

UCLA

Technical Reports

Title

Bacterium-inspired Robots for Environmental Monitoring

Permalink

<https://escholarship.org/uc/item/5085h7qt>

Authors

Amit Dhariwal
Gaurav S. Sukhatme
Aristides A. Requicha

Publication Date

2003

Peer reviewed

Bacterium-inspired Robots for Environmental Monitoring

Amit Dhariwal, Gaurav S. Sukhatme and Aristides A. G. Requicha
Department of Computer Science, Center for Robotics and Embedded Systems
University of Southern California, Los Angeles, California, USA
{dhariwal, gaurav, requicha}@usc.edu

Abstract—Locating gradient sources and tracking them over time has important applications to environmental monitoring and studies of the ecosystem. We present an approach, inspired by bacterial chemotaxis, for robots to navigate to sources using gradient measurements and a simple actuation strategy (biasing a random walk). Extensive simulations show the efficacy of the approach in varied conditions including multiple sources, dissipative sources, and noisy sensors and actuators. We also show how such an approach could be used for boundary finding. We validate our approach by testing it on a small robot (the robomote) in a phototaxis experiment. A comparison of our approach with gradient descent shows that while gradient descent is faster, our approach is better suited for boundary coverage, and performs better in the presence of multiple and dissipative sources.

I. INTRODUCTION

Several phenomena in nature induce gradients in their environment. For example, a fire induces a temperature gradient in its vicinity, an oil spill induces a concentration gradient of oil in the water etc. Detection, seeking and tracking such phenomena in-situ has received some attention recently [1], [2], [3]. The ability to autonomously detect, locate and track such phenomena (the source of the induced gradient) would give scientists a tool to monitor and study ecosystems at an unprecedented level of detail. Typical gradients of interest to scientists studying the ecosystem include temperature, light, salinity, mineral concentration, pH, etc. These problems are difficult because of the time-varying nature of the source, the dynamics of the environment, a multiplicity of (possibly interacting) sources, and finally, a paucity of sensing.

Motivated by these applications and challenges, we are interested in the development of simple, robust, energy efficient and cost-effective techniques which could be used in-situ to locate source phenomena of interest to scientists. We are motivated by a particular example from marine biology; the detection of a mineral pollutant in seawater. In this paper we focus on a 2D version of this problem, and propose a simple strategy for a mobile robot (or multiple robots) to navigate their way to such a source using gradient information and extremely rudimentary actuation. Our strategy is inspired by studies of taxis in bacteria.

Previous approaches include spiral surge [1], gradient seek [4], sensor arrays [5], swarm intelligence [6]. These approaches are good for application in nearly static environments. Most of these approaches have difficulties tracking

weaker and smaller sources, and sources which have a lot of variation. The gradient seeking strategies are susceptible to local minima and plateaus. In the presence of multiple sources they are able to locate the closest source since they enact a greedy solution.

Another related area is the detection of dynamically changing gradient source boundaries in-situ in a distributed manner [4], [5], [7]. This process involves locating the periphery of the source, a small region which has the sharpest gradient in a region close to the source. A brute force technique is to compute the gradient at every point in the search space, but such a technique is not efficient in terms of energy, time requirements and the amount of processing involved.

In this paper, we present a novel technique based on biased random walk [8] for the detection, seeking and tracking of gradient inducing source phenomena. Our approach is inspired by the way bacteria detect, locate and track nutrient sources in nature [9]. We begin by a discussion of the characteristics of bacterial motion and how it can be adapted and applied to the problem at hand. In section III we describe the simulation platform we created to evaluate our strategy along with a discussion of the results obtained from simulations. Section IV presents an implementation of our approach on the Robomote robot platform and a discussion of the results obtained and their applicability. In section V, we present a comparison between the behavior of a gradient search strategy and the behavior of a biased random walk model. This is followed by our conclusions and a brief comment on future work we plan to pursue.

II. BACTERIAL MOTION

A. Biased Random Walk as a result of Chemotaxis

Nature presents us with a wide variety of simple biological models which have evolved and refined over time. These models help in sustaining the flora and fauna on our planet and maintaining the ecological balance in diverse habitats. Bacterial motion [9] and its response to the presence of chemical concentration gradients called chemotaxis (or chemokinesis) have been well studied [10], [11], [12], [13]. The response to a chemical stimulus in their vicinity helps bacteria find sources of nutrients which are essential for their survival. Chemotaxis is also observed in various other species of animals for varied purposes including colony formation, predator avoidance and breeding ground location [9].

Bacteria sense chemical concentration using receptors. They are able to detect temporal and spatial changes in chemical concentration based on the fraction of receptors occupied at successive time intervals. An increase in the fraction of occupied receptors is called a positive gradient while a decrease is called a negative gradient. A chemical whose concentration gradient attracts the bacterial cells is called a chemo-attractant.

Bacteria produce motion by the movement of their flagellum [14]. A counter clockwise flagellar rotation results in a smooth swim motion in a straight line in a particular direction (we call this a run) while a clockwise rotation of the flagellum causes the bacterium to randomly reorient itself in a new direction (we call this a tumble), which is the direction for the next run. Motion alternates between these two stages (run and tumble).

The duration of the run (which is related to the mean free path) is dependent on the concentration gradient that is sensed in the vicinity of the bacterial cell. In the absence of a gradient, the run length is independent of the direction of motion and the bacterium executes a random walk. In the presence of a positive gradient, the frequency of tumbling is reduced resulting in a longer run length [12]. The presence of a negative gradient does not have any effect on the tumbling frequency. This change of tumbling frequency in response to concentration gradient results in chemotaxis, allowing bacteria to move towards sources of nutrients. Informally, chemotaxis is a biased random walk.

B. Analysis of Bacterial Motion

A number of mathematical models have been proposed to model bacterial chemotaxis [8], [15], [16] [17], [18]. In this section we briefly review the reaction-diffusion model proposed by Keller and Segel [16].

$$J_{chemotaxis} = \chi(c)n\nabla(c) \quad (1)$$

(1) models the flux up(or down) the gradient of chemical concentration, with the flux increasing with n . $\chi(c)$ is the chemotactic sensitivity and represents the specific form of cells response to chemical signals.

Adding the diffusive term to the model, we get the cell kinetics model as

$$\frac{\delta n}{\delta t} = \nabla[D_n \nabla(n) - \chi(c) n \nabla(c)] + f(n) \quad (2)$$

$$\frac{\delta c}{\delta t} = \nabla[D_c \nabla(c) + g(n, c)] \quad (3)$$

where D_n is the diffusion coefficient, $c(x,t)$ is the density of the nutrients in the x direction, $f(n)$ represents the cell division and death and $g(n,c)$ represents the production and degradation of the chemical. Keller and Segel used $g(n,c) = an - rc$ and $f(n)=0$ with a, r, D_n, D_c , and χ as constant. (2) is also known as the parabolic chemotactic equation. The first term models the random diffusive behavior of the bacterial motion while the second term describes the biased random walk motion in response to the chemical gradient.

Ignoring the cell growth term, $f(n) = 0$, we get

$$\frac{\delta n}{\delta t} = \nabla[D_n \nabla(n) - \chi(c) n \nabla(c)] \quad (4)$$

As can be seen from (4) switching the direction of chemical gradient can produce both forward and reverse movement of the cells. Solving (3) and (4) gives the velocity of the bacterial band as

$$\nu = nk/ac_{initial} \quad (5)$$

where $c_{initial}$ is the initial nutrient concentration, n is the number of bacteria in the band and a is the cross section area of the band.

C. A Robotic Implementation

Based on the description of the bacterial motion and its mathematical analysis presented in the last two subsections, it is clear that a strategy based on biased random walk could be used to locate and track gradient sources. This biologically-inspired algorithm can be implemented on a group of robots with simple sensing and actuation. The strategy of such a bacteria-like robot can be summarized as "sense and move". A robotic node executing a biased random walk has very little requirements in terms of memory since only the last sensor reading needs to be stored. The processing requirements are minimal since the only processing required is comparison between successive sensor readings (gradient computation). Only a minimal amount of motion control is required to hold the heading of the robot in a particular direction for a particular duration of time (depending on bias levels).

III. SIMULATION EXPERIMENTS

A. Methodology

To validate our ideas and explore the possible implications of a bacterial motion-based approach for localization and tracking of gradient sources and gradient boundaries, we designed a simulation platform.

We created a model of the world as a uniform two dimensional grid of dimension 2000 * 2000 Units (Fig. 1). We choose this size for our model because this would help us generate areas of very low/negligible concentrations within the grid. We initialize one or more sources of gradient (S_i) at randomly chosen positions in the grid and compute the gradient generated by these sources at the other grid points following an inverse square law distribution given by

$$Intensity(x, y) = \frac{1}{K} \sum_{i=0}^m \frac{Q_i}{r_i^2} \quad (6)$$

which gives the concentration that can be sensed at a point (x, y) on the grid in the presence of m sources of gradient, Q_i is the intensity of the source S_i , K is a constant of proportionality and r_i is the distance between the grid point (x, y) and center of source S_i .

All our simulations consisted of a set of 100 robots. These were either deployed randomly in the grid (using the values




Fig. 2

PURE RANDOM WALK VS BIASED RANDOM WALK




Fig. 1

SIMULATION GRID SETUP

derived from a uniform distribution over the grid region) or based on the requirements of the simulation. At each time step each of these robots could either move to one of its 8 neighboring grid points or change its direction of motion or stay at its position if it has reached the gradient source.

We simulated a biased random walk with a Mean Free Path (MFP) of 10 units, i.e., under the absence of a concentration gradient, each robot would move 10 units of distance along the grid in a particular direction before tumbling and changing its direction of motion randomly. However, if the robot senses a positive change in gradient, it decreases its tumbling frequency thus increasing the run-length resulting in a biased random walk (Fig. 2). The typical bias value we used in our simulations is 10% of the MFP which is similar to the bias values observed in nature for bacterial motion [9].

We also carried out simulations where the intensity of the source varied over time. The change in intensity could be attributed to a variety of reasons such as inherent nature of

the source to dissipate its energy over time, occlusion of the source, consumption of the source (for ex., a nutrient source being consumed by bacteria over time). We used the following model:

$$q = (q_0 - k_1 N_i t) + (q_0 e^{-k_2 t}) \quad (7)$$

where k_1 and k_2 are constants whose values are set based on the type of source we are trying to model, N_i is the number of robots at source S_i depleting its energy at time t . The first term models the decrease in source intensity due to consumption by the robots near it while the second term models the decrease in intensity due to source depletion or dissipation over time.

B. Simulation Results

We carried out several simulations modeling the biased random walk strategy under a wide variety of simulation conditions. In our trials, once a robot reaches a particular gradient source, it stays there as long as the source remains sufficiently active (i.e., intensity remains above some threshold value). We also designed a set of metrics which would help understand and evaluate our approach. The results we present in this section are averages over 10^4 trials over the whole group for various simulation setup conditions.

Fig. 3 presents results of simulations performed to study the effect of bias levels on the speed of convergence of the robots to the source. Fig. 4 shows results from simulations where we begin by initializing a single gradient source at the start of simulation ($t = 0s$). A second source is introduced at $t = 5000s$. Both sources are modeled to dissipate over time at a rate proportional to the number of robots which have reached them.

Our first metric is the average displacement of the robot from its initial position. We monitor this over time as the simulations proceed.

We used a set of different robot and source initialization techniques and observed the following results. As can be seen in Fig. 3, even when no bias is present, the robots do move and explore a lot of space. The presence of a bias




Fig. 3

THE EFFECTS OF VARYING BIAS. THE LARGER THE BIAS, THE QUICKER THE CONVERGENCE. THE DATA ARE VISUALIZED USING THREE METRICS: A. ROBOT DISPLACEMENT VS. TIME (LEFT), B. DISTANCE BETWEEN ROBOT AND SOURCE VS. TIME (MIDDLE), AND C. PERCENTAGE OF ROBOTS AT SOURCE VS. TIME (RIGHT).




Fig. 4

MULTIPLE SOURCES. A. ROBOT DISPLACEMENT VS. TIME, B. DISTANCE FROM SOURCE VS. TIME, C. PERCENTAGE OF ROBOTS AT SOURCE 1 & 2, AND D. INTENSITY OF SOURCE 1 & 2

speeds up the movement of the robots into regions of interest. The exploration phenomena was observed to be independent of initialization conditions as long as the other simulation parameters (bias etc.) remained the same.

The second metric is the average distance of the robot from the source over time. This reflects how an average robot moves with respect to a gradient source over time. We studied this parameter also under a wide range of simulation models (different types of gradient setups, different initialization strategies for robots and sources of gradients, different types of gradient sources (decaying vs. constant intensity)). Some of the results can be seen in Fig. 3(b), 4(b). As expected, in the absence of a bias, the robots wander around without showing any particular progress towards the source. An introduction of a very small amount of bias of the order of 10% results in a rapid speed up of movement of robots towards the source of the gradient.

Another question that comes up to mind is how well did the robots do in terms of reaching an individual gradient source, how fast and how many of them reached which source? This gives us a third metric to evaluate our model wherein we monitor the average number of robots reaching a particular source over time. This metric is more meaningful when plotted

alongside the source intensity variations. A look at a few results Fig. 3(c), 4(c) demonstrates how the robots reach the source as the simulation proceeds in time. They also demonstrate how the appearance of a second source attracts some of the robots thereby pulling them away from the first source. As the intensity of one of the sources starts decreasing, some of the robots start moving away from it and once it disappears, all the robots which were near it start executing a random walk again in search of other sources and move away from it.

From the above results we can come up with a measure of how much bias is ideal for the type of system we are trying to model. Higher bias values speed up the movement towards the source but might not be good for sources which vary in intensity or are mobile. Lower bias values result in slower response time to reach the source but are more effective in tracking weaker, mobile sources whose intensity varies over time. One can thus trade bias for speed vs. efficiency.

The results from the multiple source experiment (Fig. 4) also demonstrate that this technique does not show any preferential behavior of moving towards a particular source based on the time of appearance of the source. We also carried out a set




Fig. 5

EFFECT OF INTRODUCTION OF ERRORS IN THE DECISION FUNCTION

of simulations simulating a number of sources with different intensities and the results were as we expected. All sources were tracked though the weaker sources got a comparatively smaller fraction of the robots in case the gradients they set up were very small, but on an average all sources were well covered. Another set of simulations were performed to study how the performance of the algorithm varied in presence of linear, cubic and exponential gradients. The results were almost similar since the algorithm inherently works by measuring the sign of difference between successive samples rather than absolute magnitudes.

After executing a *run*, the robot essentially takes a new sensor reading and compares it with the previous sensor reading taken before it started the current run. The decision to tumble or continue is then based on if the difference was negative or positive and does not depend on the absolute sensor readings. This makes the system tolerant to static sensor errors. We carried out another set of simulations introducing a Gaussian error in the decision function itself. Fig. 5 present the results. Even in the presence of 40% error, the system still converges to the gradient source. This error models the non-static sensor errors and actuation errors (i.e., motion of the robot might not be the same as the command signal applied).

We also carried out a set of simulations to understand how well the robots spread around a gradient source i.e., do the robots reach a gradient source all at the same place? We modeled a circular source with a diameter of 45 units and studied the effect of deploying the robots randomly, uniformly and at a single location in the grid at the start of the simulation. Fig. 6 presents the result of initializing all the robots at a single location at one corner of the grid. As can be seen, the robots approach the source from all directions. Similar results were obtained from random and uniform initial robot deployments. A point to note here is that even if all the robots were deployed in very close vicinity of the source, some spread along the periphery of the source was achieved. The results from irregularly shaped sources were also in agreement with the above observations.





Fig. 6

EFFECTIVENESS AT BOUNDARY DETECTION



C. Discussion

From the above results we conclude that a strategy based on biased random-walk implemented on a group of mobile robots can effectively track multiple dynamically changing gradient sources at the same time. Moreover, such a set of robots require very minimal amount of controls to be developed. The only control element is to change the length of the run in response to the sensed gradient change.

The source boundary detection results also highlight the suitability of our approach to a set of applications where the source dimensions are comparatively large and we need to track its complete boundary (e.g. the boundary of an underwater plume or the boundary of an oil spill). For such applications the algorithm should be able to achieve a sufficient spread of robots along the entire boundary.

IV. EXPERIMENTS WITH THE ROBOMOTE ROBOT

Validation experiments were carried out on the Robomote (Fig. 7(a)) test bed developed at the Robotic Embedded Systems Lab at the University of Southern California. The capabilities of this small matchbox sized two wheeled robomote include moving along a straight line for a specified duration and/or distance and the ability to turn in place by a specified




Fig. 7
ROBOMOTE AND THE TESTBED




Fig. 8
ROBOT BEHAVIOR IN RESPONSE TO (A) SINGLE AND (B) MULTIPLE GRADIENT SOURCES

angle. For a more detailed description of the platform the reader is referred to [19].

We used Mica-I motes developed by Crossbow to give control commands to the Robomote using TinyOS [20]. A light gradient was generated using a light source placed at one end of the test-bed. A basic sensor board with a photo sensor was mounted on the Robomote (Fig. 7(b)) to sense the light gradient. The experimental setup can be seen in Fig. 7(c). The position of the robomote on the test bed was tracked using an overhead camera which captured frames and passed these to a tracker [21] for data analysis and storage. Color blobs were mounted on top of the Robomote to help in the detection of Robomote location on the test bed.

We used the two basic components *move* and *rotate* written in TinyOS for controlling the robomote to carry out the biased random walk. We positioned the robomote at a distance d ($d=40\text{cm}, 80\text{cm}, 120\text{cm}$) from the source. Note that by fixing d , the heading of the robomote still was a random variable and could be towards or away from the source. Also we considered a small circle of 5cm. around the source which we considered as the source radius. We were interested in measuring if the robomote reaches the source and if so in how much time? The speed of the robomote was set at 2cm/s with a turn time of approximately the same duration.

Once switched on at a distance d , the robomote starts off by taking a sample using the photo sensor (S_i). It starts moving along a straight line in the direction of its current heading for a distance and/or duration as specified in the random walk

parameter MFP. At this point it takes another photo sensor reading and compares it with the previous reading from the photo sensor. If it senses no change or a negative change in gradient, it randomly chooses a new heading direction (θ_i) and rotates in place to orient to that heading. If a positive change in gradient was sensed, it continues its motion for an additional distance specified by its bias value before randomly computing the new heading and making a turn ('tumbling'). In either case the procedure is repeated by moving along a straight line in the direction of its heading for a distance and/or duration as specified in the random walk parameter MFP, followed by another decision based on the photo sensor reading. The experiment terminates when the robot reaches the source.

Each of the d values constituted a circular arc on the table of radius d units from the center of the light source. We repeated the experiment 75 times for each of the d values with random starting orientations and starting locations on the arc and averaged our position readings between the gathered data for our analysis. We believe this gave us a good enough set to evaluate the effectiveness of the approach. The graphs for the metrics we proposed in the previous section can be seen in Fig. 8(a). The results from the robomote platform agree with the results we obtained from our simulation work.

We repeated the experiment with two equal intensity sources present at the same time at opposite corners of the test bed and started the robomote at distances d ($d = 25\%, 50\%$ (center) and 75% positions on the test bed) (Fig. 7(d)). Contrary to the




Fig. 9

COMPARISON OF GRADIENT DESCENT STRATEGY WITH BIASED RANDOM WALK IN PRESENCE OF A SINGLE SOURCE




Fig. 10

COMPARISON OF GRADIENT DESCENT STRATEGY WITH BIASED RANDOM WALK IN PRESENCE OF MULTIPLE SOURCES

expectation that the robomote would have moved to the closest source, we observed what we had expected i.e., the source with which the robomote finally ended at was independent of the starting orientation and location of the robomote at the start of the experiment.

As a final test for our biased random-walk strategy, under the same setup as previous experiment(Fig. 7(d)), we started with only one source switched on initially at $t = 0$ s. At $t = 180$ s. we switched on the second source located at the other end of the test bed. This was followed by switching off the first source completely at $t = 435$ s. We repeated the experiment for the different values of d ($d = 25\%, 50\%$ (center) and 75% positions on the test bed)). Again the results (Fig. 8(b)) obtained in hardware platform were in agreement with our simulation results.

V. COMPARISON WITH GRADIENT DESCENT STRATEGY

We are trying to locate a gradient source by essentially following gradients. An obvious question arises: why not just use a simple gradient descent algorithm? In this section we compare our approach to a simple gradient search based approach.

We begin by presenting an evaluation of the performance of the simple gradient descent strategy in simulation. We use the same simulation framework that we developed in section III-A. We simulated a gradient descent strategy with a Mean

Free Path(MFP) of 10 units, i.e., under the absence of a positive concentration gradient in the current direction of motion, the robot would move 10 units of distance along the grid in that direction before tumbling and changing its direction of motion randomly and then repeat the run and tumble stages. However, if the robot senses a positive change in gradient, it continues moving, in the direction it was moving, for another MFP units before checking for the gradient change again.

We repeated the simulations for over a 10^4 trials so that the effects of noise and effects of outliers, if any, can be filtered out by averaging and the results could be compared against those from the biased random walk approach. The results of the simulation are presented in Fig. 9 along with results for the biased random walk simulations (repeated here for easier comparison).

The results clearly indicate that a simple gradient search based scheme performs better than our biased random walk for small bias values. However, as the bias levels are increased, the two are comparable. But sources in nature seldom occur alone and without interference from each other. Since our technique was developed to work in real life situations, we need to evaluate the gradient search strategy in those environments also.

We performed another set of simulations with multiple sources. We started the simulation with a single source at one corner of the grid at time $t = 0$. At time $t = 1500$ s. we

initialized a second source at the opposite corner of the grid and tracked the performance of the gradient descent algorithm over time. At $t=10000s$, we turned off the first source completely and continued the tracking of the robots. The results from the simulation are presented alongside the results for biased random walk in Fig. 10. As can be seen from the results, the gradient descent strategy works better than the biased random walk approach as long as there is one source. It is better in terms of number of robots which reached the source (Fig. 10(c)) as well as the time it takes to get there. But when we introduce a second source, the gradient descent strategy just follows the first source and the number of robots at the first source keeps increasing and almost none of the robots reach the second source (Fig. 10(d)). On the other hand, the results from the biased random walk are quite impressive. Both the sources get a good share of the number of robots reaching it and both can effectively be tracked at the same time. The introduction of more sources results in some robots tracking each one of them as long as there are some available robots. Thus for the purpose of tracking multiple sources, our algorithm clearly outperforms the gradient descent approach.

Our last set of simulations verify the source coverage obtained by the gradient search algorithm. We initialized all the robots at the same location on the grid. The gradient descent algorithm terminated with all the robots on one quarter of the source boundary (nearest their initial position), whereas (as presented in the previous section), the biased random walk approach resulted in a spread all around the source. From the above set of results, clearly the biased random walk approach we propose outperforms the simple gradient search strategies for the applications we consider.

VI. CONCLUSION AND FUTURE WORK

In this paper, we presented an algorithm based on biased random walk [8] for the detection, seeking and tracking of gradient inducing source phenomena. Our approach is inspired by the way bacteria detect, locate and track nutrient sources in nature [9]. Through a wide set of simulation and experimental work on robots, we demonstrated how our strategy is well suited to varied conditions including multiple sources, dissipative sources, and noisy sensors and actuators. We also show how our approach could be used for boundary finding. We validated our approach by testing it on a small robot (the robomote) in a phototaxis experiment. A comparison of our approach with gradient descent shows that while gradient descent is faster, our approach is better suited for boundary coverage, and performs better in the presence of multiple and dissipative sources.

Our algorithm is robust in the sense that it is guaranteed to work as long as there is a gradient. The inherent randomness of the algorithm saves the robots from landing in a local minimum. Since the algorithm is insensitive to inherent sensor errors and can tolerate actuation errors, it is a good choice for low-cost sensors. The algorithm can thus be implemented on very simple robots, suitable eventually for deployment in large numbers.

In the future, we will carry out surface-water experiments for detection and tracking of one or more dissipating color dye sources. We will also extend our algorithm to perform distributed data fusion of the data from multiple sensors on the robot.

ACKNOWLEDGMENT

The authors thank Prof. David A. Caron and Beth Stauffer, Department of Biological Science, University of Southern California for their comments. This work is funded in part by grants CCR-0120778, EIA-0121141, IIS-0133947 from the National Science Foundation.

REFERENCES

- [1] A. T. Hayes, A. Martinoli, and R. M. Goodman, "Distributed odor source localization," in *Special Issue on Artificial Olfaction*, H. T. Nagle, J. W. Gardner, and K. Persaud, Eds., vol. 2, no. 3, IEEE Sensors. IEEE Press, 2002, pp. 260–271.
- [2] F. Grasso, T. Consi, and J. Mountain, D. Atema, "Locating odor sources in turbulence with a lobster inspired robot," in *Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior, From Animals to Animals 4*. Cape Cod, MA: MIT Press, 1996.
- [3] R. Russell, T. D., R. Devez, and A. Mackay-Sim, "A robotic system to locate hazardous chemical leaks," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 1995, pp. 556–561.
- [4] H. Ishida, T. Nakamoto, T. Moriizumi, T. Kikas, and J. Janata, "Plume-tracking robots: A new application of chemical sensors," *Biological Bulletin*, no. 200, pp. 222–226, Apr 2001.
- [5] H. Ishida, M. Hirayama, J. Janata, D. Webster, and T. Moriizumi, "Sensor arrays for tracking of underwater chemical plumes," *IEICE, Tech. Rep.* 253, 2001.
- [6] A. T. Hayes, A. Martinoli, and R. M. Goodman, "Swarm robotic odor localization," in *Proceedings of the IEEE Conf. on Intelligent Robots and Systems IROS-01*, oct 2001.
- [7] B. Zhang, G. S. Sukhatme, and A. A. G. Requicha, "Adaptive sampling for marine microorganism monitoring," submitted to International Conference on Robotics and Automation (ICRA 2004).
- [8] W. Alt, "Biased random walk models for chemotaxis and related diffusion approximations," *J. Math. Biol.*, vol. 9, pp. 147–177, 1980.
- [9] H. C. Berg, *Random Walks in Biology*. Princeton University Press, 1983.
- [10] J. Adler, "Chemotaxis in bacteria," *Science*, vol. 153, pp. 708–716, 1966.
- [11] H. C. Berg and D. A. Brown, "Chemotaxis in *Escherichia coli* analysed by three-dimensional tracking," *Nature*, vol. 239, pp. 500–504, Oct 1972.
- [12] D. A. Brown and H. C. Berg, "Temporal stimulation of chemotaxis in *Escherichia coli*," in *Proc. Natl. Acad. Sci. USA* 71, 1974, pp. 1388–1392.
- [13] S. D. Muller, J. Marchetto, S. Airaghi, and P. Koumoutsakos, "Optimization based on bacterial chemotaxis," in *IEEE Transactions on Evolutionary Computation*, vol. 6, no. 1, feb 2002.
- [14] H. C. Berg and R. A. Anderson, "Bacteria swim by rotating their flagellar filaments," *Nature*, vol. 245, pp. 380–382, 1973.
- [15] C. S. Patlak, "Random walk with persistence and external bias," *Bull. of Math. Biophys.*, vol. 15, pp. 311–338, 1953a.
- [16] E. F. Keller and L. A. Segel, "Travelling bands of chemotactic bacteria, a theoretical analysis," *J. Theor. Biol.*, vol. 30, pp. 235–248, 1971.
- [17] H. G. Othmer, S. R. Dunbar, and W. Alt, "Reaction-diffusion equations and their applications to biology," *J. of Math. Biol.*, vol. 26, no. 3, pp. 263–298, 1988.
- [18] I. Golding, I. C. Y. Kozlovsky, and E. Ben-Jacobs, "Studies of bacterial branching growth using reaction-diffusion models for colonial development," *Physics A*, vol. 260, pp. 510–554, 1998.
- [19] G. T. Sibley, M. H. Rahimi, and G. S. Sukhatme, "A tiny mobile robot platform for large-scale sensor networks," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA2002)*, Washington DC, May 2002.
- [20] TinyOS software for the motes. [Online]. Available: <http://webs.cs.berkeley.edu/tos/index.html>
- [21] Mezzanine - an overhead visual object tracker. [Online]. Available: <http://playerstage.sourceforge.net/mezzanine/mezzanine.html>