



NRST: Non-rigid Surface Tracking from Monocular Video

Marc Habermann¹(✉)(iD), Weipeng Xu¹(✉)(iD), Helge Rhodin²(✉)(iD),
Michael Zollhöfer³(✉)(iD), Gerard Pons-Moll¹(✉)(iD),
and Christian Theobalt¹(✉)(iD)

¹ Max Planck Institute for Informatics, 66123 Saarbrücken, Germany
{mhaberma, wxu, gpons, theobalt}@mpi-inf.mpg.de

² EPFL, 1015 Lausanne, Switzerland

helge.rhodin@epfl.ch

³ Stanford University, Stanford, CA 94305, USA

zollhoefer@cs.stanford.edu

<https://www.mpi-inf.mpg.de/home/>, <https://www.epfl.ch/>,

<https://www.stanford.edu/>

Abstract. We propose an efficient method for non-rigid surface tracking from monocular RGB videos. Given a video and a template mesh, our algorithm sequentially registers the template non-rigidly to each frame. We formulate the per-frame registration as an optimization problem that includes a novel texture term specifically tailored towards tracking objects with uniform texture but fine-scale structure, such as the regular micro-structural patterns of fabric. Our texture term exploits the orientation information in the micro-structures of the objects, e.g., the yarn patterns of fabrics. This enables us to accurately track uniformly colored materials that have these high frequency micro-structures, for which traditional photometric terms are usually less effective. The results demonstrate the effectiveness of our method on both general textured non-rigid objects and monochromatic fabrics.

1 Introduction

In this paper, we propose NRST, an efficient method for non-rigid surface tracking from monocular RGB videos. Capturing the non-rigid deformation of a dynamic surface is an important and long-standing problem in computer vision. It has a wide range of real world applications in fields such as virtual/augmented reality, medicine and visual effects. Most of the existing methods are based on multi-view imagery, where expensive and complicated system setups are required [3, 23, 25]. There also exist methods that rely on only a single depth or RGB-D camera [18, 19, 42, 44]. However, these sensors are not as ubiquitous as RGB cameras, and these methods cannot be applied on plenty of existing video

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-12939-2_23) contains supplementary material, which is available to authorized users.

footage which is found on social media like YouTube. There are also monocular RGB methods [30, 43], of course with their own limitations; e.g., they rely on highly textured surfaces and they are often times slow.




Fig. 1. We propose an efficient method for interactive non-rigid surface tracking from monocular RGB videos for general objects such as faces ((a)–(d)). Given the input image (a) our reconstruction nicely overlays with the input (b) and looks also plausible from another view point (c). The textured overlay looks realistic as well (d). Furthermore, our novel texture term leads to improved reconstruction quality for fabrics given a single video (e). Again the overlaid reconstruction (f) aligns well, and also in 3D (g) our result (red) matches the ground truth (blue). (Color figure online)

In this work, we present a method which is able to densely track the non-rigid deformations of general objects such as faces and fabrics from a single RGB video (Fig. 1). To solve this challenging problem, our method relies on a textured mesh template of the deforming object’s surface. Given the input video, our algorithm sequentially registers the template to each frame. More specifically, our method automatically reproduces a deformation sequence of the template model that coincides with the non-rigid surface motion in the video. To this end, we formulate the per-frame registration as a non-linear least squares optimization problem – with an objective function consisting of a photometric alignment and several regularization terms. The optimization is computationally intensive due to the large number of residuals in our alignment objective. To address this, we adapt the efficient GPU-based Gauss-Newton solver of Zollhoefer et al. [44] to our problem that allows for deformable object tracking at interactive frame rates.

Besides the efficiency of the algorithm, the core contribution of our approach is a novel texture term that exploits the orientation information in the micro-structures of the tracked objects, such as the yarn patterns of fabrics. This enables us to track uniformly colored materials which have high frequency patterns, for which the classical color-based term is usually less effective.

In our experimental results, we evaluate our method qualitatively and quantitatively on several challenging sequences of deforming surfaces. We use well established benchmarks, such as pieces of cloth [31, 40] and human faces [39, 43]. The results demonstrate that our method can accurately track general non-rigid objects. Furthermore, for materials with regular micro-structural patterns, such as fabrics, the tracking accuracy is further improved with our texture term.

2 Related Work

There is a variety of approaches that reconstruct geometry from multiple images, e.g., template-free methods [3], variational ones [25] or object specific approaches [23]. Although multi-view methods can produce accurate tracking results, their setup is expensive and hard to operate. Some approaches use a single RGB-D sensor instead [9, 10, 18, 19, 36, 42, 44]. They manage to capture deformable surfaces nicely and at high efficiency, some even build up a template model alongside per-frame reconstruction. The main limitations of these methods are that the sensors have a high power consumption, they do not work outdoors, the object has to be close to the camera and they cannot use the large amount of RGB-only video footage provided by social media. On these grounds, we aim for a method that uses just a single RGB video as input. In the following, we focus on related monocular reconstruction and tracking approaches.

Monocular Methods. Non-rigid structure from motion methods, which do not rely on any template, try to infer the 3D geometry from a single video by using a prior-free formulation [4], global models [37], local ones [27] or solving a variational formulation [6]. But they often either capture the deformations only coarsely, are not able to model strong deformations, typically require strongly textured objects or rely on dense 2D correspondences. By constraining the setting to specific types of objects such as faces [7], very accurate reconstructions can be obtained, but at the expense of generality. Since in recent years, several approaches [11, 22] build a 3D model given a set of images, and even commercial software¹ is available for this task, template acquisition has become easier. Templates are an effective prior for the challenging task of estimating non-rigid deformations from single images as demonstrated by previous work [1, 2, 14–17, 21, 24, 28–33, 40, 43]. But even if a template is used, ambiguities [30] remain and additional constraints have to be imposed. Theoretical results [1] show that only allowing isometric deformations [24] results in a uniquely defined solution. Therefore, approaches constrain the deformation space in several ways, e.g., by a Laplacian regularization [21] or by non-linear [32] or linear local surface models [29]. Salzmann et al. [28] argued that relaxing the isometric constraint is beneficial since it allows to model sharp folds. Moreno-Noguer et al. [17] and Malti et al. [16] even go beyond this and show results for elastic surfaces; Tsoli and Argyros [38] demonstrated tracking surfaces that undergo topological changes but require a depth camera. Other approaches investigate how to make reconstruction more robust under faster motions [33] and occlusions [20], or try to replace the feature-based data term by a dense pixel-based one [15] and to find better texture descriptors [8, 12, 26]. Brunet et al. [2] and Yu et al. [43] formulate the problem of estimating non-rigid deformations as minimizing an objective function which brings them closest to our formulation. In particular, we adopt the photometric, spatial and temporal terms of Yu et al. [43] and combine them with an isometric and acceleration constraint as well as our novel texture term.

¹ <http://www.agisoft.com/>.

Along the line of monocular methods, we propose NRST, a template-based reconstruction framework that estimates the non-rigidly deforming geometry of general objects from just monocular video. In contrast to previous work, our approach does not rely on 3D to 2D correspondences and due to the GPU-based solver architecture it is also much faster than previous approaches. Furthermore, our novel texture term enables tracking of regions with little texture.

3 Method

The goal is to estimate the non-rigid deformation of an object from T frames $I^t(x, y)$ with $t \in \{1, \dots, T\}$. We assume a static camera and known camera intrinsics. Since this problem is in general severely under-constrained, it is assumed that a template triangle mesh of the object to be tracked is given as the matrix $\hat{\mathbf{V}} \in \mathbb{R}^{N \times 3}$ where each row contains the coordinates of one of the N vertices. According to that, $\hat{\mathbf{V}}_i$ is defined as the i th vertex of the template in vector form. This notation is also used for the following matrices. The edges of the template are given as the mapping $\mathcal{N}(i)$. Given a vertex index $i \in \{1, 2, \dots, N\}$, it returns the set of indices sharing an edge with $\hat{\mathbf{V}}_i$. The F faces of the mesh are represented as the matrix $\mathbf{F} \in \{1, \dots, N\}^{F \times 3}$. Each row contains the vertex indices of one triangle. The UV map is given as the matrix $\mathbf{U} \in \mathbb{N}^{N \times 2}$. Each row contains the UV coordinates for the corresponding vertex. The color $\mathbf{C}_i \in \{0, \dots, 255\}^3$ of vertex i can be computed by a simple lookup in the texture map I_{TM} at the position \mathbf{U}_i . The color of all vertices is stored in the matrix $\mathbf{C} \in \{0, \dots, 255\}^{N \times 3}$. Furthermore, it is assumed that the geometry at time $t = 1$ roughly agrees with the true shape shown in the video so that the gradients of the photometric term can guide the optimization to the correct solution without being trapped into local minima. The non-rigidly deformed mesh at time $t + 1$ is represented as the matrix $\mathbf{V}^{t+1} \in \mathbb{R}^{N \times 3}$ and contains the updated vertex positions according to the 3D displacement from t to $t + 1$.

3.1 Non-rigid Tracking as Energy Minimization

Given the template $\hat{\mathbf{V}}$ and our estimate of the previous frame \mathbf{V}^t , our method sequentially estimates the geometry \mathbf{V}^{t+1} of the current frame $t + 1$. We jointly optimize per-vertex local rotations denoted by Φ^{t+1} and vertex locations \mathbf{V}^{t+1} . Specifically, for each time step the deformation estimation is formulated as the non-linear optimization problem

$$(\mathbf{V}^{t+1}, \Phi^{t+1}) = \arg \min_{\mathbf{V}, \Phi \in \mathbb{R}^{N \times 3}} E(\mathbf{V}, \Phi), \quad (1)$$

with

$$E(\mathbf{V}, \Phi) = \lambda_{\text{Photo}} E_{\text{Photo}}(\mathbf{V}) + \lambda_{\text{Smooth}} E_{\text{Smooth}}(\mathbf{V}) + \lambda_{\text{Edge}} E_{\text{Edge}}(\mathbf{V}) + \lambda_{\text{Arap}} E_{\text{Arap}}(\mathbf{V}, \Phi) + \lambda_{\text{Vel}} E_{\text{Vel}}(\mathbf{V}) + \lambda_{\text{Acc}} E_{\text{Acc}}(\mathbf{V}). \quad (2)$$

λ_{Photo} , λ_{Smooth} , λ_{Edge} , λ_{Vel} , λ_{Acc} , λ_{Arap} are hyperparameters set before the optimization starts and afterwards they are kept constant. $E(\mathbf{V}, \Phi)$ combines different cost terms ensuring that the mesh deformations agree with the motion in the video. The resulting non-linear least squares optimization problem is solved with the GPU-based Gauss-Newton solver based on the work of Zollhoefer et al. [44] where we adapted the Jacobian and residual implementation to our energy formulation. The high efficiency is obtained by exploiting the sparse structure of the system of normal equations. For more details we refer the reader to the approach of Zollhoefer et al. [44]. Now, we will explain the terms in more detail.

Photometric Alignment. The photometric term

$$E_{Photo}(\mathbf{V}) = \sum_{i=1}^N \left\| \sigma(I^{t+1} * \mathbf{G}_w(\Pi(\mathbf{V}_i)) - \mathbf{C}_i) \right\|^2 \quad (3)$$

densely measures the re-projection error. $\|\cdot\|$ is the Euclidean norm, $*$ is the convolution operator and \mathbf{G}_w is a Gaussian kernel with standard deviation w . We use Gaussian smoothing on the input frame for more stable and longer range gradients. $\Pi(\mathbf{V}_i) = \left(\frac{u}{w}, \frac{v}{w}\right)^\top$ with $(u, v, w)^\top = \mathbf{I}\mathbf{V}_i$ projects the vertex \mathbf{V}_i on the image plane and $I^{t+1} * \mathbf{G}_w$ returns the RGB color vector of the smoothed frame at position $\Pi(\mathbf{V}_i)$ which is compared against the pre-computed and constant vertex color \mathbf{C}_i . Here, $\mathbf{I} \in \mathbb{R}^{3 \times 3}$ is the intrinsic camera matrix. σ is a robust pruning function for wrong correspondences with respect to color similarity. More specifically, we discard errors above a certain threshold because in most cases they are due to occlusions.

Spatial Smoothness. Without regularization, estimating 3D geometry from a single image is an ill-posed problem. Therefore, we introduce several spatial and temporal regularizers to make the problem well-posed and to propagate 3D deformations into areas where information for data terms is missing, e.g., poorly textured or occluded regions. The first prior

$$E_{Smooth}(\mathbf{V}) = \sum_{i=1}^N \sum_{j \in \mathcal{N}(i)} \left\| (\mathbf{V}_i - \mathbf{V}_j) - (\hat{\mathbf{V}}_i - \hat{\mathbf{V}}_j) \right\|^2 \quad (4)$$

ensures that if a vertex \mathbf{V}_i changes its position, its neighbors \mathbf{V}_j with $j \in \mathcal{N}(i)$ are deformed such that the overall shape is still spatially smooth compared to the template mesh $\hat{\mathbf{V}}$. In addition, the prior

$$E_{Edge}(\mathbf{V}) = \sum_{i=1}^N \sum_{j \in \mathcal{N}(i)} \left(\|\mathbf{V}_i - \mathbf{V}_j\| - \|\hat{\mathbf{V}}_i - \hat{\mathbf{V}}_j\| \right)^2 \quad (5)$$

ensures isometric deformations which means that the edge length with respect to the template is preserved. In contrast to E_{Smooth} , this prior is rotation invariant.

Finally, the as-rigid-as-possible (ARAP) prior [34]

$$E_{\text{Arap}}(\mathbf{V}, \Phi) = \sum_{i=1}^N \sum_{j \in \mathcal{N}(i)} \left\| (\mathbf{V}_i - \mathbf{V}_j) - \mathcal{R}(\Phi_i)(\hat{\mathbf{V}}_i - \hat{\mathbf{V}}_j) \right\|^2, \quad (6)$$

allows local rotations for each of the mesh vertices as long as the relative position with respect to their neighborhood remains the same. Each row of the matrix $\Phi \in \mathbb{R}^{N \times 3}$ contains the per-vertex Euler angles which encode a local rotation around \mathbf{V}_i . $\mathcal{R}(\Phi_i)$ converts them into a rotation matrix.

We choose a combination of spatial regularizers to ensure that our method can track different types of non-rigid deformations equally well. For example, E_{Smooth} is usually sufficient to track facial expressions without large head rotations. But tracking rotating objects can only be achieved with rotational invariant regularizers (E_{Edge} , E_{Arap}). In contrast to Yu et al. [43], we adopt the Euclidean norm in Eqs. 4 and 6 instead of the Huber loss because it led to visually better results.

Temporal Smoothness. To enforce temporally smooth reconstructions, we propose two additional priors. The first one is defined as

$$E_{\text{Velocity}}(\mathbf{V}) = \sum_{i=1}^N \left\| \mathbf{V}_i - \mathbf{V}_i^t \right\|^2 \quad (7)$$

and ensures that the displacement of vertices between t and $t+1$ is small. Second, the prior

$$E_{\text{Acc}}(\mathbf{V}) = \sum_{i=1}^N \left\| (\mathbf{V}_i - \mathbf{V}_i^t) - (\mathbf{V}_i^t - \mathbf{V}_i^{t-1}) \right\|^2 \quad (8)$$

penalizes large deviations of the current velocity direction from the previous one.

3.2 Non-rigid Tracking of Woven Fabrics

Tracking of uniformly colored fabrics is usually challenging for classical color-based terms due to the lack of color features. To overcome this limitation, we inspected the structure of different garments and found that most of them show line-like micro-structures due to the manufacturing process of the woven threads, see Fig. 2 left. Those can be recorded with recent high resolution cameras such that reconstruction algorithms can make use of those patterns. To this end, we propose a novel texture term to refine the estimation of non-rigid motions for the case of woven fabrics. It can be combined with the terms in Eq. 2. Now, we will explain our novel data term in more detail.

Histogram of Oriented Gradient (HOG). Based on HOG [5] we compute for each pixel (i, j) of an image the corresponding histogram $\mathbf{h}_{i,j} \in \mathbb{R}^b$ where b is the number of bins that count the total number of gradient angles present in the neighborhood of pixel (i, j) . To be more robust with respect to outliers




Fig. 2. Histogram of oriented gradients of woven fabrics. Left. The neighborhood region with the center pixel in the middle. Right. The corresponding histogram.

and noise we count the number of gradients per angular bin irrespective of the gradient magnitude and only if the magnitude is higher than a certain threshold. Compared to pure image gradients, HOG is less sensitive to noise. Especially for woven fabrics, image gradients are very localized since changing the position in the image can lead to large differences in the gradient directions due to the high frequency of the image content. HOG instead averages over a certain window so that outliers are discarded.

Directions of a Texture. Applying HOG to pictures of fabrics reveals their special characteristics caused by the line like patterns (see Fig. 2). There are two dominant texture gradient angles α and $\beta = ((\alpha + 180) \bmod 360)$ perpendicular to the lines. So α provides the most characteristic information of the pattern in the image at (i, j) and can be computed as the angle whose bin has the highest frequency in $\mathbf{h}_{i,j}$. α is then converted to its normalized 2D direction, also called *dominant frame gradient* (DFG), which is stored in the two-valued image $I_{\text{Dir}}(i, j)$. To detect image regions that do not contain line patterns, we set $I_{\text{Dir}}(i, j) = (0, 0)^\top$ if the highest frequency is below a certain threshold.

Texture-based Constraint. Our novel texture term

$$E_{\text{Tex}}(\mathbf{V}) = \sum_{i=1}^F \|\rho(\mathbf{d}_{\text{M},i}, \mathbf{d}_{\text{F},i})\|^2 \quad (9)$$

ensures now that for all triangles the projected DFG $\mathbf{d}_{\text{M},i}$ parametrized on the object surface agrees with the frame’s DFG $\mathbf{d}_{\text{F},i}$ at the location of the projected triangle center. An overview is shown in Fig. 3. More precisely, by averaging $\mathbf{U}_k, \mathbf{U}_m, \mathbf{U}_l$ one can compute the pixel position $\mathbf{z}_{\text{TM},i} \in \mathbb{R}^2$ of the center point of the triangle $\mathbf{F}_i = (k, m, l)$ in the texture map. Now, the neighborhood region for HOG around $\mathbf{z}_{\text{TM},i}$ is defined as the 2D bounding box of the triangle. The HOG descriptor for $\mathbf{z}_{\text{TM},i}$ can be computed and by applying the concept explained in the previous paragraph one obtains the DFG $\mathbf{d}_{\text{TM},i}$ (see Fig. 3(a)). Next, we define $\mathbf{b}_{\text{TM},i} = \mathbf{z}_{\text{TM},i} + \mathbf{d}_{\text{TM},i}$ and express it as a linear combination of the triangles’ UV coordinates leading to the *barycentric coordinates* $\mathbf{B}_{i,1}, \mathbf{B}_{i,2}, \mathbf{B}_{i,3}$ of the face \mathbf{F}_i . They form together with the other triangles the barycentric

coordinates matrix $\mathbf{B} \in \mathbb{R}^{F \times 3}$. Each row represents the texture map’s DFG for the respective triangle of the mesh in an implicit form. Since $\mathbf{b}_{\text{TM},i}$ can be represented as a linear combination, one can compute the corresponding 3D point $\mathbf{b}_{3\text{D},i} = \mathbf{B}_{i,1} \mathbf{V}_k + \mathbf{B}_{i,2} \mathbf{V}_m + \mathbf{B}_{i,3} \mathbf{V}_l$ as well as the triangle center $\mathbf{z}_{3\text{D},i} \in \mathbb{R}^3$ in 3D (see Fig. 3(b)). The barycentric coordinates remain constant, so that $\mathbf{b}_{3\text{D},i}$ and $\mathbf{z}_{3\text{D},i}$ only depend on the mesh vertices \mathbf{V}_k , \mathbf{V}_m and \mathbf{V}_l . One can then project the DFG of the mesh $\mathbf{d}_{\text{M},i} = \Pi(\mathbf{b}_{3\text{D},i}) - \Pi(\mathbf{z}_{3\text{D},i})$ into the frame and compare it against the DFG $\mathbf{d}_{\text{F},i}$ of the frame $t + 1$ at the location of $\Pi(\mathbf{z}_{3\text{D},i})$ which can be retrieved by an image lookup in I_{Dir}^{t+1} (see Fig. 3(c)). $\rho(\mathbf{x}, \mathbf{y})$ computes the minimum of the differences between \mathbf{x}, \mathbf{y} and $\mathbf{x}, -\mathbf{y}$ iff both \mathbf{x} and \mathbf{y} are non-zero vectors (otherwise we are not in an area with line patterns) and the directions are similar up to a certain threshold to be more robust with respect to occlusions and noise. As mentioned above, there are two DFGs in the frame for the case of line patterns. We assume the initialization is close to the ground truth and choose the minimum of the two possible directions.




Fig. 3. Overview of the proposed texture term.

4 Results

All experiments were performed on a PC with an NVIDIA GeForce GTX 1080Ti and an Intel Core i7. In contrast to related methods [43], we achieve interactive frame rates using the energy proposed in Eq. 2.

4.1 Qualitative and Quantitative Results

Now, we evaluate NRST on datasets for general objects like faces where we disable E_{Tex} . After that, we compare our approach against another monocular method. Finally, we evaluate our proposed texture term on two new scenes showing line-like fabric structures, perform an ablation study and demonstrate interesting applications. More results can be found in the supplemental video.

Qualitative Evaluation for General Objects. In Fig. 4 we show frames from our monocular reconstruction results. We tested our approach on two face sequences [39, 43] where templates are provided. Note that NRST precisely reconstructs facial expressions. The 2D overlay (second column) matches the input

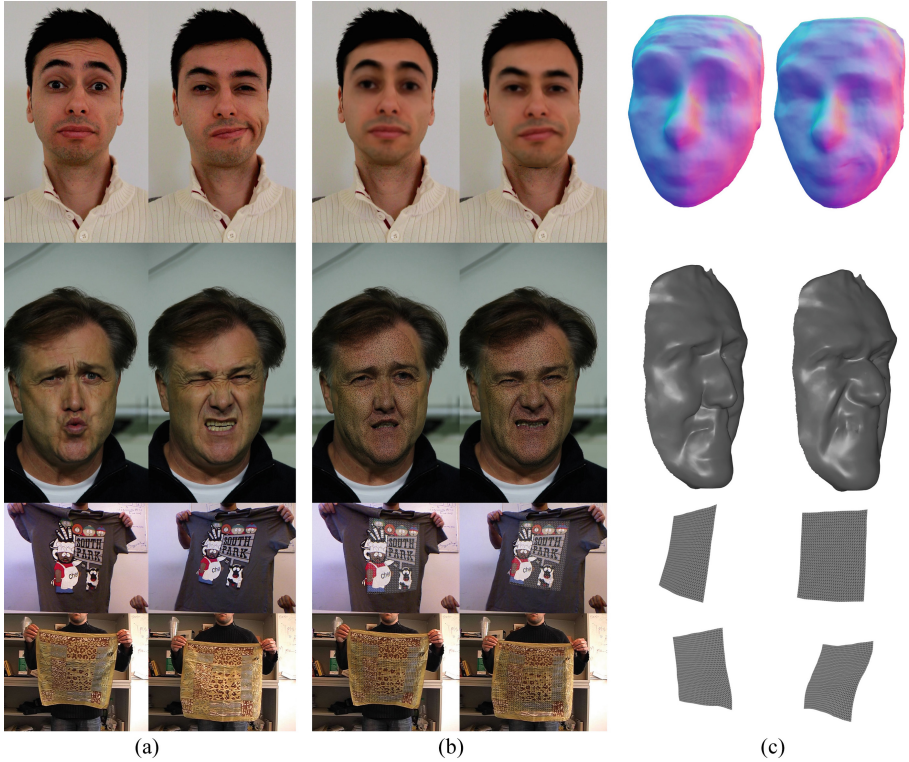


Fig. 4. Reconstructions of existing datasets [31, 39, 40, 43]. Input frames (a). Textured reconstructions overlaid on the input (b). Deformed geometries obtained by our method (rendered from a different view point) (c).

and also in 3D (third column) our results look realistic. Furthermore, we evaluated on the datasets of Varol et al. [40] and Salzmann et al. [31] showing fast movements of a T-shirt and a waving towel. Again for most parts of the surface the reconstructions look accurate in 2D since they overlap well with the input and they are also plausible in 3D. This validates that our approach can deal with the challenging problem of estimating 3D deformations from a monocular video for general kinds of objects.

Comparison to Yu et al. [43]. Figure 5 shows a qualitative comparison between our method and the one of Yu et al. [43]. It becomes obvious that both capture the facial expression, but the proposed approach is faster than the one of Yu et al. due to our data-parallel GPU implementation. In particular, on their sequence our method runs at 15 fps whereas their approach takes several seconds per frame. More side-by-side comparisons on this sequence can be found in the supplemental video.



Fig. 5. Comparison of NRST’s reconstruction (right) and the one of Yu et al. [43] (middle). It becomes obvious that both capture the facial expression shown in the input frame (left), but the proposed approach is significantly faster than the one of Yu et al. due to our data-parallel GPU implementation.

Qualitative Evaluation for Fabrics. The top row of Fig. 6 shows frames (resolution 1920×1080) of a moving piece of cloth that has the typical line patterns. Although the object is sparsely textured, our approach is able to recover the deformations due to the texture term, which accurately tracks the DFG of the line pattern. As demonstrated in the last column, the estimated angles for the frames are correct and therefore give a reliable information cue exploited by E_{Tex} . For quantitative evaluation, we created a synthetic scene that is modeled and animated in a modeling software showing a carpet that has the characteristic line pattern but is also partially textured (see bottom row of Fig. 6). We rendered the scene at a resolution of 1500×1500 . E_{Tex} helps in the less textured regions where E_{Photo} would fail. The last column shows how close our reconstruction (red) is with respect to ground truth (blue).

Ablation Analysis. Apart from the proposed texture term, our energy formulation is similar to the one of Yu et al. [43]. To validate that E_{Tex} improves the reconstruction over a photometric-only formulation, we perform an ablation study. We measured the averaged per-vertex Euclidean distance between the ground truth mesh and our reconstructions. For the waving towel shown in Fig. 6 bottom, we obtained an error of 26.8 mm without E_{Tex} and 25.5 mm if we also use our proposed texture term leading to an improvement of 4.8%. The diagonal of the 3D bounding box of the towel is 3162 mm. For the rotation sequence (resolution 800×800) shown in Fig. 7 the color variation is very limited since background and object have the same color. In contrast to E_{Photo} alone, E_{Tex} can rotate the object leading to an error of 4.1 mm for the texture-only case and 6.7 mm for the photometric-only setting. So E_{Tex} improves over E_{Photo} by 38.8%.

4.2 Applications

Our method enables several applications such as free view point rendering or re-texturing on general deformable objects or for virtual face make-up (see Fig. 8).




Fig. 6. Reconstruction of line patterns. Top from left to right. Input frames. Textured reconstructions overlaid on the frames. Color coded visualization of the estimated dominant frame angles. Regions where no line pattern was detected are visualized in black. Bottom from left to right. Input frames. Textured reconstructions overlaid on the frames. Ground truth geometries (blue) and our reconstructions (red). (Color figure online)

Since our approach estimates the deforming geometry, one can even change the scene lighting for the foreground such that the shading remains realistic.

4.3 Limitations

By the nature of the challenging task of monocular tracking of non-rigid deformations, our method has some limitations which open up directions for future work. Although, our proposed texture term uses more of the information contained in the video than a photometric-only formulation, there are still image cues that can improve the reconstruction like shading and the object contour as demonstrated by previous work [13, 41]. So, one could combine them in a unified framework. To increase robustness, the deformations could be jointly estimated




Fig. 7. Rotating object sequence. From left to right. First and last frame. Note that the object and the background have the same color. The reconstructions (red) of the last frame with either E_{Photo} or E_{Tex} overlayed on the ground truth geometry (blue). Note that E_{Tex} can recover the rotation in contrast to E_{Photo} . (Color figure online)



Fig. 8. Applications. Left. Re-textured shirt. Right. Re-textured and re-lighted face.

over a temporal sliding window as proposed by Xu et al. [41] and an embedded graph [35] could lead to improved stability by reducing the number of unknowns.

5 Conclusion

We presented an optimization-based analysis-by-synthesis method that solves the challenging task of estimating non-rigid motion, given a single RGB video and a template. Our method tracks non-trivial deformations of a broad class of shapes, ranging from faces to deforming fabric. Further, we introduce specific solutions tailored to capture woven fabrics, even if they lack clear color variations. Our method runs at interactive frame rates due to the GPU-based solver that can efficiently solve the non-linear least squares optimization problem. Our evaluation shows that the reconstructions are accurate in 2D and 3D which enables several applications such as re-texturing.

References

1. Bartoli, A., Gérard, Y., Chadebecq, F., Collins, T.: On template-based reconstruction from a single view: analytical solutions and proofs of well-posedness for developable, isometric and conformal surfaces. In: CVPR (2012)

2. Brunet, F., Hartley, R., Bartoli, A., Navab, N., Malgouyres, R.: Monocular template-based reconstruction of smooth and inextensible surfaces. In: Kimmel, R., Klette, R., Sugimoto, A. (eds.) ACCV 2010. LNCS, vol. 6494, pp. 52–66. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-19318-7_5
3. Carceroni, R.L., Kutulakos, K.N.: Multi-view scene capture by surfel sampling: from video streams to non-rigid 3D motion, shape & reflectance. In: ICCV (2001)
4. Dai, Y., Li, H., He, M.: A simple prior-free method for non-rigid structure-from-motion factorization. IJCV **107**, 101–122 (2014)
5. Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: CVPR (2005)
6. Garg, R., Roussos, A., Agapito, L.: Dense variational reconstruction of non-rigid surfaces from monocular video. In: CVPR (2013)
7. Garrido, P., Valgaerts, L., Wu, C., Theobalt, C.: Reconstructing detailed dynamic face geometry from monocular video. TOG **32**, 158 (2013)
8. Gårding, J.: Shape from texture for smooth curved surfaces in perspective projection. J. Math. Imaging Vis. **2**, 327–350 (1992)
9. Jordt, A., Koch, R.: Fast tracking of deformable objects in depth and colour video. In: BMVC, pp. 1–11 (2011)
10. Jordt, A., Koch, R.: Direct model-based tracking of 3D object deformations in depth and color video. IJCV **102**(1–3), 239–255 (2013)
11. Labatut, P., Pons, J.P., Keriven, R.: Efficient multi-view reconstruction of large-scale scenes using interest points, delaunay triangulation and graph cuts. In: ICCV (2007)
12. Liang, J., DeMenthon, D., Doermann, D.: Flattening curved documents in images. In: CVPR (2005)
13. Liu-Yin, Q., Yu, R., Agapito, L., Fitzgibbon, A., Russell, C.: Better together: joint reasoning for non-rigid 3D reconstruction with specularities and shading. In: BMVC (2016)
14. Ma, W.J.: Nonrigid 3D reconstruction from a single image. In: ISAI (2016)
15. Malti, A., Bartoli, A., Collins, T.: A pixel-based approach to template-based monocular 3D reconstruction of deformable surfaces. In: ICCV Workshops (2011)
16. Malti, A., Hartley, R., Bartoli, A., Kim, J.H.: Monocular template-based 3D reconstruction of extensible surfaces with local linear elasticity. In: CVPR (2013)
17. Moreno-Noguer, F., Salzmann, M., Lepetit, V., Fua, P.: Capturing 3D stretchable surfaces from single images in closed form. In: CVPR (2009)
18. Newcombe, R.A., Fox, D., Seitz, S.M.: DynamicFusion: reconstruction and tracking of non-rigid scenes in real-time. In: CVPR (2015)
19. Newcombe, R.A., et al.: KinectFusion: real-time dense surface mapping and tracking. In: International Symposium on Mixed and Augmented Reality (2011)
20. Ngo, D.T., Park, S., Jorstad, A., Crivellaro, A., Yoo, C.D., Fua, P.: Dense image registration and deformable surface reconstruction in presence of occlusions and minimal texture. In: ICCV (2015)
21. Östlund, J., Varol, A., Ngo, D.T., Fua, P.: Laplacian meshes for monocular 3D shape recovery. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) ECCV 2012. LNCS, vol. 7574, pp. 412–425. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-33712-3_30
22. Pan, Q., Reitmayr, G., Drummond, T.: ProFORMA: probabilistic feature-based on-line rapid model acquisition. In: BMVC (2009)
23. Perriollat, M., Bartoli, A.: A quasi-minimal model for paper-like surfaces. In: CVPR (2007)

24. Perriollat, M., Hartley, R., Bartoli, A.: Monocular template-based reconstruction of inextensible surfaces. *IJCV* **95**, 124–137 (2011)
25. Pons, J.P., Keriven, R., Faugeras, O.: Modelling dynamic scenes by registering multi-view image sequence. In: *CVPR* (2005)
26. Rao, A.R.: Computing oriented texture fields. *Comput. Vis. Graph. Image Process.: Graph. Models Image Process.* **53**, 157–185 (1991)
27. Russell, C., Fayad, J., Agapito, L.: Energy based multiple model fitting for non-rigid structure from motion. In: *CVPR* (2011)
28. Salzmann, M., Fua, P.: Reconstructing sharply folding surfaces: a convex formulation. In: *CVPR* (2009)
29. Salzmann, M., Fua, P.: Linear local models for monocular reconstruction of deformable surface. *Trans. Pattern Anal. Mach. Intell.* **33**, 931–944 (2011)
30. Salzmann, M., Lepetit, V., Fua, P.: Deformable surface tracking ambiguities. In: *CVPR* (2007)
31. Salzmann, M., Moreno-Noguer, F., Lepetit, V., Fua, P.: Closed-form solution to non-rigid 3D surface registration. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) *ECCV 2008*. LNCS, vol. 5305, pp. 581–594. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-88693-8_43
32. Salzmann, M., Urtasun, R., Fua, P.: Local deformation models for monocular 3D shape recovery. In: *CVPR* (2008)
33. Shen, S., Shi, W., Liu, Y.: Monocular template-based tracking of inextensible deformable surfaces under L_2 -norm. In: Zha, H., Taniguchi, R., Maybank, S. (eds.) *ACCV 2009*. LNCS, vol. 5995, pp. 214–223. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-12304-7_21
34. Sorkine, O., Alexa, M.: As-rigid-as-possible surface modeling. In: *SGP* (2007)
35. Sumner, R.W., Schmid, J., Pauly, M.: Embedded deformation for shape manipulation. *ACM Trans. Graph.* **26**(3) (2007). Article no. 80. <https://doi.org/10.1145/1276377.1276478>. ISSN: 0730-0301
36. Tao, Y., et al.: DoubleFusion: real-time capture of human performance with inner body shape from a depth sensor. In: *IEEE Conference on Computer Vision and Pattern Recognition* (2018)
37. Torresani, L., Hertzmann, A., Bregler, C.: Non-rigid structure-from-motion: estimating shape and motion with hierarchical priors. *Trans. Pattern Anal. Mach. Intell.* **30**, 878–892 (2008)
38. Tsoli, A., Argyros, A.: Tracking deformable surfaces that undergo topological changes using an RGB-D camera. In: *3DV*, October 2016
39. Valgaerts, L., Wu, C., Bruhn, A., Seidel, H.P., Theobalt, C.: Lightweight binocular facial performance capture under uncontrolled lighting. *SIGGRAPH Asia* (2012)
40. Varol, A., Salzmann, M., Fua, P., Urtasun, R.: A constrained latent variable model. In: *CVPR* (2012)
41. Xu, W., et al.: MonoPerfCap: human performance capture from monocular video. *TOG* **37**, 27 (2018)
42. Xu, W., Salzmann, M., Wang, Y., Liu, Y.: Nonrigid surface registration and completion from RGBD images. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) *ECCV 2014*. LNCS, vol. 8690, pp. 64–79. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10605-2_5
43. Yu, R., Russell, C., Campbell, N.D.F., Agapito, L.: Direct, dense, and deformable: template-based non-rigid 3D reconstruction from RGB video. In: *ICCV* (2015)
44. Zollhoefer, M., et al.: Real-time non-rigid reconstruction using an RGB-D camera. *TOG* **33**, 156 (2014)