

A BAYESIAN UNCERTAINTY QUANTIFICATION APPROACH FOR AGENT-BASED MODELING OF NETWORKED ANAGRAM GAMES

Xueying Liu
Zhihao Hu
Xinwei Deng

Department of Statistics
Virginia Tech
800 Washington Street Southwest
Blacksburg, VA 24061, USA

Chris J. Kuhlman

Biocomplexity Institute & Initiative
University of Virginia
1827 University Avenue
Charlottesville, VA 22904, USA

ABSTRACT

In group anagram games, players cooperate to form words by sharing letters that they are initially given. The aim is to form as many words as possible as a group, within five minutes. Players take several different actions: requesting letters from their neighbors, replying to letter requests, and forming words. Agent-based models (ABMs) for the game compute likelihoods of each player's next action, which contain uncertainty, as they are estimated from experimental data. We adopt a Bayesian approach as a natural means of quantifying uncertainty, to enhance the ABM for the group anagram game. Specifically, a Bayesian nonparametric clustering method is used to group player behaviors into different clusters without pre-specifying the number of clusters. Bayesian multi-nominal regression is adopted to model the transition probabilities among different actions of the players in the ABM. We describe the methodology and the benefits of it, and perform agent-based simulations of the game.

1 INTRODUCTION

1.1 Background and Motivation

In our variant of networked group anagram games (NGrAGs), players share their provided alphabetic letters with their neighbors so that all players have more letters with which to form words. In both individual and group games, the goal is to form as many words as possible within some time limit. Group and individual anagram games have been used in a wide range of studies. We partition these studies into two groups: (i) those that use anagram games as a cognitive task to measure the effect of some priming activity, and (ii) those that study the game itself. In the first group, Porath and Erez (2007) demonstrate through several series of experiments that small incivilities can so upset a person that is the target of the unkindness that this target reduces her individual performance on anagram tasks (i.e., reduces cognitive performance) by about 25% in terms of numbers of words formed. Erez and Isen (2002) used anagram games as a performance-based task to measure the effects of priming and positive affect (e.g., experiencing positive emotions). In the second group, Mayzner and Tresselt (1958) and Gilhooly and Johnson (1978) study the effects of letter order and word frequency (i.e., popularity of words) on anagram game performance. See Cedeno-Mieles et al. (2020) for other studies. Hence, it is of interest to model these games.

1.2 Description of Our Networked Group Anagram Game

Figure 1 provides illustrations of the four actions a player can take in the experimental NGrAG games, conducted online through players' web browsers. (Our models are faithful to this description of the game.)




Figure 1: Illustrative play in a networked group anagram game (NGrAG) between two players v_1 and v_2 , each with three initial letters within a black box. Each player may take any of the four actions: request letter, reply with letter, form word, and think. (a) v_1 is requesting letter d from v_2 . (b) v_2 is replying d to v_1 . (c) v_1 is forming word *had*. (d) v_1 and v_2 are thinking, which is a “no-op” or idle action. In the model of these games, a player executes one action per time step where a time step is one second. The game is described in the text. The actions are idle (a_1), reply to letter request (a_2), request letter (a_3), and form word (a_4).

Communication channels (light blue) are used by players to request letters and reply to letter requests. Overall, a player may take any of the four actions in Figure 1, any number of times, and in any order, throughout the 300-second game time. In the experiments, players are initially assigned three letters, as displayed in the boxes below the players. In a game with n players, each player v_i , $1 \leq i \leq n$, is connected to $d < n$ other players (its neighbors) with whom v_i can share letters; and hence the name Networked Group Anagram Game. A graph of players and their interaction channels in experiments is always a random regular graph (i.e., all players have the same degree d) so that we develop multiple data sets (i.e., per player) in one game, for the specified degree. This produces a connected network on all n players. If a player shares a letter with the requestor, then both the requestor and the player replying have a copy of the letter (see Figures 1b and 1c). This is to encourage sharing letters to form more words. Also, a person may use a letter in any number of words (letters are not lost when words are formed), and any number of times within a word. For instance, a player may form *deed* with only d and e , and she retains those letters. The team realizes earnings that are proportional to the total number of words formed, and these earnings are divided equally among players, regardless of individual performance. Each team member generally forms between 10 and 40 words in a game. Additional game details are in Ren et al. (2018).

1.3 Novelty and Contributions of Our Work

The main goal of this work is to produce a Bayesian uncertainty quantification (UQ) framework for agent-based modeling (ABM) of the NGrAG, which relies on modeling and estimation using experimental data collected from a set of online, human subjects NGrAGs. Note that the experimental data contain heterogeneous behaviors of players’ actions with different levels of data scarcity for different actions. Overall, the key novelty of this work is to adopt a Bayesian UQ framework for ABM to flexibly model the heterogeneous behaviors of game players. Specifically, the novelty of this work can be summarized in the following three aspects, which we now address with their commensurate contributions.

Our first contribution is use of a Bayesian clustering method based on the Dirichlet process to automatically partition players into several homogeneous groups based on the experimental data. The developed techniques do not need to pre-specify the number of clusters, while it makes the clustering results more meaningful with proper settings of priors. This is in contrast to state-of-the-art work on modeling anagram games, where number of clusters was an input to the k-means clustering method (Hu et al. 2021a). Our work and results are described in Sections 3.1 and 4.1.

Our second contribution is the formulation and construction of a Bayesian multi-logit model with proper priors to estimate the transition probabilities among the four possible player actions (see Figure 1). Given a player’s previous action at time t , along with her internal state, the probability of her taking each action at $(t + 1)$ is computed. The formal model and parameter estimation are presented in Sections 3.2 and 3.3, respectively. This is the first use of Bayesian models with anagram games. Furthermore, we take advantage of the Bayesian UQ methodology to overcome data scarcity in different actions (i.e., idle, requesting letters,

replying to requests, and forming words) within different clusters, which is novel. The use of appropriate priors can make the model estimation more stable. Because of the Bayesian approach, uncertainty and distributions of the estimated transition probabilities can be easily obtained from the posterior samples without the use of the asymptotic distribution (as done in previous work (Hu et al. 2021a)), which is less accurate when data sample sizes are small. Results are provided for selected model parameters in Section 4.2.

Third, we use the same Bayesian UQ framework to construct ABMs and exercise them within a simulation system, using heterogeneous agent behaviors beyond those tested. In Section 5, we illustrate representative ranges in player performance that simulations of the anagram game can produce. We show that the probabilities of a player’s “next action” can vary over time in the Bayesian model, and that these give rise to different regimes of player behavior over the 5-minute game duration. Plots of player behaviors and latent variables over games are provided.

2 RELATED WORK

Modeling networked group anagram games. Five works exist on modeling NGrAG games. An initial study, used to specify game conditions in preparation for online experiments, is provided in Hu et al. (2019). (The game description is given in Cedeno-Mieles et al. (2020).) The remaining works use game data to formulate logistic regression models, as follows. Cedeno-Mieles et al. (2020) describe the use of all data to formulate one model that explicitly accounts for a node’s (player’s) degree in the game network. Hu et al. (2021a) refine this work to account for player degree and behavior class via clustering of player behaviors. Each cluster has associated with it an average player’s behavior. Because each cluster has uncertainty in behavior, Hu et al. (2021b) provide another level of refinement by drilling into each cluster and quantifying worst, average, and best player behaviors in terms of (i) number of requests plus replies made and (ii) number of words formed. None of these models use a Bayesian approach, as is done here.

Bayesian clustering. Cluster analysis is the task of partitioning a set of distinguishable objects into groups such that the objects within a group are similar, whereas the groups themselves are different. K-means clustering (McQueen 1967) is one of the most commonly used clustering methods. Although k-means clustering is simple and fast, it has a restriction that the number of clusters has to be pre-specified. Determining the appropriate number of clusters for a given data set is a problem with no single right answer. Moreover, the k-means algorithm finds a local rather than a global optimum, and the results are sensitive to the initial value (James et al. 2013).

K-means clustering is non-model-based and relies on a measure of similarity among data. There also exists model-based clustering methods assuming that data are generated by a mixture of probability distributions (Wolfe 1963). The Expectation-maximization (EM) algorithm, developed by Dempster et al. (1977), is the most popular technique for mixture models, with the convergence property proved by Wu (1983). However, classical methods might fail in high-dimensional space because of over-parametrization. Therefore, Bayesian nonparametric models were introduced, assuming that the finite observations are sampled from an infinite number of latent clusters. The prior distribution is a probability measure on a space of density functions. The most common one is the Dirichlet Process (DP) (Ferguson 1973). Blackwell and MacQueen (1973) proposed the Polya urn scheme, also known as the Chinese Restaurant Process (CRP). Consider a Chinese restaurant with an infinite number of circular tables. Each customer can choose to either sit at an occupied table with a probability proportional to the number of customers already there or sit at a new table with a probability proportional to α . In this process, customers are more likely to sit at a table with more existing people, indicating the clustering property of the DP. The Bayesian model with the DP prior is called Dirichlet Process Mixture (DPM), and West et al. (1994) developed a Gibbs sampler algorithm to conduct inference on the posterior distribution.

In our framework, we perform Bayesian clustering of players based on the experimental data, which enables grouping the players based on their activity in the game and estimating parameters for each

cluster separately. Here we adopt multinomial logistic regression to model the transition probabilities of players' actions with the Bayesian approach. Then Markov chain Monte Carlo (MCMC) methods are applied to quantify the uncertainty of parameters by obtaining a sequence of random samples from the exact posterior distribution and calculating the appropriate characteristics. The proposed framework of uncertainty quantification is also applicable to other models, such as Gaussian process modeling.

3 BAYESIAN ANALYSIS AND MODELS

3.1 Bayesian Clustering for Players

It is known that players will act differently in NGrAGs. For effective analysis of the game data, it is useful to divide players into different groups based on their activity levels in the game. Therefore, we define two variables, *engagements* and *words*, to quantify the players' activity levels. *Engagements* is the sum of the number of requests and number of replies of a player, and *words* is the number of words a player forms in a game. Since the number of neighbors a player has in a game could also affect her activity level, we partition the players into two groups with $k = 2$ neighbors (group $g = 1$) and $k = 4, 6, 8$ neighbors (group $g = 2$) based on the result of hypotheses testing (Hu, Deng, and Kuhlman 2021a) on the game data.

To conduct clustering, we adopt DP-means clustering proposed by Kulis and Jordan (2012) in each group. The Dirichlet Process (DP) is

$$G \sim DP(\alpha, G_0) \quad \text{and} \quad x_i \sim G \quad i = 1, \dots, n,$$

where G_0 is the base distribution, α is the concentration parameter, and x_1, \dots, x_n is a sequence of n observations sampled from G . Let x_k^* be the unique elements in $\{x_1, \dots, x_{n-1}\}$:

$$x_n | x_1, \dots, x_{n-1} = \begin{cases} x_k^* & \text{with probability } \frac{\text{num}_{n-1}(x_k^*)}{\alpha + n - 1} \\ \text{a new draw from } G_0 & \text{with probability } \frac{\alpha}{\alpha + n - 1} \end{cases},$$

where $\text{num}_{n-1}(x_k^*) = |\{j : x_j = x_k^*, j < n\}|$ is the number of previous observations that are equal to x_k^* .

DP-means clustering is based on the Dirichlet Process Mixture (DPM) model under the Gaussian mixture distribution as

$$G \sim DP(\alpha, G_0) \quad \text{and} \quad \phi_i \sim G \quad i = 1, \dots, n \quad \text{and} \quad x_i \sim \mathcal{N}(\phi_i, \sigma I) \quad i = 1, \dots, n,$$

where the base distribution G_0 is assumed to be zero-mean Gaussian with covariance a diagonal matrix ρI , i.e., $\mathcal{N}(\mathbf{0}, \rho I)$. After drawing a distribution G from DP, ϕ_i is drawn independently from G and is taken as the mean of the Gaussian distribution that observations followed. After defining α to be a function of σ and ρ ($\alpha = (1 + \rho/\sigma)^{d/2} \cdot \exp(-\frac{\lambda}{2\sigma})$ for some λ) and implementing Gibbs sampling, Kulis and Jordan (2012) showed that as σ approaches 0, only the smallest value among the constant λ and the Euclidean distances between the point and centers of existing clusters (i.e., $\min\{\|x_i - \mu_1\|^2, \dots, \|x_i - \mu_k\|^2, \lambda\}$) will receive a non-zero cluster indicator. In this way, if the smallest value of $\{\|x_i - \mu_1\|^2, \dots, \|x_i - \mu_k\|^2\}$ is smaller than λ , data point x_i will be assigned to this closest cluster, which is the same as the conventional k-mean clustering method. However, when all the values are larger than λ , the algorithm will start a new cluster with x_i in it. Thus, the approach can automatically partition the data and avoid the problem caused by inaccurately pre-specifying the number of clusters. Algorithm 1 formalizes these ideas.

Algorithm 1: DP-means Algorithm**Input:** data x_1, \dots, x_n and penalty parameter λ **Output:** clusters l_1, \dots, l_k and number of clusters k Initialize $k = 1$, $l_1 = \{x_1, \dots, x_n\}$, cluster indicators $z_i = 1 (i = 1, \dots, n)$, and cluster center

$$\mu_1 = \sum_{i=1}^n x_i / n.$$

while Results are not convergent **do**1. **for** i in $1 : n$ **do** Compute the Euclidean distances $d_{ij} = \|x_i - \mu_j\|^2 \quad j = 1, \dots, k.$ **if** $\min_j d_{ij} > \lambda$ **then** | set $k = k + 1$, $z_i = k$, and $\mu_k = x_i.$ **else** | set $z_i = \operatorname{argmin}_j d_{ij}.$ 2. Generate clusters l_1, \dots, l_k based on z_1, \dots, z_k : $l_j = \{x_i | z_i = j\} \quad j = 1, \dots, k.$ 3. For each cluster l_j , compute the center $\mu_j = \frac{1}{|l_j|} \sum_{x_i \in l_j} x_i \quad j = 1, \dots, k.$ **3.2 Bayesian Estimation for Model Parameters**

There are several ABMs for the NGrAG such as the works in Ren et al. (2018) and Cedeno-Mieles et al. (2020). They modeled the game as a discrete-time stochastic process, where at each time step, a player executes one of the four actions as shown in Figure 1. The uncertainty from clustering is considered by constructing an ABM for each cluster.

For each player, a multinomial logistic regression is used to model π_{ij} —the probability of the player taking action a_j at time $t + 1$, given action a_i the player took at time t and four predictors provided in Table 1—as

$$\pi_{ij} = \frac{\exp(\mathbf{z}^T \beta_j^{(i)})}{\sum_{m=1}^4 \exp(\mathbf{z}^T \beta_m^{(i)})}, \quad i, j = 1, 2, 3, 4, \quad (1)$$

where $\mathbf{z} = (1, Z_B(t), Z_L(t), Z_W(t), Z_C(t))_{5 \times 1}$, and $\beta_j^{(i)} = (\beta_{j1}^{(i)}, \dots, \beta_{j5}^{(i)})^T$. For a given i (the index of action a_i at time t), the parameter set can be expressed as

$$\mathbf{B}^{(i)} = \begin{pmatrix} \beta_1^{(i)T} \\ \vdots \\ \beta_4^{(i)T} \end{pmatrix} = \begin{pmatrix} \beta_{11}^{(i)} & \beta_{12}^{(i)} & \cdots & \beta_{15}^{(i)} \\ \vdots & \vdots & \ddots & \vdots \\ \beta_{41}^{(i)} & \beta_{42}^{(i)} & \cdots & \beta_{45}^{(i)} \end{pmatrix}.$$

Table 1: The four temporal variables of players in the NGrAG and model. $Z_C(t)$ is used to ensure agents do not stagnate in thinking; this parameter enables agents to have a finite deliberation period before acting.

Variable	Description
$Z_B(t)$	Size of the buffer of letter requests that v has yet to reply to at time t .
$Z_L(t)$	Number of letters that v has available to use at t to form words.
$Z_W(t)$	Number of valid words that v has formed up to t .
$Z_C(t)$	Number of consecutive time steps that v has taken the same action.

Suppose that, in a cluster, there are n data points having the same “most recent” action i , then the next action for observation l , namely y_l , has a multinomial distribution with corresponding probability π_{li1} , π_{li2} ,

π_{li3} , and π_{li4} (π_{lij} is the π_{ij} in Equation 1 for observation l). That is,

$$y_l \sim \text{Multinomial}(1, \pi_{li1}, \pi_{li2}, \pi_{li3}, \pi_{li4}) \quad l = 1, \dots, n$$

$$f(y_l | \mathbf{B}^{(i)}) = \pi_{li1}^{y_{l1}} \times \pi_{li2}^{y_{l2}} \times \pi_{li3}^{y_{l3}} \times \pi_{li4}^{y_{l4}},$$

where $y_{lj} = \begin{cases} 1 & \text{if } y_l = j \\ 0 & \text{otherwise} \end{cases}, j = 1, 2, 3, 4$. The likelihood function is

$$L(\mathbf{B}^{(i)} | y_1, \dots, y_n) = \prod_{l=1}^n f(y_l | \mathbf{B}^{(i)}) = \prod_{l=1}^n \{ \pi_{li1}^{y_{l1}} \times \pi_{li2}^{y_{l2}} \times \pi_{li3}^{y_{l3}} \times \pi_{li4}^{y_{l4}} \}.$$

The frequentist approach to fit the multinomial regression is the maximum likelihood estimation (MLE) (Fisher 1922), and uses the asymptotic property of MLE to conduct inferences. However, the asymptotic normal approximation of MLE demands a large sample size. When the sample size is small, the bias can be substantial (Griffiths et al. 1987). Zellner and Rossi (1984) analyzed the exact posterior distribution using numerical integration techniques and showed that normal asymptotic approximations do not work well for small-size samples. For these reasons, Albert and Chib (1993) switched to the Bayesian approach and made it possible to obtain exact inferences for the parameters using MCMC techniques. The memorylessness property of MCMC (Van Ravenzwaaij et al. 2018) can alleviate the extreme value problem caused by data scarcity since every sample is only generated based on the previous one.

With the Bayesian approach, we can update the probability distribution of the unknown parameters given the observations. We assume a multivariate normal prior on $\mathbf{B}^{(i)}$:

$$\mathbf{B}^{(i)} \sim \mathcal{N}(b_0, B_0^{-1}),$$

where b_0 is the prior mean and B_0 is the prior precision matrix. Combining both the likelihood and prior distribution on $\mathbf{B}^{(i)}$, we can obtain the posterior distribution for $\mathbf{B}^{(i)}$. One can obtain the Bayesian estimator of parameters by maximizing the posterior distribution. That is,

$$\begin{aligned} \hat{\mathbf{B}}_{\text{bayes}}^{(i)} &= \arg \max \log P(\mathbf{B}^{(i)} | y_1, \dots, y_n) \\ &= \arg \max \{ \log L(\mathbf{B}^{(i)} | y_1, \dots, y_n) + \log P(\mathbf{B}^{(i)}) \}. \end{aligned}$$

3.3 Quantifying Uncertainty for Model Parameters

Based on the posterior distribution, we can quantify the uncertainty of parameters by calculating various characteristics such as posterior means, standard deviations and credible intervals. MCMC methods can be used to conduct sampling from the exact posterior distribution. The Metropolis-Hastings (M-H) (Metropolis et al. 1953; Hastings 1970) and Gibbs sampling (Geman 1984) algorithms are the two commonly-used MCMC algorithms. By applying the M-H algorithm, a sequence of parameter θ is sampled from the target distribution $f(\theta)$ where θ^{t+1} is generated by sampling a random variable θ^* from a proposal distribution $g(\theta | \theta^t)$ and accept $\theta^{t+1} = \theta^*$ with probability $\alpha = \min\{1, \frac{f(\theta^*)/g(\theta^* | \theta^t)}{f(\theta^t)/g(\theta^t | \theta^*)}\}$. In this paper, we choose the proposal distribution to be Gaussian distribution, making it a random walk M-H algorithm since $\mathcal{N}(\theta^* | \theta^t) = \mathcal{N}(\theta^t | \theta^*)$. Specifically, we have

$$g(\mathbf{B}^{(i)} | \mathbf{B}^{(i)(t-1)}) = \mathcal{N}(\mathbf{B}^{(i)(t-1)}, (B_0 + C^{-1})^{-1}),$$

where B_0 is the prior precision matrix, and C is the sample variance-covariance matrix of the MLEs. These steps are formalized in Algorithm 2.

Algorithm 2: Random Walk Metropolis-Hastings Algorithm for $\mathbf{B}^{(i)}$

Initialize $\mathbf{B}^{(i)(0)} = \mathbf{0}$.

for t in $1 : T$ **do**

1. Generate a sample $\mathbf{B}^{(i)*}$ from proposal distribution $g(\mathbf{B}^{(i)}|\mathbf{B}^{(i)(t-1)})$.
 2. Calculate the acceptance ratio $\alpha = \min\{1, \frac{P(\mathbf{B}^{(i)*}|y_1, \dots, y_n)}{P(\mathbf{B}^{(i)(t-1)}|y_1, \dots, y_n)}\}$.
 3. $\mathbf{B}^{(i)(t)} = \begin{cases} \mathbf{B}^{(i)*} & \text{with probability } \alpha \\ \mathbf{B}^{(i)(t-1)} & \text{with probability } 1 - \alpha \end{cases}$.
-

In practice, it could be difficult to choose the most appropriate priors and one may want to explore the performance under a variety of different priors. Jeffreys (1946) introduced non-informative priors that are invariant under the transformation of parameters. For $\mathbf{B}^{(i)}$ in our model, setting $B_0 = 0$ leads to the Jefferys’ prior (a flat prior) for the mean of normal distribution.

4 MODEL EVALUATION

4.1 DP-Means Clustering Players

Section 3.1 describes the scalable DP-based Bayesian clustering approach. We standardized the *words* and *engagements* variables to make the selection of values of cluster penalty parameter λ more reasonable. Here we choose $\lambda = 2.5$ for clustering, indicating that once the distances between a data point and every existing cluster center are greater than 2.5 after standardization, we will form a new cluster. In this way, four clusters are formed for both $k = 2$ neighbors [group $g = 1$] and 4,6,8 neighbors [group $g = 2$].

Figure 2 contains boxplots of performance for the four $g = 1$ clusters in the top row, and for the four $g = 2$ clusters in the bottom row. For the $g = 1$ players, cluster 1 has significantly fewer engagements per game than the other three clusters, and the players in clusters 1 and 2 form fewer words than do the players in clusters 3 and 4. These general observations also hold for the $g = 2$ game player behaviors, with some differences in details. Also, the magnitudes of words formed and engagements are greater in $g = 2$. These differences support the contention that clustering of players by behaviors is warranted.

4.2 Bayesian Approach for Model Parameters

In multinomial logistic regression models, given a player’s most recent action, we model four possible outcomes for the next action taken by a player: idle, reply, request, and form a word. Setting idle as a reference category, we get 15 non-zero parameters $\{\beta_2^{(i)T}, \beta_3^{(i)T}, \beta_4^{(i)T}\} = \{\beta_{21}^{(i)}, \dots, \beta_{25}^{(i)}, \beta_{31}^{(i)}, \dots, \beta_{35}^{(i)}, \beta_{41}^{(i)}, \dots, \beta_{45}^{(i)}\}$ in $\mathbf{B}^{(i)}$ for every “most recent” action. By using the M-H algorithm with 10000 iterations after 1000 burn-in, we can draw a sequence of random samples for each parameter from the posterior distribution. The Geweke diagnostic (Geweke 1991) is used to test the convergence of MCMC samples. Figure 3 shows the density plots for $\beta_{22}^{(1)}, \beta_{23}^{(1)}, \beta_{24}^{(1)}, \beta_{25}^{(1)}$ of players in cluster $c = 1$ of group $g = 1$ whose most recent action is idle ($i = 1$). The uncertainties in player actions from the experimental games are represented by the distributions of β coefficients, such as those in Figure 3.

For each parameter, we take the mean of the random samples as the Bayesian estimation of $\mathbf{B}^{(i)}$ shown in Table 2. Then the probability of players in cluster 1 of $k = 2$ neighbors [group $g = 1$] with action at time t of idle taking action j at time $(t + 1)$ is calculated using Equation 1. The probability of idle is computed as one minus the sum of the probabilities of the three “next actions” in Table 2.

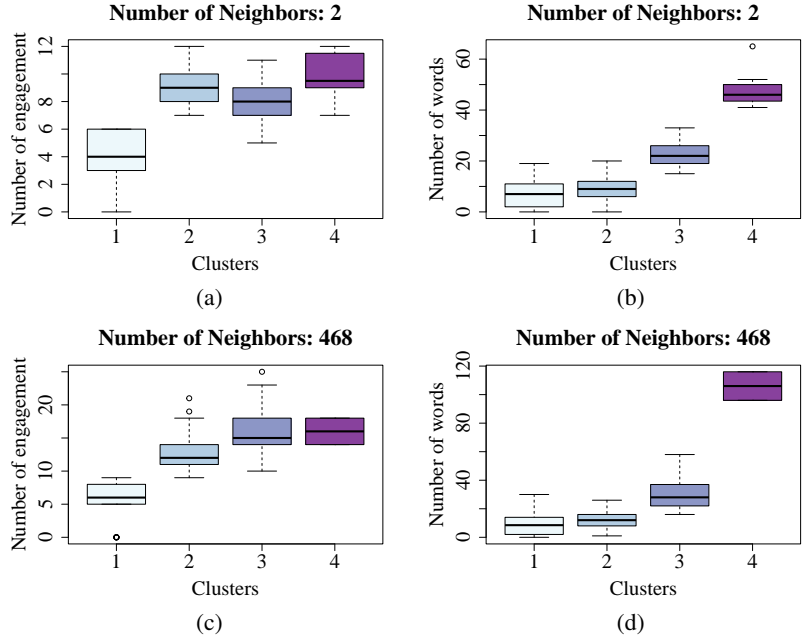


Figure 2: Results from clustering analysis of NGrAG game data. (a) Boxplot of number of replies and requests (engagements) for 2 neighbors. (b) Boxplot of number of words for 2 neighbors. (c) Boxplot of number of replies and requests (engagements) for 4, 6, and 8 neighbors. (d) Boxplot of number of words for 4, 6, and 8 neighbors. All values are per player per game.

Table 2: Posterior mean and standard deviation values of parameters in $\mathbf{B}^{(t)}$ for players in cluster 1 of $k = 2$ neighbors group with initial (or most recent) action being a_1 (idle).

Next Action	Intercept		Buffer $Z_B(t)$		Letter $Z_L(t)$		Word $Z_W(t)$		Constant $Z_C(t)$	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SD
reply	-3.1159	0.806	0.2910	0.097	-0.5800	0.246	-0.0303	0.078	-0.0153	0.004
request	-3.5242	0.499	0.1262	0.067	-0.1770	0.146	-0.0660	0.049	-0.0163	0.004
word	-5.4232	0.277	0.0337	0.037	0.5199	0.066	-0.0599	0.020	-0.0209	0.003

5 AGENT-BASED SIMULATIONS

The simulation system encodes the models developed in Section 3. Each game player is modeled as a node in a graph and has assigned to it a behavior that corresponds to a particular group g and cluster c .

NGrAG scenarios. The game configuration that we model is given in Figure 4. There are three purple nodes of $d = 2$ and hence are in $g = 1$ (where $d \leq 2$), while the two green nodes with $d = 3$ are in $g = 2$ (where $d \geq 3$). We evaluate the effects of different cluster assignments to the nodes. For each node, we identify the group and cluster pair by $[g, c]$. Each player is assigned four letters (one vowel and three consonants) that they can use to form words and to share with their neighbors. The letters are those most commonly used in words so that agents are not prevented from forming words owing to poor letters (e.g., they are not given letters such as q , x , and z).

Simulation. Each simulation instance follows (generally) the conditions of a game that were provided in Section 1.2. Each player is initially provided with four letters (in experiments, each player is given three letters). Each player tries to form words while also requesting letters from their neighbors (to form more words) and replying to letter requests (to help their neighbors form more words). The game duration is

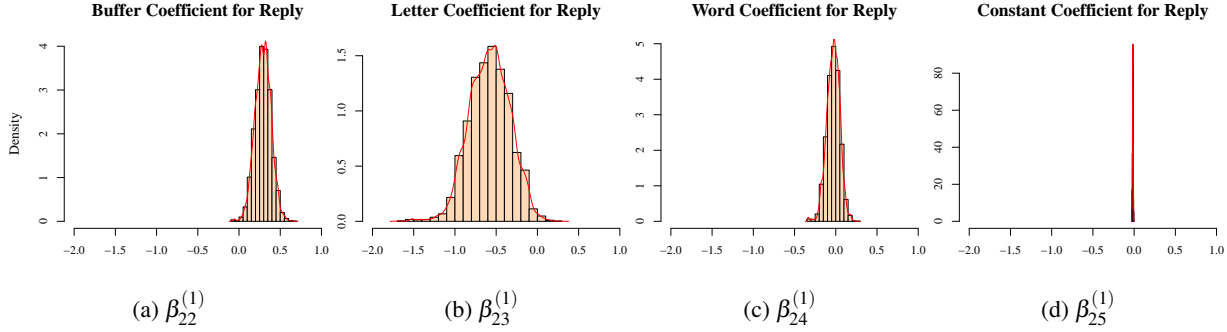


Figure 3: Density plots for $\beta_{22}^{(1)}, \beta_{23}^{(1)}, \beta_{24}^{(1)}, \beta_{25}^{(1)}$ of players in cluster $c = 1$ of group $g = 1$ with initial action being idle. The x-axis range is the same in all four plots, but the y-axis ranges vary markedly.




Figure 4: Five-node (i.e., player) network where two players (in green) have degree $d = 3$ and three nodes (in purple) have $d = 2$. Edges in the network are communication channels and conduits to share alphabetic letters between pairs of nodes. The green nodes are in group $g = 2$; purple nodes are in group $g = 1$.

300 seconds, where time is incremented in 1-second intervals (i.e., player actions take place at integer time in seconds); the experimental data support this approach because, over all games, in a negligible number of instances, a player takes two actions within one second. A simulation (i.e., one software execution) consists of 100 simulation instances. These results are then averaged over time.

Simulation parameters. Owing to space limitations, we study only one cluster c for each of two groups g : $[g, c] = [1, 1]$ for the low degree ($d = 2$) nodes and $[g, c] = [2, 2]$ for the high degree ($d = 3$) nodes.

Simulation results. Figure 5 shows two plots of the probability histories for the four actions (idle [or thinking], replying to letter requests, requesting letters, and forming words; see legends) for each of nodes 3 and 4. These data are for one out of 100 simulation instances. We denote the probability vector and these probability components as $p_{act} = (p_{think}, p_{srep}, p_{sreq}, p_{fw})$. Recall that at each time step, the probability of an agent (node) taking each action is computed, so each curve is composed of 300 data points, one for each second in the 300-second game. In both cases, the probability of thinking (e.g., deciding what to do next) is about 0.9, with the remaining probability distributed among the three actions involving letters and words. Note that each probability is dependent on the \mathbf{z} vector for each node, which changes in time, and the β vectors (based on node degree and cluster), per Equation (1), so that probabilities may evolve in different ways. This is observed in Figure 5, where p_{srep} , p_{sreq} , and p_{fw} differ in magnitude and also p_{fw} differs markedly in character (i.e., shape).

Figure 6 shows the time histories of actions for different players (agents, nodes) for one of the 100 iterations (the same iteration used in Figure 5). In Figure 6a, the counts of words formed per agent are given. Nodes 1 and 4, with the greatest degrees, ultimately form the most words, for the specified clusters of behaviors. The number of words formed for $d = 2$ players varies from one to seven. The difference in words formed between the two $d = 3$ players is less, for this particular iteration. Of note is that the character of the curves for the three $d = 2$ nodes is concave-down, while that for node 4 is concave up. This follows directly from the character of the p_{fw} curves in Figure 5. It is interesting that nodes 1 and 4 produce different regimes of behavior in time. That is, both players form relatively fewer words in roughly




Figure 5: Probability time histories for two nodes of the graph of Figure 4. One curve in each plot corresponds to the probability for an agent taking the corresponding action. These data are for one of the 100 simulation instances. Each player has four initial letters. (a) Probabilities for node 3. (b) Probabilities for node 4.

the first 200 seconds of the game, and then form disproportionately more words in the last 100 seconds. In Figures 6b and 6c, the word histories are repeated (for comparison) and the counts of the other actions are also given. From the legends, these histories are numbers of: (letter) replies received, (letter) replies sent, (letter) requests received, (letter) requests sent, and words formed. Among the many observations to be made, two of these are that node 4 is more active than node 3, and that of the ten letter requests made by node 4, it only received five replies with letters.




Figure 6: Time history simulation results for one iteration out of the 100 iterations (the same iteration as that used for Figure 5) for the 5-node graph of Figure 4. Each player has four initial letters. (a) Number of words formed by each player. (b) Counts of each action for node 3. (c) Counts of each action for node 4.

Figure 7 shows the same three plots as in Figure 6, but now the data are the averages over the 100 iterations. For many quantities plotted, the characteristics of the curves are somewhat similar to those in the preceding figure, but there are also several differences. Among the latter is that node 0 forms slightly more words than do nodes 1 and 4, on average. In general, the shape of these curves will be influenced by the behavior clusters assigned to the nodes. By contrasting these two figures, it is evident that the per-iteration data are useful in accentuating the idea that players make discrete choices at particular times.

6 CONCLUSION

In this work, we developed a Bayesian framework of uncertainty quantification for agent-based modeling (ABM) of networked group anagram games. The Bayesian approach includes Bayesian clustering to partition players based on their activity level and Bayesian analysis of experimental data to better quantify the uncertainty of model estimation through posterior sampling. The estimated Bayesian model in each cluster enables ABM to incorporate uncertainty in player behavior in networked anagram games. It also makes the corresponding ABM robust to the data scarcity issue. Future work includes more efficient




Figure 7: Average time history simulation results over 100 iterations or instances for the 5-node graph of Figure 4. Each player has four initial letters. (a) Number of words formed by each player. (b) Counts of each action for node 3. (c) Counts of each action for node 4.

posterior inference methods (e.g. approximation Bayesian computation (ABC) (Beaumont 2019) is a popular approach to tackle intractable likelihood functions) and model validation by comparing simulation results with experimental data. It is worth pointing out that the general outline of the proposed Bayesian UQ approach is not limited to ABM for the NGrAG. There are other multi-action games, e.g., Mason and Watts (2012), Coviello et al. (2012), that could use our approach for UQ. But particulars of models would be based on game actions and data available for modeling.

ACKNOWLEDGMENT

We thank the anonymous reviewers for their helpful feedback. We thank our colleagues at NSSAC and Research Computing at The University of Virginia for providing computational resources and technical support. This work has been partially supported by NSF CRISP 2.0 (CMMI Grant 1916670).

REFERENCES

- Albert, J. H., and S. Chib. 1993. “Bayesian analysis of binary and polychotomous response data”. *Journal of the American Statistical Association* 88(422):669–679.
- Beaumont, M. A. 2019. “Approximate bayesian computation”. *Annual Review of Statistics and its Application* 6:379–403.
- Blackwell, D., and J. MacQueen. 1973. “Ferguson Distributions Via Polya Urn Schemes”. *The Annals of Statistics* 1.
- Cedeno-Mieles, V., Z. Hu, Y. Ren, X. Deng, A. Adiga, C. L. Barrett, N. Contractor, S. Ekanayake, J. M. Epstein, B. J. Goode, G. Korkmaz, C. J. Kuhlman, D. Machi, M. W. Macy, M. V. Marathe, N. Ramakrishnan, S. S. Ravi, P. Saraf, and N. Self. 2020. “Networked Experiments and Modeling for Producing Collective Identity in a Group of Human Subjects Using an Iterative Abduction Framework”. *Soc. Netw. Anal. Min. (SNAM)*:43.
- Coviello, L., M. Franceschetti, M. D. McCubbins, R. Paturi, and A. Vattani. 2012. “Human Matching Behavior in Social Networks: An Algorithmic Perspective”. *PLoS ONE* 7(8):e41900.
- Dempster, A., N. Laird, and D. Rubin. 1977, 01. “Maximum Likelihood From Incomplete Data Via The EM algorithm”. *Journal of the Royal Statistical Society. Series B (Methodological)* 39:1–38.
- Erez, A., and A. M. Isen. 2002. “The Influence of Positive Affect on the Components of Expectancy Motivation”. *Journal of Applied Psychology* 87:1055–1067.
- Ferguson, T. 1973. “A Bayesian Analysis of Some Non-Parametric Problems”. *The Annals of Statistics* 1:209–230.
- Fisher, R. A. 1922. “On the Mathematical Foundations of Theoretical Statistics”. *Philosophical Transactions of the Royal Society of London. Series A* 222(594-604):309–368.
- Geman, D. 1984, 01. “Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images”. *IEEE Trans. Pattern Anal. Mach. Intell.* 6:721–741.
- Geweke, J. F. 1991. “Evaluating the Accuracy of Sampling-Based Approaches to the Calculation of Posterior Moments”. Technical report, Federal Reserve Bank of Minneapolis.
- Gilhooly, K. J., and C. E. Johnson. 1978. “Effects of Solution Word Attributes on Anagram Difficulty: A Regression Analysis”. *Quarterly Journal of Experimental Psychology* 30(1):57–70.
- Griffiths, W. E., R. C. Hill, and P. J. Pope. 1987. “Small Sample Properties of Probit Model Estimators”. *Journal of the American Statistical Association* 82(399):929–937.

- Hastings, W. 1970. "Monte Carlo Sampling Methods Using Markov Chains and Their Application". *Biometrika* 57.
- Hu, Z., X. Deng, B. J. Goode, N. Ramakrishnan, P. Saraf, N. Self, A. Adiga, G. Korkmaz, C. J. Kuhlman, D. Machi, M. V. Marathe, S. S. Ravi, Y. Ren, V. Cedeno-Mieles, and S. Ekanayake. 2019. "On the Modeling and Agent-Based Simulation of a Cooperative Group Anagram Game". In *Proceedings of the 2019 Winter Simulation Conference*, edited by N. Mustafee, K.-H. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo, P. Haas, and Y.-J. Son, 169–180. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Hu, Z., X. Deng, and C. J. Kuhlman. 2021a. "Approaches to Uncertainty Quantification in Building Models of Human Behavior". In *Proceedings of the 2021 Winter Simulation Conference*, edited by S. Kim, B. Feng, K. Smith, S. Masoud, Z. Zheng, C. Szabo, and M. Loper, 1–12. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Hu, Z., X. Deng, and C. J. Kuhlman. 2021b. "Versatile Uncertainty Quantification of Contrastive Behaviors for Modeling Networked Anagram Games". In *Complex Networks*.
- James, G., D. Witten, T. Hastie, and R. Tibshirani. 2013. *An Introduction to Statistical Learning*, Volume 112. Springer.
- Jeffreys, H. 1946. "An Invariant Form for the Prior Probability in Estimation Problems". *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* 186(1007):453–461.
- Kulis, B., and M. Jordan. 2012. "Revisiting k-means: New Algorithms via Bayesian Nonparametrics". In *ICML*.
- Mason, W., and D. J. Watts. 2012. "Collaborative Learning in Networks". *PNAS* 109(3):765–769.
- Mayzner, M. S., and M. E. Tresselt. 1958. "Anagram Solution Times: A Function of Letter Order and Word Frequency". *Journal of Experimental Psychology* 56(4):376.
- McQueen, J. 1967. "Some Methods for Classification and Analysis of Multivariate Observations". *Computer and Chemistry* 4:257–272.
- Metropolis, N., A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. 1953. "Equation of State Calculations by Fast Computing Machines". *The Journal of Chemical Physics* 21(6):1087–1092.
- Porath, C., and A. Erez. 2007. "Does Rudeness Really Matter? The Effects of Rudeness on Task Performance and Helpfulness". *Academy of Management Journal* 50(5):1181–1197.
- Ren, Y., V. Cedeno-Mieles, Z. Hu, X. Deng, A. Adiga, C. Barrett, S. Ekanayake, B. J. Goode, G. Korkmaz, C. J. Kuhlman, D. Machi, M. V. Marathe, N. Ramakrishnan, S. S. Ravi, P. Saraf, N. Selt, N. Contractor, J. Epstein, and M. W. Macy. 2018. "Generative Modeling of Human Behavior and Social Interactions Using Abductive Analysis". In *ASONAM*, 413–420.
- Van Ravenzwaaij, D., P. Cassey, and S. D. Brown. 2018. "A Simple Introduction to Markov Chain Monte Carlo sampling". *Psychonomic Bulletin & Review* 25(1):143–154.
- West, M., P. M., and M. Escobar. 1994. "Hierarchical Priors and Mixture Models, With Application in Regression and Density Estimation".
- Wolfe, J. H. 1963. *Object Cluster Analysis of Social Areas*. Ph. D. thesis, University of California.
- Wu, C. J. 1983. "On the Convergence Properties of the EM Algorithm". *The Annals of Statistics*:95–103.
- Zellner, A., and P. Rossi. 1984. "Bayesian Analysis of Dichotomous Quantal Response Models". *Journal of Econometrics* 25:365–393.

AUTHOR BIOGRAPHIES

XUEYING LIU is a Ph.D. candidate in the Department of Statistics at Virginia Tech. She received her M.S. in Statistics from Georgia Tech. Her email address is xliu96@vt.edu.

ZHIHAO HU is a recently-graduated Ph.D. student in the Department of Statistics at Virginia Tech. He received his M.S. in Materials Science & Engineering from Virginia Tech. His email address is huzhihao@vt.edu.

XINWEI DENG is a professor in the Department of Statistics at Virginia Tech. He received his Ph.D. in Industrial Engineering from Georgia Tech. His email address is xdeng@vt.edu.

CHRIS J. KUHLMAN is a Research Associate Professor at the University of Virginia. He received a Ph.D. in Computer Science at Virginia Tech. His email address is hugo3751@gmail.com.