

Learning-Based Safety-Aware Task Scheduling for Efficient Human-Robot Collaboration

Abstract—Ensuring human safety in collaborative robotics can compromise efficiency because traditional safety measures increase robot cycle time when human interaction is frequent. This paper proposes a safety-aware approach to mitigate efficiency losses without assuming prior knowledge of safety logic. Using a deep-learning model, the robot learns the relationship between system state and safety-induced speed reductions based on execution data. Our framework does not explicitly predict human motions but directly models the interaction effects on robot speed, simplifying implementation and enhancing generalizability to different safety logics. At runtime, the learned model optimizes task selection to minimize cycle time while adhering to safety requirements. Experiments on a pick-and-packaging scenario demonstrated significant reductions in cycle times.

I. INTRODUCTION

Collaborative robotics is expected to be an enabling technology of the factory of the future, offering benefits such as increased workspace efficiency and the removal of restrictive robot cages. However, robots operating alongside humans in shared workspaces pose significant safety challenges. Ensuring compliance with safety standards (e.g., ISO 10218-2 [1]) is essential to prevent harm, but often results in efficiency trade-offs. For instance, commonly used static safety zones cause robots to slow down or stop when a human enters a predefined area, leading to increased cycle times during frequent interactions.

An intuitive way to reduce safety slowdowns is to synchronize the robot and the human's operations so that the agents avoid entering low-speed regions at the same time. Existing methods to address this challenge have focused on optimizing robot scheduling [2], human-aware path planning [3, 4], or trajectory adjustments based on operator intent [5–7]. While effective, these approaches rely on a priori knowledge of the safety logic, intertwining control and safety design. For example, they require that the developer knows how the risk assessment phase will be carried out and, vice versa, that the system integrator can act on the design of the motion planner. However, risk assessment and control development typically occur independently and at a different time [8].

In practice, obtaining the exact safety configuration is often difficult even after the deployment. While in theory one could extract it from risk assessment documents or inspect controller configurations, in reality safety functions are often manufacturer-specific and hidden from the user. In other cases, the configuration can technically be retrieved, but converting it into a usable set of equations for control purposes requires significant manual effort. By contrast, learning the relationship between system state and safety-induced slowdowns directly from execution data avoids this burden and provides a more practical and generalizable alternative.

Motivated by these limitations, we propose a data-driven, safety function-agnostic approach to enhance collaborative robot efficiency by optimizing robot task scheduling without requiring explicit knowledge of safety logic. The system uses a deep-learning model to learn the relationship between human-robot interaction modes and robot speed reductions based on process data. This data-driven approach has multiple advantages: i) It allows integration with existing cells where the safety configuration is inaccessible; ii) It adapts naturally to changes in task or layout after retraining, especially considering that training data can be collected during normal operations; iii) Unlike model-based methods, which require predicting human motion at future time instants, our approach directly learns the impact of human activity on robot speed. At runtime, the learned model guides task selection to minimize cycle time. To do so, we propose two action selection algorithms that leverage the deep-learning model to minimize expected speed reductions. The first algorithm predicts the robot's slowdown for all available actions and chooses the one with the minimum impact on the robot's speed. The second algorithm uses a Monte Carlo approach to plan for longer prediction horizons.

Experiments on a pick&packaging application showed a significant improvement in the robot cycle time and mitigation of safety speed reductions for both the proposed algorithms.

II. RELATED WORKS

Safety standards such as the ISO 10218-2 [1] define frameworks for deploying collaborative robotics systems. These standards provide general guidelines for human safety, requiring system integrators to perform case-specific risk analysis and mitigation based on ISO 12100. Techniques such as Speed and Separation Monitoring (SSM) and Power and Force Limitation (PFL) ensure safety by applying speed reduction rules. However, the speed reduction significantly impacts robot cycle times, leading to inefficiencies in collaboration.

Previous works addressed this problem by acting on the robot's control pipeline [2–4, 6, 7, 9–11].

For example, [3] and [4] proposed a safety-aware path planner minimizing the expected robot's execution time considering also safety slowdowns. They define cost functions modeling the SSM and PFL speed reduction rules defined in ISO 10218-2 and search for a path minimizing such cost functions. At a lower control level, [12] uses model predictive control to find a feasible trade-off between path following and compliance with safety constraints. As for scheduling, [2] proposes a scheduler to prevent stops owed to simultaneous access to the same area. To do so, they assign penalty terms to each region of the workspace according to ISO 10218-2. Other works merge the psycho-physical state of the operator with

safety awareness [13, 14]. These methods translate nominal speed reduction values from ISO 10218-2 into an optimization problem. However, ISO 10218-2 defines lower bounds on the human-robot distance, but such values need to be inflated depending on the perception system and the case-specific risk analysis.

Safety-aware methods are tightly coupled with the prediction of the human motion as they shall anticipate the human actions [15–18]. Depending on the length of the prediction horizon, they integrate human prediction in different ways. Short-term prediction is usually addressed with physics-based models [19] or filtering techniques [20]. Long-term prediction usually leverages additional information on the context and the human goal or use agglomerated estimates of the human occupancy [21]. Recently, deep-learning approaches have become the most widespread for long-term, context-aware prediction [22–24].

Our approach differs from prior works in two key aspects. First, it does not assume prior knowledge of safety logic. Instead, it learns the correlation between system states and safety-induced speed reductions from process data, enabling generalization across different implementations. Second, rather than explicitly modeling the human movements, our approach directly predicts the robot’s slowdown, simplifying the model while providing a meaningful prediction for decision-making.

III. PROBLEM STATEMENT

We consider a robot and an operator working in a shared environment. We assume a safety system acting as follows: when the human is not in a cooperative area, the robot moves at a nominal speed; when the human enters a collaborative area, it reduces its speed to a percentage of the nominal speed. A speed reduction function computes the speed percentage based on the measured system states. For example, common implementations slow the robot down proportionally to the inverse of the human-robot relative distance and velocity or assign fixed values when the human enters pre-defined areas.

In our setup, the human and the robot take actions asynchronously by choosing online from a pool of available actions. Our goal is that the robot selects the action that mitigates the safety speed reduction and, thus, reduces cycle and idle times.

Let $x_r \in X_r$ and $x_h \in X_h$ be the robot and human’s states, $u_r \in U_r$ the robot’s action, where X_r , X_h , U_r are the corresponding state and actions apces. We denote the safety speed reduction function by $s(x_r, x_h) \in [0, 1]$. That is, the actual robot speed is equal to the nominal speed (i.e., without human presence) multiplied by $s(x_r, x_h)$. According to the scenario above, we assume (i) s has an unknown shape but its value at time \bar{t} can be observed and (ii) x_r , x_h , u_r are observable. We do not make assumptions about the shape of s (except that it depends on x_r and x_h).

A robot’s action, u_r , consists of moving the robot to a goal state, g_r . For this reason, we can define a function that associates u_r with g_r , such that:

$$g_r = \text{getRobotGoal}(u_r) \quad (1)$$

where U_r is the set of all robot actions. Similarly, we assume the human task is composed of sub-operations, each one associated

with a goal position. We denote the set of human goals by $\mathcal{G}_h = \{\mathcal{G}_{h,j}\}$ and the current human goal by G_h . To account for human variability, $\mathcal{G}_{h,j}$ is a random variable such that

$$\mathcal{G}_{h,j} \sim \mathcal{N}(\mu_j, \sigma_j) \quad (2)$$

where μ_j and σ_j are the mean and standard deviation of the Gaussian distribution. We assume it is possible to retrieve G_h for the current human operation (e.g., through a task segmentation algorithm that identifies the current human operation or explicit feedback on an HMI).

Our goal is to choose the robot’s next action, u_r^* , that minimizes the safety speed reduction intervention. This problem can be framed as a Markov Decision Process with an exogenous input (Exo-MDP). An Exo-MDP is a particular type of Markov Decision Process (MDP), where the process can be factorized in an endogenous and an exogenous part. In our problem, the robot is the endogenous part and the human is the exogenous input. In particular, we define the Exo-MDP as a tuple (X_r, U_r, X_h, P, R) where:

- X_r is the endogenous state space (i.e., the robot position and/or velocity and/or auxiliary variables);
- U_r is the robot action space;
- X_h is the exogenous input (i.e., the human position and/or velocity)
- $P(x'_r|x_r, x_h, u_r)$ is a probability function that models the transition probability from one endogenous state to another;
- $R(x_r, x_h)$ is a reward function proportional to the safety speed scaling $s(x_r, x_h)$.

If the robot selects an action $u_r(k)$ at a decision step k , it will reach the goal $g_r(k)$ at the next decision step. Therefore, we can consider the robot dynamics as deterministic (i.e., $P = 1$) and obtain

$$x_r(k+1) = \text{getRobotGoal}(u_r(k)) \quad (3)$$

Note that decision steps are incremental values that correspond to the time instant when the robot ends the ongoing operation and needs to select the next one. We can define a function $\tau(k)$ to map decision steps into actual time instants.

The safety scaling factor plays the role of the reward function. At decision step k , we aim to maximize the safety scaling factor over a time horizon w , i.e.,

$$u_r^*(\tau(k)) = \underset{u_r \in U_r(x_r(k))}{\text{argmax}} \mathbb{E} \left[\int_{\tau(k)}^{\tau(k)+w} s(x_r(t), x_h(t)) dt \right] \quad (4)$$

where $U_r(k)$ is the set of available actions at step k . Note that $s \in [0, 1]$ and the closer is s to 1, the smaller is the robot slowdown. However, in our problem, s has an unknown shape.

Our goal is to learn the correlation between the system’s state and the safety effect from data and use it to approximately solve (4) online, every time the robot has to select a new action.

IV. METHOD

We tackle the problem above in two phases. In the learning phase, we collect data from randomized task executions (i.e., let the robot take action randomly) and record x_r , x_h , g_r , G_h , and $s(x_r, x_h)$. We use the resulting dataset to train a neural

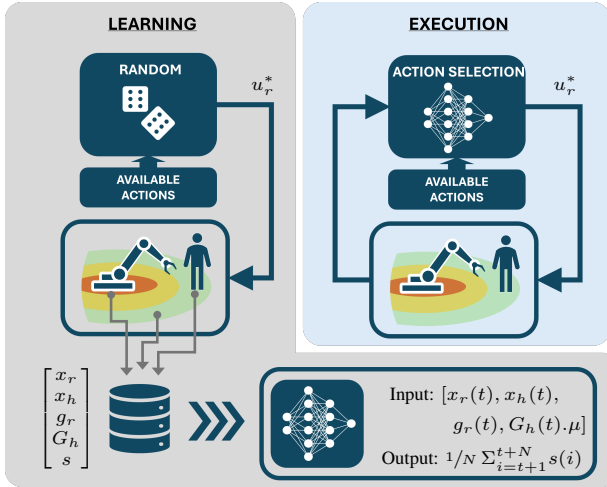


Fig. 1: Working principle of the proposed approach.

network that predicts the average values of s over a future time window.

In this work, robot and human's states and goals are Cartesian positions of the robot's end-effector and the human centroid, respectively. However, the method can accommodate for different or more complex representations (e.g., position, velocity, and other process variables).

In the execution phase, the robot has to choose its next action among the available ones. We propose two approximated methods using the neural network to predict the slowdown associated with each available action and choose the action with the less severe slowdown. The working principle of the approach is illustrated in Fig. 1.

A. Learning pipeline

We aim to predict the average value of the scaling function over a given time window. We consider the following features: the robot and human's states, x_r , x_h , and the robot and human's goals at the current time, g_r and G_h . As an example, in our experiments, x_r and x_h are the (x,y,z) position of the robot end-effector and the human centroid, g_r is the final position of the robot trajectory, and g_h is the mean value of the human task final position.

Our learning architecture uses a feed-forward neural network. The loss function is the Mean Squared Error (MSE) of the average scaling factor over N samples, with sampling period Δt , i.e., $L = \|y - \hat{y}\|^2$, where $\hat{s}(t)$ is the observed scaling value at time t , \hat{y} is the prediction of the network and

$$y = \frac{1}{N+1} \sum_{i=0}^N \hat{s}(\bar{t} + i\Delta t). \quad (5)$$

While our approach can generalize to any safety function s , it is worth analyzing the specific case where s is a staircase function (see Sec. V for an example). As mentioned in Sec. I, this choice is common in industrial applications, where the robot's speed scaling is set to fixed values based on the human

position and robot's positions. In this case, the safety function can be modeled as

$$s(x_r, x_h) = \begin{cases} \tilde{s}_1, & \text{if } d(x_r, x_h) \in D_1 \\ \tilde{s}_2, & \text{if } d(x_r, x_h) \in D_2 \\ \dots & \\ \tilde{s}_K, & \text{if } d(x_r, x_h) \in D_K \end{cases} \quad (6)$$

where $\tilde{s}_i \in [0, 1]$ are constant scaling values, $\tilde{s}_i < \tilde{s}_{i+1}$, $d(x_r, x_h)$ is a generic function in the robot and human state, and D_i is a set of states whose membership determines the scaling value. The average scaling over $[\bar{t}, \bar{t} + N\Delta T]$ is

$$\bar{s}_{[\bar{t}, \bar{t} + N\Delta T]} = \frac{1}{K} \sum_{i=1}^K \alpha_i \tilde{s}_i, \quad \text{with } \sum_{i=1}^K \alpha_i = 1 \quad (7)$$

where α_i is the fraction of $[\bar{t}, \bar{t} + N\Delta T]$ where $\hat{s}(t) = \tilde{s}_i$.

Building on this observation, we use a network whose last hidden layer has width equal to K and a Softmax activation function. Then, we use a linear output layer and obtain

$$\hat{y} = \beta_0 + \sum_{i=1}^K \beta_i \hat{y}'_i \quad (8)$$

where $\hat{y}' \in [0, 1]^K$ is the output of the last hidden layer, $\beta \in \mathbb{R}^{K+1}$ is the weight vector of the output layer, and $\sum_{i=1}^K \hat{y}'_i = 1$ thanks to the Softmax function. The intuition behind this design is that the hidden layers will learn a feature vector, \hat{y}' , proportional to α while the last hidden layer will learn how to weigh each contribution and yield the average scaling.

Because the network does not know K beforehand, we estimate it directly through data clustering as follows:

- 1) We consider the collected values of the scaling factor $\{\hat{s}_i\}$, with $\hat{s}_i \in [0, 1]$.
- 2) We apply a clustering algorithm on $\{\hat{s}_i\}$ to determine the number of clusters. Most clustering algorithms automatically output the optimal number of clusters. For those that do not (e.g., K-means), it can be computed using the Silhouette or the elbow method [25].
- 3) We approximate the number of steps K as equal to the number of clusters.

Finally, we set the width of the last hidden layer equal to K . The resulting network is exemplified in Fig. 2 for $K=5$.

This approach works only if the safety function is a staircase function. Otherwise, the learning pipeline may benefit from adjustments to the last hidden layers.

B. Execution pipeline

We propose two selection algorithms to choose the most convenient action based on the neural network's predicted scaling values. The first algorithm is a greedy selection strategy that works well for short time horizons, while the second one leverages the Exo-MDP formulation in a Monte Carlo Planning framework.

1) *Greedy Action Selection*: Given a desired time window, w , we set $N = w/\Delta t$ in (5). Because the network predicts the average value over the time interval $(\bar{t}, \bar{t} + N\Delta t]$, we can solve (4) by simply inferring the predicted scaling factor, \hat{y} , for all available actions and selecting the one with the highest




Fig. 2: Neural network architecture.

Algorithm 1: Greedy Action Selection Algorithm

Input: Observed states x_r and x_h , Current human goal G_h , Neural Network NN

Output: Selected action $best_action$

```

1 best_action ← None;
2 highest_reward ←  $-\infty$ ;
3 actions ← getAvailableActions( $x_r$ );
4 for each action in actions do
5    $g_r$  ← getRobotGoal(action);
6   reward ← NN.predict( $x_r, x_h, g_r, G_h \cdot \mu$ );
7   if reward > highest_reward then
8     best_action ← action;
9     highest_reward ← reward;
10 Return best_action;
```

value (i.e., the one that introduces the smallest average speed reduction):

$$u_r^* = \operatorname{argmax}_{u_r \in U_r(\bar{t})} \hat{y}(x_r(\bar{t}), x_h(\bar{t}), g_r(\bar{t}), G_h(\bar{t}) \cdot \mu) \quad (9)$$

where $g_r(\bar{t}) = \operatorname{getRobotGoal}(u_r(\bar{t}))$ and $G_h(\bar{t}) \cdot \mu$ is the mean value of the current human goal distribution.

A pseudo-code of the procedure is in Alg. 1. At each decision time, we predict the expected safety scaling (line 6) and return the action associated with the highest value. Note that the loop at line 4 can be easily parallelized for computational efficiency.

We refer to this strategy as greedy because it only considers the best action over a short horizon, w . In our experiments, w is approximately equal to the average duration of one action.

2) *Monte Carlo Action Selection:* The Greedy Action Selection does not work well for long horizons because the scaling prediction worsens when the time horizon encompasses several robot actions and human goals. In general, Exo-MDPs can be solved with Monte Carlo search methods like Monte Carlo Tree Search or Monte Carlo Planning.

To suit computational time requirements, we propose an anytime search algorithm inspired by Monte Carlo Planning. The algorithm works as follows (see Alg. 2). Starting from the current state, it expands all children by applying each available robot’s action. Then, it starts a thread for each child and applies random rollouts (see Alg. 3) until the sequence reaches a terminal state. In our case, a terminal state is either a state where no available actions exist or a maximum length

of the sequence is reached. Each rollout returns a reward associated with one of the children. After the timeout, the algorithm computes the average reward for each child and selects the action that generated the highest reward child.

The difficulty of applying Monte Carlo methods to Exo-MDPs is that the exogenous input’s dynamics is unknown. For this reason, it is not possible to propagate the exogenous state at a generic iteration i . A solution is to sample it from a probability distribution. In our problem, this applies both to the human state and goal. To sample meaningful values of $x_h(i)$ and $g_h(i)$, we select a goal $G_h(i) = \mathcal{G}_{h,j}$ where

$$\mathcal{G}_{h,j} \sim \mathcal{U}(\mathcal{G}_h), \quad (10)$$

and \mathcal{U} is the uniform distribution. Considering that $\mathcal{G}_{h,j} \sim \mathcal{N}(\mu_j, \sigma_j) \forall j$, we choose

$$x_h(i) \sim \mathcal{G}_{h,j} \text{ and } g_h(i) = \mu_j. \quad (11)$$

This allows us to obtain concrete values of x_h and g_h to propagate the state over the tree and estimate the reward at future steps¹. The resulting propagation function is shown in Alg. 4 and forms the basis of the random rollouts described in Alg. 3).

Alg. 2 is a simplified version of Monte Carlo Planning that allows for straightforward parallelization and a faster computation compatible with the intended online implementation. On the other hand, the search is shallower than standard MCP because it uses random rollouts instead of the iterative exploration-v.-exploitation trade-off typical of MCP.

V. EXPERIMENTS

We demonstrate our approach in a simulated pick&packaging example and a real-world scenario. Simulations evaluate the learning pipeline and provide an ablation study of the scheduling approach. The real-world scenario proves the effectiveness of the approach in execution against several baselines.

A. Simulated pick&packaging scenario

We consider a box-picking scenario with a robot and an operator in the RoboDK simulator. As RoboDK does not have a human simulator, we mimic the operator by using a humanoid robot, Motoman SDA10F, equipped with an omnidirectional mobile base. The manipulator is a 6-dof Fanuc CRX10iA with a vacuum gripper. The scenario is in Fig. 3.

The robot and the operator pick and place boxes from inbound to outbound areas. The robot’s inbound area consists of a conveyor, while the outbound areas are on the two tables on the robot’s sides and accommodate up to 5 boxes each. The human inbound area contains 6 boxes, while the outbound areas are near the robot and contain 3 boxes each. The robot slows down according to a stepwise safety function, which can be modeled according to (6) by choosing $d = \|x_r - x_h\|$, $D_i = [d_{i-1}, d_i]$, $d_0 = 0$, $d_K = +\infty$, and $\{d_i\}$ evenly spaced for $i = 1, \dots, K-1$. An example for $K = 5$ is in Fig. 4. In the experiments, we used $d_4 = 2$ m.

¹This approximation relies on the assumption that the operator will spend most of the time near their goal position. In general, x_h can be sampled from an occupancy map conditioned on the chosen goal, $G_h(i)$.

Algorithm 2: Monte Carlo Selection Algorithm

Input: Observed states x_r and x_h , Observed human goal G_h , Neural Network NN , Discount factor γ
Output: Selected action $best_action$

```

1 actions  $\leftarrow$  getAvailableActions( $x_r$ );
2 best_action  $\leftarrow$  None; best_average_reward  $\leftarrow$   $-\infty$ ;
3 for each action in actions do
4   Start a new thread;
5   Threaded function evaluateAction(action, max_time);
6   begin
7      $x_r$ , total_reward  $\leftarrow$  propagate( $x_r$ ,  $x_h$ ,  $G_h$ , action);
8      $i \leftarrow 0$ ;
9     while elapsed_time < max_time do
10      reward  $\leftarrow$  performRollout( $x_r$ );
11      total_reward  $\leftarrow$  total_reward +  $\gamma^i$  reward;
12       $i \leftarrow i + 1$ ;
13    Return total_reward / iterations ;
14 for each action thread in active threads do
15   Wait for all threads to complete;
16   average_reward  $\leftarrow$  getThreadResult(action);
17   if average_reward > best_average_reward then
18     best_average_reward  $\leftarrow$  average_reward;
19     best_action  $\leftarrow$  action;
20 Return best_action;
```

Algorithm 3: performRollout

Input: x_r
Output: total_reward

```

1 total_reward  $\leftarrow 0$ ;
2 while isTerminalState( $x_r$ ) is false do
3   actions  $\leftarrow$  getAvailableActions( $x_r$ );
4   action  $\leftarrow$  random(actions);
5    $G_h \leftarrow$  random( $G_h$ );
6    $x_r$ , reward  $\leftarrow$  propagate( $x_r$ ,  $G_h$ , action);
7   total_reward  $\leftarrow$  total_reward + reward;
8 Return total_reward;
```

Algorithm 4: propagate

Input: x_r , G_h , action
Output: x_r , reward

```

1  $g_r \leftarrow$  getRobotGoal(action) ;
2  $x_h \sim G_h$ ;
3 reward  $\leftarrow$  NN.predict( $x_r$ ,  $x_h$ ,  $g_r$ ,  $G_h \cdot \mu$ );
4  $x_r \leftarrow g_r$ ;
5 Return  $x_r$ , reward;
```




Fig. 3: Simulation scenario.




Fig. 4: Example of safety speed scaling function for $K = 5$.

Every time the robot picks a new box, we aim to estimate the expected future slowdown for all available placing points. A video of the process is available in the multimedia attachments.

1) *Data Collection and Training:* We run the process 1000 times and, with a frequency of 10 Hz, record the following data for each time \bar{t} :

- the robot’s end-effector position, $x_r(\bar{t}) \in \mathbb{R}^3$;
- the human’s centroid position, $x_h(\bar{t}) \in \mathbb{R}^3$;
- the robot’s goal position, $g_r(\bar{t}) \in \mathbb{R}^3$;
- the human’s goal position, $G_h(\bar{t}) \in \mathbb{R}^3$;
- the robot’s speed scaling, $s(\bar{t}) \in [0, 1]$.

To account for the human movements’ stochasticity, we randomize the path of the operator: the human’s goal is sampled from a Gaussian distribution centered in the nominal goal position with a standard deviation of 0.05 m in the horizontal plane; the midpoint of the path is sampled from a Gaussian distribution centered in a nominal midpoint and with standard deviation equal to 0.25 m in the horizontal plane.

We use an 80%-20% training/test split of the collected dataset and we train a feed-forward neural network. The design of the network follows the method described in Sec. IV-A for staircase safety functions. The network has 12 input features, the hidden layers’ width is equal to 64 (using batch normalization and ReLU activation), and a last hidden layer with a width equal to K and a Softmax activation function. The output layer is linear and yields a scalar.

B. Learning Results

Fig. 5a shows the accuracy of the network on the test sets with $K = 5$ and a predictive horizon of 14 seconds (equal to the average of robot tasks’ durations estimated during the data collection). Similar accuracy values were obtained for

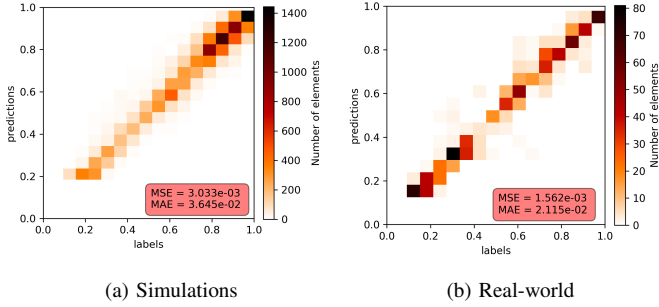


Fig. 5: Actual v. predicted average scaling values on test datasets. The heatmap represents the density of samples. A high density along the diagonal denotes a high regression accuracy.

TABLE I: Learning results for different values of K .

K	3	5	10	20
MSE $\cdot 10^3$	3.22	3.03	4.07	5.45
# hidden layers	5	5	6	6

predictive horizons in the range of 8–20 seconds. The diagonal pattern and low MSE values clearly show the good predictive capabilities of the regressor. To demonstrate that generality of the approach across safety configurations, we evaluate the prediction accuracy for different values of K . Note that the design principles of the network remain the same for different values of K , as described in Section IV-A. The only hyperparameter is the number of hidden layers, which may increase for complicated functions. Table I shows the accuracy and the number of hidden layers of the network for $K \in \{3, 5, 10, 20\}$. The results show that the accuracy is consistent for different values of K . As expected, large values of K may benefit from additional hidden layers to achieve a satisfactory accuracy. It is worth mentioning that the choice of the number of hidden layers can be easily automatized through grid search.

C. Execution Results

The main motivation behind this work is that the knowledge of the safety function is often uncertain or hard to model explicitly. For this reason, we evaluate the performance of a selection strategy when the safety function is wrongly modeled. To this end, we consider a test scenario with $K = 5$ and compare the following cases:

- greedy: our greedy action selection method (Sec. IV-B).
- greedy-inacc-1: same greedy strategy, but the scaling predictor is trained on a scenario with $K = 3$.
- greedy-inacc-2: same greedy strategy, but the scaling predictor is trained on a scenario with inflated distance thresholds d_i^* , such that $d_i^* = 1.2 d_i$ for $i = 1, \dots, 4$.

We also consider a `random` selection strategy, where the next robot’s action is randomly chosen among the available ones.

Table II shows the average execution time and safety speed slowdown for all the methods. `greedy` reduced the execution time by around 12% with respect to `random`. By contrast, `greedy-inacc-1` and `greedy-inacc-2` respectively increased execution time by 6% and 9% compared

TABLE II: Simulation results. Average values of the execution time (per task) and speed scaling for all methods. Values in parenthesis refer to std. deviation.

	exec. time	scaling
greedy	17.7(2.1)	0.82
greedy-inacc-1	18.7(3.2)	0.78
greedy-inacc-2	19.3(2.7)	0.75
random	20.1(3.1)	0.72

to `greedy`, showing that inaccurate knowledge of the safety function significantly degrades task selection performance.

D. Real scenario

We consider a pick&packaging scenario with a human operator and a collaborative robot, Universal Robots UR5e. The robot picks small boxes from a slider and places them into one of the four boxes on the table. In the meantime, the operator performs operations such as packing items, inspecting the packages, and refilling items. Thus, the robot has one picking point and four placing areas, while the human moves among three areas spanning both the shared and the non-shared workspaces. The robot slows down according to the stepwise function shown in Fig. 4 with $d_1=0.6$ m, $d_2=0.8$ m, $d_3=0.1.2$ m, and $d_4=1.6$ m. The scenario is shown in Fig. 6.

We tested two variants of the scenario.

- In the **first variant**, each box is replaced as soon as it is full to emulate a continuous flow use case. From the action selection point of view, this means that the robot can always choose any box.
- In the **second variant**, the boxes are replaced only when all of them are full. This means that the robot’s available actions decrease according to the previous choices.

1) *Data Collection and Training*: Similarly to the simulations, we collect the human’s centroid and goal, the robot’s end-effector and goal, and the speed scaling value. The robot position is computed by applying forward kinematics to the robot’s measured joint state; the robot’s goal is known as it is decided by the action selection algorithm; the robot’s speed scaling is available from the robot’s APIs. The operator’s position is estimated by a skeleton tracking algorithm using RGB-D images from a fixed Intel Realsense d435i camera. As for the operator’s goal, we approximate it with the average human’s position during an operation and a standard deviation equal to 0.1 m.

Note that the order of the operations is not fixed as the operator selects its task on the fly. To identify the ongoing human operation (and the associated goal, G_h), the operator presses a button when it starts a new operation.

The neural network structure is the same as in Sec. V-A1. To train it, we collect data for around 90 minutes with a sampling rate of 10 Hz. Fig. 5b shows the regression accuracy, which is consistent with the simulation results.

2) *Metrics and Baselines*: We aim to demonstrate that our approach reduces the cycle time and safety speed slowdowns. Thus, we measure the robot’s execution time for each time and the average speed scaling during execution.

TABLE III: Experimental results. Average values of the execution time (per task) and speed scaling for all methods. Values in parenthesis refer to std. deviation.

	variant 1		variant 2	
	exec. time	scaling	exec. time	scaling
random	24.7(6.6)	0.60	24.9(6.6)	0.59
round-robin	26.8(5.9)	0.55	27.1(6.6)	0.54
reactive	24.5(4.7)	0.60	25.1(5.5)	0.59
greedy	18.6(1.7)	0.74	20.2(3.9)	0.68
monte carlo	20.3(3.6)	0.72	20.7(3.9)	0.66

We compare the following action selection algorithms:

`random`: the next robot’s action is randomly chosen among the available ones.

`round-robin`: the robot selects the next placing box in sequential order from the first to the fourth box. Full boxes are skipped.

`reactive`: the robot selects the box furthest from the human position at the current time.

`greedy`: our greedy action selection method (Sec. IV-B).

`monte carlo`: our MCP variant (Sec. IV-B).

3) *Implementation notes*: The main tunable parameter in our method is the predictive window, w . `greedy` and `monte carlo` use the neural network to predict the average scaling value over the next robot’s action. For this reason, we set w equal to the average duration of the robot’s actions (~ 14 s) in the `greedy` method. While `greedy` has a negligible computational delay, `monte carlo` requires a large number of rollouts to find a good solution. As the robot picking task starts and ends at the same point, we run `monte carlo` in masked time during the picking operation. This strategy grants ~ 4 seconds to find a solution but introduces a mismatch in x_h . As a matter of fact, the value of x_h when the robot applies the action is different from that used by the algorithm. To account for this mismatch, we use $w=18$ s in `monte carlo` to account for the picking phase duration.

4) *Results*: For the first variant, we run one test of around 25 minutes for each method. This corresponds to a number of action selections between 60 and 80 for each method. As for the second variant, we run the whole process 15 times for each method, for a total of 180 action selections for each method.

Table III summarizes the results. In the first variant, `greedy` gives the best results, followed by `monte carlo`. Compared to `random`, they improve the robot execution time of about -25% and -18% and the average speed scaling of about +23% and +20%, respectively. The second variant yields similar results: `greedy` and `monte carlo` outperform the baseline of about -19% and -17% in terms of execution time and of about +15% and +11% in terms of average speed scaling. The proposed methods also show a smaller variance.

As an example, Fig. 7 compares the distribution of the speed scaling during one execution of the second variant. For the sake of brevity, we only show `random`, `greedy`, and `monte carlo`. The histograms show that `greedy` and `monte carlo` lead to a higher frequency of high scaling values (i.e., $s=0.75$ and $s=1.0$), which means these methods correctly choose the next action in such a way as to reduce the impact of the safety speed reduction.

We did not observe a significant difference between the `random` selection and the other baselines, except that `reactive` shows a smaller variance. Although `reactive` is human-aware, this strategy is unable to choose convenient actions because it is based only on the current operator’s position. This strategy is ineffective because the scenario requires the operator to move frequently across the workspace.

In our experiments, `greedy` and `monte carlo` gave comparable results. One reason for this is that the scenario did not involve a strong coupling between current and future actions’ rewards, which leads to good results for the `greedy` algorithm. Moreover, `monte carlo` performance is affected by implementation issues. First, the algorithm is triggered in advance as described in Sec. V-D3, meaning that the operator may be in a different position than expected at planning time; ii) Second, `monte carlo` is often unable to find optimal solutions in the allotted time. Improvements on the `monte carlo` implementation will be object of future works.

Finally, Fig. 6 shows snapshots taken from an execution using the `greedy` and the `random` methods. The pictures show how the proposed method (top pictures in Fig. 6) picks actions that do not interfere with the operator by anticipating human movements. On the contrary, the `random` method (bottom pictures in Fig. 6) often drives the robot close to the region occupied by the human (see the attached video).

VI. CONCLUSIONS

This paper presented a safety function-agnostic method for optimizing collaborative robot task selection by learning the correlation between system state and safety-induced speed reductions. One limitation of the learning-based approach is that it assumes a level of repeatability in human executions, making it best suited for structured tasks. Future work will focus on extending the method’s capabilities to partially unstructured tasks by incorporating context-aware models and using automatic task segmentation to infer the human current operation without explicit feedback.

REFERENCES

- [1] “ISO 10218-2:2025 Robots and robotic devices — Safety requirements for industrial robots,” International Organization for Standardization, Geneva, CH, Standard, 2025.
- [2] M. Faccio, I. Granata, and R. Minto, “Task allocation model for human-robot collaboration with variable cobot speed,” *J of Intelligent Manufacturing*, pp. 1–14, 2023.
- [3] M. Faroni, M. Beschi, and N. Pedrocchi, “Safety-aware time-optimal motion planning with uncertain human state estimation,” *IEEE Robot. Autom. Lett.*, vol. 7, no. 4, pp. 12 219–12 226, 2022.
- [4] R. Laha, W. Wu, R. Sun, N. Mansfeld, L. F. Figueredo, and S. Haddadin, “S*: On safe and time efficient robot motion planning,” in *IEEE ICRA*, 2023, pp. 12 758–64.
- [5] C.-M. Huang and B. Mutlu, “Anticipatory robot control for efficient human-robot collaboration,” in *ACM/IEEE HRI*, 2016, pp. 83–90.
- [6] A. Pupa, M. Minelli, and C. Secchi, “A Time-Optimal Energy Planner for Safe Human-Robot Collaboration,” in *IEEE ICRA*, 2024, pp. 17 373–17 379.



Fig. 6: Screenshots from videos of the experiments. (top) greedy. (bottom) random.




Fig. 7: Experimental results. Statistic distribution of the speed scaling factor in the second scenario variant.

[7] A. Palleschi, M. Hamad, S. Abdolshah, M. Garabini, S. Haddadin, and L. Pallottino, “Fast and safe trajectory planning: Solving the cobot performance/safety trade-off in human-robot shared environments,” *IEEE Robot. Autom. Lett.*, vol. 6, pp. 5445–5452, 2021.

[8] F. Vicentini, “Terminology in safety of collaborative robotics,” *Robotics & CIM*, vol. 63, p. 101921, 2020.

[9] A. Kanazawa, J. Kinugawa, and K. Kosuge, “Adaptive motion planning for a collaborative robot based on prediction uncertainty to enhance human safety and work efficiency,” *IEEE T. Rob.*, vol. 35, pp. 817–832, 2019.

[10] J. Flowers and G. Wiens, “A spatio-temporal prediction and planning framework for proactive human-robot collaboration,” *Journal of Manufacturing Science and Engineering*, vol. 145, no. 12, 2023.

[11] M. Li, J. Qin, Z. Wang, Q. Liu, Y. Shi, and Y. Wang, “Optimal motion planning under dynamic risk region for safe human-robot cooperation,” *IEEE/ASME T. Mech.*, pp. 1–11, 2024.

[12] A. Pupa and C. Secchi, “Efficient iso/ts 15066 compliance through model predictive control,” in *IEEE ICRA*, 2024, pp. 17 358–17 364.

[13] M. Hosseinzadeh, B. Sinopoli, and A. F. Bobick, “Robot action planning in the presence of careless humans,” *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 54, no. 9, pp. 5400–5411, 2024.

[14] M. Lagomarsino, M. Lorenzini, E. De Momi, and A. Ajoudani, “Pro-mind: Proximity and reactivity optimization of robot motion to tune safety limits, human stress, and productivity in industrial settings,” *IEEE T. Robotics*, vol. 41, pp. 2067–2085, 2025.

[15] L. Meng, L. Yang, and E. Zheng, “Hierarchical human

motion intention prediction for increasing efficacy of human-robot collaboration,” *IEEE Robotics and Automation Letters*, vol. 9, no. 9, pp. 7637–7644, 2024.

[16] W. Liu, X. Liang, and M. Zheng, “Task-constrained motion planning considering uncertainty-informed human motion prediction for human-robot collaborative disassembly,” *IEEE/ASME T. Mech.*, vol. 28, no. 4, pp. 2056–2063, 2023.

[17] L. Scalera, C. Nainer, A. Giusti, and A. Gasparetto, “Robust safety zones for manipulators with uncertain dynamics in collaborative robotics,” *Int. J. Comput. Integrated Manuf.*, vol. 37, no. 7, pp. 887–899, 2024.

[18] W. Liu, K. Eltoumy, S. Tian, X. Liang, and M. Zheng, “Integrating uncertainty-aware human motion prediction into graph-based manipulator motion planning,” *IEEE/ASME T. Mech.*, vol. 29, no. 4, pp. 3128–3136, 2024.

[19] A. Hermann, F. Mauch, K. Fischnaller, S. Klemm, A. Roennau, and R. Dillmann, “Anticipate your surroundings: Predictive collision detection between dynamic obstacles and planned robot trajectories on the gpu,” in *IEEE ECMR*, 2015, pp. 1–8.

[20] M. Ferrari, S. Sandrini, C. Tonola, E. Villagrossi, and M. Beschi, “Predicting human motion using the unscented kalman filter for safe and efficient human-robot collaboration,” in *IEEE ETFA*, 2024, pp. 1–8.

[21] S. Pellegrinelli and N. Pedrocchi, “Estimation of robot execution time for close proximity human-robot collaboration,” *Integrated Computer-Aided Engineering*, vol. 25, no. 1, pp. 81–96, 2018.

[22] M. N. Finean, L. Petrović, W. Merkt, I. Marković, and I. Havoutis, “Motion planning in dynamic environments using context-aware human trajectory prediction,” *Robotics and Auton. Syst.*, vol. 166, p. 104450, 2023.

[23] Z. Cao, H. Gao, K. Mangalam, Q.-Z. Cai, M. Vo, and J. Malik, “Long-term human motion prediction with scene context,” in *ECCV*, 2020, pp. 387–404.

[24] P. Kratzer, M. Toussaint, and J. Mainprice, “Prediction of human full-body movements with motion optimization and recurrent neural networks,” in *IEEE ICRA*, 2020, pp. 1792–1798.

[25] P. J. Rousseeuw, “Silhouettes: A graphical aid to the interpretation and validation of cluster analysis,” *J. of Comput. and Applied Math.*, vol. 20, pp. 53–65, 1987.