2020 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)

Social Capital Games as A Framework for Social Structural Pattern Emergence

Yang Chen and Jiamou Liu
School of Computer Science, The University of Auckland
Auckland, New Zealand
{yang.chen, jiamou.liu}@auckland.ac.nz

***Abstract*—Prominent structural patterns such as small-world and core-periphery structures amount to some of the most important emergent characteristics of a social network. Yet little work is done to interpret these emergent phenomena in a unified way. Towards a unified interpretation framework, we connect the establishment of social patterns with *social capital*. Social capital captures the benefits that an individual gains from its social surrounding. We argue that individuals' desire to gaining higher social capital may give rise to important network properties. To validate this claim, we propose *social capital game* that mathematically conceptualizes bonding and bridging social capital. This framework allows us to regard individuals in a social network as learning agents who gain social capital through iteratively building interpersonal ties. The link-building decisions of these agents are guided by a multi-agent reinforcement learning (MARL) algorithm which improves agents' capability through repeated game plays. We conduct a series of experiments which demonstrate (1) the collective behaviors of the agents give rise to salient social patterns, and (2) by varying agents' preferences to different forms of social capital, different types of social patterns emerge. In particular, bonding social capital plays a pivotal role in the formation of a community structure in the network while bridging social capital is instrumental to the emergence of core-periphery structure. Our work sheds light on the formation of complex network phenomena.**

***Index Terms*—Multi-agent reinforcement learning; Network formation; Social structures; Social capital**

I. Introduction

One of the key themes of social network analysis revolves around discovering salient structural patterns within a social network. For example, co-authorship networks often demonstrate *community structure*, where scholars form rings of collaboration within disciplinary boundaries [1]. Another example is that online social platforms often exhibit *small-world property*, which means that any two users are often connected through a small number of intermediate acquaintances [2]. A third example is that in human organizations one usually observes *core-periphery structure*, where a minority of nodes form a core that sits in the center of the network while others stay at the outskirts [3]. Capturing these properties not only give us important insights on the network, but also enable the design of random graph models that could simulate real-world networks through computational means. As social networks are organically grown, these structural properties also arise as an emergent, rather than a prescriptive, phenomenon. Uncovering the mechanism behind these emergent phenomena would deepen our understanding of how social networks form, function and evolve. However, no theory yet achieves a unified interpretation of the emergence of social structural patterns. Our knowledge regarding why such structural patterns emerge, and more specifically, why some networks exhibit some of the patterns while others exhibit other patterns, remain limited. A question thus naturally arises: *What would be a unified mechanism behind the emergence of social structural patterns?*

To answer this question, one needs to explore the driving force behind the formation of social structures. Evidence rooted in sociology has indicated that social structures tightly correlate with *social capital* [4]. The concept of social capital amounts to the benefits that social ties confer to individuals. These benefits can be tangible, *e.g.*, economic and human resources, or impalpable, *e.g.*, social support, information control, and social influence [5]. In social psychology, the *reward theory of attraction* claims that people tend to interact with those whose behaviors are rewarding to themselves or those who are associated with rewarding events [6]. These studies support a theory that links the formation of the social network structure with social capital, where social capital embodies the rewards gained by individuals through social networking.

To develop this theory, two challenges need to be resolved. We first need a rigorous definition of social capital. A well-known dichotomy has divided discussions on social capital into two categories: *bonding capital* and *bridging capital* [7]. The former depicts the aggregate welfare that an individual draws from its closed social circle in the form of, *e.g.*, trust and social support [8], while the latter captures the individual's capacity to acquire opportunities and information via open links and determines, *e.g.*, status and power [9]. An individual's reward in building social ties would be a combination of these two forms of social capital.

Since social capital arises from the establishment of social ties [10], it is then important to develop a model for the social instruments that enable individuals to build ties [11], [12]. Most existing model are *one-shot models* in the sense that they build ties all at once, providing limited insight into the network dynamics [13], [14]. In contrast, several recent work aim to capture the dynamics of networks [15], [16]. Individuals' access to information is also an important factor to consider. Most existing models grant individuals with complete information about the network, while recently some models stipulate that agents only have local information to reflect their restricted sights in a large and decentralized network [17], [18]. Thus, to simulate real-world societies, a combination of dynamic models and incomplete information is desirable within the scope of this work.

This work aims to develop an agent-based framework for the formation of social network from a social capital perspective.

IEEE/ACM ASONAM 2020, December 7-10, 2020
978-1-7281-1056-1/20/$31.00 © 2020 IEEE

In our framework, social actors are agents who are able to establish social links with others. Their decisions to establish links are driven by a desire to grow their social capital. The framework specifies how their decisions evolve and improve through repeated interactions with others. These improvements are reflected at the population-level as the emergence of social structural patterns. The contribution of our work are manyfold: (1) Our framework formalizes the notions of social capital using metrics in social network analysis. (2) We propose a game-based model, namely *social capital games* (SCG), that capture the dynamics of social networks and behavioral traits of agents. (3) To model the agents' behavioral changes when playing SCG, we adopt multi-agent reinforcement learning (MARL). Agents are treated as independent learners in a shared environments equipped only with local information about the network. To the best of our knowledge, we are the first to bring the multi-agent learning paradigm to the field of network formation. (4) We perform a series of experiments and observe the resulting networks formed as a result of the learning process of agents in SCG. In particular, we observe the emergence of community, small-world, and core-periphery structures as the agents acquire higher social capital through training under different configurations. Moreover, by varying agents' preferences to different forms of social capital, different types of social patterns emerge. In particular, bonding social capital plays a pivotal role in the formation of a community structure in the network while bridging social capital is instrumental to the emergence of core-periphery structure. With these analysis, our work sheds light on the formation of complex network phenomena.

Related Work. Pioneering works in sociology advanced the research on social capital. Coleman laid the foundation for the research on social capital [5], [19]. Bourdieu proposed that homophily is the source of bonding capital [8]. Granovetter, Putnam, and Burt stated that weak ties are the source of bridging capital [9], [20]. Inspired by their work, we formalize two types of social capital. Recently, a number of works have uncovered vital roles that social capital plays in a wide range of applications, such as resource management [21], disaster survival analysis [22] and unemployment analysis [23].

Network formation models aim to discover the natural emergence of social structures. Traditional approaches to network formation fall into three main paradigms: models based on *random events*, models based on *strategic decisions*, and *empirical models* distilled via mining data of real-world social networks. Models based on random events are generative models with ad-hoc designs that mimic real-world networks [24]–[26]. Agents wherein are manipulated in probabilistic manners, *e.g.*, preferential attachment [27], to produce intended degree distributions. Though such models can generate networks with desired structural properties, they pay limited attention to agents' behavioral acquisitions and are thus limited to explain the emergence of social structures over time. Models based on strategic decisions provide an explanation for how social structures emerge as equilibria of *network formation games* [13], [17]. However, results of these models are limited to *stability* and *efficiency* and neglect dynamics of networks. Empirical models can offer evidence for community detection [28], link prediction and recommendation [29] by mining real-world data, but are powerless to show how networks form in real life. Our work differs from these models as the following properties hold simultaneously in our proposed framework: (1) the SCG captures dynamics of societies; (2) MARL captures behavioral traits and acquisitions of agents; (3) our model reproduces classical types of social structures rather than being restricted to a specific class.

Multi-agent reinforcement learning has been adopted in social simulation under game-theoretic setups. For example, Sen and Airiau introduce reinforcement learning to simulate the emergence of social convention [30]. In their framework, agents learn strategies through iteratively interacting with each others. Many work follows this idea to study norm emergence in networked multi-agent systems [31]–[33]. More recently, deep learning-based MARL is used to analyze the emergence of cooperation [34]–[36]. So far, there has not been any work that applies MARL to the field of network formation. In this work, instead of studying the emergence of behavioral norms, we focus on the emergence of "structural norms", *i.e.*, social structures. We use graph neural networks to learn latent features of the network, since the social surrounding of an agent is structured data. Likewise, the reinforcement learning is adopted to learn a policy.

II. Background

Following standard conventions, a social network is viewed as a graph. Fix $N = \{1, 2, \dots, n\}$ as a set of social actors (i.e., agents), each representing a node in a graph. Edges between nodes represent social ties. We define the complete graph g^N as the set of all subsets of N of size 2. Hence $\{g \mid g \subseteq g^N\}$ denotes the set of all possible graphs on N. For any two distinct nodes $i, j \in N$, $\{i, j\} \in g$ indicates that an undirected edge exists between i and j in graph g. For simplicity, we write $ij \in g$. Now we can formalize the creation of links, for any $g' \subseteq g^N$, let $g + g'$ denote the integrated graph obtained via adding each link $ij \in g'$ into g, *i.e.*, $g + g' = g \cup g'$.

A *path* in g is a sequence of edges $i_1i_2, i_2i_3, \dots, i_Mi_{M+1}$ where $i_mi_{m+1} \in g$ for all $1 \leq m \leq M$. The *distance* $\mathsf{dist}(i, j)$ between i and j is the length of a shortest path between these two nodes. We focus on connected graphs and thus $\mathsf{dist}(i, j) < \infty$ for any $i, j \in N$. The *d-hop neighbor set* of i is $\mathcal{N}_d(i) \coloneqq \{j \in N \mid \mathsf{dist}(i, j) = d\}$. We denote $\mathcal{N}_d[i] \coloneqq \{j \in N \mid \mathsf{dist}(i, j) \leq d\}$ all i's neighbors within distance d. As stated earlier, an agent is often restricted to access complete information of the network in real life. To capture this fact, we assume an agent i's vision is restricted to $\mathcal{N}_2[i]$, *i.e.*, friends and friends of friends. Formally, we employ the notion of *2-level ego network*, which represents the social surrounding that an agent perceives and maintains.

Definition 1 (2-level Ego Network). *The* 2-level ego network *of node $i \in N$ is the subgraph $o_i \subseteq g$ induced by i, i's 1-hop and 2-hop neighbors,* i.e., $o_i \coloneqq \{jk \mid j, k \in \mathcal{N}_2[i]\}$.

Social structural patterns are topological properties in the context of graphs. We briefly present three classical and widely-studied structural patterns:

- **Community:** The community structure has been known as a prevalent property in social networks. Namely, nodes in the network can be partitioned into clusters, with a high density of edges within each cluster but low density of edges between these clusters [1], [37]. Formally, a community structure of a graph $g \subseteq g^N$ refers to a partition of N, $\mathcal{C} = \{C_1, C_2, \ldots, C_k\}$, such that each component induces a connected subgraph. A standard method to quantify the significance of community structure is *modularity*, which measures how many edges lie within clusters relative to the expected number of such edges:

$$MOD(\mathcal{C}) = \frac{1}{2|g|} \sum_{i,j \in N} \left(A_{i,j} - \frac{d_i d_j}{2|g|} \right) \delta(c_i, c_j), \quad (1)$$

where $A_{i,j} = 1$ if $ij \in g$ and $A_{i,j} = 0$ otherwise; d_i is the degree of node i; c_i denotes the community of node i; $\delta(c_i, c_j) = 1$ if $c_i = c_j$ and $\delta(c_i, c_j) = 0$ otherwise.

• **Small-World:** The small-world property is one of most important attributes of a complex network, *e.g.*, Internet and gene networks [2]. Loosely speaking, the property asserts that every pair of nodes are within a small distance apart. Small-world is reflected by a high *clustering coefficient* and a low *average shortest path length*. The clustering coefficient of a node i, $\mathrm{c}(i)$, measures the probability of two randomly chosen friends of i are also friends. The clustering coefficient of a graph g, $\mathrm{C_C}(g)$, averages clustering coefficients of all agents:

$$C_C(g) = \sum_{i \in N} \frac{\mathrm{c}(i)}{|N|} = \sum_{i \in N} \frac{2 \cdot |\{jk \in g \mid ij, ik \in g\}|}{d_i(d_i - 1)|N|}. \quad (2)$$

The average shortest path length of a graph g, $L(g)$, measures the average distance between all pairs of nodes, *i.e.*,

$$L(g) = \frac{1}{|N|(|N| - 1)} \sum\nolimits_{i \neq j \in N} \mathsf{dist}(i, j). \quad (3)$$

• **Core-Periphery:** Core-periphery (C-P) structure is a common property of real-world social networks and networks of the economy [3], [15], where a densely connected core is in the center, and other nodes are located at edges. Broadly speaking, a scale-free graph can be viewed as an instance of C-P structure, where the degree distribution follows a power law. A popular measurement is *C-P coefficient*, which measures the prominence of C-P structure of the given graph relative to the expected prominence [38]. C-P coefficient is defined based on the notion of *extended closeness centrality* for a subset of nodes:

$$\delta_g(U) = \sum\nolimits_{i \in U, j \neq i \in N} \mathsf{dist}(i, j)/|U|, \quad (4)$$

where $U \subseteq N$. The C-P coefficient is formally defined as

$$C_{CP}(g) = \frac{\delta_g(U_{k\text{-core}}(g))}{\delta_g(N)} - \mathbb{E}\left[\frac{\delta_{g'}(U_{k\text{-core}}(g'))}{\delta_{g'}(N)}\right], \quad (5)$$

where $U_{k\text{-core}}(g)$ is the set of nodes of the maximum subgraph of g with minimum degree k and maximal δ_g value, g' is a graph with the same degree sequence as g.

III. Towards a Formalization of Social Capital

Rooted in sociology, the concept of social capital aims to capture the benefits attained by individuals via social interactions. Such benefits can emerge in the form of social support, companionship, solidarity, influence, and control over information, which are closely related to network structural properties. As a result, social capital should be measured based on structural properties. However, the structure itself does not define social capital. Instead, social capital arises as a function of social interactions and information conveyed through social relations [19], [39]. Social relations have long been classified based on their functions: While strong ties link homogeneous and like-minded individuals, weak ties bridge diverse and weakly connected groups [20]. Analogously, social capital also consists of two types: *bonding capital* refers to benefits an individual draws from its closed neighborhood, in the form of, *e.g.*, trust and support, which are brought by strong ties; while *bridging capital* is an embodiment of benefits of accessibility to information and control over information flow, which are largely functions of weak ties. However, so far no consensus has been reached over the formal definitions of social capital. In this paper, we adopt two metrics to measure these two notions in the context of social networks.

Bonding capital, as it is often expressed as trust and companionship, measures the extent to which two nodes bind with each other [8]. This can be aptly captured through a measure of "social proximity". In other words, a node gains more bonding capital as it gets closer to others in its neighborhood. To this end, we adopt *personalized PageRank index*, which evaluates structural proximity between nodes through predicting the likelihood of edges between any pairs of nodes [40], [41]. The metric is adapted from PageRank: It takes as input a starting node i, and assigns a score to every node j that captures the likelihood of a random walk from i to reach j. Fix a restart probability $\beta \in (0, 1)$, random walk starts from the node i; stops moving at each node with the probability of β and restarts from the node i; or continues to walk with the probability of $1 - \beta$ by randomly selecting a node from the neighbors of the current node. The probability that each node is accessed converges in finite rounds of walking. Each entry of the personalized PageRank vector $\boldsymbol{pr}$ records the probability that the corresponding node is accessed. More formally, let $\boldsymbol{a}_j$ be the column vector in the adjacency matrix of g corresponding to node j. Denote by pr_j, the link prediction score between i and j is:

$$\mathsf{pr}_j = \beta r_j + (1 - \beta)(\boldsymbol{pr} \cdot \boldsymbol{a}_j / |\mathcal{N}_1(j)|), \quad (6)$$

where $\beta \in (0, 1)$ is the restart probability, $r_j = 1$ if $j = i$ and $r_j = 0$ otherwise, and each entry of the personalized PageRank vector $\boldsymbol{pr}$ records the probability that the corresponding node is accessed. Intuitively, assume that i holds certain amount of "goodwill" which is randomly shared with i's neighbors, and whoever that obtains such goodwill can continue to pass goodwill to their neighbors or return them to the node i, in the same manner as a random walk. Bonding capital can be viewed as the amount of goodwill eventually received by i.

Definition 2 (Bonding Capital). *Given a graph $g \subseteq g^N$ and a node $i \in N$, the* bonding capital *of i is defined by summing personalized PageRank indices between i and i's neighbors, namely,* $\mathsf{bo}_i := \sum_{j \in \mathcal{N}_1(i)} \mathsf{pr}_j$.

Occupying a central position to act as a gateway for information exchange brings an individual bridging capital [9]. Betweenness centrality is used to evaluate bridging capital in [26], [42]. The betweenness centrality of a node measures the number of the shortest paths between each pair of other nodes that pass through it, and thus reflects the agent's ability to broker interactions between different groups of agents.

Definition 3 (**Bridging Capital**). *Let $g \subseteq g^N$ be a connected graph. The* bridging capital *of $i \in N$, is defined as i's betweenness centrality:* $\mathsf{br}_i := \sum_{j \neq i \neq k \in N} \sigma_{jk}(i)/\sigma_{jk}$, *where σ_{jk} is the number of shortest paths between nodes j and k, and $\sigma_{jk}(i)$ is the number of shortest paths passing i.*

An individual may have different preferences to two types of capital. To reflect this fact, we employee a *preference weight* $w \in [0, 1]$ to define the *mixed capital*.

Definition 4 (**Mixed Capital**). *Let $g \subseteq g^N$ be a graph. For a node $i \in N$ and a preference weight $w \in [0, 1]$, the* mixed capital *is defined as* $\mathsf{mix}_{i,w} := w\mathsf{bo}_i + (1 - w)\mathsf{br}_i$.

IV. The Model: Social Capital Games

We construct a game-based model, *social capital game*, which takes social capital as utilities. Individuals wherein are regarded as social-capital-driven and self-interested agents. Before formally define social capital games, we first address following auxiliary notions.

Dynamic networks. Let $N = \{1, 2, \ldots, n\}$ be a finite set of agents. An (ℓ-length) finite *dynamic network* is a sequence of graphs $G = g^0, g^1, \ldots, g^\ell$ that evolves in finite discrete time steps $0, 1, \ldots, \ell$, where ℓ is the *termination step*. Each $g^t \subseteq g^N$ is called a *network instance* at step t. Throughout, we use superscript t and subscript i to denote the corresponding notation derived from time step t and agent i, respectively. At each step $t < \ell$, each agent $i \in N$ builds a link to another agent a_i^t from i's 2-hop neighbor set $\mathcal{N}_2^t(i)$, as i's observation is restricted to the 2-level ego network o_i^t. All agents make decisions simultaneously, resulting in the next network instance at step $t+1$. Formally, $\forall 0 \leq t < \ell : g^{t+1} = g^t + \{ia_i^t\}_{i \in N}$, where $a_i^t \in \mathcal{N}_2^t(i)$.

Remark. Two caveats exist here: Firstly, we have to clarify the evolution of real-world networks. In principle, any addition/deletion of node/edge may happen. To simplify the model, however, we only allow additive changes, that is, the only allowable change is the addition of edges. Secondly, the creation of edges is unilateral in our model, though a large literature on network formation games assume an edge must be established reciprocally, *i.e.*, an edge emerges and persists only if both involved nodes gain payoff from it [13]. However, unilateral edges also cover a large class of real-world networks such as online social networks and academic citation networks. On the other hand, many classical network formation models also adopt additive changes and unilateral edges [2], [43].

Utilities. We measure the *immediate utility* of an agent as the increment of the mixed capital between two consecutive time steps. Formally, the utility of agent i received at step $t+1$ after linking to a_i^t is defined as: $u_i^{t+1} := \mathsf{mix}_{i,w_i}^{t+1} - \mathsf{mix}_{i,w_i}^t$. As such, the *cumulative utility* of agent i at step t sums over rewards received at all elapsed time steps: $U_i^t := \sum_{i=1}^t u_i^t = \mathsf{mix}_{i,w_i}^t - \mathsf{mix}_{i,w_i}^0$. A *social capital game (SCG)* incorporates dynamic networks and rewards, which is formally defined below.

Definition 5 (**Social Capital Games**). *A* social capital game *(SCG) is a tuple (N, W, g^0, ℓ), where*

- $N = \{1, 2, \ldots, n\}$ *is a finite set of agents;*
- $W = (w_1, w_2, \ldots, w_{|N|})$ *is a* preference vector*, in which each entry w_i records the preference weight of agent i;*
- $g^0 \subseteq g^N$ *is the* initial network*;*
- $\ell \in \mathbb{N}^+$ *is the* termination step.

Conceptually, one can view an SCG as a multi-stage game played amongst agents over graphs with imperfect information. The goal of an agent i is to select actions in a way that maximizes the cumulative utility U_i^ℓ during the whole course of a game. At each stage, all agents build links simultaneously, resulting in that the game enters to the next stage. The behavioral trait of an agent is thus reflected by the *policy* to build links. Our attention here is that how an agent i take a strategy to build a relationship under an observation o_i and preference weights of within agents.

Definition 6 (**Policy**). *Let (N, W, g^0, ℓ) be a social capital game. A* policy *of an agent $i \in N$ is a function π_i defined on all possible 2-level ego networks*[1] *of i such that $\pi_i(o_i) = a \in \mathcal{N}_2(i)$ for any $o_i \subseteq g^N$.*

We put the process of learning a policy under the MARL framework. A refinement of policies yields an updated underlying dynamic network. The network instance at the termination step, g^ℓ, depicts the social structure resulted by the ensemble of agents' policies. Thus the change of g^ℓ represents the evolution of social structures as a function of behavioral acquisitions towards gaining social capital. We investigate what and how social structures emerge under different configurations of SCGs.

V. Learning: MARL

To learn a policy, an agent should have two basic abilities: (1) extract explicit and latent information from the observation; (2) based on received utilities, criticize and adjust the way to make decisions. To this end, our learning method is adapted from S2V-DQN as in [44], which is an end-to-end deep learning architecture to solve graph-based combinatorial problems. S2V-DQN incorporates *graph embedding* and reinforcement learning. Graph embedding is a technique of representation learning on graphs, learning latent features of network structures. It resolves two challenges in SCGs: (1) graphs are not fixed-size formatted data; (2) social surroundings of an agent can be too complex to learn from.

In our proposed MARL method for SCGs, all agents are independent learners that independently and synchronously use S2V-DQN to learn a policy. More formally, let (N, W, g^0, ℓ) be an SCG, each agent $i \in N$ estimates the quality of linking to another agent $a \in \mathcal{N}_2(i)$ under an observation (2-level ego network) o_i using an *evaluation function* $Q_i(o_i, a)$. The policy π_i thus naturally functions greedily with respective to Q_i, *i.e.*,

$$\pi_i(o_i) := \arg\max_{a \in \mathcal{N}_2(i)} Q_i(o_i, a). \tag{7}$$

S2V-DQN uses `structure2vec` [45] to parameterize $Q_i(o_i, a; \mathbf{\Theta}_i)$ that computes a p-dimensional feature embedding $\boldsymbol{\mu}_j$ for each node j involved in an observation o_i. $\boldsymbol{\mu}_j$ is iteratively updated. Initialized as $\mathbf{0}$, after T iterations, $\boldsymbol{\mu}_j$ will

[1] Since preference weights are fixed of a social capital game, the preferences weights of agents within agent i's observation is uniquely determined by o_i.

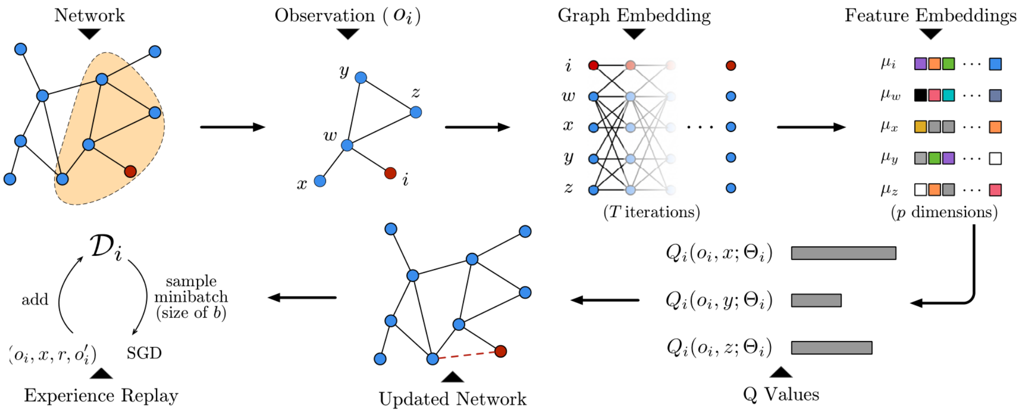

Fig. 1. An example for the learning processing for an single agent in a single time step. We run $T = 4$ iterations in the procedure of graph embedding. We set the dimension of vectors $p = 32$. The size of minibatch for experience replay is set as $b = 32$. Note that all agents execute the above procedure simultaneously while learning.

contain information about its T-hop neighbors as determined by the structure of o_i. The update rule is:

$$\boldsymbol{\mu}_j^{(t+1)} = \mathsf{ReLU}\left(\boldsymbol{\theta}_1 \boldsymbol{x}_j + \boldsymbol{\theta}_2 \sum\nolimits_{k \in \mathcal{N}_1(j)} \boldsymbol{\mu}_k^{(t)}\right), \tag{8}$$

where $\boldsymbol{\theta}_1 \in \mathbb{R}^{p \times 2}, \boldsymbol{\theta}_2 \in \mathbb{R}^{p \times p}$ are model parameters and ReLU is the rectified linear unit ($\mathsf{ReLU}(z) = \max(0, z)$). $\boldsymbol{x}_j$ is a vector that incorporates explicit features of j. In the context of SCGs, we set $\boldsymbol{x}_j = (w_j, \mathsf{dist}(i,j))^\intercal$.

The embedding $\boldsymbol{\mu}_a$ and the pooled embedding over the entire observation, $\boldsymbol{\phi}(o_i) := \sum_{j \in \mathcal{N}_2[i]} \boldsymbol{\mu}_j$, are used as the surrogates for a and o_i, respectively, *i.e.*,

$$Q_i(o_i, a; \boldsymbol{\Theta}_i) = \boldsymbol{\theta}_3^\intercal \mathsf{relu}(\boldsymbol{\theta}_4 \boldsymbol{\phi}(o_i) \oplus \boldsymbol{\theta}_5 \boldsymbol{\mu}_a), \tag{9}$$

where $\boldsymbol{\theta}_3 \in \mathbb{R}^{2p}, \boldsymbol{\theta}_4, \boldsymbol{\theta}_5 \in \mathbb{R}^{p \times p}$, and $\oplus$ is the concatenation operator. $Q_i(o_i, a; \boldsymbol{\Theta}_i)$ is based on a collection of 5 parameters $\boldsymbol{\Theta}_i = \{\boldsymbol{\theta}_m\}_{1 \le m \le 5}$, which will be learned.

The *experience replay* is used to update $\boldsymbol{\Theta}_i$ with a batch of samples drawn from a experience dataset $\mathcal{D}_i$. A dataset $\mathcal{D}_i$ is pooled over episodes such that for each $t < \ell$, a tuple $\left(o_i^t, a_i^t, r_i^{t+1}, o_i^{t+1}\right)$ is added. For the terminate step ℓ, we define $Q_i(o_i^\ell, a) \equiv 0$. For each step, a minibatch of tuples (size of b) is randomly sampled from $\mathcal{D}_i$. Then stochastic gradient descent is executed on the following squared loss:

$$\mathcal{L}(\boldsymbol{\Theta}_i) = \mathbb{E}_{(o,a,r,o') \sim \mathcal{D}_i}\left[(y - Q_i(o, a; \boldsymbol{\Theta}_i))^2\right], \tag{10}$$

where $y = r + \max_{a'} Q_i(o', a'; \boldsymbol{\Theta}_i)$ is the update target. The architecture of learning is depicted in Fig. 1. The pseudocode is illustrated in Alg. 1.

Algorithm 1: MARL for Social Capital Games

Input: A social capital game (N, W, g^0, ℓ)
Initialization: initialize experience dataset $\mathcal{D}_i$ for all $i \in N$
for *each episode* **do**
  for *step* $t = 0$ *to* $\ell - 1$ **do**
    for *agent* $i = 1$ *to* n **do**
      ▷ Choose actions subject to ϵ-greedy
$$a_i^t = \begin{cases} \text{random agent } a \in \mathcal{N}_2^t(i), & \text{w.p. } \epsilon \\ \arg\max_{a \in \mathcal{N}_2^t(i)} Q_i(o_i^t, a; \Theta_i), & \text{otherwise} \end{cases}$$
    $g^{t+1} = g^t + \{ia_i^t\}_{i \in N}$ ▷ Update Network
    for *agent* $i = 1$ *to* n **do**
      $r_i^{t+1} = \mathsf{mix}_{i,w_i}^{t+1} - \mathsf{mix}_{i,w_i}^t$ ▷ Compute rewards
    for *agent* $i = 1$ *to* n **do**
      ▷ Update experience dataset
      Add tuple $\left(o_i^t, a_i^t, r_i^{t+1}, o_i^{t+1}\right)$ to $\mathcal{D}_i$
    for *agent* $i = 1$ *to* n **do**
      ▷ Experience replay
      Sample a minibatch (of size b) $B \overset{iid.}{\sim} \mathcal{D}_i$
      Update Θ_i by SGD over Eq.(10) for B

VI. Emergence: Social Structural Patterns

We train $|N| = 100$ agents and set the initial network g^0 as a regular ring lattice, a graph with $|N|$ nodes each connected to two neighbors, one on each side. A lattice captures a homogeneous and loosely connected society which does not exhibit any meaningful community, small-world or core-periphery structure. Lattices are often used as initial configurations in network generation models, *e.g.*, small-world model [2]. We investigate the emergence of social structures via varying the terminate step ℓ and the preference vector W. For comparison, we employ random network generation models as baselines. Also, for each termination step $\ell \in \{2, 5, 8\}$, we generate 100 randomly created networks for reference, where in each step each agent randomly links to one from its 2-hop neighbors.

A. Emergence of community structure

Individuals naturally bind with each other to form groups. In a community structure, the communities form the social surrounding that provides individuals with a range of benefits such as social support, *i.e.*, bonding capital. Therefore, we make the following prediction:

Prediction 1. *Community structure emerges when all agents are in pure pursuit of bonding capital.*

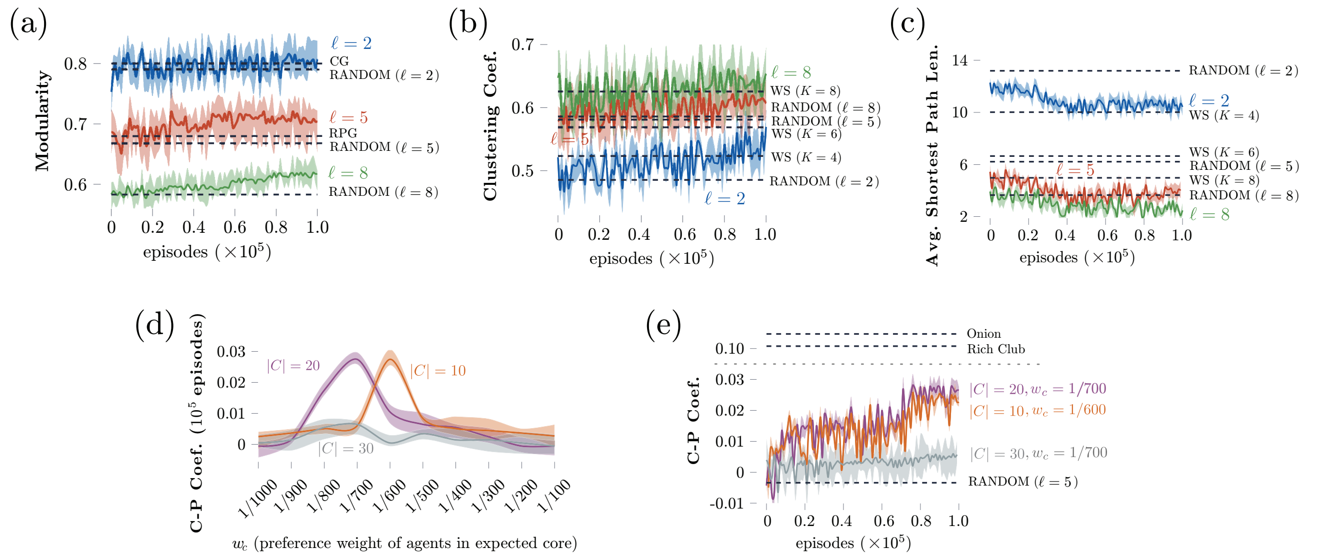

Fig. 2. Results for (a) modularity for the emergence of community structure; (b) clustering coefficients and (c) average shortest path length for the emergence of small-world structure; (d) C-P coefficients after 10^5 episodes and (e) change of C-P coefficients during learning for the emergence of core-periphery structure. The darker line shows the median over 10 independent runs and the shaded area is obtained by averaging the two extreme values.

For experimental settings, we set preference weight $w_i = 1$ for all $1 \leq i \leq 100$ and vary ℓ in $\{2, 5, 8\}$. After each training episode terminates, we use Louvain algorithm, a well-known community detection method [46], to compute a community structure of each g^ℓ. We then compute the modularity value of g^ℓ based on the obtained community strcuture. We conduct 10 independent runs.

We adopt two random graph generation models for the community structure as baselines:

- *Caveman graphs* (CG) [47]. It generates a graph by modifying a set of n isolated k-cliques (complete graphs) by removing one edge from each clique and using it to connect to a neighboring clique along a central cycle such that all n cliques form a single unbroken loop. Each modified clique forms a community.
- *Random partition graphs* (RPG). A graph is generated from n groups of isolated nodes $C_1, C_2, \ldots, C_n$, each of predefined sizes $s_1, s_2, \ldots, s_n$. Nodes in the same group are connected with probability p_{in} and nodes of different groups are connected with probability $p_{out} < p_{in}$. Each group is counted as a community.

We generate 100 instances for each model by varying model parameters ($n \cdot k = 100$ for caveman graphs; $\sum_i^n s_i = 100$, $p_{in} = 0.25$ and $p_{out} = 0.01$ for random partition graphs).

Results. See results in Fig. 2(a). Snapshots are shown in Fig. 3(a) that depict the emergence of communities. Three facts stand out: (1) Modularity grows and fluctuates as the number of training episodes increase. (2) Modularity output by our learning framework is higher than randomly generated networks. (3) Modularity resulted by our learning framework when $\ell = 2$ and $\ell = 5$ is comparable to by CG and RPG, respectivey. These results clear suggest the emergence of community structure, which verifies Prediction 1. Interestingly, the modularity shows a negative correlation with the termination step ℓ. This implies the fact that a capability to establish more relationships tends to drive agents to meet more other faraway agents, diminishing the formation of communities.

B. Emergence of small-world structure

As mentioned above, a small-world network tends to have high clustering coefficient and small average shortest path length. These properties provides agents with higher ability to coordinate communications across the network. Such abilities are conceptually consistent with bridging social capital. Therefore, we make the following prediction:

Prediction 2. *Small-world structure emerges when all agents are in pure pursuit of bridging capital.*

We set set preference weight $w_i = 0$ for all $1 \leq i \leq 100$ and vary ℓ in $\{2, 5, 8\}$. We adopt *Watts-Strogatz* (WS) model [2], a famous small-world network generation model, as the baseline. The model starts from a regular lattice, each node connected to K neighbors and $K/2$ on each side. Then edges are randomly rewired with probability p. When $p \in [0.01, 0.1]$, the graph typically demonstrates a small-world property. We fix p to 0.01 and vary K in $\{4, 6, 10\}$. For each value of K, we generate 100 instances and use the average clustering coefficient and shortest path length for comparison. Note that each value of K corresponds to a value of ℓ, as the maximum number of relationships that an agent can hold in SCGs is $\ell + 2$. Again, we generate 100 random networks for each ℓ in $\{2, 5, 8\}$.

Results. Fig. 2 (b), (c) and Fig. 3 (b) depict results and snapshots, respectively. Our training framework achieves comparable high clustering coefficients and lower average shortest path lengths for each value of ℓ, compared to the baselines with a corresponding value of K. This is a strong sound for the reproduction of the small-world structure. Moreover, both clustering coefficient and average path length show a positive

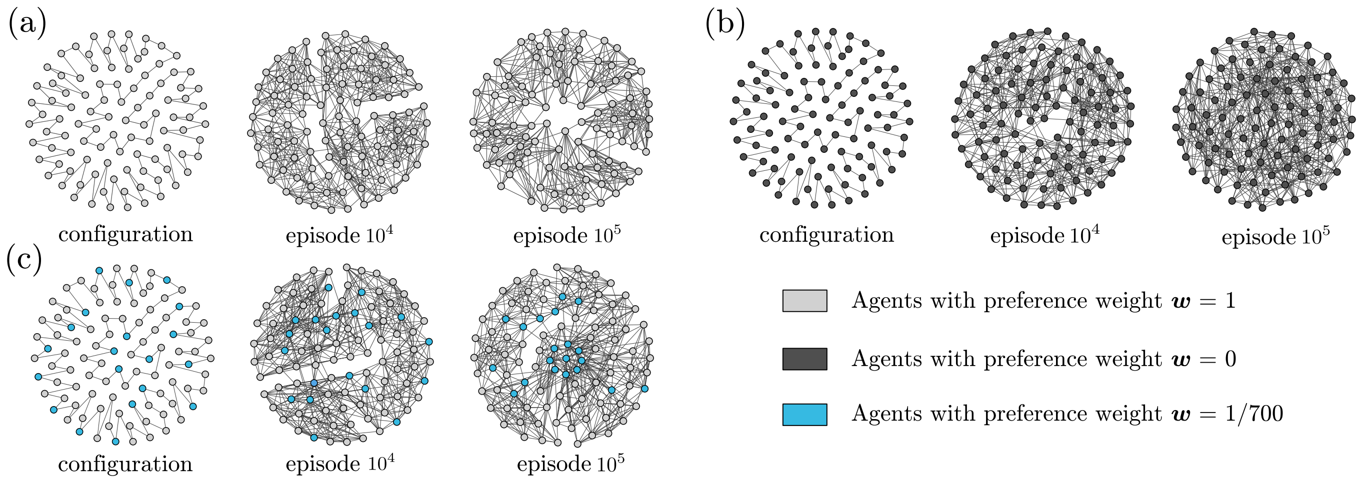

Fig. 3. Initial network configurations and snapshots during learning for (a) community structure, (b) small-world structure and (c) core-periphery structure (termination step $\ell = 5$).

correlation with ℓ. This phenomenon is easy to understand as a lager capacity to build relationships can bring agents an expanded vision in a network. As a result, agentss are more likely to get contacted with each other, thus the "smaller world" emerges.

C. Emergence of core-periphery structure

For the emergence of core-periphery (C-P) structure, we probe into the formation of two substructures, the core and the periphery. Three questions arise: (1) why do individuals in the periphery stay at the edge of society? (2) why are a part of individuals located in the center? (3) how are central individuals connected to form a core? The emergence of communities may provide an answer to the first question: a group of agents are only interested in bonding capital and they thus stay at the periphery to exploit their social surroundings. While, the emergence of small-world structure can answer the second question: the pursuit of bridging capital drives a group of individuals to bridge those at the periphery, resulting in a center. The combination of bonding and bridging capital may resolve the third question: those in the center are in pursuit of both bridging and bonding capital; the former pushes them to the center and the latter draws them together to form a core. Hence, the following prediction comes out:

Prediction 3. *Core-periphery structure emerges when a group of agents are in pure pursuit of bonding capital, while the other group of agents show mixed preferences to bonding and bridging capital.*

We randomly select a subset $C \subset N$ (expected core) with varying size in $\{10, 20, 30\}$. For all $c \in C$, we vary w_c from $1/1000$ to $1/100$[2]. For each remaining agent $p \in N \setminus C$ (expected periphery), we set $w_p = 1$. Throughout, we fix ℓ to an intermediate value, 5.

For baselines, we use two typical C-P network models, *rich club model* and *onion model* [3]. Both models develop dense cores. The former generates sparse peripheries, while the latter generates peripheries in the form of several layers surrounding the core. We generate 100 instances for each model by varying model parameters, each of size 100.

[2]Because the measured value of bonding capital is larger than of bridging capital. To balance two types of social capital, we set preference weighs to small values.

Results. Fig. 2 (d) shows the clustering coefficients after 10^5 episodes of learning, where three peaks occur under $(|C| = 10, w_c) = (10,\ 1/600)$, $(20, 1/700)$ and $(30, 1/700)$. Fig. 2 (e) plots the change of clustering coefficients during learning. Snapshots are shown in Fig. 3 (c), captured when $(|C|, w_c) = (20, 1/700)$. Results indicate that the tendency towards C-P structure rises when both the size of core and preferences to social capital are properly set. This implies that in agent societies, social polarization occurs when agents show different preferences to social capital. Moreover, the ambitious ones tend to get connected if they need to earn capital by uniting others. Another fact is that our C-P coefficients output by our framework is lower than by baseline models. This is not surprising as instances generated by C-P network generation models are with perfect C-P structure but unrealistic, limited to depict the C-P phenomenon in the real world. On the other hand, as C-P coefficients of randomly generated networks are negative, the positive values obtained in learning indicates that the tendency towards C-P is significant and meaningful for reproducing the C-P structure.

VII. Conclusions and Outlook

Towards gaining an insight on the emergence of prominent social structural patterns, we explore the potential role of social capital in driving the formation of social ties. We conceptualize bonding and bridging social capital from a structural perspective. Bonding capital represents an individual's benefits earned by exploiting social surroundings, while bridging capital captures utilities coming from exploring the far-away society. We propose an agent-based model to capture the process of relationship building as agents pursue higher social capital. In this model, agents are equipped with reinforcement learning abilities. We observe that the different preferences towards social capital play a key role in the emergence of various social structures: bonding capital drives the community structure to emerge, bridging capital leads to small-world structure and a combination of bonding and bridging capital contributes to the emergence of core-periphery structure. Overall, our framework unifies the explanations for the emergence of classical social structures.

Ideas and methods proposed in this paper represent a novel research initiative. There are several potential directions for future work. A fairly straightforward expansion is to allow an agent's preference weight to change over time. This extends the dynamics of current model and may result in the emergence of more types of social structures. Another future challenge is to add the mechanisms of node addition and edge removal to our current model. A third future work is to extend the current model through asynchronously organizing agents. These future work would help bring a comprehensive understanding of the natural emergence of social structures.

References

[1] M. Girvan and M. E. Newman, "Community structure in social and biological networks," *Proceedings of the national academy of sciences*, vol. 99, no. 12, pp. 7821–7826, 2002.

[2] D. J. Watts and S. H. Strogatz, "Collective dynamics of 'small-world' networks," *Nature*, vol. 393, no. 6684, p. 440, 1998.

[3] P. Csermely, A. London, L.-Y. Wu, and B. Uzzi, "Structure and dynamics of core/periphery networks," *Journal of Complex Networks*, vol. 1, no. 2, pp. 93–123, 2013.

[4] R. S. Burt, "The network structure of social capital," *Research in organizational behavior*, vol. 22, pp. 345–423, 2000.

[5] J. S. Coleman, "Social capital in the creation of human capital," *American journal of sociology*, vol. 94, pp. S95–S120, 1988.

[6] D. G. Myers, "Relationship rewards," *Social Psychology*, pp. 418–420, 2010.

[7] N. Lin, *Social capital: A theory of social structure and action*. Cambridge university press, 2002, vol. 19.

[8] P. Bourdieu, "The forms of capital," *Handbook of Theory and Research for the Sociology of Education*, 1986.

[9] R. S. Burt, "Structural holes and good ideas," *American journal of sociology*, vol. 110, no. 2, pp. 349–399, 2004.

[10] P. S. Adler and S.-W. Kwon, "Social capital: Prospects for a new concept," *Academy of management review*, vol. 27, no. 1, pp. 17–40, 2002.

[11] A. Moskvina and J. Liu, "How to build your network? a structural analysis," *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 2597–2603, 2016.

[12] B. Yan, Y. Chen, and J. Liu, "Dynamic relationship building: exploitation versus exploration on a social network," in *International Conference on Web Information Systems Engineering*. Springer, 2017, pp. 75–90.

[13] M. O. Jackson and A. Wolinsky, "A strategic model of social and economic networks," *Journal of economic theory*, vol. 71, no. 1, pp. 44–74, 1996.

[14] S. Currarini, M. O. Jackson, and P. Pin, "An economic model of friendship: Homophily, minorities, and segregation," *Econometrica*, vol. 77, no. 4, pp. 1003–1045, 2009.

[15] B. Yan, Y. Liu, J. Liu, Y. Cai, H. Su, and H. Zheng, "From the periphery to the core: Information brokerage in an evolving network," in *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, 2018, pp. 3912–3918.

[16] Y. Cai, H. Zheng, J. Liu, B. Yan, H. Su, and Y. Liu, "Balancing the pain and gain of hobnobbing: Utility-based network building over atributed social networks," in *AAMAS*, 2018, pp. 193–201.

[17] Y. Song and M. van der Schaar, "Dynamic network formation with incomplete information," *Economic Theory*, vol. 59, no. 2, pp. 301–331, 2015.

[18] B. Yan, K. Song, J. Liu, F. Meng, Y. Liu, and H. Su, "On the maximization of influence over an unknown social network," in *AAMAS*, 2019, pp. 2279–2281.

[19] Coleman, *Foundations of social theory*. Harvard university press, 1994.

[20] M. S. Granovetter, "The strength of weak ties," in *Social networks*. Elsevier, 1977, pp. 347–367.

[21] J. Bouma, E. Bulte, and D. van Soest, "Trust and cooperation: Social capital and community resource management," *Journal of environmental economics and management*, vol. 56, no. 2, pp. 155–166, 2008.

[22] D. P. Aldrich and M. A. Meyer, "Social capital and community resilience," *American behavioral scientist*, vol. 59, no. 2, pp. 254–269, 2015.

[23] M. Hällsten, C. Edling, and J. Rydgren, "Social capital, friendship networks, and youth unemployment," *Social Science Research*, vol. 61, pp. 234–250, 2017.

[24] M. O. Jackson and B. W. Rogers, "Meeting strangers and friends of friends: How random are social networks?" *American Economic Review*, vol. 97, no. 3, pp. 890–915, 2007.

[25] J. Leskovec, L. Backstrom, R. Kumar, and A. Tomkins, "Microscopic evolution of social networks," in *SIGKDD*. ACM, 2008, pp. 462–470.

[26] A. M. Alaa, K. Ahuja, and M. van der Schaar, "A micro-foundation of social capital in evolving social networks," *IEEE Transactions on Network Science and Engineering*, vol. 5, no. 1, pp. 14–31, 2017.

[27] M. E. Newman, S. H. Strogatz, and D. J. Watts, "Random graphs with arbitrary degree distributions and their applications," *Physical review E*, vol. 64, no. 2, p. 026118, 2001.

[28] I. M. Kloumann and J. M. Kleinberg, "Community membership identification from small seed sets," in *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2014, pp. 1366–1375.

[29] L. Backstrom and J. Leskovec, "Supervised random walks: predicting and recommending links in social networks," in *WSDM*. ACM, 2011, pp. 635–644.

[30] S. Sen and S. Airiau, "Emergence of norms through social learning." in *20th International Joint Conference on Artificial Intelligence*, vol. 1507, 2007, p. 1512.

[31] S. Hu, C.-w. Leung, H.-f. Leung, and J. Liu, "To be big picture thinker or detail-oriented? utilizing perceived gist information to achieve efficient convention emergence with bilateralism and multilateralism," in *AAMAS*, 2019, pp. 2021–2023.

[32] S. Mahmoud, N. Griffiths, J. Keppens, and M. Luck, "Efficient norm emergence through experiential dynamic punishment." in *20th European Conference on Artificial Intelligence (ECAI)*, vol. 12, 2012, pp. 576–581.

[33] M. Mihaylov, K. Tuyls, and A. Nowé, "A decentralized approach for convention emergence in multi-agent systems," *Journal of Autonomous Agents and Multi-Agent Systems*, vol. 28, no. 5, pp. 749–778, 2014.

[34] J. Z. Leibo, V. Zambaldi, M. Lanctot, J. Marecki, and T. Graepel, "Multi-agent reinforcement learning in sequential social dilemmas," in *AAMAS*, 2017, pp. 464–473.

[35] E. Hughes, J. Z. Leibo, M. Phillips, K. Tuyls, E. Dueñez-Guzman, A. G. Castañeda, I. Dunning, T. Zhu, K. McKee, R. Koster *et al.*, "Inequity aversion improves cooperation in intertemporal social dilemmas," in *NIPS*, 2018, pp. 3326–3336.

[36] J. X. Wang, E. Hughes, C. Fernando, W. M. Czarnecki, E. A. Duéñez-Guzmán, and J. Z. Leibo, "Evolving intrinsic motivations for altruistic behavior," in *AAMAS*, 2019, pp. 683–692.

[37] Y. Liu, J. Liu, Z. Zhang, L. Zhu, and A. Li, "Rem: From structural entropy to community structure deception," in *Advances in Neural Information Processing Systems*, 2019, pp. 12 938–12 948.

[38] P. Holme, "Core-periphery organization of complex networks," *Physical Review E*, vol. 72, no. 4, p. 046111, 2005.

[39] Y. Chen, J. Liu, H. Zhao, and H. Su, "Social structure emergence: A multi-agent reinforcement learning framework for relationship building," in *AAMAS*, 2020, pp. 1807–1809.

[40] H. Tong, C. Faloutsos, and J.-Y. Pan, "Fast random walk with restart and its applications," in *6th International Conference on Data Mining (ICDM)*. IEEE, 2006, pp. 613–622.

[41] J. Liu, L. Li, and K. Russell, "What becomes of the broken hearted? an agent-based approach to self-evaluation, interpersonal loss, and suicide ideation," in *AAMAS*, 2017, pp. 436–445.

[42] Q. Chen, H. Su, J. Liu, B. Yan, H. Zheng, and H. Zhao, "In pursuit of social capital: Upgrading social circle through edge rewiring," in *APWeb-WAIM*. Springer, 2019, pp. 207–222.

[43] A.-L. Barabási and E. Bonabeau, "Scale-free networks," *Scientific american*, vol. 288, no. 5, pp. 60–69, 2003.

[44] E. Khalil, H. Dai, Y. Zhang, B. Dilkina, and L. Song, "Learning combinatorial optimization algorithms over graphs," in *NIPS*, 2017, pp. 6348–6358.

[45] H. Dai, B. Dai, and L. Song, "Discriminative embeddings of latent variable models for structured data," in *International conference on machine learning (ICML)*, 2016, pp. 2702–2711.

[46] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre, "Fast unfolding of communities in large networks," *Journal of statistical mechanics: theory and experiment*, vol. 2008, no. 10, p. P10008, 2008.

[47] D. J. Watts, "Networks, dynamics, and the small-world phenomenon," *American Journal of sociology*, vol. 105, no. 2, pp. 493–527, 1999.