Strategies for Safety in Human Robot Interaction

D. Kulić   E. A. Croft

Department of Mechanical Engineering
University of British Columbia
2324 Main Mall
Vancouver, BC, V6T 1Z4, Canada

Abstract

This paper presents a hierarchical path planning and control strategy for ensuring safety during a human-robot interaction. At the planning stage, a two-step process is used where first the danger of the interaction is minimized, followed by a goal seeking optimization. This approach reduces the likelihood of encountering local minima due to conflicts between reducing danger and a demanded interaction task. At the control stage, the human intent signal is evaluated at every step to ensure safe operation of the robot. In initial simulation work, the controller drives the robot away from the planned interaction if danger is identified, and then allows the planned interaction to resume once the danger has passed.

1. Introduction

Two key issues hampering the entry of robots into unstructured environments populated by humans are safety and dependability [2,3]. In industrial applications, where the use of robotics is widespread, the safety of human–robot interaction is effected by isolating the robot from the human [2,4,5] (in effect there is no interaction). Specifically, in North America, the primary robot safety industrial standard is the ANSI/RIA 15.06 [4]. The standard prescribes that safety is achieved by defining a region around the robot that must be safeguarded. The prescribed action to be taken by the robot system upon detecting an intrusion into the safeguarding space is an emergency stop, which must remove all drive power and all other energy sources. The European standard, EN-775 contains similar provisions for robot safety [5]. However, as robots move from isolated industrial environments to interactive environments, this approach is no longer tenable [2], precluding the use of these standards. However, the concepts of risk assessment [4,6,7] contained in these standards can still be applied.

Once the potential hazards of a robotic system have been identified, there are three main approaches to reduce or eliminate the risk of the hazard: (i) by redesigning the system to eliminate the hazard, (ii) by controlling the hazard through electronic or physical safeguards, and, (iii) by warning the operator, either during operation or by training [4,6].

Industrial experience indicates that eliminating hazards by design is the most effective risk reduction strategy [4,7]. This principle has also been applied to service robotics. Examples include a whole-body robot visco-elastic covering [8,9] and the use of spherical and compliant joints [9,10]. However, in unstructured environments, mechanical design alone is not adequate to ensure safe and human friendly interaction. Additional safety measures need to be implemented through system control.

Control safeguards determine whether a hazardous configuration in the human-robot interface exists. The space around the robot to be safeguarded can be a set distance [8,11], or it can be sized by evaluating the potential danger of a human-robot interaction based on the robot configuration and the motion of the robot relative to the human [12,13]. Ikuta et al. [12] developed a danger evaluation method using the potential impact force as an evaluation measure.

Once the level of danger has been assessed, the controller must determine what corrective action (if any) is required. Bearveldt [11] defines three operating zones: no human in the work area, human in the work area but at a safe distance, and human dangerously close to robot. If no human is in the work area, the robot will operate at maximum speed. When a human is detected in the work area, but is still at a safe distance, the robot will operate at reduced speed. Once the human enters the unsafe area, an emergency stop is issued and all robot motion stops. Similar zones are also proposed in [14].

Yamada et al. [8] combine mechanical safety measures with the safeguarding concept. The robot is covered with a viscoelastic covering that both attenuates the impact force between the robot and the human and signals that the surface has been contacted. The space occupied by the viscoelastic covering is considered the safeguarded zone. If contact is detected, the robot's velocity is reduced to allow the operator to react and move away from the contact. If the robot motor torques rapidly increase, an emergency stop is generated.

The above systems [8,11] define three basic zones: the full-speed zone, the slow-down zone, and the emergency stop zone. However, the robot path is not modified. In contrast, in their "elusive robot" design, Traver et al. [13] combine a danger index with an obstacle avoidance strategy. If the distance between the robot and the human falls below a certain threshold level, the robot will deviate from its planned trajectory to avoid human contact.

Another approach to safety control is to control by minimizing the impact force during any potential human robot contact [15,16]. While this approach does reduce the potential hazard during human-robot contact, it is not sufficient to ensure safety, because the potential for hazard will also depend on the point of contact, as well as the payload the robot may be carrying. In addition, to ensure user acceptance, the robot should have additional mechanisms to avoid impact when possible.

2. Approach

Although most existing approaches implement an emergency stop when a hazardous situation (or hazard level) is detected, in unstructured environments, this may not be the safest response. In this work, the proposed system acts to minimize the hazard, both in the planning and control stages. Specifically, in the planning stage, the path is computed to minimize potential impact with humans during operation. On the control side both trajectory modification (slow down and stop) as well as avoidance strategies are utilized, based on the level of perceived danger. A key component of this approach is to improve the perceptive abilities of the robot in order to improve safety during interaction. In particular, knowledge of the user's reactions to robot movements can provide valuable information during control of the robot.

An overview of the system is presented in Figure 1. The user issues a command to the robot to initiate the interaction. The command interpreter translates the natural language command (e.g.: pick up the red cup) into a set of target locations and actions (e.g., execute a grip maneuver at coordinates [x,y,z]). The planner module then begins planning a safe path for the robot. During the interaction, the user is monitored to assess the level of approval of robot actions. This information is then used to modify the velocity of the robot along the planned path. The safety control module initiates deviation from the planned path if a change in the environment is detected which threatens the safety of the interaction. At that point, the safety module will initiate a re-assessment of the plan and initiate re-planning if necessary.

To ensure the safety and intuitiveness of the interaction, the complete system must incorporate (i) safe mechanical design, (ii) human friendly interfaces such as natural language interaction and (iii) safe control and planning strategies. Our work focuses on the design of the control and planning strategies. This work can be further divided into three key components: planning, intent, and control. The intent component of our method is described in the companion paper [17], while this paper focuses on the planning and safety aspects.

2.1. Planning Strategy

Most of the existing approaches to ensuring safety through control focus on reacting once a hazard is perceived. By including safety in the planning stage, the potential for hazard can be reduced, and the robot can be placed in a better position to respond to the hazard. For this reason, safe planning is an important component of the safety strategy.

When selecting a path planning strategy, there is a tradeoff between fast local methods that may fail to find the goal, and slow global methods [18]. To exploit advantages of both methods, recent path planning algorithms have used a hybrid approach, where global path planning is used to find a coarse region through which the robot should pass, and local methods are used to find the exact path through the region [19]. Similarly, the planning module in this approach generates a contiguous set of regions that together describe a safe path region. It is then left to the trajectory planner and the safety module to generate the exact path in the region, and the trajectory along that path.

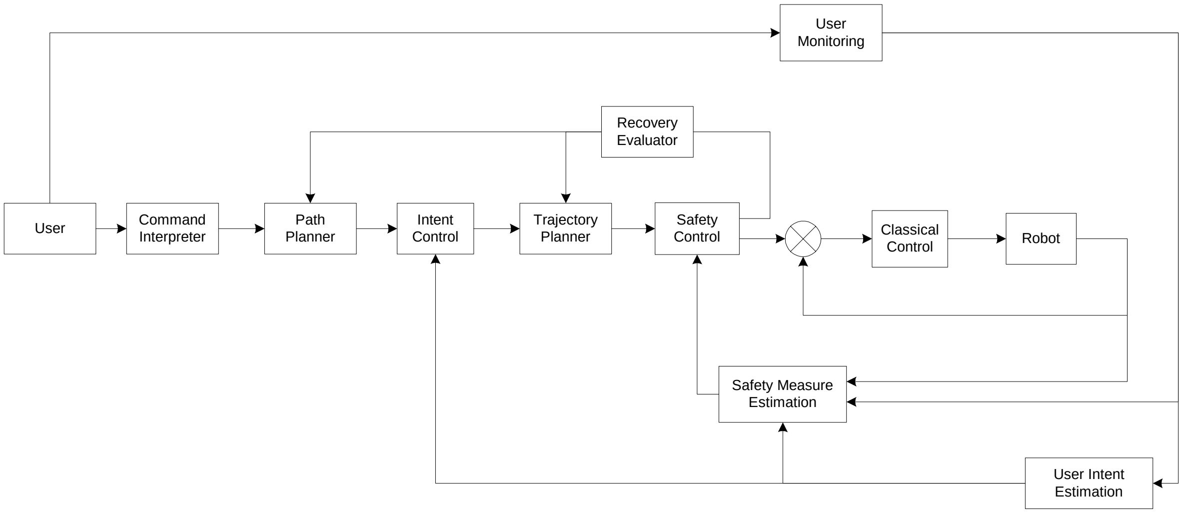

Figure 1 - System Overview Diagram

Another issue to consider when developing a planning strategy in human-robot interaction is the representation of the human. If the goal of the interaction is for the robot to approach and/or contact the human, then it is not appropriate to represent the human simply as an obstacle as in [13]. In this work, during pre-planning, each segment of the path is classified as interactive or non-interactive. If the segment is classified as non-interactive, the entire region of space occupied by the human (described by a set of spheres) is treated as an obstacle. If the segment is classified as interactive, a smaller set of spheres is used, such that the target area of the human (for example, the hand) is excluded from the obstacle area.

The safest path can be found by searching for contiguous regions that: remain free of obstacles, lead to the goal, and minimize a measure of danger (a danger criterion). Since path planning (as opposed to trajectory planning) does not consider robot velocities, a configuration-based danger criterion is required. Several measures can be considered: relative distance between the robot and the human, the robot stiffness, the robot inertia, or some combination of these measures [12]. Herein, a criterion based on inertia and distance is used for preliminary analysis (Section 5). However, regardless of the measure used, it is likely that these criteria will conflict with the goal seeking criteria during the search, leading to local minima and very long search times. To avoid this problem, a two-stage search is proposed. In the first stage, maximum priority is placed on minimizing the danger criterion. A threshold is established for determining when an acceptable level of danger is achieved. Once a path is found which places the robot below this threshold, the second stage of the search is initiated. In this stage, maximum priority is placed on the goal-seeking criterion. In the resulting overall path, the robot will spend most of its time in low danger regions. One can note that this approach will not result in the shortest distance path. The tradeoff between increased safety and reduced distance can be controlled by modifying the threshold where the switch from the first stage to the second stage occurs. A simplified version of the algorithm has been implemented for initial testing and is described in Section 3.

Several concerns exist with this approach. The first issue is the completeness of the planner. Depending on the search method chosen, the planner may not be able to find a path given the constraints of the environment and the danger criterion. A complete method will take an unacceptably long time to complete for robots with a large number of degrees of freedom (DOFs) [18], or in cases when the search constraints conflict. If randomized planning is used, a solution may be found faster in favorable conditions, but can also take indefinitely long if the search space is highly constrained. The problem is exacerbated with this system because an additional constraint (the danger criterion) is added to the search formulation. In particular, if there are several obstacles positioned close to the robot, it may not be possible to complete the stage 1 search within the given threshold. In this case, the user would be notified that a safe path cannot be found in the current environment.

2.2. Real-Time Safety Module

The safety module is responsible for controlling the danger index of the interaction in real time. As opposed to the danger criterion, this index incorporates dynamic measurements, such as velocity, distance and intent. This module is responsible for reacting to sudden changes in the environment, not anticipated during the planning stage. The inputs into the safety module consist of the proposed next configuration of the robot from the trajectory planner, which includes the velocity and acceleration information, the current user configuration, and an estimate of the user's level of intent. Based on this information, the safety module evaluates the danger index at the proposed next step. If the danger index is acceptable, the proposed plan can proceed, otherwise, a corrective decision is made and an alternate configuration is passed to the low-level controller.

The key element of the safety module is the estimation of the danger index. As suggested by Ikuta et al. [12], a "danger measure" should include distance between the robot and the human, the relative velocity between them, as well as the inertia and stiffness of the robot. In addition to these elements, the "danger measure" should include a measure of the intent of the human [17]. The danger index estimation algorithm needs to combine all of the above elements into a single estimate and at the same time manage the uncertainty associated with estimating and combining these elements. In this initiatory work, a fuzzy logic estimator is utilized for danger assessment. The implementation is described in Section 3. Once the danger index has been estimated, if corrective action is required, the safety module implements one step ahead planning to minimize the danger index. This is, in effect, a real time implementation of the potential-field approach, using the gradient of the danger index as the potential field.

After the corrective action has been initiated, the Evaluator Module is also activated to determine the cause of the corrective action, and if a recovery to complete the task is possible. Three possible outcomes of the evaluation are possible:

1. The gross path plan is still valid, but local replanning is needed. This case is handled by [19].
2. Re-planning of the path or portion of the path is required.
3. A retreat is necessary and the mission must be abandoned. In this case, new instructions would be requested from the user.

As shown in Figure 1, if local re-planning is needed, the trajectory planner is reactivated; if large-scale re-planning is needed, the planning module is reactivated.

One concern with the real-time safety module is control stability. To ensure that the system does not oscillate at the threshold, a hysteresis proportional to the expected uncertainty of the danger index is introduced around the threshold. A more thorough stability analysis of the controller will be performed once the controller design is finalized.

3. Simulations

Tests of the feasibility of this approach were carried out in a simulation environment. The simulation consists of a simple 3 link planar robot operating in two-dimensional space. The task of the robot is to pick up an object and deliver it to the person's hand, in a basic handover task.

The planning module uses the best first planning approach, which is suitable for low DOF robots [17]. The cost function being minimized consists of a quadratic goal seeking function, quadratic obstacle avoidance function, and a safety measure function. The safety measure function is a measure of the inertia of the robot configuration and a measure of the distance of the end effector from the straight line leading to the base of the robot. The preliminary cost function for the danger criterion is given by Equation 1 below:

$$S = W_i \cdot I + W_d \cdot d^2 \tag{1}$$

S is the cost measure, I is the inertia of the robot calculated around the robot base, d is the distance from the end effector to the line joining the robot base to the starting end effector position, and W_i and W_d are relative weights of the inertia and distance term. In the simulation below, $W_i = 0.2$ and $W_d = 0.8$ were used. The trajectory planner used in the simulation is a simplified version of the quintic trajectory planner proposed in [20]. The safety module uses a fuzzy estimator to implement the danger estimation task. The distance between the robot and the human, the velocity of the robot towards the human and the estimated intent level are used as inputs into the fuzzy estimator.

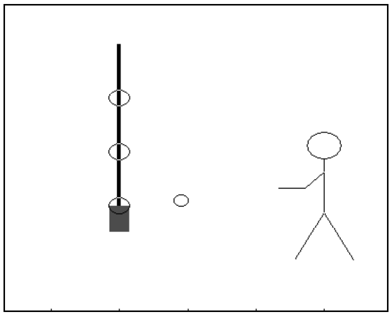

Figure 2 - t = 0s

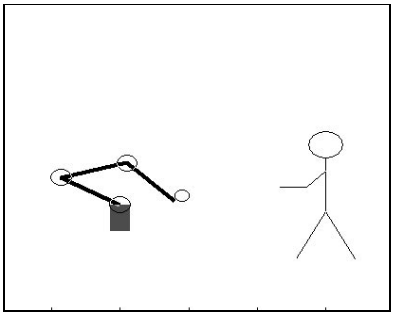

Figure 3 - t = 7.7s

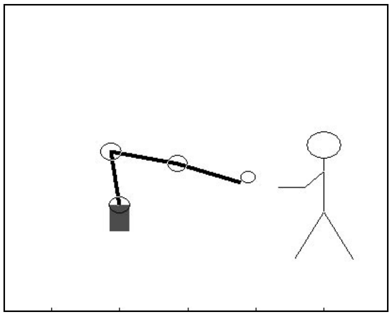

Figure 4 - t = 11.2s

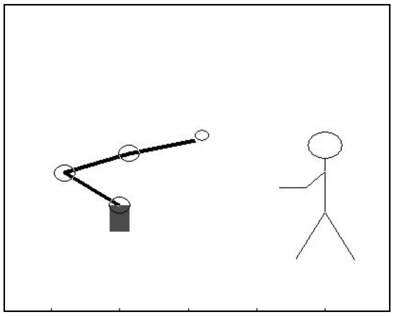

Figure 5 - t = 13.1s

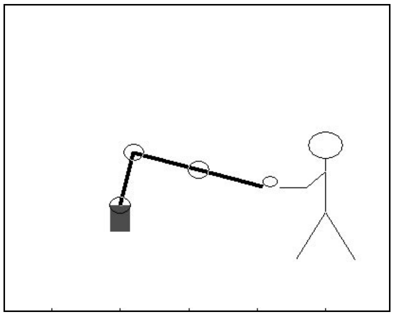

Figure 6 - t = 18.6s

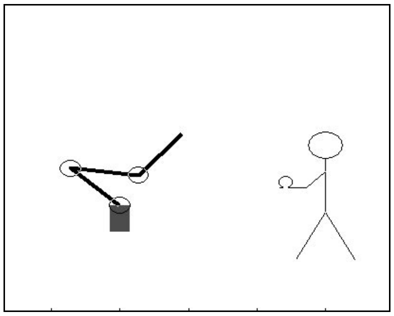

Figure 7 - t = 21.5s

The output of the fuzzy estimator is a level of danger. The rulebase relates the relative distance, velocity and intent level to the danger index estimate using linguistic rules. For example, if the distance is low, and the velocity is positive high, and the intent level is low, the danger index is high. A partial rulebase is given in Table 1. Measurement of the intent level is discussed in the companion paper [17].

Table 1 - Partial Danger Index Rule Base

If (Distance is LOW) and (Intent is LOW) then (Danger is HIGH)
If (Velocity is VPOS) and (Intent is LOW) then (Danger is HIGH)
If (Velocity is POS) and (Intent is LOW) then (Danger is HIGH)
If (Distance is MED) and (Velocity is VPOS) and (Intent is LOW) then (Danger is HIGH)
If (Distance is MED) and (Velocity is POS) and (Intent is LOW) then (Danger is HIGH)

A simplified version of the evaluator, which initiates local re-planning, is implemented for this initial simulation.

Figures 2 – 7 show a sequence from a sample simulation. Figure 8 shows the intent signal. In Figure 2 the robot is in its initial starting position. The robot first lowers its inertia, and then goes towards the object. In Figure 3, the robot picks up the object. After picking up the object, the robot moves towards the person.

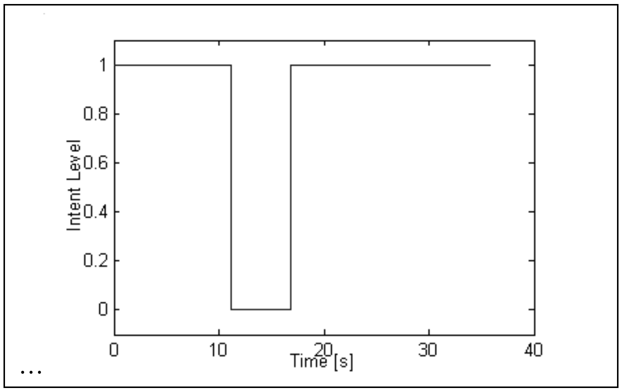

Figure 8 - Intent Signal

In Figure 4, as the robot approaches the person, the intent drops to zero. At this point, the safety module takes corrective action and retracts the robot away from the person, as shown in Figure 5. The robot will stay in the position shown in Figure 4 until the intent value is again increased. After the intent value is returned to 1, the robot resumes its mission and approaches the person, as shown in Figure 6. Figure 7 shows the robot retracting to reduce its inertia before returning to its initial position.

4. Conclusions and Future Work

The initial simulations show the feasibility of our approach. Further work is needed to develop and evaluate different safety measure estimation algorithms, for both the planning and the real-time safety modules. As well, studies of both the stability of the proposed system, and human response to the system will be undertaken.

Acknowledgements

This work is supported by the National Science and Engineering Research Council of Canada.

References

[1] J. M. Wiener et al. Measuring the Activities of Daily Living: Comparisons Across National Surveys. Journal of Gerontology: Social Sciences, 45(6), pp. 229–237, 1990.

[2] P. I. Corke. Safety of advanced robots in human environments. Discussion paper for IARP. http://www.international-robotics.org/wg/dependability/safety.pdf

[3] C. W. Lee, Z. Bien, G. Giralt, P. Corke and M. Kim. Report on the First IART/IEEE-RAS Joint Workshop: Technical Challenge for Dependable Robots in Human Environments, 2001.

[4] ANSI. American National Standard for Industrial Robots and Robot Systems – Safety Requirements. American National Standard Institute, New York, NY, RIA/ANSI R15.06 – 1999.

[5] S. P. Gaskill and S. R. G. Went. Safety issues in modern applications of robots. Reliability Engineering and System Safety, Vol. 52, pp. 301 – 307, 1996.

[6] ISO. Safety of machinery – Principles of risk assessment. International Standards Association, Geneva, Switzerland, ISO 14121:1999.

[7] IEC. Functional safety of electrical/ electronic/ programmable electronic safety-related systems – Part 1: General Requirements. IEC 61508 – 1: 1998-12.

[8] Y. Yamada et al. Human – Robot Contact in the Safeguarding Space. IEEE/ASME Transactions on Mechatronics, 2(4), pp. 230 – 236, 1997.

[9] Y. Yamada et al. FTA-Based Issues on Securing Human Safety in a Human/Robot Coexistance System. IEEE SMC'99, Vol. 2, pp. 1058 – 1063, 1999.

[10] A. Bicchi et al. Compliant design for intrinsic safety: General issues and preliminary design. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, pp. 1864 – 1869, 2001.

[11] A. J. Bearveldt. Cooperation between Man and Robot: Interface and Safety. Proc. of the IEEE Int. Workshop on Robot Human Communication, pp. 183 – 187, 1993.

[12] K. Ikuta et al. Safety Evaluation Method of Human-Care Robot Control. Proc. Int. Symp. Micro - mechatronics and Human Science, pp. 119 – 127, 2000.

[13] V. J. Traver et al. Making Service Robots Human-Safe. Poc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, pp. 696 – 701, 2000.

[14] J. Zurada et al. A Neuro-Fuzzy Approach for Robot System Safety. IEEE Trans. on Systems, Man and Cybernetics – Part C: Applications and Reviews, 31(1), pp. 49 – 64, 2001.

[15] J. Heinzmann and A. Zelinsky. Building Human – Friendly Robot Systems. Proc. Int. Symp. of Robotics Research, pp. 305 – 312, 1999.

[16] Y. Matsumoto et al. The Essential Components of Human – Friendly Robot Systems. Proc. Int. Conf. On Field and Service Robotics, pp. 43-51, 1999.

[17] D. Kulic and E. A. Croft. Intent Based Control for Human Robot Interaction. ICAR 2003.

[18] J.C. Latombe. Robot Motion Planning. Kluwer Academic Publishers, Boston, MA. 1991.

[19] O. Brock and O. Khatib. Real-Time Replanning in High-Dimensional Configuration Spaces Using Sets of Homotopic Paths. Proc. IEEE Int. Conf. on Robotics and Automation, pp. 550 – 555, 2000.

[20] S. Macfarlane and E. A. Croft. Jerk-Bounded Robot Trajectory Planning - Design for Real-Time Applications. IEEE Trans. on Robotics and Automation, 19(1), pp. 42 – 52, 2003.