Modeling agent decision and behavior in the light of data science and artificial intelligence

Abstract

Agent-based modeling (ABM) has been widely used in numerous disciplines and practice domains, subject to many eulogies and criticisms. This article presents key advances and challenges in agent-based modeling over the last two decades and shows that understanding agents' behaviors is a major priority for various research fields. We demonstrate that artificial intelligence and data science will likely generate revolutionary impacts for science and technology towards understanding agent decisions and behaviors in complex systems. We propose an innovative approach that leverages reinforcement learning and convolutional neural networks to equip agents with the intelligence of self-learning their behavior rules directly from data. We call for further developments of ABM, especially modeling agent behaviors, in the light of data science and artificial intelligence.

Keywords: agent-based modeling, modeling agent decisions and actions, artificial intelligence, machine learning, data science

Although agent-based modeling (ABM; ABMs for agent-based models) emerged as early as the 1970s (Schelling, 1971) or even earlier (W. Zhang et al., 2021), it has been extensively applied in ecology, where it is usually referred to as “individual-based modeling” (Grimm, 1999), and numerous other disciplines since the 1990s (Vincenot, 2018). Subsequently, ABM has exploded in applications (Figure 1), an indication of its usefulness across multiple sciences. Milestones in ABM development include the Overview, Design concepts, and Details (ODD) protocol for the model documentation (Grimm et al., 2010, 2020) and the Pattern-Oriented Modeling (POM) paradigm (Grimm et al., 2005). ABM received a major institutional endorsement in 2001 when it was featured by the U.S. National Academy of Sciences’ Sackler Colloquium and the resultant special issue in the *Proceedings of the National Academy of Sciences* (Bonabeau, 2002). Since then, ABM has been hailed with both enthusiasm and optimism because of its potential to create a “revolution” among the social, ecological, behavioral, and complexity sciences.

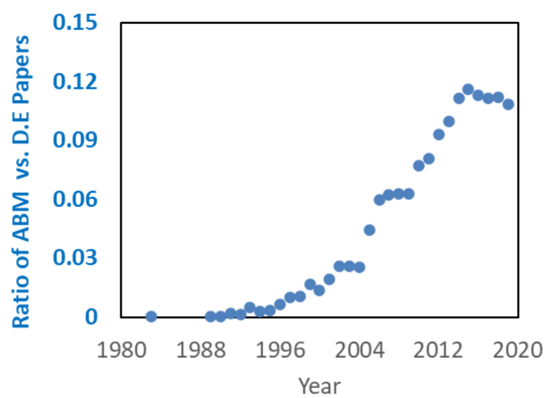

Figure 1. ABM-related papers in comparison with papers related to differential equations. The y axis is the ratio of the number of ABM papers vs. that of the differential equation (D.E) papers. For related data collection, see An, Grimm, Sullivan, Turner II, et al. (2021).

Yet, progress has been slower than initially anticipated in several critical areas of development and application (Grimm & Berger, 2016; Lorscheid et al., 2019), leading to various criticisms of ABM (Couclelis, 2002; Roughgarden, 2012). This slower progress, likely leading to the fade of enthusiasm in ABM, is embedded in the context of a big issue for ABM or any process-based modeling effort: the need to balance 1) the pattern-informed, top-down approach, which reproduces macro-level patterns without adequate explanatory power, such as the well reproduced flight patterns of hawks without an explanation of the mechanistic processes behind the patterns (Conte & Paolucci, 2014a); and 2) the theory-driven approach, which aims at generating macro-patterns from bottom-up processes. ABM developers aim to not only accurately predict or replicate the observed patterns in question, but also to understand and explain the mechanisms behind such patterns.

In this paper, we propose a third approach, based on artificial intelligence (AI) and data science, to detect, formulate, and test mechanistic processes (e.g., structures, rules, parameters) that complement the above two traditional approaches. The paper first briefly reviews historic advances in modeling human behavior, followed by an overview of the challenges of modeling agent/human behavior. Next, we show how artificial intelligence and machine learning, in combination with data science, can help reveal mechanistic processes and model agent/human behavior. In the Appendix, we provide more details about the most promising methods in relation to what we propose in this paper.

Background about modeling agent/human behavior in ABM

Historically, modeling agent/human decision-making and behavior in ABM was largely based on economic theories, such as benefit maximization or cost minimization by rational actors, and largely ignored other approaches, such as those in psychology and neurology (Groeneveld et al., 2017). Advances in relevant social and cognitive sciences have greatly enhanced the capacity of ABM to model agent (human in particular) behavior (Filatova et al., 2013; Niamir et al., 2020), generating normative models, cognitive models, and psychologically (especially emotional models) and neurologically inspired models that are instrumental for understanding and modeling human behavior (for excellent review papers, see (Balke & Gilbert, 2014), (Huber et al., 2018), and (Bourgais et al., 2018)).

Along this thread, one prominent example is the Belief-Desire-Intentions (BDI) framework. Inspired by logical and psychological principles (e.g., Michael Bratman's theory of human practical reasoning; Bratman, 1999), the BDI framework models practical reasoning and subsequent action of resource-bounded, rational agents. These agents carry 1) beliefs, which are facts agents believe about the environment, 2) desires, which involve desired end state or goals to achieve; and 3) intentions, which are intended commitments to accomplish the corresponding desires (goals). Under this framework, intentions are important elements and precursors of planned actions. Considered an improvement compared with the BDI framework, the physical, emotional, cognitive, and social factors (PECS; Conte & Paolucci, 2014; Schmidt, 2002) framework aims to explain or predict human behavior from a common deep structure, which has four fundamental elements: physical conditions, emotional state, cognitive capabilities, and social status. Once a PECS deep structure is constructed and tested as a reference model for real systems or agents, various superficial qualities can be imposed on the structure to represent a set of local, heterogenous characteristics and better predict the behavior (Schmidt, 2002).

Challenges in modeling agent behavior

ABM faces several major challenges, detailed in Appendix 1 and several assessments (An, Grimm, Sullivan, Turner, et al., 2021; Crooks et al., 2008; McDowall & Geels, 2017; O'Sullivan et al., 2016). These challenges arise from ABM's greater complexity in comparison to traditional models (Sun et al., 2016), which is the price paid for ABM's superior flexibility and capacity to capture the corresponding processes or mechanisms (Filatova et al., 2013). ABMs tend to be "data-hungry" and difficult to understand. Common solutions deployed to date include simplifying assumptions, theoretical representation of processes, and inverse parameterization using sets of observed patterns. Among all challenges that ABM faces, we highlight the seriousness of two related problems: equifinality and multifinality. ABM suffers from these two problems, although we acknowledge that they are not endemic to ABM, but involve any mechanistic modeling efforts.

The equifinality problem may blind the true pathway or mechanistic process that generates the observed macro-level pattern or outcome because the end state can be reproduced through multiple pathways (von Bertalanffy, 1968). For example, the Prisoner's Dilemma can be related to several seemingly plausible mechanisms, including group selection (Di Tosto et al., 2007), strong reciprocity (Boyd et al., 2003), tit-for-tat retaliation (Axelrod, 1997), and others (Conte & Paolucci, 2014b). While equifinality may be considered a "curse" by way of the processes involved in the outcome, it may also be viewed as a "blessing", providing a means to explore alternative explanatory pathways and falsification of existing theories. Such exploration is facilitated when data at multiple spatial, temporal, or organizational scales are available to filter out those theories that cannot explain all patterns simultaneously (the pattern-oriented modeling framework (Grimm et al., 2005)).

A related challenge is the multifinality problem, in which the same causes and/or starting conditions give rise to very different trajectories or ultimate outcomes. This problem may arise from uncertainties in key processes and/or parameters. A related challenge is that verbally formulated theories of agent behavior often leave too much room for interpretation when it comes to formalize them in an ABM so that different implementations of the same theory can lead to different results and conclusions (Muelder & Filatova, 2018).

To demonstrate ABM's unmatched potential to provide (alternative) pathways or mechanisms for explaining or predicting observed patterns as well as the need to handle the equifinality problem, we provide an exemplar ABM that aims at understanding the dynamics of firms in the U.S. (Axtell, 2001, 2015). The ABM—with absence of many parameters and assumptions included in traditional economic models—is sufficient to endogenously produce the kinds of macro-pattern dynamics in firm sizes, ages, growth rates, job tenure, wages, networks, and so on that agree with empirical data (Appendix 2).

This seemingly "successful" ABM, though providing unique, useful mechanisms, can still be questioned due to the equifinality problem. To break new ground for improving the science and art of modeling agent decision and behavior, approaches for systematic ABM development and testing, such as the "MoHuB" (Schlüter et al., 2017a) and "POM" (Grimm et al., 2005) frameworks, have been suggested. In addition, we have developed a set of guidelines for modelers and reviewers and for novices (An, Grimm, Sullivan, Turner II, et al., 2021), which include a comparison of commonly used ABM toolkits and software packages (An, Grimm, Sullivan, Turner II, et al., 2021) given the existence of 85+ platforms or toolkits for ABM (Abar et al., 2017). Also, we promote more effective ABM education and communication through several means, such as developing and sharing curricula, promoting the reusability of ABM modules (e.g., ABM "Lego" or "Mr. Potatohead" pieces), and engaging a broader ABM community in collaborative education efforts (An, Grimm, Sullivan, Turner II, et al., 2021).

These and other endeavors, though useful, have their own limitations. We therefore propose a new pathway for modeling agent decisions and behavior, which is based on artificial intelligence (AI) and data science. Over the last decade or so, AI, data science, and their applications for ABM have led to a critical mass of tools, applications, and insights so that the potential of this new pathway has become clearly visible. This paper focuses on data collection, processing, and methods or algorithms that support modelling decisions and behavior, recognizing the importance of other important issues such as verification of the mechanisms or rules thus derived.

Opportunities from artificial intelligence and machine learning

Traditional artificial intelligence leverages machines to understand and mimic human intelligence. Machine learning, an essential element of artificial intelligence, can be as simple as standard linear regression models. On the other hand, machine learning can leverage more advanced models and reveal non-linear, complex processes through, for example, neural networks, genetic algorithms, decision trees, naive Bayes, and Bayesian networks. For instance, the data-driven agent-based crowd model by Tan et al., (2019) adopts a standard differential evolution genetic algorithm to calibrate model parameters (e.g., pedestrian speed, turn angle) based on video and virtual reality (VR) experiment data. In a study that features a data-driven agent-based model (Taghikhah et al., 2022), machine learning is leveraged to identify automatically the causal relationships and derive decision rules for agents from microdata on behavior. Furthermore, machine learning can be employed to detect patterns in model output, which may help to evaluate the robustness of the model. Below we focus on illustrating the usefulness of neural networks.

Inspired by the structure of human and animal brains, neural networks have emerged as one of the most versatile algorithms in machine learning. Neural networks are increasingly employed in consequential decision-making processes in many domains such as banking, medicine, and criminal justice. By 2030, artificial intelligence could boost the global economy by $15.7 trillion, which includes massive decisions made by neural networks (West & Allen, 2018). The huge explosion of neural networks presents an unparalleled opportunity to augment individual human life, learning, intelligence, and productivity.

A neural network consists of layers of nodes that are connected by links. Here, nodes may be interpreted differently, which may be analogous to agents in the context of complex systems, variables, or decision points (Abdulkareem et al., 2019), while links could be agent-agent or agent-environment relationships (Cranmer et al., 2020; Kipf & Welling, 2016). As input data are fed into the machine learning algorithm(s), nodes receive messages from parent (sending) nodes and pass messages to their child (receiving) nodes, depending on whether some conditions are met. With sufficient data and an appropriate model structure, the trained models can offer high predictive power, offering significant opportunities to calibrate and/or validate ABMs. For instance, modelers can assign and implement each agent with its own unique regression equation or neural network (H. Zhang et al., 2016). Then the process of understanding and envisioning agent behavior entails optimizing the regression equations or neural networks for all the agents (see the example below). Models trained in this way—the behavior rules of agents in particular—are relatively rare for many reasons, such as the difficulty of independently training a large number of convolutional neural networks.

Another critical issue concerning neural networks centers on the difficulty of interpretation: such models are often like a "black box", offering little or no understanding of the mechanisms governing the processes. Below we propose a reinforcement learning (RL) plus convolutional neural network (CNN) based approach (i.e., RL-CNN approach) to equip agents with the intelligence of self-uncovering and self-learning behavior mechanisms instead of relying on the modeler to "hardcode" (W. Zhang et al., 2021) behavioral rules beforehand. Among the three ways machine learning can contribute to ABM analysis (i.e., prior to running the ABM, during the running of the ABM, and post running the ABM to analyze ABM output; (Abdulkareem et al., 2019)), the one related to empowering agents to self-learn and formulate mechanisms during the running of the ABM is most challenging (e.g., computationally intense) and promising. The most common practice is that ABM modelers hardcode agents' behavioral rules prior to running ABM (W. Zhang et al., 2021). In a recent multi-agent model integrated with reinforcement learning, effective preventive maintenance policies (i.e., rules governing agent actions) can be learned directly from data without any knowledge about the environment and maintenance strategies, ensuring smooth and efficient production for large-scale manufacturing systems (Su et al., 2022).

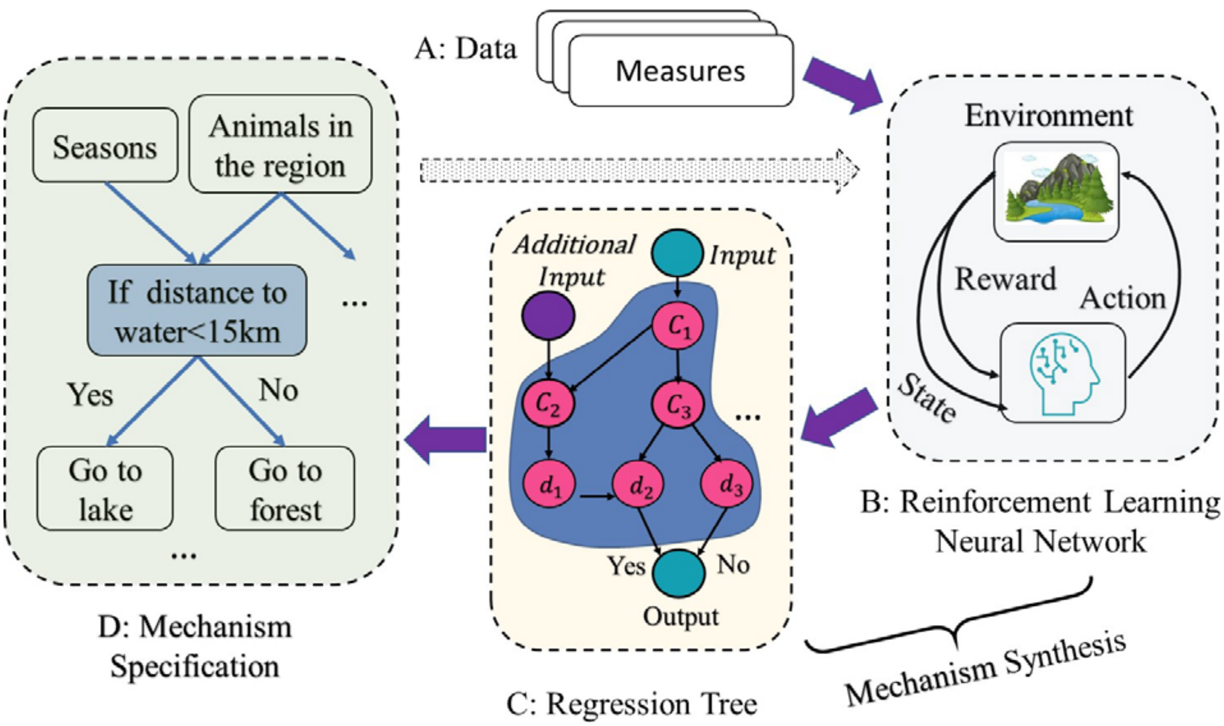

Figure 2. Illustration of reinforcement learning in deciphering animal behavior rules under various environmental conditions. The node "if distance to water < 15 km" (within the blue box in Panel D) comes from the multiple nodes and links in the blue area of Panel C.

Traditional machine learning is powerful in understanding and simulating agents decision-making and behavior, but tends to suffer from insufficient data and/or data-handling capabilities (Gil & Selman, 2019) to identify the correct model structure and parameters and therefore appropriately calibrate ABMs (Srikrishnan & Keller, 2021). The advent of data science and its methods, tools, and data infrastructures has powerfully enriched machine learning to derive processes behind patterns of interest, verifying or rebutting the underlying hypothetical mechanisms behind such patterns. Reinforcement learning (RL), through a certain set of reward and penalty rules, is a promising tool in this regard (Su et al., 2022). Specifically, RL can be assigned to the agents under investigation. With little or no pre-knowledge about such mechanisms, RL-enabled agents can “learn” the best behavioral rules from data so that the learned “rules” can maximizes the RL’s reward (or minimize the penalty) when dealing with other agents and the environment. One successful RL application is the multiagent system RL (MARL) (Buşoniu et al., 2010), under which a computer Go program called AlphaGo is developed and can beat a human professional player on a full-sized board (Silver et al., 2016); recently a newer version called KataGo can even beat world-class human Go players (Edwards, 2022).

Take an example of theorizing from (or seeking mechanisms of) animal behavioral science as shown in Figure 2. We begin with RL without pre-knowledge or hypothesis on the mechanisms (the term mechanism is often called policy in the Machine Learning domain). As data (Panel A) are used as input to train the RL neural network (a built-in capacity of each agent; Panel B), the agent’s RL neural network can then learn and establish a set of nodes and links, which can maximize the reward function with compliance to the state (for detail about state see Appendix 4). To reveal the thus established, yet hidden nodes and links, a regression tree (Panel C) can be used to “translate” them into a set of visible decision tree links (arrows in Panel C) and nodes (e.g., C_1, $C_2 \ldots d_3$ in Panel C). In turn, these nodes and links in the decision tree, with the aid of domain knowledge, can be used and interpreted as meaningful and understandable mechanisms (Panel D), helping theorize and understand the processes generating the macro patterns (e.g., data in Panel A). Alternatively, the above process may start with Panel D, where we have pre-knowledge or hypotheses regarding the mechanisms of interest that need to be verified or polished. In this case, the RL process starts from both data (Panel A) and such specified mechanisms (Panel D), where the dashed arrow indicates the “extra” input to train the RL network in Panel B. All the remaining steps remain the same as above. The outcome is that the pre-knowledge or hypotheses regarding the mechanisms—including parameters and structure—may be partially or fully modified according to the nodes and links in the decision tree. For instance, the parameter 15 km in Panel D may be changed to be 20 km, and “Go to lake” to “Stay where it is”.

The above example takes the data for granted, which may or may not reflect the actual conditions. We may leverage a so-called convolutional neural network (CNN), a data extraction method (see Appendix 3 for detail), to prepare data that are useable in the above RL procedure (Figure 2). In the above example, CNN can be leveraged to identify/detect animals based on images from different sources (e.g., GPS collars or drone imaging).

The RL-CNN approach, though promising and exciting, does not imply that AI, machine learning and data science are not unbiased, nor does it exhaust the potentials that AI and machine learning can contribute to modelling agent behavior. First, we still emphasize the importance of domain knowledge and theory that are obtained elsewhere (Taghikhah et al., 2022). The mechanism specification in Panel D of Figure 2, if employed as a starting point for RL network (Panel B), reflects this importance. The mechanisms or rules thus derived—for example, cause-effects and feedback loops in many instances—should be subject to continued examination by domain knowledge and theory. Also, as new data become available, the above RL-CNN or other approaches should be continually used to polish or revise existing rules, even establish new rules. Therefore, continuous real-time data collection is important for not only deriving, but also for validating and renewing, such rules. The concept of "Digital Twins" (DT) is based on this idea of updating, in regular intervals, the data underlying a realistic model used for forecasting. This principle is well-known from weather forecast and widely used in industry (Singh et al., 2022), but has also become the basis of large initiatives to support decision making regarding climate, ocean, and biodiversity, such as the Destination Earth program of the European Commission (Nativi et al., 2021).

While neural networks and RL are among the most flexible and powerful tools, there are many other useful AI and machine learning algorithms. For instance, it is reported that Bayesian networks (Abdulkareem et al., 2019) and artificial neural networks (van der Hoog, 2019) represent viable alternatives for small training datasets. Such alternatives are illustrated here by an example regarding Graph Neural Networks (GNNs), which have recently emerged to link nodes horizontally and improved predictive tasks. In this context, a graph is a structure (frequently a mathematical function) that models pairwise relations between nodes, in which all nodes (agents) are connected by edges or links. In one recent application, GNN was leveraged to derive successfully the closed-form, symbolic expression of Newton's law of motion based on data from the experiment. Simply put, the machine-mining approach can be used to exactly "recover" Newton's formula $F = G\frac{m_1 m_2}{r^2}$ without any previous clue or assumption regarding its form. Note that F, G, m_1, m_2, and r represent the force between Particles 1 and 2, the gravitational constant, the mass of Particle 1, the mass of Particle 2, and the distance between the two particles (Cranmer et al., 2020). This success has boosted AI's potential to recover laws or mechanisms in other domains: we present a potential way, as an example, to recover mechanisms or behavioral rules in agent-based complex systems (see Appendix 5).

Opportunities from new forms of data

Traditional AI's capability to nourish ABM rules is also constrained because new forms of data, including data in high volumes, are either unavailable or too difficult to handle using traditional data processing and analytic methods. In applications, machine learning will be much more empowered if aided with some non-traditional datasets such as big data or qualitative data. Such challenges are effectively addressed with recent advances in data science.

Big data have several unique features that distinguish them from traditional data, largely in terms of huge volume, high velocity, wide variety, variable veracity, and value. Big data are increasingly nourishing quick detection and understanding of processes or patterns in many scientific fields (De Mauro et al., 2016). On the other hand, qualitative data could provide essential insights into understanding the above processes or patterns. Qualitative data take the form of text, images, videos, audio documents, and the like. Yet both big data and qualitative data are very different when compared to such traditional data as census data and survey data (Marcus, 2018).

For example, in an instance of social-sensing analysis of the impacts of disasters (C. Zhang et al., 2020), Twitter data are used to reveal the dynamic emotions, e.g., disgust, fear, joy, sadness, anger, and surprise, in relation to a hurricane outbreak and related rescue activities in Houston, TX during August 25–30, 2017. Numeric emotion scores are derived from tweets describing certain types of events (e.g., help and rescue events) or flood-control infrastructures. These emotion scores, expressed as the relative abundance of words related to a certain emotion out of all words, can be used to verify or debut related ABM rules or outcomes (C. Zhang et al., 2020). Such data can also help in the above mechanism retrieval steps. For instance, the emotional scores can help at Step #1 (see Appendix 5) by ruling out some unrealistic functions, or at Step #3 by casting out unreasonable outcomes (and the corresponding functions at Step #1).

Conclusion

With the advent of the digital industrial revolution, new technologies and data forms are exploding in biophysical, human, Anthropocene, and many other realms. Among these, artificial intelligence and data science (machine learning in particular) should be among the top priority areas for future research, which will likely bring in revolutionary impacts on the science and technology addressing agent decisions and behaviors in complex systems. It must be pointed out that we do not downplay the importance of traditional scientific investigations and the related findings. On the contrary, the artificial intelligence and data science approach should build on and complement such traditional investigations through, for example, experimenting, fieldwork, inductive and deductive reasoning, hypothesis testing, and theorization, and vice versa. For instance, the data (Figure 2A) and pre-knowledge / hypotheses (Figure 2D) may come from traditional investigations.

At the same time, it is worth emphasizing the unique potential of this artificial intelligence and data science approach to detecting internal, theory-relevant mechanisms. For instance, the links and nodes in the decision tree (Figure 2C), "translated" from the hidden network (Figure 2B), may reveal unique factors, structures (e.g., causal relationships), and parameter values (Figure 2D) that would not be imagined and/or included in traditional scientific investigations and will likely be used to stimulate/formulate new theory development or improve existing theory. We do not intend to say that such factors, structures, and parameters are completely free of bias and "right". Instead, we seek to provide alternative (related to traditional scientific investigations) thinking and modeling choices. Therefore, these innovative approaches will likely pave unprecedented ways for not only formulating agent behavior mechanisms or rules, but also forming new, more robust theories or rebutting existing theories (thus making equifinality less problematic). This approach may also be conducive to better understanding "commonalities and differences between theories" and addressing the "degree of formalization" problems (Schlüter et al., 2017b).

There is abundant literature regarding pathways to "uncover" or formulate mechanisms or rules behind agent behavior or decisions, such as the Inverse Generative Social Science (Vu et al., 2019) and the Mr. Potatohead (Parker et al., 2008) frameworks. Correspondingly, there exist a large amount of AI and data science tools, algorithms, or models we can leverage; for good reviews in this regard, we refer to W. Zhang et al. (2021). In the context of such literature and tools, this position paper does not seek to provide a comprehensive review of them. Instead, we aim to call for more attention and efforts towards uncovering agent decision and behavior mechanisms in the light of data science and artificial intelligence. Towards using and advancing this AI and data science approach, barriers may exist for many reasons, such as its demanding computational power, difficulties in multi- and inter-disciplinary learning, conversing, and understanding, and coding some vague theoretical frameworks (Muelder & Filatova, 2018). However, we envision this approach will be increasingly recognized, used, and advanced in many aeras of research and practical applications related to understanding agent behavior and decision-making.

Appendix:

1. ABM challenges

An and colleagues have identified a set of ABM challenges in a recent publication (An, Grimm, Sullivan, Turner, et al., 2021). These challenges include developing integrated human-environment ABMs, modeling human behavior, and building spatially explicit ABMs, particularly when considering the "telecoupling" effects (Liu et al., 2014). Challenges also abound in addressing module reusability and transparency, model verification and validation (a challenge to all types of model, not ABM alone), high-performance computing, and so on.

2. Axtell's agent-based model

ABM can be used to formulate new theory. We illustrate this kind of application by explaining high levels of turnover (e.g., individual workers' job changes; firm termination or start-up) in the American private sector with stationary firm size distributions (Axtell, 2001). To explain such high dynamism, conventional economic models—which assume general equilibrium in the economy—must resort to factors such as exogenous shocks and firm-specific (e.g., technological or productivity-related) variables, along with product markets, prices, and consumption patterns.

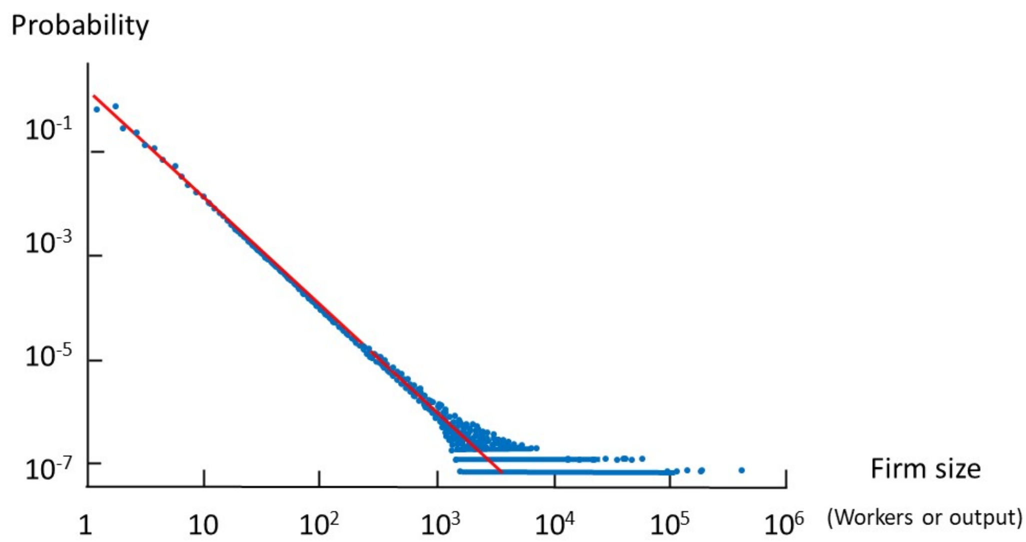

Figure S1: The Zipf's law fits almost the entire distribution of firm sizes measured by employees and output. The blue dots are produced by the ABM, which fits a straight trend line (red) (Axtell, 2015). This pattern is observed in empirical data (1997) from the U.S. Census Bureau (Axtell, 2001).

The ABM developed by Axtell (2015) simulates 120 million workers (agents) in the U.S. private sector, where each agent pursues their own self-interest by seeking information about alternative jobs within a limited social network (two to six friends). The ABM is sufficient to endogenously produce the kinds of macro-pattern dynamics in firm sizes, ages, growth rates, job tenure, wages, networks, etc. that closely resemble empirical data that follow power law distributions (Axtell, 2001) (Figure S1)—without resorting to the aforementioned conventional theories or variables. We may have some confidence in the validity of this finding under the rule of minimality, also known as Occam's razor: the simpler explanation wins if it can account for a phenomenon in the same way as a more complex one. Yet this rule *per se* needs further proof or justification—especially in the complex system domain.

3. Neural networks

Neural networks are increasingly leveraged in many social domains. In 2015, Potash and associates used a machine learning model named random forests to predict children's risk of lead poisoning, collaborating with the Chicago Department of Public Health (Potash et al., 2015). This proposed warning system is designated to detect and remediate lead hazards before the potential adverse effect emerges. In 2016, Carton et al. proposed a machine learning method to detect police officers at the risk of adverse behavior, including unjustified use of force or shootings and sustained complaints, to promote several inventions such as training and counseling (Carton et al., 2016). Athey and Wager employ a modification of random forests to predict heterogeneous treatment effects with the model trained using data from The National Study of Learning Mindsets (Athey & Wager, 2019).

While the neural network technology holds great promise, its convenience has often distracted us from its side effects. Neural network is often considered as mysterious or unknowable, called black boxes since researchers who design the architecture do not know how or why they work well. Some research efforts have been spent on demystifying a neural network to understand what it has learned. The neural network is modeled like the human brain, which has been implemented into layers of interconnected nodes, named neurons, to process data. For example, suppose the researcher built a neural network to recognize animals from the camera. Then the neural network might arrange specific neurons to detect ears of foxes and some other neurons to detect the mouth. Some researchers (e.g., Zewe, 2022) propose a method called MILAN (mutual information-guided linguistic annotation of neurons) to automatically demystify functions of all neurons in a network by generating descriptions to fulfill the requirements. It is essential because one neural network consists of hundreds of thousands of individual neurons.

The main components of the convolutional neural network (CNN) are input layers, convolution and pooling layers, and multi-layer perceptron. A layer is a group of neurons with the same operation; a perceptron is a function that combine neutrons in the earlier layer to form a set of new neutrons; put another way, a perceptron is the layer that has the full connectivity with all neurons and helps to map the representation between input and output. The convolution layer is the core block of the CNN, which performs a dot product between two matrices: input and weights. The pooling layer is applied to replace the output of the network to derive a summary statistic of the nearby output. This layer helps reduce the spatial size of the representation. There are several blocks in these layers: tensor, neuron, layer (comprised of many neurons), kernel weights, and biases. A tensor is an n-dimensional matrix to represent data in deep learning. A neuron is a function (often a linear combination of some matrices) that takes several inputs and generates a single output. A layer is group of neurons with the same operation. Kernel weights and biases are unique to each neuron to allow a classifier to yield to environment.

In Figure S2, the architecture of a neural network from the “You only look once” (YOLO, an ML algorithm that quickly detects and classifies objects) family is presented. The neural network processes the information through the process of “convolution-activation-pooling”, where convolution is a math operation that filters irrelevant information. To interpret the CNN, there are some visual explanation methods. For image classification task, a visual explanation from the model for justifying and target category requires to satisfy two properties: class discriminative and high resolution. The first property enables localization of different regions in the input image to contribute to different output classes (what does a class do?). The second property enables capturing fine-grained detail. The gradient-weighted class activation mapping (Grad-CAM) (Suresh, 2021) employs the gradients of any object concepts (say, an airplane in a classification network) input into the final convolutional layer to generate a coarse localization map marking the important areas in the image. As an example of this method, an airplane image is an input into the convolutional neural network. By applying this method, a coarse heatmap is generated to visualize the different layers' functions. In Figure S2, we present the coarse heatmap procedure by (Suresh, 2021) to visualize the functions of the neural network.

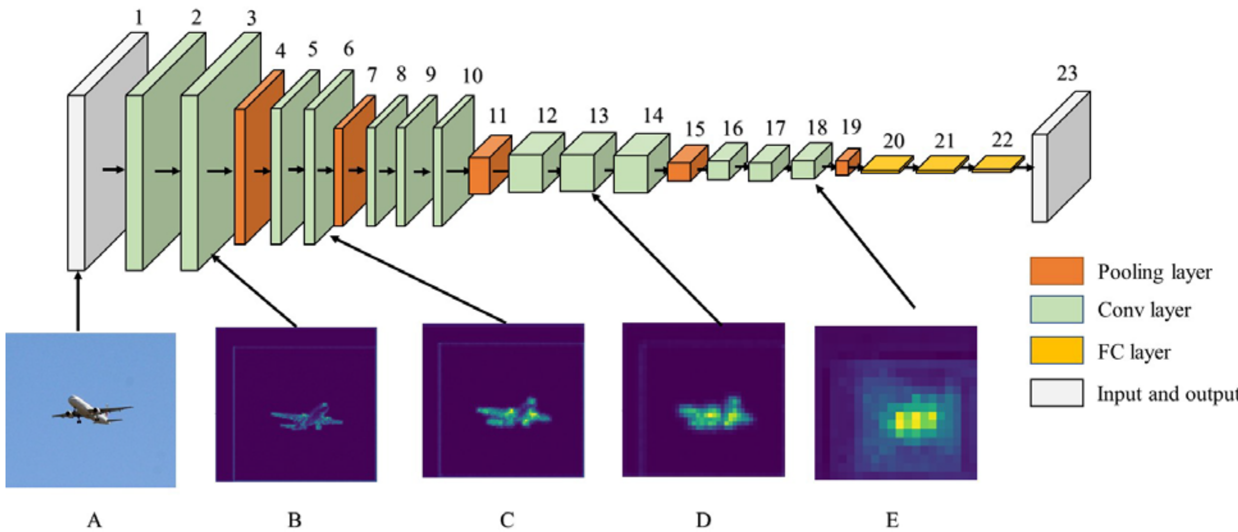

Figure S2: The architecture of convolutional neural network Yolo family.

A convolutional layer is a filter that removes some redundant (compared to the major task) information. For example, the background info (blue sky) is filtered out as it does not help identification of the target (airplane). In Figure S2, layers 2, 3, 5, 6, 8~10, 12~14, and 16~18 belong to such layers. The pooling layers (orange; layers 4, 7, 11, 15, and 19) calculate the mean, max, and min and summarize what is in the previous layer. Through the pooling process, the size of the previous image is decreased, enabling a focus on the elements that are meaningful. The fully connected (FC) layer is a multi-layer perceptron, which adds up all convolutional layers, resulting in a full representation. Images B ~E are the outcome images after implementing all the convolutional and pooling layers. Images A and B should have the same size, while images C, D, and E should be smaller in size as the pooling layers are used to focus on a smaller area (we present them at the same size for illustration purpose with a sacrifice in image resolution).

4. Reinforcement Learning

Under RL, the neural network developer does not need to develop and adjust a model to fit in the environment. Here the environment, more than natural environment, refers to all context other than itself—including all other objects or agents. Instead, the agent learns from the environment and adopts the best behavior using two types of RLs: model-based algorithm and model-free algorithm. The model-based algorithm employs experience to construct an internal model, which can calculate the agent's transitions and outputs in the environment. Then, the internal model chooses the best action to respond to the environment. The agent can externally receive the rewards from the probability function and state (state refers to components of a certain behavior or set of behaviors) transition. In contrast, the model-free algorithm uses experience from the environment to learn the policy or function of values (behavior rules of agents) directly from the data without a predefined model. The agent only needs to recognize its possible states and actions in an environment without knowing the state transition and reward function. The agent decides what to do according to some generic function or constraints, which can be determined by the data and the emergent outcomes. Finally, the modeler can unveil each agent's RL structure and parameters, extracting the common factors and relationships (from which we can develop/verify theory). Such factors and relationships may help detect hidden structures and/or variables, support or rebut existing hypotheses, and/or formulate new theories.

Reinforcement Learning (RL) may be interpreted in three different ways: using another method to generate explanations, introducing a new learning model called intrinsic model, and a combination method by changing the original model (the RL) to improve interpretability. Among them, a very popular method to allow non-experts to understand RL is through the use of a decision tree-based structure (Humbird et al., 2019), among others. This method leverages the hierarchical decision structure that can encode user experience into an interpretable structure. In this way, each rule or policy (i.e., part of the mechanism in Figure 2) can be mapped to decision nodes in a tree. Based on this method, combinational decision trees with natural language initializations and explanations have been proposed and approved within more conducive for specific policy (i.e., part of the mechanism in Figure 2). A novel collaborative framework to interpret an autonomous agent's behavior rooted in principles of human-centered design has been recently proposed (Tambwekar et al., 2022). This work employs a novel deep learning (which refers to a class of machine learning algorithms that use a cascade of multiple layers of nonlinear processing units for feature extraction and transformation) framework, named HAN2Tree, and a differential decision tree (DDT) model to represent the policy with policy gradients (connectionist) given a user-defined task completion signal (i.e., a reward function) (Tambwekar et al., 2022). Then, the unstructured natural languages from humans initialize the behavior are converted to lexical decision trees. Lastly, explanations of learned policy in multiple modules are provided to users.

5. Uncover mechanisms using a three-step GNN approach

To uncover an important law in a complex system, we can use a three-step strategy. Step #1: build an edge model to represent links or edges (i.e., messages) amongst all agents. Having formed potential candidate sets of equations or functions (e.g., those conforming to the BDI theory or PECS framework) with input from domain scientists, we may select one through a means of deep learning, such as graph neural network. Using this edge model, we conduct Step #2: develop a node model, in which each node (agent) receives messages from other agents, with the magnitude of each message calculated from the candidate function obtained in Step #1. Finally, Step #3: establish a global model to aggregate and update the status of all messages and nodes over time based on the edge and node models chosen above. Once a function (plus number of parameters as part of complexity) from the edge model (Step #1) is examined, the status of all agents can be calculated and compared to the observed time series of agent data. The function (or a set of functions) with the least measure of some index (e.g., normalized mean square error) can be chosen among the different alternative functions (plus parameter values) under the rule of minimality, such as the KISS principle, "keep it simple, stupid" (Axelrod, 1997; Kemerer, 1995), or other more sophisticated rules if necessary. Note that to handle actions of less interest or importance, black-box approaches may suffice, saving the hassle of following the three steps outlined above.

Acknowledgements

We are indebted to financial support from the National Science Foundation (NSF) through the Method, Measure & Statistics and Geography and Spatial Sciences (BCS #1638446) and the Dynamics of Integrated Socio-Environmental Systems programs (BCS 1826839 and DEB 1924111). We thank the participants of the ABM 17 Symposium (sponsored by the above NSF grant; http://complexities.org/ABM17/) for input and comments. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 757455) and through an ESRC/Alan Turing Joint Fellowship (ES/R007918/1).

References

Abar, S., Theodoropoulos, G. K., Lemarinier, P., & O'Hare, G. M. P. (2017). Agent Based Modelling and Simulation tools: A review of the state-of-art software. *Computer Science Review*, *24*, 13–33. https://doi.org/10.1016/j.cosrev.2017.03.001

Abdulkareem, S. A., Mustafa, Y. T., Augustijn, E.-W., & Filatova, T. (2019). Bayesian networks for spatial learning: A workflow on using limited survey data for intelligent learning in spatial agent-based models. *GeoInformatica*, *23*(2), 243–268.

An, L., Grimm, V., Sullivan, A., Turner, B. L. I., Malleson, N., Heppenstall, A., Vincenot, C., Robinson, D., Ye, X., Liu, J., Lindkvist, E., & Tang, W. (2021). Challenges, tasks, and opportunities in modeling agent-based complex systems. *Ecological Modelling*, *457*, 109685. https://doi.org/10.1016/j.ecolmodel.2021.109685

An, L., Grimm, V., Sullivan, A., Turner II, B. L., Wang, Z., Malleson, N., Huang, R., Heppenstall, A., Vincenot, C., Robinson, D., Ye, X., Liu, J., Lindvist, E., & Tang, W. (2021). Agent-based complex systems and agent-based modeling. *Working Paper of Center for Complex Human-Environment Systems, San Diego State University*. http://complexities.org/Photo&PDF/ACS-ABM-document.pdf

Athey, S., & Wager, S. (2019). Estimating Treatment Effects with Causal Forests: An Application. *Observational Studies*, *5*(2), 37–51. https://doi.org/10.1353/obs.2019.0001

Axelrod, R. (1997). *The complexity of cooperation: Agent-based models of competition and collaboration*. Princeton University Press.

Axtell, R. L. (2001). Zipf distribution of U.S. firm sizes. *Science*, *293*, 1818–1820.

Axtell, R. L. (2015). *Endogenous dynamics of multi-agent firms* (Working Paper No. 16-39). George Mason University. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2827059

Balke, T., & Gilbert, N. (2014). How Do Agents Make Decisions? A Survey. *Journal of Artificial Societies and Social Simulation*, *17*(4), 13. https://doi.org/10.18564/jasss.2687

Bonabeau, E. (2002). Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences*, *99*(suppl 3), 7280–7287.

Bourgais, M., Taillandier, P., Vercouter, L., & Adam, C. (2018). Emotion Modeling in Social Simulation: A Survey. *Journal of Artificial Societies and Social Simulation*, *21*(2), 5. https://doi.org/10.18564/jasss.3681

Boyd, R., Gintis, H., Bowles, S., & Richerson, P. J. (2003). The evolution of altruistic punishment. *Proceedings of the National Academy of Sciences*, *100*(6), 3531–3535. https://doi.org/10.1073/pnas.0630443100

Bratman, M. E. (1999). *Intention, Plans, and Practical Reason*. CSLI Publications.

Buşoniu, L., Babuška, R., & De Schutter, B. (2010). Multi-agent Reinforcement Learning: An Overview. In D. Srinivasan & L. C. Jain (Eds.), *Innovations in Multi-Agent Systems and Applications—1* (pp. 183–221). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-14435-6_7

Carton, S., Helsby, J., Joseph, K., Mahmud, A., Park, Y., Walsh, J., Cody, C., Patterson, C. E., Haynes, L., & Ghani, R. (2016). Identifying Police Officers at Risk of Adverse Events. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 67–76. https://doi.org/10.1145/2939672.2939698

Conte, R., & Paolucci, M. (2014a). On agent-based modeling and computational social science. *Frontiers in Psychology*. https://doi.org/10.3389/fpsyg.2014.00668

Conte, R., & Paolucci, M. (2014b). On agent-based modeling and computational social science. *Frontiers in Psychology*, *5*(000668). https://doi.org/10.3389/fpsyg.2014.00668

Couclelis, H. (2002). Why I no longer work with agents: A challenge for abms of human-environment interactions. In D. C. Parker, T. Berger, & S. M. Manson (Eds.), *Meeting the Challenge of Complexity: Vol. 1.1* (pp. 14–16). Center for Spatially Integrated Social Science University of California at Santa Barbara. http://www.csiss.org/events/other/agent-based/additional/proceedings.pdf

Cranmer, M., Sanchez-Gonzalez, A., Battaglia, P., Xu, R., Cranmer, K., Spergel, D., & Ho, S. (2020). Discovering symbolic models from deep learning with inductive biases. *ArXiv:2006.11287 [Cs.LG]*. https://arxiv.org/abs/2006.11287

Crooks, A., Castle, C., & Batty, M. (2008). Key challenges in agent-based modelling for geo-spatial simulation. *Computers, Environment and Urban Systems*, *32*(6), 417–430. https://doi.org/10.1016/j.compenvurbsys.2008.09.004

De Mauro, A., Greco, M., & Grimaldi, M. (2016). A formal definition of Big Data based on its essential features. *Library Review*. https://doi.org/10.1108/LR-06-2015-0061

Di Tosto, G., Paolucci, M., & Conte, R. (2007). Altruism among simple and smart vampires. *International Journal of Cooperative Information Systems*, *16*(1), 51–66.

Edwards, B. (2022). New Go-playing trick defeats world-class Go AI—but loses to human amateurs. *ARS TECHNICA*. https://arstechnica.com/information-technology/2022/11/new-go-playing-trick-defeats-world-class-go-ai-but-loses-to-human-amateurs/

Filatova, T., Verburg, P. H., Parker, D. C., & Stannard, C. A. (2013). Spatial agent-based models for socio-ecological systems: Challenges and prospects. *Thematic Issue on Spatial Agent-Based Models for Socio-Ecological Systems*, *45*, 1–7. https://doi.org/10.1016/j.envsoft.2013.03.017

Gil, Y., & Selman, B. (2019). A 20-year community roadmap for artificial intelligence research in the US. *ArXiv.Org*, arXiv:1908.02624.

Grimm, V. (1999). Ten years of individual-based modelling in ecology: What have we learned and what could we learn in the future? *Ecological Modelling*, *115*(2–3), 129–148.

Grimm, V., & Berger, U. (2016). Structural realism, emergence, and predictions in next-generation ecological modelling: Synthesis from a special issue. *Ecological Modelling*, *326*, 177–187.

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.

Grimm, V., Railsback, S. F., Vincenot, C., Berger, U., Gallagher, C., DeAngelis, D., Edmonds, B., Ge, J., Giske, J., Groeneveld, J., Johnston, A. S. A., Miles, A., Nabe-Nielson, J., Polhill, J. G., Radchuk, V., Rohwader, M.-S., Stillman, R. A., Theile, J., & Ayllon, D. (2020). The ODD protocol for describing agent-based and other simulation models: A second update to improve clarity, replication, and structural realism. *Journal of Artificial Societies and Social Simulation*, *23*(2). USGS Publications Warehouse. https://doi.org/10.18564/jasss.4259

Grimm, V., Revilla, E., Berger, U., Jeltsch, F., Mooij, W. M., Railsback, S. F., Thulke, H.-H., Weiner, J., Wiegand, T., & DeAngelis, D. L. (2005). Pattern-oriented modeling of agent-based complex systems: Lessons from ecology. *Science*, *310*(5750), 987–991.

Groeneveld, J. J., Müller, B., Buchmann, C. M., Dressler, G., Guo, C., Hase, N., Hoffmann, F., John, F., Klassert, C. J., Lauf, T., Liebelt, V., Nolzen, H., Pannicke, N., Schulze, J., Weise, H., & Schwarz, N. (2017). Theoretical foundations of human decision-making in agent-based land use models – A review. *Environmental Modelling & Software*, *87*, 39–48.

Huber, R., Bakker, M., Balmann, A., Berger, T., Bithell, M., Brown, C., Grêt-Regamey, A., Xiong, H., Le, Q. B., Mack, G., Meyfroidt, P., Millington, J., Müller, B., Polhill, J. G., Sun, Z., Seidl, R., Troost, C., & Finger, R. (2018). Representation of decision-making in European agricultural agent-based models. *Agricultural Systems*, *167*, 143–160. https://doi.org/10.1016/j.agsy.2018.09.007

Humbird, K. D., Peterson, J. L., & Mcclarren, R. G. (2019). Deep Neural Network Initialization With Decision Trees. *IEEE Transactions on Neural Networks and Learning Systems*, *30*(5), 1286–1295. https://doi.org/10.1109/TNNLS.2018.2869694

Kemerer, C. F. (1995). Software complexity and software maintenance: A survey of empirical research. *Annals of Software Engineering*, *1*(1), 1–22. https://doi.org/10.1007/BF02249043

Kipf, T. N., & Welling, M. (2016). Semi-supervised classification with graph convolutional networks. *ArXiv:1609.02907 [Cs.LG]*.

Liu, J., Moran, E., Hertel, T., Vina, A., & Nichols, S. (2014). *Complex dynamics of telecoupled human and natural systems*.

Lorscheid, I., Berger, U., Grimm, V., & Meyer, M. (2019). From cases to general principles: A call for theory development through agent-based modeling. *Ecological Modelling*, *393*, 153–156. https://doi.org/10.1016/j.ecolmodel.2018.10.006

Marcus, G. (2018). Deep learning: A critical appraisal. *ArXiv.Org*, arXiv:1801.00631.

McDowall, W., & Geels, F. W. (2017). Ten challenges for computer models in transitions research: Commentary on Holtz et al. *Environmental Innovation and Societal Transitions*, *22*, 41–49. https://doi.org/10.1016/j.eist.2016.07.001

Muelder, H., & Filatova, T. (2018). One Theory—Many Formalizations: Testing Different Code Implementations of the Theory of Planned Behaviour in Energy Agent-Based Models. *Journal of Artificial Societies and Social Simulation*, *21*(4), 5. https://doi.org/10.18564/jasss.3855

Nativi, S., Mazzetti, P., & Craglia, M. (2021). Digital Ecosystems for Developing Digital Twins of the Earth: The Destination Earth Case. *Remote Sensing*, *13*(11). https://doi.org/10.3390/rs13112119

Niamir, L., Ivanova, O., & Filatova, T. (2020). Economy-wide impacts of behavioral climate change mitigation: Linking agent-based and computable general equilibrium models. *Environmental Modelling & Software*, *134*, 104839. https://doi.org/10.1016/j.envsoft.2020.104839

O'Sullivan, D., Evans, T., Manson, S., Metcalf, S., Ligmann-Zielinska, A., & Bone, C. (2016). Strategic directions for agent-based modeling: Avoiding the YAAWN syndrome. *Journal of Land Use Science*, *11*(2), 177–187.

Parker, D. C., Entwisle, B., Rindfuss, R. R., Vanwey, L. K., Manson, S. M., Moran, E., An, L., Deadman, P., Evans, T. P., Linderman, M., Rizi, S. M. M., & Malanson, G. (2008). Case studies, cross-site comparisons, and the challenge of generalization: Comparing agent-based models of land-use change in frontier regions. *Journal of Land Use Science*, *3*(1), 41–72.

Potash, E., Brew, J., Loewi, A., Majumdar, S., Reece, A., Walsh, J., Rozier, E., Jorgenson, E., Mansour, R., & Ghani, R. (2015). Predictive Modeling for Public Health: Preventing Childhood Lead Poisoning. *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2039–2047. https://doi.org/10.1145/2783258.2788629

Roughgarden, J. (2012). *Individual based models in ecology: An evaluation, or how not to ruin a good thing*. Philosophy of Science Association Biennial Meeting, San Diego, CA. http://philsci-archive.pitt.edu/9434/1/RoughgardenPSA2012IBMLecture.pdf

Schelling, T. C. (1971). Dynamic Models of Segregation. *Journal of Mathematical Sociology*, *1*, 143–186.

Schlüter, M., Baeza, A., Dressler, G., Frank, K., Groeneveld, J., Jager, W., Janssen, M. A., McAllister, R. R. J., Müller, B., Orach, K., Schwarz, N., & Wijermans, N. (2017a). A framework for mapping and comparing behavioural theories in models of social-ecological systems. *Ecological Economics*, *131*, 21–35.

Schlüter, M., Baeza, A., Dressler, G., Frank, K., Groeneveld, J., Jager, W., Janssen, M. A., McAllister, R. R. J., Müller, B., Orach, K., Schwarz, N., & Wijermans, N. (2017b). A framework for mapping and comparing behavioural theories in models of social-ecological systems. *Ecological Economics*, *131*, 21–35.

Schmidt, B. (2002). *The modelling of human behaviour: The PECS reference model* (A. Verbraeck & W. Krug, Eds.). SCS Europe Bvba.

Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., & Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, *529*(7587), 484–489. https://doi.org/10.1038/nature16961

Singh, M., Srivastava, R., Fuenmayor, E., Kuts, V., Qiao, Y., Murray, N., & Devine, D. (2022). Applications of Digital Twin across Industries: A Review. *Applied Sciences*, *12*(11). https://doi.org/10.3390/app12115727

Srikrishnan, V., & Keller, K. (2021). Small increases in agent-based model complexity can result in large increases in required calibration data. *Environmental Modelling & Software*, *138*, 104978. https://doi.org/10.1016/j.envsoft.2021.104978

Su, J., Huang, J., Adams, S. C., Chang, Q., & Beling, P. A. (2022). Deep multi-agent reinforcement learning for multi-level preventive maintenance in manufacturing systems. *Expert Syst. Appl.*, *192*, 116323.

Sun, Z., Lorscheid, I., Millington, J. D., Lauf, S., Magliocca, N. R., Groeneveld, J., Balbi, S., Nolzen, H., Müller, B., Schulze, J., & Buchmann, C. M. (2016). Simple or complicated agent-based models? A complicated issue. *Environmental Modelling & Software*, *86*, 56–67. https://doi.org/10.1016/j.envsoft.2016.09.006

Suresh, S. (2021, May 22). *Interpreting CNN Models*. Medium. https://towardsdatascience.com/interpreting-cnn-models-a11b1f720097

Taghikhah, F., Voinov, A., Filatova, T., & Polhill, J. G. (2022). Machine-assisted agent-based modeling: Opening the black box. *Journal of Computational Science*, *64*, 101854. https://doi.org/10.1016/j.jocs.2022.101854

Tambwekar, P., Silva, A., Gopalan, N., & Gombolay, M. (2022). *Specifying and Interpreting Reinforcement Learning Policies through Simulatable Machine Learning* (arXiv:2101.07140). arXiv. https://doi.org/10.48550/arXiv.2101.07140

Tan, S. K., Hu, N., & Cai, W. (2019). A data-driven path planning model for crowd capacity analysis. *Journal of Computational Science*, *34*, 66–79. https://doi.org/10.1016/j.jocs.2019.05.003

van der Hoog, S. (2019). Surrogate Modelling in (and of) Agent-Based Models: A Prospectus. *Computational Economics*, *53*(3), 1245–1263. https://doi.org/10.1007/s10614-018-9802-0

Vincenot, C. E. (2018). How new concepts become universal scientific approaches: Insights from citation network analysis of agent-based complex systems science. *Proceedings of the Royal Society B: Biological Sciences*, *28*(1874), 20172360.

von Bertalanffy, L. (1968). *General System Theory: Foundations, Development, Applications*. George Braziller.

Vu, T. M., Probst, C., Epstein, J. M., Brennan, A., Strong, M., & Purshouse, R. C. (2019). Toward Inverse Generative Social Science Using Multi-Objective Genetic Programming. *Proceedings of the Genetic and Evolutionary Computation Conference*, 1356–1363. https://doi.org/10.1145/3321707.3321840

West, D. M., & Allen, J. R. (2018). *How artificial intelligence is transforming the world*. https://www.brookings.edu/research/how-artificial-intelligence-is-transforming-the-world/

Zewe, A. (2022). *Demystifying machine-learning systems*. MIT News | Massachusetts Institute of Technology. https://news.mit.edu/2022/explainable-machine-learning-0127

Zhang, C., Yao, W., Yang, Y., Huang, R., & Mostafavi, A. (2020). Semiautomated social media analytics for sensing societal impacts due to community disruptions during disasters. *Computer-Aided Civil and Infrastructure Engineering*, *n/a*(n/a). https://doi.org/10.1111/mice.12576

Zhang, H., Vorobeychik, Y., Letchford, J., & Lakkaraju, K. (2016). Data-driven agent-based modeling, with application to rooftop solar adoption. *Autonomous Agents and Multi-Agent Systems*, *30*(6), 1023–1049. https://doi.org/10.1007/s10458-016-9326-8

Zhang, W., Valencia, A., & Chang, N.-B. (2021). Synergistic Integration Between Machine Learning and Agent-Based Modeling: A Multidisciplinary Review. *IEEE Transactions on Neural Networks and Learning Systems*, 1–21. https://doi.org/10.1109/TNNLS.2021.3106777