

Learning-Based Adaptive Impedance Control Towards Safe Autonomous Transseptal Puncture

Anna Bicchi¹, *Student Member, IEEE*, Eleonora Pollini¹, Junling Fu¹, *Graduate Student Member, IEEE*, Federica Gramegna², Elena De Momi¹, *Senior Member, IEEE*

Abstract—TransSeptal Puncture (TSP) is a key step in many minimally invasive cardiac procedures, enabling access to the left atrium by crossing the interatrial septum from the right atrium. The task demands extreme precision, as excessive force may cause cardiac tamponade. Robotic platforms can improve precision and repeatability, but most of the existing systems are designed for training or teleoperation rather than autonomous execution, and do not incorporate adaptive impedance modulation. This work presents a learning-from-demonstration framework based on probabilistic impedance modeling to investigate the feasibility of learning and embedding expert force-regulation strategies for autonomous execution of the TSP under clinician-defined targets. Specifically, Gaussian Mixture Models, trained on teleoperated demonstrations, capture the relationship between contact force and operator stiffness, enabling real-time modulation of impedance through Gaussian Mixture Regression. The adaptive controller was deployed on a 7-DoF robotic platform and validated on fossa ovalis phantoms of varying compliance and on *ex vivo* porcine tissue. Performance was evaluated using clinically relevant metrics: Needle Puncture Force (NPF), Tenting Distance (TD), and Needle Stopping Space (NSS). Compared to fixed-stiffness baselines, the proposed controller reduced NPF by up to 35%, maintained TD within safe limits, and limited NSS below 1 mm in nominal anatomies, with consistent performance also observed under extreme phantom anatomies and *ex vivo* tissue. These findings demonstrate that embedding human-like impedance modulation can enable safe, anatomy-aware autonomous control, advancing robotic TSP toward clinical feasibility.

Note to Practitioners—This study proposes a learning-based adaptive impedance control scheme for TSP, a delicate step in minimally invasive cardiac procedures where excessive puncture forces can lead to severe complications. While robotic systems can enhance precision and repeatability, existing platforms cannot typically autonomously adapt to patient-specific anatomical variability, relying instead on operator input. The proposed approach consists of two main steps. First, impedance profiles are extracted from teleoperated demonstrations, capturing the relationship between applied force and operator stiffness through a probabilistic model. Then, during autonomous execution, the robot leverages this model to modulate its impedance in real time, adapting its behavior to the compliance of the cardiac tissue to perform safe and precise punctures. By embedding human-like impedance modulation into the control loop, this framework combines surgical expertise with robotic precision, improving procedural safety and consistency. Beyond TSP, the same strategy can be applied to other surgical tasks that require adaptation to different patient anatomies, such as vascular access,

soft tissue insertion, or targeted puncture, by training the model with demonstrations from experienced surgeons.

Index Terms—Adaptive Impedance Control, Learning from Demonstrations, Human Impedance Modeling, Minimally Invasive Surgical Robotics, Autonomous Transseptal Puncture

I. INTRODUCTION

Humans possess remarkable adaptive capabilities, abstracting essential task goals and dynamically adjusting their actions to accommodate changing circumstances. These qualities have long inspired robotics research aimed at endowing machines with similar flexibility [1]. In surgical settings, clinicians develop highly refined perceptual skills by integrating visual and tactile feedback to navigate unpredictable intraoperative conditions. Inspired by these human adaptive strategies, Learning from Demonstrations (LfD) approaches in surgical robotics aim to capture expert demonstrations and translate them into controllers that can generalize across diverse clinical contexts and anatomical variations.

TransSeptal puncture (TSP) is a procedure widely used in many cardiac surgeries that requires precise tactile and visual feedback to navigate complex and variable anatomical conditions and interact with cardiac tissue. Indeed, TSP enables access to the left atrium from the right atrium, playing a crucial role in a range of minimally invasive structural cardiac interventions [2], such as mitral valve repair, left atrial appendage closure, and ablation [3]. Although generally safe when performed by experienced operators, TSP remains technically demanding and carries a risk of serious complications [4], [5]. The most significant risks associated with TSP include cardiac tamponade, pericardial effusion, inadvertent aortic puncture, and, more rarely, complications such as pneumopericardium and embolic stroke [6], [7]. A key determinant of procedural success is the accurate localization of the Fossa Ovalis (FO), a thinner, more compliant region of the interatrial septum, representing the ideal puncture site. Under fluoroscopic guidance, clinicians typically advance the catheter into the right atrium from the vena cava and identify the FO by detecting a thin mechanical “step” as the catheter is retracted while in contact with the interatrial septum. This maneuver, although widely practiced, relies heavily on tactile feedback and image interpretation, underscoring the complexity and operator dependency of the procedure [8], [9]. After identifying the FO, the clinician is required to adjust the force applied to the catheter to safely perforate the tissue. This step requires fine modulation of interaction forces in response

¹Department of Electronics, Information and Bioengineering (DEIB), Politecnico di Milano, 20133 Milan, Italy (e-mail: anna.bicchi@polimi.it; eleonora.pollini@mail.polimi.it, junling.fun@polimi.it, elena.demomi@polimi.it).

²IRCCS Humanitas Research Hospital, Milan, Italy (e-mail: federica.gramegna@humanitas.it)

Corresponding author: A. Bicchi (e-mail: anna.bicchi@polimi.it).

to the local mechanical properties of the interatrial septum, without direct measurement of contact force. Excessive force may lead to perforation beyond the septum, while insufficient force can result in prolonged tenting and loss of procedural control. As a result, force regulation during puncture is highly skill-dependent and represents one of the most critical aspects of TSP. Moreover, while fluoroscopic imaging is effective for navigation, it exposes both patients and clinicians to ionizing radiation. As a result, an emerging research direction aims to reduce reliance on fluoroscopy by leveraging alternative sensing modalities for tool position feedback, such as ultrasound imaging and embedded sensing technologies, including electromagnetic tracking and fiber Bragg grating sensors [10]. Robotic systems offer a promising solution to overcome these limitations by enhancing precision, repeatability, and the integration of haptic feedback through force sensing [11], [12]. Several studies have explored robotic solutions for TSP, mainly as tools for surgical training [13], [14] or for improving accuracy in teleoperated setups [15], [16]. In a previous study, we investigated the feasibility of an autonomous TSP and how puncture velocities can affect it [17]. Despite technological advances, a fundamental limitation of current robotic systems for TSP lies in their inability to replicate the adaptive behavior exhibited by clinicians. During the procedure, experienced operators instinctively adjust arm stiffness and damping as they transition from navigating towards the FO to applying the controlled force required for septal puncturing, adjusting it in response to the specific stiffness of the patient's tissues. This skill is based on the human capability of modulating their endpoint force and impedance through coordinated action of their spring-like muscles [18], [19]. This regulation enables humans to interact with various environments skillfully, adapting both force and compliance according to the contact conditions, uncertainty, or phase of the task.

This principle has inspired the development of the Adaptive Impedance Control (AIC) strategies in robotics, particularly in teleoperation and physical human-robot interaction [20]. In this respect, recent studies have explored learning-based approaches to acquire human-like variable impedance. A skill transfer framework combining Electromyography (EMG)-based stiffness estimation with dynamic movement primitives has been proposed to enable robots to learn and generalize human-like variable impedance strategies for contact-rich tasks [21]. A shared autonomy framework has been proposed that leverages LfD to encode complex motion and wrench profiles via virtual potential fields, integrating variable impedance control and authority arbitration to enable compliant and adaptive human-robot collaboration [22]. A variable impedance control scheme has been proposed for autonomous guidewire delivery in vascular interventions, where the surgeon's insertion skill is modeled via a distal position-based impedance formulation, allowing the robot to replicate smooth and responsive manipulation while reducing operational workload [23]. Another learning-based bilateral teleoperation framework employs user demonstrations and onboard torque sensing to train a variable stiffness control policy, achieving enhanced force tracking and stability without requiring additional hardware [24]. Similarly, in the context of

soft tissue puncture, a human-in-the-loop learning framework integrates EMG-driven impedance estimation with a target-directed motion planner, enabling the robot to dynamically adjust stiffness in response to tissue variability and external forces [25]. A similar approach was proposed in [26], where human impedance capabilities were also estimated based on EMG signals in a task aimed at improving collaborative manipulation. In [27], a Human Impedance Model (HIM) was estimated based on muscle activation patterns to adjust robot impedance online during teleoperated tasks.

In this study, we propose a framework that leverages LfD to estimate skilled users' impedance profiles from demonstration data and integrates this model into an online adaptive impedance control strategy. Applied to autonomous TSP, this approach enables context-sensitive modulation of robot stiffness and damping, replicating human-like adaptation across task phases and anatomical variability. By embedding a HIM-based adaptive control strategy into the robotic system, along with a sensing framework aimed at reducing reliance on fluoroscopy and associated radiation exposure, our goal is to bridge the gap between human motor intelligence and autonomous robotic performance in high-risk surgical procedures such as TSP.

II. MATERIALS AND METHODS

This section is organized as follows. Subsection II-A summarizes the main steps of the manual TSP to provide an overview of the clinical workflow. Subsection II-B introduces the proposed robotic workflow, designed to follow the clinical procedure. Subsection II-C describes the human impedance model adopted to characterize expert stiffness modulation. The control strategies are detailed in Subsection II-D, and the experimental setup and validation are presented in Subsection II-E.

A. Standard Clinical Workflow of TSP

In clinical practice, TSP is performed by experienced interventional cardiologists as part of standard catheter-based procedures to access the left atrium. The overall workflow includes an intravascular stage and an intracardiac stage. The intravascular stage begins with venous access, typically obtained through the right femoral vein. The access sheath, dilator, and transseptal needle are advanced intravascularly through the inferior vena cava into the right atrium under fluoroscopic guidance. This stage is well established in clinical practice and is not explicitly addressed in this work.

Once the catheter-needle assembly reaches the right atrium, the procedure enters the intracardiac stage, which is the focus of this study. As schematized in Figure 1, this stage consists of two steps. During the *Navigation Phase*, the assembly is positioned against the interatrial septum and the FO is identified using anatomical landmarks, fluoroscopic projections, and indirect mechanical cues, such as a characteristic force step. After target identification, the *Puncture Phase* begins, during which the needle is advanced to perforate the interatrial septum. Successful septal crossing is typically indicated by a sudden release of septal tenting and confirmed




Fig. 1. From left: In the navigation phase, the needle is advanced through the right atrium, approaching the interatrial septum toward the Superior Vena Cava (SVC). Once contact with the septum is achieved, the needle is retracted slightly until a visible step indicates that it is positioned over the fossa ovalis (FO). At this point, the puncture phase begins, involving first a tenting step followed by the final puncture through the FO into the left atrium. Here, \bar{x}_n represents the target location above the FO, while \hat{x}_n denotes the needle position retrieved from the EM sensor.

by visualization of the needle tip within the left atrium and/or characteristic changes in pressure waveforms. Following the puncture, the sheath can be advanced into the left atrium to enable subsequent catheter-based interventions. In this work, robotic assistance is applied exclusively during the intracardiac stage, with particular emphasis on force-regulated execution of the puncture step.

B. Proposed Robotic Workflow

The proposed framework for robotic TSP is designed to replicate the clinical workflow of the intracardiac phase, depicted in Figure 1, followed by interventional cardiologists, which involves two main phases upon reaching the right atrium:

- Phase 1 – Navigation Phase, where the FO is identified,
- Phase 2 – Puncture Phase, where the septal tissue is punctured to access the left atrium.

The proposed approach, schematized in Figure 2, aims at learning directly from experts, capturing their implicit impedance parameters adaptation strategies and force modulation skills in order to design an adaptive impedance controller for autonomous TSP. Expert demonstrations were thus collected using a leader-follower teleoperation setup, [28]: three sixth-year medical students performed the TSP procedure on different anatomies of FO while we recorded end-effector trajectories and contact forces at the slave side. A HIM was then derived based on these demonstrations to extract time-varying stiffness parameters, capturing how operators modulate force in response to needle motion and haptic feedback during the puncture phase [24].

To support these experiments, a patient-specific CT scan was obtained at Ospedale San Raffaele in Milan, Italy, collected under the <https://www.artery-project.eu> ARTERY project. Subsequently, the segmented model was imported into a digital twin environment, developed in UNITY (version 2022.3), following the methodology described in [29], providing a virtual replica of the patient’s anatomy. The digital twin provides a patient-specific anatomical reference that enables navigation and target definition within a clinically interpretable spatial context. Within this environment, clinicians can select a strategic reference point located above the interatrial septum in the superior

vena cava and a target point on the FO. This step replicates the standard clinical approach, where surgeons do not immediately target the FO but instead use an intermediate landmark to refine the trajectory. This approach allows the system to recognize the characteristic force feedback associated with transitioning from the interatrial septum to the FO, a key haptic cue that clinicians rely on during manual procedures. By integrating this additional feedback, the robotic system ensures that the puncture is performed at the correct location, reducing the risk of misplacement. Moreover, to monitor the position of the needle, which is not rigid and can deform upon contact with tissue, we integrated an ElectroMagnetic (EM) sensor to provide real-time feedback on its position as depicted in Figure 3. After detecting the FO, the system executes the puncture and automatically stops upon recognizing the puncture event to prevent excessive needle advancement.

C. Human Impedance Model (HIM)

To estimate human-like stiffness modulation during TSP, we implemented a model-based framework that reconstructs time-varying stiffness profiles from teleoperated demonstrations, as proposed by [24]. Let $\{x_{z,t}\}_{t=1}^T$ denote the position of the end-effector of the slave robots along the insertion axis (z – axis), and $\{f_{e,z,t}\}_{t=1}^T$ the corresponding force at the tool tip during the whole acquisition as depicted in Figure 2. Once position and force signals are collected, time derivatives are computed via finite differences to obtain velocity $\{\dot{x}_{z,t}\}$ and acceleration $\{\ddot{x}_{z,t}\}$. We model the force–motion relationship along the insertion axis as a second-order linear system:

$$f_{e,z,t} = \ddot{x}_{z,t} + b_h \dot{x}_{z,t} + k_{h,t} x_{z,t}, \quad (1)$$

where b_h , $k_{h,t}$ are the human damping and stiffness, and where we are assuming a unit mass. Notably, the damping coefficient b_h was chosen experimentally such that $b_h > 0$ consistent with values commonly used in human-in-the-loop manipulation tasks [24]. We define a sliding window of size $2L + 1$ centered at time t , and apply regularized ridge regression to estimate the stiffness from the measured dynamics:

$$k_{h,t} = \frac{\sum_{j=t-L}^{t+L} x_{z,j} (f_{e,z,j} - \ddot{x}_{z,j} - b_h \dot{x}_{z,j})}{\sum_{j=t-L}^{t+L} x_{z,j}^2 + \lambda}, \quad (2)$$

where λ is a regularization term to improve numerical robustness and penalize high values of $k_{h,t}$. This procedure yields a stiffness profile $\{k_{h,t}(t)\}$ aligned with the puncture phase of each demonstration, and paired with the corresponding force measurements $\{f_{e,t,z}(t)\}$.

To encode the relationship between contact force and the corresponding stiffness applied by the operator, we modeled the joint distribution of the dataset $\mathcal{D} = \{(f_{e,z,i}, k_{h,i})\}_{i=1}^T$ using a Gaussian Mixture Model (GMM) [26]. This probabilistic representation captures the multimodal nature of the demonstrated behavior and enables smooth generalization across different force conditions. First, we defined $\xi = [f_{e,z}, k_h]^T$ and model its distribution as a weighted sum of M Gaussian components:

$$p(\xi) = \sum_{n=1}^M \pi_n \mathcal{N}(\xi | \mu_n, \Sigma_n), \quad (3)$$

where π_n is the prior probability of the n -th component. Each component is characterized by its mean vector and full covariance matrix:

$$\boldsymbol{\mu}_n = \begin{bmatrix} \mu_n^f \\ \mu_n^k \end{bmatrix}, \quad \boldsymbol{\Sigma}_n = \begin{bmatrix} \Sigma_n^{ff} & \Sigma_n^{fk} \\ \Sigma_n^{kf} & \Sigma_n^{kk} \end{bmatrix} \in \mathbb{R}^{2 \times 2}. \quad (4)$$

Here, Σ_n^{ff} and Σ_n^{kk} denote the variances of the force and stiffness, respectively, while $\Sigma_n^{kf} = (\Sigma_n^{fk})^\top$ represents their cross-covariance. The model parameters $\{\pi_n, \boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n\}_{n=1}^M$ are learned from demonstration data using the Expectation-Maximization algorithm with k-means initialization and regularized covariance estimation. At runtime, the desired stiffness for a given force measurement f^* is obtained via GMR, by conditioning the joint GMM on the input:

$$\begin{aligned} \hat{k}(f^*) &= \mathbb{E}[k \mid f = f^*] = \\ &= \sum_{n=1}^M h_n(f^*) \left(\mu_n^k + \Sigma_n^{kf} (\Sigma_n^{ff})^{-1} (f^* - \mu_n^f) \right), \end{aligned} \quad (5)$$

where $h_n(f^*)$ are the posterior probabilities of each component, and Σ_n^{kf} and Σ_n^{ff} are the cross- and auto-covariance blocks of the n -th component.

D. Control Strategy

Algorithm 1 Phase-dependent adaptive impedance control for TSP

```

1: state ← NAVIGATION
2: while task running do
3:   if state = NAVIGATION then
4:     K ← Kfixed Position Control + PID correction
5:     if  $\frac{F_{\max,2} - F_{\max,1}}{F_{\max,1}} < TH_f$  ▷ FO Detected Eq. (10)
6:       state ← PUNCTURE
7:     end if
8:   else if state = PUNCTURE then
9:     K ←  $\hat{k}(f(t))$  ▷ AIC strategy
10:    if  $\frac{F_{\max,2} - F_{\max,1}}{F_{\max,1}} < TH_p$  ▷ Eq. (14) then
11:      Stop motion
12:      state ← END
13:    end if
14:  end if
15: end while
    
```

The proposed control strategy integrates a Cartesian impedance controller with position correction based on EM tracking for accurate needle tip localization, a force-based event detection mechanism, and an AIC strategy. The overall control framework for TSP consists of two main phases, as illustrated in Figure 2. During *Phase 1: "Navigation Phase"*, the robot operates under a Cartesian impedance control scheme with fixed impedance parameters, resulting in position-dominated behavior. A closed-loop position controller based on a Proportional–Integral–Derivative (PID) regulator computes corrective motion commands to track a predefined anatomical target. EM tracking is used to continuously estimate the needle tip position and correct deviations due to needle flexibility. During this phase, force measurements

are monitored exclusively for event detection and are not used to regulate motion. In particular, navigation is primarily driven by position control toward the predefined target, while the observation of a characteristic force drop is used as a confirmation that the catheter tip has reached the more compliant FO region, triggering the transition to the puncture phase. Upon FO detection, the control scheme switches to AIC with a force reference (blue) for *Phase 2: "Puncture Phase"*, ensuring a controlled puncture while preventing excessive needle advancement and replicating human skills.

The control framework for TSP is based on an Adaptive Cartesian Impedance Control specifically designed to manage the dynamic interaction between the needle attached to the robotic manipulator and the compliant cardiac tissues during the TSP. For a serial robotic manipulator, the interaction with the environment is modeled as a mass-spring-damper system:

$$\mathbf{f}_e(t) = \mathbf{M}\ddot{\tilde{\mathbf{x}}} + \mathbf{D}_c(t)\dot{\tilde{\mathbf{x}}} + \mathbf{K}_c(t)\tilde{\mathbf{x}} \quad (6)$$

where $\tilde{\mathbf{x}} = \bar{\mathbf{x}}_r - \hat{\mathbf{x}}_r$. $\bar{\mathbf{x}}_r$ and $\hat{\mathbf{x}}_r \in \mathbb{R}^3$ are the desired and measured current robot position in Cartesian coordinates of the robotic system. $\dot{\tilde{\mathbf{x}}}$ and $\ddot{\tilde{\mathbf{x}}} \in \mathbb{R}^3$ are the velocity and acceleration errors. $\mathbf{M}, \mathbf{D}_c(t), \mathbf{K}_c(t) \in \mathbb{R}^{3 \times 3}$ are the positive-definite inertia, damping, and stiffness matrices, respectively. Subsequently, for a 7-degree-of-freedom manipulator, the torque $\boldsymbol{\tau}$ generated by the impedance controller for task execution is computed as:

$$\boldsymbol{\tau}_{imp}(t) = \mathbf{B}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} + \mathbf{g}(\mathbf{q}) + \mathbf{J}^T(\mathbf{q})\mathbf{f}_e(t), \quad (7)$$

where $\boldsymbol{\tau}_{imp}(t) \in \mathbb{R}^7$ is the vector of joint torques, $\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}} \in \mathbb{R}^7$ are the position, velocity, and acceleration of the joints, respectively. $\mathbf{B}(\mathbf{q}) \in \mathbb{R}^{7 \times 7}$ is the inertia matrix, $\mathbf{C}(\mathbf{q}, \dot{\mathbf{q}}) \in \mathbb{R}^{7 \times 7}$ is the Coriolis and centrifugal matrix, $\mathbf{g}(\mathbf{q}) \in \mathbb{R}^7$ is the gravity vector, and $\mathbf{f}_e(t) \in \mathbb{R}^3$ represents the external force vector. $\mathbf{J} \in \mathbb{R}^{3 \times 7}$ represents the Jacobian matrix.

1) *Navigation Phase:* During the navigation phase, we employ the impedance controller in Eq. (7) with a fixed-impedance strategy to provide the high stiffness needed to maintain continuous tissue contact and ensure precise position control. Under these conditions, Eq. (6) can be rewritten as

$$\mathbf{f}_e(t) = \mathbf{M}\ddot{\tilde{\mathbf{x}}} + \mathbf{D}_c\dot{\tilde{\mathbf{x}}} + \mathbf{K}_c\tilde{\mathbf{x}}, \quad (8)$$

where the damping and stiffness matrices are kept constant, i.e., $\mathbf{K}_c(t) = \mathbf{K}_c$ and $\mathbf{D}_c(t) = \mathbf{D}_c$. Accordingly, the total torque command supplied to the robot is given by $\boldsymbol{\tau}(t) = \boldsymbol{\tau}_{imp}(t)$. Indeed, to ensure accurate anatomical navigation, the system incorporates a preoperative planning phase in which a target point above the FO is defined in the virtual environment, which the robot reaches before initiating FO detection. Specifically, a PID controller was implemented to compensate for the flexibility of the transseptal puncture kit and its interaction with surrounding tissues, which can cause deformations of the needle itself. To correct for these effects, an EM sensor is placed directly at the tip of the dilator (see Figure 3), enabling real-time tracking of the actual needle tip position \mathbf{x}_n (see Figure 1). The position error, \mathbf{e} , is thus computed at the needle tip, ensuring that the control loop actively corrects not only the robot's positioning but also any

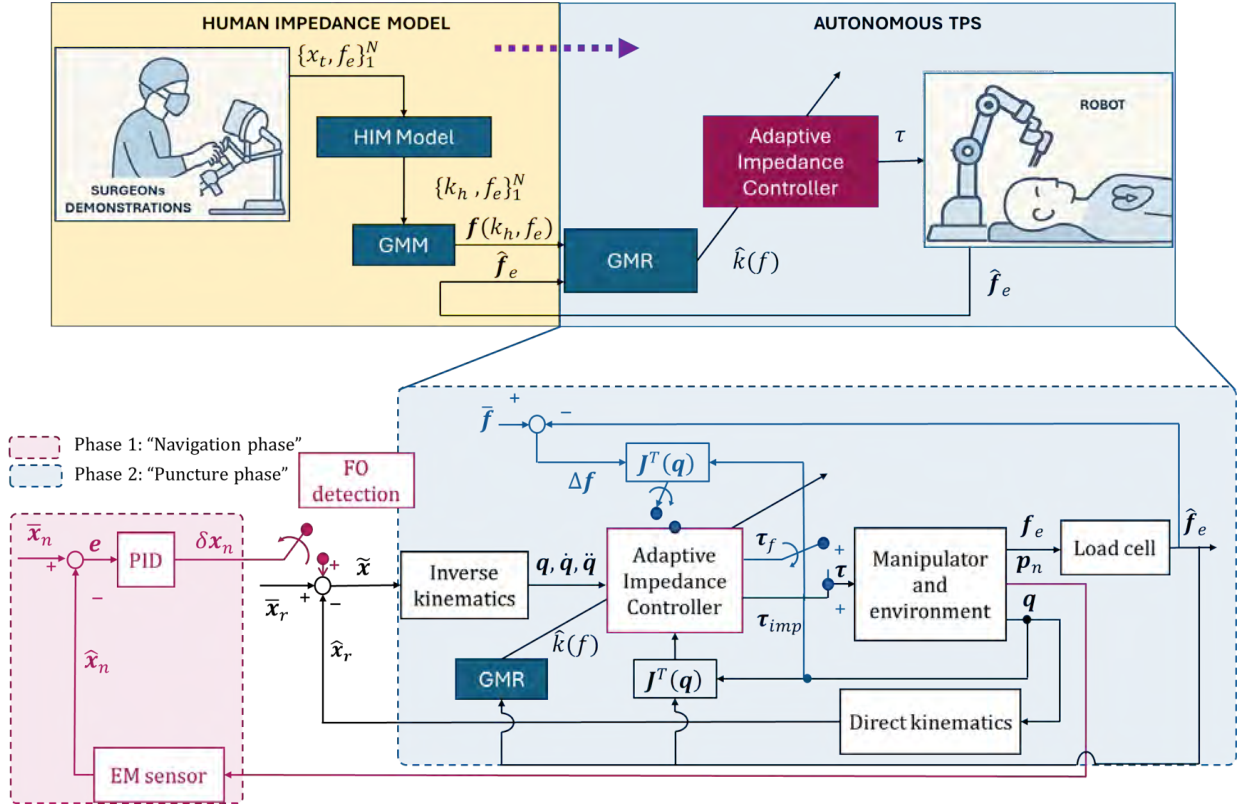


Fig. 2. From the top: Overview of the proposed teleoperated-to-autonomous transfer pipeline. The framework comprises two steps: learning a HIM from teleoperated demonstrations (left), and deploying an AIC strategy in autonomous execution (right). During learning, needle position $x_i(t)$ and contact force $f_e(t)$ are recorded and processed to extract the time-varying stiffness $k_h(t)$ via the HIM model (Eq. (2)). The resulting dataset $\{k_h(t), f_e(t)\}$ trains a Gaussian Mixture Model (GMM) that captures the relationship between contact force and stiffness. At runtime, Gaussian Mixture Regression (GMR) estimates the target stiffness \hat{k} from the measured force $f_e(t)$, enabling the AIC to compute compliant torque commands τ . In the bottom: Block diagram of the control architecture. *Phase 1 (Navigation)* uses a PID controller with EM feedback to reach the initial FO position and an Impedance Controller with fixed Impedance. After FO detection, the system switches to *Phase 2 (Puncture)*, where AIC regulates the contact force and stops motion upon puncture detection. *Note: for clarity, time dependencies are omitted in the figure.*

deformations of the needle. This allows for accurate navigation and alignment with the target region, even in presence of tissue-induced forces. The PID control law used to minimize this position error is defined as:

$$\delta x_n = K_P e + K_I \int e dt + K_D \frac{de}{dt}, \quad (9)$$

where $e = \bar{x}_n - \hat{x}_n \in \mathbb{R}^3$ is the error between the desired, \bar{x}_n , and measured needle tip position, \hat{x}_n , retrieved by the EM tracking system and transformed into the robot reference frame. K_P , K_I , and $K_D \in \mathbb{R}^{3 \times 3}$ are diagonal matrix empirically tuned. To ensure compensation for both the robot positioning error and the needle deformation, the computed correction δx_n is added to the desired end-effector position before computing the inverse kinematics.

2) *Fossa Ovalis Detection Strategy*: Instead of relying solely on positional accuracy, the system identifies the FO through force feedback. Once the catheter reaches the target reference point above the FO, the needle is gradually retracted while force variations are monitored. The FO is detected when a force drop is observed, indicating a transition from the thicker interatrial septum to the thinner FO tissue. The force-

based detection follows a threshold approach:

$$\delta F = \frac{F_{\max,2} - F_{\max,1}}{F_{\max,1}} < TH_f, \quad (10)$$

where $F_{\max,1}$ and $F_{\max,2}$ are the maximum force values measured within two consecutive windows of 100 samples each. The threshold for FO detection is set at $TH_f = 0.07$ (7%), as determined from experimental trials.

3) *Adaptive Impedance Control (AIC)*: Once the FO is detected, the *Phase 1: "navigation phase"* ends and the *Phase 2: "puncture phase"* begins. A desired force reference is then commanded to the control loop, as shown in Figure 2. In this phase, we implement the AIC strategy in which the stiffness is modulated in real time based on the force perceived. The variation of stiffness is driven by the GMR-based estimate $\hat{k}(f)$, previously introduced in Eq. (5), and scaled to map the human-like stiffness profile into a feasible range for the robot. To ensure consistent and critically damped behavior, the damping coefficient is modulated concurrently with stiffness according to the following relationship [30]: $D(t) = 2\xi\sqrt{MK}(t)$. Thus, changes in the damping matrix are directly dependent on variations in the stiffness matrix. In our framework, two key assumptions are made: first, the mass matrix M is considered constant [31]; secondly, both the stiffness matrix K and the

damping matrix D are assumed to be diagonal. Moreover, this study focuses on forces acting along the vertical (z - axis) of the end-effector, as they play a dominant role in TSP procedures. Although stiffness estimation is performed only along the z -axis, the same value $\hat{k}_r(t)$ is applied uniformly across all Cartesian directions, yielding isotropic impedance behavior. This choice is motivated by the dominance of axial interaction forces during TSP. This formulation allows the impedance control law to be simplified along the z -axis as:

$$f_{\text{ext},z}(t) = m_z \ddot{\tilde{z}} + d_z(t) \dot{\tilde{z}} + \hat{k}_z(t) \tilde{z}, \quad (11)$$

where \tilde{z} , $\dot{\tilde{z}}$, and $\ddot{\tilde{z}}$ represent the position, velocity, and acceleration errors along the insertion direction, and $f_{\text{ext},z}$ is the corresponding external force measured at the tool tip. The resulting joint torques are computed by projecting the adaptive Cartesian impedance forces into joint space using the Jacobian transpose. Hence, the τ_{imp} joint torque vector is defined as in Eq. 7 where $\mathbf{f}_e(t) = [f_{\text{ext},x}, f_{\text{ext},y}, f_{\text{ext},z}]^\top$, with $f_{\text{ext},z}$ defined in Eq. (11). The total torque vector for the puncture phase is then computed as:

$$\boldsymbol{\tau}(t) = \boldsymbol{\tau}_{\text{imp}}(t) + \boldsymbol{\tau}_f(t), \quad \text{with} \quad \boldsymbol{\tau}_f = \mathbf{J}^T(\mathbf{q})\Delta\mathbf{f}(t), \quad (12)$$

where $\Delta\mathbf{f}(t) \in \mathbb{R}^3$ is the force error between the desired puncture force vector $\mathbf{f}(t) \in \mathbb{R}^3$ and the measured interaction force $\hat{\mathbf{f}}_e(t) \in \mathbb{R}^3$. The desired force $\bar{\mathbf{f}}(t)$ is incrementally increased by 0.1 N at each control loop iteration, up to a maximum of 6 N, or until puncture is detected via a sudden drop in force feedback:

$$\delta F = \frac{F_{\text{max},2} - F_{\text{max},1}}{F_{\text{max},1}} < TH_p, \quad (13)$$

where $F_{\text{max},1}$ and $F_{\text{max},2}$ are defined as in Eq.10. This method ensures robust puncture detection through the threshold $TH_p = 0.08$ (8%), while mitigating the effect of high-frequency noise. Upon detection, the robot halts immediately to prevent over-penetration. To ensure the stability of the system under adaptive impedance control, we applied a passivity-based filter inspired by [32], and [30]. The evolution of the stiffness is governed by:

$$k_r(t) = k_0 + \gamma(t) \cdot (k_{\text{max}} - k_0), \quad (14)$$

where k_0 and k_{max} define the permissible stiffness range, and $\gamma(t) \in [0, 1]$ is a scheduling variable that modulates stiffness over time. The full formulation and implementation details are reported in the Appendix.

E. Experimental Validation

1) *Hardware Components:* The experimental setup is shown in Figure 3. It mainly includes:

- The Abbott transseptal kit (Abbott Vascular, Santa Clara, California) comprising a bent Brockenbrough needle, dilator, and sheath, trimmed in the last 150 mm.
- A 7-DoFs serial robotic manipulator (LWR 4+, KUKA, Germany) adopted for the transseptal kit actuation.
- An EM tracking system for position monitoring composed of the field generator and two 6-DoF EM sensors (Aurora, NDI, Canada) attached to the tip and the base of

the transseptal kit, providing continuous measurements at 40 Hz with an accuracy of 0.80 mm.

- An anatomical phantom derived from patient-specific CT scans simulating the interatrial septum width, and mitral valve position within the heart chamber.
- Silicone FOs representing different anatomical conditions (detailed description of the silicone FOs can be found in Section *FO Models*).
- A 7-DOFs haptic device Sigma.7 (Force Dimension, Switzerland), equipped with a switch pedal, was employed as the leader side haptic interface.
- A 6-axis force/torque sensor (M3815C-SN2537, SRI, China), with a resolution 0.03 N, mounted to the robot end-effector to measure the real-time contact force with a frequency of 200 Hz.

Interaction forces are measured using a 6-axis force/torque sensor mounted at the robot end-effector. The measured forces correspond to those transmitted through the needle-tool assembly rather than to a direct measurement at the needle tip. The needle and puncture kit are not assumed to be perfectly rigid; however, the needle is shortened to approximately 15 cm and supported along most of its length by the puncture kit/sheath, which substantially reduces the effective unsupported length and limits bending during insertion. All experiments are conducted at low insertion speeds under impedance control. Under these conditions, inertial effects are negligible compared to interaction forces, allowing the task to be treated as quasi-static. Although the needle is slightly curved and compliant, its curvature primarily introduces bending moments and minor lateral force components, without compromising the transmission of the axial force required for tissue penetration. The Ubuntu 16.04 desktop utilizes the Kinetic version of the Robot Operating System (ROS) to ensure efficient data exchange¹. At the same time, the robot is controlled using the FastResearchInterface (FRI) library provided by KUKA. PC-1 connects to the serial robot via the User Datagram Protocol (UDP) and manages its control through the FRI library. The control frequency was set to 100 Hz.

Figure 3 illustrates all the reference frames associated with the components of the system setup. Accurate spatial alignment between these components is crucial for executing the procedure correctly. To achieve this, specific calibration steps are performed:

- $T_U^A \in \mathbb{R}^{4 \times 4}$ maps the digital twin environment $\{U\}$ to the EM tracking system $\{A\}$, ensuring that the virtual and physical setups are aligned. A paired-point registration method [33] was used to calibrate the two reference frames by collecting six corresponding points, resulting in the transformation matrix T_U^A .
- $T_A^K \in \mathbb{R}^{4 \times 4}$ maps the EM tracking system frame $\{A\}$ to the robot base frame $\{K\}$, enabling the robotic manipulator to follow the planned navigation and puncture paths precisely. Specifically, the hand-eye calibration [34] was performed to calculate T_A^K .

Consequently, the transformation matrix $T_U^K \in \mathbb{R}^{4 \times 4}$ between the digital twin environment frame $\{U\}$ and the robotic system

¹<https://github.com/CentroEPIaggio/kuka-lwr>

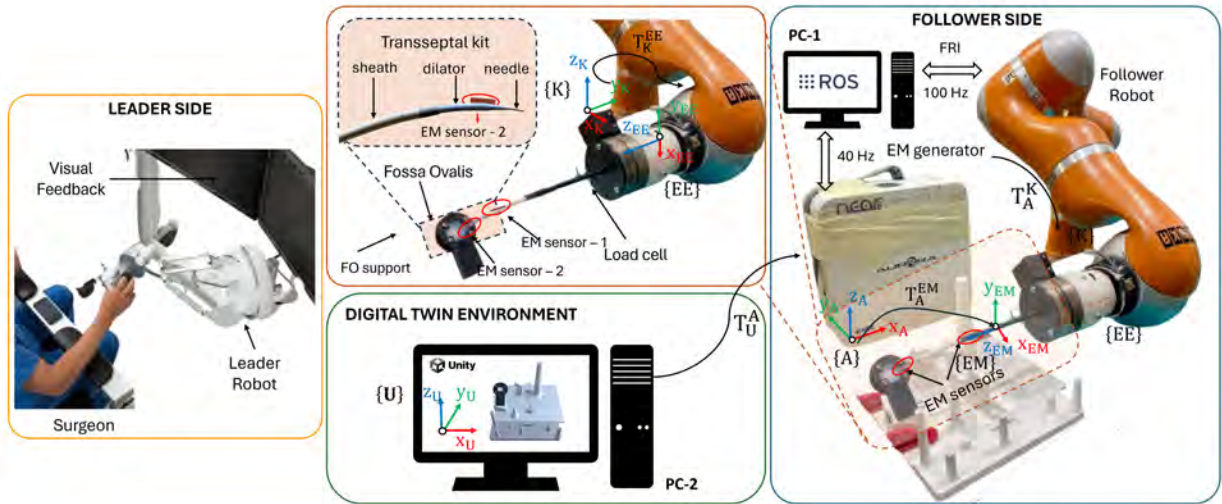


Fig. 3. Experimental Setup for Teleoperated-to-Autonomous Transseptal Puncture The system integrates teleoperation, real-time tracking, and digital twin simulation. Leader Side (yellow): a surgeon operates a Sigma.7 haptic interface with visual feedback. Follower Side (blue): a 7-DOF KUKA robot executes the puncture using a transseptal kit (sheath, dilator, needle) with two embedded EM sensors. A 6-axis load cell measures contact forces. The EM tracking system (Aurora) provides 40Hz pose data, synchronized with robot control at 100Hz via PC-1 (ROS). Digital Twin Environment (green): a Unity-based virtual model replicates patient anatomy and is aligned with the physical workspace through rigid transformations T_U^A and T_A^K . All coordinate frames—robot base $\{K\}$, end-effector $\{EE\}$, EM tracker $\{A\}$, and virtual model $\{U\}$ —are spatially calibrated for accurate execution and force-based FO detection. The setup enables seamless transfer from human demonstrations to autonomous puncture control.

frame $\{K\}$ can be obtained as

$$T_U^K = T_U^A T_A^K \quad (15)$$

2) *FO Models*: TSPs were performed on custom silicone FO models, produced using Dragon SkinTM 10 medium silicone, validated in [17]. FOs presented an oval shape of $12.1 \text{ mm} \times 14.1 \text{ mm}$. Three anatomies with different thicknesses were fabricated using purposely designed 3D-printed molds, in which silicone was poured: a standard FO [35], a thicker FO [36], and a floppy FO [37]. These three anatomies were used to collect teleoperated demonstrations and derive the HIM. To further evaluate robustness in unseen conditions, two additional anatomies were created outside the anatomical limits reported in literature: a hyper-thick FO and a hyper-floppy FO. These extreme cases were not included in the demonstration set and were instead used exclusively for testing the generalization capabilities of the proposed approach. The thickness of each FO model representing different anatomies is reported in Table I.

TABLE I
SUMMARY OF THE FO PHANTOM ANATOMIES AND THEIR THICKNESSES.

Anatomy	Thickness [mm]
Standard FO	1.04
Thick FO	1.89
Floppy FO	0.60
Hyper-thick FO	2.90
Hyper-floppy FO	0.20

3) *Parameters tuning*: Human stiffness was estimated from the collected teleoperated demonstrations using the dynamic model in Eq. (2), applied over a sliding window of 11 samples ($L=5$, stride 1). Motion and force signals were low-pass filtered using a second-order Butterworth filter with a cutoff frequency of 20 Hz. The model parameters were fixed: mass

$m = 1 \text{ kg}$, damping $b_h = 0.1$, and regularization term $\lambda = 0.005$. Stiffness–force pairs (F, k_h) were collected from teleoperated demonstrations, restricted to the puncture phase (from initial contact to breakthrough), and used to train a GMM with $M = 5$ full-covariance components, as defined in Eq. (3). The resulting probabilistic model was queried via GMR to infer a time-varying human stiffness estimate $\hat{k}_h(t)$ from real-time axial force. To make the predicted stiffness compatible with the robot’s impedance controller, we applied a linear scaling from the human range ($k_{\min}^h = 1 \text{ N/m}$, $k_{\max}^h = 25 \text{ N/m}$) to the robot range ($k_{\min}^r = 800 \text{ N/m}$, $k_{\max}^r = 4000 \text{ N/m}$), as given by:

$$\hat{k}_r(t) = k_{\min}^r + \frac{\hat{k}_h(t) - k_{\min}^h}{k_{\max}^h - k_{\min}^h} \cdot (k_{\max}^r - k_{\min}^r). \quad (16)$$

The resulting value $\hat{k}_r(t)$ was then passed through a passivity filter based on [30] to enforce stable interaction. The filter used a convergence gain $\alpha = 0.01$ and a slope-limiting coefficient $\beta = 100$. For benchmarking, we implemented two fixed-stiffness baselines: a low-stiffness condition (LS) with $k_r = 800 \text{ N/m}$ and a high-stiffness condition (HS) with $k_r = 3000 \text{ N/m}$. For the navigation phase, a high fixed stiffness was selected in line with values reported in the literature ([38]), as higher stiffness improves precision and trajectory tracking. The chosen value was $k_{\text{nav}} = 3000 \text{ N/m}$. Cartesian motion was controlled via a PID loop referenced at the needle tip, with diagonal gains set empirically to $K_P = \text{diag}(0.05, 0.05, 0.05) \text{ N/m}$, $K_I = \text{diag}(0.1, 0.1, 0.1) \text{ N/(m}\cdot\text{s)}$, and $K_D = \text{diag}(0.1, 0.1, 0.1) \text{ N}\cdot\text{s/m}$. The controller was manually tuned to ensure stability and maintain tracking errors below 2 mm throughout navigation.

4) *Experimental Protocol*: To assess the effectiveness of the proposed LfD AIC-based autonomous TSP, we designed a three-phase experimental pipeline.

- 1) Teleoperated demonstrations were collected using a leader–follower setup, in which three sixth-year medical students (2 males, 1 female) performed 9 punctures each on three different FO phantoms, representing standard, floppy, and thicker anatomical conditions, for a total of 27 demonstrations².
- 2) Based on these data, the HIM was derived and encoded with GMM, enabling real-time stiffness inference via GMR.
- 3) The learned model was then integrated into an AIC strategy and deployed on a robotic platform to perform the complete TSP workflow autonomously.

During experimental validation, we assessed the accuracy of the learned HIM, evaluated the system’s navigation performance, and analyzed puncture outcomes across varying anatomical conditions. The proposed adaptive controller was compared against two fixed-stiffness baselines: a Low-Stiffness (LS) controller, resulting in force-dominated interaction behavior, and a High-Stiffness (HS) controller, resulting in position-dominated interaction behavior, to quantify improvements in safety and adaptability. All controllers were tested across the five anatomically distinct FO phantoms: standard, floppy, thick, hyper-floppy, and hyper-thick (see Table I). For each combination of FO type and control strategy, five autonomous puncture trials were performed, resulting in a total of 75 independent executions.

5) *Performance Metrics*: The performance indices evaluated during the tests were selected as the most representative ones to assess the safety of TSP, specifically:

- *Needle Puncture Force (NPF)*: defined as the force measured at the moment of FO puncture, corresponding to the force required to perform the puncture (subfigure (A) in Figure 6).
- *Maximum Tenting Distance (TD)*: it quantifies the deformation of the FO before puncturing due to the pressure applied by the transeptal kit. Excessive displacement might bring the needle closer to cardiac walls, increasing the risk of puncturing other structures and severe complications. It is defined as the Euclidian distance between the initial position and the position of puncture, (subfigure (A) in Figure 6):

$$TD = \|\mathbf{p}_{puncture} - \mathbf{p}_{start}\|. \quad (17)$$

- *Needle Stopping Space (NSS)*: promptly halting the needle after puncture is critical to minimize the risk of damaging surrounding tissues. It is defined as the Euclidian distance between the final position after detecting the puncture and the puncture position (subfigure (A) in Figure 6):

$$NSS = \|\mathbf{p}_{stop} - \mathbf{p}_{puncture}\|. \quad (18)$$

Afterward, a statistical analysis was performed on metrics; mean and standard deviation (SD) were computed. Upon verifying the normality of the data distribution through a

²All participants provided informed consent before participation, and the experimental protocol was approved by the Ethics Committee of Politecnico di Milano, Italy (No. 45/2023).




Fig. 4. Experimental setup for *ex vivo* porcine tissue tests. The same experimental setup adopted for the silicone experiments was used, with the silicone membrane replaced by *ex vivo* porcine tissue.

Kolmogorov-Smirnov test, a parametric T-test was applied at a significance level of $\alpha = 0.05$. The analysis was carried out in MATLAB (R2021b).

6) *Ex Vivo Tissue Experiments*: *Ex vivo* porcine pericardial tissue was chosen to test the robustness of the proposed control strategy due to its mechanical characteristics. Cardiovascular soft tissues such as FO and pericardium are indeed characterized by a fibrous connective structure rich in collagen and have been shown to exhibit complex, nonlinear, and anisotropic mechanical behavior under puncture and traction loads [39], [40]. The experiments focused exclusively on the puncture phase: the tool was pre-positioned in contact with the tissue and advanced along the insertion axis under three control strategies (HS, LS, AIC). A single porcine pericardial membrane was divided into three spatially separated regions. Tissue samples were stored and hydrated in sterile saline (0.9% NaCl) at 4 °C and tested within 48 h of collection. For each tissue sample, four puncture trials per control strategy were performed. To avoid the influence of local damage, stress relaxation, or preconditioning effects, consecutive punctures were executed at distinct locations within each region, ensuring that each trial was carried out on previously unpunctured tissue. The *ex vivo* tissue tests were performed using the same experimental setup adopted for the silicone FOs, as shown in Figure 4.

III. RESULTS AND DISCUSSION

Figure 5 shows the estimated human stiffness profiles across both training and unseen data. In both cases, the GMR-based prediction (black dashed) successfully tracks the reference stiffness $k_{HIM}(t)$ (magenta), rising in response to increased contact force and decreasing after puncture. The results confirm that the learned probabilistic model captures the temporal trend of human-like impedance modulation and generalizes to previously unseen force trajectories. Minor discrepancies around force peaks are attributed to input patterns not present

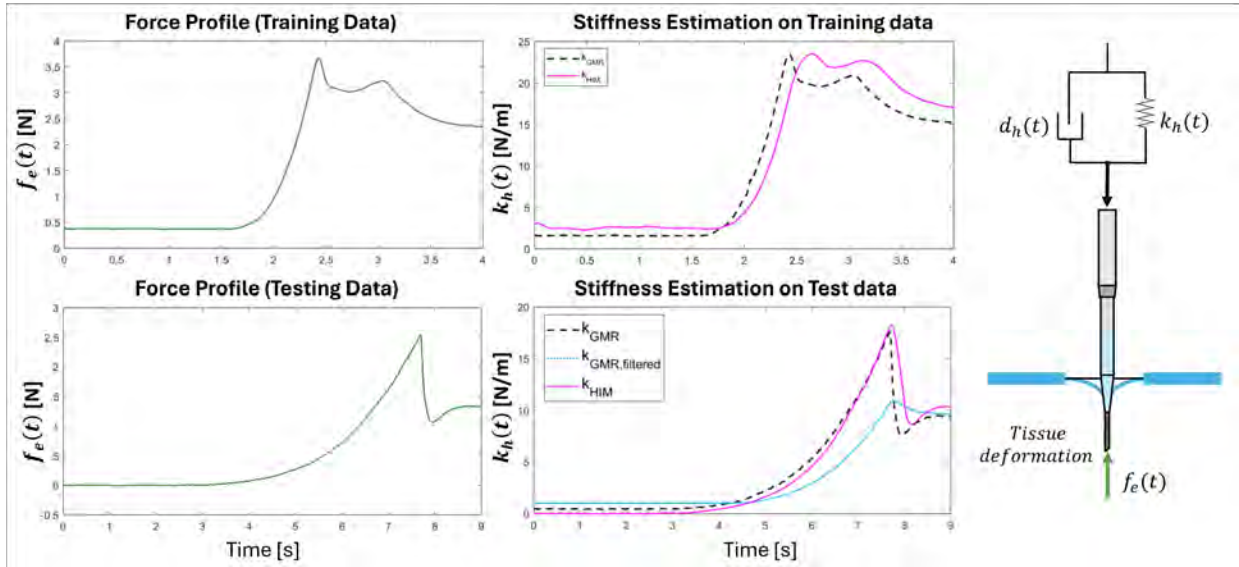


Fig. 5. Estimation of human stiffness from teleoperated demonstrations. Force and stiffness profiles extracted from training (top) and testing (bottom) datasets. Left: measured axial contact force $f_e(t)$. Right: human stiffness $k_h(t)$ estimated from the HIM model (magenta) and predicted via GMM/GMR (black dashed), with additional filtered output (blue dotted) shown in the test case. The schematic on the right represents the target estimation: a time-varying human stiffness $k_h(t)$ and damping $d_h(t)$ acting on the needle as it deforms the tissue through axial interaction.

in the training data. In the test scenario, the filtered estimate (blue dotted) is also shown, demonstrating how the passivity filter smooths high-frequency variations while preserving key trends. This filtered profile is then used in the AIC to ensure stable and compliant interaction during autonomous execution. To assess the system navigation performance, starting from a predefined “home” position located above the interatrial septum and not in contact with the tissue, the system executed a position-controlled trajectory toward a target point defined in a digital twin environment. This waypoint was selected to correspond to the typical initial contact location with the FO, mirroring clinical practice. During this navigation phase, a PID controller guided the robot using real-time feedback from an EM-tracked needle tip.

Across all trials, the Cartesian positioning error at the target location had a global median error of 1.6 mm, remaining within the predefined accuracy threshold of 2 mm. Upon reaching the waypoint contact region, the needle was retracted along its insertion axis while monitoring axial force. A characteristic drop in force reflecting the transition from the thicker septal region to the thinner, more compliant FO, was used to detect the puncture site autonomously. This process replicates the manual technique used by clinicians to identify the FO before puncture.

Following FO identification, the system transitioned into the *puncture phase*, during which the proposed AIC method was evaluated against two fixed-impedance baselines: the HS impedance controller and an LS impedance controller. Limiting the axial force during TSP is the most clinically relevant safety target.

Column (B) of Figure 6 reports the NPF obtained in standard silicone FO anatomies. In this condition, the average NPF was 3.1 N and did not differ significantly among AIC, LS, and HS ($p > 0.05$), confirming that all controllers operate

within a safe force range under nominal anatomy. However, in anatomies with altered compliance, also reported in column (B), AIC showed significant advantages. In the floppy FO, AIC reduced NPF by 28% compared to LS ($p < 0.05$) and by 18% compared to HS ($p < 0.01$), while in the thick FO it achieved a 35% reduction with respect to HS ($p < 0.01$). These trends indicate that AIC effectively limits puncture forces as tissue compliance deviates from nominal conditions. The advantages of AIC become more pronounced in extreme silicone anatomies, shown in column (C) of Figure 6. In the hyper-floppy FO, AIC reduced NPF to 2.37 N, compared with 5.34 N for LS and 4.35 N for HS. Similarly, in the hyper-thick FO, AIC lowered the puncture force to 3.38 N, compared with 5.6 N for LS and 4.8 N for HS. These results confirm that the proposed AIC-based control modality consistently limits puncture forces even under extreme anatomical conditions.

In the second row of column (B) of Figure 6 are reported the TD obtained for each FO–controller combination. In nominal FO conditions (standard, floppy, and thick), TD increased slightly with AIC in the standard and floppy FO ($p < 0.05$), while remaining unchanged in the thick FO. This behavior reflects the gradual increase in effective stiffness adopted by AIC under moderate force slopes, prioritizing force reduction at the expense of limited additional tissue deformation. In contrast, column (C) shows that in extreme anatomies, AIC drastically reduced TD. In the hyper-floppy FO, TD was reduced to 11.59 mm with AIC, compared with 31.14 mm for LS and 27.9 mm for HS. A similar reduction was observed in the hyper-thick FO, where TD decreased to 11.8 mm with AIC, compared with 28.23 mm and 25.9 mm for LS and HS, respectively. In these conditions, the steeper force–compliance dynamics drive the stiffness scheduler more rapidly, enforcing earlier stiffening and effectively clamping TD.

Post-puncture behavior is quantified by the NSS reported

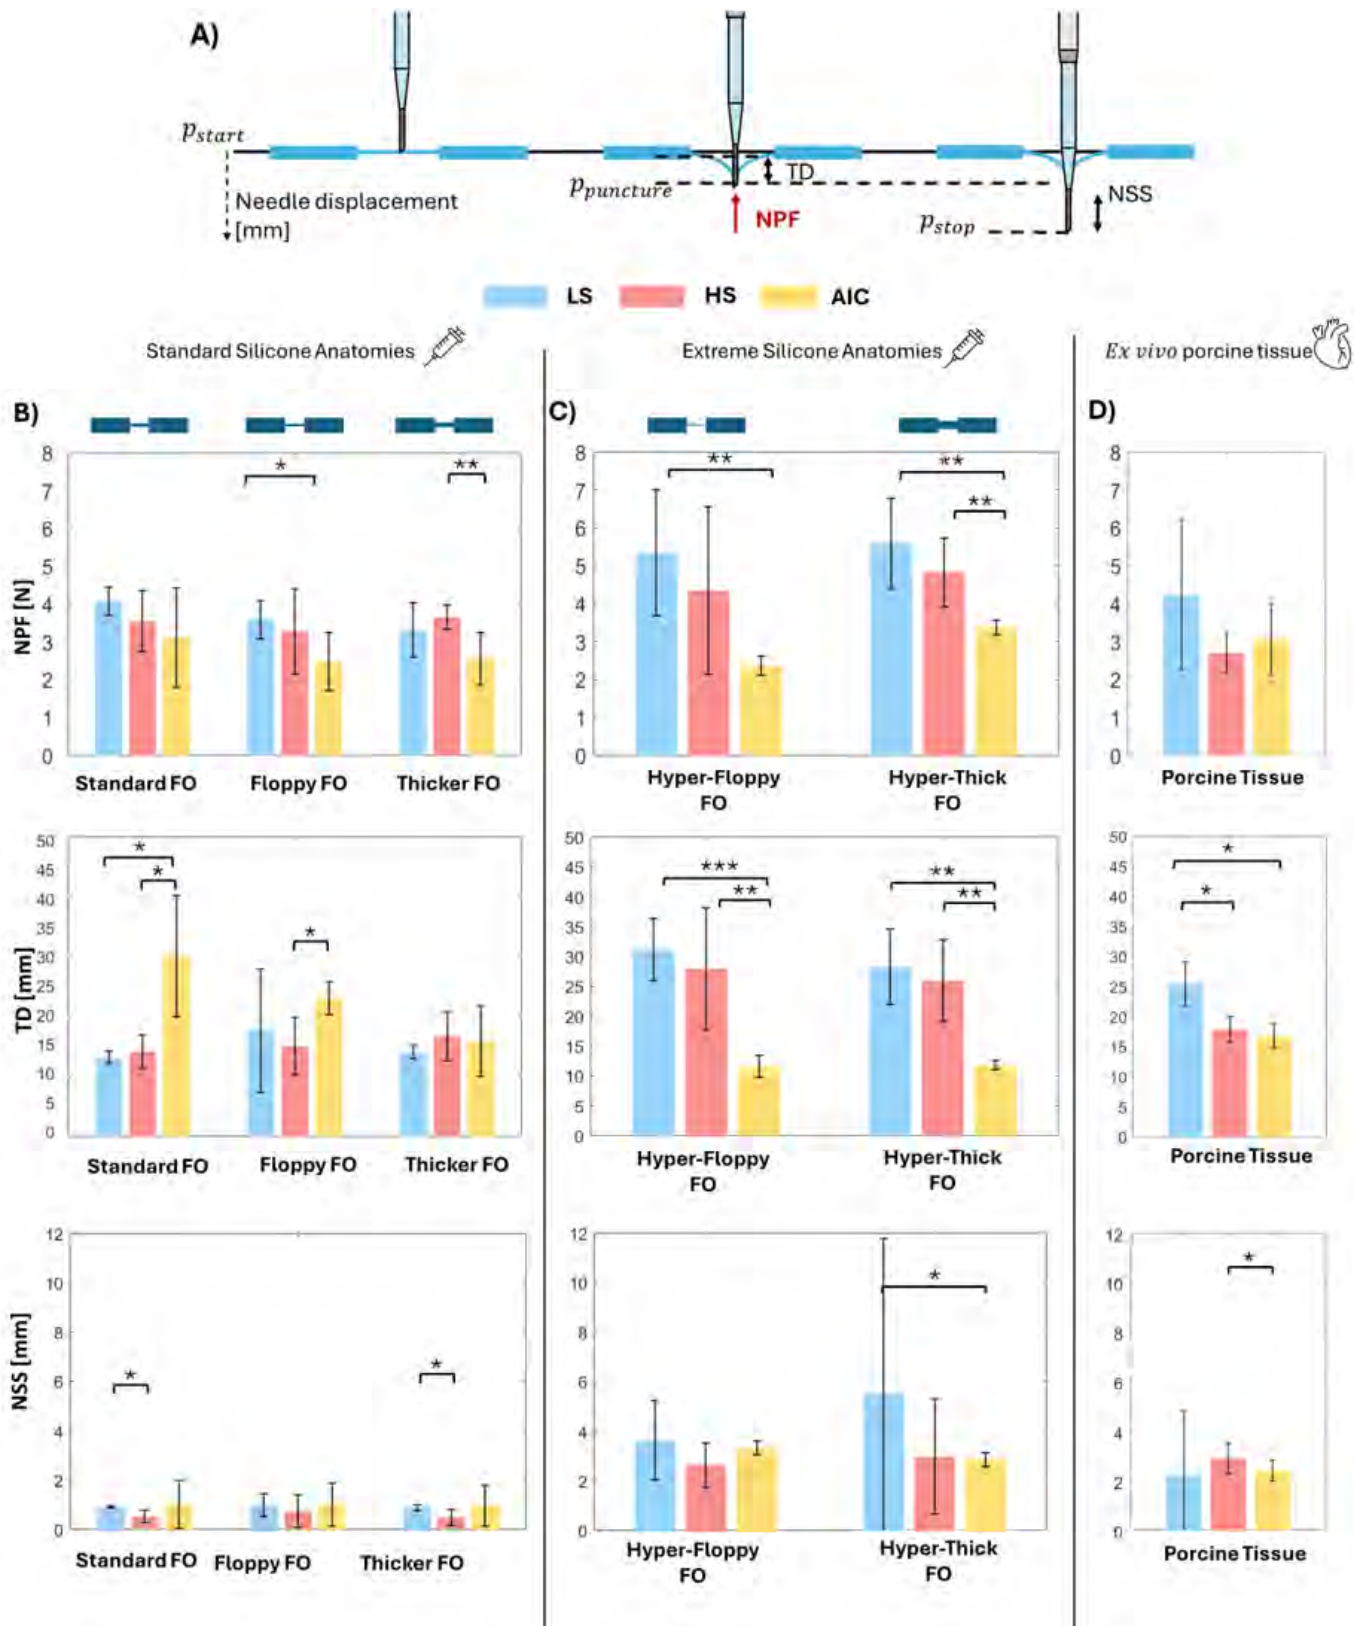


Fig. 6. Definition of puncture metrics and experimental results. (A) Schematic of the needle insertion experiment and definition of the measured quantities. The puncture event occurs at $p_{puncture}$, corresponding to the peak needle puncture force (NPF). The tenting distance (TD) is defined as the displacement from p_{start} to $p_{puncture}$, while the needle stopping space (NSS) is defined as the displacement from $p_{puncture}$ to p_{stop} . Three control strategies are compared: low stiffness (LS), high stiffness (HS), and adaptive impedance control (AIC). (B–D) Mean \pm standard deviation of NPF, TD, and NSS obtained on standard silicone anatomies, extreme silicone anatomies, and *ex vivo* porcine tissue, respectively. Statistical significance is indicated by asterisks (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$).

in the bottom row of Figure 6. In standard silicone FO anatomies (column B), all controllers exhibit limited post-puncture overshoot, with NSS remaining below approximately 2 mm. In these conditions, AIC consistently maintains NSS below 1 mm, indicating prompt stopping upon membrane breakthrough.

In extreme silicone anatomies (column C), different trends emerge depending on the FO type. In the hyper-floppy FO, NSS values remain limited for all controllers, indicating that post-puncture overshoot is intrinsically constrained by the high compliance of the material. In contrast, in the hyper-thick FO, LS exhibits a marked increase in NSS, reaching values close to 5–8 mm, while HS shows intermediate behavior. Notably, AIC consistently maintains post-puncture overshoot below approximately 4 mm across both extreme anatomies. These results indicate that, while extreme anatomical conditions may lead to larger post-puncture displacements for fixed-stiffness controllers, AIC preserves robust and controlled post-puncture behavior across all tested materials.

Figure 6 (D) shows the results obtained on *ex vivo* porcine pericardium. These results confirm the trends observed in silicone models, despite the higher variability inherent to biological samples. AIC achieves a balanced compromise between puncture force, tenting, and stopping space, whereas LS and HS controllers exhibit either excessive deformation or increased puncture force. The presence of statistically significant differences across controllers in biological tissue reinforces the relevance of adaptive impedance control for realistic surgical scenarios.

Overall, these results demonstrate that AIC provides a more robust and consistent needle insertion behavior across a wide range of tissue conditions. By simultaneously reducing puncture force, limiting pre-puncture deformation, and controlling post-puncture advancement, AIC offers clear advantages over fixed-stiffness strategies, particularly in anatomies characterized by high mechanical variability.

IV. CONCLUSIONS

This work investigates the feasibility of autonomous TSP by embedding human-like impedance modulation into a robotic control framework, under controlled experimental conditions. The proposed approach achieved promising results in terms of safety and adaptability across a range of compliant phantom anatomies and *ex vivo* tissue. By embedding human impedance modulation into the control strategy, the system effectively minimized puncture forces, controlled deformation, and reacted promptly to breakthrough events, demonstrating inherently safe behavior even under challenging anatomical conditions. Beyond TSP, this framework could serve as a foundation for anatomy-aware automation in other surgical procedures requiring controlled puncture or insertion, and more generally in applications that demand adaptive modulation of interaction with patient-specific anatomical variations. To the best of our knowledge, this is the first work to investigate how human stiffness modulation can be leveraged to perform puncturing operations in a surgical scenario. By embedding such modulation into the control framework, the

proposed approach enables anatomy-aware and adaptive interaction during delicate procedures. While demonstrated here in the context of TSP, the methodology is broadly applicable to other interventional procedures involving controlled puncture or insertion, where patient-specific anatomical variability plays a critical role.

Despite these advantages, some shortcomings remain to be addressed. Future work will address this limitation by extending the learning-from-demonstration framework to model force-related impedance modulation during both navigation and puncture, enabling a unified probabilistic representation of interaction behavior across task phases. Moreover, the current implementation does not explicitly account for the total needle length or compensate for target motion. Ongoing work is focused on extending the digital twin toward dynamic representations by incorporating automated detection of cardiac anatomical structures from echocardiographic images and predictive motion estimation, enabling anticipation of tissue motion to improve accuracy and robustness in more realistic scenarios.

APPENDIX PASSIVITY PROOF

To ensure the stability of the system with variable impedance parameters, a passivity filter was introduced ([32]). Specifically, stability is ensured if the following passivity conditions hold:

$$0 \leq K - \dot{D} \quad 0 \leq 2DM^{-1}(K - \dot{D}) - \dot{K} + \ddot{D} \quad (19)$$

Assuming diagonal matrices and focusing on the z -axis, the conditions in (19) can be simplified to:

$$d_z - k_z \leq 0 \quad (20)$$

$$\dot{k}_z + \frac{2}{m_z}d_z\dot{d}_z - \dot{d}_z - \frac{2}{m_z}k_z\dot{d}_z \leq 0 \quad (21)$$

To satisfy these constraints, a rate-of-change limit on the scheduling variable $\gamma(t)$ is derived, ensuring passivity. This is expressed as:

$$\dot{\gamma} \leq \frac{2(d_0\delta k + k_0\delta d)\gamma + 2\delta k\delta d\gamma^2 + 2k_0d_0}{\delta k + 2d_0\delta d + 2\delta d^2\gamma} \triangleq h_1(\gamma), \quad (22)$$

where $\delta k = k_1 - k_0$, $d_0 = 2\xi\sqrt{k_0}$, and $\delta d = 2\xi(\sqrt{k_1} - \sqrt{k_0})$. When $\delta d \neq 0$, an alternative passivity condition also applies:

$$\dot{\gamma} \leq \frac{\delta k}{\delta d}\gamma + \frac{k_0}{\delta d} \triangleq h_2(\gamma). \quad (23)$$

To ensure that the system remains passive and stable as the desired stiffness γ increases, a passivity filter is introduced to regulate the rate of change of $\dot{\gamma}$. The passivity filter takes as input the desired $\bar{\gamma}_i$ profile and generates an output $\hat{\gamma}_i$ that guarantees system passivity:

$$\dot{\gamma} = \min(h_1(\gamma), h_2(\gamma), \beta(\bar{\gamma} - \gamma)) \quad (24)$$

Although the derivation above considers only the z -axis, in our implementation, the same filtered stiffness $k_r(t)$ is applied isotropically across all Cartesian directions. This ensures uniform impedance behavior while maintaining passivity under the scalar formulation of the filter.

REFERENCES

- [1] S. Chernova and A. L. Thomaz, *Robot learning from human teachers*. Springer Nature, 2022.
- [2] G. Russo, M. Taramasso, and F. Maisano, “Transseptal puncture: A step-by-step procedural guide,” *Card. Interv. Today*, vol. 13, no. 3, pp. 22–26, 2019.
- [3] C. Denis, G. Clerfond, A. Chalard, C. Riocreux, B. Pereira, O. Lamallem, T. Guizani, P.-A. Catalan, A. Boudias, F. Jean *et al.*, “Safety and efficacy of mini-invasive left atrial appendage closure: A propensity-score analysis,” *Canadian Journal of Cardiology*, vol. 40, no. 11, pp. 2025–2035, 2024.
- [4] M. Almendarez, R. Alvarez-Velasco, I. Pascual, A. Alperi, C. Moris, and P. Avanzas, “Transseptal puncture: review of anatomy, techniques, complications and challenges, a critical view,” *International journal of cardiology*, vol. 351, pp. 32–38, 2022.
- [5] J. C. Díaz, C. D. Niño, O. Bastidas, J. M. Aristizábal, J. E. Marín, J. M. Martínez, and M. Duque, “How to perform a transseptal puncture safely and effectively,” *Revista Colombiana de Cardiología*, vol. 30, no. 6, pp. 361–371, 2023.
- [6] G. D. Katritsis, G. C. Siontis, E. Giazitzoglou, N. Fragakis, and D. G. Katritsis, “Complications of transseptal catheterization for different cardiac procedures,” *International journal of cardiology*, vol. 168, no. 6, pp. 5352–5354, 2013.
- [7] V. Marimuthu, K. Srinivasa, S. K. Shankar Rao, N. Alur, M. C. Nanjappa, and N. Prasad, “Pneumopericardium—an unusual complication following transseptal puncture: A series of 3 cases,” *Case Reports*, vol. 4, no. 11, pp. 671–676, 2022.
- [8] S. P. Sharma, R. Nalamasu, R. Gopinathannair, C. Vasamreddy, and D. Lakkireddy, “Transseptal puncture: devices, techniques, and considerations for specific interventions,” *Current Cardiology Reports*, vol. 21, pp. 1–7, 2019.
- [9] D. Ge, H. Zhao, Y. Wang, D. Li, X. Li, and H. Ding, “Learning compliant dynamical system from human demonstrations for stable force control in unknown environments,” *Robotics and Computer-Integrated Manufacturing*, vol. 86, p. 102669, 2024.
- [10] M. J. Hagan, T. Remacle, O. P. Leary, J. Feler, E. Shaaya, R. Ali, B. Zheng, A. Bajaj, E. Traupe, M. Kraus *et al.*, “Navigation techniques in endoscopic spine surgery,” *BioMed Research International*, vol. 2022, no. 1, p. 8419739, 2022.
- [11] N. M. Alekseeva, A. G. Viller, and A. M. Romanov, “The design of endovascular robots for cardiac interventions: A survey,” *Journal of Field Robotics*, vol. 42, no. 5, pp. 1914–1962, 2025.
- [12] R. Konda, T. A. Brumfiel, Z. L. Bercu, J. A. Grossberg, and J. P. Desai, “Robotically steerable guidewires—current trends and future directions,” *Science Robotics*, vol. 10, no. 105, p. eadt7461, 2025.
- [13] N. A. Thompson, S. Shin, A. G. Kocheril, E. T. Hsiao-Weckler, and G. Krishnan, “Design and validation of a soft robotic simulator for transseptal puncture training,” *IEEE Transactions on Biomedical Engineering*, vol. 70, no. 10, pp. 3003–3014, 2023.
- [14] Y. Yan, H. Wang, H. Yu, F. Wang, J. Fang, J. Niu, and S. Guo, “Machine learning-based surgical state perception and collaborative control for a vascular interventional robot,” *IEEE Sensors Journal*, vol. 22, no. 7, pp. 7106–7118, 2022.
- [15] H.-S. Song, B.-J. Yi, J. Y. Won, and J. Woo, “Learning-based catheter and guidewire-driven autonomous vascular intervention robotic system for reduced repulsive force,” *Journal of Computational Design and Engineering*, vol. 9, no. 5, pp. 1549–1564, 2022.
- [16] S. B. Kesner and R. D. Howe, “Force control of flexible catheter robots for beating heart surgery,” in *2011 IEEE international conference on robotics and automation*. IEEE, 2011, pp. 1589–1594.
- [17] L. Carlini, A. Bicchì, J. Fu, and E. De Momi, “Towards robotic transseptal puncture: A preliminary study investigating the influence of puncture velocity in minimally invasive cardiovascular surgery,” in *2024 10th IEEE RAS/EMBS International Conference for Biomedical Robotics and Biomechanics (BioRob)*. IEEE, 2024, pp. 581–586.
- [18] 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, pp. 368–375, 2004.
- [19] 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. 3270–3282, 2003.
- [20] M. Shariif, A. Zakerimanesh, J. K. Mehr, A. Torabi, V. K. Mushahwar, and M. Tavakoli, “Impedance variation and learning strategies in human-robot interaction,” *IEEE Transactions on Cybernetics*, vol. 52, no. 7, pp. 6462–6475, 2021.
- [21] X. Yu, P. Liu, W. He, Y. Liu, Q. Chen, and L. Ding, “Human-robot variable impedance skills transfer learning based on dynamic movement primitives,” *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 6463–6470, 2022.
- [22] S. Jadav, J. Heidersberger, C. Ott, and D. Lee, “Shared autonomy via variable impedance control and virtual potential fields for encoding human demonstrations,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 15 151–15 157.
- [23] Z. Liu, S. Wang, Y. Cao, Z. Cao, L. Lin, and L. Xie, “An autonomous guidewire delivery method with distal position-based impedance control,” *IEEE Transactions on Automation Science and Engineering*, 2024.
- [24] Y. Michel, R. Rahal, C. Pacchierotti, P. R. Giordano, and D. Lee, “Bilateral teleoperation with adaptive impedance control for contact tasks,” *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5429–5436, 2021.
- [25] X. Zhai, L. Jiang, H. Wu, H. Zheng, D. Liu, X. Wu, Z. Xu, and X. Zhou, “Learning target-directed skill and variable impedance control from interacting demonstrations for robot-assisted soft tissue puncture tasks,” *IEEE Transactions on Automation Science and Engineering*, 2024.
- [26] A. Ajoudani, C. Fang, N. Tsagarakis, and A. Bicchì, “Reduced-complexity representation of the human arm active endpoint stiffness for supervisory control of remote manipulation,” *The International Journal of Robotics Research*, vol. 37, no. 1, pp. 155–167, 2018.
- [27] Z. Du, W. Wang, Z. Yan, W. Dong, and W. Wang, “Variable admittance control based on fuzzy reinforcement learning for minimally invasive surgery manipulator,” *sensors*, vol. 17, no. 4, p. 844, 2017.
- [28] J. Fu, G. Maimone, E. Iovene, J. Zhao, A. Redaelli, G. Ferrigno, and E. De Momi, “Human-inspired active compliant and passive shared control framework for robotic contact-rich tasks in medical applications,” *IEEE Transactions on Robotics*, vol. 41, pp. 2549–2568, 2025.
- [29] A. Peloso, R. Damiano, X. Zhang, A. Bicchì, E. Votta, and E. De Momi, “Imitation learning for path planning in cardiac percutaneous interventions,” *IEEE Transactions on Biomedical Engineering*, 2025.
- [30] E. Iovene, R. Monaco, J. Fu, F. Costa, G. Ferrigno, and E. De Momi, “Emg-based variable impedance control for enhanced haptic feedback in real-time material recognition,” *IEEE Transactions on Haptics*, 2025.
- [31] K. Kronander and A. Billard, “Stability considerations for variable impedance control,” *IEEE Transactions on Robotics*, vol. 32, no. 5, pp. 1298–1305, 2016.
- [32] M. Bednarczyk, H. Omran, and B. Bayle, “Passivity filter for variable impedance control,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 7159–7164.
- [33] X. Zhang, M. C. Palumbo, F. Perico, M. Magro, A. Fortuna, T. Magni, E. Votta, A. Segato, and E. De Momi, “Robotic actuation and control of a catheter for structural intervention cardiology,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 5907–5913.
- [34] K. S. Arun, T. S. Huang, and S. D. Blostein, “Least-squares fitting of two 3-d point sets,” *IEEE Transactions on pattern analysis and machine intelligence*, no. 5, pp. 698–700, 1987.
- [35] P. Yeshwanthi, J. Chopra, P. Manik, A. Rai, S. Kumari, and K. Dande, “Morphological characteristics of interatrial septum and its clinical relevance,” *Catheterization and cardiovascular interventions: official journal of the Society for Cardiac Angiography & Interventions*. 2025.
- [36] F. Qadir, T. Ashraf, K. F. Aamir, A. S. Achakzai, S. M. Afaq, M. N. Khan, N. A. Soomro, S. S. Qureshi, and M. Karim, “Measurement of interatrial septal thickness by echocardiography in patients with moderate to severe rheumatic mitral stenosis undergoing percutaneous balloon mitral valvuloplasty,” *IJC Heart & Vasculature*, vol. 22, pp. 35–38, 2019.
- [37] J. Shirani, A. M. Zafari, and W. C. Roberts, “Morphologic features of fossa ovalis membrane aneurysm in the adult and its clinical significance,” *Journal of the American College of Cardiology*, vol. 26, no. 2, pp. 466–471, 1995.
- [38] F. Ficuciello, L. Villani, and B. Siciliano, “Variable impedance control of redundant manipulators for intuitive human-robot physical interaction,” *IEEE Transactions on Robotics*, vol. 31, no. 4, pp. 850–863, 2015.
- [39] A. Caballero, F. Sulejmani, C. Martin, T. Pham, and W. Sun, “Evaluation of transcatheter heart valve biomaterials: Biomechanical characterization of bovine and porcine pericardium,” *Journal of the mechanical behavior of biomedical materials*, vol. 75, pp. 486–494, 2017.
- [40] S. A. Howard, S. G. Quallich, M. A. Benscoter, B. C. Holmgren, C. D. Rolfes, and P. A. Iaizzo, “Tissue properties of the fossa ovalis as they relate to transseptal punctures: a translational approach,” *Journal of interventional cardiology*, vol. 28, no. 1, pp. 98–108, 2015.