Adaptive synthesis of dynamically feasible full-body movements for the humanoid robot HRP-2 by flexible combination of learned dynamic movement primitives

Albert Mukovskiy, Christian Vassallo, Maximilien Naveau, Olivier Stasse, Philippe Souères, Martin Giese

HAL Id: hal-01484935
https://hal.archives-ouvertes.fr/hal-01484935

Submitted on 8 Mar 2017

Adaptive synthesis of dynamically feasible full-body movements for the humanoid robot HRP-2 by flexible combination of learned dynamic movement primitives

Albert Mukovskiy[a], Christian Vassallo[b], Maximilien Naveau[b], Olivier Stasse[b], Philippe Souères[b], Martin A. Giese[a]

[a]*Section for Computational Sensomotorics, Department of Cognitive Neurology, Hertie Institute for Clinical Brain Research & Centre for Integrative Neuroscience, University Clinic, Tübingen, Germany*
[b]*Gepetto Lab, LAAS/CNRS, Université de Toulouse, Av. du Colonel Roche 7, F-31400, Toulouse, France*

Abstract

Skilled human full-body movements are often planned in a highly predictive manner. For example, during walking while reaching towards a goal object results in steps and body postures are adapted to the goal position already multiple steps before the goal contact. The realization of such highly predictive behaviors for humanoid robots is a challenge because standard approaches, such as optimal control, result in computation times that are prohibitive for the predictive control of complex coordinated full-body movements over multiple steps. We devised a new architecture that combines the online-planning of complex coordinated full-body movements, based on the flexible combination of learned dynamic movement primitives, with a Walking Pattern Generator (WPG), based on Model Predictive Control (MPC), which generates dynamically feasible locomotion of the humanoid robot HRP-2. A dynamic filter corrects the Zero Moment Point (ZMP) trajectories in order to guarantee the dynamic feasibility of the executed behavior taking into account the upper-body movements, at the same time ensuring an accurate approximation of the planned motion trajectories. We demonstrate the high flexibility of the chosen movement planning approach, and the accuracy and feasibility of the generated motion. In addition, we show that a naïve approach, which generates adaptive motion by using machine learning methods by the interpolation between feasible training motion examples fails to guarantee the stability and dynamic feasibility of the generated behaviors.

Keywords:
robotics, navigation, walking pattern generator, goal-directed movements, movement primitives, motor coordination, action sequences

1. Introduction

The modeling and the synthesis of the online-reactive multi-action sequences is an extremely important topic in both, computer graphics and humanoid robotics. The most challenging problem in the online control of complex whole-body behaviors, which is solved apparently effortless by humans, is the flexible coordination of goal-directed movements with the maintenance of dynamic balance during locomotion. The solution of this problem requires simultaneously the flexible adaptation of executed upper-body behaviors, e.g. to changing positions of goal objects or obstacles, combined with a control of dynamic balance during walking, in order to avoid that the robot falls. In addition, the robot's joint torques have to be kept in a feasible range. The detailed analysis of human behavior shows that their motor control is highly predictive, and often optimizes complex behaviors over long time horizons, e.g. lasting over multiple steps. This is fundamentally different from many solutions of this problem in humanoid robotics [Siciliano and Khatib (2016)].

The realization of such complex behaviors for humanoid robots with long time horizons for predictive control is a challenging problem. A standard method for the computation of control signals for such high-

_____

Email addresses: `Corresponding author:`
`albert.mukovskiy@medizin.uni-tuebingen.de` (Albert Mukovskiy), `christian.vassallo@laas.fr` (Christian Vassallo), `maximilien.naveau@laas.fr` (Maximilien Naveau), `ostasse@laas.fr` (Olivier Stasse), `philippe.soueres@laas.fr` (Philippe Souères), `martin.giese@uni-tuebingen.de` (Martin A. Giese)

dimensional robots is optimal control. However, the solution of optimal control problems over such long temporal horizons is computationally extremely costly. With state-of-the-art methods [Koschorreck and Mombaur (2012)] even the optimization of a single step for a humanoid robot can take minutes, and multi-step sequences would take even multiple hours. This renders such methods inappropriate for the real-time planning and control of such complex full-body motion sequences in humanoid robots.

In this paper, we present an approach for the solution of this problem that combines two components. The first component is an online-capable planning algorithm that is based on learned dynamic movement primitives, which generates human-like full-body motion sequences that flexibly adapt to changes in the action space, e.g. displacements of the goal object. The second component is a nonlinear Model Predictive Control (MPC) system [Naveau et al. (2016, in press)] for the humanoid robot that combines the outputs of a Walking Pattern Generator (WPG) with the panned upper-body motion in a way that guarantees the the dynamic feasibility of the resulting full-body motion. An essential element of this architecture is a *dynamic filter* that appropriately modifies the planned Center of Mass (CoM) and Zero Moment Point (ZMP) trajectories in dependence of the planned upper-body motion. We demonstrate the novel approach by the control of multi-step sequences that realize highly adaptive reaching and walking towards goal objects at different distances, where the system implements human-like highly predictive control over multiple-steps. The resulting computational complexity of this control system is not much higher than the one of the state of the art WPG algorithm [Naveau et al. (2014)].

Our article is organized as follows: In the following section, we give an overview of related work in the areas of computer graphics and humanoid robotics. Section 3 describes the developed system. This includes a short description of the underlying human trajectory data from a drawer-opening task, of the online planning algorithm that is based on a special form of dynamic movement primitives, and a more detailed description of the integration of this planning algorithm with the MPC-based control architecture of the robot. In Section 4 we present a variety of results obtained with the OpenHRP robot simulator, which evaluate the proposed approach also in comparison with simpler solutions, and a preliminary test that implemented the planned trajectories on the HRP-2 humanoid robot. Limitations and extensions of the developed approach are discussed in the final section.

2. Related Work

Work related to the developed system can be found in biological motor control and related robotics approaches, computer animation, and in humanoid robotics.

2.1. *Control of multi-step sequences in biological systems*

Biological systems effortlessly coordinate locomotion with other goal-directed tasks [Weigelt and Schack (2010)]. A relevant example are studies on the coordination of walking and reaching. The kinematics of this behavior can be approximated by two separate underlying movement components, which mainly model the periodic locomotion and the non-periodic goal-directed movement [Chiovetto and Giese (2013)]. A recent study [Land et al. (2013)] investigated in detail the underlying coordination, using a task where participants had to walk towards a drawer and to reach for an object. Participants showed highly predictive control in their motor behavior, where within multi-step sequences already the first step was adapted dependent on the position of the goal object. In addition, participants adjusted their behavior in a way that ensured comfortable reaching during in the last step. This behavior is compatible with the *maximum end-state comfort* hypothesis that has been formulated in human motor control [Rosenbaum (2008); Weigelt and Schack (2010)]. In recent work we have tried to reproduce this behavior by an algorithm for trajectory synthesis that is based on learned dynamic movement primitives [Mukovskiy et al. (2015)]. A similar problem has also been solved by [Gienger et al. (2010)], who computed optimized stance locations with respect to the position of a reaching target, using a dynamical systems approach for the generation of reaching behavior.

An influential idea in the field of biological motor control has been the concept of movement primitives [Flash and Sejnowski (2001); Flash and Hochner (2005)]. According to this hypothesis the coordination of complex movements is based on the combination of lower-dimensional control units, strongly reducing the dimensionality of the underlying control problem. Such primitives have been extracted by unsupervised learning from kinematic and EMG data. This idea has been transferred to robotics. [Taïx et al. (2013)] extracted primitives from human reaching movements using principle component analysis (PCA), successfully implementing reaching behavior on an HRP-2 humanoid robot. Movement primitives, including the use of force feedback, have also been proposed by [Gams

et al. (2009, 2013)]. A related important idea is the concept of dynamic movement primitives that generates planned trajectories by appropriately designed nonlinear dynamical systems [Schaal et al. (2003); Ijspeert et al. (2003)]. Systems based on dynamic movement primitives have been proposed for the generation of complex movements in real-time [Ijspeert et al. (2013), Ajallooeian et al. (2013)]. But all these online DMP-based methods of modeling the kinematic trajectories do not guarantee the dynamic feasibility of the resulting motion, which is a critical issue.

2.2. Modeling of whole-body movement sequences in computer graphics

The problems of kinematic synthesis of complex whole body movements has been addressed extensively in computer graphics, e.g. [Levine et al. (2012)], and many learning-based approaches have been proposed that provide low-dimensional parameterizations of classes of whole body motion [e.g. Hsu et al. (2005); Safonova et al. (2004); Wang et al. (2008); Li et al. (2002)]. The generated individual movements can be automatically concatenated into longer sequences, taking into account additional task constraints [Kovar et al. (2002)]. A relevant example is [Huang and Kallmann (2014)] who modeled the coordination between loco-motion and arm pointing in the final step, by blending and selecting arm pointing primitives dependent on the actual gait phase. All these methods model the movement kinematics without taking dynamic constraints into account. A recent example is [Feng et al. (2012)] who blended motion-captured example motion prioritized 'stack of controllers'. [Shoulson et al. (2014)] presented a method where controllers for different body parts are blended, where their prioritization is changed sequentially over time, dependent on the actual action within a longer sequence.

Other work in this domain has developed dynamic filtering techniques in order to adjust such synthesized motion to fulfill dynamic constraints derived from physical models, e.g. for the Zero Moment Point (ZMP), in order to increase the generalization regime of such learning-based methods [Liu et al. (2005)].

2.3. Related approaches in humanoid robotics

In humanoid robotics numerous approaches have been proposed for the synthesis of walking in combination with grasping movements. An example is the DARPA robotic challenge valve manipulation task. For this problem, [Ajoudani et al. (2014)] proposed a hybrid controller with a goal-driven fast foot step planner

that is combined with visual servoing for the reaching and grasping of the valve. [Kuindersma et al. (2015)] proposed a control architecture for the humanoid robot Atlas that automatically finds foot steps around and over obstacles, in order to reach for a goal object and to realize more complex actions. Other solutions for the combination of walking and vision-controlled reaching of a static and mobile targets during walking have been proposed in [Stasse et al. (2008)] and [Brandao et al. (2013)].

Some researchers have used randomized motion planning algorithms for whole-body walking combined with manipulation tasks in constrained environments [Dalibard et al. (2013)]. For example, [Kanoun et al. (2011)] proposed a method that is based on a virtual kinematic tree for the planning of foot placements, which was successfully implemented on the HRP-2 robot. A framework that decomposes reach-to-grasp human movements into sequences of kinematic tasks has been developed in [Sreenivasa et al. (2012)]. Further work applied imitation learning [Mühlig et al. (2010)], where walking and grasping were modeled as a sequence of separate actions. A task priority approach based on a generalized inverse kinematics was applied in [Yoshida et al. (2007)] in order to organize several sub-tasks, including stepping and hand motion.

The control of human-like multi-joint systems taking into account contact constraints and guaranteeing dynamic balance is a challenging approach. Current solutions range from near real-time whole body Model Predictive Control with regularized modeling of contacts in order to decrease the associated computational cost [Tassa et al. (2012); Koenemann et al. (2015)] to approaches based on optimal control with precise modeling of contact phases, requiring typically hours of off-line computation time (e.g. [Koschorreck and Mombaur (2012)]).

A solution based on prioritized IK, that integrates DMPs with MPC for individual actions has been proposed by [Vuga et al. (2013)].

3. System architecture

In the following we first give a brief overview of the human data that was used for the training of the primitives of our online planning algorithm, and which also provides evidence of the highly predictive coordination of complex human goal-directed movements. Subsequently, we describe briefly our movement primitive-based online motion planning algorithm and discuss how this planning algorithm can be integrated with the
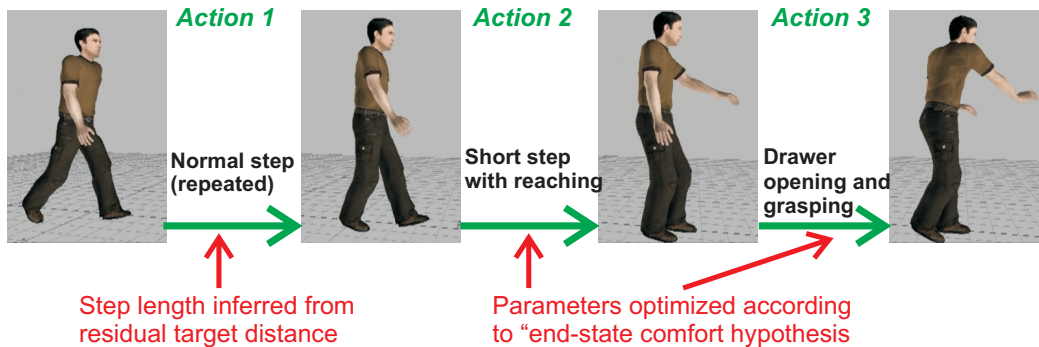
Figure 1: Illustration of the different phases of the human behavior, illustrated in terms of intermediate postures: 1) normal step; 2) step with initiation of reaching; and 3) standing with opening of the drawer and reaching for the object.

model predictive control architecture for the humanoid robot HRP-2.

3.1. Human data

3.1.1. Drawer opening task

The modeling of the coordination of walking and reaching was based on a motion capture data set from humans who opened a drawer. The participants walked towards a drawer, opened it with their left hand and reached for an object inside it with right hand. The initial distance from the drawer and the position of the object inside it were varied [Mukovskiy et al. (2015)]; see Fig. 1. The recorded sequences consist of three subsequent actions: 1) a normal walking step (starting with left heel strike and ending with left heel strike); 2) a shortened step with the left-hand reaching towards the drawer. This step showed a high degree of adaptability, and its length was typically adjusted in order to create an optimum distance from the drawer for the final reaching movement. This behavior is consistent with the *maximum end-state comfort hypothesis* in motor control, which assumes that motor planning optimizes the comfort of the end state of planned movements [Land et al. (2013)]; 3) the drawer opening combined with the reaching for the object while standing. An example trial is shown in [**movie**[1]].

The data set consists of the trajectories of ten trials of single participant, recorded in Univ. Bielefeld with optical motion capture system (Vicon Motion Systems, Oxford, UK) consisting of 12 MX-F20 CCD cameras at a frame rate of 200 fps with a spatial accuracy of about 1.5 mm. PluginGait marker set was used with 41 markers. The length of the individual steps (actions) for the

individual motion-captured sequences is shown in Fig. 2. This figure very nicely illustrates the predictive nature of the motor planning. The size of the second step (yellow) is strongly adapted dependent on the distance of the starting point from the goal position. The lengths of the other steps is much less variable and shows systematic dependence on this distance. (See [Mukovskiy et al. (2015)] for further details about this data set.)
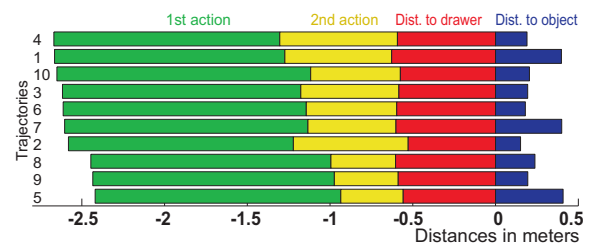

Figure 2: Predictive planning of real human trajectories. Distances from the pelvis to the front panel of the drawer (green, yellow, red), and the distance between the front panel and the object (blue) for different trials. (Reproduced from [Mukovskiy et al. (2015)].)

3.2. Preprocessing of training trajectories

The recorded motion capture data was processed and animated in MotionBuilder (Autodesk), using an 'actor' puppet whose geometric parameters were adapted to the recorded subject. The trajectories were cut, starting at the first heel strike and ending with the object reaching. A kinematic model of the HRP-2 robot was created in Maya (Autodesk), neglecting joint angle constraints. All ten trajectories were retargeted to the HRP-2 model using MotionBuilder. During retargeting the feet positions of HRP-2 were constrained to level ground, and the step sizes were reduced proportionally to the height
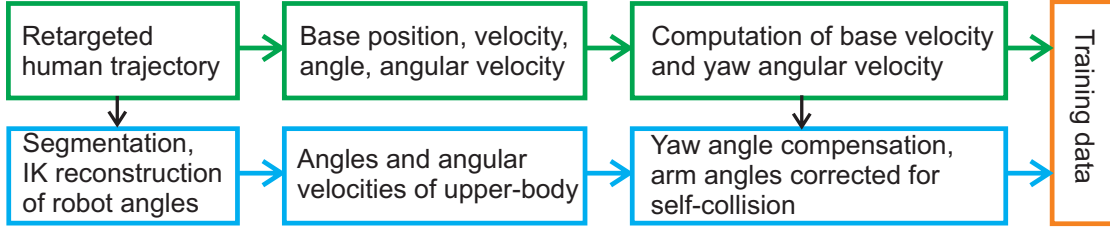
---

[1] https://goo.gl/5HKiG7

Figure 3: Offline pre-processing of motion capture data from humans.

of the robot. The resulting joint frame trajectories were exported, using the Denavit-Hartenberg (DH) convention. Trajectories were then segmented by hand, and the durations each action and the corresponding step sizes were stored separately.

The computed trajectories were further analyzed in Matlab (MathWorks) and resampled, resulting in a normalized duration of 1.6 sec for each action. The data was split into two subsets, separating the stored pelvis trajectories (time course of pelvis position and pelvis direction in the horizontal plane), and the upper body trajectories (HRP-2 joint angles extracted from DH representation). The pelvis position trajectories were rescaled, ensuring the maximally admissible propagation velocity for the HRP-2 (0.5 m/sec). The pelvis yaw-angle trajectories were rescaled by a constant factor, and a fraction of the yaw angle trajectory was added back to the trunk yaw-angle for compensation. After this compensation, customized inverse kinematics (IK) methods were applied to correct the upper body arm reaching motion in order to satisfy joint limit constraints. As input to the Walking Pattern Generator (WPG) (see below) we used the time courses of pelvis velocities in the horizontal plane, and of the pelvis yaw angular velocity. An overview of the pre-processing steps is given in Fig. 3.

An illustration of this preprocessing is given in [**movie**[2]], which shows the angular trajectories, animating a human avatar, and the corresponding retargeted trajectories for a HRP-2 kinematic model in Motion-Builder.

3.3. Primitive-based online motion planner

The first core component of our architecture is an online motion planning algorithm that is based dynamic movement primitives, which are derived from classes of human trajectories by unsupervised learning. Because of space limitations we can here only briefly sketch the

structure of the online planning algorithm. Further details can be found in [Mukovskiy et al. (2015)].

3.3.1. Learning of kinematic primitives

The first step in the construction of dynamic movement primitives is the learning of a low-dimensional representation of a set of motion-captured trajectories that spans the space of relevant behaviors, by a superposition of a small number of source or basis functions. By this form of dimension reduction the relevant behaviors can be generated by a very small number of coupled dynamic movement primitives (s.b.). Contrasting with many related approaches for the modeling of trajectories, which exploit for example PCA or ICA (e.g. [Safonova et al. (2004)]), we fit such trajectory sets by a generative model that is known in acoustics as *anechoic mixing model*. Opposed to the instantaneous mixing model that underlies PCA and ICA models for trajectory representation, the anechoic model allows for time shifts of the superposed components. We have shown elsewhere [Omlor and Giese (2011)] that the anechoic mixing model for many types of movements result in representations with a much smaller number of source functions (typically by factor two), for equal approximation quality in comparison with standard methods such as PCA or ICA. This low dimensionality is essential since it determines the dimensionality of the state space of the nonlinear dynamical system that generates behaviors adaptively online.

Mathematically, the *anechoic mixture model* is defined by the equation:

$$\underbrace{\xi_i(t)}_{angles} = m_i + \sum_j w_{ij} \underbrace{\sigma_j\left(t - \tau_{ij}\right)}_{sources}$$

The joint angle trajectories $\xi_i(t)$ were derived from the original motion capture data that is temporally segmented into the three subsequent actions (s.a.). The normalized action trajectories are approximated by a linear mixture of the source signals $\sigma_j(t)$, weighted with the mixing weights w_{ij}. The individual sources are shifted
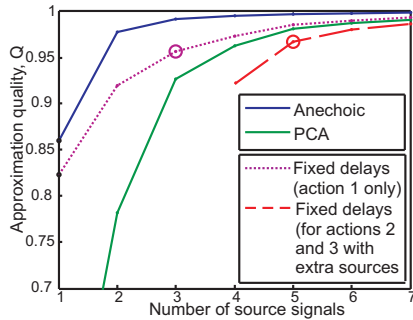
Figure 4: Approximation quality as function of the number of sources for all three actions, comparing anechoic demixing without constraints (*blue*) and PCA (*green*). The *purple dotted* line indicates approximation quality for the first action with fixed delays across trials. The *red dashed* line indicates approximation quality with two additional sources (with fixed delays).

Figure 5: Estimated source functions.

in time with the delays τ_{ij}. The means of the angle trajectories are indicated by the variables m_i. Source functions and model parameters were learned, applying anechoic demixing algorithms described in [Omlor and Giese (2011); Chiovetto et al. (2013)]. For the application presented in this paper, in addition to the learned source functions, which approximate the periodic signals components, we used an additional nonperiodic source component, which was pre-specified. This component was given by the fixed function $s_0(t) = \cos(\pi t/T)$, where T was the cycle time of the learned periodic source function with the lowest fundamental frequency.

In order to model the multi-step sequences we learned such a representation using a step-wise regression approach. The whole training data was first used to fit the mean values m_i and the weights corresponding to the non-periodic source function. The residuals of action 1, the normal walking step, were then approximated by three periodic source functions, applying a modified algorithm that constrains all time delays for the same source function to be equal across all trials (but allowing different delays for different joint angles). This constraint simplifies the spatio-temporal blending between different motion patterns, at the cost that more sources have to be introduced for an accurate approximation (cf. Fig. 4). The second and third action then were approximated using the sources introduced for the normal walking step, and two additional periodic sources that were added in order to account for the residuals, where the same constraint was applied to the estimated time delays. Fig. 4 shows the obtained approximation quality in
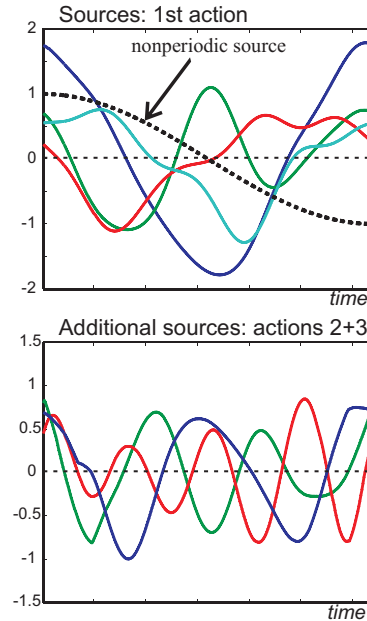
comparison with a PCA model, and with models without constraints for the time delays. It is possible to accomplish a quite accurate approximation with a total of four sources for the first action and two additional periodic sources for actions 2 and 3. And the resulting shapes of the learned source functions for the approximation scheme of 4+3 sources are shown in Fig. 5.

3.3.2. Online planning of multi-action sequences

Our online planning algorithm for whole-body movements generates the trajectories as solutions of nonlinear dynamical systems that is based on dynamic movement primitives, which are derived from the kinematic primitives described in section 3.3.1. Dynamic movement primitives have been proposed in robotics before [Schaal et al. (2003); Ijspeert et al. (2003)], and similar approaches have been described in [Gams et al. (2009); Petrič et al. (2011); Buchli et al. (2006)]. These previous approaches, however, exploit no dimension reduction for the learning of the kinematic primitives. We have previously demonstrated the suitability of our approach for the adaptive online generation of complex multi-step sequences, coordinated with arm movements, and for the animation of coordinated crowds of agents [Mukovskiy et al. (2015); Giese et al. (2009)].

We constructed dynamic movement primitives (DMPs) from the kinematic primitives described in the previous subsection. For this purpose, we mapped the
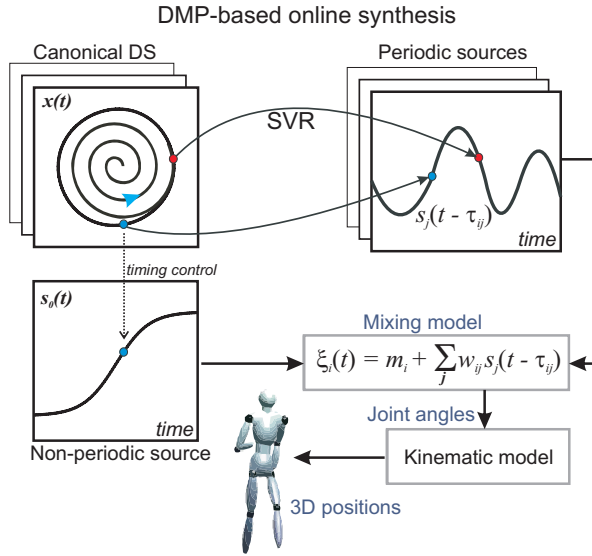
Figure 6: Architecture for the online synthesis of body movements using dynamic primitives, [Giese et al. (2009)].

state space of simple nonlinear canonical dynamical system onto the values of the learned source functions. These nonlinear mappings were learned using Support Vector Regression (using a Radial Basis Function kernel, exploiting the LIBSVM Matlab$^{®}$ library [Chang and Lin (2001)]). In this way the source signals can be generated online as solutions of a nonlinear dynamical system. The canonical dynamical systems related to different primitives were dynamically coupled in order to ensure a synchronization of the corresponding states. The resulting architecture is summarized in Fig. 6. A more detailed discussion of the design of this coupling and its relationship to the stability of the resulting dynamics is given in [Mukovskiy et al. (2013)]. The online generated source signals are then used as inputs for the anechoic mixing model, which defines the planned joint angle trajectories. An overview of the underlying architecture is shown in Fig. 6.

As canonical dynamics for the periodic DMPs we chose a limit cycle oscillator (Andronov-Hopf oscillator), which is given by the equations (ω defining the eigenfrequency), and the pair of state variables $[x(t), y(t)]$:

$$\dot{x}(t) = [1 - (x^2(t) + y^2(t))]x(t) - \omega y(t)$$
$$\dot{y}(t) = [1 - (x^2(t) + y^2(t))]y(t) + \omega x(t)$$

Since the attractors of this nonlinear systems can be mapped onto circle in the phase plane, delays can be represented by a rotation of the vectors in state space by

an angle that is proportional to the delay. In this way, we are able to model coupled networks with delays between different CPGs by a set of coupled set of differential equations without explicit delay times (see [Giese et al. (2009)] for further details). The instantaneous phase of the leading DMP, which generates the periodic solution with the lowest frequency also was used to control the timing of the non-periodic source.

In order to plan online highly flexible behaviors, with an adaptation of steps and reaching behavior to the goal position, we learned nonlinear mappings from task parameters onto the mixing weights of the anechoic mixing model. The task parameters were the were steps lengths and durations. Mappings were learned from training data, applying Locally Weighted Linear Regression (LWLR) [Atkeson et al. (1997); Mukovskiy et al. (2015)].

For the synthesis of multi-step step sequences the step lengths was computed from the actual estimated target distance. For this purpose we tried to reproduce the dependencies between the individual step sizes and the distance to the goal. In the human data, the reaching distance of the arm (action 3) is positively correlated with the distance to the object inside the box, and negatively with the length of the previous step. These dependencies were modeled by linear regression and exploited for the computation of the reaching distances while standing. For the second step, the step length was adjusted in order to realize a maximum-comfort distance for reaching. The length of the other steps then was adjusted accordingly. Step ranges were computed from the training data, and an appropriate number of additional steps was automatically introduced when the target could not be reached within three steps. A more detailed description of the algorithms for the smooth interpolation of the mixing weights, ensuring smooth transitions between the different steps is given in [Mukovskiy et al. (2015)].

Figure 7 illustrates the high degree of flexibility of this online planning algorithm. For this example, the goal (drawer) jumps away from its original position while the agent is approaching it. The algorithm adapts online to this perturbation, generating very human-like adaptive coordination, and it includes automatically an additional step in order to compensate for the suddenly increased distance to the drawer. See also **movie**[3].

3.4. Integration of online-planning with model-predictive control of the HRP-2

The central innovation of our work is the integration of the described online-planning algorithm with a con-
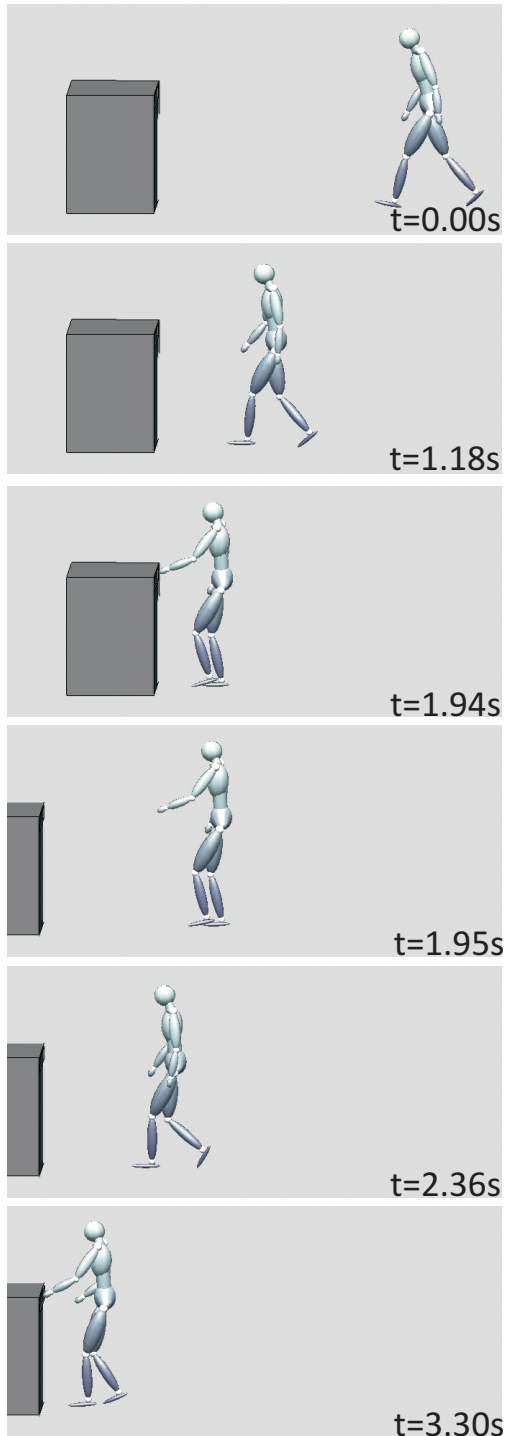
---

[3]https://goo.gl/9fLzO7

Figure 7: Adaptivity of online planning. If the goal (drawer) jumps away during the approaching behavior, automatically an additional step is introduced. Overall, a very smooth human-like whole-body coordination pattern emerges.

trol architecture for the HRP-2 humanoid robot, which is based on nonlinear model predictive control (NMPC). This does not only involve the combination of trajectories derived from human data, as described in section 3.2, but it requires specifically the approximation of human data by dynamically feasible trajectories, exploiting NMPC framework. These feasible trajectories form a novel training set, from which a new set of optimized dynamic primitives was derived.

3.4.1. Overview of control architecture

The control architecture for the HRP-2 robot is shown in Fig. 8. It consists of three main building blocks. The online kinematic synthesis algorithm, which was laid out in section 3.3, provides input to the control architecture (shaded box in Fig. 8) in terms of two sets of variables: the velocity and angular velocity of the Center of Mass (CoM) (variables $\mathbf{v}^{\text{ref}}$ and ω^{ref}), and the joint angles of the upper body $\mathbf{q}^{\text{upper body}}$.

The first building block is a Walking Pattern Generator (WPG) that computes from the variables $\mathbf{v}^{\text{ref}}$ and ω^{ref}, for one gait cycle, foot placements $\mathbf{x}_{\text{feet}}$ and trajectories of the CoM $\mathbf{x}_{\text{CoM}}$ and of the Zero Moment Point (ZMP) $\mathbf{x}_{\text{ZMP}}$ that ensure the dynamic stability of the gait [Vukobratović and Stepanenko (1972)]. This computation is based on model predictive control (MPC), and further details about the underlying computations can be found in section 3.4.2 and in [Naveau et al. (2016, in press)].

The second block is a Dynamic Filter (DF) that corrects the preplanned foot, CoM, and ZMP trajectories, taking into account the planned upper-body motion, resulting in the corrected trajectories $\mathbf{x}_{\text{feet}}^{\text{cor}}$, $\mathbf{x}_{\text{CoM}}^{\text{cor}}$ and $\mathbf{x}_{\text{ZMP}}^{\text{cor}}$. The DF operates in closed-loop together with WPG, and further details about the underlying algorithms are described in [Naveau et al. (2016, in press); Stasse (2013)].

The third building block is generalized inverse kinematics (IK) module that implements a 'Stack-of-Task' approach. This module combines the corrected CoM and ZMP trajectories, and the upper-body motion (specified by the joint angles). This module outputs joint angle trajectories for the legs and the upper-body that respect the dynamic stability constraints of the robot, at the same time approximating, as far as possible, the planned behavior of the upper body. For this purpose the task of stabilizing the locomotion is given the highest priority, and the approximation of the planned trajectories is realized in the null-space of the control signals for this prioritized task. The resulting optimization problem is solved by a sequential quadratic programming approach (QP solver).
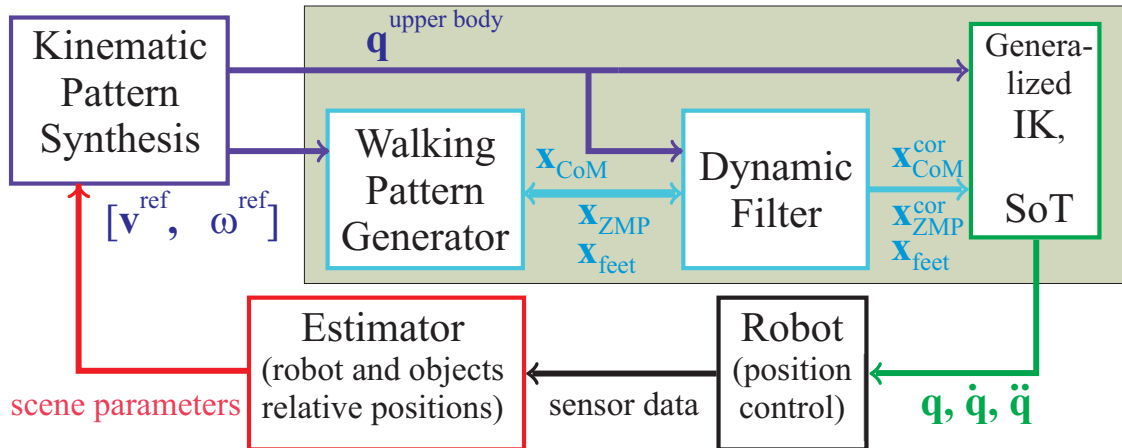
Figure 8: Control system for the humanoid robot HRP-2. The Walking Pattern Generator computes foot positions and CoM and ZMP trajectories, which are further adjusted by the Dynamic Filter, dependent on the planned upper body motion. The resulting trajectories are consistent with the dynamic stability constraints of the robot. The approximation of the planned upper body movement and dynamic stability of walking are guaranteed by a Stack of Task approach, where optimal trajectories are computed by sequential quadratic programming. (See text for further details.)

The resulting optimal trajectories $\mathbf{q}$ are dynamically feasible and can be realized by the low-level controllers of the HRP-2 robot. During motion execution, the real-world environmental and task parameters and the current state of the robot are fed back to the kinematic planner, closing the control loop for an adaptive interaction between online planning and MPC in the real world.

3.4.2. WPG based on optimal predictive control

The Walking Pattern Generators (WPG) based on Model Predictive Control (MPC). The first WPG of this class was proposed by [Kajita et al. (2003)]. This method computed the reference nominal Zero Moment Point (ZMP) trajectory from the desired placements of feet during the gait cycle. A simplified linear inverted pendulum dynamics ('Cart-Table Model') was used to link the Center of Mass (CoM) and the ZMP. Preview control was exploited for computing the CoM trajectory from the desired ZMP. Due to the model simplifications, the real ZMP trajectory deviates from the desired one. This deviation is the result of neglecting the inertial and Coriolis forces generated by the leg swing and by fast movements of the upper-body. In order to alleviate this problem, the authors ran the full body inverse dynamics in order to compute a better approximation of the real ZMP. This new ZMP can be computed for the preview horizon in real-time. The resulting ZMP error was transformed into a resulting CoM error via the Preview Control, following the approach proposed by [Kajita et al. (2003)]. This result can then be ex-

ploited to correct the CoM trajectory. The described two steps of preview control combined with an evaluation of the inverse dynamics can be repeated iteratively, successively reducing the ZMP error. This approach for the dynamic correction can be interpreted as a kind of Newton-Raphson iteration [Stasse (2013)], and was referred to as Dynamic Filter in section 3.4.1.

Another improvement of MPC-based WPG is the integration of the computation of the optimal ZMP trajectory within the constrained quadratic optimization framework that computes the optimal CoM trajectory [Wieber (2006)]. This approach requires only the specification of the preplanned foot positions as input, returning the optimal trajectories for the ZMP and the CoM. Our approach for nonlinear MPC relies in addition on another improvement of the same framework made by [Herdt et al. (2010)], which is the further extension of the approach by [Wieber (2006)]. This reformulation of the optimization framework allows to exploit positional and angular velocities of the CoM as reference trajectories (for a time horizon of the next two steps), returning the foot placements and the optimal ZMP trajectories as result of the nonlinear predictive control problem. This framework (which is described in detail in [Naveau et al. (2014, 2016, in press)]) was exploited in our system.

3.4.3. Generation of the dynamically feasible training data

In order to link the described approach for the online synthesis of movements with the NMPC approach de-

9

scribed above, we transform a set of human-compatible movement trajectories that were generated by interpolation from the original human data into trajectories that result in dynamically feasible behavior of the robot. For this purpose, we approximated the human-like trajectories by ones generated by physics-based simulations, exploiting the NMPC framework discussed in section 3.4.1. This training of our learning-based approach using dynamically feasible training data is one of the key concepts of our proposed approach. The details of retargeting and transformation in dynamically feasible trajectories of the training data are discussed in section 4.3.

4. Results

In the following we first briefly discuss the computation time of our approach. Then we present some results on the online kinematic planning algorithm. We then present results of the performance of the method in the off-line mode, where it was used to reproduce the behaviors of retargeted training trajectories without adaptation to new step sizes or goal distances. We demonstrate that the obtained behaviors indeed are dynamically feasible and can be implemented on the real HRP-2 robot. We then demonstrate the performance of the system in case when adaptive behavior is planned dependent on the actual goal positions. In the last section we compare the robustness of the proposed method that integrates MPC with learning-based online planning with a simpler machine learning-based approach, which realizes control by interpolation between learned whole-body angle trajectories which have been derived from training examples that were dynamically feasible. It turns out that such a more naïve machine learning approach in many cases results in instability and infeasibility of the produced behavior.

4.1. Computation time

The kinematic pattern synthesis algorithm has a computation time around 81-86 ms for the whole trajectory (1280 time steps, each 5 ms) on modern a CPU (Intel(R) Xeon(R) CPU E3-1241 v3, 3.50GHz, Ubuntu 14.04). The kinematic synthesis required only when target goal changes its position, and this computation time is below the buffer size for the preview control (100 ms). The average computation time of the optimization problem involving WPG-DF iterations is 4 ms (see Naveau et al. (2016, in press)), which is below the duration of control time-step (5 ms). The whole algorithm is thus realtime-capable. In contrast, the optimal control approaches typically require several hours for the off-line synthesis of multi-step sequences.

4.2. Primitive-based synthesis of kinematic trajectories

In order to validate the primitive-based online planning algorithm we generated a set of highly human-like novel full-body movement behaviors, varying the initial distance of the actor from the goal object, including a spectrum of distances that were not part of the training set. In order to judge the human-likeness of the interpolated behaviors, we did not retarget the movements to the robot kinematics and illustrate them as movies, using a human avatar. We learned 3 sources for the approximation of the first step, and another two extra sources for approximating the residuals of the second and third step ("3+2 sources"). The training data set consisted of 10 human joint-angle trajectories.

From the model trained with this data new trajectories were generated by interpolation, and the agent's propagation velocity and rotation of the base (pelvis) were computed from the feet-ground contact events. The step distances from this simulation served as task parameter for the Linearly Weighted Regression (LWLR).

For all tested novel goal distances the algorithm generated very human-like highly coordinated three-step sequences. This is illustrated in the **movie**[4] that shows step sequences for total goal distances between 2.34 and 2.94 m, all of which were not in the training set. Behaviors for goal distances above 3 meters are shown in **movie**[5], where the algorithm introduces automatically additional gait steps in order to ensure that the agents reaches the goal. Also these behaviors look amazingly human-like.

The capability of online replanning is demonstrated in **movie**[6] (see also Fig. 7). I this case, the goal jumps away from the agent during the approaching behavior towards a more distant position, where it cannot be reached anymore with the originally panned number of steps. The online planning algorithm automatically introduces additional steps and adjusts the others, so that the behavior can successfully accomplished, again resulting in quite human-like appearance of the generated behavior.

4.3. Approximation of training trajectories by robot movements

In order to validate our new architecture, we first tested the system by the realization of open-loop control, simulating a physical model of HRP-2 robot that

---

[4]https://goo.gl/Pn7atI
[5]https://goo.gl/JBz216
[6]https://goo.gl/9fLzO7

was implemented using the OpenHRP simulator, and also testing generated behaviors on the real robot.

In a first set of simulations the robot started from a parking position and makes a transition to a normal step. At the end of this step the pelvis velocities (propagation and angular) were determined and used as initial conditions for the generation of a three-action sequence. At the end of the last action, a step back to the final parking position was generated by spline interpolation of the pelvis angular and positional coordinates between the final state of the last step of the action sequence and the final position, introducing two additional steps on the spot. We also generated examples of four-action sequences. For this purpose, the retargeted trajectories were extended by an additional normal walking gait cycle. In order to augment the training data set for the learning of the mappings between the task parameters and the model parameters, we generated additional artificial kinematic data by scaling of the pelvis forward propagation velocities for all gait cycles uniformly (by the factors 0.8, 0.92, and 1.2), while keeping the upper body trajectories fixed. In this way we generated a total of 30 training examples from the original 10 motion capture trials. Examples of the generated three- and four-action sequences are shown in **movie**[7].

These trajectories were dynamically feasible for the robot, but still based on movement primitives learned from human data. In order to construct optimized primitives for the control of the robot, we generated 30 trajectories that were simulated with the OpenHRP physics simulator of the robot as novel training data and learned from this novel optimized movement primitives. For this purpose, the trajectories were approximated using 4 sources for the approximation of the first step, and 3 additional ones for the approximation of the residuals of the other steps, because this resulted in the best approximations with a small number of sources (Fig. 5).

A systematic validation of the approximation quality, dependent on the number of learned sources, is presented in Fig. 9. This figure shows histograms of the reproduction errors of the step sizes of the first two actions and the resulting arm reaching distance for the last action for different choices of the number of source functions. In all cases the spatial errors of the parameters, realized by the full control system, are small, always below 10 mm and often below 5 mm. This shows that in spite of the high complexity of the operations that are necessary to transform the original human motion
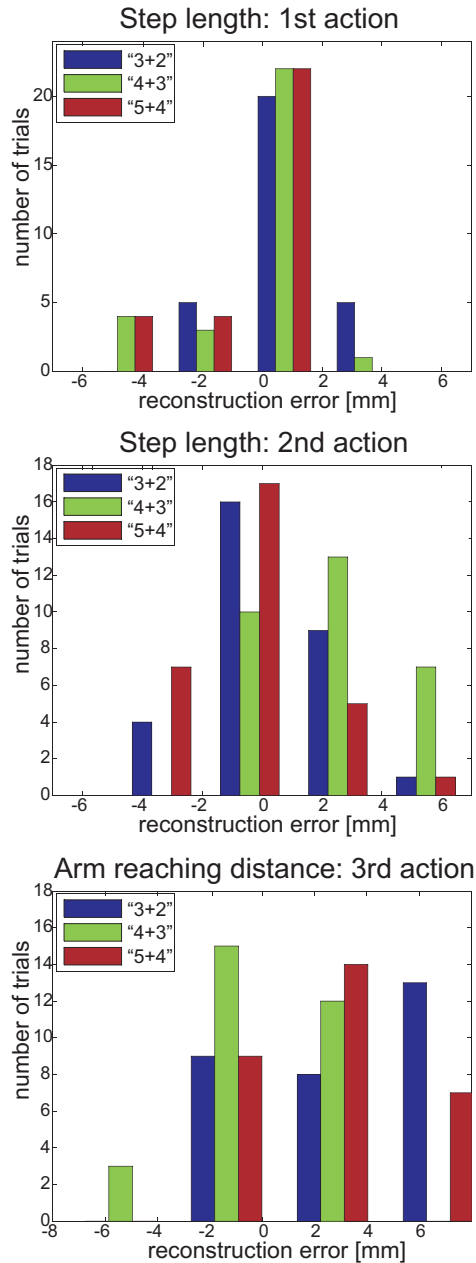

Figure 9: Reconstruction accuracy of the step-sizes and reaching goal distances for different numbers of sources. The figure shows the histograms of the spatial errors. The first number indicates the number of sources learned from the first step (action 1), and the second number the number of sources used to approximate the residuals of the other actions (2 and 3).

---

[7]https://goo.gl/7IZ0P1

into a motion sequence that is feasible for the robot the final control system produces movements that approximate the planned step sizes and reaching distances quite accurately.
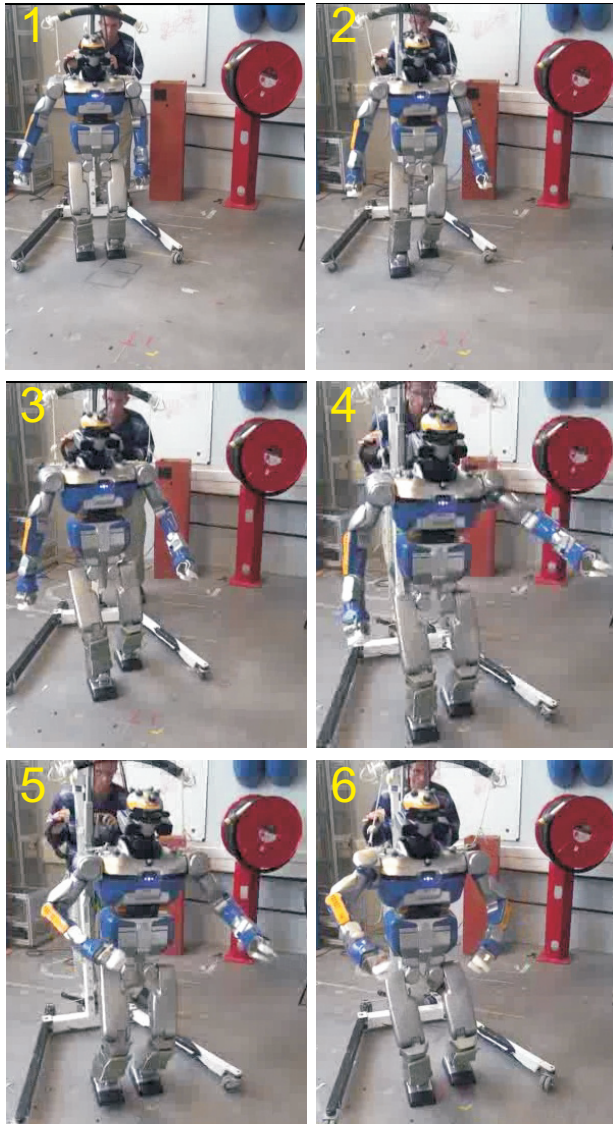

Figure 10: Real HRP-2 robot performing a 4-action walking-reaching sequence in the laboratory of LAAS/CNRS.

Some of these feasible re-synthesized trajectories were also tested using the real HRP-2 robot (cf. Fig. 10). A demonstration of the resulting behaviors for the three-action sequence is shown in **movie**[8], and a four-action sequence is shown in **movie**[9].

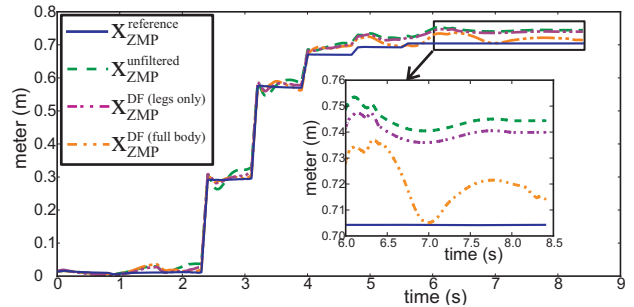[8] https://goo.gl/jjAVfT
[9] https://goo.gl/RqT6Q3

Figure 11: Trajectories of the Zero Moment Point (ZMP) (in walking direction, x coordinate) for different architectures. The blue curve $\mathbf{x}_{\text{ZMP}}^{\text{reference}}$ indicates the reference ZMP trajectory computed from the linear inverted pendulum model. The green curve $\mathbf{x}_{\text{ZMP}}^{\text{unfiltered}}$ shows the ZMP trajectory without filter correction. The trajectory with filter correction of all degrees-of-freedom is indicated in orange color $\mathbf{x}_{\text{ZMP}}^{\text{DF(full body)}}$, and the case where the dynamic filter was only applied to the lower-body degrees-of-freedom is indicated by the magenta trajectory $\mathbf{x}_{\text{ZMP}}^{\text{DF(legs only)}}$.

We also quantified the improvement of the behavior resulting from the inclusion of the dynamic filter in comparison with an architecture without this stage. Fig. 11 shows the x-coordinate trajectories of the Zero Moment Point (ZMP) for different model variants: 1) the idealized inverted pendulum model, which provides a reference trajectory for the underlying MPC approach (solid blue line $\mathbf{x}_{\text{ZMP}}^{\text{reference}}$); 2) the architecture without the dynamic filter correction (green dashed-dotted line $\mathbf{x}_{\text{ZMP}}^{\text{unfiltered}}$); 3) application of the dynamic filter only to the lower body degrees of freedom, assuming the upper body degrees-of-freedom to be frozen (magenta dashed-dotted line $\mathbf{x}_{\text{ZMP}}^{\text{DF(legs only)}}$) and 4) when the dynamic filter takes in account the full body motion (orange dashed-dotted line $\mathbf{x}_{\text{ZMP}}^{\text{DF(full body)}}$). The trajectory of a model without dynamic filter correction (green) deviates significantly from the planned reference trajectory (blue). The inclusion of the dynamic filter results in a much better approximation of the reference trajectory (orange color curve). This correcting effect of the dynamic filter is significantly reduced when it is only applied to the lower body degrees-of-freedom (magenta curve). This implies that only if the dynamic filter is applied to all degrees of freedom the robot motion is close to the planned dynamically feasible motion.

4.4. Inference of adaptive behaviors for novel gait distances

In order to test the architecture, with an online generation of new behaviors (step lengths and reaching move-

ments) dependent on the actual state of the robot, we synthesized the control signals for 30 different 4-action sequences, where a spectrum of step sizes was generated by linear morphing of the sources' weights. The first normal walking step length spanned 30 values in the range of 50.5 and 56.1cm, and the size of the second step was linearly sampled within the interval between 16.3 and 35.9cm. The reaching distance of the box in the last step varied in the interval of 66.3 to 75.5cm. The distance between the object and the front of the drawer was varied within the interval between 12.4 and 27.3 cm. The generated behaviors for the most extreme step sizes (smallest and largest) are shown in Fig. 12. **Movie**[10] shows these action sequences. For all tested intermediate step sizes that were not part of the initial training set very human-like coordinated behavior was generated.

In order to validate more precisely whether the generated closed-loop behaviors reproduce details of human grasping-reaching behavior we quantified the generated step sizes for different goal distances and starting positions. Figure 13 shows the step sizes generated for 10 different combinations of the two task parameters: distance of initial standing positions from the goal, and position of the object in the drawer. Consistent with the results in humans, the generated behavior shows a weak positive correlation between position of the object inside the drawer (blue) and the reaching distance of the arm (red). This reaching distance is almost constant, realizing the principle of the optimization of end-state comfort. The bars in the other colors indicate the durations of the starting step, the initial walking step, and the stopping step, which are changed in an adaptive manner similar to the behavior shown in Fig. 2. The starting gait cycle is not present in the human behavior, and is required in order to ensure a correct initiation of the first step from the parking position.

A more quantitative assessment of the performance is given in Fig. 14, which shows the variability of the ZMP in the lateral plane. The figure compares feasible trajectories, which are generated by the WPC from original trajectories without interpolation to novel step sizes or goal distances, with the behaviors of the system for novel goal distances that were not part of the training set and that required adaptation of the behavior using the online planning architecture. We compared again the behaviors for the choices of different number of sources for the anechoic mixing model (in total between 5 and 9 sources). The analysis is based on 30 newly synthesized four-action sequences for novel goal distances.
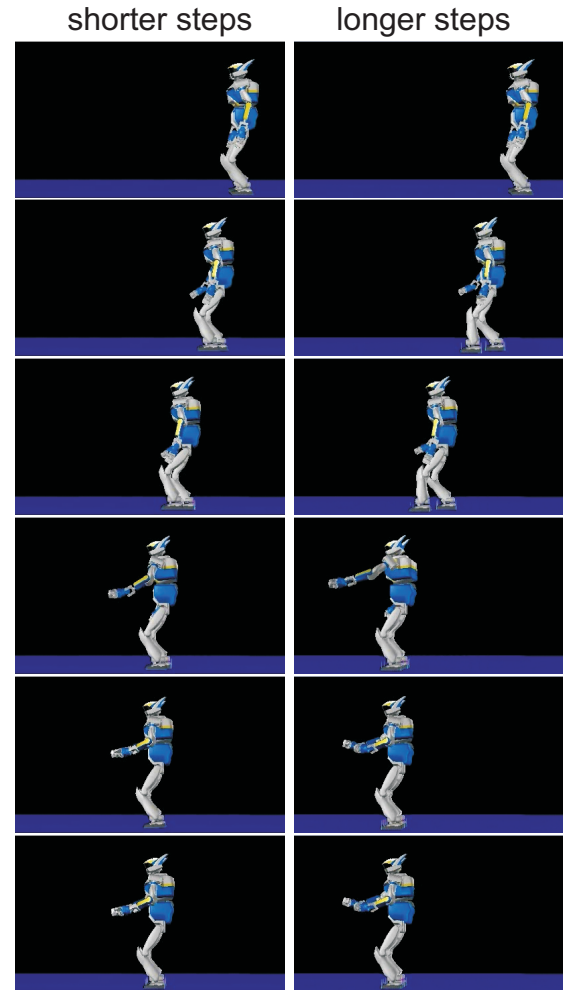
---

[10]https://goo.gl/IcwrXb

Figure 12: Synthesized behaviors with the full close-loop architecture, simulated with the OpenHRP simulator for the two most extreme goal distances.

The ZMP trajectory in the lateral plane was computed within all stance intervals, and the standard deviation (STD) of the difference between this trajectory and the reference ZMP trajectory was computed. The figure shows error bars with mean and variances as well as the maximum ranges of the variation. The ZMP variability is relatively independent of the number of sources for the reconstruction of trajectories and even for an inference of novel step distances the variability is not significantly higher that for original trajectories generated with the WPG. This shows that the closed-loop system produces highly stable behaviors in terms of the variation of the ZMP.
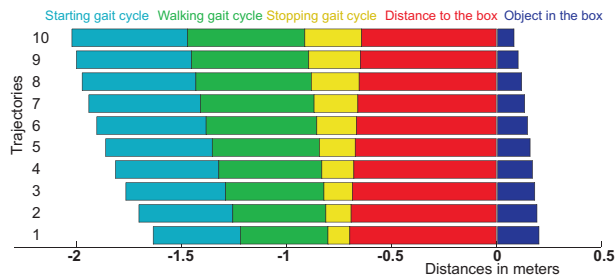
Figure 13: Simulation of multi-step sequences with the full closed-loop architecture using the OpenHRP simulator. The behavior reproduces details of the step-length relationships in humans, Specifically it reproduces adjustments which result in a relatively constant reaching distance during the last step, compatible with the *maximum end-state comfort hypothesis*. The results are shown for 10 simulated trials, aligning the positions according to the front of the drawer. Different colors refer to different steps within the sequence. (See text.)

4.5. Comparison with simple machine learning approach

One might ask if the proposed complex architecture is really necessary, and if one could not just learn dynamically feasible trajectories generated with the WPG and interpolate between the corresponding full-body kinematic angle trajectories using machine learning techniques. This approach would be based on the hope that the generated interpolations of the control signals also result in dynamically feasible behaviors when the training trajectories were dynamically feasible. We tested our method against such a simpler approach.

For this test we created training data consisting of 30 dynamically feasible walking-reaching trajectories, which were directly generated by the MPC-based WPG. Each of these trajectories results in dynamically stable behavior of the robot. The resulting full-body angle trajectories were again approximated with anechoic mixing models with different numbers of sources (between 5 and 9). Based on this training data 30 new trajectories for the new goal distances were computed, using either the simple machine learning approach discussed above, or with our method of learning upper-body and base trajectories separately.

The behaviors generated with the simple machine learning approach often result in falling of the robot, specifically during the last action (box opening and reaching for the object, where both arms are extended). The instability frequently also emerges earlier, already after the robot stops during the reaching step. A demon-
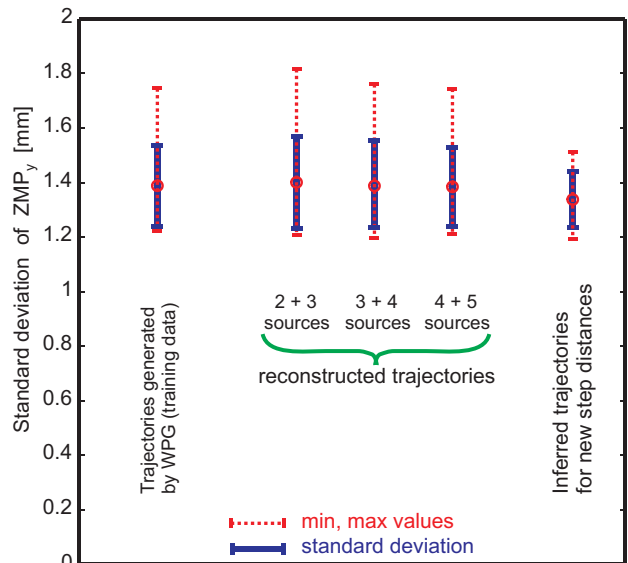

Figure 14: ZMP variability in the lateral plane (y-direction). See the text for more explanations.

stration of this behavior is given in **movie**[11], which compares the behavior generated by the naïve machine learning approach with the stable behavior obtained with our method. The parameters of the target behaviors are exactly identical for the two simulations.

A further quantitative analysis is given in Tab. 1, that shows how often the robot fell down out of the 30 novel synthesized behaviors. The simulations are grouped according to the speed of the walks. In addition, we tested interpolations generated with different numbers of source functions for the machine learning approach, and compared this with our method using 4+3 sources. For the low speed behaviors the machine learning approach leads to stable behaviors in some cases, and to falling in others, where the success of the method varies in a non-systematic manner with the number of source functions used for the approximation. For the fast speed movements the simple machine learning approach always results in falling in a significant number of cases. Opposed to this result, our method always results in stable behaviors without falling.

The superiority of our approach is also confirmed by an additional analysis of mechanical parameters that determine whether the behaviors can be realized on the real HRP-2 robot. Figure 15 shows the peak values of the ankle pitch torques for behaviors created directly using the MPC-based WPG, behaviors generated with the

---

[11]https://goo.gl/6hbX6g

14

The distribution of falling events			
Algorithm	slow speed (tr.1-10)	middle speed (tr.11-20)	fast speed (tr.21-30)
"ML 3+2"	0	10	7
"ML 4+3"	1	0	2
"ML 5+4"	0	5	10
Our method	0	0	0

Table 1: Fraction of trials with falls of the robot within 30 test trials with novel goal distances that were not part of the training set. Simple interpolation using machine learning techniques, approximating the trajectories with different numbers of sources (ML) is compared with our method that integrates online planning with the MPC control system. (For the ML conditions, the first number indicates the number of sources for the approximation of the fist action, and the second number the additional sources introduced for the approximation of the other steps).

naïve machine learning approach (ML) of approximating the full body angle trajectories, and our method. For the naïve machine learning approach almost all torque peak values exceed 30 Nm, which is infeasible for the robot (red shaded region in Fig. 15). This is especially true if this approach is used for the learning-based inference of the new trajectories. Contrasting with this result, the torques for behaviors generated directly with the WPG and the ones generated with our method are always in the feasible range. This is true for both, for the off-line reconstruction and for the learning-based inference using our method, and independent of the used number of source functions.

A similar result emerges for the analysis of the ground reaction forces (maximal normal force of the feet during the 4-action sequence). The maximum admissible force for the real HRP-2 is 800 N. Figure 16 shows that for the naïve ML approach in many cases the ground reaction force is too large compared to this limit, except for the reconstruction with 9 sources. Especially for the synthesis of new inferred behaviors, the peak ground reaction forces are always infeasible. This contrasts with the results obtained with our method. Here in all cases, for the off-line reconstruction and for the learning-based inference, the ground reaction forces are always in the feasible range and quite similar to the peak values that are obtained when the behavior is directly computed by the WPG using MPC.

Summarizing, we think that these results convincingly show that the proposed architecture provides a significant benefit over simpler approaches that just interpolate between control signals obtained from training data that corresponds to stable behaviors of the robot. The integration of online planning with the MPC-based control architecture in combination with the dynamic filter results in always stable and robust behavior, even largely independently of the accuracy of the learned trajectory model (number of source functions).

5. Conclusions

We have presented an architecture that combines the highly flexible online planning of coordinated full-body movements, based on learned dynamic movement primitives, with a control architecture that is based on a Walking Pattern Generator, which exploits nonlinear Model Predictive Control. The proposed architecture is suitable for the online generation of human-like highly coordinated full-body movements with long planning horizons. It generates dynamically feasible behavior of the robot, ensuring appropriate balance control during walking in presence of fast online replanning.

To our knowledge, no other presently existing approach allows the realization of such human-like long-term predictive motion planning in combination with a guarantee of dynamic balance during walking in combination with other tasks for the upper body of bipedal robots. Common alternative approaches, such as the optimization of such complex behavior by model-based optimal control approaches [Koschorreck and Mombaur (2012)] are presently computationally too costly to allow the online generation of such complex, where even the optimization of short multi-step sequences can take easily hours of computation time with the presently available optimization methods. The functionality and flexibility of the proposed architecture was demonstrated by simulation using the OpenHRP physics simulator and also in trials on the real HRP-2 robot. In addition, the proposed system realizes predictive motor behavior that is compatible with the end-state comfort hypothesis [Weigelt and Schack (2010); Rosenbaum (2008)]. Similar approaches have been proposed for the off-line optimization of reaching behaviors before (e.g. [Gienger et al. (2007, 2008)]).

The shown results represent only a first feasibility test of the proposed architecture, and they demonstrate that a single highly complex behavior can be robustly implemented on the HRP-2, resulting in robust behavior,
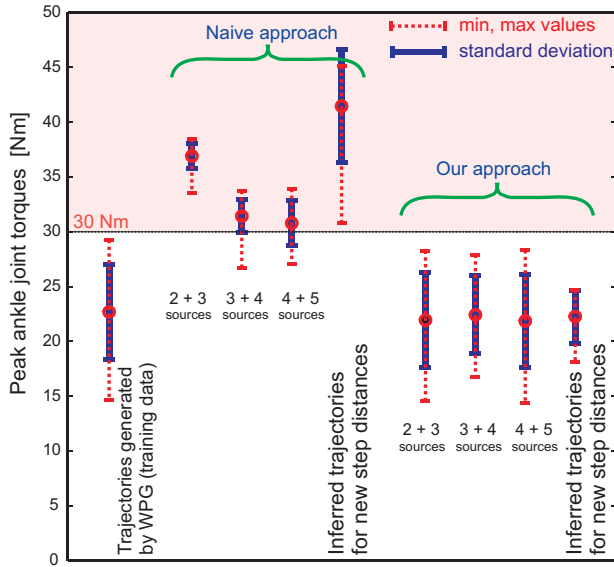
Figure 15: Peak ankle torques obtained for testing trials with different methods: WPG: trajectories generated with the WPG; naïve ML: interpolation of feasible control signals using machine learning methods; and with our method. We compare also resynthesis of training behaviors, using different numbers of sources for the approximation of the trajectories, and the synthesis of new trajectories for new target distances. (Blue error bars indicate mean and standard deviation. Red lines indicate the ranges between minimum and maximum value).
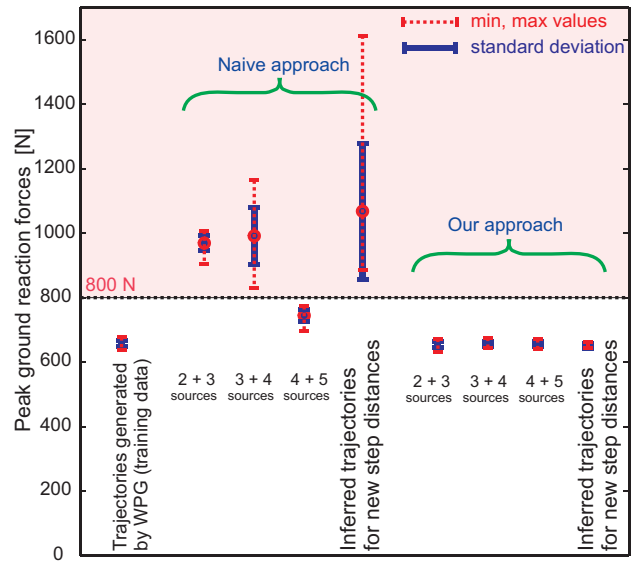

Figure 16: Peak ground reaction forces obtained for testing trials with different methods: WPG: trajectories generated with the WPG; naïve ML: interpolation of feasible trajectories using machine learning methods; and with our method. We compare also resynthesis of training behaviors, using different numbers of sources for the approximation of the trajectories, and the synthesis of new trajectories for new target distances. (Blue error bars indicate mean and standard deviation. Red lines indicate the ranges between minimum and maximum value).

where the closed-loop system so far has been only tested using the OpenHRP simulator, and presently is being implemented on the real robot. Our present work focuses on the implementation of the full adaptive algorithm on the real HRP-2 platform, testing the system in scenarios that require online replanning.

An extension of the approach to other classes of complex behaviors seems possible, since we have demonstrated elsewhere that the proposed primitive-based planning model is suitable for the highly adaptive synthesis of other types of complex behaviors, either of individual agents or even crowds [Mukovskiy et al. (2013)]. Since the computational efficiency of optimum control approaches is limited, we think that architectures like ours make a useful contribution to the control of humanoid robots, especially for the online generation of complex behaviors with longer planning horizons. The realization of the end-state-comfort hypothesis contributes to the creation of robots that realize principles of human motor control. This helps to increase the human-likeness and acceptance of humanoid robots,

and might be interesting for the realization of smoother interactions between real humans and humanoid robots.

Acknowledgements

References

Ajallooeian, M., van den Kieboom, J., Mukovskiy, A., Giese, M., Ijspeert, A., 2013. Programmable central pattern generators: an application to biped locomotion control. Physica D: Nonlin. Phen. 263, 41–56.

Ajoudani, A., Lee, J., Rocchi, A., Ferrati, M., Hoffman, E. M., Settimi, A., Caldwell, D. G., Bicchi, A., Tsagarakis, N. G., 2014. A manipulation framework for compliant humanoid COMAN: Application to a valve turning task. In: Humanoid Robots (Humanoids), 2014 14th IEEE-RAS International Conference on. pp. 664–670.

Atkeson, C. G., Moore, A. W., Schaal, S., 1997. Locally weighted learning. A.I. Review 11, 11–73.

Brandao, M., Jamone, L., Kryczka, P., Endo, N., Hashimoto, K., Takanishi, A., 2013. Reaching for the unreachable: integration of locomotion and whole-body movements for extended visually guided reaching. In Proc. of 13th IEEE-RAS Int. Conf. on Humanoid Robots (Humanoids), 28–33.

Buchli, J., Righetti, L., Ijspeert, A. J., 2006. Engineering entrainment and adaptation in limit cycle systems - from biological inspiration to applications in robotics. Biol. Cyb. 95 (6), 645–664.

Chang, C.-C., Lin, C.-J., 2001. LIBSVM: a library for support vector machines. Software available at http://www.csie.ntu.edu.tw/ cjlin/libsvm.

Chiovetto, E., d'Avella, A., Endres, D., Giese, M. A., 2013. A unifying algorithm for the identification of kinematic and electromyographic motor primitives. Bernstein Conference.

Chiovetto, E., Giese, M. A., 2013. Kinematics of the coordination of pointing during locomotion. Plos One 8 (11).

Dalibard, S., El Khoury, A., Lamiraux, F., Nakhaei, A., Taïx, M., Laumond, J., 2013. Dynamic walking and whole-body motion planning for humanoid robots: an integrated approach. Int. Journal of Robotics Research 32 (9-10), 1089–1103.

Feng, A. W., Xu, Y., Shapiro, A., 2012. An example-based motion synthesis technique for locomotion and object manipulation. Proc. of ACM SIGGRAPH I3D, 95–102.

Flash, T., Hochner, B., 2005. Motor primitives in vertebrates and invertebrates. Curr. Opin. Neurobiol. 15 (6), 660–666.

Flash, T., Sejnowski, T., 2001. Computational approaches to motor control. Curr. Opin. Neurobiol. 11, 655–662.

Gams, A., Ijspeert, A. J., Schaal, S., Lenarčič, J., 2009. On-line learning and modulation of periodic movements with nonlinear dynamical systems. Autonomous Robots 27 (1), 3–23.

Gams, A., Nemec, B., Žlajpah, L., Wächter, M., Ijspeert, A. J., Asfour, T., Ude, A., 2013. Modulation of motor primitives using force feedback: Interaction with the environment and bimanual tasks. In Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS 2013), 5629–5635.

Gienger, M., Bolder, B., Dunn, M., Sugiura, H., Janssen, H., Goerick, C., 2007. Predictive behavior generation - a sensor-based walking and reaching architecture for humanoid robots. In: Berns, K., Luksch, T. (Eds.), Informatik Aktuell (AMS). Springer-Verlag, pp. 275–281.

Gienger, M., Toussaint, M., Goerick, C., 2010. Whole-body motion planning building blocks for intelligent systems. In: Harada, K. (Ed.), Motion Planning for Humanoid Robots. Springer, pp. 67–98.

Gienger, M., Toussaint, M., Jetchev, N., Bendig, A., Goerick, C., 2008. Optimization of fluent approach and grasp motions. In Proc. of 8th IEEE-RAS Int. Conf. on Humanoid Robots (Humanoids 2008), 111–117.

Giese, M. A., Mukovskiy, A., Park, A., Omlor, L., Slotine, J. J. E., 2009. Real-time synthesis of body movements based on learned primitives. In: Cremers et al., D. (Ed.), Stat. and Geom. Appr. to Vis. Mot. Anal., LNCS5604. Springer, pp. 107–127.

Herdt, A., Diedam, H., Wieber, P.-B., Dimitrov, D., Mombaur, K., Diehl, M., 2010. Online walking motion generation with automatic foot step placement. Advanced Robotics 24 (5–6), 719–737.

Hsu, E., Pulli, K., Popovic, J., 2005. Style translation for human motion. ACM Trans. on Graphics 24, 1082–1089.

Huang, Y., Kallmann, M., 2014. Planning motions for virtual demonstrators. In: Intelligent Virtual Agents. Springer, pp. 190–203.

Ijspeert, A. J., Nakanishi, J., Hoffmann, H., Pastor, P., Schaal, S., 2013. Dynamical movement primitives: Learning attractor models for motor behaviors. Neural Computation 25 (2), 328–373.

Ijspeert, A. J., Nakanishi, J., Schaal, S., 2003. Learning attractor landscapes for learning motor primitives. Adv. in N.I.P.S. 15, 1547–1554.

Kajita, S., Kanehiro, F., Kaneko, K., Fujiwara, K., Harada, K., Yokoi, K., Hirukawa, H., 2003. Biped walking pattern generation by using preview control of zero-moment point. Proc. of Int. Conf. on Robotics and Autom., 1620–1626.

Kanoun, O., Laumond, J., Yoshida, E., 2011. Planning foot placements for a humanoid robots: a problem of inverse kinematics. The Int. J. of Robotics Research 30, 476–485.

Koenemann, J., Del Prete, A., Tassa, Y., Todorov, E., Stasse, O., Bennewitz, M., Mansard, N., 2015. Whole-body model-predictive control applied to the HRP-2 humanoid. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 3346–3351.

Koschorreck, J., Mombaur, K., 2012. Modeling and optimal control of human platform diving with somersaults and twists. Optimization and Engineering 13 (1), 29–56.

Kovar, L., Gleicher, M., Pighin, F., 2002. Motion graphs. Proc. of SIGGRAPH 2002, 473–482.

Kuindersma, S., Deits, R., Fallon, M., Valenzuela, A., Dai, H., Permenter, F., Koolen, T., Marion, P., Tedrake, R., 2015. Optimization–based locomotion planning, estimation, and control design for the Atlas humanoid robot. Autonomous Robots, 1–27.

Land, W. M., Rosenbaum, D. A., Seegelke, S., Schack, T., 2013. Whole-body posture planning in anticipation of a manual prehension task: Prospective and retrospective effects. Acta Psychologica 114, 298–307.

Levine, S., Wang, J. M., Haraux, A., Popović, Z., Koltun, V., 2012. Continuous character control with low-dimensional embeddings. ACM SIGGRAPH 2012. ACM Transactions on Graphics, Article N.28 31 (4).

Li, Y., Wang, T., Shum, H. Y., 2002. Motion texture: A two level statistical model for character motion synthesis. Proc. of SIGGRAPH 2002, 465–472.

Liu, K., Hertzmann, A., Popović, Z., 2005. Learning physics-based

motion style with nonlinear inverse optimization. ACM Trans. on Graphics 23 (3), 1071–1081.

Mühlig, M., Gienger, M., Steil, J., 2010. Human-robot interaction for learning and adaptation of object movements. In Proc. of IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS 2010), 4901–4907.

Mukovskiy, A., Land, W., Schack, T., Giese, M., 2015. Modeling of predictive human movement coordination patterns for applications in computer graphics. Journal of WSCG 23 (2), 139–146.

Mukovskiy, A., Slotine, J., Giese, M., 2013. Dynamically stable control of articulated crowds. Journal of Computational Science 4 (4), 304–310.

Naveau, M., Carpentier, J., Barthelemy, S., Stasse, O., Soueres, P., 2014. METAPOD - template META-PrOgramming applied to Dynamics: CoP-CoM trajectories filtering. Proc. of IEEE/RAS Int. Conf. on Humanoid Robot (ICHR).

Naveau, M., Kudruss, M., Stasse, O., Kirches, C., Mombaur, K., Souères, P., 2016, in press. A reactive walking pattern generator based on nonlinear model predictive control. IEEE Robotics and Automation Letters.

Omlor, L., Giese, M. A., 2011. Anechoic blind source separation using wigner marginals. J. of Machine Learning Res. 12, 1111–1148.

Petrič, T., Gams, A., Ijspeert, A. J., Žlajpah, L., 2011. On-line frequency adaptation and movement imitation for rhythmic robotic tasks. The Int. J. of Robotics Res. 30 (14), 1775–1788.

Rosenbaum, D. A., 2008. Reaching while walking: reaching distance costs more than walking distance. Psych. Bull. Rev. 15, 1100–1104.

Safonova, A., Hodgins, J., Pollard, N., 2004. Synthesizing physically realistic human motion in low-dimensional, behavior-specific spaces. ACM Trans. on Graphics 23 (3), 514–521.

Schaal, S., Ijspeert, A. J., Billard, A., 2003. Computational approaches to motor learning by imitation. Philos. Trans. R. Soc. of London, Biol. Sci. 358 (1431), 537–547.

Shoulson, A., Marshak, N., Kapadia, M., Badler, N., 2014. Adapt: The agent development and prototyping testbed. IEEE Trans. on Visualiz. and Comp. Graphics (TVCG) 99, 1–14.

Siciliano, B., Khatib, O. E., 2016. Springer Handbook of Robotics, 2nd ed. Springer International Publishing, Switzerland.

Sreenivasa, M., Souères, P., Laumond, J.-P., 2012. Walking to grasp: Modeling of human movements as invariants and an application to humanoid robotics. IEEE Trans. on System, Man and Cybernetic Part A: Systems and Humans 42 (4), 880–893.

Stasse, O., 2013. Habilitation Thesis. Paul Sabatier University, CNRS, Toulouse.

Stasse, O., Verelst, B., Davison, A., Mansard, N., Saidi, F., Vanderborght, B., Esteves, C., K., Y., 2008. Integrating walking and vision to increase humanoid autonomy. Int. Journal of Humanoid Robotics, Special Issue on Cognitive Humanoid Robots 5, 287–310.

Taïx, M., Tran, M., Souères, P. Guigon, E., 2013. Generating human-like reaching movements with a humanoid robot: A computational approach. J. of Computational Science 4, 269–284.

Tassa, Y., Erez, T., Todorov, E., 2012. Synthesis and stabilization of complex behaviors through online trajectory optimization. In: Intelligent Robots and Systems (IROS), IEEE/RSJ International Conference on. pp. 4906–4913.

Vuga, R., Ogrinc, M., Gams, A., Petrič, T., Sugimoto, N., Ude, A., Morimoto, J., 2013. Motion capture and reinforcement learning of dynamically stable humanoid movement primitives. In Proc. of IEEE Int. Conf. on Robotics and Autom. (ICRA'2013), 5284–5290.

Vukobratović, M., Stepanenko, Y., 1972. On the stability of anthropomorphic systems. Mathematical Biosciences 15, 1–37.

Wang, J. M., Fleet, D. J., Hertzmann, A., 2008. Gaussian process dynamical models for human motion. IEEE Trans. on Pattern Analysis and Machine Intelligence 30 (2), 283–298.

Weigelt, M., Schack, T., 2010. The development of end-state comfort planning in preschool children. Exper. Psych. 57 (6), 476–782.

Wieber, P.-B., 2006. Trajectory free linear model predictive control for stable walking in the presence of strong perturbations. Proc. of IEEE Int. Conf. Humanoids'06, 137–142.

Yoshida, E., Mallet, A., Lamiraux, F., Kanoun, O., Stasse, O., Poirier, M., Dominey, P.-F., Laumond, J.-P., Yokoi, K., 2007. ʿGive me the Purple Ballʾ – he said to HRP-2 N.14. In Proc. of IEEE-RAS Int. Conf. on Humanoid Robots (Humanoids'07).

18