

Movement Control Methods for Complex, Dynamically Simulated Agents: Adonis Dances the Macarena

Maja J. Matarić
Computer Science Department
University of Southern California
Los Angeles, CA 90089-0781
tel: (213) 740-4520
matarić@cs.usc.edu

Victor B. Zordan
College of Computing
Georgia Institute of Technology
Atlanta, GA 30332-0280
tel: (404) 894-4998
victor@cc.gatech.edu

Zachary Mason
Computer Science Department
Brandeis University
Waltham, MA 02254-9110
tel: (781) 736-2719
zmason@cs.brandeis.edu

Abstract

We describe and compare two implemented controllers for Adonis, a physically simulated humanoid torso, one based on joint-space torques and the other on convergent force-fields applied to the hands. The two come from different application domains: the former is a common approach in manipulator robotics and graphics, while the latter is inspired by biological limb control. Both avoid explicit inverse kinematic calculations found in standard Cartesian control, trading generality of motion for programming efficiency. The two approaches are compared on a common sequential task, the familiar dance “Macarena” and evaluated based on ease of generating new behaviors, flexibility, and naturalness of movement; we also compare them against human performance on the same task. Finally, we discuss the tradeoffs and present a more general framework for addressing complex motor control of simulated agents.

1 Introduction

Control of humanoid agents, dynamically simulated or physical, is an extremely difficult problem due to the high dimensionality of the control space, i.e., the many degrees of freedom and the redundancy of the system. In robotics, standard methods have been developed for simpler manipulators and have been gradually scaled up to more complex arms (Paul 1981, Brady, Hollerbach, Johnson, Lozano-Perez & Mason 1982) and recently to physical human-like arms (Schaal 1997, Williamson 1996). The problem of anthropomorphic control has also found a new application

area in realistic, physically-based animation, where the control of dynamic simulations of human characters, involving realistic physical models, matches the complexity of the robotics problem (Pai 1990, Hodgins, Wooten, Brogan & O’Brien 1995, Van de Panne & Lamouret 1995).

We introduce a novel approach to manipulator control, based on models from neuroscience which employ convergent force-fields at the end-points of a manipulator. We explore the feasibility of such a model for complex animation and agent control in general, and discuss how it can be generalized to high-level motor tasks and incorporated into a general control framework. Furthermore, we compare the methodology to a standard robotics approach, which has also been adopted in animation, employing joint-space control with torque actuators. Both approaches are appealing because they avoid explicit computation of inverse kinematics (IK) found in standard Cartesian control. The inputs of each controller are used explicitly, as either positions or orientations, without IK solvers converting the input data. However, this presents a tradeoff between generality of motion and programming efficiency. To compare the two approaches, we implemented them on a common motor task: a continuous sequence of smooth movements. For the purpose of evaluation, a well known, goal-driven sequence was chosen, the popular dance “Macarena.” The dance presents a non-trivial, well-defined task that can be precisely specified and evaluated, both relative to the quantitative specification and to qualitative human performance.

2 Background

2.1 Control in Robotics and Computer Animation

Computer animation and robotics are two primary areas of research into motion for artificial agents. 3D animated character motion has traditionally been cre-

ated by hand, through a time-consuming process. Recently, physical modeling has been used to generate motion by minimizing user-specified constraints while allowing the model constraints to add physical realism. Witkin & Kass (1988) pursue physical modeling through such a constraint-based approach; by choosing start and end conditions, they generate anticipation and determination in the action. Cohen (1992) extended this approach with higher DOF systems and more complex constraints. Ngo & Marks (1993) introduce a constraint approach of creating behaviors automatically using genetic algorithms.

Dynamic simulation has been used to generate graphical motion by applying dynamics to physically-based models and using forward integration. Simulation ensures physically plausible motion by enforcing the laws of physics. Pai (1990) simulates walking gaits, drawing strongly from robotics work. His torso and legs use a controller based on high-level time-varying constraints. Hand-tuned control of simulations has been applied successfully to more complex systems such as full articulated human figures. Raibert & Hodgins (1991) demonstrate rigid body dynamic simulations of legged creatures. Their hand-tuned controllers consist of state machines that cycle through rule-based constraints to perform different gaits. Hodgins et al. (1995) extend this work to human characters, suggesting a toolbox of techniques for controlling articulated human-like systems to generate athletic behaviors such as 3D running, diving, and bicycling. Van de Panne & Lamouret (1995) use search techniques to find balancing controllers for human-like character locomotion, aiming at more automatic control of simulated agents.

In robotics, manipulator control has been largely addressed for point-to-point reaching, typically by specifying Cartesian 3D goals and explicitly solving the IK for the manipulator's joint angles (Paul 1981, Brady et al. 1982). Various neural network approaches to learning IK for simple manipulators have been explored and more sophisticated learning methods for dynamic tasks and higher DOF systems are being developed (Atkeson 1989, Schaal & Atkeson 1994). The work most similar to the force-field approach we describe was performed by Williamson (1996), who presented a controller for a 6-DOF robot arm, based on the same biological evidence we describe next. It consists of four behaviors: three reaching and one resting posture; intermediate targets are achieved by linear interpolation.

2.2 Biological Inspiration

The flexibility and efficiency of biological motion provides a desirable model for complex agent control.

Our work is inspired by a specific principle derived from evidence in neuroscience. Mussa-Ivaldi & Giszter (1992), Giszter, Mussa-Ivaldi & Bizzi (1993) and related work on spinalized frogs and rats suggest the existence of force-field motor primitives that converge to single equilibrium points and produce high-level behaviors such as reaching and wiping. When a particular field is activated, the frog's leg executes a behavior and comes to rest at a position that corresponds to the equilibrium point; when two or more fields are activated, either a linear superposition of the fields (87% of tested cases), or a "winner-take-all" response (58% of remaining cases) results (Mussa-Ivaldi, Giszter & Bizzi 1994). This suggests an elegant organizational principle for motor control, in which entire behaviors are coded with low-level force-fields, and may be combined into higher-level, more complex behaviors. The idea of supplying an agent with a collection of *basis behaviors* or *primitives* representing force-fields, and combining those into a general repertoire for complex motion, is very appealing. Our previous work (Mataric 1995, Mataric 1997), inspired by the same biological results, has already successfully applied the idea of *basis behaviors* to control of planar mobile agents/robots. This paper extends the notion to agents with more DOFs.

Another inspiration comes from psychophysical data describing what people fixate on when observing human movement. Mataric & Pomplun (1997) demonstrate that when presented with videos of human finger, hand, and arm movements, observers focus on the hand, yet when asked to imitate the movements, subjects are able to reconstruct complete trajectories (even for unnatural movements involving multiple DOFs) in spite of having attended to the end-point. This could suggest some form of internal models of complete behaviors or primitives for movement, which effectively encapsulate the details of low-level control. Given an appropriately designed motor controller, tasks could be specified largely by end-point positions and a few additional constraints, and the controller could generate the appropriate corresponding postures and trajectories. The force-field approach we describe is a small step toward such an approach. To compare it with an alternative, we implemented both on a common testbed, described next.

3 The Dynamic Anthropomorphic Simulation: Adonis

Adonis is a rigid-body simulation of a human torso, with static graphical legs (Figure 1), consisting of eight rigid links connected with revolute joints of one and three DOFs, totaling 20 DOFs. The dynamic model for Adonis was created by methods described



Figure 1: The Adonis dynamic simulation testbed.

in Hodgins et al. (1995). Mass and moment-of-inertia information is generated from the graphical body parts and equations of motion are calculated using a commercial solver, *SD/Fast* (*SD/Fast User's Manual* 1990). The simulation acts under gravity, accepts other external forces from the environment, and applies low level PD-servo control (see Section 6) to keep balance at the waist and neck. No collision detection, with itself or its environment, is used in the described implementation; we are currently implementing this extension.

Adonis is particularly well suited for testing and comparing different motor control strategies; the simulation allows us to apply forces to the end-points while a physical robot implementation would require explicit calculation of the IK of the arms to solve for actuator torques from the desired forces. This, in turn, enables us to implement experimental controllers for human-like movement more easily, while having the simulation software handle the issue of IK and dynamics.

4 Task Specification

Natural, goal-driven movement relies on precise specification, coordination, and constraints imposed by implementation and evaluation. A test task should be challenging to control but familiar enough to evaluate. The Macarena is a popular dance which involves a sequence of coordinated movements that constitute natural subtasks. We used a verbal description, aimed at teaching people the dance (found on the web at <http://www.radiopro.com/macarena.html>) to implement controllers according to a common set of

meaningful instructions, as opposed to fully describing the tasks at the simulation level. The Macarena specification, omitting the hip and whole-body subtasks at the end, is given below:

1. Extend Right Arm straight out, palm down
2. Extend Left Arm straight out, palm down
3. Rotate Right Hand (palm up)
4. Rotate Left Hand (palm up)
5. Touch Right Hand to top of your left shoulder
6. Touch Left Hand to top of your right shoulder
7. Touch Right Hand to the back of your head
8. Touch Left Hand to the back of your head
9. Touch Right Hand to the left side of your ribs
10. Touch Left Hand to the right side of your ribs
11. Move Right Hand to your right hip
12. Move Left Hand to your left hip

No task-planning was necessary because the sequence is provided by the dance specification. However, the individual subtasks are not specified in the same frame of reference. The first four deal with extending the arms, best expressed in joint angles, while the rest are better described in an ego-centric Cartesian reference frame. This type of heterogeneous task specification is common in natural language descriptions, and control systems must satisfy each of the different goals regardless of the underlying representation. To address the issue of controller representation, we implemented two very different alternatives, and compared them based on common performance metrics to demonstrate the tradeoffs involved.

5 The Force-Field Approach

We introduce a biologically-inspired control approach based on applying convergent force-fields to end effectors, Adonis' hands. Inspired by the neuroscience work on motor control in frogs described in Section 2, this approach affords a more intuitive user interface at the expense of motion generality. A number of force fields are used as a set of basis functions for controlling Adonis; desired Cartesian goals are reached by applying forces to the hands to move them to the desired destinations. Each field has a single equilibrium point (EP) at a target position where the field is zero. From any point in Adonis' reachable workspace, the hand will stably move toward the EP, excluding singularities in the kinematics. EPs are chosen empirically as target positions for the hands in the 12 subtasks of the Macarena. Some intermediate postures were used, as described in Section 7.

The force fields used as basis functions depend on the position and velocity of the hand being controlled. The exerted force is proportional to the difference between the current and desired velocity, calculated

as a function of the difference between the current position and the equilibrium point. Formally:

$$F = c (v_{actual} - v_{desired} (|x - x_{EP}|))$$

where $v_{desired}$ is the desired velocity, v_{actual} is the actual velocity and c is the gain constant. The magnitude of c determines the speed of the complete movement. For well-chosen values of c , this model makes the hand position converge to the EP, following a roughly bell-shaped velocity profile, consistent with human data on regular-speed reaching motions (Flash & Hogan 1985). Linear combinations of the force-fields can produce new fields with convergence points in the convex hull of the EPs of the bases. Bases can be combined using vector summation, as in our approach, or using a winner-take-all framework, as used by Williamson (1996) for robot reaching. Our trajectories are generated by moving the equilibrium point through a set of subgoals or control points. The EP is computed from a linear weighted composition of the bases functions. We use feedback to move between control points, transitioning when the current control point is reached. This works well since the Macarena is a relatively slow task; highly dynamic behaviors (e.g., ball throwing) are better and more robustly handled with feedforward control, but non-trivial arm dynamics make it difficult to predict proper transition timing.

The force-field approach presents a convenient control interface because it effectively lowers the dimensionality of the system. This makes the control task easier, although the current implementation limits flexibility. Several of the joints, including the waist, neck and wrists, use PD-servo controllers (described in Section 6) to maintain fairly rigid behavior. Because these joints are not explicitly controlled by the user, the range of controllable motion is limited, and the naturalness of the resulting behavior can be reduced, as discussed in Section 7.

6 The Joint Torque Approach

Joint-space controllers with torque actuators have been used successfully to generate behaviors for a variety of systems (Pai 1990, Raibert & Hodgins 1991, Hodgins et al. 1995, Van de Panne & Lamouret 1995). In general, these controllers choose a set of torques for all actuated joints. We designed a hand-tuned PD-servo feedback controller for performing the Macarena on Adonis. In this approach, torques are calculated for each joint as a function of angular position and velocity errors between the feedback state and desired state:

$$T = kd (\dot{\theta}_{desired} - \dot{\theta}_{actual}) + k (\theta_{desired} - \theta_{actual})$$

where $\dot{\theta}_{actual}$ and $\dot{\theta}_{desired}$ correspond to the actual and desired joint velocities, and θ_{actual} and $\theta_{desired}$ correspond to the actual and desired joint angles.

To generate the Macarena controller, the desired angles used for the feedback error are interpolated from hand-picked key postures. The postures were derived directly from the task specification, each corresponding to one of the 12 subtasks enumerated in Section 4. Intermediate postures between subtasks help guide the joint trajectories through difficult transitions. For example, an intermediate posture was needed for swinging the hands around the head to prevent a direct yet unacceptable path through the head. The incremental desired angles use a spline to smoothly interpolate between the postures. Gains for the PD-servo were chosen by hand and remained constant throughout the behavior. The joint torque approach allows direct control of each actuated joint in the system, giving the user local control of the details of each behavior. However, the controller in turn requires a complete set of desired angles at all times. Specifying all of this information can be tedious, especially for joints that are less important for the behavior being generated. Interpolating between key postures is a reasonable method for reducing the required amount of information.

The control of actuated joints may be individually modified using their respective desired angles, giving localized control over the generated motion. All desired key postures are specified as a set of angles in joint space. In behaviors such as the Macarena, position constraints like “hands behind the head”, can be satisfied with user-level feedback. However, precise Cartesian-space constraints, like “finger on the tip of the nose”, would be difficult to control by hand-tuning using joint-space errors directly. An inverse kinematics solver could be used to generate desired angles from position constraints, although the controller has no direct measure of errors in Cartesian space, currently.

7 Performance Analysis and Comparisons

Analysis and evaluation of complex behavior is an open research challenge. As synthetic behavior becomes more complex, the issue becomes increasingly acute in animation, robotics, and AI in general. We explored several evaluation criteria before deciding on those elaborated below.

7.1 Controller Flexibility

Behavior constraints come in various forms. Complex movements can be broken down into subtasks with Cartesian-space or joint-space constraints. For

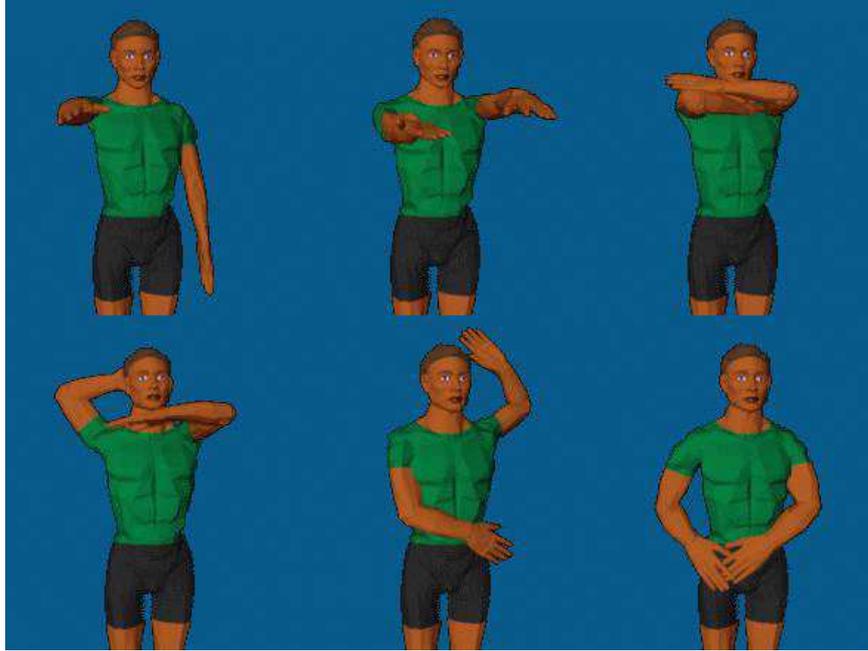


Figure 2: Adonis performing the Macarena, using the joint-space controller.

example, placing hands behind the head is best represented by a Cartesian constraint while turning the hand is better treated as a joint constraint. The issue of coordinate and representation transformation has been addressed extensively in manipulator robotics (Paul 1981, Brady et al. 1982). The two controllers implemented favor very different representations and thus the evaluation of their ability to satisfy all of the various subgoals serves as an effective performance metric.

The force-field controller employs a 3D equilibrium point as a control handle. This representation is well suited to subtasks that describe the location of the end-points/hands, such as pointing, reaching, and basic interaction with a 3D environment. However, a Cartesian equilibrium point is an indirect and, at times, unusable control handle in that complex postures, like folding the arms or turning the palms, cannot be naturally achieved. A key advantage of this approach is that it limits the amount of information required for describing a behavior, reducing the dimensionality of the problem. However, it also removes access to the individual degrees of freedom for the agent. While specializations of the force-field approach are possible, they compromise its simplicity.

The joint-space controller uses desired angles as control handles, allowing more complete control of the agent’s motion. Behaviors that include joint con-

straints such as gesturing and free-form movement are more easily controlled with this approach, and subtleties, such as head motion and elbow positioning, can be controlled (Figure 2). The richness of the control, however, requires the user to specify joint angle information for each DOF in each subtask. In this approach, the representational transformation from Cartesian to joint space is done implicitly by the programmer. Consequently, tight Cartesian constraints are difficult to achieve with the current implementation and may require an IK solver.

A control structure should easily accommodate a variety of tasks and constituent behaviors, and generating new behaviors should require a minimal amount of user input. The joint-torque approach to control is not easily generalizable between behaviors, as little or no knowledge can be transferred; timing and postures are usually specific and need to be regenerated with each new controller. In contrast, the force-field approach may be used to generate new behaviors more easily due to its reduced control dimensionality. Also, using its linear additive properties, new behaviors may be generated by combining existing ones.

7.2 Naturalness of Movement: Qualitative

Judging the naturalness of movement is important, but aesthetic judgment is difficult to quantify. Dynamic simulation constrains motion to be physically

<i>Human Data</i>			<i>Torque-driven</i>		<i>Force-driven</i>	
T	jerk	time	jerk	time	jerk	time
1	85800	1.00	38500	0.70	1100	2.87
2	89000	1.33	63400	0.57	560	1.73
3	1970	1.33	1010	0.73	*	*
4	5060	0.87	1930	0.53	*	*
5	18800	0.30	55200	0.57	360	3.40
6	21200	1.00	84100	0.47	360	3.27
7	141000	1.00	105000	0.90	22600	1.60
8	111000	1.57	64000	0.70	7300	1.70
9	68200	0.90	205000	0.60	14700	1.67
10	65900	1.13	401000	0.67	3200	1.67
11	71900	0.80	8990	0.30	20	1.67
12	12100	1.00	12800	1.00	40	1.73

Figure 3: A comparison of minimum jerk values.

plausible but it is not necessarily natural. Viewer-based evaluation is one qualitative approach to measuring naturalness. Real-time playbacks of both controllers we implemented are available from <http://www-robotics.usc.edu/~agents/macarena.html>

Rigid body simulation imposes limitations that cannot be overcome by control. For instance, its un-actuated back will necessarily appear stiff. Furthermore, the controllers have no knowledge of the body position and avoiding self-collisions was done by hand. This results in conservative, unnatural trajectories which would be improved with a controller capable of collision prediction and avoidance. Similarly, joint limits would contribute to more natural behavior. To improve appearance, the force-field method uses strong PD-servos to clamp some joints, avoiding marionette-like unconstrained movement. The joint-torque method interpolated postures with splines to smooth the resulting motion, and includes small head and hand movements that result in richer motion. Head gaze can add much in terms of naturalness of motion and its implementation is straightforward.

7.3 Naturalness of Movement: Quantitative

Minimal jerk of hand position has been proposed by Flash & Hogan (1985) as a metric for describing human arm movements in the plane, by minimization of the following index:

$$C_j = \int_0^{t_{final}} (jerk_x^2 + jerk_y^2) dt$$

where $jerk_x$ and $jerk_y$ correspond to the third derivative of the x and y positions with respect to time. We used jerk as the evaluation metric for comparing the simulated motion for the two implemented

controllers against the motion of a human performing the Macarena. We chose this metric over other alternatives, such as minimum torque change, suggested by Uno, Kawato & Suzuki (1989), for its simplicity of computation in our domain; the hand position was easily measured for both the simulation and the human movement. The comparison is performed using the metric of total square jerk in 3D Cartesian space; we added a z position term to the calculation.

Motion for 3D hand positions of a human subject performing the Macarena was recorded using a commercial Flock of Birds electro-magnetic tracking system. Each subtask was isolated and calculated separately for a finer-grain evaluation. Corresponding data were collected for the two simulations; the results are presented in Figure 3. The values in the table correspond to the square jerk over the length of the subtask, for the major moving hand in that subtask (e.g., in subtask one the right arm, in subtask two the left arm, and so on). The units for squared jerk are m^2/sec^6 , and the duration of each subtask, labeled “time” is given in seconds. Subtasks 3 and 4 could not be implemented with the force-field approach, since it lacked the ability to control hand orientation; these entries are indicated with *. Jerk is a sensitive measure, susceptible to position error unavoidable in motion capture data collected with electro-magnetic trackers. Furthermore, timing has a significant effect on the jerk; slower movements imply less jerk, as is shown in the case of the force-control approach. Therefore, the exact values above are less significant than the general trends they indicate.

In this evaluation, there is no correspondence between the goal positions and timing for the human and the simulated agent motions. Therefore, we do not expect correspondence between the controllers and the human jerk values, but instead in trends across each subtask and between the three sets of data. Rather than make a one-to-one comparison, we consider where the data are significantly consistent or inconsistent and suggest some reasons for these patterns. The hand jerk calculation is notably noisy and varies widely from task to task, so a qualitative analysis is more beneficial than an actual numeric comparison. It should be noted that the force-field controller explicitly applies forces to the hand and thus directly affects the jerk. Therefore, the magnitude of the jerk is a contrived metric for the force-controlled movement as a whole. However, the relative jerk within each of the three sets of controllers provides valuable insight.

An important consideration in this comparison is the amount of environment knowledge available to the controller (human or simulated). Human access

to knowledge about body location and effective means of avoiding self-collision allow for generating smooth movements of arbitrary complexity. In contrast, simulations have minimal environment knowledge or trajectory planning capabilities. Consequently, subtasks involving complex self-collision avoidance result in less efficient and less natural behavior. For example, subtasks 5, 6, and 9, 10 exhibit these characteristic trajectory planning difficulties.

Interestingly, while tasks 7 and 8 are similar in nature, and in fact appear most complex in terms of self-collision avoidance, their corresponding jerk values do not show the same pattern. One interpretation of this result stems from the subgoals involved in generating the complex movement. While we have little understanding of the strategies people use in reaching this subgoal, we can infer from the behavior that coordinated movement of the head and the arm, and subsequent contact with the head, is involved. In contrast, the simulated behavior proceeds through a set of intermediate positions which avoid self-collision and result in the arm behind, but not in contact with, the head. Not surprisingly, the resulting jerk pattern is quite different. To study this discrepancy further, a less familiar and intuitive set of subtasks could be used that may provide a stronger comparison between human and simulated motion.

To facilitate an honest comparison, simulated and human motion were generated independently. However, various techniques can be implemented to generate a closer fit between the data, if that is desired. Specifically, human hand positions could be used to select goal positions for the force-field equilibrium points. Similarly, an IK solver could find joint postures for the joint-space controller that achieved these hand positions. In addition, timing considerations influence the comparison. As seen from the data, the timing of subtasks was not correlated. Timing taken from human motion could be used to generate simulated motion that more closely fit the human performance. Lastly, minimization techniques could be applied to the controller parameters to find movements that minimize jerk or other performance metrics.

7.4 Other Evaluation Criteria

Other criteria were also considered. We explored evaluating the correctness of the resulting behavior, but without a set of analytical definitions of each subtask, this proved of little use. Such specific constraints could be computed but were rather inconsistent with the graphical environment where the definition of achieving the task did not rest on precise placement as much as on the overall impression of the complete movement. We also considered efficiency of

implementation as a metric. Both of the described methods avoid explicit IK computations in favor of simple but less general alternatives. Their relative efficiency is difficult to compare, however, due to the integral role of the designer in both approaches.

8 Continuing Work and Conclusion

We have compared two approaches to anthropomorphic agent control, both implemented on a dynamic torso simulation, Adonis, and tested on a 12-subgoal Macarena task. We introduced a biologically-motivated force-field approach, which sets up convergent force fields that operate as a function of the velocity and position of the hands. We then compared it to a joint-space controller which manipulates the individual torques for each DOF of the system allowing complete control over the agent but requiring more intuition on the behalf of the user. Both approaches avoid explicit IK computation by feeding un-converted user inputs to the simulation, making task-specification efficient but not fully general. We compare the controllers against each other and to human data on the same task using a minimum jerk computation as a common metric. The fundamental tradeoff between believability and control effort still remains, as the two approaches produce different results depending on subtask specification.

The described work is a part of a project aimed at developing a biologically-inspired behavior-based approach to motor control. We presented implementations that employ a single representational methodology, which are a part of an effort toward a system that encapsulates the low-level control details within primitive parametric behaviors that satisfy various motor tasks, including movement to point using a specific achievement-goal (such as reaching, and most of the Macarena subtasks) and repetitive and oscillatory movements using maintenance-goals (such as bouncing, waving, swinging, etc.). Individual behaviors may rely on different representations, but their use and performance can be seamlessly integrated by sequencing and co-activation, as in the biological model. In addition to its biological motivation and potential dimensionality reducing properties, the behavior-based model for complex motor control provides a simple and scalable interface for graphical agents. One of our goals is to develop a control architecture that allows for combining various movement primitives (possibly from different users) into a versatile and general system. As complex articulated agents become more prevalent, such a modular approach to control would use its “open architecture” to combine the advantages of various approaches by

encapsulating them into primitives.

Acknowledgments

This work is supported by the NSF Career Grant IRI-9624237 to M. Matarić. The authors thank Nancy Pollard for help with the jerk calculations, Len Norton for help with human motion data, and Stefan Schaal and Jessica Hodgins for sharing expertise and providing many insightful comments. The Adonis simulation was developed by Jessica Hodgins at Georgia Institute of Technology.

References

- Atkeson, C. G. (1989), 'Learning Arm Kinematics and Dynamics', *Annual Review of Neuroscience* **12**, 157–183.
- Brady, M., Hollerbach, J. M., Johnson, T. L., Lozano-Perez, T. & Mason, M. T., eds (1982), *Robot Motion: Planning and Control*, MIT Press, Cambridge, MA.
- Cohen, M. F. (1992), Interactive Spacetime Control for Animation, in 'Computer Graphics (Proceedings, SIGGRAPH '92)', 293–301.
- Flash, T. & Hogan, N. (1985), 'The coordination of the arm movements: an experimentally confirmed mathematical model', *Journal of Neuroscience* **7**, 1688–1703.
- Giszter, S. F., Mussa-Ivaldi, F. A. & Bizzi, E. (1993), 'Convergent force fields organized in the frog's spinal cord', *Journal of Neuroscience* **13**(2), 467–491.
- Hodgins, J. K., Wooten, W. L., Brogan, D. C. & O'Brien, J. F. (1995), Animating Human Athletics, in 'Computer Graphics (Proceedings, SIGGRAPH '95)', Annual Conference Series, ACM SIGGRAPH, Addison Wesley, 71–78.
- Matarić, M. & Pomplun, M. (1997), What do People Look at When Watching Human Movement?, Technical Report CS-97-194, Brandeis University.
- Matarić, M. J. (1995), 'Designing and Understanding Adaptive Group Behavior', *Adaptive Behavior* **4**(1), 50–81.
- Matarić, M. J. (1997), 'Behavior-Based Control: Examples from Navigation, Learning, and Group Behavior', *Journal of Experimental and Theoretical Artificial Intelligence* **9**(2-3), 323–336.
- Mussa-Ivaldi, F. A. & Giszter, S. F. (1992), 'Vector field approximations: a computational paradigm for motor control and learning', *Biological Cybernetics* **67**, 491–500.
- Mussa-Ivaldi, F. A., Giszter, S. F. & Bizzi, E. (1994), 'Linear combinations of primitives in vertebrate motor control', *Proceedings of the National Academy of Sciences* **91**, 7534–7538.
- Ngo, J. T. & Marks, J. (1993), Spacetime Constraints Revisited, in J. T. Kajiya, ed., 'Computer Graphics (Proceedings, SIGGRAPH '93)', 343–350.
- Pai, D. (1990), Programming Anthropoid Walking: Control and Simulation, Technical Report Computer Science Tech Report TR 90-1178, Cornell University.
- Paul, R. P. (1981), *Robot Manipulators: Mathematics, Programming, and Control*, MIT Press, Cambridge, MA.
- Raibert, M. & Hodgins, J. (1991), Animation of Dynamic Legged Locomotion, in 'Computer Graphics (Proceedings, SIGGRAPH '91)', 349–356.
- Schaal, S. (1997), Learning from demonstration, in M. Mozer, M. Jordan & T. Petsche, eds, 'Advances in Neural Information Processing Systems 9', The MIT Press, 1040–1046.
- Schaal, S. & Atkeson, C. C. (1994), 'Robot Juggling: An Implementation of Memory-Based Learning', *Control Systems Magazine* **14**, 57–71.
- SD/Fast User's Manual* (1990), Technical report, Symbolic Dynamics, Inc.
- Uno, Y., Kawato, M. & Suzuki, R. (1989), 'Formation and Control of Optimal Trajectory in Human Arm Movement-Minimum Torque-Change Model', *Biological Cybernetics* **61**, 89–101.
- Van de Panne, M. V. & Lamouret, A. (1995), Guided Optimization for Balanced Locomotion, in 'Proceedings, Eurographics Workshop on Computer Animation and Simulation', 165–177.
- Williamson, M. (1996), Postural Primitives: Interactive Behavior for a Humanoid Robot Arm, in 'Fourth International Conference on Simulation of Adaptive Behavior', P. Maes, M. Matarić, J.-A. Meyer, J. Pollack and S. Wilson, eds., The MIT Press, 124–131.
- Witkin, A. & Kass, M. (1988), Spacetime Constraints, in 'Computer Graphics (Proceedings, SIGGRAPH '88)', 159–168.