphs whose nodes represent variables rather than events?
- We are exploring the use of genetic methods to tune a SCAMP model (including both the agents and the structure of the environment) to match observations [44].
- Because U operates by modifying individual elements of N, it is well suited to parallelization by distributing the environment on a GPGPU [21], greatly accelerating its performance.

– The presence features over N form (up to normalization) a *probability field* indicating where we may expect to find agents. The expression "probability field" is most common in quantum physics, where it describes the amplitude of a particle's wave function, giving the probability that when the wave function is collapsed by observation, the particle will be found at each location. The quantum parallel to SCAMP goes further. Feynman's path formalization of quantum theory [10] (originally inspired by ant behavior) imagines that a particle follows every possible path, and sums its contributions to obtain this field. While ghost agents do not explore every possible path, they do sample many of them to compute SCAMP's probability field. These parallels suggest that mathematics from quantum field theory might offer powerful tools for analyzing the results of SCAMP simulations.
– One could cast the insights in this paper in the framework of category theory by defining a category of GCMs, identifying the various models we have considered (including SCAMP) as objects in this category, and exploring morphisms among them. Such an analysis would enrich our understanding of both MASs and more generic GCMs, and support the development of hybrid architectures.

Appendix 1: Details on other graphical causal models

This appendix describes each formalism (other than SCAMP) listed in Table 1, and gives an example of it. These examples build on the children's rhyme introduced in Sect. 4:

Little Miss Muffet sat on a Tuffet
Eating her curds and whey.
Along came a spider, and sat down beside her,
And frightened Miss Muffet away.

We illustrate how each formalism might represent the causal relations involving whether or not Miss Muffet is at the Tuffet. The formalisms fall into several classes, grouped by the horizontal lines within Table 1.

Non-computational

Some formalisms are intended primarily for human examination. The semantics for nodes and their values (and in one case, for edges) are quite loose, and updating U is an informal subjective assessment by the practitioners.

Factor tree analysis (root cause analysis) includes several causal graphs used in industry. These graphs are for human review rather than automatic analysis, so the nodes, values, and update mechanisms are not formalized. One example is the Ishikawa or fishbone diagram [26], used to trace the causes of quality problems in manufacturing. In this model, the top-level causes are pre-defined branches (e.g., Equipment, Process, People, Materials, Environment, Management), to lead analysts to consider different areas where quality problems may arise. Edges leading into each of these branches are primary causes; they in turn support secondary causes, and so forth. The lower level nodes are verbal descriptions of causes, and may be events (e.g.,"shipments delayed"), measurable observations ("rusty components"), or even prepositional phrases. Figure 8 shows a simple fishbone diagram for the problem "Miss Muffet not on Tuffet," This diagram suggests that the problem may

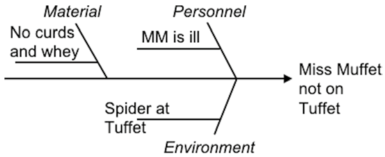

Fig. 8 Factor tree for Miss Muffet

Fig. 9 Causal loop diagram for Miss Muffet

result from difficulties with material (lack of curds and whey), personnel (Miss Muffet ill), or the environment (presence of a spider).

Such diagrams contribute greatly to the quality of manufactured products, but are too ambiguous for formal analysis.

The nodes in *causal loop diagrams* [74] are understood to have scalar values, though they are not evaluated. Edges are labeled with a + if an increase in the source promotes an increase in the destination, and a - if an increase in the source promotes a decrease in the destination. Optionally, a double slash // indicates an unspecified time delay. Figure 9 shows a causal loop diagram for Miss Muffet. The left-hand loop means that she will only sit on the Tuffet if she has something to eat, but the longer she is there, the less curds and whey will remain. The right-hand loop means that she eventually attracts the lonely spider, who scares her away. Causal loop diagrams, unlike many other formalisms, do support causal loops, and the (informal) updating U for a given node is understood to reflect the changes in the node's value over time.

While the causal loop diagram is qualitative and not quantitative, it is the basis for the stock-and-flow model, discussed in "Other analytic models" section in Appendix 1, which supports a set of ordinary differential equations and can thus yield quantitative results.

Correlation

Sewall Wright's path diagrams [77] are the basis for modern structural equation models (SEMs) [2, 76] and marginal structural models (MSMs) [63], and an inspiration for Bayesian Causal Diagrams [54]. Path diagrams compute correlations between variables connected either directly or indirectly by directed edges. The update mechanism U computes the correlation between end points of a path as the product of correlations along a single path, and as the sum of correlations entered by different paths. Extensions estimate the values of latent variables based on observed variables and compute conditional probabilities throughout the graph under assumptions of causal influence, but say nothing about how these conditional relations occur. Path diagrams do not represent time or allow cycles, and U propagates node values through the graph at a single point in time. Figure 10 shows a simple path diagram for Miss Muffet. The nodes representing the presence of Miss Muffet and of the spider at the Tuffet are binary.

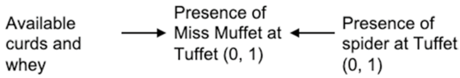

Fig. 10 Path diagram for Miss Muffet

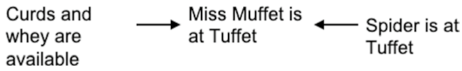

Fig. 11 Causal diagram for Miss Muffet

Probabilistic

Bayesian causal diagrams [53] are a dominant causal formalism, commonly used in planning experiments and selecting experimental variables. Figure 11 shows a simple causal diagram for Miss Muffet. Nodes carry probabilities. *U* evaluates these probabilities by multiplying the probability of cause nodes through conditional probability tables attached to each effect node. One application of such graphs uses convention rather than computation for *U*. Because of the formalism's probabilistic semantics, the pattern of connections alone can be used for experimental design, identifying what variables should and should not be observed to confirm a causal hypothesis.

These graphs have two shortcomings. First, like path diagrams, the underlying mathematics does not allow cycles or support time intrinsically (though temporal extensions have been proposed [39]). Second, computation over probabilities requires complete conditional probability tables on the nodes, and these are hard to procure, particularly for non-repeatable social situations offering limited data.. The latter problem has motivated the development of "canonical models" [9] that make simplifying assumptions about the relations among the nodes in order to reduce the parameters needed to evaluate the model.

One example of a canonical model is the Influence Net [64]: nodes are events with baseline probabilities. Edges can be supporting or inhibitory, and contain two probabilities: that the effect will obtain if the cause is true, and if the cause is false. *U* propagates these values across a net—again, at a single point in time. The method recognizes the importance of events rather than variables as nodes, but characterizes them simply by probability of occurrence, treating them as propositions of the form "Event X occurred" to which belief values can be assigned.

Another canonical simplification of the Bayesian causal graph is the Causal Influence Model [55], whose nodes can be Boolean, ordinal, or categorical. They have a baseline probability (if categorical, a baseline probability for each option), and the connection between two nodes is an influence in $[-1, 1]$ on the probability of the target. *U* updates the baseline probability of the child node based on a function (typically the mean) of the influences of its parents.

Influence diagrams

Influence diagrams are a large and important family of GCMs. Their central feature is the distinction between *decision* nodes (reflecting agent choices) and *uncertainty* or *chance* nodes (reflecting chance events). Chance nodes may model questions or experiments that the decision maker can perform, with various probabilities of outcomes. This distinction

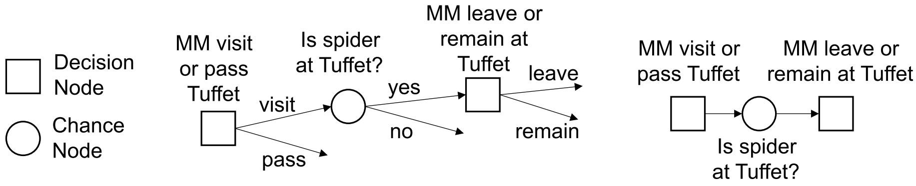

Fig. 12 Partial decision tree (left) and influence diagram (right) for Miss Muffet

emerged from an earlier formalism, decision trees, which constrained the order of evaluation of the two node types to be a tree, as alternate moves in a game between the decision maker and a partner named "chance" [60]. Influence diagrams ([24], reprinted as [25]) remove this ordering dependency, yielding directed acyclic graphs. In both decision trees and the earliest influence diagrams, the expected payoff from a node is associated with the outgoing edge corresponding to the choice that the decision-maker or chance has made. Later formalisms define *value* or *utility* nodes to record these payoffs.

Figure 12 shows a fragment of a decision tree (left) and influence diagram (right) for Miss Muffet. The outcome of the chance node "Is spider at tuffet?" depends on whether Miss Muffet decides to visit the tuffet on her walk, and the presence or absence of the spider determines whether she remains at the tuffet to finish her snack, or leaves prematurely. Influence diagrams are a specialization of probabilistic GCMs, and their update mechanisms U are based in one way or another on the chain rule and Bayes rule, with a variety of algorithmic refinements, such as message passing [33] and variable elimination [36], to name only a few.

Influence diagrams represent two advances over the formalisms considered so far.

1. They formally introduce a primitive notion of agency. Decision nodes are under the control of one distinguished agent (the decision maker), while the chance node captures both aleatoric uncertainty and epistemic uncertainty (including the actions taken by all other agents). Thus the agent modeled by the diagram has a sense of self vs. other. However, there is no notion of group affiliation, and no individual state.
2. Chance and decision nodes represent choices among different *actions*, while value or utility nodes represent *variables*. Intuitively, variables (the nouns and adjectives in a scenario) do not cause anything. Influence diagrams recognize that causality should be a relation among *events*, the verbs of the scenario.

Numerous refinements of influence diagrams (IDs) have been developed. For example:

– partial IDs (PIDs) [38]) relax the condition that the decision variables be ordered temporally;
– limited memory IDs (LIMIDs) [33] relax the assumption that previous observations are remembered and considered in all future decisions;
– unconstrained IDs (UIDs) [28] and sequential (SIDs) [27] relax the order of observations;
– dynamic IDs (DIDs) [20] allow chance nodes to change state, and incorporate cycles.

Of particular importance to the MAS community is the extension of influence diagrams to multi-agent scenarios by distinguishing multiple agents. For example:

Fig. 13 POMDP for Miss Muffet

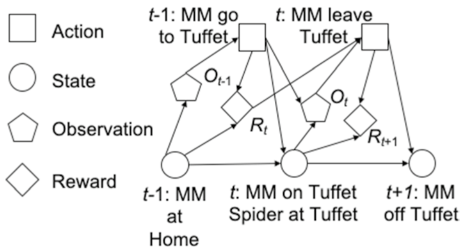

- multi-agent IDs (MAIDs) [31] incorporate separate decision and utility variables for each agent in a single influence diagram.
- ID networks (IDNs) [14] and networks of IDs (NIDs) [15] construct a digraph of MAIDs, then solve them from the leaves to yield a single MAID incorporating relevant chance nodes that represent their solutions.
- interactive dynamic IDs (I-DIDs) [56] nest DIDs, allowing agents to model one another.

One particularly important derivative of influence diagrams is the Partially Observable Markov Decision Process (POMDP) [29] (e.g., [58]), with three types of nodes. Like a Markov process, a POMDP defines a set S of *states* and transition probabilities among them. Like a Markov decision process (MDP), it augments the set of states with a set $\mathbb{A}$ of *actions* α that some agent can take, based on the current state. The agent chooses among available actions to maximize its reward function $R : Sx\mathbb{A} \rightarrow \mathbb{R}$, which typically computes the expected time-discounted future reward for each available action. A set of transition probabilities $T(s'|s,\alpha)$ computes the system's next state, conditioned on the present state and the action chosen, generalizing the transition probabilities in the simple Markov process. A POMDP further adds a set of *observations* Ω (corresponding to the chance nodes in a primitive influence diagram) derived from the world state through a set of observation probabilities $O(\omega|s,\alpha)$ conditioned on the state being observed and the action that brought the agent into that state. Thus the next action in a POMDP is driven by the state of the world probabilistically, not deterministically.

Like the other formalisms in this paper, the POMDP lends itself to representation as a directed graph. The nodes are states, actions, and observations, and edges are conditional probabilities from T and O. Though rewards are strictly speaking a function over actions and states, the common use of influence diagrams as a representation for POMDPs [75] leads to the convention of representing them as additional nodes in the graph. Figure 13 shows a fragment of a POMDP for Miss Muffet in her decision to leave the Tuffet.

Like other Markov processes, POMDPs can revisit a state more than once, and so support cycles. They offer a partial solution to representing time and agency. For tractability, POMDPs in agent-based systems use a discrete time model with a constant period of time between successive actions. Thus the model captures time, but all actions have the same duration. POMDPs also incorporate agency, because an agent performs each action. Extensions [56] assign different agents to different actions. However, agency is at the level of individual agents, without intrinsic support for groups of agents, and state nodes capture (an agent's view of) the state of the environment (including other agents), not the personal state of the agent. In addition, POMDPs require model builders to think in terms of transition probabilities, rather than psychological primitives from which probabilities are generated by the model.

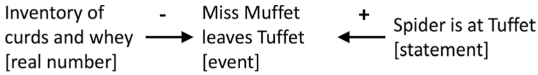

Fig. 14 Fuzzy cognitive map for Miss Muffet

$$\frac{dc}{dt} = -\mu m, \quad \frac{dm}{dt} = \kappa c - \sigma s, \quad \frac{ds}{dt} = \nu m$$

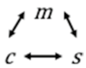

Fig. 15 ODE model for Miss Muffet

Other analytic models

Fuzzy cognitive maps (FCMs) [32] are inspired by feed-forward neural networks. Nodes are *concepts*, and may include variables, events, and entity names. Values *V* on FCMs are *activation levels* derived from the intrinsic value of the category and scaled (depending on the domain) either to [0, 1] or [− 1, 1], and edges are weights in [− 1, 1]. *U* consists of multiplying the activation of a node's causes by weights and summing, usually through a thesholding function to maintain activation bounds. This process, unlike the probabilistic theory behind causal diagrams, permits causal cycles (as in a recurrent neural network). Thus *U* is not simply propagating causality at a single time throughout the graph, but generating the dynamics exhibited by node values over time, though FCMs have no quantitative representation of time.

Figure 14 is a toy FCM for Miss Muffet illustrating the informality of concepts. This imprecision extends to the meaning of *activation*. If concepts are viewed as events, then activation is reasonably understood as probability of occurrence. If they are statements, activation becomes level of belief. The informality of its semantics and the ease with which multiple models of the same domain can be combined makes the method attractive for participatory modeling involving mathematically unsophisticated domain experts [17]. For such users, *activation* is a surrogate for node probability in indicating the prominence of a concept, but is not constrained to a formal probabilistic semantics.

Though causal loop diagrams ("Non-computational" section in Appendix 1) do not directly support computation, they are often the first step to a System Dynamics model [12]. This formalism is inspired by physical theories and is based on ordinary differential equations (ODEs) rather than Bayesian probability. Nodes evaluate to variables that are transformed into one another by the edges, using a metaphor of fluid flow often described as "stocks and flows." Stocks correspond to variables in an ODE, while flows correspond to first derivatives. Time is intrinsic to the behavior of a differential equation, so these models allow feedback loops and characterize the system's behavior over time. *U* consists of integrating the equations through time. As a continuous formalism, ODEs deal more easily with real-valued quantities than Boolean or integer values, and we adjust our running example to accommodate the following variables:

- c amount of curds and whey available
- s spider population near Tuffet
- m Miss Muffet population on Tuffet

Figure 15 shows a few ODEs and the corresponding directed graph for Miss Muffet, based on these variables. μ, κ, σ, ν are the transition rates for the ODEs.

Like probabilistic models and unlike factor trees and fuzzy cognitive maps, the values associated with the nodes of System Dynamics models have clearly defined mathematical

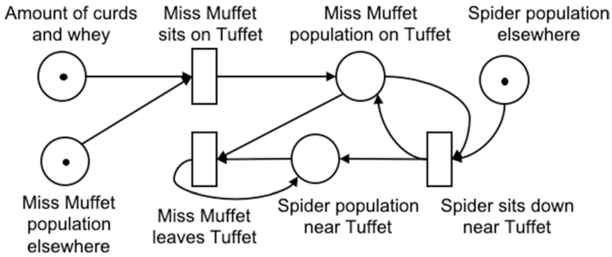

Fig. 16 Stochastic Petri Net for Miss Muffet

meaning, but with very different semantics: ODEs support cycles and quantitative time, while probabilistic models do not. The difference is due to the difference in *U*. In an ODE *U* evolves node values *through* time, but in a probabilistic model it simply propagates node values to achieve a consistent labeling *at a point in time*.

A Stochastic Petri Net (SPN) [19] is a bipartite digraph whose nodes alternate between integer-valued *places* (thus, variables) and *transitions* that can have durations (thus modeling the passage of time) and activation probabilities (thus "stochastic"). *U* is algorithmic rather than analytic: a transition is eligible to fire when all of its input places are greater than 0, and when it fires, it decrements each input place by 1 and augments each output place by 1. If a transition does not have the same number of input and output places, the total value of places in the net is not conserved. A place's value is sometimes called its *marking* and represented graphically by dots. Figure 16 is an SPN for a fragment of Miss Muffet. Circles represent places, while rectangles represent transitions.

The reader can verify from Fig. 16 that

- Miss Muffet requires curds and whey to sit on the Tuffet.
- When she sits down, the amount of curds and whey decreases.
- The spider requires Miss Muffet to be on the Tuffet in order to come near the Tuffet.
- Miss Muffet's departure from the Tuffet requires both that she is on the Tuffet, and that the spider has approached.
- Miss Muffet and the spider are conserved, while the curds and whey are not.

SPNs can be represented (within continuity constraints) as sets of ODEs [4] and thus evaluated by integration. Unlike system dynamics diagrams but like influence diagrams, they distinguish between *variables* and the *events* that change them. Unlike influence diagrams, they do not capture agency. In Fig. 16, the markings for Miss Muffet and the spider are not agents that participate in events, but semaphores that enable the events.

Discussion of conventional methods

None of the four desirable features we summarized at the beginning of Sect. 3 is uniformly supported by all methods. The most common is column 6 in Table 1, probabilistic estimate of effects resulting from causes, formally supported in seven of the 11 formalisms. Column 7 (cycles and feedback) is supported in only six of the formalisms, and a quantitative estimate of time (column 8) in only four. Only influence diagrams (including POMDPs) support agency. One can always encode a particular agent as a causal node, but without a sense of an *event*, it is difficult to formalize the *agent* responsible for the event. Even the Influence Net, whose nodes are events, simply estimates the probability of their occurrence rather than representing their action. Even influence diagrams do not capture the effect of an agent's history on its decisions.

No formalism considered so far has a clear semantics of group agency. Psychological and social features are naturally associated with the different groups of agents involved in a scenario. Again, one can always instantiate a social or psychological feature as a causal node, but the treatment is ad-hoc and not integral to the formalism. This observation suggests that agent-based models, with their explicit semantics of agency and action, can fill an important gap in modeling causality. Influence diagrams do have a natural notion of agency for individual agents, but the representation is restricted to probabilities. SCAMP provides a much richer set of modeling artifacts to capture important psychological and sociological features.

Appendix 2: Symbols

Table 5 summarizes the symbols used in this paper.

Table 5 Symbols used

Symbol	Meaning
A	Set of agents a in a SCAMP model
$\mathbb{A}$	Set of actions α in a POMDP
C	GCM, defined as $\langle N, V, E, U \rangle$
E	Set of edges e in a GCM
F	Mapping $N \to V$
F_1	F on CEG node restricted to wellbeing, urgency, and presence features
F_2	F on CEG node restricted to agency
F_3	F on CEG node restricted to nominal duration
F_4	F on HGN goal restricted to satisfaction
F_5	F on HGN goal restricted to urgency
G	Set of groups g in a SCAMP model
$\mathbb{G}$	Directed graph $\langle N, E \rangle$
Γ	Set of goals γ in SCAMP's HGNs; may be restricted to a group g
$GpIndex$	Map from groups G, agents A, or goals Γ to an integer group index
N	Set of nodes m and n in a GCM
O	Set of observation probabilities in a POMDP
Ω	Set of observations ω in a POMDP
P	Preference vector of an agent a or a group g
R	Reward function in a POMDP
$\mathbb{R}$	Real numbers
S	Set of states s in a POMDP
$Succ$	Map from CEG nodes N to their successors along agency edges
T	Set of transition probabilities t in a POMDP
U	Update mechanism $V \to V'$ in a GCM
V	Set of values that N can have in a GCM

Acknowledgements 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. This paper has benefited from the careful reading and thoughtful suggestions of multiple JAAMAS reviewers.

References

1. Albrecht, S., & Stone, P. (2018). Autonomous agents modelling other agents: A comprehensive survey and open problems. *Artificial Intelligence, 258*, 66–95.
2. Amorim, L. D. A. F., Fiaccone, R. L., Santos, C. A. S. T., Santos, TNd., Moraes, L. T. L. P. D., Oliveira, N. F., Barbosa, S. O., Santos, DNd., Santos, LMd., Matos, S. M. A., & Barreto, M. L. (2010). Structural equation modeling in epidemiology. *Cadernos de Sade Pblica, 26*, 2251–2262.
3. Argonne National Laboratory. (2007). Repast agent simulation toolkit. http://repast.sourceforge.net/
4. Baez, J. C., & Biamonte, J. D. (2018). *Quantum techniques for stochastic mechanics*. World Scientific.
5. Bernstein, D. S., Zilberstein, S., & Immerman, N. (2000) The complexity of decentralized control of Markov decision processes. In C. Boutilier & M. Goldszmidt (Eds.) *Sixteenth conference on uncertainty in artificial intelligence* (UAI2000) (pp. 32–37). Morgan Kaufmann.
6. Bonabeau, E., Dorigo, M., & Theraulaz, G. (1999). *Swarm intelligence: From natural to artificial systems. SFI studies in the sciences of complexity*. Oxford University Press.
7. Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review, 100*(3), 432–459.
8. Dickau, R. (2020). Counting lattice paths. https://www.robertdickau.com/lattices.html
9. Diez, F. J., & Druzdzel, M. (2006). Canonical probabilistic models for knowledge engineering. Technical Report. CISIAD-06-01, UNED
10. Feynman, R. (1948). Space-time approach to non-relativistic quantum mechanics. *Reviews of Modern Physics, 20*(2), 367–387. http://authors.library.caltech.edu/47756/1/FEYrmp48.pdf
11. Fisher, W. R. (1989). *Human communication as narration: Toward a philosophy of reason, value, and action*. University of South Carolina Press.
12. Forrester, J. W. (1961). *Industrial dynamics*. MIT Press.
13. Friedkin, N. E., & Johnsen, E. C. (2011). *Social influence network theory: A sociological examination of small group dynamics structural analysis in the social sciences*. Cambridge University Press.
14. Gal, Y., & Pfeffer, A. (2003). A language for modeling agents' decision making processes in games. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems* (pp. 265–272). Association for Computing Machinery
15. Gal, Y., & Pfeffer, A. (2008). Networks of influence diagrams: A formalism for representing agents' beliefs and decision-making processes. *Journal of Artificial Intelligence Research, 33*, 109–147.
16. Grassé, P. P. (1959). La reconstruction du nid et les coordinations inter-individuelles chez *Bellicositermes natalensis* et *Cubitermes sp*. La théorie de la stigmergie: Essai d'interprétation du comportement des termites constructeurs. *Insectes Sociaux, 6*, 41–84.
17. Gray, S. A., Gray, S., De Kok, J. L., Helfgott, A. E. R., O'Dwyer, B., Jordan, R., & Nyaki, A. (2015). Using fuzzy cognitive mapping as a participatory approach to analyze change, preferred states, and perceived resilience of social-ecological systems. *Ecology and Society, 20*(2), 11.
18. Grimm, V., Berger, U., DeAngelis, D. L., Polhill, J. G., Giske, J., & Railsback, S. F. (2010). The ODD protocol: A review and first update. *Ecological Modelling, 221*(23), 2760–2768. http://www.agsm.edu.au/bobm/teaching/SimSS/ODD_protocol.pdf
19. Haas, P. J. (2002). *Stochastic Petri Nets: Modelling, stability, simulation*. Springer.
20. Hazen, G. B. (2004). Dynamic influence diagrams: Applications to medical decision modeling. In M. L. Brandeau, F. Sainfort & W. P. Pierskalla (Eds.), *Operations research and health care: A handbook of methods and applications* (pp. 613–638). Springer. https://doi.org/10.1007/1-4020-8066-2_24.
21. Hermellin, E., & Michel, F. (2016). GPU delegation: Toward a generic approach for developping MABS using GPU programming. In C. M. Jonker, S. Marsella, J. Thangarajah, & K. Tuyls (Eds.), *Proceedings of the 2016 international conference on autonomous agents and multiagent systems* (pp. 1249–1258).
22. HeuerRichards, J. J., & Pherson, R. H. (2010). *Structured analytic techniques for intelligence analysis*. CQ Press.

23. Horling, B., Lesser, V., Vincent, R., Wagner, T., Raja, A., Zhang, S., Decker, K., & Garvey, A. (2004). The TÆMS white paper. http://mas.cs.umass.edu/pub/paper_detail.php/182
24. Howard, R. A., & Matheson, J. E. (1984). Influence diagrams. In R. A. Howard & J. E. Matheson (Eds.), *Readings on the principles and applications of decision analysis* (Vol. 2, pp. 719–762). Strategic Decisions Group.
25. Howard, R. A., & Matheson, J. E. (2005). Influence diagrams. *Decision Analysis, 2*(3), 127–143.
26. Ishikawa, K., & Loftus, J. H. (1990). *Introduction to quality control*. 3A Corporation
27. Jensen, F. V., Nielsen, T. D., & Shenoy, P. P. (2006). Sequential influence diagrams: A unified asymmetry framework. *International Journal of Approximate Reasoning, 42*(1), 101–118.
28. Jensen, F.V., & Vomlelov, M. (2002) Unconstrained influence diagrams. In: A. Darwiche & N. Friedman (Eds.), *Eighteenth conference on uncertainty in artificial intelligence (UAI'02)* (pp. 234–241). Morgan Kauffmann.
29. Kaelbling, L. P., Littman, M. L., & Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial Intelligence, 101*(1), 99–134. https://doi.org/10.1016/S0004-3702(98)00023-X
30. Kahneman, D., & Tversky, A. (1982). The simulation heuristic. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 201–208). ambridge University Press.
31. Koller, D., & Milch, B. (2001). Multi-agent influence diagrams for representing and solving games. In *Proceedings of the 17th international joint conference on artificial intelligence* (Vol. 2, pp. 1027–1034). Morgan Kaufmann.
32. Kosko, B. (1986). Fuzzy cognitive maps. *International Journal of Man-Machine Studies, 24*(1), 65–75.
33. Lauritzen, S. L., & Nilsson, D. (2001). Representing and solving decision problems with limited information. *Management Science, 47*(9), 1235–1251.
34. Levenshtein, V. I. (1966). Binary codes capable of correcting deletions, insertions, and reversals. *Soviet Physics Doklady, 10*(8), 707–710.
35. Lindley, C. A. (2005). Story and narrative structures in computer games. In B. Bushoff (Ed.), *Developing interactive narrative content* (Chap. 10). High Text Verlag.
36. Mauá, D. D., de Campos, C. P., & Zaffalon, M. (2012). Solving limited memory influence diagrams. *Journal of Artificial Intelligence Research, 44*, 97–140.
37. Mumford, S., & Anjum, R. L. (2013). *Causation: A very short introduction*. Oxford University Press.
38. Nielsen, T. D., & Jensen, F. V. (1999). Welldefined decision scenarios. In *The fifteenth conference on uncertainty in artificial intelligence (UAI 1999)* (pp. 502–511). Morgan Kaufmann.
39. Park, C. Y., Laskey, K., Costa, P. C. G., & Matsumoto, S. (2014). A predictive situation awareness reference model using multi-entity Bayesian networks. In *Seventeenth international conference on information fusion (FUSION 2014)*.
40. Parunak, H. V. D. (2004). Evolving swarming agents in real time. In *Genetic programming theory and practice (GPTP05)*. Springer. https://www.abcresearch.org/abc/papers/GPTP05.pdf
41. Parunak, H. V. D. (2006). A survey of environments and mechanisms for human-human stigmergy. In D. Weyns, F. Michel, & H. V. D. Parunak (Eds.), *Proceedings of E4MAS 2005*. Lecture notes on AI (Vol. 3830, pp. 163–186). Springer.
42. Parunak, H. V. D. (2020). Psychology from stigmergy. In *Computational social Science (CSS 2020)* (Vol. (forthcoming). CSSSA.
43. Parunak, H. V. D. (2020). ODD protocol for SCAMP. Report. Wright State Research Institute. https://abcresearch.org/abc/papers/ODD4SCAMP.pdf
44. Parunak, H. V. D. (2021). Learning actor preferences by evolution. In *Computational social science (CSS21)*. CSSSA.
45. Parunak, H. V. D. (2021). Social simulation for non-hackers. In K. H. Van Dam & N. Verstaevel (Eds.), *22nd International workshop on multi-agent-based simulation (MABS 2021)*. Springer.
46. Parunak, H. V. D., Belding, T., Bisson, R., Brueckner, S., Downs, E., Hilscher, R., & Decker, K. (2009). Stigmergic modeling of hierarchical task networks. In G. D. Tosto & H. V. D. Parunak (Eds.), *The tenth international workshop on multi-agent-based simulation (MABS 2009, at AAMAS 2009)* (Vol. 5683, pp. 98–109). Springer.
47. Parunak, H. V. D., & Brueckner, S. (2006). Concurrent modeling of alternative worlds with polyagents. In *The seventh international workshop on multi-agent-based simulation (MABS06, at AAMAS06)* (pp. 128–141). Springer.

48. Parunak, H. V. D., Brueckner, S., Downs, L., & Sappelsa, L. (2012). Swarming estimation of realistic mental models. In F. Giardini & F. Amblard (Eds.), *Thirteenth workshop on multi-agent based simulation (MABS 2012, at AAMAS 2012)* (Vol. 7838, pp. 43–55). Springer.
49. Parunak, H. V. D., Brueckner, S. A., Matthews, R., & Sauter, J. (2006). Swarming methods for geospatial reasoning. *International Journal of Geographical Information Science, 20*(9), 945–964.
50. Parunak, H. V. D., Greanya, J., Morell, J. A., Nadella, S., & Sappelsa, L. (2021). SCAMP's stigmergic model of social conflict. *Computational and Mathematical Organization Theory*. https://doi.org/10.1007/s10588-021-09347-8
51. Parunak, H. V. D., Morell, J. A., Sappelsa, L., & Greanya, J. (2020). SCAMP user manual. Report, Parallax Advanced Research https://www.abcresearch.org/abc/papers/SCAMPUserManual.zip
52. Parunak, H. V. D., Savit, R., & Riolo, R. L. (1998). Agent-based modeling vs. equation-based modeling: A case study and users' guide. In N. Gilbert, R. Conte, & J. S. Sichman (Eds.), *First International Workshop on Multi-agent systems and agent-based simulation*. LNCS (pp. 10–25). Springer.
53. Pearl, J. (2009). *Causality* (2nd ed.). Cambridge University Press.
54. Pearl, J., & Mackenzie, D. (2018). *The book of why*. Basic Books.
55. Pfautz, J., Cox, Z., Catto, G., Koelle, D., Campolongo, J., & Roth, E. (2007). User-centered methods for rapid creation and validation of Bayesian belief networks. In K. B. Laskey, S. M. Mahoney, & J. Goldsmith (Eds.), *Fifth Bayesian modeling applications workshop (UAI-AW 2007) at BMA '07* (pp. 37–46).
56. Polich, K., & Gmytrasiewicz, P. (2007). Interactive dynamic influence diagrams. In *6th International joint conference on autonomous agents and multi-agent systems* (pp. 1–3).
57. Pynadath, D., & Marsella, S. (2005). PsychSim: Modeling theory of mind with decision-theoretic agents. In *International joint conference on artificial intelligence* (pp. 1181–1186).
58. Pynadath, D. V., Dilkina, B., Jeong, D. C., John, R. S., Marsella, S. C., Merchant, C., Miller, L. C., & Read, S. J. (2021). Disaster world: Decision-theoretic agents for simulating population responses to hurricanes. *Computational & Mathematical Organization Theory* (forthcoming).
59. Pynadath, D. V., & Tambe, M. (2002). The communicative multiagent team decision problem: Analyzing teamwork theories and models. *Journal of Artificial Intelligence Research, 16*, 389–423.
60. Raiffa, H., & Schlaifer, R. (1961). *Applied statistical decision theory*. Harvard University.
61. Rao, A. S., & Georgeff, M. P. (1995). BDI agents: From theory to practice. In *The first international conference on multi-agent systems (ICMAS-95)* (pp. 312–319). AAAI.
62. Richards, W., Finlayson, M. A., & Winston, P. H. (2009). Advancing computational models of narrative. Tech. Rep. MIT-CSAIL-TR-2009-063, MIT CSAIL.
63. Robins, J. M., Hernn, M., & Brumback, B. (2000). Marginal structural models and causal inference in epidemiology. *Epidemiology, 11*(5), 550–560.
64. Rosen, J. A., & Smith, W. L. (1996). Influence net modeling with causal strengths: An evolutionary approach. In *Command and control research and technology symposium*.
65. Sappelsa, L., Parunak, H. V. D., & Brueckner, S. (2014). The generic narrative space model as an intelligence analysis tool. *American Intelligence Journal, 31*(2), 69–78.
66. Sauter, J. A., Matthews, R., Parunak, H. V. D., & Brueckner, S. (2002). Evolving adaptive pheromone path planning mechanisms. In *Autonomous agents and multi-agent systems (AAMAS02)* (pp. 434–440). ACM. https://www.abcresearch.org/abc/papers/AAMAS02Evolution.pdf
67. Savage, E. L., Schruben, L. W., & Ycesan, E. (2005). On the generality of event-graph models. *INFORMS Journal on Computing, 17*(1), 3–9.
68. Shapiro, B. P., van den Broek, P., & Fletcher, C. R. (1995). Using story-based causal diagrams to analyze disagreements about complex events. *Discourse Processes, 20*(1), 51–77.
69. Sheyner, O. M. (2004). Scenario graphs and attack graphs. Ph.D. thesis, Computer Science
70. Shivashankar, V. (2015). Hierarchical goal networks: Formalisms and algorithms for planning and acting. Ph.D. thesis, Computer Science.
71. Shnerb, N. M., Louzoun, Y., Bettelheim, E., & Solomon, S. (2000). The importance of being discrete: Life always wins on the surface. *Proceedings of the National Academy of Sciences of the United States of America, 97*(19), 10322–10324.
72. Simon, H. (1969). *The sciences of the artificial*. MIT Press.
73. Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics, 69*(1), 99–118.
74. Sterman, J. (2000). *Business dynamics*. McGraw-Hill.
75. Tatman, J. A., & Schachter, R. D. (1990). Dynamic programming and influence diagrams. *IEEE Transactions on Systems, Man, and Cybernetics, 20*(2), 365–379.
76. VanderWeele, T. J. (2012). Invited commentary: Sructural equation models and epidemiologic analysis. *American Journal of Epidemiology, 176*(7), 608–612.

77. Wright, S. (1934). The method of path coefficients. *Annals of Mathematical Statistics 5*(3), 161–215.
78. Xiang, Y., Jensen, F., & Chen, X. (2006). Inference in multiply sectioned Bayesian networks: Methods and performance comparison. *IEEE Systems, Man, and Cybernetics, 36*(3), 546–558.
79. Xiang, Y., Poole, D., & Beddoes, M. P. (1993). Multiply sectioned Bayesian networks and junction forests for large knowledge-based systems. *Computational Intelligence, 9*(2), 171–220.