This article has been accepted for publication in IEEE Transactions on Knowledge and Data Engineering. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TKDE.2023.3289848

Interpretable Stochastic Block Influence Model: Measuring Social Influence Among Homophilous Communities

Yan Leng[1,2], Tara Sowrirajan[2,3,4], Yujia Zhai[5], Alex Pentland[2]
[1] McCombs School of Business, The University of Texas at Austin
[2] MIT Media Lab [3] Harvard University [4] Kellogg School of Management [5] Tianjin Normal University

Abstract—Decision-making on networks can be explained by both homophily and social influences. While homophily drives the formation of communities with similar characteristics, social influences occur both within and between communities. Social influences can be reasoned through role theory, which indicates that the influences among individuals depending on their roles and the behavior of interest. To operationalize these social science theories, we empirically identify the homophilous communities and use the community structures to capture such "roles", affecting particular decision-making processes. We propose a generative model named the Stochastic Block influences Model and jointly analyzed both network formation and behavioral influences within and between different empirically-identified communities. To evaluate the performance and demonstrate the interpretability of our method, we study the adoption decisions for a microfinance product in Indian villages. We show that although individuals tend to form links within communities, there are strongly positive and negative social influences between communities, supporting the weak ties theory. Moreover, communities with shared characteristics are associated with positive influences. In contrast, communities that do not overlap are associated with negative influences. Our framework facilitates the quantification of the influences underlying decision communities and is thus a helpful tool for driving information diffusion, viral marketing, and technology adoption.

Index Terms—Social influence; Homophily; Stochastic Block Model; Community structure; Generative model

1 Introduction

We are living in an increasingly connected society [1], [2], [3]. The connections among individuals foster information diffusion and enable inter-dependencies in decisions among peers. Therefore, understanding and modeling how hidden social influences change individuals' decisions are essential and critical to many practical applications, such as viral marketing, political campaigns, and large-scale behavioral change [4], [5].

Homophily, the tendency of similar individuals to associate together, widely exists in various types of social networks and controls the outcomes of many critical network-based phenomena [6]. Salient features for homophily come from a wide range of sources, including age, race, socioeconomic status, occupation, and gender [6]. The complex nature of social relationships and the high-dimensional characteristics of individuals thus determine the multi-dimensionality of homophily [7]. Homophily can lead to locally clustered communities and may affect network dynamics, such as information diffusion and product adoption. The block model has been applied to low-dimensional, pre-defined homophilous features and provides a building block to uncover underlying community structures[1] with high-dimensional homophily empirically [8].

Social influences are widely studied in both economics and computer science literature due to their importance in understanding human behavior. In economics, researchers focus on causally disentangling social influences from homophily with randomization strategies, such as propensity score matching, behavioral matching and regression adjustment. In the computer science literature, researchers focus on maximizing the likelihood of the diffusion path of influences by proposing different generative processes. These works focus on the strength or the pathways of social influence, and they do not link social influences to the underlying homophilous communities and the network formation process.

Two theories explain how local communities affect information diffusion and contagion in decision-making. On the one hand, homophily and the requirement of social reinforcement for behavioral adoption in complex contagion theory indicate that influences tend to be localized in homophilous communities [6], [9]. In other words, behavioral diffusion and network formation are endogenous, which helps explain the phenomenon of within-community spreading [10]. On the other hand, the weak ties theory [11] implies that bridging ties between communities facilitates the spread of novel ideas. Empirically, it has been shown that reinforcement from multiple communities, rather than from the same communities, predicts a higher adoption rate [12]. Motivated by these competing theories, we seek whether social influences spread locally within each homophilous community or globally to other communities by taking advantage of long ties.

According to the role theory, "the division of labor in society takes the form of interaction among heterogeneous specialized positions" [13]. That is to say, depending on the social roles and the behavior of interest, the underlying interactions and norms for decision-making are different. Motivated by this proposition, we develop a method to associate social influences with the underlying communities, which are associated with the behavior of interest. To formalize this idea, we propose a generative model to understand how social influences impact decision-making by

1. In this paper, we use community and block interchangeably.

inferring the spread of influences across empirically-identified blocks. Using our framework, we uncover the underlying blocks and infer two types of relationships across these blocks: social interaction and social influence. Unlike the Stochastic Block Model, the observed individual decisions are used to inform the communities, complementing the observed network. In addition, we infer an influences matrix consisting of the social influences across different communities. This influences matrix reveals hidden social influences at the community level, which would otherwise be impossible to observe and generalize.

As a case study, we experiment on the diffusion of microfinance in Indian villages and perform extensive analysis on the influences matrices estimated from the model. Even though social relationships are denser within communities, social influences mainly spread across communities. This may be due to the importance of across community weak ties [11], and the strength of structural diversity [12]. Our generative framework and subsequent understanding of how social influences operate have practical applications, such as viral marketing, political campaigns, and large-scale health-related behavioral change.

Our paper makes the following contributions to the literature:

- SBIM integrates networks, decisions, and characteristics into the generative process. It jointly infers two types of relationships among empirically-identified communities: social connections and social influences. It can flexibly accommodates positive and negative social influences, which cannot be obtained from SBMs.
- Our SBIM is motivated by role theory, which posits that individuals make decisions depending on the context of the decision type [13] (e.g., adopting microfinance as opposed to adopting healthy habits). To achieve this, we allow the underlying community to vary with the behavior of interest.
- We perform a case study on the adoption of microfinance products in Indian villages. We demonstrate the interpretability of our SBIM with detailed analyses of the influences structure.

2 Related Literature

There are two prominent theories for explaining social influence propagation: simple contagion and complex contagion. According to simple contagion theory, individuals will adopt the behavior as long as they have been exposed to the information [11], which is a sensible model for epidemics and information spreading. Complex contagion theory, on the other hand, requires social reinforcement from neighbors to trigger adoption [9]. Many studies have shown that complex contagion explains behaviors such as registration for health forums [14]. However, these exposure-based models are analytically simple and do not allow social influences to be negative, i.e., the adoption decision of one's neighbors might decrease, rather than increase, the likelihood of one's adoption decision. Moreover, they typically are not able to capture the heterogeneity of social influences [5]. In this paper, we propose a model to account for negative and heterogeneous influences.

The Stochastic Block Model (SBM) is a statistical model for studying latent cluster structures in the network data [8]. The SBM generalizes the Erdos-Renyi random graph model with higher intra-cluster and lower inter-cluster probability. The traditional SBM only infers community structures from network connections. However, when contextual information on nodes is available, leveraging information from different sources facilitates inference. There has been interesting work on utilizing covariates to infer underlying block structures in recent statistics literature. For example, Binkiewicz et al. [15] present a covariate-regularized community detection method to find highly connected communities with relatively homogeneous covariates. They balance the two objectives (i.e., the node covariance matrix and the regularized graph laplacian) with tuned hyper-parameters. Yan et al. [16] propose a penalized optimization framework by adding a k-means type regularization. This framework is based on the premise that the estimated communities are consistent with latent membership in the covariate space. Although these variations to the SBM utilize auxiliary information on individual nodes, they specify the importance of recovering the network and the smoothness of covariates on the network on an ad hoc basis. Different from these models, we take advantage of role theory [13] and utilize the decision-making process on the network that could also inform community detection. More importantly, the communities that we discover are specifically relevant for the decision-making of interest, while the ones discovered by SBM is only affected by the network connects (and agnostic to decision-making). For example, professional communities are more useful for the adoption of technologies at work, while social communities are more useful for the adoption of social mobile apps. The underlying communities depend on the role and behavior of interest because social influences spread through network links in different applications. SBIM bridges the rich SBM and social contagion literature. It opens up future opportunities to adapt to other variations of SBM.

3 Methodology

3.1 Stochastic Block influences Model

Assume a random graph $G(\mathcal{V}, \mathcal{E})$ with N individuals in node set $\mathcal{V}$ and edge set $\mathcal{E}$. It is partitioned into C disjoint blocks $(\mathcal{V}_1, ..., \mathcal{V}_C)$, and the proportion of nodes in each block c is π_c, and $\sum_{c=1}^{C} \pi_c = 1$. $\mathbf{A} \in \mathbb{R}^{N \times N}$ represents the adjacency matrix. $\mathbf{A}_{ij} = 1$ if i and j are connected, and $\mathbf{A}_{ij} = 0$ otherwise. Let matrix $\mathbf{B} \in \mathbb{R}^{C \times C}$ denote block-to-block connection likelihood. Let $\mathbf{M}_i$ be the block assignment of individual i, and by summing over all C blocks, we have $\sum_{k=1}^{C} \mathbf{M}_{ik} = 1$. We combine the block vector of all individuals in the matrix $\mathbf{M} \in \mathbb{R}^{N \times C}$. Therefore, the probability of a link between v_i and v_j between two separate blocks $\mathcal{V}_k$ and $\mathcal{V}_l$ as $P\big((v_i, v_j) \in \mathcal{E} | v_i \in \mathcal{V}_k, v_j \in \mathcal{V}_l\big) = p_{ij}$. $\mathbf{y} \in \mathbb{R}^N$ is a binary vector representing individuals' adoption behaviors. Let $\mathbf{X}_i \in \mathbb{R}^D$ represent demographic features, where D is the number of covariates. We use $\mathbf{F} \in \mathbb{R}^{C \times C}$ to represent the block-to-block influences matrix. Finally, $\mathbf{h}$ is a binary vector, that captures whether or not each individual is aware of the product at the beginning of the observation period. For a new product, $\mathbf{h}$ is sparse, while for a mature product, $\mathbf{h}$ is dense.

In SBIM, we link latent communities to adoption decisions and socio-demographics. SBIM reveals the underlying nature of high-dimensional homophily in a data-driven fashion rather than using pre-defined communities using observed sociodemographics (e.g., race or occupation). Solely using pre-defined homophilous characteristics does not aptly capture the multi-faceted characteristics that define individuals and their social ties. In other words, individuals are associated with different communities, each formed by various homophilous characteristics. Neighbors belonging to different communities may influence focal individuals differently.

We illustrate this using the adoption of microfinance in Indian villages. We posit that several traits define the diverse nature of individuals: different professions, castes, education levels, and a variety of other demographic features. One individual, an educated worker of a lower social caste, belongs with varying degrees of

This article has been accepted for publication in IEEE Transactions on Knowledge and Data Engineering. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TKDE.2023.3289848

affiliation to different communities. The individual is perhaps most strongly affiliated with a group that has a certain level of education and less strongly affiliated with another group where most of the group are of a lower caste. This mixed membership captures the realistic nature of our social relationships and characteristics. Within such a village with multi-dimensional homophily, how can we understand who influences this individual and what processes are involved in that individual's decision-making? Specifically, this individual could be influenced by neighbors belonging to different communities characterized by specific educational backgrounds, professions, and castes. The data-driven multi-dimensional blocks of the model allows us to capture these critical, hidden relationships.

Next, we formalize our SBIM model. To jointly infer how influences spread within and across communities, we create a model with the following properties:

1) The model leverages both the observed friendship network structure and adoption behavior to infer the underlying communities.
2) The link formation and social influences between two individuals are jointly determined by their underlying communities.

Blocks for network connections. For each individual pair $\{i, j\}$, depending on their community assignment vectors, the predicted link $\hat{\mathbf{A}}_{ij}$ is generated according to the connection probability matrix, $\mathbf{B}$. In particular, the probability of the existence of a link between i and j is:

$$\mathbb{P}(\hat{\mathbf{A}}_{ij} = 1|\mathbf{M}, \mathbf{B}) = (\mathbf{M}\mathbf{B}\mathbf{M}^T)_{ij}. \tag{1}$$

Blocks for decision-making. Next, we discuss how we incorporate individual characteristics and adoption decisions into our SBIM. The adoption likelihood depends on individuals' characteristics and on the influences of their neighbors who have already adopted [17]. This generative model builds upon the communities an individual i and i's neighbors belong to, as well as the community-to-community influences matrix $\mathbf{F}_{ij}$. Each individual decides on whether or not to adopt to maximize her utility. The utility of i depends on what this individual prefers and the aggregated influences of neighbors. The pairwise influences depend on the communities that i and her neighbors belong to. We illustrate how influences and communities affect one's decision-making in Figure 1. We consider individual A, who has three friends, B, C, and D, belonging to a lower socioeconomic status (SES) group (as colored in red), and one friend, E, belonging to a higher SES group (as colored in blue). The adoption likelihood of individual A is a function of this individual's preferences as well as the influences of friends B, C, D, and E. The strength of the influences depends on the corresponding communities of A and her friends.

More generally, a user's adoption likelihood, $\hat{\mathbf{y}}$, is defined as:

$$\hat{y}_i = \text{logit}\big(\boldsymbol{\beta}\mathbf{X}_i + \sum_j \big((\mathbf{M}\mathbf{F}\mathbf{M}^T) \circ ((\mathbf{h} \otimes \mathbf{1}) \circ \mathbf{A})\big)_{ji} + \epsilon_i\big), \tag{2}$$

where $\circ$ is the element-wise matrix multiplication. The first term, $\boldsymbol{\beta}\mathbf{X}_i$, measures the adoption decision conditioned on i's sociodemographic features if there were no social influence, where $\boldsymbol{\beta} \in \mathbb{R}^D$ and D is the dimension of the covariates. The second term aggregates the influences of i's neighbors. ϵ_i is the idiosyncratic error term. Without loss of generality, we assume $\epsilon_i \sim \mathcal{N}(0, 1)$.

For a mature product that everyone knows, we can simplify Eq. (2) as:

$$\hat{y}_i = \text{logit}\Big(\boldsymbol{\beta}\mathbf{X}_i + \sum_{j=1}^{N} \big((\mathbf{M}\mathbf{F}\mathbf{M}^T) \circ \mathbf{A}\big)_{ji} + \epsilon_i\Big). \tag{3}$$

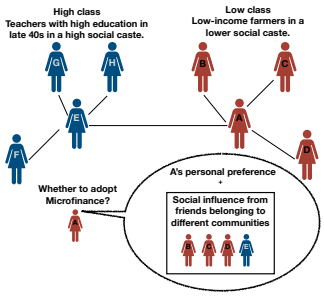

Fig. 1: **Graphical representation of the Stochastic Block influences Model (SBIM).** Assume there are two communities, a high socioeconomic status (SES) group (dark blue) and low SES group (dark red), characterized by multi-dimensional sociodemographic features. The two groups have higher intra SES group connection probability and lower inter SES group connection probabilities. The decision-making of individual A is jointly influenced by this individual's preferences, as well as her friends from the same and different communities.

Eq. (2) only accounts for the influences among direct neighbors. Note that in a small-scale network, it is reasonable to assume higher-order social influences do not exist. In a large-scale network, Leng et al. [18] show that social influences spread beyond immediate neighbors. Our model can be easily adapted to higher-order influences by summing up the powers of the adjacency matrix A to account for multiple degrees of separation, as done in [5].

We finally introduce the loss function ($\mathcal{L}$):

$$\mathcal{L} = -\sum \mathbf{y}\log(\hat{\mathbf{y}}) - \sum \mathbf{A}\log(\hat{\mathbf{A}}). \tag{4}$$

The first component is the main objective for typically SBM and the second component is SBIM's main advantage. The same inferred blocks $\mathbf{M}$ are used to compute both $\hat{\mathbf{A}}$ (Eq. (1)) and $\hat{\mathbf{y}}$ (Eqs. (2) and (3)). In the loss function, we minimize the difference between (1) predicted and observed links (the objective for existing SBM in "Blocks for network connections"); (2) predicted and observed behaviors (contribution of our paper discussed in "Blocks for decision-making").

3.2 Generative process

For the full network, we assume the following generative process in the model, which defines a joint probability distribution over N individuals, based on the node-wise membership matrix $\mathbf{M}$, block-to-block interaction matrix $\mathbf{B}$, block-to-block influences matrix $\mathbf{F}$, attributes' coefficients $\boldsymbol{\beta}$, observed friendship network $\mathbf{A}$, observed attributes $\mathbf{X}$, and observed adoption decision $\mathbf{y}$. For brevity, we denote $\mathcal{Z}$ as set of the hidden variables, $\mathcal{Z} = \{\mathbf{M}, \boldsymbol{\beta}, \mathbf{B}, \mathbf{F}\}$ and θ as the set of hyperparameters, where $\theta = \{c, a, b, \mu, \sigma, \mu_b, \sigma_b\}$.

(1) For each node $v_i \in \mathcal{V}$, draw a C-dimensional mixed membership vector $\mathbf{M}_i \sim \text{Dirichlet}(c)$.
(2) For the connection probability from community k to l in the block-to-block connectivity matrix, draw $\mathbf{B}_{kl} \sim \text{Beta}(a, b)$.
(3) For the influences from community k to l in the block-to-block influences matrix, draw $\mathbf{F}_{kl} \sim \mathcal{N}(\mu_F, \sigma_F)$.
(4) For each attribute in $\boldsymbol{\beta}_d$, draw the coefficient $\boldsymbol{\beta}_d \sim \mathcal{N}(\mu_b, \sigma_b)$.
(5) Draw the connection between each pair of nodes v_i and v_j, $\hat{\mathbf{A}}_{ij}$, according to Eq. (1).
(6) Draw the adoption decision $\hat{\mathbf{y}}_i$, according to Eq. (2).

Steps (3), (4), and (6) are unique processes of SBIM, which are relevant to the adoption behavior of our method. The blocks used in SBM in Step (1) are also used in the generation process

in Eqs. (2)-(3) in Step (6). As discussed in Section 3, the latent component M contributes both to $\hat{\mathbf{A}}$ (Eq. (1)) and $\hat{y}$ (Eqs. (2) and (3)), as it affects both the decision-making and the link formation.

The posterior distribution defined by SBIM is a conditional distribution of the hidden block structure and relationships given the observed friendship network and adoption behavior, which decomposes the agents into C overlapping blocks. The posterior will place a higher probability on configurations of the community membership that describe densely connected communities and stronger (positive or negative) influences.

The posterior of the SBIM is intractable, similar to many hierarchical Bayesian models. Therefore, we use the Markov Chain Monte Carlo (MCMC) algorithm as an approximate statistical inference method to estimate the parameters. The MCMC draws correlated samples that converge to the target distribution and are generally asymptotically unbiased. MCMC methods include Gibbs sampling, Metropolis-Hastings, Hamiltonian Monte Carlo (HMC), and No-U-Turn Sampler (NUTS). Gibbs sampling and Metropolis-Hastings methods converge slowly to the target distribution as they explore the parameter space by random walk. HMC suppresses the random walk behaviors with an auxiliary variable that transforms the problem by sampling to a target distribution using simulated Hamiltonian dynamics. However, HMC requires the gradient of the log-posterior, which is complicated in our model. Moreover, it requires a reasonable specification of the step size and several steps, which would otherwise result in a substantial drop in efficiency [19]. Therefore, we apply the NUTS, a variant of the HMC method, to eliminate the need for choosing the number of steps by automatically adapting the step size. Specifically, the NUTS builds a set of candidate points that spans the target distribution recursively and automatically stops when it starts to double back and retrace its steps [19].

4 Experiments

We study the adoption of microfinance in five villages in India collected by the Abdul Latif Jameel Poverty Action Lab (J-PAL) [20]. In 2007, a microfinance institution introduced a microfinance program to some selected villages in India. In early 2011, they collected information about whether or not the villagers had adopted the microfinance product. Because the villages are fairly small and microfinance had been on the market for four years when JPAL collected individuals' adoption decisions, it is reasonable to assume that everyone in the village was aware of microfinance, which is hence a considered mature product in this village. Therefore, we employ Eq. (3) as the decision-making function. The data contains information about self-reported relationships among households and other amenities, including village size, quality of access to electricity, quality of latrines, the number of beds, the number of rooms, the number of beds per capita, and the number of rooms per capita. These demographic features are used as the independent variables. The outcome variable is the adoption decision of microfinance. The microfinance institution asked the villagers to self-report other villagers they considered as friends.

To train our model and evaluate the performance for a particular C, the number of blocks, we cross-validated by randomly splitting the data into 75% training samples and 25% test samples. We repeat this process ten times. With NUTS, we obtain the point estimates for all latent variables in $\mathcal{Z}$. We then rerun our model (as previously described) with all latent variables fixed to the estimates on the test dataset. This step returns the predicted adoption probability for each villager in the test data. To choose the optimal number of blocks, we first tune the model for $C \in \{2, 6, 10, 14\}$ and then calculate the average loss. We observe a negative parabolic trend with the loss peaking at its lowest at $C = 10$ blocks, so we use this optimal number of blocks in additional analyses. We use a machine learning model with sociodemographics and the hidden community learned by spectral clustering[2] (i.e., blocks) on the adjacency matrix as the independent variables. In this way, we use the same information in SBIM and the benchmarks.

Since the dependent variable in our data is imbalanced, we evaluate our SBIM using the area under the Receiver-Operating-Characteristics curve (AUC) that we plot using the false positive rate and correct positive rate for different thresholds. We define a loss metric to select the best configurations during the training period. We formulate it by taking the negative of the standard improvement measure, which is the absolute improvement in performance normalized by the room for improvement. This measure captures the improvement of our SBIMcompared to the baseline model. Since we have a small test set, making predictions on a randomly-drawn test set is hard. Measuring the relative improvement ensures that the composition of the test set does not bias the performance due to sample variation. This metric is formulated by: $L = \frac{\text{Baseline}_{\text{test AUC}} - \text{SBIM}_{\text{test AUC}}}{1 - \text{Baseline}_{\text{test AUC}}}$, where the AUC of the baseline model and SBIM on the test split in cross-validation are represented as $\text{Baseline}_{\text{test AUC}}$ and $\text{SBIM}_{\text{test AUC}}$, respectively.

SBIM has seven hyperparameters in θ.Since the parameter space is large, we adapt a bandit-based approach to tune the parameters developed called hyperband [22]. Our adaptation of this algorithm allows each configuration tested to run with full resources due to our sampling procedure, allowing NUTS to run consistently across all configurations.

We compare the performance of our SBIM model with six methods benchmarks, all of which use the learned block as the extra feature, in Table 1. We find that our SBIM outperforms the best benchmark (elastic net with blocks) in the test set by 13.4% using the improvement metric mentioned above.

TABLE 1: Model and baseline performance

	Train: Mean (S.D.)	Test: Mean (S.D.)
Random forest (with blocks)	0.901 (0.010)	0.610 (0.095)
Gradient boosting (with blocks)	0.843 (0.075)	0.531 (0.058)
Adaboost (with blocks)	0.876 (0.038)	0.528 (0.058)
Elastic net (with blocks)	0.724 (0.107)	0.612 (0.079)
Lasso regression (with blocks)	0.719 (0.069)	0.607 (0.078)
Multilayer perceptron (with blocks)	0.703 (0.098)	0.536 (0.056)
SBIM	0.805 (0.022)	**0.664** (0.062)

5 Analysis and Discussions

Interpretability is a broad term in machine learning. We follow the definition of [23]. We define interpretable machine learning as the extraction of relevant knowledge from a machine-learning model concerning relationships either contained in the data or learned by the model. Our SBIM model satisfies several characteristics of the model-based interpretability methods developed in this paper, including sparsity, simulability, and modularity. We discuss how our model satisfies each criterion in the Appendix.

We can associate individuals' sociodemographics with the individuals who belong to each block to generalize block types as consisting of characteristics such as high or low SES, homogeneous

2. Spectral clustering uses the second smallest eigenvector of the graph laplacian as the semi-optimal partition [21].

or diverse, and skilled or less educated, as depicted in Table 2 (remaining examples are shown in Table A1 in the Appendix). In this example, each block is associated with a qualitative type, and the attributes within that block describe such characterizations. Caste composition, education levels, and profession types are employed to designate lower or higher SES blocks. Homogeneous or diverse blocks are designated by some professional composition, caste types, native language composition, gender imbalance, and what fraction of village inhabitants are natives. We also display the

We use normalized entropy to measure the diversity of attributes studied in this paper. Normalized entropy is a metric used to capture the number of types of characteristics within each category while accounting for the frequency of each entity type within a category. It can be formulated by, $Q = -\frac{\sum_{i=1}^{q} p_i \log(p_i)}{\sum_{i=1}^{q} \frac{1}{n_i} \log(\frac{1}{n_i})}$, where q denotes the number of types within a category, p_i denotes the probability of each type i, f denotes the number of each type n_i.

The gender ratio (R) is measured within a block and is formulated by $R = \frac{r_m}{r_f}$, where r_m and r_f denote the number of occurrences of males and females respectively. Thus, since R denotes the ratio of males to females in a block, both a high or low gender ratio corresponds to a high gender imbalance.

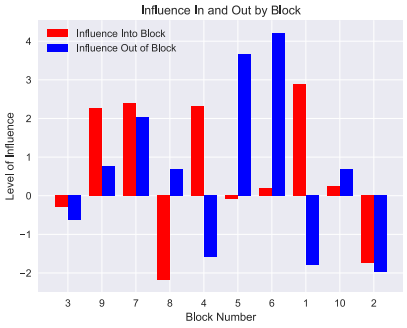

Fig. 2: **Net influences into and out of each block.**

The total influences into and out of each block are depicted in Figure 2, which allows us to evaluate the aggregated influences a block receives and spreads (net positive, negative, or neutral). For example, we can see diverse, low-SES block five and senior, low-SES block six have high output levels of positive influence, and diverse, middle-SES block eight receives a net high level of negative influence. We find that some blocks have stronger outgoing influences than others and can perceive these as positive and negative influences leaders. Similar reasoning applies to characterize blocks that receive a high level of influence as follower blocks. We also observe the difference in net incoming and outgoing influences within each block relating to its role in the block-to-block network. We refer to this to interpret different dynamics between social blocks and then pairing this information with demographics to make further evaluations about the block characteristics associated with different types of influence.

In Figure 3, a subset of the sociodemographic features is displayed for each block, where the network of blocks is connected with varying degrees of influences between them. We find that the equal gender ratio block ten positively influences the similarly equal gender ratio block nine. Block ten positively influences block nine, with similarly high caste diversity. Lower professionally diverse block one negatively influences higher professionally diverse block three. Due to space limit, we display other attributes in Figure A1 in the Appendix.

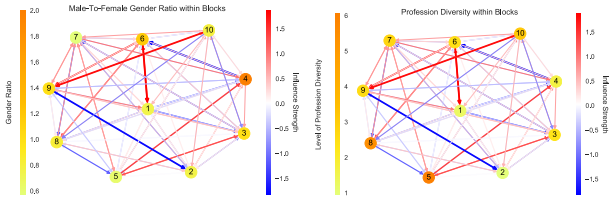

Fig. 3: **Social influences across social blocks (directed links) for gender ratio and profession diversity (node color).**

TABLE 2: **Block characteristics example.** SES shorts for socioeconomic status. Higher education refers to having education levels at PUC (pre-university course) and having a "degree or above" designation.

Block	Block Type	Attributes
1	Homogeneous, low-SES	Only one disadvantaged caste and one language spoken
		Low profession diversity and education levels
2	Diverse, skilled, highly-educated	Several different castes from many levels
		Diverse languages and diverse, high-skilled professions
3	Senior, low-SES	Majority disadvantaged caste
		Majority low skill-level professions in agriculture
4	Young, low-SES	Younger average age, gender imbalanced block
		Majority lowest caste members, mostly natives
		Higher education
5	Diverse, low-SES	Diverse number of disadvantaged castes
		Moderate language diversity, moderate education

By analyzing several examples in this manner using block characteristic composition and observing the types and patterns of influence, several general trends arise, as depicted in Table 3. The block attributes most frequently associated with different types of influence are summarized into key trends. Positive influence occurs when two blocks overlap in the following characteristics: gender distribution, majority castes and professions. Negative influence frequently occurs when two blocks have a lack of overlap in the following characteristics: gender distribution, caste composition, and profession diversity level. Furthermore, the direction of negative influence is most frequently observed from a low-SES block to a high-SES block. Additionally, we frequently observe positive self-influence, which is from a block to itself, and this occurs when a block is characterized by a younger average age, highly-educated, high job diversity, higher-skilled jobs, high language diversity, large gender imbalance, and having a large number of village natives. The remaining examples are displayed in Table A2.

TABLE 3: **Block attributes (including gender, caste, and profession) associated with different types of influence.**

Attribute	Positive influence	Negative influence	Positive self-influence
Gender	Similar gender distribution	Gender-imbalanced block is more open to negative influences	Large gender imbalance
Caste	Overlapping majority castes	Lack of overlap in caste composition	Majority village natives
Profession	Profession overlap, in specialty jobs specifically; large professions diversity	Professionally diverse block receives negative influences from a less professionally diverse block; lack of professional overlap causes a negative influences	High job diversity and higher-skilled jobs

When we analyze several examples using block characteristic composition and observing the types and patterns of influence, several general trends arise, as depicted in Table A2. The block attributes most frequently associated with different influences are summarized into key trends. We find a positive influence when two blocks overlap in the following characteristics: gender distribution, majority castes, professions, high profession diversity, highly educated, highly-skilled jobs, and native languages. Negative influences frequently occur when two blocks do not overlap in the following characteristics: gender distribution, caste composition, profession diversity level, education levels, and average age. Furthermore, negative influences are most frequently observed

from a low-SES block to a high-SES block. Additionally, we frequently observe positive self-influence, which is from a block to itself, and this occurs when a block is characterized by a younger average age, highly-educated, high job diversity, higher-skilled jobs, high language diversity, large gender imbalance, and having a large number of village natives.

When paired with block type characterizations, we find that these trends lead to interesting associations, such as block-to-block perceptions of lower or higher SES groups having influence. Blocks with the higher SES group designation more frequently receive negative influences from lower-SES blocks. Blocks of similar SES, especially higher SES, have more frequent positive influences, and High-SES blocks also have more frequent positive self-influence.

These findings suggest that firms employ marketing strategies that take into account the underlying communities. For example, the microfinance institution could organize separate information sessions for the high-SES and low-SES groups to take advantage of the positive influences between groups that share similar characteristics while avoiding the negative influences across the different communities. Moreover, suppose the microfinance institution is to introduce the product into other villages (as a new product). In that case, they should send the information to individuals with the following characteristics: (1) high-SES with fewer low-SES neighbors, (2) individuals who speak a diverse set of languages, and (3) communities with similar gender ratios.

6 Conclusion

According to the role theory, the interactions of individuals depend on their roles and behaviors of interest. To conceptualize this idea, we use the underlying community structures to capture the "roles," which affect the particular decision-making processes of individuals. Specifically, we develop the Stochastic Block influences Model, which infers two types of hidden relationships: (1) block-to-block interactions, and (2) block-to-block influences on decision-making. Moreover, our model flexibly allows for both positive and negative social influences. The latter is more common in practice but has been largely ignored in the contagion models in the literature [9], [20]. In the adoption of microfinance examples we present, the inferred block-to-block influences analysis offers insights into how different social blocks influence individuals' decision-making. Our framework has far-reaching practical impacts for understanding the patterns of influences across communities and identifying the crucial characteristics of influential individuals for several applications, including the following. First, practitioners and researchers can identify the most influential communities (e.g., leaders and followers) and understand the dynamics among different communities that are not available nor observable without our SBIM model. Second, marketers can investigate which sociodemographics predict positive or negative social influences and utilize this information when introducing a product to a new market. Lastly, Marketing firms can use the influences of each individual to decide whom to target for campaigns [5].

Our SBIM is not without limitations and opens up several directions for future studies. First, future research can easily adapt the SBIM to accommodate a more complicated stochastic block model, such as a degree-corrected SBM or a power-law regularized SBM. Second, a scalable inference method as an alternative to NUTS sampling will help to improve the efficiency and scalability of SBIM. Third, future research can extend the SBIM to a dynamic model and consider its application for new products, where the influences matrix varies with time and distances from the source of information, and consider the diffusion rate. Lastly, for computer scientists and social scientists who have access to similar types of data but in different settings (e.g., different behaviors collected in different regions), it will be interesting to apply and compare the influences matrices to see if there exists any generalizable pattern to support the contagion and decision-making theories.

Acknowledgement

Y.L. is supported by the NSF grant IIS-2153468.

References

[1] J. Travers and S. Milgram, "The small world problem," 1967.
[2] L. Backstrom, P. Boldi, M. Rosa, J. Ugander, and S. Vigna, "Four degrees of separation," in *Proceedings of the 4th Annual ACM Web Science Conference*. ACM, 2012, pp. 33–42.
[3] Y. Leng, D. Santistevan, and A. Pentland, "Familiar strangers: the collective regularity in human behaviors," *arXiv preprint arXiv:1803.08955*, 2018.
[4] J. H. Fowler and N. A. Christakis, "Dynamic spread of happiness in a large social network: longitudinal analysis over 20 years in the framingham heart study," *Bmj*, vol. 337, p. a2338, 2008.
[5] Y. Leng, Y. Sella, R. Ruiz, and A. Pentland, "Contextual centrality: going beyond network structure," *Scientific reports*, vol. 10, no. 1, pp. 1–10, 2020.
[6] M. McPherson, L. Smith-Lovin, and J. M. Cook, "Birds of a feather: Homophily in social networks," *Annual review of sociology*, vol. 27, no. 1, pp. 415–444, 2001.
[7] P. Block and T. Grund, "Multidimensional homophily in friendship networks," *Network Science*, vol. 2, no. 2, pp. 189–212, 2014.
[8] E. Abbe, "Community detection and stochastic block models: recent developments," *arXiv preprint arXiv:1703.10146*, 2017.
[9] D. Centola and M. Macy, "Complex contagions and the weakness of long ties," *American journal of Sociology*, vol. 113, no. 3, pp. 702–734, 2007.
[10] P. Pin and B. W. Rogers, "Stochastic network formation and homophily," 2016.
[11] M. S. Granovetter, "The strength of weak ties," in *Social networks*. Elsevier, 1977, pp. 347–367.
[12] J. Ugander, L. Backstrom, C. Marlow, and J. Kleinberg, "Structural diversity in social contagion," *Proceedings of the National Academy of Sciences*, p. 201116502, 2012.
[13] B. J. Biddle, "Recent developments in role theory," *Annual review of sociology*, vol. 12, no. 1, pp. 67–92, 1986.
[14] D. Centola, "The spread of behavior in an online social network experiment," *Science*, vol. 329, no. 5996, pp. 1194–1197, 2010. [Online]. Available: http://science.sciencemag.org/content/329/5996/1194
[15] N. Binkiewicz, J. T. Vogelstein, and K. Rohe, "Covariate-assisted spectral clustering," *Biometrika*, vol. 104, no. 2, pp. 361–377, 2017.
[16] B. Yan and P. Sarkar, "Covariate regularized community detection in sparse graphs," *Journal of the American Statistical Association*, no. just-accepted, pp. 1–29, 2019.
[17] Z. Katona, P. P. Zubcsek, and M. Sarvary, "Network effects and personal influences: The diffusion of an online social network," *Journal of marketing research*, vol. 48, no. 3, pp. 425–443, 2011.
[18] Y. Leng, X. Dong, E. Moro *et al.*, "The rippling effect of social influence via phone communication network," in *Complex Spreading Phenomena in Social Systems*. Springer, 2018, pp. 323–333.
[19] M. D. Hoffman and A. Gelman, "The no-u-turn sampler: adaptively setting path lengths in hamiltonian monte carlo." *Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1593–1623, 2014.
[20] A. Banerjee, A. G. Chandrasekhar, E. Duflo, and M. O. Jackson, "The diffusion of microfinance," *Science*, vol. 341, no. 6144, p. 1236498, 2013.
[21] A. Y. Ng, M. I. Jordan, and Y. Weiss, "On spectral clustering: Analysis and an algorithm," in *Advances in neural information processing systems*, 2002, pp. 849–856.
[22] L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh, and A. Talwalkar, "Hyperband: A novel bandit-based approach to hyperparameter optimization," *The Journal of Machine Learning Research*, vol. 18, no. 1, pp. 6765–6816, 2017.
[23] W. J. Murdoch, C. Singh, K. Kumbier, R. Abbasi-Asl, and B. Yu, "Definitions, methods, and applications in interpretable machine learning," *Proceedings of the National Academy of Sciences*, vol. 116, no. 44, pp. 22 071–22 080, 2019.