



University of Pennsylvania
ScholarlyCommons

Departmental Papers (ESE)

Department of Electrical & Systems Engineering

8-2015

A drift-diffusion model for robotic obstacle avoidance

Paul B. Reverdy

University of Pennsylvania, preverdy@seas.upenn.edu

B. Deniz Ilhan

University of Pennsylvania

Daniel E. Koditschek

University of Pennsylvania, kod@seas.upenn.edu

Follow this and additional works at: https://repository.upenn.edu/ese_papers



Part of the [Electrical and Computer Engineering Commons](#), and the [Systems Engineering Commons](#)

Recommended Citation

Paul B. Reverdy, B. Deniz Ilhan, and Daniel E. Koditschek, "A drift-diffusion model for robotic obstacle avoidance", *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems*. August 2015.



To appear.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/ese_papers/703
For more information, please contact repository@pobox.upenn.edu.

A drift-diffusion model for robotic obstacle avoidance

Abstract

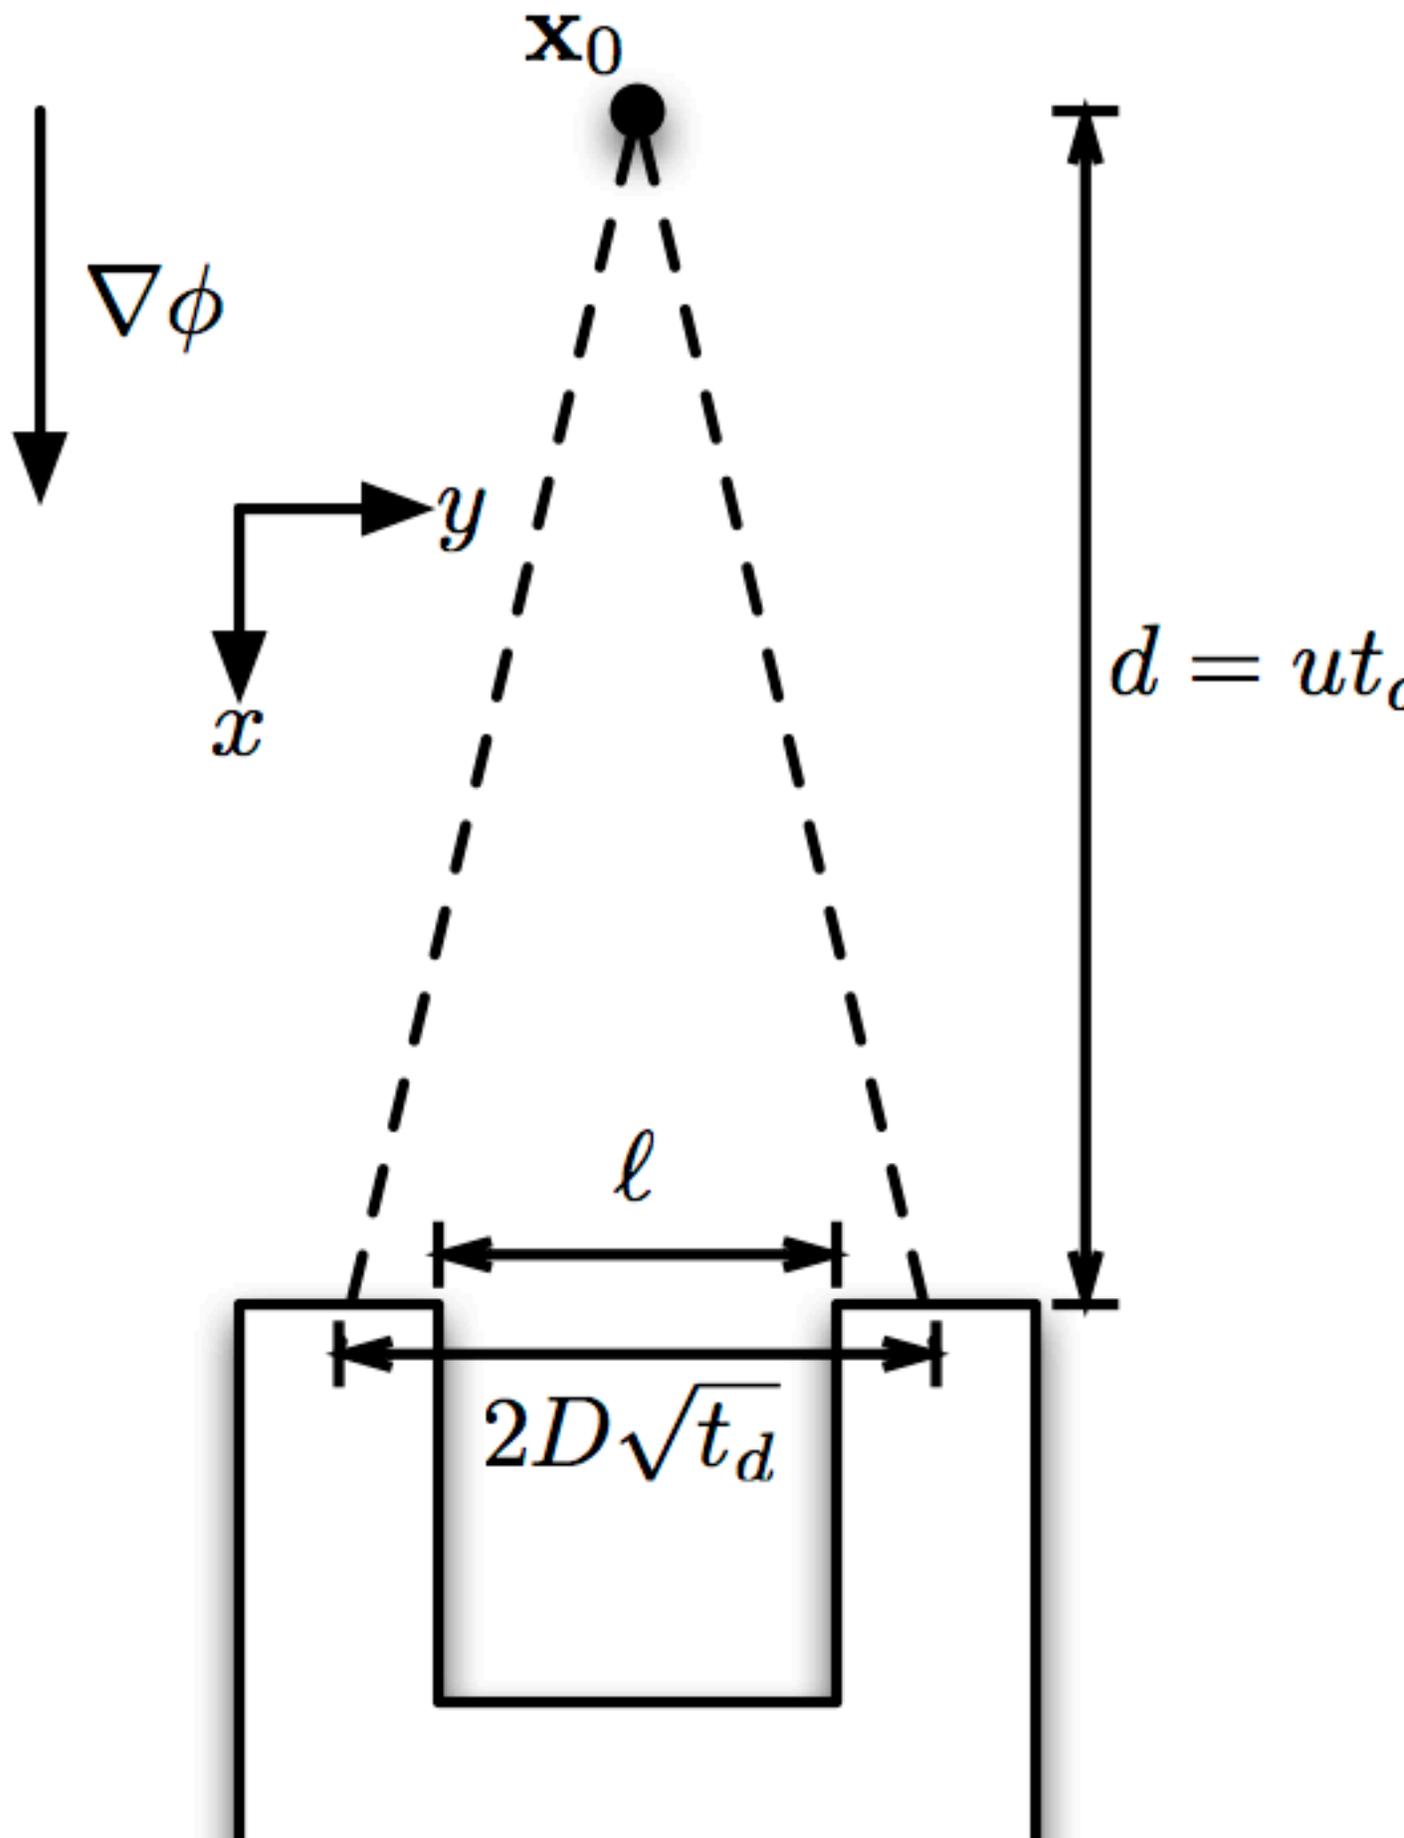
We develop a stochastic framework for modeling and analysis of robot navigation in the presence of obstacles. We show that, with appropriate assumptions, the probability of a robot avoiding a given obstacle can be reduced to a function of a single dimensionless parameter which captures all relevant quantities of the problem. This parameter is analogous to the Peclet number considered in the literature on mass transport in advection-diffusion fluid flows. Using the framework we also compute statistics of the time required to escape an obstacle in an informative case. The results of the computation show that adding noise to the navigation strategy can improve performance. Finally, we present experimental results that illustrate these performance improvements on a robotic platform.

For more information: [Kod*Lab](#)

Disciplines

Electrical and Computer Engineering | Engineering | Systems Engineering

Comments



To appear.

A drift-diffusion model for robotic obstacle avoidance

Paul Reverdy, B. Deniz İlhan, and Daniel E. Koditschek

Abstract—We develop a stochastic framework for modeling and analysis of robot navigation in the presence of obstacles. We show that, with appropriate assumptions, the probability of a robot avoiding a given obstacle can be reduced to a function of a single dimensionless parameter which captures all relevant quantities of the problem. This parameter is analogous to the Péclet number considered in the literature on mass transport in advection-diffusion fluid flows. Using the framework we also compute statistics of the time required to escape an obstacle in an informative case. The results of the computation show that adding noise to the navigation strategy can improve performance. Finally, we present experimental results that illustrate these performance improvements on a robotic platform.

The field of mobile robotics seeks to develop platforms that can autonomously navigate through spatial domains in the service of a task. Central to any spatial navigation capability intuitively deserving of the term autonomy is the ability to negotiate (avoid and escape from) obstacles and decades of research have yielded many different approaches to that problem. Inevitably, any “intelligent” navigation method must make assumptions about its environment: the more assumptions the more potential for “intelligence” — and, conversely, for foolishness when those assumptions prove ill-founded.

Of the many approaches to obstacle avoidance, we are particularly interested in artificial potential functions, e.g., [1], [2], [3], [4], [5], which encode the problem in terms of forces (acting to both repel away from undesired and attract toward desired states) since they generate controllers that readily lift to more realistic models of a robot’s mechanics [6]. These methods assume full knowledge of the environment with the benefit of developing deterministic controls amenable to formal obstacle avoidance guarantees of varying strength. Of course, real environments are never known exactly and real robots have limited perceptual capacity.

Recent efforts to confer on these dynamically effective, formally correct but undeservedly optimistic methods [2], [5] a formalizable degree of robustness against such uncertainties have yielded an approach to dynamical replanning [7] that introduces an internal model capable of inferring and reacting to the presence of an unexpected obstacle by exciting special behaviors that promote escape. Unfortunately, the problem representation suitable to sound reasoning about the dynamical implications of these methods leaves a substantial gap with respect to the implications relating to knowledge about the geometric properties of the environment—most crucially, the obstacle loci and shape. Other authors, e.g.,

[8], have considered boundary-following strategies to encode geometric information into potential field methods. However, as far as we are aware, the literature on such strategies relies on ad-hoc taxonomies of possible obstacle geometries.

In this paper we take the very earliest steps toward a fundamentally stochastic approach to reasoning about the interaction between such a system and the geometric properties of its state space that shows promise for meeting up usefully with the deterministic properties of the underlying dynamics. For now, as a first step toward a stochastically-enhanced version of the deterministic replanner [7], we simply replace it and introduce stochastic noise into the otherwise deterministic robot dynamics and reason about the statistics of the resulting interaction with the uncertain local geometric environment. Unsurprisingly, this approach allows a more natural representation of that uncertainty. However, at the same time, less obviously, it invites a representation of the deterministic aspects of the obstacle avoidance control strategy in terms of boundary interaction models treated by a growing body of literature on stochastic differential equations (SDEs).

The remainder of the paper is organized as follows. In Section I we set our notation and define the obstacle-avoidance problem in rigorous mathematical terms. In Section II we analyze the problem in the special case of a single obstacle and compute the probability of avoiding the obstacle in our framework. In Section III we present an implementation of the framework on a legged mobile robot platform and show experimental results that corroborate our theoretical predictions. Finally, in Section IV we provide avenues for future work and conclude.

I. PROBLEM STATEMENT

The starting point for our framework is the navigation function method originally proposed in [2]. We model the robot as a point mass traveling in a domain $\mathcal{D} \subseteq \mathbb{R}^2$, so its configuration at time $t \in \mathbb{R}$ is given by $\mathbf{x}(t) \in \mathcal{D}$. The domain is cluttered with obstacles, which we model as closed curves in \mathcal{D} . We assume the existence of a navigation function $\phi : \mathcal{D} \rightarrow \mathbb{R}$, which is a differentiable function with a unique maximum. The navigation function encodes the robot’s task, which is to find maxima of ϕ . The robot achieves its task if its trajectory $\mathbf{x}(t)$ obeys

$$\lim_{t \rightarrow +\infty} \mathbf{x}(t) = \arg \max_{\mathbf{x}} \phi(\mathbf{x}). \quad (1)$$

The robot carries out its task by climbing the spatial gradient $\nabla \phi$ of the task function ϕ , so its idealized dynamics are given by

$$\dot{\mathbf{x}} = u \nabla \phi, \quad u \in \mathbb{R}_+,$$

The authors are with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104, USA. {preverdy, bdeniz, kod}@seas.upenn.edu

where the quantity u controls the speed at which the robot climbs the gradient. However, there are disturbances to these idealized dynamics due to, e.g., issues measuring the gradient $\nabla\phi$, interactions with the environment, as well as disturbances introduced as part of the control scheme. Denote the coordinates on \mathcal{D} as $(x, y) = \mathbf{x}$. We model the disturbance in each coordinate as a Wiener process of strength $D(\mathbf{x}) \in \mathbb{R}_+$, and assume that the two processes are independent. The process noise intensity is the sum of two terms: $D(\mathbf{x}) = D_a(\mathbf{x}) + D_c(\mathbf{x})$, where $D_a(\mathbf{x}) \in \mathbb{R}_+$ is the ambient noise due to the environment and $D_c(\mathbf{x}) \in \mathbb{R}_+$ is the control noise added as part of the control strategy.

The noise-corrupted dynamics are described by the following Itô stochastic differential equation (SDE)

$$d\mathbf{x} = \begin{bmatrix} dx \\ dy \end{bmatrix} dt = u \begin{bmatrix} \frac{\partial\phi}{\partial x} \\ \frac{\partial\phi}{\partial y} \end{bmatrix} dt + D(\mathbf{x}) \begin{bmatrix} dW_t \\ dV_t \end{bmatrix}, \quad (2)$$

where $D(\mathbf{x})$ is the strength of the disturbance at $\mathbf{x} \in \mathcal{D}$ and dW_t and dV_t are independent Wiener increments. Dependencies in the disturbances can be modeled by making $D(\mathbf{x})$ a positive-definite matrix-valued function of \mathbf{x} . Standard references for the SDE methods used in this paper are the books [9] and [10].

Solving Equation (2) generates trajectories of a single particle. Solving the equation repeatedly from the same initial conditions generates different trajectories due to the stochastic nature of the dynamics. Alternatively, one can consider the probability distribution $p(\mathbf{x}, t)$ of the state $\mathbf{x}(t)$ as a function of time t . The probability distribution is a function that gives the probability of finding the robot in a set of states:

$$\Pr[\mathbf{x}(t) \in S] = \int_S p(\mathbf{x}, t) d\mathbf{x}, \quad (3)$$

where $S \subseteq \mathcal{D}$ is a subset of the state space.

The time evolution of the distribution p induced by the dynamics (2) is given by the following partial differential equation:

$$\frac{\partial p}{\partial t} = \frac{1}{2} \nabla \cdot (D(\mathbf{x}) D(\mathbf{x})^T \nabla p) - u \nabla\phi \cdot \nabla p. \quad (4)$$

Equation (4) is known as the *Fokker-Planck equation* [9], [10]. Equations of this form are studied in the literature on scalar transport phenomena under various names such as the advection-diffusion equation and the *drift-diffusion equation*.

The following physical analogy is illustrative. Consider a drop of dye in a fluid flow. The function $p(\mathbf{x}, t)$ measures the concentration of dye at the spatial location \mathbf{x} at time t . If the dye is initially concentrated at \mathbf{x}_0 , the initial condition for the equation (4) is the Dirac delta function $p(\mathbf{x}, t_0) = \delta(\mathbf{x} - \mathbf{x}_0)$. As time elapses, the dye moves with the fluid, which flows with local velocity $\nabla\phi(\mathbf{x})$ and diffuses with coefficient $D(\mathbf{x})$. Transport due to the local velocity is called *advection*, or *drift*, while the spreading due to the $D(\mathbf{x})$ term is called *diffusion*, and the two terms of the equation are referred to accordingly.

The equations (2) and (4) define a stochastic dynamical system where u is a control parameter. In future work, we

will leverage tools from the stochastic geometry literature to derive ways to tune u such that the robot can navigate through a spatially-distributed obstacle field. A key next step to developing this theory will be the extension of our model to the case of multiple obstacles.

II. SINGLE OBSTACLE

In this section, we analyze the interaction of a particle obeying the stochastic dynamics (2) and a single obstacle, which we model as a closed curve in \mathcal{D} . We develop a set of assumptions under which we can calculate the probability of escaping a single obstacle in closed form as a function of a single dimensionless parameter.

A. Assumptions

We make the following assumptions to develop analytical tools to study informative cases of the obstacle escape problem.

- 1) The coordinates are aligned with the local gradient $\nabla\phi$, such that $\partial\phi/\partial x = 1$ is a constant and $\partial\phi/\partial y = 0$. In other words, traveling in the $+x$ direction is equivalent to climbing the (constant) local gradient.¹
- 2) The diffusion tensor $D(\mathbf{x})$ is diagonal and constant in \mathbf{x} : $D(\mathbf{x}) = D_i \delta_{ij}$.
- 3) Diffusion only acts in the dimension orthogonal to the gradient, so $D(\mathbf{x})$ has x component $D_1 = 0$ and y component $D_2 = D$.
- 4) The particle interacts with obstacles through specular reflection: if, prior to the interaction it has momentum \mathbf{p} , after the interaction it will have momentum $\mathbf{p}' = \mathbf{p} - 2\hat{\mathbf{n}}(\hat{\mathbf{n}} \cdot \mathbf{p})$, where $\hat{\mathbf{n}}$ is the outward normal vector to the surface of the obstacle at the point of contact. This is equivalent to assuming that the obstacles have infinite mass and that the particle-obstacle interaction is an elastic collision.

Specular reflection constitutes one of the two canonical types of boundary conditions generally specified for stochastic differential equations (with absorption being the other [10]). More recent work, e.g., [11], [12], has considered more general boundary conditions that could model inelastic collisions with a coefficient of restitution $\epsilon \in (0, 1)$. However, the interpretation of these boundary conditions is more complicated and would require more detailed modeling of the physical obstacle interaction. Therefore, in this preliminary study, we adopt the abstract reflecting boundary condition as the most appropriate for developing analytical results with the particle model considered here.

With these assumptions, the dynamics (2) reduce to

$$d\mathbf{x} = \begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} dt = \begin{bmatrix} u \\ 0 \end{bmatrix} dt + \begin{bmatrix} 0 & 0 \\ 0 & D \end{bmatrix} \begin{bmatrix} dW_t \\ dV_t \end{bmatrix}. \quad (5)$$

The drift-diffusion equation (4) induced by (5) is

$$\frac{\partial p}{\partial t} = \frac{D^2}{2} \frac{\partial^2 p}{\partial y^2} - u \frac{\partial p}{\partial x}. \quad (6)$$

¹This analytical simplification (guaranteed by the ‘‘flowbox’’ theorem of dynamical systems to be possible in the neighborhood of any non-singular orbit) would not have any impact on the actual physical implementation.

We assume that the initial location of the particle is known with certainty to be $\mathbf{x}_0 = (x_0, y_0) \in \mathcal{D}$, so the initial condition for (6) is $p(\mathbf{x}, t = 0) = \delta(\mathbf{x} - \mathbf{x}_0)$. Finally, we assume that the speed u is constant.

In the absence of boundary conditions, the solution of (6) can be found in closed form, and is (cf. [9, (5.20)])

$$p(\mathbf{x}, t) = \frac{1}{\sqrt{2\pi D^2 t}} \exp\left(-\frac{(y - y_0)^2}{2D^2 t}\right) \delta(x - (x_0 + ut)). \quad (7)$$

The solution can be interpreted as follows. The particle moves deterministically along the x coordinate with a constant velocity u and moves stochastically along the y coordinate according to a random walk. At time t , the particle distribution is Gaussian with center $(x_0 + ut, y_0)$ and standard deviation $D\sqrt{t}$.

For a given evolution time t the distribution has two characteristic lengths:

- 1) Advection length scale: ut
- 2) Diffusion length scale: $D\sqrt{t}$.

Their ratio, $ut/D\sqrt{t}$, is a form of the Péclet number [13], which is a dimensionless quantity that measures the relative strength of advection and diffusion. This ratio is a function of evolution time t ; if we specify an evolution time, we get a characteristic number that captures all the dimensional variables governing the dynamical behavior.

B. Probability of escaping a single obstacle

The dynamics (5) have a clear flow in the positive x direction. We take advantage of this behavior to characterize obstacles according to their geometry relative to the flow. Intuitively, the reflecting boundary condition specified in assumption 4) allows the particle to bounce off of obstacles and in some cases escape an obstacle by moving around it. However, a particle will clearly not be able to escape all obstacles in this fashion. Consider, for example, the crescent-shaped obstacle shown in Figure 1-B. If the advection term dominates in the dynamics (5), then the particle will tend to get trapped by the obstacle.

The examples in Figure 1 illustrate that the important characteristic of an obstacle in this framework is the convexity of its footprint with respect to the local advective flow. Loosely speaking, an obstacle is *convex* with respect to the flow (5) if the obstacle appears convex to an observer looking at it from the upstream direction. An obstacle *concave* with respect to the flow is defined analogously.

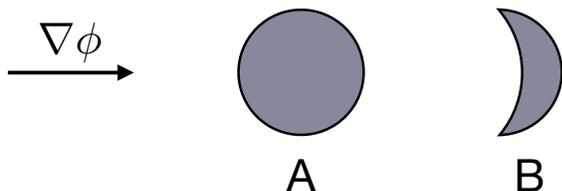


Fig. 1. Example obstacles placed in the flow described by the dynamics (5). Panel A: an obstacle that is convex with respect to the flow, and will not trap a particle with $D > 0$. Panel B: An obstacle that is concave with respect to the flow, and may trap a particle regardless of the value of D .

The definition of obstacle convexity can be made more precise in the following way. Define a section Σ transverse to the flow upstream of the obstacle. Consider the noise-free dynamics, i.e., (5) with $D = 0$. For each point $\mathbf{x} \in \Sigma$, define $g(\mathbf{x})$ as the time at which the solution to the noise-free dynamics with initial condition \mathbf{x} first touches the obstacle. The convexity property of the obstacle can now be formally defined as being inherited from that of the time-to-impact function, g .

If a particle following the dynamics (5) with $D = 0$ evolves from an initial condition upstream of a concave obstacle, it will eventually hit the obstacle and remain close to the point of impact. Conversely, we define a particle to have escaped an obstacle if its trajectory passes downstream of the obstacle. On the basis of physical intuition and numerical experiments, we argue that a particle following (5) with $D > 0$ will escape convex obstacles with probability one. This statement follows from the asymptotic behavior of solutions of (6), but that degree of formal development lies beyond the scope of the present exploratory paper.

In contrast to convex obstacles, concave obstacles can trap particles with positive probability. Therefore, we explore in somewhat greater detail the interplay between controlled drift and diffusion in the face of concavity. Figure 2 defines the quantities relevant to the interaction with a concave obstacle. The advection and diffusion length scales defined in the previous section appear, as well as two geometric length scales: d is the distance between the initial position of the particle and the front of the obstacle located downstream, and ℓ is the width of the concave section of the obstacle. Note that ℓ can be less than the width of the obstacle itself. For simplicity of exposition we assume that the obstacle has a mirror symmetry over the $y = y_0$ axis. The probability of escape thus computed is a lower bound for the probability associated with a non-symmetric obstacle with the same width ℓ .

The geometric length scales allow us to compute the probability that a particle obeying (5) will escape a given concave obstacle. We evolve the probability distribution (7) of the location of the particle until it reaches the front of the obstacle. This requires the particle to travel a distance d at a constant speed u , which takes time $t_d = d/u$. This sets the evolution time for the advection and diffusion length scales. The particle's location follows a Gaussian distribution with mean y_0 and standard deviation $\sigma = D\sqrt{t_d} = D\sqrt{d/u}$.

The particle will move into the concave region of the obstacle and get trapped if it is at the edge of the concave region at time t_d , i.e., if its y coordinate is in the range $(-\ell/2, \ell/2)$. Since the particle's location is Gaussian distributed, the probability that y is in this range, and therefore that the particle will become trapped, can be calculated in closed form. This yields π , the probability that the particle will avoid the obstacle:

$$\pi = \Pr[\text{Avoid obstacle}] = 2 \left(1 - \Phi\left(\frac{\text{Pe}}{2}\right)\right), \quad (8)$$

where $\text{Pe} = \sqrt{\frac{\ell^2 u}{D^2 d}} \geq 0$ is the Péclet number for the given

interaction and $\Phi : \mathbb{R} \rightarrow [0, 1]$ is the cumulative distribution function for the standard normal (i.e., Gaussian) distribution.

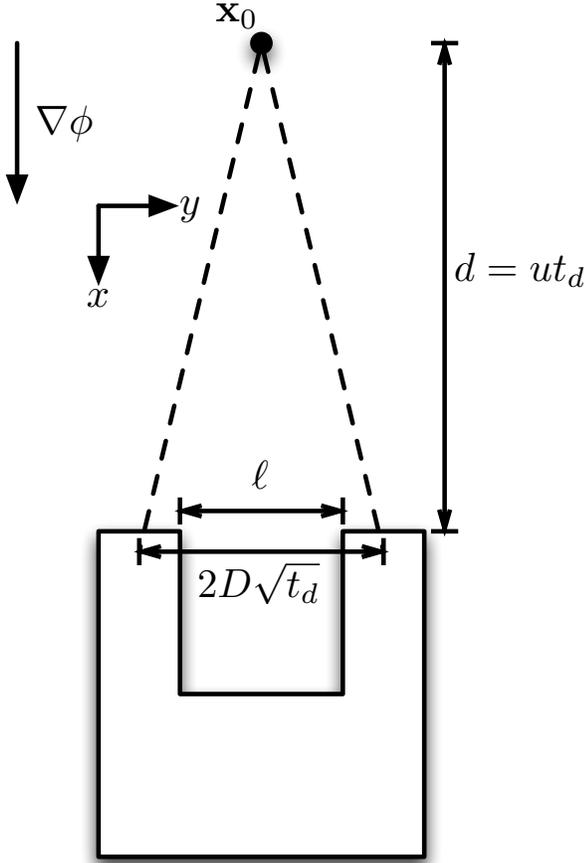


Fig. 2. The geometry of interaction with a concave obstacle. There are three characteristic lengths involved: d , $D\sqrt{t_d}$, and l . The particle starts at location \mathbf{x}_0 , which is at a distance d from the obstacle, and travels at a constant speed u . This defines the time to interaction t_d through the relationship $d = ut_d$. At the interaction time, the effects of diffusion mean the particle distribution has characteristic width $D\sqrt{t_d}$. When the particle interacts with the obstacle, it will get trapped if its location falls inside the concave footprint, which has width l .

Figure 3 compares the analytical avoidance probability (8) with the simulated avoidance probability computed from 100 numerical simulations of the particle interaction depicted in Figure 2. The two avoidance probabilities match well up to moderate values of the diffusion coefficient D ; for large values of D , there is more of a discrepancy, but this is likely due to approximation effects in the simulation code.

C. Escape time

The primary objective in the single obstacle problem is escaping the obstacle, for which the probability of avoidance (8) gives a quantitative metric. Given that the particle escapes the obstacle, a secondary objective is to do so quickly. For this objective a quantitative metric is time to escape, which can also be analyzed in our stochastic framework.

Consider again an obstacle interaction with geometry as in Figure 2. Define the random variable T as the first time at which the particle passes beyond the face of the obstacle.

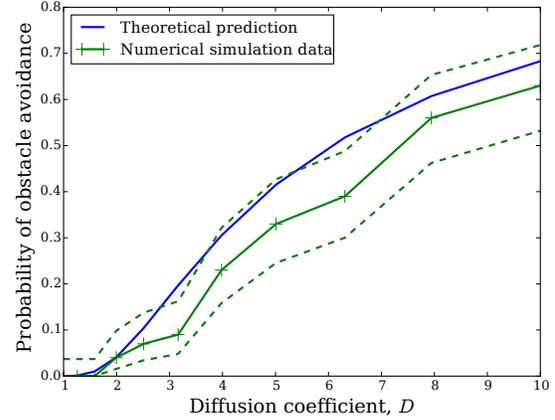


Fig. 3. Analytical vs. simulated obstacle avoidance probability for the concave obstacle depicted in Figure 2. The theoretic analytical probability is given by (8), while the simulated probability (with approximate 95% confidence interval) is computed as the empirical avoidance probability from 100 numerical simulations. The two quantities match well up to moderate values of the diffusion coefficient D .

That is,

$$T = \inf_{\tau \geq 0} \{x(\tau) > 0\},$$

where $x(\tau)$ is the x coordinate of the particle at time τ . The quantity T is a random variable because of the stochastic nature of the dynamics. In general, T can have a complicated distribution, which depends on the initial location of the particle. Let $T(\mathbf{x})$ represent the mean of T conditional on the initial location being equal to \mathbf{x} .

The function $T(\mathbf{x})$ (and the higher-order moments of T) can be computed using a partial differential equation that is closely related to the Fokker-Planck equation (4) [10, Section 6.6]. The partial differential equation can be solved analytically only in special cases, corresponding to obstacles with simple geometries. In more general cases, it can be solved numerically using finite element methods. An alternative method for finding the distribution of T is direct simulation of individual trajectories. This method is completely general and can be thought of as a type of particle filter method. In the following, we use direct simulation to study escape probability and escape time.

Figure 4 shows the probability of escape π and mean escape time $T(\mathbf{x}_0)$ as a function of diffusion coefficient D for a particle obeying dynamics (5) interacting with a circular obstacle with the geometry depicted in Figure 5. This geometry can be considered a special case of the geometry depicted in Figure 2 with the length l of the concave footprint being $l = 0$. As argued above, the details of the obstacle geometry outside the concave section of the footprint do not matter so long as they are convex with respect to the drift flow $\nabla\phi$.

When $D = 0$, the particle hits the obstacle at the point $(x, y) = (0, 0)$ and reflects directly along the direction in which it came, thereby getting trapped with probability one. For $D > 0$, the particle eventually escapes the obstacle, though the time to escape can be arbitrarily long. The figure depicts probability of escape in less than 10 time units; for

$D > 10^{-3}$, the probability of escape is effectively one. The time to escape, conditional on escaping in less than 10 time units, decreases with increasing D until it appears to reach an asymptotic value of approximately 2.5 for large D . The asymptotic value is similar to the value that would be seen if there were no obstacle and the particle were simply traversing the distance d .

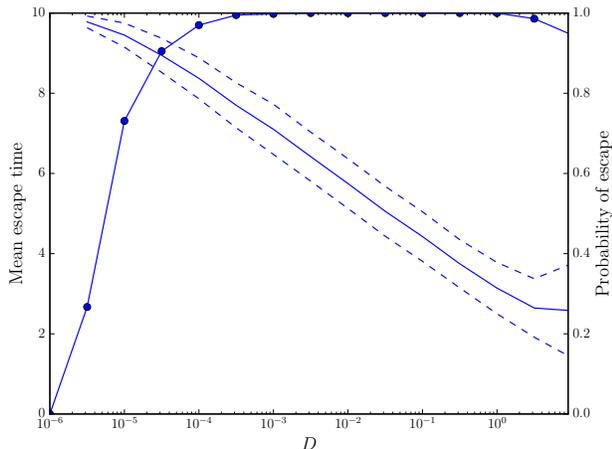


Fig. 4. Probability of escape (line with circles, right scale) and expected time to escape conditional on escaping (solid line, left scale) a circular obstacle of radius $R = 5$, as a function of diffusion coefficient D . For $D = 0$, the particle gets trapped with probability one, while for D greater than 10^{-3} , the probability of escape is effectively one. The drop in probability of escape for D greater than 1 is due to the finite time of simulation. The dashed lines indicate one standard deviation above and below the mean expected time to escape. All quantities were computed based on 1,000 simulations for each set of parameter values.

D. Implications for control

The result (8) and the escape time shown in Figure 4 have direct implications for obstacle avoidance control for a mobile robot, for which the particle model serves as a control target. We assume that the robot can control its speed u and the amount of process noise in its dynamics D by manipulating D_c . When there is an obstacle, onboard sensors, e.g., a LIDAR unit, will provide the robot with estimates of the distance d to the obstacle and its width, which serves as an upper bound for ℓ . If the sensor is sufficiently precise, it may be able to classify the obstacle as convex or concave, and provide an estimate of ℓ in the latter case. If the obstacle is convex, any positive noise will guarantee that the robot escapes the obstacle with certainty, i.e., probability one.

If the obstacle is concave, (8) implies that there is a non-zero probability that the robot will get stuck. However, we can make π , the probability of avoiding the obstacle, take any value in $(0, 1)$ by appropriately manipulating the control parameters u and D . This result provides a point of contact to recent work in the robotics literature that makes use of results from percolation theory, e.g., [14]. In this literature, the obstacle-strewn domain is modeled as a lattice graph and the probability of avoiding an obstacle is represented in

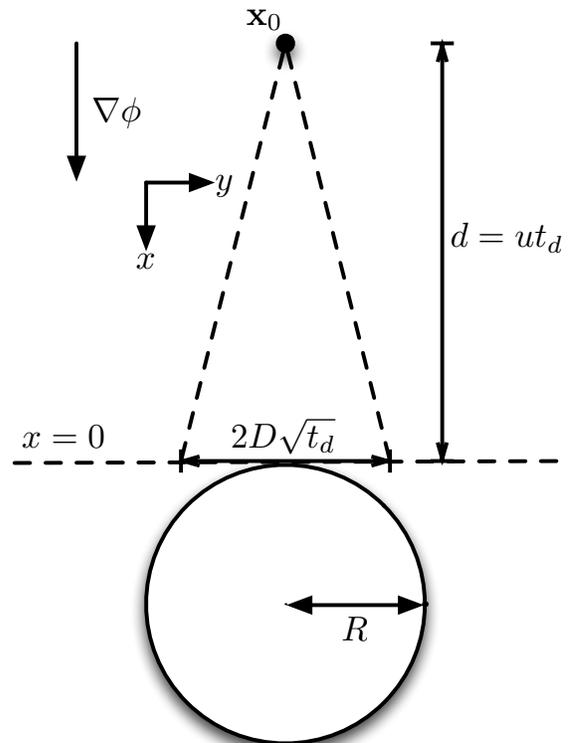


Fig. 5. The geometry of interaction with a circular obstacle. This can be thought of as a case of the geometry in Figure 2 with $\ell = 0$, as explained in detail in Section II-C. A trajectory of the particle dynamics (5) is said to escape from the obstacle if the trajectory crosses the plane $x = 0$ denoted by the dashed horizontal line.

terms of probabilities associated with the edges and vertices of the graph. Percolation theory then provides tools to find conditions under which it is feasible to travel extended distances through the graph.

For an individual obstacle, there will be tradeoffs between the control parameters because high avoidance probabilities are associated with small speeds u and large diffusion parameters D . Large values of D result in fast escape times, as seen in Figure 4. However, such large diffusion parameters result in large deviations from the desired flow along $\nabla\phi$. These deviations result in occasional large escape times T , which produce the dip in probability of escape and increased spread of escape time observed for $D > 1$. For a given interaction geometry, (8) shows that the two control parameters trade off in an inverse manner. This gives us a first step towards understanding the optimal way to trade off the parameters, which we intend to continue in future work.

III. EXPERIMENTAL RESULTS

Consider a particle interacting with a convex obstacle with geometry as in Figure 5. The qualitative prediction of the theory developed in the previous section is that in the noise-free case $D = 0$, the particle will get trapped behind the obstacle. This can be seen from Equation (8): a convex obstacle corresponds to the limit $\ell \rightarrow 0$, while the noise-free case corresponds to $D \rightarrow 0$. For the case of noise-free motion with a convex obstacle, D goes to 0 more quickly than ℓ , so this case corresponds to a Péclet number $\text{Pe} \rightarrow +\infty$. In

contrast, in the case with noise $D > 0$, the Péclet number obeys $Pe \rightarrow 0$ and Equation (8) predicts that the particle will eventually escape the obstacle. As shown in Figure 4, the theory also predicts that in this case the mean time to escape decreases sharply with increasing noise. In this section, we present results of robot experiments that corroborate these qualitative predictions.

A. Implementation on RHex

To verify the qualitative predictions of the theory developed above in the context of a physically interesting robot (rather than a more literal instantiation of the abstract point particle for which the theory and simulation are literally applicable), we implemented a version of the stochastic dynamics (5) on an X-RHex hexapedal robot [15]. The X-RHex robots have non-trivial dynamics [16] whose coarse horizontal plane motion in slow gaits (e.g., up to two body lengths per second) can be reasonably well approximated by a kinematic unicycle [17] and by a second order generalization of such nonholonomically constrained models when moving at higher speeds [18]. For purposes of the present exploratory paper, we used a gait slow enough to be usefully abstracted by the kinematic unicycle, and applied a variant of the controller in [17] whose anchoring relation [19] to the notional point-particle gradient dynamics posited in this paper can be rigorously established [20].

However, because we are disinclined to allow our robot to actually collide and bounce off a physical obstacle, our point particle gradient controller is rather more complicated than the simple constant-flow-with-elastic-collision model underlying the theoretical results presented above. Rather, we posit that the modified navigation function controller [7] implemented in these experiments introduces local deterministic interactions with obstacles that would be better modeled by the case of a plastic collision — i.e., the case $\epsilon = 0$ in Section II-A, Assumption 4. Looking ahead to assessing the efficacy of more sophisticated local replanners [7], we are pursuing the analysis of the more general “scattering” collision models discussed in that section. However, these more sophisticated analyses all lie beyond the scope of the present paper. In sum, the discrepancies between our actual robot control strategy and the abstraction used to develop the theory of Section II preclude any likelihood that quantitative predictions from the stochastic theory could be directly comparable to these experimental results. However, as we now report, the qualitative predictions are encouragingly reflected in these early empirical trials.

The implementation used for the robot experiments follows an approach that was first introduced by Khatib [1] where the task to be executed is represented by an artificial potential field composed of an attractive pole representing the goal state and repelling regions representing the obstacles. An extension to this approach was developed by Borenstein and Koren [21] where the obstacles are represented by certainty grids which enables a temporal filtering approach to obstacle detection. An alternative approach introduced by Borenstein and Koren [3] stems from some limitations

of the previous method and focuses on moving to empty regions rather than being repelled by obstacle regions. A previous implementation on the RHex platform [22] utilizes a similar approach. Our assumptions regarding obstacle shape and distribution let us disregard the limitations described by Borenstein and Koren and implement a simple local repelling field around obstacles where, with proper choice of control parameters, any spurious fixed points introduced to the force field are guaranteed to be unstable [20].

B. Experimental setup

In our experiments we used a circular obstacle in the geometry depicted in Figure 5. The effective radius of the obstacle was approximately $R = 0.75$ m, and the initial location of the center of the robot was at $\mathbf{x}_0 = (-2.0, 0.0)$ m, which is equivalent to an initial distance $d = 2.0$ m. The gradient field $\nabla\phi$ was generated by placing a point attractor in the far distance directly in front of the robot’s initial position, along with an immediate repeller located in the obstacle. The effective radius of the repeller was small, and is included in the effective radius of the obstacle. The resulting gradient field approximates the constant field assumed in the dynamics (5) to a degree of precision comparable to the other experimental uncertainties. Timing for runs was performed through manual control of logging software, which resulted in measurements of the time to escape that were accurate to within 1 s.

As defined above, the process noise D was modeled as the sum of two terms: $D = D_a + D_c$, where D_a was the ambient noise due to the environment and D_c was the control noise added as part of the control strategy. The ambient noise D_a is due to noise in the robot’s perceptual systems and various control loops. We manipulated the control noise D_c to have two values: either $D_c = 0$ or $D_c = \sqrt{40} \approx 6.3 \text{ m}\cdot\text{s}^{-1/2}$. We did not directly measure nor manipulate the ambient noise D_a , but it can reasonably be assumed to have been small and constant across the series of experiments. Importantly, the experimental results presented below imply that D_a was non-zero.

C. Results

The experiments consist of a number of obstacle interactions for the two control noise cases: the noise-free case $D_c = 0$ and the noisy case $D_c = \sqrt{40}$. For these first efforts, we focused exclusively on the single circular convex obstacle, rather than a family of obstacles including both convex and concave examples; such a family will be the subject of future work. Again, the noise values are not directly comparable to the diffusion coefficient D defined in Section II because of the details of the control strategy used on the robot. The results presented in Table I show, however, that the experiments match the qualitative trend predicted in Figure 4: adding control noise results in a higher probability of escaping the obstacle and a shorter mean time to escape for those runs that do escape.

In the noise-free case where $D_c = 0$, 50% of the runs resulted in the robot escaping the obstacle. In view of the

	Noise-free, $D_c = 0$ $N = 8$ runs	Noisy, $D_c > 0$ $N = 10$ runs
Probability of escape	50%	100%
Mean time to escape	45.08 s	8.860 s
Standard deviation	13.94 s	0.5393 s

TABLE I

EXPERIMENTAL RESULTS. THE NOISY CONTROL STRATEGY RESULTS IN AVOIDING THE OBSTACLE MUCH MORE QUICKLY AND WITH SIGNIFICANTLY HIGHER PROBABILITY.

results presented in Figure 4, this implies that the ambient noise D_a is small, resulting in an overall noise $D = D_a$ that is comparable to the value of 10^{-5} that one can interpolate from Figure 4. Adding noise ensures that 100% of the empirical runs resulted in the robot escaping the obstacle. This corresponds to pushing the system into the regime on the right hand side of Figure 4. The other benefit of the noise can be seen in the mean time to escape: adding noise results in decreasing the mean time to escape by a factor of approximately five. This represents a substantial increase in obstacle avoidance performance.

Clearly the theory does not account for all of the empirical trends: for example, the empirical standard deviation of time to escape decreases with increasing noise, while the simulations based on the particle model presented in Figure 4 show a standard deviation that is increasing with increasing noise intensity. The intuition behind this trend is as follows. In the model, when a particle interacts with an obstacle, it can be reflected into the direction opposite the desired direction of motion. When this occurs, the particle takes longer to escape the obstacle. The control noise injects momentum into the particle, so larger noise intensities result in more energetic reflecting interactions with the obstacle and larger standard deviations of the escape time. This energetic reflecting behavior does not occur in the physical experiment, showing a limitation of the reflecting boundary condition model. More detailed modeling of the robot-obstacle interaction is required to better match theory and experiment. This modeling will be the subject of future work.

IV. CONCLUSIONS

In this paper, we have presented a simple stochastic model for analyzing robot navigation in the presence of isolated obstacles. Pairing a very simple geometric model with the simple stochastic differential equation affords a characterization of the types of obstacles likely to trap a robot in terms of their convexity along with a closed form expression for the probability of escape from them. Simulation of this model yields numerical estimates of the expected time to escape (conditioned on escape) from the obstacle. Our simulation study suggests that adding noise to a simple navigation function-based obstacle avoidance approach can significantly decrease the time to escape, thereby substantially improving performance. Finally, a physical implementation on a mobile robotic platform yields empirical results whose trends match theoretical predictions.

Our stochastic approach also allows us to connect to recent work in robotics utilizing results from stochastic geometry. Notably, we aim to leverage the work of Karaman and

Frazzoli [14], who studied the motion planning problem for a vehicle obeying simple first-order integrator dynamics in a Poisson forest, i.e., an environment in which obstacles occur according to a spatial Poisson process. They proved that there is a phase change in the qualitative nature of the motion planning problem as a function of the vehicle's speed. Below a critical value of the speed, there exist planners that can navigate arbitrarily far through the forest. Above the critical value, no motion planner can avoid collisions: for any planner, a collision with an obstacle will inevitably happen in finite time. We hope to develop similar results for vehicles using navigation functions.

The work presented in this paper represents a first step towards a more general theory of navigation functions for uncertain environments. As mentioned previously, a key next step to developing this theory will be the extension of our model to the case of multiple obstacles. For this case we will rely on several tools from statistics, specifically the field of stochastic geometry. Hard-core point processes [23], [24] will provide the tools to model fields of obstacles, while percolation theory results, e.g., [14], [25] will provide the tools to develop control schemes to navigate through these obstacle fields.

As was shown by the comparison between our theoretical predictions and empirical data, further modeling needs to be done to understand how our simple model works when it is implemented in a physical robot. The methods of stochastic Lyapunov functions [26] are likely to be important tools for this work. The resulting theory of navigation functions for uncertain environments will provide tools to develop practical controllers for mobile robots that have provable navigation performance guarantees.

ACKNOWLEDGEMENT

This work was funded in part by the Air Force Office of Scientific Research under the MURI FA9550-10-1-0567. The authors thank D. Guralnik for discussions on the generalized definition of obstacle convexity developed in Section II-B, and the various reviewers for their comments, which strengthened the paper.

REFERENCES

- [1] O. Khatib. Real-time obstacle avoidance for manipulators and mobile robots. *The International Journal of Robotics Research*, 5(1):90–98, 1986.
- [2] D. E. Koditschek. Exact robot navigation by means of potential functions: Some topological considerations. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 1987.
- [3] J. Borenstein and Y. Koren. The vector field histogram-fast obstacle avoidance for mobile robots. *IEEE Transactions on Robotics and Automation*, 7(3):278–288, 1991.
- [4] J.-O. Kim and P. K. Khosla. Real-time obstacle avoidance using harmonic potential functions. *IEEE Transactions on Robotics and Automation*, 8(3):338–349, 1992.
- [5] E. Rimon and D. E. Koditschek. Exact robot navigation using artificial potential fields. *IEEE Transactions on Robotics and Automation*, 8(5):501–518, 1992.
- [6] D. E. Koditschek. The control of natural motion in mechanical systems. *Journal of Dynamic Systems, Measurement, and Control*, 113:547–551, 1991.

- [7] S. Revzen, B. D. Ilhan, and D. E. Koditschek. Dynamical trajectory replanning for uncertain environments. In *Proceedings of the 51st IEEE Annual Conference on Decision and Control*, pages 3476–3483, 2012.
- [8] M. Trevisan, M. A. P. Idiart, E. Prestes, and P. M. Engel. Exploratory navigation based on dynamical boundary value problems. *Journal of Intelligent and Robotic Systems*, 45(2):101–114, 2006.
- [9] H. Risken. *The Fokker-Planck Equation*. Springer, 1984.
- [10] C. Gardiner. *Stochastic Methods: A Handbook for the Natural and Social Sciences*, volume 13 of *Springer Series in Synergetics*. Springer, 4th edition, 2009.
- [11] P. Szymczak and A. J. C. Ladd. Boundary conditions for stochastic solutions of the convection-diffusion equation. *Phys. Rev. E*, 68:036704, 2003.
- [12] A. Singer, Z. Schuss, A. Osipov, and D. Holcman. Partially reflected diffusion. *SIAM Journal on Applied Mathematics*, 68(3):844–868, 2008.
- [13] M. Huysmans and A. Dassargues. Review of the use of Péclet numbers to determine the relative importance of advection and diffusion in low permeability environments. *Hydrogeology Journal*, 13(5-6):895–904, 2005.
- [14] S. Karaman and E. Frazzoli. High-speed flight in an ergodic forest. In *IEEE International Conference on Robotics and Automation*, pages 2899–2906. IEEE, 2012.
- [15] K. C. Galloway, G. C. Haynes, B. D. Ilhan, A. M. Johnson, R. Knopf, G. Lynch, P. Plotnich, M. White, and D. E. Koditschek. X-RHex: A highly mobile hexapedal robot for sensorimotor tasks. Technical report, University of Pennsylvania, 2010.
- [16] J. D. Weingarten, R. E. Groff, and D. E. Koditschek. A framework for the coordination of legged robot gaits. In *Robotics, Automation and Mechatronics, 2004 IEEE Conference on*, volume 2, 2004.
- [17] G. A. D. Lopes and D. E. Koditschek. Visual servoing for nonholonomically constrained three degree of freedom kinematic systems. *The International Journal of Robotics Research*, 26(7):715–736, 2007.
- [18] A. De, K.S. Bayer, and D.E. Koditschek. Active sensing for dynamic, non-holonomic, robust visual servoing. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6192–6198, May 2014.
- [19] R. Full and D. Koditschek. Templates and anchors: Neuromechanical hypotheses of legged locomotion on land. *J. of Experimental Biology*, 202(23):3325–3332, 1999.
- [20] B. D. Ilhan, A. M. Johnson, and D. E. Koditschek. Autonomous legged hill ascent. (*in preparation*).
- [21] J. Borenstein and Y. Koren. Real-time obstacle avoidance for fast mobile robots. *Systems, Man and Cybernetics, IEEE Transactions on*, 19(5):1179–1187, 1989.
- [22] A. M. Johnson, M. T. Hale, G. C. Haynes, and D. E. Koditschek. Autonomous legged hill and stairwell ascent. In *Safety, Security, and Rescue Robotics (SSRR), 2011 IEEE International Symposium on*, 2011.
- [23] D. Stoyan and H. Stoyan. On one of Matérn’s hard-core point process models. *Mathematische Nachrichten*, 122(1):205–214, 1985.
- [24] J. Teichmann, F. Ballani, and K.G. van den Boogaart. Generalizations of Matérn’s hard-core point processes. *Spatial Statistics*, 3(0):33–53, 2013.
- [25] S. N. Chiu, D. Stoyan, W. S. Kendall, and J. Mecke. *Stochastic geometry and its applications*. John Wiley & Sons, 2013.
- [26] H. J. Kushner. Stochastic stability. In *Lecture Notes in Mathematics* 294. Springer-Verlag, New York, 1972.