

Continuation Methods for Adapting Simulated Skills

KangKang Yin Stelian Coros Philippe Beaudoin Michiel van de Panne*

University of British Columbia




Figure 1: A regular walking gait is automatically adapted to tasks involving a variety of interaction with the environment, such as pushing furniture (left), taking a large step up (middle), and walking on ice (right).

Abstract

Modeling the large space of possible human motions requires scalable techniques. Generalizing from example motions or example controllers is one way to provide the required scalability. We present techniques for generalizing a controller for physics-based walking to significantly different tasks, such as climbing a large step up, or pushing a heavy object. Continuation methods solve such problems using a progressive sequence of problems that trace a path from an existing solved problem to the final desired-but-unsolved problem. Each step in the continuation sequence makes progress towards the target problem while further adapting the solution. We describe and evaluate a number of choices in applying continuation methods to adapting walking gaits for tasks involving interaction with the environment. The methods have been successfully applied to automatically adapt a regular cyclic walk to climbing a 65cm step, stepping over a 55cm sill, pushing heavy furniture, walking up steep inclines, and walking on ice. The continuation path further provides parameterized solutions to these problems.

1 Introduction

Animated characters should exhibit rich and purposeful behavior if they are to mimic human abilities. Data-driven approaches which resequence-and-interpolate captured motions have been remarkably successful at generating high-quality animated motion. However, data-driven approaches are data hungry, particularly if the animation synthesis techniques perform only a shallow analysis of the motion capture data. This motivates the need for methods that are good at generalizing or extrapolating from example motions or example controllers. Because achieving better generalization is based on a deeper understanding of the motions, such methods often model the underlying physics and are typically based on optimization.

The choices for ‘what’ and ‘how’ to optimize need to support the bewildering mix of soft and hard constraints that help shape motions, including physics, joint limits, torque limits, friction limits, and stylistic constraints. The use of a forward-dynamics simulation is an easy way to ensure that many of these constraints are satisfied. The optimization is challenging because of the high-dimensional nature of the required control and the common occurrence of discrete discontinuity-causing events such as an unexpected contact, which can play havoc with gradient-based methods.

The idea of *continuation methods* is to solve a difficult problem by starting from the known solution of a related, but easier problem [Seeger 2006]. For our problem domain, this also captures the intuition that motor skills can often be improved in a progressive fashion. A concrete example is that of learning how to climb up a 65cm step, as shown in Figure 1 (middle). A continuation-based optimization process begins by adapting the control parameters of a particular stride of a regular walking gait in order to successfully climb up a small step. In the next stage, the control parameters are further adapted for climbing a slightly higher step up, and so forth. With this simple description in mind, there remain many questions to be answered. How quickly should the continuation be advanced, i.e., what should the step height increment be? Are gradient-descent techniques sufficient to make the required adaptations? For any given step height, what objective function should drive the adaptation? This paper addresses these questions in the context of generalizing walking steps to situations that involve significant variations in the environment such as those shown in Figure 1.

An abstract view of the continuation problem is shown in Figure 2(a). Given an initial set of parameters θ_0 that produces a regular walking gait, we wish to find the parameters θ_T which achieve walking in a physical environment that has been significantly modified in some way, such as the addition of a large step, an obstacle to step over, icy terrain, or a piece of furniture to push. We define the *continuation variable* γ to represent the parameterization of this change, such as a step height. $\gamma = 0$ represents no change, i.e., level terrain with normal friction and no obstacles, and γ_T represents a target value of the modification, e.g., $\gamma_T = 65cm$ for a large step-up. The grey regions in Figure 2(a) denote settings of (γ, θ) which result in failure, examples of which include slipping or tripping. Also illustrated is that with advancing continuation, the solution often becomes more constrained and requires a significant deviation from the nominal parameters θ_0 . Because any particular value of γ may have many viable solutions, a *style objective function* helps define a unique optimal solution, $\theta^*(\gamma)$. The style objective may include criteria such as minimal deviation from the original control parameters and other user-specified features such as a desired step length or walking speed. Lastly, there may be regions that locally

*e-mail: {kkyin|scoros|beaudoin|van@cs.ubc.ca}

trap the continuation process or there may also be topologically-isolated solution regions.

Our work looks at several alternatives for solving this type of continuation problem for simulation-based animation problems. A central issue is deciding how style optimization and continuation advancement should be coupled. We investigate gradient-descent, local-stochastic search, and hybrid continuation methods, each of which offer a different degree of coupling between these two problem features. The methods must satisfy the multiple goals of advancing the continuation process, staying out of failure regions, and optimizing stylistic aspects of the motion. These multiple criteria are not found in more classical optimization problems. For all the examples presented in this paper, the continuation variables represent physical parameters of the environment, such as step height or coefficient of friction.

2 Related Work

Trajectory-optimization methods have been demonstrated to be a promising way of synthesizing animation using a set of constraining keyframes and principles defined in an objective function. One of the principles is usually that the resulting motion should adhere to the laws of physics [Witkin and Kass 1988]. These methods have been extended to human body animation with the help of simplified physical models [Popović and Witkin 1999], initializations derived from motion capture data [Sulejmanpasic and Popović 2005], and reduced dimension subspaces [Safonova et al. 2004], among others. Given a family of solutions, the resulting motions can be kinematically interpolated in order to obtain a parameterized set of motions [Abe et al. 2004]. Our work shares the use of optimization procedures to compute motions that are adapted to new situations. It differs in that we adapt control system parameters rather than motion parameters. This makes it easy to model the effect of unanticipated changes to a motion, such as slipping, tripping, or pushing a piece of furniture that is lighter than expected. It also allows for interpolation between related motions, e.g., a large step and small step, at interactive rates and in a dynamically-consistent fashion.

An alternative approach to creating physics-based animations is to develop controllers that drive forward-dynamics simulations. While they can be challenging to design by algorithm or by hand, their feedback-based nature is reflective of control strategies that govern the motions of humans, animals, and robots. Controllers have been developed for locomotion [Raibert and Hodgins 1991; Stewart and Cremer 1992; van de Panne and Fiume 1993; Auslander et al. 1995; Grzeszczuk and Terzopoulos 1995; van de Panne 1996; Laszlo et al. 1996; Sharon and van de Panne 2005], human athletics [Hodgins et al. 1995; Wooten 1998], protective falls [Faloutsos et al. 2001], and, more recently, tracking motion capture data while maintaining balance [Yin et al. 2007; Sok et al. 2007; da Silva et al. 2008]. Unpredicted perturbations and environment changes can be accommodated to some extent by feedback mechanisms built into the controllers. A goal of this paper is to demonstrate how control strategies can be automatically adapted to large observable changes in the environment.

Prior examples of the use of continuation methods in animation include the use of external guidance forces for walking optimization [van de Panne and Lamouret 1995], the smooth introduction of coupling constraints when optimizing interacting multi-character motions [Liu et al. 2006], and the introduction of blurring factors to help guide the optimization of controllable smoke simulations [Treuille et al. 2003].

The work of [Hodgins and Pollard 1997] is the closest in spirit to our own. The parameters of a running controller are automatically adapted to runners having different physical parameters, i.e.,




Figure 2: (a) Abstract view of the continuation problem. The shaded regions denote failure regions where no meaningful gradient can be computed. γ is an environment-based continuation parameter. θ is the vector of free control parameters. (b) Gradient descent with fixed-step continuation (GRAD). (c) Stochastic local search continuation (STOC).

masses, inertia tensors, and dimensions. Beginning with the physical parameters and control system settings for a running man, an adapted control system is developed for a woman and a child. This was accomplished by manually advancing the continuation parameter, followed by a simulated annealing phase for style optimization using five control parameters. It is suggested that the continuation could be advanced automatically by advancing the continuation parameter until the system can no longer run for 10 seconds without falling. The objective function was crafted to capture the quality and appearance of the running motion. Goals of our work are to demonstrate automatic advancement of the continuation process, to investigate various design choices in doing so, and to show that continuation methods can be a powerful tool for generalizing controllers for many types of interactions with the environment.

3 Continuation Methods

There are many possible ways to approach the continuation problem posed in Figure 2(a). The actual performance of any given method will depend on the true shape of the failure regions and the optimization function in the high-dimensional parameter landscape. Following optimization function gradients may or may not regularly lead into failure regions. Greedy continuation may or may not lead into regions that then result in premature termination of the continuation. Given the many difficult-to-quantify issues and unknowns that characterize the continuation problems we would like to solve, our approach is to apply three continuation methods to five walking-skill adaptations and to evaluate the results. We now describe the three methods.

GRAD: One of the simplest approaches is to advance the continuation with a fixed step size $\Delta\gamma$, as illustrated in Figure 2(b). After each advancement, gradient descent is used to optimize the style of the motion, according to the style objective function. A centered finite-difference scheme is used to compute the gradient, which is followed by an inner iteration of line search to find the best step size to take in the downhill direction of the gradient [Nocedal and Wright 1999]. The gradient computation and line search procedure repeats until no further progress is made or until a maximum of 5 outer iterations. We also experimented with Knitro, a nonlinear optimization package, and found its performance to be comparable. An advantage of GRAD is that it tracks the optimal style θ^* in a systematic fashion. However, the gradients are expensive to compute, local minima are possible, and the solution can be easily led into a failure region where meaningful gradients can no longer be computed, such as when the character trips.

HYBRID: A second approach attempts to improve upon GRAD in several ways. A first change is to use a linearly-predicted new pa-

parameter value as a starting point after a continuation advancement, $\theta_{i+1} = \theta_i + (\gamma_{i+1} - \gamma_i)(\theta_i - \theta_{i-1})/(\gamma_i - \gamma_{i-1})$, instead of assuming the previous value, θ_i , where i is the index of the continuation step. This aims to provide a better starting point for the next gradient descent phase. A second change is to allow for adaptations to the continuation step. Regular sampling is applied in a trust region for the linear prediction, defined by $\delta\gamma \in [0.2\Delta\gamma, 2\Delta\gamma]$ in order to find the largest failure-free value. This allows for more rapid advancement of the continuation when possible, while allowing for smaller changes where required by problem difficulty. The sampling is done using the prediction scheme for θ . If all of the samples for the continuation search fail, we advance the continuation by $\delta\gamma = 0.2\Delta\gamma$ and initiate a stochastic local search on θ from this point.

STOC1, STOC2, STOC3: We test three versions of local stochastic search algorithm [Spall 2003] which treat continuation advancement and style optimization in a coupled fashion. A term $c_\gamma\gamma$ is added to the objective function (eqn. 1) in order to reward advances made in the continuation parameter. STOC1, STOC2, and STOC3 reward continuation to different degrees, given by $c_\gamma = c_0$, $c_\gamma = 10c_0$, and $c_\gamma = 100c_0$, respectively. The choice of c_0 is problem specific. Figure 2(c) illustrates a solution trajectory. At any given step, the search is advanced by drawing sample points in a uniform random fashion from a given window around the current solution: $(\gamma \in [\gamma_i, \gamma_i + \Delta\gamma/2], \theta \in [\theta_i - \Delta\theta, \theta_i + \Delta\theta])$. $\Delta\theta$ is a vector defining the window sizes for each of the parameter dimensions. For all of our examples, the search is conducted in a greedy fashion. The first sample yielding an improvement to the objective function is accepted. Advantages of this method include its ability to avoid failure regions and local minima within the windowed sampling region. The search is blind as to whether θ or γ yields the improvement to the objective function, and thus can trace a flexible solution path.

Evaluation: The methods can be evaluated according to various criteria. One is to examine how quickly and how far each method advances the continuation. A second is to look at the quality of the final solutions, as measured by the objective function at a fixed value of γ . Our results provide the necessary data for evaluating the methods according to these criteria.

Interpolation: The continuation process produces a solution path that embeds a family of solutions ranging from the original unadapted motion to the fully adapted motion. Linear interpolation between neighboring values of the continuation parameter can be used to generate a parameterized control action. It is usually not necessary to keep all the continuation steps in order to obtain a functional interpolation. This is discussed further in §5.

4 Problem Representation

Our five example adaptation problems share a number of common features. In each example, the continuation parameter, γ , represents a physical parameter of the environment, e.g., step height. In all cases, directly applying the default control parameters θ_0 to γ_T , will result in failure, e.g., a fall. There is no meaningful gradient if beginning at γ_T and the solution θ_T is generally far from θ_0 .

The default controller is an implementation of the four-state finite-state machine (FSM) walk controller described in [Yin et al. 2007]. Each of the states specifies target angles for all the joints, which are used by PD-controllers in order to compute applied joint torques. Two feedback-gain parameters, c_d and c_v , linearly modulate the swing leg hip angles as a function of the center-of-mass position (as measured with respect to the stance foot) and velocity. All states have a dwell-time in our implementation, including those based on foot contacts. Instead of transitioning to the next state immediately

upon foot contact, the transition now waits a further duration ΔT before making the transition. Any of the control parameters in this representation can in principle be included in the optimization. All of our examples treat a subset of the joint target angles for one or more FSM states as part of the parameter set Ω . It is also feasible to include the balance feedback parameters for the hip, c_d and c_v , and the state dwell time, ΔT . The FSM states that participate in the parameter set is problem specific.

The objective functions for our problems share a common structure:

$$f(\theta) = g(\theta) + w_\theta \delta\theta^T W \delta\theta + c_\gamma\gamma. \quad (1)$$

The first term rewards a desired motion style, to be described shortly. The second term penalizes making large changes to the original control parameters. \mathbf{W} is a diagonal weighting matrix (we use $\mathbf{W} = \mathbf{I}$), and $\delta\theta = \theta - \theta_0$. The optional third term rewards advancing the continuation in the case of STOC1, STOC2, and STOC3. We now detail the style term $g(\theta)$ and other relevant details for each of our example problems.

STEPUP: The goal is to adapt walking steps to take a large step up from one flat terrain onto another flat terrain. The step is introduced at a fixed distance from one of the walking steps and the step height is treated as the continuation parameter γ . The parameters for four states (one complete gait cycle, i.e., a right step followed by a left step) are designated to be modifiable. The default controller resumes action after these two steps, although the evaluation function also tracks the impact on the following four steps. Each state has 10 modifiable parameters, yielding a 40-dimensional parameter vector. The per-state parameters encompass the sagittal-plane target angles for the left-and-right ankles, left-and-right knees, the swing hip, and the waist. Four additional parameters are given by c_d (lateral and sagittal), the lateral swing-hip angle, and the state dwell-time. The style evaluation function is given by:

$$g(\theta) = \sum_{i=0}^1 (|\hat{l}_i - l_i|^2 + |\hat{h}_i - h_i|^2 + w_1 * c_i) + w_2 * \sum_{i=2}^5 (|\hat{l}_i - l_i|^2)$$

where i is a step index, \hat{l}_i is a desired step length, l_i is the actual step length, \hat{h}_i is the desired height of the swing foot center-of-mass when it passes the edge of the step, h_i is the corresponding actual height, c_i is the time integral, in seconds, of the combined unwanted contact durations of the swing foot with the environment, and w_1 and w_2 are constants. We use $\hat{h}_i = \gamma + 0.2m$, $w_1 = 10$, $w_2 = 0.1$, and $\hat{l}_i = (0.4, 0.1, 0.2, 0.2, 0.2, 0.2)$. The \hat{h}_i term helps inform the motion about the current step height. The desired step lengths for the series of four steps that follow the step up provide a way to guide the system towards the solution step-up strategies that lead to regular well-balanced steps afterwards. The c_i term assigns a penalty to unwanted foot contact, which helps avoid motions that may be marginally successful. We use $\Delta\gamma = 0.02m$. The evaluation takes place using a simulation of a total of six steps.

STEPOVER: The goal is to step over a thin obstacle or sill of variable height from a fixed location. The height of the sill is the continuation parameter. The choices of optimization parameters, optimization FSM states, and other parameters are the same as for STEPUP. The style term of the objective function is the same as that used for STEPUP.

PUSH: The task is to push a heavy object forward while walking at a regular pace. The density of the object is the continuation variable. We use $\Delta\gamma = 3.0kg/m^3$. It has a fixed volume $V = 0.72m^3$. The ground friction is set to $\mu = 0.8$ for both the character's feet and the object to be pushed. When the hands encounter the object, the PUSH controller is invoked. It is defined by only two modifiable states because left-right symmetry is used to define the remaining

two states of the FSM. The choice of parameters is largely the same as for STEPUP and STEPOVER. We add a left-right symmetric sagittal elbow angle and remove the lateral hip angle and the lateral, sagittal c_d parameters. This thus results in a 16-dimensional parameter space. The style term of the objective function rewards taking regular-length steps and is given by:

$$g(\theta) = \sum_{i=2}^7 (\|\hat{l} - l_i\|^2)$$

where $\hat{l} = 0.2m$. It is evaluated during the last six steps out of a series of eight simulated steps, which allows time for the simulation to achieve a steady-state behavior.

HILL: The task is to be able to walk up steep inclines. The slope of the incline, as measured in degrees, is the continuation variable. We use $\Delta\gamma = 3^\circ$. The modifiable parameters consist of the target sagittal angles for the left-and-right ankles, the left-and-right knees, swing hip, and waist, giving six parameters per state. As for the PUSH task, there are only two modifiable states because of left-right symmetry, yielding a 12-dimensional parameter vector. The modifiable controller is invoked for the first step before it is about to encounter the incline. No separate transition step is used. The style term is identical to that for PUSH, but evaluated only over the last four of eight simulated steps, and using $\hat{l} = 0.2m$ as measured along the incline.

ICE: The task is to walk on low-friction terrain. The coefficient of friction μ is the continuation variable, and we begin at $\mu = 0.18$, which is the lowest friction supported by the initial walking gait. We use $\Delta\gamma = -0.03$. The choice of parameters is similar to that of the STEPUP problem, but without the lateral and sagittal c_d parameters. The style term of the objective function rewards taking regular-length steps and is given by:

$$g(\theta) = \sum_{i=2}^7 (\|\hat{l} - l_i\|^2 + s_i^2)$$

where $\hat{l} = 0.2m$ and s_i is the foot-slippage, measured as the movement in the ground plane of the center of mass of the stance foot. This is evaluated using the last six of a series of eight steps.

5 Results

We evaluate the five continuation methods (three plus two variations on STOC) on the five example problems. Our simulations run at approximately real-time rates on a 2.20 GHz Intel Core Duo CPU T7500 laptop. One simulated second requires 0.8 s of wall-clock time, including rendering. The wall-clock time for 1000 simulations (equivalently, objective function evaluations) varies from 27-54 minutes for the example problems.

We plot the cumulative number of simulation-based function evaluations as a means of evaluating the performance of the five continuation methods. Figures 3-6 should not be confused with the abstract view of the continuation method given in Figure 2. The abstract parameter θ was illustrated as being one-dimensional, whereas our example problems have between 12 and 40 parameters. We also report on the final quality of the obtained solutions, as measured by the objective function $f(\theta)$ and sampled at the value of γ indicated by the dashed line. The final objective function values are illustrated using the bars on the right-hand side of the graph. The solution quality (bar length) is given as a ratio with respect to the HYBRID solution. Lower is better, and the HYBRID solution provides the best quality for all of our example problems. Note that we remove the value of the continuation reward term, $c_\gamma\gamma$, from

STOC1, STOC2, and STOC3 in order to provide a fair comparison. The number on the vertical axis beside each bar indicates the number of function evaluations, not the objective function value. If present, the red circle on the HYBRID curve indicates where local stochastic search is invoked. A red \times on the GRAD curve indicates a termination and that meaningful gradients can no longer be computed after this point.

STEPUP: Figure 3 shows the comparative evaluation for the STEPUP task while Figure 7 shows two example animations. The GRAD method fails at a step height larger than 11cm, while the other methods can continue to a step height of 65cm (beyond the shown 48cm). STOC1 manages a 27cm step, but is slow to make progress at this stage, most likely because the continuation reward is not enough to overcome the required change in style. STOC2 and STOC3 make similar continuation progress over time, despite STOC3 giving a continuation reward that is 10 \times that of STOC2. While the HYBRID solution is slower than STOC2 and STOC3, it produces a solution that is measurably better in terms of the objective function and also in terms of visual quality. The STOC3 solution is visibly more erratic than that produced by HYBRID or STOC2. The STOC2 solution is comparable to the HYBRID solution in quality, both numerically and visually.

STEPOVER: The evaluation graphs for the STEPOVER task are qualitatively similar to those of STEPUP, although the advancement of all methods is slower by roughly a factor of two for achieving a comparable height. Two related animations are shown in Figure 8. The GRAD method can do no better than a 7cm step. STOC2 and STOC3 advance the continuation very quickly for this problem. STOC1 offers insufficient continuation reward to progress beyond 14cm.

PUSH: The PUSH task can be accomplished using all five methods, although with variations in final quality, as shown in Figure 4. Example animations are shown in Figure 9. STOC1 gets stuck early, while STOC2 and HYBRID produce comparable quality solutions. STOC3 shows that very fast continuation can be achieved, albeit with lower quality solutions.

HILL: All the methods have some success in adapting a simulated walk to climb inclined slopes, as shown in Figure 5. Figure 10 shows two related animations. GRAD eventually fails after 22 $^\circ$. STOC1 stalls at a slope of 26 $^\circ$, where presumably the continuation incentive is insufficient to compensate for required further deviation in style. STOC2, STOC3, and HYBRID can all achieve slopes of 30 $^\circ$ and larger, with approximately equal quality.

ICE: HYBRID, STOC2, and STOC3 are all capable of achieving walks with friction coefficients as low as $\mu = 0.08$, as shown in Figure 6. An example motion is illustrated in Figure 1. HYBRID has the best solution, although it is an order of magnitude slower. The STOC3 solution is of significantly lower quality.

Motion variations can be obtained by changing the style term in the objective function. Figure 10 shows walking up an inclined slope with desired step-lengths of 30cm. The degrees of freedom used in the optimization will also affect the resulting motion. The top two rows of Figure 9 show two pushing motions with different subsets of joints being included in the set of modifiable parameters. As a result, different strategies are discovered.

The family of controllers defined by the continuation path can be interpolated online to generate controllers that can adapt to the environment in realtime. Figure 7 shows the stepping up controllers for two steps of different heights. The controllers for these examples use a point j on the solution continuation path such that $\gamma_j \leq \gamma \leq \gamma_{j+1}$. Linear interpolation is then used to compute $\theta(\gamma)$ from θ_j and θ_{j+1} . For PUSH and HILL linear interpolation be-

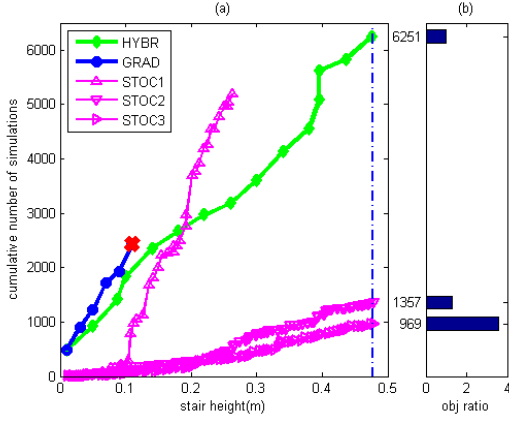


Figure 3: Evaluation for STEPUP walking.




Figure 5: Evaluation for HILL walking.

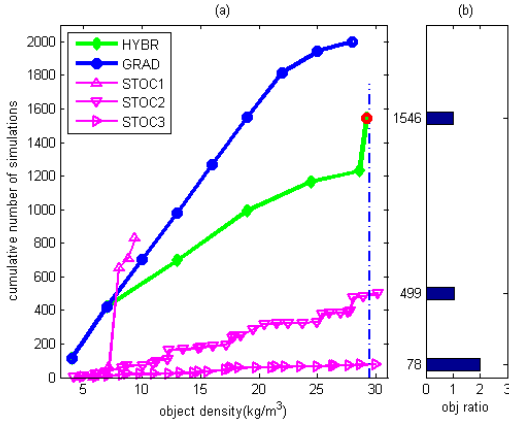


Figure 4: Evaluation for PUSH walking.

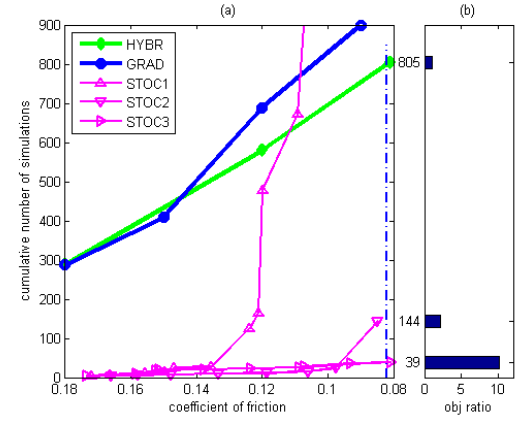


Figure 6: Evaluation for ICE walking.

tween the start and end points of the continuation is sufficient, while for the others it is not.

Mismatches between the expectations of a character and the reality of its environment can also be modeled. Figure 9(d) shows the result of applying the controller for pushing a 20 kg object to pushing a 10 kg object. The result is a faster motion with larger steps. The video also shows results for walking up slopes of 10° and 20° with a controller that expects a 30° slope.

6 Discussion

Designing objective functions and choosing the appropriate set of optimization parameters requires a user in the loop and some experimentation. A strategy that we have found to be effective is to begin with a desired stepping-length term as the style objective. Additional terms can then be designed with the help of a visual inspection of the results. For example, a stepping-width term can be added if the character is found to change stepping width in an undesired fashion. It can be useful to add a shaping term to the objective function, such as the swing-foot desired height term that is used in STEPUP and STEPOVER, and whose value is a function of γ . The HYBRID and STOC continuation techniques can be successful in its absence, although the continuation proceeds more quickly with the help of such shaping terms. The final objective functions for our five examples share a significant amount of structure.

A core issue in the design of the continuation methods is how the style optimization and continuation advancement should be coupled. First, they can be completely uncoupled, as in the case of GRAD. This effectively treats the problem as that of finding an independent optimal solution, $\theta(\gamma_i)$, for each of a regular sequential sampling of the continuation parameter, γ_i . θ_i is used as a good initial estimate of θ_{i+1} . However, using a gradient-descent method for the optimization of $\theta(\gamma_i)$ is expensive because gradients need to be evaluated using finite differences, and it is inadequate when a failure region is encountered.

Parameter optimization and continuation advancement are partially coupled in the case of HYBRID because of the use of a prediction when the continuation is advanced. The HYBRID method produces the best-quality solutions among our methods and avoids the failure problems of the gradient method. The partial coupling of the HYBRID approach works best if θ^* is at the bottom of a smooth winding valley in the optimization landscape. This is not necessarily the case, however, because the best solution to a given style objective for a fixed value of γ could be bounded by a failure region. As such, meaningful gradients and directions may not be always available and thereby necessitate some type of stochastic search.

Lastly, style optimization and advancement can be fully coupled, as is the case for STOC*. The introduction of a term that rewards continuation has the advantage of turning the problem into a single optimization problem with one extra variable. However, it can

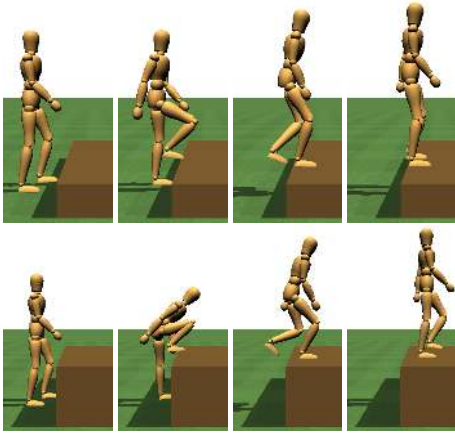


Figure 7: *STEPUP*: A 30cm step (left) and a 65cm step (right).




Figure 8: *STEPOVER*: A 29cm sill (top) and a 55cm sill (bottom).

be problematic because the results are sensitive to the weight given to the continuation term. Insufficient weight results in no advancement, while excessive weight results in greedy continuation, possibly erratic motions, and a non-smooth continuation path. For each example problem, we did initial experiments to select a value of c_0 , which is used as a continuation-reward scale parameter by each of STOC1, STOC2, and STOC3. The value was selected so that the STOC variations would span the range of too-slow and too-greedy behaviors. The experimentation to select this parameter is a disadvantage of the method. We have also experimented with using gradient descent with a continuation reward term. With a high reward, the optimization rapidly gets stuck in a local minima or failure region. With too low a reward, the optimization can actually drive the continuation backwards. Overall, it was difficult to tune and slow to advance the continuation parameter.

The choice of which continuation method to use depends on several factors. If the goal is to get the result as fast as possible, or as many styles as possible without having to touch the objective function, then STOC* should be used. This comes at the price of a solution that will potentially be less optimal in terms of the style objective and a need to choose the parameter that rewards the advancement of the continuation. The HYBRID method is recommended for offline computation, for producing more deterministic results, or when wanting to further control the style through the objective function. Based on our experience, we do not recommend pure gradient algorithms, i.e., GRAD.

The continuation methods and our evaluation methodology have a number of limitations. We have not exhaustively considered all possible variations of continuation methods. For example, an element




Figure 9: *PUSH*: Four different pushing motions. (a) Disallowing changes to the elbow. (b) Allowing changes to the elbow. (c) Pushing with a larger desired step length. (d) Pushing an object that is lighter than expected.




Figure 10: *HILL*: Uphill with a 30cm desired step length.

of prediction could be added to the stochastic methods in order to bias the sampling towards the area of the θ, γ -space that would represent a smooth extrapolation of the continuation path. This type of prediction exists in HYBRID, but not in STOC. It may be possible to achieve fast, greedy advances using STOC, and then further optimize the solution while holding the continuation parameter fixed. The parameter vector contains variables that have different units and to which the resulting motion will display varying sensitivity. It may be possible to use estimates of such sensitivities to automatically rescale the window used for the local stochastic-search elements of the HYBRID and STOC methods.

We do not currently evaluate how well the continuation path follows θ^* , as shown in Figure 2(a). Finding this curve is unfortunately non-trivial because of possible local minima in the high-dimensional parameter space as well as the presence of irregularly-shaped failure regions in this space. The effectiveness of the method may be dependent on the choice of the underlying control representation. In our case, the continuous and robust balance adaptations made by the controller [Yin et al. 2007] help to avoid falling as changes are made to the motion.

In some cases, it may be difficult to attribute the final form or style of a solution to particular parameters of the optimization problem. A particular type of solution may be favored because of an objective function term, the choice of the set of motion parameters to be optimized, random choices during STOC, or other factors such as joint limits, strength limits, balance constraints, etc. Thus, answering the question “why did it do this in this way?” will not always be easy. A related question might be “why does the continuation get stuck at some particular limiting value of the continuation parameter?” This reflects the potential complexity of the interactions between motion

constraints and the resulting motions. Another limitation is that parts of the solution space that are topologically isolated by a large enough unfeasible region will be unreachable. STOC* is limited to making jumps within its search window, which has dimensions appropriate for local search.

7 Conclusions

Generalizing controllers to new situations is an important tool in developing larger skill sets for physically-simulated characters. Continuation-based methods provide a natural way to adapt controllers to large changes in the environment. Continuation methods can be used to achieve surprisingly large motion adaptations. Problems that are seemingly very different, such as pushing furniture and walking on ice, can be treated within the same framework.

This paper has explored a number of the underlying issues and choices related to their use for problems in simulation-based animation. Results have been demonstrated for climbing large steps, stepping over tall sills, pushing heavy objects, climbing steep inclines, and walking on ice. Simple style objective functions are often adequate. For problems such as step-up and step-over, a shaping term in the objective function is helpful.

In the future, we wish to investigate multidimensional continuation methods. A simple example is that of being able to generate controllers that can climb up a step of varying height, at varying distances from the current stance foot, with varying weight in a back-pack and perhaps with varying coefficients of friction. An interesting question is the extent to which the result of the multidimensional continuation can be represented in a factored form that supports linear superposition. These types of factorizations would be of great help in coping with the infinite types of motion variations and motion constraints that may arise in practice. We also wish to make the motions as natural as possible and to optimize for robustness in addition to style.

Acknowledgements

We thank the anonymous reviewers for their comments and suggestions. Funding from NSERC (Natural Sciences and Engineering Research Council of Canada) is gratefully acknowledged.

References

ABE, Y., LIU, C. K., AND POPOVIC, Z. 2004. Momentum-based parameterization of dynamic character motion. In *Proc. Symposium on Computer Animation*, 173–182.

AUSLANDER, J., FUKUNAGA, A., PARTOVI, H., CHRISTENSEN, J., HSU, L., REISS, P., SHUMAN, A., MARKS, J., AND NGO, J. T. 1995. Further experience with controller-based automatic motion synthesis for articulated figures. *ACM Trans. on Graphics* 14, 4, 311–336.

DA SILVA, M., ABE, Y., AND POPOVIĆ, J. 2008. Simulation of human motion data using short-horizon model-predictive control. *Computer Graphics Forum (Proc. Eurographics)* 27, 2.

FALOUTSOS, P., VAN DE PANNE, M., AND TERZOPOULOS, D. 2001. Composable controllers for physics-based character animation. *Proc. ACM SIGGRAPH*, 251–260.

GRZESZCZUK, R., AND TERZOPOULOS, D. 1995. Automated learning of muscle-actuated locomotion through control abstraction. *Proc. ACM SIGGRAPH*, 63–70.

HODGINS, J. K., AND POLLARD, N. S. 1997. Adapting simulated behaviors for new characters. In *Proc. ACM SIGGRAPH*, 153–162.

HODGINS, J. K., WOOTEN, W. L., BROGAN, D. C., AND O'BRIEN, J. F. 1995. Animating human athletics. In *Proc. ACM SIGGRAPH*, 71–78.

LASZLO, J. F., VAN DE PANNE, M., AND FIUME, E. 1996. Limit cycle control and its application to the animation of balancing and walking. In *Proc. ACM SIGGRAPH*, 155–162.

LIU, C. K., HERTZMANN, A., AND POPOVIĆ, Z. 2006. Composition of complex optimal multi-character motions. In *Proc. Symposium on Computer animation*, 215–222.

NOCEDAL, J., AND WRIGHT, S. J. 1999. *Numerical Optimization*. Springer.

POPOVIĆ, Z., AND WITKIN, A. 1999. Physically based motion transformation. In *Proc. ACM SIGGRAPH*, 11–20.

RAIBERT, M. H., AND HODGINS, J. K. 1991. Animation of dynamic legged locomotion. In *Proc. ACM SIGGRAPH*, 349–358.

SAFONOVA, A., HODGINS, J. K., AND POLLARD, N. S. 2004. Synthesizing physically realistic human motion in low-dimensional, behavior-specific spaces. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 514–521.

SEEGER, A. 2006. *Recent Advances in Optimization*. Springer.

SHARON, D., AND VAN DE PANNE, M. 2005. Synthesis of controllers for stylized planar bipedal walking. In *International Conference on Robotics and Automation*.

SOK, K. W., KIM, M., AND LEE, J. 2007. Simulating biped behaviors from human motion data. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, Article 107.

SPALL, J. C. 2003. *Introduction to Stochastic Search and Optimization*. Wiley.

STEWART, A., AND CREMER, J. 1992. Beyond keyframing: An algorithmic approach to animation. In *Proceedings of Graphics Interface*, 273–281.

SULEJMANPASIC, A., AND POPOVIĆ, J. 2005. Adaptation of performed ballistic motion. *ACM Trans. on Graphics* 24, 1, 165–179.

TREUILLE, A., MCNAMARA, A., POPOVIĆ, Z., AND STAM, J. 2003. Keyframe control of smoke simulations. In *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 716–723.

VAN DE PANNE, M., AND FIUME, E. 1993. Sensor-actuator networks. In *Proc. ACM SIGGRAPH*, 335–342.

VAN DE PANNE, M., AND LAMOURET, A. 1995. Guided optimization for balanced locomotion. In *Proc. EG Workshop on Computer Animation and Simulation*, 165–177.

VAN DE PANNE, M. 1996. Parameterized gait synthesis. *IEEE Computer Graphics and Applications* 16, 2 (March), 40–69.

WITKIN, A., AND KASS, M. 1988. Spacetime constraints. In *Proc. ACM SIGGRAPH*, 159–168.

WOOTEN, W. L. 1998. *Simulation of leaping, tumbling, landing, and balancing humans*. PhD thesis, Georgia Inst. Technology.

YIN, K., LOKEN, K., AND VAN DE PANNE, M. 2007. Simbi-con: Simple biped locomotion control. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, Article 105.