

Safe, Fluent and Acceptable Motion Generation and Execution for Human–Robot Interaction in Manufacturing Environments

Thibaut Lopez¹, Sébastien Chevron¹, Olivier Aycard¹, Pierre-Brice Wieber², Philip Scales¹, Mohamed Boua³, Christine Jeoffrion³

Abstract—Robots operating in human environments must not only ensure physical safety but also exhibit behaviors that are understandable, fluent, and acceptable to human partners. This paper investigates motion generation strategies that combine safety guarantees with interaction quality considerations, such as motion smoothness and human comfort. While the design of robots capable of ensuring safety in shared human-robot environments has enabled closer and more advanced forms of interaction, these new proximity-based tasks require moving beyond purely technical considerations. In particular, robot behavior must also be addressed from psycho-cognitive and social perspectives. In this context, we argue for the relevance of integrating social-aware motion control into robotic systems. First, we identify the motion parameters that influence human perception and operator experience. Then, we implement a Model Predictive Control (MPC) framework that generates four distinct socially-informed robot behaviors. Finally, we conduct a user study to evaluate and validate these behaviors and assess their social impact on non-expert participants. The results demonstrate that variations in robot behavior significantly affect the perceived social acceptability of the system. These findings highlight the importance of incorporating human-centered considerations into motion generation strategies for robots operating in shared environments.

I. INTRODUCTION

Robotic pick-and-place in human-shared environments, that is, in spaces shared with untrained humans and evolving rules defined by users, is a key enabler for service and collaborative robotics. Such capabilities are expected to have a strong impact in domains such as healthcare, logistics, and public spaces. However, the safe deployment of robots operating in close proximity to untrained humans remains a critical bottleneck.

Thanks to advances in mechanical design and control, modern manipulator robots can now safely share their workspace with humans. In particular, these systems are designed such that potential collisions result in a low (5%) probability of serious injury, even at speeds as high as 2ms^{-1} [1].

This work is supported by the French National Research Agency in the framework of the “Investissements d’avenir” program (ANR-15-IDEX-02).

¹ Thibaut Lopez, Sébastien Chevron, Olivier Aycard and Philip Scales are with GIPSA Lab and Grenoble Institute of Technology, 38000 Grenoble, France `firstname.lastname@univ-grenoble-alpes.fr`

² Pierre-Brice Wieber is with Inria, 38000 Grenoble, France `pierre-brice.wieber@inria.fr`

³ Mohamed Boua and Christine Jeoffrion are with LIP/PC2S and Univ. Grenoble Alpes, Univ. Savoie Mont Blanc, 38000 Grenoble, France `firstname.lastname@univ-grenoble-alpes.fr`

Nevertheless, collisions should still be avoided whenever possible, not only because of the remaining risk, however small, but also because each collision disrupts the tasks of both the robot and the human.

In previous work, we combined human motion perception and short-term prediction [2] with control strategies to generate real-time robot trajectories that ensure safety while preserving task continuity and motion fluency [3]. With respect to the ISO Technical Specification 15066 on collaborative robots, and unlike existing solutions based on emergency stops or fixed safety sensors [4], our approach combines aspects of *speed and separation monitoring* when actively avoiding collisions, with aspects of *power and force limiting* if a collision eventually occurs. The focus of our approach is on active collision avoidance, while we rely on the manipulator robot’s mechanical and control design to mitigate the risk of injury in the event of contact. Alternative approaches address this latter aspect more directly, for example by explicitly controlling the forces involved during collision [5].

While such approaches significantly improve physical safety, safety alone is not sufficient for effective human–robot collaboration. In real-world shared environments, robots must not only avoid collisions but also interact with humans in ways that are understandable, predictable, and comfortable. Recent research in human–robot interaction therefore emphasizes the importance of *interaction fluency*, defined as the ability of humans and robots to coordinate their actions smoothly and continuously during a shared task [6]. Motion characteristics such as smoothness, anticipation, and consistency with human expectations are essential for enabling humans to understand robot intentions and adapt their behavior accordingly.

Beyond fluency, *acceptability* has emerged as a key factor in the deployment of collaborative robots. The perceived comfort, trust, and social perception of robot behavior strongly influence the willingness of users to interact with robotic systems [7]. Parameters such as robot speed, interpersonal distance, and trajectory smoothness can significantly impact human perception of robot behavior and influence the quality of interaction.

These observations suggest that the design of robot motion in human-shared environments should jointly address three complementary objectives: ensuring physical safety, maintaining task efficiency, and generating motion that is both

fluent and socially acceptable. In this paper, we build upon predictive safety-aware motion generation approaches and extend them toward a human-aware framework that explicitly incorporates interaction fluency and acceptability constraints.

The contributions of this paper are threefold:

- First, we identify key motion parameters that influence human perception during human–robot interaction and derive a set of representative interaction behaviors associated with different social attitudes;
- Second, we design a motion control strategy capable of implementing these interaction behaviors while preserving the safety guarantees provided by our predictive control framework;
- Third, we evaluate the impact of these behaviors on the users’ perceived workload and social perception in real-world experiments involving non-expert users.

The remainder of this paper is organized as follows. First, we analyze the motion parameters that influence human perception and define representative interaction behaviors. Then, we briefly summarize the predictive control framework previously developed to ensure safety, and describe its extension to implement these interaction behaviors. Finally, we present experimental results obtained with non-expert users interacting with the robot in realistic conditions.

II. BEHAVIOR PARAMETERIZATION

In order to parameterize robot motion during human–robot interaction, we consider two fundamental control variables: the **maximum translational velocity** v_{max} and the **maximum interaction distance** d_{max} . These parameters are selected as they directly modulate both the dynamic properties of robot motion and the proxemic relationship between the robot and the human.

The choice of these parameters is strongly supported by prior work in human–robot interaction. In shared workspaces, the kinematic and geometric properties of robot motion have a direct impact on user comfort, perceived safety, and collaboration efficiency. In particular, robot velocity and interpersonal distance are identified as primary factors influencing human perception and interaction quality [7], [8].

Several studies have shown that robot velocity directly affects perceived safety and anticipation capabilities. Lower velocities improve motion predictability and allow the human to better anticipate robot behavior, whereas higher velocities tend to increase stress and reduce perceived safety [7], [15]. Similarly, the distance maintained between the robot and the human plays a critical role in interaction comfort, as larger distances improve anticipation and reduce perceived risk [8], [16].

Moreover, the distance parameter influences the degree of overlap between human and robot workspaces, which has been shown to affect cognitive load: larger overlaps require continuous monitoring from the human, increasing mental effort and reducing trust in the system [9].

Beyond velocity and distance, motion dynamics such as acceleration profiles also impact perception. Abrupt changes in motion can induce surprise or discomfort, whereas smooth

and progressive adaptations improve legibility and facilitate coordination [6], [10]. These elements highlight that motion is not only functional but also acts as a communication channel conveying robot intent.

In this work, we deliberately focus on velocity and interaction distance as primary parameters, as they are both fundamental, interpretable, and directly controllable within our system. They also constitute key dimensions of motion expressivity and proxemic interaction, making them suitable for systematically generating distinct robot behaviors.

Formally, we define:

$$v_{max} \in \mathbb{R}^+ \quad \text{and} \quad d_{max} \in \mathbb{R}^+ \quad (1)$$

Using these two parameters, we construct a discrete set of four motion regimes by combining two levels of each parameter (low/high velocity and short/long interaction distance):

- **Low** v_{max} , **Short** d_{max} : reactive and cautious behavior with late adaptation.
- **Low** v_{max} , **Long** d_{max} : anticipatory and conservative behavior promoting smooth interaction.
- **High** v_{max} , **Short** d_{max} : efficiency-driven behavior with reduced legibility.
- **High** v_{max} , **Long** d_{max} : anticipatory yet efficient behavior balancing performance and interaction quality.

This 2×2 parameterization enables a controlled exploration of the trade-off between efficiency, safety, cognitive load, and interaction quality. It is grounded in established findings in human–robot interaction, which highlight the central role of velocity, distance, and motion dynamics in shaping human perception of robot behavior. The subsequent user study, in section II aims to verify if these intended behaviors translate into actual differentiated human perceptions.

Beyond classical evaluation criteria in robotics, such as safety, efficiency, or trajectory smoothness, recent research in human–robot interaction has highlighted that robot motion inherently conveys socio-affective cues that shape human perception. This perspective is formalized by the concept of movement prosody, which draws an analogy between motion dynamics and vocal prosody in human communication. In particular, prior work has shown that robot behavior can be assessed through bipolar perceptual dimensions derived from expressive human movement studies. These dimensions aim to capture how robot motion is interpreted in terms of attitude, intention, and social meaning, beyond purely geometric or kinematic properties.

In line with this perspective, we evaluate robot behavior through 10 bipolar adjective scales: *aggressive–gentle*, *authoritarian–polite*, *confident–hesitant*, *solid–fragile*, *strong–weak*, *smooth–rough*, *rigid–flexible*, *tender–sensitive*, *inspires confidence–does not inspire confidence* and *sympathetic–antipathetic*. Taken together, these dimensions cover both physical qualities of motion and higher-level social interpretations. Some scales primarily reflect the dynamic and embodied properties of the movement, such as *smooth–rough*, *rigid–flexible*, *solid–fragile*, or *strong–weak*. Others capture

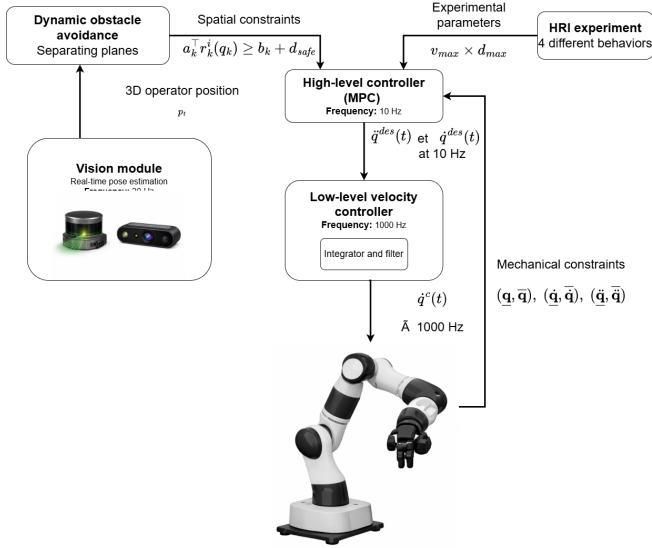


Fig. 1: Abstract flow diagram of control architecture.

interpersonal or affective impressions, such as *authoritarian-polite*, *aggressive-gentle*, *sympathetic-antipathetic*, or *tender-sensitive*. Finally, dimensions such as *confident-hesitant*, *self-assured-doubtful*, and *inspires confidence-does not inspire confidence* characterize the degree to which the robot is perceived as assured in its motion and able to elicit trust from the human partner.

This choice allows us to examine not only whether a given motion strategy is safe and efficient, but also how it is interpreted by human users. However, unlike prior work that mainly seeks to identify or validate perceptual dimensions, our objective is to analyze how controlled motion parameters, namely velocity modulation and interaction distance, influence these perceptions.

III. HIERARCHICAL CONTROL ARCHITECTURE FOR BEHAVIOR GENERATION

The goal of our scheme is to generate a collision-free trajectory for a cobot (e.g., a 7-DoF manipulator cobot) that has to perform a task in a workspace shared with a human worker. This scheme is composed of 3 parts, as shown in Figure 1: (i) a collision-free trajectory generation module (section IV), (ii) a human motion prediction module (section V) and (iii) a low-level robot motion control module (section VI).

IV. A COLLISION FREE TRAJECTORY MODULE

In this section, we summarise how to compute a collision-free trajectory based on the perception module and the cobot's dynamics. Also, we emphasise the role of the terminal constraint to guarantee safety. This terminal constraint provides a *passive motion safety* guarantee [11], which means that if a collision occurs, the cobot is at rest at the time of the collision so that it doesn't inject its own kinetic energy. In the following, we recall the main equations from the MPC approach developed in our previous work [3].




Fig. 2: An illustration of separating plane between two objects.

A. Separating plane optimisation

As illustrated in Figure 2, if there exists at the prediction time $k \in \mathbb{N}$ a plane defined by a normal vector $a_k \in \mathbb{R}^3$ and a scalar constant $b_k \in \mathbb{R}$ such that all vertices \mathbf{p}_k^j related to the human stay on one side between instants k and $k+1$ while all vertices r^i related to the cobot stay on the other side, then we have evidence that they don't collide over this interval of time. Here, $j \in \{1, \dots, N_p\}$ and $i \in \{1, \dots, N_r\}$ where N_p and N_r are the number of vertices associated with the human and the cobot. The distance $d \in \mathbb{R}$ is the geometric margin optimized during the separating-plane computation: it measures how far the robot vertices lie from the plane once the human vertices are constrained to remain on the other side. In the following, a_k^T denotes the transpose of the normal vector a_k .

$$\min_{a_k, b_k, d} -d + \alpha d^2 + \beta \|a_k - a_k^p\|^2 + \beta \|b_k - b_k^p\|^2 \quad (2a)$$

$$\text{s.t. } \forall j \in \{1, \dots, N_p\}, a_k^T \mathbf{p}_k^j \leq b_k, \quad (2b)$$

$$\forall j \in \{1, \dots, N_p\}, a_k^T \mathbf{p}_{k+1}^j \leq b_k, \quad (2c)$$

$$\forall i \in \{1, \dots, N_r\}, a_k^T r_k^i \geq b_k + d, \quad (2d)$$

$$\forall i \in \{1, \dots, N_r\}, a_k^T r_{k+1}^i \geq b_k + d, \quad (2e)$$

$$-1_3 \leq a_k \leq 1_3, \quad (2f)$$

$$1 - \varepsilon \leq a_k^T a_k^p \leq 1 \quad (2g)$$

We want to maximise the distance d between the separating plane and the cobot. Given the formulation as a minimisation problem, we include the term $-d$ in the cost function, Equation 2a. This quantity d is an intermediate result of the plane-search problem and should be distinguished from the fixed safety margin (used later in the MPC constraints). The following term in the cost function smooths the variations of separating planes, with a_k^p and b_k^p the separating plane parameters obtained at the previous sampling time, and some small weights α and β . Finally, we use constraints, Equations (2f-2g), to approximate a nonlinear constraint to bound the vector a_k to a unit norm. Where $1_3 \in \mathbb{R}^3$ is a row vector of ones.

Algorithm 1 CollisionFreeTrajectoryComputation

Input: U_k^p, S_k, a_k^p, b_k^p
Output: U_k
1: $i \leftarrow 0$
2: **while** $\|U_k - U_k^p\|^2$ or $i \leq k$ **do**
3: $U_k^p \leftarrow U_k$
4: Update robot parameters
5: $\{a, b\} \leftarrow$ solve Eq. (2) for $k \in \{0, \dots, N-1\}$
6: $\{U_k\} \leftarrow$ solve Eq. (3)
7: $i \leftarrow i + 1$
8: **end while**

B. Optimal motion generation

Once we have a sequence of separating planes parameters, we can include them in our MPC scheme to compute an optimal collision-free trajectory:

$$\min_{\underline{u}} \sum_{k=0}^{N-1} \|s_{k+1} - s_{k+1}^{des}\|_Q^2 + \|u_k - u_k^{des}\|_R^2 \quad (3a)$$

$$\text{s.t. } \forall k \in \{0, \dots, N-1\}, \underline{u} \leq u_k \leq \bar{u}, \quad (3b)$$

$$\forall k \in \{1, \dots, N\}, \underline{q} \leq q_k \leq \bar{q}, \quad (3c)$$

$$\forall k \in \{1, \dots, N-1\}, \underline{\dot{q}} \leq \dot{q}_k \leq \bar{\dot{q}}, \quad (3d)$$

$$\dot{q}_N = 0, \quad (3e)$$

$$\forall k \in \{0, \dots, N-1\}, \forall i,$$

$$a_k^T r_k^i(q_k^p) + a_k^T J(q_k^p)(q_k - q_k^p) \geq b_k + d_{safe}, \quad (3f)$$

$$\forall k \in \{0, \dots, N-1\}, \forall i,$$

$$a_k^T r_k^i(q_{k+1}^p) + a_k^T J(q_{k+1}^p)(q_{k+1} - q_{k+1}^p) \geq b_k + d_{safe} \quad (3g)$$

Where $q_k \in \mathbb{R}^n$ and $\dot{q}_k \in \mathbb{R}^n$ are respectively the joint position and velocity, with n the number of degrees of freedom. The state $s_k \in \mathbb{R}^{2n}$ includes q_k and \dot{q}_k , and the reference state $s_k^{des} \in \mathbb{R}^{2n}$ is defined similarly. The control input $u_k \in \mathbb{R}^n$ is the joint acceleration of the cobot, while u_k^{des} denotes its reference value. The safety distance d_{safe} represents a prescribed buffer zone between the robot and the predicted human position. Unlike d , which is optimized when computing the separating plane, d_{safe} is a design parameter imposed in the MPC constraints to enforce a minimum clearance during execution.

Our prediction horizon has a length $N \in \mathbb{N}$. The cost function, Equation 3a, is designed to track a reference state trajectory s_k^{des} together with a reference acceleration input u_k^{des} , while $\underline{q}, \bar{q}, \underline{\dot{q}}, \bar{\dot{q}}, \underline{u}, \bar{u}$ indicate minimum and maximum joint positions, velocities, and accelerations (we assume that $\underline{\dot{q}} \leq 0 \leq \bar{\dot{q}}$ and $\underline{u} \leq 0 \leq \bar{u}$). The terminal constraint, Equation 3e, ensures that the cobot is at rest at the end of the prediction horizon in order to provide a passive motion safety guarantee, making sure that the cobot is able to stop and stay at rest before any collision happens in the future. Equations 3f-3g introduce the collision avoidance constraints based on separating planes, which are computed by linearising the kinematics of the cobot around the previously computed trajectory:

$$r_k^i = r^i(q_k) \approx r^i(q_k^p) + J(q_k^p)(q_k - q_k^p) \quad (4)$$

C. Behavior Encoding

Interaction behaviors are encoded through two parameters v_{max} and d_{max} . These parameters directly affect the velocity bounds $\bar{q} \propto v_{max}$ and the safety distance $d_{safe} \propto d_{max}$.

The safety distance is directly proportional to the interaction distance parameter d_{max} defined in Section II. Larger values of d_{max} result in more conservative (anticipatory) motion, with increased clearance from human predictions. This safety margin is used in the collision-avoidance constraints (Equations 3f-3g) to ensure that the robot maintains a minimum separation from the human trajectory.

By selecting different combinations of (v_{max}, d_{max}) , the controller generates four distinct behaviors, modulating anticipation, reactivity, and efficiency.

V. HUMAN POSE DETECTION AND HUMAN HAND MOTION PREDICTION

The collision-avoidance scheme described in Section IV relies fundamentally on accurate predictions of future human positions to compute separating planes and enforce collision-avoidance constraints in the MPC formulation. Specifically, the predicted human hand positions feed into the separating plane optimization (Equations 2b-2c) and ultimately shape the safety constraints (Equations 3f-3g) that govern the robot's trajectory. The quality and accuracy of these predictions thus directly determine the effectiveness of the motion control strategy described above. In this section, we summarize our work on perception [2] and describe how we estimate the 3D position of the human hand from RGB-D data and use the resulting trajectory as input to a prediction model.

The collision-avoidance scheme requires a short-horizon prediction of future human positions. We therefore estimate the 3D position of the human hand from RGB-D data and use the resulting trajectory as input to a prediction model.

A. Human Pose Detection

The perception system relies on a Realsense L515 camera providing synchronized colour and depth images. Human keypoints are first detected in the RGB image with lightweight pose-estimation methods such as OpenPose [12] or MediaPipe [13]; examples are shown in Figure 3. The corresponding depth values are then associated to the detected pixels after camera calibration [14], and the 3D Cartesian position is recovered with the pinhole projection model.

B. Human Hand Motion Prediction

From the measured 3D hand positions, we predict a short sequence of future positions. To stay consistent with Section IV, we denote by $p_t \in \mathbb{R}^3$ the 3D position of the human hand at time t . The hand dynamics are modeled in discrete time as

$$p_{t+1} = g(p_t, w_t) \quad (5a)$$

$$z_t = h(p_t) \quad (5b)$$

where w_t represents unmeasured inputs and z_t is the observed hand position. Rather than predicting only one step ahead,

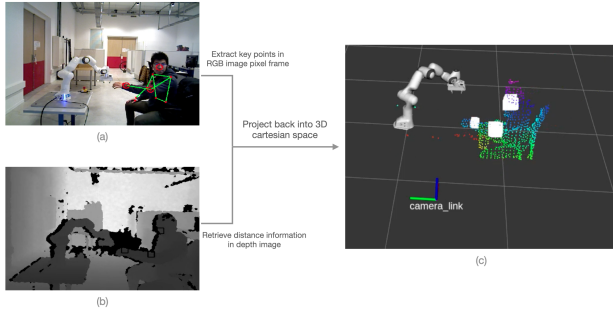


Fig. 3: Example demonstration of RGB-D image for mapping human hand in 3D space, a) shows RGB image on which hand joints are detected, b) shows the bounding box around the points in the depth image, and c) shows the mapping of points in 3D space.

we seek a mapping from a short observation window $\mathbf{p} = \{p_t, \dots, p_{t-L}\}$ to a future sequence $\hat{\mathbf{p}} = \{\hat{p}_{t+1}, \dots, \hat{p}_{t+T}\}$.

Because human motion is nonlinear and affected by unmeasured factors, the predictor is implemented as an encoder-decoder LSTM network. Past observations are encoded through stacked recurrent layers, and the decoder outputs a multi-step prediction $\hat{\mathbf{p}}$ of the future hand trajectory. This predicted trajectory is then passed to the motion planner to enforce collision-avoidance constraints in real time. In the end, the predicted positions $\hat{p}_{t+\ell}$ are provided to the separating plane computation and instantiate the human vertices \mathbf{p}_k^j used in Equations 2b and 2c.

VI. LOW-LEVEL CONTROL: VELOCITY TRACKING AND EXECUTION

The low-level controller is responsible for executing in real time the trajectory computed by the MPC. While the high-level predictive controller operates on the horizon and enforces collision-avoidance and terminal safety constraints, the low-level loop runs at a higher frequency (typically 100–1000Hz) and converts the planned joint motion into actuator commands. Its role is therefore not to re-plan the motion, but to ensure that the robot accurately tracks the admissible trajectory generated by the upper layer despite discretization, actuation delays, and small execution errors.

In order to keep the notation consistent with Section IV, we denote by $q^{des}(t)$, $\dot{q}^{des}(t)$, and $\ddot{q}^{des}(t)$ the position, velocity, and acceleration references provided by the high-level controller after interpolation of the MPC solution at a period Δt . The measured joint position and velocity are denoted by $q(t)$ and $\dot{q}(t)$. The low-level controller then computes a command $\dot{q}^c(t)$ to be sent to the robot servo loop at a period δt .

A. Velocity Integration and Command Generation

The MPC described in Section IV provides a sequence of joint accelerations u_k together with the associated state trajectory (q_k, \dot{q}_k) . At execution time, these discrete references are converted into continuous-time commands over one control period. In the velocity-controlled implementation, the commanded velocity must follow the state evolution predicted

TABLE I: Behavioral conditions used in the experiments

Condition	Behavior name	Velocity profile	Interaction distance
B1	Reactive and cautious	Low	Short / near
B2	Anticipatory and conservative	Low	Long / far
B3	Efficiency-driven	High	Short / near
B4	Anticipatory yet efficient	High	Long / far

by the MPC, in particular the desired velocity \dot{q}^{des} . To ensure consistency between the predictive model and the low-level controller, the desired acceleration \ddot{q}^{des} is integrated over the current control interval using \dot{q}^{des} as the initial condition.

In practice, \dot{q}^{des} corresponds to the acceleration command optimized by the MPC, i.e., $\dot{q}^{des}(t) \approx u_k$ over the current sampling interval. Since the MPC update rate is not strictly periodic and may deviate from the nominal period Δt , the actual integration time step δt may differ from Δt . Using \dot{q}^{des} as the initial condition ensures that the commanded velocity remains consistent with the linearized dynamics used in the MPC prediction model.

This results in the following control law:

$$\dot{q}^c(t) = \dot{q}^{des}(t) + \ddot{q}^{des}(t) \delta t. \quad (6)$$

The value computed by (6) is then fed to the robot servo-loop. Is integrated in this system a PID controller to guarantee that the robot follows the given joint speed $\dot{q}^c(t)$.

Overall, the low-level controller ensures faithful realization of the motion behaviors defined at the high level. The interaction style is determined upstream through the parameters v_{max} and d_{max} and through the MPC constraints, whereas the low-level loop guarantees that these planned behaviors are executed smoothly and accurately on the real robot. This separation between predictive planning and fast feedback control is important: it preserves the safety and interaction properties established at the planning level while providing the robustness required for physical execution.

VII. EXPERIMENT 1: PRELIMINARY VALIDATION OF THE FOUR BEHAVIORS

We first ran a preliminary robotic experiment to verify that the controller could generate the four behaviors introduced in Section II. In a fixed pick-and-place task, only v_{max} and d_{max} were varied, while the rest of the control architecture remained unchanged.

The four tested conditions are listed in Table I. Overall, the controller produced clearly differentiated motions while preserving task execution and safety. Videos of the experiment are available in the online playlist: https://youtube.com/playlist?list=PL8ZyzBKIMS520Njc_XROIVTgbS9IDw_n_.

An example of the efficiency-driven behavior is shown in Fig. 4. When the human–robot distance drops below the avoidance threshold, the robot accelerates to perform the evasive maneuver. This behavior is consistent with the velocity profile depicted in Fig. 5, where the robot velocity increases as the distance to the human decreases. This result highlights the efficiency-driven nature of the proposed control strategy, which favors rapid avoidance when the interaction distance becomes critical.




Fig. 4: Example of the robot’s efficiency-driven behavior.




Fig. 5: Velocity profile during obstacle avoidance showing acceleration as the robot approaches the obstacle.

Qualitatively, these preliminary trials showed that the controller can be used as a behavior generator and not only as a safety filter. Larger values of d_{max} produced earlier and more conservative deviations, while the velocity profile mainly affected the temporal appearance of the motion, with accelerating behaviors appearing smoother and more reactive. This first experiment therefore validates the practical relevance of the two-dimensional parameterization and provides the basis for the user study reported in the next section.

VIII. EXPERIMENT 2: HUMAN–ROBOT INTERACTION STUDY

The second experiment was designed to assess the acceptability of the generated behaviors from a human-centered perspective. The central question was no longer whether the controller could produce different trajectories, but whether these trajectories changed the human experience of interaction. More specifically, we asked whether the two control variables retained in this paper, namely velocity modulation and interaction distance, influence i) the subjective workload associated with the interaction and ii) the social and qualitative perception of the robot. This distinction is important because a motion strategy can be safe in a geometric sense while still being uncomfortable, difficult to interpret, or socially inappropriate, as emphasized in previous work on fluency, trust, and acceptability in human–robot interaction [6], [7].

A. Protocol

The study involved 111 participants and followed a repeated-measures 2×2 experimental design. Each participant was exposed to the four behavioral conditions resulting from

TABLE II: Summary of the user-study protocol

Participants	111
Design	Repeated-measures 2×2
Factors	Velocity \times interaction distance
Velocity levels	Deceleration / acceleration
Distance levels	Short / long
Conditions per participant	4
Total observations	444
Workload measure	NASA-TLX (global and weighted)
Perception measure	10 bipolar adjective scales

the combination of the two manipulated factors: acceleration versus deceleration, and short versus long interaction distance. Because each participant evaluated all conditions, the protocol made it possible to compare the perceived effects of the behaviors within the same individual, while limiting the confounding influence of inter-subject variability.

The sample included 78 women (70.27%) and 33 men (29.73%), with ages ranging from 18 to 68 years ($M = 25.16$, $SD = 9.82$).

The experimental logic was as follows. After interacting with the robot under each condition, participants were asked to evaluate both the experienced workload and their perception of the robot. Workload was assessed with the NASA-TLX questionnaire. In addition to the usual global score, a weighted NASA-TLX score was also computed, so as to account for the relative importance assigned by each participant to the six workload dimensions. This part of the protocol was intended to capture whether one type of robot behavior required more effort, more monitoring, or generated more stress than another.

Perception of the robot was assessed through 10 bipolar dimensions on a four-point scale. The evaluated dimensions were: *aggressive–gentle*, *authoritarian–polite*, *confident–hesitant*, *solid–fragile*, *strong–weak*, *smooth–rough*, *rigid–flexible*, *tender–sensitive*, *inspires confidence–does not inspire confidence* and *sympathetic–antipathetic*. These dimensions were selected to characterize not only physical comfort, but also the social interpretation of the robot motion. The associated analyses were conducted with linear mixed models in order to account for the repeated-measures structure of the data.

B. Measured variables

The two families of dependent variables correspond to two different hypotheses. The first hypothesis is related to

TABLE III: NASA-TLX scores by experimental condition

Condition	Global NASA-TLX	Weighted NASA-TLX
Deceleration + near	2.56 (1.09)	2.65 (1.20)
Deceleration + far	2.44 (1.07)	2.50 (1.14)
Acceleration + near	2.50 (1.12)	2.59 (1.22)
Acceleration + far	2.38 (1.06)	2.43 (1.16)

interaction cost: if one of the behaviors is too abrupt, too close, or too difficult to anticipate, it should increase workload scores. The second hypothesis is related to *interaction meaning*: even if all behaviors remain similarly easy to handle, they may nevertheless convey different impressions in terms of confidence, smoothness, robustness, politeness, or sympathy. This distinction is central to the contribution of the paper, because it separates what the user experiences functionally from how the user interprets the motion socially, in line with prior HRI studies showing that kinematic cues affect both workload-related coordination and social interpretation [9], [10].

C. Results on workload

At the descriptive level, workload scores remained relatively low across all conditions. The mean global NASA-TLX score over the whole study was 2.47, and the mean weighted NASA-TLX score was 2.54, indicating a globally low to moderate workload. Table III reports the descriptive statistics by condition. In both versions of the score, the near-distance conditions yielded slightly higher averages than the far-distance conditions, and decelerating conditions yielded slightly higher averages than accelerating ones. However, these differences remained small.

The inferential analysis confirmed that workload was not significantly affected by the manipulated factors. For the global NASA-TLX score, neither velocity ($p = .353$), nor interaction distance ($p = .491$), nor their interaction ($p = .340$) reached significance. The same conclusion held for the weighted score, with non-significant effects of velocity ($p = .457$), distance ($p = .232$), and interaction ($p = .486$). In practical terms, this means that none of the four behaviors made the interaction substantially more demanding for the participants. The control parameterization therefore appears able to modify the style of the robot motion without degrading the subjective ease of the interaction, which is consistent with the idea that variations in spatial overlap and motion coordination do not necessarily translate into a large cognitive cost when the interaction remains predictable and well structured [6], [9].

D. Results on robot perception

The results were more contrasted for the perceptual variables. While workload remained globally stable, several dimensions of robot perception were sensitive to the velocity profile. Table IV reports descriptive means for a subset of the most informative perceptual dimensions. A clear pattern emerges: accelerating conditions tend to receive more favorable ratings on dimensions related to confidence, solidity, strength, smoothness, and assurance.

TABLE IV: Selected perceptual ratings by condition (mean (SD))

Dimension	D+N	D+F	A+N	A+F
Conf.-hes.	2.55 (0.98)	2.52 (0.97)	2.25 (0.96)	2.23 (0.99)
Solid-frag.	2.14 (0.85)	2.08 (0.94)	1.94 (0.82)	2.05 (0.87)
Strong-weak	2.03 (0.78)	2.13 (0.79)	1.94 (0.74)	1.94 (0.81)
Smooth-rough	2.39 (0.91)	2.39 (0.91)	2.31 (0.89)	2.26 (0.83)
Polite-auth.	2.88 (0.84)	3.02 (0.81)	2.96 (0.81)	3.02 (0.83)
Symp.-anti.	2.36 (0.88)	2.22 (0.88)	2.26 (0.87)	2.27 (0.87)

TABLE V: Main inferential results for the user study

Variable	Main result
Global NASA-TLX	No effect of velocity, distance, or interaction
Weighted NASA-TLX	No effect of velocity, distance, or interaction
Confident-hesitant	Velocity effect, $p < .001$
Solid-fragile	Velocity effect, $p = .029$
Strong-weak	Velocity effect, $p = .001$
Smooth-rough	Velocity effect, $p = .037$
Polite-authoritarian	Distance tendency, $p = .080$
Sympathetic-antipathetic	Velocity \times distance tendency, $p = .066$

This descriptive pattern was supported by the mixed-model analysis. When the robot decelerated during the avoidance phase, it was perceived as significantly more hesitant on the *confident-hesitant* dimension ($p < .001$), more fragile on the *solid-fragile* dimension ($p = .029$), weaker on the *strong-weak* dimension ($p = .001$), rougher on the *smooth-rough* dimension ($p = .037$). By contrast, interaction distance had no strong main effect on most perceptual dimensions. A non-significant but interpretable tendency was nevertheless observed for *authoritarian-polite* ($p = .080$), suggesting slightly more positive ratings for larger clearances, and a marginal interaction was found for *sympathetic-antipathetic* ($p = .066$), indicating that sympathy may depend on a combination of velocity and distance rather than on either factor alone. These observations are coherent with prior studies showing that robot velocity, smoothness, and interpersonal distance strongly contribute to perceived comfort, safety, and acceptability [7], [8], [15].

E. Interpretation

Taken together, the two experiments lead to a coherent interpretation. Experiment 1 showed that the proposed control architecture can effectively generate differentiated motion regimes from the pair (v_{max}, d_{max}) . Experiment 2 then showed that these differences are not only technical: they are also perceived by human users, but mainly through the temporal dynamics of the motion rather than through geometric clearance alone.

The absence of significant workload effects is important. It indicates that none of the tested behaviors made the interaction more mentally or physically demanding in a substantial way. In other words, the proposed behavior modulation does not seem to come at the price of increased subjective difficulty. This is a useful result for robotic design, because it suggests that one can tune the style of the robot behavior without necessarily increasing the human cognitive cost of the interaction.

The perceptual results are arguably even more informative. They show that the velocity profile acts as a strong communicative cue. A robot that accelerates when avoiding the

human is interpreted as more confident, more solid, stronger and smoother, and more self-assured. Conversely, a robot that decelerates is more easily interpreted as hesitant, fragile, rough, or doubtful. This means that the temporal profile of the motion does not merely affect kinematic appearance; it changes the social meaning of the robot behavior. This interpretation is in agreement with the literature on legible and expressive motion, which highlights that humans infer robot intent and attitude from dynamic cues such as smoothness, timing, and anticipatory adaptation [6], [10].

By contrast, the weaker effect of interaction distance suggests that geometric clearance alone is not sufficient to strongly alter social perception in the considered task. From the point of view of the controller, this result is particularly interesting. It suggests that the role of d_{max} is primarily to shape anticipation and to maintain conservative safety margins, whereas the role of velocity modulation is more directly linked to legibility and acceptability. In practice, this means that a human-aware controller should not be designed only around separation distances, but also around the way motion unfolds over time. This nuance complements earlier proxemics-oriented studies, where distance was identified as an important factor of comfort and safety, but not necessarily the only or dominant determinant of social interpretation in all task settings [7], [8], [16].

Overall, these findings support the main thesis of this paper. Predictive safety-aware motion generation should be evaluated not only in terms of collision avoidance and passive safety guarantees, but also in terms of the subjective impressions it produces during interaction. The proposed parameterization offers a simple and operational way to address both aspects: d_{max} supports conservative and anticipatory avoidance, while velocity modulation influences how the robot is interpreted by the human partner. This makes the proposed framework relevant not only for safe interaction, but also for fluent and acceptable motion generation in shared workspaces, in line with broader perspectives on human-aware navigation and collaborative fluency [6]–[8].

IX. CONCLUSION AND PERSPECTIVES

This paper addressed motion generation for robots operating in shared human–robot environments, where safety must be ensured while preserving interaction quality. We proposed an MPC-based framework that combines predictive safety constraints with motion parameters related to human perception, in particular interaction distance and velocity profile.

The experiments showed that the proposed controller can generate distinct robot behaviors and that these behaviors are perceived differently by human users. More specifically, the results suggest that velocity modulation has a stronger impact on perceived social acceptability than geometric clearance alone. This highlights that safe robot motion should be evaluated not only through collision avoidance, but also through the social meaning conveyed by motion timing and smoothness.

These results support the relevance of combining safety, fluency, and acceptability in a unified control framework.

Future work will focus on a deeper understanding of the impact of these different behaviors, with a particular emphasis on perceived safety and trust, in order to design robot behaviors that can be adaptive and personalized. This will enable the extension of this work to real industrial scenarios involving collaborative robots, where social acceptability, perceived safety, and trust play a key role in long-term human–robot collaboration.

REFERENCES

- [1] A. Albu-Schäffer, S. Haddadin, C. Ott, A. Stemmer, T. Wimböck, and G. Hirzinger, “The dlr lightweight robot: design and control concepts for robots in human environments,” *Industrial Robot: an international journal*, vol. 34, no. 5, pp. 376–385, 2007.
- [2] P. Zheng, P.-B. Wieber, J. Baber, and O. Aycard, “Human arm motion prediction for collision avoidance in a shared workspace,” *Sensors journal*, 2022.
- [3] P. Zheng, P.-B. Wieber, and O. Aycard, “Online optimal motion generation with guaranteed safety in shared workspace,” in *ICRA*, 2020.
- [4] W. Li, Y. Hu, Y. Zhou, and D. T. Pham, “Safe human–robot collaboration for industrial settings: a survey,” *Journal of Intelligent Manufacturing*, vol. 35, no. 5, pp. 2235–2261, 2024.
- [5] P. Svarny, M. Tesar, J. K. Behrens, and M. Hoffmann, “Safe physical HRI: Toward a unified treatment of speed and separation monitoring together with power and force limiting,” in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Nov 2019, pp. 7580–7587.
- [6] G. Hoffman, “Evaluating fluency in human–robot collaboration,” *IEEE Transactions on Human-Machine Systems*, vol. 49, no. 3, pp. 209–218, 2019.
- [7] P. A. Lasota, T. Fong, and J. A. Shah, “A survey of methods for safe human–robot interaction,” *Foundations and Trends in Robotics*, vol. 5, no. 4, pp. 261–349, 2017.
- [8] E. A. Sisbot, L. F. Marin-Urias, R. Alami, and T. Simeon, “A human-aware mobile robot motion planner,” *IEEE Transactions on Robotics*, vol. 23, no. 5, pp. 874–883, 2007.
- [9] S. e. a. Howell, “Workspace overlap and cognitive load in human–robot collaboration,” *IEEE Transactions on Human-Machine Systems*, 2019.
- [10] A. D. Dragan, K. C. T. Lee, and S. S. Srinivasa, “Legibility and predictability of robot motion,” in *Proceedings of the 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2013, pp. 301–308.
- [11] S. Bouraine, T. Fraichard, and H. Salhi, “Provably safe navigation for mobile robots with limited field-of-views in dynamic environments,” *Autonomous Robots*, vol. 32, no. 3, pp. 267–283, 2012.
- [12] Z. Cao, G. Hidalgo, T. Simon, S.-E. Wei, and Y. Sheikh, “Openpose: realtime multi-person 2d pose estimation using part affinity fields,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, no. 1, pp. 172–186, 2019.
- [13] C. Lugaresi, J. Tang, H. Nash, C. McClanahan, E. Uboweja, M. Hays, F. Zhang, C.-L. Chang, M. G. Yong, J. Lee *et al.*, “Mediapipe: A framework for building perception pipelines,” *arXiv preprint arXiv:1906.08172*, 2019.
- [14] J. Nowak, P. Ffaisse, A. Cherubini, and J.-P. Daures, “Point clouds with color: A simple open library for matching rgb and depth pixels from an uncalibrated stereo pair,” in *2021 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*. IEEE, 2021, pp. 1–7.
- [15] M. e. a. Story, “Effects of robot motion characteristics on human perception,” *International Journal of Social Robotics*, 2018.
- [16] T. e. a. Arai, “Human-robot collaborative systems,” *IEEE Transactions on Industrial Electronics*, 2010.