PNAS RESEARCH ARTICLE | COMPUTER SCIENCES SYSTEMS BIOLOGY

Modeling feasible locomotion of nanobots for cancer detection and treatment

Noble Harasha[a], Cristina Gava[b], Nancy Lynch[a,1], Claudia Contini[c], and Frederik Mallmann-Trenn[b]

Affiliations are included on p. 11.

Contributed by Nancy Lynch; received May 3, 2025; accepted October 22, 2025; reviewed by Andrea W. Richa and Christian Scheideler

Deploying motile nanosized particles, also known as "nanobots," in the human body promises to improve selectivity in drug delivery and reduce side effects. We consider a swarm of nanobots locating a single cancerous region and treating it by releasing an onboard payload of drugs at the site. At nanoscale, the computation, communication, sensing, and locomotion capabilities of individual agents are extremely limited, noisy, and/or nonexistent. We present a general model to formally describe the individual and collective behavior of agents in a colloidal environment, such as the bloodstream, for cancer detection and treatment by nanobots. This includes a feasible and precise model of agent locomotion, inspired by actual nanoparticles that, in the presence of an external chemical gradient, move toward areas of higher concentration by means of self-propulsion. We present two variants of our model: the first assumes an endogenous chemical gradient fixed over time and centered at the cancer site; the second is a more speculative, dynamic variant in which agents themselves create and amplify a gradient centered at the cancer site. In both settings, agents sense the gradient and ascend it noisily, locating the cancer site more quickly than by simple Brownian motion. For the first variant, we present simulation and analytical results to bound the time it takes for agents to reach the cancer site. For the second, simulations highlight collective benefit from agent-generated signaling, showing significant runtime improvement via chemical signal amplification.

nanobots | cancer detection | cancer treatment | agent-based models | random walks

Significance

We present a mathematical model of nanorobots moving in a colloidal environment within the human body to locate a single, targeted cancer site and deliver localized treatment. The capabilities and behavior of individual agents are inspired by actual chemotactic nanoparticles, supporting the model's physical feasibility. We consider two scenarios: one in which a natural chemical signal guides the agents, and another in which agents generate and release their own chemical signal to establish this gradient. While the latter scenario is more speculative, it shows a marked improvement in performance. These results suggest potential design directions for future nanoparticles capable of cooperative signal amplification to enhance targeted drug delivery.

Motile nanoparticles suspended in a solution, or what are generally referred to as "nanobots," possess great potential in many medical applications as their scale allows for penetrating biological barriers to otherwise unreachable regions of the body (e.g., the blood–brain barrier), unique maneuverability within certain regions of the body, the deployment of swarms containing significantly high numbers of individual agents, and general precision. We consider the dual problem of cancer detection and treatment in the context of nanomedicine, whose use holds significant promise. Nanobots are being studied and engineered to move within the human body and target sites of interest, such as a group of cancerous cells, in order to offer treatment by releasing the appropriate drugs once the target site is located and reached (1–3).

This offers a drug delivery solution which is more precise and selective, and thus less toxic to extraneous regions, especially in comparison to existing methods of treatment such as chemotherapy. This promising approach, however, faces a number of challenges including limited, if not absent, control on the behavior of these robots as a consequence of their very small size. If externally controlled movement is to be implemented, it often comes at the cost of having larger scale (e.g., microscopic) robots.

We therefore investigate the interesting and crucial challenge of nanobots finding and treating a cancerous site, distributively and autonomously. In this work, we consider the process to be stochastic in nature, with nanobots performing a biased random walk, influenced by the chemical makeup and physics of the environment through which they navigate.

We first consider the situation—hereinafter referred to as the "passive agent model"—in which there already exists a naturally occurring chemical signal coming from the targeted cancer site which agents can sense and follow. Next, we consider a speculative scenario—hereinafter referred to as the "active agent model"—in which there is no endogenous chemical signal to follow; agents must find the cancer site initially unaided, and then eventually create a chemical signal themselves for the later agents to use. We formally describe the general models we propose for both settings in Section 2, with a particular focus on a model of agent locomotion that is based upon experimental data

Author contributions: N.H., C.G., N.L., C.C., and F.M.-T. designed research; N.H., C.G., and F.M.-T. performed research; N.H., C.G., N.L., C.C., and F.M.-T. contributed new analytic tools; N.H., C.G., and F.M.-T. analyzed data; and N.H., C.G., N.L., C.C., and F.M.-T. wrote the paper. All authors reviewed and approved the final manuscript.

Reviewers: A.W.R., Arizona State University; and C.S., Paderborn University.

The authors declare no competing interest.

Copyright © 2025 the Author(s). Published by PNAS. This article is distributed under Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 (CC BY-NC-ND).

[1]To whom correspondence may be addressed. Email: lynch@csail.mit.edu.

This article contains supporting information online at https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2510036122/-/DCSupplemental.

Published November 24, 2025.

Downloaded from https://www.pnas.org by 92.41.158.49 on July 17, 2026 from IP address 92.41.158.49.

on actual motile nanoparticles. We then present simulation and analytical results for the passive agent model in Section 3 and simulation results for the active agent model in Section 4. The results for the active agent model show a significant runtime improvement as the total number of agents in the swarm increases. We conclude the work with a brief discussion and suggestions for future work. The full simulation code is available on GitHub.

1.1. Related Work. At the end of the 20th century, the first ideas on nanotechnology and its potential applications in medicine emerged and laid the groundwork for the development of the field of nanomedicine. The works in refs. 4 and 5 present some of the first introductions to the concept of nanotechnology and the potential gains of its use for medical purposes, including drug delivery, diagnostics, and tissue engineering. The review in ref. 6 presents, instead, the major advancements up to date in the general field of colloidal robotics, exploring application space (i.e. medicine, biology, environmental engineering), the organized behavior of these robots and their design principles.

Here, we study the application of nanoparticles to drug delivery for cancer treatment, as explored in refs. 1, 2, and 7. The main motivation behind the use of nanoparticles here is the reduction of systemic side effects in favor of a treatment with the same level of efficacy, if not higher. While immune cells, for example, are capable of long-distance sensing, most nanoparticles can be assumed to only have access to a very local subset of information about their environment. The general problem of having a collective of agents locate an unknown target site when long-distance sensing is impossible is well studied in the field of swarm robotics, including works which model social insects and other animal groups; at nanoscale, though, unique challenges arise.

Beyond limited sensing capabilities, precise control over the locomotion of individual nanobots is difficult, if not impossible. Authors in ref. 8 consider nanobots to have zero locomotion capabilities with their position changing by passive displacement, simply following the flow of the circulatory system. On the other hand are a series of works (9, 10), operating from a centralized perspective, which achieve precise control over the locomotion of agents by employing external rotating magnetic fields. Our study, however, assumes nanobots are able to move distributively and autonomously in space, by performing some sort of random walk, one that is not entirely determined by the external forces at play in the given environment (e.g., blood flow).

The Lévy flight random walk (11–13) is a popular model for the movement of simple biological agents, as it is mathematically elegant, has been observed among many actual animal species, and has been proven to produce nearly optimal results in the problem of foraging. The Lévy flight random walk can be very loosely characterized by long strides interspersed with random self reorientation. When considering particles moving in a medium with a weak chemical gradient, the Lévy flight random walk is a framework not fully dissimilar from chemotaxis, described as the movement of organisms in response to a chemical stimulus, particularly, a chemical gradient. Studies involving chemotaxis typically consider a colloidal solution—a mixture where one substance, made of nanoscopic or microscopic insoluble particles, is dispersed in another substance, which is usually termed the medium of the solution. Studies show how different types of cancer cells are able to create their own chemical gradient to facilitate cancer growth through chemotaxis (14–16). The very same chemical gradient can therefore be used by other particles to localize the cancer site via a similar chemotactic behavior. Indeed, as presented by a number of works, such as refs. 17 and 18 the chemotactic behavior of nanoparticles proves pivotal to localized drug delivery. When considering a clearly sensed, uniform gradient such as our case, we will show that the random walk our nanoparticles perform is a biased random walk. The studies in refs. 19–23 are deeply relevant for chemotaxis by nanoparticles. The works in refs. 19, 20, and 22 look at the self-propulsion of nanoparticles in which the medium in which they are suspended does not present any particular gradient to follow; particles are able to move by the mechanism of chemotaxis, but this does not result in a consistent or specific direction of movement by the collective. Authors in ref. 20 present a detailed model of the walk being performed, with the precise characterization of a drift velocity and a rotational velocity.

In contrast, in refs. 21 and 23, the authors present novel nanoparticles which are suspended in a solution with a nonuniform chemical gradient; via chemotaxis, the nanoparticles self-propel and follow the gradient, moving in a directed fashion. In ref. 21, authors present an in-vitro empirical analysis of nanoparticles' movement following the concentration gradient of a chemical dissolved in the colloidal solution. The chemical dissolved in the solution is glucose. The authors show how it is precisely the presence of this gradient that favors a specific type of walk from these particles, in which they ultimately tend toward areas of higher chemical concentration, noisily ascending the gradient. The nanoparticles are asymmetrical, biocompatible vesicles containing glucose oxidase and catalase. In the presence of an external glucose gradient, these encapsulated enzymes react with the glucose, producing molecules that escape from a more permeable region of the membrane. The subsequent asymmetric distribution of product molecules outside the nanoparticle induces a slip velocity on the surface of the particle in the opposite direction of the gradient, propelling it forward in the direction of the gradient (21).

Based on the empirical results in ref. 21, we define a mathematical model, our passive agent model, which rigorously characterizes the movement of these nanoparticles. We extend this to an active agent model, which proposes additional agent capabilities and environment assumptions. We note that the polymersomes in the colloidal solution examined in ref. 21 are in the range of 50 nm, supporting our models at the nano scale. This model is also relevant to cancer detection, and is consistent with the locomotion of agents in our passive agent model based on ref. 21. Basic signal amplification ideas similar to those in our active agent model were described and studied in refs. 24 and 25.

2. Model

We present a continuous space, discrete time general model for the problem of cancer detection and treatment by nanobots in the human body, including two variants of the model with distinct environment and agent capabilities assumptions to be outlined below. In particular, we include a precise model for agent locomotion. Table 1 is provided as a lookup table for the parameter notation used for the formal mode description of this general model (and its two variants).

In both variants of the model, a set of n identical agents—nanobots—move in a two-dimensional Euclidean space $\mathbb{R}^2$. Time is discretized. We consider one, single cancer site that is concentrated in a single point in space. The cancer site naturally produces some unique surface cell marker, which agents can bind

Downloaded from https://www.pnas.org by 92.41.158.49 on July 17, 2026 from IP address 92.41.158.49.

to via the appropriate antibody which we assume they possess. This allows agents to detect the presence of a nearby cancer site once they are within distance $\varepsilon > 0$, with perfect accuracy; we consider them to have "reached the cancer site" at this point. In both variants, agents have a step size of α. We define the dimensionless quantities $\tilde{\alpha} = \alpha/\ell, \tilde{\varepsilon} = \varepsilon/\ell$, where $\ell = 1$ m.

2.1. Passive Agent Model. In the passive agent variant of the model, each agent carries a payload of a drug for cancer treatment. Once an agent is this ε-distance away from the cancer site, it immediately releases its payload of drug, delivering its marginal treatment, and then, for all practical purposes, effectively ceases to exist in the environment.

In this "passive agent model," we also importantly assume there exists a specific global chemical gradient centered at the cancer site, which agents can sense. We call this chemical the "signal chemical," which can be imagined to be some naturally occurring entity which increases in concentration at closer distances to tumors and cancerous cells. Note that this gradient is time-constant, i.e., for any location x in the space, the chemical concentration at x remains constant over time. The role of the signal chemical will be formalized by the movement model in Section 2.3. No direct interaction or communication occurs between agents. An agent's movement is a function of only its previous state and the signal chemical that it is currently sensing in the vicinity of its current location.

We investigate the agents' ability to distributively and autonomously locate the cancer site and release their payloads (i.e., the treatment). Given the random and noisy behavior of individual agents, our goal will be for at least 75%* of the agents to deliver their marginal treatment. The following quantities are involved in the process.

The location of the cancer site is indicated with x^* and we assume without loss of generality that $x^* = \vec{0}$ (the two-dimensional null-vector). The location of an agent i (arbitrary index $i \in \mathbb{N}, i \in [0, n-1]$) at time t is $x_i^{(t)} \in \mathbb{R}^2$, and its orientation is the vector $\theta_i^{(t)} \in \mathbb{R}^2$. The location of the initial site from which all agents begin is indicated with $x^0 \in \mathbb{R}^2$, i.e., $x_i^{(0)} = x^0 \ \forall i$. Define the dimensionless positions by $\tilde{x}^* := \frac{x^*}{\ell}$, $\tilde{x}_i^{(t)} := \frac{x_i^{(t)}}{\ell}$, $\tilde{x}^0 := \frac{x^0}{\ell}$, with $\ell = 1$ m.

The number of agents that have released their drug payload up to time t is $y^{(t)} \in \mathbb{N}$, with $y^{(0)} = 0$.

The concentration of the "signal chemical" at location $x \in \mathbb{R}^2$ is given by $\gamma(x) = \frac{P}{4\pi D t^*} \cdot \exp\left(-\frac{(||x-x^*||_2)^2}{4Dt^*}\right)$, where P, D, and t^* are constants. Imagining the signal chemical gradient to be the result of instantaneous point source diffusion frozen at a specific point in time, P, D, and t^* are, respectively, the point mass amount, diffusion coefficient, and time since point mass introduction (i.e., time diffusing).†

To be concrete, all parameters values hereafter are in *SI Appendix* units, i.e., P is in kg, D is in $\mathrm{m}^2 \cdot \mathrm{s}^{-1}$, and t^* is in s, yielding $\gamma(x)$ to have units $\mathrm{kg} \cdot \mathrm{m}^{-2}$. Define the dimensionless analogue

$$\tilde{\gamma}(\tilde{x}) = \frac{\ell^2}{P}\gamma(\ell\tilde{x}) = \frac{1}{\pi\tau^*}\exp\left(-\frac{(||\tilde{x}-\tilde{x}^*||_2)^2}{\tau^*}\right),$$

where $\tau^* = 4Dt^*/\ell^2$. To be consistent with ref. 21, we can think of one timestep in our model to be equivalent to one second. If we assumed that agents moved faster, we could consider larger total areas of operation while still achieving reasonable runtimes.

2.2. Active Agent Model. Next we consider a speculative and more dynamic variant of the general model as just introduced, which we call the "active agent model." We begin by removing the assumption that there already naturally exists some "signal chemical" gradient in the environment which is centered at the cancer site. However, agents can still perfectly detect the presence of the cancer site within an ε-distance and subsequently bind there to release their payload. Now, half ($\lceil \frac{n}{2} \rceil$) of the agents carries a payload of the cancer treating drug, while the other half ($\lfloor \frac{n}{2} \rfloor$) carries a payload of some artificial "signal chemical" which all agents can sense. The artificial signal chemical diffuses perfectly radially out from the cancer site via instantaneous point-source diffusion (26) (see the new $\gamma(\cdot)$ below), forming a global signal chemical gradient centered at the cancer site. This chemical gradient will serve the same purpose as the signal chemical gradient in the passive agent model, being followed by agents to locate the cancer site more efficiently. Before, the signal chemical gradient was endogenous and constant over time; now, it is created by the agents themselves and changes over time "passively" as a result of diffusion and "actively" as a result of more agents releasing their payloads.

We redefine $y^{(t)}$ and introduce the quantity $z^{(t)}$ as follows: The number of agents that have released their drug payload up to time t is $y^{(t)} \in \mathbb{N}$, with $y^{(0)} = 0$. The number of agents that have released their signal chemical payload up to time t is $z^{(t)} \in \mathbb{N}$, with $z^{(0)} = 0$. Imagine one agent releases its signal chemical payload of amount P at the cancer site at time t^*. It will immediately begin diffusing, with the concentration of this individual payload, at location x at time t, being a function of the distance from the cancer site $||x - x^*||_2$ and the time since its release $(t - t^*)$: $\frac{P}{4\pi D(t-t^*)}\exp\left(-\frac{(||x-x^*||_2)^2}{4D(t-t^*)}\right)$ (26), where D is the diffusion coefficient. We simplify by assuming that the diffusion of each agent's payload is independent—so that the concentrations are additive. Thus, the concentration of the signal chemical at time t at location $x \in \mathbb{R}^2$ is therefore given by $\gamma^{(t)}(x) = \frac{P}{4\pi D}\sum_{j=1}^{z^{(t)}} \frac{1}{t-t_j^*}\exp\left(-\frac{(||x-x^*||_2)^2}{4D(t-t_j^*)}\right)$ where t_j^* is the time at which the j'th agent with a signal chemical payload to reach the cancer site released their payload.

Define the dimensionless, time-varying concentration

$$\tilde{\gamma}^{(t)}(\tilde{x}) = \frac{\ell^2}{P}\gamma^{(t)}(\ell\tilde{x}) = \frac{1}{\pi}\sum_{j=1}^{z^{(t)}}\frac{1}{\tau_j}\exp\left(-\frac{\|\tilde{x}-\tilde{x}^*\|_2^2}{\tau_j}\right),$$

where $\tau_j := \frac{4D(t-t_j^*)}{\ell^2}$.

Our goal now will be for 75% of the agents with a cancer drug payload to deliver their marginal treatment. Besides the specific items noted above, the rest of the general model is left unchanged from the passive agent model in Section 2.1. As more agents reach the cancer site, they maintain, if not amplify, the signal chemical gradient, thus helping other agents find the cancer site more

*This proportion is arbitrarily chosen such that the total treatment delivered is effective, while still being sufficiently less than 100%. In this way, drug overdosing is avoided with high probability and runtime is reasonably finite.

†The constant t^* here is a parameter representing the steepness of the gradient. It can be viewed as the amount of time that passed between the point mass introduction and when time is frozen, at which point we stop the diffusion process from evolving further and the gradient from changing.

Downloaded from https://www.pnas.org by 92.41.158.49 on July 17, 2026 from IP address 92.41.158.49.

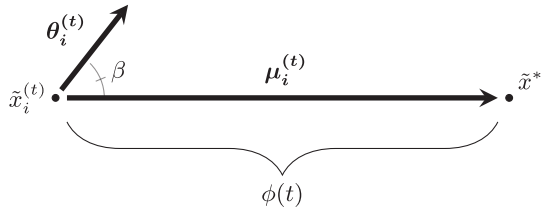

Fig. 1. Visual representation of quantities defined for the orientation-biased movement model for an agent i. Precisely, the agent's location at time t $x_i^{(t)}$, the cancer's location x^*, the orientation vector $\theta_i^{(t)}$, the agent-target location vector $\mu_i^{(t)}$, and the angle formed β. Recall that $\phi(t) = ||\mu_i^{(t)}||_2$.

efficiently over time. We investigate the impact of this collective behavior.

2.3. Orientation-Biased Movement Model. We describe the update step for the locomotion of an individual agent, which is the same for the passive and active agent models. Notably, in the active agent model, all agents follow the same movement model regardless of which specific chemical payload they possess. The only difference in the implementation of this update step between the two model variants is the specific signal chemical concentration function $\gamma(\cdot)$ being used, which is defined in Sections 2.1 and 2.2, respectively.

All of the steps traverse the same distance, but there is a bias of moving toward the cancer site. As a result, we call this movement model the "Orientation-Biased Model." Consider some agent i and let

$$\mu_i^{(t)} \coloneqq \tilde{x}^* - \tilde{x}_i^{(t)}$$

be the vector pointing from the current agent location to the target location. Let

$$\phi(t) \coloneqq ||\mu_i^{(t)}||_2$$

be the current Euclidean distance between the agent and the cancer site. Let $\theta_i^{(t)}$ be the orientation vector representing the direction agent i moves in, relative to the cancer site. We define β to be the angle formed by $\theta_i^{(t)}$ and $\mu_i^{(t)}$. We model the noise in the movement direction by drawing $\beta \sim \mathcal{N}(0, \sigma^2)_{[-\pi,\pi)}$, i.e., truncated normal distribution, where $\beta \in [-\pi, \pi)$—in expectation, the agent moves toward the cancer cell. See Fig. 1 for a visual representation of the scalars and vectors involved.

Recall that Table 1 lists all parameters. We assume the underlying normal distribution has variance

$$\sigma^2 = \left(b \cdot \left| \frac{\mathrm{d}}{\mathrm{d}\,\phi(t)} \left(\tilde{\gamma}^{(t)} \left(\tilde{x}_i^{(t)} \right) \right) \right| \right)^{-1}$$

for some constant $b \in \mathbb{R}_{\geq 0}$, where we note that the variance scales inversely with the magnitude of the derivative of $\tilde{\gamma}(\cdot)$ with respect to the distance to the cancer site $\phi(t)$: $\frac{\mathrm{d}}{\mathrm{d}\phi}\tilde{\gamma}(\cdot)$. The steeper the signal chemical gradient, the more biased toward zero β is. Since the agent's movement direction is rotated by β, we can calculate the orientation vector $\theta_i^{(t)}$ as a function of β as follows (representing $\mu_i^{(t)}$ as a column vector): $\theta_i^{(t)} = \begin{bmatrix} \cos\beta & -\sin\beta \\ \sin\beta & \cos\beta \end{bmatrix} \mu_i^{(t)}$. That is, there is always a bias to orient toward the cancer site, with a greater bias as the local chemical gradient is steeper. This "orientation-bias" also increases as b increases. In fact, as b approaches infinity, the movement becomes taking the shortest, straight-line path to $\tilde{x}^*$, and as b approaches zero, the movement converges to the simple random walk.[‡] Agent i's position is then updated as follows, taking a step of displacement/length in the direction of its orientation vector:

$$\tilde{x}_i^{(t)} = \tilde{x}_i^{(t-1)} + \tilde{\alpha} \cdot \frac{\theta_i^{(t)}}{||\theta_i^{(t)}||_2}. \quad [1]$$

For the rest of this paper, we will assume a bounded space, such that agents are always within $\phi_{\max}$ distance from the cancer site x^*, i.e., we can consider the space to be a disk of radius $\phi_{\max}$ centered at x^*.[§] We define $\tilde{\phi}_{\max} = \phi_{\max}/\ell$ to be the dimensionless counterpart, where $\ell = 1$ m.

At each timestep, each agent attempts to follow the above update step until its new location is indeed within the given boundary.[¶] Despite space being bounded, note that $\gamma(\cdot)$ models diffusion in an unbounded space; imposing this $\phi_{\max}$ boundary

Table 1. Summary table of notation for relevant parameters

Notation	Parameter description
n	Total number of nanobot agents
$\varepsilon, \tilde{\varepsilon}$	Cancer site detection distance (w/ and w/o dimension)
$x^*, x_i^{(t)}, x^0$	Locations $\in \mathbb{R}^2$ of cancer site, agent i at time t, and initial site, respectively
$\tilde{x}^*, \tilde{x}_i^{(t)}, \tilde{x}^0$	Dimensionless locations $\in \mathbb{R}^2$ of cancer site, agent i at time t, and initial site, respectively
$\mu_i^{(t)}$	Vector pointing directly toward cancer site $\left(x^* - x_i^{(t)}\right)$
$\phi(t)$	Current distance to cancer site at time t $\|\|\mu_i^{(t)}\|\|_2$
$\theta_i^{(t)}$	Orientation vector $\in \mathbb{R}^2$ of agent i at time t
β	Angle $\in [-\pi, \pi)$ formed between $\mu_i^{(t)}$ and $\theta_i^{(t)}$
$y^{(t)}$	Number of drug payloads dropped up to time t
$z^{(t)}$	Number of signal chemical payloads dropped up to time t (for active agent model)
$\gamma(\cdot), \tilde{\gamma}(\cdot)$	Signal chemical concentration functions (w/ and w/o dimension)
P, t^*	Signal chemical gradient magnitude and steepness, respectively (for passive agent model)
P, t_j^*	Signal chemical payload size and time jth signal chemical payload dropped, respectively (for active agent model)
D	Diffusion coefficient
b	Orientation-bias parameter
$\alpha, \tilde{\alpha}$	Step length (w/ and w/o dimension)
$\phi_{\max}, \tilde{\phi}_{\max}$	Boundary distance (w/ and w/o dimension)
$\delta, d_1, \cdots, d_5$	Quantities used for the analysis of Passive Agent Model

[‡] Hereafter, the "simple," "fully," or "unbiased" random walk in $\mathbb{R}^2$ all always refer to the movement model in which an agent's scalar orientation is sampled uniformly at random, i.e., $\beta \sim \mathcal{U}(-\pi, \pi)$.

[§] A disk is a reasonable model for drug delivery because of the inherent radial symmetry that chemical diffusion models have. Moreover, modeling irregular cancer areas as a circle captures the average spread while keeping calculations tractable, yet accurate.

[¶] For sufficiently small $\tilde{\alpha}$ a valid new location will eventually be achieved with probability 1. We assume that all of these attempts begin from the same starting location, and happen within one timestep.

Downloaded from https://www.pnas.org by 92.41.158.49 on July 17, 2026 from IP address 92.41.158.49.

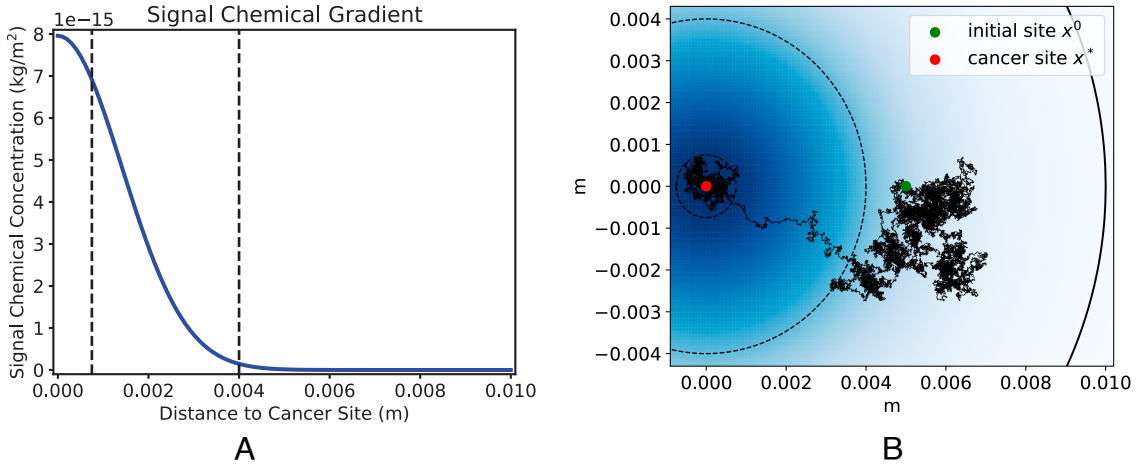

Fig. 2. (Passive agent model). Example simulation run of one agent following our Orientation-Biased movement model with parameters $b = 10^{11}$, $x^0 = (0.005, 0)$m, $\alpha = \varepsilon = 2 \cdot 10^{-5}$ m, and $P = 10^{-19}, D = 10^{-10}, t^* = 10^4, \phi_{\max} = 0.01$ (SI units). (*A*) shows the concentration of signal chemical $\gamma(x)$ plotted as function of distance to cancer site $||x - x^*||_2$, given chemical gradient's radial symmetry. (*B*) shows the agent's position over ≈70,000 timesteps plotted as black trajectory. Signal chemical concentration shown in blue color map. The rightmost solid black curve is $\phi_{\max}$ boundary. Dashed lines in both figures show the same reference distances of 0.00075 m and 0.004 m from the cancer site, respectively, and approximate the notable distances d_2 and d_4 defined in Section 3.1.2.

is a simplification we make simply in order to aid convergence/termination in simulations.

Fig. 2 gives a simulated example of an agent moving under our model. Recall that an agent's orientation-bias, as we call it, is a function of the derivative of the global signal chemical gradient $\gamma(\cdot)$, which is plotted in Fig. 2*A* for this particular example simulation, with respect to the distance to the cancer site. For $||x - x^*||_2 \in [0.004, 0.01]$ m, the derivative is close to zero; for $||x - x^*||_2 \in [0.00075, 0.004]$ m , the derivative is nonnegligible in magnitude; and for $||x - x^*||_2 \in [0, 0.00075]$ m, the derivative again flattens to zero. This is reflected directly in the resulting agent's motion as plotted in Fig. 2*B*: The motion begins as something close to a simple random walk or Brownian motion, becomes more biased/favorable once the agent reaches a distance of around 0.004 m, and finally returns to being mostly fully random once extremely close to the cancer site. This trend in the derivative and resulting behavior is a direct consequence of the form of the signal chemical concentration function $\gamma(\cdot)$ which is very roughly $\exp(-x^2)$ and will always be seen to varying extents regardless of parameter settings.

2.4. Model Validation. Our model, regarding the locomotion of individual agents or nanobots, is based upon the actual nanoscopic vesicles in ref. 21. Therefore, it is likely that some version of our model can be implemented, and can have real-world impact. We now argue that our model resembles the locomotion of the actual nanoparticles. Fig. 3 provides both a quantitative and qualitative comparison between the actual particles and the agents under our model, albeit in a single, specific environment setting.

For sufficiently small windows, or ranges, of distances, the chemical concentration function $\gamma(\cdot)$ can be approximated by a linear function; we make this approximation for our model in Fig. 3. This yields a constant variance σ^2 at all points in time and locations in space according to our Orientation-Biased Model of movement. The experimental data from ref. 21 also involves a chemical signal that changes linearly with respect to distance, making our comparison direct. The histogram(s) on the *Left* of Fig. 3*A* show that the distribution of scalar orientations matches extremely closely between the actual particles in ref. 21 and our model. The *Right* panel of Fig. 3*A* shows the average step sizes (effective velocities) of the agents. In contrast to the actual particles studied in ref. 21, where more favorable orientations correlate with larger average step sizes between 2.8 and 5.5 (abstract) units, our model, for the sake of elegance and forthcoming mathematical analysis, assumes a fixed step size of 4.6 (abstract) units. Fig. 3*B* shows the corresponding agent trajectories in which our model exhibits agent movement which is reasonably similar to the actual particles, with both having agents noisily ascending the external chemical gradient.

3. Results for the Passive Agent Model

We present analytical results for the time required for a single agent to reach the cancer site in Section 3.1. Still looking at the hitting time of a single agent, we present simulation results in Section 3.3. We then consider a swarm with multiple agents and present both analytical and simulation results in Section 3.4. For all of our simulations, we fix $\phi_{\max} = 0.01$m, $\varepsilon = \alpha = 2{\cdot}10^{-5}$ m, and $D = 10^{-10}\text{m}^2/\text{s}$.

3.1. Analytical Results for a Single Passive Agent. We first consider the case of a single agent ($n = 1$). This is a simpler, related problem that gives some intuition for the behavior of the larger swarm/system, as agents' random walks in the passive agent model are independent.

Fig. 4*B* shows the expected value of the progress made in a single step in the distance remaining to the cancer site, under different signal chemical gradients. The exact expected value calculations were derived from Eq. **2**. Across all settings, the expected progress is close to zero from farther away, increases as the agent approaches the cancer site, and then decreases toward zero once more when extremely close. This trend mirrors our discussion of Fig. 2; we are in fact looking at the same phenomenon, just now using more precise analytical results to come to the same conclusion. With larger values of t^*, i.e., as the signal chemical is less steep (or diffuses for longer), agents begin making nonnegligible, positive progress from farther away as the gradient now reaches further, but their absolute maximum expected progress decreases as the gradient

Downloaded from https://www.pnas.org by 92.41.158.49 on July 17, 2026 from IP address 92.41.158.49.

is flattening. Imagining the signal chemical gradient here to actually be dynamic and diffusing, the chemical signal expands but weakens over time.

3.1.1. Expected progress in a single step. We begin by deriving the expected progress made in a single step in Eq. **2**, for use in the analysis to come. For some agent i, with location $\tilde{x}_i^{(t)}$ at time t, we let $\phi\left(\tilde{x}_i^{(t)}\right)$ be its current distance to the cancer site, i.e.,

$$\phi\left(\tilde{x}_i^{(t)}\right) = ||\tilde{x}_i^{(t)} - \tilde{x}^*||_2.$$

When the agent in question is clear from the context, we abuse notation and simply write $\phi(t)$. During timestep t, the agent moves with scalar orientation β relative to the cancer site (Fig. 4*A*). Let $\Delta\phi(t+1) \coloneqq \phi(t) - \phi(t+1)$, i.e., $\Delta\phi(t+1) > 0$ implies the agent made positive progress toward the cancer site during timestep t and is now closer. The Law of Cosines gives

$$\phi(t+1)^2 = \phi(t)^2 + \tilde{\alpha}^2 - 2\phi(t)\tilde{\alpha}\cos(\beta),$$

which yields $\Delta\phi(t+1) = \phi(t) - \sqrt{\phi(t)^2 - 2\phi(t)\tilde{\alpha}\cos(\beta) + \tilde{\alpha}^2}$. We write $\Delta\phi_\beta$ to express the change conditioned on β. Recall, we assume $\beta \sim \mathcal{N}(0, \sigma^2)_{[-\pi,\pi]}$, i.e., truncated normal distribution, where $\beta \in [-\pi, \pi)$ and $\phi = \phi(\tilde{x}_i^{(t)})$, we get

$$\sigma^2 = \sigma_t^2 = \frac{1}{b \cdot \left|\frac{d}{d\,\phi(t)}\tilde{\gamma}\left(\tilde{x}_i^{(t)}\right)\right|} = \frac{1}{b}\frac{\pi(\tau^*)^2}{2\phi}\exp(\phi^2/\tau^*),$$

since $\frac{d}{d\,\phi(t)}\tilde{\gamma}\left(\tilde{x}_i^{(t)}\right) = -\frac{2\phi}{\pi(\tau^*)^2}\exp\left(-\frac{\phi^2}{\tau^*}\right)$. Note that, abusing notation slightly, we have $\lim_{\phi(t)\to 0^+}\sigma(\phi(t)) = \infty$ and $\lim_{\phi(t)\to\infty}\sigma(\phi(t)) = \infty$, but $\sigma(\phi(t))$ attains a global minimum at a finite, positive value of $\phi(t)$. As we will see in our analysis, the consequence of this is that the agent first (as a function of $\phi(t)$) behaves almost like a simple random walk, then like a biased random walk, and finally again like a simple random walk. In Eq. **2**, we derive the expected progress made in a single step in the distance remaining to the cancer site. Recall $\Delta\phi(t+1) = \phi(t) - \phi(t+1)$. By symmetry, we can restrict $\beta \in [0, \pi]$. For some appropriately chosen normalization constant c, we have

$$\begin{aligned}\mathbb{E}[\Delta\phi(t+1)] &= c\int_0^\pi \frac{\Delta\phi_\beta}{\sqrt{2\pi\sigma^2}}\exp\left(-\frac{\beta^2}{2\sigma^2}\right)d\beta \\ &= \frac{c}{\sqrt{2\pi\sigma^2}}\int_0^\pi \left(\phi(t) - \sqrt{\phi(t)^2 - 2\phi(t)\alpha\cos(\beta) + \alpha^2}\right) \\ &\quad \cdot \exp\left(-\frac{\beta^2}{2\sigma^2}\right)d\beta \qquad [2]\end{aligned}$$

Recall from Figs. 2*B* and 4*B* the following general trend in agent motion in the passive agent model: fully random when far away from the cancer site, more biased when approaching, and finally fully random once again when extremely close. We will break up the analysis into three phases, in accordance with the above progression, as illustrated in Fig. 5.

3.1.2. Notable distances. We introduce the key distances $d_5, d_4, \ldots, d_1$ that govern the behavior of an agent. In order to do this, we also pick some fixed constant $\delta \in \mathbb{R}_{>0}$ that is a given desired bias that every agent has in moving toward the cancer site; our key distances will be chosen to then guarantee a bias of δ in the "blue" phase (Fig. 5).[#]

1. $d_5 \coloneqq \tilde{\phi}_{\max}$ is the maximum distance, i.e., $\phi(t) \le d_5$ always holds.
2. d_4 represents the largest distance from which the expected progress toward the cancer site in one step $\mathbb{E}[\Delta\phi]$ is at least δ. In symbols, $d_4 \coloneqq \arg\max_\phi \mathbb{E}[\Delta\phi] \ge \delta$.
3. d_2 represents the smallest distance from which the potential change of Eq. **2** is at least δ throughout $[d_2, d_4]$. In symbols, $d_2 \coloneqq \arg\min_\phi \forall d \in [\phi, d_4] : \mathbb{E}[\Delta\phi(d)] \ge \delta$.
4. $d_3 = \frac{d_2+d_4}{2}$.
5. $d_1 \coloneqq \tilde{\varepsilon}$ is the (maximum) cancer site detection distance, i.e., hitting the ball around $\tilde{x}^*$ of radius d_1 is deemed sufficient for having "reached the cancer site" by our model definition.

A consequence of our definitions is that $\delta \le \tilde{\alpha}$.

We also make the following additional assumptions.

Assumption 3.1. *Assume $\alpha \in (0, 1]$ (therefore, $\delta \le \tilde{\alpha} \le 1$);* $d_3 - d_2 > \frac{4}{\delta^5}$; $\frac{d_3^2 + d_3 d_5\sqrt{2} + d_5^2}{2} \ge 1$; $\frac{d_3 - d_2}{\tilde{\alpha}} \ge 1$

All of our assumptions are for the ease of presentation and, by extending the proofs and statements, all of them can be dropped. We discuss our assumptions at the end of Section 3.2.

We now present the main theorem of this paper, which bounds the time required for a single agent in the passive agent model to reach the cancer site and deliver its treatment. Recall, Let $\phi(t) \coloneqq ||\tilde{x}_i^{(t)} - \tilde{x}^*||_2$.

Theorem 3.2. *Working within the passive agent model, consider a single agent i. Assume $\tilde{\alpha} = \tilde{\varepsilon}$ and that Assumptions 3.1 hold.*

Let T be the random variable denoting the first point in time t for which $\phi(t) \le \tilde{\varepsilon}$, i.e., when the agent reaches the cancer site. Let

$$s^{-1} = \frac{2d_1^2}{\pi\tilde{\alpha}(d_3 - d_2)}\exp\left(-\frac{d_1^2 + d_1 d_2\sqrt{2} + d_2^2}{\tilde{\alpha}(d_3 - d_2)}\right) \qquad [3]$$

and

$$s' = \frac{\exp\left(-1/\delta^3\right)}{1 - e^{-\delta^3/4}}. \qquad [4]$$

We have

$$\begin{aligned}\mathbb{E}[T] \le &\left(1 + \frac{s+1}{s'}\right)\frac{\pi e(d_3^2 + d_3 d_5\sqrt{2} + d_5^2)^2}{2\tilde{\alpha}^2 d_3^2} \\ &+ (s+1)\frac{d_3 - d_2}{\delta} + s\frac{d_3 - d_2}{\tilde{\alpha}}.\end{aligned}$$

3.2. Proof Idea for Theorem 3.2. Without loss of generality, we let $x^0 = (\phi(0), 0) \in \mathbb{R}^2$. By our assumptions, $\phi(0) \le d_5$. For any $d, d' \in \mathbb{R}$ with $d > d' \ge 0$ we define $T_{d,d'}$ as the random variable denoting the first point in time t for which $\phi(t) \le d'$, where $\phi(0) = d$. In particular, let $T = T_{\phi(0),d_1}$ be the time the (passive) agent first reaches the target cancer site. Since $\phi(0) \le d_5$, we then have

$$\mathbb{E}[T] \le \mathbb{E}[T_{d_5,d_3}] + \mathbb{E}[T_{d_3,d_2}] + \mathbb{E}[T_{d_2,d_1}]. \qquad [5]$$

[#] It might well be that δ is too large, resulting in an empty interval $[d_2, d_4]$. However, for a broad range of parameters-particularly when the chemical signal is sufficiently strong-the interval remains nonempty.

Downloaded from https://www.pnas.org by 92.41.158.49 on July 17, 2026 from IP address 92.41.158.49.

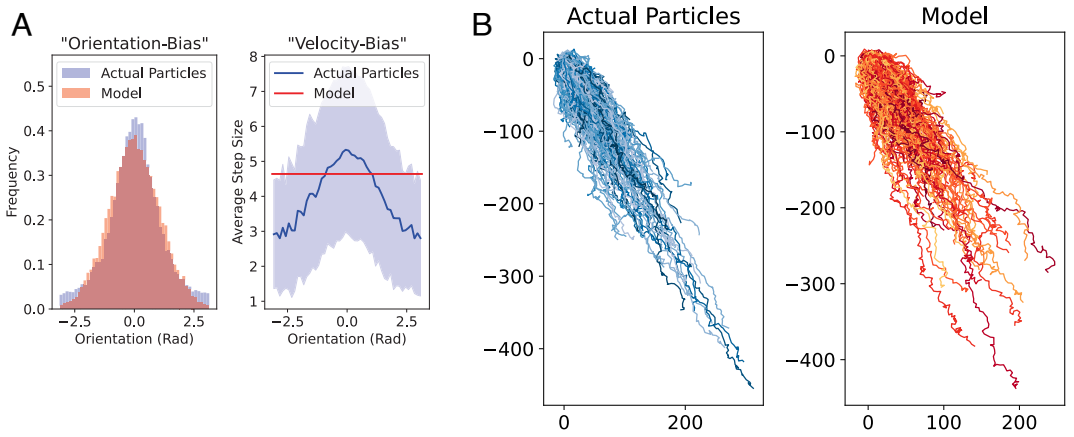

Fig. 3. Comparison between the actual nanoparticles of ref. 21 in vitro experiment and agents under our model in simulation. In both, agents are moving in response to an external chemical gradient which is linear with respect to distance and time-constant. Note that units of distance are arbitrary/abstract here. (*A*)-*Left*: histograms of agent orientations over all timesteps. Orientation of zero points in the same direction as chemical gradient increases. (*B*)-*Right*: average units of distance traveled by agents in one timestep (i.e., step size) over all timesteps. The shaded region is SD. (*B*): single agent trajectories normalized to start at $(0, 0)$. Chemical gradient increases in direction of black arrow.

Each term on the right-hand side can and needs to be bound separately, since the "drift"[‖] of the agent is different in each of these three regimes. We now give an overview of how we treat the analyses of these quantities in the different regimes. See Fig. 5 for an illustration. Our results hinge on a crucial property: The distance of the agent is majorized by the distance of a two-dimensional unbiased random walk. We prove this formally in *SI Appendix*. We now give an overview of the different regimes and their analysis.

1. Between d_5 and d_4, the absolute value of drift of the agent is, by definition of d_4, less than δ. Using the fact that the agent is no slower (majorization) than a slightly biased random walk, we can estimate the time it takes to be within the d_3 using results for unbiased random walks.
2. Between d_3 and d_2, the agent is majorized by a 1D random walk (with respect to ϕ) with a negative bias of $-\delta$. The reason we do not simply assume the random walk for this phase/regime starts at d_4 is that after just one step outside d_4, it would no longer be a $|\delta|$-biased random walk. Thus, we use the regime between d_3 and d_2 and allow the random walk to go all the way to d_4 (which we show will happen very rarely).
3. Between d_2 and d_1, the bias becomes weak again, and we use our majorization with an unbiased random walk that allows for a clean analysis: We are interested in the time it takes to hit the ball at $(0, 0)$ with radius $d_1 = \tilde{\varepsilon}$. Each attempt has a probability of success p. With the remaining probability, the random walk has still failed to hit the ball, but, by construction, it will still be within d_3, and thus, return quickly to d_2. After $1/p$ trials, we expect the random walk to hit the d_1-radius ball.

Proof details appear in *SI Appendix*.

3.2.1. Justification of Assumptions 3.1. As noted above, each assumption can be removed at the cost of a more involved analysis. We retain them for clarity: Without them, the time bounds become regime-dependent in the parameters and harder to interpret.

We now discuss the assumptions one by one. The first one ensures that all steps are at most of size 1 which is useful in our concentration inequalities but can be avoided. The second ensures that the blue regime where the biased random walk governs the process is long enough that the bias matters, in the

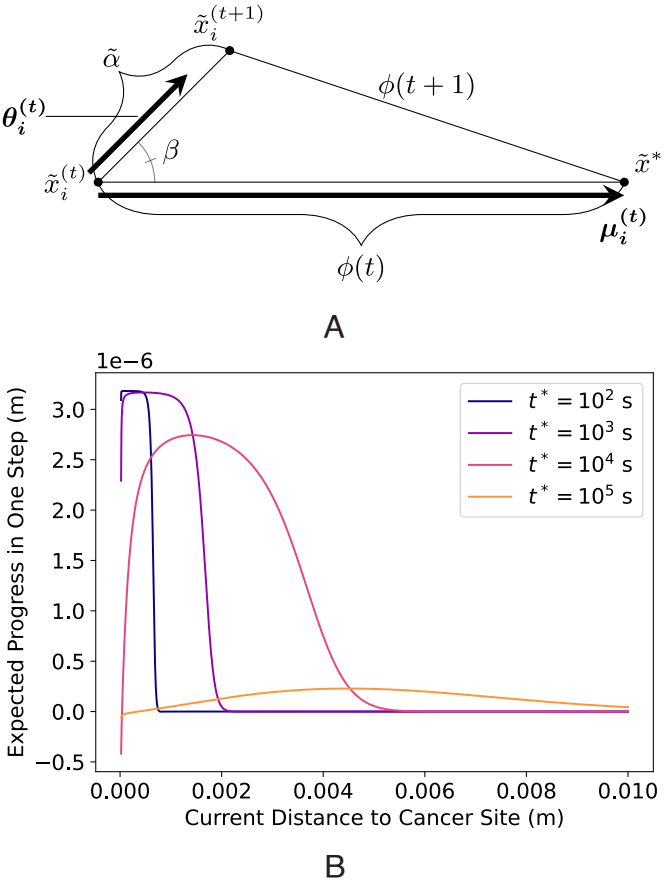

Fig. 4. (*A*)–(Passive agent model). A single step, from time t to time $t+1$, of an agent i in the passive agent model with fixed scalar orientation β and step length α. Recall that $\phi(t) \in \mathbb{R}$ is the distance the agent is away from the cancer site at time t ($\phi(t+1)$ is analogously defined, for time $t+1$). Bold symbols represent vectors, nonbold symbols represent scalar quantities or points. (*B*)–(Passive agent model). Expected progress made in a single step toward the cancer site ($\mathbb{E}[\phi(t) - \phi(t+1)]$) by an agent under our passive agent model from varying current distances, and under different signal chemical gradients. Fixed parameters include $\alpha = \varepsilon = 2 \cdot 10^{-5}, P = 10^{-19}, D = 10^{-10}, b = 10^{12}, \phi_{\max} = 0.01$ (SI units).

[‖]Here, we are considering the one-dimensional random walk with values in $[0, d_5]$ tracking the agent's current distance $\phi(t)$ from the cancer site x^*. So the "drift" or "bias" in this context then denotes the expected change in the distance to x^* in one step. Note that an unbiased random walk in $\mathbb{R}^2$ has a drift away from the cancer cite, unlike in $\mathbb{R}^1$, where there is no drift.

Downloaded from https://www.pnas.org by 92.41.158.49 on July 17, 2026 from IP address 92.41.158.49.

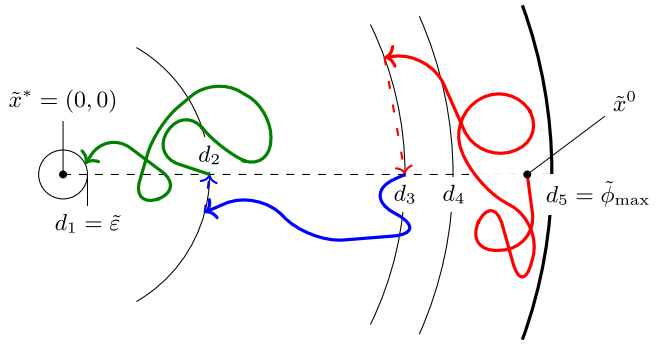

Fig. 5. (Passive agent model). Depiction of space as considered and notated in the proof of Theorem 3.2, with example trajectories for phases/regimes 1, 2, and 3 of the analysis shown in red, blue, and green, respectively. Note that this analysis mirrors the general trend in agent motion in the passive agent model (Figs. 2*B* and 4*B*): "fully" random when far away from the cancer site, more biased as approaching, and finally "fully" random once again when extremely close.

sense that the agent very likely goes in the right direction. Should said regime be much shorter, then the entire process can be treated as an unbiased random walk (resulting in a worse bound on the hitting time). The third ensures that our round length L in the lemma bounding T_{d_5,d_3} (*SI Appendix*) is at least 1 timestep. This can be avoided by using $L = \max\{1, \cdots\}$. The final assumption ensures that the random walk in the lemma bounding $\mathbb{E}[T_{d_2,d_1}]$ (*SI Appendix*) lands in the interval $[d_2, d_3]$, and not farther away, after each (failed) attempt of L' time steps. This assumption can be dropped, by assuming the random walk ends up at d_5 after each attempt (resulting in a worse bound of the hitting time). The final assumption also ensures that our round length L' in the same lemma is at least 1 timestep.

3.3. Simulation Results for a Single Passive Agent. Moving beyond analysis, we now present simulation results for the passive agent model. Since the upper bound presented in *Theorem* 3.2 simply assumes an initial distance of $\phi_{\max}$, the simulation results provide a more accurate assessment of runtime across a range of varying initial distances. In actuality, the agents usually reach the cancer site much faster than the bound presented in *Theorem* 3.2;** The simulation results also better illustrate the effect of certain model parameters on performance.

Still considering only a single agent under the passive agent model with orientation-bias parameter b, Fig. 6 shows the hitting time results from simulation, plotted in comparison to the simple random walk for a baseline of comparison. Recall that larger values of b correspond to agents being more likely to orient and move toward the cancer site (i.e., greater amounts of orientation-bias).

In Fig. 6*A*, without loss of generality—since the signal chemical gradient is radially symmetric—the agent has initial location $x^0 = (\phi(0), 0)$, where $\phi(0)$ is the initial distance between the agent and the cancer site x^*. We fix $P = 10^{-19}$ kg. As $\phi(0)$ increases, hitting times increase, with this decline in performance worsening with less orientation-bias. However, under these specific environment settings $\gamma(\cdot)$, once within a close enough distance ($\approx$0.004 m) and thus a steep enough signal chemical gradient, agents reach the cancer site very efficiently regardless of the amount of bias. Fig. 6*B* shows the effect of

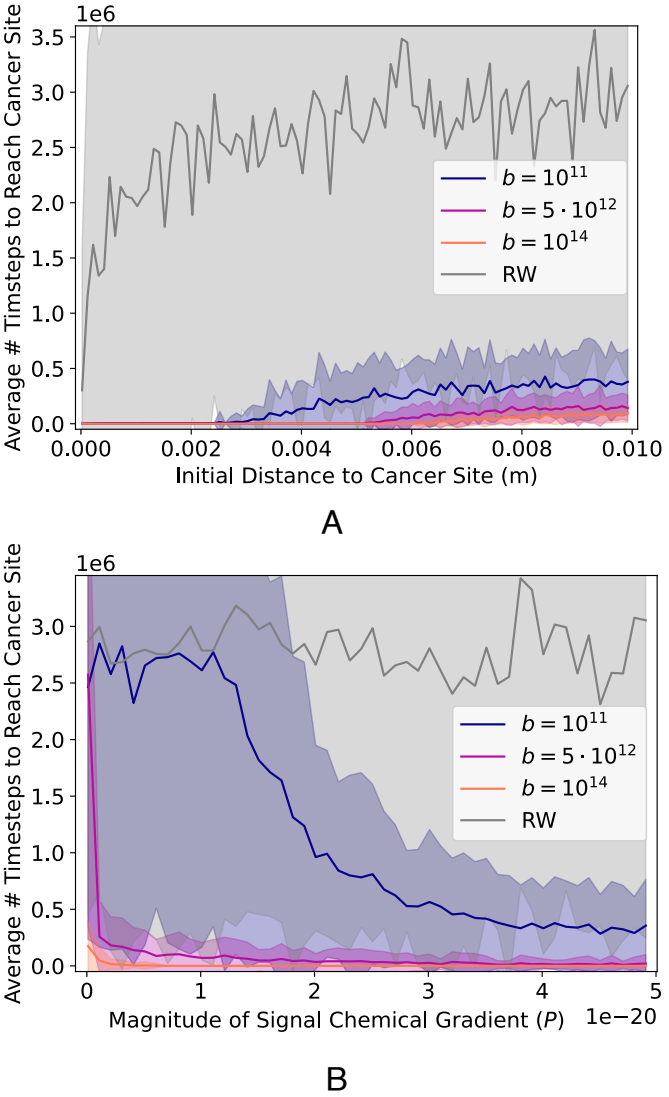

Fig. 6. (Passive agent model). Number of timesteps for a single agent under our passive agent model to reach the cancer site (i.e., hitting time) under varying degrees of orientation-bias, and compared to a simple random walk ("RW"). Plotted lines are average hitting times over 100 trials and shaded regions are SDs. (*A*): How hitting time scales with initial distance to cancer site. Fix signal chemical gradient, $P = 10^{-19}$ kg. (*B*): Effect of the magnitude of the signal chemical gradient on hitting time. Fix initial distance to cancer site at 0.005 m. For both: Fixed parameters include $\alpha = \varepsilon = 2 \cdot 10^{-5}, D = 10^{-10}, t^* = 10^4, \phi_{\max} = 0.01$ (SI units). Hitting times were capped at 10^7 timesteps.

the magnitude of the signal chemical gradient on hitting times. We fix $\phi(0) = 0.005$ m. As P increases, hitting time decreases since there is a more noticeable chemical signal farther out from the cancer site, which the agent senses and consequently orients/moves more favorably in expectation. The effect of P on hitting time is more noticeable under weaker amounts of orientation-bias. As shown in both subfigures, our movement model significantly outperforms the simple random walk across all settings.

3.4. Results for Multiple Passive Agents. We now return to the main application of interest in which at least 75% of the n (>1) nanobot agents are tasked with locating the cancer site in a reasonable amount of time such that all of their marginal drug payloads together provide the desired amount of treatment.

**Even when the initial distance is $\phi_{\max}$, the analysis techniques used for *Theorem* 3.2 likely still produce a bound whose lack of tightness is significant; i.e., we expect agents to usually reach the cancer site much faster the analytical bound, which these hitting times being reflected in simulation results.

Downloaded from https://www.pnas.org by 92.41.158.49 on July 17, 2026 from IP address 92.41.158.49.

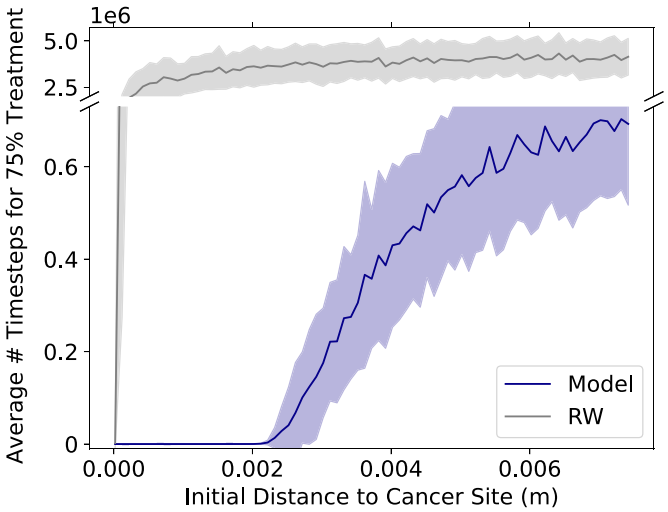

Fig. 7. (Passive agent model). Runtime results for 25 agents in the passive agent model for varying initial distances to the cancer site, and in comparison to the pure RW. Plotted are average runtimes over 100 trials and shaded regions are SDs. Runtimes were capped at 10^7 timesteps. Fixed parameters include $\alpha = \varepsilon = 2 \cdot 10^{-5}, P = 10^{-18}, D = 10^{-10}, t^* = 10^3, \phi_{\max} = 0.01, b = 10^{12}$ (SI units).

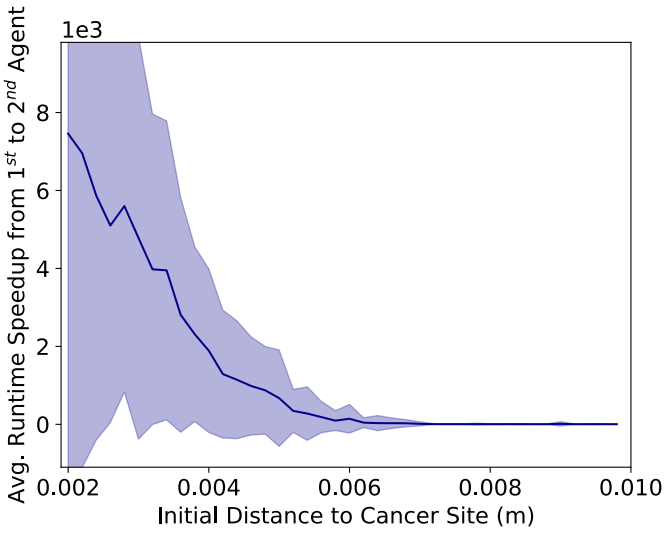

Fig. 8. (Active agent model). Runtime speedup (ratio) from an agent following a simple random walk—"1st agent"—to the first agent with a nonzero chemical signal to follow in the active agent model—"2nd agent." Both agents start from the same initial distance. The signal chemical gradient for this "2nd agent" is given by $z^{(0)} = 1 = z^{(t)}\ \forall t$, $t_1^* = 0$. Fixed parameters include $\alpha = \varepsilon = 2 \cdot 10^{-5}, P = 10^{-19}, D = 10^{-9}, b = 10^{12}, \phi_{\max} = 0.01$ (SI units).

We can easily obtain the following bound, based on *Theorem* 3.2. The proof appears in *SI Appendix*.

Corollary 3.3. *Let $T_1, \ldots, T_n$ be i.i.d. random variables representing the drug delivery times of n agents, and let t^+ be the upper bound derived in Theorem 3.2 such that $\mathbb{E}[T_i] \leq t^+\ \forall i \in [1, n]$. Then, for any number of agents $n \geq 1$, at least 75% of the agents complete delivery by time $10t^+$ with probability at least $1 - 2^{-n/4}$.*

We complement our theoretical result of *Corollary* 3.3 with simulations. Runtime results are shown in Fig. 7. We fix $n = 25$. All n agents have the same initial location $x^0 = (\phi(0), 0)$, where $\phi(0)$ is the initial distance between the agents and the cancer site. Similar to the single agent hitting time results, we see runtimes increasing nonlinearly as the initial distance to the cancer site increases. Again, agents under the passive agent model significantly outperform the simple random walk as expected (by approximately one to two orders of magnitude, depending on the initial distance).

Note that we cannot simply extend our single-agent analytical result to the multiagent case, despite the agents' walks being independent in the passive agent model. That is, given only the expected time for an individual agent to reach the cancer site (and not knowing more about the distribution), we cannot conclude the expected time for multiple agents.[††]

4. Results for the Active Agent Model

Now we present our simulation results for the active agent model. We first present results in Fig. 8 on the runtime speedup seen once the very first signal chemical payload is dropped and there is a nonzero chemical gradient to follow. We then consider the full swarm of n agents and present simulation results in Fig. 9 to show the collective benefit of having more total agents.

Beyond agents dropping payloads in order to have some chemical signal for others to follow and essentially recreate the environment conditions of the passive agent model, one would hope that there would be some additional collective benefit and runtime improvement in having agents be in control of their own dynamic chemical signal—this is indeed what we observe. Breaking down this advantageous collective behavior into contributing factors, we begin by looking at the immediate hitting time improvement seen after the very first signal chemical payload is dropped, shown in Fig. 8. The first agent with a chemical signal to follow reaches the cancer site around three orders of magnitude faster than the simple random walk, proving that the time it takes for the very first agent to drop its signal chemical payload is crucial to the overall runtime.

Fig. 9 shows runtime results with different numbers of agents. The total runtime is effectively the sum of the time it takes for the very first signal chemical payload to be dropped—shown in Fig. 9*A*—and the time it takes for the desired 75% of drug payloads to be dropped once the first signal chemical payload is already dropped—shown in Fig. 9*B*. The former portion of the total runtime is improved with more agents as the time for the first of $\lfloor n/2 \rfloor$ simple random walks to hit the cancer site trivially decreases in expectation as n increases. The latter portion of the total runtime is also improved with more agents as the global signal chemical gradient is amplified, resulting in more favorable agent behavior particularly by those that reach the cancer site later on. Recall Fig. 4*B*, which demonstrated that after long enough, the signal chemical gradient flattens out and starts to become useless for agents to follow. That is, over time, the chemical signal weakens and yields fully random agent motion; however, with more agents, the time between successive dropped payloads decreases which combats this phenomenon, maintaining the chemical gradient at a useful (steep) state, even if it is not

[††]This can be demonstrated by a simple toy example with two agents ($n = 2$). Here, the time for at least 75% of the agents to reach the cancer site is simply the time for both agents to reach the site. Say the expected hitting time for a single agent (in some given environment/model setting) is 1. Consider two distinct possibilities that would both correctly yield this expected hitting time of 1: a) each agent always takes time 1, and b) each agent takes either time 0 or 2, each with probability 1/2. In case a), the total expected hitting time for both agents to reach the cancer site is 1, while in case b), the total expected hitting time for both agents is 11/2. This proves knowing the expected hitting time for a single agent is not, in general, enough knowledge in order to make a statement on the total expected hitting time of multiple agents.

Downloaded from https://www.pnas.org by 92.41.158.49 on July 17, 2026 from IP address 92.41.158.49.

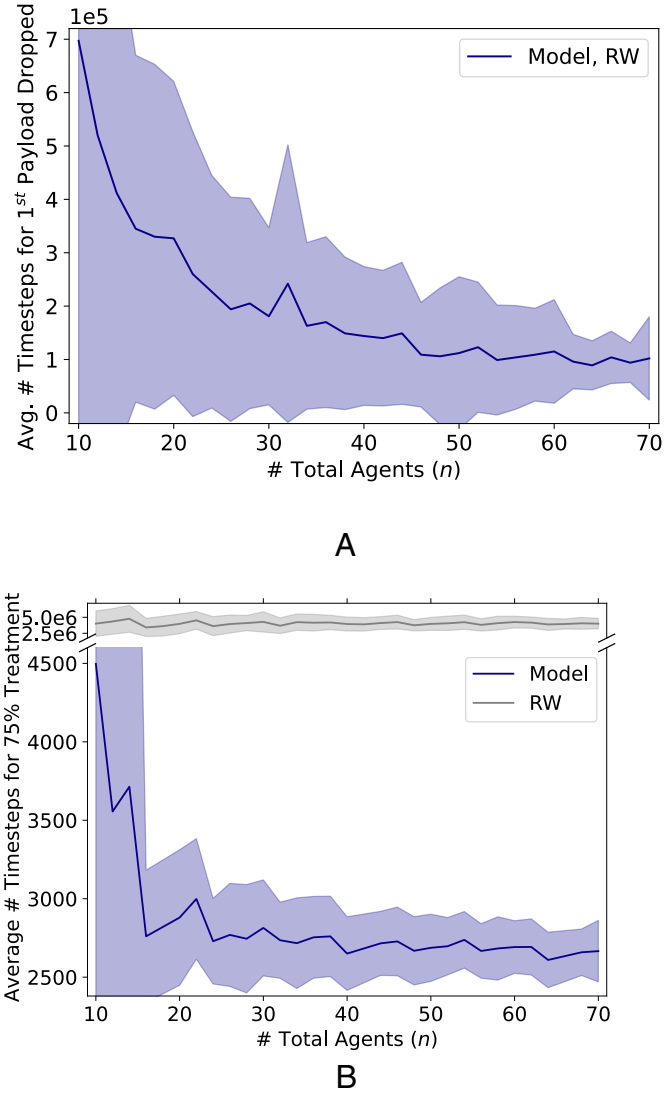

Fig. 9. (Active agent model). Runtime results for the active agent model, showing the collective benefit of a larger swarm, and in comparison to the simple random walk ("RW"). (*A*): Time for first signal chemical payload to be dropped with n total agents, equivalent to the time for first of $\lfloor n/2 \rfloor$ independent simple random walks to hit cancer site. (*B*): Time for 75% of $\lceil n/2 \rceil$ drug-carrying agents to reach cancer site once first signal chemical payload is already dropped, for n total agents. Note that the combined times of (*A* and *B*) is exactly the total runtime for 75% of the agents with drug payloads to deliver their treatment starting from the initial site. Fixed parameters include $\alpha = \varepsilon = 2 \cdot 10^{-5}, P = 10^{-19}, D = 10^{-9}, b = 10^{12}, \phi_{\max} = 0.01$, $\phi(0) = 0.005$ (SI units). Plotted lines are average runtimes over 100 trials and shaded regions are SDs. Runtimes were capped at 10^7 timesteps.

amplified. Ultimately, the active agent model outperforms the simple random walk by approximately three orders of magnitude.

5. Discussion

We present a formal model of nanobots for cancer detection and treatment, with agent locomotion behavior representative of actual nanoparticles. Our results demonstrate that for the Orientation-Biased movement model, agents' movement becomes more biased as they get closer to the cancer site, before returning to fully random movement once they are extremely close. This trend is a result of the specific signal chemical concentration function $\gamma(\cdot)$, not just the movement model in general.

We consider two distinct scenarios, and accompanying variants of the general model: First, the passive agent model in which there already exists an endogenous chemical gradient centered at the targeted cancer site; and second, the active agent model in which agents must form the chemical signal themselves and maintain it over time by dropping payloads at the cancer site. We have simulation results for both versions of the model, showing the behavior of both the individual agents as well as the collective swarm. Across all settings, our model significantly outperforms the fully random walk. We also provide analytical results for the passive agent model, bounding the time it takes for a single agent to reach the cancer site starting from a given initial distance, as well as the time for 75% of the agents to reach the site with high probability. Simulation results for the passive agent model show how runtime scales nonlinearly with the initial distance to the cancer site as well as the extent to which a stronger signal chemical gradient benefits performance. While all walks are independent in the passive agent model, this is not the case in the active agent model.

Simulation results for the active agent model demonstrate a runtime speedup of multiple orders of magnitude with more agents. With more total agents, the time when the chemical signal first appears is drastically reduced as this coincides with the time when the first of $\lfloor n/2 \rfloor$ simple random walks hits the cancer site. Further, as more and more agents release their signal chemical payloads over time, the gradient is able to be maintained at a steep—and thus useful—state before fully diffusing/dispersing and becoming uniform—and thus useless. That is, here, agent movement becomes more biased both when closer to the cancer site and as time progresses. The second agent to reach the cancer site does so faster in expectation than the first agent because of the introduction of a chemical signal to follow, the third agent reaches the site faster in expectation than the second because of an amplified gradient, and so on.

If only in passing, we acknowledge that we did not consider the average clearance time for nanoparticles in the given colloidal environment; for the above mechanism to be indeed useful in practice, the clearance time would need to be sufficiently long such that the later agents have enough time to sense the gradient before being expelled. Although speculative in the assumptions made on individual agent capabilities, we hope that the results shown for the active agent model are compelling enough to motivate experimentalists when designing and testing new nanoparticles.

5.1. Future Work. Future work could vary further environment parameters such as the diffusion coefficient D, individual payload size P (in the active agent model), cancer site detection distance ε, and maximum boundary distance $\phi_{\max}$, and evaluate the resulting performance of the swarm (under both variants of the model). Such an exercise, given interesting simulation results, could motivate new experiments to run with actual nanoparticles. The 3-dimensional setting also remains to be investigated. Future work could include studying the effect of different forms of noise on swarm behavior, such as imperfect detection of the cancer site and agent motion impacted by blood flow.

Future work could also consider multiple cancer sites, introducing the problem of how to evenly allocate the agents' treatment across all sites even if certain sites are more easily located. This setting could call for more involved strategies such as multiple distinct classes of agents, multiple distinct signal chemicals, and different injection times (i.e., different agents introduced into the environment at different points in time).

Downloaded from https://www.pnas.org by 92.41.158.49 on July 17, 2026 from IP address 92.41.158.49.

Last, in the active agent model, it is possible that analytical bounds on the expected time for some percentage of the total agents to reach the cancer site can be derived.

Data, Materials, and Software Availability. Simulation code data have been deposited in GitHub (DOI: 10.5281/zenodo.15278454). (27)

ACKNOWLEDGMENTS. Special thanks to Adithya Balachandran, Sabrina Drammis, and Mien Brabeeba Wang for their suggestions that helped lead our earlier research to the presented work. Also thanks to Sangeeta Bhatia for her suggestions and encouragement. This work was supported by NSF under grants CCR-2139936 and CCR-2003830, and by the UK Engineering and Physical Sciences Research Council under grant EP/W005573/1.

Author affiliations: [a]Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139; [b]Department of Informatics, King's College London, London WC2R 2LS, United Kingdom; and [c]Department of Bioengineering, Imperial College London, London SW7 2AZ, United Kingdom

1. I. Brigger, C. Dubernet, P. Couvreur, Nanoparticles in cancer therapy and diagnosis. *Adv. Drug Deliv. Rev.* **64**, 24–36 (2012).
2. D. Zhang *et al*., Advanced medical micro-robotics for early diagnosis and therapeutic interventions. *Front. Robot. AI* **9**, 1086043 (2023).
3. A. Ghosh *et al*., Stimuli-responsive soft untethered grippers for drug delivery and robotic surgery. *Front. Mech. Eng.* **3**, 7 (2017).
4. B. Crandall, *Nanotechnology: Molecular Speculations on Global Abundance* (Mit Press, 1996).
5. R. A. Freitas, *Nanomedicine, Volume I: Basic Capabilities* (Landes Bioscience Georgetown, TX, 1999), vol. 1.
6. A. T. Liu *et al*., Colloidal robotics. *Nat. Mater.* **22**, 1453–1462 (2023).
7. K. Kostarelos, Nanorobots for medicine: How close are we? *Nanomedicine* **5**, 341–342 (2010).
8. J. T. Gómez *et al*., "Markov model for the flow of nanobots in the human circulatory system" in *Proceedings of the Eight Annual ACM International Conference on Nanoscale Computing and Communication*, L. Galluccio, U. Mitra, *Eds*. (Association for Computing Machinery (ACM), New York, NY, 2021), pp. 1–7.
9. T. Gwisai *et al*., Magnetic torque-driven living microrobots for increased tumor infiltration. *Sci. Robot.* **7**, eabo0665 (2022).
10. S. Martel, M. Mohammadi, O. Felfoul, Z. Lu, P. Pouponneau, Flagellated magnetotactic bacteria as controlled MRI-trackable propulsion and steering systems for medical nanorobots operating in the human microvasculature. *Int. J. Robot. Res.* **28**, 571–582 (2009).
11. A. Clementi, F. d'Amore, G. Giakkoupis, E. Natale, "Search via parallel lévy walks on $mathbbZ^2$" in *Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing* (, A. Miller, K. Censor-Hillel, *Eds*. (Association for Computing Machinery (ACM), New York, NY, 2021), pp. 81–91.
12. R. Fujisawa, S. Dobata, "Lévy walk enhances efficiency of group foraging in pheromone communicating swarm robots" in Proceedings of the 2013 IEEE/SICE International Symposiumon System Integration , H. Kobayashi, *Ed*. (IEEE, 2013).
13. A. M. Reynolds, C. J. Rhodes, The lévy flight paradigm: Random search patterns and mechanisms. *Ecology* **90**, 877–887 (2009).
14. E. T. Roussos, J. S. Condeelis, A. Patsialou, Chemotaxis in cancer. *Nat. Rev. Cancer* **11**, 573–587 (2011).
15. M. J. Oudin, V. M. Weaver, "Physical and chemical gradients in the tumor microenvironment regulate tumor cell invasion, migration, and metastasis" in *Cold Spring Harbor Symposia on Quantitative Biology*, J. Darnell, *Ed*. (Cold Spring Harbor Laboratory Press, 2016), vol. 81, pp. 189–205.
16. A. J. Muinonen-Martin *et al*., Melanoma cells break down LPA to establish local gradients that drive chemotactic dispersal. *PLoS Biol.* **12**, e1001966 (2014).
17. F. Peng, Y. Tu, J. C. Van Hest, D. A. Wilson, Self-guided supramolecular cargo-loaded nanomotors with chemotactic behavior towards cells. *Angew. Chem. Int. Ed.* **54**, 11662–11665 (2015).
18. J. Wang *et al*., Self-propelled PLGA micromotor with chemotactic response to inflammation. *Adv. Healthc. Mater.* **9**, 1901710 (2020).
19. R. Golestanian, T. B. Liverpool, A. Ajdari, Propulsion of a molecular machine by asymmetric distribution of reaction products. *Phys. Rev. Lett.* **94**, 220801 (2005).
20. J. R. Howse *et al*., Self-motile colloidal particles: From directed propulsion to random walk. *Phys. Rev. Lett.* **99**, 048102 (2007).
21. A. Joseph *et al*., Chemotactic synthetic vesicles: Design and applications in blood-brain barrier crossing. *Sci. Adv.* **3**, e1700362 (2017).
22. M. Serra-Casablancas *et al*., Catalase-powered nanobots for overcoming the mucus barrier. *ACS Nano* **18**, 16701–16714 (2024).
23. D. Xu *et al*., Enzyme-powered liquid metal nanobots endowed with multiple biomedical functions. *ACS Nano* **15**, 11543–11554 (2021).
24. D. Simberg *et al*., Biomimetic amplification of nanoparticle homing to tumors. *Proc. Natl. Acad. Sci. U.S.A.* **104**, 932–936 (2007).
25. G. von Maltzahn *et al*., Nanoparticles that communicate in vivo to amplify tumour targeting. *Nat. Mater.* **10**, 1476–4660 (2011).
26. J. Shi, *Diffusion of Point Source and Biological Dispersal, Notes* (William & Mary, Math 490–01 Partial Differential Equations and Mathematical Biology, 2006).
27. N. Harasha, C. Gava, N. Lynch, C. Contini, F. Mallmann-Trenn, Simulation code for "Modeling feasible locomotion of nanobots for cancer detection and treatment." *Zenodo*. https://doi.org/10.5281/zenodo.15278454. Deposited 25 April 2025.

Downloaded from https://www.pnas.org by 92.41.158.49 on July 17, 2026 from IP address 92.41.158.49.