A. Juarez-Perez and M. Kallmann, Modeling Data-Based Mobility Controllers – To appear in the Proceedings of Digital Human Modeling 2016, Montreal. This is the manuscript of the authors and not the final publication.

Modeling Data-Based Mobility Controllers with Known Coverage and Quality Properties

ALAIN JUAREZ-PEREZ *† and MARCELO KALLMANN†

† University of California, Merced

Abstract

Synthesizing controllable and realistic locomotion and stepping motions for the support of manipulation is a challenging computational problem that is important for simulating a variety of whole-body tasks. We propose a methodology for parameterizing primitive data-based mobility controllers with known coverage and quality characteristics. Our approach is computationally efficient and achieves realism and high controllability of stepping behaviors, thus addressing key properties for supporting a variety of whole-body manipulation tasks. The proposed mobility controllers are based on deformation operations applied to a captured motion clip of one step cycle, such that a small number of deformations achieve full parameterization of the end location of one step. Controllers are then quantified with respect to end location coverage and motion quality, in order to allow higher level decision mechanisms to select and employ the best mobility controllers to support the intended manipulation task.

Keywords: *locomotion and stepping control, whole-body manipulation.*

1. Introduction

While reaching and grasping constitute the main components of an object manipulation, there are several situations in which precise coordination between manipulation and body mobility is crucial for achieving a virtual human model capable of positioning itself in a suitable location for performing the intended manipulation.

Correct body placement is critical for performing a number of manipulation tasks and its importance can be observed even in very simple situations. For example, if a virtual human model is slightly far away from a target to be reached, the spine will need to bend towards the target in order for the target to be reachable. However, a simple step forward would probably eliminate the need to move the spine and would bring the model to a more comfortable location to perform the action. Moving to a suitable position for manipulation is often preferable than performing a strenuous manipulation motion.

Body positioning for manipulation is particularly important in whole-body manipulation tasks. The problem can be observed in a number of situations such as when operating wide control panels or large machinery, collecting items in warehouses, assembling large structures, opening doors, etc. A number of factors influencing a chosen position have been observed [Land et al. 2013], indicating a complex whole-body posture planning process in anticipation to a manipulation task.

In this paper we focus on the particular problem of parameterizing primitive data-based mobility controllers with known coverage and quality properties, in order to allow higher level body planning mechanisms to select and employ the best mobility controllers to support a given whole-body manipulation task.

The proposed mobility controllers are based on deformation operations applied to a captured motion clip of a single step motion. The motion can be one cycle of a patterned locomotion clip, such as walking, or a single step towards a rest pose, such as by a lateral or frontal step. We use the term mobility to refer to both locomotion and stepping behaviors.

A small number of motion deformations are then defined in order to achieve full parameterization of the end location of each mobility controller. Controllers are quantified with respect to end location coverage and motion quality, providing valuable information for controller selection during whole-body manipulation planning. Initial results

*Corresponding author. Email: ajuarez-perez@ucmerced.edu

obtained from applying the proposed mobility controllers for the task of opening and passing through a door are also presented.

2. Related Work

Significant previous work in computer animation has been dedicated to developing controllable animation based on motion capture data. One possible approach for achieving realistic results among constraints is to rely on global search methods that reuse the captured data in several ways [Kovar et al. 2002] [Choi et al. 2003] [Arikan et al. 2003] [Mahmudi and Kallmann 2015]. The main drawback of such methods is that they involve time-consuming search procedures with limited applicability to real-time applications. Another difficulty faced by some of these methods is that a significant amount of motion capture data may be required in order to be able to address complex environments or other types of constraints.

Methods suitable for real-time applications will typically rely on a collection of motion clips specifically organized for being concatenated in order to achieve a continuous stream of motions [Gleicher et al. 2003] [Park et al., 2004], and the overall approach can also incorporate parameterization strategies for achieving precise control of the modelled behaviors [Mukai et al. 2005] [Heck et al. 2007]. Such approaches can be suitable for controlling the mobility of virtual characters; however, motion parameterization still requires a significant amount of well-designed motion clips to be interpolated in order to achieve an effective parameterization of the involved motions. A related approach for parameterization is to rely on reinforcement learning [Treuille et al. 2007], however requiring significant learning time over a suitable collection of motions.

Our proposed approach is based on simple deformation strategies applied to a single motion capture clip. The simplicity of the approach allows it to scale well to multiple controllers integrated as sub-components under a higher-level whole-body motion planner. The proposed approach is based on motion adaptation techniques which can be efficiently computed and thus are suitable for achieving real-time results.

In this paper we are not only interested in achieving a real-time data-based mobility controller but also in defining a methodology for quantifying control coverage and quality of the produced motions. While previous work has not yet specifically addressed quantification metrics for mobility controllers, the area of evaluating perception of motion quality has received increased attention in recent years. For instance, Prazak et al. [2010] and Ryall et al. [2012] have evaluated the perception of time-warping motion operations, and Prazak et al. [2011] have investigated the perception of *footskating* artifacts. While in this paper our current quality metrics are limited to validity tests, our methodology can be easily extended to also include perceptual metrics.

The primary interest of the proposed mobility controller is to offer precise control with known quality for whole-body manipulation planners. In computer animation different whole-body motion planners have been proposed [Bai et al. 2012] [Huang and Kallmann, 2016] which require the existence of a suitable locomotion controller. Previous work in digital human modeling has also recognized the importance of mobility in manipulation tasks and the difficulty in addressing the mobility component [Wagner 2006]. This paper proposes a methodology with the potential to improve generic approaches to full-body motion planners. The principles of our approach were previously presented in a poster [Juarez-Perez et al. 2014]. In this paper we fully present our method together with several results.

In summary, while several previous methods have investigated different techniques capable of achieving realistic mobility controllers, our proposed approach consists of a simple methodology that includes precise determination and quantification of coverage and quality properties in order to support the needs of whole-body manipulation planners.

3. Data-Based Mobility Controllers

Using a full-body motion capture suit we have collected a series of stepping motions and we have devised a methodology for parameterizing captured stepping motions or locomotion cycles. Each step cycle of interest is segmented out of the motion and processed in order to obtain a parameterized primitive for stepping towards any point inside a neighborhood around the original end-point of the captured clip. Independently if the clip represents one locomotion cycle or a single stepping motion, clips can be self-concatenated, which means that the final frame of the clip should be similar to the first frame of the clip. We refer to the clip as the motion cycle.

Cycle Deformation and Parameterization

Each motion capture clip fully specifies the motion cycle. It describes the topology of the skeleton, the location of its joints, and the joint values for each frame of the motion. The joint values are the joint rotations at each frame, and a special joint, called the *root* joint, also contains the position of the skeleton in world coordinates. The root joint is usually placed between the hips of the character and its rotation values describe the overall skeleton orientation.

The root position and its rotation component around the vertical axis are the main elements manipulated by our deformation operations. By controlling these parameters we are able to deform the overall path that the animation follows. Any undesired artifacts (like feet sliding) that are eventually introduced in the motion are corrected algorithmically, as later described in this section.

Each captured cycle can be deformed in 3 different ways: forward stretching or compression, left-right trajectory turning, or left-right final body orientation change. These three operations introduce a parameterization of the original motion, but at the cost of potentially reducing the realism of the original motion. Our analysis methodology, presented later, will represent the tradeoff between coverage and quality. During real-time deformation operations, the system will always correct the root joint trajectory such that a continuous self-concatenated motion cycle is always achieved.

Deformations are applied over the root position and orientation of the character skeleton hierarchy. We consider the root to be the joint to which the hips and spine are connected to in the humanoid skeleton representation. Consider the motion cycle ground trajectory being represented by M as a vector containing indexed root position and orientation values $\rho_i = \{x_i, z_i, \theta_i\}$ for every frame i of the motion, where θ_i represents the yaw rotation (or the body orientation) of the character and (x_i, z_i) is the projected position of the root joint on the floor plane. Therefore the full motion trajectory is represented by $M = \{\rho_i\}$, i ={1, ..., N}, where N is the number of frames in the motion cycle.

We perform three kinds of deformations:

- **Trajectory Turning.** To achieve turning, we update the motion frames with the result of applying a 2D rotation by a given rotation parameter φ to $v_i = (\mathrm{x_i}, \mathrm{z_i}) - (\mathrm{x_{i-1}}, \mathrm{z_{i-1}})$, and adding φ to θ_i. The 2D rotation by φ is a rotation around the vertical axis of the character. The trajectory of the character will turn to the left or to the right according to the rotation parameter.

- **Compression or Stretching**. Translation operations are appied at each frame i of amount s over the displacement vector $v_i = (\mathrm{x_i}, \mathrm{z_i}) - (\mathrm{x_{i-1}}, \mathrm{z_{i-1}})$, by updating the position of each $(\mathrm{x_i}, \mathrm{z_i})$ such that the length of the displacement vector is updated to become $\|v_i\|$+s after the displacement is applied. This is achieved by replacing $(\mathrm{x_i}, \mathrm{z_i})$ with $(\mathrm{x_i}, \mathrm{z_i}) + s \cdot \frac{v_i}{||v_i||}$. The resulting deformation will either compress or strecht the original motion. This transformation includes an optional speed adjustment, where the timing of the keytimes associated with the frames can be multiplied by a factor proportional to the deformation. This allows to maintain the speed as the original motion speed while streching or compressing the motion, or to allow to change the speed according to the performed deformation.

- **Final Orientation**. In order to change the body orientation at the end of the clip, at every frame, the value of ψ is added to θ_i. The final orientation of the character changes but the trajectory of the motion remains the same.

The result of the described deformations is illustrated in Figure 1. In order to achieve a parameterization independent of the number of frames in a clip, the parameters sent to the entry point of each deformation procedure are considered to correspond to the whole motion and thus they are multiplied by a factor of $\frac{1}{N}$, where N is the number of frames in the motion clip, in order to obtain the per-frame deformation parameters that are used by each described deformation procedure.

These deformations can be applied to a large variety of stepping motions and together they are able to fully control the final location of the step cycle while maintaining the same style of the captured cycle. Note that the control is on the final position and orientation of the skeleton root joint position projected on the floor, and not on individual footstep locations, what would impose a higher number of parameters to be controlled.

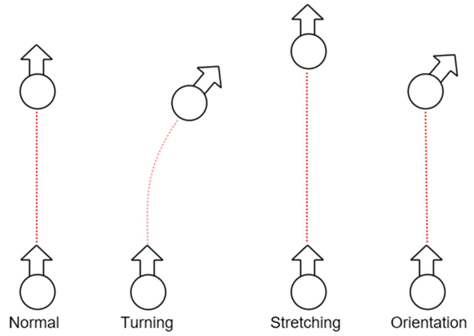

Figure 1: Deformations over a motion cycle. The left-most diagram corresponds to the original clip cycle. The next ones represent the outcome of each deformation operation.

For example, with a 90 degrees turning angle and a −90 degrees orientation angle, we end up with a lateral stepping motion cycle, although not realistic and with self-collisions introduced. If we apply a compression deformation such that the total ground displacement becomes 0, and also apply a given trajectory turning deformation value, we obtain in-place turning. By providing a constant turning rotation, we achieve the character performing a circular trajectory. Clearly, "extreme" motions will not look realistic and we later introduce coverage-quality maps in order to determine quality boundaries for the control space.

Enforcement of Feet Constraints

The presented deformation operations provide parameterization of the motion cycle but at the same time introducing several undesirable effects, in particular related to breaking original feet-ground constraints. These discontinuities are corrected with a Inverse Kinematics (IK) and blending operations.

We first annotate the moment at which the ball of the foot reaches a stable ground support position in the motion (achieving foot contact) and the moment at which it loses this property (foot release). In this way, the parameterization deformations will take advantage of a possible small twist motion introduced over the first contact point with the floor. Each foot will have four different states: Transitioning, Landing, Supporting and Departing.

By using IK corrections and blending operations over the stages we ensure continuity over the motion while maintaining the capability of deforming the motions as desired. Each foot cycle is processed as follows:

- **Transitioning Phase**. During this phase one foot is transitioning in the air between two ground contact points while the other foot remains fixed on the ground as the support foot. During this phase no constraint correction is needed for the transitioning leg; however, IK correction for the upcoming contact will gradually update the leg joints, in an ease-in blending fashion, as it approaches its next ground contact point. This will ensure that the transitioning foot lands exactly at the expected next ground contact point, which is re-computed at every pose based on the current user-controlled deformation parameters.

- **Landing Phase**. This phase starts at the first pose initiating ground-foot contact, right after the corresponding foot finishes its transitioning phase. During this phase the target end-effector for the IK continues to be updated according to any real-time deformations that may happen, and IK continues to be applied to correct the degrees of freedom of the leg until the foot stops and becomes completely fixed on the ground by IK.

- **Supporting Phase**. During this phase the supporting foot is completely enforced by IK correction to the position and orientation of the established ground contact. The original leg posture in the motion cycle data is used to inform the IK solver to produce a solution as close as possible to the original data.

- **Departing Phase**. This phase starts at the time of foot ground release, when the system gradually interpolates out of the IK-corrected configuration and returns towards the original data from the captured motion clip.

These four phases represent natural stages observed in human motion. Whenever we suddenly change speed or direction, the foot on the landing phase applies some extra force to maintain the balance while the position and orientation of the root is deformed, making a slight movement over the initial contact of the step. The Supporting phase is necessary to displace the body, and the Departing phase completes the cycle making the motion ready for concatenation. See Figure 2.

Continuous Cycles with Blending Operations

The process of concatenating the motion cycle in order to achieve continuous stepping control will naturally introduce discontinuities at the connection point because the original unprocessed motion cycle clip will never perfectly self-concatenate. To address this problem we apply motion blending at the transition points.

We perform a special blending operation for each step during the concatenation phase. The transitioning leg in the transitioning phase only requires an ease-in and ease-out of the motion but the leg that is in support/contact requires an extra correction. We compute the global matrix $\mathbf{F}$ of the contact foot at the end of the clip, and the global matrix $\mathbf{G}$ of the same foot at the start of the following clip. If concatenation blending is being performed, there is going to be a twisting and sliding motion at the foot during the concatenation, and so we use matrix $\tilde{\mathbf{F}} = \mathbf{F}\mathbf{G}^{-1}$ over the following clip in order to correct the global position of that foot and eliminate any undesired foot sliding or twisting effect.

Figure 2 illustrates the main phases and the blending operations that are performed. Blending is used to perform two operations: concatenation and feet constraint enforcement. All corrections are performed in real-time so that an input motion cycle can be directly used by the system without the need of special pre-processing. IK correction is performed when there is a transition from and to a contact foot, and is always gradually applied, via blending with the original data values.

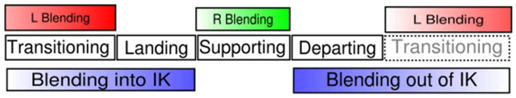

Figure 2 – Blending operations according to cycle phases.

4. Coverage-Quality Maps

Clearly, when large deformations are imposed on the motion cycle, a number of problems may arise. We perform several measurements in order to quantify the quality of the motion after a deformation is performed.

Constraint Enforcement

If the distance from the IK end-effector goal to the hip joint, when maintaining feet constraints, is greater than the maximum reachable distance, a solution will never be found. However, small errors might be acceptable and an approximate IK solution may be good enough for cases where the difference is small. Because of the nature of the error, the best way to quantify it is to measure the sum of the distance errors whenever the IK cannot compute the exact solution. This gives us a quality metric inversely proportional to the IK error.

Continuity

The second possible problem when applying IK is to ensure continuity over subsequent frames which were independently corrected. Every frame is corrected by IK seeking to be as close as possible to the original data; however, it is still possible that consecutive frames have a noticeable difference between them. Instead of relying on expensive IK formulations for improving smoothness, our IK operations are kept analytical and highly efficient, and our overall approach will discard the extreme deformations that generate discontinuous motions.

Collisions

Finally, some of the deformations will cause the legs to collide with each other, and a collision detection test is required in order to reject these cases and achieve correct results. In our system each rigid segment of the leg is represented by a geometric capsule primitive for the purpose of collision detection. Mesh-based collision detection can also be employed when working with skinned characters.

The metrics described above are then used to quantify coverage and quality. Given a motion clip to be parameterized and deformed, the entire parameter space can be sampled and the corresponding quality for each 3-tuple deformation parameter can be quantified according to the metrics described above.

5. Results and Discussion

Figure 3 illustrates a typical coverage-quality map obtained for a step forward mobility controller with respect to the stretching/compression and turning deformations. The dark blue region represents the coverage of the fully valid region according to the employed metrics. The original motion clip is a single locomotion cycle forward step, where the starting position is at the center of the diagram with the virtual human model facing upwards. The dark blue region therefore indicates that a large frontal region can be reached by the presented deformation operations without imposing any significant loss of quality, according to the employed metrics.

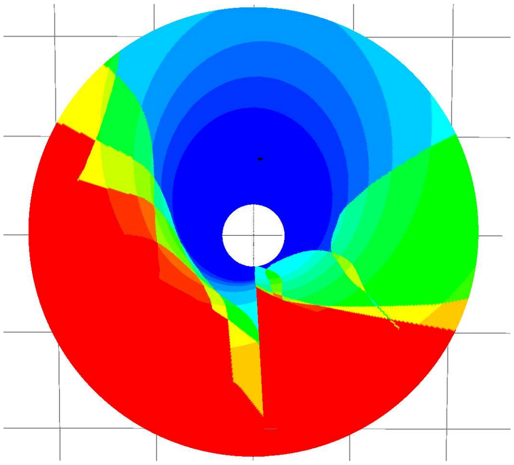

Figure 3: Coverage-quality map for one forward step cycle motion.

The colors that are not dark blue in Figure 3 represent regions where too much deformation was applied causing invalid IK operations or self-collisions to occur. The range of blue colors, from dark to light blue, represents the amount of IK errors detected, quantified as the number of frames where either a foot constraint or continuity between frames was not perfectly enforced. The range of colors green, yellow, orange and red represent the number of frames detected to have collisions. Green, yellow and orange represent 1, 2, and 3 frames with collisions respectively, and red means that 4 or more collisions were detected. The designed color visualization scheme allows a straight forward visualization of the coverage achieved for the desired motion quality output.

Figure 3 visualizes the coverage of the first two deformation operations. In Figure 4 we present the coverage achieved by the body orientation deformation.

The colors in Figure 4, for each final (x,y) body position, evaluate the quality of the full range of possible final body orientations. Dark blue represents positions at which the range of orientations between -75 and +75 degrees, sampled at 1 degree, are fully valid. Lighter blue to green are areas that have some invalid deformations but mostly valid ones. Red areas are the ones in which most of the deformation parameters return an invalid clip.

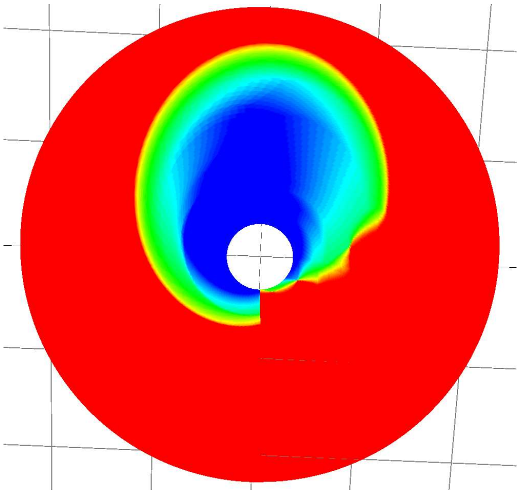

Figure 4: Final body orientation (after forward step) coverage-quality map.

Figure 5 uses the same color scheme as in Figure 3, but applied to a lateral step. The original motion starts from the center of the diagram and performs a single lateral step to the right, with the end location being the point marked in black. It is possible to note that the valid deformation region is not symmetric and is larger in the area ahead of the character than in the area behind of the character. This probably corresponds to the body posture in the original motion clip slightly bending forward, in order to favor capacity to reach supporting location targets in the frontal region.

Figure 6 presents a volumetric visualization of the coverage-quality volume produced by the three deformation operations considered together.

One first utility of the presented coverage-quality maps is to inform how much control space can be addressed by a single mobility clip. The analysis easily allows us to determine new mobility clips in order to achieve a collection of clips fully covering the entire region around the character with high-quality (dark blue) motions. It is possible to observe that a small set of forward, backward, and lateral steps will often fully cover the entire region around the character with dark blue.

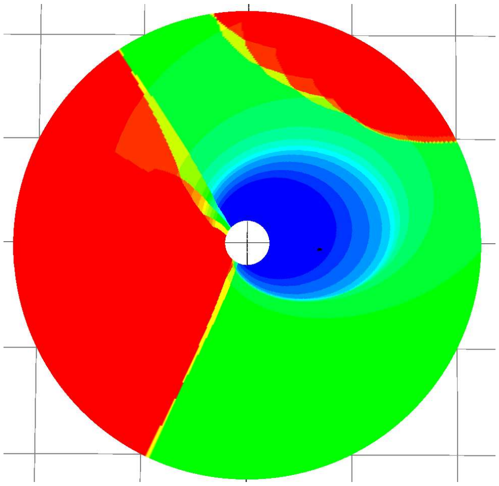

Figure 5: Coverage-quality map of one right lateral step.

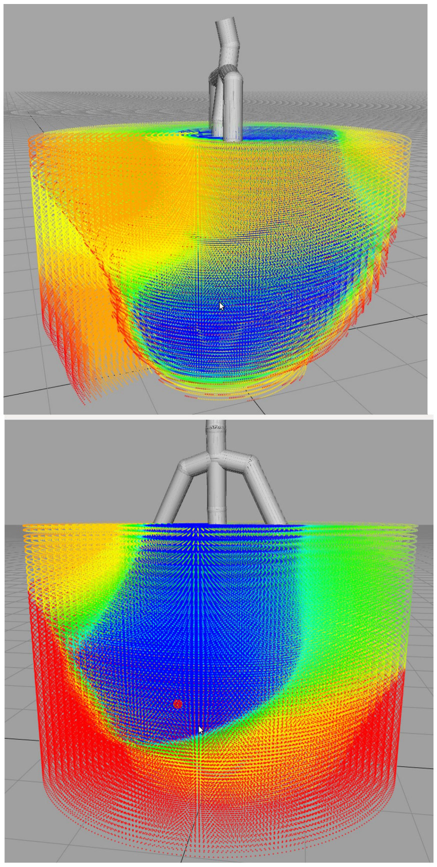

Figure 6: Volumetric coverage-quality map for a frontal step. The horizontal layers are planar maps equivalent to the one in Figure 3, with the vertical axis representing changes to the final body orientation deformation parameter.

The proposed set of deformation strategies enables not only parameterization of the end location of discrete steps, but also full on-line control of locomotion. For instance, Figure 7 illustrates our system being used to control a character turning around following keyboard commands executed by the user during simulation. Key presses are mapped to amounts of deformation in order to steer the character. Every time a key is pressed the current amount of deformation is incremented or decremented, and the corresponding deformation operation is applied instantly starting at the next frame being executed, with the corresponding blending operations being always activated.

The achieved flexibility makes the proposed mobility controller suitable for integration with higher-level planners for achieving whole body manipulations. Figure 8 presents initial results obtained during the synthesis of whole-body motions for opening and passing through a door. The mobility controller was used to generate a path passing through the door that is both feasible and close enough to the handle of the door in order for the handle to be successfully reached by IK applied on the arm supposed to manipulate the door handle.

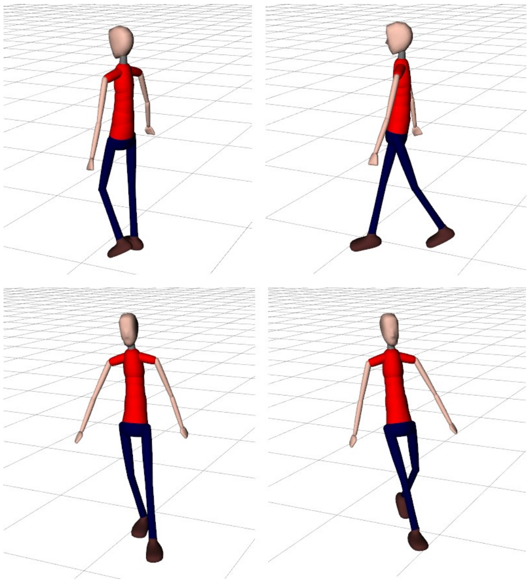

Figure 7: Example of real-time control of locomotion forming a circular trajectory.

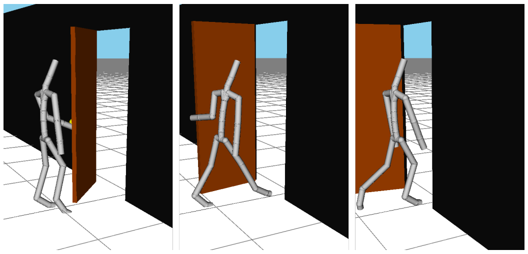

Figure 8: Example of autonomously opening a door.

Our current whole-body manipulation algorithm works as follows. Paths connecting the current location of the character to a location on the other side of the door are randomly generated. For each path, the locomotion mobility controller is applied along the path and collision tests with the environment are performed along with reachability tests with respect to the door handle. After a few iterations a successful motion is achieved. The coverage-quality map is used to control the maximum curvature accepted by the generated paths, such that the forward walking motion stays in its high-quality dark blue region. During motion testing, turning deformations are applied in order to ensure that the generated path is closely followed.

The deformation operations are fast and therefore allow the planner to generate complex motions under a second of computation time. The described approach represents a promising application of the proposed overall mobility synthesis and analysis methodology. Several optimizations and improvements are currently being developed, with the goal of achieving complex whole-body manipulations coordinated with locomotion and autonomous body positioning.

A number of additional directions are being explored in terms of future work. The ability to visualize spatial coverage and quality not only informs building a collection of mobility controllers achieving high-quality coverage, but also introduces a generic methodology for analyzing quality metrics spatially. We intend to include in our methodology both perceptual and ergonomic metrics to the quality evaluation of the produced motions. Perceptual metrics can be specified with the help of human studies, where bounds on each deformation operation can be specified perceptually in order to guarantee perceptually-validated high-quality motions. Ergonomic metrics can be incorporated with the inclusion of movement energy expenditure metrics and as well stability or body balance metrics. In addition, similar metrics computed over the upper-body motions generated by IK during manipulation can also be incorporated in order to extend the coverage-quality analysis to address manipulation motions.

6. Conclusions

The presented mobility controllers are able to successfully provide real-time control with known coverage and quality. The employed techniques represent computationally efficient solutions making the mobility controllers fast enough to be used to evaluate several possible trajectories for supporting whole-body manipulations. Our initial results with the problem of manipulating doors are promising and we are currently developing improved whole-body planners based on the described mobility controllers.

Acknowledgements

This work was partially supported by NSF Award CNS-1305196.

References

Arikan O, Forsyth D A, O'Brien, J F, 2003. Motion synthesis from annotations. Proceedings of SIGGRAPH, 22(3):402-408.

Bai Y, Siu K, Liu C K, 2012. Synthesis of concurrent object manipulation tasks. ACM Transactions on Graphics, 31(6):156:1-156:9.

Choi M, Lee J, Shin S, 2003. Planning biped locomotion using motion capture data. ACM Transactions on Graphics, 22(2):182-203.

Gleicher M, Shin H J, Kovar L, Jepsen A, 2003. Snap-together motion: assembling run-time animations. In Proceedings of the symposium on Interactive 3D graphics and Games (I3D), pages 181-188, NY, USA.

Heck R, Gleicher M, 2007. Parametric motion graphs. In Proceedings of the 2007 symposium on Interactive 3D graphics and games, I3D '07, pages 129-136, New York, NY, USA. ACM.

Huang Y, Kallmann M, 2016. Planning Motions and Placements for Virtual Demonstrators. IEEE Transactions on Visualization and Computer Graphics (TVCG), 2016

Juarez-Perez A, Feng A, Kallmann M, Shapiro A, 2014. Deformation, parameterization and analysis of a single locomotion cycle. In Proceedings of the Seventh International Conference on Motion in Games, MIG '14, pages 182-182, New York, NY, USA. ACM.

Kovar L, Gleicher M, and Pighin F, 2002. Motion Graphs. ACM Transactions on Graphics 21(3):473-482, Proceedings of SIGGRAPH 2002.

Land W M, Rosenbaum A, Seegelke C, Schack T, 2013. Whole-body posture planning in anticipation of a manual prehension task: Prospective and retrospective effects. Acta Psychologica 144, 298-307.

Mahmudi M, Kallmann M, 2015. Multi-Modal Data-Driven Motion Planning and Synthesis. Motion in Games 2015

Mukai T, Kuriyama S, 2005. Geostatistical motion interpolation. ACM Transactions on Graphics. 24, 3, 2005, 1062-1070, Proceedings of SIGGRAPH 2005.

Park S I, Kim T, Shin H J, 2004. Online motion blending for real-time locomotion generation. Computer animation and virtual worlds 15.

Prazák M, Hoyet L, O'Sullivan C, 2011. Perceptual evaluation of footskate cleanup. In Proceedings of the 2011 ACM SIGGRAPH/Eurographics Symposium on Computer Animation, SCA '11, pages 287-294, New York, NY, USA. ACM.

Prazák M, McDonnell R, O'Sullivan C, 2010. Perceptual evaluation of human animation timewarping. In ACM SIGGRAPH ASIA 2010 Sketches, SA '10, pages 30:1-30:2, New York, NY, USA.

Ryall K, Hoyet L, Hodgins J, O'Sullivan C, 2012. Exploring sensitivity to time-warped biological motion. Perception ECVP abstract, 2012.

Treuille A, Lee Y, Popović Z, 2007. Near-optimal character animation with continuous control. ACM Transaction on Graphics 26(3), July 2007, Proceedings of SIGGRAPH 2007.

Wagner D, Reed M, Chaffin D, 2006. A Task-Based Stepping Behavior Model for Digital Human Models. University of Michigan, DHM 2006.