



Towards real-time whole-body human dynamics estimation through probabilistic sensor fusion algorithms

A physical human–robot interaction case study

Claudia Latella¹ · Marta Lorenzini² · Maria Lazzaroni² · Francesco Romano¹ · Silvio Traversaro¹ · M. Ali Akhras¹ · Daniele Pucci¹ · Francesco Nori¹

Received: 30 December 2016 / Accepted: 3 October 2018 / Published online: 31 October 2018
© The Author(s) 2018

Abstract

Physical human–robot interaction is receiving a growing attention from the scientific community. One of the main challenges is to understand the principles governing the mutual behaviour during collaborative interactions between humans. In this context, the knowledge of human whole-body motion and forces plays a pivotal role. Current state of the art methods, however, do not allow one for reliable estimations of the human dynamics during physical human–robot interaction. This paper builds upon our former work on human dynamics estimation by proposing a probabilistic framework and an estimation tool for online monitoring of the human dynamics during human–robot collaboration tasks. The soundness of the proposed approach is verified in human–robot collaboration experiments and the results show that our probabilistic framework is able to estimate the human dynamic, thereby laying the foundation for more complex collaboration scenarios.

Keywords Real-time human dynamics estimation · Human–robot physical collaboration · Probabilistic sensor fusion algorithm

Electronic supplementary material The online version of this article (<https://doi.org/10.1007/s10514-018-9808-4>) contains supplementary material, which is available to authorized users.

✉ Claudia Latella
Claudia.Latella@iit.it

Marta Lorenzini
Marta.Lorenzini@iit.it

Maria Lazzaroni
Maria.Lazzaroni@iit.it

Francesco Romano
francesco.romano.1987@gmail.com

Silvio Traversaro
Silvio.Traversaro@iit.it

M. Ali Akhras
ali.akhras@yahoo.fr

Daniele Pucci
Daniele.Pucci@iit.it

Francesco Nori
Francesco.Nori@iit.it

1 Introduction

In the last two decades, the scientific community has shown a great interest in understanding and shaping the interaction mechanisms between humans and robots. Human–robot interaction, in fact, implies communication between two dyadic agents and, if the type of interaction is ‘physical’, the communication is represented by the set of forces mutually exchanged. Within this context, the step of quantifying these forces becomes of pivotal importance for understanding the interaction mechanisms. At the current scientific stage, classical robots are built to act for humans, but in order to adapt their functionality to the current technological demand, the new generation of robots are required to collaborate with humans. This implies that the robots will be endowed with the capability to control physical collaboration through inten-

¹ Dynamic Interaction Control, Istituto Italiano di Tecnologia, via Morego 30, 16163 Genoa, Italy

² Advanced Robotics Department, Istituto Italiano di Tecnologia, via Morego 30, 16163 Genoa, Italy

tional interaction with humans. Future robots have to be aware of the human partner kinematics and dynamics (contact and exchanged forces, internal forces, joint torques). While several well-established technologies make nowadays possible the estimation of the real-time human motion (e.g., marker-based motion capture as Vicon or marker-less systems as Xsens and Microsoft Kinect), the estimation of exchanged forces with the external environment is still an outstanding problem. Most of the studies on the physical human–robot interaction (pHRI) derive from the intrinsic mutual adaptive nature that automatically occurs when two humans are cooperating together to accomplish a common task. Based on the pioneering study on the minimum-jerk model for human manipulation tasks (Flash and Hogan 1985), a method based on the minimization of the jerk can be used as a suitable approximation for estimating the human partner motions during a human–robot cooperative manipulation in Maeda et al. (2001). Here the authors' attempt is to incorporate human characteristics into the control strategy of the robot. The weakness of this type of approach lies in the pre-determination of the task and in the role that the robot has to play in the task execution. Furthermore, the minimum-jerk model reliability decreases considerably if the human partner decides to apply non-scheduled trajectory changes during the task (Miossec and Kheddar 2009).

Another suitable route for pHRI is the imitation learning approach, justified by a wide variety of computational algorithms that have been recently developed [see an exhaustive overview in Schaal et al. (2003)]. This approach yields to the *teacher–student* concept and consists in (i) collecting movements of a human agent (i.e., the teacher in the imitation process), gathered by motion capture techniques and clustered in motion databases (Guerra-Filho and Biswas 2011; Kuehne et al. 2011; Mandery et al. 2015; Wojtusich and Stryk 2015); (ii) classifying motions as a set of primitive actions (i.e., motion primitives) for model abstraction (Kulic et al. 2008); (iii) mapping human models into a robot platform, i.e., the student in the imitation process. The re-targeting implies the development of robot control algorithms capable of learning motion models from these primitives (Amor et al. 2014; Terlemez et al. 2014; Mandery et al. 2016) in order to emulate human-like movements. In general, although in the imitation approach the task pre-determination does not represent a drawback anymore, a relevant problem (termed, correspondence problem) concerns the comparability of the models (teacher vs. student). Human motion primitives are exploitable only if their knowledge could be properly mapped into a compatible robot kinematic/dynamic model. This inevitably narrows down the potential student-role candidate to anthropomorphic robots and it requires an accurate analysis on the imitation mechanism assumptions.

Unlike the current leaning, the focus is here shifted on the key role that a proper sensing technology for human

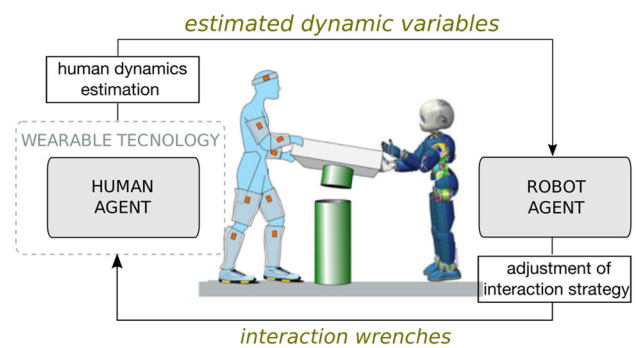


Fig. 1 A possible pHRI scenario: the human agent is provided with a wearable technology and an estimation algorithm allows to compute information about his dynamics. By properly embedding human estimations in the control loop of a human-aware robot, the intentional collaboration may be enhanced (Color figure online)

beings (body-mounted IMUs, force plates on the ground, force/torque sensors on the iCub robot) together with a dynamics estimation algorithm may offer for retrieving whole-body motion and interaction forces. The importance in retrieving this set of information can be exemplified by the scenario in Fig. 1: once the human dynamics is completely known, the human feedback may be provided to the robot controllers. As a straightforward consequence, the robot may accordingly adjust the strategy of the interaction to enhance in this way the common aim of the task. This kind of online awareness in the robot controllers will be the key to bridge the current limitation of robots in observing the human dynamics. This specific case falls within the scope of assistive collaboration (e.g., a humanoid robot is helping a human in removing a tabletop) but in general this framework can be applied in any field where a proficient collaboration between humans and robots is required (e.g., in an industrial scenario where a manipulator is helping an assembly-line worker, or in a rehabilitative field where an actuated exoskeleton is assisting an injured person).

The paper is built on the theoretical framework described in our previous work (Latella et al. 2016) from which it inherits both the notation and formulation. This paper represents the real-time evolution of Latella et al. (2016) and a first attempt at advancing the current state of the art in pHRI by designing an estimation tool for monitoring the dynamics of the human while physically interacting with a robot, in real-time.

The paper is structured as follows. Section 2 introduces the theoretical framework which the paper is based on. Section 3 presents the modelling of the human body as an articulated multi-body system. In Sect. 4 the adopted Gaussian probabilistic domain for the sensor fusion methodology is briefly recalled. Section 5 outlines the experimental setup of the framework followed by the offline algorithm validation in Sect. 6. In Sect. 7 the methodology for the real-time estima-

tion is defined. Conclusions, current limitations and several considerations on the pivotal role of further control and estimation developments are discussed in Sect. 8.

2 Background

The humanoid robots formalism is adopted throughout the paper for modelling the human body as an articulated rigid multi-body system. The advantage of this choice lies in the possibility to handle both the systems (i.e., human and robot) with the same mathematical tool. In this domain, the application of the Euler–Poincaré formalism (Marsden and Ratiu 1999) leads to three sets of equations describing: (i) the motion of the human, (ii) the motion of the robot, (iii) the linking equations characterizing the contacts between the systems.

$$\begin{aligned} (i) \quad & \mathbf{M}(\mathbf{q})\dot{\mathbf{v}} + \mathbf{C}(\mathbf{q}, \mathbf{v})\mathbf{v} + \mathbf{G}(\mathbf{q}) = \begin{bmatrix} \mathbf{0} \\ \boldsymbol{\tau} \end{bmatrix} + \mathbf{J}^\top(\mathbf{q})\mathbf{f} \\ (ii) \quad & \mathbb{M}(\bar{\mathbf{q}})\dot{\bar{\mathbf{v}}} + \mathbb{C}(\bar{\mathbf{q}}, \bar{\mathbf{v}})\bar{\mathbf{v}} + \mathbb{G}(\bar{\mathbf{q}}) = \begin{bmatrix} \mathbf{0} \\ \bar{\boldsymbol{\tau}} \end{bmatrix} + \mathbb{J}^\top(\bar{\mathbf{q}})\mathbf{f} \\ (iii) \quad & [\mathbf{J}(\mathbf{q}) \quad -\mathbb{J}(\bar{\mathbf{q}})] \begin{bmatrix} \dot{\mathbf{v}} \\ \dot{\bar{\mathbf{v}}} \end{bmatrix} + [\dot{\mathbf{J}}(\mathbf{q}) \quad -\dot{\mathbb{J}}(\bar{\mathbf{q}})] \begin{bmatrix} \mathbf{v} \\ \bar{\mathbf{v}} \end{bmatrix} = \mathbf{0} \end{aligned}$$

Equations (i) and (ii) are floating-base representations of the dynamics of the human and robot models, respectively. In general, a floating-base system is defined by its configuration $\mathbf{q} = (\mathcal{I}\mathbf{H}_B, \mathbf{q}_j) \in SE(3) \times \mathbb{R}^n$ and its velocity $\mathbf{v} = (\mathcal{I}\mathbf{v}_B, \dot{\mathbf{q}}_j) \in \mathbb{R}^{n+6}$, where n is the (internal) number of Degrees of Freedoms (DoFs) of the system, \mathcal{I} is an inertial frame and B the base frame. $\mathcal{I}\mathbf{H}_B$ is the homogeneous transformation from B to \mathcal{I} and $\mathcal{I}\mathbf{v}_B$ is the base velocity with respect to (w.r.t.) \mathcal{I} . The configuration and the velocity of the internal DoFs are denoted with \mathbf{q}_j and $\dot{\mathbf{q}}_j$, respectively. The matrices \mathbf{M} , \mathbf{C} , \mathbf{G} and \mathbb{M} , \mathbb{C} , \mathbb{G} denote the mass matrix, the Coriolis matrix and the gravity bias term for the human and the robot systems, respectively. The exchanged forces between the two systems are denoted by the vector $\mathbf{f} \in \mathbb{R}^{6k}$, where k is the number of the wrenches¹ exchanged during the interaction.² It is worth remarking that while all the variables are different for the two systems, the exchanged wrenches \mathbf{f} are mandatorily the same quantities. Matrices $\mathbf{J}(\mathbf{q})$ and $\mathbb{J}(\bar{\mathbf{q}})$ are the proper Jacobians for the human and the robot, respec-

tively. Equation (iii) is due to the rigid contacts assumption between the two systems. Since in this paper the human performs the experiments with the feet fixed on the force plates, we consider the human system to be fixed base, i.e., $\mathcal{I}\mathbf{H}_B = \text{const}$ and known a priori and $\mathcal{I}\mathbf{v}_B = \mathbf{0}$. This choice implies $\mathbf{q} = \mathbf{q}_j$ and reduces remarkably the framework complexity.

3 Human body modelling

Inspired by the biomechanical model developed for the Xsens MVN motion capture system (Roetenberg et al. 2009), we propose a human body model as a mechanical system represented by an oriented kinematic tree with $N_B = 23$ moving links and n internal DoFs. Note that $n = \sum_{i=1}^{N_B} n_i$ is the total number of DoFs of the system. The generic i th link and its parent are coupled with a i th joint by following the topological Denavit–Hartenberg convention for joint numbering (Denavit and Hartenberg 1955). The origin of each link is located at the parent joint origin (i.e., the joint that connects the link to its parent). We model the links with simple geometric shapes (parallelepiped, cylinder, sphere) whose dimensions are estimated by readings coming from the motion capture acquisition (Fig. 2).

The dynamic properties, such as center of mass (CoM) and inertia tensor for each link, are not embedded in the Xsens output data since they are usually computed in a post-processing phase. Since our aim is to go towards the direction of a real-time estimation for the human dynamic variables, the knowledge of the dynamic properties during the acquisition phase is mandatory (Drillis et al. 1964). Since it is impractical to retrieve these quantities in-vivo for humans, we relied on the available anthropometric data concerning the relation between the total body mass and the mass of each body (Winter 2009; Herman 2007), under the assumption of geometric approximations and homogeneous density for the rigid bodies (Hanavan 1964; Yeadon 1990).

4 Probabilistic sensor fusion algorithm

In this section we briefly recall the probabilistic method for estimating dynamic variables of an articulated mechanical system by exploiting the so-called sensor fusion information, already presented in our previous work [the reader should refer to Latella et al. (2016) for a more thorough description]. The following description is applied within the context of the fixed-base representation of equations in Sect. 2.

¹ With an abuse of notation, we define as wrench a quantity that is not the dual of a twist but a vector $\in \mathbb{R}^6$ containing both the forces and the related moments.

² For the sake of simplicity, we omit the forces the two systems exchange with the external environment (i.e., the ground) from the formulation of (i) and (ii). As a straightforward consequence of this, the linking equations between each system with the external environment are not considered.




Fig. 2 Human body reference model with labels for links and joints. 17 IMUs are distributed along the body, embedded in a lycra suit. Reference frames are shown using RGB (Red–Green–Blue) convention for x – y – z axes (Color figure online)

Here we are interested in computing an estimation of a vector of dynamics variables \mathbf{d} defined as follows:

$$\mathbf{d} = [d_1^\top \ d_2^\top \ \dots \ d_{N_B}^\top]^\top \in \mathbb{R}^{24N_B+2n}, \tag{1a}$$

$$d_i = [a_i^\top \ f_i^{B^\top} \ f_i^\top \ \tau_i \ f_i^{x^\top} \ \ddot{q}_{ji}]^\top \in \mathbb{R}^{24+2n_i}, \tag{1b}$$

where a_i is the i th body spatial acceleration, f_i^B is the net wrench, f_i is the internal wrench exchanged from the parent link to the i th link through the i th joint, τ_i , $\ddot{q}_{ji} \in \mathbb{R}^{n_i}$ are the torque and the acceleration at the i th joint, respectively. The system can interact with the surrounding environment, and the result of this interaction is reflected in the presence of the external wrenches f_i^x acting on the i th link.

The dynamics of the mechanical system is represented by the Newton–Euler equations³ (Featherstone 2008) whose matrix form yields to the following linear system of equations in the variable \mathbf{d} :

$$D(q_j, \dot{q}_j)\mathbf{d} + \mathbf{b}_D(q_j, \dot{q}_j) = \mathbf{0}, \tag{2}$$

³ Here we prefer to adopt the Newton–Euler formalism as an equivalent representation of the system dynamics. More details about this choice in Sect. 3.3 of Latella et al. (2016).

where D is a block matrix⁴ $\in \mathbb{R}^{(18N_B+n) \times d}$ and the bias vector $\mathbf{b}_D \in \mathbb{R}^{18N_B+n}$. We now consider the presence of N_S measurements coming from different sensors (e.g., accelerometers, force/torque sensors) and we denote with $\mathbf{y} \in \mathbb{R}^y$ the vector containing all the measurements. The dynamic variables and the values measured by the sensors can be coupled with \mathbf{d} by the following set of equations:

$$Y(q_j, \dot{q}_j)\mathbf{d} + \mathbf{b}_Y(q_j, \dot{q}_j) = \mathbf{y}, \tag{3}$$

with the block matrix $Y \in \mathbb{R}^{N_S \times d}$ and the bias vector $\mathbf{b}_Y \in \mathbb{R}^{N_S}$. By stacking together (2) and (3) we obtain a linear system of equations in the variable \mathbf{d} , such that

$$\begin{bmatrix} Y(q_j, \dot{q}_j) \\ D(q_j, \dot{q}_j) \end{bmatrix} \mathbf{d} + \begin{bmatrix} \mathbf{b}_Y(q_j, \dot{q}_j) \\ \mathbf{b}_D(q_j, \dot{q}_j) \end{bmatrix} = \begin{bmatrix} \mathbf{y} \\ \mathbf{0} \end{bmatrix}. \tag{4}$$

Equation (4) describes, for most cases, an overdetermined linear system. The bottom part [corresponding to (2)] represents the Newton–Euler equations, while the upper part contains the information coming from the, possibly noisy or redundant, sensors. It is possible to compute the whole-body dynamics estimation by solving the system in (4) for \mathbf{d} . One possible approach is to solve (4) in the least-square sense, by using a Moore–Penrose pseudo-inverse or a weighted pseudo-inverse.

In the following we perform a different choice. We frame the estimation of \mathbf{d} given the knowledge of \mathbf{y} and the prior information about the model and the sensors in a Gaussian domain by means of a *Maximum-A-Posteriori* (MAP) estimator⁵ such that

$$d_{MAP} = \arg \max_{\mathbf{d}} p(\mathbf{d}|\mathbf{y}).$$

Since in this framework the probability distributions are associated to both the measurements and the model, it suffices to compute the expected value and the covariance matrix of \mathbf{d} given \mathbf{y} , i.e.,

$$\Sigma_{d|\mathbf{y}} = \left(\bar{\Sigma}_D^{-1} + Y^\top \Sigma_y^{-1} Y \right)^{-1}, \tag{5a}$$

$$\mu_{d|\mathbf{y}} = \Sigma_{d|\mathbf{y}} \left[Y^\top \Sigma_y^{-1} (\mathbf{y} - \mathbf{b}_Y) + \bar{\Sigma}_D^{-1} \bar{\mu}_D \right], \tag{5b}$$

where $\bar{\mu}_D$ and $\bar{\Sigma}_D$ are the mean and covariance of the probability distribution $p(\mathbf{d}) \sim \mathcal{N}(\bar{\mu}_D, \bar{\Sigma}_D)$ of the model,

⁴ Note on the notation adopted for $D \in \mathbb{R}^{(18N_B+n) \times d}$. Within this notation, D is a matrix with $(18N_B + n)$ rows and d columns, i.e., number of rows of \mathbf{d} in (1a), namely $(24N_B + 2n)$. This form will be recurring throughout the paper.

⁵ The benefits of the MAP estimator choice are explained in Sect. 4 of Latella et al. (2016).

respectively; Σ_y is the covariance matrix of the measurements distribution $p(y) \sim \mathcal{N}(\mu_y, \Sigma_y)$. In the Gaussian framework, (5b) corresponds to the estimation of d_{MAP} .

5 Experimental setup

This section describes the experimental setup for the algorithm analysis. The MAP framework requires: (1) the computation of the human kinematics (possible through whole-body distributed sensor measurements), (2) the estimation of the human dynamics (i.e., the solution [5b]).

To this purpose, the experimental setup encompassed the following different sensor modules: (i) a wearable suit for the motion tracking, (ii) two standard force platforms for the ground reaction wrenches tracking, (iii) the force/torque sensors of an iCub robot arms. Ten healthy adult subjects (7 female, 3 male) have been recruited in the experimental session, height (166.6 ± 4.5 cm) and mass (61.14 ± 5.76 kg). Data were collected at Istituto Italiano di Tecnologia (IIT), Genoa, Italy. Each subject was provided of a written informative consent before starting the experiment. Kinematics data were acquired by using a full-body wearable lycra suit provided by Xsens Technologies. The wearable suit is composed of 17 wired trackers (i.e., inertial sensor units-IMUs including an accelerometer, a gyroscope and a magnetometer). The suit has signal transmitters that send measurements to the acquisition unit through a wireless receiver which collects data at a frequency of 240 Hz. The human subject performed the required task standing with the feet on two standard AMTI OR6 force platforms mounted on the ground, while interacting with the robot. Each platform acquired a wrench sample at a frequency of 1 kHz by using AMTI acquisition units.

Experiments were conducted on the iCub (Metta et al. 2010), a full-body humanoid robot (Fig. 3a) with 53 DoFs: 6 in the head, 16 in each arm, 3 in the torso and 6 in each leg. The iCub is endowed with whole-body distributed force/torque sensors, accelerometers, gyroscopes and tactile sensors. Specifically, the limbs are equipped with six force/torque sensors placed in the upper arms, in the upper legs and in the ankles (Fig. 3b). Internal joint torques and external wrenches are estimated through a whole-body estimation algorithm (Nori et al. 2015). Robot data were collected at a frequency of 100 Hz.

Each subject was asked to wear the sensing suit and to stand on the two force plates by positioning one foot per platform. The robot was located in front of the subject facing him at a known distance from the human foot location (Fig. 4). The mutual feet position was fixed for all the trials. The interaction implied that the human grasped and pushed down both the robot arms while was performing a five-repetition block of a bowing task (BT).




Fig. 3 **a** The humanoid iCub robot. **b** Model of the iCub with the force/torque sensors embedded in the limbs structure (Color figure online)

6 Algorithm validation

In this section we discuss about the validation of the MAP as a tool for estimating quantities related to the joints and links of the human model. The offline validation is mainly performed in Matlab by implementing the algorithm for the framework in Fig. 5. Data coming from the force plates and from the robot (f^x) are considered as acquired from a particular class of net external wrench sensors. Each link linear acceleration a and angular velocity ω are acquired by Xsens IMUs. Xsens data are used as input for the OpenSim (Delp et al. 2007). Inverse Kinematics (IK) toolbox that allowed to retrieve the joint angles q_j . Joint velocities \dot{q}_j and accelerations \ddot{q}_j are computed by using a weighted sum of windows of elements, with a third-order polynomial Savitzky–Golay filtering (Savitzky and Golay 1964). Also joint accelerations are considered as acquired from a class of ‘fictitious’ DoF acceleration sensors. Since the acquisition sampling rate was different among the sources, data are linearly interpolated in order to guarantee the synchronization. In general, by considering the data acquired from all above-mentioned sensors and the state (q_j, \dot{q}_j) , the MAP estimator provides the estimation of d given y . The reader should refer to the Algorithm 1 for the description of the algorithm procedure.

It is worth remarking that the estimated vector d , i.e., the solution (5b), contains variables that can be directly measured (a , f^x , \ddot{q}_j) and variables that can not be directly measured in humans (f^B , f , τ) but only estimated via algorithms.

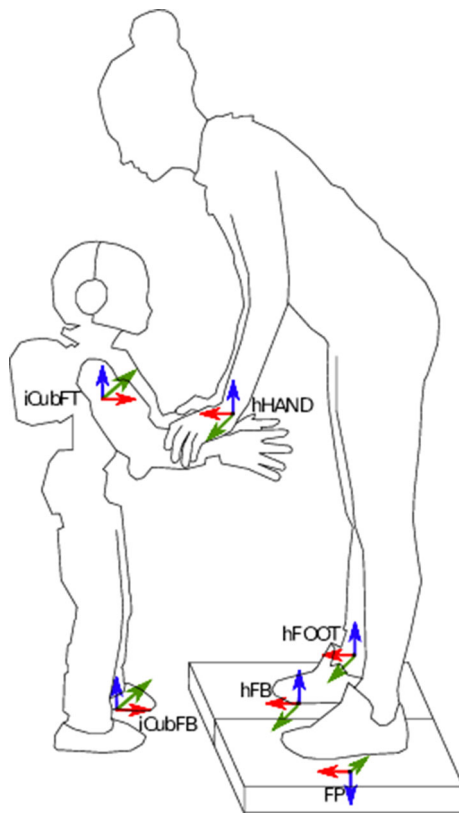


Fig. 4 Subject grasps and pushes down the robot arms. The figure shows the reference frames (in RGB convention) for the force/torque sensor of the robot (iCubFT), the robot fixed base (iCubFB), the force plate (FP), the human fixed base (hFB), the human foot and hand (hFOOT, hHAND) respectively. The mutual feet position between human and robot is fixed during the experiment (Color figure online)

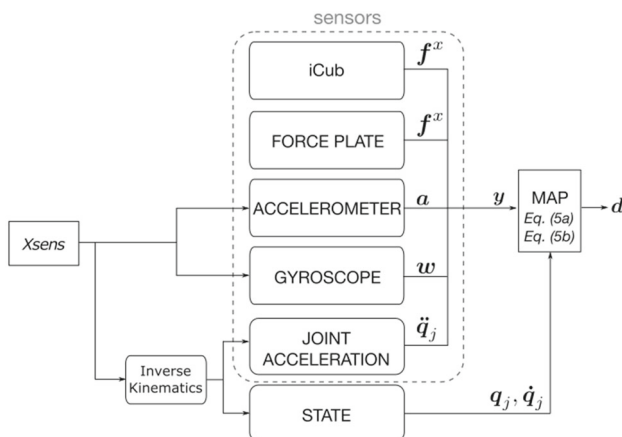


Fig. 5 Schematic pipeline of the MAP estimation framework

The MAP algorithm represents, in a sense, the probabilistic way to estimate those quantities for which a direct measure does not exist. The leading validation idea is to compare the same variables (measured and then estimated) in order to prove the quality of the proposed algorithm for certain⁶ links/joints of the model. Figure 6a shows the comparison between the linear acceleration measured by IMUs (in red) and the related MAP estimation (in blue) for the right foot, upper leg and hand, respectively. These results show that the MAP algorithm is a reliable tool for estimating the human link acceleration. We also investigated the reliability for the external forces (Fig. 6b) by considering only the links (right foot, right hand and left hand) where external forces are applied. Even in this case we obtained a suitable estimation of the quantities measured by the force/torque sensors. The same analysis was computed for the joint acceleration (in Fig. 6c). By relying on the previous results we exploited the MAP approach for estimating the joint torques in humans. Figure 6d shows the estimation of the joint torque during a *BT* for the right ankle, hip and shoulder, respectively.

Algorithm 1 Algorithm validation

Require: IMUs Xsens acquisition, human models templates

- 1: **procedure** MAP OFFLINE ESTIMATION PROCEDURE
- 2: $S \leftarrow$ number of subjects
- 3: $T \leftarrow$ number of tasks
- 4: *main loop:*
- 5: **for** $s = 1 \rightarrow S$ **do**
- 6: *nested loop:*
- 7: **for** $t = 1 \rightarrow T$ **do**
- 8: \Rightarrow parse Xsens data : $\text{suit}^s\{t\}$
- 9: **if** $t ==$ initial task **then**
- 10: \Rightarrow create model for the MAP : $\text{model}^s\{t\}.\text{urdf}$
- 11: \Rightarrow create model for the IK : $\text{model}^s\{t\}.\text{osim}$
- 12: **end**
- 13: \Rightarrow compute IK : $q^s\{t\}$
- 14: \Rightarrow compute Savitzky–Golay filtering : $\dot{q}^s\{t\}, \ddot{q}^s\{t\}$
- 15: \Rightarrow wrap available measurements in $y^s\{t\}$
- 16: \Rightarrow compute MAP : $d^s\{t\}$ **goto** Eq. (5a),(5b)
- 17: **end**
- 18: **end**

6.1 Robustness test

To test the robustness of the approach w.r.t. modelling errors, one subject was asked to perform the *BT* in two different configurations, i.e., with and without an additional mass (W) of 6 kg roughly positioned in correspondence of the torso CoM. The MAP computation was performed by considering the following cases as algorithm inputs (see Table 1):

⁶ Our approach is able to estimate d for every link/joint in the model, but for the sake of simplicity we prefer to show only those results that are the most relevant in the specific-task analysis.

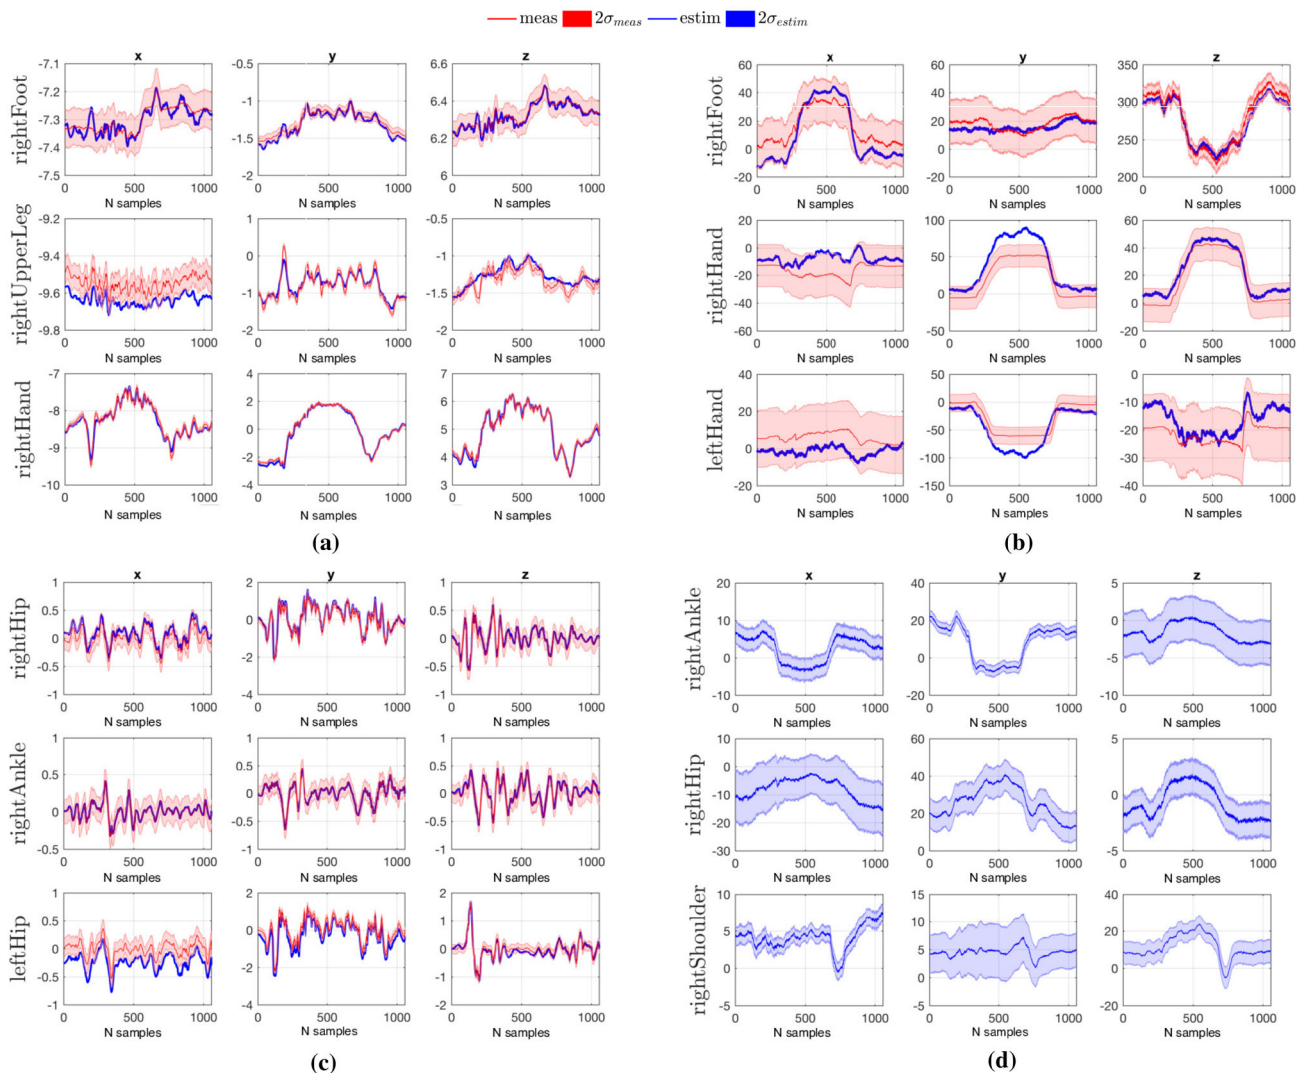


Fig. 6 Comparison between the variables measured (mean and standard deviation, in red) and their MAP estimation (mean, in blue) for (a) the linear acceleration \mathbf{a} , (b) the external force \mathbf{f}^x and (c) the joint acceleration $\dot{\mathbf{q}}_j$. The analysis shows the goodness of the MAP algorithm

in estimating those variables that can be directly measured by sensors. By relying on this result, the estimation of the joint torque $\boldsymbol{\tau}$ (mean and standard deviation) is shown in (d) (Color figure online)

Table 1 Cases for the MAP validation

| | CASE A | CASE B |
|--------------------------------|--|-----------------------------|
| Model | Without W | Without W |
| Measurements | With W | Without W |
| Σ model | 10^{-1} | 10^{-4} |
| $\boldsymbol{\tau}$ estimation | $\boldsymbol{\tau}_{(model+6\text{kg})}$ | $\boldsymbol{\tau}_{model}$ |

Both the cases are performed with the model of the subject without W . In order to highlight a lower reliability for the model used in *CASE A*, it is assigned to the model a variance⁷

⁷ We refer here to covariance associated to the model as a diagonal matrix where diagonal elements are the variance value.

equal to 10^{-1} (differently from the 10^{-4} value assigned for the *CASE B*).

By exploiting the linearity property of the system we consider the following expression for the torques:

$$\boldsymbol{\tau}_{(model+6\text{kg})} - \boldsymbol{\tau}_{model} = \boldsymbol{\tau}_{6\text{kg}}, \tag{6}$$

where $\boldsymbol{\tau}_{6\text{kg}}$ is the theoretical torque due to the additional W positioned on the torso.⁸ Given (6), it is possible to retrieve

⁸ We consider a simple 2-DoF system [see Latella et al. (2016)] in which the position of W and the hip joint angle are known.




Fig. 7 Mean and standard deviation of τ (i.e., the sum of the hips torque estimation (i.e., the sum of the torques estimated at the hips) among five repetitions of *BT* performed by a subject computed for *CASE A* (red) and *CASE B* (blue) by means of (a) the MAP algorithm and (b) the OpenSim ID toolbox. (c) Box plots of the torque estimation error ϵ_τ

computed in (7) with MAP (on the right) and with OpenSim (on the left). It shows that MAP is a method more robust to the modelling errors since it gives the possibility of weighting the reliability of the model by properly setting the related covariance matrix (Color figure online)

the error ϵ_τ on the τ estimation for the subject with W , due to the *CASE A*:

$$\epsilon_\tau = |\tau_{(model+6\text{ kg})} - \tau_{model}| - \tau_{6\text{ kg}}. \tag{7}$$

We compute (7) by using the OpenSim Inverse Dynamics (ID) toolbox as well, in order to evaluate its effectiveness w.r.t. the modelling errors. Figure 7a, b provide the mean and the standard deviation of the torque estimation by means of the MAP algorithm and the OpenSim ID, respectively. Figure 7c shows the comparison between the computation of the error in (7) for both the above methods: the error is higher in OpenSim computation than in MAP since OpenSim does not offer the possibility of setting the model reliability in the computation.

6.2 Incremental sensor fusion analysis

The novelty of the MAP framework consists in replacing the classical boundary conditions for the Newton–Euler equations with multi-sensor readings. In this section, we discuss the effects on the estimation when different measurements are progressively added into the algorithm. Let us consider all the sensor sources involved in this analysis, such that

$$y_{FP} = {}^{FP}X_{hFOOT} f_{hFOOT}, \tag{8a}$$

$$y_{IMU} = {}^{IMU}X_L a_L, \tag{8b}$$

$$y_{iCubFT} = {}^{iCubFT}X_{hHAND} f_{hHAND}. \tag{8c}$$

Equation (8a) is built for each force plate and it is used the transformation matrix⁹ X between the human foot reference frame (*hFOOT*) and the force plate frame (*FP*). Equation

⁹ See Featherstone (2008) for the definition of the transformation matrix between two reference frames.

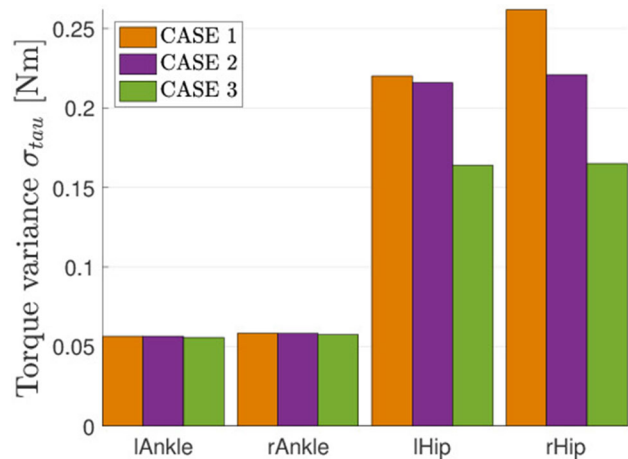


Fig. 8 10-subjects mean torque variance [Nm] of five repetitions of the *BT* by using three different combinations of the measurements equation: *CASE 1* for (9a), *CASE 2* for (9b) and *CASE 3* for (9c), for both left and right ankle, hip joints, respectively (Color figure online)

(8b) is built for each IMU and it exploits the transformation between each human link frame L on which the IMU is attached and that particular *IMU* reference frame. Equation (8c) is used for the *iCub* force/torque sensors for which it is necessary the transformation from each human hand frame (*hHAND*) to the robot sensor frame (*iCubFT*). A detailed frames view in Fig. 4. In particular, we build (3) for three different cases:

$$CASE\ 1 \quad y = [\ddot{q}_j, y_{FP}] \tag{9a}$$

$$CASE\ 2 \quad y = [\ddot{q}_j, y_{FP}, y_{IMUs}] \tag{9b}$$

$$CASE\ 3 \quad y = [\ddot{q}_j, y_{FP}, y_{IMUs}, y_{iCubFT}] \tag{9c}$$

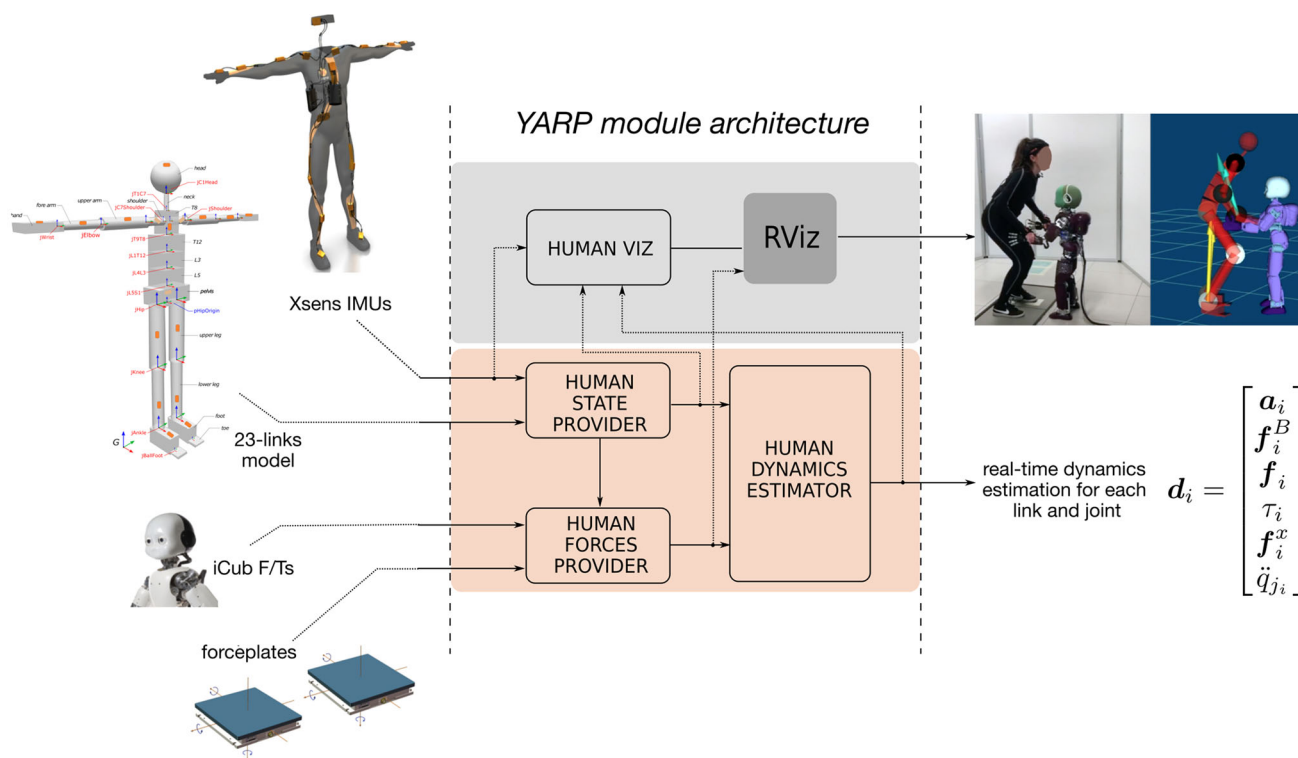


Fig. 9 The YARP architecture for estimating and visualizing real-time human dynamics during a physical interaction with a robot. The software architecture is able to estimate every 10 ms the dynamic variables

in d for each i th link/joint in the model (pink area) and to visualize the information about the kinematics and the dynamics of the human via RViz (gray area) (Color figure online)

The MAP algorithm is performed thrice by including (as set of sensors) Eqs. (9a), (9b), (9c), respectively. This analysis is to prove that, by progressively adding (i.e., increasing) sensors data, the variance associated to the estimated dynamic variables consequently decreases at each MAP computation, by making the estimation more reliable. Among the estimated variables in d we consider again the torque τ (e.g., both the right and left ankle and hip) and the variance along the axis of major relevance (i.e., y axis). By passing progressively from CASE 1 to CASE 3 the variance associated to the mean torques decreases.¹⁰ Figure 8 shows the decreasing behaviour of the mean variance of the torque at the hips and the ankles computed among ten subjects. The variance values at the ankles do not vary significantly among the three sensor configurations since the ankle torque estimation depends mostly on the contribution of the force plates that are included in all the cases of the computation. Conversely, a significant decreasing behaviour is evident in the values associated to the hips. In this case the contribution of the three sensors becomes

important since the torque estimation at the hips is affected by the complete set of sensors (i.e., contribution of the force plates weights lower, contribution of IMUs + iCub sensors weights more).

7 Towards the real-time estimation

We exploit in this section the YARP middleware (Metta et al. 2006) to perform in a real-time domain the human dynamics estimation. The YARP interface is mainly used to minimize the effort in the infrastructure-level software development by facilitating code reusability and modularity. The algorithm of Fig. 5 is hereafter arranged to fit the online requirements, Fig. 9. In this specific case, YARP allows to split the algorithm into three different modules, such that

1. The *human-state-provider* module reads the links information from a human model (built in a pre-acquisition phase) and returns the human state (q_j, \dot{q}_j) . Note that the results will depend on the complexity of the model (e.g., few-DoF model may lead to worst approximations) but the formulation does not depend on the model complexity.

¹⁰ In order to substantiate the statistical significance of the results, a paired-samples t test is performed firstly between CASE 1 and CASE 2 (2 sensors vs. 3 sensors) and then between CASE 2 and CASE 3 (3 sensors vs. all sensors). Torque variances statistically significant, p value < 0.05 .

2. The *human-forces-provider* module reads forces coming from different YARP-based sources and transforms them into quantities expressed in human reference frames.
3. The *human-dynamics-estimator* module estimates the dynamics vector \mathbf{d} every 10 ms given the output data of the previous two modules.

Furthermore, YARP allows to visualize the real-time estimation (whether a pHRI is occurring or not) by means of some structured modules and a ROS-based visualizer (gray part in Fig. 9). The *human-viz* module is composed of different sub-modules that are in charge of reading information from the Xsens system, the human state, the estimated vector \mathbf{d} , and of publishing information in the form of ROS-topic messages to be visualized in the RViz toolit (Kam et al. 2015). The advantage in adopting this tool is due to its versatility and agnosticism to the type of the data structure or algorithm. Figure 10 shows two different interaction tasks with their real-time visualizations, for a bowing (*on top* sequence) and a squat task (*on middle* sequence), respectively. In addition to providing a visual feedback to the estimation, the visualizer gives us an important information about how much the human joint ‘effort’ is. The joint torques (estimated via MAP) are represented with balls whose colour can be considered as an indicator of the effort: in a gray scale, a light ball means a high effort, a dark ball a minor effort.

8 Conclusions and future works

The paper endeavours to design an estimation tool for monitoring the real-time dynamics of a human being physically involved in a collaboration task with a robot. Indeed, it introduces a description of the steps needed *to get towards* the real-time estimation: starting from an offline validation of the algorithm (performed in Matlab, Sect. 6), we designed a (C++ based) tool for the estimation of the online human joints torques (Sect. 7). Within the real-time framework any kind of pHRI can be visualized as long as it obeys the initial assumptions (i.e. fixed base and rigid contact conditions). Even if a mathematical formalism already exists for the floating-base representation (Sect. 2), the existing software tool does not support it yet. The upgrade to a floating-base representation is mandatory for the implementation of the algorithm in more complex experimental setups where both humans and robots could move while interacting.

In this work, we applied the proposed approach to human models composed of 1-DoF revolute joints, by using the classical formalism widespread in robotics (Featherstone 2008). In particular, we combined this type of joints to obtain a series of joints with a high number of DoFs, that however are only a rough approximation of the complexity exhibited by real-life biomechanical joints. Despite we chose this joint model for




Fig. 10 Subject performing three different tasks with the related real-time RViz sequence visualization: a bowing (*on top*), a squat (*on middle*), a task where the human is helping the iCub to stand up from a rigid support (Romano et al. 2018) (*on bottom*). The visualizer shows the forces expressed in human frames: the ground reaction forces (in yellow) and the forces measured by the force/torque sensors of the robot at the human hands (in light blue). It visualizes also how much effort (in terms of joint torques) the human is doing during the pHRI by means of gray-scale balls placed at the joints: a light ball means a high effort, a dark ball a minor effort (Color figure online)

an initial exploration, the proposed algorithm is not limited to this particular choice. In particular, the properties of any joint (with an arbitrary number of DoFs) can be encapsulated in an interface where the relative position, velocity and

acceleration of the two bodies connected by it are described by arbitrary functions of joint coordinates variables and their derivatives that can then be directly inserted in (2) and (3). A possible implementation to generalize our method to arbitrarily complex musculoskeletal models is presented in Seth et al. (2010), where it is used for biomechanical simulations.

Another important investigation will concern the estimation of the human inertial parameters (mass, CoM, inertias). The use of the anthropometric tables is currently the most widespread tool for estimating such values. Even though they allow to scale the model along with the subject, this could be a rough approximation for the non-standard population samples (children, elderly, obese, individual with prostheses) and this is pushing the scientific community towards the development of new alternative ways to estimate these parameters directly from data (Venture et al. 2009; Robert et al. 2017).

The long-term objective of the entire work is represented by the possibility to extend the human dynamics estimation framework to a new framework that encompasses an *active* collaboration with a robot. In the experimental setup of Sect. 5, the robot is considered only as a *passive* forces measurer, i.e., as a tool for estimating the external forces acting on the human hands. Since the human dynamics is of pivotal importance for a control design aimed at considering the *human in the loop*, forthcoming developments will provide (online) the robot with the human force feedback. This information could be used as a tool for *reactive* human–robot collaboration (implying a robot reactive control) and, in a long-term perspective, for *predictive* collaboration. A first investigation towards this direction has been done in Romano et al. (2018). The paper attempts to answer the question “*How can we predict human intentions so as to synthesize robot controllers that are aware of and can react to the human presence?*” by considering an interaction scenario between a human and a human-aware iCub robot. In the paper, the momentum-based balancing controller of the robot has been modified [see Sect. 2 of Romano et al. (2018)] to take into account and exploit the human forces. A task for the robot stand-up from a rigid support has been performed with (bottom sequence of Fig. 10) and without the human help showing that the robot needs to provide less torque when helped by the human since it is capable to exploit the human assistance.

Acknowledgements This paper was supported by EU An.Dy Project (An.Dy has received funding from the European Union’s Horizon 2020 Research and Innovation Programme, No. 731540). This paper was also supported by the FP7 EU Projects CoDyCo (No. 600716 ICT 2011.2.1 Cognitive Systems and Robotics). The content of this publication is the sole responsibility of the authors. The European Commission or its services cannot be held responsible for any use that may be made of the information it contains.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution,

and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Amor, H.B., Neumann, G., Kamthe, S., Kroemer, O., & Peters, J. (2014). Interaction primitives for human–robot cooperation tasks. In *2014 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 2831–2837).
- Delp, S. L., Anderson, F. C., Arnold, A. S., Loan, P., Habib, A., John, C. T., et al. (2007). OpenSim: Open-source software to create and analyze dynamic simulations of movement. *IEEE Transactions on Bio-medical Engineering*, *54*(11), 1940–1950.
- Denavit, J., & Hartenberg, R. (1955). A kinematic notation for lower pair mechanism based on matrices. *Transactions of the ASME. Journal of Applied Mechanics*, *22*, 215–221.
- Drillis, R., Contini, R., & Bluestein, M. (1964). Body segment parameters; a survey of measurements techniques. *Artificial Limbs*, *8*, 44–66.
- Featherstone, R. (2008). *Rigid body dynamics algorithms*. New York: Springer.
- Flash, T., & Hogan, N. (1985). The coordination of arm movements: An experimentally confirmed mathematical model. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *5*(7), 1688–1703.
- Guerra-Filho, G., & Biswas, A. (2011). The human motion database: A cognitive and parametric sampling of human motion. *Face and Gesture*, *2011*, 103–110.
- Hanavan, E. P. (1964). A mathematical model of the human body. *AMRL-TR. Aerospace Medical Research Laboratories (U.S.)* (pp. 1–149).
- Herman, I. (2007). *Physics of the Human Body | Irving P. Herman | Springer*. Berlin: Springer.
- Kam, H. R., Lee, S.-H., Park, T., & Kim, C.-H. (2015). RViz: A toolkit for real domain data visualization. *Telecommunication Systems*, *60*(2), 337–345.
- Kuehne, H., Jhuang, H., Garrote, E., Poggio, T., & Serre, T. (2011). HMDB: A large video database for human motion recognition. In *2011 International Conference on Computer Vision* (pp. 2556–2563).
- Kulic, D., Lee, D., Ott, C., & Nakamura, Y. (2008). Incremental learning of full body motion primitives for humanoid robots. In *Humanoids 2008 - 8th IEEE-RAS International Conference on Humanoid Robots* (pp. 326–332).
- Latella, C., Kuppuswamy, N., Romano, F., Traversaro, S., & Nori, F. (2016). Whole-body human inverse dynamics with distributed micro-accelerometers, gyros and force sensing. *Sensors*, *16*(5), 727.
- Maeda, Y., Hara, T., & Arai, T. (2001). Human–robot cooperative manipulation with motion estimation. In *Proceedings 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems. Expanding the Societal Role of Robotics in the the Next Millennium (Cat. No. 01CH37180)* (Vol. 4, pp. 2240–2245).
- Mandery, C., Terlemez, Do, M., Vahrenkamp, N., & Asfour, T. (2015). The KIT whole-body human motion database. In *2015 International Conference on Advanced Robotics (ICAR)* (pp. 329–336).
- Mandery, C., Terlemez, Do, M., Vahrenkamp, N., & Asfour, T. (2016). Unifying representations and large-scale whole-body motion databases for studying human motion. *IEEE Transactions on Robotics*, *32*(4), 796–809.
- Marsden, J. E., & Ratiu, T. (1999). *Introduction to mechanics and symmetry: A basic exposition of classical mechanical systems. texts in applied mathematics* (2nd ed.). Berlin: Springer.

- Metta, G., Fitzpatrick, P., & Natale, L. (2006). YARP: Yet another robot platform. *International Journal of Advanced Robotic Systems*, 3(1), 8.
- Metta, G., Natale, L., Nori, F., Sandini, G., Vernon, D., Fadiga, L., et al. (2010). The iCub humanoid robot: An open-systems platform for research in cognitive development. *Neural Networks: The Official Journal of the International Neural Network Society*, 23(8), 1125–1134.
- Miossec, S., & Kheddar, A. (2009). Human motion in cooperative tasks: Moving object case study. In *2008 IEEE International Conference on Robotics and Biomimetics* (pp. 1509–1514).
- Nori, F., Traversaro, S., Eljaik, J., Romano, F., Del Prete, A., Pucci, D. (2015). iCub whole-body control through force regulation on rigid non-coplanar contacts. *Frontiers in Robotics and AI*, 2.
- Robert, T., Leborgne, P., Abid, M., Bonnet, V., Venture, G., & Dumas, R. (2017). Whole body segment inertia parameters estimation from movement and ground reaction forces: A feasibility study. *Computer Methods in Biomechanics and Biomedical Engineering*, 20, 175–176.
- Roetenberg, D., Luinge, H., & Slycke, P. (2009). *Xsens MVN: Full 6dof human motion tracking using miniature inertial sensors*. Enschede: Xsens Motion Technologies BV.
- Romano, F., Nava, G., Azad, M., amernik, J., Dafarra, S., Dermay, O., et al. (2018). The CoDyCo project achievements and beyond: Toward human aware whole-body controllers for physical human robot interaction. *IEEE Robotics and Automation Letters*, 3(1), 516–523.
- Savitzky, A., & Golay, M. (1964). Smoothing and differentiation of data by simplified least squares procedures—(ACS publications). *Analytical Chemistry*, 36(8), 1627–1639.
- Schaal, S., Ijspeert, A., & Billard, A. (2003). Computational approaches to motor learning by imitation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 358(1431), 537–547.
- Seth, A., Sherman, M., Eastman, P., & Delp, S. (2010). Minimal formulation of joint motion for biomechanisms. *Nonlinear Dynamics*, 62(1), 291–303.
- Terlemez, Ulbrich, S., Mandery, C., Do, M., Vahrenkamp, N., & Asfour, T. (2014). Master motor map (MMM). framework and toolkit for capturing, representing, and reproducing human motion on humanoid robots. In *2014 IEEE-RAS International Conference on Humanoid Robots* (pp 894–901).
- Venture, G., Ayusawa, K., Nakamura, Y. (2009). A numerical method for choosing motions with optimal excitation properties for identification of biped dynamics—An application to human. In *2009 IEEE International Conference on Robotics and Automation* (pp. 1226–1231).
- Winter, D. (2009). *Biomechanics and motor control of human movement* (4th ed.). Hoboken: Wiley.
- Wojtusch, J., Stryk, O. V. (2015). HuMoD—A versatile and open database for the investigation, modeling and simulation of human motion dynamics on actuation level. In *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)* (pp. 74–79).
- Yeadon, M. R. (1990). The simulation of aerial movementII. A mathematical inertia model of the human body. *Journal of Biomechanics*, 23(1), 67–74.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Claudia Latella received the B.Eng. and the M.Eng. degrees in Bioengineering at the University of Genova, Italy, in 2009 and 2011, respectively. In 2014 she started her Ph.D. in Cognitive Robotics at the Italian Institute of Technology where she received the title in 2018. Since then, she has been a post-doctoral fellow within the An.Dy project (H2020-ICT-2016-1, No. 731540). Her research interests include motion and force capture estimation for human movements, collaborative physical human–machine interaction, rehabilitation robotics. She received a Best Paper Award in 2015 for the pilot paper on the simultaneous human force and motion estimation at the 2nd International Electronic Conference on Sensors and Applications (ECSA).



Marta Lorenzini received the B.Eng. and M.Eng. degrees in Biomedical Engineering at the Politecnico di Milano, Italy, in 2014 and 2016, respectively. In 2016 she was a research fellow at the Dynamic Interaction Control Lab, Italian Institute of Technology. In 2017 she started the Ph.D. in Human–Robot Interfaces and Physical Interaction Lab, Advanced Robotics Department, Italian Institute of Technology. Marta's interests include human inverse dynamics real-time estimation, development of tools for human–robot interaction.



Maria Lazzaroni received the B.Eng. and M.Eng. degrees in Biomedical Engineering at the Politecnico di Milano, Italy, in 2014 and 2016, respectively. In 2016 she collaborated with the Dynamic Interaction Control Lab as research fellow, Italian Institute of Technology. She is currently a Ph.D. candidate at Advanced Robotics Department, Italian Institute of Technology, where she is dealing with the development of robotic exoskeletons to enhance factory workers and reduce the incidence of occupational injuries, in partnership with the Italian Workers' Compensation Authority (INAIL).



Francesco Romano received his B.Eng. and M.Eng. degrees in Computer Engineering (with specialisation in Robotics and Automation) with highest honours from University of Genova in 2008 and 2011, respectively. In April 2016 he received a Ph.D. title in Robotics at the Italian Institute of Technology under the supervision of Francesco Nori with a thesis on whole-body control methods for humanoid robots. His research interests include nonlinear control and optimal control applied to

humanoid robotics and whole body motion.



Silvio Traversaro received the B.Eng. and M.Eng. degrees in Computer Science and Robotics Engineering with highest honors from the University of Genova, in 2011 and 2013, respectively. In 2017 he received the Ph.D. in Robotics from the University of Genova working at the Italian Institute of Technology on dynamics modelling, estimation and identification applied to the iCub robot. Since then, he has been a post-doctoral fellow at the Italian Institute of Technology,

working on industrial application of robotics.



M. Ali Akhras received his B.Eng. in Automatic Control and Industrial Electronics (2008) and M1 in Control and Automation Engineering from University of Aleppo (2009); international MSc in Bioinformatics from University of Bologna (2011), specializing master in Automatica and Control Technologies from Politecnico di Torino (2012), and Ph.D. in Robotics, Cognition and Interaction Technologies from the University of Genova under a joint program with the Italian Institute

of Technology (2017). His research interests focus on combining engineering and biology knowledge with the aim to explore biological principles of biomechanical systems and neural control strategies. He is currently involved in the R&D of industrial applications at an Italian global high-tech company, LEONARDO, which is one of the key players in aerospace, defense and security.



Daniele Pucci received the B.Eng. and M.Eng. degrees in Control Engineering with highest honors from 'Sapienza', University of Rome, in 2007 and 2009, respectively. In 2013 he received the Ph.D. title in Information and Communication Technologies from University of Nice Sophia Antipolis, France, and in Control Engineering from Sapienza University of Rome, Italy. Since then, Daniele has been a post-doctoral fellow at the Italian Institute of Technology, Dynamic Interaction

Control laboratory. Within the CoDyCo project (FP7-ICT-2011.2.1, No. 600716), he has been the principal scientific contributor and he has developed innovative control techniques for whole-body motion control with tactile and force/torque sensing. Daniele's research interests include control of nonlinear systems and its applications to aerial vehicles and robotics.



Francesco Nori received his D.Eng. degree from the University of Padova (Italy) in 2002. In 2002 he was a member of the UCLA Vision Lab as a visiting student, University of California, Los Angeles, where he started a research activity in the field of computational vision and human motion tracking. He received his Ph.D. in Control and Dynamical Systems from the University of Padova in 2005. He moved to the University of Genova in 2006 and started his Post.Doc. at the lab for

integrated advanced robotics (LiraLab). In 2007 he has moved to the Italian Institute of Technology where in 2015 he was appointed Tenure Track Researcher of the Dynamic and Interaction Control research line. He was one of the key researchers involved in the iCub development, with specific focus on control and whole-body force regulation exploiting tactile information. Francesco is currently coordinating the An.Dy project (H2020-ICT-2016-1, No. 731540). In the past he has been involved in two FP7-EU projects: CoDyCo as coordinator and Koroibot as principal investigator. In 2017 Francesco joined Deepmind where his interests seamlessly span robotics and reinforcement learning, with applications in both manipulation and locomotion.