

Emergent Universality from Symmetric Rules: Rule-110 Emulation by Alternating Trivial Cellular Automaton Rules

Emergent Universality from Symmetric Rules: Rule-110 Emulation by Alternating Trivial Cellular Automaton Rules

Abstract

Elementary cellular automaton (ECA) rule 110 is Turing complete, yet its local rule is chiral: it treats its left and right neighbours differently, and this handedness underlies the glider dynamics that carry its computation. A recent construction shows that rule 110 can nevertheless be emulated by three constituent ECA rules — rules 36, 37, and 5 — each of which is individually mirror-symmetric and individually produces only trivial frozen or oscillatory patterns; the chirality resides entirely in their fixed spatial arrangement. The open question this paper addresses is quantitative: under what noise and scale conditions does such an arrangement behave as a coherent universal system, and how robust is that behaviour? We formalize the construction, show that parity breaking arises from the ordering map $c(i) = i \bmod 3$ rather than from any constituent rule, and analyze domain walls between desynchronized regions with a two-parameter mean-field rate equation. With explicit arithmetic we derive the steady-state wall fraction $\rho^* = \sqrt{a\varepsilon/b}$, reproducing the $\varepsilon^{1/2}$ scaling reported for the construction; illustrative values are $\rho^* = 0.1$ at $\varepsilon = 10^{-2}$ and $\rho^* = 0.01$ at $\varepsilon = 10^{-4}$. We derive the relaxation time $\tau = 1/\sqrt{ab\varepsilon}$, a lock-on compatibility threshold $\rho_0 \geq 5 \times 10^{-4}$ walls per site for a 10^3 -step window, consistent with the reported observation that random initial conditions lock on to rule 110 over 90% of the time, and a noise-tolerance bound $\varepsilon_{\max} = b/(aN_c^2)$, equal to 10^{-4} for one tolerated defect per 100 sites. Because the emulation runs at a time factor $\lambda = 3$, the physical per-step noise budget tightens to $\varepsilon_{\max} \approx 3.3 \times 10^{-5}$. The analysis changes the design question for self-organizing computation from "how sophisticated must each unit be?" to "what ordering of trivial units suffices, and what noise budget does the ordering buy?".

1. Introduction

Engineers who want computation to emerge from large ensembles of simple, unreliable units — amorphous computing, synthetic biology gene networks, self-configuring materials — face a persistent design barrier. The classical route to universal computation in a spatially extended system is to install one sophisticated, asymmetric rule at every site: rule 110 in elementary cellular automata is the canonical example, and it is rigorously proven to be computationally universal. That route forces every unit to carry the full complexity of the rule table, and it offers no guidance on how much of that complexity can be offloaded onto the organization of the ensemble rather than the unit. Where the units are biological or physical components that cannot be individually programmed with high precision, the barrier is practical: the design either demands per-unit sophistication the substrate cannot deliver, or it abandons universality altogether. A demonstration that a repeating sequence of individually trivial, individually symmetric rules can emulate a Turing-complete chiral rule changes the design space: it shows that asymmetry can be supplied by arrangement, so the per-unit specification can be reduced to nearly nothing provided the schedule is fixed and correct. If the robustness analysis of this paper holds, a designer of such a substrate gains a concrete noise budget and a concrete initial-disorder threshold instead of a qualitative hope that the system "usually" self-organizes.

The concrete construction is due to [1, 2]. Three elementary rules — rule 36, rule 37, and rule 5 — are assigned to sites by position modulo 3, and three steps of the mixture equal one step of rule 110 under a block encoding. Each constituent is mirror-symmetric and, alone, evolves to trivial frozen or oscillatory patterns; the composite emulates the Turing-complete rule 110. Two statistical regularities are reported with it: random initial conditions lock on to rule 110 behaviour over 90% of the time, explained through the motion of domain walls between desynchronized regions, and under random noise of strength ε the fraction of the lattice occupied by noise-seeded domain walls scales as $\varepsilon^{1/2}$, explained through a simple rate equation [1, 2].

What is missing from the reported result, and what this paper supplies, is the quantitative skeleton underneath those claims. The reported scaling is an observation; the rate equation that explains it is stated but not evaluated; the lock-on fraction is an empirical number without a stated density threshold; and the noise budget that a would-be user of such a system could tolerate is nowhere computed. This paper derives each of these quantities with every arithmetic step shown, from a two-parameter mean-field model whose assumptions are stated explicitly, and it separates carefully between (i) results derived here from stated assumptions, (ii) illustrative numerical evaluations under labeled parameter choices, and (iii) empirical results reported in [1, 2] and cited as such. No new simulations are presented.

The conceptual stakes extend beyond cellular automata. Turing's machine model, in which a finite control manipulates a tape through a single access pointer, remains the operational definition of computability and the basis for the major complexity classes such as P and NP [3]; Turing's work separated software from hardware by isolating "local information" and "state of mind" as the two relevant ingredients of calculation [7]. The construction of [1, 2] instantiates a Turingian theme in a sharpened form: the "program" of the universal machine is redistributed from the local rule table into the spatial ordering of the rules, in the same spirit as the program-size versus time-complexity trade-offs studied for small universal Turing machines [8]. Whether such redistribution genuinely saves descriptive complexity is contested — the source work itself concludes that diversity does not save information in this construction [1] — and we examine that objection directly in Section 6.

The contributions are:

1. A formal statement that the composite update map of the three-rule mixture breaks parity through the ordering map $c(i) = i \bmod 3$ while every constituent satisfies $f(x, y, z) = f(z, y, x)$, with the constituent rules identified as rules 36, 37, and 5 and the emulation carrying time factor $\lambda = 3$ (Section 3).
2. A mean-field derivation, with explicit arithmetic, of the steady-state noise-seeded wall fraction $\rho^* = \sqrt{a\varepsilon/b}$, reproducing the $\varepsilon^{1/2}$ scaling reported in [1, 2], with illustrative values 0.1, 0.01, and 0.5 at $\varepsilon = 10^{-2}$, 10^{-4} , and 0.25 (Section 4.1).
3. The exact transient solution of the rate equation, giving relaxation time $\tau = 1/\sqrt{ab\varepsilon}$, and the noiseless coarsening law $\rho(t) = \rho_0/(1 + b\rho_0 t)$ (Section 4.2).
4. A lock-on compatibility threshold $\rho_0 \geq 5 \times 10^{-4}$ walls per site for a 10^3 -step window under ballistic wall motion, shown to be consistent with the >90% lock-on rate reported in [1, 2] (Section 4.3).
5. A noise-tolerance bound $\varepsilon_{\max} = b/(aN_c^2) = 10^{-4}$ for one tolerated defect per $N_c = 100$ sites, corrected for the emulation time factor $\lambda = 3$ to a physical per-step budget $\varepsilon_{\max} \approx 3.3 \times 10^{-5}$ (Section 4.4).

2. Background and Related Work

The source construction. The primary object of study is the demonstration in [1, 2] that the Turing-complete ECA rule 110 can be emulated by alternating three simple symmetric rules in space. The constituent rules are given explicitly there: rule 36 ("become active only if both neighbours are the opposite color of the cell"), rule 37 ("the same as rule 36, but also activate if all sites are inactive"), and rule 5 ("become active only if both outer neighbours

are inactive"); all three are mirror-symmetric and individually evolve to trivial frozen or oscillatory patterns [1, 2]. Rule 36 is assigned to sites with $i \bmod 3 = 0$, rule 37 to $i \bmod 3 = 1$, and rule 5 to $i \bmod 3 = 2$; logical 0 is encoded as the block [100] and logical 1 as the block [000], and three steps of the mixture equal one step of rule 110 [1, 2]. The same work reports that the mixture usually self-organizes into rule 110 from random initial conditions, with failures explained by domain walls between desynchronized regions, that lock-on occurs over 90% of the time, and that under noise the wall density scales as $\sqrt{\epsilon}$ via a rate equation [1, 2]. Its discussion adds two points this paper builds on: the chirality of the composite is "explicitly seen in the speeds of domain walls", and the hope that diversity saves information — several simple rules replacing one complex rule — is explicitly not realized in this construction [1]. Our analysis takes the reported regularities as empirical inputs and supplies the missing quantitative derivations.

Turing machines and complexity. The computational backdrop is the Turing machine, described in [3] as a survey of the Turing machine and its many influences on the theory of computation and complexity. The relevance here is precise: because rule 110 is Turing complete [1, 2], the three-rule mixture inherits the full class of Turing-computable functions, including their undecidability structure, so no general prediction procedure for the emulated dynamics can exist.

Turing's legacy and the origins of complexity. The anniversary assessment in [7] reviews and assesses Turing's contributions to the origins of complexity in the light of new developments. The construction studied here is a concrete instance of the software–hardware separation that Turing's framework enables: the "software" of rule 110 is carried partly by the spatial schedule rather than by the local hardware of each cell. The practical reading for the present system is that once lock-on is achieved, long-term prediction of the emulated dynamics is bounded by the computational power of rule 110's behaviour, not by the simplicity of the constituent rules.

Small machines and complexity trade-offs. The study of small Turing machines in [8] introduces computability and universality concepts and explores trade-offs between different measures of complexity, particularly algorithmic (program-size) and computational (time) complexity, as a means of explaining these measures in a novel manner [8]. The three-rule construction is a cellular-automaton analogue of exactly this trade: the description of the universal machine is redistributed from an 8-bit chiral rule table into three simpler rule tables plus a period-3 spatial schedule plus a block encoding, paying a time factor $\lambda = 3$. Whether the redistribution reduces total description length is examined adversarially in Section 6, in light of the source work's own conclusion that diversity does not save information here [1].

Variants of Turing computation. Infinite time Turing machines are obtained by extending the Turing machine concept to transfinite ordinal time, providing a natural model of infinitary computability and a setting for analyzing supertask computation [5]. The present construction sharpens the contrast: even the finite-time dynamics of a system whose local rules are all parity-symmetric is already universal, so infinitary extensions concern computational strength over time, not local rule complexity. Along a different axis, [6] compares Turing's framework with Kleene's higher-type schemes and poses the question whether there is an approach that marries the best of both frameworks, noting that both have their pros and cons [6]. That program concerns which higher-order objects are computable; the three-rule construction is orthogonal to it, showing that the substrate of ordinary Turing universality can be made locally homogeneous without changing the computable functions.

Turing, learning, and intelligence. The analysis of the Turing Test in [4] distinguishes a macro-level, post-hoc test of human intelligence from the micro-level, a priori definition of a Turing machine, and argues that an "out-of-the-box" Turing machine will not pass the Turing Test because intelligence involves repeated interactive learning [4]. The distinction matters here: the three-rule mixture is universal in Turing's micro-level sense, yet universality does

not entail adaptive behaviour. The lock-on phenomenon analyzed in Section 4.3 is a robustness property — the system recovers a coherent dynamical identity from perturbation — which is a precondition for, but by [4]'s argument not a route to, adaptive behaviour.

Scope of the related-work base. The quantitative analysis of this paper rests on the eight works [1]–[8]; the derivation sections use only [1, 2] as the source of the construction and its reported regularities, and the remaining works supply the computational-theory context in which the result is interpreted.

3. Methods

3.1 Elementary cellular automata and parity symmetry

An elementary cellular automaton (ECA) consists of a bi-infinite binary lattice $a_i(t) \in \{0, 1\}$, $i \in \mathbb{Z}$, updated synchronously by a radius-1 rule:

$$a_i(t+1) = f(a_{i-1}(t), a_i(t), a_{i+1}(t)),$$

with $f : \{0, 1\}^3 \rightarrow \{0, 1\}$ encoded as an 8-bit rule number $R = \sum_{k=0}^7 f_k 2^k$. Rule 110 has $R_{110} = 110 = 2^6 + 2^5 + 2^3 + 2^1$, i.e., binary 01101110. Explicitly, $f_{110}(1, 1, 0) = 1$ while $f_{110}(0, 1, 1) = 0$: the rule distinguishes the left neighbour from the right neighbour. This chirality is essential to the glider dynamics that carry information in rule 110's universal computation [1, 2].

A rule f is *parity-symmetric* (mirror-symmetric) if

$$f(x, y, z) = f(z, y, x) \quad \text{for all } (x, y, z) \in \{0, 1\}^3,$$

equivalently, its rule number is invariant under the mirror involution on the 8-bit table. The three constituent rules of [1, 2] each satisfy this condition, so no local update carries a preferred direction. Their verbal definitions, as given in [1, 2], are:

- Rule 36: become active only if both neighbours are the opposite color of the cell.
- Rule 37: the same as rule 36, but also activate if all sites are inactive.
- Rule 5: become active only if both outer neighbours are inactive.

Each is individually trivial in the sense of [1, 2]: alone, each eventually results in frozen or oscillatory patterns.

3.2 The alternating construction

Following [1, 2], the three symmetric rules are applied in a fixed repeating *spatial* sequence: site i is assigned rule 36 if $i \bmod 3 = 0$, rule 37 if $i \bmod 3 = 1$, and rule 5 if $i \bmod 3 = 2$. Writing $f_1 = f_{36}$, $f_2 = f_{37}$, $f_3 = f_5$, and $c(i) = i \bmod 3$, every site updates simultaneously with its assigned rule:

$$a_i(t+1) = f_{c(i)+1}(a_{i-1}(t), a_i(t), a_{i+1}(t)).$$

The composite update map T is not parity-symmetric. For the reflection operator σ defined by $(\sigma a)_i = a_{-i}$, each constituent satisfies $f_j(\sigma a) = \sigma f_j(a)$, but the class pattern is not reflection-invariant: defining $c^-(i) = c(-i)$, we have $c(i) \neq c^-(i)$ for $i \not\equiv 0 \pmod{3}$. Concretely, at $i = 1$: $c(1) = 1$, so the forward update at site 1 uses f_2 (rule 37), while $c^-(1) = 2$, so the reflected update uses f_3 (rule 5). Hence

$$T(\sigma a)_i = f_{c(i)+1}((\sigma a)_{i-1}, (\sigma a)_i, (\sigma a)_{i+1}) \neq f_{c^-(i)+1}((\sigma a)_{i-1}, (\sigma a)_i, (\sigma a)_{i+1}) = (\sigma T a)_i$$

generically. The composite is chiral even though no constituent is: the chirality is carried by the ordering map $c(i)$, an organizational property. This is the precise content of the statement in [1, 2] that "the chirality of rule 110 is created by their arrangement", and it is why the source work observes the chirality "explicitly" in the speeds of domain walls [1]: walls separating regions that run the three-rule cycle in opposite spatial phase inherit the handedness of the schedule.

3.3 Emulation of rule 110

Following [1, 2], the composite dynamics T emulates rule 110 in the block-encoding sense: logical 0 is encoded as the block [100] and logical 1 as the block [000], and three steps of the mixture equal one step of rule 110 [1, 2]. Formally, there is a coarse-graining map Φ from configurations of the three-rule system to configurations of rule 110 and a time-scale factor $\lambda = 3$ such that

$$\Phi(T^{\lambda n}(A)) = F_{110}^n(\Phi(A))$$

for initial conditions A in the emulating subset, where F_{110} is the global map of rule 110. The period-3 spatial structure of T reproduces the asymmetric neighbourhood sampling that rule 110 performs with its chiral rule table. Since rule 110 is Turing complete [1, 2], the three-rule system is Turing complete as well. We note that rule 36 can be replaced by rule 54, as the two rules only differ in positions that are never visited [1, 2]; nothing in the analysis below depends on which of the two is used.

3.4 Domain walls, lock-on, and noise

The emulating dynamics does not enforce global phase coherence: different regions of the lattice may run the three-rule cycle in different relative phases (desynchronized regions). Boundaries between phase domains behave as *domain walls*: localized defects that move through the lattice and annihilate on collision, as described in [1, 2]. Lock-on to rule 110 behaviour is the absorption of all phase boundaries, after which the lattice evolves as a coherent rule-110-equivalent system.

We model walls as point particles on a line with density $\rho(t)$ (walls per site), moving with characteristic speed v (sites per step) and annihilating on contact. Noise of strength ε (probability per site per step of a spontaneous bit flip, applied after the deterministic update) creates new walls at rate $a\varepsilon$ per site per step, where a is a susceptibility constant depending on how often a flip actually destabilizes the local phase. The mean-field rate equation, of the same form used in [1, 2] to explain the $\sqrt{\varepsilon}$ scaling, is

$$\frac{d\rho}{dt} = a\varepsilon - b\rho^2,$$

where the annihilation term $b\rho^2$ encodes pairwise wall collisions (each collision removes two walls, giving the quadratic loss term standard in coarsening theory) and $b > 0$ is the collision-rate constant.

4. Analysis

4.1 Steady-state wall fraction and the $\varepsilon^{1/2}$ law

Set $d\rho/dt = 0$ in the rate equation:

$$0 = a\varepsilon - b\rho^2 \implies \rho^2 = \frac{a\varepsilon}{b} \implies \rho^* = \sqrt{\frac{a\varepsilon}{b}} = \sqrt{\frac{a}{b}} \varepsilon^{1/2}.$$

This reproduces the square-root scaling of the noise-seeded wall density reported in [1, 2]. The exponent $1/2$ follows solely from the structure *creation linear in noise, annihilation quadratic in density*; it is independent of the constants a and b . The prefactor $\sqrt{a/b}$ is a model normalization, not a measured quantity.

Numerical illustration (labeled illustrative parameters). Take $a = 1$ (every noise event seeds a wall, the maximal-susceptibility upper bound) and $b = 1$ (unit collision rate), in units where ρ is per site and t per step. Inputs: $a = 1, b = 1$ (normalization choices, Section 3.4); ε values are parameters of the thought experiment.

- For $\varepsilon = 10^{-2}$: $\rho^* = \sqrt{1 \cdot 10^{-2}/1} = \sqrt{10^{-2}} = 10^{-1} = 0.1$. That is, 10% of sites host a wall.
- For $\varepsilon = 10^{-4}$: $\rho^* = \sqrt{10^{-4}} = 10^{-2} = 0.01$, i.e., 1%.
- For $\varepsilon = 0.25$: $\rho^* = \sqrt{0.25} = 0.5$, i.e., 50%.

Check of the scaling: reducing ε by a factor of 10^2 (from 10^{-2} to 10^{-4}) reduces ρ^* by a factor of 10^1 (from 0.1 to 0.01), consistent with $\rho^* \propto \varepsilon^{1/2}$ since $(10^2)^{1/2} = 10$.

4.2 Transient solution and relaxation time

The rate equation is separable. With initial condition $\rho(0) = \rho_0$:

$$\int \frac{d\rho}{a\varepsilon - b\rho^2} = t \quad \Longrightarrow \quad \frac{1}{\sqrt{ab\varepsilon}} \operatorname{artanh}\left(\rho\sqrt{\frac{b}{a\varepsilon}}\right) = t + C,$$

so that

$$\rho(t) = \sqrt{\frac{a\varepsilon}{b}} \tanh\left(\sqrt{ab\varepsilon} t + \operatorname{artanh}\left(\rho_0\sqrt{\frac{b}{a\varepsilon}}\right)\right).$$

The approach to steady state occurs on the timescale

$$\tau = \frac{1}{\sqrt{ab\varepsilon}}.$$

With $a = b = 1$ and $\varepsilon = 10^{-2}$: $\tau = 1/\sqrt{10^{-2}} = 1/10^{-1} = 10$ steps. For $\varepsilon = 10^{-4}$: $\tau = 1/\sqrt{10^{-4}} = 1/10^{-2} = 100$ steps. Weaker noise therefore both lowers the steady wall fraction and lengthens the relaxation time. (A linearization around ρ^* gives $\tau = 1/(2\sqrt{ab\varepsilon})$, a factor of 2 smaller; that value governs only the late-stage exponential tail, and we adopt the exact-solution value as the convention for the full approach.)

4.3 Noiseless coarsening and lock-on from random initial conditions

Without noise ($\varepsilon = 0$), the rate equation gives pure annihilation, $d\rho/dt = -b\rho^2$. Derivation: separating, $d\rho/\rho^2 = -b dt$; integrating, $-1/\rho = -bt + C$; with $\rho(0) = \rho_0$, $C = -1/\rho_0$; hence $1/\rho = bt + 1/\rho_0$, giving

$$\rho(t) = \frac{\rho_0}{1 + b\rho_0 t}.$$

Walls coarsen algebraically, not exponentially.

Lock-on criterion (projection). Take the reported lock-on fraction of over 90% of random initial conditions [1, 2] and ask what initial wall density ρ_0 is compatible with full coarsening within a window of $T = 1000$ steps on a lattice of length $L = 10^3$ sites, assuming walls move at speed $v = 1$ site per step and annihilate in pairs. Full

lock-on requires every wall to meet a partner: the mean distance between walls is $\ell = 1/\rho_0$, and two approaching walls close the gap at relative speed $2v$, so the mean annihilation time is

$$t_{\text{ann}} = \frac{\ell}{2v} = \frac{1}{2v\rho_0}.$$

Requiring $t_{\text{ann}} \leq T$ with $v = 1, T = 1000$:

$$\frac{1}{2\rho_0} \leq 1000 \implies \rho_0 \geq \frac{1}{2000} = 5 \times 10^{-4}.$$

Lock-on within 10^3 steps is therefore expected whenever the initial wall density exceeds 5×10^{-4} walls per site. As an order-of-magnitude assumption (explicitly labeled as such, not a measurement): random binary initial conditions on $L = 10^3$ sites typically contain of order 10^2 or more phase mismatches between adjacent regions, i.e., ρ_0 of order 10^{-1} walls per site, which exceeds the threshold by two orders of magnitude. This is a compatibility argument, not a derivation of the reported 90% figure: the threshold explains why typical random configurations coarsen successfully, while the exact lock-on probability depends on wall pinning and rare phase-locked configurations that the mean-field model does not resolve.

4.4 Noise threshold for sustained computation

If a universal computation running on the emulating lattice tolerates at most one defect per N_c sites, the steady-state condition $\rho^* \leq 1/N_c$ gives

$$\sqrt{\frac{a\varepsilon}{b}} \leq \frac{1}{N_c} \implies \varepsilon_{\text{max}} = \frac{b}{aN_c^2}.$$

With $a = b = 1$ and $N_c = 100$:

$$\varepsilon_{\text{max}} = \frac{1}{100^2} = \frac{1}{10^4} = 10^{-4}.$$

Correction for the emulation time factor. The mixture produces one rule-110 step per $\lambda = 3$ of its own steps (Section 3.3). A noise event applied per mixture step therefore occurs λ times per emulated step, so the noise budget per mixture step must be divided by λ to obtain the budget per emulated rule-110 step:

$$\varepsilon_{\text{max}}^{(\text{emulated step})} = \frac{b}{aN_c^2\lambda} = \frac{1}{10^4 \times 3} = \frac{1}{3 \times 10^4} \approx 3.3 \times 10^{-5}.$$

The physical per-step noise tolerance of the three-rule substrate is thus a factor of three tighter than the naive bound 10^{-4} , a direct cost of paying the time factor $\lambda = 3$.

5. Results

-
- **R1 (scaling law, derived).** $\rho^* = \sqrt{a\varepsilon/b}$, reproducing the $\varepsilon^{1/2}$ scaling reported in [1, 2] (Section 4.1).
 - **R2 (illustrative steady states).** With $a = b = 1$: $\rho^* = 0.1$ at $\varepsilon = 10^{-2}$; $\rho^* = 0.01$ at $\varepsilon = 10^{-4}$; $\rho^* = 0.5$ at $\varepsilon = 0.25$ (Section 4.1).
 - **R3 (relaxation times, illustrative).** $\tau = 1/\sqrt{ab\varepsilon}$ gives $\tau = 10$ steps at $\varepsilon = 10^{-2}$ and $\tau = 100$ steps at $\varepsilon = 10^{-4}$ (Section 4.2).
 - **R4 (transient law, derived).** $\rho(t) = \rho_0/(1 + b\rho_0 t)$ for $\varepsilon = 0$, showing algebraic coarsening (Section 4.3).

- **R5 (lock-on compatibility, projection).** The threshold $\rho_0 \geq 5 \times 10^{-4}$ walls per site ensures lock-on within 10^3 steps under the stated ballistic assumption ($v = 1$, $T = 1000$); typical random initial conditions, at an assumed ρ_0 of order 10^{-1} , exceed this threshold by two orders of magnitude, consistent with the >90% lock-on rate reported in [1, 2] (Section 4.3).
- **R6 (noise tolerance, derived).** $\varepsilon_{\max} = b/(aN_c^2) = 10^{-4}$ for $N_c = 100$; corrected for the emulation time factor $\lambda = 3$, the per-mixture-step budget is $\varepsilon_{\max} \approx 3.3 \times 10^{-5}$ (Section 4.4).
- **R7 (empirical lock-on, input).** Random initial conditions lock on to rule 110 behaviour over 90% of the time [1, 2]; this is an empirical input used for compatibility analysis, not a prediction of this paper.
- **R8 (universality, inherited).** Because rule 110 is Turing complete and the three-rule system emulates it with $\lambda = 3$, the three-rule system is also Turing complete, inheriting the undecidability structure of Turing computation [1, 2, 3].

6. Discussion

Mean-field limitations. The rate equation assumes spatially homogeneous mixing of walls, which in one dimension can be violated by fluctuations that produce larger empty intervals; coarsening in one dimension is known to admit logarithmic corrections to mean-field exponents in related particle systems. Such corrections could modify the exponent or introduce logarithmic factors in $\rho^*(\varepsilon)$. The reported $\varepsilon^{1/2}$ scaling in [1, 2] suggests the mean-field picture captures the dominant behaviour, but the prefactor $\sqrt{a/b}$ has not been measured, and a and b are normalization choices here, not fitted constants.

Lock-on figure. The >90% lock-on statistic is taken directly from [1, 2]. Our wall-density analysis explains why typical random configurations exceed the required ρ_0 threshold, but does not predict the exact 90% value; deviations could arise from pinned walls, walls that bind to gliders and are carried without annihilation, or rare phase-locked configurations in which two domains coexist indefinitely.

Falsifiability. The $\varepsilon^{1/2}$ law would be falsified if direct simulation of the three-rule system under imposed noise showed a systematically different exponent (for example, linear scaling $\rho^* \propto \varepsilon$ or an exponent with logarithmic corrections). The claim that chirality originates solely from ordering would be falsified if reversing the spatial order of the three rules produced identical rule-110 dynamics, or if a temporally alternating composition of the same three rules (one rule per time step, uniformly in space) also emulated rule 110 — either outcome would show that the asymmetry is not carried by the spatial schedule alone.

Failure modes. (i) Walls might bind to gliders, preventing annihilation and invalidating the pairwise-collision term. (ii) Noise could generate persistent localized structures that do not behave as point walls, breaking the single-density description. (iii) The emulation might hold only on a restricted set of initial conditions, so that the coarse-graining map Φ is not defined on the configurations actually reached after noise events. Each failure mode would weaken the quantitative predictions derived here, though not the qualitative emulation result of [1, 2].

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. The direct users at present are researchers designing substrates for emergent or self-organizing computation: amorphous-computing groups, synthetic-biology circuit designers, and designers of cellular-automaton-based hardware or unconventional computing fabrics. For them, the result changes one design decision: the per-unit specification. Instead of engineering a single sophisticated, asymmetric update rule into every unit, the designer can distribute three trivial, mirror-symmetric rules across the substrate in a fixed period-3 spatial schedule and obtain rule-110-equivalent universal dynamics, provided two quantitative conditions derived here hold: the initial phase disorder stays above the lock-on threshold $\rho_0 \geq 5 \times 10^{-4}$ walls per site

(Section 4.3) and the per-step noise stays below $\varepsilon_{\max} = b/(aN_c^2)$, equal to 10^{-4} for one tolerated defect per $N_c = 100$ sites, or $\approx 3.3 \times 10^{-5}$ per mixture step after correcting for the time factor $\lambda = 3$ (Section 4.4). A second class of users is theoretical: anyone studying whether descriptive complexity can be reduced by organization will use the negative finding of [1] — that diversity does not save information in this construction — together with the quantitative noise and disorder budgets derived here as the terms of that trade.

(b) Current readiness. This work is purely analytical: it formalizes a published construction, derives consequences of a stated rate equation, and presents no simulations or hardware. Against the NASA Technology Readiness Level definitions, the construction itself has been demonstrated in the source work's own cellular-automaton experiments [1, 2], but the quantitative robustness analysis of this paper is analytical and has not been peer reviewed, which places this paper's contributions at TRL 1–2; no analytical result of this kind can exceed TRL 3. No manufacturing readiness is claimed or claimable for a result of this type.

(c) The smallest next test. The minimal de-risking calculation is a direct simulation of the three-rule mixture on a lattice of $L = 10^3$ sites with imposed per-site-per-step noise at several strengths ε , measuring the steady-state wall density $\rho^*(\varepsilon)$ and fitting the exponent. The derived law predicts $\rho^* = \sqrt{a\varepsilon/b}$, i.e., a slope of $1/2$ on a log-log plot of ρ^* against ε . An outcome with a systematically different exponent — for example $\rho^* \propto \varepsilon$ or $\rho^* \propto \varepsilon^{1/2}$ multiplied by a strong logarithm — refutes the mean-field derivation of Section 4.1 as a description of this system. A second cheap test measures the lock-on fraction from random initial conditions against the predicted threshold: configurations seeded at ρ_0 below 5×10^{-4} walls per site should fail to lock on within 10^3 steps under the ballistic assumption $v = 1$; lock-on at densities well below the threshold would falsify the annihilation-time estimate of Section 4.3.

(d) Binding constraints. The constraints that would bind first, in the order the analysis fixes them: (i) the noise budget, which is the tightest derived number — $\varepsilon_{\max} \approx 3.3 \times 10^{-5}$ per mixture step for $N_c = 100$, and it tightens quadratically as the tolerated defect spacing N_c shrinks, since $\varepsilon_{\max} \propto N_c^{-2}$; (ii) the time factor $\lambda = 3$, which triples the wall-clock cost of every emulated rule-110 step and triples the exposure to noise per logical step; (iii) the schedule fidelity — the spatial assignment $c(i) = i \bmod 3$ must be exact everywhere, since a single mis-assignment site is precisely a domain wall of the kind the lock-on analysis counts; and (iv) the constants a , b , and v , which are normalization choices here, not measured values, and which scale every numerical budget linearly or as a square root. Whether a physical or biological substrate can hold a fixed period-3 spatial schedule with the required fidelity, and what defect rate such a substrate actually exhibits, are open questions the available sources do not quantify.

(e) Order-of-magnitude performance bounds (projections). Under the stated assumptions $a = b = 1$, $v = 1$ site per step, and the emulation time factor $\lambda = 3$: a computation of n rule-110 steps requires $3n$ mixture steps; the wall population relaxes on the timescale $\tau = 1/\sqrt{ab\varepsilon}$, so at $\varepsilon = 10^{-4}$ the system recovers from a perturbation in roughly 100 mixture steps, i.e., about 33 emulated rule-110 steps; and sustained computation tolerates of order one defect per $N_c = 100$ sites at noise $\varepsilon = 10^{-4}$ per mixture step. These are projections from the mean-field model with its stated normalization assumptions, not measurements of any physical substrate.

(f) Risks and misuse paths. No dual-use risk is apparent: the construction is a mathematical equivalence between cellular automaton rules, and the derived quantities are robustness budgets for idealized lattice dynamics. The main misuse is interpretive rather than practical — citing the construction as evidence that "diversity alone generates complexity" without the qualification, stated in [1] itself, that the arrangement is a structured object and that the source work concludes diversity does not save information here. Overstating the lock-on figure as a guarantee would be a second misuse: the >90% statistic is an empirical report of [1, 2], and the threshold of Section 4.3 is a compatibility argument, not a probability prediction.

8. Conclusion

This paper has supplied the quantitative skeleton under the construction of [1, 2], in which the Turing-complete, chiral elementary cellular automaton rule 110 is emulated by three individually mirror-symmetric and individually trivial rules — rules 36, 37, and 5 — arranged in a fixed period-3 spatial schedule. The derivations establish: (i) that parity breaking resides entirely in the ordering map $c(i) = i \bmod 3$, since every constituent satisfies $f(x, y, z) = f(z, y, x)$ (Section 3.2); (ii) the mean-field steady-state law $\rho^* = \sqrt{a\varepsilon/b}$, reproducing the $\varepsilon^{1/2}$ scaling reported in [1, 2], with illustrative values $\rho^* = 0.1$ at $\varepsilon = 10^{-2}$ and $\rho^* = 0.01$ at $\varepsilon = 10^{-4}$ (Section 4.1); (iii) the exact transient solution and relaxation time $\tau = 1/\sqrt{ab\varepsilon}$ (Section 4.2); (iv) a lock-on compatibility threshold $\rho_0 \geq 5 \times 10^{-4}$ walls per site for a 10^3 -step window, consistent with the >90% lock-on rate reported in [1, 2] (Section 4.3); and (v) a noise-tolerance bound $\varepsilon_{\max} = b/(aN_c^2) = 10^{-4}$ for $N_c = 100$, tightened to $\approx 3.3 \times 10^{-5}$ per mixture step by the emulation time factor $\lambda = 3$ (Section 4.4).

The broader thesis survives in a qualified form. Universal computation can be carried by units that are individually trivial and symmetric, with the asymmetry supplied by arrangement; but the arrangement is itself a structured object, and the source work's own conclusion that diversity does not save information in this construction [1] stands: the construction redistributes descriptive complexity from the rule table into the schedule and the block encoding, paying a time factor $\lambda = 3$ and a noise budget that tightens as N_c^{-2} . What the analysis adds is that this trade has computable terms. Future work should measure the constants a , b , and v by direct simulation, verify the exponent of the $\varepsilon^{1/2}$ law against fluctuation-corrected one-dimensional coarsening, and test the falsification conditions stated in Section 6 — in particular, whether a temporally alternating composition of the same three rules also emulates rule 110, which would overturn the claim that the spatial schedule alone carries the chirality.

Appendix A. Divergence report

D1 (spatial versus temporal composition). Drafts A and B described the construction as a fixed repeating *spatial* sequence of rules (site i uses rule $f_{c(i)+1}$ with $c(i) = i \bmod 3$); draft C described it as a *temporal* composition $F_3 \circ F_2 \circ F_1$ of global update operators. The convention behind the disagreement: the source work states the rules alternate in space, so the main text adopts the spatial convention; the temporal framing is retained only as a remark in Section 3.2.

D2 (relaxation-time convention). Draft A computed the relaxation time from the exact solution of the rate equation, $\tau = 1/\sqrt{ab\varepsilon}$; draft C linearized around ρ^* and obtained $\tau = 1/(2\sqrt{ab\varepsilon})$, a factor of 2 smaller. The disagreement is a convention: the exact-solution timescale governs the full approach from any ρ_0 , while the linearized value governs only the late-stage exponential tail. The main text adopts the exact-solution value and notes the factor of 2 in Section 4.2.

D3 (unverifiable measured data). Draft C included a table of "measured" wall densities and lock-on convergence probabilities attributed to simulation. These numbers could not be verified against [1, 2] or reproduced from stated inputs, and were rejected during reconciliation; the main text uses only derived values, clearly labeled illustrative parameters, and results explicitly reported in [1, 2].

D4 (identity of the constituent rules). Draft B treated the constituent rules as abstract symmetric rules f_1, f_2, f_3 , stating their identities were unspecified; drafts A and C named them as rules 36, 37, and 5 with their verbal definitions. The source work gives the rules explicitly, so the main text adopts the named convention.

Appendix B. Claim attribution

Claim	Substance	Drafts	Status
C1	Rule 110 is Turing complete and is emulated by three individually parity-symmetric ECA rules [1, 2]	A, B, C	CONVERGENT
C2	Chirality of the composite arises from the ordering map $c(i) = i \bmod 3$, not from any constituent rule	A, B, C	CONVERGENT
C3	Random initial conditions lock on to rule 110 behaviour over 90% of the time [1, 2]	A, B, C	CONVERGENT
C4	Noise-induced wall fraction scales as $\varepsilon^{1/2}$ [1, 2]	A, B, C	CONVERGENT
C5	Rules alternate in space (A, B) versus in time (C)	A, B, C	DIVERGENT (D1)
C6	Relaxation time $\tau = 1/\sqrt{ab\varepsilon}$ (A) versus $\tau = 1/(2\sqrt{ab\varepsilon})$ (C)	A, C	DIVERGENT (D2)
C7	Tabulated measured wall densities and lock-on probabilities	C	SINGLE (rejected, D3)
C8	Lock-on compatibility threshold $\rho_0 \geq 5 \times 10^{-4}$ walls per site for a 10^3 -step window	A	SINGLE (retained as projection)
C9	Noise-tolerance bound $\varepsilon_{\max} = b/(aN_c^2)$	A, B	CONVERGENT
C10	Rate-equation derivation $\rho^* = \sqrt{a\varepsilon/b}$ with illustrative parameters $a = b = 1$	A, B	CONVERGENT
C11	Constituent rules named as rules 36, 37, and 5 with verbal definitions (A, C) versus left abstract (B)	A, B, C	DIVERGENT (D4)
C12	Emulation time factor $\lambda = 3$ and corrected noise budget $\approx 3.3 \times 10^{-5}$	A	SINGLE (retained, derived in Section 4.4)

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.18.1-parallel-8). Language models were tools in that process; they are not authors. The named author is accountable for the content.
- Models: drafting by openai/gpt-oss-120b, zai-org/glm-5.3-flash; reconciliation, completion, review, revision and verification by zai-org/glm-5.3-flash, zai-org/glm-5.3, openai/gpt-oss-120b, with the QNFO model router as a fallback.
- Sources: every reference is drawn from a bibliography retrieved from arXiv and the QNFO corpus before drafting; every cited work's full text was retrieved (arXiv HTML, or the QNFO paper's body) and the writers were given excerpts of its own sections; statements about a cited work are restricted to its retrieved title, abstract and those excerpts.
- Checks passed before publication: required sections present and finished, cross-references resolve, reference list rendered from the retrieved bibliography, no assistant residue or placeholders, adversarial model review with its blocking findings resolved, and a citation existence check at publication.
- Independent recomputation: a separately generated script recomputed 8 quantitative claims from the paper's stated inputs; 0 differed from the stated value at the stated precision.
- Not checked: whether each cited source supports the sentence that cites it beyond the retrieved full-text excerpts; derivation steps other than the recomputed claims. This paper has not been peer reviewed, and this version was published by the pipeline without a prior manual review by the author.

References

[1] TITLE: arXiv Query: `searchquery=&idlist=2610.09879&start=0&max_results=1` [2] A repeating sequence of simple rules creates a universal computer. arXiv:2610.09879v1. <https://arxiv.org/abs/2610.09879v1> [3] Turing Machines and Understanding Computational Complexity. arXiv:1201.1223v1. <https://arxiv.org/abs/1201.1223v1> [4] Learning, Social Intelligence and the Turing Test - why an "out-of-the-box" Turing Machine will not pass the Turing Test. arXiv:1203.3376v1. <https://arxiv.org/abs/1203.3376v1> [5] Infinite Time Turing Machines: Supertask Computation. arXiv:math/0212047v1. <https://arxiv.org/abs/math/0212047v1> [6] Between Turing and Kleene. arXiv:2111.05052v1. <https://arxiv.org/abs/2111.05052v1> [7] Alan Turing and the Origins of Complexity. arXiv:1110.0271v1. <https://arxiv.org/abs/1110.0271v1> [8] Complejidad descriptiva y computacional en maquinas de Turing pequenas. arXiv:1010.1328v2. <https://arxiv.org/abs/1010.1328v2>