

# Cross-Entropy Randomized Motion Planning

Marin Kobilarov

**Abstract**—This paper is concerned with motion planning for nonlinear robotic systems operating in constrained environments. Motivated by recent developments in sampling-based motion planning and Monte Carlo optimization we propose a general randomized path planning method based on sampling in the space of trajectories. The idea is to construct a probability distribution over the set of feasible paths and to perform the search for an optimal trajectory through importance sampling. At the core of the approach lies the cross-entropy method for estimation of rare-event probabilities. The algorithm recursively approximates the optimal sampling distribution which guides the set of sampled trajectories towards regions of progressively lower cost until converging to a delta distribution at the optimum. Our main goal is to provide a framework for consistent adaptive sampling correlating the spatial structure of trajectories and their computed costs. The approach is illustrated with two simple examples—a point mass vehicle and the Dubins car, and is then applied to a simulated helicopter flying optimally in a 3-D terrain.

## I. INTRODUCTION

Consider a robotic vehicle with state  $x \in \mathcal{X}$  controlled using actuator inputs  $u \in \mathcal{U}$ , where  $\mathcal{X}$  is the state space and  $\mathcal{U}$  denotes the set of controls. The vehicle dynamics is described by the function  $f : \mathcal{X} \times \mathcal{U} \rightarrow T\mathcal{X}$  defined by

$$\dot{x}(t) = f(x(t), u(t)), \quad (1)$$

which is used to evolve the vehicle state forward in time. In addition, the vehicle is subject to constraints arising from actuator bounds and obstacles in the environment. These constraints are expressed through the vector of inequalities

$$F(x(t)) \geq 0, \quad (2)$$

for all  $t \in [0, t_f]$ , where  $t_f > 0$  is the final time of the trajectory. The goal is to compute the optimal controls  $u^*$  and time  $t_f^*$  driving the system from its initial state  $x_0 \in \mathcal{X}$  to a given goal region  $\mathcal{X}_f \subset \mathcal{X}$ , i.e.

$$\begin{aligned} (u^*, t_f^*) &= \operatorname{argmin}_{u, t_f} \int_0^{t_f} C(u(t), x(t)) dt, \\ \text{subject to } \dot{x}(t) &= f(x(t), u(t)), \\ F(x(t)) &\geq 0, \quad x(0) = x_0, \quad x(t_f) \in \mathcal{X}_f \end{aligned} \quad (3)$$

for all  $t \in [0, t_f]$  and where  $C : \mathcal{U} \times \mathcal{X} \rightarrow \mathbb{R}$  is a given cost function. A typical cost function includes a time component and a control effort component, i.e.  $C(u, x) = 1 + \lambda \|u\|^2$  where  $\lambda \geq 0$  is a chosen weight.

The problem (3) has no closed-form solution since both the dynamics (1) and constraints (2) are nonlinear. Gradient-based optimization is not suitable unless a good starting guess is chosen since the constraints (2) impose many local minima. In addition, constraints corresponding to arbitrary obstacles can be non-smooth and require special differentiation [3] to guarantee convergence. An alternative is to discretize the vehicle

state space  $\mathcal{X}$ , e.g. using a grid and generate candidate paths by transitioning between adjacent cells. Such an approach is generally limited to few dimensions, e.g.  $\dim(\mathcal{X}) \leq 3$  and to systems with very simple dynamics, e.g. an unconstrained point mass in the plane. This is due to the exponential (both in state dimension and planning horizon  $T$ ) size of the search space, also known as the *curse of dimensionality*.

Since the motion planning problem (3) is computationally NP-complete in general [13] one has to resort to approximation algorithms. *Sampling-based motion planning* has become an established methodology in this context. The basic idea is to construct a graph structure with nodes corresponding to states and with edges satisfying the dynamics (1) and constraints (2). In essence, the graph is regarded as a finite approximation of the infinite set of feasible trajectories. The optimal control problem is then solved approximately through graph search. The two main families of such methods are *rapidly-exploring dense trees* (RDT) [13] and *probabilistic roadmaps* (PRM) [2]. In recent years considerable attention has been given to augmenting the basic sampling-based motion planners with strategies for accelerating the search. A common realization is that the complex problem can be solved effectively through a proper balance of *exploration and exploitation* [16, 17] and by guiding the algorithm based on acquired information.

Motivated by these developments we propose a technique for adaptive sampling which uses the cost of computed solutions to guide the search. In contrast to standard motion planning, we consider *sampling from the space of trajectories* rather than sampling nodes from the state space  $\mathcal{X}$  alone. In essence, the problem is solved as a random search by optimally exploiting the gathered information from prior samples. At the core of our approach lies the *cross-entropy* (CE) method [19, 9] initially developed for estimation of rare-event probabilities and later employed as a general optimization framework. The CE method is widely applicable and is used to successfully solve complex combinatorial problems such as the minimum graph cut or the traveling salesman problems. The basic idea behind applying the CE approach to motion planning is to recursively iterate the two steps:

- 1) Generate trajectories sampled from a distribution and compute their costs
- 2) Update the distribution using a subset of “high-quality” trajectories

until the set of sampled trajectories becomes concentrated around the optimum, or equivalently until the distribution has converged to a delta function. The scheme is general and converges to an optimum assuming that enough feasible trajectories can be sampled. As with most randomized methods there is no guarantee that this would be the global optimum but

with high likelihood the global optimum can be approached if enough samples are generated.

We first reformulate the motion planning problem into a finite dimensional constrained optimization in §II. In §III we provide a quick introduction to the required probabilistic techniques for the CE approach. The main cross-entropy algorithm is then given in §IV. Two simple examples are used to illustrate the theory in §V. The link to standard sampling-based motion planning is then elaborated in §VI. The method is finally applied to a complex high-dimensional problem: a simulated helicopter navigating optimally in an urban terrain.

## II. PROBLEM FORMULATION

A trajectory recording the controls and states over the time interval  $[0, t_f]$  is denoted by the function  $\pi : [0, t_f] \rightarrow \mathcal{U} \times \mathcal{X}$ , i.e.  $\pi(t) = (u(t), x(t))$  for all  $t \in [0, t_f]$ . A given control curve  $u : [0, t_f] \rightarrow \mathcal{U}$  determines a unique state trajectory  $x : [0, t_f] \rightarrow \mathcal{X}$  by evolving the dynamics from an initial state  $x_0 \in \mathcal{X}$ . This is more formally defined through the *integral flow* of the ODE (1) denoted by  $\Phi_u : \mathcal{X} \times \mathbb{R} \rightarrow \mathcal{X}$  in terms of

$$x(t) = \Phi_u(x_0, t), \forall t > 0. \quad (4)$$

The space of all trajectories originating at point  $x_0$  and satisfying the dynamics is denoted

$$\mathcal{P} = \{\pi : t \in [0, +\infty] \rightarrow (u(t), x(t)) \mid x(t) = \Phi_u(x_0, t)\}.$$

Consider a finite-dimensional parametrization of trajectories in terms of vectors  $z \in \mathcal{Z}$  where  $\mathcal{Z} \subset \mathbb{R}^{n_z}$  is the *parameter space*. Assume that the parametrization is given by the function  $\varphi : \mathcal{Z} \rightarrow \mathcal{P}$  according to

$$\pi = \varphi(z) \equiv \varphi_z.$$

The (control, state) tuples along a trajectory parametrized by  $z$  are written as  $\pi(t) = \varphi_z(t)$ . Note that we do not impose any restrictions on  $\mathcal{Z}$  other than  $\varphi(\mathcal{Z}) \subset \mathcal{P}$ . Ideally  $\varphi(\mathcal{Z})$  should be a “good” approximation of  $\mathcal{P}$ , i.e. for any given  $\pi \in \mathcal{P}$  there should always exist a  $z \in \mathcal{Z}$  for which  $\varphi(z)$  is “close” to  $\pi$ .

Let  $\kappa : \mathcal{U} \times \mathcal{X} \rightarrow \mathcal{X}$  project onto the state component, i.e.  $\kappa(u, x) = x$ . The *constrained parameter space*  $\mathcal{Z}_{\text{con}} \subset \mathcal{Z}$  is the set of parameters satisfying the boundary conditions and constraints and is defined by

$$\mathcal{Z}_{\text{con}} = \{z \in \mathcal{Z} \mid F(\kappa(\varphi_z(t_f))) \geq 0, \kappa(\varphi_z(t_f)) \in \mathcal{X}_f\}, \quad (5)$$

for some  $t_f > 0$  and all  $t \in [0, t_f]$ . Define the *cost function*  $J : \mathcal{Z} \rightarrow \mathbb{R}$  according to

$$J(z) = \int_0^{t_f} C(\varphi_z(t)) dt.$$

Problem (3) can now be equivalently restated as finding the optimal  $t_f^*$  and  $(x^*, u^*) = \varphi(z^*)$  such that

$$z^* = \arg \min_{z \in \mathcal{Z}_{\text{con}}} J(z). \quad (6)$$

We consider problems over a continuous set  $\mathcal{Z}$  since the discrete case is too constraining for problems involving dynamics. In addition, for systems with very complicated dynamics  $\mathcal{Z}$  can be a product of discrete and continuous spaces as in the hybrid system abstraction considered in §VII.

## III. BACKGROUND ON MONTE CARLO OPTIMIZATION

We provide a quick introduction to the probabilistic framework for solving the optimization (6) through sampling trajectories from  $\mathcal{Z}_{\text{con}}$ . Our development follows closely [18] using notation adapted to our setting.

### A. Importance Sampling

Consider the estimation of the following expression

$$\ell = \mathbb{E}_p[H(Z)] = \int H(z)p(z)dz, \quad (7)$$

where  $H : \mathcal{Z} \rightarrow \mathbb{R}$  is some non-negative performance metric and  $p$  is the probability density of  $Z$ . Assume that there is another dominating<sup>1</sup> probability density  $q$  which is easy to evaluate and sample from, such as a Gaussian. The integral (7) can be expressed as

$$\ell = \int H(z) \frac{p(z)}{q(z)} q(z) dz = \mathbb{E}_q \left[ H(Z) \frac{p(Z)}{q(Z)} \right]. \quad (8)$$

The density  $q$  is called the *importance density* and can be used to evaluate the integral using i.i.d. random samples  $Z_1, \dots, Z_N$  from  $q$  so that

$$\hat{\ell} = \frac{1}{N} \sum_{i=1}^N H(Z_i) \frac{p(Z_i)}{q(Z_i)} \quad (9)$$

is an unbiased estimator of  $\ell$ . An important question is then how to select a good density  $q$ . The most natural choice is the density that minimizes the variance of the estimator  $\hat{\ell}$ , i.e.

$$q^* = \arg \min_q \mathbb{V}_q \left( H(Z) \frac{p(Z)}{q(Z)} \right),$$

the solution to which is

$$q^*(z) = \frac{H(z)p(z)}{\ell} \quad (10)$$

since  $\mathbb{V}_{q^*}(\ell) = 0$ . The density  $q^*$  is called the *optimal importance sampling density*. Of course, this density is only hypothetical and cannot be implemented in practice since it involves the value of  $\ell$  which is what is being estimated in the first place!

A natural way to find a density  $q$  that is closest to  $q^*$  is in the Kullback-Leibler (KL) sense, i.e. with minimum cross-entropy distance between  $q^*$  and  $q$ . The KL distance between any two given distributions  $q$  and  $p$  is defined by

$$\mathfrak{D}(p, q) = \int p(z) \ln p(z) dz - \int p(z) \ln q(z) dz. \quad (11)$$

and the required  $q$  solves the following optimization

$$\min_q \mathfrak{D}(q^*, q). \quad (12)$$

We next consider the case when  $Z$  has a pdf  $p(\cdot; \bar{v})$  belonging to some parametric family  $\{p(\cdot; v), v \in \mathcal{V}\}$  where  $\bar{v}$  is the true or nominal parameter. For instance, this could be a mixture of Gaussians. It is natural to consider an importance density  $q$  from the same family. Its optimal parameter  $v$  is found through the parametric optimization

$$\min_v \mathfrak{D}(q^*, p(\cdot, v))$$

<sup>1</sup>the density  $q$  dominates  $Hp$  when  $q(z)=0 \Rightarrow H(z)p(z)=0$ .

This is equivalent to maximizing with respect to  $v$

$$\int H(z)p(z, \bar{v}) \ln p(z, v) dz,$$

which is obtained using (10) and (11). In other words, the optimal importance density parameter  $v^*$  can be found as

$$v^* = \operatorname{argmax}_v \mathbb{E}_{\bar{v}}[H(Z) \ln p(Z, v)]. \quad (13)$$

Finally, the optimal parameter can be approximated numerically by

$$\hat{v}^* = \operatorname{argmax}_{v \in \mathcal{V}} \frac{1}{N} \sum_{i=1}^N H(Z_i) \ln p(Z_i, v), \quad (14)$$

where  $Z_1, \dots, Z_N$  are i.i.d. samples from  $p(\cdot, \bar{v})$ .

### B. Estimation of Rare-Event Probabilities

Consider the estimation of the probability that a trajectory  $z \in \mathcal{Z}_{\text{con}}$  sampled from  $p(\cdot; \bar{v})$  has a cost  $J(z)$  smaller than a given constant  $\gamma$ . It is defined as

$$\ell = \mathbb{P}_{\bar{v}}(J(Z) \leq \gamma) = \mathbb{E}_{\bar{v}}[I_{\{J(Z) \leq \gamma\}}], \quad (15)$$

where  $I_{\{\cdot\}}$  is the indicator function. This can be computed approximately using (9) according to

$$\hat{\ell} = \frac{1}{N} \sum_{i=1}^N I_{\{J(Z_i) \leq \gamma\}} \frac{p(Z_i; \bar{v})}{p(Z_i; v)},$$

where  $Z_1, \dots, Z_N$  are i.i.d. samples from  $p(\cdot, v)$ . In order to determine the optimal  $v$  for this computation we can employ (14) to obtain

$$\hat{v}^* = \operatorname{argmax}_{v \in \mathcal{V}} \frac{1}{N} \sum_{i=1}^N I_{\{J(Z_i) \leq \gamma\}} \ln p(Z_i, v), \quad (16)$$

where  $Z_1, \dots, Z_N$  are i.i.d. samples from  $p(\cdot, \bar{v})$ . The problem is that when  $\{J(Z) \leq \gamma\}$  is a rare event, this approximation is meaningless because there will be almost no samples  $z$  with  $J(z) \leq \gamma$  and  $\hat{\ell}$  will be incorrectly estimated as zero!

The idea behind the CE method is to employ a multilevel approach using a sequence of parameters  $\{v_j\}_{j \geq 0}$  and levels  $\{\gamma_j\}_{j \geq 1}$ . At the end the sequence converges to the optimal  $v^*$  which then can be used to estimate the integral  $\hat{\ell}$  correctly. The procedure starts by drawing  $N$  samples  $Z_1, \dots, Z_N$  using an initial parameter  $v_0$ , for instance  $v_0 = \bar{v}$ . Let  $\varrho$  be a small number, e.g.  $10^{-2} \leq \varrho \leq 10^{-1}$ . The value  $\gamma_1$  is set to the  $\varrho$ -th quantile of  $H(Z)$ , i.e.  $\gamma_1$  is the largest real number for which

$$\mathbb{P}_{v_0}(H(Z) \leq \gamma_1) = \varrho.$$

The level  $\gamma_1$  can be computed approximately by sorting the costs of the samples  $J(Z_1), \dots, J(Z_N)$  in an increasing order, say  $J_1 \leq \dots \leq J_N$ , and setting  $\hat{\gamma}_1 = J_{\lceil \varrho N \rceil}$ . The optimal parameter  $v_1$  for level  $\hat{\gamma}_1$  is then estimated using (14) by replacing  $\gamma$  with  $\hat{\gamma}_1$ .

Note that the samples with costs  $J_1, \dots, J_{\lceil \varrho N \rceil}$  will also be the samples used to estimate  $v_1$ . They form the *elite set*, i.e. the  $\varrho$ -fraction of the  $N$  samples with the best costs. The procedure then iterates to compute the next  $\gamma_i$  and  $v_i$  and terminates when  $\gamma_i \leq \gamma$ . At this point we set  $v = v_i$  as the optimal

parameter corresponding to the originally given level  $\gamma$  and the probability of  $J(Z) \leq \gamma$  is computed using  $v$ . In summary, each iteration of the algorithm perform two steps, starting with  $v_0$ ,

- 1) *Sampling and updating of  $\gamma_j$* : Sample  $Z_1, \dots, Z_n$  from  $p(\cdot, \hat{v}_{i-1})$  and compute the  $\varrho$ -th quantile  $\hat{\gamma}_t$ .
- 2) *Adaptive Updating of  $v_j$* : Compute  $\hat{v}_j$  such that

$$\hat{v}_j = \operatorname{argmin}_{v \in \mathcal{V}} \frac{1}{|\mathcal{E}_j|} \sum_{Z_k \in \mathcal{E}_j} \ln p(Z_k; v), \quad (17)$$

where  $\mathcal{E}_j$  is the *elite set* of samples, i.e. samples  $Z_k$  for which  $J(Z_k) \leq \hat{\gamma}_j$ .

### C. Cross-Entropy Optimization

The idea behind the cross-entropy method is to treat the optimization (6) as an *estimation problem of rare-event probabilities*. Define the cost function optimum  $\gamma^*$  by

$$\gamma^* = \min_{z \in \mathcal{Z}_{\text{con}}} J(z).$$

Finding the optimal trajectory then amounts to iterating the rare-event simulation steps defined in §III-B until the cost  $\gamma_j$  approaches  $\gamma^*$ . Typically, after a finite number of iterations  $p(\cdot, v_j)$  will approach a delta distribution and all samples  $Z_i$  will become almost identical. This signifies that the optimum has been found and  $z^*$  is set to the sample with lowest cost. Note that the term ‘‘optimal’’ should be used with caution because although the method explores the state space globally it might still converge to a local solution if, for instance, no samples were obtained near the true global value.

## IV. MOTION PLANNING

We now construct a general motion planning algorithm based on the cross-entropy method. We choose to represent the importance density using a Gaussian mixture model since it is a compact way to encode a rich set of trajectories across multiple homotopy classes.

The parameter space is  $\mathcal{V} = (\mathbb{R}^{n_z} \times \mathbb{R}^{(n_z^2 + n_z)/2})^K \times \mathbb{R}^K$  with elements  $v = (\mu_1, \Sigma_1, \dots, \mu_K, \Sigma_K, w_1, \dots, w_K)$  corresponding to  $K$  mixture components with means  $\mu_k$ , covariance matrices  $\Sigma_k$  (excluding identical elements due to the matrix symmetry) and weights  $w_k$ . The density is defined as

$$p(z; v) = \sum_{k=1}^K \frac{w_k}{\sqrt{(2\pi)^{n_z} |\Sigma_k|}} e^{-\frac{1}{2}(z - \mu_k)^T \Sigma_k^{-1} (z - \mu_k)},$$

where  $\sum_{k=1}^K w_k = 1$ . The number of mixture components  $K$  is chosen adaptively (see e.g. [4]). Even the simplest case with  $K = 1$  is capable of solving complex multi-extremal problems globally. The complete algorithm is summarized below.

#### Algorithm 4.1: CE Motion Planning

- 1) Choose initial trajectory samples  $Z_1, \dots, Z_N$  so that the set  $\hat{\mathcal{P}} = \{\varphi_{Z_i}\}_{i=1}^N$  approximates (sparsely) the set of feasible trajectories over  $\mathcal{X}$ ; Set  $j = 0$  and  $\hat{\gamma}_0 = \infty$
- 2) Update  $\hat{v}_j$  using (17) (e.g. by EM) over the elite set  $\mathcal{E}_j = \{Z_i \mid J(Z_i) \leq \hat{\gamma}_j\}$

- 3) Generate samples  $Z_1, \dots, Z_N$  from  $p(\cdot, v_j)|_{\mathcal{Z}_{\text{con}}}$  and compute the  $\varrho$ -th quantile  $\hat{\gamma}_{j+1} = J_{\lceil \varrho N \rceil}$
- 4) If a stopping criteria is met then finish, otherwise set  $j = j + 1$  and goto step (2)

There are several important points determining the success of the approach.

*a) Initialization.:* The initial set of samples should be generated (step (1)) to achieve a good coverage of  $\mathcal{X}$ . The samples do not need to have good quality in terms of the cost  $J$ , they only need to achieve a reasonable exploration of  $\mathcal{X}$ . They can be generated using an initial parameter  $v_0$ , i.e.  $Z_i \sim p(\cdot, v_0)$  where  $v_0$  can encode some a priori knowledge about the problem. A more general approach that does not rely on any problem assumptions is to perform uniform sampling over  $\mathcal{Z}_{\text{con}}$ . This corresponds to a high-entropy uninformative random start also proven effective in the context of mixture models [15]. An alternative is also to generate the initial set using a classical motion planning roadmap [13, 2].

*b) Parameter Update.:* The optimal parameter update (step (2)) is accomplished using an Expectation-Minimization (EM) algorithm [15] for  $K \geq 2$ . In the case  $K = 1$  the components of  $\hat{v}_j$  are updated simply as

$$\mu = \frac{1}{|\mathcal{E}_j|} \sum_{Z_k \in \mathcal{E}_j} Z_k, \quad \Sigma = \frac{1}{|\mathcal{E}_j|} \sum_{Z_k \in \mathcal{E}_j} (Z_k - \mu)(Z_k - \mu)^T.$$

As the algorithm iterates, the uncertainty volume (i.e. the determinants of the covariances  $\Sigma_k$ ) is shrinking towards a delta distribution. This might happen prematurely before reaching a good quality solution if, for instance, too few samples were used. To prevent such degeneracy it is useful to inject a small amount of noise with variances  $\nu \in \mathbb{R}^{n_z}$  (see [1]) i.e. by setting  $\Sigma_k = \Sigma_k + \text{diag}(\nu)$  for each  $k = 1, \dots, K$ .

*c) Sampling.:* The sampling step (3) in Algorithm (4.1) is accomplished using a standard accept-reject argument and is given below only for convenience.

#### Generating Samples over $\mathcal{Z}_{\text{con}}$

- 1) Compute  $A_k = \sqrt{\Sigma_k}$  for all  $k = 1, \dots, K$  and set  $i = 1$
- 2) Choose  $k \in \{1, \dots, K\}$  proportional to  $w_k$
- 3) Sample  $r \sim \mathcal{N}_{n_z}(0, 1)$  and set  $Z_i = \mu_k + A_k r$
- 4) if  $Z_i \notin \mathcal{Z}_{\text{con}}$  then go to line (3)
- 5) if  $i = N$  then finish; else set  $i = i + 1$  and goto line (2)

Note that the only expensive operation is the square-root (using e.g. Cholesky decomposition) on line (1) but it is performed only  $K$  times *before* drawing all  $N$  samples. It is also possible to handle constrained sampling more efficiently through a local search to detect the boundary of  $\mathcal{Z}_{\text{con}}$  and escape from  $\mathcal{Z} \setminus \mathcal{Z}_{\text{con}}$ .

## V. EXAMPLES

Consider a physical workspace (i.e. an environment) denoted by  $\mathcal{W}$ , where  $\mathcal{W} = \mathbb{R}^2$ , or  $\mathcal{W} = \mathbb{R}^3$  [12]. The workspace contains a number of *obstacles* denoted by  $\mathcal{O}_1, \dots, \mathcal{O}_{n_o} \subset \mathcal{W}$  with which the vehicle must not collide. Assume that the vehicle at state  $x \in \mathcal{X}$  is occupying a region  $\mathcal{A}(x) \subset \mathcal{W}$  and that the function  $\mathbf{prox}(\mathcal{A}_1, \mathcal{A}_2)$  returns the closest Euclidean distance between two sets  $\mathcal{A}_{1,2} \subset \mathcal{W}$  and is negative if they

intersect. One of the constraints defined in (2) is then to avoid obstacles, generally expressed as

$$F_1(x(t)) = \min_i \mathbf{prox}(\mathcal{A}(x(t)), \mathcal{O}_i), \text{ for all } t \in [0, \infty]. \quad (18)$$

We demonstrate the cross-entropy method to two simple types of vehicles navigating in a planar obstacle field—a point mass double integrator and a Dubins car.

### A. Piece-wise Cubic Spline

Consider a point mass vehicle operating in the plane with state space  $\mathcal{X} = \mathbb{R}^2 \times \mathbb{R}^2$  and state  $x = (q, v)$ . It has a double-integrator dynamics given by  $\dot{q} = v$  and  $\dot{v} = u$  where the controls  $u \in \mathcal{U} = \mathbb{R}^2$  are simply the accelerations. Assume that the cost function is of the form

$$C(u, x) = \|v\| + \lambda \|u\|^2,$$

i.e. a combination of the length of the path and its acceleration where  $\lambda > 0$  controls the curve smoothness. The goal set consists of a single point  $x_f \in \mathcal{X}$ , i.e.  $\mathcal{X}_f = \{x_f\}$ .

*1) Trajectory Generation:* A trajectory is parametrized using  $m$  “knots”  $(q_i, v_i) \in \mathbb{R}^4$ ,  $i = 1, \dots, m$ . The parameter space is  $\mathcal{Z} = \mathbb{R}^{4m}$  with state  $z = (q_1, v_1, \dots, q_m, v_m)$ . Denote the given start state by  $x_0 = (q_0, v_0)$  and denote the goal state by  $x_{m+1} \equiv x_f$ . In essence, the trajectory will be represented by  $m + 1$  cubic splines joining the pairs of states  $x_i \rightarrow x_{i+1}$ , for  $i = 0, \dots, m$ . Since the cost function does not involve time we can use simply set  $t_f = 1$  and generate all trajectories over the time interval  $[0, 1]$ . Let  $t_i = i/(m + 1)$  denote the starting time of the segment beginning at state  $x_i$ . Over the interval  $[t_i, t_{i+1}]$  for each  $i = 0, \dots, m$  the function  $\varphi_z = (u, q, v)$  takes the form

$$\begin{aligned} u(t) &= 6a_i \bar{t} + 2b_i, \\ v(t) &= 3a_i \bar{t}^2 + 2b_i \bar{t} + c_i, \\ q(t) &= a_i \bar{t}^3 + b_i \bar{t}^2 + c_i \bar{t} + d_i, \end{aligned}$$

where  $\bar{t} := t - t_i$ ,  $c_i = 3q_{i+1} - 3q_i - 2v_i - v_{i+1}$ ,  $d_i = -2q_{i+1} + 2q_i + v_i + v_{i+1}$ .

*2) Numerical Results:* Consider an environment depicted on Fig. 1. It contains one concave obstacle and several homotopy classes which create an interesting planning scenario. Algorithm 4.1 is run using  $N = 100$  and  $\varrho = .1$  and one mixture component  $K = 1$ . Initial sampling is achieved using a “nominal” trajectory computed by ignoring the obstacles and consisting of equally spaced knots along a straight line between start and goal. The first sampling iteration is performed using a distribution centered at this nominal parameter. Fig. 1 shows that it quickly converges to the globally optimal solution (known in advance) in a few iterations. Empirically Fig. 2 shows that the cost function decreases nearly exponentially in the number of iterations.

### B. Dubins Car

The Dubins car is one of the simplest vehicle models with non-holonomic constraints. Its state space is  $\mathcal{X} = SE(2)$  with state  $x = (r_x, r_y, \theta)$  denoting the car position  $(r_x, r_y)$  and orientation  $\theta$ . The dynamics is defined as

$$\dot{r}_x = v \cos \theta, \quad \dot{r}_y = v \sin \theta, \quad \dot{\theta} = u, \quad (19)$$

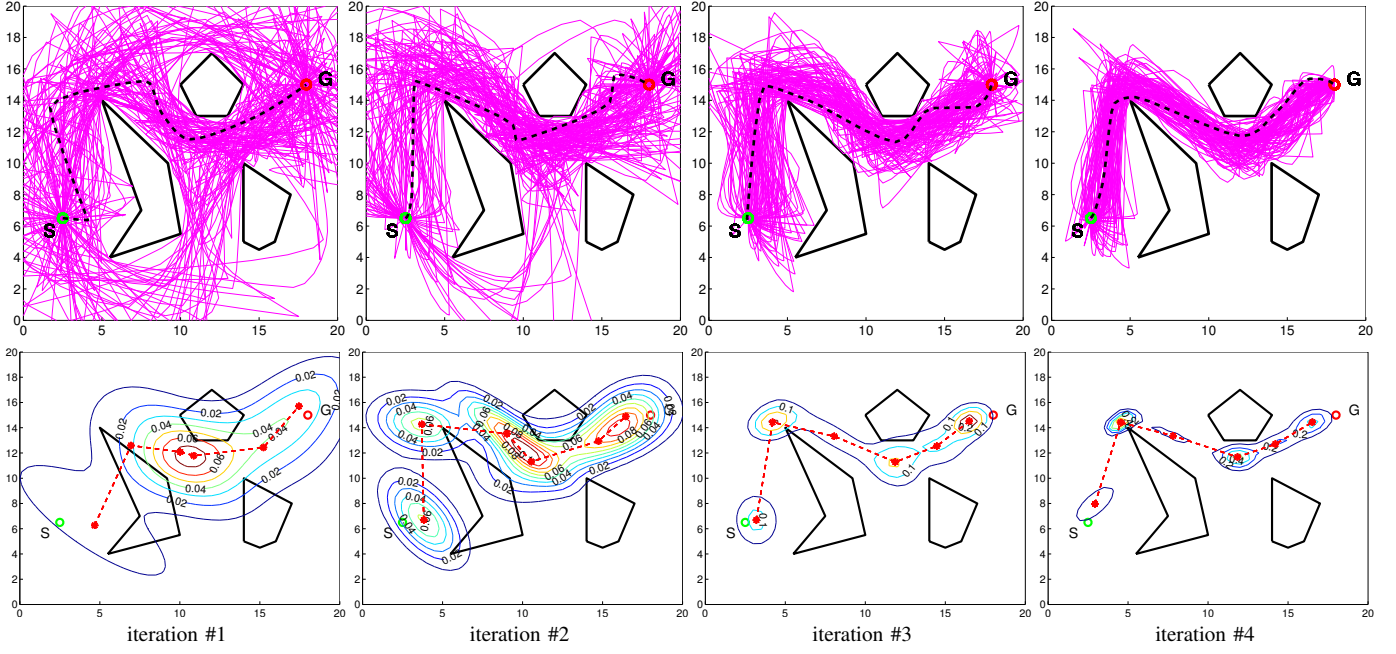


Fig. 1. The first four iterations of the CE algorithm 4.1, i.e.  $j = 1, \dots, 4$ , applied to a point mass vehicle among obstacles. The upper plots show the sampled trajectories  $\varphi(Z_1), \dots, \varphi(Z_N)$  and the current optimal path  $\varphi(z^*)$  (dashed). The lower plots visualize  $p(\cdot, \hat{v}_j)$  as the level sets of another density over the  $(x, y)$ -position space. The density of each point corresponds to the density of the best trajectory passing through it. The algorithm starts from a nearly uniform state space exploration and converges towards a delta distribution at the global optimum.




Fig. 2. The same scenario as in Fig. (1) with a) showing the cost  $J(z^*)$  as a function of the algorithm iterations and b) showing the resulting optimal path  $\varphi(z^*)$ .

where  $v > 0$  is a given constant speed and  $u$  is the controllable turn rate. The control space is  $\mathcal{U} = [-\tan \phi_{max}, \tan \phi_{max}]$ , where  $\phi_{max} \in (0, \pi/2)$  is the maximum steering angle. The Dubins car fails the small-time local controllability (STLC) test which creates interesting non-trivial control problems 13. Consider computing the shortest time path arriving at a prescribed goal  $x_f$ . A simple way to define the problem is to set  $\mathcal{X}_f = \mathcal{X}$  and define the cost function as

$$C(u, x) = 1 + \lambda \|x - x_f\|^2,$$

where  $\lambda$  is a positive constant.

1) *Trajectory Generation*: A trajectory satisfying (19) can be represented using a set of connected *primitives* which are either straight lines with velocity  $v$  or arcs with radius  $v/u$  in position space. A primitive is represented by its constant control  $u$  and time duration, denoted by  $\tau > 0$ . Assume that

a trajectory is parametrized using  $m$  such primitives, i.e.

$$z = (u_1, \tau_1, \dots, u_m, \tau_m) \in \mathbb{R}^{2m}.$$

The ending time of the execution of the  $i$ -th primitive, denoted  $t_i$ , is computed as  $t_i = \sum_{k=1}^i \tau_k$  for  $i = 1, \dots, m$  with  $t_0 = 0$ . During time interval  $t \in [t_i, t_{i+1}]$  the parametrization  $\varphi_z = (u, r_x, r_y, \theta)$  takes the form

$$u(t) = u_{i+1}$$

$$r_x(t) = \begin{cases} r_x(t_i) + \frac{v}{u_{i+1}} (\sin(\theta_i + \Delta t_i u_{i+1}) - \sin \theta_i), & \text{if } u_{i+1} \neq 0 \\ r_x(t_i) + v \Delta t_i \cos \theta_i, & \text{otherwise} \end{cases}$$

$$r_y(t) = \begin{cases} r_y(t_i) + \frac{v}{u_{i+1}} (\cos \theta_i - \cos(\theta_i + \Delta t_i u_{i+1})), & \text{if } u_{i+1} \neq 0 \\ r_y(t_i) + v \Delta t_i \cos \theta_i, & \text{otherwise} \end{cases}$$

$$\theta(t) = \theta_i + (\Delta t_i) u_{i+1},$$

where  $\Delta t_i = t - t_i$  and  $\theta_i = \theta(t_i)$ .

2) *Numerical Results*: Consider the same environment as in Fig. 1 and the same algorithm parameters as in §V-A. A trajectory is represented using  $m = 6$  primitives. Fig. 3 shows the algorithm progression and the cost evolution and final solution are shown on Fig. 4. In order to set the nominal sampling parameter  $v_0$  we first compute the shortest kinematic path with maximum velocity ignoring any obstacles and then split it into primitives by properly selecting the coasting times. The set of primitives is then used as the mean component in  $v_0$  while the standard deviation is chosen to achieve coverage of the workspace.

Note that the number of primitives  $m$  can be adaptively chosen. For instance, if the condition  $\alpha_i = \alpha_{i+1}$  persists during iteration then we simply remove  $\alpha_{i+1}$ , set  $m = m - 1$  and merge the times  $\tau_i = \tau_i + \tau_{i-1}$ . A primitive  $\alpha_i$  can also

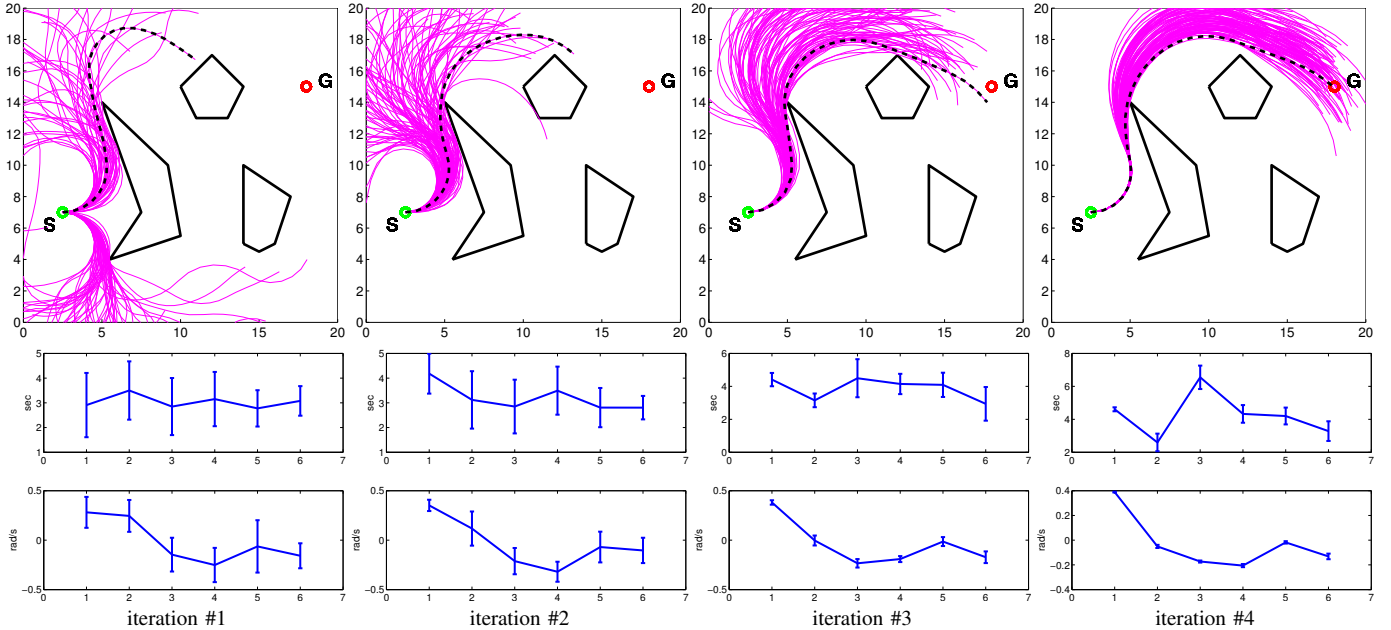


Fig. 3. The first four iterations of the CE algorithm 4.1 applied to the Dubins car. The upper plots show the sampled trajectories  $\varphi(Z_1), \dots, \varphi(Z_N)$  and the current optimal path  $\varphi(z^*)$  (dashed). The lower plots visualize  $p(\cdot, \hat{v}_j)$  by plotting the components of the mean  $\mu$  and standard deviations extracted from the diagonal of  $\Sigma$ . A total of  $m = 6$  primitives were used and the lower plots shows the graphs of the time durations ( $i, \tau_i$ ) and commanded velocities ( $i, \omega_i$ ).




Fig. 4. The resulting cost  $J(z^*)$  of the Dubins car path optimization (see Fig. (3)) is shown on a) while the resulting optimal path  $\varphi(z^*)$  on b).

be removed if the condition  $\tau_i \approx 0$  persists. In that respect the algorithm has the capacity to determine an *optimal number* of different primitives.

## VI. LINKS TO SAMPLING-BASED MOTION PLANNING

It is interesting to note that the CE approach effectively constructs a sequence of parametric models of sets of progressively “better” trajectories, i.e. with decreasing costs. Since a trajectory is a set of states this model can also be regarded as a density over the state space  $\mathcal{X}$ . A state  $x \in \mathcal{X}$  has a high importance density, say  $p_{\mathcal{X}}(x)$ , if it belongs to a trajectory  $\varphi_z \in \mathcal{P}$  with high density  $p(z)$ . Such a sampling density is defined as

$$p_{\mathcal{X}}(X; v) = \eta \cdot \max_{Z \in \mathcal{Z}_{\text{con}}} \{p(Z; v) \mid X \in \kappa(\varphi_Z)\}, \quad (20)$$

where  $X$  is a random variable over  $\mathcal{X}$  and  $\eta > 0$  is a normalizing constant. For instance, the lower plots of Fig. 1 show the evolution of  $p_{\mathcal{X}}$  restricted to the position space.

The issue of how to properly *sample and connect nodes* lies at the core of almost all classical motion-planning methods. While various successful methods for improved sampling have been proposed (e.g. [10, 11, 7, 8, 17, 14]) the *optimal nodes* selection remains a hard and unsolved problem. In that respect, the density  $p_{\mathcal{X}}$  is in fact an approximation to the optimal density for sampling nodes. A fruitful future direction is therefore to employ (20) for adaptive sampling in the context of probabilistic roadmap methods.

## VII. AERIAL VEHICLE APPLICATION

Consider a small autonomous helicopter depicted in Fig. 5 operating in a 3-D terrain. The vehicle is modeled as a single underactuated rigid body with position  $\mathbf{r} \in \mathbb{R}^3$  and orientation matrix  $R \in SO(3)$ . Its *body-fixed* angular and linear velocities are denoted by  $\omega \in \mathbb{R}^3$  and  $\mathbf{v} \in \mathbb{R}^3$ , respectively. The vehicle has mass  $m$  and principal moments of rotational inertia  $J_1, J_2, J_3$  forming the inertia tensor  $\mathbb{J} = \text{diag}(J_1, J_2, J_3)$ .

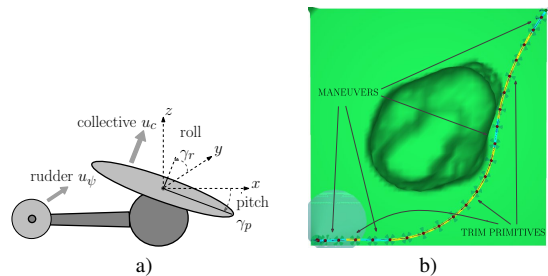


Fig. 5. Simplified helicopter model used in our tests (a). An example of a computed sequence of four maneuvers and three trim primitives is shown (b), connecting two zero-velocity states in the corners of the workspace and avoiding an obstacle in the center.

The vehicle is controlled through a *collective*  $u_c$  (lift produced by the main rotor) and a *yaw*  $u_\psi$  (force produced by the rear rotor), while the direction of the lift is controlled by tilting the main blades forward or backward through a *pitch*  $\gamma_p$  and sideways through a *roll*  $\gamma_r$ . The four control inputs then consist of the two forces  $u = (u_c, u_\psi)$  and the two shape variables  $\gamma = (\gamma_p, \gamma_r)$ . The state space of the vehicle is  $\mathcal{X} = SE(3) \times \mathbb{R}^6 \times \mathbb{R}^2$  with  $x = ((R, \mathbf{r}), (\omega, \mathbf{v}), \gamma)$ .

The equations of motion are

$$\begin{bmatrix} \dot{R} \\ \dot{\mathbf{r}} \end{bmatrix} = \begin{bmatrix} R \hat{\omega} \\ R \mathbf{v} \end{bmatrix}, \quad (21)$$

$$\begin{bmatrix} \mathbb{J} \dot{\omega} \\ m \dot{\mathbf{v}} \end{bmatrix} = \begin{bmatrix} \mathbb{J} \omega \times \omega \\ m \mathbf{v} \times \omega + R^T(0, 0, -9.81m) \end{bmatrix} + F(\gamma)u, \quad (22)$$

where the map  $\hat{\cdot} : \mathbb{R}^3 \rightarrow \mathfrak{so}(3)$  and control matrix  $F$  are defined by

$$\hat{\omega} = \begin{bmatrix} 0 & -\omega^3 & \omega^2 \\ \omega^3 & 0 & -\omega^1 \\ -\omega^2 & \omega^1 & 0 \end{bmatrix}, \quad F(\gamma) = \begin{bmatrix} d_t \sin \gamma_r & 0 \\ d_t \sin \gamma_p \cos \gamma_r & 0 \\ 0 & d_r \\ \sin \gamma_p \cos \gamma_r & 0 \\ -\sin \gamma_r & -1 \\ \cos \gamma_p \cos \gamma_r & 0 \end{bmatrix}.$$

Since gravity is the only configuration-dependent term in the dynamics and is invariant to translations and rotations around the  $z$ -axis, i.e. to transformations by the symmetry group  $G = SE(2) \times \mathbb{R}$ . This property can be exploited to generate trajectories composed of two types of simple motions—constant velocity (trim) curves and transitions between them (maneuvers), in terms of a maneuver automaton [5].

3) *Trim Primitive*: Trim motions correspond to helices, circles, or straight lines in position space. Such curves are traced using velocity with components  $(\omega'_z, v'_x, v'_y, v'_z)$  which are related to the body-fixed velocities  $v$  and  $\omega$  through

$$\omega = R_{(\phi, \theta, 0)}^T \omega', \quad \mathbf{v} = R_{(\phi, \theta, 0)}^T \mathbf{v}', \quad (23)$$

where  $R_{(\phi, \theta, \psi)}$  denotes rigid body rotation parametrized by roll  $\phi$ , pitch  $\theta$ , and yaw  $\psi$  and where  $\mathbf{v}' = (v'_x, v'_y, v'_z)$  and  $\omega' = (0, 0, \omega'_z)$ .

An invariant trajectory has a constant body-fixed velocity. Therefore, in order to determine the invariance conditions for desired  $\omega'$  and  $\mathbf{v}'$  one sets  $\dot{\mathbf{v}} = 0$  and  $\dot{\omega} = 0$  in (22) and substitutes (23) into (22). In order to simplify the derivation of the trim conditions assume two identical inertia components  $J_2 = J_3$  and no vertical velocity, i.e.  $v'_z = 0$ . A closed-form solution can be computed as

$$\begin{aligned} \theta &= 0, & u_y &= 0, & \gamma_p &= 0, & \gamma_r &= 0, \\ \phi &= \arctan(-\omega'_z v'_x / g), \\ u_c &= m(g \cos \phi - \omega'_z v'_x \sin \phi). \end{aligned} \quad (24)$$

Finally, the motion along a trim primitive with initial state  $x_0 = (R_{(\phi, \theta, \psi)}, \mathbf{r}, \omega, \mathbf{v}, \gamma)$  at time  $t_0$  evolves according to

$$x(t_0 + t) = (R_{(\phi, \theta, \psi + \Delta\psi)}, R_{(0, 0, \Delta\psi)} \mathbf{r} + \Delta \mathbf{r}, \omega, \mathbf{v}, \gamma),$$

where  $\Delta\psi = t\omega_z$  and  $\Delta \mathbf{r} = (\Delta r_x, \Delta r_x, t v_z)$  with

$$\begin{aligned} \Delta r_x &= \begin{cases} (v'_x \sin t\omega'_z - v'_y(1 - \cos t\omega'_z)) / \omega'_z, & \text{if } \omega'_z \neq 0 \\ t v'_x, & \text{otherwise,} \end{cases} \\ \Delta r_y &= \begin{cases} (v'_x(1 - \cos t\omega'_z) + v'_y \sin t\omega'_z) / \omega'_z, & \text{if } \omega'_z \neq 0 \\ t v'_y, & \text{otherwise.} \end{cases} \end{aligned}$$

While in our example we consider only planar trims with  $v'_z = 0$  the connecting maneuvers evolve in 3-D position space. Trims are stored by discretizing the space of allowable velocities  $(\omega', v'_x, v'_y)$ .

4) *Maneuvers*.: Maneuvers are computed to connect two trimmed states. Let the map  $\varpi : X \rightarrow X \setminus G$  subtract out the invariant coordinates from a given state, i.e.  $\varpi((R_{(\phi, \theta, \psi)}, \mathbf{r}, \omega, \mathbf{v}, \gamma)) = (R_{(\phi, \theta, 0)}, \mathbf{0}, \omega, \mathbf{v}, \gamma)$ . A maneuver between two given primitives with states  $x_\alpha$  and  $x_\beta$  is computed through nonlinear optimization of a trajectory whose start and end satisfy the given trim conditions. This is defined through the following procedure:

*Compute*:  $T; x : [0, T] \rightarrow \mathcal{X}; u : [0, T] \rightarrow \mathcal{U}$

*minimizing*:  $J(x, u, T) = \int_0^T (1 + \lambda \|u(t)\|^2) dt,$

*subject to*:  $\varpi(x(0)) = x_\alpha, \varpi(x(T)) = x_\beta,$

dynamics eq. (21),(22) for all  $t \in [0, T]$ .

for some  $\lambda > 0$  which adds regularity to the problem. The optimizations were performed offline and all trims and maneuvers are organized and saved in a library which is loaded at run-time providing instant look-up during planning.

5) *Motion Planning*.: The cross-entropy algorithm 4.1 is implemented using  $m$  primitives forward from the start state  $x_0$  and  $m$  primitives backward in reverse from the goal state  $x_f$ . A total of 5 primitives are used for closing the gap between the evolved states (since  $\dim(SE(2)) = 3$  and two maneuvers are required between the three trims). A trajectory consisting of the minimum number of five primitives is computed instantly in closed form through inverse kinematics. Fig. 5 shows an example of such a sequence of primitives. The terrain is represented using a digital elevation map loaded from a file. Collision checking and avoidance is performed using the Proximity Query Package (PQP) [6] that compute closest distance between arbitrary polyhedra and was used to implement the function **prox** defined in (18). Trims are sampled from a continuous Gaussian mixture and the samples are projected onto the discrete grid of predefined trims.

6) *Numerical Results*.: The algorithm is tested in a digital terrain corresponding to the training facility in Ft. Benning. The helicopter is not permitted to fly above buildings. Fig. 6 shows several iterations of the algorithm and the final optimized path between two given states. The first iteration took five seconds of computation time while the remaining iterations took less than a second. CPU time is dominated by attempting to sample trajectories which collide with obstacles. As the informedness of  $p(Z; v)$  increases fewer infeasible trajectories are tried during sampling and each iteration completes in a few milliseconds.






Fig. 6. Several iterations of the CE algorithm applied to the helicopter scenario. As the optimization proceeds the set of samples concentrates in the homotopy classes with minimum trajectory cost.

## VIII. CONCLUSION

The paper addresses the motion planning problem through a randomized optimization in the continuous space of feasible trajectories. Adaptive importance sampling using the cross-entropy method is employed to guide the search towards optimal solution regions. Empirical results show that the proposed motion planning algorithm can handle non-trivial problems as long as feasible trajectories can be sampled successfully. Future work will focus on deriving formal convergence rates. A promising direction in that respect is to establish a relationship between existing motion planning probabilistic completeness results [2] and existing bounds on the complexity of stochastic programming [19]. As a result it would be interesting to consider augmenting existing graph-based planners with the proposed cross-entropy adaptive sampling for optimally guiding the exploration.

## REFERENCES

- [1] Zdravko Botev and Dirk P. Kroese. Global likelihood optimization via the cross-entropy method with an application to mixture models. In *Winter Simulation Conference*, pages 529–535, 2004.
- [2] Howie Choset, Kevin M. Lynch, Seth Hutchinson, George A Kantor, Wolfram Burgard, Lydia E. Kavraki, and Sebastian Thrun. *Principles of Robot Motion: Theory, Algorithms, and Implementations*. MIT Press, June 2005.
- [3] Francis H. Clarke, Yuri S. Ledyaev, Ronald J. Stern, and Peter R. Wolenski. *Nonsmooth Analysis and Control Theory*. Springer, 1998.
- [4] M. A. F. Figueiredo and A. K. Jain. Unsupervised learning of finite mixture models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(3):381–396, 2002. doi: 10.1109/34.990138.
- [5] E. Frazzoli, M. A. Dahleh, and E. Feron. Maneuver-based motion planning for nonlinear systems with symmetries. *IEEE Transactions on Robotics*, 21(6):1077–1091, dec 2005.
- [6] S. Gottschalk, M. C. Lin, and D. Manocha. OBBTree: A hierarchical structure for rapid interference detection. *Eurographics/ACM SIGGRAPH Symposium on Computer Animation*, 30:171–180, 1996.
- [7] D. Hsu, G. Sanchez-Ante, and Zheng Sun. Hybrid prm sampling with a cost-sensitive adaptive strategy. In *Proc. IEEE Int. Conf. Robotics and Automation ICRA 2005*, pages 3874–3880, 2005. doi: 10.1109/ROBOT.2005.1570712.
- [8] David Hsu, Jean-Claude Latombe, and Hanna Kurniawati. On the probabilistic foundations of probabilistic roadmap planning. *The International Journal of Robotics Research*, 25(7):627–643, 2006.
- [9] Dirk Kroese, Sergey Porotsky, and Reuven Rubinstein. The cross-entropy method for continuous multi-extremal optimization. *Methodology and Computing in Applied Probability*, 8:383–407, 2006. ISSN 1387-5841.
- [10] H. Kurniawati and D. Hsu. Workspace importance sampling for probabilistic roadmap planning. In *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS 2004)*, volume 2, pages 1618–1623, 2004.
- [11] A. M. Ladd and L. E. Kavraki. *Fast Tree-Based Exploration of State Space for Robots with Dynamics*, pages 297–312. Springer, STAR 17, 2005.
- [12] J.-C. Latombe. *Robot Motion Planning*. Kluwer Academic Press, 1991.
- [13] S. M. LaValle. *Planning Algorithms*. Cambridge University Press, Cambridge, U.K., 2006.
- [14] Yanbo Li and K. E. Bekris. Balancing state-space coverage in planning with dynamics. In *Proc. IEEE Int Robotics and Automation (ICRA) Conf*, pages 3246–3253, 2010. doi: 10.1109/ROBOT.2010.5509481.
- [15] Geoffrey McLachlan and David Peel. *Finite Mixture Models*. John Wiley & Sons, Inc., 2002.
- [16] Warren B. Powell. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. Wiley Series in Probability and Statistics, 2007.
- [17] M. Rickert, O. Brock, and A. Knoll. Balancing exploration and exploitation in motion planning. In *Proc. IEEE Int. Conf. Robotics and Automation ICRA 2008*, pages 2812–2817, 2008. doi: 10.1109/ROBOT.2008.4543636.
- [18] Reuven Y. Rubenstein and Dirk P. Kroese. *Simulation and the Monte Carlo Method*. Wiley, 2008.
- [19] Reuven Y. Rubinstein and Dirk P. Kroese. *The cross-entropy method: a unified approach to combinatorial optimization*. Springer, 2004.