POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

Executive Summary of the Thesis

Predictive Control with Online Jacobian Update for Dynamic Mitral Valve Targeting in Continuum Robots

Laurea Magistrale in Biomedical Engineering - Ingegneria Biomedica

Author: Filippo Manini

Advisor: Prof. Elena De Momi

Co-advisor: Anna Bicchi

Academic year: 2025-2026

1. Introduction

Mitral Valve Regurgitation (MVR) is a Structural Heart Disease related to the Mitral Valve (MV) with incidence increasing due to population aging and reported annual mortality rates approaching 34%. It is characterized by the abnormal backflow of blood from the left ventricle to the left atrium during systolic contraction. In severe cases, MVR requires surgical transcatheter intervention to prevent progressive cardiac dysfunction and heart failure.
Transcatheter Mitral Valve Repair (TMVR) using the MitraClip® (MC) system has emerged as a minimally invasive alternative to open chest surgery for high-risk surgical patients. Despite its clinical success, the procedure exposes clinicians to a high radiation dose, in addition to being technically demanding, requiring precise catheter manipulation and accurate clip positioning, resulting in a steep learning curve [4]. To address these challenges, this work proposes an integrated control pipeline for autonomous TMVR. Starting from processed transesophageal echocardiography (TEE) images, the framework identifies a dynamic model of the mitral valve, predicts its future position, plans an optimal interception trajectory, and computes the required actuation commands through an Online-Estimated Jacobian.The aim is to achieve accurate and robust targeting in the presence of valve motion and modeling uncertainties.

2. State of the Art

Due to its flexible and continuous structure, the MC delivery system can be modeled as a Continuum Cobot (CR). Compared to rigid-link manipulators, CRs offer intrinsic compliance, high dexterity, and the ability to navigate confined anatomical environments, reducing surgical trauma and recovery time. These mechanical advantages come at the cost of increased modeling and control complexity, as accurately mapping task space to actuation space remains challenging. Classical model-based approaches, including constant-curvature (CC), piecewise constant-curvature (PCC), and higher-fidelity elastic or Cosserat rod models, improve accuracy at the expense of significant computational burden, limiting real-time applicability [3]. Data-driven methods such as neural networks and Gaussian processes alleviate this limitation by enabling efficient kinematic learning and compensation of unmodeled effects,

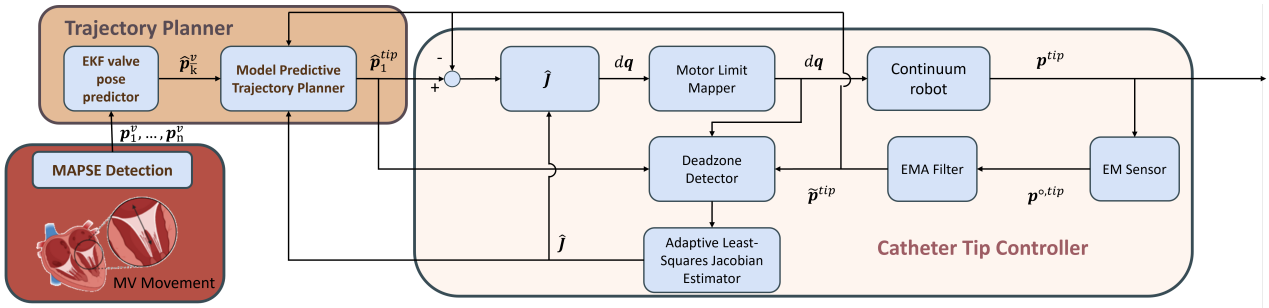

Figure 1: Proposed closed-loop control architecture integrating EKF-based MV modeling from TEE images, Model Predictive Trajectory Planning (MPTP), and Online-Estimated Jacobian (JUPD) for robust task–actuation mapping of the continuum robot.

but often suffer from limited interpretability and strong dependence on training data quality, reducing robustness in dynamic or unseen conditions. In TMVR, an additional challenge arises from the interaction with a dynamically moving anatomical target. The catheter tip must reach the MV, which varies in position during the cardiac cycle due to the Mitral Annular Plane Systolic Excursion (MAPSE). This scenario is neither pure static positioning nor continuous tracking, but rather an interception problem involving a time-varying target. From a control perspective, this can be formulated as a rendezvous problem, where two systems with independent dynamics must meet at a common point in space and time. Rendezvous strategies are well established in aerospace and autonomous systems, often relying on motion prediction and Model Predictive Control (MPC). In surgical robotics, however, cardiac motion is typically treated as a disturbance to be compensated for. Explicit rendezvous formulations between a CR and the dynamically moving mitral annulus remain largely unexplored.

3. Materials and Methods

Fig. 1 summarizes the proposed control pipeline. An Extended Kalman Filter (EKF) estimates a dynamic model of the mitral valve from TEE images (Subsec. 3.1). The estimated model and the measured CR tip position are provided to a Model Predictive Trajectory Planner (MPTP) (Subsec. 3.2), which generates a smooth and feasible trajectory.

An Online Jacobian Update (JUPD) strategy (Subsec. 3.3) estimates the task–actuation mapping from measured actuation–displacement data. The updated Jacobian is fed back to the

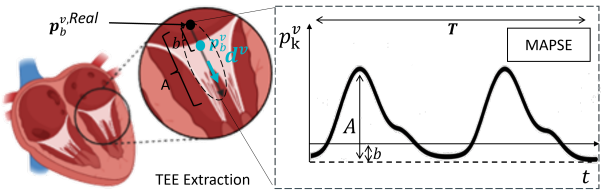

Figure 2: MAPSE identification. MV motion is described by the displacement vector $\mathbf{d}^v$ and basal point $\mathbf{p}_\mathrm{b}^v$ (cyan). A is the motion amplitude along the MV direction and b is the axial offset between $\mathbf{p}_\mathrm{b}^v$ and $\mathbf{p}_\mathrm{b}^{v,\mathrm{real}}$.

MPTP, improving trajectory accuracy and robustness against modeling uncertainties.

3.1. Extended Kalman Filter

The dominant oscillation direction of the MV plane is identified via Singular Value Decomposition (SVD) applied to a finite window of sampled valve positions. The principal eigenvector of the covariance matrix defines the unit direction vector $\mathbf{d}^v \in \mathbf{R}^3$. The valve base position $\mathbf{p}_\mathrm{b}^v$ is estimated as the mean of the 3D points corresponding to local minima of the projected motion along $\mathbf{d}^v$ (Fig 2). The valve motion is subsequently modeled using a classical EKF framework, comprising nonlinear state prediction, error covariance propagation, measurement update, and correction steps.

State Definition: The MV motion is modeled as a periodic displacement along $\mathbf{d}^v$. The state vector is defined as $\mathbf{x}_k = \left[p_k^v,\ b_k,\ A_k\right]^T \in \mathbb{R}^3$, where:

- p_k^v is the scalar valve displacement along $\mathbf{d}^v$.
- b_k is the offset term.
- A_k is the oscillation amplitude.

The oscillation frequency f is assumed known from heart-rate monitoring, with period $T = 1/f$.

State Transition Model: The nonlinear state transition is defined as

$$p^v_{k+1} = \begin{cases} A_k \sin(2\pi f t_k) + b_k, & t_k < \frac{T}{2}, \\ -A_k \sin(2\pi f t_k) + b_k, & t_k \geq \frac{T}{2}. \end{cases} \quad (1)$$

The remaining states evolve as

$$b_{k+1} = b_k, \qquad A_{k+1} = A_k. \quad (2)$$

Compactly,

$$\mathbf{x}_{k+1} = f(\mathbf{x}_k, t_k) + \mathbf{w}_k, \ \mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}), \ \mathbf{Q} \in \mathbb{R}^{3\times 3} \quad (3)$$

.

Jacobian Matrix: The Jacobian of the nonlinear transition function is

$$\mathbf{F}_k = \frac{\partial f}{\partial \mathbf{x}} = \begin{bmatrix} 0 & 1 & \alpha(t_k) \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (4)$$

where

$$\alpha(t_k) = \begin{cases} \sin(2\pi f t_k), & t_k < \frac{T}{2}, \\ -\sin(2\pi f t_k), & t_k \geq \frac{T}{2}. \end{cases} \quad (5)$$

Measurement Model: Only the scalar valve displacement is measured: $z_k = p^v_k + n_k$, $n_k \sim \mathcal{N}(0, R)$. consequently, the measurement matrix is $\mathbf{H} = \begin{bmatrix} 1 & 0 & 0 \end{bmatrix}$.

Future Valve Position Prediction: For trajectory planning, only the prediction step is executed. The predicted 3D valve plane position is obtained as: $\overline{\mathbf{p}}^v_{k+1} = \mathbf{p}^v_{\mathrm{b}} + \hat{p}^v_{k+1}\, \mathbf{d}^v$.

3.2. Model Predictive Trajectory Planning

Starting from an imperfect estimate of the task–actuation mapping $\hat{\mathbf{J}}$, the use of the same model for both trajectory prediction and cost formulation may result in suboptimal convergence and increased target-reaching time. To mitigate these effects, the proposed MPTP directly optimizes the predicted task-space tip positions $\mathbf{p}^{\mathrm{tip}}_0, \ldots, \mathbf{p}^{\mathrm{tip}}_K$, hereafter denoted as $\mathbf{x}_0, \ldots, \mathbf{x}_K$, over a finite prediction horizon K.

Cost Function Formulation: The overall cost function is defined as the sum of four components (6, where λ_t is the task cost, which penalizes the deviation of the system state from the desired target position in task space, λ_s enforces temporal consistency between successive trajectory points, λ_a penalizes task-space displacements that would require large or inefficient actuation changes through the Jacobian and λ_e acts as a soft constraint to promote convergence toward the target at the end of the horizon:

$$\begin{aligned} \theta(\mathbf{x}_0, \ldots, \mathbf{x}_K) = \\ = \sum_{k=1}^{K} \Big[\underbrace{(\mathbf{x}_k - \overline{\mathbf{p}}^v_{k+1})^\top \mathbf{T} (\mathbf{x}_k - \overline{\mathbf{p}}^v_{k+1})}_{\lambda_t} + \\ + \underbrace{(\mathbf{x}_k - \mathbf{x}_{k-1})^\top \mathbf{S} (\mathbf{x}_k - \mathbf{x}_{k-1})}_{\lambda_s} + \\ + \underbrace{\big(\hat{\mathbf{J}}(\mathbf{x}_k - \mathbf{x}_{k-1})\big)^\top \mathbf{A} \big(\hat{\mathbf{J}}(\mathbf{x}_k - \mathbf{x}_{k-1})\big)}_{\lambda_a} \Big] + \\ + \underbrace{(\mathbf{x}_K - \overline{\mathbf{p}}^v_{k+1})^\top \mathbf{W} (\mathbf{x}_K - \overline{\mathbf{p}}^v_{k+1})}_{\lambda_e} \end{aligned} \quad (6)$$

The weighting matrices $\mathbf{T}$, $\mathbf{S}$, $\mathbf{W} \in \mathbb{R}^{3\times 3}$ are constant design parameters. The actuation weighting matrix $\mathbf{A}$ is defined as a diagonal, insertion-adaptive matrix.

The adaptation is governed by a sigmoid scheduling function $\alpha(\iota) \in [0, 1]$, defined as

$$\alpha(\iota) = \frac{1}{1 + \exp\left(-\gamma \left(\frac{\iota}{\iota_{\max}} - \frac{\iota_{\max} - \iota_c}{\iota_{\max}}\right)\right)}, \quad (7)$$

where ι denotes the insertion level, $\iota_{\max}$ the maximum admissible insertion, ι_c a cutoff threshold, and γ a steepness parameter.

Using $\alpha(\iota)$, the adaptive actuation weighting matrix is defined as

$$\mathbf{A}(\iota) = \mathrm{diag} \begin{bmatrix} (1-\alpha)\, w^-_{INS} + \alpha\, w^+_{INS} \\ (1-\alpha)\, w^-_{\mathrm{ML}} + \alpha\, w^+_{\mathrm{ML}} \\ (1-\alpha)\, w^-_{\mathrm{AP}} + \alpha\, w^+_{\mathrm{AP}} \end{bmatrix}. \quad (8)$$

This formulation biases the optimization toward insertion(INS)-dominant motion during the early phase of the trajectory, while allowing increased Medio-Lateral (ML) and Antero Posterior (AP) actuation as the target is approached. See Fig3.

3.3. Online Jacobian Estimation

To compensate for configuration-dependent nonlinearities of the tendon-driven catheter without explicit Jacobian inversion, an adaptive Least-Squares task–actuation mapping is employed.

Incremental Model: The desired task-space update is defined as $\mathbf{p}^{\mathrm{tip}}_{k+1} = \mathbf{p}^{\mathrm{tip}}_k + \delta \overline{\mathbf{p}}_k$, Under a

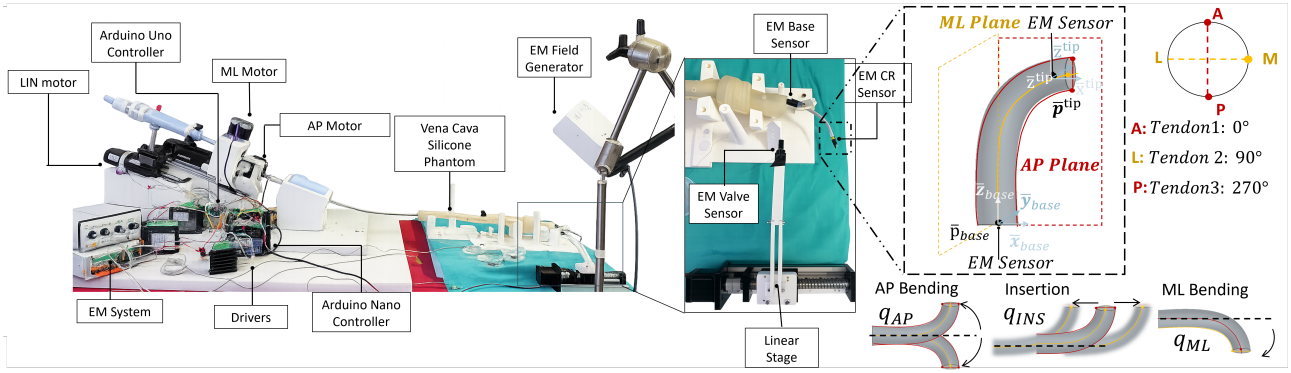

Figure 3: Left: Experimental testbed for validation of the proposed predictive control framework, including the motorized tendon-driven continuum robot (LIN, ML, AP actuation), motor drivers, and Arduino-based control units; a patient-specific vena cava silicone phantom; and an electromagnetic (EM) tracking system with three 6-DoF sensors (base, tip, and valve-mounted) providing real-time robot and target pose measurements. A motorized linear stage reproduces physiologically inspired valve motion. Right: Kinematic structure and actuation layout of the tendon-driven continuum robot, showing backbone reference frames and EM sensor placement, tendon cross-section (ML at 0°, antagonistic AP pair at 90° and 270°), and representative motions generated by the 3 Degrees of Actuation: AP bending, insertion, and ML bending, .

local linearity assumption, the incremental kinematics is approximated as

$$\delta\mathbf{q}_k \approx \mathbf{J}\,\delta\overline{\mathbf{p}}_k,\ \ \delta\mathbf{q} = [\text{INS, ML, AP}]^\top \tag{9}$$

where INS, ML and AP are the 3 Degrees of Actuation (see Fig.3)
Given the current estimate $\hat{\mathbf{J}}_k$, the prediction error is

$$\mathbf{e}_k = \delta\mathbf{q}_k - \hat{\mathbf{J}}_k\,\delta\mathbf{p}^\circ_k. \tag{10}$$

Weighted Least-Squares Update: Stacking the last H observations h, the Jacobian correction $\delta\mathbf{J}$ is obtained by minimizing

$$\min_{\delta\mathbf{J}} \sum_{h=0}^{H} \rho_h \left\| \delta\mathbf{q}_{k-h} - (\hat{\mathbf{J}}_k + \delta\mathbf{J})\delta\mathbf{p}^\circ_{k-h} \right\|_2^2 + \|\delta\mathbf{J}\|_F^2. \tag{11}$$

The forgetting factor ρ_h combines temporal decay and noise sensitivity:

$$\rho_h = \lambda_{\text{forget}}^h \frac{\sigma^2_{\text{forget}}}{\sigma^2_{\text{forget}} + \|\mathbf{e}_{k-h}\|_2^2}. \tag{12}$$

The Frobenius norm $\|\delta\mathbf{J}\|_F$ acts as a convex regularization term, ensuring numerical stability and smooth Jacobian adaptation. The initial Jacobian estimate was obtained by solving the Weighted Least-Squares problem in (11), enforcing $\rho_h = 1$ for all samples, thus removing temporal weighting after having collected 90 samples.
Sub-Actuation Strategy: To preserve local linearity, large actuation commands $\delta\mathbf{q} = [\text{INS, ML, AP}]^\top$ are decomposed into n smaller sub-actuations applied sequentially. Each sub-actuation produces an independent observation $(\delta\mathbf{q}_i, \delta\mathbf{p}^\circ_i)$, thereby improving data efficiency and enhancing robustness of the Jacobian estimation. After each sub-actuation, the Jacobian update algorithm is iterated for a maximum of 1000 iterations or until convergence, defined by $\|\delta\mathbf{J}\|_F^2 < \varepsilon_J$, where ε_J is a predefined stability threshold, enabling local stability while avoiding unnecessary computational effort.

3.4. Implementation Details

Dead-Zone Detection: To avoid corruption of the Jacobian estimation by zero-displacement samples, a dead-zone mechanism was introduced. Position measurements were filtered using an exponential moving average (EMA, $\alpha = 0.1$), and the displacement magnitude $d = \|\delta\mathbf{p}^\circ\|_2$ was evaluated after each actuation. Samples with $d < d_{\text{th}}$ were discarded.
For the AP tendon, affected by a central dead-zone and insertion-dependent relaxation, repeated pure-AP commands are applied until $d \geq d_{\text{th}}$, ensuring operation outside inactive regions.
Motor Limit Mapping: Motor saturation bounds were identified offline by commanding each actuation direction until displacement ceased, storing the maximum admissible steps as limits. During online control, the optimizer output is clipped accordingly to prevent infeasible commands while preserving stability.

4. Experiments

Hardware Components: The proposed method was validated using a comprehensive experimental testbed. The setup includes the following components:

- The tendon-driven continuum robot, actuated through stepper motors driving the AP and ML tendons, and an independent linear actuator providing insertion motion.
- Two 6-DoF electromagnetic (EM) sensors (Aurora, NDI, Canada) mounted at the robot base and distal tip, providing pose measurements at 40 Hz
- A patient-specific anatomical phantom derived from CT data, reproducing the inferior vena cava, interatrial septum, and mitral valve geometry to emulate the transseptal access pathway.
- A motorized linear stage.
- An additional 6-DoF EM sensor mounted on the moving stage, providing real-time measurements of the target position and emulating ultrasound-based annular localization.
- Dedicated motor drivers and microcontroller units operating at a control frequency of 10 Hz.
- A ROS-based architecture managing communication and synchronization between sensing, actuation, and control modules.

Experimental Protocol: The proposed control (named JUPD form now on) strategy was compared against state-of-the-art approaches for mobile target tracking, namely a classical MPC controller [1] and a model-based PID correction scheme [2]. Six task-space target points were defined at physiologically feasible distances from the valve. For each target, all three controllers were evaluated over four independent runs, yielding a total of 72 trials (*Performance Test*).

A second experimental phase assessed robustness under more realistic conditions. Valve motion was physically reproduced using a motorized linear stage equipped with an EM sensor providing position feedback, emulating ultrasound-based measurements in a clinical setting. Robustness to target-motion variations was further evaluated by commanding periodic trajectories with different amplitudes and beating frequencies within physiological ranges on 3 different points (*Robustness Test*). Variations in both MV amplitude (A) and oscillation frequency were introduced within physiological ranges. For each point–amplitude–frequency triplet, the control experiment was repeated four times, resulting in a dataset of 36 observations.

Evaluation Metrics: The following performance metrics were used to assess the proposed control strategy:

- **Success Rate (SR) [%]:** A trial is considered successful if the target is reached within $T_{\max} = 600$ s and with a positioning error $\varepsilon = 0.98$ mm (EM sensor accuracy). The success rate is defined as $\mathrm{SR} = \dfrac{N_{\mathrm{succ}}}{N} \times 100$, where N_{succ} is the number of successful trials and N the total number of trials.
- **Trajectory Length (TL) [mm]:** For each successful trial, the tip trajectory length up to the reaching time is $\mathrm{TL} = \sum_{k=1}^{K-1} \left\| \mathbf{p}_{k+1}^{\mathrm{tip}} - \mathbf{p}_{k}^{\mathrm{tip}} \right\|_2$, where K corresponds to t_{reach}.
- **Time to Target (TT) [s]:** For each successful trial, TT is defined as the elapsed time between trial start and the first instant at which the target is reached.
- **Time to Threshold (TUT_X) [s]:** For each trial, TUT_X is the time required for the Euclidean distance between the target position and the measured tip position to fall below a predefined threshold X: $\left\| \overline{\mathbf{p}}_{k+1}^{\mathrm{tip}} - \mathbf{p}_{k+1}^{\circ,\mathrm{tip}} \right\|_2 < X$.

Due to its formulation, the classical MPC controller cannot be evaluated using TUT_X. Since the MPC directly outputs control actions rather than an explicit predicted tip trajectory, this metric is not consistently definable for that method.

Statistical analysis was performed on the considered performance metrics. Results are reported using the mean and interquartile range (IQR). For SR, statistical significance was assessed using Fisher's exact test. For all other metrics, data normality was first evaluated using the Shapiro-Wilk test; since normality could not be assumed, the non-parametric Kruskal-Wallis test was applied. All tests were conducted at a significance level of $\alpha = 0.05$ or $\alpha = 0.01$. The analysis was conducted using Python.

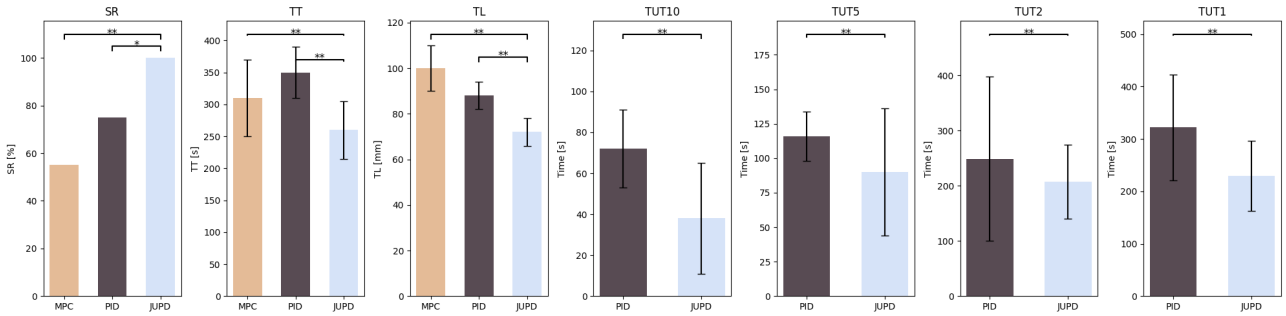

Figure 4: Performance comparison among MPC, PID, and the proposed J_{update} approach in the first set of experiments. Reported metrics include SR (%), TL (mm), TT (s), and TUT_{10} (s), TUT_5 (s), TUT_2 (s) and TUT_1 (s). Error bars denote variability across trials. Statistical significance is indicated by * ($p < 0.05$) and ** ($p < 0.01$).

Table 1: Robustness analysis of the proposed control approach under increasing frequency of the sinusoidal input applied to the linear stage. Results are reported as mean ± standard deviation.

f [Hz]	**SR [%]**	**TT [s]**	**TL [mm]**
1.0	92 ± 8	185 ± 28	52 ± 5
1.15	100 ± 0	200 ± 22	54 ± 7
1.3	100 ± 0	180 ± 35	55 ± 8

5. Results and Discussion

In the *Performance Test*, JUPD achieved a SR of 100%, outperforming MPC (54.17%) and PID (75%), with an average TT of 242.26 s versus 360.03 s and 374.98 s, respectively. Fisher's exact test confirmed the SR improvement between JUPD ($p < 0.01$) and PID ($p < 0.05$), while no significant difference was observed between PID and MPC, highlighting the impact of the online Jacobian update.

Trajectory analysis showed that JUPD produced shorter and more efficient paths than MPC ($p < 0.01$) and PID ($p < 0.01$), reducing total time to target despite slightly higher computation compared to PID. Kruskal–Wallis on all TUT thresholds revealed faster convergence for JUPD across all thresholds ($p < 0.01$), with consistently lower mean values and narrower interquartile ranges at 2 and 1 mm (Fig 4), indicating more precise and stable control near the target.

Regarding *Robustness Test* (Tab. 1), the proposed strategy failed in only one trial. No statistically significant differences were observed across amplitude-frequency variations, confirming robustness of both valve position estimation and trajectory planning under physiological oscillations.

6. Conclusions and Future Developments

This work investigated the feasibility of position control on a dynamically moving target using a continuum robot. The proposed task-space control based on an Online-Estimated Jacobian demonstrated more accurate kinematic mapping and more efficient trajectories, ultimately improving task-level performance coupled with robustness under non-ideal scenarios, including variations in target dynamics. Still, the proposed approach shows some limitations. Future developments may include improving the MV modeling, as the current approach provides only a first-order approximation. Additionally, the MPTP could be enhanced with a cost function promoting orthogonality between the robot tip orientation and the valve motion direction, for instance, by constraining the final trajectory points to be orthogonal to the MAPSE direction. While this may not guarantee perfect orthogonality along the entire trajectory, it can guide the planner toward a more physiologically favorable approach angle.

Finally, preliminary tests suggest that a Variable Prediction Horizon in the MPTP could improve performance in terms of TT. By dynamically adjusting the horizon based on the current state of the valve and catheter, the planner can better anticipate changes in target motion, adapt to uncertainties, and generate more feasible and efficient trajectories without over-constraining the optimization problem.

References

[1] Julian Berberich and Frank Allgöwer. An overview of systems-theoretic guarantees in data-driven model predictive control. *Annual Review of Control, Robotics, and Autonomous Systems*, 8(1):77–100, May 2025.

[2] Anna Bicchi, Xiu Zhang, Benjamín Ignacio Fortuño Jara, Vanessa Cannizzaro, Angela Peloso, and Elena De Momi. Model-based position control of a tendon-driven variable-length continuum robot for minimally invasive mitral valve repair. *IEEE Transactions on Medical Robotics and Bionics*, 7(2):562–571, 2025.

[3] Jessica Burgner-Kahrs, D. Rucker, and Howie Choset. Continuum robots for medical applications: A survey. *IEEE Transactions on Robotics*, 31:1–20, 11 2015.

[4] A. K. Chhatriwalla, S. Vemulapalli, D. R. Holmes Jr, D. Dai, Z. Li, G. Ailawadi, D. Glower, S. Kar, M. J. Mack, J. Rymer, and et al. Institutional experience with transcatheter mitral valve repair and clinical outcomes: insights from the tvt registry. *Cardiovascular Interventions*, 12(14):1342–1352, 2019.