

Model predictive impedance control with Gaussian processes for human and environment interaction

Haninger, Kevin; Hegeler, Christian; Peternel, Luka

DOI

[10.1016/j.robot.2023.104431](https://doi.org/10.1016/j.robot.2023.104431)

Publication date

2023

Document Version

Final published version

Published in

Robotics and Autonomous Systems

Citation (APA)

Haninger, K., Hegeler, C., & Peternel, L. (2023). Model predictive impedance control with Gaussian processes for human and environment interaction. *Robotics and Autonomous Systems*, 165, Article 104431. <https://doi.org/10.1016/j.robot.2023.104431>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy


Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.



Model predictive impedance control with Gaussian processes for human and environment interaction

Kevin Haninger^{a,*}, Christian Hegeler^a, Luka Peternel^b

^a Department of Automation, Fraunhofer IPK, Berlin, Germany

^b Cognitive Robotics Department, Delft University of Technology, Delft, The Netherlands

ARTICLE INFO

Article history:

Received 6 December 2022

Received in revised form 6 February 2023

Accepted 13 April 2023

Available online 23 April 2023

Keywords:

Impedance control
Model predictive control
Human–robot interaction
Gaussian processes
Intention detection

ABSTRACT

Robotic tasks which involve uncertainty – due to variation in goal, environment configuration, or confidence in task model – may require human input to instruct or adapt the robot. In tasks with physical contact, several existing methods for adapting robot trajectory or impedance according to individual uncertainties have been proposed, e.g., realizing intention detection or uncertainty-aware learning from demonstration. However, isolated methods cannot address the wide range of uncertainties jointly present in many tasks.

To improve generality, this paper proposes a model predictive control (MPC) framework which plans both trajectory and impedance online, can consider discrete and continuous uncertainties, includes safety constraints, and can be efficiently applied to a new task. This framework can consider uncertainty from: contact constraint variation, uncertainty in human goals, or task disturbances. An uncertainty-aware task model is learned from a few (≤ 3) demonstrations using Gaussian Processes. This task model is used in a nonlinear MPC problem to optimize robot trajectory and impedance according to belief in discrete human goals, human kinematics, safety constraints, contact stability, and frequency-domain disturbance rejection. This MPC formulation is introduced, analyzed with respect to convexity, and validated in co-manipulation with multiple goals, a collaborative polishing task, and a collaborative assembly task.

© 2023 Elsevier B.V. All rights reserved.

1. Introduction

Manipulation with physical contact is essential for robotics, where impedance control¹ is often paramount to robust and safe execution. This is especially critical in tasks that involve physical human–robot interaction (HRI), where the robot compliance allows the human co-worker to safely affect robot motion for teaching or online collaboration. Impedance control defines the characteristics of physical interaction by stiffness, damping, and inertia parameters [1]. These parameters must be appropriately adapted to an application, with a large number of proposed methods to determine appropriate impedance parameters. Nevertheless, in practice, the tuning of these parameters is often done with ad-hoc or application-specific methods.

Some recent works propose generalizable approaches to adapting impedance parameters based on machine learning and human-in-the-loop approaches. A common theme in impedance adaptation is the degree of uncertainty. From the literature, we

identify five key categories related to uncertainty and impedance control. (1) In learning from demonstration (LfD), uncertainty derived from the *variability in demonstrations* of trajectories can be used to infer lower robot stiffness, since the human demonstrator does not seem to care about precision in that section of motion [2,3]. (2) In LfD, the robot's available *confidence in the skill* is used, where stiffness can be reduced when entering unexplored regions as actions might be unsafe [4]. (3) When faced with uncertainty in *environment constraints* in task execution, the robot should have lower stiffness in order not to produce dangerously high forces when the contact is unexpected [5–7]. (4) In case of uncertain *external perturbations* the robot should increase the impedance to ensure the accuracy of position tracking task [8,9]. (5) When there is uncertainty in *human goal* during collaborative tasks, the robot should be compliant to let them determine a specific degree of freedom (DoF). For example, when the human should lead the collaborative carrying of an object [10], or when the human should guide a polishing machine that is carried by the robot [11].

Nevertheless, scaling up impedance control to different tasks, conditions, and environments while accounting for human goals in real-time is still a major challenge in physical HRI. The existing methods typically focus on either one or a couple of these

* Corresponding author.

E-mail address: kevin.haninger@ipk.fraunhofer.de (K. Haninger).

¹ Impedance is used to refer to general physical interaction control, i.e., both impedance and admittance control.

uncertainty categories, although many tasks feature several of them. We aim to develop a framework which can account for multiple types of uncertainties for a unified physical HRI approach. The goal is for robots to be able to plan trajectories and adapt impedance for semi-structured tasks in physical contact with humans and environments, without needing an analytical model of environment/human or large amounts of data.

To attain this goal, we propose a planning and optimization framework with probabilistic models of external force. Gaussian processes are used to model force from the environment or human, and used as input to an admittance model of the robot to project stochastic trajectories. These trajectories can be evaluated with cost functions and constraints, to allow online optimization of admittance parameters and force reference trajectory. This acts as an outer-loop controller, setting parameters and force reference to an inner-loop admittance controller which maintains real-time responsiveness to external forces. The main contributions of this work are:

- A taxonomy of uncertainty in physical interaction, including their sources, consequences for safety and performance, and a unified problem formulation to address these with adaptive impedance control.
- Unlike existing approaches that use GPs as direct robot reference behavior generators (e.g., [12]), we use them as intermediate task models which combine with the robot impedance in a Model Predictive Control (MPC) problem, allowing explicit objectives which depend on the trajectory uncertainty, intention detection, safety constraints, and frequency-domain costs.
- Unlike existing sampling-based approaches to adapting impedance online (e.g., [13]), we propose the first gradient-based optimization of impedance, which allows constraints in the optimization problem such as maximum contact force and contact stability.
- The approach is experimentally validated, inferring discrete human goal online, increasing impedance damping for contact stability, and reducing impedance in directions of task variation, with 10–20 Hz MPC rates. The code² and video³ are available.

This approach handles various types of uncertainty as seen in Fig. 1. The uncertainty category (1) related to variability in demonstrations can be accounted for in the proposed framework through GP covariance, but we note this can produce unsafe impedance skills when the low demonstration variance arises from contact with the environment [7]. The proposed framework also addresses the uncertainty category (2) related to confidence in skill by leveraging on GP model uncertainty in the cost function and optimization constraints. The uncertainty category (3) related to environment constraints is handled by the increase in GP covariance induced by fitting hyperparameters on data with contact variance. An additional frequency-domain perturbation model is used to resolve the uncertainty category (4) related to external perturbations. The uncertainty category (5) related to human goal estimation is addressed by taking the expectation over the belief in human goal, inferred from the GPs.

A preliminary study was presented at the 2022 IEEE International Conference on Robotics and Automation [14]. This paper significantly expands upon the previous study by an extended theoretical formulation of the method that accounts for all identified uncertainties, adding two types of MPC constraints, a new theoretical analysis of impedance optimization, and a new experimental evaluation with three additional practical tasks involving

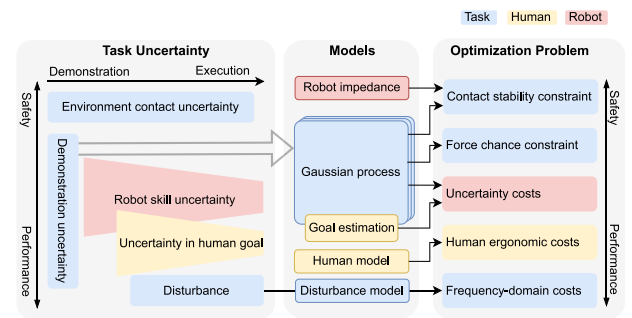


Fig. 1. Sources of uncertainty during both demonstration and execution (left), how they are modeled (center), and how they enter the MPC framework (right).

physical HRI (collaborative rail assembly, double peg-in-the-hole, and polishing) where we examine new aspects (task uncertainty, goal uncertainty, and human ergonomics).

The remainder of the paper is structured as follows; also see Fig. 2 for an overview. The related work is given in Section 2. The modeling using GP is described in Section 3. The MPC problem formulation is provided in Section 4. This is followed by a theoretical analysis in Section 5. Finally, experimental validation on practical tasks is described in Section 6.

2. Related work

Here we examine the proposed method with respect to related work in the literature. We first outline the work done in physical human–robot interaction where human intention modeling plays a critical role in seamless collaboration. Then we proceed to robot impedance adaptation, which is a fundamental part of the robot's low-level actions as a result of the inferred intention. Finally, we review planning in contact tasks.

2.1. Physical human–robot interaction

An important challenge in physical HRI is the ability of robots to infer human co-worker's goals and states. To address this challenge, several recent works have proposed inference techniques and applications. The study in [15] showed that robot adaptation to the human task execution can improve task-related performance. Robot adaption to the human state can also improve ergonomics [11]. Finally, user satisfaction can also be improved by robot adaptation [16]. Nevertheless, to adapt to these different aspects, the robot should be able to infer the human state, which often requires complex sensory setups [17,18]. However, complex measurement systems are not always available or feasible for practical applications, thus some specific states can be modeled, such as the desired motion of the robot/payload [19–22], human actions [23], preferences [24,25], human physical fatigue [11], cognitive stress [26], and suitability of shared workspace [27–30].

For seamless collaboration, the adaption of the robot to the human should be prompt, and thus the human intent or goal must be inferred during the task execution [11,31]. Human intent is often considered as a continuous variable [15,22], however, it can also reflect discrete changes in the task [32], e.g., a collection of possible goals. Discrete human state has been considered for virtual fixtures [33], Dynamic Movement Primitives (DMPs) [32], and impedance control [34].

When the robot has the knowledge of the human state and intent, it must respond promptly with appropriate actions to facilitate the collaborative task execution. This typically means generation of motion trajectories and often simultaneous impedance

² <https://gitlab.cc-asp.fraunhofer.de/hanikevi/gp-mpc-impedance-control>

³ <https://youtu.be/Of203mHfM94>

Table 1

Selected prior work in adaptation of impedance, organized according to motivating principle. Robotic applications to co-manipulation and contact are shown, as well as the targeted degree of robot autonomy.

	Principle	Robot impedance should:	Co-manip.	Env. Contact	Autonomy
Dynamics	Complementarity [46]	Be inverse of environment	–	–	–
	Stability [9,47]	Stabilize unstable dyn	[8]	[48]	Full
	Well-damped [39]	Reduce oscillation or jerk	[10,39]	[49]	Interact, full
Optimal Control	Min intervention [2]	stiff where low variance	[43,50,51]	[51] safety ^a [7]	Interact, full
	game theory [52]	Min coupled cost function	[53]	–	Interact
	Reinf. learning	Min cost function	[36]	[54,55]	Full
Transfer Hum Imp	Human hand imp [3]	Match human impedance	–	[56,57]	Full
	Co-contraction [58]	Prop or inv to co-contraction	–	[58]	Interact
	Demo dynamics [34]	Which explains demos	[34]	–	Interact

^aThis relates to the robot being stiff while being rigidly constrained by a potentially high-stiffness environment, which can cause contact instability.

adaptation. The trajectories can be chosen to optimize task-related objectives: reducing trajectory jerk [35,36], minimize positioning error [8,37,38], or render appropriate velocity response [39,40]. In addition, human operator-related objectives can be optimized as well, such as minimizing interaction forces [41] or metabolic cost [42]. A common rule is that the robot can take over aspects of the task that require precision, while the human can handle adaption to variations. The ‘minimum intervention principle’ [43] follows this rule, where uncertain DOF should have a lower stiffness [44,45], which can also be interpreted as a risk-sensitive control. Nevertheless, such an approach is not reasonable for all interactive tasks that involve physical constraints [7].

The literature can address the human-related properties (e.g., actions, preferences, ergonomics, etc.), real-time intention processing, and uncertainty in an isolated manner. However, a unified framework that can effectively combine all these aspects for practical manufacturing tasks is still missing.

2.2. Impedance adaptation

Physical HRI requires a safe and robust low-level control approach to account for the uncertainty of practical tasks [58]. Classic position or force controllers often fail to perform well in such complex settings since they only focus on controlling either one of the variables, but not the characteristics of the interaction that define their relationship, i.e., impedance. Impedance adaption is empirically important for safety and robustness in manipulation, and has been shown to improve sample efficiency of higher-level learning control [59]. Accordingly, a wide range of techniques has been proposed for adjusting impedance parameters. Recent surveys provide a thorough overview of this field [60,61], so we present only an abbreviated summary in Table 1.

Dynamics-based objectives are often based on continuous-time dynamic models, considering factors like stability and damping of the coupled system. A challenge is that environment dynamics (e.g., stiffness) are application-specific. In some cases, the environment dynamics are known or controlled [47], in other cases properties of the dynamics (e.g., damping) can be found from trajectories [62].

Other approaches plan to minimize a general cost function, which may include the human or environment state [53]. Such approaches allow consideration of uncertainty, but typically also require models of humans or the environment, which are typically simplified such as assuming human force is a PD controller [43]. Data-driven methods such as reinforcement learning can be applied [54,55], but often require lots of trial and error.

An approach in [4] uses GP models to encode demonstrated desired robot movements and then leverages on the measure of uncertainty in the modeled skill to modulate the stiffness through a direct expression. If the robot experiences unexpected external perturbations, the impedance controller forces the motion into the regions of higher certainty to be more confident in the given skill. If the robot still ends up in a region of high uncertainty, the stiffness is reduced as there might be insufficient confidence in the skill to safely perform the movements there. Nevertheless, there was no higher-level planner (e.g., MPC) to exploit given learned trajectories to generate new ones.

From this, we conclude that there is a wide range of objectives and models—it seems unlikely that a single principle would be suited for all or a large majority of these areas. Accordingly, we propose a framework for optimizing impedance, capable of accommodating a range of models and objectives.

2.3. Planning in contact tasks

Intention modeling and low-level impedance control alone can work well for simplified tasks. However, for a more flexible robot behavior, a higher-level planner is required. MPC is becoming increasingly popular for manipulation [13,63], and is already a standard tool for locomotion and whole-body control, capable of planning through contact models with either linear complementarity [64] or switching dynamic modes [65]. Contact can be considered as fully rigid constraints, using time-stepping approaches to resolve the contact impulse over an integrator time step [66], avoiding the issue of unbounded contact forces in rigid contact. A dynamic model for the environment has also been proposed for manipulation [67], relaxing kinematic constraints into a stiffness. In most contact planning situations, signed distance functions are needed—typically, this requires geometry information typically given by CAD, although it has been learned on simplified objects [68].

We make two observations from the perspective of manipulation: (1) signed distance functions are difficult with the complex geometries typical in assembly and manipulation tasks, and (2) compliance is both common and critical—be it physical or by control. To address this, we use a compliant contact model here (regressing contact force over robot position) which does not require a signed distance function, and we consider the optimization of impedance parameters in the MPC problem.

3. Models

This section presents the models used: GP models of external force (from the environment or human), human kinematics, robot admittance, and force disturbance. This aspect is highlighted by with orange in Fig. 2.




Fig. 2. Overview of the proposed approach, with corresponding paper sections in red. Initially, human demonstrations are used to model human and task forces with Gaussian Processes. These models are then used in a model predictive controller that accounts for robot impedance and trajectories, and includes costs and constraints related to the task and human. The MPC controller runs online, updating trajectory and impedance parameters for a lower-level impedance controller.

3.1. Force models

External forces, arising from contact with the human or environment, are modeled as Gaussian processes over robot pose. These models are used to capture two types of uncertainty: continuous uncertainty in force via the covariance of the GP, and discrete variation by having a discrete collection of GPs. Discrete variation arises from, e.g., having multiple discrete goals in a task, where the human forces will depend on the current goal. Concretely, for a robot with a tool-center point (TCP) pose \mathbf{x} , we consider three types of external forces:

Human forces $\mathbf{f}^h(\mathbf{x}, n)$ are a function of robot state and human mode n , where $n \in [1 \dots N]$ indexes discrete possible human behavior such as goals.

Contact forces $\mathbf{f}^c(\mathbf{x}, \mathbf{x}^e)$ depend on an environment configuration \mathbf{x}^e which is typically unmeasured and may vary between iterations, such as a contact or constraint position.

Disturbance forces $\mathbf{f}^d(t)$ such as polishing, cutting, or other process forces which should not affect robot position. We assume they can be represented in a frequency-domain model $\mathbf{F}^d(j\omega)$, e.g., they are predominantly high-frequency.

We use one force/torque (F/T) sensor, such that demonstrations measure total force $\mathbf{f} = \mathbf{f}^h + \mathbf{f}^c + \mathbf{f}^d$. While multiple F/T sensors can separate human and environment forces [55,56], a single F/T sensor reduces complexity. A disadvantage of a single F/T sensor is that in co-manipulation with contact, contact may be more difficult to identify, although this can be partly addressed by how the data is collected. Another disadvantage of this approach is that forces are regressed over position, which may be challenging in stiff contact. We validate both these aspects in Section 5.4 and Fig. 6. Velocity-dependent terms of human forces are not captured with this method, and human force may vary between data collection and online execution, due to factors such as the human adapting their impedance. We show the proposed method is still capable of inferring human goal from online forces in Section 6.4.

3.1.1. GP models

GPs are used to regress force over pose for each mode n , based on a dataset of measured pose and force $\mathcal{D} = \{\{\mathbf{x}_1, \mathbf{f}_1\} \dots, \{\mathbf{x}_M, \mathbf{f}_M\}\}$. We fit a GP model [69, p13] that gives a Gaussian variable $\mathbf{f}(\mathbf{x}) \sim \mathcal{N}(\boldsymbol{\mu}_f(\mathbf{x}), \boldsymbol{\Sigma}_f(\mathbf{x}))$ for a pose \mathbf{x} , given a mean function $\mathbf{m}(\mathbf{x}) \in \mathbb{R}^6$, kernel function $k(\mathbf{x}, \mathbf{x}') \in \mathbb{R}$, and observed data \mathcal{D} . The GP models for each dimension of \mathbf{f} are fit independently and the covariance $\boldsymbol{\Sigma}_f$ is therefore diagonal. As a result, correlations between the force dimensions are not modeled, which could arise from, e.g., changes in contact plane orientation or in the direction of desired motion from the operator.

In the sequel, we consider the fitting of a single dimension of \mathbf{f} with corresponding scalars f, μ_f, Σ_f . Here, we use a squared exponential kernel of the form

$$k(\mathbf{x}, \mathbf{x}') = \sigma_n^2 \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{x}')^T \text{diag}(\mathbf{I}^{-2})(\mathbf{x} - \mathbf{x}')\right), \quad (1)$$

with hyperparameters σ_n^2 for the noise variance and $\mathbf{I} \in \mathbb{R}^6$ as the length scale. The GP then defines a marginal Gaussian distribution for a test point \mathbf{x} as

$$p(\mathbf{f} | \mathbf{x}, \mathcal{D}, n, k) \sim \mathcal{N}(\boldsymbol{\mu}^f, \boldsymbol{\Sigma}^f), \quad (2)$$

where mean $\boldsymbol{\mu}^f = \mathbf{m}(\mathbf{x}) + \mathbf{k}_*^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} [f_1, \dots, f_M]^T$, covariance $\boldsymbol{\Sigma}^f = \sigma_n^2 - \mathbf{k}_*^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{k}_*$ are built from kernel matrices $\mathbf{k}_* = [k(\mathbf{x}_1, \mathbf{x}), \dots, k(\mathbf{x}_M, \mathbf{x})]^T \in \mathbb{R}^M$ and $\mathbf{K} \in \mathbb{R}^{M \times M}$, where the i, j^{th} element is $k(\mathbf{x}_i, \mathbf{x}_j)$.

A straightforward way to incorporate prior knowledge is a mean function. A common mean function is the zero function, $\mathbf{m}(\mathbf{x}) = 0$. To better model contact forces \mathbf{f}^c , a hinge function is also used for the mean:

$$\mathbf{m}(\mathbf{x}) = \begin{cases} c_1, & \text{if } x \leq c_3 \\ c_1 + c_2 x, & \text{otherwise} \end{cases} \quad (3)$$

where x is the element of \mathbf{x} which matches the dimension of the force being fit f , and constants $c_{1,2,3} \in \mathbb{R}^6$ are hyperparameters which are fit from data.

The data for these models are collected from initial demonstrations where the robot is manually guided by the human. These models are built and the hyperparameters are fit (according to negative log-likelihood) independently for each mode n and dimension of \mathbf{f} . By implementing them in an automatic differentiation framework, the gradient of the mean and covariance with respect to pose can be efficiently calculated for the MPC problem.

3.2. Robot admittance dynamics

This subsection develops the dynamics of a robot with a Cartesian admittance controller. An admittance is used as the hardware implementation here uses a position-controlled industrial robot. As impedance and admittance are just inverse mathematical viewpoints of the interactive dynamics, the dynamics here can be readily adapted to impedance-controlled robots. The admittance dynamics are used to model how the external forces affect the distribution of robot trajectories, to allow the optimization problem to consider how admittance and nominal trajectory affect the trajectory distribution.

Consider a rendered admittance in the TCP frame which approximates the continuous-time dynamics of

$$\mathbf{f} - \mathbf{f}^r = \mathbf{M}\ddot{\mathbf{x}} + \mathbf{D}\dot{\mathbf{x}} + \mathbf{K}(\mathbf{x} - \mathbf{x}_0), \quad (4)$$

with pose $\mathbf{x} \in \mathbb{R}^6$, velocity $\dot{\mathbf{x}} \in \mathbb{R}^6$, acceleration $\ddot{\mathbf{x}} \in \mathbb{R}^6$, measured force $\mathbf{f} \in \mathbb{R}^6$ and force reference \mathbf{f}^r . The impedance parameters are inertia $\mathbf{M} \in \mathbb{R}^{6 \times 6}$, damping $\mathbf{D} \in \mathbb{R}^{6 \times 6}$, and stiffness $\mathbf{K} \in \mathbb{R}^{6 \times 6}$ matrices, which are here all diagonal, and \mathbf{x}_0 which is the rest position of the spring. The rotational elements of \mathbf{x} and \mathbf{f} are the angles and torques, respectively, about the three axes of the




Fig. 3. At time step t , the stochastic dynamics are rolled out from $\xi_t = [\mathbf{x}_t, \dot{\mathbf{x}}_t]$, with a trajectory for each mode n . Here, two modes are shown, resulting in the green and red trajectories of mean and covariance.

TCP frame. In the sequel, \mathbf{K} and \mathbf{x}_0 are dropped as it was found empirically that the quasi-static tracking can be handled by the MPC adjusting \mathbf{f}^r .

3.2.1. Integrator

To formulate an MPC problem, the dynamics in (4) have to be discretized. Let T_s be the sample time, subscript t denote the value at time $t_0 + tT_s$, and state vector $\xi = [\mathbf{x}^T, \dot{\mathbf{x}}^T]^T$.

A first-order explicit Euler integrator of (4) has oscillation when $1 - T_s D_i M_i^{-1} < 0$ [14], where \bullet_i denotes the i , i -th element of a matrix. Instead, we use the semi-implicit integration of [66] on (4) to find

$$\xi_{t+1} = \begin{bmatrix} \mathbf{I} & T_s \mathbf{I} \\ \mathbf{0} & \mathbf{M}(\mathbf{M} + T_s \mathbf{D})^{-1} \end{bmatrix} \xi_t + \begin{bmatrix} \mathbf{0} \\ T_s \mathbf{M}^{-1} \end{bmatrix} (\mathbf{f}_t - \mathbf{f}_t^r), \quad (5)$$

where $\mathbf{I} \in \mathbb{R}^{6 \times 6}$ and $\mathbf{0} \in \mathbb{R}^{6 \times 6}$ are the identity and zero matrix, respectively.

3.2.2. Stochastic dynamics

The robot's admittance affects the sensitivity to force, and thus the covariance of future trajectories when the forces are stochastic. To calculate the trajectory distribution for a single mode and dropping the argument n for compactness, we can rewrite (5) in the form of

$$\xi_{t+1} = \mathbf{A}_\phi \xi_t + \mathbf{B}_\phi (\mathbf{f}_t - \mathbf{f}_t^r), \quad (6)$$

where system matrices $\mathbf{A}_\phi, \mathbf{B}_\phi$ depend on impedance parameters $\phi = \{\mathbf{M}, \mathbf{D}\}$. Consider a Gaussian distribution over state $\xi_t \sim \mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$, we find

$$\boldsymbol{\mu}^+ = \mathbf{A}_\phi \boldsymbol{\mu} + \mathbf{B}_\phi (\boldsymbol{\mu}^f(\boldsymbol{\mu}) - \mathbf{f}^r) \quad (7)$$

$$\boldsymbol{\Sigma}^+ = \mathbf{A}_\phi \boldsymbol{\Sigma} \mathbf{A}_\phi^T + \mathbf{B}_\phi \boldsymbol{\Sigma}^f(\boldsymbol{\mu}) \mathbf{B}_\phi^T, \quad (8)$$

where \bullet^+ denotes the value of \bullet at time step $t + 1$. Let the vectorization of the covariance matrix be denoted $\boldsymbol{\sigma}^+ = \text{vec}(\boldsymbol{\Sigma}^+)$ and similarly $\boldsymbol{\sigma}^f = \text{vec}(\boldsymbol{\Sigma}^f)$. Vectorizing (8) gives [70, (520)]

$$\boldsymbol{\sigma}^+ = \mathbf{A}_\phi \otimes \mathbf{A}_\phi^T \boldsymbol{\sigma} + \mathbf{B}_\phi \otimes \mathbf{B}_\phi^T \boldsymbol{\sigma}^f, \quad (9)$$

where \otimes is the Kronecker product.

As there is a GP force model for each mode, the trajectory is rolled out for each mode by applying the corresponding $\boldsymbol{\mu}^f(\boldsymbol{\mu}, n)$ and $\boldsymbol{\sigma}^f(\boldsymbol{\mu}, n)$ from the current state ξ_t , as illustrated in Fig. 3. These trajectories use the same admittance parameters and \mathbf{f}_t^r , but result in trajectories ξ_{t+1}^n and $\boldsymbol{\Sigma}_{t+1}^n$.

3.3. Human kinematics

Ergonomic costs are more naturally expressed in terms of human joint torques [18,30], so we also consider a 4-DOF kinematic model of the human arm with three rotational DOF at the

shoulder and one rotational joint for the elbow extension. This model has parameters of l_1 and l_2 , for the length of the upper and lower arm, and $\mathbf{x}^{sh} \in \mathbb{R}^3$ for the spatial position of the human shoulder. Using these parameters and the human joint angles $\mathbf{q} \in \mathbb{R}^4$, we can calculate the human hand position $\mathbf{x}^H \in \mathbb{R}^3$ with forward kinematics as

$$\mathbf{x}^H = \text{FK}(\mathbf{q}, \mathbf{x}^{sh}), \quad (10)$$

Forces measured at the end-effector can be translated into human joint torques $\boldsymbol{\tau}_H$ as

$$\boldsymbol{\tau}_H = \mathbf{J}_H^T(\mathbf{q}) \mathbf{f}_H, \quad (11)$$

where $\mathbf{f}_H \in \mathbb{R}^3$ are linear forces at the human hand, and $\mathbf{J}_H(\mathbf{q}) \in \mathbb{R}^{3 \times 4}$ is the standard Jacobian matrix of (10) with respect to \mathbf{q} .

Task-specific knowledge can be used to model \mathbf{x}^H or \mathbf{q} . In most cases, we simplify the human kinematics by assuming the internal/external rotation of the arm is zero, $q_3 = 0$, and that $\mathbf{x}^H = \mathbf{x}$, i.e., the human hand is at the robot TCP, allowing us to write closed-form inverse kinematics and evaluate (11).

3.4. Inference of mode n

For tasks which have discrete variation in human goal, denote the current goal $n \in [1, \dots, N]$. When a GP is fit for each mode, then n can be inferred online by Bayesian estimation. Let the belief in mode n at time t be denoted $\mathbf{b}_t \in \mathbb{R}^N$, $\mathbf{b}_t[n] = p(n | \mathbf{f}_{1:t}, \mathbf{x}_{1:t})$, and $\bullet_{1:t} = [\bullet_1, \dots, \bullet_t]$. This belief is updated recursively as

$$\mathbf{b}_{t+1}[n] \propto \mathbf{b}_t[n] \det(2\pi \boldsymbol{\Sigma}_n^f) \exp((\mathbf{f}_t - \boldsymbol{\mu}_n^f)^T (\boldsymbol{\Sigma}_n^f)^{-1} (\mathbf{f}_t - \boldsymbol{\mu}_n^f)), \quad (12)$$

where $\boldsymbol{\mu}_n^f$ and $\boldsymbol{\Sigma}_n^f$ are the mean and covariance of the GP model in mode n evaluated at \mathbf{x}_t , and \mathbf{b}_t is normalized after the update such that $\sum \mathbf{b}_t[n] = 1$.

4. MPC problem

This section presents the MPC optimization problem: the constraints, cost functions, and multiple shooting formulation. This aspect is highlighted by the green part in Fig. 2.

4.1. Optimization problem

The MPC problem is formulated using multiple-shooting transcription with a generic problem statement of

$$\begin{aligned} \mathbf{u}_{t:t+H} = & \text{Section 4.1} \\ & \arg \min_{\mathbf{u}} \sum_t^{t+H} \sum_{n=1}^N b_t[n] c_n(\xi_t^n, \mathbf{u}_t) \quad \text{Section 4.2} \\ \text{s.t. } & \forall n \in [1, \dots, N], \tau \in [t, \dots, t+H-1]: \\ & \boldsymbol{\mu}_\tau^n = \xi_\tau, \boldsymbol{\Sigma}_\tau^n = \mathbf{0} \\ & |\boldsymbol{\mu}_{\tau+1}^n - \mathbf{f}(\boldsymbol{\mu}_\tau^n, \mathbf{u}_\tau)| \leq \varrho \quad (7) \\ & |\boldsymbol{\Sigma}_{\tau+1}^n - \mathbf{g}(\boldsymbol{\mu}_\tau^n, \boldsymbol{\Sigma}_\tau^n, \mathbf{u}_\tau)| \leq \varrho \quad (9) \\ & \mathbf{h}(\xi_t, \mathbf{u}_t) \geq 0 \quad \text{Section 4.3} \\ & \mathbf{u} \in U \end{aligned} \quad (13)$$

where H is the planning horizon, U is the range of allowed inputs, ϱ the slack for the continuity constraints (the inequality is applied element-wise), f is from (7), g following (9), and \mathbf{h} are inequality constraints to be presented in 4.3. The cost function c_n is shown with its parameter argument ξ_t and decision variable \mathbf{u} , the detailed cost function provided in (4.2.1). Initial covariance is here simplified by setting $\boldsymbol{\Sigma}_t = \mathbf{0}$ as it is assumed that the robot state is fully observed, but this can be initialized with, e.g. the covariance from an observer.

The constraints in (13) are nonlinear, so an interior-point nonlinear optimization solver is used. While nonlinear, the problem is written in an automatic differentiation framework, allowing calculation of the gradient and Hessian of the objective and constraints, significantly improving convergence.

The MPC framework here allows for different choices of decision variable \mathbf{u} , which can include any of the following:

$\mathbf{f}_{t:t+H}^r$ Robot trajectory
 Δ_t^M, Δ_t^D Change in robot impedance

where Δ_t^M and Δ_t^D are a change in the impedance parameters \mathbf{M} and \mathbf{D} .

4.2. Costs

4.2.1. Stage cost function c_n

The general stage cost function is defined as

$$c_n(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\mu}^{f,n}, \boldsymbol{\Sigma}^{f,n}, \boldsymbol{\tau}_H, \mathbf{u}) = \boldsymbol{\mu}^T \mathbf{Q}_\mu \boldsymbol{\mu} + \text{Tr}(\mathbf{Q}_\Sigma \boldsymbol{\Sigma}) + (\boldsymbol{\mu}^{f,n})^T \mathbf{Q}_f \boldsymbol{\mu}^{f,n} + \text{Tr}(\mathbf{Q}_{\Sigma_f} \boldsymbol{\Sigma}^{f,n}) + \boldsymbol{\tau}_H^T \mathbf{Q}_\tau \boldsymbol{\tau}_H + \mathbf{u}^T \mathbf{Q}_u \mathbf{u}, \quad (14)$$

where $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are mean and covariance for the predicted state $\boldsymbol{\xi}$ in mode n , $\boldsymbol{\mu}^{f,n}$ and $\boldsymbol{\Sigma}^{f,n}$ are the mean and covariance of the GP forces in mode n , $\boldsymbol{\tau}_H$ are the human joint torques from (11), \mathbf{u} are control inputs, and \mathbf{Q}_\cdot are the related weight-matrices for each cost factor. State cost \mathbf{Q}_μ includes only a small penalty on velocity (position terms are 0), the trajectory following is implicit in GP mean and covariance. In some applications (e.g. the rail assembly in Fig. 6), replacing $(\boldsymbol{\mu}^{f,n})^T \mathbf{Q}_f \boldsymbol{\mu}^{f,n}$ with $(\boldsymbol{\mu}^{f,n} + \mathbf{f}_t^r)^T \mathbf{Q}_H (\boldsymbol{\mu}^{f,n} + \mathbf{f}_t^r)$ is used to match the human demonstrations.

4.2.2. Human joint torque cost

Human arm manipulability plays a key role in ergonomics [28,30], as it relays the information about how well the joint torques can be transferred into the end-effector force where the task is produced. Therefore, we defined the ergonomic cost as

$$\boldsymbol{\tau}_H^T \mathbf{Q}_\tau \boldsymbol{\tau}_H = \mathbf{f}_H^T \mathbf{J}_H(\mathbf{q}) \mathbf{Q}_\tau \mathbf{J}_H(\mathbf{q})^T \mathbf{f}_H, \quad (15)$$

where, when $\mathbf{Q}_\tau = \mathbf{I}$, $\mathbf{J}_H \mathbf{J}_H^T$ defines the standard manipulability ellipsoid. This is a general formulation and the composition of the human end-effector force \mathbf{f}_H can be specified based on the task. For example, in the case of a polishing task, the structure to be polished itself supports the gravity force, which thus does not play a key role. Here, the dominant force component is related to the movement of the robot along a path, therefore we use the damping forces of the robot $\mathbf{f}_H = \mathbf{D}\dot{\mathbf{x}}$, where $\dot{\mathbf{x}}$ provides a weighting according to the direction of the expected forces. This term explicitly includes the impedance parameter, allowing the damping to be optimized to direction of expected motion and ergonomics. Nevertheless, the proposed formulation is general and \mathbf{f}_H can include gravity force and other static forces for tasks such as carrying and drilling.

4.2.3. Disturbance rejection

In many tasks, certain forces are *external perturbations* and should not influence the robot position, e.g., forces arising from the rotation of a polishing disk in a polishing task. In some cases, these forces may be more easily distinguished in the frequency domain, where typical human forces are low-frequency, and disturbance forces may be medium- to high-frequency. Frequency-dependent costs can be considered through an augmented state model [71], but this increases state dimension (and thus computational complexity) and requires the cost can be expressed from

the inputs or state. Here, measured force is not an input to the MPC problem, thus a new approach is taken.

We consider a disturbance force model in Laplace domain as $F^d(s) = \frac{\alpha s}{s + \omega_p}$, where α and ω_p adjust the magnitude and cut-off frequency. This model is a simple high-pass filter, distinguishing from lower-frequency human input. This disturbance results in robot position deviation according to the admittance dynamics $X^d(s) = A(s)F^d(s)$, where $A(s) = (M_i s^2 + D_i s + K_i)^{-1}$. Recall the equivalency of the two-norm in time and Laplace domain, $\|x^d(t)\|_2^2 = \|X^d(s)\|_2^2$ by Parseval's theorem [72], and that $\|X^d(s)\|_2^2$ can be found by the observability grammian as

$$\|X^d(s)\|_2^2 = B_d^T X_0 B_d, \quad (16)$$

$$\text{where } A_d X_0 + X_0 A_d^T = C_d^T C_d, \quad (17)$$

where A_d, B_d, C_d are a state-space realization of $X^d(s)$ [72]. The Lyapunov equation in (17) is differentiable [73] with respect to admittance parameters M_i, D_i, K_i , allowing the efficient calculation of the Jacobian and Hessian of this objective for the optimization problem. Note that the naïve solution of (17), by vectorizing and a linear solve may have numerical issues when the matrices are poorly conditioned, in such case using specialized solvers like Bartels–Stewart can improve performance.

4.3. Constraints

Impedance parameters affect contact safety. To improve contact safety, constraints for well-damped contact and force limit are formulated.

4.3.1. Force limit chance constraint

A fundamental concern in contact is avoiding excessive force; with a model of force over robot pose, this constraint can be directly transcribed. This force model may have uncertainty arising from:

1. Variation in *environment constraint* location
2. Exploration of a new region or *confidence in the skill*

Both these uncertainties are captured in the GP regression here. For (1), the GP will fit hyperparameters which increase the uncertainty in the force model (validated in Section 5.4). For (2), when farther from the training data, the GP will revert to a baseline variance given by hyperparameter σ_n^2 .

A chance constraint is a constraint applied to a random variable which holds for a specified degree of certainty. For a scalar Gaussian variable $f \sim \mathcal{N}(\boldsymbol{\mu}^f, \boldsymbol{\Sigma}^f)$, a chance constraint can be re-written over mean and covariance as

$$\text{Prob}\{f \leq \bar{f}\} > 1 - \epsilon \iff \boldsymbol{\mu}^f < \bar{f} - \text{erf}^{-1}(1 - \epsilon) \|\sqrt{\boldsymbol{\Sigma}^f}\|, \quad (18)$$

where \bar{f} is an upper bound on the force, erf^{-1} is the inverse of the error function and ϵ is a tuning parameter for the confidence that the chance constraint should hold [74].

It can be seen that the uncertainty adds a conservatism to the force bound; as $\boldsymbol{\Sigma}^f$ gets larger, the mean force $\boldsymbol{\mu}^f$ must be kept smaller. This constraint is applied element-wise on the linear forces, evaluating the GP at the final MPC trajectory point $\boldsymbol{\xi}_{t+H}$ to reduce computational complexity.

4.3.2. Well-damped contact constraint

One challenge in admittance control is contact stability with *environment constraints*, where a higher-stiffness environment typically requires higher virtual damping [75]. Ad hoc methods can increase damping near contact [76], but make assumptions such as the robot is slowing down before contact.

If we assume that the forces in demonstration were locally dominated by a stiffness of $\mathbf{f} = \mathbf{K}_e(\mathbf{x} - \mathbf{x}_0)$, where \mathbf{x}_0 is the rest

position of the environmental stiffness (i.e. contact position), and further assume that this function is well-modeled by the posterior GP mean function $\mu^f(\mathbf{x})$, we can find

$$K_{e,i}(\mathbf{x}_t) \approx \left. \frac{d\mu_i^f(\mathbf{x})}{dx_i} \right|_{\mathbf{x}=\mathbf{x}_t}, \quad (19)$$

where $K_{e,i}$ is the estimated environment stiffness in the i^{th} direction at pose \mathbf{x}_t . With this estimated stiffness, the dynamic parameters can be directly constrained using the condition for well-damped 2nd order systems

$$D_i \geq 2\zeta \sqrt{M_i K_{e,i}(\mathbf{x}_t)}, \quad (20)$$

where $\zeta = 1$ is critically-damped, and $\zeta > 1$ introduces a safety margin which can help account for the non-idealized admittance dynamics. This constraint is applied to the entire MPC trajectory $\xi_{t+1:t+H}$.

5. Analysis

This section, shown in purple in Fig. 2 shows some properties of the MPC problem: how trajectory and impedance optimization differ, necessary conditions for convexity, and investigates the ability to model contact with GPs.

5.1. Optimization of impedance vs trajectory

Why should we optimize impedance parameters ϕ instead of just the trajectory μ_\bullet ? Taking the differential of the linear-Gaussian dynamics (7) and (9),

$$\begin{bmatrix} d\mu^+ \\ d\sigma^+ \end{bmatrix} = \begin{bmatrix} \mathbf{B}_\phi & d_\phi \mathbf{A}_\phi \mu + d_\phi \mathbf{B}_\phi \mu^f \\ 0 & d_\phi (\mathbf{A}_\phi \otimes \mathbf{A}_\phi^T) \sigma + d_\phi (\mathbf{B}_\phi \otimes \mathbf{B}_\phi^T) \sigma^f \end{bmatrix} \begin{bmatrix} d\mathbf{f}^r \\ d\phi \end{bmatrix}, \quad (21)$$

where prefix d indicates total differential and d_\bullet the partial derivative with respect to \bullet . We note two properties: (i) input \mathbf{f}^r does not affect the covariance σ^+ , and (ii) state μ^+ is determined by both control input \mathbf{f}^r and impedance parameters ϕ , i.e., different combinations of force and impedance parameters can realize a certain change in μ^+ , as previously noted in [54].

When the robot changes admittance, the mean and variance of the robot's trajectory are affected. However, changing the robot's force reference only affects the mean trajectory. Additionally, a change in trajectory can be realized by a combination of force and admittance, thus the optimization of trajectory alone is underdetermined. Empirically, this leads to variation in optimized admittance parameters, requiring additional regularization. Thus, explicit consideration of trajectory covariance is important for impedance, and the MPC trajectories should be stochastic when optimizing impedance.

5.2. Convexity over impedance parameters ϕ

For the optimization problem to be efficient, the objective should be convex with respect to the decision variables [77]. As optimizing with respect to parameters of the dynamics (e.g., impedance parameters) is not studied, we check the convexity of a toy cost function with respect to impedance parameters, denoted $\phi = [\mathbf{M}, \mathbf{D}]$. As impedance parameters ϕ mostly affect covariance, we consider a simplified objective in a single step shooting problem of $\min_\phi \text{Tr}(\mathbf{Q}\Sigma^+)$, where by results from [70],

$$\begin{aligned} \text{Tr}(\mathbf{Q}\Sigma^+) &= \text{Tr}(\mathbf{Q}\mathbf{A}_\phi \Sigma \mathbf{A}_\phi^T + \mathbf{Q}\mathbf{B}_\phi \Sigma^f \mathbf{B}_\phi^T) \\ &= \sigma^T \text{vec}(\mathbf{A}^T \mathbf{Q} \mathbf{A}) + (\sigma^f)^T \text{vec}(\mathbf{B}^T \mathbf{Q} \mathbf{B}) \\ &= (\sigma^T (\mathbf{A}^T \otimes \mathbf{A}^T) + (\sigma^f)^T (\mathbf{B}^T \otimes \mathbf{B}^T)) \text{vec}(\mathbf{Q}). \end{aligned} \quad (22)$$




Fig. 4. Minimum eigenvalue of the Hessian $\nabla^2 \text{Tr} \Sigma^+$ and $\nabla^2 \text{In det} \Sigma^+$ with respect to impedance parameters as Σ_f varies, where the minimum eigenvalue must be positive for this objective to be convex. The integrator schemes (Euler/implicit/exponential) and objective (Tr vs In det) affect the convexity, where the implicit and exponential integrator with trace performs best over a range of force uncertainty.

The convexity of this objective will depend on the impedance parameter values, the choice of the integrator (which affects \mathbf{A} and \mathbf{B}), and the relative magnitude of σ and σ^f . Note that $\mathbf{B}^T \otimes \mathbf{B}^T = [0, 0, 0, T_s \mathbf{M}^{-2}]$, which gives $\nabla_{\mathbf{M}}^2 \mathbf{B}^T \otimes \mathbf{B}^T \geq 0$. Thus increasing Σ^f can increase the minimum eigenvalue of the Hessian of the objective (22). The terms for damping are more complex, thus numerical analysis is carried out.

These results are validated numerically to see what matrix cost and integrator have better convexity. The minimum eigenvalue of the Hessian $\nabla_{\phi}^2 \text{Tr}(\mathbf{Q}\Sigma^+)$ and $\nabla_{\phi}^2 \text{In det}(\mathbf{Q}\Sigma^+)$ are compared for three different integrator schemes: explicit Euler $\dot{x}^+ = (1 + T_s D/M)\dot{x}$, implicit $\dot{x}^+ = M(M + T_s D)^{-1}\dot{x}$, and exponential $\dot{x}^+ = \exp(-T_s D/M)\dot{x}$. We can see in Fig. 4 that the minimum eigenvalue generally increases as $|\Sigma^f|$ increases. Using a $\text{Tr}(\mathbf{Q}\Sigma^+)$ is unilaterally better than $\text{In det}(\mathbf{Q}\Sigma^+)$, and is thus used in the cost function here. The implicit integrator has better average performance over the range of Σ^f —notably, the performance of the explicit Euler integrator is poor; the Hessian is indefinite for typical parameter values. Thus, here we use the trace of matrix costs with an implicit integrator. Also note that convexity improves as Σ^f increases, as suggested in the above analysis.

5.3. Computational costs

One limitation of the GP approach is the higher online computational costs. The direct evaluation of a GP with a squared exponential kernel has computational complexity of M^2 , where M is the number of data points used in the dataset [69]. However, the overall computational efficiency of an optimization problem also depends on the gradient evaluation efficiency, in addition to the effects of the GP on the loss landscape, a high complexity which requires empirical analysis [74].

To do this empirical analysis, we consider a fixed MPC problem, using the parameters used in Section 6.2, to compare the effect of the number of GP points M and the number of modes N on the average solve time, evaluated offline over a replay of the same trajectory. The results can be seen in Fig. 5, along




Fig. 5. Computational cost as number of GP points M or number of modes N varies, with default parameters from Section 6.2 and baseline $M = 50$, $N = 1$. The fit is approximately quadratic and linear, respectively.

with a polynomial fit. As the number of GP points increases, the average MPC solve time increases quadratically. Furthermore, as the number of modes increases, the average MPC solve time increases linearly. This suggests that the number of GP points is critical, and the number of points needed to describe a task may vary on its complexity and size within the state space.

5.4. Contact modeling

We investigated the ability of the GPs to model contact by verifying (1) that GPs can identify the contact in human-guided contact scenarios, (2) the impact of variation in the contact position, and (3) the impact of environment stiffness on the ability to regress the models. These together allow the direct modeling of *environment constraint* uncertainty in realistic contact. The hyperparameters in the following were all fit to minimize negative log-likelihood from the same initial values.

When the human guides the robot into contact, the sum of human and environmental forces is measured. We collect human-guided contact and autonomous contact data (i.e. the robot makes contact without human forces), and compare the GPs which are fit to both of them. In Fig. 6(a), we see that both models identify the contact location with the hinge mean function, but the human

data identifies a lower stiffness and a contact position 1 cm too high. In the applications here, this error was acceptable, but this is application-specific.

Next, the contact position was varied between trials, with a total variation of 2 cm over three trials. The result can be seen in Fig. 6(b), where the three contact positions can be seen and the higher covariance in the fit GP (green) can be seen compared to the constant contact position (blue).

Finally, contact with environments of various stiffnesses was made, with corresponding stiffness of [15, 45, 126] N/mm. The resulting quality of the fit can be seen in Fig. 6(c), where the hinge function can readily identify the changes in contact position and stiffness.

6. Validation

This section presents four sets of experiments to validate the various features of the proposed framework. These validation experiments should provide useful pointers for the application of the proposed method in real-world scenarios (the blue part in Fig. 2). Videos are attached as multimedia and can be found at <https://youtu.be/Of2O3mHfM94>. MPC and GP parameters used for each experiment can be found at the repository <https://gitlab.cc-asp.fraunhofer.de/hanikevi/gp-mpc-impedance-control>.

6.1. Contact uncertainty and stiffness

We first used a contact task to validate safe contact with the environment. In particular, we examined the effects of force chance constraint and well-damped constraint as contact variance or stiffness increases. The robot was taught a task moving from free space into contact, which resulted in GP models as seen in Section 5.4. The contact location could be varied by adding blocks (Fig. 6(b)), and the environment made stiff by locking out a compliant element (Fig. 6(c)). For each contact condition – soft, high variance, and stiff – the corresponding model from Section 5.4 was used, and only the MPC parameters relating to the chance constraint ($\bar{f} = 12$, $\epsilon = 0.5$, chosen by hand tuning) and well-damped constraint ($\xi = 1.2$) are changed. The low impedance test condition kept the minimum mass 5 Kg and damping 500 Ns/m.

The force trajectories during contact are shown in Fig. 7, averaged over three trials, where variance is indicated by the




Fig. 6. Mean and covariance of GP models for various contact situations where contact occurs on the right side at -0.34 m, real data in dots and corresponding model in matching color. Learning from human-guided contact (a) introduces acceptable error in contact location and stiffness, (b) variation in contact location induces higher variance GP models, and (c) the hinge function can represent a range of environment stiffnesses.

Table 2
Objectives and constraints of the MPC framework, with definition, validation, and average MPC loop rate as evaluated by playing back the data recorded from the rail assembly problem.

	Feature	Functionality	Defined	Experiment	Rate (Hz)
	Baseline	–			26.2
Human	Discrete modes	Human goal estimation		Dual peg-in-hole Section 6.1	10.6
	Human joint torque	Human ergonomics	(11)	Polishing Section 6.3	16.6
	Model variance	Repeatable task DOF Constraint uncertainty	(1) (1)	Rail Section 6.2 Contact Section 6.1	24.2
Task	Disturbance reject	Task performance	(16)	Polishing Section 6.3	14.6
	Well-damped	Contact stability	(20)	contact Section 6.1	22.2
	Chance constraint	Safety	(18)	Contact Section 6.1	13.0




Fig. 7. Contact forces under contact conditions of soft (top, 45 N/mm), higher variance (middle, ± 1 cm) and stiff (bottom, 126 N/mm). The two contact constraints are compared, chance constraint (18) and well-damped (20). When the contact location has higher variance, the chance constraint is more conservative (middle). The well-damped constraint stabilizes contact with the high-stiffness environment by increasing damping (bottom), but does not interfere with the low-stiffness contact where increased damping is not needed (top).

shaded area. In the soft contact condition, the difference is small as the constraints were not violated. As the variance in the environment increased, the low impedance MPC had a higher contact force, but both constraints reduced this peak force. The chance constraint had more oscillation, but the well-damped constraint was smoother. As the stiffness of the environment increased, the low impedance MPC had a higher peak contact force, and had unstable contact (i.e., lost contact due to oscillation). The well-damped constraint had more stable contact with fewer oscillations, compared to the chance-constrained MPC.

In the supplementary video, it can be seen that these constraints also affected the behavior approaching contact. Once the chance constraint was activated, the robot pulled back from the contact approach. This was typically accompanied by infeasible optimization problems (i.e., no solution meeting constraints could be found). This problem was empirically more pronounced when using a GP with a hinge function, thus the poor conditioning of stiff contact was identified as a contributing cause. Alternatively, the well-damped constraint paused to allow the damping to increase before making contact. Also, the well-damped constraint was found empirically to have lower computational cost, with an average MPC rate of 22 compared with 13 for the chance constraint (Table 2).




Fig. 8. Rail assembly and polishing tasks, where the DOF which the human should determine are shown in green.




Fig. 9. Impedance values in the rail assembly task, showing principle axes of damping and mass in task space along three trajectories. The DOF with task variation (green arrow) has consistently lower damping values, allowing easier human manipulation along this axis.

6.2. Rail assembly

Here we present a rail assembly task, where the location of the switch along the rail may vary according to factors not known to the robot, such as existing parts. However, the rail mounting itself is repeatable. This task was taught in three demonstrations from various initial conditions, with varying goal locations along the rail. The desired behavior is that this division of repeatable task DOF and variation can be automatically extracted from data, such that the robot allows the human to more easily manipulate this DOF, i.e., has a lower impedance in the DOF that varies.

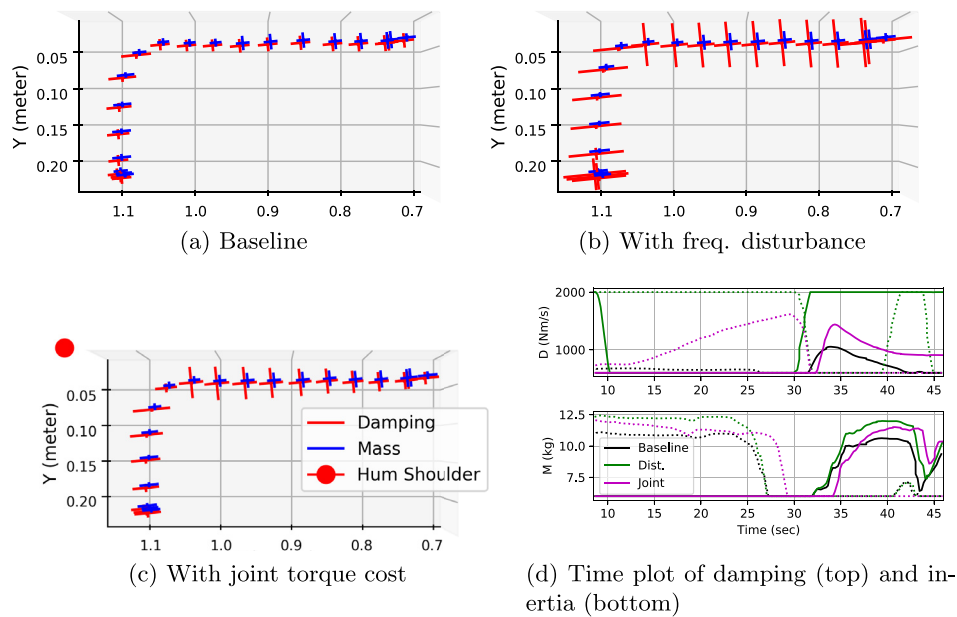


Fig. 10. Impedance values in the polishing task, showing that (b) frequency disturbance increases the mass across from the path, and (c) the human joint torque cost reduces impedance when farther from human, where more joint torque would be required. In (d) impedance over time is shown, where dotted is y-axis, solid x. The transition at 31 seconds corresponds to the corner in the trajectory.

Table 3

RMS velocity below and above 15 Hz, approximating the intentional motion and vibration, respectively. The high impedance and disturbance term suppress high-frequency vibration, but the high impedance results in a slow system.

MPC type	High imp.	Baseline	Dist.	Dist. + Jt.
High-freq RMS	5.4e-6	7.3e-6	6.5e-6	7.6e-6
Low-freq RMS	0.8e-4	4.2e-4	4.2e-4	4.1e-4

The results can be seen in Fig. 9, which shows the linear impedance parameters of mass and damping over space in three validation experiments. It can be seen that the impedance initialized at low values, but as the rail was approached, the impedance in the repeatable DOF was increased substantially, while kept low in the DOF which varied. The supplementary video shows that this results in a human easily intervening to adjust position along the rail. The task could also be achieved autonomously after the demonstrations when a cost was added to keep a positive force reference in the approach direction $\|15 - f_2^r\|$.

6.3. Polishing

We also examined a polishing task where the human had to guide a robot-supported polishing spindle along a path. In this case, the human was determining speed along the path, while polishing vibration should not affect the robot’s trajectory. To balance the rejection of polishing force disturbance against the need for the human to easily manipulate the robot, a force disturbance cost was added to the MPC problem. An ergonomic cost is also added, based on the assumed human arm geometry, and tested. The supplementary video shows polishing on a flat surface, corresponding to Fig. 8(b), and a round object to show the ability to handle orientation.

As seen in Fig. 10, the different cost function terms affected the rendered impedance along the polishing path. When the frequency-domain disturbance was added, the directionality of the impedance was significantly affected, resulting in a higher impedance cross to the direction of motion, reducing oscillation cross from the direction of motion. This resulted in a reduction of high-frequency velocity, as seen in Fig. 11 and Table 3. Adding




Fig. 11. Velocity realized in polishing experiments, colors match coordinate system in Fig. 8(b). The standard cost function is compared with adding disturbance cost. It can be seen that the disturbance result decreases the high-frequency (> 10 Hz) velocity of the robot.

a penalty on human joint torques reduced the impedance when farther from the human, as an equivalent force at a greater distance requires more shoulder joint torques due to poor force manipulability. However, the human joint cost did not improve the vibration suppression as significantly as the disturbance cost as seen in Table 3.

6.4. Double peg-in-hole

This task involved multiple possible goals from the human, with the two goals indicated in Fig. 12(a). While the human objective was not known in advance, the belief was updated




Fig. 12. Dual peg-in-hole task.

online as seen in Fig. 12(b). This update in the belief resulted in different force trajectories, as seen in Fig. 12(c). This allowed the planned robot trajectory to respond online to changes in the human intent, providing 6-DOF assistance to the human.

7. Conclusion

This paper presented an MPC framework for impedance control in contact with humans and the environment. The use of Gaussian Processes to model external force allows explicit consideration of *variability in demonstrations* (rail assembly), *confidence in the skill* (higher covariance outside demonstrations), and uncertain *environment constraints* (Fig. 6(b)). By using a GP model per human goal, a discrete goal can be estimated online to handle uncertainty in *human goal*. This approach can also consider frequency-domain disturbance models for *external perturbations*. By using stochastic trajectories, the impedance parameters can be directly optimized. Furthermore, it is empirically found that trajectory covariance in the cost function is helpful for admittance parameters to converge.

While flexible, the approach has some limitations, many arising from contact. Learning contact forces from human-led demonstrations is limited as human forces are not distinguished from the environment—autonomous data collection is also needed. While the hinge function in GPs allows better modeling of stiff contact, the integration over these stiff contact models requires a small step size, even with implicit integrators. Without suitable integrator parameters, the contact may be lost in the future trajectory and constraints such as chance or well-damped lose effectiveness. These inequality constraints are effectively solved in an interior-point method, but care must be taken to keep the problem feasible, e.g., allowing fast changes in damping to react to fast changes in projected environment stiffness.

Further limitations arise from the GP models, which cannot model certain types of uncertainty—this includes variation in noise over the state space, variation in human behavior due to changes in the robot policy, or correlations between DOF in the force signals. Variation in demonstrations is only captured through optimization of the hyperparameters, which can only affect the GP covariance per DOF, not locally. Heteroscedastic GPs, or fitting non-diagonal elements can address these issues, but computationally-efficient methods are needed. Optimization of the GP parameters needs strong regularization, especially to keep comparable between DOF or different GP models. Additionally, GPs as implemented here cannot capture frequency-domain information, so adapting to suppress vibration required a separate model.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Kevin Haninger reports financial support was provided by European Research Council.

Data availability

Code and data in Gitlab link.

Acknowledgments

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement Nos 820689–SHERLOCK and 101058521–CONVERGING.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.robot.2023.104431>.

References

- [1] N. Hogan, Impedance control: An approach to manipulation: Part II—implementation, *J. Dyn. Syst. Meas. Control* 107 (1) (1985) 8–16.
- [2] S. Calinon, I. Sardellitti, D.G. Caldwell, Learning-based control strategy for safe human-robot interaction exploiting task and robot redundancies, in: *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS*, IEEE, 2010, pp. 249–254.
- [3] K. Kronander, E. Burdet, A. Billard, Task transfer via collaborative manipulation for insertion assembly, in: *Workshop on Human-Robot Interaction for Industrial Manufacturing, Robotics, Science and Systems*, Citeseer, 2014.
- [4] G. Franzese, A. Mészáros, L. Peternel, J. Kober, ILoSA: Interactive learning of stiffness and attractors, in: *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS*, IEEE, 2021, pp. 7778–7785.
- [5] D.S. Walker, J.K. Salisbury, G. Niemeyer, Demonstrating the benefits of variable impedance to tele robotic task execution, in: *2011 IEEE International Conference on Robotics and Automation, IEE*, 2011, pp. 1348–1353.
- [6] A. Ajoudani, N. Tsagarakis, A. Bicchi, Tele-impedance: Teleoperation with impedance regulation using a body-machine interface, *Int. J. Robot. Res.* 31 (13) (2012) 1642–1656.
- [7] L. Peternel, T. Petrič, J. Babič, Robotic assembly solution by human-in-the-loop teaching method based on real-time stiffness modulation, *Auton. Robots* 42 (1) (2018) 1–17.
- [8] C. Yang, G. Ganesh, S. Haddadin, S. Parusel, A. Albu-Schaeffer, E. Burdet, Human-like adaptation of force and impedance in stable and unstable interactions, *IEEE Trans. Robot.* 27 (5) (2011) 918–930.
- [9] A. Naceri, T. Schumacher, Q. Li, S. Calinon, H. Ritter, Learning optimal impedance control during complex 3D arm movements, *IEEE Robot. Autom. Lett.* 6 (2) (2021) 1248–1255.

- [10] V. Duchaine, C.M. Gosselin, General model of human-robot cooperation using a novel velocity based variable impedance control, in: EuroHaptics Conference, 2007 and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. World Haptics 2007. Second Joint, IEEE, 2007, pp. 446–451.
- [11] L. Peternel, N. Tsagarakis, D. Caldwell, A. Ajoudani, Robot adaptation to human physical fatigue in human–robot co-manipulation, *Auton. Robots* (2018) 1–11.
- [12] G. Franzese, A. Mészáros, L. Peternel, J. Kober, Ilosa: Interactive learning of stiffness and attractors, in: 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS, IEEE, 2021, pp. 7778–7785.
- [13] L. Roveda, J. Maskani, P. Franceschi, A. Abdi, F. Braghin, L.M. Tosatti, N. Pedrocchi, Model-based reinforcement learning variable impedance control for human-robot collaboration, *J. Intell. Robot. Syst.* (2020) 1–17.
- [14] K. Haninger, C. Hegeler, L. Peternel, Model predictive control with Gaussian processes for flexible multi-modal physical human robot interaction, 2021, arXiv preprint arXiv:2110.12433.
- [15] A. Takagi, G. Ganesh, T. Yoshioka, M. Kawato, E. Burdet, Physically interacting individuals estimate the Partner's goal to enhance their movements, *Nat. Hum. Behav.* 1 (3) (2017) 0054.
- [16] B. Corteveille, E. Aertbeliën, H. Bruyninckx, J. De Schutter, H. Van Brussel, Human-inspired robot assistant for fast point-to-point movements, in: Proceedings 2007 IEEE International Conference on Robotics and Automation, IEEE, 2007, pp. 3639–3644.
- [17] A. Ajoudani, A.M. Zanchettin, S. Ivaldi, A. Albu-Schäffer, K. Kosuge, O. Khatib, Progress and prospects of the human–robot collaboration, *Auton. Robots* (2017) 1–19.
- [18] P. Maurice, A. Malaisé, C. Amiot, N. Paris, G.-J. Richard, O. Rochel, S. Ivaldi, Human movement and ergonomics: An industry-oriented dataset for collaborative robotics, *Int. J. Robot. Res.* 38 (14) (2019) 1529–1537.
- [19] Y. Demiris, Prediction of intent in robotics and multi-agent systems, *Cognit. Process.* 8 (3) (2007) 151–158.
- [20] Y. Li, S.S. Ge, Human–Robot collaboration based on motion intention estimation, *IEEE/ASME Trans. Mechatronics* 19 (3) (2014) 1007–1014.
- [21] C. Wang, Y. Li, S.S. Ge, T.H. Lee, Reference adaptation for robots in physical interactions with unknown environments, *IEEE Trans. Cybern.* 47 (11) (2017) 3504–3515.
- [22] G. Kang, H.S. Oh, J.K. Seo, U. Kim, H.R. Choi, Variable admittance control of robot manipulators based on human intention, *IEEE/ASME Trans. Mechatronics* 24 (3) (2019) 1023–1032.
- [23] S. Jain, B. Argall, Recursive Bayesian human intent recognition in shared-control robotics, in: 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS, IEEE, 2018, pp. 3905–3912.
- [24] R.K. Jain, S. Majumder, B. Ghosh, S. Saha, Design and manufacturing of mobile micro manipulation system with a compliant piezoelectric actuator based micro gripper, *J. Manuf. Syst.* 35 (2015) 76–91.
- [25] A. Bajcsy, D.P. Losey, M.K. O'Malley, A.D. Dragan, Learning robot objectives from physical human interaction, 2017, p. 10.
- [26] C. Messeri, G. Masotti, A.M. Zanchettin, P. Rocco, Human-robot collaboration: Optimizing stress and productivity based on game theory, *IEEE Robot. Autom. Lett.* 6 (4) (2021) 8061–8068.
- [27] N. Vahrenkamp, H. Arnst, M. Wächter, D. Schiebener, P. Sotiropoulos, M. Kowalik, T. Asfour, Workspace analysis for planning human-robot interaction tasks, in: 2016 IEEE-RAS 16th International Conference on Humanoid Robots, Humanoids, IEEE, 2016, pp. 1298–1303.
- [28] S. Gopinathan, S. Otting, J. Steil, A user study on personalized adaptive stiffness control modes for human-robot interaction, in: The 26th IEEE International Symposium on Robot and Human Interactive Communication, 2017, pp. 831–837.
- [29] N. Mansfeld, M. Hamad, M. Becker, A.G. Marin, S. Haddadin, Safety map: A unified representation for biomechanics impact data and robot instantaneous dynamic properties, *IEEE Robot. Autom. Lett.* 3 (3) (2018) 1880–1887.
- [30] L. Peternel, D.T. Schøn, C. Fang, Binary and hybrid work-condition maps for interactive exploration of ergonomic human arm postures, *Front. Neurobot.* (2021) 114.
- [31] H. Ben Amor, G. Neumann, S. Kamthe, O. Kroemer, J. Peters, Interaction Primitives for Human-Robot Cooperation Tasks, IEEE, 2014, pp. 2831–2837.
- [32] M. Khoramshahi, A. Billard, A dynamical system approach to task-adaptation in physical human–robot interaction, *Auton. Robots* 43 (4) (2019) 927–946.
- [33] G. Raiola, S.S. Restrepo, P. Chevalier, P. Rodriguez-Ayerbe, X. Lamy, S. Tliba, F. Stulp, Co-manipulation with a library of virtual guiding fixtures, *Auton. Robots* 42 (5) (2018) 1037–1051.
- [34] L. Roza, S. Calinon, D.G. Caldwell, P. Jimenez, C. Torras, Learning physical collaborative robot behaviors from human demonstrations, *IEEE Trans. Robot.* 32 (3) (2016) 513–527.
- [35] Y. Maeda, T. Hara, T. Arai, Human-robot cooperative manipulation with motion estimation, in: Proceedings 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems. Expanding the Societal Role of Robotics in the the Next Millennium (Cat. No. 01CH37180), Vol. 4, Ieee, 2001, pp. 2240–2245.
- [36] F. Dimeas, N. Aspragathos, Reinforcement learning of variable admittance control for human-robot co-manipulation, in: Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on, IEEE, 2015, pp. 1011–1016.
- [37] E. Gribovskaia, A. Kheddar, A. Billard, Motion learning and adaptive impedance for robot control during physical interaction with humans, in: Robotics and Automation (ICRA), 2011 IEEE International Conference on, IEEE, 2011, pp. 4326–4332.
- [38] H. Lee, B. Lee, W. Kim, M. Gil, J. Han, C. Han, Human-robot cooperative control based on pHRI (Physical Human–Robot Interaction) of exoskeleton robot for a human upper extremity, *Int. J. Precis. Eng. Manuf.* 13 (6) (2012) 985–992.
- [39] T. Tsumugiwa, R. Yokogawa, K. Hara, Variable impedance control based on estimation of human arm stiffness for human-robot cooperative calligraphic task, in: Robotics and Automation, 2002. Proceedings. ICRA'02. IEEE International Conference on, Vol. 1, IEEE, 2002, pp. 644–650.
- [40] V. Duchaine, B.M. St-Onge, D. Gao, C. Gosselin, Stable and intuitive control of an intelligent assist device, *IEEE Trans. Haptics* 5 (2) (2012) 148–159.
- [41] X. Lamy, F. Colledani, F. Geffard, Y. Measson, G. Morel, Achieving efficient and stable comanipulation through adaptation to changes in human arm impedance, in: Robotics and Automation, 2009. ICRA'09. IEEE International Conference on, IEEE, 2009, pp. 265–271.
- [42] J.R. Koller, D.H. Gates, D.P. Ferris, C.D. Remy, 'Body-in-the-Loop' optimization of assistive robotic devices: A validation study., in: Robotics: Science and Systems, 2016.
- [43] J.R. Medina, D. Lee, S. Hirche, Risk-sensitive optimal feedback control for haptic assistance, in: Robotics and Automation (ICRA), 2012 IEEE International Conference on, IEEE, 2012, pp. 1025–1031.
- [44] S. Calinon, D. Bruno, M.S. Malekzadeh, T. Nanayakkara, D.G. Caldwell, Human–robot skills transfer interfaces for a flexible surgical robot, *Comput. Methods Programs Biomed.* 116 (2) (2014) 81–96.
- [45] E. Pignat, S. Calinon, Learning adaptive dressing assistance from human demonstration, *Robot. Auton. Syst.* 93 (2017) 61–75.
- [46] N. Hogan, On the stability of manipulators performing contact tasks, *Robot. Automat., IEEE J.* 4 (6) (1988) 677–686.
- [47] E. Burdet, R. Osu, D.W. Franklin, T.E. Milner, M. Kawato, The central nervous system stabilizes unstable dynamics by learning optimal impedance, *Nature* 414 (6862) (2001) 446–449.
- [48] Y. Li, G. Ganesh, N. Jarrassé, S. Haddadin, A. Albu-Schaeffer, E. Burdet, Force, impedance, and trajectory learning for contact tooling and haptic identification, *IEEE Trans. Robot.* 34 (5) (2018) 1170–1182.
- [49] F. Ferraguti, C. Talignani Landi, L. Sabatini, M. Bonfè, C. Fantuzzi, C. Secchi, A variable admittance control strategy for stable physical human–robot interaction, *Int. J. Robot. Res.* 38 (6) (2019) 747–765.
- [50] P. Kormushev, S. Calinon, D.G. Caldwell, Robot motor skill coordination with EM-based reinforcement learning, in: Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on, IEEE, 2010, pp. 3232–3237.
- [51] S. Calinon, D. Bruno, D.G. Caldwell, A task-parameterized probabilistic model with minimal intervention control, in: 2014 IEEE International Conference on Robotics and Automation, ICRA, IEEE, 2014, pp. 3339–3344.
- [52] Y. Li, J.-Y. Zhu, R. Tedrake, A. Torralba, Connecting touch and vision via cross-modal prediction, 2019, arXiv:1906.06322 [cs].
- [53] Y. Li, G. Carboni, F. Gonzalez, D. Campolo, E. Burdet, Differential game theory for versatile physical human–robot interaction, *Nat. Mach. Intell.* 1 (1) (2019) 36–43.
- [54] F. Stulp, J. Buchli, A. Ellmer, M. Mistry, E.A. Theodorou, S. Schaal, Model-free reinforcement learning of impedance control in stochastic environments, *IEEE Trans. Auton. Men. Dev.* 4 (4) (2012) 330–341.
- [55] C. Chang, K. Haninger, Y. Shi, C. Yuan, Z. Chen, J. Zhang, Impedance adaptation by reinforcement learning with contact dynamic movement primitives, in: 2022 IEEE Intl. Conf. on Advanced Intelligent Mechatronics, AIM, 2022.
- [56] T. Tang, H.-C. Lin, M. Tomizuka, A learning-based framework for robot peg-hole-insertion, in: Dynamic Systems and Control Conference, vol. 57250, American Society of Mechanical Engineers, 2015, V002T27A002.
- [57] E. Caldarelli, A. Colomé, C. Torras, Perturbation-based stiffness inference in variable impedance control, *IEEE Robot. Autom. Lett.* 7 (4) (2022) 8823–8830.
- [58] L. Peternel, N. Tsagarakis, A. Ajoudani, A Human–Robot co-manipulation approach based on human sensorimotor information, *IEEE Trans. Neural Syst. Rehabil. Eng.* 25 (7) (2017) 811–822.
- [59] M. Ulmer, E. Aljalbout, S. Schwarz, S. Haddadin, Learning robotic manipulation skills using an adaptive force-impedance action space, 2021, arXiv:2110.09904 [cs].
- [60] F.J. Abu-Dakka, M. Saveriano, Variable impedance control and learning – A review, 2020, arXiv:2010.06246 [cs].
- [61] M. Sharifi, A. Zakerimanesh, J.K. Mehr, A. Torabi, V.K. Mushahwar, M. Tavakoli, Impedance variation and learning strategies in human-robot interaction, *IEEE Trans. Cybern.* (2021) 1–14.

- [62] F. Ferraguti, N. Preda, M. Bonfe, C. Secchi, Bilateral teleoperation of a dual arms surgical robot with passive virtual fixtures generation, in: *Intelligent Robots and Systems (IROS)*, 2015 IEEE/RSJ International Conference on, IEEE, 2015, pp. 4223–4228.
- [63] M. Rubagotti, T. Taunyazov, B. Omarali, A. Shintemirov, Semi-autonomous robot teleoperation with obstacle avoidance via model predictive control, *IEEE Robot. Autom. Lett.* 4 (3) (2019) 2746–2753.
- [64] Z. Manchester, S. Kuindersma, Variational contact-implicit trajectory optimization, in: N.M. Amato, G. Hager, S. Thomas, M. Torres-Torriti (Eds.), in: *Robotics Research*, vol. 10, Springer International Publishing, Cham, 2020, pp. 985–1000.
- [65] J.-P. Sleiman, F. Farshidian, M.V. Minniti, M. Hutter, A unified mpc framework for whole-body dynamic locomotion and manipulation, *IEEE Robot. Autom. Lett.* 6 (3) (2021) 4688–4695.
- [66] D.E. Stewart, Rigid-body dynamics with friction and impact, *SIAM Rev.* 42 (1) (2000) 3–39.
- [67] M.V. Minniti, R. Grandia, K. Föh, F. Farshidian, M. Hutter, Model predictive robot-environment interaction control for mobile manipulation tasks, 2021, [arXiv:2106.04202](https://arxiv.org/abs/2106.04202) [cs].
- [68] S. Pfrommer, M. Halm, M. Posa, ContactNets: Learning of discontinuous contact dynamics with smooth, implicit representations, 2020, [arXiv:2009.11193](https://arxiv.org/abs/2009.11193).
- [69] C.E. Rasmussen, *Gaussian Processes for Machine Learning*, MIT Press, 2006.
- [70] K.B. Petersen, M.S. Pedersen, et al., *The Matrix Cookbook*, Vol. 7, Technical University of Denmark, 2008, p. 15.
- [71] R. Grandia, F. Farshidian, A. Dosovitskiy, R. Ranftl, M. Hutter, Frequency-Aware model predictive control, *IEEE Robot. Autom. Lett.* (2019) 1.
- [72] K. Zhou, J.C. Doyle, K. Glover, *Robust and Optimal Control*, Vol. 40, Prentice hall New Jersey, 1996.
- [73] T.-C. Kao, G. Hennequin, Automatic differentiation of Sylvester, Lyapunov, and algebraic Riccati equations, 2020, [arXiv:2011.11430](https://arxiv.org/abs/2011.11430) [cs, math].
- [74] L. Hewing, J. Kabzan, M.N. Zeilinger, Cautious model predictive control using Gaussian process regression, *IEEE Trans. Control Syst. Technol.* 28 (6) (2019) 2736–2743.
- [75] A. Albu-Schäffer, C. Ott, U. Frese, G. Hirzinger, Cartesian impedance control of redundant robots: Recent results with the dlr-light-weight-arms, in: *Robotics and Automation, 2003. Proceedings. ICRA'03. IEEE International Conference on*, Vol. 3, IEEE, 2003, pp. 3704–3709.
- [76] F. Ficuciello, L. Villani, B. Siciliano, Variable impedance control of redundant manipulators for intuitive human-robot physical interaction, *IEEE Trans. Robot.* 31 (4) (2015) 850–863.

- [77] J.B. Rawlings, D.Q. Mayne, M. Diehl, *Model Predictive Control: Theory, Computation, and Design*, Vol. 2, Nob Hill Publishing Madison, 2017.



Kevin Haninger researches physical, interactive systems, developing algorithms and design methods for impedance control, compliant mechanisms, and observers for contact-rich manipulation and human-robot interaction. He graduated with a Ph.D. in Mechanical Engineering from the University of California, Berkeley in 2016, and is currently affiliated with Fraunhofer IPK in Berlin, Germany.



Christian Hegeler is a researcher in the area of robot planning and control, developing data-efficient machine learning methods for robots to handle uncertain tasks. He has finished his Master's degree at the Technical University of Berlin in 2023.



Luka Peternel received a Ph.D. in robotics from the Faculty of Electrical Engineering, University of Ljubljana, Slovenia in 2015. He conducted my Ph.D. studies at the Department for Automation, Biocybernetics and Robotics, Jožef Stefan Institute in Ljubljana from 2011 to 2015, and at the Department of Brain-Robot Interface, ATR Computational Neuroscience Laboratories in Kyoto, Japan in 2013 and 2014. He was with the Human-Robot Interfaces and Physical Interaction Lab, Advanced Robotics, Italian Institute of Technology in Genoa, Italy from 2015 to 2018. Since 2019, he is an

Assistant Professor at the Department of Cognitive Robotics, Delft University of Technology in the Netherlands.