

Integrating Predictive Motion Uncertainties with Distributionally Robust Risk-Aware Control for Safe Robot Navigation in Crowds

Kanghyun Ryu¹ and Negar Mehr¹

Abstract—Ensuring safe navigation in human-populated environments is crucial for autonomous mobile robots. Although recent advances in machine learning offer promising methods to predict human trajectories in crowded areas, it remains unclear how one can safely incorporate these learned models into a control loop due to the uncertain nature of human motion, which can make predictions of these models imprecise. In this work, we address this challenge and introduce a distributionally robust chance-constrained model predictive control (DRCC-MPC) which: (i) adopts a probability of collision as a pre-specified, interpretable risk metric, and (ii) offers robustness against discrepancies between actual human trajectories and their predictions. We consider the risk of collision in the form of a chance constraint, providing an interpretable measure of robot safety. To enable real-time evaluation of chance constraints, we consider conservative approximations of chance constraints in the form of distributionally robust Conditional Value at Risk constraints. The resulting formulation offers computational efficiency as well as robustness with respect to out-of-distribution human motion. With the parallelization of a sampling-based optimization technique, our method operates in real-time, demonstrating successful and safe navigation in a number of case studies with real-world pedestrian data.

I. INTRODUCTION

A major challenge in incorporating autonomous robots into real-world scenarios is guaranteeing safe navigation among crowds. Humans present unique challenges due to the stochastic and potentially multi-modal, nature of their motion. However, the inherent stochasticity in human motion makes robust control approaches that account for worst-case scenarios overly cautious. Conversely, conventional stochastic control methods that focus solely on optimizing expected performance do not offer assurances of safety. To bridge this gap, in this work, we consider characterizing robot safety in the form of chance constraints which equip robots with a means to evaluate the risk associated with their actions. Notably, chance constraints also offer an intuitive interpretation of robot safety, enabling a more nuanced navigation strategy.

Evaluating chance constraints relies on understanding the distribution of human motion. Thanks to recent advancements in machine learning (ML), numerous studies explore human trajectory predictions in crowded environments [1], [2], [3], [4]. Nevertheless, while these prediction modules provide reasonable models of human motion distribution, the resulting distribution of human motion is often complicated

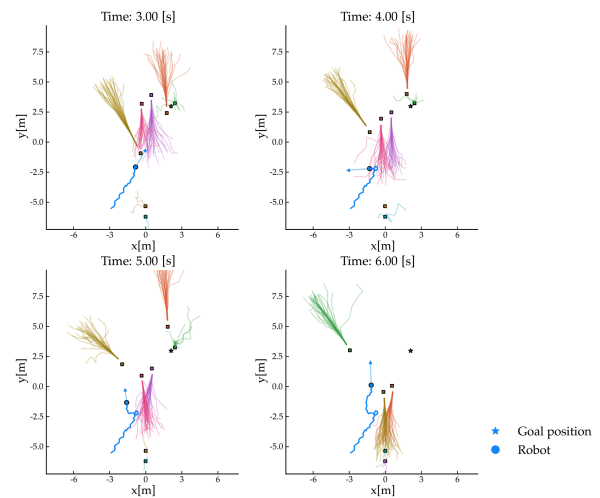
¹Kanghyun Ryu and Negar Mehr are with the Department of Mechanical Engineering, University of California, Berkeley, Berkeley, CA 94720, USA {kanghyun.ryu, negar}@berkeley.edu

This work is supported by the National Science Foundation, under grants CNS-2218759 and CCF-2211542.

The code is available at <https://github.com/labicon/DRCC-MPC.git>



(a) Real-world image of ETH/Hotel sequence.



(b) Example trajectory of a robot navigating among pedestrians in ETH/Hotel sequence.

Fig. 1: Trajectory forecasting model enables a robot to estimate the distribution of future human trajectories. In response, the robot modifies its route to reduce the chance of collision.

and potentially multi-modal. Hence, how to incorporate these forecasts into the robot's control loop in a safe manner remains a challenging task. In addition, their prediction may be inaccurate because of the stochastic nature of human trajectories [5] and the potential uncertainties stemming from sensory inputs or limited data. Relying solely on predictions, which might be imprecise, can lead to system failures, which is particularly critical in environments shared with humans. Thus, the robot's control algorithm must account for the potential inaccuracy in the human trajectory forecasts to ensure safety.

To address these challenges, we propose distributionally robust chance-constrained model predictive control (DRCC-MPC). Our proposed DRCC-MPC aims to satisfy chance constraints not only for an estimated human motion distribution but also for all human trajectory distributions within an ambiguity set. By considering distributionally robust chance

constraints, we can achieve two primary enhancements: 1) it allows us to develop an approximation method for the real-time calculation of chance constraints, and 2) our controller can accommodate potential discrepancies between the actual distribution of human motion and its predicted distribution. We consider robustness against an ambiguity set to which the true distribution of human motion belongs. We construct this ambiguity set assuming that the first and second-order moments of the human motion distribution, which are estimated via Monte Carlo samples of ML-based trajectory forecasts, are known. We exploit the equivalence between chance constraints and Value-at-Risk (VaR) to create a conservative but tractable convex Conditional Value-at-Risk (CVaR) approximation of our safety constraint. Our model predictive controller employs the cross-entropy method (CEM) [6], a sampling-based optimization technique, to identify control sequences that meet the chance constraint. CEM optimization is parallelized using a GPU for real-time implementation. We evaluate the performance of our proposed controller in a number of navigation case studies using real-world datasets and demonstrate that our method outperforms the state-of-the-art in terms of safety in navigating human crowds.

II. RELATED WORK

A. Safe Crowd Navigation

When a robot navigates in a crowded space with other agents, accurately modeling the behavior these agents is crucial. Earlier works employed reciprocal assumption [7], [8] or potential fields [9] for collision avoidance when other agents' velocities are given. However, they were unable to model complex interaction of agents as in crowd navigation. Social force models [10], [11] try to model more complex interaction using repulsion and attraction forces based on each agent's characteristics. Still, these approaches do not account for the uncertainty in other agents' behaviors, making them less suitable for safe navigation. While game theoretic approaches have been another popular method for modeling agent interactions [12], [13], [14], [15], they assume that agents have optimal policy, which may not accurately represent pedestrian movements.

Focusing on the safety aspect of crowd navigation, Control Barrier Functions [16], [17] and Hamilton-Jacobi reachability analysis [18] can guarantee safety against the worst-case behavior of other agents. However, they tend to be overly conservative. To address this, probabilistic constraints have been proposed [19], assuming Gaussian distribution for the state random variable [20], [21]. However, evaluating chance constraints for general multi-modal distributions, such as those from ML-based models, presents computationally challenges [22].

Reinforcement learning (RL) methodologies have shown success in handling complex interactions between multiple agents by leveraging graph structures [23], [24] or attention mechanisms [25], [26], [27]. Yet, these methods often lack the interpretability and explainability required for safety assessment, which is especially crucial when deploying robots in real-world settings around humans.

Alternatively, modular approaches decouple a learning-based trajectory predictor and solve an optimal control problem for planning purposes using exhaustive search [28] or Sequential Action Control (SAC) [29]. However, they rely on soft constraints, which cannot give formal guarantees against collisions and fail to provide an interpretable notion of robot safety. Moreover, an overreliance on potentially flawed machine-learned predictions may pose significant vulnerabilities. To overcome these drawbacks, our methodology seeks interpretable safety assurance while providing robustness against distribution shifts in the trajectory forecasting module.

B. Distributionally Robust Control

Distributionally robust control (DRC) acknowledges that the true distribution of a stochastic parameter might not be accurately captured. From this observation, DRC formulates an ambiguity set, encompassing potential probability distributions of the system [30]. The optimal control action is determined by ensuring that constraints are met for all distributions within the ambiguity set. Owing to its capacity to guarantee safety even with imperfect system modeling, DRC has been extensively employed for control in safety-critical domains [31], [32]. The choice of ambiguity set is pivotal to the performance of DRC. Although several studies have defined the ambiguity set using the Wasserstein metric [33] or ϕ -divergence [34], [35], determining an appropriate radius for these sets is problem-specific and not straightforward to determine in advance. As a result, moment-based ambiguity set [36], [37] has emerged as a popular alternative to modeling moving obstacles [38]. However, their computation time complexity makes real-time operations in crowd navigation unfeasible. While Output Feedback Distributionally Robust RRT* [39] presents a distributionally robust sampling-based motion planning algorithm for computational tractability, it is specifically tailored for Kalman filter-based perception modules, whereas our work can incorporate general ML-based trajectory forecasts.

III. PRELIMINARIES

A. Chance Constraints

Conventional robust constraints [40], [41] which account for worst-case uncertainties cannot handle uncertainties with unbounded support such as the stochastic pedestrian motion. In such scenarios, it is preferable to impose chance constraints that limit the probability of collision.

Definition 1 (Chance constraint over a collision-free set): For a random state x where the underlying distribution is \mathbb{P}^* , we require

$$\text{Prob}^{\mathbb{P}^*}(x \in \mathcal{X}_{free}) \geq 1 - \varepsilon, \quad (1)$$

where \mathcal{X}_{free} is the collision-free space, and ε is the allowable probability of collision.

This chance constraint provides an interpretable measure of safety requiring the system to be collision-free with at least $1 - \varepsilon$ probability. However, assessing a chance constraint, even with a known distribution, remains NP-hard [22]. We




Fig. 2: VaR and CVaR in a normal distribution with allowable probability $\varepsilon = 0.1$. VaR is a $1 - \varepsilon$ quantile of the safety loss and CVaR is the conditional expectation of safety loss over the VaR (blue area).

choose to approximate our chance constraint via *Conditional Value-at-Risk (CVaR)*, which is well known as a tight convex approximation of chance constraints [42]. We reach this approximation through the notion of *Value-at-Risk (VaR)* by considering a notion of a *safety loss* function.

Assumption 1 (Safety loss): We assume that there exists a safety loss function $l : \mathbb{R}^n \rightarrow \mathbb{R}$, which characterizes our collision-free set \mathcal{X}_{free} through its zero level sets:

$$\mathcal{X}_{free} = \{x \in \mathbb{R}^n \mid l(x) \leq 0\}.$$

With this assumption, we can translate chance constraint (1) into a VaR constraint.

Definition 2 (Value-at-Risk): For $x \sim \mathbb{P}^*$, and the safety loss function $l(x)$, we have:

$$\text{VaR}_\varepsilon^{\mathbb{P}^*}(l(x)) := \inf\{\gamma \in \mathbb{R} \mid \text{Prob}^{\mathbb{P}^*}(l(x) > \gamma) \leq \varepsilon\}.$$

VaR represents the potential safety loss $l(x)$ with the allowable probability ε . Since chance constraint (1) requires $\text{Prob}^{\mathbb{P}^*}(l(x) \geq 0) \leq \varepsilon$, in order for the chance constraint to be satisfied, VaR should be less or equal to zero by its definition. We can then relate the resulting VaR Constraint to Conditional Value-at-Risk (CVaR).

Definition 3 (Conditional Value-at-Risk): For a random state x and safety loss function $l(x)$, CVaR is defined as:

$$\text{CVaR}_\varepsilon^{\mathbb{P}^*}(l(x)) := \inf_{\beta \in \mathbb{R}} \left\{ \beta + \frac{1}{\varepsilon} \mathbb{E}_{\mathbb{P}^*} \{ (l(x) - \beta)^+ \} \right\},$$

where $(\cdot)^+ = \max\{\cdot, 0\}$.

CVaR is interpreted as the expected value over VaR, as illustrated in Figure 2. Due to this trait, it naturally can be a conservative, convex approximation of VaR, effectively addressing VaR's non-convex nature [22]:

$$\begin{aligned} \text{CVaR}_\varepsilon^{\mathbb{P}^*}(l(x)) \leq 0 &\Rightarrow \text{VaR}_\varepsilon^{\mathbb{P}^*}(l(x)) \leq 0 \\ &\Leftrightarrow \text{Prob}^{\mathbb{P}^*}(x \in \mathcal{X}_{free}) \geq 1 - \varepsilon. \end{aligned} \quad (2)$$

B. Distributionally Robust Constraints

Although CVaR offers a convex approximation for chance constraints, it still requires an accurate representation of the true distribution \mathbb{P}^* , which can be hard to obtain for pedestrians that exhibit complex and multimodal motions. Therefore, we compute CVaR for a set of distributions \mathcal{P} referred to as the *ambiguity set* with the assumption that $\mathbb{P}^* \in$

\mathcal{P} . Then, we consider constraint (1) to be distributionally robust if and only if

$$\inf_{\mathbb{P} \in \mathcal{P}} \text{Prob}^{\mathbb{P}}(x \in \mathcal{X}_{free}) \geq 1 - \varepsilon.$$

Intuitively, the above constraint requires the satisfaction of chance constraint for all possible distributions \mathbb{P} within the ambiguity set \mathcal{P} . Similarly, we can define distributionally robust CVaR constraint for a random state x as:

$$\sup_{\mathbb{P} \in \mathcal{P}} \text{CVaR}_\varepsilon^{\mathbb{P}}(l(x)) \leq 0.$$

Now, the connection between chance constraints, VaR, and CVaR constraints in (2) can be extended to their distributionally-robust counterparts:

$$\begin{aligned} \sup_{\mathbb{P} \in \mathcal{P}} \text{CVaR}_\alpha^{\mathbb{P}}(l(x)) \leq 0 &\Rightarrow \sup_{\mathbb{P} \in \mathcal{P}} \text{VaR}_\alpha^{\mathbb{P}}(l(x)) \leq 0 \\ &\Leftrightarrow \inf_{\mathbb{P} \in \mathcal{P}} \text{Prob}^{\mathbb{P}}(l(x) \leq 0) \geq 1 - \alpha. \end{aligned} \quad (3)$$

IV. DISTRIBUTIONALLY ROBUST NAVIGATION IN CROWDED ENVIRONMENTS

A. Notation and Problem Formulation

In our notation, we use subscript r to refer to the robot and i for the index of humans. Superscript k denotes time step. Let \mathcal{I}^k be the set of pedestrian indices present in k . Humans and a robot operate in a state space \mathcal{X} , while \mathcal{X}_{robot} indicates the region occupied by a robot. We model the robot and humans as discrete-time dynamical systems.

$$x_i^{k+1} = f_i(x_i^k, u_i^k), \quad \forall i \in \mathcal{I}^k, \quad (4a)$$

$$x_r^{k+1} = f_r(x_r^k, u_r^k). \quad (4b)$$

Here, the robot does not have access to human control u_i^k , rendering x_i^k stochastic from the perspective of the robot.

We employ a distributionally robust chance constraint to ensure robot safety. We define the system to be collision-free when all pedestrians remain outside \mathcal{X}_{robot} . As a result, we propose to formulate our safety constraint as:

$$\inf_{\mathbb{P}_i^k \in \mathcal{P}_i^k} \text{Prob}^{\mathbb{P}_i^k}(x_i^k \in \mathcal{X}_{i,free}^k) \geq 1 - \varepsilon, \quad \forall i \in \mathcal{I}^k, \quad (5)$$

where $\mathcal{X}_{i,free}^k$ represents the collision-free region for human i at time step k . Finally, we formulate a safe robot navigation task as a model predictive control (MPC) problem with constraint (5), which we call DRCC-MPC.

Problem 1 (DRCC-MPC):

$$\min_{\mathbf{x}_r, \mathbf{u}_r} \mathcal{J} := \sum_{k=0}^{K-1} J(x_r^k, u_r^k, x_{goal}) + J_K(x_r^K, x_{goal}) \quad (6a)$$

$$s.t. \quad x_r^{k+1} = f_r(x_r^k, u_r^k), \quad (6b)$$

$$x_r^k \in \mathcal{X}, u_r^k \in \mathcal{U}, \quad (6c)$$

$$\inf_{\mathbb{P}_i^k \in \mathcal{P}_i^k} \text{Prob}^{\mathbb{P}_i^k}(x_i^k \in \mathcal{X}_{i,free}^k) \geq 1 - \varepsilon, \quad \forall i \in \mathcal{I}^k, \quad (6d)$$

where $\mathbf{x}_r := (x_r^0, \dots, x_r^K)$, $\mathbf{u}_r := (u_r^0, \dots, u_r^{K-1})$ are the robot state and input trajectory for horizon K , and \mathcal{U} is the control space of the robot. In our DRCC-MPC, J denotes the stage-wise cost function of the robot and J_K is its terminal cost function.




Fig. 3: For human i , collision-free set $\mathcal{X}_{i,free}^k$ is an ellipsoid centered on μ_i^k which belongs to $\mathcal{X} \setminus \mathcal{X}_{robot}$. We constraint the random human state x_i^k to belong to such a set.

B. Reformulation of Chance Constraints

We assume that every human $i \in \mathcal{I}_k$ satisfies Assumption 1 for $\mathcal{X}_{i,free}^k$ with some safety loss function $l_i^k(x_i^k)$. Constraint (6d) can be approximated by the following distributionally robust CVaR constraint (3), which is a sufficient condition for satisfaction of (6d).

$$\sup_{\mathbb{P}_i^k \in \mathcal{P}_i^k} \text{CVaR}_{\varepsilon}^{\mathbb{P}_i^k}(l_i^k(x_i^k)) \leq 0, \quad \forall i \in \mathcal{I}^k. \quad (7)$$

To make distributionally robust CVaR constraint tractable, we require each safety loss function $l_i^k(x_i^k)$ to be convex and quadratic [36], which leads to an ellipsoidal collision-free set $\mathcal{X}_{i,free}^k$. Considering that x_i^k centers on its mean μ_i^k , we represent $\mathcal{X}_{i,free}^k$ as an ellipsoid centered at μ_i^k outside of \mathcal{X}_{robot} .

$$\begin{aligned} \mathcal{X}_{i,free}^k &:= \{x \in \mathbb{R}^n | l_i^k(x_i^k) \leq 0\} \\ &= \{x \in \mathbb{R}^n : (x - \mu_i^k)^T E_i^k (x - \mu_i^k) + e_i^k \leq 0\}, \end{aligned} \quad (8)$$

with $E_i^k \in \mathbb{R}^{n \times n}$ where n is the dimension of $x_i^k \in \mathbb{R}^n$, $E_i^k \succ 0$, $e_i^k \in \mathbb{R}$, and $\mathcal{X}_{i,free}^k \subset \mathcal{X} \setminus \mathcal{X}_{robot}$. Assuming \mathcal{X}_{robot} is circle in 2D space, we can define an ellipsoid $\mathcal{X}_{i,free}^k$ whose minor axis is tangent to \mathcal{X}_{robot} , as illustrated in Fig 3.

We define the ambiguity set for the i th human \mathcal{P}_i^k as all probability distributions \mathbb{P}_i^k that have a given mean μ_i^k and covariance Σ_i^k .

$$\begin{aligned} \mathcal{P}_i^k &:= \{\mathbb{P}_i^k | \mathbb{E}_{\mathbb{P}_i^k}\{x_i^k\} = \mu_i^k, \\ &\quad \mathbb{E}_{\mathbb{P}_i^k}\{[x_i^k - \mu_i^k] \cdot [x_i^k - \mu_i^k]^T\} = \Sigma_i^k\}. \end{aligned} \quad (9)$$

An advantage of the moment-based ambiguity set \mathcal{P}_i^k is that it requires only μ_i^k and Σ_i^k , which are easily obtained from ML-based forecasting modules with Monte Carlo sampling. Therefore, we will utilize the estimated mean $\hat{\mu}_i^k$ and covariance $\hat{\Sigma}_i^k$ of x_i^k to construct an ambiguity set.

The ellipsoidal collision-free set (8), combined with the moment-based ambiguity set (9), makes the distributionally robust CVaR constraints tractable, which we further simplify with the following result.

Theorem 1: If $\mathcal{X}_{i,free}^k$ is defined as (8) and the random state x_i^k has mean $\hat{\mu}_i^k$ and covariance $\hat{\Sigma}_i^k$, then

$$\sup_{\mathbb{P}_i^k \in \mathcal{P}_i^k} \text{CVaR}_{\varepsilon}^{\mathbb{P}_i^k}(l_i^k(x_i^k)) = e_i^k + \frac{1}{\varepsilon} \text{Tr}\{\hat{\Sigma}_i^k E_i^k\}, \quad (10)$$

where $\text{Tr}\{\cdot\}$ is the trace of a matrix.

Proof: See Proof of Corollary 1.3 in [36] \blacksquare

Theorem 1 provides an aid to evaluate distributionally robust CVaR constraints only using the structure of the collision-free set (8), estimated mean $\hat{\mu}_i^k$, and the estimated covariance $\hat{\Sigma}_i^k$ of x_i^k . Therefore, we can arrive at a tractable sufficient condition for (6d) using (3).

Corollary 1: If $\mathcal{X}_{i,free}^k$ is defined as (8) and the random human state x_i^k has mean $\hat{\mu}_i^k$ and covariance $\hat{\Sigma}_i^k$, then

$$\begin{aligned} e_i^k + \frac{1}{\varepsilon} \text{Tr}\{\hat{\Sigma}_i^k E_i^k\} \leq 0 \Rightarrow \\ \inf_{\mathbb{P}_i^k \in \mathcal{P}_i^k} \text{Prob}^{\mathbb{P}_i^k}(x_i^k \in \mathcal{X}_{i,free}^k) \geq 1 - \varepsilon. \end{aligned} \quad (11)$$

Thus, we can reformulate Problem 1 as:

$$\min_{\mathbf{x}_r, \mathbf{u}_r} \mathcal{J} = \sum_{k=0}^{K-1} J(x_r^k, u_r^k, x_{goal}) + J_K(x_r^K, x_{goal}) \quad (12a)$$

$$\text{s.t. } x_r^{k+1} = f_r(x_r^k, u_r^k), \quad (12b)$$

$$x_r^k \in \mathcal{X}, u_r^k \in \mathcal{U}, \quad (12c)$$

$$e_i^k + \frac{1}{\varepsilon} \text{Tr}\{\hat{\Sigma}_i^k E_i^k\} \leq 0, \quad \forall i \in \mathcal{I}^k, \quad (12d)$$

where e_i^k and E_i^k are constructed from (8).

C. Constrained Cross-Entropy Method

Identifying a feasible solution for the constrained optimization problem (12) remains challenging, particularly due to the non-trivial process involved in retrieving e_i^k and E_i^k in (12d). To address this issue, we adopt the Cross-Entropy Method (CEM), a sampling-based optimization technique, to enable the straightforward evaluation of (12d). CEM recursively iterates between generating a set of robot control sequences $\{\mathbf{u}_r\}^{\text{sampled}}$ from a CEM sampling distribution and refining this distribution based on an elite set $\{\mathbf{u}_r\}^{\text{elite}}$ that achieves optimal performance. This process continues until the maximum number of iterations is reached.

As introduced in [43], we should evaluate whether the sampled trajectories are feasible to solve a constrained optimization problem with CEM. Note that (12d) can be rewritten as:

$$\max_{i \in \mathcal{I}^k} [e_i^k + \frac{1}{\varepsilon} \text{Tr}\{\hat{\Sigma}_i^k E_i^k\}] \leq 0, \quad \forall k \in \{1, \dots, K\}, \quad (13)$$

which indicates that if the pedestrian presenting the maximal risk satisfies the constraint, then all pedestrians will comply with the constraint. Consequently, it is computationally efficient and memory-saving to consider only the pedestrian who poses the greatest risk.

When implementing CEM, during its initial iterations, CEM might fail to sample trajectories that satisfy constraints since its sampling distribution has not yet converged. In such cases, we should select the elites set from the trajectories deemed safest. This ensures that CEM will sample safer trajectories in subsequent iterations. We use a discounted sum of (13) as the *Risk Score*, which evaluates the overall risk of

a robot trajectory in a single scalar value.

$$\text{Risk Score}(\mathbf{x}_r) = \sum_{k=1}^K \gamma^k \max_{i \in \mathcal{I}^k} [e_i^k + \frac{1}{\varepsilon} \text{Tr}\{\hat{\Sigma}_i^k E_i^k\}], \quad (14)$$

where γ is the discount factor. When CEM fails to sample any feasible trajectory, we will use control sequences with the lowest Risk Score as an elite set. This aids CEM to converge to less risky actions and sample safer trajectories in the next iterations. Our algorithm is summarized in Algorithm 1. In our implementation, we used Gaussian distributions for CEM sampling, where μ_{CEM} and σ_{CEM} are the mean and standard deviation of the CEM sampling distribution. After each iteration, they are updated using the mean and standard deviation of an elite set.

V. EXPERIMENTS

A. Implementation Details

We model our robot as a 2-D single integrator $x_r^{k+1} = x_r^k + u_r^k \Delta t$. We employ Trajectron++ [44] as our trajectory forecasting module for human motion. The time step of trajectory generated by Trajectron++ is $\Delta t_{Trajectron} = 0.4[s]$, while the time step for our controller and dynamics is $\Delta t = 0.1[s]$. Therefore, we use linear interpolation on the trajectory estimation to align it with our controller time step. For our evaluations, we use the ETH [45] dataset, which captures various real-world human trajectories.

We consider the following quadratic objective for the robot:

$$J(x_r^k, u_r^k, x_{goal}) = \gamma^k [(x_r^k - x_{goal})^T Q (x_r^k - x_{goal}) + (u_r^k)^T R u_r^k],$$

$$J_K(x_r^K, x_{goal}) = (x_r^K)^T Q_K x_r^K,$$

where $Q = Q_K = 0.5I_{2 \times 2}$, $R = 0.05I_{2 \times 2}$, and $\gamma = 0.99$. A collision is flagged if the human-robot distance is

Algorithm 1 CEM optimization for DRCC-MPC

Input Current robot state x_r^0 , Estimated mean $\hat{\mu}_i^{1:K}$ and covariance $\hat{\Sigma}_i^{1:K}$ of human $i \in \mathcal{I}^k$

Output Control input for robot u_r^0

- 1: **while** stop criteria not satisfied **do**
 - 2: $\{\mathbf{u}_r\}^{sampled} \leftarrow \text{Sample}(\mu_{CEM}, \sigma_{CEM})$
 - 3: $\{\mathbf{x}_r\}^{sampled} \leftarrow \text{dynamics}(x_r^0, \mathbf{u}_r^{sampled})$
 - 4: Evaluate constraint (13)
 - 5: Evaluate objective function (12a)
 - 6: **if** All trajectories not feasible **then**
 - 7: Sort with *RiskScore* (14)
 - 8: $\{\mathbf{u}_r\}^{elite} \leftarrow \text{SelectElite}(\{\mathbf{u}_r\}^{sampled})$
 - 9: **else**
 - 10: $\{\mathbf{x}_r\}^{feasible} \leftarrow \text{SelectFeasible}(\{\mathbf{x}_r\}^{sampled})$
 - 11: Sort with \mathcal{J} (12a)
 - 12: $\{\mathbf{u}_r\}^{elite} \leftarrow \text{SelectElite}(\{\mathbf{u}_r\}^{feasible})$
 - 13: **end if**
 - 14: $\mu_{CEM}, \sigma_{CEM} \leftarrow \text{Update}(\{\mathbf{u}_r\}^{elite})$
 - 15: **end while**
 - 16: Get first control sequence in $\{\mathbf{u}_r\}^{elite}$
 - 17: Output first control input u_r^*
-

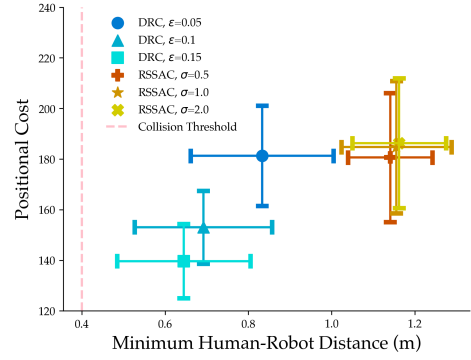


Fig. 4: Sensitivity analysis in Hotel sequence with 100 experiments. As the DRCC-MPC becomes more conservative by reducing ε , the robot tends to keep a larger distance from the humans and decides to take a longer path. On the contrary, RSSAC showed similar behavior for different choices of its risk-sensitivity parameter σ , showing an inability to display a straightforward interpretation of risk and safety.

below $r = 0.4[m]$. We consider the planning horizon to be $K = 40$. The robot operates within a maximum input bound $u_{max} = 2.0[m/s]$. $M = 30$ trajectories are drawn from the Trajectron++ module to construct ambiguity set. During the 5 iterations of CEM optimization, we sample 400 control sequences and select 40 elite set. CEM optimization used GPU parallel computing to achieve real-time performance in a desktop computer with an AMD Ryzen 5 5600 CPU and an NVIDIA GeForce RTX 3060 GPU.

We use the Buffered Input Cell (BIC) method [46] as a reference for reciprocal collision avoidance and the CrowdNav [25] model as RL baseline. Among the modular control strategies, we considered exhaustive search with motion primitives [28] and Risk-Sensitive Sequential Action Control (RSSAC) [29] since they both account for the stochasticity of human behavior. They incorporate a collision cost, measured by the distance between humans and the robot within their optimization objective. While exhaustive search and RSSAC with risk-sensitive parameter $\sigma = 0$ identifies the optimal control by assessing average cost, RSSAC with $\sigma > 0$ optimizes the entropic risk measure [47]. The codes and hyperparameters for these baselines are from RSSAC’s code repository. Lastly, we consider different collision probabilities $\varepsilon = 0.05, 0.1, 0.15$, to observe the effect of the collision probability ε on performance of our method.

B. Real-World Pedestrian Dataset

Our evaluation using the ETH dataset considers two unique sequences: the ”ETH/Hotel” and ”ETH/ETH” sequences. Over a span of 10 seconds, the Hotel sequence has a total of 8 pedestrians, whereas the ETH sequence includes 16 pedestrians. Detailed results can be found in Table I.

In both sequences, DRCC-MPC outperformed the baseline methods in reducing collisions and improving success rate. We note that the poor performance of BIC is attributed to the violation of its reciprocal assumption, primarily because we relied on a fixed real-world pedestrian dataset. This implies that BIC highly relies on a reciprocal assumption, which

(a) Hotel sequence

Method	Collision	Minimum Distance	Success	Positional Cost	Computation Time(ms)
BIC	2.7986 ± 3.6774	0.4253 ± 0.3208	48.46	204.19 ± 153.25	2.80 ± 5.17
CrowdNav	0.3368 ± 1.2555	1.8846 ± 1.7901	92.59	617.61 ± 629.79	59.10 ± 6.02
Exhaustive	0.0519 ± 0.3504	1.2522 ± 0.3287	96.59	240.55 ± 168.66	273.57 ± 101.78
RSSAC($\sigma = 0$)	0.0109 ± 0.1330	0.9473 ± 0.1547	99.32	196.04 ± 120.09	24.08 ± 22.23
RSSAC($\sigma = 1$)	0.0403 ± 0.2687	0.9645 ± 0.2351	97.27	242.33 ± 137.57	23.40 ± 21.62
DRCC-MPC($\varepsilon = 0.05$)	0.0000 ± 0.0000	0.7166 ± 0.1874	100	239.68 ± 142.21	39.58 ± 22.03
DRCC-MPC($\varepsilon = 0.1$)	0.0102 ± 0.1008	0.5706 ± 0.0869	98.97	198.83 ± 122.80	39.66 ± 22.04
DRCC-MPC($\varepsilon = 0.15$)	0.0068 ± 0.0825	0.5312 ± 0.0809	99.32	184.89 ± 116.67	39.72 ± 23.16

(b) ETH sequence

Method	Collision	Minimum Distance	Success	Positional Cost	Computation Time(ms)
BIC	2.5717 ± 2.2883	0.4030 ± 0.3512	31.31	255.96 ± 186.43	2.38 ± 1.41
CrowdNav	0.4878 ± 1.7498	1.5893 ± 1.4401	91.25	599.36 ± 491.81	59.51 ± 5.60
Exhaustive	0.0815 ± 0.3328	1.0024 ± 0.3988	91.92	232.72 ± 169.18	455.46 ± 193.66
RSSAC($\sigma = 0$)	0.0154 ± 0.1886	1.0893 ± 0.5462	99.33	169.62 ± 144.92	36.99 ± 44.48
RSSAC($\sigma = 1$)	0.0208 ± 0.2604	1.0589 ± 0.4661	99.33	195.34 ± 169.11	35.76 ± 49.52
DRCC-MPC($\varepsilon = 0.05$)	0.0034 ± 0.0580	0.8690 ± 0.3310	99.66	285.41 ± 193.76	51.33 ± 31.53
DRCC-MPC($\varepsilon = 0.1$)	0.0000 ± 0.0000	0.6926 ± 0.2723	100	223.97 ± 168.24	51.31 ± 31.43
DRCC-MPC($\varepsilon = 0.15$)	0.0067 ± 0.0819	0.6878 ± 0.3770	99.33	201.52 ± 156.01	50.59 ± 31.71

TABLE I: Simulation result from 300 experiments with different initial and goal states. Average and standard deviation is reported while excluding experiments with initial state collisions. DRCC-MPC achieved safer results compared to other baselines. “Collision”: Represents the collision rate normalized over a 10-second duration. We report the normalized value because of the different control frequencies across methods. ”Minimum Distance”: Minimum distances between the human and robot throughout the navigation. “Success”: Percentage of experiments that were concluded without a collision. “Positional Cost”: Sum of the squared distances to the goal throughout the trajectory $\sum_{k=0}^{end} \|x_r^k - x_{goal}\|^2 \Delta t$. “Computation Time”: Wall time of each method for control and prediction update.

might not be true in a crowd-dense environment. CrowdNav often failed to reach the goal, resulting in a significant performance cost. A potential reason for this could be its initialization through a suboptimal policy [24], making it challenging to adapt to more complex scenarios. Exhaustive search with $9^4 = 6561$ motion primitives failed to run in real-time. Notably, DRCC-MPC ensured fewer collisions and maintained a closer minimum distance than RSSAC, the best-performing baseline. This is because RSSAC’s collision cost includes an added margin to ensure safety. In addition, the risk sensitivity parameter σ in RSSAC did not show a significant difference in navigation safety and it only resulted in higher positional costs compared to the risk-neutral setting. This is because σ is more oriented towards determining yielding behaviors to pedestrians [29]. This interpretation can be non-intuitive, especially when trying to map it directly to real-world safety. In contrast, our collision probability ε in DRCC-MPC is more straightforward and can be directly related to navigation scenarios, offering a clearer and potentially more reliable safety parameter. This confirms our hypothesis that chance constraints provide a more intuitive and interpretable measure of risk.

C. Sensitivity to Chance Constraint Probability

Now, we investigate the impact of collision probability ε on performance. The results from Table I clearly illustrate the safety-efficiency trade-off determined by the choice of ε . For instance, in the Hotel sequence, selecting $\varepsilon = 0.05$ over $\varepsilon = 0.15$ led to a 34% increase in the minimum distance. In the ETH sequence, this increment was 26%. However, opting for the smaller ε value also resulted in substantial increase in positional cost: 29% in the Hotel sequence and 41% in

the ETH sequence.

For a more in-depth examination of the effect of ε , we carried out a controlled experiment, holding the robot’s starting and goal positions constant. Fig 4 illustrate the mean and standard deviation of minimum human-robot distance and positional cost, for various ε . We observed a direct relationship between ε and minimum human-robot distance; smaller ε increased the minimum human-robot distance. Conversely, a larger ε led the robot to navigate through riskier paths, thereby reducing the positional cost. As a result, the direct relationship between ε , trajectory uncertainty, and safe distance becomes readily comprehensible. This relationship provides system designers with the significant advantage of being able to adjust the safety-performance trade-off easily. This contrasts to the risk sensitivity parameter σ in RSSAC, which did not influence the outcome of the safety evaluation.

VI. CONCLUSIONS AND FUTURE WORKS

In this paper, we addressed the problem of safe robot navigation in the presence of moving pedestrians. We incorporated a distributionally robust chance constraint as a risk metric in our MPC problem. We modeled a moment-based ambiguity set from Monte Carlo samples of human trajectory predictions. Then, distributionally robust chance constraint was reformulated using an ellipsoidal collision-free set. Finally, a control sequence that satisfied the proposed constraint was found with CEM, which was implemented with GPU for real-time computation. Our control algorithm successfully demonstrated the ability to navigate safely using a human motion forecaster which was trained on a real-world dataset. Looking forward, we intend to extend our work to environments where humans change their behavior in response to the robot’s action.

REFERENCES

- [1] Y. Yuan, X. Weng, Y. Ou, and K. M. Kitani, "Agentformer: Agent-aware transformers for socio-temporal multi-agent forecasting," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9813–9823, 2021.
- [2] S. Ettinger, S. Cheng, B. Caine, C. Liu, H. Zhao, S. Pradhan, Y. Chai, B. Sapp, C. R. Qi, Y. Zhou, *et al.*, "Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9710–9719, 2021.
- [3] Y. Huang, J. Du, Z. Yang, Z. Zhou, L. Zhang, and H. Chen, "A survey on trajectory-prediction methods for autonomous driving," *IEEE Transactions on Intelligent Vehicles*, vol. 7, no. 3, pp. 652–674, 2022.
- [4] K. Mangalam, Y. An, H. Girase, and J. Malik, "From goals, waypoints & paths to long term human trajectory forecasting," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 15233–15242, 2021.
- [5] R. Cheng, R. M. Murray, and J. W. Burdick, "Limits of probabilistic safety guarantees when considering human uncertainty," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3182–3189, IEEE, 2021.
- [6] Z. I. Botev, D. P. Kroese, R. Y. Rubinstein, and P. L'Ecuyer, "The cross-entropy method for optimization," in *Handbook of statistics*, vol. 31, pp. 35–59, Elsevier, 2013.
- [7] J. Van den Berg, M. Lin, and D. Manocha, "Reciprocal velocity obstacles for real-time multi-agent navigation," in *2008 IEEE international conference on robotics and automation*, pp. 1928–1935, Ieee, 2008.
- [8] J. Van Den Berg, S. J. Guy, M. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," in *Robotics Research: The 14th International Symposium ISRR*, pp. 3–19, Springer, 2011.
- [9] N. Pradhan, T. Burg, and S. Birchfield, "Robot crowd navigation using predictive position fields in the potential function framework," in *Proceedings of the 2011 American control conference*, pp. 4628–4633, IEEE, 2011.
- [10] D. Yang, Ü. Özgüner, and K. Redmill, "Social force based microscopic modeling of vehicle-crowd interaction," in *2018 IEEE Intelligent Vehicles Symposium (IV)*, pp. 1537–1542, IEEE, 2018.
- [11] D. Yang, Ü. Özgüner, and K. Redmill, "A social force based pedestrian motion model considering multi-pedestrian interaction with a vehicle," *ACM Transactions on Spatial Algorithms and Systems (TSAS)*, vol. 6, no. 2, pp. 1–27, 2020.
- [12] Z. Williams, J. Chen, and N. Mehr, "Distributed potential ilqr: Scalable game-theoretic trajectory planning for multi-agent interactions," *arXiv preprint arXiv:2303.04842*, 2023.
- [13] M. Wang, N. Mehr, A. Gaidon, and M. Schwager, "Game-theoretic planning for risk-aware interactive agents," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 6998–7005, IEEE, 2020.
- [14] M. Bhatt, Y. Jia, and N. Mehr, "Efficient constrained multi-agent trajectory optimization using dynamic potential games," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 7303–7310, IEEE, 2023.
- [15] T. Kavuncu, A. Yaraneri, and **Negar Mehr**, "Potential iLQR: A Potential-Minimizing Controller for Planning Multi-Agent Interactive Trajectories," in *Proceedings of Robotics: Science and Systems (RSS)*, (Virtual), July 2021.
- [16] Z. Jian, Z. Yan, X. Lei, Z. Lu, B. Lan, X. Wang, and B. Liang, "Dynamic control barrier function-based model predictive control to safety-critical obstacle-avoidance of mobile robot," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3679–3685, IEEE, 2023.
- [17] W. Luo, W. Sun, and A. Kapoor, "Multi-robot collision avoidance under uncertainty with probabilistic safety barrier certificates," *Advances in Neural Information Processing Systems*, vol. 33, pp. 372–383, 2020.
- [18] M. Chen, J. C. Shih, and C. J. Tomlin, "Multi-vehicle collision avoidance via hamilton-jacobi reachability and mixed integer programming," in *2016 IEEE 55th Conference on Decision and Control (CDC)*, pp. 1695–1700, IEEE, 2016.
- [19] O. Andersson, M. Wzorek, P. Rudol, and P. Doherty, "Model-predictive control with stochastic collision avoidance using bayesian policy optimization," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4597–4604, IEEE, 2016.
- [20] D. Althoff, M. Althoff, D. Wollherr, and M. Buss, "Probabilistic collision state checker for crowded environments," in *2010 IEEE international conference on robotics and automation*, pp. 1492–1498, IEEE, 2010.
- [21] 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, pp. 51–76, 2013.
- [22] A. Nemirovski and A. Shapiro, "Convex approximations of chance constrained programs," *SIAM Journal on Optimization*, vol. 17, no. 4, pp. 969–996, 2007.
- [23] Q. Li, F. Gama, A. Ribeiro, and A. Prorok, "Graph neural networks for decentralized multi-robot path planning," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 11785–11792, IEEE, 2020.
- [24] S. Liu, P. Chang, W. Liang, N. Chakraborty, and K. Driggs-Campbell, "Decentralized structural-rnn for robot crowd navigation with deep reinforcement learning," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3517–3524, IEEE, 2021.
- [25] C. Chen, Y. Liu, S. Kreiss, and A. Alahi, "Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning," in *2019 international conference on robotics and automation (ICRA)*, pp. 6015–6022, IEEE, 2019.
- [26] S. Liu, P. Chang, Z. Huang, N. Chakraborty, K. Hong, W. Liang, D. L. McPherson, J. Geng, and K. Driggs-Campbell, "Intention aware robot crowd navigation with attention-based interaction graph," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 12015–12021, IEEE, 2023.
- [27] E. Leurent and J. Mercat, "Social attention for autonomous decision-making in dense traffic," *arXiv preprint arXiv:1911.12250*, 2019.
- [28] E. Schmerling, K. Leung, W. Vollprecht, and M. Pavone, "Multimodal probabilistic model-based planning for human-robot interaction," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3399–3406, IEEE, 2018.
- [29] H. Nishimura, B. Ivanovic, A. Gaidon, M. Pavone, and M. Schwager, "Risk-sensitive sequential action control with multi-modal human trajectory forecasting for safe crowd-robot interaction," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 11205–11212, IEEE, 2020.
- [30] H. Rahimian and S. Mehrotra, "Distributionally robust optimization: A review," *arXiv preprint arXiv:1908.05659*, 2019.
- [31] M. Schuurmans, P. Sotasakis, and P. Patrinos, "Safe learning-based control of stochastic jump linear systems: a distributionally robust approach," in *2019 IEEE 58th Conference on Decision and Control (CDC)*, pp. 6498–6503, IEEE, 2019.
- [32] K. Long, Y. Yi, J. Cortes, and N. Atanasov, "Safe and stable control synthesis for uncertain system models via distributionally robust optimization," *arXiv preprint arXiv:2210.01341*, 2022.
- [33] A. Hakobyan and I. Yang, "Wasserstein distributionally robust motion control for collision avoidance using conditional value-at-risk," *IEEE Transactions on Robotics*, vol. 38, no. 2, pp. 939–957, 2021.
- [34] A. Sinha, M. O'Kelly, H. Zheng, R. Mangharam, J. Duchi, and R. Tedrake, "Formulazero: Distributionally robust online adaptation via offline population synthesis," in *International Conference on Machine Learning*, pp. 8992–9004, PMLR, 2020.
- [35] H. Nishimura, N. Mehr, A. Gaidon, and M. Schwager, "Rat ilqr: A risk auto-tuning controller to optimally account for stochastic model mismatch," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 763–770, 2021.
- [36] B. P. Van Parys, D. Kuhn, P. J. Goulart, and M. Morari, "Distributionally robust control of constrained stochastic systems," *IEEE Transactions on Automatic Control*, vol. 61, no. 2, pp. 430–442, 2015.
- [37] S. Zymler, D. Kuhn, and B. Rustem, "Distributionally robust joint chance constraints with second-order moment information," *Mathematical Programming*, vol. 137, 2013.
- [38] B. HomChaudhuri, "Distributionally robust model predictive control for collision avoidance in the presence of stochastic obstacles," *Proceedings of the Institution of Mechanical Engineers, Part I: Journal of Systems and Control Engineering*, p. 09596518221122721, 2023.
- [39] V. Renganathan, I. Shames, and T. H. Summers, "Towards integrated perception and motion planning with distributionally robust risk constraints," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 15530–15536, 2020.
- [40] R. Soloperto, J. Köhler, F. Allgöwer, and M. A. Müller, "Collision avoidance for uncertain nonlinear systems with moving obstacles using

- robust model predictive control,” in *2019 18th European Control Conference (ECC)*, pp. 811–817, IEEE, 2019.
- [41] P. Bouffard, A. Aswani, and C. Tomlin, “Learning-based model predictive control on a quadrotor: Onboard implementation and experimental results,” in *2012 IEEE International Conference on Robotics and Automation*, pp. 279–284, IEEE, 2012.
- [42] Z. Chen, D. Kuhn, and W. Wiesemann, “On approximations of data-driven chance constrained programs over wasserstein balls,” *arXiv preprint arXiv:2206.00231*, 2022.
- [43] Z. Liu, H. Zhou, B. Chen, S. Zhong, M. Hebert, and D. Zhao, “Constrained model-based reinforcement learning with robust cross-entropy method,” *arXiv preprint arXiv:2010.07968*, 2020.
- [44] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, “Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data,” in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVIII 16*, pp. 683–700, Springer, 2020.
- [45] S. Pellegrini, A. Ess, K. Schindler, and L. Van Gool, “You’ll never walk alone: Modeling social behavior for multi-target tracking,” in *2009 IEEE 12th international conference on computer vision*, pp. 261–268, IEEE, 2009.
- [46] M. Wang, Z. Wang, S. Paudel, and M. Schwager, “Safe distributed lane change maneuvers for multiple autonomous vehicles using buffered input cells,” in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4678–4684, IEEE, 2018.
- [47] A. Majumdar and M. Pavone, “How should a robot assess risk? towards an axiomatic theory of risk in robotics,” in *Robotics Research: The 18th International Symposium ISRR*, pp. 75–84, Springer, 2020.