Chapter 1

Time to Contact for Robot Safety Stop in Close Collaborative Tasks

Alberto Olivares-Alarcos[1], Sergi Foix[1] and Guillem Alenyà[1]

Within the next years, industrial collaborative robots will collaborate with human operators, whose safety might be compromised. To ensure a safe collaboration, robots should estimate the risk of collision with respect to the pose and motion of operators within the shared workspace. A common approach is to compute and maintain a minimum distance between humans and robots during tasks' execution. Nevertheless, separation-based solutions do not capture the real dynamics of the human-robot interaction, and tend to be rather conservative, avoiding a close human-robot collaboration. In this chapter, we explore the concept of time-to-contact (TTC) as a softer trigger of safety stops in collaborative scenarios where humans and robots are in constant closeness. Particularly, we propose a TTC formulation and study its advantages with respect to two approaches based on the protective distance proposed by the ISO standards. We compared the three methods in some representative cases extracted from an example of collaborative task. The evaluation is firstly done in simulation, and then, in a more realistic setup with a simulated human, aiming for repeatability, and a real robot. Furthermore, we showcased our approach in a demo of a complete collaborative task. TTC allows robots to operate closer to humans and for longer times before a safety stop is issued, which benefits long-term productivity. This later stop produces shorter human-robot distances, which might affect safety. However, the increment in time that the robot moves before stopping (productivity) is far greater than the reduction in distance (safety). Hence, we can state that TTC greatly improves productivity while slightly compromising safety. In conclusion, our work demonstrates that TTC is a smoother, but still safe, collision risk estimator for close human-robot collaboration.

1.1 INTRODUCTION

Over the last years, there has been a growing need in the industrial sector for more flexible manufacturing processes where humans and robots are expected to work together. Nevertheless, freeing robots from their restrictive current work cells would

[1]Institut de Robòtica i Informàtica Industrial, CSIC-UPC, Llorens i Artigas 4-6, 08028 Barcelona, Spain; {aolivares,sfoix,galenya}@iri.upc.edu

compromise humans' safety. Collaborative robots, or co-bots, are robots specifically designed for direct interaction with humans within a defined collaborative workspace. Collaborative robots shall, among others, be able to perceive human operators' pose and motion, estimate possible risks, and adapt their behavior to ensure humans' safety and a reliable and productive collaboration.

In 2011, the International Organization for Standardization released the ISO 10218.1 [1] and the ISO 10218.2 [2], which presented safety guidelines for industrial robots. In 2016, the ISO/TS 15066 [3] extended the previous standards providing specific guidance for safety in collaborative robotics, where a formulation to compute the minimum protective distance was proposed. One of the main limitations of this formulation is that the real direction of motion of the robot and the human is not taken into account. Hence, it results in an over-conservative risk estimator, and prevents a proper collaboration in applications where humans and robots constantly and closely share the workspace. Indeed, there is not a standard way to address this issue yet, and the actual implementation of the formula is still greatly left to the discretion of the integrator [4].

Inspired by the aforementioned ISO standards, several works about safety in collaboration have been published during the last years [5, 6, 7, 8]. Indeed, some of them discussed and aimed to overcome different ISO's formulation drawbacks [9, 10]. However, there is still room for improvement, especially in collaborative tasks in which the human-robot closeness is regular (such as the one in Fig. 1.1). Hence, our aim is to explore an alternative to the ISO's formulation well adapted to intensive human-robot collaboration.

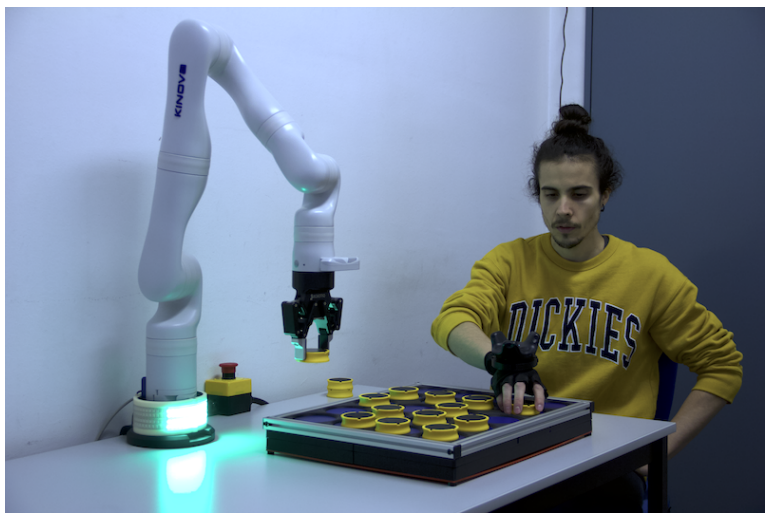

Figure 1.1: Collaboratively filling a tray: example of an industrial task where the human and the robot continuously share both the workspace and the execution of the task (pick and place).

In this chapter, we explore the concept of time-to-contact (TTC) [11] as an indicator of the likelihood of the risk of collision during close human-robot collaborations. We show how TTC can be computed and used to stop the robot to avoid collisions while closely sharing the execution of a task with the human. Unlike most of the approaches within the literature, TTC naturally captures the dynamics of the human-robot interaction, considering the actual pose, speed, and direction of both agents. This makes TTC a more precise risk indicator, delaying safety stops and allowing the robot to operate for longer times before stopping. The contributions of this work are:

- an extension of a two-dimensional formulation to compute TTC in a three-dimensional space;
- an algorithm to stop the robot's motion based on the computed TTC;
- a twofold evaluation, first, in simulation, and second, using a real robot and a simulated human (aiming for repeatability). In both, using prototypical motions to validate the improvement of our approach in delaying the safety stop compared to two ISO-based methods;
- and a qualitative validation through the implementation of our approach in a real use case: a complete collaborative task in which a human and a robot, closely interacting, filled a tray with tokens.

1.2 RELATED WORK

The ISO standards for safety in collaborative robotics [1, 2, 3], proposed to use the Speed and Separation Monitoring (SSM) approach, maintaining a robot's speed and a minimum distance between the robot and the human. SSM has been extensively applied in the collaborative robotics domain, and it is utterly aligned to the scope of this work. In some works, the robot's speed is adapted to the distance using discrete regions of the space [5, 6, 7]. Aiming to enhance the collaboration's fluency, several authors proposed to continuously adapt the robot's speed w.r.t. the current human-robot distance, but still without considering the human's and robot's motion direction [12, 8, 13, 14]. Looking for a further improvement, more complex formulations to compute the human-robot distance considered the human's and/or the robot's motion direction [9, 15, 10]. In all those articles, as in our approach, the robot's behavior is adapted based on the estimation of a possible risk of collision with a human. The estimation is based on the distance between the human and the robot, and the robot adapted its speed. In our work, we propose to simplify the formulation using the time-to-contact concept, as it naturally embeds the directions and velocities of the two agents.

Time-to-contact has been widely used in the literature for automotive collision estimation, warning, and avoidance [16, 17, 18, 19]. Authors proposed different approaches: vehicle to vehicle and vehicle to driver warning, obstacle avoidance, autonomous emergency braking system, etc. In those works, TTC was computed in a two-dimensional space, since cars move in a plane. In this work, we explored how

the ideas discussed in those articles might be extrapolated to collaborative industrial scenarios, where humans and robots move in a three-dimensional space.

We have opted for stopping the robot, because that is the most compliant strategy with the regulations of industrial environments (e.g., ISO). However, TTC has also been used to control other robot's reactions such as adapting the speed or modifying the robot's motion plan for different applications: robot docking and landing [20, 21], UAS maneuvers [22], and obstacle avoidance and target chasing [23]. Those works presented interesting approaches showing the potential use of TTC in different applications, which might be considered as ideas for our future work.

1.3 TIME-TO-CONTACT-BASED SAFETY STOP FOR CLOSE HUMAN-ROBOT COLLABORATION IN INDUSTRIAL TASKS

1.3.1 Background

TTC is a biologically inspired measure typically used for obstacle detection and reactive control of motion. It can be defined as the time that an observer will take to make contact with a surface assuming constant relative velocity. Hence, TTC is usually expressed in terms of the speed and the distance of the considered obstacle:

$$ttc = -\frac{Z}{\frac{dZ}{dt}}, \tag{1.1}$$

where Z is the distance between the observer and the obstacle, and dZ/dt is the velocity of the observer with respect to the obstacle. It is possible to compute TTC from a pure computer vision perspective, by just detecting the deformation of objects in consecutive RGB images without calibration [24, 25, 26]. These methods are very sensitive to error detection, so we will compute TTC utilizing the actual observer's and obstacle's poses and velocities as we have access to these measurements. In our work, the observer is the end effector of a robot and the obstacle is the hand of a human.

1.3.2 Time-to-contact computation

Inspired by Hou et al. [27] and their TTC formulation for 2D collisions between circles, we propose here a 3D variation of their formula. Hence, the robot's end effector and the human's hand are considered as spheres, and we compute the TTC as the time that it would take the two spheres to collide. Determining when two spheres collide is a matter of determining the moment at which the distance between their centers is equal to the sum of their radii.

Let $\vec{r'} = (r'_x, r'_y, r'_z)$ and $\vec{h'} = (h'_x, h'_y, h'_z)$ denote the positions of the robot and the human respectively at the moment of contact. Hence, the distance between them is

$$d = ||(r'_x, r'_y, r'_z) - (h'_x, h'_y, h'_z)||. \tag{1.2}$$

Knowing that when in contact the distance is equal to the sum of their radii and expanding Eq. 1.2 we get

$$rad_r + rad_h = \sqrt{(r'_x - h'_x)^2 + (r'_y - h'_y)^2 + (r'_z - h'_z)^2}, \tag{1.3}$$

where rad_r and rad_h are the radii of the spheres representing the robot and the human respectively.

Assuming that the human and the robot are not currently colliding and that both move with constant linear velocity, we can rewrite their positions at the moment of contact based on their current position and velocity:

$$\vec{r'} = \vec{r} + \vec{v_r} ttc \tag{1.4}$$

$$\vec{h'} = \vec{h} + \vec{v_h} ttc, \tag{1.5}$$

where $\vec{r} = (r_x, r_y, r_z)$ and $\vec{h} = (h_x, h_y, h_z)$ denote the current positions of the robot and the human, $\vec{v_r} = (v_{r_x}, v_{r_y}, v_{r_z})$ and $\vec{v_h} = (v_{h_x}, v_{h_y}, v_{h_z})$ their current Cartesian velocity respectively, and ttc is the corresponding time-to-contact. Substituting Eqs. 1.4 and 1.5 in Eq. 1.3, we obtain a quadratic equation where, if real positive roots exist, the smallest value is the pursued TTC.

1.3.3 TTC-based safety stop algorithm

Given a TTC between the human and the robot, we can use it so that the robot adapts its current state to avoid possible collisions. In this work, we propose to follow one of the strategies proposed in the ISO 10218.1 [1], issuing a safety-rated monitored stop. Hence, the robot would continue its motion until a certain TTC threshold is violated, from now on, time-to-stop (*ttstop*). In Alg. 1, we show the decision-making process to compute TTC and stop the robot given that TTC value. If the value of TTC is equal or smaller than *ttstop*, the robot will stop its motion and a safety-rated monitored stop will be issued (see line 6 in Alg. 1). When the value of TTC is greater than *ttstop*, the robot will continue its motion at the task's nominal speed (see line 9 in Alg. 1). We decided that once a safety stop was issued, the robot would remain stopped and we would exit the TTC-based safety stop loop (see line 2 in Alg. 1). Since our focus is on very close human-robot applications, the robot would only resume the motion after a human's command. We think that this recovery strategy is the most appropriate one for industrial scenarios similar to our case.

1.4 STATE-OF-THE-ART APPROACHES: ISO AND FUZZY ISO

In order to evaluate our proposal, we compared it against two methods to trigger a safety stop based on computing the minimum protective distance. First, we used the linear version of the formulation defined in ISO/TS 15066 [3]:

$$S = (v_h T_r + v_h T_s) + (v_r T_r + v_s T_s) + (C + Z_r + Z_s), \tag{1.6}$$

Algorithm 1: decision-making loop to compute TTC, stop the robot and select the robot's speed

Input: Time-to-stop ($ttstop$), nominal robot's speed (V_{r_n}), robot and human radii (rad_r, rad_h)

$safety\_stop \leftarrow false$
while ***not****(safety_stop)* **do**
  $info_r \leftarrow$ GetRobotPoseVelocity()
  $info_h \leftarrow$ GetHumanPoseVelocity()
  $ttc \leftarrow$ ComputeTTC($info_r$, $info_h$, rad_r, rad_h)
  if $ttc \leq ttstop$ **then**
    $V_r \leftarrow 0$
    $safety\_stop \leftarrow true$
  else
    $V_r \leftarrow V_{r_n}$
  end
  Publish computed robot's desired speed V_r
end

where S is the protective distance, v_h is the 'directed speed' of the operator (i.e., the rate of travel of the operator toward the robot), v_r is the directed speed of the robot in the direction of the operator, and v_s is the directed speed of the robot in the course of stopping. T_r is the time for the robot system to respond to the operator's presence, while T_s is the time to bring the robot to a safe, controlled stop. The remaining terms capture measurement uncertainty, where C is an intrusion distance safety margin based on the expected human reach, Z_r is the robot position uncertainty, and Z_s is the operator position (sensor) uncertainty. There is not a standard way to measure the human's and the robot's speeds, nor to compute the times. Indeed, setting the values of the uncertainty constants might also be challenging [4]. Since a discussion of the proper values to choose is out of the scope of this chapter, we will use a simplified version of the formula. First, we will assume that the speed of the robot while stopping, v_s, is equal to the motion speed, v_r. Finally, we will obviate the effect of the uncertainty constants because they would affect all the methods compared in this work. The simplified equation is:

$$S = (v_h + v_r)T_b, \tag{1.7}$$

where T_b is the sum of T_r and T_s, thus the total time to brake, including the time to respond to the human's presence and to stop the robot. During the different evaluations, we decided that the value used for T_b was also used to set the time-to-stop ($ttstop$). Hence, we established a correlation between our method (see Alg. 1) and the ones presented in this section. Actually, it makes sense to use the total time to brake as the minimum TTC used to stop ($ttstop$). It would ensure that the robot would stop right before contacting the human, just as the minimum protective distance would.

The second state-of-the-art method used for the evaluations, Fuzzy ISO, was proposed by Campomaggiore et al. [10]. They presented a fuzzy-logic system to

merge the protective distance formulation with information of the current human's and robot's motion direction. The fuzzy rules are used to scale the effect of the human's and the robot's velocities: e.g. when they are going away in opposite directions their method allows them to relax the ISO's formula. The resulting formula would be the Eq. 1.7 multiplied by the output of the fuzzy-logic system $\alpha \in (0,1)$:

$$S = \alpha(v_h + v_r)T_b. \tag{1.8}$$

Using the aforementioned two equations, we compute the minimum protective distance (S). When the Euclidean 3D distance between the spheres representing the human and the robot is smaller or equal to S the robot stops, similarly to our proposal (Alg. 1).

1.5 VALIDATION OF TIME-TO-CONTACT FOR ROBOT SAFETY STOP IN CLOSE COLLABORATIVE TASKS

We aimed at being able to evaluate all the pros and cons between the different safety stop methods. Hence, we found necessary to create the taxonomy of all possible situations in tasks such as ours (Fig. 1.1). After a comprehensive study, we summarized all possible robot-human collaborative dynamic states in just 7 prototypical cases (see Fig. 1.2). For the evaluation, we were only interested in the cases that would cause a safety stop. Hence, the cases in which a collision might occur (Figs. 1.2a, 1.2c, and 1.2e). From now on, we will refer to them as cases: A, C, and E respectively. It is important to remark that we do not need to evaluate the rest of situations, neither the complete task execution. The reason is that there will be no change of behaviour among methodologies.

1.5.1 Evaluation I - Simulation and statistical analysis of the results' significance

First, we evaluated in simulation the performance of our approach against the two methods presented in Sec. 1.4 (ISO and Fuzzy ISO). We used the already mentioned most interesting cases from the taxonomy: A, C, and E.

For each case, we simulated two spheres that moved close to each other. They represented the end effector of a robot and the hand of a human, from now on just robot and human. Once the safety stop was triggered, we finished the simulation. The first hypothesis is that the proposed approach would allow the robot to move during more time before issuing a safety stop. The second hypothesis is that our approach would let the robot to get closer to the human but would still be safe, issuing a safety stop before any possible collision.

In order to validate our hypotheses, we evaluated the approaches for several human speeds in each of the three selected situations. Other parameters such as the human's and robot's initial position or the robot's nominal speed, would affect the triggering of the safety stop. However, we focused on the human's velocity parameter for two reasons. First, it has a direct effect on the dynamics of the interaction, playing a fundamental role in our hypotheses. Secondly, since it is a parameter we

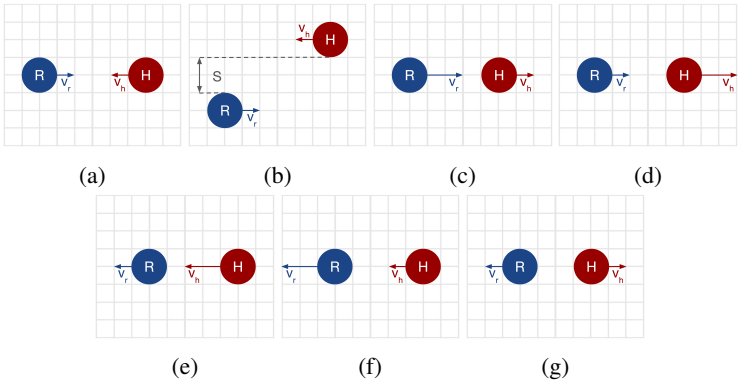

Figure 1.2: 2D symbolic representation of the prototypical human-robot collaborative situations. The module of the vector represents the robot's (R) and human's (H) velocity's magnitude (v_r and v_h). (a) Both agents approaching with probable contact. (b) Both agents approaching without probable contact. (c) Robot following the human with probable contact (the robot moves faster than the human). (d) Robot following the human without probable contact. (e) Human following the robot with probable contact (the human moves faster). (f) Human following the robot without probable contact. (g) Both agents getting away. Of all seven cases only three can trigger a safety stop due to a potential contact: (a), (c), and (e).

could not control in real scenarios, we wanted to see how the three approaches reacted when it changed. Note that the human and the robot moved in one axis. We randomly generated 1000 different human's velocities uniformly distributed within an interval of interest for each case. We also added at each simulation timestamp, different white Gaussian noise to the three axes of each of the velocities.

The power of the noise for all the three axes motion was -20 dBW. The noise allowed us to simulate a realistic human motion, not just a straight line movement, as it can be seen in Fig. 1.3. The simulation time (10s) was selected to match the robot's speed (0.2m/s) and the maximum reach (1m) of the robot used in this work, Kinova Gen3. The frequency was 100Hz, and the $ttstop$ (see line 6 in Alg. 1) and the T_b (see Eq. 1.7) were both set to 0.5s. The radii of the spheres representing the human and the robot were fixed to 0.05m. The robot's nominal speed was 0.2m/s. The initial distance between the robot and the human was 0.9m, for case A, and 0.4m for cases C and E. The intervals for the randomly generated human's speeds for the cases A, C and E were: $[0.3, 0.6]$, $[0.0, 0.15]$, and $[0.3, 0.6]$, respectively. Recall that the motion was mainly in one axis, although we added noise to it. Each case's main direction is depicted in Fig. 1.2.

Before a safety stop was issued, we evaluated the time the robot was moving, and the final Euclidean distance between the human and the robot. The time before stopping allowed us to validate whether our method delayed the safety stop or not. We reported the distance to show that our method allows the robot to get closer while still being safe (no collision).

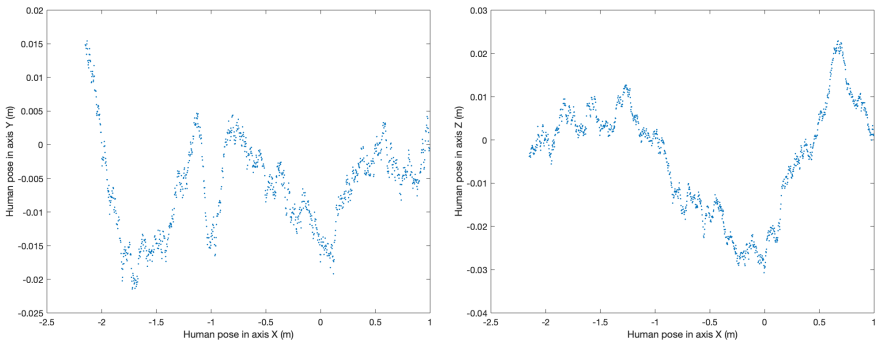

(a) (b)

Figure 1.3: Example of the human pose evolution for a single simulated noisy motion. (a) and (b) show the evolution of the human pose along the planes XY and XZ, respectively. Note that we simulated a noisy human velocity at each simulation timestamp aiming for a realistic human motion, therefore, the pose evolution did not follow a straight trajectory.

1.5.1.1 Evaluation I - First hypothesis: the robot moves more time

We conducted a statistical analysis to evaluate the significance of the proposed approach's improvement in delaying the safety stop w.r.t. the ISO and the Fuzzy ISO methods. We measured the time the robot moved before the safety stop was issued. Fig. 1.4 shows the distributions of measured time for each case and evaluated approach.

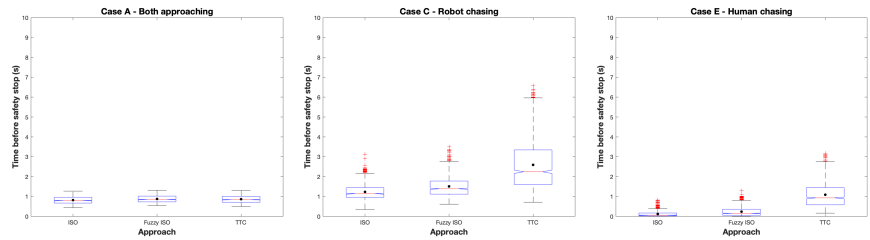

(a) (b) (c)

Figure 1.4: Distributions of the time the robot moved (a, b, c) for each simulated case and approach. The mean value is represented by a dot. TTC clearly outperformed the others in (b) and (c).

We manipulated the three different methods (independent variable) and assessed them with respect to the time before stopping (dependent variable), for each of the prototypical cases. First, we used Kruskal-Wallis to evaluate if there was a statistically significant difference in group mean, obtaining $\chi^2(2) = 55.08, p < 0.001$, for case A, $\chi^2(2) = 1113, p < 0.001$, case C, and $\chi^2(2) = 1812.8, p < 0.001$, case E.

Second, as the results rejected the null hypothesis, we carried out a post-hoc analysis to find out where the differences occurred between the groups. A Dunn & Sidák post-hoc multiple comparison test revealed a significant pairwise difference between our method and the other two in two of the cases: C and E. In case C, the time before braking using TTC was significantly different from the ones using ISO and Fuzzy ISO with a p-value of 0 in both cases. In case E, TTC was also significantly different to ISO and Fuzzy ISO, with a p-value of 0 for both comparisons. These results proved that the differences in the values depicted in Figs. 1.4b and 1.4c are statistically significant. Hence, in these two cases, TTC allowed the robot to move during a larger time before stopping. This fact decreased the human-robot protective distance. Specifically, in case C, the robot using TTC moved a 110.52% more time on average than with ISO, and a 71.81% more than with Fuzzy ISO. In case E, the improvement was far greater, a 802.68% w.r.t. using ISO and a 358.83% w.r.t. Fuzzy ISO. These improvements would be notorious in long-term collaborations at industries, especially in tasks that would imply medium and high levels of interaction (see Fig. 1.5). In case A, our method was significantly different from the ISO method with a p-value of $9.010e^{-8}$, while no significant differences were found between the Fuzzy ISO and TTC. However, the mean difference between TTC and ISO was really small (5.85%), as we can see in the values depicted in Fig. 1.4a. Note that in case A, the human and the robot approached each other, which is the default assumption that ISO's formulation does. Hence, we expected our method to capture the high risk of the situation and behaved similarly to the ISO.

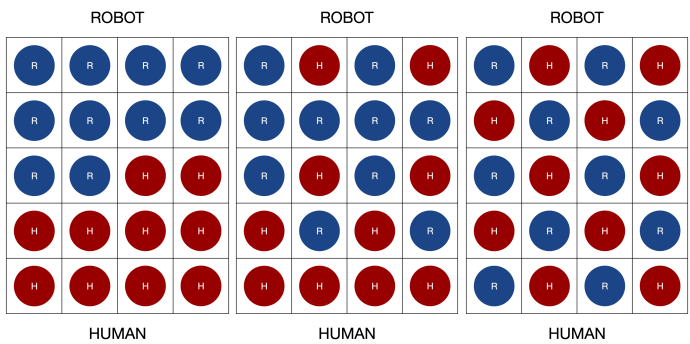

Figure 1.5: Examples of our task's distribution between the robot and the human with different levels of interaction. Blue circles (R) indicate the tray's compartments that the robot would fill. Red circles (H) would be filled by the human. (a) Low level of interaction with nearly zero possible crossing trajectories. (b) Medium level of interaction where a few of the robot's and human's targets might involve trajectories' intersections. (c) High level of interaction where the distribution of the task's targets would potentially cause crossing trajectories, leading to several probable contacts and safety stops.

1.5.1.2 Evaluation I - Second hypothesis: the robot gets closer but it is still safe

A second statistical analysis evaluated the significance of our approach's improvement in allowing closer but still safe human-robot distances. The analysis was again performed w.r.t. the ISO and the Fuzzy ISO methods. We measured the final human-robot distance once the safety stop was issued. Fig. 1.6 shows the distributions of measured distance for each case and evaluated approach.

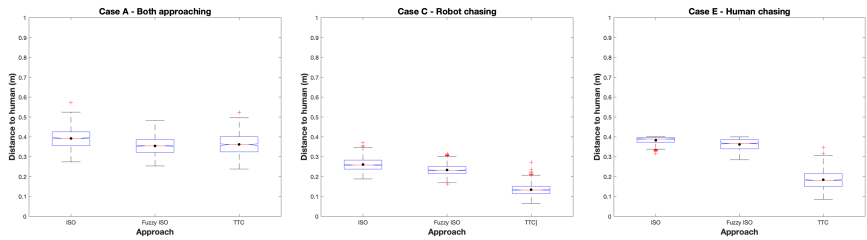

(a) (b) (c)

Figure 1.6: Distributions of the final distance to the human (a, b, c) for each simulated case and approach. The mean value is represented by a dot. TTC produces shorter distances in (b) and (c), but always ensuring safety and avoiding collisions.

We manipulated the three different methods (independent variable) and assessed them with respect to the human-robot distance after stopping (dependent variable). This is for each of the three prototypical cases. First, we used Kruskal-Wallis to evaluate if there was a statistically significant difference in group mean, obtaining $\chi^2(2) = 290.71, p < 0.001$, for case A, $\chi^2(2) = 2138, p < 0.001$, case C, and $\chi^2(2) = 2132.5$, $p < 0.001$, case E. Second, as the results rejected the null hypothesis, we carried out a post-hoc analysis to find out where the differences occurred between the groups. A Dunn & Sidák post-hoc multiple comparison test revealed a significant pairwise difference between our method and the other two in the three prototypical cases: A, C and E. In case A, the distance after braking using TTC was significantly different from the ones using ISO and Fuzzy ISO with a p-value of 0 and 0.0018, respectively. In cases C and E, TTC was also significantly different to ISO and Fuzzy ISO, with a p-value of 0 for both comparisons.

These results proved that the differences in the values depicted in Fig. 1.4 are statistically significant. However, in case A, the mean difference in the human-robot distance produced by TTC was really small w.r.t. the other two methods. Specifically, using TTC the average final distance was a 7.69% smaller than with ISO and a 2.19% larger than with Fuzzy ISO. In cases C and E, the differences were larger and TTC allowed the robot to get closer to the human before stopping. Specifically, in case C, the robot using TTC produced a reduction in the final human-robot of a 48,51% and a 42.46% w.r.t. ISO and Fuzzy ISO respectively. In case E, the reduction was a bit larger, a 51.82% w.r.t. using ISO and a 49.06% w.r.t. Fuzzy ISO. Reducing the final human-robot distance might affect safety, but TTC produced no collisions during the simulation. Furthermore, the percentage reduction in the final distance is

far smaller than the increase in the time the robot moves before stopping. Hence, we can say that TTC greatly improves productivity while slightly compromising safety.

1.5.2 Evaluation II - Real robot and simulated human (aiming for repeatability)

In this case, we contextualized the previous simulation-based evaluation, showing how our proposal may be useful in the task shown in Fig. 1.1. The objective was to evaluate our approach implemented in a real robot, avoiding the problem of repeatability of the human operator. Hence, we implemented a realistic setup, where a real Kinova Gen3 robot moved towards a specific target pose. Meanwhile, the system was fed with the position and the velocity of a simulated human. This simulated human moved accordingly to the three cases evaluated before. Specifically, for the situations shown in Figs. 1.2a and 1.2c, the real robot moved from the grasp pose to the release pose of one of the tray's compartments. This would be the case of a robot picking a token and trying to place it on a compartment. For the remaining situation, Fig. 1.2e, the real robot moved along the opposite trajectory. In this case, emulating when the robot would have already placed the token and it would go to pick a new one. In this evaluation, we simulated four human's velocities, and we compared the final Euclidean distance from the robot to the target pose after stopping. Note that the distance to the target is related to the two variables studied in Sec. 1.5.1: the robot's motion time and the distance to the human before stopping. Considering the previous evaluation's results, our hypothesis was that the robot would clearly be able to get closer to the target pose using our method in cases C and E. We used the robot Kinova Gen3, equipped with the 2F-85 two-finger gripper from Robotiq, and the same parameters as in Sec. 1.5.1. The four different simulated human's speeds were selected from the intervals we used in Sec. 1.5.1 for the cases A, C and E: $[0.3, 0.6]$, $[0.0, 0.15]$, and $[0.3, 0.6]$, respectively. Some videos of this evaluation are shown in the additional material[2].

Fig. 1.7 depicts the evolution of the distance to the target for the four human speeds and each case. These results corroborated, in this case using a real robot, what we already obtained in simulation in Sect 1.5.1. In case A (see Fig. 1.7a), the differences between the three approaches remained minimal, around 2%. While in cases C and E (see Figs. 1.7b and 1.7c) TTC again outperformed the other two approaches, allowing the robot to get closer to its target pose. Hence, our method let the robot get closer to finish its task (e.g. placing a token) before stopping. Specifically, in the case C, the final distance to the target using TTC was a 23.22% and a 19.92% shorter on average than with ISO, and Fuzzy ISO, respectively. In case E, the improvement is even greater, a 31.53% and a 29.16% shorter distance to target on average w.r.t. ISO and Fuzzy ISO, respectively. Fig. 1.8 depicts a comparison of the final robot's pose w.r.t. the target for the three methods in one of the simulated examples of the case E.

[2] www.iri.upc.edu/groups/perception/TTC

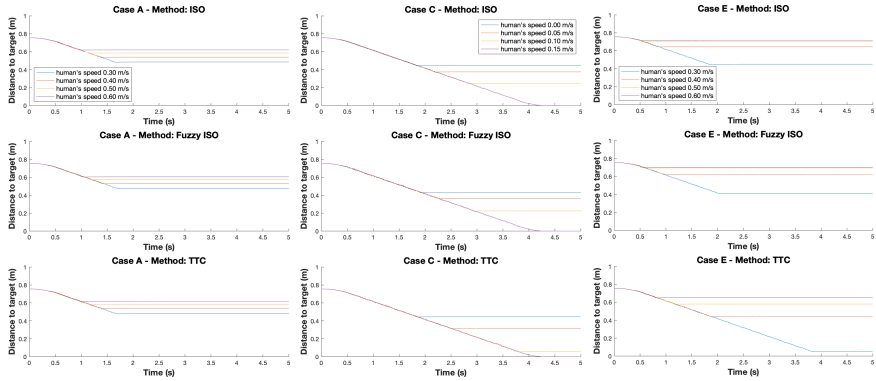

(a) (b) (c)

Figure 1.7: The plots show the evolution of the robot's distance to a target pose during the second evaluation: real robot and simulated human. (a) Both agents approaching with probable contact. (b) Robot following the human with probable contact. (c) Human following the robot with probable contact. Once the safety stop was triggered, the human's simulation finished and the robot remained stopped. This is why the distance to the target becomes constant in the plots, implicitly representing the time the robot was moving before the safety stop.

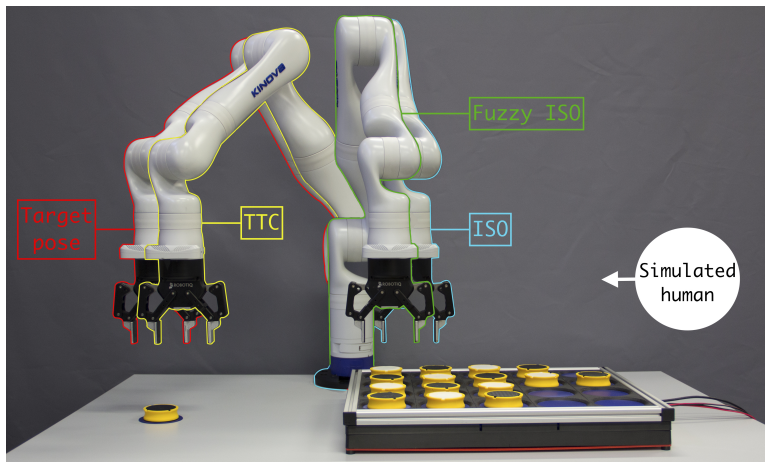

Figure 1.8: Exemplification of the final distance to the target pose during the second evaluation (case E): real robot and simulated human. The robot is trying to reach the target pose at 0.2m/s while the human is following the robot at a faster speed (0.3m/s in this case). In the image, we can see the target pose (red) and the final robot's pose after the safety stop issued by ISO (cyan), Fuzzy ISO (green) and TTC (yellow). As we can see, TTC allows the robot to get closer to the target before stopping.

1.5.3 Qualitative validation - Demo of a collaborative task with the real robot and a human

Finally, we present a qualitative validation of the implementation of our approach in a realistic scenario where a robot and a human share the task of filling the compartments of a tray [28]. The video of one of the task executions can be found in the additional material[2]. In this implementation, when the human filled one of the compartments the robot was meant to fill, the robot modified its plan and continued with the other free targets. Once a safety stop was issued, the robot was put in joint admittance mode (compliant), and it resumed the motion only after the human's command. The human's pose and velocity were detected using an HTC Vive tracker on their hand, and the tokens using an RFID-based board for fast and precise detection. The measurements' rate of the HTC and the robot was 100Hz. The robot shared with the operator its interpretation of the collision's risk using the lights on the robot's base (see Fig. 1.9).

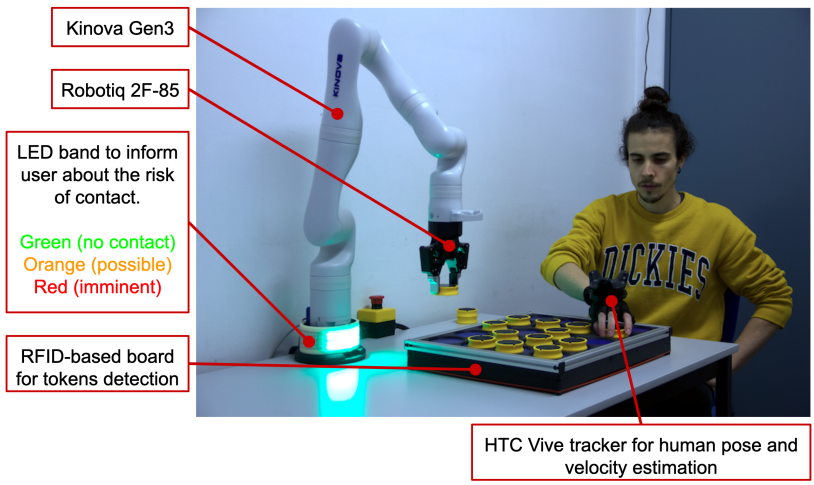

Figure 1.9: Setup for the demo of a collaborative task: filling a tray.

1.6 CONCLUSION

In this chapter, we studied the concept of time-to-contact (TTC) to issue a safety stop in close collaborative robotic scenarios. We proposed a TTC formulation and an algorithm to activate a robot safety stop when there is a potential contact. We evaluated our approach against two state-of-the-art methods in a set of prototypical cases. First, in simulation, where a statistical analysis was performed to study the significance of the results. Second, with a real robot and a simulated human, aiming for a more realistic evaluation while avoiding the issue of human repeatability. In both evaluations and two out of three cases, our approach clearly produced

a later safety stop than the other standard methods, increasing the time the robot is moving/working. In the remaining case, the differences between the three methods were too small to be relevant. Later stops resulted in shorter final distances between the human and the robot. However, the distance was always large enough to avoid collisions. Furthermore, the increment in the time the robot moves before stopping (productivity) is higher than the reduction in the final human-robot distance (safety). This work is a step forward to enabling robots to continue working more time before a safety stop is activated, and to get closer to humans while sharing the execution of tasks with them. In the future, we want to extend the quantitative evaluations done in this work to other tasks, robot's speeds, and users. Furthermore, we would like to qualitatively evaluate with users the three compared approaches in terms of the perceived safety. We also plan to consider all the human and the robot joints for the computation of TTC. This might be done re-using and adapting previous work on visual human pose estimation [29]. Finally, it might be interesting to implement an algorithm to continuously adapt the robot's speed using TTC, something that has already shown potential in landing operations for drones.

ACKNOWLEDGMENT

This work has been partially funded by MCIN/ AEI /10.13039/501100011033 under the project CHLOE-GRAPH (PID2020-119244GB-I00); by the "European Union (EU) NextGenerationEU/PRTR and by MCIN/ AEI /10.13039/501100011033 under the project COHERENT (PCI2020-120718-2); and by the European Commission NextGenerationEU, through CSIC's Thematic Platforms (PTI+ Neuro-Aging). A. Olivares-Alarcos is supported by the European Social Fund and the Ministry of Business and Knowledge of Catalonia through the FI 2020 grant.

References

[1] ISO I. 10218 Robots and robotic devices-Safety requirements for industrial robots - Part 1: Robots. The International Organization for Standardization. 2011;.

[2] ISO I. 10218 Robots and robotic devices-Safety requirements for industrial robots - Part 2: Robot systems and integration. The International Organization for Standardization. 2011;.

[3] ISO I. TS 15066 Robots and robotic devices-Collaborative robots. The International Organization for Standardization. 2016;.

[4] Marvel JA, Norcross R. Implementing speed and separation monitoring in collaborative robot workcells. Robotics and computer-integrated manufacturing. 2017;44:144–155.

[5] Vogel C, Elkmann N. Novel Safety Concept for Safeguarding and Supporting Humans in Human-Robot Shared Workplaces with High-Payload Robots in Industrial Applications. In: Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction.

HRI '17. New York, NY, USA: Association for Computing Machinery; 2017. p. 315–316.

[6] Nikolakis N, Maratos V, Makris S. A cyber physical system (CPS) approach for safe human-robot collaboration in a shared workplace. Robotics and Computer-Integrated Manufacturing. 2019;56:233 – 243.

[7] Magrini E, Ferraguti F, Ronga AJ, et al. Human-robot coexistence and interaction in open industrial cells. Robotics and Computer-Integrated Manufacturing. 2020;61:101846.

[8] Zanchettin AM, Ceriani NM, Rocco P, et al. Safety in human-robot collaborative manufacturing environments: Metrics and control. IEEE Transactions on Automation Science and Engineering. 2015;13(2):882–893.

[9] Vicentini F, Giussani M, Tosatti LM. Trajectory-dependent safe distances in human-robot interaction. In: Proceedings of the 2014 IEEE Emerging Technology and Factory Automation (ETFA); 2014. p. 1–4.

[10] Campomaggiore A, Costanzo M, Lettera G, et al. A Fuzzy Inference Approach to Control Robot Speed in Human-robot Shared Workspaces. In: 16th International Conference on Informatics in Control, Automation and Robotics, ICINCO 2019. vol. 2. SciTePress; 2019. p. 78–87.

[11] Hecht H, Savelsbergh G. Time-to-contact. vol. 135. Elsevier; 2004.

[12] Lasota PA, Rossano GF, Shah JA. Toward safe close-proximity human-robot interaction with standard industrial robots. In: 2014 IEEE International Conference on Automation Science and Engineering (CASE). IEEE; 2014. p. 339–344.

[13] Joseph L, Pickard JK, Padois V, et al. Online velocity constraint adaptation for safe and efficient human-robot workspace sharing; 2019. Working paper or preprint.

[14] Rosenstrauch MJ, Pannen TJ, Krüger J. Human robot collaboration - using kinect v2 for ISO/TS 15066 speed and separation monitoring. Procedia CIRP. 2018;76:183 – 186. 7th CIRP Conference on Assembly Technologies and Systems (CATS 2018).

[15] Byner C, Matthias B, Ding H. Dynamic speed and separation monitoring for collaborative robot applications – Concepts and performance. Robotics and Computer-Integrated Manufacturing. 2019;58:239 – 252.

[16] Li Y, Zhang L, Song Y. A vehicular collision warning algorithm based on the time-to-collision estimation under connected environment. In: 2016 14th International Conference on Control, Automation, Robotics and Vision (ICARCV); 2016. p. 1–4.

[17] Qu C, Qi W, Wu P. A High Precision and Efficient Time-to-Collision Algorithm for Collision Warning Based V2X Applications. In: 2018 2nd International Conference on Robotics and Automation Sciences (ICRAS); 2018. p. 1–5.

[18] Song W, Yang Y, Fu M, et al. Real-Time Obstacles Detection and Status Classification for Collision Warning in a Vehicle Active Safety System. IEEE Transactions on Intelligent Transportation Systems. 2018;19(3):758–773.

[19] Yang W, Zhang X, Lei Q, et al. Research on longitudinal active collision avoidance of autonomous emergency braking pedestrian system (AEB-P). Sensors. 2019;19(21):4671.

[20] Kendoul F. Four-dimensional guidance and control of movement using time-to-contact: Application to automated docking and landing of unmanned rotorcraft systems. The International Journal of Robotics Research. 2014;33(2):237–267.

[21] Zhang H, Cheng B, Zhao J. Extended tau theory for robot motion control. In: 2017 IEEE International Conference on Robotics and Automation (ICRA); 2017. p. 5321–5326.

[22] Eguíluz AG, Rodríguez-Gómez JP, Martínez-de Dios JR, et al. Asynchronous Event-based Line Tracking for Time-to-Contact Maneuvers in UAS. In: 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS); 2020. p. 5978–5985.

[23] Kaneta Y, Hagisaka Y, Ito K. Determination of time to contact and application to timing control of mobile robot. In: 2010 IEEE International Conference on Robotics and Biomimetics; 2010. p. 161–166.

[24] Alenyà G, Nègre A, Crowley JL. A comparison of three methods for measure of time to contact. In: 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems. IEEE; 2009. p. 4565–4570.

[25] Kaneta Y, Hagisaka Y, Ito K. Determination of time to contact and application to timing control of mobile robot. In: 2010 IEEE international conference on robotics and biomimetics. IEEE; 2010. p. 161–166.

[26] Garcia AJS, Figueroa HVR, Hernandez AM, et al. Estimation of time-to-contact from Tau-margin and statistical analysis of behavior. In: 2016 International Conference on Systems, Signals and Image Processing (IWSSIP). IEEE; 2016. p. 1–6.

[27] Hou J, List GF, Guo X. New algorithms for computing the time-to-collision in freeway traffic simulation models. Computational intelligence and neuroscience. 2014;2014.

[28] Olivares-Alarcos A, Foix S, Borgo S, et al. OCRA – An ontology for collaborative robotics and adaptation. Computers in Industry. 2022;138:103627.

[29] Gassó Loncan Vallecillo JC, Olivares-Alarcos A, Alenyà G. Visual feedback for humans about robots' perception in collaborative environments. Institut de Robòtica i Informàtica Industrial, CSIC-UPC; 2020. IRI-TR-20-03.