A Model of Minority Influence in Preferential Norm Formation

Qiang Liu[1], Hong Zheng[1], Weidong Li[2](✉), Jiamou Liu[2], Bo Yan[1], and Hongyi Su[1]

[1] Beijing Lab of Intelligent Information Technology, School of CS and Technology, Beijing Institute of Technology, Beijing, China
hongzheng@bit.edu.cn
[2] School of Computer Science, The University of Auckland, Auckland, New Zealand
wli916@aucklanduni.ac.nz

Abstract. Minority influence takes place when a minority group influences the majority group, which has been widely studied in psychology and sociology. Most existing works are, however, qualitative studies or based on field experiments; large-scale quantitative experiments are hard to conduct due to the lack of a model. Thus, we propose a new agent-based framework, Minority Influence Model, for modelling the interactions among individuals to explore minority influence in the context of norm emergence. What makes the model novel is that it is an integration of three key aspects. First, the model considers multiple player interactions; second, the actions of agents are driven by a combination of endogenous preference and exogenous social surrounding; third, the decision-making strategy for each agent is based on the evolution of gist. With Minority Influence Model, considerable experiments have been performed. The experimental results suggest that the size of the minority group, the extent to which the agents of the minority group prefer an action to the other, and the degree of disagreement between the majority and the minority can be the three main factors of minority influence. We also found some tipping points under some specific scenarios.

Keywords: Minority influence · Social convention · Social norm

1 Introduction

Emerson wrote that "all history is a record of the power of minorities and of minorities of one"[1]. The power of minorities, called *minority influence*, takes place when one or more members of a minority group influence the members of the majority group to accept or conform to the behaviour of the minority group [6]. Understanding minority influence gives us insight into social events and social movements. For example, it can help us to understand the feminist

[1] The Later Lectures of Ralph Waldo Emerson, 1843–1871. University of Georgia Press. 2010.

J. Chen et al. (Eds.): KSS 2019, CCIS 1103, pp. 106–121, 2019.
https://doi.org/10.1007/978-981-15-1209-4_8

movement in the last century, which was led by a minority group and persuaded the majority. Also, exploring the factor of minority influence may give us ideas of how to promote or prevent the influence of the minority. For instance, for companies planning to launch new products, it may be helpful to analyse how to attract more users and then dominate the market.

There are considerable studies regarding factors affecting minority influence in the human sciences [13,19,21,27]. most of which are qualitative research. It has been claimed that people's actions are influenced by both their attitudes and perceived social pressure [1,14,18,20]. In groups where individuals are able to interact with each other, individuals under real or imagined social pressure may change their behaviours to be consistent with social conventions [19]. At the same time, individuals holding attitude different from the majority can influence the whole group, if they are consistent and self-confident, and individuals by the minority defect the majority group [19]. However, due to the lack of a model that simulates communication among individuals and the decision-making process of each individual, quantitative studies are hard to conduct.

Consider a scenario with three individuals, Alex, Bob and Charlie. Charlie prefers playing football, while Alex and Bob rate playing football and shopping the same. Additionally, they do not want to go out alone, so they expect some others can choose the same activity as they do. Ideally, they choose the preferred activity and let others choose the same action. For instance, Charlie wants to play football and persuade Alex and Bob to join him. If others do not join him, Charlie might rather go shopping with them. In other words, individuals consider both their preferences and the actions of others when they make decisions. In this example, if Charlie, as the minority, can affect the action of others, the minority influence occurs. The question lies in how to model this scenario.

This scenario requires a framework which is able to model the interactions among multiple individuals, and take the preference of each and the influence of the group on each into consideration. To achieve this, we build the Minority Influence Model and then investigate minority influence in the context of norm emergence. Norm emergence is a well-established field of research in the multi-agent system that describes the emergence of coordinated actions among multiple interacting agents. Through engaging in repeated coordination games, agents may be able to carry out social learning, updating individual actions according to self-interest, while eventually arriving at unity in their action. Minority influence can be regarded as the emergence of the social norm consistent with the action of the minority group. For example, after interacting with each other, Alex and Bob can be affected by the minority Charlie, and choose to play football with him, then a social norm, as well as minority influence, occurs. Thus, a model of norm emergence can be used to investigate minority influence. For any model of norm emergence, there are two main ingredients. The first is the way how agents interact, and the second is the way how agents make decisions. In this paper, the Minority Influence Model is built as follows:

To define the way in which agents interact, we use a *multi-player coordination game with preference*. This game involves a number of players each of which having the same 2 actions to choose from, and their payoffs are determined by (1) their own preference over these two actions and (2) the percentage of agents in the game who had the same action as theirs. The underlying motivation for the separation between preference and actions, as well as the concept of social conformity, is backed by social psychological studies [1,19,20].

To define the way in which agents make decisions, we use a *gist-based reinforcement learning* paradigm. The paradigm assumes the following: (1) agents are able to observe the actions from others that they interact with in the past and (2) such past information is not precisely stored, but rather, is abstracted into a form of "gist", general belief regarding the public's opinion. Inspired by the fuzzy trace theory [2], the notion of gist-based Q-learning was recently proposed in the work of Hu et al. [8], where the Q-value of each action can be determined by current perceived gist. It is a major novelty in this paper that this paradigm is used in the context of the multi-player coordination games with preference.

The combination of the two ingredients above has resulted in a new model that enables the investigation of minority influence. The significance of this new model is that it is the first which integrates multiple aspects: (1) multiple player interaction (2) agents actions are driven by a combination of endogenous preference and exogenous social surrounding (3) the evolution of gist in explaining the formation of social norms. To explore minority influence, we conduct a set of experiments to reveal potential correlations between the many factors involved and the outcome of the social norm. The experimental results suggest that minority influence can be correlated with the size of the minority group, the degree to which the minority prefers some certain action, and the degree of disagreement between the majority and the minority. Tipping points can be observed under two scenarios. The first is that the majority has a neutral attitude, and a contrary action to the minority. The second is that the minority holds a contrary preference to the majority.

The paper is organised as follows: In Sect. 2, we review related literature; in Sect. 3, we propose the Minority Influence Model; in Sect. 4, we present a set of experiments based on the proposed model; in Sect. 5, we draw a conclusion and suggest some possible future work.

2 Related Work

Numerous studies in psychology and sociology have analysed minority influence. Myers regarded minority influence as the fact that individuals can influence their groups [19]. Milgram et al. claimed that a group with a larger size could be more influential [16]. Moscovici argued that the minority who are consistent in its attitude could be more influential than a wavering minority [17]. Allen and Levine claimed that when the minority doubt the attitude of the majority consistently, some of the members of the majority can defect to the minority. Once switch to the minority, they can be more influential than initial members of the minority [13]. Furthermore, once defections start, members of the majority can soon follow the defector and switch to the minority [19].

However, most of the studies above are qualitative studies. To conduct quantitative experiments, a model is required. Minority influence can be viewed as the process of the majority conforming to the action of the minority, which in fact establishes a social norm. Thus, models of norm emergence can be used to explore minority influence. Early studies regarding norm emergence introduced the notion of emergent conventions [23]. In 1993, Kittock adopted 90% convergence metric and claimed that a convention emerges in a system if at least 90% of the agents chooses the same action [12], which is widely used in studies of norm emergence as well as this paper. Models of norm emergence consist of two main ingredients. The first is the way that agents interact, and the second is the method that agents make decisions.

In terms of the interaction ingredient, there are extensive research works. In 1997, Shoham and Tennenholtz defined standard game-theoretic notions such as n-k-g stochastic game [24]. n-k-g stochastic game illustrates the repeated process of randomly selecting k agents from n given agents and letting them play a certain base game. This framework has been widely used in the domain of norm emergence and also adopted in this paper.

Coordination game [12] can be used as the base game of n-k-g stochastic game. In a coordination game, all agents can receive the same payoff if they choose the same action, which indicates that the goal of the agents is to achieve uniformity [5,24]. Coordination game ignores the fact that individuals can have different preferences on different actions in real life. To enrich this model, several variants have been introduced, one of which is unbalanced coordination game [5]. Compared with standard coordination game, agents can receive different payoffs when they converge to different actions. It considers the difference between actions, while it is still not suitable for scenarios where individuals have different preferences. Another variant is competitive-coordination game [9], also known as the battle of the sexes [5], which is proposed to simulate the situation where agents tend to build different conventions, although they intend to convention emergence. In other words, agents prefer coordination than dis-coordination, but different agent prefers different coordination. Competitive-coordination game considers the different preference of each agent, but it is not suitable for modelling the situations where more than two agents are involved. One more variant is n-player coordination game [8], which generalises the standard coordination game and can be used to model the interactions among more than two individuals. The payoff each agent receiving is based on the proportion of other agents choosing the same action. For each agent, the more other agents choosing the same action as it, the more payoff it receives. However, n-player coordination game does not consider the different preference of each agent. Therefore, to take both endogenous preference and exogenous social surrounding into consideration, we combine competitive-coordination game and n-player coordination game to establish the interaction ingredient of the Minority Influence Model.

Decision-making strategy is another essential part of norm emergence models. As individuals learn from experience [24], the decision-making strategy essentially represents the "brain" for each individual. In early studies, memory-based strategies such as the Highest Cumulative Reward (HCR) update rule has been used as a decision-making method [24]. Recently, researchers apply reinforcement learning algorithms to decision-making strategies [8]. Reinforcement learning allows an agent to learn a sequence of actions through trial-and-error interactions with a dynamic environment [10]. Q-learning is a model-free reinforcement learning which dynamically evaluates the available actions [26]. It has been shown that Q-learning with ϵ-greedy exploration leads to faster convention emergence compared to other methods [22]. Thus, many works try to extend Q-learning. Hao and Leung extended the idea of joint action learner [4] to norm emergence problem, and proposed a decision-making method that agents learn the q-values of joint actions, i.e., a combination of their own actions and the actions of their neighbours [7]. Hu et al. modified Q-learning to establish gist-based Q-learning [8], based on the fuzzy trace theory [2].

Gist is a psychological concept that refers to a type of vague and high-level mental representation of events [2]. In real life, individuals tend to memorise events by gist, i.e., memorise only general information instead of details. Inspired by this, gist-based Q-learning has been proposed [8]. At each iteration, agents update their Q-values based on the perceived prevalence of actions and choose the action with the higher Q-value. The perceived prevalence is implied from the actions of other agents it interacts with at this iteration. Gist-based Q-learning can reflect the influence of the actions of other individuals on one single agent, which is consistent with our assumption in this paper. Thus, we borrow and modify gist-based Q-learning in the Minority Influence Model, which is introduced next.

3 Minority Influence Model

Imagine a situation, say, that a few individuals are discussing whether to buy a new mobile phone, and there are only two decisions available, to buy or not to buy. Due to different background, personality, and position, people can hold different opinions about this mobile phone. Some people might prefer to buy the phone because they need a new phone or like the manufacturer, while others can prefer not to buy it because they do not like the design of this phone or they are fans of another brand. Additionally, people can not only favour different decisions, but they can prefer different decisions to different extents. Some people may feel it is acceptable to either purchase the phone or not purchase. In contrast, some people extremely want to own the new product, and some individuals can hate the phone.

On the other hand, people tend to follow the action of the majority to avoid punishment and get a benefit [3,15,19,20,28]. In this scenario, a possible consequence can be isolation from friends, when one person decides not to buy the phone, while most of his friends want to buy it. Thus, people need to balance their preferences and the influence of the majority to make the right decision, which can bring more benefit. However, people in the real world cannot know the opinion of each person in society. Instead, they can only get the opinion of

Algorithm 1. The Framework of the Minority Influence Model

```
Input agents N, number of iterations T, number of peers k
initialisation (N)
t = 0
while t < T do
    for agent ∈ N do
        randomly select k agents from N \ {agent}, set them to P
        players = {agent} ∪ P
        n-player coordination game with preference (players)
    end for
    t = t + 1
end while
```

people whom they have met and communicated. Their knowledge of the majority is derived from their observation of the people around them. As a result, their perception of the opinion of the majority can change as they interact with new people or people who change their ideas.

What we are interested in is that under the circumstance discussed above, whether a minority of people with extreme opinions can influence the whole society. In order to explore this problem, we propose a new model, the Minority Influence Model, to simulate how people get influenced by their own preferences and others' opinions. We first introduce the framework of the Minority Influence Model, then explain two mechanisms applied in the model, namely n-player coordination game with preference that simulates how people interact with each other, and decision-making method that mimics how people make decisions in this particular scenario.

3.1 The Framework of the Proposed Model

This framework (Algorithm 1) is an overview of the model which simulates how people get influenced by their own preferences and others' behaviours. Assume that there are n people in the society and they can interact with some others every day. Here we apply repeated game, which means some base game, i.e., n-player coordination game with preference is repeatedly executed. After t days' interaction, we observe whether the minority can influence the majority.

To formalise this process, we use N to denote the set of individuals $a_1, \ldots, a_n$, T to denote the number of iterations, $A = \{x, y\}$ to represent the two actions available to each agent. Firstly, we initialise N. Each agent a_i is given two values α and β, which denote the extent to which the agent prefers action x and y, respectively. Then, for each agent, we randomly select k agents as its peers or neighbours, and let the $k + 1$ agents play the n-player coordination game with preference. We repeat this process for T times. More details about the n-player coordination game with preference are given in the next section. Finally, we observe the actions of each agent.

3.2 N-Player Coordination Game with Preference

N-player coordination game with preference is the core part of the Minority Influence Model. It reflects how a group of people share their own opinions or decision about one topic, influence each other, and get influenced by each other, then evaluate their own opinions or decisions. It formalises the interactions between individuals, considering both the power of the preference of each individual and the opinion of the majority.

For example, let us consider a 4-player coordination game with preference. Assume that these four individuals discuss whether to buy a newly-listed mobile phone. In the beginning, only one person wants to buy this phone, because he is a fan of the company which designed the product. Other three people are fans of another company, so decide not to buy it. For the person who wants to purchase the phone, he evaluates his decision by considering both his own taste and others' opinions. If he purchases the phone, he can be happy because of making a decision he preferred, but others may judge him and make him feel isolated. If he does not buy the phone, he can integrate into the group, but cannot get the product he loves. After consideration, he still wants to purchase the phone, because he is loyal to the company. For other people who do not want buy the phone, their tastes are consistent with the opinion of the majority, so they do not change their decisions.

To formalise this process, we propose n-player coordination game with preference. We use agents to denote individuals, actions to denote the opinion or decision of individuals, and rewards to represent the result of the evaluation on decisions of each individual. Individuals prefer holding the same opinion as others than holding a different idea, so we say that one individual prefers the convention λ_x if he prefers the action x.

Let N be the set of all individuals, $A = \{x, y\}$ be the set of decisions available. N_x be the set of agents prefer convention λ_x, and N_y be the set of agents prefer convention λ_y, such that N_x and N_y are disjoint sets, and $N_x \cup N_y = N$.

Definition 1. *An n-player coordination game with preference is a 4-tuple*

$$\langle N, A, (p_i), (r_i) \rangle,$$

where $N = \{1, 2, \cdots, n\}$ is the set of n agents; $A = \{x, y\}$ is the set of actions that agents can perform; p_i is the convention that an agent i prefers, i.e., $p_i = x$ iff $i \in N_x$; (r_i) is the reward function which is defined as follows:

$$r_i = \begin{cases} \alpha \times \dfrac{\sum_{j=1}^{n} I(a_i, a_j) - 1}{n-1}, & \text{if } a_i = p_i, \\[2ex] \beta \times \dfrac{\sum_{j=1}^{n} I(a_i, a_j) - 1}{n-1}, & \text{otherwise}, \end{cases} \tag{1}$$

where α and β are the preference values of the preferred action and the other action of the agent i respectively, $\alpha > \beta > 0$, and $I\colon A \times A \to \{0, 1\}$ is a function such that:

$$I(a_i, a_j) = \begin{cases} 1 & \textit{if } a_i = a_j, \\ 0 & \textit{otherwise.} \end{cases} \tag{2}$$

For each agent of this coordination game, the preference values α and β are the maximum rewards that it can receive by performing the preferred action and the ill-favoured action, respectively. They denote to what extent an agent prefers some action or convention. The actual reward it receives is calculated by the corresponding preference value multiplies the proportion of its neighbours who are performing the same action as it. The reward depends on both the action it chooses and the decisions of other agents involved in this game, which is supported by psychological studies [1,20]. If an agent decides to choose action x, the more other agents choose action x, the higher reward this agent receives. Similarly, if an agent can predict that most of its peers will choose action x, it can then compute the possible rewards of performing each action, then choose the action with a higher reward.

It is not necessarily true that the agent performing the preferred action, which results in a higher preference value, leads to a higher reward. For example, consider a 4-player coordination game with preference. Let agent a_1 prefers action x with a preference value 8, and a preference value 2 for another action y. Let other agents a_2, a_3 and a_4 perform action y. If agent a_1 performs action x, the reward it receives is $8 \times (0/3) = 0$, and if it chooses action y, the reward it receives is $2 \times (3/3) = 2$. Similarly, it is also not necessarily true that a convention where an agent chooses the action choosing by the majority of the group can bring a higher reward than a non-convention.

N-player coordination game with preference models the process of people influencing each other during the interaction. At the end of each game, people can receive some feedback or evaluation of their decisions (the rewards). People tend to learn from each game to make a better decision in the future. Next, we introduce the mechanism used by agents to make better decisions.

3.3 Decision-Making Under Gist-Based Q-Learning

As we discussed before, people's decisions can be influenced by their own preferences and the opinions of other people. As gist-based Q-learning [8] illustrates how people make decisions under the influence of others, we modify this algorithm to build the decision-making mechanism of the proposed model.

The same as previous, each individual has two possible actions x and y. The reward they received is computed by the mechanism introduced before. We assume all people are gist-neutral, which means they are objective so that their perceived prevalence of actions is the actual frequency of actions [8].

Let $A = \{x, y\}$ be the set of actions that available for each agent, following gist-based Q-learning, the Q-value of an action x at time t is calculated as follows:

$$Q_x^t = b_x^t + w_y^t p_y^t \tag{3}$$

where b_x^t is a bias parameter that describes how well to perform action x without considering others' decisions, w_y^t is a weight that reflects the correlation between

Algorithm 2. The Modified Gist-based Q-learning

Input the performed action x, the perceived prevalence of the other action p_y, reward r, learning rate η

$Q_x = b_x + w_y p_y$

$b_x = b_x + \eta(r - Q_x)$

$w_y = w_y + \eta(r - Q_x)p_y$

$Q_x = b_x + w_y p_y$

the Q-value of action x the perceived prevalence of action y, and p_y^t denotes the proportion of agents choosing action y in the neighbourhood. As we assume that all agents are gist-neutral, p_y^t can be used to represent the perceived prevalence of action y.

Q-values for each action denote expect rewards of performing these actions. Thus, at each iteration, agents perform the action with a higher Q-value. During the coordination game, they observe others' actions, and then update their q-table. The method to update Q-values is introduced as follows.

In gist-based q-learning, instead of updating Q-values directly, agents learn the bias parameter b and the weight w to minimise the deviation between the Q-values of actions and the actual reward the agents receive by choosing these actions. Gradient descent is used here to learn the values of b and w. Assume that an agent chooses action x at this iteration, gets the reward r, and observes the proportion of people choosing the other action p_y. Algorithm 2 shows how we update the values of b and w, and then update the Q-value for the agent.

For instance, an agent chooses the action x, gets a reward 5, and observes 50% of his peers choose the other action y. Set the learning rate η to 0.1, and assume the initial value of b_x and w_y are 4 and 1, respectively. We show the procedures as follows. Firstly, compute the Q-value: $Q_x = b_x + w_y p_x = 4 + 1 \times 0.5 = 4.5$. Secondly, update the value of b_x: $b_x = b_x + \eta(r - Q_x) = 4 + 0.1(5 - 4.5) = 4.05$. Thirdly, update the value of w_y: $w_y = w_y + \eta(r - Q_x)p_y = 1 + 0.1(5 - 4.5) \times 0.5 = 1.025$. Finally, update the Q-value: $Q_x = b_x + w_y p_y = 4.05 + 1.025 \times 0.5 = 4.5625$. After this process, the agent updates its expected reward for action x. At the next iteration, the agent can choose the action with a higher Q-value, get a reward, then update its q-table again.

To sum up, the Minority Influence Model consists of a base game and a decision-making strategy. By repeating the base game, agents can learn from interactions and may switch to another action. The action they choose is based on their own preferences and the actions of others. Based on this framework, in the next section, we conduct experiments to explore minority influence.

Table 1. Parameters for experiments

Parameter	Description	Default value
n	The number of agents	100
T	The number of iterations	1000
E	The number of repeated experiments	100
p	The proportion of the minority agents	0.15
α_0	The preference value of action x for the minority agents	8
β_0	The preference value of action y for the minority agents	2
α_1	The preference value of action x for the majority agents	5
β_1	The preference value of action y for the majority agents	5
k	The number of peers in coordination games	7
η	The learning rate	0.01
ϵ	The explore rate	0.05

4 Experiments

4.1 Experiment Setup

Following Minority Influence Model, first we define all the parameters in Table 1. In each set of experiments, we may only modify several parameters, and others that are not mentioned remain the default value.

In each set of experiments, agents are divided into two groups, the minority, and the majority. Agents of the minority prefer action x than y, and they always choose action x at the beginning. Initially, both the minority and the majority groups achieve Nash equilibrium, where no agent can gain more by changing its strategy [25]. Then individuals of the two groups begin to interact with each other. After a number of iterations, we observe whether the minority can influence the majority. Since the preference of the majority can be consistent, neutral, or contrary compared to the minority, we consider three scenarios, which are introduced in more details next.

Scenario 1: Consistent Preference Scenario. The first scenario assumes that the agents of the majority share the same preference with the minority but act differently at the beginning. We set $(\alpha_0, \beta_0) = (\alpha_1, \beta_1)$, where $\alpha_0 > \beta_0$, the initial action of the minority to x, the initial action of the majority to y.

Scenario 2: Neutral Majority Scenario. Different from the first scenario, the majority may have a neutral attitude towards the actions, i.e., they prefer the two actions at the same level. Thus, we set $(\alpha_1, \beta_1) = (5, 5)$. There are three possible initial states for the majority. The first is that all agents of the majority choose action y. The second is that half the agents choose action x, and others choose y. The third is that all agents choose action x. All three states are under the Nash equilibrium. As the agents of the minority group choose action x, we only consider the first and the second states.

Scenario 3: Contrary Preference Scenario. This scenario assumes all agents of majority have a contrary preference of action compared to the minority group, and the majority will choose their action according to their own preference at the beginning of the iteration. Under this situation, the minority faces higher difficulty to affect the majority. We'd like to explore the case when the minority group has an extreme belief of their preference.

In order to explore the minority influence under different scenarios, we conduct E simulations for each different set of parameters for each scenario. Part of the experiment results is shown and discussed next.

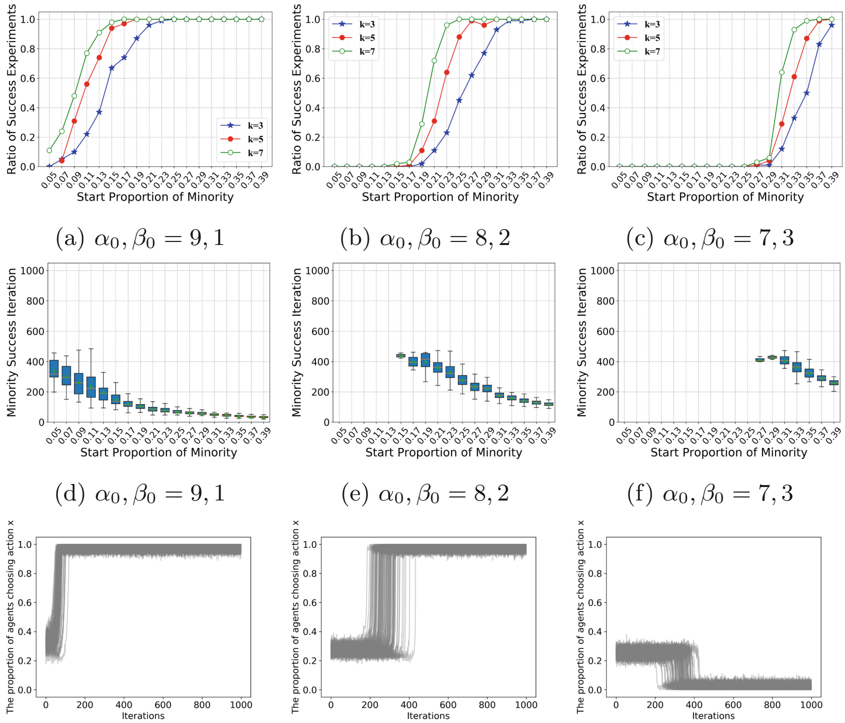

(a) $\alpha_0, \beta_0 = 9, 1$ (b) $\alpha_0, \beta_0 = 8, 2$ (c) $\alpha_0, \beta_0 = 7, 3$

(d) $\alpha_0, \beta_0 = 9, 1$ (e) $\alpha_0, \beta_0 = 8, 2$ (f) $\alpha_0, \beta_0 = 7, 3$

(g) $\alpha_0, \beta_0 = 9, 1; p = 0.25$ (h) $\alpha_0, \beta_0 = 8, 2; p = 0.25$ (i) $\alpha_0, \beta_0 = 7, 3; p = 0.25$

Fig. 1. Scenario 1. Consistent preference scenario

4.2 Scenario 1: Consistent Preference Scenario

The experiment results are displayed in Fig. 1. In Figs. 1a, b and c, the x-axis denotes the proportion of minority to the whole society, the y-axis denotes the proportion of experiments out of n experiments that the minority has successfully affected the society, which means over 90% of agents in the society ends up with choosing action x (p_x) which is preferred by the minority agents. In Figs. 1d, e and f, the x-axis denotes the proportion of minority, and the y-axis denotes the iteration when the minority successfully influence the majority, which also represents the emergence time of the social norm in each independent experiment under the given parameters. The box-plot shows the mean value and variance of all the emergence time of the succeed experiments of the parameter set. It indicates the speed and stability of the emergence of minority influence. In Figs. 1g, h and i, the x-axis denotes the number of iterations, the y-axis denote the proportion of agents choosing action x (p_x) which is preferred by the minority agents, and each line denotes one single repeat experiment. And from left to right in each row of figures, the preference of agents are different. For other scenarios, the figures are organised in the same way.

It can be observed from Fig. 1 that with fixed preference values, as p increases, the ratio of the succeed experiments increases, and the norm emerges faster and more stable. This indicates that the larger the population of the minority, the easier the minority can affect the majority. Meanwhile, a larger k leads to a higher success ratio under the same p. Finally, the minority group with more extreme preference values can be more influential.

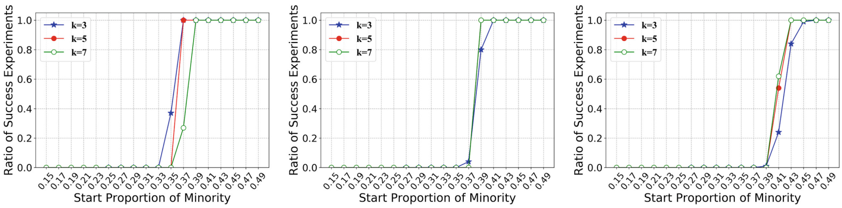

(a) $\alpha_0, \beta_0 = 9, 1$ (b) $\alpha_0, \beta_0 = 8, 2$ (c) $\alpha_0, \beta_0 = 7, 3$

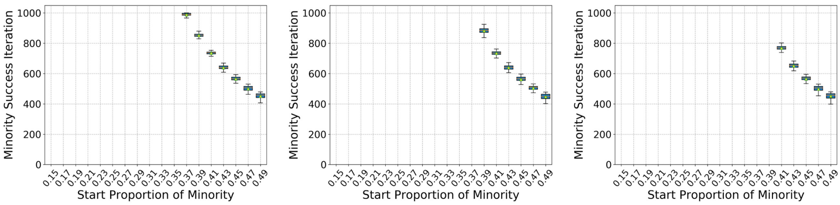

(d) $\alpha_0, \beta_0 = 9, 1$ (e) $\alpha_0, \beta_0 = 8, 2$ (f) $\alpha_0, \beta_0 = 7, 3$

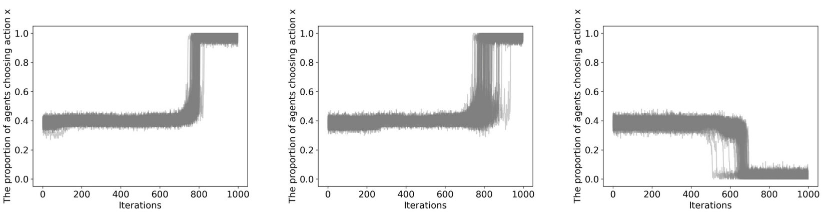

(g) $\alpha_0, \beta_0 = 9, 1; p = 0.39$ (h) $\alpha_0, \beta_0 = 8, 2; p = 0.39$ (i) $\alpha_0, \beta_0 = 7, 3; p = 0.39$

Fig. 2. Scenario 2-1. Neutral majority with contrary initial action

4.3 Scenario 2: Neutral Majority

Scenario 2-1: Neutral Majority with Contrary Initial Action. This scenario can be much harder, compared to the case above, for the minority to influence the majority, i.e., it requires a larger p and more iterations (Fig. 2). Consistent with Scenario 1, a higher p and a larger ratio between α_0 and β_0 can also lead to faster norm emergence. Moreover, we can observe some tipping points. For example, when we set α_0 to 9, β_0 to 1, k to 5, we can see the tipping point $p = 0.36$. This is consistent with the defection effect [13,19], which indicates that when the minority affect a few individuals of the majority, the rest of the majority can soon follow these defectors and trigger minority influence. The experimental results suggest that once p reaches the tipping point, minority influence can occur, and a social norm can emerge.

Scenario 2-2: Neutral Majority with Random Initial Action. Figure 3 shows that a relatively small minority group with extreme preference can lead the population to conform to a new social norm. This may indicate that in a society that the majority have no certain preference, a minority group with a small population can trigger the minority influence. In this scenario, we do not observe any significant correlation between k, (α_0, β_0) and the emergence of minority influence.

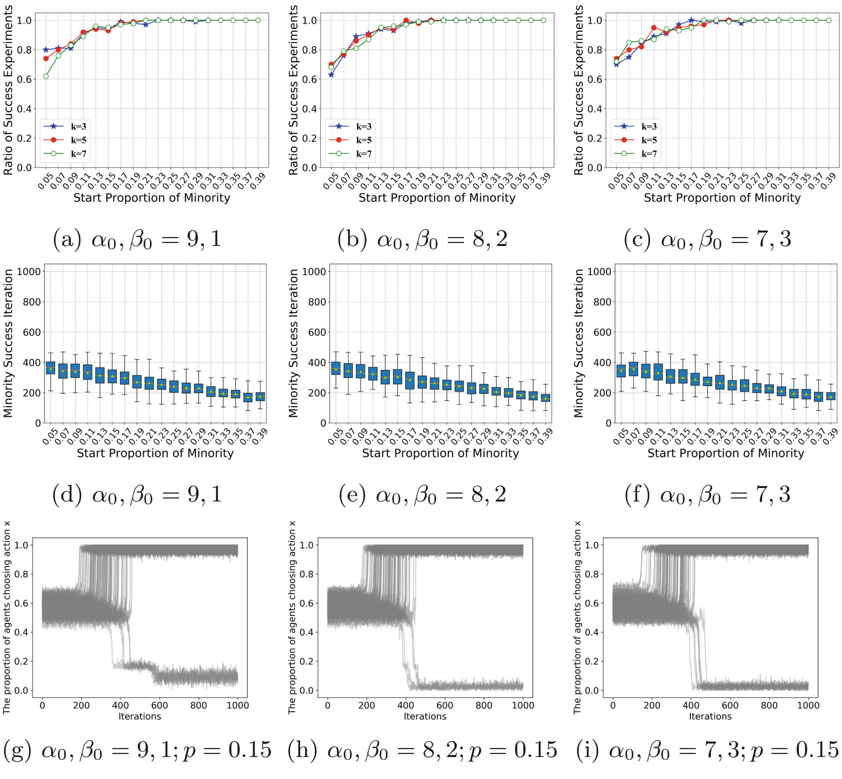

(a) $\alpha_0, \beta_0 = 9, 1$ (b) $\alpha_0, \beta_0 = 8, 2$ (c) $\alpha_0, \beta_0 = 7, 3$

(d) $\alpha_0, \beta_0 = 9, 1$ (e) $\alpha_0, \beta_0 = 8, 2$ (f) $\alpha_0, \beta_0 = 7, 3$

(g) $\alpha_0, \beta_0 = 9, 1; p = 0.15$ (h) $\alpha_0, \beta_0 = 8, 2; p = 0.15$ (i) $\alpha_0, \beta_0 = 7, 3; p = 0.15$

Fig. 3. Scenario 2-2. Neutral majority with random initial action

4.4 Scenario 3: Contrary Preference Scenario

Unlike other scenarios, we only analyse experiments with parameters $\alpha_0, \beta_0 = 9, 1$, $k = 7$. This is because it can be tough for the minority with gentle preference to affect the majority under this scenario. The experimental results show that minority influence occurs within 1000 iterations when p is larger than 0.45. Consistent with other scenarios, the larger the value of p, the faster the norm emerges. When p is less than 20%, we cannot observe minority influence even after 6000 iterations. This may indicate that it can be very hard for the minority to influence the majority with an opposing attitude, even when the minority is extreme and with a moderate population.

5 Conclusion and Future Works

In this paper, we proposed a novel framework, Minority Influence Model, to explore minority influence in the context of norm emergence. The model considers multiple agent interactions, where agents perform actions based on both their preferences and perceived social norm, which is reflected by gist. Based on the proposed model, extensive experiments have been done. The experimental results suggest minority influence can be correlated to the initial proportion of the minority, the extent to which the minority prefers some action, and the degree of disagreement between the minority and the majority. Furthermore, some tipping points have been observed under some scenarios. In our future work, we will extend the Minority Influence Model to structured networks [11]. We will also enrich our model to explore minority influence in more complex situations where agents can have different social positions.

References

1. Ajzen, I., Fishbein, M.: The influence of attitudes on behavior. In: The Handbook of Attitudes, vol. 173, no. 221, p. 31 (2005)
2. Brainerd, C.J., Reyna, V.F.: Gist is the grist: fuzzy-trace theory and the new intuitionism. Dev. Rev. **10**(1), 3–47 (1990)
3. Cai, Y., Zheng, H., Liu, J., Yan, B., Su, H., Liu, Y.: Balancing the pain and gain of hobnobbing: utility-based network building over attributed social networks. In: Proceedings of the 17th AAMAS, pp. 193–201 (2018)
4. Claus, C., Boutilier, C.: The dynamics of reinforcement learning in cooperative multiagent systems. In: AAAI/IAAI 1998, pp. 746–752 (1998)
5. Easley, D., Kleinberg, J., et al.: Networks, Crowds, and Markets, vol. 8. Cambridge University Press, Cambridge (2010)
6. Gardikiotis, A.: Minority influence. Soc. Pers. Psychol. Compass **5**(9), 679–693 (2011)
7. Hao, J., Leung, H.f.: The dynamics of reinforcement social learning in cooperative multiagent systems. In: Twenty-Third International Joint Conference on Artificial Intelligence (2013)

8. Hu, S., Leung, C.w., Leung, H.f., Liu, J.: 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 (2019)
9. Hu, S., Leung, H.: Compromise as a way to promote convention emergence and to reduce social unfairness in multi-agent systems. In: Mitrovic, T., Xue, B., Li, X. (eds.) AI 2018. LNCS (LNAI), vol. 11320, pp. 3–15. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-03991-2_1
10. Kaelbling, L.P., Littman, M.L., Moore, A.W.: Reinforcement learning: a survey. J. Artif. Intell. Res. **4**, 237–285 (1996)
11. Khoussainov, B., Liu, J., Khaliq, I.: A dynamic algorithm for reachability games played on trees. In: Královič, R., Niwiński, D. (eds.) MFCS 2009. LNCS, vol. 5734, pp. 477–488. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-03816-7_41
12. Kittock, J.E.: Emergent conventions and the structure of multi-agent systems. In: Proceedings of the 1993 Santa Fe Institute Complex Systems Summer School, vol. 6, pp. 1–14. Citeseer (1993)
13. Levine, J.M.: Reaction to opinion deviance in small groups. Psychol. Group Influ. **2**, 187–231 (1989)
14. Liu, J., Li, L., Russell, K.: What becomes of the broken hearted?: an agent-based approach to self-evaluation, interpersonal loss, and suicide ideation. In: Proceedings of the 16th AAMAS, pp. 436–445 (2017)
15. Liu, J., Wei, Z.: Network, popularity and social cohesion: a game-theoretic approach. In: Thirty-First AAAI Conference on Artificial Intelligence (2017)
16. Milgram, S., Bickman, L., Berkowitz, L.: Note on the drawing power of crowds of different size. J. Pers. Soc. Psychol. **13**(2), 79 (1969)
17. Moscovici, S.: Innovation and minority influence. In: Perspectives on minority influence, pp. 9–52 (1985)
18. Moskvina, A., Liu, J.: Togetherness: an algorithmic approach to network integration. In: ASONAM 2016, pp. 223–230. IEEE (2016)
19. Myers, D.G.: Social Psychology, 10th edn. McGraw Hill, New York (2010)
20. Nail, P.R., MacDonald, G., Levy, D.A.: Proposal of a four-dimensional model of social response. Psychol. Bull. **126**(3), 454 (2000)
21. Nemeth, C.J.: Differential contributions of majority and minority influence. Psychol. Rev. **93**(1), 23 (1986)
22. Sen, S., Airiau, S.: Emergence of norms through social learning. In: IJCAI, vol. 1507, p. 1512 (2007)
23. Shoham, Y., Tennenholtz, M.: Emergent conventions in multi-agent systems: Initial experimental results and observations. In: Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning, pp. 225–231 (1992)
24. Shoham, Y., Tennenholtz, M.: On the emergence of social conventions: modeling, analysis, and simulations. Artif. Intell. **94**(1–2), 139–166 (1997)
25. Shubik, M.: Game Theory in the Social Sciences: Concepts and Solutions, vol. 155. MIT Press Cambridge, Cambridge (1982)
26. Watkins, C.J., Dayan, P.: Q-learning. Mach. Learn. **8**(3–4), 279–292 (1992)

27. Wood, W., Lundgren, S., Ouellette, J.A., Busceme, S., Blackstone, T.: Minority influence: a meta-analytic review of social influence processes. Psychol. Bull. **115**(3), 323 (1994)
28. Yan, B., Liu, Y., Liu, J., Cai, Y., Su, H., Zheng, H.: From the periphery to the center: information brokerage in an evolving network. In: Proceedings of the 27th IJCAI, pp. 3912–3918 (2018)