Cage-Based Tracking for Performance Animation

Yann Savoye and Jean-Sébastien Franco

INRIA Bordeaux University, France

Abstract. Full body performance capture is a promising emerging technology that has been intensively studied in Computer Graphics and Computer Vision over the last decade. Highly-detailed performance animations are easier to obtain using existing multiple views platforms, markerless capture and 3D laser scanner. In this paper, we investigate the feasibility of extracting optimal reduced animation parameters without requiring an underlying rigid kinematic structure. This paper explores the potential of introducing harmonic cage-based linear estimation and deformation as post-process of current performance capture techniques used in 3D time-varying scene capture technology. We propose the first algorithm for performing cage-based tracking across time for vision and virtual reality applications. The main advantages of our novel approach are its linear single pass estimation of the desired surface, easy-to-reuse output cage sequences and reduction in storage size of animations. Our results show that estimated parameters allow a sufficient silhouette-consistent generation of the enclosed mesh under sparse frame-to-frame animation constraints and large deformation.

1 Introduction

Modeling dynamic 3D scene across time from multiple calibrated views is a challenging problem that has gained full attention of the Computer Vision and Computer Graphics communities in recent years. Nevertheless, the relentless increase in demand of 3DTV industry has inspired researchers to exhibit new approaches able to produce reusable contents. Current pipelines try to achieve deformable mesh tracking using one or many linear or non-linear numerical optimizations. Video-based approaches are proven to be more convenient to acquire human performances. The field of targeted applications is very large, ranging from content generation for the 3D entertainment and motion picture industries, video game to sport analysis. Even if the human motion can be abstracted by the kinematic of rigid parts, the observed surface to track is purely non rigid because of small detail variations induced by clothes, hair motion and certain degrees of flexibility by virtue of natural properties.

A major challenge is to exhibit an efficient framework to achieve plausible boneless tracking that produces pleasing deformations, preserves the global appearance of the surface and offers flexible reusable outputs for animation. For this reason, we propose a new fully linear framework to track mesh models with full correspondence from multiple calibrated views.

In this paper, we introduce a new cage-based mesh parametrization for tracked surfaces. The surface model, initially acquired by a laser or initial dense reconstruction method, is smoothly and volumetrically embedded in a coarse but topologically

identical mesh, called *the cage*, whose vertices serve as control points. The resulting reduction in control parameters and space embedding yields an interesting new trade off to tackle the mesh tracking problem. More precisely, we focus on the estimation of desired enclosed models in a linear manner preserving the smoothness of the cage under sparse linear subspace constraints. Such constraints are defined over the enclosed mesh surface itself. We take advantage of optimal reduced parameters offered by the given coarse cage surrounding the surface. In order to avoid artifacts induced by the large number of degrees of freedom, the cage layer is enhanced with laplacian regularization. In other words, we embed to-be-tracked models in an adapted bounding polytope cage using generalized barycentric coordinates having local smooth properties, drawing inspiration from Computer Graphics animation-oriented model parametrizations.

The rest of the paper is organized in the following manner. Relevant works concerning markerless full body performance capture and non-rigid surface deformation are briefly reviewing and discussing in section 2. The problem statement and contributions are presented in section 3. Background and notations are introduced in section 4. We give an overview of our system in section 5. The core of the proposed method is detailed in section 6. Section 7 presents some results and evaluations of our novel technique in the context of multiple views performance animation. We show the effectiveness of our method with several examples. This paper is concluded in section 8 and limitations are discussed and 9.

2 Related Works

In this section, we will briefly describe relevant recent works in the field of 3D Video and Vision-based Graphics. We mainly focus on previous approaches addressing the problem of markerless full body performance capture from multiple views and interactive mesh deformation.

Performance Capture and 3D video Markerless performance capture can be defined as the process of generating spatio-temporally coherent and connectivity preserving geometry. Recent years have seen strong interest for video-based performance capture dealing with video driven laser-scanned template or template-free deformation as well as skeleton-based or boneless tracking. One of the first pioneering work in surface capture for performance-based animation is proposed in [1] with a fully automated system to capture shape and appearance based on visual hull extraction, feature matching and dense reconstruction.

De Aguiar *et al.* have proposed in [2] a markerless approach to capture human performances from multi-view video that produces a novel dense and feature-rich output. This approach is enhanced in [3] by a novel optimization scheme for skeleton-based pose estimation and surface local estimation. Moreover Vlasic *et al.* present in [4] a framework for articulated mesh animation from multi-view silhouettes. They also proposed a new method in [5] to dynamically perform shape capture using multi-view photometric stereo normal maps. Another approach presented in [6] introduces a tracking algorithm to realize markerless motion capture of skinned models in a four camera setup using optical flow and skeletal subspace deformation into a nonlinear minimization framework.

3D Video processing becomes a promising visual media by enabling interactive viewpoint control. 3D video technologies are able to capture high-fidelity full 3D shape, motion and texture to offer free-and-rotate special effects. The full 3D video pipeline is

presented in [7]. In the context of human motion synthesis from 3D video [8] proposed a method where surface motion graph is constructed to concatenate repeated motion.

Non-Rigid Deformation High quality shape deformation has gained a lot of attention in recent years. For instance, non rigid mesh evolution can be performed using intrinsic surface deformation or extrinsic space deformation techniques.

A lot of effort has been done on surface-based deformation. There are several types of approaches exploiting a differential descriptor of the edited surface in terms of laplacian and differential coordinates for mesh editing [9, 10]. Differential information as local intrinsic feature descriptors has been massively used for mesh editing in various frameworks over the decade [11–13]. Observing the local behaviour of the surface has been proposed recently in [14], where "as-rigid-as-possible" surface modeling is performed.

There has been a great deal of work done in the past on developing techniques for deforming a mesh with generalized barycentric coordinates. Inspired from the pioneering work presented in [15], cage-based methods are ideal for deforming a surface coherently by improving space deformation technique. The cage parametrization allows model vertices to be expressed as a linear combination of cage vertices for generating realistic deformation. This family of methods has important properties: quasi-conformal mappings, shape preservation and smoothness. To animate the model, cage vertices are displaced and the vertices of the model move accordingly through a linear weighted combination of cage geometry parameters. An approach to generalize mean value coordinates is introduced in [16]. The problem of designing and controlling volume deformations used to articulate characters are treated in [17], where the introduction of harmonic coordinates significantly improves the deformation stability thanks to a volumetric heat diffusion process respecting the connectivity of mesh volume.

A non linear coordinates system proposed in [18] called *Green Coordinates* leads space deformation with shape preserving properties. However such approaches require to obtain automatically a fairly coarse control mesh approximating enough a given surface [19, 20]. Furthermore, there has been a great deal of work made feasible with respect to the work presented in [21–23], where the authors use an analogy to the traditional use of skeleton-based inverse kinematics.

3 Problem statement and Contributions

Even if lots of methods reconstruct continuous surfaces from several viewpoints, we notice a lack of global, flexible and reusable parametrisation. However, finding suitable non-rigid performance animation model with reduced control parameters is a key problem in video-based mesh tracking. Unlike methods based on an underlying rigid skeleton, we aim to estimate the subspace deformation corresponding to the time-varying non-rigid surface evolution. Even if reconstruction is not part of our contribution, we deal with the peculiarities of surfaces that have been reconstructed from real video footage. Thus, we propose a new approach estimating an optimal set of cage vertices position allowing the template to fit the silhouettes, in a single-pass linear method ensuring cage smoothness under sparse subspace constraints.

Our method estimates cage parameters using animation sequence generated from calibrated multi-view image sequences and cage-based deformation. We exhibit an external parametrization involving a reduced number of parameters comparing to the number of enclosed mesh vertices. However, the key contribution is to solve a sparse

linear system to estimate the best cage parameters reproducing the desired deformation of the enclosed model, given sparse constraints expressed on the enclosed model itself. This paper shows that cage parametrization can be used for video-based acquisition. To the best of our knowledge, this is the first approach using cage-based animation for performance capture.

4 Background and Notations

Multiple View Setup We assume that a scene, composed of a single non-rigid object (Fig. 1), observed by a camera network composed of a set of *i* views where each of them corresponds to a fixed pinhole camera model. In addition, we assume that the static background can be learned in advance. The 4×3 projection matrix of the *i*th camera is denoted by \mathbf{P}_i . The mapping from a 3D point $\mathbf{p} = (x, y, z)$ of the surface in the world coordinate space to its image coordinates (u_i, v_i) by projection in camera *i* is given by:

$$u_{i}(\mathbf{p}) = \frac{\mathbf{P}_{i}^{1}(x, y, z, 1)^{T}}{\mathbf{P}_{i}^{3}(x, y, z, 1)^{T}} \quad ; \quad v_{i}(\mathbf{p}) = \frac{\mathbf{P}_{i}^{2}(x, y, z, 1)^{T}}{\mathbf{P}_{i}^{3}(x, y, z, 1)^{T}}$$
(1)

where \mathbf{P}_{i}^{j} is the j^{th} rows of P_{i} associated to the i^{th} camera. The projection matrix is obtained easily from the intrinsic and extrinsic camera. This projective formulation is used for qualitative evaluation.

Cage-Based Animation In the rest of the paper, we use the following terminology. The coarse morphable bounding mesh C and the enclosed mesh \mathcal{M} are respectively called *the cage* and *the model*. We assume that both entities are two-manifold triangular meshes fully-connected, with fixed toplogy, even if the model can be multi-component. The set of n cage vertices is denoted $\mathcal{V}_C = \{c_1, \cdots, c_n\}$ where c_k is the position coordinates of the k^{th} cage vertex, and the set of m model vertices with $\mathcal{V}_{\mathcal{M}} = \{v_1, \cdots, v_m\}$ where v_i is the location of the i^{th} model vertex.



Fig. 1: Cage-based multiple-view setup.

5 Pipeline Overview

Before introducing technical details, we briefly describe the procedure of our method in an overview. Our method retrieves the video-based space warping of the observed surface across time. We assume that our framework exploits already reconstructed mesh sequences with full correspondence for the moment. We cast the problem as a minimization problem for cage recovery. The main challenge is to deal with the high number of degrees of freedom provided by the coarse cage.

As input of our pipeline (Fig. 2), we give a collection of images captured from the observed scene from calibrated views, and the reconstructed mesh sequence. As output, the system produces a sequence of cage's geometry parameters for each frame. In our system, we employ laplacian operator on the cage and harmonic coordinates to perform a space deformation surfacically regularized that allows us to obtain a coherent cage estimation. We estimate a sequence of cage parameters expressing the mesh at each



Fig. 2: Cage-based tracking: pipeline overview.

animation frame using cage-based inverse kinematics process. The optimization process retrieves the pose in a least-squares sense from sparse motion constraints expressed directly over the enclosed surface and transfered into the subspace domain using its indirection. This new technique produces suitable outputs for animation compression and re-using.

6 Linear Least Squares Cage Fitting

This section presents the laplacian-based regularization applied exclusively on the cage structure, instead of current methods which regularize the full mesh itself. We also introduce the association of harmonic subspace deformation with cage-based dual laplacian to perform space tracking using a cage.

6.1 Harmonic Subspace Deformation

A cage is a coarse closed bounding polyhedral volume, not a lattice. This flexible low vertex-count polytope, topologically similar to the enclosed object, can efficiently control its deformation and produce realistic looking deformation. Model vertices are expressed as a linear combination of cage vertices. The weights are given by a set of generalized barycentric coordinates stored in a $m \times n$ deformation weight matrix denoted by \mathcal{H} . We also denote by $h_k(i)$ the normalized blend weights representing the deforming influence of the k^{th} cage vertex on the i^{th} model vertex. In another words, it is possible to deform an arbitrary point on the enclosed mesh expressed as a linear combination of the coarse mesh vertex position via a constant weight deformation. The cage-based forward kinematic function is expressed as follows:

$$v'_{i} = \sum_{k=1}^{n} h_{k} (i) \cdot c'_{k}$$
(2)

where v'_i is the deformed cartesian coordinates according to a vector of cage geometry $\{c'_1, \dots, c'_n\}$. In order to produce as-local-as possible topological changes on the enclosed surface, the model is rigged to the cage using harmonic coordinates. The harmonic rigging is the pre-computed solution of Laplace's equation with Dirichlet boundary condition obtained by a volumetric heat diffusion in the cage interior volume. The resulting matrix corresponds to the matrix \mathcal{H} .



Fig. 3: Harmonic Rigging process: (from left to right hand side) laser scanned template, cage-based model embedding, cage voxelization, voxel tagging and harmonic weights computation.

The subspace domain is the volume enclosed in the cage interior defined by control points. For each control point c_k , we define a harmonic function $h_k(\cdot)$ inside the cage by enforcing the interpolation property $h_k(c_j) = 1$, if k = j, and 0 if not, and asking that h_k be linear on the edges of the cage.

6.2 Laplacian-based Subspace

Generally laplacian mesh editing techniques are interactive, but not real-time because the per-vertex laplacian is defined for the whole mesh and thus computationally expensive. A better idea to ensure the smoothness of the laser scan surface independently of its high-detail resolution is to define a Dual Laplacian operator. We propose to define the Laplacian operator on the cage instead of the model to improve the computation process and to keep model detail properties good enough.

Let's denote by $\mathcal{L}_{\mathcal{C}}(\cdot)$ the Dual Laplacian operator defined at each cage vertex domain by as the weighted sum of the difference vectors between the vertex and its adjacent neighbors. Given the fact that the cage is highly irregular, we prefer to use the cotangent weighing system in the computation of laplacian matrix. We also denote the differential coordinates of the cage by $\hat{\delta}$. Encoding each control vertex relatively to its neighborhood preserves the local geometry using differential coordinates.

6.3 Linear Subspace Constraints

The subspace topology is preserved across time because the default cage is seen as a connectivity mesh only and feature constraints are seen as external deformation. This surface-and-space based deformation technique preserves the mesh spatial coherence.

The geometry of the cage can be reconstructed efficiently from its indirect harmonic coordinates and relative coordinates by solving a system of sparse linear equations. We cast the problem of deformation as a least-squares laplacian cage reconstruction process using a consistent minimization approach of an objective function. The cage parameters recover the sparse pose features by minimizing an objective function in a least-squares sense in order to fit a continuous volume. Then, the geometry of the desired model is simply obtained by generating its vertex position according to the expressed cage parameters obtained on the concept of least-squares cage.

Given the differential coordinates, laplacian operator of the default cage, the harmonic weights $h_k(i)$ according to the cage and the model at the default pose, and several 3D sparse surface constraints, the absolute coordinates of the model geometry can be reconstructed by estimating the absolute coordinates of the cage geometry. The combination of the differential coordinates and harmonic coordinates allows the reconstruction of the surface by solving a linear system. We can formulate overall energy to lead an overdetermined linear system to extract the cage parameters.

This least-squares minimization problem can be expressed exclusively in term of cage geometry (for frame-to-frame animation) as demonstrated in the following formula:

$$\min_{c'_{k}} \left(\sum_{k=1}^{n} \left\| \mathcal{L}_{C} \left(c'_{k} \right) - \hat{\delta'}_{k} \right\|_{2}^{2} + \sum_{v_{i} \in \mathcal{S}} \left\| v'_{i} - \sum_{k=1}^{n} c'_{k} \cdot h_{k} \left(i \right) \right\|_{2}^{2} \right)$$
(3)

Note that the first term of the energy preserves the cage smoothness and ensures a pleasant deformation under sparse constraints. The space-based distortion energy is measured by the laplacian on the cage. The total local distortion measure for a deformation process is given by a quadratic energy term.

The second term of the energy enforces the position of vertices to fit the desired model defined by positional constraints. To our best knowledge, the simple global optimization component of our framework with such formulated constraints to minimize does not already exist in the literature. Overall energy performed by our technique reproduces harmonic space deformation recovery under indirected dual laplacian mesh editing. After the cage retrieval process, the geometry of the desired enclosed model is reconstructed in linear combination function of cage geometry parameters related to the new estimation, preserving the fixed connectivity.

In order to deform the bounding cage, positional constraints are defined on the model using anchor points. We denote by v'_i the cartesian coordinates position of the target point at t + 1 associated to v_i to create a positional constraint. S is the subset composed of a irregularly distributed collection of positional constraints (for each selected v_i that form S) over the enclosed surface. The second term of the objective function measures how much the cage enforces sparse positional constraints. The transfer of surfacic constraints into the subspace domain exploiting the cage indirection is expressed by this energy term. In other words, the last formulation enables to express surface constraints directly in terms of cage parameters linearly using an inverse quasi-conformal harmonic mapping.

Given the control mesh for the previous frame in a deformation mesh sequence, we need to exploit frame-to-frame deformation of the finer mesh to automatically constructed an altered control mesh for every frames in the sequence. As shown on the results, the cage retrieval process only requires a small number of corresponding input vertices and their displacement to form sparse linear subspace constraints to be able to estimate a cage expressing a surface fitting to the silhouette good enough.



Fig. 4: Qualitative evaluation of the cage-based reconstructed surface: (a) estimated cage, (b) textured enclosed surface, (c) cage-based model reprojected, (d) silhouette overlap error.



Fig. 5: Influence of the number of positional constraints on the quality of enclosed mesh expressed by the estimate cage (the number of positional constraints are expressed in percentage of enclosed mesh vertices) (dataset courtesy of [4]).

7 Results and Evaluation

This section describes our experiments using this system. Our framework, implemented with OpenGL and C/C++, proposes a robust mechanism to extract a cage for various potential applications. The entire process takes less than two seconds per frame without any code optimization, and uses solvers running on CPU. The algorithm performance is validated by both qualitative and quantitative evaluations. We show that the cage reproduces the 3D motion of life-like character accurately without requiring a dense mapping.

The performance of our system was tested on multi-view video sequences that were recorded with eight cameras at a resolution of 1600x1200 pixels. The template is composed of 10002 vertices and the cage is composed of 141 vertices (80% of parameter reduction of the enclosed model). To validate our method, some experimental results are shown on real datasets (Fig. 4). Qualitative evaluations are based upon visual comparisons (silhouette consistency) of each reconstructed frame with the original one and various percentage of vertex constraints randomly selected (Fig. 5). We also provide rendering feedback to allow qualitative evaluation on two different sequences with a total of 348 frames (Fig. 7).

We run our cage-based tracking method to measure how much the estimated cagebased deformation of the template can fit the observed silhouettes without applying



Fig. 6: Quantitative evaluation on MIT Crane dataset.

an additional silhouette regularization on the enclosed surface. For our evaluation as shown on Fig. 6, we measure the fidelity of our output with several error metrics such as edge length variation, silhouette overlap error, root mean square volume variation comparing to the input dataset models.

We claim the feasibility of generating digital mesh animation for visualizing real scene of realistic human in motion using cage capture. In addition, the deformation driven by the cage offers an affordable silhouette-consistency with respects to all images recorded by all cameras. Because the fixed connectivity of the cage is preserved across the sequence our technique well suited for encoding and compression. To show the accuracy of cage-based tracking we have developed a 3D video player that displays in real-time the cage-based performance animation. To increase the realism of the rendering, the enclosed model is rendered using an omnidirectional texture mapping generated from the multiple views video stream. Cage-based deformation allows the 3D video player to produce a smooth and accurate frame-to-frame playback of time-varying mesh.

8 Discussion

We have shown the feasibility of using cage-based parametrization for multiple-views video-based acquisition. The main advantage of our framework is its linear form, as



Fig. 7: Projection of the wireframe mesh generated by the estimated cage with 5% of constraints randomly selected $(1^{st}$ and 3^{rd} rows), silhouette overlap between the rendered silhouette and the extracted silhouette $(2^{nd}$ and 4^{th} rows.) (crane and handstand dataset)

well as the reduction of mesh parameters, which is independent of the surface resolution making possible reuse. Our harmonic cage-based deformation allows mesh rim vertices to fit the silhouette more precisely comparing to skeleton based deformation only because of more degrees of freedom. Our techniques can efficiently facilitated the extraction of the deformation subspace of mesh animations by retrieving the cage for all frames using a minimization framework fully independent of the model resolution. In addition, our technique drastically decreases the size of dataset without any lost of quality. Nevertheless our method suffer of some drawbacks directly derived from laplacian and space deformation properties. For example, the volume shrinking can provoke interior folding and potential fold over under non-silhouette consistent target points. Another major limitation of our method is that the deformation result depend on the shape and the tessellation of the cage. Moreover automatic cage generation for the setup process is also an opened hard problem.

9 Conclusion and Future Works

Even if there has been seen a strong interest for template-based approaches and multiview photometric for performance capture, no previous work tried to use cage-based parametrization for mesh sequence tracking and animation retrieval. In this paper, we have investigated the opportunities in-between cage-based tracking and multiple-views spatio-temporal reconstructed shape. We have developed a framework incorporating cage-based optimization in the context of the multi-view setup and captured the space deformation.

This cage-based deformation technique is a useful tool to improve the incremental reconstruction across time, because this method can provide a better control over the surface to allow rim vertices to fit the silhouette without prior knowledge of rigid parts. In this paper, we demonstrate the strength of harmonic coordinates used inside a linear minimization framework to reconstruct an enclosed mesh fitting the silhouette better than a skeleton. We show that our method is adapted in the context of a multiple-views setting with a proper experimental validation Finally, our novel approach is also very interesting for 3D video compression and animation reuse.

This work opens up a lot of new and interesting research directions. This algorithm is simple enough to be widely implemented and tested with previous framework. In the future, we plan to investigate and explore the possibility of achieving incremental 4D reconstruction, not relying on pre-shape dense sequences. Our method could be improved easily by integrating several image-based reconstruction cues such as sparse features like surface texture, silhouette and motion features observed in multiple viewpoint images.

References

- Starck, J., Hilton, A.: Surface capture for performance-based animation. IEEE Comput. Graph. Appl. 27 (2007) 21–31
- de Aguiar, E., Stoll, C., Theobalt, C., Ahmed, N., Seidel, H.P., Thrun, S.: Performance capture from sparse multi-view video. In: SIGGRAPH '08: ACM SIG-GRAPH 2008 papers, New York, NY, USA, ACM (2008) 1–10
- Gall, J., Stoll, C., de Aguiar, E., Theobalt, C., Rosenhahn, B., Seidel, H.P.: Motion capture using joint skeleton tracking and surface estimation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR'09), Miami, USA, IEEE Computer Society (2009) 1–8
- Vlasic, D., Baran, I., Matusik, W., Popović, J.: Articulated mesh animation from multi-view silhouettes. In: SIGGRAPH '08: ACM SIGGRAPH 2008 papers, New York, NY, USA, ACM (2008) 1–9
- Vlasic, D., Peers, P., Baran, I., Debevec, P., Popović, J., Rusinkiewicz, S., Matusik, W.: Dynamic shape capture using multi-view photometric stereo. ACM Trans. Graphics (Proc. SIGGRAPH Asia) 28 (2009)

- Ballan, L., Cortelazzo, G.M.: Marker-less motion capture of skinned models in a four camera set-up using optical flow and silhouettes. In: 3DPVT, Atlanta, GA, USA (2008)
- T. Takai, S. Nobuhara, H.Y., Matsuyama, T.: 3d video technologies: Capturing high fidelity full 3d shape, motion, and texture. In: International Workshop on Mixed Reality Technology for Filmmaking (in cooperation with ISMAR 2006). (2006)
- Huang, P., Hilton, A., Starck, J.: Human motion synthesis from 3d video. Computer Vision and Pattern Recognition, IEEE Computer Society Conference on 0 (2009) 1478–1485
- Lipman, Y., Sorkine, O., Cohen-Or, D., Levin, D., Rössl, C., Seidel, H.P.: Differential coordinates for interactive mesh editing. In: Proceedings of Shape Modeling International, IEEE Computer Society Press (2004) 181–190
- Sorkine, O.: Differential representations for mesh processing. Computer Graphics Forum 25 (2006) 789–807
- Lipman, Y., Sorkine, O., Alexa, M., Cohen-Or, D., Levin, D., Rössl, C., Seidel, H.P.: Laplacian framework for interactive mesh editing. International Journal of Shape Modeling (IJSM) 11 (2005) 43–61
- Au, O.K.C., Tai, C.L., Liu, L., Fu, H.: Dual laplacian editing for meshes. IEEE Transactions on Visualization and Computer Graphics 12 (2006) 386–395
- Luo, Q., Liu, B., Ma, Z.g., Zhang, H.b.: Mesh editing in roi with dual laplacian. In: CGIV '07: Proceedings of the Computer Graphics, Imaging and Visualisation, Washington, DC, USA, IEEE Computer Society (2007) 195–199
- Sorkine, O., Alexa, M.: As-rigid-as-possible surface modeling. In: Proceedings of Eurographics/ACM SIGGRAPH Symposium on Geometry Processing. (2007) 109–116
- Meyer, M., Lee, H., Barr, A., Desbrun, M.: Generalized barycentric coordinates on irregular polygons. Journal of Graphics Tools 7 (2002) 13–22
- Ju, T., Schaefer, S., Warren, J.: Mean value coordinates for closed triangular meshes. In: SIGGRAPH '05: ACM SIGGRAPH 2005 Papers, New York, NY, USA, ACM (2005) 561–566
- Joshi, P., Meyer, M., DeRose, T., Green, B., Sanocki, T.: Harmonic coordinates for character articulation. ACM Trans. Graph. 26 (2007) 71
- Lipman, Y., Levin, D., Cohen-Or, D.: Green coordinates. In: ACM SIGGRAPH 2008 papers, New York, NY, USA, ACM (2008) 78:1–78:10
- Ju, T., Zhou, Q.Y., van de Panne, M., Cohen-Or, D., Neumann, U.: Reusable skinning templates using cage-based deformations. ACM Trans. Graph. 27 (2008) 1–10
- Chuhua Xian, Hongwei Lin, S.G.: Automatic generation of coarse bounding cages from dense meshes. In: IEEE International Conference on Shape Modeling and Applications (Shape Modeling International 2009). (2009)
- Sumner, R.W., Zwicker, M., Gotsman, C., Popović, J.: Mesh-based inverse kinematics. ACM Trans. Graph. 24 (2005) 488–495
- Der, K.G., Sumner, R.W., Popović, J.: Inverse kinematics for reduced deformable models. In: SIGGRAPH '06: ACM SIGGRAPH 2006 Papers, New York, NY, USA, ACM (2006) 1174–1179
- Shi, X., Zhou, K., Tong, Y., Desbrun, M., Bao, H., Guo, B.: Mesh puppetry: cascading optimization of mesh deformation with inverse kinematics. ACM Trans. Graph. 26 (2007) 81