Speed-Weighted Flocking for Sailing Swarms under Dynamic Environmental Forcing

Pranav Kedia[1,2][0000−0003−4614−9540] ⋆, Aaron Gan[3], Hannah J. Williams[1,4,5][0000−0002−6338−529X], Andreagiovanni Reina[1,2,5][0000−0003−4745−992X], and Heiko Hamann[1,2][0000−0002−2458−8289]

[1] Centre for the Advanced Study of Collective Behaviour, University of Konstanz, Konstanz, Germany
{pranav.kedia, hannah.williams, andreagiovanni.reina, heiko.hamann}@uni-konstanz.de
[2] Department of Computer and Information Science, University of Konstanz, Konstanz, Germany
[3] University of Pittsburgh, Pittsburgh, PA, USA
[4] Department of Biology, University of Konstanz, Konstanz, Germany
[5] Max Planck Institute of Animal Behavior, Konstanz, Germany

Abstract. Collective behavior models, such as aggregation and flocking, typically assume self-propelled robots can directly execute desired speeds and headings. However, autonomous sailing robots violate this assumption. Their motion is shaped by wind-dependent propulsion, restricted headings, and spatially varying wind conditions. In particular, maneuverability is coupled to wind speed: in weak wind, sailboats may turn only slowly or not at all, whereas stronger wind enables faster turns. This introduces transient heterogeneity in speed and maneuverability across the flock. We focus on this fast-slow coordination problem in sailing robot flocks. To study this problem, we introduce *SailSwarmSwIM*, a reduced-order simulator that captures wind-dependent speed and maneuverability, no-go zones, tacking behavior, and steady or gusty wind fields. To design our novel flocking technique, we start from the Couzin model and introduce a speed-weighted social interaction rule that accounts for each robot's transient motion constraints. Across four wind conditions, moderate slow-neighbor weighting combined with speed matching under sailing constraints improves compactness, reduces unsafe events, and improves polarization under sailing constraints. This effect arises from a balance between attraction to fast neighbors, which helps maintain movement, and cohesion around slow neighbors, which prevents the flock from fragmenting. Together, SailSwarmSwIM and the speed-weighted interaction rule provide a framework for studying collective behavior in robotic fleets whose motion capabilities are continuously shaped by wind.

Keywords: Autonomous Sailing Robots · Sailing Swarm Simulator · Flocking · Collective Behavior · Wind-powered Locomotion

⋆ Corresponding author

1 Introduction

Classical models of collective motion show how coherent group-level behavior emerges from local interactions, such as repulsion, orientation, and attraction, without centralized control. Many foundational models of collective motion assume fully homogeneous system [10,16,1]. Homogeneous social interactions and homogeneous locomotion capabilities are strong assumptions that can complicate the design of controllers for robots whose motion is constrained by actuation limits, morphology, and environmental forcing. More recent work has relaxed these assumptions by studying how individual heterogeneity shapes collective movement [8,5]. Two-species Vicsek models [2] and [6] show that variable individual speed and speed-dependent turning can qualitatively alter collective motion, affecting polarization, disorder transitions, spatial extent, and group-size effects.

Our focus on motion-constrained heterogeneous collective motion originates from the challenge of building a swarm of autonomous sailing robots. Wind-driven sailing robots cannot move freely in arbitrary directions. Their speed and maneuverability depend on wind speed, wind direction, sail trim, rudder control, and their heading relative to the wind [13,14]. In weak winds, turns can become slow, unreliable, or even temporarily impossible; conversely, stronger winds can enable faster and more reliable maneuvers. Sailboats also cannot sail directly upwind but must perform tacking maneuvers to reach targets inside the upwind "no-go" zone [7]. Any desired heading produced by a flocking controller must be transformed into a feasible sailing direction. The mapping from social interaction rules to realized motion is therefore indirect, state-dependent, and environmentally mediated.

These constraints become more important in a swarm. In a fleet of sailing robots, each robot may experience different yet spatially correlated local wind conditions, apparent wind angles, and tacking constraints [13]. Even when the robots are physically homogeneous, the environment introduces temporary heterogeneity in speed and maneuverability. Similar differences are common in animal collectives, where individuals exhibit phenotypic variation and may differ in speed, energetic state, age, morphology [4], or maneuverability [5]. Cohesion matters because a fragmented sailing swarm may lose local interaction links, leave wind-constrained robots behind, and cease to operate as a coherent sensing or monitoring team. We therefore treat reduced flock spread as improved cohesion, while evaluating unsafe proximity separately to avoid equating cohesion with crowding. Maintaining cohesion often requires locomotor compromise through adaptive behaviors: faster or less constrained individuals may slow down, deviate, or otherwise adjust their movement to remain with slower or more constrained group members [11]. For artificial swarms, this raises an important design question: should social influence be weighted towards fast neighbors that can make rapid progress, or towards slower neighbors that risk being left behind?

Starting from a classical zonal flocking controller based on attraction, alignment, and repulsion [1], we introduce a speed-weighted social interaction rule. This rule defines a behavioral continuum ranging from following the fastest-moving neighbor to anchoring the flock around the slowest neighbors that are

currently more constrained. Slow-neighbor weighting may improve cohesion and safety by preventing faster robots from overstretching or fragmenting the flock. A central question is therefore whether an intermediate weighting regime can reduce proximity risks while preserving group polarization.

Before implementing our flocking method on autonomous sailboat hardware, we test it in *SailSwarmSwIM*, our dedicated reduced-order simulator for autonomous sailing robot swarms. Prototyping and testing in simulation is essential due to the high cost of deployment (e.g., unpredictable weather, high operational cost, and safety risks) [13,14]. This approach is also consistent with prior work on single-boat controllers, which are often tested in simulation before deployment [15,12]. SailSwarmSwIM focuses on sailing-specific wind-dependent speed, no-go zones, tacking behavior, and configurable steady- or gusty-wind fields.

We formulate sailing robot flocks as collective-behavior systems shaped by wind-driven locomotor constraints. We compare the baseline Couzin-type controller [1] against speed-weighted variants that combine modified social heading selection with sail luffing, a velocity-modulation mechanism that reduces propulsion by easing the sail. Our results show that moderate slow-neighbor weighting improves alignment, safety, and cohesion.

2 Wind-Driven Sailing Robot Model

We use a reduced-order planar model. The model does not reproduce full six-degree-of-freedom hydrodynamics; instead, it captures the locomotor constraints relevant for swarm coordination: wind-dependent speed, infeasible upwind headings, tacking, and spatially or temporally varying wind. These constraints are sufficient to render the mapping from the desired social direction to the realized motion indirect and environment-dependent.

Sailing terminology. We use the following sailing terms throughout. The *no-go zone* is the range of headings too close to the upwind direction for useful propulsion. *Tacking* is the maneuver of alternating between feasible headings on either side of the wind to make upwind progress toward a target inside the no-go zone. A *close-hauled* heading is the closest feasible heading to the wind direction. *Luffing* reduces aerodynamic drive; we use it as an optional speed-matching mechanism (Sect. 3), separate from the speed-weighted social rule itself.

Reduced-order sailing dynamics. Each robot i has state $\mathbf{x}_i = (x_i, y_i, \psi_i)$, where (x_i, y_i) is its position in the horizontal plane and ψ_i is its heading. The controller outputs a desired heading, which is mapped to feasible motion through simplified sailing dynamics. Forward speed depends on heading relative to local wind, reflecting that sailboats cannot directly command arbitrary velocities. Both sail trim and rudder action must be coordinated with the wind [14]. The position update is simply

$$\begin{pmatrix} x_i(t+\Delta t) \\ y_i(t+\Delta t) \end{pmatrix} = \begin{pmatrix} x_i(t) \\ y_i(t) \end{pmatrix} + \Delta t \left[v_i(t) \begin{pmatrix} \cos\psi_i(t) \\ \sin\psi_i(t) \end{pmatrix} + \mathbf{u}_{\mathrm{drift},i}(t) \right], \tag{1}$$

for the wind-dependent forward speed $v_i(t)$, surface drift $\mathbf{u}_{\text{drift},i}(t)$, and simulation time step Δt. We focus on wind-driven effects and use calm-water conditions. The heading updates by $\psi_i(t + \Delta t) = \psi_i(t) + \Delta t\, \omega_i(t)$, where $\omega_i(t)$ is the commanded yaw rate after applying sailing feasibility constraints.

No-go constraint and tacking. Headings inside an upwind no-go cone are infeasible and must be replaced. A sailing robot can reach any waypoint virtually under any wind condition, but needs to tack. The robot does so by alternating between port and starboard headings to make upwind progress [7]. When the desired direction lies within the no-go cone, the robot selects a nearby close-hauled heading and proceeds with tacking. Let $\theta_w(\mathbf{x}_i, t)$ denote the local wind direction at robot i, and let α_{ng} be the half-angle of the no-go cone. We define the signed angular difference between desired heading and upwind direction as

$$\delta_i = \operatorname{atan2}\left(\sin(\psi_i^\star - \theta_w(\mathbf{x}_i, t)), \cos(\psi_i^\star - \theta_w(\mathbf{x}_i, t))\right). \tag{2}$$

The desired heading is kept if $|\delta_i| \geq \alpha_{\text{ng}}$, and otherwise projected to the nearest close-hauled heading $\tilde{\psi}_i = \theta_w(\mathbf{x}_i, t) + \operatorname{sgn}(\delta_i)\, \alpha_{\text{ng}}$ and $\tilde{\psi}_i = \psi_i^\star$ otherwise.

This projection preserves the controller's intended direction as much as possible while enforcing the basic locomotor constraint of sailing. As a result, two robots receiving similar controller outputs may produce different motions depending on local wind conditions.

Wind field. The wind field is modeled as a base wind plus optional gusts: $\mathbf{w}(x, y, t) = \mathbf{w}_{\text{base}}(t) + \mathbf{w}_{\text{gust}}(x, y, t)$. Here, $\mathbf{w}(x, y, t) \in \mathbb{R}^2$ denotes the local wind-velocity vector. Its magnitude $\|\mathbf{w}(x, y, t)\|$ gives the local wind speed, and its direction defines the local wind direction $\theta_w(x, y, t)$. The base wind $\mathbf{w}_{\text{base}}$ sets the nominal speed and direction. Gusts are modeled as a coherent sinusoidal disturbance that propagates along the base-wind direction, $\mathbf{w}_{\text{gust}}(\mathbf{p}, t) = A \sin\big(\omega t - k\,(\hat{\mathbf{u}} \cdot \mathbf{p})\big)\, \hat{\mathbf{u}}$, where $k = 2\pi/\lambda$ is the spatial wavenumber, and $\omega = k\, c$ is the temporal angular frequency induced by propagation speed U_g, $\mathbf{p} = (x, y)$. Here, A, λ, c and $\hat{\mathbf{u}}$ denote amplitude, wavelength, propagation speed, and base-wind direction. In steady conditions, we set $A = 0$; in gusty conditions we use $A = 2\,\text{m/s}$, $\lambda = 50\,\text{m}$, and $c = 10\,\text{m/s}$ (temporal period $\lambda/c = 5\,\text{s}$). These gusts introduce the spatial and temporal variation that induces transient locomotor heterogeneity: mechanically identical robots differ in instantaneous ability to follow the flocking rule depending on local wind, point of sail, tacking requirements, and gusts. We evaluate four wind conditions: steady and gusty variants of 5 m/s and 10 m/s base wind.

Simulation loop. At each step, the simulator evaluates local wind, computes the social heading, applies sailing-feasibility constraints, and advances the robot state using Eq. (1) and the heading update.

3 Flocking Controllers

We use a Couzin-type controller whose repulsion, orientation, and attraction zones represent collision avoidance, alignment, and cohesion. Speed-dependent influence is applied only to the social terms; short-range repulsion remains unchanged. The controller computes a desired social heading, which is passed through the sailing controller described in Sect. 2. We compare a baseline Couzin-type flocking controller with our novel speed-weighted variants.

Baseline: Couzin-type flocking. The baseline controller follows the zonal interactions of Couzin-type flocking models [1,3]. Let $\mathbf{p}_i = (x_i, y_i)^\top$ denote the planar position of robot i. Each robot i observes neighboring robots $j \in \mathcal{N}_i$ within a finite interaction range and partitions them into three zones according to their distance $d_{ij} = \|\mathbf{p}_j - \mathbf{p}_i\|$. Robots within the repulsion zone ($d_{ij} \le r_{\text{rep}}$) are too close and trigger collision-avoidance behavior. Robots inside the orientation zone ($r_{\text{rep}} < d_{ij} \le r_{\text{ori}}$) influence heading alignment. Robots inside the attraction zone ($r_{\text{ori}} < d_{ij} \le r_{\text{att}}$) pull toward the local group. Here, r_{rep}, r_{ori}, and r_{att} denote the radii of the repulsion, orientation, and attraction zones, respectively. For each neighbor j, we define $\hat{\mathbf{e}}_j = (\cos\psi_j, \sin\psi_j)^\top$. The repulsion, orientation, and attraction terms are then

$$\mathbf{r}_i = -\sum_{\substack{j\in\mathcal{N}_i \\ d_{ij}\le r_{\text{rep}}}} \frac{\mathbf{p}_j - \mathbf{p}_i}{d_{ij}}, \quad \mathbf{o}_i = \sum_{\substack{j\in\mathcal{N}_i \\ r_{\text{rep}}< d_{ij}\le r_{\text{ori}}}} \hat{\mathbf{e}}_j, \quad \mathbf{a}_i = \sum_{\substack{j\in\mathcal{N}_i \\ r_{\text{ori}}< d_{ij}\le r_{\text{att}}}} \frac{\mathbf{p}_j - \mathbf{p}_i}{d_{ij}}. \tag{3}$$

The controller uses the standard priority order: repulsion, then orientation and attraction [1,3]. If at least one neighbor lies inside the zone of repulsion, the desired social direction is determined by $\mathbf{r}_i$. Otherwise, the robot combines the orientation and attraction terms $\mathbf{d}_i^{\text{social}} = \mathbf{o}_i + \mathbf{a}_i$. This direction-only Couzin controller is our baseline. It computes a desired social heading, while speed changes arise passively from wind-dependent sailing dynamics.

Speed-weighted social influence. The baseline controller assigns equal influence to all neighbors within the same interaction zone. In wind-driven robot flocks, this ignores the fact that robots may differ in their instantaneous locomotor capacity. A robot on a favorable point of sail may move and respond quickly, while an unfavorable apparent wind angle, a no-go constraint, or a tacking maneuver may slow another. We introduce a speed-weighted variant in which a neighbor's influence depends on their current speed. This information could be communicated between autonomous sailboats. For considered robot i and each of its neighboring robots $j \in \mathcal{N}_i$, we define a regularized inverse-power law

$$w_{ij} = \frac{1}{(v_j + \varepsilon)^\gamma}, \tag{4}$$

where $v_j > 0$ is the current speed of neighbor j, ε is a small regularization constant, and γ is the speed-weighting exponent. The weighted orientation and

attraction terms are

$$\mathbf{r}_i = -\sum_{\substack{j\in\mathcal{N}_i \\ d_{ij}\le r_{\mathrm{rep}}}} \frac{\mathbf{p}_j-\mathbf{p}_i}{d_{ij}}, \quad \mathbf{o}_{\gamma,i} = \sum_{\substack{j\in\mathcal{N}_i \\ r_{\mathrm{rep}}<d_{ij}\le r_{\mathrm{ori}}}} w_{ij}\hat{\mathbf{e}}_j, \quad \mathbf{a}_{\gamma,i} = \sum_{\substack{j\in\mathcal{N}_i \\ r_{\mathrm{ori}}<d_{ij}\le r_{\mathrm{att}}}} w_{ij}\frac{\mathbf{p}_j-\mathbf{p}_i}{d_{ij}} \tag{5}$$

and the corresponding social vector is $\mathbf{d}^{\mathrm{social}}_{\gamma,i} = \mathbf{o}_{\gamma,i} + \mathbf{a}_{\gamma,i}$.

The exponent γ in Eq.(4) is important as it defines the slow-fast behavioral axis of our novel flocking algorithm. For positive γ, the rule gives larger weights to smaller v_j, whereas for negative γ this preference is reversed. The offset $\varepsilon > 0$ removes the pole at $v_j = 0$, so the weighting remains finite over the admissible domain. We have several qualitatively different options: $\gamma < 0$ gives greater influence to faster neighbors, $\gamma = 0$ recovers uniform social weighting ($\forall i, j \in \mathcal{N}_i : w_{ij} = 1$), and $\gamma > 0$ gives greater influence to slower neighbors. Importantly, $\gamma = 0$ does not remove wind-dependent speed variation from the robots; it only removes explicit speed-dependent social weighting. Here, γ changes the social weighting of neighbors rather than acting as a direct throttle command. The uniform Couzin baseline uses equal social weights and therefore does not exploit speed differences when determining the desired social direction. Sweeping γ allows us to find parameters balancing between cohesion, collision reduction, and polarization. The purpose is not to find a universally optimal flocking controller, but to expose a slow-fast tradeoff. If wind-favored robots dominate the social direction, the flock may stretch or fragment as slower robots fall behind. If slow robots dominate, the group may remain compact but lose directional alignment.

Safety and boundary handling. The repulsion term is not speed-weighted: local collision avoidance should not depend on whether a neighbor is fast or slow. If a neighbor enters the repulsion zone, the controller prioritizes separation using Eq. (3). Although SailSwarmSwIM can accommodate COLREGs-based vessel right-of-way rules (international collision-avoidance rules for vessels at sea) for maritime encounters [9], we use a lightweight hard-repulsion layer to focus on flocking under wind-driven locomotor constraints. Candidate headings that would immediately violate the threshold $d_{\mathrm{rep}} = 1.5\,\mathrm{m}$ are rejected in favor of the nearest feasible alternative. While d_{rep} is used only to activate the hard-repulsion controller, safety is measured by the cumulative number of unsafe proximity events, defined as pairwise distances below a smaller threshold $d_{\mathrm{near}} = 1.0\,\mathrm{m}$.

Speed Equalization by Sail Luffing. The speed-weighted rule sets heading but not speed, so a boat on a beam-reach may outrun one sailing close-hauled or stalled near the no-go zone ($\beta_i \lesssim \alpha_{\mathrm{ng}}$). To reduce this heading-induced speed dispersion, our speed-weighted controllers add a sail-luffing layer that depowers fast boats toward the neighborhood speed. Sail trim is $\tau_i \in [0, 1]$ (with 1 fully powered). Luffing is asymmetric: reducing sail force but not increasing it. With apparent-wind angle $\beta_i \in [0°, 180°]$ sail easing is most effective near head-to-wind ($\beta_i \to 0°$, lift-driven) and does little near dead downwind ($\beta_i \to 180°$, drag-driven). Therefore, the depower factor D depends both on τ_i and β_i as

Table 1. Default simulation and controller parameters.

param.	value	meaning	param.	value	meaning
N	10	number of robots	α_{ng}	45°	no-go half-angle
L_h	35 m	arena half-size	r_{rep}	4 m	social repulsion rad.
d_{rep}	1.5 m	hard-repulsion threshold	r_{ori}	10 m	orientation radius
d_{near}	1.0 m	unsafe-event threshold	r_{att}	18 m	attraction radius
γ	$[-2, 10]$	speed-weighting exponent	k_p	0.4	luffing gain
$\Delta^{\tau}_{\max}$	0.05	luffing trim rate limit			

$D(\beta_i, \tau_i) = 1-(1-\beta_i/180^\circ)^2(1-\tau_i)$. Each robot i sets its trim τ_i by a rate-limited proportional law driving it to the mean neighbor speed $\bar{v}_i = \frac{1}{|\mathcal{N}_i|}\sum_{j\in\mathcal{N}_i} v_j$:

$$\tau_i \leftarrow \mathrm{clip}_{[0,1]}(\tau_i - \Delta\tau_i)\ , \qquad \Delta\tau_i = \mathrm{clip}_{[-\Delta^{\tau}_{\max},\Delta^{\tau}_{\max}]}\left(k_p\,\frac{v_i-\bar{v}_i}{\max(\bar{v}_i,\varepsilon_v)}\right). \quad (6)$$

with proportional gain k_p, regularization ε_v, and per-step rate limit $\Delta^{\tau}_{\max}$. A boat faster than its neighbors eases and slows; one with no neighbors sets to $\tau_i = 1$. Since trim only depowers, the controller tends to reduce local speed dispersion and may approach the slowest locally sustainable pace.

Because speed determines social influence under the speed-weighting rule, luffing changes not only the physical motion of a boat but also its subsequent contribution to its neighbors' social decisions. The coupling is therefore bidirectional: speed-dependent weighting changes headings and, thereby, wind-induced speeds, while luffing changes speeds and, thereby, social weights. Consequently, the combined response is generated by a closed feedback loop rather than by two independent, additive mechanisms. Separate ablations remain possible, but they do not isolate context-independent contributions of weighting and luffing.

4 Experimental Design and Statistical Analysis

All experiments use a square surface arena with half-size $L_{\mathrm{h}} = 35\,\mathrm{m}$, so that robot positions lie in $x, y \in [-L_{\mathrm{h}}, L_{\mathrm{h}}]$. For flocking, we use a toroidal interpretation of positions when computing neighbor relationships and flock metrics, as in classical flocking studies. This avoids artificial fragmentation or any wall effects when a coherent flock crosses an arena boundary, while the sailing dynamics and wind field are still evaluated in the square domain.

For each run, a swarm of N robots is initialized by sampling positions uniformly at random inside a central disk of radius R_0, with headings sampled uniformly from $[-\pi, \pi)$. Unless otherwise stated, we use $N = 10$ robots and repeat each controller–environment configuration over 50 random seeds. The simulator advances with a time step of $\Delta t = 1\,\mathrm{s}$ up to a horizon of $T = 300\,\mathrm{s}$(Table 1).

Metrics. We evaluate three quantities: flock compactness, alignment, and safety. Flock compactness is measured by the convex-hull area $A_{\mathrm{hull}}(t)$ of the robot positions. Lower area indicates a more compact flock. Alignment is measured

by the polarization $\Phi(t) = \left\| N^{-1} \sum_{i=1}^{N} \hat{\mathbf{e}}_i(t) \right\|$, where $\Phi \approx 1$ indicates aligned headings and $\Phi \approx 0$ indicates disordered headings (under a maximum entropy assumption). This definition is consistent with the literature [16,1,3]. Safety is measured by the cumulative number of unsafe proximity events recorded during the analysis window.

Statistical analysis. For each independent simulation run, we ignore a transient phase ($t \in [0, 100)$), and compute one steady-state summary over the interval $t \in [100, 300]$ s. For flock area and polarization, we use the median value over this interval. For unsafe proximity events, we use the cumulative count. This produces one value per metric and simulation run.

We sweep γ over 25 values $\gamma \in [-2, 10]$, with denser sampling near the exponent sign change at $\gamma = 0$ to resolve the transition between fast-neighbor following, uniform Couzin weighting, and slow-neighbor anchoring. We evaluate each controller under four wind conditions (each under a calm-water condition): 5 m/s steady wind, 10 m/s steady wind, 5 m/s wind with gusts, and 10 m/s wind with gusts.

Each speed-weighted setting is compared against the baseline Couzin controller using seed-matched paired differences $\Delta_{\gamma,s} = M_{\gamma,s} - M_{\text{base},s}$, where $M_{\gamma,s}$ is the metric value for tested exponent γ and seed s, and $M_{\text{base},s}$ is the corresponding baseline value for the same seed and environment. For flock area and unsafe proximity events, negative $\Delta_{\gamma,s}$ indicates improvement (smaller metric values are better). For polarization, positive $\Delta_{\gamma,s}$ indicates improvement (higher alignment is better). Because unsafe-proximity events, are discrete and skewed, and because the paired differences for some metrics show tail deviations from normality, we use Wilcoxon signed-rank tests for pairwise comparisons against the baseline. For each metric and environment, p-values are corrected across all tested γ values using Holm correction. We report a change as a significant improvement only if Holm-corrected $p < 0.05$ and the median paired difference is in the improvement direction.

5 Results

Fig. 1 shows the results for alignment and safety for the two 10 m/s-wind settings based on our sweep of the speed-weighting exponent γ (Eq. 4). Statistical significance and improvements over the baseline are given in Fig. 2 and Table 2. Results in the other conditions (5 m/s wind) and for the compactness metric are available in the online supplementary material,[6] along with code and videos.

Moderate Slow-Neighbor Weighting under Luffing Provides the Best Compromise. The exponent γ parameterizes a continuous social-influence axis: $\gamma < 0$ gives faster neighbors greater authority over heading (fast-neighbor following) and $\gamma > 0$ anchors social influence toward slower neighbors.

[6] https://praked.github.io/publications/SailSwarmSim-SAB2026

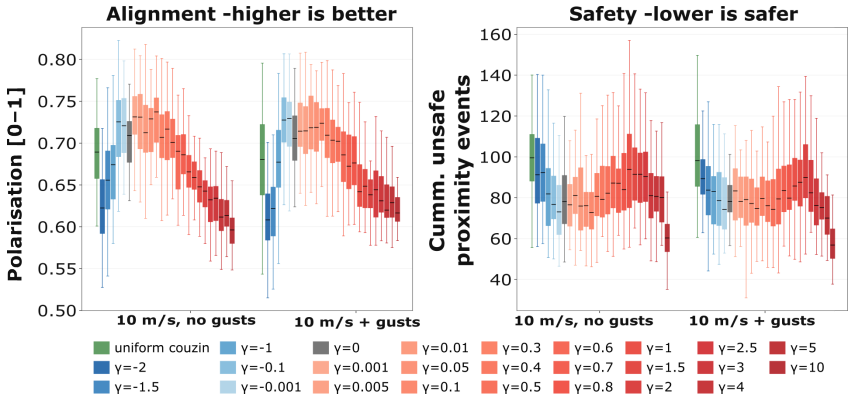

Fig. 1. Steady-state alignment and safety at 10 m/s across 50 runs per γ. Higher polarization and fewer cumulative unsafe-proximity events indicate better performance.

Alignment exhibits a clear inverted-U shape, with a central region above our Couzin model baseline. Safety also shows clear improvements over the baseline with a relatively complex system behavior for $\gamma > 1$. Within a narrow band $\gamma \in [-0.1, 0.3]$ every cell of Fig. 2 registers a significant improvement over the baseline, without any significant degradations. Outside this band, the picture deteriorates asymmetrically. Strong fast-neighbor following ($\gamma \leq -1.5$) significantly stretches the flock and reduces polarization: in the 10 m/s environments, at $\gamma = -2$, median flock area increases by 121.6 and 136.9 m^2 under steady and gusty wind, respectively, and polarization drops by 0.04–0.07. Strong slow-neighbor anchoring ($\gamma \geq 1$) preserves or improves cohesion but progressively erodes alignment, culminating at $\gamma = 10$ with a polarization loss of 0.09 and an area increase of 194.7 m^2 under steady 10 m/s wind. The two failure modes are visible as opposite limbs of Fig. 2.

A Robust Compromise across Cohesion, Safety, and Alignment. Within the central $\gamma \in [-0.1, 0.3]$ plateau for safety, small positive exponents γ provide significant gains in both safety and also cohesion as seen in Table 2 for the representative case $\gamma = 0.01$. The convex flock area shrinks by 45.6 m^2 to 106.2 m^2, corresponding to a 15 to 29% reduction relative to the baseline median (272 m^2 to 376 m^2). Cumulative unsafe proximity events drop by between 13.6 and 24.7 events, which is a 22 to 32% reduction relative to the baseline. Polarization rises modestly by 0.028 to 0.080. The improvement is largest in the highest-energy environment (10 m/s with gusts).

Within this coupled feedback loop, two effects plausibly contribute to the joint improvement in cohesion–safety. First, increased weighting of slower neighbors pulls the flock's centroid toward stalled or windward robots. Hence, we do not observe long-tail stretching following uniform alignment in gusty wind. Second, the sail-luffing layer reduces the relative speeds between leaders and trailers, lowering approach speed and the rate at which the unsafe-proximity layer is triggered; removing luffing erases most of this safety gain and can push

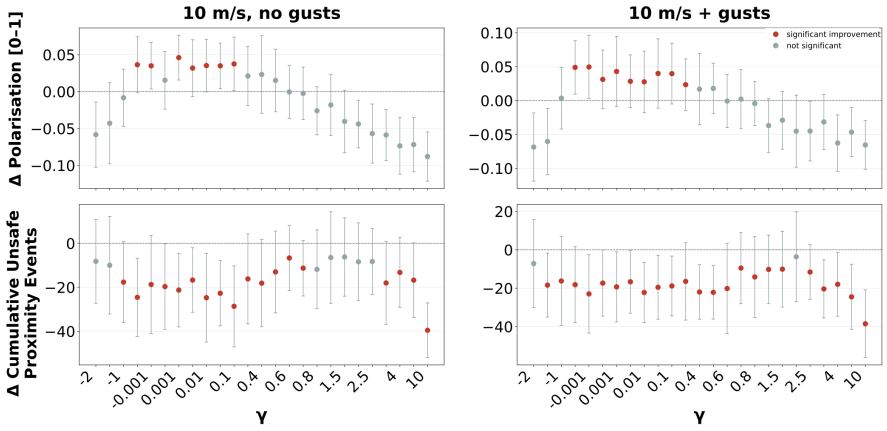

Fig. 2. Effect of speed weighting at 10 m/s: median paired differences from the uniform Couzin baseline over 50 seed-matched runs, with interquartile ranges. Positive $\Delta\Phi$ and negative ΔC indicate improved alignment and safety, resp.

Table 2. Effect of the slow-anchoring setting $\gamma = 0.01$ relative to the baseline (paired differences, 50 seed-matched runs; brackets show Cohen's d_z). Negative values improve area and unsafe proximity events; positive values improve polarization. Asterisks denote significance ($^{*}\,p < 0.05$; $^{**}\,p < 10^{-3}$; $^{***}\,p < 10^{-5}$).

Environment	γ	ΔA_{hull} [m^2]	$\Delta\Phi$	ΔC
5 m/s, steady	0.01	-61.4^{***} [−0.72]	$+0.080^{***}$ [+0.99]	-16.9^{***} [−0.73]
10 m/s, steady	0.01	-93.8^{***} [−0.93]	$+0.035^{*}$ [+0.49]	-24.7^{**} [−0.64]
5 m/s, gusts	0.01	-45.6^{*} [−0.52]	$+0.052^{*}$ [+0.53]	-13.6^{***} [−0.68]
10 m/s, gusts	0.01	-106.2^{***} [−0.98]	$+0.028^{**}$ [+0.60]	-22.2^{***} [−0.75]

unsafe events back to baseline levels (Fig. S10). Comparing the behaviors based only on the luffing layer ($\gamma = 0$) with those incorporating both speed-weighting and luffing ($\gamma \neq 0$; see Figs. S4–S10), shows that speed weighting significantly increases polarization but has no significant effect on the other metrics.

In summary, we find that the speed-weighting layer must be combined with luffing to improve flocking performance, and that $\gamma = 0.01$ provides a robust default across metrics, but not as an optimum for any single metric. Within the plateau, per-metric optima differ: in the 10 m/s environments, area is minimized near $\gamma \approx 1$ and unsafe proximity events near $\gamma = 10$, but both settings significantly reduce polarization. By contrast, $\gamma = 0.01$ sacrifices a few percent of per-metric improvement in exchange for simultaneous gains on all metrics across all environments, without significant degradation.

Over-Anchoring Reveals an Alignment Tradeoff. Pushing γ beyond the plateau exposes a sharp alignment tradeoff. For $\gamma \geq 1$ in the 10 m/s environments, alignment drops significantly relative to baseline. Under steady wind, $\Delta\Phi$ decreases from -0.018 at $\gamma = 1$ to -0.088 at $\gamma = 10$, with qualitatively the

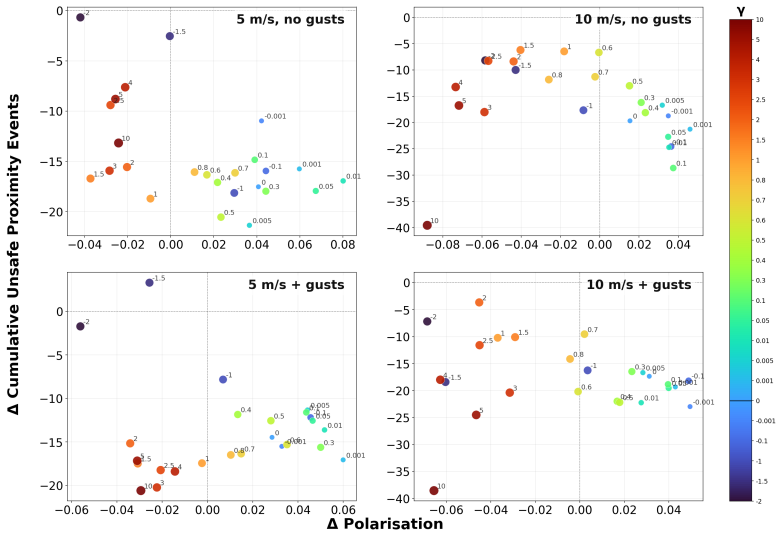

Fig. 3. Alignment–safety tradeoff across wind conditions. Markers are γ settings as a median paired effect relative to the uniform Couzin baseline. The favorable region is lower right: higher polarization and fewer unsafe proximity events.

same monotone degradation under gusts. Cohesion continues to improve over this range (ΔA_{hull} remains negative until $\gamma = 10$, at which point it reverses to $+194.7\,\text{m}^2$ under steady 10 m/s wind). This can be seen as over-anchoring: an overly strong influence of slow robots preserves cohesion but degrades alignment. In Fig. 3, we show this tradeoff by plotting changes in polarization against changes in unsafe proximity events. The desirable region is the lower-right quadrant (increased polarization and reduced unsafe proximity). Across environments, small positive γ settings occupy this region.

6 Conclusion

We studied flocking in wind-driven robot swarms where social interaction rules are filtered through sailing-specific locomotor constraints. Because no-go zones, tacking, and local wind conditions make some robots temporarily faster or more constrained than others, flocking becomes a slow-fast coordination problem. This study shows the effect of speed-weighted social influence under luffing using $N = 10$ robots, calm-water conditions, and toroidal boundary handling for flock metrics. Safety is represented by a simplified hard-repulsion layer, and robots are assumed to have accurate access to neighboring positions, headings, and speeds. In physical sailing robotic swarms, these estimates may be affected by limited sensing and communication range, localization noise, delay, and packet loss. We introduced a speed-weighted Couzin controller with luffing in which $\gamma < 0$ implements fast-neighbor following and $\gamma > 0$ implements slow-neighbor anchoring. Across four wind conditions, moderate slow-neighbor anchoring improved com-

pactness, increased safety, and preserved or improved polarization. Our analysis shows that luffing conditionally contributes to improvements in safety and cohesion, while the speed-weighting exponent governs the alignment–cohesion tradeoff across the sweep. Stronger anchoring revealed an alignment tradeoff. Future work should therefore test goal-related tasks, highly temporally variable wind conditions, non-toroidal environments, and robustness to imperfect neighbor-state information.

Acknowledgments. This work has been supported by the DFG under Germany's Excellence Strategy, EXC 2117-422037984.

Disclosure of Interests. The authors have no competing interests to declare.

References

1. Couzin, I.D., Krause, J., James, R., Ruxton, G.D., Franks, N.R.: Collective memory and spatial sorting in animal groups. Journal of Theoretical Biology (2002)
2. Dutta, A.K., 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 **112**(2), 024137 (2025)
3. Giardina, I.: Collective behavior in animal groups: Theoretical models and empirical studies. HFSP Journal **2**(4), 205–219 (2008)
4. Harel, R., Loftus, J.C., Crofoot, M.C.: Locomotor compromises maintain group cohesion in baboon troops on the move. Proc. R. Soc. B (2021)
5. Jolles, J.W., King, A.J., Killen, S.S.: The role of individual heterogeneity in collective animal behaviour. Trends in Ecology & Evolution (2020)
6. Klamser, P.P., et al.: Impact of variable speed on collective movement of animal groups. Frontiers in Physics (2021)
7. Le Bars, F., Jaulin, L.: The world robotic sailing championship, a competition to stimulate the development of autonomous sailboats. In: OCEANS 2015 - Genova. pp. 1–5. IEEE (2015)
8. del Mar Delgado, M., et al.: The importance of individual variation in the dynamics of animal collective movements. Phil. Trans. R. Soc. B: Biol. Sciences (2018)
9. Naeem, W., Irwin, G.W., Yang, A.: COLREGs-based collision avoidance strategies for unmanned surface vehicles. Mechatronics **22**(6), 669–678 (2012)
10. Reynolds, C.W.: Flocks, herds and schools: A distributed behavioral model. In: Proceedings of SIGGRAPH '87. pp. 25–34. ACM (1987)
11. Sankey, D.W.E., et al.: Speed consensus and the 'Goldilocks principle' in flocking birds (*Columba livia*). Animal Behaviour (2019)
12. Stenersen, H.S., et al.: Construction and control of an autonomous sail boat. In: IFAC-PapersOnLine. vol. 49. Elsevier (2016)
13. Sun, Q., Qi, W., Liu, H., Ji, X., Qian, H.: Toward long-term sailing robots: State of the art from energy perspectives. Front. in Robotics and AI (2022)
14. Tipsuwan, Y., Sanposh, P., Techajaroonjit, N.: Overview and control strategies of autonomous sailboats—a survey. Ocean Engineering **281**, 114879 (2023)
15. Vautier, U., et al.: Development and test of an open-source autonomous sailing robot with accessibility, generality and extendibility. In: IRSC (2018)
16. 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)