

Nanobot Algorithms for Treatment of Diffuse Cancer

Extended Abstract

Noble Harasha

Massachusetts Institute of Technology
Cambridge, USA
nharasha@mit.edu

Nancy Lynch

Massachusetts Institute of Technology
Cambridge, USA
lynch@csail.mit.edu

ABSTRACT

We consider the problem of a swarm of *nanobots* detecting and treating human cancer that is *diffuse*, dispersed in a region with multiple separate cancer sites in need of treatment. We present a mathematical model of nanobots and their colloidal environment that is inspired by actual chemotactic nanoparticles, involving agents noisily following chemical gradients (both attractively and repellingly, depending on the chemical). We present three incrementally sophisticated algorithms that describe additional chemical payloads that agents carry onboard, beyond the cancer-treating drug K , as well as the rules for when agents drop their payloads: *Algorithm KM*, in which agents simply ascend naturally existing chemical M signals that surround cancer sites; *KMA*, in which agents themselves amplify these natural signals by dropping chemical A payloads upon reaching a site; and *KMAR*, in which agents choose to either amplify the signal by dropping chemical A or counteract/reduce the signal by dropping chemical R , according to the current unsatisfied *demand* of the site. We present simulation results for all of the algorithms, across a set of distinct cancer arrangements, that track both the achieved treatment success as well as the time/duration of the treatment. *KM* has generally successful treatment unless the natural M -signals are weak, in which case the treatment progresses too slowly. *KMA* demonstrates a significant speedup in treatment time (over *KM*), but also a drop in success except for the most concentrated cancer patterns. *KMAR* has relatively optimal performance across all types of cancer patterns, demonstrating robustness and adaptability in its mechanisms for nanobot coordination.

KEYWORDS

Nanobots; Biological Modeling; Distributed Algorithms

ACM Reference Format:

Noble Harasha and Nancy Lynch. 2026. Nanobot Algorithms for Treatment of Diffuse Cancer: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 3 pages. <https://doi.org/10.65109/WHCQ7191>

1 INTRODUCTION

Motile nanoparticles suspended in a solution, or *nanobots*, possess unique potential in various medical applications [2, 3]; however, their small size yields limited capabilities in sensing, locomotion,

and computation. Nanobots are being engineered to navigate within the human body to locate sites of interest and deliver drugs in a more targeted manner with fewer side effects [1, 9, 17]; we investigate this targeted drug delivery problem here, specifically for cancer. We consider a stochastic process, with nanobots performing *chemotaxis* [4, 7, 8, 11, 16] that we model as a biased random walk. This work builds on our previous work [5], which defined a mathematical model to characterize the movement of the nanoparticles in [8] (similar to Algorithm *KM* here) and included a signal amplification mechanism similar to [13, 15] and *KMA* here. We extend [5] to consider the *multi-site* case—not only a single target site—as well as to include negative—not only positive—chemotaxis [10, 14] (in *KMAR*). Having multiple, distinct cancerous sites introduces the problem of optimally allocating treatment according to *demands*.

Results show that *KM* has successful treatment but is slow for weak endogenous chemical signals. *KMA* significantly improves efficiency over *KM* with its amplified attractive signals, but can be unsuccessful if the amplification mechanism too strongly favors one site, leaving the others untreated (this is a reduced issue when one site has most of the demand). *KMAR*'s added mechanism of dynamically favoring chemical signals towards sites still needing treatment helps fix *KMA*'s above issues, while still being efficient.

The full version of this paper is available online [6].

2 MODEL

We present a continuous space \mathbb{R}^2 , discrete time model for multi-site cancer detection and treatment by nanobots in the human body. There are a number of *cancer sites*, i.e., distinct clusters of contiguous cancerous cells, each concentrated at a single point in \mathbb{R}^2 . An agent can detect the presence of a nearby cancer site once it is within an ϵ -distance, and plant there. Each agent carries a payload of a drug for cancer treatment, *chemical K*, that it drops upon planting. *Signal chemicals* are sensed by all agents and their gradients can be followed directionally: ascended if *attractive*, descended if *repellent*. No direct interaction or communication occurs between agents.

Centered at each cancer site is an endogenous gradient of attractive signal chemical M , that is persistent/time-constant. The strength of a site's M -signal is directly proportional to its *demand* P_{M_j} , i.e., the amount of treatment it requires. One agent is needed per unit of demand. The concentration of M at position x is $\gamma_M(x) := \frac{10^6}{\pi} \sum_j P_{M_j} \exp(-10^6(\|y_j - x\|_2)^2)$, where y_j is cancer site j 's location. The other signal chemicals we consider are both artificial chemicals that are carried and electively dropped as payloads: an attractive signal chemical A , and a repellent R . Unlike for chemical M , A and R are not persistent; they dynamically dissipate and diffuse over time, which we model via instantaneous point-source diffusion [12]. The concentration of A at time t at position



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/WHCQ7191>




Figure 1: Simulation results with 55 agents, comparing algorithms’ performance for different cancer arrangements (depicted by small scatterplots). Fixed: $\alpha = \epsilon = 2 \cdot 10^{-5}$, $\phi_{max} = 0.005$, $P_A = 10$, $P_R = 50$, $r_{A,M} = 10^7$. Plotted lines show average success over 20 trials; shaded regions are standard deviations. Points on main plots show when treatment stabilizes/finishes (treatment time).

x is $\gamma_A^{(t)}(x) := \frac{10^9 P_A}{4\pi} \sum_j \sum_i (\frac{1}{t-t_{j,i}^*} \exp(-10^9 \|y_j - x\|^2 / (4(t-t_{j,i}^*))))$, where P_A is the payload size and $t_{j,i}^*$ is when the i ’th A-payload was dropped at site j . $\gamma_R^{(t)}(x)$ is defined analogously.

We now describe the update step for the locomotion of an individual agent, i.e., the movement model. Consider agent i at time t , with position $x_i^{(t)}$ and orientation vector $\theta_i^{(t)}$. Let $\gamma_{TOT}^{(t)}(x_i^{(t)}) := \gamma_M(x_i^{(t)}) + \gamma_A^{(t)}(x_i^{(t)}) - \gamma_R^{(t)}(x_i^{(t)})$. Let $\mu := \nabla \gamma_{TOT}^{(t)}(x_i^{(t)})$, and let $\beta \sim \mathcal{N}(0, \sigma^2)$ with variance $\sigma^2 = 10^{11} / (b \|\mu\|)$. If $\gamma_{TOT}^{(t)}(x_i^{(t)}) = 0$, $\beta \sim U(-\pi, \pi)$ and $\mu = (1, 0)$. $\theta_i^{(t+1)}$ is updated such that β is the angle formed between μ and $\theta_i^{(t+1)}$. Larger b values yield movement that is more biased toward following the given chemical signal(s), i.e., farther from standard Brownian motion. Taking a step of length α in the direction of its orientation vector, agent i ’s position is updated as $x_i^{(t+1)} = x_i^{(t)} + \alpha(\theta_i^{(t+1)} / \|\theta_i^{(t+1)}\|_2)$. To summarize, if the attractive chemical signal (M plus A) dominates the repellent signal (R) in steepness, then the agent is biased to move *towards* the nearest local maximum of the attractive chemical gradient (e.g., the nearest cancer site). If the repellent signal dominates, the agent is biased to move (roughly) *away* from the nearest cancer site.

We assume a bounded space $[0, \phi_{max}]^2 \subset \mathbb{R}^2$; the above update step is repeated until a valid new location is produced. Nanobots have a finite lifetime; we define a runtime cutoff/clearance time T^* .

3 ALGORITHMS

The least sophisticated algorithm KM, in which agents simply ascend the stable, natural M-signals, is feasible and aligned with current nanobot technologies. From KM to KMA, in an effort to improve treatment time, we add A-payloads to amplify the M-signals. From KMA to KMAR, in an effort to improve success, we also add R-payloads which uniquely allow for agents to descend chemical gradients and thus, be driven away from sites that are already treated. We acknowledge that each algorithm (in that order) is more speculative than the next regarding individual nanobot capabilities. An agent can only drop its payload(s) at one cancer site in total.

Algorithm KM: Agents only have payloads of chemical K. When an agent reaches a cancer site, it releases its K-payload to deliver treatment. There is no chemical A nor R anywhere.

Algorithm KMA: Each agent has a payload of chemical K and a payload of chemical A. When an agent reaches a site, it releases both payloads. There is no chemical R anywhere.

Algorithm KMAR: Each agent has one payload of each chemical, K, A, and R. When an agent reaches some site j at time t^* , it immediately releases its K-payload. Then, if $\gamma_A^{(t^*)}(y_j) / P_{M_j} < r_{A,M}$, it also releases its A-payload, but does not release its R-payload. Otherwise, it instead releases its R-payload, but not its A-payload. That is, when the A-signal at a site is too strong, agents release R in an effort to encourage agents to explore and administer treatment elsewhere; chemical A is used here as a proxy to *estimate* the current amount of treatment (K) already administered.

4 SIMULATION RESULTS

We fix a specific parameter setting for each algorithm (chosen based upon preliminary results [6]) and compare their performance across a set of distinct cancer site and demand arrangements; see Figure 1. Agents have uniformly random initial positions.

For the more diffuse cancer arrangements (a) and (b), KMAR and KM achieve the same, highest eventual success rate, though KMAR is faster and thus outperforms KM for many clearance times. KMA has lower achieved success than KM (and KMAR), but, because of its fast progressing treatment, outperforms KM for some clearance times. The sparse (a) yields higher overall treatment success for KM and KMAR than the dense (b). For the concentrated cancer arrangement (c), all algorithms perform similarly very well in both success and treatment time, showing the largest improvement over the simple random walk RW. Although KM achieves high success here for all arrangements, it has increasingly slow progressing treatment under weaker M-signals (a feature of the natural environment). We consider only a single pass of treatment here, but having multiple, repeated passes is also reasonable given nanobots’ nontoxicity. In this case, KMA has increased usefulness as it quickly and reliably kills the main cancer site (with the leftover(s) handled in later passes) via a less sophisticated algorithm than KMAR.

ACKNOWLEDGMENTS

Work is supported by NSF Awards CCR-2003830 and CCR-2139936.

REFERENCES

- [1] Irène Brigger, Catherine Dubernet, and Patrick Couvreur. 2012. Nanoparticles in cancer therapy and diagnosis. *Advanced drug delivery reviews* 64 (2012), 24–36.
- [2] BC Crandall. 1996. *Nanotechnology: molecular speculations on global abundance*. MIT Press.
- [3] Robert A Freitas. 1999. *Nanomedicine, volume I: basic capabilities*. Vol. 1. Landes Bioscience Georgetown, TX.
- [4] Ramin Golestanian, Tanniemola B Liverpool, and Armand Ajdari. 2005. Propulsion of a molecular machine by asymmetric distribution of reaction products. *Physical review letters* 94, 22 (2005), 220801.
- [5] Noble Harasha, Cristina Gava, Nancy Lynch, Claudia Contini, and Frederik Mallmann-Trenn. 2025. Modeling feasible locomotion of nanobots for cancer detection and treatment. *Proceedings of the National Academy of Sciences (PNAS)* 122, 48 (2025). <https://doi.org/10.1073/pnas.2510036122>
- [6] Noble Harasha and Nancy Lynch. 2025. Nanobot Algorithms for Treatment of Diffuse Cancer. arXiv:2509.06893 [cs.MA] <https://arxiv.org/abs/2509.06893>
- [7] Jonathan R Howse, Richard AL Jones, Anthony J Ryan, Tim Gough, Reza Vafabakhsh, and Ramin Golestanian. 2007. Self-motile colloidal particles: from directed propulsion to random walk. *Physical review letters* 99, 4 (2007), 048102.
- [8] Adrian Joseph, Claudia Contini, Denis Cecchin, Sophie Nyberg, Lorena Ruiz-Perez, Jens Gaitzsch, Gavin Fullstone, Xiaohe Tian, Juzaili Azizi, Jane Preston, Giorgio Volpe, and Giuseppe Battaglia. 2017. Chemotactic synthetic vesicles: Design and applications in blood-brain barrier crossing. *Science Advances* 3, 8 (2017).
- [9] Kostas Kostarelos. 2010. Nanorobots for medicine: how close are we? *Nanomedicine* 5, 3 (2010), 341–342.
- [10] Mihail N. Popescu, William E. Uspal, Clemens Bechinger, and Peer Fischer. 2018. Chemotaxis of Active Janus Nanoparticles. *Nano Letters* 18, 9 (2018), 5345–5349.
- [11] Meritxell Serra-Casablanca, Valerio Di Carlo, David Esporrin-Ubieto, Carles Prado-Morales, Anna C. Bakenecker, and Samuel Sánchez. 2024. Catalase-Powered Nanobots for Overcoming the Mucus Barrier. *ACS Nano* 18, 26 (2024), 16701–16714. <https://doi.org/10.1021/acsnano.4c01760> arXiv:<https://doi.org/10.1021/acsnano.4c01760>
- [12] Junping Shi. 2006. Diffusion of point source and biological dispersal, Notes. William & Mary, Math 490-01 Partial Differential Equations and Mathematical Biology.
- [13] Dmitri Simberg, Tasmia Duza, Ji Ho Park, Markus Essler, Jan Pilch, Lianglin Zhang, Austin M. Derfus, Meng Yang, Robert M. Hoffman, Sangeeta Bhatia, Michael J. Sailor, and Erkki Ruoslahti. 2007. Biomimetic amplification of nanoparticle homing to tumors. *Proceedings of the National Academy of Sciences* 104, 3 (2007), 932–936. <https://doi.org/10.1073/pnas.0610298104> arXiv:<https://www.pnas.org/doi/pdf/10.1073/pnas.0610298104>
- [14] Ambika Somasundar, Subhadip Ghosh, Farzad Mohajerani, Lynnica N. Massenburg, Tinglu Yang, Paul S. Cremer, Darrell Velegol, and Ayusman Sen. 2019. Positive and negative chemotaxis of enzyme-coated liposome motors. *Nature Nanotechnology* 14 (2019), 1129–1134.
- [15] Geoffrey von Maltzahn, Ji-Ho Park, Kevin Y. Lin, Neetu Singh, Christian Schwöppe, Rolf Mesters, Wolfgang E. Berdel, Erkki Ruoslahti, Michael J. Sailor, and Sangeeta N. Bhatia. 2011. Nanoparticles that communicate in vivo to amplify tumour targeting. *Nature Materials* 10, 7 (2011), 1476–4660. <https://doi.org/10.1038/nmat3049>
- [16] Dandan Xu, Jing Hu, Xi Pan, Samuel Sánchez, Xiaohui Yan, and Xing Ma. 2021. Enzyme-Powered Liquid Metal Nanobots Endowed with Multiple Biomedical Functions. *ACS Nano* 15, 7 (2021), 11543–11554. <https://doi.org/10.1021/acsnano.1c01573> arXiv:<https://doi.org/10.1021/acsnano.1c01573>
- [17] Dandan Zhang, Thomas E Gorochowski, Lucia Marucci, Hyun-Taek Lee, Bruno Gil, Bing Li, Sabine Hauert, and Eric Yeatman. 2023. Advanced medical micro-robotics for early diagnosis and therapeutic interventions. *Frontiers in Robotics and AI* 9 (2023), 1086043.