

Equilibrium flow: FROM SNAPSHOTS TO DYNAMICS

Anonymous authors

Paper under double-blind review

ABSTRACT

Scientific data, from cellular snapshots in biology to celestial distributions in cosmology, often consists of static patterns from underlying dynamical systems. These snapshots, while lacking temporal ordering, implicitly encode the processes that preserve them. This work investigates how strongly such a distribution constrains its underlying dynamics and how to recover them. We introduce the *Equilibrium flow* method, a framework that learns continuous dynamics that preserve a given pattern distribution. Our method successfully identifies plausible dynamics for 2-D systems and recovers the signature chaotic behavior of the Lorenz attractor. For high-dimensional Turing patterns from the Gray-Scott model, we develop an efficient, training-free variant that achieves high fidelity to the ground truth, validated both quantitatively and qualitatively. Our analysis reveals the solution space is constrained not only by the data but also by the learning model’s inductive biases. This capability extends beyond recovering known systems, enabling a new paradigm of inverse design for Artificial Life. By specifying a target pattern distribution, we can discover the local interaction rules that preserve it, leading to the spontaneous emergence of complex behaviors, such as life-like flocking, attraction, and repulsion patterns, from simple, user-defined snapshots.

1 INTRODUCTION

Computer programs are often designed to accept an input, execute a sequence of hidden computations, and return a concise output, discarding most of the intermediate state along the way. Many natural systems work differently. Because they evolve through continuous ordinary-differential-equation (ODE) dynamics, the *process* and the *output* are inseparable: every microscopic configuration is itself a legitimate input / output of the system. When we sample local patterns in space – whether fossil distributions in palaeontological strata or luminosity-color distributions of planets (Russell, 1979; Babusiaux et al., 2018) – we capture static snapshots of an underlying flow. In other words, the observed pattern distribution is almost invariant under the underlying dynamics.

Each pattern distribution therefore encodes at least one minimal set of dynamics that could preserve it. The ability to recover these dynamics is a central challenge in many areas of science and engineering, as it underpins our capacity to understand, predict, and control the behavior of complex systems. For example, in the life sciences, the search for biomedical interventions is shaped by our ability to understand the algorithms and strategies guiding cellular decision-making, while snapshots of behavior are used to infer underlying properties of cognitive agents. Indeed, developmental and evolutionary biology are rife with inverse problems linking genomic information to system-level form and function (Lobo et al., 2014).

Kolmogorov complexity renders a direct search for such minimal dynamics for *patterns* uncomputable (Chaitin, 1975). We instead pose a more tractable inverse problem: *given a distribution of patterns, find dynamical rules that preserve that distribution*. A key empirical clue guides this reformulation. In many physical, biological, and geological systems the macroscopic distribution drifts slowly, whereas microscopic states fluctuate rapidly. Hence, we ask two questions: 1) given a distribution of patterns, how many different dynamics can preserve it? and 2) how close are the potential dynamics to the ground truth?

A rich literature exists on learning dynamics from distributions with low time resolution. The Fokker-Planck-Kolmogorov equation is typically used to model the time evolution of distributions, as demonstrated by works such as (Zhang et al., 2025; Tong et al., 2020; Chardès et al., 2023).

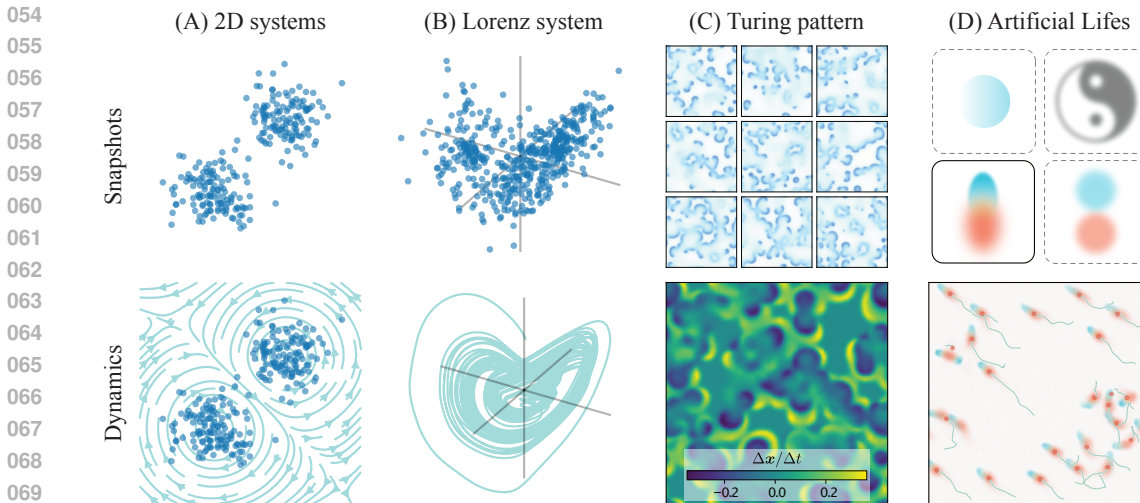


Figure 1: By learning distribution-preserving dynamics, we propose the *Equilibrium flow* method that can infer possible underlying dynamics that could preserve the observed patterns. (A) We show-
the application on simple 2D systems. Given a static 2D distribution, we can infer dynamics
that do not change the distribution. (B) Beyond this, we show the recovered dynamics of the chaotic
Lorenz system. The recovered dynamics can preserve the chaotic behavior of the original Lorenz
system. (C) By applying our method to Turing patterns, we can recover dynamics that preserve
both qualitative and quantitative features of the ground truth. (D) Furthermore, our method also
provides a new way to design artificial life forms. By specifying the target pattern, our method can
design local dynamics that preserve the target pattern and generate emergent collective behaviors.
The moving trajectories are highlighted with green lines.

However, these approaches either still require some temporal information or rely on computationally expensive neural ODEs to model dynamics and fit distributions via maximum mean discrepancy (MMD) (Gretton et al., 2006) or normalization flows. Some researchers have attempted to address this challenge by adopting spatial grids, but this unfortunately limits applicability to high-dimensional systems (Botvinick-Greenhouse et al., 2023; Yang et al., 2023). Although Li (2023) proposed a method that eliminates the need for temporal information by estimating an explicit pseudo-time for each sample, this approach is limited to periodic systems. Similarly, the work by Arts et al. (2023) directly uses the score function of a density distribution as the force field, but this inevitably misses a significant portion of dynamics that have curl. While distribution-preserving dynamics have been discussed in (Hwang et al., 2005; Rey-Bellet & Spiliopoulos, 2015), these approaches are primarily limited to accelerating sampling procedures for systems with known density functions.

Given these limitations, we propose a new, concise method for learning dynamics from static patterns without requiring temporal information or restricting the form of dynamics. Our approach, which we call *Equilibrium flow*, represents dynamics as a vector field parameterized by a neural network. The training process forces the learned flow to preserve the empirical distribution. Across both simple and complex systems, our method recovers families of dynamics consistent with the data and, crucially, retains qualitative features of the ground truth: reconstructions of chaotic systems remain chaotic, and the reconstructed 2D patterns exhibit similar movement and propagation behaviors (Figure 1).

Our contributions are threefold. Methodologically, we present *Equilibrium flow*, a flexible and broadly applicable framework for inferring dynamics from static data. Theoretically, we use this framework to show that the pattern distribution, along with model biases, imposes strong constraints on the solution space of dynamics. Finally, we demonstrate the generative potential of our method by using it to design novel forms of Artificial life (Langton, 2019).

2 LEARNING DYNAMICS FROM STATIC PATTERN DISTRIBUTION

Algorithm 1 Equilibrium flow

Require: N samples of d -dimensional data $X = \{x_i\}$. Number of epochs T . Number of trace estimates k . Diffusion time τ (by default, we choose the τ that equivalent to set $\alpha \approx 0.95$).

- 1: $s \leftarrow$ train diffusion model on X
- 2: $v_\theta \leftarrow$ initialize a neural network with parameter θ
- 3: **for** $i \in [1, 2, \dots, T]$ **do**
- 4: **for** $x \sim p_\tau(x)$ **do**
- 5: $v \leftarrow v_\theta(x)$
- 6: $\nabla \cdot v \leftarrow \frac{1}{k} \sum_{j=1}^k z_j^\top \nabla_x (z_j^\top v)$, $z_j \sim \mathcal{N}(0, I_d)$ \triangleright Hutchinson’s Trace Estimator
- 7: $\mathcal{L} \leftarrow [\nabla \cdot v + v^\top s(x)]^2$ \triangleright Minimize equation 1
- 8: Take a gradient descent step on \mathcal{L}
- 9: **end for**
- 10: **end for**

To find a continuous dynamics that preserves a given distribution $p(x)$, we can consider the probability density as a mass density and map this problem to local mass conservation. By ignoring the noise component, this approach leads us to the continuity equation:

$$\nabla \cdot [p(x, t)v(x, t)] = -\frac{\partial p(x, t)}{\partial t} = 0. \quad (1)$$

Because $p(x)$ is a scalar function, we expand it in this form:

$$p(x) \nabla \cdot v(x) + v(x)^\top \nabla p(x) = 0. \quad (2)$$

Obtaining the probability density $p(x)$ directly is often a challenging problem. To eliminate the $p(x)$ term, we note that $\nabla p(x) = p(x) \nabla \log p(x)$, and obtain:

$$\begin{aligned} \nabla \cdot [p(x)v(x)] &= p(x) \nabla \cdot v(x) + v(x)^\top p(x) \nabla \log p(x) \\ &= p(x) [\nabla \cdot v(x) + v(x)^\top s(x)], \end{aligned} \quad (3) \quad (4)$$

where $s(x) = \nabla \log p(x)$ is the *score function*, a quantity that can be effectively estimated with pre-trained diffusion models (Ho et al., 2020; Song et al., 2020; Dhariwal & Nichol, 2021) (see Appendix A.3). While the well-known Stein identity (Liu et al., 2016) requires this term to be zero only in expectation, our approach derives a much stronger, local condition by requiring the expression to be zero pointwise for all x . This gives us our target equation:

$$\nabla \cdot v(x) + v(x)^\top s(x) = 0. \quad (5)$$

The term $v(x)^\top s(x)$ is straightforward to compute, but the divergence $\nabla \cdot v(x)$ is computationally expensive for high-dimensional x . We can, however, efficiently approximate it using Hutchinson’s Trace Estimator (Hutchinson, 1989):

$$\nabla \cdot v(x) = \mathbb{E}_{z \sim \mathcal{N}(0, I_d)} [z^\top \nabla_x (z^\top v(x))]. \quad (6)$$

Thus, equation 1 can be approximated as:

$$\nabla \cdot [p(x)v(x)] \approx p(x) \left[\frac{1}{k} \sum_{j=1}^k z_j^\top \nabla_x (z_j^\top v(x)) + v(x)^\top s(x) \right], \quad (7)$$

where k is a hyper-parameter that determines the number of trace estimations. In this way, we eliminate the need to explicitly compute $p(x)$, requiring only the score function s , which can be obtained from a pre-trained diffusion model. Since we are locally optimizing $\nabla \cdot [p(x)v(x)]$ to zero, we can discard the $p(x)$ terms and define the training loss function as:

$$\mathcal{L}(v_\theta) = \mathbb{E}_{x \sim p(x)} \left[\frac{1}{k} \sum_{j=1}^k z_j^\top \nabla_x (z_j^\top v_\theta(x)) + v_\theta(x)^\top s(x) \right]^2. \quad (8)$$

162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215

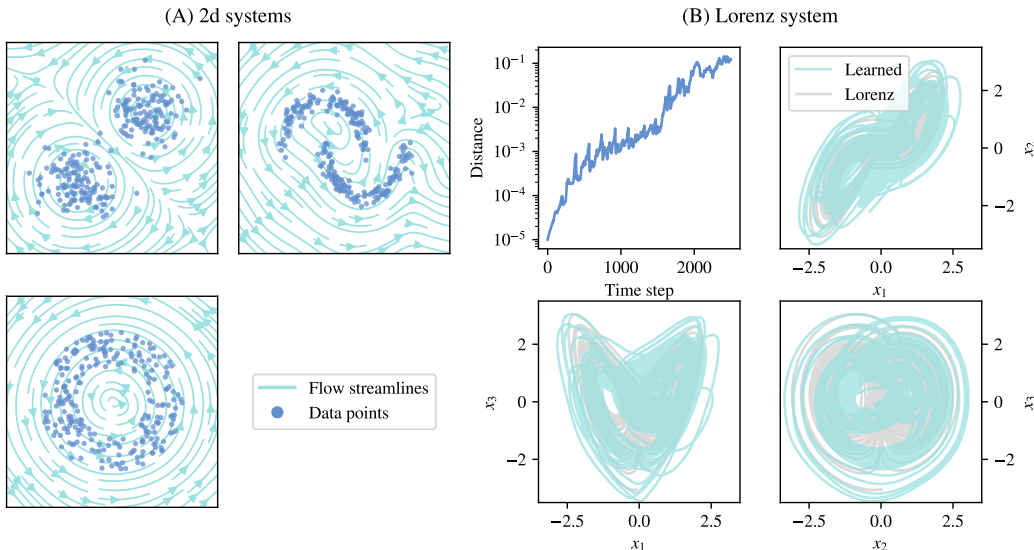


Figure 2: Dynamics learned from static pattern distribution. **(A)** After being trained on samples from statistical 2-D distributions (dots), this method identifies dynamics (arrows) that preserve the distributions. **(B)** After being trained on points sampled from the Lorenz system, the identified dynamics also exhibit chaotic behavior. The top-left figure shows how a small initial difference grows over time, indicating a positive Lyapunov exponent. The other figures compare the evolution of the original Lorenz system with that of the identified dynamics. Although they are not identical, the identified dynamics exhibit a similar overall structure to the Lorenz system and preserve its chaotic behavior.

In practice, diffusion models do not provide the exact score of $p(\mathbf{x})$ but rather that of a noise-perturbed (diffused) distribution, $p_\tau(\mathbf{x})$, where $\tau \in [0, 1]$ represents the noise level. We empirically find that using a small but non-zero noise level strikes an effective balance between training stability and the fidelity of the learned dynamics (see Appendix A.3.1 for details).

Furthermore, a key property of a system with an invariant distribution $p(\mathbf{x})$ is that its mean position, $\langle \mathbf{x} \rangle = \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})}[\mathbf{x}]$, must remain constant. This implies that the mean velocity over the distribution must be zero: $d\langle \mathbf{x} \rangle / dt = \langle \mathbf{v}(\mathbf{x}) \rangle = 0$. To enforce this physical constraint on our model, we apply a batch normalization layer to the output of the neural network \mathbf{v}_θ , forcing it to be zero-mean over each mini-batch. This technique improves training stability and helps prevent the model from converging to trivial solutions. The complete training procedure is detailed in Algorithm 1.

Once the model \mathbf{v}_θ is trained, we can use it to generate continuous dynamics for a given boundary condition. We then apply our method to 2-dimensional data distributions to demonstrate and validate its effectiveness. As shown in Figure 2A, the learned dynamics successfully preserve (or closely approximate) the underlying distribution while enabling individual samples to evolve rapidly. For instance, when the distribution is a mixture of two Gaussian distributions, the learned dynamics is a flow with two vortices in opposite directions. When the distribution is a ring, the learned dynamics is a rotational flow. For more complicated distributions, such as the two-moon distribution in the top-right corner of Figure 2A, the learned dynamics approximately follows the boundary of the distribution.

Subsequently, we extend our approach to samples generated by the Lorenz system. The learned dynamics accurately preserve the original distribution over long time horizons, as demonstrated in Appendix B. Notably, despite having no access to temporal information, our method discovers dynamics that exhibit chaotic behavior – a hallmark of the Lorenz system where trajectories with small initial differences quickly diverge. Although the learned dynamics do not exactly replicate the original Lorenz system, they capture its characteristic shape and fundamental chaotic properties, see Figure 2B. On the top-left figure of Figure 2B, we show how a small initial difference grows over

time, indicating a positive Lyapunov exponent. We applied this method 100 times with different random seeds, and all of them exhibit positive Lyapunov exponents.

3 FROM 2D PATTERN TO DYNAMICS

The relationship between static patterns and their underlying dynamics is a fundamental question in science. This problem is often approached from two perspectives. One, rooted in computational complexity theory, quantifies this relationship using measures like Kolmogorov complexity, which is theoretically elegant but generally uncomputable. A complementary line of work, in the pattern-formation literature, seeks to discover the dynamical processes capable of generating an observed pattern. A classic example is Alan Turing’s reaction-diffusion model (Turing, 1952), which used differential equations to explain how complex biological patterns can emerge from simple rules. This foundational work inspired subsequent models, such as the Gray-Scott model (Gray & Scott, 1984), that further explore this connection.

A common, often implicit, assumption in pattern formation studies is that an observed pattern strongly constrains its underlying dynamics. Otherwise, a single pattern could be explained by a multitude of theories, limiting their predictive power. This raises a critical question: how many distinct continuous dynamics can generate a given pattern distribution? To investigate this, we apply our method to patterns generated by the Gray-Scott model. This model describes the interaction of two chemical species, \mathbf{u} and \mathbf{v} , through a system of reaction-diffusion equations:

$$\frac{\partial \mathbf{u}}{\partial t} = D_u \nabla^2 \mathbf{u} - \mathbf{u}\mathbf{v}^2 + F(1 - \mathbf{u}), \quad \frac{\partial \mathbf{v}}{\partial t} = D_v \nabla^2 \mathbf{v} + \mathbf{u}\mathbf{v}^2 - (F + k)\mathbf{v}, \quad (9)$$

where D_u, D_v are diffusion coefficients, F is the feed rate, and k is the kill rate. We denote the state of the system as $\mathbf{x} = (\mathbf{u}, \mathbf{v})$ and its dynamics as $d\mathbf{x}/dt = \mathbf{v}_{\text{GS}}(\mathbf{x})$. By varying these parameters, the model can produce a wide variety of patterns. Some examples are showing in the first row of Figure 3.

To construct our pattern distributions, we selected four parameter sets that generate distinct life-like (solitons), spiral, wave, and static maze-like patterns (see Appendix C for their parameter settings). For each set, we initiated simulations from random conditions and ran them for 10,000 steps to ensure the system reached a stationary distribution. The final state of each simulation was then added to our dataset, resulting in 8,192 pattern samples for each of the four categories. However, applying our original method to such high-dimensional data is challenging due to the computational cost of Hutchinson’s trace estimator. To address this, we developed a training-free method that finds a valid subset of solutions using a linear transformation of the score function.

3.1 TRAINING-FREE METHOD

For high-dimensional distributions like Turing patterns, training a neural network \mathbf{v}_θ with Hutchinson’s trace estimator is computationally prohibitive due to its high memory usage and reliance on second-order derivatives. As an alternative, we propose a training-free approach that can identify a valid subset of solutions. This method constructs a dynamics \mathbf{v}_S from a linear transformation of the score function (Hwang et al., 2005):

$$\mathbf{v}_S(\mathbf{x}) = \mathbf{S} \nabla \log p(\mathbf{x}), \quad (10)$$

where $\mathbf{S} \in \mathbb{R}^{d \times d}$ is a skew-symmetric matrix (i.e., $\mathbf{S}^\top = -\mathbf{S}$). As proven in Appendix D.1, \mathbf{v}_S is a valid, divergence-free vector field that is orthogonal to the score function. Intuitively, it represents a flow that moves along surfaces of equal probability density.

From the perspective of Helmholtz decomposition, any vector field can be decomposed into a curl-free (gradient) part and a divergence-free (curl) part. Our \mathbf{v}_S represents only the curl component. To account for the missing gradient component, which pulls the system towards high-density regions, we add the score function back into the dynamics. This yields an irreversible Langevin SDE (Hwang et al., 2005; Rey-Bellet & Spiliopoulos, 2015):

$$d\mathbf{x} = [\mathbf{v}_S(\mathbf{x}) + \eta \nabla \log p(\mathbf{x})] dt + \sqrt{2\eta} dW_t, \quad (11)$$

where η is a hyperparameter controlling the strength of the gradient term and dW_t is the Wiener process.

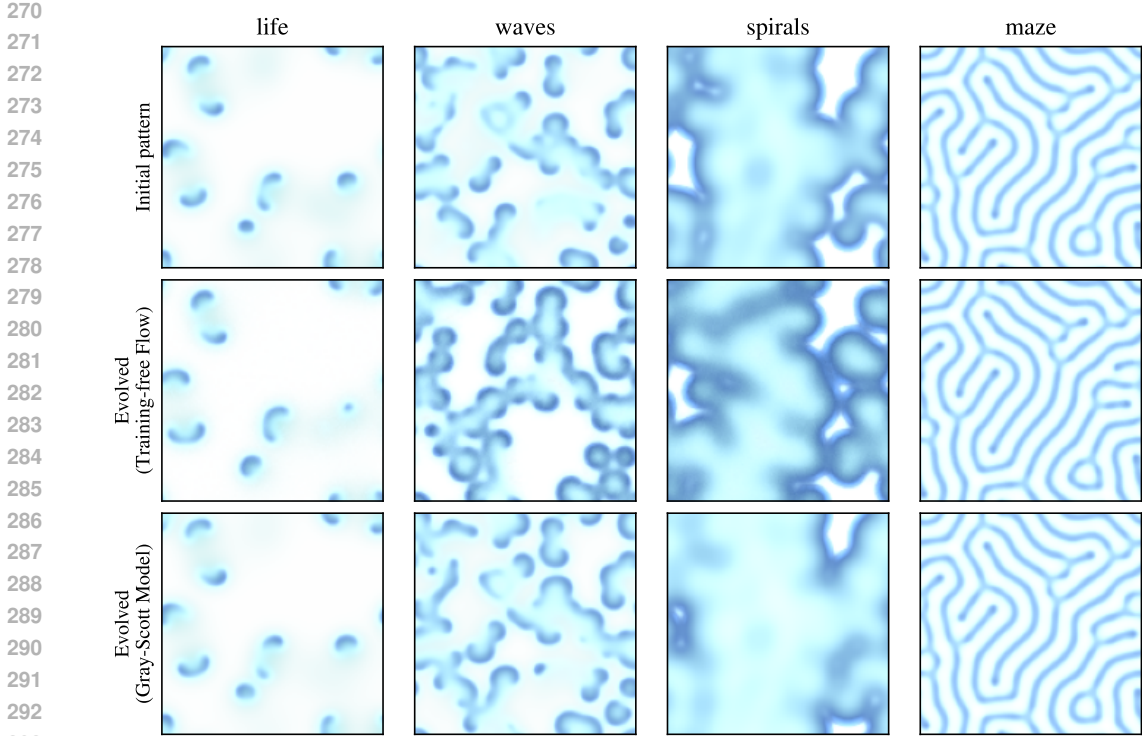


Figure 3: Comparison of dynamics recovered by our training-free method against the ground-truth Gray-Scott dynamics, starting from identical initial conditions. We display four representative pattern types: life-like (solitons), spiral, wave, and static maze-like. Although not identical, the recovered dynamics successfully reproduce key qualitative features of the ground truth, including soliton movement, wave and spiral propagation, and the static nature of maze-like patterns.

For structured data like images, the skew-symmetric transformation \mathcal{S} can be efficiently implemented as a convolutional layer with a kernel $K \in \mathbb{R}^{c \times c \times (2r+1) \times (2r+1)}$. The operator is skew-symmetric if the kernel satisfies the condition (see Appendix D.2 for proof):

$$K_{o,i,u,v} = -K_{i,o,-u,-v}, \tag{12}$$

where o, i are channel indices and u, v are spatial indices. For the two-channel Gray-Scott model, we use a simple 1×1 convolution ($r = 0$) to avoid introducing spatial asymmetries. The kernel K thus becomes a 2×2 matrix:

$$K_{::,0,0} = \gamma \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}, \tag{13}$$

where the scalar γ controls the evolution speed and its sign determines the rotational direction of the flow. In our experiments, we test both $\gamma = 1$ and $\gamma = -1$ and select the one whose resulting dynamics are most similar to the ground truth. We set the stabilization coefficient to $\eta = 0.1$ and compute the score with a noise scale of $\tau = 0.1$. The results in Figure 3 show that this approach successfully captures the key qualitative features of the ground-truth dynamics, demonstrating its effectiveness.

4 THE UNIQUENESS AND FIDELITY OF LEARNED DYNAMICS

Despite the success of qualitatively recovering complex dynamics (Figures 2 and 3), a fundamental question remains: how unique is the learned solution? That is, how many distinct continuous dynamics can preserve a given pattern distribution, and how closely does a learned solution resemble the ground truth? To investigate this, we conduct an empirical study of uniqueness (consistency among solutions) and fidelity (similarity to ground truth) on two types of systems: the Lorenz attractor and the Turing patterns from the Gray-Scott model.

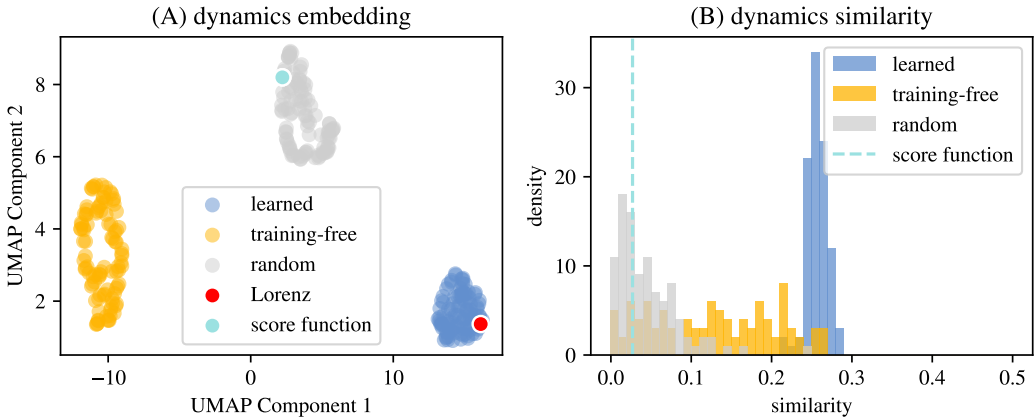


Figure 4: Uniqueness and fidelity of dynamics for the Lorenz system. Representations are formed from vector fields computed at 1024 sample points. (A) A 2D UMAP embedding of the dynamics representations. The learned dynamics (v_θ , blue dots) cluster tightly near the ground truth (orange cross), while training-free solutions (green) are more dispersed and random dynamics (gray) are scattered. (B) Cosine similarity to the ground-truth Lorenz dynamics. The trained models (v_θ) show the highest and most consistent similarity.

To compare different dynamics, we construct a high-dimensional representation for each one. For a given dynamics v , its representation e_v is formed by concatenating its vector field values at N fixed points $\{x_i^*\}_{i=1}^N$ sampled from the data distribution $p(x)$: $e_v = [v(x_1^*) \quad v(x_2^*) \quad \dots \quad v(x_N^*)]$.

We generate these representations for several dynamics: 100 independently trained models (v_θ), solutions from our training-free method, the ground-truth dynamics, the score function (s), and a baseline of 100 randomly initialized networks (v_R). Since if v is a solution, so is $-v$, we account for this inherent sign ambiguity by aligning each representation e_v with the ground-truth representation e_{GT} by ensuring a positive dot product: $e_v \cdot e_{GT} \geq 0$.

For the Lorenz system, we used $N = 1024$ sample points. We first visualize the space of solutions by applying the UMAP dimension reduction method (McInnes et al., 2018) to the representations, using a cosine metric. As shown in Figure 4A, the 100 learned dynamics (v_θ , blue dots) form a tight cluster near the ground truth, indicating a highly constrained solution space. In contrast, the random dynamics are widely scattered. Figure 4B quantifies this, showing that v_θ has a much higher cosine similarity to the ground truth than the baselines.

Table 1 provides a detailed quantitative analysis. The learned dynamics v_θ exhibit exceptionally high self-similarity (0.99 ± 0.01), confirming that our training procedure consistently finds nearly identical solutions. They also achieve the highest fidelity to the Lorenz system (0.25 ± 0.01). Counter-intuitively, the training-free method, which theoretically samples from a constrained set of solutions, shows significantly lower self-similarity (0.50 ± 0.29) and fidelity (0.13 ± 0.08), indicating a wider practical solution space than the neural network approach.

Table 1: Average cosine similarities between dynamics representations for the Lorenz system. The trained models (v_θ) show the highest self-similarity (uniqueness) and the highest similarity to the Lorenz dynamics (fidelity).

Similarity	Learned (v_θ)	Training-free	Random	Lorenz	Score (s)
Learned (v_θ)	0.99 ± 0.01	0.37 ± 0.20	0.08 ± 0.06	0.25 ± 0.01	0.01 ± 0.004
Training-free	—	0.50 ± 0.29	0.12 ± 0.08	0.13 ± 0.08	0.02 ± 0.01
Random	—	—	0.38 ± 0.24	0.04 ± 0.04	0.09 ± 0.07
Lorenz	—	—	—	1.00	0.017
Score (s)	—	—	—	—	1.00

Table 2: Absolute cosine similarity to ground truth for different methods. We compare our training-free method against a modified training-free method with random K rather than skew-symmetric ones, and a re-weighted neural network using s as the dynamics function. Similarities are shown with standard deviations in brackets. Training-free methods do not show standard deviations since there are only two possible K matrices for two channels. Our training-free method achieves the highest similarities across all patterns studied.

	Life	Waves	Spirals	Maze
Training-free	0.271	0.589	0.250	0.270
Random K	0.021 (0.022)	0.182 (0.221)	0.083 (0.090)	0.072 (0.104)
Random Re-weight	0.052 (0.032)	0.105 (0.044)	0.180 (0.106)	0.031 (0.022)

This surprising result suggests that the uniqueness observed in the trained models is largely attributable to the *inductive biases* of neural networks, such as smoothness (Rahaman et al., 2019) and a preference for simple functions (Valle-Perez et al., 2019). The training-free method, lacking these biases, may be more sensitive to noise in the score function. The skew-symmetric matrix transformation, which computes weighted differences between state variables, can amplify this noise, especially when dimensions are correlated, leading to a more varied set of solutions.

For the high-dimensional Turing patterns, we assess the fidelity of the dynamics recovered by the training-free method. Due to the stochastic Langevin term in the dynamics, we define the representation as the net change in the pattern after a short evolution from a single initial state x_0 to x_T : $e_v = \text{flatten}(x_T - x_0)$. The similarities reported in Table 2 show that for each pattern type, the recovered dynamics is most similar to its corresponding ground-truth dynamics compared to random sampled dynamics. Importantly, both the training-free method and random dynamics use the same $\eta = 0.1$ Langevin term, suggesting that the similarity is primarily due to the learned dynamics rather than the score function. This confirms that the method captures pattern-specific dynamics.

Interestingly, these results also support our hypothesis about inductive bias. The training-free method for these patterns used a 1×1 convolution to implement the skew-symmetric transformation. This structure forces the dynamics at each pixel to be a linear combination of the score function values at that same pixel. Consequently, the recovered dynamics inherits the strong inductive biases of the score function’s convolutional architecture, such as locality and translation invariance, which contributes to the successful recovery of structured dynamics.

The added Langevin dynamics also plays a crucial role. Although the learned dynamics may deviate from a specific ground truth, the score function acts as a corrective force, pulling the system back toward the data manifold. This stabilization is particularly important in sparse distributions, where the dynamics are constrained to move within the narrow “tunnels” of the manifold.




Figure 5: Artificial life forms created via our inverse design approach. We first define a target pattern (top-left corners of each panel) and create a distribution by scattering it with random positions and rotations. Our method then finds the dynamics that preserve this distribution. Trajectories are highlighted with green lines. More complex movements are visualized by using an exponential moving average (EMA). Several systems display emergent properties, such as collective flocking behavior (see “firefly”) and pattern-based attraction (see “blue-red dots”).

432 This ability to robustly find dynamics that preserve a given pattern distribution opens up new appli-
433 cations beyond recovering existing systems. It is particularly useful for designing new dynamics,
434 such as *Artificial Life*, a field that studies not only what life is, but *what life could be* (Langton, 2019).
435 Current methods for creating artificial life often rely on discovery through brute-force search (Chan,
436 2018) or evolution with genetic algorithms (Reinke et al., 2020; Kumar et al., 2025). While power-
437 ful, these approaches often require extensive trial and error and offer limited direct control over the
438 final forms and behaviors.

439 Our method, in contrast, offers a new paradigm for designing such systems through *inverse design*.
440 We can specify a desired pattern, a photograph or a simple hand drawing, and discover the local
441 interaction rules that sustain its distribution. To do this, we scatter instances of the pattern randomly
442 in a 2D space and apply our method to find the dynamics that preserve it. As shown in Figure 5,
443 this process successfully generates artificial life forms that match the desired patterns. The resulting
444 dynamics link a pattern’s morphology to its motile behavior. For instance, the highly asymmetric
445 “firefly” pattern leads to smooth movement, while patterns with greater symmetry like “taiji” and
446 “blue dots” have slower motion. Our method thus provides a tool to study the relationship between
447 form and emergent behavior. These systems also exhibit complex *emergent properties* that were not
448 explicitly programmed, such as collective flocking reminiscent of the boids model and attraction or
449 repulsion behaviors.

451 5 DISCUSSION

452
453 In this paper, we addressed the fundamental, yet often intractable, problem of inferring process
454 from pattern. Rather than attempting the uncomputable task of finding dynamics for a single static
455 pattern, we posed a more tractable problem: discovering continuous dynamics that preserve an en-
456 tire distribution of patterns. To this end, we introduced the *Equilibrium flow* method, a framework
457 that successfully recovers plausible dynamics for systems ranging from simple 2D distributions to
458 the chaotic Lorenz attractor and high-dimensional Turing patterns. Our results consistently demon-
459 strate that the learned dynamics exhibit high self-similarity across independent trials and significant
460 fidelity to the ground truth, suggesting the solution space is highly constrained.

461 Our central finding is that this constraint arises from the combined influence of data sparsity and
462 model inductive bias. High-dimensional patterns from natural systems often lie on low-dimensional
463 manifolds, which inherently restrict the possible trajectories of evolution. The architectural and
464 optimization biases of our learning framework further narrow the space of potential solutions. This
465 dual constraint enables a new paradigm for studying the relationship between patterns and dynamics,
466 with applications in both scientific discovery and creative design. For recovering dynamics, our
467 method opens promising avenues in fields like systems biology and cosmology, where temporal
468 data is scarce but snapshot data is abundant. For designing dynamics, the goal shifts from finding a
469 single ground truth to discovering a set of plausible rules. Our work on Artificial Life exemplifies
470 this, where specifying a target morphology allows us to inverse-design the local interaction rules
471 that sustain it and give rise to emergent collective behaviors.

472 This work also opens numerous avenues for future research. While we have qualitatively shown that
473 the solution space is constrained, a key theoretical direction is to quantify precisely how properties of
474 the data distribution, such as its sparsity, symmetries, and topology, affect the uniqueness and behav-
475 ior of the learned dynamics. A deeper understanding here could forge a more tangible, computable
476 link to the problem’s abstract framing in terms of Kolmogorov complexity. Methodologically, fur-
477 ther investigation is needed to understand the boundaries of our approach and improve its efficiency.
478 The current two-stage process, which relies on a pre-trained score function, could be replaced by
479 more efficient one-stage, end-to-end models. Exploring alternatives to Hutchinson’s Trace Estimator
480 or developing ways to deliberately impose beneficial inductive biases on the training-free method
481 are also critical next steps.

482 In conclusion, we offer a dual contribution. First, we have introduced a novel computational tool
483 for investigating the fundamental relationship between static patterns and the dynamic processes that
484 form them. Second, we provide a practical methodology for inferring temporal evolution from snap-
485 shot data, a common challenge across numerous scientific and engineering disciplines. By bridging
the gap between static observation and dynamic understanding, we have laid the computational
groundwork for a future theory of how process is encoded in pattern.

6 REPRODUCIBILITY STATEMENT

We have included the code for our experiments in the supplementary material. To ensure reproducibility, all data for the Turing patterns discussed in Section 3 are generated using programs with fixed random seeds. For the Artificial Life experiments discussed in Section 4, we have included the original PNG files of the target patterns in the supplementary material. The generated Artificial Life videos are also provided in the supplementary material.

REFERENCES

- Marloes Arts, Victor Garcia Satorras, Chin-Wei Huang, Daniel Zugner, Marco Federici, Cecilia Clementi, Frank Noé, Robert Pinsler, and Rianne van den Berg. Two for one: Diffusion models and force fields for coarse-grained molecular dynamics. *Journal of Chemical Theory and Computation*, 19(18):6151–6159, 2023.
- Carine Babusiaux, F Van Leeuwen, Martin Adrian Barstow, C Jordi, Antonella Vallenari, D Bossini, Alessandro Bressan, T Cantat-Gaudin, M Van Leeuwen, Anthony GA Brown, et al. Gaia data release 2-observational hertzsprung-russell diagrams. *Astronomy & Astrophysics*, 616:A10, 2018.
- Jonah Botvinick-Greenhouse, Robert Martin, and Yunan Yang. Learning dynamics on invariant measures using pde-constrained optimization. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 33(6), 2023.
- Gregory J Chaitin. A theory of program size formally identical to information theory. *Journal of the ACM (JACM)*, 22(3):329–340, 1975.
- Bert Wang-Chak Chan. Lenia-biology of artificial life. *arXiv preprint arXiv:1812.05433*, 2018.
- Victor Chardès, Suryanarayana Maddu, and Michael J Shelley. Stochastic force inference via density estimation. *arXiv preprint arXiv:2310.02366*, 2023.
- Ricky TQ Chen, Jens Behrmann, David K Duvenaud, and Jörn-Henrik Jacobsen. Residual flows for invertible generative modeling. *Advances in Neural Information Processing Systems*, 32, 2019.
- Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.
- Stefan Elfving, Eiji Uchibe, and Kenji Doya. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural networks*, 107:3–11, 2018.
- P. Gray and S.K. Scott. Autocatalytic reactions in the isothermal, continuous stirred tank reactor: Oscillations and instabilities in the system $a + 2b \rightarrow 3b$; $b \rightarrow c$. *Chemical Engineering Science*, 39(6):1087–1097, 1984. ISSN 0009-2509. doi: [https://doi.org/10.1016/0009-2509\(84\)87017-7](https://doi.org/10.1016/0009-2509(84)87017-7). URL <https://www.sciencedirect.com/science/article/pii/0009250984870177>.
- Arthur Gretton, Karsten Borgwardt, Malte Rasch, Bernhard Schölkopf, and Alex Smola. A kernel method for the two-sample-problem. *Advances in neural information processing systems*, 19, 2006.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Michael F Hutchinson. A stochastic estimator of the trace of the influence matrix for laplacian smoothing splines. *Communications in Statistics-Simulation and Computation*, 18(3):1059–1076, 1989.
- Chii-Ruey Hwang, Shu-Yin Hwang-Ma, and Shuenn-Jyi Sheu. Accelerating diffusions. *The Annals of Applied Probability*, 15(2), May 2005. ISSN 1050-5164. doi: [10.1214/105051605000000025](https://doi.org/10.1214/105051605000000025). URL <http://dx.doi.org/10.1214/105051605000000025>.

- 540 Akarsh Kumar, Chris Lu, Louis Kirsch, Yujin Tang, Kenneth O Stanley, Phillip Isola, and David Ha.
541 Automating the search for artificial life with foundation models. *Artificial Life*, 31(3):368–396,
542 2025.
- 543 Christopher G Langton. Artificial life. In *Artificial life*, pp. 1–47. Routledge, 2019.
- 544 Qian Li. sctour: a deep learning architecture for robust inference and accurate prediction of cellular
545 dynamics. *Genome Biology*, 24(1):149, 2023.
- 546 Qiang Liu, Jason Lee, and Michael Jordan. A kernelized stein discrepancy for goodness-of-fit tests.
547 In *International conference on machine learning*, pp. 276–284. PMLR, 2016.
- 550 Daniel Lobo, Mauricio Solano, George A Bubenik, and Michael Levin. A linear-encoding model
551 explains the variability of the target morphology in regeneration. *Journal of The Royal Society*
552 *Interface*, 11(92):20130918, 2014.
- 553 Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and
554 projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.
- 555 Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and
556 Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications*
557 *of the ACM*, 65(1):99–106, 2021.
- 558 Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models.
559 In *International conference on machine learning*, pp. 8162–8171. PMLR, 2021.
- 560 Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua
561 Bengio, and Aaron Courville. On the spectral bias of neural networks. In *International conference*
562 *on machine learning*, pp. 5301–5310. PMLR, 2019.
- 563 Prajit Ramachandran, Barret Zoph, and Quoc V Le. Swish: a self-gated activation function. *arXiv*
564 *preprint arXiv:1710.05941*, 7(1):5, 2017.
- 565 Chris Reinke, Mayalen Etcheverry, and Pierre-Yves Oudeyer. Intrinsically motivated discovery of
566 diverse patterns in self-organizing systems. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=rkg6sJHYDr>.
- 567 Luc Rey-Bellet and Konstantinos Spiliopoulos. Irreversible langevin samplers and variance reduction:
568 a large deviations approach. *Nonlinearity*, 28(7):2081, 2015.
- 569 Henry Norris Russell. Relations between the spectra and other characteristics of stars. In *A Source*
570 *Book in Astronomy and Astrophysics, 1900–1975*, pp. 212–220. Harvard University Press, 1979.
- 571 Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In
572 *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=TIIdIXIpzhoI>.
- 573 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv*
574 *preprint arXiv:2010.02502*, 2020.
- 575 Alexander Tong, Jessie Huang, Guy Wolf, David Van Dijk, and Smita Krishnaswamy. Trajectory
576 Net: A dynamic optimal transport network for modeling cellular dynamics. In Hal Daumé III
577 and Aarti Singh (eds.), *Proceedings of the 37th International Conference on Machine Learning*,
578 volume 119 of *Proceedings of Machine Learning Research*, pp. 9526–9536. PMLR, 13–18 Jul
579 2020. URL <https://proceedings.mlr.press/v119/tong20a.html>.
- 580 AM Turing. The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society*
581 *of London. Series B, Biological Sciences*, pp. 37–72, 1952.
- 582 Guillermo Valle-Perez, Chico Q. Camargo, and Ard A. Louis. Deep learning generalizes because
583 the parameter-function map is biased towards simple functions. In *International Confer-*
584 *ence on Learning Representations*, 2019. URL <https://openreview.net/forum?id=rye4g3AqFm>.

594 Yunan Yang, Levon Nurbekyan, Elisa Negrini, Robert Martin, and Mirjeta Pasha. Optimal transport
595 for parameter identification of chaotic dynamics via invariant measures. *SIAM Journal on Applied*
596 *Dynamical Systems*, 22(1):269–310, 2023.
597
598 Zhenyi Zhang, Tiejun Li, and Peijie Zhou. Learning stochastic dynamics from snapshots through
599 regularized unbalanced optimal transport. In *The Thirteenth International Conference on Learn-*
600 *ing Representations*, 2025. URL <https://openreview.net/forum?id=gQlxd3Mtru>.
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647

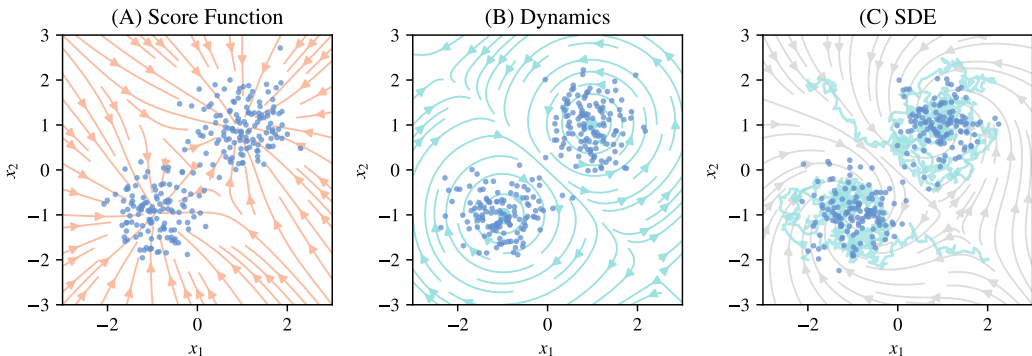


Figure 6: **(A)** The estimated score function (orange streams) trained from sampled data (blue dots) using diffusion model (with $\tau = 0.2$). **(B)** The learned dynamics (blue streams) inferred from the estimated score function and sampled data (blue dots). **(C)** The learned dynamics can be combined with the score function to generate stochastic dynamics ($\eta = 0.25$), shown as blue lines.

A NEURAL NETWORK ARCHITECTURE

In this section of appendix, we describe the neural network architectures used in our experiments. We first outline the structure of our dynamic model, which infers the dynamics from statical patterns, followed by details of the score function model that estimates the gradient of the log probability density.

A.1 DYNAMIC MODEL FOR 2-D AND LORENZ SYSTEMS

We parameterize the dynamics as ordinary differential equations (ODEs):

$$\frac{d\mathbf{x}}{dt} = \mathbf{v}_\theta(\mathbf{x}). \quad (14)$$

where \mathbf{v}_θ is implemented using a multilayer perceptron (MLP) neural network that maps d -dimensional input vectors to d -dimensional output vectors.

To enhance the network’s ability to capture high-frequency dynamics, we incorporate positional encoding:

$$\mathbf{v}_\theta(\mathbf{x}) = f_\theta(\mathbf{x}, \gamma^n(\mathbf{x})), \quad (15)$$

where the positional encoding function γ^n ($n = 4$ by default) from (Mildenhall et al., 2021) expands the input dimensions as follows:

$$\gamma^n(\mathbf{x}) = [\sin(2^0 \mathbf{x}), \cos(2^0 \mathbf{x}), \sin(2^1 \mathbf{x}), \cos(2^1 \mathbf{x}), \dots, \sin(2^n \mathbf{x}), \cos(2^n \mathbf{x})]. \quad (16)$$

For the activation function in f_θ , we use SiLU (Ramachandran et al., 2017; Elfving et al., 2018) instead of ReLU. This choice is critical because our gradient descent process involves vector-Jacobian products that require non-zero second-order gradients (Chen et al., 2019).

A.2 SCORE FUNCTION FOR GRAY-SCOTT MODELS

We adopt a minimal 2D U-Net, based on the `diffusers` library’s `UNet2DModel`, to learn the score function for the Turing patterns. The network receives a two-channel 128×128 grid, representing the chemical concentrations, and outputs a two-channel grid of the same size.

To better capture the system’s dynamics, we customized the model to use circular padding, which respects the periodic boundary conditions of the Gray-Scott simulation. The specific architecture, detailed in Listing 1, is a lightweight U-Net with a symmetric encoder-decoder structure. The encoder consists of three `DownBlock2D` blocks that progressively downsample the input, while the number of feature channels increases from 16 to 32, and finally to 64. The decoder mirrors this

structure with three `UpBlock2D` blocks to reconstruct the output. This streamlined design reduces computational overhead while effectively modeling the spatial features of the patterns.

```

702
703
704
705 def get_unet_diffusion():
706     unet = UNet2DModelWithPadding(
707         sample_size=128,
708         in_channels=2,
709         out_channels=2,
710         down_block_types=
711             ("DownBlock2D", "DownBlock2D", "DownBlock2D"),
712         up_block_types=
713             ("UpBlock2D", "UpBlock2D", "UpBlock2D"),
714         block_out_channels=(16, 32, 64),
715         norm_num_groups=8,
716         padding_mode='circular',
717         only_when_effective=True,
718         log_changed=True
719     )
720     return unet

```

Listing 1: Python code for configuring the 2D U-Net model. We customize the model to use circular padding, which respects the periodic boundary conditions of the Gray-Scott simulation.

A.3 SCORE FUNCTION ESTIMATION

A.3.1 DIFFUSION MODELS

To estimate the score function $s(\mathbf{x}) = \nabla \log p(\mathbf{x})$ for data sampled from $p(\mathbf{x})$, we utilized diffusion models (Ho et al., 2020; Song et al., 2020). These models work by training neural networks to recover the original data from noisy versions.

The training process involves minimizing the following loss function:

$$L = \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x}), \tau \sim \mathcal{U}(0,1), \epsilon \sim \mathcal{N}(0, I_d)} \|\epsilon_\theta(\sqrt{\alpha_\tau} \mathbf{x} + \sqrt{1 - \alpha_\tau} \epsilon, \tau) - \epsilon\|^2 \quad (17)$$

In this formulation, $\mathbf{x}_\tau = \sqrt{\alpha_\tau} \mathbf{x} + \sqrt{1 - \alpha_\tau} \epsilon$ represents a noisy version of the original data \mathbf{x} , where α_τ controls the noise level based on τ . The noise parameter τ interpolates between the original distribution ($\tau = 0$, giving $\mathbf{x}_0 \sim p(\mathbf{x})$) and pure Gaussian noise ($\tau = 1$, giving $\mathbf{x}_1 \sim \mathcal{N}(0, I^d)$) (Nichol & Dhariwal, 2021).

Once trained, the model ϵ_θ can predict the noise component added to the original data. From this noise prediction, we can derive the score function s of $p(\mathbf{x}_\tau)$ (Dhariwal & Nichol, 2021) using:

$$s(\mathbf{x}_\tau) = -\frac{1}{\sqrt{1 - \alpha_\tau}} \epsilon_\theta(\mathbf{x}_\tau, \tau). \quad (18)$$

For practical implementation, we set $\tau > 0$ rather than $\tau = 0$: at $\tau = 0$, the model would be untrained, while high τ values would deviate significantly from the original distribution. Our experiments show that $\alpha_\tau \approx 0.95$ provides good performance across most scenarios, which corresponds to $\tau = 0.2$ in 2D and Lorenz systems, and $\tau = 0.1$ for Gray-Scott patterns.

A.3.2 v-PREDICTION

While the original diffusion model training objective works well on toy models, it shows limitations on larger, more complex tasks. Specifically, the estimated score functions are not accurate enough for downstream applications. To address this, we adopt the v -prediction method (Salimans & Ho, 2022). Instead of training a model to predict ϵ , which is the noise added to the data, we train it to predict the “velocity” defined as:

$$\mathbf{v} = \sqrt{\alpha_\tau} \epsilon - \sqrt{1 - \alpha_\tau} \mathbf{x}_0. \quad (19)$$

The model v_θ is trained to predict this velocity instead of the noise directly:

$$L = \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x}), \tau \sim \mathcal{U}(0,1), \epsilon \sim \mathcal{N}(0, \mathbf{I}_d)} \|\mathbf{v}_\theta(\mathbf{x}_\tau, \tau) - \mathbf{v}\|^2. \quad (20)$$

This approach can significantly accelerate the training process and achieve higher accuracy in score estimation. Based on the design of diffusion models, we can derive:

$$\epsilon = \frac{\mathbf{x}_\tau - \sqrt{\alpha_\tau} \mathbf{x}_0}{\sqrt{1 - \alpha_\tau}}. \quad (21)$$

Hence, we can replace ϵ in equation 18 with the formula for \mathbf{x}_0 :

$$\mathbf{s}(\mathbf{x}_\tau) = -\frac{\epsilon}{\sqrt{1 - \alpha_\tau}} \quad (22)$$

$$= \frac{\mathbf{x}_0 \sqrt{\alpha_\tau} - \mathbf{x}_\tau}{1 - \alpha_\tau}. \quad (23)$$

Therefore, we can unify the score function estimation via the estimated \mathbf{x}_0 , which is independent of which prediction objective we use.

A.3.3 BATCH NORMALIZATION

Since the distribution $p(\mathbf{x})$ is invariant under the flow $\mathbf{v}(\mathbf{x})$, the average position $\langle \mathbf{x} \rangle$ is constant. The time evolution of the average position is given by the average velocity, which must therefore be zero:

$$\frac{d}{dt} \langle \mathbf{x} \rangle = \langle \mathbf{v} \rangle = \mathbf{0}. \quad (24)$$

Based on this property, we enforce this zero-mean condition on our neural network \mathbf{v}_θ by applying a batch normalization layer to its output. A standard batch normalization layer includes a learnable affine transformation that could produce the trivial solution $\mathbf{v} \equiv \mathbf{0}$ (e.g., by learning a zero scaling factor). To prevent this, we use a layer that only performs normalization. For a mini-batch of raw network outputs $\{\hat{\mathbf{v}}(\mathbf{x}_k)\}$, the final output $\mathbf{v}(\mathbf{x}_k)$ for each input \mathbf{x}_k is computed as:

$$\mathbf{v}(\mathbf{x}_k) = \frac{\hat{\mathbf{v}}(\mathbf{x}_k) - \mathbb{E}_{\text{batch}}[\hat{\mathbf{v}}]}{\sqrt{\text{Var}_{\text{batch}}[\hat{\mathbf{v}}] + \epsilon}}, \quad (25)$$

where $\mathbb{E}_{\text{batch}}[\hat{\mathbf{v}}]$ and $\text{Var}_{\text{batch}}[\hat{\mathbf{v}}]$ are the mean and variance calculated over the current mini-batch. During inference, these are replaced by their corresponding running averages maintained throughout the training process.

B LONG-TERM DISTRIBUTION PRESERVATION

To validate our method’s ability to preserve distributions over extended time periods, we evaluate distribution preservation over time on the Lorenz system. We simulate the dynamics for 100 steps with step size $\Delta t = 0.1$ using the Runge-Kutta method and compute the Maximum Mean Discrepancy (MMD) (Gretton et al., 2006) between the initial distribution and distributions at different time steps.

As shown in Figure 7, our methods (both learned dynamics and the training-free approach) exhibit significantly lower MMD values, demonstrating that our methods can effectively preserve the distribution over long time horizons.

C TURING PATTERNS

We use Gray-Scott model (Gray & Scott, 1984) to generate Turing patterns. The model is defined by the following reaction-diffusion equations:

$$\begin{aligned} \frac{\partial \mathbf{u}}{\partial t} &= D_u \nabla^2 \mathbf{u} - \mathbf{u} \mathbf{v}^2 + F(1 - \mathbf{u}), \\ \frac{\partial \mathbf{v}}{\partial t} &= D_v \nabla^2 \mathbf{v} + \mathbf{u} \mathbf{v}^2 - (F + k) \mathbf{v}. \end{aligned} \quad (26)$$

810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863




Figure 7: Maximum Mean Discrepancy (MMD) evolution between the initial distribution and distributions at different time steps. The MMD is computed using a Radial Basis Function (RBF) kernel with bandwidth $\sigma = 0.5$. Blue and orange represent the learned dynamics and training-free method dynamics, respectively. Compared to random dynamics (gray) and the score function approach (light blue), our methods exhibit significantly lower MMD values, demonstrating superior long-term distribution preservation.

By changing the hyperparameters D_u , D_v , F , and k , we can generate a variety of Turing patterns. Here, we select three sets of parameters that produce distinct patterns, with $D_u = 0.16$, $D_v = 0.08$, and other parameters varying:

- “life”: Life-like pattern: $F = 0.006$, $k = 0.045$.
- “wave”: Wave pattern: $F = 0.018$, $k = 0.049$.
- “spirals”: Spiral pattern: $F = 0.007$, $k = 0.028$.
- “maze”: Static pattern: $F = 0.029$, $k = 0.057$.

D TRAINING-FREE METHOD

D.1 SKEW-SYMMETRIC SOLUTION

Theorem 1. *The vector field $\mathbf{v}_S(\mathbf{x}) = \mathbf{S}\nabla \log p(\mathbf{x})$, constructed with any constant skew-symmetric matrix \mathbf{S} ($\mathbf{S}^\top = -\mathbf{S}$), is a solution to the steady-state continuity equation, $\nabla \cdot [p(\mathbf{x})\mathbf{v}(\mathbf{x})] = 0$.*

Proof. $\nabla \cdot [p(\mathbf{x})\mathbf{v}(\mathbf{x})] = 0$ can be simplified by noting that $\nabla p(\mathbf{x}) = p(\mathbf{x})\nabla \log p(\mathbf{x})$, which leads to:

$$\nabla \cdot [p(\mathbf{x})\mathbf{S}\nabla \log p(\mathbf{x})] = \nabla \cdot \left[p(\mathbf{x})\mathbf{S} \frac{\nabla p(\mathbf{x})}{p(\mathbf{x})} \right] = \nabla \cdot [\mathbf{S}\nabla p(\mathbf{x})].$$

Using Einstein notation, this divergence is a contraction between \mathbf{S} and the Hessian of $p(\mathbf{x})$. This term vanishes due to the interplay between the skew-symmetry of \mathbf{S} (i.e., $S_{ij} = -S_{ji}$) and the symmetry of the second partial derivatives of $p(\mathbf{x})$ (i.e., $\partial_i \partial_j p = \partial_j \partial_i p$). The derivation is as follows:

$$\begin{aligned} \nabla \cdot [\mathbf{S}\nabla p(\mathbf{x})] &= S_{ij} \partial_i \partial_j p(\mathbf{x}) \\ &= \frac{1}{2} (S_{ij} \partial_i \partial_j p(\mathbf{x}) + S_{ji} \partial_j \partial_i p(\mathbf{x})) \\ &= \frac{1}{2} (S_{ij} \partial_i \partial_j p(\mathbf{x}) - S_{ij} \partial_i \partial_j p(\mathbf{x})) = 0. \end{aligned}$$

In the second line, we rewrite the term by averaging it with a copy where the indices (i, j) have been swapped. In the third line, we substitute the aforementioned properties. \square




Figure 8: An example of a kernels with $c = 3$ channel, 5×5 convolutional layer satisfying the skew-symmetry condition. Note how the off-diagonal kernel K_{12} is the negated and spatially flipped version of K_{21} , and the diagonal kernel K_{22} is point-symmetric with negation (e.g., the top-left value is the negative of the bottom-right value).

D.2 SKEW-SYMMETRIC CONVOLUTIONAL LAYERS

Theorem 2. A convolutional layer with kernel $K \in \mathbb{R}^{c \times c \times (2r+1) \times (2r+1)}$ corresponds to a skew-symmetric matrix if and only if its weight tensors satisfy the condition

$$K_{o,i,u,v} = -K_{i,o,-u,-v} \quad \forall o, i \in \{1, \dots, c\}, \quad \forall (u, v) \in \{-r, \dots, r\}^2. \quad (27)$$

Proof. A linear transformation on a c -channel image of size $h \times w$ can be represented by a large block matrix $S \in \mathbb{R}^{(chw) \times (chw)}$ with $c \times c$ blocks. When this transformation is a standard convolution with a kernel $K \in \mathbb{R}^{c \times c \times (2r+1) \times (2r+1)}$, the matrix S is a block matrix given by:

$$S = \begin{bmatrix} C(K_{1,1}) & C(K_{1,2}) & \cdots & C(K_{1,c}) \\ C(K_{2,1}) & C(K_{2,2}) & \cdots & C(K_{2,c}) \\ \vdots & \vdots & \ddots & \vdots \\ C(K_{c,1}) & C(K_{c,2}) & \cdots & C(K_{c,c}) \end{bmatrix}, \quad (28)$$

where $K_{o,i} \in \mathbb{R}^{(2r+1) \times (2r+1)}$ is the kernel slice that maps the i -th input channel to the o -th output channel, and $C(K_{o,i}) \in \mathbb{R}^{(hw) \times (hw)}$ is the corresponding matrix for this single-channel convolution.

For the transformation S to be skew-symmetric, it must satisfy the condition $S = -S^\top$. Since the transpose of a block matrix is given by $(S^\top)_{o,i} = S_{i,o}^\top$, the skew-symmetry condition requires that each block satisfies:

$$C(K_{o,i}) = -C(K_{i,o})^\top \quad \forall o, i \in \{1, \dots, c\}. \quad (29)$$

To translate this matrix equation into a direct constraint on the kernel weights, we examine the structure of $C(K)$. Under standard convolution definitions, the element of $C(K_{o,i})$ corresponding to the influence of an input at spatial location q on an output at location p is given by the kernel value at the displacement $p - q$. That is, $[C(K_{o,i})]_{\phi(p), \phi(q)} = (K_{o,i})_{p-q}$, where ϕ is an arbitrary bijection that maps the spatial location to the index of the matrix.

Consequently, the corresponding element of the transposed matrix is:

$$[C(K_{i,o})^\top]_{\phi(p), \phi(q)} = [C(K_{i,o})]_{\phi(q), \phi(p)} = (K_{i,o})_{q-p}. \quad (30)$$

918 Substituting these expressions into Equation 29 yields:

$$919 \quad (K_{o,i})_{p-q} = -(K_{i,o})_{q-p}.$$

921 By defining the kernel coordinate vector as $\mathbf{u} = p - q$, we have $-\mathbf{u} = q - p$. The condition on the
922 kernel weights thus simplifies to:

$$923 \quad (K_{o,i})_{\mathbf{u}} = -(K_{i,o})_{-\mathbf{u}} \quad (31)$$

924 for all kernel coordinates $\mathbf{u} \in \{-r, \dots, r\}^2$. In index notation, with $\mathbf{u} = (u, v)$, this is written as:

$$925 \quad K_{o,i,u,v} = -K_{i,o,-u,-v}. \quad (32)$$

928 This simple relation enforces skew-symmetry. For diagonal blocks ($i = o$), the kernel $K_{i,i}$ must
929 be point-symmetric with negation ($K_{i,i,u,v} = -K_{i,i,-u,-v}$). For off-diagonal blocks ($i \neq o$), the
930 kernel $K_{o,i}$ must be the negated and spatially flipped version of the kernel $K_{i,o}$. Figure 8 illustrates
931 an example of a kernel that satisfies this property. \square

933 E ARTIFICIAL LIFE

934 E.1 GENERATING TRAINING DATA

935 We generate training data by randomly scattering target patterns in a 128×128 2D space with
936 random rotations and positions. The space is circular to avoid boundary effects. To prevent pattern
937 overlap, we use PNG format with alpha channels to represent pattern transparency, employing a
938 trial-and-reject approach for pattern placement. We accept a pattern only if its alpha channel does
939 not overlap with existing patterns. For each 128×128 image, we attempt to add patterns multiple
940 times, with 2 placement attempts per pattern.

943 F THE USE OF LLM

944 We use LLMs to help us proofread the manuscript, including grammar correction and sentence
945 structure improvement. We also use LLMs to assist with writing code that can be easily verified.
946 For example, in the Artificial Life experiments, we use LLMs to write the code for generating
947 training data, which we can easily verify for correctness by inspecting both the results and the code.
948 LLMs also help us with well-defined mathematical problems. However, all proofs and mathematical
949 derivations are verified by the authors.