Delft University of Technology

Model predictive contouring control for collision avoidance in unstructured dynamic environments

Brito, Bruno; Floor, Boaz; Ferranti, Laura; Alonso-Mora, Javier

DOI

[10.1109/LRA.2019.2929976](https://doi.org/10.1109/LRA.2019.2929976)

Publication date

2019

Document Version

Final published version

Published in

IEEE Robotics and Automation Letters

Citation (APA)

Brito, B., Floor, B., Ferranti, L., & Alonso-Mora, J. (2019). Model predictive contouring control for collision avoidance in unstructured dynamic environments. *IEEE Robotics and Automation Letters*, 4(4), 4459-4466. <https://doi.org/10.1109/LRA.2019.2929976>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

Model Predictive Contouring Control for Collision Avoidance in Unstructured Dynamic Environments

Bruno Brito , Boaz Floor , Laura Ferranti , and Javier Alonso-Mora

Abstract—This letter presents a method for local motion planning in unstructured environments with static and moving obstacles, such as humans. Given a reference path and speed, our optimization-based receding-horizon approach computes a local trajectory that minimizes the tracking error while avoiding obstacles. We build on nonlinear model-predictive contouring control (MPCC) and extend it to incorporate a static map by computing, online, a set of convex regions in free space. We model moving obstacles as ellipsoids and provide a correct bound to approximate the collision region, given by the Minkowsky sum of an ellipse and a circle. Our framework is agnostic to the robot model. We present experimental results with a mobile robot navigating in indoor environments populated with humans. Our method is executed fully onboard without the need of external support and can be applied to other robot morphologies such as autonomous cars.

Index Terms—Collision avoidance, motion and path planning.

I. INTRODUCTION

APPLICATIONS where autonomous ground vehicles (AGVs) closely navigate with humans in complex environments require the AGV to safely avoid static and moving obstacles while making progress towards its goal. Motion planning and control for AGVs are typically addressed as two independent problems [1]. In particular, the motion planner generates a collision-free path and the motion controller tracks such a path by directly commanding the AGV's actuators. Our method combines both motion planning and control in one module, relying on constrained optimisation techniques, to generate kinematically feasible local trajectories with fast replanning cycles. More specifically, we rely on a Model Predictive Controller (MPC) to compute an optimal control command for the controlled system, which directly incorporates the predicted intentions of dynamic obstacles. Consequently, it reacts in advance to smoothly avoid moving obstacles. We propose a practical reformulation of the Model Predictive Contouring Control (MPCC) approach [2], namely, a Local Model Predictive

Fig. 1. The faculty corridor was the scenario used to evaluate the capabilities of our method to avoid dynamic obstacles.

Contouring Control (LMPCC) approach, suitable for real-time collision-free navigation of AGVs in complex environments with several agents. Our design is fully implemented on board of the robot including localization and environment perception, i.e., detection of static obstacles and pedestrians. Our design runs in real time thanks to its lightweight implementation. The method relies on an open-source solver [3] and will be released with this letter as a ROS module. Our method can be adapted to other robot morphologies, such as cars.

A. Related Work

Collision avoidance in static and dynamic environments can be achieved via reactive methods, such as time-varying artificial potential fields [4], the dynamic window [5], social forces [6] and velocity obstacles [7]. Although these approaches work well in low speed scenarios, or scenarios of low complexity, they produce highly reactive behaviours. More complex and predictive behaviour can be achieved by employing a motion planner. Our method relies on model predictive control (MPC) [8], [9] to obtain smooth collision-free trajectories that optimize a desired performance index, incorporate the physical constraints of the robot and the predicted behavior of the obstacles.

Due to the complexity of the motion planning problem, path planning and path following were usually considered as two separate problems. Many applications of MPC for path-following control are found in the literature, e.g., [10], [11]. These methods assume the availability of a collision-free smooth path to follow. In contrast, we use MPC for local motion planning and control in the presence of static and dynamic obstacles.

Several approaches exist for integrated path following and control for dynamic environments. These include: employing a

Manuscript received February 24, 2019; accepted June 26, 2019. Date of publication July 22, 2019; date of current version October 24, 2019. This letter was recommended for publication by Associate Editor J.-M. Lien and Editor N. Amato upon evaluation of the reviewers' comments. This work was supported by the Amsterdam Institute for Advanced Metropolitan Solutions and the Netherlands Organisation for Scientific Research (NWO) domain Applied Sciences. (Corresponding author: Bruno Brito.)

The authors are with the Department of Cognitive Robotics, Delft University of Technology, 2628 CD Delft, The Netherlands (e-mail: brunofbrito0@gmail.com; boazfloor@gmail.com; l.ferranti@tudelft.nl; j.alonsomora@tudelft.nl).

This letter has supplementary downloadable material available at <http://ieeexplore.ieee.org>, provided by the authors.

Digital Object Identifier 10.1109/LRA.2019.2929976

set of motion primitives and optimizing the control commands to execute them [12] and offline computation of tracking error bounds via reachability analysis that ensures a safety region around the robot for online planning [13], [14]. These approaches, however, do not allow to incorporate the predicted intentions of the dynamic obstacles and consequently, can lead to reactive behaviors. To overcome the latter issue, [15] proposed the model predictive contouring control (MPCC) that allows to explicitly penalize the deviation from the path (in terms of contouring and lateral errors) and include additional constraints. Later, [16] designed a MPC for path following within a stable handling envelope and an environmental envelope. While [16] is tailored to automotive applications without dynamic obstacles, MPCC has been employed to handle static and dynamic obstacles in structured driving scenarios [2]. There, static collision constraints were formulated as limits on the reference path and thus limited to on-the-road driving scenarios. The previous approaches do not account for the interaction effects between the agents and may fail in crowded scenarios, a problem known as the freezing robot problem (FRP) [17]. Interaction Gaussian Processes (IGP) [18] can be used to model each individual's path. The interactions are modeled with a nonlinear potential function. The resulting distribution, however, is intractable, and sampling processes are required to approximate a solution, which requires high computational power and is only real-time for a limited number of agents. Learning-based approaches address this issue by learning the collision-avoidance strategy directly from offline simulation data [19], or the complex interaction model from raw sensor readings [20]. Yet, both methods learn a reactive collision avoidance policy and do not account for the kinodynamic constraints of the robot.

Without explicitly modeling interaction, we propose a method for local motion planning, which is real-time, incorporates the robot constraints and optimizes over a prediction horizon. Our approach relies on MPCC and extends it to mobile robots operating in unstructured environments. Similar to [21], among others, we employ polyhedral approximations of the free space, which can provide larger convex regions than safety bubbles [22].

B. Contribution

We build on [2], with the following contributions to make the design applicable to mobile robots navigating in unstructured environments with humans:

- A static obstacle avoidance strategy that explicitly constrains the robot's positions along the prediction horizon to a polyhedral approximation of the collision-free area around the robot.
- A closed-form bound to conservatively approximate collision avoidance constraints that arise from ellipsoidal moving obstacles.
- A fully integrated MPCC approach that runs in real-time on-board of the robot and with on-board perception.

We present experimental results with a mobile robot navigating in indoor environments among static and moving obstacles and compare them with three state-of-art planners, namely the dynamic window [5], a classical MPC for tracking [11] and

a socially-aware motion planner [23]. Finally, to illustrate the generality of the proposed MPCC framework, we present results with a simulated car.

II. PRELIMINARIES

A. Robot Description

Let B denote an AGV on the plane $\mathcal{W} = \mathbb{R}^2$. The AGV dynamics are described by the discrete-time nonlinear system

$$\mathbf{z}(t+1) = f(\mathbf{z}(t), \mathbf{u}(t)), \quad (1)$$

where $\mathbf{z}(t)$ and $\mathbf{u}(t)$ are the state and the input of the robot, respectively, at time $t \geq 0$.¹ For the case of our mobile robot we consider the state to be equal to the configuration, that is, $\mathbf{z}(t) \in \mathcal{C} = \mathbb{R}^2 \times \mathcal{S}$. For the case of a car (Section IV-E), the state space includes the speed of the car. The area occupied by the robot at state $\mathbf{z}$ is denoted by $\mathcal{B}(\mathbf{z})$, which is approximated by a union of n_c circles, i.e., $\mathcal{B}(\mathbf{z}) \subseteq \bigcup_{c \in \{1, \dots, n_c\}} \mathcal{B}_c(\mathbf{z}) \subset \mathcal{W}$. The center of each circle, in the inertial frame, is given by $\mathbf{p} + R_B^W(\mathbf{z})\mathbf{p}_c^B$. Where $\mathbf{p}$ is the position of the robot (extracted from $\mathbf{z}$), $R_B^W(\mathbf{z})$ is the rotation matrix given by the orientation of the robot, and $\mathbf{p}_c^B$ is the center of circle c expressed in the body frame.

B. Static Obstacles

The static obstacle environment is captured in an occupancy grid map, where the area occupied by the static obstacles is denoted by $\mathcal{O}^{\text{static}} \subset \mathcal{W}$. In our experiments we consider both a global map, which is built a priori and used primarily for localization, and a local map from the current sensor readings. Therefore, the static map is continuously updated locally. Dynamic obstacles, such as people, that are recognized and tracked by the robot are removed from the static map and considered as moving obstacles.

C. Dynamic Obstacles

Each moving obstacle i is represented by an ellipse of area $\mathcal{A}_i \subset \mathcal{W}$, defined by its semi-major axis a_i , its semi-minor axes b_i , and a rotation matrix $R_i(\psi)$. We consider a set of moving obstacles $i \in \mathcal{I} := \{1, \dots, n\}$, where n can vary over time. The area occupied by all moving obstacles at time instant t is given by $\mathcal{O}_t^{\text{dyn}} = \bigcup_{i \in \{1, \dots, n\}} \mathcal{A}_i(\mathbf{z}_i(t))$, where $\mathbf{z}_i(t)$ denotes the state of moving obstacle i at time t . In this work, and without a loss of generality, we assume a constant velocity model with Gaussian noise $\omega_o(t) \sim \mathcal{M}(0, Q_o(t))$ in acceleration, that is, $\ddot{\mathbf{p}}_i(t) = \omega_o(t)$, where $\mathbf{p}_i(t)$ is the position of obstacle i at time t . Given the measured position data of each obstacle, we estimate their future positions and uncertainties with a linear Kalman filter [24].

D. Global Reference Path

Consider that a reference path is available. In its simplest form, the reference path can be a straight line to the goal

¹In the remainder of the letter we omit the time dependency when it is clear from the context.

position or a straight line in the direction of preferred motion. But it could also be given by a global planner. We consider a global reference path $\mathcal{P}$ consisting of a sequence of path segments connecting M way-points $\mathbf{p}_m^r = [x_m^p, y_m^p] \in \mathcal{W}$ with $m \in \mathcal{M} := \{1, \dots, M\}$. For smoothness, we consider that each path segment $\varsigma_m(\theta)$ is defined by a cubic polynomial. We denote by θ a variable that (approximately) represents the traveled distance along the reference path, and which is described in more detail in Section III-C3. We do not require the reference to be collision free, therefore, the robot may have to deviate from it to avoid collisions.

E. Problem Formulation

The objective is to generate, for the robot, a collision free motion for N time-steps in the future, while minimizing a cost function J that includes a penalty for deviations from the reference path. This is formulated in the optimization problem

$$\begin{aligned} J^* = & \min_{\mathbf{z}_{0:N}, \mathbf{u}_{0:N-1}, \theta_{0:N-1}} \sum_{k=0}^{N-1} J(\mathbf{z}_k, \mathbf{u}_k, \theta_k) + J(\mathbf{z}_N, \theta_N) \\ \text{s.t. } & \mathbf{z}_{k+1} = f(\mathbf{z}_k, \mathbf{u}_k), \theta_{k+1} = \theta_k + v_k \tau, \\ & \mathcal{B}(\mathbf{z}_k) \cap (\mathcal{O}^{\text{static}} \cup \mathcal{O}_k^{\text{dyn}}) = \emptyset, \\ & \mathbf{u}_k \in \mathcal{U}, \mathbf{z}_k \in \mathcal{Z}, \mathbf{z}_0, \theta_0 \text{ given.} \end{aligned} \quad (2a) \quad (2b) \quad (2c)$$

where v_k is the forward velocity of the robot (for the mobile robot it is part of the input and for the car it is part of the state), τ is the time-step and $\mathcal{Z}$ and $\mathcal{U}$ are the set of admissible states and inputs, respectively. $\mathbf{z}_{1:N}$ and $\mathbf{u}_{0:N-1}$ are the set of states and control inputs, respectively, over the prediction horizon T_{horizon} , which is divided into N prediction steps. θ_k denotes the predicted progress along the reference path at time-step k . By solving the optimization problem, we obtain a locally optimal sequence of commands $[\mathbf{u}_t^*]_{t=0}^{t=N-1}$ to guide the robot along the reference path while avoiding collisions with static and moving obstacles.

III. METHOD

The proposed method consists of the following steps, which are executed in every planning loop.

- 1) Search for a collision-free region in the updated static map centered on the robot and constrain the control problem such that the robot remains inside (Section III-A).
- 2) Predict the future positions of the dynamic obstacles and use the corrected bound to ensure dynamic collision avoidance (Section III-B).
- 3) Solve a modified MPCC formulation applicable to mobile robots (Section III-C).

A. Static Collision Avoidance

Given the static map of the environment, we compute a set of convex four-sided polygons in free space. This representation can provide larger collision-free areas compared with other approximations such as circles.

To obtain the set of convex regions at time t , we first shift the optimal trajectory computed at time $t-1$, namely, $\mathbf{q}_{0:N} =$

Fig. 2. Representation of the convex free space (orange squares) around each prediction step on the prediction horizon (purple line) with respect to the inflated static environment and the collision space of the dynamic obstacle (green ellipses) with respect to the vehicle representing discs (blue).

$[\mathbf{p}_{1:N|t-1}^*, \mathbf{q}_N]$, where $\mathbf{q}_N$ is an extrapolation of the last two points, that is, $\mathbf{q}_N = 2\mathbf{p}_{N|t-1}^* - \mathbf{p}_{N-1|t-1}^*$. Then, for each point $\mathbf{q}_k$ ($k = 1, \dots, N$) we compute a convex region in free space, given by a set of four linear constraints $c_k^{\text{stat}}(\mathbf{p}_k) = \bigcup_{l=1}^4 c_k^{\text{stat},l}(\mathbf{p}_k)$. This region separates $\mathbf{q}_k$ from the closest obstacles. In our implementation we compute a rectangular region aligned with the orientation of the trajectory at $\mathbf{q}_k$, where each linear constraint is obtained by a search routine and reduced by the radius of the robot circles r_{disc} . Fig. 2 shows the collision-free regions along the prediction horizon defined as yellow boxes. At prediction step k for a robot with state $\mathbf{z}$ the resulting constraint for disc j and polygon side l is

$$c^{\text{stat},l,j}(\mathbf{z}) = h^l - \tilde{n}^l \cdot (\mathbf{p} - R_B^W(\mathbf{z})\mathbf{p}_j^B) > 0, \quad (3)$$

where h^l and $\tilde{n}^l$ define each side of the polygons.

B. Dynamic Collision Avoidance

Recall that each moving obstacle i is represented by its position $\mathbf{p}_i(t)$ and an ellipse of semi-axis a_i and b_i and a rotation matrix $R_i(\psi)$. For each obstacle $i \in \{1, \dots, n\}$, and prediction step k , we impose that each circle j of the robot does not intersect with the elliptical area occupied by the obstacle. Omitting i for simplicity, the inequality constraint on each disc of the robot with respect to the obstacles is

$$c_k^{\text{obst},j}(\mathbf{z}_k) = \begin{bmatrix} \Delta x_k^j \\ \Delta y_k^j \end{bmatrix}^T R(\psi)^T \begin{bmatrix} \frac{1}{\alpha^2} & 0 \\ 0 & \frac{1}{\beta^2} \end{bmatrix} R(\psi) \begin{bmatrix} \Delta x_k^j \\ \Delta y_k^j \end{bmatrix} > 1, \quad (4)$$

where the distance between disc j and the obstacle is separated into its Δx_k^j and Δy_k^j components (Fig. 2). The parameters α and β are the semi-axes of an enlarged ellipse that includes the union of the original ellipse and the circle.

While previous approaches approximated the Minkowski sum of the ellipse with the circle as an ellipse of semi-major $\alpha = a + r_{\text{disc}}$ and semi-minor axis $\beta = b + r_{\text{disc}}$ [2], this assumption is not correct and collisions can still occur [25]. We now describe how to compute the values for α and β such that collision avoidance is guaranteed, represented by the larger in light red ellipsoid in Fig. 3.

Fig. 3. Safety boundary computation. The ellipsoid with semi-major axis and semi-minor axis, a and b respectively, is represented in green, the ellipsoid with axis enlarged by r_{disc} is represented in gray, the Minkowsky sum is represented in black and the ellipsoid enlarged by δ is represented in red.

Consider two ellipsoids $E_1 = \text{Diag}(\frac{1}{a^2}, \frac{1}{b^2})$ and $E_2 = \text{Diag}(\frac{1}{(a+\delta)^2}, \frac{1}{(b+\delta)^2})$. E_1 is an ellipsoid with a and b as semi-major and semi-minor axes, respectively. E_2 represents the ellipsoid E_1 enlarged by δ in both axis. The goal is to find the smallest ellipsoid that bounds the Minkowsky sum. This is equivalent to find the minimum value of δ such that the minimum distance between ellipsoid E_1 and E_2 is bigger than r^2 , the radius of the circle bounding the robot.

Lemma 1 ([26]): Let $X^T A_1 X = 1$ and $X^T A_2 X = 1$ be two quadratics in $\mathbb{R}^n$. If the matrix $A_1 - A_2$ is sign definite, then the square of the distance between the quadratic $X^T A_1 X = 1$ and the quadratic $X^T A_2 X = 1$ equals the minimal positive zero of the polynomial.

$$F(z) = D_\lambda(\det \lambda A_1 + (z - \lambda A_2) - \lambda(z - \lambda A_1 A_2))$$

where D stands for the discriminant of the polynomial treated with respect to λ .

Considering $A_1 = E_1$, $A_2 = E_2$ and $\{\delta, a, b\} \in \mathbb{R}^+$ this ensures that $E_1 - E_2$ is sign definite. Hence, we can apply Lemma 1 to determine the polynomial $F(z)$ and its roots. For the two ellipsoids E_1 and E_2 , the roots λ of $F(z)$ are:

$$\lambda \in \left\{ \begin{array}{l} \frac{(2a(r+\delta)^3 + 2b(r+\delta)^3 + 4ab(r+\delta)^2)}{(a^2 + 2ab + 2ra + b^2 + 2rb)}, \\ (r+\delta)^2, \\ 4a^2 + 4a(r+\delta) + (r+\delta)^2, \\ 4b^2 + 4b(r+\delta) + (r+\delta)^2 \end{array} \right\}. \quad (5)$$

The first two roots have multiplicity two. The minimum distance equation is the square root of the minimal positive zero of $F(z)$. Thus, the minimum enlargement factor is found by solving for the value of δ that satisfies $\min_j \lambda(j) = r^2$.

A closed form formula can be obtained by noting that the $\min_j \lambda(j)$ is achieved for the first root and solving this equation. Due to its length, it is not presented in this letter but can be found

at.³ This value of semi-axis $\alpha = a + \delta$ and $\beta = b + \delta$ guarantee that the constraint ellipsoid entirely bounds the collision space.

C. Model Predictive Contouring Control

MPCC is a formulation specially tailored to path-following problems. This section presents how to modify the baseline method [2] and make it applicable to mobile robots navigating in unstructured environments with on-board perception.

1) *Progress on Reference Path:* Eq. (2) approximates the evolution of the path parameter by the travelled distance of the robot. In each planning stage we initialize θ_0 . We find the closest path segment, denoted by m , and compute the value of θ_0 via a line search in the neighborhood of the previously predicted path parameter.

2) *Selecting the Number of Path Segments:* As detailed in Section II-D, the *global* reference is composed of M path segments. To lower the computational load, only $\eta \leq M$ path segments are used to generate the *local* reference that is incorporated into the optimization problem. The number of path segments η cannot be arbitrarily small, and there is a minimum number of segments to be selected to ensure the robot follows the reference path along a prediction horizon. The number of path segments η in the local reference path is a function of the prediction horizon length, the individual path segment lengths, and the speed of the robot at each time instance. We select a conservative η by considering the maximum longitudinal velocity v_{max} and imposing that the covered distance is lower than a lower bound of the travelled distance along the reference path, namely,

$$\underbrace{\tau \sum_{j=0}^{N-1} v_j}_{\text{Traveled dist.}} \leq \tau N v_{\text{max}} \leq \underbrace{\sum_{i=m+1}^{m+\eta} \|\mathbf{p}_{i+1}^r - \mathbf{p}_i^r\|}_{\text{Waypoints dist.}} \leq \underbrace{\sum_{i=m+1}^{m+\eta} s_i}_{\text{Ref. path length}}, \quad (6)$$

where m is the index of the closest path segment to the robot, τ is the length of the discretization steps along the horizon, and s_i the length of each path segment.

3) *Maintaining Continuity Over the Local Reference Path:* We concatenate the η reference path segments into a differentiable local reference path $\mathbf{L}^r$, which will be tracked by the LMPCC, as follows

$$\bar{\mathbf{p}}^r(\theta_k) = \sum_{i=m}^{m+\eta} \sigma_{i,+}(\theta_k) \sigma_{i,-}(\theta_k) \mathbf{s}_i(\theta_k), \quad (7)$$

where $\sigma_{i,-}(\theta_k) = 1/(1 + e^{(\theta - \sum_{j=m}^i s_j)/\epsilon})$ and $\sigma_{i,+}(\theta_k) = 1/(1 + e^{(-\theta + \sum_{j=m}^{i-1} s_j)/\epsilon})$ are two sigmoid activation functions for each path segment and ϵ is a small design constant. This representation ensures a continuous representation of the local reference path needed to compute the solver gradients.

4) *Cost Function:* For tracking of the reference path, a contour and a lag error are defined, see Fig. 4 and combined in an

²Note we use r instead of r_{disc} in the reminder of the section to simplify the notation.

³A Mathematica notebook with the derivation of the bound and a Matlab script as example of its computation can be found in <http://www.alonsomora.com/docs/19-debritto-boundcomputation.zip>.

error vector $\mathbf{e}_k := [\tilde{\epsilon}^c(\mathbf{z}_k, \theta_k), \tilde{\epsilon}^l(\mathbf{z}_k, \theta_k)]^T$, with

where $\phi(\theta_k) = \arctan(\partial y^r(\theta_k)/\partial x^r(\theta_k))$ is the direction of the path. Consequently, the LMPCC tracking cost is

where Q_ϵ is a design weight.

The solution which minimizes the quadratic tracking cost defined in Eq. (9) drives the robot towards the reference path. To make progress along the path we introduce a cost term that penalizes the deviation of the robot velocity v_k from a reference velocity v_{ref} , i.e., $J_{\text{speed}}(\mathbf{z}_k, \mathbf{u}_k) = Q_v (v_{\text{ref}} - v_k)^2$ with Q_v a design weight. This reference velocity is a design parameter given by a higher-level planner and can vary across path reference segments.

To increase the clearance between the robot and moving obstacles, we introduce an additional cost term similar to a potential function,

$$J_{\text{repulsive}}(\mathbf{z}_k) = Q_R \sum_{i=1}^n \left(\frac{1}{(\Delta x_k)^2 + (\Delta y_k)^2 + \gamma} \right), \quad (10)$$

where Q_R is a design weight, and $\Delta x_k, \Delta y_k$ represent the components of the distance from the robot to the dynamic obstacles. A small value $\gamma \geq 0$ is introduced for numerical stability. Eq. (10) adds clearance with respect to obstacles and renders the method more robust to localization uncertainties.

Additionally, we penalize the inputs with $J_{\text{input}}(\mathbf{z}_k, \theta_k) = \mathbf{u}_k^T Q_u \mathbf{u}_k$, where Q_u is a design weight.

The LMPCC control problem is then given by a receding horizon nonconvex optimization, formally,

$$J^* = \min_{\mathbf{z}_{0:N}, \mathbf{u}_{0:N-1}, \theta_{0:N}} \sum_{k=0}^{N-1} J(\mathbf{z}_k, \mathbf{u}_k, \theta_k) + J(\mathbf{z}_N, \theta_N) \quad (11a)$$

$$\text{s.t. : } (2a), (2c), \quad (11b)$$

$$c_k^{\text{stat},l,j}(\mathbf{z}_k) > 0, \forall j \in \{1, \dots, n_c\}, l \in \{1, \dots, 4\} \quad (11c)$$

$$c_k^{\text{obst},j}(\mathbf{z}_k) > 1, \forall j \in \{1, \dots, n_c\}, \forall \text{obst} \quad (11d)$$

where the stage cost is $J(\mathbf{z}_k, \mathbf{u}_k, \theta_k) := J_{\text{tracking}}(\mathbf{z}_k, \theta_k) + J_{\text{speed}}(\mathbf{z}_k, \mathbf{u}_k) + J_{\text{repulsive}}(\mathbf{z}_k) + J_{\text{input}}(\mathbf{u}_k)$ and the terminal

Algorithm 1: Local Model Predictive Contouring Control.

- 1: Given $\mathbf{z}_{\text{init}}$, $\mathbf{z}_{\text{goal}}$, $\mathcal{O}_{\text{static}}$, $\mathcal{O}_{\text{dyn}}$, and N
- 2: **for** $t = 0, 1, 2, \dots$ **do**
- 3: $\mathbf{z}_0 = \mathbf{z}_{\text{init}}$
- 4: Estimate θ_0 according to Section III-C1
- 5: Select η according to Eq. (6)
- 6: Build $\bar{\mathbf{p}}^r(\theta_k)$, $k = 1, \dots, N$, according to Eq. (7)
- 7: Compute $c_k^{\text{stat},j}(\mathbf{p}_k)$ along $\mathbf{q}_{0:N}$
- 8: Get dynamic-obstacles predicted pose
 (Section III-B)
- 9: Solve the optimization problem of Eq. (11)
- 10: Apply $\mathbf{u}_0^*$
- 11: **end for**

cost is $J(z_N, \theta_N) := J_{\text{tracking}}(z_N, \theta_N) + J_{\text{repulsive}}(z_N)$. Eq. (11c) and Eq. (11d) are defined by Eq. (3) and Eq. (4), respectively. Algorithm 1 summarizes our method. Note that at each control iteration, we solve (using [3]) Problem (11) until either a Karush-Kuhn-Tucker condition [27] or the maximum number of iterations ($\text{iter}_{\max}$) is satisfied (line 1). The selection of $\text{iter}_{\max}$ is empirically based on the maximum number of iterations allowed within the sampling time of our system to guarantee real-time performance.

IV. RESULTS

This section presents experimental and simulation results for three scenarios with a mobile robot. We evaluate different settings for the parameters in our planner (Section IV-B), as well as compare its performance in static (Section IV-C) and dynamic (Section IV-D) environments against state of art motion planners [5], [11], [23]. A video demonstrating the results accompanies this letter.

A. Experimental Setup

1) *Hardware Setup*: Our experimental platform is a fully autonomous Clearpath Jackal ground robot, for which we implemented on-board all the modules used for localization, perception, motion planning, and control. Our platform is equipped with an Intel i5 CPU@2.6GHz, which is used to run the localization and motion planning modules, a Lidar Velodyne for perception, and an Intel i7 NUC mini PC to run the pedestrian tracker.

2) *Software Setup*: To build the global reference path we first define a series of waypoints and we construct a smooth global path by connecting the waypoints with a clothoid. We then sample intermediate waypoints from this global path and connect them with 3rd order polynomials, which are then used to generate the local reference path.

The robot localizes with respect to a map of the environment, which is created before the experiments. For static collision avoidance the robot utilizes a map that is updated online with data from its sensors and we employ a set of rectangles to model the free space. Our search routine expands the sides of a vehicle-aligned rectangle simultaneously in the occupancy grid

Fig. 5. Experiment for evaluation of performance.

environment with steps of $\Delta_{\text{search}} = 0.05$ m, until either an occupied cell is found or the maximum search distance $\Delta_{\text{max}}^{\text{search}} = 2$ m is reached. Once an expanding rectangle side is fixed as a result of an occupied cell, the rest of the rectangle sides are still expanded to search for the largest possible area. This computation runs in parallel to the LMPCC solver, in a different thread, and with the latest available information. Our experiments employ the open-source SPENCER Pedestrian tracker and 2d laser data [28] for detection and tracking of dynamic obstacles. If a pedestrian is detected, it is removed from the static map and treated as a moving ellipse. Our simulations use the open-source ROS implementation of the Social Forces model [29] for pedestrian simulation.

The LMPCC problem of Eq. (11) is nonconvex. Our planner solves this problem online in real time using ACADO [3] and its C-code generation tool. We use a continuous-time kinematic unicycle motion model [30] to describe the robot's kinematics. The model is then discretized directly in ACADO using a multiple-shooting method combined with a Gauss-Legendre integrator of order 4, no Hessian approximations, and a sampling time of 50 ms. We select qpOASES [31] to solve the resulting QP problem and set a KKT tolerance of 10^{-4} and a maximum of 10 iterations. If no feasible solution is found within the maximum number of iterations, then the robot decelerates. The planner computes a new solution in the next cycle. Based on our experience, this allowed to recover the feasibility of the planner quickly. Our motion planner is implemented in C++/ROS and will be released open source.

B. Parameter Evaluation

To evaluate the performance of the planner we performed several experiments at different reference speeds, $v_{\text{ref}} \in \{1 \text{ m/s}, 1.25 \text{ m/s}, 1.5 \text{ m/s}\}$, and prediction-horizon lengths, $T_{\text{Horizon}} \in \{1 \text{ s}, 3 \text{ s}, 5 \text{ s}\}$. The robot follows a Fig. 8 path (red line in Fig. 5(a)) while avoiding two pedestrians (green ellipses) and staying within the collision-free area (yellow rectangles). We use one circle to represent the planned position of the robot (light blue circles). Each pedestrian is bounded by an ellipse of semi-axis 0.3 m and 0.2 m. The predicted positions of the pedestrians are represented by a green line. We align the semi-minor axis to the pedestrian walking direction. For this experiment we rely on a motion capture system to obtain the

Fig. 6. Computation time required to solve the LMPCC problem online for different velocity references and horizon lengths. The central mark indicates the median. The bottom and top edges of the box indicate the 25th and 75th percentiles, respectively. The red crosses represent the outliers.

TABLE I
CLEARANCE BETWEEN THE ROBOT AND THE DYNAMIC OBSTACLES FOR DIFFERENT VELOCITY REFERENCES AND HORIZON LENGTHS

Prediction Horizon [s]	Clearance Mean (1st percentile) [m]		
	$v_{\text{ref}} = 1$	$v_{\text{ref}} = 1.25$	$v_{\text{ref}} = 1.5$
1	1.54 (0.077)	1.60 (0.025)	1.69 (0.003)
3	1.66 (0.077)	1.69 (0.072)	1.82 (0.065)
5	1.69 (0.062)	1.68 (0.055)	1.92 (0.043)

Fig. 7. Static collision avoidance scenario. The red crosses depict the path to follow (waypoints) and the colored points are the laser scan data. The blue, green, and magenta lines represent the trajectory obtained with the LMPCC, the MPC tracking controller, and the Dynamic Window, respectively.

position of the obstacles and robot. Fig. 6 shows the computation time to solve the optimization problem. For $v_{\text{ref}} \in \{1, 1.25, 1.5\}$ m/s and $T_{\text{Horizon}} \in \{1, 3\}$ s the computation times are under 50 ms, which is lower than the cycle-time defined for the planner. But, for $T_{\text{Horizon}} = 5$ s the 99th percentile is above the 50 ms, not respecting the real-time constraint. The cases in which the planner exceeds the sampling time of the system are the situations in which the pedestrians suddenly step in front of the robot or change their direction of motion, requiring the solver more iterations to find a feasible solution. If not all the constraints can be satisfied, our problem becomes infeasible, and no solution is found. In this case, we reduce the robot velocity, allowing the solver to recover the feasibility after few iterations. Table I summarizes the behavior of the planner in terms of clearance, that is, the distance from the border of the circle of the robot to the border of the ellipse defining the obstacles. It demonstrates that a short horizon leads to lower safety distances and more importantly, that our method was able to keep a safe clearance in most cases.

Based on these results, we selected a reference speed of 1.25 m/s and a horizon length of 3 s for the following experiments.

TABLE II

STATISTIC RESULTS OF MINIMUM DISTANCE TO THE PEDESTRIANS (WHERE CLEARANCE IS DEFINED AS THE BORDER TO BORDER DISTANCE), TRAVELING DISTANCE AND PERCENTAGE OF FAILURES OBTAINED FOR 100 RANDOM TEST CASES OF THE DYNAMIC COLLISION AVOIDANCE EXPERIMENT FOR $n \in \{2, 4, 6\}$ AGENTS. THE PEDESTRIANS FOLLOW THE SOCIAL FORCES MODEL [29]

# agents	Clearance Mean (1st percentile) [m]			% failures (% collisions / % stuck)			Traveled distance Mean (Std.) [m]		
	DW	CADRL	LMPCC	DW	CADRL	LMPCC	DW	CADRL	LMPCC
2	0.28 (0)	0.15 (0)	0.29 (0.015)	20 (19 / 1)	4 (0 / 0)	2 (0 / 2)	17.99 (3.9)	19.27 (3.7)	15.81 (3.2)
4	0.46 (0)	0.32 (0)	0.25 (0.026)	35 (32 / 3)	31 (0 / 0)	5 (2 / 3)	19.43 (5.6)	21.34 (4.4)	15.77 (4.4)
6	0.06 (0)	0.33 (0)	0.38 (0.013)	43 (41 / 2)	51 (51 / 0)	7 (5 / 2)	21.09 (4.8)	18.97 (5.6)	16.13 (1.9)

C. Static Collision Avoidance

In this experiment we compare the proposed planner with two baseline approaches:

- *A MPC tracking controller* [11]: We minimize the deviation from positions on the reference path up to 1 m ahead of the robot.
- *The Dynamic Window (DW)* [5]: We use the open-source ROS stack implementation. The DW method receives the next waypoint once the distance to the current waypoint is less than 1 m.

For this and the following experiment, we fully rely on the onboard localization and perception modules. A VLP16 Velodyne Lidar is used to build and update a local map centered in the robot for static collision avoidance. The prebuilt offline map is only used for localization.

In this experiment the mobile robot navigates along a corridor while tracking a global reference path (red waypoints in Fig. 7). When it encounters automatic doors, the robot must wait for them to open. When it encounters the obstacle located near the third upper waypoint from the left, which was not in the map, the robot must navigate around it. Fig. 7 shows the results of this experiment. All the three approaches are able to follow the waypoints and interact with the automatic door. When they encounter the second obstacle the classical MPC design fails to proceed towards the next waypoint. The DW and the LMPCC approaches are able to complete the task. Both methods showed similar performance (e.g., the traveled distance was 30.59 m for our LMPCC and 30.61 m for the DW). The time to the goal was relatively higher for the DW (50 s) compared to the LMPCC (41 s). This difference is mainly due to the ability of the LMPCC to follow a reference velocity. In addition, we tested DW and LMPCC in a scenario where half of the first door does not open due to malfunction. In such scenario, the LMPCC was able to traverse around the broken door while the DW gets into a deadlock state due to the narrow opening. A video of the experiments accompanies the letter.

D. Dynamic Collision Avoidance

Our simulations compare our planner with an additional baseline: *a socially-aware motion planner* named Collision Avoidance with Deep RL (CADRL) [23]. We employ the open-source ROS stack implementation. The CADRL method receives the same waypoints as the DW method.

1) *Simulation Results*: We compare the performance of the MPCC planner with the DW and CADRL baselines in the presence of pedestrians. Fig. 1 shows the setup of our experiment.

Fig. 8. Dynamic collision avoidance scenario. The red crosses represent the global path to follow (waypoints). The blue and magenta lines represent the trajectory executed by our LMPCC (top) and by the Dynamic Window (bottom), respectively. The trajectories of the two pedestrians are represented by the green and magenta circles. In the lower case (dynamic window) the robot reacts late and the pedestrians must actively avoid it.

To avoid the overlap of the static and dynamic collision constraints, the detected pedestrians are removed from the updated map used for static avoidance and modeled as ellipses with a constant velocity estimate. The global path, as described in Section IV-A, consists of a straight line along the corridor. The robot has to follow this path while avoiding collisions with several pedestrians moving in the same or opposite direction. In this experiment, we do not evaluate the MPC tracking controller since it was unable to complete the previous experiment. Aggregated results in Table II show that the LMPCC outperforms the other methods. It achieves a considerably lower failure rate, smaller traveling distances, and maintains larger safety distances to the pedestrians. Only for the four pedestrians case, the DW achieved larger mean clearance, but with larger standard deviation. By accounting for the predictions of the pedestrians, our method can react faster and thus generate safer motion plans. Table II also shows that the number of failures grows with the number of agents. Yet, our method can scale up to six pedestrians with low collision probability and perform real-time. For larger crowds our method would select the closer 6 pedestrians.

2) *Experimental Results*: We use the previous setup to compare the LMPCC and DW methods on a real scenario with two pedestrians. We do not test the CADRL method because the current open-source implementation does not allow static collision avoidance for unconstrained scenarios such as the faculty corridor depicted in Fig. 8. Fig. 8 shows one representative run of our method (top) and the DW (bottom). We observe that the proposed method reacts in advance to avoid the pedestrians, resulting in a larger clearance distance. In contrast, the DW reacts

late to avoid the pedestrian, which has to avoid the robot himself actively. Our proposed method was able to navigate safely in all of our experiments with static and two pedestrians.

E. Applicability to an Autonomous Car

To validate the applicability of our method to more complex robot models, we have performed a simulation experiment with a kinematic bicycle model [32] of an autonomous car.

The planner commands the acceleration and front steering. The car follows a global reference path while staying within the road boundaries (i.e., the obstacle-free region) and avoiding moving obstacles (such as a simulated cyclist proceeding in the direction of the car and a pedestrian crossing the road in front of the car). The accompanying video shows the results where the autonomous vehicle successfully avoids the moving obstacles, while staying within the road limits. We refer the reader to [33] for more details and results.

V. CONCLUSIONS & FUTURE WORK

This letter proposed a local planning approach based on Model Predictive Contouring Control (MPCC) to safely navigate a mobile robot in dynamic, unstructured environments. Our local MPCC relies on an upper bound of the Minkowski sum of a circle and an ellipse to safely avoid dynamic obstacles and a set of convex regions in free space to avoid static obstacles. We compared our design with three baseline approaches (classical MPC, Dynamic Window, and CADRL). The experimental results demonstrate that our method outperforms the baselines in static and dynamic environments. Moreover, the light implementation of our design shows the scalability of our method up to six agents and allowed us to run all algorithms on-board. Finally, we showed the applicability of our design to more complex robots by testing the design in simulation using the model of an autonomous car. As future work, we intend to expand our approach for crowded scenarios, by accounting for the interaction effects between the robot and the other agents.

REFERENCES

- [1] W. Schwarting, J. Alonso-Mora, and D. Rus, "Planning and decision-making for autonomous vehicles," *Annu. Rev. Control, Robot., Auton. Syst.*, vol. 1, no. 1, pp. 187–210, 2018.
- [2] W. Schwarting, J. Alonso-Mora, L. Paull, S. Karaman, and D. Rus, "Safe nonlinear trajectory generation for parallel autonomy with a dynamic vehicle model," *IEEE Trans. Intell. Transp. Syst.*, vol. 19, no. 9, pp. 2994–3008, Sep. 2018.
- [3] B. Houska, H. Ferreanu, and M. Diehl, "ACADO toolkit—An open source framework for automatic control and dynamic optimization," *Optimal Control Appl. Methods*, vol. 32, pp. 298–312, 2011.
- [4] O. Khatib, "Real-time obstacle avoidance for manipulators and mobile robots," in *Proc. Auton. Robot. Veh.*, 1986, pp. 396–404.
- [5] 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, Mar. 1997.
- [6] G. Ferrer, A. Garrell, and A. Sanfeliu, "Robot companion: A social-force based approach with human awareness-navigation in crowded environments," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2013, pp. 1688–1694.
- [7] J. Alonso-Mora, P. Beardsley, and R. Siegwart, "Cooperative collision avoidance for nonholonomic robots," *IEEE Trans. Robot.*, vol. 34, no. 2, pp. 404–420, Apr. 2018.
- [8] J. M. Maciejowski, *Predictive Control: With Constraints*. London, U.K.: Pearson, 2002.
- [9] F. Borrelli, A. Bemporad, and M. Morari, *Predictive Control for Linear and Hybrid Systems*. Cambridge, U.K.: Cambridge Univ. Press, 2017.
- [10] F. Borrelli, P. Falcone, T. Keviczky, J. Asgari, and D. Hrovat, "MPC-based approach to active steering for autonomous vehicle systems," *Int. J. Veh. Auton. Syst.*, vol. 3, no. 2–4, pp. 265–291, 2005.
- [11] C. M. Kang, S. Lee, and C. C. Chung, "On-road path generation and control for waypoints tracking," *IEEE Intell. Transp. Syst. Mag.*, vol. 9, no. 3, pp. 36–45, Fall 2017.
- [12] T. M. Howard, C. J. Green, and A. Kelly, "Receding horizon model-predictive control for mobile robot navigation of intricate paths," in *Field and Service Robotics*, 2010, pp. 69–78.
- [13] D. Fridovich-Keil, S. L. Herbert, J. F. Fisac, S. Deglurkar, and C. J. Tomlin, "Planning, fast and slow: A framework for adaptive real-time safe trajectory planning," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2018, pp. 387–394.
- [14] J. F. Fisac *et al.*, "Probabilistically safe robot planning with confidence-based human predictions," *Robot.: Sci. Syst.*, 2018.
- [15] D. Lam, C. Manzie, and M. Good, "Model predictive contouring control," in *Proc. 49th IEEE Conf. Decis. Control*, 2010, pp. 6137–6142.
- [16] M. Brown, J. Funke, S. Erlien, and J. C. Gerdes, "Safe driving envelopes for path tracking in autonomous vehicles," *Control Eng. Pract.*, vol. 61, pp. 307–316, 2017.
- [17] T. Kruse, A. K. Pandey, R. Alami, and A. Kirsch, "Human-aware robot navigation: A survey," *Robot. Auton. Syst.*, vol. 61, pp. 1726–1743, 2013.
- [18] P. Trautman, "Sparse interacting Gaussian processes: Efficiency and optimality theorems of autonomous crowd navigation," in *Proc. IEEE 56th Annu. Conf. Decis. Control*, 2017, pp. 327–334.
- [19] B. Lütjens, M. Everett, and J. P. How, "Safe reinforcement learning with model uncertainty estimates," in *Proc. Int. Conf. Robot. Autom.*, 2019, pp. 8662–8668.
- [20] P. Long, T. Fanl, X. Liao, W. Liu, H. Zhang, and J. Pan, "Towards optimally decentralized multi-robot collision avoidance via deep reinforcement learning," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2018, pp. 6252–6259.
- [21] J. Alonso-Mora, S. Baker, and D. Rus, "Multi-robot formation control and object transport in dynamic environments via constrained optimization," *Int. J. Robot. Res.*, vol. 36, no. 9, pp. 1000–1021, Jul. 2017.
- [22] Z. Zhu, E. Schmerling, and M. Pavone, "A convex optimization approach to smooth trajectories for motion planning with car-like robots," in *Proc. 54th IEEE Conf. Decis. Control*, 2015, pp. 835–842.
- [23] M. Everett, Y. F. Chen, and J. P. How, "Motion planning among dynamic, decision-making agents with deep reinforcement learning," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2018, pp. 3052–3059.
- [24] G. Welch *et al.*, "An introduction to the Kalman filter," 2006.
- [25] J. Chen, W. Zhan, and M. Tomizuka, "Constrained iterative LQR for on-road autonomous driving motion planning," in *Proc. 20th Int. Conf. Intell. Transp. Syst.*, 2017, pp. 1–7.
- [26] A. Y. Uteshev and M. V. Yashina, "Metric problems for quadrics in multidimensional space," *J. Symbolic Comput.*, vol. 68, pp. 287–315, 2015.
- [27] H. W. Kuhn and A. W. Tucker, "Nonlinear programming," in *Proc. 2nd Berkeley Symp. Math. Statist. Probability*, 1951, pp. 481–492.
- [28] T. Linder, S. Breuers, B. Leibe, and K. O. Arras, "On multi-modal people tracking from mobile platforms in very crowded and dynamic environments," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2016, pp. 5512–5519.
- [29] D. Helbing and P. Molnar, "Social force model for pedestrian dynamics," *Phys. Rev. E*, vol. 51, no. 5, 1995, Art. no. 4282.
- [30] R. Siegwart and I. R. Nourbakhsh, *Introduction to Autonomous Mobile Robots*, 2nd ed. Cambridge, MA, USA: MIT Press, 2011.
- [31] H. J. Ferreanu, C. Kirches, A. Potschka, H. G. Bock, and M. Diehl, "qpOASES: A parametric active-set algorithm for quadratic programming," *Math. Program. Comput.*, vol. 6, pp. 327–363, 2014.
- [32] J. Kong, M. Pfeiffer, G. Schildbach, and F. Borrelli, "Kinematic and dynamic vehicle models for autonomous driving control design," in *Proc. IEEE Intell. Veh. Symp.*, 2015, pp. 1094–1099.
- [33] L. Ferranti *et al.*, "SafeVRU: A research platform for the interaction of self-driving vehicles with vulnerable road users," in *Proc. IEEE Intell. Veh. Symp.*, 2019.