

Speed-Weighted Adaptive Flocking for Sailing Swarms under Dynamic Environmental Forcing

Pranav Kedia^{1,2}, Aaron Gan³, Hannah J. Williams^{1,4,5}, Andreagiovanni Reina^{1,2,4}, and Heiko Hamann^{1,2}

¹ Centre for the Advanced Study of Collective Behaviour, Konstanz, Germany

² Dep. of Computer and Information Science, University of Konstanz, Germany

³ University of Pittsburgh, Pittsburgh, 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, usually assume self-propelled robots that can directly execute their desired speed and direction of motion without fundamental constraints. 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 for autonomous sailing robot swarms 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. A key result is that increasing the social influence of slower robots improves polarization and reduces close encounters. 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, our simulator, *SailSwarmSwIM*, and the speed-weighted interaction rule provide a modeling framework for studying adaptive collective behavior in robotic fleets whose motion capabilities are continuously shaped by wind.

Keywords: Autonomous Sailing Robots · Robotic Swarm Simulator · Flocking

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 [1]. Many foundational models of collective motion assume fully homogeneous systems, including the well-known models by

Reynolds [2], Vicsek [3], and Couzin [4]. 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 [5,6]. For example, two-species Vicsek models [7] and [8] 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 [9,10]. In weak winds, turns can become slow, unreliable, or even temporarily impossible; conversely, stronger winds can enable faster and more reliable maneuvers. Sailboats can also not sail directly upwind but must perform tacking maneuvers to reach targets inside the upwind “no-go” zone [11]. 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 (while probably correlated) local wind conditions, apparent wind angles, and tacking constraints [9]. 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 [12], or maneuverability [6]. 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 [13]. 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 [4], 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) [9,10]. This approach is also consistent with prior work on single-boat controllers, which are often tested in simulation before deployment [14,15]. 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 [4] 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.

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 [10]. The position update is simply

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

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

$$\psi_i(t + \Delta t) = \psi_i(t) + \Delta t \omega_i(t), \quad (2)$$

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 [11]. 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 the desired heading and the upwind direction as

$$\delta_i = \text{atan2}(\sin(\psi_i^* - \theta_w(\mathbf{x}_i, t)), \cos(\psi_i^* - \theta_w(\mathbf{x}_i, t))). \quad (3)$$

The desired heading is feasible if $|\delta_i| \geq \alpha_{\text{ng}}$. Otherwise, it is projected to the nearest close-hauled heading,

$$\tilde{\psi}_i = \theta_w(\mathbf{x}_i, t) + \text{sgn}(\delta_i)\alpha_{\text{ng}}, \quad \text{if } |\delta_i| < \alpha_{\text{ng}}, \quad (4)$$

and $\tilde{\psi}_i = \psi_i^*$ 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 defines the nominal wind speed and direction. Gusts introduce spatial and temporal variation. We use four wind conditions: 5 m/s steady wind, 10 m/s steady wind, 5 m/s wind with gusts, and 10 m/s wind with gusts. These wind conditions induce 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.

Simulation loop. At each simulation step, SailSwarmSwIM updates the wind field, evaluates the local wind at each robot, computes a desired heading direction from the selected flocking controller, projects this direction through the sailing-feasibility layer, and advances the robot state using Eqs. (1) and (2). We separate social interaction from locomotor-feasibility and, hence, the same flocking rule can produce different realized motions depending on the wind.

3 Flocking Controllers

We compare a baseline Couzin-type flocking controller with our novel speed-weighted variant, which modulates the influence of faster and slower neighbors in the alignment calculation. Both controllers compute a desired alignment direction from local information. This direction is then mapped to a target heading and passed to the low-level sailing controller (Sec. 2).

Baseline: Couzin-type flocking. The baseline controller follows the zonal interactions of Couzin-type flocking models [4,16]. 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} . Robots within the repulsion zone ($d_{ij} \leq r_{\text{rep}}$) are too close and trigger collision-avoidance behavior. Robots inside the orientation zone ($r_{\text{rep}} < d_{ij} \leq r_{\text{ori}}$) influence heading alignment. Robots inside the attraction zone ($r_{\text{ori}} < d_{ij} \leq r_{\text{att}}$) pull toward the local group.

Let $\mathcal{N}_i$ be the set of neighbors observed by robot i , and let $\mathbf{p}_i = (x_i, y_i)^\top$ denote its planar position. For each neighbor j , we define

$$d_{ij} = \|\mathbf{p}_j - \mathbf{p}_i\|, \quad \hat{\mathbf{r}}_{ij} = \frac{\mathbf{p}_j - \mathbf{p}_i}{d_{ij}}, \quad \hat{\mathbf{e}}_j = (\cos \psi_j, \sin \psi_j)^\top.$$

The repulsion, orientation, and attraction terms are then

$$\mathbf{r}_i = - \sum_{\substack{j \in \mathcal{N}_i \\ d_{ij} < r_{\text{rep}}}} \hat{\mathbf{r}}_{ij}, \quad \mathbf{o}_i = \sum_{\substack{j \in \mathcal{N}_i \\ r_{\text{rep}} \leq d_{ij} < r_{\text{ori}}}} \hat{\mathbf{e}}_j, \quad \mathbf{a}_i = \sum_{\substack{j \in \mathcal{N}_i \\ r_{\text{ori}} \leq d_{ij} < r_{\text{att}}}} \hat{\mathbf{r}}_{ij}, \quad (5)$$

here r_{rep} , r_{ori} , and r_{att} denote the radii of the repulsion, orientation, and attraction zones, respectively. The controller uses the standard priority order: repulsion, then orientation and attraction [4,16]. 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 rather than active speed matching or luffing (loss of sail drive when too close to the wind).

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

$$w_{ij} = \frac{1}{(v_j + \varepsilon)^\gamma}, \quad (6)$$

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 exponent γ 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 γ -sweep, therefore, changes the social weighting of neighbors, while the luffing layer provides the corresponding speed-equalization mechanism. The uniform Couzin baseline does not actively luff for speed matching; consequently, comparisons against this baseline measure the combined effect of speed-aware social weighting and sail-based speed modulation.

The weighted orientation and attraction terms are

$$\mathbf{r}_i = - \sum_{\substack{j \in \mathcal{N}_i \\ d_{ij} < r_{\text{rep}}}} \hat{\mathbf{r}}_{ij}, \quad \mathbf{o}_{\gamma,i} = \sum_{\substack{j \in \mathcal{N}_i \\ r_{\text{rep}} \leq d_{ij} < r_{\text{ori}}}} w_{ij} \hat{\mathbf{e}}_j, \quad \mathbf{a}_{\gamma,i} = \sum_{\substack{j \in \mathcal{N}_i \\ r_{\text{ori}} \leq d_{ij} < r_{\text{att}}}} w_{ij} \hat{\mathbf{r}}_{ij} \quad (7)$$

The corresponding social vector is $\mathbf{d}_{\gamma,i} = \mathbf{o}_{\gamma,i} + \mathbf{a}_{\gamma,i}$. The purpose of Eq. (6) is not to set the stage to find a universally optimal flocking controller, but to expose a slow-fast tradeoff. If fast or wind-favored robots dominate the social direction, the flock may stretch or fragment as slower robots fall behind. If slow robots dominate too strongly, the group may remain compact but lose directional alignment. Sweeping γ therefore allows us to test where the useful regime lies between cohesion, collision reduction, and polarization.

Speed Equalization by Sail Luffing. Our speed-weighted approach changes the desired collective heading but is agnostic to speed modulation. A robot on a favorable beam reach may move faster than a robot sailing close-hauled or stalled near the no-go zone, because sailboat speed depends on the point of sail. To reduce this heading-induced speed dispersion, our speed-weighted controllers include a simple sail-luffing layer that depowers fast boats toward the neighborhood speed. We represent sail trim by $\tau_i \in [0, 1]$, where $\tau_i = 1$ is fully powered and $\tau_i = 0$ is fully eased. Luffing is asymmetric: it can reduce the available sail force but cannot increase it beyond the wind. The depower factor is

$$1 - \left(\frac{\beta_i}{180^\circ} \right)^2 (1 - \tau_i), \quad (8)$$

where β_i is the apparent-wind angle (0° is dead downwind and 180° is head-to-wind). Hence, easing the sail is constrained downwind ($\beta_i \approx 180^\circ$) but strongly reduces lift-driven propulsion near head-to-wind ($45^\circ < \beta_i \ll 180^\circ$).

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. (5). Although SailSwarmSwIM can accommodate COLREGs-based vessel right-of-way rules (international collision-avoidance rules for vessels at sea) for maritime encounters [17], we use a lightweight hard-repulsion layer here

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	repulsion radius
d_{near}	1.0 m	unsafe proximity threshold	r_{ori}	10 m	orientation radius
γ	$[-2, 10]$	speed-weighting exponent	r_{att}	18 m	attraction radius
k_p	0.4	luffing proportional gain			

to keep the focus on flocking under wind-driven locomotor constraints. Candidate headings that would immediately violate the safety distance are rejected in favor of the nearest feasible alternative. A safety-threshold violation is recorded whenever any pair of boats comes within $d_{\text{safe}} = 1.5$ m. This metric captures repeated close encounters and provides a conservative proxy for safety.

4 Experimental Design and Statistical Analysis

All experiments use a square surface arena with half-size $L_h = 35$ m, so that robot positions lie in $x, y \in [-L_h, L_h]$. 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 $\Delta t = 1$ s up to a horizon of $T = 300$ s. Metrics are logged at each step, but statistical comparisons are performed on time-median over $t \in [100, 300]$ s per seed rather than on tick-level samples.

Controller and environment sweep. We compare the baseline Couzin-type flocking controller against our speed-weighted controller from Sec. 3. The speed-weighted controller uses $\varepsilon = 0.1$ and sweeps γ over both negative and positive values. Negative γ values correspond to fast-neighbor following; $\gamma = 0$ does not fully recover the Couzin baseline because we do luffing, and positive γ values correspond to slow-neighbor anchoring. 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: 5 m/s steady wind, 10 m/s steady wind, 5 m/s wind with gusts, and 10 m/s wind with gusts. All experiments reported here use calm-water conditions, so the observed differences arise from wind-driven locomotor constraints rather than wave-induced drift.

Metrics. We evaluate three quantities: flock compactness, alignment, and safety. Flock compactness is measured by the convex-hull area $A_{\text{hull}}(t)$ 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 [3,4,16]. 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.

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 collision counts 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

We report 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. 6) over 25 values $\gamma \in [-2, 10]$, see Fig. 1. Statistical significance and improvements over the baseline are given in Fig. 2 and Table 2. Supplementary results are available online.⁶

Intermediate Speed Weighting Improves Flocking Performance. 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. 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 the

⁶ <https://praked.github.io/publications/SailSwarmSim-SAB2026>




Fig. 1: Per-seed steady-state distributions for alignment and safety in the 5 m/s environments. Each box summarizes one value per seed over $t \in [100, 300]$ s across 50 runs. Higher polarization indicates stronger alignment; lower cumulative unsafe proximity events indicate safer flocking. Formal comparisons use paired seed-level summaries against the uniform Couzin baseline.

4×3 (env, metric) grid 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 $+136.9 \text{ m}^2$ under steady and gusty 10 m/s wind respectively, and polarization drops between 0.04 to 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 $+194.7 \text{ m}^2$ 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 $\gamma = 0.01$. All twelve cells reach significance, with paired effect sizes $|d_z| \in [0.49, 0.99]$. The convex flock area shrinks by 45.6 to 106.2 m^2 , corresponding to a 15 to 29% reduction relative to the baseline median (272 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), where the baseline controller is most stressed: $d_z = -0.98$ for area.

We identify two mechanisms that plausibly account for the joint improvement in cohesion–safety. First, increased weighting of slower neighbors pulls the flock’s




Fig. 2: Effect of speed-weighting on alignment and safety in the 10 m/s environments. Each panel shows the median paired difference relative to the uniform Couzin baseline over 50 seed-matched runs; error bars show the interquartile range. Positive $\Delta\Phi$ indicates improved alignment, while negative ΔC indicates fewer unsafe proximity events.

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]

centroid toward stalled or windward robots. Hence, we do not observe long-tail stretching following uniform alignment in gusty wind. Second, by reducing the relative speeds between leaders and trailers, slow-neighbor anchoring lowers the approach speed and the rate at which the unsafe proximity layer is triggered.

In summary, we find $\gamma = 0.01$ as a robust default, but not as a per-metric optimum. Within the plateau, per-metric optima occur at different points across the sweep: in the 10 m/s environments, area is minimized near $\gamma \approx 1$ and unsafe proximity events near $\gamma = 10$, but those same settings significantly reduce polarization. Choosing $\gamma = 0.01$ sacrifices the last few percent of per-metric improvement in exchange for simultaneous gains on all three metrics in all four environments, with no significant degradation in any cell.




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.

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 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 increases, 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. We introduced a speed-weighted Couzin controller in which $\gamma < 0$ implements fast-neighbor following and $\gamma > 0$ implements slow-neighbor anchoring. Across four

wind conditions, moderate slow-neighbor anchoring improved compactness, increased safety, and preserved or improved polarization. Stronger anchoring revealed an alignment tradeoff. Future work should extend this analysis to heterogeneous fleets, explicit energy costs, wave and current disturbances, and field experiments with physical sailing robots.

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

References

1. M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo, “Swarm robotics: a review from the swarm engineering perspective,” *Swarm Intelligence*, 2013.
2. C. W. Reynolds, “Flocks, herds and schools: A distributed behavioral model,” in *Proceedings of SIGGRAPH ’87*. ACM, 1987, pp. 25–34.
3. T. Vicsek, A. Czirók, E. Ben-Jacob, I. Cohen, and O. Shochet, “Novel type of phase transition in a system of self-driven particles,” *Physical Review Letters*, 1995.
4. I. D. Couzin, J. Krause, R. James, G. D. Ruxton, and N. R. Franks, “Collective memory and spatial sorting in animal groups,” *Journal of Theoretical Biology*, 2002.
5. M. del Mar Delgado *et al.*, “The importance of individual variation in the dynamics of animal collective movements,” *Phil. Trans. R. Soc. B: Biol. Sciences*, 03 2018.
6. J. W. Jolles, A. J. King, and S. S. Killen, “The role of individual heterogeneity in collective animal behaviour,” *Trends in Ecology & Evolution*, 2020.
7. A. K. Dutta, M. Mangeat, H. Rieger, R. Paul, and S. Chatterjee, “Stability of flocking in the reciprocal two-species Vicsek model: Effects of relative population, motility, and noise,” *Physical Review E*, vol. 112, no. 2, p. 024137, 2025.
8. P. P. Klamser, L. Gómez-Nava, T. Landgraf, J. W. Jolles, D. Bierbach, and P. Romanczuk, “Impact of variable speed on collective movement of animal groups,” *Frontiers in Physics*, 2021.
9. Q. Sun, W. Qi, H. Liu, X. Ji, and H. Qian, “Toward long-term sailing robots: State of the art from energy perspectives,” *Front. in Robotics and AI*, 2022.
10. Y. Tipsuwan, P. Sanposh, and N. Techajaroontjit, “Overview and control strategies of autonomous sailboats—a survey,” *Ocean Engineering*, vol. 281, p. 114879, 2023.
11. F. Le Bars and L. Jaulin, “The world robotic sailing championship, a competition to stimulate the development of autonomous sailboats,” in *OCEANS 2015 - Genova*. IEEE, 2015, pp. 1–5.
12. R. Harel, J. C. Loftus, and M. C. Crofoot, “Locomotor compromises maintain group cohesion in baboon troops on the move,” *Proceedings of the Royal Society B: Biological Sciences*, 2021.
13. D. W. E. Sankey, E. L. C. Shepard, D. Biro, and S. J. Portugal, “Speed consensus and the ‘Goldilocks principle’ in flocking birds (*Columba livia*),” *Animal Behaviour*, Nov. 2019.
14. U. Vautier, J. Wan, M. Dai, C. Viel, and R. Hone, “Development and test of an open-source autonomous sailing robot with accessibility, generality and extendibility,” in *IRSC*, 2018.
15. H. S. Stenersen *et al.*, “Construction and control of an autonomous sail boat,” in *IFAC-PapersOnLine*, vol. 49, no. 23. Elsevier, 2016, pp. 284–289.
16. I. Giardina, “Collective behavior in animal groups: Theoretical models and empirical studies,” *HFSP Journal*, vol. 2, no. 4, pp. 205–219, 2008.
17. W. Naeem, G. W. Irwin, and A. Yang, “COLREGs-based collision avoidance strategies for unmanned surface vehicles,” *Mechatronics*, vol. 22, no. 6, pp. 669–678, 2012.