Autonomous Robots (2023) 47:963–977
https://doi.org/10.1007/s10514-023-10121-9

An empirical characterization of ODE models of swarm behaviors in common foraging scenarios

John Harwell[1] · Angel Sylvester[1] · Maria Gini[1]

Received: 14 January 2022 / Accepted: 16 June 2023 / Published online: 23 July 2023

Abstract
There is a large class of real-world problems, such as warehouse transport, at different scales, swarm densities, etc., that can be characterized as Central Place Foraging Problems (CPFPs). We contribute to swarm engineering by designing an Ordinary Differential Equation (ODE) model that strives to capture the underlying behavioral dynamics of the CPFP in these application areas. Our simulation results show that a hybrid ODE modeling approach combining analytic parameter calculations and post-hoc (i.e., after running experiments) parameter fitting can be just as effective as a purely post-hoc approach to computing parameters via simulations, while requiring less tuning and iterative refinement. This makes it easier to design systems with provable bounds on behavior. Additionally, the resulting model parameters are more understandable because their values can be traced back to problem features, such as system size, robot control algorithm, etc. Finally, we perform real-robot experiments to further understand the limits of our model from an engineering standpoint.

Keywords Swarm robotics · Foraging · ODE model · Diffusion

1 Introduction

Swarm Robotics (SR) is the study of the coordination of large numbers of simple robots (Şahin, 2005). SR systems can be homogeneous, i.e., single robot type and identical control software, or heterogeneous, i.e., multiple robot types and control software (Ramachandran et al., 2020). The main differences between SR systems and multi-robot systems stem from the mechanisms on which SR systems are based. Historically, those were principles of biological mimicry or problem solving techniques inspired by natural systems of agents such as bees, ants, and termites (Labella and Dorigo, 2006), though modern SR systems typically incorporate mathematically rigorous elements of multi-robot system design (Castello et al., 2016).

In this work, we study the Central Place Foraging Problem (CPFP) in which robots gather objects (blocks) across a finite operating arena and bring them to a central location (nest) under various conditions and constraints. Foraging is one of the most studied applications of SR, due to its straightforward mapping to real-world applications (Hecker and Moses, 2015); for an extensive discussion of the state of the art, see (Lu et al., 2020). The complexity of the foraging task frequently gives rise to behavioral dynamics caused by inter-robot interactions, in which the *events* robots experience, such as encountering other robots and objects, do not form a homogeneous Poisson process; i.e., non-negligible higher moments are present. We seek to understand—from an engineering perspective—the *practical* limits of using Poisson-based models to model foraging swarms. That is, we do not seek to derive a generalized model that encompasses *every* possible variant of the foraging problem, but rather to derive a practical model which is accurate on scenarios with high real-world utility; e.g., those that map to important problems such as warehouse transport or supply chain logistics in large outdoor spaces. Specifically, we are interested in understanding under what conditions behavioral dynamics in the foraging problem can and cannot be captured by using ODE modeling, regardless of whether the actual characteristics of the dynamics are strictly amenable to it.

✉ John Harwell
harwe006@umn.edu

Angel Sylvester
sylve057@umn.edu

Maria Gini
gini@umn.edu

[1] Department of Computer Science and Engineering, University of Minnesota, 200 Union St SE, Minneapolis, MN 55455, USA

Our results, using a hybrid ODE model we developed, show that behavioral dynamics can be captured in many scenarios across scales, swarm densities, and other parameters, demonstrating the utility of ODE modeling as an engineering tool on many important problems.

2 Motivation and related work

Many SR systems have been designed around imitating natural systems, and use heuristic decision making (Castello et al., 2016) rather than combining natural principles with a mathematical grounding (Talamali et al., 2020). Nevertheless, heuristic approaches to swarm control have been effective for robots that operate with incomplete information and limited computing power. Researchers using this approach average large numbers of simulation runs to develop accurate models of swarm behavior and obtain empirical insights into real-world problems (e.g., (Harwell and Gini, 2019; Castello et al., 2014; Matthey et al., 2009)). Alternatively, collective swarm behaviors can be characterized by mathematically deconstructing robot control algorithms; this is more difficult, but provides the means to more precisely predict behavior *a priori–without* the need of repetitive experiments.

Under the *mean-field approximation*, a system of N robots can be replaced by a single robot that is placed in an external field equivalent to the average effect of the other robots. The canonical question when using this simplifying approximation is how large N must be for it to be valid. The mean-field approximation has been used to model SR systems as asynchronous and event-driven (Mather et al., 2013; Hsieh et al., 2008; Lerman et al., 2006). Robots experience events in an interval t according to a Poisson distribution:

$$P(k \text{ events in interval } t) = e^{-\lambda}\frac{(\lambda t)^k}{k!} \tag{1}$$

where λ is the rate parameter, which is the mean. The validity of this approach is predicated on (a) all robot events are independent, (b) the average rate of events per t is constant and the same for all robots, and (c) two robots cannot experience events at the same time. When these conditions hold, Eq. (1) captures the first moment (λ) of the system by applying the master equation used in statistical physics (Kampen, 2007).

Within this paradigm, macro- and microscopic models for group dynamics and individual behavior over time are used to model the *change* in the average behavior of the swarm, which is easier than modeling the average behavior of the swarm directly (Lerman and Galstyan, 2002; Berman et al., 2007; Galstyan et al., 2005; Sugawara and Sano, 1997). In the microscopic model, discrete difference equations are used to model the stochastic transitions between robot states and state transitions for all robots. In the macroscopic model, the swarm is conceptualized as a differentiable, continuous quantity, and its dynamics modeled with a set of ODEs whose variables are the population fractions associated with the different roles. Using this ODE modeling, the *average* behavior of the swarm in the steady state can be recovered (Berman et al., 2007).

Extensions to this line of research to capture higher order moments of behavioral dynamics include ensemble time-delay filtering (Mather et al., 2013; Hsieh et al., 2008), and explicit ensemble modeling (Mather and Hsieh, 2012) inspired by work in polynomial stochastic hybrid systems (Klavins, 2010). For a further discussion of the state of the art and current challenges in mean-field modeling, see (Elamvazhuthi and Berman, 2019).

In the exhibited *linear* collective behaviors of an SR system, the behavior of N robots is a linear function of the behavior of 1 robot; this is not true for systems exhibiting *non-linear* collective behaviors. It has been established that above a certain swarm density ρ_S (ratio of swarm size to arena size) (Sugawara and Sano, 1997; Hamann, 2013), interactions between robots can produce a non-linear mapping from the behaviors of individual robots to the observed collective behavior. This is because density governs the frequency at which interactions between robots occur (Cotsaftis, 2009; Georgé and Gleizes, 2005; Hunt, 2020; De Wolf and Holvoet, 2005). The exact ρ_S at which a given linear model of swarm behavior breaks down is influenced by many factors, including the control algorithm of the robots, the number of robots, and characteristics of the problem being solved, so in general it cannot be determined *a priori*.

Clearly, an interplay exists between (a) the swarm size N, (b) the spatial distribution of ρ_S, and (c) the characteristics of the problem being modeled (e.g., the number of objects and their distribution). Thus, we investigate swarms of both high and low ρ_S *and* both large and small N. We are interested in understanding the limits of ODE modeling from an engineering, rather than theoretical perspective. The scope of our investigation and what we will be able to infer about any results is defined by the following considerations.

First, because we are interested in a *practical*, empirical characterization, we employ homogeneous Poisson processes in our ODE modeling. As a result, *any* ODE model we create cannot directly capture the variability of the robot distribution beyond the mean. Thus, our chosen ODE modeling technique may not be sufficient if higher moments are present in collective behavioral distributions; infinitesimal generators or other modeling approaches could be employed in that case (Krishnarajah et al., 2005; Mather and Hsieh, 2012; Klavins, 2010) as an extension.

Second, in small swarms (low N, any ρ_S) fluctuations from *finite size effects* can be of order N, resulting in inaccurate ODE models because the requirements are not met (e.g, the mean-field approximation is not valid or large

higher moments are present in behavioral distributions). The master equation (Kampen, 2007) can be used to calculate those effects, but such computations are often intractable or algebraically difficult. Thus, our analytical expressions for model parameters to approximate these higher moments are effectively empirical approximations from a theoretical perspective; however, sufficiently accurate approximations will still provide substantial utility.

Third, we must consider sparse swarms (low ρ_S, any N). Recent results on the challenges of moving swarms into the real world argue that research towards sparse swarms is critical (Tarapore et al., 2020). In sparse foraging problems, each robot is responsible for a large area on the order of 100m^2, which could be indoor warehouses ($\approx 64\text{m}^2$), outdoor search and rescue, precision agriculture, or field monitoring ($600000\text{m}^2 \approx 150$ acres).

Fourth, in CPFPs ρ_S is rarely uniform, and depends on the object distributions and arena geometry, often resulting in behavioral distributions with non-negligible higher moments. This difficulty persists even with non-linear ODEs (Lerman et al., 2001).

Models in the ODE paradigm operate on both the forward problem, i.e., predicting collective behavior from features of the control algorithm of each robot (Lerman and Galstyan, 2002), and the inverse problem, i.e., incorporating design constraints into algorithm design in order to produce a desired collective behavior (Berman et al., 2007; Hsieh et al., 2008). However, the models have some limitations. First, they make simplifying assumptions such as homogeneous robot spatial distributions and environments (e.g., no obstacles, a completely visible arena), and Markov/semi-Markov scenarios (Berman et al., 2007). Second, to the best of our knowledge they are generally only evaluated within the context of a specific problem variant (e.g., a single spatial object distribution for a foraging task when the mean-field assumption holds), so their applicability to other problems is unclear. Nevertheless, many notable applications of various forms of ODE models appear in the literature, demonstrating their practical utility. These include the stick pulling experiment (Ijspeert et al., 2001), foraging of green/red pucks using agent memory (Lerman and Galstyan, 2003), the "house hunting model" (Hsieh et al., 2008; Berman et al., 2007), and ant-inspired models that collaborate with or without communication (Sugawara and Sano, 1997).

3 Background

We use the individual robot Finite State Machine (FSM) shown in Fig. 1, identical to previous work (Lerman et al., 2001; Lerman and Galstyan, 2002). Each state maps directly to a single or a set of robot behaviors that together make up the robot controller for executing a foraging task. It is a coarse-grained model of robot behavior, which omits controller details such as sensing and actuation and contains the *minimum* number of states needed to describe the system dynamics. Additional states can be added if necessary (Lerman and Galstyan, 2002).

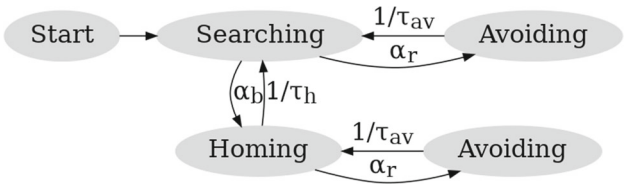

Fig. 1 FSM state diagram for a single robot. The *Avoiding* state is duplicated to uniquely identify the context of collision avoidance: avoidance while *Homing* or avoidance while *Searching*. Transition rates and notation are described in Table 1. Note that the inverse of the amount of time a robot spends in a given state (e.g., τ_h for *Homing*) is the rate of robots leaving the state

Table 1 Summary of ODE model components.

Term	Description
α_r	Robot encounter rate of robots *anywhere* in the arena
$\alpha_{r'}\hat{}$	Robot encounter rate *near* the nest
$\alpha_r\hat{}$	Robot encounter rate of robots *far away* from the nest
α_b	Robot encounter rate of blocks
τ_h	Mean robot homing time
τ_{av}	Mean robot time spent avoiding collision, per occurrence
N	Swarm size
ρ_S	Swarm density; ratio of N to operating area
$N_h(t)$	Mean # of robots returning to nest with blocks
$N_s(t)$	Mean # of robots searching for blocks
$N_{av}^h(t)$	Mean # of robots avoiding collision while homing
$N_{av}^s(t)$	Mean # of robots avoiding collision while searching
$B(t)\hat{}$	Mean # of blocks in the arena
$B_j(t)*$	Mean # of blocks in area j of the arena

Components with a $\hat{}$ are only in previous work, the one with a * is only in our model

Each of the states in Fig. 1 maps directly to an ODE describing this change, where the transition rates for states become the ODE terms. We refer the reader to (Lerman et al., 2001; Kampen, 2007) for a proof of the validity of this translation. The quantities modeled in Fig. 1 are listed in Table 1.

Our ODE model is inspired by, but distinct from previous work (Lerman et al., 2001; Lerman and Galstyan, 2002); for a discussion of crucial differences, see Sect. 4. We provide a brief summary of the previous model as follows. The change in the number of robots in the searching state, which decreases as those robots pickup blocks or encounter other

robots and switch to collision avoidance, is described by

$$\begin{aligned}\frac{dN_s(t)}{dt} = & -\alpha_b N_s(t)\left[B(t) - N_h(t) - N^h_{av}(t)\right] \\ & -\alpha_r N_s(t)\left[N_s(t) + N\right] \\ & + \frac{1}{\tau_h}N_h(t) + \frac{1}{\tau_{av}}N^s_{av}(t). \end{aligned} \tag{2}$$

A detailed explanation for the rate at which robots leave the searching state and switch to collision avoidance, $\alpha_r N_s(t)[N_s(t) + N]$, can be found in (Lerman et al., 2001). Finally, searching robots encounter other robots at rate α_r, regardless of the state of the other robot. Equation (2) increases as homing robots deposit blocks in the nest or as searching robots exit the collision avoidance state.

The change in the number of robots in the homing state, which increases as robots acquire blocks or leave the collision avoidance state, and decreases as robots enter the collision avoidance state or deposit their block in the nest, is modeled as

$$\begin{aligned}\frac{dN_h(t)}{dt} = & \alpha_b N_s(t)\left[B(t) - N_h(t) - N^h_{av}(t)\right] \\ & -\alpha_{r'} N_h(t)\left[N_h(t) + N\right] \\ & -\frac{1}{\tau_h}N_h(t) + \frac{1}{\tau_{av}}N^h_{av}(t). \end{aligned} \tag{3}$$

Eq. (3) assumes that the encounter rate for homing robots is different than for searching robots, since there will be more congestion near the nest, and uses a separate parameter $\alpha_{r'}$ to account for this. The change in the number of robots avoiding collision with other robots is

$$\frac{dN^s_{av}(t)}{dt} = \alpha_{r'} N_h(t)\left[N_h(t) + N\right] - \frac{1}{\tau_{av}}N^s_{av}(t), \tag{4}$$

and the change in the number of blocks available for robots to find is

$$\frac{dB(t)}{dt} = -\frac{1}{\tau_h}N_h(t). \tag{5}$$

Eq. (5) decreases whenever a searching robot acquires a block; once deposited in the nest, blocks are not re-distributed in the arena.

In addition to using ODEs to model collective behavior, we also draw on diffusion theory. In diffusing systems composed of homogeneous particles undergoing *normal diffusion*, the average particle displacement is proportional to the diffusion time. This linear relationship assumes the particle is moving in an infinite, structureless medium close to equilibrium. However, there are important cases in which the relationship between average particle displacement and time is non-linear, referred to as *anomalous diffusion* (Oliveira et al., 2019; Metzler et al., 2014; Vlahos et al., 2008). For instance, in biological systems, interactions with other particles could influence the diffusion for macro-proteins as they move through biological media (Santamaria-Holek and Vainstein, 2009; Weiss et al., 2003; Nicolau Jr et al., 2007). Anomalous subdiffusion arises due to crowding in a finite space, making the system heterogeneous and disordered (Ghosh et al., 2016).

4 Assumption analysis

We analyze the assumptions made in previous work vs. those we made, explicitly stating commonalities and noting differences to further motivate our model.

4.1 Common assumptions

1. N is sufficiently large so that S can be approximated using the mean field model.
2. S has reached steady-state; this maps naturally to swarms with long-running autonomy.
3. The population density ρ_S is relatively low (motivated by Tarapore et al. (2020)), so the behavior of N robots can be approximated by a linear function of the behavior of a single robot. Hence, the effect of a robot avoiding collision and encountering *another* robot during avoidance can be ignored.

None of these assumptions reduces the utility of foraging ODE models since they are consistent with the constraints of the CPFPs themselves (e.g., the need for *some* collision avoidance strategy near a common drop-off point for object transport). Furthermore, the results presented in Sect. 7 show that these assumptions, while helpful for derivation, are not essential for practical applications of the model.

4.2 Assumptions in previous work

Previous work made the following assumptions, which negatively impacted its broader utility:

1. The swarm is uniformly distributed in 2D space, with blocks scattered randomly. Such *Random* (RN) block distributions (Sugawara and Sano, 1997; Castello et al., 2016; Hecker and Moses, 2015; Lerman et al., 2001) (Fig. 2b), are appropriate in scenarios such as order fulfillment in a warehouse, but not in other scenarios. For example, when modeling evacuation of civilians from a disaster zone, the block distribution cannot be inferred *a priori*, and a *power law* (PL) distribution (Fig. 2a), with

blocks clustered in groups of various sizes, is more appropriate (Hecker and Moses, 2015; Harwell et al., 2020). Modeling the transfer of material from one side of a building to another location requires a *single source* (SS) (Fig. 2c) or a *dual source* (DS) (Fig. 2d) block distribution model, as in (Harwell and Gini, 2019; Ferrante et al., 2015; Pini et al., 2011; Harwell et al., 2020).
2. The number of objects to gather is finite; this conflicts with the steady-state assumption above.

4.3 Our assumptions

1. Robots use Correlated Random Walk (CRW) (Renshaw and Henderson, 1981; Harwell et al., 2020), which is a random walk where the direction of the next step is biased based on the direction of the previous step. The bias angle θ is drawn from a probability distribution $f(\theta)$ (Codling et al., 2010). Robots move at a mean velocity v_s, wandering until they acquire a block, which they then transport *directly* to the nest using phototaxis, i.e., motion in response to light, at a constant speed v_h. Robots have no memory.
2. Robots do not go to the exact nest center to drop their carried block. They choose a random point along their trajectory from where they enter the nest.
   This shortens the homing distance and can be computed as shown later in Sect. 5.1.
3. The number of objects to gather is infinite, i.e., when an object is collected a new one appears in the environment, resulting in steady-state foraging dynamics and collective behaviors.

We close this section with two important notes. First, the parameters of the model in Eqs. (2), (3), (4), (5) from previous work are free and must be determined post-hoc. Free parameters limit the model's reuse by requiring iterative parameter refinement via experimentation. Second, the assumptions made in the previous model and the assumptions we make are not compatible: infinite vs. finite number of objects, steady-state vs. not steady-state. This means that meaningful comparisons of the two models are not possible—see the supplementary material for videos illustrating why these differing assumptions result in different spatial distributions of robot behaviors and incompatible comparisons.

5 Generalized ODE foraging model

We now present our analytical ODE model. It relaxes the assumption of spatial uniformity in object distribution from (Lerman and Galstyan, 2002), and reduces the need

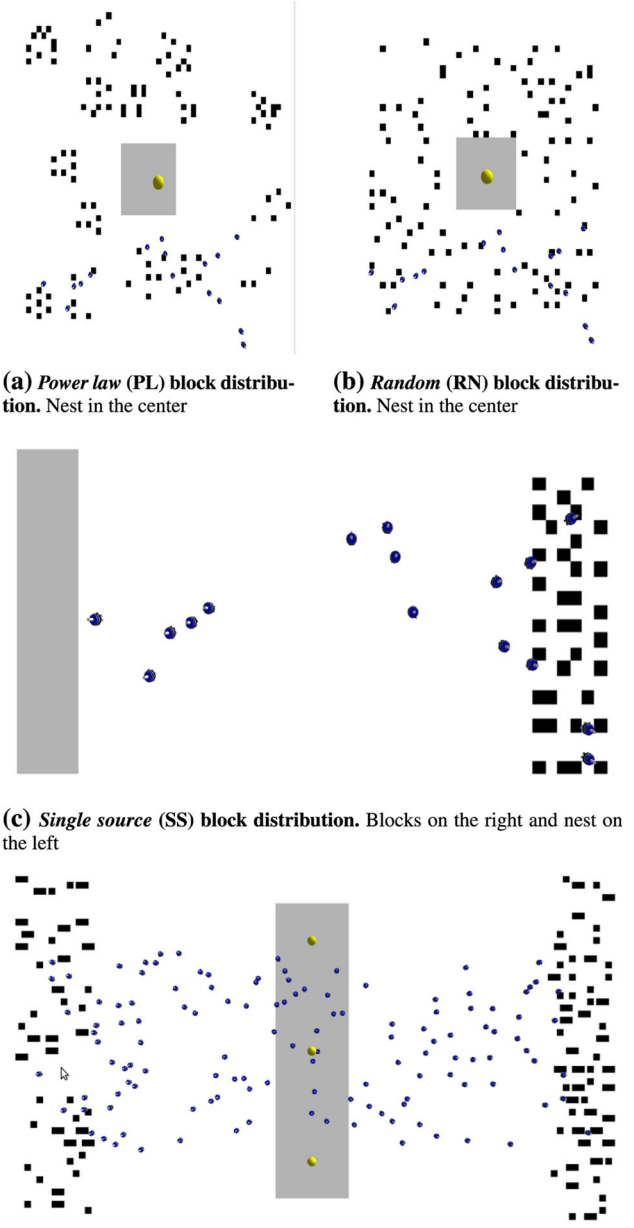

(a) ***Power law*** **(PL) block distribution.** Nest in the center

(b) ***Random*** **(RN) block distribution.** Nest in the center

(c) ***Single source*** **(SS) block distribution.** Blocks on the right and nest on the left

(d) ***Dual source*** **(DS) block distribution.** Blocks on the right and left and nest in the center

Fig. 2 Example of simulated foraging scenarios in ARGoS. Multiple robots (blue blobs), blocks to be collected (black squares), the nest (gray), and the lights robots use for phototaxis and localization (yellow) (Color figure online)

for experimental tuning, requiring only two free parameters: $C_{ca}(m)$, $C_h(m)$ (see Table 2). Together, these improvements greatly increase the usefulness of ODE modeling in foraging problems.

Let $M = \{SS, DS, PL, RN\}$ be the set of scenarios based on the block distributions shown in Fig. 2. In each $m \in M$, let the area where blocks can be distributed be a subset of the overall area A of the arena. Let $j = 1, \ldots, J$ be the sub-areas

Table 2 Summary of notation

Term	Description
S	The swarm of N robots
N	Total number of robots in the swarm
A	The operating arena S is deployed in, with dimensions $A_x \times A_y$
ρ_S	swarm density, given by the ratio of N to arena size (A)
$(A_j, \mathbf{c}_j, \rho_j)$	A sub-area j within A where blocks can be distributed. A_j is the area of j, $\mathbf{c}_j$ is the center of the sub-area, and ρ_j is the mean steady-state block density. There are J sub-areas within A
$\mathbf{x}$	An arbitrary (X,Y) location within A
$\mathbf{x}_n$	The center of the nest
$E[\mathbf{x}_{acq}]$	The expected location within A at which a robot will acquire a block
κ	The control algorithm that all robots run; in our case a Correlated Random Walk (CRW) which is a random walk where the direction of the next step is biased based on the direction of the previous step
v_h	Mean robot homing velocity
v_s	Mean robot searching velocity
θ	The bias angle governing each robot's CRW
$F(S)$	The diffusion constant of swarm S of N CRW robots
$C_{ca}(m)$	A post-hoc characterization constant for the inter-robot interference for scenario m
$C_h(m)$	A post-hoc characterization constant for environment heterogeneity (i.e., how dissimilar it is from an "ideal" scenario of a given type) for a scenario m

Additional terms are in Table 1

within A in which blocks can be distributed, each described by a tuple $(A_j, \mathbf{c}_j, \rho_j)$ (see Table 2). The area within A where blocks can be found is the union of these disjoint subsets: $A_d = \cup A_j$. The value of ρ_j varies across sub-areas in our model, so that the block encounter rate α_b can be captured accurately even in extreme non-homogeneous block distributions. We are now ready to introduce the equations for our ODE model:

$$\frac{dN_s(t)}{dt} = -\alpha_b - \alpha_r + \frac{1}{\tau_h}N_h(t) + \frac{1}{\tau_{av}}N^s_{av}(t), \tag{6}$$

$$\frac{dN_h(t)}{dt} = \alpha_b - \alpha_r - \frac{1}{\tau_h}N_h(t) + \frac{1}{\tau_{av}}N^h_{av}(t), \tag{7}$$

$$\frac{dN^s_{av}(t)}{dt} = \alpha_r - \frac{1}{\tau_{av}}N^s_{av}(t). \tag{8}$$

Equations (6), (7), (8) describe our behavioral ODE model. They replace Eqs. (2), (3), (4) described earlier in Sect. 3. We simplified them by removing α'_r and replacing $\tau_h, \alpha_b, \alpha_r$ with the mathematical derivations shown later in Sects. 5.1, 5.2, 5.3.

Improving on Eq. (5), Eq. (9) allows the number of blocks in the arena to increase *and* decrease:

$$\frac{dB_j(t)}{dt} = \frac{A_j}{A_d}\left[\frac{1}{\tau_h}N_h(t) - \alpha_b\right] \quad j = 1, \ldots, J. \tag{9}$$

We made the following additional assumptions about the block distribution. First, whenever a block is redistributed in the arena all j subareas are selected with probability proportional to the fraction of distributable area they contain. Second, blocks are distributed uniformly within a given j. Third, every j can hold any number of blocks, allowing two blocks to occupy the same location (i.e., stacking).

Next we derive analytical models for τ_h, α_b, and α_r from the arena geometry, number of blocks, block distribution, etc. We do not derive τ_{av}, because it depends intrinsically on the interference avoidance strategy employed by S and therefore cannot be derived independently from κ without additional assumptions.

5.1 Derivation of homing time τ_h

We use the following intuition to build the block acquisition probability for a robot. Since searching begins from the nest, the density of $N_s(t)$ must be *greater* near the nest because robots do CRWs starting from a common point. Consequently, this non-uniform spatial distribution means that the mean distance from the nest at which a searching robot encounters a block is *not* the same as the mean distance of a block from the nest. From (Codling et al., 2010), the spatial occupancy distribution from the central point falls off linearly as a result of a biased random walk with bias distribution $Uniform(-\theta, \theta)$. Thus, we would intuitively expect the following:

1. The fall off is moderated by $\rho_j = B_j(t)/A_j$, because the rate of decay of the mean block acquisition distance as a function of distance from the nest within a given A_j is slower for low ρ_j. Consequently, a j sparsely populated with blocks will have minimal effect on the overall swarm block acquisition probability distribution, while a j densely packed with blocks will create an area of higher acquisition probability.
2. ρ_j would play an exponentially moderating role only when the block acquisition location $\mathbf{x}$ is close to the center of the nest $\mathbf{x}_n$, such as for RN or PL block distributions. For SS and DS distributions, where the mean distance from a block to the nest is large, the effect of ρ_j on block acquisition locations should be minimal.

We approximate the occupancy distribution of a single walker doing a CRW starting from the nest center $\mathbf{x}_n$, using

the results of (Codling et al., 2010) and our intuitions:

$$p_{acq_j}(\mathbf{x}) = \frac{C}{\left(\sqrt{\|\mathbf{x} - \mathbf{x}_n\|} - \frac{\ln(\rho_j)}{2\rho_j}\right)^2}, \tag{10}$$

where C is a normalization constant to ensure $p_{acq_j}(\mathbf{x})$ integrates to 1 over J. Having defined the probability density function, we now derive the expected acquisition location by finding the expected values of the marginal density functions in x by integrating Eq. (10) and summing over all j:

$$E[x_{acq}] = \sum_J \int_x \int_y p_{acq_j}(\mathbf{x}) x dx dy, \tag{11}$$

and similarly for y. We now write an expression for τ_h^1:

$$\tau_h^1 = \frac{\left\| E[\mathbf{x}_{acq}] - \mathbf{x}_n \right\| - d_{cr}}{v_h} \tag{12}$$

where v_h is the robot phototaxis velocity, specified in the input configuration, and $\left\| E[\mathbf{x}_{acq}] - \mathbf{x}_n \right\|$ is the expected distance an acquired block will be from the center of the nest. d_{cr} is the distance the homing path is shortened due to the employed congestion reduction strategy, which is straightforward to calculate from the arena geometry.

An example calculation for a square arena, which in this work corresponds to the RN and PL scenarios, is shown below. Similar calculations can be done for other arena shapes. The mean distance to the origin from a randomly selected point in a region $R{:=}\{(x, y)0 \leq y \leq x \leq L/2\}$ (a triangle 1/8 of the nest) is calculated for a square of length L. The uniform density on this region is then

$$f(x, y) = \begin{cases} 8/L^2 & \text{if } (x, y) \in R \\ 0 & \text{otherwise.} \end{cases} \tag{13}$$

The mean distance d_{cr} is computed via

$$\begin{aligned} d_{cr} &= \frac{8}{L^2} \int_0^{L/2} \int_0^x \sqrt{x^2 + y^2} dy dx \\ &= \frac{L}{6}(\sqrt{2} + \ln(1 + \sqrt{2})). \end{aligned} \tag{14}$$

Finally, to derive τ_h, we note that under our assumption of low to moderate ρ_S, the homing time increases linearly with N according to the expected value of time lost due to inter-robot interference (Lerman and Galstyan, 2002), averaged across all robots, giving us

$$\tau_h = \tau_h^1 \left[1 + \frac{\alpha_r \tau_{av}}{N}\right]. \tag{15}$$

5.2 Derivation of block acquisition rate α_b

Using our mean-field assumption, S can be approximated as a fluid composed of robot particles, and be considered to obey many of the same laws; in the long-term limit, the governing equation for the CRW used in this work is the advection–diffusion equation (Codling et al., 2010). Given sufficiently simple robots (i.e., those which are reactive, memory-less), this approximation gives good results (Codling et al., 2010; Pang et al., 2019).

Using this intuition, we obtain α_b by computing the mean time it takes a robot "particle" starting at $\mathbf{x}_n$ to "diffuse" within the operating area A to the expected acquisition location $E[\mathbf{x}_{acq}]$. Viewing $\left\| E[\mathbf{x}_{acq}] - \mathbf{x}_n \right\|$ as the Root Mean Square (RMS) displacement distance and assuming a linear relationship between displacement and diffusion time, we obtain

$$\frac{1}{\alpha_b} = \frac{\left\| E[\mathbf{x}_{acq}] - \mathbf{x}_n \right\|^2}{2F(N)}, \tag{16}$$

where $F(N)$ is the diffusion constant for a swarm of N robots, and α_b is the expected time to diffuse from $\mathbf{x}_n$ to $E[\mathbf{x}_{acq}]$. We note that while the calculation for α_b does not directly consider the distribution of blocks in the arena and their densities within a given block cluster, it depends on $E[\mathbf{x}_{acq}]$. We approximate $F(N)$ using the results of the RMS diffusion for CRW (Codling et al., 2010) and the idea of *tortuosity*, i.e., the amount of turning in a given space or time, which states that highly tortuous paths will spread more slowly than straight paths (Codling et al., 2008), giving us

$$F(N) = N \frac{L_{df}(m) D_{xy}}{D_\theta}, \tag{17}$$

where $L_{df}(m)$ is a per-scenario parameter characterizing the linearity of the diffusion rate. D_{xy} is defined as follows, where v_s is the robot searching velocity and t is the timestep size: (both scenario parameters):

$$D_{xy} = \frac{{v_s}^2}{4t} \underbrace{\int_{-\pi}^{\pi} (1 \pm \cos 2\theta) f(\theta) d\theta}_{D_\theta}. \tag{18}$$

Under "normal" circumstances with inert particles, Eqs. (16), (17) give sufficient results because system diffusion varies linearly with time, which these equations require. However, many foraging environments are heterogeneous and have non-uniform distributions of blocks. As a result, the swarm may experience non-linear *anomalous diffusion* (Hasnain et al., 2018; Woringer et al., 2020) (i.e., crowding and heterogeneity in some areas of the environment) giving rise to the need for $L_{df}(m)$ to capture these artifacts. In the

derivations below, we account for this anomalous diffusion to "restore" linear diffusion in our model.

We define the *environment heterogeneity* $H(m)$ in reference to that of an environment with a perfectly uniform block distribution; that is, in reference to the RN scenario. Thus, $H(RN) = 0$. In general:

$$H(m) = \frac{1}{J}\begin{cases} \sum_{J} \left\| \mathbf{c}_j - \mathbf{x}_n \right\| & \text{if } m \in \{SS, DS\} \\ \sum_{J} E(\mathbf{c}_j - E[\mathbf{c}_j])(\mathbf{c}_j - E[\mathbf{c}_j])^T & \text{else.} \end{cases} \tag{19}$$

For RN and PL scenarios this is calculated as the variance of the mean inter-cluster distance, as shown in Eq. 19 (bottom). For SS and DS scenarios with $|J| \leq 2$, the mean inter-cluster distance is not meaningful, so $H(m)$ is calculated as the mean cluster distance from $\mathbf{x}_n$ (Eq. 19), top.

We refine our estimate of environmental heterogeneity by including the influence of crowding to obtain $H^*(m)$. Let the set of inter-cluster distances be $X = \{\left\| \mathbf{c}_i - \mathbf{c}_j \right\| \; i \in J, j \in J\}$, the minimum inter-cluster distance be $X_1 = \min(X)$, and the "second minimum element" be $X_2 = \min(X \setminus X_1)$, and so on. We have:

$$H^*(m) = \begin{cases} H(m) & \text{if } m \in \{SS, DS\} \\ 2^{H(m)} \frac{\frac{1}{3}\sum_{k=1}^{3} X_k}{\sqrt{A_x^2 + A_y^2}} & \text{else.} \end{cases} \tag{20}$$

For RN and PL scenarios, more heterogeneity will result in exponentially more crowding, as reflected in Eq. (20), (bottom). $H(m)$ is normalized by a factor of (a) the largest possible inter-cluster distance (the arena diagonal) and (b) the mean inter-cluster distance of the three closest clusters. For SS and DS scenarios (Eq. (20), top), no additional tuning for crowding was performed, using the intuition that most robots will spread out in the environment searching, thus resulting in relatively little crowding around the nest.

Finally, using a hybrid first-principle and post-hoc modeling approach we obtain $L_{df}(m)$. Post-hoc terms are shown in [ ] in Eqs. (21), (22), (23), (24).

$$L_{df}(RN) = D_\theta H^*(RN)[{A_x}^{0.1} J^{D_\theta}], \tag{21}$$

$$L_{df}(PL) = D_\theta H^*(PL)[{A_x}^{0.75} 0.7^J], \tag{22}$$

$$L_{df}(SS) = D_\theta H^*(SS)[{A_x}^{0.1}] + [0.1], \tag{23}$$

$$L_{df}(DS) = D_\theta H^*(DS)[{A_x}^{0.5}]. \tag{24}$$

Note that each block distribution has a slightly different linearity constant, placing higher precedence on some environmental variables over others, to provide a "glue" to combine the different components of our anomalous subdiffusion model. This aspect of our ODE model still requires post-hoc tuning, and is not purely derived from first principles; however, we again emphasize that the semantic meaning of these parameters makes their values more understandable, and suggests the possibility of future derivation.

5.3 Derivation of robot encounter rate α_r

To derive α_r, we use our assumption of low to moderate ρ_S, to model the robot encounter rate as a function of α_r^1 and $F(N)$. We view Fig. 1 as a queueing network, where robots are either performing collision avoidance maneuvers or not. The input/output transition rates for a state are summed to form the arrival and service rates for the collision avoidance queue Q_{ca} ($\lambda_{ca} = \alpha_r^1$ and $\mu_{ca} = 1/\tau_{av}$, respectively). Modeling Q_{ca} as a M/M/1 queue, i.e., at most one robot exits collision avoidance per Δt, which is reasonable if Δt is small, we can write the following using Little's Law (Šeda et al., 2017):

$$\alpha_r = \frac{N_{av}(t)}{\tau_{av}} - \alpha_r^1 N_{av}(t). \tag{25}$$

The second term in Eq. (25) is a corrective factor accounting for robots that experience interference due to encountering arena walls, *not* other robots, which is simply the scaled rate at which a single robot experiences interference. $N_{av}(t)$ cannot be used directly, as α_r needs to be computed *a priori*, but it can be estimated as a function of α_r^1 and $N_{av}^1(t)$, using an intuitive formulation of swarm diffusion:

$$\hat{N_{av}}(t) = N_{av}^1(t) \frac{F(N)}{D_\theta} C_{ca}(m), \tag{26}$$

where α_r^1 and $N_{av}^1(t)$ can be computed from κ using the results of (Codling et al., 2010). We increase the influence of θ in Eq. (26) by introducing another D_θ in the denominator; smaller θ will result in more inter-robot interference due to straight line motion. $C_{ca}(m)$ characterizes the sub- or super-linearity of Eq. (26) for a scenario m in relation to the RN scenario.

6 Experimental setup

Scenarios have no obstacles, except for robots that act as obstacles to each other. The parameters used in the scenarios are summarized in Table 3.

6.1 Simulation experiments

We use the ARGoS simulator (Pinciroli et al., 2012) with a dynamical physics model of the marXbot robot in a 3D space for maximum fidelity. Robots are restricted to motion in the

Table 3 Summary of experimental parameters

Parameter	Value/Notes
θ	$\pi/36$
τ_{av}	Computed from the robot control algorithm κ (details omitted)
$C_{ca}(m)$	Computed post-hoc for each scenario
t	0.2 seconds, i.e., the robot control algorithm runs at 5 Hz

XY plane.[1] We use the foraging scenarios from Fig. 2, and robots run the FSM in Fig. 1. Four sets of experiments are conducted, characterized by different densities, ρ_S. Experiments use $N = 1 \ldots 6,006$ robots in environments with SS and DS source block distributions, and $N = 11 \ldots 11837$ robots for RN and PL block distributions (see Fig. 2 for block distributions). Experimental results shown are from 32 runs of 200000 seconds each.

6.2 Real-robot experiments

We chose a robotic platform similar to the marXbot to minimize the simulation-reality gap: the TurtleBot3 (Amsters and Slaets, 2020) running ROS. TurtleBots have approximately the same physical size as the marXbot, and use the onboard lidar sensor for collision avoidance, which functions nearly identically to the proximity sensors on the marXbot. We extended TurtleBot3 with the following, as shown in Fig. 3:

1. *An ultrasonic sensor* used to detect blocks within a short distance in front of the robot.
2. *Four light sensors* used for localization with respect to a light source and to perform phototaxis. The four light sensors are mounted at angles $\pi/4, 3\pi/4, 5\pi/4, 7\pi/4$. This configuration always provides the robot with at least two sensor readings during phototaxis so it can form a reliable vector to the light source.
3. *Two static "arms"* that robots can use to "pickup" and hold onto nearby blocks by pushing them, once they are close enough to be sensed with the ultrasonic sensor.

Experiments were run in the space shown in Fig. 4. Experimental results are calculated from 4 runs of $T = 300$ seconds each, with $N = 1 \ldots 6$. Videos of the experiments are available.[2]

[1] Our is code open-source at https://github.com/jharwell/fordyca.

[2] Supplementary videos for both simulated and real-robot experiments can be found at https://www-users.cse.umn.edu/~harwe006/showcase/auro-2022-extra/.

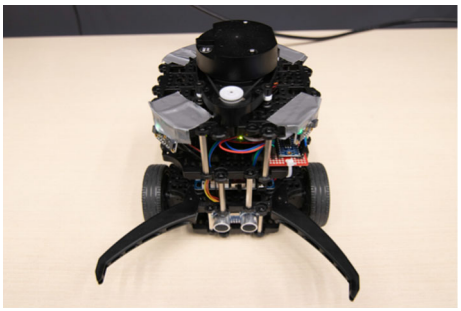

Fig. 3 Extended turtlebot3 robot. Light (mounted to gray squares on each corner) and ultrasonic (bottom center) sensors added. Static "arms" have been added to enable robots to "hold" onto blocks as they push them

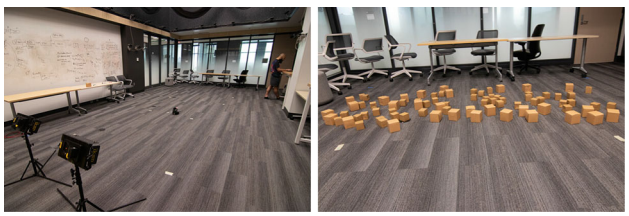

Fig. 4 The arena used for real robot experiments. Left: empty arena. Right: the blocks in the arena. The arena is 8m × 4m = 32m^2, and mirrors the layout of the single source distribution scenario shown in Fig. 2c

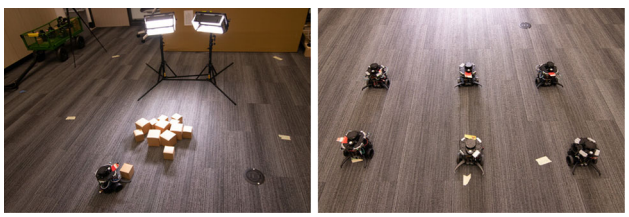

Fig. 5 The robots used for the experiments. Left: a robot bringing a block to the nest, which is in front of the lights. Right: the 6 robots used

7 Simulation experiments

Across the four sets of experiments, core assumptions we made when deriving our model are relaxed. For all the scenarios, predictions for $N_s(t)$ are omitted since it can be computed from the conservation of robots via $N_s(t) = N - N_{av}(t) - N_h(t)$. Results are presented with a non-scaled X-axis, with one data point per experiment, and a logarithmic Y-axis for larger N, to improve readability. We use $\rho_S = 0.1$

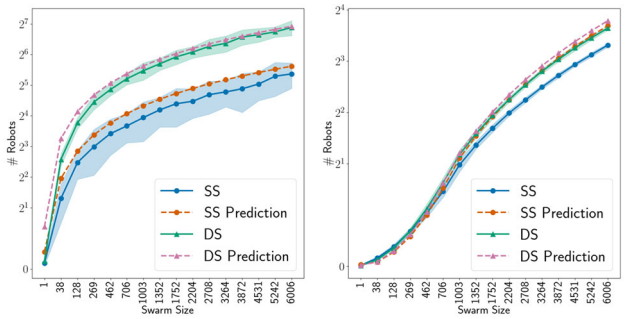

(a) Comparison of predictions of $N_h(t)$ (homing) behavior with experimental data for SS and DS scenarios

(b) Comparison of predictions of $N_{av}(t) = N^h_{av}(t) + N^s_{av}(t)$ (avoidance) behavior with experimental data for SS and DS scenarios

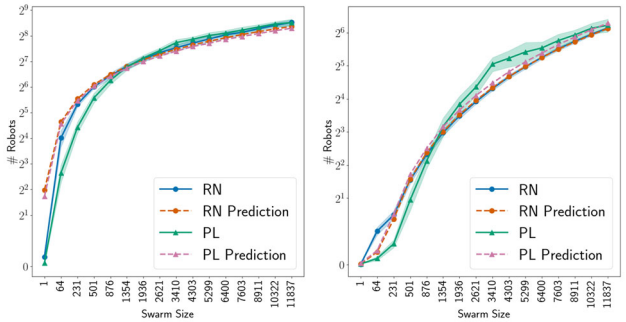

(c) Comparison of predictions of $N_h(t)$ (homing) behavior with experimental data for RN and PL scenarios

(d) Comparison of predictions of $N_{av}(t) = N^h_{av}(t) + N^s_{av}(t)$ (avoidance) behavior with experimental data for RN and PL scenarios

Fig. 6 Predictions of swarm behavior at large scales with constant population density $\rho_S = 0.1$. Single source (SS), dual source (DS), random (RN), and power law (PL) scenarios

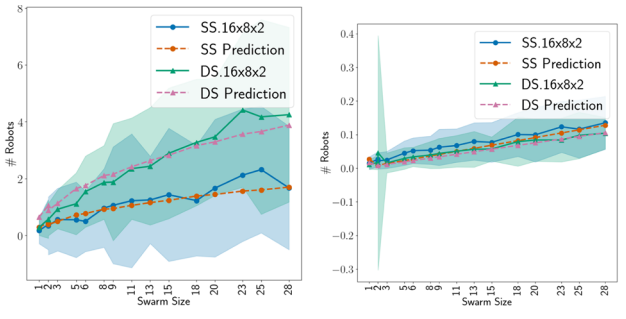

(a) Comparison of predictions of $N_h(t)$ (homing) behavior with experimental data for SS and DS scenarios

(b) Comparison of predictions of $N_{av}(t) = N^h_{av}(t) + N^s_{av}(t)$ (avoidance) behavior with experimental data for SS and DS scenarios

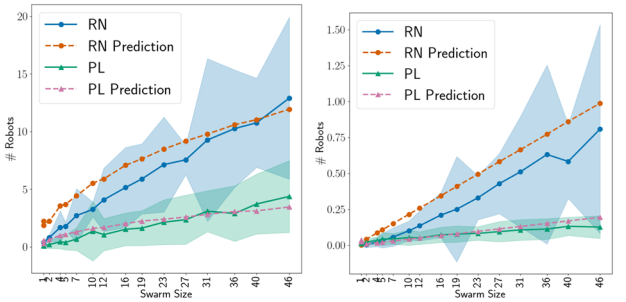

(c) Comparison of predictions of $N_h(t)$ (homing) behavior with experimental data for RN and PL scenarios

(d) Comparison of predictions of $N_{av}(t) = N^h_{av}(t) + N^s_{av}(t)$ (avoidance) with experimental data for RN and PL scenarios

Fig. 7 Predictions of swarm behavior at small scales with constant density $\rho_S = 0.1$. Single source (SS), dual source (DS), random (RN), and power law (PL) scenarios

as the upper limit for ρ_S for our experiments with constant swarm density; this mirrors the upper limit for real-world swarms of practical utility, as it corresponds to 1 robot/10m^2 which is a high density in the context of feasible real-world applications.

7.1 Constant ρ_S, large N

Uniform density ρ_S and large number of robots N are the "ideal" conditions in which to test our model.

In Fig. 6, we see strong agreement between the predictions of our model and experimental results for all scenarios, providing compelling evidence that our model is capturing the underlying dynamics of inter-robot interference and searching accurately. PL scenarios are the least favorable of all foraging environments, since they are asymmetrical and do not contain easily exploitable block clusters. Our model struggles to predict $N_h(t)$, $N_{av}(t)$ within the 95% confidence interval for PL scenarios, but does track the overall trend reasonably well, showing that our underlying diffusion model and assumptions about linearity of α_r are generally accurate even in this difficult case. The divergence between predictions and experiments for PL scenarios suggests that Eq. (10) is moderately inaccurate; this is further supported by slight differences between experimental data and predictions for $N_h(t)$ in the RN and SS scenario environments. In Fig. 6b, when our predictions of $N_{av}(t)$ are inaccurate in terms of confidence intervals for some DS scenarios, we find that the inaccuracy should not be of practical concern as the predictions only differ from the experimental results by < 1 robot/6000.

7.2 Constant ρ_S, small N

We test the mean-field assumption, and evaluate our model's ability to capture finite-size effects when S cannot be reasonably approximated using the mean-field model (with $N \leq 50$).

Figure 7 shows strong agreement between experimental results and the predictions of our model for all tested scenarios with small N. Even when the mean-field assumption is relaxed, our model accurately captures finite-size effects. The greatest discrepancy between predictions and results is shown in Fig. 7c for RN scenarios, which have the highest moments in the behavioral distribution because of the placement of the nest relative to the objects to be gathered. This results in non-uniform spatial distributions of the swarm. In

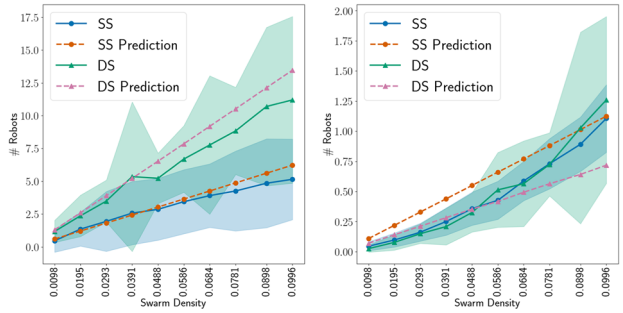

(a) Comparison of predictions of $N_h(t)$ (homing) behavior with experimental data for SS and DS scenarios

(b) Comparison of predictions of $N_{av}(t) = N^h_{av}(t) + N^s_{av}(t)$ (avoidance) behavior with experimental data for SS and DS scenarios

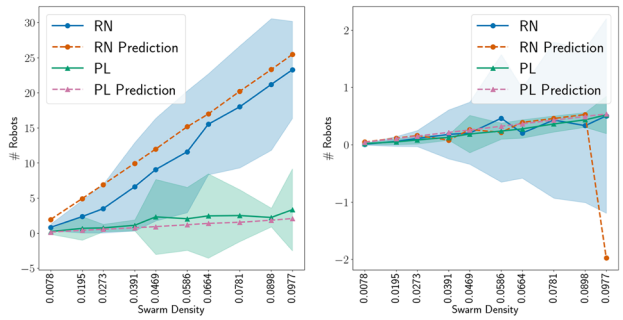

(c) Comparison of predictions of $N_h(t)$ (homing) behavior with experimental data for RN and PL scenarios

(d) Comparison of predictions of $N_{av}(t) = N^h_{av}(t) + N^s_{av}(t)$ (avoidance) behavior with experimental data for RN and PL scenarios

Fig. 8 Predictions of swarm behavior at small scales with variable population density $\rho_S = 0.01 - 0.1$. Single source (SS), dual source (DS), random (RN), and power law (PL) scenarios. Swarm size N varies $5 \ldots 51$ for SS and DS scenarios, and $2 \ldots 25$ for RN and PL scenarios

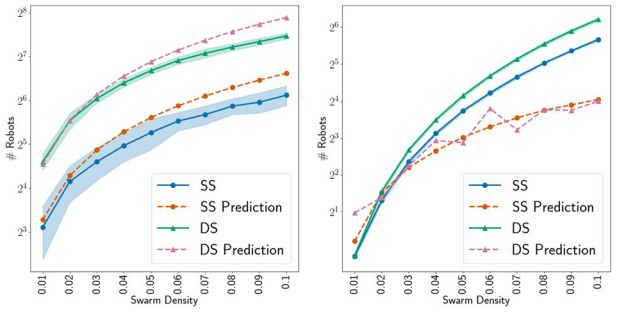

(a) Comparison of predictions of $N_h(t)$ (homing) behavior with experimental data for SS and DS scenarios

(b) Comparison of predictions of $N_{av}(t) = N^h_{av}(t) + N^s_{av}(t)$ (avoidance) behavior with experimental data for SS and DS scenarios

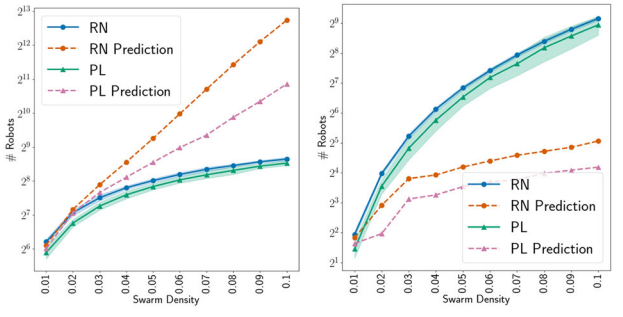

(c) Comparison of predictions of $N_h(t)$ (homing) behavior with experimental data for RN and PL scenarios

(d) Comparison of predictions of $N_{av}(t) = N^h_{av}(t) + N^s_{av}(t)$ (avoidance) behavior with experimental data for RN and PL scenarios

Fig. 9 Predictions of swarm behavior at large scales with variable population density $\rho_S = 0.01 - 0.1$. Single source (SS), dual source (DS), random (RN), and power law (PL) scenarios. Swarm size N varies $327 \ldots 3276$ for SS and DS scenarios, and $655 \ldots 6553$ for RN and PL scenarios

comparison to Sect. 7.1 stronger agreement between predictions and results with constant ρ_S is seen when N is small; as N increases, the size of the higher moments increases as well, as expected.

7.3 Variable ρ_S, small N

Continuing with our trend of relaxation, we now test our model's ability to capture finite-size effects with $\rho_S = 0.01 - 0.1$, and $N \leq 50$. These scenarios are much more adverse to our ODE modeling.

In Fig. 8 we see that our model generally performs well even when the low ρ_S assumption is violated. We see strong agreement between the model and the experimental data for all scenarios for $\rho_S = 0.01 - 0.05$ for both $N_h(t)$ and $N_{av}(t)$. In Fig. 8d, there is a striking divergence between predictions and results at $\rho_S = 0.09$. It may be a numerical anomaly or the point at which the ODE solver begins to struggle.

7.4 Variable ρ_S, large N

Finally, we now restore the mean-field assumption and test our model's ability to capture collective dynamics across scales on the order of thousands of robots when the spatial distribution of S is extremely non-uniform.

Figure 9 shows that our model generally does not model swarm behavior at large scales well at high densities. Substantial divergence between model predictions and results is observed for $N_{av}(t)$, $N_h(t)$ for all scenarios beyond $\rho_S = 0.02$, and beyond $\rho_S = 0.01$ for RN, PL scenarios. This maps directly to our intuition: with high robot densities and a centrally located nest (RN) or multiple irregularly spaced block clusters, the distribution of swarm behaviors has strong higher moments. For RN and PL scenarios, a strong correlation is observed between the ρ_S at which our model's predictions for both $N_h(t)$ and $N_{av}(t)$ break down, which happens at $\approx \rho_S = 0.01$. overall, the model fails to provide meaningful predictions due to the highly non-uniform spatial distributions, as expected.

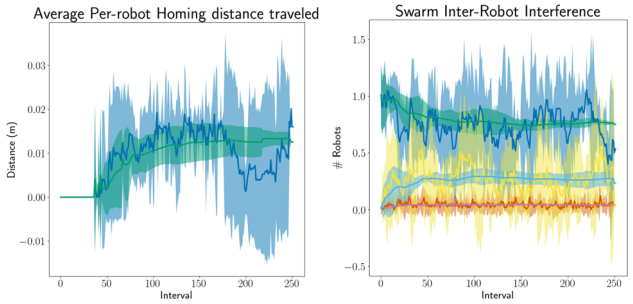

Fig. 10 Time series graphs of real robot behaviors. Homing behavior (left) and {entering, exiting, experiencing} types of inter-robot interference events (right) over time. Both instantaneous and running cumulative averages are shown for $N = 6$; other graphs for $N = 1, \ldots, 5$ omitted

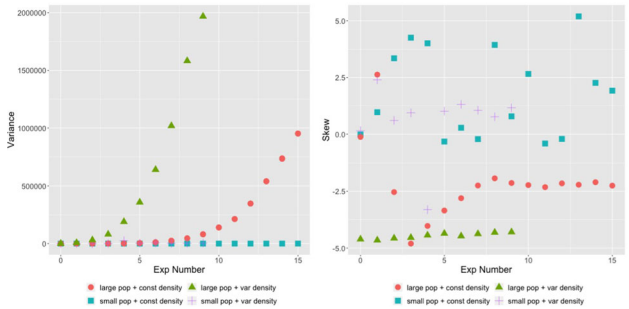

Fig. 11 Variance (left) and skew (right) of behavioral distributions. For {●,■} as X increases swarm size increases. For {+, ▲} as X increases both swarm size and swarm density increase (Color figure online)

8 Real-robot experiments

We describe some initial efforts to verify the simulation results from Sect. 7 using real robots.

In Fig. 10, we see that even under the reduced experiment times the swarm reaches steady-state, validating a crucial assumption of our ODE model. However, in order to cross the simulation-reality gap, we had to make the following adjustments to the real robot controllers. First, we had to remove the congestion-reducing behavioral assumption from Sect. 4, because it was not feasible to implement it without precise localization, which our simple robots did not have. Implementing this behavior is necessary for predictive modeling with real robots; however, we emphasize that even without it our model may still provide rough order-of-magnitude predictions which have utility in guiding design decisions. Second, we had to extensively tweak parameters such as maximum robot speed, collision-avoidance proximity threshold, etc., in order to minimize the human intervention in the experiments. In our initial experiments, robots would often turn too quickly upon sensing a block, resulting in the detected block slipping out of range of the arms. In addition, if robots got too close, their arms could get caught on each other, resulting in a deadlock.

9 Discussion

Our ODE-based model is *not* the first model in the literature capable of capturing the behavioral dynamics of swarms for large as well as small N. However, to the best of our knowledge it *is* the first such model that does so accurately on multiple block distributions, while also minimizing the use of post-hoc parameters.

When we consider the swarm behavior with constant ρ_S and large N (Fig. 6), we see that our model successfully predicts behavior at $N \geq 100$ (mean-field approximation holds). However, it struggles at $N < 100$ for SS, DS, and RN scenarios, underscoring the difficulty in creating an ODE model that can capture both mean-field and finite-size behavioral effects simultaneously. PL scenarios are the most challenging to model because of the increased spatial non-uniformity of the ρ_S distribution. Our model's predictions are usually outside the confidence interval range; however, this corresponds to accuracy within ≈ 16 robots/11837, which is sufficient in most cases. Furthermore, while our model will break down at some $\rho_S > 0.1$ that threshold is above any realistic scenario (densities > 0.1 are infeasible for real-world problems). For swarms in RN scenarios, large 95% confidence intervals reflect the greater variability in behavior.

These results highlight the main assumptions our model relies on. First, swarm behavior is a linear function of the behavior of a single robot. Second, robot events are distributed *close enough* to Eq. (1) to be amenable to ODE modeling. Clearly, these two assumptions are not valid in all contexts across scales or scenarios for similar N. Nevertheless, they are valid in *many* contexts (e.g., Fig. 7, 8, parts of Fig. 6). Third, the number of objects to be found in the arena is roughly constant; this is because our ODE model requires steady state. This maps to real-world scenarios, such as warehouse transport in which objects are transported from one side of the warehouse to the other, or arrive according to a different process, such as delivery trucks. There are also real-world scenarios, such as search and rescue, for which this assumption is not valid.

These assumptions limit the applicability of our model in the following ways. First, it cannot capture behaviors beyond the mean, because it uses homogeneous Poisson processes; non-homogeneous processes could have alternatively been employed to model *both* steady-state and transient behaviors with time-varying λ, as well as higher moments. If we calculate the higher moments of behavioral distributions present in a given scenario, as shown in Fig. 11 and compare to our results from Sect. 7, we see that the cases where our model struggles or fails are precisely where there are non-negligible higher moments. We further see that experi-

ments that are more difficult to model also frequently have negative skew: this matches our intuition, as heavy-tailed behavioral distributions across time run counter to our steady state assumption.

Our model could be extended by explicitly modeling higher moments (Mather and Hsieh, 2012; Klavins, 2010) or by adding time delays (Mather et al., 2013; Hsieh et al., 2008; Krishnarajah et al., 2005) to reduce behavioral variability and enhance model usability. Because we chose to target steady-state behaviors our model assumes and requires steady-state to work. Our model's parameters *also* assume steady-state. Specifically, they require that the number of objects to be found is constant, no matter how many have been picked up and brought to the nest. Thus, our model *cannot* be used in scenarios with a finite number of blocks, such as the ones studied in (Lerman and Galstyan, 2002).

Second, as a macroscopic model of foraging behavior, the choice of what to model affects the results. If the foraging behaviors are Poisson distributed with minimal skew, as in previous works (Lerman and Galstyan, 2002; Lerman et al., 2001), then the accuracy of the model is good. However, as shown in Sect. 7.4, when non-negligible higher moments are present in behavioral distributions, the predictions reveal the fundamental limits of the model.

If the macroscopic states of the model instead represent robot tasks corresponding to specific regions in the arena, then the model accuracy is not subject to the distribution of robot events, but depends on the distribution of task execution times (Berman et al., 2007; Mather et al., 2013; Hsieh et al., 2008). Clearly, the choice of what states the macroscopic model represents is arbitrary, but has profound consequences for what can be modeled, and how accurately.

As a result of our investigation, we can state the following about the limits of applying ODE modeling to foraging scenarios from an engineering standpoint: it is *possible* to capture behaviors using ODE models across scales, swarm densities, and scenarios, and general behavioral curves can be tracked with reasonably accuracy. Furthermore, doing so is also *practical*: our model's predictions are accurate within 95% confidence for $N \leq 50$ in all tested scenarios. It can be used to provide useful bounds on swarm behaviors in the tested scenarios—all of which are commonly used to model real-world problems. Thus, even though our model has no theoretical guarantees of applicability to *arbitrary* foraging scenarios, its demonstrated usefulness on *many* common foraging scenarios makes it of practical utility in designing SR systems.

Our model uses diffusion theory to model robots as particles in a liquid as a convenient abstraction, given that robots do not actually move as particles would. Diffusion theory is an ideal tool because the complex mathematical relationships it captures simplify to simple equations. Our model still relies on two free parameters: $C_{ca}(m)$ and $L_{df}(m)$, which have strong semantic meaning: they are connected to the levels of interference robots experience and the linearity of swarm diffusion, respectively, suggesting future mathematical derivation.

This is an improvement over previous work (Lerman and Galstyan, 2002). All the parameters in our model (free or calculated analytically), have strong semantic contexts, making it easier to understand *why* specific values give good predictions. The number of free parameters has been reduced from five to two.

When considering the utility of our ODE model beyond foraging, the specific values for the ODE parameters clearly depend on the problem configuration (e.g., N, block distribution). In that sense our model cannot be reused as-is. However, this does not diminish its utility for foraging tasks. Among the variants of foraging we can mention collaborative transport (Lerman et al., 2001) and reconnaissance (Hsieh et al., 2008). Consider a reconnaissance task: we could use diffusion theory to model swarms spreading over an environment as they search for a suitable location for a base of operations, modeling the searching behaviors using ODEs.

Most of the presented results were obtained in simulation, similar to other mathematical approaches to modeling SR systems (Guerrero-Bonilla et al., 2020; Tarapore et al., 2017). Our real robot experiments in Sect. 8, while preliminary, strongly suggest that our model is viable with real robots, opening the door to first-principle predictive modeling that is not only feasible, but that also has real-world utility.

We have shown that our model is accurate with swarms of non-trivial size (i.e., up to 46 robots in Fig. 7, up to 51 robots in Fig. 8, up to 6006 and 11837 robots in Fig. 6). In such swarms, non-deterministic transient behaviors will arise even in simulation, e.g., floating point representation errors; and we can see similar errors in small swarms of real robots, as visible in Fig. 10.

10 Conclusions and future work

We have investigated the ODE modeling paradigm to better understand its limits for CPFPs, and have contributed to swarm engineering by showing how it can be used in designing multi-robot systems. We developed a hybrid ODE model combining analytic derivation of parameters with post-hoc parameter fitting, and showed it is accurate in a variety of scenarios with widely varying behavioral dynamics that map to real-world problems. Furthermore, the model is sufficiently accurate even in scenarios for which the assumptions used to derive our model do not hold. Initial real-robot experiments

validate our model, with promising results suggesting that it is robust enough to be used in the real-world.

Declarations

Conflict of interest Partial financial support for this work was provided by the MnDRIVE initiative, the Minnesota Robotics Institute, and the University of Minnesota Informatics Institute. The Minnesota Supercomputing Institute provided computing resources. No other competing interests exist.

References

Amsters, R., & Slaets, P., (2020). Turtlebot 3 as a robotics education platform. In: Merdan, M., Lepuschitz, W., Koppensteiner, G., et al. (eds) *Robotics in education*. Springer International Publishing, pp 170–181

Berman, S., Halász, Á., Kumar, V., et al. (2007). Algorithms for the analysis and synthesis of a bio-inspired swarm robotic system. *Swarm Robotics, vol LNCS 4433* (pp. 56–70). Berlin: Springer-Verlag.

Castello, E., Yamamoto, T., & Nakamura, Y. et al. (2014). Foraging optimization in swarm robotic systems based on an adaptive response threshold model. In: *Proceedings of the IEEE Int'l Conference on Robotics and Automation*, pp 1343–1356

Castello, E., Yamamoto, T., Libera, F. D., et al. (2016). Adaptive foraging for simulated and real robotic swarms: the dynamical response threshold approach. *Swarm Intelligence, 10*(1), 1–31.

Codling, E., Plank, M., & Benhamou, S. (2008). Random walks in biology. *Journal of the Royal Society, Interface/the Royal Society, 5*, 813–34.

Codling, E. A., Bearon, R. N., & Thorn, G. J. (2010). Diffusion about the mean drift location in a biased random walk. *Ecology, 91*(10), 3106–3113.

Cotsaftis, M. (2009). An emergence principle for complex systems. *Complex sciences* (pp. 1105–1117). Heidelberg: Springer.

De Wolf, T., & Holvoet, T., et al. (2005). Emergence versus self-organisation: different concepts but promising when combined. In S. A. Brueckner, Serugendo G. Di Marzo, & A. Karageorgos (Eds.), *Engineering self-organising systems* (pp. 1–15). Berlin: Springer.

Elamvazhuthi, K., & Berman, S. (2019). Mean-field models in swarm robotics: a survey. *Bioinspiration and Biomimetics*, *15*(1), 015001

Ferrante E, Turgut AE, Duéñez-Guzmán E, et al. (2015). Evolution of self-organized task specialization in robot swarms. *PLOS Computational Biology*, *11*(8)

Galstyan, A., Hogg, T., & Lerman, K. (2005). Modeling and mathematical analysis of swarms of microscopic robots. In: *Proceedings 2005 IEEE Swarm Intelligence Symposium, 2005*. SIS 2005., pp. 201–208

Georgé, J. P., & Gleizes, M. P. (2005). Experiments in emergent programming using self-organizing multi-agent systems. *Multi-Agent Systems and Applications IV, 3690*, 450–459.

Ghosh, S.K., Cherstvy, A.G., & Grebenkov, D.S. et al. (2016). Anomalous, non-gaussian tracer diffusion in crowded two-dimensional environments. *New Journal of Physics*

Guerrero-Bonilla, L., Saldaña, D., & Kumar, V. (2020). Dense r-robust formations on lattices. In: *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6633–6639

Hamann, H. (2013). Towards swarm calculus: urn models of collective decisions and universal properties of swarm performance. *Swarm Intelligence, 7*, 145–172.

Harwell, J., & Gini, M. (2019). Swarm engineering through quantitative measurement of swarm robotic principles in a 10,000 robot swarm. In: *Proceedings of the Twenty-Eighth Int'l Joint Conference on Artificial Intelligence*, IJCAI-19, pp. 336–342

Harwell J, Lowmanstone L, Gini M (2020) Demystifying emergent intelligence and its effect on performance in large robot swarms. In: *Proceedings of the Autonomous Agents and Multi-Agent Systems (AAMAS)*, pp. 474–482

Hasnain, S., Harbola, U., & Bandyopadhyay, P. (2018). A memory-based random walk model to understand diffusion in crowded heterogeneous environment. *International Journal of Modern Physics B*, *32*(16):1850193

Hecker, J. P., & Moses, M. E. (2015). Beyond pheromones: evolving error-tolerant, flexible, and scalable ant-inspired robot swarms. *Swarm Intelligence, 9*(1), 43–70.

Hsieh, M. A., Halász, Á., Berman, S., et al. (2008). Biologically inspired redistribution of a swarm of robots among multiple sites. *Swarm Intelligence, 2*(2), 121–141.

Hunt, E. R. (2020). Phenotypic plasticity provides a bioinspiration framework for minimal field swarm robotics. *Frontiers in Robotics and AI, 7*, 23.

Ijspeert, A. J., Martinoli, A., Billard, A., et al. (2001). Collaboration through the exploitation of local interactions in autonomous collective robotics: The stick pulling experiment. *Autonomous Robots, 11*(2), 149–171.

Kampen, N. V. (2007). *Stochastic processes in physics and chemistry*. North Holland.

Klavins, E. (2010). Proportional-integral control of stochastic gene regulatory networks. In: *49th IEEE Conference on Decision and Control (CDC)*, pp. 2547–2553

Krishnarajah, I., Cook, A., Marion, G., et al. (2005). Novel moment closure approximations in stochastic epidemics. *Bulletin of Mathematical Biology, 67*(4), 855–873.

Labella, T. H., & Dorigo, M. (2006). Division of labor in a group of robots inspired by ants' foraging behavior. *ACM Trans on Autonomous and Adaptive Systems (TAAS), 1*(1), 4–25.

Lerman, K., & Galstyan, A. (2002). Mathematical model of foraging in a group of robots: Effect of interference. *Autonomous Robots, 13*(2), 127–141.

Lerman, K., & Galstyan, A. (2003). Macroscopic analysis of adaptive task allocation in robots. In: *Proceedings of the 2003 IEEE/RSJ Int'l Conference on Intelligent Robots and Systems (IROS 2003)*, pp. 1951–1956

Lerman, K., Galstyan, A., Martinoli, A., et al. (2001). A macroscopic analytical model of collaboration in distributed robotic systems. *Artificial Life, 7*(4), 375–393.

Lerman, K., Jones, C., Galstyan, A., et al. (2006). Analysis of dynamic task allocation in multi-robot systems. *International Journal of Robotics Research, 25*(3), 225–241.

Lu, Q., Fricke, G., Ericksen, J., et al. (2020). Swarm foraging review: Closing the gap between proof and practice. *Curr Robot Reports, 1*, 215–225.

Mather, T. W., Hsieh, M. A. (2012). Ensemble modeling and control for congestion management in automated warehouses. In: *2012 IEEE Int'l Conference on Automation Science and Engineering (CASE)*, pp. 390–395

Mather, T. W., Braun, C., & Hsieh, M. A., et al. (2013). Distributed filtering for time-delayed deployment to multiple sites. In A. Martinoli, F. Mondada, & N. Correll (Eds.), *The 10th International Symposium Distributed Autonomous Robotic Systems* (pp. 299–312). Berlin: Springer.

Matthey, L., Berman, S., & Kumar, V. (2009). Stochastic strategies for a swarm robotic assembly system. In: *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 1953–1958

Metzler, R., Jeon, J. H., & Cherstvy, A. G., et al. (2014). Anomalous diffusion models and their properties: non-stationarity, non-ergodicity, and ageing at the centenary of single particle tracking. *Physical Chemistry Chemical Physics*, *16*, 24128–24164

Nicolau, D. V., Jr., Hancock, J. F., & Burrage, K. (2007). Sources of anomalous diffusion on cell membranes: A Monte Carlo study. *Biophysical Journal, 92*(6), 1975–1987.

Oliveira, F. A., Ferreira, R. M. S., & Lapas, L. C., et al. (2019). Anomalous diffusion: A basic mechanism for the evolution of inhomogeneous systems. *Frontier Physics*

Pang, B., Song, Y., & Zhang, C., et al. (2019). A swarm robotic exploration strategy based on an improved random walk method. *Journal of Robotics*, 2019

Pinciroli, C., et al. (2012). ARGoS: A modular, parallel, multi-engine simulator for multi-robot systems. *Swarm Intelligence, 6*, 271–295.

Pini, G., Brutschy, A., Frison, M., et al. (2011). Task partitioning in swarms of robots: An adaptive method for strategy selection. *Swarm Intelligence, 5*(3–4), 283–304.

Ramachandran, R. K., Fronda, N., & Sukhatme, G. S. (2020). Resilience in multi-robot target tracking through reconfiguration. In: *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp 4551–4557

Renshaw, E., & Henderson, R. (1981). The correlated random walk. *Journal of Applied Probability, 18*(2), 403–414.

Şahin, E. (2005). Swarm robotics: From sources of inspiration to domains of application. In: *Swarm Robotics*. Springer, LNCS 3342, pp 10–20

Santamaria-Holek, I., & Vainstein, M. (2009). Protein motors induced enhanced diffusion in intracellular transport. *Physica A: Statistical Mechanics and its Applications, 388*(8), 1515–1520.

Šeda, M., Šedová, J., & Horký, M. (2017). Models and simulations of queueing systems. In: Matoušek, R. (ed) *Recent Advances in Soft Computing*. Proceedings of the 22nd international conference on soft computing (MENDEL 2016), vol. 576. Springer

Sugawara, K., & Sano, M. (1997). Cooperative acceleration of task performance: Foraging behavior of interacting multi-robots system. *Physica D: Nonlinear Phenomena, 100*(3–4), 343–354.

Talamali, M. S., Bose, T., Haire, M., et al. (2020). Sophisticated collective foraging with minimalist agents: A swarm robotics test. *Swarm Intelligence, 14*(1), 25–56.

Tarapore, D., Christensen, A. L., & Timmis, J. (2017). Generic, scalable and decentralized fault detection for robot swarms. *PLOS One, 12*(8), 1–29.

Tarapore, D., Gross, R., & Zauner, K. P. (2020). Sparse robot swarms: Moving swarms to real-world applications. *Frontiers in Robotics and AI, 7*, 83.

Vlahos, L., Isliker, H., & Kominis, Y., et al. (2008). Normal and anomalous diffusion: A tutorial. arXiv:0805.0419 [nlin.CD]

Weiss, M., Hashimoto, H., & Nilsson, T. (2003). Anomalous protein diffusion in living cells as seen by fluorescence correlation spectroscopy. *Biophysical Journal, 84*(6), 4043–4052.

Woringer, M., Izeddin, I., Favard, C., et al. (2020). Anomalous subdiffusion in living cells: Bridging the gap between experiments and realistic models through collaborative challenges. *Frontiers in Physics, 8*, 134.

John Harwell received his PhD at the University of Minnesota from the Department of Computer Science under Dr. Maria Gini in 2022. His research interests include multi-agent systems, mathematical modeling, and bio-inspired algorithm design.

Angel Sylvester is a graduate student studying Artificial Intelligence and Robotics with Maria Gini at the University of Minnesota. She has primarily worked on projects related to ODE modeling in SWARM robotics alongside senior graduate student John Harwell in addition to outreach projects dedicated to promoting diversity in STEM based courses through SEISMIC and the Minnesota Robotics Institute.

Maria Gini is a College of Science & Engineering Distinguished Professor in the Department of Computer Science and Engineering at the University of Minnesota. She works on decision making for autonomous agents in application domains, ranging from swarm robotics to distributed methods for allocation of tasks, to methods for robots to explore an unknown environment, and navigation in dense crowds. She is a Fellow of AAAI, ACM, and IEEE. She has published more than 60 journal articles and more than 300 conference papers and book chapters. She is Editor in Chief of Robotics and Autonomous Systems, and is on the editorial board of numerous journals, including Autonomous Agents and Multi-Agent Systems, Current Robotics Reports, and Integrated Computer-Aided Engineering.