Doctoral Dissertation
Doctoral Program in Mechanical Engineering

Effective and safe framework for human-robot interaction

Michele Polito

* * * * * *

Supervisors
Prof. Stefano Pastorelli, Supervisor
Prof. Laura Gastaldi, Co-Supervisor

Politecnico di Torino
January 31, 2026

Abstract

Human–robot collaboration requires robotic systems capable of operating safely and efficiently in close proximity to human operators, while adapting to dynamic and uncertain environments. Although current collaborative robots comply with international safety standards, their capabilities are often limited to conservative and reactive behaviors, which restrict the exploitation of true human–robot synergy. This doctoral thesis addresses these limitations by proposing an integrated, human-centered framework that tightly couples perception, safety assessment, and motion planning to enable adaptive and context-aware collaboration.

The main contributions of this work focus on the perception and cognition layers of the interaction loop. A multimodal human arm tracking system is developed by fusing wearable inertial measurement units with an event-based vision sensor through a Kalman filter–based framework. The complementary characteristics of inertial sensing and event-driven vision enable accurate and low-latency upper-limb tracking while reducing computational load. Experimental results demonstrate improved accuracy and stability compared to unimodal tracking approaches.

The same inertial sensing infrastructure is further exploited to assess the operator's motion safety state. A machine learning–based classifier is trained to detect abrupt and potentially hazardous movements in real time, allowing the system to distinguish between nominal task execution and safety-critical behaviors. The network achieves recognition accuracy above 93% during testing. When integrated into the complete system, the detection pipeline enables timely safety responses, including robot stopping actions, with an overall reaction time below 0.5 s.

At the planning level, a reinforcement learning–based framework is proposed to generate human-aware robot motions. Two learning environments are designed: a joint-control environment, used to progressively train the robot to accurately follow target trajectories, and a path-planning environment, in which the agent learns to generate collision-free Cartesian paths in the presence of a human operator and environmental constraints. Safety, kinematic feasibility, and collision avoidance are embedded directly into the reward functions, allowing the learned policies to internalize safety constraints rather than relying on external supervisory

logic. The path planner proved to be an effective approach for generating collision-free trajectories with a success rate exceeding 90% within the specified workspace.

All components are integrated within a modular ROS 2–based architecture that enables real-time communication between perception, safety assessment, and motion planning modules. The complete framework is validated on a representative collaborative assembly task, where the robot and the human operator share the workspace. The experimental results demonstrate that the proposed architecture can safely manage human presence, react to abrupt operator movements, and generate adaptive robot trajectories consistent with the interaction context.

Overall, this thesis shows how multimodal perception, data-driven safety assessment, and reinforcement learning–based planning can be cohesively integrated to move beyond purely reactive safety mechanisms, contributing toward adaptive and robust human–robot collaboration.