

Exploiting Variable Stiffness in Explosive Movement Tasks

David J. Braun, Matthew Howard and Sethu Vijayakumar

University of Edinburgh, School of Informatics, IPAB, Scotland, UK

E-mail: david.braun@ed.ac.uk, matthew.howard@ed.ac.uk, sethu.vijayakumar@ed.ac.uk

Abstract—It is widely recognised that compliant actuation is advantageous to robot control once high-performance, explosive tasks, such as throwing, hitting or jumping are considered. However, the benefit of intrinsic compliance comes with high control complexity. Specifically, coordinating the motion of the system through a compliant actuator and finding a task-specific impedance profile that leads to better performance is non-trivial. Here, we utilise optimal control to devise time-varying torque and stiffness profiles for highly dynamic movements in compliantly actuated robots. The proposed methodology is applied to a ball-throwing task where we demonstrate that: (i) the method is able to tailor impedance strategies to specific task objectives and system dynamics, (ii) the ability to vary stiffness leads to better performance in this class of movements, (iii) in systems with variable physical compliance, our methodology is able to exploit the energy storage capabilities of the actuators. We illustrate these in several numerical simulations, and in hardware experiments on a device with variable physical stiffness.

I. INTRODUCTION

Recently, significant research effort has focused on development of variable stiffness actuators (VSAs). Numerous designs for VSAs have been proposed, with the motivation of (i) improving safety of robots and humans (by providing an intrinsic compliance) [26], (ii) adding additional redundancy in the control (allowing robots to be stiff and accurate or compliant to the environment according to task demands) [12], and (iii) improving the energy efficiency and dynamic range of existing actuators (e.g., by exploiting the energy storage capabilities of such devices) [5].

Due to the last point, a particularly promising area in which VSAs may be deployed are applications involving highly dynamic, *explosive movements*. Such movements are characterised by a large release of energy over a short time frame, with typical examples including throwing, hitting, kicking, or jumping. Additionally, rhythmic movements as walking and running can also contain an explosive component [24]. Achieving such movements with traditional, joint torque actuators presents significant difficulties, particularly in respect of the size and power limitation of the motors. In contrast, by incorporating a physical elastic element, novel VSAs offer the possibility of achieving a much higher dynamic range, with much smaller, less expensive motors, by instead exploiting the compliant nature of the actuators.

However, one of the difficulties of using such actuators is the considerably increased complexity in the planning and control, particularly in the context of such highly dynamic movements. From the perspective of mechanical design, VSAs are often

developed with an antagonistic architecture [10, 19], but have also been developed using more practical (technical) objectives [21, 25]. Typically, this introduces additional complexity to the dynamics in the form of non-linearities, coupling of stiffness and motion characteristics, and increased dimensionality of the control input. As a result, it becomes increasingly difficult to design or hand-tune control strategies for such actuators, and to exploit the benefits of variable stiffness.

Hogan [4] suggested that if the objective is to minimise the tracking error and the interface force, then the manipulator impedance should be inversely proportional to the environmental impedance. This rule, by which the manipulator should act as the dual of the environment, is known as a *duality principle* in impedance control [2]. It is, however, not entirely clear how to select a desirable, possibly time-varying target impedance for a given non-linear system and a generic dynamic task. As such, it is usually the case that the control of the VSAs presented in the literature to date is realised with 'tuned' constant impedances that may not be optimal with respect to the task.

There is growing interest in addressing this limitation. For example, the benefit of optimal time-varying stiffness control has been investigated in [9]. The utility of damping variation was shown in [6], where a robot hand controlled by an optimal variable damper effectively supported a human in a cooperative lifting task. The optimisation of the passive system properties has also been discussed under more dynamic conditions. In this context, [22, 23], demonstrated that tracking a small amplitude oscillatory motion of a pendulum can be made efficient by matching the (constant) stiffness of the actuators with the natural stiffness of the reference trajectory. With a similar objective in [20], the authors combined trajectory tracking with adaptation of the constant joint stiffness to its optimal value.

In this paper, we investigate the use of model-based optimal control for devising strategies for VSAs in highly dynamic, explosive tasks. Specifically, we seek *optimal torques* and *time-varying optimal stiffness profiles* for an explosive task such as throwing. Our goal is to assess the benefits of variable stiffness for such tasks, and to investigate how, by taking an optimal control approach, we may exploit the physical properties of a variable impedance system to improve performance. Our approach has the benefit of providing behavioural predictions that are non-trivial to obtain by other methods or hand-tuning, for example, strategy changes due to modification

of the system dynamics. Furthermore, predictions made by our approach could be used as a tool to comprehend explosive task execution strategies in humans. We illustrate our approach with a number of simulation studies, and in a throwing experiment on robotic hardware.

II. OPTIMAL TORQUE/STIFFNESS CONTROL

In this section, we present an optimal torque and stiffness control methodology applied to compliantly actuated systems. The basic ingredients of such control approach are (i) the dynamic equation of the robot, (ii) the model of the compliant actuator, (iii) the physical constraints on the actuator inputs, and (iv) the performance index that defines the control task.

A. Robot dynamics

Consider an n -degree-of-freedom robotic system, the configuration of which is uniquely specified by $\mathbf{q} \in \mathbb{R}^n$ joint angles. Let the equation of motion of the system be represented as

$$\mathbf{M}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}}) + \mathbf{G}(\mathbf{q}) = \boldsymbol{\tau}(\mathbf{q}, \mathbf{u}) + \boldsymbol{\tau}_f, \quad (1)$$

where $\mathbf{M} \in \mathbb{R}^{n \times n}$ is a symmetric and positive definite mass matrix, $\mathbf{C} \in \mathbb{R}^n$ represents centrifugal and Coriolis terms, $\mathbf{G} \in \mathbb{R}^n$ are the gravitational terms, $\boldsymbol{\tau} \in \mathbb{R}^n$ are the joint torques from the actuators, $\mathbf{u} \in \mathbb{R}^m$ are the control inputs and $\boldsymbol{\tau}_f \in \mathbb{R}^n$ represents the torques due to joint friction (e.g., due to Coulomb and viscous friction). Usually, for industrial manipulators, $\boldsymbol{\tau}$ is considered to be the control input (i.e., $\boldsymbol{\tau} \equiv \mathbf{u}$) [17] since it can be directly commanded through stiff (highly-g geared) joint-torque actuators. In the present paper however, we consider robots equipped with compliant actuators, a model of which is introduced in the following.

B. Compliant actuation

If the actuators have in-built compliance, the joint torques $\boldsymbol{\tau}$ cannot be directly commanded. In such cases, the torque function is in general complicated, position dependent (see (1)), and can only be indirectly modulated with the control inputs \mathbf{u} . For example, a popular approach is to use servo-control to adjust the length and the moment-arm of the compliant element (i.e., spring) embedded in the actuator [21]. In the following, we consider a model of VSAs which allow simultaneous torque and stiffness control.

1) *Torque control*: On compliantly actuated systems, the relation between the joint torques, joint angles and the control inputs $\mathbf{u} \in \mathbb{R}^m$ ($m > n$), is given in the following form¹

$$\boldsymbol{\tau}(\mathbf{q}, \mathbf{u}) = \mathbf{A}^T(\mathbf{q}, \mathbf{u})\mathbf{F}(\mathbf{q}, \mathbf{u}), \quad (2)$$

where $\mathbf{A} \in \mathbb{R}^{p \times n}$ ($p \geq n$) is the moment-arm matrix, defined by the geometric attributes of the actuators, and $\mathbf{F} \in \mathbb{R}^p$ are the corresponding forces due to the elastic elements (characterised by the physical attributes of these elements). As indicated in (2), both the moment arm and the associated

¹While in the present paper, the actuator torque is assumed to be position dependent (as is the case in the majority of VSA actuators), the formulation remains valid for cases where the torque is velocity dependent (e.g., due to viscoelastic forces for example).

forces may explicitly depend on the control inputs \mathbf{u} . As will be subsequently discussed, this dependence may allow one to modulate not only the joint torques, but also the (passive) joint stiffness of the actuators.

2) *Stiffness control*: Using (2), the passive joint stiffness of the actuators $\mathbf{K} := -\partial\boldsymbol{\tau}/\partial\mathbf{q} \in \mathbb{R}^{n \times n}$, can be computed as

$$\mathbf{K}(\mathbf{q}, \mathbf{u}) = -\frac{\partial\mathbf{A}^T}{\partial\mathbf{q}}\mathbf{F} - \mathbf{A}^T\frac{\partial\mathbf{F}}{\partial\mathbf{q}}. \quad (3)$$

Relation (3) allows one to identify the conditions under which the control inputs \mathbf{u} can change the joint stiffness. Specifically, the first term in (3) indicates that the joint stiffness can be directly changed by controlled modulation of the elastic force \mathbf{F} , if the moment arm \mathbf{A} is position dependent (see the antagonistic design in [12] for example). If the moment arm is constant, then joint-level stiffness modulation requires a non-linear force-angle relation (i.e., the stiffness defined by $\mathbf{k}_F = -\partial\mathbf{F}/\partial\mathbf{q}$ should be control dependent). This is the mechanism used in many antagonistic actuators, see [10] for example. One can however modulate the joint stiffness through the second term, even if \mathbf{k}_F is constant. In that case however, the moment arm should be control dependent, see [21] for example. Regardless of which of these mechanisms is employed, we assume that the actuators under consideration have the necessary *control redundancy* (i.e., $m > n$) required for simultaneous torque and stiffness modulation on one or more of the joints.

3) *Admissible control actions*: In many compliant actuator designs control constraints exist. In order to formally introduce these constraints, we define the set of *admissible controls* as

$$U = \{\mathbf{u} \in \mathbb{R}^m : \mathbf{u} \in [\mathbf{u}_{min}, \mathbf{u}_{max}]\}, \quad (4)$$

where \mathbf{u}_{min} and \mathbf{u}_{max} is the lower and the upper bound on the control inputs respectively. This set defined with box-constraints is well suited to express the physical restriction inherently present on compliant actuators.

C. Optimal control formulation

An optimal control problem is defined with a performance criterion which is minimised (or maximised) with respect to the control actions. There are two types of physical constraints that apply to this minimisation in general. The first comes from the plant dynamics (1) and (2), which specifies the state output for a given control input, while the second is due to the physical restrictions on the realisable control actions (4). In the following, we introduce the computational framework used in this work and discuss application of the result to compliantly actuated robots in the context of optimal torque and stiffness control.

1) *Problem formulation*: Let us consider a state-space representation of the controlled system dynamics (1) and (2), given as

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}), \quad (5)$$

where $\mathbf{x} = (\mathbf{x}_1^T, \mathbf{x}_2^T)^T = (\mathbf{q}^T, \dot{\mathbf{q}}^T)^T \in \mathbb{R}^{2n}$ is a state vector, while

$$\mathbf{f} = \begin{bmatrix} \mathbf{x}_2 \\ -\mathbf{M}(\mathbf{x}_1)^{-1}(\mathbf{C}(\mathbf{x}) + \mathbf{G}(\mathbf{x}_1) - \boldsymbol{\tau}(\mathbf{x}_1, \mathbf{u}) - \boldsymbol{\tau}_f) \end{bmatrix}, \quad (6)$$

is a vector field which is in general non-linear in states \mathbf{x} and actions \mathbf{u} .

Within the framework of optimal control, a wide variety of tasks can be represented using the following formulation: For a given finite time interval $t \in [0, T]$, and for a given initial state of the system $\mathbf{x}(0) = \mathbf{x}_0$, find an admissible control law $\mathbf{u} = \mathbf{u}(t, \mathbf{x}) \in U$ that minimises the *optimisation criterion*

$$J(\mathbf{x}_0) = h(\mathbf{x}(T)) + \int_0^T c(\mathbf{x}(t), \mathbf{u}(t, \mathbf{x}(t))) dt \in \mathbb{R}, \quad (7)$$

where $h(\mathbf{x}(T)) \in \mathbb{R}$ is the *terminal cost*, while $c(\mathbf{x}, \mathbf{u}) \in \mathbb{R}$ is the *running cost* used to encode the control objectives within the formulation [15]. Once the control law is devised, the control sequence can be computed by state feedback as $\mathbf{u}_{opt}(t) = \mathbf{u}(t, \mathbf{x}(t))$, and applied to (5) for optimal task execution.

2) *Solution method*: For non-linear plant dynamics (5) and non-quadratic cost in (7), a globally valid optimal control law $\mathbf{u}_{opt} = \mathbf{u}(t, \mathbf{x})$ could be derived by means of *dynamic programming*². However such a solution method is computationally expensive (often intractable), and accordingly not attractive for practical implementations. In order to circumvent this issue, we utilise the iterative linear quadratic regulator (iLQR) method [8] to efficiently compute the optimal control actions for the non-linear optimal control problem (5) and (7), as described below.

The iLQR method is initialised with a nominal control sequence and the corresponding state trajectories $(\hat{\mathbf{x}}, \hat{\mathbf{u}})$. These are then iteratively improved by means of sequentially solving a set of local LQR sub-problems. The sub-problems are formed by linear approximation of the system dynamics (5),

$$\delta \dot{\mathbf{x}} = \mathbf{f}_x \delta \mathbf{x} + \mathbf{f}_u \delta \mathbf{u}, \quad (8)$$

and quadratic approximation of the objective functional (7),

$$\begin{aligned} \Delta J = & h_x^T \delta \mathbf{x}(T) + \delta \mathbf{x}^T(T) h_{xx} \delta \mathbf{x}(T) + \int_0^T (c_x^T \delta \mathbf{x} + c_u^T \delta \mathbf{u}) dt + \\ & \frac{1}{2} \int_0^T (\delta \mathbf{x}^T c_{xx} \delta \mathbf{x} + \delta \mathbf{x}^T c_{xu} \delta \mathbf{u} + \delta \mathbf{u}^T c_{uu} \delta \mathbf{u}) dt, \quad (9) \end{aligned}$$

evaluated along the current state and control sequence (lower indices in (8) and (9), denote partial derivatives with respect to the corresponding variables). This sub-problem, (equations (8) and (9)), is solved for $(\delta \mathbf{x}, \delta \mathbf{u})$ via a modified Ricatti-like system [8], and a new (improved) sequence is formed by $\hat{\mathbf{x}} \leftarrow \hat{\mathbf{x}} + \delta \mathbf{x}$ and $\hat{\mathbf{u}} \leftarrow \hat{\mathbf{u}} + \delta \mathbf{u}$.

When the method converges (i.e., $\Delta J \approx 0$ achieved numerically), it returns the optimal state and control trajectories $(\mathbf{x}^*(t), \mathbf{u}^*(t))$ together with a set of feedback-gains³ $\mathbf{L}^*(t) \in \mathbb{R}^{m \times 2n}$. The feedback control law for optimal task execution can then be defined as $\mathbf{u}_{opt}(t, \mathbf{x}) = \mathbf{u}^*(t) + \mathbf{L}^*(t)(\mathbf{x} - \mathbf{x}^*(t))$.

²This method requires one to solve the non-linear Hamilton-Jacobi-Bellman partial differential equation to define a *general solution* and the optimal feedback control law, see [18].

³If the control inputs are saturated (i.e., restricted with ‘‘hard constraints’’ (7)), the corresponding feedback gains may be set to zero, as proposed in [8]. Alternatively, one could utilise penalty terms to embed the inequality constraints in the objective functional (9), see [18].

While the feedback correction part may not be of a central importance for short explosive movements, it may become beneficial when the motion is long and/or when the system dynamics are not well identified.

3) *Application to compliant control*: Using the optimal solution given by $\mathbf{x}^* = (\mathbf{q}^{*T}, \dot{\mathbf{q}}^{*T})^T$ and \mathbf{u}^* , the optimal joint torques $\boldsymbol{\tau}^*$, and the optimal joint stiffness profiles \mathbf{K}^* can be obtained from (2) and (3) by substitution. Moreover, if the optimal feedback gains $\mathbf{L}^* = (\mathbf{P}^*, \mathbf{D}^*)$ (where $\mathbf{P}^* \in \mathbb{R}^{m \times n}$ and $\mathbf{D}^* \in \mathbb{R}^{m \times n}$) can also be computed, then the torque and stiffness will be defined in a form of a *locally valid* feedback control law: $\boldsymbol{\tau}_{opt} = \boldsymbol{\tau}(\mathbf{q}, \mathbf{u}_{opt})$ and $\mathbf{K}_{opt} = \mathbf{K}(\mathbf{q}, \mathbf{u}_{opt}) + \frac{\partial \boldsymbol{\tau}}{\partial \mathbf{u}} \mathbf{P}^*$, where $\mathbf{u}_{opt} = \mathbf{u}^* + \mathbf{P}^*(\mathbf{q} - \mathbf{q}^*) + \mathbf{D}^*(\dot{\mathbf{q}} - \dot{\mathbf{q}}^*)$. Note that the feedback correction on \mathbf{u}^* is a joint-level PD-control performed with optimal position and velocity gains.

In order to obtain not only the torques, but also the optimal time-varying stiffness profile, the present framework resolves the control redundancy (i.e., $\mathbf{u} \in \mathbb{R}^m$, $m > n$) in a system dependent and task specific way [14]. In this light, the control redundancy is optimally exploited to devise the best stiffness profiles for the system (1), actuator (2) and task (7) considered.

4) *Implementation and limitation*: In the iLQR implementation, we utilise finite differences to derive the linear approximation of the system dynamics (8), and the quadratic approximation of the cost (9) wherever an exact analytical approximation is not feasible. Furthermore, we use a fixed-step fourth-order Runge-Kutta method for numerical integration of the dynamics during the iterations. The admissible domain for the control actions is set by box constraints on the inputs as defined in (4). It is also important to point out that the iLQR method is local, and in order to circumvent local minima issues, we utilise multiple (random) initialisations to select the best possible solution. The model-based iLQR method is a viable optimization tool when the model of the system dynamics is reasonably well identified. If such model is not available analytically or too complex to estimate accurately, one could employ any model-free method for optimal control that is able to optimize over the inputs \mathbf{u} which non-linearly enter into the dynamics. Alternately, one could also use iLQG with learned dynamics (iLQG-LD, see [11]) where the model is acquired from data.

III. CASE STUDY: OPTIMAL BALL THROWING

In this section, we investigate the use of an optimal control approach to devising appropriate controllers for highly dynamic, explosive movements. As an example, we look at the problem of throwing a ball using a compliant two-link arm. Following the problem formulation, we analyse the control strategies devised by our approach in the light of varying the objective function and the dynamical properties of the system. Our goal is to characterise the optimal solutions found by our approach, before going on to analyse the exploitation of variable stiffness in greater depth.

A. Problem formulation

Here we formulate the problem of ball-throwing with a two-link, compliantly actuated robot arm. We present the system

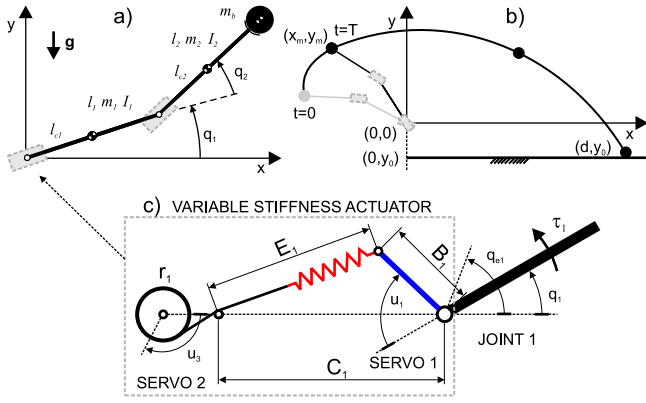


Fig. 1. a) Configuration of the arm with the inertial and geometric parameters: $l_1=0.25m$, $l_2=0.25m$, $l_{c1}=0.125m$, $l_{c2}=0.125m$, $m_1=0.1kg$, $m_2=0.1kg$, $m_b \in [0.01, 0.1]kg$, $I_1 = m_1 l_1^2 / 12 kgm^2$, $I_2 = m_2 l_2^2 / 12 kgm^2$. b) Ball throwing. c) Prototype of the variable stiffness actuator (used on the first joint). $u_1 = q_{e1}$ (and $u_2 = q_{e2}$) adjusts the equilibrium point of the springs, while u_3 (and u_4) adjusts the spring pretension. Actuator parameters: maximum limitations on the servos $u_{1,2} \in [-\pi/2, \pi/2]rad$ and $u_{3,4} \in [0, \pi]rad$, spring stiffness $\kappa_{1,2} = (400, 200)N/m$, radius of the winding drum (for spring pretension) $r_{1,2} = (0.01, 0.01)m$, other geometric parameters $B_{1,2} = (0.05, 0.05)m$, $C_{1,2} = (0.15, 0.15)m$. Joint damping: $\tau_f = -(b_1 \dot{q}_1, b_2 \dot{q}_2)^T$, where $b_{1,2} = (0.025, 0.025)^T Nms$.

dynamics, introduce the variable stiffness redundant actuation mechanism, and define the optimisation criterion.

1) *System dynamics*: The system dynamics of the arm, depicted in Fig. 1a, (left hand side of the equation of motion (1)), is specified as

$$\mathbf{M} = \begin{bmatrix} m_{11} & m_{12} \\ m_{21} & m_{22} \end{bmatrix}, \mathbf{C} = \begin{bmatrix} c_1 \\ c_2 \end{bmatrix}, \mathbf{G} = \begin{bmatrix} g_1 \\ g_2 \end{bmatrix},$$

where $\mathbf{q} = (q_1, q_2)^T$ define the joint angles, $m_{11} = I_1 + m_1 l_{c1}^2 + I_2 + m_2 (l_1^2 + 2l_1 l_{c2} \cos(q_2) + l_{c2}^2) + m_b (l_1^2 + 2l_1 l_2 \cos(q_2) + l_2^2)$, $m_{12} = m_{21} = I_2 + m_2 (l_{c2}^2 + l_1 l_{c2} \cos(q_2)) + m_b (l_2^2 + l_1 l_2 \cos(q_2))$, $m_{22} = I_2 + m_2 l_{c2}^2 + m_b l_2^2$ specify the mass matrix, $c_1 = -(m_2 l_{c2} + m_b l_2) l_1 \sin(q_2) (2\dot{q}_1 \dot{q}_2 + \dot{q}_2^2)$, $c_2 = (m_2 l_{c2} + m_b l_2) l_1 \sin(q_2) \dot{q}_1^2$ specify the Coriolis and normal inertial terms, while $g_1 = (m_1 l_{c1} + m_2 l_1 + m_b l_1) \cos(q_1) + (m_2 l_{c2} + m_b l_2) \cos(q_1 + q_2)$ and $g_2 = (m_2 l_{c2} + m_b l_2) \cos(q_1 + q_2)$ are the gravitational terms. The geometric and inertial parameters of the model are given in the caption of Fig. 1.

2) *Actuation model*: As an example of a physically realisable VSA, we investigate throwing using an arm with mechanically adjustable series elastic actuators (MACCEPA) introduced in [21]. On this actuator (depicted in Fig. 1c), two servo inputs are employed for direct control of the equilibrium point of the actuator, and the pretension of the (linear) spring respectively. The relations between the control inputs \mathbf{u} , joint torque $\boldsymbol{\tau}$, and joint stiffness⁴ \mathbf{K} , is given by

$$\boldsymbol{\tau} = \begin{bmatrix} \frac{B_1 C_1 \sin \alpha_1}{E_1} & 0 \\ 0 & \frac{B_2 C_2 \sin \alpha_2}{E_2} \end{bmatrix} \begin{bmatrix} \kappa_1 (l_{s1} - l_{01}) \\ \kappa_2 (l_{s2} - l_{02}) \end{bmatrix}, \quad (10)$$

⁴In the present case the stiffness matrix has diagonal elements only. Off-diagonal elements in the stiffness matrix appear if the configuration change on one joint induces torque change on another joint (see a human arm model in [13] which incorporates bi-articular muscles).

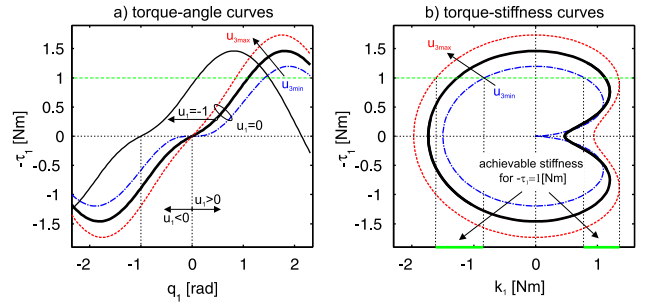


Fig. 2. Depicted is: a) the torque-angle ($q_0 = 0$) and b) the torque-stiffness relation on the (first) joint. By changing the equilibrium point $u_1 = q_{e1}$, the torque-angle curve translates in Fig.2a without affecting the torque-stiffness characteristic depicted in Fig.2b. The torque-stiffness curve can be altered with u_3 by modulating the pretension of the spring. In this way, u_1 and u_3 together allow simultaneous torque and stiffness control. Note that, if u_3 is kept constant, the torque-stiffness curve is fixed in Fig.2b and simultaneous torque-stiffness control is not possible. (The decreasing branch of the torque-angle curve, in Fig.2a, corresponds to the negative stiffness of the joint.)

$$\mathbf{K} = \text{diag}(\mathbf{k}), \mathbf{k} = [-\partial \tau_1 / \partial q_1, -\partial \tau_2 / \partial q_2]^T, \quad (11)$$

where $u_{1,2}$ and $u_{3,4}$ are the servo inputs associated with equilibrium point modulation and the spring tensioning on the first (base), and the second joint, respectively; \mathbf{q}_0 is the equilibrium configuration under no external force and no control input; $r_{1,2}$, $B_{1,2}$, $C_{1,2}$, $\alpha_i = u_i - q_i + q_{0i}$ and $E_i = \sqrt{B_i^2 + C_i^2} - 2B_i C_i \cos \alpha_i$, $i \in \{1, 2\}$, specify the geometry of the actuators along (see Fig. 1); $l_{si} = r_i u_{i+2} + E_i$ is the length of the springs, $l_{0i} = B_i - C_i$ is the corresponding equilibrium length and $\kappa_{1,2}$ are the spring constants. The specific parameter values can be found in the caption of Fig. 1.

Utilising (10) and (11), the torque and the stiffness of the joints can be simultaneously modulated, however, these relations are highly non-linear (see Fig. 2) which makes (motion, torque and stiffness) control of this system non-intuitive. Moreover, the range of motion of the adjuster servos is limited, restricting the achievable value of the joint-stiffness for a given torque. Such physical limitation, incorporated into the formulation with (4), may be a necessary ingredient in the optimal planning process, since it disallows unrealistic (e.g., infinitely large) controls to be employed for the task.

3) *Performance criterion*: The performance criterion for the optimal ball throwing task considered is defined by

$$J_w = -d(\mathbf{q}(T), \dot{\mathbf{q}}(T)) + \frac{w}{2} \int_0^T \|\mathbf{F}(\mathbf{q}, \mathbf{u})\|^2 dt, \quad (12)$$

where d is the distance thrown (measured at the horizontal ground level y_0 , see Fig. 1b), $w \in (0, \infty)$ weighs the effort minimisation term (here, chosen as the time integral of the squared spring forces⁵, \mathbf{F}), while T is the time permitted for task execution, $t \in [0, T]$.

The relation between the distance thrown and the configuration of the arm (evaluated at the time of the ball release in (12)) $d = d(\mathbf{q}, \dot{\mathbf{q}})$, is computed from the ballistic equation of the flying ball as

$$d = x_m(\mathbf{q}) + \dot{x}_m(\mathbf{q}, \dot{\mathbf{q}}) T_m(\mathbf{q}, \dot{\mathbf{q}}), \quad (13)$$

⁵Alternatively, one could use other objectives as in [15] for example.

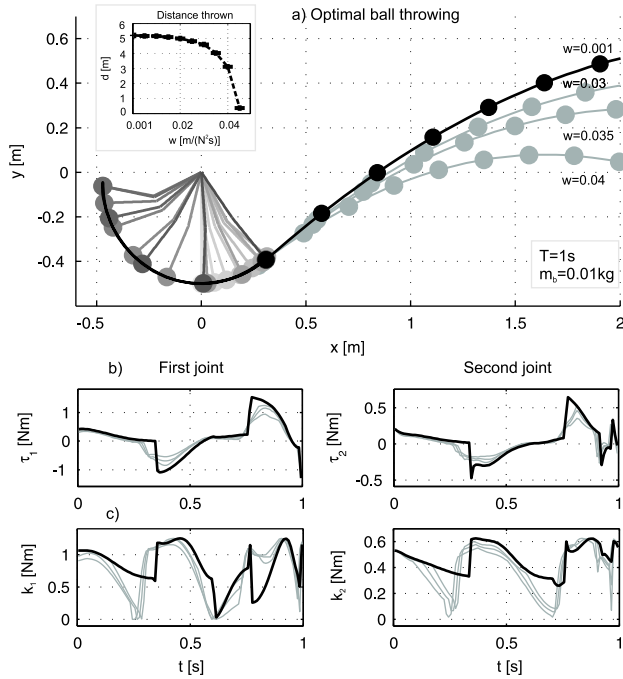


Fig. 3. a) Stroboscopic view of the optimal throwing motion. b) Joint torques. c) Joint stiffness. The arm is initialised from the vertical resting configuration. The limits on the control inputs are set to: $u_{1,2} \in [-\pi/6, \pi/6]rad$ and $u_{3,4} \in [0, 2\pi/3]rad$. The equilibrium of the springs (at no actuation) correspond to $\mathbf{q}_0 = [-\pi/2, 0]rad$.

where $x_m = l_1 \cos(q_1) + l_2 \cos(q_1 + q_2)$ and $\dot{x}_m = -l_1 \sin(q_1) \dot{q}_1 - l_2 \sin(q_1 + q_2) (\dot{q}_1 + \dot{q}_2)$ denote the horizontal position and velocity of the ball, and T_m is the ball flight time measured from release until the time impact of the ball with the ground. The latter is computed as

$$T_m = \frac{1}{g} (\dot{y}_m(\mathbf{q}, \dot{\mathbf{q}}) + \sqrt{\dot{y}_m(\mathbf{q}, \dot{\mathbf{q}})^2 + 2g(y_m(\mathbf{q}) - y_0)}), \quad (14)$$

where g is the gravitational constant, and $y_m = l_1 \sin(q_1) + l_2 \sin(q_1 + q_2)$, $\dot{y}_m = l_1 \cos(q_1) \dot{q}_1 + l_2 \cos(q_1 + q_2) (\dot{q}_1 + \dot{q}_2)$ denote the vertical position and velocity of the ball, respectively.

In the following, we confirm that our methodology is able to find optimal solutions adapted to this problem setup. In particular, we first look at how varying the objective functional (i.e., varying w) affects the solutions found by our framework. We then look at how changes to the system dynamics (specifically, changes to the mass of the ball thrown) affects the solution found by our approach.

4) *Variation of the objective function*: In this investigation, we looked at how the solutions found by our framework depend on the choice of weighting parameter w . The throwing motion was performed with a ball of mass $m_b = 0.01 kg$ and time duration $T = 1 s$.

Typical simulation results for choices of $w \in [0.001, 0.04]$ are depicted in Fig. 3. As can be seen, depending on the relative importance of the distance and effort terms in (12) (i.e., magnitude of w) a variety of optimal behaviours are predicted (Fig. 3a) characterised by different torque and stiffness profiles (Fig. 3b,c).

As expected, lower w (decreased penalisation of effort) results in longer distance throws. More interestingly, we note




Fig. 4. Throwing movement predicted for two different masses. The optimisation selects an under-arm strategy for the heavier ball (as in Sec. III-A4), however, for the light ball an over-arm strategy is selected. The arm is initialised from the vertical resting configuration, and the limits on the control inputs are set to: $u_{1,2} \in [-\pi/6, \pi/6]rad$ and $u_{3,4} \in [0, 2\pi/3]rad$; $\mathbf{q}_0 = [\pi/6, 0]rad$.

that a common pattern emerges in the movements, irrespective of the w chosen. This is a characteristic *counter-movement* strategy whereby there is an initial back-swing prior to the rapid forward acceleration before release. We note that such a strategy is often used by humans during fast, explosive movements (i.e., the “stretch shortening cycle” during throwing, hitting, jumping, kicking) [7]. The numerical predictions depicted in Fig. 3, obtained for a simple robotic device, are consistent with this biologically plausible strategy.

5) *Variation of the dynamics*: In this numerical investigation, we looked at how changes to the dynamics of the system affect the solutions found by our framework. Here, we analysed the effect of changing the mass of the ball on the throwing strategy employed. For these simulations, the duration was fixed to $T = 1 s$, and we selected $w = 0.001$. We compared throwing for $m_b = 0.1 kg$ and $m_b = 0.01 kg$.

A representative result is shown in Fig. 4. Interestingly, we see that for the two masses, very different throwing strategies emerge. In particular, we see that for the heavy ball an under-arm movement is used, while for the lighter ball an over-arm strategy is employed.

Emergence of the two strategies can be explained by considering the dynamic effects during the corresponding task execution. Specifically, if the weight of the ball is large, under-arm throwing is more beneficial from the optimisation point of view since lifting a heavy ball requires significant effort penalised by the second term in the control objective (12). On the other hand, if the ball is light, the effort required for over-arm throwing can pay off through a larger distance thrown (first term in (12)). In this case the motion is fast, dominated by the inertial dynamics, and executed through a fast counter-movement action. Again we note that both of these strategies are similar to those naturally employed by humans. Given a heavy ball, humans prefer to throw under-arm (as, for example, in ten-pin bowling), while for lighter balls they more commonly throw over-arm (as, for example, when fielding in cricket or base-ball) when attempting to send the ball over a large distance. The result presented in Fig. 4 demonstrates that such *strategy change in task execution* can be predicted by optimisation, depending on the weight of the ball.

IV. EXPLOITING VARIABLE STIFFNESS THROUGH OPTIMAL CONTROL

It is often argued that variable stiffness actuation is beneficial in order to achieve a human-like performance in highly dynamic, explosive tasks⁶. It is predicted that such benefits will arise from the ability of VSAs to simultaneously modulate joint torque and stiffness, and to amplify power (store energy). Using optimal control, we can exploit these abilities in a system dependent and task-specific way. This is demonstrated by comparing a *variable stiffness device* with a *fixed stiffness device*, where the joint torque-stiffness function cannot be altered with the control inputs.

A. The benefit of stiffness variation

In order to investigate the benefit of optimal control with variable stiffness, we performed a comparative study between (i) the redundantly actuated system (10), and (ii) a system with identical dynamics, but with the actuation redundancy removed.

Specifically, the latter is a variant on the MACCEPA design, but with a fixed pre-tensioning of the springs (i.e., the pre-tensioning servos were commanded to a fixed position, $u_{3,4} = \text{const.}$), leaving the actuation to the equilibrium motors, $u_{1,2}$, alone. Commands for this system were selected by optimising the constants $u_{3,4}$, and seeking the optimal command sequence for $u_{1,2}$. Note that, keeping $u_{3,4}$ constant does not ensure constant joint stiffness, but it does ensure that joint torque and stiffness cannot be independently optimised, see Fig. 2b in Sec. III-A2.

For the purpose of demonstration, we consider a throwing task specified with $T = 3s$, $m_b = 0.05kg$ and $w = 0.001$. In Fig. 5a,b we show that the throwing performance is improved for the case when both torque and stiffness are simultaneously optimised ($d = 5.7m$, $J_w = -5.04$) as opposed to using a fixed torque-stiffness relation ($d = 3.7m$, $J_w = -3.04$). Note that by choosing a longer time horizon ($T = 3s$), an interesting cyclic movement emerges which is investigated in more detail in the next section.

B. Energy storage with passive compliance

One of the proposed benefits of passive VSAs is their ability to store mechanical energy, which can be later used to amplify the power output of an actuator. This is particularly important for explosive tasks [1, 24, 16] where, by exploiting this property, we may significantly enhance the peak joint performance without the need for increasing the size and power (and thereby cost) of our actuators. In this section, we analyse the solutions found by the optimal control framework to determine, if they are indeed able to exploit energy storage effects.

Fig. 6a illustrates the strategy selected by our approach for throwing with $m_b = 0.01kg$, $T = 3s$ and $w = 0.001$. As can be seen, a sequential *energy pumping* strategy emerges during

⁶Human peak performance characterised with the rotation speed of the shoulder during a baseball pitch of a professional pitcher, is between $6900 - 9800^\circ/s$, [3]. This kind of high-performance task execution is not in the scope of present robotic systems.

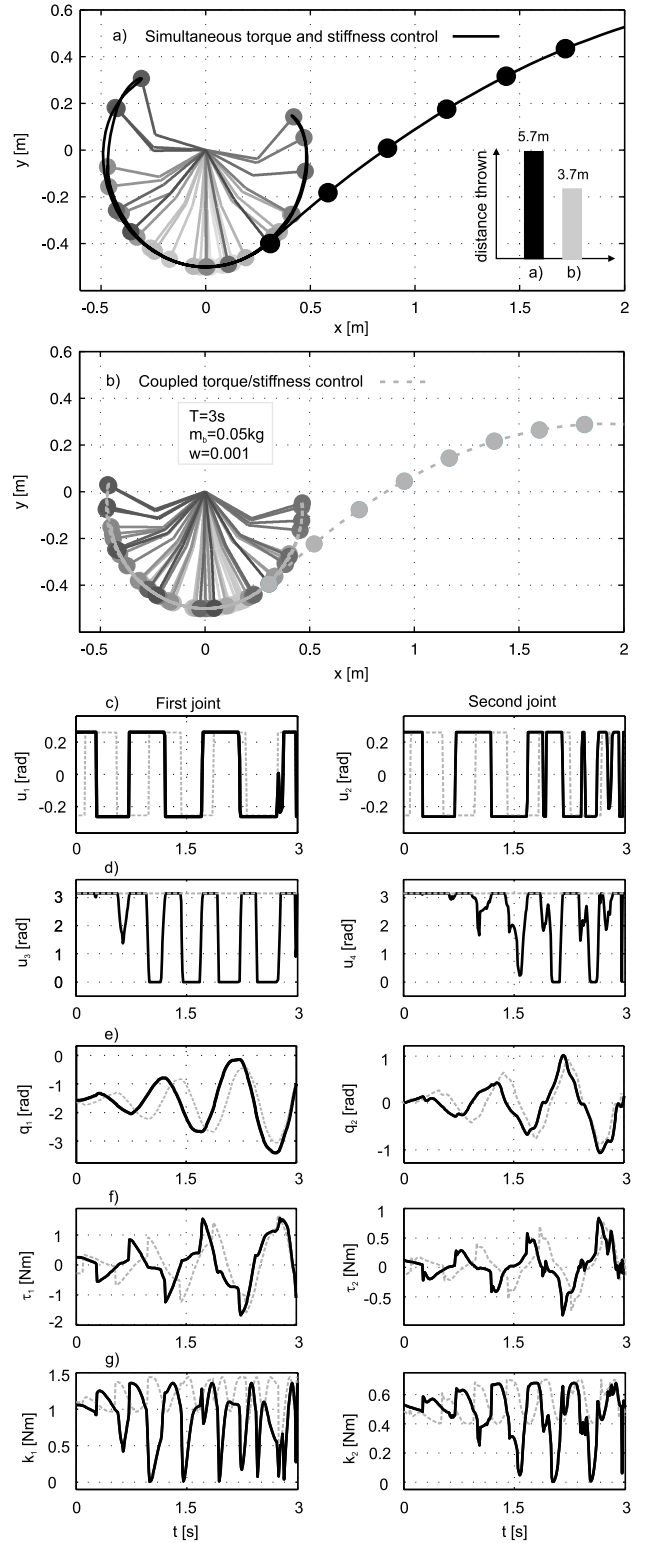


Fig. 5. Exploitation of variable stiffness actuation. Stroboscopic view of the throwing motion: a) simultaneous torque/stiffness control, results plotted in black and b) control with fixed torque-stiffness relation ($u_{3,4} = \text{const.}$), results plotted in gray. c) Optimal equilibrium point commands, d) spring pretension, (i.e., the *optimal constant values*, used in b), corresponds to the maximal pretension, $u_3 = u_4 = \pi \text{rad}$, e) motion, f) torques, and g) time-varying stiffness profiles. The arm is initialised from a vertical resting configuration. Limits on the control inputs are set to: $u_{1,2} \in [-\pi/12, \pi/12] \text{rad}$ and $u_{3,4} \in [0, \pi] \text{rad}$; $\mathbf{q}_0 = [-\pi/2, 0] \text{rad}$.

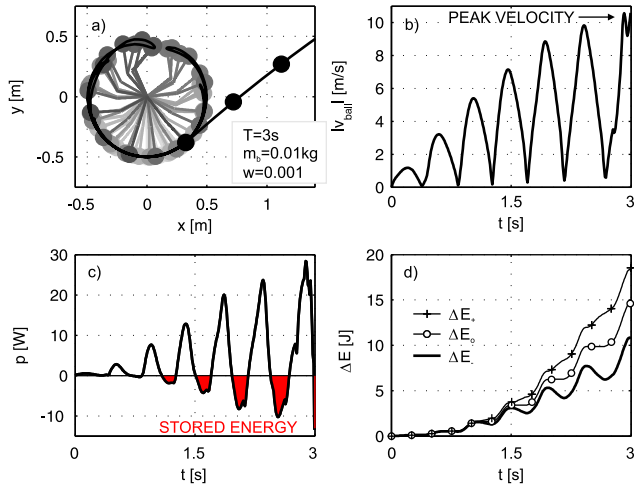


Fig. 6. a) Stroboscopic view of the optimal throwing motion. The arm is initialised from a vertical resting configuration. The limits on the control inputs are set to: $u_{1,2} \in [-\pi/6, \pi/6]rad$ and $u_{3,4} \in [0, 2\pi/3]rad$; $\mathbf{q}_0 = [-\pi/2, 0]rad$. b) Ball velocity during the throwing motion. c) Output mechanical power $p = \boldsymbol{\tau}^T \dot{\boldsymbol{\varphi}}$. d) Output mechanical energy calculated for: 1) a compliant actuator $\Delta E_- = \int_0^T p dt$, 2) an actuator which has no ability to store energy $\Delta E_0 = \int_0^T (1/2)(p + |p|) dt$, and 3) a stiff actuator where both acceleration and braking is actively generated $\Delta E_+ = \int_0^T |p| dt$.

which the velocity of the ball is cyclically increased, and the output power of the actuators is progressively amplified (see Fig. 6c).

Implemented on systems with *passively* compliant actuators, this energy pumping strategy would be realised by using the actuators to cyclically accelerate the arm (which corresponds to positive power), followed by braking the motion by storing energy (negative power) (see Fig. 6c). This is in contrast to some of the sophisticated *actively controlled* VS robots that have the ability to modulate joint stiffness in software without actually using passive elements in the design. While temporal stiffness modulation is feasible, they are unable to realise energy storage or power amplification during the motion. This is illustrated in Fig. 6d, where we demonstrate the difference in the output mechanical energy requirement for the same throwing motion realised on the present passively compliant system where energy storage is possible, and other (two) actuators which does not allow bidirectional power flow, and as such cannot store energy. As indicated, in Fig. 6d, the output mechanical energy requirement would be considerably higher if the system were to be actuated without passive elements. However, whether the passive compliance mechanism will significantly improve task performance depends on the system dynamics, the capacity of the actuator to store energy, but also the task considered.

V. BALL THROWING EXPERIMENT

In this section, we present a ball-throwing experiment performed with a one-link variable stiffness robot (see Fig. 7a). Simultaneous torque and stiffness control on this device is achieved with a MACCEPA actuator. The actuator is realised using two servomotors (Hitec HSR-5990TG) which are controlled with 50 Hz PWM signals from a micro-controller

(Atmel ATmega328). The joint angle is measured with a rotary encoder (Melexis MLX90316GO). The experiment is performed with a tennis ball equipped with a small metallic plate. Using an electromagnet (Magnet-Schultz, GMHX025X00A01 mounted at the end of the moving link, see Fig. 7a) the ball is held during the movement and released at the final instant.

The optimisation for variable stiffness ball throwing is performed using $w \in \{0.001, 0.01, 0.012, 0.015\}$, $T = 2.5s$, $m_b = 0.075kg$, with a model of the robot that incorporates a third-order (linear) electro-mechanical model of the servomotors. The corresponding experimental videos⁷ demonstrate that with increasing effort weight the distance thrown decreases (i.e., $d = 2.3m$, $d = 1.7m$, $d = 1.4m$, $d = 0.9m$). A similar trend was predicted numerically for the two link arm (see Fig. 3).

In Fig. 7 we present: b) the stroboscopic view, c,d) the command sequences, and e,f) the recorder (and simulated) joint angle and velocity data for a representative experiment which corresponds to $w = 0.01$. As can be seen, in Fig. 7e,f, there is a good match between the simulation and the data recorded from the robot, in particular, with respect to the synchronisation, ensuring near to optimal timing of the ball release (i.e., at the maximum speed). We note also that these results were obtained with *open-loop execution of the optimal commands*: such a strategy is preferable for fast, explosive movements in the presence of sensory feedback delays and motor delays originating from the third-order dynamics of the servos.

In addition to the above result, we also performed a fixed stiffness throwing experiment (with optimal constant pre-tensioning $u_2 = 1.1rad$) for the task defined with $w = 0.01$, $T = 2.5s$, $m_b = 0.075kg$, (see supplementary video). As predicted in the paper, variable stiffness throwing ($d = 1.7m$) provided a clear performance benefit compared to the corresponding fixed stiffness case ($d = 1.4m$). However, unlike the two-link example (see Fig. 5a,b), the performance gain with variable stiffness on this one-link experimental device is relatively modest. This highlights the fact that the benefits of stiffness variation depend critically on the dynamics and the actuator utilised. Reflected to the present case, due to space constraints, the stiffness range in our implementation of the MACCEPA actuator is relatively small. With larger stiffness range better exploitation of variable stiffness is expected.

Furthermore, comparing the results in this section with those predicted for the two-link devices (ref. Fig. 5), it appears that stiffness variation may play an important role in exploitation of the interaction-torques (i.e., torques given by $\mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})$ in (1)), not present during single joint movements. Exploration of this effect is part of our future work.

VI. CONCLUSION

In this article, we demonstrate the utility of an optimal control framework applied to compliantly actuated robotic systems in

⁷<http://homepages.inf.ed.ac.uk/svijayak/pub/braun-RSS2011.wmv>




Fig. 7. a) One-link variable stiffness arm. b) Stroboscopic view of the throwing motion. c,d) Optimal control sequences. e,f) Simulated and experimentally realised trajectories of variable stiffness throw. The geometric and inertial parameters of the robot are: length $l = 0.295m$, link mass $m = 0.2kg$, mass of the ball $m_b = 0.075kg$. The design parameters of the MACCEPA actuator (see Fig.1c) are: $\kappa = 323Nm$, $r = 0.01m$, $B = 0.03m$ and $C = 0.13m$.

the context of highly dynamic, explosive tasks. Using ball-throwing as a prototypical example, we have characterised the optimal strategies in settings with differing task demands, both in terms of (i) the trade-off between performance and effort costs, and (ii) changes to the dynamics induced, for example, by changes to ball mass. In particular, we have seen the emergence of torque/stiffness profiles that exploit the system dynamics, often in a non-intuitive way, that would be difficult to obtain through hand-tuning. Finally, we have presented an analysis of these results in the light of the energy-storage ability of compliant actuators, and demonstrated the benefits of stiffness variation over fixed stiffness, both in simulation and experiment. In future work, we intend to (i) further validate our numerical predictions in the context of multi-link throwing, and (ii) investigate the role of other impedance terms, such as damping, in the context of optimal explosive movements.

ACKNOWLEDGMENTS

This work was funded by the EU Seventh Framework Programme (FP7) as part of the STIFF project. The authors gratefully acknowledge this support.

REFERENCES

- [1] R. M. Alexander and H. C. Bennet-Clark. Storage of elastic strain energy in muscle and other tissues. *Nature*, 265:114–117, 1977.
- [2] R. Anderson and M. Spong. Hybrid impedance control of robotic manipulators. *IEEE J. Robotics and Automation*, 4(5):549–556, 1988.
- [3] I. P. Herman. *Physics of the Human Body*. Springer Verlag, 2007.
- [4] N. Hogan. Impedance control: An approach to manipulation. *ASME J. Dyn. Syst., Meas., Control*, 107:124, 1985.
- [5] J. W. Hurst, J. Chestnutt, and A. A. Rizzi. The actuator with mechanically adjustable series compliance. *IEEE Trans. Robotics*, 26(4): 597–606, 2010.
- [6] R. Ikeura, T. Moriguchi, and K. Mizutani. Optimal variable impedance control for a robot and its application to lifting an object with a human. In *Proc. IEEE Int. W.S. Robot and Human Interactive Comm.*, 2002.
- [7] P. V. Komi. *Stretch-shortening cycle. The encyclopaedia of sports medicine*, volume In: Strength and power in sport. Blackwell Scientific, Oxford, 1992.
- [8] W. Li and E. Todorov. Iterative linearization methods for approximately optimal control and estimation of non-linear stochastic system. *Int. J. Control*, 80(9):1439–1453, 2007.
- [9] M. Matinfar and K. Hashtrudi-Zaad. Optimisation-based robot compliance control: Geometric and linear quadratic approaches. *IJRR*, 24(8): 645–656, 2005.
- [10] S. A. Migliore, E. A. Brown, and S. P. DeWeerth. Novel nonlinear elastic actuators for passively controlling robotic joint compliance. *J. Mechanical Design*, 129(4):406–412, 2007.
- [11] D. Mitrovic, S. Klanke, and S. Vijayakumar. *From Motor Learning to Interaction Learning in Robots: Adaptive Optimal Feedback Control with Learned Internal Dynamics Models*, volume SCI 264. Springer-Verlag, 2010.
- [12] D. Mitrovic, S. Klanke, and S. Vijayakumar. Learning impedance control of antagonistic systems based on stochastic optimization principles. *IJRR*, 30(2):1–18, 2011.
- [13] F. A. Mussa-Ivaldi, N. Hogan, and E. Bizzi. Neural, mechanical, and geometric factors subserving arm posture in humans. *J. Neuroscience*, 5:2732–2743, 1985.
- [14] Y. Nakamura and H. Hanafusa. Optimal redundancy control of robot manipulators. *IJRR*, 6(1):32–42, 1987.
- [15] W. L. Nelson. Physical principles for economies of skilled movements. *Biological Cybernetics*, 46(2):135–147, 1983.
- [16] D. Paluska and H. Herr. The effect of series elasticity on actuator power and work output: Implications for robotic and prosthetic joint design. *Robotics & Auton. Sys.*, 54:667–673, 2006.
- [17] B. Siciliano and O. Khatib. *Handbook of Robotics*. Springer, 2008.
- [18] R. F. Stengel. *Optimal control and estimation*. Dover Publications, New York, 1994.
- [19] G. Tonietti, R. Schiavi, and A. Bicchi. Design and control of a variable stiffness actuator for safe and fast physical human/robot interaction. In *IEEE ICRA*, pages 526–531, Barcelona, Spain, 2005.
- [20] M. Uemura and S. Kawamura. Resonance-based motion control method for multi-joint robot through combining stiffness adaptation and iterative learning control. In *IEEE ICRA*, pages 1543 – 1548, Kobe, Japan, 2009.
- [21] R. van Ham et. al. MACCEPA, the mechanically adjustable compliance and controllable equilibrium position actuator: Design and implementation in a biped robot. *Robotics & Auton. Sys.*, 55(10):761–768, 2007.
- [22] B. Vanderborght, B. Verrelst, R. Van Ham, M. Van Damme, D. Lefeber, B. M. Y. Duran, and P. Beyl. Exploiting natural dynamics to reduce energy consumption by controlling the compliance of soft actuators. *IJRR*, 25(4):343–358, 2006.
- [23] B. Verrelst, R. Van Ham, B. Vanderborght, J. Vermeulen, D. Lefeber, and F. Daerden. Exploiting adaptable passive behaviour to influence natural dynamics applied to legged robots. *Robotica*, 23(2):149–158, 2005.
- [24] A. M. Wilson, J. C. Watson, and G. A. Lichtwark. A catapult action for rapid limb protraction. *Nature*, 421:35–36, 2003.
- [25] S. Wolf and G. Hirzinger. A new variable stiffness design: Matching requirements of the next robot generation. In *IEEE ICRA*, pages 1741 – 1746, Pasadena, CA, USA, 2008.
- [26] M. Zinn, O. Khatib, B. Roth, and J. K. Salisbury. Playing it safe. *IEEE Rob. & Autom. Mag.*, 11(2):12–21, 2004.