Proactive Model Predictive Control with Multi-Modal Human Motion Prediction in Cluttered Dynamic Environments

Lukas Heuer[1], Luigi Palmieri[2], Andrey Rudenko[2], Anna Mannucci[2], Martin Magnusson[1] and Kai O. Arras[2]

***Abstract*— For robots navigating in dynamic environments, exploiting and understanding uncertain human motion prediction is key to generate efficient, safe and legible actions. The robot may perform poorly and cause hindrances if it does not reason over possible, multi-modal future social interactions. With the goal of enhancing autonomous navigation in cluttered environments, we propose a novel formulation for nonlinear model predictive control including multi-modal predictions of human motion. As a result, our approach leads to less conservative, smooth and intuitive human-aware navigation with reduced risk of collisions, and shows a good balance between task efficiency, collision avoidance and human comfort. To show its effectiveness, we compare our approach against the state of the art in crowded simulated environments, and with real-world human motion data from the THÖR dataset. This comparison shows that we are able to improve task efficiency, keep a larger distance to humans and significantly reduce the collision time, when navigating in cluttered dynamic environments. Furthermore, the method is shown to work robustly with different state-of-the-art human motion predictors.**

I. INTRODUCTION

Planning legible and safe yet efficient trajectories in human-shared environments is challenging due to the intrinsically stochastic nature of human motion. Recent works show that keeping a safe and conservative distance from humans, based on the reachable space, is a workable approach for collision avoidance that can provide strong safety guarantees [1]–[5]. However, this strategy may yield very cautious robot behavior since the robot will always assume the worst-case scenario instead of reasoning on future human intentions and behaviors.

Using more informed motion predictions may reduce this conservativeness by providing the robot with an educated guess on the future movement of humans and thereby letting it plan its motion in a less constrained space.

Nonlinear model predictive control (NMPC) is a widely used method for local motion planning and control of autonomous mobile robots [7]–[9]. Several NMPC approaches use constraints on the maneuverable space according to the current and predicted position of humans [10], [11]. However, constraining NMPC can make it hard to find a feasible solution to the motion planning problem in cluttered dynamic environments [10].

[1]L. Heuer, and M. Magnusson are with the School of Science and Technology, Örebro University, Sweden {lukas.heuer, martin.magnusson}@oru.se.

[2]L. Palmieri, A. Rudenko, A. Mannucci and K. O. Arras are with Robert Bosch GmbH, Corporate Research, Stuttgart, Germany {luigi.palmieri, andrey.rudenko, anna.mannucci, kaioliver.arras}@de.bosch.com.

This work was partly supported by the EU Horizon 2020 research and innovation program under grant agreement No. 101017274 (DARKO).

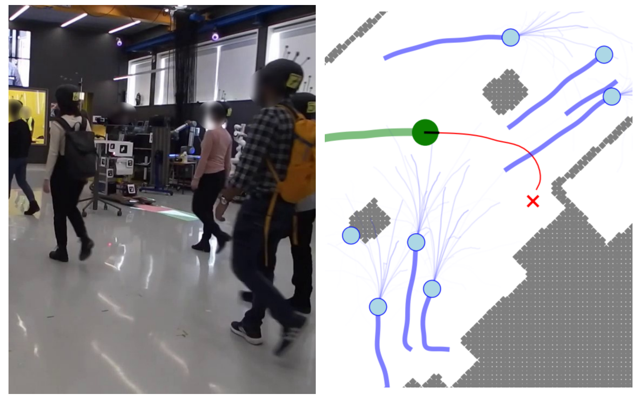

(a) (b)

Fig. 1: **(a)** Image from the the THÖR dataset [6] which is used to evaluate our multi-modal prediction aware motion planning approach. **(b)** An exemplary scenario from our experiments with the robot (**green**) navigating amongst multiple humans (**blue**). The local motion plan (**red**) is generated by the NMPC method using our novel collision avoidance cost for stochastic multi-modal predictions.

A recent trend in human-motion forecasting is to adopt multi-modal predictions [12] to better account for the uncertainty in intent and direction of movement. While there are methods that include multi-modal human motion predictions in NMPC [11], successfully doing so in cluttered dynamic environments remains a challenge. This is because complex situations potentially cause constraint-based NMPC approaches to freeze the robot (i.e., stopping the robot) due to a large number of collision avoidance constraints being active at the same time (e.g., due to optimizers returning infeasible solutions). This is even more common when multi-modal human motion predictions are considered. In this work, we improve on the state of the art by utilizing a multi-modal-aware collision-avoidance-cost, without imposing constraints based on predictions. Our hypothesis is that this greatly reduces the risk of a freezing robot when using multi-modal predictions and increases its ability to actively avoid collisions.

To this end, we introduce the Multi-Modal Collision Avoidance NMPC (MMCA-NMPC), a novel model-based predictive planning approach for robot navigation in crowded spaces. MMCA-NMPC utilizes a new collision avoidance cost formulation based on stochastic multi-modal human motion predictions and is agnostic to the human motion prediction method. We systematically compare our approach against state-of-the-art methods using a real-world human motion dataset and interactive simulations in complex environments. The results show that our method significantly outperforms the baselines if one considers the combined performance in task efficiency, human awareness, and collision avoidance.

II. Related Work

Early methods for collision avoidance include model-based techniques that do not consider predictions of humans, one the most popular being Dynamic Window Approach (DWA) [13]. Lately there has been an increasing interest in MPC techniques for collision avoidance [7], [9], [14], [15].

In this work, building on our previous results [7], [12], [15], we shift our attention towards the utilization of stochastic, multi-modal predictions, which are better suited to represent the complexity of erratic and dynamic environments. We formulate a cost term by applying ideas from Chance-Constrained NMPC (CC-NMPC) to the concept of multi-modal human motion prediction. CC-NMPC was first investigated by Blackmore et al. [16] for static obstacle avoidance and extended more recently in [10], [17] to handle dynamic obstacles like humans. Zhu et al. [10] apply CC-NMPC for motion planning and collision avoidance of micro aerial vehicles. A similar approach is presented by Castillo-Lopez et al. [17]. While CC-NMPC achieves good results in terms of efficiency and safety, these methods work with uni-modal Gaussian predictions and are not evaluated in complex, cluttered, and dynamic 2D environments.

Several approaches have exploited multi-modal human motion predictions for generating robot motion [8], [11], [18], [19]. Schmerling et al. [18] and Schaefer et al. [8] use NMPC formulations which include interactions between the agents in the scene. Schaefer et al. [8] specifically penalize predicted human-robot interactions. However, differently from our approach, their method is limited to multi-modal predictions that depend on, and are differentiable with respect to, the robot's future controls and only perform explicit collision avoidance in specific situations. Nair et al. [11] use a multi-modal version of the chance constraint, but they do not evaluate their approach in cluttered environments. Similarly to Schaefer et al. [8] and Kamel et al. [20], our approach is designed for minimally-invasive robot navigation, and it can be fine-tuned to reduce the conservativeness of the robot behavior. Differently from other methods that use multi-modal predictions, our method directly considers the spacial relation between the predicted modes of motion.

Several Deep Reinforcement Learning (DRL) approaches learn a policy for navigation in crowded environments (e.g., GA3C-CADRL [21] or SA-CADRL [22]). Those approaches learn human-aware robot behaviors, but differently from our approach they often adopt a limited amount of predefined discrete actions, thus do not completely account for the dynamics of the robot and the environment.

III. Background

In this section, we briefly review NMPC (Section III-A) with different collision avoidance strategies (Section III-B, and methods for multi-modal human motion predictions (Section III-C).

A. Nonlinear Model Predictive Control

NMPC formulates a short-term motion planning problem for a prediction horizon T as a nonlinear program. This program is solved iteratively, and only the first control action of the optimal solution is applied to the system at each iteration. An overview of the numerical methods used to solve optimization problems in NMPC is given by Rawlings et al. [23] in Chapter 8. The stochastic nonlinear optimal control problem solved at every NMPC iteration is defined as:

$$
\begin{aligned}
&\min_{\mathbf{x},\mathbf{u}} \quad J(\mathbf{x},\mathbf{u}) && \qquad (1)\\
\text{s.t.: } & \mathbf{x}_{t+1} = f(\mathbf{x}_t,\mathbf{u}_t) + \boldsymbol{\omega} && \forall t \in [0, T-1]\\
& \mathbf{x}_0 = \mathbf{x}_m \\
& h_t(\mathbf{x}_t,\mathbf{u}_t) \leq 0 && \forall t \in [0, T]\\
& \mathbf{x}_t \in \mathbb{X} && \forall t \in [0, T]\\
& \mathbf{u}_t \in \mathbb{U} && \forall t \in [0, T-1],
\end{aligned}
$$

where $\mathbf{x}_t \in \mathbb{X}$ and $\mathbf{u}_t \in \mathbb{U}$ are the state and control of the robot for a given timestep t, respectively. $\mathbb{X} \subseteq \mathbb{R}^{N_x}$ and $\mathbb{U} \subseteq \mathbb{R}^{N_u}$ denote the state and control spaces. The vector $\boldsymbol{\omega} \in \mathbb{R}^{N_\omega}$ denotes the process disturbances with zero-mean Gaussian probability distributions and variance $\boldsymbol{\Sigma}_\omega$. $J : \mathbb{R}^{N_x} \times \mathbb{R}^{N_u} \rightarrow \mathbb{R}^1$ is the chosen objective function. The first constraint in Eq. 1 represents the (possibly nonlinear) robot kinematics $f(\mathbf{x},\mathbf{u})$. The measured initial state is given as $\mathbf{x}_m \in \mathbb{R}^{N_x}$. The function $h_t(\mathbf{x},\mathbf{u}) : \mathbb{R}^{N_x} \times \mathbb{R}^{N_u} \rightarrow \mathbb{R}^{N_h}$ collects arbitrary inequality constraints.

B. Collision Avoidance in NMPC

To avoid collisions, NMPC needs to include information about obstacles surrounding the robot. In dynamic environments, this information is time-dependent. To leverage this information, one can define a collision avoidance constraint, or include it as costs in the objective function.

1) Constraint-Based: Most recent work uses constraints to enforce a minimum safety distance between the robot and the humans at all times [10], [17], [20], [24], [25]. Such constraint-based approaches guarantee that the optimal solution does not violate the defined constraints. However, a high number of constraints also increases the computational complexity of the nonlinear program. Most importantly, constraints also introduce the possibility that the optimizer is not able to find a solution to the program, thus resulting in an infeasible iteration.

2) Cost-Based: While some of the above-mentioned approaches use slack variables to tie the constraints to the objective functions, using a cost term to achieve collision avoidance is less popular. Schaefer et al. [8] and Kamel et al. [20] use a cost term to drive the optimal solution away from the predicted obstacle positions. However, they do not provide a cost formulation that explicitly uses generic multi-modal predictions. A cost term in the objective function does not impose any hard limits but results in a "the less, the better" behavior. Thus, a cost term cannot provide formal safety guarantees but cannot cause the nonlinear program to be infeasible either.

C. Multi-Modal Human Motion Prediction

Rudenko et al. [12] show that the literature does not provide an entirely consistent definition for multi-modal motion predictions. In some cases, the multiple modes of human motion reflect the possible intentions or paths towards a goal [26]. Other authors use the term to refer to a multi-modal probability distribution e.g., a Gaussian mixture model (GMM) [27]–[29].

Remark 1: In this work, we consider a stochastic multi-modal prediction of human states over the time horizon τ. Thus, for each discrete timestep $t \in [0, \tau]$ the predicted human position $\mathbf{p}_t = (x, y)$ is a random variable, drawn from a multivariate GMM: $\mathbf{p}_t \sim \sum_{z=1}^{Z} \beta_z \, \mathcal{N}(\hat{\mathbf{p}}_{z,t}, \mathbf{\Sigma}_{z,t})$ with mixture components $z \in Z$, mixture probability β_z, mixture component mean $\hat{\mathbf{p}}_z$ and covariance matrix $\mathbf{\Sigma}_z$. Consequently, each mode of motion can be represented by the sequence of the corresponding multivariate Gaussian distributions.

It can generally be assumed, that the spatial configuration of the mixture components in the GMMs is important when interpreting the prediction and using it in the motion planning problem. For instance, if the means of the different mixture components have a large distance to each other but similar probabilities, this corresponds to higher uncertainty in the human intent.

IV. Multi-Modal Collision Avoidance

Using constraint-based collision avoidance is a popular approach when integrating human motion prediction into NMPC [10], [17], [20], [24], [30]. However, this approach may prove impractical for cluttered environments with a high number of humans. Especially when navigating in a cluttered and crowded 2D plane, having many active collision avoidance constraints can easily lead to an infeasible NMPC iteration. Infeasible iterations are typically handled by using controls that stop the robot [10] as there is no safe, optimal control action available. Using multi-modal predictions (as in [11]) requires even more constraints on the state space, causing a decrease in computational performance and a higher likelihood of a freezing robot.

Importantly, the main benefit of constraints is that they are able to give guarantees for the resulting solution to the motion planning problem. This is important when discussing the safety of an NMPC method. However, the guarantees given by a collision avoidance constraint only hold if the underlying prediction is absolutely accurate. If we do not assume to have a perfect predictor, we are still left with an unknown risk of collision which stems from the probability that our prediction is wrong.

Because of the two arguments presented above, our proposed NMPC formulation uses constraints only on the current measured position of the humans at each timestep. Additionally, we use stochastic multi-modal human motion prediction to formulate a collision cost, therefore accounting for the predicted human positions and actively avoid possible collisions. Because we use a constraint only on the current position of the humans, our approach is less effected by an errornous prediction. In other words, we utilize predictions in a way that cannot cause an infeasible NMPC iteration. This ensures that the NMPC method will provide valid control actions even in complex, cluttered, dynamic situations and the robot will be able to actively participate in avoiding collisions instead of being frozen. Our novel formulation of a collision avoidance cost is inspired by the convex chance constraint from [17] and uses multi-modal human motion predictions, formulated as a sequence of GMMs, to find the optimal robot trajectory.

A. Multi-Modal Collision Avoidance NMPC

Our novel Multi-Modal Collision Avoidance (MMCA-) NMPC is described in Alg. 1 and 2. At each iteration we solve the following nonlinear program (NLP):

$$\begin{aligned} \min_{\mathbf{x}_R, \mathbf{u}_R} \quad & J_g(\mathbf{x}_R, \mathbf{u}_R) + J_{\text{col}}(\mathbf{x}_R, \mathbf{x}_H) && (2) \\ \text{s.t.: } & \mathbf{x}_{R,t+1} = f(\mathbf{x}_{R,t}, \mathbf{u}_{R,t}) + \boldsymbol{\omega} && \forall t \in [0, T-1] \\ & \mathbf{x}_{R,0} = \mathbf{x}_{R,m} \\ & \|\hat{\mathbf{p}}_{R,t} - \hat{\mathbf{p}}^i_{H,0}\|_2 \geq d^i && \forall t \in [0, T], \forall i \in I \\ & \mathbf{x}_{R,t} \in \mathbb{X} && \forall t \in [0, T] \\ & \mathbf{u}_{R,t} \in \mathbb{U} && \forall t \in [0, T-1], \end{aligned}$$

where $\mathbf{x}_R, \mathbf{u}_R$ are the robot's state and control vectors, T is the prediction horizon and I is a set of all considered humans. The subscripts $\cdot_R$ and $\cdot_H$ denote the association of a variable $\cdot$ with the robot or a human, respectively. Specifically, we use $\cdot_H = [\cdot^i_H]_{i \in H}$ with $\cdot^i_H$ being the value of the variable $\cdot$ for human i. The symbol $\mathbf{p}$ describes Cartesian coordinates (x, y) of a state vector $\mathbf{x}$ and is given as a multivariate random variable $\mathbf{p} \sim \mathcal{N}(\hat{\mathbf{p}}, \mathbf{\Sigma})$. Thus, the robot and humans positional distributions are denoted as $\mathbf{p}_R$ and $\mathbf{p}_H$ respectively. Finally, we refer to d^i as the distance between the robot and the human i.

The objective function in Eq. 2 is a linear combination of the collision cost $J_{\text{col}}(\mathbf{x}_R, \mathbf{x}_H)$, detailed in Section IV-B, and the goal cost:

$$\begin{aligned} J_g(\mathbf{x}_R, \mathbf{u}_R) = \quad & \frac{\|\mathbf{x}_{R,T} - \mathbf{x}_{R,g}\|_S^2}{\rho_g} + \\ & \sum_{t=0}^{T-1} \frac{\|\mathbf{x}_{R,t} - \mathbf{x}_{R,g}\|_Q^2}{\rho_g} + \|\mathbf{u}_{R,t}\|_W^2, \end{aligned} \tag{3}$$

with Q, W, and S as diagonal matrices, weighting the stage error, the control effort, and the terminal error, and $\mathbf{x}_{R,g}$ as the goal state. The scalar $\rho_g = \|\mathbf{x}_{R,0} - \mathbf{x}_{R,g}\|^2$ normalizes by the current squared distance from the robot to the goal state.

Before solving the NLP, we use a motion prediction method to generate a multi-modal motion prediction for every human (see Alg. 2, line 1). We finalize the NMPC iteration by applying the first element of the optimal control sequence (see Alg. 1, line 4).

Since the collision cost term J_{col} cannot guarantee safety, we employ a constraint on each current human positions. Importantly, we only constrain on the current position of the humans, and not their predicted trajectory. As the positions and therefore the constraints update frequently (e.g. 10 Hz),

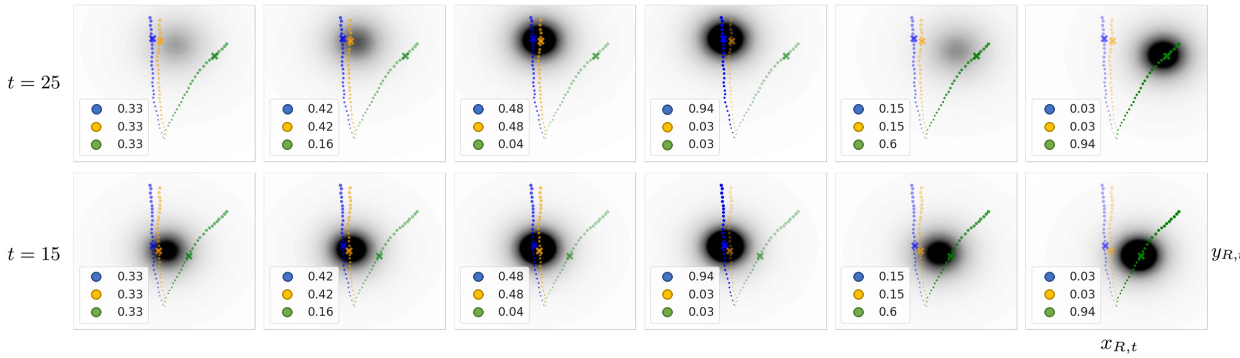

Fig. 2: The cost maps resulting from $J_{\text{col}}(\mathbf{x}_R, \mathbf{x}_H)$ in Eq. 6 at predicted timesteps $t = 25$ (**Top**) and $t = 15$ (**Bottom**) in a scenario with only one human with a three-modal human motion prediction of $\mathbf{x}_H$. The six **columns** show the cost maps for different triples of the mode probability $\beta_{\{1,2,3\}}$ of $\mathbf{x}_H$ with the specific values given in the legends. The three possible human modes are shown in yellow, blue and green. The predicted position of the human $\mathbf{x}_H$ at timestep t is marked with an $\times$. The cost maps are given over $\mathbf{x}_R = (x, y)$, and darker shade corresponds to higher cost. The different **columns** illustrate how our cost formulation allocates cost, dependent on the individual mode probabilities. Note that a higher cost is allocated if more modes predict the human to be in the same area. If the modes diverge but are equally likely, this uncertainty about the human's intent and future trajectory results in a lower cost overall. The presence of multiple humans will result in summing up their individual costs.

this means the robot is not able to actively cause a collision by directly driving into a human. The constraint for all timesteps $t \in T$ is written as

$$\|\hat{\mathbf{p}}_{R,t} - \hat{\mathbf{p}}^{i}_{H,0}\|_2 \geq d^i. \tag{4}$$

Similarly to [10], if the optimizer does not find a feasible solution for an NMPC iteration, we use decelerate-to-stop controls for that given iteration.

B. Multi-Modal Collision Avoidance Cost

Given a multi-modal human motion prediction with modes $z \in Z$, as described in Section III-C, with β^i_z as the probability of mixture-component z from the prediction of human i. We formulate a convex function that depends on the state of the robot and a multi-modal human motion prediction as follows:

$$c(\mathbf{x}_{R,t}, \mathbf{x}^i_{H,t}) = \sum_{z \in Z} \beta^i_z \sum_{j=1}^{2} \left(\frac{\hat{p}^j_{R,t} - \hat{p}^{i,j}_{H,z,t}}{d^i + \sqrt{\left(\sigma^j_{R,t}\right)^2 + \left(\sigma^{i,j}_{H,z,t}\right)^2}} \right)^2, \tag{5}$$

where index j runs over the two Cartesian coordinates (x, y) and with $\left(\sigma^j\right)^2$ as the variance in the respective dimension.

We then define the final collision avoidance cost as:

$$J_{\text{col}}(\mathbf{x}_R, \mathbf{x}_H) = \sum_{t=0}^{T} \sum_{i=1}^{I} \frac{g}{c(\mathbf{x}_{R,t}, \mathbf{x}^i_{H,t})}, \tag{6}$$

with g as a gain that can be used to influence the conservativeness of the robot when engaging with humans.

A visual illustration of J_{col} is shown in Fig 2. For highly ambiguous situations, one planning iteration might yield an optimal trajectory that can be potentially unsafe. However, as only the first timestep of the trajectory is executed before the entire situation is reevaluated, such ambiguity is resolved by the subceeding NMPC-iterations.

Algorithm 1 MMCA-NMPC

Require: $\mathbf{x}_{R,g},\ \Delta t$ ▷ goal state, sampling time
1: **while** $\mathbf{x}_{R,m} \neq \mathbf{x}_{R,g}$ **do** ▷ goal reached?
2: $\quad \mathbf{x}_{R,m} \leftarrow$ GETCURRENTSTATE()
3: $\quad \mathbf{u}_R \leftarrow$ MMCA-NLP-ITERATION($\mathbf{x}_{R,g},\ \mathbf{x}_{R,m},\ \Delta t$)
4: $\quad$ CONTROLROBOT($\mathbf{u}_{R,0}$) ▷ apply control $\mathbf{u}_0$
5: **end while**

Algorithm 2 MMCA-NLP-iteration

Require: $\mathbf{x}_{R,g},\ \mathbf{x}_{R,m},\ b,\ \Delta t$ ▷ goal state, measured state, decelerate-to-stop controls, sampling time
1: $\mathbf{x}_H \leftarrow$ GETMULTIMODALPREDICTIONS()
2: NLP $\leftarrow$ PREPARENLP($\mathbf{x}_{R,g},\ \mathbf{x}_{R,m}, \mathbf{x}_H,\ \Delta t$)
3: **if** NLP.ISFEASIBLE() **then**
4: $\quad \mathbf{u}_R \leftarrow$ NLP.SOLVE() ▷ Solve NLP in Eq. 2
5: **else**
6: $\quad \mathbf{u}_R \leftarrow b$ ▷ use decelerate-to-stop controls
7: **end if**
8: **return** $\mathbf{u}_R$

V. EVALUATION

In this section, we present the experimental setup and the scenarios used in our evaluation. Furthermore, we detail relevant performance metrics and the baselines (Section V-D) we compare our methods against.

A. Setup

In all experiments, the robot drives according to differential-drive kinematics: $\mathbf{x}_R = (x,\ y,\ \theta,\ v)$ and $\dot{\mathbf{x}}_R = (v\ \cos\theta,\ v\ \sin\theta,\ \phi,\ a)$, where x, y, θ and v are the Cartesian position, the heading and the speed of the robot, respectively. The control vector $\mathbf{u}_R = (\phi,\ a)$ contains the angular velocity and the longitudinal acceleration respectively. We discretize the robot dynamics, using the RK4 method [31]. For all methods, the following limits are enforced: $v \in [0, 1.3]\frac{\text{m}}{\text{s}}$, $\phi \in [-0.5\pi, 0.5\pi]\frac{\text{rad}}{\text{s}}$, $a \in [-10, 10]\frac{\text{m}}{\text{s}^2}$. We use a planning horizon $T = 30$ steps with an integration timestep of $0.1\,\text{s}$. Q, W and S in Eq. 3 are diagonal matrices with entries $(1.5, 1.5, 0, 0)$, $(0.0005, 0.0005)$ and

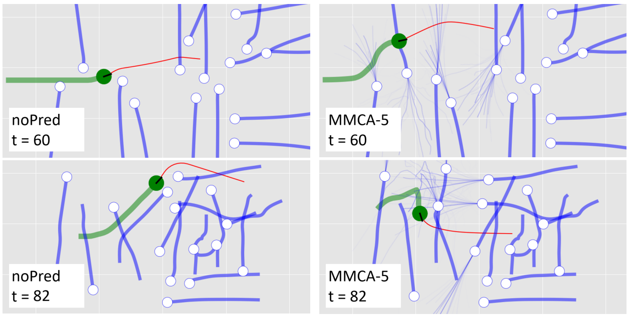

Fig. 3: Example frames from a simulation with human-robot interactions, showing the difference in robot behavior with (**right**) and without (**left**) using human motion prediction when planning. The rows show different timesteps of the simulation (**Top**: $t = 60$, **Bottom**: $t = 82$). As can be seen, planning without predictions leads to more risky and assertive robot behavior that pushes human from the original trajectories whereas anticipating human motion makes the robot avoid humans e.g. by actively getting behind them.

$(50, 50, 0, 0)$, respectively. As the different summands of the objective function are responsible for different behaviors, Q, W, S and g from Eq. 6, can be changed to influence the overall robot behavior.

We use Taylor expansion for propagating the uncertainty. If not explicitly noted otherwise, we consider only the 12 most probable modes of motion from the multi-modal prediction and compose the set I in Eq. 2 with up to 6 humans who are closest to the robot. As initial guess for the NLP we use the current state of the robot. The humans and robot are considered to have a circular footprint with radii of $0.2\,\mathrm{m}$ and $0.3\,\mathrm{m}$, respectively.

Human motion predictions are generated using one of the pre-trained, open-source, data-driven, state-of-the-art predictors Trajectron++ [29] and S-GAN [28], or the constant velocity model. The output format of the respective models was modified to provide mean, variance and the mixture weights of the respective predictions. If not stated otherwise, Trajectron++ [29] was used in the experiment. To simulate human motions and interactions, we use an open-source implementation of the social force model [32].

We adopt CasADi [33] to implement the NMPC controllers (with just-in-time compilation) and IPOPT [34] as numerical optimization back-end. All experiments run on a single core of a Xeon E5-1620 CPU and 32 GB of RAM. Methods and experiments are implemented in Python3.

B. Testing Scenarios

We run our experiments in two scenarios. Both run for a specific amount of time and the robot is given a random goal location upon successfully navigating to the latest one. The first scenario is a simulation inspired by the Edinburgh pedestrian dataset [35] seen in Figure 4a. We have replicated the original dataset in simulation by defining start and goal locations that create motion patterns that are consistent with the dataset. This simulation involves 50 humans moving between the target areas, as visualized in Figure 4b. The humans interact according to the social force model [32] and get a new goal assigned as soon as they reach their current one. In some experiments the robot also interacts with the humans. The scenario runs for 2000 seconds and allows us to evaluate and compare navigation capabilities in large, open environments, cluttered with dynamic obstacles.

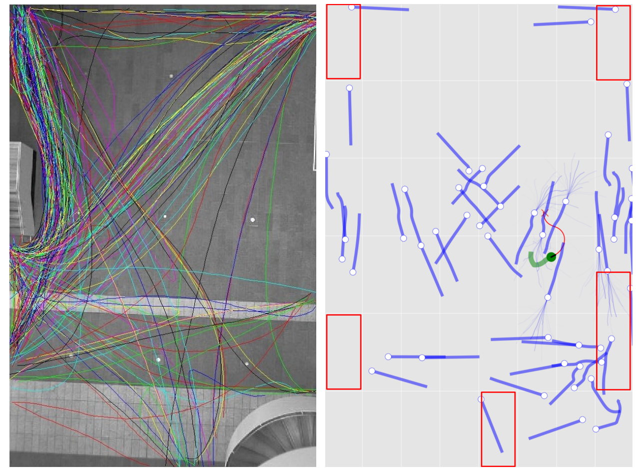

(a) (b)

Fig. 4: **(a)** Visualization of the Edinburgh pedestrian dataset [35], and **(b)** our analogous simulation. The goals of the humans in the simulation are randomly sampled within the **red** marked areas. The robot's goal is assigned randomly. The robot and its past path are shown in **green**. The humans, their past paths and motion predictions are shown in **blue**.

In the second scenario, we use the THÖR dataset [6] with three static obstacles (see Figure 1). We replay 1000 seconds of the recorded real-world human trajectories. Therefore, we do not rely on interactive motion simulation from the social force model. Importantly, this dataset was recorded in an environment with walls and polygonal obstacles. To account for these, we extend the NMPC formulation in Eq. 2 with dual variables and constraints according to Section 3-A of [36].[1] This scenario allows us to assess the navigation capabilities amongst real humans. Importantly it also lets us verify the results we obtain from the first scenario.

C. Metrics

Relevant metrics to assess the performance of the proposed approach include a combination of navigation/computational efficiency and safety indices. Efficiency indices are: n_{goals}, the number of goals reached by the robot in the experiment; $T_{g,\text{avg}}$ [s], average time the robot needs to reach a goal; $l_{p,\text{avg}}$ [m], average path length; $T_{\text{opt,avg}}$ [s], average optimization time for one NLP iteration; $T_{v=0}$ [%], time in which the robot is stopped. Safety indices include: $d_{h,\text{avg}}$ [m], average distance the robot keeps to the closest human; r_{feas} [%], iterations for which a feasible solution is available; T_{col} [%], time in which the robot is in collision with a human. Metrics labeled with [%] are given as percentages of the total scenario duration.

D. Baselines

We compare our cost-based approach with different gains g in Eq. 6 (denoted as MMCA-*gain*) to the following state-of-the-art methods for local motion planning and control:

[1]For space reasons we do not write out the full modifications and refer the interested reader to the article.

Method	$n_{\text{goals}} \uparrow$	$T_{g,\text{avg}}$ $[s] \downarrow$	$l_{p,\text{avg}}$ $[m] \downarrow$	$T_{v=0}$ $[\%] \downarrow$	$d_{h,\text{avg}}$ $[m] \uparrow$	r_{feas} $[\%] \uparrow$	T_{col} $[\%] \downarrow$
noPred	234	8.48	8.10	22.0	1.43	91.3	6.00
Single-MCA-5	204	9.79	10.45	13.8	2.02	98.8	0.80
Single-MCA-10	122	16.27	17.44	12.0	2.34	99.4	0.45
MMCA-5	**204**	**9.78**	**11.03**	**9.3**	**2.07**	**99.6**	**0.21**
MMCA-10	**127**	**15.67**	**17.06**	**10.9**	**2.53**	**99.7**	**0.21**
MMCA-15	**106**	**18.80**	**20.03**	**11.7**	**3.39**	**99.9**	**0.03**
CC	120	16.52	8.93	51.0	1.47	51.2	4.76
MultiCC	39	45.74	7.84	82.6	1.35	17.7	5.72
GA3C-CADRL	194	10.27	11.88	7.9	2.04	n.a.	4.46
DWA	27	66.91	47.50	18.2	1.82	n.a.	3.54

TABLE I: Performance of **our MMCA** approach (in bold) against other state-of-the-art controllers in an invisible robot experiment. Arrows indicate for which columns a higher (↑) or lower (↓) value is better.

Method	$n_{\text{goals}} \uparrow$	$T_{g,\text{avg}}$ $[s] \downarrow$	$l_{p,\text{avg}}$ $[m] \downarrow$	$T_{v=0}$ $[\%] \downarrow$	$d_{h,\text{avg}}$ $[m] \uparrow$	r_{feas} $[\%] \uparrow$
noPred	253	7.90	8.39	14.0	1.38	95.8
Single-MCA-5	190	10.42	11.48	10.7	1.99	99.9
Single-MCA-10	132	15.10	16.27	12.0	2.39	100.0
MMCA-5	**219**	**9.13**	**10.19**	**9.6**	**2.04**	**100.0**
MMCA-10	**163**	**12.13**	**13.03**	**11.8**	**2.29**	**100.0**
MMCA-15	**82**	**23.93**	**25.48**	**12.5**	**3.10**	**100.0**
CC	110	18.17	8.86	55.2	1.40	47.1
GA3C-CADRL	179	11.17	13.75	2.2	2.58	n.a.
DWA	20	99.53	71.87	17.7	2.27	n.a.

TABLE II: Performance of **our MMCA** approach (in bold) against other state-of-the-art controllers in a human-robot interaction experiment.

Prediction	**Control**	$n_{\text{goals}} \uparrow$	$d_{h,\text{avg}}$ $[m] \uparrow$	T_{col} $[\% \downarrow$
CVM	Single-MCA-5	189	2.05	0.40
Trajectron++	Single-MCA-5	204	2.02	0.80
Trajectron++	MMCA-5	204	2.07	0.21
S-Gan	MMCA-5	201	2.05	0.18

TABLE III: Comparison of our cost-based method using different human motion predictors in an invisible robot experiment.

No. humans	$n_{\text{goals}} \uparrow$	$d_{h,\text{avg}}$ $[m] \uparrow$	$T_{\text{opt,avg}}$ $[s] \downarrow$
4	218	1.99	0.020
6	205	2.05	0.026
8	195	2.06	0.031
10	198	2.17	0.034

TABLE IV: Performance of MMCA-5 when varying the number of humans in the invisible robot experiment.

Method	$n_{\text{goals}} \uparrow$	$d_{h,\text{avg}}$ $[m] \uparrow$	T_{col} $[\%] \downarrow$
noPred	169	2.47	4.98
Single-MCA-5	123	2.94	1.66
MMCA-5	**137**	**2.91**	**1.08**
MMCA-10	**112**	**3.21**	**0.61**
MMCA-15	**89**	**3.56**	**0.98**
CC	99	2.50	4.34

TABLE V: Experiment with 1000 seconds of data replayed from the THÖR dataset. Results of **our MMCA** cost are in bold.

a deep reinforcement learning approach based on an actor-critic architecture (GA3C-CADRL) [21]; the Dynamic Window Approach (DWA) [13]; a uni-modal chance constraint formulation (CC) [17] with 3% risk of collision; and a multi-modal version of chance constraint NMPC (MultiCC, as in Nair et al. [11]) which constrain on every mode of the prediction as if it were an individual trajectory. We choose these baselines so we can appropriately compare our approach with navigation methods from different domains. To investigate the influence that multi-modal predictions have on the collision avoidance capabilities in the first place, we include two more approaches in our comparisons: An NMPC controller that uses no predictions, by considering the humans as static obstacles (noPred), only imposing collision avoidance constraints for the current position of the humans; and a uni-modal version of our approach (denoted as Single-MCA) by only considering the most likely predicted mode in the cost formulation.

VI. Results and Discussion

We conduct 5 experiments to investigate different properties and compare our approach to the baselines in different ways. The first four experiments are performed in the first scenario described in Section V-B (i.e. Edinburgh simulated data) and the fifth experiment is done in the second scenario (i.e. real-world THÖR data).

A. Experiments and Results

1) Invisible Robot: The first experiment is an invisible robot experiment. That is, the robot is aware of the humans locations and will try to avoid them, but not vice-versa. Thus collisions can occur if the robot is not avoiding the humans. Results are reported in Table I.

2) Human Robot Interaction: In the second experiment, the robot is visible to the humans and thus humans will actively try to keep a minimum safe distance to the robot. Thus, collisions happen very rarely (none were observed in our experiment). This makes it harder to evaluate the collision avoidance performance of the methods but nevertheless provides insights into their capabilities to navigate amongst humans. Results are presented in Table II.

3) Different Predictors: In the third experiment, we investigate the influence of different prediction methods on collision avoidance in an invisible robot experiment. For this, we run the experiment with Trajectron++ [29], S-Gan [28], and the constant velocity model (CVM). We also investigate the effects of using only the most likely mode of motion from Trajectron++ (Single-MCA-5). Results are given in Table III.

4) NMPC Efficiency: In the fourth experiment, we study optimization time for one NMPC iteration, which translates to the runtime of the approach, in an invisible robot experiment. As explained in Section V, we only consider the n closest humans in the NLP. This experiment evaluates our method for different values of n to see how this number influences optimization time and safety (results in Table IV).

5) THÖR Dataset: The fifth experiment evaluates and compares our method in the second scenario, using the THÖR dataset [6] with real-world human trajectories. As we are replaying recorded data, this is also an invisible robot experiment. Results are given in Table V.

B. Discussion

In the following, we discuss the performance of our method compared to the baselines described in Section V-D.

MMCA-NMPC: Our method shows a notable improvement over the baselines. Table I shows that we are able to notably reduce the time the robot spends in collision with humans to 0.03% in the best case. At worst, our method improves the time in collision by a factor of 20, from around 4.5% for the baselines, to 0.21%. A qualitative evaluation show that these collisions occur if humans change directions unexpectedly. Then, it is possible that the robot cannot avoid a collision in time, since it follows more restrictive kinematics than the humans. The overall findings confirm the hypotheses we outline in the beginning of Section IV. MMCA-NMPC competes with noPred and GA3C-CADRL in terms of task efficiency, even reducing the time the robot is standing still or being frozen. The use of g in Eq. 6 works as intended, improving safety and increasing the average distance to the humans (3.39 m for MMCA-15 vs. 2.07 m for MMCA-5), at the expense of less efficient navigation (108 goals reached for MMCA-15 vs. 204 for MMCA-5). Our approach arguably performs best overall in the human-robot interaction experiment (see Table II), considering its good performance in all metrics. The results of Single-MCA NMPC suggest that, while the use of multi-modal predictions seems to benefit performance and safety, the choice of not defining constraints on predictions improves performance notably over the other approaches. However, during testing we find that constraints on the current human position are still helpfull to prevent the robot from colliding in very crowded and ambiguous situations. Table III shows that our method is agnostic to the type of predictor that is used. The improvements gained from using multi- over single-modal predictions are not big in the measure of absolute improvement. For this chain of arguments it is important to point out that the single-mode prediction is actually the most likely sub-component of the multi-modal one. Additionally, the predictions made by Trajectron++ or S-Gan are not goal driven, and span most of the space in front of the predicted human. We believe that this is a clear limitation of many pre-trained open-source available prediction methods that have not been explicitly designed for a robot navigation use-case.

The NLP formulation in Eq. 2 is solved, on average, in 0.02–0.04 s (Table IV). The number of humans considered influences the computational effort to solve the NLP, and the trade-off between n_{goals} and $d_{h,\mathrm{avg}}$ as expected.

The experiment on the THÖR dataset only involves NMPC-based approaches, as the others do not have a straightforward way of considering polygonal obstacle. This experiment confirms the results from the other experiments in a more complex and cluttered environment. The higher values for T_{col} can be attributed to the more restrictive environment, making it harder to avoid collisions, and the fact that humans "appear" in the dataset which cannot be foreseen by the prediction. The experiment shows that our approach works in real scenarios with recorded human motion data, and can be used in synergy with other NMPC based trajectory planning methods like static obstacle avoidance.

CC-NMPC: Chance-Constrained NMPC seems to struggle in 2D cluttered dynamic environments, as it is only able to find a feasible solution in about 50% of iterations ($r_{\mathrm{feas}} = 51.2\%$ in Table I). This problem is reflected in the other metrics as well. The accepted risk of collision is set to 3%, as suggested by Zhu et al. [10], which corresponds to the time in collision we observe in our experiments ($T_{\mathrm{col}} = 4.95\%$ in Table I). The increase in collisions can be attributed to the mismatch between the prediction and the actual trajectory, as well as infeasible iterations for which the robot is receiving decelerate-to-stop controls.

MultiCC-NMPC: The MultiCC approach has two major problems and is only investigated in the first experiment. The approach is rarely able to find a valid solution, leaving the robot frozen for most of the time ($T_{v=0} = 82.6\%$ in Table I). Moreover, the time it takes to solve one iteration of the NMPC is around 1.2 seconds. This is most likely due to the number of constraints added to the NLP, which makes it harder to find a feasible solution.

DWA: DWA performs poorly in the efficiency metrics (n_{goals}, $T_{g,\mathrm{avg}}$, $l_{p,\mathrm{avg}}$), but performs well in terms of collision avoidance performance ($T_{\mathrm{col}} = 3.45\%$ in Table I). The results indicate its inability to plan more complex local trajectories. In challenging situations, DWA is not able to find local trajectories that avoid collisions with all surrounding humans, while also moving closer to the goal state.

noPred-NMPC: The NMPC controller that does not use predictions (noPred) navigates efficiently as it works with no information besides its goal state and the current position of the humans. However, different from our approach, the time in collision, as well as the average distance to humans, show that the robot is much more prone to collisions (a factor of 200 for T_{col} in the worst case, see Table I). Also, when evaluating the behavior qualitatively, the approach shows undesirable, assertive behavior and does not anticipate human movement. Figure 3 shows how the trajectories generated by the noPred-NMPC compare to the ones from MMCA-NMPC in an identical scenario with human-robot interaction. To reach its goal, noPred collides and pushes a human out of its original trajectory.

GA3C-CADRL: For GA3C-CADRL we observe similar results as reported in [21]. Its high time in collision compares to the one from CC or noPred while the distance to humans is notably higher. The similarities to noPred make sense as this approach also does not explicitly consider human motion prediction. This means the method has to consider future human movement implicitly when generating controls. Importantly, GA3C-CADRL cannot be infeasible and therefore stays maneuverable even in difficult situations. Even though the efficiency is comparable to MMCA-5, our method is still notably better at avoiding collisions (0.21% collision time for MMCA-5 vs. 4.46% collision time for GA3C-CADRL).

VII. Conclusions

In this paper, we have introduced a novel Multi-Modal Collision Avoidance method for NMPC-based motion planning in cluttered dynamic environments. Our novel formulation uses stochastic multi-modal human motion predictions to achieve efficient and smooth trajectory planning while actively and reliably avoiding collisions. Our approach significantly outperforms state-of-the-art methods in realistic cluttered dynamic environments. The experiments show that our approach achieves at least 20 times better collision avoidance while being as least as efficient as the compared baselines. The experiments confirm that the method can solve the underlying NLP fast enough for online use and easily integrates with different human motion predictors. In future work, we aim to incorporate intention-driven prediction methods and explore the applicability of our approach with different parameters and robot kinematics (e.g., higher maximum velocity). We also want to systematically evaluate our approach in real-world experiments with human-robot interaction to better analyze its legibility. In doing so we plan to release the approach as open source software to integrate with common robotics and simulation frameworks.

References

[1] S. Bansal, A. Bajcsy, E. Ratner, A. D. Dragan, and C. J. Tomlin, "A hamilton-jacobi reachability-based framework for predicting and analyzing human motion for safe planning," in *IEEE Int. Conf. Robot. & Autom. (ICRA)*. IEEE, 2020.

[2] N. Malone, H.-T. Chiang, K. Lesser, M. Oishi, and L. Tapia, "Hybrid dynamic moving obstacle avoidance using a stochastic reachable set-based potential field," *IEEE Trans. on Robotics (TRO)*, 2017.

[3] A. Pereira and M. Althoff, "Safety control of robots under computed torque control using reachable sets," in *IEEE Int. Conf. Robot. & Autom. (ICRA)*. IEEE, 2015.

[4] S. Kousik, S. Vaskov, F. Bu, M. Johnson-Roberson, and R. Vasudevan, "Bridging the gap between safety and real-time performance in receding-horizon trajectory design for mobile robots," *Int. J. Robot. Research*, vol. 39, no. 12, 2020.

[5] K. Leung, E. Schmerling, M. Chen, J. Talbot, J. C. Gerdes, and M. Pavone, "On infusing reachability-based safety assurance within probabilistic planning frameworks for human-robot vehicle interactions," in *Int. Symp. Exp. Robot.* Springer, 2018.

[6] A. Rudenko, T. P. Kucner, C. S. Swaminathan, R. T. Chadalavada, K. O. Arras, and A. J. Lilienthal, "Thör: Human-robot navigation data collection and accurate motion trajectories dataset," *IEEE Robot. & Autom. Letters*, vol. 5, no. 2, pp. 676–682, 2020.

[7] T. Schoels, L. Palmieri, K. O. Arras, and M. Diehl, "An NMPC approach using convex inner approximations for online motion planning with guaranteed collision avoidance," in *Int. Conf. Robot. & Autom. (ICRA)*. IEEE, 2020.

[8] S. Schaefer, K. Leung, B. Ivanovic, and M. Pavone, "Leveraging neural network gradients within trajectory optimization for proactive human-robot interactions," in *Int. Conf. Robot. & Autom. (ICRA)*. IEEE, 2021.

[9] C. Rösmann, A. Makarow, and T. Bertram, "Online motion planning based on nonlinear model predictive control with non-euclidean rotation groups," in *Europ. Contr. Conf. (ECC)*, 2021, pp. 1583–1590.

[10] H. Zhu and J. Alonso-Mora, "Chance-constrained collision avoidance for mavs in dynamic environments," *IEEE Robot. & Autom. Letters*, vol. 4, no. 2, 2019.

[11] S. H. Nair, V. Govindarajan, T. Lin, C. Meissen, H. E. Tseng, and F. Borrelli, "Stochastic mpc with multi-modal predictions for traffic intersections," in *2022 IEEE 25th Int. Conf. on Intelligent Transportation Systems (ITSC)*, 2022, pp. 635–640.

[12] A. Rudenko, L. Palmieri, M. Herman, K. M. Kitani, D. M. Gavrila, and K. O. Arras, "Human motion trajectory prediction: a survey," *Int. Journal of Robotics Research*, vol. 39, no. 8, 2020.

[13] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance," *IEEE Robot. & Autom. Mag.*, vol. 4, no. 1, pp. 23–33, 1997.

[14] S. Singh, A. Majumdar, J.-J. Slotine, and M. Pavone, "Robust online motion planning via contraction theory and convex optimization," in *Int. Conf. Robot. & Autom. (ICRA)*. IEEE, 2017.

[15] T. Schoels, P. Rutquist, L. Palmieri, A. Zanelli, K. O. Arras, and M. Diehl, "Ciao∗: Mpc-based safe motion planning in predictable dynamic environments," *IFAC-PapersOnLine*, vol. 53, no. 2, 2020.

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

[17] M. Castillo-Lopez, P. Ludivig, S. A. Sajadi-Alamdari, J. L. Sanchez-Lopez, M. A. Olivares-Mendez, and H. Voos, "A real-time approach for chance-constrained motion planning with dynamic obstacles," *IEEE Robot. & Autom. Letters*, vol. 5, no. 2, 2020.

[18] E. Schmerling, K. Leung, W. Vollprecht, and M. Pavone, "Multimodal probabilistic model-based planning for human-robot interaction," in *Int. Conf. Robot. & Autom. (ICRA)*. IEEE, 2018.

[19] Y. Chen, U. Rosolia, C. Fan, A. Ames, and R. M. Murray, "Reactive motion planning with probabilistic safety guarantees," in *Conf. Robot Learn. (CoRL)*, vol. 155, 2020, pp. 1958–1970.

[20] M. Kamel, J. Alonso-Mora, R. Siegwart, and J. Nieto, "Robust collision avoidance for multiple micro aerial vehicles using nonlinear model predictive control," in *IEEE/RSJ Int. Conf. Intell. Robot. Sys. (IROS)*, 2017, pp. 236–243.

[21] M. Everett, Y. F. Chen, and J. P. How, "Collision avoidance in pedestrian-rich environments with deep reinforcement learning," *IEEE Access*, vol. 9, pp. 10 357–10 377, 2021.

[22] Y. F. Chen, M. Liu, M. Everett, and J. P. How, "Decentralized non-communicating multiagent collision avoidance with deep reinforcement learning," in *Int. Conf. Robot. & Autom. (ICRA)*. IEEE, 2017.

[23] J. B. Rawlings, D. Q. Mayne, and M. Diehl, *Model predictive control: theory, computation, and design*. Nob Hill Publishing Madison, WI, 2017, vol. 2.

[24] H. Zhu, F. M. Claramunt, B. Brito, and J. Alonso-Mora, "Learning interaction-aware trajectory predictions for decentralized multi-robot motion planning in dynamic environments," *IEEE Robot. & Autom. Letters*, vol. 6, no. 2, pp. 2256–2263, 2021.

[25] J. Park and H. J. Kim, "Online trajectory planning for multiple quadrotors in dynamic environments using relative safe flight corridor," *IEEE Robot. & Autom. Letters*, vol. 6, no. 2, pp. 659–666, 2021.

[26] P. Dendorfer, A. Osep, and L. Leal-Taixé, "Goal-gan: Multimodal trajectory prediction based on goal position estimation," in *Proc. Asian Conf. Computer Vision*, 2020.

[27] Y. Yao, E. Atkins, M. Johnson-Roberson, R. Vasudevan, and X. Du, "Bitrap: Bi-directional pedestrian trajectory prediction with multi-modal goal estimation," *IEEE Robot. & Autom. Letters*, vol. 6, no. 2, 2021.

[28] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi, "Social gan: Socially acceptable trajectories with generative adversarial networks," in *Proc. of the IEEE Conf. on Comp. Vis. and Pat. Rec. (CVPR)*, 2018.

[29] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, "Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data," in *Proc. of the European Conf. on Computer Vision (ECCV)*. Springer, 2020.

[30] A. Wang, A. Jasour, and B. C. Williams, "Non-Gaussian chance-constrained trajectory planning for autonomous vehicles under agent uncertainty," *IEEE Robot. & Autom. Letters*, vol. 5, no. 4, pp. 6041–6048, 2020.

[31] N. S. Bakhvalov, *Numerical methods: analysis, algebra, ordinary differential equations*. Mir Publishers, 1977.

[32] D. Helbing and P. Molnar, "Social force model for pedestrian dynamics," *Physical review E*, vol. 51, no. 5, 1995.

[33] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "CasADi – A software framework for nonlinear optimization and optimal control," *Mathematical Programming Computation*, vol. 11, 2019.

[34] A. Wachter, "An interior point algorithm for large-scale nonlinear optimization with applications in process engineering," Ph.D. dissertation, Carnegie Mellon University, 2002.

[35] B. Majecka, "Statistical models of pedestrian behaviour in the forum," Master's thesis, School of Informatics, University of Edinburgh, 2009.

[36] X. Zhang, A. Liniger, and F. Borrelli, "Optimization-based collision avoidance," *IEEE Trans. Contr. Sys. Tech.*, vol. 29, no. 3, pp. 972–983, 2021.