CS 420/594 Fall 2007 Project 2

Creation of Spatial Structure by Activation/Inhibition Cellular Automaton

Revised 2007-10-19

Due October 25, 2007

Background

Activation/Inhibition Cellular Automaton

In this project you will investigate and measure the creation of spatial structure by an activation/inhibition CA (AICA) such as we discussed in class (Lecture 8). Structure will be quantified in terms of *correlation length*, *joint entropy*, and *mutual information* (discussed below), which you will measure after the CA has converged to a stable state.

Recall that the state transition rule of an activation/inhibition network is given by a formula such as this:

$$s_{\mathbf{i}}(t+1) = \text{sign}\left[h + J_1 \sum_{r_{\mathbf{ij}} < R_1} s_{\mathbf{j}}(t) + J_2 \sum_{R_1 \le r_{\mathbf{ij}} < R_2} s_{\mathbf{j}}(t)\right]. \quad (1)$$

Since this is a 2D CA, the cell indices are two-dimensional vectors, $\mathbf{i} = (i_1, i_2)$, $\mathbf{j} = (j_1, j_2)$. As usual, we will also assume that the space is a torus, that is, the indices wrap around on the top and bottom and on the left and right edges. For the purposes of this project you may assume that $J_1 \ge 0$ and $J_2 \le 0$ (which are the usual cases). Also, you should assume that the R_1 neighborhood includes the cell ($\mathbf{i}$) at its center.

The distance $r_{\mathbf{ij}}$ between cells can be measured in many different ways, for example by Euclidean distance (also known as the L_2 metric). For the purpose of this project it will be more convenient to measure distance by the L_1 metric, which is defined:

$$r_{\mathbf{ij}} = |i_1 - j_1| + |i_2 - j_2|. \quad (2)$$

This gives neighborhoods that are diamond-shaped, rather than circular, but that doesn't really matter.

Correlation Length

We are interested in the extent to which that states of cells at various distances are correlated to each other. Therefore, we define the *absolute correlation at distance* ℓ as

$$\rho_\ell = \left|\langle s_\mathbf{i} s_\mathbf{j}\rangle - \langle s_\mathbf{i}\rangle\langle s_\mathbf{j}\rangle\right| \text{ for all } \mathbf{i},\ \mathbf{j} \text{ such that } r_{\mathbf{ij}} = \ell.^1$$

The angle brackets mean "average value of." To understand this formula, suppose that the states +1 and −1 are equally frequent in the space. Then the individual state averages are 0: $\langle s_\mathbf{i}\rangle = \langle s_\mathbf{j}\rangle = 0$. Therefore $\langle s_\mathbf{i} s_\mathbf{j}\rangle$ is the average value of the product $s_\mathbf{i} s_\mathbf{j}$ for cells that are a distance ℓ apart. If these cells tend to be +1 at the same time, or −1 at the same time, this average will be greater than zero (positive correlation). If they tend to have opposite signs, then it will be less than zero (negative correlation). If they tend to have the same sign as often as the opposite sign, then the average will be near zero (no correlation). By subtracting $\langle s_\mathbf{i}\rangle\langle s_\mathbf{j}\rangle$ we compensate for an overall bias toward positive or negative states (such as we get when $h \neq 0$). We take the absolute value, because we are not interested in whether the correlation is positive or negative, only its magnitude. Note that $\rho_0 = 1 - \langle s_\mathbf{i}\rangle^2$.

Often correlation decreases approximately exponentially with distance, $\rho_\ell \propto e^{-\ell/\lambda}$, where λ, the characteristic *correlation length*, measures how quickly correlation decreases with distance. By assuming that correlation is exponentially decreasing, we can estimate λ. Let α be an arbitrary constant of proportionality $\rho_\ell = \alpha e^{-\ell/\lambda}$. Then,

$$\rho_0 = \alpha e^0 = \alpha,$$

$$\rho_\lambda = \alpha e^{-\lambda/\lambda} = \alpha e^{-1} = \rho_0 \big/ e.$$

Therefore we can estimate $\lambda \approx \ell$ such that $\rho_\ell = \rho_0 / e$. That is, the correlation length is that length at which the correlation has decreased to $1/e$ of its maximum value ρ_0. (*Correlation time* can be defined similarly.)

---

[1] Note: A true correlation coefficient is normalized to [–1, 1] by dividing by standard deviations of the variables: $\rho_\ell = \left|\left(\langle s_\mathbf{i} s_\mathbf{j}\rangle - \langle s_\mathbf{i}\rangle\langle s_\mathbf{j}\rangle\right) \big/ \sigma_{s_i}\sigma_{s_j}\right|$ for all $\mathbf{i}$, $\mathbf{j}$ such that $r_{\mathbf{ij}} = \ell$. You can compute it this way if you want, but make sure to tell us.

Next let's consider more explicitly how to compute ρ_ℓ. Suppose there are N^2 cells in the space, and let C_ℓ $(1 \le \ell < N/2)$ be the circumference (in number of cells) of a neighborhood of radius ℓ. For the L_1 metric, $C_\ell = 4\ell$. Thus there are C_ℓ cells at a distance ℓ from a given cell. Then we can see (make sure you really do see it!) that:

$$\rho_\ell = \left| \frac{2}{N^2 C_\ell} \sum_{\substack{\langle \mathbf{ij} \rangle \\ r_{\mathbf{ij}} = \ell}} s_\mathbf{i} s_\mathbf{j} - \left(\frac{1}{N^2} \sum_\mathbf{i} s_\mathbf{i} \right)^2 \right|.$$

The notation $\langle \mathbf{ij} \rangle$ under the summation means "all pairs of distinct cells **i** and **j**" (taking each pair just once); therefore there are $N^2/2$ of these pairs. Thus we are averaging over all pairs at a distance of ℓ. For purposes of computation, this can be made more explicit:

$$\rho_\ell = \left| \left(\frac{1}{4\ell N^2} \sum_{i_1, i_2} \left(s_{i_1, i_2} \sum_{|j_1|+|j_2|=\ell} s_{i_1+j_1, i_2+j_2} \right) \right) - \left(\frac{1}{N^2} \sum_{i_1, i_2} s_{i_1, i_2} \right)^2 \right|.$$

Notice that the second summation is over positive and negative j_1, j_2 in the range $-\ell \le j_1, j_2 \le \ell$. The coefficient on the first summation is:

$$\frac{1}{2} \frac{2}{N^2} \frac{1}{C_\ell} = \frac{1}{4\ell N^2}.$$

The $\frac{1}{2}$ factor is to avoid double counting the pairs $\langle \mathbf{ij} \rangle$. Make sure that you understand the formula for ρ_ℓ. Remember that cell indices wrap around both vertically and horizontally.

Mutual Information Between Distant Cells

As discussed in class (Lecture 4), another way to measure the correlation between cell states is by *average mutual information*. The average mutual information between cells at a distance ℓ is related to the joint entropy between cells at this distance. Therefore, first define the average entropy $H(S)$ of the cellular space S:

$$H(S) = -\sum_{s \in \{-1,+1\}} \Pr\{s\} \lg \Pr\{s\}.$$

Remember that we take $0 \lg 0 = 0$. For state counting convenience, let $\beta(s) = (1+s)/2$ so that $\beta(+1) = 1$ and $\beta(-1) = 0$ (β converts a *bipolar* number $\in \{-1,+1\}$ into a *binary* number $\in \{0,1\}$). Then the probabilities of the states can be computed:

$$\Pr\{+1\} = \frac{1}{N^2} \sum_\mathbf{i} \beta(s_\mathbf{i}),$$

$$\Pr\{-1\} = 1 - \Pr\{+1\}.$$

The probability that two cells at distance ℓ have state +1 is:

$$P_\ell\{+1,+1\} = \frac{2}{N^2 C_\ell} \sum_{\substack{\langle \mathbf{ij} \rangle \\ r_{\mathbf{ij}} = \ell}} \beta(s_\mathbf{i})\beta(s_\mathbf{j}).$$

Similarly,

$$P_\ell\{-1,-1\} = \frac{2}{N^2 C_\ell} \sum_{\substack{\langle \mathbf{ij} \rangle \\ r_{\mathbf{ij}} = \ell}} \beta(-s_\mathbf{i})\beta(-s_\mathbf{j}).$$

These summations can be computed in the same way as ρ_ℓ. Finally,

$$P_\ell\{+1,-1\} = P_\ell\{-1,+1\} = 1 - P_\ell\{+1,+1\} - P_\ell\{-1,-1\}.$$

(Notice that there are only three distinct possibilities for cells at distance ℓ: both +1, both –1, or opposite signs.) The joint entropy between cells at a distance ℓ, H_ℓ, is then defined in terms of the probabilities in the usual way:

$$H_\ell = -\left(P_\ell\{+1,+1\}\lg P_\ell\{+1,+1\} + P_\ell\{-1,-1\}\lg P_\ell\{-1,-1\} + P_\ell\{+1,-1\}\lg P_\ell\{+1,-1\}\right).$$

The *average mutual information* between two sources A and B is defined $I(A,B) = H(A) + H(B) - H(A,B)$. Therefore, the average mutual information between cells at distance ℓ is defined:

$$I_\ell = 2H(S) - H_\ell.$$

Experiments

In this project you will be investigating how spatial structure (as measured by correlation length, mutual information, etc.) is affected by the parameters (J_1, J_2, R_1, R_2, h) of an AICA.

It is difficult to anticipate exactly how long these computations will take. Therefore I think it is important to remain flexible and keep the objectives in mind. The goal is to get the measures of structure for (hopefully) 10–15 different combinations of parameters for each of the following three experiments. Therefore, try to allocate your runs intelligently so that you can test your hypotheses about the relation of the parameters to these measures.

Your CA space should be $N = 30$ in each dimension. Use asynchronous updating (i.e., update the cells one at a time rather than all at once).

For each set of parameters you will make several runs, with different random initial states, so that you can average your results. Do a few trial runs to see how much the measures vary from run to run. If they don't vary much, then you don't need to do many runs for each set of parameters (maybe one run will be enough). If they do vary, then you will have to make several runs (less than 5) and average them together, and you won't be able to investigate so many parameter values.

For each run you should allow the CA to stabilize, and then compute $H(S)$, estimated λ, and ρ_ℓ, H_ℓ, and I_ℓ for $\ell = 0,\ldots,N/2$. Do not compute these

measures for every ℓ value in this range; try to identify the ranges at which interesting things are happening, and do more measurements in those areas. In any case, 10 ℓ values for each parameter set are probably enough. For each set of parameter values, compute the average of all of these and print them out (generate plots for ρ_ℓ, H_ℓ, and I_ℓ). You don't need to hand in a graph for every set of parameters, just those that are relevant to your hypotheses.

Following are some descriptions of specific experiments that you should try. In each case, discuss and explain your results.

1. Set $J_2 = 0$ to disable the inhibition system and set $J_1 = 1$. Quantify the spatial structure for a range of R_1 values (in the range $0 < R_1 < N/2$ and at least three h values in the range $-6 \le h \le 6$ including $h = 0$.
2. Set $J_1 = 0$ to disable the activation system and set $J_2 = -0.1$. Quantify spatial structure as in Experiment (1) but varying R_2.
3. Set $J_1 = 1$ and $J_2 = -0.1$. Investigate, as in the preceding experiments, the spatial structure for a variety of values of R_1, R_2, and h.

Based on all the preceding experiments, draw conclusions about the dependence of correlation length, joint entropy, and mutual information on the parameters of the AICA. See if you can draw any quantitative conclusions.

Convergence of AICA

This part of the project is **required for CS 594** students and **extra-credit for CS 420 students**. Consider again the state update rule (Eq. 1, p. 1). Recall that r_{ij} represents the distance between cells $\mathbf{i}$ and $\mathbf{j}$, so the first summation is over all cells within a distance of R_1 to cell $\mathbf{i}$, and the second summation is over all cells with a distance from R_1 to R_2 (Eq. 2, p. 1). For simplicity in this part of the project, assume that the R_1 neighborhood does *not* include the center cell $\mathbf{i}$.

The state of a CA can be updated either *synchronously* or *asynchronously*. Recall that with synchronous updating all the states are updated simultaneously. With asynchronous updating the cells are updated one at a time (usually in some random order).

This part of the project explores the convergence of this AICA; that is, does it inevitably reach a stable state?

Problems

Problem 1

Prove that if the states are updated asynchronously, then the AICA must reach a stable state.

Hint: Define the following function (called an *energy* or *Lyapunov* function) of the total state of an AICA:[2]

$$E\{\mathbf{s}(t)\} = -\tfrac{1}{2}\sum_{\mathbf{i}} s_{\mathbf{i}}(t)\left[2h + J_1 \sum_{r_{ij}<R_1} s_{\mathbf{j}}(t) + J_2 \sum_{R_1 \le r_{ij} < R_2} s_{\mathbf{j}}(t)\right].$$

Show that updating any single cell, according to the state update rule, cannot increase this function (that is, $\Delta E \le 0$). What else do you need to show in order to guarantee convergence to a stable state?

Extra Credit: Assume that the R_1 neighborhood *does* include the center cell, and explore any additional assumptions that might be needed to guarantee convergence.

Problem 2

Prove, by exhibiting a counter-example, that if synchronous updating is used, then the AICA may not reach a stable state.

Hint: Construct a very simple AICA, obeying the above state update equation (Eq. 1), that cycles between two different states.

[2] For this energy function, look in Bar-Yam on p. 630 (section 7.2.2) and p. 170 (sec. 1.6.6).