What Can the LGA Universe Compute?

Andrew Haddock and Scott Wisdom

Electrical Engineering Department, University of Washington, Seattle, USA

***Abstract*— The second law of thermodynamics dictates that an isolated system will evolve towards a state of maximum entropy. Simulation engines such as the lattice gas automata (LGA) allow for the study of systems moving towards equilibrium by simulating micro- and macro-scale dynamics. In the last few decades, many problems in fields such as image processing have been formulated as constrained maximum entropy optimization problems, suggesting that LGA methods could provide a new method for solving such problems. However, the proper methods for representing arbitrary data in the LGA and how the LGA can be used to solve constrained maximum entropy optimization problems is not clear. This project aims to answer these questions through a combination of passive and active simulation methods in the LGA. In the passive scenario, chambers are initialized with an amount of particles corresponding to some data value; we then want to see if by running the system to equilibrium we can infer that data value. In the active scenario, chambers are initialized with an amount of particles proportional to gray-scale pixel value, and we develop a method for constraining maximum entropy to reconstruct images from blurred measurements with missing data.**

I. INTRODUCTION

In this paper we will present two applications of the principle of maximum entropy applied to problems in the LGA. First we will leverage the LGA to investigate what a system approaching equilibrium, or maximum entropy, can tell us about its initial state. This leads to an investigation of diffusion processes and flow rates for a series of connected chambers. We will show that by measuring flow rates we can determine both the initial state of the system as well as its maximum entropy solution. An application of this is, for example, determining the relationship between a set of data points. By measuring flow rates between chambers we know something about the initial state, and by running the system to equilibrium we can say something about how the initital state of data corresponds to the average state of all data.

We will refer to this first method as ”passive” because the system itself is not being perturbed; we are just measuring its properties. In our second method, we take an ”active” approach, so named because we take measurements and then perturb the system depending on the values of these measurements. Such an active approach allows us to apply constraints to the system, which is useful for performing constrained optimization problems when the objective function to be maximized is entropy. Using an active approach, we develop and show results for a maximum entroy image reconstruction in the LGA.

II. PASSIVE METHODS FOR DATA ESTIMATION

The LGA, absent of modeling physical properties such as temperature or gravity, is capable of representing the maximum entropy solution of a uniform random variable. For data encoded as particles in a series of connected chambers, as in Figure 1, the maximum entropy solution will then be obtained as the average number of particles in each chamber. Therefore, maximum entropy can encode the average of a set of data, but can it in some way encode individual data values?

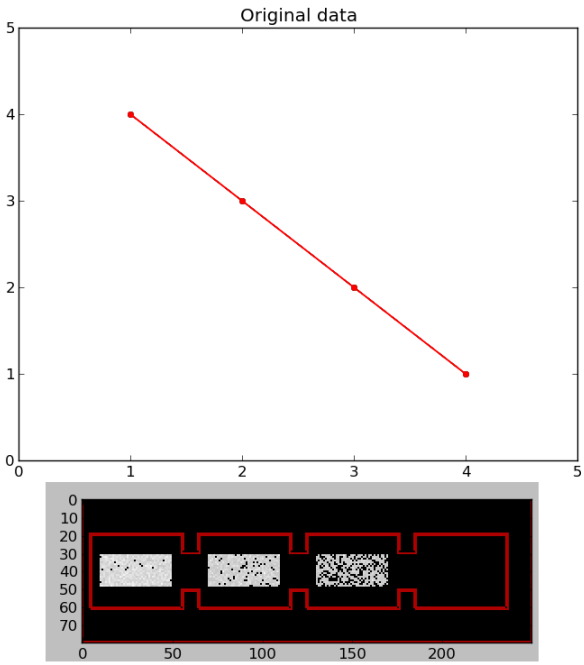

Fig. 1: Conceptual illustration of how arbitrary data could be encoded in particles initialized in seperate chambers in LGA. In each chamber, the number of particles initially is proportional to the value of the data relative to other data in the set.

To answer this question, we can look at flow rates between interconnected chambers as a system approaches equilibrium (equal density in each chamber). The relationship between particle flow rates and concentration gradients is captured by Fick's Laws of Diffusion [6], which state that the two measures are related by a proportional constant, D:

$$\frac{\mathrm{d}\phi}{\mathrm{d}t} = D\frac{\mathrm{d}^2\phi}{\mathrm{d}x^2} \tag{1}$$

where ϕ is the concentration of the diffusing substance. Therefore, if we are able to measure the flow rates of particles between a set of chambers, we should be able to determine how many particles are in one chamber relative to the number of particles in the other chamber. Such a local measurement that is representative of the overall state of the system could be useful in constructing a "demon" or reversible computation engine wherein initial states are preserved through an irreversible process. However, whether or not this is possible in the LGA is not clear, so the following experiments were performed to confirm Fick's Laws.

A. Flow Rate Experiments

We tested Fick's Laws of Diffusion in the LGA using the simple 2-chamber setup shown in Figure 2. In this experiment, we varied the number of particles initialized in the left chamber and ran the system to equilibrium, estimating the flow rate along the way. We estimated the flow rate by counting the number of particles in a chamber between two time steps and dividing by the time interval. It is worth noting that a shorter statistical computing step size actually increases the noise in the flow rate estimate. In our experiments we set the statistical step size as Δt = 25 seconds so as to implicitly smooth out the flow rate estimate.

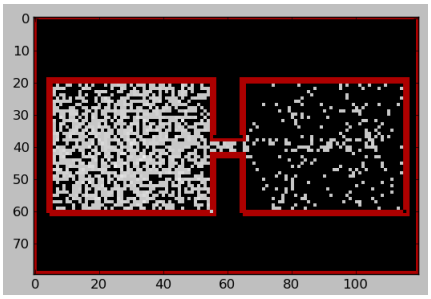

Fig. 2: The 2-chamber LGA simulation that was used to validate Fick's Laws of Diffusion. The cross-sectional area (or, vertical dimension) of the tube connecting chambers is denoted as w.

Plotting the estimated flow rate against the difference in number of particles between each chamber we obtain Figure 3, which shows a linear relationship with the slope given by D in equation 1. By performing this same experiment and setting a smaller cross-sectional area of the tube, w, we find that the slope, D, also decreases. For a smaller opening we have that over the same range of concentration gradients the flow rate is lower; thus D captures how difficult it is for particles to move from one chamber to the next.

Next, we considered whether the same method of estimating flow rates could be applied to a more practical statistical computing area: the tube connecting the chambers. Running the same experiment for an area with w = 10, we found that the Fick's Law relationship plotted in Figure 4 still holds but with a less tight correlation.

The results of these experiments show that flow rate measurements can be used to estimate particle gradients in

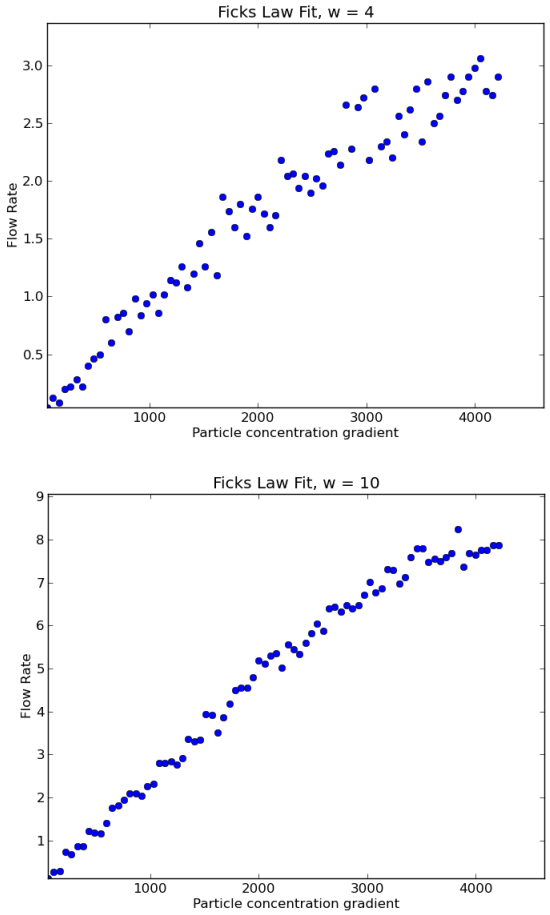

Fig. 3: Fick's Law fits for w = 4 (top) and w = 10 (bottom). This plot shows that there is a linear relationship between flow rate and concentration gradient, and that the slope, D, of this relationship can be increased by increasing the cross-sectional area of the tube connecting the chambers.

the LGA based on Fick's Laws of Diffusion. These results have a direct analog to describing current and voltage in resistive elements in electric circuits. The plots in Figure 3 can be viewed as I-V curves, where flow rate and concentration gradient are analogous to current and voltage, respectively. However, the possible saturation effect at higher concentrations may suggest that particles flowing through a constrained volume behave in a way more analogous to the operation of MOSFET devices, which exhibit saturation voltages. The existence of this saturation region also implies that there will be a practical limit to the range of data values that can be encoded in this manner.

B. Data Inference in the LGA

We return to the question of inferring encoded data values by measuring flow rates between connected chambers in a system approaching equilibrium. We have shown that by measuring the flow rates we can estimate the particle concentration gradient. An expression for determining the

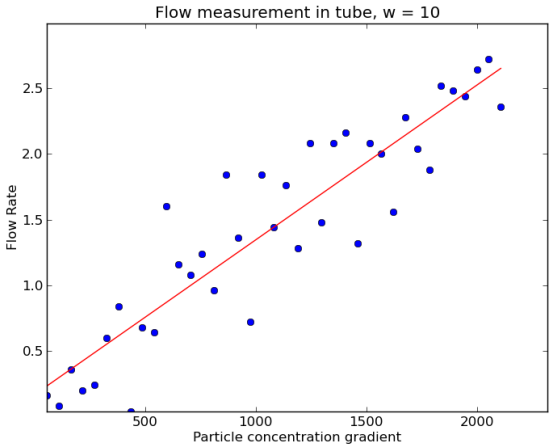

Fig. 4: Fick's Law fit with flow rate measurements performed in the tube connected the chambers. The linear relationship predicted holds but with a less correlated structure.

relative concentration gradient between two chambers is the following:

$$\frac{\mathrm{d}n_1}{\mathrm{d}t} = K(C_2 - C_1) \tag{2}$$

To expand this expression to encompass a series of connected chambers as in Figure 1 we can note that flow is conserved across the system and that the flow in each chamber is a net flow determined by flow between each adjacent chamber. We can then express the flow out of the jth chamber in Figure 1 with the following equation:

$$\frac{\mathrm{d}n_i}{\mathrm{d}t} = \frac{\mathrm{d}n_{i,i+1}}{\mathrm{d}t} - \frac{\mathrm{d}n_{i-1,i}}{\mathrm{d}t} = K(C_{i+1} - C_i) - K(C_i - C_{i-1}) \tag{3}$$

Now, if we take flow measurements from each chamber we can set up a system of equations in the form $\mathrm{A}\mathbf{x} = \mathbf{b}$ and find $\mathbf{x}$ the particles initialized in each chamber. However, it ends up being the case that A is a singular matrix for the case where n connected chambers are encoding n unique data points, which renders the system without a unique solution for $\mathbf{x}$. This is due to the fact that the densities shown in figure 1 have no absolute reference for baseline concentration. To remedy this problem, we used a different method for encoding data in chambers, with one central reference chamber to which all concentration gradients can be compared; this is shown in figure 5. This setup allows for unique estimation of $\mathbf{x}$, the data encoded in each chamber, which is provided in figure 5 with respect to the original data shown in figure 1. Thus we have shown how one can use the LGA to encode arbitrary data values and naively identify them using the relationship provided by Fick's Laws.

C. Discussion

A natural question to consider at this point is whether this method for estimating data values encoded in the LGA by way of measuring flow rates represents a computing operation that is, in some sense, reversible. Consider the

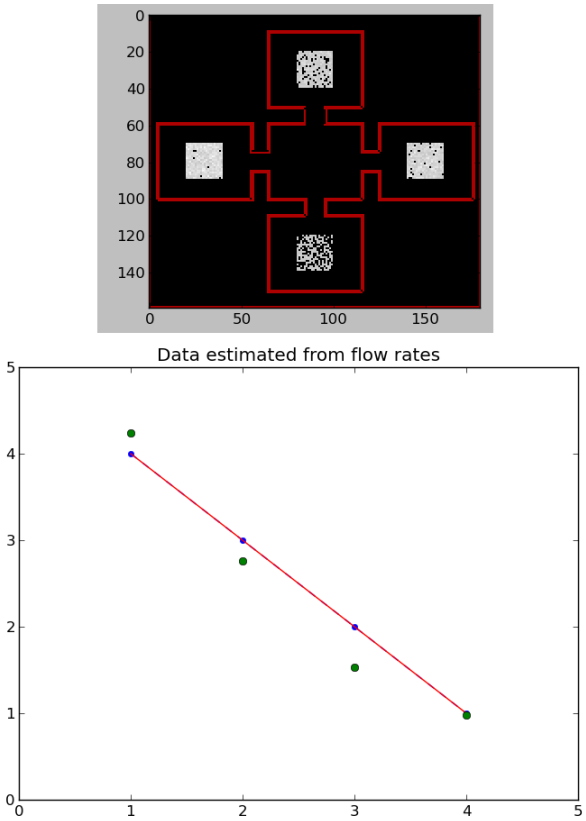

Fig. 5: Data encoding in chambers with a chamber for absolute reference. This arrangement allows for the unique estimation of $\mathbf{x}$, the data encoded in each chamber, the result of which is shown at the bottom.

2-chamber model presented in Figure 2. By the previous methods we can estimate the initial conditions of the system (particle concentration gradients) and also the maximum entropy, steady state solution (averaged density) by making local flow rate measurements. Hence, although the maximum entropy solution is not one-to-one (i.e., the same average density could be obtained by a number of initialized particle configurations), we have preserved information by making local measurements [7], [8]. Estimating the particle gradients is, then, not strictly reversible, but we still have all the information on both ends of the computation. How is this possible?

One possible thermodynamic interpretation of these results is obtained by considering the relationship between flow rate, concentration gradient, and entropy. At the start of a simulation, the system begins to move from a minimal entropy state towards a maximal one, and those initial flow rate measurements are most indicative of the particle concentration gradient. At the end of the simulation, the system is in a steady state with maximum entropy and zero net flow between chambers, which is uninformative of the initial concentration gradient. Thus, in evolving towards a state of maximum entropy, the system has encoded its useful work into the flow rates, which are informative of the initial

conditions of the system. This is what allows us to do estimation of data encoded in the LGA.

III. Active methods for data estimation

In this section, we consider active methods of inference in the LGA. These methods are "active" because particle statistics are explicitly measured and particles are dynamically added or removed depending on the measured statistics. This violates the reversible computing properties of the simulation, but allows the LGA to solve more interesting problems.

A. Maximum entropy distributions

The maximum entropy method originated with Edwin T. Jaynes's seminal 1957 paper [4] that connected statistical mechanics with information theory. Entropy itself is a measure of uncertainty, often computed in information theory for a discrete random variable X with probability mass function $p_y = p(X = x_i)$ as

$$H(p) = E\left[\frac{1}{\log p_i}\right] = \sum_i p_i \frac{1}{\log p_i}. \tag{4}$$

The base of the logarithm is often chosen to be 2, which means the entropy has units of bits. Other bases can be used. For example, if e is used as a base, the units of entropy are in nats. More detail can be found in [2]. Here, we satisfy ourselves with the intuition that entropy measures uncertainty.

In information theory, a maximum entropy probability distribution is one that is "least committal" given prior information [4]. For example, if we have a discrete random variable X with cardinality $|\mathcal{X}|$, then if we have no measurements of the values that X can take on, then the maximum entropy distribution is a uniform distribution: $X \sim U(|\mathcal{X}|)$. That is, all values of X are equally probable. As another example, when we only know the first moment μ and second moment σ^2 of a continuous random variable Y with sample space $\mathcal{Y} = (-\infty, \infty)$, the maximum entropy distribution is a normal distribution with mean μ and variance σ^2.

Likewise, statistical mechanics says that physical systems tend become more disordered over time. Jaynes showed that thermodynamic entropy is equivalent to information theoretic entropy. Thus, since the LGA is simulating a physical system, it will always tend towards the maximum entropy solution. If we somehow encode prior information from data into the LGA, it is possible to exploit the maximum entropy property to perform inference.

B. Maximum entropy image reconstruction

Here, we consider using the LGA to solve the interesting problem of image reconstruction from blurred and/or missing measurements using the maximum-entropy method. Gull and Skilling present an overview of this method in [3], and Skilling and Bryan present a general gradient-based algorithm to solve this problem in [5]. Maximum entropy has found wide application in image processing, tomography, crystallography, and astronomy. To our knowledge, no one seems to have tried using a lattice gas automata for maximum entropy reconstruction.

The basic idea behind maximum entropy is to find the image that is "least committal" given measured data. That is, the maximum entropy image has the least amount of correlation between the pixels. This property gives maximum entropy the ability to fill in missing data as well as perform deconvolution of a "true" image from a blurring response.

Let D_k be measured image data. The index k is one-dimensional here for notational simplicity, and corresponds to all 2-D locations in the image that have been measured. That is, $k \in \mathcal{K} = \{(y, x) \mid \text{data is measured at } (y, x)\}$. If the 2-D indices in $\mathcal{K}$ span the entire image, i.e. $y = 1, 2, ..., M$ and $x = 1, 2, ..., N$, then there is no missing data. Also let R_k be a blurring kernel, which could be the response of the measurement system. [1]

The goal is to estimate the "true" $M \times N$ image data f_i. Again, i is a one-dimensional index for notational simplicity, where $i \in \mathcal{I} = \{(y, x) \mid y = 1, 2, ..., M;\ x = 1, 2, ..., N\}$. To compute the entropy of a candidate image f_i, we will normalize it to be a joint probability distribution. Thus, define

$$p_i = \frac{f_i}{\sum_{i \in \mathcal{I}} f_i}. \tag{5}$$

We can write maximum entropy image reconstruction as a constrained optimization problem:

$$\begin{aligned} \underset{\mathbf{f}}{\text{maximize}} \quad & H(\mathbf{p}) = \sum_i p_i \log \frac{1}{p_i} \\ \text{subject to} \quad & p_i = \frac{f_i}{\sum_i f_i} \\ & D_k = \sum_j R_{k,j} f_j + n_k \end{aligned} \tag{6}$$

where n_k is some noise that provides a tolerance on how much measured data points (i.e., f_k for $k \in \mathcal{K}$) can deviate from the original measured data values. Entropy is a concave function, which means that it possesses a global maximum.

C. Maximum entropy image reconstruction in the LGA

Here we detail how image data is encoded in the LGA, how constraints are enforced, and practical implementation issues.

1) Encoding image data in the LGA: A $M \times N$ image with measured pixel values D_k, $k \in \mathcal{K}$, can be encoded into the LGA as follows.

This is shown schematically in figure 6. The particles clouds are initialized using the `rectNFill` function, which randomly creates particles in the specified region with a density between 0.0 and 1.0. Since all regions are the same size, this is the same as initializing the region particles counts to be proportional to measured data. The initial implementation of this scheme also drew containing reflecting walls around

[1] For example, maximum entropy was successful in reconstructing images from the Hubble space telescope when it first went up and had problems with its measurement system. Using knowledge of the response the faulty measurement system caused, maximum entropy was able to provide satisfying reconstructions of the images [1].

each region, with openings to chambers adjacent horizontally and vertically. However, it soon became apparent during testing that not using any chamber walls produced superior results. This is because the particles are less constrained to move between regions, which is important as the LGA approaches a maximum entropy state subject to constraints.

2) Regulating region particle counts in the LGA: Regulation of the chambers is performed using a simple proportional control algorithm. After the particles propogate using the LGA ruleset at time step t, the control algorithm is applied to each regulated region. In a regulated region, let $N_k(t)$ be the particle count in region k at discrete time step t. The error between current particle count and desired particle count in region k is $E_k(t) = N_k^{\text{des}} - N_k(t)$, where N_k^{des} is the number of particles that satisfies the constraint

$$N_k^{\text{des}} = \sum_{j \in \mathcal{S}(R_k)} R_{k,j} N_j \tag{7}$$

where $\mathcal{S}(R_k)$ is the support of the blurring kernel R_k. This enforces the second constraint in (6). Said in words, this constraint ensures that the estimated "true" data f_i, when passed through the blurring kernel, equals the measured data D_k.

The control algorithm adds or removes a number of particles proportional to the error as

$$\bar{N}_k(t) = N_k(t) + K_u \cdot \text{round}[E_k(t)] \tag{8}$$

where K_u is the proportional gain and $\bar{N}_k(t)$ is the adjusted number of particles for the next propogation step.

Regulation of region particle counts in the LGA essentially implements the n_k term in the second constraint in (6). The variance of the noise is proportional to the "looseness" of the controller, parameterized by a proportional gain K_u.

The actual addition or removal of particles is performed as follows. For each particle location in the regulated region, a random particle direction is chosen. The LGA encodes the six particle directions as 6-bit one-hot binary strings, so choosing a random direction consists of choosing a random integer b uniformly from $[0, 5]$, and the random binary string is 2^b. If particles are to be removed, and the particle with the specified direction exists, the particle is cleared by taking the logical-XOR of 2^b with the overall bit string and a counter is incremented. If particles are to be added, and the particle in the specified direction does not exist, the particle is added by taking the logical-OR of 2^b with the overal bit string and a counter is incremented. The control algorithm loops over locations in the region, adding/removing particles until the counter equals the desired number of added/removed particles.

D. Results for maximum entropy reconstruction using the LGA

In this subsection we present two examples of maximum entropy image reconstruction using the LGA: (1) reconstructing blurred text with missing data and (2) reconstructing an everyday image with missing data. We use these examples to perform an initial characterization of the K_u parameter, as well as the relative performance of image reconstruction using the LGA.

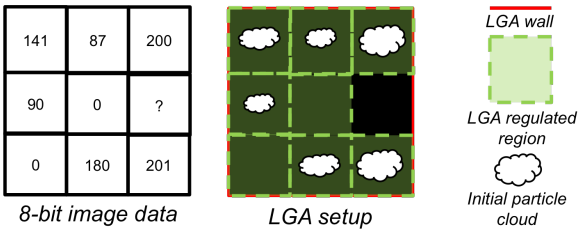

Fig. 6: Schematic of encoding image data in the LGA. Image data, shown on the left, is encoded as the number of particles initially in a region. For measured data, the region corresponding to the data is actively regulated. Missing data is encoded as an empty region that is unregulated.

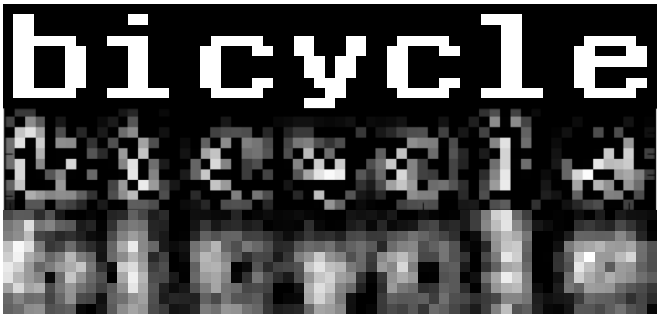

Fig. 7: Maximum entropy reconstruction of a severely degraded text string. Top is the original image, middle is Gaussian-blurred with 50% missing data, and bottom is LGA maximum entropy reconstruction.

1) Blurred text with missing data: For our first example, we take the image in the top panel of figure 7. A 5×5 Gaussian blurring kernel of radius 3 is applied, and 50% of the data is removed. The result is shown in the middle panel of figure 7. The text is now unreadable. Using the algorithm described in the previous section, the text is reconstructed using maximum entropy in the LGA. For this example, $K_u = 0.8$, and the best result occurs after about 20 steps. These parameters were determined experimentally to yield the best results. The result of the reconstruction, which is shown in the bottom panel of figure 7, is now readable, proving the usefulness of maximum entropy in the LGA.

There are several explanations for why the LGA does not achieve excellent results on this example. One reason is the locality enforced by the LGA model. Within one time step, particles can only drift between adjacent chambers. However, in the usual implementation of maximum entropy reconstruction, such locality is not present, and all pixels can interact with all other pixels equally. Thus, the LGA is limited in its ability to reconstruct images by this local correlation. In fact, in a sense the LGA is not performing true maximum entropy reconstruction, but only an approximation of it. A potential fix for this would be implementing some kind of "teleportation" between chambers. However, this would make the LGA depart further from the physical system

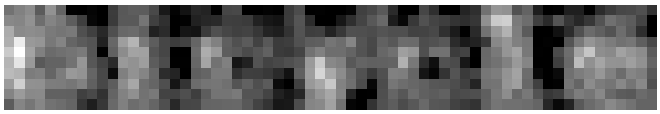

Fig. 8: Bicycle reconstruction after 100 time steps, showing blurriness.

it is supposed to simulate.

Another undesirable quality of the reconstructed images using the LGA is the inherent noisiness. This is caused by particles continuously jittering around, even in a state with relatively high entropy. This problem may be solved in the future by averaging LGA reconstructions over a number of trials.

Yet another undesirable aspect of this example is that the asymptotic performance is not ideal. Figure 8 shows the LGA reconstruction output after 100 time steps. The resulting image is quite blurry, though arguably readable. The local correlation mentioned in the previous paragraph may be causing the blurriness, preventing the LGA from finding a better approximate maximum entropy solution. Another way of improving this aspect may be averaging over multiple reconstruction trials.

2) Everyday image with missing data: To better quantify how well the LGA reconstruction is performing, we consider the 64×64 Cameraman image in the upper left panel of figure 9. This image has a nice mix of large scale and small scale features.

For a range of K_u values, we ran the LGA maximum entropy reconstruction for $T = 80$ time steps per K_u and measured the mean-squared error (MSE) and peak signal-to-noise ratio (PSNR) between the reconstructed image and the orignal image. MSE is computed as

$$MSE(\hat{f}_j) = \sum_j (f_j - \hat{f}_j)^2 \tag{9}$$

and PSNR is computed as

$$PSNR(\hat{f}_j) = 10\log_{10}\left(\frac{255^2}{MSE(\hat{f}_j)}\right) \tag{10}$$

where $\hat{f}_i$ is the estimated image and f_i is the true original image. Results are shown in table 1. It can be easily seen that a K_u close to unity gives ideal results, which makes sense.

TABLE I: Objective metrics for reconstructions

K_u	MSE (lower is better)	PSNR (dB, higher is better)
0	2803.1	13.65
0.4	594.6	20.39
0.8	424.5	21.85
1.2	410.8	21.99
1.6	1297.2	17.00

One nice aspect of this example is that the asymptotic state after many time steps is stable, unlike the first example. Again, the output could be improved by averaging over many trials.

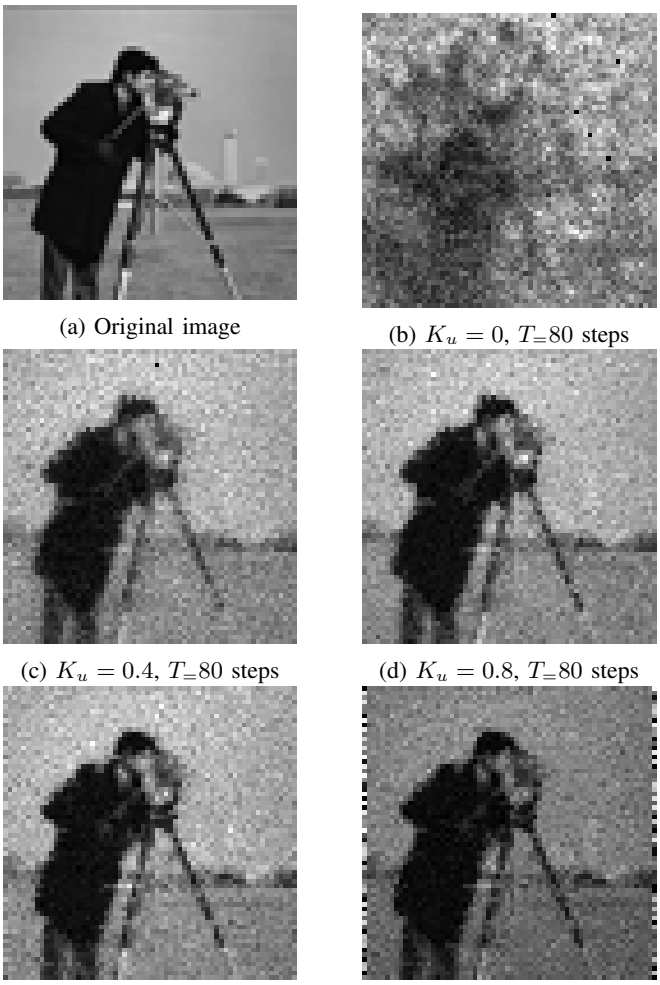

(a) Original image

(b) $K_u = 0$, T=80 steps

(c) $K_u = 0.4$, T=80 steps

(d) $K_u = 0.8$, T=80 steps

(e) $K_u = 1.2$, T=80 steps

(f) $K_u = 1.6$, T=80 steps

Fig. 9: Original Cameraman image (top left) and LGA maximum entropy reconstructions for different K_u (others).

E. Improvements and future work

There are several improvements we propose for maximum entropy image reconstruction in the LGA.

First, it would be interesting to experiment with different annealing schedules for K_u. That is, to make K_u a function of time step t. This would allow the control algorithm to be "looser" at the beginning, allowing particles to flow more easily between regions. As the LGA hones in on a solution, K_u could increase, locking in more optimal solutions. This would be a simulated annealing approach. Second, it would be interesting to allow more flow between different regions and see what the LGA maximum entropy reconstructions look like. Third, the effect of averaging over mutliple trials should be attempted to see if better reconstructions with less noise result. Fourth, currently noise is not being taken into account in LGA maximum entropy image reconstruction. The LGA may be suitable for image denoising in addition to deblurring and reconstructing from missing data. Finally, it would be interesting to extend LGA maximum entropy reconstruction to blind image deconvolution, which would

be similar to the method mentioned by Gull and Skilling in [3]. For blind deconvolution, there would be two large chambers, one encoding the "true" image data and the other encoding R_k, the blurring kernel. Currently knowledge of R_k is assumed, but in the blind setup the LGA may be able to learn the maximim entropy solution for both f_i and R_k.

IV. Conclusion

In this paper, we have described two methods of doing inference in the lattice gas automata. Examination of passive methods led to insights about how flow rates connect to initial densities, and we explored how these properties can be used for data estimation. Next, using an active method of actually perturbing the LGA system, we demonstrated proof-of-concept for maximum entropy image reconstruction using the LGA, and proposed several avenues of further exploration that could improve the results.

We hope that this unconventional approach to inference may inspire other attempts, and that it may encourage researchers to look at problems in new ways.

V. ACKNOWLEDGMENTS

The authors thank Prof. Josh Smith for the very interesting course, and James Youngquist for helping develop the code!

References

[1] N L Bonavito, J E Dorband, and T Busse. Maximum entropy restoration of blurred and oversaturated hubble space telescope imagery. *Applied optics*, 32(29):5768–5774, October 1993. PMID: 20856397.

[2] T. M Cover and Joy A , Thomas. *Elements of information theory*. Wiley-Interscience, Hoboken, N.J., 2006.

[3] S.F. Gull and J. Skilling. Maximum entropy method in image processing. *Communications, Radar and Signal Processing, IEE Proceedings F*, 131(6):646–659, 1984.

[4] Edwin T. Jaynes. Information theory and statistical mechanics. *Physical review*, 106(4):620, 1957.

[5] J. Skilling and R. K. Bryan. Maximum entropy image reconstruction - general algorithm. *Monthly Notices of the Royal Astronomical Society*, 211:111, November 1984.

[6] E. L. Cussler, *Diffusion: Mass Transfer in Fluid Systems*, Cambridge University Press, Cambridge, UK; 2009.

[7] C. H. Bennett, "Logical reversibility of computation", *IBM J. Res. Develop.*, pp.525 -532 1977

[8] E. Fredkin, T Toffoli, "Conservative Logic", *International Journal of Theor. Physics*, vol. 21 , pp.219-253, 1982.