



Whole-body multi-contact motion in humans and humanoids: Advances of the CoDyCo European project

Vincent Padois, Serena Ivaldi, Jan Babič, Michael Mistry, Jan Peters,
Francesco Nori

► To cite this version:

Vincent Padois, Serena Ivaldi, Jan Babič, Michael Mistry, Jan Peters, et al.. Whole-body multi-contact motion in humans and humanoids: Advances of the CoDyCo European project. *Robotics and Autonomous Systems*, 2016, 10.1016/j.robot.2016.08.017 . hal-01399360

HAL Id: hal-01399360

<https://hal.science/hal-01399360v1>

Submitted on 18 Nov 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Whole-body multi-contact motion in Humans and Humanoids: Advances of the CoDyCo European project

Vincent Padois

*Sorbonne Universités, UPMC Univ Paris 06, CNRS, Institut des Systèmes Intelligents et de Robotique (ISIR), 4 place Jussieu, 75252
PARIS cedex 05, France. Email: vincent.padois@upmc.fr*

Serena Ivaldi

*Inria, Villers-lès-Nancy, France.
CNRS & Université de Lorraine, UMR 7503, Loria, Vandoeuvre-lès-Nancy, France. Email: serena.ivaldi@inria.fr*

Jan Babič

Jožef Stefan Institute, Ljubljana, Slovenia. Email: jan.babic@ijs.si

Michael Mistry

University of Birmingham, Birmingham, United Kingdom. Email: m.n.mistry@bham.ac.uk

Jan Peters

Max Planck Institute for Intelligent Systems and TU Darmstadt, Darmstadt, Germany. Email: mail@jan-peters.net

Francesco Nori

Robotics, Brain and Cognitive Science Department, Istituto Italiano di Tecnologia, Genova, Italy. Email: francesco.nori@iit.it

Abstract

Traditional industrial applications involve robots with limited mobility. Consequently, interaction (*e.g.* manipulation) was treated separately from whole-body posture (*e.g.* balancing), assuming the robot firmly connected to the ground. Foreseen applications involve robots with augmented autonomy and physical mobility. Within this novel context, physical interaction influences stability and balance. To allow robots to surpass barriers between interaction and posture control, forthcoming robotic research needs to investigate the principles governing whole-body motion and coordination with contact dynamics. There is a need to investigate the principles of motion and coordination of physical interaction, including the aspects related to unpredictability. Recent developments in compliant actuation and touch sensing allow safe and robust physical interaction from unexpected contact including humans. The next advancement for cognitive robots, however, is the ability not only to cope with unpredictable contact, but also to exploit predictable contact in ways that will assist in goal achievement. Last but not least, theoretical results needs to be validated in real-world scenarios with humanoid robots engaged in whole-body goal-directed tasks. Robots should be capable of exploiting rigid supportive contacts, learning to compensate for compliant contacts, and utilising assistive physical interaction from humans. The work presented in this paper presents state-of-the-art in these domains as well as some recent advances made within the framework of the CoDyCo European project.

Keywords: whole-body, control, free-floating, interaction, contacts, compliance.

1. Introduction

For cognitive agents, such as humanoid robots, to persist and act in natural human environments, contact and physical interaction become necessary and unavoidable. Everyday tasks involve making and breaking contact, among all areas of the body, whether the contacts are accidental disturbances or intentional support for dynamic movement. Critically, robots should be robust enough to cope

with unpredictable contact, via safe control mechanisms and compliance. Moreover, cognitive goal directed robots need the ability to exploit predictable contact, to aid in goal achievement, as well as learn dynamics of contact in order to generalise to novel tasks and domains.

Physical interaction has been studied in robotics, extensively under the umbrella of manipulation. For historical reasons, these studies have assumed a fixed-base as cur-

rent industrial applications do not necessitate extended mobility. Foreseen robotic applications will demand an increasing level of autonomy, including physical mobility. These applications call for extending studies on interaction to cases where the robot has a mobile-base. Remarkably and differently from the fixed-base case, interaction in these situations may compromise system balance, and goal directed actions require proper whole-body coordination and use of contact. However, the principles governing whole-body coordination in humans are far from being understood and implementations on complex systems, such as humanoids, are missing, especially besides walking.

Within this context one of the major challenges of robotic research is to advance the current control and cognitive understanding about robust, goal-directed whole-body motion execution with multiple contacts. Remarkably, focus should be posed on complex systems, such as humans and humanoids. In a crescendo of complexity, as illustrated in Figure 1, the current state-of-the-art in these domains (state-of-art 1 and 2) should be advanced to address more complex scenarios (challenges 1 and 2).

State-of-art 1: balancing with multiple rigid contacts. The robot is standing and balancing with its hands supported by a rigid table in front of its body. However, the table is unstable, and unexpectedly the contact with the table breaks. A contact state change is sensed, and the robot’s control architecture automatically adjusts posture control parameters to maintain balance in light of the reduced support. The unexpected breaking of contact makes it more challenging.

State-of-art 2: goal directed actions involving contacts. The robot is standing with its hands at its side, and intends to reach for an object on a table in front. The robot recognises that the distance is sufficiently far away, and the task cannot be achieved without compromising balance. The robot decides to initiate a new contact with its left hand on the table, providing sufficient support for reaching the object with its right hand.

Challenge 1: learning non-rigid contacts. The robot sits down on a chair with a soft cushion, however the cushion has a particular stiffness quality not experienced before. The robot tries to reach for an object on a table, but it fails as it did not adequately compensate for the unexpected dynamics of the soft cushion. After a few attempts, the robot adapts its model of the contact interaction, and is able to infer new control action to successfully reach the goal.

Challenge 2: human assistive contacts. The robot is seated in a chair, and a person comes to assist the robot to stand. He/she grabs both hands of the robot and starts pulling upwards. The robot senses the new contact, and recognising from the interaction force that it is an external agent, allows its arms to be compliant. When the force becomes sufficient to enable standing, the robot recognises




Figure 1: The four main scenarios involving whole-body motion with multiple contacts, addressed in the CoDyCo project.

the intended action and stiffens its arms while pushing its legs to rise from the chair. Finally once standing, but still in contact with the human, the robot returns compliance to its arms to allow for safe interaction while retaining overall control of its posture.

1.1. Compliance in whole-body motion

Present day robots are still far from the human capabilities in exploiting predictable events and in coping with uncertainty. The gap between humans and robots is particularly apparent when in tasks involving unstructured physical interaction with the environment or other agents. Recent behavioural experiments yielded a new perspective on modelling the way humans deal with both predictable and unpredictable motor control tasks. In early experiments, it has been shown [90] that humans learn and adapt internal dynamical models of their own arm in interaction with the environment. Such internal models appear to be crucial in predicting how muscle activations produce hand movements and therefore may play an essential predictive role in movement planning. However, Burdet et al. [10] have shown that when prediction is not a viable strategy, humans can rely on arm compliance regulation (by means of muscle co-activation) to cope with the unpredictability that naturally arises from feedback delays when performing arm-reaching movements in unstable environments. Basic research and robotics technology are ready to extend such insights from single limb movements to whole-body interaction and the validation of these models appears feasible. In contrast to manipulation scenarios with static base robot systems, dynamic whole-body interaction concerns the analysis of phenomena at a higher scale (bigger interaction forces, bigger muscle activations, etc.). Whole-body compliance regulation with force/impedance control is not only favoured by current theoretical progress and available technologies, but may actually be ready for widespread use instead of being limited to just a few prototypes.

1.2. Roadmap beyond state-of-the-art

With reference to Figure 2 and following, we propose a classification that relies on the well known concept of compliance (or the inverse concept of stiffness and more




Figure 2: Classification of whole-body tasks based on external-compliance. The complexity increases from top to bottom, *i.e.*, with the need of exploiting the compliance of the contacts.

generally impedance), to be understood as the force-displacement characteristic of a contact. Interaction scenarios can be classified by quantitatively measuring two essential components of contacts: external and internal compliance (internal here refers to the agent or “the self”). The first scenarios classification (Figure 2) is based on the external-compliance; it includes scenarios that involve non-compliant (rigid) external contacts and scenarios with compliant external contacts. This second category is extremely wide in consideration of the multitude of possible compliant behaviours that can be experienced: from the linear force-displacement characteristic of a linear spring to the complex non-linear characteristic of a pillow. Scenarios within this category practically overlap with the first category but rigid contacts are replaced by non-rigid contacts. In these two categories the agent (or “the self”, represented with a human silhouette) is always interacting with inanimate objects (the external contacts: a chair, a sofa, the floor, etc.). In the last category, “the self” and “the other” are both humans. In these scenarios the external-compliance is not a well-defined relationship between force and displacement but depends on the active intention of “the other”.

External-compliance is only one side of the interaction, and the agent has limited control over it. The other side of the interaction is what we call the “self” (internal) compliance, which is instead fully under control of the cognitive agent. Self-compliance needs to be adapted to the environment compliance and the ability to actively regulate the internal compliance has been only recently implemented on multi-degrees-of-freedom robots. The self-compliance regulation represents the pro-active and cognitive component of the interaction and therefore gives the robot an enhanced degree of autonomy to be exploited in handling situations not anticipated at design time. In this sense, the self-compliance level and actuation range can be used




Figure 3: Classification of whole-body tasks according to an increasing self-compliance level and actuation range.

to classify different scenarios as shown by Figure 3. At the very first level of this classification we consider scenarios that do not require significant self-compliance regulation as they typically involve dynamically stable situations. Such situations involve for example dynamically stable tasks, which substantially require direct control of stable postures. The second level of the classification includes tasks that require a certain level of active compliance either to stabilise unstable systems (*e.g.* balancing) or to compensate for unpredictable interaction characteristics (*e.g.* standing hand in hand with another agent). Finally at the highest level of this classification we consider highly complex tasks characterised by strong requirements in terms of “self”-compliance planning and regulation.

External and self-compliance are two fundamental aspects of any interaction. It is therefore crucial to understand how these two concepts become intertwined once contacts are established. The concept of contact-compliance indeed corresponds to the overall compliance obtained once the external and the self-compliance become coupled with the contact establishment. A contact can be seen as the serial connection of two compliances, one representing the external-compliance, the other representing the self-compliance. The compliance of a serial interconnection is simply the linear sum of the individual compliances. Roughly speaking, the contact-compliance does not significantly change when the external and self-compliance are changed simultaneously by an equal and opposite quantity. No advancement can be associated to situations which correspond to augmenting the self-compliance at the cost of diminishing the external-compliance or vice versa, as in these situations the overall contact-compliance does not change. This fundamental procedural principle is well sketched in Figure 4. The horizontal axis sorts possible

scenarios according to a progressively increasing external-compliance level. The vertical axis instead orders the same scenarios by means of increasing self-compliance levels and actuation ranges: tasks involving minimal self-compliance regulation or low levels of compliance are shown at the bottom; tasks involving wide self-compliance regulation ranges including high compliance levels are at the top. The grey-colour-valued function shown in the space defined by these two axes is a qualitative evaluation of progress beyond the state-of-the-art: dark grey is the state-of-the-art, increasing levels of blue represent step-by-step progress beyond state-of-the-art. Progress in handling whole body contacts can be achieved only by simultaneously increasing the external and the self-compliance levels. Conversely, little advances are achieved when increasing the environmental compliance but reducing the active compliance component. Vice versa, a dual way to achieve little progress beyond the state-of-the-art corresponds to scenarios that involve a strong self-compliance regulation but reduced external-compliance.




Figure 4: The metric space to evaluate the progress work beyond the current state-of-the-art. Interaction is the inter-twined combination of two components, external and self-compliance, both contributing to the concept of contact-compliance. Whole-body scenarios should be evaluated in a metric space that takes into account how self and external-compliance contribute to contact-compliance. Contact-compliance is the sum of self and external-compliance. Remarkably the major advances can be obtained by simultaneously advancing the external and the self-compliance requirements. The vertical axis represents both self-compliance levels and actuation ranges in consideration of the fact we are mainly interested in self-compliance regulation, actuation and control. The four proposed scenarios have increasing complexity with respect to current state-of-the-art.

1.3. Scope of the paper

Looking at whole-body multi-contact motion in Humans and Humanoids through the prism of external and

self compliance, and more generally contact compliance, provides an original and insightful light to the addressed question. It is a good starting point to define the problems to be addressed in this domain. These problems can be divided in four closely intertwined domains: control theory, machine learning, human behavioural experiments and software development. Each of them serve the objectives of the CoDyCo project:

1. develop a general software toolkit for whole-body dynamics computation with multiple external contacts;
2. conduct human behavioural studies for understanding human use of external contact with environments exhibiting different external compliances¹, including interpersonal cooperative contacts in natural whole body tasks;
3. develop a control architecture for whole-body coordination and regulation of whole body compliance;
4. leverage machine learning methods for acquiring models of compliant contact with the environment and provide the humanoid robot control architecture with the core abilities for the adaptation, generalization and self-improvement of both control laws and tasks related to these types of physical interaction.

In these domains, numerous technological and scientific developments have been led over the last decades. A brief overview of these related works is provided in the next section. The strength of the CoDyCo approach with respect to some of these existing works is to gather expertises in different fields around a very challenging scientific question. Interesting synergies between these different domains are much more likely to emerge in these conditions. The following sections are directly related to three scientific domains mentioned here-before and describe the work led within the framework of the CoDyCo project in human postural control and whole-body motion in contact with the environment, whole-body controllers and learning respectively. Future works are described in the conclusion which also summarizes the major contributions of the project.

2. State-of-the-art

In this section, a brief overview of some of the key aspects in the state-of-the-art related to whole-body motion in contact for humans and humanoids are presented.

2.1. Technological state-of-the-art

Among the recent achievements in the field of robotics, there are two major technological prerequisites that will play a fundamental role in enhancing whole-body motion capabilities: distributed force and touch sensing. Both technologies have been only recently integrated and used

¹and more generally *impedances*.

in (humanoid) robots, including the iCub [56]. Force control is a fundamental property for any autonomous agent in interaction with the environment. First attempts to regulate interaction forces relied on active force and compliance control schemes, typically coupled with custom mechanical designs such as the ones proposed in [83] and [29], which were eventually implemented on successful commercial manipulators. Similar solutions have been eventually implemented on some humanoid platforms [12] [21], including the iCub [27]. Recent theoretical and technological advances have revealed the importance of intentionally introducing mechanical compliance in the design [78] and (even more recently) the necessity of actively regulating the actuator passive compliance [44] [94] [57]. It is to be expected that within the next years robots such as iCub will be equipped with variably compliant actuation technologies at some (if not all) of the main joints [96] [95].

Touch is another fundamental sensing capability for autonomous agents willing to interact with an unstructured environment or humans [6]. Whole-body distributed touch sensing has been only recently embedded on humanoid robots, but there already exist quite a few examples: Robovie-IV [59], RI-MAN [62], Macra [29] and Meka [38], just to cite a few. The iCub already integrates a mature technology [87] covering the upper body, legs and feet soles. Finally, several open-source software libraries have been developed in the last years to support research in whole-body dynamics and contact simulation. Several dynamics simulators have been developed for robotics (see [37] for a survey). The most interesting physics engines for our purposes are the ones with Featherstone-like forward dynamics calculation [23] [99], built-in collision detection and stable numerical contact forces computations [64]. Among the kinematic and dynamic libraries it is worth citing HuMANs, a toolbox for analysis and control of both human and humanoid motion, and iDyn, a generic software library for computation of whole-body dynamics and external contact forces of complex manipulators and humanoids [36]. iDyn has been extensively used in iCub to compute whole-body dynamics as it enables to reinforce these computations with measurements coming from inertial, force and tactile sensors embedded in the robot [27].

2.2. Human motor control state-of-the-art

Human whole-body motion control has been studied within tasks such as reaching on a supporting surface and sit-to-stand. These movements involve coordination of multiple joints, significant shift of the centre of mass, and control of equilibrium, either in static or dynamic conditions. These are skills learnt early in human childhood but also studied extensively in the context of motor disability, *e.g.* after neurological insults like stroke, or in the elderly with reduced muscle power, joint flexibility and sensory loss. However, almost nothing is yet known about when healthy subjects choose to make use of contacts with support surfaces. It has been shown that in standing posture, this contact provides augmented sensory informa-

tion reducing sway [40], and how in some circumstances, non-weight bearing but informative “light-touch” between two standing subjects can cause coupling that leads to increased sway, emphasising that knowledge about the stability and compliance of the contact surface is vital.

Reach using supports. Human reaching with arm support has not been extensively studied. There is almost no literature on the issue of how humans use one hand to extend their reach space. For example, to lean forwards requires a shift of the trunk and a shift of the centre of mass [30]. At some point it becomes advantageous to use a supporting surface, allowing a reduction in anticipatory postural adjustments and a simplified control strategy [92]. But the decisions about when to implement support using one arm, which will depend on the availability, reliability and compliance of a support surface are almost unstudied [45].

Sit-to-stand. The postural adjustments that contribute to a sit-to-stand action are well documented. The action requires a shift of centre of mass, development of momentum, and precisely timed hip and knee extension, combining with maintaining stability with ankle control. As motor ability lessens, *e.g.* in the elderly, compensatory foot placement with increased momentum generation using hip flexion and arm movement is often employed [39]. Support from the chair arm or from a cane [49] increases stability in the forward axis. Again, decisions about when the support surface would be used, depending on its stability and compliance, are unstudied. The effects of unstable foot support in the sit-to-stand action are studied in [30] the authors suggests a clear trade-off between support surface stability and manoeuvrability, and argue that adapting to the added uncertainty could help individuals become more manoeuvrable. Finally, there is little work on how the sit-to-stand action changes with elastic support - this has been studied in locomotion and jumping [5], but not in inter-actions with support surfaces.

Dimensionality reduction. Complex multi-joint movements call for control strategies that simplify and reduce degrees of freedom. There are various competing theories of how this can be achieved [47]. Perhaps most relevant is the uncontrolled manifold hypothesis [88] that demonstrates that it is highly effective to allow some parameters to be uncontrolled, if task irrelevant, and to control only a fewer task relevant parameters. In [88] Scholz & Schoner applied this to the sit-to-stand task, and show that the centre of mass in the forward axis is well controlled, head and hand position are less controlled, and vertical head position appears little controlled. How these behaviours change with support is an open question. Equally important is the issue of how high dimensionality whole body motion of human models can be reduced to extract principles of action applicable to robots with different geometries. These are implemented by muscle and joint synergies that reduce the functional degrees of freedom during

a given action. There has been very effective use of principal or independent components analyses to capture such human whole body movement and reduce dimensionality (*e.g.* as in [25]). Recent developments include extracting functional components, which treat joint-kinematics data as functions instead of as a series of independent samples, and are comparable across groups of subjects [13].

2.3. Robot control state-of-the-art

In complex scenarios, when the robot and the environment are assumed to be perfectly known, planning approaches explore the possible states of the robot (*e.g.* configurations of the robot in its environment) in probabilistic graph-like manners [48] to determine the sequence of commands to provide to the robot to perform a certain action in free space [15] [46] or in complex contact situations [8] [9]. Such methods are usually computationally demanding and difficult to apply online. Conversely, when the global goal of the robot is relatively simple, the high-level planner can be almost disregarded because the goal to be achieved can “easily” be described *a priori* in terms of operational tasks [41] to be activated and combined. This falls into “the simultaneous management of multiple operational objectives”, a well-known problem in model-based reactive control. The most popular method to deal with a set of objectives is a hierarchical framework, where operational tasks are typically prioritized in a “stack” [55], which found several applications to humanoids [89] [58]. QP (Quadratic Program) solvers have recently gained popularity in humanoid robotics as they do not require the explicit inversion of any model of the system [2] [14] [21] [82]. This corpus of reactive methods mostly succeeds in overcoming the “complexity and uncertainty” factor thanks to the use of feedback. However the proposed solutions are only locally optimal and the overall decision-making process cannot be addressed in the most general cases (*i.e.* without scripted scenarios). There is obviously a need for approaches where planning and reactive control are combined in a strongly intertwined way. This is not a simple problem: there are very few works where such a combination has actually been tested in a non ad hoc manner. The work of [3] contributed to describe the necessary control architecture but did not propose any general control solution for such a combination to exist in practice. More recently [77] introduced an architecture combining a whole body control level and a reactive symbolic planning, while [100] focused on dedicated mission-level planning methods for humanoids, coupled to task-level controllers. [82] have also proposed an architecture where sequences of operational tasks are generated on the fly based on a fuzzy-logic, rule-based decision engine. This approach, even though efficient in various specific applications, fails to scale-up as the number of required rules explodes with the growing complexity of the considered scenarios.

2.4. Learning

Real-world environments are often hard to capture perfectly with physical models. The uncertainty in model

predictions is important during controlled physical contacts between a (humanoid) robot system and its environment. Large errors either in the environmental model or in the task will lead to drastic failures and therefore need to be limited as much as possible by model adaptation. Human-inhabited worlds will never allow perfect modelling and instead require that the system generalises the tasks in such a manner that they work in a wide variety of different uncertain scenarios where there is contact between robots and either humans or physical objects. Machine learning approaches are therefore needed. Particularly in whole-body motion they are necessary for the successful implementation of the control architecture, and its implementation and application to the real-worlds scenarios. However, off-the-shelf machine learning methods are concerned with static data sets and require massive amount of computations, often rendering real-time learning in-feasible. To date, a variety of robot learning approaches have been suggested. The most important being model learning, operational space control learning, learning of elementary tasks and hierarchical combinations of tasks, which are briefly evaluated hereinafter.

Model learning. High model accuracy and constant model adaptation may be key for low torque interaction during contact. Models of the robot dynamics have been learned by real-time regression, *e.g.*, locally weighted projection regression [86] and local Gaussian process regression [65]. Learning force/torque models in iCub has been investigated with different regression techniques, such as SVM and Neural Networks [26]. Nevertheless, if any of these approaches would be given the data from a robot in contact with the environment, it would fit the model to this particular case, as the contact forces would just be treated as an additional non-linearity. As a result, the model will not generalise to new contact models and instead it would be necessary to learn a new model for each type of contact.

Operational space control learning. Control in operational space has been approached both as a direct policy learning problem [76] as well as an indirect learning problem via forward models [80]. Visuo-motor models from scratch via iterative and incremental learning have been computed, for iCub [20], exploiting its active compliance [36] to deal with self-body collisions and contacts with the environment. Here, the problem may be even more drastic as changing the contact formulation will alter the problem in its essence. As a result, an operational space control law may not transfer at all but rather become highly problematic under new circumstances.

Learning of elementary tasks and hierarchical combinations of tasks. While learning of contact-free elementary tasks by the combination of imitation learning and policy search [1] [43] is a well-explored topic, no general approaches to date can tackle the exact same problem and allow for different contact combination. Furthermore, learning of hierarchical elementary task combinations is still in

its infancy. Several interesting approaches have been suggested [93] [63] [61] [16] in literature, relying on substantially different insights. Further exploration in this area is clearly needed, especially in unexplored multi-contact scenarios. While all of these frameworks are well motivated in their domains, they have two major shortcomings from the viewpoint of whole-body motion control: they do not explicitly incorporate contact, and they do not leverage on the analytical robotics and control knowledge surrounding them.

3. Human postural control and whole body motion in contact with the environment

The aim of the work performed on human postural control and whole body motion in contact with environment is to provide a solid multidisciplinary base for future research work. We made a thorough review and summary of the recent relevant literature on human postural control and whole body motion in contact with environment². The review examines postural control strategies without and with additional support contacts, types of perturbations that are commonly used to study neuromuscular functions involved in postural control and reviews the methods for stability evaluation of bipedal systems. The review is concluded with examination of stability metrics that can be applied for non-planar contacts. Based on this review an experimental protocol has been designed to explore human strategies used when non-coplanar assistive contacts are made.

3.1. Design of models for human whole body motion in contact

Some work has been performed on understanding how to derive simplified models of whole-body balance that will encapsulate the task relevant parameters of posture control with multiple contacts.

3.1.1. Postural stability with multiple contacts

By emulating situations when balance of an individual is challenged, we examined functional role of supportive hand contact at different locations where balance of an individual was perturbed by translational perturbations of the support surface. The experimental methods are depicted on the left side of Figure 5. We found that an additional supportive hand contact significantly reduced the maximal displacement of the subject's centre of pressure (CoP) regardless of the position of the handle and the type of the perturbation. On the other hand, the position of the handle had no effects on the maximal CoP displacement (top right diagram on Figure 5) which is against the previous belief that the quality of postural control depend on the location of the hand contact [84] and supports the idea that maintaining postural stability is the task of the




Figure 5: Examining the functional role of a supportive hand contact. The subjects were standing on a force plate mounted on top of the Stewart platform that generated translational perturbations. The subjects were holding the handle with a built-in force sensor in four different positions. Major results of the study are shown on the two diagrams on the right side. Adapted from [7].

highest priority and that the central nervous system does whatever necessary to keep the body balanced [98]. Specifically, subjects always generated the required hand force, no matter where the location of the handle was, to keep the body balanced to the same extent. To get a better understanding of the functional role of supportive hand contacts, we examined the handle forces exerted by the subjects during the perturbation. In contrast with the effects on CoP, we found significant effects of perturbation direction, perturbation intensity and handle position on the maximal force in the handle (bottom right diagram on Figure 5). A detailed description of these results can be found in [7]. A 3D dynamic model of a human holding to a stable object during continuous perturbations of stance was also created using OpenSim [19]. The model was devised from measurements on 13 male subjects. Using the kinematical data recorded with frequency of 100Hz, forces that the subjects exerted on the ground, and the forces in the handle, we performed an inverse dynamic procedure and obtained joint torques produced by muscles during the experiment. An illustration of the model is shown in the left panel in Figure 6. Using this modelling approach we are now able to efficiently study the biomechanics of humans in contact with the environment.

3.1.2. Metric for postural stability with multiple contacts

A work on defining a suitable metric to measure the effects of the environmental contacts on the robot's stability was also undertaken. It used the basic concept of end-effector manipulability (for manipulators) in the literature and introduced a new tool to analyse the ability of balance for legged robots which we called manipulability of the center of mass. This tool relates the actuated joint velocities (all of them, *e.g.* the ankle, the knee, the hip, the shoulder and the elbow ones for the planar humanoid robot described on the right side in Figure 6) to

²Cf. [CoDyCo deliverable 2.1](#)




Figure 6: Left: three dimensional model of a human subject holding a handle. Right: center of mass velocity ellipses for a planar humanoid robot.

the linear velocity of the center of mass. It defines three different types of ellipsoids which are called 1) velocity ellipsoid, 2) instantaneous velocity ellipsoid and 3) instantaneous velocity ellipsoid due to the unit impulse. The first one shows the velocity of the CoM in different directions due to the unit norm of the joint velocities. The second and third types of the ellipsoids, which are obtained by using impulsive dynamics, show instantaneous changes of the CoM velocity due to the unit norm of instantaneous changes at the joint velocities and the unit norm of impulse at the actuated joints, respectively. By involving the motion equations into the calculations for the second and third types of ellipsoids (via impulsive dynamics), these ellipsoids allow us to study the effect of under-actuation as well as kinematic constraints on the robot's stability.

As an example, right panel in Figure 6 shows a planar humanoid robot (with its hand is fixed) and instantaneous velocity ellipses for the robot in the specified configuration. Since the robot is fully actuated, the first and second types of ellipses (type 2 and type 3) are the same. This ellipse (type 2) shows how the velocity of the CoM changes when the instantaneous change of the joint velocities due to the impulse has the unit norm. This shows the ability to move the CoM in different directions by a certain amount of movements at the actuated joints. The ellipse type 3 shows how the velocity of the CoM changes due to the unit impulse at the joints. In other words, it shows how a certain amount of impulse at the actuated joints can accelerate the CoM in different directions. All of the ellipses are independent from the controller and they are dependent only on the physical parameters of the robot and its kinematic constraints.

In balancing in a plane, the CoM movement in the horizontal direction is an important measure. By projecting a velocity ellipse on x -axis, we obtain a line which its length equals to the maximum change of velocity of the CoM in the horizontal direction. Figure 7 (left side) shows maximum instantaneous change of the CoM velocity in the horizontal direction for different constrained hand loca-





Figure 7: Maximum instantaneous change of the CoM velocity in the x direction due to the unit norm of instantaneous change of the joint velocities for (top) the constrained robot and (bottom) for both constrained and unconstrained robots.

tions (*i.e.* different elbow and shoulder angles). This is due to the unit norm of instantaneous change of the joint velocities. In the right side of this figure, the graph at the left side is compared with the case that the hand is not constrained. It is obvious that the movement of the CoM is limited due to the kinematic constraint at the hand.

Figure 8 shows maximum instantaneous change of the CoM velocity due to the unit impulse at the joints for different hand locations and for both constrained and unconstrained hands. As it can be seen in this figure, the graph for the constrained robot is always higher than the other one. This implies that the same amount of impulse can cause bigger changes at the CoM velocity in the constrained robot rather than the unconstrained one. The reason is that, in the constrained case, the robot exploits the contact force to accelerate the CoM and therefore less (impulse) torque is needed for the same change at the CoM




Figure 8: Maximum instantaneous change of the CoM velocity in x direction for the constrained and unconstrained robots due to the unit norm of impulse at the joints.

velocity.

3.2. Strategies of dealing with uncertainties in contact

A novel method to study human strategies of dealing with contacts with uncertain environment was developed. In this method a human subject was made to perform psychical contacts with the environment through the robot. The human was included into the robot control loop through human-robot interfaces. The idea is that the human sensorimotor system and cognitive system controls a novel mechanical system, *i.e.* the robot, in physical interaction with the environment. This implies additional human motor control learning and adaptation that can potentially provide us with a deeper insight into how humans deal with a novel environment.

Another advantage of this approach is that the human sensorimotor system does not use its own limbs to directly make the contacts with the environment, but uses the robotic limb to do so. Compared to pure biomechanical studies, where the measured human behaviour must be further interpreted, adjusted or transformed before it can be used on the robots, in this approach the measured human behaviour can be directly captured and used in the robot control. This study therefore provides a good complement to our conventional biomechanical studies.

The block scheme of the proposed approach is shown in Figure 9. The human controlled the motion of the robotic limb with the motion of his/her own limb. In addition to controlling the motion, the human also controlled the impedance of the robot. Primary information about the robot state was relayed to the human through a visual feedback. A haptic device was used to provide the human with an additional feedback about the forces sensed by the robot. While controlling the robot in the proposed human-in-the-loop approach, the human central nervous




Figure 9: Block diagram of proposed human-in-the-loop robot control framework for study of human behaviour in contacts with environment. During the learning and adaptation stage, the human performs the contacts with the environment through the robot (blue section). The acquired data was used to observe and study the human behaviour. When the human learning process and observation is complete, the learnt skill can be directly captured and used in the autonomous robot control (red section). This is the main advantage compared to the conventional biomechanical studies.




Figure 10: Human-robot interfaces. First developed interface (left) measured human limb motion via optical motion capture system and mapped it to the motion of the robotic limb. The human muscle activity was measured by sEMG and was used as an interface to control the robot impedance. Second developed interface (right) consisted of *HapticMaster* robot and impedance control handle. *HapticMaster* robot measured the human limb position and provided the force feedback. Impedance control handle was based around a spring-return linear potentiometer and was held in the human hand.

system had to adapt to a new mechanism through sensorimotor learning to perform the desired contact with the environment.

The main goal of studies of human behaviour in contacts with environment is to offer a basis from which we can devise equivalent humanoid robot behaviour. The most appealing prospect of the proposed approach to study human motion in contacts with the environment is that the data from the study can be used to directly form skills for autonomous robot control. The sensorimotor data was collected while the human was making the desired physical contacts with the environment through robotic mechanism. This data was then used to form the trajectories. The trajectories were encoded with Dynamical Movement Primitives (DMPs) [35]. The parameters of DMPs were learned by locally weighted regression [85]. The learned trajectories represented the robot skill for dealing with the contacts with the environment according to the human strategy. The trajectories can be included into the robot control system and used for autonomous execution of the learnt task.

One of the key features of the proposed approach is the ability of the human to directly control the impedance of the robot limb in an equivalent way that he/she controls his/her own. For this purpose we developed two




Figure 11: The subject was standing on a force plate, connected to the motorised waist-pull system that generated translational perturbations. The subject was holding the handle with a built-in force sensor mounted on a vertical pole. EMG electrodes were positioned on the major body muscles of the subject’s right-hand side.

novel human-robot interfaces [74, 73] that allow the human to modulate the stiffness of the robotic limb in real-time. The first interface (see Figure 10, left) was based on measuring human muscle activity by surface electromyography (sEMG). The current measured muscle activity was mapped to the robot stiffness. The second interface (see Figure 10, right) was based around a linear potentiometer inside a handle held in the human hand. The human controlled the position of the potentiometer knob with a finger position. The finger position is then mapped to the robot stiffness via measured potentiometer voltage.

3.3. Human contact choice and learning through physical interaction

In order to understand how humans make contact choice decisions (*e.g.* whether or not to initiate a hand contact, and where to place the hand), we need an estimation of joint torques as well as a metric of stability in various multi-contact situations.

3.3.1. Motor adaptation with supportive hand contacts

In continuation of this work, we studied how additional hand contact with the surrounding objects influences whole-body balance conditions. The experiments were performed on multiple subjects where we challenged their balance. The experiments were divided into two main stages. Each stage had 15 sessions in which the subject’s balance was perturbed for 5 minutes. In one stage the subjects did not use supportive hand contact. In the other stage they were holding a handle in front of them. We used




Figure 12: Effect of holding a handle after adaptation stabilised in the last session. The graphs show representative power spectral density (PSD) profiles of CoP_{AP} and muscle activations measured in trunk and lower leg muscles. After the adaptation, effect of additional supportive hand contact stabilised to the perturbation in the last session. All EMG and CoP_{AP} values are presented in a frequency domain, ranging from 0.25 Hz - 1 Hz. The blue (solid) lines represent the power in no-handle and the orange (dashed) lines in handle stage. The grey (dotted) line is the power of the perturbation signal. All signals are normalized to the peak value in the last session. The effect of handle is shown as reduced muscle activation in all muscles in the handle session, except in the trunk flexor muscle (OE), where there is an opposite effect.

a motorised wait-pull mechanism [72] to continuously perturb the balance of the standing subjects in either stage by exerting external forces on the approximate position of centre of mass. See Figure 11 for the experimental setup. The perturbation waveform of the waist-pull mechanism was constructed in a way that the possible muscle reactions associated with reflexes were eliminated. These reactions could potentially mask the actual role of the hand muscles as the reflex would activate the muscles unrelated to the magnitude of the perturbation. To avoid that, the perturbation waveform was continuous, had relatively low frequency and low pulling forces. During the experiment, we measured muscle activation of the subject’s lower leg, trunk and arm muscles, forces in the handle and the anteroposterior movement of CoP (CoP_{AP}). The results of muscle activation analysis showed that when the subjects were holding to the handle, the activation of the leg muscles was minimal (see Figure 12). Based on this we can conclude that the subjects mainly used their arm muscles to maintain postural stability. The trunk flexor muscle (Obliques Externus, OE) was more active in the stage when the subjects were holding the handle compared to when they were not. This indicates that a synergy between the arm and trunk muscles was established when additional hand contact was utilised to maintain the equilibrium.

The analysis of the CoP_{AP} movement showed that the displacement of the CoP_{AP} was progressively dropping throughout the repeated sessions of the experiment (see Figure 13). This was true both in case when supportive hand contact was used and in case when no supportive




Figure 13: Adaptation of movement of CoP_{AP} is shown on the left graph. Experimental stage with handle (WH) is shown in red, while condition without handle (NH) is shown in blue. Full markers indicate statistically significant differences between the first session and each of the following sessions. In both stages the adaptation is statistically confirmed ($p < 0.001$). In the stage where the subjects were holding the handle, the adaptation appeared right after the first session. In no-holding stage it appeared after the third session. The superimposed best-fit curves are shown on the right graphs with orange solid lines. A calculated session number at 3τ of the fitted curve (vertical dotted line) indicates faster stabilisation of adaptation in the handle stage, compared to no-handle stage.

hand contact was used. These results give a strong hint that a learning and adaptation mechanism was present through the sessions of the experiment, as the subject gradually improved the balance control.

We further analysed whether there are any effects of repeated sessions on adaptation of muscle activation and movement of CoP_{AP} , and whether there are any differences between the two stages of the experiment. The results show that the effect of human adaptation in lower leg muscles was statistically significant in the stage when the subjects were not using the additional hand support. However, this was not the case for the stage when the subjects were holding to the handle. The activation of the trunk extensor muscle (MF) was almost the same in both stages and throughout all sessions. On the other hand, the activation of the trunk flexor OE remained unchanged throughout the sessions only in the stage when subjects held the handle. The activation of OE was much higher in this stage compared to stage when subject did not use supportive hand contact.

We performed an analysis of differences in EMG activation levels between the two experimental stages in the frequency spectrum of the perturbation waveform (low = 0.25 - 0.5 Hz, medium = 0.5 - 0.75 Hz, high = 0.75 - 1.0 Hz). A paired samples analysis between the two stages for low, medium and high frequency range revealed that there was an influence of additional hand contact on both lower leg muscles. There were confirmed statistically significant differences between the two stages in all frequency ranges and for all sessions. For the MF muscle these differences were not significant in any of the frequency range nor session. However, there were significant differences between the two stages for the OE muscle. These differences occurred in the medium and




Figure 14: Absence of effect of the repetition of sessions on the handle forces.

high frequency range but only in the last session. When the subjects were holding to the handle, we recorded the forces exerted on the handle during the continuous postural perturbations. Statistical analysis of handle forces revealed that the repetition of sessions had no significant effects (see Figure 14). Even though the activation of arm extensor muscle changed (decreased) during sessions, there was no significant change in forces applied on the handle.

We performed an analysis of differences in EMG activation levels between the two experimental stages in the frequency spectrum of the perturbation waveform (low = 0.25 - 0.5 Hz, medium = 0.5 - 0.75 Hz, high = 0.75 - 1.0 Hz). A paired samples analysis between the two stage for low, medium and high frequency range revealed that there was an influence of additional hand contact on both lower leg muscles. There were confirmed statistically significant differences between the two stages in all frequency ranges and for all sessions. For the MF muscle these differences were not significant in any of the frequency range nor session. However, there were significant differences between the two stages for the OE muscle. These differences occurred in the medium and high frequency range but only in the last session. When the subjects were holding to the handle, we recorded the forces exerted on the handle during the continuous postural perturbations. Statistical analysis of handle forces revealed that the repetition of sessions had no significant effects. Even though the activation of arm extensor muscle changed (decreased) during sessions, there was no significant change in forces applied on the handle.

3.3.2. Planned vs reactive contact models

We studied whether supporting contacts in human arm reaching tasks are planned or an effect of a reactive controller. Investigations on human motor learning has focused on adaptation experiments with fixed contact points leaving research on the computational role of contacts as




Figure 15: Trunk trajectories predict wrist trajectories. (a) 600ms of trunk trajectories are observed. These observations can predict the wrist trajectories. Shown are predictions for the two exterior targets on the screen. For training 10 trials for each target are used starting from trial 240 backwards in time (before the catch trials). For testing the first perturbed trial after trial number 240 were used. (b) The effect of the observation horizon on the target prediction error is shown for a representative subject. The mean of the training data denotes the base line (BL). (c) Average statistics (mean and 95 percent confidence bound) over 20 subjects.

a free control variable unexplored. In perturbed target reaching experiments sketched in Figure 15, we studied whether supporting contacts are planned or reactive. Subjects had to reach for distant targets on a screen with their right hand. For reaching the target additional support through contacts with a table using the left hand was inevitably. If the contacts are planned then the left hand’s motion can predict the right hand reaching. We studied how probabilistic inference in learnt models can be used to answer this question. Evidence for planned contacts could be provided through learning probabilistic models of trajectory distributions and using the models to generate predictions, Figure 15 (a). We found that the target on the screen could be predicted from both, the left hand (mse: $10.4\text{cm} \pm 2\text{cm}$ over 20 subjects) and the trunk movement (mse: $6.7\text{cm} \pm 1.4\text{cm}$ over 20 subjects), which is illustrated in Figure 15 (b-c). The learnt probabilistic model could also be used to analyse the rate of adaptation of the left hand and the trunk kinematics, where the trunk trajectories converged faster than the left hand motion. This is intuitively explained by the strong need for corrective trunk movements in balancing. A report on the findings is currently in progress of writing.

3.3.3. Time and precision trade-off in supportive hand contacts

Driven by the question on how human CNS optimizes arm reaching motions when the supportive hand contact has to be reached in order to maintain postural balance, a combined experimental and computational study was started where the aim is to challenge two well-established but conceptually separated motor control phenomena: (i) Humans tend to reach faster to a target that looks more rewarding, despite the additional muscular cost of a faster movement [24], and (ii) when humans have to be precise, movements take longer to perform [91]. The aim of our study is to experimentally disclose both phenomena and evaluate a novel computational model designed to join them. We obtained several very promising prelimi-




Figure 16: Experimental setup to understand how humans optimize arm reaching motions when the supportive hand contact has to be reached in order to maintain postural balance. The task of the subject was to obtain as high reward as possible in the given time by hitting a target on the virtual wall without knowing its actual size. In effect, the subjects had to find the optimal balance between precision, speed of motion and its cost in order to maximise the reward. To amplify the effect of cost of motion, haptic robot emulated a viscous media through which the subject had to move the hand.

nary results indicating a general mechanism that can unify both phenomena and point out a global trade-off arising from the interactions between movement time, cost and accuracy. The experimental set-up is shown in Figure 16.

4. Whole-body controllers

The overall objective of the work on controllers is to provide a control architecture dedicated to humanoid robots involved in personal/service applications that imply physical interactions, *i.e.* contacts, with the environment. Such a control architecture is a requirement to bridge the existing gap between state-of-the-art methods in humanoid robots control and real-world applications. This gap is, at the control level, mostly due to two factors. The first one is the intrinsic complexity of the robotic system itself as well as the complexity of the environment. This complexity induces uncertainties in the knowledge of the models. The second factor is related to the complexity of the decision making process which, in real world applications, can be very challenging especially when dealing with missions implying the sequenced and/or parallel realization of complex actions by the robot. The combination of those two complexities results in third factor, related to the large computation times necessary to take control decisions. The state-of-the-art methods in humanoid robot are mostly two: pure motion planning and pure reactive control. The former tries to solve off-line the overall decision-making process but the actual action execution phase (typically open-loop) tends to fail because of the “complexity and uncertainty” factors. The latter succeeds in overcoming the “complexity and uncertainty” factor mostly thanks to the use of feedback. However the proposed solutions are only locally optimal and the overall decision-making process cannot be addressed in the most general cases (*i.e.* without scripted scenarios). The path

followed by the CoDyCo project to achieve both globally optimal and locally feasible control policies is based on a control architecture featuring two intertwined levels. The first one is the central node of the work described here: given a set of elementary operational tasks to achieve and their respective importance, it provides a framework to compute the torques to be produced by the actuators at each time in order to achieve at best the prescribed set of elementary tasks given some constraints acting on the system (limits, saturations, local obstacles, contacts...). We call this level the “local controller” level. The second level directly impacts the objectives to be achieved by the local controller, temporally sequences and parametrizes the use of elementary operational tasks in order to achieve some complex goal (*e.g.* “grabbing an object on a table while standing and balancing using several contact points”) in a globally optimal fashion. We call this level the “global control policy” level.

4.1. Formulating and solving the local control problem

The work performed on formulating and solving the control problem has led to the definition of what a task can be considered to be in the context of the reactive formulation of a multi-task whole body control problem. Among the different characteristics of a task (physical frame, task variable, forward model, desired target trajectory, local controller, priority), the notion of task priority has been largely modified with respect to the classical lexicographic task ordering met in the robotics literature and which is particularly appropriate for cascade resolution approaches such as the one recently proposed in [22]. A partial order has been defined such that task priorities can be described for any pair of task i and j . This leads to a richer formulation which includes the original one but is also particularly appropriate for describing task insertion and removal processes as well as priority switching between tasks. Furthermore, this new prioritization paradigm provides a unique way of defining strict and soft hierarchies between tasks. Associated to this work, the notion of generalized task projector has been introduced. Each task is associated to a projector which is built based on the tasks priorities. The interest of this projector is that it filters the joint space motion associated to a task so that all priorities are respected, being them soft or strict.

The control problem has been formulated as an LQP [82] which can be solved by any convex optimization solver dealing with linear constraints. Despite the task hierarchy, the introduction of a generalized task projector per task allows to solve only one LQP. This can be done by introducing as many virtual joint space variables as the number of tasks and using the generalized projector of each task in the expression of the constraints. The resulting problem can be solved by standard convex optimization tools and the cost of introducing virtual joint space variables is compensated for by the fact that only one optimization problem has to be solved. Details regarding this work, so-called

Generalized Hierarchical Control (GHC), are provided in [51] and a humanoid implementations is described in [50].

A more classical hierarchical controller has also been derived and is the one currently in use on the real robot [18], [66] and is described in the next section related to validation scenarii for the CoDyCo project.

We also started to investigate scenarios where the robot is interacting with the environment through rigid and non-rigid contacts. Assuming that no information is *a priori* available regarding the nature of the contact surface, a first control strategy has been proposed in [52] where the desired contact force is adapted online as a function of the velocity of the contact point. Indeed, the risk with an unknown contact surface is to assume that it will almost instantaneously provide the required contact force to maintain the robot balance. If the surface is non-rigid, the contact point will actually move while being pushed and stable support forces will only be provided to the robot once the contact is properly established. The goal of the adaptation of the desired value for the contact force is to accelerate the attainment of a stable contact force supporting the robot. The desired trajectory for the center of mass of the robot is also adapted to account for the non-rigidity of the contact surface. One of the advantages of this approach is that it does not actually requires the knowledge of the contact surface impedance. Figure 17 provides a view of the types of considered scenarii and the structure of the considered controller. In this work, the local control problem is solved using the solver described in [81] rather than the one developed in [51]. This choice is related to the fact that the computation cost of the GHC approach remains important and is too high to be actually used in a real-time reactive control architecture for a humanoid robot.




Figure 17: Scenarios of interaction with a non-rigid environment (top). Structure of the adaptive control architecture (left).

4.2. Bootstrapping and validating the control approach in rigid world and compliant cases

In coordination with the local controller, we also have explored the contribution of MPC (Model Predictive Control) approaches to handle the postural balancing problem under varying contact conditions. The hybrid nature of the problem, where varying contact conditions can be accommodated either by adapting the internal forces distribution given a set of contact or by modifying the set of contacts itself, requires control approaches where the desired task trajectories performed through the local, reactive, whole-body controller have to be optimally planned ahead of time in order to provide robust behaviours. The contributions in this domain are mostly related to the work of A. Ibanez [31], [33] and [32]. The originality of these contributions lies in:

- an augmented ZMP (Zero Moment Point, [97]) model including external forces exerted directly on indirectly on the center of mass;
- a distributed optimization approach that provides a way of generating reference trajectories for the center of mass representing a good compromise given some antagonistic balance and task;
- a non scripted foot step placement optimization.

As a continuation of these works, in order to compute optimal time, duration and position of footsteps along with the center of mass trajectory of a humanoid, a novel mixed-integer model of the system is introduced in [34]. The introduction of this model in a predictive control problem brings the definition of a Mixed-Integer Quadratic Program, subject to linear constraints. Simulation results demonstrate the simultaneous adaptation of the gait pattern and posture of the humanoid, in a walking activity under large disturbances, to efficiently compromise between task performance and balance. In addition, a push recovery scenario displays how, using a single balance-performance ratio, distinct behaviours of the humanoid can be specified. Results have been obtained in simulation³ and are being implemented on the TORO robot developed at DLR. Two simple and novel approaches to solve for 3D locomotion with multiple non-coplanar contacts have also been explored in [71]. Both formulations use model predictive control to generate dynamically balanced trajectories with no restrictions on the center of mass height trajectory. The first formulation treats the balance criterion as an objective function, and solves the control problem using a sequence of alternating convex quadratic programs, while the second formulation considers the criterion as constraints to the problem, and solves a succession of convex quadratically constrained quadratic programs. Preliminary results have been obtained in a scenario




Figure 18: The robot performs a circle with its left arm. The forearm collides alternatively with the left, the right or both contacts.




Figure 19: Prediction of torques with multiple contacts and the corresponding activation of the gating network. Our mixture-of-experts model combines the learned single-contact models into a multiple-contact model which outperform the analytic approach.

where a hand contact on a vertical wall is used to improve balance. A staircase climbing scenario has also been studied.

Bootstrapping between the local controller and a more global reasoning approach also lies in the capability to incrementally learn and adapt the models used for control. Thus, we continued research in inverse dynamics model learning in situations with contacts. A mixture of experts approach combined with Gaussian Processes was proposed in [11], to learn the torque contributions due to contact exploiting the iCub's tactile and force/torque sensors. This approach was evaluated on the iCub robot, where the learned model accurately predicts contact forces, is robust to changes in the environment and outperforms existing analytic dynamic models that make use of force/torque sensor data. The interest in the use of such learned models over analytical ones lies also in the fact that learned models do not require a spatial calibration of the skin taxels, a procedure that in complex robots such as iCub is often prone to errors that significantly impact the torque estimation [17]. An exemplary task is illustrated in Figure 18 when obstacles are introduced on both sides of a planned circular motion. In Figure 19, it can be seen that the mixture-of-experts recognizes the presence of the two different contacts and opportunely active the corresponding expert to compensate for the contact. As a result, the torques predicted from this approach (red curve) closely follow the ground truth (blue curve) and outperform the analytic model (green curve).

4.3. Validation scenarios

This section presents the whole-body control framework implemented on the humanoid robot iCub for one foot balancing and motion control. This framework ensures a degree of compliance for the multi-body system, which allows for safe human robot interaction.

³A video associated to this work can be found [here](#).

4.3.1. System modelling

The system dynamics are characterized by the following differential equations:

$$M(q)\dot{\nu} + h(q, \nu) - J^\top(q)f = S\tau, \quad (1)$$

$$J(q)\dot{\nu} + \dot{J}(q, \nu)\nu = 0, \quad (2)$$

where $q \in SE(3) \times \mathbb{R}^n$ represents the configuration of the free floating system, which is given by the pose of a base-frame and n generalized coordinates q_j characterizing the joint angles. The vector $\nu \in \mathbb{R}^{n+6}$ represents the robot velocity (it includes both $\dot{q}_j \in \mathbb{R}^n$ and the linear and angular velocity of the base-frame $v_b \in \mathbb{R}^6$), the system acceleration is denoted as $\dot{\nu}$, the derivative of ν , the control input $\tau \in \mathbb{R}^n$ is the vector of joint torques, $M \in \mathbb{R}^{(n+6) \times (n+6)}$ is the mass matrix, $h \in \mathbb{R}^{n+6}$ contains both gravitational and Coriolis terms, $S \in \mathbb{R}^{n \times (n+6)} := (0_{n \times 6}, I_n)^\top$ is the matrix selecting the actuated degrees of freedom, k is the number of constraints acting on the system, $f \in \mathbb{R}^k$ are the generalised forces associated to the constraints, and $J \in \mathbb{R}^{k \times (n+6)}$ is the constraint Jacobian (2).

4.3.2. Problem statement

The control objective is the asymptotic stabilization of a desired centroidal dynamics [67]. Let H denote the centroidal momentum of the robot. Then, the time derivative of H is equal to the summation of the external wrenches acting on the multi-body system. By expressing the centroidal momentum with respect to the center of mass, we have:

$$\dot{H} = Xf + mg = \begin{pmatrix} m\ddot{x} \\ \dot{H}_\omega \end{pmatrix} \quad (3)$$

where m is the mass of the robot, $g \in \mathbb{R}^6$ is the gravitational acceleration, $\ddot{x} \in \mathbb{R}^3$ is the acceleration of the center of mass, $H_\omega \in \mathbb{R}^3$ is the *angular momentum* of the robot, and the matrix X maps the contact wrenches on the center of mass.

The control objective is to find a control law for the inputs τ such that $x \rightarrow x(0)$ and $H_\omega \rightarrow 0$. This choice is sufficient for balancing purposes. Also, while achieving this control objective, the system shall have a degree of compliance.

4.3.3. The control strategy

The control strategy is composed of two steps. We first choose the external force f such that $x \rightarrow x(0)$ and $H_\omega \rightarrow 0$. Then, we generate this force through the internal torques. Since iCub possesses more than six degrees-of-freedom, which are necessary to generate the contact force f , we choose the remaining control inputs so that to have compliance at the joint level.

4.3.4. The choice of the contact force

Being the matrix X invertible, the contact force f achieving the control objective may be chosen as follows:

$$f = -X^\dagger \left[k_d H + k_p \begin{pmatrix} x - x(0) \\ 0_{3 \times 1} \end{pmatrix} + mg \right] + (I - X^\dagger X) f_0, \quad (4)$$

with k_d and k_p two positive constants and f_0 arbitrarily chosen to obtain a solution f as similar as possible to a desired value.

In order to keep the motion constraints satisfied, f must satisfy some constraints, *e.g.*, the contact forces must belong to the associated friction cones. In general, the contact constraints can be represented by inequalities of the form $Cf < d$, with the matrix C and the vector d properly chosen. Then, we choose the contact wrench as follows:

$$f = \underset{\xi \in \mathbb{R}^6}{\operatorname{argmin}} \|\xi - f_d\|^2 \quad (5a)$$

$$\text{s.t. } C\xi < d, \quad (5b)$$

with the desired wrench f_d given by (4).

4.3.5. The choice of the joint torques

The control input τ must generate the force f . The relationship between the contact wrench and the joint torques can be obtained by using the constraint equation along with the free-floating dynamics, *i.e.* Eq. (1). One can show that the torques generating f are given by the summation of two terms, *i.e.*,

$$\tau = \tau_f + N\tau_0, \quad (6)$$

where τ_f ensures $f = f_d$, the matrix $N \in \mathbb{R}^{n \times n}$ is the null space projector of $JM^{-1}S$, and τ_0 is a vector that can be chosen at will. To obtain compliance at joint level, we choose τ_0 similar to a gravity and external force compensation, plus a term of the form

$$-k(q_j - q_d),$$

which ensures compliance at joint level.

4.3.6. Experiment

We implemented the proposed control strategy on the iCub platform as illustrated in Figure 20. The control framework is composed of two loops. The inner loop is in charge of stabilizing *desired* joint torques, while the outer loop is governed by Eq. (6). Both loops runs at the same frequency of 100 Hz.

The experiment consists in two phases. In the first phase, we change the desired q_d in order to generate *internal* motions, which do not perturb the stability of the robot momentum thanks to the prioritization of tasks described in the previous section. In the second phase, we apply external perturbations by interacting with the robot




Figure 20: Screenshot of the video showing the full experiment. iCub balances by controlling the foot wrench.

as illustrated in Figure 21. This interaction results to be safe thanks to the compliance at joint level⁴.

For more detailed information and description of the system architecture (comprising torque and forces estimation and low level torque control) see [66].




Figure 21: The picture shows the iCub while performing compliant single foot balancing. Details on the controller can be found in [66]. A video of the task is available on [youtube](#).

5. Learning

The goal of the work on Learning in CoDyCo is to endow humanoid robots control architectures with the core abilities for the adaptation, generalization and self-improvement of both control laws and tasks that involve physical interaction with humans, and the environment. In this context, we propose learning approaches that work in conjunction with the control architecture devised in the previous section and rather complement analytical robotic approaches with on-policy learning than starting from scratch. A core idea behind this work is that Learning should complement classical approaches and not supersede them.

⁴A video of the experiment is available [here](#) for the interested reader.

5.1. Inferring the Operational Space and Appropriate Controls with Multiple Contacts

For controlling high-dimensional robots, most stochastic optimal control algorithms use approximations of the system dynamics and of the cost function (*e.g.*, using linearisations and Taylor expansions). These approximations are typically only locally correct, which might cause instabilities in the greedy policy updates, lead to oscillations or the algorithms diverge. To overcome these drawbacks, we added a regularization term to the cost function that punishes large policy update steps in the trajectory optimization procedure. We applied this concept to the Approximate Inference Control method (AICO), where the resulting algorithm guarantees convergence for uninformative initial solutions without complex hand-tuning of learning rates.

The new algorithm was evaluated on two simulated robotic platforms. A robot arm with five joints was used for reaching multiple targets while keeping the roll angle constant. On the humanoid robot Nao, we show how complex skills like reaching (see Figure 22) and balancing can be inferred from desired center of gravity or end effector coordinates. This work was published at the international conference on humanoid robots [79]⁵.

5.2. Generalizing and Improving Elementary Tasks with Contacts

We aim to generate new skills from data, where elementary skills are acquired by imitation learning and transferred to novel situations using dynamic systems. To do so, we developed a novel representation of movement primitives that can be used for imitation learning from noisy observations. Uncertainty of observed trajectories is explicitly modelled and used to generate new skills. This movement representation has state-of-the-art capabilities in generalization, coupling between the degrees of freedom of the robot, and moreover, a time varying feedback controller can be derived in closed form. These features are partially illustrated in Figure 23. More details on this work can be found in [68].

The advent of robots in our every day life can only be accomplished with reliable mechanisms for movement generation. Movement Primitives (MP) are a well-established approach for representing modular and re-usable robot movement generators that can be composed into complex movements. An easy-to-learn representation of the primitive is, additionally, the key of recent imitation and reinforcement learning successes. Current MP's approaches offer viable properties such as concise representations of the inherently continuous and high dimensional space of robot movements, generalization capabilities to novel situations, temporal modulation of the primitive, sequencing of primitives, coupling between the degrees of freedom of the robot, and controllers for real time execution. However, no single MP framework exists that offers all these

⁵Supplemental Matlab demo code is available [online](#).




Figure 22: Reaching task with the humanoid robot *Nao*. The robot has to reach a desired end effector position with the right arm while maintaining balance. Eight snapshots of the inferred movement are shown in (A). In (B), the convergence of the costs of the optimization procedure is shown, where we compare *iLQG*, the standard implementation of AICO and the regularized variant. The movement objectives for the right arm are shown in the left panel in (C). To balance the robot lifts its left hand and bends the head back.




Figure 23: (a) Conditioning on different target states. The blue shaded area represents the learned trajectory distribution. We condition on different target positions, indicated by the x -markers. The produced trajectories exactly reach the desired targets while keeping the shape of the demonstrations. (b) Combination of two ProMPs. The trajectory distributions are indicated by the blue and red shaded areas. Both primitives have to reach via-points at different points in time, indicated by the x -markers. We co-activate both primitives with the same activation factor. The trajectory distribution generated by the resulting feedback controller now goes through all four via-points. (c) Blending of two ProMPs. We smoothly blend from the red primitive to the blue primitive. The activation factors are shown in the bottom. The resulting movement (green) first follows the red primitive and, subsequently, switches to following the blue primitive.

properties. We extended previous results on modelling stochastic movements [69, 70].

We incorporated all the desirable properties current approaches offer into a single framework and, additionally, we introduced new operations on the primitives, such as continuous blending and co-activation of multiple primitives. Most importantly, in this approach, the novel co-activation operator is capable of solving multiple tasks concurrently as illustrated in Figure 24. Furthermore, our approach is capable of reproducing exactly the demonstrated variability of the movement and the coupling between the degrees of freedom of the robot. In this approach, called Probabilistic Movement Primitives (ProMPs) [69], we derived all operations in closed form. In order to use the ProMPs for online feedback control, we also derived a stochastic

feedback controller that reproduces exactly the encoded primitive. We evaluated and compared this approach on several simulated and real robot scenarios.

Probabilistic movement primitives are a promising approach for learning, modulating, and re-using movements in a modular control architecture. To effectively take advantage of such a control architecture, ProMPs support simultaneous activation, match the quality of the encoded behavior from the demonstrations, are able to adapt to different desired target positions, and efficiently learn by imitation. ProMPs meets all of the aforementioned requirements. The desired trajectory distribution of the primitive is parametrized by a hierarchical Bayesian model with Gaussian distributions. The trajectory distribution can be obtained from demonstrations and simultaneously




Figure 24: Robot Hockey. The robot shoots a hockey puck. We demonstrate ten straight shots for varying distances and ten shots for varying angles. The pictures show samples from the ProMP model for straight shots (b) and angled shots (c). Learning from the union of the two data sets yields a model that represents variance in both, distance and angle (d). Multiplying the individual models leads to the combined a model that only reproduces shots where both models had probability mass, in the center at medium distance (e). The last picture shows the effect of conditioning on only left and right angles (f).

defines a feedback controller which is used for movement execution. Currently, we are investigating extensions of the ProMPs framework to tasks that involve contacts with the environment. In addition, we started to investigate the improvement of elementary skills encoded in ProMPs with reinforcement learning, where a conference paper was submitted for review.

5.3. Learning the Prioritization of Tasks

We have been leading research on computed torque control leveraging low-gain control. In computed torque control, robot dynamics are predicted by dynamic models. This enables more compliant control, as the gains of the feedback term can be lowered, because the task of compensating for robot dynamics is delegated from the feedback to the feed-forward term. We already showed that Gaussian Process regression is an effective method for learning computed torque control, by setting the feed-forward torques to the mean of the Gaussian Process. During the second year of the project, we extended this work by also exploiting the variance predicted by the Gaussian Process, by lowering the gains if the variance is low [4]. This enables an automatic adaptation of the gains to the uncertainty in the computed torque model, and leads to more compliant low-gain control as the robot learns more accurate models over time. On a simulated 7-DOF robot manipulator, we demonstrated how accurate tracking can be achieved, despite the gains being lowered over time, which is illustrated in Figure 25.

We also studied how to deal with interferences between tasks using machine learning tools. Whole-Body Control methods offer the potential to execute several simultaneous tasks on highly redundant robots, such as humanoids. Unfortunately, task combinations often result in interferences or incompatibilities which generate undesirable behaviours. Prioritization schemes between tasks, such as strict and soft hierarchies, are typically used to manage




Figure 25: Mean and variance of the Gaussian process ($\mu \pm 2\sigma$) on the same test trajectory of 8 seconds, after having been trained with 5, 10, 20 and 40 training trajectories. With an increasing number of training data, the mean of the GP approaches the true function (black dashed line). The known values for u_{ff} are plotted as a black dotted line.

these interferences but generally require a deal of time consuming and arbitrary tuning.

To circumvent these issues, we presented a novel framework for defining and optimizing multiple tasks in order to resolve potential interferences prior to task execution. In a first study [53] the tasks are parametrized with Dynamical Movement Primitives, whose parameters are optimized based on a general compatibility principle, which is independent of the robot's topology, tasks or environment. Two test cases on a simulation of a humanoid robot are used to demonstrate the successful optimization of initially interfering tasks⁶.

In a second study [54], we studied how task variability can be used to modulate task priorities during their execution, to temporarily deviate certain tasks in the presence of incompatibilities. A method for mapping from task variance to task priority was presented as well as

⁶A video summarizing the outcome of this work can be viewed [here](#).




Figure 26: Three common multi-task incompatibility scenarios. The desired hand task trajectories are indicated by the green markers. Medium size spheres represent way-points, and large transparent spheres represent the final waypoints or goals.

an approach for calculating task variance for generated trajectories. The method successfully resolved three common task conflict scenarios online illustrated on Figure 26.

We also addressed the problem of learning the temporal profile of soft task priorities and null-space projectors for the multi-task controllers developed in [51].

The *Prioritized Task-Space Inverse Dynamics* controller from [18] is based on *strict task hierarchies*, where a hierarchical ordering of the tasks is set, such that critical tasks (or tasks that are considered as more important) are fulfilled with higher priorities, and the low-priority tasks are solved in the null-space of the higher priority tasks. The strict control approach requires the pre-specification of the task hierarchy. However, in many contexts it is difficult to organize the tasks in a stack and define their relative importance in forms of priorities. When priorities are strict, a higher task can completely block lower tasks, which can result in movements that are not satisfactory for the robot mission (*e.g.*, its “global” task). Swapping tasks can result in discontinuities of the control law. Moreover, the task priority must be defined *a priori*, which is typically done by an expert able to define the tasks and their relative priorities to produce smooth behaviours.

Another set of whole-body controllers is based on *soft task hierarchies*, where the solution is typically given by a combination of weighted tasks (see for example [82]). The importance or “soft priority” of each individual task is represented by a scalar function, often named “weight”. By tuning the vector of scalar weights, evolving in time, the global robot behaviour can be optimised. Liu et al. [51] propose a generalised projector (GHC) that handles strict and non-strict priorities with smooth transitions when tasks priorities are swapped. They show that adapting these weights may result in a seamless transition between tasks (*i.e.*, reaching for an object, staying close to a resting posture and avoiding an obstacle) and in continuous task sequencing. Despite the elegant framework, their controller needs again a lot of manual tuning: particularly,




Figure 27: This scheme briefly describes the proposed method. The control torques are computed by a combination of tasks weighted by soft priorities, represented as parametrized activation policies, that are multiplied by a Null-space projector, where some activation functions for different projectors are joined. The global task execution is evaluated and a fitness function is computed: a policy search method is then used to optimize the parameters of the activation policies, both for tasks and projector.

the evolution of the tasks priorities in time, the timing and the tasks transitions need to be designed by hand. While this approach could still be easy for few tasks and simple robotic arms, it can quickly become infeasible for complex robots such as humanoids performing whole-body movements that usually require multiple tasks and constraints (*e.g.*, control balance, posture, end-effectors, stabilise head gaze, prevent slipping, control interaction forces etc.).

An intuitive solution to the problem of defining the weights is to learn or optimize them through trial-and-error. We study how the *temporal* profiles of the task weights can be learned from a reward function, which is assumed to be given⁷. As a learning algorithm, we choose CMA-ES [28], a stochastic optimization derivative-free algorithm that is suitable for the optimization of non-linear and non-convex objective functions. The main reason for using CMA-ES is that it is relatively simple to use and with a small number of free parameters. However, any other optimization algorithm with similar properties could be used. As a first step towards a controller that is capable of handling multiple tasks and constraints on a complex robot, while allowing us to efficiently learn the task priorities, we propose a regularized version of the Unified Framework (RUF) proposed by Peters et al [75], where the tasks weights and Null space projectors weights are represented by parametrized functional approximators that can be automatically determined through a stochastic optimization procedure. The concept is presented in Figure 27.

As a first results, we show that the optimization process generates weights profiles that cannot be designed manually in advance, see panels (a) and (b) in Figure 28.

We then compare the performance of our controller with the state-of-the-art method GHC proposed by Liu et al. [51]. We consider the following experimental scenario: a 7-DOF KUKA Light Weight Robot arm, starting from a vertical position, must reach a desired point with its end-effector, while avoiding to collide with a table, represented

⁷For many robotic task, *e.g.*, tracking desired center-of-mass or end-effector trajectories while avoiding obstacles, such reward functions have been defined in [42].




Figure 28: The panels in (a) and (b) show the mean and standard deviation of the temporal profile of the activation functions α, β , optimized by RUF+CMA-ES, computed over $R = 50$ replications of the same experiment of the table scenario. (c) Comparison of our method (blue line) to the generalized projector method (GHC).

by a surface parallel to the z-axis in between the robot and the goal. The aim of the experiment is to bring the end-effector as close as possible to the desired position, while avoiding collisions with the obstacle. We define $T = 3$ tasks: a regulation task in the joint space, and two reaching tasks for the elbow and the end-effector. For both methods, we find the optimal profiles for the weighting functions with CMA-ES. Our second result is that our controller generally performs better than GHC, even if we optimize the policies in both methods: on an average of 20 replicates, our RUF+CMA-ES finds 90% of the solutions found by our RUF+CMA-ES satisfies the constraints, while only 75% of the solutions of GHC+learning are acceptable. Furthermore, the final best solution found by RUF+CMA-ES outperforms the one of GHC+CMA-ES, as shown in Figure 28 (c). In [60] we showed on a real and simulated Jaco arm that the proposed learning method improves the movement performances in terms of fitness values, over existing tasks priorities that were manually tuned, even when the learning process starts from random initial values.

6. Conclusions

In this paper we report on the current state-of-the-art in whole-body dynamics studies concerning human movement analysis, robot control and learning. We outline a roadmap of experiments and research questions that are currently explored by the consortium of the European Project CoDyCo.

We also briefly present some of the recent CoDyCo advances in these domains. These numerous intertwined contributions cover four domains: control theory, machine learning, human behavioural experiments and software. They can be summarized as follows.

From a technological and software point of view, a whole-body control software infrastructure has been de-

signed and implemented. This infrastructure consists of several modules which significantly improve controllers accuracy and robustness thanks to a module for whole-body torques estimation, a module for force/torque sensors calibration, a module for whole-body dynamics identification and a module for dynamics estimation.

On the human side, several experimental protocols have been designed and data of experiments for studying humans in multi-contact interaction with the environment have been collected and analysed. This includes experiments on hand-contact assisted balancing for examining the functional role of supportive hand contact, a metric for whole-body stability characterisation, experiments of human-robot physical interaction to study strategies of dealing with uncertainties in contact and a study on human contact choice learning through physical interaction.

From a control point of view, state of the art control strategies for whole-body motion with multiple contacts have been implemented through the development of a solver for the whole-body reactive control that provides an expressive and rich description of the control problem as well as an efficient way of solving it. Original whole-body control reactive and predictive strategies in presence of non-rigid contacts and experiments on postural control under multiple environmental contacts while controlling the operational space dynamics have also been developed and tested in simulation. A validation scenario consisting in whole-body motion control subject to postural, contact and goal-directed (Cartesian) constraints has also been implemented.

Finally, on the learning side a theoretical framework for representing movement primitives within probabilistic contexts has been developed to improve skill learning from noisy observations. Studies on the problem of learning strategies to adapt temporal activation of low-level primitives and to deal with interferences in combining multiple

whole-body tasks have also been initiated.

Future works will be focused on the two challenges defined in introduction. More specifically, we would like to show that the iCub is able to learn how to exploit compliant contacts during whole-body tasks. Suitable control actions will be learnt (either on-the-fly or after few trials) and adapted to the compliance of the perceived contact, to be learnt on-line and a-priori unspecified.

We also aim at showing the iCub capability to successfully take advantage of the help of a naive caregiver willing to give support in performing a task otherwise impossible to perform. Proper task execution will require a certain amount of training/learning (ideally boot-strapped using the findings of the human experiments), to be performed either off-line (i.e. prior to the demonstration) or on-line (i.e. during successive trials). The naive caregiver will be given a set of instructions to leave a natural degree of autonomy. The iCub will be in charge of properly tuning its own compliance flowing the caregiver intentions. The expected validation scenario will involve a caregiver helping the iCub to stand (either from the ground or from a chair).

These works and future works on dealing with compliant and unknown environments will provide significant advances in the understanding of the use of contacts both in human motor control and humanoid robot control.

7. Acknowledgement

This work was supported by the FP7 EU project CoDyCo (No. 600716 ICT 2011.2.1 Cognitive Systems and Robotics).

8. References

- [1] P. Abbeel and A. Y. Ng. Exploration and apprenticeship learning in reinforcement learning. In *Proceedings of the International Conference on Machine Learning*, 2005.
- [2] Y. Abe, M. Da Silva, and J. Popovic. Multiobjective Control with Frictional Contacts. In *Proceedings of the Symposium on Computer Animation*, 2007.
- [3] R. Alami, R. Chatila, S. Fleury, M. Ghallab, and F. Ingrand. An Architecture for Autonomy: Introduction. *The International Journal Of Robotics Research*, pages 1–40, 1997.
- [4] N. T. Alberto, M. Mistry, and F. Stulp. Computed torque control with variable gains through gaussian process regression. In *Proceedings of the IEEE International Conference on Humanoid Robots*, 2014.
- [5] A. Arampatzis, G. P. Brüggemann, and G. M. Klapsing. Leg stiffness and mechanical energetic processes during jumping on a sprung surface. *Medicine & Science in Sports & Exercise*, 33(6):923–931, 2001.
- [6] B. D. Argall, E. L. Sauser, and A. Billard. Tactile guidance for policy refinement and reuse. In *Proceedings of the IEEE International Conference on Development and Learning*, 2010.
- [7] J. Babič, T. Petrič, L. Peternel, and N. Šarabon. Effects of supportive hand contact on reactive postural control during support perturbations. *Gait & posture*, 40(3):441–446, 2014.
- [8] K. Bouyarmane, A. Escande, F. Lamiriaux, and A. Kheddar. Potential field guide for humanoid multicontacts acyclic motion planning. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1165–1170, May 2009.
- [9] K. Bouyarmane and A. Kheddar. FEM-based static posture planning for a humanoid robot on deformable contact support. In *Proceedings of the IEEE International Conference on Humanoid Robots*, pages 487–492, 2011.
- [10] E. Burdet, R. Osu, D. W. Franklin, T. E. Milner, and M. Kawato. The central nervous system stabilizes unstable dynamics by learning optimal impedance. *Nature*, 414(6862):446–9, 2001.
- [11] R. Calandra, S. Ivaldi, M. Deisenroth, E. Rueckert, and J. Peters. Learning inverse dynamics models with contacts. In *Proceedings of the International Conference on Robotics and Automation*, 2015.
- [12] G. Cheng, S.-H. Hyon, J. Morimoto, A. Ude, G. Colvin, W. Scroggin, and S. Jacobsen. CB: A Humanoid Research Platform for Exploring NeuroScience. In *Proceedings of the IEEE/RAS International Conference on Humanoid Robots*, pages 182–187, 2006.
- [13] N. Coffey, A. J. Harrison, O. A. Donoghue, and K. Hayes. Common functional principal components analysis: A new approach to analyzing human movement data. *Human Movement Science*, 30(6):1144–1166, 2011.
- [14] C. Collette, A. Micaelli, P. Lemerle, and C. Andriot. Robust Balance Optimization Control of Humanoid Robots with Multiple non Coplanar Grasps and Frictional Contacts. In *Proceedings of the IEEE International Conference on Robotics & Automation*, pages 3187–3193, Pasadena, USA, 2008.
- [15] S. Dalibard, A. Nakhaei, F. Lamiriaux, and J.-P. Laumond. Whole-body task planning for a humanoid robot: a way to integrate collision avoidance. In *Proceedings of the IEEE/RAS International Conference on Humanoid Robots*, pages 355–360, December 2009.
- [16] C. Daniel, G. Neumann, and J. Peters. Hierarchical Relative Entropy Policy Search. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*, 2012.
- [17] A. Del Prete, S. Denei, L. Natale, F. Mastrogiovanni, F. Nori, G. Cannata, and G. Metta. Skin spatial calibration using force/torque measurements. In *Proceedings of the IEEE/RAS International Conference on Intelligent Robots and Systems*, pages 3694–3700, 2011.
- [18] A. Del Prete, F. Nori, G. Metta, and L. Natale. Prioritized motion-force control of constrained fully-actuated robots: “task space inverse dynamics”. *Robotics and Autonomous Systems*, 63:150–157, Jan 2015.
- [19] S.L. Delp, F.C. Anderson, A.S. Arnold, P. Loan, A. Habib, C.T. John, E. Guendelman, and D.G. Thelen. Opensim: open-source software to create and analyze dynamic simulations of movement. *IEEE Transactions on Biomedical Engineering*, 54(11):1940–1950, 2007.
- [20] A. Droniou, S. Ivaldi, V. Padois, and O. Sigaud. Autonomous online learning of velocity kinematics on the icub: a comparative study. In *Proceedings of the IEEE/RAS International Conference on Intelligent Robots and Systems*, 2012.
- [21] A. Escande, N. Mansard, and P.-B. Wieber. Fast Resolution of Hierarchized Inverse Kinematics with Inequality Constraints. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 3733–3738, 2010.
- [22] A. Escande, N. Mansard, and P.-B. Wieber. Hierarchical Quadratic Programming: fast online humanoid-robot motion generation. *The International Journal of Robotics Research*, 33(7):1006–1028, October 2014.
- [23] R. Featherstone and D. E. Orin. *Handbook of Robotics*, chapter Dynamics, pages 35–65. B. Siciliano and O. Khatib Eds., Springer, 2008.
- [24] P. M. Fitts. The information capacity of the human motor system in controlling the amplitude of movement. *Journal of experimental psychology*, 47(6):381–391, 1954.
- [25] A. Forner-Cordero, O. Levin, Y. Li, and S. P. Swinnen. Principal component analysis of complex multijoint coordinative movements. *Biological Cybernetics*, 93(1):63–78, 2005.
- [26] M. Fumagalli, A. Gijsberts, S. Ivaldi, L. Jamone, G. Metta, L. Natale, F. Nori, and G. Sandini. Learning to exploit prox-

- mal force sensing: a comparison approach. *From Motor Learning to Interaction Learning in Robots*, pages 149–167, 2010.
- [27] M. Fumagalli, S. Ivaldi, M. Randazzo, L. Natale, G. Metta, G. Sandini, and F. Nori. Force feedback exploiting tactile and proximal force/torque sensing. theory and implementation on the humanoid robot icub. *Autonomous Robots*, 33(4):381–398, 2012.
- [28] N. Hansen and A. Ostermeier. Completely derandomized self-adaptation in evolution strategies. *Evolutionary computation*, 9(2):159–195, 2001.
- [29] M. Hayashi, T. Sagisaka, Y. Ishizaka, T. Yoshikai, and M. Inaba. Development of functional whole-body flesh with distributed three-axis force sensors to enable close interaction by humanoids. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3610–3615, 2007.
- [30] H. J. Huang and A. A. Ahmed. Tradeoff between Stability and Maneuverability during Whole-Body Movements. *PLoS ONE*, 6(7):10, 2011.
- [31] A. Ibanez, P. Bidaud, and V. Padois. Previewed impedance adaptation to coordinate upper-limb trajectory tracking and postural balance in disturbed conditions. In *Proceedings of the 16th CLAWAR International Conference*, pages 519–528, Sydney, Australia, July 2013.
- [32] A. Ibanez, P. Bidaud, and V. Padois. Automatic optimal biped walking as a mixed-integer quadratic program. In *Proceedings of the 14th International Symposium on Advances in Robot Kinematics*, Ljubljana, Slovenia, July 2014.
- [33] A. Ibanez, P. Bidaud, and V. Padois. A distributed model predictive control approach for robust postural stability of a humanoid robot. In *Proceedings of the IEEE International Conference on Robotics and Automation*, Hong-Kong, China, June 2014.
- [34] A. Ibanez, P. Bidaud, and V. Padois. Emergence of humanoid walking behaviors from mixed-integer model predictive control. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Chicago, USA, September 2014.
- [35] A. J. Ijspeert, J. Nakanishi, and S. Schaal. Learning rhythmic movements by demonstration using nonlinear oscillators. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 958–963, September 2002.
- [36] S. Ivaldi, M. Fumagalli, M. Randazzo, F. Nori, G. Metta, and G. Sandini. Computing robot internal/external wrenches by means of inertial, tactile and F/T sensors: theory and implementation on the iCub. In *Proceedings of the IEEE/RAS International Conference on Humanoid Robots*, pages 521–528, 2011.
- [37] S. Ivaldi, J. Peters, V. Padois, and F. Nori. Tools for simulating humanoid robot dynamics: a survey based on user feedback. In *Proceedings of the IEEE/RAS International Conference on Humanoid Robots*, Madrid, Spain, Nov 2014.
- [38] A. Jain, M. Killpack, A. Edsinger, and C. Kemp. Reaching in clutter with whole-arm tactile sensing. *The International Journal Of Robotics Research*, 32(4), April 2013.
- [39] W. G. M. Janssen, H. B. J. Bussmann, and H. J. Stam. Determinants of the sit-to-stand movement: a review. *Physical Therapy*, 82(9):866–79, September 2002.
- [40] L. Johannsen, A. M Wing, and V. Hatzitaki. Contrasting effects of finger and shoulder interpersonal light touch on standing balance. *Journal of Neurophysiology*, 107(1):216–25, 2012.
- [41] O. Khatib. Real-Time Obstacle Avoidance for Manipulators and Mobile Robots. *The International Journal of Robotics Research*, 5(1):90–98, 1986.
- [42] J. Kober, J. A. Bagnell, and J. Peters. Reinforcement learning in robotics: a survey. *The International Journal Of Robotics Research*, 32(11):1238–1274, 2013.
- [43] Jens Kober and Jan Peters. Policy search for motor primitives in robotics. *Machine Learning*, 21(September):1–8, 2010.
- [44] K. Koganezawa. Mechanical stiffness control for antagonistically driven joints. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1544–1551, 2005.
- [45] V. Krishnamoorthy, M. L. Latash, J. P. Scholz, and V. M. Zatsiorsky. Muscle modes during shifts of the center of pressure by standing persons: effect of instability and additional support. *Experimental Brain Research*, 157(1):18–31, 2004.
- [46] J. Kuffner, K. Nishiwaki, S. Kagami, M. Inaba, and H. Inoue. Motion planning for humanoid robots under obstacle and dynamic balance constraints. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 692–698, 2001.
- [47] M. L. Latash, M. F. Levin, J. P. Scholz, and G. Schöner. Motor control theories and their applications. *Medicina Kaunas Lithuania*, 46(6):382, 2010.
- [48] S. M. LaValle. *Planning Algorithms*. Cambridge University Press, 2006.
- [49] C.-Y. Leung and C.-S. Chang. Strategies for posture transfer adopted by elders during sit-to-stand and stand-to-sit. *Perceptual and motor skills*, 109(3):695–706, 2009.
- [50] M. Liu, R. Lober, and V. Padois. Whole-Body Hierarchical Motion and Force Control for Humanoid Robots. *Autonomous Robots*, 2015.
- [51] M. Liu, Y. Tan, and V. Padois. Generalized hierarchical control. *Autonomous Robots*, 40(1):17–31, 2015.
- [52] Mingxing Liu and Vincent Padois. Reactive whole-body control for humanoid balancing on non-rigid unilateral contacts. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3981–3987, 2015.
- [53] R. Lober, V. Padois, and O. Sigaud. Multiple task optimization using dynamical movement primitives for whole-body reactive control. In *Proceedings of the IEEE/RAS International Conference on Humanoid Robots*, pages 193–198, Madrid, Spain, November 2014.
- [54] R. Lober, V. Padois, and O. Sigaud. Variance modulated task prioritization in whole-body control. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3944–3949, Hamburg, Germany, September 2015.
- [55] N. Mansard and F. Chaumette. Task sequencing for high-level sensor-based control. *IEEE Transactions on Robotics*, 23(1):60–72, 2007.
- [56] Metta G., Natale L., Nori F., Sandini G., Vernon D., Fadiga L., von Hofsten C., Rosander K., Santos-Victor J., Bernardino A., and Montesano L. The iCub Humanoid Robot: An Open-Systems Platform for Research in Cognitive Development. *Neural Networks*, 23(Special issue on Social Cognition: From Babies to Robot):1125–1134, 2010.
- [57] S. A. Migliore, E. A. Brown, and S. P. DeWeerth. Biologically Inspired Joint Stiffness Control. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 4508–4513, April 2005.
- [58] M. Mistry and L. Righetti. Operational Space Control of Constrained and Underactuated Systems. In *Proceedings of Robotics: Science and Systems*, Los Angeles, USA, 2011.
- [59] N. Mitsunaga, T. Miyashita, H. Ishiguro, K. Kogure, and N. Hagita. Robovie-IV: A Communication Robot Interacting with People Daily in an Office. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5066–5072. IEEE, 2006.
- [60] V. Modugno, G. Neumann, E. Rueckert, G. Oriolo, J. Peters, and S. Ivaldi. Learning soft task priorities for control of redundant robots. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2016.
- [61] U. Muico, J. Popovic, and Z. Popovic. Composite control of physically simulated characters. *ACM Transactions on Graphics*, 30(3), 2011.
- [62] T. Mukai, M. Onishi, T. Odashima, S. Hirano, and Z. Luo. Development of the Tactile Sensor System of a Human-Interactive Robot “RI-MAN”, 2008.
- [63] K. Mülling, J. Kober, and J. Peters. Learning Table Tennis with a Mixture of Motor Primitives. In *Proceedings of the*

- IEEE/RAS International Conference on Humanoid Robots, pages 411–416, Piscataway, USA, 2010.
- [64] S. Nakaoka, S. Hattori, F. Kanehiro, S. Kajita, and H. Hirukawa. Constraint-based dynamics simulator for humanoid robots with shock absorbing mechanisms. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3641–3647, 2007.
- [65] D. Nguyen-tuong and J. Peters. Using Model Knowledge for Learning Inverse Dynamics. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 2677–2682, 2010.
- [66] F. Nori, S. Traversaro, J. Eljaik, F. Romano, A. Del Prete, and D. Pucci. iCub whole-body control through force regulation on rigid noncoplanar contacts. *Frontiers in Robotics and AI*, 2(6), 2015.
- [67] D. E. Orin, A. Goswami, and S.-H. Lee. Centroidal dynamics of a humanoid robot. *Autonomous Robots*, 35(2-3):161–176, 2013.
- [68] A. Paraschos, C. Daniel, J. Peters, and G. Neumann. Probabilistic movement primitives. In *Advances in Neural Information Processing Systems (NIPS)*, Cambridge, MA: MIT Press., 2013.
- [69] A. Paraschos, C. Daniel, J. Peters, and G. Neumann. Probabilistic movement primitives. In *Proceedings of the Conference on Neural Information Processing Systems*, 2013.
- [70] A. Paraschos, G. Neumann, and J. Peters. A probabilistic approach to robot trajectory generation. In *Proceedings of the IEEE/RSJ International Conference on Humanoid Robots*, 2013.
- [71] N. Perrin, D. Lau, and V. Padois. Effective generation of dynamically balanced locomotion with multiple non-coplanar contacts. In *Proceedings of the International Symposium on Robotics Research*, September 2015.
- [72] L. Peternel and Jan Babič. Learning of compliant human-robot interaction using full-body haptic interface. *Advanced Robotics*, 27(13):1003–1012, Jun 2013.
- [73] L. Peternel, T. Petrič, and J. Babič. Human-in-the-loop approach for teaching robot assembly tasks using impedance control interface. In *Proceedings of the IEEE International Conference on Robotics and Automation*, May 2015.
- [74] L. Peternel, T. Petrič, E. Oztop, and J. Babič. Teaching robots to cooperate with humans in dynamic manipulation tasks based on multi-modal human-in-the-loop approach. *Autonomous robots*, 36(1):123–136, jan 2014.
- [75] J. Peters, M. Mistry, F. E. Udawadia, J. Nakanishi, and S. Schaal. A unifying framework for robot control with redundant dofs. *Autonomous Robots*, 1:1–12, 2008.
- [76] J. Peters and S. Schaal. Learning to control in operational space. *The International Journal Of Robotics Research*, 27:197–212, 2008.
- [77] R. Philippsen, N. Nejati, and L. Sentis. Bridging the Gap Between Semantic Planning and Continuous Control for Mobile Manipulation Using a Graph-Based World Representation. In *Proceedings of the International Workshop on Hybrid Control of Autonomous Systems*, Pasadena, USA, 2009.
- [78] G. Pratt and M. Williamson. Series elastic actuators. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 399–406, Los Alamitos, CA, USA, 1995.
- [79] E. Rueckert, M. Mindt, J. Peters, and G. Neumann. Robust policy updates for stochastic optimal control. In *Proceedings of the International Conference on Humanoid Robots*, Madrid, Spain, 2014.
- [80] C. Salaün, V. Padois, and O. Sigaud. Learning Forward Models for the Operational Space Control of Redundant Robots. In O. Sigaud and J. Peters, editors, *From Motor Learning to Interaction Learning in Robots*, chapter 2, pages 169–192. Springer Verlag, 1st edition, 2010.
- [81] J. Salini. *Dynamic control for the task/posture coordination of humanoids: toward synthesis of complex activities*. PhD thesis, Université Pierre et Marie Curie, Paris, France, June 2012.
- [82] J. Salini, V. Padois, and P. Bidaud. Synthesis of complex humanoid whole-body behavior: A focus on sequencing and tasks transitions. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1283–1290, Shanghai, China, May 2011.
- [83] K. Salisbury, W. Townsend, B. Ebrman, and D. DiPietro. Preliminary design of a whole-arm manipulation system (WAMS). In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 254–260, 1988.
- [84] T. A. Sarraf, D. S. Marigold, and S. N. Robinovitch. Maintaining standing balance by handrail grasping. *Gait & posture*, 39(1):258–64, January 2014.
- [85] S. Schaal and C. G. Atkeson. Constructive incremental learning from only local information. *Neural Computation*, 10(8):2047–2084, November 1998.
- [86] S. Schaal, C. G. Atkeson, and S. Vijayakumar. Scalable techniques for nonparametric statistics for real-time robot learning. *Applied Intelligence*, 17(1):49–60, 2002.
- [87] A. Schmitz, P. Maiolino, M. Maggiali, L. Natale, G. Cannata, and G. Metta. Methods and Technologies for the Implementation of Large Scale Robot Tactile Sensors. *IEEE Transactions on Robotics*, 27(3), 2011.
- [88] J. P. Scholz and G. Schöner. The uncontrolled manifold concept: identifying control variables for a functional task. *Experimental Brain Research*, 126(3):289–306, 1999.
- [89] L. Sentis, J. Park, and O. Khatib. Compliant Control of Multicontact and Center-of-Mass Behaviors in Humanoid Robots. *IEEE Transactions on Robotics*, 26(3):483–501, 2010.
- [90] R. Shadmehr and F. A. Mussa-Ivaldi. Adaptive representation of dynamics during learning of a motor task. *Journal of Neuroscience*, 14(5/2):3208–3224, 1994.
- [91] R. Shadmehr, J.-J. Orban de Xivry, and W. Xu. Temporal discounting of reward and the cost of time in motor control. *The Journal of neuroscience*, 30(31):10507–10516, 2010.
- [92] H. Slijper and M. Latash. The effects of instability and additional hand support on anticipatory postural adjustments in leg, trunk, and arm muscles during standing. *Experimental Brain Research*, 135(1):81–93, 2000.
- [93] F. Stulp and S. Schaal. Hierarchical reinforcement learning with movement primitives. In *Proceedings of the IEEE/RAS International Conference on Humanoid Robots*, pages 231–238, 2011.
- [94] G. Tonietti, R. Schiavi, and A. Bicchi. 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*, pages 526–531, April 2005.
- [95] N. G. Tsagarakis, Z. Li, J. Saglia, and D. G. Caldwell. The design of the lower body of the compliant humanoid robot “ccub”. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2011.
- [96] N. G. Tsagarakis, B. Vanderborght, M. Laffranchi, and D. G. Caldwell. The mechanical design of the new lower body for the child humanoid robot “iCub”. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4962–4968, St. Louis, USA, 2009.
- [97] M. Vukobratović and B. Borovac. Zero-moment point thirty five years of its life. *International Journal of Humanoid Robotics*, 1(01):157–173, 2004.
- [98] D. A. Winter. Human balance and posture control during standing and walking. *Gait & Posture*, 3(4):193–214, December 1995.
- [99] K. Yamane and Y. Nakamura. Parallel $O(\log N)$ Algorithm for Dynamics Simulation of Humanoid Robots. In *Proceedings of the IEEE/RAS International Conference on Humanoid Robots*, pages 554–559, 2006.
- [100] E. Yoshida, I. Belousov, C. Esteves, and J.-P. Laumond. Humanoid Motion Planning for Dynamic Tasks. In *Proceedings of the IEEE/RAS International Conference on Humanoid Robots*, Tsukuba, Japan, 2005.