

Optimal Navigation Functions for Nonlinear Stochastic Systems

Matanya B. Horowitz, Joel W. Burdick

Abstract—This paper presents a new methodology to craft navigation functions for nonlinear systems with stochastic uncertainty. The method relies on the transformation of the Hamilton-Jacobi-Bellman (HJB) equation into a linear partial differential equation. This approach allows for optimality criteria to be incorporated into the navigation function, and generalizes several existing results in navigation functions. It is shown that the HJB and that existing navigation functions in the literature sit on ends of a spectrum of optimization problems, upon which tradeoffs may be made in problem complexity. In particular, it is shown that under certain criteria the optimal navigation function is related to Laplace’s equation, previously used in the literature, through an exponential transform. Further, analytical solutions to the HJB are available in simplified domains, yielding guidance towards optimality for approximation schemes. Examples are used to illustrate the role that noise, and optimality can potentially play in navigation system design.

I. INTRODUCTION

This paper presents a new method to construct navigation functions that bring a robot operating under stochastic uncertainty in its control inputs to a desired configuration while avoiding collision with obstacles. This is done through analysis of the Hamilton Jacobi Bellman Equation (HJB) for nonlinear stochastic optimal control problems. It has recently been found that under mild conditions the HJB is related through a logarithmic transform to a linear partial differential equation (PDE) [26], [29], [14]. This paper specializes this analysis to robot navigation functions, resulting in new planning methods for robots with stochastic nonlinear dynamics.

Related Work. This paper touches upon two classical problems: navigations functions (see Section II for a review), and stochastic optimal control problems. Navigation functions were introduced by Koditschek and Rimon [17], [18], [25] to remedy the local minima problem in the classical potential field method of robot motion planning [15], and their early work on navigation functions focused primarily on the existence and discovery of potential functions whose gradient would lead a point mass model of a robot from any point in the robot’s configuration space to a desired goal. Later work extended the navigation function concept to incorporate multiple agents [11], and sensory input [23].

However, little consideration has been given to including notions of optimality into a navigation function. This paper shows how to directly incorporate optimality criteria that can be modeled as the sum of a (possibly) nonlinear state dependent cost and a quadratic control cost. This contrasts

with the existing focus of the navigation approach, which implicitly defines a decomposition of the the problem into a trajectory generation method (the solution of the navigation function) and a local feedback trajectory following control method. This may lead to suboptimality in the path planning or control law, or even instability [20]. By formulating the problem relative to the Hamilton Jacobi Bellman problem, our method allows the impact of the dynamic model to be directly incorporated into the navigation function, if desired.

Specifically, we examine the presence or absence both dynamics and a state dependent cost function in an optimal control problem. We find that if these are neglected, methods similar to those used to generate navigation functions that exist in the literature are recovered [16], [21]. The result is that a spectrum of problems, ranging from the full HJB to the classical potential-based navigation function, are made explicit and the tradeoffs in modeling complexity becomes visible. The analysis we present makes it apparent how dynamics may then be incorporated to a navigation function if desired. To our knowledge, this paper also represents the first attempt to formally include stochastic uncertainty into the construction of the navigation function.

Solutions to stochastic optimal control problems pose inherently difficult computational problems. Classical work on addressing these problems have pursued several approaches. Systems with linear dynamics perturbed by Gaussian noise have well known closed form solutions [33]. For systems which do not fit these assumptions, several approximation methods, such as receding horizon control [12] or roll-out methods [1] have proven to be very useful. Another direction to solving these problems has been through discretization of the system’s state space, leading to a Markov Decision Process. When framed as such, it is possible to solve the problem exactly and globally through techniques such as value iteration or linear programming [1]. If a measure of error is acceptable, these methods have been made applicable to arbitrarily large state spaces through Approximate Dynamic Programming [5].

This paper follows an alternative line of work—*linearly solvable* stochastic optimal control (SOC) algorithms. The study of linearly solvable SOC problems has recently been studied from two avenues. Work begun by Kappen [13] examined the HJB and found that, using a transformation borrowed from quantum mechanics and assumptions on stochasticity, it was possible to find a linear partial differential equation (PDE) whose solution is the SOC solution. Independently, Todorov [28] found an alternate model for control in Markov Decision Processes that allowed for optimal MDP policies to be found as a solution to a set of linear equations,

Matanya Horowitz and Joel Burdick are with the Department of Control and Dynamical Systems, Caltech, 1200 E California Blvd., Pasadena, CA. The corresponding author is available at mhorowit@caltech.edu
Matanya Horowitz is supported by a NSF Graduate Research Fellowship.

as compared to value iteration. It was subsequently found that by taking the continuous limit of the discrete Markov model, this same linear HJB PDE is obtained. The HJB can then be understood as the limit process of an MDP.

Recent research along these lines has tended towards developing sampling based approaches for solving the resulting linear PDE. This is done through the use of the Feynman-Kac Lemma, which allows for a linear PDE to be solved by examining the diffusion of a stochastic process. This has been further developed by Theodorou et al. [27], [26] into the Path Integral framework. Other research into the study of the computational benefits, and analytic properties, of the linear HJB are available in [30], [8]. Our work begins from the same point, that of Stochastic Optimal Control. However, whereas [26] uses the linearity of the HJB to develop a novel reinforcement learning algorithm with emphasis on numerical solution methodologies, the focus of the present paper is on the application of this flavor of Stochastic Optimal Control to the problem of navigation functions. Connections between these two research fields are emphasized over the specifics of the numerical calculation of the PDE solution, for which we discuss available methods from the literature.

II. REVIEW: NAVIGATION FUNCTIONS

Let \mathcal{CS} denote the robot's *configuration space* (or *c-space*)—the possible configurations that a robot can occupy. As is standard, let the subset of \mathcal{CS} where the robot collides with an obstacle define the set of *configuration-space obstacles*, \mathcal{CO} , while the *free configuration space*, $\mathcal{F} \subset \mathcal{CS}$, is the complement of \mathcal{CO} in \mathcal{CS} . Under the assumption of perfect sensory information, and traditionally also perfect actuation control, the motion planning task is to move the robot from its starting configuration, $q_{init} \in \mathcal{F}$ to a desired goal position $q_d \in \mathcal{F}$. One approach to solve this problem is to construct a *navigation function*:

Definition 1. (From [18]) Let q_d be a goal configuration in \mathcal{F} , the free c-space. A map $\varphi : \mathcal{F} \rightarrow [0, 1]$ is a navigation function if it is

- 1) smooth on \mathcal{F} (at least a $C^{(2)}$ function);
- 2) polar at q_d , i.e., has a unique minimum at q_d on the path-connected component of \mathcal{F} containing q_d ;
- 3) admissible on \mathcal{F} , i.e., uniformly maximal on the boundary of \mathcal{F} ;
- 4) a Morse function.

Given a navigation function, $\varphi(q)$, the robot's path to the goal from any starting configuration in \mathcal{F} can be realized by following the gradient $\nabla\varphi(q)$ at each q . The definition assures that the robot will achieve the goal while remaining in \mathcal{F} , and not become trapped in a local minima of $\varphi(q)$.

Navigation functions may be constructed in several forms. In the well-known classical approach of [25], a navigation function may be calculated analytically when the when the bounded problem domain, the obstacle shapes, and the goal region are all diffeomorphic to spheres. Similarly, if the boundary, obstacles, and goal region are star-shaped

sets (which are homeomorphic to spheres), then one can compute the navigation function by transforming the problem to a sphereworld, find the sphereworld navigation function, and transforming the function back to the original problem domain.

A. Navigation Function Considerations

Navigation functions have been successful in part due to their rapid computability and its transparent nature. Furthermore, the navigation function provides both a global plan, as well as a feedback controller that follows the gradient of the navigation function. This allows the total motion planning and execution problem to be abstracted into a trajectory planner (the path followed by the gradient of the navigation function) and a path following controller [25]. The price to pay for this convenience is optimality: ignoring dynamics in the quest to follow the gradient will rarely result in a procedure that minimizes control effort. This is in part due to the fact that the system dynamics do not enter into the navigation functions calculation, and may result in unexpected and unstable behavior in some contexts [20]. The work presented below shows the connection between navigation functions and more general optimal control theory, allowing for system dynamics to be included if this is desirable. It also becomes possible to include more sophisticated weighting of various goals, such as the desire for minimum-time trajectories.

In the construction and intuition behind navigation functions, implicit is the idea of robustness. Since navigation functions are defined over the entire free configuration space, small deviations from the desired path place the robot in nearby locations where the desirable behavior is similar. Indeed, smoothness of the solution is typically enforced. This work furthers the understanding of robustness properties of navigation functions which has largely heretofore been only analyzed ad-hoc.

Finally, some approaches for finding navigation functions require difficult calculations and may not extend to complex obstacle geometries.

III. A LINEAR HAMILTON JACOBI BELLMAN PDE

We begin by reviewing results in optimal control theory that lead to the construction of the HJB. This paper considers nonlinear systems that evolve with the following nonlinear stochastic dynamics

$$dx = (f(x) + G(x)u) dt + B(x)L d\omega \quad (1)$$

where $x_t \in \mathbb{R}^n$ is the system state at time t , $u_t \in \mathbb{R}^m$ denotes the control input, ω is a unit variance Brownian motion stochastic process, and L is a state independent noise scaling factor. The functions $f(x)$, $G(x)$, $B(x)$ are assumed to be smooth functions of the state and may be nonlinear. We assume that the problem accrues costs r according to

$$r(x, u) = q(x) + \frac{1}{2}u^T Ru \quad (2)$$

where $q(x)$ is any (potentially nonlinear) non-negative function of state. The model for cost is limited to a form quadratic

in the controls in order to simplify the HJB equation in a future step of the analysis. The goal is to minimize the following expected cost functional

$$J(x_{0:T}, u_{0:T}) = \mathbb{E} \left[\phi(x_T) + \int_0^T r(x_t, u_t) dt \right] \quad (3)$$

where ϕ represents a state-dependent terminal cost, T is the final time of the trajectory, $\mathbb{E}[\cdot]$ is the expectation operators, and the symbols $x_{0:T}$ and $u_{0:T}$ denote the state and control over the interval $[0, T]$.

We consider the first-exit problem, wherein the state of the system exists in a compact domain Ω . The system continues to operate, and accrue cost, until it reaches the boundary, $\partial\Omega$, of the domain at time T whereupon the terminal cost $\phi(x(T))$ is accrued. In the navigation problem, this boundary consists of goals and obstacles in the robot workspace.

A common construction in the optimization literature is the value function, $V(x_t)$, which captures the ‘‘cost-to-go’’ from a given state. The optimal action follows the gradient of the value function, bringing the agent into the states with lower cost over the remaining time horizon. The solution to the optimization problem is, beginning from an initial point x_t at time t ,

$$V(x_t) = \min_{u(\cdot)} \mathbb{E}[J(x_t)] \quad (4)$$

Using a dynamic programming argument it is possible to derive the Hamilton-Jacobi-Bellman (HJB) equation associated with this problem [7], which is found to be

$$0 = \min_u \left(r + (\nabla_x V)^T f + \frac{1}{2} Tr((\nabla_{xx} V) G \Sigma_\epsilon G^T) \right) \quad (5)$$

where $\Sigma_\epsilon = LL^T$ and the dependency on state is suppressed for brevity. Since the control effort enters quadratically into the cost, the optimal control takes the form:

$$u^* = -R^{-1}G^T(\nabla_x V)$$

Substituting the optimal u into (5) yields the following nonlinear, second order PDE in the cost-to-go $V(\cdot)$:

$$\begin{aligned} 0 &= q + (\nabla_x V)^T f - \frac{1}{2} (\nabla_x V)^T G R^{-1} G^T (\nabla_x V) \\ &+ \frac{1}{2} Tr((\nabla_{xx} V) B \Sigma_\epsilon B^T) \end{aligned}$$

The difficulty of solving this nonlinear, second order PDE often prevents practitioners of optimal control from attempting to solve for the value function directly. However, it has recently been found [13], [30], [26] that with the assumption

$$\lambda G(x)R^{-1}G(x)^T = B(x)\Sigma_\epsilon B(x)^T \triangleq \Sigma_t \quad (6)$$

and the logarithmic transformation

$$V = -\lambda \log \Psi \quad (7)$$

one can obtain, after substitution and simplification, the following *linear* PDE

$$0 = -\frac{1}{\lambda} q \Psi + f^T (\nabla_x \Psi) + \frac{1}{2} Tr((\nabla_{xx} \Psi) \Sigma_t) \quad (8)$$

This transformation of the value function, which we call here the *desirability*, provides an additional, computationally appealing, method through which to calculate the value function. The solution to the desirability may readily be transformed by (7) to obtain the value function, which may then be used for execution. Note that condition (6) can roughly be interpreted as a controllability-type condition: the system must have sufficient control to span (or counterbalance) the effects of input noise on the system dynamics. Furthermore, it must be ‘‘cheap’’ for the system to push in directions where noise is high, and expensive where noise is low. Additional discussion may be found in [30].

Note that Eq. (8) is in particular an *elliptic* PDE, and therefore obeys the maximum principle for elliptic PDEs [6]. This implies that there exist no local minima or maxima in the interior Ω of these HJB solutions, satisfying the Morse property of navigation functions in Definition 1.

IV. NAVIGATION FUNCTIONS THROUGH OPTIMAL CONTROL

This section will first reduce the SOC problem introduced above to the standard setting of navigation functions by sequentially incorporating the assumptions which hold in the classical navigation function setting. These successive eliminations of terms will then illuminate some connections between our approach and classical navigation function approaches. Finally, we will suggest how an approximate minimum time problem can be formulated in this approach.

A. Reduction to the Navigation Function

Dynamics. Since the classical navigation function approach implicitly decouples the trajectory generation problem from the trajectory following control design, the dynamics of a specific system are ignored. Hence, in our parallel development of the Navigation HJB equation, the dynamic term may be dropped $f := 0$. This results in the *Navigation PDE*:

$$0 = -\frac{1}{\lambda} q \Psi + \frac{1}{2} Tr((\nabla_{xx} \Psi) \Sigma_\epsilon).$$

Similarly, the classical navigation function setting does not consider spatially dependent costs. Thus, the state-dependent term in the cost function, $q(x)$, may be simplified to a free scalar parameter $q := \alpha$, producing the PDE

$$0 = -\frac{\alpha}{\lambda} \Psi + \frac{1}{2} Tr((\nabla_{xx} \Psi) \Sigma_\epsilon) \quad (9)$$

We will term this PDE as the *Augmented Navigation PDE*, as it incorporates additional cost information as compared to traditional navigation functions, but does not include the effects of system dynamics. The effect is that those states that appear only in the dynamics, and are not the workspace states, may be neglected as well. In reverse, if one wishes to include dynamics, their presence in $f(x)$ will require the additional of these states as dimensions in the HJB PDE.

Interestingly, this PDE is well known as the homogeneous *Screened Poisson Equation*, and has found applications in image processing [2]. Of interest here are the observations

that this is a second order PDE with isotropic diffusion and mass terms, a situation which has been well studied [24].

Boundary Costs. The boundary conditions for the PDE correspond to the penalty accrued as the robot exits the configuration domain and collides with an obstacle or reaches the goal state. In (3) this is represented as the terminal cost ϕ . Recall that by (7), this terminal cost must be transformed along with the value function. Thus, we have for the boundary condition

$$\Psi|_{\partial\Omega} = e^{-\frac{\phi}{\lambda}} \quad (10)$$

where $\partial\Omega$ is the boundary of the operating domain, Ω . Classically, the cost assigned to a collision has been modeled as uniform over all obstacles (Property 3 of Definition 1), and we may thus set $\phi(x_T) = c$ for an arbitrary constant c . Other choices are certainly possible.

Remark 2. We note that the boundary conditions of a linear PDE, as we have here, obey the principle of superposition. Solutions to problems may be composed at essentially zero cost, a potentially significant savings for systems that require many plans over a common domain, or require a change in the workspace. This topic is explored in [29] and [10].

The free variables $q(x)$ and R define a notion of cost, and therefore a notion of optimality. The inclusion of these variables allows us to compare navigation functions according to their perceived cost, and furthermore to declare navigation functions optimal with respect to a choice of criteria. Such criteria has traditionally been eschewed in favor of simplicity in construction of the navigation function, and hence our framework may be said to be a slight generalization, bringing notions of optimality into consideration.

Control-dependent costs. Recall that our initial definition of cost (2) includes a control dependent term. Navigation functions have traditionally been unconcerned with the control effort. Recall that the assumption on control effort and noise (6) needed to realize a linearly solvable HJB PDE is:

$$\lambda G(x)R^{-1}G(x)^T = \Sigma_t \quad (11)$$

where Σ_t is fixed as a function of the known control vector field matrix, $G(x)$, and noise characteristics, $B(x)$ and Σ_e . The control effort penalty R cannot be brought to zero naively without violating this assumption. It is possible to compensate for this limitation by using the free parameter λ to maintain the underlying relation in this assumption. That is, set $\lambda = \beta$ and define $R = \beta\tilde{R}$, yielding expressions

$$\begin{aligned} \lambda G(x) \left(\beta\tilde{R} \right)^{-1} G(x)^T &= \Sigma_t \\ G(x)\tilde{R}^{-1}G(x)^T &= \Sigma_t \end{aligned} \quad (12)$$

which is independent of β , allowing the control penalty cost to be reduced to zero. The difficulty is that as $\lambda \rightarrow 0$, (9) becomes nonsensical in the limit. Fortunately, we have assumed no cost over the states and set $\alpha = 0$ to produce the *Navigation PDE*

$$0 = Tr((\nabla_{xx}\Psi)\Sigma_t) \quad (13)$$

which is recognized to be Laplace's equation scaled according to the system noise characteristics. The practical cost incurred by this reduction of the complete SOC HJB, Equation (8), to Equation (13) is that we have eliminated consideration of control effort and state dependent penalties, which is natural in the robotics setting. Interestingly, Laplace's equation has been used previously in the generation of navigation functions [4], [16], [21]. In this prior work, the authors suggested the use of Laplace's equation, with the motivation that solutions to Laplace's equations can be shown to have no local minima over their domain. The following theorem justifies this from an optimality perspective, albeit through the transformation (7).

Theorem 3. *The optimal robust desirability function absent costs over state is given by $V = -\lambda \log \Psi$ where λ is according to (11), and Ψ is the solution to Laplace's equation over the domain Ω :*

$$\begin{aligned} 0 &= Tr((\nabla_{xx}\Psi)\Sigma_t) \\ \Psi|_{\partial\Omega} &= e^{-\frac{\phi}{\lambda}} \end{aligned}$$

Remark 4. *There is an interesting trade-off resulting from Eq. (11). Define $\tilde{\Sigma}_t \triangleq \gamma\Sigma_t$ in order that the noise may be scaled by γ . Define $\lambda = \beta$, as in (12) in order to scale the control penalty. Then $\lambda = a\beta\gamma$ for some fixed constant a . The result is that a scaling of the control effort has the same effect on the solution as a scaling of the noise, and these manifest only through the transformation (7) and the boundary conditions, and surprisingly not through the differential constraints on Ψ as γ is simply cancelled in Eq. (13). What does have an effect, however, is the directional influence of Σ_t , i.e. the solution will incorporate paths that have beneficial drift.*

Due to the exponential dependence of Eq. (7), one cannot realize the limit $\beta \rightarrow 0$. Instead, β is chosen based on magnitude of the noise, or the level of control penalty, depending on the perspective.

Remark 5. *Laplace's equation has been justified in the presence of noise here, whereas it was previously justified due to its lack of local minima. It is notable that in this case the deterministic and stochastic case share their solution, subject to a scaling.*

Two PDEs (9), (13) have been produced in this initial analysis. The first of these allows one to naturally incorporate several optimality criteria into the concept of a navigation function, while the second is especially simple and creates a connection to previously inspired navigation function formulae.

B. Approximate Time-Optimal Navigation Functions

The freedom afforded by the optimality parameters allows for the solution to be biased towards time-optimal navigation functions by penalizing time spent away from the goal. This is accomplished by setting $q(x) = c$ for some constant c . Control effort may also be adjusted through the free parameter λ and decreased relative to the state cost. The

robustness of the navigation function to noise may also be controlled through the noise characteristics defined by Σ_ϵ , and again reduced. As the value of the constant c is increased while parameters λ and Σ_ϵ are increased, cost is accrued only when the system remains outside the goal region. The optimal action during this time is to take the quickest path to the goal, ignoring the amount of control effort used.

C. Analogy with electro-statics

It is interesting that early researchers on potential field navigation methods were naturally drawn towards analogies with electro-statics. Khatib’s seminal work [15] conceptually frames the collision avoidance problem as a process of adding potential fields that would repulse or attract the point robot mass in much the same way as electrostatic fields might. We now see that this intuitive notion of attractive and repulsive forces can also be grounded in notions of optimality. Indeed, for the Navigation PDE of Eq. (13), the analogy is exact in desirability-space, with the representation of obstacles and goals manifesting identically to static charges on their surfaces. However, a logarithmic transform improves the solution from an intuitive one to one which is optimal.

D. Convergence of the solution trajectories

Due to the presence of the control cost in the solution to the HJB, it isn’t possible to directly determine the probability the system successfully reaches a goal region, where failure is dictated by the modeling assumption of stochastic forcing in (1). As mentioned in Remark 4, entirely removing the affect of this cost component isn’t possible as the necessary formulae break down in the limit. However, it is possible to take a small value for the control cost R . The cost-to-go is then predominantly governed by the cost of colliding with an inadmissible boundary. By setting the boundary conditions for obstacles to $\phi = 1$, the cost-to-go then becomes a *conservative* approximation of probability of success in reaching the goal, i.e. the cost-to-go will overstate the probability of failure by also including the cost of future trajectories’ control effort in its value. For even moderately small values of R , this approximation may not be overly conservative.

It is important to note that of the various PDEs proposed, only (13) guarantees non-collision. Specifically, the penalty accrued on the way to the goal when $q(x) \neq 0$ may be greater than the boundary penalty (for instance if a high time-penalty was applied), and the system may simply plan to hit the nearest obstacle. This may be ameliorated by simply increasing the obstacle penalty to a sufficiently high value, but calculation of such a value a-priori remains an open question.

E. Range of Navigation Functions

We briefly review the results presented up to this point. Three distinct PDEs have been presented, beginning from the HJB and then constructed by neglecting a term. Our goal was not the construction of the Laplacian, but showing its justification as in fact the optimal solution given the

specific assumptions of matched noise (6), and a neglect of state dependent cost and dynamics. It is also now clear how existing Laplacian techniques may be augmented towards optimality with a simple exponential transformation (7), as well as through the inclusion of dynamics to form the Navigation PDE (8), which may also have the state cost simply set to zero. This informs the discussion as navigation functions are brought to bear in novel domains [22].

V. NUMERICAL SOLUTIONS

The linear HJB is a PDE, and therefore appropriate for conventional numerical tools. These include techniques such as the Finite Difference Method (FDM), the Finite Element Method (FEM), or even the Boundary Element-Fast Multipole Method (FMM) [3], [24]. In a number of papers Todorov [29], [31], [32] has looked into a number of alternatives based on the structure of the discretized problem as well. In this work, we use the FDM due to its simplicity as the focus of this work is not on numeric computation.

The mentioned methods are quite fast in practice and may scale up to approximately six dimensions. However, in high dimensions such techniques fall prey to the curse of dimensionality. Recent work using sparse tensor decompositions [9] have allowed for these HJBs to be solved in dimensions up to twelve, with higher dimensions possible. Alternatively, the Feynman-Kac Lemma may be used to solve the problem through sampling [26], [13] by simulating a Brownian motion with the PDE related to its generator. Such methods scale independent of dimension (albeit with a constant that is dimension-dependent) but face the drawback that each execution of the algorithm is only able to return the solution to the PDE at an individual point.

VI. EXAMPLES

The approach is also illustrated several numerical examples. Each of these is solved with an $h = 0.1$ discretization of the domain and the use of a simple Finite Difference Method.

A. Problems with Simple Solutions

In simple domains it is possible to find an analytical or simply computed solution to (9), (13). For instance, suppose a point robot is commanded to move to a goal location located at the origin of a two-dimensional configuration space with no obstacles, while it is perturbed with noise whose characteristics are uniform across the configuration space. The solution to the associated Navigation PDE (13) corresponds to the solution of Laplace’s equation for a point potential, i.e. the fundamental solution, which is

$$\Psi = -\frac{\log(r)}{2\pi}$$

where r is the distance from the robot’s current configuration to the origin [6]. The solution to the Augmented Navigation PDE (9) for this problem may be found as follows. Taking the Fourier Transform of the equation yields

$$\left(k^2 + \frac{\alpha}{\lambda}\right) \tilde{\Psi} = 1.$$




Fig. 1. Navigation function for a two dimensional point-mass robot calculated according to Eq. (13) with varying corridor widths.

The fundamental solution is then found by solving for $\tilde{\Psi}$ and then finding the inverse Fourier Transform, which yields

$$\Psi = \frac{1}{2\pi} K_0 \left(\sqrt{\frac{\alpha}{r}} \lambda \right)$$

with K_0 the modified Bessel function of the second kind.

B. Effect of Noise on Corridor Navigation

Next, we demonstrate the method for a two dimensional robot whose task is to reach the top right corner of a square configuration space. The domain has two obstacles, creating a pair of corridors that the robot must traverse if it begins in lower portion of the configuration space. For this example $\Sigma_t = 2I_{2 \times 2}$, $\lambda = 1$, and the width of the thinner corridor is set to 1.5 and 2 distance units for comparison. The resulting solutions are shown in Fig. 1. This example shows why it can be important to include noise in the construction of a navigation function. Consider the situation when the robot starts near the bottom of the figure. In both environments, robot can travel through two different corridors to reach the goal. In both environments, the navigation functions lack local minima, as expected. In the left-hand environment, the robot can potentially choose between a wide corridor, and a medium width corridor. The choice of the corridor will depend upon the robot's specific starting configuration, as it is safe to traverse both corridors. In the right hand figure, the robot can potentially choose between the same large corridor, or a very narrow corridor (whose width is $\frac{3}{4}$ the magnitude of the noise variance). For almost all starting positions, the navigation function guides the robot away from the potentially dangerous narrow corridor, unless the robot happens to start positioned well into the narrow corridor. This intuitively logical result occurs because the potential for collision in the narrow corridor places too high a cost on that potential path.

Remark 6. *The solution to Eqs. (9), (13) take place in exponentiated coordinates, and for many examples tend to be close to zero for large regions of the state space. It is therefore usually more useful to consult the value function directly rather than examining the desirability.*

C. Maze

The second example shows that complicated environments can be well handled by this method, and also highlights




Fig. 2. Navigation function for a two dimensional point-mass robot in a maze-like environment with equi-dimensional noise. On the top-left is the standard navigation function calculated according to (13). On the top-right, the minimum time-criteria is approximated by taking $\alpha = 100$, $\lambda = .04$ in (9). The bottom-left and -right have $\Sigma_t = 4I_2, 4.6I_2$, and $\lambda = .4, .46$ respectively. The obstacles and boundary are chosen to have penalty of 20 units, while the goal region has a penalty of 0 units. The goal is located at the origin.

the effects of including additional cost criteria into the navigation function. The same robot dynamics and same noise distribution of the previous example are used, however the obstacles are placed in a more complicated maze-like pattern. The Augmented Navigation PDE of Eq. (9) is used in order to incorporate the additional optimality criteria. The resulting navigation functions are shown in Figure 2 with several noise and cost configurations..

The results compare the use of (13) with the additional cost criteria. It is seen that the solutions of (13) and (9) are qualitatively similar. As λ is decreased in the general case, this has the interpretation of either increasing the control cost weighting, R , or increasing the noise covariance. The solutions do not change qualitatively, only the magnitude of the cost-to-go. In contrast, the approximated minimal-time solution is characterized by shortest-path level sets. These level sets are characterized by straight lines near corners, in contrast to the circular level sets of other examples, as would be expected for a shortest path solution amongst piecewise-linear obstacles.

Remark 7. *For some choices of costs q, R the navigation function may bring the robot directly into an obstacle. This is no contradiction, as framing the problem through the lens of optimal control allows for freedom on the placement of boundary conditions. The penalty for hitting an obstacle, if chosen improperly, may be less than the cost to traverse the domain and enter the goal region, and thus the most economical choice is for the robot to simply collide with the boundary. This isn't a problem when using (13).*

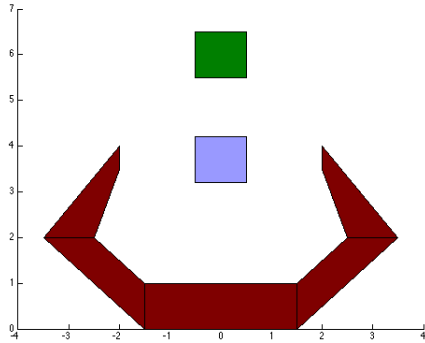


Fig. 3. An illustration of a nut grasping task. The square nut is shown in green while the gripper is shown in red. The blue region denotes a range of acceptable nut locations relative to the gripper. Each of these colored regions is transformed into configuration space, with the intersection of gripper and nut an obstacle, and intersection between the acceptable locations and the nut a goal.

D. Grasping

A final example is of a simple planar grasping task wherein a gripper must be positioned in the plane so that it may close and grasp a small nut, an illustration of which is shown in Figure 3. The goal of the problem is to move the end effector to one of a continuum of desirable locations surrounding the object in such a way that the gripper’s orientation places the jaws around the nut, whereupon a simple closing action will reliably grasp the nut. The problem is transformed into the system’s configuration space, and then solved in the Optimal Navigation framework, with the results shown in Figure 4. As the method treats the goal states as a set of boundary conditions, the relatively large goal region is handled easily.

The example illustrates the approach for a non-point mass robot, and illustrates the smoothness of the navigation function over the domain in a typical manipulation task. It is easily to see the optimal path of the gripper, wherein it smoothly rotates, following the basin of low cost in blue, until the nut is captured.

VII. DISCUSSION

This paper introduced a generalization of the navigation function framework to include system noise, dynamics, and cost criteria. Philosophically, this paper links the classical robotics subject of navigation functions with recent advances in Stochastic Optimal Control. Previous results that were developed on a somewhat ad-hoc basis have been shown to be related to optimality considerations, and the intuition which led to their development well placed. Remarkably, many existing results using harmonic potentials can be shown to be optimal for a particular configuration of noise and cost models when the solution is simply adjusted by a logarithmic transformation.

We also showed how navigation functions can approximately incorporate minimum time task requirements. From a practical point of view, the methods developed here to construct the navigation functions can be applied to en-




Fig. 4. Cross sections of navigation function for the nut grasping task illustrated in Figure 3. Parameters used are $\alpha = 0.02$, $R = .02I_{3 \times 3}$, $\Sigma_t = 5I_{3 \times 3}$ and the boundary costs are set to $\phi = 1$. Spatial discretization in the x - and y -coordinates are $h_x = h_y = 0.25$ and in the angular direction $h_\theta = 20^\circ$. The goal region isosurface is displayed in dark blue.

vironments, obstacles, and robots with arbitrary smooth geometries. Furthermore, it allows for the results of existing methods (i.e. [19]) to be compared against the underlying *optimal* solution to the problem, if the system’s noise characteristics are captured by our model. The ability to solve for the navigation function in terms of a linear PDE also expands the space of algorithms for calculating solutions.

Our numerical experiments show that solutions in configurations space having dimension up to 5 can be computed on a desktop PC computer. Future work will explore specialized numerical algorithms in an attempt to expand the practical computational limits of the approach. Additionally, it may be possible to use the results presented here to inform the choice of artificial potential fields, yielding navigation functions that may be assembled quickly, but that better approximate the optimal navigation function for the problem.

REFERENCES

- [1] D. P. Bertsekas. *Dynamic programming and optimal control*. Athena Scientific, 1995.
- [2] P. Bhat, B. Curless, M. Cohen, and C. L. Zitnick. Fourier analysis of the 2D screened Poisson equation for gradient domain problems. In *European Conference on Computer Vision (ECCV)*, pages 114–128. Springer, 2008.
- [3] R. Coifman, V. Rokhlin, and S. Wandzura. The fast multipole method for the wave equation: A pedestrian prescription. *Antennas and Propagation Magazine, IEEE*, 35(3):7–12, 1993.
- [4] C. I. Connolly, J. B. Burns, and R. Weiss. Path planning using Laplace’s equation. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2102–2106, 1990.
- [5] D. P. de Farias and B. Van Roy. The linear programming approach to approximate dynamic programming. *Operations Research*, 51(6):850–865, 2003.
- [6] L. C. Evans. *Partial Differential Equations (Graduate Studies in Mathematics, V. 19)*. American Mathematical Society, 1998.
- [7] W. H. Fleming and H. M. Soner. *Controlled Markov processes and viscosity solutions*, volume 25. springer New York, 2006.
- [8] M. B. Horowitz. *Efficient methods for stochastic optimal control*. PhD thesis, California Institute of Technology, 2014.

- [9] M. B. Horowitz, A. Damle, and J. W. Burdick. Linear Hamilton Jacobi Bellman equations in high dimensions. In *IEEE Conference on Decision and Control (CDC)*, 2014. Arxiv:1404.1089v1.
- [10] M. B. Horowitz, E. Wolff, and R. M. Murray. A compositional approach to stochastic optimal control with co-safe temporal logic specifications. *International Conference on Intelligent Robots and Systems (IROS)*, 2014.
- [11] A. Howard, M. J. Matarić, and G. S. Sukhatme. Mobile sensor network deployment using potential fields: A distributed, scalable solution to the area coverage problem. In *Distributed Autonomous Robotic Systems*, pages 299–308. Springer, 2002.
- [12] A. Jadbabaie, J. Yu, and J. Hauser. Unconstrained receding-horizon control of nonlinear systems. *IEEE Transactions on Automatic Control*, 46(5):776–783, 2001.
- [13] H. J. Kappen. Linear Theory for Control of Nonlinear Stochastic Systems. *Physical Review Letters*, 95(20):200201, Nov. 2005.
- [14] H. J. Kappen. Path integrals and symmetry breaking for optimal control theory. *Journal of Statistical Mechanics: Theory and Experiment*, 2005(11):P11011–P11011, Nov. 2005.
- [15] O. Khatib. Real-time obstacle avoidance for manipulators and mobile robots. In *IEEE International Conference on Robotics and Automation*, pages 500–505, 1985.
- [16] 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, June 1992.
- [17] D. Koditschek. Exact robot navigation by means of potential functions: Some topological considerations. In *IEEE International Conference on Robotics and Automation*, pages 1–6, 1987.
- [18] D. E. Koditschek and E. Rimon. Robot navigation functions on manifolds with boundary. *Advances in Applied Mathematics*, 11(4):412–442, Dec. 1990.
- [19] K. Konolige. A gradient method for realtime robot control. In *IEEE International Conference on Intelligent Robots and Systems (IROS)*, volume 1, pages 639–646, 2000.
- [20] Y. Koren and J. Borenstein. Potential field methods and their inherent limitations for mobile robot navigation. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 1398–1404, 1991.
- [21] S. G. Loizou. Closed form navigation functions based on harmonic potentials. In *IEEE Conference on Decision and Control and European Control Conference (CDC-ECC)*, pages 6361–6366, 2011.
- [22] S. G. Loizou. Navigation functions in topologically complex 3-D workspaces. In *American Control Conference (ACC)*, pages 4861–4866, 2012.
- [23] V. J. Lumelsky and T. Skewis. Incorporating range sensing in the robot navigation function. *IEEE Transactions on Systems, Man and Cybernetics*, 20(5):1058–1069, 1990.
- [24] C. Pozrikidis. *Introduction to finite and spectral element methods using Matlab*. CRC Press, 2005.
- [25] E. Rimon and D. E. Koditschek. Exact robot navigation using artificial potential functions. *IEEE Transactions on Robotics and Automation*, 8(5):501–518, 1992.
- [26] E. Theodorou. *Iterative path integral stochastic optimal control: Theory and applications to motor control*. PhD thesis, University of Southern California, June 2011.
- [27] E. A. Theodorou, J. Buchli, and S. Schaal. Path integral-based stochastic optimal control for rigid body dynamics. In *IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL)*, pages 219–225, 2009.
- [28] E. Todorov. Linearly-solvable Markov decision problems. *Advances in neural information processing systems (NIPS)*, 19:1369, 2007.
- [29] E. Todorov. Compositionality of optimal control laws. *Advances in neural information processing systems (NIPS)*, 22:1856–1864, 2009.
- [30] E. Todorov. Efficient computation of optimal actions. *Proceedings of the national academy of sciences*, 106(28):11478–11483, 2009.
- [31] E. Todorov. Eigenfunction approximation methods for linearly-solvable optimal control problems. *IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL)*, pages 161–168, 2009.
- [32] M. Zhong and E. Todorov. Aggregation Methods for Linearly-solvable Markov Decision Process. *World Congress of the International Federation of Automatic Control (IFAC)*, 18(1):11220–11225, 2011.
- [33] K. Zhou, J. C. Doyle, K. Glover, et al. *Robust and optimal control*, volume 40. Prentice Hall New Jersey, 1996.