Encounter-Aware Optimal Control Model for Wheelchair-Obstacle Interactions

Damla Leblebicioğlu[1,2], Ege Özgül[1,2], and Taşkın Padır[1,2]

***Abstract*— Autonomous navigation and shared-control systems offer a pathway to improve independent mobility and quality of life for older adults. A wide range of path-planning algorithms have been developed for smart wheelchairs operating in environments shared with pedestrians. Conventional approaches often treat pedestrians and other wheelchair users as dynamic obstacles to be avoided, while interaction-aware methods may rely on pedestrian trajectory prediction or learned interaction models that are difficult to estimate reliably during close encounters. This work formulates wheelchair–pedestrian and wheelchair–wheelchair encounters as a two-agent optimal control problem. The proposed formulation computes a minimum-effort, minimum-risk co-maneuver while promoting recovery of the agents' pre-encounter velocity direction, speed, and path. The joint solution serves as a normative reference for efficient cooperative behavior. Optimal-control solutions across varying encounter conditions will then be used to train a multi-layer perceptron controller to approximate the optimal wheelchair response without repeatedly solving the optimization problem online.**

I. INTRODUCTION

According to the World Health Organization (WHO), the number of people aged 80 and older is expected to triple between 2020 and 2050, reaching 426 million [1], [2], highlighting the growing need for mobility assistive technologies such as powered wheelchairs (PWs). For many older adults, age-related declines in motor control, reduced upper-limb strength, hand tremor, and cognitive impairments can make PW operation physically and cognitively demanding [3], [4]. Autonomous navigation and shared-control systems offer a pathway to improve independent mobility and quality of life for this population [5]. Smart wheelchair (SW) systems (PW platforms equipped with embedded computers, sensors, and controllers capable of autonomous navigation or shared control) have seen remarkable advances in recent years [6], [7]. Despite this progress, navigating around humans in outdoor and indoor environments remains a fundamental challenge for SWs [8]. A wide range of path planning algorithms have been developed for SW navigation in environments shared with pedestrians, including A* [9], artificial potential fields (APF) [10], [11], sampling-based planners such as rapidly exploring random tree (RRT) [12], and optimization-based approaches including model predictive control (MPC) and optimal control [13], [14], [15], [8]. To account for moving pedestrians, the social force model (SFM) [16], [17], [15] and optimal reciprocal collision avoidance (ORCA) [18] have been used to predict pedestrian trajectories, which are then incorporated into the wheelchair control framework.

Conventional path planning approaches treat pedestrians as dynamic obstacles to be avoided [17], [14], [8], and often fail to generate smooth or continuous maneuvers in close human proximity [6]. Furthermore, approaches that rely on pedestrian trajectory prediction require assumptions or models of how other agents will move or respond. This may not always hold in practice. The proposed framework instead makes a single behavioral assumption about the interacting agent, whether a pedestrian or another wheelchair user: the agent is collision-averse. The wheelchair is therefore expected to act optimally without requiring the interacting agent to execute an optimal response. The interaction between the wheelchair and surrounding individuals affects not only the safety and comfort of the wheelchair user, but also those sharing the environment [19]. Therefore, human-aware autonomous maneuvering is an essential capability for SW systems.

In this work, pedestrians and other wheelchair users are treated as interaction-aware moving agents rather than obstacles. This shifts the core question from: *"How should a wheelchair avoid an obstacle?"* to: *"How can the encounter between a wheelchair and a pedestrian or another wheelchair be modeled, learned and integrated into autonomous navigation?"*. This is particularly relevant in environments such as nursing homes and assisted living facilities, where a SW may need to maneuver not only around pedestrians but also around other wheelchair users, in shared spaces that are often narrow and densely occupied.

We formulate the two-agent (wheelchair-pedestrian or wheelchair-wheelchair) encounter as an optimal control problem. The desired behavior of the navigation system is described as follows: upon entering the encounter zone, the wheelchair does not stop or reroute, but follows a minimum-energy, minimum-risk maneuver obtained from a joint two-agent optimal control formulation, with both agents returning to their original trajectories as closely as possible after the interaction. The joint optimal solution provides a normative reference for efficient collision-avoidance behavior and does not require the interacting agent to execute the corresponding optimal maneuver. In this way, the wheelchair can generate a continuous and safe trajectory with minimal changes to its original motion behavior.

This research was funded, in part, by the Advanced Research Projects Agency for Health (ARPA-H) Agreement No. 140D042590012. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government.

[1]Department of Electrical and Computer Engineering, Northeastern University, Boston, MA 02115, USA.

[2]Institute for Experiential Robotics, Northeastern University, Boston, MA 02115, USA.

Emails: {leblebicioglu.d, ozgul.e, t.padir}@northeastern.edu.

The contributions of this paper are as follows:

1. The mathematical model of a wheelchair–pedestrian or wheelchair-wheelchair encounter is formulated as a two-agent dynamic interaction problem, where both agents are represented by their initial positions and velocities at the encounter boundary. The wheelchair is able to generate an interaction-aware maneuver rather than simply treating the interacting agent as a dynamic obstacle.

2. Unlike social navigation approaches that attempt to model or predict human behavior, the proposed framework does not require the other agent to follow a predefined intent, trajectory, or optimal strategy. The other agent is assumed to be only collision-averse, and the wheelchair does not rely on the interacting agent to execute the optimal maneuver computed by the joint formulation. Instead, optimal behavior is used as a normative benchmark. This allows the wheelchair to generate an efficient risk-reducing maneuver while also providing a reference for interpreting what an energetically efficient avoidance behavior would look like if both agents responded optimally.

3. The cost function combines collision-risk minimization with recovery of the agents' pre-encounter motion. The proposed cost function is composed of collision-risk reduction (L_{risk}), including a minimum-separation penalty, control-effort minimization ($L_{control}$), return to the initial velocity direction ($L_{direction}$), return to the initial velocity magnitude (L_{speed}), return to the initial path line (L_{line}), and terminal recovery terms ($J_{terminal}$). As a result, the wheelchair is not merely driven away from the obstacle to avoid collision in an undesired direction.

4. The encounter is formulated as a continuous co-maneuvering problem, where both agents remain in motion and adjust their trajectories smoothly to reduce collision risk, rather than stopping, yielding or reversing.

The rest of this paper is organized as follows: Section II is the related work, Section III presents the mathematical formulation of the two-agent encounter; Section IV describes the simulation setup and results; Section V details the multi-layer perceptron (MLP) training process; Section VI presents the physical wheelchair set-up for hardware experiments; Section VII concludes the paper and discusses directions for future work.

II. RELATED WORK

Early work on path-planning in dense crowds introduced the concept of *joint collision avoidance*, recognizing that pedestrians adapt their trajectories to one another to create space for movement [20]. This observation motivated the Interacting Gaussian Process (IGP) framework, which jointly models possible robot and pedestrian trajectories while penalizing trajectory combinations that bring agents into close proximity. [20], [21]. Accurately predicting pedestrian interactions remains challenging, as interactions inferred from neighboring agents' past motions may be too complex to predict reliably and may have limited influence on pedestrian trajectories in many cases [22]. Pedestrian trajectory decision models have been identified using inverse optimal control and Gaussian process regression [14], while predicted pedestrian trajectories have been incorporated directly into MPC-based wheelchair planning [15]. MPC-based wheelchair planning has also been developed to account for uncertainty in pedestrian behavior in constrained environments [8]. Learning-based approaches instead learn collision-avoidance policies or representations of agent interactions from training data [23], [24]. Collectively, these approaches account for human behavior through prediction, modeling, uncertainty-aware planning, or learned interaction representations.

III. MATHEMATICAL MODEL FOR TWO-AGENT ENCOUNTER INTERACTIONS

In this section, the mathematical model for the encounter between two agents: a wheelchair (Agent 1) and a dynamic obstacle, which may be a pedestrian or another wheelchair (Agent 2) is developed. Each agent $i \in \{1, 2\}$ is characterized by a position vector $\bar{p}_i(t) = [x_i(t),\ y_i(t)]^\top$ and a velocity vector $\bar{v}_i(t) = [v_{xi}(t),\ v_{yi}(t)]^\top \in \mathbb{R}^2$, both defined with respect to a fixed global reference frame and evolving over time. Given the initial velocity (encoding both magnitude and direction) and position vectors of both agents at the encounter boundary, the central question is: *what is the optimal behavior for minimizing the risk and returning back to original behavior with minimum energy, given that both move simultaneously, neither stopping nor yielding, but maneuvering together continuously?*

A. Encounter-Boundary

The encounter is modeled as a *local, event-triggered* interaction problem. The inter-agent distance is defined as $r(t) = \|\bar{p}_1(t) - \bar{p}_2(t)\|_2$. The encounter is initialized when two conditions are simultaneously satisfied: the agents enter a predefined encounter region of distance $d_{\text{enc}} = 2.5$ m ($r(t) \leq d_{\text{enc}}$), and the collision risk function $h(\phi(t))$ exceeds a predefined threshold ($h(\phi(t)) \geq h_{\text{thresh}}$). Here, ϕ denotes the angle between the agent's relative position ($\bar{p}_{\text{rel}}(t) = \bar{p}_2(t) - \bar{p}_1(t)$) and relative velocity vectors ($\bar{v}_{\text{rel}}(t) = \bar{v}_2(t) - \bar{v}_1(t)$), (see Fig. 1). The vector $\bar{p}_{\text{rel}}(t)$ points from the wheelchair to the obstacle, defining where the obstacle is relative to the wheelchair. The vector $\bar{v}_{\text{rel}}(t)$ is the velocity of the obstacle relative to the wheelchair. The collision risk function (given in Eq. 1, Fig. 2) is defined as a Gaussian centered at $\phi = 180°$:

$$h(\phi(t)) = \exp\left(-\frac{1}{2}\frac{(\phi - 180°)^2}{\sigma^2}\right) \tag{1}$$

where $\sigma = 30°$ is the standard deviation controlling the width of the risk region, and $h_{\text{thresh}} = 0.01$ is the activation threshold. With these parameters, the high-risk regime spans approximately $\phi \in [90°, 270°]$.

The full system state vector is 8-dimensional, $\bar{x}(t) = [\bar{p}_1(t), \bar{p}_2(t), \bar{v}_1(t), \bar{v}_2(t)]^\top \in \mathbb{R}^8$. The control input for each agent is defined as $\bar{u}_i(t) = [u_{ix}(t), u_{iy}(t)]^\top \in \mathbb{R}^2$, $i \in \{1, 2\}$, representing the force, in Newton, applied to each agent along the x and y axes, respectively. The total control input vector is then $\bar{u}(t) =$

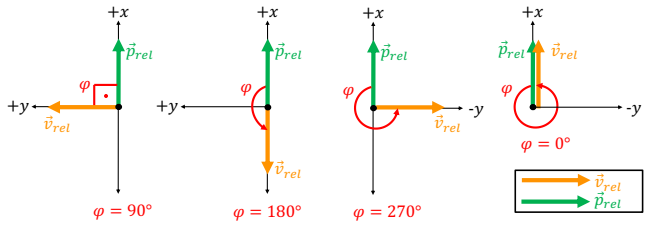


Fig. 1. The angle ϕ is measured from $\bar{p}_{\text{rel}}$ in the counter-clockwise (CCW) direction as positive. It captures whether the interacting agent is approaching the wheelchair ($\phi = 180°$), moving away ($\phi = 0°$), or passing laterally ($\phi = 90°, 270°$).

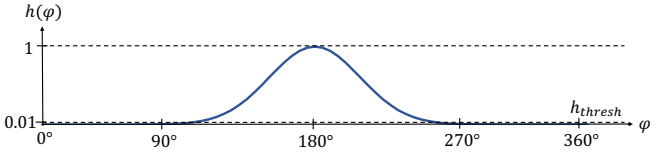


Fig. 2. The collision risk function $h(\phi(t))$: $h(\phi) = 1$ corresponds to maximum risk when the agents are moving directly toward each other, and $h(\phi) = 0$ when they are moving away. The threshold h_{thresh} at which the collision risk starts to increase corresponds to $h(\phi) = 0.01$.

$[u_{1x}(t), u_{1y}(t), u_{2x}(t), u_{2y}(t)]^\top \in \mathbb{R}^4$. The optimal controller is active only during the local encounter interval. Specifically, the system operates in four modes as described in Algorithm 1. Upon encounter detection, the optimal control problem is activated over the planning horizon $[t_0, t_f]$, where t_0 is the moment the encounter condition is first satisfied and t_f is the end of the planned trajectory. N is the number of intervals over the planning horizon and U_{plan} is the computed optimal control sequence.

The basic objective is to determine the optimal control inputs for both agents that produce a natural encounter behavior while minimizing deviations from their *pre-encounter behavior* with low control effort.

B. Optimal Control Problem Formulation

The trajectory of each agent is governed by Newton's second law: $m\,\ddot{\bar{x}} = \bar{F}$, where $\bar{F} = \bar{u}$ is the force vector generated by the optimal controller. Introducing the generalized coordinates $\bar{x}_1 = \bar{p}_1$, $\bar{x}_2 = \bar{p}_2$, $\bar{x}_3 = \bar{v}_1$, and $\bar{x}_4 = \bar{v}_2$, hence the state equations are:

$$\dot{\bar{x}}_1 = \bar{x}_3, \tag{2}$$
$$\dot{\bar{x}}_2 = \bar{x}_4, \tag{3}$$
$$\dot{\bar{x}}_3 = \frac{1}{m_1}\,\bar{u}_1, \tag{4}$$
$$\dot{\bar{x}}_4 = \frac{1}{m_2}\,\bar{u}_2, \tag{5}$$

where m_1 and m_2 are the masses of the wheelchair and the obstacle, respectively. Equations of motion (EoM) in the state-space form can be written as $\dot{\bar{x}}(t) = A\,\bar{x}(t) + B\,\bar{u}(t)$,

Algorithm 1 Encounter-Aware Optimal Controller Activation

Initialize: $k \leftarrow 0$, $U_{\text{plan}} \leftarrow \emptyset$
for $i = 0, 1, 2, \ldots, T$ **do**
  Compute $r(t;i)$, $h(\phi(t;i))$ from current state $\bar{x}(t;i)$
  if $U_{\text{plan}} = \emptyset$ **then**
    if $r(t;i) > d_{\text{enc}}$ **then**
      Mode 1: Outside encounter region
      $\bar{u}(t;i) \leftarrow \bar{0}$ ▷ no encounter detected, continue at nominal velocity
    else if $r(t;i) \le d_{\text{enc}}$ **and** $h(\phi(t;i)) < h_{\text{thresh}}$ **then**
      Mode 2: Inside encounter region, low risk
      $\bar{u}(t;i) \leftarrow \bar{0}$ ▷ not risky, maintain nominal velocity
    else if $r(t;i) \le d_{\text{enc}}$ **and** $h(\phi(t;i)) \ge h_{\text{thresh}}$ **then**
      Mode 3: Inside encounter region, high risk
      $k \leftarrow 0$, solve optimal control over the horizon $[t_0, t_f]$, store U_{plan} ▷ compute $\bar{u}^*(0), \bar{u}^*(1), \ldots, \bar{u}^*(N-1)$
    end if
  end if
  if $U_{\text{plan}} \neq \emptyset$ **and** $k < N$ **then**
    $\bar{u}(t;i) \leftarrow \bar{u}^*(k)$, $k \leftarrow k+1$ ▷ execute planned control step k
  end if
  if $k \ge N$ **and** $U_{\text{plan}} \neq \emptyset$ **then**
    Mode 4: Post-encounter
    $U_{\text{plan}} \leftarrow \emptyset$, $k \leftarrow 0$
    $\bar{u}(t;i) \leftarrow \bar{0}$ ▷ deactivate optimal controller, continue with last velocity computed by $\bar{u}^*(N-1)$
    Check $r(t;i)$, $h(\phi(t;i))$
    if $r(t;i) \le d_{\text{enc}}$ **and** $h(\phi(t;i)) \ge h_{\text{thresh}}$ **then**
      $k \leftarrow 0$, solve optimal control, store U_{plan} ▷ Mode 3: optimal controller reactivated
    else
      $\bar{u}(t;i) \leftarrow \bar{0}$ ▷ Mode 1 or 2: resume nominal behavior
    end if
  end if
  $\dot{\bar{x}}(t;i) \leftarrow A\bar{x}(t;i) + B\bar{u}(t;i)$ ▷ state update via roll-out dynamics
end for

with A and B are given as:

$$A = \begin{bmatrix} 0_{4\times4} & I_{4\times4} \\ 0_{4\times4} & 0_{4\times4} \end{bmatrix}, \qquad B = \begin{bmatrix} 0_{4\times4} \\ \operatorname{diag}\left(\frac{1}{m_1}, \frac{1}{m_1}, \frac{1}{m_2}, \frac{1}{m_2}\right) \end{bmatrix} \tag{6}$$

The optimal control problem [25] is defined as:

$$\begin{aligned} \min_{\bar{u}} \quad & J(\bar{u}) \\ \text{s.t.} \quad & \dot{\bar{x}}(t) = A\,\bar{x}(t) + B\,\bar{u}(t), \quad \bar{x}(t_0) = \bar{x}_0 \\ & -u_{\max} \le u_i(t) \le u_{\max}, \quad i = 1,2,3,4 \end{aligned} \tag{7}$$

The optimal control problem is supposed to minimize a cost function $J(\bar{u})$ subject to the state equations (Eqs. 2–6) as dynamic constraints (constraints enter the problem implicitly via the calculation of states as a function of control inputs). By solving this problem, we obtain the optimal policy of what both the wheelchair and the obstacle should do under different encounter conditions. The overall cost function is defined over the planning horizon $[t_0, t_f]$ as:

$$\begin{aligned} J(\bar{u}) = \int_{t_0}^{t_f} \Big(& L_{\text{risk}}(t) + L_{\text{direction}}(t) \\ & + L_{\text{speed}}(t) + L_{\text{line}}(t) + L_{\text{control}}(t) \Big)\, dt \\ & + J_{\text{terminal}}. \end{aligned} \tag{8}$$

The proximity penalty function $f(r(t))$ is a piecewise constant function such that it increases as the agents get closer and rises sharply when the distance falls below $d_{\text{enc}}/1.5$, with $c_2 \gg c_1 > 0$:

$$f(r(t)) = \begin{cases} 0 & \text{if } r > d_{\text{enc}} \\ c_1 & \text{if } \dfrac{d_{\text{enc}}}{1.5} < r(t) \le d_{\text{enc}} \\ c_2 & \text{if } r(t) \le \dfrac{d_{\text{enc}}}{1.5} \end{cases} \tag{9}$$

The collision risk cost ($L_{risk}(t)$) is then:

$$L_{\text{risk}}(t) = f(r(t)) \cdot h(\phi(t)) + w_{\text{col}} \max\big(0,\, d_{\text{safe}} - r(t)\big)^2 \tag{10}$$

where $w_{\text{col}} \gg 0$ is the collision penalty weight. The velocity direction cost ($L_{direction}(t)$) penalizes deviation from the initial velocity direction.

$$\begin{aligned} L_{\text{direction}}(t) = w(t) \Big[& (\hat{v}_1(t) - \hat{v}_1(t_0))^\top Q_{\text{dir}} (\hat{v}_1(t) - \hat{v}_1(t_0)) \\ & + (\hat{v}_2(t) - \hat{v}_2(t_0))^\top Q_{\text{dir}} (\hat{v}_2(t) - \hat{v}_2(t_0)) \Big] \end{aligned} \tag{11}$$

where $\hat{v}_i(t) = \bar{v}_i(t)/\|\bar{v}_i(t)\|$ is the unit direction vector, $Q_{\text{dir}} \in \mathbb{R}^{2\times 2}$ is a diagonal, positive definite matrix used for both agents. $Q_{\text{dir}} = \text{diag}(\rho,\, \rho)$ with $\rho > 0$ applied along both x and y axes equally. The recovery weight $w(t)$ (given in Eq. 12), increases quadratically over time and $k_{\text{recovery}} \ge 0$ controls the recovery rate. So that, in the beginning agents can maneuver more freely and towards the end they will return to the pre-encounter velocity direction. At $t = t_0$, $w = 1$; at $t = t_f$, $w = 1 + k_{\text{rec}}$.

$$w(t) = 1 + k_{\text{recovery}} \left(\frac{t - t_0}{t_f - t_0} \right)^2 \tag{12}$$

The velocity magnitude recovery cost (L_{speed}) penalizes deviation from the initial speed of each agent, ensuring that both agents return to their pre-encounter speeds:

$$\begin{aligned} L_{\text{speed}}(t) = w(t)\, q_{\text{speed}} \Big[& \big(\|\bar{v}_1(t)\| - \|\bar{v}_1(t_0)\| \big)^2 \\ & + \big(\|\bar{v}_2(t)\| - \|\bar{v}_2(t_0)\| \big)^2 \Big] \end{aligned} \tag{13}$$

where $q_{\text{speed}} > 0$ is a scalar weighting parameter and $w(t)$ is the same recovery weight as in Eq. 12. The cross-track error cost ($L_{line}(t)$) penalizes lateral deviation from the initial trajectory line. The perpendicular distance of agent i from its initial trajectory line is:

$$e_i(t) = (\bar{p}_i(t) - \bar{p}_i(t_0))^\top \hat{n}_i \tag{14}$$

where $\hat{n}_i = [-\hat{v}_{i,y}(t_0),\ \hat{v}_{i,x}(t_0)]^\top$ is the unit normal vector to the initial velocity direction. The cross-track cost is then:

$$L_{\text{line}}(t) = w(t)\, q_{\text{line}} \left[e_1(t)^2 + e_2(t)^2 \right] \tag{15}$$

where $q_{\text{line}} > 0$ is a scalar weighting parameter and $w(t)$ is the same recovery weight as in Eq. 12. The control effort cost (L_{control}) penalizes large control efforts to force minimum energy maneuvers:

$$L_{\text{control}}(t) = \alpha\, \bar{u}(t)^\top D\, \bar{u}(t) \tag{16}$$

where $\alpha > 0$ is a scaling parameter and $D \in \mathbb{R}^{4\times 4}$ is a diagonal, positive definite matrix, $D = \text{diag}(d_1, d_2, d_3, d_4)$ with $d_i > 0$, penalizing the control forces along each axis of both agents. The terminal cost forces that both agents return to their initial velocity directions and trajectory lines at the end of the planning horizon t_f:

$$\begin{aligned} J_{\text{terminal}} = & (\hat{v}_1(t_f) - \hat{v}_1(t_0))^\top Q_{\text{term}} (\hat{v}_1(t_f) - \hat{v}_1(t_0)) \\ & + (\hat{v}_2(t_f) - \hat{v}_2(t_0))^\top Q_{\text{term}} (\hat{v}_2(t_f) - \hat{v}_2(t_0)) \\ & + q_{\text{terminal}} \left[e_1(t_f)^2 + e_2(t_f)^2 \right] \end{aligned} \tag{17}$$

where $Q_{\text{term}} \in \mathbb{R}^{2\times 2}$ is a diagonal, positive definite matrix, $Q_{\text{term}} = \text{diag}(\gamma,\, \gamma)$ with $\gamma > 0$, and $q_{\text{terminal}} > 0$ is a scalar weighting parameter for the terminal cross-track error.

C. Numerical Solution of the Optimal Control Problem

The optimal control problem is infinite-dimensional in its original form. To solve it numerically, the planning horizon $[t_0, t_f]$ is divided into $N = 20$ equal intervals of length $\Delta t = (t_f - t_0)/N$, yielding the time grid $t_i = t_0 + i\,\Delta t$ for $i = 0, 1, \ldots, N$. The control input $\bar{u}(t)$ is assumed to be constant over each interval, i.e., $\bar{u}(t) = \bar{u}_{id}$ for $t \in [t_i, t_{i+1}]$. The unknown vector $\bar{U}_d$ is constructed by concatenating individual $\bar{u}_{id}$'s as:

$$\bar{\mathbf{U}}_{\mathbf{d}} = [\bar{u}_0^\top, \bar{u}_1^\top, \ldots, \bar{u}_{N-1}^\top]^\top \in \mathbb{R}^{4N} \tag{18}$$

resulting in $4 \times 20 = 80$ unknowns in total. The state vector is a dependent variable, computed by solving state equations. Using the forward difference approximation:

$$\bar{x}_{i+1} = \bar{x}_i + \Delta t\, (A\bar{x}_i + B\bar{u}_i) \tag{19}$$

The continuous cost integral is approximated using the trapezoid rule where:

$$\int_{t_0}^{t_f} f(t)\, dt \approx \frac{\Delta t}{2} \sum_{i=0}^{N-1} (f_i + f_{i+1}) \tag{20}$$

The resulting finite-dimensional optimization problem is solved using Particle Swarm Optimization (PSO), a well-known gradient-free method suitable for non-linear optimization problems.

IV. SIMULATION RESULTS

In simulations, $m_1 = 100$ kg for the wheelchair and $m_2 = 80$ kg for the interacting agent, modeled as a human. The optimal control problem was solved using a global-best PSO algorithm. The optimization horizon was set to $t_f = 10$ s and discretized into $N = 20$ intervals, resulting in a time step of $\Delta t = t_f/N = 0.5$ s. Only the control inputs are treated as optimization variables; the state trajectories of the wheelchair and the interacting agent are generated from the initial conditions through the roll-out dynamics. The control dimension is $n_u = 4$, corresponding to the x- and y-direction control forces of both agents. Therefore, the PSO decision vector contains $Nn_u = 20 \times 4 = 80$ variables.

Each PSO particle represents one candidate control sequence over the complete optimization horizon. The global-best PSO was run using 50 particles for 2000 iterations, with cognitive and social coefficients of 1.5 and an inertia weight of 0.7. The control forces of both agents were bounded to ± 500 N throughout the optimization. The resulting agent trajectories are shown in Figs. 3 and 4 for the maximum-risk ($\phi = 180°$) and high-risk side-approach ($\phi = 166.7°$) encounter conditions, respectively. In Fig. 4, the minimum-control-effort objective leads the wheelchair to reduce its speed with minimal lateral deviation, while the lighter interacting agent (human) accelerates and performs a larger avoidance maneuver to prevent the inter-agent distance from falling below d_{safe}.

In practice, the interacting-agent (i.e., human) may not execute the joint-optimal maneuver. The human's optimal controls are therefore replaced with a model-predictive controller (MPC) that minimizes only the human's own cost while anticipating the wheelchair's plan (Fig. 5). To compare the agents, the cost in Eq. 8 is split into J_1 and J_2, each consisting of the shared risk term (L_{risk}) and the remaining terms of the corresponding agent. As the prediction horizon grows, the human's cost (J_2) approaches the joint-optimal value, reaching within 2% at 2.0 s, showing that a self-interested agent adopts a near-optimal response without being instructed to cooperate (see Tab. I). From 2.5 s onward, J_2 falls slightly below the optimum, since the best-response problem (at most 12 variables) is solved more accurately than the 80-variable joint problem. This comes at the price of incomplete terminal recovery, as the human's terminal cost remains above that of the optimal maneuver.

V. MLP CONTROLLER DESIGN

Solving the optimal control problem in real time is computationally demanding. Therefore, the optimal control problem will be solved offline over a 19×15 grid of encounter conditions, sweeping the angle ϕ from 90° to 270° in 10° increments, and the human's speed up to 6 km/h in 15 equal steps. For each scenario, the optimal control inputs

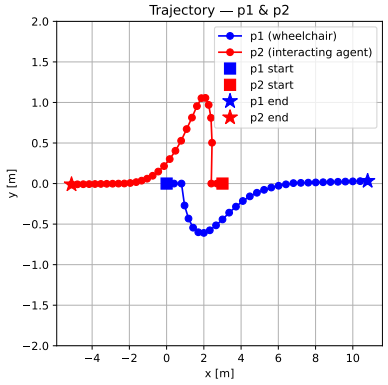


Fig. 3. Maximum-risk encounter with $\phi = 180°$. The initial states are $\bar{p}_1 = [0, 0]$ m, $\bar{p}_2 = [3, 0]$ m, $\bar{v}_1 = [0.8, 0]$ m/s, and $\bar{v}_2 = [-0.6, 0]$ m/s. The resulting relative velocity is $\bar{v}_{\text{rel}} = [-1.4, 0]$ m/s, which is opposite to the relative position vector $\bar{p}_{\text{rel}} = [3, 0]$ m, corresponding to a head-on, maximum-risk encounter.

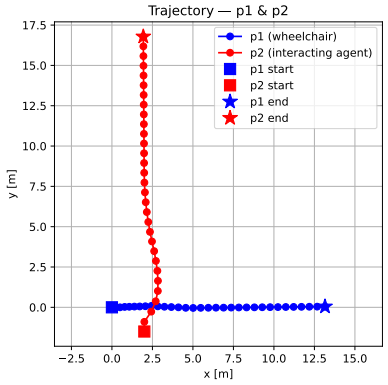


Fig. 4. High-risk side-approach encounter with $\phi \approx 166.7°$. The initial states are $\bar{p}_1 = [0, 0]$ m, $\bar{p}_2 = [2, -1.5]$ m, $\bar{v}1 = [1, 0]$ m/s, and $\bar{v}2 = [0, 1.2]$ m/s. The resulting relative position and velocity vectors are $\bar{p}\text{rel} = [2, -1.5]$ m and $\bar{v}\text{rel} = [-1, 1.2]$ m/s, respectively. The optimal maneuver reduces the wheelchair speed while increasing the interacting agent's speed to reduce collision risk while both agents remain in motion.

will be discretized into 20 time steps and recorded, together with the entry and exit times of the risk region, to train a MLP [26]. During deployment, the MLP will provide a continuous mapping from the encounter-entry state to the control input in real time.

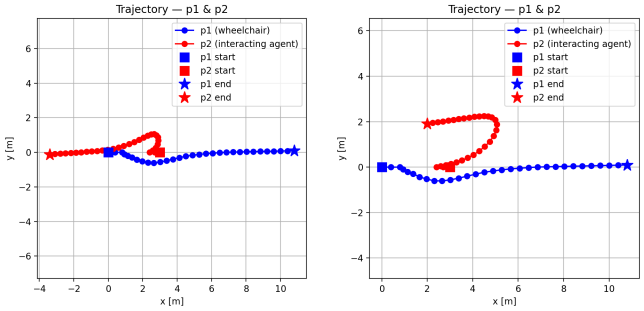
VI. EXPERIMENTAL SET-UP

The LUCI SW system [27] will be used for the hardware experiments, as shown in Fig. 6. The wheelchair will operate in AutonomousRemote mode with the LUCI collision-avoidance system disabled. Two complementary sets of hardware experiments will be conducted to evaluate the proposed framework. The first set of experiments will be conducted in a $5.4 \text{ m} \times 5.3 \text{ m}$ cage equipped with an OptiTrack motion-capture (MoCap) system. Eight overhead cameras provide coverage of the experimental area. Reflective markers will be mounted on the wheelchair and on beanies worn by the human participants to track their positions and orientations throughout the experiments. The MoCap system will provide ground-truth position and orientation measurements for both the wheelchair and the human participants. The second set

TABLE I

HEAD-ON ENCOUNTER: OPTIMAL HUMAN VS. HUMAN MPC ANTICIPATING THE WHEELCHAIR'S PLAN FOR DIFFERENT PREDICTION HORIZONS H. THE WHEELCHAIR EXECUTES THE SAME JOINT-OPTIMAL PLAN IN ALL RUNS. J_1, J_2: WHEELCHAIR AND HUMAN COSTS (SHARED RISK PLUS OWN TERMS); J_2^{term}: HUMAN TERMINAL COST; JOINT: $J_1 + J_2 - J_{\text{risk}}$; J_2^*: HUMAN COST IN THE OPTIMAL CASE; $t_{r<d_{\text{safe}}}$: TIME WITH $r < d_{\text{safe}}$.

Human	H [s]	$t_{r<d_{\text{safe}}}$ [s]	J_1 ($\times 10^3$)	J_2 ($\times 10^3$)	J_2^{term}	Joint ($\times 10^3$)	J_2/J_2^*
Optimal	–	0.5	33.7	52.5	0.1	62.2	1.00
MPC	1.0	1.0	34.0	125.7	24600	135.4	2.39
MPC	1.5	1.0	33.9	67.1	4100	76.8	1.28
MPC	**2.0**	**1.0**	**33.8**	**53.5**	**808**	**63.3**	**1.02**
MPC	2.5	0.5	33.6	49.9	175	59.6	0.95
MPC	3.0	0.5	33.6	49.0	54	58.7	0.93



(a) $H = 3.0$ s (b) $H = 1.0$ s

Fig. 5. Trajectories with the human MPC anticipating the wheelchair's plan. The wheelchair executes the same joint-optimal plan in both runs. (a) With a 3.0 s horizon, the human performs a near-optimal avoidance maneuver. (b) With a 1.0 s horizon, the human avoids the wheelchair but deviates considerably further from its original direction and path than in (a).

of experiments will be conducted in an open area outside the MoCap environment using onboard sensors. In this configuration, LiDAR-based localization will be used to estimate the wheelchair pose, while RGB-D cameras will be used to detect and track pedestrians and estimate their relative positions and velocities with respect to the wheelchair.

VII. FUTURE WORK AND CONCLUSION

As future work, the MLP controller will be trained and evaluated through hardware experiments on the physical wheelchair platform. In parallel, the optimal control parameters (Q_{dir}, Q_{term}, D, α, σ) used in this study will be calibrated and validated using motion capture data collected from real human-wheelchair encounter experiments. Multiple encounter configurations will be tested by varying the initial angle between the velocity vectors of the two agents as well as their initial speeds. The encounter-entry states (position and velocity vectors recorded at the moment both encounter conditions are satisfied, i.e., $r(t) \leq d_{\text{enc}}$ and $h(\phi) \geq h_{\text{thresh}}$) will serve as the initial conditions for the optimal control simulation. The resulting simulated trajectories will be compared against the measured experimental trajectories collected under motion capture. The parameters will be

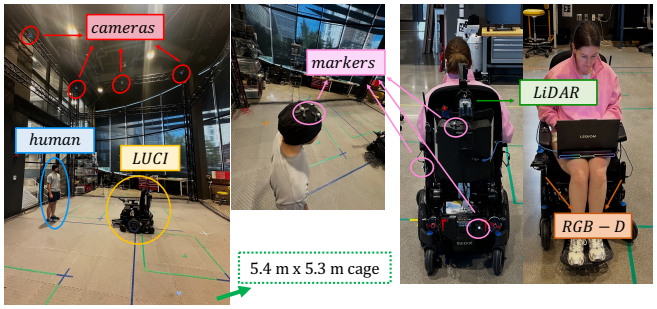


Fig. 6. LUCI SW platform and OptiTrack MoCap system.

iteratively tuned until convergence, with 70% of trajectories used for calibration and 30% reserved for generalization.

REFERENCES

[1] World Health Organization, "Ageing and health," 2022, accessed: 8 July 2025. [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/ageing-and-health

[2] United Nations, Department of Economic and Social Affairs, Population Division, "World population ageing 2019," United Nations, Tech. Rep., 2019, accessed: 8 July 2025. [Online]. Available: https://digitallibrary.un.org/record/3846855

[3] D. Labbé, W. B. Mortenson, P. W. Rushton, L. Demers, and W. C. Miller, "Mobility and participation among ageing powered wheelchair users: using a lifecourse approach," *Ageing & Society*, vol. 40, no. 3, pp. 626–642, 2020.

[4] R. Patusco, K. Gyasi, and A. Kaufmann, "Exploring the nexus between inflammation and mobility through the lens of healthy aging: current scenario and future perspectives," *Aging Clinical and Experimental Research*, 2026.

[5] J. Leaman and H. M. La, "A comprehensive review of smart wheelchairs: past, present, and future," *IEEE Transactions on Human-Machine Systems*, vol. 47, no. 4, pp. 486–499, 2017.

[6] G. Atulan, K. R. Chowdhury, F. Nusrat, and M. P. Schoen, "Evolution and emerging trends in intelligent wheelchair control: A comprehensive review," *Machines*, vol. 14, no. 1, p. 33, 2025.

[7] Y. Morales, N. Kallakuri, K. Shinozawa, T. Miyashita, and N. Hagita, "Human-comfortable navigation for an autonomous robotic wheelchair," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 2737–2743.

[8] H. Oinuma, K. Sekiguchi, and K. Nonaka, "A self-driving wheelchair navigating narrow and winding paths accepting potential soft collisions with pedestrians," in *2025 American Control Conference (ACC)*. IEEE, 2025, pp. 2681–2688.

[9] M. Hosny, R. Alsarrani, and A. Najjar, "Indoor wheelchair navigation for the visually impaired," in *International Conference on Human-Computer Interaction*. Springer, 2015, pp. 411–417.

[10] K. Shibata, N. Shibata, K. Nonaka, and K. Sekiguchi, "Model predictive obstacle avoidance control for vehicles with automatic velocity suppression using artificial potential field," *IFAC-PapersOnLine*, vol. 51, no. 20, pp. 313–318, 2018.

[11] O. Khatib, "Real-time obstacle avoidance for manipulators and mobile robots," *The international journal of robotics research*, vol. 5, no. 1, pp. 90–98, 1986.

[12] J. Rios-Martinez, A. Spalanzani, and C. Laugier, "Understanding human interaction for probabilistic autonomous navigation using risk-rrt approach," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2011, pp. 2014–2019.

[13] E. Kawaguchi, K. Sekiguchi, and K. Nonaka, "Self-driving electric wheelchair in crowded environments using a fuzzy potential model predictive control," *IFAC-PapersOnLine*, vol. 56, no. 2, pp. 11 827–11 833, 2023.

[14] H. Murayama, E. Uchiyama, and G. Venture, "Comfortable collision avoidance control for pedestrians based on optimal control and inverse optimal control," in *2025 IEEE International Conference on Advanced Robotics and its Social Impacts (ARSO)*. IEEE, 2025, pp. 305–310.

[15] X. Zuo, M. Mukai, M. A. S. Kamal, and K. Yamada, "Model predictive control of an autonomous wheelchair considering pedestrian reaction using social force model for collision avoidance," *SICE Journal of Control, Measurement, and System Integration*, vol. 18, no. 1, p. 2596874, 2025.

[16] D. Helbing and P. Molnar, "Social force model for pedestrian dynamics," *Physical review E*, vol. 51, no. 5, p. 4282, 1995.

[17] M. Kalenberg, M. Heine, S. Ghanate, C. Hofmann, J. Walter, J. Franke, and J. Fürst, "Social force-based proxemics for mutual and comfortable navigation of an intelligent wheelchair," in *2026 IEEE International Conference on Advanced Robotics and its Social Impacts (ARSO)*. IEEE, 2026, pp. 1–6.

[18] J. Van Den Berg, S. J. Guy, M. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," in *Robotics research: the 14th international symposium ISRR*. Springer, 2011, pp. 3–19.

[19] S. Udupa, V. R. Kamat, and C. C. Menassa, "Shared autonomy in assistive mobile robots: a review," *Disability and Rehabilitation: Assistive Technology*, vol. 18, no. 6, pp. 827–848, 2023.

[20] P. Trautman and A. Krause, "Unfreezing the robot: Navigation in dense, interacting crowds," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2010, pp. 797–803.

[21] P. Trautman, J. Ma, R. M. Murray, and A. Krause, "Robot navigation in dense human crowds: Statistical models and experimental studies of human–robot cooperation," *The International Journal of Robotics Research*, vol. 34, no. 3, pp. 335–356, 2015.

[22] C. Schöller, V. Aravantinos, F. Lay, and A. Knoll, "What the constant velocity model can teach us about pedestrian motion prediction," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1696–1703, 2020.

[23] Y. F. Chen, M. Liu, M. Everett, and J. P. How, "Decentralized non-communicating multiagent collision avoidance with deep reinforcement learning," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 285–292.

[24] C. Chen, Y. Liu, S. Kreiss, and A. Alahi, "Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2019, pp. 6015–6022.

[25] D. E. Kirk, *Optimal control theory: an introduction*. Courier Corporation, 2004.

[26] S. Haykin, *Neural Networks: A Comprehensive Foundation*, 2nd ed. Prentice Hall, 1999.

[27] LUCI Mobility, "LUCI Smart Wheelchair Technology," https://www.luci.com, accessed: Sep. 4, 2026.