Evolutionary Discovery of the Behavioral Repertoire of Self-Propelled Particle Models

Roujia Wen[a], Eric Bonabeau[b,*]

[a] *University of Cambridge, UK (currently with Tracebit, London, UK)*
[b] *School of Complex Adaptive Systems, Arizona State University, Tempe, AZ, USA*

Abstract

Self-propelled-particle models can produce many collective states, but standard phase diagrams usually examine only one or two control parameters at a time. We use quality-diversity search to map the behavior of a nine-parameter model and learn a two-dimensional representation from translation- and rotation-invariant trajectory features. On an independent test set, the learned map preserves behavioral neighborhoods better than two-dimensional PCA. Local response measurements show that only two to four parameter combinations are resolved as stiff in each regime, and the stiff combinations change from one regime to another. In a two-species extension, reciprocal interactions are sufficient for coherent translation, whereas persistent rotation in the tested regime disappears when antisymmetric couplings are removed. We also locate regions where small parameter changes produce large changes in trajectory descriptors. These high-sensitivity regions do not always coincide with order-parameter transitions, and which regions are detected depends on the descriptor. The approach provides a practical way to map behavior and parameter sensitivity in high-dimensional stochastic active-matter models.

1. Introduction

Collective motion is a canonical example of organization far from equilibrium. Local interactions among self-propelled units can generate coherent flocks, rotating mills, compact swarms, bands, lanes, density waves, segregated mixtures, and disordered gases. The Vicsek model established that alignment and noise alone can produce collective order in two dimensions [1], while cohesive extensions introduced attraction and repulsion and generated aggregation and milling [2,3]. These models now occupy a central place in active-matter physics and in quantitative descriptions of animal groups, cell collectives, active colloids, and robotic swarms [4-8]. The ingredients used here are deliberately phenomenological: local alignment represents orientational or social coupling, short-range repulsion and longer-range attraction represent exclusion and cohesion, a taxis term represents response to an external cue, and angular noise represents unresolved directional variability. We do not calibrate the model to a single material or organism; it is used as a finite-system test bed for exploring how several physically interpretable controls interact.

The usual way to characterize such a model is to vary one or two parameters while fixing the rest. This works when the important control plane is known in advance. It is less reliable when

*Corresponding author: ebonabea@asu.edu

density, speed, inertia, interaction range and strength, adhesion, alignment, and noise can all vary independently. A behavior confined to a narrow or oblique region of parameter space can be missed by coordinate sweeps.

Order parameter optimization addresses a different question. One can maximize alignment, milling, or clustering and recover a parameter set that produces the chosen behavior. This shows that the behavior is reachable, but it does not show what other behaviors are reachable or how they are related.

Quality-diversity (QD) optimization maps many behaviors in a single search. Behavior measures define the coordinates of an archive, and each occupied cell stores a stable parameter set that realizes that behavior [9-11]. Covariance-adapting variants such as CMA-ME and CMA-MAE improve search efficiency in continuous parameter spaces [12,13]. The result is a repertoire rather than a single optimum.

Predefined order parameters are useful but incomplete. Alignment, milling, and clustering can assign similar values to states that differ in density modulation, local packing, spatial anisotropy, persistence, or wall localization. Learned descriptors provide a second view by extracting coordinates directly from trajectory features [14,15]. Related methods have been used to identify phases of matter from configurations [16-18].

We also ask which parameter combinations control behavior locally. Multiparameter models often have a few directions to which the output is sensitive and many directions to which it is weakly sensitive [19-22]. We estimate this local sensitivity directly from behavioral response Jacobians and distinguish it from the covariance of parameter vectors stored in a QD archive.

Finally, we extend the model to two particle types. Reciprocal and non-reciprocal cross-species interactions can produce similar observed motion, so behavior alone does not identify mechanism. We therefore compare matched reciprocal and non-reciprocal systems, use long trajectories, and analyze species-resolved observables [23-30].

The paper proceeds from the simplest case to the most general. We first map and validate the single-species repertoire. We then expand the density range, learn a trajectory-based representation, and measure local parameter sensitivity. We next introduce two particle types and isolate the effect of antisymmetric couplings. Finally, we search the full parameter space for high-sensitivity regions where small parameter changes reorganize behavior.

The main contributions are:

1. a direct QD map and multi-seed validation of the classical single-species repertoire;
2. a density-expanded trajectory corpus and an independently tested learned behavior representation;
3. a local parameter-sensitivity analysis with an explicit stochastic noise floor;
4. matched reciprocal and non-reciprocal tests in a two-species extension; and
5. a full-space search for high-sensitivity regions, including tests of how the chosen descriptor changes what is detected.

The active-matter phenomena themselves are established: the model reproduces the Vicsek noise transition [1,2], milling and cohesive aggregation [2,3], high-density hexatic order [31], and

non-reciprocity-driven chirality [28]. The contribution here is methodological. We ask how much of a high-dimensional behavioral repertoire can be mapped, what a learned trajectory representation adds to classical order parameters, which parameter combinations are locally important, and how the chosen representation affects the boundaries that are found.

2. Model and observables

2.1. Single-species dynamics

Model overview. At each synchronous time step, every particle first forms a deterministic drive direction from four ingredients: persistence of its previous heading, alignment with nearby headings, an optional fixed taxis cue, and pairwise attraction/repulsion. An independent angular perturbation is then added to that direction, and the particle advances at constant speed. The Figure 1 schematic separates these contributions before the formal definitions.

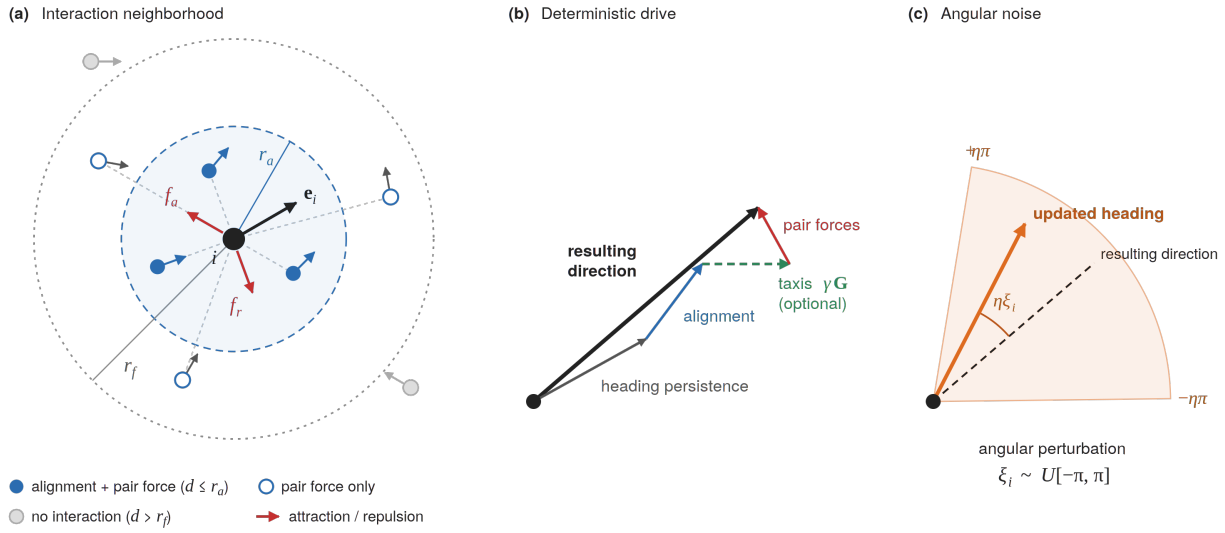


Figure 1. Model schematic. A focal particle combines heading persistence, local alignment inside r_a, pair interactions up to r_1, and an optional taxis cue. The angular variable xi_i is a direction perturbation rather than a random translational force; eta sets its amplitude.

We consider N identical particles in a square arena of side length S. Particle i has position $\mathbf{x}_i(t)$ and heading angle $\theta_i(t)$, with unit heading vector

$$\mathbf{u}_i(t) = (\cos\theta_i(t),\ \sin\theta_i(t))\,.$$

The speed v_0 is constant. Positions are updated synchronously,

$$\mathbf{x}_i(t+1) = \mathbf{x}_i(t) + v_0\mathbf{u}_i(t+1).$$

The new heading is the argument of a drive vector containing heading inertia, local alignment, an optional external gradient, and pairwise force, followed by angular noise:

$$\theta_i(t+1) = \arg\left[\epsilon\mathbf{u}_i(t) + \alpha\mathbf{a}_i(t) + \gamma\mathbf{G} + f_0\mathbf{F}_i(t)\right] + \eta\xi_i(t),$$

Here ε is a dimensionless heading-persistence weight: increasing ε makes the new direction more strongly inherit the previous heading, so it should not be read as a mechanical moment of

inertia. The coefficient α weights local alignment, γG is an optional fixed taxis vector, and f_0 weights the pair force. The random variable $\xi_i(t)$ is an angular deviate, independently drawn from a uniform distribution on [-π, π] at every particle and step; because it enters as $\eta\xi_i(t)$, η sets the half-width of the angular perturbation, from zero noise at $\eta = 0$ to a full $\pm\pi$ range at $\eta = 1$. Uniform angular noise is used because the model is a discrete-time Vicsek-type heading model, not a Langevin model of thermal translational forcing. The main experiments set $\gamma = 0$ so that collective alignment cannot be imposed trivially by a global field.

The alignment direction is the normalized sum of headings within radius r_a,

$$\mathbf{a}_i(t) = \frac{\mathbf{u}_i(t) + \sum_{j\neq i} \mathbf{1}(r_{ij} \leq r_a)\mathbf{u}_j(t)}{\| \mathbf{u}_i(t) + \sum_{j\neq i} \mathbf{1}(r_{ij} \leq r_a)\mathbf{u}_j(t) \|}.$$

The pairwise force is

$$\mathbf{F}_i(t) = \sum_{j\neq i} f(r_{ij})\widehat{\mathbf{r}}_{ij}.$$

The scalar radial force law is

$$f(r) = \begin{cases} -\infty, & r \leq 2r_0, \\ (\beta+1)\frac{r-2r_0}{r_1-2r_0} - 1, & 2r_0 < r < r_1, \\ 0, & r \geq r_1, \end{cases} \qquad \beta = exp\left[\tan\left(\frac{\pi B}{2}\right)\right],$$

Here r_0 is the particle radius and r_1 is the interaction cutoff. The adhesion parameter B controls the zero crossing of the finite branch of the radial force and therefore the relative extent of its repulsive and attractive portions: changing B changes where the interaction switches from exclusion-like repulsion to cohesion-like attraction. The numerical implementation replaces the divergent hard-core branch with a large finite repulsion and performs a short midpoint-preserving overlap projection, so that no sampled final state retains center separations below $2r_0$.

Particle number is specified by a density fraction ρ of a hexagonal close-packing estimate,

$$N = round\left[\rho\frac{\sqrt{3}}{6}\left(\frac{S}{r_0}\right)^2\right].$$

The primary classical QD searches use $S = 4$ and $0.05 \leq \rho \leq 0.30$, yielding approximately 23-139 particles. The representation-learning and high-density experiments use $S = 3.2$ and $0.05 \leq \rho \leq 0.70$, yielding approximately 15-206 particles (the $\rho = 0.70$ close-packing value is 207) while preserving feasible non-overlapping initialization.

The bounds in Table 1 are computational exploration bounds rather than estimates for one experimental system. Length is measured relative to the particle radius and time in simulation steps. Density spans dilute conditions to near close packing; the interaction and alignment radii range from near-contact neighborhoods to broadly nonlocal neighborhoods; force, alignment, and persistence weights span negligible to dominant contributions to the drive vector; and η spans deterministic headings to the largest angular-noise range used in the model. The density-expanded study deliberately pushes farther toward crowding while reducing the maximum neighborhood radii to remain compatible with the finite box. All repertoire and boundary statements are therefore conditional on these stated search boxes.

Table 1. Active single-species parameters and computational bounds.

Parameter	Symbol	Classical QD bound	Density-expanded bound	Interpretation
Density fraction	ρ	0.05-0.30	0.05-0.70	Fraction of close-packing estimate
Speed	v_0	0.01-0.12	0.01-0.12	Distance advanced per step
Angular inertia	ε	0-5	0-5	Persistence of current heading
Interaction-range ratio	r_1/r_0	2.1-18	2.1-14	Force cutoff relative to radius
Force strength	f_0	0-5	0-5	Weight of pairwise force
Adhesion	B	-0.8-0.8	-0.8-0.8	Shape of the pair force
Alignment-range ratio	r_a/r_0	2.1-18	2.1-14	Alignment radius
Alignment strength	α	0-5	0-5	Weight of local alignment
Noise	η	0-1	0-1	Angular-noise amplitude

2.2. Classical collective-state measures

Polar alignment is

$$\phi = \parallel \frac{1}{N} \sum_i \mathbf{u}_i \parallel .$$

Nematic order is

$$S_2 = \left| \frac{1}{N} \sum_i e^{2i\theta_i} \right| .$$

Milling is the normalized angular momentum around the group center,

$$m = \frac{|\sum_i \mathbf{r}_i \times \mathbf{u}_i|}{\sum_i \parallel \mathbf{r}_i \parallel},$$

where $\mathbf{r}_i$ is relative to the arithmetic center for reflecting boundaries and to a circular coordinate-wise center for periodic boundaries. The latter avoids the spurious large radii generated when a cohesive group crosses a periodic boundary.

Clustering compares the observed mean number of neighbors within r_1 with the expectation for uniformly distributed particles,

$$\Gamma = \frac{\langle k_i(r_1) \rangle}{(N-1)p(r_1)},$$

where $p(r_1) = \pi r_1^2/S^2$ for a periodic square, subject to the finite-box radius constraint. For reflecting boundaries we use the exact first branch of the square line-picking distance distribution. Thus $\Gamma \approx 1$ denotes a spatially uniform reference state, $\Gamma > 1$ local aggregation, and $\Gamma < 1$ local depletion or large-scale spacing.

For density-expanded analysis we additionally compute: the largest low-wave-number density-mode amplitude, a local-neighbor-count coefficient of variation, and sixfold bond order

$$\psi_6 = \left| \frac{1}{N} \sum_i \frac{1}{k} \sum_{j \in \mathcal{N}_k(i)} e^{6i\alpha_{ij}} \right|,$$

using the six nearest neighbors. These are diagnostics rather than archive objectives.

2.3. Direct quality-diversity search and scalar baseline

The experiments use a 12 by 12 MAP-Elites archive [9-11]. For a physics reader, the procedure can be read simply as an adaptive way to populate a behavioral plane. Each trial selects a parameter vector, simulates it, computes the two behavior coordinates, and assigns the result to one of 144 bins. A bin stores the parameter vector with the highest temporal stability found so far for that behavioral type. Mutation of stored parameter vectors then generates new trials, so the search spends evaluations both improving occupied bins and reaching previously empty ones. Unlike scalar optimization, the output is therefore a repertoire of stable examples distributed across behavior space rather than one optimum. Temporal stability is one minus the mean within-trajectory variation of alignment, milling, and scale-normalized clustering.

Each classical archive is filled by a process-parallel batched MAP-Elites run of 64,000 simulator evaluations, made affordable because the embarrassingly parallel evaluations are distributed across cores by a batched search that is numerically equivalent to the serial routine. The alignment-milling archive uses reflecting boundaries because walls stabilize rotating states. The alignment-clustering archive uses periodic boundaries to prevent wall accumulation from dominating aggregation. Representative rules are then evaluated for longer trajectories and over 20 independent seeds.

A global optimizer, Differential Evolution (DE/rand/1/bin), provides the historical scalar baseline [32]. The DE runs use a population of 16 and eight generations after initialization, totaling 144 evaluations. The milling objective is $m(1 - \phi)$ to suppress translating flocks.

2.4. Density-expanded behavior corpus

To expose high-density behavior and provide data for learned representations, we generated 360 space-filling parameter vectors using scrambled Sobol designs. The corpus contains 120 low-density rules ($0.05 \leq \rho < 0.20$), 120 medium-density rules ($0.20 \leq \rho < 0.45$), and 120 high-density rules ($0.45 \leq \rho \leq 0.70$). Four fifths use periodic boundaries and one fifth reflecting boundaries. Each rule is simulated for 160 steps, with behavior sampled over the second half and configurations recorded every eight steps. All statistical comparisons use rules as independent observations.

2.5. Learned behavior descriptor

The learned descriptor is computed from recorded positions and headings, not from ϕ, m, Γ, or the other hand-designed order parameters. It is translation-invariant, rotation-invariant, and compatible with periodic boundaries. For each of the final ten recorded frames it includes:

1. a 12-bin radial distribution estimate normalized by the ideal-gas shell expectation,

2. a 12-bin radial heading-correlation profile,

3. nearest-neighbor distance quantiles,
4. local-neighbor-count heterogeneity,
5. radius of gyration and shape anisotropy,
6. eight low-wave-number density-spectrum shells,
7. short-lag autocorrelations of the complex mean heading,
8. global-heading angular-increment statistics, and
9. frame-to-frame mobility.

In this list, the radial distribution records excess pair density in concentric shells; radial heading correlation averages the dot product of particle headings for pairs in each shell; nearest-neighbor quantiles summarize local spacing; local-neighbor-count heterogeneity is the coefficient of variation of counts inside the chosen local radius; radius of gyration and shape anisotropy summarize overall spatial extent and elongation; low-wave-number spectral shells summarize large-scale density modulation; heading autocorrelations measure short-time persistence of the complex mean heading; global-heading increments summarize its angular changes; and frame-to-frame mobility is the mean particle displacement between successive recorded configurations after applying the same periodic unwrapping used for trajectory measurements. For radial bin b, n_b is the number of particle pairs whose separation lies in that bin and A_b is the area of the corresponding annulus, so the radial statistic is normalized against the ideal-gas shell expectation.

For radial bin b,

$$g_b = \frac{n_b}{T\,N(N-1)A_b/(2S^2)}, \qquad C_b = \frac{1}{n_b}\sum_{(i,j)\in b}\cos(\theta_i - \theta_j),$$

where A_b is the annular area and T is the number of frames. The low-wave-number density spectrum uses

$$S(\mathbf{k}) = \left|\frac{1}{N}\sum_j \exp\left(\frac{2\pi i}{S}\mathbf{k}\cdot\mathbf{x}_j\right)\right|^2.$$

The resulting descriptor $\mathbf{h} \in \mathbb{R}^{51}$ is the raw input to the representation. Because the components differ in scale by orders of magnitude, they are standardized with statistics computed on the training partition only,

$$\widetilde{\mathbf{h}} = (\mathbf{h} - \mu_{\text{train}}) \oslash \sigma_{\text{train}}, \qquad \sigma_{\text{train},d} \leftarrow max(\sigma_{\text{train},d}, 10^{-6}),$$

Here μ_train and σ_train are, respectively, the componentwise mean and standard deviation of each descriptor feature computed from the training partition only; division is elementwise, and the floor on σ_train prevents division by near-constant features. Standardization uses only training statistics so that neither validation nor independent test data influence the representation.

The autoencoder architecture, training procedure, model-selection protocol, and independent post-selection validation are detailed in Supplementary Section S4. The main text uses the fixed

two-dimensional representation for visualization and archive construction; higher-dimensional audits are reported in the Supplementary Materials.

Detecting learned behaviors. For archiving and region discovery the two learned coordinates are rescaled to the unit square by their first and ninety-ninth training percentiles,

$$z_k^{01} = clip\left(\frac{z_k - q_1(z_k)}{q_{99}(z_k) - q_1(z_k)},\, 0,\, 1\right), \qquad k \in \{1, 2\},$$

This percentile rescaling limits the effect of a few outlying embeddings. The 20 by 20 grid is used only for coverage accounting; it does not enter autoencoder training or the independent-test neighborhood metrics. Likewise, k-means with twelve centers is a descriptive partition used to choose a compact set of representative examples, not a phase-identification step. Twelve was selected as a human-inspectable summary and the silhouette coefficient (0.426) is reported rather than optimized; fifty restarts reduce dependence on a poor k-means initialization. Changing these display choices would alter bin counts or cluster labels, but not the learned coordinates, test-set comparisons, local Jacobians, or the reciprocal-control conclusions. Each region is represented by the simulation whose rescaled code is nearest its centroid. Finally, each learned axis is correlated (Spearman) against the withheld physical measures φ, m, Γ, S_2, ψ_6, P, density, drift, and chirality, so that the data-driven axes can be interpreted without having been trained on them. These clusters are descriptive partitions of the learned manifold, not claims of thermodynamic phases.

2.6. Local sensitivity and parameter geometry

For six representative regimes (aligned, milling, clustered, density-wave, high-density aligned, and high-density aggregate, of which the aligned and high-density aligned representatives are the same dense-flock rule, giving five distinct parameter points), each active parameter is perturbed by ± 0.02, ± 0.04, and ± 0.08 of its normalized range, with sixteen seeds evaluated at each side and twenty-four further replicates at the unperturbed point used to set a seed-to-seed noise floor (Section 6.2). The behavioral output vector contains

$$\mathbf{y} = (\phi, m, log(1 + \Gamma), S_2, z_1, z_2, \psi_6, S_{\max}),$$

standardized by its corpus-wide standard deviation. The local Jacobian is

$$J_{a\mu} = \frac{\partial \widetilde{y}_a}{\partial u_\mu},$$

where $u_\mu \in [0, 1]$ is a normalized parameter. The local behavioral information matrix is

$$F = J^{\mathsf{T}} J.$$

This matrix measures how the selected behavioral observables respond locally to the model parameters. It is not a Fisher information matrix, because it is not derived from the curvature of a trajectory likelihood [19-22]. Large eigenvalues identify parameter combinations that strongly change the measured behavior and small eigenvalues identify weakly responsive directions. With eight outputs and nine parameters, at least one direction is rank-deficient.

For comparison with the local response analysis, we also build a 20 by 20 archive from the 360-rule corpus in the learned plane and compute the covariance of the normalized elite parameter vectors. This archive covariance describes how selected parameter sets are distributed across the repertoire and should not be interpreted as local stiffness.

2.7. Two-species extension

Particles are assigned type A or B. Same-type alignment remains one. Cross-species alignment is decomposed into symmetric and antisymmetric components,

$$\mathbf{J} = \begin{pmatrix} 1 & J_s + J_a \\ J_s - J_a & 1 \end{pmatrix}.$$

The additive cross-species radial coupling is

$$\mathbf{C} = \begin{pmatrix} 0 & C_s + C_a \\ C_s - C_a & 0 \end{pmatrix}.$$

Rows denote focal species and columns neighbor species. $J_a = C_a = 0$ defines the reciprocal subfamily. Nonzero antisymmetric components permit alignment frustration and chase-like forces. The single-species model is recovered when all particles have one type or when cross-species couplings reproduce same-species interactions.

We fix the dimensional accounting once here. The full model has eleven parameters (Table 1 plus the two gradient parameters γ and the gradient angle). The main experiments set the external gradient to zero, leaving the 9 active single-species parameters of Table 1. The two-species extension adds five coupling parameters (f_B, J_s, J_a, C_s, C_a), giving 14 parameters for the full non-reciprocal family; the reciprocal subfamily fixes $J_a = C_a = 0$ and therefore has twelve. Statements elsewhere in the paper of a "nine- to fourteen-dimensional" model refer to this range from the single-species (nine) to the full two-species (fourteen) case.

Two trajectory measures describe the principal dynamical outcomes. Center-of-mass drift is computed from unwrapped positions,

$$d = \frac{\| \mathbf{X}_{CM}(t_1) - \mathbf{X}_{CM}(t_0) \|}{v_0(t_1 - t_0)}.$$

Persistent chirality combines the signed slope ω of the unwrapped global-heading angle, the weighted regression coefficient R^2, and mean polarization,

$$C \;=\; min(|\omega|/0.05,\; 1)\; R^2\; mean(\varphi)$$

The first factor saturates once the fitted angular speed reaches 0.05 rad per step; 0.05 is therefore a reporting scale, not a fitted critical threshold. Multiplication by R^2 suppresses trajectories whose unwrapped global heading does not rotate approximately linearly in time, so a large net angular excursion produced by irregular reversals is not scored as persistent rotation. Multiplication by mean polarization further suppresses cases in which the global heading is poorly defined. The score is thus a bounded screening coordinate for archive construction; the mechanistic tests in Sections 7.2-7.3 also report the underlying angular velocity directly and do not rely on the composite score alone.

The reciprocal and non-reciprocal drift-chirality archives are each filled by a 64,000-evaluation process-parallel search, as in the single-species case. Selected rules are validated over 20 seeds and up to 1,200 time steps. A species-aware learned descriptor adds AA, BB, and AB pair-distance and heading-correlation profiles, species polarizations, relative heading, and inter-species center separation. It is trained on 126 reciprocal and non-reciprocal archive or random rules.

3. Classical single-species repertoire

Simulator validation, including the Vicsek-like noise response and numerical checks of hard-core handling and periodic-boundary observables, is reported in Supplementary Section S1. The main text focuses on the behavioral repertoire recovered by illumination.

3.1. Alignment-milling and alignment-clustering archives

At 64,000 evaluations the alignment-milling archive occupies 140 of 144 cells (97.2% coverage) and the alignment-clustering archive fills all 144 cells (100% coverage). The alignment-milling map contains stable rotating states with m approaching 0.50 at low-to-moderate polarization, together with nearly perfectly aligned flocks. In the alignment-clustering map, the bottom edge (Γ near one or below) contains approximately uniform states, the right side contains strongly polarized flocks, and the upper part contains increasingly aggregated states, reaching Γ of about 8. These labels describe regions of the displayed behavior plane; representative configurations are shown in Figure 2.

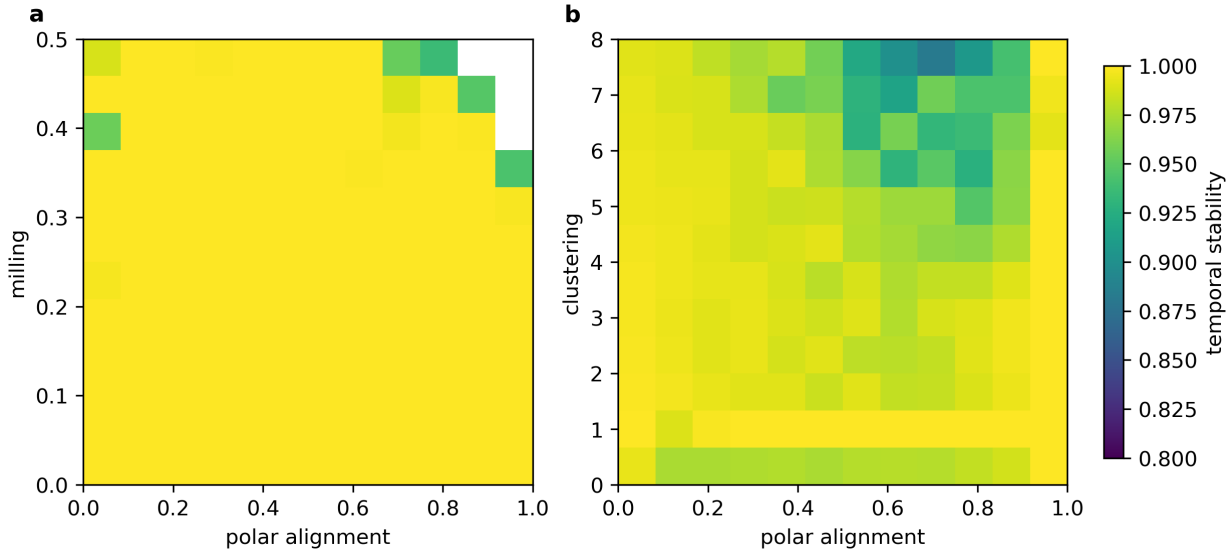


Figure 2. Direct quality-diversity archives for one particle type. (a) Alignment-milling archive with reflecting boundaries. (b) Alignment-clustering archive with periodic boundaries. Each archive is a single 64,000-evaluation process-parallel batched MAP-Elites run; the alignment-milling archive reaches 97.2% coverage (only the four high-alignment/high-milling corner cells stay empty) and the alignment-clustering archive reaches full coverage. Color is temporal stability.

The archive also shows a clear constraint. Strong milling occurs away from the fully aligned edge, and the four cells with both very high alignment and very high milling remain empty. The rest of the alignment-milling archive is essentially filled. After 64,000 evaluations this persistent empty corner is strong evidence that such states are not reached within the tested parameter bounds and archive resolution, although it is not a proof that they are impossible.

Coverage rises rapidly and then saturates (Supplementary Figure S17). The alignment-clustering archive reaches 144 of 144 cells, and the alignment-milling archive plateaus at 140 of 144. The unfilled high-alignment/high-milling corner is therefore not an ordinary search gap at this resolution.

3.2. Multi-seed representative validation

Four classical representatives were selected: maximum alignment, maximum $m(1-\phi)$, maximum clustering, and a low-order state. Figure 3 shows final configurations, and Table 2 reports 20-seed validation.

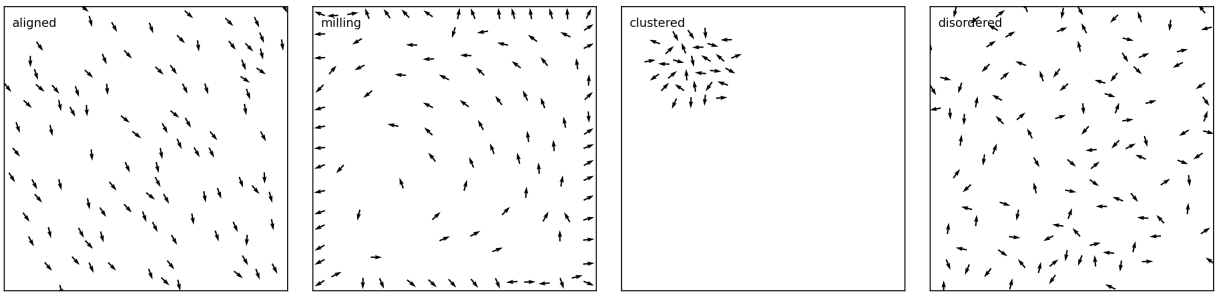


Figure 3. Validated classical single-species representatives: aligned flock, reflecting-boundary mill, compact cluster, and disordered gas. Arrow direction indicates particle heading.

Table 2. Validation of classical single-species representatives over 20 seeds. Values are mean ± standard deviation.

State	N	φ	m	Γ	Seed reproducibility
Aligned	100	0.960 ± 0.000	0.033 ± 0.007	0.749 ± 0.032	0.987
Milling	104	0.259 ± 0.032	0.463 ± 0.024	0.803 ± 0.009	0.978
Clustered	31	0.112 ± 0.010	0.060 ± 0.017	6.317 ± 0.770	0.950
Disordered	114	0.068 ± 0.014	0.049 ± 0.011	0.142 ± 0.027	0.983

The aligned state is highly reproducible across seeds. The mill is also reproducible but is more sensitive to box size and wall conditions. The compact cluster has a much larger local-density excess, but it contains only 31 particles because density was an evolved parameter. It demonstrates strong aggregation in the model family; it does not show that low density is intrinsically optimal for clustering.

System-size scope. The representative states in Table 2 are intentionally modest finite systems because the paper maps many parameter combinations and seeds. The conclusions are not based only on N = 50-100, however. The density-expanded corpus reaches N = 206, the noise-driven order-disorder check in Section 8.2 reaches approximately N = 320, and the long-horizon rotation scan in Section 7.2 spans N = 15-244. Supplementary Section S2 additionally changes box size at fixed density for aligned, milling, and clustered representatives. Alignment and clustering remain robust over that scan, while milling is more size- and boundary-sensitive. We therefore treat all results as finite-system statements and avoid thermodynamic-limit claims where the available scaling tests do not support them.

4. Density dependence and high-density states

4.1. Behavior changes with density

The density-expanded corpus contains 120 low-, 120 medium-, and 120 high-density rules. Across this space-filling sample, density is negatively correlated with polar order (r_s = -0.579),

milling (r_s = -0.768), and nematic order (r_s = -0.775); all three correlations have $p < 10^{-23}$. Clustering also decreases with density, but more weakly (r_s = -0.216, $p = 3.7 \times 10^{-5}$), because its normalization removes the trivial increase in neighbor number.

Median polarization falls from 0.183 in the low-density group to 0.069 in the high-density group, and median milling falls from 0.098 to 0.044. Large-scale density modulation changes even more strongly: the median low-wave-number density peak falls from 0.102 to 0.0037. All reported diagnostics differ significantly among the three density groups by Kruskal-Wallis tests. See Figure 4.

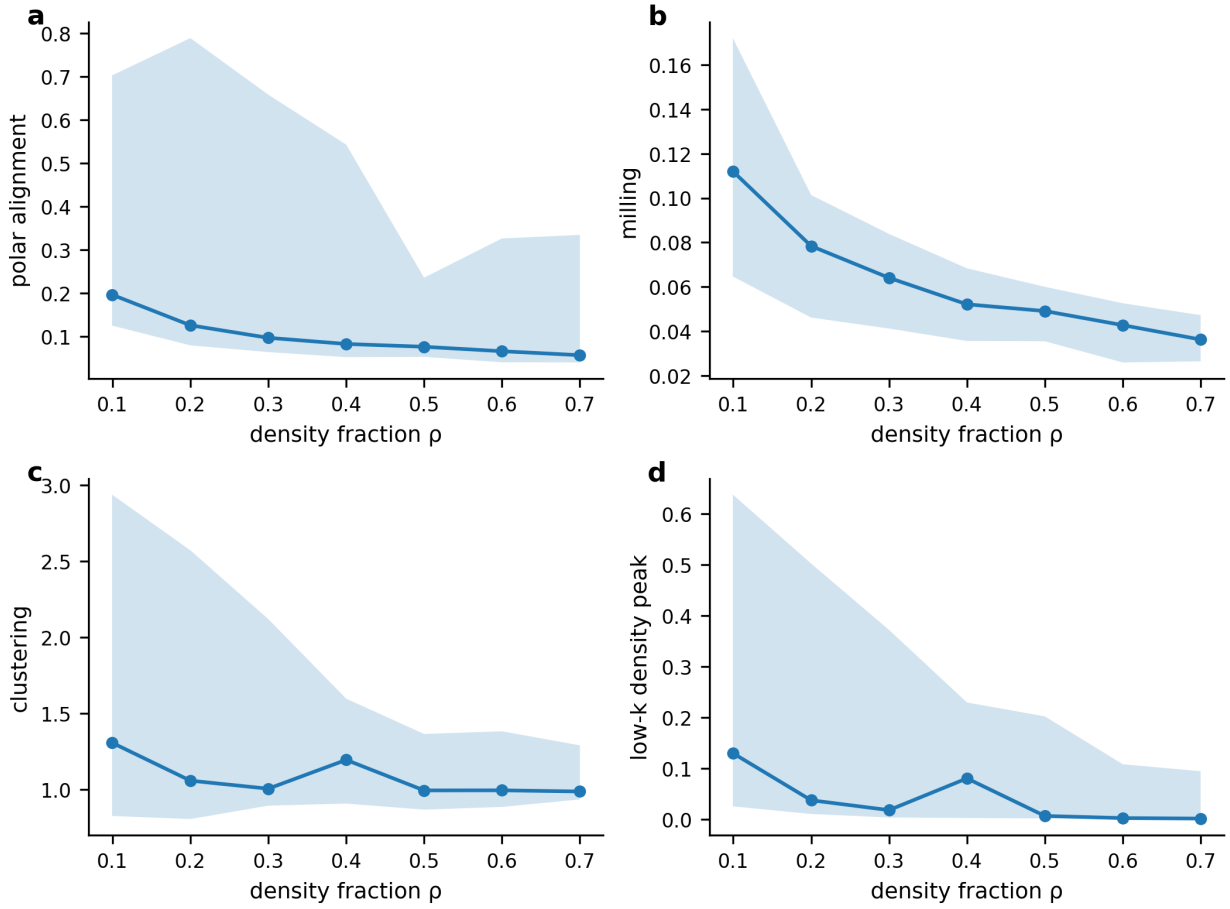


Figure 4. Density dependence in the 360-rule corpus. Lines are medians within density bins and shaded bands are the 10th to 90th percentiles. Ordered states become less common at high density, but the upper tails show that they remain accessible**.**

Because the other eight parameters vary at the same time, these curves are not a thermodynamic phase diagram. They show a simpler repertoire-level result: under broad parameter sampling, strongly ordered states become less common as density increases.

The density dependence here should also not be read as a test for motility-induced phase separation (MIPS). MIPS in active Brownian systems is normally diagnosed with coexistence, structural, or transport signatures; recent work, for example, shows that diffusion itself can carry transient signatures across the MIPS region [35]. Our density corpus instead asks a repertoire-level question while all other model parameters are allowed to vary.

4.2. Ordered states still occur at high density

The average trend does not exclude ordered high-density states. The high-density group contains a periodic state with $\rho = 0.595$, $N = 176$, $\varphi = 0.99999$, and $S_2 = 0.99997$. It also contains a reflecting-wall aggregate with $N = 190$, $\Gamma = 2.67$, and low-wave-number density peak 0.322; a periodic density-modulated state with $N = 201$ and density peak 0.391; and a state with $N = 194$ and sixfold bond order $\psi_6 = 0.744$.

The dense aligned state is a counterexample to a simple density-suppresses-order rule. Strong alignment can maintain a coherent flock even near close packing. High-density aggregation and density modulation, in contrast, depend strongly on wall conditions and force parameters. These differences are not captured by polarization, milling, and clustering alone.

5. Learned representation of single-species behavior

5.1. Independent test performance

We evaluate the fixed two-dimensional autoencoder on 72 additional rules generated only after the representation had been selected. On this independent test set, its standardized reconstruction error is 0.250, trustworthiness is 0.890, and ten-neighbor overlap is 0.556. Two-component PCA, fit only on the original training data, gives 0.402, 0.829, and 0.458 on the same three measures. Thus the learned map preserves local neighborhoods better than two-dimensional PCA on data that played no role in model selection. Training and model-selection statistics are reported in Supplementary Section S4.

The three reported scores measure different aspects of compression. Reconstruction error is the mean squared error after mapping a trajectory descriptor into two coordinates and decoding it back. Trustworthiness asks whether points that appear close in the two-dimensional map were also close in the original 51-dimensional descriptor space; values closer to one indicate fewer false neighbors. Ten-neighbor overlap is the fraction of each rule's ten nearest descriptor-space neighbors that remain among its ten nearest neighbors after projection. These quantities make the comparison with PCA interpretable as a neighborhood-preservation test rather than a claim that the two learned axes are unique physical order parameters.

The geometry is also reasonably stable across random initializations. Across four independently trained models, the median Spearman correlation between pairwise latent distances is 0.795 on the independent test set, with a range of 0.661 to 0.945. The coordinates themselves can rotate or deform, but nearby and distant behaviors are largely preserved.

Two dimensions are useful for visualization, but they are not sufficient to represent all descriptor information. On the independent test set, reconstruction error falls from 0.250 in two dimensions to 0.133 in four and 0.111 in six. Neighborhood preservation improves most by four dimensions and changes little thereafter. We therefore use two dimensions as a display and archive coordinate system, not as a claim that the repertoire is intrinsically two-dimensional.

5.2. Physical interpretation of the learned axes

Figure 5 compares the learned coordinates with physical quantities that were not given to the autoencoder. Learned axis 2 is strongly associated with polarization ($r_s = 0.719$) and center-of-mass drift ($r_s = 0.649$). Learned axis 1 is negatively associated with large-scale density modulation ($r_s = -0.698$) and positively associated with density ($r_s = 0.583$). Clustering depends on both axes, and milling is most strongly associated with axis 1 ($r_s = -0.456$).

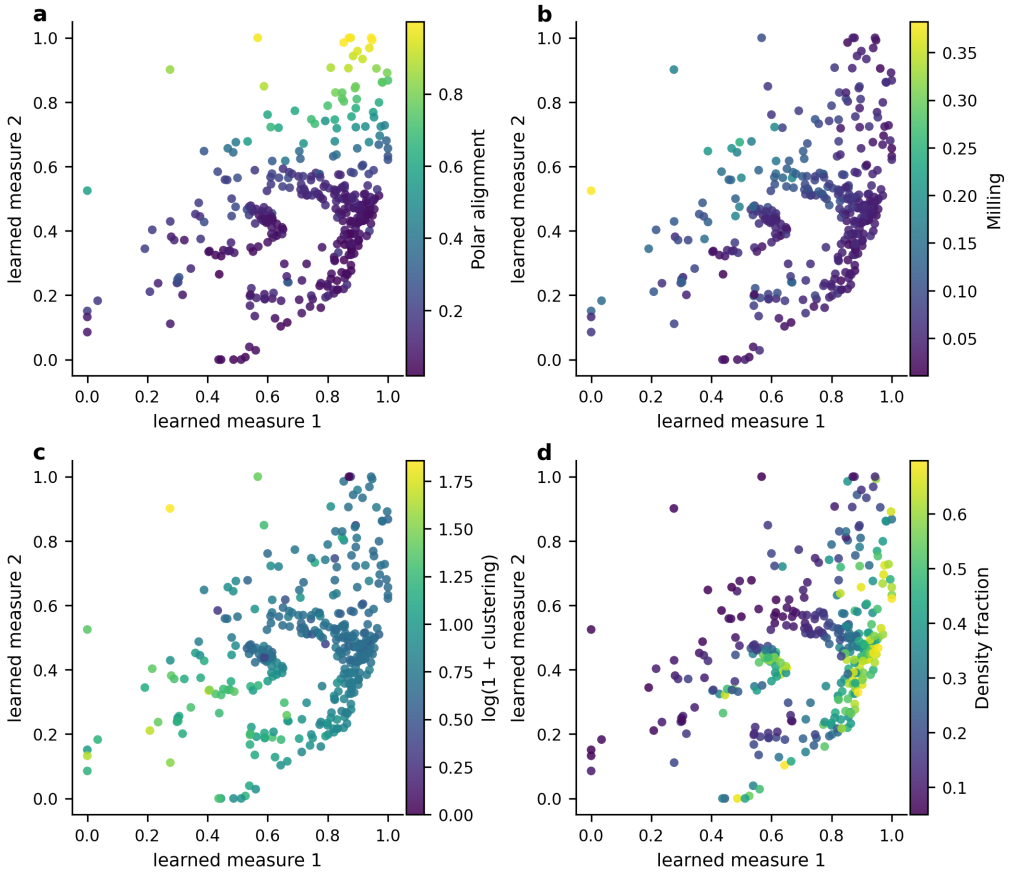


Figure 5. Two-dimensional learned behavior space for 360 single-species simulations. Points are colored by physical quantities that were not used to train the autoencoder. The learned coordinates recover major variation in polarization, milling, aggregation, density modulation, and density.

The learned map is not simply a relabeling of the classical order parameters. States with similar low polarization and similar clustering can still differ in density, sixfold order, spatial modulation, or wall localization. It also separates compact periodic aggregates from wall-supported aggregates even though boundary condition was not provided as an input feature.

5.3. Information beyond the classical measures

To test whether the learned map contains structure not captured by the classical measures, we compare it with the eight-dimensional vector (φ, m, $\log(1+\Gamma)$, S_2, ψ_6, P, $c\rho$, ρ). We ask two questions: how well can these eight measures predict each learned axis, and how well does a two-dimensional map of the eight measures preserve neighborhoods in the full 51-dimensional trajectory descriptor?

The second learned axis is almost fully predicted by the classical measures. On the independent test set, the best of four regressors explains 93.5% of its variance. For axis 1, the best regressor explains 66.5%, leaving 33.5% unexplained. The remaining structure is mainly associated with spatial modulation, local packing, and wall localization. This result does not prove that axis 1 is mathematically irreducible to the eight measures; it shows that four substantially different predictors fail to recover about one third of its test-set variance. See Figure 6.

The neighborhood comparison gives the same ordering. Against the full 51-dimensional descriptor, the autoencoder has test-set trustworthiness 0.890 and ten-neighbor overlap 0.556. Descriptor PCA gives 0.829 and 0.458, and a two-component PCA of the eight classical measures gives 0.765 and 0.374. This comparison is specific to preservation of the trajectory descriptor, but it shows that the learned map retains structure that a two-dimensional classical map loses.

Individual examples make the difference concrete. Among the 10% of rule pairs that are closest in the eight classical measures, most are also close in the learned plane, but some are separated by as much as 0.5 in unit-square coordinates. Such pairs can have nearly identical clustering, density peak, and sixfold order while differing in their spatial-spectrum or heading-correlation profiles.

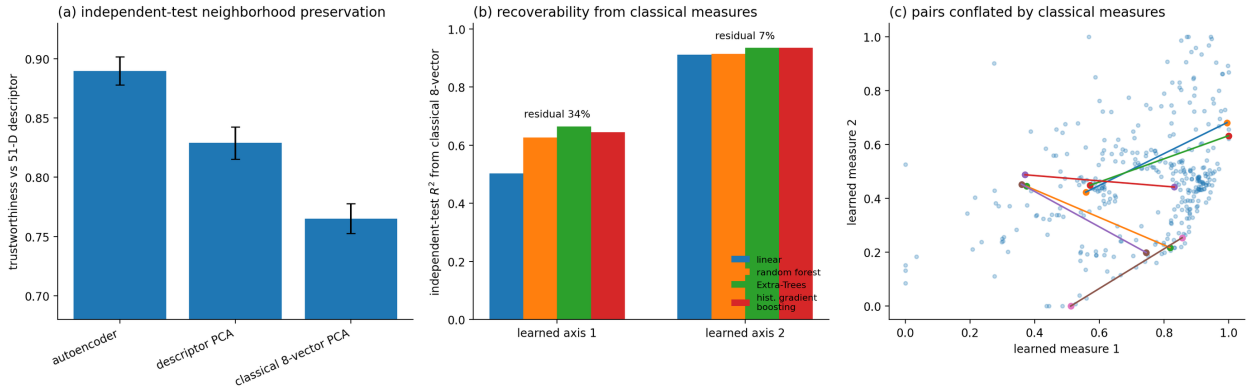


Figure 6. Comparison of the learned map with the classical eight-vector. (a) Independent-test trustworthiness against the full 51-dimensional trajectory descriptor. (b) Independent-test R^2 for predicting each learned axis from the eight classical measures. (c) Examples of rule pairs that are close in classical-measure space but well separated in the learned plane.

Supplementary Section S5 gives a twelve-region descriptive clustering and representative configurations. The clusters summarize the map for inspection; they are not thermodynamic phase labels.

6. Local parameter sensitivity

6.1. State-dependent stiff directions

Raw two-seed sensitivity spectra span 3.3 to 7.3 decades, but most of that range is below stochastic simulation noise. After increasing the number of seeds and applying the noise threshold described in Section 6.2, only two to four local singular directions are resolved in each regime. We therefore interpret the resolved directions, not the full raw decade span.

The dominant resolved directions depend on the state (Supplementary Table S6). For the mill, the leading direction is mainly a combination of density (0.65 [0.51, 0.77]), force strength (0.56 [0.36, 0.72]), and adhesion (0.32 [0.15, 0.49]). For the clustered state it combines interaction range (0.44 [0.19, 0.68]), adhesion (-0.41), and noise (-0.33). The density-wave direction is dominated by density (-0.97 [-0.99, -0.94]), and the high-density aggregate by adhesion (-0.88 [-0.90, -0.87]). The dense aligned rule has no well-resolved leading eigenvector at this seed count, so we do not assign one.

Figure 7 shows the normalized local spectra. Each regime has a small number of large eigenvalues followed by directions that quickly fall into the noise-dominated range. The archive-wide

covariance spectrum is much narrower, which already suggests that local sensitivity and global parameter dispersion measure different things.

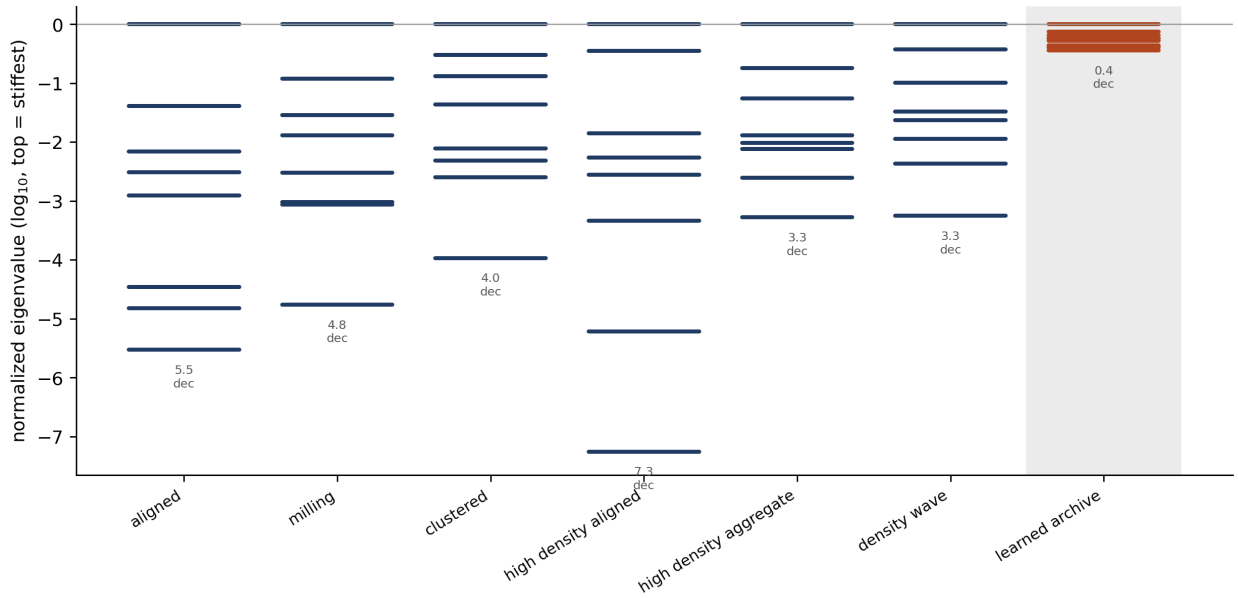


Figure 7. Local sensitivity spectra. Eigenvalues are normalized by the largest value in each regime. Only the top two to four directions lie above the seed-to-seed noise floor. The shaded final column shows the much narrower inverse-variance spectrum of parameters stored in the learned archive.

6.2. Resolving sensitivity above simulation noise

We estimate the noise floor directly. At each unperturbed rule, the eight standardized outputs are measured over 24 independent seeds. The local Jacobian is then recomputed with 16 seeds on each side of the perturbation at three step sizes, $\Delta u = 0.02$, 0.04, and 0.08. Seed-to-seed variance gives the uncertainty of each finite-difference derivative. We propagate those uncertainties by Monte Carlo and use the expected largest singular value of a noise-only Jacobian as the threshold. A singular direction is counted only when it exceeds that threshold.

At $\Delta u = 0.04$, two directions are resolved for the dense aligned and milling rules, three for the clustered and high-density aggregate rules, and four for the density wave. The count changes by one or two directions when the perturbation size is changed, so sub-leading directions are not stable enough to interpret. The multi-seed raw spectra also shrink to about four to five decades, already narrower than the two-seed estimates.

The conclusion is that near each tested regime, behavioral change is dominated by two to four parameter combinations, and those combinations vary by regime. The deeper spectrum is not resolved at the present seed count. Supplementary Table S6 reports the raw spans, the number of resolved directions, and bootstrap intervals for the leading loadings.

6.3. Archive covariance is not local stiffness

The learned archive contains 134 occupied cells. The inverse-variance spectrum of its normalized elite parameter covariance spans only 0.44 decades, far less than the local response spectra. The archive therefore covers a broad parameter volume even though individual behaviors can be locally sensitive to a few parameter combinations.

This distinction is important. Covariance describes how selected parameter vectors are distributed across an archive or an optimizer population. Local stiffness describes how behavior changes when parameters are perturbed around one state. The two quantities need not agree, especially when different states have different stiff directions.

7. Two-species dynamics

Short 200-step drift-chirality archives for reciprocal and non-reciprocal mixtures look similar and are shown in Supplementary Section S7. Long trajectories reveal that many apparently rotating reciprocal states are transient, so the main text uses matched long-horizon controls.

7.1. Translation is not specific to non-reciprocity

The highest-drift non-reciprocal rule is an almost perfectly aligned translating flock with $d \approx 1$. Setting $J_a = C_a = 0$ while leaving every other parameter unchanged still gives $d \approx 1$ and $\varphi \approx 1$ over 20 seeds. Coherent translation therefore does not require broken reciprocity in this model.

This is consistent with earlier work showing that non-reciprocity can generate traveling states when the reciprocal reference is stationary [24]. Our reciprocal reference is already self-propelled and aligning, so translation is not a unique signature of non-reciprocal interactions.

7.2. Persistent rotation depends on antisymmetric couplings

The strongest non-reciprocal rotating rule has low center-of-mass drift and a polarized mean heading that rotates coherently. Over 20 seeds at 800 steps, its median chirality is 0.700 and median $|\omega|$ is 0.163 rad per step. The matched reciprocal control gives 0.013 and 0.00166, and the best rule from an independent reciprocal search gives 0.0060 and 0.00165. The differences are highly significant by paired one-sided Wilcoxon tests (chirality $p = 2.86 \times 10^{-6}$ and 1.91×10^{-6}; angular velocity $p = 9.54 \times 10^{-7}$ for both comparisons) (Figure 8).

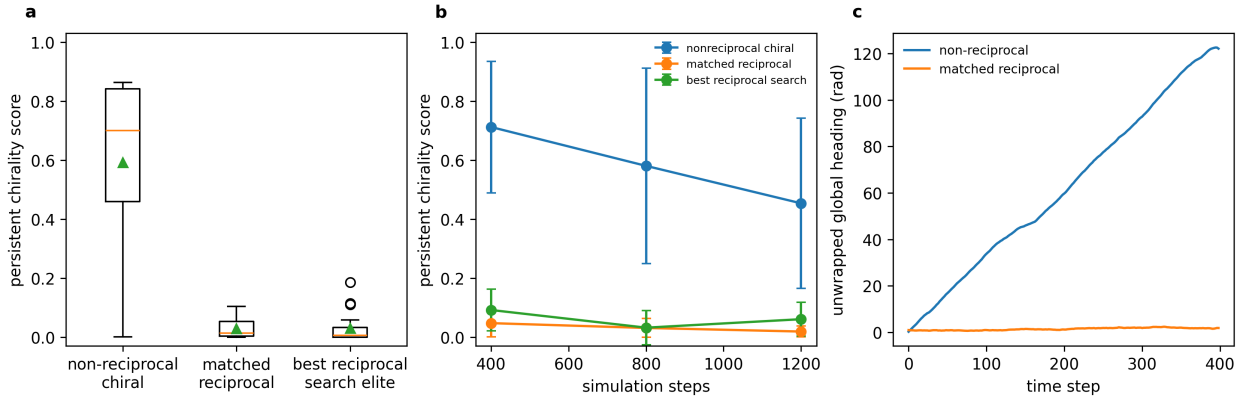


Figure 8. Long-horizon validation of the rotating non-reciprocal state against a matched reciprocal control and the best reciprocal search result. Persistent rotation remains much stronger in the non-reciprocal system.

Windowed 2,000-step trajectories show that the rotation rate remains stationary, and a size scan from 15 to 244 particles shows no systematic decay of rotation. The lower composite chirality score at long horizons comes from its global-linearity term, not from slowing of the motion (Supplementary Figure S9).

We can describe this as a long-lived rotating finite-system regime. The simulations do not establish a stable limit cycle in the thermodynamic limit.

7.3. Rotation turns on with antisymmetric coupling

To measure how rotation depends on non-reciprocity, we fix the base model, species fraction, and symmetric couplings at the rotating rule and sweep the antisymmetric alignment J_a and force C_a. The reciprocal case is $J_a = C_a = 0$. A 13 by 13 grid over $[-1.5, 1.5]^2$, with eight seeds per point, shows the resulting chirality and rotation rate (Figure 9).

At the reciprocal point, mean chirality is 0.035 and $|\omega|$ is indistinguishable from zero. Rotation grows as either antisymmetric coupling moves away from zero, and the sign of the rotation reverses approximately under $(J_a, C_a) \to (-J_a, -C_a)$. Thus, along these two coupling directions, antisymmetry controls both the onset and handedness of rotation.

This result is also consistent with the broader observation that nonreciprocal, information-like interactions can generate chiral collective currents and density organization in active suspensions; a recent visual-quorum-sensing model reports hyperuniform states and edge currents generated by nonreciprocal chirality switching [36]. The microscopic interaction rule is different from ours, so the comparison is qualitative rather than a claim of the same mechanism.

Along the cut $C_a = 0$, $|\omega|$ remains near zero around $J_a = 0$ and rises sharply once $|J_a|$ is of order one. The coarse 13-point scan places the across-seed susceptibility maximum at $J_a \approx 1.33$. A finer 25-point scan places it in the range 1.4 to 1.5. We therefore treat 1.4 to 1.5 as the resolved finite-system onset/crossover range, not as a thermodynamic critical point.

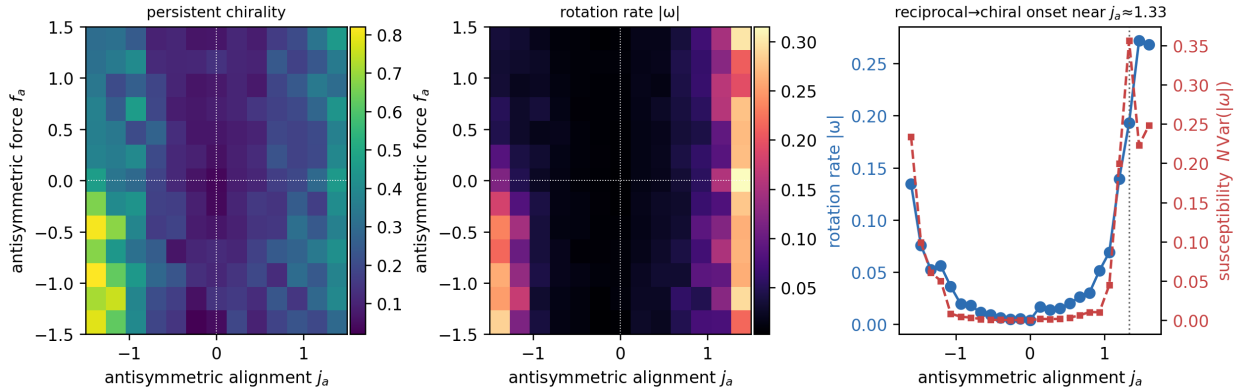


Figure 9. Response to antisymmetric two-species coupling. (a) Persistent chirality and (b) rotation rate over J_a and C_a, with all other parameters fixed. Rotation vanishes at the reciprocal point $J_a = C_a = 0$. (c) Along $C_a = 0$, rotation rises when $|J_a|$ is of order one; the coarse-grid susceptibility peak is near $J_a = 1.33$.

8. Finding transitions and crossovers along parameter cuts

The archives show where behaviors occur, but they do not by themselves locate where behavior changes rapidly. In a 9- to 14-dimensional parameter space, exhaustive phase-diagram sampling is impractical [33]. We therefore use one-dimensional cuts to test three complementary indicators of reorganization before extending the method to the full parameter space.

8.1. Three estimators

Along a parameter path λ, we evaluate many independent seeds at each point. When an order parameter O(λ) is available, we measure its across-seed fluctuation susceptibility and its Binder cumulant. We also compute a descriptor-based rate of change that does not require a predefined order parameter.

$$\chi(\lambda) = N\,Var_s\left[O(\lambda)\right],$$

The susceptibility peaks where the chosen order parameter varies most strongly across seeds. The second estimator is the Binder cumulant,

$$U(\lambda) = 1 - \frac{\langle O^4 \rangle}{3\,\langle O^2 \rangle^2},$$

Binder-cumulant crossings across system sizes can help localize a transition. The third estimator is the normalized rate of change of the 51-dimensional invariant trajectory descriptor Ψ,

$$r(\lambda) = \frac{\parallel\ (\lambda + \Delta) -\ (\lambda - \Delta)\ \parallel}{2\,\Delta\ \parallel sd(\ )\ \parallel},$$

This central difference measures how rapidly the full trajectory descriptor changes along the parameter path. It can therefore flag a reorganization that is not visible in a chosen scalar order parameter.

These quantities have different interpretations in a finite system. The susceptibility used here is a fluctuation indicator, not a thermodynamic susceptibility. A Binder cumulant is informative only through comparisons across system sizes. The descriptor rate is a general sensitivity measure, not a calibrated transition locator. We therefore use finite-size checks when an order parameter exists and describe other sharp changes as finite-system crossovers or candidate boundaries.

8.2. Noise-driven order-disorder change

We first test the method in a Vicsek-like limit with broad alignment and no radial force [1,34]. As noise η increases from 0 to 1, polar order falls from nearly 1 to about 0.1. At a single size, the susceptibility has a broad maximum around η = 0.5 to 0.6. We therefore repeat the cut at five system sizes, N ≈ 29 to 320, and test whether the apparent peak is caused by incomplete relaxation at larger N (Figure 10).

The peak location stays near η = 0.50 across all five sizes and remains there when the trajectory length is doubled. The growth in peak height is largely a finite-window effect associated with longer autocorrelation times, and the polarization distribution remains unimodal. Supplementary Section S9 gives the run-length and autocorrelation checks.

The stable peak location is consistent with a noise-driven order-disorder transition near η* = 0.50. We do not use the peak height to estimate a critical exponent or the order of the transition, because the measured susceptibility is based on window-averaged trajectories and does not obey a standard finite-size scaling form.

The descriptor-rate maximum occurs earlier, around η = 0.4 to 0.5. It is detecting a change in trajectory structure before polarization reaches its largest across-seed fluctuations. This is useful for finding candidate reorganizations, but it also shows that descriptor sensitivity and an order-parameter transition need not occur at the same parameter value.

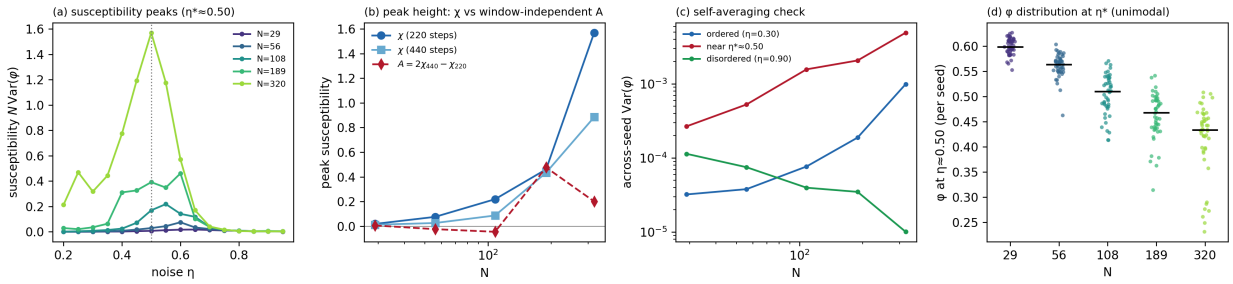


Figure 10. Finite-size check of the noise-driven order-disorder cut. (a) The susceptibility peak stays near $\eta = 0.50$ as N increases. (b) Doubling the run length shows that most peak-height growth is a finite-window effect. (c) The number of effectively independent samples falls as autocorrelation time grows. (d) Polarization remains unimodal near the peak. The evidence supports a transition location near $\eta = 0.50$ but does not determine its order.

8.3. From parameter cuts to full-space search

One-dimensional cuts are useful only where a likely transition direction is already known. To search the full parameter space, we use the trajectory descriptor itself: first identify points with unusually large local descriptor sensitivity, then refine and trace those high-sensitivity sets. A behavioral boundary is claimed only after a physical observable and appropriate controls confirm that behavior changes across the set.

9. Sensitivity ridges in the full parameter space

9.1. Definition and tracing

We call a connected set of unusually large descriptor sensitivity a sensitivity ridge. Such a ridge is only a geometric feature of the chosen descriptor. It becomes a candidate behavioral boundary only if crossing it produces a reproducible change in an interpretable physical observable. This distinction is central because different descriptors can emphasize different changes.

Operationally, at a parameter point u* we estimate the local Jacobian of the invariant trajectory descriptor. A point is accepted as high sensitivity only when the largest singular value exceeds the Monte Carlo seed-noise threshold from Section 6.2 and remains above threshold at at least two perturbation sizes with consistent normal direction. The leading right singular vector gives the local direction of greatest descriptor change.

For boundary validation we use trajectories long enough to exceed 50 integrated autocorrelation times of the defining observable, subject to a step cap. We also avoid the composite chirality score in this stage because its long-horizon bias was identified in Section 7.2. Instead we use separate species-resolved measures of alignment, relative rotation, mixing, demixing, and local sublattice order.

Candidate points are drawn from existing archives and the behavior corpus, supplemented when needed by a QD search that favors high descriptor sensitivity. Each candidate is refined along one dimension and then traced with predictor-corrector continuation. Supplementary Section S10 gives the continuation algorithm, local patch tests, and operating characteristics.

9.2. Controls used to classify a ridge

Each traced ridge is tested with three controls. First, a reciprocal control sets $J_a = C_a = 0$ while keeping composition and the other parameters fixed. Second, a single-species control asks whether the same reorganization is present in the nine-parameter single-species subspace. Third, swapping the species labels should not change the classification beyond seed noise. A ridge that disappears under the reciprocal control and is absent from the single-species subspace is classified as non-reciprocity-dependent. A ridge that survives both is a single-species reorganization. A ridge that survives reciprocity but disappears in the single-species subspace is a reciprocal mixture boundary. If a required control is missing or gives a conflicting result, we leave the case unclassified.

9.3. Descriptor sensitivity and order-parameter transitions

The known noise and non-reciprocity cuts provide checks on the method. On the noise cut, descriptor sensitivity peaks near $\eta = 0.4$, before the polarization-fluctuation maximum near 0.5 to 0.6. On the non-reciprocity cut, descriptor sensitivity peaks near $J_a = 1.0$, close to the first clear onset of rotation, whereas the rotation-rate susceptibility peaks later, around 1.4 to 1.5, when seed-to-seed variability is largest. Across six tested cuts, the signed difference between the two peak locations changes sign and has a spread of about 0.37 of the parameter range (Supplementary Figure S16a). Thus the two measures can identify different reorganizations. In a deliberately flat region, the general detector exceeds threshold at only 0.5% of 200 points.

9.4. Detected ridges depend on the descriptor

With the general trajectory descriptor, most strong ridges found in the two-species search turn out, after controls, to be single-species structural changes. Even the ridge near the reciprocal-to-chiral onset is ambiguous because a single-species reorganization is present at the same location. We therefore repeat the search with a descriptor restricted to mixture-specific quantities: relative rotation, Warren-Cowley mixing, low-wave-number demixing power, and nearest-neighbor sublattice order. This restricted descriptor identifies five additional ridges among about two dozen candidate points. Two survive the reciprocal control but disappear in the single-species subspace and therefore satisfy the criterion for reciprocal mixture boundaries. Three disappear when antisymmetric couplings are removed but have not yet been tested in the single-species subspace, so they remain reciprocity-dependent candidates rather than validated non-reciprocity boundaries. The mixture-specific detector has a 4.5% false-positive rate on 200 identical-species configurations. The main conclusion is simple: the descriptor determines which reorganizations are easy to see.

Supplementary Table S4 lists every ridge and its control status. The present mixture observables distinguish mixing from demixing but do not yet separate different demixed morphologies such as core-shell and side-by-side states. Figure 11 summarizes these full-space results.

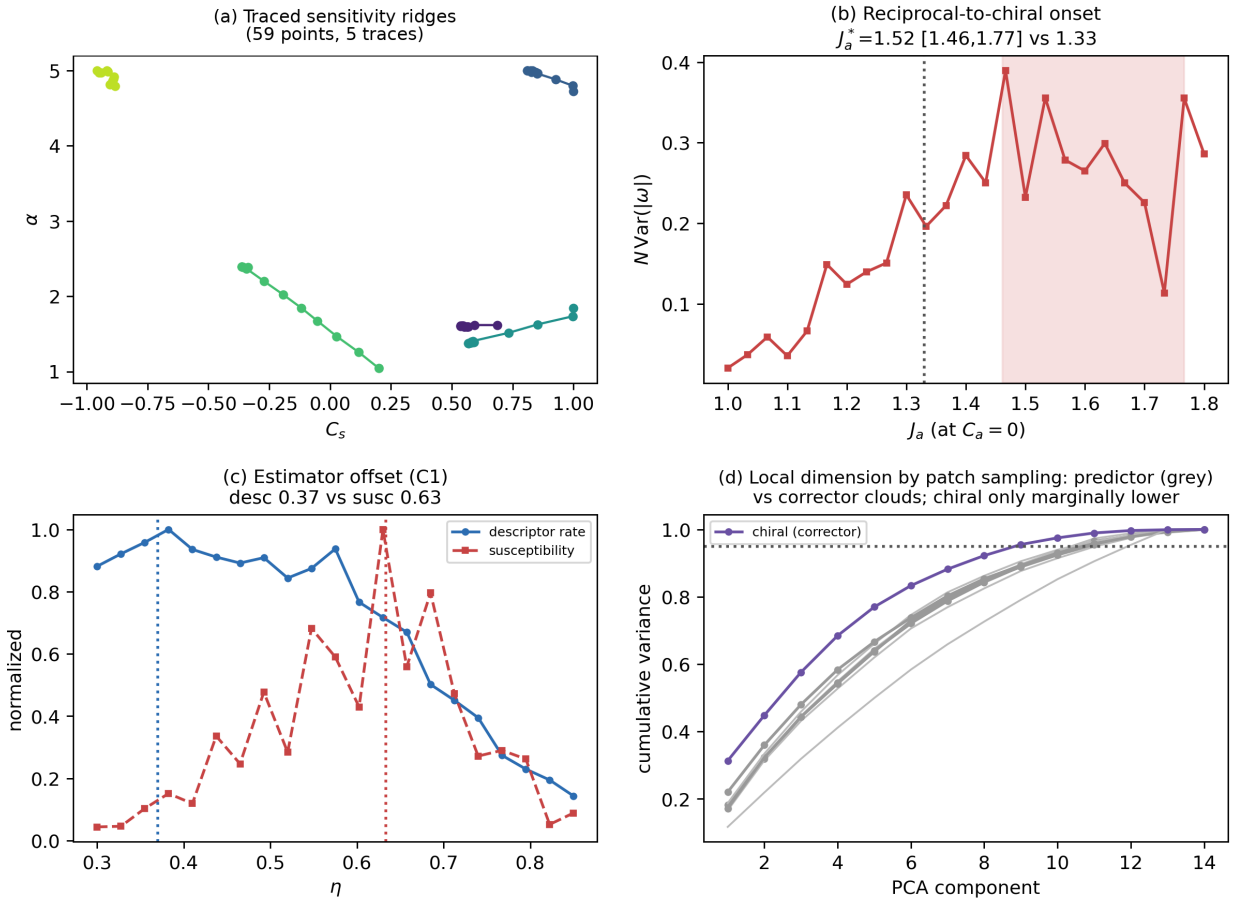


Figure 11. Full-space sensitivity-ridge analysis. (a) Example traced ridges. (b) On the non-reciprocity cut, the descriptor-sensitivity maximum occurs near the onset of rotation, while the rotation-rate susceptibility peaks later. (c) The same displacement occurs on the noise cut. (d) Patch sampling shows that the single-species ridges behave as extended hypersurfaces, whereas the chiral ridge terminates in most tangent directions. Details and controls are in Supplementary Section S10.

10. Discussion

10.1. Main physical observations

The classical QD archives recover aligned, milling, clustered, and disordered states, but the density-expanded corpus shows that these labels do not describe the full repertoire. Density, wall localization, large-scale density modulation, and local sixfold order distinguish states that have similar polarization, milling, and clustering.

High density changes what is typical without imposing a strict constraint. Polarization and milling usually decrease as density rises, yet nearly perfectly aligned dense flocks still occur. Dense aggregates and density-modulated states also differ from dilute aggregates with similar classical clustering.

The learned map provides a compact way to retain some of this extra structure. Its independent test performance is better than two-dimensional PCA, and its first axis is only partly predicted by the eight classical measures. At the same time, four- and six-dimensional versions preserve more information. The two-dimensional map should therefore be viewed as a useful summary, not as a unique physical coordinate system.

10.2. Local parameter sensitivity

The local response analysis supports a limited form of sloppiness. Only two to four parameter combinations are resolved above the noise floor near each tested state, and the dominant combinations depend on the state. The mill is most sensitive to density, force strength, and adhesion; the cluster to interaction range, adhesion, and noise; the density wave mainly to density; and the dense aggregate mainly to adhesion. The dense aligned state has no stable leading direction at the present seed count.

Archive covariance tells a different story because it pools parameter vectors from many behaviors with different local sensitivities. The broad archive distribution therefore does not contradict strong local anisotropy. Local response and ensemble covariance answer different questions and should not be used interchangeably.

Why map the full repertoire if only a few local directions are stiff? The two analyses answer different questions. A local Jacobian can identify the sensitive combinations near one preselected state, but it cannot reveal which other states exist, where to place the next local analysis, or whether the sensitive subspace rotates across behavior space. Here it does rotate: the leading combinations for milling, clustering, density waves, and dense aggregation are different. The global map is therefore not replaced by two to four targeted parameter directions; rather, it supplies the states and neighborhoods on which those low-dimensional local descriptions are conditioned.

10.3. Learned and predefined descriptors

The learned single-species map recovers major physical variation without being trained on the classical order parameters. In the two-species data, however, common features such as drift and polarization dominate the two-dimensional representation, while rare chirality is compressed. This is why learned and hand-designed observables are complementary: learned descriptors summarize common structure, whereas targeted order parameters can preserve rare phenomena that matter physically.

We keep the representation fixed throughout the present analysis so that downstream comparisons use stable coordinates. An adaptive AURORA-style procedure [14] could discover additional regions, but repeated retraining would move the archive coordinates and would require explicit realignment and convergence checks.

10.4. What two species add

The two-species extension separates heterogeneity from non-reciprocity. Reciprocal mixtures can already translate coherently, so translation is not evidence of broken reciprocity. In the tested rotating regime, by contrast, persistent rotation disappears when the antisymmetric couplings are removed.

The learned two-species map also shows why mechanism cannot be inferred from behavior alone. Reciprocal and non-reciprocal rules can occupy the same behavioral region when they generate similar trajectories. Mechanistic claims therefore require interventions or matched controls, not clustering alone.

The three layers of analysis are not required to agree automatically. A learned map is descriptive, a local response calculation is differential, and a matched physical control tests mechanism. If they disagree, we treat that disagreement as diagnostic. In particular, a descriptor-sensitivity ridge is not promoted to a behavioral boundary when the relevant observable or control does not support it; such cases remain ambiguous or unclassified. Conversely, a physical control can reveal a mechanism that is compressed in a learned representation, as occurs for rare chirality in the two-species map. Mechanistic interpretation therefore rests on the targeted controls, with representation geometry used to guide discovery rather than to overrule them.

10.5. Limitations

All transition and boundary statements in this paper are about finite systems. The noise cut is consistent with an order-disorder transition near $\eta = 0.50$, but we do not determine its order or critical exponents. The two-species rotation result establishes a long-lived finite-system regime, not a thermodynamic-limit cycle.

Search is also finite. The single-species archives saturate at the tested 64,000-evaluation budget and resolution, but the empty high-alignment/high-milling corner is an empirical exclusion within those bounds, not a proof of impossibility. Reflecting walls are part of the dynamics and stabilize some mills and aggregates.

The learned representations are specific to the chosen trajectory features, training distribution, and bottleneck dimension. The single-species map has an independent 72-rule test, whereas the two-species audit is based on only 26 held-out rules. Other descriptors could reveal different structure, and higher-dimensional bottlenecks retain information that the two-dimensional maps lose.

The local sensitivity analysis resolves only the leading two to four directions. Deeper directions remain below seed-to-seed noise and depend on perturbation size and the chosen output vector. More precise spectra would require many more seeds, variance reduction, or a likelihood-based treatment.

Finally, high descriptor sensitivity is not itself a behavioral boundary. Descriptor-sensitivity peaks can be displaced from order-parameter susceptibility peaks, and the detected ridge catalog changes when the descriptor is changed. We therefore require a physical observable and the relevant reciprocal and single-species controls before assigning a boundary interpretation.

The model is also intentionally generic. Its parameters are dimensionless computational controls, not fitted material constants, and different experimental systems would require different coarse-graining and calibration. The framework is most directly transferable to settings in which trajectories are available but the relevant collective coordinates are uncertain: animal or cell collectives with several interacting cues, active-colloid mixtures, and programmable or robotic swarms. In such applications, the archive can organize observed or simulated behaviors, while local perturbations and intervention-style controls can test which couplings are actually responsible for a regime.

11. Conclusion

Quality-diversity search provides a practical view of a high-dimensional self-propelled-particle model as a repertoire rather than as a set of isolated optima. It recovers the classical states, shows how their prevalence changes with density, and supplies the trajectory corpus needed to learn a broader behavioral representation. The fixed two-dimensional autoencoder preserves descriptor

neighborhoods better than two-dimensional PCA on an independent test set, while the classical measures explain much of one learned axis but only about two thirds of the other.

The same analysis clarifies parameter control and mechanism. Near each tested state, only two to four parameter combinations are resolved as locally stiff, but those combinations depend on the state, explaining why a repertoire-level map remains useful. Reciprocal interactions already support coherent translation, whereas persistent rotation in the tested two-species regime depends strongly on antisymmetric couplings. Full-space sensitivity search then shows that different descriptors can reveal different reorganizations, and matched controls determine which of those reorganizations merit a mechanistic interpretation. For experimental or engineered active systems, the practical use is not the numerical parameter box itself but the workflow: map a broad repertoire, learn where local responses are concentrated, and use controlled perturbations to distinguish visually similar states generated by different mechanisms.

12. Data and code availability

The code for reproducing all the experiments described in this paper will be made available on a public GitHub repository before publication.

13. AI Statement

During preparation of this work, the authors used AI tools to support simulation coding and to check parts of the manuscript. The authors reviewed and edited all such output and take full responsibility for the published content.

Funding

This work received no external funding.

Competing Interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

14. References

1. Vicsek T, Czirók A, Ben-Jacob E, Cohen I, Shochet O. Novel type of phase transition in a system of self-driven particles. Physical Review Letters. 1995;75:1226-1229. doi:10.1103/PhysRevLett.75.1226.

2. Grégoire G, Chaté H. Onset of collective and cohesive motion. Physical Review Letters. 2004;92:025702. doi:10.1103/PhysRevLett.92.025702.

3. D'Orsogna MR, Chuang YL, Bertozzi AL, Chayes LS. Self-propelled particles with soft-core interactions: patterns, stability, and collapse. Physical Review Letters. 2006;96:104302. doi:10.1103/PhysRevLett.96.104302.

4. Vicsek T, Zafeiris A. Collective motion. Physics Reports. 2012;517:71-140. doi:10.1016/j.physrep.2012.03.004.

5. Marchetti MC, Joanny JF, Ramaswamy S, Liverpool TB, Prost J, Rao M, Simha RA. Hydrodynamics of soft active matter. Reviews of Modern Physics. 2013;85:1143-1189. doi:10.1103/RevModPhys.85.1143.

6. Toner J, Tu Y. Long-range order in a two-dimensional dynamical XY model: how birds fly together. Physical Review Letters. 1995;75:4326-4329. doi:10.1103/PhysRevLett.75.4326.

7. Couzin ID, Krause J, James R, Ruxton GD, Franks NR. Collective memory and spatial sorting in animal groups. Journal of Theoretical Biology. 2002;218:1-11. doi:10.1006/jtbi.2002.3065.

8. Gompper G, et al. The 2025 motile active matter roadmap. Journal of Physics: Condensed Matter. 2025;37:143501. doi:10.1088/1361-648X/adac98.

9. Mouret JB, Clune J. Illuminating search spaces by mapping elites. arXiv:1504.04909. 2015.

10. Cully A, Clune J, Tarapore D, Mouret JB. Robots that can adapt like animals. Nature. 2015;521:503-507. doi:10.1038/nature14422.

11. Pugh JK, Soros LB, Stanley KO. Quality diversity: a new frontier for evolutionary computation. Frontiers in Robotics and AI. 2016;3:40. doi:10.3389/frobt.2016.00040.

12. Fontaine MC, Togelius J, Nikolaidis S, Hoover AK. Covariance matrix adaptation for the rapid illumination of behavior space. In: Proceedings of the Genetic and Evolutionary Computation Conference (GECCO). 2020:94-102. doi:10.1145/3377930.3390232.

13. Fontaine MC, Nikolaidis S. Covariance Matrix Adaptation MAP-Annealing. In: Proceedings of the Genetic and Evolutionary Computation Conference (GECCO). 2023:456-465. doi:10.1145/3583131.3590389.

14. Grillotti L, Cully A. Unsupervised behavior discovery with quality-diversity optimization. IEEE Transactions on Evolutionary Computation. 2022;26(6):1539-1552. doi:10.1109/TEVC.2022.3159855.

15. Hedayatian S, Nikolaidis S. AutoQD: Automatic discovery of diverse behaviors with quality-diversity optimization. International Conference on Learning Representations (ICLR). 2026. arXiv:2506.05634.

16. Carrasquilla J, Melko RG. Machine learning phases of matter. Nature Physics. 2017;13:431-434. doi:10.1038/nphys4035.

17. Wetzel SJ. Unsupervised learning of phase transitions: from principal component analysis to variational autoencoders. Physical Review E. 2017;96:022140. doi:10.1103/PhysRevE.96.022140.

18. Dulaney AR, Brady JF. Machine learning for phase behavior in active matter systems. Soft Matter. 2021;17:6808-6816. doi:10.1039/D1SM00266J.

19. Transtrum MK, Machta BB, Sethna JP. Geometry of nonlinear least squares with applications to sloppy models and optimization. Physical Review E. 2011;83:036701. doi:10.1103/PhysRevE.83.036701.

20. Quinn KN, Abbott MC, Transtrum MK, Machta BB, Sethna JP. Information geometry for multiparameter models: new perspectives on the origin of simplicity. Reports on Progress in Physics. 2023;86:035901. doi:10.1088/1361-6633/aca6f8.

21. Gutenkunst RN, Waterfall JJ, Casey FP, Brown KS, Myers CR, Sethna JP. Universally sloppy parameter sensitivities in systems biology models. PLoS Computational Biology. 2007;3:e189. doi:10.1371/journal.pcbi.0030189.

22. Transtrum MK, Machta BB, Brown KS, Daniels BC, Myers CR, Sethna JP. Perspective: sloppiness and emergent theories in physics, biology, and beyond. Journal of Chemical Physics. 2015;143:010901. doi:10.1063/1.4923066.

23. Ivlev AV, Bartnick J, Heinen M, Du CR, Nosenko V, Löwen H. Statistical mechanics where Newton's third law is broken. Physical Review X. 2015;5:011035. doi:10.1103/PhysRevX.5.011035.

24. You Z, Baskaran A, Marchetti MC. Nonreciprocity as a generic route to traveling states. Proceedings of the National Academy of Sciences USA. 2020;117:19767-19772. doi:10.1073/pnas.2010318117.

25. Fruchart M, Hanai R, Littlewood PB, Vitelli V. Non-reciprocal phase transitions. Nature. 2021;592:363-369. doi:10.1038/s41586-021-03375-9.

26. Dinelli A, O'Byrne J, Curatolo A, Zhao Y, Sollich P, Tailleur J. Non-reciprocity across scales in active mixtures. Nature Communications. 2023;14:7035. doi:10.1038/s41467-023-42713-5.

27. Kreienkamp KL, Klapp SHL. Nonreciprocal alignment induces asymmetric clustering in active mixtures. Physical Review Letters. 2024;133:258303. doi:10.1103/PhysRevLett.133.258303.

28. Chen J, Lei X, Xiang Y, Duan M, Peng X, Zhang HP. Emergent chirality and hyperuniformity in an active mixture with nonreciprocal interactions. Physical Review Letters. 2024;132:118301. doi:10.1103/PhysRevLett.132.118301.

29. Chatterjee S, Mangeat M, Woo CU, Rieger H, Noh JD. Flocking of two unfriendly species: the two-species Vicsek model. Physical Review E. 2023;107:024607. doi:10.1103/PhysRevE.107.024607.

30. Dutta AK, Mangeat M, Rieger H, Paul R, Chatterjee S. Stability of flocking in the reciprocal two-species Vicsek model: effects of relative population, motility, and noise. Physical Review E. 2025;112:024137. doi:10.1103/PhysRevE.112.024137.

31. Digregorio P, Levis D, Suma A, Cugliandolo LF, Gonnella G, Pagonabarraga I. Full phase diagram of active Brownian disks: from melting to motility-induced phase separation. Physical Review Letters. 2018;121:098003. doi:10.1103/PhysRevLett.121.098003.

32. Storn R, Price K. Differential evolution - a simple and efficient heuristic for global optimization over continuous spaces. Journal of Global Optimization. 1997;11:341-359. doi:10.1023/A:1008202821328.

33. Dai C, Glotzer SC. Efficient phase diagram sampling by active learning. Journal of Physical Chemistry B. 2020;124:1275-1284. doi:10.1021/acs.jpcb.9b09202.

34. Chaté H, Ginelli F, Grégoire G, Raynaud F. Collective motion of self-propelled particles interacting without cohesion. Physical Review E. 2008;77:046113. doi:10.1103/PhysRevE.77.046113.

35. Nayak S, Bag P, Ghosh PK, Li Y, Zhou Y, Yin Q, Marchesoni F, Nori F. Diffusion transients in motility-induced phase separation. Physical Review Research. 2025;7:013153. doi:10.1103/PhysRevResearch.7.013153.

36. Zhou Y, Yin Q, Nayak S, Bag P, Ghosh PK, Li Y, Marchesoni F. Visual quorum sensing in chiral suspensions: Hyperuniformity and edge currents. PNAS Nexus. 2025;4(12):pgaf373. doi:10.1093/pnasnexus/pgaf373.

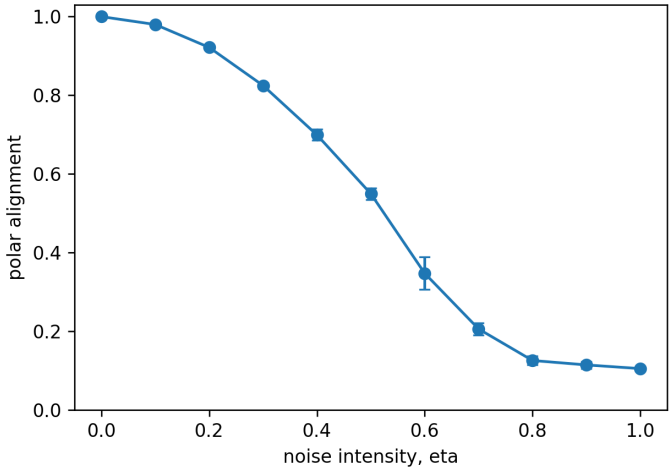
Supplementary Materials

Evolutionary Discovery of the Behavioral Repertoire of Self-Propelled Particle Models

Roujia Wen and Eric Bonabeau

S1. Numerical validation of the simulator

As a basic simulator check, we use a Vicsek-like limit with no pair force or gradient, broad alignment range, and strong alignment. Mean polar order ϕ decreases monotonically from 1.00 at $\eta = 0$ to 0.92 at $\eta = 0.2$, 0.70 at $\eta = 0.4$, 0.35 at $\eta = 0.6$, and 0.10 at $\eta = 1$. This is not a critical-point estimate; it checks that the implementation reproduces the expected loss of alignment with increasing noise.



Supplementary Figure S1. Polar alignment in the Vicsek-like validation limit as a function of noise. Points are means over eight seeds; error bars are one standard deviation.

We also checked hard-core handling and the periodic definition of the group center. After overlap projection, no sampled final state has a persistent center separation below $2r_0$. Translating a cohesive group across a periodic boundary leaves its milling value unchanged to numerical precision, whereas using the arithmetic centroid produces a discontinuity.

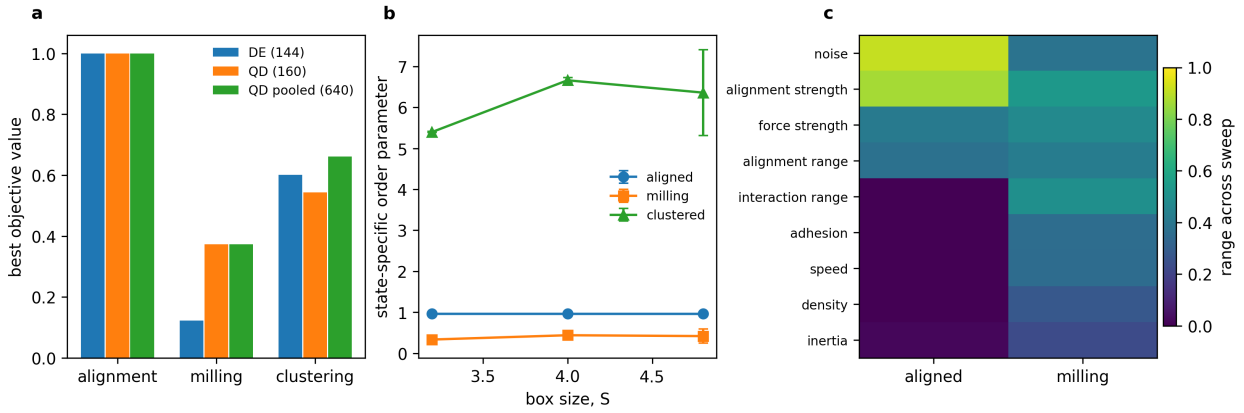
S2. Search convergence, scalar baselines, and robustness checks

For a budget-matched comparison, we use small runs rather than the 64,000-evaluation archives of main text Section 3.1. All methods find nearly perfect alignment. With comparable evaluation counts, a 160-evaluation QD run finds a higher milling objective than 144-evaluation Differential Evolution (0.373 versus 0.123). Differential Evolution gives a slightly higher clustering objective at that budget (0.601 versus 0.543), but increasing QD to 640 evaluations raises clustering to 0.662. The main advantage of QD is not its scalar optimum: one run returns many stable behaviors and their parameter vectors.

CMA-MAE provides a more sample-efficient QD baseline than the simple MAP-Elites emitter used for the main archives. With the same simulator and archive resolution, CMA-MAE reaches alignment-milling coverage of 0.90 by 4,000 evaluations and 0.97 by 8,000, compared with 0.72 and 0.89 for vanilla MAP-Elites. On alignment-clustering, CMA-MAE reaches full coverage by 16,000 evaluations, when vanilla MAP-Elites is at 0.94 (Supplementary Figure S17). Vanilla MAP-Elites is slightly ahead only during the first thousand evaluations. Because the 64,000-evaluation MAP-Elites archives already saturate, CMA-MAE changes the cost of reaching the map, not the final coverage.

The aligned representative remains close to $\phi = 0.96$ as the box size increases from $S = 3.2$ to 4.8 at fixed density. The compact cluster also persists, with normalized clustering between about 5.4 and 6.7. Milling is more size-sensitive: mean m is 0.334 at $S = 3.2$, 0.439 at $S = 4.0$, and 0.418 at $S = 4.8$, with larger seed-to-seed variation at the largest size. The mill is therefore reproducible in the finite system but less robust to size and boundary effects than the aligned state.

These checks complement the larger-N tests in the main text. They are not presented as finite-size scaling to a thermodynamic limit: the purpose is to distinguish conclusions that are already stable to a substantial change in finite system size from those, especially milling, that remain sensitive to walls and box size.

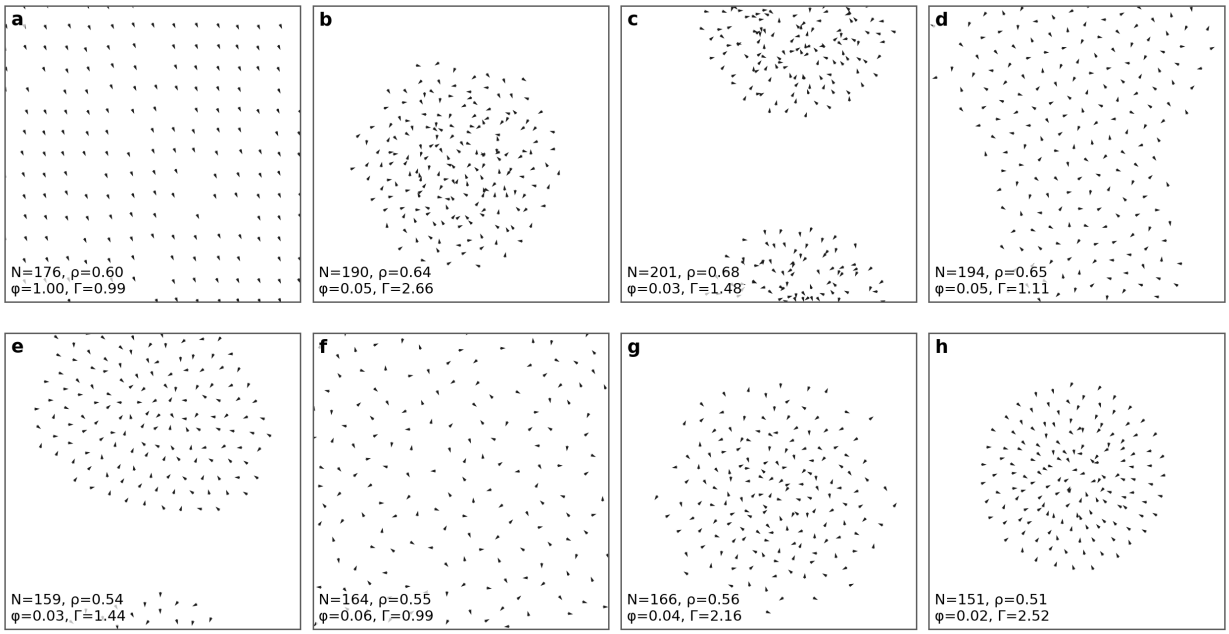


Supplementary Figure S2. Classical single-species search and robustness checks. (a) Best scalar objective from Differential Evolution, a small QD run, and the pooled QD archive. (b) Defining order parameters across box sizes at fixed density. (c) One-at-a-time parameter sweeps.

One-at-a-time sweeps give the same qualitative picture as the local-sensitivity analysis. Alignment responds mainly to noise and alignment strength. Milling depends on a broader combination of alignment range and strength, interaction range, force strength, density, adhesion, speed, and inertia. Main text Section 6 tests these dependencies with the full local response Jacobian rather than isolated parameter sweeps.

S3. High-density repertoire examples

The high-density sample contains several ordered exceptions to the overall density trend. One periodic rule has $\rho = 0.595$, $N = 176$, $\phi = 0.99999$, and $S_2 = 0.99997$. Reflecting-wall aggregates reach $N = 190$ and $\Gamma = 2.67$, with a low-wave-number density peak of 0.322. A periodic density-modulated state has $N = 201$ and density peak 0.391, and another state has $N = 194$ with sixfold bond order $\psi_6 = 0.744$.



Supplementary Figure S3. Eight high-density examples from the density-expanded corpus. Arrows show particle headings. Density fractions range from 0.51 to 0.68 and particle counts from 151 to 201. The panels include a dense flock, wall-supported aggregates, density-modulated states, a sixfold-rich state, a nematic-rich state, and dense near-uniform states.

Supplementary Table S1. Selected high-density examples from the discovery trajectories.

Example	ρ	N	Boundary	φ	m	Γ	S_2	ψ_6	Density peak
Dense aligned	0.595	176	Periodic	1.000	0.020	0.994	1.000	0.039	0.001
Dense wall aggregate	0.640	190	Reflecting	0.044	0.036	2.665	0.065	0.033	0.322
Dense density wave	0.679	201	Periodic	0.027	0.026	1.478	0.038	0.017	0.391
Sixfold-rich	0.654	194	Periodic	0.057	0.030	1.107	0.063	0.744	0.067
Nematic-rich dense	0.538	159	Periodic	0.036	0.058	1.349	0.247	0.193	0.155
Dense near-uniform	0.555	164	Periodic	0.048	0.047	0.990	0.070	0.008	0.001

The dense aligned state shows that high density does not exclude collective order. Strong, long-range alignment can overcome crowding. High-density aggregation and density modulation, by contrast, depend strongly on wall geometry and force parameters. These examples also show why ϕ, m, and Γ alone are not enough to describe the full repertoire.

S4. Training and validation of the learned representation

S4.1 Architecture, training, and independent test protocol

The learned representation uses a deterministic feed-forward autoencoder with a two-dimensional bottleneck. The encoder E maps the 51 standardized trajectory features through hidden layers of width 64 and 24 with hyperbolic-tangent activations to a linear two-dimensional code. The decoder D mirrors this structure back to 51 outputs,

$$\mathbf{z} = E(\widetilde{\mathbf{h}}) = W_3 \tanh\Big(W_2\, tanh(W_1\widetilde{\mathbf{h}} + \mathbf{b}_1) + \mathbf{b}_2\Big) + \mathbf{b}_3, \qquad \widehat{\widetilde{\mathbf{h}}} = D(\mathbf{z}),$$

The encoder widths are $51 \to 64 \to 24 \to 2$ and the decoder widths are $2 \to 24 \to 64 \to 51$. The nonlinear hidden layers allow a nonlinear compression of the descriptor, while the linear bottleneck leaves the two archive coordinates unconstrained. The network contains no batch normalization, dropout, or skip connections, so a trained model maps a given input deterministically to its code.

Training minimizes the mean squared reconstruction error of the standardized descriptor,

$$\mathcal{L} = \frac{1}{M}\sum_{n=1}^{M} \parallel \widetilde{\mathbf{h}}_n - D(E(\widetilde{\mathbf{h}}_n)) \parallel^2,$$

We train with full-batch Adam, learning rate 2×10^{-3}, weight decay 10^{-5}, and a maximum of 450 epochs. The original corpus is split once into 78% training and 22% validation data, stratified over the six density-band $\times$ boundary-condition combinations. Training stops after 50 epochs without an improvement of at least 10^{-5}, and the best validation weights are restored. Four random initializations are trained; the model with the lowest validation error is selected, with trustworthiness used to break ties. We report reconstruction error, trustworthiness, k-nearest-neighbor overlap, comparison with two-component PCA on the same standardized split, and pairwise-distance agreement across random initializations. Bottlenecks of 3, 4, and 6 dimensions test how much information is lost by the two-dimensional display. Because the 22% partition is used for both early stopping and initialization selection, it is a validation set, not an independent test set.

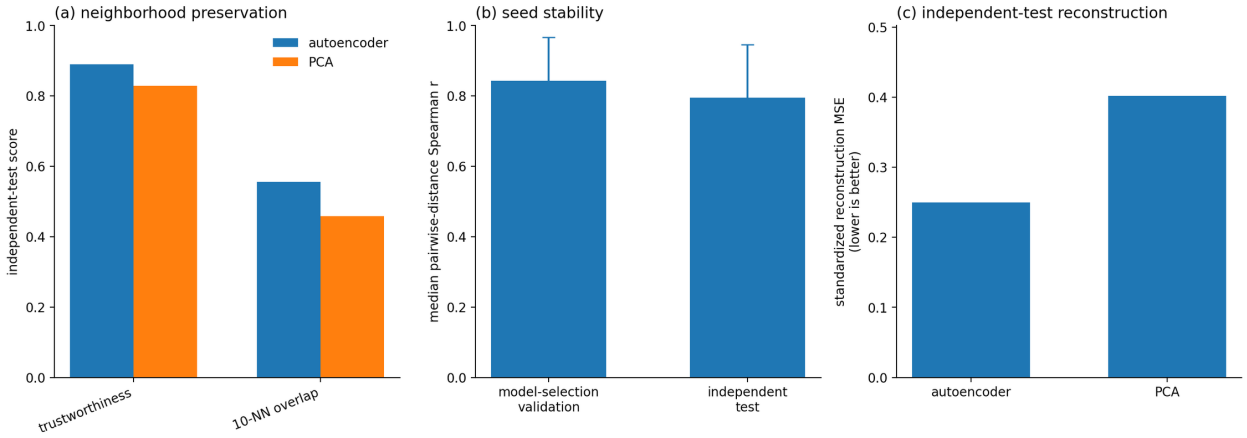
The independent test was generated only after the architecture, training procedure, model-selection rule, and two-dimensional representation had been fixed. It contains 72 new Sobol rules with the same nine parameter bounds, simulator settings, three equal density strata, and 4:1 periodic-to-reflecting allocation as the original 360-rule corpus, but with a new Sobol seed (20260825) and disjoint simulation seeds. Re-running the four original training pipelines on the original 78/22 split again selects seed 2, and the resulting coordinates on the original corpus are numerically identical to the archived coordinates. We then score that selected model once on the 72 new rules. PCA is fit only on the original training partition and evaluated on the same test set. For main text Section 5.3, predictors from the classical eight-vector are fit to the original 360 rules after the learned representation is fixed and are evaluated once on the 72-rule test set. No representation or predictor is selected using the test results.

S4.2 Independent validation and latent dimensionality

On the independent 72-rule test corpus, the selected two-dimensional autoencoder has standardized reconstruction error 0.250, trustworthiness 0.890, and ten-neighbor overlap 0.556. The

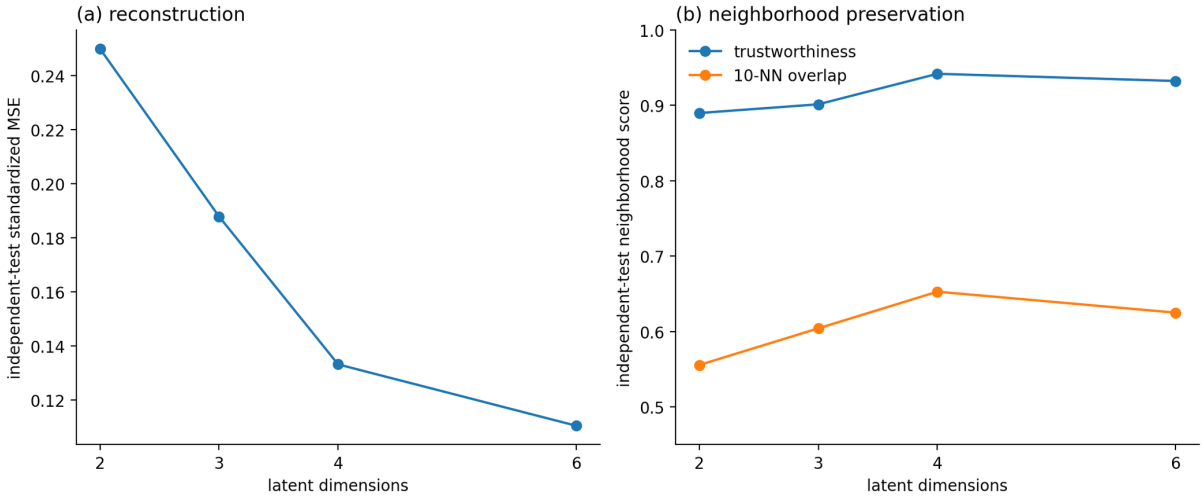
corresponding model-selection values were 0.336, 0.898, and 0.588; the training reconstruction error was 0.168. Two-component PCA, fit only on the original training partition, explains 57.2% of the training descriptor variance and gives test reconstruction error 0.402, trustworthiness 0.829, and ten-neighbor overlap 0.458. The autoencoder therefore keeps its neighborhood-preservation advantage on data that were not used for model selection.

The geometry is also reasonably stable across random initializations. On the model-selection validation set, the median Spearman correlation between pairwise latent distances is 0.843, with a range of 0.720-0.967 across the four runs. On the independent 72-rule test set, the median is 0.795 with range 0.661-0.945. The two-dimensional embedding is not unique, but nearby and distant rules are ordered similarly across independently trained models.



Supplementary Figure S4. Independent validation of the two-dimensional single-species representation. (a) Trustworthiness and ten-neighbor overlap for the autoencoder and PCA on the 72-rule test set. (b) Pairwise-distance agreement across four autoencoder initializations on the model-selection validation and independent test sets; error bars show the observed range. (c) Test reconstruction error. The 72 test rules were generated after all representation choices were fixed.

Higher-dimensional bottlenecks retain more information. On the independent test set, reconstruction error at latent dimensions 2, 3, 4, and 6 is 0.250, 0.188, 0.133, and 0.111. Trustworthiness is 0.890, 0.901, 0.942, and 0.932, and ten-neighbor overlap is 0.556, 0.604, 0.653, and 0.625. Most of the gain in neighborhood preservation is reached by four dimensions. Six dimensions improve reconstruction further but not the neighborhood metrics. We therefore use two dimensions for visualization and archiving, not because the behavioral repertoire is intrinsically two-dimensional.



Supplementary Figure S5. Effect of latent dimension on the independent 72-rule test set. Reconstruction error decreases from two to six dimensions, while neighborhood preservation improves mainly through four dimensions. Model initialization at each dimension is selected only on the original validation partition.

S4.3 Interpretation of validation metrics and display choices

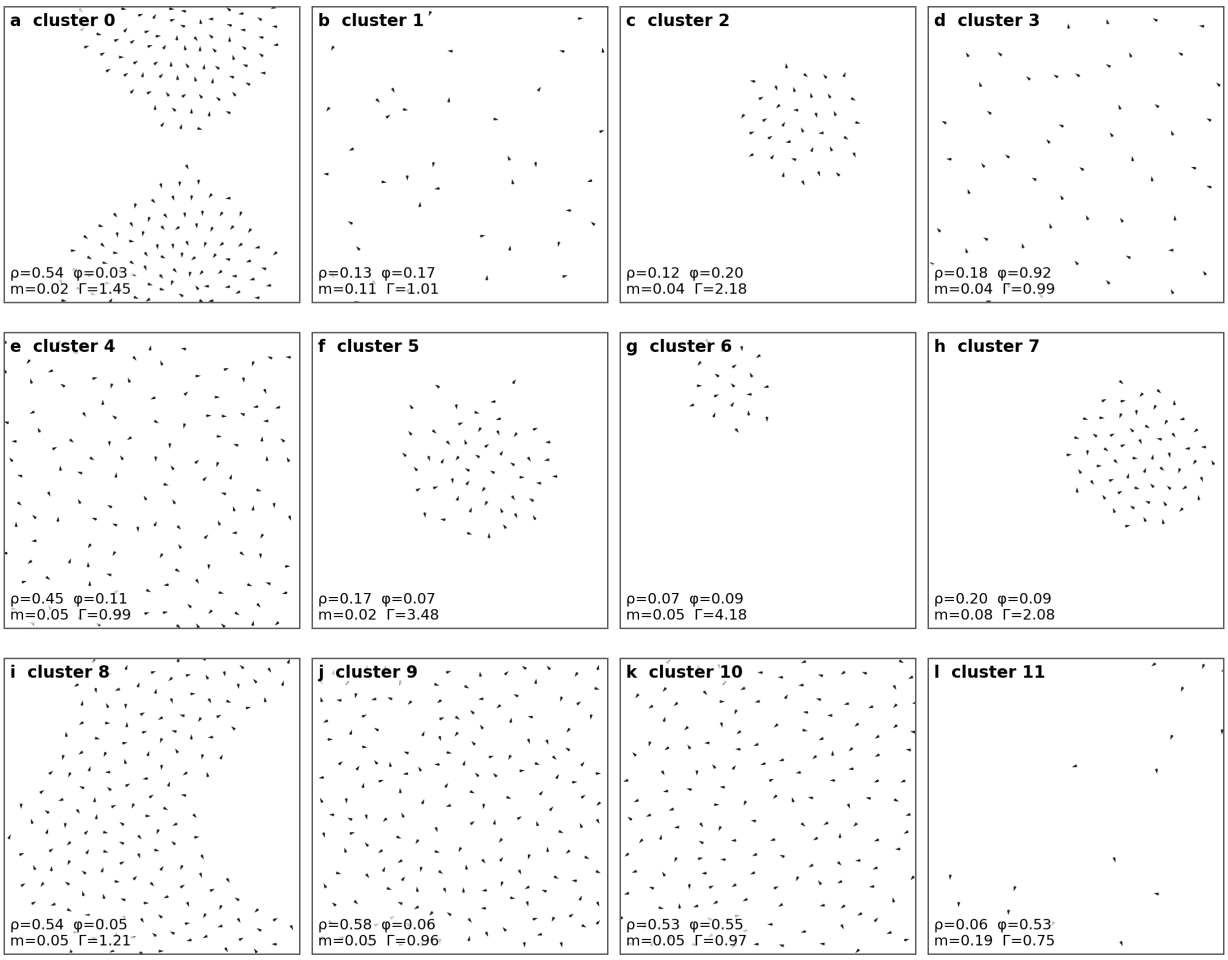
Trustworthiness measures whether neighbors introduced by the low-dimensional embedding were genuinely nearby in the original descriptor space; values approach one when false neighbors are rare. Ten-neighbor overlap is the mean fraction of the ten nearest neighbors in the full descriptor that remain among the ten nearest neighbors in the embedding. Reconstruction error measures how much standardized descriptor information is lost after encode-decode compression. These metrics therefore test complementary properties and are evaluated on the independent 72-rule corpus after representation choices are fixed.

The 20 by 20 grid and the twelve-region k-means partition are downstream visualization choices. The grid is used for coverage accounting only. The twelve clusters are used to select a manageable set of examples and are explicitly not thermodynamic phases; fifty k-means restarts reduce sensitivity to an unfavorable centroid initialization. Neither choice changes the autoencoder coordinates, the independent-test metrics in Table S10, the local-sensitivity Jacobians, or the reciprocal versus non-reciprocal controls. Robustness of the learned geometry itself is addressed more directly by the independent test, the four-initialization pairwise-distance agreement, and the latent-dimension audit: neighborhood preservation improves mainly through four dimensions, while the two-dimensional map remains a deliberately compressed display.

S5. Single-species regions in the learned map

We use twelve k-means regions as a compact summary of the learned plane; the silhouette coefficient is 0.426. Regions with fewer than five rules are tabulated but not interpreted individually. Region 3 contains coherent flocks with mean $\phi = 0.871$, and region 10 contains mixed-density polar states with mean $\phi = 0.498$. Region 2 has sixteen low-density, strongly density-modulated aggregates with mean $\Gamma = 2.61$, while region 5 contains reflecting-wall aggregates with $\Gamma = 2.45$. Region 4 is a broad, dense, nearly uniform weak-order region, and region 8 is denser and has

stronger mean sixfold order. Regions 1 and 11 contain dilute or wall-influenced moving states with relatively high milling.



Supplementary Figure S6. One representative configuration from each of twelve learned single-species regions. Arrows show headings. The regions include coherent flocks, diffuse moving states, compact and wall-supported aggregates, density-modulated states, dense homogeneous states, dense structured states, and mixed polar states.

Supplementary Table S2. Selected single-species learned regions. Values are cluster means.

Learned region	Count	Mean ρ	Mean φ	Mean m	Mean Γ	Mean density peak	Interpretation
2	16	0.111	0.184	0.110	2.608	0.608	Compact/density-modulated aggregates
3	20	0.251	0.871	0.048	0.934	0.042	Coherent flocks
4	90	0.455	0.095	0.056	0.943	0.009	Dense, near-uniform weak order
5	16	0.272	0.111	0.076	2.453	0.321	Reflecting-wall aggregates

Learned region	Count	Mean ρ	Mean φ	Mean m	Mean Γ	Mean density peak	Interpretation
8	36	0.518	0.068	0.043	1.213	0.091	Dense locally structured states
10	31	0.414	0.498	0.057	1.071	0.044	Mixed-density polar states
11	17	0.108	0.489	0.148	1.609	0.163	Dilute/wall-influenced moving states

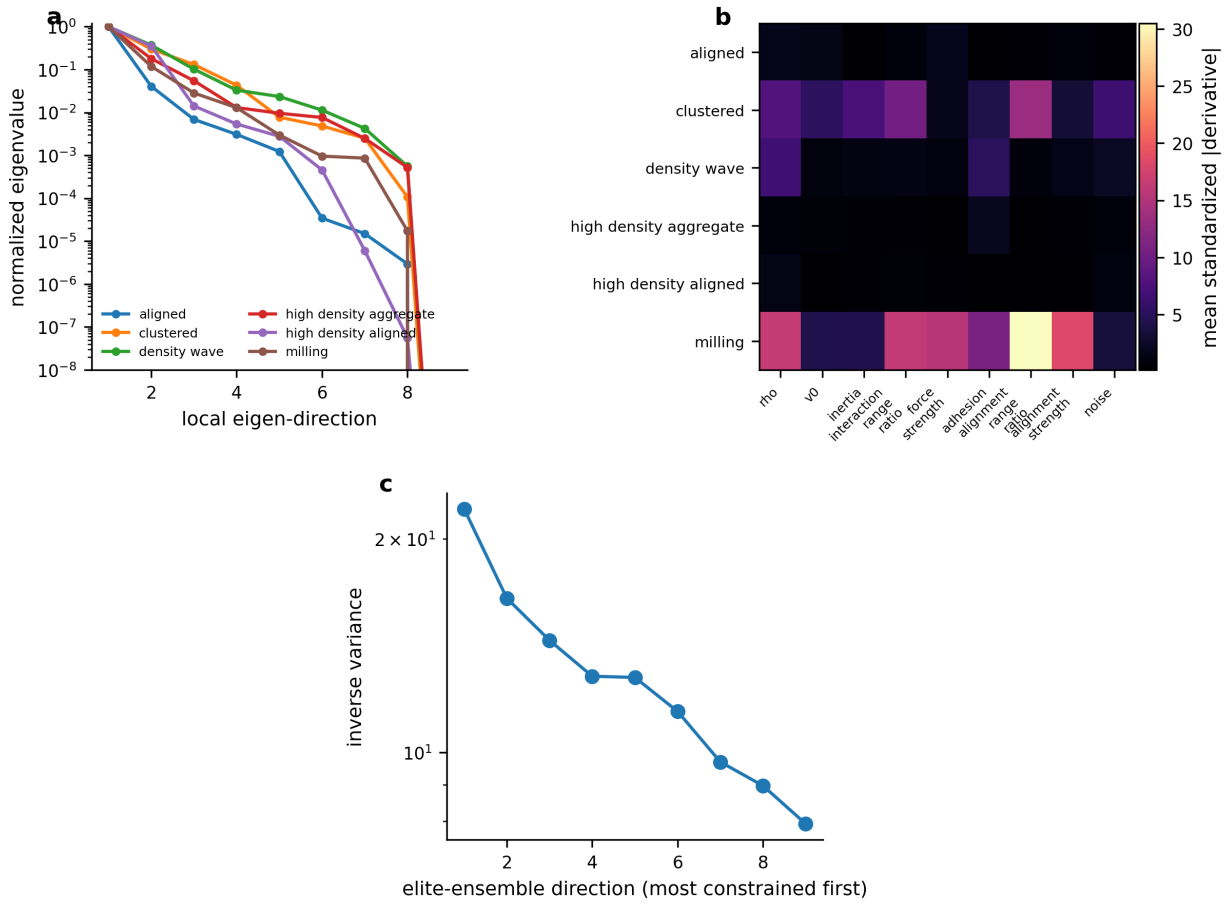
The learned map separates differences in morphology and density that are compressed by the classical scalar measures. It also provides representative parameter sets across the repertoire without requiring a phase list to be specified in advance.

S6. Additional local-sensitivity results

S6.1 Resolved local sensitivity directions

With only two matched seeds on each side of a perturbation, the local sensitivity spectra appear to span 3.3-7.3 decades (Supplementary Table S6). Most of that apparent range is not resolved. Main text Section 6.2 repeats the calculation with more seeds and an explicit seed-noise threshold; only the leading part of each spectrum remains distinguishable from stochastic simulation noise. We therefore interpret the resolved directions, not the raw decade span.

The leading resolved directions differ by regime. We report a loading only when its 90% bootstrap interval, based on sixteen seeds per perturbation side, excludes zero (Supplementary Table S6). For the mill, the leading direction loads on density ($0.65\,[0.51, 0.77]$), force strength ($0.56\,[0.36, 0.72]$), and adhesion ($0.32\,[0.15, 0.49]$). For the clustered state it loads on interaction range ($0.44\,[0.19, 0.68]$), adhesion (-0.41), and noise (-0.33). The density-wave direction is dominated by density ($-0.97\,[-0.99, -0.94]$), and the high-density aggregate by adhesion ($-0.88\,[-0.90, -0.87]$). The dense aligned rule is different: its leading eigenvector is not reliably determined at this seed count, so we do not assign it a stiff parameter combination. This is also why the two nominal aligned representatives are merged: the same dense-flock rule gave different leading eigenvectors in separate two-seed calculations.

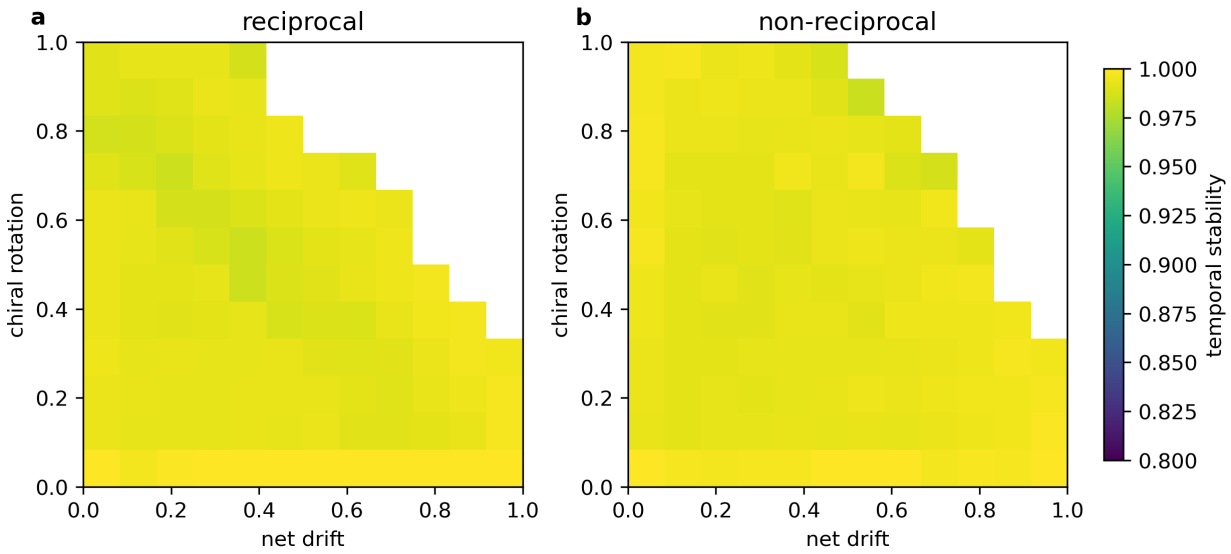


Supplementary Figure S7. Local parameter sensitivity. (a) Normalized local sensitivity spectra. (b) Mean absolute standardized derivatives by regime and parameter. (c) Inverse-variance spectrum of elite parameters in the learned archive.

S7. Supporting two-species analyses

S7.1 Short-horizon reciprocal and non-reciprocal archives

The reciprocal two-species family is the necessary baseline because reciprocal alignment can already produce flocking [29,30]. After 64,000 evaluations, the reciprocal drift-chirality archive occupies 111 of 144 cells (77.1%) and the non-reciprocal archive 118 cells (81.9%). Both contain short trajectories with chirality near one and both populate the low-drift, high-chirality region. At this horizon, archive coverage does not distinguish reciprocal from non-reciprocal dynamics.



Supplementary Figure S8. Drift-chirality archives for two particle types after 64,000 evaluations. (a) Reciprocal cross-species couplings, 77.1% coverage. (b) Full family with antisymmetric alignment and force components, 81.9% coverage. Both short-horizon archives populate the high-chirality region; persistent rotation is distinguished only by the longer matched-control tests in main text Section 7.2.

The reason is the 200-step chirality measure used by the archive. Long-horizon validation in main text Section 7.2 shows that many apparently chiral reciprocal rules are transient: over twenty seeds and 800-1,200 steps, the best reciprocal rule has median chirality near 0.01, while the selected non-reciprocal rule remains near 0.7. A short-trajectory archive can therefore place transient and persistent rotation in the same behavior cells. The mechanism is resolved only by the longer matched-control experiments, not by archive coverage.

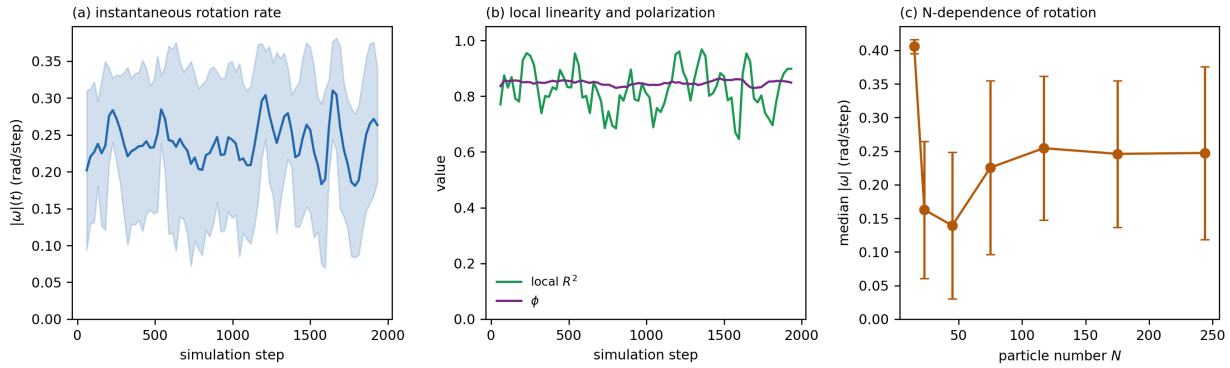
S7.2 Long-horizon and finite-size diagnosis of rotation

The composite chirality score used for archive construction combines three factors: absolute angular speed, saturated at 0.05 rad per step; the regression coefficient of determination R-squared; and mean polarization. The angular-speed scale is a saturation scale rather than an inferred threshold. R-squared penalizes irregular or reversing accumulated-heading trajectories that can have a nonzero fitted slope without persistent rotation, and mean polarization downweights trajectories whose global heading is poorly defined. For this reason, all mechanistic conclusions below are checked against the underlying windowed or fitted angular velocity rather than the composite score alone.

The composite chirality score falls at long horizons even though the underlying rotation persists. To identify the source of the decline, we run the selected rule for 2,000 steps and measure the windowed rotation rate $\mid \omega \mid (t)$, local linearity $R^2(t)$, and polarization $\phi(t)$ (Supplementary Figure S9). All three are stationary. The windowed $\mid \omega \mid$ remains near 0.24 radians per step from the beginning to the end of the run (early and late means 0.239 and 0.243), R^2 fluctuates around 0.83 without a downward trend, and ϕ remains near 0.85. The system therefore does not slow, depolarize, or lose local rotational coherence. The windowed rate is higher than the 0.163 radians per step obtained from a single regression over the full 800-step trajectory because the global fit averages

over slow drift and occasional changes in the accumulated heading. Those effects reduce the full-trajectory R^2 term in the composite score even while local rotation continues. The long-horizon drop in the composite chirality score is therefore a property of the measurement, not a decay of the rotating state.

The same rule does not show a systematic loss of rotation as particle number increases. At fixed parameters, varying the box size gives N from 15 to 244. Median $\mid \omega \mid$ is 0.41, 0.16, 0.14, 0.23, 0.25, 0.25, and 0.25 radians per step at $N = 15, 23, 45, 75, 117, 175, 244$, respectively, and is approximately 0.24 radians per step for $N \gtrsim 75$ (Supplementary Figure S9c). The spread across seeds remains large, so this supports a persistent finite-system regime rather than a thermodynamic-limit claim.



Supplementary Figure S9. Long-horizon rotation of the selected non-reciprocal rule. (a) Windowed rotation rate remains stationary to 2,000 steps. (b) Local linearity and polarization show no downward trend. (c) Median rotation rate does not decrease toward zero as particle number increases from 15 to 244 and approaches about 0.24 rad/step at larger N.

S7.3 Species-aware learned representation

S7.3.1 Validation of the species-aware representation

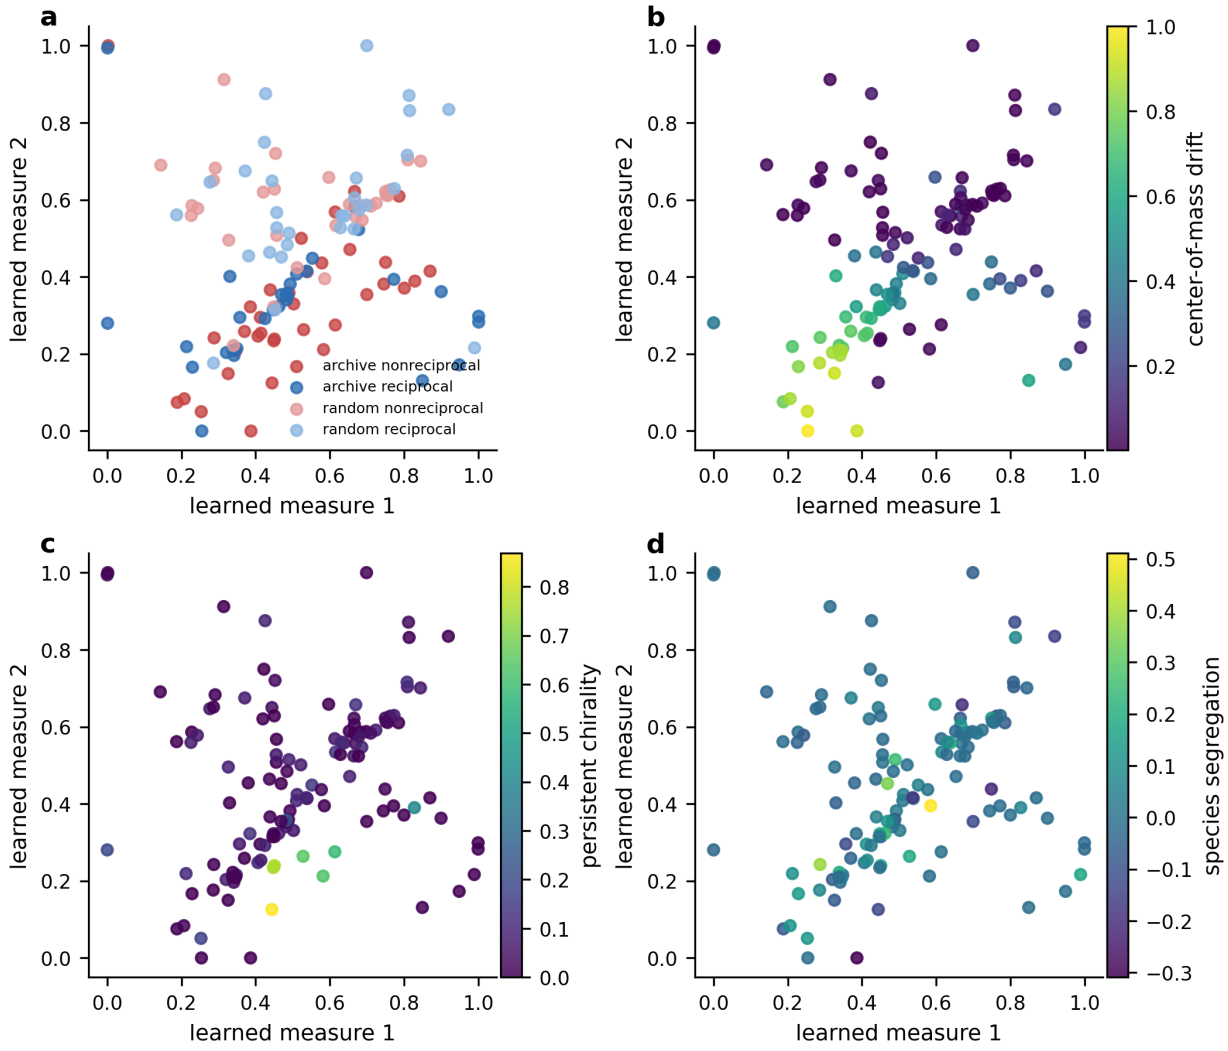
The species-aware descriptor adds species-resolved pair distributions and heading correlations, species polarizations, relative heading, and inter-species center separation to the 51 single-species features. In the original model-selection split, the selected two-dimensional autoencoder has validation reconstruction error 0.577, trustworthiness 0.930, and ten-neighbor overlap 0.636. A separate source-stratified 60/20/20 audit again selects seed 2 using only its validation partition. On the untouched 26-rule test partition, reconstruction error is 0.407, trustworthiness 0.845, and ten-neighbor overlap 0.627. Pairwise latent-distance stability across the four audit initializations has median Spearman correlation 0.918, with range 0.877-0.957. The test set is small, so these estimates are less precise than the single-species results, but the representation is reproducible on held-out rules.

The clean-split audit also shows that two dimensions are incomplete. Increasing the bottleneck to six dimensions lowers test reconstruction error from 0.407 to 0.258 and raises ten-neighbor overlap from 0.627 to 0.723; trustworthiness changes only from 0.845 to 0.847. We therefore use two dimensions as a display, not as a complete coordinate system for mixture behavior.

S7.3.2 The learned map captures behavior, not reciprocity

The two-dimensional mixture map is dominated by collective translation. Learned measure 2 correlates with polarization at $\mid r_s \mid = 0.839$ and with center-of-mass drift at $\mid r_s \mid = 0.762$. Learned

measure 1 correlates with clustering at $|\ r_s\ | = 0.438$. Chirality and segregation have only weak marginal correlations with either axis.



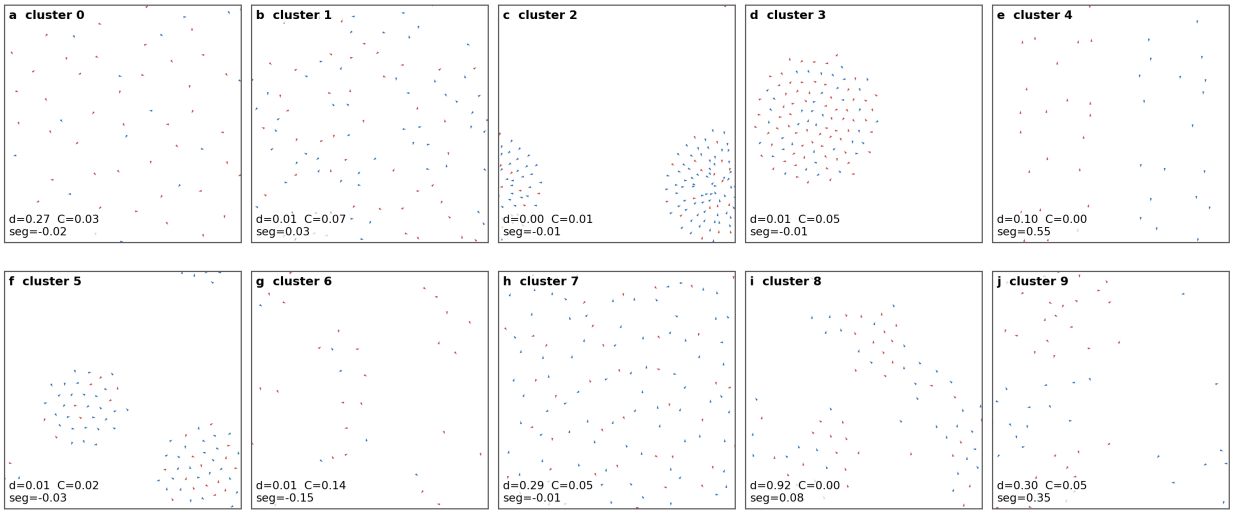
Supplementary Figure S10. Species-aware learned behavior space. Panels show (a) reciprocal versus non-reciprocal source, (b) center-of-mass drift, (c) persistent chirality, and (d) species segregation. Reciprocal and non-reciprocal rules overlap when they produce similar trajectories.

The learned coordinates do not identify reciprocity. A logistic classifier based only on the two coordinates has mean five-fold AUC 0.441, essentially chance; using the full descriptor raises AUC only to 0.611. This is expected because reciprocity is a property of the interaction matrix, and different interaction laws can generate similar trajectories.

Chirality is easier to detect, although still weakly represented in two dimensions. The two learned coordinates classify the top chirality quintile with mean AUC 0.637, and the full descriptor reaches 0.715. With a six-dimensional bottleneck, the largest chirality-axis correlation increases from 0.059 to 0.235. Rare dynamical modes can therefore be compressed when the reconstruction objective is dominated by more common variation such as drift and polarization.

S7.3.3 Mixture regions and examples

Ten regions summarize the mixture corpus. Region 8 contains strongly translating flocks (mean drift 0.819, mean $\phi = 0.855$) and includes both reciprocal and non-reciprocal rules. Region 0 contains moderate translation and more chirality (mean drift 0.432, mean chirality 0.212) and is enriched, but not exclusively populated, by non-reciprocal rules. Regions 2, 3, and 5 contain low-drift clustered states. Region 4 is entirely reciprocal in this sample, but its defining behavioral feature is moderate drift rather than reciprocity itself.



Supplementary Figure S11. One representative configuration from each of ten learned two-species regions. Blue and red indicate species. The regions include translating mixed flocks, weakly moving gases, clustered mixtures, partially segregated aggregates, and curved or chiral motion.

The learned mixture map and the targeted drift-chirality archive answer different questions. Reconstruction-based coordinates summarize the dominant behavioral variation, whereas hand-designed measures can retain rare dynamics that would otherwise receive little representation capacity. In practice, the two approaches are complementary [14,15].

S8. Directed searches for extreme behaviors

The QD archives map many behaviors at once. Here we ask a different question: for each chosen objective, what is the most extreme state found within the search domain? We consider eleven objectives across the single- and two-species models, extend the density range into the dense regime, optimize several pattern measures that were only diagnostics in main text Section 4, and record a representative animation for each resulting champion.

S8.1 Objective definitions

Two trajectory-level quantities used in these objectives are defined here. Temporal stability, the archive quality measure used throughout the paper, is one minus the mean within-trajectory variation of alignment, milling, and scale-normalized clustering,

$$T = 1 - \frac{1}{3}\left[\sigma_\phi + \sigma_m + min\left(\frac{\sigma_\Gamma}{max(\langle\Gamma\rangle, 1)}, 1\right)\right],$$

where σ_x is the standard deviation of measure x over the sampled second half of the trajectory. Large-scale density modulation is measured by the low-wavenumber structure factor and its maximum,

$$S(\mathbf{k}) = \left| \frac{1}{N} \sum_{j=1}^{N} e^{\,i\,\mathbf{k}\cdot\mathbf{x}_j} \right|^2, \qquad \mathbf{k} = \frac{2\pi}{L}\,(k_x, k_y)\,,$$

$$P = \max_{1 \le k_x^2 + k_y^2 \le K} \langle S(\mathbf{k}) \rangle_{\text{shell}},$$

where the maximum is taken over the eight lowest-wavenumber shells. Local-density heterogeneity is the coefficient of variation of the neighbor count,

$$c_\rho = \frac{sd_i\, k_i(r_c)}{\langle k_i(r_c) \rangle}, \qquad r_c = 4\, r_0,$$

where $k_i(r_c)$ is the number of particles within r_c of particle i. The eight single-species objectives are

$$\mathbf{F}_{\text{single}} = \left(\phi, \quad m(1-\phi), \quad \frac{1}{10} min(\Gamma, 10), \quad S_2(1-\phi), \quad \psi_6, \quad P, \quad \min\left(\frac{c_\rho}{2}, 1\right), \quad (1-\phi)(1-m)\,T \right),$$

These objectives target an aligned flock, a milling vortex, a dense cluster, a nematic-order state, a hexatic state, strong low-wavenumber density modulation, high local-density heterogeneity, and persistent low global order. The factor $1-\phi$ suppresses translating flocks in the milling and nematic objectives, while $(1-\phi)(1-m)T$ favors states that remain disordered over time rather than passing briefly through low order. The three two-species objectives use the drift, chirality, and segregation measures defined in the main text,

$$\mathbf{F}_{\text{two}} = (\, D, \quad \chi, \quad \sigma\,),$$

for coherent translation, persistent rotation, and species demixing. Four of the single-species objectives use measures that were not archive coordinates in the main analysis: $S_2(1-\phi)$, ψ_6, P, and c_ρ, corresponding to nematic order, sixfold bond order, low-wavenumber density power, and local-density heterogeneity.

The labels used for these objectives are convenient shorthand, not physical diagnoses. In particular, the objectives called "nematic lanes," "density waves," "phase separation," and "active turbulence" measure nematic order with low polarization, low-wavenumber density modulation, local-density heterogeneity, and persistent low global order. Demonstrating lane topology, propagating-wave dispersion, phase coexistence, or active-turbulent statistics would require additional diagnostics.

S8.2 Search and validation of objective champions

The directed search uses the density range $0.05 \le \rho \le 0.70$, with the other bounds unchanged from the density-expanded corpus. For each objective, we first score candidates already present in the classical archives and the 360-rule corpus. The best candidates seed an elitist hill climb in normalized parameter space: Gaussian mutation steps shrink over generations, candidates are

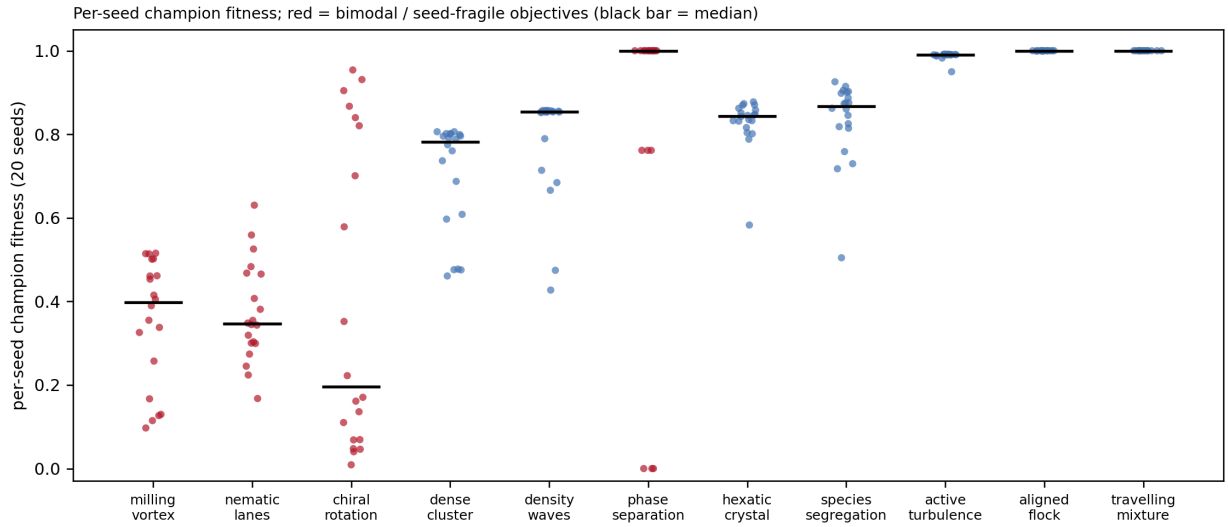
evaluated on independent seeds, and the best population is retained. The champion is the highest-scoring parameter vector found by either stage. It is then evaluated on independent validation seeds. For visualization, we use the seed whose fitness is closest to the validation median rather than the best-looking seed.

Supplementary Figure S13 and Table S3 show the resulting champions. Most objectives produce a clear extreme state, but their seed-to-seed robustness differs substantially.

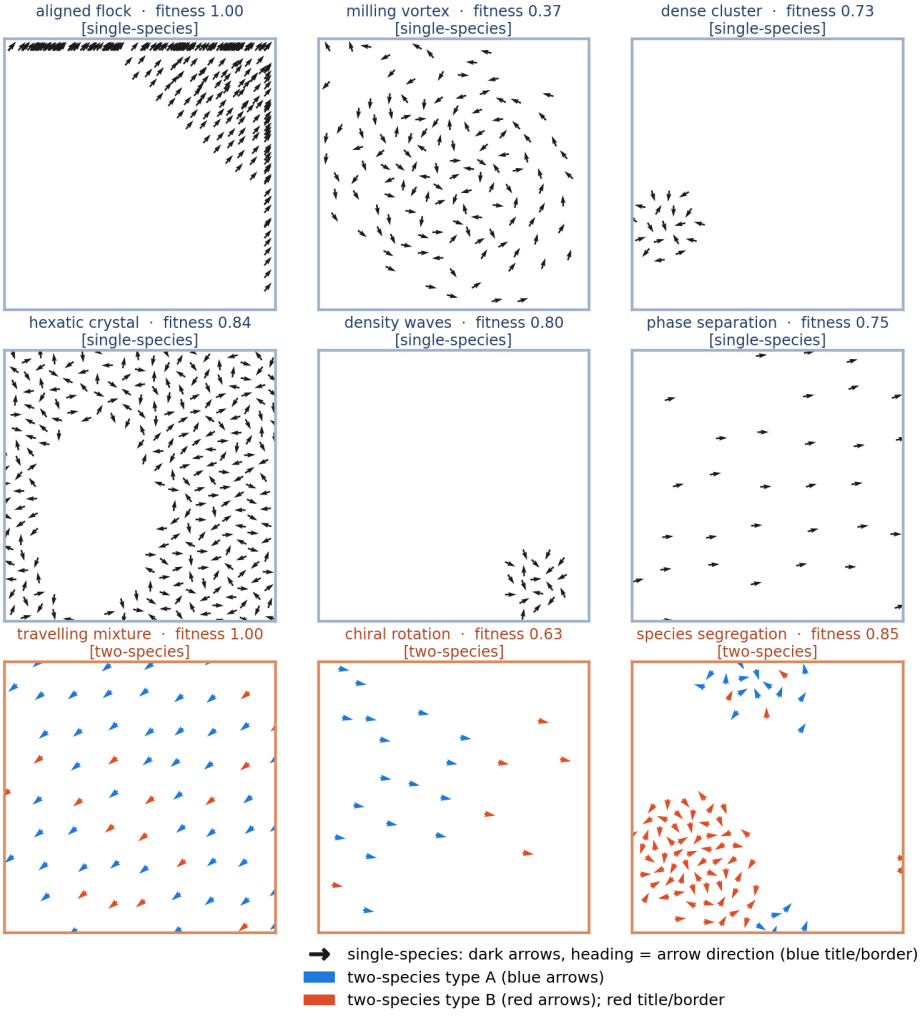
Several objective functions are capped, so a reported value at one can represent either a naturally saturated measure or an imposed ceiling. The aligned flock reaches $\phi = 1.000$ and the translating mixture reaches $D = 1.000$, both genuine saturation of bounded quantities. The dense-cluster champion has $\Gamma \approx 7.9$, so its normalized fitness $\Gamma/10$ remains below one. The density-heterogeneity champion has $c_\rho \gtrsim 2$, beyond the cap in $min(c_\rho/2, 1)$, so its reported fitness of one hides additional heterogeneity.

Many champions also lie near the edge of the parameter box. Single-species champions have, on average, 2.1 of nine parameters within 1% of a bound; two-species champions have 2.7 of fourteen (Supplementary Table S8). The clearest cases are the $\phi = 1.000$ aligned flock and the $D = 1.000$ translating mixture, both with $\eta = 0$; force strength is also at its lower bound, and mixture inertia is at its lower bound. Their zero seed variance reflects a nearly deterministic corner of the search space rather than exceptional dynamical robustness. The active-turbulence and chiral-rotation champions also have $\eta = 0$. By contrast, most of the newly optimized pattern objectives are not pinned to many bounds. The hexatic champion reaches $\psi_6 = 0.84$ at $\rho = 0.70$, its only bound, and forms close-packed crystalline domains separated by voids. Representative trajectories for the low-wavenumber-density, density-heterogeneity, and persistent-low-order objectives show a traveling band, dense and sparse regions, and multi-vortex-like flow, respectively. These visual descriptions do not by themselves establish a density-wave phase, thermodynamic phase separation, or active turbulence.

Seed-to-seed distributions reveal an important limitation of several champions. Over twenty validation seeds, milling, the nematic objective, the density-heterogeneity objective labeled phase separation, and especially chiral rotation are not realized reliably on every seed (Supplementary Figure S12). The chiral champion has median fitness 0.20, interquartile range 0.07-0.83, and individual values from 0.01 to 0.95; 60% of seeds fall below half its best value. The phase-separation objective is strongly bimodal: most seeds reach one, while roughly one in six falls near zero. Milling and the nematic objective are more broadly distributed. For these reasons Table S3 reports medians and interquartile ranges, and Table S9 gives the full distributions. The aligned flock, translating mixture, persistent-low-order objective, hexatic state, dense cluster, low-wavenumber-density objective, and species-segregation objective are realized on essentially every seed. All animations use the median-fitness seed.



Supplementary Figure S12. Champion fitness over twenty independent validation seeds. Black bars are medians. Milling, the nematic objective, the phase-separation objective, and chiral rotation have substantial low-fitness tails; the other objectives are reproduced on nearly every seed.



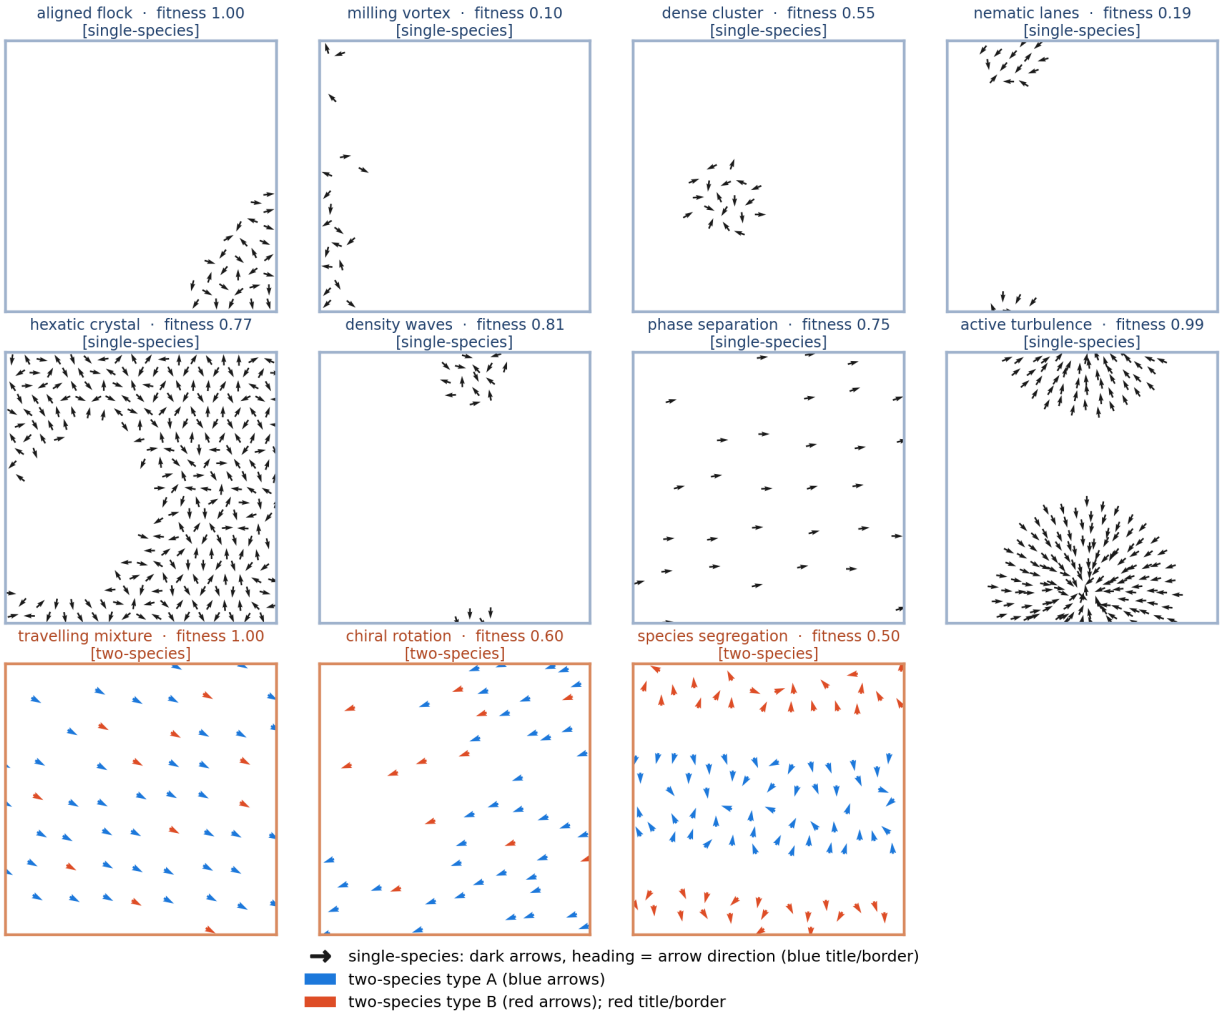
Supplementary Figure S13. Objective champions at representative validation seeds. Single-species panels use dark heading arrows; two-species panels color arrows by particle type. Titles show validated fitness. The panels labeled hexatic crystal, density waves, phase separation, and active turbulence are named for their objective functions and should be interpreted through those defining measures.

Objective	Type	Default boundary	Median [IQR] (20 seeds)	Defining measure
Aligned flock	single	reflecting	1.000 [1.000, 1.000]	φ
Milling vortex	single	reflecting	0.397 [0.235, 0.471]	$m(1-\varphi)$
Dense cluster	single	periodic	0.782 [0.606, 0.800]	$\Gamma/10$
Nematic lanes	single	periodic	0.346 [0.300, 0.466]	$S_2(1-\varphi)$
Hexatic crystal	single	periodic	0.844 [0.828, 0.859]	ψ_6
Density waves	single	periodic	0.853 [0.771, 0.855]	P
Phase separation	single	periodic	1.000 [0.762, 1.000]	c_ρ
Active turbulence	single	periodic	0.990 [0.990, 0.991]	$(1-\varphi)(1-m)T$
Travelling mixture	two	periodic	1.000 [1.000, 1.000]	drift D
Chiral rotation	two	periodic	0.196 [0.069, 0.825]	chirality χ
Species segregation	two	periodic	0.868 [0.817, 0.899]	segregation σ

Supplementary Table S3. Directed-search champions, reported as median and interquartile range over twenty independent validation seeds. Full seed distributions are in Supplementary Figure S12 and Table S9. Capped objectives report the capped fitness; uncapped defining measures are discussed in Section S8.2. Objective names are shorthand for the corresponding functionals.

S8.3 Boundary-condition dependence

Only the aligned-flock and milling objectives use reflecting walls by default, so these are the two objectives for which wall removal is a direct ablation. We re-optimize both under periodic boundaries and validate over twenty seeds (Supplementary Figure S14). Alignment is unchanged: the median fitness remains 1.00, although the reflecting system packs against a corner while the periodic flock spreads across the box. Milling drops from a median of 0.40 to 0.10 under periodic boundaries. The rotating vortex is therefore strongly wall-stabilized, consistent with the boundary choice used for the classical milling archive. The other objectives already use periodic boundaries and do not rely on reflecting walls by construction.



Supplementary Figure S14. Boundary-condition test for the two objectives that use reflecting walls by default. Alignment remains near one under periodic boundaries, while milling drops sharply, confirming its wall dependence. The other objectives are periodic by construction.

S9. Additional transition and crossover checks

S9.1 Finite-window and autocorrelation checks for the noise cut

The first check is the location of the fluctuation maximum. η^* remains near 0.50 across all five system sizes and is unchanged when the trajectory length is doubled from 220 to 440 steps. This lack of detectable drift is evidence consistent with a transition rather than a strongly size-dependent crossover, although it does not rule out every smooth crossover. The location, rather than the peak height, is the main finite-size evidence used in the paper.

We next ask whether the growth of the susceptibility peak is physical or caused by the finite observation window. For a window-averaged order parameter we write $\chi = A + B/T$, where T is the run length. Using $T = 220$ and 440 steps gives the window-independent estimate $A = 2\chi(440) - \chi(220)$. This term is small: at the largest size it is about one quarter of the measured peak, and at intermediate sizes it is near zero or slightly negative. The peak ratio $\chi(440)/\chi(220)$ lies between about 0.35 and 0.65, close to 0.5 for a dominant finite-window contribution and well

below one. A third run at 880 steps confirms the $A + B/T$ form: the measured peak is 0.58, compared with 0.54 predicted from the two shorter runs. We therefore do not fit a power law in N. Doubling T would not remove the ambiguity because a B/T contribution can itself vary with N through the relaxation time $\tau(N)$, leaving an apparent exponent in N unchanged while halving the amplitude. Because the first half of every run is discarded as burn-in, doubling T also doubles the burn-in interval; B/T therefore combines residual relaxation and sampling variance. Both are finite-window effects, and we do not attempt to separate them within B.

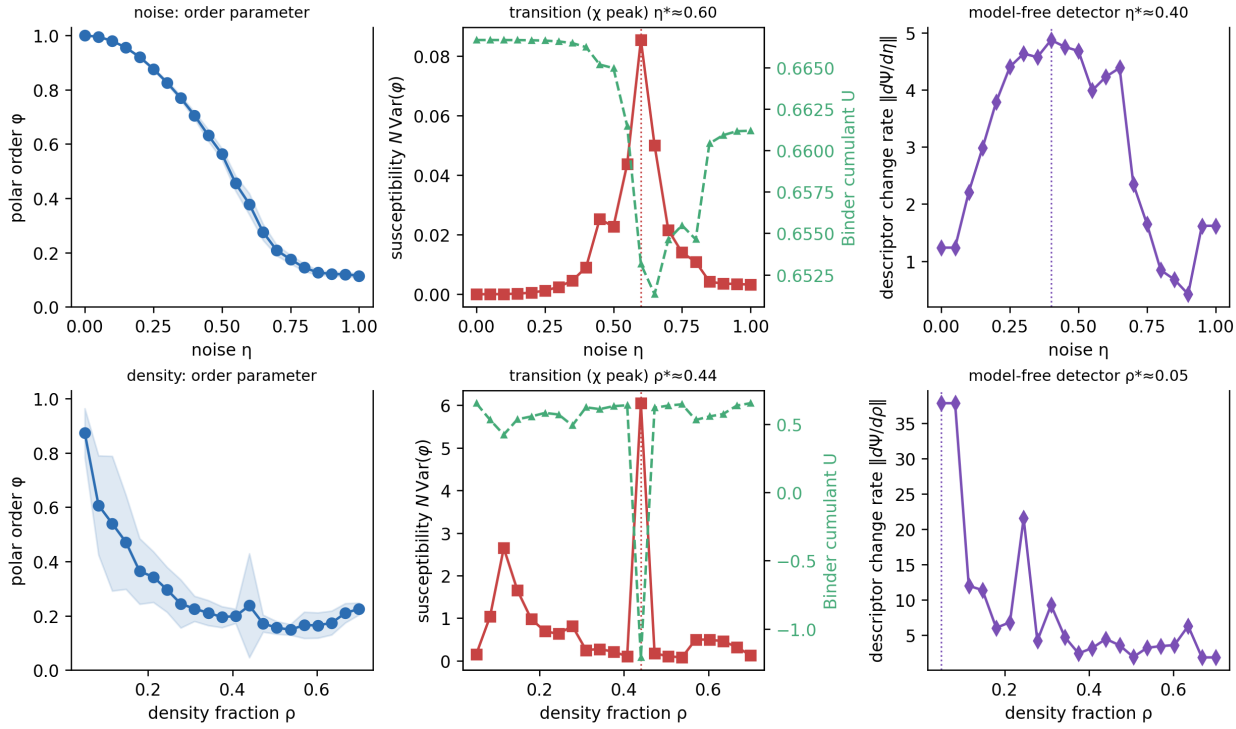
The autocorrelation data explain this finite-window contribution. In the disordered regime, the across-seed variance decreases with N at $\eta = 0.90$. Deep in the ordered regime it instead increases with N at fixed run length. The integrated autocorrelation time of ϕ rises from about 2 steps at $N = 29$ to about 22 steps at $N = 320$. The 110-step measurement window therefore contains roughly forty-five autocorrelation times at the smallest size but only about five at the largest. Large systems provide fewer effectively independent samples per trajectory, which inflates the fixed-window variance.

Finally, the per-seed distribution of ϕ at η^* is unimodal for every size. At the sizes tested, we therefore find no coexistence bimodality of the type associated with the first-order scenario in the cohesive Vicsek model [2,34].

S9.2 Density crossover and multi-estimator comparison

For the density cut, we hold a coherent-flock rule fixed and increase ρ from 0.05 to 0.70 (Supplementary Figure S15, bottom row). Polar order falls from about 0.88 at low density. Its susceptibility has a dominant peak near $\rho^* \approx 0.44$ and a secondary low-density feature. The descriptor rate, however, is largest at the lowest densities, where local structure changes rapidly as initially isolated particles begin to interact. The two detectors are therefore responding to different reorganizations along this cut.

The same susceptibility measure also locates the reciprocal-to-chiral onset in a coupling direction. Main text Section 7.3 gives the coarse estimate $J_a^* \approx 1.33$, while the finer scan in main text Section 9 places the relevant range near 1.4 to 1.5.



Supplementary Figure S15. Transition and crossover diagnostics along parameter cuts. Top: noise-driven loss of polar order. Bottom: density-driven loss of polar order. Columns show the order parameter, its across-seed fluctuation measure with Binder cumulant, and the descriptor change rate. Different peak locations indicate that the measures can respond to different reorganizations.

S10. Sensitivity-ridge search and controls

S10.1 Candidate selection, localization, and tracing

Candidate high-sensitivity regions are found in two stages. We first search the existing archives and trajectory corpus for nearby parameter vectors with a large descriptor change per unit parameter distance, which requires no new simulation. If this does not provide enough candidates, we run a QD search that favors large descriptor sensitivity $b(\mathbf{u})$ while using the learned coordinates as archive axes. Because noisy QD scores can favor lucky evaluations, incumbents are re-evaluated and compared using re-evaluated means. Each candidate is then localized along one parameter direction by a coarse descriptor-rate scan, golden-section refinement, and the same multi-scale seed-noise test used in the main text. A tanh fit to the local descriptor profile is used only to describe the width of the change as abrupt or gradual; it is not used to infer transition order. We trace an accepted ridge with predictor-corrector continuation, periodically recomputing the full Jacobian. Tracing stops when the ridge leaves the parameter box, falls below $b_{\min}$, or closes on itself. Local ridge geometry is evaluated separately by patch sampling rather than by the one-dimensional trace itself.

Each traced ridge is classified with three controls. The reciprocal control sets $J_a = C_a = 0$ while holding composition and the other parameters fixed. The single-species control asks whether the same reorganization occurs in the corresponding nine-parameter single-species model. The species-swap control exchanges the species labels; a physical classification should be unchanged

apart from seed noise. If a ridge disappears under the reciprocal control and is absent in the single-species subspace, it is classified as non-reciprocity-dependent. If it survives both controls, it is a single-species reorganization. If it survives the reciprocal control but disappears in the single-species subspace, it is a reciprocal mixture boundary. Missing or conflicting controls leave the ridge unclassified. The $f_B \to 0$ limit is used only as a consistency check because mixture observables necessarily collapse there. Before interpreting newly found ridges, we also recover the known noise reorganization near $\eta^* \approx 0.5$ and the reciprocal-to-chiral onset near $J_a^* \approx 1.33$, and check symmetry, composition limits, identical-species behavior, and false positives.

The search is intentionally focused on directions that are less well covered by existing work: symmetric versus antisymmetric radial coupling, composition versus non-reciprocity, adhesion versus cross-species coupling, and density versus antisymmetric coupling at high density. The alignment plane is used mainly as a reference because it has already been studied extensively at larger system sizes. When density changes during continuation, the box size is held fixed and particle number is recorded, making explicit that a density change also changes system size in this parameterization.

S10.2 Patch geometry and ridge termination

Patch sampling distinguishes an extended ridge from a bounded patch. Around a ridge point, we draw random directions in the local tangent space, correct each point back to the high-sensitivity set, and compare the principal-component spectrum of the corrected cloud with an isotropic null of the same sample size. This sample-size comparison matters because n points can span at most $n-1$ dimensions. For the single-species ridges, all thirty corrections succeed and the corrected clouds are consistent with generic codimension-one hypersurfaces: about eleven of thirteen components are needed for 95% of the variance, comparable with the matched null. The reciprocal-to-chiral ridge behaves differently. Only about one in five tangent directions can be corrected back to the ridge, so the high-sensitivity set terminates in most directions. When the number of successful points is matched by drawing more directions, its principal-component spectrum is only slightly below the same-size null, nine components versus ten. The robust conclusion is therefore that this ridge is a bounded local patch, not a low-dimensional filament. Earlier low-dimensional estimates were caused by analyzing the one-dimensional trace or too few successful correction points.

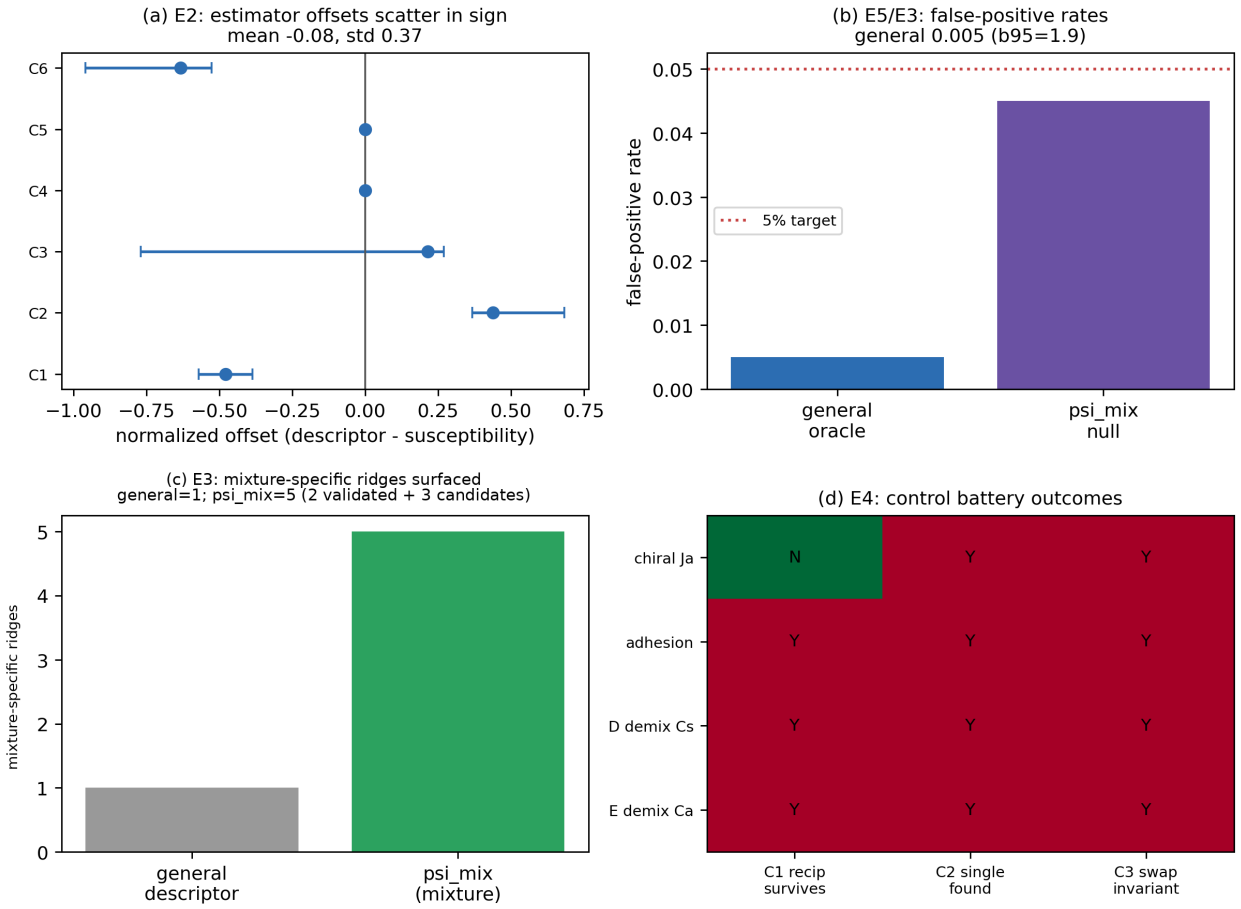
S10.3 Catalog and operating characteristics

Supplementary Table S4 summarizes the final catalog. With the general descriptor, the search finds an adhesion-driven single-species ridge class, two high-density ridges that the controls reclassify as single-species packing reorganizations, and an ambiguous ridge near the reciprocal-to-chiral onset. Restricting the descriptor to mixture-specific quantities adds five ridges. Two satisfy the full criterion for reciprocal mixture boundaries; three disappear when antisymmetric couplings are removed but have not been tested in the single-species subspace, so they remain reciprocity-dependent candidates. The full study required roughly forty thousand simulations, about half the cost of one 64,000-evaluation archive. Most of that cost comes from full-Jacobian sensitivity calculations rather than the cheaper directional updates used during tracing. Run length was at most 440 steps; only a small number of points failed the requirement of at least fifty integrated autocorrelation times, mainly in the dense, low-noise region where relaxation is slowest.

Supplementary Table S4. Sensitivity-ridge catalog and control status. A behavioral-boundary classification is assigned only when the required controls are complete. The reciprocal control tests

dependence on antisymmetric couplings; the single-species control tests whether the reorganization is mixture-specific.

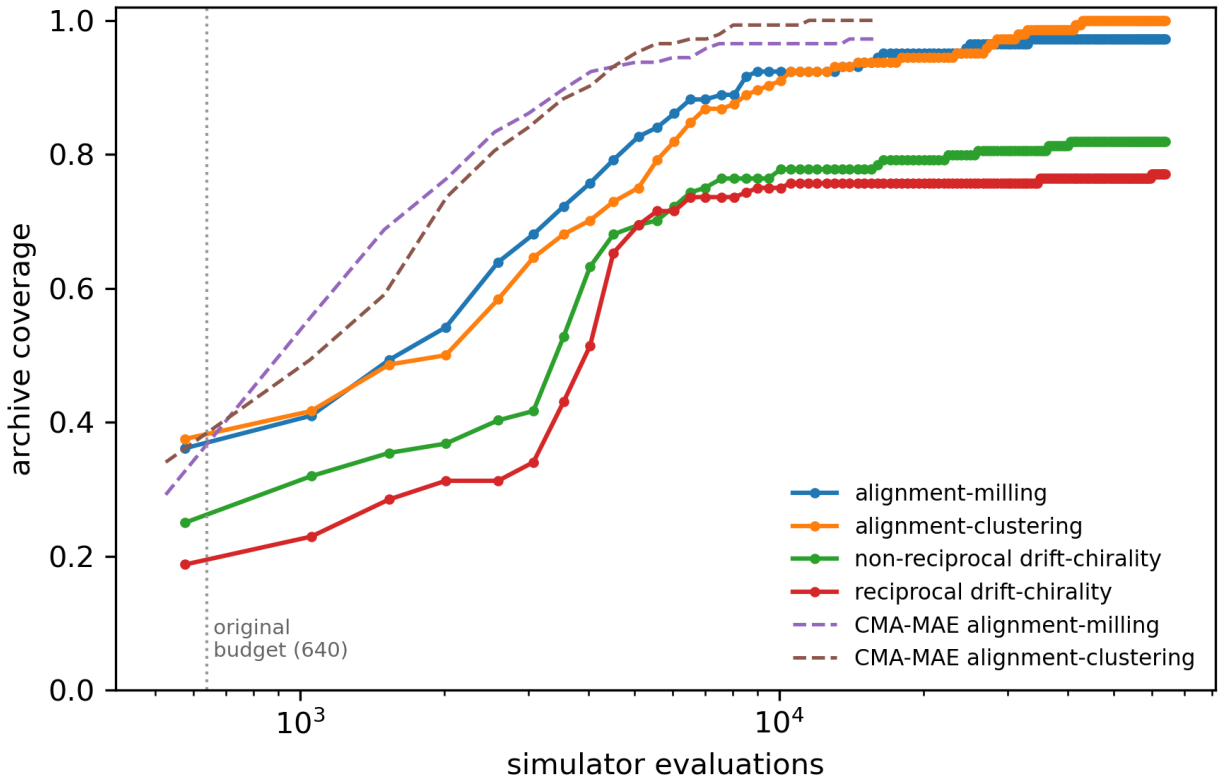
Ridge / candidate boundary	Dominant observable	Reciprocal control	Single-species control	Classification / evidence status
Under the general learned descriptor				
Reciprocal-to-chiral ridge	co-alignment P_+	collapses	reorganization persists	ambiguous reciprocity-dependent ridge; coincident single-species reorganization
Adhesion aggregation ridge	Warren-Cowley α	survives	reorganization persists	single-species structural ridge
High-density ridge, C_s	drift D	survives	reorganization persists	single-species packing ridge
High-density ridge, C_a	co-alignment P_+	survives	reorganization persists	single-species packing ridge
Under the mixture-restricted descriptor				
Reciprocity-dependent demixing ridge 1	demixing Φ	collapses	not run	candidate mixture boundary; incompletely classified
Reciprocity-dependent demixing ridge 2	demixing Φ	collapses	not run	candidate mixture boundary; incompletely classified
Reciprocity-dependent demixing ridge 3	demixing Φ	collapses	not run	candidate mixture boundary; incompletely classified
Composition demixing boundary 1	demixing Φ	survives	collapses	validated reciprocal mixture boundary
Composition demixing boundary 2	demixing Φ	survives	collapses	validated reciprocal mixture boundary



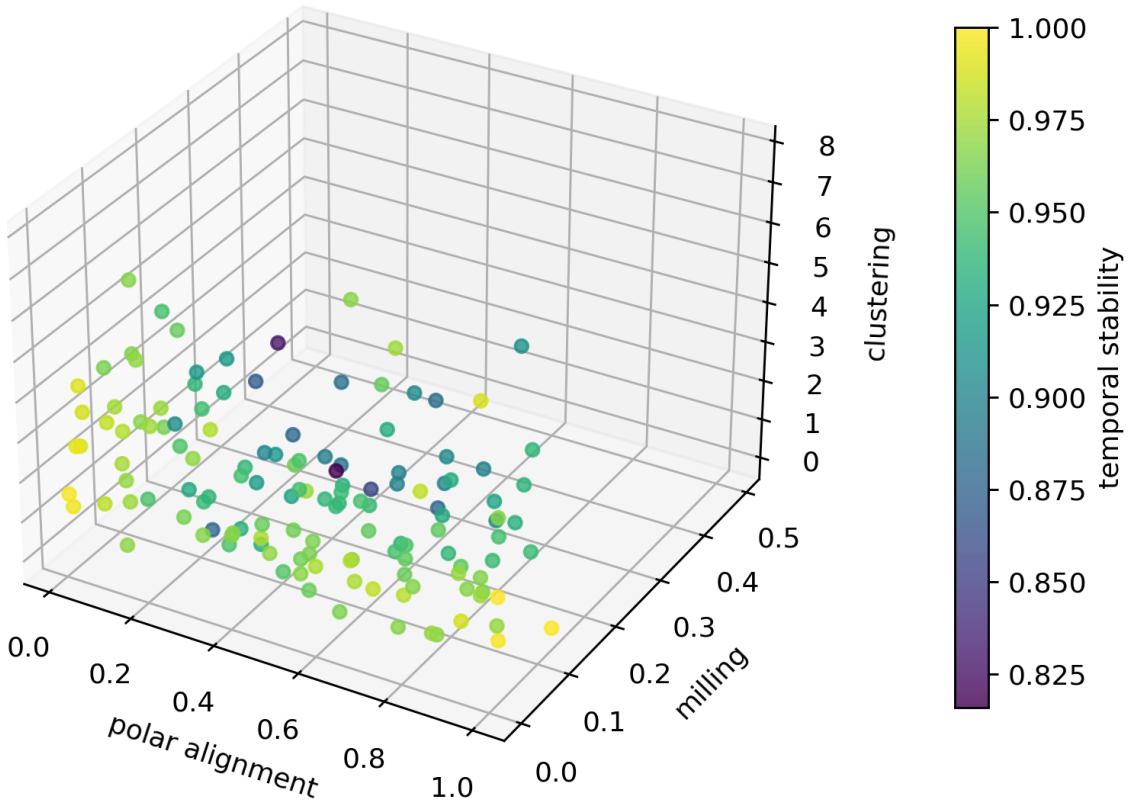
Supplementary Figure S16. Performance and classification of the sensitivity-ridge search. (a) Descriptor-sensitivity and order-parameter peak locations differ by a sign-varying amount across six cuts. (b) False-positive rates for the general sensitivity detector and the mixture-specific detector. (c) The mixture-specific descriptor reveals two fully classified reciprocal-mixture boundaries and three additional reciprocity-dependent candidates. (d) Outcomes of the reciprocal, single-species, and species-swap controls for each ridge type.

S11. Additional archive figures and complete supplementary tables

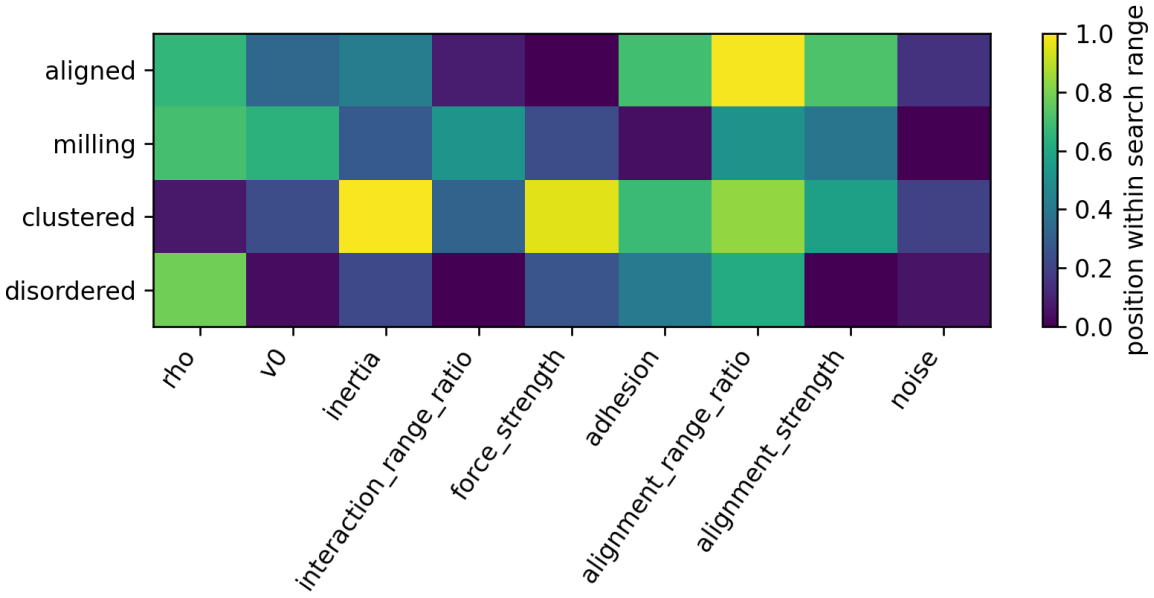
S11.1 Additional archive figures



Supplementary Figure S17. Archive coverage versus simulator evaluations. The single-species archives saturate at 144/144 cells for alignment-clustering and 140/144 for alignment-milling. CMA-MAE reaches the same high coverage with fewer evaluations. The two-species reciprocal and non-reciprocal archives reach 77.1% and 81.9% coverage.



Supplementary Figure S18. Elite single-species rules in alignment, milling, and clustering coordinates.



Supplementary Figure S19. Normalized parameter values of the aligned, milling, clustered, and disordered representatives.

S11.2 Complete tabulations

Supplementary Table S5. Learned single-species region summaries.

Cluster	Count	Mean ρ	Mean φ	Mean m	Mean Γ	Mean ψ_6	Mean density peak	High-density fraction	Reflecting fraction
0	28	0.328	0.097	0.058	1.628	0.177	0.260	0.179	0.000
1	39	0.131	0.192	0.108	0.981	0.069	0.049	0.000	0.000
2	16	0.111	0.184	0.110	2.608	0.162	0.608	0.000	0.063
3	20	0.251	0.871	0.048	0.934	0.069	0.042	0.100	0.000
4	90	0.455	0.095	0.056	0.943	0.049	0.009	0.522	0.000
5	16	0.272	0.111	0.076	2.453	0.094	0.321	0.250	1.000
6	4	0.072	0.139	0.099	2.643	0.171	0.770	0.000	0.000
7	22	0.291	0.079	0.056	2.039	0.164	0.460	0.182	0.000
8	36	0.518	0.068	0.043	1.213	0.204	0.091	0.722	0.000
9	41	0.406	0.110	0.064	1.183	0.110	0.052	0.439	0.976
10	31	0.414	0.498	0.057	1.071	0.066	0.044	0.452	0.161
11	17	0.108	0.489	0.148	1.609	0.087	0.163	0.000	0.588

Supplementary Table S6. Local sensitivity spectra after applying the seed-to-seed noise threshold.

The six nominal representatives correspond to five distinct parameter vectors because the aligned and high-density-aligned cases use the same dense-flock rule. In separate two-seed calculations that rule produced raw spans of 5.53 and 7.26 decades and different leading eigenvectors. We therefore bootstrap the loadings over sixteen seeds per perturbation side and report them only

when the 90% interval excludes zero. "Resolved directions" are the singular directions of J that lie above the seed-to-seed noise floor defined in main text Section 6.2.

Regime	Raw span, 2-seed (dec.)	Resolved dirs	Leading resolved eigenvector (loading [90% interval])
Aligned (dense flock)	5.53 / 7.26	2	not determined (only noise -0.37 [-0.58,-0.12], adhesion 0.39 [0.00,0.75])
Milling	4.76	2	density 0.65 [0.51,0.77], force strength 0.56 [0.36,0.72], adhesion 0.32 [0.15,0.49]
Clustered	3.97	3	interaction range 0.44 [0.19,0.68], adhesion -0.41 [-0.66,-0.17], noise -0.33 [-0.51,-0.15]
Density wave	3.26	4	density -0.97 [-0.99,-0.94]
High-density aggregate	3.28	3	adhesion -0.88 [-0.90,-0.87]

Supplementary Table S7. Learned two-species region summaries.

Cluster	Count	Reciprocal fraction	Mean chirality	Mean drift	Mean segregation	Mean φ	Mean Γ
0	25	0.360	0.212	0.432	0.029	0.641	0.992
1	27	0.407	0.041	0.052	0.004	0.133	1.025
2	2	0.500	0.006	0.002	-0.014	0.012	2.254
3	9	0.222	0.032	0.017	-0.029	0.095	1.829
4	5	1.000	0.002	0.283	0.023	0.288	1.037
5	8	0.500	0.030	0.023	0.000	0.088	2.149
6	7	0.714	0.042	0.039	-0.055	0.161	1.146
7	8	0.250	0.056	0.278	-0.042	0.329	0.959
8	17	0.529	0.041	0.819	0.031	0.855	1.075
9	18	0.667	0.039	0.238	0.053	0.305	1.250

Supplementary Table S8. Directed-search champions and proximity to parameter bounds.

For each champion, the table reports how many of the nine single-species or fourteen two-species parameters lie within 1% of a search-box bound and identifies those parameters. The ceiling solutions ϕ and $D = 1.000$ occur at $\eta = 0$.

Objective	Median fitness	Noise η	Params at bound	Parameters on a bound
Aligned flock	1.000	0.000	3/9	force strength, noise (lo); ρ (hi)
Milling vortex	0.397	0.027	1/9	interaction range (hi)
Dense cluster	0.782	0.214	1/9	ρ (lo)
Nematic lanes	0.346	0.087	4/9	v_0, alignment strength (lo); inertia, adhesion (hi)

Objective	Median fitness	Noise η	Params at bound	Parameters on a bound
Hexatic crystal	0.844	0.197	1/9	ρ (hi)
Density waves	0.853	0.126	3/9	ρ (lo); inertia, adhesion (hi)
Phase separation	1.000	0.054	1/9	adhesion (hi)
Active turbulence	0.990	0.000	3/9	noise (lo); interaction range, adhesion (hi)
Travelling mixture	1.000	0.000	4/14	inertia, force strength, noise (lo); alignment range (hi)
Chiral rotation	0.196	0.000	4/14	ρ, noise, J_a (lo); v_0 (hi)
Species segregation	0.868	0.147	0/14	none

Supplementary Table S9. Per-seed champion fitness distributions over twenty validation seeds.

The table reports the median, interquartile range, minimum, maximum, and fraction of seeds below half the best observed value. Milling, the nematic objective, the phase-separation objective, and chiral rotation have large low-tail fractions and are therefore poorly summarized by a mean and standard deviation.

Objective	Median	IQR	Min	Max	Frac. seeds < ½ max
Aligned flock	1.000	1.000-1.000	0.999	1.000	0.00
Milling vortex	0.397	0.235-0.471	0.097	0.515	0.30
Dense cluster	0.782	0.606-0.800	0.461	0.806	0.00
Nematic lanes	0.346	0.300-0.466	0.168	0.630	0.35
Hexatic crystal	0.844	0.828-0.859	0.583	0.878	0.00
Density waves	0.853	0.771-0.855	0.427	0.857	0.05
Phase separation	1.000	0.762-1.000	0.000	1.000	0.15
Active turbulence	0.990	0.990-0.991	0.950	0.992	0.00
Travelling mixture	1.000	1.000-1.000	1.000	1.000	0.00
Chiral rotation	0.196	0.069-0.825	0.009	0.954	0.60
Species segregation	0.868	0.817-0.899	0.505	0.926	0.00

Supplementary Table S10. Independent representation-validation results.

Single-species results use the independent 72-rule post-selection corpus and the fixed training and selection procedure used for the main representation. Two-species results use the untouched 26-rule test partition from the separate source-stratified 60/20/20 audit; those audit coordinates are not used elsewhere in the analysis. All neighbor-overlap values use ten nearest neighbors.

Audit	Model / latent dimension	Std. MSE	Trustworthiness	10-NN overlap	Additional result
Single independent	Autoencoder, 2D	0.250	0.890	0.556	Seed stability 0.795 [0.661, 0.945]
Single independent	Descriptor PCA, 2D	0.402	0.829	0.458	57.2% train variance explained
Single independent	Classical 8-vector PCA, 2D	-	0.765	0.374	Scored against 51-D descriptor
Single independent	Autoencoder, 3D	0.188	0.901	0.604	Validation-selected seed 1
Single independent	Autoencoder, 4D	0.133	0.942	0.653	Validation-selected seed 0
Single independent	Autoencoder, 6D	0.111	0.932	0.625	Validation-selected seed 0
Two-species clean split	Autoencoder, 2D	0.407	0.845	0.627	Seed stability 0.918 [0.877, 0.957]
Two-species clean split	Autoencoder, 6D	0.258	0.847	0.723	Validation-selected seed 1

For the independent recoverability test, predictors are fit on the original 360 single-species rules and evaluated once on the 72-rule test set. Test R^2 for learned axis 1 is 0.503 for linear regression, 0.626 for random forest, 0.665 for Extra-Trees, and 0.644 for histogram gradient boosting. For learned axis 2 the corresponding values are 0.911, 0.914, 0.935, and 0.935.

Supplementary References

1. Vicsek T, Czirók A, Ben-Jacob E, Cohen I, Shochet O. Novel type of phase transition in a system of self-driven particles. Physical Review Letters. 1995;75:1226-1229. doi:10.1103/PhysRevLett.75.1226.

2. Grégoire G, Chaté H. Onset of collective and cohesive motion. Physical Review Letters. 2004;92:025702. doi:10.1103/PhysRevLett.92.025702.

3. D'Orsogna MR, Chuang YL, Bertozzi AL, Chayes LS. Self-propelled particles with soft-core interactions: patterns, stability, and collapse. Physical Review Letters. 2006;96:104302. doi:10.1103/PhysRevLett.96.104302.

4. Vicsek T, Zafeiris A. Collective motion. Physics Reports. 2012;517:71-140. doi:10.1016/j.physrep.2012.03.004.

5. Marchetti MC, Joanny JF, Ramaswamy S, Liverpool TB, Prost J, Rao M, Simha RA. Hydrodynamics of soft active matter. Reviews of Modern Physics. 2013;85:1143-1189. doi:10.1103/RevModPhys.85.1143.

6. Toner J, Tu Y. Long-range order in a two-dimensional dynamical XY model: how birds fly together. Physical Review Letters. 1995;75:4326-4329. doi:10.1103/PhysRevLett.75.4326.

7. Couzin ID, Krause J, James R, Ruxton GD, Franks NR. Collective memory and spatial sorting in animal groups. Journal of Theoretical Biology. 2002;218:1-11. doi:10.1006/jtbi.2002.3065.

8. Gompper G, et al. The 2025 motile active matter roadmap. Journal of Physics: Condensed Matter. 2025;37:143501. doi:10.1088/1361-648X/adac98.

9. Mouret JB, Clune J. Illuminating search spaces by mapping elites. arXiv:1504.04909. 2015.

10. Cully A, Clune J, Tarapore D, Mouret JB. Robots that can adapt like animals. Nature. 2015;521:503-507. doi:10.1038/nature14422.

11. Pugh JK, Soros LB, Stanley KO. Quality diversity: a new frontier for evolutionary computation. Frontiers in Robotics and AI. 2016;3:40. doi:10.3389/frobt.2016.00040.

12. Fontaine MC, Togelius J, Nikolaidis S, Hoover AK. Covariance matrix adaptation for the rapid illumination of behavior space. In: Proceedings of the Genetic and Evolutionary Computation Conference (GECCO). 2020:94-102. doi:10.1145/3377930.3390232.

13. Fontaine MC, Nikolaidis S. Covariance Matrix Adaptation MAP-Annealing. In: Proceedings of the Genetic and Evolutionary Computation Conference (GECCO). 2023:456-465. doi:10.1145/3583131.3590389.

14. Grillotti L, Cully A. Unsupervised behavior discovery with quality-diversity optimization. IEEE Transactions on Evolutionary Computation. 2022;26(6):1539-1552. doi:10.1109/TEVC.2022.3159855.

15. Hedayatian S, Nikolaidis S. AutoQD: Automatic discovery of diverse behaviors with quality-diversity optimization. International Conference on Learning Representations (ICLR). 2026. arXiv:2506.05634.

16. Carrasquilla J, Melko RG. Machine learning phases of matter. Nature Physics. 2017;13:431-434. doi:10.1038/nphys4035.

17. Wetzel SJ. Unsupervised learning of phase transitions: from principal component analysis to variational autoencoders. Physical Review E. 2017;96:022140. doi:10.1103/PhysRevE.96.022140.

18. Dulaney AR, Brady JF. Machine learning for phase behavior in active matter systems. Soft Matter. 2021;17:6808-6816. doi:10.1039/D1SM00266J.

19. Transtrum MK, Machta BB, Sethna JP. Geometry of nonlinear least squares with applications to sloppy models and optimization. Physical Review E. 2011;83:036701. doi:10.1103/PhysRevE.83.036701.

20. Quinn KN, Abbott MC, Transtrum MK, Machta BB, Sethna JP. Information geometry for multiparameter models: new perspectives on the origin of simplicity. Reports on Progress in Physics. 2023;86:035901. doi:10.1088/1361-6633/aca6f8.

21. Gutenkunst RN, Waterfall JJ, Casey FP, Brown KS, Myers CR, Sethna JP. Universally sloppy parameter sensitivities in systems biology models. PLoS Computational Biology. 2007;3:e189. doi:10.1371/journal.pcbi.0030189.

22. Transtrum MK, Machta BB, Brown KS, Daniels BC, Myers CR, Sethna JP. Perspective: sloppiness and emergent theories in physics, biology, and beyond. Journal of Chemical Physics. 2015;143:010901. doi:10.1063/1.4923066.

23. Ivlev AV, Bartnick J, Heinen M, Du CR, Nosenko V, Löwen H. Statistical mechanics where Newton's third law is broken. Physical Review X. 2015;5:011035. doi:10.1103/PhysRevX.5.011035.

24. You Z, Baskaran A, Marchetti MC. Nonreciprocity as a generic route to traveling states. Proceedings of the National Academy of Sciences USA. 2020;117:19767-19772. doi:10.1073/pnas.2010318117.

25. Fruchart M, Hanai R, Littlewood PB, Vitelli V. Non-reciprocal phase transitions. Nature. 2021;592:363-369. doi:10.1038/s41586-021-03375-9.

26. Dinelli A, O'Byrne J, Curatolo A, Zhao Y, Sollich P, Tailleur J. Non-reciprocity across scales in active mixtures. Nature Communications. 2023;14:7035. doi:10.1038/s41467-023-42713-5.

27. Kreienkamp KL, Klapp SHL. Nonreciprocal alignment induces asymmetric clustering in active mixtures. Physical Review Letters. 2024;133:258303. doi:10.1103/PhysRevLett.133.258303.

28. Chen J, Lei X, Xiang Y, Duan M, Peng X, Zhang HP. Emergent chirality and hyperuniformity in an active mixture with nonreciprocal interactions. Physical Review Letters. 2024;132:118301. doi:10.1103/PhysRevLett.132.118301.

29. Chatterjee S, Mangeat M, Woo CU, Rieger H, Noh JD. Flocking of two unfriendly species: the two-species Vicsek model. Physical Review E. 2023;107:024607. doi:10.1103/PhysRevE.107.024607.

30. Dutta AK, Mangeat M, Rieger H, Paul R, Chatterjee S. Stability of flocking in the reciprocal two-species Vicsek model: effects of relative population, motility, and noise. Physical Review E. 2025;112:024137. doi:10.1103/PhysRevE.112.024137.

31. Digregorio P, Levis D, Suma A, Cugliandolo LF, Gonnella G, Pagonabarraga I. Full phase diagram of active Brownian disks: from melting to motility-induced phase separation. Physical Review Letters. 2018;121:098003. doi:10.1103/PhysRevLett.121.098003.

32. Storn R, Price K. Differential evolution - a simple and efficient heuristic for global optimization over continuous spaces. Journal of Global Optimization. 1997;11:341-359. doi:10.1023/A:1008202821328.

33. Dai C, Glotzer SC. Efficient phase diagram sampling by active learning. Journal of Physical Chemistry B. 2020;124:1275-1284. doi:10.1021/acs.jpcb.9b09202.

34. Chaté H, Ginelli F, Grégoire G, Raynaud F. Collective motion of self-propelled particles interacting without cohesion. Physical Review E. 2008;77:046113. doi:10.1103/PhysRevE.77.046113.

35. Nayak S, Bag P, Ghosh PK, Li Y, Zhou Y, Yin Q, Marchesoni F, Nori F. Diffusion transients in motility-induced phase separation. Physical Review Research. 2025;7:013153. doi:10.1103/PhysRevResearch.7.013153.

36. Zhou Y, Yin Q, Nayak S, Bag P, Ghosh PK, Li Y, Marchesoni F. Visual quorum sensing in chiral suspensions: Hyperuniformity and edge currents. PNAS Nexus. 2025;4(12):pgaf373. doi:10.1093/pnasnexus/pgaf373.