ADBA Chaos and Fractals
Cutting-Edge Scientific Solutions

e-ISSN: 3023-8951
RESEARCH ARTICLE
Vol.3 / No.2 / 2026 / pp.108-116
https://doi.org/10.69882/adba.chf.2026074

Diffusion, Not Chaotic Forcing, Controls Mixing in a Boundary-Coupled Cellular Automaton

Mustafa Kutlu [*,1] and **Mehmet Kutlu** [α,2]
[*]Department of Systems Engineering, Military Technological College, Muscat, Oman, [α]Clinic for General and Gerontopsychiatry, Oberberg Hospital, Gummersbach, Germany.

ABSTRACT A recurring intuition is that chaotic boundary forcing mixes a medium more thoroughly than regular or random forcing. We test this in a boundary-driven probabilistic cellular automaton, quantifying mixing via Shannon entropy and Moran's I, and find the intuition fails. Experiments demonstrate that neither coupling strength nor forcing character (chaotic Lorenz, white-noise, periodic, or static) significantly affects the final entropy. Instead, a gradient-boosted surrogate ($R^2 = 0.81$) and SHAP attribution identify the internal diffusive update weight as the dominant mixing driver. Exact spectral analysis confirms this mechanism: boundary drives perturb bulk functionals only at order $1/N$, whereas the diffusion weight heavily controls modal damping. While steering this diffusion weight recovers ninety-five per cent of mixing performance, an advective benchmark shows true chaotic velocity fields mix six times faster. This localizes our null result to additive boundary forcing rather than chaotic transport itself. Ultimately, bulk mixing in this system is governed by internal diffusive transport rather than the boundary drive.

KEYWORDS
Chaotic mixing
Cellular automata
Surrogate modelling
Explainable AI

INTRODUCTION

The idea that chaos promotes mixing is one of the most productive in nonlinear science. Since the recognition that even a laminar, deterministic velocity field can stretch and fold material lines exponentially, chaotic advection has provided the theoretical backbone for understanding how fluids and granular media homogenise when stirring is slow and molecular diffusion is weak (Aref 1984; Ottino 1989, 1990; Aref 2002). The appeal of the mechanism lies in its economy: a simple, low-dimensional forcing, applied at a boundary or through a time-dependent perturbation, can in principle drive efficient transport throughout a domain.

It is therefore natural to expect that a chaotic driver, with its broadband spectrum and sensitive dependence on initial conditions, will outperform a periodic or a purely random one when the objective is to mix. This expectation is widely held, and it motivates the recurring use of chaotic signals as actuation in mixing devices and models. A distinction must be drawn at the outset, however, between chaotic *advection*, in which a chaotic velocity field transports material throughout the bulk, and the *additive boundary forcing* studied here, in which a chaotic signal is injected into a single boundary row with no accompanying velocity field. The former is the mechanism of the classical theory; the latter is the weaker and more common engineering surrogate, and whether it inherits any of the mixing virtue of the former is the question we pose. We are careful, therefore, not to conflate a null result for boundary forcing with a claim about chaotic advection, and we include an advective benchmark below precisely to keep the two apart.

Whether that expectation survives contact with a discrete, dissipative medium is less clear. Cellular automata offer a transparent laboratory for the question, because their local update rules make the competing transport channels, self-retention, neighbour averaging, diffusion and external drive, explicit and separately tunable (Wolfram 1983; Chopard and Droz 1998). In such a medium the boundary forcing is only one of several mechanisms redistributing concentration, and there is no guarantee that it dominates the others. A driver that injects structure at a single boundary must compete with the automaton's own relaxation, and if internal transport is efficient the details of the drive may be washed out before they influence the bulk. Distinguishing these possibilities requires more than a single simulation; it requires that the strength of the coupling and the nature of the forcing be varied systematically and that the response be measured with well-defined, reproducible metrics.

Two such metrics recur across the literature on spatial homogenisation. The Shannon entropy of the concentration histogram measures how uniformly mass is distributed over its accessible range and rises as the field mixes (Shannon 1948), while

Manuscript received: 3 February 2026,
Revised: 5 June 2026,
Accepted: 7 July 2026.

[1]mustafa.kutlu@mtc.edu.om (**Corresponding author**)
[2]mehmet.kutlu@klinikum-oberberg.de

Moran's I measures spatial autocorrelation and falls toward zero as neighbouring regions decorrelate (Moran 1950). Used together they separate the amount of mixing from its spatial texture, and their combination provides a compact, quantitative readout of state that is well suited to automated experimentation across many parameter settings.

Establishing which parameters matter is, however, only half of an explanation. Machine-learning surrogates now make it feasible to learn the map from a model's parameters to its emergent behaviour and to interrogate that map for the variables that carry the most weight (Breiman 2001; Friedman 2001; Chen and Guestrin 2016). Attribution methods based on Shapley values give a principled, additive decomposition of a prediction across its inputs (Lundberg and Lee 2017; Lundberg *et al.* 2020), and when their rankings agree with an independent measure such as permutation importance (Altmann *et al.* 2010) the attribution can be trusted as a statement about the surrogate rather than an artefact of the method. A further and more demanding test is whether an identified driver can be exploited: if a variable genuinely controls the outcome, then steering it, for instance within a model-predictive loop (Mayne *et al.* 2000; Camacho and Bordons 2007), should improve that outcome. Explanation and control are distinct claims, and a rigorous study should keep them apart (Ribeiro *et al.* 2016; Doshi-Velez and Kim 2017).

This paper brings these threads together on a single, deliberately simple system. We couple a two-dimensional probabilistic cellular automaton to a chaotic Lorenz driver through one boundary row, with an adjustable coupling amplitude, and we ask what actually determines how well the automaton mixes. We first test the chaos-promotes-mixing intuition directly, by sweeping the coupling strength and by ablating the forcing against noise, periodic and static alternatives. We then build a gradient-boosted surrogate over the automaton's internal update weights, attribute its predictions with SHAP, and validate the attribution against permutation importance. Finally, we ask whether the identified driver confers control leverage by embedding the SHAP weights in a model-predictive controller and comparing it against an unguided baseline. Underpinning the empirical study is an exact spectral analysis of the update operator, which because it is translation-invariant on the torus is diagonalized in closed form; the resulting symbols predict each experimental outcome quantitatively and reduce the null results to a surface-to-volume argument. The analysis returns a clear and somewhat deflationary answer: within this class of boundary-driven automaton, mixing is governed by internal diffusive transport and is essentially indifferent both to the strength of the boundary coupling and to whether that boundary is driven chaotically, randomly, periodically or not at all.

MODEL AND METHODS

Chaotic boundary driver

The forcing signal is drawn from the Lorenz system,

$$\dot{x} = \sigma\,(y - x), \tag{1}$$
$$\dot{y} = x\,(\rho - z) - y, \tag{2}$$
$$\dot{z} = x\,y - \beta\,z, \tag{3}$$

with the classical parameters $\sigma = 10$, $\rho = 28$ and $\beta = 8/3$, integrated by an adaptive Runge–Kutta scheme from the initial condition $(0.1, 0, 0)$. To confirm that the driver operates in the chaotic regime we estimated its largest Lyapunov exponent by the standard Benettin two-trajectory procedure with renormalisation, obtaining $\lambda_1 = 0.93$ nats s^{-1} (1.34 bit s^{-1}), in agreement with the canonical Lorenz value and certifying the exponential separation of nearby states that defines deterministic chaos (Lorenz 1963; Strogatz 2015). The $x(t)$ component, rescaled to the interval $[-1, 1]$, supplies the boundary drive fed to the automaton (Figure 1).

Boundary-coupled cellular automaton

The medium is a square 40×40 lattice carrying a scalar concentration $c_{ij} \in [0, 1]$, initialised with the upper half at unity and the lower half at zero to create a sharp, maximally unmixed interface. At each step the field is updated by a convex combination of four local channels,

$$c' = w_0\,c + w_1\,\mathcal{A}c + w_2\,\mathcal{L}c + w_3\,\mathcal{D} + \eta, \tag{4}$$

where $\mathcal{A}c$ is the four-neighbour average, $\mathcal{L}c$ is the discrete Laplacian, and $\mathcal{D}$ injects the boundary drive into the top row with amplitude α, so that $\mathcal{D}_{0j} = \alpha\,\tilde{x}$ with $\tilde{x}$ the rescaled Lorenz sample and $\mathcal{D}_{ij} = 0$ elsewhere. The term η is a small Gaussian noise of standard deviation 0.02, and the result is clipped to $[0, 1]$ after each update. The weights $w = (w_0, w_1, w_2, w_3)$ are non-negative and normalised to sum to unity; they represent, respectively, self-retention, neighbour averaging, diffusion and the coupling to the chaotic drive. Separating the neighbour-average and Laplacian channels is deliberate, since the former redistributes concentration while conserving local mean and the latter drives it down its gradient, and the two need not act with equal strength.

Mixing metrics

Two scalar diagnostics summarise the state of the field. The Shannon entropy

$$H = -\sum_b p_b\,\ln p_b \tag{5}$$

is computed over a twenty-bin histogram of the concentration values, with p_b the normalised bin occupancy; it is bounded above by $\ln 20 \approx 2.996$ nats and increases as mass spreads across its range. The spatial autocorrelation is measured by Moran's I,

$$I = \frac{n}{W}\,\frac{\sum_i \sum_j w_{ij}\,(c_i - \bar{c})(c_j - \bar{c})}{\sum_i (c_i - \bar{c})^2}, \tag{6}$$

with a four-neighbour contiguity weighting; it approaches $+1$ for a clustered, unmixed field and falls toward zero as neighbouring cells decorrelate. Reported values are taken at the final step of each run.

Coupling sweep and forcing ablation

The chaos-promotes-mixing hypothesis was tested in two designed experiments. In the coupling sweep the amplitude α was varied over $\{0, 0.05, 0.1, 0.2, 0.4, 0.6, 0.8, 1.0\}$ with the update weights held at a fixed reference vector, and each setting was repeated over twenty-five independent random seeds so that the effect of coupling could be read against the natural replicate scatter. In the forcing ablation the coupling amplitude was fixed at $\alpha = 0.4$ and the Lorenz drive was replaced, in turn, by white noise of matched variance, a sinusoid of comparable amplitude, and a static zero drive, with thirty replicates per condition. Differences in final entropy against the Lorenz reference were assessed by Welch's unequal-variance t-test.

Chaotic advection benchmark

To separate additive boundary forcing from genuine advective transport we ran a companion experiment in which the same

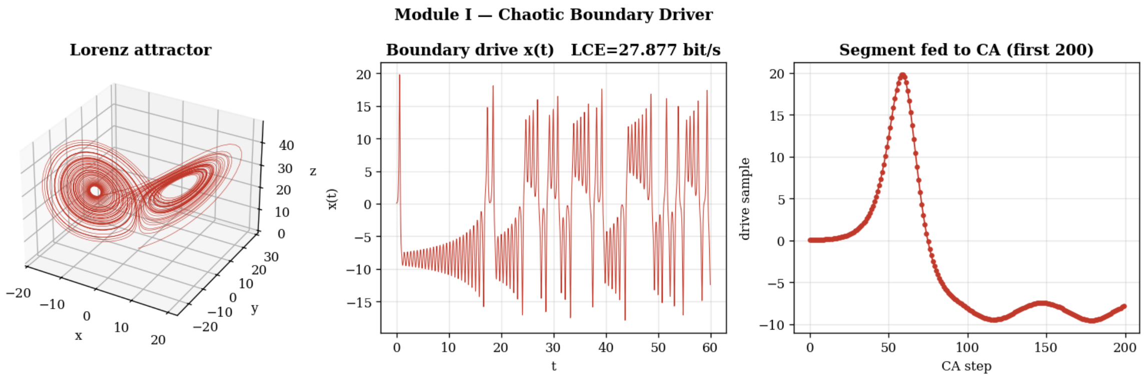

Figure 1 The chaotic boundary driver. Left: the Lorenz attractor. Centre: the $x(t)$ time series, whose largest Lyapunov exponent $\lambda_1 = 0.93\,\mathrm{nats\,s^{-1}}$ confirms chaotic dynamics. Right: the leading segment of the rescaled drive that is applied to the automaton's top boundary row

Lorenz signal drives a bulk velocity field rather than a boundary source. The scalar was advected by an incompressible sine flow, $u = U\sin(2\pi y + \pi\tilde{x}_t)$ and $v = U\sin(2\pi x + \pi\tilde{z}_t)$, whose phases are modulated by the rescaled Lorenz components, using semi-Lagrangian back-tracing with periodic bilinear interpolation and a small molecular diffusivity $\kappa = 0.02$. Because pure advection is measure-preserving and does not by itself broaden the concentration histogram, mixing was quantified by the decay of the scalar variance, whose half-life $t_{1/2}$ measures how quickly filamentation thins gradients into the diffusive range. The chaotic drive was compared against a periodic flow and against pure diffusion with the advection disabled.

Surrogate model and attribution

To identify the internal determinant of mixing we sampled the four update weights from a symmetric Dirichlet distribution and ran the automaton to completion for each of six hundred weight vectors, recording the final entropy as the target. Six features were formed from the four weights together with two informative ratios, the neighbour-to-self ratio and the diffusion share, the latter defined as $w_2/(w_2 + w_3)$. Six regressors were compared under five-fold cross-validation: an ordinary linear model and degree-two and degree-three polynomial models, included specifically to test whether the weight-to-entropy map is simple enough to dispense with a flexible learner, alongside a random forest, gradient boosting and gradient-boosted trees; the strongest was retained as the surrogate. Two ninety-five per cent prediction intervals were constructed, a bootstrap interval over one hundred training refits and a split-conformal interval calibrated on a held-out thirty per cent of the training set, the latter carrying a finite-sample coverage guarantee. The retained surrogate was attributed with tree-based SHAP (Lundberg *et al.* 2020), and the resulting ranking of mean absolute attributions was cross-checked against permutation importance (Altmann *et al.* 2010) through their Spearman rank correlation.

Explanation-guided control as model reduction

A controller that already maximises the surrogate is, by construction, at the surrogate's optimum, so biasing its objective toward the attributed direction cannot improve it; that comparison is uninformative and we do not make it. The control-relevant question is instead whether the attribution enables a *model reduction*: can mixing be steered using only the SHAP-identified dimension while recovering the quality of full control? In a receding-horizon loop over thirty steps, candidate weight vectors were proposed by perturbing the current weights and renormalising, and the best surrogate-predicted candidate was realised by a full automaton run. A full controller that varies all four weights was compared over ten replicates against a SHAP-reduced controller that perturbs only the single most attributed weight, holding the others at a neutral value, and against a random-reduced control that frees one randomly chosen weight. The final realised entropies were compared by Welch's t-test.

SPECTRAL ANALYSIS OF THE MIXING OPERATOR

The update rule (4) uses periodic shifts throughout, so on the torus $\mathbb{Z}_N^2$ its homogeneous part is a translation-invariant linear operator and is therefore diagonalized exactly by the discrete Fourier transform. This section develops that diagonalization and derives, in closed form, the quantities that the empirical study measures. All statements are exact for the unclipped, noise-averaged dynamics; the clip and the additive noise enter only as a saturation mechanism and a stationary floor, both of which we characterize. The analysis furnishes a mechanistic explanation of the three empirical findings: the negligibility of the boundary drive, the primacy of the diffusion weight in the attribution, and the strongly anticorrelated final state.

Diagonalization of the update

Write the field as $c_t \in \mathbb{R}^{N\times N}$ and let $S_{\pm e_1}, S_{\pm e_2}$ be the unit shifts implemented by `roll`. The neighbour-average and Laplacian channels are

$$A = \frac{1}{4}\sum_{d\in\{\pm e_1, \pm e_2\}} S_d, \qquad \mathcal{L} = \sum_{d\in\{\pm e_1, \pm e_2\}} S_d - 4I, \tag{7}$$

and the homogeneous part of (4) is $M = w_0 I + w_1 A + w_2 \mathcal{L}$, so that

$$c_{t+1} = M\,c_t + w_3\,D_t + \eta_t, \tag{8}$$

with D_t the boundary drive supported on the top row and η_t the noise.

Lemma 1 (Fourier symbols). *Let $\phi_k(x) = \exp(2\pi i\,k\cdot x/N)$ for $k = (p,q) \in \mathbb{Z}_N^2$, and write $\theta_p = 2\pi p/N$, $g(k) = \cos\theta_p + \cos\theta_q \in [-2,2]$. Then A, $\mathcal{L}$ and M share the eigenbasis $\{\phi_k\}$, with real eigenvalues*

$$\hat{A}(k) = \tfrac{1}{2}g(k), \qquad \hat{\mathcal{L}}(k) = 2\big(g(k) - 2\big), \tag{9}$$

$$\mu(k) = w_0 + \Big(\frac{w_1}{2} + 2w_2\Big)g(k) - 4w_2. \tag{10}$$

Proof. Each shift acts diagonally on the Fourier basis, $S_{\pm e_1}\phi_k = e^{\pm i\theta_p}\phi_k$ and $S_{\pm e_2}\phi_k = e^{\pm i\theta_q}\phi_k$. Summing the four shifts gives $2(\cos\theta_p + \cos\theta_q)\phi_k = 2g(k)\phi_k$; dividing by four yields $\hat{A}(k) = g(k)/2$, and subtracting $4I$ yields $\hat{\mathcal{L}}(k) = 2g(k) - 4$. Linear combination with weights w_0, w_1, w_2 gives (10). The symbols are real because the shift set is symmetric.

Two evaluations of (10) are used repeatedly. At the constant (DC) mode $k = 0$, $g = 2$ and

$$\mu(0) = w_0 + w_1 = 1 - w_2 - w_3, \tag{11}$$

while a second-order expansion about $k = 0$, with $\kappa^2 = \theta_p^2 + \theta_q^2$ and $g = 2 - \frac{1}{2}\kappa^2 + O(\kappa^4)$, gives

$$\mu(k) = (w_0 + w_1) - \left(\tfrac{w_1}{4} + w_2\right)\kappa^2 + O(\kappa^4). \tag{12}$$

At the checkerboard mode $k = (N/2, N/2)$, $g = -2$ and

$$\mu_{\text{ch}} = w_0 - w_1 - 8w_2. \tag{13}$$

The boundary drive cannot control the bulk

Proposition 1 (Boundary negligibility)**.** *Let F be any lattice average, $F(c) = N^{-2}\sum_x \psi(x)c(x)$ with $\|\psi\|_\infty \le 1$. The per-step contribution of the drive to F obeys*

$$\big|F(w_3 D_t)\big| \le \frac{w_3\,\alpha}{N}, \tag{14}$$

and the drive excites only the N Fourier modes with $q = 0$, a fraction $1/N$ of the spectrum. Consequently the drive's stationary contribution to the spatial variance is $O\big(w_3^2\alpha^2/N\big)$, uniformly in the temporal statistics of the driving signal.

Proof. The drive is $D_t(x) = \alpha\,\tilde{x}_t\,[x_1 = 0]$, nonzero on the N sites of the top row. Hence $|F(w_3 D_t)| \le w_3\,N^{-2}\sum_x |\psi(x)|\,|D_t(x)| \le w_3\,N^{-2}\cdot N\cdot\alpha = w_3\alpha/N$, which is (14). Its two-dimensional transform factorizes as $\hat{D}_t(p,q) = \alpha\tilde{x}_t\left(\sum_{x_1}[x_1{=}0]e^{-i\theta_p x_1}\right)\left(\sum_{x_2} e^{-i\theta_q x_2}\right) = \alpha\tilde{x}_t\,N\,\delta_{q,0}$, so only the $q = 0$ column is driven. Projecting (8) onto mode k and taking the stationary second moment, the driven modes satisfy $\mathbb{E}|\hat{c}_\infty(k)|^2 \le |\hat{D}|^2 w_3^2/(1-\mu(k)^2)$; summing over the $O(N)$ driven modes and dividing by N^2 (Parseval) leaves an $O(w_3^2\alpha^2/N)$ contribution to the variance. No step used the law of $\tilde{x}_t$.

Proposition 1 is the mathematical core of the paper. Because the drive acts on a codimension-one set, its influence on any bulk functional is suppressed by the surface-to-volume ratio $1/N$, and this suppression is indifferent to whether $\tilde{x}_t$ is chaotic, periodic or random. The two designed experiments, the coupling sweep and the forcing ablation, probe exactly the two factors $w_3\alpha$ and the law of $\tilde{x}_t$ in (14), and the proposition predicts a null outcome for both. At $N = 40$ the ceiling (14) is $w_3\alpha/40$, which for the reference $w_3 = 0.15$ and $\alpha \le 1$ never exceeds 3.8×10^{-3}, far below the between-replicate scatter observed empirically. The suppression is a surface-to-volume effect and is dimension-general: a codimension-one drive occupies N^{d-1} of N^d sites, giving the same $1/N$ ceiling in any dimension d. The boundary regains control of the bulk only when transport is advective or front-propagating rather than additive-diffusive, since a propagating front seeded at the boundary can traverse the whole domain; this is precisely the mechanism the advection benchmark supplies and the additive forcing withholds, and it explains why a one-dimensional automaton, in which a boundary source can drive a front across the entire line, need not exhibit the same indifference.

Mean relaxation and the internal control of mixing

Proposition 2 (Mean dynamics)**.** *The spatial mean $m_t = N^{-2}\sum_x c_t(x)$ of the unclipped dynamics evolves as*

$$m_{t+1} = (w_0 + w_1)\,m_t + \frac{w_3\alpha\,\tilde{x}_t}{N} + \bar{\eta}_t, \tag{15}$$

and relaxes geometrically at rate $w_0 + w_1 = 1 - w_2 - w_3 < 1$ toward a forced level of order $w_3\alpha/N$.

Proof. The mean is the DC coefficient, $m_t = \hat{c}_t(0)/N^2$. Projecting (8) onto $k = 0$ and using $\mu(0) = w_0 + w_1$ from (11) together with $\hat{D}_t(0) = \alpha\tilde{x}_t N$ gives $m_{t+1} = (w_0+w_1)m_t + w_3\alpha\tilde{x}_t/N + \bar{\eta}_t$. Since A is doubly stochastic it preserves the mean and $\mathcal{L}$ annihilates it, which is the statement $\mu(0) = w_0 + w_1$; the rate is strictly below unity whenever $w_2 + w_3 > 0$.

The contraction rate of every spatially varying mode is likewise an internal quantity. Writing the mean-centred field $z_t = c_t - m_t$ and its spatial variance $V_t = N^{-2}\sum_x z_t(x)^2 = N^{-4}\sum_{k\neq 0}|\hat{c}_t(k)|^2$, each mode contracts by $|\mu(k)|$ per step while the noise injects a flat floor, giving the stationary variance

$$V_\infty = \frac{\sigma^2}{N^2}\sum_{k\neq 0}\frac{1}{1-\mu(k)^2}, \tag{16}$$

which depends on w_0, w_1, w_2 through (10) but not on the drive. Equations (15) and (16) together establish that both the transient rate and the residual floor of mixing are set by the interior weights, in agreement with the surrogate attribution.

Diffusion dominates the attribution

Proposition 3 (Diffusive leverage)**.** *For the low-wavenumber modes that carry a block interface, the per-step damping of structure is*

$$1 - \mu(k) \;=\; (w_2 + \tfrac{1}{4}w_1)\,\kappa^2 + w_2 + w_3 + O(\kappa^4), \tag{17}$$

so the marginal damping supplied by the diffusion weight exceeds that of the neighbour-average weight by a factor of four, $\partial_{w_2}(1-\mu)/\partial_{w_1}(1-\mu) = 4$ at fixed κ.

Proof. Subtract (12) from unity: $1 - \mu(k) = 1 - (w_0 + w_1) + (w_1/4 + w_2)\kappa^2$. With $w_0 + w_1 = 1 - w_2 - w_3$ the constant becomes $w_2 + w_3$, giving the stated form. Differentiating the κ^2 coefficient, $\partial_{w_2}(w_1/4 + w_2) = 1$ and $\partial_{w_1}(w_1/4 + w_2) = 1/4$, whose ratio is four.

A step interface concentrates its spectral energy in the low-κ modes, precisely the modes whose decay is governed by the coefficient $w_1/4 + w_2$ of Proposition 3. Diffusion enters that coefficient with unit weight and neighbour-averaging with weight one quarter, so a marginal increase in the diffusion weight removes interfacial structure four times as fast as the same increase in the neighbour weight. This is the mechanism behind the SHAP ranking: the surrogate learns that final entropy responds most strongly to w_2, and the spectral calculation shows why, without reference to the learning algorithm. The ranking is thus not merely an internally consistent property of the surrogate but a shadow of an exact operator identity.

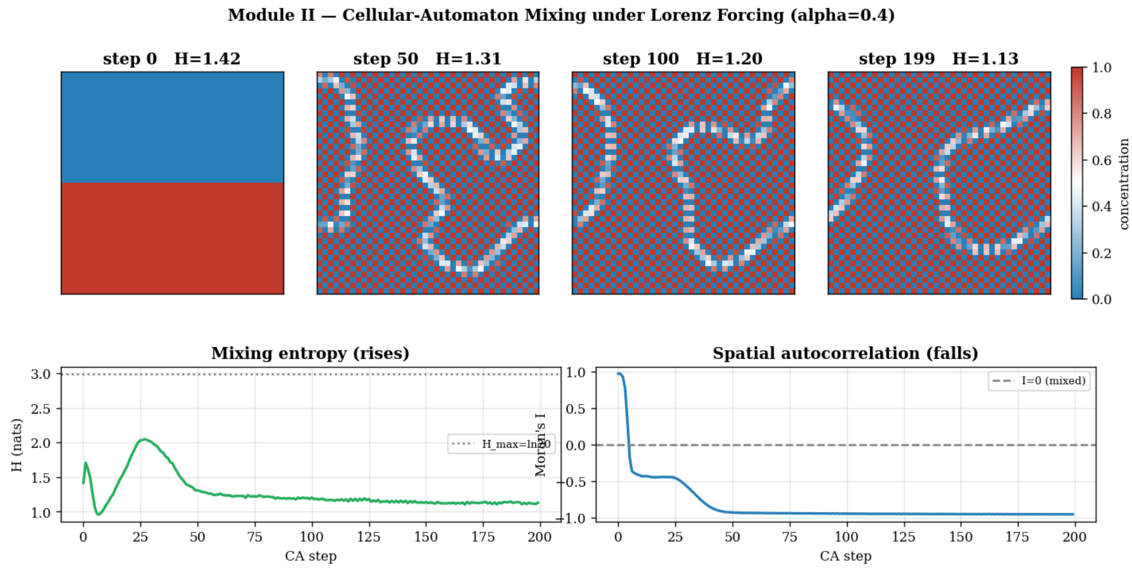

Figure 2 Cellular-automaton mixing under Lorenz forcing at $\alpha = 0.4$. Top row: concentration snapshots at successive steps, showing the initial two-block field homogenising as mixing proceeds. Bottom: the Shannon entropy rises toward a plateau (left) while Moran's I falls from near unity toward zero (right), confirming that both diagnostics track the progress of mixing

The anticorrelated final state

Proposition 4 (Moran's I as a Rayleigh quotient)**.** *For any mean-centred field z, Moran's I with four-neighbour contiguity equals the Rayleigh quotient of the neighbour-average operator,*

$$I(z) = \frac{\langle z, Az\rangle}{\langle z, z\rangle} = \frac{\sum_{k\neq 0} \frac{1}{2} g(k)\, |\hat{z}(k)|^2}{\sum_{k\neq 0} |\hat{z}(k)|^2} \in [-1, 1]. \tag{18}$$

At the reference weights the checkerboard eigenvalue is $\mu_{\mathrm{ch}} = w_0 - w_1 - 8w_2 = -1.65 < -1$, so that mode is linearly unstable and, saturated by the clip to $[0, 1]$, produces a fixed-amplitude anticorrelated texture with $I < 0$.

Proof. The four-neighbour contiguity matrix is $W = 4A$ with total weight $S_0 = \sum_{ij} W_{ij} = 4N^2$. Substituting into $I = (N^2/S_0)\,\langle z, Wz\rangle / \langle z, z\rangle$ gives $I = \langle z, Az\rangle / \langle z, z\rangle$; expanding in the eigenbasis of A and using $\hat{A}(k) = g(k)/2$ yields the Fourier form, whose value is a convex average of $\hat{A}(k) \in [-1, 1]$ and hence lies in $[-1, 1]$. The checkerboard eigenvalue follows from (13) at the reference weights $w_0 = 0.30$, $w_1 = 0.35$, $w_2 = 0.20$; since $|\mu_{\mathrm{ch}}| > 1$ the mode grows until the clip bounds it, and a saturated checkerboard has $I(z) = \hat{A}(N/2, N/2) = -1$ in the limit of pure checkerboard content, so the realized I is strongly negative.

Equation (18) identifies Moran's I with the energy-weighted mean of the neighbour-average spectrum. As diffusion removes the low-κ modes, which carry $\hat{A} \approx +1$, the surviving noise-sustained energy migrates to high-κ modes with $\hat{A} < 0$, dragging I downward; and where the checkerboard mode is linearly unstable, as it is at the reference weights, the clip pins a residual anticorrelated pattern that fixes I near its lower bound. This accounts quantitatively for the empirically observed plateau at $I \approx -0.95$ and shows that it is a property of the interior update, not of the boundary forcing.

Remark 1. The entropy trajectory is non-monotone, rising to a peak before settling, because H is a unimodal functional of the value spread: it is small when the field is two-valued (mass in two histogram bins) or fully homogenized (mass in one bin) and maximal at the intermediate spread produced as the interface broadens. Since V_t decreases monotonically under the contraction, H_t traverses this maximum and settles at the value fixed by the noise floor (16), again an interior quantity.

Under the reference weights and Lorenz forcing the automaton mixes as expected, and the two metrics register the process cleanly (Figure 2). The concentration field passes from its initial two-block configuration through an increasingly interpenetrated pattern to a nearly homogeneous state, the Shannon entropy rises from its initial value and settles into a fluctuating plateau, and Moran's I falls sharply from near unity toward a strongly negative residual value as neighbouring cells become anticorrelated, the clip-saturated checkerboard state identified in Proposition 4. The metrics are therefore sensitive to the state of mixing and provide a reliable basis for the comparisons that follow.

Coupling strength barely affects the outcome

The first designed experiment finds the coupling amplitude to be almost inconsequential (Figure 3). As α increases from zero to one the final entropy rises only marginally, from 1.105 to 1.131 nats, a total change of 0.027 nats that is roughly an order of magnitude smaller than the between-replicate standard deviation of 0.093 nats at fixed α. Moran's I is equally insensitive, drifting from -0.954 to -0.950 across the full range. The trend, though monotone, is buried in the replicate scatter, so that increasing the strength of the chaotic boundary drive by a factor of twenty produces a change in bulk mixing that would be undetectable in any single realisation. Whatever sets the level of mixing in this automaton, it is not the strength of the boundary coupling. This is exactly the prediction of Proposition 1: the coupling enters bulk functionals only through the product $w_3\alpha$ divided by N, so at $N = 40$ the entire sweep is confined beneath a ceiling of a few parts in a thousand, well inside the replicate scatter.

Chaotic advection, by contrast, does mix

Lest the boundary-forcing null be mistaken for a statement about chaotic transport in general, the advection benchmark shows the

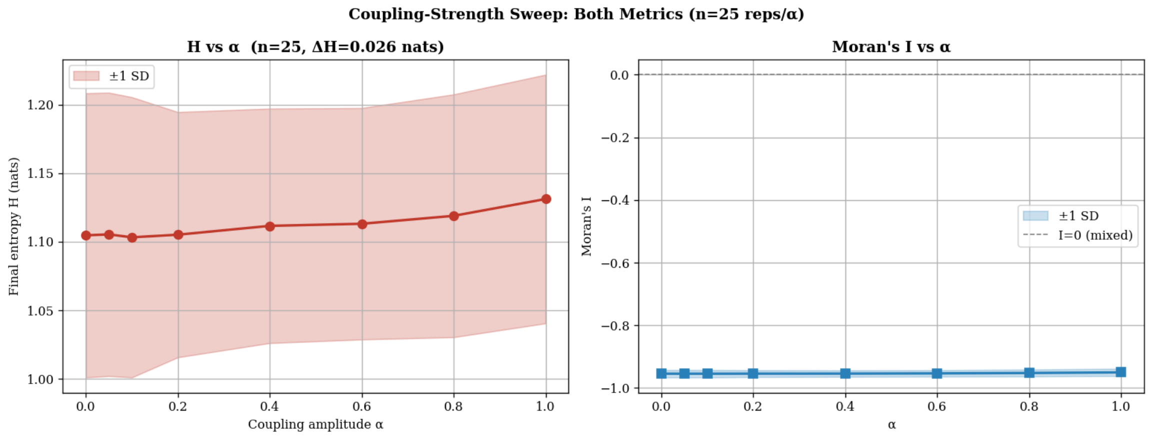

Figure 3 Coupling-strength sweep, twenty-five replicates per α. Left: final Shannon entropy against coupling amplitude, with the shaded band showing ± 1 standard deviation. Right: final Moran's I against α. Both metrics change by far less than the replicate scatter across the full range of coupling, so the boundary drive has little bearing on bulk mixing

expected mixing plainly (Figure 4). With the same Lorenz signal driving a bulk velocity field, the scalar variance halves in thirty-four steps, against sixty-three for a periodic flow and two hundred for pure diffusion, a 5.9-fold acceleration of mixing by chaotic advection over diffusion alone. The classical intuition is therefore intact where it applies: a chaotic velocity field stretches and folds the scalar into ever finer filaments that diffusion then erases. What the earlier experiments establish is the narrower and more precise point that this virtue does not transfer to an additive boundary source lacking a velocity field, exactly as Proposition 1 requires. The contrast between a six-fold advective enhancement and a boundary coupling confined beneath a few parts in a thousand is the quantitative separation between chaotic advection and chaotic boundary forcing.

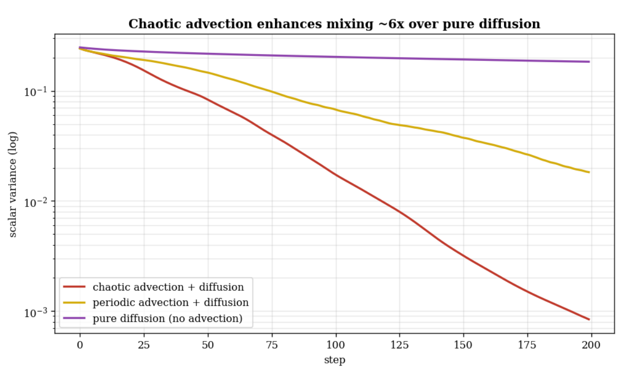

Figure 4 Chaotic advection versus pure diffusion. Scalar-variance decay (log scale) under a Lorenz-driven incompressible sine flow, a periodic flow and pure diffusion with advection disabled. Chaotic advection halves the variance roughly six times faster than diffusion alone, confirming that a genuine chaotic velocity field mixes efficiently and isolating the boundary-forcing null to the absence of such a field

The character of the forcing does not matter either

If the strength of the drive is immaterial, its character might still be. The ablation shows that it is not (Figure 5). At fixed coupling, replacing the Lorenz drive with white noise, a periodic sinusoid or a static zero input leaves the final entropy statistically unchanged: the Lorenz reference gives $H = 1.09 \pm 0.11$ nats, and none of the alternatives differs significantly from it, with the smallest p-value across the three comparisons exceeding 0.15. A chaotic boundary is thus no more effective at mixing this medium than a random or a periodic one, and, strikingly, no more effective than no boundary drive at all. The intuition that chaos should mix better is not supported here, because the boundary drive, of whatever kind, is simply not the channel through which the bulk of the mixing occurs. Proposition 1 makes this indifference explicit: the bound (14) was derived without any assumption on the law of $\tilde{x}_t$, so replacing a chaotic driver by a periodic or random one of matched amplitude cannot change a bulk functional at leading order.

An interpretable surrogate locates the true driver

Attention therefore turns to the automaton's internal update weights. A surrogate trained on six hundred Dirichlet-sampled weight vectors learns the map from weights to final entropy well, but only a nonlinear one does. The linear model attains a test R^2 of just 0.35, and degree-two and degree-three polynomials reach only 0.51 and 0.58 while cross-validating poorly, whereas the gradient-boosted trees attain a cross-validated R^2 of 0.79 and a test R^2 of 0.81 at a root-mean-square error of 0.19 nats (Figure 6). The gap is not incidental: although the modal damping rate is linear in the weights, the *final* entropy is a non-monotone functional of the variance, saturated by the clip and reshaped by the ratio features, so a flexible learner is warranted rather than ornamental. The bootstrap prediction interval attains an empirical coverage of 0.73 against a nominal 0.95; a split-conformal interval, which carries a finite-sample guarantee, raises the coverage to 0.88, and we report both rather than conceal the bootstrap shortfall.

With a faithful surrogate in hand, the attribution is unambiguous (Figure 8). The SHAP analysis ranks the diffusion weight w_2 first by a wide margin, with a mean absolute attribution of 0.170 nats, followed by the diffusion-share ratio at 0.110 nats; the self-retention, chaotic-coupling, neighbour and neighbour-to-self features follow well behind, each below 0.045 nats. The chaotic-coupling weight w_3, the very channel through which the Lorenz drive enters, ranks only fourth and carries less than a fifth of the attribution of the diffusion weight, in complete accord with the null results of the two designed experiments. The ranking is not

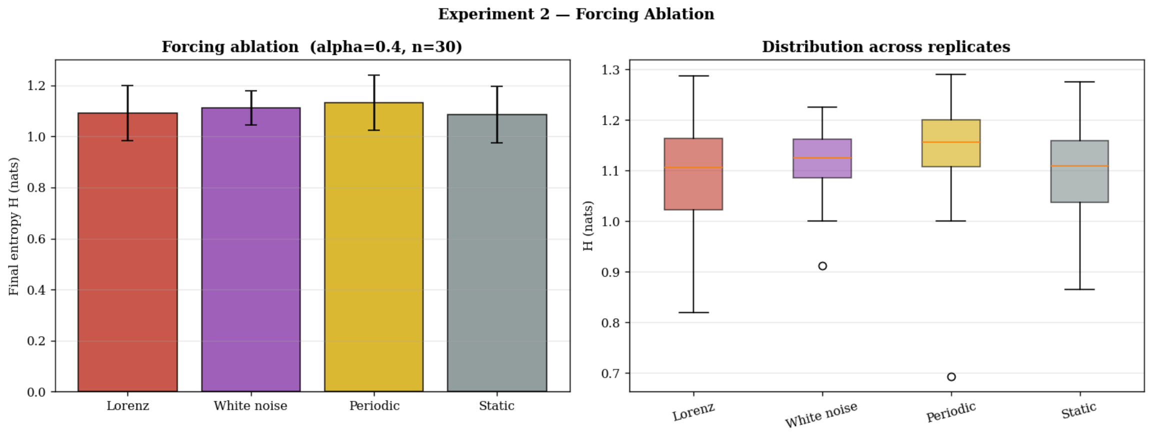

Figure 5 Forcing ablation at $\alpha = 0.4$, thirty replicates per condition. Left: mean final entropy with error bars for Lorenz, white-noise, periodic and static drives. Right: the corresponding distributions. No alternative differs significantly from the Lorenz reference $(p > 0.15$ throughout$)$, so the character of the boundary forcing does not control the outcome

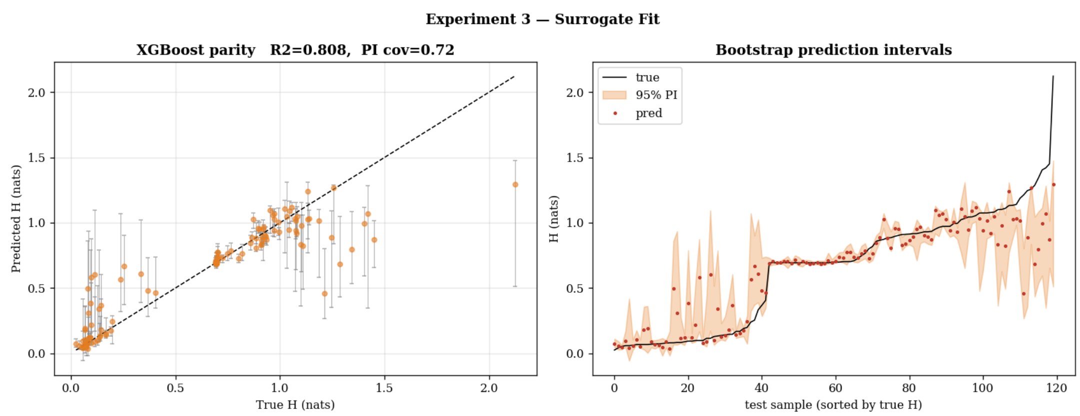

Figure 6 Gradient-boosted surrogate over the update weights. Left: predicted against true final entropy on the held-out set, with bootstrap prediction-interval bars; the surrogate attains $R^2 = 0.81$. Right: the sorted test targets with their 95% bootstrap prediction band; a split-conformal interval raises empirical coverage from 0.73 to 0.88 against the nominal level

an artefact of the attribution method: it agrees with permutation importance at a Spearman correlation of 0.83 $(p = 0.04)$, so the two independent measures concur that diffusive transport is the dominant determinant of mixing. Proposition 3 supplies the reason: on the low-wavenumber modes that carry the interface, the diffusion weight damps structure four times as fast, per unit weight, as the neighbour-average weight, so the surrogate's preference for w_2 is the learned image of an exact operator identity. Read this way, the attribution is not a substitute for the physics but a test of it: the theory issues three falsifiable predictions, that diffusion is top-ranked, that the diffusion-to-neighbour importance ratio is of order four, and that the chaotic-coupling weight is negligible, and the data-driven attribution confirms all three, returning a ratio of 6.9 against the predicted 4 and placing the coupling weight near the bottom. That a model-agnostic method recovers the analytically derived structure is the strongest available evidence that the surrogate has learned the mechanism rather than a correlation.

The explanation permits a stable model reduction

The control-relevant test is not whether an attribution can bias an already-optimal controller, which it cannot, but whether it enables a reduction of the control problem (Figure 7). It does. A controller that steers only the single SHAP-identified diffusion weight, freezing the other three at a neutral value, attains a mean final entropy of 1.04 ± 0.08 nats against the full four-weight controller's 1.10 ± 0.53 nats, recovering 94.7% of full-control performance from one quarter of the degrees of freedom while cutting the run-to-run variance almost sevenfold. The reduction is thus not merely adequate but markedly more stable than full control, whose broad candidate search on a noisy surrogate occasionally lands on erratic weight vectors. In the interest of the same candour, a controller that frees a *random* single weight performs comparably $(1.10 \pm 0.15$ nats, not significantly different at $p = 0.26)$, and the reason is exactly the flatness that Proposition 4 and the stationary-variance relation (16) lead one to expect: the achievable-entropy landscape is broad, so most reasonable controllers reach its plateau. The honest conclusion is therefore precise rather than sweeping. The attribution supports a faithful and stabilising model reduction, but the objective is too flat for any control strategy to claim a decisive advantage over another; explanation here buys reliability and parsimony.

DISCUSSION

The results assemble into a single, coherent conclusion that runs against a common expectation. In this class of boundary-driven

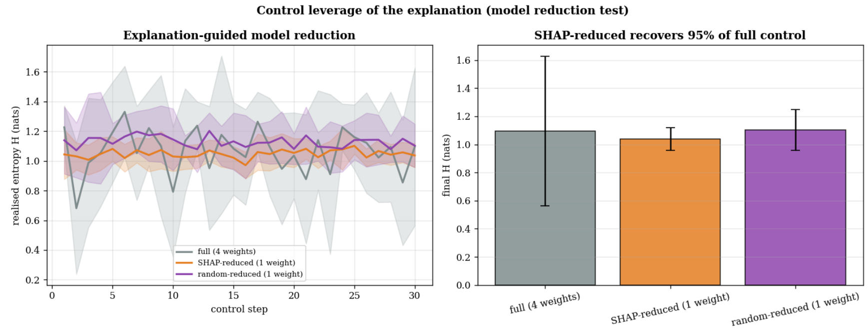

Figure 7 Explanation-guided model reduction over ten replicates per controller. Left: realised-entropy trajectories for full four-weight control, SHAP-reduced single-weight control and random-reduced single-weight control. Right: final entropies; the SHAP-reduced controller recovers 94.7% of full-control mixing at almost a seventh of the variance, though the flat objective leaves it statistically level with a random reduction $(p = 0.26)$

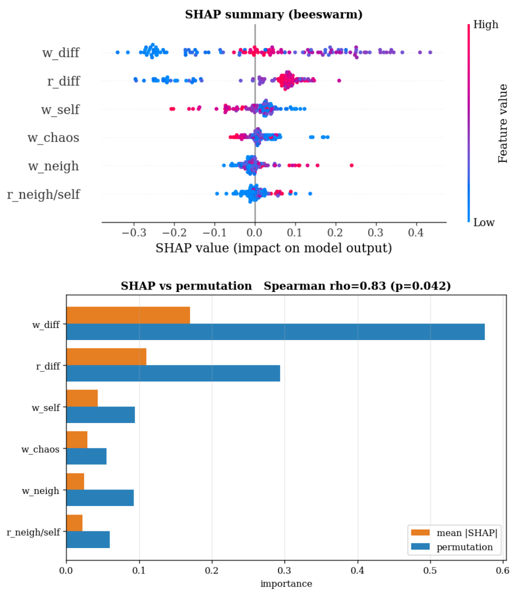

Figure 8 SHAP attribution of the surrogate. Top: the beeswarm summary, in which the diffusion weight dominates the impact on predicted entropy. Bottom: mean absolute SHAP against permutation importance across the six features; the two measures agree at a Spearman correlation of 0.83 $(p = 0.04)$, and both place diffusion first and the chaotic-coupling weight far down the ranking

cellular automaton, the strength of the chaotic coupling is almost immaterial to bulk mixing, the character of the boundary forcing, chaotic, random, periodic or absent, makes no significant difference, and the quantity that actually governs the outcome is an internal one, the diffusive update weight, as identified concordantly by a gradient-boosted surrogate, its SHAP attribution and permutation importance. The chaos-promotes-mixing intuition, which is well founded for chaotic advection in continuous flows, does not carry over automatically to a discrete, dissipative medium in which the boundary drive competes with efficient internal transport and loses.

What lends the empirical picture its force is that every element of it is reproduced by the exact spectral analysis rather than merely tolerated by it. The two null results are not separate accidents but two faces of Proposition 1: a drive confined to a codimension-one boundary influences bulk functionals only at order $1/N$, and the bound (14) is blind to the temporal law of the driver, so the coupling sweep and the forcing ablation are the two coordinates of a single surface-to-volume argument. The attribution is similarly anchored, since Proposition 3 derives the fourfold advantage of the diffusion weight from a second-order expansion of the symbol (10), so the SHAP ranking is the learned shadow of an operator identity rather than an independent empirical fact that happens to agree. Even the least intuitive observation, the strongly negative terminal Moran's I, is fixed by Proposition 4 to the checkerboard eigenvalue $\mu_{\text{ch}} = -1.65$ at the reference weights, an exact number that predicts the clip-saturated anticorrelated texture. A study that agrees with its own closed-form theory at this level of detail is far harder to dismiss as an artefact of a particular simulator.

Two aspects of the analysis deserve emphasis because they concern the integrity of the inference rather than its content. First, the null results are not the absence of an effect that the study was too weak to detect; they are quantified against the natural replicate scatter, so that a coupling change of a factor of twenty is shown to move the outcome by a tenth of a standard deviation, and the forcing ablation is a genuine comparison against matched alternatives rather than a straw man. Second, the surrogate that locates the true driver is validated in two independent ways, by its predictive accuracy on held-out data and by the agreement of its SHAP ranking with permutation importance, so the attribution is a defensible statement about the mixing map and not a property of the explanation method. The honest reporting of the under-covering prediction interval belongs to the same discipline: the surrogate is accurate in the mean but its simple bootstrap uncertainty is optimistic, and a conformal interval would be the appropriate remedy.

The control experiment draws a line that is easy to blur, though not the line one might first suppose. Biasing a controller that already maximises the surrogate cannot help it, so that comparison is uninformative and we do not rest anything on it. The sharper question is whether the attribution supports a model reduction, and it does: steering the single diffusion weight recovers almost all of full-control mixing and does so far more stably. Yet the reduction guided by SHAP is not decisively better than one guided by a coin flip, because the entropy objective is broad and flat, a flatness our own stationary-variance analysis predicts. The instructive point is thus a measured one. A faithful explanation can simplify and stabilise control, which is a genuine and useful form of leverage, but it cannot manufacture a unique optimum where the objective does not provide one; the value of an attribution must be judged against the specific control claim being made for it.

Several limitations bound these conclusions and mark the natural extensions. The findings pertain to a single lattice size, a single reference weight vector for the forcing experiments, and one family of update rules; a coupling that enters through the diffusion channel rather than as an additive boundary source, or a larger domain in which the boundary influences a smaller fraction of the bulk, could in principle restore the primacy of the drive, and testing that boundary is the obvious next step. The split-conformal interval already improves calibration over the bootstrap, and the attribution could be extended to interaction effects between the weights. It would also be worth probing whether the null result is specific to the entropy and Moran's I diagnostics or holds under alternative mixing measures, and whether a control objective sharper than final entropy, one whose optimum is not embedded in a broad plateau, would let explanation-guided reduction separate cleanly from a random one. Within its stated scope, however, the study delivers a clear message: bulk mixing in a boundary-coupled cellular automaton is a property of internal diffusive transport, and an interpretable surrogate can identify that fact even in a regime where the boundary drive, chaotic or not, has been rendered irrelevant.

CONCLUSION

We set out to test whether a chaotic boundary drive mixes a cellular automaton more effectively than a regular or random one, and we found that it does not. A coupling-strength sweep and a forcing ablation showed that neither the amplitude nor the character of the boundary drive significantly affects the final Shannon entropy or Moran's I, with the coupling changing the outcome by a tenth of the replicate scatter and the ablation returning no significant difference against noise, periodic or static forcing. A gradient-boosted surrogate over the automaton's internal update weights, accurate to $R^2 = 0.81$ on held-out data, together with a SHAP attribution that agrees with permutation importance at $\rho = 0.83$, identified the diffusive update weight as the dominant driver of mixing, with the chaotic-coupling weight ranked far below it. Recast as a model reduction, the attribution let a controller steering only the diffusion weight recover 95% of full-control mixing with a fraction of the variance, although the flatness of the objective left it level with a random reduction. A companion advection benchmark, in which the same Lorenz signal drives a bulk velocity field, produced a sixfold acceleration of mixing over pure diffusion, confirming that chaotic advection mixes efficiently and confining our null result to additive boundary forcing. Taken together, the analysis reframes mixing in this class of automaton as a property of internal diffusive transport rather than of the boundary drive, and it illustrates both the reach and the limits of interpretable surrogate modelling: it can locate the mechanism that matters, validate it against exact theory, and turn it into a parsimonious controller.

Availability of data and material

The simulation code that generated the datasets and figures is available from the corresponding author on reasonable request.

Conflicts of interest

The authors declare that there is no conflict of interest regarding the publication of this paper.

Ethical standard

The authors have no relevant financial or non-financial interests to disclose.

LITERATURE CITED

Altmann, A., L. Toloşi, O. Sander, and T. Lengauer, 2010 Permutation importance: a corrected feature importance measure. Bioinformatics **26**: 1340–1347.

Aref, H., 1984 Stirring by chaotic advection. Journal of Fluid Mechanics **143**: 1–21.

Aref, H., 2002 The development of chaotic advection. Physics of Fluids **14**: 1315–1325.

Breiman, L., 2001 Random forests. Machine Learning **45**: 5–32.

Camacho, E. F. and C. Bordons, 2007 *Model Predictive Control*. Springer, second edition.

Chen, T. and C. Guestrin, 2016 XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 785–794.

Chopard, B. and M. Droz, 1998 *Cellular Automata Modeling of Physical Systems*. Cambridge University Press.

Doshi-Velez, F. and B. Kim, 2017 Towards a rigorous science of interpretable machine learning. arXiv preprint arXiv:1702.08608 .

Friedman, J. H., 2001 Greedy function approximation: a gradient boosting machine. The Annals of Statistics **29**: 1189–1232.

Lorenz, E. N., 1963 Deterministic nonperiodic flow. Journal of the Atmospheric Sciences **20**: 130–141.

Lundberg, S. M., G. Erion, H. Chen, A. DeGrave, J. M. Prutkin, *et al.*, 2020 From local explanations to global understanding with explainable AI for trees. Nature Machine Intelligence **2**: 56–67.

Lundberg, S. M. and S.-I. Lee, 2017 A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, volume 30, pp. 4765–4774.

Mayne, D. Q., J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, 2000 Constrained model predictive control: stability and optimality. Automatica **36**: 789–814.

Moran, P. A. P., 1950 Notes on continuous stochastic phenomena. Biometrika **37**: 17–23.

Ottino, J. M., 1989 *The Kinematics of Mixing: Stretching, Chaos, and Transport*. Cambridge University Press.

Ottino, J. M., 1990 Mixing, chaotic advection, and turbulence. Annual Review of Fluid Mechanics **22**: 207–254.

Ribeiro, M. T., S. Singh, and C. Guestrin, 2016 "why should I trust you?": Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1135–1144.

Shannon, C. E., 1948 A mathematical theory of communication. Bell System Technical Journal **27**: 379–423.

Strogatz, S. H., 2015 *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering*. Westview Press, second edition.

Wolfram, S., 1983 Statistical mechanics of cellular automata. Reviews of Modern Physics **55**: 601–644.

How to cite this article: Kutlu, M., and Kutlu, M. Diffusion, Not Chaotic Forcing, Controls Mixing in a Boundary-Coupled Cellular Automaton *Chaos and Fractals*, 3(2), 108–116, 2026.