

Stable (Variable) Impedance Control Learning: from Diagonal to Full Matrices

Paolino De Risi¹, Alberto Arezzo², Bruno Siciliano¹, and Fanny Ficuciello¹

¹Department of Information Technology and Electrical Engineering, Università degli Studi di Napoli Federico II

²Department of Surgical Sciences, University of Torino

INTRODUCTION

Impedance control [1] is a cornerstone in robotics, enabling manipulators to interact safely and compliantly with their environment, especially in unstructured settings. However, relying on fixed impedance gains is limiting in dynamic or unpredictable scenarios. Variable Impedance Control (VIC) addresses this by adjusting parameters online, often leveraging machine learning (ML) approaches—like Learning from Demonstration (LfD) and Reinforcement Learning (RL)—to learn task-specific variable impedance profiles. A critical challenge arises: impedance data consists of Symmetric Positive Definite (SPD) matrices, and failing to respect this geometric structure during learning or adaptation can lead to closed-loop instability. This study tackles two main stability challenges related to generating and modifying impedance parameters: I) we propose a control policy that restricts the search space for ML-tuned impedance gains (for both decoupled and coupled systems) to regions that ensure system stability, mitigating concerns related to their SPD nature; II) we employ Riemannian geometry to naturally enforce the validity of learned full symmetric positive definite (SPD) impedance matrices; III) we derive a provably stable transition duration for increasing stiffness, ensuring safe behavior during adaptation.

MATERIALS AND METHODS

This section examines stability issues related to impedance parameters: (I) stability of ML-generated diagonal impedance matrices with low controller frequency, (II) stability of ML-generated full impedance matrices, and (III) stability of varying impedance behaviors under dynamic decoupling.

Stability of ML-generated fixed diagonal impedance matrices of low frequency impedance controllers

The closed-loop dynamics of a robotic manipulator under impedance control are given by:

$$\Lambda_d \ddot{\mathbf{e}} + \mathbf{K}_v \dot{\mathbf{e}} + \mathbf{K}_p \mathbf{e} = \mathbf{h}_A, \quad (1)$$

where $\mathbf{e} = \mathbf{x}_d - \mathbf{x}$ is the Cartesian tracking error, $\mathbf{h}_A$ is the external force, and Λ_d , $\mathbf{K}_v$, $\mathbf{K}_p$ are the desired inertia, damping, and stiffness matrices, respectively. We assume $\mathbf{K}_p$ and $\mathbf{K}_v$ are diagonal and positive definite.

*The research leading to these results has been supported by the TIREN "Towards Intelligent Robotic Endoscopic Dissection" project, PRIN 2022- PNRR, identification code Prot. 20229N5YW8.

Assuming each diagonal element is generated by an ML algorithm (e.g., Cross Entropy Method as in [2]), ensuring positivity (e.g., by squaring θ —the output hyperparameter for $\mathbf{K}_p$, and computing $\mathbf{K}_v = 2\zeta\sqrt{\mathbf{K}_p\Lambda_d}$), the system is stable in continuous time. However, in discrete implementations, aliasing may cause instability with high gains at low control frequencies. To address this, we constrain the parameter search space using the Shannon-Nyquist criterion. Since the system is linear and decoupled, each DoF behaves as an independent second-order system. Assuming a fixed damping ratio ζ , its bandwidth is approximated by:

$$\omega_{bw_i} = \omega_{n_i} \sqrt{1 - 2\zeta^2 + \sqrt{2 - 4\zeta^2 + 4\zeta^4}}, \quad (2)$$

with $\omega_{n_i} = \sqrt{k_{p_i}/\lambda_{d_i}}$, k_{p_i} and λ_{d_i} representing the i -th element on the respective matrix diagonal. Bandwidth depends only on k_{p_i} for given ζ and λ_{d_i} .

To avoid aliasing, bandwidth must satisfy the Shannon-Nyquist theorem:

$$\omega_{bw_i} \leq N\omega_s \quad (3)$$

with $N > 2 \in \mathbb{N}$ (or $N = 2$ for limited bandwidth systems) and ω_s the angular sampling rate. This yields a maximum allowable stiffness:

$$k_{i_{MAX}} = \left(\frac{N}{\sqrt{1 - 2\zeta^2 + \sqrt{2 - 4\zeta^2 + 4\zeta^4}}} \omega_s \right)^2 \lambda_{d_i} \quad (4)$$

which defines the constrained search interval for the ML algorithm as:

$$-\sqrt{k_{i_{MAX}}} \leq \theta_i \leq \sqrt{k_{i_{MAX}}} \quad (5)$$

Stability of ML-generated full impedance matrices

In this section, we address the generation of full symmetric positive definite (SPD) impedance matrices via machine learning, moving beyond the common restriction to diagonal forms. Unlike diagonal matrices, generating full SPD matrices requires particular care to ensure system stability. Prior work, such as [3], used a Wishart-based variant of the Cross Entropy Method (CEM) to maintain SPD structure, replacing standard Gaussian sampling. However, this approach lacks a closed-form update rule, relying instead on empirical adjustments. We propose an alternative based on the Riemannian geometry of SPD matrices, interpreting them as elements of the manifold $\mathcal{S}_n^{++}$. This view enables us to leverage intrinsic geometric properties unavailable in Euclidean space. We begin by

introducing key Riemannian concepts: a manifold is a space equipped with smoothly varying inner products on tangent spaces, allowing definitions of geodesic distances, angles, and curvature. For SPD matrices, the geodesic distance between $X, Y \in \mathcal{S}_n^{++}$ is:

$$d(X, Y) = \|\log(X^{-1/2} Y X^{-1/2})\|_F,$$

with $\|\cdot\|_F$ denoting the Frobenius norm. This framework allows local linearization via projection onto the tangent space $T_X \mathcal{S}_n^{++}$, and reconstruction via the exponential map:

$$\begin{aligned} \exp_X(S) &= X^{1/2} \exp(X^{-1/2} S X^{-1/2}) X^{1/2}, \\ \log_X(Y) &= X^{1/2} \log(X^{-1/2} Y X^{-1/2}) X^{1/2}, \end{aligned} \quad (6)$$

where $\exp(\cdot)$ and $\log(\cdot)$ denote the matrix exponential and logarithm, respectively. To adapt Gaussian-based learning to the manifold, we use the Riemannian Gaussian distribution. Given a mean $\mu \in \mathcal{S}_n^{++}$ and covariance Σ , we sample $v \sim \mathcal{N}(0, \Sigma)$ in the tangent space $T_\mu \mathcal{S}_n^{++}$ and map it back via $\exp \mu(v)$. The probability density is:

$$p(x) = \frac{1}{Z} \exp\left(-\frac{1}{2} \log_\mu(x)^T \Sigma^{-1} \log_\mu(x)\right) |\det D \log_\mu(x)|^{-1}$$

where Z is the normalization constant and $D \log_\mu(x)$ is the differential of the logarithmic map at x . The use of Riemannian geometry thus enables the formulation of learning algorithms that respect the curved nature of the underlying SPD manifold, providing a principled framework for analysis and generation of such structured data, thus guaranteeing generation of SPD matrices.

Stability of Varying Impedance Behaviors Under Dynamic Decoupling

Assuming time-varying diagonal impedance matrices from an initial stiffness k_{s_i} to a target k_i^* , an instantaneous switch is unstable [4], necessitating a monotonic C^1 impedance profile over a duration T_K . While energy tank approaches (e.g., [5]) exist, they lack rigorous guarantees for tracking predefined stiffness profiles [6], potentially compromising task performance. Therefore, alternative strategies are needed. Due to the decoupled closed-loop dynamics in (1), we distinguish two cases:

(I) *Decreasing stiffness* ($k_{s_i} > k_i^*$): These profiles are stable as the controller dissipates energy [4]. Considering the Lyapunov candidate:

$$V_i(\tilde{q}_i, \dot{\tilde{q}}_i, t) = \frac{1}{2} \dot{\tilde{q}}_i^T b_{d_i} \dot{\tilde{q}}_i + \frac{1}{2} \tilde{q}_i^T k_i(t) \tilde{q}_i \quad (7)$$

Its derivative is:

$$\dot{V}_i(\tilde{q}_i, \dot{\tilde{q}}_i, t) = -\dot{\tilde{q}}_i^T d_i(t) \dot{\tilde{q}}_i + \frac{1}{2} \dot{\tilde{q}}_i^T \dot{k}_i(t) \tilde{q}_i \quad (8)$$

Since $\dot{k}_i(t) < 0$, $\dot{V}_i < 0$, ensuring stability.

(II) *Increasing stiffness* ($k_{s_i} < k_i^*$): In this case, the controller injects energy, and stability is not guaranteed. Using the *theorem for stability conditions under dynamic decoupling* from [4], we determine the minimum duration T_K . By assembling increasing impedance parameters in $\bar{\mathbf{B}}_d$, $\bar{\mathbf{K}}(t)$, $\bar{\mathbf{K}}_v(t)$ (where $\bar{\mathbf{B}}_d$ is constant, symmetric, positive definite, and others are symmetric, positive definite, and

C^1), the system is uniformly globally asymptotically stable if an $\alpha > 0$ exists such that $\forall t \geq 0$:

$$\alpha \bar{\mathbf{B}}_d - \bar{\mathbf{K}}_v(t) < \mathbf{0}, \quad \dot{\bar{\mathbf{K}}}_p(t) + \alpha \bar{\mathbf{K}}_v(t) - 2\alpha \bar{\mathbf{K}}_p(t) < \mathbf{0} \quad (9)$$

The first condition is satisfied by selecting:

$$\alpha = \min_t \frac{\lambda(\bar{\mathbf{K}}_v(t))}{\bar{\lambda}(\bar{\mathbf{B}}_d)} \quad (10)$$

where $\lambda(\cdot)$ and $\bar{\lambda}(\cdot)$ are the minimum and maximum eigenvalues. The full procedure for computing a safe T_K is outlined in Algorithm 1. While accurate with small ΔT_K , its recursive design is unsuitable for real-time control; larger ΔT_K reduces computation time at the cost of precision.

Algorithm 1 Pseudocode for the computation of T fulfilling the *stability conditions under dynamic decoupling* in [4]

Initialize: $\mathbf{K}_{p_s}$ (initial stiffness), $\mathbf{K}^*$ (target stiffness from GPs), $\mathbf{B}_d$ (shaped inertia), $T_K > 0$, $\Delta_t > 0$;

Define: impedance profile matrices $\bar{\mathbf{B}}_d$, $\bar{\mathbf{K}}_p(t)$, $\bar{\mathbf{K}}_v(t)$;

repeat

 Compute $\bar{\mathbf{K}}_p(t)$, $\bar{\mathbf{K}}_v(t)$, and derivatives $\dot{\bar{\mathbf{K}}}_p(t)$, $\dot{\bar{\mathbf{K}}}_v(t)$;

 Compute $\alpha = \min_t \frac{\lambda(\bar{\mathbf{K}}_v(t))}{\bar{\lambda}(\bar{\mathbf{B}}_d)}$ to satisfy (9);

 Compute $\mathbf{C}(t) = \dot{\bar{\mathbf{K}}}_p(t) + \alpha \bar{\mathbf{K}}_v(t) - 2\alpha \bar{\mathbf{K}}_p(t)$;

$T_K \leftarrow T_K + \Delta_t$;

until $\mathbf{C}(t) < \mathbf{0} \quad \forall t \leq T_K$

return T_K

CONCLUSIONS AND DISCUSSION

In this work, we presented a comprehensive framework to ensure stability in learning-based variable impedance control. We introduced formally grounded methods to constrain ML-generated diagonal impedance gains at low control frequencies, leverage Riemannian geometry to guarantee the generation of valid full SPD matrices; and computed a provably stable transition time for dynamically increasing stiffness. These contributions provide a rigorous methodology for developing safe and adaptive impedance controllers, addressing stability from parameter generation to real-time adaptation.

REFERENCES

- [1] N. Hogan, "Impedance control: An approach to manipulation," in *1984 American Control Conference*, 1984, pp. 304–313.
- [2] P. De Risi, F. Amadio, G. Garofalo, F. Ficuciello, and P. Falco, "Variable impedance control combining reinforcement learning and gaussian process regression," in *2024 IEEE-RAS 23rd International Conference on Humanoid Robots (Humanoids)*, 2024, pp. 343–350.
- [3] S. A. Khader, H. Yin, P. Falco, and D. Kragic, "Stability-guaranteed reinforcement learning for contact-rich manipulation," *IEEE Robotics and Automation Letters*, vol. 6, no. 1, pp. 1–8, 2021.
- [4] K. Kronander and A. Billard, "Stability considerations for variable impedance control," *IEEE Transactions on Robotics*, vol. 32, no. 5, pp. 1298–1305, 2016.
- [5] F. Ferraguti, C. Secchi, and C. Fantuzzi, "A tank-based approach to impedance control with variable stiffness," in *2013 IEEE International Conference on Robotics and Automation*, 2013, pp. 4948–4953.
- [6] K. Kronander and A. Billard, "Stability considerations for variable impedance control," *IEEE Transactions on Robotics*, vol. 32, no. 5, pp. 1298–1305, 2016.