

Optimal variable stiffness control: formulation and application to explosive movement tasks

David Braun · Matthew Howard · Sethu Vijayakumar

Received: 1 September 2011 / Accepted: 5 June 2012 / Published online: 27 June 2012
© Springer Science+Business Media, LLC 2012

Abstract It is widely recognised that compliant actuation is advantageous to robot control once dynamic tasks are considered. However, the benefit of intrinsic compliance comes with high control complexity. Specifically, coordinating the motion of a system through a compliant actuator and finding a task-specific impedance profile that leads to better performance is known to be non-trivial. Here, we propose an optimal control formulation to compute the motor position commands, and the associated time-varying torque and stiffness profiles. To demonstrate the utility of the approach, we consider an “explosive” ball-throwing task where exploitation of the intrinsic dynamics of the compliantly actuated system leads to improved task performance (i.e., distance thrown). In this example we show that: (i) the proposed control methodology is able to tailor impedance strategies to specific task objectives and system dynamics, (ii) the ability to vary stiffness can be exploited to achieve better performance, (iii) in systems with variable *physical compliance*, the present formulation enables exploitation of the energy storage capabilities of the actuators to improve task performance. We illustrate these in numerical simulations, and in hardware experiments on a two-link variable stiffness robot.

Keywords Variable impedance control · Optimal stiffness control · Dynamic task · Power amplification

1 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 humans when interacting with compliantly actuated robots (Bicchi and Tonietti 2004; Zinn et al. 2004), (ii) *adding functionality* to enable adaptation to task requirements, e.g., robots can be stiff and accurate, but also compliant to the environment if required, and (iii) *improving the dynamic range of existing systems* by exploiting the energy storage capabilities of VSAs (Hurst et al. 2010). Due to the last point, a particularly promising area in which variable stiffness actuators may be deployed are applications involving highly dynamic movements (Braun et al. 2011). Typical examples of such movements include: throwing, hitting, jumping or kicking, often referred to in human studies as *explosive movements* (Newton et al. 1996; Putnam 1993; van Soest and Bobbert 1993). Additionally, rhythmic movements such as walking and running may also contain an explosive component (Wilson et al. 2003).

Explosive movement tasks may be intuitively characterised by “large release of energy over a short time frame”. Achieving such movements with traditional, joint torque actuators presents significant difficulties, particularly due to the size and power limitation of the motors. In contrast, by incorporating a physical elastic element, variable stiffness actuators offer the possibility of achieving a much higher dynamic range with smaller motors by instead exploiting the compliant nature of the actuators.

Electronic supplementary material The online version of this article (doi:10.1007/s10514-012-9302-3) contains supplementary material, which is available to authorized users.

D. Braun (✉) · M. Howard · S. Vijayakumar
School of Informatics, IPAB, University of Edinburgh,
Edinburgh, Scotland, UK
e-mail: david.braun@ed.ac.uk

M. Howard
e-mail: matthew.howard@ed.ac.uk

S. Vijayakumar
e-mail: sethu.vijayakumar@ed.ac.uk

One of the difficulties of using such actuators, however, is the considerably increased complexity in the planning and control, particularly in the context of dynamic movements. Indeed, VSAs (English 1999a, 1999b; Ham et al. 2007; Hurst et al. 2010; Koganezawa et al. 1999; Laurin-Kovitz et al. 1991; Migliore et al. 2007; Morita and Sugano 1995; Shen and Goldfarb 2007; Tonietti et al. 2005; Wolf and Hirzinger 2008), often introduce non-linearities, coupling of motion, torque and stiffness characteristics, and redundancy (i.e., increased dimensionality of the control input). As a result, it becomes increasingly difficult to hand-tune or design control strategies for such actuators in order to exploit the benefits of variable stiffness. To address this issue, Hogan (1985) pointed out that if the objective is to minimise the tracking error and the interface force, then the manipulator impedance (e.g., stiffness, damping and inertia) 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 (Anderson and Spong 1988). 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 issue for various tasks. For example, in contact tasks, trajectory tracking with optimal constant impedance control has been proposed by Johansson and Spong (1994). The benefit of optimal time-varying stiffness control has been investigated by Matinfar and Hashtrudi-Zaad (2005) and Mitrovic et al. (2011) for contact tasks and compliant motion control tasks, respectively. The utility of damping variation was shown by Ikeura et al. (2002), 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, Vanderborght et al. (2006) and Verrelst et al. (2005) 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 (Uemura and Kawamura 2009) combined trajectory tracking with adaptation of the constant joint stiffness to its optimal value. By considering velocity maximisation, optimal control on variable stiffness devices has been recently considered by Garabini et al. (2011) and Haddadin et al. (2011). In these works, analytical predictions to optimal bang-bang control were obtained for idealised variable stiffness systems.

Here we present a non-linear optimal control formulation for variable stiffness control that is applicable to multi-

link robots, performing complex dynamic tasks, under real-world conditions. The premise of this formulation is to: provide a physically realistic *control parametrisation* and a *dynamic representation* that allow the natural dynamics of the robot and the intrinsic compliance of the VSAs to be exploited during the motion. To achieve this, we propose to directly optimise on the level of the redundant input commands to the VSAs while taking into account the: (i) possibly non-linear model of the VSA, (ii) the bandwidth limitations of the actuator dynamics and (iii) the range limitations of the input commands respectively. When implemented, the optimal input commands provide the elastic joint torque and joint stiffness profiles together with the link trajectories that are *not planned in traditional sense* but are the direct consequence of the mechanical properties of the system dynamics and the actuators. In this way, the proposed formulation enables the dynamics of the compliantly actuated system to be exploited to achieve better task performance. This may be difficult to obtain by alternative formulations that do not take the physical constraints imposed by the VS actuators into account.

The present formulation is tested on a highly dynamic ball throwing task.¹ In this task we first show that (i) increased penalisation of effort results in shorter distance thrown and that (ii) over-arm versus under-arm throwing may emerge depending on the weight of the ball. Both of these results are intuitive but non-trivial to obtain by hand-tuning or other heuristic methods. Furthermore, we investigate *how*, by taking an optimal control approach, one may exploit the physical properties of a variable stiffness system to improve performance. To this end, we show by example, that in variable stiffness control the ability to *independently* control the joint torque and the joint stiffness can be leveraged to achieve larger distance thrown, compared to an alternative strategy that does not employ independent stiffness modulation. In addition, we discuss *why* (and under what condition) may stiffness modulation emerge as an optimal strategy in the case of highly dynamic (explosive) tasks, despite the energy cost it requires.

The present control methodology is illustrated with a number of simulation studies, and with throwing experiments on a variable stiffness robot. These experiments demonstrate the viability of the proposed variable stiffness control framework under real-world conditions.

2 Model of a variable stiffness robot

In this section, we present a model of a variable stiffness robotic system that is the basis for the optimal control for-

¹Regarding the ball throwing task, the reader may also refer to recent works that consider underactuated robots, and use optimal motion control, or employ unstable zero dynamics to achieve fast and accurate throw; see Mettin et al. (2010), Shoji et al. (2010).

mulation introduced later in this paper. The basic ingredients of this model are: (i) the dynamic equation of the robot, (ii) the torque and stiffness function associated with the compliant actuators and (iii) the performance index that defines the control task.

2.1 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

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

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, $\mathbf{D} \in \mathbb{R}^n$ represents the dissipative terms due viscous friction, $\mathbf{G} \in \mathbb{R}^n$ are the gravitational terms, $\boldsymbol{\tau} \in \mathbb{R}^n$ are the joint torques and $\boldsymbol{\theta} \in \mathbb{R}^m$ are the motor positions, that are the inputs to the actuators. If the motors were directly (rigidly) connected to the links (i.e., $\boldsymbol{\theta} = \mathbf{q}$) the joint torques $\boldsymbol{\tau}$ could be considered as control input (Siciliano and Khatib 2008). In the present paper, however, we consider robots equipped with *compliant actuators*, a model of which is introduced in the following.

2.2 Compliant actuators

If the actuators have in-built compliance, the joint torques $\boldsymbol{\tau}$ cannot be directly commanded.² In such cases, the torque function is in general complicated, position dependent, see (1), and can only be indirectly modulated through control of the motor positions $\boldsymbol{\theta}$. A viable approach for this, is to use servo-control on $\boldsymbol{\theta}$ to adjust the length and the moment-arm of the compliant elements (i.e., springs) embedded in the actuator (Ham et al. 2007; Hurst et al. 2010). In the following, we consider a model of VSAs that allow simultaneous torque and stiffness modulation. This model includes the dynamics of the servo-controlled motors and the compliant mechanism.

2.2.1 Dynamics of the motors

In this paper, we consider the motor positions $\boldsymbol{\theta}$ to be servo-controlled through critically-damped second-order dynamics,

$$\ddot{\boldsymbol{\theta}} + 2\alpha\dot{\boldsymbol{\theta}} + \alpha^2\boldsymbol{\theta} = \alpha^2\boldsymbol{\theta}_d, \quad (2)$$

²This is because on compliant actuators the joint torques do not correspond to the motor torques directly.

where³ $\boldsymbol{\alpha} = \text{diag}(\alpha_1, \dots, \alpha_m) \in \mathbb{R}^{m \times m}$ are positive gains and $\boldsymbol{\theta}_d \in \mathbb{R}^m$ are the desired angles (i.e., motor angles reflected through gear reduction). The equation above models the servo-controlled closed-loop dynamics of the position input to the variable stiffness actuators. The admissible values for $\boldsymbol{\alpha}$ (i.e., one parameter per equation) are limited by the bandwidth of the motor-gearbox unit under the implemented servo control loop. Using $\boldsymbol{\alpha} \in (\mathbf{0}, \boldsymbol{\alpha}_{max}]$, (2) provides motor positions that can be tracked by the *real actuators*. In practice, $\boldsymbol{\alpha}_{max}$ is estimated from velocity limits of the servo system, or identified by fitting the response of the real actuators with the response of (2) obtained under the same excitation $\boldsymbol{\theta}_d$. Setting $\boldsymbol{\alpha} = \boldsymbol{\alpha}_{max}$ provides full exploitation of the dynamic range of the actuators, while setting $\boldsymbol{\alpha} < \boldsymbol{\alpha}_{max}$ allows one to limit the dynamic range (maximum speed and acceleration) of the actuators.

In addition to the *dynamic constraints* represented by (2), there are often also *range constraints* on the motor positions. The admissible set defined by these constraints is given by

$$\Theta = \{\boldsymbol{\theta} \in \mathbb{R}^m : \boldsymbol{\theta}_{min} \leq \boldsymbol{\theta} \leq \boldsymbol{\theta}_{max}\}, \quad (3)$$

where $\boldsymbol{\theta}_{min}$ and $\boldsymbol{\theta}_{max}$ are the lower and the upper bounds to the admissible motor positions. In this paper we will explicitly incorporate (3) into the optimal control formulation (see Sect. 2.3.1) when devising the optimal motor commands and the corresponding torque/stiffness modulation.

2.2.2 Torque modulation

On compliantly actuated systems, the relation between the joint torques, joint angles and the motor positions is given in the following form⁴

$$\boldsymbol{\tau}(\mathbf{q}, \boldsymbol{\theta}) = \mathbf{A}^T(\mathbf{q}, \boldsymbol{\theta})\mathbf{F}(\mathbf{q}, \boldsymbol{\theta}), \quad (4)$$

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 (4), both the moment arm and the associated forces may explicitly depend on the motor positions $\boldsymbol{\theta}$. As will be subsequently discussed, this dependence may allow one to independently modulate the joint torques and the (passive) joint stiffness of the actuators.

³Where $\alpha^2 = \text{diag}(\alpha_1^2, \dots, \alpha_m^2)$.

⁴While in the present paper, the actuator torque is assumed to be position dependent (as is the case in the majority of VS actuators), the formulation remains valid for cases where the torque is velocity dependent (e.g., due to viscoelastic forces).

2.2.3 Stiffness modulation

Using (4), 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

$$\mathbf{K}(\mathbf{q}, \boldsymbol{\theta}) = -\frac{\partial \mathbf{A}^T}{\partial \mathbf{q}} \mathbf{F} - \mathbf{A}^T \frac{\partial \mathbf{F}}{\partial \mathbf{q}}. \tag{5}$$

Relation (5) allows one to identify the conditions under which changing the motor positions $\boldsymbol{\theta}$ leads to active modulation of \mathbf{K} . Specifically, the first term in (5) indicates that the joint stiffness can be directly changed by modulation of the elastic force \mathbf{F} through $\boldsymbol{\theta}$, if the moment arm \mathbf{A} depends on the link-side position (e.g., Mitrovic et al. 2011). If the moment arm does not depend on \mathbf{q} , then joint-level stiffness modulation requires either (i) a non-linear force-angle relation (i.e., the stiffness, defined by $\mathbf{K}_F = -\partial\mathbf{F}/\partial\mathbf{q}$, should be motor position dependent) or (ii) a motor position dependent moment arm. The former is the mechanism used in many antagonistic actuators (e.g., English 1999a; Migliore et al. 2007), while the latter allows joint stiffness modulation, even if \mathbf{K}_F is constant (e.g., Ham et al. 2007; Kim and Song 2010). Regardless of which of these mechanisms is employed, we assume that the joint torque function (4) is redundant with respect to the motor positions (i.e., $m > n$) enabling (but not ensuring) simultaneous and independent torque and stiffness modulation on one or more of the joints.

2.3 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 while the second is due to the physical restrictions on the realisable control actions.

In the following, we define the control inputs, propose a state-space representation of the dynamics, introduce the computational framework used in this work and discuss application of the present formulation to compliantly actuated robots.

2.3.1 System dynamics and control constraints

In this paper, we propose a state-space representation of a variable stiffness system that includes the rigid body dynamics of the robot (1) and the closed-loop dynamics of the motors (2). In this representation,

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}), \tag{6}$$

$\mathbf{x} = (\mathbf{x}_1^T, \mathbf{x}_2^T, \mathbf{x}_3^T, \mathbf{x}_4^T)^T = (\mathbf{q}^T, \dot{\mathbf{q}}^T, \boldsymbol{\theta}^T, \dot{\boldsymbol{\theta}}^T)^T \in \mathbb{R}^{2(n+m)}$ is the state vector,

$$\mathbf{f} = \begin{bmatrix} \mathbf{x}_2 \\ \mathbf{M}(\mathbf{x}_1)^{-1}(\boldsymbol{\tau}(\mathbf{x}_1, \mathbf{x}_3) - \mathbf{C}(\mathbf{x}_1, \mathbf{x}_2) - \mathbf{D}(\mathbf{x}_2) - \mathbf{G}(\mathbf{x}_1)) \\ \mathbf{x}_4 \\ -2\boldsymbol{\alpha}\mathbf{x}_4 - \boldsymbol{\alpha}^2\mathbf{x}_3 + \boldsymbol{\alpha}^2\mathbf{u} \end{bmatrix},$$

is the non-linear vector field while $\mathbf{u} = \boldsymbol{\theta}_d \in \mathbb{R}^m$ are the control inputs associated with the desired motor positions, see (2).

In the model above, the diagonal matrix $\boldsymbol{\alpha} \in \mathbb{R}^{m \times m}$ provides a simple way to incorporate the bandwidth limitations of the motor dynamics (see Sect. 2.2.1). In addition, the constraints on the motor positions (3), defined by $\boldsymbol{\theta} \in \Theta$ (where $\Theta = \{\boldsymbol{\theta} \in \mathbb{R}^m : \boldsymbol{\theta}_{min} \leq \boldsymbol{\theta} \leq \boldsymbol{\theta}_{max}\}$), can also be incorporated through control constraints defined by

$$\mathbf{u} \in U = \{\mathbf{u} \in \mathbb{R}^m : \boldsymbol{\theta}_{min} \leq \mathbf{u} \leq \boldsymbol{\theta}_{max}\}. \tag{7}$$

This is because $\mathbf{u} \in U$ implies $\boldsymbol{\theta} \in \Theta$ under (2).

The three important ingredients in this formulation are:

- (a) the intrinsically non-linear state dependent torque model (4), that is specific to the variable stiffness actuator,
- (b) the bandwidth limitation of the actuators, that is incorporated by (2) and
- (c) the range constraints on the (desired) motor positions enforced by (7).

Formulations that utilise idealistic VSA models (e.g., linear torque models), instead of the design specific non-linear one (4), may not be able to exploit the intrinsic compliance of the system to improve performance. Moreover if (b) and (c) are not included in the formulation, than the planned optimal strategy may not be applicable under real-world conditions. This is because disregarding (b) allows *instantaneous* modulation of the elastic joint torque and the joint stiffness, while not including (c) would enable *unconstrained* joint torque and stiffness modulation. Neither of these is possible on a realistic VS actuator.

2.3.2 Control objective

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*

$$J = h(\mathbf{x}(T)) + \int_0^T c(\mathbf{x}(t), \mathbf{u}(t, \mathbf{x}(t))) dt, \tag{8}$$

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 (Nelson 1983).

It is known that defining the objective function is non-trivial for *many tasks* (Anderson and Pandy 2001; Bobrow et al. 1985; Flash and Hogan 1985; Pandy et al. 1995; Uno et al. 1989). This is mainly because the relation between

the cost and the optimal behaviour are often non-intuitive (Todorov 2004), but also because different costs may lead to similar behaviour for a given task (Collins 1995). Despite these issues, representation of “explosive movement tasks” with an objective function is in many cases *unambiguous* (Pandy et al. 1990). *One of the main reasons for this is that such objective functions are defined by the kinematic attributes of the movement with no emphasis on minimising the effort cost associated with the task.*⁵ This is exemplified in Sect. 3.1.3 where we demonstrate how to define (8) to represent an explosive ball throwing task.

2.3.3 Solution method

For non-linear plant dynamics (6) and non-quadratic cost (8), a globally valid optimal control law could be derived by means of *dynamic programming* (Bellman 1957). This method requires one to find the general solution to the non-linear Hamilton-Jacobi-Bellman (HJB) partial differential equation in order to define the optimal feedback control law $\mathbf{u} = \mathbf{u}(t, \mathbf{x})$ (Bryson and Ho 1975; Stengel 1994). While the HJB equation provides a *sufficient condition* to a global optimal feed-back control, it is computationally expensive and as such less attractive for complex non-linear systems.

To circumvent this issue, one may utilise trajectory optimisation to define an open-loop optimal control $\mathbf{u} = \mathbf{u}(t)$ provided by *Pontryagin’s maximum principle* (PMP) (Kirk 1970; Pontryagin et al. 1962). According to this method, the optimal control input can be computed from the solution of a non-linear *two-point boundary value problem*.⁶ While application of PMP is not intractable, it often requires good initialisation, and a sophisticated numerical treatment, to converge to the optimal solution (Betts 1998). Another alternative is to employ differential dynamic programming (Jacobson and Mayne 1970), or the iterative linear quadratic regulator/Gaussian (iLQR/G) framework (Li and Todorov 2004, 2007) that uses a local (quadratic or linear) approximation of the system dynamics and (quadratic approximation of) the objective function to improve computational efficiency. In the present paper we compute the control actions for the non-linear optimal control problem (6) and (8) using the iLQR framework.

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

⁵This is indeed the case in many athletic disciplines (e.g., shot put, discus throw, hammer throw, javelin throw, high jump) executed through explosive actions.

⁶This two-point boundary value problem defines the *stationary point* of the optimal control problem (6) and (8). In this way, it provides a *necessary condition* to find a *locally* optimal solution.

sub-problems are formed by linear approximation of the system dynamics (6),

$$\delta\dot{\mathbf{x}} = \mathbf{f}_{\mathbf{x}}\delta\mathbf{x} + \mathbf{f}_{\mathbf{u}}\delta\mathbf{u}, \quad (9)$$

and quadratic approximation of the objective function (8),

$$\begin{aligned} \Delta J &= h_{\mathbf{x}}^T \delta\mathbf{x}(T) + \delta\mathbf{x}^T(T)h_{\mathbf{xx}}\delta\mathbf{x}(T) \\ &+ \int_0^T (c_{\mathbf{x}}^T \delta\mathbf{x} + c_{\mathbf{u}}^T \delta\mathbf{u}) dt \\ &+ \int_0^T \frac{1}{2} (\delta\mathbf{x}^T c_{\mathbf{xx}} \delta\mathbf{x} + \delta\mathbf{x}^T c_{\mathbf{xu}} \delta\mathbf{u} + \delta\mathbf{u}^T c_{\mathbf{uu}} \delta\mathbf{u}) dt, \end{aligned} \quad (10)$$

evaluated along the current state and control sequence.⁷ This sub-problem, (9) and (10), is solved for $(\delta\mathbf{x}, \delta\mathbf{u})$ via a modified Riccati-like system and a new (improved) sequence is formed by updating the nominal trajectory $\hat{\mathbf{x}} \leftarrow \hat{\mathbf{x}} + \delta\mathbf{x}$ 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-gains⁸ $\mathbf{L}^*(t) \in \mathbb{R}^{m \times 2(n+m)}$. The locally valid⁹ feedback control law for optimal task execution can then be defined as $\mathbf{u} = \mathbf{u}^*(t) + \mathbf{L}^*(t)(\mathbf{x} - \mathbf{x}^*(t))$. By disregarding the feedback corrections, the method provides a feed-forward optimal control sequence defined by: $\mathbf{u} = \mathbf{u}^*(t)$.

2.3.4 Application to compliant control

Using the optimal solution given by $\mathbf{x}^* = (\mathbf{q}^{*T}(t), \dot{\mathbf{q}}^{*T}(t), \boldsymbol{\theta}^{*T}(t), \dot{\boldsymbol{\theta}}^{*T}(t))^T$, the optimal joint torques $\boldsymbol{\tau}^* = \boldsymbol{\tau}(\mathbf{q}^*(t), \boldsymbol{\theta}^*(t))$, and the optimal joint stiffness $\mathbf{K}^* = \mathbf{K}(\mathbf{q}^*(t), \boldsymbol{\theta}^*(t))$ can be obtained from (4) and (5) by substitution. Moreover, if the optimal feedback gains¹⁰ $\mathbf{L}^* = (\mathbf{P}_q(t), \mathbf{D}_q(t), \mathbf{P}_\theta(t), \mathbf{D}_\theta(t))$ can also be computed, then the control inputs will be given in a form of a *locally valid* feedback control law:¹¹ $\mathbf{u} = \mathbf{u}^*(t) + \mathbf{P}_q(t)(\mathbf{q} - \mathbf{q}^*(t)) + \mathbf{D}_q(t)(\dot{\mathbf{q}} - \dot{\mathbf{q}}^*(t)) + \mathbf{P}_\theta(t)(\boldsymbol{\theta} - \boldsymbol{\theta}^*(t)) + \mathbf{D}_\theta(t)(\dot{\boldsymbol{\theta}} - \dot{\boldsymbol{\theta}}^*(t))$. This control law would lead to a local feed-back control of the elastic torque and the associated stiffness properties of the VS actuators. While the feedback correction part may not be

⁷The lower indices in (9) and (10), denote partial derivatives with respect to the corresponding variables.

⁸If the control inputs are saturated (i.e., restricted with “hard constraints” (7)), the corresponding feedback gains may be set to zero, as suggested in Li and Todorov (2007). Alternatively, one could utilise penalty terms to embed the inequality constraints in the objective function (8), see Stengel (1994).

⁹This control law is only optimal in the neighbourhood of the optimal open-loop motion generated by: $\mathbf{u} = \mathbf{u}^*(t)$.

¹⁰ $\mathbf{P}_q \in \mathbb{R}^{m \times n}$, $\mathbf{D}_q \in \mathbb{R}^{m \times n}$, $\mathbf{P}_\theta \in \mathbb{R}^{m \times m}$ and $\mathbf{D}_\theta \in \mathbb{R}^{m \times m}$.

¹¹Note that the feedback correction on \mathbf{u}^* is a PD-control performed with optimal position and velocity gains.

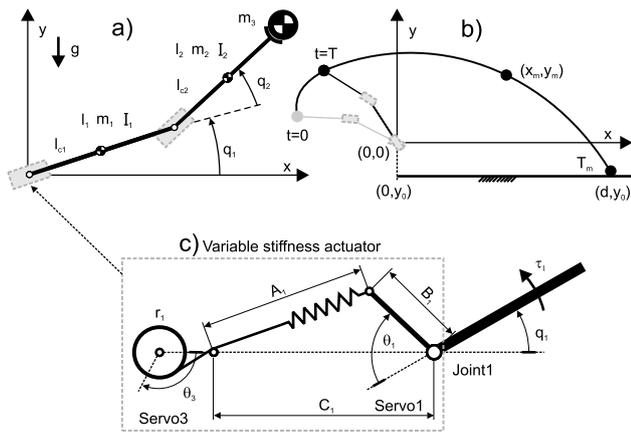


Fig. 1 (a) Configuration of the arm with the inertial (m_1, m_2, m_3, I_1, I_2) and geometric (l_1, l_2, l_{c1}, l_{c2}) parameters. (b) Ball throwing. (c) Prototype of the variable stiffness actuator (identical for both joints). θ_1 (and θ_2) can be used to change the equilibrium position of the link, while θ_3 (and θ_4) adjusts the spring pretension. The physical parameters of the dynamics and the actuators are reported in Table 1

of a central importance for short movements, it may become beneficial when the motion is long and/or when the system dynamics are not well identified. Such feedback introduces additional energy cost and in practise it may also lead to stability issues due to noise and the limited bandwidth of the control loop. In this work we focus on feed-forward motor position control: $\mathbf{u} = \mathbf{u}^*(t)$, to generate the desired joint torque and stiffness, supported with the zero-time-delay feedback given by the intrinsic compliance of the VSAs.

In order to obtain the motor commands and the associated torques and time-varying stiffness profiles, the present framework optimally resolves the *actuation redundancy* (i.e., $\mathbf{u} \in \mathbb{R}^m, m > n$) in a system dependent and task specific way (Todorov 2004). In this light, the actuation redundancy is not only resolved, but optimally exploited to devise the best stiffness profiles for the system (1), actuator (4) and task (8) considered.

2.3.5 Implementation and limitations

In the iLQR implementation, we utilise finite differences to derive the linear approximation of the system dynamics (9), and the quadratic approximation of the cost (10) 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. 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 find the best possible solution.

The model-based iLQR method is a viable optimisation tool when the model of the system dynamics is reasonably

Table 1 Geometric and inertial parameters of the two-link arm and the variable stiffness actuators. The viscous frictional parameters at the joints are defined by b_1 and b_2

(*)	l_* [m]	l_{c*} [m]	m_* [kg]	I_* [kg m ²]	b_* [$\frac{Nms}{rad}$]
1	0.250	0.135	0.42	0.0022	0.01
2	0.305	0.115	0.23	0.0017	0.01
(*)	κ_* [$\frac{N}{m}$]	r_* [m]	B_* [m]	C_* [m]	q_{0*} [rad]
1	771	0.01	0.03	0.125	$-\pi/2$
2	771	0.01	0.03	0.125	0

The mass at the moving end of the arm is defined by: $m_3 = m_b + m_m$ where $m_b = 0.075$ kg is the mass of the ball while $m_m = 0.05$ kg is the mass of the magnet. The range limits on the control inputs/motor positions are $\mathbf{u}_{min} = \theta_{min} \in [-2\pi/5, -\pi/2, 0, 0]^T$ rad and $\mathbf{u}_{max} = \theta_{max} \in [2\pi/5, \pi/2, \pi/2, \pi/2]^T$ rad. The “servo gain parameter” in (6) is set to: $\alpha = \text{diag}(\sqrt{75}, \sqrt{75}, \sqrt{250}, \sqrt{250}) \text{ s}^{-1}$. Unless redefined, these parameters are used in the simulations and experiments

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 over \mathbf{u} , (e.g., see Lagoudakis and Parr 2003; Peters and Schaal 2006). Alternatively, one could also use dynamics learning (e.g., iLQG-LD, see Mitrovic et al. 2010) where the model is acquired from data.

3 Optimal variable stiffness control

In this section, we investigate the proposed optimal control formulation to devise appropriate controllers for dynamic, explosive movements. As an example, we look at the problem of throwing a ball using a two-link arm equipped with variable stiffness actuators. Following the problem formulation, we analyse the control strategies obtained by our approach in the light of varying the objective function and the dynamical properties of the system. Our goal is to verify that the proposed approach is able to predict intuitive but non-trivial behaviours, before going on to analyse the exploitation of variable stiffness in greater depth.

3.1 Problem formulation

Here we present the system dynamics, introduce the variable stiffness actuation mechanism, and define the optimisation criterion for a ball-throwing task.

3.1.1 System dynamics

The configuration of the arm, depicted in Fig. 1a, is specified with joint angles $\mathbf{q} = (q_1, q_2)^T$, while the system dynamics (left hand side of the equation of motion (1)), is specified by the symmetric and positive definite mass matrix \mathbf{M} , the Coriolis and normal inertial terms \mathbf{C} , viscous friction \mathbf{D} and the gravitational terms \mathbf{G} defined by (11):

$$\begin{aligned}
 \mathbf{M} &= \begin{bmatrix} I_1 + m_1 l_{c1}^2 + I_2 + m_2(l_1^2 + 2l_1 l_{c2} \cos(q_2) + l_{c2}^2) + m_3(l_1^2 + 2l_1 l_2 \cos(q_2) + l_2^2) & \times \\ I_2 + m_2(l_{c2}^2 + l_1 l_{c2} \cos(q_2)) + m_3(l_2^2 + l_1 l_2 \cos(q_2)) & I_2 + m_2 l_{c2}^2 + m_3 l_2^2 \end{bmatrix}, \\
 \mathbf{C} &= \begin{bmatrix} -(m_2 l_{c2} + m_3 l_2) l_1 \sin(q_2) (2\dot{q}_1 \dot{q}_2 + \dot{q}_2^2) \\ (m_2 l_{c2} + m_3 l_2) l_1 \sin(q_2) \dot{q}_1^2 \end{bmatrix}, \quad \mathbf{D} = \begin{bmatrix} b_1 \dot{q}_1 \\ b_2 \dot{q}_2 \end{bmatrix}, \\
 \mathbf{G} &= \begin{bmatrix} g(m_1 l_{c1} + m_2 l_1 + m_3 l_1) \cos(q_1) + g(m_2 l_{c2} + m_3 l_2) \cos(q_1 + q_2) \\ g(m_2 l_{c2} + m_3 l_2) \cos(q_1 + q_2) \end{bmatrix}.
 \end{aligned} \tag{11}$$

The geometric and inertial parameters of the model are given in Table 1.

3.1.2 Actuation model

As an example of a physically realisable VSA, we investigate throwing using an arm equipped with variable stiffness actuators (MACCEPA¹²) introduced by Ham et al. (2007). In this actuator (depicted in Fig. 1c), two servo motors are employed at each joint for (i) direct control of the equilibrium point of the actuator $\theta_{1,2}$ and (ii) control of the pretension of the (linear) spring $\theta_{3,4}$. The relations between the motor side positions $\theta = [\theta_1, \theta_2, \theta_3, \theta_4]^T$, joint torque τ and stiffness \mathbf{K} are given by

$$\tau = \underbrace{\begin{bmatrix} \frac{B_1 C_1 \sin \alpha_1}{E_1} & 0 \\ 0 & \frac{B_2 C_2 \sin \alpha_2}{E_2} \end{bmatrix}}_{\mathbf{A}^T} \underbrace{\begin{bmatrix} \kappa_1 (l_{s1} - l_{01}) \\ \kappa_2 (l_{s2} - l_{02}) \end{bmatrix}}_{\mathbf{F}}, \tag{12}$$

$$\mathbf{K} = \begin{bmatrix} k_1 & 0 \\ 0 & k_2 \end{bmatrix} = \begin{bmatrix} -\frac{\partial \tau_1}{\partial q_1} & 0 \\ 0 & -\frac{\partial \tau_2}{\partial q_2} \end{bmatrix}, \tag{13}$$

where $B_i, C_i, \alpha_i = \theta_i - q_i + q_{0i}$ 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, $l_{si} = r_i \theta_{i+2} + E_i$ and $l_{0i} = C_i - B_i$ are the extended and unextended length of the springs while κ_i are the spring constants (see Fig. 1c for an illustration and Table 1 for the specific parameter values).

By changing the motor positions θ , the torque (12) and the stiffness (13) of the joints can be simultaneously modulated, however, these relations are highly non-linear (see Fig. 2) and, 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 limitations (7),

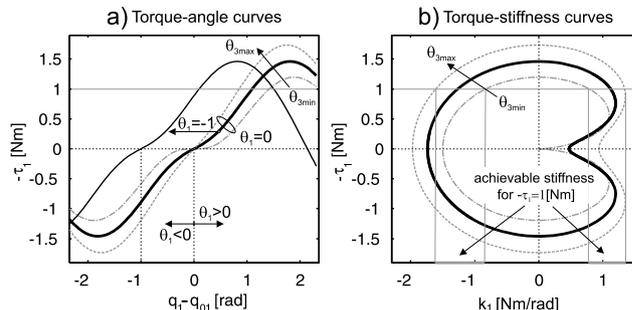


Fig. 2 Depicted is: (a) the torque–angle and (b) the torque–stiffness characteristic of the (first) joint. By changing θ_1 (through control input u_1), the torque–angle curve translates in figure (a) without affecting the torque–stiffness characteristic of the joint depicted in (b). The torque–stiffness curve can be altered by modulating the pretension of the spring, using θ_3 (controlled through u_3). This makes the present actuator a variable stiffness actuator. Note that, if θ_3 is kept constant (i.e., $u_3 = \theta_{d3} = \theta_3$), the torque–stiffness curve in (b) is fixed and simultaneous torque–stiffness modulation is not possible

along with the non-linearity of the dynamics, make (motion, torque and stiffness) control of this system non-trivial.

3.1.3 Performance criterion

The performance criterion for the optimal ball throwing task is defined by

$$J_w = -d + \frac{1}{2} \int_0^T (w \|\mathbf{F}\|^2 + \epsilon \|\mathbf{u}\|^2) dt, \tag{14}$$

where $d = d(\mathbf{q}(T), \dot{\mathbf{q}}(T))$ is the distance thrown (measured at the horizontal ground level y_0 , see Fig. 1b), $\|*\|$ denotes the Euclidean norm, $\mathbf{F} = \mathbf{F}(\mathbf{q}, \theta)$ is the spring force, $w \in [0, \infty)$ defines the relative importance of the distance maximisation and effort minimisation terms, $\epsilon \|\mathbf{u}\|^2$ is a small regularisation term (i.e., $0 < \epsilon \ll 1$) while T is the time permitted for task execution (i.e., $t \in [0, T]$). In many dynamic tasks, effort (here represented by the squared spring forces) plays a significant role in selecting the optimal solution, however in *explosive movements* its importance is highly diminished. To show this, we utilise the weighting parameter w to investigate the asymptotic solution of the

¹²Mechanically Adjustable Compliance and Controllable Equilibrium Position Actuators.

¹³In 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 (e.g., see a human arm model of Mussa-Ivaldi et al. (1985) that incorporates bi-articular muscles).

optimisation as the effort term vanishes, i.e., as $w \rightarrow 0$ and $J_w \approx -d$ (see Sect. 3.2.1).

The distance thrown depends on the configuration of the arm at release, and is computed from the ballistic equation of the flying ball

$$d = x_m(\mathbf{q}) + \dot{x}_m(\mathbf{q}, \dot{\mathbf{q}})T_m(\mathbf{q}, \dot{\mathbf{q}}), \tag{15}$$

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 time until the ball hits the ground. The latter is computed as

$$T_m = \frac{1}{g} (\dot{y}_m(\mathbf{q}, \dot{\mathbf{q}}) + \sqrt{\dot{y}_m(\mathbf{q}, \dot{\mathbf{q}})^2 + 2g(y_m(\mathbf{q}) - y_0)}) \tag{16}$$

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.

3.2 Optimal solutions

In the following, we confirm that our methodology is able to find optimal solutions adapted to the problem setup. In particular, we first look at how varying the objective function (i.e., varying w in (14)) 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) modulates the behaviour found by the present approach.

3.2.1 Variation of the objective function

First, we look at how the solutions found by the proposed framework depend on the choice of the relative importance of the distance and effort terms (defined by the weighting parameter w in (14)).

Typical simulation results are depicted in Fig. 3. As can be seen, depending on w , a variety of optimal behaviours are obtained (Fig. 3a) that are characterised by different velocity, torque and stiffness profiles (Fig. 3d–f). As expected, lower w (decreased penalisation of effort) results in larger ball velocity and longer distance throws (Fig. 3b, c). This exemplifies a *performance-effort trade-off* that characterises many dynamic movement tasks. However, it is interesting to note that, as the penalisation of effort diminishes (i.e., $w \rightarrow 0$) there is an asymptotic behaviour of the distance thrown and the release velocity (Fig. 3b, c) to the maximal values. The optimal solution associated to this limiting case corresponds to the *explosive ball-throwing task*, see Fig. 3a (black line, $w = 10^{-6}$).

It is also interesting to note that a common strategy emerges in the movements, irrespective of the w chosen. This is a characteristic *counter-movement* action (Cho 2004)

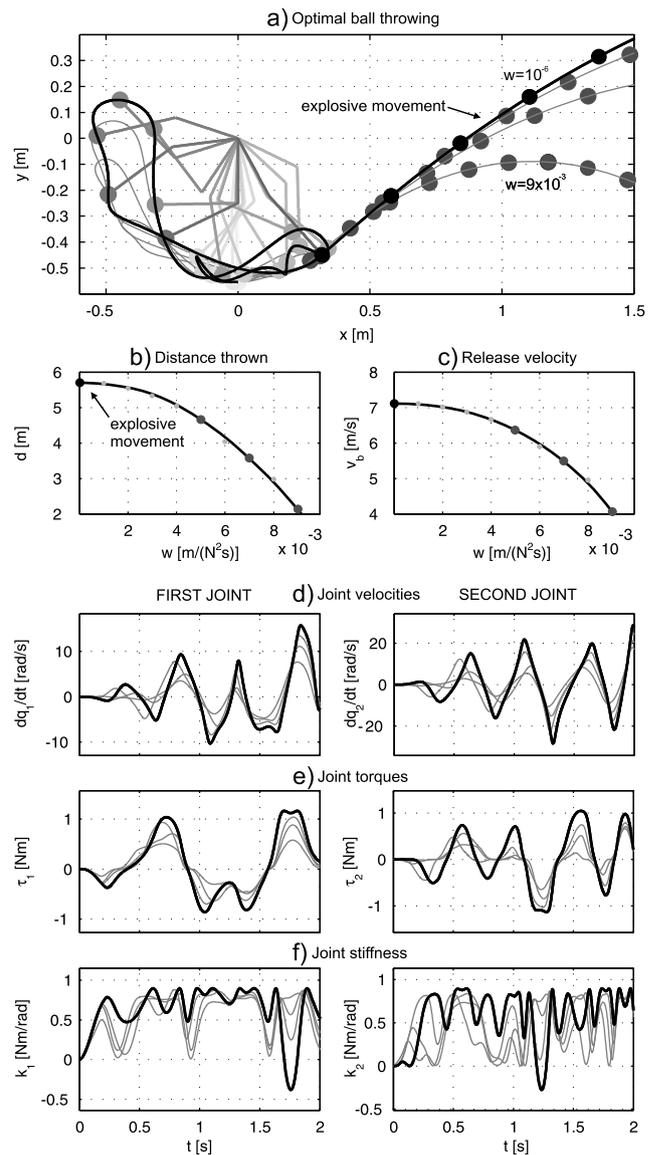


Fig. 3 (a) Stroboscopic view of the optimal throwing motion: explosive movement $w = 10^{-6}$ (black line) and $w = 10^{-3} \times \{5, 7, 9\}$ (grey lines). (b) Distance thrown. (c) Ball velocity. (d) Joint velocities. (e) Joint torques. (f) Joint stiffness. The throwing motions are performed for a fixed time duration $T = 2$ s. The arm is initialised from the vertical resting configuration, $\mathbf{q}(0) = (-\pi/2, 0)^T$, $\dot{\mathbf{q}}(0) = \mathbf{0}$, $\boldsymbol{\theta}(0) = \mathbf{0}$, $\dot{\boldsymbol{\theta}}(0) = \mathbf{0}$

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) (Komi 1992; Schenau et al. 1997). The numerical predictions depicted in Fig. 3, obtained for a simple robotic device, are consistent with this biologically plausible realisation.

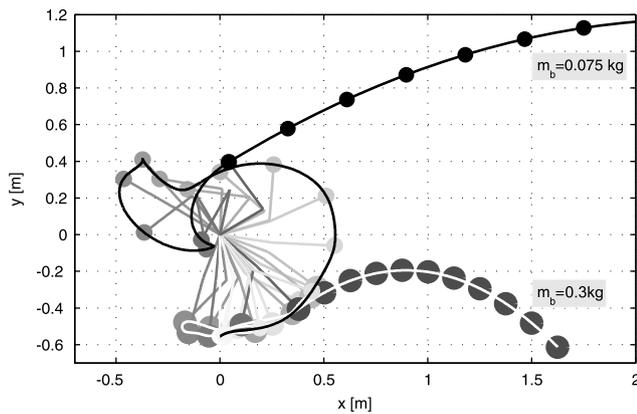


Fig. 4 Throwing movement predicted for two different masses: $m_b = 0.075$ kg (black line) and $m_b = 0.3$ kg (white line). The optimisation selects an under-arm strategy for the heavier ball and an over-arm strategy for the light ball. The arm is initialised from the vertical resting configuration. The duration of the motions is fixed to $T = 1.5$ s, and we selected $w = 10^{-6}$ to specify the relative importance of the effort during the task execution. In order to make the system capable (i.e., strong enough) of over-arm throwing, the basic parameters given in Table 1 are modified such that: the equilibrium configuration and the stiffness of the springs is set to: $\mathbf{q}_0 = (\pi/6, \pi/6)^T$ rad and $\kappa = (1542, 1542)^T$ N/m respectively

3.2.2 Variation of the dynamics

In this numerical investigation, we look at how changes to the dynamics of the system affect the solutions found by the proposed framework. Here, we analyse the effect of changing the mass of the ball on the throwing motion. A representative result is shown in Fig. 4. As can be observed, for two different masses (i.e., $m_b = 0.075$ kg and $m_b = 0.3$ kg) very different throwing strategies emerge. In particular, for the heavy ball an under-arm movement is predicted, while for the lighter ball an over-arm strategy is obtained.

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, and if w is non-negligible in (14), then under-arm throwing may be preferable from the optimisation point of view. This is because lifting a heavy ball requires significant effort penalised by the second term in the control objective (14). This is also the case when the actuators are weak and, as such, not capable of lifting the ball up. In that case, under-arm throwing is the only viable strategy, even in the case of explosive movements (i.e., when effort is not penalised). Following the same argument, if the ball is light, over-arm throwing is expected to lead to larger distance thrown. 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 (Bingham 1988). 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 when attempting to send the ball over a large distance (as, for example, when fielding in cricket or baseball). The result presented in Fig. 4 demonstrates that the present optimal control approach is able to find such *strategy change in task execution* depending on the dynamics (i.e., weight of the ball).

4 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 tasks.¹⁴ Here we explore whether such benefits can arise from the ability of VSAs to simultaneously modulate joint torque and stiffness, and to amplify power (store energy).

4.1 Benefit of stiffness variation

To investigate the benefit of optimal control with variable stiffness, we compare the throwing performance of the same device using:

- optimal variable stiffness control*—where the joint torque and the joint stiffness are independently and optimally modulated through control, and
- optimal fixed stiffness control*—where the joint torque and the joint stiffness cannot be independently modulated but are simultaneously optimised during the motion.

Specifically, while in the former case all components in $\mathbf{u} = [u_1(t), u_2(t), u_3(t), u_4(t)]^T$ are optimised in time to independently control the joint torques and the joint stiffness, in the latter case commands to the pre-tensioning servos are fixed to optimal constant values: $\mathbf{u} = [u_1(t), u_2(t), u_3, u_4]^T$ (i.e., $u_{3,4} = \text{const.}$). Note that, on the MACCEPA keeping $u_{3,4}$ constant does not ensure *constant joint stiffness*, but it does ensure that the joint torque and stiffness cannot be *independently* optimised (see Fig. 2b in Sect. 3.1.2).

A comparison between optimal variable and optimal fixed stiffness control is provided in Fig. 5. In Fig. 5a1, a2 we see that the throwing performance is improved for the case when both torque and stiffness are simultaneously optimised ($d = 5.3$ m, $J_w \approx -5.3$) as opposed to using an optimal fixed torque–stiffness relation ($d = 4.3$ m, $J_w \approx -4.3$). The difference between the activation patterns can be clearly seen by comparing Fig. 5c1, c2 where, in the former case,

¹⁴As an example, human peak performance as characterised by the rotation speed of the shoulder during a baseball pitch of a professional pitcher, is between 6900–9800°/s (Herman 2007). This kind of high-performance task execution is not in the scope of present robotic systems.

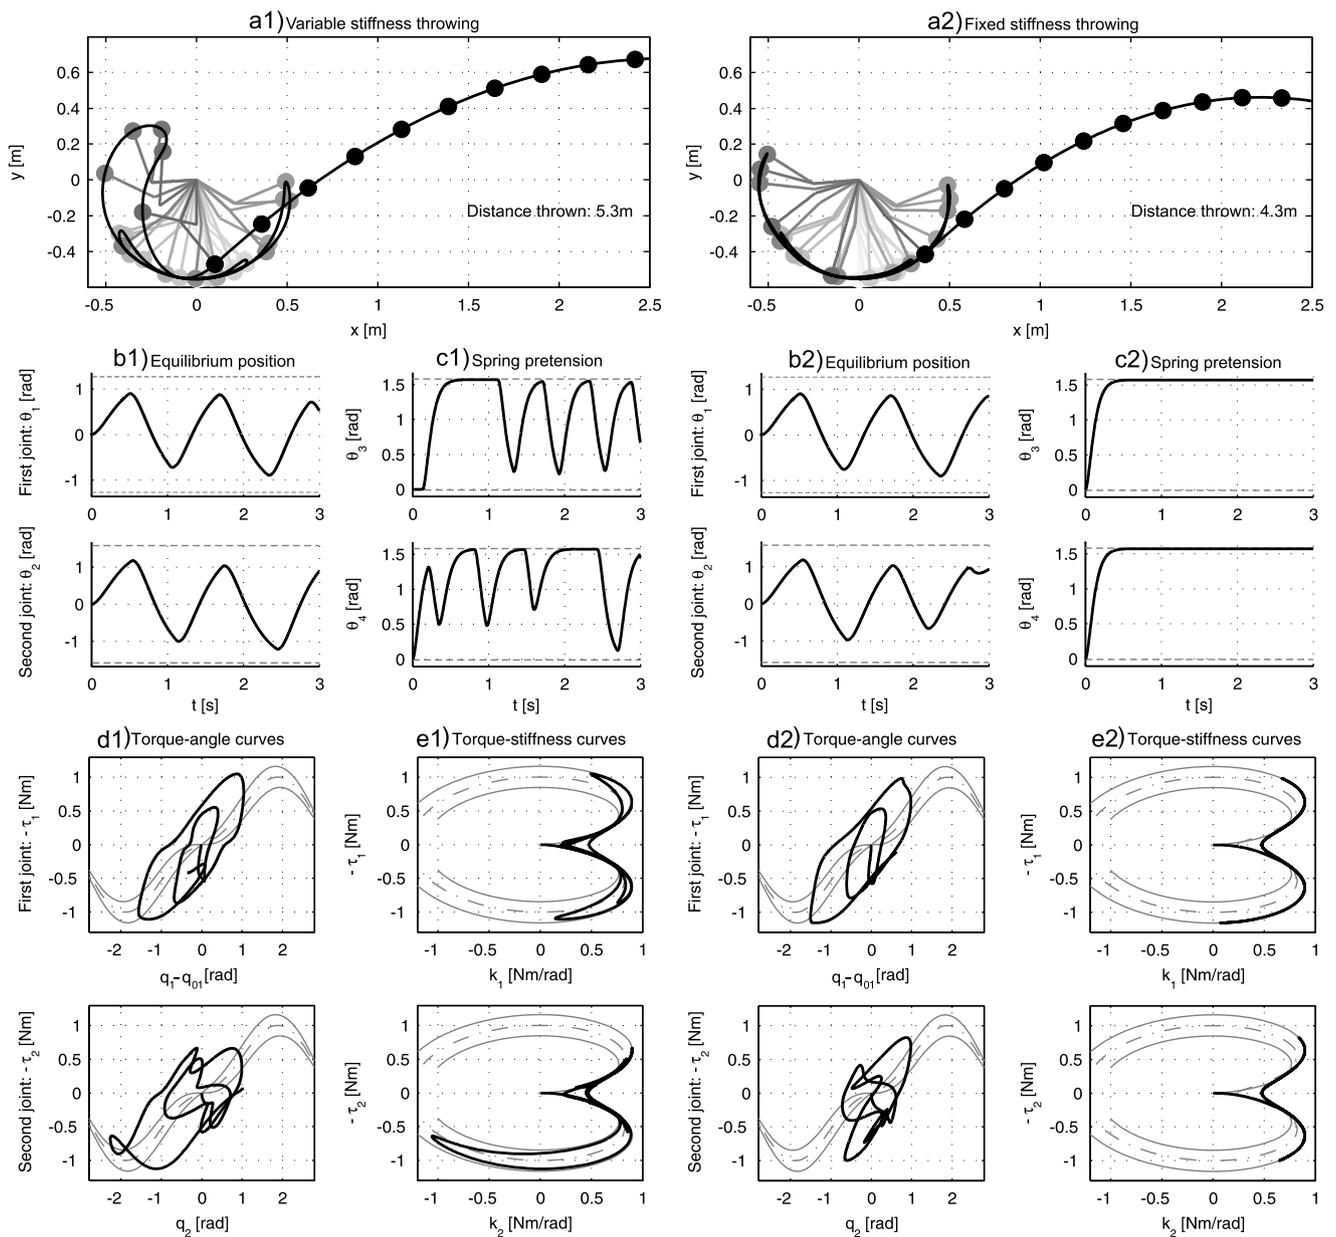


Fig. 5 Exploitation of variable stiffness actuation. Stroboscopic view of the throwing motion: **(a1)** optimal variable stiffness control (i.e., independent torque/stiffness control) and **(a2)** control with fixed torque–stiffness relation ($u_{3,4} = \text{const.}$ —non-independent torque/stiffness control). **(b1, b2)** Optimal equilibrium position commands. **(c1, c2)** Optimal spring pretension commands. (Note, the optimal constant values, used in **(c2)**, correspond to the maximal pretension, $u_3 = u_4 = \pi/2$ rad). **(d1, d2)** Optimal torque–angle curves. **(e1, e2)** Optimal torque–stiffness curves. The *grey lines* in **(e1, e2)**

correspond to minimum (inner loop) mid-range (*dot/dashed lines*) and maximum (outer loop) pre-tension (stiffness), respectively. The corresponding *grey lines* in **(d1)** and **(d2)** correspond to the passive torque–angle curves on the joints at zero equilibrium position commands (i.e., $u_{1,2} = 0$). For the purpose of demonstration, we consider an explosive task specified with $w = 10^{-6}$, $T = 3$ s and $\alpha = \text{diag}(5, 5, \sqrt{250}, \sqrt{250}) \text{ s}^{-1}$ for these two systems. The arm is initialised from a vertical resting configuration with fully relaxed springs

the spring pre-tension (and thereby the stiffness) is modulated by control throughout the movement, while in the latter, it is not modulated by control.¹⁵

Looking at the optimal solution with no independent torque–stiffness modulation, it appears that stiffness is maximised: in Fig. 5e2 we see that the solution remains on the constant-command iso-line where the stiffness is greatest. On the MACCEPA this means that the torque–stiffness curve selected by optimisation is the one that gives the

¹⁵There is a transient at the start of the movement as the pre-tensioning motors move to the optimal (but fixed) commanded positions.

largest torque range to the joint. In contrast, if the pretension is allowed to vary, the functional relation between torque and stiffness is modulated periodically to improve task performance (see Fig. 5c1, e1). During this modulation, the torque range, and more importantly the stiffness of the joint, is reduced, allowing the arm to be “more decoupled” from the actuators and move more freely during the task execution. This leads to larger motion range and improved task performance (i.e., larger distance thrown).

During the task execution, power is amplified by cyclically extending and compressing the springs while moving the arm back and forth. The physical implication of such strategy is analysed in greater depth in the next section.

4.2 The benefit of passive compliance

One of the distinct feature of VSAs that incorporate passive elasticity into the system by design, is their ability to store mechanical energy and to utilise the stored energy to enhance the power output of the actuators. This may be achieved by making the elastic components absorb the energy generated by the motors at a low rate, and then releasing this energy at a high rate to drive the link-side motion (Bingham 1988; Paluska and Herr 2006). This is particularly important for explosive movements where it can significantly enhance the peak joint performance (Alexander and Bennet-Clark 1977; Jöris et al. 1985; Wilson et al. 2003).

Here, we show that the present optimal control formulation naturally exploits this physical mechanism to improve task performance. For this purpose, we first define the conditions required for energy storage and power amplification on compliantly actuated systems and then identify the presence of these physical effects along the optimal solution.

4.2.1 Conditions for power amplification and energy storage

During the motion of the system, the input power delivered by the motors is given by

$$p_{in} = p_{out} + \dot{E}_s = \boldsymbol{\tau}^T \dot{\mathbf{q}} + \dot{E}_s(\mathbf{q}, \boldsymbol{\theta}), \quad (17)$$

where p_{out} is the mechanical power output of the VSA, while \dot{E}_s is the time derivative of the elastic energy accumulated in the springs.¹⁶ Using (17), we can define two distinct operation modes depending on the energy flow during the motion; the first is

- (a) *power amplification*, that takes place if the (positive) output power delivered by the VSA, is higher than the

- input power provided by the motors i.e., $p_{out} > 0$ and $p_{out} > p_{in}$ ($\dot{E}_s < 0$), while the second is
 (b) *energy storage*, which occurs when the VSA is back-driven by the rigid body dynamics. In this latter case, the output power is negative i.e., $p_{out} < 0$ and the energy given by the link-side motion is stored by the actuators $\dot{E}_s > 0$.

By alternating between these two modes, energy stored by the springs in an earlier phase of motion can be re-utilised to enhance the output power of the actuators in a latter phase (e.g., when releasing the ball).

4.2.2 Optimal power flow and energy storage during the motion

The cyclic motion observed in Fig. 5 belies a strategy of alternating between the power amplification (light gray) and energy storage (black) modes of operation as indicated in Fig. 6. There are two characteristic power peaks just before ball-release indicating: (i) *significant power amplification* at the end of the movement and (ii) a *proximal-to-distal power flow*, see Fig. 6b1, b2. Such (an optimal) proximal-to-distal power flow was suggested to be connected with the characteristic sequential action of body segments from larger proximal to smaller distal links observed in humans (Jöris et al. 1985; Putnam 1993). In our passively compliant system, this serves to gradually increase the peak ball velocity with each swing until the final release (Fig. 6a).

Looking at Fig. 6c, we see that throughout the movement, the *mechanical energy* input to the actuators (E_{in}), is lower than that required for the *same motion* realised: (i) without exploiting the energy storage effects (e.g., using *non-compliant actuators* E_o) and (ii) without exploiting the inertial and gravitational effects during the motion (e.g., using *non-backdrivable actuators* E_+). This example demonstrate the benefit passive compliance may provide by enabling the kinetic energy and the gravitational potential energy to be stored during the movement. While this energy storing mechanism may contribute to improve *mechanical efficiency* of the system, it does not imply that compliant actuators would consume less *electrical energy* compared to their non-compliant counterparts. Indeed, this may heavily depend on how much the actuators are actively used, but also the efficiency of the gear trains, motors and the power electronics employed. On the other hand, it is clear that non-compliant actuators cannot provide power amplification (i.e., the output power cannot exceed the input power provided by the motors: $\dot{E}_s = 0$, $p_{out} = p_{in}$). This highlights the benefit that *passively* compliant actuators can provide during explosive movements, where high power output (possibly obtained by optimal power amplification) is necessary for effective task execution (Jöris et al. 1985; Newton et al. 1996).

¹⁶The potential energy stored by the linear springs in the present actuators is computed as: $E_s(\mathbf{q}, \boldsymbol{\theta}) = \frac{1}{2} \mathbf{F}(\mathbf{q}, \boldsymbol{\theta})^T \mathbf{K}_s^{-1} \mathbf{F}(\mathbf{q}, \boldsymbol{\theta})$ where \mathbf{F} is the spring force while $\mathbf{K}_s = \text{diag}(\kappa_1, \kappa_2)$ is the matrix of the spring stiffness constants.

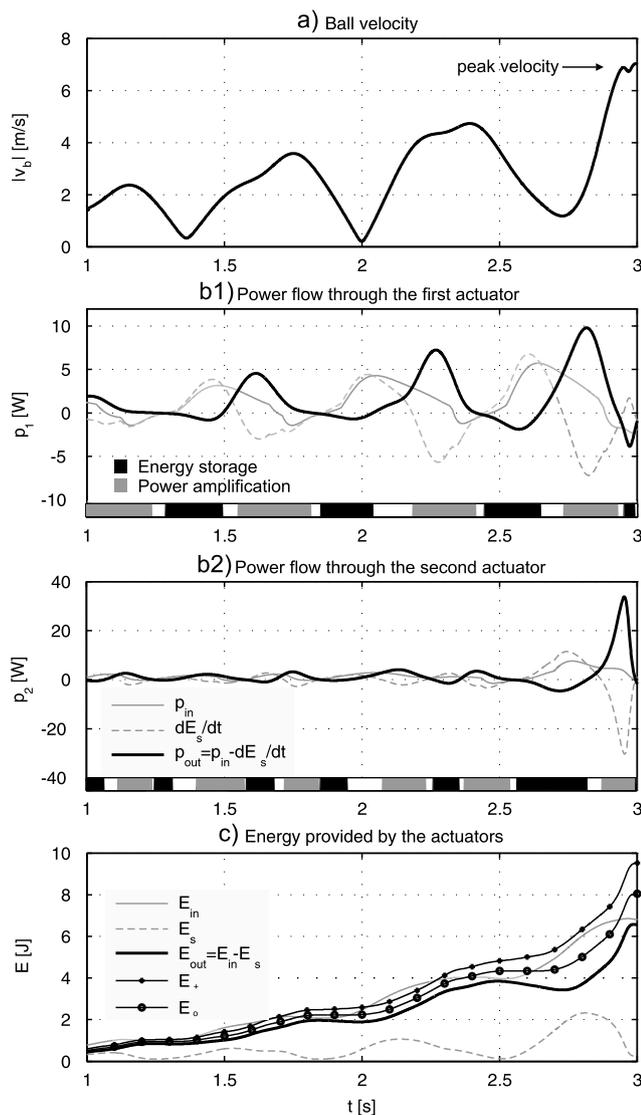


Fig. 6 (a) Velocity of the ball during corresponding to the throwing motion depicted in Fig. 5. (b1, b2) Mechanical power flow through the actuators. The power flow is decomposed into the input power p_{in} , output power p_{out} and the power flow through the springs \dot{E}_s (i.e., $p_{out} = p_{in} - \dot{E}_s$). (c) Mechanical energy supplied by the motors E_{in} , energy stored in the actuators E_s (i.e., potential energy of the springs), and output energy delivered by the VS actuators E_{out} (i.e., $E_{in} = E_s + E_{out}$ where $E_{out} = \int_0^t p_{out}(s)ds$). The same plots also contain the input mechanical energy that would be ideally required to generate the same motion by: (i) non-compliant but back-drivable actuators that are not able to store energy $E_o(t) = \frac{1}{2} \int_0^t (p_{1out} + |p_{1out}| + p_{2out} + |p_{2out}|)ds$, and (ii) non-back-drivable actuator where both acceleration and braking are actively generated $E_+(t) = \int_0^t (|p_{1out}| + |p_{2out}|)ds$

It is important to note that the present objective function (14), used to generate the throwing motion, does not directly encode the observed sequential energy-storing and power-amplification strategy. This strategy emerges from the optimisation that is able to exploit the coupling between the dynamics and the actuators in the present formulation.

5 Experiments: optimal variable stiffness versus optimal fixed stiffness control

In this section, we present ball-throwing experiments using a *two-link variable stiffness robot*, see Fig. 7a. This device is capable of simultaneous and independent joint torque and joint stiffness modulation using VS actuators (Ham et al. 2007). Each actuator is realised with two servomotors (Hitec HSR-5990TG) per joint, controlled with 50 Hz PWM signals from a micro-controller (ATmega2560). The joint angles are measured by rotary potentiometers (Alps RDC503013A). The throwing experiments are performed with a tennis ball (fitted with a magnetic plate) that weighs $m_b = 0.075$ kg. During motion, the ball is held with an electromagnet (Magnet-Schultz, GMHX030X00D02) mounted at the end of the arm (see Fig. 7a) and released at the final instant $t = T$.

The purpose of the experimental is to: illustrate the principles of variable stiffness control, demonstrate applicability of the present formulation to real-world problems, and provide evidence that supports the numerical predictions obtained by simulations. With regard to the last point, our aim is to confirm that variable stiffness control can be used to improve task performance (versus fixed stiffness control) under experimental conditions.

The experimental results reported below correspond to the simulations presented in Fig. 5. Specifically, the experiment corresponding to variable stiffness control is depicted in Fig. 7b, c. In Fig. 7b, we observe a reasonable match between the simulated and the real behaviour, both in motion and synchronisation ensuring near-to-optimal timing of the ball release. In Fig. 7b2 we can see that stiffness modulation takes place on both of the joints. Note that by increasing the stiffness (increasing the stiffness commands $\theta_{3,4}$) the actuators can couple the motion of the links, while by decreasing the stiffness the actuators will not impede the motion of the robot. Both of these can be beneficial during the movement, namely, while the former allows the actuators to transfer torques more effectively, the latter enables the rigid body dynamics to extend the motion range. The throwing performance obtained on the device (distance thrown: $d = 5.1$ m) is reasonably close to that predicted by the simulation ($d = 5.3$ m, see Fig. 5a1). The difference is mainly due to the minor mismatches between the real hardware and the idealised modelling assumptions. These issues, however, neither adversely affected the coordination pattern during the motion (please see the experimental video), nor significantly altered the throwing performance. Moreover, despite the natural sensitivity of the thrown distance to delays in timing (Chowdhary and Challis 1999), the experimental performance is close to the one obtained under idealistic conditions in simulation. We note that this result was obtained with *open-loop execution of the optimal commands*, using

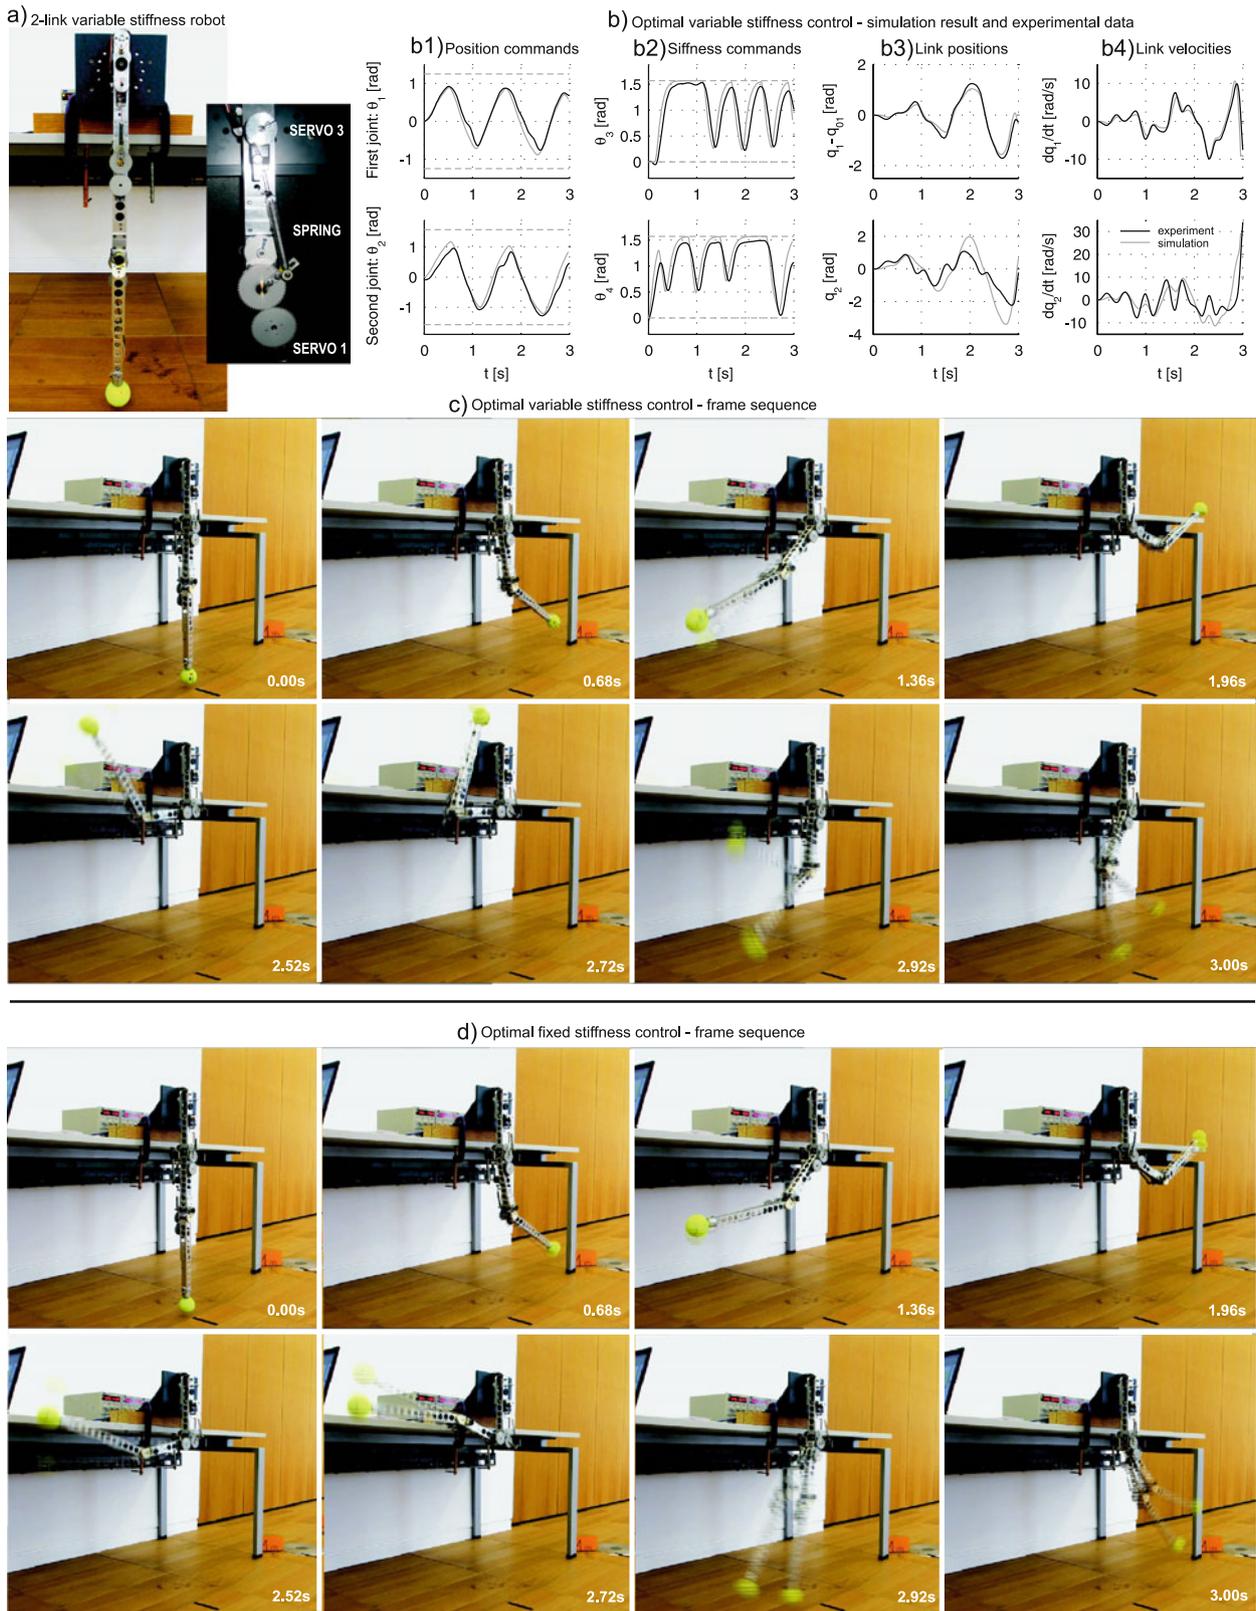


Fig. 7 (a) Two-link variable stiffness arm. (b) Simulated (grey lines) and measured (black lines): (b1, b2) motor positions θ , (b3) joint angles q and (b4) velocities \dot{q} , of variable stiffness throw. (c) Optimal variable stiffness control—snapshots of the throwing motion. (d) Opti-

mal fixed stiffness control. The presented experiments for variable and fixed stiffness control correspond to the simulation results presented in Fig. 5. For the corresponding videos please see the supplementary material

no active feedback. This strategy is often argued to be employed by humans during fast movements (van Soest and Bobbert 1993), and may also be preferred on compliantly actuated robots having slow actuator dynamics.

In addition to the above experiment, we have performed a fixed stiffness throwing experiment, see Fig. 7d. In Fig. 7c, d we can compare the variable stiffness and the corresponding fixed stiffness control experiments. As predicted by the simulation study (see Sect. 4.1, Fig. 5a1, a2), variable stiffness control provides a clear performance benefit (experimental: $d = 5.1$ m, simulated: $d = 5.3$ m) compared to the corresponding fixed stiffness control (experimental: $d = 4$ m, simulated: $d = 4.3$ m). For qualitative assessment of the realised motion, the reader may refer to the frame sequence depicted in Fig. 7c, d and the corresponding experimental video provided in the supplemental material.

6 Discussion: implications for explosive movement tasks

Hogan (1984) suggested that humans may increase the stiffness of their limb (by antagonistic co-activation) to maintain an unstable upright posture in an uncertain environment. While this strategy is often argued to be *energetically expensive*,¹⁷ it may turn out to be *essential* if, for e.g., active torque (stiffness) control is not viable under feedback delays. Although the utility of such a stiffness control strategy has been shown in humans (Burdet et al. 2001; Mussa-Ivaldi et al. 1985) for (unstable) *static tasks*, the benefit of independent torque and stiffness modulation *during movement* is unclear and has yet to be demonstrated.

In this paper, we predict through simulation studies (see Sect. 4.1), and demonstrate in hardware experiments (see Sect. 5), that *independent* torque and stiffness modulation provides performance improvement in a highly dynamic *explosive movement task*, compared to the alternative strategy where torque and stiffness are not modulated independently, see also Braun et al. (2011). It is important to note, however, that such controlled stiffness modulation is often energetically expensive, and that unlike in unstable and static tasks, it is *not essential* to realise explosive dynamic movements.¹⁸ Due to these reasons, one of the central questions to be answered is *whether the task performance improvement provided by independent stiffness modulation justifies the associated effort (energy) cost*. This may not be the case for task where *effort minimisation* is an essential part of the control objective (e.g., $w \gg 1$ in (14)). In explosive movements,

however, details of the (optimal) control strategy may *not* be devised by minimising an effort cost (e.g., $w \rightarrow 0$ in (14)). Accordingly, in explosive tasks, the benefit provided by independent stiffness modulation is *not conditioned* on the effort cost it requires.

It is important to note however that the above is *not sufficient* to make independent stiffness modulation beneficial. This is because the benefit provided by this control modality, *during movement*, will also depend on *dynamic features of the actuators*: for e.g., how fast can stiffness be changed. This additional consideration may not be important for static tasks, but could be crucial to achieve a *rapidly varying (desired) stiffness profile* during fast movements. For this reason, the answer to whether independent torque stiffness modulation is beneficial during movement does not only depend on whether it can be achieved, but also on *how* (fast) it can be realised. This implies that before any theories can be postulated about the benefit of stiffness modulation in biological actuation based on this study, the appropriate modelling of the corresponding actuator dynamics (Hill 1938; Winters and Stark 1985) and their characteristics are in order.

7 Conclusion

In this article, we demonstrate the utility of an optimal control formulation applied to compliantly actuated robotic systems. Using this formulation, we devised optimal variable stiffness control strategies that exploit the system dynamics, often in a non-intuitive way, that would be difficult to obtain through hand-tuning and other non-algorithmic methods. In addition, we have presented an analysis of these results in the light of the energy-storage and power amplification ability of compliant actuators, and demonstrated the benefits of independent torque and stiffness control both in simulation and experiment. Related to this latter result, we discussed: why is stiffness modulation justified despite its inherent cost, and under what condition could it emerge during fast movements.

In future work, we intend to: (i) extend this framework by considering other impedance terms, such as damping, and (ii) further investigate the role of variable impedance during dynamic tasks. The proposed computational framework will be a key tool towards extending these results.

Acknowledgements This work was funded by the EU Seventh Framework Programme (FP7) as part of the STIFF project. The authors gratefully acknowledge this support. We would like to thank Alexander Enoch for his work on the hardware design and Andrius Sutas for his contribution to the control interface. In addition, we thank Dr. Jun Nakanishi and Dr. Takeshi Mori for fruitful discussions regarding this work.

¹⁷During antagonistic co-activation in humans, muscles do no mechanical work but consume metabolic energy.

¹⁸Indeed, an explosive movement can be executed by feed-forward joint torque modulation, e.g., realised by optimal fixed stiffness control, as shown in Sects. 4.1 and 5.

References

- Alexander, R. M., & Bennet-Clark, H. C. (1977). Storage of elastic strain energy in muscle and other tissues. *Nature*, 265, 114–117.
- Anderson, F. C., & Pandy, M. G. (2001). Dynamic optimization of human walking. *Journal of Biomechanical Engineering*, 123, 381–390.
- Anderson, R., & Spong, M. (1988). Hybrid impedance control of robotic manipulators. *IEEE Journal of Robotics and Automation*, 4(5), 549–556.
- Bellman, R. (1957). *Dynamic programming*. Princeton: Princeton University Press.
- Betts, J. T. (1998). Survey of numerical methods for trajectory optimization. *AIAA Journal of Guidance, Control and Dynamics*, 21(2), 193–207.
- Bicchi, A., & Toniatti, G. (2004). Fast and soft arm tactics: dealing with the safety-performance trade-off in robot arms design and control. *IEEE Robotics and Automation Magazine*, 11, 22–33.
- Bingham, G. P. (1988). Task-specific devices and the perceptual bottleneck. *Journal of Human Movement Science*, 7, 255–264.
- Bobrow, J. E., Dubowsky, S., & Gibson, J. S. (1985). Time-optimal control of robotic manipulators along specified paths. *International Journal of Robotics Research*, 4(3), 3–17.
- Braun, D. J., Howard, M., & Vijayakumar, S. (2011). Exploiting variable stiffness in explosive movement tasks. In *Proceedings of robotics: science and systems*, Los Angeles, CA, USA.
- Bryson, A. E., & Ho, Y. C. (1975). *Applied optimal control*. Washington: Hemisphere/Wiley.
- Burdet, E., Osu, R., Franklin, D. W., Milner, T. E., & Kawato, M. (2001). The central nervous system stabilizes unstable dynamics by learning optimal impedance. *Nature*, 414, 446–449.
- Cho, A. (2004). To throw farther, waste energy. *Science*, 306(5693), 42–43.
- Chowdhary, A., & Challis, J. (1999). Timing accuracy in human throwing. *Journal of Theoretical Biology*, 201(4), 219–229.
- Collins, J. J. (1995). The redundant nature of locomotor optimization laws. *Journal of Biomechanics*, 28(3), 251–267.
- English, C. E. (1999a). Implementation of variable joint stiffness through antagonistic actuation using rolamite springs. *Mechanism and Machine Theory*, 341, 27–40.
- English, C. E. (1999b). Mechanics and stiffness limitations of a variable stiffness actuator for use in prosthetic limbs. *Mechanism and Machine Theory*, 341, 7–25.
- Flash, T., & Hogan, N. (1985). The coordination of arm movements: an experimentally confirmed mathematical model. *Journal of Neuroscience*, 5, 1688–1703.
- Garabini, M., Passaglia, A., Belo, F. A. W., Salaris, P., & Bicchi, A. (2011). Optimality principles in variable stiffness control: the VSA hammer. In *Proceedings of the IEEE/RSJ international conference on intelligent robots and systems*, San Francisco, USA.
- Haddadin, S., Weis, M., Wolf, S., & Albu-Schäffer, A. (2011). Optimal control for maximizing link velocity of robotic variable stiffness joints. In *Proceedings of the 18th IFAC world congress, Part 1* (Vol. 18).
- Ham, R. V., Vanderborght, B., Damme, M. V., Verrelst, B., & Lefeber, D. (2007). MACCEPA, the mechanically adjustable compliance and controllable equilibrium position actuator: design and implementation in a biped robot. *Robotics and Autonomous Systems*, 55(10), 761–768.
- Herman, I. P. (2007). *Physics of the human body*. Berlin: Springer.
- Hill, A. V. (1938). The heat of shortening and the dynamic constants of muscle. *Proceedings of the Royal Society B*, 126, 136–195.
- Hogan, N. (1984). Adaptive control of mechanical impedance by coactivation of antagonist muscles. *IEEE Transactions on Automatic Control*, AC-29(8), 681–690.
- Hogan, N. (1985). Impedance control: an approach to manipulation. *ASME Journal of Dynamic Systems, Measurement and Control*, 107, 1–24.
- Hurst, J. W., Chestnutt, J., & Rizzi, A. A. (2010). The actuator with mechanically adjustable series compliance. *IEEE Transactions on Robotics*, 26(4), 597–606.
- Ikeura, R., Moriguchi, T., & Mizutani, K. (2002). Optimal variable impedance control for a robot and its application to lifting an object with a human. In *Proceedings of the IEEE international workshop on robot and human interactive communication*.
- Jacobson, D. H., & Mayne, D. Q. (1970). *Differential dynamic programming*. New York: Elsevier.
- Johansson, R., & Spong, M. (1994). Quadratic optimisation of impedance control. In *Proceedings of the IEEE international conference on robotics and automation*, San Diego, CA, USA (pp. 616–621).
- Jöris, H. J. J., van Muyen, A. J. E., van Ingen Schenau, H. C. G., & Kemper, G. J. (1985) Force, velocity and energy flow during the overarm throw in female handball players. *Journal of Biomechanics*, 18(6), 409–414.
- Kim, B. S., & Song, J. B. (2010). Hybrid dual actuator unit: A design of a variable stiffness actuator based on an adjustable moment arm mechanism. In *Proceedings of the IEEE international conference on robotics and automation*, Anchorage, Alaska, USA (pp. 1655–1660).
- Kirk, D. E. (1970). *Optimal control theory: an introduction*. New York: Prentice-Hall.
- Koganezawa, K., Watanabe, Y., & Shimizu, N. (1999). Antagonistic muscle-like actuator and its application to multi-dof forearm prosthesis. *Advanced Robotics*, 12(7–8), 771–789.
- Komi, P. V. (1992). Stretch-shortening cycle. The encyclopaedia of sports medicine. In *Strength and power in sport*, Oxford: Blackwell Scientific.
- Lagoudakis, M. G., & Parr, R. (2003). Least-squares policy iteration. *Journal of Machine Learning Research*, 4, 1107–1149.
- Laurin-Kovitz, K. F., Colgate, J. E., & Carnes, S. D. R. (1991). Design of components for programmable passive impedance. In *Proceedings of the IEEE international conference on robotics and automation* (Vol. 2, pp. 1476–1481).
- Li, W., & Todorov, E. (2004). Iterative linear-quadratic regulator design for nonlinear biological movement systems. In *Proceedings of the 1st international conference on informatics in control, automation and robotics* (Vol. 1, pp. 222–229).
- Li, W., & Todorov, E. (2007). Iterative linearization methods for approximately optimal control and estimation of non-linear stochastic system. *International Journal of Control*, 80(9), 1439–1453.
- Matinfar, M., & Hashtrudi-Zaad, K. (2005). Optimisation-based robot compliance control: Geometric and linear quadratic approaches. *International Journal of Robotics Research*, 24(8), 645–656.
- Mettin, U., Shiriaev, A. S., Freidovich, B., & Sampei, M. (2010). Optimal ball pitching with an underactuated model of a human arm. In *Proceedings of the IEEE international conference on robotics and automation*, Anchorage, Alaska, USA (pp. 5009–5014).
- Migliore, S. A., Brown, E. A., & DeWeerth, S. P. (2007). Novel nonlinear elastic actuators for passively controlling robotic joint compliance. *Journal of Mechanical Design*, 129(4), 406–412.
- Mitrovic, D., Klanke, S., & Vijayakumar, S. (2010). *From motor learning to interaction learning in robots: adaptive optimal feedback control with learned internal dynamics models*. *SCI* (Vol. 264). Berlin: Springer.
- Mitrovic, D., Klanke, S., & Vijayakumar, S. (2011). Learning impedance control of antagonistic systems based on stochastic optimization principles. *International Journal of Robotics Research*, 30(2), 1–18.

- Morita, T., & Sugano, S. (1995). Design and development of a new robot joint using a mechanical impedance adjuster. In *Proceedings of the IEEE international conference on robotics and automation*, Nagoya, Japan (Vol. 3, pp. 2469–2475).
- Mussa-Ivaldi, F. A., Hogan, N., & Bizzi, E. (1985). Neural, mechanical, and geometric factors subserving arm posture in humans. *Journal of Neuroscience*, 5, 2732–2743.
- Nelson, W. L. (1983). Physical principles for economies of skilled movements. *Biological Cybernetics*, 46(2), 135–147.
- Newton, R. U., Kraemer, W. J., Hakkinen, K., Humphries, B. J., & Murphy, A. J. (1996). Kinematics, kinetics and muscle activation during explosive upper body movements. *Journal of Applied Biomechanics*, 12, 31–43.
- Paluska, D., & Herr, H. (2006). The effect of series elasticity on actuator power and work output: implications for robotic and prosthetic joint design. *Robotics & Autonomous Systems*, 54, 667–673.
- Pandy, M., Zajac, F., Sim, E., & Levine, W. (1990). An optimal control model for maximum-height human jumping. *Journal of Biomechanical Engineering*, 23, 1185–1198.
- Pandy, M., Garner, B., & Anderson, F. (1995). Optimal control of non-ballistic muscular movements: a constraint-based performance criterion for rising from a chair. *Journal of Biomechanical Engineering*, 117, 15–26.
- Peters, J., & Schaal, S. (2006). Policy gradient methods for robotics. In *Proceedings of the IEEE/RSJ international conference on intelligent robots and systems*, Beijing, China (pp. 2219–2225).
- Pontryagin, L. S., Boltyanskii, V. G., Gamkrelidze, R. V., & Mishchenko, E. F. (1962). *The mathematical theory of optimal processes*. New York: Wiley.
- Putnam, C. (1993). Sequential motions of body segments in striking and throwing skills: descriptions and explanations. *Journal of Biomechanics*, 26(1), 125–135.
- Schenau, G., Bobbert, M. F. & de Haan, A. (1997). Mechanics and energetics of the stretch-shortening cycle: a stimulating discussion. *Journal of Applied Biomechanics*, 13, 484–496.
- Shen, X., & Goldfarb, M. (2007). Simultaneous force and stiffness control of a pneumatic actuator. *Journal of Dynamic Systems, Measurement, and Control*, 129(4), 425–434.
- Shoji, T., Nakaura, S., & Sampei, M. (2010). Throwing motion control of the springed pendubot via unstable zero dynamics. In *Proceedings of the IEEE international conference on control applications: multi-conference on systems and control*, Yokohama, Japan (pp. 1602–1607).
- Siciliano, B., & Khatib, O. (2008). *Handbook of robotics*. Berlin: Springer.
- van Soest, A. J., & Bobbert, M. F. (1993). The contribution of muscle properties in the control of explosive movements. *Biological Cybernetics*, 69, 195–204.
- Stengel, R. F. (1994). *Optimal control and estimation*. New York: Dover.
- Todorov, E. (2004). Optimality principles in sensorimotor control. *Nature Neuroscience*, 7(9), 907–915.
- Tonietti, G., Schiavi, R., & Bicchi, A. (2005). Design and control of a variable stiffness actuator for safe and fast physical human/robot interaction. In *Proceedings of the IEEE international conference on robotics and automation*, Barcelona, Spain (pp. 526–531).
- Uemura, M., & Kawamura, S. (2009). Resonance-based motion control method for multi-joint robot through combining stiffness adaptation and iterative learning control. In *Proceedings of the IEEE international conference on robotics and automation*, Kobe, Japan (pp. 1543–1548).
- Uno, Y., Kawato, M., & Suzuki, R. (1989). Formation and control of optimal trajectories in human multi-joint arm movements: minimum torque-change model. *Biological Cybernetics*, 61, 89–101.
- Vanderborght, B., Verrelst, B., Ham, R. V., Damme, M. V., Lefeber, D., Duran, B. M. Y., & Beyl, P. (2006). Exploiting natural dynamics to reduce energy consumption by controlling the compliance of soft actuators. *International Journal of Robotics Research*, 25(4), 343–358.
- Verrelst, B., Ham, V., Vanderborght, B., Vermeulen, J., Lefeber, D., & Daerden, F. (2005). Exploiting adaptable passive behaviour to influence natural dynamics applied to legged robots. *Robotica*, 23(2), 149–158.
- Wilson, A. M., Watson, J. C., & Lichtwark, G. A. (2003). A catapult action for rapid limb protraction. *Nature*, 421, 35–36.
- Winters, J. M., & Stark, L. (1985). Analysis of fundamental human movement patterns through the use of in-depth antagonistic muscle models. *IEEE Transactions on Biomedical Engineering*, 32, 826–839.
- Wolf, S., & Hirzinger, G. (2008). A new variable stiffness design: matching requirements of the next robot generation. In *Proceedings of the IEEE international conference on robotics and automation*, Pasadena, CA, USA (pp. 1741–1746).
- Zinn, M., Khatib, O., Roth, B., & Salisbury, J. (2004). Playing it safe. *IEEE Robotics & Automation Magazine*, 11(2), 12–21.



David Braun is a postdoctoral research fellow at the Institute for Perception, Action and Behaviour within the School of Informatics at the University of Edinburgh. Prior to this he was affiliated with the Centre for Intelligent Mechatronics at Vanderbilt University. He received his PhD in mechanical engineering from Vanderbilt University (2006–2009) and an M.Sc. degree in applied mechanics from the University of Novi Sad, Serbia (2001–2003). His research interests span issues in system dynamics and control, mechatronics and robotics.



Matthew Howard is a JSPS postdoctoral research fellow at the Nakamura Lab, in the Department of Mechano-informatics at the University of Tokyo. Prior to this, he was at the Institute for Perception, Action and Behaviour, a part of the School of Informatics at Edinburgh University. He received his PhD in Machine Learning and Robotics from Edinburgh University in 2009, working in collaboration with the Honda Research Institute Europe, in Offenbach, Germany. He obtained an M.Sc. in Artificial Intelligence (2005) from Edinburgh University, and his M.Sc. in Physics from Imperial College, London in 2004. His research interests span various topics in machine learning, robotics and control, especially with regard to imitation learning, optimal control and variable impedance actuation.



Sethu Vijayakumar is the Professor of Robotics and the Director of the Institute of Perception, Action and Behaviour in the School of Informatics at the University of Edinburgh. Since 2007, he holds a fellowship of the Royal Academy of Engineering in Learning Robotics, co-sponsored by Microsoft Research Cambridge. He also holds additional appointments as an Adjunct Faculty of the University of Southern California, Los Angeles, a Research Scientist of the ATR Computational Neuroscience

Labs, Kyoto-Japan and a Visiting Research Scientist at the RIKEN Brain Science Institute, Tokyo. He has a Ph.D. ('98) in Computer Science and Engineering from the Tokyo Institute of Technology. His research interests span a broad interdisciplinary curriculum ranging from statistical machine learning, robotics, planning and optimization in autonomous systems to motor control and computational neuroscience.