

Stochastic Cellular Automata Solve the Density Classification Problem with an Arbitrary Precision

Nazim A. Fatès

► To cite this version:

Nazim A. Fatès. Stochastic Cellular Automata Solve the Density Classification Problem with an Arbitrary Precision. 2010. inria-00521415v1

HAL Id: inria-00521415 https://inria.hal.science/inria-00521415v1

Preprint submitted on 27 Sep 2010 (v1), last revised 7 Jan 2011 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Stochastic Cellular Automata Solve the Density Classification Problem with an Arbitrary Precision

Nazim Fatès INRIA Nancy – Grand Est, LORIA, Nancy Université nazim.fates@loria.fr

September 27, 2010

Abstract

The density classification problem consists in using a binary cellular automaton (CA) to decide whether an initial configuration contains more 0s or 1s. This problem is known for having no exact solution in the case of deterministic one-dimensional CA. Stochastic cellular automata have been studied as an alternative for solving the problem. This paper is amied at presenting techniques to describe stochastic CA rules as a "blend" of deterministic CA rules. We analyse two previously studied rules and present a new rule that solves the problem with an arbitrary precision. Using analytical calculations and numerical simulations, we estimate how the quality of classification and the average time of classification varies as a function of the blending parameter. From a practical point of view, simulations show that for a great range of values, the new rule exhibits a quality of classification never attained so far.

Introduction

The density classification problem is one of the most studied inverse problems in the field of cellular automata. Informally, it requires that a binary cellular automaton — or more generally a discrete dynamical system — decides whether an initial binary string contains more 0s or more 1s. In its classical formulation, the cells are arranged in a ring and each cell can only read its own state and the states of the neighbouring cells. The challenge is to design a behaviour of the cells that drives the system to a uniform state, that consists of all 1s if the initial configuration contained more 1s and all Os otherwise. In short, the convergence of the cellular automaton should decide whether the initial density of 1s was greater or lower than 1/2.

Although the task of the problem looks trivial, it has attracted a considerable amount of research since its formulation [13]. The difficulty of finding a solution

comes from the impossibility to centralise the information or to use any classical counting technique. Instead, the convergence to a uniform state should be obtained by using only *local* decisions, that is, by using an information that is limited to the close neighbours of a cell. Moreover, as CA are homogeneous by nature (the cells obey the same law), there can be no specialisation of the cells for a partial computation. Solving the problem efficiently requires to find the right balance between making local decisions within the cells and following other cells' decision to attain a global consensus.

The quest for efficient rules has been conducted on two main directions: mandesigned rules and rules obtained with large space exploration techniques such as genetic algorithms (e.g., [11]). The Gacs-Kurdymov-Levin (GKL) rule, which was originally designed the purpose of resisting small amounts of noise [14, 5], proved to be a good candidate (~80% of the initial conditions well-classified on rings of 149 cells) and remained unsurpassed for a long time. In 1995, after observing that outperforming this rule was difficult, Land and Belew issued a key result: no perfect (deterministic) density classifier that uses only two states exist [9]. However, this did not stop the search for efficient CA as nothing was known about how well a rule could perform. In particular, it was asked whether an upper bound on a rule quality would exist. The search for rules with an increasing quality has been carried on until now, with genetic algorithms as the main investigation tool (see e.g. [4, 12] and references therein).

On the other hand, various modifications to the classical problem were proposed, allowing one to solve the problem exactly. For instance, Cappacarere *et al.* proposed to modify the output specification of the problem to find a solution that classifies the density perfectly [3]. Fukś showed that running two CA rules successively would also provide an acceptable solution [7]. This issue was further explored by Martins and Oliveira, who discovered various couples and triples of rules that solve the problem when applied sequentially and for a given number of steps [10]. Some authors also proposed to embed a memory in the cells, which is another method for enhancing the abilities of the rules [1, 16].

However, all these solutions break the original specification of the problem, where the cells have only two states and obey a homogeneous rule. The use of stochastic (or probabilistic)¹ CA is an interesting alternative that complies with these two conditions. Indeed, in stochastic CA, the only modification to the CA structure is that the outcome of the local transitions of the cells are no longer deterministic, they are specified by a probability to be a given state. This research path was opened by Fukś whose rule acts as a "stochastic copy" of the state of the neighbouring cells [8]. However, the mechanism has no force that drives the system towards its goal; the convergence is mainly attained with a random drift of the system (see Sec. 2). Recently, Schüle et al. proposed a stochastic rule that implements a local majority calculus [15]. This allows the system to converge to its goal more efficiently, but the convergence rates remain bounded by some intrinsic limitations (see Sec. 3).

 $^{^1}$ Both terms 'stochastic' and 'probabilistic' CA are found in literature. We prefer to employ the former as etymologically the Greek word 'stochos' implies the idea of goal, aim, target or expectation.

We propose to follow this path and present a new stochastic rule that solves the density classification problem with an arbitrary precision, that is, with a probability of success arbitrarily close to 1. This result answers negatively the open question to whether there exists an upper bound on the success rate one can reach. The idea is to use randomness to solve the dilemma between the local majority decisions and the propagation of a consensus state. A tradeoff is obtained by tuning a single parameter, δ , that weights two well-known deterministic rules, namely the majority rule and the "traffic" rule. We show that the probability of making a good classification approaches 1 as δ is set closer to 0. To evaluate the "practical" use of our rule, we perform numerical simulations. Results show that this rule attains qualities of classification that have been out of reach so far.

1 Formalisation of the Problem

In this section, we define the deterministic Elementary Cellular Automata and their stochastic counterpart. We introduce the main notations for studying our problem.

1.1 Elementary Cellular Automata

Let $\mathcal{L} = \mathbb{Z}/n\mathbb{Z}$ represent a set of *n* cells arranged in a ring. Each cell can hold a state in $\{0, 1\}$ and we call a *configuration* the state of the system at a given time ; the configuration space is $\mathcal{E}_n = \{0, 1\}^{\mathcal{L}}$, it is finite and we have $|\mathcal{E}_n| = 2^n$. We denote by $|x|_P$ the number of occurrences of a pattern *P* in *x*. The *density* $\rho(x)$ of a configuration $x \in \mathcal{E}_n$ is the ratio of 1s in this configuration: $\rho(x) = |x|_1/n$. We denote by $\mathbf{0} = \mathbf{0}^{\mathcal{L}}$ and $\mathbf{1} = \mathbf{1}^{\mathcal{L}}$ the two special uniform configurations. For $q \in \{0, 1\}$, a configuration *x* is a *q*-archipelago if all the cells in state *q* are isolated, *i.e.*, if *x* does not contain two adjacent cells in state *q*.

In all the following, we assume that n is odd. This will prevent us from dealing with configurations that have an equal number of 0s and 1s.

An Elementary Cellular Automaton (ECA) is a one-dimensional binary CA with nearest neighbour topology, defined by its local transition rule, a function $\phi : \{0,1\}^3 \to \{0,1\}$ that specifies how to update a cell using only nearest-neighbour information. For a given ring size n, the global transition rule $\Phi : \mathcal{E}_n \to \mathcal{E}_n$ associated to ϕ is the function that maps a configuration x^t to a configuration x^{t+1} such that:

$$\forall c \in \mathcal{L}, \ x_c^{t+1} = \phi(x_{c-1}^t, x_c^t, x_{c+1}^t)$$

A Stochastic Elementary Cellular Automaton (s-ECA) is also defined by a local stochastic transition rule f, but the next state of a cell is known only with a given probability. In the binary case, we define $f : \{0, 1\}^3 \rightarrow [0, 1]$ where f(x, y, z) is probability that the cell updates to state 1 given that its neighbourhood has the state (x, y, z).

We define the global transition rule F associated to a s-ECA f as the stochastic function that assigns to a random configuration x^t the random configuration x^{t+1} characterised² by:

$$\forall c \in \mathcal{L}, \ x_c^{t+1} = \mathcal{B}_c^t \left(f(x_{c-1}^t, x_c^t, x_{c+1}^t) \right) \tag{1}$$

where x_c^{t+1} denotes the random variable that is given by observing the state of cell c and where $(\mathcal{B}_c^t)_{c \in \mathcal{L}, t \in \mathbb{N}}$ is a series of independent Bernoulli random variables, *i.e.*, $\mathcal{B}_c^t(p)$ is a random variable that equals to 1 with probability pand 0 with probability 1 - p.

1.2 Density Classifiers

We say that a configuration x is a *fixed point* for the global function F if we have F(x) = x with probability 1. A global function F associated to a local function f is a *(density) classifier* if **0** and **1** are its two only fixed points.

For a classifier C, let T(x) be the random variable that takes its values in $\mathbb{N} \cup \infty$ defined as:

$$T(x) = \min\{t : x^t \in \{0, 1\}\}$$

We say that C correctly classifies a configuration x if T(x) is finite and if $x^{T(x)} = \mathbf{1}$ for $\rho(x) > 1/2$ or if $x^{T(x)} = \mathbf{0}$ for $\rho(x) < 1/2$. The probability of good classification G(x) of a configuration x is the probability that C correctly classifies x.

To evaluate quantitatively the quality of a classifier requires to choose a distribution of the initial configurations. Various such distributions are found in literature, often without an explicit mention, and this why one may read different quality evaluations for the same rule (for instance compare the results given for the GKL rule: 82% in Ref. [3] and 97.8% in Ref. [9]). In order to avoid ambiguities, we re-define here the three main distributions of initial configurations that have been used by authors:

- (a) The *d*-binomial distribution $\mu_{\rm b}$ is obtained by choosing a configuration uniformly in \mathcal{E}_n .
- (b) The *d*-uniform distribution μ_d is obtained by choosing an initial probability ρ uniformly in [0, 1] and then building the configuration x by assigning to each cell a probability ρ to be in state 1 and a probability 1 - ρ to be state 0.
- (c) The k-uniform distribution μ_k is obtained by choosing a number k uniformly in $\{0, \ldots, n\}$ and then by choosing uniformly a configuration x in the set of configurations of \mathcal{E}_n that contain exactly k ones.

²Note that defining rigorously the series of random variables x^t obtained from F would require to introduce advanced tools from the probability theory. In particular, one should define a space of realisation Ω and always consider probability measures on Ω and for all $\omega \in \Omega$, define the random variables with respect to the configurations $x^t(\omega) \in \mathcal{E}_n$. For the sake of simplicity, and as it is frequently done, the parameter ω is omitted and the random variables are defined only with regard to their probability of realisation on Ω .

Formally,

$$\forall x \in \mathcal{E}_n, \quad \mu_{\rm b}(x) = \frac{1}{2^n} \quad ; \quad \mu_{\rm d}(x) = \int_0^1 p^k (1-p)^{n-k} dp \quad ; \quad \mu_{\rm k}(x) = \frac{1}{n+1} \cdot \frac{1}{\binom{n}{k}}$$

where $k = |x|_1$ is the number of 1s in x.

Proposition 1 The d-uniform distribution μ_d and the k-uniform distribution μ_k are equivalent.

This equality is established by identifying $\mu_{\rm d}(x)$ to the values of the so-called 'Beta function' or 'Euler's integral of first kind'. An original proof is given in Appendix A.

Given a ring size n, a distribution μ on \mathcal{E}_n , the quality of a classifier \mathcal{C} is defined by:

$$Q(n) = \sum_{x \in \mathcal{E}_n} \mathsf{G}(x) \cdot \mu(x)$$

In this paper, we evaluate the performance of a classifier using a d-binomial quality $Q_{\rm b}$, defined with the distribution $\mu_{\rm b}$ and a d-uniform quality $Q_{\rm d}$, defined with the distribution $\mu_{\rm c}$. Intuitively, we see that for most classifiers, we will have $Q_{\rm d} > Q_{\rm b}$. Indeed, when we take the d-binomial distribution, as n grows to infinity, most initial configurations of \mathcal{E}_n have a density close to 1/2 and are generally more difficult to classify that configuration with densities close to 0 or 1. The d-uniform distribution avoids this difficulty by assigning an equal chance to appear to all the initial densities.

Similarly, we define the average classification times with regard to distribution μ as

$$T_{\mu} = \sum_{x \in \mathcal{E}_n} \mathbb{E}\big\{\,T(x)\,\big\} \mu(x)$$

We denote by $T_{\rm b}$ and $T_{\rm d}$ the average classification time obtained with the $\mu_{\rm b}$ and $\mu_{\rm d}$ distributions, respectively. As for most classifiers, we have $T_{\rm d} < T_{\rm b}$, we will be only interested in estimating $T_{\rm b}$.

1.3 Structure of the s-ECA space

Obviously, the classical deterministic ECA are particular s-ECA with a local rule that takes its values in $\{0, 1\}$. The space of s-ECA can be described as an eight-dimensional hypercube with the ECA in its corners. This can be perceived intuitively if we see s-ECA rules as points of a hypercube, to which we apply the operations of addition and multiplication. More formally, taking p s-ECA F_1, \ldots, F_k and w_1, \ldots, w_k real numbers in [0, 1] such that $\sum_{i=0}^k w_i = 1$, the weighted average of the s-ECA (F_i) with weights w_i is the s-ECA g defined with:

$$\forall x, y, z \in \{0, 1\}, g(x, y, z) = \sum_{i=0}^{k} w_i \cdot f(x, y, z)$$

Table 1: Table of the 8 active transitions and their associated letters. The transition code of an ECA is the sequence of letters of its active transitions.

A	В	С	D	E	F	G	Н
000	001	100	101	010	011	110	111
1	1	1	1	0	0	0	0

As a consequence, one may choose to express an s-ECA as a weighted average of other ECAs. The most intuitive basis of the s-ECA is formed by the 8 ECA that have only one transition that leads to 1: the coordinates correspond to the values f(x, y, z). Equally, one may express a s-ECA as a weighted average of the 8 (deterministic) ECA that have only one *active* transition, *i.e.*, only one change of state in their transition table. Such ECA are labelled A, B, ..., H according to the notation introduced summed up in Tab. 1. Formally, for every s-ECA f, there exists a 8-tuple $(p_A, p_B, \ldots, p_H) \in [0, 1]^8$ such that: $f = p_A \cdot A + p_B \cdot B + \cdots + p_H \cdot H$. We denote this relationship by $f = [p_A, p_B, \ldots, p_H]_T$, where the subscript T stands for (active) "transitions".

This basis many advantages for studying the random evolution of configurations (see Ref. [6]). For instance, the group of symmetries of a rule can easily be obtained: the left-right symmetry permutes $p_{\rm B}$ and $p_{\rm C}$, and $p_{\rm F}$ and $p_{\rm G}$, whereas the 0-1 symmetry permutes $p_{\rm A}$ and $p_{\rm H}$, $p_{\rm B}$ and $p_{\rm G}$, etc. As we will now illustrate, this notation also allows us to easily write the conservation laws of a stochastic CA and to estimate some aspects of its global behaviour.

2 Fukś Density Classifier

To start examining how stochastic CA solve the density classification problem, let us first consider the probabilistic density classifier proposed by Fukś [8]. For $p \in [0, 1/2]$, the local rule **C1** is defined with the following transition table:

This rule is a density classifier as 0 and 1 are its only fixed points. With notations introduced above, we write:

$$C1 = [0, p, p, 2p, 2p, p, p, 0]_{T}$$
$$= p \cdot BDEG + p \cdot CDEF$$

where the rules³ BDEG(170) and CDEF(240) are the left and right shift respectively. This means that Fuks' rule can be interpreted as applying, for each cell independently: (a) a left shift with probability p, (b) a right shift with probability p, and (c) staying in the same state with probability 1 - 2p (see Fig. 1).

 $^{^{3}}$ We give the "classical" rule code into parenthesis ; it is obtained by converting the series of 8 bits of the transition table (000 to 111) to the corresponding decimal number.



Figure 1: Space-time diagrams showing the evolution of the **C1**, **C2**, **C3** classifiers with n = 39, and same initial density ~ 0.4. Time goes from bottom to top ; white cells are 0-cells and blue cells are 1-cells. (left & middle) evolution will most probably end with a good classification (0); (right) evolution will with end with a good classification with probability 1 (an archipelago has been reached).

We also note that this rule is invariant under both the left-right and the 0-1 permutations (as $p_{\rm B} = p_{\rm C} = p_{\rm F} = p_{\rm G}$, $p_{\rm A} = p_{\rm H}$ and $p_{\rm D} = p_{\rm E}$).

Theorem 1 For the classifier C1 set with $p \in (0, 1/2]$,

$$\forall x \in \mathcal{E}_n, \ \mathcal{G}(x) = \max \left\{ \rho(x), 1 - \rho(x) \right\} \quad and \quad T_{\rm b} \le n^2/2p$$

The relationship on G(x) was observed experimentally with simulations and partially explained by combinatorial arguments [8]. As for the classification time of the system, no predictive law was given. We now propose a proof that uses the analytical tools developed for asynchronous ECAs [6] and completes the results established by Fukś. The proof stands on the following lemma:

Lemma 1 For a sequence of random variables $(x^t)_{t\in\mathbb{N}}$ that describes the evolution of a stochastic CA with the initial condition $x \in \mathcal{E}_n$, let M be a mapping $M : \mathcal{E}_n \to \{0, \ldots, m\}$ and let (X_t) be the sequence of random variables defined by $\forall t, X_t = M(x^t)$. If X_t and $\Delta X_{t+1} = X_{t+1} - X_t$ verify:

- The stochastic process (X_t) is a martingale on $\{0, \ldots, m\}$, that is, for a filtration \mathcal{F}_t adapted to (X_t) , $\mathbb{E}\{\Delta X_{t+1} | \mathcal{F}_t\} = 0$.
- $X_t \in \{1, \dots, m-1\} \implies \operatorname{var}\{\Delta X_{t+1}\} > v$

then:

$$\Pr\{X_T = m\} = \frac{q}{m}$$

and the absorbing time of the process $T = \min\{t : X_t = 0 \text{ or } X_t = m\}$ is finite and obeys:

$$\mathbb{E}\{T\} \le \frac{q(m-q)}{v} \le \frac{m^2}{2v}$$

where $q = \mathbb{E}\{X_0\} = M(x)$.

Proof. [Sketch] A similar lemma was formulated for studying asynchronous CA [6]. The main elements of its proof are: (1) to note that T is stopping time, (2) to use the Optional Stopping Time theorem to calculate $E[X_T]$, (3) to note that the process $Y_t = X_t^2 - v \cdot t$ is a submartingale and use again the Optional Stopping Time theorem. (See Appendix B). \Box

Proof. [Proof of Theorem 1] We simply take $X_t = |x^t|_1$ and show that Lemma 1 applies to X_t . For sake of simplicity, we write $a(x) = |x|_{000}$, $b(x) = |x|_{001}$, ..., $h(x) = |x|_{111}$ (see Tab. 1) and drop the argument x when there is no ambiguity. The following equalities hold [6]:

$$b + d = e + f; \quad c + d = e + g \quad ; b = c = f = g$$
 (2)

This allows us to write:

$$\mathbb{E}\left\{\Delta X_{t+1} \left| \mathcal{F}_t \right.\right\} = p.b + p.c + 2p.d - 2p.e - p.f - p.g$$
$$= p.(b + d - e - f) + p.(c + d - e - g)$$

Using Eq. (2), we obtain $\mathbb{E}\left\{\Delta X_{t+1} | \mathcal{F}_t\right\} = 0$.

Second, we assume that $X_t \in \{1, \ldots, n-1\}$. It implies that $x^t \notin \{0, 1\}$, that is, x^t is not a fixed point. Denoting by \tilde{A} , \tilde{B} , the cells where transitions A, B, etc. apply, and given that transitions B, C, D (resp. E, F, G) increase (resp. decrease) ΔX_{t+1} by one, we write:

$$\Delta X_{t+1} = \sum_{c \in \tilde{B}, \tilde{C}} \mathcal{B}_c^t(p) + \sum_{c \in \tilde{D}} \mathcal{B}_c^t(2p) - \sum_{c \in \tilde{E}} \mathcal{B}_c^t(2p) + \sum_{c \in \tilde{F}, \tilde{G}} \mathcal{B}_c^t(p)$$
(3)

where (\mathcal{B}_c^t) is the series of the Bernoulli random variables of Eq. (1). As $\operatorname{var} \{ \mathcal{B}(p) \} = p(1-p)$, Eq. (3) gives:

$$\operatorname{var} \left\{ \Delta X_{t+1} \right\} = (b+c+f+g) \cdot p(1-p) + (d+e) \cdot 2p(1-2p) = p \cdot [(s_1+2s_2) - (s_1+4s_2) \cdot p]$$

with $s_1 = b + c + f + g$ and $s_2 = d + e$. Using Eq. (2) and noting that the value of n is odd, we remark that there exists a 00 or 11 pattern and that $s_1 = b + c + f + g \ge 2$. From $p \le 1/2$, we obtain $(s_1 + 2s_2) - (s_1 + 4s_2) \cdot p \ge 1$ and $\operatorname{var} \{ \Delta X_{t+1} \} \ge p$.

The lemma thus applies by taking v = p and m = n. We finally find that the probability that the process stops on $X_T = n$, that is, on the fixed point **1**, is equal to the initial density $\rho(x) = |x|_1/n$. We also find that :

$$\forall x \in \mathcal{E}_n, T(x) \le \frac{|x|_1 (n - |x|_1)}{p}$$
 and $T_{\rm b} \le n^2/2p$

From this result, we derive that the probability of good classification of any configuration x is equal to $G(x) = \max\{\rho(x), 1 - \rho(x)\}$. The *d*-uniform quality of **C1** is thus equal to $Q_d(n) = 3/4$ (obtained y a simple integration). For n = 2k + 1, the *d*-binomial quality of **C1** is equal to (see Appendix C) $Q_b(n) = 1/2 + \binom{2k}{k}/2^{2k+1}$. This formula explains why the quality of classification of **C1** quickly decreases as the ring size n increases. For instance for n = 49, we have: $Q_b(n) = 0.557$, that is, the gain compared to a random guess is less than 6%. For the reference value n = 149, the gain drops down to 3.3% (see Tab. 2 p. 13).

3 Schüle Density Classifier

We now consider the probabilistic density classifier proposed by Schüle et al [15]. For $\varepsilon \in (0, 1]$, the local rule **C2** is defined with the following transitions:

> (x, y, z)000 101 110 111 001 010 011 100 f(x, y, z)0 $1-\varepsilon$ $1-\varepsilon$ ε $1-\varepsilon$ ε ε 1

This rule is a density classifier as 0 and 1 are its only fixed points. With notations introduced above, we write:

$$\begin{aligned} \mathbf{C2} &= [0, 1 - \varepsilon, 1 - \varepsilon, \varepsilon, \varepsilon, 1 - \varepsilon, 1 - \varepsilon, 0] \\ &= (1 - \varepsilon) \cdot \mathsf{BCFG} + \varepsilon \cdot \mathsf{DE} \end{aligned}$$

where rule $\mathsf{BCFG}(150)$ is the rule that implements a XOR function with three neighbours and DE is the majority rule. This means that Schüle's rule can be interpreted as applying for each cell independently: (a) a XOR with probability $1-\varepsilon$ (b) a majority with probability ε (see Fig. 1). This rule is invariant under both the left-right and the 0-1 symmetries (as we have: $p_{\rm B} = p_{\rm C} = p_{\rm F} = p_{\rm G}$, $p_{\rm A} = p_{\rm H}$ and $p_{\rm D} = p_{\rm E}$).

Theorem 2 For the classifier C2, for $\varepsilon = 2/3$,

$$\forall x \in \mathcal{E}_n, \ \mathcal{G}(x) = \max \left\{ \rho(x), 1 - \rho(x) \right\} \quad and \quad T_{\rm b} \le 9/2 \cdot n^2$$

This property was proved under the mean-field approximation in [15]. We now propose to use the same techniques as for Fuks' rule.

Proof. Let us take $X_t = |x^t|_1$ and show that Lemma 1 applies to X_t . We have:

$$\mathbb{E}\left\{\Delta X_{t+1}\right\} = (1-\varepsilon)(b+c-f-g) + \varepsilon(d-e)$$
$$= \varepsilon(b+c-d+e-f-g) + d-e$$

Using Eq. (2), we obtain:

$$\mathbb{E}\left\{\Delta X_{t+1}\right\} = (3\varepsilon - 2)(d - e) \tag{4}$$

which leads to $\mathbb{E}\left\{\Delta X_{t+1}\right\} = 0$ for $\varepsilon = 2/3$.

Let us now assume that $X_t \in \{1, \ldots, n-1\}$. This implies that x^t is not a fixed point and that $b + c + d + e + f + g \ge 1$. Recall that we denote by \tilde{A} , \tilde{B} , the cells where transitions A, B, etc. apply. We have:

$$\Delta X_{t+1} = \sum_{c \in \tilde{B}, \tilde{C}} \mathcal{B}_c^t(1-\varepsilon) + \sum_{c \in \tilde{D}} \mathcal{B}_c^t(\varepsilon) - \sum_{c \in \tilde{E}} \mathcal{B}_c^t(\varepsilon) + \sum_{c \in \tilde{F}, \tilde{G}} \mathcal{B}_c^t(1-\varepsilon)$$

where (\mathcal{B}_c^t) is the series of Bernoulli random variables of Eq. (1). This results in:

$$\operatorname{var}\left\{\Delta X_{t+1}\right\} = (b+c+d+e+f+g) \cdot \varepsilon(1-\varepsilon)$$

$$\geq \varepsilon(1-\varepsilon)$$

The lemma thus applies by taking $v = \varepsilon(1 - \varepsilon)$ and m = n. Consequently, we find that the probability that the process stops on the fixed point **1** (given by $X_T = n$) is equal to $X_0/n = \rho(x)$ and that $\forall x \in \mathcal{E}_n, T(x) \leq \frac{x_1(n-|x|_1)}{\epsilon(1-\epsilon)}$, which implies $T_{\rm b} \leq \frac{n^2}{2\epsilon(1-\epsilon)} \leq \frac{9}{2} \cdot n^2 \square$

Equation (4) also allows us to understand the general behaviour of Schüle's classifier **C2** for $\varepsilon \neq \frac{2}{3}$. Informally, let us consider a configuration x with a density close to 1. For such a configuration, we most likely have more isolated 0s than isolated 1s, that is, d-e > 0 and the sign of ΔX_{t+1} is the same as $3\epsilon - 2$. As for such configuration, we want the density to *increase*, we see that setting $\epsilon > 2/3$ drives the system more rapidly towards its goal. This also explains why for $\epsilon < 2/3$, it was no longer possible to observe the system convergence within "reasonable" simulation times. In fact, the system is then in a metastable state: although the classification time is finite, the system is always attracted towards configurations with a density close to 1/2. This phenomenon was observed by the authors [15]. Last, but not least, we think that for $\epsilon > 2/3$, only isolated 0s or 1s of the *initial* configuration contribute to driving the system to its goal. This lead us to formulate the following statement:

Proposition 2 For the classifier C2 set with $\varepsilon > 2/3$, the quality of classification $Q_{\rm b}(n)$ is bounded. More precisely:

$$\forall \epsilon > 2/3, \forall x \in \mathcal{E}_n : |x|_{\mathsf{010}} = |x|_{\mathsf{101}} = 0, \, \mathcal{G}(x) = \max\{\rho(x), 1 - \rho(x)\}$$

and

$$\forall \epsilon > 2/3, \forall x \in \mathcal{E}_n, \mathcal{G}(x) \le \max\{\rho^*(x), 1 - \rho^*(x)\}$$

where $\rho^*(x) = (\Phi_{MAJ}^{\infty}(x))$ is the density attained by an asymptotic evolution of x under the majority rule.

Theses hypotheses are confirmed by numerical simulations (see Tab. 2 and Appendix E). We also verified experimentally that for $\epsilon \to 1$, the quality approaches an asymptotic limit while the average classification time diverges. We leave a rigorous proof this statement for future work (see Appendix D) and present a rule that does not suffer from such limitations.

4 A New Rule for Density Classification

For $\delta \in (0, 1]$, let us consider the following s-ECA:

(x, y, z)000 001 010 011 100 101 110 111 f(x, y, z)0 0 0 1 $1 - \delta$ 1 δ 1

With the transition code, this writes:

C3 =
$$[0, 0, 1 - \delta, 1, 1, 0, 1 - \delta, 0]_T$$

= δ .DE + $(1 - \delta)$.CDEG

For $\delta = 0$ we have CDEG(184), which is a well-known rule, often called the "traffic" rule. This rule is number conserving, *i.e.*, the number of 1s is conserved as the system evolves (see *e.g.*, [2]). Observing the evolution of the rule, we see that a 1 with a 0 at its right moves to right while a 0 with a 1 at its left is moved to the left. So all happens if the 1s were cars that tried to go to the right, with possible traffic jams. These jams resorb by going in the inverse directions of the cars (when possible). For $\delta = 1$, we have the majority rule DE(232). For $\delta \in (0, 1)$, the effect of the rule is the same as applying, for each cell and at each time step, the traffic rule with probability $1 - \delta$ and the majority rule with probability δ (see Fig. 1). This combination combination of rules has a surprising property: although the system is stochastic, there exists an infinity of configurations that can be classified with no error.

Lemma 2 An archipelago is well-classified with probability 1.

Proof. The proof is simple and relies on two observations.

First, let us note that the successor of a q-archipelago is a q-archipelago. To see why this hold, let us denote by y the successor of a configuration $x \in \mathcal{L}$. Let $C = \{c \in \mathcal{L} : y_c = 1\}$ be the 1-cells in y. If we look in x at the local predecessor pattern of a cell $c \in C$, we have $(x_{c-1}, x_c, x_{c+1}) \in \{100, 101, 011, 111\}$ by examining the transition function of C3, and, as x is an archipelago, $(x_{c-1}, x_c, x_{c+1}) \in \{100, 101\}$. As these two patterns do not overlap, it is not possible to have two successive cells of \mathcal{L} contained in C and y is an archipelago.

Second, without loss of generality, let us assume that x is a 1-archipelago, we then have d(x) < 1/2. We remark that the number of 1s of x^t is a non-increasing function of t. At each time step, each isolated 1 can "disappear" if transition C is not applied, which happens with probability $\delta > 0$. As a result, all the 1s will eventually disappear and the system will attain the fixed point $\mathbf{0}$, which corresponds to a good classification. \Box

The second interesting property with C3 is its ability to make any configuration evolve to an archipelago with a given probability.

Lemma 3 For every $p \in [0,1)$, there exists a setting δ of the classifier C3 such that for every configuration $x \in \mathcal{E}_n$ such that $\rho(x) < 1/2$ (resp. $\rho(x) > 1/2$), the probability to evolve to a 1-archipelago (resp. to a 0-archipelago) is greater than p.

Proof. The proof relies on the well-known property of the traffic rule to evolve to an archipelago in at most n/2 steps. Let us denote by Φ the global transition function of CDEG and write $y^t = \Phi^t(x)$, that is, (y_t) is the series of configurations obtained with x as an initial condition. From the properties of the traffic rule, we have that $d(y_t) = d(x)$ and that $y^{\lceil n/2 \rceil}$ is an archipelago (see *e.g.*, Ref. [3] Lemma 4).

For a given p and given n, without loss of generality, let us consider a configuration x such that d(x) < 1/2. Let us now evaluate the probability that rule **C3** does *not* behave like the traffic rule in the first $T = \lceil n/2 \rceil$ steps. Formally let $D^t = card\{c \in \mathcal{L} : x_c^t \neq y_c^t\}$.

Comparing the transition rules of CDEG and C3, we see that differences in the evolution of the two rules can only occur for cells where transitions B and F apply, that is, cells that have a 100 and 110 neighbourhood. As we have b = f = g = c, and $b + f + g + c \leq n$, we write $b + f \leq \lceil n/2 \rceil$. For such cells, differences of evolution occur with a probability δ , which implies that, at each time step, the probability $p_{\text{diff}} = \Pr\{D^t > 0\}$ that the evolution of C3 and CDEG differ on T steps is upper-bounded by: $p_{\text{diff}} \leq \delta^{T \cdot \lceil n/2 \rceil} \leq \delta^{T^2}$. The probability $P_{\text{eq}} = \Pr\{D^1 = 0, \ldots, D^T = 0\}$ that the two rules evolve identically on T steps is thus greater or equal than $1 - p_{\text{diff}}$ and we find that it is sufficient to set: $\delta < 1 - p^{\frac{1}{T^2}}$ to guarantee that $P_{\text{eq}} > p$, *i.e.*, that the probability to reach a 1-archipelago is greater than p. \Box

This inequality shows that, by taking δ small enough, the probability that a configuration x with d(x) < 1/2 evolves to a 1-archipelago can be made arbitrarily small. This allows us to state our main result.

Theorem 3 For all $p \in [0,1)$, there exists a setting δ of the classifier C3 such that $\forall x \in \mathcal{E}_n$, $G(x) \ge 1-p$. As a consequence, for all ring sizes $n \in 2\mathbb{Z}+1, \delta \to 0$ implies $Q_{\rm b}(n) \to 1$.

Proof. Combining the two previous lemmas to prove the theorem is straightforward: for δ small enough, the system evolves to an archipelago that has the same density as the initial condition (Lemma 3). It is then necessarily well-classified as it will progressively "drift" towards the appropriate fixed point (Lemma 2). However, we remark that the time taken to reach the fixed point increases exponentially as δ decreases. \Box

The analytical estimation of the quality of **C3** and its time of convergence is more complex than for Fukś and Schüle classifiers. Table 2 shows the values of $Q_{\rm b}$, $Q_{\rm d}$ and $T_{\rm b}$ estimated by numerical simulations. We can observe that the *d*-uniform quality rapidly increases to high values, even when keeping the average convergence time to a few thousand steps. In particular for $\delta < 1\%$, the quality goes above the symbolic rate of 90%, which to our knowledge, has not been yet reached for one-dimensional systems (see *e.g.* [4, 12]). Another major point regards the classification time of **C3**: it is experimentally determined as varying linearly with the ring size *n* (see Appendix E).

siniulations.									
model	setting	$Q_{\rm b}~({\rm in}\%)$	$Q_{\rm d} \ ({\rm in}\%)$	$T_{\rm b}$					
C1	p = 0.25	53.3	75.0	4638					
C1	p = 0.48	53.3	75.0	2652					
C1	p = 0.5	53.3	75.0	8985					
C2	$\epsilon = 0.7$	54.0	80.1	4061					
C2	$\epsilon = 0.8$	55.1	83.8	6223					
C2	$\epsilon = 0.9$	56.6	85.8	11887					
C3	$\delta = 0.1$	82.4	98.1	517					
C3	$\delta = 0.01$	91.0	99.1	4950					
C3	$\delta = 0.005$	93.4	99.3	9981					

Table 2: Synthesis of results for n = 149; averages obtained with 10 000 simulations.

Conclusion

This paper presented an analysis of three parametrised stochastic cellular automata that solve the density classification problem. We first re-examined the behaviour of two previously studied rules using the analytical techniques originally developed for asynchronous cellular automata. This allowed us to analytically estimate their quality of classification and average convergence time.

We then presented a new parametrised rule that solves this problem with an arbitrary precision. This shows that no upper bound exists on how well a rule can perform, at least for stochastic cellular automata. Numerical simulations were performed to estimate the behaviour of the rules for various settings. It was observed that for the standard reference value n = 149, quality of classification of over 90% could be attained, with an average convergence time of the order of a few thousand steps.

Far giving a definitive answer to the problem, the existence of this solution suggests that the quality of classification cannot be taken as the unique criterion for evaluating the classifiers. Instead, it is some trade-off between quality and time to classify that has now to be searched for. We are particularly interested in knowing whether there are other couples (or t-uples) of rules that can be "blended" to obtain similar or even better results. The "blending" technique could also be used with rules that have a larger radius, and on lattices with two or more dimensions.

References

- Ramón Alonso-Sanz and Larry Bull, A very effective density classifier twodimensional cellular automaton with memory, Journal of Physics A 42 (2009), no. 48, 485101.
- [2] Nino Boccara and Henryk Fuks, Number-conserving cellular automaton rules, Fundamenta Informaticae 52 (2002), no. 1-3, 1–13.

- [3] Mathieu S. Capcarrere, Moshe Sipper, and Marco Tomassini, Two-state, r = 1 cellular automaton that classifies density, Phys. Rev. Lett. 77 (1996), no. 24, 4969–4971.
- [4] Pedro P.B. de Oliveira, José C. Bortot, and Gina M.B. Oliveira, The best currently known class of dynamically equivalent cellular automata rules for density classification, Neurocomputing 70 (2006), no. 1-3, 35 – 43.
- [5] Paula Gonzaga de Sá and Christian Maes, The Gacs-Kurdyumov-Levin automaton revisited, Journal of Statistical Physics 67 (1992), 507–522.
- [6] Nazim Fatès, Michel Morvan, Nicolas Schabanel, and Eric Thierry, Fully asynchronous behavior of double-quiescent elementary cellular automata, Theoretical Computer Science 362 (2006), 1–16.
- Henryk Fukś, Solution of the density classification problem with two cellular automata rules, Physical Review E 55 (1997), no. 3, R2081–R2084.
- [8] _____, Nondeterministic density classification with diffusive probabilistic cellular automata, Physical Review E 66 (2002), no. 6, 066106.
- [9] Mark Land and Richard K. Belew, No perfect two-state cellular automata for density classification exists, Physical Review Letters 74 (1995), no. 25, 5148–5150.
- [10] Claudio L.M. Martins and Pedro P.B. de Oliveira, Evolving sequential combinations of elementary cellular automata rules, Advances in Artificial Life (Mathieu S. Capcarrere, Alex A. Freitas, Peter J. Bentley, Colin G. Johnson, and Jon Timmis, eds.), Lecture Notes in Computer Science, vol. 3630, Springer Berlin Heidelberg, 2005, pp. 461–470.
- [11] Melanie Mitchell, James P. Crutchfield, and Peter T. Hraber, Evolving cellular automata to perform computations: Mechanisms and impediments, Physica D 75 (1994), 361–391.
- [12] Gina M. B. Oliveira, Luiz G. A. Martins, Laura B. de Carvalho, and Enrique Fynn, Some investigations about synchronization and density classification tasks in one-dimensional and two-dimensional cellular automata rule spaces, Electronic Notes in Theoretical Computer Science 252 (2009), 121–142.
- [13] Norman H. Packard, Dynamic patterns in complex systems, ch. Adaptation toward the edge of chaos, p. 293301, World Scientific, Singapore, 1988.
- [14] Leonid A. Levin Peter Gács, Georgii L. Kurdiumov, One-dimensional homogeneous media dissolving finite islands, Problemy Peredachi Informatsii 14 (1987), 92–96.

- [15] Martin Schüle, Thomas Ott, and Ruedi Stoop, Computing with probabilistic cellular automata, ICANN '09: Proceedings of the 19th International Conference on Artificial Neural Networks (Berlin, Heidelberg), Springer-Verlag, 2009, pp. 525–533.
- [16] Christopher Stone and Larry Bull, Evolution of cellular automata with memory: The density classification task, BioSystems 97 (2009), no. 2, 108– 116.

A Equivalence of μ_d and μ_k

Let us show the equality:

$$\forall x \in \mathcal{E}_n, \mu_{\mathrm{d}}(x) = \int_0^1 \rho^k (1-\rho)^{n-k} d\rho = \frac{1}{n+1} \cdot \frac{1}{\binom{n}{k}} = \mu_{\mathrm{k}}(x)$$

Proof. The first technique is to remark that $\mu_d(x)$ corresponds to the calculus of the Beta function $\mathcal{B}(k+1, n-k+1)$, which is defined on real numbers. The analytical form of the function was studied by Euler and Legendre and it is generally obtained by applying a change of variable and using trigonometric relations. We propose here an original technique that exploits the fact that we want too evaluate this function only on integer numbers. For $i \in \{0, \ldots, n\}$, let us introduce $M_i = \int_0^1 \binom{n}{i} p^i (1-p)^{n-i} dp$. By noting that:

$$1^{n} = (1 + (1 - p))^{n} = \sum_{i=0}^{n} \binom{n}{i} p^{i} (1 - p)^{n-i}$$

we find that

$$\sum_{i=0}^{n} M_{i} = \int_{0}^{1} 1 dp = 1$$

On the other hand, using the integration by parts of M_i gives

$$M_{i} = \int_{0}^{1} {n \choose i} p^{i} (1-p)^{n-i} dp$$

= ${n \choose i} \cdot \left\{ \left[\frac{-1}{i+1} p^{i} (1-p)^{n-i-1} \right]_{0}^{1} + \int_{0}^{1} \frac{n-i}{i+1} p^{i+1} (1-p)^{n-(i+1)} dp \right\}$

By noting that the first term is equal to 0 and that $\binom{n}{i} \cdot \frac{n-i}{i+1} = \binom{n}{i+1}$, we obtain the equality $M_i = M_{i+1}$, which gives $M_i = \frac{1}{n+1}$ and confirms the result $\mu_d = \mu_k$. \Box

B Convergence Lemma 1

Proof. First, let us show that T is finite. We remark that, as ΔX_{t+1} is a process which expectancy is 0 and which variance is greater than v, there exists

a constant p such that $t < T \implies \Pr{\{\Delta X_{t+1} \leq 1\} > p}$, that is, the probability that X_t decreases by at least one is non zero. For all t < T, if we look msteps ahead, the probability that X_t makes m successive decreases or attains the fixed point 0 before is then non-zero, say α . The time of convergence to a fixed point is thus upper-bounded by the time of first success of a geometric law of parameter α , it is thus finite.

As T is finite, we apply Doob's optional stopping theorem on the martingale (X_t) and write

$$\mathbb{E}\left\{X_T\right\} = \mathbb{E}\left\{X_0\right\} = q$$

On the other hand we have:

$$\mathbb{E}\left\{X_T\right\} = 0 \cdot \Pr\{X_T = 0\} + m \cdot \Pr\{X_T = n\}$$

Combining the two equations gives us : $\Pr{X_T = n} = q/m$.

To calculate an upper bound on T, let us consider the following process: $Y_t = X_t^2 - v \cdot t$. For t < T, we have:

$$\mathbb{E}\left\{Y_{t+1} - Y_t \left|\mathcal{F}_t\right.\right\} = \mathbb{E}\left\{X_{t+1}^2 - X_t^2 \left|\mathcal{F}_t\right.\right\} - v$$
$$= \mathbb{E}\left\{\left(\Delta X_{t+1} + X_t\right)^2 - X_t^2\right) \left|\mathcal{F}_t\right.\right\} - v$$
$$= \mathbb{E}\left\{\Delta X_{t+1}^2 \left|\mathcal{F}_t\right.\right\} + 2 \cdot \mathbb{E}\left\{\Delta X_{t+1} \cdot X_t \left|\mathcal{F}_t\right.\right\} - v$$

We also have :

$$\operatorname{var}\left\{\Delta X_{t+1}\right\} = \mathbb{E}\left\{\Delta X_{t+1}^{2}\right\} - \mathbb{E}\left\{\Delta X_{t+1}\right\}^{2} = \mathbb{E}\left\{\Delta X_{t+1}^{2}\right\}$$

and

$$\mathbb{E}\left\{\Delta X_{t+1} \cdot X_t \left| \mathcal{F}_t \right.\right\} = \mathbb{E}\left\{\Delta X_{t+1} \left| \mathcal{F}_t \right.\right\} \cdot X_t = 0$$

Finally, using the Lemma hypothesis var $\{\Delta X_{t+1}\} \ge v$, we obtain: $\mathbb{E}\{Y_{t+1} - Y_t\} \ge 0$, that is, (Y_t) is a submartingale. Again, applying the Optional Stopping Theorem gives:

$$\mathbb{E}\left\{Y_T\right\} \ge \mathbb{E}\left\{Y_0\right\}$$

which reads:

$$\mathbb{E}\left\{X_T^2 - v \cdot T\right\} \ge q^2$$

or

$$\mathbb{E}\left\{T\right\} \le \frac{\mathbb{E}\left\{X_T^2\right\} - q^2}{v}$$

To calculate $\mathbb{E}\left\{X_T^2\right\}$, we write:

$$\mathbb{E}\{X_T^2\} = 0^2 \cdot \Pr\{X_T = 0\} + m^2 \cdot \Pr\{X_T = m\} = m^2 \frac{q}{m} = m \cdot q$$

which gives us the expected result: $\mathbb{E}\left\{T\right\} \leq \frac{q(m-q)}{v}$

C *d*-binomial quality of C1

Proof. We take an odd ring size n = 2k + 1. By definition of the *d*-binomial quality $Q_{\rm b}$:

$$Q_{\rm b}(n) = \frac{1}{2^n} \sum_{x \in \mathcal{L}} \max\{d(x), 1 - d(x)\}$$

Grouping the configurations with equal number of 1s gives:

$$Q_{\rm b}(n) = \frac{1}{2^n} \sum_{i=0}^n \binom{n}{i} \max\{\frac{i}{n}, 1 - \frac{i}{n}\}$$

which we decompose into:

$$Q_{\rm b}(n) = \frac{1}{2^n} \sum_{i=0}^k \binom{n}{i} \frac{n-i}{n} + \sum_{i=k+1}^n \binom{n}{i} \frac{i}{n} n$$

Using the symmetry $\binom{n}{i} = \binom{n}{n-i}$, we obtain:

$$Q_{\rm b}(n) = \frac{1}{2^n} \cdot 2 \cdot \sum_{i=k+1}^n \binom{n}{i} \frac{i}{n} = \frac{1}{2^{n-1}} \sum_{i=0}^k \binom{n-1}{i}$$

To calculate the sum, we write n-1 = 2k and: $\sum_{i=0}^{k} \binom{2k}{i} = \sum_{i=0}^{k} \binom{2k}{2k-i} = \sum_{i=k}^{2k} \binom{2k}{i}$. which gives:

$$2.\sum_{i=0}^{k} \binom{2k}{i} = \sum_{i=0}^{k} \binom{2k}{i} + \binom{2k}{k} + \sum_{i=k+1}^{n} \binom{2k}{i} = 2^{2k+1} + \binom{2k}{k}$$

We then obtain:

$$Q_{\rm b}(n) = \frac{1}{2^{2k}} \frac{2^k + \binom{2k}{k}}{2} = \frac{1}{2} + \frac{\binom{n-1}{(n-1)/2}}{2^n}$$

D Justification of Proposition 2

Intuitively, let us consider a configuration x^t that has no isolated 0 or 1. Equation (4) then reads $\mathbb{E}\left\{\Delta X_{t+1} | \mathcal{F}_t\right\} = 0$. At time t + 1, isolated 0s or 1s may appear (if transitions B and F, or C and G, are applied simultaneously) but there is an equal chance to make an isolated 0 than an isolated 1. As this

analysis can be repeated recursively, this shows that once the first isolated 0s or 1s have disappeared (with the effect of the majority rule), the boundary of the homogeneous regions randomly drift until they merge. As a consequence, the probability of good classification of a configuration with no isolated 0 or 1 is the same as for Fukś classifier. This behaviour implies that the average convergence times scales as n^2 , which is confirmed experimentally (see above).

E Scaling law of $T_{\rm b}$



Figure 2: Average classification time $T_{\rm b}$ as a function of ring size n for different settings of C2 and C3. inset: same data with logscale