Human-intention-aware skill modulation using energy tanks for collaborative tasks

Edoardo Fiorini[1], **Markus Knauer**[1,2] **and João Silvério**[1]
[1] German Aerospace Center (DLR), Robotics and Mechatronics Center (RMC),
Münchener Str. 20, 82234 Weßling, Germany
[2] Technical University of Munich (TUM), School of Computation, Information and Technology (CIT)
Arcisstr. 21, 80333 Munich, Germany
edoardo.fiorini@dlr.de

Abstract: Human-Robot interaction (HRI) is a key requirement to allow robotic systems to cooperate with humans in various daily scenarios. There are different methods for interacting with a robot, but physical contact offers the human the best sense of collaboration. However, to make the best use of this input, the robot's cognitive abilities need to distinguish which of the contacts it detects through force-torque sensing are human intentions and which are task-related environmental contacts. In this paper, we propose an energy-tank based method that detects human intention in three different Degrees of Freedom (DoF) of Cartesian space, allowing the human operator to correct or provide input to a skill in the desired direction. During the interaction, a key role is played by the controller, which is responsible for the robot's compliant behavior. In our novel approach, we modulate the stiffness and reference force of an impedance controller according to an intention index, making the collaboration smoother and more sensitive. We demonstrate our approach with a learned-by-demonstration pick-and-place manufacturing task on a torque-controlled, 7-DoF robot. Thanks to the force and joint torque capabilities of the robot, we exploit an external force observer to allow the user to interact with any part of the robot's body instead of just the end-effector, as is usually the case. Overall, the user is able to interact naturally and intuitively with the system, adding via-points in different DoF to the learned skill through a Kernelized Motion Primitives (KMP) model.

1 Introduction

Recent advancements in robotics and artificial intelligence are enabling robotic systems to develop increasingly sophisticated cognitive capabilities, positioning them to play a growing role in various human activities. The fusion of robotic precision, speed, and repeatability with human perception and decision-making is expected to yield significant advancements in areas such as cooperative manufacturing, medical and rehabilitation assistance [1], and precision agriculture [2]. A critical attribute of such systems is collaboration, enabling robots to work alongside humans and extend their capabilities. This raises the question: what does it mean for a robot to be "collaborative"? In current research, this term is closely tied to human-robot interaction (HRI), where the key elements of successful interaction are compliance and safety [3]. Compliance refers to the robot's ability to modulate its stiffness based on the level of contact required with the human operator, while safety mechanisms ensure the system halts immediately in response to unexpected contact that could pose a risk to the human collaborator. These factors are essential for enabling humans to share a workspace with a cobot (collaborative robot).

Beyond safety and compliance, another crucial aspect of collaboration is the robot's ability to learn and adapt to tasks. Robots can acquire tasks either through demonstration [4] or programmed instructions, with human input provided at the start or during task execution [5]. Various modalities

8th Conference on Robot Learning (CoRL 2024), Munich, Germany.

enable this input in HRI, including graphical user interfaces, voice recognition, vision-based systems, and physical-contact interaction. While the former are more "virtual" forms of communication, physical contact provides a more direct and intuitive sense of collaboration, closely resembling human-to-human interaction. In the context of physical-contact interaction, it is critical that input from the human operator, typically measured through the robot's force/torque sensors, is not mistaken for task-related interactions with the environment. Therefore, the robot must be able to accurately interpret human intention to facilitate this type of interaction.

Various approaches have been proposed to differentiate between human intention and sensor noise. For example, machine learning methods can be trained to classify the source of input [6], or thresholds can be established to distinguish between sensor noise and forces exerted by an operator intending to collaborate on a task [7]. Achieving compliant behavior, as discussed earlier, depends largely on how the system is controlled. One of the most effective and widely adopted solutions is impedance control, which adjusts the robot's response through its stiffness matrix. In [8], the stiffness matrix is modulated based on detected forces, while [9] adjusts the matrix according to the variance of the input data. Additionally, in [10], a system is proposed where position or impedance control is activated based on user requirements, allowing the human operator to take control of the robot when necessary. Typically, such systems alternate between phases of autonomy and phases where the user has direct control over the robot.

Figure 1: The experimental setup is composed of a torque-controlled, 7-DoF, robot mounted on a manufacturing workstation, boxes A and B, which contain inside them the rings that have to be picked from A and placed to B.

In this work, we present a novel energy-tank-based method, inspired by [11], that effectively detects human intention across the three positional degrees of freedom (DoF) of the Cartesian space. By leveraging an external observer [12], integrated with the force-torque and joint-torque sensors of our torque-controlled 7-DoF robot (Fig. 1), we enable interaction along the entire robot body—not just at the end-effector. Additionally, we dynamically modulate the stiffness matrix and force reference of the robot's impedance controller using an intention index, ensuring a smooth and intuitive interaction experience. We validate our approach on a torque-controlled, 7-DoF robot performing a pick-and-place task, learned via demonstration using a kernelized motion primitives (KMP) model. Once the task is learned, the user can interact with the system by introducing new via-points to refine the skill further. Our method allows these via-points to be seamlessly integrated into the task trajectory across multiple DoF, according to the detected intention.

The structure of this paper is as follows: Section 2 outlines the relevant theoretical concepts, Section 3 details the proposed approach, Section 4 presents an analysis of the experimental results, and Section 5 outlines the conclusions and discusses potential avenues for future research.

2 Background

2.1 Energy-tank-based intention detection and control

Energy-tank-based approaches are a common choice when designing controllers to ensure system passivity and optimize energy efficiency [13]. In [11], Khoramshahi et al. proposed a human-

guidance detection method which behaves as a switch. Representing the external and the estimated human-intended forces as $\boldsymbol{F}_e$ and $\boldsymbol{F}_h$ respectively, they are related as follows

$$\boldsymbol{F}_h = h\boldsymbol{F}_e, \tag{1}$$

where $h \in [0, 1]$ describes the interaction ratio. A value of $h = 0$ means that no intention is detected when an external force is applied, while $h = 1$ is reached when the external force is the result of an intentional human contact. The ìntensity of the interaction can be estimated using the input P_i and output P_o powers of the system

$$P_i = \dot{\boldsymbol{x}}^T \boldsymbol{F}_e, \qquad P_o = \dot{\boldsymbol{x}}^T \boldsymbol{F}_h, \tag{2}$$

where $\dot{\boldsymbol{x}}$ is the end-effector Cartesian velocity. An energy tank, with size $E_{\max}$ and state E, with dynamics

$$\dot{E} = P_i - P_o - (1 - h)P_d, \tag{3}$$

is considered, where $P_d > 0$ is the dissipative rate which lowers the tank energy when no input power is applied. The energy in the tank is computed at t by integrating (3) as $E = \int_0^t \dot{E}dt$. The human interaction rate h is computed according to the energy stored in the tank as

$$h = \begin{cases} 0, & \text{if} \quad E \leqslant E^* \\ (E - E^*)/(E_{max} - E^*), & \text{otherwise} \end{cases} \tag{4}$$

where E_{max} and E^* are the maximum size of the tank and the energy value that triggers the start of a possible intention detection, respectively. As per (1)–(4), externally applied forces increase the energy in the tank through P_i, which is positive when the robot moves along the direction of the applied force. In addition, the dissipative factor P_d helps reject high frequency, low magnitude forces. The combination of these terms and the resulting dynamics into (4) provide a measure of when an external force is the result of a purposeful human interaction.

2.2 Impedance control

Given a manipulator with n DoF, its rigid-body dynamics is represented by the following expression [14]

$$\boldsymbol{B}(\boldsymbol{q})\ddot{\boldsymbol{q}} + \boldsymbol{C}(\boldsymbol{q}, \dot{\boldsymbol{q}})\dot{\boldsymbol{q}} + \boldsymbol{g}(\boldsymbol{q}) = \boldsymbol{\tau} + \boldsymbol{\tau}^{ext} \tag{5}$$

with the generalized coordinate vecotr $\boldsymbol{q} \in \mathbb{R}^n$ and the inertia, Coriolis and centrifugal, and gravity matrices $\boldsymbol{B} \in \mathbb{R}^{n\times n}$, $\boldsymbol{C} \in \mathbb{R}^{n\times n}$, and $\boldsymbol{g} \in \mathbb{R}^n$ respectively; $\boldsymbol{\tau}$, $\boldsymbol{\tau}^{ext} \in \mathbb{R}^n$ represent the input command and external forces.

Considering the full Cartesian space, let us define $m = 6$. Thus, the compliant end-effector robot behavior required in HRI can be achieved using a cartesian impedance control [15]

$$\boldsymbol{\tau}_{imp} = \boldsymbol{J}(\boldsymbol{q})^T \boldsymbol{F}_{imp} + \boldsymbol{g}(\boldsymbol{q}) \tag{6}$$

$$\boldsymbol{F}_{imp} = \boldsymbol{K}_P \tilde{\boldsymbol{x}} + \boldsymbol{K}_D \dot{\tilde{\boldsymbol{x}}} + \boldsymbol{F}_d \tag{7}$$

where $\boldsymbol{F}_{imp}$ is the task-space wrenched needed to obtain a desired spring-damper behavior starting from a virtual spring deflection $\tilde{\boldsymbol{x}} = \hat{\boldsymbol{x}} - \boldsymbol{x}$, with $\hat{\boldsymbol{x}}$ and $\boldsymbol{x} \in \mathbb{R}^m$ the desired and actual operational-space coordinates, respectively; $\boldsymbol{K}_P$ and $\boldsymbol{K}_D \in \mathbb{R}^m$ define the stiffnes and damping matrices, which determine how force must be applied in response to deviations in both position and velocity; $\boldsymbol{F}_d$ is the the force the robot should desirably apply on the environement, since in many tasks it is essential to be able to control the interaction force as well. Here, the robot tracks the desired positions with a force overley, adapting the compliant behavior thanks to $\boldsymbol{K}_P$.

2.3 Kernelized movement primitives

KMPs [16] are used in LfD to predict the distribution of $\boldsymbol{\xi}$ given observations of $\boldsymbol{s}$. A KMP is initialized with a *reference trajectory distribution* comprised of N Gaussians with parameters $\{\boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n\}_{n=1}^N$, computed from human demonstrations for inputs $\boldsymbol{s}_{n=1,\dots,N}$ using a GMM.

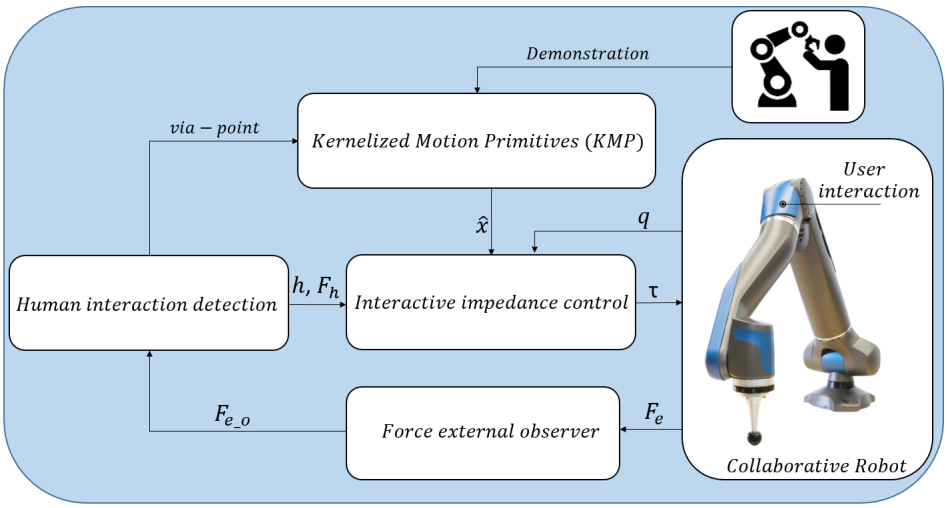

Figure 2: A schematic representation of the proposed approach. $\boldsymbol{F}_e$ is the redundant set of force-torque measurements, $\boldsymbol{F}_{e_o}$ is the generalized momentum along the body of the robot, h and $\boldsymbol{F}_h$ are the intention rate coming from the detected intention and human-intended force respectively, $\boldsymbol{\tau}$ is the final command input, $\boldsymbol{q}$ is the robot joint state, and $\hat{\boldsymbol{x}}$ is the predicted point coming from the KMP model.

For a test input $\boldsymbol{s}^*$, the expectation and covariance of $\boldsymbol{\xi}(\boldsymbol{s}^*)$ are given by

$$\mathbb{E}\left[\boldsymbol{\xi}(\boldsymbol{s}^*)\right] = \boldsymbol{k}^* \left(\boldsymbol{K} + \lambda_1 \boldsymbol{\Sigma}\right)^{-1} \boldsymbol{\mu}, \tag{8}$$

$$\operatorname{cov}\left[\boldsymbol{\xi}(\boldsymbol{s}^*)\right] = \frac{N}{\lambda_2}\left(\boldsymbol{k}^{**} - \boldsymbol{k}^* \left(\boldsymbol{K} + \lambda_2 \boldsymbol{\Sigma}\right)^{-1} \boldsymbol{k}^{*\top}\right), \tag{9}$$

where $\boldsymbol{K} = [\hat{\boldsymbol{k}}(\boldsymbol{s}_1)^\top, \ldots, \hat{\boldsymbol{k}}(\boldsymbol{s}_N)^\top]$, $\boldsymbol{k}^* = \hat{\boldsymbol{k}}(\boldsymbol{s}^*)$, with $\hat{\boldsymbol{k}}(\boldsymbol{s}_i) = [\boldsymbol{k}(\boldsymbol{s}_i, \boldsymbol{s}_1), \ldots, \boldsymbol{k}(\boldsymbol{s}_i, \boldsymbol{s}_N)]$, $\boldsymbol{k}^{**} = \boldsymbol{k}(\boldsymbol{s}^*, \boldsymbol{s}^*)$, $\boldsymbol{k}(\boldsymbol{s}_i, \boldsymbol{s}_j) = k(\boldsymbol{s}_i, \boldsymbol{s}_j)\boldsymbol{I}$ and $k(\boldsymbol{s}_i, \boldsymbol{s}_j)$ is a kernel function. Moreover, $\boldsymbol{\mu} = \left[\boldsymbol{\mu}_1^\top \ldots \boldsymbol{\mu}_N^\top\right]^\top$, $\boldsymbol{\Sigma} = \operatorname{blockdiag}\left(\boldsymbol{\Sigma}_1, \ldots, \boldsymbol{\Sigma}_N\right)$ and λ_1, λ_2, are hyperparameters. The kernel matrices are denoted as $\boldsymbol{K}, \boldsymbol{k}^*$ and $\boldsymbol{k}^{**}$. From (8)–(9) it follows that if, for a certain $\boldsymbol{\mu}_n$, the covariance $\boldsymbol{\Sigma}_n$ is small, the expectation at $\boldsymbol{s}_n$ will be close to $\boldsymbol{\mu}_n$. This provides a principled way for trajectory modulation. Indeed, if, for a new input $\bar{\boldsymbol{s}}$, one wants to ensure that the expectation passes through a desired $\bar{\boldsymbol{\mu}}$, it suffices to manually add the pair $\{\bar{\boldsymbol{\mu}}, \bar{\boldsymbol{\Sigma}}\}$ to the reference distribution provided that $\bar{\boldsymbol{\Sigma}}$ is small enough. This both makes (8) closely match $\bar{\boldsymbol{\mu}}$ and lowers the covariance (9) to match $\bar{\boldsymbol{\Sigma}}$.

3 Proposed Approach

The proposed approach combines the concepts outlined in Section 2 to enable the human operator to interact directly with the robot at any point along its kinematic chain. This interaction allows the operator to convey their intentions across various DoFs in the task space. The detected intention modulates the underlying impedance controller, ensuring compliant contact and facilitating an intuitive interaction method for correcting learned models – specifically, KMPs – by incorporating new via-points. An overview of the implemented workflow is depicted in Fig. 2.

3.1 Human Intention Detection

The main limitation of the energy-tank-based method described in Section 2, is that all dimensions of the external force $\boldsymbol{F}_e$ contribute to the computation of h. Specifically, when the power of the system P_i is calculated according to Eq. (2), these are combined with the Cartesian velocity resulting in a single value, which fills the level of the energy tank. This leads to a single interaction index for all DoFs in the operational space. As a result, it is challenging to determine which DoF the user intends to correct.

Let us consider only positional Cartesian coordinates $\boldsymbol{F}_e \in \mathbb{R}^3$, to allow the operator to freely express their intention along the desired task-space axis, we propose to replicate the presented architecture three times, one for each DoF $j = 1, \ldots, 3$. This decouples all the elements and creates three different energy sub-tanks. The input of each sub-architecture is the force and Cartesian velocity of the corresponding DoF. In this way we are able to obtain three distinct intention rates h, which we define as h_j. Furthermore, within our new architecture we also propose to change the dynamics of how h is calculated from the tank energy. The solution in Eq. (4) is, on the one hand, potentially discontinuous, depending on the choice of E^*, and, on the other hand, its rate of change is difficult to modulate. To achieve a smoother behavior we introduce a logistic sigmoid function

$$h_j = \frac{1}{1 + e^{-\alpha(E_j - E_j^*)}} \tag{10}$$

where α is the logistic growth rate representing the steepness of the curve, while E_j is the E_j^* value at the function's midpoint. The parameter α provides the means to regulate the smoothness of the change of h, making it better suited for human interaction. Equations (1)–(3) are thus reformulated as:

$$F_h^j = h_j F_e^j \tag{11}$$

$$P_i^j = \dot{x}_j F_e^j \qquad P_o^j = \dot{x}_j F_h^j \tag{12}$$

$$\dot{E}_j = P_i^j - P_o^j - (1 - h_j) P_d^j. \tag{13}$$

Another fundamental elements of our human intention detection approach is the learning of the nominal energy task. It is well-known that the robot accumulates energy during his movement since it has a velocity and can register forces due to the imperfect compensation of gravity, the gripper or the weight of the object being transported. Indeed, to find only human intention and not be influenced by these common factors, the power expressed as Eq. (13) is integrated during an execution without interaction with the user, $\tilde{E}_j = \int_0^t \dot{E}_j dt$, learned by the KMP, and then subtracted from the energy calculated during the phase of possible interaction with the operator $\hat{E}_j = \int_0^T \dot{E}_j dt - \tilde{E}_j$, compensating for the possible noises explained previously.

Unlike [11], who used a force-torque sensor at the end-effector to measure $\boldsymbol{F}_e$, we here rely on a external force observer [12] which allows us to accurately measure external forces on any point of the robot kinematic chain. This way users are not constrained to interact with only the end-effector, when trying to correct a skill, but are able to interact with any point on the robot, making our approach more intuitive and versatile.

3.2 Interactive impedance control

The intention captured by the architecture plays a key role in making the impedance control that governs the robot more interactive and in making the behaviour more compliant when required. The estimated human-intended force F_h^j and the interaction rate h_j modulate the impedance control Section 2.2 at each loop, the obtained control command input τ is given by

$$\boldsymbol{\tau} = \boldsymbol{J}(\boldsymbol{q})^T \boldsymbol{F}_{imp} + \boldsymbol{g}(\boldsymbol{q}) \tag{14}$$

$$\boldsymbol{F}_{imp} = [\boldsymbol{I} - \mathrm{diag}(h_1, \ldots, h_j)] \boldsymbol{K}_P \tilde{\boldsymbol{x}} + \boldsymbol{K}_D \dot{\tilde{\boldsymbol{x}}} + \boldsymbol{F}_h \tag{15}$$

where $h_1, \ldots, h_m$ and $\boldsymbol{F}_h = [F_h^1, \ldots, F_h^m]$ modulate the position elements of the diagonal stiffness matrix $\boldsymbol{K}_P$ and the force elements respectively. While the orientation and torque part are kept constant because we do not deal with them in our architecture. Overall, the robot becomes softer when the intention starts to be detected or, in other words, h reaches 1. Meanwhile, the force that is passed as a reference within the controller, which also grows proportionally to the intention, further facilitates the interaction with the robot because it pushes in the same direction as the user, further cancelling out the robot's inertia, but at the same time following the desired position target.Here, it is clear why a regular dynamic of h is necessary, this way we are sure that the robot does not jump during its change of behaviour, making everything as natural as possible.

3.3 Intention-aware via-point definition

We employ the DoF-specific intention detection to selectively modulate a demonstrated skill along the relevant DoFs. For this we use the KMP via-point definition mechanism, described in Section 2.3, to add new via-points based on the values of h_j. Consider a via-point $\bar{\boldsymbol{\mu}} \in \mathbb{R}^m$ as introduced in Section 2.3. If an intention is detected on the j-th DoF at time t, i.e. $h_{j,t} = 1$, a via-point is created with

$$\bar{\boldsymbol{s}} = t, \qquad \bar{\mu}_j = \bar{x}_j, \qquad \bar{\boldsymbol{\Sigma}} = \gamma \boldsymbol{I} \tag{16}$$

where $\bar{x}_j$ is the value of the j-th DoF measured when $h_{j,t} = 1$ and γ is a small factor. The remaining DoFs $\bar{\mu}_i$, $i \neq j$, are set to the values of $\boldsymbol{\mu}_n$ in the *reference trajectory distribution* whose input $\boldsymbol{s}_n$ is the closest to t.

4 Results and Discussion

We evaluate our approach on a torque-controlled, 7-DoF, robot in a pick-and-place production task learned by demonstration. This type of task often requires skill corrections due to changes in the environment (e.g., obstacles) or variations in object tolerances that increase precision demands. Figure 1 shows the experimental setup, where the goal is to take a ring from box A and place it in box B. The user provided 5 demonstrations, and a KMP model was trained with hyperparameters $l = 2 \times 10^{-1}$, $N = 500$, $\lambda_1 = 1 \times 10^{-5}$. Note that in this experiment we did not employ covariance prediction (9) as the focus was on generating a trajectory for the robot to follow. The proposed interactive framework (Section 3) was tested by physically interacting with the robot along its body to correct the learned trajectory during task execution. When the interaction indices reached $h_j = 1$, signaling human intention, via-points were added at the corresponding DoF at that time instant. In our experiments demonstrations, forces and corrections are all represented in the robot base, thus the human can interact with the desired DoF with respect to this reference frame. For the energy tank parameters we used $E_j^* = 0.002$ and $\alpha = 4800$. While this experiment validated the core features of our approach, it was conducted in a simplified environment. Future work will focus on applying the method to more realistic scenarios with dynamic environments and greater task complexity.

Figure 3 presents the obtained results in this experiment. On the left-hand side, the Cartesian position of the end-effector is shown, with the dark blue trajectory representing the path learned through demonstration and the green trajectory showing the re-generated path after the desired via-points, marked with red stars, have been added. In the middle plots we display the intention index h_j, for the different DoFs, throughout the task. It can be observed that the two trajectories on the left-most plots follow a similar pattern, except in the regions where via-points are introduced. Here, the green line deviates from the original trajectory, passing through the via-points and then rejoining the initial path, effectively performing the intended task correction. Each via-point added to the trajectory corresponds to a peak in the corresponding h_j at the same time instant, indicating that the new point is only added when human intention is detected by the proposed method.

On the right-hand side of the figure, we show the recorded forces during task execution. Several increases in force are observed throughout the task, but only those that contribute to increasing the system's energy are detected as intentions. Force increases that do not result in intention detection are attributed to routine interactions with the environment, such as picking up or dropping the ring, and are correctly ignored by our system. Additionally, note that the force along the vertical axis, F_3, is not zero during the ring transportation due to the ring's weight. However, thanks to our energy compensation mechanism described in 3.1 this force does not trigger an intention.

5 Conclusions

For a human operator to collaborate effectively with a robot in executing a skill, it is essential that the robot can understand human intention and adapt its movements to ensure safe shared operation. In this paper, we proposed a method for detecting human intention through physical contact along

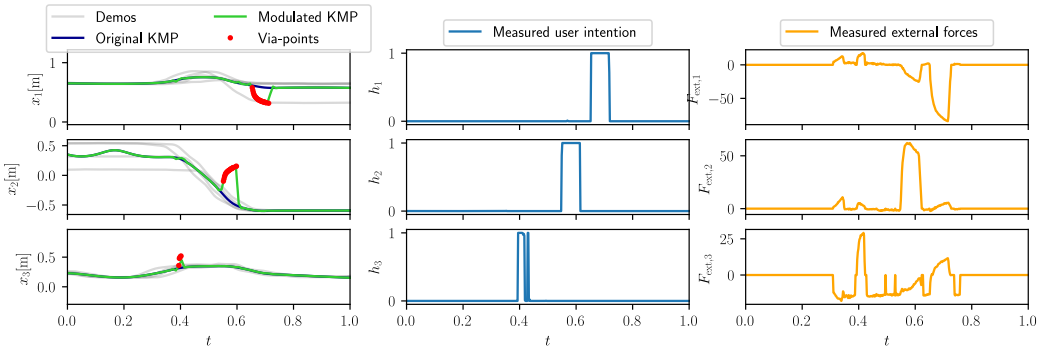

Figure 3: The plot shows the obtained results during the correction of a manufacturing pick and place task performed by a torque-controlled robot along different DoF. **Left**: dark blue line is the trajectory predicted by KMP model, green line is the one generated by KMP model with added via points, red star are desired via-points, and grey the given demontrations. **Center**: The interaction index h. **Right**: The forces measured by external observer during the human-robot body contact.

the robot's entire body. Our approach generates an intention index, which modulates the stiffness and force reference of the impedance controller that drives the robot. We evaluated this method on a manufacturing task learned via demonstration using a KMP model, where intention was used to add new via-points, correcting the trajectory.

The proposed method can also be applied in scenarios where intention is needed to activate sub-tasks, restore a learned method to its original state after perturbations, change control policies, or activate safety layers. As future work, we plan to extend our method to correct skills not only in position but also in force profiles. Additionally, we aim to integrate a visualization tool, such as a web interface, to allow real-time monitoring of skills, robot movements, detected intentions, and the ability to adjust hyperparameters.

Acknowledgments

The authors kindly thank Maged Iskandar and Thomas Eiband for their feedback and insights during the development of our approach. This work was partially funded by the DLR project "Factory of the Future Extended" and the European Union's Horizon Research and Innovation Programme under Grant 101136067 (INVERSE).

References

[1] P. E. Dupont, B. J. Nelson, M. Goldfarb, B. Hannaford, A. Menciassi, M. K. O'Malley, N. Simaan, P. Valdastri, and G.-Z. Yang. A decade retrospective of medical robotics research from 2010 to 2020. *Sci Robot*, 6(60):eabi8017, Nov. 2021.

[2] A. Botta, P. Cavallone, L. Baglieri, G. Colucci, L. Tagliavini, and G. Quaglia. A review of robots, perception, and tasks in precision agriculture. *Applied Mechanics*, 3(3):830–854, 2022. ISSN 2673-3161. doi:10.3390/applmech3030049. URL https://www.mdpi.com/2673-3161/3/3/49.

[3] T. B. Sheridan. Human–robot interaction: Status and challenges. *Human Factors*, 58(4):525–532, Jun 2016. ISSN 0018-7208. doi:10.1177/0018720816644364. URL https://doi.org/10.1177/0018720816644364.

[4] B. D. Argall, S. Chernova, M. Veloso, and B. Browning. A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57(5):469–483, May 2009. ISSN 0921-8890. doi:10.1016/j.robot.2008.10.024.

[5] C. Celemin, R. Pérez-Dattari, E. Chisari, G. Franzese, L. de Souza Rosa, R. Prakash, Z. Ajanović, M. Ferraz, A. Valada, J. Kober, et al. Interactive imitation learning in robotics: A survey. *Foundations and Trends® in Robotics*, 10(1-2):1–197, 2022.

[6] F. Franzel, T. Eiband, and D. Lee. Detection of collaboration and collision events during contact task execution. In *2020 IEEE-RAS 20th International Conference on Humanoid Robots (Humanoids)*, pages 376–383, 2021.

[7] M. Knauer, A. Albu-Schäffer, F. Stulp, and J. ao Silvério. Interactive incremental learning of generalizable skills with local trajectory modulation. *arXiv*, 2024. URL https://arxiv.org/abs/2409.05655.

[8] S. Jadav, J. Heidersberger, C. Ott, and D. Lee. Shared autonomy via variable impedance control and virtual potential fields for encoding human demonstration, 2024. URL https://arxiv.org/abs/2403.12720.

[9] J. Silvério, Y. Huang, L. Rozo, F. Abu-Dakka, and D. G. Caldwell. Uncertainty-aware imitation learning using kernelized movement primitives. page (to appear), Macau, China, November 2019.

[10] J. Silvério, Y. Huang, L. Rozo, S. Calinon, and D. G. Caldwell. Probabilistic learning of torque controllers from kinematic and force constraints. pages 6552–6559, Madrid, Spain, October 2018.

[11] M. Khoramshahi and A. Billard. A dynamical system approach for detection and reaction to human guidance in physical human–robot interaction. *Autonomous Robots*, 44(8):1411–1429, Nov 2020. ISSN 1573-7527. doi:10.1007/s10514-020-09934-9. URL https://doi.org/10.1007/s10514-020-09934-9.

[12] M. Iskandar, O. Eiberger, A. Albu-Schäffer, A. De Luca, and A. Dietrich. Collision detection, identification, and localization on the dlr sara robot with sensing redundancy. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3111–3117, 2021. doi:10.1109/ICRA48506.2021.9561677.

[13] F. Califano, R. Rashad, C. Secchi, and S. Stramigioli. On the use of energy tanks for robotic systems. In P. Borja, C. Della Santina, L. Peternel, and E. Torta, editors, *Human-Friendly Robotics 2022*, pages 174–188, Cham, 2023. Springer International Publishing. ISBN 978-3-031-22731-8.

[14] B. Siciliano, L. Sciavicco, L. Villani, and G. Oriolo. *Robotics: Modelling, Planning and Control*. Springer Publishing Company, Incorporated, 2010. ISBN 1849966346.

[15] F. Caccavale, C. Natale, B. Siciliano, and L. Villani. Six-dof impedance control based on angle/axis representations. *IEEE Transactions on Robotics and Automation*, 15(2):289–300, 1999. doi:10.1109/70.760350.

[16] Y. Huang, L. Rozo, J. Silvério, and D. Caldwell. Kernelized movement primitives. *The International Journal of Robotics Research (IJRR*, 38:833–852, 05 2019. doi:10.1177/0278364919846363.