Forty-Nine Years of the Garden-of-Eden Theorem

Silvio Capobianco

Institute of Cybernetics at TUT

Theory Days at Tõrve
October 7-8-**9**, 2011

Joint work with Pierre Guillon [1] and Jarkko Kari [2]

Revision: October 11, 2011

[1]CNRS & IML Marseille
[2]University of Turku

Outline of the talk

1. The context of the theorem
2. The theorem
3. Similar theorems
4. Range of validity
5. Our results

Conway's Game of Life

Invented by John Horton Conway, popularized by Martin Gardner.[3]
The checkboard is an infinite square grid.
Each case (cell) of the checkboard is "surrounded" by those within a chess' king's move, and can be "living" or "dead".

1. A dead cell surrounded by exactly three living cells, becomes living.
2. A living cell surrounded by two or three living cells, survives.
3. A living cell surrounded by less than two living cells, dies of isolation.
4. A living cell surrounded by more than three living cells, dies of overpopulation.

[3] *Sci. Am.* **223**, October 1970)

Cellular automata

A cellular automaton (CA) on a group G is a triple $\mathcal{A} = \langle Q, \mathcal{N}, f \rangle$ where:

- Q is a finite set of states.
- $\mathcal{N} = \{n_1, \ldots, n_k\} \subseteq G$ is a finite neighborhood index.
- $f : Q^k \to Q$ is a finitary local function

The local function induces a global function $F : Q^G \to Q^G$ via

$$\begin{aligned} F(c)(x) &= f\left(c(x \cdot n_1), \ldots, c(x \cdot n_k)\right) \\ &= f\left(c|_{x\mathcal{N}}\right) \end{aligned}$$

The same rule induces a function over patterns with finite support:

$$f(p) : E \to Q \ , \ f(p)(x) = f\left(p|_{x\mathcal{N}}\right) \ \ \forall p : E\mathcal{N} \to Q$$

In a Garden of Eden

A Garden of Eden (briefly, GoE) for a given CA
is either an infinite configuration or a finite pattern
which cannot be produced by the CA
from another configuration or pattern.

A GoE is thus a sort of "paradise lost" which can be started from, but not returned to.

A CA has a GoE configuration
— *i.e.*, it is non-surjective —
if and only if it has a GoE pattern.

The Flower of Eden

Beluchenko, 2009. Smallest known GoE pattern for the Game of Life.

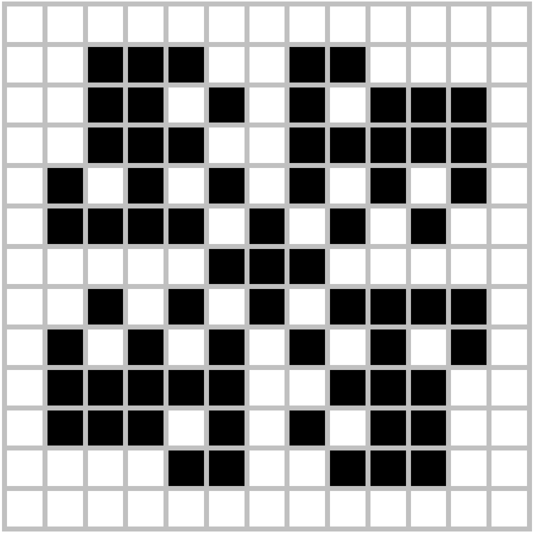

“Not injectivity, but almost”

Two distinct patterns $p, p' : E \to Q$ are mutually erasable for a CA with global rule F, if any two configurations c, c' with

$$c|_E = p \,, \quad c'\big|_E = p' \,, \text{ and } \; c|_{G \setminus E} = c'\big|_{G \setminus E}$$

satisfy $F(c) = F(c')$.

A cellular automaton without mutually erasable patterns
is called pre-injective

The Garden-of-Eden theorem (Moore, 1962)

If a cellular automaton over $\mathbb{Z} \times \mathbb{Z}$
has two mutuably erasable patterns,
then it also has a Garden of Eden pattern

Myhill's converse to Moore's theorem (1962)

If a cellular automaton over $\mathbb{Z} \times \mathbb{Z}$
has a Garden of Eden pattern,
then it also has two mutuably erasable patterns

From finite to infinite

Suppose the group of the CA is finite. Then:

$\frac{\text{pattern}}{\text{mutually erasable}}$	is the same as	$\frac{\text{configuration}}{\text{same image}}$

So Moore's GoE theorem, and its converse by Myhill, together mean that:

cellular automata on an infinite space
behave, with regard of surjectivity,
more or less as they were finitary functions.

Not completely, however: pre-injectivity is strictly weaker than injectivity.

Counterexample: XOR with the right neighbor

The thing that makes it work

Moore's and Myhill's theorems work because in $\mathbb{Z} \times \mathbb{Z}$ — and in fact, in $\mathbb{Z}^d$ for every $d \geq 1$,

the orange grows faster than the peel

- The volume of the hypercube is polynomial of degree d.
- The surface of the hypercube is polynomial of degree $d - 1$.

Consequenty:

- If a CA has two mutually erasable patterns of side ℓ,
  then it has a GoE pattern of side $M \times \ell$.
- If a CA has a GoE pattern of side ℓ,
  then it has two mutually erasable patterns of side $G \times \ell$.

The constants M and G depend on the CA.

So, what other properties are linked to GoE?

Balancedness

A cellular automaton is balanced if for any given shape E, every pattern $p : E \longrightarrow Q$ has the same number of preimages.

- For 2D CA with Moore neighborhood:
  every square pattern of side ℓ has $|Q|^{4\ell+4}$ preimages.

A balanced CA has no Garden of Eden.

The balancedness theorem (Maruoka and Kimura, 1976)

An unbalanced d-dimensional CA has a Garden of Eden.

A measure-theoretic version of balancedness

The "basic" open subsets of $Q^{\mathbb{Z}^d}$ are the cylinders of the form

$$C(p) = \left\{ c : \mathbb{Z}^d \to Q \mid c|_E = p \right\} \ , \ p : E \to Q$$

The product measure is defined by

$$\mu_\Pi(C(p)) = |Q|^{-|E|} \ , \ p : E \to Q$$

on the σ-algebra generated by the cylinders.

A cellular automaton is balanced
if and only if
it preserves the product measure,
i.e., $\mu_\Pi(F^{-1}(U)) = \mu_\Pi(U)$ for every measurable set U.

Computing opens

- Consider a computable bijection $\phi : \mathbb{N} \to \mathbb{Z}^d$.
- ϕ induces a computable, bijective enumeration B' of the cylinders.
- A family $\mathcal{U} = \{U_n\}_{n \geq 0}$ of open subsets of $Q^{\mathbb{Z}^d}$ is computable if there is a *recursively enumerable* set $A \subseteq \mathbb{N}$ such that

$$U_n = \bigcup_{\pi(n,k) \in A} B'_k \ \forall n \geq 0 ,$$

where $\pi(x, y) = \dfrac{(x+y)(x+y+1)}{2} + x$: that is, if $\mathcal{U}$ is
computably constructible from the cylinders
uniformly in the elements' index

The importance of being random

Random configurations

- A computable family $\mathcal{U} = \{U_n\}_{n \geq 0}$ of open sets is a Martin-Löf test if $\mu_\Pi(U_n) < 2^{-n}$ for every $n \geq 0$.
- A configuration c fails a M-L test $\mathcal{U}$ if $c \in \bigcap_{n \geq 0} U_n$.
- $c : \mathbb{Z}^d \to Q$ is M-L random if it does not fail any M-L test.

The world is random, almost surely

- For the set U of M-L random configurations, $\mu_\Pi(U) = 1$.
- Every pattern has an occurrence in any M-L random configuration.

Theorem (Calude *et al.*, 2001)

If a d-dimensional cellular automaton sends a M-L random configuration into one which is not, then it has a Garden of Eden.

A collection of the classical Garden-of-Eden theorems

Let $\mathcal{A}$ be a d-dimensional CA. The following are equivalent.

- $\mathcal{A}$ has a Garden of Eden.
- $\mathcal{A}$ has two mutuably erasable patterns.
- $\mathcal{A}$ is unbalanced.
- $\mathcal{A}$ does not preserve the product measure.
- $\mathcal{A}$ sends some M-L random configurations into some that are not.

Towards infinity... and beyond

- We have seen that the Garden-of-Eden theorem, and several analogous statements, hold in arbitrary dimension.
- We then trust it to be a general principle, holding for cellular automata in general, even on meshes more complicated that $\mathbb{Z}^d$.
- **. . . or do we?**

For the rest of the talk, we will work with finitely generated groups. This is not restrictive for what we want to prove.
However, we can only talk about M-L random configurations on groups that are computably bijective to $\mathbb{N}$. This is true, for instance, when the word problem is decidable.

A counterexample on the free group

Consider the following CA on the free group on two generators a, b:

- $Q = \{0, 1\}$.
- $\mathcal{N} = \{1, a, b, a^{-1}, b^{-1}\}$.
- $$f(t, x, y, z, w) = \begin{cases} 1 & \text{if } x + y + z + w = 3 \\ & \text{or } t = 1, x + y + z + w \in \{1, 2\} \\ 0 & \text{otherwise} \end{cases}$$

Theorem (Ceccherini-Silberstein et al., 1999)

- This CA does not have any GoE.
- This CA does have mutually erasable patterns.
- This CA is not balanced.

What the free group on two generators looks like

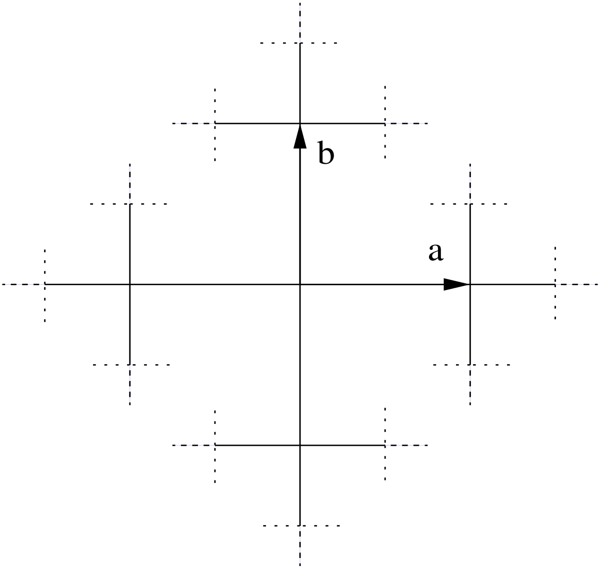

Amenable groups

A group G is amenable if it satisfies any of the following, equivalent conditions:

1. There exists a finitely additive probability measure $\mu : \mathcal{P}(G) \to [0, 1]$ such that $\mu(gA) = \mu(A)$ for every $g \in G$, $A \subseteq G$.
2. For every finite $U \subseteq G$ and every $\varepsilon \geq 0$ there exists a finite $K \subseteq G$ such that $|UK \setminus K| < \varepsilon|K|$.

We then clearly see that the free group is not amenable!

- On the other hand, amenability is still a condition of the type:

  a peel of any shape can be made
  arbitrarily proportionally small
  by choosing a suitable orange

- And in fact, $\mathbb{Z}^d$ is amenable for every $d \geq 1$.

Fact: A group is amenable iff every finitely generated subgroup is.

The importance of being amenable

Theorem (Ceccherini-Silberstein, Machì and Scarabotti, 1999)

Let G be an amenable group.

- (Moore) Every surjective CA on G is pre-injective.
- (Myhill) Also, every pre-injective CA on G is surjective.

But there are counterexamples to both in some non-amenable groups.

Theorem (Bartholdi, 2010)

Let G be a group. The following are equivalent.

- Every surjective CA on G is pre-injective.
- Every surjective CA on G preserves the product measure.
- G is amenable.

Mutual implications (2010)

property	implies	amenable	non-amenable
surjectivity	pre-injectivity	yes	no
pre-injectivity	surjectivity	yes	
surjectivity	balancedness	yes	no
surjectivity	μ_Π preserved	yes	no
balancedness	μ_Π preserved	yes	yes
μ_Π preserved	balancedness	yes	yes
surjectivity	M-L random to M-L random	yes on $\mathbb{Z}^d$	

So, what else can be said?

Special maps for special groups

Bartholdi's proof is based on a smart use of

bounded propagation 2:1 compressing maps

A b.p. 2:1 compressing map on a group G with propagation set S is a transformation $\phi : G \to G$ such that:

1. For every $g \in G$, $(\phi(g))^{-1} g \in S$.
2. For every $g \in G$, $|\phi^{-1}(g)| = 2$.

Groups with such maps are precisely those that are not amenable.

The key counterexample (Guillon, 2011)

Let G be a non-amenable group,
ϕ a bounded propagation 2:1 compressing map with propagation set S.
Define on S a total ordering $\preceq$.
Define a CA $\mathcal{A}$ on G by $Q = (S \times \{0,1\} \times S) \sqcup \{q_0\}$, $\mathcal{N} = S$, and

$$f(u) = \begin{cases} q_0 & \text{if } \exists s \in S \mid u_s q_0, \\ (p, \alpha, q) & \text{if } \exists (s,t) \in S \times S \mid s \prec t, u_s = (s, \alpha, p), u_t = (t, 1, q), \\ q_0 & \text{otherwise.} \end{cases}$$

Theorem (Capobianco, Guillon and Kari, 2011)

1. $\mathcal{A}$ has no GoE.
2. $\mathcal{A}$ is not nonwandering.
   This means that there is an open set U such that $F^t(U)$ never intersects U after $t = 0$.
3. $\mathcal{A}$ sends M-L random configurations into nonrandom ones.
   (if the group has a decidable word problem)

Mutual implications (2011)

property	implies	amenable	non-amenable
surjectivity	pre-injectivity	yes	no
pre-injectivity	surjectivity	yes	open problem[1]
balancedness	surjectivity	yes	yes
surjectivity	balancedness	yes	no
surjectivity	μ_Π preserved	yes	no
balancedness	μ_Π preserved	yes	yes
μ_Π preserved	balancedness	yes	yes
nonwandering	surjectivity	yes	yes
surjectivity	nonwandering	yes[2]	no
random to random	surjectivity	yes[3]	yes[3]
surjectivity	random to random	yes[3]	no[3]

1. Not satisfied for groups with a free subgroup on two generators.
2. Because of the *Poincaré recurrence theorem*.
3. For groups with decidable word problem.

Conclusions

- Moore's Garden of Eden theorem was the first rigorous result of cellular automata theory.
- It is a beautiful statement on its own.
- It opened the way to other insightful statements.
- It actually characterizes an important class of groups!
- What can be said about its converse by Myhill?
- What can be said about the other statements?

Thank you for attention!

Any questions?