Inverse Generative Social Science Workshop 2021

Abstracts

The following is a master list of extended abstracts accepted to the workshop.

[Keynote 1] Beyond Allocation: How ABMs Can Transform The Study and Design of Institutional Ensembles

Jenna Bednar & Scott E Page

University of Michigan, Ann Arbor, USA

> “Was it not only the timidity of an age which had lost all confidence in ultimate values which led us to attempt to claim ‘scientific’ justifications for attitudes which in the nature of things could not be justified (or refuted) by appeal to laboratory methods” - *Lionel Robbins (1938 pg 639)*

ABMs offer the potential to transform the study of institutions by expanding the breadth and depth of analyses. Societies rely on ensembles of institutions, both formal and informal, to make decisions, allocate resources, and establish order across a variety of domains (Arrow 1974, Ostrom 2009). Mathematical institutional theory (mechanism design) considers institutions in isolation and takes a limited view of context, characterizing it as technology, preferences, and information, and focuses on allocations and decisions as outcomes. ABMs enable the study of ensembles of institutions, and can reveal how an existing set of institutions influences the performance and design of new institutions. In other words, ABMs model institutions as part of a system as opposed to in isolation.

To make that methodological advance ABMs expand the conception of context to include **civic capacity:** social capital (Putnam 1993), behavioral repertoires (Dolan and Galizzi 2015, Henrich et al 2001, 2004), tool kits (Swidler 1986), belief systems (Ingelhart 1997), network structures (Jackson, Rogers, and Zenou 2017) and the distribution of knowledge (Chwe 2003, Easterly and Easterly 2006). These types of civic capacity, like allocations and decisions, can be thought of as outcomes of a mecahnism. It follows that a mechanism might produce desirable allocations but degrade civic capacity. For example, a market might ensure Pareto efficient allocations but result in self-interested behaviors, network ties based on mutual and self interest, and specialized knowledge, which could undermine democracy.

Challenges of three-objective bi-level model discovery: Application to alcohol use in New York State, 1985-2015

Tuong Manh Vu, Charlotte Buckley, Alan Brennan, & Robin Purshouse

University of Sheffield, United Kingdom

This study aims to explain the population trends in alcohol use in New York (NY) State, 1985-2015. Over this period, male drinking in NY dropped during the 1980s before recovering during the 1990s, while female alcohol use was quite stable. Based on social norms theory, we developed a mechanism-based agent-based simulation and calibrated the model to alcohol use behavior over the 1985-2015 period. Within the parameter calibration, the baseline model structure produces only flat trajectories for both male and female drinking or fits well to male alcohol use targets but performs not so well for female targets.

We proposed using model discovery to explore different model structures in order to find an explanatory model that can adequately fit both male and female alcohol use. Specifically, we use grammar-based genetic programming with three optimization objectives: male alcohol use, female alcohol use, complexity. The study also attempts to perform bi-level optimization by including the parameters in the model discovery process.

The model discovery process demonstrated the ability to discover structures and parameters with better fit to the targets, but still with a trade-off between male and female goodness-of-fit. The importance of descriptive and injunctive norms was found to differ between males and females, with compromise models focusing on injunctive norms. In addition, we identified three challenges in the study. The first challenge is balancing the trade-offs between the three targets and deciding which models are credible. Secondly, it is possible to include parameters in the model discovery process, but the number of parameters and their range should be kept to the minimum to prevent an abundant number of discovered candidates. Lastly, theoretical interpretability is always a difficult task requiring several rounds of discussion with domain experts and potentially grammar modification to improve the model discovery process.

An ABM generating realistic epidemic dynamics and a genetic algorithm to implement vaccinations minimizing the number of symptomatic people

Pietro Terna[15], *Gianpiero Pescarmona*[1], *Alberto Acquadro*[1], *Paolo Pescarmona*[2], *Giuseppe Russo*[3], *Emilio Sulis*[1], *& Stefano Terna*[4]

[1] *University of Torino, Italy*

[2] *University of Groningen, Netherlands*

[3] *Centro Einaudi, Torino, Italy*

[4] *tomorrowdata.io*

[5] *Collegio Carlo Alberto, Torino, Italy*

The communication revolves around a micro-based model of interacting agents, following plausible behavioral rules into a world where the Covid-19 epidemic is affecting the actions of everyone. A short paper about the model is at https://rofasss.org/2020/10/20/sisar/. The model works with (i) infected agents categorized as symptomatic or asymptomatic and (ii) the places of contagion specified in a detailed way, thanks to agent-based modeling capabilities. The infection transmission is related to three factors: the infected person's characteristics and the susceptible one, plus those of the space in which contact occurs. The model includes the structural data of Piedmont, an Italian region, but we can easily calibrate it for other areas. The simulation reproduces a realistic calendar (e.g., national or local government decisions) via a dedicated script interpreter.

The micro-based structure of the model allows factual, counterfactual, and conditional simulations to investigate both the spontaneous or controlled development of the epidemic. Examples of counterfactual situations are those both (i) of different timing in the adoption of the non-pharmaceutical containment measures or (ii) of alternative strategies focusing exclusively on the defense of fragile people.

In this way, the model is generative not of social situations but of complex epidemic dynamics emerging from the consequences of agents' actions and interactions, with high variability in outcomes, but frequently with a stunning realistic reproduction of the successive contagion waves that occurred in the reference region. We take charge of the variability of the epidemic paths within the simulation, running batches of executions with 10,000 occurrences for each experiment.

The inverse generative side of the model comes from constructing a meta-agent optimizing the vaccine distribution among people groups–characterized by age, fragility, work conditions–to minimize the number of symptomatic people. We can characterize the action of the planner both (i) introducing ex-ante rules following “plain” or “wise” strategies that we imagine as observers or (ii) evolve those strategies via the application of a genetic algorithm. The genome is a matrix of vaccination quotas with the time interval of validity. The inverse generative side of the model comes from constructing a meta-agent optimizing the vaccine distribution among people groups–characterized by age, fragility, work conditions–to minimize the number of symptomatic people.

Inverse generated identity from interactions in self-organizing groups

Juan Salamanca & Daniela Delgado

University of Illinois, USA

Although individuals in homogeneous social groups share similar mental models, their concurrent action flows experience mutual frictions exhibiting a systemic social viscosity. In this paper, we examine how a blank-slate actor introduced in a self-organizing social system can develop an approximation of the dominant mental model with minimal disruption in social viscosity. The new actor uses reinforcement learning --inspired in theories of identity formation-- to evolve social models from

situated interactions with others. The result is twofold: first, the models are imperfect yet socially apt, and second, the impact of the learning agent's action on social viscosity depends on the breadth of its interaction domain and agents' tolerance to rule-compliance.

Our test bench is an agent-based model (ABM) representing people organizing themselves by color proximity. It has two independent variables: scope of interaction and tolerance to imperfect proximity. All agents (N=20) observe others' colors within their relative network seeking to be close to those with similar hues and far from dissimilar ones. They adjust their position with a degree of tolerance from one state to the next. This results in orderly arrangements of agents replicating Newton's innate color wheel, as described in the literature and confirmed empirically in our lab. The learning agent lacks a color model to estimate color proximity but gradually develops one that allows it to participate in the ongoing organization process and converge around one of the emerging attractors.

The learning agent learns how to choose the best model from a pool of color permutations instead of mapping situations to actions. It implements an adaptation of the Q-learning algorithm that resembles Identity Control Theory. It accounts for the inter-subjective construction of social identity, reaffirmed by constant feedback from others present within the social setting. In practice, each repositioning produces a reward that reflects how well others perceive the learning agent's actions and reassures or punishes a registry of model quality values (q-values). This technique is highly demanding in computing resources. More efficient ones need to be explored.

Social viscosity evaluates the compatibility of the learning agent's action flow with those of other agents at every time step. It aggregates the discrepancy between each agent's expected situation relative to other agents and the actual realization of the world. In self-organizing systems, social viscosity gradually decreases as agents maximize their proximity function. Interactions with a learning agent exacerbate agents' action modulation, but it also decreases as the learner acquires a stable social model. It seems that this metric is a proxy for the assessment of uncoordinated social action.

In conclusion, we present an ABM in which an agent inverse generates a model of the environment that enables it to make situated inferences about how others within a social domain will behave and inform its action modulation. Our methodological contribution is the use of identity control theory in conjunction with multi-agent reinforcement learning to generate models for action rather than action choices. The generated models could account for alternative world-views unforeseen during the ABM design.

Meta-brain Models: Biologically complex agent-based models

Bradly Alicea & Jesse Parent

Orthogonal Research and Education Lab, USA

Agent-based models present a sophistication challenge on the part of an agent's behavior. While agent-based models can produce behavioral complexity without a complex brain, models that simulate animal behavior are nevertheless assumed to do so with a brain. At the same time, Artificial Intelligence presents a similar challenge involving representational complexity: while minimal representations can produce behavioral outputs like locomotion or simple decision-making, more elaborate internal representations might offer a richer variety of behaviors. While Artificial Intelligence is more concerned with individual behaviors, the consequences of rich behavioral repertoires and flexible internal models has consequences for constructing more realistic and informative agent-based models.

We propose that these issues can be addressed with a computational approach we call meta-brain models. Meta-brain models are hybrid models of Artificial Intelligence that includes layered components of varying degrees of representational complexity. Rather than taking a conventional neural network or deep learning approach, layering in a meta-brain is instead inspired by biological neuroanatomy. To demonstrate this, we propose a two-layer meta-brain consisting of one representation-free (standard connectionist model) and one representation-rich (analogical system) component. We will propose combinations of layers composed using specialized types of models. Rather than using a generic black box approach to unify each component, this relationship mimics systems like the neocortical-thalamic system relationship of the mammalian brain, which utilizes both feedforward and feedback connectivity to facilitate functional communication. Importantly, the relationship between layers can be made anatomically explicit. For example, the representation-rich component can "encase" or be "layered" on top of the representation-free component. This allows for structural specificity that can be incorporated into the model's function in interesting ways.

We further propose that the meta-brained agent architecture can be

utilized in agent-based models of self-organized and population-level phenomena, as well as social behaviors. In this talk, we will introduce several model systems to demonstrate this compatibility, including chemical morphogenesis, the ethnocentrism game (Schelling's segregation model), and flocking/swarming. The application of meta-brain models to specific systems also allows us to consider how developmental and evolutionary approaches might be used to build our layered heterogeneous representations in a specific environmental context. Of particular interest is the dynamic behavior of connections between each layer along with the establishment of new layers.

The meta-brains project currently involves three aims: creating models with different degrees of representational complexity, creating a layered meta-architecture that mimics the structural and functional heterogeneity of biological brains, and an input/output methodology that is flexible enough to accommodate behavioral and social phenomena alike. We will conclude by proposing next steps in the development of this flexible and open-source approach.

Evolving the Optimal Transmission Network

Aaron Bramson[1] *& Patrick Grim*[2]

[1] *RIKEN, Japan*

[2] *Stony Brook University, USA*

Network propagation is a well-studied phenomena, with models reflecting hundreds of variations. Typical models in a variety of fields concentrate on only a small handful of standard network structures (small-work networks, preferential attachment networks, lattices and the like), tracking how those network structures affect propagation. This project takes the reverse approach; assume a simple mechanism and evolve networks that optimize propagation. This evolutionary approach provides a far wider sampling of network characteristics.

Our population consists of 50 directed networks each with 50 nodes and initially 100 non-reflexive edges randomly assigned. The propagation dynamics follow a probabilistic susceptible-infectious (SI) epidemiological model: in each generation and for each network we perform 5 trials by changing which agent is initially infected. Our measure for network fitness is 1 divided by the average (over the five trials) of the number of steps the model runs until one of two halting conditions is met: (1) all agents are infected or (2) no agent changes state for S steps. The value of S effectively acts as a fitness penalty to networks that do not manage to infect all agents.

We evolve networks by applying genetic algorithm techniques to their adjacency matrices. For each of 100 generations, the 25 networks with the highest fitness are kept as is, and are joined by 25 altered copies formed using two alteration mechanisms: transcription (copying a random splice of one of the top 25 network's adjacency matrix) and mutation (flipping the (0,1) value for m in {0.0, 0.5, 1.0 } percent of the edges).

Our results show a divergence of solutions: (1) one in which connectivity increases so as to reduce the time to saturation and another (2) in which there is no saturation and connectivity decreases to reduce the transmission steps. Although possible, in none of our initial networks are the agents connected so as to guarantee that all agents become infected (i.e. contain an embedded tree graph with the initially infected node at the root). Although this structure could arise through transcription alone, it did not occur in the course of our experiments.

However, even with a mutation rate as low as 0.005, in every case networks emerge which manage to infect all agents. The dynamic starts with the emergence of one mutated network in one trial that manages to infect all agents. Because it does not incur the halting penalty S in {5, 25}, that trial raises that network's fitness high enough for it to be duplicated, together with further mutations that typically result in networks that infect all agents in a higher percentage of trials and/or in fewer steps. Very quickly these network variants take over the population, after which there is a slower-pace adaptation to networks that saturate faster. The networks that perform the best are then analyzed to determine the degree to which they match any of the standard networks and their associated attributes (small world property, scale free degree distribution, etc.).

Evolving Networks Created by Preferential Attachment and Decay

Justin Downes

George Mason University, USA

Many growth models have been proposed to explain the formation of a variety of types of real networks. These real networks often display a scale free structure where the distribution of node degrees follows a power-law. Often these models grow networks by adding nodes one at a time with a variable number of edges introduced from the new node to existing nodes. This growth structure makes sense for situations such as citation networks where created edges do not change over time, but fail to adequately describe the evolution of many social networks where edges are created and removed constantly even while new nodes are added.

In this study I seek to demonstrate how existing growth models can be used in situations where a network is evolved, that is existing nodes create and lose edges at each time step. Additionally I will show what the implications are when using those models are with regards to the power law distribution of node degrees. Finally I will try to show what modifications to a model are needed to evolve a graph that is generated through preferential attachment growth models.

We know that many real networks follow a scale free structure and that those networks do not have fixed edges. By developing models that can simulate this situation we can augment the incredibly difficult task of collecting
real world social network data and increase the pace at which computational social science models are developed.

DeepABM: Scalable and Efficient Agent-based Simulations via Graph Convolutional Frameworks

Ayush Chopra[1]*, Jayakumar Subramanian*[2]*, Balaji Krishnamurthy*[2]*,*
Ramesh Raskar[1]*, & Anil Kamath*[2]

[1] *Massachusetts Institute of Technology, USA*

[2] *Media and Data Science Research Lab, Adobe Experience Cloud, India/USA*

Conventional ABM frameworks such as Mesa, NetLogo, etc. have an object-oriented architecture. While conceptually appealing, these are inefficient for large populations and do not differentiate between agent transition modeling and agent behavior modeling, which makes behavior learning in these frameworks challenging.
To overcome these problems, we have developed the DeepABM framework which takes a network-centric functional architecture and is built using the concepts of graph neural networks from deep learning frameworks. Using graph convolutional networks has enabled the following key benefits in DeepABM: i) scale ABMs to large agent populations in real-time, ii) run ABMs efficiently on GPUs, and iii) enable

more efficient calibration of ABMs using gradient-based supervised machine learning instead of the status-quo randomized search methods.

To validate the framework, we adapted DeepABM for simulating COVID-19 spread (DeepABM-COVID henceforth) and studied the effect of behavioral and clinical interventions on the same. DeepABM-COVID supports various interventions (quarantine, digital exposure notification, vaccination, testing) and can scale to large agent populations (> 100,000 agents) in real-time on a GPU. In particular, DeepABM-COVID models 200 million interactions across 100,000 agents over 180 time-steps in 90 seconds as against 5 hours in an equivalent Mesa implementation. We augment DeepABM-COVID with an interactive plug-n-play interface to compare the effects of different interventions through simulations with custom choices of infection, transmission, and demographic parameters.

Finally, using DeepABM-COVID, we collaborated with leading clinicians to study the Public Health impact of delaying the second dose of the COVID-19 vaccine, and the work was recently accepted at the British Medical Journal [1].
A brief overview of the DeepABM framework can be found at [2] and an interactive demo of the framework at [3] (at min 44)

[1]: https://www.medrxiv.org/content/10.1101/2021.02.23.21252299v1
[2]: https://drive.google.com/file/d/1BQKqtRhVz4nq2zsMBD4kEMilnd72YdTq/view?usp=sharing
[3]: https://www.youtube.com/watch?v=4ZwH-KliJos

Does the Representation of Managerial Search Behavior in Agent-based Models of Organizations Matter: Satisficing vs. Hill-Climbing?

Friederike Wall

University of Klagenfurt, Austria

Agent-based models in management science build on the understanding that human decision-makers are not as gifted as traditional schools of economics suggest. Hence, decision-making agents search stepwise for potentially superior solutions, and often hill-climbing algorithms are employed for representing experiential learning and search behavior of decision-makers. However, based on experimental evidence, it has been argued that hill-climbing algorithms may be inappropriate representations of managerial search behavior (e.g., Tracy et al. 2017). Hence, the question arises in how far macro patterns emerging in agent-based models of decision-making organizations are subject to the familiar hill-climbing type of representation of managerial behavior. In other words: Would certain macro patterns also show up for another and potentially more realistic representation of managerial search behavior?

For addressing these questions, the presentation provides three components. First, the presentation introduces an algorithm (Wall 2021) that captures the core elements of Herbert A. Simon's concept of “satisficing” (Simon 1955; Simon 1959). According to empirical and experimental evidence, satisficing is a relevant representation of human search behavior (Caplin et al. 2011; Gigerenzer 2002; Güth 2010). Satisficing means a process of sequential search for options until a satisfactory level of utility is achieved. What the decision-maker regards as satisfactory is captured by the aspiration level. Depending on the difficulty of the problem as perceived by the decision-maker, the aspiration level may be subject to adaptation which also applies to the maximum number of options searched.

Second, the presentation shows that hill-climbing algorithms – as familiar in agent-based models of managerial search – could be regarded as borderline cases (or parametrizations) of the satisficing algorithm.

Third, the presentation highlights the effects of different parametrizations of the satisficing algorithm, including parameter settings that capture different variants of the hill-climbing algorithm. For this, an exemplary agent-based model based on NK fitness landscapes (Kauffman 1993; Kauffman and Levin 1987) is employed. This particular choice is made since the NK model serves as a kind of “quasi-standard” in studies on managerial search (e.g., Baumann et al. 2019; Wall 2016). In the model, artificial organizations search for superior performance to a multi-dimensional binary decision problem decomposed into distinct sub-problems, each of which is exclusively delegated to a manager. The managers seek to contribute to the organization's objective, but - due to specialization - each has a different view on the organization's decision problem. The model controls for the intra-organizational complexity meaning the interactions among subordinate managers' decision problems.

For the emerging macro patterns, the simulations reveal two-fold observations. At a relatively abstract level, the effects of complexity on organizational performance and performance enhancements go in the same direction for all parametrizations studied. However, in more detail, models' behavior may be remarkably shaped by parametrization. From a more general perspective for inverse generative social science, these two-fold observations may pose the question of when we would regard an emergent macro pattern robust to a change in the parameter settings and, thus, when a comprehensive set of rules for managerial search behavior is identified.

Machine Learning Stimulates Agent-Based Model Towards Policy-Making

Bernardo Alve Furtado

IPEA, Brazil

This presentation focuses on an empirically-based spatial-economic existing agent-based model (ABM) (Furtado 2021), specifically in a 244-unique combination of eight parameter-based rules and 37 actual parameters that are associated to 5,594 individual-run results as the basis for a discussion on the vision of Inverse Generative Social Sciences. We argue that even (a) if we were able to sufficiently describe an initial and target point for a given phenomena, (b) if we knew what the trajectory was and (c) if we could guarantee that a slightly different (previous) starting point would not lead to a different pathway and target, there exists a dense space of policy alternative interventions that lead to a bundle of solutions that are consistently and relatively more socially beneficial when compared to baseline scenarios. We aim to show that although a full description and understanding of society may be fuzzy, a superior dominant dense policy space may be identifiable, thus leading to the design of sound policy normative framework. Such line of reasoning leads to outputs that identify possible causal mechanisms that perform comparatively better despite the imperfections and incompleteness of both description and trajectories. The procedure is intended to be as follows: we will train machine learning models with a variety of methodological alternatives, but especially neural networks and genetic algorithm, that map our initial configuration of parameters and rules towards the obtained results. Once a mapping is completed, that is, once the structure of the literature-driven ABM has been codified and learned, we can amplify the space of reason of initial parameters and rules to its full extent (thus much larger than the original 244 tried configurations) and associate each new configuration to results. Among those results we can filter out the ones that construe a space of mostly social benefit – with higher output and less inequality being a prime. These selected top results will point to relevant configuration of rules (policies), thus in practice

constituting itself in a dominant space of policy effective results despite the lacking of full knowledge on description and fitness
to past, eventually random, trajectories.

Adapting an existing NetLogo radar model for Evolutionary Model Discovery compatibility

Mariah Yelenick, Matthew Koehler, David Slater, Jason Veneman, & Alex Isherwood

The MITRE Corporation, USA

The Evolutionary Model Discovery (EMD) toolkit uses genetic programming to search for agent behaviors that explain existing phenomena. EMD requires modular factor design but writing modular code may be a skill that many agent-based modelers have not learned. Modular programming is a technique in which code is divided into the smallest functional parts necessary to build large-scale functionality. The focus of this presentation is the process of converting an existing non-modular NetLogo model into a modular model that is compatible with EMD. The model simulates unmanned aerial vehicles (UAVs) flying near a spinning radar. The goal is to get the UAVs close to the location of the radar without being detected by the radar. EMD was used to find behaviors that help the UAVs achieve this goal.

Designing the factors used to compose potential agent behaviors can be difficult even when creating a model specifically for this process, and even harder if the user has existing NetLogo models they wish to use with the EMD tool. Through the radar model, best practices will be demonstrated in how to write modular code and how to convert non-modular code – skills important for writing good code outside of the EMD tool as well.

EMD results show every behavior that was tried and how well it performed. This output can be clustered to find behaviors that perform similarly or analyzed to determine which factors are most important to the success of the experiment.

This presentation will address the following questions:

- How to determine which part of a model to evolve
- How to determine useful metrics and write an appropriate fitness

function

- How to create rules/factors that are useful and flexible
    - ○ How to determine parameter and return types
    - ○ How to write modular code and why this is important to EMD
- What was learned from the UAV/radar experiment results

Taking agent-based model robustness seriously: Some ways forward

Gabriel Istrate

West University of Timisoara, Romania

"Can we find a more complete set of rules? Relatedly, how robust is the solution to a small change in the agent rules? If we could discover a "neighborhood of" agent models whose members all generate the target, the result would seem less ad hoc and unstable. These are the core concerns of the nascent field of Inverse Generative Social Science." The above description, taken from the webpage of the IGSS workshop, makes apparent the crucial role that robustness of models has in making the conclusions derived from them reliable.

In this talk, related to a Blue Sky Paper that has just appeared in the Proceedings of the 20 International Conference on Autonomous Agents and Multiagent Systems (AAMAS'2021) I will attempt to offer some ways of making the meaning of the terms in the statement above somewhat more precise.

- I will argue that formal methods based on logics (such as model checking and formal verification) are the only way to lead to principled, practical solutions to the verification and validation problem for agent-based social simulations. However, to be successful such an approach should not simply import the logical tools from the theory of formal methods, but attempt to address concerns of agent-based modelers, and instead consist of a highly unusual mix of logic and formal methods, on one hand, and sociological theories, on the other.

- I will argue that committing to such a mix will lead to new research problems. One such problem is that of defining formal tools for quantifying the "robustness" of statements in parametric logical theories with respect to parameter changes, and, more ambitiously, "the distance" between two logical statements. Such concerns seem largely absent from contemporary applications of logical methods.

- Alternatively, one may study the existence of "phase transitions" in the validity of stylized facts. I will argue that results in related theoretical areas (the study of threshold properties of random graphs using finite model theory, and that of phase transitions) provide hope for a more principled theory, that may have applications even to agent-based models.

- I will advocate the development of a theory of robustness of results about agent-based models with respect to adversarial perturbations. I will highlight some possible principles of such a theory by briefly presenting the adversarial scheduling approach to game-theoretic models of social dynamics. I will argue that such methods may exist for agent-based simulations as well.

Testing hypotheses in Agent-based models

Hang Xiong[1] *& Georgiy Bobashev*[2]

[1]*Macro Agriculture Research Institute, Huazhong Agricultural University, China*

[2]*Center for Data Science, RTI International, USA*

Computer simulations have the potential to provide a new approach to testing empirical hypotheses for social sciences studies. We present a method for conducting computational hypotheses using agent-based and microsimulation models. The computational hypothesis could be expressed in the form of structural parameters being equal to zero. Considering computational uncertainty intervals from simulation models instead of traditional confidence intervals in data-based hypothesis setting we can infer whether certain computational structures are indistinguishable from noise. The Occam razor principle becomes important especially when the number of validation observations is not very big. This approach allows one to develop efficient and simpler simulation models. We discuss the usefulness of this approach in applications to theoretical settings and practical simulation models in social science.

Specifically, for the parameter indicating a hypothesis, values that each produce the simulation result best fitting a set of observed data can be found. A confidence interval will be constructed using these values.

Whether the estimate of the parameter is significantly different from zero at a certain significance level can then indicate whether the hypothesis can be accepted. As the number of sets of observed data available is often small in social sciences studies, we also present a method to construct confidence intervals adequate for significance tests out of a small number of datasets. The legitimacy of these methods is discussed from a statistical point of view. The hypothesis testing process is illustrated using a case that examines peer effects in the diffusion of an innovation. Last, we discuss the usefulness of this approach in particular for testing interactive effects in social processes.

Generating Agent Based Models From Scratch With Genetic Programming

Rory Greig & Jordi Arranz

Improbable, United Kingdom

Program synthesis (PS) and genetic programming (GP) allow non-trivial programs to be generated from example data. Agent-based models (ABMs) are a promising field of application as their complexity at a macro level arises from simple agent-level rules. Previous attempts at using evolutionary algorithms to learn the structure of ABMs have focused on modifying and recombining existing models targeted to the domain in question, which requires prior domain knowledge. Our approach is closely related to inverse generative social science, which is a discovery process for testing multiple hypotheses automatically from data, in order to explain and illuminate observed phenomena. This technique enables us to generate models which are relatively free from existing domain priors and human preconceptions, and may shed light on completely new dynamics which have been overlooked because they are unintuitive or non-obvious. The resulting output is an interpretable symbolic model which can be understood and extended by a human modeller, so this could be used for automatically building quick prototype models before a modeller refines them. Automatically learning the model logic for ABMs is not a new idea, how- ever the existing literature has mainly focused on learning model structure by mutating an existing model which is al- ready adapted to the modelling problem in question, and works by recombining a set of “primitives” which is tailored to the domain. This is sometimes referred to as “structural calibration’.

We successfully induce the agent update rule of two different models — flocking and opinion dynamics — we employ an evolutionary algorithm which evolves a population of individuals consisting of programs expressed in a domain-specific language (DSL). In turn, these programs are combinations of the DSL operators and operands. The fitness of each individual in the population is evaluated by executing its program and comparing the output to the output of a reference model. The best programs are then copied and randomly changed — i.e. mutated. This process is repeated until the average fitness of the population reaches a desired threshold. We also perform several additional experiments to find out how well the evolved solutions perform with out-of-sample data — i.e., trajectories of the model which have not previously seen during training/evolution. We employ a flexible D S L which consists of basic mathematical building blocks. The flexibility of our method is demonstrated by learning symbolic models in two different domains: flocking and opinion dynamics, targeting data produced from reference models.

Our results clearly show how PS and GP can be used beyond model calibration to learn full symbolic representations of core model logic, by only providing reference data and without encoding previous domain knowledge. We have successfully synthesised an opinion dynamics model and a flocking model. Despite we employed a single trajectory of the reference model, in both cases the resulting models were able to generate identical macro behaviours. More importantly, we have also shown that the evolved solutions in both cases (i.e., the agent update rules) generalise very well and are highly interpretable. This improves over existing work in this area by learning models from scratch (starting from empty behaviours) and employing a generic and flexible DSL consisting of basic mathematical and conditional operations. This level of accountability offers a huge advantage over most deep learning (DL) techniques which suffer from opacity, opening the door to applications that go beyond modelling such as inverse generative social science in which the synthetic model is employed to explain and illuminate the phenomenon being modelled.

Evolutionary model discovery of behavioral factors driving decision-making in irrigation experiments

Lux Miranda, Ozlem Ozmen, & Jacopo Baggio

University of Central Florida, USA

The precise mechanisms for how and why people choose to cooperate in social-ecological commons dilemmas has garnered continual scholarly attention for decades (Dawes, 1980), yet explanations often remain incomplete or contested. The study of such phenomena grows increasingly important as more and more communities become ecologically vulnerable due to climate change (Folke, 2006). Irrigation systems are one such variety of commons dilemma characterized by actors having asymmetric access to the system's resources yet requiring collective contributions from all actors to maintain the physical irrigation infrastructure.

In Janssen & Baggio (2017), a set of agent-based models were constructed in an attempt to explain experimental data collected from a series of irrigation games with university students. In these experiments, participants were given charge of crop fields in need of watering via a river-based irrigation system. Participants were required to utilize the river's limited resources and invest tokens towards maintaining the communal irrigation infrastructure in order to maximize their individual token earnings in the round-based game.

Each agent-based model was constructed using rules derived from a different alternative behavioral theory. Each model was validated against a variety of macro-scale metrics produced by the experimental data such as the average infrastructure efficiency achieved or the average Gini coefficients for collected tokens each round. However, the study found that no single model was the best fit. Rather, several models performed roughly the same overall with each model having various strengths and weaknesses.

Evolutionary model discovery is a framework for inverse generative social science that has shown an ability to uncover causal factors driving agent decisions in agent-based models validated against empirical data (Gunaratne, 2019; Gunaratne & Garibay, 2020). In this study, we will decompose each irrigation model into the base behavioral factors they use to explain decision-making processes, namely: tendencies towards naive selfishness, naive cooperation, conditional cooperation, utility maximization, and random behavior. This creates a theoretically infinite space of possible models composed of different combinations of these factors. Evolutionary model discovery shall be used to search within this space for a model that produces agent behavior with the best fit for each of the metrics defined in Janssen & Baggio (2017). We hypothesize that utilizing evolutionary model discovery in this way will allow us to discover a model (or set of models) with greater fitness than models in the original set.

Given the "crystal box" nature of evolutionary model discovery, the resulting model has the potential to shed a more precise light on which processes agents are using to make decisions within irrigation games.

Such a result would have implications for social theory of cooperation in commons dilemmas as well as potentially provide practical insights for methodology in inverse generative social science.