

Bilateral Teleoperation with Adaptive Impedance Control for Contact Tasks

Youssef Michel, Rahaf Rahal, Claudio Pacchierotti, Paolo Robuffo Giordano, Dongheui Lee

Abstract—This paper presents an adaptive impedance control architecture for robotic teleoperation of contact tasks featuring continuous interaction with the environment. We use Learning from Demonstration (LfD) as a framework to learn variable stiffness control policies. Then, the learnt state-varying stiffness is used to command the remote manipulator, so as to adapt its interaction with the environment based on the sensed forces. Our system only relies on the on-board torque sensors of a commercial robotic manipulator and it does not require any additional hardware or user input for the estimation of the required stiffness. We also provide a passivity analysis of our system, where the concept of energy tanks is used to guarantee a stable behavior. Finally, the system is evaluated in a representative teleoperated cutting application. Results show that the proposed variable-stiffness approach outperforms two standard constant-stiffness approaches in terms of safety and robot tracking performance.

Index Terms—Compliance and Impedance Control, Learning from Demonstration, Telerobotics and Teleoperation

I. INTRODUCTION

MODERN-day robotics has witnessed a significant advancement both from the mechanical and control points of view. These promising results have enabled a technology paradigm shift, from the rigid high-gain position controlled robots used in structured industrial environments to more compliant and light-weighted systems, bringing robots one step closer to an everyday interaction and collaboration with humans. Nevertheless, the advanced human cognitive abilities and problem-solving skills are still required in many scenarios. Bilateral telemanipulation allows robotic systems to benefit from such human abilities, enabling humans to accurately perform critical tasks in otherwise inaccessible environments. In addition to their cognitive capabilities, humans are also able to skillfully interact with a variety of environments by continuously modulating their end-point force and impedance through coordinated action of their spring-like muscles [1], [2]. This human capability has inspired variable-impedance control methods, which aim at endowing teleoperated robots with such force adaptation capabilities.

Manuscript received: October 15, 2020; Revised: January 11, 2021; Accepted: February 23, 2021.

This paper was recommended for publication by Editor Clement Gosselin upon evaluation of the Associate Editor and Reviewers' comments.

This research was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – SPP 2134 and the Helmholtz Association

Y. Michel is with the Human-centered Assistive Robotics, Technical University of Munich, Munich, Germany. youssef.abdelwadoud@tum.de

R. Rahal is with Univ Rennes, Inria, CNRS, IRISA, Rennes, France. rahaf.rahaf@irisa.fr

C. Pacchierotti and P. Robuffo Giordano are with CNRS, Univ Rennes, Inria, IRISA, Rennes, France. claudio.pacchierotti.prg@irisa.fr

D. Lee is with the Human-centered Assistive Robotics, Technical University of Munich and the Institute of Robotics and Mechatronics, German Aerospace Center (DLR), Munich, Germany. dhlee@tum.de

Digital Object Identifier (DOI): see top of this page.

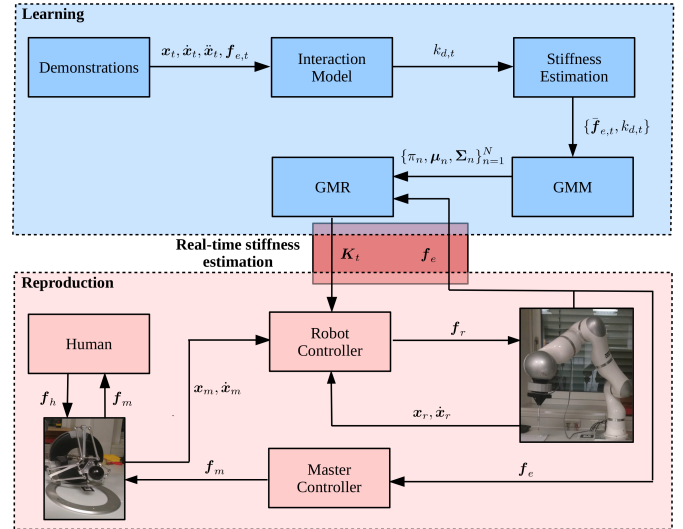


Fig. 1. Block diagram summarizing our system. The learning framework allows the dynamic estimation of the desired stiffness K_t for the teleoperation task. This estimation is based on the measured force f_e on the robot.

In this respect, Buchli et al. [3] proposed a reinforcement learning approach to learn a variable impedance gain schedule based on a user defined objective. Kronander and Billard [4] proposed a teaching interface to allow a human teacher to transfer variable stiffness control policies to a robot. This was achieved by using increased grasp pressure as a sign to stiffen up, while a wiggling gesture indicates the robot shall become more compliant. A framework for motion learning and adaptive impedance control in human-robot interaction was proposed in [5], allowing the robot to anticipate its human partner's motion and adapt to uncertainty. This has inspired the work by Rozo et al. [6], where physical collaborative robot behaviors are learnt from human demonstrations, including both motion and stiffness profiles. Along the same lines, Abu Dakka et al. [7] proposed a variable impedance control framework learnt from human demonstrations and used it in contact manipulation tasks. Inspired by human motor control theory, an adaptive impedance control that adapts the robot force, trajectory and impedance was developed in [8]. This study was extended later on into an iterative learning control framework dedicated for contact tooling tasks such cutting and drilling, as well as haptic exploration [9].

In the context of teleoperation, several works have aimed at endowing teleoperation architectures with variable impedance control capabilities. For example, Ajoudani et al. [10] proposed the tele-impedance approach, where the robot receives motion and stiffness commands from the human master, who in turn receives only visual feedback. While motion commands are

obtained via optical tracking, stiffness is estimated based on muscle activity measured via EMG. When compared to constant high and constant low stiffness executions for a ball catching and peg-in-hole tasks, the tele-impedance approach performed the best in terms of robustness and safety. Similarly, Yang et al. [11] used muscle activation signals to command the remote robot stiffness, however in a more conventional teleoperation system, showing again the superiority of their approach compared to a constant stiffness one. The advantages of variable impedance teleoperation over fixed high/low gains were also shown in [12], where grip force sensing was used to modulate the stiffness of the remote robot. The same authors combined the benefits of teleoperation with variable impedance actuation and control in [13]. The user input was there used to modulate stiffness of a variable impedance remote actuator, where low frequency modulation was controlled via software, while high frequency impedance was regulated via the actuator. Peternel et al. [14] proposed a teleoperation system for teaching robots both motion and stiffness profiles. This was achieved using a joystick-like device that allows the user to modulate the robot stiffness in real-time during demonstrations. Interestingly, the authors argued that the redundancy in the force control problem should be exploited whenever possible and achieve the task using a low stiffness. This becomes especially relevant in teleoperation, as high stiffness amplifies the so-called induced master motion, which affects safety and makes it difficult to maintain a stable contact with the environment [15]. As analyzed in [16], this induced motion results eventually in violent recoiling of the remote robot during contact, leading to an internal control loop that is unstable under high gains.

This work aims to combine the benefits of variable stiffness control with the capabilities of bilateral haptic teleoperation. It endows a robotic teleoperation system with an additional input channel conveying an automatically-estimated desired user stiffness, in addition to the desired motion. In the literature, most of the work in this topic requires additional external hardware, such as grip force sensing or EMG measurements. Furthermore, in works such as [14], additional cognitive load is imposed on the users as, in addition to controlling the position, they also need to manually control the stiffness. In contrast to previous works, we instead propose a variable stiffness teleoperation architecture that only uses the torque sensors embedded in commercial robotic manipulators, not requiring any additional hardware. In particular, we use user demonstrations to learn a variable stiffness control policy, envisioned for contact tasks that require the user to maintain a continuous interaction with the environment, e.g., drilling, cutting, or exploration tasks. Inspired by [7], we argue that, for such tasks, sensed external forces act primarily as a disturbance (e.g., friction), but can also convey valuable information for adapting the robot stiffness. We use LfD to learn the stiffness policies, relying on the stiffness information derived from the task dynamics, together with Gaussian mixture models (GMM) to encode a task model, and Gaussian mixture regression (GMR) for real-time execution. The learnt policy keeps the manipulator stiffness as low as possible for a safe and compliant interaction, increasing it only when needed to compensate for environmental disturbances. With respect to standard constant

high-stiffness approaches, we mitigate the effects of induced master motion that might lead to dangerous oscillations of the system. On the other hand, with respect to standard constant low-stiffness approaches, we maintain low tracking errors and completion times.

To the best of our knowledge, incorporating variable impedance learning schemes into bilateral teleoperation has not been considered before in the literature. The contributions of this paper can be summarized as follows:

- We present a variable impedance learning scheme applied to robotic teleoperation which, unlike previous works, does not require any external hardware (such as grip force or EMG sensors), or any cognitive effort from the side of the user in choosing the right stiffness.
- We ensure our system is passive by augmenting it with a passivity layer in the form of a global energy tank [17]. This guarantees additional safety and robustness, during interaction with unknown environments.

Fig. 1 shows a block diagram summarizing our approach. We demonstrate the effectiveness of our method in a representative cutting experiment, similar to [18]. However our algorithm can be used in any type of task where a continuous contact with the environment is required.

II. TASK LEARNING

In this section, we describe our procedure for learning the stiffness profiles. Similar to [6], [7], we use the mass-spring-damper model (MSD) to derive the stiffness. We also propose a method for computing the attractor path of the MSD, consistent with human-motor control theory. Then, we show the GMM/GMR approach and finally describe a method to construct a symmetric positive definite (SPD) stiffness matrix. In the following, for simplicity, we will only consider translational motion tasks. Future work will also consider rotational tasks as well. In the rest of this paper, unless stated otherwise, we use lowercase symbols to denote scalar quantities, lowercase bold symbols to denote vectors and uppercase bold symbols to denote matrices.

A. Task Dynamics

In order to learn the stiffness from a user demonstrated task, we first need to evaluate the stiffness profile from the corresponding task dynamics. To achieve this, the user provides a set of J demonstrations for Q task situations (a task situation can be for instance a certain material to cut with a specific hardness, in the case of a cutting task). During demonstrations, we collect the robot end effector positions $\{\mathbf{x}_t\}_{t=1}^T$ and the sensed external forces at the end effector $\{\mathbf{f}_{e,t}\}_{t=1}^T$, where t is a time parametrization and T the length of each demonstration. We augment these time series by the velocities $\{\dot{\mathbf{x}}_t\}_{t=1}^T$ and accelerations $\{\ddot{\mathbf{x}}_t\}_{t=1}^T$ obtained via finite differences.

For simplicity, in the following we will restrict the analysis along only one degree-of-freedom. For deriving the stiffness, we assume that along the motion direction, the end-effector behaves as a unit mass driven by a spring and a damper, subject to the external environment force. This can be expressed as

$$\ddot{x}_t = k_t(x_{t,d} - x_t) - b\dot{x}_t + f_{e,t}, \quad (1)$$

where k_t is the time varying stiffness, b is the damping chosen experimentally such that $b > 0$, while $x_{t,d}$ is the unknown spring equilibrium trajectory.

To obtain $x_{t,d}$, we assume that during demonstrations, the human teacher uses a stiffness composed of two parts: an unknown time varying part $k_{d,t}(f_{e,t})$ we seek to derive that continuously adapts to compensate for the external environment force, while a constant part k_c generates the observed motion dynamics, and such that $k_t = k_{d,t}(f_{e,t}) + k_c$. This assumption is also consistent with the way humans control their endpoint stiffness during interaction tasks. As stated in [19], one part of the human arm stiffness is constant and typically low (k_c), ensures stability during free motion, while the other part ($k_{d,t}(f_{e,t})$) increases to compensate for the stiffness of the environment. This implies that (1) can be approximated as

$$\ddot{x}_t = k_c(x_{t,d} - x_t) - b\dot{x}_t, \quad (2)$$

where k_c is set empirically, to a reasonable low value. We can then obtain $x_{t,d}$ as

$$x_{t,d} = k_c^{-1}(\ddot{x}_t + b\dot{x}_t) + x_t. \quad (3)$$

Similar to [7], we proceed to define a sliding window L that moves along the demonstrations. We define $\tilde{x}_t = x_{t,d} - x_t$, $y_t = \ddot{x}_t + b\dot{x}_t - f_{e,t}$, and concatenate adjacent data points from all the demonstrations that correspond to the specific task situations defining

$$\begin{aligned} \mathbf{X}_t &= [\tilde{x}_{t-L,1} \ \tilde{x}_{t-L+1,2} \ \dots \ \tilde{x}_{t+L,j}], \\ \mathbf{Y}_t &= [y_{t-L,1} \ y_{t-L+1,2} \ \dots \ y_{t+L,j}]. \end{aligned} \quad (4)$$

We then obtain k_t for each time instant via regularized regression as

$$k_t = (\mathbf{X}_t^T \mathbf{X}_t + \lambda \mathbf{I})^\dagger (\mathbf{X}_t \mathbf{Y}_t^T), \quad (5)$$

where λ is a tuning parameter to avoid ill-conditioning and penalize high values of k_t . Having obtained $k_{d,t} = k_t - k_c$, we construct the dataset $\{\bar{f}_{e,1:T,q}, k_{d,1:T,q}\}_{q=1}^Q$ where $\{\bar{f}_{e,1:T,q}\}$ is average sensed forces trajectory, from all demonstrations corresponding to the j th task situation.

B. Stiffness Learning

So far, we are able to derive stiffness profiles from user demonstrations. However, we still need a framework that would allow us to encode and learn a functional relationship between the sensed external force and the desired stiffness, and use it eventually for real-time reproduction of the learnt skill. To achieve this, we resort to GMM to learn a joint probability distribution between the sensed external force and the desired stiffness. First, let us define $\boldsymbol{\xi} = [\xi_i, \xi_o]^T$ where $\xi_i = f_{e,t}$ and $\xi_o = k_{d,t}$. We then model the joint distribution of the multi-dimensional variable $\boldsymbol{\xi}$ such as

$$\begin{aligned} \mathcal{P}(\xi_i, \xi_o) &= \sum_{n=1}^N \pi_n \mathcal{N}(\boldsymbol{\xi}, \boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n), \\ \boldsymbol{\mu}_n &= \begin{pmatrix} \mu_n^i \\ \mu_n^o \end{pmatrix}, \quad \boldsymbol{\Sigma}_n = \begin{bmatrix} \Sigma_n^i & \Sigma_n^{i,o} \\ \Sigma_n^{o,i} & \Sigma_n^o \end{bmatrix} \end{aligned} \quad (6)$$

where N is the number of gaussian models, which can be set empirically, or determined optimally from the data via e.g the

Bayesian-Information-Criterion (BIC). π_n , $\boldsymbol{\mu}_n$ and $\boldsymbol{\Sigma}_n$ are the priors, means and covariances of the gaussians, respectively. In order to learn these parameters from demonstrations, we apply the well-known iterative Expectation-Maximization (EM) algorithm. During reproduction, we seek to obtain the distribution of the output conditioned on the input. This can be achieved via GMR as

$$\begin{aligned} \mathcal{P}(\xi_o | \xi_i) &= \sum_{n=1}^N h_i(\xi_i) \mathcal{N}(\xi_o, \hat{\mu}_n^o, \hat{\Sigma}_n^o), \\ \hat{\mu}_n^o &= \mu_n^o + \Sigma_n^{o,i} \Sigma_n^{i,i}{}^{-1} (\xi_i - \mu_n^i), \\ \hat{\Sigma}_n^o &= \Sigma_n^o - \Sigma_n^{o,i} \Sigma_n^{i,i}{}^{-1} \Sigma_n^{i,o}, \\ h_i(\xi_i) &= \frac{\pi_n \mathcal{N}(\xi_i | \mu_n^i, \Sigma_n^i)}{\sum_{p=1}^P \pi_p \mathcal{N}(\xi_i | \mu_p^i, \Sigma_p^i)}. \end{aligned} \quad (7)$$

The desired output stiffness $k_{d,t}$ is then obtained as the mean of the conditional distribution.

C. Constructing a SPD stiffness

In the aforementioned analysis, we have only considered one cartesian direction. Indeed, we aim to adapt the stiffness along the direction of motion subject to the external environmental disturbance. Nevertheless, we still need to construct a symmetric positive definite stiffness matrix for the robot impedance controller. We achieve this by reconstructing a matrix that provides adaptation selectively along the direction of motion, computed in real-time. Formally, let us define a time window W that contains the measured end-effector velocities $\dot{\mathbf{x}}$ over the past W samples up to the current time instant t . We derive the current motion direction as

$$\mathbf{u}_t = \frac{\bar{\dot{\mathbf{x}}}_t}{\|\bar{\dot{\mathbf{x}}}_t\|}, \quad \bar{\dot{\mathbf{x}}}_t = \frac{1}{W} \sum_{t-W}^t \dot{\mathbf{x}}_w. \quad (8)$$

We also project the current sensed external force along the motion direction

$$f_{e,\mathbf{u}_t} = \mathbf{f}_e \cdot \mathbf{u}_t. \quad (9)$$

Let $\mathbf{v}_1 = \mathbf{u}_t$ and the vectors $\mathbf{v}_1 \dots \mathbf{v}_n$ span an orthonormal basis of \mathbb{R}^n where $\mathbf{v}_2, \dots, \mathbf{v}_n$ are normalized and orthogonal to \mathbf{v}_1 , and the matrix $\mathbf{V}(\mathbf{x}_t)$ whose columns are the vectors $\mathbf{v}_1 \dots \mathbf{v}_n$, we compute the stiffness at each time instant as

$$\mathbf{K}_t = \mathbf{V}(\mathbf{x}_t) \mathbf{A}(f_{e,\mathbf{u}_t}) \mathbf{V}(\mathbf{x}_t)^T, \quad (10)$$

where $\mathbf{A}(f_{e,\mathbf{u}_t})$ is a diagonal matrix with the first eigenvalue $a_1 = k_c + k_{d,t}(f_{e,\mathbf{u}_t})$ that adapts depending on the external force, and with $k_{d,t}(f_{e,\mathbf{u}_t})$ computed via the learnt GMM through regression, while the remaining eigen values can be set such that $a_2 \dots a_n \geq 0$.¹

¹For the tasks we consider, contact forces mainly arise along the direction of motion, such that the assignment of the stiffness eigenvalues in the remaining directions has little effect on performance.

III. VARIABLE IMPEDANCE TELEOPERATION

After learning the GMR model in the previous section, we present here the variable impedance teleoperation system that will adapt the stiffness of the remote robot impedance controller in real-time based on the sensed force from the environment.

Let $\mathbf{x}_r \in \mathbb{R}^3$ define the position of the remote robot, and $\mathbf{x}_m \in \mathbb{R}^3$ define that of the master device. The master device is modeled as a generic (gravity pre-compensated) mechanical system,

$$\mathbf{M}_m(\mathbf{x}_m)\ddot{\mathbf{x}}_m + \mathbf{C}_m(\mathbf{x}_m, \dot{\mathbf{x}}_m)\dot{\mathbf{x}}_m = \mathbf{f}_m + \mathbf{f}_h, \quad (11)$$

where $\mathbf{M}(\mathbf{x}_m) \in \mathbb{R}^{3 \times 3}$ is the positive-definite and symmetric inertia matrix, $\mathbf{C}_m(\mathbf{x}_m, \dot{\mathbf{x}}_m) \in \mathbb{R}^{3 \times 3}$ represents Coriolis/centrifugal terms, and $\mathbf{f}_m, \mathbf{f}_h \in \mathbb{R}^3$ are the control and operator forces, respectively. Similarly, on the level of the remote manipulator,

$$\mathbf{M}_r(\mathbf{x}_r)\ddot{\mathbf{x}}_r + \mathbf{C}_r(\mathbf{x}_r, \dot{\mathbf{x}}_r)\dot{\mathbf{x}}_r = \mathbf{f}_r + \mathbf{f}_e, \quad (12)$$

where $\mathbf{M}(\mathbf{x}_r) \in \mathbb{R}^{3 \times 3}$ is the positive-definite and symmetric inertia matrix, $\mathbf{C}_r(\mathbf{x}_r, \dot{\mathbf{x}}_r) \in \mathbb{R}^{3 \times 3}$ accounts for Coriolis/centrifugal terms, and $\mathbf{f}_r, \mathbf{f}_e \in \mathbb{R}^3$ are the control command and external forces exerted on the robot, respectively. For both the master and remote systems, we have the valuable passivity feature that lies in the skew symmetry of $(\dot{\mathbf{M}}_i - 2\mathbf{C}_i)$ for $i = \{r, m\}$, which implies that the remote robot and master dynamics are passive with respect to the ports $(\mathbf{f}_e + \mathbf{f}_r, \dot{\mathbf{x}}_r)$ and $(\mathbf{f}_h + \mathbf{f}_m, \dot{\mathbf{x}}_m)$, respectively.

The control command on the remote robot is designed as an impedance controller, such that the position of the robot end effector follows that of the master interface

$$\mathbf{f}_r = \mathbf{K}_t(\mathbf{x}_m - \mathbf{x}_r) - \mathbf{D}_r\dot{\mathbf{x}}_r, \quad (13)$$

where $\mathbf{K}_t \in \mathbb{R}^{3 \times 3}$ is our learnt state varying stiffness term and $\mathbf{D}_r \in \mathbb{R}^{3 \times 3}$ is a positive definite damping matrix. As for the master controller, it is designed as to reflect the environment force on the remote side as

$$\mathbf{f}_m = -\mathbf{f}_e - \mathbf{D}_m\dot{\mathbf{x}}_m, \quad (14)$$

where \mathbf{D}_m is a positive definite damping matrix, added to improve contact stability.

IV. PASSIVITY ANALYSIS

While the proposed approach benefits from some intrinsic robustness as the stiffness is increased only when needed and as dictated by the interaction with the remote environment, the stable interaction with arbitrary passive environments is still not guaranteed. This is due to the fact the the proposed teleoperator system is non-passive. In general, it is well known that a force-reflecting teleoperation system is non-passive even in the non-delayed case [20]. The problem is further aggravated in case of a varying stiffness, as stiffness variations represent an active action [21]–[23]. Since environments can be assumed passive and the human operator hand impedance is often characterized by a passive velocity to force map, it suffices to passify the controlled teleoperator system to guarantee overall stable operation. To achieve that, we will use the concept of energy tanks which were featured prominently in many

robotic applications such as variable stiffness control [21], passification of null space projections [24] and delayed bilateral tele-manipulation [25], [26]. Similarly, we will augment our controllers with a passivity layer that consists of a global energy tank. An energy tank can be regarded as a virtual storage element with a certain allocated energy budget that can be used for the execution of potentially non-passive control actions [17]. This budget, reflected by the choice of the initial tank energy and consequently its initial state, should be chosen to provide a reasonable compromise between safety and performance, since very high initial values can mask unstable behaviors. The choice of one global tank compared to two local tanks on master and remote robot levels as done in [25] is motivated by the work in [24], where the authors noted that a global tank is less conservative than local tanks.

First, let us analyze the passivity of our system with the following storage function

$$S = \frac{1}{2}\tilde{\mathbf{x}}_r^T \mathbf{K}_t \tilde{\mathbf{x}}_r + \frac{1}{2}\dot{\mathbf{x}}_m^T \mathbf{M}_m \dot{\mathbf{x}}_m + \frac{1}{2}\dot{\mathbf{x}}_r^T \mathbf{M}_r \dot{\mathbf{x}}_r, \quad (15)$$

where $\tilde{\mathbf{x}}_r = \mathbf{x}_m - \mathbf{x}_r$. The derivative of (15) along the closed loop dynamics of the system can be expressed as

$$\begin{aligned} \dot{S} = & \dot{\mathbf{x}}_r^T \mathbf{f}_e + \dot{\mathbf{x}}_m^T \mathbf{f}_h - \dot{\mathbf{x}}_r^T \mathbf{D}_r \dot{\mathbf{x}}_r - \dot{\mathbf{x}}_m^T \mathbf{D}_m \dot{\mathbf{x}}_m \\ & - \dot{\mathbf{x}}_m^T \mathbf{f}_e + \dot{\mathbf{x}}_m^T \mathbf{K}_t (\mathbf{x}_m - \mathbf{x}_r) + \frac{1}{2}\tilde{\mathbf{x}}_r^T \dot{\mathbf{K}}_t \tilde{\mathbf{x}}_r, \end{aligned} \quad (16)$$

where it is required to prove passivity with respect to the ports $(\dot{\mathbf{x}}_m, \mathbf{f}_h)$ and $(\dot{\mathbf{x}}_r, \mathbf{f}_e)$ which represent the ports of interaction through which the teleoperator system interacts with the operator and the remote environment. Unfortunately however, the sign of the last three terms is indefinite implying that the system can be active. In order to solve this problem, we will augment the controllers by a global energy tank with a state $z \in \mathbb{R}$ that exchanges energy with both the master and robot controllers, and regulates their output depending on the level of tank energy. We define our tank dynamics as

$$\begin{aligned} \dot{z} = & \frac{\sigma_m \dot{\mathbf{x}}_m^T \mathbf{D}_m \dot{\mathbf{x}}_m}{z} + \frac{\sigma_r \dot{\mathbf{x}}_r^T \mathbf{D}_r \dot{\mathbf{x}}_r}{z} + \frac{\lambda_m(E(z)) \dot{\mathbf{x}}_m^T \mathbf{D}_{m,h} \dot{\mathbf{x}}_m}{z} \\ & + \frac{\lambda_r(E(z)) \dot{\mathbf{x}}_r^T \mathbf{D}_{r,h} \dot{\mathbf{x}}_r}{z} - \frac{\alpha_r \dot{\mathbf{x}}_r^T \mathbf{K}_t (\mathbf{x}_m - \mathbf{x}_r)}{z} + \frac{\alpha_m \dot{\mathbf{x}}_m^T \mathbf{f}_e}{z}, \end{aligned} \quad (17)$$

and with tank energy $E(z) = \frac{1}{2}z^2$. As for σ_m and σ_r , they control the amount of energy filled in the tank through dissipation. They are defined as

$$\begin{cases} 0 < \sigma_r, \sigma_m \leq 1 & \text{if } E(z) < \bar{E}, \\ 0 & \text{otherwise,} \end{cases} \quad (18)$$

where \bar{E} is an absolute upper limit for the maximum energy allowed in the tank to avoid practically unstable behavior. As for the valves α_m and α_r , they have the role of modifying the control actions if the tank is depleted. They are defined as

$$\begin{cases} 0 < \alpha_r, \alpha_m \leq 1 & \text{if } E(z) > \underline{E}, \\ 0 & \text{otherwise,} \end{cases} \quad (19)$$

where \underline{E} is a lower threshold for the minimum energy allowed in the tank, to prevent singularities. In order to avoid complete stoppage of the task whenever the tank becomes depleted, we




Fig. 2. Experimental setup used for validating our approach on a cutting task. Left: Our Omega3 haptic device used as master. Right: Our 7 DOF KUKA lightweight robot with a mounted cutting scalpel. The black clay is placed on a piece of foam on the table and serves as the material to cut.

propose to add the valves $\lambda_m(E(z))$ and $\lambda_r(E(z))$ that have the role of harvesting additional dissipation energy to refill the tank by injecting local damping on the master and robot controllers. In order to avoid compromising task performance, this damping energy is only requested whenever the energy in the tank drops below a threshold E_{thresh} and where $\underline{E} \leq E_{thresh} \leq \bar{E}$. $\lambda_m(E(z))$ and $\lambda_r(E(z))$ are designed as smoothly rising functions from 0 to 1, while $D_{r,h}$ and $D_{m,h}$ are design parameters. We now modify the master and robot controllers as

$$\begin{aligned} \mathbf{f}_r &= \alpha_r \mathbf{K}_t (\mathbf{x}_{r,d} - \mathbf{x}_r) - \mathbf{D}_r \dot{\mathbf{x}}_r - \lambda_r \mathbf{D}_{r,h} \dot{\mathbf{x}}_r, \\ \mathbf{f}_m &= -\alpha_m \mathbf{f}_e - \mathbf{D}_m \dot{\mathbf{x}}_m - \lambda_m \mathbf{D}_{m,h} \dot{\mathbf{x}}_m, \end{aligned} \quad (20)$$

and analyze passivity with the following storage

$$S = E(z) + \frac{1}{2} \dot{\mathbf{x}}_m^T \mathbf{M}_m \dot{\mathbf{x}}_m + \frac{1}{2} \dot{\mathbf{x}}_r^T \mathbf{M}_r \dot{\mathbf{x}}_r. \quad (21)$$

We can now make the following statement regarding the passivity of the system:

Proposition 1: Consider the teleoperator system given by (11) and (12), with corresponding master and robot controllers (20). The system represents a passive map with respect to the input pair $(\mathbf{f}_e, \mathbf{f}_h)$ and the output pair $(\dot{\mathbf{x}}_r, \dot{\mathbf{x}}_m)$. The passivity claim can be validated by taking the rate of (21), which can be expressed as

$$\begin{aligned} \dot{S} &= \dot{\mathbf{x}}_r^T \mathbf{f}_e + \dot{\mathbf{x}}_m^T \mathbf{f}_h - (\sigma_r - 1) \dot{\mathbf{x}}_r^T \mathbf{D}_r \dot{\mathbf{x}}_r - \\ &\quad (\sigma_m - 1) \dot{\mathbf{x}}_m^T \mathbf{D}_m \dot{\mathbf{x}}_m, \end{aligned} \quad (22)$$

where the last two terms are negative definite, according to the definitions of σ_m and σ_r in (18), which concludes our passivity proof.

V. EXPERIMENTAL EVALUATION

To validate our approach, we conducted a series of experiments to verify both the learning and task execution. We chose cutting as a representative task. First, we briefly describe our learning phase, followed by a detailed experimental evaluation where we tested our method and compared it with constant high- and low-stiffness baselines.

As a master device, we used the Omega 3 haptic device commanded in force control mode, while the remote robot consisted of the 7 DOF KUKA lightweight robot, which was commanded in cartesian impedance control mode. For all




Fig. 3. Results of the GMR on training and test data. Left column: Force profiles used as input to the GMR. Right column: Stiffness profiles obtained via regression (red) and learnt stiffness profiles (dotted black).

experiments, the orientation of the robot end-effector was kept constant via active control. In order to perform the cutting task, a custom tool holder was designed, in which a cutting scalpel was fitted, and attached to the robot via a Quick-change adapter system, as shown in Fig. 2. A material to cut is placed on the table in front of the robot. The master device is 3 DoF and it is thus free to move along the 3 translation directions. The same constraint is imposed on the remote robot, as its orientation is fixed at the beginning of the experiment through control to make sure the scalpel is always perpendicular to the material surface.

This setup is used both for the task learning and the reproduction phase. Of course, the learning part could also be done through kinesthetic teaching. However, direct interaction is not always possible with the robot in real settings (e.g., in nuclear applications). Furthermore, the physical coupling between the human and the robot induces additional dynamics during the teaching stage [14]. To avoid these drawbacks, we performed our demonstrations using direct teleoperation with haptic feedback.

A. Task Learning

For the learning phase, we conducted teleoperated cutting experiments on four materials with different stiffness properties. For each material, we collected five demonstrations. Data from all demonstrations were segmented to identify the contact phase, based on a pre-defined threshold of the contact force along the motion direction and subjected to low-pass filtering to remove sensor noise, prior to obtaining the higher order derivatives. This was followed by the stiffness estimation phase, where the algorithms described in Sec. II were run over the data. We used $k_c = 90$ N/m, while the damping was set as $b = 2\sqrt{k_c}$ N/m². As explained in Sec. II, the k_c describes the constant low component of the human arm stiffness. Therefore, the chosen value of k_c represents a relatively low stiffness in our robotic system (typically below 100 N/m), and was also used for the constant low stiffness experiments. We specified a window

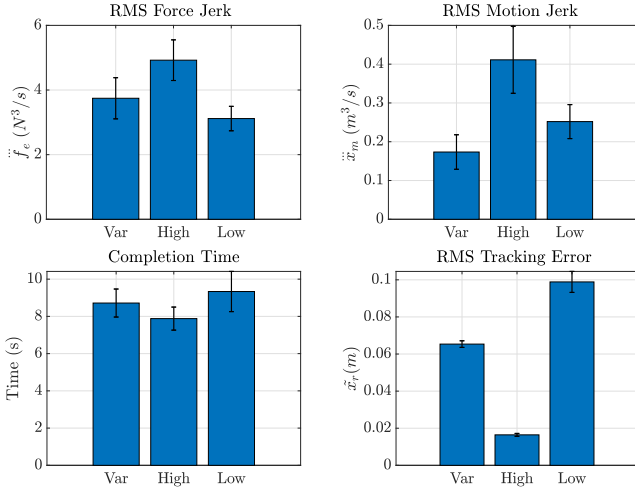


Fig. 4. Cutting one material, no delay. Bar plot comparison of the different metrics among variable, constant low and constant high stiffness, where the X-axis represents the condition, while the y axis represent the mean of the corresponding metric. The vertical black lines in each bar indicate the standard deviation.

length of $L = 3$, a regularization value $\lambda = 0.005$ and used a GMM parameterized with 4 states.

Fig. 3 shows the results of our learning. The right column shows the stiffness profiles obtained by GMR (dotted black) both on force data used in the training phase (upper), and on unseen force data (lower), compared to the stiffness profiles obtained via regularized regression (red). The left column shows the force profiles used as input to the regression process, which represent the mean from all demonstrations of a material.

B. Robot Validation

We validate our approach on the teleoperation system described above and shown in Fig. 2. The following comparisons were considered:

- 1) The proposed approach vs. two baseline conditions using constant low and constant high stiffness for teleoperated cutting on an unseen material with no communication delay.
- 2) The proposed approach vs. one baseline condition using constant high stiffness for teleoperated cutting on a given material with a communication delay of 20 ms.
- 3) The proposed approach vs. one baseline condition using constant low stiffness for teleoperated cutting through two different materials with no communication delay.

1) *Cutting one material, no delay*: In the first series of experiments, we aim to test our proposed solution when cutting a material which was not used during the demonstrations. We compare our approach against baseline executions of constant low stiffness (90 N/m) and constant high stiffness (800 N/m). For each condition, fifteen trials were conducted where a human operator commands the remote robot with similar motions. For fairness of comparison, the depth of cut was not controlled by the user but was maintained constant in all trials. We compare the three conditions with respect to the following performance metrics:




Fig. 5. Cutting one material, with delay. The upper row shows the motion trajectories of the master (blue) and remote manipulator (orange) both for the variable case (left) and the high stiffness case (right). The left-bottom figure shows the sensed external forces for the high stiffness case (blue) and for the variable case (orange), while the bottom right figure shows the stiffness profile for the variable stiffness. All variables are displayed along the axis of motion.

- Root mean square (RMS) of motion tracking error between the master interface and the remote robot;
- Task completion time;
- RMS of the remote robot motion jerk;
- RMS of the remote robot force jerk.

While the first two metrics are an indication of the task execution speed, we choose the last two metrics as a measure of smoothness in the task execution [27]. Jerk metrics have also been used to characterize physical fatigue levels [28]. The results of this study are shown in Fig. 4 as a bar plot showing the mean of the analyzed metrics over trials, together with the standard deviation. As expected, the RMS tracking error is the highest for the low stiffness case, with a mean across trials of 0.0989 ± 0.0057 m compared to a mean of 0.0164 ± 0.0008 m for the high stiffness case and 0.0654 ± 0.0018 m for the proposed variable case. Similarly, the completion time is highest in the low stiffness case (9.3359 ± 1.0792 s) when compared to the variable stiffness and high stiffness cases, which have relatively lower completion times of 8.7171 ± 0.7522 s and 7.8807 ± 0.6193 s. On the other hand, in terms of smoothness, our variable stiffness approach outperforms the high stiffness case. The former shows a RMS motion jerk and RMS force jerk of 0.1735 ± 0.0446 ms^{-3} and 3.7448 ± 0.6368 Ns^{-3} , respectively, compared to 0.4112 ± 0.0863 ms^{-3} and 4.9233 ± 0.6291 Ns^{-3} for the high stiffness case.

2) *Cutting one material, with delay*: The difference in performance between the variable and high stiffness cases can be better seen in the presence of disturbances, e.g., communication delays. It is well known that for force reflecting teleoperation, induced master motion can result in violent recoiling of the remote robot during contact, a problem further aggravated with delays. Therefore, we conducted a representative experiment where an artificial delay of 20 ms between master and remote systems was simulated. This delay




Fig. 6. Cutting two materials, no delay. Results of the transition from soft to hard materials while cutting. The left figure shows the tracking error for the low (orange) and variable (blue) stiffness case. The right figure shows the variable stiffness profile generated from learnt GMM. The dotted vertical lines denote the moment of transition from soft to hard, where the black line indicates the moment of transition for the low stiffness case ($t \approx 5.2$), while the dotted green denotes the transition for the variable case ($t \approx 4.2$). For better readability of the graphs, in the left figure we show the moment of transition for the low stiffness case only.

was chosen based on typically used values in the literature on delayed bilateral teleoperation [20]. In this case, we kept the energy tank off, allowing us to showcase the difference in performance/safety solely from the considered stiffness method. Results are shown in Fig. 5. It depicts the motion profiles of the master and the remote robot, for both the variable and the constant high stiffness cases (upper row), the external forces recorded using the robot force sensors, and the stiffness profile for the variable stiffness case. Clearly, for the high stiffness case, the task becomes quickly unstable, inducing high oscillations in the motion and force profiles, endangering both the task performance and the robot safety. As expected, for the variable stiffness case, oscillations are significantly smaller in the motion, force and stiffness profiles, indicating a much smoother task execution. It is also worth noting that execution time is also shorter for the variable case, where the task is completed in around 1.8 s less.

3) *Cutting two materials, no delay*: While the above situation shows the merit of our method compared to constant high stiffness in terms of safety, here we analyze the benefit of our approach compared to a constant low stiffness approach. More specifically, we test our approach in a situation where it is required to cut through a piece composed of a soft part and then a hard part. The results of this experiment are shown in Fig. 6. On the right, it shows the stiffness profile for the variable stiffness case. On the left, it shows the absolute tracking error between the master and robot motion, along the y axis (the motion direction), both for the variable stiffness case (blue) and the low stiffness case (orange). The stiffness profile starts increasing smoothly at the moment of transition, at $t \approx 4.2$ s, to account for the increased hardness of the materials, which is reflected in higher sensed external forces. This adaptation results in the motion tracking error between master and remote robot to remain almost constant throughout the task. On the other hand, in the low stiffness case, the tracking error becomes higher after the transition at $t \approx 5.2$ s, due to the inability of the robot controller to adapt to the higher resistance force of the harder material.

VI. DISCUSSION

Previous works have analyzed the potential benefits of incorporating variable impedance control into a teleoperation

architecture. Our experimental results further consolidate these findings. In fact, while using a constant high stiffness (with respect to the environment) results in better tracking accuracy and faster execution, it poses potentially-dangerous stability issues. On the other hand, using a constant low stiffness results in degraded tracking accuracy and higher completion times, but guarantees higher safety margins. Our approach aims at combining the best of the two worlds. By giving preference to a low stiffness that increases only when needed, we are able to achieve a good compromise between task execution speed, tracking accuracy, smoothness, and safety. Compared to previous works on the topic, our system does not need any cognitive effort from the user, as the stiffness is chosen automatically at runtime, or any additional sensor on the user or robot sides.

The efficiency of our method was confirmed in the first experimental study of Sec. V-B1, where we achieved task completion times comparable to those of the high stiffness method but with the advantage of enhanced smoothness, both on motion and force levels. This difference can be attributed to the oscillations resulting from the induced master motion, which become higher with a high stiffness gain, making it harder for the human operator to maintain a stable and smooth contact. This problem becomes significantly aggravated even in presence of a delay in the communication channel. As expected, our approach is much more robust in this regard, even without applying any stability related-solutions dedicated for delayed bilateral telemanipulation (see Sec. V-B2). It could be argued that a slightly lower constant high stiffness might improve the performance. However, the choice of this parameter to provide a compromise between accuracy and robustness is task specific and not always easy to determine. On the contrary, our algorithm determines this desired stiffness automatically. Furthermore, we validated the ability of our approach to adapt to situations that feature force variations typical in a transition from a soft to hard material in cutting, or variations induced by a gravitational load during a lifting task [7]. The stiffness increases smoothly when the hard material is detected, maintaining the tracking error constant and not requiring the operator to exert any additional effort to compensate for it by, e.g., stiffening up his/her arm. This result implies that the human operator can focus on the careful execution of the task, without any distraction, while the robot automatically adapts its interaction force with the environment (see Sec. V-B3). It is worth noting that such an adaptation is made possible thanks to the fact that our GMM is not time-driven, but instead it performs the regression on the perceived external forces, alleviating the necessity of any additional vision module.

VII. CONCLUSION

In this paper, we presented a variable impedance teleoperation system for contact tasks. It automatically adapts to the environment by changing the control stiffness between the master interface and remote robot. It uses multiple demonstrations on different materials to learn the stiffness model with respect to the applied force at the remote environment. This variable stiffness estimation at run-time is then incorporated in the teleoperation architecture and allows the robot to successfully

perform the task in any type of material and adapt to new environment. We performed a passivity analysis on our system and ensure its passivity using energy tanks. We then tested our algorithm in three cutting tasks, including a situation with communication delay and one where the robot had to cut through two different materials. Results showed the advantages of our approach compared to a constant high or a constant low stiffness.

In the future, we would like to perform a more comprehensive user study to test our system viability across different operators and materials, vs. other state-of-the-art variable stiffness methods, as well as to register the subjective evaluation of our method from the users. In addition, we would like to study the effect of the proposed approach on the user comfort and effort put in the task. In fact, our method showed to have a better motion jerk compared to the high stiffness trials, and since [28] showed that jerk could be an indication of physical fatigue, we would like to study the ergonomic effect of our method on the user, using, e.g., EMG data. Since our approach proved robust to small delays, another interesting future direction would be to compare the performance of our variable impedance approach to standard energy tanks applied for the high stiffness case in a delayed teleoperation scenario. Furthermore, we will extend our approach to teleoperation tasks that require a change in orientation, such as screwing.

REFERENCES

- [1] K. P. Tee, E. Burdet, C. M. Chew, and T. E. Milner, "A model of force and impedance in human arm movements," *Biological Cybernetics*, vol. 90, no. 5, pp. 368–375, May 2004.
- [2] D. W. Franklin, R. Osu, E. Burdet, M. Kawato, and T. E. Milner, "Adaptation to stable and unstable dynamics achieved by combined impedance control and inverse dynamics model," *Journal of neurophysiology*, vol. 90, no. 5, pp. 3270–82, 2003.
- [3] J. Buchli, E. Theodorou, F. Stulp, and S. Schaal, "Variable impedance control - a reinforcement learning approach," in *Robotics: Science and Systems*, 2010, pp. 153–160.
- [4] K. Kronander and A. Billard, "Learning compliant manipulation through kinesthetic and tactile human-robot interaction," *IEEE Transactions on Haptics*, vol. 7, no. 3, pp. 367–380, 2014.
- [5] E. Gribovskaya, A. Kheddar, and A. Billard, "Motion learning and adaptive impedance for robot control during physical interaction with humans," in *IEEE International Conference on Robotics and Automation*, 2011, pp. 4326–4332.
- [6] L. Rozo, S. Calinon, D. G. Caldwell, P. Jiménez, and C. Torras, "Learning physical collaborative robot behaviors from human demonstrations," *IEEE Transactions on Robotics*, vol. 32, no. 3, pp. 513–527, 2016.
- [7] F. J. Abu-Dakka, L. Rozo, and D. G. Caldwell, "Force-based variable impedance learning for robotic manipulation," *Robotics and Autonomous Systems*, vol. 109, pp. 156 – 167, 2018.
- [8] G. Ganesh, N. Jarrassé, S. Haddadin, A. Albu-Schaeffer, and E. Burdet, "A versatile biomimetic controller for contact tooling and haptic exploration," in *IEEE International Conference on Robotics and Automation*, May 2012.
- [9] Y. Li, G. Ganesh, N. Jarrassé, S. Haddadin, A. Albu-Schaeffer, and E. Burdet, "Force, impedance, and trajectory learning for contact tooling and haptic identification," *IEEE Transactions on Robotics*, vol. 34, no. 5, pp. 1170–1182, 2018.
- [10] A. Ajoudani, N. Tsagarakis, and A. Bicchi, "Tele-impedance: Teleoperation with impedance regulation using a body-machine interface," *The International Journal of Robotics Research*, vol. 31, no. 13, pp. 1642–1656, 2012.
- [11] C. Yang, C. Zeng, P. Liang, Z. Li, R. Li, and C. Su, "Interface design of a physical human-robot interaction system for human impedance adaptive skill transfer," *IEEE Transactions on Automation Science and Engineering*, vol. 15, no. 1, pp. 329–340, 2018.
- [12] D. S. Walker, J. K. Salisbury, and G. Niemeyer, "Demonstrating the benefits of variable impedance to telerobotic task execution," in *IEEE International Conference on Robotics and Automation*, 2011, pp. 1348–1353.
- [13] D. S. Walker, R. P. Wilson, and G. Niemeyer, "User-controlled variable impedance teleoperation," in *IEEE International Conference on Robotics and Automation*, 2010, pp. 5352–5357.
- [14] L. Pernel, T. Petrić, and J. Babić, "Robotic assembly solution by human-in-the-loop teaching method based on real-time stiffness modulation," *Autonomous Robots*, vol. 42, pp. 1–17, 2018.
- [15] C. Pacchierotti, *Cutaneous haptic feedback in robotic teleoperation*. Springer, 2015.
- [16] K. J. Kuchenbecker and G. Niemeyer, "Induced Master Motion in Force-Reflecting Teleoperation," *Journal of Dynamic Systems, Measurement, and Control*, vol. 128, no. 4, pp. 800–810, 04 2006.
- [17] S. Stramigioli, "Energy-aware robotics," in *Mathematical Control Theory I*, M. K. Camlibel, A. A. Julius, R. Pasumarthy, and J. M. Scherpen, Eds. Cham: Springer International Publishing, 2015, pp. 37–50.
- [18] R. Rahal, F. Abi-Farraj, P. R. Giordano, and C. Pacchierotti, "Haptic shared-control methods for robotic cutting under nonholonomic constraints," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 8151–8157.
- [19] Y. Li, N. Jarrassé, and E. Burdet, "Versatile interaction control and haptic identification in humans and robots," in *Geometric and Numerical Foundations of Movements*. Springer, 2017, pp. 187–206.
- [20] D. Heck, "Delayed bilateral teleoperation : a direct force-reflecting control approach," Ph.D. dissertation, Technische Universiteit Eindhoven, 2015.
- [21] F. Ferraguti, C. Secchi, and C. Fantuzzi, "A tank-based approach to impedance control with variable stiffness," in *IEEE International Conference on Robotics and Automation*, May 2013, pp. 4948–4953.
- [22] Y. Michel, C. Ott, and D. Lee, "Passivity-based variable impedance control for redundant manipulators," in *IFAC World Congress*, 2020.
- [23] C. Pacchierotti, A. Tirmizi, G. Bianchini, and D. Prattichizzo, "Enhancing the performance of passive teleoperation systems via cutaneous feedback," *IEEE Transactions on Haptics*, vol. 8, no. 4, pp. 397–409, 2015.
- [24] A. Dietrich, X. Wu, K. Bussmann, C. Ott, A. Albu-Schäffer, and S. Stramigioli, "Passive hierarchical impedance control via energy tanks," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 522–529, April 2017.
- [25] M. Franken, S. Stramigioli, S. Misra, C. Secchi, and A. Macchelli, "Bilateral telemanipulation with time delays: A two-layer approach combining passivity and transparency," *IEEE Transactions on Robotics*, vol. 27, no. 4, pp. 741–756, Aug 2011.
- [26] G. Bianchini, J. Bimbo, C. Pacchierotti, D. Prattichizzo, and O. A. Moreno, "Transparency-oriented passivity control design for haptic-enabled teleoperation systems with multiple degrees of freedom," in *IEEE Conference on Decision and Control*, 2018, pp. 2011–2016.
- [27] S. Balasubramanian, A. Melendez-Calderon, A. Roby-Brami, and E. Burdet, "On the analysis of movement smoothness," *Journal of NeuroEngineering and Rehabilitation*, vol. 12, 12 2015.
- [28] L. Zhang, M. M. Diraneyya, J. Ryu, C. T. Haas, and E. M. Abdel-Rahman, "Jerk as an indicator of physical exertion and fatigue," *Automation in Construction*, vol. 104, pp. 120–128, 2019.