

Dense and Accurate Spatio-Temporal Multi-View Stereovision

Jérôme Courchay¹, Jean-Philippe Pons²,
Pascal Monasse¹, and Renaud Keriven¹

¹ IMAGINE , 6 Av Blaise Pascal - Cité Descartes, Marne-la-Vallée, FR
{courchay,monasse,keriven}@imagine.enpc.fr

² CSTB , 290 route des Lucioles , Sophia Antipolis, FR
jean-philippe.pons@cstb.fr

Abstract. In this paper, we propose a novel method to simultaneously and accurately estimate the 3D shape and 3D motion of a dynamic scene from multiple-viewpoint calibrated videos. We follow a variational approach in the vein of previous work on stereo reconstruction and scene flow estimation. We adopt a representation of a dynamic scene by an animated mesh, i.e. a polygonal mesh with fixed connectivity whose time-varying vertex positions sample the trajectories of material points. Interestingly, this representation ensures a consistent coding of shape and motion by construction. Our method accurately recovers 3D shape and 3D motion by optimizing the positions of the vertices of the animated mesh. This optimization is driven by an energy function which incorporates multi-view and inter-frame photo-consistency, smoothness of the spatio-temporal surface and of the velocity field. Central to our work is an image-based photo-consistency score which can be efficiently computed and which fully handles projective distortion and partial occlusions. We demonstrate the effectiveness of our method on several challenging real-world dynamic scenes.

Key words: Spatio-temporal stereovision, Scene flow, Motion capture

1 Introduction

In recent years, several methods for automatic generation of complete spatio-temporal models of dynamic scenes from multiple videos have been proposed [1–15]. In particular, the most recent ones have proven effective for full-body marker-less motion capture, yielding visually impressive results. However, when taking a closer look at the aforementioned techniques, it becomes apparent that very few of them achieve a desirable *coupled, dense and accurate 3D shape and 3D motion estimation*.

Accurate 3D shape. Many recent techniques still produce an approximate geometry: free-form deformation of a template body model [2, 11, 15], visual hull [1, 3, 15], Laplacian deformation of a laser scan of the initial pose [4, 5]. These methods are unable to recover genuine geometric details such as facial expressions and clothing folds and wrinkles.

Accurate 3D motion estimation is crucial in some applications like motion transfer and time interpolation. Also, a coarse motion estimation precludes the enforcement of temporal consistency constraints during coupled shape and motion estimation. However, in most existing performance capture techniques, 3D scene flow [16], i.e. the dense 3D motion field of the scene, is not accurately estimated. Often, it is interpolated from sparse 3D correspondences [3, 4, 12]. Some methods do not address 3D motion estimation whatsoever: [7] uses a four-dimensional level set representation which, beyond its very high computational and memory requirements, does not encode 3D correspondence. [10, 15] produce animated meshes but, despite appearances, the underlying 3D correspondences are purely artificial.

Coupled 3D shape and 3D motion estimation allows to exploit their redundancy, and has long been recognized [17] as a desirable way to improve their performance. However, most marker-less motion capture methods fail to integrate spatio-temporal consistency constraints. In [3, 12, 13], shape is computed independently in each time frame, prior to motion estimation. In [9], shape and motion are estimated sequentially, not simultaneously. In [5], an initial mesh is propagated by 3D scene flow, under silhouette constraints, but without any stereo cues; as a result, this method suffers from temporal drift. The latter is circumvented in [4] by substituting sparse 3D correspondences for dense 3D scene flow, but neither shape or motion are accurate enough to allow enforcing spatio-temporal consistency. In [1, 7], a certain degree of spatio-temporal coherence is obtained through four-dimensional representations, but as these representations do not encode temporal correspondence, they cannot exploit inter-frame matching constraints. In [14], shape and motion are estimated simultaneously using a plane-sweep carving algorithm in a 6D space, but this approach has a very high computational and memory cost, is limited to two frames, and is unable to enforce the smoothness of the recovered shape and motion.

Thus, to our knowledge, two methods [6, 8] achieve this highly desirable coupled, dense and accurate 3D shape and 3D motion estimation. In [8], shape and motion are represented through the detail coefficients of a time-varying subdivision surface. The latter coefficients are estimated by simultaneously optimizing multi-view and inter-frame photo-consistency. However, the non-linearity of the chosen multi-resolution representation makes this optimization intricate. Also, the required motion initialization relies on the spatio-temporal derivatives of the input images, thereby making it applicable mainly to slowly-moving Lambertian scenes under constant illumination.

[6] is the only work to date which can handle complex real-world dynamic scenes. Despite the effectiveness of this method, we believe that the expansion framework used does not allow to take into account the full visibility depending on occluding patch not computed yet.

In this paper, we propose a novel method to simultaneously and accurately estimate the 3D shape and 3D motion of a dynamic scene from multiple-viewpoint videos. First, we follow a **variational approach** in the vein of previous work on stereo reconstruction and scene flow estimation [9, 17–21]. None of these meth-

ods fits our applications in their current state: most are limited to a single time-varying depth map of the scene [17–21], while others do not enforce spatio-temporal consistency constraints [9, 19].

Second, we adopt a representation of a dynamic scenes by an **animated mesh**, i.e. a polygonal mesh with fixed connectivity whose time-varying vertex positions sample the trajectories of material points. Interestingly, this representation ensures a consistent coding of shape and motion by construction. It is widely used in computer graphics, especially in computer animation. It is also popular for performance capture from video [3–6, 10, 11, 15] or from time-varying point clouds [22, 23] (the latter being obtained from video or from fast 3-D scanning hardware).

Our method accurately recovers 3D shape and 3D motion by optimizing the positions of the vertices of the animated mesh. This optimization is driven by an energy function which incorporates multi-view and inter-frame photo-consistency, smoothness of the spatio-temporal surface and of the velocity field. Central to our work is an image-based photo-consistency score which can be efficiently computed and which fully handles projective distortion and partial occlusions, in the spirit of [9].

The rest of this article is organized as follows. In Section 2, we describe in detail the discrete geometric representation, the variational formulation, the energy function and the associated minimization procedure which constitute our approach. In Section 3, we discuss implementation aspects and we demonstrate the effectiveness of our method on several challenging real-world dynamic scenes.

2 Our Approach

2.1 Discretize then optimize

An overwhelming majority of variational methods in this area [9, 17–19, 21] and more generally in computer vision, rely on an *optimize then discretize* approach: an energy functional depending on a continuous infinite-dimensional spatio-temporal representation is considered, the gradient of this energy functional is computed analytically, then the obtained evolution flow is discretized.

In contrast, we adopt a *discretize then optimize* approach: we define an energy function depending on a discrete finite-dimensional spatio-temporal representation, and we use standard non-convex optimization tools. The benefits of this approach have long been recognized in mesh processing, but have seldom been demonstrated in computer vision [24–26]. Thus, the choice of an adequate discrete spatio-temporal representation is crucial in our work.

2.2 Animated mesh representation

In our context, animated polygonal meshes present many significant advantages. Compared to unrelated meshes at different time instants, they are more compact, easier to store and to manipulate. They provide a direct access both to the shape

of the scene at a given time instant, and to motion trajectories. 3D shape and 3D motion are mutually consistent by construction.

Their fixed topology may be regarded as a limitation, as argued in [12]. We believe that it is not, since the human body has a constant - spherical, if disregarding pierces - topology. It is questionable to treat a character with hands on hips as a genus-2 torus. It should rather be regarded as a topological sphere with some temporary contact regions.

Furthermore, let us mention that our method is not limited to a spherical topology: while the topology of the animated mesh is constant *across time*, we are able to modify it *across our optimization process* using a spatio-temporal version of Delaunay deformable models [27].

2.3 Variational formulation

In the following, we consider a dynamic scene, imaged by N calibrated and synchronized video sequences composed of T frames, and represented by an animated polygonal mesh with K vertices. We note:

- $I_{i,t} : \Omega_i \subset \mathbb{R}^2 \rightarrow \mathbb{R}^d$, $i \in \{1..N\}$, $t \in \{1..T\}$ the input images. In practice $d = 1$ for grayscale images and $d = 3$ for color images.
- $\mathbf{X} = \{\mathbf{x}_{k,t}, k \in \{1..K\}, t \in \{1..T\}\}$ the 3D positions of the vertices of the animated mesh at the different time instants,
- $\mathbf{X}_t = \{\mathbf{x}_{k,t}, k \in \{1..K\}\}$ the t^{th} temporal slice of the animated mesh.

In the sequel, by a slight abuse of notation, we indistinctly use \mathbf{X} and \mathbf{X}_t to refer to the animated mesh and to the positions of its vertices.

The energy to minimize with respect to \mathbf{X} is composed of a data attachment term, of a regularization term for the spatio-temporal surface and of a regularization term for the velocity field:

$$E(\mathbf{X}) = E_D(\mathbf{X}) + \lambda_S E_S(\mathbf{X}) + \lambda_V E_V(\mathbf{X}) . \quad (1)$$

E_D encourages multi-view and frame-to-frame matching consistency. It is defined as the sum over camera pairs (i, j) and pairs of time frames (t, u) of a dissimilarity measure between image $I_{i,t}$ and the reprojection of $I_{j,u}$ via the animated mesh. The detailed description of this term is left to Section 2.4.

E_S favors the regularity of the spatio-temporal surface. We use the total area of the animated mesh. The minimization of this term by gradient descent yields a discrete version of the well known mean curvature motion, which we implement as described in [28].

E_V penalizes rapid variations of the velocity field along the animated mesh. It is the total squared L^2 norm over the animated mesh of the gradient of the velocity field. The detailed description of this term is left to Section 2.5.

We minimize the above energy function using a standard gradient descent on the spatio-temporal positions \mathbf{X} . In order to avoid unwanted local minima, we resort to a multi-resolution and chronological scheme. We first optimize the first two frames of a low-resolution animated mesh using low-resolution versions




Fig. 1. Finite element representation over a facet (k, l, m) of the animated mesh.

of input images. Then we initialize an additional time frame by extrapolating 3D position from speed and acceleration of previous frames. We iteratively add time frames, and optimize the sequence using a sliding time window of a few frames, until we reconstruct the whole temporal sequence at low resolution. We then refine the obtained spatio-temporal mesh with increased image and mesh resolutions, until we reach the desired accuracy.

2.4 Data attachment term

The formal definition of E_D and of its gradient requires some additional notations. The perspective projection performed by camera i is denoted by $\Pi_i : \mathbb{R}^3 \rightarrow \mathbb{R}^2$. Our method takes into account the visibility of the surface points. We refer to $\mathbf{X}_{i,t}$ as the part of the temporal slice \mathbf{X}_t visible in image i . The back-projection of a point of camera i on the animated mesh at frame t is denoted by $\Pi_{i,\mathbf{X}_t}^{-1} : \Pi_i(\mathbf{X}_t) \rightarrow \mathbf{X}_{i,t}$.

We also define 3D transport functions $T_{\mathbf{X}_t \rightarrow \mathbf{X}_u}$ that map points in \mathbf{X}_t to points in \mathbf{X}_u . This can be written formally using the linear finite-element representation depicted in Figure 1. For each vertex k of the animated mesh at some time frame t , we define a basis function ϕ_{k,\mathbf{X}_t} such that (i) $\phi_{k,\mathbf{X}_t}(\mathbf{x}_{k,t}) = 1$ (ii) $\forall l \neq k, \phi_{k,\mathbf{X}_t}(\mathbf{x}_{l,t}) = 0$ (iii) ϕ_{k,\mathbf{X}_t} varies linearly inside the triangular facets adjacent to the k^{th} vertex, and cancels outside this ring. We then have at pixel p_i in image i :

$$T_{\mathbf{X}_t \rightarrow \mathbf{X}_u} = \sum_k \mathbf{x}_{k,u} \phi_{k,\mathbf{X}_t}. \quad (2)$$

In a simpler way we can say that the back-projection Y_t of pixel p_i lies on a triangular facet f and has barycentric coordinates $\phi_{l,\mathbf{X}_t}(Y_t)$ at time t , l being a vertex of f . So the position of this particle at time u is $Y_u = \sum_{l \in f} \mathbf{x}_{l,u} \phi_{l,\mathbf{X}_t}(Y_t)$, that is $Y_u = \sum_{k \in \mathbf{X}} \mathbf{x}_{k,u} \phi_{k,\mathbf{X}_t}(Y_t)$ since $\phi_{k,\mathbf{X}_t}(Y_t)$ cancels if vertex k is outside facet f .

Finally, we define image transport functions $T_{(i,\mathbf{X}_t) \rightarrow (j,\mathbf{X}_u)}$ which map positions in $I_{i,t}$ to positions in $I_{j,u}$ via the animated mesh:

$$T_{(i,\mathbf{X}_t) \rightarrow (j,\mathbf{X}_u)} = \Pi_j \circ T_{\mathbf{X}_t \rightarrow \mathbf{X}_u} \circ \Pi_{i,\mathbf{X}_t}^{-1}. \quad (3)$$




Fig. 2. Reprojection of image j at time u in image i at time t via the animated mesh.

With these notations in hand the reprojection of image j at time u in image i at time t via the animated mesh writes $I_{j,u} \circ T_{(i, \mathbf{X}_t) \rightarrow (j, \mathbf{X}_u)}$. This is illustrated in Figure 2.

The data attachment term is the sum over oriented camera pairs (i, j) and oriented pairs (t, u) of time frames of a dissimilarity measure M between image $I_{i,t}$ and the above defined reprojection of $I_{j,u}$ via the animated mesh. The dissimilarity is computed only over the region of image plane i where both images are defined, i.e. after discarding semi-occluded regions. This image region writes $\Pi_i(\mathbf{X}_{i,t} \cap T_{\mathbf{X}_u \rightarrow \mathbf{X}_t}(\mathbf{X}_{j,u}))$. More clearly, pixel p_i in image i is visible in both images, if its back-projection lies on the surface at time t , and this point on the surface once transported at time u is visible (none occluded, none outside the image frame) in image $I_{j,u}$. This visible image region is computed before each optimization step on graphics hardware. For conciseness, we will omit it in the equations below:

$$E_D(\mathbf{X}) = \sum_{i,j} \sum_{t,u} M [I_{i,t}, I_{j,u} \circ T_{(i, \mathbf{X}_t) \rightarrow (j, \mathbf{X}_u)}] . \quad (4)$$

We now compute the partial derivative of this energy term with respect to the variation of a single position $\mathbf{x}_{k,t}$ of the animated mesh. First, we note that the only oriented pairs of time frames affected by such a variation are (u, t)

and (t, u) , $u \in \{1..T\}$. Second, when the animated mesh moves, the reprojected image changes. Hence the partial derivative of E_D involves the derivative of the similarity measure M with respect to its second argument, denoted by $\partial_2 M$.

Using the chain rule, and after some index manipulations, we get:

$$\begin{aligned} \frac{\partial E_D}{\partial \mathbf{x}_{k,t}} &= \sum_{i,j} \sum_u \\ &\int_{\Omega_i} \partial_2 M [I_{i,t}, I_{j,u} \circ T_{(i,\mathbf{X}_t) \rightarrow (j,\mathbf{X}_u)}] DI_{j,u} \frac{\partial T_{(i,\mathbf{X}_t) \rightarrow (j,\mathbf{X}_u)}(\mathbf{p}_i)}{\partial \mathbf{x}_{k,t}}(\mathbf{p}_i) d\mathbf{p}_i \\ &+ \int_{\Omega_j} \partial_2 M [I_{j,u}, I_{i,t} \circ T_{(j,\mathbf{X}_u) \rightarrow (i,\mathbf{X}_t)}] DI_{i,t} \frac{\partial T_{(j,\mathbf{X}_u) \rightarrow (i,\mathbf{X}_t)}(\mathbf{p}_j)}{\partial \mathbf{x}_{k,t}}(\mathbf{p}_j) d\mathbf{p}_j, \quad (5) \end{aligned}$$

where $DI_{..}$ denotes the Jacobian matrices of the input images. For conciseness, we have omitted the points where the latter are evaluated in the above equation.

As regards the quantities $\frac{\partial T_{..}}{\partial \mathbf{x}_{k,t}}$, we can make several observations. First, they are purely geometric, i.e. independent of image data. Second, they cancel outside the ring of triangular facets adjacent to the k^{th} vertex. Hence, despite appearances, integration is performed only over the visible projection of this ring in the different images, not over the full image domains. Third, these quantities involve the normal of the triangular facet visible at pixel \mathbf{p}_i , and the barycentric coordinate of $\mathbf{x}_{k,t}$ in this facet. Complete expressions can be obtained using a non trivial geometric reasoning. We refer the interested reader to [29], where the detailed numerical computation, but also an additional intuitive solving are proposed. The numerical solving, mainly consist in computing how barycentric coordinates change for a small perturbation of the surface.

2.5 Velocity field regularization term

The velocity field is unambiguously encoded by the animated mesh \mathbf{X} . Specifically, it is a continuous and piecewise linear vector field $\mathbf{X}_t \rightarrow \mathbb{R}^3$ defined by

$$\mathbf{v}_{\mathbf{X},t}(\mathbf{x}) = T_{\mathbf{X}_t \rightarrow \mathbf{X}_{t+1}}(\mathbf{x}) - \mathbf{x}, \quad (6)$$

or equivalently by

$$\mathbf{v}_{\mathbf{X},t} = \sum_k (\mathbf{x}_{k,t+1} - \mathbf{x}_{k,t}) \phi_{k,\mathbf{X}_t}. \quad (7)$$

The velocity field regularization term writes:

$$E_V(\mathbf{X}) = \sum_t \int_{\mathbf{X}_t} \|\nabla \mathbf{v}_{\mathbf{X},t}(\mathbf{x})\|^2 d\mathbf{x}. \quad (8)$$

To simplify this expression, we use the fact that $\nabla \phi_{k,\mathbf{X}_t}$ is constant in each triangular facet f of \mathbf{X}_t and equals $\frac{\mathbf{h}_{k,f}}{\|\mathbf{h}_{k,f}\|^2}$, where $\mathbf{h}_{k,f}$ is triangle's height going through vertex k . A_f being the area of f , the energy term becomes:

$$E_V(\mathbf{X}) = \sum_t \sum_{f \in \mathbf{X}_t} A_f \left\| \sum_{k \in f} \frac{\mathbf{h}_{k,f}}{\|\mathbf{h}_{k,f}\|^2} (\mathbf{x}_{k,t+1} - \mathbf{x}_{k,t}) \right\|^2. \quad (9)$$

If we neglect the variation of $\mathbf{h}_{k,f}$ with respect to vertex displacement, the partial derivatives $\frac{\partial E_V}{\partial \mathbf{x}_{k,t}}$ of this energy term can now easily be derived.

3 Numerical Experiments

3.1 Implementation aspects

The computation of image reprojections via the animated mesh and of the gradient of the data attachment term are the most expensive parts of our algorithm. Hence, they are implemented on GPU using the OpenGL API and the Cg shading language.

In all our experiments, we choose the opposite of normalized cross correlation as the image dissimilarity measure M , in order to accommodate moving shadows and time-varying lighting conditions.

The storage of the animated mesh and the computation of spatio-temporal smoothing terms are based on the C++ Computational Geometry Algorithms Library (CGAL)³.

The resolution of the mesh is controlled by a lower and an upper edge length thresholds, that are applied to the whole time sequence: an edge is bisected if it is longer than the upper threshold in *at least one* time frame; an edge is collapsed if it is shorter than the lower threshold in *all* time frames. The topology of the mesh is automatically corrected when needed by applying Delaunay deformable models [27] to the coordinates of the animated mesh at a reference time frame. The user chooses a reference frame that reflects the actual topology of the scene: e.g a pose with arms and legs slightly apart for human motion.

3.2 Experimental results

We have tested our algorithm on two challenging multi-view video sequences of non-rigid scenes.

The “Pants” dataset is composed of 8 cameras 480×640 pixels. It is courtesy of R. White, K. Crane and D.A. Forsyth [30]. We have successfully applied our algorithm to the first 60 frames of this dataset. Due to the high image resolution, four multi-resolution scales have been used to obtain the accurate spatio-temporal reconstruction shown in Figure 3. It has taken 24 hours to reconstruct the 60 frames spatio-temporal model.

Figure 4 demonstrates the superiority of our spatio-temporal approach compared to a frame-by-frame multi-view stereovision method [9], on the “Pants” dataset. The improvements are three-fold: (i) our approach exploits speed and acceleration to make better initial guesses of the subsequent time frames, thus being less prone to unwanted local minima (ii) thanks to the enforcement of temporal coherence, our approach is less likely to fail in regions with low photo-consistency evidence (iii) our approach simultaneously and consistently estimates 3D shape and 3D scene flow.

³ <http://www.cgal.org/>




Fig. 3. Our results on the “Pants” dataset. See text for more details.

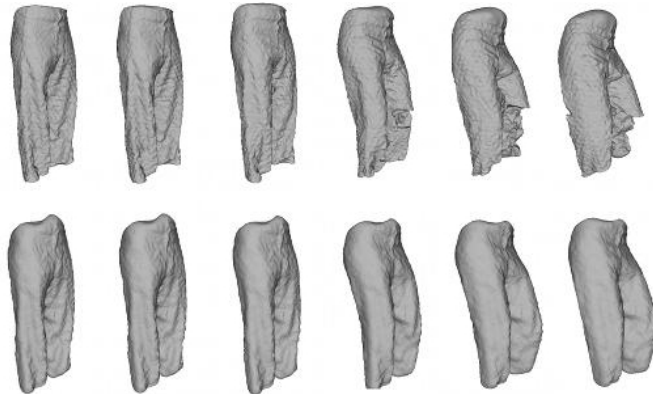


Fig. 4. Comparison between a frame-by-frame multi-view stereovision approach (top) and our spatio-temporal approach (bottom) on the “Pants” dataset. See text for details.




Fig. 5. Our results on the “Dancer” dataset. See text for more details.

The “Dancer” dataset was made available to us by the 4Dviews company⁴. It was acquired by 14 calibrated and synchronized video cameras 1000×1000 pixels. We have applied our algorithm to the first 10 frames of this dataset. To bootstrap our multi-resolution and chronological optimization procedure, we have used a standard stereo-vision algorithm at the first time frame. The obtained reconstruction after processing three multi-resolution levels is displayed in Figure 5. We insist on the fact that we have not used silhouette information and that stereovision on such a dataset is quite challenging: because it was design for visual hull based techniques, many parts of the subject are textureless. It has taken 10 hours to teconstruct the full spatio-temporal model.

4 Conclusion

We have presented a novel variational approach to dense and accurate 3D shape and motion reconstruction from multi-view video sequences. Our method leverages the benefits of the animated mesh representation, of image-based photo-consistency, of discrete geometric optimization and of GPU computation. We have validated our algorithm on two challenging real datasets, and obtained results that rival state-of-the-art techniques.

⁴ <http://4dviews.com>

References

1. Aganj, E., Pons, J.P., Ségonne, F., Keriven, R.: Spatio-temporal shape from silhouette using four-dimensional Delaunay meshing. In: *IEEE International Conference on Computer Vision*. (2007)
2. Ahmed, N., de Aguiar, E., Theobalt, C., Magnor, M., Seidel, H.P.: Automatic generation of personalized human avatars from multi-view video. In: *Proc. VRST'05*. (2005) 257–260
3. Ahmed, N., Theobalt, C., Rössl, C., Thrun, S., Seidel, H.P.: Dense correspondence finding for parametrization-free animation reconstruction from video. In: *IEEE Conference on Computer Vision and Pattern Recognition*. (2008)
4. de Aguiar, E., Stoll, C., Theobalt, C., Ahmed, N., Seidel, H.P., Thrun, S.: Performance capture from sparse multi-view video. In: *ACM SIGGRAPH*. (2008)
5. de Aguiar, E., Theobalt, C., Stoll, C., Seidel, H.P.: Marker-less deformable mesh tracking for human shape and motion capture. In: *IEEE Conference on Computer Vision and Pattern Recognition*. (2007)
6. Furukawa, Y., Ponce, J.: Dense 3D motion capture from synchronized video streams. In: *IEEE Conference on Computer Vision and Pattern Recognition*. (2008)
7. Goldlücke, B., Magnor, M.: Space-time isosurface evolution for temporally coherent 3D reconstruction. In: *IEEE Conference on Computer Vision and Pattern Recognition*. Volume 1. (2004) 350–355
8. Neumann, J., Aloimonos, Y.: Spatio-temporal stereo using multi-resolution subdivision surfaces. *The International Journal of Computer Vision* **47**(1–3) (2002) 181–193
9. Pons, J.P., Keriven, R., Faugeras, O.: Multi-view stereo reconstruction and scene flow estimation with a global image-based matching score. *The International Journal of Computer Vision* **72**(2) (2007) 179–193
10. Starck, J., Hilton, A.: Surface capture for performance based animation. *IEEE Computer Graphics and Applications* **27**(3) (2007) 21–31
11. Theobalt, C., Ahmed, N., Lensch, H., Magnor, M., Seidel, H.P.: Seeing people in different light-joint shape, motion, and reflectance capture. *IEEE Transactions on Visualization and Computer Graphics* **13**(4) (2007) 663–674
12. Varanasi, K., Zaharescu, A., Boyer, E., Horaud, R.P.: Temporal surface tracking using mesh evolution. In: *European Conference on Computer Vision*. Volume 2. (2008) 30–43
13. Vedula, S., Baker, S., Kanade, T.: Image-based spatio-temporal modeling and view interpolation of dynamic events. *ACM Transactions on Graphics* **24**(2) (2005) 240–261
14. Vedula, S., Baker, S., Seitz, S., Kanade, T.: Shape and motion carving in 6D. *IEEE Conference on Computer Vision and Pattern Recognition* (2000)
15. Vlastic, D., Baran, I., Matusik, W., Popović, J.: Articulated mesh animation from multi-view silhouettes. *ACM Transactions on Graphics* **27**(3) (2008)
16. Vedula, S., Baker, S.: Three-dimensional scene flow. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **27**(3) (2005) 475–480
17. Zhang, Y., Kambhamettu, C.: Integrated 3D scene flow and structure recovery from multiview image sequences. In: *IEEE Conference on Computer Vision and Pattern Recognition*. Volume 2. (2000) 674–681
18. Huguet, F., Devernay, F.: A variational method for scene flow estimation from stereo sequences. In: *IEEE International Conference on Computer Vision*. (2007)

19. Pons, J.P., Keriven, R., Faugeras, O., Hermosillo, G.: Variational stereovision and 3D scene flow estimation with statistical similarity measures. In: IEEE International conference on Computer Vision. Volume 2. 597–602
20. Wedel, A., Rabe, C., Vaudrey, T., Brox, T., Franke, U., Cremers, D.: Efficient dense scene flow from sparse or dense stereo data. In: European Conference on Computer Vision. (2008) 739–751
21. Zhang, Y., Kambhamettu, C.: On 3D scene flow and structure estimation. In: IEEE Conference on Computer Vision and Pattern Recognition. Volume 2. (2001) 778–785
22. Süßmuth, J., Winter, M., Greiner, G.: Reconstructing animated meshes from time-varying point clouds. *Computer Graphics Forum (Proc. Eurographics Symposium on Geometry Processing)* **27**(5) (2008) 1469–1476
23. Wand, M., Jenke, P., Huang, Q., Bokeloh, M., Guibas, L., Schilling, A.: Reconstruction of deforming geometry from time-varying point clouds. In: Eurographics Symposium on Geometry Processing. (2007) 49–58
24. Delaunoy, A., Prados, E., Gargallo, P., Pons, J.P., Sturm, P.: Minimizing the multi-view stereo reprojection error for triangular surface meshes. In: British Machine Vision Conference. (2008)
25. Slabaugh, G.G., Unal, G.B.: Active polyhedron: Surface evolution theory applied to deformable meshes. In: IEEE Conference on Computer Vision and Pattern Recognition. Volume 2. (2005) 84–91
26. Vu, H.H., Keriven, R., Labatut, P., Pons, J.P.: Towards high-resolution large-scale multi-view stereo. In: IEEE Conference on Computer Vision and Pattern Recognition. (2009)
27. Pons, J.P., Boissonnat, J.D.: Delaunay deformable models: Topology-adaptive meshes based on the restricted Delaunay triangulation. In: IEEE Conference on Computer Vision and Pattern Recognition. (2007)
28. Kobbelt, L., Campagna, S., Vorsatz, J., Seidel, H.P.: Interactive multi-resolution modeling on arbitrary meshes. In: International Conference on Computer Graphics and Interactive Techniques. (1998) 105–114
29. Courchay, J., Pons, J.P., Keriven, R., Monasse, P.: Dense and accurate spatio-temporal multi-view stereovision. Technical Report 09-43, IMAGINE (July 2009)
30. White, R., Crane, K., Forsyth, D.: Capturing and animating occluded cloth. In: SIGGRAPH. (2007)