Normalized Euclidean Distance Matrices for Human Motion Retargeting

Antonin Bernardin, Ludovic Hoyet, Antonio Mucherino, Douglas S. Gonçalves, Franck Multon

Normalized Euclidean Distance Matrices for Human Motion Retargeting

Antonin Bernardin
University of Limoges, Inria
France
antonin.bernardin@inria.fr

Ludovic Hoyet
Inria
France
ludovic.hoyet@inria.fr

Antonio Mucherino
University of Rennes 1, Inria
France
antonio.mucherino@irisa.fr

Douglas Gonçalves
Federal University of Santa Catarina
Florianòpolis, Brazil
douglas@mtm.ufsc.br

Franck Multon
University Rennes2, Inria
France
franck.multon@irisa.fr

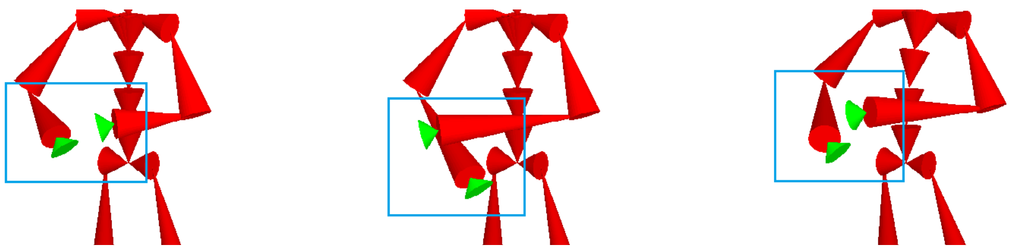

Figure 1: (a) Example pose where the actor's hands come to a close distance. The same pose retargeted on a skeleton with longer forearms by (b) simply transferring joint angles or (c) using our normalized Euclidean distance matrix approach.

ABSTRACT

In character animation, it is often the case that motions created or captured on a specific morphology need to be reused on characters having a different morphology while maintaining specific relationships such as body contacts or spatial relationships between body parts. This process, called *motion retargeting*, requires determining which body part relationships are important in a given animation. This paper presents a novel frame-based approach to motion retargeting which relies on a normalized representation of body joints distances. We propose to abstract postures by computing all the inter-joint distances of each animation frame and store them in Euclidean Distance Matrices (EDMs). They 1) present the benefits of capturing all the subtle relationships between body parts, 2) can be adapted through a normalization process to create a morphology-independent distance-based representation, and 3) can be used to efficiently compute retargeted joint positions best satisfying newly computed distances. We demonstrate that normalized EDMs can be efficiently applied to a different skeletal morphology by using a Distance Geometry Problem (DGP) approach, and present results on a selection of motions and skeletal morphologies. Our approach opens the door to a new formulation of motion retargeting problems, solely based on a normalized distance representation.

MiG'17, November 2017, Barcelona, Spain

ACM ISBN 978-1-4503-5541-4/17/11…$15.00
https://doi.org/10.1145/3136457.3136466

CCS CONCEPTS

• **Computing methodologies → Animation**; **Optimization algorithms**;

KEYWORDS

Character Animation, Motion Retargeting, Distance Geometry

ACM Reference format:
Antonin Bernardin, Ludovic Hoyet, Antonio Mucherino, Douglas Gonçalves, and Franck Multon. 2017. Normalized Euclidean Distance Matrices for Human Motion Retargeting. In *Proceedings of Motion in Games, Barcelona, Spain, November 2017 (MiG'17),* 6 pages.
https://doi.org/10.1145/3136457.3136466

1 INTRODUCTION

Character animation is nowadays largely used in movie and video game industries. Typically, a 3D model and its associated skeletal structure are designed by an artist, then animated either manually or using recorded motion capture data. However, it is often the case that motions created for a specific character, or captured from a given actor, need to be reused on characters with a different morphology, a process called *motion retargeting*. Examples include adapting motions to preserve important relationships between body parts (e.g., a character's hand touching its chin when nodding) or between body parts and the environment (e.g., ensuring that feet remain planted on the ground during locomotion support phases). Motion retargeting is especially important when using motion capture data, where the differences between morphologies of the human actor and of the character to animate raise adaptation issues.

A classical human motion representation consists in sequences of local rotations for each joint of the skeleton. Without specific retargeting techniques, these local rotations are simply transferred from the original to the target skeleton, which often leads to incorrect results (see Figure 1, center). While multiple retargeting techniques have been proposed, most require the prior identification of body part relationships to be maintained, a step that remains difficult to accurately and automatically determine. For this reason, these relationships are still mostly manually specified by animators, and Inverse Kinematics is therefore commonly used in video games to solve simple issues such as adapting the feet of a character according to both ground shape and its morphology.

Instead of relying on this classical representation for motion retargeting, we propose in this paper to explore a novel approach solely based on joint distances. More precisely, we propose to compute all the inter-joint distances for each frame of a motion in order to abstract human postures in a structure called a Euclidean Distance Matrix (EDM) (see Figure 2), which possesses relevant properties for motion retargeting. First, because an EDM contains all the inter-joint distances of a human posture, it accurately captures all the subtle relationships between body parts, which is particularly important to accurately retarget motions. Second, computing the set of joint positions best representing a given EDM can be efficiently performed using approaches from Distance Geometry (DG), a field of research which explores Euclidean distance solving problems [Liberti et al. 2014]. Finally, given that an EDM is specific to the morphology of a given character, we propose a novel method to normalize EDMs to create a morphology-independent distance-based representation. We then demonstrate that these normalized EDMs can be efficiently combined with a new skeletal morphology to retarget motions using an existing Distance Geometry Problem (DGP) approach [Mucherino and Gonçalves 2017].

The remainder of this paper is organized as follows. Section 2 presents related work on character animation and Distance Geometry. Section 3 details our approach. Experimental results are then presented in Section 4, and further discussed in Section 5. Finally, Section 6 provides the concluding remarks.

2 RELATED WORK

Character Animation has been an active field of research for decades, exploring techniques ranging from example-based adaptations [Kovar et al. 2002; McCann and Pollard 2007] to motion simulation [Hodgins et al. 1995; Popović and Witkin 1999]. In this Section, we will be focusing on approaches most related to our work, i.e., related to motion retargeting, and therefore refer the reader to a more general review of the literature on human motion simulation [Guo et al. 2015]. It is also important to mention that some approaches explore retargeting human motions onto non-humanoid characters [Abdul-Massih et al. 2017; Hecker et al. 2008], or to transfer the style of a particular individual on the motions of another [Hsu et al. 2005; Yumer and Mitra 2016], but lie outside the scope of this paper.

Motion retargeting is the process of adapting the motion of a source character to a target character with a different morphology, i.e., usually with the same skeletal structure but different bone lengths. Early motion retargeting approaches relied on space-time constraints in order to preserve desirable qualities of the original motion [Gleicher 1998]. Others proposed to use Inverse Kinematics, combined for instance with prioritized constraints [Le Callennec and Boulic 2004], end-effector importances [Shin et al. 2001] or intermediate/normalized skeletons [Kulpa et al. 2005; Monzani et al. 2000] to handle both motion adaptation and retargeting. Recently, a new type of approach has also been proposed by Ho. et al [2010], which is based on a new structure called Interaction Mesh. This new structure is particularly efficiently to represent implicit spatial relationships between body parts, and has also been extended to account for relationships between human motions and its environment [Al-Asqhar et al. 2013; Ho and Shum 2013].

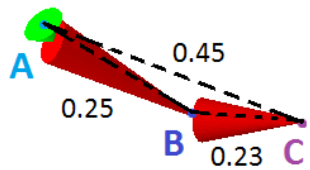

$$\begin{array}{c} \\ A \\ B \\ C \end{array} \begin{array}{c} \begin{array}{ccc} A & B & C \end{array} \\ \begin{bmatrix} 0 & 0.25 & 0.45 \\ 0.25 & 0 & 0.23 \\ 0.45 & 0.23 & 0 \end{bmatrix} \end{array}$$

Figure 2: Example of an Euclidean Distance Matrix (EDM) containing the inter-joint distances of a two-segment skeleton (e.g., arm). In our method, each EDM contains inter-joint distances for the entire skeleton (27 joints in our examples).

However, space-time constraints and Inverse Kinematics approaches depend upon determining or manually specifying which constraints are important for retargeting motions, e.g., that the hands are in contact. Instead of relying on such constraints, which are difficult to automatically identify, our approach is based on representing motions by computing all the inter-joint distances for each frame and storing them in a Euclidean Distance Matrix, which provides relevant information about joint configurations for every frame of the motion. For instance, close interactions like clasping hands will lead to small inter-joint distances. In a certain manner, our approach is closer to the Interaction Mesh presented by Ho. et al [2010], which relies on a Delauney tetrahedralization of the joint positions to compute geometrical spatial relationships. However, their approach prunes the inter-joint relationships that are not required to create the Interaction Mesh, while our approach conserves all the subtle inter-joint relationships in our EDM instead.

Because inter-joint distances depend on the character's morphology, these EDMs need to be adapted before retargeting motions. We therefore propose to normalize and denormalize these EDMs according to the length of each kinematic chain. Unlike [Kulpa et al. 2005; Molla et al. 2017] who also propose to use a similar normalization process, our approach involves the normalization of all the inter-joint distances representing a given posture.

Finally, in order to retarget human motions using these EDMs, we propose to use an optimization technique from the Distance Geometry (DG) community to compute retargeted joint positions. Technically, one of the main problems in DG consists in identifying the coordinates of a set of points in K dimensions that best represent a given distance matrix ($K = 3$ in our case). While such distance-based approaches have been applied to several fields, including sensor network localization [Biswas et al. 2006] or protein structure identification [Lavor et al. 2013], it is to our knowledge the first time it is applied to the problem of retargeting human motions.

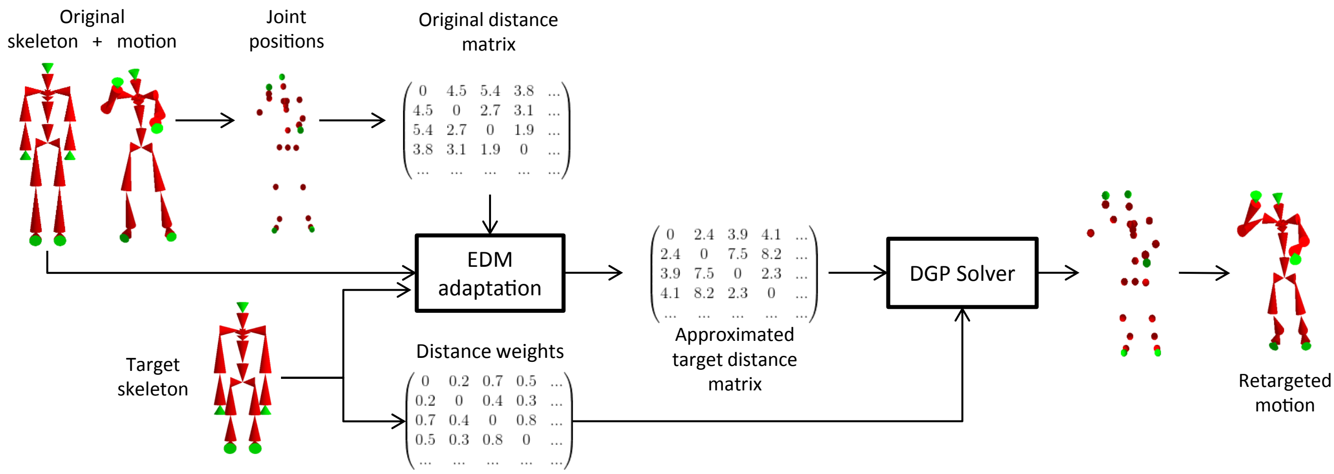

Figure 3: Overview of our method

3 DISTANCE-BASED MOTION RETARGETING

In this Section, we detail our normalized EDM-based representation, together with our retargeting pipeline (summarized in Figure 3).

The cornerstone of our approach is a morphology-independent representation of human motion solely based on distances through a normalized Euclidean Distance Matrix (EDM). All the inter-joint distances are first extracted from the original motion (Section 3.1), therefore representing both skeletal information (joint pairs corresponding to bones) and motion information (other joint pairs). These inter-joint distances are then stored in an EDM for each frame of the original motion. In order to adapt these EDMs to the target skeleton, we then propose a normalization/denormalization procedure based on kinematic chain lengths.

Then, in order to solve for retargeted joint positions best satisfying the adapted EDM, we show how the retargeting of motions can be expressed as a DGP for each frame of the animation (Section 3.2). This optimization is based on a spectral-gradient method coupled with a line-search technique to improve performance in terms of CPU cost. Because all the distances do not have the same importance in retargeting motions, we also assign a weight to each element of the EDM, which is automatically computed from the target skeleton morphology. For example, it is crucial to respect bone lengths, but it may not be as important to respect exactly the distance between a foot and a hand when they are quite far apart.

Finally, joint positions computed by our DGP solver are used to reconstruct the local rotation of the skeletal joints, which are then displayed for visualisation (Section 3.3).

3.1 Normalized Euclidean Distance Matrix

For each frame of the original motion, we compute all the inter-joint Euclidean distances, and store them in our EDM $D^{original}$ (see example in Figure 2), of size $N \times N$ (where N is the number of joints in the skeleton).

Assuming that X is our set of joint positions, each distance $d_{uv}^{original} \in D^{original}$ between joint u and v is defined as:

$$d_{uv} = \|x_u - x_v\|, \; x_u, x_v \in X \tag{1}$$

Given only the distances specified in $D^{original}$, it is possible to reconstruct a posture X' satisfying all distance constraints, an inverse problem known in the scientific literature as the DGP, which is at the core of our approach and detailed in the Section 3.2.

By definition, $D^{original}$ contains both information about the structure of the skeleton and information about the pose of the character (motion), i.e., $d_{uv}^{original}$ corresponds to the bone length if joints u and v are connected by a bone and otherwise provides information about the posture of the character, e.g., how close are the hands from one other or from the head.

In order to retarget motions on a skeleton differing from the original skeleton, it is therefore necessary to adapt distances in $D^{original}$ to reflect the new skeletal structure. Because posture-based distances (i.e., non-bone inter-joint distances) are not completely independent from the skeletal structure, only adapting bone distances would introduce large artifacts and incompatibilities. To tackle this problem, we propose to normalize $D^{original}$ using the properties of the original skeleton, then to denormalize it using the properties of the target skeleton. In particular, we propose to normalize and denormalize distances based on the actual length τ_{uv} of the kinematic chain S_{uv} between joints u and v in the original and target skeletons, in a way similar to [Molla et al. 2017]:

$$\tau_{uv} = \sum_{i=1}^{|S_{uv}|} s_i \tag{2}$$

where $|S_{uv}|$ is the number of segments in the kinematic chain connecting joints u and v, and s_i is the length of the i^{th} segment of the kinematic chain. A normalized distance matrix $D^{normalized}$ is then computed according to:

$$\forall \{u, v\} \in X, \quad d_{uv}^{normalized} = \frac{d_{uv}^{original}}{\tau_{uv}^{original}} \tag{3}$$

In particular, $d_{uv}^{normalized} = 0$ represents joint-to-joint contacts, while $d_{uv}^{normalized} = 1$ represents a fully extended kinematic chain. While the concept is similar to [Hecker et al. 2008; Kulpa et al. 2005; Molla et al. 2017], our approach normalizes all the inter-joint distances representing a given posture.

Finally, we compute a denormalized distance matrix D^{target} according to the target skeleton morphology, using:

$$\forall \{u,v\} \in X, \quad d_{uv}^{target} = \tau_{uv}^{target} \cdot d_{uv}^{normalized} \tag{4}$$

This normalisation process presents the advantage of correctly reconstructing the bone lengths of the target skeleton, while simultaneously adapting the other inter-joint distances based on the morphology of the target skeleton.

3.2 Distance Geometry Problem Approach

Once we have approximated the distance matrix for our target skeleton, our goal is to reconstruct the target joint positions satisfying these inter-joint Euclidean distances. This type of problem can be efficiently solved using approaches from the field of Distance Geometry. In short, DGP solvers search for the set of coordinates in K dimensions (in our case, $K = 3$ for joint positions in 3D space) that best satisfy a given distance matrix.

In our DGP formulation, we consider an objective function $\sigma(X)$ that measures the violation of the distance constraints as follows:

$$\sigma(X) = \sum \omega_{uv} \cdot (\|x_u - x_v\| - d_{uv})^2 \tag{5}$$

where ω_{uv} is a weight associated with each distance. Intuitively, in this iterative optimization process, distances between joints that are close in the skeletal structure (i.e corresponding to kinematic chains with fewer bones) must be strongly satisfied. Conversely, other distances, such as between the hand and a foot, are of less importance. For this reason, a nonnegative weight ω_{uv}, representing the "importance" of the distance with respects to the others, is associated to each distance d_{uv}. Importance is expressed as a relation on the number of joints $|P_{uv}|$ encountered along the kinematic chain from joint u to v, such as:

$$\omega_{uv} = \frac{|P_{\max}| - |P_{uv}| + 2}{|P_{\max}|}, \ \omega_{uv} \in [0,1] \tag{6}$$

where $|P_{\max}|$ is the number of joints of the longest kinematic chain. In particular, ω_{uv} is defined so that it is maximal in the case of distances representing bones ($|P_{uv}| = 2$), in order to avoid distortion of these highly-constrained distances. In the contrary, ω_{uv} decreases with the length of the kinematic chain.

To conduct the optimization, we use an existing spectral gradient algorithm [Glunt et al. 1993], which was previously used to solve DG problems. In particular, it was demonstrated to provide faster local convergence than classical gradient methods. It is further combined with a non-monotone line-search strategy [Zhang and Hager 2004] in order to improve convergence (ensuring convergence in linear computational time). In the following, we will refer to our implementation of the non-monotone spectral gradient as our DGP solver. More details about this implementation are given in our Supplemental Material.

Finally, because our DGP solver consists in an optimization process, results can depend on the initialization of our retargeted joint positions. However, human postures do not vary dramatically from a keyframe to the next, therefore we perform a warm start of our frame-based optimization process and initialize it using the retargeted joint positions computed for the previous frame. Because such a previous frame is not available for the first retargeted frame, a rough initial guess of the target joint positions is provided using joint positions of the first frame of the original motion.

3.3 From joint positions to joint transforms

The output of the DGP solver is a point cloud corresponding to a set of retargeted joint positions for each frame of the animation. For visualisation and animation purposes, we therefore recompute the global transformation of each body segment using standard animation techniques. In particular, we deduce from the positions of successive joints in the kinematic chain the direction of each bone, and compute normals on a bone-specific basis. For instance, the normal vector of the shoulder and elbow joints are defined by the normal of the half plane containing the shoulder, elbow and wrist joints, similarly to [Kulpa et al. 2005]. While this approach is effective for most bones, it is however currently not possible to accurately determine the global rotation of the root joint, as well as of the end-effector joints (hands, toes and head), a point further discussed in Section 5.

4 EXPERIMENTAL RESULTS

Our method was implemented in a C++/OpenGL framework, and tested on a selection of motion captured examples acquired with a 24-camera vicon system (27 joints per skeleton). Experiments were run on an Intel Xeon @3.0GHz with 6GB RAM, running Windows 7.

In the following experiments, we compare our distance-based retargeting approach with the traditional local rotation transfer method. Results are presented in Figure 4 on selected poses, and retargeted motions are presented in the companion video. In particular, results presented in Figure 4 focus on a number of selected poses retargeted using our approach on a target skeleton with 35%-longer forearms (right). For visual comparison, we also include the captured poses on the original skeleton (left), and a "traditional" local rotation transfer on a target skeleton with 35%-longer forearms. Our results show that motions involving close interactions between body parts are accurately retargeted using our approach, in comparison with a traditional local rotation transfer. For instance, contacts between both hands are preserved in Example C, as well as shoulder-hand relationships in Examples A and B. Moreover, traditional local rotation transfer can lead to closer body-part relationship in some cases (e.g., Examples D, E and F), which are also better preserved using our distance-based approach. These examples show that our approach is able to preserve close interactions, while also avoiding incorrect close contacts (or even body intersections) that could result from a naive method.

In terms of computation time, our frame-base retargeting approach takes on average 5±1ms to retarget a pose (max: 10ms, min: 0.5ms). This average was computed on two dance motions captured at 100Hz (total number of frames: 8568 frames) which was retargeted on skeletons with longer forearms or upperarms (+30%, +15%, -15%, -30%), and was not influenced by the percentage of skeletal differences. Our fast computation times are due to both using a spectral gradient method, which was demonstrated to have

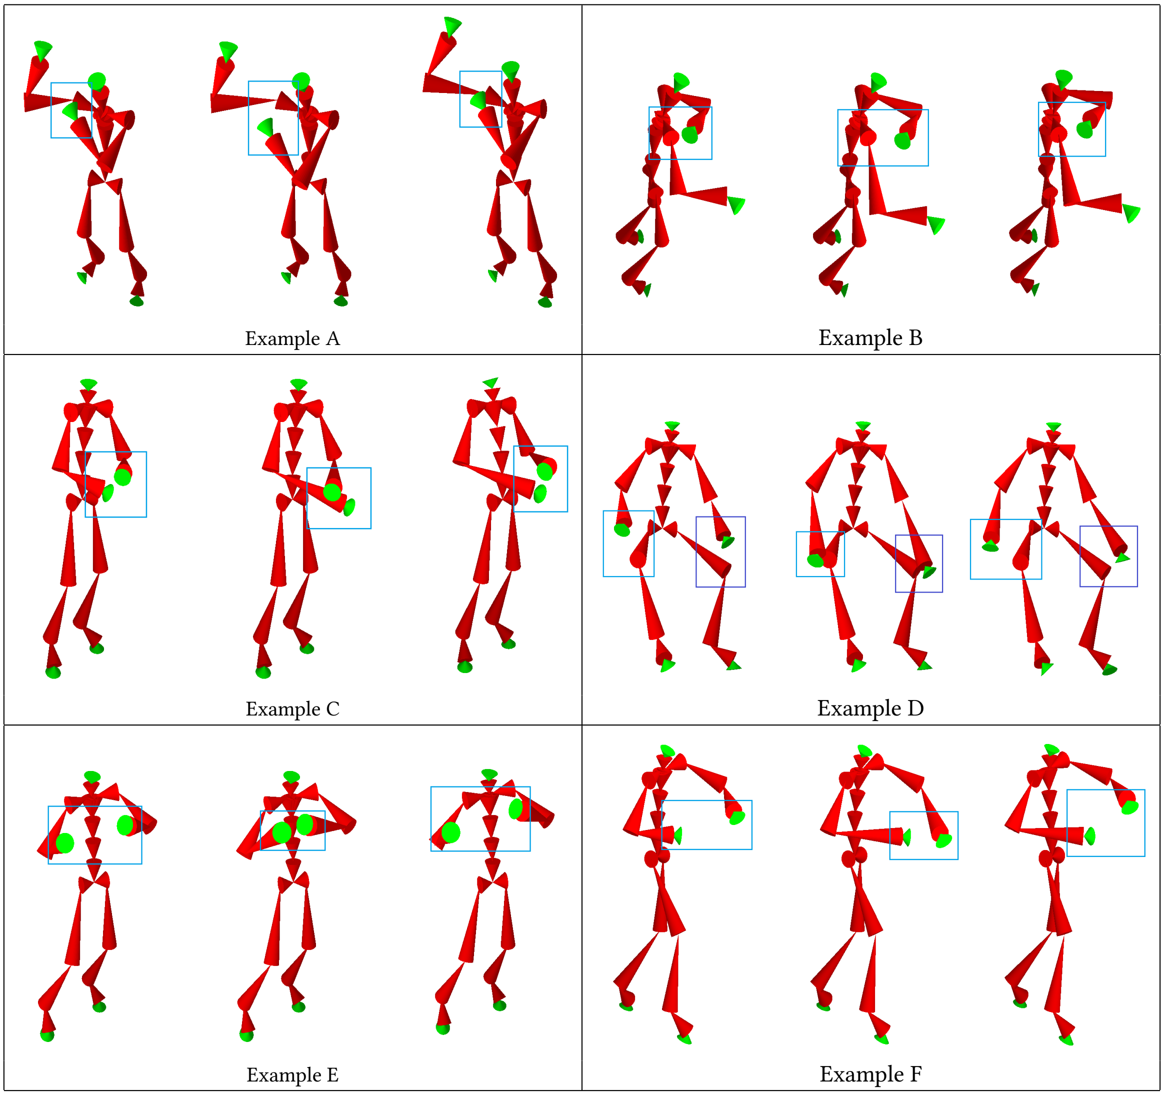

Figure 4: Results on a number of example poses from a dancing captured motion using: (left) the captured pose on the original skeleton, (center) local rotations transfered on a target skeleton with 35%-longer forearms, and (right) a pose retargeted using our approach on a target skeleton with 35%-longer forearms.

large benefits in terms of CPU time, and to the warm start using the retargeted joint positions from the previous frame, which provides a close starting point for the optimization. However, it is important to mention that computation time is significantly higher from the first frame (50ms to 150ms), because the first initialization more significantly differs from the target skeleton morphology. From the same reason, while our optimization process is limited to 50 iterations per frame, this limit is set to 1000 iterations for the first frame to retarget in order to ensure a correct first retargeted posture.

5 DISCUSSION

Our experimental results demonstrate that our distance-based approach efficiently preserves important body-part relationships while retargeting human motions, without manually specifying the relationships. In particular, it both preserves close interactions from the original motion, while also avoiding incorrect contacts or body part intersections that could result from a naive method. While we demonstrate our results only for skeletal modifications of the upper limbs, because these body parts most commonly present important relationships with the rest of the body, our method is general and can be applied to any modifications of the skeletal morphology. Moreover, one of the strengths of our approach is that it is solely based on inter-joint distances, which provides an abstraction of human postures which is easy to implement, to compute and to understand. Such an abstraction would also be valuable for general motion editing in order to specify motion constraints in an more intuitive manner. For instance, a hand touching another body part would simply be obtained by specifying a small distance between the corresponding joints, which is intuitive to specify.

Because of the complexity of human motion, and the formulation of the DGP, our method however suffers from some limitations. In particular, a DGP solver computes a set of retargeted joint positions which is invariant in translation and rotation, i.e., which does not provide any information about the global position and orientation of our character in the 3D space. However, the root global transformation of the source motion could provide a good initial guess, and be adapted using existing animation techniques. Similarly, our approach cannot currently compute end-effector orientations, as they are represented by only two positions in the skeletons we used for experiments, which could be solved by using the original end effector orientations, or by using additional reference points on our end-effector segments (which could easily be included in our approach given its generality).

Our distance-based representation proved to be efficient in capturing relationships between body parts. However, computing retargeted joint positions relies on defining which distances are more important. Our approach favors distances for kinematic chains with fewer joints, because of inaccuracies introduced by the normalization/denormalisation of longer kinematic chains. While we found that this solution produced the best results, it still fails to compute correct retargeted joint positions in some cases. In particular, inversions can occur for highly symmetrical postures, because two different sets of joint positions will lead to a similar distance matrix. Examples of such problems are presented in the supplementary video. Possible solutions involve accounting for joint limits in the optimization process, or defining metrics measuring when a retargeted posture differs excessively from the original posture in order to readapt the optimization process. Also, we found that extremely large skeletal differences often produced non-natural configurations, e.g., with a spine extremely bended, in particular when differences were applied on several body parts simultaneously. While such large differences are always difficult to take into account, we think that smaller successive retargeting steps could be considered in order to produce natural retargeted motions.

6 CONCLUSION AND FUTURE WORK

To conclude, we have presented in this paper a novel approach based on joint distances to retarget human motions. More precisely, it relies on abstracting human postures by computing all of the inter-joint distances, which are stored in an EDM for each frame. Such a distance matrix is simple to compute, while simultaneously presenting relevant properties for motion retargeting. For instance, it captures all the subtle relationships between body parts, and is also used to efficiently compute retargeted joint positions using approaches from Distance Geometry. We also proposed a manner of normalizing these EDMs to account for differences in skeletal morphologies, which is crucial in the retargeting process.

Our approach shows promising results, opening the door to a new manner of tackling motion retargeting problems while also raising a number of challenges. For instance, solely basing motion retargeting on a distance representation raises new challenges to represent the relationships between a human posture and its environment, or to further detail the subtle surface to surface relationships between body parts. In the future, we are particularly interested in including such additional information into the distance matrix in order to further generalize our approach. The quality of our retargeted motions should also be evaluated in order to validate our results. In particular, we are now interested in comparing our results with existing approaches such as [Ho et al. 2010; Kulpa et al. 2005] using objective and subjective evaluations. This approach could also be relevant to the design of new motion capture systems using sensors providing only distances between body parts, especially given that it runs in real time at approximately 200fps.

REFERENCES

M. Abdul-Massih, I. Yoo, and B. Benes. 2017. Motion Style Retargeting to Characters With Different Morphologies. *Computer Graphics Forum* 36, 6 (2017).

R.A. Al-Asqhar, T. Komura, and M.G. Choi. 2013. Relationship Descriptors for Interactive Motion Adaptation. In *Proceedings of the 12th ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA '13)*. 45–53.

P. Biswas, T.-C. Liang, K.-C. Toh, Y. Ye, and T.-C. Wang. 2006. Semidefinite programming approaches for sensor network localization with noisy distance measurements. *IEEE Trans. on automation science and engineering* 3, 4 (2006).

M. Gleicher. 1998. Retargetting Motion to New Characters. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '98)*. 33–42.

W. Glunt, T.L. Hayden, and M. Raydan. 1993. Molecular conformations from distance matrices. *Journal of Computational Chemistry* 14, 1 (1993), 114–120.

S. Guo, R. Southern, J. Chang, D. Greer, and J.J. Zhang. 2015. Adaptive motion synthesis for virtual characters: a survey. *The Visual Computer* 31, 5 (2015), 497–512.

C. Hecker, B. Raabe, R.W. Enslow, J. DeWeese, J. Maynard, and K. van Prooijen. 2008. Real-time Motion Retargeting to Highly Varied User-created Morphologies. *ACM Trans. Graph.* 27, 3 (2008).

E.S.L. Ho, T. Komura, and C.-L. Tai. 2010. Spatial Relationship Preserving Character Motion Adaptation. *ACM Trans. Graph.* 29, 4 (2010).

E.S.L. Ho and H.P.H. Shum. 2013. Motion adaptation for humanoid robots in constrained environments. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on*. 3813–3818.

J.K. Hodgins, W.L. Wooten, D.C. Brogan, and J.F. O'Brien. 1995. Animating Human Athletics. In *Proceedings of the 22Nd Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '95)*. 71–78.

E. Hsu, K. Pulli, and J. Popović. 2005. Style Translation for Human Motion. *ACM Trans. Graph.* 24, 3 (2005), 1082–1089.

L. Kovar, M. Gleicher, and F. Pighin. 2002. Motion Graphs. *ACM Trans. Graph.* 21, 3 (2002).

R. Kulpa, F. Multon, and B. Arnaldi. 2005. Morphology-independent representation of motions for interactive human-like animation. *Computer Graphics Forum* (2005).

C. Lavor, L. Liberti, and A. Mucherino. 2013. The interval Branch-and-Prune algorithm for the discretizable molecular distance geometry problem with inexact distances. *Journal of Global Optimization* (2013), 1–17.

B. Le Callennec and R. Boulic. 2004. Interactive Motion Deformation with Prioritized Constraints. In *Proceedings of the 2004 ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA '04)*. 163–171.

L. Liberti, C. Lavor, N. Maculan, and A. Mucherino. 2014. Euclidean distance geometry and applications. *Siam Review* 56, 1 (2014), 3–69.

J. McCann and N. Pollard. 2007. Responsive Characters from Motion Fragments. *ACM Trans. Graph.* 26, 3 (2007).

E. Molla, H.G. Debarba, and R. Boulic. 2017. Egocentric mapping of body surface constraints. *IEEE Transactions on Visualization and Computer Graphics* (2017).

J.-S. Monzani, P. Baerlocher, R. Boulic, and D. Thalmann. 2000. Using an intermediate skeleton and inverse kinematics for motion retargeting. In *Computer Graphics Forum*, Vol. 19. Wiley Online Library, 11–19.

A. Mucherino and D.S. Gonçalves. November 2017. An Approach to Dynamical Geometry. In *Proceedings of Geometric Science of Information (GSI17)*. To appear in Lecture Notes in Computer Science, F. Nielsen, F. Barbaresco (Eds.), 8 pages.

Z. Popović and A. Witkin. 1999. Physically Based Motion Transformation. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '99)*. 11–20.

H.J. Shin, J. Lee, S.Y. Shin, and M. Gleicher. 2001. Computer Puppetry: An Importance-based Approach. *ACM Trans. Graph.* 20, 2 (2001).

M.E. Yumer and N.J. Mitra. 2016. Spectral Style Transfer for Human Motion Between Independent Actions. *ACM Trans. Graph.* 35, 4 (2016).

H. Zhang and W.W. Hager. 2004. A nonmonotone line search technique and its application to unconstrained optimization. *SIAM journal on Optimization* 14, 4 (2004), 1043–1056.