SIMULATE, DON'T INTERPOLATE: RECOVERING CONTINUOUS SPATIAL DYNAMICS FROM DISCRETE, UNPAIRED SNAPSHOTS

Anonymous authors
Paper under double-blind review

ABSTRACT

Spatiotemporal processes are central to many scientific questions, from how structure emerges in physical systems to how organisms form, or diseases develop. However, measuring such systems often relies on destructive experiments that preclude repeated time-series measurements, especially in biology and medicine. The alternative, unpaired snapshots, generally cannot fully determine the underlying dynamical processes. Recovering them requires constraining the solution space; implementing this at the level of the rules governing the system remains an open problem. We introduce REGIS[1,2] (REcovering Generative dynamics from Independent Snapshots), a framework that recovers spatial dynamics from time-labelled images by learning a unified update rule that is spatially local, Markovian, and time-invariant, and that is iterated to reach every observed time. We implement these constraints with neural cellular automata, trained adversarially by matching generated rollouts to the observed distributions. On synthetic data, REGIS outperforms state-of-the-art baselines in trajectory quality within and beyond the training horizon, data efficiency, and recovery of the underlying dynamics. Applied to a real biological system, a spatial atlas of zebrafish heart regeneration, REGIS recovers a rule that simulates regenerative dynamics. When used in an *in silico* functional screen, REGIS identifies critical cell types and hierarchical dependencies consistent with known biology, giving the learned rule a mechanistic interpretation. The result is a general recipe for building simulators from unpaired spatial snapshot data by sharing weights across space and time under a locality constraint. The resulting simulators are mechanistic and perturbable, offering a practical route to digital twins for biological systems.

1 INTRODUCTION

Many complex systems can only be understood through their change over time. Examples range from cell-fate decisions in biology (Purvis et al., 2012) to chemical pattern formation (Ouyang & Swinney, 1991) and grain growth in materials (McKenna et al., 2014). Modelling these systems requires understanding the processes that generate the observed states and govern their subsequent evolution. Yet complete trajectories are often difficult or impossible to observe. Destructive measurements, for example, provide detailed snapshots while preventing repeated observation of the same sample (Schiebinger et al., 2019; McKenna et al., 2014). These observations provide a fragmented view of the system: they describe its marginal distributions at different times but do not directly reveal individual trajectories. Our aim is to learn a dynamical process from these unpaired snapshots, so that we can simulate how the system evolves and investigate its behaviour, including beyond the observed times or under local perturbations.

Recovering dynamics from snapshots is an ill-posed problem. Distinct processes can produce the same observed distributions, so fitting these distributions alone cannot identify the underlying dynamics (Weinreb et al., 2018; Albergo et al., 2025). Existing approaches draw on optimal transport, continuous flows, and stochastic bridges to infer evolution from unpaired observations (Tong et al.,

[1]Code available at GitHub `https://github.com/anon03752/regis`
[2]Interactive demos available at `https://anon03752.github.io/regis-demo/`

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

2020; 2024; Gravier et al., 2025; Bunne et al., 2022). Within this landscape, flow matching provides a way to learn dynamics from chosen probability paths, typically relying on linear or spline-based interpolants (Lipman et al., 2023). Recent multi-marginal methods construct paths through several observed distributions (Rohbeck et al., 2025; Lee et al., 2025), while adversarially learned interpolants make the path construction itself trainable (Kviman et al., 2026). Although these models can reproduce the observed marginal distributions, the evolution connecting them is itself an object of scientific interest. Their utility for understanding how structures form, evolve, and respond to perturbations, thus, depends on whether the assumptions governing that evolution reflect the system being studied.

Rather than selecting a particular interpolant, we propose to constrain the space of possible dynamics using scientific priors on the underlying update rules. We focus on spatial processes where locality provides a natural prior: larger-scale effects arise through the propagation of local changes, as illustrated by tissue morphogenesis, collective cell migration, chemical pattern formation, and grain growth (Turing, 1952; Serra-Picamal et al., 2012; Ouyang & Swinney, 1991; McKenna et al., 2014). Additionally, we expect the same physical state to evolve by the same rule, regardless of when the change occurs. This motivates a time-independent update rule, unlike standard flow matching. We therefore restrict the model architecturally to local, Markovian, time-independent updates, requiring a shared rule to explain all observed time points through repeated application.

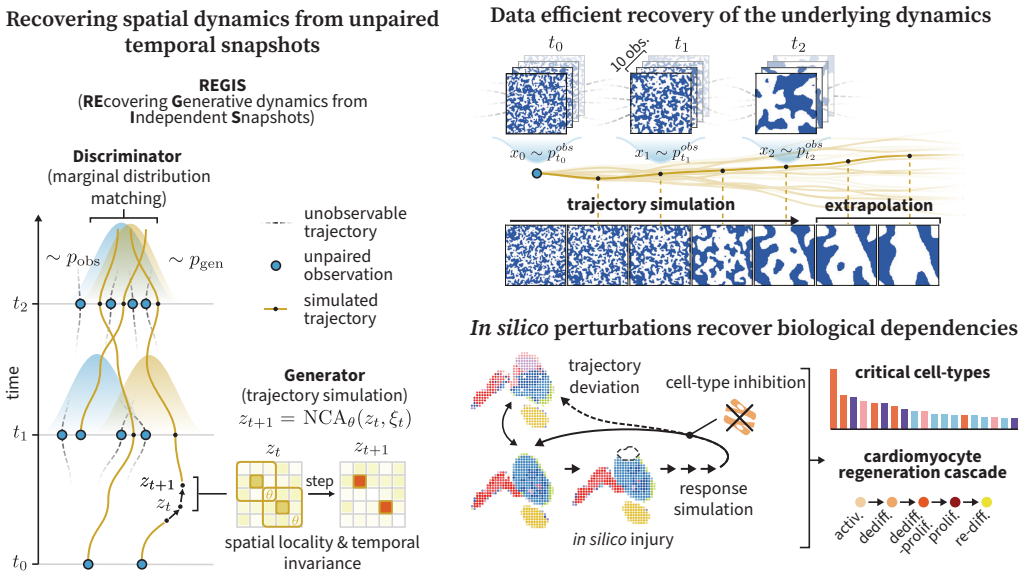

Figure 1: **REGIS learns spatial dynamics from unpaired snapshots.** *Left:* Observations at different times (blue points) provide marginal distributions, while individual trajectories remain unobserved (grey dashed curves). REGIS generates otherwise unobservable trajectories (gold) through repeated local updates, using a local and time-invariant rule. Adversarial training matches the generated distributions to the observations at the available times. *Top right:* Ising coarsening illustrates learning from limited data and simulating the process between and beyond the observed times. *Bottom right:* In zebrafish heart regeneration, simulated injuries and cell-type inhibition probe the learned dynamics, revealing dependencies consistent with the known biology.

We implement this idea using neural cellular automata (NCAs) (Mordvintsev et al., 2020), which learn spatially shared local rules, apply them repeatedly, and have been used to model growth and spatiotemporal patterns (Mordvintsev et al., 2020; Richardson et al., 2024), including generative settings (Palm et al., 2022; Otte et al., 2021; Luu et al., 2026). To train from unpaired observations, we use an adversarial objective to match the distributions produced by NCA to those observed at the corresponding times (Goodfellow et al., 2014). Unpaired image translation offers a useful precedent: CycleGAN and CUT similarly combine adversarial training with constraints on the learned mapping (Zhu et al., 2017; Park et al., 2020). Thus, adversarial training supplies supervision at the observed times, while the NCA architecture constrains the evolution between them (Figure 1, left).

The resulting model can extrapolate trajectories beyond the observation window, operate on spatial domains of different sizes, and continue evolving after local modifications of its state, allowing for interventional studies. Sharing an update rule across space and time also provides an inductive bias for learning from limited observations, improving the data efficiency of the method.

Our main contributions are as follows:

1. We introduce REGIS, a method for learning image dynamics from unpaired, time-labelled snapshots using neural cellular automata and adversarial training.
2. Using synthetic and real datasets, we show that REGIS outperforms state-of-the-art baselines in trajectory quality within and beyond the observation window, data efficiency, performance in low-data regimes, and recovery of the underlying dynamics.
3. We demonstrate that REGIS can be successfully applied to real-world experimental data, recovering both the cellular dynamics of zebrafish heart regeneration and known underlying biological mechanisms through *in silico* perturbation experiments of the trained simulator.

2 METHOD

2.1 PROBLEM DEFINITION

We are given a dataset $\mathcal{D} = \{(x_i, t_i)\}_{i=1}^{N}$ of time-labelled snapshots of a dynamical system. Each snapshot is represented by a multi-channel image $x_i \in \mathbb{R}^{H\times W\times C}$ and comes from a different realisation of the process. Let p_t denote the distribution of images at time t, so that $x_i \sim p_{t_i}$. The data provide samples from these marginal distributions, but do not provide correspondence between observations at different times.

As motivated previously, our aim is to learn a stochastic evolution rule whose repeated application generates trajectories consistent with these observations. To impose spatial locality and time independence, we use a modified NCA – a type of model that builds on the concept of a cellular automaton, where at each time step, the state of each cell is updated based on the states of its neighbours. In an NCA, such an update rule is parametrised by a neural network. We train it adversarially to match the image distributions at the observed times. The following subsections describe the architecture and its training procedure.

2.2 NEURAL CELLULAR AUTOMATON

At update t, the NCA state is

$$z_t = [x_t, h_t] \in \mathbb{R}^{H\times W\times (C+C_h)}, \tag{1}$$

where brackets denote concatenation along the channel dimension:

- **Visible channels** x_t: C channels representing the system at time t.
- **Hidden channels** h_t: C_h zero-initialised additional writeable channels that expand the model's state space (Mordvintsev et al., 2020).

The full state evolves according to

$$z_{t+1} = \mathrm{NCA}_\theta(z_t, \xi_t), \tag{2}$$

where $\mathrm{NCA}(\cdot)_\theta$ denotes the complete stochastic update and ξ_t is an optional noise input. Each update consists of three stages:

1. Perception. To gather the information about the neighbourhood of each pixel, a local depthwise convolution $\mathcal{K}_\theta$ is applied independently for each state channel, with the kernels shared between channels. The noise input ξ_t is concatenated to the resulting features to introduce randomness:

$$v_t = [\mathcal{K}_\theta(z_t), \xi_t]. \tag{3}$$

The filters may be learned or fixed; the original NCA uses fixed identity and Sobel filters (Mordvintsev et al., 2020). The noise may be sampled anew at each update or kept constant for the entire trajectory.

2. Increment prediction. A shared multilayer perceptron (MLP) acts independently at each pixel, mixing its channel features to predict an increment for the visible and hidden channels:

$$\Delta_t = \mathrm{MLP}_\theta(v_t). \tag{4}$$

3. Post-processing. The predicted increments are clipped and applied according to a firing mask. With random firing and an optional alive mask, the update is

$$z_{t+1} = A_t \odot \big(z_t + M_t \odot \mathrm{clip}(\Delta_t, -c, c)\big), \tag{5}$$

where $c > 0$ bounds the magnitude of each increment component and $\odot$ denotes elementwise multiplication. The firing mask has independent entries $M_t \sim \mathrm{Bernoulli}(p_{\mathrm{fire}}) \in \mathbb{R}^{H\times W}$, with one entry shared by all channels of a pixel. The alive mask A_t is computed by thresholding a neighbourhood maximum of a visible channel before the update and is used to prevent the updates from propagating outside of active regions. Either mask can be disabled. Experiment-specific inputs and update schedules are detailed in the appendix.

2.3 Adversarial distribution matching

We train the NCA to match the observed image distributions through an adversarial game with a time-conditioned convolutional discriminator $D_\phi(x, t)$. The discriminator compares observed and generated images associated with the same time point, without access to the NCA's hidden channels. Training consists of three steps:

1. Trajectory generation. Choose successive observation times $t_a < t_b$, expressed in NCA update units. Starting from z_{t_a}, repeatedly apply Equation (2) to generate a trajectory segment:

$$z_{t_a} \xrightarrow{t_b - t_a \text{ NCA updates}} z_{t_b} = [x_{t_b}, h_{t_b}]. \tag{6}$$

Starting states are initialised from observations or a known initial distribution, with hidden channels set to zero, or drawn from a pool of previously generated states (Mordvintsev et al., 2020). For temporal regularisation, we sample a readout time $\tilde{t}_b$ uniformly from the last δ updates of the segment; $\delta = 1$ gives a fixed endpoint.

2. Distribution comparison. Draw an observed image $x^{\mathrm{obs}}_{t_b} \sim p_{t_b}$ independently of the starting state. The discriminator scores this image and the generated visible channels $x_{\tilde{t}_b}$, both conditioned on the target time t_b:

$$d_{\mathrm{obs}} = D_\phi(x^{\mathrm{obs}}_{t_b}, t_b), \qquad d_{\mathrm{gen}} = D_\phi(x_{\tilde{t}_b}, t_b). \tag{7}$$

3. Adversarial optimisation. We alternate discriminator and NCA updates to minimise the regularised relativistic losses of R3GAN (Huang et al., 2024):

$$\mathcal{L}_D = \mathbb{E}[\mathrm{softplus}(d_{\mathrm{gen}} - d_{\mathrm{obs}})] + \frac{\gamma}{2}(R_1 + R_2), \qquad \mathcal{L}_{\mathrm{NCA}} = \mathbb{E}[\mathrm{softplus}(d_{\mathrm{obs}} - d_{\mathrm{gen}})], \tag{8}$$

where $\mathrm{softplus}(v) = \log(1 + e^v)$. The penalties R_1 and R_2 are the squared norms of the discriminator-input gradients on observed and generated samples, respectively (Mescheder et al., 2018); γ controls their strength. Problem-specific auxiliary losses may be added to $\mathcal{L}_{\mathrm{NCA}}$.

We backpropagate through the full trajectory segment or, for long segments, through a fixed window of its final updates. Following R3GAN, we maintain an exponential moving average (EMA) of the NCA parameters and use the averaged model for evaluation (Yazıcı et al., 2019).

3 Experiments

During the evaluation, we focus on answering these three questions, listed in the order of increasing complexity:

- Does REGIS reproduce the observed marginal distributions?
- Does the learned rule capture the underlying dynamics?
- Can the learned rule provide scientific insight?

No standard benchmark exists for spatial multi-marginal snapshot data, and commonly used non-spatial datasets, such as single-cell RNA-seq data (Schiebinger et al., 2019; Rohbeck et al., 2025; Klein et al., 2025; Yeo et al., 2021), do not reflect the challenges of high-dimensional, spatially resolved data. We therefore selected a diverse set of problems that illuminate different aspects of our questions of interest.

3.1 CYCLING MNIST

The easily recognisable structure of handwritten digits makes MNIST (LeCun et al., 1998) a useful setting for assessing the quality of generated images. Since the digits have no natural temporal evolution, we define a synthetic cycle, $0 \to 1 \to \cdots \to 9 \to 0$, and learn its transitions from one digit to the next. Closing the sequence into a cycle extends the sequential MNIST benchmark of Gravier et al. (2025), which terminates, into one that can be run indefinitely, letting us test whether repeated application of the learned rule preserves image quality over long trajectories.

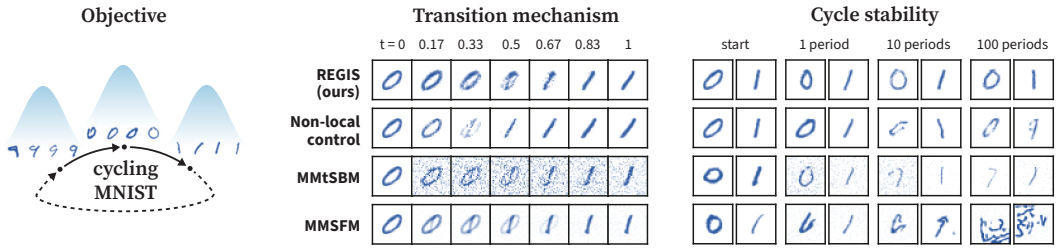

Figure 2: **REGIS produces stable, gradual digit transitions.** *Left:* The schematic illustration of the cycling MNIST objective. *Centre:* REGIS's transition between 0 and 1 is gradual and preserves the recognisable structure of the digit, while for other methods, the transition includes interpolation and background noise. *Right:* REGIS shows the most stable transition at long horizons.

We compare REGIS with a non-local control model that has a larger receptive field, trained with the same procedure. We also evaluate multi-marginal stochastic flow matching (MMSFM; Lee et al., 2025) and multi-marginal temporal Schrödinger bridge matching (MMtSBM; Gravier et al., 2025), adapted to the cyclic task (Appendix A.3). Starting from real test images, we evaluate the quality of generated trajectories up to 1000 consecutive digit transitions, using the Fréchet distance.

		Digit transitions completed				
Model	Parameters	1	5	10	100	1000
REGIS (ours)	88 K	**37.2** 28.9–54.4	**28.1** 21.9–44.4	**45.7** 28.8–49.9	**20.8** 19.1–25.3	**21.4** 18.9–24.7
Non-local control	83 K	53.9 39.9–81.1	131 55.7–270	208 161–294	1,304† 1,017–1,575	1,512† 1,376–2,533
MMSFM	18.6 M	81.0 42.6–122	159 86.7–215	46.6 26.9–66.6	262† 160–1,272	34,929† 5,498–64,359
MMtSBM	3.6 M	285 195–429	303 204–2,059	848 184–6,177	1,596 176–64,561	1,852† 176–55,284

Table 1: **Image quality over repeated digit transitions.** Fréchet distance in a frozen MNIST classifier's feature space (lower is better), using 5000 generated and 5000 real samples at each horizon. Cells: median over seven seeds, with the range across seeds beneath. REGIS performs best at all horizons and is the only model that maintains high image quality at 1000 transitions. †: at least one seed diverged and is excluded from the median and the range. Per-seed results appear in Table 2.

Table 1 shows that separation between REGIS and the non-local control grows over time: the distances are similar after one transition (37.2 versus 53.9), but reach 21.4 versus 1512 after 1000 transitions. The transport baselines (MMSFM and MMtSBM) also show deterioration or divergence in long rollouts. At horizons spanning more than 100 transitions, REGIS is the only model that maintains high image quality while reproducing the correct digit sequence (Figures 2, 7, 9).

Aside from the differences in long-term stability, the methods also take distinct approaches to transitioning between digits, as illustrated in Figure 2 by the example $0 \to 1$ transition. REGIS reshapes the initial stroke, whereas the non-local control largely fades out the zero as the one appears. MMSFM merely interpolates between the two digits, while the MMtSBM examples contain pronounced background noise. For MNIST, there is no natural transition between digits, but we argue that the trajectory produced by REGIS, reusing and gradually deforming the existing structures rather than noising-denoising or linearly interpolating, more closely resembles real-world processes, especially in biology.

3.2 ISING MODEL

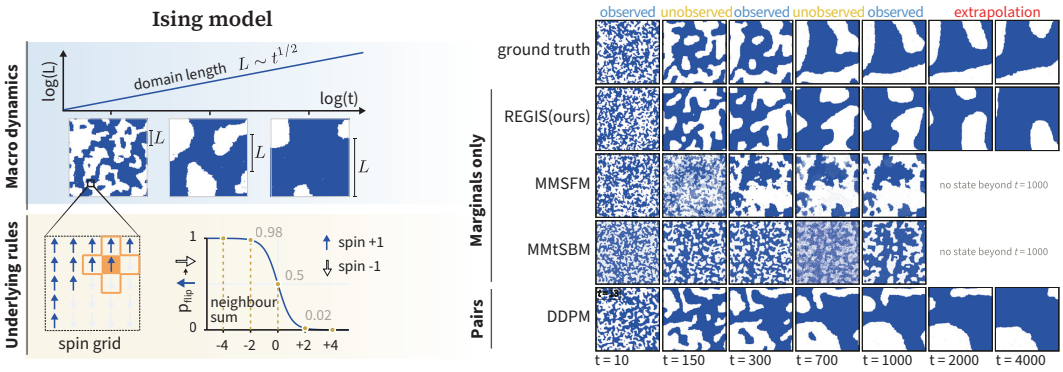

Figure 3: **Ising rollouts from one shared start.** *Left:* A schematic illustration of the Ising model. *Right:* Every row evolves the same random initial configuration at 512^2; columns are simulation times t: observed (training marginals), unobserved (between them) and extrapolation (beyond the last marginal, $t = 1000$). The transport models (MMSFM and MMtSBM) have no state past their last marginal. REGIS and DDPM show the gradual coarsening of the domains along the entire trajectory, while other models experience growth arrest. More examples are in Figure 14.

The MNIST experiment demonstrates that REGIS can reproduce a sequence of marginal distributions. However, no natural process governs transitions between digit images, so we cannot assess how closely the learned dynamics resemble a real-world process. We therefore turn to the Ising model, a widely studied model of ferromagnetism in statistical physics, which is well characterised (Bray, 1994) and provides a reference against which to assess whether REGIS reproduces the system's true evolution. We consider a two-dimensional grid of spins, each having one of two possible orientations. At each simulation step, each spin can flip with a probability determined by its alignment with its neighbours (Glauber, 1963), resulting in the gradual emergence of connected regions of the same orientation (see Figure 3, *left*). Appendix B.1 provides further details of the Ising model and experimental setup.

We generate independent realisations of the Ising model at $t \in \{3, 10, 17, 300, 1000\}$, each obtained by evolving a random initial state. We train REGIS using the distributions at each time point as training marginals. Alternative selection of marginals is possible, as REGIS is robust to the choice of their number and positions (Figure 13). Following the experimental setup of the MNIST experiment, we compare it with a non-local convolutional control, MMSFM and MMtSBM. Additionally, we train several models with a relaxed assumption that paired data is available – namely a state-conditioned DDPM (Ho et al., 2020) and a convolutional GAN. All models have been modified to incorporate system-specific inductive biases (Appendix B.2, B.3).

To evaluate the models, we first quantify the main trend of this system – the growth of the linear size of the domains, measured as the domain length $L(t)$ (Appendix B.4). Our primary metric, $L(t)/L_{\mathrm{GT}}(t)$, compares this length with that of the ground truth simulator at the same timepoint; a value of 1 indicates agreement. REGIS closely follows the reference growth at the observed times and at the unobserved time points between them, and accurately extrapolates to $t = 4000$, four times the last observation time (Figure 4). At this final time, its median $L(t)/L_{\mathrm{GT}}(t)$ ratio is 1.01, with a range of 1.00–1.05 across three training seeds (Table 5). All other models trained in the marginals-only setting qualitatively capture the growth trend, but numerically deviate from the reference for some or all seeds. Representative rollouts are shown in Figure 3.

The paired DDPM provides the closest comparison, also capturing domain growth beyond the training window. We further compare the two methods as the amount of training data is reduced (Figure 5). At $t = 4000$, REGIS's median domain length remains close to the reference across all tested data budgets, including just 50 unpaired images. The DDPM underestimates domain length, with larger deviations at the smaller data budgets.

Importantly, REGIS not only matches the system's macroscopic behaviour, but also recovers the underlying microscopic rules, learning the most accurate spin-flip probabilities for different neighbour

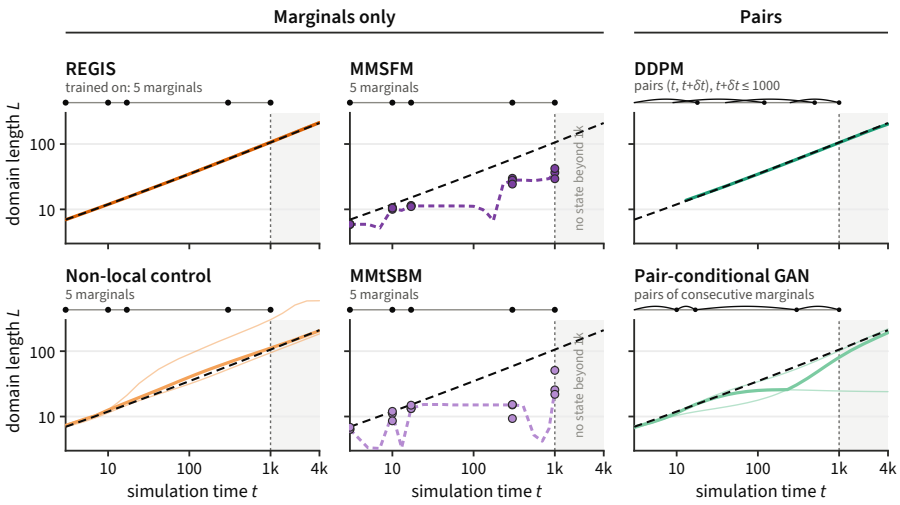

Figure 4: **Domain growth against the ground truth.** Mean domain length $L(t)$ over simulation time t (log–log). The dashed black line indicates the growth of the ground truth. The strip above each panel marks the marginals observed during training or the structure of data used for pair-trained models. $n = 3$ seeds are used for each model; the bold line indicates the median. Transport models are evaluated at their marginals (filled circles) and part-way along a transition (Appendix B.5); they have no state beyond $t = 1000$ (grey). Numeric values are shown in Table 5.

configurations across all evaluated methods (Figures 11, 12). This can be attributed in part to the locality constraint: models with larger receptive fields perform worse on all metrics (Figures 10, 11).

Despite receiving only sparse, unpaired snapshots, REGIS discovers the ground-truth rules that govern the system's evolution and applies them to extrapolate beyond the training window. These results suggest that REGIS-like models are a promising tool for learning predictive models of spatial processes from limited observations, even when individual trajectories cannot be measured.

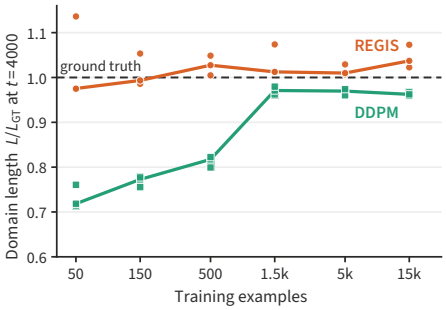

Figure 5: **Data efficiency.** Mean domain length relative to the ground truth, $L(4000)/L_{\mathrm{GT}}(4000)$, as a function of training set size. The closer to 1 the better. Markers show individual training runs; solid lines connect the medians across training seeds. The dashed line marks the ground-truth median. REGIS learns the correct growth trend under all tested data budgets, including as little as 50 unpaired views – 10 samples per marginal.

3.3 Recovering dynamics from real-world data

To test whether the framework recovers meaningful dynamics from real-world data, we apply it to a published spatiotemporal atlas of zebrafish heart regeneration (Li et al., 2025). The atlas profiles heart cross-sections by spatial transcriptomics, giving the spatial location of each annotated cell type in the uninjured heart and at seven time points after induced injury to the ventricular area (105 sections from 35 hearts, 3 sections per heart, 4–6 hearts per time point; Section C.1, Figure 6). Because the measurement destroys the sample, no heart is observed twice and the data are unpaired snapshots rather than trajectories. The dataset has the typical obstacles of biological data: limited sample sizes, destructive measurement, and high variability in heart shape, composition and response (Table 9, Figure 16), matching the challenges REGIS was designed to address.

We define the target dynamics such that the uninjured state persists until injury, introduced at random, after which the trajectory passes through the observed marginal states before returning to the

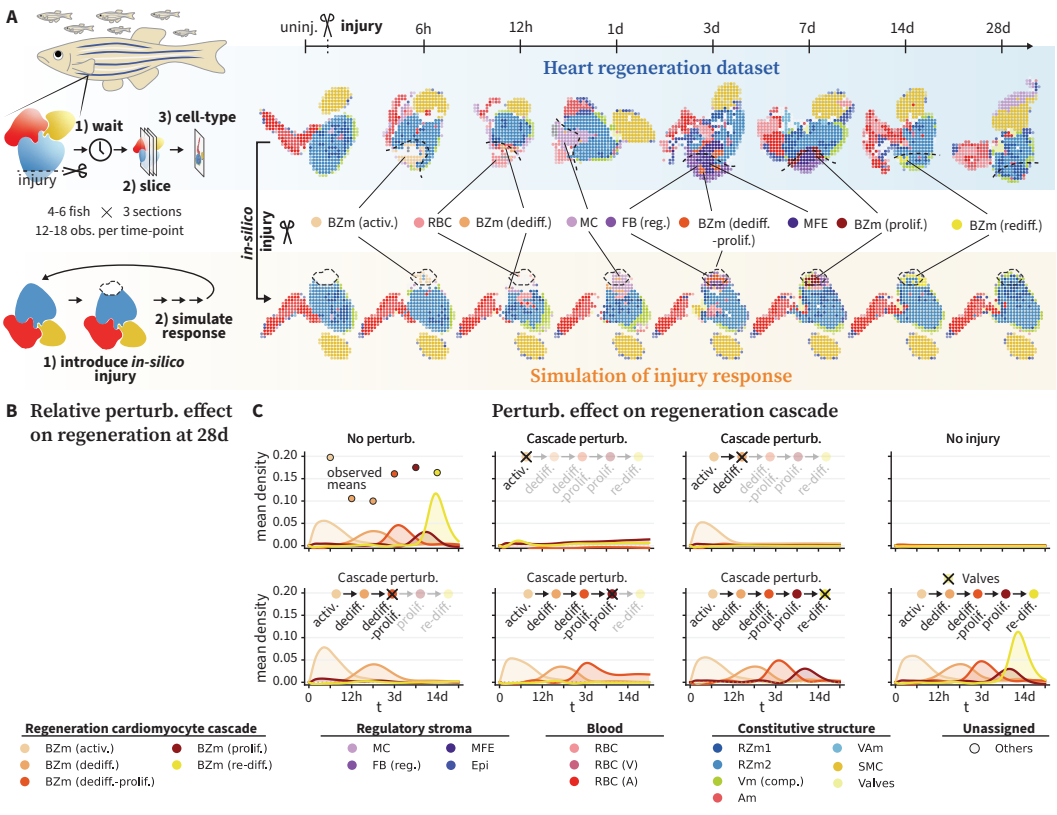

Figure 6: **Simulating zebrafish heart regeneration.** **A** *Top:* Experimental scheme. After ventricular apex amputation, hearts are collected at fixed times, sectioned and profiled by spatial transcriptomics. Regeneration progresses through cardiomyocyte (BZm) states (Section C.1) alongside stromal populations (MC, FB). Variation between panels is both technical (heart size and shape) and biological (injury response). *Bottom:* Regeneration simulated for an individual heart after an *in silico* injury. **B** Perturbation effects on regeneration. Implicated groups have stronger effects than structural groups. **C** Abundance of the five cardiomyocyte states during regeneration: unperturbed (dots, observed levels in the cohort), each state zeroed in turn (cross), and two controls: no injury, and a valve knockout (not implicated). Removing an early state suppresses downstream states while preserving upstream ones, recovering the cascade's known hierarchical dependencies.

uninjured state. Unlike Ising, this data doesn't have a simple ground truth. After training, we evaluate REGIS by introducing previously unobserved injuries into the uninjured heart samples. To probe whether the model captures causally relevant biology, we perform *in silico* perturbation experiments analogous to biological knockouts, zeroing cell states and measuring the effect on the resulting trajectory and state at the time of settled recovery (after 28d) (Section C.4; Figure 6B,C).

The simulator accurately generates unobserved regeneration trajectories following simulated injury, achieving $r = 0.848 \pm 0.066$ across seven seeds (Table 11 and Figure 17). When one fish is held out entirely during training, the median temporal correlation on its sections remains $r = 0.900$, demonstrating generalisation to unseen hearts and injuries (Section C.3).

We next test whether the learnt dynamics recover known biological dependencies. Comparing knockouts across cell-types grouped by their potential functional roles shows that removing regeneration-programme and regulatory-stroma populations has a stronger effect on the final state than removing blood or structural cells ($77 \pm 9\%$ of cross-group pairs, one-sided $p = 0.008$ across seeds; Table 10, Figure 6B, Figure 20A). This agrees with the known role of these populations in regeneration (Table 8). Heart regeneration proceeds through a downstream cascade involving border-zone cardiomyocytes (BZm) (Li et al., 2025). Knockout (zeroing out) of the involved cell types in the simulation disrupts propagation to downstream stages ($31 \pm 10\%$) while leaving upstream stages largely unaffected ($99 \pm 12\%$), consistent with the documented dependency structure

(Table 11; Figure 6C). In contrast, negative-control knockouts of cell types with no known role in regeneration, such as red blood cells and valve cells, leave the trajectory largely unchanged, retaining $87 \pm 14\%$ of the unperturbed amplitude across seeds (Table 11 and Figure 6C).

Both architectural constraints – locality and temporal independence – are required for this result (Tables 10 and 11, Section C.3, and Figure 20). Providing the generator with an explicit time label improves agreement with the observed marginals ($r = 0.956 \pm 0.004$), but weakens the learned cascade, retaining $63 \pm 7\%$ after knockout compared with $31 \pm 10\%$ for REGIS, and recovers the cell-type groups basically at chance ($60 \pm 11\%$ of pairs). Similarly, a non-local control reproduces the observed profiles ($r = 0.877 \pm 0.038$), but retains $74 \pm 4\%$ of the cascade and recovers the groups at chance ($52 \pm 11\%$ of pairs). Notably, the time-conditioned model appears to fit sampling fluctuations as dynamics: it removes and regenerates valve cells in response to marginal variation, whereas REGIS preserves their presence once established, yielding more consistent dynamics (Figure 17). MMSFM, the only baseline whose chained rollout fits the marginals, recovers neither the cascade (101% retained) nor the groups (21% of pairs; Tables 10 and 11 and Section C.5).

The knockout experiments are interventional tests within the learned simulator. Their agreement with the documented regeneration cascade shows that REGIS recovers hierarchical dependencies consistent with known biology that are not apparent in the snapshots themselves, and that models without temporal invariance or spatial locality constraints fail to recover.

4 Discussion

In this work, we have presented REGIS, a method for learning spatial dynamics from unpaired snapshots by combining neural cellular automata with adversarial training. On the synthetic cycling MNIST task, it preserves the quality of the generated samples over long rollouts. The Ising and heart experiments carry this idea further, showing that the model can reproduce coarsening beyond the observed times and capture biologically meaningful dependencies in a regenerating tissue.

We attribute these results in part to constraints and inductive biases of the model. Learning simple local rules has several advantages. The model only needs to encode changes within a neighbourhood, allowing a compact update network and faster inference, which could additionally be applied to domains larger than those seen during training. Sharing this rule spatially improves data efficiency, as each sample contains many local neighbourhoods, each providing a training signal.

Time independence provides a further constraint, preventing spurious overfitting: the heart regeneration ablation shows that making the model time-dependent improves agreement with observed cell-types, but weakens the dependencies revealed by knockouts (Table 11). The simple formulation of the trajectory generation – sequential application of updates – enables simulation between and beyond the observed marginals, and also makes it easy to intervene in the learned process, facilitating the *in silico* experiments on digital twins.

Finally, we expect the growing availability of methods of spatiotemporal profiling to drive new approaches to learning dynamics from unpaired observations. Single-cell genomics has already motivated methods based on optimal transport and continuous flows (Schiebinger et al., 2019; Tong et al., 2020). Spatial assays bring neighbourhoods and tissue organisation into view, giving locality a concrete role in modelling how these systems evolve. REGIS demonstrates how local rules learned from unpaired snapshots can provide a dynamical model of an evolving system. We hope the approach explored in REGIS and its underlying principles will help shape future methods for studying spatial processes and, as spatiotemporal datasets accumulate, become increasingly useful for scientific investigation.

AI USE STATEMENT

In this work, we used generative AI tools for improving the grammar and language of the manuscript, writing the research code and the manuscript LaTeX code, and finding the missing references. The ideas behind the study, all research decisions, choice of methods, interpretations of the results, writing the manuscript, creating the figures and choice of the title were all the work of the authors. All code generated by generative AI tools was reviewed by the authors. All suggested references were checked manually for not being hallucinated. We take responsibility for the final content of this work, including text, claims or artefacts produced with the aid of generative AI.

ETHICS STATEMENT

This work focuses on methodological aspects of building models for biology, and does not involve human subjects or data. In the scope of this study, we did not perform any biological experiments. The zebrafish heart regeneration data was collected from a previously published dataset. We do not identify any inherent ethical concerns in the method itself.

REPRODUCIBILITY STATEMENT

In this section, we provide the materials and information to reproduce the results presented in the paper:

- **Code:** The code for REGIS is available at GitHub `https://github.com/anon03752/regis`. The repository includes the code for the experiments and figures, and contains detailed instructions for reproducing the paper.
- **Datasets:** The two external datasets used in the paper – MNIST and the zebrafish heart regeneration data (Li et al., 2025) – are available online, with links provided in the corresponding appendices and the GitHub repository.
- **Computational resources:** All experiments were performed using NVIDIA H100 GPUs. Training of a single REGIS model requires approximately 3 hours on a single GPU for the MNIST and Ising models, and approximately 10 hours for the zebrafish experiment.

REFERENCES

Michael Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *Journal of Machine Learning Research*, 26(209):1–80, 2025. URL `https://jmlr.org/papers/v26/23-1605.html`.

A. J. Bray. Theory of phase-ordering kinetics. *Advances in Physics*, 43(3):357–459, 1994. doi: 10.1080/00018739400101505. URL `https://arxiv.org/abs/cond-mat/9501089`.

Charlotte Bunne, Laetitia Papaxanthos, Andreas Krause, and Marco Cuturi. Proximal optimal transport modeling of population dynamics. In *International Conference on Artificial Intelligence and Statistics*, pp. 6511–6528. PMLR, 2022.

Roy J. Glauber. Time-dependent statistics of the Ising model. *Journal of Mathematical Physics*, 4(2):294–307, 1963. doi: 10.1063/1.1703954. URL `https://doi.org/10.1063/1.1703954`.

Juan Manuel González-Rosa, Caroline E. Burns, and C. Geoffrey Burns. Zebrafish heart regeneration: 15 years of discoveries. *Regeneration*, 4(3):105–123, 2017. doi: 10.1002/reg2.83. URL `https://doi.org/10.1002/reg2.83`.

Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, volume 27, 2014. URL `https://proceedings.neurips.cc/paper_files/paper/2014/hash/f033ed80deb0234979a61f95710dbe25-Abstract.html`.

Thomas Gravier, Thomas Boyer, and Auguste Genovesio. Multi-marginal temporal Schrödinger Bridge Matching from unpaired data, 2025. URL https://arxiv.org/abs/2510.01894.

Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. In *Advances in Neural Information Processing Systems*, volume 30, 2017. URL https://proceedings.neurips.cc/paper/2017/hash/8a1d694707eb0fefe65871369074926d-Abstract.html.

Jonathan Ho, Ajay N. Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html.

Hessel Honkoop, Dennis E. M. de Bakker, et al. Single-cell analysis uncovers that metabolic reprogramming by erbb2 signaling is essential for cardiomyocyte proliferation in the regenerating heart. *eLife*, 8:e50163, 2019. doi: 10.7554/eLife.50163.

Bo Hu, Sara Lelek, Bastiaan Spanjaard, Hadil El-Sammak, Mariana Guedes Simoes, Janita Mintcheva, Hananeh Aliee, Ronny Schaefer, Alexander M. Meyer, Fabian J Theis, Didier Y R Stainier, Daniela Panakova, and Jan Philipp Junker. Origin and function of activated fibroblast states during zebrafish heart regeneration. *Nature Genetics*, 54(8):1227–1237, 2022. doi: 10.1038/s41588-022-01129-5.

Yiwen Huang, Aaron Gokaslan, Volodymyr Kuleshov, and James Tompkin. The GAN is dead; long live the GAN! a modern GAN baseline. In *Advances in Neural Information Processing Systems*, volume 37, 2024. doi: 10.52202/079017-1402. URL https://papers.nips.cc/paper_files/paper/2024/hash/4e2acb1e1c8e297d394ae29ed9535172-Abstract-Conference.html.

Guillaume Huguet, Daniel Sumner Magruder, Alexander Tong, Oluwadamilola Fasina, Manik Kuchroo, Guy Wolf, and Smita Krishnaswamy. Manifold interpolating optimal-transport flows for trajectory inference. In *Advances in Neural Information Processing Systems*, volume 35, 2022.

Chris Jopling, Eduard Sleep, Marina Raya, Mercè Martí, Angel Raya, and Juan Carlos Izpisúa Belmonte. Zebrafish heart regeneration occurs by cardiomyocyte dedifferentiation and proliferation. *Nature*, 464:606–609, 2010. doi: 10.1038/nature08899. URL https://www.nature.com/articles/nature08899.

Kazu Kikuchi, Jennifer E. Holdway, Andreas A. Werdich, Ryan M. Anderson, Yi Fang, Gregory F. Egnaczyk, Todd Evans, Calum A. MacRae, Didier Y. R. Stainier, and Kenneth D. Poss. Primary contribution to zebrafish heart regeneration by gata4^{+} cardiomyocytes. *Nature*, 464(7288):601–605, 2010. doi: 10.1038/nature08804.

Kazu Kikuchi, Jennifer E Holdway, Robert J Major, Nicola Blum, Randall D Dahn, Gerrit Begemann, and Kenneth D Poss. Retinoic acid production by endocardium and epicardium is an injury response essential for zebrafish heart regeneration. *Developmental Cell*, 20(3):397–404, 2011. doi: 10.1016/j.devcel.2011.01.010.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015. URL https://arxiv.org/abs/1412.6980.

Diederik P. Kingma and Max Welling. Auto-encoding variational Bayes. In *International Conference on Learning Representations*, 2014. URL https://arxiv.org/abs/1312.6114.

Dominik Klein, Giovanni Palla, Marius Lange, Michal Klein, Zoe Piran, Manuel Gander, Laetitia Meng-Papaxanthos, Michael Sterr, Lama Saber, Changying Jing, et al. Mapping cells through time and space with moscot. *Nature*, 638(8052):1065–1075, 2025.

Oskar Kviman, Kirill Tamogashev, Nicola Branchini, Víctor Elvira, Jens Lagergren, and Esmeralda S. Whitammer. Multi-Marginal Flow Matching with Adversarially Learnt Interpolants. In *International Conference on Learning Representations*, March 2026. doi: 10.48550/arXiv.2510.01159. URL https://arxiv.org/abs/2510.01159. arXiv:2510.01159, version 3.

Shih-Lei Lai, Rubén Marín-Juez, Pedro Luís Moura, Carsten Kuenne, Jason Kuan Han Lai, Ayele Taddese Tsedeke, Stefan Guenther, Mario Looso, and Didier Y R Stainier. Reciprocal analyses in zebrafish and medaka reveal that harnessing the immune response promotes cardiac regeneration. *eLife*, 6:e25605, 2017. doi: 10.7554/eLife.25605.

Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

Justin Lee, Behnaz Moradijamei, and Heman Shakeri. Multi-marginal stochastic flow matching for high-dimensional snapshot data at irregular time points. *arXiv preprint arXiv:2508.04351*, 2025. doi: 10.48550/arXiv.2508.04351. URL https://arxiv.org/abs/2508.04351.

Lei Li, Meina Lu, Lidong Guo, Xuejiao Zhang, Qun Liu, Meiling Zhang, Junying Gao, Mengyang Xu, Yijian Lu, Fang Zhang, Yao Li, Ruihua Zhang, Xiawei Liu, Shanshan Pan, Xianghui Zhang, Zhen Li, Yadong Chen, Xiaoshan Su, Nannan Zhang, Wenjie Guo, Tao Yang, Jing Chen, Yating Qin, Zhe Zhang, Wei Cui, Lindong Yu, Ying Gu, Huanming Yang, Xun Xu, Jianxun Wang, Caroline E Burns, C Geoffrey Burns, Kai Han, Long Zhao, Guangyi Fan, and Ying Su. An organ-wide spatiotemporal transcriptomic and cellular atlas of the regenerating zebrafish heart. *Nat. Commun.*, 16(1):3716, April 2025. doi: 10.1038/s41467-025-59070-0. URL https://www.nature.com/articles/s41467-025-59070-0.

Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. In *International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=PqvMRDCJT9t.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. URL https://arxiv.org/abs/1711.05101.

Anh Thi Luu, Nick Lemke, and Anirban Mukhopadhyay. Coarse to fine: Iterative adversarial neural cellular automata for medical image synthesis. *arXiv preprint arXiv:2608.28909*, 2026.

Rubén Marín-Juez, Michele Marass, Sebastien Gauvrit, Andrea Rossi, Shih-Lei Lai, Stefan C Materna, Brian L Black, and Didier Y R Stainier. Fast revascularization of the injured area is essential to support zebrafish heart regeneration. *Proceedings of the National Academy of Sciences*, 113 (40):11237–11242, 2016. doi: 10.1073/pnas.1605431113.

I. M. McKenna, S. O. Poulsen, E. M. Lauridsen, W. Ludwig, and P. W. Voorhees. Grain growth in four dimensions: A comparison between simulation and experiment. *Acta Materialia*, 78:125–134, 2014. doi: 10.1016/j.actamat.2014.06.028. URL https://www.sciencedirect.com/science/article/pii/S1359645414004480.

Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for GANs do actually converge? In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 3481–3490. PMLR, 2018. URL https://proceedings.mlr.press/v80/mescheder18a.html.

Takeru Miyato and Masanori Koyama. cGANs with projection discriminator. In *International Conference on Learning Representations*, 2018. URL https://openreview.net/forum?id=ByS1VpgRZ.

Alexander Mordvintsev, Ettore Randazzo, Eyvind Niklasson, and Michael Levin. Growing neural cellular automata. *Distill*, 2020. doi: 10.23915/distill.00023. URL https://distill.pub/2020/growing-ca/.

Maximilian Otte, Quentin Delfosse, Johannes Czech, and Kristian Kersting. Generative adversarial neural cellular automata. *arXiv preprint arXiv:2108.04328*, 2021. doi: 10.48550/arXiv.2108.04328. URL https://arxiv.org/abs/2108.04328.

Q. Ouyang and Harry L. Swinney. Transition from a uniform state to hexagonal and striped Turing patterns. *Nature*, 352:610–612, 1991. doi: 10.1038/352610a0. URL https://www.nature.com/articles/352610a0.

Rasmus Berg Palm, Miguel González-Duque, Shyam Sudhakaran, and Sebastian Risi. Variational neural cellular automata. *arXiv preprint arXiv:2201.12360*, 2022.

Taesung Park, Alexei A. Efros, Richard Zhang, and Jun-Yan Zhu. Contrastive learning for unpaired image-to-image translation. In *European Conference on Computer Vision*, 2020. URL https://www.ecva.net/papers/eccv_2020/papers_ECCV/html/792_ECCV_2020_paper.php.

Jeremy E. Purvis, Kyle W. Karhohs, Caroline Mock, Eric Batchelor, Alexander Loewer, and Galit Lahav. p53 dynamics control cell fate. *Science*, 336(6087):1440–1444, 2012. doi: 10.1126/science.1218351. URL https://doi.org/10.1126/science.1218351.

Alex D. Richardson, Tibor Antal, Richard A. Blythe, and Linus J. Schumacher. Learning spatio-temporal patterns with neural cellular automata. *PLOS Computational Biology*, 20(4): e1011589, 2024. doi: 10.1371/journal.pcbi.1011589. URL https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1011589.

Martin Rohbeck, Edward De Brouwer, Charlotte Bunne, Jan-Christian Huetter, Anne Biton, Kelvin Chen, Aviv Regev, and Romain Lopez. Modeling complex system dynamics with flow matching across time and conditions. In *International Conference on Learning Representations*, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/hash/206b764b00195b8c896c5af8bde80726-Abstract-Conference.html.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, volume 9351 of *Lecture Notes in Computer Science*, pp. 234–241. Springer, 2015. doi: 10.1007/978-3-319-24574-4_28. URL https://link.springer.com/chapter/10.1007/978-3-319-24574-4_28.

Pauline Sallin, Anne-Sophie de Preux Charles, Vincent Duruz, Catherine Pfefferli, and Anna Jaźwińska. A dual epimorphic and compensatory mode of heart regeneration in zebrafish. *Developmental Biology*, 399(1):27–40, 2015. doi: 10.1016/j.ydbio.2014.12.002.

Héctor Sánchez-Iranzo, María Galardi-Castilla, Andres Sanz-Morejon, Juan Manuel González-Rosa, Ricardo Costa, Alexander Ernst, Julio Sainz de Aja, Xavier Langa, and Nadia Mercader. Transient fibrosis resolves via fibroblast inactivation in the regenerating zebrafish heart. *Proceedings of the National Academy of Sciences*, 115(16):4188–4193, 2018. doi: 10.1073/pnas.1716713115.

Geoffrey Schiebinger, Jian Shu, Marcin Tabaka, Brian Cleary, Vidya Subramanian, Aryeh Solomon, Joshua Gould, Siyan Liu, Stacie Lin, Peter Berube, Lia Lee, Jenny Chen, Justin Brumbaugh, Philippe Rigollet, Konrad Hochedlinger, Rudolf Jaenisch, Aviv Regev, and Eric S. Lander. Optimal-Transport Analysis of Single-Cell Gene Expression Identifies Developmental Trajectories in Reprogramming. *Cell*, 176(4):928–943.e22, February 2019. ISSN 0092-8674, 1097-4172. doi: 10.1016/j.cell.2019.01.006. URL https://www.cell.com/cell/abstract/S0092-8674(19)30039-X.

Xavier Serra-Picamal, Vito Conte, Romaric Vincent, Ester Anon, Dhananjay T. Tambe, Elsa Bazellieres, James P. Butler, Jeffrey J. Fredberg, and Xavier Trepat. Mechanical waves during tissue expansion. *Nature Physics*, 8:628–634, 2012. doi: 10.1038/nphys2355. URL https://www.nature.com/articles/nphys2355.

Jisheng Sun, Elizabeth A Peterson, Annabel Z Wang, Jianhong Ou, Kieko E Smith, Kenneth D Poss, and Jinhu Wang. hapln1 defines an epicardial cell subpopulation required for cardiomyocyte expansion during heart morphogenesis and regeneration. *Circulation*, 146(1):48–63, 2022. doi: 10.1161/CIRCULATIONAHA.121.055468.

Alexander Tong, Jessie Huang, Guy Wolf, David van Dijk, and Smita Krishnaswamy. TrajectoryNet: A dynamic optimal transport network for modeling cellular dynamics. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 9526–9536. PMLR, 2020. URL https://proceedings.mlr.press/v119/tong20a.html.

Alexander Y. Tong, Nikolay Malkin, Kilian Fatras, Lazar Atanackovic, Yanlei Zhang, Guillaume Huguet, Guy Wolf, and Yoshua Bengio. Simulation-free Schrödinger bridges via score and flow matching. In *Proceedings of the 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pp. 1279–1287. PMLR, 2024. URL https://proceedings.mlr.press/v238/tong24a.html.

A. M. Turing. The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 237(641):37–72, 1952. doi: 10.1098/rstb.1952.0012. URL https://doi.org/10.1098/rstb.1952.0012.

Ke-Hsuan Wei, I-Ting Lin, Kaushik Chowdhury, Khai Lone Lim, Kuan-Ting Liu, Tai-Ming Ko, Yao-Ming Chang, Kai-Chien Yang, and Shih-Lei Lai. Comparative single-cell profiling reveals distinct cardiac resident macrophages essential for zebrafish heart regeneration. *eLife*, 12:e84679, 2023. doi: 10.7554/eLife.84679.

Caleb Weinreb, Samuel Wolock, Betsabeh K. Tusi, Merav Socolovsky, and Allon M. Klein. Fundamental limits on dynamic inference from single-cell snapshots. *Proceedings of the National Academy of Sciences*, 115(10):E2467–E2476, 2018. doi: 10.1073/pnas.1714723115. URL https://doi.org/10.1073/pnas.1714723115.

Yasin Yazıcı, Chuan-Sheng Foo, Stefan Winkler, Kim-Hui Yap, Georgios Piliouras, and Vijay Chandrasekhar. The unusual effectiveness of averaging in GAN training. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=SJgw_sRqFQ.

Grace Hui Ting Yeo, Sachit D Saksena, and David K Gifford. Generative modeling of single-cell time series with prescient enables prediction of cell trajectories with interventions. *Nature communications*, 12(1):3222, 2021.

Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. In *International Conference on Learning Representations*, 2016. URL https://arxiv.org/abs/1511.07122.

Yuhui Zhang, Yuchang Su, Chenyu Wang, Tianhong Li, Zoe Wefers, Jeffrey Nirschl, James Burgess, Daisy Ding, Alejandro Lozano, Emma Lundberg, and Serena Yeung-Levy. CellFlux: Simulating cellular morphology changes via flow matching. In *Proceedings of the 42nd International Conference on Machine Learning*, 2025.

Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2223–2232, 2017. doi: 10.1109/ICCV.2017.244. URL https://openaccess.thecvf.com/content_iccv_2017/html/Zhu_Unpaired_Image-To-Image_Translation_ICCV_2017_paper.html.

A MNIST

A.1 Data and preprocessing

We use the standard MNIST split: 60,000 training images and 10,000 test images (LeCun et al., 1998). Every image, for training and for evaluation alike, is resized to 22×22 and then zero-padded by 5 pixels on each side to a 32×32 canvas to create a margin for the digit to grow into.

The task is a cyclic digit sequence. A trajectory starts from a real held-out image of some class k, which is then transformed by the model into classes $(k+1) \bmod 10$, $(k+2) \bmod 10$, and so on indefinitely. One transformation from k to $(k+1) \bmod 10$ is called a transition.

A.2 Evaluation protocol

We evaluate the quality of the generated digits using a Fréchet distance (FID analogue; Heusel et al., 2017) computed in the feature space of a small, self-trained convolutional classifier (98.5% test accuracy, 545k parameters, $1 \to 32 \to 64$ conv, 128-dim linear head). We report distances between Gaussians fitted to the 128-dimensional penultimate-layer features for 5,000 generated samples vs 5,000 real MNIST samples, all classes pooled together. If any generated pixel exceeds 10 in absolute value, the seed is considered to have diverged, and the Fréchet distance is not evaluated from that transition on.

A.3 Models and hyperparameters

REGIS. The state of the model consists of 1 visible channel and 16 hidden channels. Perception is a fixed depthwise 3×3 convolution – identity, Sobel-x and Sobel-y per channel, totalling 51 features. At each step, a 16-dimensional noise vector is concatenated to the features. Unlike the configurations of our model used for other experiments, this noise is not resampled every step, but is kept constant for the whole trajectory, serving a role similar to the latent code in the variational autoencoder (VAE) (Kingma & Welling, 2014). Additionally, we have found that the local nature of the model results in the spatial drift of the digits on the horizon of hundreds of transitions. To mitigate this, we concatenate a fixed positional feature, with the value of the feature being a Gaussian centred on the domain, $\exp\big(-(x^2+y^2)/0.5\big)$ on coordinates normalised to $[-1, 1]$.

The features are then passed to a pointwise MLP of widths $68 \to 256 \to 256 \to 17$ with ReLU activations and a zero-initialised final layer. Increments are clipped to $[-1, 1]$ and applied under a firing mask with $p_{\text{fire}} = 0.5$ shared across channels at a pixel. Finally, an alive mask, the 3×3 maximum of the visible channel thresholded at 0.1, is applied to the entire updated state. One transition is 30 such updates.

The non-local control. To measure the effect of locality, we train a model that is identical to REGIS, but uses a non-local update network. The update network is a stack of dilated convolutions (Yu & Koltun, 2016) at dilations $(1, 2, 5, 9)$, whose receptive field is $1 + 2(1 + 2 + 5 + 9) = 35$ pixels, which allows it to reach the entire image in a single update. The parameter count is similar to that of REGIS; the same training hyperparameters are used.

Training. Both models are trained for $100{,}000$ iterations at batch size 128 with Adam (Kingma & Ba, 2015) ($\eta = 2 \times 10^{-4}$ for both networks, $\beta = (0.0, 0.99)$), under the relativistic objective of Equation (8) with $\gamma = 1.0$ and zero-centred penalties on both real and generated inputs. Starting states are drawn from a pool of 2048 previously generated states; after the transition the generations are returned to the pool or replaced with real MNIST digits after reaching the age of 6 transitions. The discriminator is a convolutional network with four stages of 32 channels. Additionally, we use a squared hinge penalty on the norm of the hidden states above 1. The evaluated timestep is sampled uniformly from $\{26, \ldots, 30\}$. Gradients are clipped at norm 1.0. Evaluation uses an exponential moving average of the generator weights (decay $(1+n)/(10{,}000+n)$ at iteration n).

MMtSBM. We use a reimplementation of MMtSBM (Gravier et al., 2025) using the U-Net configuration used for MNIST in the original work: per transport direction, a U-Net (Ronneberger et al., 2015) with block widths $32/64/128$, two layers per block and a sinusoidal time embedding (3.6 M

parameters); AdamW (Loshchilov & Hutter, 2019), learning rate 2×10^{-4}, $\beta = (0.9, 0.999)$, weight decay 0.01, gradient clipping 1.0, EMA decay 0.9999. We extend their five-digit transition task to a ten-digit cyclic objective, which we train with $\sigma = 1.2$, batch size 256, $50{,}000$ warm-up steps and two IMF iterations of 12,500 steps per direction. Inference uses the forward EMA network only, integrated by Euler–Maruyama at 60 steps per transition.

MMSFM. We use the authors' official implementation (Lee et al., 2025). The flow and score networks use the same U-Net architecture, with channel widths $64/128/128/128$, two residual blocks per encoder level, and attention at spatial resolutions 8 and 4 (9.3 M parameters per network). We use Adam with $\beta = (0.9, 0.999)$, without weight decay or EMA. We adapt the model to the ten-digit cycle using ten overlapping triplets $(d, (d+1) \bmod 10, (d+2) \bmod 10)$ and a periodic time embedding. Training uses monotone cubic transport splines through minibatch exact-OT couplings, with 16 samples per marginal in each window and noise level $\sigma = 0.15$ on images scaled to $[-1, 1]$. We train for $40{,}000$ optimiser updates with two gradient-accumulation steps, giving an effective batch of 320. The learning rate increases linearly from 10^{-8} to 10^{-4} over the first 500 updates, then decreases linearly toward 10^{-8}. Inference uses both trained networks (18.6 M parameters in total), integrating the SDE with 160 stochastic Runge–Kutta steps per transition.

A.4 EXEMPLAR SAMPLES FOR DIFFERENT METHODS, SEEDS AND HORIZONS

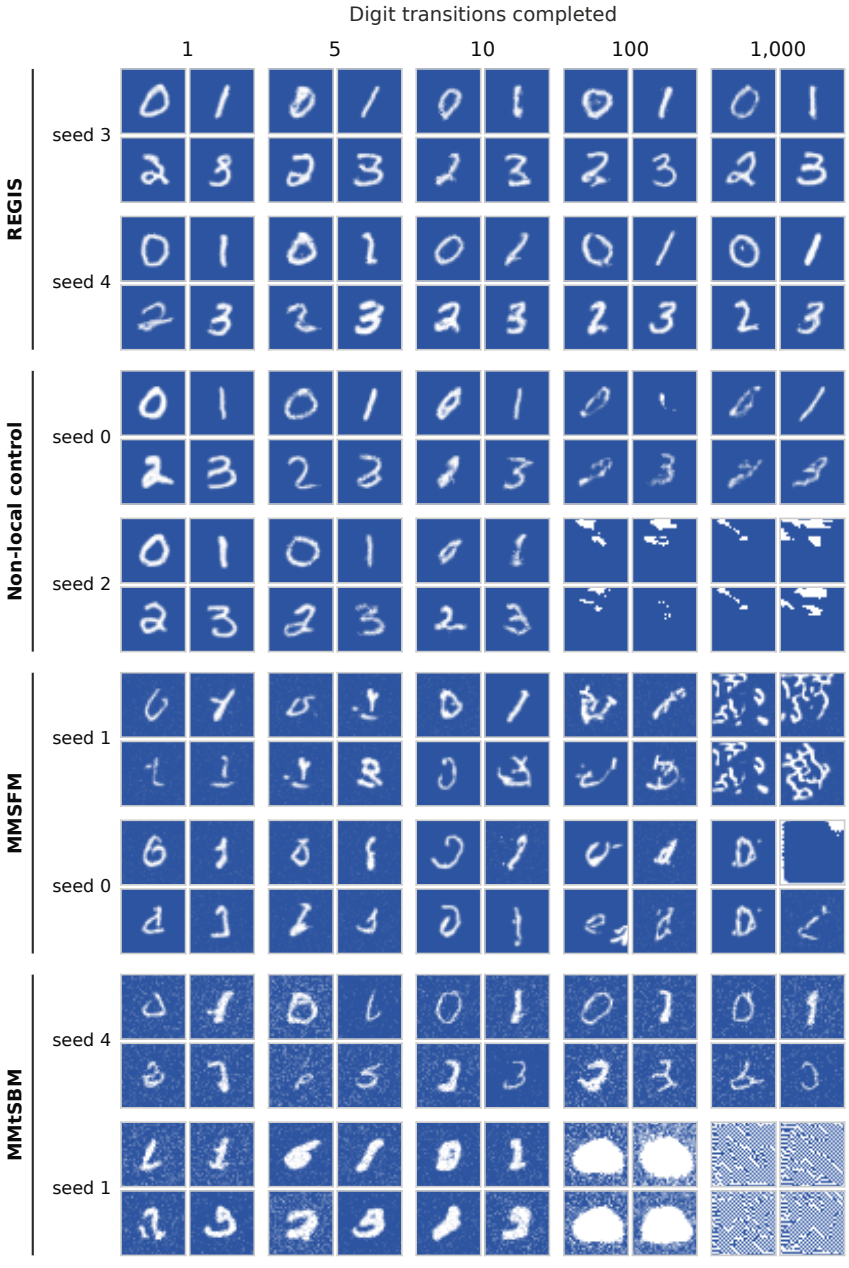

Figure 7: **Generated digits after** 1 **to** 1000 **transitions**, two seeds per method. While each method produces high-quality digits in the first few transitions, only REGIS is able to maintain the quality along the entire trajectory.

A.5 FULL PER-SEED RESULTS

Model	Parameters	ms / transition	Seed	Digit transitions completed: 1	5	10	100	1000
REGIS (ours)	88 K	0.117	0	**37.8** ± 4.4	**28.0** ± 2.9	**49.9** ± 3.6	**20.8** ± 2.3	**21.6** ± 2.6
			1	**34.4** ± 2.9	**44.4** ± 3.9	**28.8** ± 2.6	**22.7** ± 2.7	**23.3** ± 2.9
			2	**54.4** ± 5.6	**21.9** ± 3.0	**39.4** ± 3.3	**19.1** ± 2.6	**18.9** ± 2.6
			3	**37.2** ± 3.7	**28.1** ± 3.3	**45.7** ± 3.5	**20.5** ± 2.7	**19.6** ± 2.5
			4	**33.0** ± 2.7	**29.4** ± 2.9	**48.8** ± 4.0	**25.3** ± 3.0	**24.7** ± 3.0
			5	**49.1** ± 3.8	**23.1** ± 3.3	**45.8** ± 4.2	**20.5** ± 2.9	**21.4** ± 3.0
			6	**28.9** ± 2.0	**43.4** ± 4.4	**43.7** ± 4.1	**22.1** ± 2.9	**20.0** ± 2.6
Non-local control	83 K	0.116	0	67.0 ± 7.6	94.3 ± 5.8	265 ± 12	1,017 ± 22	1,512 ± 34
			1	39.9 ± 4.8	219 ± 10	161 ± 9	1,304 ± 19	1,376 ± 19
			2	81.1 ± 7.2	131 ± 7	208 ± 9	—	—
			3	53.9 ± 5.7	270 ± 10	189 ± 9	1,575 ± 17	1,581 ± 14
			4	59.2 ± 6.3	144 ± 7	167 ± 8	1,088 ± 20	2,533 ± 38
			5	41.8 ± 5.5	131 ± 7	291 ± 10	—	—
			6	50.2 ± 5.9	55.7 ± 5.3	294 ± 12	1,334 ± 22	1,505 ± 23
MMSFM	18.6 M	16.2	0	52.7 ± 4.5	144 ± 9	60.9 ± 5.3	—	—
			1	70.5 ± 5.0	148 ± 7	44.2 ± 3.5	284 ± 9	5,498 ± 47
			2	122 ± 7	174 ± 8	46.6 ± 3.5	—	—
			3	92.4 ± 5.7	159 ± 8	66.6 ± 4.5	241 ± 9	—
			4	42.6 ± 3.0	86.7 ± 5.0	26.9 ± 2.9	1,272 ± 47	64,359 ± 213
			5	81.0 ± 5.7	215 ± 10	60.3 ± 5.2	—	—
			6	104 ± 6	189 ± 9	46.0 ± 4.0	160 ± 7	—
MMtSBM	**3.6 M**	**1.07**	**0**	**287** ± 7	**222** ± 6	**211** ± 5	**236** ± 7	**243** ± 7
			1	363 ± 6	591 ± 13	1,378 ± 34	64,561 ± 459	—
			2	226 ± 6	303 ± 10	848 ± 22	2,125 ± 23	2,129 ± 23
			3	285 ± 6	264 ± 7	259 ± 9	1,028 ± 63	55,284 ± 1,869
			4	**254** ± 6	**204** ± 6	**184** ± 5	**176** ± 5	**176** ± 5
			5	429 ± 7	2,059 ± 27	6,177 ± 46	39,108 ± 90	39,362 ± 91
			6	195 ± 6	378 ± 14	975 ± 24	1,596 ± 19	1,575 ± 20

Table 2: **Per-seed results behind Table 1.** Fréchet distance of the generated digits to the ground truth for individual seeds of different models. ±: half-width of a 95% percentile bootstrap interval over the 5,000 generated samples, the real reference held fixed. —: the seed has diverged by this horizon. REGIS and several seeds of MMtSBM maintain low Fréchet distance even after 1000 transitions. Notably, while for MMtSBM time is enforced externally, REGIS has to rely on its own learned "internal clock".

A.6 LATENT STYLE



Figure 8: **Persistent noise encodes style.** Eight draws of the noise ξ, one per row, selected to show diverse styles. When noise is fixed while the trajectory is being generated, it results in consecutive digits within this trajectory having the same style.

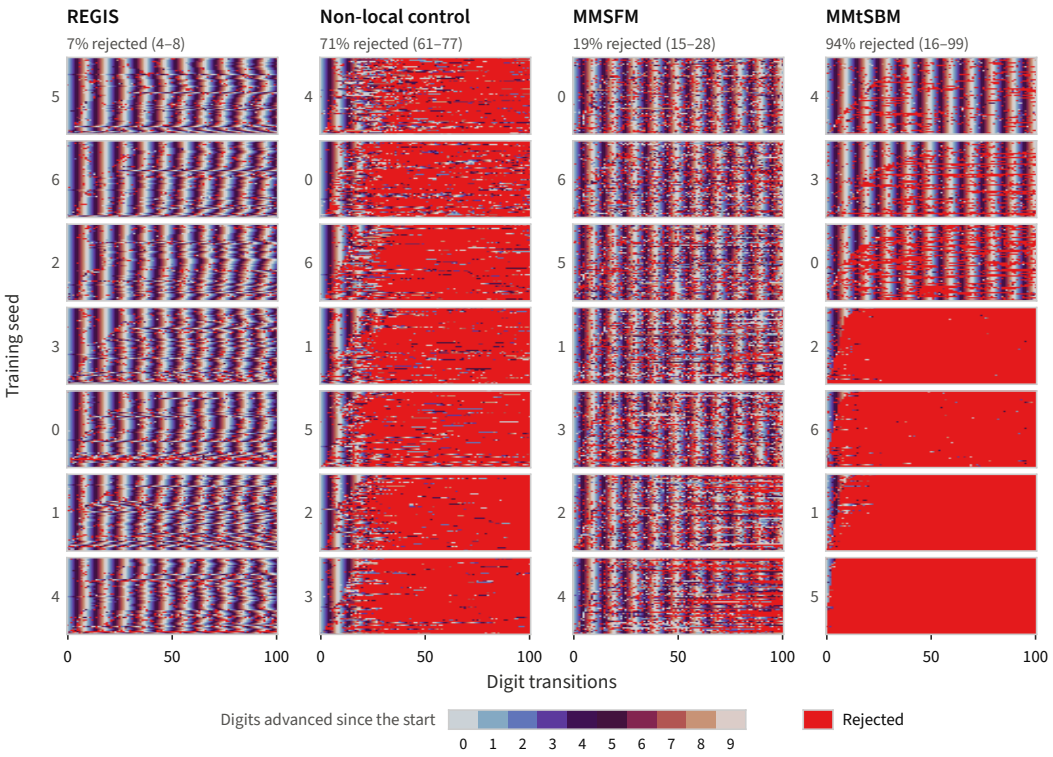

Figure 9: **Cycling and image validity over the first** 100 **transitions, every seed.** Each small panel is one training seed, 60 sampled trajectories. Colour: digits advanced since the start, modulo ten, as read by a MNIST digit classifier. Red: the image is rejected, by classifier confidence below 0.62, stroke area outside the central 98% of real digits of the class read, or a pixel with value beyond 10; together, these criteria reject 2.8% of real test digits. All models start producing valid digits, but only REGIS retains the quality in the long run for all seeds. For MMSFM and MMtSBM, the transition rules are time-dependent and the time is provided externally, resulting in a more consistent transition schedule. REGIS doesn't use time explicitly, but is able to maintain the current digit order, albeit its schedule is more irregular.

B ISING

B.1 THE ISING MODEL AND THE DATA

We study the two-dimensional Ising model on a periodic square lattice, where each site carries a binary spin $s_i \in \{0, 1\}$ and evolves under Glauber dynamics (Glauber, 1963). The model describes ferromagnetic interactions, for which aligned neighbouring spins are energetically favourable. As a result, when started from a random state, the system forms increasingly large domains of aligned spins. A spin's probability of flipping depends only on how many of its four nearest neighbours currently align with it:

$$P_{\text{flip}} = \frac{1}{1 + \exp(\Delta E/T)}, \qquad \Delta E = 2\big(n_{\text{aligned}} - n_{\text{non-aligned}}\big), \tag{9}$$

where ΔE is the energy change from flipping the spin and T is the temperature. The counts n_{aligned} and $n_{\text{non-aligned}}$ refer to the four nearest neighbours, measured relative to the central spin's orientation. For temperature $T = 1$, the spin flip probabilities are given in Table 3.

Aligned neighbours	0	1	2	3	4
Flip probability	0.9997	0.9820	0.5000	1.80×10^{-2}	3.35×10^{-4}

Table 3: **True flip probabilities.** Probability that a spin flips in one step, depending on the number of aligned neighbours.

Updates are applied on a checkerboard schedule, meaning that each update is split into two half-steps, one for each checkerboard sublattice. That way, no two interacting sites update simultaneously. Every trajectory starts from its own random state; each spin is independently set to 1 or 0 with equal probability. Gradually, the system coarsens, forming domains of aligned spins of increasing size. The characteristic size of the domains grows as $t^{1/2}$ (Bray, 1994), where t is the number of full updates. As described in Section B.4, we use this property to evaluate how well the models reproduce the real dynamics.

Training data consists of five marginal distributions at $t \in \{3, 10, 17, 300, 1000\}$. As the growth happens linearly in log time, we chose log-spaced time points to ensure that the marginals represent different stages of domain formation. As in our other experiments, all datapoints are unpaired, meaning that no more than one snapshot is taken from each generated trajectory. All models are trained on these same five marginals, but not on the same amount of data; the training set sizes are given with each model in Section B.3.

B.2 REGIS ARCHITECTURE AND TRAINING

For the perception stage, REGIS uses a single depthwise 3×3 convolution producing three maps per channel. The state has one visible channel, the spin itself, with no hidden channels, and two channels of standard Gaussian noise redrawn at every half-step (which is analogous to half-updates during the data generation; see below). The perceived vector is passed through a per-pixel three-layer network MLP 1×1 ($5 \to 128 \to 128 \to 1$) with ReLU activations, whose output layer is zero-initialised so that the model begins as the identity. The generator is 17,436 parameters in total. It is trained on 100 independent 512×512 fields per marginal, from which a random 128×128 periodic window is drawn each batch.

Given the current field z_t, holding one value in $\{0, 1\}$ per site, and fresh noise ξ, the network produces $\Delta_\theta(z_t, \xi)$, and the field becomes

$$z_{t+0.5} = z_t + m \odot \operatorname{clip}\big(\Delta_\theta(z_t, \xi),\, -c,\, +c\big), \qquad c = 10, \tag{10}$$

where m is the checkerboard mask selecting the half of the lattice that updates. A full step applies this twice, once for each sublattice, so that neighbouring sites never update together.

Additionally, we utilise the fact that REGIS's rules work at the level of individual updates and incorporate priorly known symmetries of the dynamics into the model. First, the dynamics is unchanged

by inverting every spin ($z \mapsto 1 - z$), so we antisymmetrise the raw network f_θ,

$$\Delta_\theta(z_t, \xi) = \tfrac{1}{2}\Big[f_\theta(z_t, \xi) - f_\theta(1 - z_t, -\xi)\Big], \tag{11}$$

which gives $\Delta_\theta(1 - z_t, -\xi) = -\Delta_\theta(z_t, \xi)$ exactly, for any weights learned.

Second, the dynamics is also unchanged by rotating the lattice by $180°$, so the perception kernel W is replaced by its symmetrised form $\frac{1}{2}(W + W^{\mathrm{rot}})$ at every forward pass. To ensure a fair comparison, we also apply these symmetrisations to the baseline models where possible.

To stabilise the recurrent dynamics, along with the update clipping, typically used for NCAs, we also apply a penalty on update magnitude.

$$\mathcal{L}_{\mathrm{upd}} = \lambda\, \mathbb{E}\big[\,\mathrm{relu}(|\Delta_\theta| - 1)^2\,\big], \qquad \lambda = 10^{-4}, \tag{12}$$

We use the hinge loss to allow for updates of up to the magnitude of 1, which is the ground truth single-step update for a spin flip.

Training is adversarial, using the relativistic loss and zero-centred gradient penalty of Huang et al. (2024) and Mescheder et al. (2018). The recipe, shared with the non-local control and the pair-conditional GAN, is given in Table 4. As in the MNIST experiment, rollouts can also start from a pool of previously generated states: for each observed marginal we keep the last 192 generated states, and 30% of the rollouts from that marginal start from this pool. In addition to the generator, a CNN discriminator of roughly 389K parameters is trained using the standard GAN game. While the rollouts for one transition can span hundreds of steps, we only backpropagate through the last 45 steps. We have found this has no negative impact on the training while reducing the memory footprint and computational cost. While for this experiment, the latter is the main motivation behind such truncation, during the work on other REGIS-based projects we have noticed that it also can greatly improve the convergence stability.

Optimiser	AdamW, lr $= 2 \times 10^{-4}$, $\beta = (0,\ 0.9)$, weight decay 0.01, constant
Gradient clipping	norm 1
Batch / crop	8 / 128×128 random crops
Steps	30,000
Adversarial loss	relativistic, zero-centred gradient penalty $\gamma = 1$
Update clip c / hinge weight λ	10 / 10^{-4}
Gradient window	last ≤ 45 steps of a rollout
Weights used	exponential moving average, decay $(1 + n)/(1000 + n)$ at step n (Yazıcı et al., 2019)

Table 4: **Training recipe** shared by REGIS, the non-local control and the pair-conditional GAN.

B.3 Baseline models

Non-local control. To measure the effect of locality, we use a non-local model, trained under the same framework as REGIS. A stack of four 3×3 convolutions with dilations $1, 2, 5, 9$ and 38 channels, followed by two 1×1 layers and a zero-initialised head, gives each step a 35-pixel receptive field. The non-local model has 43,169 parameters and uses the same symmetrisation and update penalty as REGIS.

Pair-conditional GAN. The motivation behind this model is answering the question of how much one would gain from having access to paired training data. To achieve this, we train a GAN where the generator is similar to the non-local control described above, but the discriminator has access not only to the current state of the system, but also its state at the previous available observed timestep (for example, the state at $t = 300$ if the current timestep is $t = 1000$). The critic therefore judges temporal correspondence directly, which potentially can produce a better training signal compared to the marginal-only models.

DDPM. This model represents the strongest supervision setting in the comparison: it is trained directly on data pairs taken from the same trajectory. A state-conditioned denoising diffusion model (Ho et al., 2020) learns to produce the state at time $t + \delta t$ given the state at time t; the conditioning state and the requested jump size (as $\log \delta t$) are passed to the denoiser model along with the noised state as additional inputs. The denoiser is a four-level convolutional U-Net (Ronneberger et al., 2015) of 4.65M parameters with widths 48–192, without attention or normalisation layers. It is trained with the v-prediction objective on a cosine schedule with 1000 diffusion steps, and sampled with 250 ancestral steps during inference. The training set consists of $15{,}000$ pairs, with the start time drawn log-uniformly from $[3, 1000)$ and the jump δt drawn log-uniformly from $[10, 1000)$, rejecting pairs that would end beyond $t = 1000$. Note that this formulation allows us to sample trajectory snapshots at all intermediate times, not only at the five marginals available – similar to REGIS-like models and unlike MMSFM and MMtSBM. Additionally, since the model is not conditioned on the absolute Ising timestep (do not confuse with the denoising timestep), it is possible to autoregressively continue the trajectories beyond the last training marginal. We train for $40{,}000$ steps with AdamW at 10^{-4} and batch size 16. At evaluation, long trajectories are produced by chaining jumps, each conditioned on the binarised output of the previous one. We attempted to train the DDPM with the symmetrisations described in Section B.2, but found that they did not improve the results – likely because the performance was already very good without them.

MMSFM (Lee et al., 2025) learns a continuous-time stochastic interpolation along the sequence of marginals: a flow network and a score network are trained so that integrating the resulting SDE carries samples of one marginal into samples of the next. Since the data contain no correspondence between snapshots at different times, the method builds its own: within each minibatch of 16, samples of consecutive marginals are paired by exact optimal transport, and a monotone cubic spline serves as the interpolation path, surrounded by Brownian-bridge noise of scale $\sigma = 0.6$. Training slides a window of three consecutive marginals along the sequence, with the anchor times spaced equidistantly on the model's internal time axis. The flow and the score network are each a 20.9M-parameter U-Net with attention, 41.7M in total, and both are antisymmetrised over spin inversion exactly as in Equation 11. We train for $20{,}000$ steps with Adam at a peak learning rate of 10^{-4} and an effective batch of 64. Samples are generated by integrating the SDE with a stochastic Runge–Kutta solver, 160 steps per transition.

MMtSBM (Gravier et al., 2025) models each transition between consecutive marginals as a Schrödinger bridge, trained by bridge matching: stochastic bridges are drawn between states independently sampled from the two marginals, and a drift network regresses onto them. The drift is a three-level U-Net with a single attention block at the bottleneck, 3.6M parameters, conditioned on the transition's position in the chain and symmetrised over spin inversion and 180° rotation in the same trained-through fashion as REGIS. The noise scale $\sigma = 2.4$ was selected by a sweep. The full method also trains a backward drift and alternates bridge matching with refinement iterations that re-couple the marginals through the model's own trajectories; in our experiments these iterations consistently reduced the marginal fidelity, so the reported model uses the bridge-matching stage alone, which needs only the forward drift: $50{,}000$ steps with AdamW at 2×10^{-4} and batch size 32. Samples are generated by Euler–Maruyama integration, 120 steps per transition.

Both MMSFM and MMtSBM generate by traversing the chain of marginals in order, so in addition to the five training marginals shared by all models they are given the initial distribution at $t = 0$ — independent random spins — as the starting anchor of the chain: a sample at $t = 1000$ is the end state of five consecutive transitions. Both are trained on $6{,}000$ windows per marginal, cut in pairs from $3{,}000$ independent fields — thirty times the number of fields given to REGIS. Neither method can produce trajectories beyond the last training marginal: each transition is learned explicitly and cannot be reused.

Window and margin. MMSFM, MMtSBM and the DDPM are all trained on 192×192 windows cut from the 512×512 fields, with the training loss evaluated only on the central 128×128 region; the surrounding 32-pixel margin is just a buffer between the content of interest and the image edge. Without this margin, a network trained on crops of a periodic field learns to create a rectangular border around the area of its original crop when extrapolated to larger lattices. The margin is not used at evaluation time.

B.4 DOMAIN LENGTH METRIC

We summarise the high-level structure of the generated samples using the mean domain length, calculated as $L(t) = 1/\rho$, where ρ is the fraction of nearest-neighbour bonds joining opposite spins. This statistic grows as $t^{1/2}$ for the true dynamics.

However, this metric is disproportionately sensitive to "thermal speckles" – the random fluctuations, resulting in isolated non-aligned spins inside solid domains. With the ground truth probability of 3.35×10^{-4}, such fluctuations are the rarest event in the system, which no marginal-only model can reliably reproduce. To focus on the domain geometry, we therefore remove these isolated spins before measuring. Throughout the paper, we report the domain length $L(t)$ calculated after the "speckles" have been removed.

			Observed t		Interpolation t		Extrapolation t	
	Model	Parameters	300	1,000	150	700	2,000	4,000
Marginals only	**REGIS (ours)**	17 K	**1.00** 1.00–1.04	**1.00** 1.00–1.05	**1.00** 0.99–1.04	**1.00** 1.00–1.05	**1.01** 1.00–1.05	**1.01** 1.00–1.05
	Non-local control	43 K	1.14 0.90–2.74	1.06 0.90–2.87	1.16 0.89–2.66	1.08 0.90–2.80	1.00 0.88–3.41	0.98 0.88–2.82
	MMSFM	41.7 M	0.48 0.41–0.51	0.35 0.28–0.40	0.20 0.19–0.24	0.32 0.24–0.35	–	–
	MMtSBM	3.6 M	0.26 0.16–0.26	0.24 0.21–0.48	0.36 0.09–0.36	0.05 0.05–0.15	–	–
Pairs	**DDPM**	4.7 M	**0.99** 0.99–0.99	**0.99** 0.99–1.00	**0.99** 0.99–0.99	**0.99** 0.99–0.99	**0.98** 0.98–0.99	0.96 0.96–0.97
	Pair-conditional GAN	43 K	0.52 0.43–0.91	0.76 0.23–0.96	0.61 0.49–0.94	0.68 0.28–0.94	0.87 0.16–1.01	0.92 0.11–1.05

Table 5: **Domain growth relative to the ground truth.** Mean magnetic domain length L divided by that of the ground-truth process, at 512^2 from the same initial states, 512 fields per cell. The ground truth is the median of three independent runs of the simulator. Columns: the two largest training marginals (observed), two times between training marginals (interpolation) and two beyond the last marginal (extrapolation). Cells: median over three training seeds with the seed range beneath. Bold: mean absolute deviation from 1 across seeds below 2.5%. Pairs: models with access to paired training data. –: the model has no state past its last marginal.

			Observed t		Interpolation t		Extrapolation t	
Model	Reach	Parameters	300	1,000	150	700	2,000	4,000
REGIS, 3×3	3 px	17.4 K	**1.00** 1.00–1.04	**1.00** 1.00–1.05	**1.00** 0.99–1.04	**1.00** 1.00–1.05	**1.01** 1.00–1.05	**1.01** 1.00–1.05
REGIS, 5×5	5 px	17.5 K	0.94 0.94–0.94	0.94 0.93–0.94	0.94 0.94–0.95	0.94 0.93–0.94	0.94 0.93–0.94	0.93 0.92–0.96
REGIS, 7×7	7 px	17.6 K	0.93 0.90–0.99	0.92 0.89–0.99	0.94 0.91–1.00	0.92 0.89–0.99	0.92 0.89–0.98	0.92 0.89–0.99
Non-local control	35 px	43.2 K	1.14 0.90–2.74	1.06 0.90–2.87	1.16 0.89–2.66	1.08 0.90–2.80	1.00 0.88–3.41	0.98 0.88–2.82

Table 6: **Perception radius and domain growth.** The same measurement as Table 5 for REGIS at three perception radii and the non-local control. The three neural cellular automata are parameter-matched to within 0.7%, so the ladder varies reach, not capacity. Cells: median over three training seeds with the seed range beneath; bold as in Table 5. The non-local control repeats its Table 5 row as the non-local end of the ladder.

B.5 INTERMEDIATE STAGES OF TRANSPORT METHODS

By construction, MMSFM and MMtSBM produce samples from the training marginals: integrating the learned flow maps samples from one marginal to the next. Unlike REGIS, for which the choice of evaluation step can be done arbitrarily, neither method formally makes a prediction between the marginals. In flow matching, the interpolant connecting two distributions is a design choice – and different choices yield different intermediate states with the same endpoints. For the Schrödinger

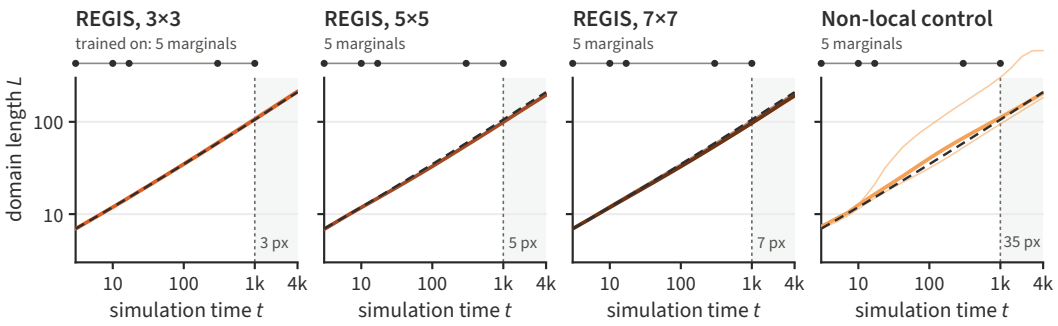

Figure 10: **Perception radius.** Domain length over time for REGIS at three perception radii and the non-local control. Three training seeds each; the ground truth is dashed.

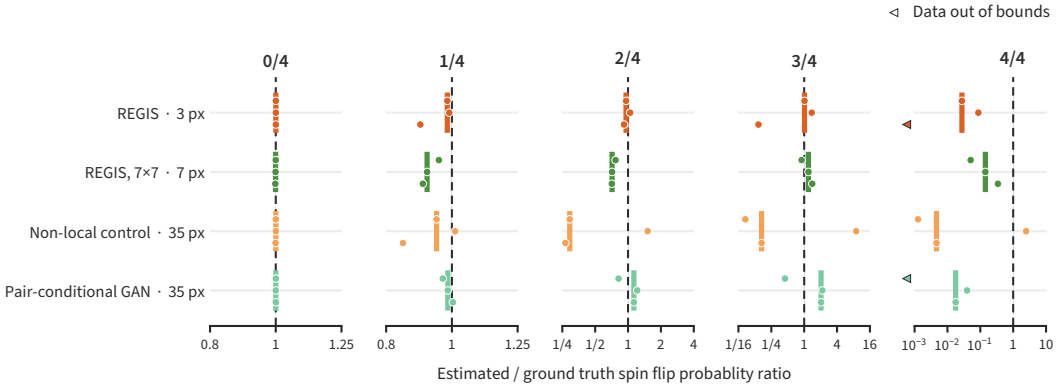

Figure 11: **Learned flip rates.** Flip rate divided by the heat-bath rate, by number of aligned neighbours. One point per training seed, bar is the seed median, shading is the 95% binomial interval at that row's sample size.

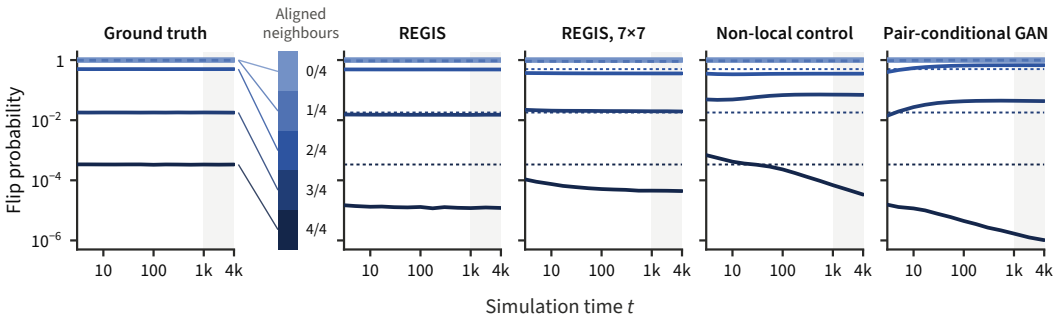

Figure 12: **Learned flip rates over time.** The flip probability of each class, resolved in simulation time and pooled over three training seeds. Dotted: the heat-bath probability of each class. The $1/4$ curve is dashed because it coincides with $0/4$. Shaded: beyond the last training marginal.

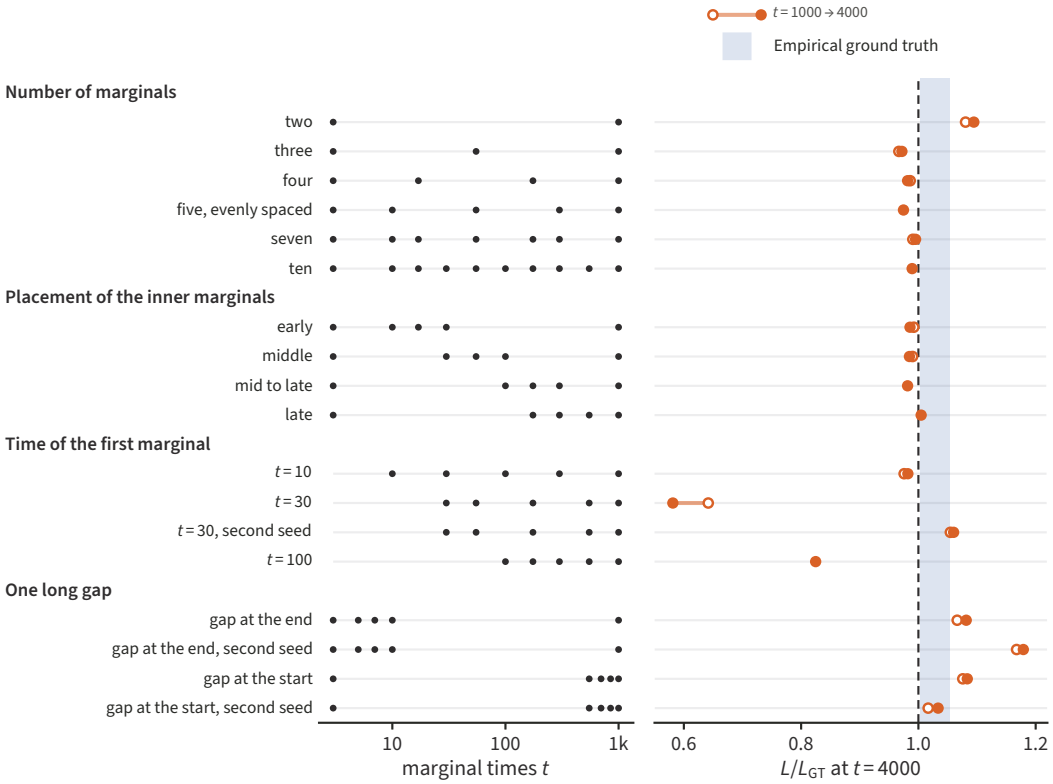

Figure 13: **Choice of training marginals.** L/L_{GT} at $t = 4000$ for fifteen training grids; three of them were retrained with a second seed, giving eighteen runs. The strip on the left marks where each grid's marginals sit.

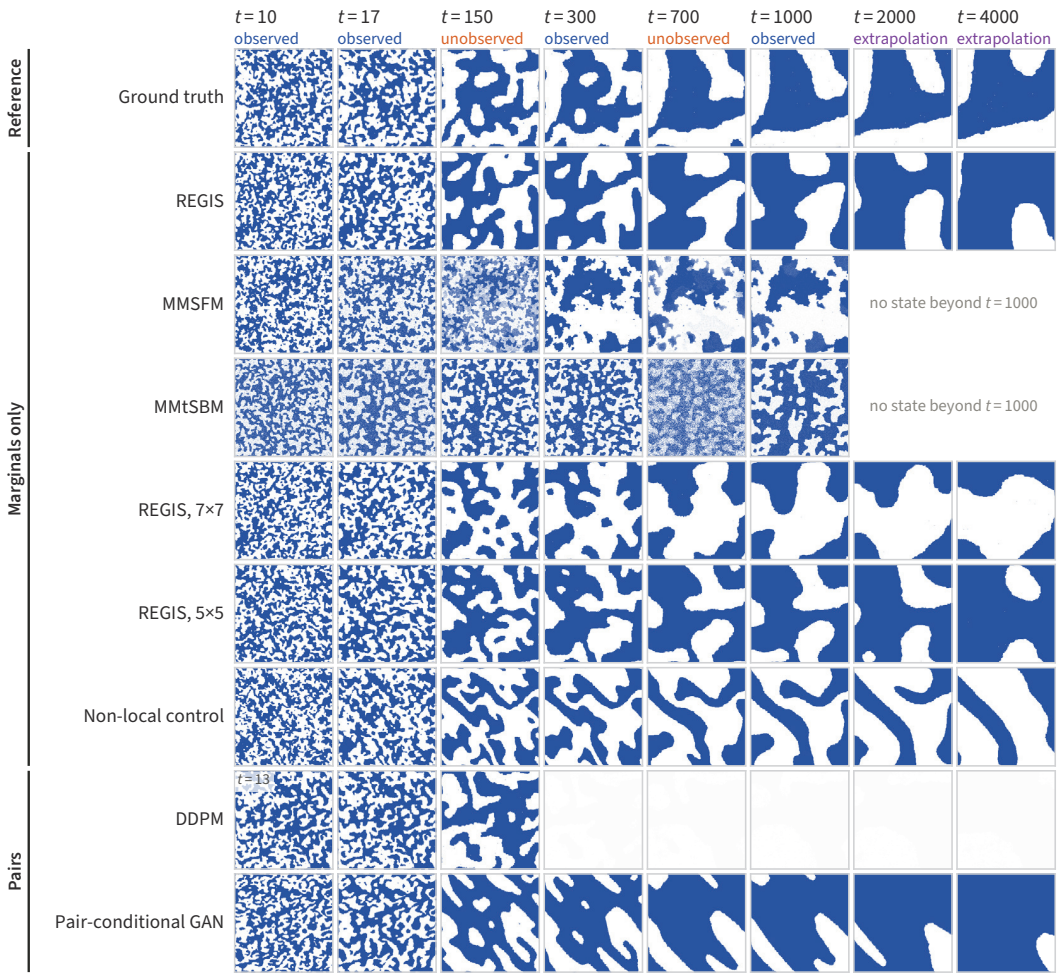

Figure 14: **Generated fields, all methods.** The top-left 256^2 window of the 512^2 field, from one shared start, at the observed, unobserved and extrapolated times, for all models including the perception-radius ladder, the non-local control and the pair-conditional GAN. Transport and DDPM frames are shown un-thresholded; transport rows have no state past $t = 1000$, and the DDPM row begins at its earliest reachable time $t = 13$.

bridges, the intermediate law is a genuine estimate only under the assumption that the true process is the reference diffusion, which does not hold for the Ising dynamics. Nevertheless, in many works, learned interpolations of this kind are interpreted as models of the underlying continuous dynamics (Tong et al., 2020; Huguet et al., 2022; Zhang et al., 2025). Therefore, while technically the states between the marginals are not inferrable from MMSFM and MMtSBM, we still evaluate the quasi-intermediate states obtained by only partially integrating the sampler, as shown in Figures 4, 3 and 14. Unsurprisingly, the resulting states are not faithful approximations of the true intermediate states of the Ising dynamics, but merely interpolation between the training marginals or high-noise bridge-matching intermediate states.

C ZEBRAFISH HEART REGENERATION

C.1 THE ATLAS AND THE DATA

We study regeneration of the adult zebrafish heart after amputation of the ventricular apex, using the Stereo-seq spatial transcriptomics atlas of Li et al. (2025). The atlas profiles the heart at eight stages: uninjured, and 6 h, 12 h, 1 d, 3 d, 7 d, 14 d and 28 d after amputation. The cohort is 105 cross-sections from 35 fish. Each fish was sectioned onto three Stereo-seq chips and so contributes exactly three sections, with 6 fish at the uninjured stage, 5 at 14 d and 4 at each of the remaining six (18, 12, 12, 12, 12, 12, 15 and 12 sections respectively). The measurement is destructive, so no heart is observed more than once and the data contain no measured trajectory. Sections from one fish are serial slices of a single heart and are therefore correlated, so the effective sample size per stage is 4–6 hearts.

We use the authors' annotations, which assign each spot to one of 20 mutually exclusive categories: 19 cell types and *Unassigned*. Each spot is treated as a one-hot vector and binned by the grid index in its barcode, so that on the native grid each bin's composition sums to 1. We then sum-pool 2×2 blocks onto a 48×48 canvas, so that a fully occupied interior bin sums to 4 and, unlike the raw data, bins hold mixtures of cell types. Heart sections and their variability are shown in Figures 15 and 16 and Table 9.

Prior biological knowledge. Unlike the Ising model, the heart has no simulator to compare against; the reference is prior biology instead. The atlas annotates border-zone myocardium (BZm) as five states — activated, dedifferentiating, dedifferentiating-proliferative, proliferative and redifferentiating — each enriched in a narrow window of the arc (6 h; 12 h and 1 d; 3 d; 7 d; 14 d), and describes them as a cascade (Li et al., 2025). At coarser resolution this matches the dedifferentiation–proliferation–redifferentiation sequence established independently of this atlas (Jopling et al., 2010; González-Rosa et al., 2017). Published loss-of-function evidence further separates cell types with an established role in regeneration from bystanders. The resulting grouping, fixed before any model was trained, is given in Table 8; we use this prior information to validate underlying mechanism recovery, similar to how true flip rates (Table 3) were used for Ising.

Target dynamics. Our objective is a simulator that, once injured, passes through the stages of regeneration and returns to the uninjured state. The atlas carries no annotation of injury location, so we hand-annotated the approximate injured area in each of the 87 injured sections. A training trajectory starts from a real section at stage k and is transformed into stage $k + 1$; after 28 d it is transformed back into the uninjured state and remains there. Unlike the digit cycle of Section 3.1, the uninjured state is a fixed point, and the only way to leave it is a new injury: at every iteration an uninjured sample is wounded with probability $\alpha = 0.3$ and otherwise holds its state. A wound is generated around the ventricular apex, the site of amputation in the experiment, and the sample is then tasked with the transition to 6 h.

Time schedule. Each stage transition is allotted the same number of updates (15), whatever the physical interval between stages (6 h to 14 d). NCA time is therefore a pseudotime that indexes progress through the regeneration programme rather than elapsed time, analogous to pseudotime in single-cell trajectory inference. Time-invariance holds in the sense that matters here: the rule reads only the state, never a clock or stage label, and the schedule only fixes where along the rollout each marginal is read out. A schedule linear in physical time is impractical, since the intervals span a

factor of over 50 and the early transitions would receive a single update. The cost is that rates cannot be compared across stages in physical units, so we make no claims about durations.

Data per model. All models are trained on the same 105 sections; the held-out-heart model is trained on 102 (Section C.3). In contrast to Ising, where the transport baselines received thirty times more fields than REGIS, the data budget here is identical across models as it is limited by the dataset size.

C.2 REGIS ARCHITECTURE AND TRAINING

The state has $C = 21$ visible channels (20 annotation categories and damage), one alive channel and $C_h = 6$ hidden channels initialised to zero, 28 channels in total. Perception is a learned 3×3 depthwise convolution producing three maps per channel, and the per-pixel MLP of Equation (4) is $87 \to 128 \to 128 \to 28$ with ReLU activations and a zero-initialised output layer, so that the model begins as the identity. The 87 inputs are the 3×28 perception features and 3 noise channels ξ_t, redrawn at every update and every pixel. The generator has 31,892 parameters.

The update departs from Equation (5) in two ways. First, the alive mask gates the increment rather than the state,

$$z_{t+1} = z_t + \eta\, A_t \odot M_t \odot \mathrm{clip}(\Delta_t, -c, c), \qquad \eta = 0.1, \quad c = 10, \quad p_{\mathrm{fire}} = 0.5, \tag{13}$$

so that no channel moves by more than 1 per update. A wounded bin is not alive by construction, so clearing dead bins would erase the wound's own damage channel on the first update; freezing them instead lets the wound close inward from the living rim. A_t thresholds a 3×3 neighbourhood maximum of the alive channel at 0.05, and the alive channel is clamped to $[0, 1]$. Second, damage is constrained to be non-increasing, $d_{t+1} \le d_t$, and capped at the bin capacity 4, so the model can close a wound but never deepen it.

The discriminator is an R3GAN residual network (Huang et al., 2024) with base width 64, four stages of two blocks, cardinality 4 and expansion 2, conditioned on the target stage by class projection (Miyato & Koyama, 2018) with a 32-dimensional embedding, 538,816 trainable parameters in total. It reads 23 channels: the 21 visible ones, the alive channel, and a reference damage layer that the generator never sees. Training is adversarial, using the relativistic loss and zero-centred gradient penalty of Section 2.3; the full recipe is given in Table 7. We read out at a fixed endpoint and back-propagate through the whole 15-update segment.

Generator	31,892 parameters; 21 visible, 1 alive and 6 hidden channels
Perception / MLP	learned 3×3 depthwise / $87 \to 128 \to 128 \to 28$, ReLU
Noise	3 channels, redrawn every update
Firing	random, $p_{\mathrm{fire}} = 0.5$
Update clip c / step size η	10 / 0.1
Updates per stage	15 (pseudotime)
Discriminator	R3GAN residual, base width 64, projection-conditioned, 538,816 parameters
Optimiser	Adam, lr $= 2 \times 10^{-4}$, $\beta = (0,\ 0.99)$, gradient clip 1.0
Batch / crop	32 / full 48×48 canvas
Steps	100,000
Adversarial loss	relativistic, zero-centred gradient penalty $\gamma = 1$
Auxiliary losses	six terms; see *Auxiliary losses* below
Gradient window	full 15-update segment
Pool	2048 slots, lifetime ramped from 5 to 20 segments
Weights used	exponential moving average, decay 0.999 after a 1000-update warm-up
Compute	$\sim$335 ms per iteration on one H100

Table 7: **Training recipe** for REGIS on the heart data, shared by the non-local and time-conditioned controls.

Auxiliary losses. The problem-specific terms added to $\mathcal{L}_{\mathrm{NCA}}$ are below, all normalised per heart so a large section does not dominate a small one. Here a is the alive channel, d damage, m the section mask and κ_b the bin capacity.

Term	Definition	λ
Silhouette	$\|\max(a/s,\, d/s) - m\|_1$, saturation $s = 0.05$	1.0
Capacity	$\left\|\sum_c x_c - \kappa_b\right\|$ per bin, damage included	0.4
Heal	$\mathrm{relu}(m - a)$, stage-weighted $(0, 0, 0, .2, .4, .7, 1, 1)$	10.0
Zero	asymmetric penalty on $(t, \text{channel})$ pairs the cohort has at exactly zero	5.0
Homeostasis	L_1 between consecutive states, on uninjured-to-uninjured segments only	1.0
Hidden norm	confinement of h_t to radius $r_{\max} = 1.0$	1.0

These terms empirically improved stability and convergence, and each targets a specific failure mode. The need for auxiliary losses is also motivated by the low sample number in the dataset, so some of the biological priors, e.g., "don't change the heart shape excessively," are imposed more strongly than they would be by the discriminator.

The *silhouette* term counters the model's tendency to change shape. Heart size is a large technical bias across time points – global size varies with a coefficient of variation of 23% across sections, and the uninjured cohort is some 50% larger than the 12h one – so the adversarial signal alone rewards shrinking or spreading the tissue towards whichever size is easiest to pass off. Comparing the union of alive and damage against the section mask pins the footprint instead. The *capacity* term pushes in the same direction, at the level of density rather than outline: holding each bin at its own κ_b keeps the heart from thinning or hollowing out while retaining its silhouette.

The *heal* term primes clearance of the wound. Since a damaged bin is not alive, the neighbourhood alive gate switches the wound interior off: updates do not flow through it and gradients do not reach it, so regeneration can otherwise only proceed inward from the living rim. Penalising $\mathrm{relu}(m - a)$ drives the alive field back into the hole and re-opens that path, with weights that rise over the arc so the early time points remain free to hold a wound open.

The *zero* term checks channels that the cohort has at exactly zero and penalises any mass there, so the model is explicitly punished for producing a cell type never observed at that stage.

Removing terms one at a time (Table 12, one seed each) shows that the two shape-preserving terms matter most: without the silhouette term the cascade is largely lost (62% downstream retention), and without the capacity term the *A* vs *D* contrast falls to chance (51%). Removing any other term changes downstream retention and the group contrast by no more than the variation between REGIS seeds.

Pool and curriculum. The pool holds 2048 slots. Each iteration samples 32 and advances each by one segment of 15 updates, so a slot traverses the arc over successive iterations rather than within one. A slot's lifetime before reseeding ramps from 5 to 20 segments over the first 80k iterations, and the reseed distribution over stages shifts across the same window from uninjured-heavy (0.60) to more even (0.45).

Augmentation. Augmentation is geometric only: rotation uniform in $[-180°, 180°]$, independent horizontal and vertical flips at $p = 0.5$, and an elastic deformation ($\alpha = 50$, $\sigma = 5$), all with nearest-neighbour interpolation and zero fill, so compositions are never blended between bins.

C.3 Ablations and baseline models

Non-local control. To measure the effect of locality, we train a model identical to REGIS in every respect except its reach. The per-pixel MLP is replaced by a stack of six width-28 convolutions — the first four 3×3 with dilations $1, 2, 5, 9$, the last two 1×1 — followed by a zero-initialised 1×1 head, giving each update a 35-bin receptive field instead of REGIS's 3. Padding is zero rather than circular, because a heart sits on a canvas with an empty margin. At $31{,}528$ parameters it is marginally smaller than REGIS ($31{,}892$), so any deficit is not one of capacity. Three training seeds.

Time-conditioned control. REGIS takes the state as its only input: there is no clock, step index or stage label, so any difference in behaviour between 6 h and 28 d must be carried by the state itself, in its visible composition or its six hidden channels. Our hypothesis is that this constraint acts as a regulariser: a time-invariant rule cannot memorise what each stage looks like and must instead learn

dynamics that produce the stages in order, which we expect to cost accuracy on the marginals and to buy dependency between the states. To test it, we add an eight-dimensional one-hot encoding of the target stage of the current segment to the MLP input — the same variable the discriminator is conditioned on. If the hypothesis holds, this model should fit the marginals better while losing the cascade structure (Section C.4), since it can read the phase instead of inferring it. Three training seeds.

Held-out heart. The metrics of Section C.4 are otherwise computed on sections the model was trained on, so a model that memorised particular sections rather than learning dynamics could still score well. We therefore retrain from scratch with one fish removed entirely — T1_S1 and all three of its sections, leaving 102 — and evaluate it only on the three held-out sections, whose injuries were never seen in training. The held-out model fits the marginals comparably to the full model (median temporal r 0.900, against 0.848 ± 0.066 across the seven REGIS seeds) and retains the dependency structure: 11% downstream retention against $31 \pm 10\%$, and 83% of cross-group pairs ranked correctly ($p = 0.008$). One fish was held out, with one training seed.

MMSFM (Lee et al., 2025) learns a stochastic flow between consecutive marginals. A flow network and a score network — each a 9.3M-parameter U-Net with attention, 18.6M in total — are trained on sliding triplets of marginals, with minibatch optimal transport over 16 candidates providing the pairing within each batch and a spline reference path with diffusion scale 0.05. We train the closed loop uninjured $\to$ 6 h $\to \cdots \to$ 28 d $\to$ uninjured, eight transitions on a period-one time coordinate with integer Fourier harmonics, for $40{,}000$ updates with fused Adam at 10^{-4}, linear warm-up then linear decay, and an effective batch of 256. Samples are generated by integrating the SDE with a stochastic Runge–Kutta solver, 160 steps per transition. Two training seeds.

MMtSBM (Gravier et al., 2025) models each transition between consecutive marginals as a Schrödinger bridge, training a forward and a backward drift network by bridge matching. Each drift is a three-level U-Net without attention of 14.4M parameters, 28.9M in total, and samples are generated by Euler–Maruyama integration, 120 steps per bridge. We train with AdamW at 2×10^{-4}, batch size 64 and diffusion scale 1.2, for $50{,}000$ warm-up updates followed by three iterative Markovian fitting iterations of $12{,}500$ updates in each direction. We report the checkpoint before the fitting iterations, since they degraded the cascade on this cohort; because this selection uses the evaluation criterion, it favours the baseline. One training seed.

Shared settings. MMSFM and MMtSBM are trained on the eight stages plus the return to the uninjured state, nine marginals in total, with 9–15 sections per marginal — two to three orders of magnitude fewer independent fields per marginal than in the Ising comparison. Both define a trajectory as a chain of transitions between consecutive marginals, with loop closure (28 d $\to$ uninjured) imposed as a training transition, as it is for REGIS. The baselines receive none of the auxiliary losses, which act on the alive and damage mechanics specific to the NCA.

C.4 Evaluation metrics

Trajectory fit. The variability documented in Table 9 and Figure 16 limits what a goodness-of-fit score can mean here. Absolute compartment sizes have coefficients of variation of 34–88% across sections, and thin or rare structures are missed depending on where the section falls: valves, for instance, are absent from every section at two of the eight stages. A strict marginal-matching score would mostly measure which sections happened to be cut, so we instead ask whether each cell type rises and falls when the cohort does. For cell type c, let $\bar{o}_c(k)$ be the cohort-mean abundance at stage k, and $\bar{y}_c(k)$ the mean abundance over simulated rollouts at the update corresponding to stage k. The temporal correlation is

$$r_c = \mathrm{corr}_k\big(\bar{y}_c(k),\, \bar{o}_c(k)\big), \tag{14}$$

taken over the measured stages, and we report the median of r_c over all cell types and over the five BZm states.

Knock-out effect by functional group. Lacking a ground-truth trajectory, the remaining metrics exploit the fact that a simulator can be perturbed. Both zero a chosen cell-type channel at every update of the rollout — the *in silico* analogue of a loss-of-function experiment — and measure what

happens to the rest of the trajectory. The model is stochastic, so the comparison is made within each section: we run the same section three times from the same initial state — twice unperturbed, to measure the model's own run-to-run variability, and once with cell type c held at zero throughout — and ask how much larger the knock-out difference is than that variability (Figure 18).

Let $x^{(i)}, x'^{(i)}, x_c^{(i)} \in \mathbb{R}^{C\times P}$ be section i's settled state unperturbed, unperturbed with a different noise draw, and with c knocked out, over C channels and P bins. Let Ω_i be the set of occupied bins of section i, and $\mathcal{C}_c = \{1,\dots,C\} \setminus \{c\}$ the channels other than the knocked-out one. Then

$$\Delta^{(i)}(u,v) \;=\; \frac{1}{|\mathcal{C}_c|\,|\Omega_i|} \sum_{c'\in\mathcal{C}_c} \sum_{p\in\Omega_i} \left|u_{c'p} - v_{c'p}\right| \tag{15}$$

is the mean absolute difference over those bins and channels. Excluding the knocked-out channel means the statistic measures the effect on the *remaining* types rather than the trivial fact that the zeroed channel is zero. The knock-out effect is

$$\rho(c) \;=\; \frac{1}{N}\sum_{i=1}^{N} \frac{\Delta^{(i)}\left(x^{(i)}, x_c^{(i)}\right)}{\Delta^{(i)}\left(x^{(i)}, x'^{(i)}\right)}, \tag{16}$$

so each section is its own control, and $\rho = 1$ means a perturbation indistinguishable from rerunning the model. States are clipped at zero before differences are taken. We use $N = 100$ wounded sections per condition.

Groups are compared on ρ, one value per cell type, by a one-sided Mann–Whitney test in the pre-specified direction. We report the test statistic as the share of cross-group pairs, $U/(n_1 n_2)$, in which the cell type from the first group has the larger effect, ties counting a half, so that 50% is the null and 100% is complete separation. Two contrasts were fixed before any model was trained, and both are reported: *A+B* vs *C+D* (9 against 10 types, 90 pairs) and *A* vs *D* (5 against 7, 35 pairs), with groups as in Table 8. Intervals on this share are 95% percentile bootstraps over 4000 resamples of cell types within each side. Across seeds we report mean $\pm$ sample s.d. and a one-sided sign test against 50% (Figure 19).

Cascade retention. Any model that fits the marginals reproduces the succession of the five BZm states; what perturbation adds is the *dependency* between them. If the model has learnt the cascade as a chain, zeroing a state should suppress the states downstream of it, whereas a model that has learnt only when each state appears should be unaffected. We zero each state b in turn and, for every other state j, record

$$R(j \mid b) = 100 \times \frac{\max_t \bar{y}_j^{(-b)}(t)}{\max_t \bar{y}_j(t)}, \tag{17}$$

the peak abundance of j with b blocked, as a percentage of its unperturbed peak. We pool R over pairs in which j follows b (*downstream*), in which j precedes b (*upstream*), and in which b is a control cell type (*control*). A directional cascade shows low downstream retention, with upstream and control retention near 100%.

Controls and caveats. Two controls moderate the interpretation. An unwounded rollout, in which no cascade state should rise, establishes that the response is triggered by injury rather than by the passage of time. Knock-outs of cell types that prior biology does not implicate in the wound response — valves, smooth muscle, atrial myocardium and blood — should leave the cascade intact. Zeroing any channel drives the state off-distribution, so without this control a general derailment would be indistinguishable from a captured dependency. For the same reason, upstream retention near 100% shows that a knock-out does not derail the rollout, but it is not in itself evidence of a dependency, since any forward simulator leaves earlier states unchanged.

Two further caveats apply. The temporal correlation is insensitive to amplitude, so a model can score well while under- or over-producing a cell type. The motivation is: high variability in wound and heart sizes impacts absolute amplitude measurements. Second, blood cell types may partly mark vasculature, whose function is required for regeneration (Marín-Juez et al., 2016), so their use as controls carries the caveat noted in Table 8. Thus in addition to A+B vs C+D we also made a direct comparison of A vs D which excludes blood (Table 10). This test doesn't change our conclusions.

C.5 EVALUATING THE TRANSPORT BASELINES

MMSFM and MMtSBM are defined as chains of transitions between consecutive marginals, so their evaluation differs from that of REGIS in three respects. First, a trajectory is read out at the end of each transition, which gives the measured stages directly; temporal correlations therefore use the same eight stages as REGIS. Second, a knock-out is applied by holding the chosen channel at zero during inference, as for REGIS, and cell composition is then recorded in the same way. Third, MMtSBM is evaluated in two modes: *chained*, the free rollout used for every other model, and *one-step*, which restarts each transition from a real section of the preceding stage. The one-step model fits the marginals comparably to REGIS (median $r = 0.842$), confirming that each learned bridge works, but chaining the bridges does not ($r = 0.495$). Since the chained rollout does not reproduce the BZm cascade, and a knock-out cannot propagate across transitions in the one-step mode, knock-out metrics are not reported for MMtSBM.

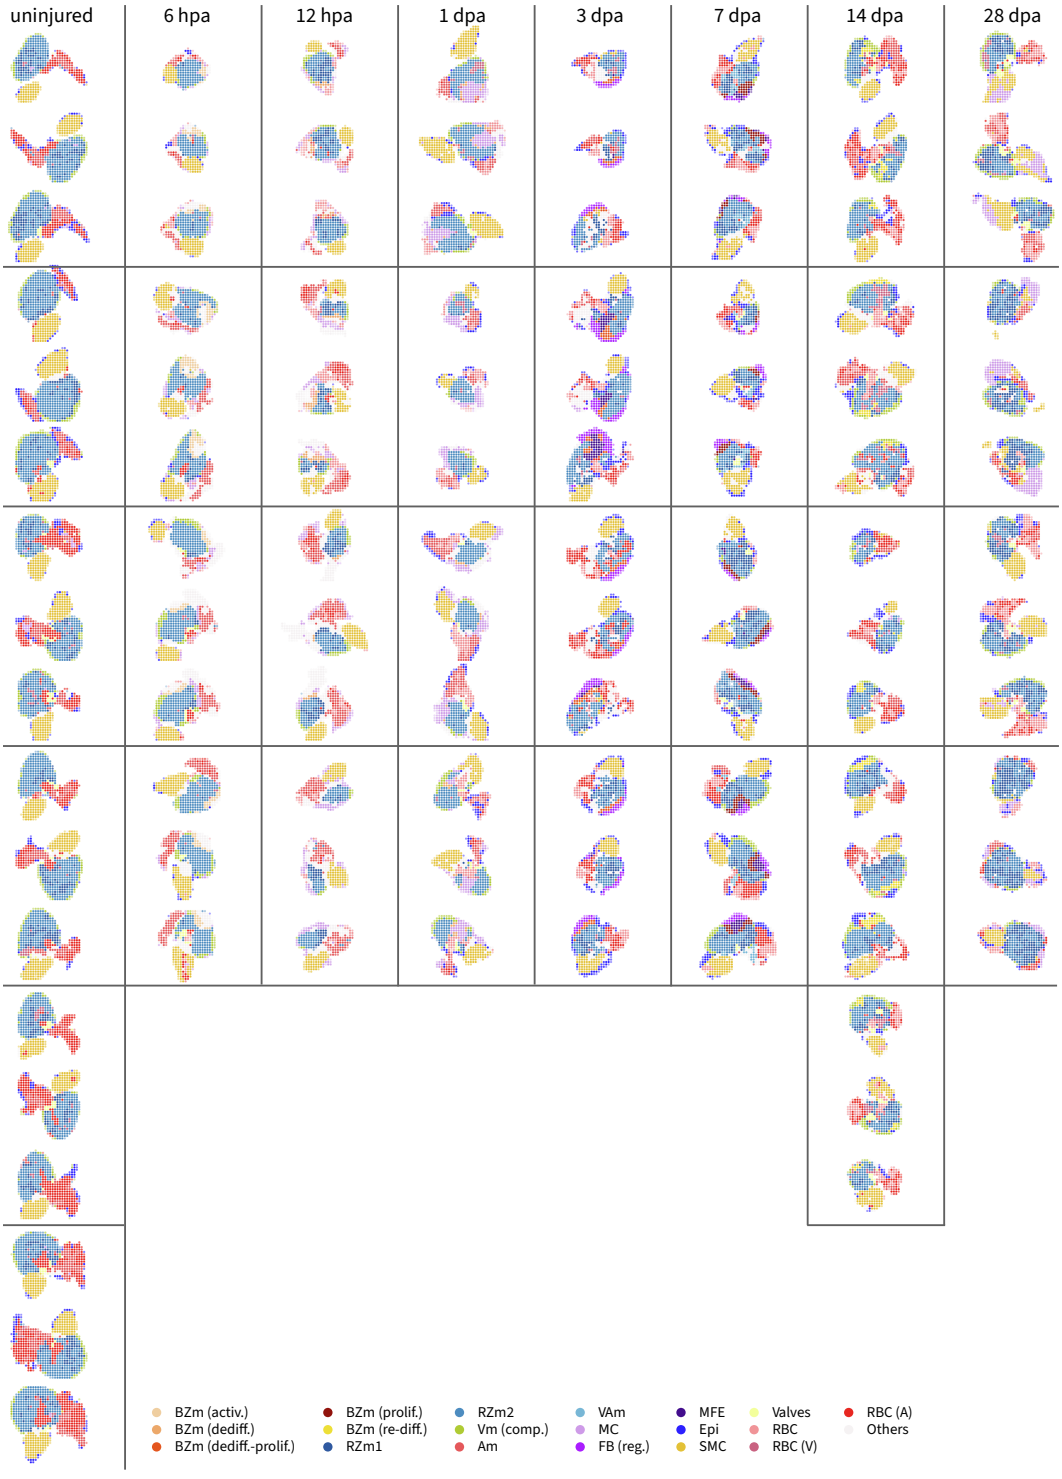

Figure 15: **Observed zebrafish heart sections across sampling stages.** Columns show the eight stages in the atlas; colours identify the annotated cell types.

Table 8: **A-priori grouping of the 19 cell types.** Groups *A* and *B* form the implicated side of the contrast, *C* and *D* the unimplicated side. The source atlas (Li et al., 2025) only defines the cell types; the evidence for the assignment comes from other sources. The assignment was fixed before any model results were seen. Absence of evidence is not evidence of absence; the contrast therefore tests only the weaker hypothesis that cell types with established roles in regeneration perturb the outcome detectably more than those without. The erythrocyte populations (RBC) may partly mark vasculature, which has been shown to also play a role in the regeneration (Marín-Juez et al., 2016); endothelium is not annotated separately here, so RBC are put into the unimplicated group with that caveat.

Cell type	Full name	Group	Implicated	Support
BZm (activ.)	Border-zone myocardium, activated	*A*	Yes	(Jopling et al., 2010; Kikuchi et al., 2010)
BZm (dediff.)	Border-zone myocardium, dedifferentiating	*A*	Yes	(Jopling et al., 2010; Kikuchi et al., 2010)
BZm (dediff.-prolif.)	Border-zone myocardium, dedifferentiating and proliferating	*A*	Yes	(Jopling et al., 2010; Honkoop et al., 2019)
BZm (prolif.)	Border-zone myocardium, proliferating	*A*	Yes	(Jopling et al., 2010; Honkoop et al., 2019)
BZm (re-diff.)	Border-zone myocardium, re-differentiating	*A*	Yes	(Jopling et al., 2010; Sallin et al., 2015)
MC	Macrophages	*B*	Yes	(Lai et al., 2017; Wei et al., 2023)
FB (reg.)	Fibroblasts, pro-regenerative	*B*	Yes	(Hu et al., 2022; Sánchez-Iranzo et al., 2018)
MFE	Mixed immune, fibroblast and endocardial	*B*	Yes	(Lai et al., 2017; Hu et al., 2022; Kikuchi et al., 2011)
Epi	Epicardium	*B*	Yes	(Kikuchi et al., 2011; Sun et al., 2022)
RBC	Erythrocytes, wound	*C*	No	—
RBC (V)	Erythrocytes, ventricular	*C*	No	—
RBC (A)	Erythrocytes, atrial	*C*	No	—
RZm1	Remote-zone myocardium 1	*D*	No	—
RZm2	Remote-zone myocardium 2	*D*	No	—
Vm (comp.)	Ventricular compact myocardium	*D*	No	—
Am	Atrial myocardium	*D*	No	—
VAm	Ventricular and atrial shared myocardium	*D*	No	—
SMC	Smooth muscle, bulbus arteriosus	*D*	No	—
Valves	Cardiac valves	*D*	No	—

Table 9: Sample diversity of the Stereo-seq data (105 sections from 35 hearts). Size is the summed channel value over the section's occupied grid bins (cell units). Chambers include the blood in their own lumen; residual collects VAm, unqualified RBC and the reactive types (MC, FB (reg.), MFE, Epi), so the rows sum to the global size. CV is the coefficient of variation across sections.

	Absolute (cell units)			Relative to global	
Compartment	Mean	SD	CV	Mean	CV
Global	1517	346	22.8%	–	–
Ventricle	613	208	33.9%	0.399	20.3%
Atrium	229	125	54.5%	0.146	38.6%
Bulb (SMC)	201	72	35.7%	0.131	25.6%
Valves	20	18	87.5%	0.013	78.6%
Unassigned	166	83	49.8%	0.114	47.4%
Residual	288	118	41.0%	0.197	39.9%

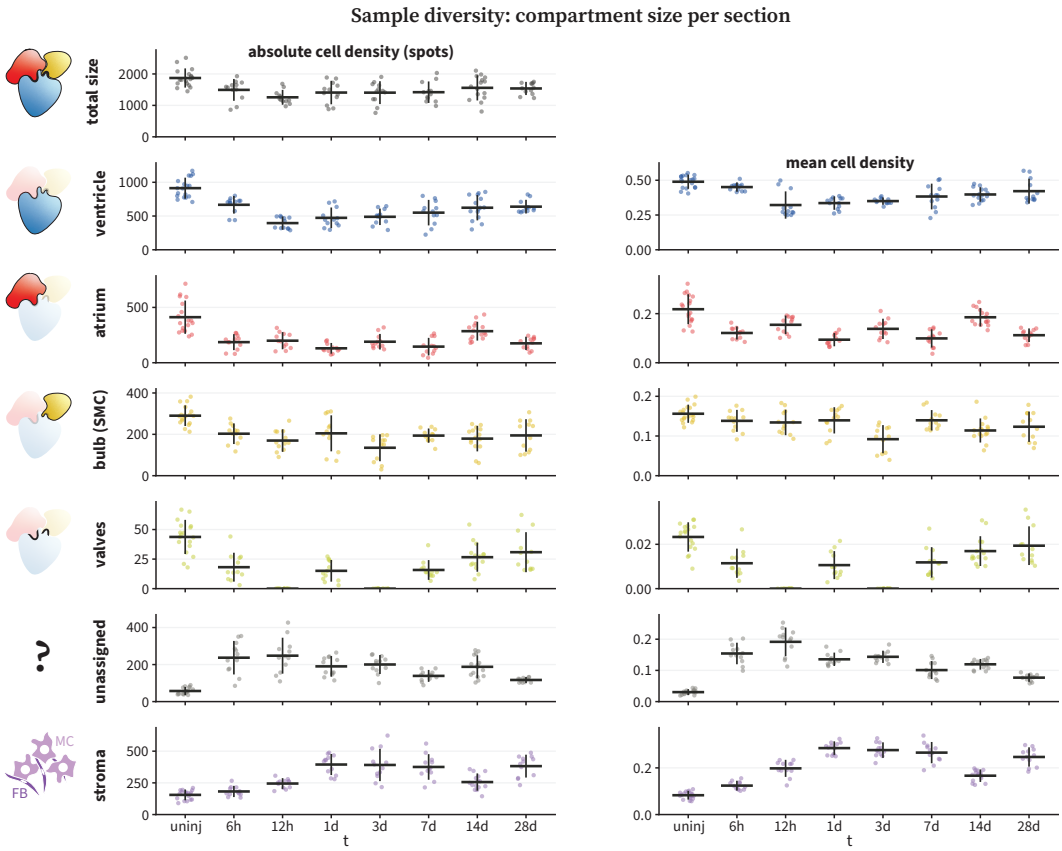

Figure 16: **Variation in heart size and composition across sampling stages.** Points show individual sections. Absolute compartment sizes are shown on the left and mean cell densities on the right. Variation in the absolute and relative sizes of atrium, bulb and valves partly reflects natural differences in heart size but mostly depends on where the cross-section falls. The stromal cells are expected to increase during regeneration. The Unassigned type is a limitation of the dataset annotations as some of the transcriptional signal was not clearly attributed to a definitive cell type.

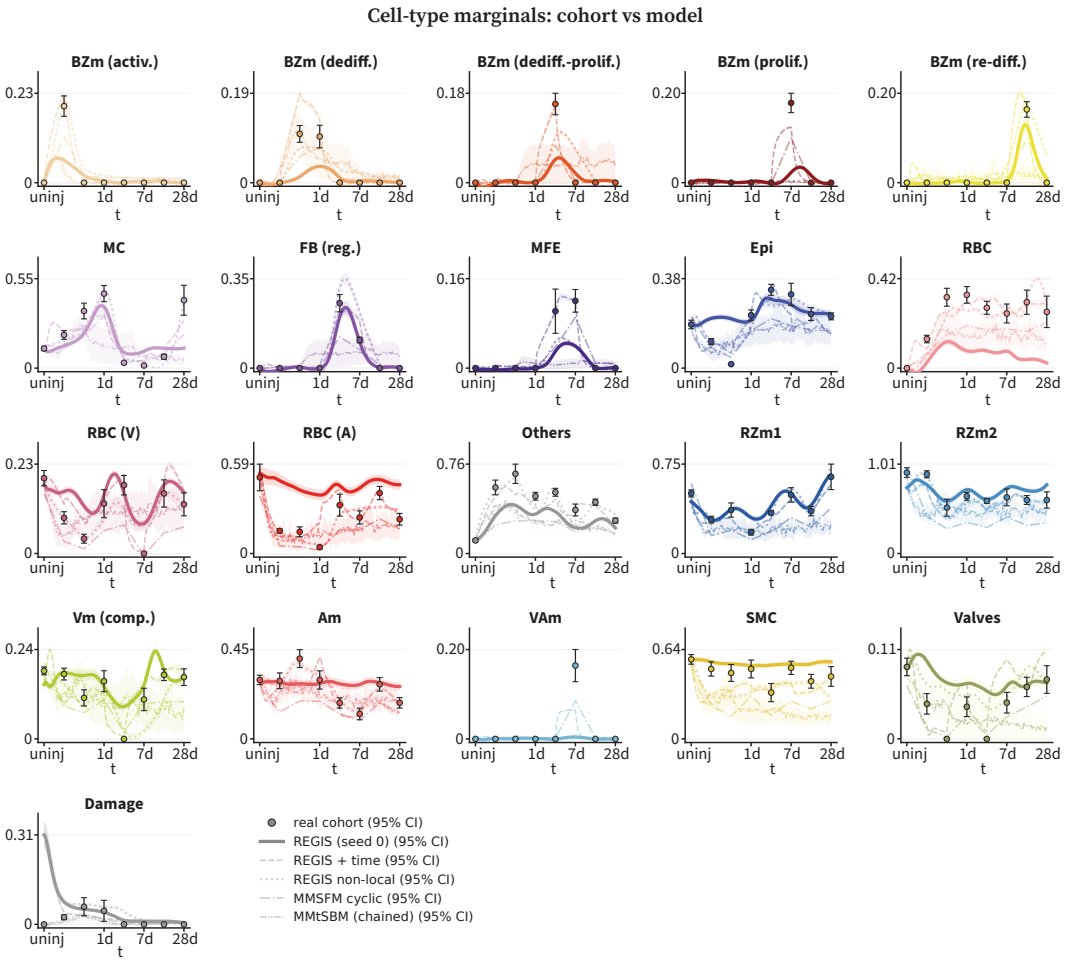

Figure 17: **Observed and simulated cell-type profiles during regeneration.** Points show cohort means, solid curves show REGIS, other shaded curves show time-conditioned ablation (REGIS + time), non-local ablation (REGIS non-local), as well as baselines MMSFM and MMtSBM. Error bars and shaded bands indicate 95% confidence intervals. Uninjured is the starting state for all models. REGIS and its ablations are shown from the *in silico* injury onward, so damage is highest at the first point; MMtSBM is shown as the chained rollout.

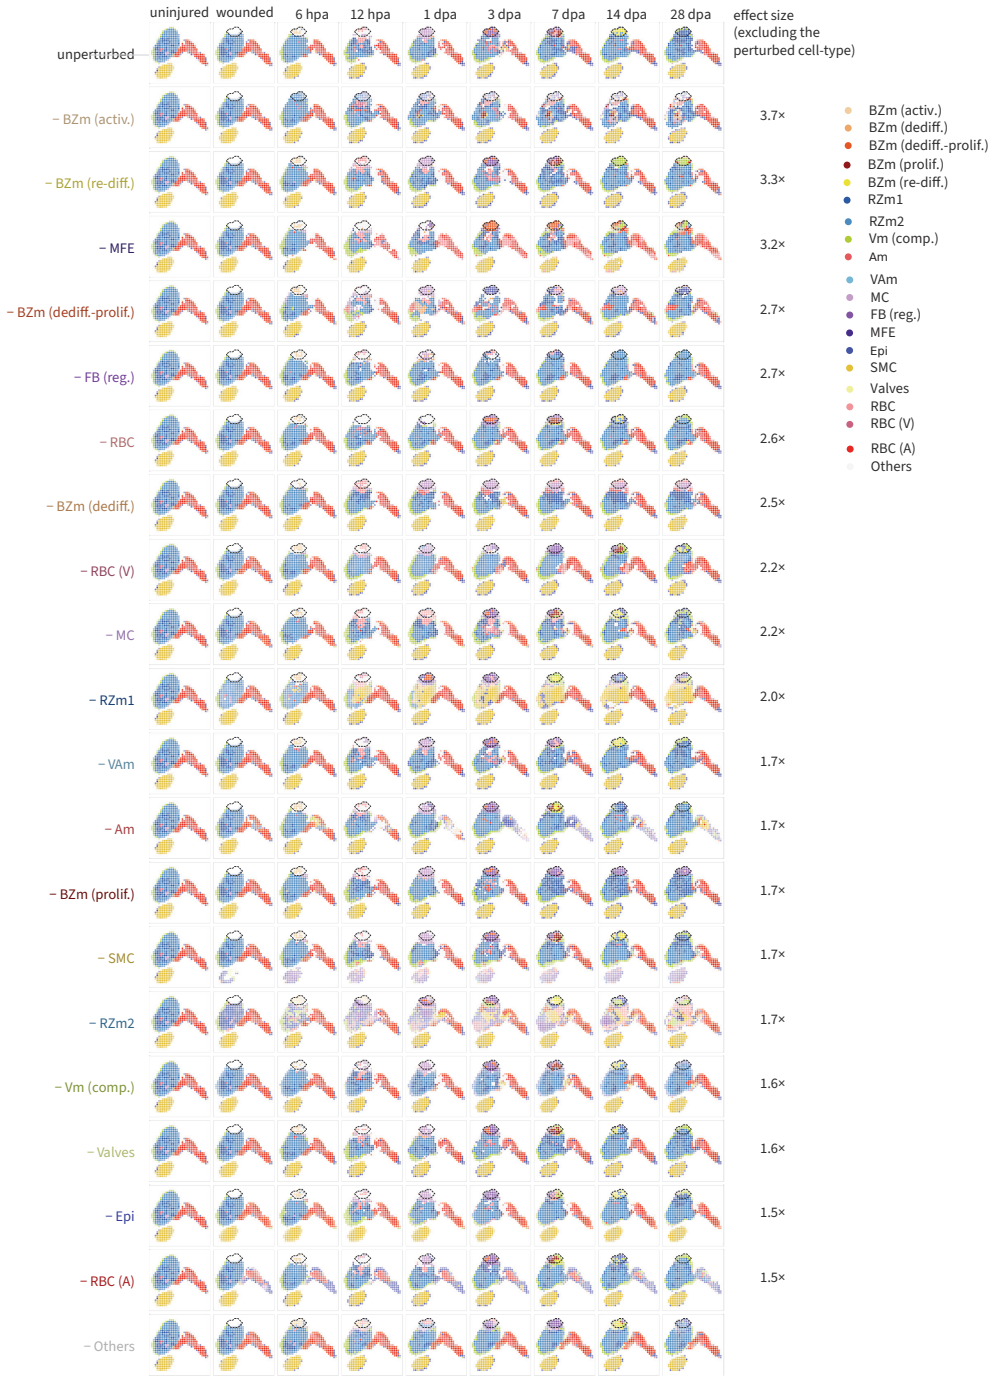

Figure 18: **Simulated wounded heart regenerating under each single-domain knock-out.** Columns are the uninjured seed, the synthetic wound at $t = 0$, and the seven measured timepoints; the dashed outline marks the wound footprint. Each bin is coloured by its dominant domain and its opacity set by density, under one display rule pooled over all panels. Rows are the unperturbed rollout followed by each knock-out, ordered by measured effect (right-hand value). The effect is the mean distance between the unperturbed and perturbed states at the last measured marginal, excluding the perturbed channel itself; see Section C.4 for the definition.

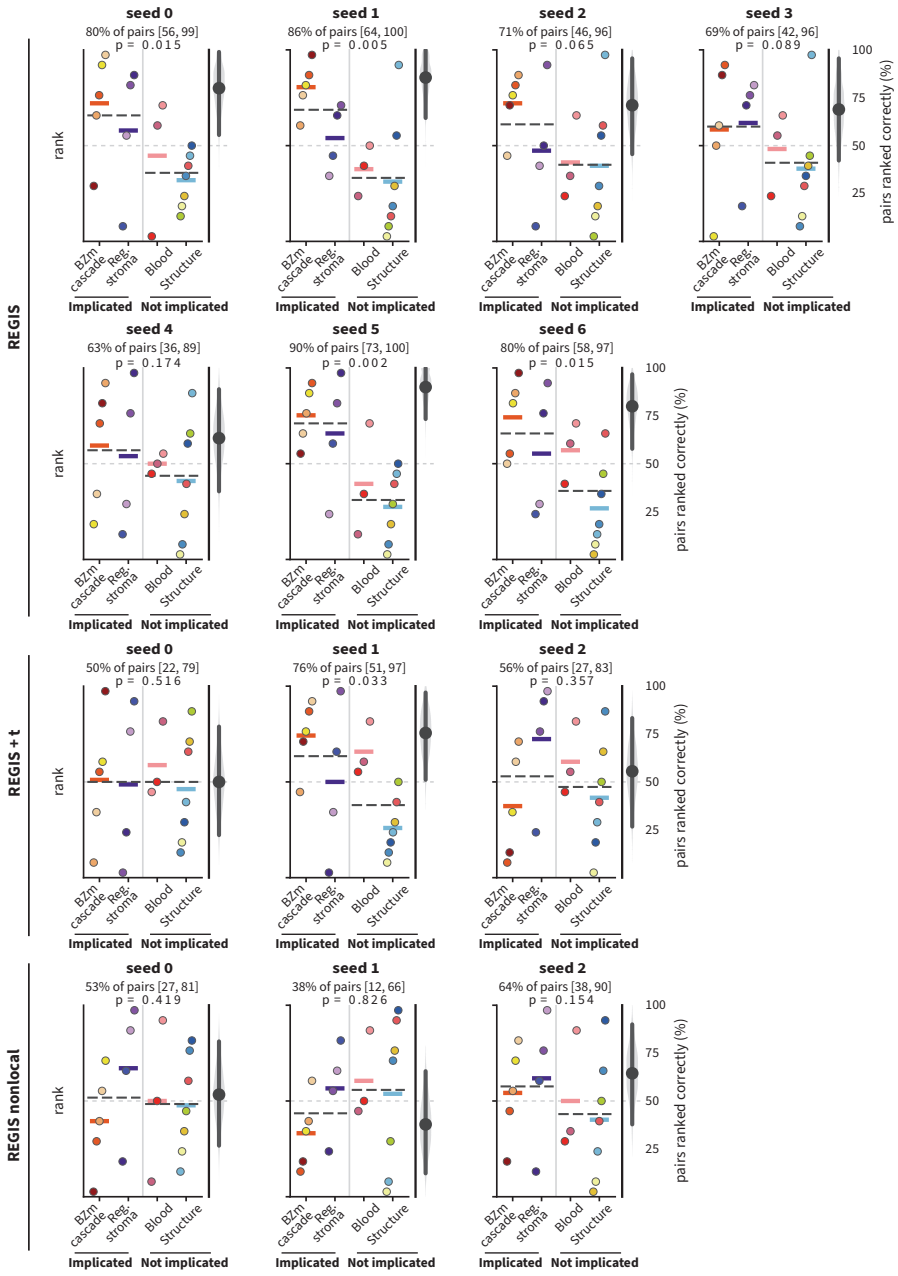

Figure 19: **Rank of each cell type's knock-out effect, by a-priori group.** Points are the 19 cell types by rank of knock-out effect (cell type removed, rollout repeated, shift in the remaining cell types at the settled end state, in units of the model's own run-to-run variability). Coloured lines are group means, dashes the means of the two tested sides. Groups are fixed a priori from literature external to the annotation as detailed in Table 8. Right-hand axis: share of implicated-vs-unimplicated pairs ranked correctly (50% = chance) with its bootstrap distribution and 95% interval; the value and one-sided Mann–Whitney p are given above each panel. Across the seven REGIS seeds, the share of pairs ranked in the pre-specified direction is $77 \pm 9\%$ (sign test across seeds, one-sided $p = 0.008$; Table 10).

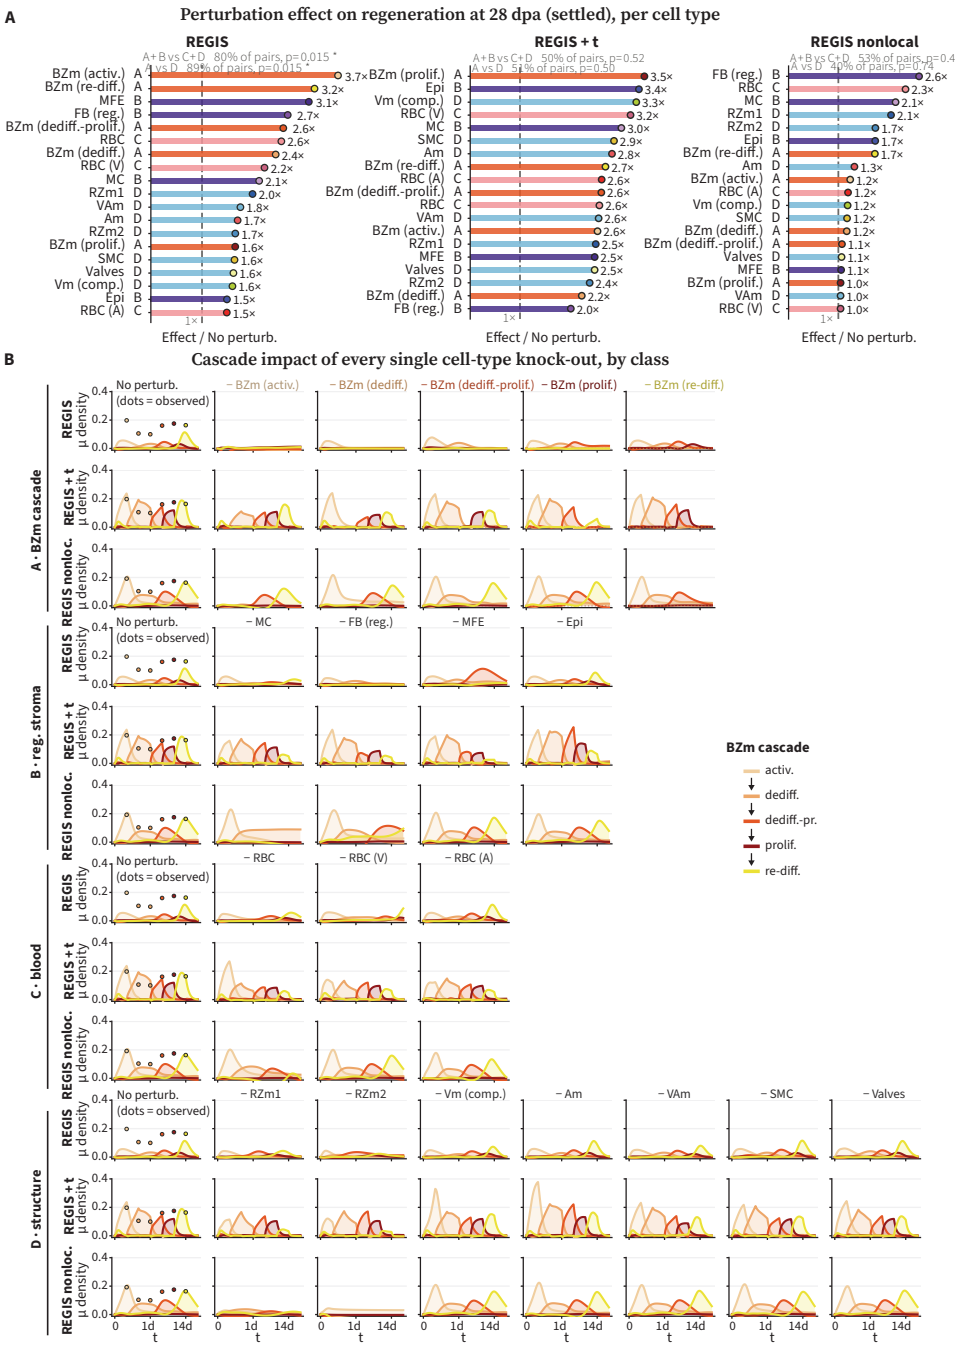

Figure 20: **Cell-type knockout effects without temporal invariance or spatial locality. A** Effects of individual knockouts in REGIS, REGIS + time and the non-local control, relative to variation between unperturbed rollouts; seed 0 of each model is shown; other seeds are in Figure 19 and Table 10. **B** Cardiomyocyte-state trajectories after individual cell-type knockouts, grouped by functional category. Dots in the unperturbed panels show observed cohort means.

Table 10: **A-priori group recovery.** Knock-out effects ρ (Section C.4) compared between the pre-specified groups of Table 8: *A* regeneration programme, *B* regulatory stroma, *C* blood, *D* constitutive structure. Pairs: the share of cross-group pairs, $U/(n_1 n_2)$, in which the cell type from the first group has the larger effect (50% is the null, 100% complete separation); p is reported for the one-sided Mann–Whitney in the pre-specified direction. The italic line under each family gives the mean $\pm$ s.d. over its seeds and a one-sided sign test across seeds against 50%.

Model	Seed	A+B vs C+D		A vs D	
		pairs	p	pairs	p
REGIS (ours)	seed 0	80%	0.015	89%	0.015
	seed 1	86%	0.005	89%	0.015
	seed 2	71%	0.065	80%	0.053
	seed 3	69%	0.089	69%	0.172
	seed 4	63%	0.174	69%	0.172
	seed 5	90%	0.002	100%	0.001
	seed 6	80%	0.015	94%	0.005
across seeds		$\mathbf{77 \pm 9\%}$	**0.008**	$\mathbf{84 \pm 11\%}$	**0.008**
REGIS + time	seed 0	50%	0.516	51%	0.500
	seed 1	76%	0.033	97%	0.003
	seed 2	56%	0.357	46%	0.622
across seeds		$60 \pm 11\%$	0.500	$65 \pm 23\%$	0.500
REGIS non-local	seed 0	53%	0.419	40%	0.735
	seed 1	38%	0.826	37%	0.784
	seed 2	64%	0.154	66%	0.216
across seeds		$52 \pm 11\%$	0.500	$48 \pm 13\%$	0.875
REGIS held-out heart	T1_S1	83%	0.008	83%	0.037
MMSFM (cyclic)	seed 0	21%	0.985	23%	0.947
	seed 1	21%	0.985	17%	0.976
across seeds		$21 \pm 0\%$	1.000	$20 \pm 4\%$	1.000

Table 11: **Trajectory fidelity and cascade retention.** Temporal r and cascade retention as defined in Section C.4. Temporal r is the median over cell types, for all cell types and for the five BZm states. Rows are individual training seeds; the italic line under each family is the mean $\pm$ s.d. over its seeds. A directional cascade shows low downstream retention, with upstream and control retention near 100%. MMtSBM is shown in its *one-step* and *chained* modes (Section C.5); knock-out metrics are not reported due to the poor fit of the *chained* model.

Model	Seed	Median temporal r		Retained after one knock-out (%)		
		all cell types	BZm cascade	downstream	upstream	control
REGIS (ours)	seed 0	0.823	0.924	13	107	100
	seed 1	0.745	0.611	31	115	93
	seed 2	0.926	0.930	27	102	63
	seed 3	0.834	0.934	31	90	72
	seed 4	0.916	0.975	47	97	91
	seed 5	0.801	0.814	36	78	97
	seed 6	0.889	0.953	33	103	91
mean $\pm$ s.d.		0.848 ± 0.066	0.877 ± 0.128	$\mathbf{31 \pm 10}$	99 ± 12	87 ± 14
REGIS + time	seed 0	0.960	1.000	67	102	94
	seed 1	0.954	0.999	55	107	100
	seed 2	0.953	1.000	67	102	87
mean $\pm$ s.d.		$\mathbf{0.956 \pm 0.004}$	$\mathbf{0.999 \pm 0.000}$	63 ± 7	104 ± 3	94 ± 7
REGIS non-local	seed 0	0.918	0.935	79	99	95
	seed 1	0.843	0.787	71	101	95
	seed 2	0.870	0.864	73	99	97
mean $\pm$ s.d.		0.877 ± 0.038	0.862 ± 0.074	74 ± 4	100 ± 1	96 ± 1
REGIS held-out heart	T1_S1	0.900	0.932	11	97	86
MMSFM	seed 0	0.841	0.993	100	102	101
	seed 1	0.828	0.994	102	101	103
mean $\pm$ s.d.		0.835 ± 0.007	0.993 ± 0.001	101 ± 1	102 ± 1	102 ± 1
MMtSBM, one-step	seed 0	0.842	0.922	—	—	—
MMtSBM, chained	seed 0	0.495	0.511	—	—	—

Table 12: **Auxiliary-loss ablations.** Each row removes one auxiliary term and retrains from scratch with a single seed; all other terms, the architecture, the schedule and the data are unchanged. The homeostasis and hidden-norm terms were not ablated. The first row is the REGIS seed with the lowest downstream retention (seed 0 in Table 11); across the seven REGIS seeds, downstream retention ranges over 13–47% and the *A+B* vs *C+D* contrast over 63–90%. Metrics as in Section C.4; $\pm$ in the retention columns is over (blocked, affected) state pairs within one run.

	Median temporal r		Cascade retention (%)			A+B vs C+D		A vs D
Terms active	all	BZm	down	up	control	pairs	p	pairs
All six (REGIS)	0.823	0.924	$\mathbf{13 \pm 14}$	107 ± 14	100 ± 9	**80%**	**0.015**	89%
− silhouette	0.792	0.677	62 ± 40	109 ± 17	89 ± 40	66%	0.135	63%
− capacity	0.878	0.978	30 ± 27	82 ± 15	88 ± 10	61%	0.219	51%
− heal	0.898	0.993	27 ± 25	99 ± 28	59 ± 36	77%	0.028	69%
− zero	0.909	0.934	29 ± 11	95 ± 27	67 ± 27	90%	0.002	91%

D MISCELLANEOUS

D.1 REGIS NAMING

The name REGIS refers to Regis III, the fictional planet in Stanisław Lem's novel *The Invincible* (*Niezwyciężony*). Regis III is a home world for tiny, insect-like micromachines (*muszki*). Despite their individual simplicity, these machines evolved the ability to assemble into large swarms (*chmura*), whose collective behaviour exhibits remarkable complexity and dynamics. The authors wholeheartedly recommend reading this novel to all interested in science fiction and alternative evolution.