PNAS RESEARCH ARTICLE | ENGINEERING

Noise-enabled goal attainment in crowded collectives

Lucy Liu[a], Justin Werfel[a], Federico Toschi[b,c], and L. Mahadevan[a,d,e,1]

Edited by John Rogers, Northwestern University, Evanston, IL; received July 18, 2025; accepted December 3, 2025

In crowded environments, individuals must navigate around other occupants to reach their destinations. Understanding and controlling traffic flows in these spaces is relevant for coordinating robot swarms and designing infrastructure for dense populations. Here, we use simulations, theory, and experiments to study how adding stochasticity to agent motion can reduce traffic jams and help agents travel more quickly to prescribed goals. A computational approach reveals the collective behavior. Above a critical noise level, large jams do not persist. From this observation, we analytically approximate the swarm's goal attainment rate, which allows us to solve for the agent density and noise level that maximize the goals reached. Robotic experiments corroborate the behaviors observed in our simulated and theoretical results. Finally, we compare simple, local navigation approaches with a sophisticated but computationally costly central planner. A simple reactive scheme performs well up to moderate densities and is far more computationally efficient than a planner, motivating further research into robust, decentralized navigation methods for crowded environments. By integrating ideas from physics and engineering using simulations, theory, and experiments, our work identifies new directions for emergent traffic research.

collective behavior | robotics | active matter

Significance

How many cooks in a kitchen make it too crowded to function? Here we tackle the mathematical challenge of predicting how quickly individuals can reach goal locations in a crowded environment. Having agents travel in less direct, more random paths can break up jams and improve the goal attainment rate. We develop estimates for travel times, then use them to find the optimal crowd density and level of randomness in the path that maximize the goals reached. We deploy our navigation strategies using robots to determine their efficacy in crowded environments. Our theoretical findings inform the effective design of robot teams and pedestrian spaces, and our experiments demonstrate new possibilities for studying crowd dynamics using robots.

Consider a robot team with a time-sensitive distributed task such as assembling a machine, fulfilling orders in a warehouse, or cleaning up hazardous debris. Robots must transport items to specific goal locations. When the space is relatively empty, adding robots is advantageous: Several robots together work faster than a lone one. However, adding too many robots will lead to traffic that slows the entire team down. Systems of many interacting agents are challenging to analyze mathematically. Approximation of a global metric, like how quickly a collective reaches goals, has been an elusive result in emergent traffic. Achieving it would enable optimized system design, including the ability to determine the optimal team size for a task.

Emergent traffic patterns like jam formation, laning, and various transitions between ordered and disordered behavior have been studied in diverse settings spanning car traffic (1, 2), colloids and bacteria (3), robots (4, 5), ants (6, 7), and humans (8–10). In these systems, the simple constraint that two agents cannot occupy the same location at the same time, so that agents must stop or slow down in high-traffic regions, produces a set of rich and interesting phenomena. For instance, it is known that collectives of random walkers with exclusion constraints alone and no attraction can self-organize into large jams (3), and that ants or pedestrians following simple rules can self-organize into lanes (7, 9). More recently, interest has grown in studying active particles with minimal intelligence, such as agents with sensing cones (11) and individual goal locations (12). A natural question is how to maximize the efficiency of target attainment in a crowded environment. How many agents should be used, and what navigation strategy should they follow? How effectively does noisy motion, the simplest mechanism for exiting jams, enable flow? How do complex global planners compare to a simple reflexive rule each agent follows on its own?

These questions sit at the nexus of decision-making, robotics, animal behavior, and statistical physics. In robotics, the problem of computing noncolliding paths between multiple agents and their goals has led to the Multi-Agent Pathfinding (MAPF) problem, for which finding time-optimal paths is NP-hard (13). One baseline algorithm for MAPF is Cooperative A* (14), a greedy algorithm where robots plan paths one after the other and can access full information about already-planned paths. Cooperative A* does not guarantee optimal solutions, so subsequent improvements have focused on speed and completeness (15). Variations of MAPF may allow robots to swap goals or may extend the problem to a lifelong version where new goals are continually added (16, 17). In animal behavior, it is now well accepted that simple local rules for carrying out collective tasks are often far more robust than global planning; this is most clearly visible in social

Author affiliations: [a]John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138; [b]Department of Applied Physics and Science Education, Eindhoven University of Technology, Eindhoven 5600 MB, The Netherlands; [c]Istituto per le Applicazioni del Calcolo "M. Picone", Consiglio Nazionale delle Ricerche, Rome I-00185, Italy; [d]Department of Physics, Harvard University, Cambridge, MA 02138; and [e]Department of Organismic and Evolutionary Biology, Harvard University, Cambridge, MA 02138

Author contributions: L.L., J.W., F.T., and L.M. designed research; L.L. performed research; L.L., J.W., F.T., and L.M. contributed new reagents/analytic tools; L.L., J.W., F.T., and L.M. analyzed data; and L.L., J.W., F.T., and L.M. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

[1]To whom correspondence may be addressed. Email: lmahadev@g.harvard.edu.

This article contains supporting information online at https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2519032123/-/DCSupplemental.

Published February 13, 2026.

insects that work as superorganisms (7, 18, 19). In physics, recent work on active particles with noisy trajectories and independent random goals (12) has shown that they reach goals most efficiently when the noise level is just above "the edge of jamming." Motivated by these and similar studies, we take interest in approximating system-level outcomes to help select optimal strategies for target attainment in crowded environments.

Our setting draws inspiration from decentralized robot collectives where task execution is possible without any central planner (20–23). A decentralized and noisy approach is relevant to agents which are hard to precisely control, like bacteria, insects, humans, microrobots, or robots in low-communication and low-visibility settings. Because the setting is highly stochastic, we are able to derive analytical results using probability theory.

Robot paths to the goal are entirely determined by local sensing and random noise. Each robot starts at a random location and receives a random goal to travel to. The robot always knows its goal's location and direction, and it homes in toward its goal in a random walk with rotational noise. Robots cease forward motion when a neighbor is detected within their forward sensing cones. The noisy walks allow robots to move stochastically around each other, and the swarm can be tuned by adjusting the swarm size and the rotational noise level. Noisy motion helps unjam traffic, but it also makes robots' paths indirect and inefficient. We measure a swarm's success using the goal attainment rate G, the average goals reached per second by all team members combined.

To address the question of what team size and noise level will maximize G, we begin with a qualitative understanding of how system behavior changes as we vary swarm size n and noise level σ. We use our observations to derive estimates for the length of a noisy walk, the jam escape time, the critical noise level, and other useful quantities in a theoretical setting. We assemble our derivations into an analytical approximation for G as a function of n. Numerical optimization of G yields an estimate for the optimal n and σ. Next, we implement the noisy navigation method on a swarm of physical robots. Finally, we compare our minimal, noisy navigation method to two more intelligent methods: a slightly more intelligent walk where noise is selectively applied only when robots encounter obstacles and a centralized planner with full global knowledge.

Agentic Model Simulations

In our noisy navigation setting, n agents $a_1, a_2, ..., a_n$ begin at random start positions and move toward random goal positions, which are all drawn independently and uniformly from a 2D square of side-length L (Fig. 1*A*). When an agent reaches within tolerance ϵ of its goal, it receives a new random goal. Arena boundary conditions vary as specified below (periodic boundary conditions simplify calculations for theoretical analysis, but are not applicable to physical experiments).

Agents move toward their goals in homing random walks, in which step lengths and heading angles are randomized. Animations of robot behavior are presented in Movie S1. The heading angle θ at each step is drawn from a Gaussian distribution $\mathcal{N}(\theta^*(t), \sigma^2)$, whose mean is the optimal direction θ^* toward the goal at the beginning of the step (Fig. 1*B*). (We note that the correctly invariant distribution for periodic variables is the von Mises-Fisher distribution (24). For small variances, the difference between the von Mises-Fisher and Gaussian distributions is small. We work with a wrapped Gaussian.) Random step lengths are drawn from a uniform distribution in the interval $(\frac{1}{2}b, \frac{3}{2}b)$, where the constant b represents average step length (if the agent does not stop for traffic during the step). Examples of homing walks with $\sigma = 0$ (no noise) and $\sigma = 0.75$ are shown in Fig. 1 *C* and *D*.

Each agent has a forward sensing cone, a $\gamma = 120°$ sector of a circle with radius r (Fig. 1*B*). When the sensing cone is unoccupied, agents move forward with speed v. When an agent senses a neighbor within the sensing cone, it ceases forward motion but can still rotate to the random angles drawn for new steps. The SD σ of the Gaussian angular noise is the same for all agents, and is the key parameter for tuning the noisiness of the team. The nondimensional parameters, which are also plotted in Fig. 2*C*, are $\frac{r}{v}\sigma \propto Pe_r^{-1}$ (the inverse Peclet number) and $\frac{r^2}{L^2}n \propto \rho$ (density). Jams resulting from simulation and experiment are shown in Fig. 1 *E* and *F*. If the robot turns to an angle where its vision cone is free again, it resumes moving forward at the new angle. When σ is higher and the motion is noisier, robots turn to a larger range of angles and are more likely to escape jams.

We begin by observing simulations to understand the typical behavior of teams with different values of n and σ. Simulation code is informed by the Stage simulator (25). Agents proceed to their goals until they encounter another robot and stop moving forward. When there is no noise to unjam collisions ($\sigma = 0$), robots remain jammed in small clusters (leftmost column of Fig. 2*A*). As σ increases (moving toward the *Right* in Fig. 2*A*), the noise allows robots to escape some jam configurations. "Easy" jams to escape usually contain few robots or have squiggly, irregular shapes; the persistent jams that remain are larger and rounder. As σ increases farther, the steady state eventually becomes a single large jam containing essentially all the agents. When σ is higher than a critical value $\sigma^*(n)$, the noise melts even this all-encompassing jam. Persistent jams no longer appear, and collisions are unjammed quickly by the high noise. We call the $\sigma < \sigma^*$ regime "the jammed phase," and the $\sigma > \sigma^*$ regime where large jams do not form "the dilute phase." The steady state goal attainment rate in the jammed phase is approximately 0 (*Left* sides of solid curves in Fig. 2*B*). In agreement with ref. 12, we observe that noise levels just above $\sigma^*(n)$ maximize goal attainment for a given n (peaks of solid curves in Fig. 2*B*). Intuitively, once σ is high enough to break up jams, adding more noise only reduces goal attainment by making travel paths more inefficient (*Right* sides of solid curves in Fig. 2*B*). The same information is presented as a heatmap in Fig. 2*C*, where the sharp transition between the jammed and dilute phases is clearly visible.

Analytic Approximations

With this qualitative understanding of the jammed and dilute phases, and the observation that goal attainment rate G peaks near the boundary between the two, we next develop quantitative approximations of steady-state behavior in both phases. We then use this model to predict the optimal settings of σ and n that maximize G.

We first approximate $G(n, \sigma)$ in the dilute phase. We then estimate the critical value of noise $\sigma^*(n)$ marking the edge of jamming for a team of n robots. The optimal goal attainment rate for a team of size n can then be estimated as $G(n, \sigma^*(n))$. Finally, we find the n maximizing $G(n, \sigma^*(n); L, r, \gamma, b, v)$ by numerical optimization. This gives the predicted optimal team size for an arena of given side length and robots with given geometric parameters.

Our primary assumptions are that the system is at a steady state, the arena has periodic boundary conditions, step size b is

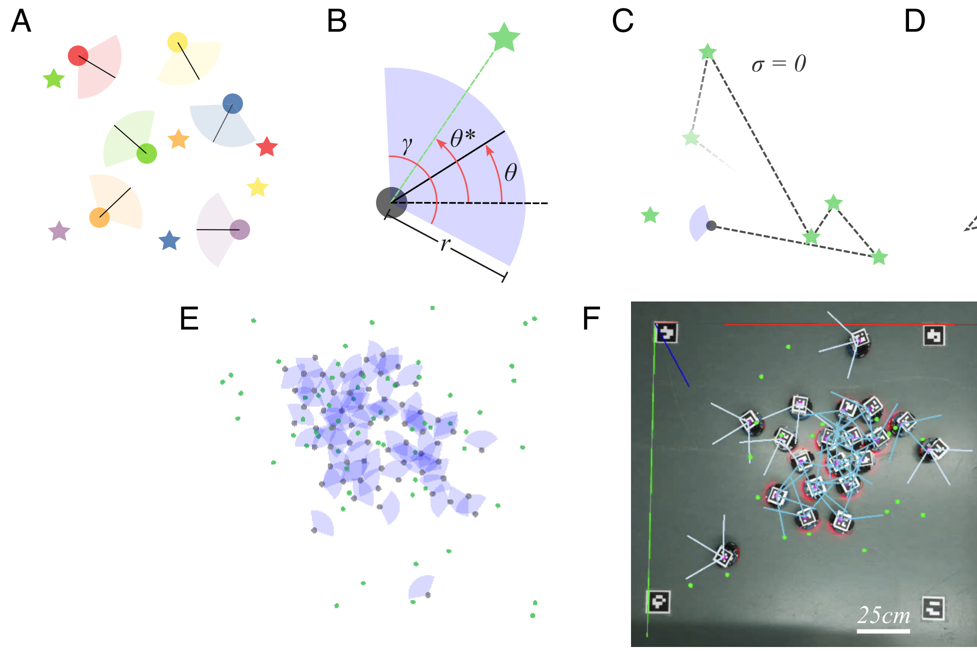

Fig. 1. Noisy swarm setting. (*A*) Robots begin in random positions and individually receive random goal positions (star shapes of matching color), which represent locations for dropping off deliveries or construction parts. See Movie S1 for animated examples of robot trajectories. (*B*) Robots cease forward motion, but can still rotate in place, when an occupant is detected in their sensing cones (blue). The sensing cone is a $\gamma = 120°$ sector of a circle with radius r. At time t, robots have heading $\theta(t)$ and an optimal travel angle $\theta^*(t)$ to the goal (green star). (*C*) A robot with no rotational noise moves directly from one goal position to the next. (*D*) A robot with rotational noise σ moves in steps. For each one, it draws the step's random length from $Unif(\frac{1}{2}b, \frac{3}{2}b)$ and its heading from the Gaussian distribution $\mathcal{N}(\theta^*(t), \sigma^2)$. (*E* and *F*) At high densities, visible traffic jams occur in simulation and experiment. Green dots mark robots' goals. In (*F*), which shows a view of our robotic experiments from the overhead camera, red lights indicate that a robot's vision cone is currently occupied. Robot positions, headings, goals, and sensing cone edges are annotated.

small, robots turn instantaneously, and (in the dilute regime where large jams do not form) all collisions are two-robot collisions. Derivation details, along with discussion of when and why approximation errors arise, are presented in *SI Appendix*, section S2.

Approximating *G* in the Dilute Regime. The average initial straight-line distance between a robot and its random goal position is (see *SI Appendix*, section S2.1, also in ref. 12)

$$\mathbb{E}(d) = \frac{\ln(3 + 2\sqrt{2}) + 2^{3/2}}{12}L \approx 0.3826L. \qquad [1]$$

Though adding noise to an agent's motion helps it escape jams, it also makes its path to the goal less efficient. To capture this tradeoff, we consider the *walk extension* $\mathcal{E}(\sigma)$, the average ratio between the start-to-goal length of a homing random walk with noise σ and a straight line. This quantity is also useful for polymer physics (26) and for other studies of noisy searching processes, such as animal foraging (27). In the limit of small b (*SI Appendix*, section S2.2),

$$\mathcal{E}(\sigma) \approx e^{\sigma^2/2}. \qquad [2]$$

In the absence of obstacles, the expected travel time t_0 to reach the goal is therefore

$$\mathbb{E}(t_0) \approx \frac{0.3826L \cdot e^{\sigma^2/2}}{v}. \qquad [3]$$

Next, we consider how many collisions C a robot in the dilute regime experiences en route to a goal.

We approximate the average collisions per goal as [robot density × steps to reach goal × vision cone area coverage per step] (*SI Appendix*, section S2.3):

$$\mathbb{E}(C) \approx \frac{n-1}{L^2} \times \frac{0.3826L\mathcal{E}(\sigma)}{b} \times \left[\frac{\gamma}{2\pi}\pi r^2 + 2\sqrt{2}br\sin(\frac{\gamma}{2})\right]. \qquad [4]$$

To estimate the time t_{jam} that robots spend jammed in each collision, we make the simplifying assumptions that all collisions are two-body and head-on, and that the jam ends when at least one robot turns enough to free its sensing cone (*SI Appendix*, Fig. S4). This occurs if a robot draws a random heading θ for which $|\theta - \theta^*| > \frac{\gamma}{2}$, the probability of which can be obtained by integrating over the Gaussian distribution of the rotational noise. Thus (*SI Appendix*, section S2.4)

$$\mathbb{E}(t_{\text{jam}}) \approx \frac{b/v}{1 - \text{erf}(\frac{\gamma}{2\sqrt{2}\sigma})^2} \qquad [5]$$

The total time $t_g(n, \sigma)$ for a robot to reach one goal is then given by the sum of the time spent traveling and the time spent in jams: $t_g(n, \sigma) \approx \mathbb{E}(t_0) + \mathbb{E}(C) \cdot \mathbb{E}(t_{\text{jam}})$. Therefore, in the dilute regime where $\sigma > \sigma^*(n)$, the total goal attainment rate $G(n, \sigma)$ for the robot team is

$$G(n, \sigma) \approx \frac{n}{t_g(n, \sigma)} \qquad [6]$$

Approximating Critical Noise Level σ^*. To solve for the noise level σ^* at the edge of jamming for a given n, we approximate the

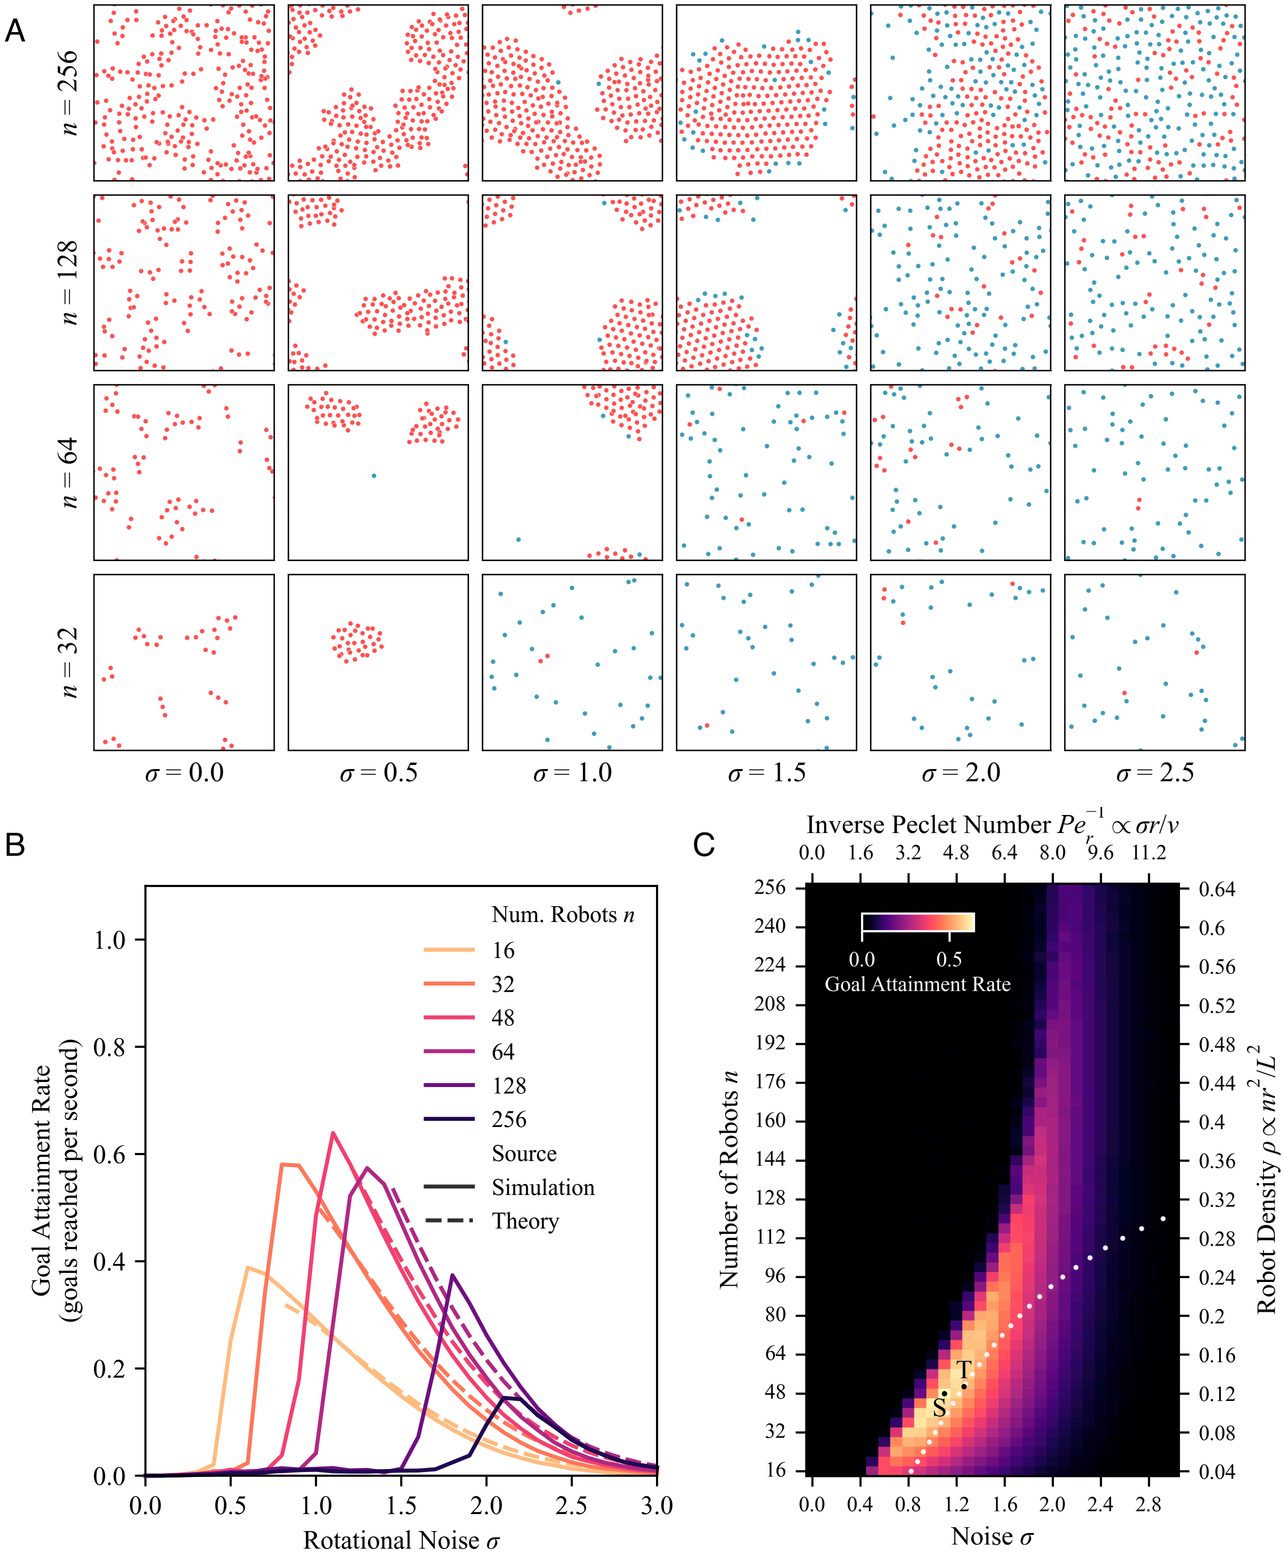

Fig. 2. Simulations and theory. (*A*) Jammed and dilute regimes. Dots show robot positions for various (n, σ) after 8,000 s of simulated time. Color indicates whether robots are blocked (red) or free to move (blue). All subplots in this figure use periodic boundary conditions, $L = 40$, $r = 2$, $\gamma = \frac{2\pi}{3}$, $v = 0.5$, $b = 0.5$, $\epsilon = 0.6$. See *SI Appendix*, Table S1 for more parameter details, and Movie S2 for an animated version of this phase diagram. (*B*) Solid lines plot simulation data for goal attainment rate G, the total goals reached per second by the entire swarm. Dashed lines show our theoretical approximations of G for the dilute regime. The dashed lines end where our theory predicts the dilute regime ends. (The absence of dashed lines for $n \in \{128, 256\}$ reflects Eq. **9** diverging at high densities; the theory remains valid in the regime of interest where the optimal parameters lie.) Goal attainment rate is computed over the last 25% of the simulation time to capture steady state behavior. Simulations are run for 50 trials in the high-variance regime ($n \leq 128$ and $\sigma \leq 2.0$), and for 20 trials otherwise. (*C*) Goal attainment across (n, σ) space. The "S" dot marks the optimal (n, σ) observed in simulations, the "T" dot marks the optimal (n, σ) predicted by theory, and white dots mark the predicted optimal σ predicted by theory for each n.

expected time it takes for a robot to enter a jam from a random location, and the time to exit the jam. At the edge of jamming, these times will be equal.

Based on Fig. 2*A*, we make the approximation that in the jammed regime at the edge of jamming, all robots except one are in one large circular jam. We approximate the robots as circles of radius $r/2$ and area $\pi r^2/4$; therefore the large jam has area $A_j \approx n\pi r^2/4$ and radius $r_j \approx r\sqrt{n}/2$.

We approximate the probability that a robot encounters the jam while traveling to a goal as $A_j/L^2 = n\pi r^2/4L^2$. Thus the robot will reach about $4L^2/n\pi r^2$ goals before getting stuck in the jam. Using Eq. **3**, the time this takes is (*SI Appendix*, section S2.5)

$$\mathbb{E}(\text{jam entry time}) \approx \frac{1}{v}\frac{4 \cdot 0.3826 \cdot L^3 \cdot \mathcal{E}(\sigma)}{n\pi r^2} \qquad [7]$$

To estimate the jam escape time, we assume: 1) in an average case, the robot needs to travel 1/4 of the jam perimeter to escape the jam; 2) the robot is in an "average" position (shown in *SI Appendix*, Fig. S5) where it needs to turn $\frac{\pi}{4} + \frac{\gamma}{2}$ to free its sensing cone, which occurs with probability $\frac{1}{2} - \frac{1}{2}\operatorname{erf}(\frac{x}{\sqrt{2}\sigma})$ per update step; 3) the 1/4-perimeter distance the robot needs to travel is scaled by both $\mathcal{E}(\sigma)$ and $\frac{1}{\frac{1}{2}-\frac{1}{2}\operatorname{erf}(\frac{x}{\sqrt{2}\sigma})}$. This overestimates the difficulty of exiting the jam, so the final calculated σ^* will be an overestimate (*SI Appendix*, section S2.6). The jam exit time can then be approximated as

$$\mathbb{E}(\text{jam exit time}) \approx \frac{1}{2}\frac{\pi r\sqrt{n}\mathcal{E}(\sigma)}{v(1-\operatorname{erf}((\frac{\pi}{4}+\frac{\gamma}{2})\frac{1}{\sqrt{2}\sigma}))}. \qquad [8]$$

By setting Eqs. **7** and **8** equal, we solve for the optimal noise level

$$\sigma^*(N) \approx (\frac{\pi}{4}+\frac{\gamma}{2})\frac{1}{\sqrt{2}}(\operatorname{erf}^{-1}(1-\frac{1}{8}\frac{\pi^2 r^3 n^{3/2}}{0.3826 L^3}))^{-1} \qquad [9]$$

For a given value of n, the optimal goal attainment rate is then $G(n, \sigma^*(n)) = n/t_g(n, \sigma^*(n))$. To find the optimal set of values (n, σ) that maximizes goal attainment rate overall, we numerically maximize $G(n, \sigma^*(n))$ as a function of n.

Approximation Accuracy. The parameters selected are close to the optimal ones observed in simulations, as shown in Fig. 2*C*. In Fig. 2 *B* and *C*, we see that predictions for attainment rate and $\sigma^*(n)$ diverge as n grows large. This aligns with several assumptions we made that break down at high n or σ. For example, we assumed that all collisions are two-body collisions; the multibody interactions that we neglected grow more frequent at higher densities. Using this assumption to obtain analytical results is the current state of the art. It is also employed, for example, in ref. 9. See *SI Appendix*, section S2 for detailed discussion of error sources.

For further testing, *SI Appendix*, Fig. S6 compares our predicted optimal (n, σ) with the optimal parameters observed in additional simulation settings varying one of r, γ, or b. The approximations are accurate except when the optimal n is large (in settings that make the sensing cone smaller, so more agents can effectively share the space). This again reflects the approximation's divergence at high densities.

Robotic Experiments

The results we presented so far use ideal assumptions, such as the sensing cone being fully reliable and the robot speed being fully controllable. To evaluate the behavior of physically instantiated systems, we repeat the above experiments on teams of Alphabot2-Pi robots matching those used in ref. 28 to study single-file motion. The physical experiments investigate system behavior under more realistic conditions, in contrast to idealized assumptions used by the theory and simulations, such as periodic boundary conditions and small step sizes. For the relatively simple Alphabots, robot speed is neither constant nor homogeneous, and sensing neither perfect nor instantaneous.

We couple the robot experiments with a new set of simulations whose parameters and conditions match those of the physical experiments. While robots are programmed to have $r = 0.2$ m and $\gamma = \frac{2\pi}{3}$, in practice they stop at a closer distance than 0.2m due to communication lag. We compute an "effective radius" ($r = 0.156$ m) based on experimental data to use as the sensing range in the matching simulations. Similarly, we use experimental data to estimate robots' effective linear and turning speeds for each (n, σ) setting, and we use the corresponding values in the simulations. See *SI Appendix*, section S3 for more experimental details.

The physical and simulated experiments show qualitative and quantitative agreement. Qualitatively, goal attainment rate curves are nonmonotonic in both density and noise level, with peaks moving to higher noise levels as density increases (Fig. 3*B*). Quantitatively, the simulations and experiments agree on the σ where goal attainment peaks (*SI Appendix*, Fig. S7).

Adding Intelligent Behavior to Agents

Having seen what we can achieve in the absence of intelligence, a natural question is how much we can gain with more intelligent

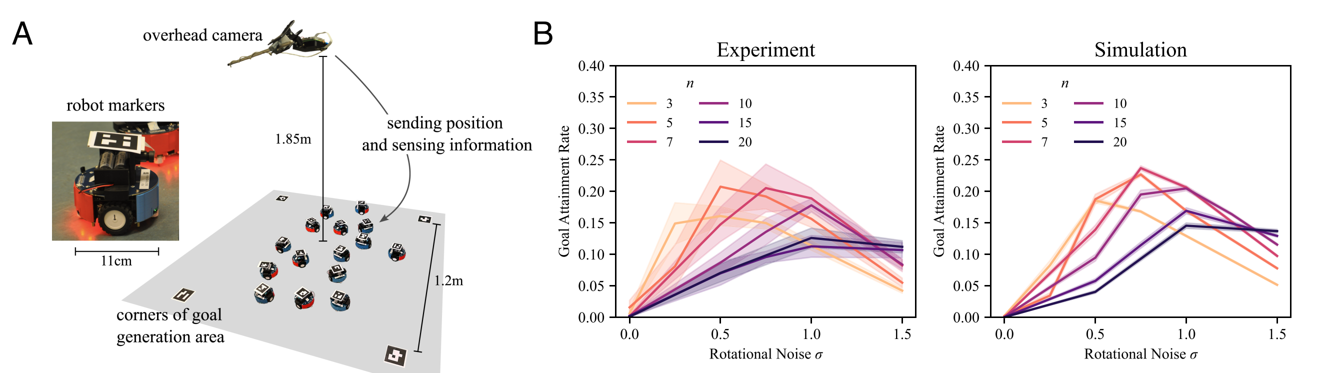

Fig. 3. Robotic experiments. (*A*) Individual robot and schematic of experimental setup. (*B*) Goal attainment rate for experiments vs. simulations with matching parameters. Experimental data have 6 trials per point; simulation data have 100 trials per point. Data include the last 75% of 300-s trials. Shaded error bands show SE. See *SI Appendix*, Table S1 for parameter details and Movie S3 for experimental footage.

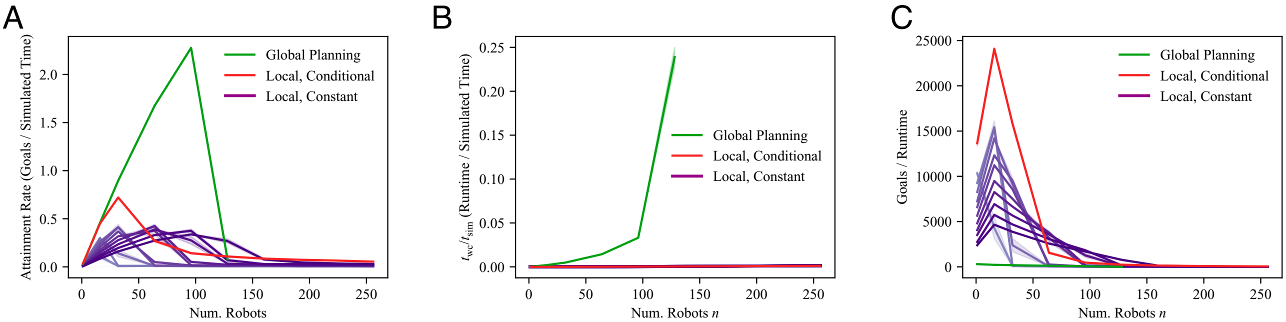

Fig. 4. More intelligent controllers. In all subplots, the purple lines for local, constant noise represent different values of σ ranging from 0.5 to 1.5. Shaded error bands show SE. Data contain 20 trials per point. See *SI Appendix*, Table S1 for more parameter details. (*A*) Comparing the attainment rate of different kinds of noisy navigation controllers. See Movie S4 for animated examples of the different controllers. (*B*) Comparing the ratio of runtime to simulated time for local and global controllers. (*C*) Comparing the goals reached per wall-clock time of local and global controllers.

agents. We compare the noisy motion controller with constant Gaussian noise against two more intelligent control methods. The first method is noisy motion with conditional noise, another distributed approach where robots react to local sensing. Robots using conditional noise move directly toward the goal unless they detect a neighbor in their sensing cone. When a robot's sensing cone is occupied, it draws a uniform random angle for its next step. Conditional noise is still distributed and depends entirely on local sensing, but adds "one bit" of intelligence. The second method we test is Cooperative A* (14), a greedy global planning algorithm for finding noncolliding paths on a graph. In the graph, each node represents occupying a cell at a certain time. The algorithm looks for agent paths to their goals in individual, sequential searches; Cooperative A* does not perform a joint search that returns all agents' paths at once. See *SI Appendix*, section S4 for implementation details.

We see in Fig. 4*A* that the simple Gaussian noise controller studied earlier reaches about half the attainment of the conditional noise method. In turn, the conditional noise method reaches about 30% of the global A* planner's maximum attainment. At low densities, the attainment of the conditional noise tracks closely with that of the global planner. However, the global planner is computationally expensive, and the wall-clock time t_{wc} needed to simulate an in-simulation time t_{sim} increases rapidly, which can be seen by plotting $t_{\mathrm{wc}}/t_{\mathrm{sim}}$ in Fig. 4*B*. All simulations are run on a 2021 M1 Macbook Pro. Because the local sensing methods are simulated with finer timesteps, their runtimes are in fact overestimated compared to the global planner. Furthermore, the local mechanism is inherently parallelizable while the global planner relies on sequential planning. To assess how efficiently the different controllers utilize computation, in Fig. 4*C*, we plot goals reached per t_{wc}. The local planner with conditional noise reaches the highest goal attainment per unit of computation: 80 times that of the maximum attained with A* global planning and 1.6 times the maximum attained with constant Gaussian noise.

Discussion

Inspired by the problem of steric constraints that can hinder task execution by agents, we studied how noisy motion can help simple agents reach individual goal locations in a crowded environment. Our integrative approach combines theory, simulation, and experiments to corroborate key system behaviors from multiple perspectives.

We found that a noisy local approach attained high efficiency relative to how much computation it used. Furthermore, the stochasticity of the Gaussian noise controller allowed us to perform approximate optimizations over global parameters (agent density and noise level). This suggests that looking for other systems where stochasticity produces predictable steady-state behavior may help us optimize the parameters for more sophisticated multibody collectives. A natural further direction is to study how maximum swarm capability increases as intelligence and computational complexity is added to a collective. How much complexity can be shed without losing significant performance?

Our robotics setting shares some parallels with the dynamics of active matter (29), especially in the context of an effective field theory in the hydrodynamic limit associated with long wavelength, slow dynamics of the agent density and its variations in space-time. Natural variables for this are robot density and speed, which are both functions of space and time. Directed motion toward a target can be described with a nonlocal gradient term that serves to orient robots toward their individual targets, subject to keeping the total number of robots constant. The jam formation process can potentially be modeled with a continuum framework to describe the likelihood of different sized jams appearing, while jam dissolution is a function of diffusive processes at play. Finally, adding intelligence is tantamount to including additional nonlocal terms that describe the locations (or probability distributions) of agents. Working toward such a theory is likely to involve generalizations of paradigms such as the formation and breakup of traffic jams (1, 2) and motility-induced phase separation (3), but with additional terms arising from global considerations, questions worthy of future study.

Materials and Methods

The Alphabot robots used in experiments are made by Waveshare Electronics. They are circular two-wheeled robots measuring 11 cm in diameter and controlled by Raspberry Pis. Their random goals are drawn from a square with side length $L = 1.2$m (no periodic boundary conditions). To provide the global positioning available in the simulation, we mount a unique ArUco marker (31) on top of each robot and in each corner of the goal generation area. An overhead Raspberry Pi with a camera, positioned 1.85 m above the arena as shown in Fig. 3*A*, continuously captures images of the robots. It extracts position and sensing cone information (Fig. 1*E*) using OpenCV (32). Robots receive their current positions, headings, and sensing cone occupancy status from the overhead camera via MQTT messaging (33) at $\sim$7 updates per second. Robots use the information about their own positions to identify the direction to their goals, and follow the same noisy motion protocol described in earlier sections. Experiment trials are five minutes long.

Data, Materials, and Software Availability. Simulation code and data from the robotic experiments is available at https://doi.org/10.5281/zenodo.17954636 (30).

ACKNOWLEDGMENTS. We thank Estevao Pereira for help with the design of the robots. L.L. thanks the NSF for financial support, F.T. thanks the Eindhoven Artificial Intelligence Systems Institute for partial financial support, and L.M. thanks the Simons Foundation and the Henri Seydoux Fund for partial financial support. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE 2140743. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

1. M. J. Lighthill, G. B. Whitham, On kinematic waves II. A theory of traffic flow on long crowded roads. *Proc. R. Soc. Lond. Ser. A Math. Phys. Sci.* **229**, 317–345 (1955).
2. G. F. Newell, A simplified theory of kinematic waves in highway traffic, part I: General theory. *Trans. Res. Part B Methodol.* **27**, 281–287 (1993).
3. M. E. Cates, J. Tailleur, Motility-induced phase separation. *Annu. Rev. Condens. Matter Phys.* **6**, 219–244 (2015).
4. L. S. Marcolino, L. Chaimowicz, "Traffic control for a swarm of robots: Avoiding target congestion" in *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems* (IEEE, 2009), pp. 1955–1961.
5. L. S. Marcolino *et al*., Avoiding target congestion on the navigation of robotic swarms. *Auton. Robot.* **41**, 1297–1320 (2017).
6. L. A. Poissonnier, S. Motsch, J. Gautrais, J. Buhl, A. Dussutour, Experimental investigation of ant traffic under crowded conditions. *eLife* **8**, e48945 (2019).
7. I. D. Couzin, N. R. Franks, Self-organized lane formation and optimized traffic flow in army ants. *Proc. R. Soc. Lond. B Biol. Sci.* **270**, 139–146 (2003).
8. A. Corbetta, J. Meeusen, C. M. Lee, F. Toschi, Continuous measurements of real-life bidirectional pedestrian flows on a wide walkway. arXiv [Preprint] (2016). http://arxiv.org/abs/1607.02897.
9. K. A. Bacik, B. S. Bacik, T. Rogers, Lane nucleation in complex active flows. *Science* **379**, 923–928 (2023).
10. P. Iyer, R. S. Negi, A. Schadschneider, G. Gompper, Directed motion of cognitive active agents in a crowded three-way intersection. *Commun. Phys.* **7**, 379 (2024).
11. R. S. Negi, R. G. Winkler, G. Gompper, Emergent collective behavior of active Brownian particles with visual perception. *Soft Matter* **18**, 6167–6178 (2022).
12. M. Casiulis, D. Hexner, D. Levine, Self-propulsion and self-navigation: Activity is a precursor to jamming. *Phys. Rev. E* **104**, 064614 (2021).
13. J. Yu, S. LaValle, "Structure and intractability of optimal multi-robot path planning on graphs" in *Proceedings of the AAAI Conference on Artificial Intelligence* (2013), vol. 27, pp. 1443–1449.
14. D. Silver, "Cooperative pathfinding" in *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment* (2005), vol. 1, pp. 117–122.
15. T. Standley, "Finding optimal solutions to cooperative pathfinding problems" in *Proceedings of the AAAI Conference on Artificial Intelligence* (2010), vol. 24, pp. 173–178.
16. H. Ma, J. Li, T. S. Kumar, S. Koenig, "Lifelong multi-agent path finding for online pickup and delivery tasks" in *Proceedings of the 16th Conference on Autonomous Agents and Multiagent Systems* (2017), pp. 837–845.
17. J. Li *et al*., "Lifelong multi-agent path finding in large-scale warehouses" in *Proceedings of the AAAI Conference on Artificial Intelligence* (2021), vol. 35, pp. 11272–11281.
18. O. Peleg, J. M. Peters, M. K. Salcedo, L. Mahadevan, Collective mechanical adaptation of honeybee swarms. *Nat. Phys.* **14**, 1193–1198 (2018).
19. A. Heyde, L. Guo, C. Jost, G. Theraulaz, L. Mahadevan, Self-organized biotectonics of termite nests. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2006985118 (2021).
20. L. Giomi, N. Hawley-Weld, L. Mahadevan, Swarming, swirling and stasis in sequestered bristle-bots. *Proc. R. Soc. A Math. Phys. Eng. Sci.* **469**, 20120637 (2013).
21. S. G. Prasath *et al*., Dynamics of cooperative excavation in ant and robot collectives. *eLife* **11**, e79638 (2022).
22. M. Rubenstein, A. Cornejo, R. Nagpal, Programmable self-assembly in a thousand-robot swarm. *Science* **345**, 795–799 (2014).
23. J. Werfel, K. Petersen, R. Nagpal, Designing collective behavior in a termite-inspired robot construction team. *Science* **343**, 754–758 (2014).
24. R. A. Fisher, Dispersion on a sphere. *Proc. R. Soc. Lond. Ser. A Math. Phys. Sci.* **217**, 295–305 (1953).
25. R. Vaughan, Massively multi-robot simulation in stage. *Swarm Intell.* **2**, 189–208 (2008).
26. A. Marantan, L. Mahadevan, Mechanics and statistics of the worm-like chain. *Am. J. Phys.* **86**, 86–94 (2018).
27. S. Hartman, S. D. Ryan, B. R. Karamched, Walk this way: Modeling foraging ant dynamics in multiple food source environments. *J. Math. Biol.* **89**, 41 (2024).
28. L. Alonso-Llanes, A. Garcimartín, I. Zuriguel, Single file motion of robot swarms. *Phys. Rev. Res.* **6**, L022037 (2024).
29. J. Tailleur, G. Gompper, M. C. Marchetti, J. M. Yeomans, C. Salomon, *Active Matter and Nonequilibrium Statistical Physics: Lecture Notes of the Les Houches Summer School: Volume 112, September 2018* (Oxford University Press, 2022), vol. 112.
30. L. Liu, lliu12/ministage: v1.0.1 - documentation improvements (2025).
31. S. Garrido-Jurado, R. Muñoz-Salinas, F. J. Madrid-Cuevas, M. J. Marín-Jiménez, Automatic generation and detection of highly reliable fiducial markers under occlusion. *Pattern Recognit.* **47**, 2280–2292 (2014).
32. G. Bradski, The OpenCV library. *Dr. Dobb's J. Softw. Tools* **120**, 122–125 (2000).
33. R. A. Light, Mosquitto: Server and client implementation of the MQTT protocol. *J. Open Sour. Softw.* **2**, 265 (2017).