Fear-Driven Collective Topology: Comparing Smart-Boid Vietoris-Rips Graphs to Animal Communication Networks via Persistent Features

Guilherme S. Y. Giardini
Northern Arizona University
gsg95@nau.edu

Carlo R. da Cunha
Northern Arizona University
Carlo.Cunha@nau.edu

Abstract

We test whether adaptive agents, a.k.a. "Smart-Boids" governed by neural networks under evolutionary pressure, can generate topologies resembling animal communication networks. Using Vietoris-Rips filtrations and persistent homology, we compare $1000+$ empirical networks to simulations via feature-based correlations. Minimal ingredients (fear of isolation, limited perception, inertia, exclusion, noise) reproduce both sparse and small-world topologies observed in diverse animal systems. Results suggest that ecological constraints, rather than complex cognition, drive the emergence of communication networks.

1 Introduction

Animal communication networks reveal how collective coordination emerges across species. Classic self-organization models such as Reynolds' Boids Reynolds [1987] and the Vicsek model Vicsek et al. [1995] demonstrated that simple attraction–alignment–repulsion rules can generate coherent group motion. Later extensions incorporated signaling, information suppression, or evolutionary adaptation Witkowski and Ikegami [2016], Mitri et al. [2011, 2009], yet few works have systematically related such simulations to real animal networks. Meanwhile, empirical tracking now yields detailed interaction graphs for ants, primates, birds, dolphins, rodents, and other taxa Mersch et al. [2013], McGregor [2005], Reichert et al. [2024], but we still lack a unified framework connecting these biological networks to generative swarm models. Here we extract persistent topological features from Smart-Boid communication graphs and compare them directly to 1,087 empirical networks, producing the first large-scale atlas of topological similarity between adaptive swarms and real animal communication systems.

2 Methodology

2.1 Empirical Datasets

We analyzed 1,087 animal communication and interaction networks from the Animal Social Network Repository (ASNR v3.0)Collier et al. [2024]. Each dataset represents nodes as individual animals and edges as observed interactions. They span proximity networks, co-occurrence graphs, and behavioral interaction networks across taxa; for all datasets we used the curated adjacency matrices exactly as released by the ASNR team. Typical network sizes range from N = 10–200 individuals. All empirical networks were preassembled from raw field observations by the curators of the Animal Social Network Repository (ASNR v3.0). We used these ASNR-provided adjacency matrices

Preprint.

exactly as released, converting them to weighted undirected graphs with igraph 0.10.12 (R/C++), preserving all original edge definitions and weights.

2.2 Smart-Boid Model and Simulation Parameters

Smart-Boids are adaptive agents whose motion in a 2D space is governed by a 1-hidden-layer neural controller evolved under fear-driven selective pressure, wherein agents that, over a 300-step relaxation window, remain farther from their nearest neighbors than the population mean are eliminated, while the closest 20% form the reproducing elite. Offspring are generated via uniform crossover on neural-network parameters, each gene inherited randomly from either parent with added mutation noise, producing controllers biased toward maintaining cohesion and avoiding isolation as described in Giardini et al. [2024]. At each timestep, agent i identifies visible neighbors within its field of view FoV (angular cone), selects its k nearest neighbors (NN), and receives their relative positions as input. The controller outputs a turning increment $\Delta\theta_i = \theta_i(t+1) - \theta_i(t)$ which is applied subject to a maximum rotational constraint ω_{max} (max turning angle), with θ_i determining the direction of motion. Motion additionally includes stochastic perturbation ξ (communication noise) and hard volume-exclusion for realism.

In the simulation we systematically vary five parameters: while keeping the number of agents fixed

Parameter	Symbol	Values Explored	Interpretation
Communication noise	ξ	0-1.0 (log-spaced)	variability in decisions
Field of view	FoV	$\pi/2,\ \pi,\ 3\pi/2,\ 2\pi$	sensory coverage
Max turning angle	ω	$\pi/2,\ \pi,\ 3\pi/2,\ 2\pi$	maneuverability
Nearest neighbors	NN	1-9	interaction range
Duration	T	10^5 timesteps	simulation length

at $N = 100$ agents. This produced 3,200 simulations, one for each parameter combination. For each simulation we retained the final 1,000 timesteps to compute steady-state structure. At every timestep we computed pairwise distances and extracted all agent positions for topological analysis.

2.3 Vietoris–Rips Graph Construction

To compare simulations and empirical networks under a unified framework, we constructed Vietoris–Rips (VR) proximity graphs from agent positions. For a radius parameter $\varepsilon \geq 0$, an edge is added between agents i and j if $\parallel x_i - x_j \parallel \leq \varepsilon$. We sweep ε across a set of radii $\varepsilon \in 0.4, 0.5, 0.5, \ldots, 5.0$ in increments of 0.1. We perform the full sweep to capture persistent topological structure. However, when comparing simulations to real animal networks we restrict the analysis to radii smaller than the distance to the farthest nearest neighbor $\varepsilon \leq r_{NN}$ since agents would not communicate directly beyond this range. Thus the final VR graphs encode only directly feasible communication links between agents.

2.4 Graph and Topological Feature Extraction

All evolutionary simulations were implemented in Julia, which output Smart-Boid position archives over the last 10^3 steps. The archives, processed in Python, had their extracted positions, with which we constructed Vietoris–Rips graphs over the radius sweep and computed all structural and topological features. Our analysis pipeline combined custom Python code with NumPy, pandas, and NetworkX for graph metrics; python-louvain for modularity; and matplotlib/tqdm for visualization and parallelization. Persistent lifetimes were obtained using our own birth–death detector applied directly to the radius-indexed feature curves. These Python tools were used only for graph and TDA analysis; all agent dynamics and evolutionary updates ran exclusively in Julia. For every VR graph we computed a comprehensive set of 64 graph-theoretic and persistence-derived features, including:

- Degree structure: avg.degree.norm, std.degree.node.norm, CV.degree
- Clustering and motifs: triangles.norm, triangles.density, transitivity
- Weighted connectivity: avg.node.strength.norm, std.node.strength.norm
- Global structure: mean.path.length.norm, network.density, modularity
- Topological invariants: normalized Betti-1 lifetime H1.norm

2.5 Persistence-Guided Feature Selection (80% Rule)

Many graph features vary substantially with ε or with sampling noise. To filter out unstable descriptors, we introduced a persistence-guided stability criterion:

A feature f is retained if it is present and stable in at least 80% of replicate simulations. We selected the 80% cutoff at the stability-curve elbow: higher thresholds leave too few descriptors, while lower ones admit noise-driven features. This cutoff provides a balance between persistency and discriminability, yielding a stable and informative subset of features for all subsequent comparisons.

2.6 Similarity Computation

For each empirical and simulated network, we created normalized feature vectors (v^{real} and v^{sim} respectively) of the persistent features of the interaction graphs and calculated their Pearson similarity:

$$r(v^{\text{real}}, v^{\text{sim}}) = \frac{\sum_i (v_i^{\text{real}} - \bar{v}^{\text{real}})(v_i^{\text{sim}} - \bar{v}^{\text{sim}})}{\sum_i (v_i^{\text{real}} - \bar{v}^{\text{real}})^2 \sum_i (v_i^{\text{sim}} - \bar{v}^{\text{sim}})^2}$$

We computed a full pairwise comparison between all simulated and all empirical networks. Because feature magnitudes vary across species and simulations, we used a rank-based similarity that compares the relative prominence of features rather than their scale, something common for high dimensional spaces. For each empirical network, the simulation with the highest rank-similarity is reported as its best match.

2.7 Anomaly-Based Cutoff for Correlation Ranking

To identify where similarity rankings become noise-dominated, we compute the rolling z-score of local slopes in the ordered correlation curve and apply a Shewhart-style threshold ($\| z \| > 1$). The first detected anomaly, appearing as an elbow in the correlation-vs-ranking data, defines the cutoff index. Similarities below this elbow were discarded for two reasons: 1. lower-rank matches are dominated by noise. 2. absolute magnitudes are unreliable in high-dimensional feature spaces. Thus only the highest-ranked graphs retained biological meaningfulness.

3 Results

3.1 Collective regimes from Smart-Boid dynamics

Across roughly 3,000 evolutionary simulations, Smart-Boids produced a diverse repertoire of collective states. K-means clustering of structural and dynamical descriptors revealed robust pattern classes: bands, flocks, swarms, lane formation, and symmetry-broken variants (horizontal vs vertical lanes) Fig. (1b). Classification was performed directly on the spatial and topological descriptors of the final configurations, and k-means consistently returned stable, well-separated formation groups across random seeds. Fig. (1a) shows the resulting cluster structure and representative examples of each learned formation type.

3.2 Correlations to real animal graphs

Using the similarity procedure described in Section 2, we generated a ranked list of simulation–empirical matches for every dataset. Table 1 reports the highest-confidence matches, i.e., those that remain above the stability cutoff identified during ranking, together with the simulation parameters most consistently associated with each species' communication topology.

3.3 Qualitative Discussion

To identify which simulation parameters most strongly explain cross-species differences, we computed a correlation matrix over all parameters appearing in the high-confidence matches of Table 1. Each entry reflects how consistently a parameter co-varies with others across the best simulation–empirical pairings. This analysis isolates the parameters that most reliably shape the matched communication topologies, revealing a small set of dominant behavioral and sensory drivers.

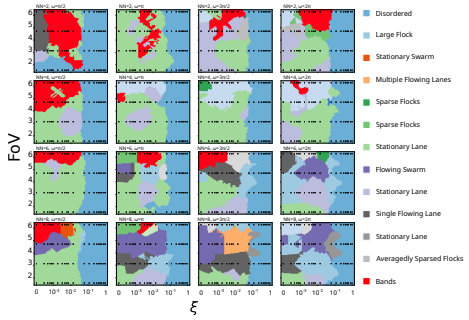

(a) K-means classification into 24 clusters using structural and dynamical features. Clusters with fewer than three simulations were removed.

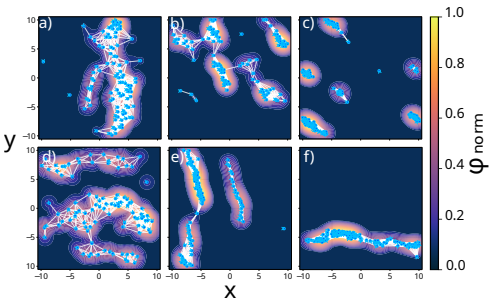

(b) Examples of observed patterns with NN=6. Lines connect agents inside the 6th nearest-neighbor distance: (a) bands, (b) dense flocks, (c) sparse flocks, (d) swarms, (e) multiple lanes, (f) single lane.

Figure 1: Phase-space classification and representative simulation outcomes.

Network ID	**noise_out**	**FoV**	ω	**NN**	**D**	**Pearson corr.**
Vole (Network_1087)	0.0004	4.71	3.14	5	1.6	0.9964
Ant (Network_244)	0.001	3.14	3.14	8	2.6	0.9959
Elephant (Network_260)	0.01	1.57	4.71	9	1.1	0.9949
Baboon (Network_263)	0.1	4.71	1.57	9	1.1	0.9977
Dolphin (Network_339)	0.1	3.14	3.14	6	1.1	0.9938
Raccoon (Network_405)	0.001	3.14	3.14	8	2.5	0.9953
Tortoise (Network_453)	0.2	4.71	6.28	7	0.8	0.9918
Bat (Network_505)	0.001	3.14	3.14	8	2.5	0.9949
Weaverbird (Network_636)	1.0	6.28	1.57	7	1.6	0.9925
Cricket (Network_666)	0.0001	6.28	1.57	7	1.1	0.9969
Weevil (Network_669)	0.2	3.14	6.28	7	0.4	0.9972
Macaque (Network_686)	0.003	4.71	6.28	4	0.4	0.9955
Mouse (Network_727)	0.003	3.14	6.28	2	1.1	0.9986

Table 1: Representative simulation-to-animal matches with parameters, Vietoris-Rips distances, and Pearson correlations.

Analysis of the top-ranked simulation–animal pairs shows that communication noise (ξ) and the maximum turning angle (ω_{max}) are the primary axes of variation structuring collective topology across species. These parameters determine how quickly individuals adjust headings, how much uncertainty they tolerate, and how efficiently local information spreads through the group.

A strong negative correlation between FoV and ω_{max} ($\rho \approx -0.41$) indicates a trade-off between spatial awareness and maneuverability. Species with broad visual fields (birds, bats, dolphins) display smoother, anticipatory trajectories, whereas species with narrower Fields-of-View (primates, rodents) compensate with sharper turns, consistent with sensory ecology.

A second trade-off between ω_{max} and nearest neighbors (NN) ($\rho \approx -0.49$) reflects a balance between agility and social cohesion. Small, fast species (crickets, weevils, mice) maintain small, sparse neighbor sets, while slower-moving species (ants, elephants, dolphins) support larger, denser groups-consistent with established findings that maneuverability scales inversely with group density.

Using a linearity analysis based on the coefficient of determination R^2 (i.e., fitting the neural-network–driven turning responses to a linear model), we quantified the extent to which each species' collective decision dynamics follow a simple reflex-like rule (linear thinking). This showed that ω_{max} (turning agility, $\rho = 0.51$) and NN (social neighborhood size, $\rho = -0.83$) are the strongest predictors of decision linearity: species capable of sharp, agile maneuvers tend to rely on fast, linear reflex responses (high R^2), while species forming larger, cohesive groups integrate multiple cues, producing more nonlinear (low R^2) and context-dependent decision making.

Noise levels further separate taxa: insects, birds, and bats show very low noise, enabling precise coordination, while larger mammals tolerate moderate noise, reflecting slower dynamics and long-range acoustic or visual signaling. Terrestrial species (macaques, voles, baboons) tend to exhibit high $\omega_{\max}$ for obstacle avoidance, while aerial/aquatic species depend on broad FoV and smooth trajectories-consistent with flocking/schooling literature.

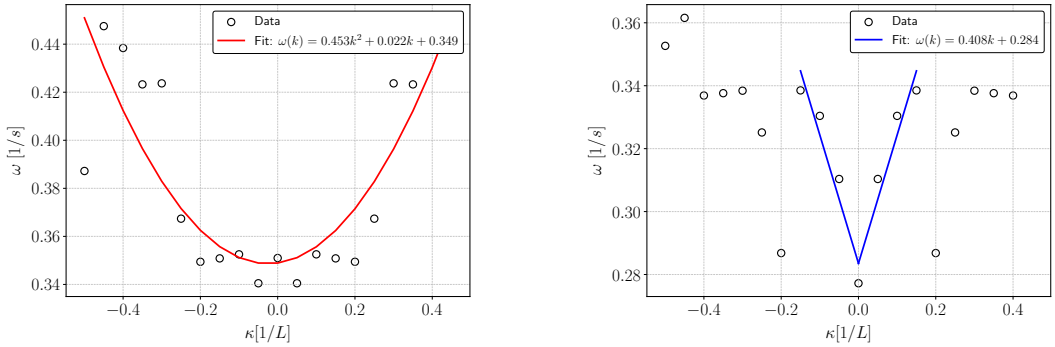

(a) Ordered regime

(b) Disordered regime

Figure 2: Dispersion relations of an ordered regime vs a disordered regime.

4 Conclusion

Fear-driven evolution alone was sufficient to generate a wide repertoire of collective movement patterns, and the addition of simple sensory constraints (e.g., field of view, turning limits) further diversified this repertoire by inducing symmetry breaking between otherwise similar groups. By applying Vietoris–Rips filtrations, we embedded both simulated and empirical interaction networks into a common persistent-feature manifold, enabling a principled, topology-preserving comparison of communication structures across taxa and across modeling paradigms.

Focusing on the highest-confidence simulation–animal matches revealed clear, reproducible relationships between dynamical parameters and biological attributes such as body size, motility, visual aperture, and turning agility, consistent with established theory[1] on the sensory and ecological determinants of group structure. Dispersion analyses further showed that the emergent formations are dynamically stable: fear and local interaction rules readily produce distinct attractors, but transitions between regimes require external drivers such as predation, obstacle pressures, resource availability, or higher-order cognition.

Finally, several matched species exhibit latent modular or hierarchical social layers that are captured by the persistent features but not necessarily by raw interaction data alone, layers such as dominance structure in elephants, grooming-linked food-sharing in bats, or sparse yet functionally cohesive nest-sharing in weavers. Together, these results establish persistent-feature analysis as a scalable, domain-agnostic tool for linking generative swarm models to real animal communication networks, and show that minimal ecological constraints can both generate and help decode the topological signatures underlying collective behavior.

References

Melissa Collier, Nitara Wijayatilake, Pratha Sah, Sania Ali, Jose Mendez, Matthew Silk, and Shweta Bansal. Animal Social Network Repository 3.0. Zenodo, November 2024.

Guilherme S. Y. Giardini, John F. Hardy II, and Carlo R. da Cunha. Evolving Neural Networks Reveal Emergent Collective Behavior from Minimal Agent Interactions, October 2024.

Peter K McGregor. *Animal Communication Networks*. Cambridge University Press, 2005.

[1] see Ward and Webster [2016]

Danielle P Mersch, Allegra Crespi, and Laurent Keller. Tracking individuals shows spatial fidelity is a key regulator of ant social organization. *Science*, 340(6136):1090–1093, 2013.

Sara Mitri, Dario Floreano, and Laurent Keller. The evolution of information suppression in communicating robots with conflicting interests. *Proceedings of the National Academy of Sciences*, 106(37):15786–15790, 2009.

Sara Mitri, Dario Floreano, and Laurent Keller. Relatedness influences signal reliability in evolving robots. *Proceedings of the Royal Society B*, 278(1704):378–383, 2011.

Michael S Reichert, Barney Luttbeg, and Elizabeth A Hobson. Collective signalling is shaped by feedbacks between signaller variation, receiver perception and acoustic environment in a simulated communication network. *Philosophical Transactions of the Royal Society B*, 379(1905): 20230186, 2024.

Craig W Reynolds. Flocks, herds and schools: A distributed behavioral model. *ACM SIGGRAPH Computer Graphics*, 21(4):25–34, 1987.

Tamás Vicsek, András Czirók, Eshel Ben-Jacob, Inon Cohen, and Ofer Shochet. Novel type of phase transition in a system of self-driven particles. *Physical Review Letters*, 75(6):1226, 1995.

Ashley Ward and Mike Webster. *Sociality: The Behaviour of Group-Living Animals*. Springer International Publishing, Cham, 2016. ISBN 978-3-319-28583-2 978-3-319-28585-6. doi: 10.1007/978-3-319-28585-6.

Olaf Witkowski and Takashi Ikegami. Emergence of swarming behavior: Foraging agents evolve collective motion based on signaling. *PLOS ONE*, 11(4):e0152756, 2016.

A Persistent Feature Stability and VR Filtration

A.1 Birth–death structure of simulated VR filtrations

Figure 3 shows the birth-death structure of a representative simulation as the Vietoris-Rips radius is increased from 0.4 to 5.0 in increments of 0.1.

Features close to the diagonal correspond to short-lived fluctuations, noise or transient neighborhood reconfigurations, while long-lived features away from the diagonal represent stable structural elements that repeatedly appear across simulations.

We observe the typical two-regime shape: a dense cloud of short-persistence features followed by a sparse band of long-persistence features. The transition between these regimes is an **empirical stability elbow**, which we use to identify features robust enough to generalize across both artificial and empirical animal networks.

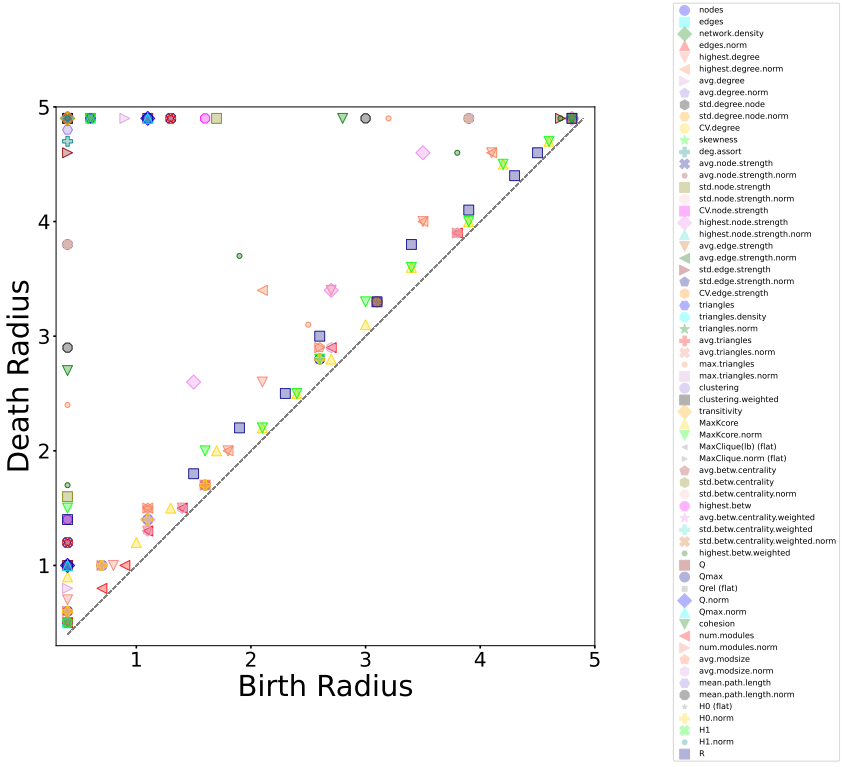

Figure 3: Example persistence diagram showing birth and death radii of β_0 and β_1 features for one simulated final frame.

A.2 Why the 80% stability cutoff?

When counting how many simulations exhibit a persistent signal for each feature, the stability curve shows a clear elbow: features that persist in $< 80\%$ of simulations sharply decay in frequency and are dominated by noise-driven fluctuations. Higher thresholds (e.g., 90–95%) leave only 3–5 features, which is insufficient for discriminating between species.

Thus, 80% provides the optimal trade-off between:

- stability across simulations, and
- maintaining a sufficiently expressive feature space for cross-species comparison.

B Rank-Based Anomaly Detection and High-Dimensional Filtering

B.1 Why correlation magnitude is unreliable in high-dimensional spaces

Similarity between empirical and simulated networks is computed in a high-dimensional feature space (64 graph-theoretic and topological descriptors, counted before the persistency analysis). In such spaces, pairwise correlations are known to *concentrate*: unrelated vectors often achieve deceptively high values due to the curse of dimensionality. Thus, absolute correlation magnitudes are not reliable discriminants of genuine structural similarity.

What *is* reliable is the **ordering** of correlations. Meaningful matches typically occupy a small region at the head of the ranking, after which the correlation curve drops sharply once the comparisons enter the noise regime.

B.2 Rank-ordered correlation curves

For each empirical animal network we:

1. compute Pearson correlations against all $\sim$3000 simulated persistent-feature vectors,
2. sort the correlations in descending order,
3. analyze the sorted curve for structural *change points*.

Figure 4 illustrates a typical case: a smooth high-correlation plateau, followed by a discontinuity ("first elbow") where spurious matches begin dominating the rank order.

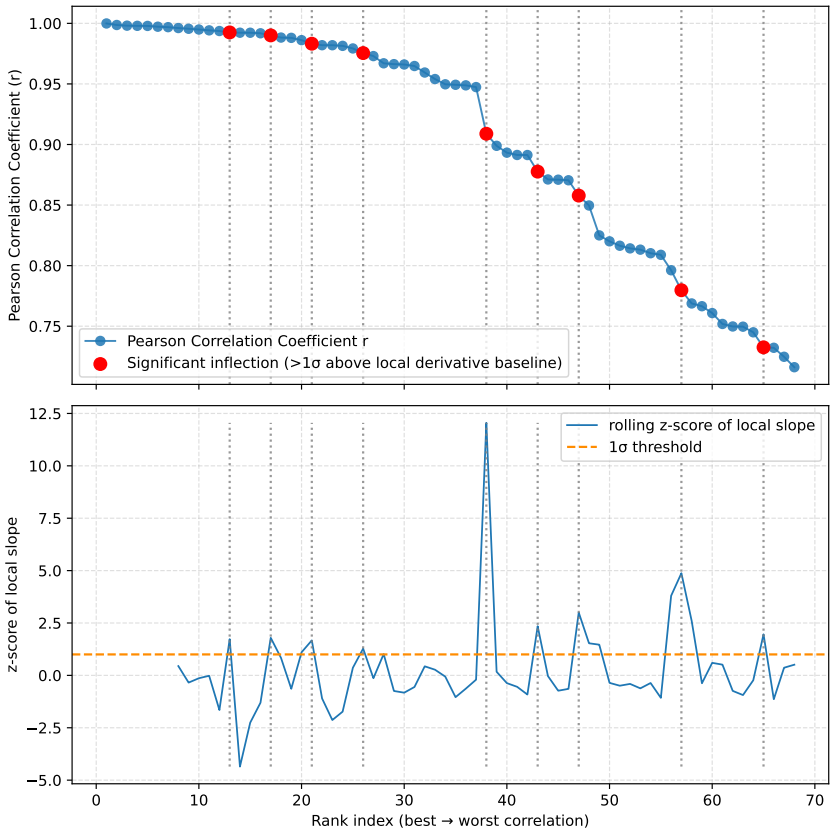

Figure 4: Rank-ordered Pearson correlations (z-scored). The first significant drop ("first anomaly") defines the truncation point for valid matches.

B.3 Shewhart-style anomaly detector

Our method is *inspired* by Shewhart's change-detection principle, but is adapted to slope irregularities in a rank-ordered curve rather than raw observations.

Let the sorted correlations be $r_1 \geq r_2 \geq \cdots$. We compute their first-order differences:

$$\Delta r_k = r_k - r_{k+1}.$$

Because the top of the curve is smooth, Δr_k is small and stable. A transition into the noise regime appears as a sudden increase in Δr_k—a discontinuity in slope.

We z-score these differences along the rank axis:

$$z_k = \frac{\Delta r_k - \mu_{\Delta r}}{\sigma_{\Delta r}},$$

and declare the first anomaly at the smallest k such that:

$$z_k > 1.$$

We use one standard deviation (1σ, not the classical 3σ) because Δr_k are differences of ordered correlations, they are not independent noise but locally correlated slopes. Thus, a 1σ threshold is appropriate for detecting the first structural discontinuity rather than extreme outliers. Also:

- the goal is to detect the *first* structural change, not an extreme outlier;
- the rank curve exhibits narrow variance near its maximum, making 1σ an effective and robust indicator of discontinuity.

All simulations with ranks above this first anomaly are retained; the remainder are treated as belonging to the high-dimensional noise floor.

B.4 Why this method works

This procedure:

- prevents false positives from the inherent correlation inflation in high dimensions,
- grounds match selection in structural continuity rather than magnitude,
- yields a reliable subset of simulations that genuinely reproduce the empirical network's topological signature.

Although conceptually similar to Shewhart control charts, our detector is specifically designed for rank-ordered similarity curves and does not assume independent or identically distributed observations.

C Principal Component Analysis (PCA)

To identify which dynamical ingredients contribute most strongly to the variability in collective topology across simulations, we performed a principal component analysis (PCA) on the z-scored simulation parameters:

$$\{\text{FoV},\ \omega,\ \xi_{\text{out}},\ NN,\ \text{motility},\ \text{max motility},\ \langle R^2 \rangle,\ \sigma(R^2)\}.$$

PCA was used solely as an exploratory tool: it reveals how parameters covary and which combinations contribute most to the variance in the simulation ensemble. Importantly, *no principal component regression (PCR)* was performed; we did not define a response variable nor regress PCs onto persistent features. The analysis therefore quantifies the intrinsic structure of the parameter space rather than predicting topological outcomes.

The PCA loadings and explained variances are presented in Table 2. The first four components explain 96.7% of the total variance, with:

- **PC1 (52.3%)** – dominated by motility and communication noise, capturing activity-level and uncertainty-driven variability.
- **PC2 (16.8%)** – dominated by decision-linearity metrics ($\langle R^2 \rangle$, $\sigma(R^2)$), separating linear vs. nonlinear reasoning regimes.
- **PC3 (15.8%)** – couples field-of-view and motility, contrasting wide-vision smooth movers with narrow-vision agile species.
- **PC4 (11.9%)** – primarily reflects variation in maximum turning angle (ω), capturing maneuverability differences.

Table 2: PCA loadings for z-scored simulation parameters. The first four components explain 96.7% of the variance.

Feature	**PC1**	**PC2**	**PC3**	**PC4**	**PC5**
fov z-score	-0.315	-0.381	0.604	-0.327	0.385
omega z-score	-0.334	0.123	-0.049	0.813	0.378
noise_out z-score	0.566	0.233	-0.020	-0.125	0.749
motility z-score	0.447	0.060	0.342	0.244	-0.094
max_motility z-score	0.387	-0.012	0.510	0.274	-0.355
mean_R2 z-score	-0.332	0.486	0.505	0.010	0.023
std_R2 z-score	-0.096	0.738	-0.025	-0.286	-0.113
Explained variance (%)	**52.3**	**16.8**	**15.8**	**11.9**	**2.1**

Findings

The PCA identifies several dominant axes of variability that complement the correlation and matching analyses in the main text:

- **Agility and communication precision dominate PC1.** High loadings from motility, max-motility, and communication noise (ξ_{out}) indicate that activity level and uncertainty strongly shape the global variation across simulations. These same parameters appear as primary discriminators across species in Table 1.
- **Decision-making complexity forms an independent axis PC2.** The strong contributions of $\langle R^2 \rangle$ and $\sigma(R^2)$ show that linear vs. nonlinear reasoning constitutes a major dimension of variation, orthogonal to basic kinematic quantities.
- **Field-of-view and motility jointly structure PC3.** This component separates wide-view species (birds, dolphins, bats) from narrow-view, highly agile species (primates, rodents), mirroring the trade-offs described in the Results.
- **Turning-angle variability is captured cleanly by PC4.** The strong loading of ω suggests that maneuverability forms its own dimension, consistent with the observed negative correlations between ω and both FoV and NN.

Together, these components reveal that **maneuverability, sensory aperture, communication noise, and reasoning linearity span the principal axes of variation** in the Smart-Boid parameter space. This structure helps explain why specific species map to distinct parameter regions and reinforces the finding that these few ingredients are sufficient to generate the diversity of real animal communication topologies.

D Pattern Classification: Feature Vector and Clustering Procedure

Emergent collective-motion patterns were classified using a custom Python pipeline written specifically for this project. The full implementation (over 1,000 lines of feature engineering, PCA reduction, clustering routines, and faceted visualization tools) is available in the project repository. The analysis relies on standard scientific Python libraries (`numpy`, `pandas`, `networkx`, `scikit-learn`, `scipy`, `matplotlib`, and `hdbscan`/`joblib`), together with our own feature-extraction and topology-aware routines tailored to Smart-Boids data.

Each simulation was converted into a feature vector using the pipeline shown in the code excerpt below (see project repository for full code):

```
# Snapshot-only phase clustering for Smart-Boids simulations
requirements:  numpy, pandas, networkx, scikit-learn, hdbscan,
matplotlib, scikit-image
```

The clustering methodology proceeds in three steps:

1. **Feature extraction.** For each simulation, structural, spatial, and heading-based descriptors were computed (polar/nematic alignment, milling circulation, kNN-graph connectivity, component spans, aspect ratio, structure-factor anisotropy, density CV, etc.), then averaged over the last 100 timesteps to capture the steady-state pattern. All features lie in $[0, 1]$ by design.
2. **Dimensionality reduction.** The feature matrix was z-scored with `RobustScaler` and projected onto its leading principal components (6 components retained), suppressing noise while preserving the dominant geometric degrees of freedom.
3. **Unsupervised clustering.** K-means was run on the PCA-reduced embedding with

   $$k = 24, \qquad \text{n\_init} = 200.$$

   The large value of k intentionally over-partitions the phase space, allowing finer substructure to be discovered before small clusters are merged. Clusters containing fewer than three simulations were removed using a nearest-neighbor reassignment rule.

After filtering, this procedure yielded a figure similar to the **11 stable pattern classes** manually classfied in Giardini et al. [2024], which are reported in Fig. (1b) of the main text. These classes include flocks, milling states, polarized bands, lane formations, stationary swarms, symmetry-broken flows, and mixed-mode transitions. The agreement between the automated classifier and human visual categorization confirms that the extracted feature representation reliably captures the qualitative structure of collective states across the Smart-Boids parameter space.

E Dispersion Relation Analysis

To assess the dynamical stability of the collective patterns observed in Smart-Boid simulations, we computed dispersion relations from the velocity field. Although commonly used in statistical physics, dispersion curves serve here as a compact summary of how small perturbations propagate through a coordinated group.

Construction of the dispersion relation

For each long-run simulation, we sampled agent velocities $\mathbf{v}_i(t)$ over time and constructed a coarse-grained velocity field on a uniform grid. The temporal Fourier transform of each spatial mode,

$$\tilde{\mathbf{v}}(\mathbf{k}, \omega) = \sum_t \sum_i \mathbf{v}_i(t)\, e^{-i(\mathbf{k}\cdot\mathbf{r}_i(t) - \omega t)},$$

yields a power spectrum $S(\mathbf{k}, \omega) = \|\tilde{\mathbf{v}}(\mathbf{k}, \omega)\|^2$. For visualization and analysis, we azimuthally average the spectrum to obtain a one-dimensional dispersion curve:

$$S(k, \omega) = \langle S(\mathbf{k}, \omega) \rangle_{\|\mathbf{k}\| = k}.$$

Interpretation

A dispersion curve reveals which temporal frequencies dominate at each spatial scale. In coordinated groups, a stable formation produces a ridge or band of high spectral intensity, indicating that perturbations at wavenumber k oscillate at a characteristic frequency $\omega(k)$. Conversely, a disordered group yields a broad, diffuse spectrum.

In our simulations, patterns such as flocks, lanes, or laminar flows showed narrow, well-defined modes, whereas disordered swarms exhibited broadband structure. The presence of sharp modes indicates that the pattern behaves like a stable dynamical attractor: small disturbances do not grow or propagate chaotically but decay into the dominant mode.

Stability of collective states

Across all 3,000 simulations, the dominant modes remained stable for long durations (up to 10^5 timesteps), and the linewidths of the spectral peaks remained narrow. This demonstrates that, once formed, collective states are dynamically rigid: transitioning out of a formation requires a large perturbation.

This result complements the pattern-matching results in the main text:

- fear-driven cohesion and sensory asymmetries generate the formation;
- but dispersion analysis shows that *maintaining* or *changing* formation requires external forcing;
- implying that ecological drivers (predators, resource pulses, leadership events, reproductive state changes) likely govern transitions between formation types in nature.

Why the dispersion analysis is included here

Because dispersion relations are not standard tools in animal-communication studies, but provide a rigorous measure of dynamical stability, we include the detailed construction in this appendix rather than the main text. The qualitative interpretation that after being trained, the Smart-Boid formations remain stable unless externally disrupted is summarized in the Results section, while this appendix provides the methodological detail for readers familiar with spectral analysis of dynamical systems.

F Additional Figures

We include additional figures omitted from the 5-page main paper for space reasons:

- **Fig. A1** — Full persistence diagram example.
- **Fig. B1** — Correlation rank curve and anomaly detection.

We also included the additional principal component analysis table 2, which shows the dimensions corresponding to highest data variability in the animal vs simulated social networks comparison dataset.

These provide complete transparency for readers reproducing or extending the analysis.