A Cellular Automaton SIS Epidemiological Model with Spatially Clustered Recoveries

David Hiebeler

Dept. of Mathematics and Statistics,
333 Neville Hall, University of Maine,
Orono, ME 04469-5752 USA
hiebeler@math.umaine.edu

Abstract. A stochastic two-state epidemiological cellular automaton model is studied, where sites move between susceptible and infected states. Each time step has two phases: an infectious phase, followed by a treatment or recovery phase. During the infectious phase, each infected site stochastically infects its susceptible neighbors. During the recovery phase, contiguous blocks of sites are reset to the susceptible state, representing spatially clustered treatment or recovery. The spatially extended recovery events are coordinated events over groups of cells larger than standard local neighborhoods typically used in cellular automata models. This model, which exhibits complex spatial dynamics, is investigated using simulations, mean field approximations, and local structure theory, also known as pair approximation in the ecological literature. The spatial scale and geometry of recovery events affects the equilibrium distribution of the model, even when the probability of block recovery events is rescaled to maintain a constant per-site recovery probability per time step. Spatially clustered treatments reduce the equilibrium proportion of infected invididuals, compared to spatially more evenly distributed treatment efforts.

1 Introduction

Consider a discrete-time lattice-based epidemiological model, where each site can be in one of two states: susceptible and infected. Infection and recovery parameters are ϕ and μ, respectively. During each time step, the following two things occur, in this order:

- Infection: every infected site will infect each of its susceptible neighbors, independently with probability ϕ each. The standard von Neumann neighborhood was used, consisting of the four orthogonal neighbors of a site. If an infected site tries to infect an already-infected neighbor, there is no effect.
- Recovery: contiguous blocks of sites recover simultaneously. Parameters b_1 and b_2 specify the dimensions of recovery blocks. Each block will consist of a $b_1 \times b_2$ (rows $\times$ columns) block of sites or a $b_2 \times b_1$ block of sites, each with

V.S. Sunderam et al. (Eds.): ICCS 2005, LNCS 3515, pp. 360–367, 2005.

probability 0.5. During the recovery phase, each site, independently with probability γ (computed from μ as described in the section "Pair Approximations" below), will have a recovery block placed so that its upper-left corner is located at the target site being considered. Note that multiple recovery blocks within a time step may spatially overlap.

This is a discrete-time analogue of a continuous-time population model investigated in [1]. While qualitative results are similar in the two models, many of the details of the analysis differ, being more complex for the discrete-time model.

2 Simulations

Simulations were performed on a 300×300 lattice with wraparound (toroidal) boundary conditions. After each time step, the proportions of sites in the susceptible and infected states were recorded. Based on methods used by [2] to test for equilibrium, beginning on time step 1000, on every time step a least-squares regression line was fit to the most recent 100 measurements of the proportion of sites infected. When the slope of this line was less than 0.001, and the difference between the minimum and maximum proportion of infected sites over the previous 100 time steps was less than 0.03, the system was considered to have reached equilibrium. After it was determined that equilibrium was reached, the simulation was run for another 50 time-steps, and the proportion of infected sites was averaged over those final 50 time steps and recorded as the final proportion of infected sites for the simulation. Although exploration showed that the model was not sensitive to initial conditions, in order to reduce the time needed to reach equilibrium, the equilibrium predicted by the local-dispersal mean-field approximation [3] was used as the initial proportion of infected individuals. However, if this initial proportion was less than 0.1, then 0.1 was used instead, to prevent fixation to a lattice completely in the susceptible state solely due to fluctuations from an initial small population of infected sites.

3 Pair Approximation

Let 0 represent the susceptible state, and 1 the infected state. The state of the lattice can be approximately described by the probabilities $P[ij]$ (where $i, j \in \{0, 1\}$) that a pair of adjacent sites are in the state configurations 00, 01, 10, and 11. Assuming rotational symmetry, $P[01] = P[10]$ (as well as $P\begin{bmatrix} i \\ j \end{bmatrix} = P[ij]$), and using the fact that the four probabilities must sum to one, only two independent probabilities are needed to describe the state of the system, for example $P[00]$ and $P[01]$. The other two probabilities may then be computed as $P[10] = P[01]$ and $P[11] = 1 - P[00] - P[01] - P[10] = 1 - P[00] - 2P[01]$. Marginal probabilities of the possible states for a single site can be recovered by summing over block probabilities, $P[i] = P[i0] + P[i1] \quad \text{for } i \in \{0, 1\}$.

As described in [3], based on ideas explored in [4, 5], the block probabilities $P_{t+1}[ij]$ at time $t+1$ can be estimated using the current probabilities $P_t[ij]$ by first estimating the probabilities of all pre-images of a pair of sites, and then applying the cellular automaton rule to those pre-images and using the law of total probability. A pre-image of a pair of sites is a set of state configurations of the group of sites which the target pair of sites depends on when updating its states, as shown in Fig. 1.

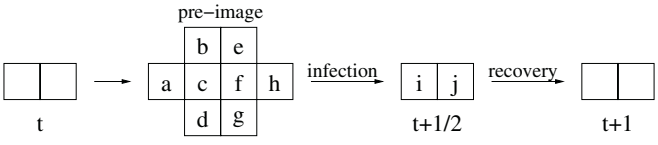

Fig. 1. The group of sites in a pre-image of a pair of sites is shown. A pre-image is the set of states of all sites which the target pair of sites depend on when updating their states, i.e. all neighbors of the pair of target sites. The probabilities of all pre-images are estimated, and then used to compute probabilities of all state configurations of a pair of sites after the infectious phase of a time step, $P_{t+\frac{1}{2}}[ij]$. The probabilities after the recovery phase, $P_{t+1}[ij]$ are then computed

The pair approximation used here assumes that non-adjacent sites are independent when conditioned on any shared neighbors, i.e.

$$P\begin{bmatrix} j & k \\ i & \cdot \end{bmatrix} = P[ijk] = P[ij\cdot]P[\cdot\,k|ij\cdot] = P[ij]P[\cdot\,k|\,\cdot\,j\cdot] = \frac{P[ij]P[jk]}{P[j]}. \quad (1)$$

In the expression above, the conditional probability that the third site is in state k given the states of the other two sites are i and j does not depend on the first site's state because the first and third sites are not adjacent. See e.g. [6, 7] for a discussion of these methods applied to continuous-time epidemiological models. Note that hereafter, $0/0$ is defined to be 0 in (1) when extending block probabilities, since if $P[j] = 0$, then $P[ij] = 0$ and $P[ijk] = 0$ for all $i, k \in \{0, 1\}$.

Heuristically, the probability of a 3×1 block may be built up by covering it with two overlapping 2×1 blocks, multiplying the 2×1 block probabilities and dividing by the overlapping single-site probability. The 2×1 probabilities may be repeatedly extended in this manner to build up probabilities of ever-larger blocks [3, 8]. However, as also seen with many information-theoretic measures of spatial complexity [9], in two or more dimensions, there can be more than one way to cover larger blocks of sites with 2×1 sites [3], and thus there is not a unique way to estimate the probabilities of pre-images. This can be seen when trying to compute the probability of a 2×2 block:

$$\begin{aligned} P\begin{bmatrix} a & b \\ c & d \end{bmatrix} &= P\begin{bmatrix} a & b \\ \cdot & d \end{bmatrix} P\left[\begin{array}{cc|cc} \cdot & \cdot & a & b \\ c & \cdot & \cdot & d \end{array}\right] = \frac{P[ab]P[bd]}{P[b]} P\left[\begin{array}{cc|cc} \cdot & \cdot & a & \cdot \\ c & \cdot & \cdot & d \end{array}\right] \\ &= \frac{P[ab]P[bd]}{P[b]} \frac{P\begin{bmatrix} a & \cdot \\ c & d \end{bmatrix}}{P\begin{bmatrix} a & \cdot \\ \cdot & d \end{bmatrix}} = \frac{P[ab]P[bd]}{P[b]} \frac{P\begin{bmatrix} a & \cdot \\ c & d \end{bmatrix}}{\sum_{i\in\{0,1\}} P\begin{bmatrix} a & \cdot \\ i & d \end{bmatrix}} \end{aligned} \quad (2)$$

where the first probability has been expanded using (1), and the second probability has been approximated by assuming that the site labelled c does not depend on the non-adjacent site labelled b, and then expanding the results using the elementary definition of conditional probability. The sum in the denominator of the final expression may be calculated using (1). The non-uniqueness of this method may be seen by observing that in the calculation above, c was the "last" site added into the block when expanding the 2×2 block probability using a conditional probability; if instead d were the last site considered, a different expression would result. There is no clear way to choose one method over the other; one could choose whichever term maximizes entropy of the resulting block probabilities, but (2) was used in this study.

Because of the nature of the update rule used, computation of the new probabilities $P_{t+1}[ij]$ in terms of the current probabilities $P_t[ij]$ may also be broken into the two phases of infection and recovery.

For the infectious phase, the probabilities of all possible pre-images as shown in the center of Fig. 1 are needed. Following the discussion above, the probability extension used here was

$$P\begin{bmatrix} \cdot & b & e & \cdot \\ a & c & f & h \\ \cdot & d & g & \cdot \end{bmatrix} = \frac{P[cf]P[bc]P[fg]P[ac]P[fh]P[be]P[ef]P[cd]P[dg]}{(P[c])^2(P[f])^2P[e]P\begin{bmatrix} b & \cdot \\ \cdot & f \end{bmatrix}P[d]P\begin{bmatrix} c & \cdot \\ \cdot & g \end{bmatrix}} \tag{3}$$

where

$$P\begin{bmatrix} b & \cdot \\ \cdot & f \end{bmatrix} = \sum_{i\in\{0,1\}} P\begin{bmatrix} b & \cdot \\ i & f \end{bmatrix} = \sum_{i\in\{0,1\}} P[bif] = \sum_{i\in\{0,1\}} \frac{P[bi]P[if]}{P[i]}$$

and similarly for $P\begin{bmatrix} c & \cdot \\ \cdot & g \end{bmatrix}$. The probabilities of pre-images given by (3) may therefore be estimated using only the current 2×1 block probabilities.

Once the pre-image probabilities have been estimated, the probabilities $P_{t+\frac{1}{2}}[ij]$ after the infectious phase of the time step may then be computed, by conditioning on the pre-image at time t:

$$P_{t+\frac{1}{2}}[ij] = \sum_{G\in\mathcal{G}} P_t(G)P(G \to [ij]) \tag{4}$$

where $\mathcal{G}$ is the set of all pre-images, $P_t(G)$ is the probability of pre-image G, and $P(G \to [ij])$ is the probability that pre-image G results in the state $[ij]$ for the target pair of sites after the infectious phase. Because there are 8 sites in the pre-image, and two states per site, there are $2^8 = 256$ pre-images in total. Because only infections occur during this phase, if $c = 1$ and $i = 0$, or $f = 1$ and $j = 0$ in Fig. 1, then $P(G \to [ij]) = 0$. Otherwise, the probability will be based on binomial distributions. Let $k_L(G) = a+b+d+f$ be the number of neighbors of the left site c which are occupied in the pre-image G, and $k_R(G) = c+e+g+h$ be the number of neighbors of the right site f which are occupied.

- If $c = 0$ and $f = 0$, then $P(G \to [00]) = (1-\phi)^{k_L(G)}(1-\phi)^{k_R(G)}$, and $P(G \to [01]) = (1-\phi)^{k_L(G)}(1-(1-\phi)^{k_R(G)})$.
- If $c = 0$ and $f = 1$, then $P(G \to [00]) = 0$, and $P(G \to [01]) = (1-\phi)^{k_L(G)}$.
- If $c = 1$ and $f = 0$, or if $c = 1$ and $f = 1$, then $P(G \to [00]) = P(G \to [01]) = 0$.

Once the probabilities $P_{t+\frac{1}{2}}[ij]$ have been estimated using (4) together with the above information, the final probabilities $P_{t+1}[ij]$ may then be estimated by applying the recovery phase of the cellular automaton rule. Because the application of recovery blocks is externally imposed and does not depend on the current states of cells or their neighbors, extension of block probabilities is not needed for this phase. Only the application of basic probability is needed, to compute the probabilities that among a pair of sites, neither, one, or both sites are contained within a recovery block. Sites may be part of a recovery block if any of several nearby sites are the target of such a block. For example, consider the case where $b_1 = 2$ and $b_2 = 3$, i.e. 2×3 and 3×2 recovery blocks are used. For the pair of sites drawn in bold in Fig. 2, recovery blocks at any of the labelled sites will affect one or both sites in the pair, as follows:

- Both 2×3 and 3×2 blocks targetted at sites A will affect only the right site of the pair.
- Both 2×3 and 3×2 blocks targetted at sites B will affect both sites of the pair.
- 3×2 blocks at site C will affect the right site of the pair, but 2×3 blocks will not affect the pair.
- 3×2 blocks at site D will affect both sites of the pair, but 2×3 blocks will not affect the pair.
- 2×3 blocks at sites E will affect both sites of the pair, but 3×2 blocks will affect only the left site of the pair.
- 2×3 blocks at sites F will affect the left site of the pair, but 3×2 blocks will not affect the pair.
- 3×2 blocks at site G will affect the left site of the pair, but 2×3 blocks will not affect the pair.

Similar enumerations can be performed for any values of b_1 and b_2. This information may then be used to calculate the probabilities that particular sites in a pair are affected by one or more blocks. Such calculations show that the probability that both sites in a pair will be affected by one or more recovery blocks is

$$P([11] \to [00]) = 1 + A^{c_1} B^{c_2} \left(A^{c_3} B^{c_4} - 2\right) \tag{5}$$

where $A = 1-\gamma$, $B = 1-\gamma/2$, $c_1 = (b_{\min})^2$, $c_2 = 2(b_{\max} - b_{\min})b_{\min}$, $c_3 = b_{\min}$, and $c_4 = b_{\max} - b_{\min}$, $b_{\min} = \min(b_1, b_2)$ and $b_{\max} = \max(b_1, b_2)$. Similarly, the probability that the left site in a pair will be affected by one or more recovery blocks, but that the right site will not be affected by any blocks, is

$$P([11] \to [01]) = A^{c_1} B^{c_2} \left(1 - A^{c_3} B^{c_4}\right) \tag{6}$$

Combining the two, the probability that any single site will be affected by one or more recovery blocks is

$$P([1] \to [0]) = 1 - A^{c_1} B^{c_2} \tag{7}$$

We wish the single-site recovery probability to be μ, but because a site may recover due to being hit by recovery blocks targeted at any number of neighboring sites, the recovery probability is altered. To correct for this, we use (7) to define $f(\gamma) = 1 - (1-\gamma)^{c_1}(1-\gamma/2)^{c_2}$, and then numerically solve for the value γ satisfying $f(\gamma) = \mu$. This value γ is the adjusted recovery probability used in all simulations and approximations of the model; it is the recovery block probability which yields a single-site recovery probability of μ per time step.

	G	D	C
F	E	B	A
F	E	**B**	**A**

Fig. 2. The set of all sites where a 2×3 recovery block could be targeted and affect a specific pair of sites (shown in bold at lower-right corner). See text for explanation of site labels

This adjustment of recovery rates is simpler in continuous-time models, where the recovery rate merely needs to be divided by $b_1 b_2$, the size of the recovery blocks [1]. In the discrete-time model, however, if this rescaling is used, the number of recovery blocks affecting a single site approaches a Poisson distribution as the block sizes become large, and the single-site recovery probability approaches $1 - e^{-\mu}$, thus making the more complex rescaling above necessary.

The recovery probabilities above in (5)–(7) may be combined with the infection probabilities given by (4) to compute the updated probabilities on the next time step, as follows:

$$P_{t+1}[00] = P_{t+\frac{1}{2}}[00] + 2P_{t+\frac{1}{2}}[01]P([1] \to [0]) + P_{t+\frac{1}{2}}[11]P([11] \to [00])$$
$$P_{t+1}[01] = P_{t+\frac{1}{2}}[01](1 - P([1] \to [0])) + P_{t+\frac{1}{2}}[11]P([11] \to [01])$$

4 Results

Equilibrium proportions of sites infected are shown for $\phi = 0.5$, as the per-site recovery rate μ was varied between 0 and 1, for $n \times n$ recovery blocks in Fig. 3 and $1 \times n$ recovery blocks in Fig. 4. Results are shown for simulations, pair approximations, and the mean-field approximation, for which the rate of recovery only depends on the single-site rate given by (7). Note that the mean-field approximation is not dependent on the size of recovery blocks because it is a spatially implicit method which ignores all spatial correlations; thus only one mean-field curve appears in the figures. Errors in the predictions, i.e. the pair

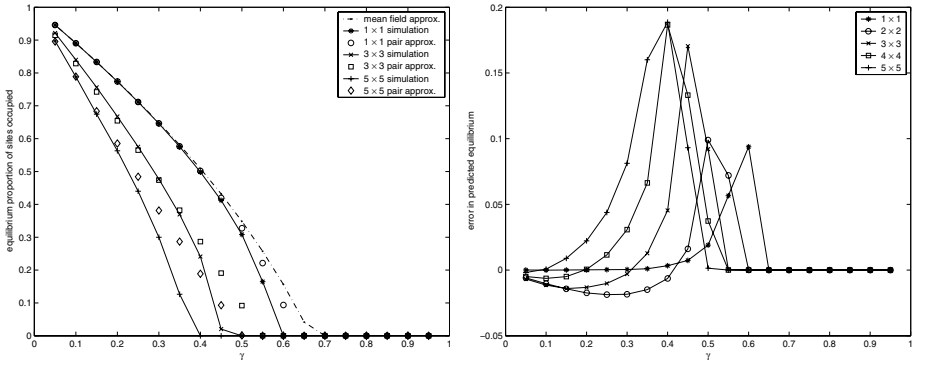

(a) (b)

Fig. 3. Results using square $n \times n$ recovery blocks of various sizes, as the per-site recovery rate μ varies between 0 and 1 on the x-axis, with $\phi = 0.5$. (a) The equilibrium proportion of infected sites is shown, from simulations, pair approximations, and the mean field approximation. (b) Prediction error, i.e. predictions from pair approximations minus measurements from simulations

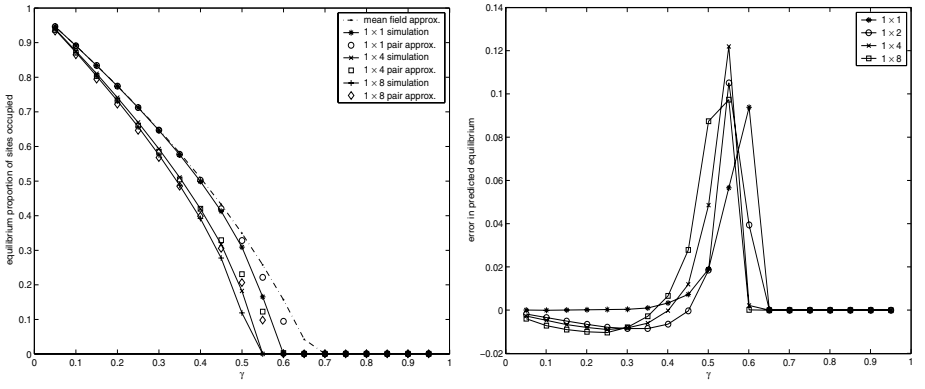

(a) (b)

Fig. 4. Results using long $1 \times n$ recovery blocks of various sizes. Compare with Fig. 3. (a) Equilibrium proportion of infected sites. (b) Pair approximation prediction error

approximation minus the simulation measurements, are also shown in Figs. 3 and 4.

It can be seen from the figures that the geometry of recovery events does affect the equilibrium distribution, even when the single-site recovery probability is held constant. This effect is more pronounced for square recovery blocks than for long, narrow blocks. For a disease with only local infection and a regular treatment regime, the long-term prevalence of the disease would be reduced by focusing treatment in fewer contiguous areas, rather than distributing treatment more evenly throughout the population.

As with the continuous-time version of the model, the pair approximations do fairly well at predicting simulation results. They do most poorly near the critical value of the recovery rate at which the equilibrium proportion of infected individuals transitions between 0 and a positive value, when spatial correlations decay more slowly with distance over spatial scales beyond that reflected by the pair approximation [1]. Also, as can be seen in Fig. 3b, the pair approximations become less accurate as the spatial scale of the recovery blocks becomes larger. The pair approximations are more accurate for long $1 \times n$ blocks as compared with square $n \times n$ blocks (compare the scales of the y-axes in Figs. 3b and 4b), and are also more accurate over a wider range of values of the rescaled recovery rate γ. Although in continuous time, pair approximations tend to almost always overestimate the equilibrium proportion of infected sites, in the discrete-time model it can be clearly seen from the figures that the pair approximation underestimates this value over a significant range of the parameter space. Further investigation is needed to determine exactly why the more complex interactions in the discrete-time model give rise to this behavior.

References

1. Hiebeler, D.: Spatially correlated disturbances in a locally dispersing population model. Journal of Theoretical Biology **232** (2005) 143–149
2. Caswell, H., Etter, R.J.: Ecological interactions in patchy environments: From patch occupancy models to cellular automata. In Levin, S., Powell, T., Steele, J., eds.: Patch Dynamics, Springer-Verlag (1993) 93–109
3. Hiebeler, D.: Stochastic spatial models: From simulations to mean field and local structure approximations. Journal of Theoretical Biology **187** (1997) 307–319
4. Gutowitz, H.A., Victor, J.D.: Local structure theory in more than one dimension. Complex Systems **1** (1987) 57–68
5. Wilbur, W.J., Lipman, D.J., Shamma, S.A.: On the prediction of local patterns in cellular automata. Physica D **19** (1986) 397–410
6. Levin, S.A., Durrett, R.: From individuals to epidemics. Philosophical Transactions: Biological Sciences **351** (1996) 1615–1621
7. Filipe, J., Gibson, G.: Comparing approximations to spatio-temporal models for epidemics with local spread. Bulletin of Mathematical Biology **63** (2001) 603–624
8. Gutowitz, H.A., Victor, J.D., Knight, B.W.: Local structure theory for cellular automata. Physica D **28** (1987) 18–48
9. Feldman, D.P., Crutchfield, J.P.: Structural information in two-dimensional patterns: Entropy convergence and excess entropy. Physical Review E **67** (2003)