

Mikiko Matsunaga, Victor B. Zordan
University of California, Riverside

A Dynamics-based Comparison Metric for Motion Graphs

the date of receipt and acceptance should be inserted later

Abstract Motion graph approaches focus on the re-use of character animation contained in a motion-capture repository by connecting similar frames in the database with transitions. Because the output animation of a motion graph comes directly from the motion capture data except for the transitions, the quality of the motion depends largely on the transition points selected. In this paper, we investigate comparison metrics for choosing transition points, aiming at improving the visual quality of animations generated using motion graphs. Specifically, we focus on the weight assigned to each body part, which reflects the relative significance of the body part on the quality of the generated motion. We introduce a novel weighting scheme, based on an estimation of the character's dynamics, which assigns weights for each body according to displaced mass and simplified friction terms. To assess the quality of the transitions selected by our proposed dynamic metric, we compare its results to previous methods, looking at both visual quality and quantitative analysis.

Keywords animation, motion capture, character motion

1 Introduction

Realistic animation of humanlike characters is an active research area with important applications in the movie and game industries. One approach for generating realistic, controllable motion is through the use of a motion graph, which gained a great deal of visibility at Siggraph 2002 [1, 8, 10, 12, 16]. Motion graphs are defined as directed graphs built by cutting and connecting segments of motion capture data. They retain visual quality by using motion capture data whenever possible and allowing transitions to take place when motions snippets (or frames) are similar enough. Because these transition points are selected based on the costs returned by a comparison metric, the choice of the metric has direct impact on the quality of the constructed motion graph.

In this paper, we aim at improving the visual quality of motion graphs by focusing on the comparison metric used to find the transition points.

While there are several components that make up a comparison metric, we focus on the weighting parameters assigned to individual body parts that are used in pairwise frame comparisons. In particular, we set out to define an assignment scheme that automatically selects weight values for individual body parts based on pertinent, quantifiable dynamic properties that affect the quality of the transitions. To meet our goal, we approximate the character's dynamics by accounting for mass displacement across a transition as well as simple constraints derived from ground contact and friction. Our assertion is that a good transition is one which matches circumstances and minimizes violations based on the dynamic conditions seen by the character over the duration of the proposed transition period. Intuitively, our dynamics based distance metric penalizes large (mass) displacements and unnatural sliding contacts with the ground (e.g. "foot skate".)

While we focus on motion comparison weights in this paper, we strive for mechanisms for evaluating both the metric and the quality of the motion produced using the motion graph approach. Thus, a secondary leg of this effort emphasizes analysis and evaluation of the produced motion transitions. In a straightforward manner, our algorithm's selected transitions are evaluated in comparison to two previous methods based on visual quality and quantitative analysis. In addition, we introduce findings that provide insights in regards to assessment of distance metrics in general as well as a quantification of the number of transitions one might expect to find based on a given database. The latter leads to a method for selecting metric thresholds.

The major contributions of this paper are:

- An automatic, dynamics-based assignment scheme for comparison.
- Introduction of body mass displacement as a method for transition evaluation.

- Quantitative reduction in amount of “foot skate” resulting from motion transitions.
- Assessment and practical suggestions related to the selection and use of comparison metrics for motion graphs

2 Related Work

Several researchers have proposed ways to reuse motion capture data by treating it as an input to build new animations. One such technique is to create animations by taking disjoint segments of motion capture data and connecting them through a simple transition [14, 1, 8, 10, 12, 16, 13, 2]. Recently, Ikemoto et al. also introduced a method for generating transitions automatically between any frame and any other frame in a motion database using multi-way blends [7].

In our work, we focus on the distance metric used to evaluate motion transitions and highlight a few distinct methods here for reference. In Kovar et al. [8], a motion graph’s transition points are selected by searching for similar frames in the motion database using a comparison metric which is the Euclidean distance between representative trace points distributed over the character’s body model. A ‘window’ of action is incorporated into the metric by taking the weighted sum of the distances for several frames before and after the times of the test frames. We use this method as our baseline (Heuristic 1) in this paper. Arikan and Forsyth construct a hierarchical motion graph and perform a randomized search on it to generate a motion that fulfills user constraints [1]. In their effort, the comparison of frames is performed by aligning the roots and using the joint position, joint velocity, torso velocity, and torso acceleration. The joint weights are varied depending on the significance, as in our metric, but the values are chosen empirically in their metric. We construct a second heuristic measure based on their technique for comparison as well. Lee et al build a two-layer graph structure using a Hidden Markov Model [10]. The comparison metric used by Lee et al is based on the difference in joint angles and joint velocities, where the sum is taken over manually selected “important” joints. They also introduce the notion of contact states, suggesting that the transitions should take place only when there is a similar contact change occurring at the frame. We also account for contact changes in our weight definition.

Several studies have been conducted to evaluate synthesized motions including edits or transitions for motion capture data. Evaluations of motion graphs can be divided into two categories: analyzing the quality of individual transitions and analyzing the resulting animation as a whole, for example, based on the coverage within an environment. Several investigations rely on human input to assess naturalness of synthesized motions - where some assess human perception of segment quality via user studies [17, 22], while others build statistical models to automatically assess quality based on hand-labelled examples [21, 19, 3, 6]. Another study evaluates the physical correctness of interpolated motion seg-

ments and makes specific suggestions in regards to generating more natural looking motion [20]. Evaluation targeted at the results from the motion graphs in its entirety has been presented by Reitsma and Pollard [18]. They introduced a metric which measures the capability of the motions generated by a motion graph to navigate a character within a given environment.

Of these efforts, the work described by Wang and Bodenheimer [21] is related to ours in that they evaluate transition cost (following the metric introduced by Lee et al [10]) and focus on the weight assignment scheme. They present an algorithm for assigning joint weights using a training set of transitions that are classified manually as visually pleasing or not. The weights are computed via optimization to minimize the contribution from the undesirable transitions. Their work is similar to ours in that we both focus on weighting schemes that lead to natural transitions, but instead of using example transitions to optimize the weights, we determine the weights based on the dynamics of a given motion segment. This distinction enables the weight assignment scheme to be automated and does not require a manually generated training set. In this respect, our work is more similar to that described by Safonova and Hodgins [20] because we compute a dynamics based error that penalizes physically unrealistic transitions.

3 Comparison Evaluation

There are various techniques that can be used to compare motion segments. In this paper, we base the structure of our comparison metric on the one presented by Kovar et al. [8] and investigate the effect of using different joint weights. Thus, we begin the description of our method with an overview of their comparison metric. We also re-implement their metric and use this as a basis for comparison.

According to their comparison method, to compute the distance errors for a pair of frames, the distance for a window of frames is extracted and compared. This window is a sequence of frames containing the frame itself and the nearby frames, whose size is defined as $L + R + 1$, where L and R denote the sizes of the left and right windows, respectively. The cost between two windows is defined as the weighted sum of the costs for individual pairs of frames in the window:

$$C(A_i, B_j) = \sum_{k=-L}^R [W_k * c(A_{i+k}, T_{\theta, (x,y)} B_{j+k})]$$

Here, $C(A_i, B_j)$ denotes the cost for transitioning from a window centered around frame A_i to another around frame B_j . W_k is a changing frame weight for the k th frame in the window. In our implementation, we set this value based on a quadratic function where the highest weight is given to the central frame ($k = 0$) and decreasing weight is given toward the ends of the window. $c(A_{i+k}, T_{\theta, (x,y)} B_{j+k})$ denotes the cost for a single frame pair, A_{i+k} , and B_{j+k} , given by the frame-frame comparison function c , and $T_{\theta, (x,y)}$ is the alignment

transformation applied to the frame B_{j+k} . The alignment transformation is a rotation along the vertical axis followed by the translation along the ground plane, and it is computed as an analytical minimum of the final cost.

Frame-frame comparison assesses the similarity of two postures, returning the quantitative difference between the two. The cost between two postures is defined as the weighted sum of the cost for each body part:

$$c(f_i, f_j) = \sum_{n=1}^N \frac{w_{pos}(f_{i,n}, f_{j,n})}{M} \sum_{m=1}^M c_{pos}(f_{i,n,m}, f_{j,n,m})$$

Here, the cost is defined for postures at frames f_i and f_j , each having N body parts. For each body, the positional difference of M representative points (termed ‘point clouds’) are computed using c_{pos} , which denotes the squared Euclidean distance between the corresponding points. The computed costs are scaled by weight w for each body, which is uniform in their approach. The representative points for each body part are weighted equally, so the average of the sum of M costs is used as the cost for the body part represented by the points. For the sake of consistency, unless otherwise specified, we follow their technique exactly except for our replacement of the weights in w .

4 Dynamics-based Weighting

To account for the dynamics, we estimate the value of the mass for each body part. This approximation can be done based on volume and density estimates or set directly based on known values for human body parts. For further details on this topic, see [5]. From these masses, we could compute the amount force or change in momentum required to move a body part based on a given input trajectory (say, the proposed motion transition). However, unless we are willing to make assumptions in regards to the external forces are acting, these calculations become intractable. Instead, we propose that the dynamics can be captured by computing two mass-based properties that account for the unconstrained mass displacement and constrained, *friction-based* sliding present in the compared motions.

Thus, the comparison metric we introduce is comprised of two distinct components: mass displacement weights and ‘friction’ weights. Then, the weight for a given body part b is defined as the sum of these two types of weights:

$$w(b) = w_{mass}(b) + w_{friction}(b)$$

where $w_{mass}(b)$ is a measure of the amount of mass that is displaced by moving a body a given amount and $w_{friction}(b)$ measures the amount of resistance felt by moving a body against friction, imposed by ground contact.

4.1 Mass Displacement Weight

Mass displacement has been used to assess motion edits previously [15]. In our scheme, the mass displacement weight




Fig. 1 Inward/outward directions with respect to the root. Root body is shown in dark.

is determined by the amount of mass that gets displaced by moving the body part assuming that the character is unconstrained (i.e. not in contact with the ground). Note, kinematic constraints are upheld, but we do not differentiate at this point based on the environment.

Because the human body is connected by joints, body parts do not move independently. For example, if a person tries to move the upper arm and the hand will move together, and if a person tries to move the lower arm, the hand will move together in a similar manner. We note the upper arm displaces more mass than moving the lower arm, which suggests that differences in the shoulder joint require more work than the differences in the elbow in general. To incorporate such differences, we define our mass displacement weight as:

$$w_{mass}(b) = Mass(b) + \sum_{b' \in \{OutParts(b)\}} Mass(b')$$

where mass displacement weight for body b is defined as the sum of the masses of the part itself plus all of the outward parts b' , where the outward direction is defined as the direction away from the root. In other words, the outward parts include all body parts that lie between b and the branch tips in the outward direction (see Figure 1). In this fashion, we account for the relative cost of moving the body parts, for example the hand vs. the pelvis based on the total mass displaced.

Note, the values of w_{mass} remain constant throughout the simulation. The value of $w_{mass}(b)$ only depends on the masses of the body parts in the model, which does not change from frame to frame. Thus, $w_{mass}(b)$ can be pre-computed for each body part in the model to avoid redundant calculations.

4.2 Friction Weight

The mass displacement weight accounts for movement in an unconstrained environment. Obviously, it is also important to account for external forces, especially for ground contact which supports the character’s body and provides traction for movement. One solution is to use the body parts that are in contact with the ground, which suffice as an approximation to avoid complex calculations [23]. However, body

parts in direct contact with the ground are not the only ones that support the body or feel the result of friction resistance. Thus, in addition to the body parts in direct contact with the ground, the inward parts of the supporting limb's branch should be penalized as well.

$$InPartsInSameBranch(b) = \{b' : b' \in InParts(b) \text{ and } Branch(b') = Branch(b)\}$$

Here, the inward bodies, *InParts*, are all of the body parts from *b* toward the root, which is the direct opposite of the outward parts mentioned in the previous subsection. The *Branch(b)* returns the branch that body part *b* belongs to, which is either the head or one of the four limbs. The second half of the check ensures that the body part *b* only affects other parts of the body that belongs to the same branch. The supporting body for the frame is then defined as all of the body parts in ground contact and all inward parts of them that belong to the same branch:

$$SupportingBody(f) = \{b : IsContactingGround(b)\} \cup \{b' : b' \in InPartsInSameBranch(b)\}$$

The friction weight is defined to be the amount of mass displaced, in this case, based on the amount of mass the body part supports. To compute the amount of mass a body supports, we draw off of our observation that a body in contact can be treated as the root and then, as we did before, we can walk out the tree from that root collecting the masses of each body displaced - up to the entire mass of the body as in the case of a single foot support. Then, as we move up the chain to the other bodies identified as support bodies (but not in contact), like the shank, we apply the same algorithm. One way to calculate this approximation is to take the total mass of the body and subtract out the masses of the outward bodies associated with that support body as

$$\begin{aligned} TotalMass &= \sum_{b \in \{BodyParts\}} Mass(b) \\ BranchMass(b) &= \\ &Mass(b) + \sum_{b' \in \{OutPartsInSameBranch(b)\}} Mass(b') \\ w_{friction}(b) &= TotalMass - BranchMass(b) \end{aligned}$$

The *OutPartsInSameBranch(b)* returns the set of body parts that are in the outward direction from the body part *b* and also belongs to the same branch, which is the opposite of the *InPartsInSameBranch* introduced earlier. The branch mass is defined to be the sum of the masses of the outward parts in same branch plus the mass of the body part *b* itself, and $w_{friction}(b)$ is found by subtracting branch mass from the total mass of the body.

Note, the friction weight is computed in the same units as the mass displacement weight and therefore no careful tuning must be done to combine these two weights. Also, because the friction weight depends on a single posture, the weight need only be computed once for each body part for each frame in the repository.

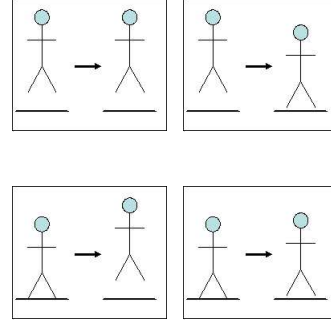


Fig. 2 Possible transitions. From top to bottom, left to right: a) unconstrained to unconstrained, b) unconstrained to constrained, c) constrained to unconstrained, d) constrained to constrained.

4.3 Weight for Frame Pair

The remaining question on weight assignment is to determine how to assign weights for pairs of postures. So far, the weights are assigned for just a single posture, but the cost needs to be computed between pairs of postures. One solution is to use the weights for the source posture, so if transitioning from f_i to f_j is being considered, the weights for f_i are used and vice versa for the opposite direction (from f_j to f_i). However, this does not differentiate between the cases where the transition destination is constrained (Figure 2(b), 2(d)) and unconstrained (Figure 2(a), 2(c)). It is important to make this distinction, since the transition from constrained to constrained (Figure 2(d)) is more restrictive than the transition from constrained to unconstrained (Figure 2(c)), making the former case more prone to undesirable visual artifacts such as foot skate.

To reflect the effect of both postures, the weight for a posture pair is computed as:

$$w_{pair}(f_i, f_j) = w_{mass} + w_{friction, f_i} + w_{friction, f_j}$$

Here, the weight is computed for frames f_i and f_j , which is defined as the sum of the mass displacement weight and the two supporting weights for the two frames. In the definition, the mass displacement weight is added only once, noting that its value is equivalent for both postures. By using the sum, w_{pair} gets contribution from both postures, resulting in assigning highest weight to the transition that goes from constrained frame to constrained frame.

5 Implementation

As a means of comparison, we implemented two heuristic-based weight metrics in addition to our dynamics-based one. The three metrics differ only in the weight assignment scheme of the body-part weights. Heuristic 1 uses uniform weight and follows the metric suggested in [8]. It makes no distinction between the body parts. Heuristic 2 uses ad-hoc, non-uniform weights, where higher weights are assigned to the


torso than the limbs. This metric is more similar to those described by Lee et al. [2002] and Arikan and Forsyth [2002], although we still incorporate the point cloud and windowing concepts throughout. For this metric, we also assign an extra weight to the single body part that is supporting most of the mass, where the extra amount is a hand-selected constant. The difference between our metric and Heuristic 2 is that our weights are determined based on mass values and our metric computes values for supporting bodies per frame based on any combination of contacts in an automated fashion, whereas the weights for Heuristic 2 are static and defined manually.

The comparison metrics are used to construct motion graphs for two motion databases, one consisting of a single long sequence of 4145 frames (about 34 seconds) of martial arts fighting and the other consisting a 3720 frames (about 31 seconds) of concatenated, shorter reaction behaviors such as stumbling, stepping, and falling. We also tested a data sequence for walking but found that each metric performed indistinguishably well because of the regular periodicity of the motion, so we chose to leave this dataset out of our analysis. All of the results are derived using a character model having 45 DOF, 3 each for 13 joints and 6 for the position and orientation of the root. The reaction based animations were not included in the final video because, while the transitions found by our metric were visually more appealing, the motions concatenated across the board lead to fairly nonsensical animations due to the nature of the behaviors (not the transitions.) As such, we chose to include the database for the analysis of the transitions, but we do not show any animations from this dataset so as not to taint the reader's perception of the results.


To view animations, blending is performed by applying linear interpolation to the joint positions and spherical linear interpolation to the joint orientations using an ease-in ease-out weighting function. Simple interpolation is used here to ensure that the visual quality of the blended sequence depends mainly on the closeness of the selected transition points and not on the blending technique used. Applying foot-skate cleanup will result in better-looking motion than what is being presented. The supplement video shows a collection of the example transitions from each set difference, ordered by their ranks. By looking at the transitions, it can be seen that more foot-skating is happening on those selected by the Heuristic 1. These transitions are obviously selected because the source and the target frames are close in terms of the positional differences of their body. But they are not chosen by our dynamics scheme because the transition happens between two constrained, double support postures and the metric assigns a high penalty from friction weights to both of the legs.

6 Results and Analysis


Figure 6 shows the normalized 2D cost matrices for the fighting database, comparing our metric and Heuristic 1. A 2D



(a) Costs from dynamics-based metric



(b) Costs from Heuristic 1



(c) Difference (Dynamics - Heuristic 1)

Fig. 3 2D cost matrices for martial arts fighting data. Intensities represent normalized costs, where low intensity (dark spots) represent a close match and high intensity (bright spots) represent a poor match. The difference matrix (c) shows difference in normalized costs. Positive values are shown in blue, negative values are shown in red, and zero difference is shown in black. Points on (c) denote the transition points chosen by: our dynamics-based metric (green), Heuristic 1 metric (magenta), and both metrics (white).

cost matrix is a symmetric matrix whose entries denote the costs, so an entry at $[i, j]$ having a value c means that the cost to transition from i th frame of source clip to j th frame of target clip is c . The computed costs are normalized by the median, and these values are used as the intensities in the graph. The grid pattern in both reveals the structure of the sequence as the fighter throws fast attacks and then returns to a neutral protective stance. The dark spots having low intensity correspond to low cost (a good match), and the bright spots having high intensity correspond to high cost. One distinction is that the dynamics metric leads to a crisper, higher contrast distance matrix and hints at the metric’s ability to distinguish between what it constitutes as good and bad transitions. The last matrix (c) shows the difference between the normalized costs for the two metrics, where the areas in red denote the transitions favored by our dynamics-based metric, and the areas in blue denote the transitions favored by Heuristic 1 (again intensity is used to show the size of the difference).

The 2D difference matrix reveals large discrepancies as well as interesting patterns in the relative costs assigned, yet it alone is not sufficient as a means for assessing quality. Because the cost metric is used to determine the transition points, it is important to observe the difference in the transition points selected by the two metrics. To find the difference, T best transition points are selected for each of the metrics, returning the sets $TransPt_{dynamics}$ and $TransPt_{Heuristic1}$. Then, set differences are taken in both directions, excluding all nearby transitions that represent the transitions which connect the same motions (within a quarter of a second.) These excluded transitions are selected by both metrics and appear in the difference matrix as white dots. The remaining two sets, shown in magenta and green in the difference matrix, reveal the transitions selected by one but not the other. These transitions allow us to most strongly distinguish between the visual aspects of the transitions selected. Thus, we include examples of these transitions in the accompanying video for the reader viewing. For fighting, the 67 “best” transitions were selected and of these 11 were unique to each metric. For the reactions, 30 transitions were allowed and 14 were unique to each.

Selecting the threshold for the cut-off between acceptable and unacceptable transitions can be tricky and varies with the same distance metric from database to database in our experience. Thus, we came up with a helpful mechanism for selecting this threshold. For our results, the number of best transition points, T , is determined based on the rank-ordered costs returned by the comparison metrics. Figure 4 shows the normalized costs for fighting sorted in the increasing order of the rank (duplicate transitions within a short time window are thrown out, hence the total number of transitions varies and is less than the total number of frames). The value T is defined to be the point at which the slope of the graph(s) change(s) after an initial ramp up. We reached this practice empirically based on the observation that independent of the weighting scheme (and other factors) the database revealed a singularity point (around rank 67) after which the slope

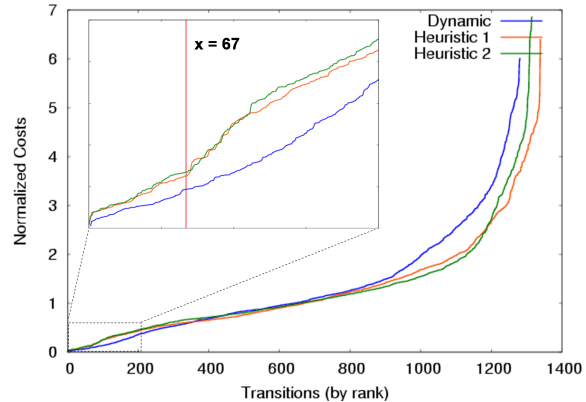


Fig. 4 Shows the normalized costs for fighting using the metrics. The red vertical line shows the average point where the slopes of the graph changes.

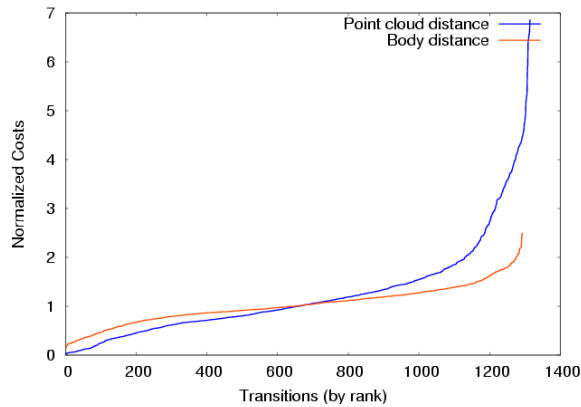
	Start to End		Arc Length	
	Fighting	Reacting	Fighting	Reacting
Dynamics	9.0	8.2	10.6	9.9
Heuristic 1	12.2	16.8	14.0	19.8
Heuristic 2	12.8	17.0	14.2	19.3

Table 1 Average foot-skate distance (in cm) for all the transitions chosen by the comparison metrics. Start-to-end measures only the distance of the end points for the transitions. Arc-length approximates the path length taken over the duration of each transition.

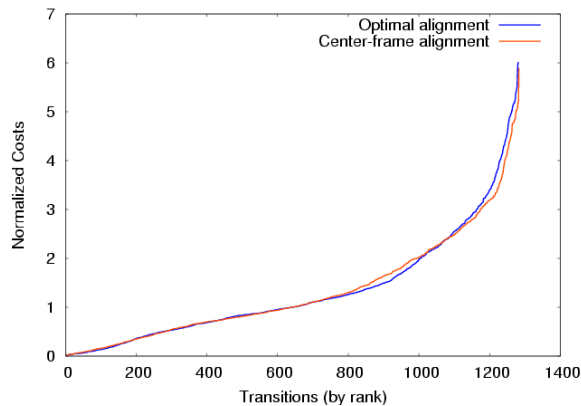
of the error values changed. Testing the transitions on both sides of this singularity revealed it as a reasonable cut-off between good and bad transitions. A similar exercise lead us to our cut-off for the reaction database, set at rank 30. Our interpretation of this result is that a given database only has so many places where segments repeat themselves within some proximity and that after the best transitions are picked off, the values of the for the ranks change in a noticeable manner.

Note, although it may seem that the transitions selected between metrics have a good deal of overlap and indeed many of them do connect logically similar motion segments, the order of the rankings of individual transitions is quite different across the metrics and the precise characteristics of the transitions selected are also quantifiably different even though they connect similar segments. One measure of this difference is the amount of foot skate present in the transitions selected. The average foot skate distance of the T best transitions are computed and the results are shown in Table 1. From this analysis we can see that our dynamics-based metric selects transitions with reduced the amount of foot-skating on average.

We include two additional graphs that attempt to reveal the effect of other components of the metric. Figure 5(a) shows the graphs of the costs computed by using the position of the point clouds and by using the position and the orientation of the body (center of mass.) The costs computed using the point clouds has greater deviation, which makes the transitions selected based on such cost to be more reliable



(a) Point cloud vs spatial position



(b) Comparison of alignment

Fig. 5 Comparison of different components of a comparison metric. Graphs show the normalized costs by differing one of the components.

because the transitions are more distinct. Thus, we conclude that point clouds are preferable, potentially only because they combine the position and rotation information into a single unified value. We also performed an analysis of two different alignment schemes. Figure 5(b) shows the comparison between the two techniques for aligning the frame windows, one using the optimal alignment as described by Kovar et al. [8] and another using the center frame of the frame window for alignment. The resulting graphs have similar structure, suggesting that the choice between the two alignment techniques is insignificant in terms of the costs that they return.

7 Discussion and Conclusions

As one can observe from the associated animations, a notable trend in the transition points chosen only by Heuristic 1 is that they involve a displacement of one or more highly

constrained (support) body parts. Such transition points are not chosen by our dynamic weighting scheme because weights assigned to the constrained body part makes the resulting cost high, discouraging the transition from being selected compared to others with less costly displacements (for example where only arm movement is required in a standing transition.) In contrast, the uniform weighting of Heuristic 1 makes no distinction between the two, since equal weight is assigned to both constrained and unconstrained body parts. Thus, it is no surprise that the foot skate observed is significantly smaller using our method.

A less obvious trend that results from our approach is that the transitions selected by our method tend to take place during periods of high activity (rather than during a static pose), see the filmstrips in Figures 6(a) and 6(b). This is likely because they include a shift in (body) weight which our weighting scheme exploits. That is, when an action such as a kick or leap takes place, there is a break in double support and this yields a ‘sweet spot’ when a transition can be made with minimal foot skate. We also believe based on experience that transitioning during active periods is preferable because viewers tend to be more forgiving of flaws during this period. One subtlety with the presented weight-assignment scheme is that it does not take the direction into account. For example, lifting an arm up is more difficult than letting it fall, suggesting that the direction of the movement affects the definition of weights in addition to what is already incorporated. We leave this improvement for future work.

In conclusion, we present a dynamics based weighting scheme for comparison of transitions selected for a motion graph and fit our scheme into an existing framework for assessment. We compare results of our algorithm to that of two alternatives designed to be similar to previous approaches for two databases and show a reduction of foot skate artifacts across all cases. In addition, we present some general analysis and suggestions related to selecting thresholds for metrics and the sensitivity of metrics to factors such as alignment and the method of computing distances between body parts. While human perception will remain the ultimate test of visual quality, we believe our method leads to quantitative (and qualitative) improvement based on inherent computable dynamic characteristics associated with motion transitions (and natural motion, in general) and that our results yield higher quality motion without the need for human intervention.

References

1. Arıkan, O., Forsyth, D.A.: Interactive motion generation from examples. In: Proceedings of ACM SIGGRAPH ’02, *ACM Transactions on Graphics*, vol. 21(3), pp. 483–490. ACM Press (2002)
2. Arıkan, O., Forsyth, D.A., O’Brien, J.F.: Motion synthesis from annotations. In: Proceedings of ACM SIGGRAPH ’03, *ACM Transactions on Graphics*, vol. 22(3), pp. 402–408. ACM Press (2003)
3. Arıkan, O., Forsyth, D.A., O’Brien, J.F.: Pushing people around. In: ACM SIGGRAPH / Eurographics Symposium on Computer Animation, pp. 59–66 (2005)

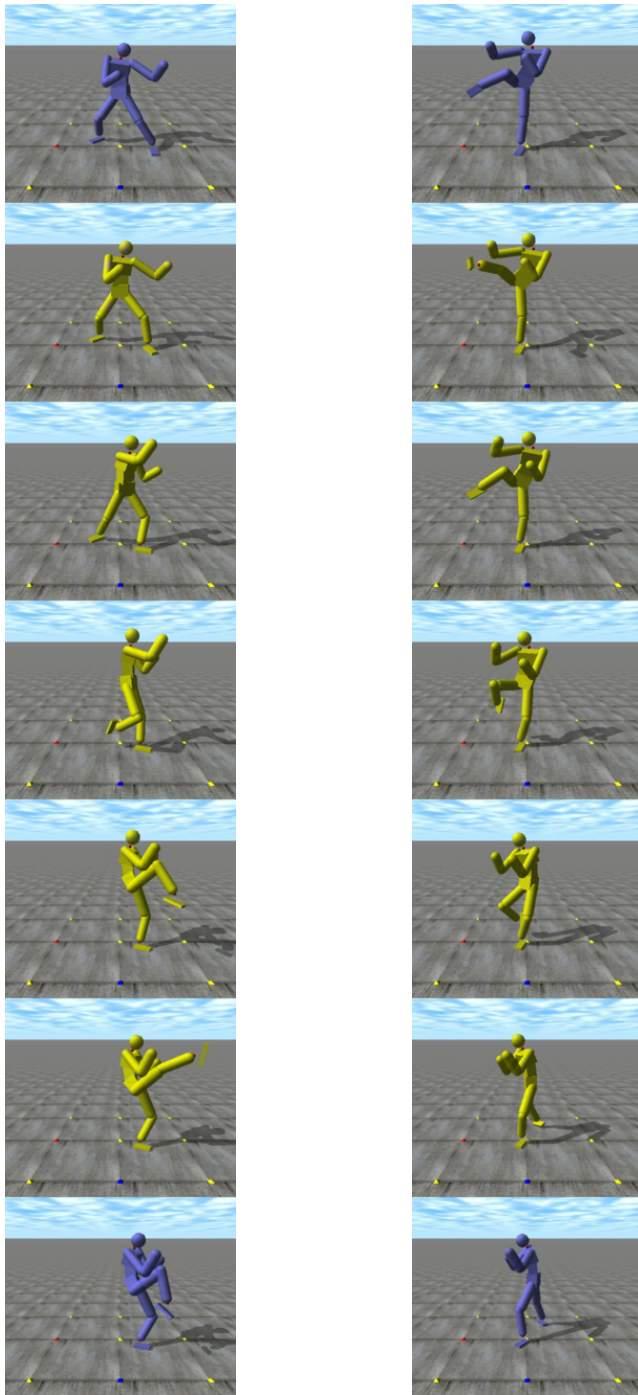


Fig. 6 Sample transitions selected by our dynamics based metric. Yellow segments show the blended sequences.

4. Galata, A., Johnson, N., Hogg, D.: Learning variable-length Markov models of behavior. *Computer Vision and Image Understanding: CVIU* **81**(3), 398–413 (2001)
5. Hodgins, J.K., Wooten, W.L., Brogan, D.C., O'Brien, J.F.: Animating human athletics. In: *Proceedings of SIGGRAPH '95*, pp. 71–78. ACM SIGGRAPH (1995)
6. Ikemoto, L., Arikan, O., Forsyth, D.A.: Knowing when to put your foot down. In: *Symposium on Interactive 3D Graphics and Games (I3D)* (2005)
7. Ikemoto, L., Arikan, O., Forsyth, D.A.: Quick motion transitions with cached multi-way blends. In: *EECS UC Berkeley Technical Report, UCB/EECS-2006-14* (2006)
8. Kovar, L., Gleicher, M., Pighin, F.: Motion graphs. In: *Proceedings of ACM SIGGRAPH '02, ACM Transactions on Graphics*, vol. 21(3), pp. 473–482. ACM Press (2002)
9. Kovar, L., Schreiner, J., Gleicher, M.: Footskate cleanup for motion capture editing. In: *ACM SIGGRAPH Symposium on Computer Animation*, pp. 97–104 (2002)
10. Lee, J., Chai, J., Reitsma, P.S.A., Hodgins, J.K., Pollard, N.S.: Interactive control of avatars animated with human motion data. In: *Proceedings of ACM SIGGRAPH '02, ACM Transactions on Graphics*, vol. 21(3), pp. 491–500. ACM Press (2002)
11. Lee, J., Shin, S.Y.: A hierarchical approach to interactive motion editing for human-like figures. In: *Proceedings of Siggraph '99*, pp. 39–48. ACM Press (1999)
12. Li, Y., Wang, T., Shum, H.Y.: Motion texture: A two-level statistical model for character motion synthesis. In: *Proceedings of ACM SIGGRAPH '02, ACM Transactions on Graphics*, vol. 21(3), pp. 465–471. ACM Press (2002)
13. Metoyer, R.: Building behaviors with examples (2002). *Ph.D. Dissertation, Georgia Institute of Technology*
14. Molina-Tanco, L., Hilton, A.: Realistic synthesis of novel human movements from a database of motion capture examples. In: *Workshop on human motion*, pp. 137–142 (2000)
15. Popović, Z., Witkin, A.: Physically based motion transformation. In: *Proceedings of Siggraph '99*, pp. 11–20. ACM Press (1999)
16. Pullen, K., Bregler, C.: Motion capture assisted animation: Texturing and synthesis. In: *Proceedings of ACM SIGGRAPH '02, ACM Transactions on Graphics*, vol. 21, 3, pp. 501–508. ACM Press (2002)
17. Reitsma, P.S.A., Pollard, N.S.: Perceptual metrics for character animation: sensitivity to errors in ballistic motion. In: *Proceedings of ACM SIGGRAPH '03, ACM Transactions on Graphics*, vol. 22(3), pp. 537–542. ACM Press (2003)
18. Reitsma, P.S.A., Pollard, N.S.: Evaluating motion graphs for character navigation. In: *ACM Siggraph/Eurographics Symposium on Computer Animation*, pp. 89–98 (2004)
19. Ren, L., Patrick, A., Efros, A.A., Hodgins, J.K., Rehg, J.M.: A data-driven approach to quantifying natural human motion. In: *Proceedings of ACM SIGGRAPH '05, ACM Transactions on Graphics*, vol. 24(3), pp. 1090–1097. ACM Press (2005)
20. Safonova, A., Hodgins, J.K.: Analyzing the physical correctness of interpolated human motion. In: *ACM SIGGRAPH /Eurographics Symposium on Computer Animation*, pp. 171–180 (2005)
21. Wang, J., Bodenheimer, B.: An evaluation of a cost metric for selecting transitions between motion segments. In: *ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pp. 232–238 (2003)
22. Wang, J., Bodenheimer, B.: Computing the duration of motion transitions: An empirical approach. In: *ACM Siggraph/Eurographics Symposium on Computer Animation*, pp. 337–346 (2004)
23. Zordan, V.B., Majkowska, A., Chiu, B., Fast, M.: Dynamic response for motion capture animation. In: *Proceedings of ACM SIGGRAPH '05, ACM Transactions on Graphics*, vol. 24(3), pp. 697–701. ACM Press (2005)