

Evolving Transition Rules for Multi Dimensional Cellular Automata

Ron Breukelaar¹ and Thomas Bäck^{1,2}

¹ Universiteit Leiden, LIACS, P.O. Box 9512, 2300 RA Leiden, The Netherlands
{rbreukel,baeck}@liacs.nl

² Nutech Solutions GmbH, Martin-Schmeisser-Weg, 44227 Dortmund, Germany

Abstract. Genetic Algorithms have been used before to evolve transition rules for one dimensional Cellular Automata (CA) to solve e.g. the majority problem and investigate communication processes within such CA [3]. In this paper, the principle is extended to multi dimensional CA, and it is demonstrated how the approach evolves transition rules for the two dimensional case with a von Neumann neighborhood. In particular, the method is applied to the binary AND and XOR problems by using the GA to optimize the corresponding rules. Moreover, it is shown how the approach can also be used for more general patterns, and therefore how it can serve as a method for calibrating and designing CA for real-world applications.

1 Introduction

According to [6] Cellular Automata (CA) are mathematical idealisations of physical systems in which space and time are discrete, and physical quantities take on a finite set of discrete values. The simplest CA is one dimensional and looks a bit like an array of ones and zeros of width N , where the first position of the array is linked to the last position. In other words, defining a row of positions $C = \{a_1, a_2, \dots, a_N\}$ where C is a CA of width N and a_N is adjacent to a_1 .

The neighborhood s_n of a_n is defined as the local set of positions with a distance to a_n along the connected chain which is no more than a certain radius (r). This for instance means that $s_2 = \{a_{148}, a_{149}, a_1, a_2, a_3, a_4, a_5\}$ for $r = 3$ and $N = 149$. Please note that for one dimensional CA the size of the neighborhood is always equal to $2r + 1$.

The values in a CA can be altered all at the same time (synchronous) or at different times (asynchronous). Only synchronous CA are considered in this paper. In the synchronous approach at every timestep (t) every cell state in the CA is recalculated according to the states of the neighborhood using a certain transition rule $\Theta : \{0, 1\}^{2r+1} \rightarrow \{0, 1\}, s_i \rightarrow \Theta(s_i)$. This rule basically is a one-to-one mapping that defines an output value for every possible set of input values, the input values being the 'state' of a neighborhood. The state of a_n at time t is written as a_n^t , the state of s_n at time t as s_n^t and the state of the entire CA C at time t as C^t so that C^0 is the initial state and $\forall n = 1, \dots, N a_n^{t+1} = \Theta(s_n^t)$. Given $C^t = \{a_1^t, \dots, a_N^t\}$, C^{t+1} can be defined as $\{\Theta(s_1^t), \dots, \Theta(s_N^t)\}$.

Because $a_n \in \{0, 1\}$ the number of possible states of s_n equals $2^{2^{r+1}}$. Because all possible binary representations of m where $0 \leq m < 2^{2^{r+1}}$ can be mapped to a unique state of the neighborhood, Θ can be written as a row of ones and zeros $R = \{b_1, b_2, \dots, b_{2^{2^{r+1}}}\}$ where b_m is the output value of the rule for the input state that maps to the binary representation of $m - 1$. A rule therefore has a length that equals $2^{2^{r+1}}$ and so there are $2^{2^{2^{r+1}}}$ possible rules for a binary one dimensional CA. This is a huge number of possible rules (if $r = 3$ this sums up to about $3,4 \times 10^{28}$) each with a different behaviour.

One of the interesting things about these and other CA is that certain rules tend to exhibit organisational behaviour, independently of the initial state of the CA. This behaviour also demonstrates there is some form of communication going on in the CA over longer distances than the neighborhood allows directly. In [3] the authors examine if these simple CA are able to perform tasks that need positions in a CA to work together and use some form of communication. One problem where such a communication seems required in order to give a good answer is the Majority Problem (as described in section 3.1). A genetic algorithm is used to evolve rules for one dimensional CA that do a rather good job of solving the Majority Problem [3] and it is shown how these rules seem to send “particles” and communicate by using these particles [4]. These results imply that even very simple cells in one dimensional cellular automata can communicate and work together to form more complex and powerful behavior.

It is not unthinkable that the capabilities of these one dimensional CA are restricted by the number of directions in which information can “travel” through a CA and that using multiple dimensions might remove these restriction and therefore improve performance. Evolving these rules for the Majority Problem for two dimensional CA using a Moore neighborhood (explained in section 3) is reported in [2] showing that the GA did not clearly outperform random search.

The goal of the research is to find a generalization and report phenomena observed on a higher level, with the future goal to use this research for identification and calibration of higher-dimensional CA applications to real world systems like parallel computing and modelling social and biological processes. The approach is described and results are reported on simple problems such as the Majority Problem, AND, XOR, extending into how it can be applied to pattern generation processes.

2 The Genetic Algorithm

As mentioned before, this research was inspired by earlier work [3,4] in which transition rules for one dimensional CA were evolved to solve the Majority Problem (as defined in section 3.1). The GA is a fairly simple algorithm using binary representation of the rules, mutation by bit inversion, truncation selection, and single-point crossover. The algorithm determined the fitness by testing the evolved rules on 100 random initial states. Every iteration the best 20% of the rules (the ‘elite’ rules) were copied to the next generation and the other 80% of

the rules were generated using single-point crossover with two randomly chosen ‘elite’ rules and then mutated by flipping exactly 2 bits in the rule.

To be able to compare two dimensional CA with one dimensional CA the GA used in section 3.1 is a copy of the the GA used in [3, 4]. The GAs in section 3.2 and 4 on the other hand are modified to fit the different problem demands, as will be explained in these sections.

3 Experimental results for two dimensional CA

The two dimensional CA in this document are similar to the one dimensional CA discussed so far. Instead of a row of positions, C now consists of a grid of positions. The values are still only binary (0 or 1) and there still is only one transition rule for all the cells. The number of cells is still finite and therefore CA discussed here have a width, a height and borders.

The big difference between one dimensional and two dimensional CA is the rule definition. The neighborhood of these rules is two dimensional, because there are not only neighbors left and right of a cell, but also up and down. That means that if $r = 1$, s_n would consist of 5 positions, being the four directly adjacent plus a_n . This neighborhood is often called “the von Neumann neighborhood” after its inventor. The other well known neighborhood expands the Neumann neighborhood with the four positions diagonally adjacent to a_n and is called “the Moore neighborhood” also after its inventor.

Rules are defined with the same rows of bits (R) as defined in the one dimensional case. For a von Neumann neighborhood a rule can be defined with $2^5 = 32$ bits and a rule for a Moore neighborhood needs $2^9 = 512$ bits. This makes the Moore rule more powerful, for it has a bigger search space. Yet, this also means that searching in that space might take more time and finding anything might be a lot more difficult. In [2] the authors discourage the use of the Moore neighborhood, yet in section 3.2 and section 4 results clearly show successes using the Moore neighborhood, regardless of the larger search space.

In a one dimensional CA the leftmost cell is connected to the rightmost cell. In the two dimensional CA this is also common such that it forms a torus structure.

3.1 Majority problem

The Majority Problem can be defined as follows: *Given a set $A = \{a_1, \dots, a_n\}$ with n odd and $a_m \in \{0, 1\}$ for all $1 \leq m \leq n$, answer the question: ‘Are there more ones than zeros in A ?’.*

The Majority Problem first does not seem to be a very difficult problem to solve. It seems only a matter of counting the ones in the set and then comparing them to the number of zeros. Yet when this problem is converted to the dimensions of a CA it becomes a lot more difficult. This is because the rule in a CA does not let a position look past its neighborhood and that is why the cells all have to work together and use some form of communication.




Fig. 1. This figure displays the number of rules that have a certain fitness value in the two dimensional experiment and compares this to the one dimensional experiment. The fitness bins are 0.01 in width and for both algorithm $F_{169,10^3}$ is calculated for 300 rules.

Given that the relative number of ones in C^0 is written as λ , in a simple binary CA the Majority Problem can be defined as: *Find a rule that, given an initial state of a CA with N odd and a finite number of iterations to run (I), will result in an ‘all zero’ state if $\lambda < 0.5$ and an ‘all one’ state otherwise.*

The fitness (f) of a rule is therefore defined as the relative number of correct answers to 100 randomly chosen initial states, where a ‘correct answer’ corresponds to an ‘all zero’ state if $\lambda < 0.5$ and an ‘all one’ state otherwise. In [3] the authors found that using a uniform distribution over λ for the initial states enhanced performance greatly; this is used here as well. The best runs will be tested using randomly chosen initial states with a normal distribution over the number of ones. The relative number of correct classifications on these states is written as $F_{n,m}$ where n is the width of the CA and m is the number of tests conducted.



Fig. 2. This figure shows a correct classification of the Majority Problem by a two dimensional CA with both width and height equal to 13 and $\lambda = 84/169$. The transition rule was one of the best tested in the experiment and scored $F_{169,10^3} = 0.715$.

Preliminary experiments showed that it took much more time to evolve rules for the Moore neighborhood than for the von Neumann neighborhood. The tests that were done with the Moore neighborhood also did not result in any encouraging results, this being in line with [2]. That is why the von Neumann neighborhood was chosen for this experiment. Because this neighborhood consists of five positions, the search space for CA rules is a lot smaller than in the one dimensional experiment. Instead of the $2^7 = 128$ bits in the rule, R now consists of $2^5 = 32$ bits, thus drastically decreasing the search space. This means that the search space decreased from 2^{128} to 2^{32} and is now $2^{(128-32)} = 2^{96}$ times smaller!

For this experiment we used a CA with width = 13 and height = 13. This means that these CA have $13 \times 13 = 169$ cells (N) and are $169 - 149 = 20$ cells larger than the one dimensional CA used in the original experiment.

This algorithm was run 300 times and each winning rule was tested by calculating $F_{N,M}$ using $F_{169,10^3}$. These results are plotted against results of our own one dimensional experiment (not reported here, analogue to [3, 4]) in Figure 1. The striking difference between this distribution of fitness and the distribution of fitness in the one dimensional experiment is the absence of the peak around $F_{N,M} \approx 0.5$ in the two dimensional results. In those results almost all the evolved rules have a fitness above 0.58. A fitness around 0.66 seems to be average and the best rules have a fitness above 0.7. That is all very surprising taking into account that the von Neumann neighborhood only consists of 5 cells.

The Majority Problem is a good example of a problem that forces cells in a CA to ‘communicate’ with another. The communication ‘particles’ can be seen in the one dimensional experiment, but are not easily spotted in the two dimensional experiment. That does not mean there are no ‘particles’ traveling in the two dimensional CA, because it might be very hard to identify these particles. In a two dimensional CA ‘particles’ are no longer restricted to traveling in only one direction, but can travel to multiple directions at the same time. Traveling particles in two dimensional CA can therefore look like expanding areas with a distinct border. But there might be multiple particles traveling at the same time, meeting each other and thereby creating new particles. This is why communication between cells in a two dimensional CA is not very visible in the Majority Problem, although results show that this communication is present.

3.2 AND and XOR Problem

To show the communication between cells in a two dimensional CA a different experiment was conducted. A genetic algorithm was used to evolve rules for two dimensional CA that could solve the simple binary operators AND and XOR. These operators both have two input values and one output value which can only be determined if both input values are known. This is unlike the OR operator for example where the output value is always one if one or more of the input values is one, so if only one input value is known to be one then the value of the other input value is not needed. This may look very trivial, but it is very important in order to force the CA to combine the two values and thereby communicate.




Fig. 3. This figure displays the iterations of a CA solving the AND problem. Every row shows the iteration of the rule using a different initial state. Note that in the first column ($t = 0$) the initial states are clearly visible and in the last column the coloring matches the output of an AND port.

The AND Problem To show the communications in a CA the information that needs to be combined must be initialized as far apart as possible. The following problem definition takes this into account: *Given a square CA with two ‘input cells’, one top left and one bottom right: find a rule that iterates the CA so that after I iterations the CA is in an ‘all one’ state if both the ‘input cells’ were one in the initial state and in an ‘all zero’ state otherwise.*

Small two dimensional CA were used with a width and a height of 5 cells and I was set to 10. The borders of the CA were unconnected to allow a larger virtual distance between the two corner cells. This means that the leftmost cell in a row was not connected to the rightmost cell in the same row and the topmost cell was not connected to the bottommost cell as was done with the Majority Problem experiment. Instead every cell on the border of the CA was connected to so called ‘zero-cells’. These ‘zero-cells’ stay zero whatever happens.

When using two input cells, there are four different initial states. These states are written as $S_{(v_1, v_2)}$ where v_1 and v_2 are the two input values. All cells other than the two input cells are initialized with zero.

Table 1. Fitness values found in the AND problem.

Fitness	Number of runs			
	Neumann		Moore	
	with crossover	without crossover	with crossover	without crossover
100	0	0	31	21
98-99	0	0	41	54
95-97	0	0	14	25
90-94	77	93	14	0
80-89	23	7	0	0
70-79	0	0	0	0
< 70	0	0	0	0

The fitness of a rule is defined as the total number of cells that have the correct values after I iterations. The number of ones in iteration t is written as

$O_{(v_1, v_2)}^t$. The total fitness of the AND problem is defined as $f = (N - O_{(0,0)}^I) + (N - O_{(0,1)}^I) + (N - O_{(1,0)}^I) + O_{(1,1)}^I$. This makes the maximum fitness equal to $4 \times 5 \times 5 = 100$.

In this experiment another variation of the simple genetic algorithm was used. A generation step starts by sorting the rules according to their fitness. Then it selects the top ten% of the rules as ‘elite’ rules and copies them without changes to the next generation. Every ‘elite’ rule is then copied nine times or is used in single-point crossover to make the other 90% of the population. Both methods were tested and compared. The generated rules are mutated and also moved to the next generation. Mutation is done by flipping every bit in the rule with a probability p_m . The algorithm stops if it finds a rule with $f = 100$ or it reaches 1000 generations. In preliminary experiments a number of different values of p_m were tested. Setting p_m to a rather high value of 0.05 turned out to be the most effective choice, confirming our insight that with increasing selection strength higher mutation rates than the usual $\frac{1}{l}$ (l being the the length of the binary string) are performing better [1].

The algorithm was run 100 runs with and without single-point crossover and using both the von Neumann and the Moore neighborhoods. The results are shown in Table 1.

Although rules evolved with the von Neumann neighborhood are not able to solve the problem perfectly, it is already surprising that it finds rules which work for 93%, for such a rule only misplaces 7 cells in the final state. All the other 93 cells have the right value. This suggests that the information was combined, but the rule could not fill or empty the whole square using the same logic.

The Moore neighborhood is clearly more powerful and was able to solve the problem perfectly. The rules that are able to do this clearly show communicational behaviour in the form of “traveling” information and processing this information at points where information “particles” meet.

It is also surprising that using crossover in combination with a Neumann neighborhood does not outperform the same algorithm without the crossover. This may be due to the order of the bits in the transition rule and their meaning. This is worth exploring in future work. Maybe using other forms of crossover might give better results in combination with multi dimensional CA.

The XOR Problem The XOR Problem is not much different from the AND problem. We used the same genetic algorithm and the same CA setup. The only difference is the fitness function. We defined the XOR problem as follows: *Given a square CA with two ‘input cells’, one top left and one bottom right: find a rule that iterates the CA so that after I iterations the CA is in an ‘all one’ state if only one of the ‘input cells’ was one in the initial state and in an ‘all zero’ state otherwise.* This means that the total fitness of the XOR problem is defined as $f = (N - O_{(0,0)}^I) + O_{(0,1)}^I + O_{(1,0)}^I + (N - O_{(1,1)}^I)$.

The algorithm was run with $p_m = 0.05$ for a maximum of 1000 generations for 100 runs with both Neumann and Moore neighborhoods with and without single point crossover. The results are shown in Table 2.




Fig. 4. This figure displays the iterations of a CA solving the XOR problem. Every row shows the iteration of the rule using a different initial state. Note that in the first column ($t = 0$) the initial states are clearly visible and in the last column the coloring matches the output of an XOR port.

These results support earlier finding in suggesting that single-point crossover doesn't really improve the performance when used in a two dimensional CA. The results show that the algorithm using only mutation has found ways to solve this rather difficult communicational problem. The Neumann neighborhood seemed unable to perform for 100%, yet it came rather close with one rule classifying the problem for 92%. The algorithm found one transition rule using the Moore neighborhood that is able to solve the problem for the full 100%. This rule depicted in Figure 4 shows clear signs of "traveling particles" and is another example of how a local rule can trigger global behaviour.

Table 2. Fitness values found in the XOR problem.

Fitness	Number of runs			
	Neumann		Moore	
	with crossover	without crossover	with crossover	without crossover
100	0	0	0	1
98-99	0	0	4	4
95-97	0	0	7	6
90-94	2	1	19	21
80-89	76	96	69	66
70-79	18	3	1	2
< 70	4	0	0	0

4 Evolving bitmaps

Now that it is shown that two dimensional CA's can communicate, it is time to increase the challenge for the CA a bit. The aim of this experiment is to evolve rules for two dimensional CA that generate patterns (or bitmaps).

The Bitmap Problem is defined as follows: *Given an initial state and a specific desired end state: find a rule that iterates from the initial state to the desired state*



Fig. 5. The bitmaps used in the pattern generation experiment.

in less than I iterations. Note that this does not require the number of iteration between the initial and the desired state to be fixed.

The CA used in this experiment is not very different from the one used in the AND/XOR experiment (section 3.2). In preliminary experiments we tried different sizes of CA, but decided to concentrate on small square bitmaps with a width and a height of 5 cells (as done in section 3.2). To make the problem harder and to stay in line with earlier experiments the CA have unconnected borders like in section 3.2. The von Neumann neighborhood was chosen instead of the Moore neighborhood and therefore s_n consist of 5 cells ($r = 1$) and a rule can be described with $2^5 = 32$ bits. The searchspace therefore is $2^{32} = 4294967296$.

After testing different initial states, the ‘single seed’ state was chosen and defined as the state in which all the positions in the CA are zero except the position $(\lfloor \text{width}/2 \rfloor, \lfloor \text{height}/2 \rfloor)$ which is one. For the GA we used the same algorithm as we used in the AND and XOR experiments. Because this experiment uses a Neumann neighborhood and the AND and XOR experiments suggested that the combination between the von Neumann neighborhood and single point crossover was not a good idea, this experiment used only mutation. Like in section 3.2 mutation is performed by flipping every bit in the rule with a probability p_m . In this experiment $p_m = 1/32 = 0.03125$.

In trying to be as diverse as possible five totally different bitmaps were chosen, they are shown in Figure 5. The algorithm was run 100 times for every bitmap for a maximum of 5000 generations. The algorithm was able to find a rule for all the bitmaps, but some bitmaps seemed a bit more difficult than others. Table 3 shows the number of successful rules for every bitmap. Note that symmetrical bitmaps seem to be easier to generate then asymmetric ones.

Although this experiment is fairly simple, it does show that a GA can be used to evolve transition rules in two dimensional CA that are able to generate patterns even with a simple von Neumann neighborhood. Ongoing experiments with bigger CA suggest that they don’t differ much from these small ones, although the restrictions on what can be generated from a single-seed state using only a von Neumann neighborhood seem to be bigger when size of the CA increases.

Table 3. Number of successful rules found per bitmap.

<i>Bitmap</i>	<i>Successful rules (out of a 100)</i>
“square”	80
“hourglass”	77
“heart”	35
“smiley”	7
“letter”	9




Fig. 6. This figure shows some iteration paths of successful transition rules.

5 Conclusions

This paper shows how two dimensional CA are able to solve the majority problem with similar results compared to one dimensional CA used in [3, 4]. Using the same GA as in [3, 4] a better average fitness was achieved suggesting that evolving two dimensional CA is easier and more reliable.

The paper shows that two dimensional CA can show communicational behaviour in the form of the AND and XOR problems and that this behaviour can be evolved using a GA. The document also shows that a more generic behaviour can be evolved using a GA by showing how different patterns can be iterated from the same initial state. These results all suggest that a multi dimensional CA is a very powerful tool and in combination with GAs they can be evolved to exhibit specific behaviour. It is therefore feasible that this combination can be used to solve all sorts of real world problems.

References

1. Bäck, T.: *Evolutionary Algorithms in Theory and Practice*. Oxford University Press, NY (1996)
2. Inverso, S., Kunkle, D., Merrigan, C.: *Evolutionary Methods for 2-D Cellular Automata Computation*. www.cs.rit.edu/~drk4633/mypapers/gacaProj.pdf (2002)
3. Mitchell, M., Crutchfield, J.P.: *The Evolution of Emergent Computation*. Proceedings of the National Academy of Sciences, SFI Technical Report 94-03-012
4. Mitchell, M., Crutchfield, J.P., Hraber, P.T.: *Evolving cellular automata to perform computations: Mechanisms and impediments*. *Physica D*, (1994), 75:361 - 391
5. Li, W., Packard, N. H., Langton, C. G.: *Transition phenomena in cellular automata rule space*. *Physica D*, (1990), 45:77-94
6. Wolfram, S.: *Statistical mechanics of Cellular Automata*. *Reviews of Modern Physics* volume 55 (1983)
7. Wolfram, S.: *Theory and Applications of Cellular Automata*. World Scientific, Singapore (1986)