

Efficient Multi-view Performance Capture of Fine-Scale Surface Detail

Nadia Robertini^{1,2}, Edilson De Aguiar^{1,3}, Thomas Helten^{1,4}, Christian Theobalt¹

¹MPI Informatik, Germany ²Intel VCI, Germany ³CEUNES / UFES, Brazil ⁴Pixargus / GoalControl, Germany

{nroberti, edeaguia, thelten, theobalt}@mpi-inf.mpg.de

Abstract

We present a new effective way for performance capture of deforming meshes with fine-scale time-varying surface detail from multi-view video. Our method builds up on coarse 4D surface reconstructions, as obtained with commonly used template-based methods. As they only capture models of coarse-to-medium scale detail, fine scale deformation detail is often done in a second pass by using stereo constraints, features, or shading-based refinement. In this paper, we propose a new effective and stable solution to this second step. Our framework creates an implicit representation of the deformable mesh using a dense collection of 3D Gaussian functions on the surface, and a set of 2D Gaussians for the images. The fine scale deformation of all mesh vertices that maximizes photo-consistency can be efficiently found by densely optimizing a new model-to-image consistency energy on all vertex positions. A principal advantage is that our problem formulation yields a smooth closed form energy with implicit occlusion handling and analytic derivatives. Error-prone correspondence finding, or discrete sampling of surface displacement values are also not needed. We show several reconstructions of human subjects wearing loose clothing, and we qualitatively and quantitatively show that we robustly capture more detail than related methods.

Acknowledgement. This project was supported by the ERC Starting Grant 335545 CapReal and CAPES Grant 11012-13-7.

1. Introduction

Performance capture methods enable the reconstruction of the motion, the dynamic surface geometry, and the appearance of real world scenes from multiple video recordings, for example, the deforming geometry of body and apparel of an actor, or his facial expressions [5, 7, 2, 18]. Many methods to capture space-time coherent surfaces reconstruct a coarse-to-medium scale 4D model of the scene in a first step, e.g. by deforming a mesh or a rigged template such that it aligns with the images [5, 18]. Finer




Figure 1. Given as input a multi-view video sequence (left - particular input frame) and a coarse mesh animation (middle - input mesh), our method is able to efficiently reconstruct fine scale surface details (right - refined mesh). Note the wrinkles and folds reconstructed on the skirt.

scale shape detail is then added in a second refinement step. In this second step, some methods align the surface to a combination of silhouette constraints and sparse image features [7]. But such approaches merely recover medium scale detail and may suffer from erroneous feature correspondences between images and shape. Photo-consistency constraints can also be used to compute smaller scale deformations via stereo-based refinement [5, 14]. However, existing approaches that follow that path often resort to discrete sampling of local displacements, since phrasing dense stereo based refinement as a continuous optimization problem has been more challenging [9]. Some recent methods resort to shading-based techniques to capture small-scale displacements, such as shape-from-shading or photometric stereo [22, 21, 18]. However, the methods either require controlled and calibrated lighting, or complex inverse estimation of lighting and appearance when they are applied under uncontrolled recording conditions.

In this paper, we contribute with a new effective solution to the refinement step using multi-view photo-consistency constraints. As input, our method expects synchronized and calibrated multiple video of a scene and a reconstructed

coarse mesh animation, as it can be obtained with previous methods from the literature. Background subtraction or image silhouettes are not required for refinement.

Our first contribution is a new shape representation that models the mesh surface with a dense collection of 3D Gaussian functions centered at each vertex and each having an associated color. A similar decomposition into 2D Gaussian functions is applied to each input video frame.

This scene representation enables our second contribution, namely the formulation of dense photo-consistency-based surface refinement as a global optimization problem in the position of each vertex on the surface. Unlike previous performance capture methods, we are able to phrase the model-to-image photo-consistency energy that guides the deformation as a closed form expression, and we can compute its analytic derivatives. Our problem formulation has the additional advantage that it enables implicit handling of occlusions, as well as spatial and temporal coherence constraints, while preserving a smooth consistency energy function. We can effectively minimize this function in terms of dense local surface displacements with standard gradient-based solvers. In addition to these advantages, unlike many previous methods, our framework does not require a potentially error-prone sparse set of feature correspondences or discrete sampling and testing of surface displacements, and thus provides a new way of continuous optimization of the dense surface deformation.

We used our approach for reconstructing full-body performances of human actors wearing loose clothing, and performing different motions. Initial coarse reconstructions of the scene were obtained with the approaches by Gall *et al.* [7] and Starck and Hilton [14]. Our results (Fig. 1 and Sect. 6) show that our approach is able to reconstruct more of the fine-scale detail that is present in the input video sequences, than the baseline methods, for instance the wrinkles in a skirt. We also demonstrate these improvements quantitatively.

2. Related Work

Marker-less performance capture methods are able to reconstruct dense dynamic surface geometry of moving subjects from multi-view video, for instance of people in loose clothing, possibly along with pose parameters of an underlying kinematic skeleton [16]. Most of them use data from dense multi-camera systems and recorded under controlled studio environments. Some methods employ variants of shape-from-silhouette or active or passive stereo [23, 11, 14, 20, 17], which usually results in temporally incoherent reconstructions. Space-time coherency is easier to achieve with model-based approaches that deform a static shape template (obtained by a laser scan or image-based reconstruction) such that it matches the subject, *e.g.* a person [4, 5, 18, 1, 7] or a person’s apparel [2]. Some

of them jointly track a skeleton and the non-rigidly deforming surface [18, 1, 6]; also multi-person reconstruction has been demonstrated [10]. Other approaches use a generally deformable template without embedded skeleton to capture 4D models, *e.g.* an elastically deformable surface or volume [5, 13], or a patch-based surface representation [3]. Most of the approaches mentioned so far either only reconstruct coarse dynamic surface models that lack fine scale detail, or coarse reconstruction is a first stage. Fine scale detail is then added to the coarse result in a second refinement step.

Some methods use a combination of silhouette constraints and sparse feature correspondences to estimate, at best, a medium scale non-rigid 4D surface detail [7]. Other approaches use stereo-based photo-consistency constraints in addition to silhouettes to achieve denser estimates of finer scale deformations [14, 5]. It is an involved problem to phrase dense stereo-based surface refinement as a continuous optimization problem, as it is done in variational approaches [9]. Thus, stereo-based refinement in performance capture often resorts to discrete surface displacement sampling which are less efficient, and with which globally smooth and coherent solutions are harder to achieve.

In this paper, we propose a new formulation of stereo-based surface refinement as a continuous optimization problem, which is based on a new surface representation with Gaussian functions. In addition, our refinement method also succeeds if silhouettes are not available, making the approach more generally applicable.

An alternative way to recover fine-scale deforming surface detail is to use shading-based methods, *e.g.* shape-from-shading or photometric stereo [21]. Many of these approaches require controlled and calibrated lighting [8, 19], which reduces their applicability. More recently, shading-based refinement of dynamic scenes captured under more general lighting was shown [22], but these approaches are computationally challenging as they require to solve an inverse rendering problem to obtain estimates of illumination, appearance and shape at the same time.

The method we propose has some similarity to the work of Sand *et al.* [12] who capture skin deformation as a displacement field on a template mesh; however, they require marker-based skeleton capture, and only fit the surface to match the silhouettes in multi-view video. Our problem formulation is inspired by the work of Stoll *et al.* [15] who used a collection of Gaussian functions in 3D and 2D for marker-less skeletal pose estimation. Estimation of surface detail was not the goal of that work. Our paper extends their basic concept to the different problem of dense stereo-based surface estimation using continuous optimization of a smooth energy that can be formulated in closed form, and that has analytic derivatives.




Figure 2. Overview of our framework. Our approach refines the input coarse mesh animation by maximizing the color consistency between the collection of 3D surface Gaussians, associated to the input vertices, and the set of 2D image Gaussians, assigned to image patches.

3. Overview

An overview of our approach is shown in Fig. 2. The input to our algorithm is a calibrated and synchronized multi-view video sequence showing images of the human subject. In addition, we assume as input a spatio-temporally coherent coarse animated mesh sequence, reconstructed from multi-view video related approaches [7, 14].

Our method refines the initial coarse animation such that the fine dynamic surface details are incorporated to the meshes. First, we create an implicit representation of the input mesh using a dense collection of 3D Gaussian functions on the surface with associated colors. The input images are also represented as a set of 2D Gaussian associated to image patches in each camera view. Thereafter, continuous optimization is performed to maximize the color consistency between the collection of 3D surface Gaussians and the set of 2D image Gaussians. The optimization displaces the 3D Gaussians along the associated vertex normal of the coarse mesh which yields the necessary vertex displacement.

Our optimization scheme has a smooth energy function, that, thanks to our Gaussians-based model, can be expressed in closed form. It further allows us to analytically compute derivatives, enabling the possibility of using efficient gradient-based solvers.

4. Implicit Model

Our framework converts the input coarse animation and input multi-view images into implicit representations using a collection of Gaussians: 3D surface Gaussians on the mesh surface with associated colors and 2D image Gaussians, with associated colors, assigned to image patches in each camera view.

4.1. 3D Surface Gaussian

Our implicit model for the input mesh is obtained by placing a 3D Gaussian at each mesh vertex v_s , $\forall s \in \{0 \dots n_s - 1\}$, n_s being the number of vertices. A 3D unnormalized isotropic Gaussian function on the surface is defined simply with a mean $\hat{\mu}_s$, that coincides with the vertex location, and a standard deviation $\hat{\sigma}_s$ (equally set to 7 mm for all 3D Gaussians on surface) as follows:

$$G_s(\hat{x}) = \exp\left(-\frac{\|\hat{x} - \hat{\mu}_s\|^2}{2\hat{\sigma}_s^2}\right) \quad (1)$$

with $\hat{x} \in \mathbb{R}^3$. Note that although $G_s(\hat{x})$ has infinite support, for visualization purposes we represent its projection as a square having center (*i.e.* diagonals intersection) in $\hat{\mu}_s$ and side length equal to $2\hat{\sigma}_s$ mm (see Fig. 3).

We further assign a HSV color value η_s to each surface Gaussian. In order to derive the colors we choose a reference frame where the initial coarse reconstruction is as close as possible to the real shape. This is typically the first frame in each sequence. For each vertex v_s of the input mesh, we first choose the camera view that sees vertex v_s best, *i.e.* where normal and camera viewing direction align best. Thereafter, the 3D Gaussian associated to v_s is projected to the image from the best camera view and the underlying pixel color average is assigned as a color attribute.

4.2. 2D Image Gaussian

Our implicit model for the input images of all cameras $c \in \{0 \dots n_c - 1\}$, n_c being the number of cameras, is obtained by assigning 2D Gaussian functions $G_i(x)$, $x \in \mathbb{R}^2$, to each image patch, $i \in I(c)$, of all camera views. Similar to Stoll *et al.* [15] we decompose each input frame into squared regions of coherent color by means of quad-tree decomposition (with maximal depth set to 8). A 2D Gaussian is assigned to each patch (Fig. 4), such that its mean




Figure 3. A representation of our collection of 3D Gaussian on the surface. Our surface Gaussians, here illustrated as tiny red-bordered squares, are assigned to vertices of the input coarse mesh (only in the skirt region in this example).




Figure 4. The input image (left) and the estimated collection of 2D image Gaussians (right). The image Gaussians are assigned to patches of coherent color in the input image and the underlying average pixel color is assigned to the Gaussians as additional attribute.

$\mu_i \in \mathbb{R}^2$ corresponds to the patch center, and its standard deviation σ_i to half of the square patch side length. The underlying average HSV color η_i is also assigned to the 2D Gaussians as additional attribute.

4.3. Projection of 3D Surface Gaussians

In order to evaluate the similarity between the 3D surface Gaussians G_s and the 2D image Gaussians G_i , we project each G_s to the 2D image space. The 3D surface Gaussian mean $\hat{\mu}_s$ is projected using the camera projection matrix P , similarly to any 3D point in the space, as follows:

$$\mu_s = \begin{pmatrix} \frac{[P\hat{\mu}_s^h]_x}{[P\hat{\mu}_s^h]_z} \\ \frac{[P\hat{\mu}_s^h]_y}{[P\hat{\mu}_s^h]_z} \end{pmatrix} \in \mathbb{R}^2 \quad (2)$$

with $[P\hat{\mu}_s^h]_{x,y,z}$ being the respective coordinates of the projected mean in homogeneous coordinates (*i.e.* the 4th dimension is set to 1). The 3D standard deviation is projected using the following formula:

$$\sigma_s = \frac{\hat{\sigma}_s f}{[P\hat{\mu}_s^h]_z} \in \mathbb{R} \quad (3)$$

where f is the camera focal length.

5. Surface Refinement

We employ an analysis-by-synthesis approach to refine the input coarse mesh animation, at every frame, by optimizing the following energy $\mathbf{E}(M)$ with respect to the collection of 3D surface Gaussian means $M = \{\hat{\mu}_0, \dots, \hat{\mu}_{n_s-1}\}$:

$$\mathbf{E}(M) = E_{sim} - w_{reg} E_{reg} \quad (4)$$

The term E_{sim} measures the color similarity of the projected collection of 3D surface Gaussians with the 2D image Gaussians obtained from each camera view. The additional term E_{reg} is used to keep the distribution of the 3D surface Gaussians geometrically smooth, whereas w_{reg} is a user defined smoothness weight, typically set to 1. Since we constrain the 3D Gaussians to move along the corresponding vertex (normalized) normal direction N_s :

$$\hat{\mu}_s = \hat{\mu}_s^{orig} + N_s k_s \in \mathbb{R}^3 \quad (5)$$

aiming at maintaining a regular distribution of 3D Gaussians on the surface, we only need to optimize for single scalar values $k_s, s \in \{0 \dots n_s - 1\}$.

5.1. Similarity Term

We exploit the power of the implicit Gaussian representation of both input images and surface in order to derive a closed-form analytical formulation for our similarity term. In principle, one pair of image Gaussian and projected surface Gaussian should have high similarity measures when they show similar properties in terms of color and their spacial localization is sufficiently close. This measure can be formulated as the integral of the product of the projected surface Gaussian $G_s(x)$ and image Gaussian $G_i(x)$, weighted by their color similarity $T(\delta_{is})$, as follows:

$$E_{is} = T(\delta_{is}) \int_{\Omega} G_i(x) G_s(x) dx \quad (6)$$

In the above equation $\delta_{is} = \|\eta_i - \eta_s\| \in \mathbb{R}^+$ measures the Euclidean distance between the colors, while $T(\delta) : \mathbb{R} \rightarrow \mathbb{R}$ is the Wendland radial basis function modeled by:

$$T(\delta) = \begin{cases} \left(1 - \frac{\delta}{\Delta}\right)^4 \left(4\frac{\delta}{\Delta} + 1\right) & \text{if } \delta < \Delta \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

where Δ is experimentally set to 0.05 for all test sequences. The main advantage of using a Gaussian representation is that the integral in Eq. 6 has a closed-form solution, namely another Gaussian with combined properties:

$$E_{is} = T(\delta_{is}) 2\pi \frac{\sigma_s^2 \sigma_i^2}{\sigma_s^2 + \sigma_i^2} \exp\left(-\frac{\|\mu_i - \mu_s\|^2}{\sigma_s^2 + \sigma_i^2}\right) \quad (8)$$

We first calculate the similarity for all components of the two models for each camera view. Then, we normalize the result considering the maximum obtainable overlap $E_{ii} = \pi\sigma_i^2$, of an image Gaussian with itself, and the number of cameras n_c as follows:

$$E_{sim} = \frac{1}{n_c} \sum_{c=0}^{n_c-1} \sum_{i \in I(c)} \frac{\min\left(\sum_{s=0}^{n_s-1} E_{is}, E_{ii}\right)}{E_{ii}} \quad (9)$$

In this equation, the inner minimization implicitly handles occlusions on the surface as it prevents occluded Gaussians projections into the same image location to contribute multiple times to the energy. This is an elegant way for handling occlusion while preserving at the same time energy smoothness. In fact, exact occlusion detection and handling algorithms are non-smooth or hard to express in closed-form.

In order to improve computational efficiency, we evaluate E_{is} only for visible surface Gaussians from each camera view. The Gaussian overlap is then computed against visible projected Gaussians and 2D image Gaussians in a local neighborhood.

5.2. Regularization Term

Our regularization term constraints the 3D surface Gaussians in the local neighborhood and each Gaussian such that the final reconstructed surface is sufficiently smooth. This is accomplished by minimizing the following equation:

$$E_{reg} = \sum_{s=0}^{n_s-1} \sum_{j \in \Psi(s)} T(\delta_{sj}) (k_s - k_j)^2, \quad (10)$$

where $\Psi(s)$ is a set of surface Gaussian indices that are neighbors of G_s , $\delta_{sj} \in \mathbb{R}^+$ is the geodesic surface distance between G_s and G_j measured in number of edges, and $T(\delta)$ is defined in Eq. 7, where $\Delta = 2 \text{ edges}$.

5.3. Optimization

Our formulation allows us to compute analytic derivatives of the energy (Eq. 4), for which we provide complete derivation in an additional document. The derivative of the similarity term, with respect to each k_s , $s \in \{0 \dots n_s - 1\}$ is:

$$\frac{\partial}{\partial k_s}(E_{sim}) = \frac{1}{n_c} \sum_{c=0}^{n_c-1} \sum_{i \in I(c)} \begin{cases} \frac{\frac{\partial}{\partial k_s}(E_{is})}{E_{ii}} & \text{if } \sum_{s=0}^{n_s-1} E_{is} < E_{ii} \\ 0 & \text{otherwise} \end{cases} \quad (11)$$

The derivative of the overlap E_{is} is defined as:

$$\begin{aligned} \frac{\partial}{\partial k_s}(E_{is}) &= T(\delta_{is}) 4\pi \frac{\sigma_s^2 \sigma_i^2}{\sigma_s^2 + \sigma_i^2} \exp\left(-\frac{\|\mu_i - \mu_s\|^2}{\sigma_s^2 + \sigma_i^2}\right) \\ &\cdot \left[[P_c N_s^h]_z \left(-1 + \frac{\sigma_s^2}{\sigma_s^2 + \sigma_i^2} - \frac{\|\mu_i - \mu_s\|^2 \sigma_s^2}{(\sigma_s^2 + \sigma_i^2)^2} \right) \right. \\ &\cdot \left. \frac{1}{[\mu_s]_z} + \frac{(\mu_i - \mu_s) \frac{\partial}{\partial k_s}(\mu_s)}{\sigma_s^2 + \sigma_i^2} \right] \end{aligned} \quad (12)$$

where P_c is the projection matrix of camera c , N_s^h is the vertex normal associated to the model gaussian G_s in homogeneous coordinates (*i.e.* the 4th dimension is set to 0), $[\mu_s]_z$ is the z-component of the projected mean, and

$$\frac{\partial}{\partial k_s}(\mu_s) = \begin{pmatrix} \frac{[P N_s^h]_x - [\mu_s]_x [P N_s^h]_z}{[P(\hat{\mu}_s^h + N_s^h k_s)]_z} \\ \frac{[P N_s^h]_y - [\mu_s]_y [P N_s^h]_z}{[P(\hat{\mu}_s^h + N_s^h k_s)]_z} \end{pmatrix} \in \mathbb{R}^2. \quad (13)$$

The derivative of the regularization term E_{reg} is given by:

$$\frac{\partial}{\partial k_s}(E_{reg}) = 4 \sum_{j \in \Psi(s)} T(\delta_{sj}) (k_s - k_j) \quad (14)$$

We efficiently optimize our energy function $\mathbf{E}(M)$ using a *conditioned gradient ascent* approach. The general gradient ascent method is a first-order optimization procedure that aims at finding local maxima by taking steps proportional to the energy gradient. The conditioner is a scalar factor associated to the analytical derivatives that increases (resp. decreases) step-by-step when the gradient sign is constant (resp. fluctuating). The use of the conditioner brings three main advantages: it allows for faster convergence to the final solution, it prevents typical zig-zag-ing while approaching local maxima, and it constraints at the same time the analytical derivative size.

6. Results

We tested our approach on three different datasets: *skirt*, *dance* and *pop2lock*. Input multi-view video sequences, as well as camera settings and initial coarse mesh reconstruction were provided by Gall *et al.* [7] and Starck and Hilton [14]. All the sequences are recorded with 8 synchronized and calibrated cameras and number of frame ranging between 250 and 721 (see Table 1). The input provided coarse mesh are obtained utilizing low-quality refining technique based on sparse feature matching, shape-from-silhouette and multi-view 3D reconstruction, and therefore lack of surface details.

In order to refine the input mesh sequences, we first subdivide the input coarse topology, by inserting additional

Sequence	Frames	G_s	Iter/s	Frame/min
<i>skirt</i>	721	3053	2.01	0.8
<i>dance</i>	573	3430	1.90	0.76
<i>pop2lock</i>	250	3880	1.67	0.66

Table 1. Computation time for the input sequences. The table shows the amount of frames for each sequence, as well as amount of 3D surface Gaussians G_s , iteration per second and frames per minute.

triangles and vertices, aiming at increasing the scale level of detail. Then we generate a collection of Gaussians on the surface as explained in Sect. 3. Since for the input sequences most of the fine-scale deformations happen on the clothing, we decided to focus on the refinement of those areas, generating surface Gaussians only for the correspondent vertices. Table 1 shows the amount of 3D surface Gaussians created for each sequence.

When rendering the final resulting mesh sequences, we added an extra epsilon to the computed vertex displacements equal to the standard deviation of the surface Gaussians used. This is needed in order to compensate for the small surface bias (shrink along the normal during optimization) that is due to the spatial extent of the Gaussians.

Evaluation. Our results (Fig. 1, Fig. 5 and the accompanying video) show that our approach is able to plausibly reconstruct more fine-scale details, *e.g.* the wrinkles and folds in the skirt, and produces closer model alignment to the images than the baseline methods ([7, 14]).

In order to verify the quantitative performance of our approach, we textured the model by assigning surface Gaussians colors to the correspondent mesh vertices. Then, we used optical flow to generate displacement flow vectors between the input images and the reprojected textured mesh models (original and refined) for all time steps. Fig. 6 plots the average optical flow displacement error difference between the input and the resulting animation sequences over time for a single camera view. As shown in the graphs, our method decreases the average flow displacement error, leading to quantitatively more accurate results.

We created an additional experiment to verify the performance of our refinement framework. For this experiment, we first spatially-smooth the input mesh sequence aiming at eliminating most of the baked-in surface details, if any. The smooth mesh animation is then used as input to our system. As we show in Fig. 7 and in the accompanying video, our approach is able to plausibly refine the input smooth mesh animation, reconstructing fine-scale details in the skirt, t-shirt and shorts. Quantitative evaluation for the smooth input sequence is provided in an additional document.

We evaluated the performance of our system on an Intel Xeon Processor E5-1620, Quad-core with Hyperthreading and 16GB of RAM. Table 1 summarizes the performances



Figure 5. Comparison of the results of our refinement capture method against the baseline provided by [7, 14] for the *pop2lock* (top), *dance* (middle) and *skirt* (bottom) sequences. From left to right: input image, color-coded normals of the input mesh and color-coded normals of the rendered output refined mesh.

we obtained for the three tested sequences. We believe we can further reduce the computation time by parallelizing orthogonal steps and implementing our method on GPU.

Limitations. Our approach is subject to a few limitations. We assume the input mesh sequence to be sufficiently accurate, such that smaller details can be easily and correctly captured by simply displacing vertices along their correspondent vertex normals. In cases where the input reconstructed meshes present misalignments with respect to the images (*e.g.* *pop2lock*) or if it is necessary to reconstruct stronger deformations, then our method is unable to perform adequately. In this respect, our refinement should be reformulated allowing more complex displacements, *e.g.* without any normal constraint. However such weaker prior on vertices motion requires more complex regularization formulation in order to maintain smooth surface, also to handle unwanted self-intersections and collapsing vertices. On top of that the increased number of parameters to optimize for (*i.e.* 3 times more, when optimizing for all 3 vertices dimensions, x , y and z) would spoil computational efficiency and raise the probability of getting stuck in local maxima solutions. The risk of returning local maxima solutions is still high when employing local solvers (*e.g.* gradient ascent) on non-convex problems as in our case. A possible solution is to use more advanced solvers, *e.g.* global solvers, when computational efficiency is not a requirement.

Another limitation of our approach is the inability to




Figure 7. Results of our refinement capture method for the smoothed *dance* (top), *pop2lock* (middle) and *skirt* (bottom) animation sequences. From left to right: input image, rendered input mesh, rendered output refined mesh, zoom at the rendered input mesh, and zoom at the rendered output refined mesh.

densely refine plain colored surfaces with few texture (*e.g.* *pop2lock* and *dance*). A solution here is to employ a more complex color model that takes into account *e.g.* illumination and shading effects, at the cost of increased computational expenses. We would like to investigate these limitations as a future work.

7. Conclusions

We presented a new effective framework for performance capture of deforming meshes with fine-scale time-varying surface detail from multi-view video recordings. Our approach captures the fine-scale deformation of the mesh vertices by maximizing photo-consistency on all vertex positions. This can be done efficiently by densely optimizing a new model-to-image consistency energy function that uses our proposed implicit representation of the deformable mesh using a collection of 3D Gaussians for the surface and a set of 2D Gaussians for the input images. Our proposed formulation enables a smooth closed-form energy with implicit occlusion handling and analytic derivatives. We qualitatively and quantitatively evaluated our refinement strategy on 3 input sequences, showing that we are able to capture and model finer-scale details.

References

- [1] L. Ballan and G. M. Cortelazzo. Marker-less motion capture of skinned models in a four camera set-up using optical flow and silhouettes. In *3DPVT*, June 2008. 2
- [2] D. Bradley, T. Popa, A. Sheffer, W. Heidrich, and T. Boubekeur. Markerless garment capture. *ACM Trans. Graph.*, 27(3):19, 2008. 1, 2
- [3] C. Cagniard, E. Boyer, and S. Ilic. Free-form mesh tracking: a patch-based approach. In *Proc. IEEE CVPR*, 2010. 2
- [4] J. Carranza, C. Theobalt, M. Magnor, and H.-P. Seidel. Free-viewpoint video of human actors. In *ACM TOG (Proc. SIGGRAPH '03)*, page 22, 2003. 2
- [5] E. de Aguiar, C. Stoll, C. Theobalt, N. Ahmed, H.-P. Seidel, and S. Thrun. Performance capture from sparse multi-view video. *ACM Trans. Graph.*, 27(3), 2008. 1, 2
- [6] J. Gall, C. Stoll, E. Aguiar, C. Theobalt, B. Rosenhahn, and H.-P. Seidel. Motion capture using joint skeleton tracking and surface estimation. In *Proc. IEEE CVPR*, pages 1746–1753, 2009. 2
- [7] J. Gall, C. Stoll, E. de Aguiar, C. Theobalt, B. Rosenhahn, and H. Seidel. Motion capture using joint skeleton tracking and surface estimation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2009. 1, 2, 3, 5, 6




Figure 6. Difference between original and refined sequences in terms of the average optical flow displacement error (in pixels) per frame. We register an average improvement equal to 16.4% for the *pop2lock* sequence (top), 8.7% for the *skirt* sequence (middle) and 1.4% for the *dance* sequence (bottom).

[8] C. Hernandez, G. Vogiatzis, G. J. Brostow, B. Stenger, and

- R. Cipolla. Non-rigid photometric stereo with colored lights. In *Proc. ICCV*, pages 1–8, 2007. 2
- [9] K. Kolev, M. Klodt, T. Brox, and D. Cremers. Continuous global optimization in multiview 3d reconstruction. *International Journal of Computer Vision*, 84(1):80–96, 2009. 1, 2
- [10] Y. Liu, J. Gall, C. Stoll, Q. Dai, H.-P. Seidel, and C. Theobalt. Markerless motion capture of multiple characters using multiview image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(11):2720–2735, 2013. 2
- [11] W. Matusik, C. Buehler, R. Raskar, S. J. Gortler, and L. McMillan. Image-based visual hulls. In *SIGGRAPH*, pages 369–374, 2000. 2
- [12] P. Sand, L. McMillan, and J. Popović. Continuous capture of skin deformation. *ACM TOG*, 22(3):578–586, July 2003. 2
- [13] Y. Savoye. Iterative cage-based registration from multi-view silhouettes. In *Proceedings of the 10th European Conference on Visual Media Production, CVMP '13*, pages 8:1–8:10. ACM, 2013. 2
- [14] J. Starck and A. Hilton. Surface capture for performance-based animation. *IEEE Computer Graphics and Applications*, 27(3):2131, 2007. 1, 2, 3, 5, 6
- [15] C. Stoll, N. Hasler, J. Gall, H.-P. Seidel, and C. Theobalt. Fast articulated motion tracking using a sums of gaussians body model. In D. N. Metaxas, L. Quan, A. Sanfeliu, and L. J. V. Gool, editors, *ICCV*, pages 951–958. IEEE, 2011. 2, 3
- [16] C. Theobalt, E. de Aguiar, C. Stoll, H.-P. Seidel, and S. Thrun. Performance capture from multi-view video. In R. Ronfard and G. Taubin, editors, *Image and Geometry Processing for 3D-Cinematography*, page 127ff. Springer, 2010. 2
- [17] T. Tung, S. Nobuhara, and T. Matsuyama. Complete multi-view reconstruction of dynamic scenes from probabilistic fusion of narrow and wide baseline stereo. In *Proc. IEEE ICCV*, pages 1709–1716, 2009. 2
- [18] D. Vlasic, I. Baran, W. Matusik, and J. Popović. Articulated mesh animation from multi-view silhouettes. *ACM TOG (Proc. SIGGRAPH)*, 2008. 1, 2
- [19] D. Vlasic, P. Peers, I. Baran, P. Debevec, J. Popović, S. Rusinkiewicz, and W. Matusik. Dynamic shape capture using multi-view photometric stereo. In *ACM TOG (Proc. SIGGRAPH Asia '09)*, 2009. 2
- [20] M. Waschbüsch, S. Würmlin, D. Cötting, F. Sadlo, and M. Gross. Scalable 3D video of dynamic scenes. In *Proc. Pacific Graphics*, pages 629–638, 2005. 2
- [21] C. Wu, Y. Liu, Q. Dai, and B. Wilburn. Fusing multiview and photometric stereo for 3d reconstruction under uncalibrated illumination. *IEEE TVCG*, 17(8):1082–1095, 2011. 1, 2
- [22] C. Wu, K. Varanasi, Y. Liu, H.-P. Seidel, and C. Theobalt. Shading-based dynamic shape refinement from multi-view video under general illumination. In *Proc. ICCV, ICCV '11*, pages 1108–1115. IEEE, 2011. 1, 2
- [23] C. L. Zitnick, S. B. Kang, M. Uyttendaele, S. Winder, and R. Szeliski. High-quality video view interpolation using a layered representation. *ACM Trans. Graph.*, 23(3):600–608, Aug. 2004. 2