



Modeling of coordinated human body motion by learning of structured dynamic representation

Albert A Mukovskiy, Nick Taubert, Dominik Endres, Christian Vassallo, Maximilien Naveau, Olivier Stasse, Philippe Souères, Martin A Giese

► To cite this version:

Albert A Mukovskiy, Nick Taubert, Dominik Endres, Christian Vassallo, Maximilien Naveau, et al.. Modeling of coordinated human body motion by learning of structured dynamic representation. Geometric and Numerical Foundations of Movements, Springer, pp.237-267, 2017, Star 117. hal-02002507

HAL Id: hal-02002507

<https://laas.hal.science/hal-02002507>

Submitted on 31 Jan 2019

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.

Modeling of coordinated human body motion by learning of structured dynamic representations

Albert Mukovskiy¹, Nick Taubert¹, Dominik Endres², Christian Vassallo³, Maximilien Naveau³, Olivier Stasse³, Philippe Souères³, and Martin A. Giese¹

Abstract The modeling and online-generation of human-like body motion is a central topic in computer graphics and robotics. The analysis of the coordination structure of complex body movements in humans helps to develop flexible technical algorithms for movement synthesis. This chapter summarizes work that uses learned structured representations for the synthesis of complex human-like body movements in real-time. This work follows two different general approaches. The first one is to learn spatio-temporal movement primitives from human kinematic data, and to derive from this Dynamic Movement Primitives (DMPs), which are modeled by nonlinear dynamical systems. Such dynamical primitives are then coupled and embedded into networks that generate complex human-like behaviors online, as self-organized solutions of the underlying dynamics. The flexibility of this approach is demonstrated by synthesizing complex coordinated movements of single agents and crowds. We demonstrate that Contraction Theory provides an appropriate frame-

¹ Albert Mukovskiy (corresp. author), e-mail: albert.mukovskiy@medizin.uni-tuebingen.de · Martin A. Giese, e-mail: giese@uni-tuebingen.de, · Nick Taubert, e-mail: nick.taubert@uni-tuebingen.de
Section for Computational Sensomotorics, Department of Cognitive Neurology, Hertie Institute for Clinical Brain Research & Center for Integrative Neuroscience, University Clinic Tübingen, Otfried-Müller Str. 25, 72076 Tübingen, Germany.

² Dominik Endres, e-mail: dominik.endres@uni-marburg.de
Theoretical Neuroscience Group, Section for General and Biological Psychology, Department of Psychology, Philipps-University Marburg, Gutenbergstr. 18, 35032 Marburg, Germany.

³ Maximilien Naveau, e-mail: maximilien.naveau@laas.fr, · Olivier Stasse, e-mail: ostasse@laas.fr, · Christian Vassallo, e-mail: christian.vassallo@laas.fr, · Philippe Souères, e-mail: philippe.soueres@laas.fr
Gepetto Lab, LAAS/CNRS, Université de Toulouse, Av. du Colonel Roche 7, F-31400, Toulouse, France.

J.P. Laumond et al. (Eds.): Geometric and Numerical Foundations of Movements, Springer STAR Series, 2016. ©Springer-Verlag Berlin Heidelberg 2016

work for the design of the stability properties of such complex composite systems. In addition, we demonstrate how such primitive-based movement representations can be embedded into a model-based predictive control architecture for the humanoid robot HRP-2. Using the primitive-based trajectory synthesis algorithm for fast on-line planning of full-body movements, we were able to realize flexibly adapting human-like multi-step sequences, which are coordinated with goal-directed reaching movements. The resulting architecture realizes fast online planning of multi-step sequences, at the same time ensuring dynamic balance during walking and the feasibility of the movements for the robot. The computation of such dynamically feasible multi-step sequences using state-of-the-art optimal control approaches would take hours, while our method works in real-time.

The second presented framework for the online synthesis of complex body motion is based on the learning of hierarchical probabilistic generative models, where we exploit bayesian machine learning approaches for nonlinear dimensionality reduction and the modeling of dynamical systems. Combining Gaussian process Latent Variable Models (GPLVMs) and Gaussian process Dynamical Models (GPDMs), we learned models for the interactive movements of two humans. In order to build an online reactive agent with controlled emotional style, we replaced the state variables of one actor by measurements obtained by real-time motion capture from a user and determined the most probable state of the interaction partner using Bayesian model inversion. The proposed method results in highly believable human-like reactive body motion.

Key words: Dynamic Movement Primitives, Animation, Machine Learning, Gaussian Process Latent Variable Model, Gaussian Process Dynamical Model, Navigation, Walking Pattern Generator, Goal-directed Movements, Motor Coordination, Action Sequences

1 Introduction

The generation of realistic human movements in reactive fashion is a difficult task with high relevance for computer graphics and robotics. An especially challenging task in this domain is the online-synthesis of complex behaviors that consist of sequences of individual actions, which adapt to continuously changing environmental constraints.

The whole body movements of humans and animals are organized in terms of muscle synergies or movement primitives [4, 17]. Such primitives characterize the coordinated involvement of subsets of the available degrees of freedom in different actions. An example is the coordination of periodic and non-periodic components of the full-body movements during reaching while walking, where behavioral studies reveal a mutual coupling between these components [12, 8, 68, 47]. The realism and human-likeness of synthesized movements in robotics and computer graphics can be improved by taking such biological constraints into account [18, 15, 73].

In this chapter we present two learning-based frameworks that make such biological properties applicable to the realtime synthesis of human-like movements in technical systems, one that is based on *Dynamic Movement primitives (DMPs)*, and another one that exploits unsupervised Bayesian learning methods.

The chapter is organized into three main sections. Section 2 introduces a framework that approximates complex human movements by combining learned dynamic movement primitives. Highly adaptive coordinated full-body movements, and even the coordination of the movements of multiple agents, can be generated online by networks of such dynamic primitives, which are mutually coupled. Section 3 discusses how the same methods can be exploited for the movement planning of humanoid robots. We present an architecture that embeds such an online synthesis model into a control architecture of a real humanoid robot, which is based on model predictive control. The proposed solution ensures the dynamic balance of the robot, so that it is prevented from falling, while realizing highly flexible online planning of movements. The last section 4 introduces a completely different approach for the learning-based representation of reactive human movements, which is based on Bayesian machine learning methods for dimension reduction and model inversion. Space constraints allow us only to give the outline of these different approaches, and we refer to the cited original publications with respect to many technical and mathematical details.

2 Modeling of human movements based on learned primitives

Human full-body movements involve typically a large number of degrees of freedom. It has been a classical idea in biological motor control that such complex body movements might be composed from lower-dimensional control units, often referred to as *movement primitives* or *synergies*. Substantial work in motor control has been dedicated to the identification of such primitives from kinematic and EMG data, applying unsupervised learning techniques for dimension reduction [14, 31, 71]. Different techniques have been applied, including Principle Component Analysis (PCA), Independent Component Analysis (ICA), or more sophisticated methods that include time shifts of the superpositioned components. Such methods approximate a set of time-dependent signals by a superposition of learned source functions, which have been interpreted as movement primitives or (muscle) synergies.

Work in computer graphics shows that the accurate approximation of motion capture data from complex full-body movements using PCA requires typically more than 8 principal components (e.g.[70]). In the following section we describe a method that often leads to more compact representations with less components or primitives. Such compact representations are important especially if parameters of the learned models have to be interpreted, e.g. in order to characterized motion styles [65]. Compact models tend to concentrate the data variance on a low number of interpretable parameters. Compact primitive-based representations are also beneficial if they are embedded into control systems or dynamic architectures for the online

generation of motion. In this case, the number of primitives determines the dimensionality of the underlying system dynamics, and systems with lower dimensionality often are easier to control and more robust against perturbations.

In the following, after reviewing some related methods in Section 2.1, we give first a short introduction in the method that we apply to learn primitives from trajectory data (Section 2.2). The resulting kinematic primitives are given by basis functions or trajectories, which by appropriate combination can approximate complex joint angle trajectories. We then discuss how from such kinematic primitives dynamic movement primitives can be constructed that generate the learned trajectories online (cf. Section 2.3). These dynamic primitives are nonlinear dynamical systems that produce the learned basis trajectories as stable solutions. In the following Section 2.4 we demonstrate how such learned dynamical generative models can be augmented by controllers that make the behaviors adaptable, realizing for example navigation through space or the control of step length or emotional style. It is demonstrated that the developed approach is suitable for the online generation of quite complex coordinated behaviors, either of single agents or even of whole crowds of agents that execute coordinated collective behaviors. In Section 2.5 we discuss finally, how such complex generative dynamical models can be designed, guaranteeing the robustness of their solutions. Contraction Theory, a special type of mathematical stability analysis, which is especially applicable to nonlinear systems which are composed of many components, makes it possible to ensure that the desired behavior is the only stable solution of the resulting nonlinear dynamical system.

2.1 *Related work*

The synthesis of the kinematics of sequences of human full-body movements has been treated extensively in computer graphics [42]. The prominent classical approach for the synthesis of human motion in computer graphics is the adaptive interpolation between motion-captured example actions, which is typically realized off-line [93, 23, 22, 2]. Other approaches are based on learned low-dimensional parameterizations of whole body motion that are embedded in mathematical frameworks for the online generation of motion [9, 66, 43, 70, 27, 88, 44]. In addition, a variety of methods for the segmentation of action streams into individual actions have been proposed, where the models for individual actions can be adapted online in order to fulfill additional constraints, such obstacle avoidance or the correct positioning of end-effectors [34, 67, 62, 16, 28]. Only very few of this work has focused on the modeling of the flexible coordination of groups of degrees of freedom, similar to synergies in biological systems [70, 75].

2.2 Approximation of human movement data by anechoic mixtures

Many standard approaches use principle component analysis (PCA) or independent component analysis (ICA) for the reduction of the dimensionality of motion data. A set of trajectories is represented as a linear combination of a limited number of basis components or source functions. In our work we used a more sophisticated mixture model for the approximation of the joint angle trajectories that contains time shifts for the superposed components or sources [56, 11]. This model is known from acoustics as *anechoic mixture* and superposes source functions s_j that are temporally shifted with the time delays τ_{ij} in order to approximate a set of trajectories $\xi_i(t)$. The corresponding model is characterized mathematically by the equation

$$\underbrace{\xi_i(t)}_{\text{angles}} = m_i + \sum_j w_{ij} \underbrace{s_j(t - \tau_{ij})}_{\text{sources}} + \text{noise} \quad (1)$$

The parameters w_{ij} specify the mixing weights, and the variables m_i signify constant offsets (means) of the approximated trajectories. Learning of an anechoic mixture model requires the estimation of these parameters, the source functions s_j , and the delays τ_{ij} . In our case, the trajectories were given by the angle trajectories of 17 joints expressed as quaternions. We have shown in previous work for different classes of human movements that this anechoic mixture model results in very accurate approximations of complex human movement data, often with as few as 3-4 source functions, and typically with factor 2 less sources than classical approaches using PCA or ICA [60, 50].

2.3 Online synthesis by networks of dynamic primitives

The discussed mixture model can be applied for an off-line analysis and synthesis of classes of trajectories. Movement types or styles can be characterized by the mixing weights (and delays) of the model, and the movement can be analyzed using these weights as features. In addition, novel movement trajectories can be generated off-line by specifying or interpolating these parameters, and using the equation (1) as a generative model [65]. However, this approach is not sufficient for applications that require an online synthesis of complex movements.

In order to make the learned structured model applicable for real-time synthesis we associated each learned source function (*kinematic primitive*) with an associated dynamic primitive [29, 19]. The dynamic primitives are defined by dynamical systems whose stable solutions approximate the learned source functions. Each dynamical primitive is defined by a *canonical dynamical system*, which has an attractor solution with well-defined mathematical properties. We used limit cycle oscillators (Andronov-Hopf oscillators) for the approximation of periodic source functions, and a ramp-like solution, which is derived from the state of a limit cycle oscillator, for the non-periodic ones. We then learned nonlinear functions that map the state spaces

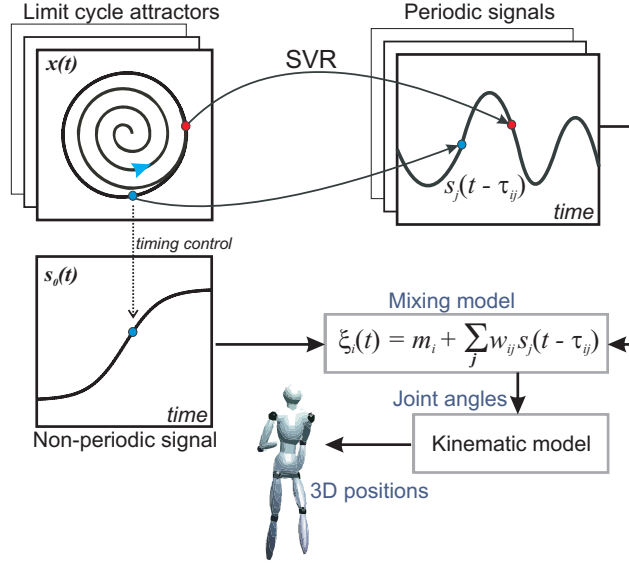


Fig. 1 Architecture for the online synthesis of body movements using dynamic primitives. The solutions $x_j(t)$ of a canonical dynamical systems (limit cycle oscillators) are mapped by Support Vector Regression (SVR) onto the values of the periodic source functions $s_j(t)$. In addition, a non-periodic source function $s_0(t)$ is constructed from these solutions. From these online generated source functions joint angle trajectories are computed using the learned anechoic mixing model.

of the canonical dynamics onto the values of the source functions $s_j(t)$ using Support Vector Regression (SVR) [10]. Figure 1 shows an overview of the developed architecture for real-time synthesis.

By insertion of couplings between the different canonical dynamical systems it is possible to synchronize their dynamics, so that the corresponding source signals are evolving in synchrony. Such couplings can be used either to model coordinated behavior between the movement primitives within a single agent, or by introduction of couplings between the dynamics of multiple agents, for the simulation of coordinated interactive behavior of multiple agents.

2.4 Style morphing and navigation

The proposed method for the online generation of body motion trajectories can be combined with a dynamic variations of motion style. For this purpose multiple examples of the same motion were motion-captured that realize different styles, and intermediate styles were generated by online interpolation (motion blending). For this purpose, we linearly interpolated the average angles m_i , the mixing weights




Fig. 2 a) Reactive online control of locomotion. Agents avoid the obstacles (poles) and other agents in the scene. Trajectories are generated by morphing between steps with different length, and curvatures of the walking path (left, straight, right), where blending weights are controlled by a navigation dynamics that controls the heading direction dependent on obstacle and goal positions. b) Folk dance of two couples, one forming a bridge, and the other couching beneath it. The behavior is fully self-organized, where the behavior of the agents depends on the relative positions with respect to the other agents.

w_{ij} , and the delays τ_{ij} of mixing models that were learned from training trajectories representing different motion styles. (See [21, 50] for further details.)

The blending weights were modulated by controllers that depend on task parameters, such as the position or orientations of agents in the scene, distances of agents to goal points, etc. One example is the generation of walking steps that realize locomotion along curved paths by morphing between straight and curved walking steps to the right or to the left. In this case, the morphing weights of the three walking patterns were determined by a controller that determines the heading direction dependent on obstacles and desired goal points (cf. [74, 91]). Likewise, movements with different emotional styles can be generated by blending between models that realize the same motion with different emotional styles, or steps with different length can be generated by morphing between long and short steps.

We worked out an application of this approach for the simulation of locomoting and navigating agents. Blending weights were controlled by a simplified version of a dynamic navigation model that had been applied successfully in robotics before [74], and which we extended by inclusion of a prediction of expected collision points with obstacles in order to make the navigation behavior more human-like [59]. The heading direction is controlled by a nonlinear first-order differential equation that depends on the actual positions of the agents and of obstacles in the scene. (See [21, 59] for details.)

An example for the navigation behavior that can be simulated with the described architecture is illustrated in Fig. 2 a), where six agents avoid the obstacles and the other agents. **Demo Movie I**¹ shows this example and other obstacle avoidance scenarios. The same method can also be exploited in order to model interactions between multiple agents that realize more complex behaviors, which integrate periodic and non-periodic movement primitives. An example is shown in Fig. 2 b) that shows a figure from a folk dance that requires one couple of agents to walk beneath a bridge that is formed by the arms of another couple. Both both couples

¹ <http://tinyurl.com/hvww9ra>

take turn and change places. This whole complex behavior with highly human-like appearance was completely self-organized using only 10 prototypical movements (normal walk, walks with left or right arm lifted, crouching walk, left and right forward turnings with two different angular velocities, left and right backward stepping turns). See also **Demo Movie I**. The proposed approach thus can be used to simulate highly complex full-body coordination patterns, and even patterns that include multiple agents. The underlying architecture is very simple, consisting only of a low-dimensional nonlinear dynamical system and some linear and nonlinear mappings. This makes it possible to generate the behavior, even of larger groups of agents in realtime.

2.5 Dynamic stability design exploiting Contraction Theory

Effectively, the proposed method synthesizes desired motion trajectories online by generating them as stable solutions of a complex dynamical system, which can be characterized as a 'network' of dynamic movement primitives. The elements of such networks are highly nonlinear: the canonical dynamics, the mappings from the state space of the canonical state variables to the source functions, and the kinematic relationship between the joint angles and the behavior of the agent in the external space. This raises the question whether for such systems any guarantees can be given that the desired behavior is the only stable solution of the system. This question is of particular importance because for nonlinear systems, and even more for complex ones, multiple stable solutions may exist.

An interesting control-theoretical approach for the analysis of the stability of composite dynamical systems that consist of coupled nonlinear elements is *Contraction Theory (CT)* [45]. We were able to show that this method is suitable to guarantee the stability of highly coordinated behaviors of crowds of locomoting avatars, where our dynamical models included the full complexity and nonlinearity that is generated by the body articulation of the locomoting agents.

The question of the dynamic stability of the created behaviors has been rarely addressed in traditional work on crowd animation. Like in our work, some approaches also tried to learn rules of interactive behaviors from human crowds [13, 58, 41], while other approaches tried to optimize the interaction within crowds by numerical optimization of appropriate cost functions (e.g. [25]). Most existing approaches for the control of group motion in computer graphics neglect the effects of the body articulation during locomotion on the control dynamics [63, 54, 36]. Another field that typically pays attention to the dynamic stability of solutions is control theory. Some work in this area has studied the temporal and spatial self-organization of crowds, typically assuming highly simplified and partly even linear agent models (e.g. [72, 57]). This shows that for more detailed models of the agent dynamics systematic methods for the design of a stable system dynamics are largely lacking.

Contraction Theory is a special type of nonlinear stability analysis that has been introduced by J.-J. Slotine and coworkers ([45, 89, 64]). The special property of this

framework, which makes it possible to simplify the analysis of complex composite systems, is that it permits to transfer stability results from parts to composite systems. In general, such a transfer is not possible for nonlinear dynamical systems, which typically renders the analysis of composite nonlinear dynamical systems impossible, even for moderately-sized systems.

Opposed to the classical approach for stability analysis that computes first a stationary solution and then linearizes about it, Contraction Theory analyzes differences between trajectories with different initial conditions. If these differences vanish exponentially over time, all solutions converge towards a single trajectory, independent from the initial states. In this case, the system is called *contracting*, and at the same time is globally asymptotically stable. More specifically, for a general dynamical system of the form

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, t) \quad (2)$$

assume that $\mathbf{x}(t)$ is one solution of the system, and $\tilde{\mathbf{x}}(t) = \mathbf{x}(t) + \delta\mathbf{x}(t)$ a neighboring one with a different initial condition. The function $\delta\mathbf{x}(t)$ is also called *virtual displacement*. With the Jacobian of the system $\mathbf{J}(\mathbf{x}, t) = \frac{\partial \mathbf{f}(\mathbf{x}, t)}{\partial \mathbf{x}}$ it can be shown [45] that any nonzero virtual displacement decays exponentially to zero over time if the symmetric part of the Jacobian $\mathbf{J}_s = (\mathbf{J} + \mathbf{J}^T)/2$ is uniformly negative definite, denoted as $\mathbf{J}_s < 0$. This implies that it has only negative eigenvalues for all relevant state vectors \mathbf{x} (within a contraction region). In this case, it can be shown that the norm of the virtual displacement decays at least exponentially to zero, for $t \rightarrow \infty$. If the virtual displacement is small enough, one can also prove the inequality: $\|\delta\mathbf{x}(t)\| \leq \|\delta\mathbf{x}(0)\| e^{\int_0^t \lambda_{\max}(\mathbf{J}_s(\mathbf{x}, s)) ds}$. This implies that the virtual displacements decay with a *convergence rate* (inverse timescale) that is bounded from below by the quantity $\rho_c = -\sup_{\mathbf{x}, t} \lambda_{\max}(\mathbf{J}_s(\mathbf{x}, t))$, where $\lambda_{\max}(\cdot)$ signifies the largest (negative) eigenvalue.

Contraction analysis can be generalized to systems with individual non-contracting directions (partial contraction) [89]. This is important, for example, for limit cycle oscillators, where the directions tangential to the stable oscillatory solution are non-contracting, but the system is contracting in all other directions orthogonal (transversally) to these trajectories. Contraction analysis can be applied to hierarchically coupled systems [45], where the systems on higher hierarchy levels do not feed back into the lower levels. Such systems can be shown to be contracting if each component system is contracting for all bounded inputs. In addition, one can derive constraints for the coupling between two contracting systems that are reciprocally connected (i.e. in a non-hierarchical forward-backward fashion) that guarantee that the resulting system also is contracting. This makes it possible to design contracting systems from contracting system components, by appropriate design of hierarchical and reciprocal connections of the modules. We applied this framework to a simplified model of the dynamics that generates coordinated behavior of crowds.

In order to apply Contraction Theory for the stability analysis of locomoting crowds, we used a model that integrated the following control levels: (i) Control of heading direction (as described before); (ii) step-size control by morphing between long and short steps; (iii) control of the gait phase in order to achieve a synchro-




Fig. 3 Self-organized reordering of a crowd. Control dynamics affects direction, row and column distances, and gait phases. a) When the sufficient contraction conditions of the system dynamics are satisfied the agents organize into an ordered formation where all agents synchronize their steps. b) For a violation of the contraction conditions the behavior becomes unstable, and the agents diverge and do not synchronize their behaviors. See [52].

nization between all agents; and (iv) control of step frequency by adaptation of the frequency parameters of the limit cycle oscillators. (See [21, 61, 51] for further details.) The resulting dynamics can be approximated by a simplified nonlinear differential equation system that depends on a nonlinear function that describes the relationship between the propagation speed of the characters and the corresponding state variable of the canonical system. This approximative system dynamics is accessible for an application of tools from contraction theory. This allows to derive sufficient contraction conditions that ensure that the generated behavior is stable and that no other attractors of the system dynamics exist. (Further details about this analysis are laid out in [51, 52].)

Fig. 3 a) shows a crowd with 36 avatars generated with a dynamics that fulfills the derived contraction conditions. By self-organization the group evolves into a spatially ordered configuration with a synchronization of gait phase, and step frequency. This behavior is robustly approached from different initial conditions and placements of the agents within the scene. Figure 3 b) shows the situation of the relevant contraction condition is violated. In this case, the crowd diverges and the dynamics becomes unstable. This example demonstrates the applicability of CT for stability design even for systems that model quite complex coordinated behaviors. (See also **Demo Movie II**²).

3 Planning of movements for humanoid robots

Standard approaches for kinematic planning in robotics model complex sequential activities by concatenations of elementary motions, each one accomplishing a spe-

² <http://tinyurl.com/jxgpptb>

cific sub-task. Differing from this, skilled human behavior is highly predictive, and behaviors are adapted to task constraints even far in the future. An example for this is the *maximum end-state comfort principle* [69] that has been demonstrated for the human coordination of walking and reaching ([92, 37]):

It seems desirable to transfer such flexible human-like planning strategies to robots, e.g. for the generation of locomotion behaviors that are coordinated with hand or arm actions. The mathematical framework presented in Section 2 is suitable for the modeling of such highly predictive coordinations strategies. For this purpose, the desired behavior of the robot is synthesized online by a network of dynamic primitives, exploiting the architecture described in the previous section and specifying a virtual kinematic trajectory that the robot should follow. However, real robots are associated with additional constraints, e.g. for the joint angles or realizable torques. In addition the behavior of bipedal walking robots has to ensure specific constraints to ensure dynamic balance, in order to prevent the robot from falling. This part of the chapter describes how the framework presented in Section 2 can be embedded as online motion planning system in the control architecture of the humanoid robot HRP-2.

The core of this control architecture is a Walking Pattern Generator (WPG), which is based on nonlinear Model Predictive Control [55]. The underlying algorithm is based on a simplified model of a bipedal walker and synthesizes a dynamically feasible behavior of the legs that prevents the robot from falling. This lower body motion is then combined with the desired motion of the arms, correcting the lower body motion by a special *Dynamic Filter* [94], in order to ensure that the overall behavior is always dynamically feasible and thus realizable on the real robot without falling. We demonstrated the functionality of this architecture for the example of coordinated walking and reaching. The developed system models flexible and very human-like behaviors for the online replanning after perturbations of the behavior, which realizes the *maximum end-state comfort principle* of human motor control.

Compared to a direct computation of dynamically feasible multi-step movements using optimal control approaches (c.f. [33]), our method is characterized by a much lower computational complexity. Optimal control methods using an accurate model of the robot require typically hours of computation time for the generation of multi-step sequences that ensure that the robot does not fall. The same goal can be achieved with our method with a computational complexity that is of the same order as the one of standard real time-capable WPG algorithms [26].

3.1 Related work

Some work in computer graphics has integrated prioritized control and stack-of-task approaches in the synthesis of trajectories from training trajectories [16, 75]. In robotics, numerous architectures which combine walking and grasping have been proposed that are not directly inspired by human behavior [1, 35, 78, 7]. Human-

inspired frameworks for the decomposition of human reach-to-grasp movements into sequential actions were proposed in [76] and [48]. An algorithm for the computation of optimal stance locations with respect to the reaching target within a dynamical systems approach was proposed in [20]. In [96] a task priority approach was applied for the integration of several sub-tasks, including stepping, hand motion, and gaze control. Other work has exploited global path planning in combination with walking pattern generators (WPGs) [32] in order to generate collision-free dynamically stable gait paths. A first attempt to transfer human reaching movements to humanoid robots by using motion-primitives was proposed in [79].

3.2 Drawer opening task

Human motor sequences have been shown to be highly predictive. Our implementation of such predictive strategies on a humanoid robot is based on recent study on the coordination of walking and reaching in humans [37]. Participants had to walk towards a drawer and to grasp an object, which was located at different positions in the drawer. Participants optimized their behavior already multiple steps before the object contact, consistent with the hypothesis of *maximum end-state comfort* during the reaching action [92, 68]. This implies that the steps prior to the reaching were modulated in a way that optimized the distance for the final reaching action in a way that simplified the reaching and grasping.

The initial distance from the drawer and the position of the object inside it were varied in the data set. The participants walked towards a drawer, opened it with their left hand and reached for an object inside the drawer with their right hand [49] (see Fig. 4). Each recorded sequence included three subsequent actions: 1) a normal walking step; 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 (consistent with the *maximum end-state comfort hypothesis*); 3) the drawer opening combined with the reaching for the object while standing. **Demo Movie III**³ shows an example for the recorded human behavior.

3.3 Adaptive model of the kinematics of multi-step sequences

In order to make the recorded motion capture data useful for a transfer of the behavior to the robot, it was retargeted on a kinematic model of the human-sized humanoid robot HRP-2 using the commercial software MotionBuilder. During retargeting the feet positions of the HRP-2 were constrained to level ground, and the step sizes were reduced proportionally to the height of the robot. This made the joint angle ranges

³ <http://tinyurl.com/he3dhhb2>




Fig. 4 Important intermediate postures from the human behavior: step with initiation of reaching, standing while opening the drawer, and reaching for the object.

compatible with the ones that can be realized by the robot. The data was split, separating the stored pelvis trajectories (pelvis position and pelvis direction angles in the horizontal plane), and the upper body trajectories, approximating the human trajectories by a kinematic model of the HRP-2. The pelvis position trajectories were also rescaled in order to match the maximally admissible propagation speed limit of the HRP-2 (0.5 m/sec). In addition, corrections were applied to the pelvis and trunk yaw-angle trajectories. Figure 5 shows a comparison between an original human and the retargeted pose, illustrated using the corresponding avatar models. (See also **Demo Movie IV⁴**.)




Fig. 5 Retargeting of the movements from a human to the unconstrained skeleton of the HRP-2 robot.

⁴ <http://tinyurl.com/j8qnbtp>

In order to model step sequences with a human-like coordination of the periodic walking and the non-periodic reaching behavior, we approximated the training data by anechoic mixtures (see Section 2.2). For this specific application, we used a step-wise regression approach. We introduced a total of five sources in order to model the three different actions (steps) within the sequence. The first action is the normal walking gait cycle. After extraction of the the mean value and non-periodic component from this step, we approximated the remaining residuals by anechoic mixtures with three sources, applying a modified demixing algorithm that constrained all time delays belonging to the same source function and joint angle to be equal across all trials. These additional constraints make it necessary to introduce more sources in order to reach the same approximation quality, but significantly simplify the motion morphing. The second highly adaptive step was approximated by the same sources, and the remaining residuals were modeled by introducing two additional periodic sources. The same set of five sources were then used to model the last action.

In order to control the styles of the actions online, we learned nonlinear mappings between task parameters (step length and duration) and the weights of the source functions in our mixing model, applying Locally Weighted Linear Regression (LWLR) [3, 49]. For the synthesis of multi-step sequences the step lengths was computed from the actual estimated target distance. Based on the training data, we computed the achievable step ranges. Additional steps were automatically introduced if the target could not be reached within three steps. For the second step, the step length was adjusted in order to realize a maximum comfort distance for reaching, and the planning distance of the other steps was adjusted accordingly. A more detailed description of the algorithms for the smooth interpolation of the weights of the kinematic primitives at the transition points between the different steps is given in [49].

For the learned parameters the system generates very natural-looking coordinated three-step sequences for total goal distances between 2.3 and 3 m, which were not included in the training data set. This is illustrated in Figure 6. When the specified goal distance exceeds this interval, the system automatically introduces additional gait steps, adapting the behavior for goal distances above 3 meters. Clips illustrating the highly flexible synthesis of multi-step sequences are shown in **Demo Movie V**⁵. Figure 7 shows the highly adaptive online replanning if the goal (drawer) jumps away while the agent is approaching it, requiring the introduction of an additional step.

3.4 *Embedding in the robot control architecture*

The algorithm described in section 3.3 generates trajectories for human-like coordinated behavioral sequences. However, these sequences are not guaranteed to result in dynamically stable behavior of the robot, and the robot just may fall due to a

⁵ <http://tinyurl.com/gktjxre>




Fig. 6 Two synthesized behaviors for two conditions with different initial distance of the character from the drawer. Both distances were not present in the training data set. Adopted from [49].

loss of dynamic balance. To solve this problem, we integrated the described algorithm for the online planning of multi-step-sequences with the control architecture of the HRP-2 robot that ensures the dynamic feasibility of the executed behaviors. An overview of the developed architecture is given in Fig. 8. The online planning module is called 'Kinematic Pattern Synthesis' in the figure. The planned gait cycle trajectory is transmitted to a Walking Pattern Generator (WPG) that is based on model predictive control, which computes the foot placements and the trajectory of the Zero Moment Point (ZMP) for the current step from the desired Center of Mass (CoM) velocity and the pelvis angular velocity of the planned gait cycle [55]. It can be shown that the gait of the robot is dynamically stable if the projection of the Zero Moment Point to the ground plane is within the support polygon on the floor, which surrounds the feet that are in contact with the ground [86].

The generated preplanned CoM and ZMP trajectories are corrected, taking into account the planned upper-body joint angles by a Dynamic Filter (DF), which operates in closed-loop together with the WPG. Both, the planned CoM and ZMP trajectories, and the upper-body joint angles are then combined in an inverse kinematics module that implements 'Stack-of-Task' approach (SoT) [46, 77]. This module outputs angular trajectories for legs and upper-body and ensures that the executed behavior respects the dynamic stability constraints of the robot, at the same time




Fig. 7 Online perturbation experiment. The goal (drawer) jumps away while the agent is approaching it, requiring online replanning of the multi-step sequence. The algorithm introduces automatically an action of type 2 (short step) in order to adjust for the increased distance to the goal.

approximating the desired behavior of the upper body, as far as possible. These resulting trajectories $\mathbf{q}(t)$ can then feasibly be realized by the low-level controllers of 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 adaptive interaction in the real world. For the successful realization of the system it is important to retrain the primitives on example trajectories that are feasible for the robot, which are generated by a robot physics simulator.

3.5 Experiments on the robot

The synthesis architecture was first tested by simulating 'open-loop' control, using the OpenHRP simulator to realize a physical model of HRP-2 robot. In the open loop simulations the robot replays the training movements, but does not create on-line adapted movements with adjusted step sizes and sequences dependent on the distance of the robot from the reaching target. In the simulations, the robot starts from the 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




Fig. 8 Control architecture of HRP-2 humanoid robot. The online kinematic pattern synthesis module is linked to a Walking Pattern generator, which computes foot placements and the ZMP trajectory (see text). The Dynamic Filter corrects the walking trajectory dependent on the joint angles of the upper body. Both, gait control parameters and upper body joint angles are integrated in a generalized kinematic planning module using a Stack-of-Task (SOT) approach, which computes the control torques for the robot. The variables $[\mathbf{v}^{ref}, \boldsymbol{\omega}^{ref}]$ signify the linear and angular velocity of CoM, and $\mathbf{q}^{upperbody}$ are the upper-body joint trajectories computed from the kinematic pattern synthesis. The variables $\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}}$ are the generalized position, velocity and acceleration vectors computed by the Stack of Tasks (SoT) approach.

action a spline interpolation of pelvis angular and positional coordinates was used to change the robots state back to the parking position (introducing two additional steps on the spot). A snapshot of the executed behavior is shown in Fig. 9. Examples of full three- and four-action sequences are provided in **Demo Movie VI**⁶.

As final step of this validation, the architecture was also tested using the real HRP-2 robot (Fig. 10). The behavior could be successfully realized, maintaining the balance of the robot. Examples of the corresponding behaviors of the real HRP-2 robot are provided in **Demo Movie VI**.

After these tests of the 'open loop behavior' of the system, without an adaptation of step and reaching parameters using the movement planning algorithm, we tested the full system including such online planning in extensive simulations using the OpenHRP simulator. As result, we found that the proposed architecture really works robustly also in the case of online adaptation and replanning. In addition, we tested our architecture in comparison with a simpler machine learning-based approach, where one learns the output trajectories $\mathbf{q}(t)$ from many training examples that produce dynamically behavior of the robot, and where one tries to interpolate between them using learning techniques. It turns out that this simplistic strategy works for only a subset of the training trajectories and fails completely for the generation of adaptive behavior online planning of new adapted step sizes and reaching movements [53].

An example of the quantitative validation of the method is shown in Fig. 11, that illustrates the ground reaction forces (maximal normal force of the feet over

⁶ <http://tinyurl.com/jxwmwnt>




Fig. 9 Off-line synthesised trajectories generated with the OpenHRP simulator using a realistic physical model of the robot.

the whole action sequence, based on 30 simulations with parameters that different from the training data). The maximum admissible ground reaction force for the real HRP-2 is 800 N. The figure compares the peak forces for trajectories directly created by the WPG without approximation of human behavior, the results from the naïve machine learning approach, and the ones obtained with our method. For the synthesis methods, the figure compares the results of the reconstruction of the training trajectories, using different numbers of source functions of the anechoic mixing model, and the case with an optimum number of sources with an inference of novel step sizes and reaching distances in the closed-loop system that includes online planning. For the naïve machine learning approach except for the case of 9 source functions, the force limit of the robot gets violated. Even with this optimum number of sources, the force limit is violated when the system is operating in closed loop. Consequently, the robot falls sometimes during the execution of such behaviors [53]. Contrasting with this result, for our methods the peak ground reaction forces remain always in the feasible region, and they are extremely similar to the ones when the movement was directly planed using the WPG without training



Fig. 10 Real HRP-2 robot performing walking-reaching sequences at LAAS-CNRS.

to approximate human behavior. It is remarkable that even for the most difficult case, the closed-loop inference of adaptive behavior, the ground reaction forces do not significantly increase. Similar behavior is observed for other critical mechanical parameters, like the joint torques. (See [53] for details.) This demonstrates that the special form of the integration of the planning algorithm in the control system is critical in order to obtain dynamically feasible behavior of the robot that prevents it from falling.

4 Probabilistic model for the online synthesis of stylized reactive movements

In the last section of this chapter we describe a completely different approach for the generation of reactive complex body movements that exploits state-of-the-art Bayesian approaches in machine learning. We applied this approach in order to simulate a reactive avatar in Virtual Reality (VR) that reacts to the movements of the user with gradually controlled emotional style. Reactive motions are generated by a dynamical extension of hierarchical Gaussian process latent variable model (GPLVM). (See [81, 80] for details.) This probabilistic model includes latent variables that encode the emotional style of the executed actions, where these variable can be adjusted at run-time. We have verified by psychophysical experiments that this method generates human motion that is almost indistinguishable from real human trajectories. In addition, it allows to control precisely and continuously the emotional style of the executed actions [83, 82]. This makes the developed method interesting for many applications, including experiments in neuroscience and psychology, computer graphics, and for the realization of human-machine interactions.

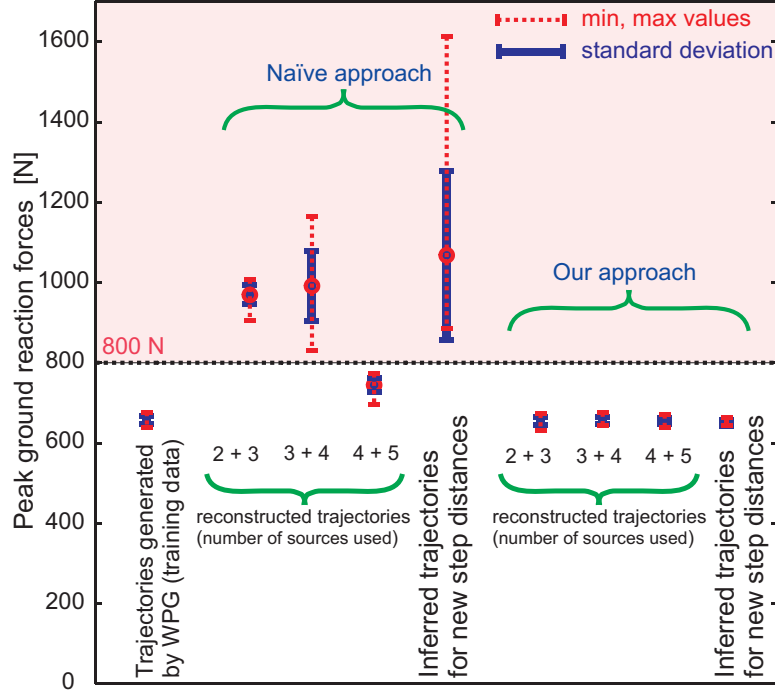


Fig. 11 Peak ground reaction forces obtained for simulated test trials. Comparison of three different synthesis methods: WPG: trajectories generated with the WPG without approximation of human behavior; naïve ML: interpolation of feasible trajectories using machine learning methods; and with our method. In addition, the figure compares the results for the resynthesis of the training trajectories with different numbers of sources, and the full closed-loop behavior with an adaptive synthesis of novel step sizes and reaching distances. (Blue error bars indicate mean and standard deviation, and red lines indicate the ranges between minimum and maximum values)

4.1 Related work

The modeling of emotional styles is a classical problem in computer graphics (e.g. [6, 90]). A variety of statistical motion models have been proposed for style interpolation [6, 30], the editing of motions styles [38], and for the analysis and synthesis of human motion data in general (e.g. [9]). However, many of these techniques result either in off-line models that cannot react in real-time to external inputs, such as other characters in the scene, or they are strongly simplified, resulting in movements that are not completely believable when compared with real human motion.

More recent approaches have tried to learn highly accurate models of human motion in an unsupervised manner from motion capture data bases. A very successful approach has been the use of Gaussian Process Latent Variable Models (GPLVMs), a nonlinear dimension reduction technique. GPLVMs have been applied in com-

puter graphics for the modeling of kinematics and motion interpolation [24], for the realization of inverse kinematics [42], and for the learning of low-dimensional dynamical models [95]. A related approach are Gaussian process dynamical models (GPDM), a method that uses the same framework for the learning on nonlinear dynamical systems that generate highly realistic human motion [88]. In our previous work [80] we introduced a dynamical mapping similar to a GPDM in a hierarchical generative model to learn the dynamics of stylized interactive movements and interpolate between them. The major problem of these models for real-time synthesis is the associated computational cost, which requires additional approximations, such as the introduction of sparsified representations, to accomplish synthesis in real-time.

4.2 Probabilistic model for interactive movements

In order to learn a generative model for two-person interactions we use motion capture data from couples of actors that executed interactive behaviors, such as handshakes or high-five movements with different emotional styles [83, 82].

The learned probabilistic model is depicted in Figure 12. It has a hierarchical structure and consists of three levels, an *agent layer* which encodes the kinematics (joint angles) of either agent, an *interaction layer* that encodes the interaction between the two agents on the level of individual time points (frames), and a *dynamics layer* that encodes a dynamic sequence of states in the interaction layer. Along the hierarchy a strong dimension reduction is realized, with a reduction from 159 joint angles to a two-dimensional latent space at the agent layer, and a further reduction from four to three dimensions in the interaction layer. The dynamics that is modeled by the dynamics layer runs in a three-dimensional state space. The whole model can be interpreted as a probabilistic graphical model, and inference techniques for such models can be applied to determine the state of the latent variables [5]. Specifically, we used a maximum-a-posteriori approximation to determine the most probable settings of the latent \mathbf{x}_t^j and \mathbf{i}_t (see Fig. 12) In the following, the individual layers are described in more detail.

The **agent layer** approximates, separately for the two agents $j \in \{1, 2\}$, a set of training trajectories by nonlinear dimensionality reduction using a GPLVM. For this purpose, we learn a nonlinear mapping from a two-dimensional latent variable \mathbf{x}_t^j and emotional style \mathbf{e} onto the 159-dimensional joint angle vectors \mathbf{y}_t^j . The nonlinear functions f_j that realize this mapping are drawn from a Gaussian Process with a composite kernel that combines a radial basis functions (RBF) kernel for the joint angle variables, and a linear kernel for an additional style variable \mathbf{e} that controls the emotional style of the movements. This defines a multi-factor model [87], where the kernel function for the GPLVM is constructed by a product of different kernel functions for motion and style. In addition, we engineered a special prior that promotes factorization of the latent variables into motion dimensions and style dimensions during learning via back constraints [40], expressing the approx-




Fig. 12 Hierarchical probabilistic model for interactive movements. The graphical model comprises three layers. At the bottom are the observable joint angle vectors \mathbf{y}_t^j of each actor $j \in \{1, 2\}$ and point in time t . The means of the \mathbf{y}_t^j are generated by the latent states \mathbf{x}_t^j of the agent layer and the emotion style variable \mathbf{e} via functions $f_j(\mathbf{x}_t^j, \mathbf{e})$ drawn from Gaussian processes. The \mathbf{x}_t^j have a much smaller dimensionality than the \mathbf{y}_t^j . The means of the \mathbf{x}_t^j of both actors are generated by a function $g(\mathbf{i}_t)$ from a yet lower-dimensional interaction-layer state \mathbf{i}_t , whose time evolution is controlled by a mapping $h(\mathbf{i}_{t-2}, \mathbf{i}_{t-1})$ in the dynamics layer. Both $g()$ and $h()$ are drawn for a Gaussian processes, too. The plates denote the assumption of replicated independent identically distributed (i.i.d.) draws across R trials per E many emotions from C couples of actors. For details, see text. Figure adopted from [82].

imately periodic nature of the movements (i.e. the avatar begins and ends a trial in approximately the same pose). This step stabilizes the highly ill-posed factorization problem in a way that results in relatively simple manifolds representing the data in the latent space. Mathematically, the mapping from latent into joint angle space is given by the equation (time index t omitted):

$$\mathbf{y}^j = f_j(\mathbf{x}^j, \mathbf{e}) + \varepsilon_j, \quad f_j(\mathbf{x}^j, \mathbf{e}) \sim GP(\mathbf{0}, k_V^j([\mathbf{x}^j, \mathbf{e}], [(\mathbf{x}')^j, \mathbf{e}'])), \quad j \in \{1, 2\}, \quad (3)$$

where k_V^j is an appropriate kernel function, and where ε_j is isotropic Gaussian noise. The linear kernel component makes it possible to morph easily along this dimension. A GPLVM can be seen as a nonlinear extension of probabilistic dual PCA that learns a low-dimensional latent space and a mapping from this space to the data space. Unlike PCA, this mapping is nonlinear. It turned out that already two latent dimensions plus a dimension per emotion were sufficient to achieve a highly accurate approximation of the data. See [80, 83] for further details.

The latent variable of the agent layer represents the behavior of the two agents as a trajectory in a four-dimensional space. The dimensionality of this high-dimensional trajectory is further reduced in the **interaction layer**, which learns a mapping from a three-dimensional latent space (variable \mathbf{i}) onto this trajectory. This mapping is again

realized by a GPLVM that is trained with the aforementioned data basis. Consequently, the individual points of the latent space of the interaction layer are mapped by the two lower layers of the model onto a pair of postures of both agents for each moment of the evolving interaction. The temporal evolution of the interaction corresponds to a three-dimensional trajectory \mathbf{i}_t in the latent space of this layer.

This time course is modeled by the **dynamics layer** by learning of an autonomous dynamical system that generates this trajectory as stable solution. This dynamical system was modeled using a Gaussian Process Dynamical model (GPDM) [88], which can be interpreted as a nonlinear generalisation of an Autoregressive (AR) model in time series analysis. Mathematically, this model is defined by the equations:

$$\begin{aligned} \mathbf{i}_t &= h(\mathbf{i}_{t-1}, \mathbf{i}_{t-2}) + \xi, \\ h(\mathbf{i}_{t-1}, \mathbf{i}_{t-2}) &\sim GP(\mathbf{0}, k_h([\mathbf{i}_{t-1}, \mathbf{i}_{t-2}], [\mathbf{i}_{t-1}, \mathbf{i}_{t-2}])) \end{aligned} \quad (4)$$

where the function h is drawn from a Gaussian process with the kernel function k_h , and where ξ signifies Gaussian noise.

The described model allows the generation pair interactions with controllable emotional style by variation of the emotion parameter \mathbf{e} . The accuracy of the synthesized movements was validated in psychophysical experiments that show that the generated motion is perceptually almost indistinguishable from real motion capture data from human pair interactions [80].

4.3 Inference for the generation of reactive movements

The described generative probabilistic model can be exploited for the simulation of the interactive behaviors of reactive human virtual agents, who react to the movements of a real human user in a human-like fashion. For this purpose, we exploited the fact that probabilistic generative models can be 'inverted' by conditioning. More precisely, using standard techniques [5], such models allow to make inference of unobserved nodes in the network, dependent on given (observed) information on a subset of the nodes. For this application, we strongly simplified the model for one agent and modeled only its hand position \mathbf{y}^2 . We then replaced the corresponding random variables by online motion capture data from the user of the system. It is then possible to infer the distributions of all other nodes in the network by conditioning on the available hand position information. Maximizing the joint probability of the latent variables and the observed hand trajectory, one can then in principle find the most probable values of all other variables in the probabilistic network. Specifically, this allows to determine the most probable joint angles of the other agent (agent 1) that correspond best to the observed trajectories of the user. The resulting way of inferring a likely posture sequence from a generative probabilistic model, using the observations as constraints to find the most probable trajectory can be interpreted as a special form of 'style-based inverse kinematics' [24].

The straightforward practical implementation of this idea suffers from a very high computational cost, like many naïve Bayesian machine learning approaches. The inversion of the described probabilistic model using straight-forward methods results in a system that is way too slow for real-time synthesis of reactive movements. In order to solve this problem, we implemented in addition the following two approximations: (1) Direct mappings from the data variables to the latent variables were explicitly learned and modeled by Gaussian process regression. This is much faster than to determine the latent variables by conditioning on the observed input, since we only need to evaluate a Gaussian process prediction. We are learning these direct mappings from latent/observed pairs during an off-line training phase. (2) To learn the model from large data sets we applied sparse approximations techniques, which approximate the data manifolds in latent space by a small set of inducing points, resulting in much fewer effective model parameters and computational cost for the evaluation of the kernel-dependent functions: This approximation makes computational cost per learning step effectively linear in the number of data-points, as opposed to cubic for the exact solution [39]. With these two additional approximations, which simplify both learning and latent inference, a speed increase of more than factor 100 was obtained, making the resulting architecture suitable for the real-time synthesis of interactive behaviors.

4.4 Application results

The proposed architecture was tested with different types of interactive human movements. One data set consisted of 'high fives' with four emotional styles (neutral, happy, angry and sad) executed by different actors. A total of 105 motion-captured trajectories was learned, which were performed on a imaginary 3×3 spatial grid for hand contact positions.

The synthesized movements look so natural that human observers were not able to distinguish them from original motion-captured trajectories, the method passes effectively the 'Turing test of computer graphics' [80]. Demonstration movies are also provided as part of **Demo Movie VII**⁷.

For testing of the reactive movement generation the architecture was embedded in a virtual reality setup that is described in detail in [83]. The underlying animation pipeline integrated a Vicon (Nexus) motion capture system, the game engine Ogre 3D, and the proposed learning-based architecture. The reconstruction took place in real-time with Ogre rendering at a frame rate of 68 fps. The functioning of this system, including the variation of emotional style, are also shown in **Demo Movie VII**. A second data set for which the novel architecture was tested were handshakes, which were executed with different emotional styles. Snapshots from the generated stylized motions are shown in Figure 13. Movies can be found in **Demo Movie VII**.

⁷ <http://tinyurl.com/j3d9xtk>

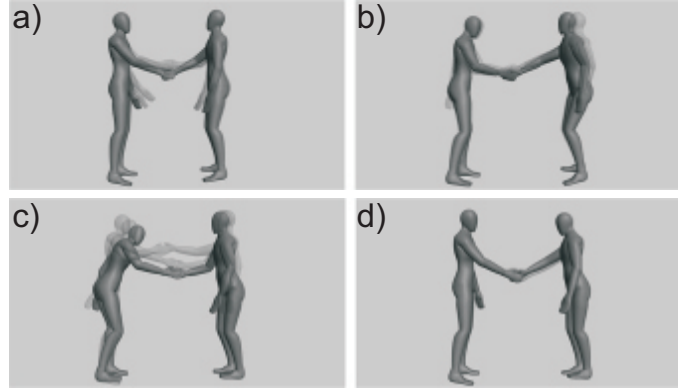


Fig. 13 Motion sequences of synthesized emotional handshakes. a) neutral, b) angry, c) happy and d) sad. Different emotions are associated with different postures and ranges of joint movements.

5 Conclusions and future work

In this chapter we reviewed two approaches that approximate complex human full-body motion by structured models that can be embedded in architectures that require a real-time synthesis of complex human motion. Opposed to the many available methods for the off-line synthesis of human motion, online synthesis requires an embedding of the synthesis process into dynamical systems that can be integrated in control architectures. We have shown two different approaches how such models can be generated, exploiting concepts from machine learning and the theory of nonlinear dynamical systems.

The first approach approximated human joint angle trajectories by highly compact (anechoic) mixture models. The resulting source functions were then synthesized on-line by mapping the solutions of canonical nonlinear dynamical systems onto them, defining a special form of dynamic movement primitive (DMP). This allowed to synthesize highly complex coordinated full-body movements by networks of dynamically coupled dynamic primitives. We showed that one can systematically design the stability of such 'primitive networks' exploiting tools from Contraction Theory. We also demonstrated how this framework can be used to model complex coordinated behaviors of individual agents, and of whole crowds of interacting individuals. In addition, we demonstrated that this method is suitable for the online planning of multi-step sequences which are coordinated with arm movements in humanoid robots in real-time, accomplishing dynamically feasible behaviors on a real humanoid robot including the control of dynamically stable walking. The advantage of the chosen approach is that it is computationally more efficient than the synthesis of the same behaviors using straight-forward optimal control approaches, since the computational complexity of the underlying optimization problems with the presently available computational power would not permit an adaptive planning of such multi-step sequences in real-time.

The second approach for the learning of such real-time capable synthesis models was based on established methods in Bayesian machine learning. We demonstrated that an approach that learns a hierarchical ('deep') architecture by combining GPLVMs and GPDMs was suitable for the synthesis of highly natural-looking human movements. In addition, the resulting probabilistic graphical model can be inverted (Bayesian model inversion), i.e. conditioned on observable data. This allowed us to learn interactions from pairs of actors, and to use these learned models then to generate online the maximally probable reactive behavior of a virtual agent that responds directly to a human, whose movements were motion-captured online.

To make this system working in real-time required a substantial amount of engineering work, due to the high computational cost of the chosen Bayesian machine learning approach. This shows that it is a non-trivial step to make such methods work in real-world applications, especially with real-time constraints. It seems likely that it will be even less trivial to embed such methods in complex control architectures, such as the one shown in Fig. 8. This illustrates limitations of these popular approaches which cannot be ignored when dealing with real technical control systems. An advantage of the described probabilistic architectures is that they can be integrated with other probabilistic systems, e.g. in computer vision or pattern analysis.

Another interesting challenge is to link the discussed hierarchical probabilistic architectures to spatial movement primitives that, similar to the source functions discussed in section 2, allow the modeling of separately coordinated clusters of degrees of freedom. First work in this direction has been successfully performed [85, 84], and it seems an exciting avenue for future research to see how far such approaches can be extended in the context of real-world problems.

Acknowledgements

The work supported by EC FP7 under grant agreements FP7-611909 (Korobot), H2020 ICT-644727 (CogIMon), FP7-604102 (HBP), PITN-GA-011-290011 (ABC), DFG GI 305/4-1, DFG GZ: KA 1258/15-1, DFG IRTG-GRK 1901 'The brain in action', BMBF, FKZ: 01GQ1002A, and DFG SFB/TRR 135 Cardinal Mechanisms of Perception, project C06.

References

1. A. Ajoudani, J. Lee, A. Rocchi, M. Ferrati, E.M. Hoffman, A. Settini, D.G. Caldwell, A. Bicchi, and N.G. Tsagarakis. A manipulation framework for compliant humanoid COMAN: Application to a valve turning task. In *14th IEEE-RAS Int. Conf. on Humanoid Robots (Humanoids)*, 2014, pages 664–670, 2014.
2. O. Arikan, D.A. Forsyth, and J. F. O'Brien. Motion synthesis from annotations. *ACM Trans. on Graphics, SIGGRAPH '03*, 22(3):402–408, 2003.

3. C. G. Atkeson, A. W. Moore, and S. Schaal. Locally weighted learning. *A.I. Review*, 11:11–73, 1997.
4. N.A. Bernstein. *The coordination and regulation of movements*. Pergamon Press, N.Y., Oxford, 1967.
5. Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2007.
6. M. Brand and A. Hertzmann. Style machines. *Proc. of SIGGRAPH Conf.*, pages 183–192, 2000.
7. M. Brandao, L. Jamone, P. Kryczka, N. Endo, K. Hashimoto, and A. Takanishi. 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)*, pages 28–33, 2013.
8. H. Carnahan, B. J. McFadyen, D. L. Cockell, and A. H. Halverson. The combined control of locomotion and prehension. *Neurosci. Res. Comm.*, 19:91–100, 1996.
9. J. Chai and J. K. Hodgins. Performance animation from low-dimensional control signals. *ACM Trans. on Graphics, SIGGRAPH '05*, 24(3):686–696, 2005.
10. C.-C. Chang and C.-J. Lin. *LIBSVM: a library for support vector machines*, 2001. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
11. E. Chiovetto, A. d’Avella, D. Endres, and M. A. Giese. A unifying algorithm for the identification of kinematic and electromyographic motor primitives. *Bernstein Conference*, 2013.
12. E. Chiovetto and M. A. Giese. Kinematics of the coordination of pointing during locomotion. *Plos One*, 8(11), 2013.
13. W. Daamen and S. P. Hoogendoorn. Controlled experiments to derive walking behaviour. *European Journal of Transport and Infrastructure Research*, 3(1):39–59, 2003.
14. A. d’Avella and E. Bizzi. Shared and specific muscle synergies in neural motor behaviours. *Proc. Natl. Acad. Sci. U.S.A.*, 102(8):3076–3081, 2005.
15. S. Degallier, L. Righetti, S. Gay, and A. J. Ijspeert. Towards simple control for complex, autonomous robotic applications: combining discrete and rhythmic motor primitives. *Auton. Robots*, pages 155–181, 2011.
16. A. W. Feng, Y. Xu, and A. Shapiro. An example-based motion synthesis technique for locomotion and object manipulation. *Proc. of ACM SIGGRAPH I3D*, pages 95–102, 2012.
17. T. Flash and B. Hochner. Motor primitives in vertebrates and invertebrates. *Curr. Opin. Neurobiol.*, 15(6):660–666, 2005.
18. A. Fod, M.J. Mataric, and O.C. Jenkins. Automated derivation of primitives for movement classification. *Autonomous Robots*, 12(1):39–54, 2002.
19. A. Gams, B. Nemec, L. Zlajpah, M. Wächter, A. J. Ijspeert, T. Asfour, and A. Ude. 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)*, pages 5629–5635, 2013.
20. M. Gienger, M. Toussaint, and C. Goerick. Whole-body motion planning building blocks for intelligent systems. In K. Harada, editor, *Motion Planning for Humanoid Robots*, pages 67–98. Springer, 2010.
21. M. A. Giese, A. Mukovskiy, A. Park, L. Omlor, and J. J. E. Slotine. Real-time synthesis of body movements based on learned primitives. In D. Cremers et al., editor, *Stat. and Geom. Appr. to Vis. Mot. Anal., LNCS5604*, pages 107–127. Springer, 2009.
22. M. Gleicher. Motion path editing. *Proc. of 2001 ACM Symp. on Interactive 3D Graphics*, pages 195–202, 2001.
23. M. Gleicher, H. J. Shin, L. Kovar, and A. Jepsen. Snap-together motion: Assembling run-time animation. *ACM Trans. on Graphics, SIGGRAPH '03*, 22(3):702–702, 2003.
24. K. Grochow, S.L. Martin, A. Hertzmann, and Z. Popovic. Style-based inverse kinematics. *ACM Trans. Graphics*, 23(3):522–531, 2004.
25. D. Helbing, P. Molnár, I. J. Farkas, and K. Bolay. Self-organizing pedestrian movement. *Environment and Planning B: Planning and Design*, 28:361–383, 2001.
26. A. Herdt, H. Diedam, P.-B. Wieber, D. Dimitrov, K. Mombaur, and M. Diehl. Online walking motion generation with automatic foot step placement. *Adv. Robotics*, 24(5–6):719–737, 2010.

27. E. Hsu, K. Pulli, and J. Popovic. Style translation for human motion. *ACM Trans. on Graphics*, 24(3):1082–1089, 2005.
28. Y. Huang and M. Kallmann. Planning motions for virtual demonstrators. In *Intelligent Virtual Agents*, pages 190–203. Springer, 2014.
29. A. J. Ijspeert, J. Nakanishi, H. Hoffmann, P. Pastor, and S. Schaal. Dynamical movement primitives: Learning attractor models for motor behaviors. *Neural Comp.*, 25(2):328–373, 2013.
30. L. Ikemoto, O. Arikan, and D.A. Forsyth. Generalizing motion edits with Gaussian processes. *ACM Trans. Graph.*, 28(1), 2009.
31. Y. Ivanenko, R. Poppele, and F. Lacquaniti. Five basic muscle activation patterns account for muscle activity during human locomotion. *Journal of Physiology*, 556:267–282, 2004.
32. S. Kajita, F. Kanehiro, K. Kaneko, K. Fujiwara, K. Harada, K. Yokoi, and H. Hirukawa. Biped walking pattern generation by using preview control of zero-moment point. *Proc. of Int. Conf. on Robotics and Autom.*, pages 1620–1626, 2003.
33. J. Koschorreck and K. Mombaur. Modeling and optimal control of human platform diving with somersaults and twists. *Optimization and Engineering*, 13(1):29–56, 2012.
34. L. Kovar, M. Gleicher, and F. Pighin. Motion graphs. *Proc. of SIGGRAPH 2002*, pages 473–482, 2002.
35. S. Kuindersma, R. Deits, M. Fallon, A. Valenzuela, H. Dai, F. Permenter, T. Koolen, P. Marion, and R. Tedrake. Optimization-based locomotion planning, estimation, and control design for the Atlas humanoid robot. *Autonomous Robots*, pages 1–27, 2015.
36. T. Kwon, K. H. Lee, J. Lee, and S. Takahashi. Group motion editing. *ACM Trans. on Graphics, SIGGRAPH 2008*, 27(3):80–87, 2008.
37. W. M. Land, D. A. Rosenbaum, S. Seegelke, and T. Schack. Whole-body posture planning in anticipation of a manual prehension task: Prospective and retrospective effects. *Acta Psychologica*, 114:298–307, 2013.
38. M. Lau, Z. Bar-Joseph, and J. Kuffner. Modeling spatial and temporal variation in motion data. *ACM Trans. Graph.*, 28(5):Art.No.171, 2009.
39. N.D. Lawrence. Learning for larger datasets with the Gaussian process latent variable model. *J. Mach. Learn. Res. - Proc. Track*, 2:243–250, 2007.
40. N.D. Lawrence and R. Court. Local distance preservation in the gp-lvm through back constraints. In *ICML*, pages 513–520, 2006.
41. A. Lerner, E. Fitusi, Y. Chrysanthou, and D. Cohen-Or. Fitting behaviors to pedestrian simulations. *Proc. Eurographics/ACM SIGGRAPH Symp. on Comp. Animation*, pages 199–208, 2009.
42. S. Levine, J. M. Wang, A. Haraux, Z. Popović, and V. Koltun. Continuous character control with low-dimensional embeddings. *ACM SIGGRAPH 2012. ACM Trans. on Graphics, Art.No.28*, 31(4), 2012.
43. Y. Li, T. Wang, and H. Y. Shum. Motion texture: A two level statistical model for character motion synthesis. *Proc. of SIGGRAPH 2002*, pages 465–472, 2002.
44. G. Liu, M. Xu, Z. Pan, and A. El Rhalibi. Human motion generation with multifactor models. *J. of Computer Animation and Virtual Worlds*, 22(4):351–359, 2011.
45. W. Lohmiller and J.J.E. Slotine. On contraction analysis for nonlinear systems. *Automatica*, 34(6):683–696, 1998.
46. N. Mansard, O. Stasse, P. Evrard, and A. Kheddar. A versatile generalized inverted kinematics implementation for collaborative working humanoid robots: The stack of tasks. In *Proc. of Int. Conf. on Adv. Robotics (ICAR)*, number art.119, 2009.
47. R.G. Marteniuk and C. P. Bertram. Contributions of gait and trunk movement to prehension: Perspectives from world- and body centered coordinates. *Motor Control*, 5:151–164, 2001.
48. M. Mühlig, M. Gienger, and J.J. Steil. Human-robot interaction for learning and adaptation of object movements. In *Proc. of IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS 2010)*, pages 4901–4907, 2010.
49. A. Mukovskiy, W. Land, T. Schack, and M.A. Giese. Modeling of predictive human movement coordination patterns for applications in computer graphics. *J. of WSCG*, 23(2):139–146, 2015.

50. A. Mukovskiy, A.-N. Park, L. Omlor, J.-J. Slotine, and M.A. Giese. Self-organization of character behavior by mixing of learned movement primitives. *In Proc. of the 13th Fall Workshop on Vision, Modeling, and Visualization (VMV)*, pages 121–130, 2008.
51. A. Mukovskiy, J.J.E. Slotine, and M.A. Giese. Analysis and design of the dynamical stability of collective behavior in crowds. *J. of WSCG*, 19(1-3):69–76, 2011.
52. A. Mukovskiy, J.J.E. Slotine, and M.A. Giese. Dynamically stable control of articulated crowds. *J. of Comput. Science*, 4(4):304–310, 2013.
53. A. Mukovskiy, C. Vassallo, M. Naveau, O. Stasse, P. Souères, and M.A. Giese. Adaptive synthesis of dynamically feasible full-body movements for the humanoid robot hrp-2 by flexible combination of learned dynamic movement primitives. *submitted to Robotics and Autonomous Systems J. of Comput. Science*, 2016.
54. R. Narain, A. Golas, S. Curtis, and M. Lin. Aggregate dynamics for dense crowd simulation. *ACM Trans. on Graphics, Art.122*, 28(5):1–8, 2009.
55. M. Naveau, M. Kudruss, O. Stasse, C. Kirches, K. Mombaur, and P. Souères. A reactive walking pattern generator based on nonlinear model predictive control. *IEEE Robotics and Automation Letters*, 2016, in press.
56. L. Omlor and M. A. Giese. Anechoic blind source separation using wigner marginals. *J. of Machine Learning Res.*, 12:1111–1148, 2011.
57. D.A. Paley, N.E. Leonard, R. Sepulchre, D. Grunbaum, and J.K. Parrish. Oscillator models and collective motion: Spatial patterns in the dynamics of engineered and biological networks. *IEEE Control Systems Magazine*, 27:89–105, 2007.
58. S. Paris, J. Pettré, and Donikian S. Pedestrian reactive navigation for crowd simulation: a predictive approach. *Proc. Eurographics 2007*, 26(3):665–674, 2007.
59. A. Park, A. Mukovskiy, L. Omlor, and M. A. Giese. Self organized character animation based on learned synergies from full-body motion capture data. *In Proc. Int. Conf. on Cognitive Systems, CogSys 2008*, 2008.
60. A. Park, A. Mukovskiy, L. Omlor, and M. A. Giese. Synthesis of character behaviour by dynamic interaction of synergies learned from motion capture data. *The 16-th Int. Conf in Central Europe on Comp. Graphics, Visualization and Computer Vision'2008, WSCG'08*, pages 9–16, 2008.
61. A. Park, A. Mukovskiy, J. J. E. Slotine, and M. A. Giese. Design of dynamical stability properties in character animation. *Proc. of VRIPHYS 09*, pages 85–94, 2009.
62. S.I. Park, H.J. Shin, and S.Y. Shin. On-line locomotion generation based on motion blending. *Proc. of the 2002 ACM SIGGRAPH/Eurographics Symp. on Comp. Animation*, pages 105–111, 2002.
63. N. Pelechano, J. M. Allbeck, and N. I. Badler. Controlling individual agents in high-density crowd simulation. *Proc. Eurographics/ ACM SIGGRAPH Symp. on Comp. Animation*, pages 99–108, 2007.
64. Q.C. Pham and J.J.E. Slotine. Stable concurrent synchronization in dynamic system networks. *Neural Networks*, 20(3):62–77, 2007.
65. C.L. Roether, L. Omlor, A. Christensen, and M. A. Giese. Critical features for the perception of emotion from gait. *Journal of Vision*, 6, 2009.
66. C. Rose, M. Cohen, and B. Bodenheimer. Verbs and adverbs: Multidimensional motion interpolation using radial basis functions. *IEEE Comp. Graphics and Appl.*, 18(5):32–40, 1998.
67. C. Rose, B. Guenter, B. Bodenheimer, and M. Cohen. Efficient generation of motion transitions using spacetime constraints. *Int. Conf. on Comp. Graph. and Interactive Techniques, Proc. ACM SIGGRAPH'96*, 30:147–154, 1996.
68. D. A. Rosenbaum. Reaching while walking: reaching distance costs more than walking distance. *Psych. Bull. Rev.*, 15:1100–1104, 2008.
69. D. A. Rosenbaum, R. G. Cohen, S. A. Jax, D. J. Weiss, and R. van der Wel. The problem of serial order in behavior: Lashley's legacy. *Human Movement Science*, 26(4):525–554, 2007. Europ. Workshop on Mov. Sci., 2007.
70. A. Safonova, J. Hodgins, and N. Pollard. Synthesizing physically realistic human motion in low-dimensional, behavior-specific spaces. *ACM Trans. on Graphics*, 23(3):514–521, 2004.

71. M. Santello, M. Flanders, and J.F. Soechting. Postural hand synergies for tool use. *Journal of Neuroscience*, 18(23):10105–15, 1998.
72. L. Scardovi and R. Sepulchre. Collective optimization over average quantities. *Proc. of the 45th IEEE Conf. on Decision and Control, San Diego, California*, pages 3369–3374, 2006.
73. S. Schaal, S. Kotosaka, and D. Sternad. Nonlinear dynamical systems as movement primitives. In *Proc. of 1st IEEE-RAS Int. Conf. on Humanoid Robots, Humanoids, 2000*, pages 117–124, Berlin, 2000. Springer-Verlag.
74. G. Schöner, M. Dose, and C. Engels. Dynamics of behavior: Theory and applications for autonomous robot architectures. *Robotics and Autonomous Systems*, 16(2–4):213–245, 1995.
75. A. Shoulson, N. Marshak, M. Kapadia, and N.I. Badler. Adapt: The agent development and prototyping testbed. *IEEE Trans. on Visualiz. and Comp. Graphics (TVCG)*, 99:1–14, 2014.
76. M. Sreenivasa, P. Souères, and J.-P. Laumond. 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, 2012.
77. O. Stasse. *Habilitation Thesis*. Paul Sabatier University, CNRS, Toulouse, 2013.
78. O. Stasse, B. Verelst, A. Davison, N. Mansard, F. Saidi, B. Vanderborght, C. Esteves, and Yokoi K. Integrating walking and vision to increase humanoid autonomy. *Int. Journal of Humanoid Robotics, Special Issue on Cognitive Humanoid Robots*, 5:287–310, 2008.
79. M. Taïx, M.T. Tran, and E. Souères, P. Guigon. Generating human-like reaching movements with a humanoid robot: A computational approach. *J. of Comput. Sci.*, 4:269–284, 2013.
80. N. Taubert, A. Christensen, D. Endres, and M.A. Giese. Online simulation of emotional interactive behaviors with hierarchical Gaussian Process Dynamical Models. In *Proc. SAP’12*, pages 25–32, New York, New York, USA, 2012. ACM Press.
81. N. Taubert, D. Endres, A. Christensen, and M.A. Giese. Shaking hands in latent space: modeling emotional interactions with Gaussian Process Latent Variable Models. In S. Edelkamp and J. Bach, editors, *Proc. KI 2011: Advances in Artificial Intelligence, LNAI*, pages 330–334. Springer, 2011.
82. N. Taubert, D. Endres, and M.A. Giese. Reactive virtual reality avatar with controllable emotional style based on hierarchical gaussian process dynamical models. In *In Proc. ICANN 2014*, page Art.No.25, 2014.
83. N. Taubert, M. Löffler, N. Ludolph, A. Christensen, D. Endres, and M.A. Giese. A virtual reality setup for controllable, stylized real-time interactions between humans and avatars with sparse Gaussian process dynamical models. In *Proc. SAP’13*, pages 41–44, 2013.
84. D. Velychko and D. Endres. The variational Gaussian process dynamical model. In *Proceedings of the Workshop on Advances in Approximate Bayesian Inference*, pages 1–6. NIPS, Montreal, Canada, 2015.
85. D. Velychko, D. Endres, N. Taubert, and M.A. Giese. Coupling Gaussian process dynamical models with product-of-experts kernels. In *Proceedings of the 24th International Conference on Artificial Neural Networks, LNCS 8681*, pages 603–610. Springer, 2014.
86. M. Vukobratović and Yu. Stepanenko. On the stability of anthropomorphic systems. *Mathematical Biosciences*, 15:1–37, 1972.
87. J.M. Wang, D.J. Fleet, and A. Hertzmann. Multifactor Gaussian process models for style-content separation. In *Proc. ICML*, 2007.
88. J.M. Wang, D.J. Fleet, and A. Hertzmann. Gaussian Process Dynamical Models for human motion. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 30(2):283–298, 2008.
89. W. Wang and J.J.E. Slotine. On partial contraction analysis for coupled nonlinear oscillators. *Biological Cybernetics*, 92(1):38–53, 2005.
90. Y. Wang, Z.-Q. Liu, and L.-Z. Zhou. Learning style-directed dynamics of human motion for automatic motion synthesis. In *Proc. IEEE Int. Conf. on Systems, Man, and Cybernetics*, pages 4428–4433, 2006.
91. W.H. Warren. The dynamics of perception and action. *Psych. Rev.*, 113(2):358–389, 2006.
92. M. Weigelt and T. Schack. The development of end-state comfort planning in preschool children. *Exper. Psych.*, 57(6):476–782, 2010.
93. A.P. Witkin and Z. Popović. Motion warping. *Proc. ACM SIGGRAPH’95*, 29:105–108, 1995.

94. K. Yamane and Y. Nakamura. Dynamics filter - concept and implementation of on-line motion generator for human figures. In *Proc. of IEEE Int. Conf. on Robotics and Autom.*, pages 688–695, 2000.
95. Y. Ye and C.K. Liu. Synthesis of responsive motion using a dynamic model. *Computer Graphics Forum (Proc. of Eurographics)*, 29(2):555–562, 2010.
96. E. Yoshida, A. Mallet, F. Lamiriaux, O. Kanoun, O. Stasse, M. Poirier, P-F. Dominey, J.-P. Laumond, and K. Yokoi. 'Give me the Purple Ball' – he said to HRP-2 N.14. In *Proc. of IEEE-RAS Int. Conf. on Humanoid Robots (Humanoids'07)*, 2007.