EMERGENT COORDINATION: ADAPTATION, OPEN-ENDEDNESS,
AND COLLECTIVE INTELLIGENCE

By

Honglin Bao

A THESIS

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

Computer Science – Master of Science

2022

ABSTRACT

EMERGENT COORDINATION: ADAPTATION, OPEN-ENDEDNESS,
AND COLLECTIVE INTELLIGENCE

By

Honglin Bao

Agent-based modeling is a widely used computational method for studying the micro-macro bridge issue by simulating the microscopic interactions and observing the macroscopic emergence. This thesis begins with the fundamental methodology of agent-based models: how agents are represented, how agents interact, and how the agent population is structured. Two vital topics, the evolution of cooperation and opinion dynamics are used to illustrate methodological innovation. For the first topic, we study the equilibrium selection in a coordination game in multi-agent systems. In particular, we focus on the characteristics of agents (supervisors and subordinates versus representative agents), the interactions of agents (reinforcement learning in the games with fixed versus adaptive learning rates according to the supervision and time-varying versus supervision-guided exploration rates), the network of agents (single-layer versus multi-layer networks), and their impact on the emergent behaviors. Regarding the second topic, we examine how opinions evolve and spread in a cognitively heterogeneous agent population with sparse interactions and how the opinion dynamics co-evolve with the open-ended society's structural change. We then discuss the rich insights into collective intelligence in the two proposed models viewed from the interaction-based adaptation and open-ended network structure. We finally link collective emergent intelligence to diverse applications in the realm of computing and other scientific fields in a cross-multidisciplinary manner.

This thesis is dedicated to my parents, for their endless love, support, and encouragement.

ACKNOWLEDGEMENTS

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ALGORITHMS

KEY TO ABBREVIATIONS

SCOOE model   **S**parse **CO**-evolutionary **O**pen-**E**nded model

Simulated **A**nnealing-based Exploration **SA** Exploration

Neural **C**ellular **A**utomata   NCA

INTRODUCTION [1]

An agent-based model is a computational method that simulates the characteristics, behaviors, and interactions of entities (the smallest units) in a system to comprehend the system's dynamics and the factors that govern such dynamics. At the macroscopic level, we can observe the rich emergence of "novelties" by assigning basic behavioral rules or communication protocols to agents to represent individual behaviors and interactions, e.g., game-theoretic, reinforcement learning-based, and behavioral science-inspired interactions. We cannot fully predict these novelties using "reductive" analysis of microscopic rules, such as differential equation-based analysis, due to their emergent nature. These novelties can be viewed as the collective behavior of simple-rule-following intelligence. This "bottom-up" modeling methodology is referred to as agent-based modeling, which has been widely applied in a variety of fields, most notably biology and social science.

The idea of "bottom-up emergence" has a long history in computer science. An early work is the von Neumann Machine (Aspray, 1989). The machine creates a replica of itself based on simple rules. John Conway, a British mathematician, was inspired by this and created Game of Life by introducing simple interactions between agents and their neighbors in a two-dimensional lattice environment. We can observe novel dynamics emerging in the Game of Life continuously: A disordered agent population gradually evolves into a structure with multiple delicate and tangible shapes under various rules and initialized states. Several of them maintain symmetry while undergoing continuous shape changes. Certain well-formed shapes are disrupted

---

[1] This section is based in part on the ongoing paper coauthored with Wolfgang Banzhaf, "*Social Inspiration in Computational Models: A Survey.*"

and destroyed when disordered agents "invade." Several of them remain stable over time. Generally, we frequently observe that "order" emerges from "chaos."

In the 1980s, Craig Reynolds used agent-based simulation to model the group behavior of social organisms and animals (Reynolds, 1987). Christopher Langton then coined the term "artificial life" or "ALife" to refer to this methodology in order to emphasize its significance in the study of the evolution of digital organisms (Langton, 1997). Artificial life now is a growing field at the intersection of evolutionary biology, complex systems, and computational cognitive science (Attenberg, 2005). Sociologists have modeled humans and their behavior as a collection of autonomous agents that operate in parallel and communicate with one another for the purpose of simulating real-world social phenomena. Schelling's segregation model is a pioneering social simulation model (Schelling, 1971), which states that people's "mild" preference for group membership can result in a highly segregated society. Epstein and Axtell have coined the term "bottom-up" social science or generative social science to refer to this simulation methodology (Epstein & Axtell, 1996). Agent-based models have thriving commercial applications in the real world, notably for robotics and self-driving cars, accompanied by the rapid advancement of artificial intelligence and machine learning. They have a profound theoretical influence on task allocation and coordination in real-world e-Commerce robotic picking systems (Barbati et al., 2012; Yuan et al., 2021; Wang et al., 2018). Waymo has proposed Carcraft, an agent-based platform for testing self-driving algorithms. It models the complex interactions between self-driving cars, pedestrians, and humans in traffic by creating different agents and calibrates the simulated human interactions with cars using massive real-world human behavioral data (Connors et al., 2018).

The purpose of this thesis is to study agent-based models, with a focus on their computational characteristics and multidisciplinary applications. The rest of the thesis is organized as follows. Section I gives an overview of the methodology. Section II and III use two examples, the evolution of cooperation and opinion dynamics, to illustrate how to incorporate the methodology into the computational model design. Section IV reflects two proposed models and their insights into collective intelligence. Section V concludes the thesis with final remarks.

I.   METHODOLOGY [2]

What is the critical methodology for incorporating agent-based models into multidisciplinary

applications via advanced computing techniques? To answer this question, we must first examine

the foundations of a "system." A system is, in brief, the sum of the relationships between entities

within the system and the structure within which they exist (Burns & Helena, 1987). In agent-

based models, an agent acts on behalf of an entity. As are composed of numerous agents to

simulate human-made or natural systems by simulating both entities and their interactions and

structures. We divide the methodology into three perspectives mirroring how a system operates:

how agents are represented, how agents interact, and how agent populations are structured.

The agent property is a critical concept in agent-based models. Economists refer to the

model that contains identical agents of the same type as a representative agent model. More

broadly, the agents in a representative model may be distinct. Nonetheless, their respective

characteristics are irrelevant to the problem at hand, i.e., the agent behaves in such a way that the

overall decision is practically the same as the decision of a single agent or a group of similar

agents. Thus, we can simplify this to a model of representative agents. On the other hand, if

distinct agent characteristics are relevant to the problem, we must consider agent heterogeneity,

i.e., the differences between agents. The heterogeneity of agents is a fascinating subject. The

majority of existing research assigns distinct functions or roles to agents to capture real-world

behaviors, such as opinion leaders and followers in social networks (Zhao & Kou, 2014),

ethnically diverse agents in studies of residential segregation (Fossett & Waren, 2005), different

species in ecological systems (Filatova et al., 2013), and trustworthy and untrustworthy agents in

---

[2] This section is based in part on the ongoing paper coauthored with Wolfgang Banzhaf, "*Social Inspiration in Computational Models: A Survey.*"

reputation computing, a sub-field of computer security (Wang et al., 2010). Numerous elements in agent property design are inspired by psychological and behavioral studies. The most well-known example is the BDI (belief-desire-intention) agent model (Georgeff, 1998). It is inspired by Michael Bratman's theory of reasoning about human practices (Bratman, 1987). A BDI agent is defined by executing programmed beliefs, desires, and intentions and utilizes these concepts to design and evolve "smarter" agents. It has been widely used in planning and scheduling problems in software engineering (Rao & Georgeff, 1995).

A critical aspect of the agent-based computing model is microscopic mechanisms of interaction and their impact on global emergence. Numerous perspectives are studied in the design of interaction protocols. For instance, game-theoretic interactions are used to simulate human strategic reasoning and behaviors (Adami et al., 2016). Reinforcement learning is used to model the interaction of agents with their environments (e.g., other agents) by maximizing cumulative rewards (Barbati et al., 2012). The concept of feedback in reinforcement learning, i.e., receiving rewards for good behaviors and punishments for creating problems, is fundamental to the study of decision-making. It is critical guidance for animals and humans acting and surviving in unknown environments. This concept permeates computational and experimental models in psychology, neuroscience, behavioral science, and ethology (Dayan & Daw, 2008). Learning and game theory are two widely used means of interaction with other agents or the environment. Agents gradually adapt to the environment in the process of interaction. In Section IV, we will examine systematically how the interaction among agents affects their adaptation and the macroscopic evolutionary dynamics.

Some work takes a structural view of the holistic-level patterns. It conceptualizes the population as a social network, with each node representing an individual and each edge

representing an interaction between individuals. We can study the macroscopic patterns, network statistics, and dynamics of social networks using graph theoretic approaches. An example is social network-inspired optimization and its wide applications, e.g., information retrieval (Nasution & Noah, 2012; Matsuo et al., 2007). Improving the recall and precision rates in the retrieval process can be modeled by detecting increasingly segregated communities in the document population represented as a network. In addition, when the agent population is structured as a network, the statistical characteristics of the network, such as connectivity and open-endedness, will have a significant effect on the macroscopic emergent patterns. This issue will be discussed systematically in Section IV based on the models proposed in the thesis.

On the other hand, some works argue that structured methods cannot be used effectively to analyze emergent dynamics (Villani et al., 2021). Otherwise, it would not be called "emergent." These scholars employ a bottom-up perspective of macroscopic phenomena, in which even if the microscopic structure is chaotic, this has no effect on the emergence of the subtle macroscopic structure. A notable example is a large number of self-organized multiagent clustering algorithms (Thrun & Ultsch, 2021; Bu et al., 2017).

Following these three lines of inquiry, Figure 1 illustrates the overview of methodology. The following sections will study how agent-based models with novel representations, interactions, and structures can be used to address computational problems related to coordination, cooperation, opinion dynamics, and collective intelligence.
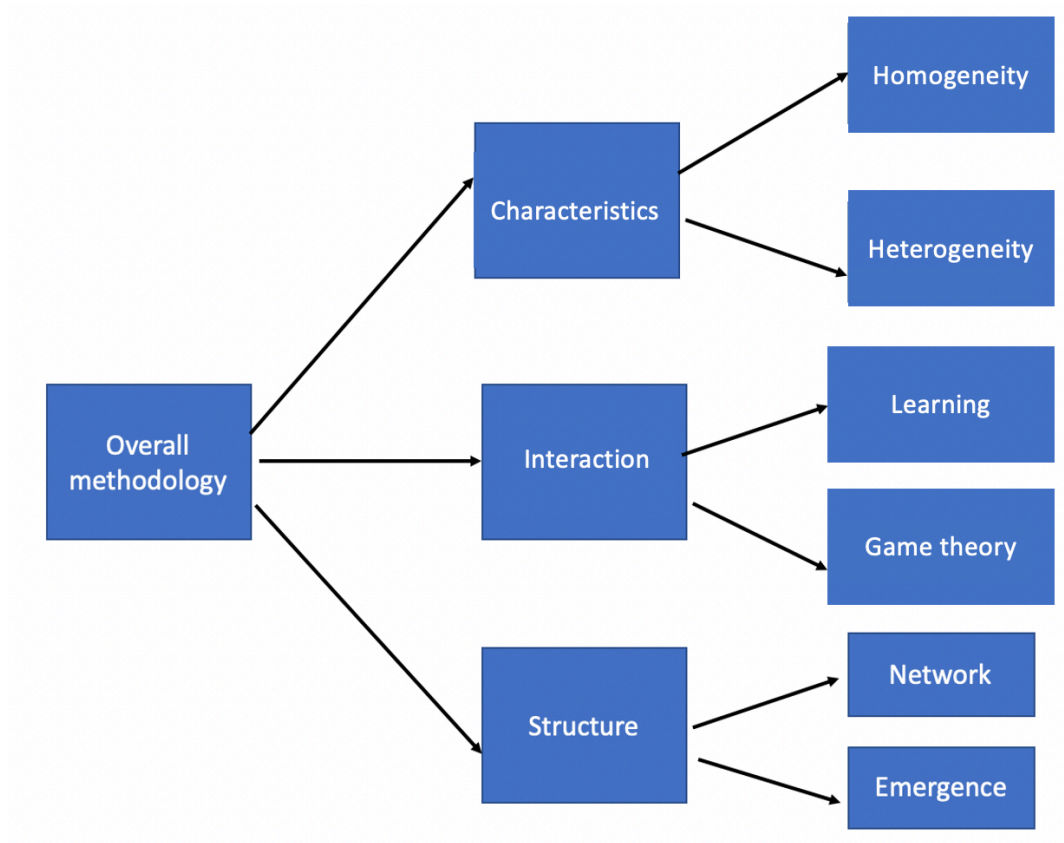
Figure 1: The overall methodology.

II.   EVOLUTION OF COOPERATION [3]

A. Coordination and Cooperation

The study of the evolution of cooperation has a long history in the application of game theory (Axelrod & Hamilton, 1981). It also assumes great importance in the field of multiagent systems. In multiagent societies, cooperation represents an interaction among agents that can be evolutionarily advantageous to improve the performance of individual agents or the overall behavior of the society they belong to. Therefore, one of the main goals in multiagent societies is to achieve efficient cooperation among agents to jointly solve tasks or to maximize a utility function.

In this section, we adopt the "Rules of the Road Game," a typical coordination game, as an example to study the evolution of cooperation (Young, 1996). Consider two carriages meeting on a narrow road from opposite directions, having no context to decide on which side of the road to pass the other. If they choose differently, it will cause a head-on collision between them, and they receive a negative payoff. Only if they choose the same way can they avoid a collision and receive a positive payoff. To abstract from this realistic situation to virtual multiagent societies, agents are striving to establish a convention/law of coordinated action by choosing from an action space without any central controller. The payoff matrix is shown in Table 1.

---

[3] This section is based in part on the paper: Honglin Bao, Qiqige Wuyun, Wolfgang Banzhaf; July 23–27, 2018. "Evolution of Cooperation through Genetic Collective Learning and Imitation in Multiagent Societies." *Proceedings of the ALIFE 2018: The 2018 Conference on Artificial Life*. ALIFE 2018: The 2018 Conference on Artificial Life. Tokyo, Japan. (pp. pp. 436-443). ASME. https://doi.org/10.1162/isal_a_00082

Table 1: Payoff matrix of an *n*-action 2-player coordination game.

	Action 1	Action 2	...	Action n
Action 1	1,1	-1,-1	...	-1,-1
Action 2	-1,-1	1,1	...	-1,-1
...	...	...	...	...
Action n	-1,-1	-1,-1	...	1,1

There are multiple Nash-equilibria in this diagonal situation. Both two players choose the same action, i.e., coordinated action. However, even purely rational players cannot choose the specific coordinated action without negotiation because they have no information to differentiate between strictly the same multiple equilibria. In reality, people can survive such social dilemmas because there are laws or social norms for them to refer to. Our goal is to train agents of a virtual society to choose cooperative action without upper-level steering and regulation.

B. Related Work

In order to realize such coordination and cooperation, some techniques developed in the field of machine learning have been introduced into various multiagent systems (Kapetanakis & Kudenko, 2002). Machine learning has been proven to be a popular approach to solving multiagent system coordination problems (Savarimuthu et al., 2011). Among machine learning techniques, reinforcement learning has gained much attention in the field of multiagent systems since it can be used to model agent learning by trial-and-error interaction with the dynamic environment. However, several new challenges arise for reinforcement learning in multiagent systems. Foremost among these is that the performance of reinforcement learning is unsatisfactory in many real-world applications. The learning algorithm may not converge to the optimal action. Some researchers showed that an adaptive strategy, called evolutionary

reinforcement learning, which combines reinforcement learning with a genetic algorithm, could reach a better performance than either strategy alone (Ackley & Littman, 1991). Some new forms of learning, e.g., observational, imitational, and communication-based learning (Taylor et al., 2006, Savarimuthu et al., 2011), also significantly promote information proliferation (Dittrich & Banzhaf, 2002) in more complex environments and can be used to solve complex distributed multiagent coordination problems better than pure reinforcement learning approaches. Furthermore, ensemble methods are used to combine the advantages of multiple learning algorithms to obtain better performance than what could be obtained from any of them alone (Polikar, 2006). More recently, Yu and colleagues (Yu et al., 2017) studied the role of reinforcement learning, collective decision-making, social structure, and information diffusion in the process of the evolution of cooperation in the networked society.

C. The Proposed Model

Although previous work provided a strong basis to study the mechanisms behind the evolution of cooperation, existing work in this area has drawbacks. Individual reinforcement learners often fail to develop globally coordinated behavior and can be trapped in local sub-optimal dilemmas. Using an evolutionary approach for strategy selection can produce optimal behavior but may require significant computational efforts. Behavior imitation always creates weak local coordination in society, caused by local interactions between agents. This study is significantly different from other frameworks for the evolution of cooperation in previous studies because of the hybrid policy of agents' decision-making. Here we design a genetic algorithm-based cooperation framework, which takes into account evolutionary selection, collective learning, and imitation, in order to solve some particular non-cooperative games in complex multiagent

10

networks, overcome previous shortcomings, and produce an acceptable tradeoff in convergence rate and computation effort.

The final decision of an agent (both parent and offspring agent) is influenced by the three kinds of processes shown below. Note that all agents in the model are identical in terms of interaction, learning, and decision-making. This kind of model is referred to as the "representative agent-based models" in economics (Gallegati & Kirman, 2012).

- Evolutionary Selection (with inheritance and mutation): A population of agents plays a game with their neighbors (i.e., the agents which are directly connected with the focal agent) on the network for several iterations. The offspring generation will be reproduced from the parent generation according to the cumulative payoff distribution, and the most successful agents will pass on action to their offspring. Mutation will occur with a small probability during the inheritance process to create novelty.

- Collective Learning: Agents on the network improve on their parents' actions and their original actions through a collective reinforcement learning algorithm with exploration and exploitation.

- Imitation: Agents update the cumulative payoff, compare their cumulative payoff to neighbors, and adopt the actions of more successful agents as their own actions with a particular probability.

These three processes interact with each other and can have a significant influence on the evolution of cooperation in the entire society.

The overview of the proposed algorithmic framework is shown in Algorithm 1. It constitutes a genetic algorithm-based cooperation framework for multiagent systems with collective decision-making, learning and imitation to facilitate the evolution of cooperation used

11

in some particular coordination games. This framework is set in a network structure such as a

small-world network (Watts & Strogatz, 1998) or a scale-free network (Barabási & Albert,

1999). A population of agents plays the coordination game with their neighbors repeatedly and

simultaneously in the network for several generations. Offspring generation i_o will be reproduced

from the parent generation i_p according to their cumulative payoff distribution E. That means, the

probability of reproduction is proportional to the cumulative payoff; The agents with higher

cumulative payoffs are more likely to reproduce their offspring. The most successful agents pass

on behavior to their offspring i_o, and mutation will change this behavior with a small probability

η during inheritance. The society information regarding nodes and edges will be updated at each

round with new offspring nodes and parent-offspring edges. Then agents will improve their

actions through a collective reinforcement learning algorithm with exploration and exploitation

and a social learning process.

Algorithm 1: The proposed cooperation framework.

	Algorithm 1: The proposed cooperation framework
1	Initialize multiagent network and parameters;
2	**for** *each step t (t=1,...,T)* **do**
3	**for** *each agent i (i=1,...,n_0)* **do**
4	**for** *each neighbor $j \in N(i)$ of agent i* **do**
5	Agent i plays the game with neighbor agent j and receives corresponding payoff r_i^j;
6	**end**
7	Agent i calculates the cumulative payoff E_i;
8	Offspring generation i_o will be reproduced from parent generation i_p according to E_i;
9	**end**
10	**for** *each parent agent i_p (i_p =1,...,n_0)* **do**
11	Parent i_p passes on behavior to the offspring i_o;
12	Mutation will change it with a small probability η during inheritance;
13	**end**
14	The society information regarding nodes and edges will be updated;
15	**for** *each agent i in a new network* **do**
16	**for** *each neighbor $j \in N(i)$ of agent i* **do**
17	Agent i improves the behavior with a collective learning method with exploration and exploitation regarding neighbor j;
18	Agent i and neighbor j update the cumulative payoff E_i' and E_j';
19	Agent i imitates the action of neighbor agent j with a probability W;
20	**end**
21	**end**
22	**end**

We adopt a widely used reinforcement learning algorithm, *Q*-learning, to model this interaction. Its one-step updating rule is given by Equation 1. Here $\alpha \in (0, 1]$ is a learning rate, and $\lambda \in [0, 1)$ is a discount factor.

$$Q(s,a) \leftarrow Q(s,a) + \alpha[R(s,a) + \lambda \max_{a'} Q(s',a') - Q(s,a)] \tag{1}$$

As shown in Equation 1, an agent has a set of states and a set of actions. An agent performs an action *a*, transitions from state *s* to another new state *s'* and receives an immediate

reward $R(s, a)$. $Q(s, a)$ is the expected reward of choosing action a in state s at time step t. During the interaction, agents want to maximize the expected discounted reward $Q(s', a')$ to make decisions in the new state s' at time step $t + 1$. The Q-function is learned during an agent's lifetime inherited to choose a best-response action based on the Q-value regularly. Each agent needs to aggregate all the best-response actions regarding its neighbors into an overall action. This is inspired by the opinion aggregation process in that people usually have to seek suggestions from many other people before making a final decision. The opinion aggregation process can be realized by an ensemble learning method which combines multiple single-learning algorithms to obtain better performance than what could be obtained from any of them alone (Polikar, 2006). The foremost method of collective voting is inspired by a simple political principle, majority rule. Consider that in a simple society (e.g., an undirected simple graph which represents the multiagent network we adopt in this paper), human beings are more keen to decide as the majority of their neighbors. So, when agents make final decisions, they consider the action which quantitatively dominates in the best-response action pool. More complex and realistic methods to make a final decision consider the weight of each neighbors, such as performance-based weighted voting method and structure-based weighted voting method. For structure-based weighted voting, the weight of each neighbor is related to the degree of each neighbor. The focal agent will give higher weight to a neighbor with more connections. For performance-based weighted voting, the focal agents will consider previous interaction experience and will give higher weight to neighbors they trust. In this study, we adopt majority voting as the opinion aggregation method (Bao et al., 2018).

On the other hand, for the pure greedy processes, such as the reinforcement learning-based interaction, agents can be trapped easily in local sub-optima, and thus fail to learn the

optimal behavior. During learning, an agent needs to strike a balance between exploitation of learned knowledge and the exploration of unexplored environments in order to try more actions, escape from local sub-optima, and learn optimal behavior. In this paper, we propose *time-varying exploration inspired by simulated annealing* (**SA exploration**) for dealing with exploitation and exploration during learning by treating the exploration rate as temperature. One step of SA exploration is given by Equation 2.

$$\mu_t = \mu_0 / \lg(1 + t) \tag{2}$$

In Equation 2, μ_t is the exploration rate in the t^{th} round of simulation, and μ_0 is the initial exploration rate. At the beginning (t is small), exploration should be given higher weight to explore the unknown environment. As the algorithm continues (t increases), the probability of exploitation (i.e., 1- μ_t) increases determining that the agent will focus more on exploitation of learnt knowledge.

The collective learning framework is illustrated in Algorithm 2: during the interaction with neighbors, agents need to find a best-response action regarding each neighbor with a *Q*-learning method. At each time step t, regarding each neighbor j, agent i chooses the best-response action with the highest *Q*-value with a probability of 1- μ_t (i.e., exploitation), or chooses an action randomly with a probability of μ_t (i.e., exploration). This occurs in the process of local interaction with neighbors. We call this process **Local SA Exploration**. When agents use specific ensemble methods to aggregate all the best-response actions into an overall action, agents choose the overall action under ensemble methods with a probability of 1- μ_t (i.e., exploitation), or choose an action randomly with a probability of μ_t (i.e., exploration). This occurs in the process of overall aggregation. We call this process **Global SA Exploration**.

15

In general, a small exploration rate (such as 10%) is kept throughout to conserve a small probability to explore.

Algorithm 2: The collective learning framework.

Algorithm 2: The collective learning framework

1 **for** *each step t (t=1,...,T)* **do**
2   **for** *each agent i (i=1,...,n)* **do**
3   **for** *each neighbor $j \in N(i)$ of agent i* **do**
4   Agent i has a Q function for each of its neighbours j;
5   Agent i chooses a best-response action $a_{i \to j}$ regarding neighbor j using a Q-learning algorithm;
6   //Local exploration;
7   **end**
8   Agent i aggregates all the actions $a_{i \to j}$ into an overall action a_i using ensemble learning methods;
9   //Global exploration;
10   **end**
11   **for** *each agent i (i=1,...,n)* **do**
12   Agent i plays action a_i with its neighbors and receives corresponding payoff $r_i^{j'}$ for each interaction;
13   Agent i updates learning information towards each neighbor using action-payoff pair $(a_i, r_i^{j'})$;
14   **end**
15 **end**

Social learning theory is connected with social behavior and learning and proposes that new behavior can be obtained by observing and imitating others' behavior (Bandura and Walters, 1977). In real life, people not only can learn through their individual trial-and-error experiences (i.e., individual Q-learning to determine best-response actions), but also seek suggestions or advice from others in a society (as mentioned in opinion aggregation). Furthermore, they can also learn from the information directly provided by others through communication, observation, and imitation (Polikar, 2006). We are inspired by social learning

theory to add an imitation process after learning to promote the evolution of cooperation. After reproduction and learning, there is a new population with better performance in multiagent societies. In every time step, when agent i updates the cumulative payoff E_i, agent i in this new population adopts neighbor agent j's behavior, replacing its heritable behavior, with a probability W. Following Szabo and Toke (Szabo & Toke, 1998), we set:

$$W = \frac{1}{1 + e^{-(E'_j - E'_i)/K}}$$

(3)

Here, E'_i and E'_j are the cumulative payoff of agent i and neighbor j after updating actions and payoffs. K represents some noise which is introduced to consider irrational choices. For $K = 0$ agent i adopts neighbor j's strategy if $E'_j > E'_i$. Here we set $K = 0.1$.

D. Model Performance

To test the model, we use the Watts-Strogatz model (Watts & Strogatz, 1998) to generate a small-world network, and use the Barabasi-Albert model (Albert & Barabasi, 2002) to generate a scale-free network. In order to use the Barabasi-Albert model, we start with 2 agents and add a new agent with 1 edge to the network at every step. Because of the rewiring probability ρ, this approach generates a scale-free network following a power law distribution with an exponent $\gamma = 3$. We set the maximum number of edges to 1,000,000 for network evolution. Mutation rate η in inheritance is 0.01. Individual Q-learning rate α is 0.1. Average exploration rate in SA exploration is 0.1 for 100 time steps, so the initialized exploration rate μ_0 is 0.144 after a simple mathematical calculation. Noise in imitation is set to 0.1. In this study, unless stated otherwise, we use the small-world network as the default network topology because it can evolve into many kinds of networks, and local SA exploration as the exploration mode. Experiments will be run in 100 time steps.

1. Comparison of Mutation and Two Types of Exploration

We test the situation under 4-action space, i.e., action 0,..., action 3 respectively. Figure 2 shows the asymptotic percentage of cooperative actions (action 0) adopted by the agents when cooperation evolves in the entire society. Initially, each agent randomly chooses an action from action space, so there are about 25% of all agents to choose each action respectively. As our framework moves on, the number of agents who choose action 0 as the cooperative action finally reaches more than 90% in the situation with SA exploration (both local and global). This result means that more and more agents have reached a consensus on that action 0 should be the cooperative action. From Figure 2, we can see that the fraction of cooperators in the society using collective learning with local SA exploration mode is almost 100% which means that almost all the agents have reached a consensus on which action should be the cooperative action. The framework works in the entire society.
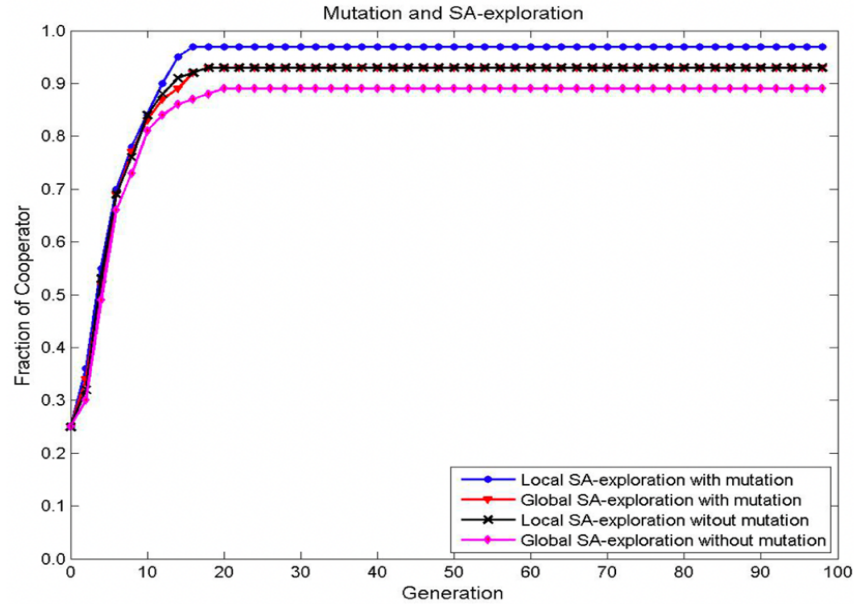
Figure 2: Fraction of cooperators under different exploration and mutation methods.

We further study Figure 2 and we can draw these conclusions:

1). Local exploration is better than global exploration. The fraction of cooperators using collective learning with the global exploration mode is lower than that using collective learning with the local exploration mode. This is because agents explore the environment with a probability of 0.1. However, as agents using local exploration to explore the environment locally (i.e., choosing irrational action during local interaction) and aggregate to an overall action collectively, the randomness caused by the exploration can be removed. In global exploration, agents explore globally when they aggregate all best-response actions into an overall action, the randomness will be kept.

2). Mutation is necessary. The fraction of cooperators with mutation is higher than that without mutation. Although sometimes mutation has a bad influence, indeed, it is the source of novelty.

2.   Comparison with Previous Work

We mainly compare the performance of our model with Yu et al., 2017. Yu's work is mainly based on collective reinforcement learning and information diffusion (i.e., communication-based social learning, agents sharing Q table to communicate). As shown in Figure 3, we set the action space as $N_a = 10$ and follow all other parameter settings. Our framework has better performance than the previous study. We additionally test other situations with different action spaces; the results show the same trends. It indicates that our model works for the evolution of cooperation in the entire society. It is indeed effective for robust evolution by combining evolutionary selection, individual learning, collective voting, and social imitation.
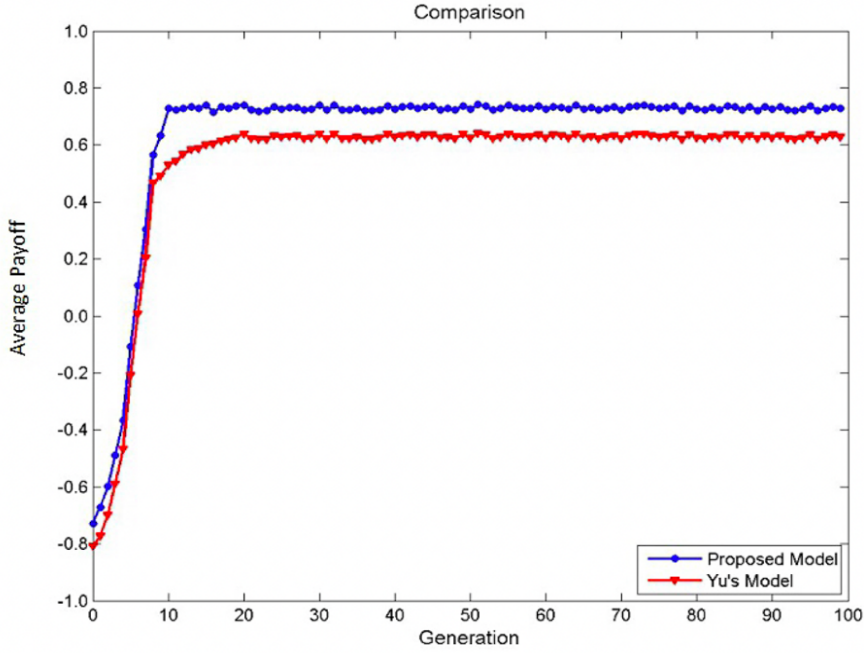
Figure 3: Comparison with Yu et al., 2017.

Through our experimental analysis, we find that there is not much difference in the efficiency of the evolution of cooperation in different sizes of agent population, different opinion aggregation methods, and different network structures. We additionally test models with separated mechanisms. We find: 1). Collective decision-making (opinion aggregation) and imitation will significantly facilitate the evolution of cooperation, especially collective decision-making; 2). Evolutionary selection does cause influence both on the convergence speed and convergence rate, but not as dramatic as collective decision-making or imitation; 3). We could not get any convergence curves in 100 generations during experiments without reinforcement learning-based interactions.

To summarize, for robust cooperation evolving in networked agent systems, the potential key factors are:

- the way how agents interact with each other. This is also called interaction protocol. For instance, interacting randomly in a population or interacting with neighbors in a network; what game-theoretical situations the interaction is based on.

- the way how agents update their learning information through interaction, i.e., what learning strategies (e.g., collective Q-learning (Yu et al., 2017), WoLF-PHC (Win or Learn Fast Policy Hill-Climbing, Bowling & Veloso, 2001), and fictitious play (Monderer & Shapley, 1996)) do agents use to update their learning information?

- the way how agents diffuse their learnt information, e.g., communication-based social learning, imitation-based social learning, and observation-based social learning.

- whether the entire population evolves in a better direction. Evolving to improve the entire fitness (e.g., reproducing offspring with better performance to increase the entire average fitness) represents an enhancement in the evolution of cooperation.

E. Multilayer Networks

Our previous model is built on a single-layer social network with representative agents. Let us further consider a variant of the basic model: how does the multilayer structure, for instance, supervisor and subordinates, shape the emergent dynamics of cooperation? Multilayer networks are networks composed of multiple layers of sub-networks, which can be widely found in nature and man-made systems. Each layer represents a distinct interaction, social circle, or timestamp. The analysis of modern networked social and physical systems like online social networks, transportation systems, metabolic and regulatory networks can all benefit from this new paradigm of network science. It can better distinguish and deal with heterogeneous relationships

between levels like supervisor/subordinate, opinion leader/follower, and hierarchical ecosystems than the previous flattening paradigms.

The formalized social structure is a multilayered networked structure with dominance between higher and lower layers. This networked structure is divided into multiple small groups $G_x \in (1, ..., X)$, where X indicates the number of small groups in the society, and x indicates the supervisor of a particular group. Supervisors of lower-level subordinates can also be subordinates of higher-level supervisors. Notice that supervisor agents are not global controllers. They make decisions based on local information reported by subordinates. So, we combine up-level supervision with bottom-level individual learning in the model.

An agent i *interacts* with its rivals who connect with it directly and uses Q-learning with an exploration algorithm with the SA exploration rate μ_t, which have been documented in the original model, to choose a best-response action a_i. Then agent i receives the corresponding payoff r_i according to Table 1. After action-selection, focal agent i stores action a_i, cumulative payoff R_i (calculated by the sum of corresponding payoffs interacting with all rivals) and learning parameters in a table. Agent i then *reports* all of them to agent i's supervisor x and recognizes the rivals' previous actions to determine whether to withdraw with a defector with a small probability of 10%. Supervisor x *combines* all reported best-response actions of its subordinate agents into an overall action a_x through the collective learning methods, which has been documented in the original model (Algorithm 2). Supervisor x then interacts with a randomly selected rival in the same cluster of the supervisor layer, and *imitates* to update a_x into a'_x according to the performance difference between the overall actions of supervisor x and this rival (e.g., average cumulative payoff), following Equation 3 (Szabo & Toke, 1998). After aggregating the subordinates' actions and imitation, supervisor x has a final action a'_x which will

be used as *guidance* to teach its subordinate agents to act better. Two basic parameters of the learning process, i.e., learning rate α and SA exploration probability μ_t, will be adjusted based on supervisor x's steering information among subordinate agent i and the peers within the same subordinate cluster. Finally, all agents update their learning information using the new learning rate and exploration rate. This closed-loop process is iterated for T time steps. The entire picture of the overall algorithmic framework is shown in Figure 4.
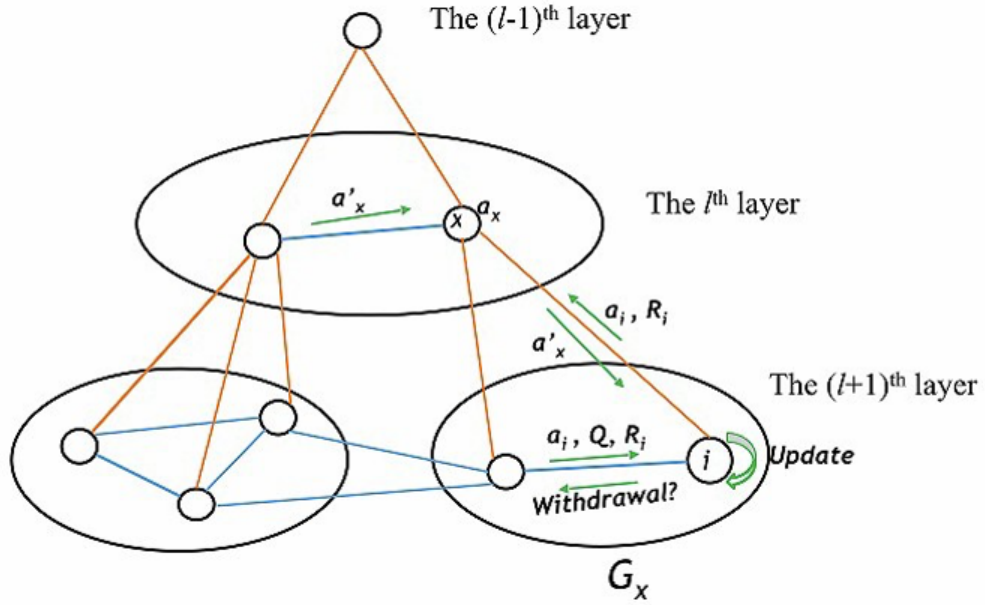
Figure 4: The overview of the multilayered algorithmic framework.

The processes of collective decision-making-based aggregation and imitation have been reported in the original model. While how do supervisors utilize the reported information to guide behavioral change among subordinates? We apply the *Adaptive Learning Method.* Supervisor x passes down the action a_x, after aggregation and imitation, to its subordinate agents i. At each time step t, based on this guidance information from supervisor x, agent i adjusts its

behavior. "Adaptive learning" means that the subordinate agents adaptively adjust their learning information based on the supervised information from supervisors.

Some previous algorithmic frameworks with high time/space complexity can be used to model this supervised process with the adaptive self-adjustment, such as MiniMax-Q Algorithm (Littman, 1994), Nash Q-learning Algorithm (Hu & Wellman, 2003), and Friend-or-Foe Q-learning Algorithm (Littman, 2001). In this thesis, we introduce a simple but insightful philosophy, i.e., win stay, lose shift (Nowak & Sigmund, 1993), to build this adaptive adjustment method. First, we should define two situations, "win" and "lose," respectively. If the reported action a_i of subordinate i is identical with the supervisor x's final action a_x after aggregation and imitation, this situation is approved by the supervisor and regarded as "win" and "lose" otherwise. The adjustment is conducted in the process of comparison of guidance information and the current situation of subordinates. Two primary parameters, i.e., learning rate α and SA exploration rate μ_t in Q-learning-based interactions, will be adjusted by accepting guidance. If "win," the focal agent will decrease both learning rate α and exploration rate μ_t to stay in the winning situation; otherwise, increase these two parameters to escape from sub-optima, i.e., "shift." In each step, the degree of adjustment is set to 10%.

We show how multilayer structure re-shapes the evolution of cooperation. We first study the influence of cluster size. We fix the subordinate population to a 32×32 grid network and vary the cluster size to 2×2 (with 16×16 supervisors), 4×4 (with 8×8 supervisors), 8×8 (with 4×4 supervisors), and 16×16 (with 2×2 supervisors), respectively, to study the influence of cluster size on the evolution of cooperative hunting, which is shown in Figure 5. We also assume there are two layers where supervisors share ½ of the corresponding payoff, but subordinates share only ¼, introducing unequal payoff sharing in the hierarchical system.

In general, we find that with a larger cluster size, a higher level of cooperation can evolve (such as in the cases of cluster size 16 × 16, 8 × 8, and 4 × 4). Both small population size and large cluster size will lead to a broader view of a partial system for the focal supervisor, which facilitates the evolution of cooperation. Additionally, the smaller cluster size leads to more local groups distributed in the entire system and significantly increases the diversity of the system. It will take more effort to step across multiple sub-optima to evolve global cooperation. This will bring a negative influence on the evolution of global cooperation. It is unusual for the case of cluster size 2×2. We find that the curve of cluster size 2 × 2 seems to violate the general trends. That is because of the large number of supervisors. In our game-theoretical settings, supervisors share a higher fraction of payoff than subordinates, and the number of supervisors is fewer than subordinates. In the case of cluster size 2 × 2, there are 256 supervisors and 1024 subordinates. The number of supervisors is significantly higher than that in other cases. This will lead to an increase in average payoff (both initial and convergence value) and compensate for the negative influence caused by the small cluster size.
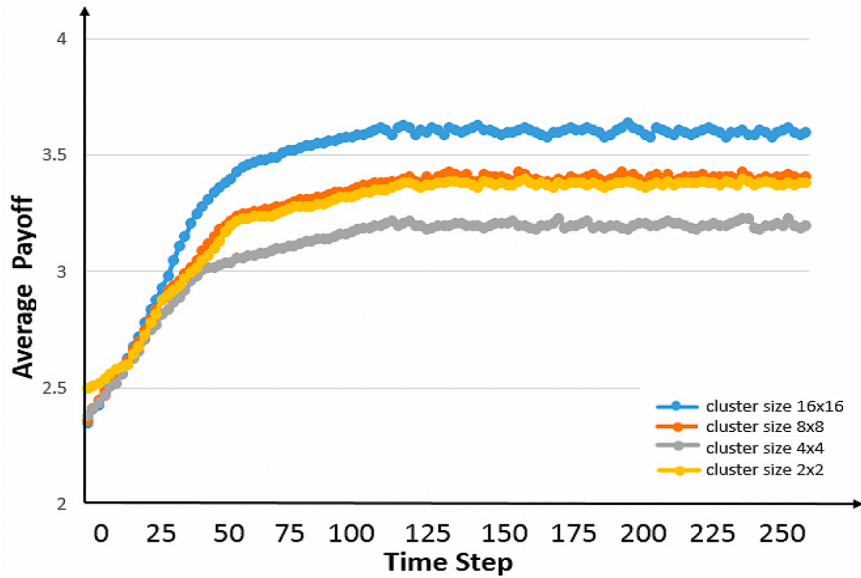
Figure 5: Influence of cluster size on the evolution of cooperation.

We then study the dynamics of learning rate and exploration rate. In our model, dynamics in both learning rate and exploration rate are introduced by the "win stay, loose shift" rule and SA exploration. Based on a hierarchical grid network with 12×12 subordinates and 3×3 supervisors, we study parameter dynamics in reinforcement learning-based interactions shown in Figure 6. The start of the exploration rate μ_t is slightly higher than the learning rate α because of the difference between initialized values. An initialized increasing but decreasing to almost 0 afterward can be seen in both learning rate α and exploration rate μ_t. It indicates that agents initially do not know which action they can adopt, then they interact to try ("trial and error"); hence both α and μ_t increase. As the system evolves, agents realize which action they should adopt, in turn, both α and μ_t decrease until almost 0. Notice that the increase at the initial stage is more significant for learning rate than the exploration rate. SA exploration introduces a continuous decrease in the exploration rate μ_t as the system evolves. As a result, the difference between the learning rate α and exploration rate μ_t appears.
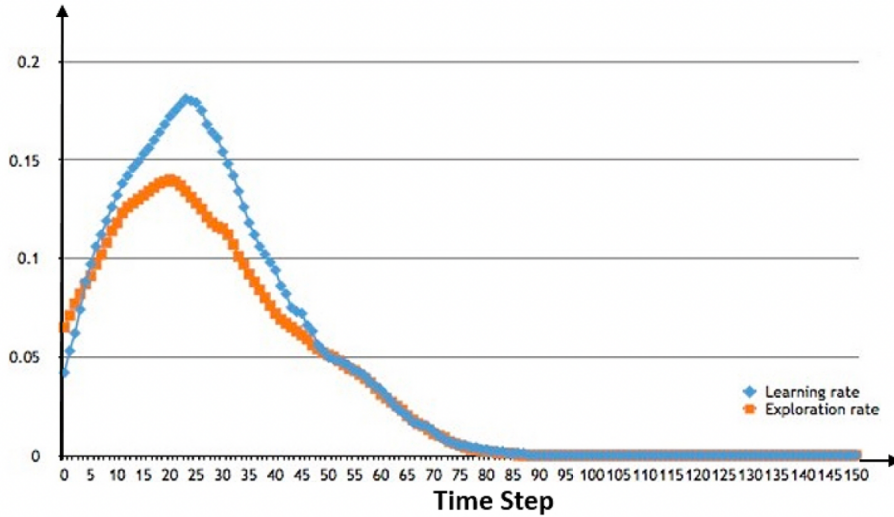
Figure 6: Evolutionary dynamics of learning rate α and exploration rate μ_t.

Many extensions can be done based on this general framework. For example, more layers can be introduced to build a more complex adaptive learning structure. Many network structures, e.g., small-world and scale-free networks, can be introduced to represent agent-based artificial societies. Some interesting factors, such as kin selection which can be intuitively understood that agents have a higher chance to cooperate with kinship, can be investigated based on this framework.

III. OPINION DYNAMICS AND STRUCTURAL CODYNAMICS [4]

We have shown the dynamics of cooperation. In this section, we study the dynamics of coordination within sparse-interaction agents based on an open-ended network.

A. Opinion Dynamics

The study of opinion dynamics, i.e., the study of the formation and dynamics of public opinions, is a crucial research topic in complex systems and social networks. The topic has been widely explored for several decades with theoretical models and real-world applications among different disciplines, including social science, control engineering, statistical physics, and computer science. Elucidation of the mechanisms behind macro-level opinion dynamics is vital for understanding social interactions/dynamics, complexity, distributed control, and decision-making. It also holds valuable lessons to apply to real-world empirical studies and applications like marketing and social media (Mastroeni et al., 2019).

Many classic agent-based models have been explored under various assumptions to study opinion dynamics from different perspectives. For example, the Hegselmann-Krause model (Hegselmann & Krause, 2002) studies opinion polarization with the bounded confidence assumption, i.e., agents interact only if their opinions are sufficiently close to each other by falling within a confidence interval. The Sznajd model and its variations (Sznajd & Jozef, 2000) study the evolution of consensus in a closed society through majority voting. In that model, a focal agent polls its complete neighborhood (i.e., the group of agents sharing connections with the focal agent in the social network) and selects the opinion of the majority. However, some

---

[4] This section is based in part on the paper coauthored with Zachary Neal and Wolfgang Banzhaf, "Coevolutionary opinion dynamics with sparse interactions in open-ended societies." *Complex & Intelligent Systems* (2022). https://doi.org/10.1007/s40747-022-00810-w

assumptions in existing models, e.g., polling the complete neighborhood, seem to be no longer suitable, notably when people with bounded rationality only have a partial view and cannot access the complete neighborhood information in their social networks. Meanwhile, when we interact with neighbors, the literature from psychology suggests, we are mainly concerned with the overall opinion of neighbors (e.g., a joint opinion through collective decision-making), and we adjust our own opinions according to this feedback (Forsyth, 2018). Some other work uses the bounded-confidence assumption (Wang & Shang, 2015) applying dense interactions and serial opinion updates through interacting with all selected neighbors. While existing models thoroughly describe the dynamics of opinions and interactions, they ignore the built-in structural dynamics caused by opinion dynamics and open-endedness, e.g., through the effects of newcomers, leavers, and their impact on structure-opinion coevolution.

B. The Proposed Model

We address the limitations of prior works in the proposed SCOOE model, which takes into account the bounded rationality (e.g., incomplete information) of agents and the coevolution of structure and opinion in an open-ended society, as shown in Figure 7. Opinion dynamics with two components of sparse interactions (extrinsic and intrinsic forms): The focal agent with a limited view can only access a partial neighborhood. It aggregates a joint opinion of the incomplete neighborhood by collective decision-making. Then the focal agent only takes this joint opinion from the extrinsic incomplete neighborhood into account by the interaction with the joint opinion, rather than by dense interactions with all neighbors or certain neighbors with similar opinions selected by polling the neighborhood. Imitation is also introduced to drive opinion intrinsic adjustments based on observation of the incomplete neighborhood environment

without direct interactions with neighbors. The open-ended structure and opinion coevolve where the opinion dynamics affect the leaver exiting society and associated structural dynamics; A joiner with a random opinion joins. It changes the structural features and the neighborhood settings, and the neighborhood settings in turn affect the opinion dynamics.
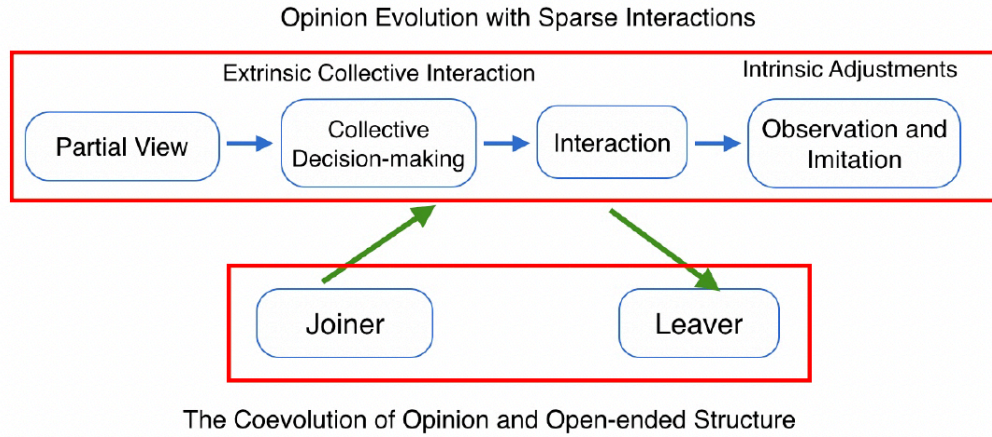
Figure 7: The overview of the SCOOE model.

For agent i being exposed to a new opinion, we assume that agent i has a built-in probability of sticking to its opinion, i.e., a stubbornness probability w_i. It describes the degree to which an agent relies on its original opinion. In contrast, the complement of stubbornness, i.e., an openness probability $1-w_i$, quantifies the degree to which agent i is willing to adopt a new opinion derived from the interaction with other agents. Heterogeneity is produced when agents hold different built-in cognitive features represented by different stubbornness (or openness) probabilities. We assume that stubbornness w_i follows a probability distribution in the population, like a Poisson or Gaussian distribution. In the experimental section below, we report on the influence of different stubbornness distributions. The opinion O_i of an agent i is represented by a real number in the continuous interval [0,1]. It describes the degree to which an agent believes the propagated information, e.g., news or rumors. A higher value of opinion O_i

means that agent i believes the propagated information more strongly. Initially, each agent is assigned a random opinion, i.e., a random number in the range [0,1].

1. Opinion Dynamics with Sparse Interactions

We first discuss the opinions dynamics. The critical theme of opinion dynamics is sparse interaction and incomplete information. People do not serially poll the neighborhood in their social networks to update the opinion but are mainly influenced through considering the joint opinion of others as their feedback (Forsyth, 2018). This contrasts with most agent-based models which are formulated with such complete information assumptions, e.g., polling the entire neighborhood to select neighbors and conduct dense interactions serially to update opinions (Forsyth, 2018; Noorazar, 2020).

The literature from psychology suggests two types of motivations for humans to change behavior, extrinsic motivation (people are assimilated into extrinsic environments) and intrinsic motivation (people are motivated by internal desire) (Deci et al., 2001). We take inspiration from this and assume two types of actions forming the sparse interaction, extrinsic collective interactions and an intrinsic observation mechanism. The interplay between these two actions enhances group opinion evolution. However, they play different roles in various stages of the model dynamics reported in the experimental section.

We first study the extrinsic collective interaction with incomplete information -- we introduce a collective decision-making approach to incorporate the sparse joint opinion formation and interaction based on a limited neighborhood (i.e., incomplete information). Though several collective decision-making approaches have been proposed in discrete-opinion models, e.g., majority voting (Choi & Goh, 2018), this approach in continuous-opinion models has not been fully developed so far. Suppose a focal agent i with a connection degree d_i in its

31

social network is able to only access a random subset of neighbors $i_1, i_2, ..., i_j$, where j is randomly

chosen and satisfies $1 \leq j \leq d_i$. This assumption means incomplete information by a limited view

and only partial access to neighbors, and it allows more dynamic interactions, e.g., an agent will

not interact with its entire neighborhood. Agent i generates a joint opinion O_i^{Joint} of its random

partial neighborhood, rather than by interacting with all its neighbors or certain neighbors with

similar opinions serially. An intuitive way to generate a joint opinion is by taking the weighted

average of the selected neighbors' opinions (Friedkin & Johnsen, 1990). The weights assigned to

different neighbors are proportional to their relative connection degree strength, as shown in

Equation 4, where d_{i_k} is the degree of neighbor i_k, $k \in [1, j]$.

$$O_i^{Joint} = \sum_{k=1}^{j}(d_{i_k}/\sum_{k=1}^{j}d_{i_k}) \times O_{i_k}$$

(4)

Thus, the more a neighbor is connected in the local network (measured by its relative

connection strength), the greater its weight and influence on the joint opinion in the collective

decision-making process.

Another critical factor in designing an interaction protocol is confirmation bias. That is,

people collect and interpret information selectively by trying to follow their original bias (e.g.,

their original opinions) (Plous, 1993). The most widely adopted interaction protocol with

confirmation bias is a bounded confidence model where rational agents owning the perfect

information poll their entire neighborhoods and select others to interact only if their opinions fall

within a confidence interval (Deffuant et al., 2000; Gomez-Serrano et al., 2012; Wang & Shang,

2015). Here we take inspiration from game theory and model this as an opinion interaction game

with confirmation bias among bounded-rational agents with limited information. So, after

generating a weighted joint opinion based on limited neighborhood information, agent i with opinion O_i receives a payoff R_i represented by Equation 5.

$$R_i = 1 - |O_i - O_i^{Joint}| \qquad (5)$$

Equation 5 means that if the opinion O_i of agent i is very different from the joint opinion in its selected neighborhood (the local environment), it receives a low payoff. Neighborhoods with more similar opinions are considered more trustworthy, thus, resulting in a higher payoff. After considering the collective interaction by the game-playing and interaction with the joint opinion, the focal agent i adapts to the neighborhood. Suppose the stubbornness of i is ω_i and its openness is $1-\omega_i$, then the adapted opinion $O_i^{Adapted}$ of agent i is calculated by Equation 6. It represents a combination of relying on its original opinion and accepting a new opinion (Centola & Macy, 2007; Friedkin & Johnsen, 2009).

$$O_i^{Adapted} = O_i \times \omega_i + O_i^{Joint} \times (1 - \omega_i) \qquad (6)$$

We have now seen how agents take advantage of extrinsic collective information within their incomplete neighborhoods. Agents also observe the local environment to adjust their opinions to seek a higher payoff. This is driven by intrinsic motivation. People sometimes engage in an activity just because they are drawn to do it (Golman & Loewenstein, 2018; Ryan & Deci, 2000). We apply the imitation rule here that does not need direct interactions but transforms information in the population through observation and self- adjustment (Szabo & Toke, 1998; Pan et al., 2018). Again, a focal agent i only accesses a random partial neighborhood as its observation environment. For these random neighbors, the focal agent i holds a probability W_{i,i_r} to imitate the local best-performing neighbor i_r (i.e., the neighbor with the highest cumulative payoff) by adopting i_r's opinion as its own opinion. We still follow Szabo and Toke (Szabo & Toke, 1998), as shown in Equation 3, to model the imitation probability, which is

proportional to the gap between cumulative payoffs of the focal agent and the selected neighbor, and we set the noise to $\mu = 1.5$. Thus, agents observe the environment and keep a close eye on the cumulative payoff gap. They then adjust their opinions voluntarily without direct interactions to achieve a greater payoff and a better position in society.

In summary, we introduce (i) sparse opinion updates by taking incomplete information-based collective decision-making into account, and (ii) observation and self-adjustment of opinions without direct interaction with neighbors. Sparse interactions are achieved.

2.   Open-ended Structural Dynamics

This section presents the open-ended structural dynamics with leaving and joining agents and the opinion-structure coevolution.

At each time step, the agent with the lowest cumulative payoff leaves the society, which models an intention to exit a society where most individuals have fairly different positions (e.g., opinions). The stubbornness and openness of the leaver are recorded. All adjacent edges of this agent are removed from the society upon leaving. As society evolves, leaver-driven structural dynamics will demonstrate the confirmation bias more strongly because stubborn agents with opinions fairly different from others will have a low payoff leading to their removal from the model. Opinion dynamics affect the cumulative payoff, influence which agents become leavers, and thus drive the structural coevolution of the system.

At each time step, a newcomer v will join. As society evolves, the community structure constantly changes. Agent v has incomplete information about different communities. It detects the real-time community structure and attempts to join a random $community_v$ by connecting to randomly selected nodes within $community_v$. We assign a random opinion O_v to v and the

recorded stubbornness/openness of the leaver (see above) to v in order to keep the cognitive ability distribution stable within the society. Note that the cumulative payoff E_v of the incoming agent v is not comparable to that of existing agents when calculating the imitation probability and removing leavers, especially for long-term experiments. We accordingly assume that given v joining at time step t_v, agent v's initialized cumulative payoff E_v is adjusted by the corresponding payoff R^{t_v} at time t_v multiplied by the number of completed interactions t_v. After initialization, the cumulative payoff E_v is calculated by regularly adding the corresponding payoff R^t at each time step t until v is removed or the system terminates globally.

After joining a community, the newcomer v chooses and connects to a node u in another community. We apply the preferential attachment principle (i.e., nodes with a higher connection degree have a stronger ability to attract new nodes added to the network) because "the rich getting richer" phenomenon is widely observed in real-world societies (Barabasi & Albert, 1999). Thus, the probability $p_{v,u}$ for v choosing u to connect to follows Equation 7.

$$\forall u \in (G - community_v): p_{u,v} \propto \frac{d_u}{\sum d_u} \tag{7}$$

$G-community_v$ represents all of the other communities except for $community_v$, which the new node v joins. d_u represents the degree of node u. If only one community exists as the society evolves, the new node joins by connecting to only one node following preferential attachment. An algorithmic view of the SCOOE model is shown in Algorithm 3.

Algorithm 3: The SCOOE model

Algorithm 3   The SCOOE Model

---

1   Initialize society, opinions, stubbornness, openness, payoffs, and
    neighbors;

2   **for** *each time step* $t(t = 1, \cdots, T)$ **do**

3   //Opinion Dynamics with Sparse Interactions;

4   **for** *each agent* i *in the society* **do**

5   Agent i accesses partial random neighbors and generates a
       joint opinion O_i^{Joint} (see Equation 4);

6   Agent i receives the corresponding payoff R_i and updates the
       cumulative payoff E_i by O_i^{Joint} (see Equation 5);

7   Agent i is assimilated into the neighborhood with the adapted
       opinion $O_i^{Adapted}$ (see Equation 6);

8   Agent i voluntarily adjusts its opinion with an imitation
       probability W_{i,i_r} (see Equation 3);

9   **end**

10   //The Coevolution of Open-ended Structure and Opinions;

11   The agent with the lowest cumulative payoff leaves the society;

12   Assign a random opinion O_v to a newcomer v;

13   Assign the stubbornness/openness of the leaver to v;

14   Adjust the cumulative payoff E_v of v;

15   v detects the real-time community structure;

16   **if** *multiple communities ($\geq$ 2) can be found in the society* **then**

17   v joins a random community by connecting to partial random
       nodes within this community;

18   An edge is generated connecting v to a node u in another
       community, following the preferential attachment (see
       Equation 7);

19   **else**

20   v connects to a node u in the society following the preferential
       attachment (see Equation 7);

21   **end**

22 **end**

C. Model Performance

To test the model, we create two small-world networks holding 500 nodes each to model two

physically separated groups of people interacting to a certain degree. So, randomly chosen edges

connect the two small-world networks. We apply the Watts–Strogatz model to generate an

individual small-world network (Watts & Strogatz, 1998). Each node is connected to four nearest

36

neighbors. The rewiring probability is set to 0.05. This structure constitutes the agent society, with each node representing an agent. A focal agent will only consider those agents connected by edges as neighbors and conduct sparse interactions based on the neighborhood. For initialization, we follow prior work from the psychology and computing realms (Das et al., 2014; Meehl, 1992) and set the stubbornness distribution to a Gaussian distribution with a mean of 0.5 and a standard deviation of 0.25. These parameters are chosen so that most values lie between 0 and 1. In addition, we apply a cut-off so that generated random numbers can only lie between 0 and 1, i.e., we constrain stubbornness to the interval between 0 to 1, as shown in Figure 8 (a). Imitation noise is set to $\mu = 1.5$. The simulation is run for 450 Monte Carlo time steps.

1. How Do Group Opinions Evolve?

We first study how far the group opinions evolve away from their initial states, measured by the variance dynamics shown in Figure 8. The opinions of agents are reasonably different at the start because agents are assigned random opinions initially. As society evolves, we find two stages of evolution: a fast-decay phase (i.e., the variance of group opinions dramatically decreases from 0.084 to 0.005) and a slow-decrease phase (i.e., the variance slowly continues dropping to 0.003 at the end of the simulation). It is interesting to find that the final opinions are in a relatively narrow band and less polarized without firmly believing or strongly unbelieving the rumors among the agent population, even with some agents never changing their opinions (stubbornness =1) but being removed by the model. The Gaussian stubbornness distribution is also evenly distributed. The majority of the population keeps a balance between maintaining their original opinions and accepting a new opinion. Mirroring reality, we find that agents are more likely to stay open-minded to propagated news/rumors during in an *open-ended* society.

(a) Default Gaussian stubbornness

(b) Opinion initialization

(c) Evolved opinions under Gaussian stubbornness distribution

(d) Variance dynamics
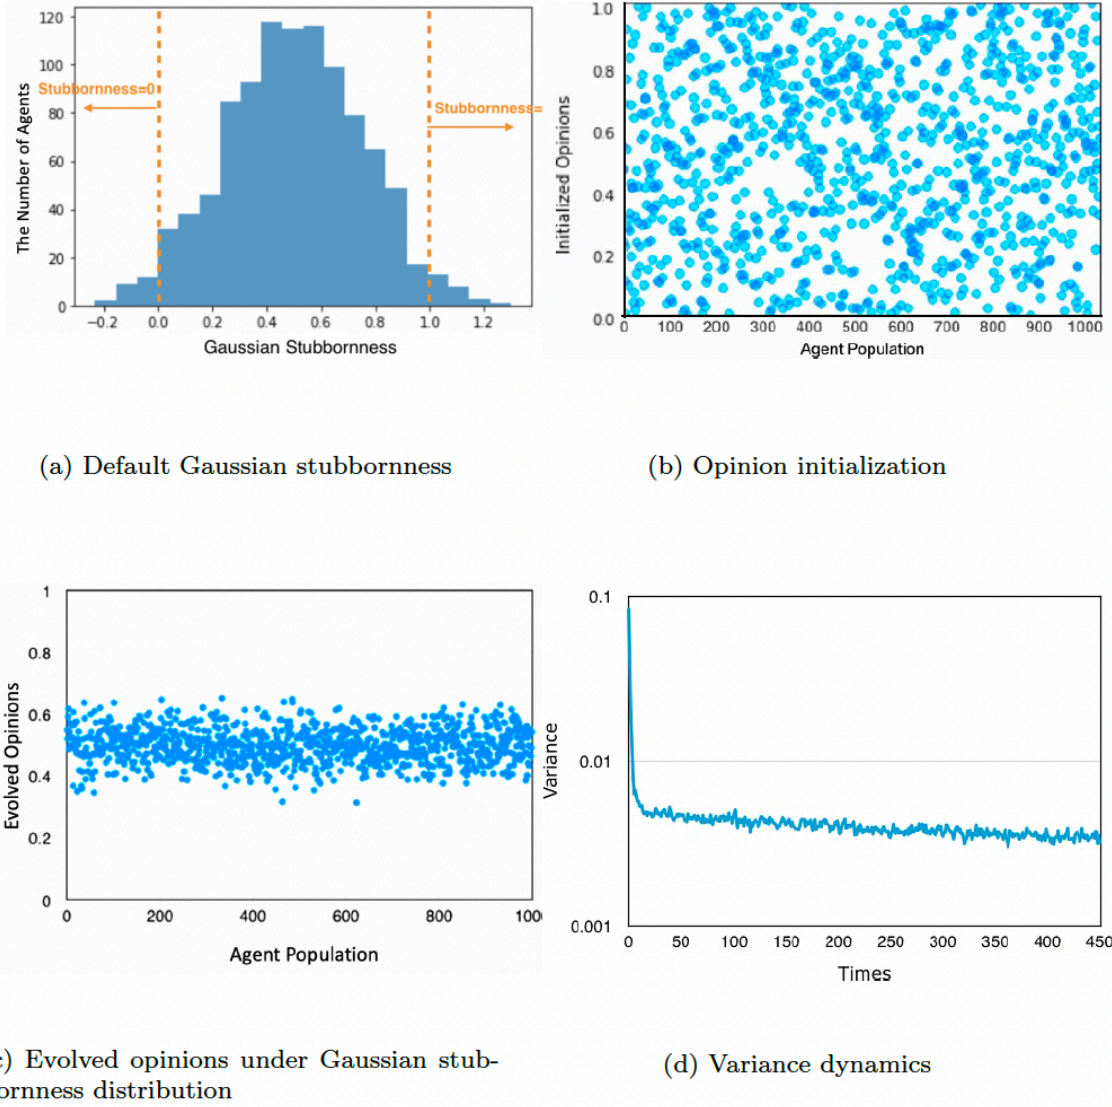
Figure 8: The opinion dynamics under default Gaussian stubbornness.

2. How Do the Opinion Dynamics Shape the Structural Co-dynamics?

For this question, we primarily focus on the dynamics of the clustering coefficient, average path length, degree distribution, and community structure. For real-time community detection, we use the most widely used method, namely the modularity-based method (Clauset et al., 2004).

The society coevolves to be a holistically dense small-world structure with a heavy-tailed degree distribution. As shown in Figures 9 and 10, we find an increase in the clustering

coefficient and average degree, as well as a decrease in average path length and the number of communities. We initialize the society as two interconnected small-world networks. The random edges between them change the initialized small-world features by randomizing them to a certain degree. Thus, we find a chaotic society initially with 26 detected tiny communities and a coevolved society with nine segregated communities by the modularity-based method (Clauset et al., 2004). We also observe that the coevolved society has a small-world feature with a high clustering coefficient. It coevolves to be a more tightly knit group with dense connection degrees, high information transmission efficiency, and a low average path length due to network homophily. That is, the final opinions of the population are relatively consistent, leading to an increase in payoff and a decrease in conflicts (e.g., confirmation bias for fairly different opinions) in the game-playing upon interactions. This coevolutionary trend of the structure in turn boosts the evolution of a global opinion (McPherson et al., 2001)

Although some small-world generation models, e.g., the Kleinberg model (Kleinberg, 2002), do not generate heavy-tailed degree distributions, it is not surprising to find a heavy-tailed degree distribution appearing in the SCOOE model. The advantages of "the rich" become significant eventually because of preferentially added joiners. Specifically, we calculate the proportion $P(d)$ of nodes with connection degree d. We find that the relationships between node proportion $P(d)$ and node degree d can be approximated by a linear relationship $\log[P(d)] \sim (-\gamma)$ * $\log(d)$ with a negative slope $-\gamma = -2.758$ through linear regression within a 95% confidence interval. Note that the data points in Figure 9 (f) represent the average degree and the node proportion in different degree ranges. We only study the nodes in these degree ranges because they fill most of the network. These nodes are enough to illustrate a linear relationship.
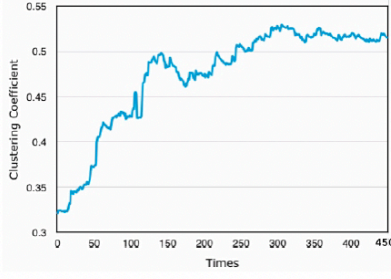
(a) Clustering coefficient
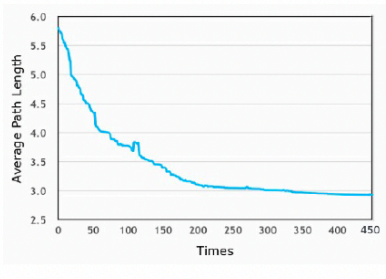
(b) Average path length
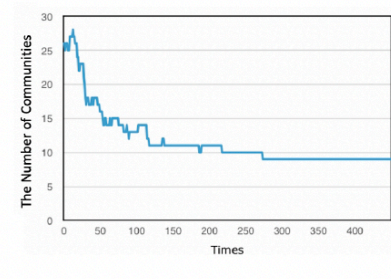
(c) The Number of communities
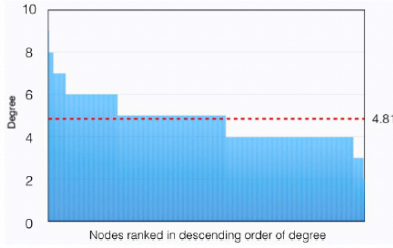
(d) Initialized degree distribution
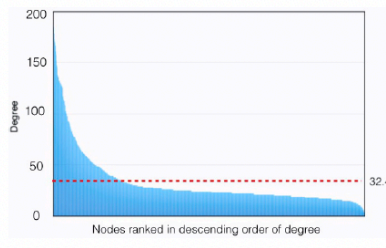
(e) Coevolved degree distribution
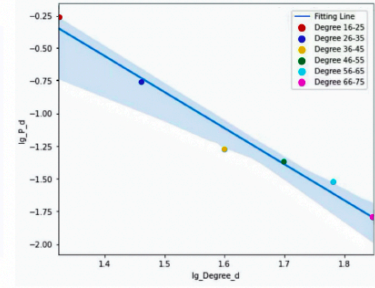
(f) The relationships between node degree and node proportion

Figure 9: The structural dynamics.

We also find an emergent dialectic relationship between community segregation and cohesion. Cohesion is a concept of togetherness and connectedness among nodes within a network. There is no unified definition of cohesion because it depends on the context. Previous literature has referred to it as cliques/communities, clusters, or average degree (Kolaczyk & Csardi, 2014). Figure 10 shows our assessment of the community segregation and cohesion.
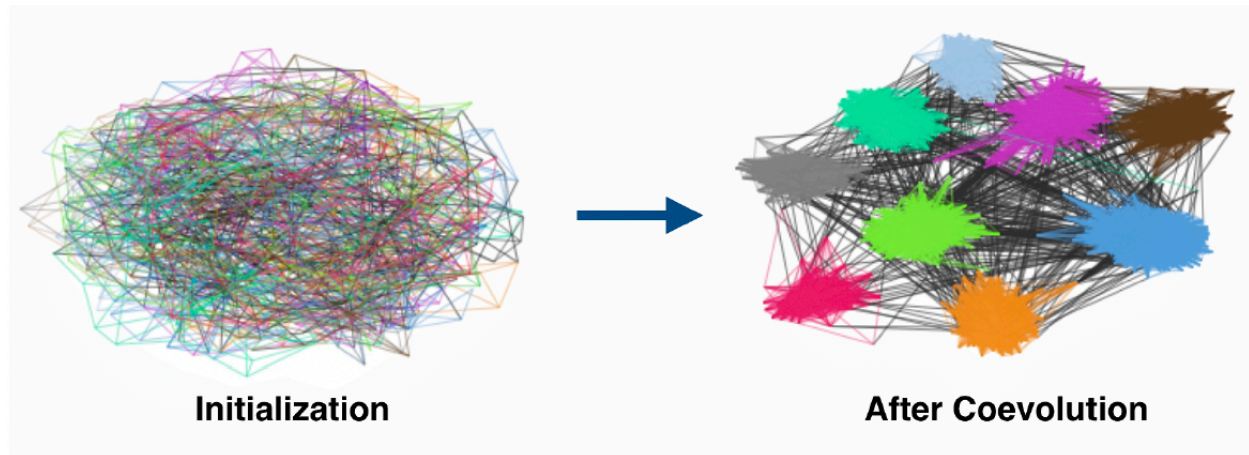
Figure 10: The community structure.

As shown in Figure 10, the initialized society is desegregated and chaotic with a low level of cohesion (i.e., with a low average degree and clustering coefficient). As society coevolves, we find it has a clear pattern of fewer segregated communities that become densely clustered (i.e., with a high average degree and clustering coefficient). Agents have disconnected social networks initially but highly cohesive social clusters eventually. It is interesting to note that society becomes segregated but dense spontaneously and simultaneously with a global consensus and cohesion, but without multiple local-opinion "barycenters" that might emerge aligned with segregated communities (Gomez-Serrano et al., 2012) Mirroring reality, as previous work suggests (Neal & Neal, 2014), a widely observed example in the real world is policy-making to reduce detrimental residential segregation. A widely adopted approach to introduce desegregated neighborhoods and reduce residential segregation is to improve cohesion, e.g., dense connections. However, a paradox exists between community segregation and cohesion. The society evolves to be dense with segregated communities, whereas a desegregated society is not as cohesive as we would expect.

41

3.   What Factors Affect the Evolved Global Opinion?

It is reasonable to suspect that the degree of stubbornness affects the emergence of a global opinion. Additionally, the SCOOE model incorporates multiple types of dynamics. What effect do these dynamics have on the evolution of a final opinion? This section will address these questions. To study the influence of stubbornness distributions, we also test a Beta distribution and a Poisson distribution. We initialize the Beta distribution with two positive shape parameters $\alpha = 7$ and $\beta = 1$, and the Poisson distribution with the expected rate of occurrences $\lambda = 1$. We normalize the two generated distributions with the maximum value representing stubbornness=1. The evolved opinions in these two cases are shown in Figure 11. The variance comparison with different stubbornness distributions is shown in Figure 12 (a). The variance comparison is defined as the ratio of the opinion variance for the Poisson/Beta stubbornness distributions to that for the baseline Gaussian stubbornness distribution at each time step t, t in $[0, 450]$. The variance dynamics with different sparse interaction mechanisms in a population with a Gaussian stubbornness distribution are shown in Figure 12 (b).

Stubbornness is generally small in a population with Poisson stubbornness. Agents are very flexible to become followers of the propagated news/rumor. As a result, it will be easier to pass the fast-decay phase, and we observe an initial lower variance than the baseline shown in Figure 12 (a). Because of the flexibility in updating opinions, evolved opinions are still inconsistent at the end of 450 time steps, and the final variance is relatively large. In contrast, the agent population generally has much higher Beta stubbornness. Accordingly, we find an initial increase in the variance ratio to pass the fast-decay phase shown in Figure 12 (a). Because of the high stubbornness, final opinions are stable with few changes, and lower final variance than the baseline can be observed.

It is challenging to drive the global opinion evolution among a stubborn population, e.g., the initially weak emergence of the global opinion in the population with high Beta stubbornness. However, it is interesting to find the most unified global consensus in such a society with many agents only weakly changing opinions. This unusual phenomenon is due to the open-endedness of society. The most stubborn agents will be considered maladapted to the environment and removed as society evolves. Agents will be assimilated by agents who surround them. No matter the initial opinions they hold in stubborn crowds, they will finally have a relatively unified group consensus after the *long-term* interactions and the slow assimilation of opinions crowding out dissidents in an *open-ended* society. We can say that these high stubbornness values serve as a "wall" — newcomers with similar opinions will be accepted, while newcomers with opinions out of this range will be removed quickly.

We additionally test the model without the intrinsic self-adjustment mechanism as shown in Figure 12 (b). A widely studied contagion phenomenon in social networks is that the chance to adopt a contested "innovation" (e.g., firmly believing a piece of rumor) will be smaller for an individual with more neighbors (Centola & Macy, 2007; Granovetter, 1973). When a focal agent aggregates the joint opinion by collective decision-making, extreme opinions (e.g., a strong endorsement) of selected neighbors are neutralized by weighted averaging. This effect will be more significant for high-degree nodes, given the larger share of their neighbors. On the other hand, high-degree nodes with a fewer likelihood of being extreme have a more substantial impact on the weighted aggregation method and a more extensive influence range. At the same time, collective interactions decrease the probability of interacting directly with extreme agents and being affected by them. So, the extrinsic collective interaction mechanism boosts the emergence of a global consensus, as shown in the similar trends of the fast-decay phase in the

two cases in Figure 12 (b). It plays fewer roles when the population rapidly reaches a pre-consensus (the start of the slow-decrease phase in variance dynamics), given the constantly adapted local interaction environment with the randomness to select neighbors, the joiners/leavers, and a constant injection of new opinions. The intrinsic adjustment mechanism continues to further the emergence of a global consensus and weakens conflicts by direct imitation. It can be said that extrinsic collective interactions primarily play a role in the fast-decay phase of the variance dynamics, whereas intrinsic adjustments mainly play a role in the slow-decrease phase. Their interplay works to enhance the evolution of a global opinion. Note that when we set the self-adjustment noise μ to a very large value, we can observe similar results to the case of removing the self-adjustment mechanism.

(a) Poisson stubbornness distribution

(b) Beta stubbornness distribution
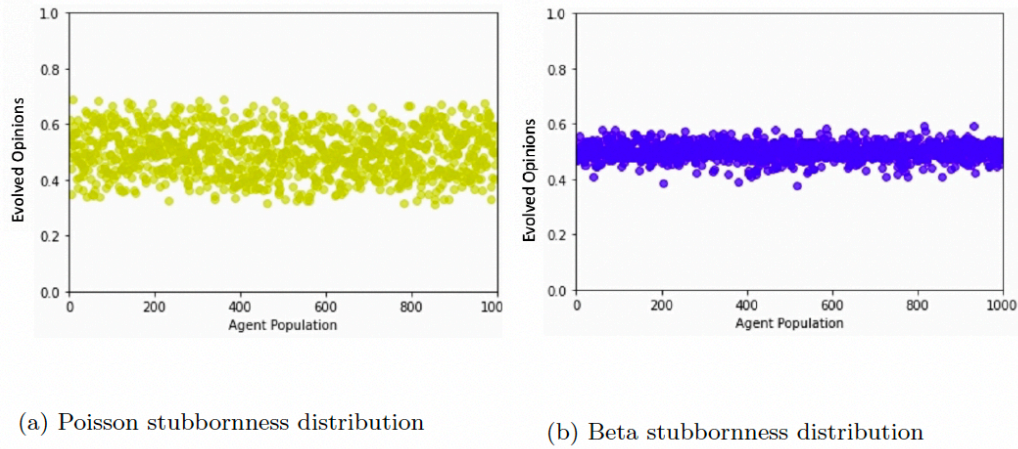
Figure 11: The evolved opinions with different stubbornness distributions.

(a) Variance comparison with different stubbornness distributions

(b) The dynamics of the variance of opinions with/without self-adjustment mechanism
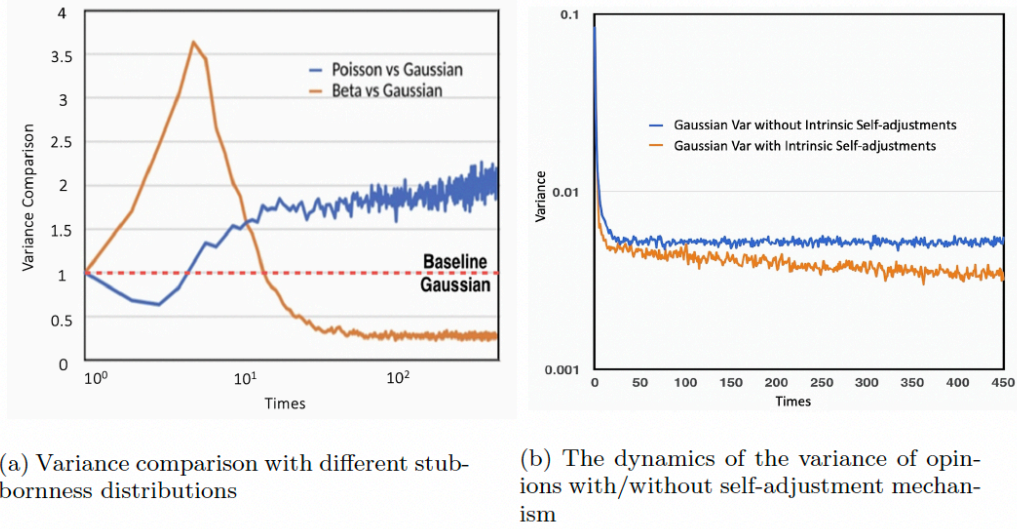
Figure 12: Opinion dynamics with different distributions and mechanisms.

D. Discussions

We revisit and discuss the proposed mechanisms by focusing on their effects on the consensus evolution within groups.

Lean and fast decision strategies can be produced with incomplete information. A broad assumption in the widely cited bounded confidence model is that rational agents owning the perfect information poll their neighborhood and select neighbors to interact with only if their opinions are sufficiently close to their own. This assumption facilitates polarization and global conflicts (Gomez-Serrano et al., 2012). It has been widely recognized that it is difficult to evolve a global consensus for large population sizes (Iniguez et al., 2009; Kou et al., 2012; Yu et al., 2015), because local consensus might be distributed in a society. As a result, such a system needs more bottom-level interactions to pass the formation of the local consensus. Our results validate several earlier findings with different mechanisms and remarkably boost the evolution even in a stubborn population (Lim et al., 2014; Semonsen et al., 2018; Weisbuch et al., 2003). Unlike some bounded confidence models, e.g., Deffuant et al., 2000 and Gomez-Serrano et al., 2012,

here we start by assuming that bounded-rational agents only access a partial neighborhood (incomplete information) to aggregate a joint opinion. Confirmation bias is represented by the stipulation that adopting more similar opinions will bring a higher payoff. We find that conflicts among bounded-rational agents are weakened globally and rapidly. Bounded rationality with incomplete information forms lean and fast decision strategies to reduce conflicts under uncertainties, whereas complete information weakens group coordination, as suggested by some literature from psychology (Gigerenzer & Selten, 2002).

Open-endedness enables permanently novel opinions. We find that eventually evolved opinions are wholly unified in some closed-society models (Semonsen et al., 2018). The continuous addition and removal of agents and the structure/opinion dynamics they bring with them influence neighbors and their surroundings in a cascading fashion. Though the designed mechanisms strongly facilitate the evolution, it is impossible to reach a highly unified global consensus. One can only approach it no matter whether the randomness or noise exists, as the slow-decrease phase in variance dynamics. It can also be said that the SCOOE model is robust to boost and enhance the evolution of a global opinion as it successfully defends against the interference of a constant injection of novel opinions.

The interplay between sparse interaction and open-ended structure reduces the echo chamber effect. The echo chamber effect in social media studies describes a situation where local opinions are reinforced by repetition inside a closed society and insulated from rebuttal or different opinions (confirmation bias). Surprisingly, a substantial body of research indicates that people are not as polarized as we would expect in the echo chamber, both empirically (Balietti et al., 2021; Bar-Gill et al., 2020; Hosseinmardi et al., 2021; Shore et al., 2016) and theoretically (Lim et al., 2014; Semonsen et al., 2018; Weisbuch et al., 2003). We offer two possible

theoretical justifications for this apparent discrepancy between evidence and intuition: From the perspective of opinion dynamics, the focal agent considers the collective opinion based on a limited view of the neighborhood, which reduces polarization quickly. From a structural dynamic perspective, the society in our model (and also in the real world) is open-ended and constantly changing. It imparts persistent dynamics on the neighborhood structure, resulting in neighbors with whom the focal agent interacts being neither isolated nor static. When we examine previous models based on a closed structure, some work has shown global/local polarization and extreme opinions (Banisch & Olbrich, 2019; Hegselmann et al., 2002; Mathias et al., 2017). The open-endedness feature with a constant injection of novel opinions in the SCOOE model helps a population defend against the echo chamber effect and stay open-minded. It reduces the chances of extreme results because extreme agents are likely to be removed from society. It also mirrors the findings of a global consensus formation in a population with high Beta stubbornness. In general, we believe that opinion and structural mechanisms are inextricably linked and that their interplay helps reduce polarization.

IV.   TOWARDS COLLECTIVE INTELLIGENCE

This section provides a thorough analysis of proposed models and relates their insights into collective intelligence to widespread applications in other subfields of computer science.

Our first model focuses on the evolution of cooperation with two variants. The first variant discusses a framework based on evolutionary selection: the parent agent produces offspring based on its performance in the entire society (cumulative payoff). The offspring inherit the parent's behavior with mutation. Additionally, they improve their inherited behavior via two methods: collective learning and social learning. Collective learning signifies that the focal agent generates an overall best action via collective voting from the set of best actions learned through reinforcement learning with different neighbors. For social learning, we employ one of the simplest methods, imitation, in which the agent imitates others' actions according to the difference in the cumulative payoff. In the second variant, we transfer this evolutionary selection framework to a hierarchical network, with the subordinate agent reporting to the supervisor after learning the overall best action through collective learning. The supervisor obtains the best action of the upper layer by social learning with its upper-layer neighbors and communicates this piece of supervised information to the lower-layer subordinates. Subordinates then adjust the learning and exploration rates in their reinforcement learning protocols according to the supervision. We examine how cooperation emerges from selfish agents using these two variants.

In the second model, we focus on opinion dynamics, i.e., how the agent adjusts its opinion in order to better integrate into society. We do not use reinforcement learning as a means of agent interaction. The agent's objective in the first model is to learn the optimal action, a typical Markov decision-making problem - deciding what action to adopt in the next step based

on the action, state, and payoff of the previous step. In contrast, for the study of opinion dynamics, the agent attempts to adapt to the environment by evolving and updating the opinion. We use a Bayesian method to update the existing opinion of the focal agent by a weighted average of existing opinion and the collective opinion from neighbors. Payoffs are lower when arguing with those who have the most divergent opinions. The degree to which agents adapt to the environment, represented by the cumulative payoff, drives them to leave and join society selectively. Opinions and social structure co-evolve, and we examine how coordination emerges through the coevolution from initial opinion chaos.

These two models, despite their different contexts, both focus on a fundamental issue in multi-agent systems, artificial life, and complexity science: the emergence of coordination. Section I divides principal methodologies in building agent-based computational models into three related categories: how agents are represented, how agents interact, and how the agent population is structured. Two proposed models introduce rich dynamics in these three aspects. In this section, we continue to follow these three lines and discuss two proposed models by examining the following three issues:

- First, we adopt an interaction-based perspective on adaptation, i.e., how agents with different representations adapt to the environment through different interaction protocols, such as reinforcement learning and opinion evolution, and how different adaptations impact the macroscopic evolutionary dynamics.
- Second, we adopt a structural view to discuss open-endedness and its effect on the robust emergence of macroscopic coordination.

- Third, we take a broader view of how research on emergent coordination and collective intelligence in the fields of artificial life, evolutionary computation, and complex systems should be exploited to inspire other subfields of computer science.

A. Adaptation

Adaptation is an intricate concept. It describes how a species, an individual, or an agent gradually becomes fitted to the environment, say, the ecological system, society, and artificial system, respectively (Pimm et al., 2016). According to different time scales, adaptation has been classified into three categories (Jablonka & Lamb, 2006; Gershenson, 2010; Aguilar et al., 2014). According to Aguilar et al. (Aguilar et al., 2014), a slow adaptation that occurs over many lifetimes is referred to as evolution. The adaptation that occurs at a moderate rate (one lifetime, for example) is referred to as development (including morphogenesis and cognitive developments). A very rapid adaptation that requires only a small portion of a lifetime is referred to as learning. Note that some forms of adaptation may be counterproductive, hindering the population's ability to survive in its environment. Some examples include the behaviors acquired through social or normative learning, such as adopting unrelated children or altruistic behaviors that do not favor relatives or kin (Staddon, 2016). In our models, agents either "learn" the optimal action in a short time (tens of time steps, as shown in Figures 2, 3, 5, and 6 in the cooperation models) or adapt to the environment in a relatively slow manner. If we consider the time an agent spends in the environment to be its lifetime, it may take a large number of lifetimes for the agent population to "evolve" toward the group coordination in an open-ended environment, as shown in the open-ended opinion dynamics model (the slowly decreasing but non-convergent phase of opinion variance, Figures 8 and 12).

The evolutionary dynamics of a system are profoundly influenced by different modes of adaptation. In this thesis, we first demonstrate that adaptation through collective interaction facilitates evolution. As shown in Figure 2 of model 1, the "bad" effects of individual-level mutations (a small probability of not adopting the best action) that lead to suboptimal collective decision-making derived from the suboptimal action set are mitigated by voting methods such as majority voting. It is also worth noting that agent representation and structure both influence adaptation. The former, agent representation, aims to enhance the model's capacity to represent reality in the multidisciplinary context. In our models, for instance, we implement cognitively heterogeneous agents. This diversity is reflected in agents' capacity to accept and incorporate new perspectives from others. The hierarchical system with unequal payoff distribution is also introduced to represent the sharing between supervisors and subordinates. Different payoff matrices among agents will have an effect on how they learn optimal actions and adapt to society. Although the cooperation among subordinates is more challenging to emerge because they receive a lower payoff for cooperation than their supervisors, we still observe the global cooperation as supervised information influences how subordinates behave. In biology, ecology, and social sciences, there is a broad practice of using heterogeneous agents to represent different roles and functions of entities within the system (Epstein & Axtell, 1996; Macal & North, 2005). The latter is more concerned with the network structure and neighbor composition. When considering a focal agent without a global view, different niches in which the focal agent exists have different effects on its adaptation. When the number of niches is large (as shown in large cluster sizes in the hierarchical cooperation model), it is possible to form some local coordination distributed in the whole society. Crossing local coordination to form global coordination will be more challenging, as shown in Figure 5.

51

One of the most prominent criticisms of contemporary artificial intelligence and machine learning research is its lack of adaptability, as it has sought to predict and control rather than adapt (Aguilar et al., 2014; Gershenson, 2013). This is the root cause of numerous problems in current deep learning research, ranging from insufficient robustness, inflexible and rigid configuration assumptions (say, fixed input and configuration), and inadequate performance in adapting to novel task settings (Ha & Tang, 2022). It is evident that combining evolution and learning (e.g., Lazaridou & Baroni, 2020) or development and learning (e.g., Chrol-Cannon & Jin, 2014) is a promising approach, particularly for programming adaptive agent behaviors. In this thesis, an attempt is made to create adaptive learning. In the hierarchical cooperation model, we propose that the subordinates' learning and exploration rates are adapted based on their supervisors' social learning information (Figures 4 and 6). There is a vast amount of future work worthy of further investigation in real-world scenarios. Recent efforts have been made on this topic across the fields of machine learning, robotics, and evolutionary computation, especially for studying real-time, online, complex strategic interactions, such as automated trading in stock markets and real-time team formation/task allocation in multi-robot systems (Bloembergen et al., 2015).

B. Open-endedness

Evolution is unstable. Much research in evolutionary biology and artificial life indicates that the complexity of evolution exhibits an open-ended, limitless expansion (Bedau, 2007). Artificial life as a field in nature aims to create an artificial system to demonstrate how unrestricted evolutionary progress can be achieved (Bedau, 2007). At the OEE (Open-Ended Evolution) workshop in York, Tylar and colleagues provided a summary of the characteristics of open-

ended evolution: the ongoing generation of adaptive novelty and the ongoing increase in complexity (Tylar et al., 2016). Based on how we define "end," Banzhaf et al. provided three primary classifications of open-endedness: (1) processes that do not terminate; (2) processes without a specific objective; (3) processes that do not terminate *and* have no specific objective (Banzhaf et al., 2016).

This thesis adopts a structural perspective on open-endedness: In models for the evolution of cooperation, complexity and novelty are increased by continuously generating offspring agents (new nodes) connected to their parents via newly generated edges in the network. In the opinion dynamics model, we introduce joiners and leavers. The joiner detects communities in real time and chooses one to join. Dissidents with significantly divergent opinions leave the network. Together, they influence the evolution of the structure. In both models, the newcomer's new behaviors are subsequently integrated into society. Novelty and complexity are constantly introduced in terms of both behavior and structure. As agents with new behaviors join the open-ended evolutionary system, the new behavior is accompanied by a reorganization of the structure, such as a change in edge configuration. Subsequent novelties are not generated solely by newly added agents and behaviors but also by the simultaneous structural dynamics. We can observe a gradual decrease in the variance of opinions that never converge (Figures 8 and 12). In both models, we find that the open-ended structural characteristic does not necessitate a coordination disruption. Instead, the dominant norm is stable even if it constantly evolves in an open-ended manner. Newcomers actively or passively choose to assimilate into it, which somewhat validates prior works from migration acculturation and cultural evolution (see Mesoudi, 2018).

The interplay between open-endedness and adaptation provides a means to improve the model's robustness. Adaptivity of the system enhances the evolution of the global opinion as it successfully defends against the interference of a constant injection of novel opinions. The open-ended society "repairs" the global opinion by expelling dissidents who may disrupt it. Recently, a large body of research in artificial life and evolutionary computation (e.g., Mordvintsev et al., 2020; Sandler et al., 2020) has shown that cellular automata-based neural networks with adaptive interactions between neurons can resist the interference of noise and self-repair when damaged. These works will be discussed in greater detail in the subsequent section on linking the emergent models with other subfields in computer science.

The open-ended evolution in our models deserves further investigation because it is incomplete in its current standing. The ongoing generation of new kinds of entities and their behaviors do exist in our models, but they do not result in a new type of adaptation. That is, new agents continue to adapt, learn, or evolve in the same manner as their predecessors. Simultaneously, new agents select actions from the established action set without generating novelty in the strategy complexity. We should not solely focus on evolving particular coordination from a given set of behaviors. Rather than that, we expect the emergence with an active complexity and boundary adjustment. How to design a comprehensive, open-ended evolutionary system with the coexistence of new adaptations, new types of entities, major transitions (e.g., emergence), and the evolved open-endedness (open-endedness as a consequence of evolution) while keeping the model elegant is a fascinating direction (Tylar et al., 2016; Pattee & Sayama, 2019).

C. Collective Intelligence

Our core in this thesis, group coordination as a form of collective intelligence, is one of the central issues in complex systems, multi-agent systems, and artificial life. The rational choice for an individual agent in the first work is to be uncooperative. The agent population learns to cooperate through the processes of learning and evolutionary selection, as cooperation is unattractive to the individual but advantageous to the population. In the second work, the agent's opinion is very different initially, and there is a constant injection of newcomers with different opinions. Through the evolutionary process, the conflict among agents with different opinions is diminished across the whole society. In our models, the emergence of collective altruism from individual selfishness and the group coordination from initial conflicts are observed frequently, benefitting the population. We hope that the insights of collective intelligence presented in this thesis will serve as a catalyst for other communities in computer science. We list some potential applications below.

- Unsupervised learning: Unsupervised learning is a self-organizing process in nature due to the absence of supervision. Artificial life and collective intelligence will unquestionably benefit the development of robust and adaptive unsupervised learning algorithms (Rand, 2006). Clustering, for instance, is a typical unsupervised learning task. A general framework is to model a data instance as an agent in a multi-agent system and implement clustering techniques based on agent attribute similarity (say, distance in the vector space). The process of data clustering is the emergence of agent communities (Chaimontree et al., 2012; Bu et al., 2017).

- Robust deep learning: An early idea of this is the Neural Cellular Automata (NCA) (Wulff & Hertz, 1992). A unit in the CA lattice is represented by a neuron of a neural network. Instead of pre-defined CA rules, the neural network learns local interaction rules and updates each neuron's state based on the interaction with local neighbors. Some recent work improves the NCA from different perspectives. Evolutionary algorithms have been used to train the learning of neural architecture, weights, and local interaction rules (Nichele et al., 2017). Some work is inspired by the NCA and creates noise-resistant deep learning algorithms by adaptively learning a coherent representation of features (Zhou et al., 2022). Adaptive interactions are introduced between neurons using the attention mechanism. The most relevant vectors to the output are shared between neurons and can be flexibly attributed to the highest weights in a collaborative manner (Tang & Ha, 2021; Jian et al., 2019). High resilience to damage is illustrated when almost the entire input image is removed, and the deep learning system is still able to regenerate it (Mordvintsev et al., 2020). Ha and Tang developed a systematical survey discussing the inspiration on robustness, generalizability, and adaptivity from collective intelligence in the current deep learning architectures from four perspectives: (1) image processing, (2) deep reinforcement learning, (3) multi-agent learning, and (4) meta-learning (Ha & Tang, 2022).
- Collective robotics: The research on collective robotics, sometimes referred to as "hard" artificial life, transits a large number of ideas in agent interaction, scheduling, and learning to real-world robotic teams (Ferrante et al., 2015). We can program robots for a variety of roles and establish a series of interaction protocols to study tasks like team formation. The emergent property permits robotic teams to manage real-time, online

tasks more effectively than centralized control (Pitonakova et al., 2014). For instance, centralized collaboration can emerge from the information sharing between decentralized modules to collaboratively optimize a shared global reward in soft-bodied robot teams (Huang et al., 2020).

- Optimization: When agents are viewed as independent problem solvers, the collaboration that emerges between them can solve complex optimization problems that a single agent cannot. A relatively early optimization algorithm based on the concept of agents' emergent collaboration and communication is the cultural algorithm. It is a collective search process in which the agent population and belief space (different knowledge of individual agents in the search space) coevolve by generating a more knowledgeable population and simultaneously updating individual agents' belief space (Reynolds, 1994). The cultural algorithm has inspired substantial research into metaheuristic algorithms and numerous real-life applications, particularly in search-based software engineering. Some instances include optimally assigning different programmers to different tasks (Harman, 2007), bug locating and auto-repair with heuristic approaches (Le Goues et al., 2012), and the co-evolution of programs (prey) and unit tests (predator) (Arcuri & Yao, 2008).

Agent-based models are rooted in the fields of artificial life/evolutionary computing and fundamentally interplay with learning, behavioral modeling, cognitive science, and network science. There is a substantial amount of significant work on these topics that not only provides insights into evolutionary, collective intelligence but also inspires other subfields of computer science and even other disciplines. As the reader can see, the context of our thesis closely relates to a variety of social science topics, including opinion polarization, social networks, cooperative behavior, and voting. Since their inception, agent-based models and artificial life have been a

natural fit with evolutionary biology. They have produced a large body of classic work (see the review paper by Murphy et al., 2020). In studies of human society, another evolutionary, open-ended, complex adaptive system, similar evolutionary models to artificial life have not received the same level of attention (Kim & Cho, 2006; Arthur, 2014). Our thesis ultimately aims to contribute to the growth of digital evolution, artificial life, and their integration with other topics, such as network science, evolutionary game theory, and machine learning in a cross-disciplinary manner.

V.   CONCLUSION

This thesis studies the evolution of cooperation and opinion dynamics viewed from agent representation, agent interaction, and agent group structure. We are primarily concerned with elucidating agent heterogeneity, different modes of interaction, network structure, and their profound effects on evolutionary dynamics. In addition, we establish a connection between two models and a broader topic, collective intelligence, in terms of interaction-based adaptation and structural open-endedness. We hope that the concepts of evolution, emergence, and coordination presented in the thesis will be beneficial and closely integrated into other fields in order to increase the adaptability, flexibility, and robustness of models.

BIBLIOGRAPHY

BIBLIOGRAPHY

Adami, Christoph, Jory Schossau, and Arend Hintze. "Evolutionary game theory using agent-based methods." *Physics of Life Reviews* 19 (2016): 1-26.

Aguilar, Wendy, Guillermo Santamaría-Bonfil, Tom Froese, and Carlos Gershenson. "The past, present, and future of artificial life." *Frontiers in Robotics and AI* 1 (2014): 8.

Arcuri, Andrea, and Xin Yao. "A novel co-evolutionary approach to automatic software bug fixing." In *2008 IEEE Congress on Evolutionary Computation (IEEE World Congress on Computational Intelligence)*, pp. 162-168. IEEE, 2008.

Arthur, W. Brian. *Complexity and the economy.* Oxford University Press, 2014.

Aspray, William. "John von Neumann's contributions to computing and computer science." *Annals of the History of Computing* 11.3 (1989): 189-195.

Attenberg, Lee. *Modularity: Understanding the development and evolution of natural complex systems*. MIT Press, 2005.

Axelrod, Robert, and William D. Hamilton. "The evolution of cooperation." *Science* 211.4489 (1981): 1390-1396.

Bandura, Albert, and Richard H. Walters. "Social learning theory." Vol. 1. *Prentice Hall: Englewood Cliffs*, 1977.

Banisch, Sven, and Eckehard Olbrich. "Opinion polarization by learning from social feedback." *The Journal of Mathematical Sociology* 43.2 (2019): 76-103.

Banzhaf, Wolfgang, Bert Baumgaertner, Guillaume Beslon, René Doursat, James A. Foster, Barry McMullin, Vinicius Veloso de Melo, Thomas Miconi, Lee Spector, Susan Stepney, and Roger White. "Defining and simulating open-ended novelty: Requirements, guidelines, and challenges." *Theory in Biosciences* 135, no. 3 (2016): 131-161.

Bao, Honglin, Qiqige Wuyun, and Wolfgang Banzhaf. "Evolution of cooperation through genetic collective learning and imitation in multiagent societies*." In Proceedings of ALIFE 2018: The 2018 Conference on Artificial Life.* MIT Press, 2018.

Bar-Gill, Sagit, Yael Inbar, and Shachar Reichman. "The impact of social vs. nonsocial referring channels on online news consumption." *Management Science* 67.4 (2021): 2420-2447.

Barabási, Albert-László, and Réka Albert. "Emergence of scaling in random networks." *Science* 286.5439 (1999): 509-512.

Barbati, Maria, Giuseppe Bruno, and Andrea Genovese. "Applications of agent-based models for optimization problems: A literature review." *Expert Systems with Applications* 39.5 (2012): 6020-6028.

Bedau, Mark A. "Artificial life." In *Philosophy of Biology*, pp. 585-603. North-Holland, 2007.

Bloembergen, Daan, Karl Tuyls, Daniel Hennes, and Michael Kaisers. "Evolutionary dynamics of multi-agent learning: A survey." *Journal of Artificial Intelligence Research* 53 (2015): 659-697.

Bowling, Michael, and Manuela Veloso. "Rational and convergent learning in stochastic games." *International Joint Conference on Artificial Intelligence*. Vol. 17. No. 1. Lawrence Erlbaum Associates Ltd, 2001.

Bratman, Michael. *Intention, plans, and practical reason.* CSLI Publications, 1987.

Bu, Zhan, Guangliang Gao, Hui-Jia Li, and Jie Cao. "CAMAS: A cluster-aware multiagent system for attributed graph clustering." *Information Fusion* 37 (2017): 10-21.

Burns, Tom R., and Helena Flam. *The shaping of social organization: Social rule system theory with applications*. Sage Publications, 1987.

Centola, Damon, and Michael Macy. "Complex contagions and the weakness of long ties." *American Journal of Sociology* 113.3 (2007): 702-734.

Chaimontree, Santhana, Katie Atkinson, and Frans Coenen. "A framework for multi-agent based clustering." *Autonomous Agents and Multi-Agent Systems* 25, no. 3 (2012): 425-446.

Choi, Jeehye, and Kwang-Il Goh. "Majority-vote dynamics on multiplex networks with two layers." *New Journal of Physics* 21.3 (2019): 035005.

Chrol-Cannon, Joseph, and Yaochu Jin. "Computational modeling of neural plasticity for self-organization of neural networks." *Biosystems* 125 (2014): 43-54.

Clauset, Aaron, Mark EJ Newman, and Cristopher Moore. "Finding community structure in very large networks." *Physical Review E* 70.6 (2004): 066111.

Connors, Jacob, Scott Graham, and Logan Mailloux. "Cyber synthetic modeling for vehicle-to-vehicle applications." *In Proceedings of 2018 International Conference on Cyber Warfare and Security*, pp. 594-XI. Academic Conferences International Limited, 2018.

Dayan, Peter, and Nathaniel D. Daw. "Decision theory, reinforcement learning, and the brain." *Cognitive, Affective, & Behavioral Neuroscience* 8, no. 4 (2008): 429-453.

de Arruda, Guilherme Ferraz, Emanuele Cozzo, Tiago P. Peixoto, Francisco A. Rodrigues, and Yamir Moreno. "Disease localization in multilayer networks." *Physical Review X* 7, no. 1 (2017): 011014.

Deci, Edward L., Richard Koestner, and Richard M. Ryan. "Extrinsic rewards and intrinsic motivation in education: Reconsidered once again." *Review of Educational Research* 71.1 (2001): 1-27.

Deffuant, Guillaume, David Neau, Frederic Amblard, and Gérard Weisbuch. "Mixing beliefs among interacting agents." *Advances in Complex Systems* 3, no. 01n04 (2000): 87-98.

Dittrich, Peter, Thomas Kron, and Wolfgang Banzhaf. "On the scalability of social order-modeling the problem of double and multi-contingency following Luhmann." *Journal of Artificial Societies and Social Simulation* 6.1 (2002).

Epstein, Joshua M., and Robert Axtell. *Growing artificial societies: Social science from the bottom up*. Brookings Institution Press, 1996.

Ferrante, Eliseo, Ali Emre Turgut, Edgar Duéñez-Guzmán, Marco Dorigo, and Tom Wenseleers. "Evolution of self-organized task specialization in robot swarms." *PLoS Computational Biology* 11, no. 8 (2015): e1004273.

Filatova, Tatiana, Peter H. Verburg, Dawn Cassandra Parker, and Carol Ann Stannard. "Spatial agent-based models for socio-ecological systems: Challenges and prospects." *Environmental Modelling & Software* 45 (2013): 1-7.

Forsyth, Donelson R. *Group dynamics*. Cengage Learning, 2018.

Fossett, Mark, and Warren Waren. "Overlooked implications of ethnic preferences for residential segregation in agent-based models." *Urban Studies* 42.11 (2005): 1893-1917.

Friedkin, Noah E., and Eugene C. Johnsen. "Social influence and opinions." *Journal of Mathematical Sociology* 15.3-4 (1990): 193-206.

Gallegati, Mauro, and Alan Kirman. "Reconstructing economics: Agent-based models and complexity." *Complexity Economics* 1.1 (2012): 5-31.

Georgeff, Michael, Barney Pell, Martha Pollack, Milind Tambe, and Michael Wooldridge. "The belief-desire-intention model of agency." *In Proceedings of the 1998 International Workshop on Agent Theories, Architectures, and Languages*, pp. 1-10. Springer, Berlin, Heidelberg, 1998.

Gershenson, Carlos. "Computing networks: A general framework to contrast neural and swarm cognitions." *Paladyn, Journal of Behavioral Robotics* 1, no. 2 (2010): 147-153.

Gershenson, Carlos. "Facing complexity: Prediction vs. adaptation." In *Complexity Perspectives on Language, Communication and Society*, pp. 3-14. Springer, Berlin, Heidelberg, 2013.

Gigerenzer, Gerd, and Reinhard Selten, eds. *Bounded rationality: The adaptive toolbox*. MIT Press, 2002.

Golman, Russell, and George Loewenstein. "Information gaps: A theory of preferences regarding the presence and absence of information." *Decision* 5.3 (2018): 143-164.

Gómez-Serrano, Javier, Carl Graham, and Jean-Yves Le Boudec. "The bounded confidence model of opinion dynamics." *Mathematical Models and Methods in Applied Sciences* 22.02 (2012): 1150007.

Granovetter, Mark S. "The strength of weak ties." *American Journal of Sociology* 78.6 (1973): 1360-1380.

Ha, David, and Yujin Tang. "Collective intelligence for deep learning: A survey of recent developments." *arXiv preprint* arXiv:2111.14377 (2021).

Harman, Mark. "The current state and future of search-based software engineering." In *Future of Software Engineering (FOSE'07)*, pp. 342-357. IEEE, 2007.

Hegselmann, Rainer, and Ulrich Krause. "Opinion dynamics and bounded confidence models, analysis, and simulation." *Journal of Artificial Societies and Social Simulation* 5.3 (2002).

Hosseini, Fatemeh Beigom, Saeed Ghorbani, and Reza Rezaeeshirazi. "Effects of perceived autonomy support in the physical education on basic psychological needs satisfaction, intrinsic motivation and intention to perform physical activity in high school students." *International Journal of School Health* 7.4 (2020): 39-46.

Hosseinmardi, Homa, Amir Ghasemian, Aaron Clauset, Markus Mobius, David M. Rothschild, and Duncan J. Watts. "Examining the consumption of radical content on YouTube." *Proceedings of the National Academy of Sciences* 118, no. 32 (2021): 166-177.

Hu, Junling, and Michael P. Wellman. "Nash Q-learning for general-sum stochastic games." *Journal of Machine Learning Research* 4. Nov (2003): 1039-1069.

Huang, Wenlong, Igor Mordatch, and Deepak Pathak. "One policy to control them all: Shared modular policies for agent-agnostic control." In *International Conference on Machine Learning*, pp. 4455-4464. PMLR, 2020.

Iniguez, Gerardo, János Kertész, Kimmo K. Kaski, and R. A. Barrio. "Phase change in an opinion-dynamics model with separation of time scales." *Physical Review E* 83, no. 1 (2011): 016111.

Jablonka, Eva, and Marion J. Lamb. "Soft inheritance: challenging the modern synthesis." *Genetics and Molecular Biology* 31, no. 2 (2008): 389-395.

Jian, Muwei, Quan Zhou, Chaoran Cui, Xiushan Nie, Hanjiang Luo, Jianli Zhao, and Yilong Yin. "Assessment of feature fusion strategies in visual attention mechanism for saliency detection." *Pattern Recognition Letters* 127 (2019): 37-47.

Kapetanakis, Spiros, and Daniel Kudenko. "Improving on the reinforcement learning of coordination in cooperative multi-agent systems." *In Proceedings of the Second AISB Symposium on Adaptive Agents and Multi-agent Systems* II (2002): 119–131.

Kim, Kyung-Joong, and Sung-Bae Cho. "A comprehensive overview of the applications of artificial life." *Artificial Life* 12, no. 1 (2006): 153-182.

Kleinberg, Jon. "Small-world phenomena and the dynamics of information." *Advances in Neural Information Processing Systems* 14 (2001).

Langton, Christopher G., ed. *Artificial life: An overview*. The MIT Press, 1997.

Lazaridou, Angeliki, and Marco Baroni. "Emergent multi-agent communication in the deep learning era." *arXiv preprint arXiv:2006.02419* (2020).

Le Goues, Claire, Westley Weimer, and Stephanie Forrest. "Representations and operators for improving evolutionary software repair." In *Proceedings of the 14th Annual Conference on Genetic and Evolutionary Computation*, pp. 959-966. 2012.

Lim, Dongwon, Hwansoo Lee, Hang-Jung Zo, and Andrew Ciganek. "Opinion formation in the digital divide." *Journal of Artificial Societies and Social Simulation* 17, no. 1 (2014).

Littman, Michael L. "Friend-or-foe Q-learning in general-sum games." In *Proceedings of the 2nd International Conference on Machine Learning*, vol. 1, pp. 322-328. 2001.

Littman, Michael L. "Markov games as a framework for multi-agent reinforcement learning." *Machine Learning Proceedings*. Morgan Kaufmann, 1994. 157-163.

Littman, Michael L., and David H. Ackley. "Adaptation in constant utility non-stationary environments." In *Proceedings of the 6th International Conference on Genetic Algorithms*, pp. 136-142. 1991.

Macal, Charles M., and Michael J. North. "Tutorial on agent-based modeling and simulation." In *Proceedings of the Winter Simulation Conference*, pp. 14-pp. IEEE, 2005.

Mastroeni, Loretta, Pierluigi Vellucci, and Maurizio Naldi. "Agent-based models for opinion formation: A bibliographic survey." *IEEE Access* 7 (2019): 58836-58848.

Mathias, Jean-Denis, Sylvie Huet, and Guillaume Deffuant. "An energy-like indicator to assess opinion resilience." *Physica A: Statistical Mechanics and its Applications* 473 (2017): 501-510.

Matsuo, Yutaka, Junichiro Mori, Masahiro Hamasaki, Takuichi Nishimura, Hideaki Takeda, Koiti Hasida, and Mitsuru Ishizuka. "Polyphonet: An advanced social network extraction system from the web." *Journal of Web Semantics* 5, no. 4 (2007): 262-278.

McPherson, Miller, Lynn Smith-Lovin, and James M. Cook. "Birds of a feather: Homophily in social networks." *Annual Review of Sociology* 27.1 (2001): 415-444.

Meehl, Paul E. "Factors and taxa, traits and types, differences of degree and differences in kind." *Journal of Personality* 60.1 (1992): 117-174.

Mesoudi, Alex. "Migration, acculturation, and the maintenance of between-group cultural variation." *PLoS One* 13, no. 10 (2018): e0205573.

Monderer, Dov, and Lloyd S. Shapley. "Fictitious play property for games with identical interests." *Journal of Economic Theory* 68.1 (1996): 258-265.

Mordvintsev, Alexander, Ettore Randazzo, Eyvind Niklasson, and Michael Levin. "Growing neural cellular automata." *Distill* 5, no. 2 (2020): e23.

Murphy, Kilian J., Simone Ciuti, and Adam Kane. "An introduction to agent-based models as an accessible surrogate to field-based research and teaching." *Ecology and Evolution* 10, no. 22 (2020): 12482-12498.

Nasution, Mahyuddin KM, and Shahrul Azman Noah. "Information retrieval model: A social network extraction perspective." In *Proceedings of the International Conference on Information Retrieval & Knowledge Management*, pp. 322-326. IEEE, 2012.

Neal, Zachary P., and Jennifer W. Neal. "The (in)compatibility of diversity and sense of community." *American Journal of Community Psychology* 53.1-2(2016):1-12.

Nichele, Stefano, Mathias Berild Ose, Sebastian Risi, and Gunnar Tufte. "CA-NEAT: Evolved compositional pattern producing networks for cellular automata morphogenesis and replication." *IEEE Transactions on Cognitive and Developmental Systems* 10, no. 3 (2017): 687-700.

Noorazar, Hossein. "Recent advances in opinion propagation dynamics: A 2020 survey." *The European Physical Journal Plus* 135.6 (2020): 1-20.

Nordhausen, Klaus. *Statistical analysis of network data with R.* Springer Nature, 2015.

Nowak, Martin, and Karl Sigmund. "A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game." *Nature* 364.6432 (1993): 56-58.

Pan, Qiuhui, Xuesong Liu, Honglin Bao, Yu Su, and Mingfeng He. "Evolution of cooperation through adaptive interaction in a spatial prisoner's dilemma game." *Physica A: Statistical Mechanics and its Applications* 492 (2018): 571-581.

Pattee, Howard H., and Hiroki Sayama. "Evolved open-endedness, not open-ended evolution." *Artificial Life* 25, no. 1 (2019): 4-8.

Pimm, Stuart L., Gareth J. Russell, John L. Gittleman, and Thomas M. Brooks. "The future of biodiversity." *Science* 269, no. 5222 (1995): 347-350.

Pitonakova, Lenka, Richard Crowder, and Seth Bullock. "Understanding the role of recruitment in collective robot foraging." In *ALIFE 14: The Fourteenth International Conference on the Synthesis and Simulation of Living Systems*, pp. 264-271. MIT Press, 2014.

Plous, Scott. *The psychology of judgment and decision making.* Mcgraw-Hill Book Company, 1993.

Polikar, Robi. "Ensemble based systems in decision-making." *IEEE Circuits and Systems Magazine* 6.3 (2006): 21-45.

Rand, William. "Machine learning meets agent-based modeling: When not to go to a bar." In *Proceedings of the Conference on Social Agents: Results and Prospects.* 2006.

Rao, Anand S., and Michael P. Georgeff. "BDI agents: From theory to practice." In *Proceedings of the First International Conference on Multiagent Systems -AAAI 1995*, vol. 95, pp. 312-319. 1995.

Reynolds, Craig W. "Flocks, herds and schools: A distributed behavioral model." *In Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*, pp. 25-34. 1987.

Reynolds, Robert G. "An introduction to cultural algorithms." In *Proceedings of the Third Annual Conference on Evolutionary Programming*, Vol. 24, pp. 131-139. River Edge: World Scientific, 1994.

Sandler, Mark, Andrey Zhmoginov, Liangcheng Luo, Alexander Mordvintsev, and Ettore Randazzo. "Image segmentation via cellular automata." *arXiv preprint* arXiv:2008.04965 (2020).

Savarimuthu, Bastin Tony Roy, Rexy Arulanandam, and Maryam Purvis. "Aspects of active norm learning and the effect of lying on norm emergence in agent societies." *In Proceedings of the 13rd International Conference on Principles and Practice of Multi-Agent Systems*, pp. 36-50. Springer, Berlin, Heidelberg, 2011.

Schelling, Thomas C. "Dynamic models of segregation." *Journal of Mathematical Sociology* 1.2 (1971): 143-186.

Semonsen, Justin, Christopher Griffin, Anna Squicciarini, and Sarah Rajtmajer. "Opinion dynamics in the presence of increasing agreement pressure." *IEEE Transactions on Cybernetics* 49, no. 4 (2018): 1270-1278.

Shore, Jesse, J. Baek, and C. Dellarocas. "Network structure and patterns of information diversity on Twitter." *MIS Quarterly* 42.3(2018): 849-872.

Staddon, John Eric Rayner. *Adaptive behavior and learning.* Cambridge University Press, 2016.

Szabó, György, and Csaba Tőke. "Evolutionary prisoner's dilemma game on a square lattice." *Physical Review E* 58.1 (1998): 69-73.

Sznajd-Weron, Katarzyna, and Jozef Sznajd. "Opinion evolution in closed community." *International Journal of Modern Physics C* 11.06 (2000): 1157-1165.

Tang, Yujin, and David Ha. "The sensory neuron as a transformer: Permutation-invariant neural networks for reinforcement learning." *Advances in Neural Information Processing Systems* 34 (2021).

Taylor, Matthew E., Shimon Whiteson, and Peter Stone. "Comparing evolutionary and temporal difference methods in a reinforcement learning domain." *In Proceedings of the 8th Annual Conference on Genetic and Evolutionary Computation*, pp. 1321-1328. 2006.

Taylor, Tim, Mark Bedau, Alastair Channon, David Ackley, Wolfgang Banzhaf, Guillaume Beslon, Emily Dolson, Tom Froese, Simon Hickinbotham, Takashi Ikegami, Barry McMullin, Norman Packard, Steen Rasmussen, Nathaniel Virgo, Eran Agmon, Edward Clark, Simon McGregor, Charles Ofria, Glen Ropella, Lee Spector, Kenneth O. Stanley, Adam Stanton, Christopher Timperle, Anya Vostinar, and Michael Wiser. "Open-ended evolution: Perspectives from the OEE workshop in York." *Artificial Life* 22, no. 3 (2016): 408-423.

Thrun, Michael C., and Alfred Ultsch. "Swarm intelligence for self-organized clustering." *Artificial Intelligence* 290 (2021): 103237.

Villani, Marco, Stefano Benedettini, Andrea Roli, David Lane, Irene Poli, and Roberto Serra. "Identifying emergent dynamical structures in network models." In *Recent Advances of Neural Network Models and Applications*, pp. 3-13. Springer, Cham, 2014.

Wang, Hanfu, Weidong Chen, and Jingchuan Wang. "Heterogeneous multi-agent routing strategy for robot-and-picker-to-good order fulfillment system." In *Proceedings of International Conference on Intelligent Autonomous Systems*, pp. 237-249. Springer, Cham, 2018.

Wang, Huanjing, and Lihui Shang. "Opinion dynamics in networks with common-neighbors-based connections." *Physica A: Statistical Mechanics and its Applications* 421 (2015): 180-186.

Wang, Yao, Jie Zhang, and Julita Vassileva. "Super-agent-based reputation management with a practical reward mechanism in decentralized systems." *In Proceedings of the 2010 IFIP WG 11.4 International Conference on Open Research Problems in Network Security* Vol. 11. 2010.

Watts, Duncan J., and Steven H. Strogatz. "Collective dynamics of 'small-world' networks." *Nature* 393.6684 (1998): 440-442.

Weisbuch, Gérard, Guillaume Deffuant, Frederic Amblard, and J-P. Nadal. "Interacting agents and continuous opinions dynamics." In *Heterogenous Agents, Interactions and Economic Performance*, pp. 225-242. Springer, Berlin, Heidelberg, 2003.

Wulff, N., and J. A. Hertz. "Learning cellular automaton dynamics with neural networks." *Advances in Neural Information Processing Systems* 5 (1992).

Young, H. Peyton. "The economics of convention." *Journal of Economic Perspectives* 10.2 (1996): 105-122.

Yu, Chao, Zhen Wang, Hongtao Lv, Honglin Bao, and Yapeng Li. "Collective learning and information diffusion for efficient emergence of social norms." In *Multi-agent and Complex Systems*, pp. 193-210. Springer, Singapore, 2017.

Yuan, Ruiping, Juntao Li, Xiaolin Wang, and Liyan He. "Multirobot task sllocation in e-Commerce robotic mobile fulfillment systems." *Mathematical Problems in Engineering* 2021 (2021).

Zhao, Yiyi, and Gang Kou. "Bounded confidence-based opinion formation for opinion leaders and opinion followers on social networks." *Studies in Informatics and Control* 23.2 (2014): 153-162.

Zhou, Ryan, Christian Muise, and Ting Hu. "Permutation-invariant representation of neural networks with neuron embeddings." In *Proceedings of the European Conference on Genetic Programming* (Part of EvoStar), pp. 294-308. Springer, Cham, 2022.