

Modeling of Predictive Human Movement Coordination Patterns for Applications in Computer Graphics

Albert Mukovskiy
Section for
Computational
Sensomotrics,
Department of
Cognitive Neurology,
Hertie Institute for
Clinical Brain Research
& Centre for Integrative
Neuroscience,
University Clinic,
Tübingen, Germany
albert.mukovskiy@medizin.uni-
tuebingen.de

William M. Land
Department of
Neurocognition and
Action, Bielefeld
University, Germany.
College of Education
and Human
Development, The
University of Texas at
San Antonio, One UTSA
Circle, San Antonio, TX
78249, USA.
william.land@utsa.edu

Thomas Schack
CITEC Center of
Excellence "Cognitive
Interaction Technology",
CoR Lab Research
Institute of Cognition and
Robotics, Department of
Neurocognition and
Action, Bielefeld
University, Germany
thomas.schack@uni-
bielefeld.de

Martin A. Giese
Section for
Computational
Sensomotrics,
Department of
Cognitive Neurology,
Hertie Institute for
Clinical Brain Research
& Centre for Integrative
Neuroscience,
University Clinic,
Tübingen, Germany
martin.giese@uni-
tuebingen.de

Abstract

The planning of human body movements is highly predictive. Within a sequence of actions, the anticipation of a final task goal modulates the individual actions within the overall pattern of motion. An example is a sequence of steps, which is coordinated with the grasping of an object at the end of the step sequence. Opposed to this property of natural human movements, real-time animation systems in computer graphics often model complex activities by a sequential concatenation of individual pre-stored movements, where only the movement before accomplishing the goal is adapted. We present a learning-based technique that models the highly adaptive predictive movement coordination in humans, illustrated for the example of the coordination of walking and reaching. The proposed system for the real-time synthesis of human movements models complex activities by a sequential concatenation of movements, which are approximated by the superposition of kinematic primitives that have been learned from trajectory data by anechoic demixing, using a step-wise regression approach. The kinematic primitives are then approximated by stable solutions of nonlinear dynamical systems (dynamic primitives) that can be embedded in control architectures. We present a control architecture that generates highly adaptive predictive full-body movements for reaching while walking with highly human-like appearance. We demonstrate that the generated behavior is highly robust, even in presence of strong perturbations that require the insertion of additional steps online in order to accomplish the desired task.

Keywords

computer animation, movement primitives, motor coordination, action sequences, prediction.

1 INTRODUCTION

A central problem in computer animation is the online-synthesis of complex behaviors that consist of sequences of individual actions, which have to adapt to continuously changing environmental constraints. An example is the online planning of coordinated walking and reaching, when the position of the reaching goal is dynamically changing.

A prominent approach for the solution of this problem in computer graphics is the adaptive inter-

polation between motion-captured example actions [WP95, GSKJ03, AFO03]. Other approaches are based on learned low-dimensional parameterizations of whole body motion, which are embedded in mathematical frameworks for the online generation of motion (e.g. [HPP05, SHP04, RCB98, WFH08, LWS02]). Several methods have been proposed that segment action streams into individual actions, where models for the individual actions are adapted online in order to fulfill additional constraints, such as obstacle avoidance or the correct positioning of end-effectors ([KGP02, RGB96, PSS02]). The dependencies between constraints in such action sequences have been recently exploited to generate more realistic animations. In [FXS12] captured motion examples are blended according to a prioritized "stack of controllers". In [SMKB14] the instantaneous blending weights of controllers are pre-specified differently for different body parts involved in the current action and

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

the priority of the different controllers is governed by their sequential order. In [HK14] the synthesis of locomotion plus arm pointing at the last step is carried out by blending of captured actions determining the weights by "inverse blending optimization". In this study arm pointing was blended with the arm swinging motion of the last step. The choice of the arm pointing primitives depended on the gait phase, according to an empirical rule introduced by authors.

Physics-based animation is another approach for the online generation of motion (e.g. [ST05, FP03]). Complex action sequences are segmented into individual actions, which are characterized by solutions of optimization problems, derived from mechanics and additional constraints (contact, friction, or specified via-points) ([AMJ07, LHP05, MLPP09]). While these approaches generate highly adaptive behavior for individual actions, the problem to generate natural-looking transitions between the individual actions is non-trivial. As consequence, artifacts (e.g. hesitation, jerky movement) can emerge at transition points, (e.g. [WZ10]).

Opposed to these approaches skilled human motor behavior has been shown to be highly predictive. Within complex activities, action goals and the associated constraints influence actions that appear already a long time before the constraint within the behavioral stream, and thus allows the generation of smooth and optimized behaviors over complex action sequences. This was investigated, for example, in a recent study on the coordination of walking and reaching. Human subjects had to walk towards a drawer and to grasp an object, which was located at different positions in the drawer. Humans optimized their behavior already significantly before object contact, consistent with the hypothesis of maximum end-state comfort during the reaching action [WS10, Ros08], and steps prior to the reaching were modulated in order to accomplish the goal.

Whole body movements of humans and animals are organized in terms of muscle synergies or movement primitives [Ber67, FH05]. 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 [CG13, CMCH96, Ros08, MB01]. The realism and human-likeness of synthesized movements in robotics and computer graphics can be improved by taking such biological constraints into account [FMJ02].

We present a learning-based framework that makes some of these properties applicable for realtime animation in computer graphics. The underlying architecture is simple and approximates complex full-body movements by dynamic movement primitives

that are formulated in terms of nonlinear dynamical systems [GMP⁺09, PMSG09]. These primitives are constructed from kinematic primitives, that are learned from trajectory sets by anechoic demixing in an unsupervised manner. Similar to the related approaches in robotics [GRIL08, BRI06], the method generates complex movements by the combination of a small number of learned dynamical movement primitives [OG11, GMP⁺09]. We demonstrate this approach by the highly adaptive online generation of multi-step sequences with coordinated arm movements.

The paper is structured as follows: After the description of the animation system in section 2, we present some example results section 3, followed by a conclusion.

2 SYSTEM ARCHITECTURE

Our work is based on motion capture data from a single human subject performing a drawer opening task. In the following, this data set is described briefly. Then the different key elements of the proposed algorithm are introduced: movement generation by dynamic primitives, modeling of coordination by step-wise regression, and the algorithms for online blending and control.

2.1 Motion capture data

Our system was based on motion capture data from a single human subject that executed a drawer opening task, walking towards a drawer and then reaching for an object in the drawer. The distance of the subject from the drawer and the position of the object was varied [LRSS13] (Fig. 1). These training sequences consisted of three subsequent actions or movements: 1) a normal walking step; 2) a shortened step with the left-hand starting to reach towards the drawer. This step showed a high degree of adaptability, and was typically adjusted in order to create an optimum distance from the drawer (maximum comfort) for the reaching movement during the last action; 3) the drawer opening and the reaching of the object while standing. The object position in the drawer was indicated to the participants at the beginning of each trial. (See [LRSS13] for further details). (See video [Demo¹].)

The analysis of the distances between the pelvis and the drawer or the object in these action sequences reveals the predictive nature of human movement planning, as shown in Fig. 2 where the distances ordered according to the initial walking distance to the drawer. While the length of the first step and the distance from the drawer in the last step are relatively constant, a major distance adjustment is made in the second step.

¹ www.uni-tuebingen.de/uni/knv/ar/avi/wscg15/v1.avi




Figure 1: Illustration of the human behavior. The figure illustrates important intermediate postures (normal walking step, step with initiation of reaching, standing while drawer opening, and object reaching).




Figure 2: Predictive planning in 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. Mainly the second action is adjusted as function of the initial distance from the goal.

The length of the first step is not significantly correlated with the initial distance to the drawer (linear regression: $R^2 = 0.08, p = 0.429$), while the correlations with the distance to the drawer after first step, and the length of the second step are highly significant ($R^2 = 0.95, p = 1.4 \cdot 10^{-6}$).

2.2 Real-time synthesis of movements by learned dynamic primitives

The modeling of the individual actions within the sequence exploits a learning-based approach, which we implemented successfully before for locomotion as well as to other complex human body movements [GMP⁺09]. The system architecture is illustrated in Fig. 3.

Based on the motion capture data, we learned spatio-temporal components of the three actions in an unsupervised way, applying anechoic demixing [OG11, CdEG13]). We have shown before that this method leads to highly compact approximations of human trajectories, reaching almost perfect approximations of often with less than five learned source functions. The skeleton model of the animated characters had 17 joints. The joint angle trajectories were represented by normalized quaternions (exploiting an exponential map representation, c.f. [Mai90], with 3 variables specifying each




Figure 3: Architecture for the online synthesis of body movements using dynamic primitives.

quaternion). The angles were approximated by an anechoic mixture model of the form:

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

The index i specifies the joint-angle component, and the index j the source signals s_j . The parameters w_{ij} and τ_{ij} specify the mixing weights and time delays of the source decomposition model, which are estimated together with the other parameters by the demixing algorithm. The parameters m_i specify the means of the joint trajectories.

In order to generate movements online, the source functions are generated by mapping the solutions of a non-linear dynamical system (canonical dynamics) onto the source functions s_j . For mathematical convenience, we chose a limit cycle oscillator (Hopf oscillator) as canonical dynamics. It can be characterized by the differential equation system (with ω defining the eigenfrequency), for the pair of state variables $[x(t), y(t)]$:

$$\begin{aligned} \dot{x}(t) &= [1 - (x^2(t) + y^2(t))]x(t) - \omega y(t) + k(x_p(t) - x(t)) \\ \dot{y}(t) &= [1 - (x^2(t) + y^2(t))]y(t) + \omega x(t) + k(y_p(t) - y(t)) \end{aligned}$$

The last terms specify coupling terms to a pair of input signals $x_p(t)$ and $y_p(t)$, and k is the coupling strength. For $k = 0$ this equation produces a stable limit cycle. The state space variables x and y are mapped onto the source functions s_j by nonlinear mapping functions $f_j(x, y)$, which were learned by support vector regression (using a radial basis function kernel and the LIBSVM Matlab[®] library [CL01]). The learned source functions $s_j(t)$ and corresponding states $[x(t), y(t)]$ from the attractor solution of the limit cycle oscillator were used as training data.




Figure 4: Comparison of approximation quality for different methods for blind source separation as function of the number of sources, using a step-wise regression approach (residuals after subtraction of the contribution of the non-periodic source signal). *Solid lines*: Approximation quality for trajectories of all three actions as a function of the number of (periodic) source functions for anechoic demixing (*blue*) and principle component analysis (PCA) (*green*). The *purple dotted* line shows the approximation quality for the first action, fixing the delays across trials. The *red dashed* line shows approximation quality when 2 additional sources (with fixed delays) were included in order to model the remaining residuals. Circles mark the chosen numbers of sources in our implementation.

The coupling term (for $k > 0$) allows the coupling of different dynamic primitives, if they are specified by the state variables of another oscillator. We have discussed elsewhere that this form of coupling, with appropriate constraints for the parameters, allows to guarantee the stability of the solutions of networks of such primitives. The relevant stability conditions can be derived using *Contraction theory* [LS98, PMSG09].

In our architecture we used one leading oscillator, and the other oscillators were coupled to this leading oscillator in the described form (star topology of the coupling graph, where couplings are unilateral from the center to the leaves of the star). The stability properties of this form of coupling were studied in detail in [PMSG09], and it can be shown that this dynamics has only a single exponentially stable solution. The state of the leading oscillator was also used for the control of the non-periodic source functions.

From the source signals that were generated online, the joint angles were computed using equation (1). Exploiting the fact that the attractor solution of the Hopf oscillator lies on a circle in state space, the delays can be replaced by an appropriate rotations of the variables of the state space (x, y) . In this way, we obtained a dynamics without explicit time delays, avoiding difficulties with the design of appropriate controllers. Different motion styles were generated by blending of the mixing weights w_{ij} and the trajectory mean values m_i .

2.3 Stepwise regression approach for the modeling of the individual actions

In order to model the step sequences with coordinated walking and reaching we approximated the training

data by the described anechoic mixtures, using a step-wise regression approach that introduced different types of source functions for the three different component actions.

Reaching is a non-periodic movement and therefore requires the introduction of a non-periodic source function. In order to generate such a function online, the phase of the leading Hopf oscillator was derived from the state variables according to the relationship $\phi(t) = \text{mod}_{2\pi}(\arctan(y(t)/x(t)))$, (ensuring $0 \leq \phi < 2\pi$). The non-periodic source signal was defined by $s_0(t) = \cos(\phi(t)/2)$, and the corresponding delay was set to zero.

The three actions of the training sequences were modeled as follows:

1st action: The weights of the non-periodic sources were determined in order to account for the non-periodic part of the training trajectory. Then this component was subtracted from the trajectory data, and the periodic source functions were determined by anechoic demixing, using an algorithm from [CdEG13], which had been modified in order to constrain all time delays belonging to the same source function to be equal. This constraint simplifies the blending between different motion styles, since then the delays of the sources are identical over styles, so that they do not have to be blended. Compared to the unconstrained anechoic model, this constraint requires the introduction of more sources for the same approximation quality (see Fig. 4). The first step could be modeled with sufficient accuracy using three periodic sources in addition to the non-periodic one.

2nd action: In order to model the second highly adaptive step, five periodic sources were required. The first three periodic sources were identical with the ones used for the approximation of the first action, and also the corresponding delays. The weights were optimized in order to minimize the remaining approximation error. The contributions of these three periodic sources (and of the non-periodic sources), then were subtracted from the training data, and two additional periodic sources were learned from the residuals (with constant delays across trials).

3rd action: In order to approximate this action, we used the same non-periodic and five periodic source signals, with the same time delays, that were identified for the modeling of the second action, while the weights of these sources were re-estimated.

The estimated source functions are shown in Fig. 5. The dotted curve illustrates the non-periodic source.




Figure 5: The source signals extracted by anechoic demixing algorithm. **a)**: three periodic source signals extracted from the first action and non-periodic source signal (dashed line). **b)**: two additional periodic source signals that were used for the modeling of the second and the third actions.

The source functions illustrated in the upper panel were used for the approximation of all three actions, and the two in the lower panel only for actions two and three.

Fig. 4 shows the approximation quality as a function of the number of source functions for the first and the second action, comparing normal anechoic demixing [OG11], our algorithm with constant delays over the different conditions [CdEG13], and a reconstruction using PCA. The measure for approximation quality was defined as $Q = 1 - (\|X - \hat{X}\|_F^2) / \|X\|_F^2$, where X is the matrix with the samples of the original signal, and \hat{X} is the reconstructed signal, $\|\cdot\|_F^2$ is the squared Frobenius norm. Especially, the model without constraints for the delays still achieves significantly better approximation quality than PCA. The reconstruction error for the first action (purple circle on Fig. 4) is 95.6%, while the one with the two additional sources, used for actions 2 and 3, is 96.7% for the whole dataset (red circle).

The absolute values of the amplitudes of the weights for a single trajectory are depicted at Fig. 6, separately for the two source signals that carried the maximum amount of variance. This is the non-periodic source and the periodic source with the lowest frequency. The figure shows that the primitives clearly contribute to the different degrees of freedom of the human body. The non-periodic source primarily contributes to the joint angles of the arm, while the periodic source function strongly influences the hip and the leg joints. This clearly reflects the organization of human full body movements in terms of movements primitives. The figure also shows that the contribution of the sources changes between the steps. In the first action the contribution of the first periodic source is dominant, while in the second and last action the non-periodic source function makes a dominant contribution, reflecting the non-periodic reaching movement.

2.4 Online blending of the mixing weights

As illustrated in Fig. 6, the mixing weights change between the different actions within the sequence. For the modeling of a smooth transitions between the different actions the mixing weights thus had to be smoothly

The distribution of the amplitudes of sources weights

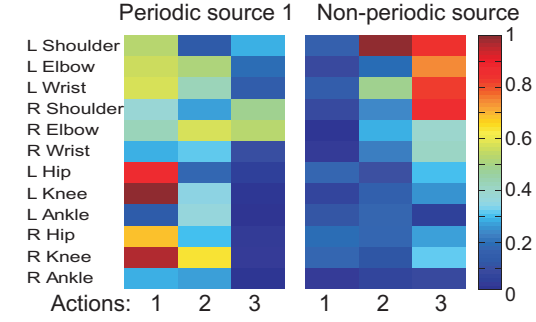


Figure 6: Absolute values of the weights for an example trajectory of the data set. The computed mixing weights are shown from the different actions within the sequence for the periodic source function with minimum frequency and for the non-periodic source. The color code is the same for both panels.

interpolated in an online fashion at the transitions between the individual actions.

For the weights associated with the periodic sources, the corresponding weight matrices were linearly blended according to the relationship $W(t) = (1 - \alpha(t))W_{\text{prev}} + \alpha(t)W_{\text{post}}$, where W_{prev} is the weight matrix in the step prior to the transition and W_{post} the one after the transition. The mean values for each of the angle trajectories were morphed accordingly: $m(t) = (1 - \alpha(t))m_{\text{prev}} + \alpha(t)m_{\text{post}}$, where m_{prev} is the mean value in the step prior to the transition and m_{post} is the one after the transition. The time-dependent blending weight $\alpha(t)$ was constructed from the phase variable $\phi(t)$ of the leading oscillator. Identifying the transition point, where the weights switch between the subsequent actions with the phase $\phi = 0$, the blending weight was given by the equation (here, regarding only two adjunct actions, we use convention: $\phi \in [-2\pi; 0[$ for a previous action, and $\phi \in [0; 2\pi[$ for a next one):

$$\alpha(t) = \begin{cases} 0 & \phi < -\beta, \\ (1 + \sin(\frac{\pi\phi(t)}{2\beta}))/2 & \phi \in [-\beta; \beta], \\ 1 & \phi > \beta \end{cases} \quad (2)$$

The parameter $\beta = \pi/5$ determines the width of the interpolation interval and was chosen to guarantee natural-looking transitions. This value was derived in previous work, optimizing transitions for other scenarios [GMP⁺09].

The weights associated with the non-periodic source had to be treated separately since they can have different signs before and after the transition. Since the timing of this source is completely determined by the phase $\phi(t)$ of the leading oscillator, we constrained the blending by allowing sign changes for these weights only at the point where this phase crosses zero ($\phi(t) = 0$). The ramp-like non-periodic source is normalized in a way so that $s_0(0) = 1$ and $s_0(T) = -1$




Figure 7: Learned nonlinear mappings between action length and duration and the mixing weight of the 1st source for hip flexion angle: **a)** 1st action, **b)** 2nd action.

(T being the duration of an oscillation of the leading oscillator in the attractor state). The following morphing rule $W(t) = \text{sign}(\phi(t))[(\alpha(t) - 1)W_{\text{prev}} + \alpha(t)W_{\text{post}}]$ ensures a smooth transition that make the weights for this source converge at the boundaries between the actions against the value $\xi_{\text{trans}} = (m_{\text{prev}} + m_{\text{post}})/2 + (W_{\text{post}} - W_{\text{prev}})/2$.

2.5 Learning of mappings between step parameters and mixing weights

In order to make the generated behavior highly adaptive for conditions that were not in the training data and for dynamic changes of the environment, we devised an online control algorithm for the blending of the weights W , separately for each action. For this purpose, we learned nonlinear functions that map the step lengths and the duration of the steps onto the mixing weights. For the learning of this highly nonlinear mapping we used locally weighted linear regression (LWLR, [AMS97]). Fig. 7 shows some example for the weights of the first periodic source.

The required step lengths are computed online from the total distance to the drawer. The length of the step of the second action was optimized in order to generate an optimum (maximally comfortable) distance for the third action, which was estimated from the human data to be about $0.6m$. The total distance between the start position and the drawer D was then redistributed between the first two actions using a linear weighting scheme, specifying the relative contributions by the weight parameter γ . The remaining distance $D - 0.6m$ was then distributed according to the relationships $D_1 = (D - 0.6m)\gamma$ and $D_2 = (D - 0.6m)(1 - \gamma)$, where we fitted $\gamma = 0.385$ based on the human data. This approach is motivated by the hypothesis that in humans predictive planning optimizes *end-state comfort*, i.e. the distance of the final reaching action [LRSS13].

We extended the algorithm in addition by a method that introduces additional normal steps (corresponding to action 1), in cases where the goal distance exceeds the distance that can be modeled without artifacts by a three-action sequence. If the distance between the goal




Figure 8: Two synthesized trajectories, illustrated in parallel for two conditions with different initial distance of the character from the drawer. Both animations look highly natural even though these goal distances were not present in the training data set.

and the agent was too short for the introduction of long steps, instead a variable number of short steps as in action 2 were introduced.

3 RESULTS

Two example sequences of concatenated actions generated by our algorithm, for distances to the goal object that were not in the training set are shown in Fig. 8. An example video can be downloaded from [Demo²].

A more systematic evaluation shows that the algorithms can, without introducing additional steps, create natural looking coordinated sequences for goal distances between 2.34 and 2.94 m [Demo³]. If the specified goal distance exceeded this interval our system introduced automatically additional gait steps, making the system adaptive for goal distances beyond 3 meters. This is illustrated in [Demo⁴] that presents two examples of generated sequences for goal distances 3.84 and 4.62 m. With 3 actions the largest achievable range of goal distances without artifacts was about 60 cm, while adding another step increases this range to about 78 cm. Adding two or more normal gait steps our method is able to simulate natural-looking actions even for goal distances longer than 5 m. The next [Demo⁵] illustrates the sequence of three actions of first type followed by actions 2 and 3 for the goal distance 5.3 m.

Fig. 9 illustrates that, like in humans, the posture at the transition between the second and third action depends

² www.uni-tuebingen.de/uni/knv/arl/avi/wscg15/v2.avi

³ www.uni-tuebingen.de/uni/knv/arl/avi/wscg15/v3.avi

⁴ www.uni-tuebingen.de/uni/knv/arl/avi/wscg15/v4.avi

⁵ www.uni-tuebingen.de/uni/knv/arl/avi/wscg15/v5.avi




Figure 9: Postures at the transition between actions 2 and 3 for different lengths of the second action (red: 0.53 m, green: 0.39 m). Even though the distances to the drawer are the same in the last action the postures differ due to the predictive planning of the second action.




Figure 10: Online perturbation experiment. The goal (drawer) jumps away during the approaching of the character. The online planning algorithm introduces automatically an action of type 2 (short step) to adjust for the large distance to the goal.

on the previous step. In one case the step lengths for action 2 were 0.53m and 0.39m, while the distance in the last step was identical (0.6m). This illustrates that in fact the posture for the reaching is modified in a predictive manner over multiple steps, where the predictive planning modifies the posture at the beginning of the last action even if the distance to the goal object for this action is identical. A planning scheme that is not predictive would predict here the same behaviors for the last action since the relevant control variable (distance from the object) is identical for both cases.

An even more extreme demonstration of this online adaptivity is shown in movie [Demo⁶]. Here the drawer jumps away during the approaching behavior by a large distance so that it can no longer be reached with the originally planned number of steps. (Fig. 10). The online planning algorithm adapts to this situation by automatically introducing an additional step so that the behavior is successfully accomplished. Again the behavior has a very natural appearance even though this scenario was not part of the training data set.

⁶ www.uni-tuebingen.de/uni/knv/ar/avi/wscg15/v6.avi

4 CONCLUSIONS

We have presented a method for the online animation of multi-step human movements that was inspired by concepts derived from biological systems. The proposed system realizes a predictive planning of multi-step sequences, including periodic and non-periodic movements that reproduce critical properties observed in experiments on human motor planning. The planning is predictive and optimizes the 'comfort' during the execution of the final action. The proposed system exploits the concept of movement primitives in order to implement a flexible and highly natural-looking coordination of periodic and non-periodic behaviors of the upper and lower limbs, and to realize smooth transitions between subsequent actions within the sequence. For the first time, our architecture is implemented for generation of goal-directed movements. Our approach differs from the whole-body motion blending approach presented in [HK14], where, in order to increase naturalness of the transitions, it was necessary to introduce empirical rules that depend on the gait phase. Future work will extend our approach to other classes of movements, including, for instance, adaptive arm reaching movements accomplished while walking. In addition, we plan a systematic evaluation of the realism of the generated motions, including psychophysical studies.

ACKNOWLEDGEMENTS

The work supported by EC FP7 under grant agreements FP7-248311 (AMARSi), FP7-611909 (Koroibot), H2020 ICT-644727 (CogIMon), FP7-604102 (HBP), PITN-GA-011-290011 (ABC), DFG GI 305/4-1, DFG GZ: KA 1258/15-1, BMBF, FKZ: 01GQ1002A. Authors thank Biwei Huang for help with data segmentation and animation.

REFERENCES

- [AFO03] O. Arikian, D.A. Forsyth, and J. F. O'Brien. Motion synthesis from annotations. *ACM Trans. on Graphics, SIGGRAPH '03*, 22(3):402–408, 2003.
- [AMJ07] Y. Abe, Da Silva M., and Popović J. Multiobjective control with frictional contacts. *ACM SIGGRAPH/Eurograph. Symp. on Comp. Anim.*, 2007.
- [AMS97] C. G. Atkeson, A. W. Moore, and S. Schaal. Locally weighted learning. *A.I. Review*, 11:11–73, 1997.
- [Ber67] N.A. Bernstein. *The coordination and regulation of movements*. Pergamon Press, N.Y., Oxford, 1967.
- [BRI06] J. Buchli, L. Righetti, and A. J. Ijspeert. Engineering entrainment and adaptation in limit cycle systems - from biological inspiration to applications in robotics. *Biol. Cyb.*, 95(6):645–664, 2006.
- [CdEG13] 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.

- [CG13] E. Chiovetto and M. A. Giese. Kinematics of the coordination of pointing during locomotion. *Plos One*, 8(11), 2013.
- [CL01] 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>.
- [CMCH96] 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.
- [FH05] T. Flash and B. Hochner. Motor primitives in vertebrates and invertebrates. *Curr. Opin. Neurobiol.*, 15(6):660–666, 2005.
- [FMJ02] A. Fod, M. J. Mataric, and O. C. Jenkins. Motor primitives in vertebrates and invertebrates. *Auton. Robots*, 12(1):39–54, 2002.
- [FP03] A. Fang and N. S. Pollard. Efficient synthesis of physically valid human motion. *ACM Trans. on Graphics*, 22(3):417–426, 2003.
- [FXS12] A. W. Feng, Y. Xu, and A. Shapiro. An example-based motion synthesis technique for locomotion and object manipulation. *Proc. of ACM SIGGRAPH 13D*, pages 95–102, 2012.
- [GMP⁺09] 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.
- [GRIL08] A. Gams, L. Righetti, A. J. Ijspeert, and J. Lenarcic. A dynamical system for online learning of periodic movements of unknown waveform and frequency. *Proc. of the IEEE RAS / EMBS Int. Conf. on Biomed. Robotics and Biomechatronics*, pages 85–90, 2008.
- [GSKJ03] 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.
- [HK14] Y. Huang and M. Kallmann. Planning motions for virtual demonstrators. In *Intelligent Virtual Agents*, pages 190–203. Springer, 2014.
- [HPP05] E. Hsu, K. Pulli, and J. Popovic. Style translation for human motion. *ACM Trans. on Graphics*, 24:1082–1089, 2005.
- [KGP02] L. Kovar, M. Gleicher, and F. Pighin. Motion graphs. *Proc. of SIGGRAPH 2002*, pages 473–482, 2002.
- [LHP05] K. Liu, A. Hertzmann, and Z. Popović. Learning physics-based motion style with nonlinear inverse optimization. *ACM Trans. on Graphics*, 23(3):1071–1081, 2005.
- [LRSS13] 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.
- [LS98] W. Lohmiller and J. J. E. Slotine. On contraction analysis for nonlinear systems. *Automatica*, 34(6):683–696, 1998.
- [LWS02] 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.
- [Mai90] P.-G. Maillot. Using quaternions for coding 3D transformations. In A. S. Glassner, editor, *Graphic Gems*, pages 498–515. Academic Press, Boston, MA, 1990.
- [MB01] 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.
- [MLPP09] U. Muico, Y. Lee, J. Popović, and Z. Popović. Contact-aware nonlinear control of dynamic characters. *ACM Trans. on Graphics*, 28(3):Art.No.81., 2009.
- [OG11] L. Omlor and M. A. Giese. Anechoic blind source separation using wigner marginals. *J. of Machine Learning Res.*, 12:1111–1148, 2011.
- [PMSG09] 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.
- [PSS02] 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.
- [RCB98] C. Rose, M. Cohen, and B. Bodenheimer. Verbs and adverbs: Multidimensional motion interpolation. *IEEE Comp. Graphics and Appl.*, 18(5):32–40, 1998.
- [RGBC96] C. Rose, B. Guenter, B. Bodenheimer, and M. Cohen. Efficient generation of motion transitions using space-time constraints. *Int. Conf. on Comp. Graph. and Interactive Techniques, Proc. ACM SIGGRAPH '96*, 30:147–154, 1996.
- [Ros08] D. A. Rosenbaum. Reaching while walking: reaching distance costs more than walking distance. *Psych. Bull. Rev.*, 15:1100–1104, 2008.
- [SHP04] 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.
- [SMKB14] 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.
- [ST05] W. Shao and D. Terzopoulos. Artificial intelligence for animation: Autonomous pedestrians. *Proc. ACM SIGGRAPH '05*, 69(5-6):19–28, 2005.
- [WFH08] 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.
- [WP95] A. Witkin and Z. Popović. Motion warping. *Proc. ACM SIGGRAPH '95*, 29:105–108, 1995.
- [WS10] M. Weigelt and T. Schack. The development of end-state comfort planning in preschool children. *Exper. Psych.*, 57(6):476–782, 2010.
- [WZ10] C.-C. Wu and V. Zordan. Goal-directed stepping with momentum control. *ACM SIGGRAPH/Eurographics Symp. on Comp. Animation (SCA) 2010*, 2010.