

Geographical influences of an emerging network of gang rivalries

Rachel A. Danson^{a,*}, Laura M. Smith^a, Alethea Barbaro^a, Andrea L. Bertozzi^a, Shannon E. Reid^b, George E. Tita^b

^a*UCLA Mathematics Department
520 Portola Plaza
Box 951555
Los Angeles, CA 90095-1555
USA*

^b*Department of Criminology, Law and Society
2340 Social Ecology II
Irvine, CA 92697-7080
USA*

Abstract

We propose an agent-based model to simulate the creation of street gang rivalries. The movement dynamics of agents are coupled to an evolving network of gang rivalries, which is determined by previous interactions among agents in the system. Basic gang data, geographic information, and behavioral dynamics suggested by the criminology literature are integrated into the model. The major highways, rivers, and the locations of gangs' centers of activity influence the agents' motion. We use a policing division of the Los Angeles Police Department as a case study to test our model. We use common metrics from graph theory to analyze our model, comparing networks produced by our simulations and an instance of a Geographical Threshold Graph to the network existing in the criminology literature.

Keywords: Agent-based Model, Complex System, Dynamic Network, Gang Rivalry, Biased Lévy Walk, Geographical Threshold Graph

1. Introduction

Street gangs are a growing problem around the world [12, 28, 27]. In fact, recent statistics from The National Gang Intelligence Center estimate there are 1 million active gang members in the United States alone [49]. Violence is intrinsic to street gangs, and rival gangs battle to gain respect and street reputation [61, 13]. Criminal activities perpetrated by gang members, including armed robbery, homicide, drug dealing, and auto theft, drain cities and governments of tight resources and also pose safety threats to community members. Much of the research on street gangs has been conducted within the United States, though there have been some efforts to understand the phenomenon in Europe and other parts of the world [12, 28, 27].

Violence perpetrated by gang members is frequently against members of a different gang. In areas with numerous gangs, it is common for gangs to have multiple violent interactions with many of the other gangs. Further, street gang members typically have locations, known as set spaces, where they spend large quantities of time [62, 47]. It is therefore reasonable to think of each gang as a node embedded in Euclidean space [50, 64]. Within this framework, the existence of persistent violence between two gangs becomes an edge connecting two nodes. From this construction, one can view a collection of gangs as a spatially embedded network [63]. The Hollenbeck policing division of eastern Los Angeles is marked by a particularly high degree of violent crimes involving gang members, including homicides and aggravated assaults [50, 25]. It is for this reason and others listed in Section 1.4 we consider Hollenbeck as a test case for our model.

*Corresponding author

Email addresses: dansonr@math.ucla.edu (Rachel A. Danson), lsmith@math.ucla.edu (Laura M. Smith), alethea@math.ucla.edu (Alethea Barbaro), bertozzi@math.ucla.edu (Andrea L. Bertozzi), reids@uci.edu (Shannon E. Reid), gtita@uci.edu (George E. Tita)

1.1. Network Models

General network models and the corresponding analysis are useful for describing the behavior of complex systems and have played an increasingly active role [40, 41, 43]. One way networks are treated in the literature is by analyzing the statistical properties of a given network. Another approach is to consider the construction of a network. There are many instances where the network of interest is not known, but there is some knowledge of the processes by which the network is formed. One popular method to construct a network is to view it as a random graph. Each edge is added with a predetermined probability, often dependent on the weight of the nodes [1, 44, 45].

In some applications, including gang rivalry networks, the geographic location of the nodes influences the structure of the network. In such cases, geographic features should be considered as part of the random network model. For example, interstate highways have been shown to be structurally different from scale-free networks such as Internet and airline flight networks [20]. The importance of geography is also seen in friendship networks [68]. In their paper, Liben-Nowell et al. use the publicly accessible location of 495,836 bloggers in the Live Journal to investigate effects of spatial proximity on friendships [30]. The study found that an estimated 69% of a person's friends can be described by geography.

One method for incorporating geographical information into the random graph construction is by using a Geographical Threshold Graph [34, 7, 8]. This is a random graph on a set of randomly weighted nodes, where the nodes are located in a metric space and the connections are determined by thresholding a function of the distance and the weights. This provides a computationally efficient way to construct a rivalry network while incorporating some geographic information. We use an instance of a Geographical Threshold Graph as a baseline against which we compare our model.

1.2. Agent-Based Models

Though using a randomly constructed network may produce a reasonable simulation of an observed network, other phenomena of interest beyond the structure of the network are not obtained from this type of model. An alternative is to use an agent-based approach. This has become a widely used tool in the area of complex systems [66, 60, 69, 18, 17, 32, 58]. For example, agent-based models have been used for modeling many types of cooperative behavior [31, 56, 14, 11, 24]. In a network context, this approach enables exploration of how changing dynamics of individual agents can affect the evolution of the network, providing control parameters which would be inaccessible in a graph-based model. This method can easily incorporate environmental and spatial information inherent to the system, e.g. in [4], which uses environmental cues to reproduce and predict fish migration.

In particular, we are interested in the coupling between the network and the underlying system. There has been some exploration of this in the literature. For example, Schweitzer and Tilch provide one example of model that uses an agent-based approach to form an emerging network [52, 53]. They model the chemical trail formed by ants searching for food at an unknown location. As the ants search their environment, networks of chemical trails form with which the ants interact. Another example is that of the EpiSims model [65, 36]. Here, the contact networks of the populations are evolving over time and depend on the internal attributes of the people in the population. In turn, as a disease is spread through the contact network, the movements of the people change in response to the disease, producing a non-trivial interaction between the system and the network.

1.3. Previous Work on Crime Modeling

Various models have been created to address criminal activity [23, 48, 46, 9, 56]. One such work uses an agent-based model to understand the formation of crime hot spots [55]. In the model proposed by Egesdal et al., an agent-based approach was used to simulate the location of violent interactions and gang retaliations in Hollenbeck [15]. Embedded in the model was a rivalry network. Though the model recreated similar features to the violence data, the model did not incorporate geographic features. In addition, agents targeted specific gangs based on probabilities corresponding to the current rivalry network. However, according to the criminology literature, this retaliatory behavior is only seen on short time scales [37]. Gang members tend to avoid the territory of rival gangs [2, 29].

Although not an agent-based model, Mohler examines the short term retaliatory behavior of the rivalries based on between gang violence data from the LAPD [37]. Each violent event between two gangs is considered

an instance of a point process associated with that pair of gangs. The intensity of the rivalry depends directly on the network of unidirectional violent interactions. This provides a top-down approach to understanding immediate consequences of violence among gangs within a system. In our work, we wish to understand factors associated with the long term gang rivalry structure.

1.4. Hollenbeck

Hollenbeck is a policing division located in eastern Los Angeles, surrounded by downtown Los Angeles to the west, Pasadena to the northeast, Vernon to the south, and to the east the unincorporated area of East Los Angeles, see Figure 1. Hollenbeck provides a diverse geography with many highways cutting through the region and is bounded by the Los Angeles River. It encompasses an area of roughly 39.4 km². Hollenbeck is home to approximately twenty-nine active gangs with sixty-nine rivalries among them [50, 64]. The set spaces for the gangs and the corresponding observed rivalry network are displayed in Figure 1, as given in [50].

Certain properties of Hollenbeck make it accessible to modeling the gang rivalry networks outlined in [64, 50]. First, it is a closed system in that the gang activity within Hollenbeck is generally isolated from gang activity outside of Hollenbeck. Further, the motivation for violence between gangs is largely characterized by disputes over geographical gang territories, as opposed to drug and racially motivated violence. Data on the geography of Hollenbeck is easily accessible, and there has been explicit documentation of the observed rivalry network.

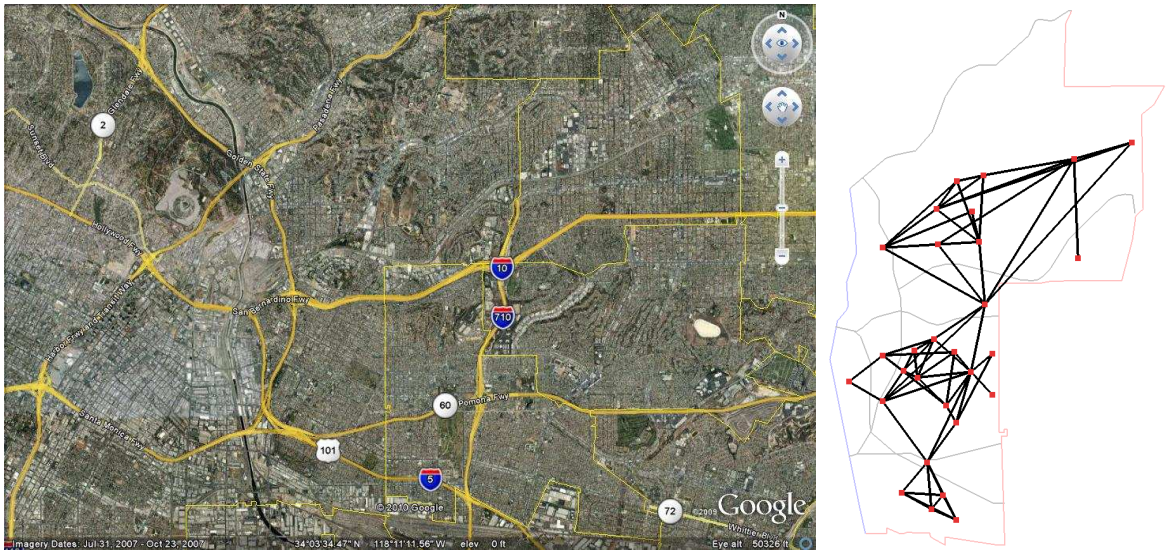


Figure 1: Google EarthTM Image of the Hollenbeck area (left). Map of the Hollenbeck area with the location of the gang set spaces and the corresponding rivalry network approximated by [50] (right). Major roads, highways, the Los Angeles river, and division lines are also seen in both images.

We propose an agent-based model that incorporates geographical features and is coupled to a dynamically evolving network. This model simulates the mobility of gang members and the resulting interactions. We compare the resulting simulated network to the gang rivalry network observed in the eastern Los Angeles division of Hollenbeck [64, 50]. In Section 2, we outline the proposed model. In Section 3, we describe two baseline models, one instance of a Geographical Threshold Graph and a network derived from Brownian Motion, to which we compare our model. In Section 4, we describe a series of metrics from network theory, examine long term behavior of the model, and compare the networks against the metrics. Section 5 provides a sensitivity analysis of our model. We conclude and give future directions in Section 6.

2. Our Model

The purpose of this model is to understand the extent to which simple behavioral rules and geographical factors, such as road density, highways, and locations of gangs' set spaces, could influence the structure of gang rivalry networks.

2.1. Motivation for model construction

The intent of this model is to capture the broad statistical features of human mobility with an emphasis on gang members' movements. Empirical data on the location and individual movements of each gang member is inaccessible, so we characterize the movements of the individual gang members in a statistical sense based on the literature on human mobility. Several studies give compelling evidence that when people move in an unconstrained environment, the jump lengths between movements is distributed like a power law [10, 51]. Further, in the presence of obstacles such as roads and buildings, the jump lengths more accurately follow a bounded power law distribution [21].

However, determining the statistical properties of the jump length is only one aspect of movement dynamics. In their paper, Rhee et al. discuss the need to incorporate geographical features and the tendency for people to go home [51]. Gonzalez et al. confirmed in their data that humans do tend to frequent a small number of locations often [21]. For these reasons, the agents in our model pick their jump length from a Bounded Pareto distribution and have a directional choice in movement.

In the case of gangs in Hollenbeck, it is reasonable to assume that the gang members have a clear sense of the location of their home territory, or set space, as well as the location of their rival gangs' set spaces [62]. Literature on gang activity suggests that, in general, gang members tend to stay away from their rival gangs' set spaces [29]. Unlike other criminal groups, such as organized crime syndicates and insurgency groups that strive for secrecy, street gangs are social organizations that proudly demarcate their territory and announce their enemies through the use of graffiti. Gangs create social boundaries and therefore areas of avoidance [2]. Our model incorporates this social geography into agents' movement dynamics.

One aspect of modeling human mobility that was touched on, but not fully explored, by the previous literature is the role of physical features specific to urban areas that may constrain agents' movement. The first consideration is the ease with which an agent can move through a city. We posit that in areas where there is a dense street network, the likelihood of an agent to move long distances is small due to such obstacles as the high density of people and cars, as well as traffic lights. On the other hand, areas where the road density is lower, agents should be able to move longer distances. A second physical consideration that affects human mobility in a city are the highway systems and rivers that can cut across the region. These features are not impassible, in that there are underpasses and bridges. However, they do provide an obstacle that may make it difficult to cross. Therefore, in our model, these are seen as semi-permeable boundaries that effect the agents' movements.

2.2. Model summary

Agents in the model move based on their location with respect to their and other gangs' set spaces and interact with agents of different gang affiliations. We count the number of interactions between gangs, and when agents of different gangs move within a certain distance of each other, the number of interactions between those gangs increases by one. As the simulation progresses, a network structure emerges. The weighted network of interactions in turn influences the directional decisions of the agents.

2.3. Entities, state variables, and scales

2.3.1. Agents

The agents of this model are gang members in a city. Each agent is associated with exactly one gang. For simplicity we assume agents' directional choice is dictated only by the location of the gang set spaces. All agents know the location of their home and rivals' set spaces. We divide the city into regions based on geographical boundaries, such as rivers and highways. An agent knows which region it is currently in as well as the region of any prospective new locations. When two agents are within interaction range, we consider them to have interacted and the corresponding element of the rivalry matrix, R , is updated. There are no immediate changes to the location of the gang members. Refer to Section 2.3.3 for details on R .

2.3.2. Environment

The environment of interest is on the scale of a small city. Agents and gang set spaces in the environment have a coordinate location in Euclidean space. Further, the set spaces provide the spatially embedded nodes of the gang rivalry network. The physical geographical features of the city are approximated by an NxM grid. Agents move in Euclidean space, and each point in the region is identified with the nearest grid element. The size and number of grid elements are constant throughout the simulation and will be limited by the available data and the memory of the computer.

Two features encoded in this NxM grid are the road density and semi-permeable boundaries represented by a region map. The road density is estimated for this project using the Weighted H¹ Maximum Penalized Likelihood Estimation method with a road map as the initial data, as in [59]. Other methods for density estimation, such as kernel density estimation or other Maximum Penalized Likelihood Estimation (MPLE) methods, could also be used to construct the road density [57, 16, 38, 22]. Each element of the NxM density map contains a number between 0 and 1. A value of 0 implies a low road density whereas a value of 1 implies high road density. The semi-permeable boundaries, corresponding to such objects as highways and rivers, are assumed to split the environment into distinct regions. Therefore, each element of the region grid corresponds to a specific region. Paired with this region grid is a matrix storing the associated probability of an agent to cross from one region to another. This is implemented to discourage agents from crossing freeway boundaries.

2.3.3. Rivalries

The network structure of the rivalries is encoded in a weighted adjacency matrix, R . Each element R_{ij} contains the current history of interactions between gang i and gang j . At the end of a simulation, we construct a thresholded rivalry graph where an edge between gang i and j exists if either $\rho_i(j)$ or $\rho_j(i)$ is larger than a given threshold T , where

$$\rho_i(j) = \frac{R_{ij}}{\sum_{k=1}^N R_{ik}} \quad \text{and} \quad \rho_j(i) = \frac{R_{ji}}{\sum_{k=1}^N R_{jk}}. \quad (1)$$

The quantity, $\rho_i(j)$, represents the proportion of gang i 's interactions which have occurred with gang j . Note that $\rho_i(j)$ is not necessarily equal to $\rho_j(i)$; however, this thresholding yields a bidirectional network or, equivalently, a symmetric adjacency matrix.

2.4. Process overview and Scheduling

At each iteration an agent is chosen from the set of all agents with equal probability. The agent then preforms one step of a biased truncated Lévy walk. In particular, the jump length is chosen from the Bounded Pareto probability distribution,

$$P(x; k, x_m, x_M) = \frac{kx_m^k x^{-k-1}}{1 - \left(\frac{x_m}{x_M}\right)^k} \quad k > 0, \quad x_M \geq x \geq x_m > 0. \quad (2)$$

For all agents the minimum jump length, x_m , and scale, k , are fixed. To determine the maximum jump length, x_M , the agent uses the approximated road density of the agent's corresponding location from the environment grid. The road density at this location, δ , is between 0 and 1. The maximum jump length is then calculated via

$$x_M = (1 - \delta) \cdot A + a, \quad (3)$$

where A is the largest maximum jump length and a is the smallest maximum jump length.

Given an agent in gang i , the bias direction, μ_i , incorporates the agent's location with respect to its home set space and the location of its rival gangs' set spaces via

$$\langle x, y \rangle = H_i(\|\vec{G}_i\|_2) \frac{\vec{G}_i}{\|\vec{G}_i\|_2} + \sum_{j \neq i} D_{ij}(\|\vec{G}_j\|_2) \frac{\vec{G}_j}{\|\vec{G}_j\|_2} \quad (4)$$

$$\mu_i = \tan^{-1} \left(\frac{y}{x} \right).$$

Here, \vec{G}_l is the vector that points to the set space of gang l from the location of the agent. When $l = i$, this vector points towards the agent's home set space, and when $l \neq i$, it points towards a different gang's set space. This concept is shown in the cartoon example in Figure 2.

In Equation 4, H_i gives the rules for weighting towards a gang member's own home set space. The weightings toward or away from different gangs' set spaces are determined by D_{ij} . Our H_i and D_{ij} take the following form:

$$H_i(\|\vec{G}_i\|_2) = h_i \|\vec{G}_i\|_2, \quad (5)$$

$$D_{ij}(\|\vec{G}_j\|_2) = w_{ij}(R) \frac{1}{\|\vec{G}_j\|_2}. \quad (6)$$

One notable feature about these equations is that $H_i(\cdot)$ is large when an agent in gang i is far from his or her gang's set space, but the $D_{ij}(\cdot)$ function is large when the agent is close to a rival gang j 's set space. The factors h_i and $w_{ij}(R)$ of the weighting functions are chosen according to the rules for agent movement. In our implementation, the factor $w_{ij}(R)$ depends on the current state of the rivalry network. Negative values of these functions result in repulsion and positive values result in attraction.




Figure 2: Cartoon example of the direction vectors incorporated in the direction of bias formula, Equation 4. The agent in this example is located at the dot. Here G_1 , G_2 , G_3 , and G_4 show the vectors pointing toward the set spaces of gangs 1 through 4, respectively. Depending on the choices of H_i and D_{ij} , different movement dynamics are possible.

After determining the direction of bias from Equation 4, we must choose in which direction the agent will move. The direction, θ , is drawn from a von Mises distribution (also known as the Circular Normal distribution) [33, 26, 5]. For $\theta \in [-\pi, \pi]$, the von Mises distribution is given by

$$f(\theta|\mu, \kappa) = \frac{\exp(\kappa \cos(\theta - \mu))}{2\pi I_0(\kappa)}.$$

Here I_0 is a modified Bessel function of order zero. The von Mises distribution requires two parameters, one for the angle of bias, μ , and one for the strength of the bias, κ . We can think of μ as being the mean of the distribution, and $\frac{1}{\kappa}$ as being comparable to the variance. The larger κ is, the stronger the bias is for the direction μ . If $\kappa = 0$, this is a uniform distribution on a circle.

From the direction and jump length, a prospective location is calculated. The new location is then checked to see if the result would move the agent into a different region. If it does not, the agent moves, meaning movement within a region is not restricted. However, if its next move would result in a region change, i.e. it is crossing a semi-permeable boundary, it has a given probability of crossing into that region. If the agent moves, it searches the other agents to see if it is close enough to interact with agents of other gangs. When an interaction does occur, the rivalry matrix, R , is updated.

The final network is observed after 20,000,000 iterations and then thresholded to ignore infrequent interactions. The location of interactions is also recorded and could be of interest to other applications, see discussion in Section 6 and Figure 12.

2.5. Initialization and Input Data

Before the simulation begins, the region map and an estimated density of the road networks must be provided in matrix form on the same grid. The probability of crossing each boundary must also be provided. Additionally, parameter values must be specified. Table 1 describes the full list of parameters needed for implementation. At the start of the simulation all of the agents are located at their gang’s set space. The size of each gang must also be specified.

Parameters	Acceptable Values	Hollenbeck Values	Tested Range	Description
Agent				
x_m	$0 < x_m < a$	0.1		Minimum jump length
k	$0 < k$	1.1	[1, 1.9]	Bounded Pareto scaling parameter
κ	$0 \leq \kappa$	3.5	[1.5, 5]	Von Mises scaling parameter
h_i	$-\infty < h_i < \infty$	1		Home weighting
$w_{ij}(R)$	$-\infty < w_{ij}(R) < \infty$	$-\rho_i(j)$		Rival gang weighting
Environment				
N_i	$N_i \in \mathbb{Z}^+$	$14 \leq N_i \leq 598$		Number of gang members in gang i
S_i	$S_i \in \mathbb{R}^2$	see Figure 1		Location of gang i set space
A	$a < A$	200	[100, 400]	Largest maximum jump length
a	$x_m < a < A$	100	[100, 200]	Smallest maximum jump length
B	$0 \leq B \leq 1$	0.2	[0, .5]	Permeability of boundaries
Network				
T	$0 \leq T$	0.04	[0, 0.6]	Threshold for existence of an edge

Table 1: Parameters needed for model implementation are listed in the first column. The second column lists theoretically acceptable parameter values. The values corresponding to the SBLN are displayed in the Hollenbeck Values Column. The Tested Range column provides the range for each variable for simulations run. The last column provides a description of each of the parameter values.

2.6. Hollenbeck parameters

The grid of environment features of Hollenbeck was approximated from the Google EarthTM image in Figure 1. Hollenbeck is about 39.4 km² [50, 64]. In our implementation, one Hollenbeck city block corresponds to approximately six grid elements. The interaction radius between agents is 3 units, or roughly half a city block. The approximated road density and region grids are show in Figure 3. The boundaries of the Hollenbeck region were approximated using points from the geographic features visible from Google EarthTM. These boundaries were used to construct the region grid. To approximate the road density of Hollenbeck we used a Weighted H¹ Maximum Penalized Likelihood Estimation method with a road map as the initial data [59]. To extend the approximated road density to the same sized grid as the region grid, the average value of the density over Hollenbeck was computed and used for the extended regions. The number of agents in each gang reflects historical information obtained from the LAPD.

The boundary crossing probability between the regions was calculated by the minimum number of boundaries one must cross to get from one region to the next. For instance, if region 1 and region 2 were separated by one boundary, the agent would have a probability, B , of accepting a move from region 1 to region 2. If region 1 and 2 were separated by α boundaries, then the agent would have a B^α probability accepting the move.

3. Baseline Comparison Models

3.1. Geographical Threshold Graphs (GTG)

For comparison to the networks produced by our simulations, we constructed an instance of a Geographical Threshold Graph (GTG). Geographical Threshold Graphs are random graphs that use spatial proximity




Figure 3: The image on the left shows the location of Hollenbeck in the $N \times M$ environment grid. The semi-permeable boundaries encoded in the model are displayed in the center image. The shades of gray of this image are used to distinguish among regions. On the right, we used a Weighted H^1 Maximum Penalized Likelihood Estimation method with a road map as the initial data to approximate the road density of Hollenbeck [59]. The scale, seen on the far right, gives the approximated road density intensity. Light shades of gray correspond to high density values near one and dark shades correspond to low densities near zero.

to assist in determining whether or not two nodes are connected with an edge [34, 7, 8]. Geographical Threshold Graphs randomly assign weights η_i to the N nodes. Then, using an interaction function $F(\eta_i, \eta_j)$, an edge between nodes n_i and n_j exists only if

$$\frac{F(\eta_i, \eta_j)}{d(n_i, n_j)^\beta} \geq \text{Threshold},$$

where $d(n_i, n_j)$ is the distance between nodes n_i and n_j . Constructing an instance of this graph is fast and computationally inexpensive. In our case, we take the multiplicative weight function $F(\eta_i, \eta_j) = \eta_i \cdot \eta_j$, since this is the number of possible pairings between members of gang i and gang j . We use Euclidean distance for the $d(n_i, n_j)$ function. The weights η_i are taken to be the size of each gang, and we choose $\beta = 2$. The threshold was chosen to give the same number of rivalries as the observed rivalry network.

3.2. Brownian Motion Network (BMN)

Another model we use to compare with the simulated network is a simplified version of the proposed model using Brownian Motion and unbiased movement rules. The semi-permeable boundaries of the model are incorporated also in this model. Specifically, each agent chooses the next prospective location from a standard normal distribution, ignoring any directional decisions. These simplifications reduce the number of variables to the threshold, T , and the permeability, B , while still incorporating the geographic boundaries. The parameter space around the Hollenbeck values was explored and run for $2 \cdot 10^7$ iterations. *A priori*, it was unclear how many iterations to run the simulation. We observed that the accuracy of the Brownian Motion networks peaked around $1.2 \cdot 10^7$ iterations and then decreased as the simulations progressed. The parameters and number of iterations that produced the highest accuracy were used for analysis. We will refer to the resulting network as the Brownian Motion Network (BMN).

Inherent in the BMN is a level of stochasticity. To understand how this stochasticity influences the final rivalry network and the resulting metrics, the BMN simulation was run for 100 different seed values. The resulting collection of final networks will be called the Ensemble BMN.

3.3. Baseline Network Graphs

Figure 4 displays the resulting GTG and BMN as compared to the observed rivalry network. The lower portion of the GTG graph has similar shape to the observed network, but contains more connections. The GTG does not make long connections. This is particularly evident in the upper half of Hollenbeck. The BMN picks up many of the longer connections, but includes far too many connections.

4. Results

The results of our network were obtained by searching the parameter space within the ranges specified in the fourth column of Table 1, allowing for dependencies between parameters. The 34,128 simulated networks




Figure 4: A visual comparison of the observed rivalry network (left), GTG (center), and BMN (right).

were then sorted according to accuracy, defined in Equation 7. Because each of the gangs in Hollenbeck are active, the graph with the highest accuracy with all non-zero degree nodes was chosen as a showcase of the model. The parameter values for the optimal run are found in the third column of Table 1. We will refer to this as Simulated Biased Lévy walk Network (SBLN). Figure 5 displays the network with our optimal parameters. The SBLN has a shape and structure similar to the observed network, but does not capture all of the longer edges. We also verified that all of the metrics we use to evaluate our model have reached a statistical equilibrium for the SBLN.




Figure 5: Comparison of the observed rivalry network (left) and the SBLN (right). The SBLN has a shape and structure similar to the observed network, but does not capture many of the longer edges.

4.1. Stochastic Effects Observed in the Simulated Biased Lévy walk Network (SBLN)

Implicit in the model is a degree of stochasticity intended to capture the gross features of human movement. In particular, the jump length and direction choice are sampled from probability distributions, and the directional bias is determined by the (inherently stochastic) current rivalry structure. These elements affect the inclusion and exclusion of rivalry network edges. To understand the effect of stochasticity on the network produced by the model, each simulation was run 100 times with different random seed values with the same SBLN parameter values. We refer to the collection of runs as the Ensemble SBLN. Each simulation was run independently and evaluated with several metrics. The resulting metrics were then averaged for analysis.

We also recorded the persistence of each edge in the ensemble of networks, and this is denoted as the *percent edge agreement*. For example, an ensemble network with 10% edge agreement refers to a network consisting of all edges that appear in at least 10% of the runs. Figure 6 displays the Ensemble SBLN with 100%, 50%, and 1% edge agreement next to the observed rivalry network. As expected, increasing the percent edge agreement decreases the number of edges present in the network. The network constructed with 100% edge agreement does not give a close representation of the observed network, because there are too few edges. However, allowing for 50% edge agreement produces a similar shape to the observed network. The Ensemble SBLN 1% edge agreement network shows all possible edges observed in the ensemble of simulation runs. Taken together, these images demonstrate the stochastic effects inherent in the model.



Figure 6: Percent edge agreement for the ensemble of runs for the SBLN parameter values. These four images give a comparison of, from left to right, the observed rivalry network, the Ensemble SBLN 1% edge agreement, the Ensemble SBLN 50% edge agreement, and the Ensemble SBLN 100% edge agreement.

For comparison, we simulated a random model that incorporates only the distance between nodes. In particular, we constructed a collection of randomly weighted Geographical Threshold Graphs by fixing the locations of the nodes and sampling the weights, η_i , independently from a uniform distribution. We selected a threshold to yield a median of 69 edges. Figure 7 displays the percent agreement of each possible edge for the Ensemble SBLN, a collection of randomly weighted Geographical Threshold Graphs, and the Ensemble BMN. For visualization, the edges for each ensemble were sorted separately in descending order based on percent edge agreement. In the Ensemble SBLN, there is 100% edge agreement for the existence of 39 of the edges (corresponding to the first 39 edges of the Ensemble SBLN along the horizontal axis in Figure 7). The 100% edge agreement network in Figure 6 shows these edges. All runs in the Ensemble SBLN consistently agree on the nonexistence of 309 edges (corresponding to the last 309 edges of the Ensemble SBLN in Figure 7). These are the edges not appearing in the 1% edge agreement network in Figure 6.

The transition between edge existence and nonexistence in the Ensemble SBLN is marked by a steep drop over 58 edges. The collection of randomly weighted Geographical Threshold Graphs displays a large degree of stochasticity indicated by fewer edges with 100% edge agreement and the more gradual decline of edge agreement. The Ensemble BMN appears to have a smaller degree of stochasticity with more edges

with 100% edge agreement and a steeper decline than the Ensemble SBLN and the collection of randomly weighted Geographical Threshold Graphs. Despite the stochasticity observed in these models, there is agreement among the edges of the Ensemble BMN and Ensemble SBLN, maintaining some structure within the simulated networks.

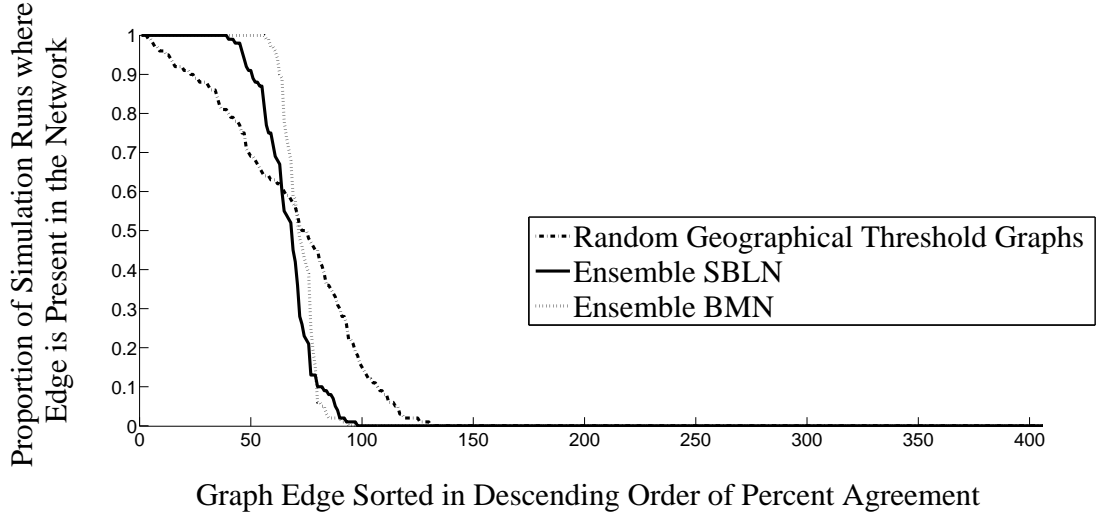


Figure 7: Plot of the edge persistence for the Ensemble SBLN (solid), Ensemble BMN (thin-dash), and an ensemble of random Geographical Threshold Graphs (thick-dash). The randomly weighted Geographical Threshold Graphs were constructed with random weights and have a median of 69 edges present. The edges were sorted in descending order according to the proportion of simulation runs where the edge is present in the network. Each ensemble of runs were sorted separately, yielding different edge numbers among ensembles.

4.2. Long Term Behavior of the SBLN

The simulated network, through the movements of each of the agents, evolves as the simulation progresses. Because of this evolution, it is natural to ask if any sort of steady state is achieved. Keeping in mind the stochasticity of the model and the interaction between the network and the agents' movements, an equilibrium in the strictest sense cannot be obtained. Despite this, the results indicate there is limiting behavior of the observed metrics as the simulation progresses. Figure 8 displays the density and accuracy over the progression of the simulations for the Ensemble SBLN; for the definition of these metrics, refer to Section 4.3. Each run is observed every 1,000 iterations and the results of each simulation are shown as a thin line. The average metric value at each iteration is calculated and plotted as the thick line. For visual investigation the vertical axis on the accuracy plot has been refined to include only the area of interest. Accuracy values can range from 0 to 1. Both of these plots suggest that after a short phase of initialization, the metrics of each run appear to stabilize. For the average values of the density and accuracy of the last iteration, refer to Table 2 and 3.

The Ensemble SBLN is shown to exhibit stable long term behavior the simulated rivalry network, with some variation due to stochasticity. Despite this variation, the network emerging from the model results in metrics with a small deviation from the average. Further, the stochasticity observed may provide a more realistic model of the true rivalry structure. Research has demonstrated that the rivalry networks that link gangs tend to be stable over time [63, 50, 47], and that the activity spaces of gangs are anchored to specific places [62, 39]. However, over longer periods of times, the membership ranks of gangs may ebb and flow due to incarceration, individuals "aging out" of active status, or other forms of incapacitation [64]. Thus, gangs may lay dormant and, though identified in the rivalry network, not actually participate in violence. In extreme cases, either through high levels of victimization at the hands of rival gangs or through the focused

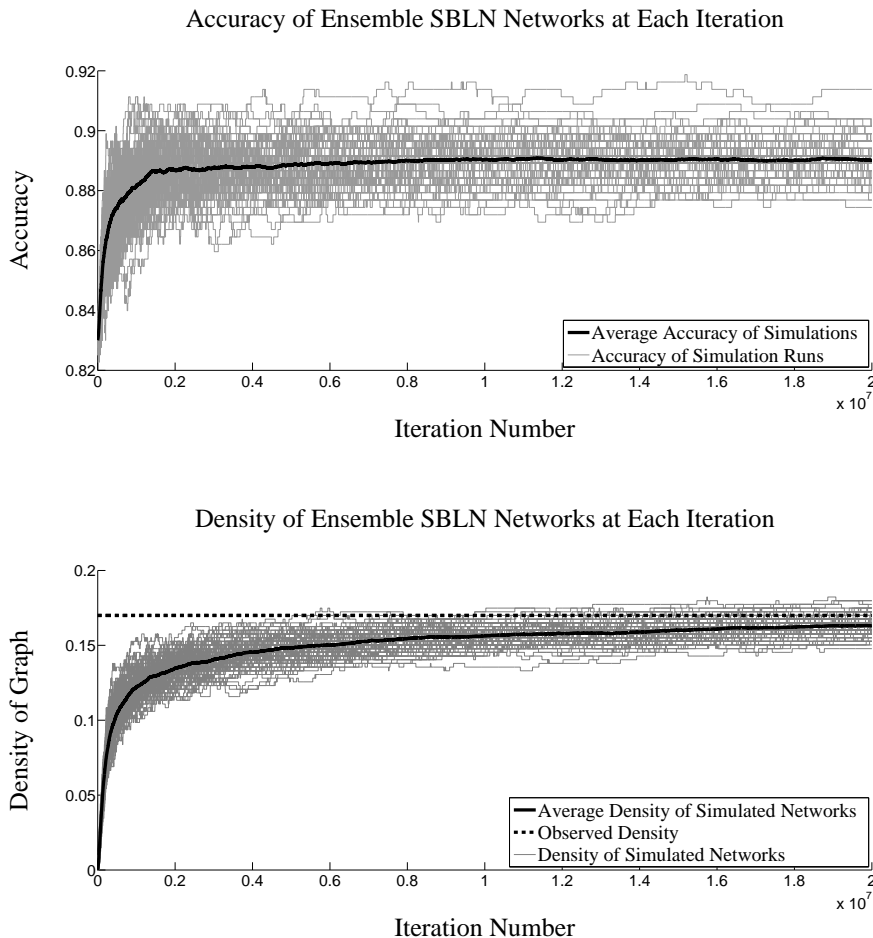


Figure 8: Plots of the accuracy (top) and the density (bottom) of the SBLN over the $2 \cdot 10^7$ iterations. Each of the 100 Ensemble SBLN runs are plotted by thin lines. The average over all the runs at each sampled iteration is shown with the solid, thick line. The density of the observed network is shown in the thick, dashed line. For visual investigation the vertical axis on the accuracy plot has been refined to include the area of interest. Accuracy values can range from 0 to 1.

enforcement of law enforcement agencies, a gang may simply disappear altogether. As more data become available, inherent stochasticity in the model may allow for further understanding of the rivalry structure.

4.3. Metrics Used for Analysis

We analyze our model according to several common metrics of accuracy, shape, and community structure. These statistics are compared to the observed rivalry network in [50], in which there are 69 rivalries among 29 active gangs in the Hollenbeck policing precinct.

4.3.1. Accuracy Metrics

The first measures of interest are the raw values for the number of correct and incorrect edges. These values provide a means for evaluating the performance of the model. However, when comparing the observed network with the constructed network, each edge can be correct in two ways and incorrect in two ways. First, the constructed network can correctly identify an edge, *true positive* (TP), and correctly identify the lack of an edge, *true negative* (TN). The constructed network can also be wrong in two different ways. It can place an edge where there is none, *false positive* (FP), and also fail to place an edge where there is one, *false negative* (FN).

There are three quantities that are of particular interest that summarize the TP, TN, FP, and FN values. First is the *accuracy* of the model. The accuracy in the context of edges on a graph is defined by

$$\text{ACC} = \frac{TP + TN}{TP + TN + FP + FN}. \quad (7)$$

The ACC ranges between 0 and 1, with 1 being a perfect reproduction of the observed network. This measure is proportional to the Q_α measure discussed in [3]. The *F1 score* provides another measure to analyze the accuracy of the predicted network, [54, 70], and is defined as

$$\text{F1} = \frac{2TP}{2TP + FP + FN}. \quad (8)$$

An exact replication of the network would have an F1 score of 1. The other summary statistic for the raw closeness to the network is the *Matthews Correlation Coefficient (MCC)* [35, 3]. This measurement varies between -1 and 1 , where a value of 1 is a perfect prediction. The MCC is defined as follows:

$$\text{MCC} = \frac{TP \cdot TN - FP \cdot FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}}. \quad (9)$$

The measurements of the TP, TN, FP, and FN provide one means by which to determine the success of the model. However, they do not describe how these correct or incorrect measurements affect the overall network structure. A strong model would create a network that is not only accurate but also, maintains the same network structure, even in the event that the individual connections are the not same.

4.3.2. Shape Metrics

We would like to verify that the simulated network has a similar shape to that of the true network. To do this, we calculate the graph *density*, standard *variance* of nodal degree and Freeman’s *centrality measure* of the graph. For definition of these metrics, see [67, 19]. The density of a network provides a normalized average of the degrees of the network. Networks with the same number of edges and nodes have the same density measure. The centrality measure of the graph is defined to be

$$\sum_{i=1}^N \frac{\text{maxDegree} - \text{degree}(i)}{(N - 1)(N - 2)}. \quad (10)$$

The centrality measure and the variance of the nodal degree provide measures for the spread of the degrees.

4.3.3. Metrics of Community Structure

Another class of measurements which are of interest to observers of social networks are those that describe a network’s community structure. Here, we use the idea of community structure strictly as a network property. The *degree distribution* has been widely used to understand the overall network structure [42, 44, 43, 1]. We compare the nodal degree cumulative distribution function (CDF) of our simulations with the observed network. The nodal *clustering coefficient* is another popular metric to analyze the community structure [1, 42, 67]. Intuitively, this is the proportion of a node’s neighbors that are also neighbors with one another to the total possible connections of this type. This measure is calculated for each node yielding a distribution of clustering coefficients. From this distribution, the mean clustering coefficient over all nodes is computed.

4.4. Evaluating Models using Graph Metrics

4.4.1. Accuracy Metric Results

Table 2 provides the accuracy measures for the GTG, BMN, Ensemble BMN, SBLN, and Ensemble SBLN. The SBLN outperforms all of the other networks on all of the accuracy metrics. Observe that the GTG also performs well on these metrics. The Ensemble SBLN metrics are comparable to the GTG and BMN metrics. In particular the average number of true negatives (TN) and false positives (FP) perform slightly better for the Ensemble SBLN than for the GTG, BMN, and Ensemble BMN. The Ensemble SBLN average of the true positives (TP) and false negatives (FN) performs slightly worse than the GTG and BMN. Only the GTG and SBLN have higher accuracy, F1 Score, and MCC values than the Ensemble SBLN average.

		TP	TN	FP	FN	ACC	F1 Score	MCC
SBLN		50	320	17	19	0.9113	0.7353	0.6822
Ensemble	Average	45.50	316.1	20.90	23.50	0.8906	0.6722	0.6069
SBLN	$\pm \sigma$	± 1.269	± 2.424	± 2.424	± 1.269	± 0.0077	± 0.020	± 0.025
GTG		48	316	21	21	0.8966	0.6957	0.6333
BMN		47	313	24	22	0.8867	0.6714	0.6031
Ensemble	Average	43.61	309.2	27.76	25.39	0.8691	0.6213	0.5424
BMN	$\pm \sigma$	± 1.380	± 1.390	± 1.390	± 1.380	± 0.0051	± 0.016	± 0.019

Table 2: Accuracy measures for the SBLN, Ensemble SBLN, GTG, BMN, and Ensemble BMN. The σ denotes the standard deviation of the ensemble metric values.

4.4.2. Shape Metric Results

Table 3 provides the shape measures for the observed network, GTG, BMN, Ensemble BMN, SBLN, and Ensemble SBLN. Note that the density of the GTG is exactly the same as the observed rivalry network by construction, but it does not perform well for the nodal degree variance. The density for the BMN, Ensemble BMN, SBLN, and Ensemble SBLN are all close to the observed network. The BMN and the Ensemble BMN average have the closest nodal degree variance to the observed network’s nodal degree variance. The centrality measure for the SBLN is the closest to that of the observed network.

		Density	Variance of Nodal Degree	Centrality	Mean Clustering Coefficient
Observed		0.16995	4.32105	0.20106	0.4921
SBLN		0.16503	3.54578	0.16799	0.6325
Ensemble	Average	0.16355	3.66423	0.15040	0.6364
SBLN	$\pm \sigma$	± 0.005593	± 0.48395	± 0.01883	± 0.02718
GTG		0.16995	9.97622	0.27778	0.6719
BMN		0.17488	3.88585	0.15741	0.7540
Ensemble	Average	0.17579	3.93926	0.16065	0.7009
BMN	$\pm \sigma$	± 0.004546	± 0.41351	± 0.02635	± 0.02681

Table 3: This table provides the shape measures for the observed network, SBLN, Ensemble SBLN, GTG, BMN, and Ensemble BMN. The σ denotes the standard deviation of the ensemble metric values. Note that the density of the GTG is exactly the same as the observed rivalry network by construction.

4.4.3. Community Structure Results

The cumulative distribution function (CDF) of nodal degree for the observed network, GTG, BMN, normalized Ensemble BMN, SBLN, and the normalized Ensemble SBLN are shown in Figure 9. A normalized ensemble CDF shows the CDF of the degree distribution of all runs divided by the number of runs. The SBLN and the normalized Ensemble BMN have the most similar distributions as the observed network. The normalized Ensemble SBLN performs better than the GTG and the BMN. In the same figure, the normalized Ensemble BMN and SBLN are plotted with two standard deviations above and below together with the observed network distribution. Here we see that there is a smaller standard deviation for the normalized BMN than the normalized SBLN. Even with the standard deviations, the degree distributions of both classes of networks are close to that of the observed degree distribution.

The mean clustering coefficient for each of the networks is seen in the last column of Table 3. On this measure of community structure, the SBLN and Ensemble SBLN average outperform all other networks. The Ensemble BMN average has the farthest mean clustering coefficient from the observed network.

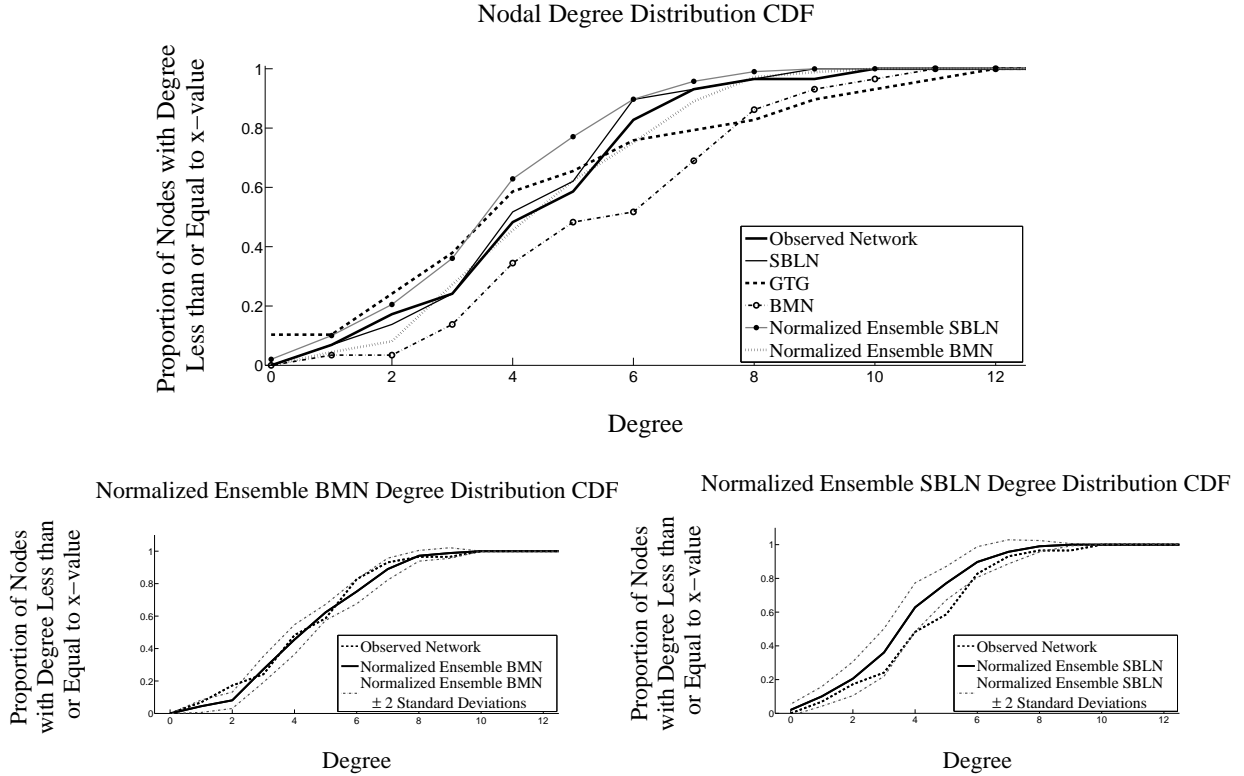


Figure 9: The top figure plots together the cumulative distribution functions of the degree distribution for the observed network (thick-solid), GTG (thick-dashed), BMN (dot-dash), normalized Ensemble BMN (thin-dash), SBLN (thin-solid), and normalized Ensemble SBLN (dot-solid). A normalized ensemble CDF shows the CDF of the degree distribution of all runs divided by the number of runs. The normalized Ensemble BMN (bottom left) and SBLN (bottom right) are plotted with two standard deviation above and below (thin-dash) with the observed network distribution (thick-dash).

4.5. Summary of Results

In all metrics except the density, the SBLN performs better than the GTG (note that the density measure of the GTG is exactly the same as the observed network by construction). Although the GTG is unable to closely replicate the standard shape measures, it has fairly high accuracy values. The Ensemble SBLN average performs similarly to the GTG in the accuracy, but performs better with shape measures, even with the stochastic considerations. On average, the Ensemble SBLN produces a slightly more accurate degree distribution than the GTG. The BMN is able to reproduce the degree distribution fairly well, however, the BMN and Ensemble BMN average have lower values for the accuracy (ACC), Matthews Correlation Coefficient (MCC), and F1 Score when compared to the other models. Our analysis demonstrates that the SBLN is the strongest model in reproducing the observed rivalry network.

5. Sensitivity Analysis

Our objective in this section is to understand the effects of the input parameters on the system by comparing the different metrics of the resulting networks as the parameters change. Due to computational constraints, we perform a local analysis of the parameter space around the SBLN parameters specified in column 3 of Table 1.

In particular, we perturb one parameter at a time by 30% from the SBLN parameter values in 10% increments. To account for the stochasticity inherent in the model, each perturbation was run using the same 25 seed values for the random number generator. The range of each parameter examined is listed in Table 4.

Bounded Pareto Scaling Parameter	$k \in [0.77, 1.43]$
Von Mises Scaling Parameter	$\kappa \in [2.45, 4.55]$
Largest Maximum Jump Length	$A \in [140, 260]$
Smallest Maximum Jump Length	$a \in [70, 130]$
Boundary Permeability	$B \in [0.14, 0.26]$
Network Threshold	$T \in [0.028, 0.052]$

Table 4: Ranges of the parameters used in the sensitivity analysis. Each parameter was changed 30% from the SBLN parameters in 10% increments. For SBLN parameter values refer to the Hollenbeck column of Table 1

For each simulation run, we compute the accuracy, Matthews Correlation Coefficient, F1 score, centrality measure, variance of nodal degree, density, and mean clustering coefficient for the resulting network. Plots of each combination of metric versus parameter values were created for the general analysis. Three examples of parameter and metric combinations with more dramatic results are plotted in Figure 10. In this figure, we display the variance of nodal degree versus the smallest maximum jump length, a , and the network threshold, T . We also display the density versus the Bounded Pareto scaling parameter, k , where the vertical axis has been rescaled for visualization. The dots represent the metric values of the simulation run at the specified parameter. The solid curve indicates the average metric value over all runs at each parameter value.




Figure 10: Plots of the nodal degree variance versus the smallest maximum jump length (top left), and the network threshold (top right). We also display the density versus the Bounded Pareto scaling parameter (bottom), where the vertical axis has been rescaled for visualization. The solid curve indicates the average metric value over all runs at each parameter value. The dots represent the metric values of the simulation run at the specified parameter.

As seen in Figure 10, the plots varying the network threshold and Bounded Pareto scaling parameters have a negative trend on average. The smallest maximum jump length, however, shows a positive trend. The stochastic effects can also be observed by the range of metric values associated with each parameter input, as illustrated by the dots in Figure 10. These plots suggest that stochasticity may influence the metric values for a particular run, and on average the resulting metric output is sensitive with respect to these parameters.

These plots give a view of how the particular metric and parameter value interact. We changed all of the parameter values by the same 30% from the SBLN parameters, and so we can compare plots with the

same metric. For example in Figure 10, we can see that in general nodal degree variance for the smallest maximum jump length has a steeper trend than the nodal degree variance for the threshold, but we can not compare the trend of the nodal degree variance plots directly to that of the density plot.

To compare the effects of all the parameters on all metrics, we rescale the data points to percent deviation from the SBLN parameter values. For example, when considering the affects of the Bounded Pareto scaling parameter, k , on the density metric, we rescaled the observed data points

$$(k_i, \text{density}_i) \mapsto \left(\frac{k_i - k_{\text{SBLN}}}{k_{\text{SBLN}}}, \frac{\text{density}_i - \overline{\text{density}}_{\text{SBLN}}}{\overline{\text{density}}_{\text{SBLN}}} \right),$$

where k_{SBLN} is the SBLN Bounded Pareto scaling parameter. Here, $\overline{\text{density}}_{\text{SBLN}}$ is the average density at the k_{SBLN} value for all 25 runs. A line was fitted to the rescaled data points, and the slope of this line was recorded. This process was repeated for each parameter and metric value combination.

The results are recorded in Table 5 and visualized in Figure 11. In Table 5, negative values indicate a negative slope of the best fit line to the scaled data, and positive values indicate a positive slope. Slopes with a greater magnitude indicate a stronger correlation between the metric and parameter. To get a clearer impression of overall sensitivity of the system, this information is displayed in Figure 11. The dark, and light, intensities of the color map represent large positive, and negative, values of the best fit line slope.

	k	κ	A	a	B	T
Accuracy	-0.0120	-0.0031	0.0011	0.0001	0.0031	0.0000
MCC	-0.2066	-0.0161	0.0023	0.1458	0.0293	-0.0000
F1 Score	-0.2149	-0.0146	0.0018	0.1562	0.0278	-0.0000
Centrality	-0.1705	-0.0100	-0.0131	0.7119	0.0195	-0.1751
Nodal Degree Variance	-0.4385	-0.0489	-0.0146	0.9456	0.0154	-0.5412
Density	-0.7410	-0.0080	-0.0114	0.6131	0.0640	-0.2460
Mean Clustering Coefficient	-0.3812	-0.0011	0.0005	0.1484	0.0084	-0.0614

Table 5: Slope of the best fit to the rescaled data for each metric and parameter combination. For reference, coefficients that correspond to the images in Figure 10 are highlighted in bold font. Figure 11 displays this information in a color map.

In general, the metrics are not very sensitive to the von Mises parameter, κ , the largest maximum jump length, A , and the boundary permeability, B , within the parameter space investigated. On the other hand, the Bounded Pareto scaling parameter, k , the smallest maximum jump length, a , and the network threshold, T , have the most influence on the metrics. As seen in the table and figure, the accuracy measures are fairly robust to changes in all parameter values. Further, note that nodal degree variance and density measures appear to be the most affected by the changes in these parameters.

The Bounded Pareto scaling parameter values result in negative slopes for all metrics. This is to be expected because an increase in the Bounded Pareto scaling parameter will decrease the likelihood of larger jumps and result in fewer edges. This phenomena is particularly evident in the density metric. Also this parameter appears to have the most effect on the accuracy measures, in particular the MCC and F1 score.

Increasing the network threshold parameter also has a negative effect on the shape and community metrics. By increasing the network threshold, the number of connections decreases. This in turn decreases the density, nodal degree variance, centrality, and mean clustering coefficient. On the other hand, increasing the smallest maximum jump length increases the connectivity of the network by allowing for larger jumps in areas of high road density. The effect of changing this parameter is more significant than changing the largest maximum jump length. Interestingly, as the largest maximum jump length increases, the connectivity decreases. This could be where attempts to cross boundaries are more likely to occur. The lower portion of Hollenbeck is approximately 300 units wide and has many boundaries. When varying the largest maximum jump length between 140 to 260, it becomes very probable that at least one boundary cross would be attempted. At this point, the boundary permeability is expected to play a stronger role in the simulation.

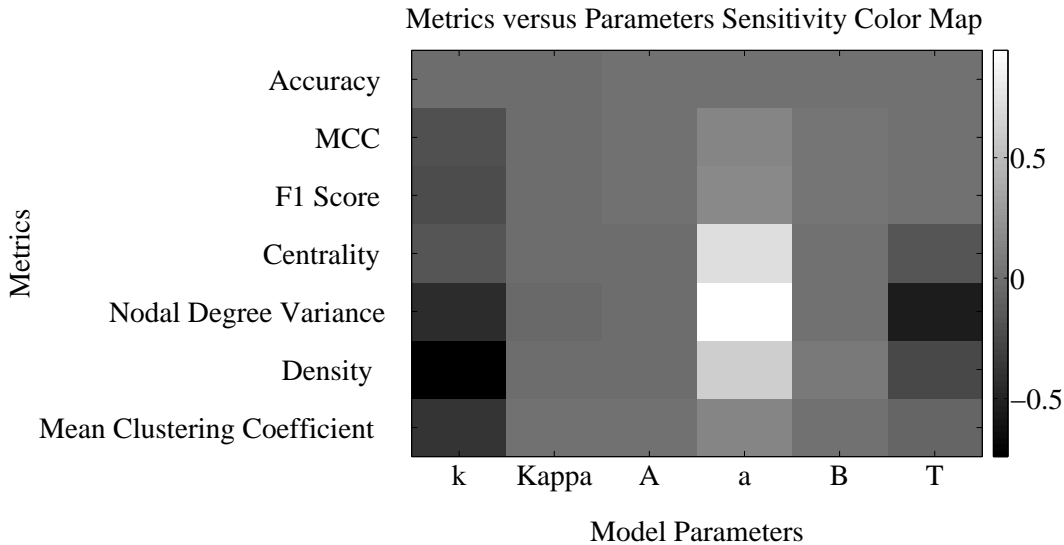


Figure 11: Slopes of the best fit line to the rescaled data for each of the parameter and metric combinations depicted in a color map. The parameters varied include the Bounded Pareto scaling parameter, k , the von Mises scaling parameter, Kappa , the largest maximum jump length, A , the smallest maximum jump length, a , the boundary permeability, B , and the network threshold, T . The scale to the right of the image gives the slope values. Tones close to the center of the scale represent combinations where the metrics are not very sensitive to the respective parameter. Combinations with tones at the ends of the spectrum (black and white) represent metrics that are sensitive to the respective parameter. The numerical values are also stored in Table 5.

Depending on the network, changes in the number of connections could be more or less beneficial in terms of accuracy. Further, small changes in the connectivity, i.e. the existence or non-existence of an edge, could have small effects on the accuracy measures and large effects on the shape measures, as seen for our simulations in the case of the network threshold parameter.

6. Discussion

Using biased truncated Lévy walks with semi-permeable boundaries, we have designed an agent-based model for gang members in Hollenbeck that incorporates quasi-realistic movement rules as well as physical geographic features existing in Hollenbeck. We have shown that it is able to simulate a gang rivalry network similar to the one observed in [64, 50]. The Simulated Biased Lévy walk Network (SBLN), the Brownian Motion Network (BMN), and an instance of a Geographical Threshold Graph (GTG) were compared to the observed rivalry network using measures of accuracy, shape, and community structure.

The GTG method performs well on the accuracy metrics and provides an alternative, computationally inexpensive method to construct the rivalry network. One could extend this model to incorporate boundary information by increasing the distance function $d(n_i, n_j)$ if n_i and n_j are in distinct regions, see Section 3.1. However, the GTG is limited to reproducing only the rivalry network and does not lend itself to understanding other phenomena, such as the gang member mobility and the locations of interactions between gang members. On the other hand, agent-based models such as the BMN and SBLN provide a reasonable approximation to the observed network and can be easily extended to include policing strategies, the location of violence, retaliatory behavior, and effects of injunctions.

Although the BMN accuracy results were not as strong as the GTG and SBLN results, this method was able to reproduce a similar shape and community structure as the observed rivalry network. This model was able to incorporate geographical features, but ignored directional decisions of the agents. One major problem with this model is that the stopping iteration for the model was artificial, in that we chose to stop it at the observed peak in accuracy. In general, there may not be an observed network, and so it would be difficult to determine stopping criteria. However, our proposed SBLN model exhibits long term stabilization of the accuracy and density metrics.

The SBLN is the best model in replicating the observed network. Further, it allows for easy incorporation of geographic features and alternate movement dynamics, while maintaining a high level of accuracy and allowing for evolution in the observed system. The success of this model and the flexibility of the method leads us to believe that the SBLN could also facilitate in understanding other social phenomena of interest related to gang violence. In fact, this model is able track the location of the agents' interactions during the simulation. This can be compared to violence data for the Hollenbeck area, and preliminary work has been done in this direction. Figure 12 shows the locations of the interactions among agents for one of the Ensemble SBLN simulation runs and the density of gang-related violent crimes in Hollenbeck from 1998 through 2000. The juxtaposition of these two plots emphasizes the similarities between the two and illustrates the potential predictive capabilities of this kind of approach. Though movement and interaction rules may need to be slightly altered to provide a closer match to the data, the current model provides a baseline model for further analysis and investigation of the gang rivalry violence in Hollenbeck. Another potential avenue for future work is to use this model to investigate territoriality and respect as a key elements in the motivation for violence [64].

The current model does not account for the difference between positive, negative, and neutral interactions. Instead the SBLN records interactions between agents with the implicit assumption that these are negative interactions. It is possible that there exists an alliance or truce between two gangs, and such phenomena has been observed in street gangs in Chicago [6]. The proposed SBLN model could be extended to include this.

Now that some of the influences of the geography and the interactions between the agents and the network are better understood, it could be beneficial to reformulate the agent-based model as a PDE. This alternative approach may allow for a deeper understanding of the model and may provide a rigorous analysis of the network dynamics.

Pursuing a model that accurately describes the violent behavior in Hollenbeck is of great value, since Hollenbeck is one of the most violent areas in Los Angeles [64, 25]. The advantage of approaching this serious problem using a computer simulation is twofold. First, these simulations may help us understand the underlying mechanisms that are involved in producing violent behavior among the gangs in Hollenbeck. Second, if the simulation can accurately model the social phenomena of interest, then we might gain some insight into how intervention strategies could alter the existing gang rivalry system. The costs of implementing these changes in the simulation are relatively small compared to those costs of public funds needed to implement experimental interventions. If the Hollenbeck area can be well understood by this approach, there may be hope in understanding, and potentially mitigating, other areas of intense violent behavior.

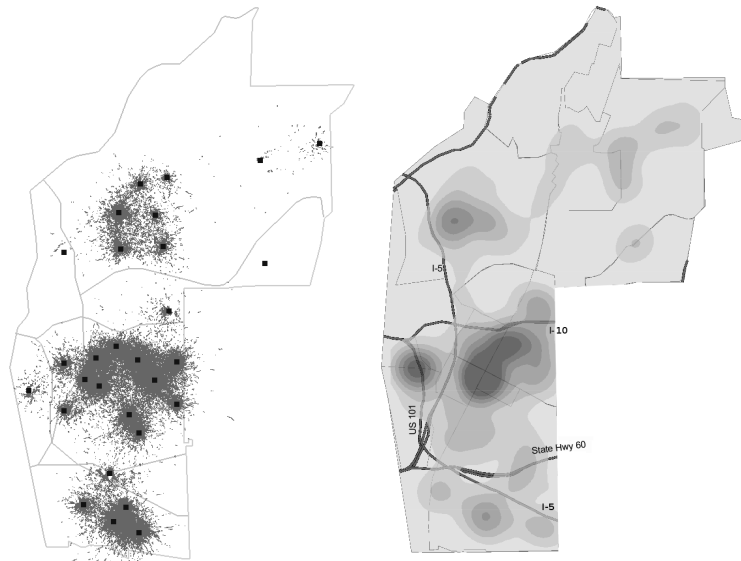


Figure 12: Locations of all the interactions between agents for one of the Ensemble SBLN runs (left). Density map of gang-related violent crimes in Hollenbeck between 1998 and 2000 (right).

7. Acknowledgments

We would like to thank Allon Percus for his useful feedback and Peterson Trethewey for his generous programming advice. We also would like to thank the LAPD, in particular Sean Malinowski, Gabriel Barboza, Arnold Suzukamo, Javier Macias, and the Hollenbeck Gang Division for their time and the enlightening conversations. In our work we used Google EarthTM images. This work was sponsored by NSF grant DMS-0907931, NSF grant DMS-0968309, AFOSR MURI grant FA9550-10-1-0569, ARO grant 58344-MA, ONR grant N000141010221, and ARO MURI grant 50363-MA-MUR.

References

- [1] R. Albert and A.-L. Barabási. Statistical mechanics of complex networks. *Rev. Mod. Phys.*, 74(1):47–97, January 2002.
- [2] E. Anderson. *Code of the Street: Decency, Violence, and the Moral Life of the Inner City*. 2000.
- [3] P. Baldi, S. Brunak, Y. Chauvin, C. A. Andersen, and H. Nielsen. Assessing the accuracy of prediction algorithms for classification: an overview. *Bioinform. Rev.*, 16(5):412–424, May 2000.
- [4] A. Barbaro, B. Einarsson, B. Birnir, S. Sigurdsson, H. Valdimarsson, O. K. Palsson, S. Sveinbjornsson, and T. Sigurdsson. Modelling and simulations of the migration of pelagic fish. *ICES J. Mar. Sci.*, 66(5):826–838, 2009.
- [5] D. J. Best and N. I. Fisher. Efficient simulation of the von Mises distribution. *J. R. Stat. Soc. Ser. C (Appl. Stat.)*, 28(2):152–157, 1979.
- [6] C. R. Block and R. Block. Street gang crime in Chicago. Technical report, December 1993.
- [7] M. Bradonjić, A. Hagberg, and A. G. Percus. Giant component and connectivity in geographical threshold graphs. In *Algorithms and Models for the Web-Graph (WAW 2007)*, pages 209–216, San Diego, CA USA, 2007.
- [8] M. Bradonjić, A. Hagberg, and A. G. Percus. The structure of geographical threshold graphs. *Internet Math.*, 5:113–140, 2009.
- [9] P. J. Brantingham and G. Tita. Offender mobility and crime pattern formation from first principles. In L. Liu and J. Eck, editors, *Artificial Crime Analysis Systems: Using Computer Simulations and Geographic Information Systems*, pages 193–208. Idea Press, 2008.
- [10] D. Brockmann, L. Hufnagel, and T. Geisel. The scaling laws of human travel. *Nature*, 439:462–465, 2006.
- [11] J. A. Canizo, J. A. Carrillo, and J. Rosado. A well-posedness theory in measures for some kinetic models of collective motion. 2010.
- [12] H. C. Covey. *Street Gangs Throughout the World*. Charles C Thomas, Publisher, 2010.
- [13] S. H. Decker and B. Van Winkle. *Life in the Gang: Family, Friends, and Violence*. Cambridge University Press, 1996.
- [14] M. R. D’Orsogna, Y.-L. Chuang, A. L. Bertozzi, and L. Chayes. Self-propelled particles with soft-core interactions: patterns, stability, and collapse. *Phys. Rev. Lett.*, 96, 2006. 104302.
- [15] M. Egesdal, C. Fathauer, K. Louie, and J. Neuman. Statistical modeling of gang violence in Los Angeles. *SIAM Undergrad. Res. Online*, 2010.
- [16] P. P. B. Eggermont and V. N. LaRiccia. *Maximum Penalized Likelihood Estimation*. Springer, 2001.
- [17] S. Eubank, H. Guclu, V. S. A. Kumar, M. V. Marathe, A. Srinivasan, Z. Toroczkai, and N. Wang. Modelling disease outbreaks in realistic urban social networks. *Nature*, 429:180–184, 2004.

- [18] J. D. Farmer and D. Foley. The economy needs agent-based modelling. *Nature*, 460:685–686, August 2009.
- [19] L. C. Freeman. Centrality in social networks conceptual clarification. *Soc. Netw.*, 1(3):215 – 239, 1978-1979.
- [20] M. T. Gastner and M. E. Newman. The spatial structure of networks. *Eur. Phys. J. B*, 49:247–252, 2006. 10.1140/epjb/e2006-00046-8.
- [21] M. González, C. A. Hidalgo, and A.-L. Barabási. Understanding individual human mobility patterns. *Nature Lett.*, 453:779–782, 2008.
- [22] I. J. Goodd and R. A. Gaskins. Nonparametric roughness penalties for probability densities. *Biometrika*, 58(2):255–277, 1971.
- [23] M. B. Gordon. A random walk in the literature on criminality: A partial and critical view on some statistical analyses and modelling approaches. *Eur. J. Appl. Math.*, 21(Special Double Issue 4-5):283–306, 2010.
- [24] C. K. Hemelrijk and K. Hanspeter. Density distribution and size sorting in fish schools: an individual-based model. *Behavioral Ecology*, 16(1):178–187, 2004.
- [25] Information Technology Division Management Report Unit. Statistical digest. *Los Angeles Police Department*, 2008.
- [26] S. R. Jammalamadaka and A. SenGupta. *Topics in Circular Statistics*. World Scientific Publishing Co. Pte. Ltd., 2001.
- [27] M. W. Klein, H.-J. Kerner, C. L. Maxson, and E. G. M. Weitekamp. *Euro Gang Paradox: Street Gangs and Youth Groups in the U.S. and Europe*. Kluwer Academic Publisher, 2001.
- [28] M. W. Klein, F. M. Weerman, and T. P. Thornberry. Street gang violence in Europe. *Eur. J. Criminol.*, 3(4):413–437, 2006.
- [29] D. Ley and R. Cybriwsky. Urban graffiti as territorial markers. *Ann. Assoc. Am. Geogr.*, 64(4):491–505, 1974.
- [30] D. Liben-Nowell, J. Novak, R. Kumar, P. Raghavan, and A. Tomkins. Geographic routing in social networks. *Proc. Natl. Acad. Sci.*, 102(33):11623–11628, August 2005.
- [31] W. Liu, M. B. Short, Y. E. Taima, and A. L. Bertozzi. Multiscale collaborative searching through swarming. *Proc. 7th Int. Conf. on Inform. Control, Autom., Robot. (ICINCO)*, June 2010.
- [32] M. W. Macy and R. Willer. From factors to actors: computational sociology and agent-based modeling. *Annu. Rev. Sociol.*, 28:143–166, 2002.
- [33] K. V. Mardia and P. E. Jupp. *Directional Statistics*. John Wiley & Sons Ltd, 2000.
- [34] N. Masuda, H. Miwa, and N. Konno. Geographical threshold graphs with small-world and scale-free properties. *Phys. Rev. E*, 71(3):036108, March 2005.
- [35] B. W. Matthews. Comparison of the predicted and observed secondary structure of T4 phage lysozyme. *Biochim. Biophys. Acta*, 405(2):442–451, 1975.
- [36] S. M. Mniszewski, S. Y. Del Valle, P. D. Stroud, J. M. Riese, and S. J. Sydoriak. EpiSims simulation of a multi-component strategy for pandemic influenza. In *SpringSim '08: Proceedings of the 2008 Spring simulation multiconference*, pages 556–563, San Diego, CA, USA, 2008. Society for Computer Simulation International.
- [37] G. Mohler. Gang rivalry dynamics via coupled point process networks. Unpublished results.

- [38] G. O. Mohler, A. L. Bertozzi, T. A. Goldstein, and S. J. Osher. Fast TV regularization for 2D maximum penalized likelihood estimation. *To appear in the J. Comput. Graph. Stat.*, 2009.
- [39] J. Moore and R. Pinderhughes, editors. *In the Barrios: Latinos and the Underclass Debate*. Russel Sage Foundation, 1993.
- [40] M. E. J. Newman. Scientific collaboration networks. I. network construction and fundamental results. *Phys. Rev. E*, 64(1):016131, June 2001.
- [41] M. E. J. Newman. Scientific collaboration networks. II. shortest paths, weighted networks, and centrality. *Phys. Rev. E*, 64(1):016132, June 2001.
- [42] M. E. J. Newman. The structure of scientific collaboration networks. *Proc. Natl. Acad. Sci.*, 98(2):404–409, 2001.
- [43] M. E. J. Newman. The structure and function of complex networks. *SIAM Rev.*, 45(2):167–256, 2003.
- [44] M. E. J. Newman, S. H. Strogatz, and D. J. Watts. Random graphs with arbitrary degree distributions and their applications. *Phys. Rev. E*, 64(2):026118, July 2001.
- [45] M. E. J. Newman, D. J. Watts, and S. H. Strogatz. Random graph models of social networks. *Proc. Natl. Acad. Sci.*, 99(suppl. 1):2566–2572, February 2002.
- [46] M. O’Leary. The mathematics of geographic profiling. *J. Invest. Psychol. Offender Profiling*, 6:253–265, 2009.
- [47] A. V. Papachristos. Murder by structure: Dominance relations and the social structure of gang homicide. *Am. J. Sociol.*, 115(1):74–128, 2009.
- [48] A. B. Pitcher. Adding police to a mathematical model of burglary. *Eur. J. Appl. Math.*, 21(Special Double Issue 4-5):401–419, 2010.
- [49] Product No. 2009-M0335-001. National gang threat assessment. *Natl. Gang Intell. Cent.*, January 2009.
- [50] S. Radil, C. Flint, and G. Tita. Spatializing social networks: Using social network analysis to investigate geographies of gang rivalry, territoriality, and violence in Los Angeles. *Ann. Assoc. Am. Geogr.*, 100(2):307–326, 2010.
- [51] I. Rhee, M. Shin, S. Hong, K. Lee, and S. Chong. On the Lévy-walk nature of human mobility: Do humans walk like monkeys? In *IEEE INFOCOM 2008 - IEEE Conference on Computer Communications*, pages 924–932. IEEE, April 2008.
- [52] F. Schweitzer. *Brownian Agents and Active Particles: Collective Dynamics in the Natural and Social Sciences*. Springer, 2003.
- [53] F. Schweitzer and B. Tilch. Self-assembling of networks in an agent-based model. *Phys. Rev. E*, 66(2):026113, August 2002.
- [54] F. Sebastiani. Machine learning in automated text categorization. *ACM Comput. Surv.*, 34(1):1–47, 2002.
- [55] M. B. Short, P. J. Brantingham, A. L. Bertozzi, and G. E. Tita. Dissipation and displacement of hotspots in reaction-diffusion models of crime. *Proc. Natl. Acad. Sci.*, 107(9):3961–3965, 2010.
- [56] M. B. Short, M. R. D’Orsogna, V. B. Pasour, G. E. Tita, P. J. Brantingham, A. L. Bertozzi, and L. B. Chayes. A statistical model of criminal behavior. *Math. Models Methods Appl. Sci.*, 18(suppl.):1249–1267, 2008.
- [57] B. W. Silverman. *Density Estimation for Statistics and Data Analysis*. Chapman & Hall/CRC, April 1986.

- [58] H. M. Singer, I. Singer, and H. J. Herrmann. Agent-based model for friendship in social networks. *Phys. Rev. E*, 80(2):026113, August 2009.
- [59] L. Smith, M. Keegan, T. Wittman, G. Mohler, and A. Bertozzi. Improving density estimation by incorporating spatial information. *EURASIP Adv. Signal Process: spec. issue Adv. Image Process. Def. Secur. Appl.*, 2010.
- [60] L. Tesfatsion. Chapter 16 agent-based computational economics: A constructive approach to economic theory. volume 2 of *Handbook of Computational Economics*, pages 831–880. Elsevier, 2006.
- [61] F. M. Thrasher. *The Gang: A Study of 1313 Gangs in Chicago*. University of Chicago Press, 1927.
- [62] G. Tita, J. Cohen, and J. Engberg. An ecological study of the location of gang “set space”. *Soc. Probl.*, 52(2):272–299, 2005.
- [63] G. Tita and S. Radil. Making space for theory: The challenges of theorizing space and place for spatial analysis in criminology. *J. Quant. Criminol.*, pages 1–13, 2010. 10.1007/s10940-010-9115-5.
- [64] G. Tita, K. Riley, G. Ridgeway, C. Grammich, A. Abrahamse, and P. Greenwood. Reducing gun violence: Results from an intervention in East Los Angeles. *Natl. Inst. Justice, RAND*, 2003.
- [65] Z. Toroczkai and H. Guclu. Proximity networks and epidemics. *Physica A: Stat. Mech. Appl.*, 378(1):68 – 75, 2007.
- [66] J.-J. Tseng, C.-H. Lin, C.-T. Lin, S.-C. Wang, and S.-P. Li. Statistical properties of agent-based models in markets with continuous double auction mechanism. *Physica A: Stat. Mech. Appl.*, 389(8):1699–1707, 2010.
- [67] S. Wasserman and K. Faust. *Social Network Analysis: Methods and Applications*. Cambridge University Press, 2009.
- [68] B. Wellman. Are personal communities local? A Dumptarian reconsideration. *Soc. Netw.*, 18(4):347 – 354, 1996.
- [69] A. Wilhite. Chapter 20 economic activity on fixed networks. volume 2 of *Handbook of Computational Economics*, pages 1013–1045. Elsevier, 2006.
- [70] Y. Yang and X. Liu. A re-examination of text categorization methods. In *SIGIR '99: Proceedings of the 22nd annual international ACM SIGIR conference on research and development in information retrieval*, pages 42–49, New York, NY, USA, 1999. ACM.