

ARMCHAIR

Integrated Inverse Reinforcement Learning and Model Predictive Control for Human-Robot Collaboration

Caregnato-Neto, Angelo; Siebert, Luciano Cavalcante; Zgonnikov, Arkady; Maximo, Marcos R.O.A.; Afonso, Rubens J.M.

DOI

[10.1007/s12369-025-01332-4](https://doi.org/10.1007/s12369-025-01332-4)

Publication date

2025

Document Version

Final published version

Published in

International Journal of Social Robotics

Citation (APA)

Caregnato-Neto, A., Siebert, L. C., Zgonnikov, A., Maximo, M. R. O. A., & Afonso, R. J. M. (2025). ARMCHAIR: Integrated Inverse Reinforcement Learning and Model Predictive Control for Human-Robot Collaboration. *International Journal of Social Robotics*, 17(12), 3173-3189. <https://doi.org/10.1007/s12369-025-01332-4>

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](#)
as part of the Taverne amendment.**

More information about this copyright law amendment
can be found at <https://www.openaccess.nl>.

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.



ARMCHAIR: Integrated Inverse Reinforcement Learning and Model Predictive Control for Human-Robot Collaboration

Angelo Caregnato-Neto¹ · Luciano Cavalcante Siebert³ · Arkady Zgonnikov⁴ · Marcos R. O. A. Maximo² · Rubens J. M. Afonso¹

Received: 29 February 2024 / Revised: 22 August 2025 / Accepted: 29 September 2025
© The Author(s), under exclusive licence to Springer Nature B.V. 2025

Abstract

One of the key issues in human-robot collaboration is the development of computational models that allow robots to predict and adapt to human behavior. Much progress has been achieved in developing such models, as well as control techniques that address the autonomy problems of motion planning and decision-making in robotics. However, the integration of computational models of human behavior with such control techniques still poses a major challenge, resulting in a bottleneck for efficient collaborative human-robot teams. In this context, we present a novel architecture for human-robot collaboration: Adaptive Robot Motion for Collaboration with Humans using Adversarial Inverse Reinforcement learning (ARMCHAIR). Our solution leverages adversarial inverse reinforcement learning and model predictive control to compute optimal trajectories and decisions for a mobile multi-robot system that collaborates with a human in an exploration task. During the mission, ARMCHAIR operates without human intervention, autonomously identifying the necessity to support and acting accordingly. Our approach also explicitly addresses the network connectivity requirement of the human-robot team. Extensive simulation-based evaluations demonstrate that ARMCHAIR allows a group of robots to safely support a simulated human in an exploration scenario, preventing collisions and network disconnections, and improving the overall performance of the task.

Keywords Cyber-physical-human systems · Human-robot collaboration · Inverse reinforcement learning · Trajectory planning · Task allocation · Model predictive control

✉ Angelo Caregnato-Neto
caregnato.neto@ieee.org

Luciano Cavalcante Siebert
l.cavalcantesiebert@tudelft.nl

Arkady Zgonnikov
a.zgonnikov@tudelft.nl

Marcos R. O. A. Maximo
mmaximo@ita.br

Rubens J. M. Afonso
rubensjm@ita.br

¹ Electronic Engineering Division, Instituto Tecnológico de Aeronáutica, São José dos Campos, São Paulo, Brazil

² Autonomous Computational Systems Lab (LAB-SCA), Computer Science Division, Instituto Tecnológico de Aeronáutica, São José dos Campos, São Paulo, Brazil

³ Intelligent Systems Department, Interactive Intelligence Group, Delft University of Technology, Delft, The Netherlands

⁴ Department of Cognitive Robotics, Delft University of Technology, Delft, The Netherlands

1 Introduction

A cyber-physical-human system (CPHS) is an emerging concept of control schemes that emphasizes the interaction between machines and humans [1]. An exciting application of CPHS is the design of collaborative multi-robot systems (MRS) devised to support humans in a variety of tasks, from construction [2] to manufacturing [3] and transport [1]. One of the core challenges in the conception of such MRS is the development of their motion planning and decision-making algorithms, which must enable the group of robots to harmoniously operate with humans towards an aligned objective [4].

A particularly attractive technique used to this end is model predictive control (MPC), a receding horizon control scheme that has been extensively demonstrated for multi-robot real-time motion planning and decision-making [5–7]. As a model-based approach, MPC typically requires accurate prediction models of the dynamic elements in the

environment [8]. This requirement is particularly relevant in the context of human-robot teaming, where appropriate group coordination depends on the capability of predicting and adapting to human decisions and motion in real-time [9].

The human behavior prediction models are not only a fundamental component of most predictive control schemes designed for human-robot collaboration but also outlined as a key element in the development of general CPHS [1]. Traditionally, fields such as psychology rely on descriptive predictions of human behavior that are based on first principles [1, 10, 11]. Alternatively, recent advances in machine learning techniques have enabled an increasing number of sophisticated data-driven models for human motion and decisions [12, 13]. As a quantitative methodology, machine learning offers the benefit of providing mathematical or algorithmic models that are naturally encoded into machines [14], facilitating their integration with powerful control techniques. As a result, machine-learned human behavior models and MPC have been successfully integrated to design efficient algorithms for CPHS [6, 7].

Another concrete motivation for this integration is the possibility of enforcing connectivity between the elements of the human-robot team. Often, proper operation of MRS requires the maintenance of network connectivity [5]. This property holds particular importance when humans are integrated into the team since typical interactions, in the form of direct human commands, negotiations, or adaptations, also entail some form of communication. By properly predicting human motion within the MPC framework through machine-learned behavior models, hard constraints on connectivity can be enforced, allowing the robots to receive commands, perform negotiations, or adapt to human actions by continually monitoring them. The maintenance of communication has also been demonstrated as an important factor in improving human-robot trust [15].

1.1 Background

Originally developed for process control in chemical plants [16], model predictive control (MPC) evolved over the last decades into a flexible tool that can be fine-tuned for a wide range of applications. In the field of robotics, it has been already demonstrated for motion planning of complex platforms such as autonomous cars [6] and MRS [17]. The inherent robustness associated with the closed-loop formulation alongside its capability to perform constrained optimization is ideal in addressing major issues in robot motion planning [18], such as active collision avoidance in environments with moving obstacles [7], safe operation in unknown environments [19], and coordinated rendezvous maneuvering [20]. The use of MPC alongside mixed-integer

programming (MIP) further enhances the technique, allowing the optimization of integer and continuous variables. This extension allows MPC-MIP optimization models to encode a variety of additional requirements such as minimum-time maneuvering, task allocation, and connectivity maintenance [5, 21, 22].

To devise human prediction models, this work leverages inverse reinforcement learning (IRL) [23], a machine learning method that has been extensively employed to model human behavior in the context of sequential decision-making, from aircraft pilots [24] to car drivers [25]. Unlike other popular machine learning methods used for this purpose, such as imitation learning [14], IRL algorithms aim to recover a reward function given a dataset of demonstrations. This distinction offers substantial benefits since the ensuing reward functions can be used to infer human intentions and also provide enhanced flexibility, being used for optimizations considering scenarios that deviate from the ones present in the original data [23]. These traits have been leveraged for many applications, such as collaborative autonomous driving [25], selecting trajectories for a manipulator robot that closely resemble observed human motion [26], and inferring objectives of rogue drones [27].

Adversarial Inverse Reinforcement Learning (AIRL) [28] is an alternate rendition of this method that leverages the principle of maximum-entropy IRL [29] and adversarial deep networks. This approach has been shown to outperform established deep imitation learning algorithms in terms of generalization [28, 30], demonstration-efficiency [30], and precision [31]. Similarly to traditional IRL, AIRL has found use in a multitude of applications, from the design of agents that can remain under human meaningful control [32, 33] to improving lane-changing [34] and lane-keeping maneuvers [35] of autonomous vehicles.

1.2 Related Work

We start the presentation of the selected literature with works that employ MPC and machine learning to design algorithms for mobile human-robot teams. In [6], an MPC-based formation controller was designed for human-leading truck platooning. A model of the driver's behavior was computed using IRL and then integrated into decentralized MPC controllers in each autonomous truck. The inherent uncertainty in the behavior of the drivers was handled using chance constraints.

The problem of coordinating a group of robots, humans, and moving obstacles without communication networks is addressed in [7]. A Recurrent Neural Network (RNN) is trained using a dataset of robot demonstrations to develop prediction models. The proposed MPC trajectory planners

are able to steer independent robots in swap maneuvers without *a priori* exchange of their intended paths.

A lane-changing problem in mixed-traffic scenarios is considered in [36]. The proposed approach relies on the recognition of six driver actions that are modeled using hidden Markov models. Probability models are then trained to predict the acceleration of the vehicle given the predicted human actions. An MPC-MIP formulation is employed to safely and efficiently steer the vehicle in the lane-changing maneuver. A strategy based on Interaction-Aware Model Predictive Path Integral (IA-MPPI) is proposed in [37] for trajectory planning of a team of Unmanned Surface Vessels (USV) in urban canals. The system is able to predict trajectories and goals using pedestrian models and social variational recurrent neural networks.

In addition to MPC and machine learning methods, another popular methodology for the development of human-robot collaboration algorithms for mobile platforms are hierarchical architectures that leverage multiple distinct techniques that individually address the problems of decision-making, motion planning, and connectivity. For example [38], proposes an architecture with three layers: a global planner leverages the fast marching square algorithm to provide proper trajectories to the goal, while a local planner addresses collision avoidance. A Social Force Model (SFM) with parameters estimated by a neural network is employed for human prediction. A Lloyd-based controller handles the multi-agent coordination problem and provides further safety guarantees.

A similar approach is proposed in [39], where the task of path planning is divided into global and local planners, while collision avoidance is handled by a low-level controller. Human demonstrations are employed by an IRL algorithm to learn proper social navigation behaviors, which are then used to compute human-like trajectories for the robot. A B-spline path generation algorithm is proposed in [40] to address the problem of humanoid robot navigation alongside humans. A combination of SFM and Bayesian human motion models is used to predict human trajectories. The integration between adaptive control strategies and social proximity potential fields has also been proposed as a solution for human-robot navigation under specific social conventions [41].

In [42] a semi-autonomous approach is proposed where the human performs a search-and-rescue operation with the support of Unmanned Aerial Vehicles (UAVs). The tasks of the team are assigned by the operator while collision avoidance is handled by a trajectory planner based on replanning using linear spatial separations. Different levels of autonomy in human-robot teaming are studied in [43] considering the task of robot-assisted docking of a ship conducted by a human.

The problem of human-robot collaborative search has been successfully demonstrated in [44] through a series of thorough experiments. The proposal leverages the concept of Social Reward Sources (SRS) to model the environment and mission, while an architecture using Monte Carlo Tree Search (MCTS) and Rapidly-exploring Random Trees (RRT) is used to collaboratively navigate a group of robots. The interactions occur through messages with the humans transmitting their intended goals and positions using a smartphone.

1.3 Contributions

The main contribution of this work is a novel architecture, integration strategy, and MILP formulation that enables deep neural networks and MPC-MIP to be jointly leveraged to solve a collaboration and coordination problem considering a human-robot team. The proposed algorithm is entitled Adaptive Robot Motion for Collaboration with Humans using Adversarial Inverse Reinforcement learning (ARMCHAIR). Specifically, ARMCHAIR differentiates itself from other solutions by:

- (1) providing both movement coordination and collaboration in task allocation for an MRS that supports a human in an exploration mission;
- (2) not requiring human intervention, with the robot team identifying the necessity of support and acting as needed;
- (3) handling the joint decision-making and motion planning optimization problems in a unified MPC-MIP framework;
- (4) actively enforcing network connectivity.

Unlike most of the proposals found in the literature, ARMCHAIR addresses: a) the typical movement coordination problems associated with collision avoidance and connectivity maintenance that are required for proper coexistence and cooperation between humans and robots collaborating on the same task; and b) it provides harmonious task allocation with the machines updating their objectives to conform with human decisions, avoiding conflicts and effectively allowing the group to *collaborate* towards an aligned global goal in spite of the uncertainty in the human's decisions.

ARMCHAIRs architecture is presented in Fig. 1. The trajectories are computed with the proposed MPC-MIP scheme, which handles both the motion planning and task allocation issues in the same framework, being able to jointly optimize them to certified global optimality in finite-time [45]. In addition, using the AIRL-based prediction model, ARMCHAIR is able to operate without human intervention by inferring their intentions and adapting to them.

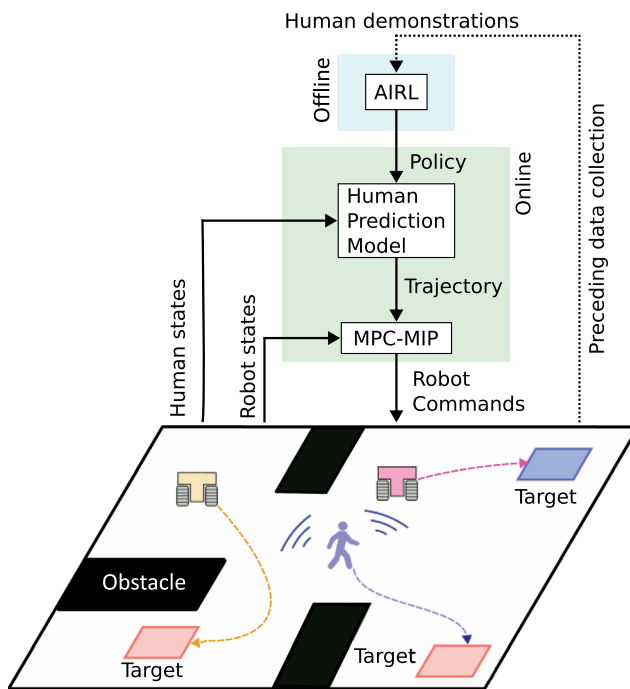


Fig. 1 ARMCHAIR control architecture. The offline AIRL training provides a human prediction model that is used by the MPC-MIP algorithm to compute proper trajectories and decisions for the MRS in a closed loop

This is made possible, in part, by the active enforcement of network connectivity through the constraints in the MPC-MIP problem. The optimization is performed in a closed-loop fashion using the current feedback of the human and robot states, which are sensed and communicated to the planner through the network. The planner provides periodical updates on the trajectories and decisions of the MRS in a receding horizon fashion.

1.4 Notation and Definitions

The prediction of a variable “ o ” for the time step $k + \ell$ made at time step k is written as $\hat{o}(\ell|k)$. Bold letters represent matrices and vectors, whereas scalars are written in plain text. The sets of all integers in the interval $[m, n]$ are written as $\mathcal{I}_m^n = \{m, m + 1, \dots, n\}$. Sets are denoted by calligraphic letters as in \mathcal{B} . All polytopes mentioned in this work are convex. The Minkowski sum and Pontryagin difference operators are denoted by \oplus and \sim , respectively [46].

2 Problem Description

We consider a group of $n_a \in \mathbb{N}$ robots and one human in a task of visiting targets that represent regions of interest in a planar environment. This setup represents a general exploration mission, which is relevant in multiple real-world

contexts such as search-and-rescue and surveillance tasks [44]. The discrete-time dynamics of the robots considering a sampling period of $T > 0$ are described by the following equations

$$\mathbf{x}_i(k+1) = \mathbf{A}_i \mathbf{x}_i(k) + \mathbf{B}_i \mathbf{u}_i(k) \quad (1)$$

$$\mathbf{y}_i(k) = \mathbf{C} \mathbf{x}_i(k), \quad \forall i \in \mathcal{I}_1^{n_a} \quad (2)$$

where $\mathbf{x}_i \in \mathbb{R}^{n_x}$, $\mathbf{u}_i \in \mathbb{R}^{n_u}$, and $\mathbf{y}_i \in \mathbb{R}^2$ are the state, input, and position of robot i , respectively; $\mathbf{A}_i \in \mathbb{R}^{n_x \times n_x}$, $\mathbf{B}_i \in \mathbb{R}^{n_x \times n_u}$, and $\mathbf{C} \in \mathbb{R}^{2 \times n_x}$ are the corresponding matrices of the state-space representation. States and inputs are bounded by the polytopic sets $\mathcal{X} \subset \mathbb{R}^{n_x}$ and $\mathcal{U} \subset \mathbb{R}^{n_u}$, respectively.

The group performs the task within the boundaries of a region represented by the polytope $\mathcal{A} \subset \mathbb{R}^2$ which contains $n_o \in \mathbb{N}$ polytopic obstacles $\mathcal{O}_g \subset \mathbb{R}^2, \forall g \in \mathcal{I}_1^{n_o}$. The space occupied by the body of the robots is defined as

$$\mathcal{R}_i(k) \triangleq \mathbf{y}_i(k) \oplus \mathcal{R}_i, \quad \forall i \in \mathcal{I}_1^{n_a} \quad (3)$$

where $\mathcal{R}_i \subset \mathbb{R}^2$ is a polytope centered at the origin representing the shape of robot i . Similarly, the space occupied by the human is defined as

$$\mathcal{R}^h(k) \triangleq \mathbf{y}^h(k) \oplus \mathcal{R}^h \quad (4)$$

with $\mathbf{y}^h \in \mathbb{R}^2$ being the position of the human and $\mathcal{R}^h \subset \mathbb{R}^2$ representing the human body. Define $\bar{\mathcal{O}} \triangleq \bigcup_{g=1}^{n_o} \mathcal{O}_g$. Then, the free operation region of the i -th robot at time step k is written as, $\forall i \in \mathcal{I}_1^{n_a}$,

$$\bar{\mathcal{A}}_i(k) = \left(\mathcal{A} \setminus \left(\bar{\mathcal{O}} \cup \mathcal{R}^h(k) \cup \bigcup_{\substack{j=1 \\ j \neq i}}^{n_a} \mathcal{R}_j(k) \right) \right) \sim \mathcal{R}_i \quad (5)$$

The tasks in the environment are represented by particular regions denominated targets; we consider that an agent (human or robot) services a task associated with a target by physically reaching it, as established in [22]. We also note that the agents are allowed to move through targets to reach another objective without servicing them.

These targets are represented by the polytopes $\mathcal{T}_e \subset \mathbb{R}^2, \forall e \in \mathcal{I}_1^{n_t}$, where $n_t \in \mathbb{N}$ is the total number of targets. These regions are further divided into two types: A and B. From the perspective of the robots both types are equally valuable, whereas the human may be biased toward a particular type. In the design of the proposed algorithms, the preferences of the human are not explicitly known and can only be inferred from data. The number of targets of

type A and B is denoted by $n_A \geq 0$ and $n_B \geq 0$, respectively, such that $n_A + n_B = n_t$. The mission is finished when the human reaches a terminal region $\mathcal{F} \subset \mathbb{R}^2$, which represents the physical exit (e.g. door) of the environment. The introduction of a spatial terminal condition allows for a more complex problem, where the decisions of the human can be impacted by the positioning of the targets w.r.t. the terminal region, \mathcal{F} . In addition, this choice of terminal condition is consistent with the MILP multi-robot motion planning and decision-making problems addressed in [5, 22]. Figure 2 provides an example of the envisioned scenario.

In our setup, the group of robots must also maintain connectivity at each time step. We consider proximity-based links, where two agents are connected if they are close enough to each other, i.e., if their relative position is within a communication region represented by the polytope $\mathcal{C} \subset \mathbb{R}^2$. Under this condition for individual connections, the whole group can communicate if the underlying network, represented by an undirected graph, is connected [22].

In order to cooperate with the human, the robots make use of a behavior model of that human. This behavior model is initially unknown, although the following assumptions are made:

- A1) A set of demonstrations \mathcal{D} containing human trajectories in similar tasks is available;
- A2) The human operates trying to maximize a compromise between visiting the maximum number of targets and finishing the mission as soon as possible.
- A3) The ensuing behavior is not necessarily optimal nor deterministic;
- A4) The human moves with constant velocity.




Fig. 2 Illustration of environment with a target of type B in \mathcal{T}_1 and two targets of type a represented by the polytopes \mathcal{T}_2 and \mathcal{T}_3 ; three obstacles \mathcal{O}_1 , \mathcal{O}_2 , and \mathcal{O}_3 ; and a terminal region \mathcal{F} . The polytopes \mathcal{R}_1 and \mathcal{R}_2 are outer approximations of the robots' bodies. Connectivity regions are depicted as circles

The set of human demonstrations assumed in A1 can be obtained by monitoring and collecting data directly during their activities. In this proof-of-concept work, we generated synthetic human behavior data rather than collecting it from real humans (see Appendix A.1 for details). This approach allows for preliminary evaluation of algorithms for human-robot collaboration [24, 30, 32, 37, 47] since it can generate human-like data in a scalable way, although it is limited in that it does not exactly represent actual human behavior. Assumptions A2 and A3 imply that the human is aware of their task and has some degree of competency in completing it. However, optimal performance is typically not achieved, as humans are susceptible to bounded rationality and uncertainty over objectives [47–49]. We employed reinforcement learning to train a synthetic agent to emulate these characteristics (Appendix A.1). The resulting policy is optimized to maximize a compromise between collecting rewards (by visiting targets) and finishing the task quickly but is also stochastic and not necessarily globally optimal. Thus, both A2 and A3 are satisfied by the proposed synthetic agent. For the sake of simplicity, Assumption A4 constrains the problem to humans moving with constant velocity. However, the methods herein proposed can be expanded to accommodate a varying velocity. In this context, the main problem addressed by ARMCHAIR is now enunciated.

Problem 1 Design a motion planning and task allocation algorithm that allow the robots to collaborate with the human, complementing their efforts in the task when necessary without the necessity of human intervention. Specifically, the decision-making problem is associated with

- (1) Deciding which regions are likely to be neglected by the human given his observed behavior;
- (2) Deciding whether it is beneficial and viable for the robot team to service the neglected regions considering a compromise between time, fuel expenditure, and the rewards associated with reaching a region and the maximum mission horizon;
- (3) Assigning robots from the group to service such regions if necessary.

The motion planning problem is to compute trajectories such that the robots reach their assigned regions; maintain connectivity between them and with the human; and avoid collisions.

Remark 1 This work focuses on designing a system that allows the robots to passively support the human in the task, i.e., there is no consensus nor competition between the operator and the robot group regarding the tasks. The operator performs his mission as it typically would, disregarding the

robot team. The latter only acts towards tasks if it is able to infer, from the observed human behavior, that a target is likely to be neglected by the operator.

3 Adaptive Robot Motion for Collaboration with Humans Using Adversarial Inverse Reinforcement Learning (ARMCHAIR)

The ARMCHAIR scheme, presented in Fig. 1, addresses Problem 1 by leveraging AIRL to model human behavior through a dataset of their demonstrations while performing the mission in the environment. The time-demanding learning process of AIRL is carried offline and only the resulting policy is used, alongside the feedback of current human and robot states, in the online closed-loop scheme. The MPC-MIP motion planning and decision-making algorithm performs an optimization considering the current state of the group and employs the policy computed by the AIRL method to predict future motion and decisions of the human given their current state. The inherent uncertainty in human actions is handled by a straightforward robustness approach using safety regions.

Each element of the ARMCHAIRs architecture is thoroughly discussed in the subsequent sections.

3.1 Adversarial Inverse Reinforcement Learning

Typically, MPC requires a model to predict the dynamic elements of the environment. Considering the problem description discussed in Sect. 2, ARMCHAIR learns a prediction model for the human given only the set of demonstrations \mathcal{D} under the




Fig. 3 Human position, target of type A, and terminal region T in a grid encoded as the feature maps \mathbf{F}_{pos} , \mathbf{F}_A , \mathbf{F}_{ter} , respectively

aforementioned assumptions. To this end, we leverage AIRL [28] to recover a reward function so that ARMCHAIR can train a policy that imitates the behavior observed in the dataset. This policy will provide predictions of the human’s high-level decisions, i.e., the targets that are likely to be visited, and the corresponding path through the environment. Henceforth, we followed the AIRL approaches from [28, 32].

In particular, we follow the scheme presented in [32] and divide the environment presented in Sect. 2 into $n_c \in \mathbb{N}$ cells of a grid. The ensuing scenario is represented by a Markov Decision Process (MDP) $\langle \mathcal{S}, \mathcal{N}, p, r, \gamma \rangle$ with

$$\mathbf{s} = [\mathbf{F}_{\text{pos}}, \mathbf{F}_A, \mathbf{F}_B, \mathbf{F}_{\text{obs}}, \mathbf{F}_{\text{ter}}] \quad (6)$$

being a state belonging to the set of states \mathcal{S} with the feature maps $\mathbf{F}_{\text{pos}}, \mathbf{F}_A, \mathbf{F}_B, \mathbf{F}_{\text{obs}}, \mathbf{F}_{\text{ter}} \in \{0, 1\}^{n_c \times n_c}$ representing the positions of the human, type A targets, type B targets, obstacles, and terminal region, respectively. Examples of feature maps are presented in Fig. 3. The set of $n_{\text{act}} \in \mathbb{N}$ actions is denoted as $\mathcal{N} = \{a_1, \dots, a_{n_{\text{act}}}\}$, whereas $p(\mathbf{s}'|\mathbf{s}, a)$ is the transition distribution associated with the dynamics of the environment. The reward function and discount factor are denoted by $r : \mathcal{S} \times \mathcal{N} \rightarrow \mathbb{R}$ and $\gamma \in [0, 1]$, respectively.

The two main components of the AIRL framework are the reward and policy networks represented by f_θ and π_ϕ , which are parameterized by θ and ϕ , respectively. The discriminator

$$D_\theta(\mathbf{s}, a, \mathbf{s}') = \frac{\exp(f_\theta(\mathbf{s}, a, \mathbf{s}'))}{\exp(f_\theta(\mathbf{s}, a, \mathbf{s}')) + \pi_\phi(a|\mathbf{s})} \quad (7)$$

previously proposed in [28] outputs the confidence that a particular state-action pair (\mathbf{s}, a) belongs to the demonstration dataset rather than being generated from π_ϕ . The discriminator is trained to improve this data classification using binary logistic regression, further refining the reward approximator f_θ at each iteration. The policy network is updated using PPO [32], to generate trajectories that better match the demonstrations, this way “confusing” the classification process of the discriminator [28]. The policy recovered by AIRL is denoted by $\hat{\pi} : \mathcal{S} \rightarrow \Delta_{\mathcal{N}}$ and is a probability distribution over all actions conditioned by a state, where $\Delta_{\mathcal{N}} \in \mathbb{R}^{n_{\text{act}}+1}$ is the probability simplex over the actions.

3.2 Human Prediction Model

Once $\hat{\pi}$ is properly learned, we can determine sequences of states representing the potential paths and decisions of the human given an initial state. However, due to the stochasticity of the policy, many distinct sequences can be computed. This uncertainty can be addressed by robust MPC planners

with a multitude of techniques e.g., chance constraints [6], contingency [50], or constraint tightening. We opt for a straightforward approach where the predictions are the most likely sequence of states given by the policy, with the uncertainty in the human’s actions being handled by the robust formulation presented in Sect. 3.4. The sequence of states is computed through the propagation of the (8) and (9) until the terminal state is reached:

$$\hat{a}(n) = \operatorname{argmax}_{a' \in \mathcal{N}} x\hat{\pi}(a'|\hat{s}(n)) \tag{8}$$

$$\hat{s}(n+1) = \operatorname{argmax}_{s' \in \mathcal{S}} p(s'|\hat{s}(n), \hat{a}(n)) \tag{9}$$

with \hat{a} and \hat{s} representing action and state predictions, respectively.

Algorithm 1 summarizes the process of computing the human predictions. It starts by associating the first predicted state $\hat{s}(0)$ with the measured state of the human at time step k , $\mathbf{s}^{\text{meas}}(k)$. The corresponding waypoint $\hat{\mathbf{w}}(n) \in \mathbb{R}^2$, representing the predicted Cartesian position of the human, is then extracted from the feature map $\mathbf{F}_{\text{pos}}(n)$ using the function GETWAYPOINTS. Considering the current state, the procedures in lines 5 and 6 compute the most likely action and next state, respectively. This process is repeated in a loop until the predictions reach the terminal state. Then, assuming that the human moves with a constant linear velocity $v > 0$, the predicted trajectories at time step k are computed by the function GETTRAJ yielding a sequence of predicted human positions, $\hat{\mathbf{y}}^h(\ell|k) \in \mathbb{R}^2$, $\forall \ell \in \mathcal{I}_0^{\bar{N}}$ with \bar{N} being the maximum prediction horizon considered.

Algorithm 1 runs periodically, employing updated measurements of the human’s state, with a frequency dictated by the sampling period T .

Algorithm 1 Human prediction model

Input: $\mathbf{s}^{\text{meas}}(k)$ \triangleright Human state at time step k
Input: $\hat{\pi}$ \triangleright Policy recovered with AIRL

```

1:  $n \leftarrow 0$ 
2:  $\hat{\mathbf{s}}(0) \leftarrow \mathbf{s}^{\text{meas}}(k)$ 
3: while terminal region not reached do
4:    $\hat{\mathbf{w}}(n) \leftarrow \text{GETWAYPOINTS}(\hat{\mathbf{s}}(n))$ 
5:    $\hat{a}(n) \leftarrow \operatorname{argmax}_{a' \in \mathcal{N}} \hat{\pi}(a'|\hat{\mathbf{s}}(n))$ 
6:    $\hat{\mathbf{s}}(n+1) \leftarrow \operatorname{argmax}_{s' \in \mathcal{S}} p(s'|\hat{\mathbf{s}}(n), \hat{a}(n))$ 
7:    $n \leftarrow n + 1$ 
8: end while
9:  $\{\hat{\mathbf{y}}^h(\ell|k)\}_{\ell=0}^{\bar{N}} \leftarrow \text{GETTRAJ}(\{\hat{\mathbf{w}}(j)\}_{j=0}^n)$ 

```

Output: $\{\hat{\mathbf{y}}^h(\ell|k)\}_{\ell=0}^{\bar{N}}$

Remark 2 Since the AIRL training is both computationally and time expensive, it must currently be performed offline. Thus, the models are only valid to predict the humans that generated the training datasets, and the robot team can only properly collaborate with these individuals. The study of the generalization capabilities of the models is left for future research.

3.3 MPC-MIP Formulation

The MPC-MIP formulation designed for the human-robot motion planning and decision-making problem is similar to [5], with the distinction of accounting for an independent agent, i.e. the human, over which the group of robots has no command. Within the multi-agent team, the human is always assigned to the last index $n_a + 1$, where n_a is the number of robots, without loss of generality. Implementation details are available in the provided source code.¹

Let the tuple $\lambda = \langle \mathbf{x}_i, \mathbf{u}_i, b^{\text{hor}}, b_{i,e}^{\text{tar}}, b_{i,j}^{\text{con}} \rangle$ represent all variables to be optimized. Then, the human-robot teaming problem can be formalized as the following mixed-integer linear programming (MILP) model.

Human-robot teaming optimization model

$$\min_{\lambda} \sum_{\ell=0}^{\bar{N}} \left(\sum_{i=1}^{n_a} \gamma_u \mathbf{u}_i(\ell|k)_1 - \sum_{e=1}^{n_t} \sum_{i=1}^{n_a+1} \gamma_e b_{i,e}^{\text{tar}}(\ell|k) \right) \tag{10a}$$

s.t.,

Human

$$\forall \ell \in \mathcal{I}_0^{\bar{N}}, \tag{10b}$$

$$\mathbf{y}_{n_a+1}(\ell|k) = \hat{\mathbf{y}}^h(\ell|k)$$

Robots

$$\forall i \in \mathcal{I}_1^{n_a}, \forall \ell \in \mathcal{I}_0^{\bar{N}} \tag{10c}$$

$$\mathbf{x}_i(0|k) = \mathbf{x}_i(k),$$

$$\mathbf{x}_i(\ell+1|k) = \mathbf{A}_i \mathbf{x}_i(\ell|k) + \mathbf{B}_i \mathbf{u}_i(\ell|k) \tag{10d}$$

$$\mathbf{y}_i(\ell|k) = \mathbf{C}_i \mathbf{x}_i(\ell|k) \tag{10e}$$

$$\mathbf{x}_i(\ell+1|k) \in \mathcal{X}_i \tag{10f}$$

$$\mathbf{u}_i(\ell|k) \in \mathcal{U} \tag{10g}$$

$$n_a + 1 \tag{10h}$$

¹ https://gitlab.com/caregnato_netto_open/ARMCHAIR.

Collaboration

$$\forall i \in \mathcal{I}_1^{n_a+1}, \forall e \in \mathcal{I}_1^{n_t}, i < j \leq n_a + 1, \forall \ell \in \mathcal{I}_0^{\bar{N}} \quad (10i)$$

$$b_{i,e}^{\text{tar}}(\ell|k) \Rightarrow \mathbf{y}_i(\ell|k) \in \mathcal{T}_e$$

$$\sum_{\ell=0}^{\bar{N}} \sum_{h=1}^{n_a+1} b_{i,e}^{\text{tar}}(\ell|k) \leq 1 \quad (10j)$$

$$b_{i,j}^{\text{con}}(\ell|k) \Rightarrow \mathbf{y}_i(\ell|k) - \mathbf{y}_j(\ell|k) \in \mathcal{C} \quad (10k)$$

$$\text{deg}_i(\ell|k) = \sum_{i=1}^{i-1} b_{i,i}^{\text{con}}(\ell|k) + \sum_{i=i+1}^{n_a} b_{i,i}^{\text{con}}(\ell|k) \quad (10l)$$

$$\text{deg}_i(\ell|k) + \text{deg}_j(\ell|k) \geq n_a, j \neq n_a + 1 \quad (10m)$$

$$\sum_{i=1}^{n_a} b_{i,n_a+1}^{\text{con}}(\ell|k) \geq 1 \quad (10n)$$

$$b^{\text{hor}}(\ell|k) \Rightarrow \mathbf{y}_{n_a+1}(\ell|k) \in \mathcal{F} \quad (10o)$$

$$\sum_{\ell=0}^{\bar{N}} b^{\text{hor}}(\ell|k) = 1 \quad (10p)$$

The cost (10a) is a trade-off between minimizing the control effort of the robots (weighted by $\gamma_u > 0$) and maximizing the collection of rewards through target visitation. The value assigned to each target collection is given by the weights $\gamma_e > 0, \forall e \in \mathcal{I}_1^{n_t}$. Notice that the activation of the target binaries $b_{i,e}^{\text{tar}} \in \{0, 1\}$, which implies that an agent visited a target due to constraint (10i), results in a discount that contributes to the minimization of (10a). Thus, the cost drives the group towards target visitation while also considering the total control effort (represented by the inputs \mathbf{u}_i) of the maneuvers.

Constraint (10b) performs the assignment of the predictions coming from the human prediction model (Algorithm 1) to the appropriate variables in the optimization model, \mathbf{y}_{n_a+1} .

Constraints (10c) - (10h) are related to the robots. Constraint (10c) establishes the receding horizon scheme of the MPC technique. The dynamics are encoded by (10d) and (10e), whereas (10f) and (10g) enforce the robot's state and input bounds, respectively. Avoidance of collisions between elements of the group and with obstacles is enforced by (10h).

Finally, collaboration between the robots in the group is achieved through constraints (10i) - (10p). The relationship between target binaries, the positions of the agents, and the

target polytopes is described by the implication in constraint (10i). Constraint (10j) establishes that the reward associated with each target can be collected only once. Constraint (10k) conditions the activation of the connectivity binary $b_{i,j}^{\text{con}} \in \{0, 1\}$ of agents i and j to their relative position being within the connectivity region polytope \mathcal{C} . The degree of the vertices that represent the agents in the communication network is encoded by (10l). The condition for connectivity between robots is enforced by constraint (10m), whereas (10n) guarantees that there is always at least one robot connected to the human, forming a network with all elements of the team.

The mission is finished when the horizon binary $b^{\text{hor}} \in \{0, 1\}$ is active. Constraint (10o) conditions the activation of this binary to the human predicted position reaching the terminal region \mathcal{F} , whereas (10p) enforces that the horizon binary can only be activated at exactly one time step.

3.4 Robust Formulation

As discussed in 3.2, the predictions employed by the motion planning algorithm take into account only the most likely actions of the human. We address the ensuing potential prediction errors by making the method more robust through an approach based on safety regions.

To prevent potential collisions, we construct a safety region $\bar{\mathcal{R}}^h \subset \mathbb{R}^2$, that must be avoided by the robots, by enlarging the polytope representing the body of the human to account for any possible direction of movement at the next step (Fig. 4). Define, $\forall i \in \mathcal{I}_1^{n_a}, \forall \ell \in \mathcal{I}_0^{\bar{N}}$,

$$\mathcal{P}_i(k+1) \triangleq \bar{\mathcal{O}} \cup \bigcup_{\substack{j=1 \\ j \neq i}}^{n_a} \mathcal{R}_j(k+1) \quad (11)$$

as the union of the sets representing the obstacles and the bodies of robots for $i \neq j$, i.e., the regions that robot i is not allowed to occupy over time. Then, we adapt the free operation region (5) used in constraint (10h) as, $\forall i \in \mathcal{I}_1^{n_a}, \forall k \in \mathcal{I}_0^{\bar{N}}$,

$$\bar{\mathcal{A}}_i(\ell+1|k) = (\mathcal{A} \setminus (\bar{\mathcal{R}}^h(k+1) \cup \mathcal{P}_i(\ell+1|k))) \sim \mathcal{R}_i \quad (12)$$

A similar approach is employed to prevent network disconnections. A shrunken connectivity region is computed considering the safety region

$$\bar{\mathcal{C}} = \mathcal{C} \sim \bar{\mathcal{R}}^h. \quad (13)$$

The resulting motion planner steers the robots such that connectivity is maintained (for the next step) considering any potential direction that the human can move. Constraint




Fig. 4 Safety region around the human to be avoided by the MRS

(10k) is rewritten to employ the shrunk region when considering connectivity with the human, $\forall i \in \mathcal{I}_1^{n_a}, \forall \ell \in \mathcal{I}_0^{N-1}, j = n_a + 1,$

$$b_{i,j}^{con}(\ell + 1|k) \implies \mathbf{y}_i(\ell|k) - \mathbf{y}_j(\ell|k) \in \bar{C} \tag{14}$$

4 Results

We evaluated the proposed scheme considering two environment configurations with four obstacles, one terminal region, and a maximum of four targets. Environment 1 was designed to investigate the functionality of ARMCHAIR in a situation where the targets are sparsely distributed, whereas environment 2 contains closely grouped targets that represent a more challenging problem in terms of human behavior prediction and appropriate robot response.

In both scenarios, the simulated human is supported by a group of $n_a = 2$ robots with discrete-time dynamics described by (1) and (2) with following state-space matrices

$$\mathbf{A}_i = \begin{bmatrix} 1 & T & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & T \\ 0 & 0 & 0 & 0 \end{bmatrix}, \mathbf{B}_i = \begin{bmatrix} 0.5T^2 & 0 \\ T & 0 \\ 0 & 0.5T^2 \\ 0 & T \end{bmatrix}, \tag{15}$$

and a sampling period of $T = 1$ s. The states are $\mathbf{x}_i = [r_{x,i}, v_{x,i}, r_{y,i}, v_{y,i}]^T, \forall i \in \mathcal{I}_1^{n_a},$ where $r_{x,i}$ and $r_{y,i}$ are the positions of robot i in the x and y axes, respectively, and $v_{x,i}$ and $v_{y,i}$ are the corresponding velocities; the input $\mathbf{u}_i = [a_{x,i}, a_{y,i}]^T, \forall i \in \mathcal{I}_1^{n_a},$ is comprised of the corresponding accelerations. Although simplified, this choice of dynamics model for motion planning has been successfully demonstrated in experiments with real robots using MPC-MIP setups [5]. The velocity and acceleration of the robots are bounded to the intervals $[0, 2.5]$ m/s and $[-5, 5]$ m/s², respectively. A circular connectivity region with radius 4 m

is approximated by a regular octagonal polytope inscribed in the circle $C \subset \mathbb{R}^2,$ following [5]. The width and length of the robots, human, and operation region are 0.5×0.5 m, 0.41×0.29 m, and 7×7 m, respectively. The weights of (10a) were selected as $\gamma_u = 0.01$ and $\gamma_e = 5, \forall e \in \mathcal{I}_1^4,$ i.e., all targets are equally valuable for the robot team.

As stated in Sect. 2, we employed a synthetic agent to represent the human, which operates independently from the robots. Its use allows for a deeper evaluation of the proposed scheme as data can be generated to study the method in a multitude of circumstances. The actions of the synthetic agent are determined by a policy $\pi^* : \mathcal{S} \rightarrow \Delta_{\mathcal{N}}.$ The dataset \mathcal{D} used in the AIRL training process discussed in Sect. 3.1 is generated using $\pi^*.$

The effectiveness of ARMCHAIR was evaluated by comparing it to two baselines:

- (a) open-loop MIP: An open-loop version of our algorithm, where the optimization is solved once at the start of the mission considering the most likely trajectory of the human given by $\hat{\pi};$
- (b) No robot support: The human performs the task without robot support.

A comparison to the baseline without any robot support allowed us to evaluate the influence of proper robot support on the task. Furthermore, by comparing ARMCHAIR with an open-loop MIP approach, we aimed to investigate the impact of the closed-loop strategy employed by the former method, where trajectories and decisions are periodically updated, as opposed to the latter, where a single optimization is performed at the start of the mission.

Since the actions of the simulated human are stochastic, we performed a Monte Carlo evaluation considering 1000 simulations in each environment. To illustrate the performance of ARMCHAIR compared to the baseline methods in each environment, we first highlight several representative simulations in detail (recordings of all 1000 simulations are available in the supplementary material).² We then present statistical analysis across all simulations considering the following metrics: average number of collisions, network disconnections, targets visited, and redundant target visitations, where a target is visited twice in the same simulation due to a conflict in target allocation.

The MILP optimization problem was formulated using the Yalmip [51] and MPT [52] toolboxes. The Gurobi [53] solver was used to compute the solutions.

² https://gitlab.com/caregnato_netto_open/ARMCHAIR.

4.1 Environment 1: Sparse Target Distribution

We found that in the sparse environment, open-loop MIP produces frequent redundant visits to targets in cases when the human deviated from the initially most likely course of action. For instance, in simulation 2 out of 1000 at the initial time step (Fig. 5a) the learned policy $\hat{\pi}$ predicted a trajectory where the human visits only the preferred targets 1 and 2, ignoring the remaining ones. In response, the open-loop MIP dispatched robot 2 to visit targets 3 and 4, while robot 1 maintained the connectivity between the human and robot 2. However, after 21 s (Fig. 5b) the human drifted from the initial prediction by visiting all targets, resulting in redundant visits of targets 3 and 4 which were reached both by the human and robot 2.

Conversely, we observed that the redundant visits were completely prevented when the agents operated using the ARMCHAIR algorithm, as illustrated by simulations 41 and 111 out of 1000 in Fig. 6a to 6h. Simulation 41, shows a significant detour (Fig. 6b and 6c) where the human demonstrates its intention to visit target 3 by moving downwards instead of upwards. ARMCHAIR promptly updates the trajectories of robot 1, stopping it for a few time steps allowing the human to move away from the region around target 4, and then dispatching robot 2 to visit it. We note that

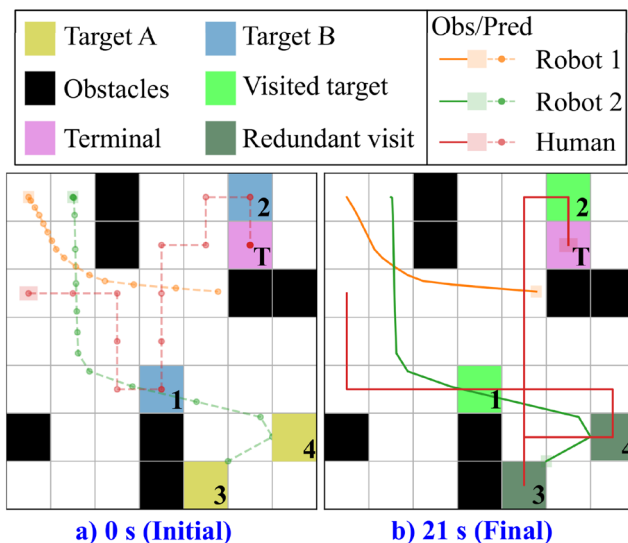


Fig. 5 Example simulation of the open-loop MIP baseline in environment 1 (Sparse target distribution): initial (a) and final (b) time steps for simulation 2 out of 1000. Solid and dashed lines represent the actual motion (observed) and predictions of each agent, respectively. The human prediction is computed using only the initial conditions (open-loop) and $\hat{\pi}$. The prediction suggests that only targets of type B will be visited by the human and the robots are dispatched to the remaining ones. In (b) we observe redundant visits to targets 3 and 4 (type A) since the human deviates from the prediction

the robots can move over targets without collecting them, as it was done in simulation 41 by robot 2, when it moved through target 1 to reach other regions without collecting the corresponding target reward. Simulation 101 depicts a similar situation, where the human visits target 3 but also decides to capture target 4 (Fig. 6f and 6g). As a result, ARMCHAIR updates the trajectory of robot 2, steering it away from the bottom right-hand side region, clearing the path for the human to finish the mission while robot 1 preserves the connectivity of the group.

Confirming the insights based on representative simulations, ARMCHAIR demonstrated superior performance compared to both baselines across all 1000 simulations (Table 1). In particular, ARMCHAIR provided proper responses to the uncertainty in human behavior, as no network disconnections or redundant visits occurred, whereas such issues were present in the open-loop MIP case. The average number of redundant visits was particularly high (0.43) for open-loop MIP, due to the incapability of the open-loop planner to update trajectories and allocation of tasks in real-time. ARMCHAIR also produced few collisions (on average 0.001 collisions per simulation) compared to the open-loop MIP (0.58 collisions per simulation) which was unable to adapt to the human deviating from the initially estimated trajectory. ARMCHAIR was unable to prevent all collisions even though a safety region was employed. This occurs in situations where the optimization problem becomes unfeasible due to the uncertain behavior of the human. Such circumstances trigger an emergency stop and robot team is not allowed to move until the problem becomes feasible again. Potential solutions to this issue are discussed in Section 5. In terms of efficiency, all targets were visited when the open-loop and ARMCHAIR planners were used, as opposed to the case where the human operated alone (2.7). However, only ARMCHAIR prevented redundant visits.

Remark 3 The scenarios where the human “collided” with a robot after is emergency stop due to infeasibility were considered as collisions in our analysis for the sake of rigor. However, the fact that a human enters a cell occupied by a robot does not necessarily means that a collision occur since the grid representation is an abstraction used by ARMCHAIR, i.e., in reality the human can be anywhere in the cell (not necessarily in its center), and it is reasonable to assume that the operator has the basic capability of avoiding a collision with a static object.”

4.2 Environment 2: Grouped Target Distribution

In the environment with grouped targets, the initial prediction from the learned policy (Fig. 7a) indicates that the

Fig. 6 Example simulation of ARMCHAIR in environment 1 (Sparse target distribution): four time steps of simulations 41 and 111 out of 1000. Solid and dashed lines represent the observed motion and predictions of each agent, respectively. The ARMCHAIR algorithm allows the replanning of the robot’s trajectories and target assignments, preventing any redundant visits when the human deviates from the initial prediction

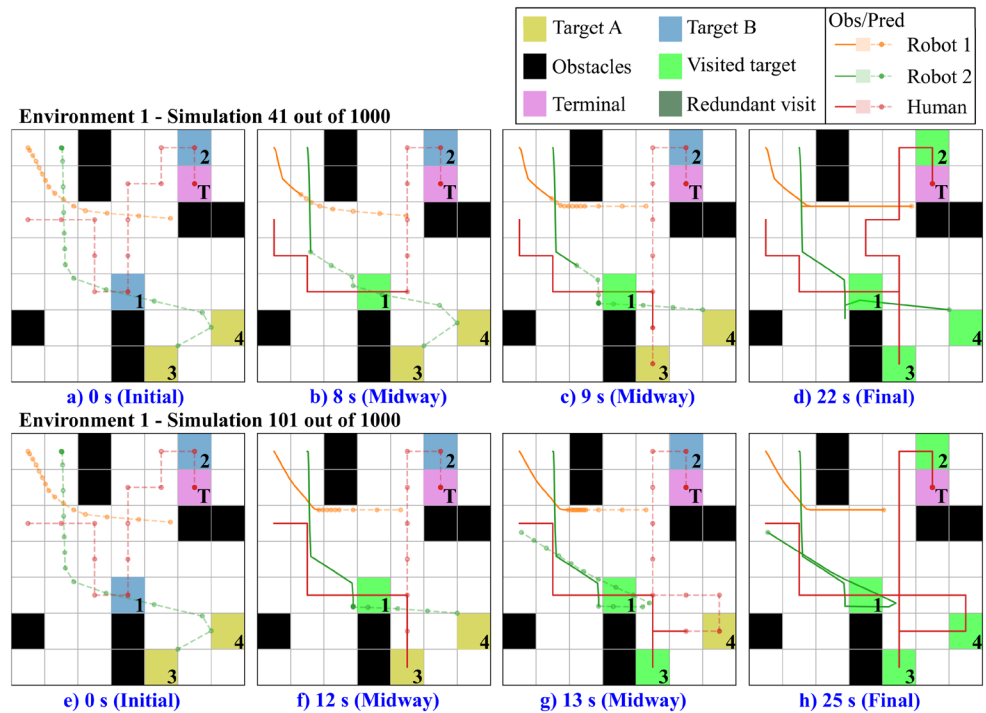


Table 1 Performance of ARMCHAIR compared to the baseline methods in environment 1 (Sparse target distribution). Each metric is represented by mean and 95% bootstrap confidence intervals (CI) for 1000 simulations

		Collisions	Disconnections	Targets	Redundant visits
ARM-CHAIR	Mean	0.001	0.0	4.0	0.0
	CI	(0.0,0.003)	(0.0,0.0)	(4.0,4.0)	(0.0,0.0)
Open-loop MIP	Mean	0.58	0.02	4.0	0.43
	CI	(0.53,0.63)	(0.01,0.03)	(4.0,4.0)	(0.40,0.47)
No robot support	Mean	N/A	N/A	2.7	N/A
	CI			(2.6,2.7)	

human is likely to visit all targets in the environment since they are close to each other. Consequently, the robots are dispatched only to maintain connectivity. After 12 seconds (Fig. 7b) we observed that the human deviates from the predictions, choosing to visit only targets 1 and 2 and finish the mission by reaching the terminal region. This result illustrates a distinct issue with the open-loop MIP approach: if in the initial prediction all targets are assigned to be likely visited by the human and a deviation occurs, the robot team is unable to update its decisions and motion to compensate for unexpected human behavior, yielding worse performance in the mission.

The simulation results in Fig. 8 show how ARMCHAIR is able to enhance the performance in the mission even

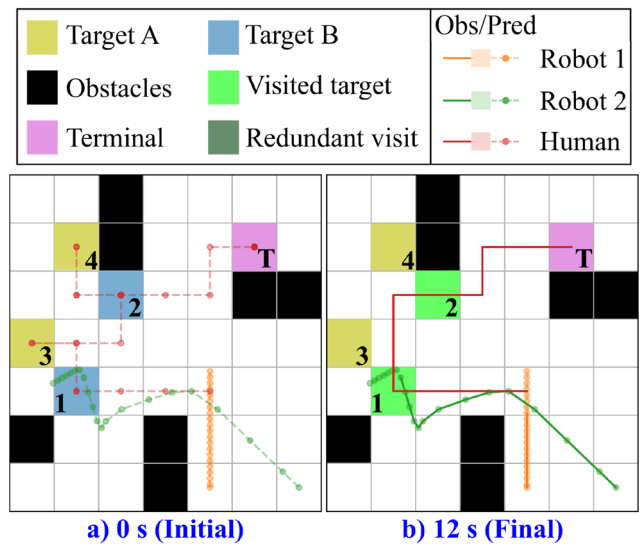


Fig. 7 Example simulation of the open-loop MIP baseline in environment 2 (Grouped target distribution): initial (a) and final (b) time steps for simulation 1 out of 1000. Solid and dashed lines represent the actual motion (observed) and predictions of each agent, respectively. The human prediction is computed using only the initial conditions (open-loop) and $\hat{\pi}$. The prediction suggests that all targets will be visited by the human and the robots are dispatched only to maintain connectivity. In (b) we observe that the human decides to ignore the targets of type a deviating from the initial prediction. As a result, the human-robot team provides inferior performance in the mission

when the human behavior predictions are imperfect. In simulation 25 out of 1000, the prediction that the human will visit all targets is maintained until time step 14 (Fig. 8b), and the robots operate very similarly to the open-loop

MIP case. Unexpectedly, the human decides not to capture target 2 even though it was properly positioned to do so. By moving upwards at time step 15 (Fig. 8c) the human indirectly declares its intention to finish the mission without visiting the remaining targets, and the predictions are readily updated with the human going directly towards the terminal cell. Even with this unexpected behavior of the human and the small window of time remaining in the mission, ARMCHAIR was able to update trajectories for robots 1 and 2 such that targets 4 and 2 were visited (Fig. 8d). The observed behavior illustrates how the ARMCHAIR algorithm enables the robot group to identify in real time whether the human intends to finish the mission alone or if their support is necessary to improve the attained performance in the task.

In simulation 184 out of 1000, the human presented a more conventional behavior, visiting the preferred targets of type B (Fig. 8f) but again deviating from the initial prediction. Fig. 8g shows that after visiting target 2, the human moves towards the terminal cell and the predictions are

updated accordingly. ARMCHAIR reacts by planning a trajectory for robot 2 in which it visits the remaining targets 3 and 4, while robot 1 maintains connectivity.

These observations are supported by the statistical results presented in Table 2. As in the first environment, ARMCHAIR was able to maintain the network connected in all simulations, whereas an average of 0.181 disconnections occurred when the open-loop MIP scheme was used. Although no redundant visits were registered with either approach, ARMCHAIR provides a higher average number of target reward collections (3.444) than the open-loop algorithm (2.384). Better performance in terms of average reward collection is achieved with the robot team as opposed to the case without support, for which the average was (2.393).

We also observed that ARMCHAIR performed worse in terms of collision avoidance in environment 2. This result stems from the increased complexity of the motion planning problem in this scenario due to the tight placement of

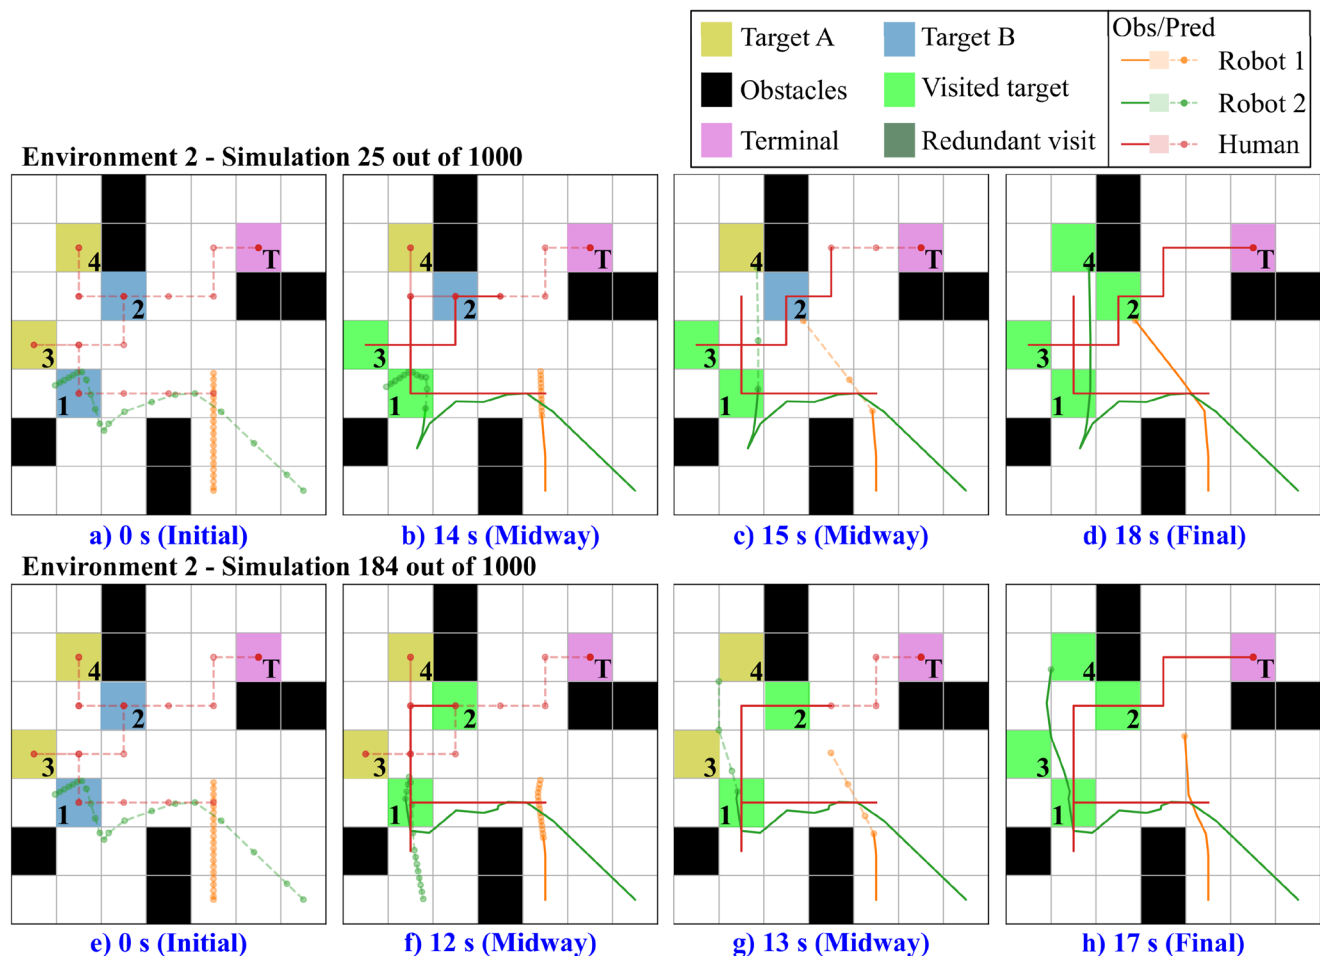


Fig. 8 Example simulation of ARMCHAIR in environment 2 (Grouped target distribution): four time steps of simulations 25 and 184 out of 1000. Solid and dashed lines represent the observed motion and predictions of each agent, respectively. In this case, the replanning capa-

bilities provided by the ARMCHAIR algorithm allow the robots to quickly adapt to compensate for an unexpected change in the behavior of the human, who decides to ignore targets that are easily accessible. As a result, all targets are visited at the end of the mission

Table 2 Performance of ARMCHAIR compared to the baseline methods in environment 2 (Grouped target distribution). Each metric is represented by mean and 95% bootstrap confidence intervals (CI) for 1000 simulations

			Collisions	Disconnections	Targets	Redundant visits
ARMCHAIR	Mean		0.013	0.0	3.444	0.0
	CI		(0.007,0.02)	(0.0,0.0)	(3.398,3.489)	(0.0,0.0)
Open-loop MIP	Mean		0.004	0.181	2.384	0.0
	CI		(0.001,0.008)	(0.157,0.205)	(2.354,2.415)	(0.0,0.0)
No robot support	Mean		N\A	N\A	2.393	N\A
	CI				(2.362,2.424)	

targets, which requires the group to operate in close proximity. Furthermore, the results indicate that the open-loop MIP scheme outperformed ARMCHAIR in terms of collision avoidance in this particular environment. This occurs because the robot team is not dispatched to visit any targets when the open-loop approach is employed (See Fig. 7). As a result, the robots typically remain far away from the human and collisions are less likely. Conversely, ARMCHAIR consistently updates the trajectories and decisions of the robots to improve performance. This is illustrated by its superior average of visited targets: 3.444 from ARMCHAIR against 2.384 from open-loop MIP. Consequently, they are occasionally steered towards targets closer to the human, increasing the likelihood of collisions.

5 Discussion and Conclusion

This paper presented ARMCHAIR, a novel architecture for human-robot collaboration that leverages the integration of AIRL and MPC. Through extensive Monte-Carlo simulation trials, we demonstrated that the scheme provides proper adaptive trajectories and decision-making, in the form of target allocations, for a group of robots supporting a human in an exploration mission.

ARMCHAIR was successful in both: a) coordinating the motion of the robot team, preventing most collisions and network disconnections; and b) autonomously identifying if the operator requires support based on the observed behavior and current mission circumstances. This was made possible by the adaptive decision-making provided by the MPC-MIP scheme, which improved performance by preventing overlaps in the allocation of targets to robots and the human.

The achieved results illustrate that ARMCHAIR provides distinct capabilities compared to most of the literature on human-robot collaboration with mobile robots. Excellent solutions for social navigation (coordination) between robots and humans are found in [6, 7, 36, 38–41]. However, the scenarios considered in this paper require the extra capability of harmonious collaboration towards the same global objective. The ensuing joint problem of human-robot trajectory planning and decision-making has been successfully addressed in [42] and [44], with the latter work thoroughly

demonstrating its effectiveness through experiments with humans. However, both of these solutions are still dependent on frequent human inputs. In [44], for example, the necessity of the human operator to divide their attention between sending messages to the robot team and working on the mission was reported as an open issue. ARMCHAIR addresses this gap by actively predicting and adapting to the operator based solely on the observed human behavior. Thus, our approach has an extra degree of autonomy that allows the operator to concentrate on performing the mission, without the necessity of explicitly directing the actions of the robot team.

Due to the novelty of the proposed integration, ARMCHAIR is still bound by a few limitations. For example, its receding horizon formulation lacks a formal demonstration of recursive feasibility and has no active method to handle unfeasibility, simply commanding the robots to stop moving until the problem becomes feasible again. Regarding computational complexity, since the human is introduced into the MILP as another agent (over which the optimization has no authority), the model's complexity increases similarly to the cases where pure robot teams were used. The tractability of such problems has been experimentally demonstrated for teams of five robots [5]. However, this centralized formulation may quickly become intractable for receding horizon schemes due to poor scalability capabilities. The simulation-based evaluations presented illustrate that the integration of MPC-MIP and AIRL for human-robot teaming algorithms is promising. However, further investigation with thorough experimental demonstrations involving real humans, as performed in [44] considering similar exploration missions, is still required.

In future works, the investigation of conditions for formal demonstrations of theoretical properties, in particular recursive feasibility, should be addressed. Furthermore, theoretical guarantees of collision avoidance should be investigated with the introduction of methods such as constraint tightening [54] or contingency MPC [50, 55]. Reactive strategies, such as inner potential field control laws, could be devised to handle the occasional collisions that occur due to infeasible optimization problems. The human prediction model can be refined to address more complex concepts, such as intentions and visual perceptions; the modeling of the human motion could also be improved, allowing for

variable velocities and smoother movement predictions that would reduce the conservatism in the MRS trajectories. Alternative methodologies for robustness could also be explored, reducing conservatism by tolerating risk [56, 57]. Distributed MIP schemes could be investigated to address the scalability issues.

The use of a) demonstrations with direct interaction between robots and humans and b) insights from psychology and cognitive modeling can enable the development of a model with enhanced capabilities, such as predictions considering learned human biases. Longitudinal trials using real humans and robots are necessary to further validate the effectiveness of the proposed architecture due to issues such as co-adaptation [58, 59].

Overall, this work demonstrates that the integration of AIRL and MPC-MIP in ARMCHAIR yields a platform that already provides notable capabilities for human-robot collaboration and can be further refined to become an enabling technology for safe, harmonious, and efficient human-robot teaming in the future.

Appendix A Synthetic Human Agent

As a surrogate for the human, we employ a reinforcement learning agent trained using Proximal Policy Approximation (PPO) on the ground truth rewards detailed in Table A1. Its actions represent movement in the *north*, *south*, *west*, and *east* directions, as well as a *collection* that allows the agent to capture target rewards. The relative value of the targets to the human is represented by the rewards $R_A > 0$ and $R_B > 0$, with $R_B > R_A$. Thus, it has a preference for the type B. The division of targets into two types without a specific meaning was deliberate, since we aimed to achieve a general model where the particular meaning of a region can be determined depending on the context or application given by the user. They can be interpreted in many ways considering a human: one region is a risky task which the human prefers to avoid, whereas the other is a safe one; one can be a laborious task that the human also avoids if in a hurry, while the other requires less work. The rewards can only be collected at most once. There are three types of penalization: collisions with walls and obstacles, movement, and failure in reaching the terminal region in the required number of steps, they are represented by the quantities $R_C < 0$, $R_M < 0$, and $R_F < 0$, respectively.

Table A1 Ground-truth rewards and penalizations for the synthetic human agent

	R_A	R_B	R_C	R_M	R_F
Value	0.5	1.0	-1.0	-0.1	-20

Notice that by penalizing movement we intend to introduce time pressure into the cost that drives the behavior of the synthetic human such that its decisions are not to trivially reach all targets in all situations. This penalization models a sense of urgency that can be potentially present in the demeanor of the human depending on the environment and circumstances.

We employ the same PPO setup as in [32] with $\gamma = 0.999$. The training employs new randomly generated environments similar to the one depicted in Fig. 6. For each environment, the maximum number of targets is $n_t \sim \mathcal{U}(0, 4)$. The number of target of type B and A is computed as $n_B \sim \mathcal{U}(0, n_t)$ and $n_A = n_t - n_B$, respectively. The initial position is also randomized for each scenario, being uniformly distributed over all cells that are not targets or obstacles. The position of the obstacles and terminal region remains fixed. The human does not consider the robot's support explicitly but rather reacts indirectly, through their impact on the environment, e.g., capture of targets.

After running the algorithm for 16×10^5 steps, the synthetic agent's policy $\pi^* : \mathcal{S} \rightarrow \Delta_{\mathcal{N}}$ is computed, where $\Delta_{\mathcal{A}}$ is a probability distribution over all actions. Finally, π^* is used to generate a set of 8000 trajectories, \mathcal{D} , representing the human navigation in different environments.

Appendix B AIRL training

Table B2 presents the hyperparameters used in the AIRL training. The employed AIRL hyperparameters are slight modifications, determined through trial-and-error, from the parameters employed in [32]. We evaluate the performance of the recovered policy by comparing the trajectories generated by $\hat{\pi}$ and π^* in 1000 new environments built randomly using the same procedure discussed in Appendix A. Table B3 summarizes the results in terms of average length of the trajectory, number of targets of type A and B visited, collisions, and return over the ground truth rewards considering the trajectories achieved using the synthetic human agent and the recovered policies.

The results show that the policy recovered by the AIRL algorithm generates useful predictions of the human agent. The resulting average trajectory lengths are very similar,

Table B2 AIRL hyperparameters

Hyperparameter	Value
Learn. rate discr.	$4e-4$
Learn. rate gen.	$4e-4$
Batch size discr.	256
Batch size gen.	8
Environment steps	$3.5e6$
ϵ -clip	0.1
γ	0.999
PPO epochs	5

Table B3 Comparison between recovered and surrogate human policies

Synthetic human agent		
	Average	Confidence Interval
Trajectory length	12.7	(12.3, 13.0)
Tar. A collection	0.68	(0.63, 0.73)
Tar. B collection	1.09	(1.02, 1.16)
Collisions	0.0	(0.00, 0.00)
Return	0.05	(0.01, 0.09)
AIRL recovered		
Trajectory length	12.3	(11.9, 12.7)
Tar. A collection	0.61	(0.56, 0.66)
Tar. B collection	0.81	(0.75, 0.87)
Collisions	0.0	(0.00, 0.00)
Return	-0.85	-(1.01, 0.70)

suggesting that the prediction model accurately captures the “urgency” of the human in finishing the mission. The obstacle avoidance capability is also appropriately emulated with no collisions occurring in both cases. However, there are mismatches in the target collection averages, in particular targets of type B. The recovered model shows a less greedy behavior, with smaller target collection averages.

Real-world deployment of a prediction model based on $\hat{\pi}$ would require its validation as a generalizable model of human behavior through trials with humans. This experiment is beyond the scope of this paper and is left for future work. In terms of real-world assets required to train the AIRL network, a sizable dataset of human operator’s experiences in the task at hand would have to be available. As discussed in Appendix A, 8000 distinct trajectories were generated with π^* (the surrogate human model) to train our AIRL prediction network. It must be noted though that the literature presents demonstration-efficient formulations of AIRL [30] that could substantially decrease the number of required trajectories.

Author Contributions ACN devised the research question, performed the literature review, prepared all figures and tables, wrote the new code required, devised the simulation environments, evaluated results, and wrote all sections. LCS, AZ, MROAM, and RJMA provided support in determining the research question and performing the literature review, suggested the addition of figures, helped in the determination of simulation environments and evaluation methodology, and contributed to the writing and proofreading of all sections.

Funding This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001. Marcos Maximo is partially funded by CNPq – National Research Council of Brazil through the grant 307525/2022–8. The participation of Luciano Cavalcante Siebert and Arkady Zgonnikov in this research was partially supported by TAILOR, a project funded by EU Horizon 2020 research and innovation programme under GA No 952215.

Data Availability All the algorithms and datasets related to the results in this paper are available at the repository.³

Declarations

Ethics Approval Not applicable.

Conflict of Interest The authors declare that they have no conflict of interest or competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

1. Annaswamy AM, Johansson KH, Pappas GJ (2023) Control for societal-scale challenges: road map 2030. IEEE Control Systems Society Publication
2. Liang C-J, Wang X, Kamat VR, Menassa CC (2021) Human–robot collaboration in construction: classification and research trends. *J Constr Eng Manag* 147(10):03121006
3. Gervasi R, Barravecchia F, Mastrogiacomo L, Franceschini F (2023) Applications of affective computing in human-robot interaction: state-of-art and challenges for manufacturing. *Proc Inst Mech Eng, Part B: J Eng Manuf* 237:815–832
4. Mavrogiannis C et al. (2023) Core challenges of social robot navigation: a survey. *J Hum-Robot Interact* 12
5. Caregnato-Neto A, Maximo MROA, Afonso RJM (2023) Real-time motion planning and decision-making for a group of differential drive robots under connectivity constraints using robot MPC and mixedinteger programming. *Adv Rob* 37:356–379
6. Ozkan MF, Ma Y (2022) Distributed stochastic model predictive control for human-leading heavy-duty truck platoon. *IEEE Trans Intell Transp Syst Intell Transp Syst* 23:16059–16071
7. Zhu H, Claramunt FM, Brito B, Alonso-Mora J (2021) Learning interaction-aware trajectory predictions for decentralized multi-robot motion planning in dynamic environments. *IEEE Robot Autom Lett* 6:2256–2263
8. Camacho EF, Bordons C, Camacho EF, Bordons C (2007) Model predictive controllers. Springer
9. Liu Y, Liu K, Wang G, Sun Z, Jin L (2023) Noise-tolerant zeroing neurodynamic algorithm for upper limb motion intention-based human–robot interaction control in non-ideal conditions. *Expert Syst Appl* 213:118891
10. Zheng H, Smereka MJ, Mikulski D (2023) Bayesian optimization based trust model for human multi-robot collaborative motion tasks in offroad environments. *Int J Soc Robot* 1181–1201
11. Schumann JF, Srinivasan AR, Kober J, Markkula G, Zgonnikov A (2023) Using models based on cognitive theory to predict human behavior in traffic: a case study. In *IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*, pp 5870–5875
12. Liu W, Liang X, Zheng M (2023) Taskconstrained motion planning considering uncertainty-informed human motion prediction for human–robot collaborative disassembly. In *IEEE/ASME Transactions on Mechatronics*, pp 1–8
13. Jin Z, Liu A, Zhang W-A, Yu L, Su C-Y (2023) A learning based hierarchical control framework for human–robot collaboration. *IEEE Trans Autom Sci Eng* 20:506–517
14. Fuchs A, Passarella A, Conti M (2023) Modeling, replicating, and predicting human behavior: a survey. *ACM Trans Auton Adapt Syst* 18

³ https://gitlab.com/caregnato_netto_open/ARMCHAIR.

15. Ullman D, Malle BF (2017) Human-robot trust: just a button press away. In Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction, pp 309–310
16. Cutler CR, Ramaker BL (1980) Dynamic matrix control: a computer control algorithm. In No. 17 in Joint Automatic Control Conference, vol 72
17. Ferranti L, Lyons L, Negenborn RR, Keviczky T, Alonso-Mora J (2023) Distributed nonlinear trajectory optimization for multi-robot motion planning. *IEEE Trans Control Syst Technol* 31:809–824
18. Ioan D, Prodan I, Olaru S, Stoican F, Niculescu S (2020) Mixed-integer programming in motion planning. *Annu Rev Control*
19. Nascimento IBP, Ferramosca A, Pimenta LCA, Raffo GV (2019) NMPC strategy for a quadrotor UAV in a 3D unknown environment. In 19th International Conference on Advanced Robotics (ICAR), pp 179–184
20. Ramalho GM, Carvalho SR, Finardi EC, Moreno UF (2018) Trajectory optimization using sequential convex programming with collision avoidance. *J Control Autom Electr Syst* 29:318–327
21. Richards A, How JP (2002) Aircraft trajectory planning with collision avoidance using mixed integer linear programming. In Vol. 3 of Proceedings of the 2002 American Control Conference, 1936–1941, Anchorage, USA
22. Afonso RJM, Maximo MROA, Galvão RKH (2020) Task allocation and trajectory planning for multiple agents in the presence of obstacle and connectivity constraint with mixed-integer linear programming. *Int J Robust Nonlinear Control* 30:5464–5491
23. Ng A, Russell S (2000) Algorithms for inverse reinforcement learning, vol. 1 of ICML. p 2
24. Town J, Morrison Z, Kamalapurkar R (2023) Pilot performance modeling via observerbased inverse reinforcement learning, pp 2307.13150
25. Sadigh D, Landolfi N, Sastry SS, Seshia SA, Dragan AD (2018) Planning for cars that coordinate with people: leveraging effects on human actions for planning and active information gathering over human internal state. *Auton Robots* 42:1405–1426
26. Avaei A, van der Spaa L, Peternel L, Kober J (2023) An incremental inverse reinforcement learning approach for motion planning with separated path and velocity preferences. *Robotics* 12
27. Perrusquía A, Guo W (2023) Drone’s objective inference using policy error inverse reinforcement learning. In *IEEE Transactions on Neural Networks and Learning Systems*, pp 1–12
28. Fu J, Luo K, Levine S (2018) Learning robust rewards with adversarial inverse reinforcement learning. In *International Conference on Learning Representations*
29. Ziebart BD, Maas AL, Bagnell JA, Dey AK et al. (2008) Maximum entropy inverse reinforcement learning. In Vol. 8 of Proc AAAI Conf Artif Intell, Chicago, IL, USA, 1433–1438
30. Sestini A, Kuhnle A, Bagdanov AD (2021) Demonstration-efficient inverse reinforcement learning in procedurally generated environments. In 2021 IEEE Conference on Games (CoG), IEEE, 1–8
31. Yu L, Song J, Ermon SC, K (2019) Multi-agent adversarial inverse reinforcement learning. In: Salakhutdinov R (eds) Chaudhuri K & Salakhutdinov R (eds) Proceedings of the 36th International conference on machine Learning, vol. 97 proceedings of machine learning research. PMLR, pp 7194–7201
32. Peschl M, Zgonnikov A, Oliehoek FA, Siebert LC (2022) MORAL: aligning AI with human norms through multi-objective reinforced active learning
33. Cavalcante Siebert L et al. (2023) Meaningful human control: actionable properties for AI system development. *AI Ethics* 3:241–255
34. Wang P, Liu D, Chen J, Li H, Chan C-Y (2021) Decision making for autonomous driving via augmented adversarial inverse reinforcement learning. In 2021 IEEE International Conference on Robotics and Automation (ICRA), 1036–1042, IEEE
35. Lee K, Isele D, Theodorou EA, Bae S (2022) Spatiotemporal costmap inference for MPC via deep inverse reinforcement learning. *IEEE Robot Autom Lett* 7:3194–3201
36. Lu J, Hossain S, Sheng W, Bai H (2023) Cooperative driving in mixed traffic of manned and unmanned vehicles based on human driving behavior understanding. In 2023 IEEE International Conference on Robotics and Automation (ICRA), pp 3532–3538
37. Jansma W, Trevisan E, Serra-Gómez Á, Alonso-Mora J (2023) Interaction-aware sampling-based MPC with learned local goal predictions 2309.14931
38. Boldrer M, Antonucci A, Bevilacqua P, Palopoli L, Fontanelli D (2022) Multiagent navigation in human-shared environments: a safe and socially-aware approach. *Robot Auton Syst* 149:103979
39. Kim B, Pineau J (2016) Socially adaptive path planning in human environments using inverse reinforcement learning. *Int J Soc Robot* 8:51–66
40. Galvan M, Repiso E, Sanfeliu A (2020) Robot navigation to approach people using bspline path planning and extended social force model. In *Robot 2019: Fourth Iberian Robotics Conference*, pp 15–27
41. Wang C, Li Y, Ge SS, Lee TH (2016) Adaptive control for robot navigation in human environments based on social force model. In 2016 IEEE International Conference on Robotics and Automation (ICRA), pp 5690–5695
42. Hooi Chan T et al. (2023) A robotic system of systems for human-robot collaboration in search and rescue operations. In 2023 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM), pp 878–885
43. Wampler J, Li B, Moscicki T, von Ellenrieder KD (2017) Towards adjustable autonomy for human-robot interaction in marine systems. *Oceans 2017 - Aberd* 1–7
44. Dalmaso M et al. (2023) Shared task representation for human-robot collaborative navigation: the collaborative search case. *Int J Soc Robot*
45. Richards A, How J (2005) Mixed-integer programming for control. In Proceedings of the 2005 American Control Conference, vol 4. pp 2676–2683
46. Baotic M (2009) Polytopic computations in constrained optimal control. *Automatika* 50:119–134
47. You Y, Thomas V, Colas F, Alami R, Buffet O (2023) Robust robot planning for humanrobot collaboration 2302.13916
48. Kahneman D (2003) Maps of bounded rationality: psychology for behavioral economics. *The Am Econ Rev* 93:1449–1475
49. Shah R, Gundotra N, Abbeel P, Dragan A (2019) On the feasibility of learning, rather than assuming, human biases for reward inference. In *International Conference on Machine Learning*, PMLR, 5670–5679
50. Schweidel KS, Koehler SM, Desraj VR, Barić M (2022) Driver-in-the-loop contingency MPC with invariant sets. In 2022 European Control Conference (ECC), pp 808–813
51. Löfberg J (2004) Yalmip: a toolbox for modeling and optimization in MATLAB. In Proc. of 2004 IEEE Int. Conf. on Robotics and Automation, Taipei, Taiwan
52. Kvasnica M, Grieder P, Baotic M (2004) Multi-Parametric Toolbox (MPT). <http://control.ee.ethz.ch/%20mpt/>
53. Gurobi Optimization, L (2021) Gurobi optimizer reference manual. <http://www.gurobi.com>
54. Richards A, How JP (2006) Robust variable horizon model predictive control for vehicle maneuvering. *Int J Robust Nonlinear Control* 16:333–351
55. Alsterda JP, Brown M, Gerdes JC (2019) Contingency model predictive control for automated vehicles 717–722

56. Pinto SC, Afonso RJ (2017) Risk constrained navigation using milp-mpc formulation. *ifac-Papersonline* 50:3586–3591. 20th IFAC World Congress
57. Blackmore L, Ono M, Williams BC (2011) Chance-constrained optimal path planning with obstacles. *IEEE Trans Robot* 27:1080–1094
58. van Zoelen EM, van den Bosch K, Rauterberg M, Barakova E, Neerinx M (2021) Identifying interaction patterns of tangible coadaptations in human-robot team behaviors. *Front Psychol* 12:645545
59. Melman T, de Winter J, Abbink D (2017) Does haptic steering guidance instigate speeding? a driving simulator study into causes and remedies. *Accid Anal Prev* 98:372–387

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

Angelo Caregnato-Neto received the M.Sc. (2018) and Ph.D. (2024) degrees in electronic and computer engineering from the Aeronautics Institute of Technology (ITA), São José dos Campos, Brazil. During his Ph.D., he was a visiting researcher with the Interactive Intelligence Group at Delft University of Technology, Delft, The Netherlands. He is currently a postdoctoral fellow at the State University of Campinas (UNICAMP), Campinas, Brazil. His research interests include optimization-based control, motion planning, and decision-making for multi-robot systems.

Luciano Cavalcante Siebert is Assistant Professor of Responsible AI at TU Delft's Interactive Intelligence group. His interdisciplinary research focuses on enabling meaningful human control over AI systems by integrating ethical and human behaviour theories into interactive approaches to align AI with human values. He is co-founder and director of the Centre for Meaningful Human Control and co-director of the AiBLE lab. He was a task leader in the EU TAILOR Network on Trustworthy AI, editor of the first Research Handbook on Meaningful Human Control (Elgar Publishing), and is a member of the ELLIS Society.

Arkady Zgonnikov is an assistant professor in the Department of Cognitive Robotics, TU Delft. He received his M.Sc. in applied mathematics from Saint-Petersburg State University, Russia, and Ph.D. in computer science and engineering from the University of Aizu, Japan. His work centers around modeling human behavior in traffic interactions and applying these insights to the development and evaluation of automated vehicles. He is also actively advancing the concept of meaningful human control over intelligent systems, both as a theoretical framework and as a practical principle for AI design and assessment.

Marcos R. O. A. Maximo received the BSc degree in Computer Engineering (with Summa cum Laude honours) and the MSc and PhD degrees in Electronic and Computer Engineering from Aeronautics Institute of Technology (ITA), Brazil, in 2012, 2015 and 2017, respectively. Maximo is currently a Professor at ITA, where he is a member of the Autonomous Computational Systems Lab (LAB-SCA) and leads the robotics competition team ITAndroids. He is especially interested in humanoid robotics. His research interests also include mobile robotics, dynamical systems control, and artificial intelligence.

Rubens J. M. Afonso Rubens Junqueira Magalhães Afonso received his Ph.D. (2015) and M.Sc. (2012) degrees in Electronic and Computer Engineering and his B.Sc. (2009) degree in Electronic Engineering from the Aeronautics Institute of Technology (ITA), Brazil. His research interests include model predictive control theory and practice, robust and optimal control, and mixed-integer programming.