Distributionally Robust Sampling-Based Motion Planning Under Uncertainty

Tyler Summers

Abstract— **We propose a distributionally robust incremental sampling-based method for kinodynamic motion planning under uncertainty, which we call distributionally robust RRT (DR-RRT). In contrast to many approaches that assume Gaussian distributions for uncertain parameters, here we consider moment-based ambiguity sets of distributions with given mean and covariance. Chance constraints for obstacle avoidance and internal state bounds are then enforced under the worst-case distribution in the ambiguity set, which gives a coherent assessment of constraint violation risks. The method generates risk-bounded trajectories and feedback control laws for robots operating in dynamic, cluttered, and uncertain environments, explicitly incorporating localization error, stochastic process disturbances, unpredictable obstacle motion, and uncertain obstacle location. We show that the algorithm is probabilistically complete under mild assumptions. Numerical experiments illustrate the effectiveness of the algorithm.**

I. Introduction

With robots operating in increasingly dynamic and uncertain environments, more sophisticated motion planning and control algorithms are needed for ensuring safe and efficient autonomous behavior. Many important advances in motion planning have been made in deterministic settings [1]–[7], and recently several extensions have been developed to explicitly incorporate various types of uncertainty into motion planning and control algorithms [8]–[23].

Research on motion planning under uncertainty can be categorized according to how the uncertainty is parameterized and the approach for searching the state space for feasible trajectories and feedback controllers. Two broad approaches to uncertainty parameterization use probabilistic and robust quantifications. In probabilistic approaches, uncertain quantities are modeled as random variables that have given probability distributions, and typically chance constraints are enforced so that nominal constraints hold with prescribed probability [9]–[20]. In robust approaches, only the support of the distribution is assumed to be known, and constraints are enforced for the worst-case disturbances in the support set [21], [22]. Two broad approaches to searching for feasible trajectories are sampling-based [9]–[18] and optimization-based methods [19], [20], [24], and both types have been developed to explicitly incorporate uncertainty. Other approaches handle uncertainty implicitly by designing feedback controllers based on deterministic models [23].

Although Gaussian assumptions and chance constraints are commonly made to ensure computational tractability, such assumptions are rarely justifiable based on actual data from nonlinear robotic systems. This can cause significant miscalculations of risk, and the underlying risk metrics do not necessarily possess desirable properties such as coherence [25]. However, recent results from the emerging area of distributionally robust optimization have shown that a much more sophisticated treatment of stochastic uncertainty is possible without sacrificing computational tractability [26]–[32]. These approaches consider ambiguity sets of distributions based on data or estimated quantities, rather than making strong assumptions about the form of the distribution. However, to the author's best knowledge, this framework has not been utilized within robotic motion planning and control algorithms. More sophisticated and nuanced approaches for quantifying risks can enable robots to operate successfully in increasingly complex, real-world environments.

We propose a distributionally robust incremental sampling-based method for kinodynamic motion planning under uncertainty, which we call distributionally robust RRT (DR-RRT). In contrast to many approaches that assume Gaussian distributions for uncertain parameters, here we consider moment-based ambiguity sets of distributions with given mean and covariance. Chance constraints for obstacle avoidance and internal state bounds are then enforced under the worst-case distribution in the ambiguity set, which gives a coherent assessment of constraint violation risks. The method generates risk-bounded trajectories and feedback control laws for robots operating in dynamic, cluttered, and uncertain environments, explicitly incorporating localization error, stochastic process disturbances, unpredictable obstacle motion, and uncertain obstacle location. This approach may be particularly effective for nonlinear systems, where moments are only approximations of non-Gaussian distributions. We show that the algorithm is probabilistically complete under mild assumptions. Numerical experiments illustrate the effectiveness of the algorithm.

The rest of the paper is organized as follows. Section II formulates the model and distributionally robust motion planning problem. Section III presents our incremental sampling-based algorithm for generating distributionally robust trajectories and feedback control policies and provides analysis of probabilistic completeness. Section IV presents illustrative numerical experiments. Section V concludes and summarizes ongoing and future research.

II. Problem Formulation

We begin by formulating a distributionally robust motion planning under uncertainty problem. The formulation follows an analogous approach for chance-constrained motions planning with Gaussian distributions [11].

The author is with the University of Texas at Dallas, E-mail: tyler.summers@utdallas.edu. This work was sponsored by the Army Research Office and accomplished under Grant Number: W911NF-17-1-0058.

A. Model

The robot dynamics are modeled by the stochastic discrete-time linear time invariant system

$$x_{t+1} = Ax_t + Bu_t + w_t, \tag{1}$$

where $x_t \in \mathbf{R}^n$ is the system state at time t, $u_t \in \mathbf{R}^m$ is the input at time t, A is the dynamics matrix, B is the input matrix, and $w_t \in \mathbf{R}^n$ is a zero-mean random vector independent and identically distributed across time. The distribution P_w of w_t is unknown (and not necessarily Gaussian) and will be assumed to belong to an *ambiguity set* $\mathcal{P}^w$ of distributions, which will be discussed in detail shortly. The initial condition x_0 may be known exactly or subject to a similar uncertainty model, with the distribution of x_0 belonging to an ambiguity set, $P_{x_0} \in \mathcal{P}^x$.

The system is nominally subject to constraints on the state and input of the form

$$x_t \in \mathcal{X}_t = \mathcal{X} \backslash \mathcal{X}_{1t} \backslash \cdots \backslash \mathcal{X}_{n_o t}, \qquad u_t \in \mathcal{U}, \tag{2}$$

where $\mathcal{X}, \mathcal{X}_{1t}, \cdots, \mathcal{X}_{n_o t} \subset \mathbf{R}^n$, $\mathcal{U} \subset \mathbf{R}^m$ are assumed to be convex polytopes, and the operator $\backslash$ denotes set subtraction. The sets $\mathcal{X}$ and $\mathcal{U}$ represent state and input constraints on the robot, and $\mathcal{X}_{it}$ represent a set of n_o obstacles in the environment to be avoided. We assume the shape and orientation of obstacles to be known, but that their placement and motion is subject to uncertainty:

$$\mathcal{X}_{it} = \mathcal{X}_i^0 + \hat{c}_{it} + c_{it}, \quad i = 1, ..., n_o, \tag{3}$$

where $+$ here denotes set translation. In particular, $\mathcal{X}_i^0 \subset \mathbf{R}^n$ represents the known shape of obstacle i, $\hat{c}_{it} \in \mathbf{R}^n$ represents a known nominal translation, and $c_{it} \in \mathbf{R}^n$ is a random vector that represents an unknown location or unpredictable obstacle motion, with unknown distribution $P_{it}^c \in \mathcal{P}_{it}^c$.

Since constraint and obstacle sets are assumed polytopic, they can be represented by finite sets of linear inequalities

$$\begin{aligned} &\mathcal{U} = \{u_t \mid A_u u_t \le b_u\} \\ &\mathcal{X} = \{x_t \mid A_0 x_t \le b_0\}, \quad \mathcal{X}_{it} \quad = \{x_t \mid A_i x_t \le b_{it}\}, \end{aligned} \tag{4}$$

where $b_u \in \mathbf{R}^{n_u}$, $b_0 \in \mathbf{R}^{n_E}$, $b_{it} \in \mathbf{R}^{n_i}$, and A_u, A_0, and A_i are matrices of appropriate dimension. The non-convex obstacle avoidance constraints can be expressed as the disjunction

$$\begin{aligned} &\neg(A_i x_t \le b_{it}), \quad \forall i = 1, ..., n_o \\ \Leftrightarrow \quad &\vee_{j=1}^{n_i} \, (a_{ij}^T x_t \ge a_{ij}^T c_{ijt}), \end{aligned} \tag{5}$$

where $\vee$ denotes disjunction and $c_{ijt} = \hat{c}_{ijt} + c_{it}$ is a point nominally on the jth constraint of the ith obstacle whose covariance is the same as that of c_{it}.

The model incorporates three distinct types of uncertainty: internal predictive and localization uncertainty from the process noise w_t and the initial state x_0, and external environmental uncertainty from the unknown obstacle locations.

B. A Distributionally Robust Motion Planning Problem

The goal of the motion planning problem is to find a feedback control policy $\pi = [\pi_0, ..., \pi_{T-1}]$ with $u_t = \pi_t(x_t)$ that yields a feasible and minimum cost trajectory and from the initial state x_0 to a goal set $\mathcal{X}_{\text{goal}} \subset \mathcal{R}^n$. Accordingly, we seek to (approximately) solve the distributionally robust constrained stochastic optimal control problem

$$\begin{aligned} \underset{\pi \in \Pi}{\text{minimize}} \quad & \sum_{t=0}^{T-1} \ell_t(\mathbf{E}x_t, \mathcal{X}_{\text{goal}}, u_t) + \ell_T(\mathbf{E}x_T, \mathcal{X}_{\text{goal}}) \\ \text{subject to} \quad & x_{t+1} = Ax_t + Bu_t + w_t \\ & x_0 \sim P_{x_0} \in \mathcal{P}^x, \quad w_t \sim P_w \in \mathcal{P}^w \\ & u_t \in \mathcal{U} \\ & \inf_{P_{x_t} \in \mathcal{P}} P_{x_t}(x_t \in \mathcal{X}_t) \ge 1 - \alpha \\ & \mathcal{X}_t = \mathcal{X} \backslash \mathcal{X}_{1t} \backslash \cdots \backslash \mathcal{X}_{n_o t} \\ & \mathcal{X}_{it} = \mathcal{X}_i^0 + \hat{c}_{it} + c_{it}, \quad i = 1, ..., n_o \\ & c_{it} \sim P_{it}^c \in \mathcal{P}_{it}^c, \quad t = 0, ..., T-1, \end{aligned} \tag{6}$$

where $\mathcal{P}$ is an ambiguity set of marginal state distributions and $\alpha \in (0, 0.5]$ is a user-prescribed risk parameter. The stage cost functions $\ell_t(\cdot)$ quantify the robot's distance to the goal set and actuator effort, and are assumed to be expressed in terms of the state mean $\mathbf{E}x_t$, so that all stochasticity appears in the constraints. A key distinction of the present work is that the state constraints are expressed as *distributionally robust chance constraints*. Standard chance constraints require the nominal state constraints $x_t \in \mathcal{X}_t$ to be satisfied with probability $1 - \alpha$. In contrast, here we enforce distributionally robust chance constraints, which require the nominal state constraints $x_t \in \mathcal{X}_t$ to be satisfied with the same probability, but under the worst case probability distribution in the ambiguity set[1]. In the next section, we present a probabilistically complete sampling-based motion planning algorithm for generating coherent distributionally robust risk constrained trajectories and feedback control laws in the presence of both internal and environmental uncertainties.

C. Ambiguity Sets and Distributionally Robust Optimization

Most motion planning algorithm do not explicitly incorporate uncertainty, and even the ones that do almost always assume a functional form (often Gaussian) for probability distributions of uncertain quantities. Assumptions of Gaussianity are rarely justifiable based on actual data from nonlinear robotic systems, but are made in the name of computational tractability. However, recent results from the emerging area of distributionally robust optimization have shown that a much more sophisticated treatment of stochastic uncertainty is possible without sacrificing computational tractability [26]–[32].

Distributionally robust optimization approaches can be categorized based on the form of the ambiguity set. There

[1]It is straightforward to extend our approach to handle pathwise state constraints of the form $\inf_{P_{x_t} \in \mathcal{P}} P_{x_t}(\wedge_t (x_t \in \mathcal{X}_t)) \ge 1 - \alpha_p$, but we restrict attention to stage wise constraints to simplify the exposition.

are several different parameterizations, including those based on moments, support, directional derivatives [26], [27], f-divergences [28], [31], and Wasserstein balls [30]. For example, a moment-based ambiguity set includes all distributions with a fixed moments up to some order (e.g., fixed first and second moments), and Wasserstein-based ambiguity sets include a ball of distributions within a given Wasserstein distance from some base distribution (such as an empirical distribution on a training dataset). We will focus here on moment-based ambiguity sets, though other parameterizations are interesting and relevant for robotic motion planning and will be pursued in future work.

Suppose the feedback control policy is a fixed affine function $u_t = K_t x_t + k_t$ and that the second moments of the initial state and disturbance are denoted by $\Sigma_{x_0} = \mathbf{E}(x_0 - \hat{x}_0)(x_0 - \hat{x}_0)^T$ and $W = \mathbf{E} w_t w_t^T$ and known (or estimated based on historical data). Then the state mean $\hat{x}_t = \mathbf{E} x_t$ and covariance matrix $\Sigma_{x_t} = \mathbf{E}(x_t - \hat{x}_t)(x_t - \hat{x}_t)^T$ evolve according to

$$\begin{aligned} \hat{x}_{t+1} &= (A + BK_t)\hat{x}_t + Bk_t \\ \Sigma_{x_{t+1}} &= (A + BK_t)\Sigma_{x_t}(A + BK_t)^T + W. \end{aligned} \tag{7}$$

Since the primitive distributions P_{x_0} and P_w are not assumed to be Gaussian, then neither are the marginal state distributions P_{x_t}. In contrast to previous work with Gaussianity assumptions where the moment dynamics (7) completely characterize the marginals, e.g., [11], here we consider the moment-based ambiguity set

$$\mathcal{P} = \{P_{x_t} \mid \mathbf{E} x_t = \hat{x}_t,\ \mathbf{E}(x_t - \hat{x}_t)(x_t - \hat{x}_t)^T = \Sigma_{x_t}\}. \tag{8}$$

Furthermore, suppose that $\mathbf{E} c_{ijt} = \hat{c}_{ijt}$ and $\mathbf{E}(c_{ijt} - \hat{c}_{ijt})(c_{ijt} - \hat{c}_{ijt})^T = \Sigma^c_{jt}$, which defines a corresponding moment-based ambiguity set for obstacle motion.

Under the moment-based ambiguity set (8), a constraint on the worst-case probability of violating the jth constraint of obstacle i

$$\sup_{P_{x_t} \in \mathcal{P}} P_{x_t}(a_{ij}^T x_t \geq a_{ij}^T c_{ijt}) \leq \alpha_i \tag{9}$$

is equivalent (see, e.g., [26]) to the linear constraint on the state mean $\hat{x}_t$

$$a_{ij}^T \hat{x}_t \geq a_{ij}^T \hat{c}_{ijt} + \sqrt{\frac{1-\alpha_i}{\alpha_i}} \|(\Sigma_t + \Sigma^c_{jt}) a_{ij}\|_2. \tag{10}$$

This represents a deterministic tightening of the nominal constraint constraint to enforce the corresponding distributionally robust constraint. This tightening has the general same form as the one involving Gaussian distributions, but the scaling constant $\sqrt{\frac{1-\alpha_{ij}}{\alpha_{ij}}}$ is larger than the Gaussian one, leading to a stronger tightening that reflects the weaker assumptions about the uncertainty distributions.

Since the robot collides with the obstacle if any one of these constraints is violated, then $\sup_{P_{x_t} \in \mathcal{P}} P_{x_t}(x_t \in \mathcal{X}_{it}) \leq \alpha_i$ is equivalent to

$$\vee_{j=1}^{n_j} \left(a_{ij}^T \hat{x}_t \geq a_{ij}^T \hat{c}_{ijt} + \sqrt{\frac{1-\alpha_i}{\alpha_i}} \|(\Sigma_t + \Sigma^c_{jt}) a_{ij}\|_2 \right). \tag{11}$$

Similarly, a constraint on the worst-case probability of violating jth state constraint defining $\mathcal{X}$, viz. $\sup_{P_{x_t} \in \mathcal{P}} P_{x_t}(a_{0j}^T x_t \geq a_{0j}^T c_{0j}) \leq \alpha_{0j}$ is equivalent to

$$a_{0j}^T \hat{x}_t \leq a_{0j}^T c_{0j} - \sqrt{\frac{1-\alpha_{0j}}{\alpha_{0j}}} \|\Sigma_t a_{0j}\|_2, \tag{12}$$

where c_{0j} is a point on the jth state constraint.

Putting all of this together gives the following result.

Theorem 1: Consider a trajectory $(\hat{x}_0, \Sigma_{x_0}), ..., (\hat{x}_T, \Sigma_{x_T})$ of the first and second moments of the state distribution given by (7). Suppose the state means satisfy the constraints (11) and (12) for each time step, each obstacle, and each environmental constraint, and that the constraint satisfaction parameters are chosen such that $\sum_{i=1}^{n_E} \alpha_{0i} + \sum_{i=1}^{n_o} \alpha_j = \alpha$. Then the trajectory satisfies the distributionally robust state constraints

$$\inf_{P_{x_t} \in \mathcal{P}} P_{x_t}(x_t \in \mathcal{X}_t) \geq 1 - \alpha \tag{13}$$

Proof: The distributionally robust state constraint (13) can equivalently be written in terms of a probabilistic upper bound on constraint violation: $\sup_{P_{x_t} \in \mathcal{P}} P_{x_t}(x_t \notin \mathcal{X}_t) \leq \alpha$. Since the constraint sets are polytopic, Boole's inequality and (11) and (12) can be used to obtain the bound

$$\begin{aligned} &\sup_{P_{x_t} \in \mathcal{P}} P_{x_t}(x_t \notin \mathcal{X}_t) \\ &\leq \sup_{P_{x_t} \in \mathcal{P}} P_{x_t}(x_t \notin \mathcal{X}) + \sum_{i=1}^{n_o} \sup_{P_{x_t} \in \mathcal{P}} P_{x_t}(x_t \in \mathcal{X}_{it}) \\ &\leq \sum_{j=1}^{n_E} \sup_{P_{x_t} \in \mathcal{P}} P_{x_t}(a_{0j}^T x_t \geq a_{0j}^T c_{0j}) + \sum_{i=1}^{n_o} \sup_{P_{x_t} \in \mathcal{P}} P_{x_t}(x_t \in \mathcal{X}_{it}) \\ &\leq \sum_{i=1}^{n_E} \alpha_{0i} + \sum_{i=1}^{n_o} \alpha_j = \alpha. \end{aligned}$$

■

In the next section, we will present a sampling-based motion planning algorithm that incrementally generates state trajectories and feedback control laws that satisfy the distributionally robust constraints. It is possible to reduce conservatism associated with applying Boole's inequality by exploiting the trajectory-wise incremental constraint checking of sampling-based motion planning algorithms to dynamically allocate risk amongst the various constraints as in [11]; this will be pursued in future work.

III. Distributionally Robust RRT (DR-RRT)

This section describes Distributionally Robust RRT, an incremental sampling-based method for kinodynamic motion planning under uncertainty. The algorithm grow trees of state *distributions*, rather than trees of states, and incorporates distributionally robust probabilistic constraints to certify feasibility of nominal state mean trajectories.

A. Algorithm

Tree Expansion. Pseudocode for distributionally robust RRT tree expansion is shown in Algorithm 1. In the first step, a random sample is taken from the feasible state set.

Algorithm 1 DR-RRT: Tree Expansion

Inputs: current tree $\mathcal{T}$ and timestep t
x_s = sample($\mathcal{X}_t$)
$(\hat{x}_{nearest}, \Sigma_{nearest})$ = nearestMnodes(x_s, $\mathcal{T}$)
for $i \in \{1, ..., M\}$ **do**
  $(\hat{x}_{traj}, \Sigma_{traj})$ = steer($\hat{x}_{nearest}, \Sigma_{nearest}, x_s$)
  while DRCollisionCheck($\hat{x}_{traj}, \Sigma_{traj}$) **do**
    $\mathcal{T}$.AddNodes($\hat{x}_{traj}, \Sigma_{traj}$)
    $\mathcal{T}$.AddEdges($x_{nearest}, \hat{x}_{traj}$)

Then a set of $M \geq 1$ tree nodes nearest to the sample are selected according to some appropriate distance metric and returned in ascending order. It is typically advantageous to use a dynamic control-based distance metric, such as an optimal cost-to-go function, instead of a purely geometric distance, in order to efficiently explore the reachable set of the dynamics and increase likelihood of generating collision-free trajectories [6]. Attempts are then made to steer the robot from the nearest tree nodes to the random sample. For this step, we use an unconstrained optimal feedback control policy obtained via a finite horizon linear quadratic dynamic programming algorithm. This policy is then used to propagate the state mean and covariance matrix, and the entire trajectory is returned by the steer function. However, many different control design methods may be used; in future work we will explore more sophisticated steering methods that also explicitly incorporate some of the nearby obstacle constraints. Each state distribution in the trajectory is then checked for distributionally robust probabilistic constraint satisfaction, and feasible portions of the trajectory are added to the tree. This expansion is then repeated until a node from the goal set is added to the tree. At that point, a distributionally robust feasible trajectory is obtained from the tree root to the goal set.

Dynamic execution. For dynamic and stochastic environments, the tree expansion can be updated dynamically as the robot moves and new information about the environment is obtained. The minimum cost feasible path currently available in the tree can be executed, and then new tree branches can be grown from the new state estimate and infeasible branches can be pruned.

B. Analysis

Assumption 1: We assume the following:

(i) There exists a sequence of feasible control inputs $u_0, ..., u_{T-1}$ (possibly generated by a control policy $\pi = [\pi_0, ..., \pi_{T-1}]$ with $u_t = \pi_t(x_t)$) such that all constraints in (6) are satisfied by the state distribution trajectory $(\hat{x}_0, \Sigma_{x_0}), ..., (\hat{x}_T, \Sigma_{x_T})$ obtained via (7).

(ii) Each node added during tree expansion is separated from its parent node by a distance of at least $\epsilon > 0$.

(iii) When attempting to connect a random sample from the feasible set to the tree by the steering function, inputs are randomly selected from a finite set, which contains all inputs in the feasible sequence from (i).

In practice, inputs are not chosen randomly as in (iii), and the exploration will be most effective when advanced control methods are used to steer the robot toward random samples. We have the following result, which follows along the lines of the analogous result for chance constrained RRT in [11].

Theorem 2: Under Assumption 1, the proposed distributionally robust RRT algorithm (DR-RRT) in Algorithm 1 is probabilistically complete, i.e., it returns a feasible solution with probability approaching 1 as the number of samples tends to infinity.

Proof: By induction. Suppose that the tree contains the feasible state distribution (x_t, Σ_{x_t}) generated by applying the feasible control input sequence $u_0, ..., u_{t-1}$ from Assumption 1(i). Since all tree nodes have a non-empty Voronoi region due to Assumption 1(ii), there is a strictly positive probability that the node corresponding to (x_t, Σ_{x_t}) will be selected for expansion. By Assumption 1(iii) there is a strictly positive probability that the steering input u_t is selected and results in the feasible state distribution $(x_{t+1}, \Sigma_{x_{t+1}})$ being added to the tree and connected to (x_t, Σ_{x_t}). Thus, as the number of samples tends to infinity, the probability of generating the feasible sequence approaches 1. Basing the tree at the initial node (x_0, Σ_{x_0}) completes the proof by induction. ∎

C. Variations

Here we list several variations of the basic DR-RRT algorithm that are being pursued in future research, some of which are straightforward to implement.

Asymptotic Optimality. It is known that the standard RRT algorithm is not asymptotically optimal, i.e., the feasible solution will not in general minimize the cost function as the number of samples tends to infinity. The recently proposed RRT* algorithm achieves asymptotic optimality by rewiring the tree for lower cost paths around a set of nodes near the sample [7]. A similar approach can be taken here to obtain a distributionally robust RRT* algorithm (DR-RRT*) that is both probabilistically complete and asymptotically optimal.

Nonlinear Dynamics. The DR-RRT algorithm can be applied to systems with nonlinear dynamics by linearizing the system and propagating first and second order moments. In this case, the distributionally robust framework is particularly compelling, since the first and second moments are only approximations of the true state distribution, even when the primitive distributions are Gaussian.

Higher Order Moment Propagation. Higher order moments beyond first and second order can be propagated, in order to sharpen risk estimates. This can be especially effective for nonlinear systems. For example, third and fourth order moments can capture skewness and heaviness of tails and lead to far less conservative trajectories than those obtained by only using first and second moments.

Output feedback with recursive filtering. In many robotic systems, the state must be estimated based on noisy sensor data. A state estimator can be used to propagate a joint distribution of the true state and state estimate based on the dynamic model and sensor data. For example, the Kalman filter and its many variations can be used, and the distributionally robust approach may be an effective way to handle risks associated with state estimation, particularly for extended Kalman filters for nonlinear systems.

Alternative ambiguity sets. As mentioned previously, other DR-RRT variations can be obtained using different parameterizations of the ambiguity set. For example, Wasserstein balls centered on empirical data-based distributions offer a natural way to incorporate knowledge about distributions that come from observed finite training datasets. Such an approach may be effective for particle simulations and particle filters of nonlinear systems. Furthermore, a combination of moment- and data-based distribution parameterizations could be used to combine their relative advantages.

IV. Numerical Experiments

We now present numerical experiments to illustrate the proposed distributionally robust RRT algorithm (DR-RRT).

A. Double Integrator in Cluttered Environment

We consider a unit mass robot with discrete-time stochastic double integrator dynamics moving in a bounded and cluttered two-dimensional environment $[0,1]^2$. The system dynamics matrices are

$$A = \begin{bmatrix} 1 & 0 & dt & 0 \\ 0 & 1 & 0 & dt \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad B = \begin{bmatrix} \frac{dt^2}{2} & 0 \\ 0 & \frac{dt^2}{2} \\ dt & 0 \\ 0 & dt \end{bmatrix}, \tag{14}$$

where $dt = 0.1s$ and the states are the two dimensional position and velocity with two dimensional force inputs. The initial position is $[0.5, 0]$, the initial velocity is zero, and the initial state and disturbance covariance matrices are

$$\Sigma_0 = 10^{-3} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \; W = 10^{-3} \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 2 & 1 \\ 0 & 0 & 1 & 2 \end{bmatrix}. \tag{15}$$

The environment contains randomly located and sized rectangular obstacles, which are static and treated by the robot as deterministic, so that all uncertainty in this example comes from the unknown initial state and process disturbance. The robot is treated as a point mass (without loss of generality, since a known geometry can be easily handled by a fixed tightening of the state constraints), and the environmental boundaries are not treated probabilistically.

To steer the robot from a tree node state x_t to a random feasible sample x_s, we solve a discrete time linear quadratic optimal control problem to compute the affine state feedback policy that minimizes the cost function

$$\mathbf{E} \sum_{t=0}^{T-1} (x_t - x_s)^T Q (x_t - x_s) + u_t^T R u_t + (x_T - x_s)^T Q (x_T - x_s) \tag{16}$$

with $T = 10$, $Q = \begin{bmatrix} 40I & 0 \\ 0 & 0.1I \end{bmatrix}$, and $R = 0.2I$. The distributionally robust state constraints are enforced with probabilistic satisfaction parameter $\alpha = 0.1$. The quadratic optimal cost-to-go function is also used as the distance metric to select tree nodes closest to the sample.

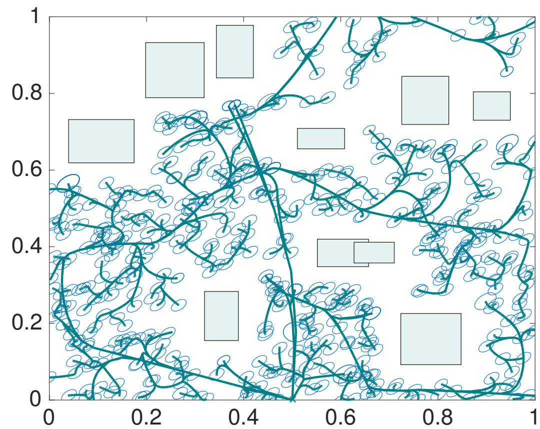

Fig. 1. Illustration of Distributionally Robust RRT (DR-RRT) for the stochastic double integrator robot. The thick lines represent trajectories distributionally robust trajectories of the state mean $\hat{x}_t$, and the ellipses show one standard deviation uncertainty regions of the state position derived from the position block of the covariance matrix Σ_t at sampled positions. The DR-RRT algorithm generates more conservative trajectories around the obstacles, explicitly incorporating the uncertainty in the state due to the initial localization and system dynamics uncertainties.

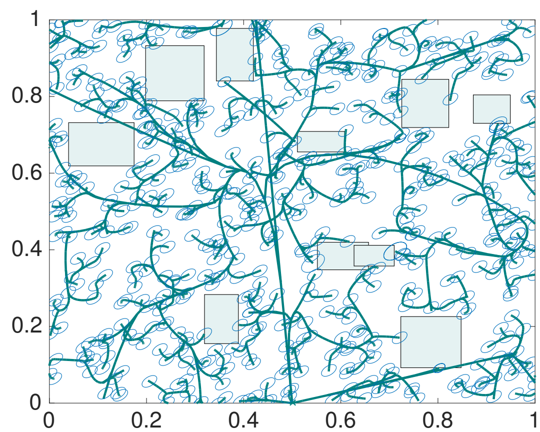

Fig. 2. RRT for the stochastic double integrator robot *without* the distributionally robust constraint tighetening. Since uncertainty is not explicitly incorporated into collision checking, trajectories with high risk of collision are generated.

B. Results and Discussion

Figures 1 and 2 show RRTs with and without the distributionally robust constraint tightening, respectively. It can be easily seen that the DR-RRT generates more conservative trajectories around the obstacles, explicitly incorporating the uncertainty in the state due to the initial localization and system dynamics uncertainties. Due to the use of closed-loop feedback control policies in the steering function, the size of the uncertainty ellipses remains bounded as the tree is constructed. The feasible set remains effectively explored with DR-RRT, although certain regions may be deemed to risky to reach, such as the top left corner that can only be reached through narrow gaps between obstacles and the environment boundary. The algorithm can also be used for higher order dynamics and state distributions propagated from linearized dynamics of nonlinear systems. Computation times are the same as chance-constrained RRT [11], since the problems have the same form.

The sizes of the uncertainty ellipses depend strongly on the steering control method. In the current algorithm, the probabilistic constraint checking is decoupled from the steering control design; the robot is steered from the tree toward random samples, and then the trajectory is check for probabilistic feasibility. Much more effective exploration of the space and generation of collision free trajectories could be obtained by coupling the local trajectory generation and control design with the constraints. For example, affine disturbance feedback with stochastic model predictive control could compute steering feedback policies that explicitly account for a subset of nearby obstacle constraints [33]. This will be pursued in future work.

V. Conclusions

We have proposed a distributionally robust incremental sampling-based algorithm for kinodynamic motion planning under uncertainty, DR-RRT. It effectively generates risk-bounded trajectories for robots operating in uncertain and dynamic environments. The distributionally robust framework for explicitly incorporating stochastic uncertainty handles risk in a more sophisticated way then state-of-the-art methods by considering a (moment-based) ambiguity set of distributions for uncertain parameters, rather than making strong assumptions about the specific form of the distributions. The algorithm is shown to be probabilistically complete under mild assumptions. Numerical experiments illustrate the effectiveness of the algorithm.

In ongoing and future work, we are exploring several variations mentioned above, including an asymptotically optimal version DR-RRT*, distribution propagation for nonlinear systems and with higher order moments, output feedback with recursive filtering to incorporate sensing uncertainty, and alternative ambiguity sets, and more sophisticated steering laws that couple feedback steering controller design with probabilistic collisions obstacle avoidance more tightly.

References

[1] J.-C. Latombe, *Robot motion planning*. Springer, 2012, vol. 124.

[2] S. M. LaValle, *Planning algorithms*. Cambridge Univ. Press, 2006.

[3] J. J. Kuffner and S. M. LaValle, "RRT-connect: An efficient approach to single-query path planning," in *IEEE International Conference on Robotics and Automation*, vol. 2. IEEE, 2000, pp. 995–1001.

[4] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE Transactions on Robotics and Automation*, vol. 12, no. 4, pp. 566–580, 1996.

[5] S. LaValle and J. Kuffner, "Randomized kinodynamic planning," *International Journal of Robotics Research*, vol. 20, no. 5, pp. 378–400, 2001.

[6] E. Frazzoli, M. A. Dahleh, and E. Feron, "Real-time motion planning for agile autonomous vehicles," *AIAA Journal of Guidance, Control, and Dynamics*, vol. 25, no. 1, pp. 116–129, 2002.

[7] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *The International Journal of Robotics Research*, vol. 30, no. 7, pp. 846–894, 2011.

[8] S. M. LaValle and R. Sharma, "On motion planning in changing, partially predictable environments," *The International Journal of Robotics Research*, vol. 16, no. 6, pp. 775–805, 1997.

[9] N. A. Melchior and R. Simmons, "Particle RRT for path planning with uncertainty," in *IEEE International Conference on Robotics and Automation*, 2007, pp. 1617–1624.

[10] S. C. Ong, S. W. Png, D. Hsu, and W. S. Lee, "Planning under uncertainty for robotic tasks with mixed observability," *The International Journal of Robotics Research*, vol. 29, no. 8, pp. 1053–1068, 2010.

[11] B. Luders, M. Kothari, and J. How, "Chance constrained RRT for probabilistic robustness to environmental uncertainty," in *AIAA Guidance, Navigation, and Control Conference*, 2010, p. 8160.

[12] A. Bry and N. Roy, "Rapidly-exploring random belief trees for motion planning under uncertainty," in *IEEE International Conference on Robotics and Automation*, 2011, pp. 723–730.

[13] B. Luders and J. How, "Probabilistic feasibility for nonlinear systems with non-gaussian uncertainty using RRT," in *InfotechAerospace 2011*.

[14] J. Van Den Berg, P. Abbeel, and K. Goldberg, "LQG-MP: Optimized path planning for robots with motion uncertainty and imperfect state information," *The International Journal of Robotics Research*, vol. 30, no. 7, pp. 895–913, 2011.

[15] J. Van Den Berg, S. Patil, and R. Alterovitz, "Motion planning under uncertainty using iterative local optimization in belief space," *International Journal of Robotics Research*, vol. 31, no. 11, pp. 1263–78, 2012.

[16] N. E. Du Toit and J. W. Burdick, "Robot motion planning in dynamic, uncertain environments," *IEEE Transactions on Robotics*, vol. 28, no. 1, pp. 101–115, 2012.

[17] G. S. Aoude, B. D. Luders, J. M. Joseph, N. Roy, and J. P. How, "Probabilistically safe motion planning to avoid dynamic obstacles with uncertain motion patterns," *Autonomous Robots*, vol. 35, no. 1, pp. 51–76, 2013.

[18] A.-A. Agha-Mohammadi, S. Chakravorty, and N. M. Amato, "Firm: Sampling-based feedback motion-planning under motion uncertainty and imperfect measurements," *The International Journal of Robotics Research*, vol. 33, no. 2, pp. 268–304, 2014.

[19] L. Blackmore, M. Ono, A. Bektassov, and B. C. Williams, "A probabilistic particle-control approximation of chance-constrained stochastic predictive control," *IEEE Transactions on Robotics*, vol. 26, no. 3, pp. 502–517, 2010.

[20] L. Blackmore, M. Ono, and B. C. Williams, "Chance-constrained optimal path planning with obstacles," *IEEE Transactions on Robotics*, vol. 27, no. 6, pp. 1080–1094, 2011.

[21] B. D. Luders, S. Karaman, and J. P. How, "Robust sampling-based motion planning with asymptotic optimality guarantees," in *AIAA Guidance, Navigation, and Control (GNC) Conference*, 2013, p. 5097.

[22] J. Ding, E. Li, H. Huang, and C. J. Tomlin, "Reachability-based synthesis of feedback policies for motion planning under bounded disturbances," in *International Conference on Robotics and Automation*. IEEE, 2011, pp. 2160–2165.

[23] R. Tedrake, I. Manchester, M. Tobenkin, and J. Roberts, "LQR-trees: Feedback motion planning via sums-of-squares verification," *International Journal of Robotics Research*, vol. 29, no. 8, pp. 1038–52, 2010.

[24] A. Jasour and C. Lagoa, "Convex relaxations of a probabilistically robust control design problem," in *IEEE Conference on Decision and Control*. IEEE, 2013, pp. 1892–97.

[25] R. T. Rockafellar, "Coherent approaches to risk in optimization under uncertainty," *Tutorials in operations research*, vol. 3, pp. 38–61, 2007.

[26] G. C. Calafiore and L. El Ghaoui, "On distributionally robust chance-constrained linear programs," *Journal of Optimization Theory and Applications*, vol. 130, no. 1, pp. 1–22, 2006.

[27] J. Goh and M. Sim, "Distributionally robust optimization and its tractable approximations," *Operations research*, vol. 58, no. 4-part-1, pp. 902–917, 2010.

[28] Z. Hu and L. Hong, "Kullback-leibler divergence constrained distributionally robust optimization," *Available: Optimization Online*, 2013.

[29] W. Wiesemann, D. Kuhn, and M. Sim, "Distributionally robust convex optimization," *Operations Research*, vol. 62, no. 6, pp. 1358–76, 2014.

[30] P. M. Esfahani and D. Kuhn, "Data-driven distributionally robust optimization using the wasserstein metric: Performance guarantees and tractable reformulations," *Math. Programming*, pp. 1–52, 2015.

[31] H. Namkoong and J. C. Duchi, "Stochastic gradient methods for distributionally robust optimization with f-divergences," in *Advances in Neural Information Processing Systems*, 2016, pp. 2208–2216.

[32] B. P. Van Parys, P. J. Goulart, and M. Morari, "Distributionally robust expectation inequalities for structured distributions," *Mathematical Programming*, pp. 1–30, 2017.

[33] P. J. Goulart, E. C. Kerrigan, and J. M. Maciejowski, "Optimization over state feedback policies for robust control with constraints," *Automatica*, vol. 42, no. 4, pp. 523–533, 2006.