

Scalable Dense Monocular Surface Reconstruction

Mohammad Dawud Ansari¹

md.dawud.ansari@dfki.de

Vladislav Golyanik^{1,2}

vladislav.golyanik@dfki.de

Didier Stricker^{1,2}

didier.stricker@dfki.de

¹ Department of Augmented Vision, DFKI ² University of Kaiserslautern

Abstract

This paper reports on a novel template-free monocular non-rigid surface reconstruction approach. Existing techniques using motion and deformation cues rely on multiple prior assumptions, are often computationally expensive and do not perform equally well across the variety of data sets. In contrast, the proposed Scalable Monocular Surface Reconstruction (SMSR) combines strengths of several algorithms, i.e., it is scalable with the number of points, can handle sparse and dense settings as well as different types of motions and deformations. We estimate camera pose by singular value thresholding and proximal gradient. Our formulation adopts alternating direction method of multipliers which converges in linear time for large point track matrices. We integrate trajectory space constraints by the measurement matrix smoothing and not for the whole optimisation. In the extensive experiments, SMSR is demonstrated to consistently achieve state-of-the-art accuracy on a wide variety of data sets.

1. Introduction

The class of methods aiming at 3D surface reconstruction by factorising given 2D point tracks across monocular views is known as non-rigid structure from motion (NRSfM). Monocular surface recovery has been tackled for perspective and orthographic cases. We assume an uncalibrated set of views and adopt orthographic camera model in this paper. Unlike its rigid counterpart [32], NRSfM is a highly ill-posed problem with several inherent ambiguities such as basis, global rotation and reflection ambiguities. Lots of research has been done in this field so far, and it remains very active [35, 33, 3, 28, 19, 12, 23].

One difficulty in NRSfM arises due to the basis ambiguity [35]. Most of the current approaches use additional constraints to handle this problem. Many methods assume the motion of the camera to be small and smooth [33, 3, 28, 19, 12] which limits their applicability in practice. Another assumption found in the literature is smooth-

ness prior on point trajectories [3] which, however, often results in oversmoothing of sudden structure changes. Besides, only several methods can handle dense sequences so far [29, 12, 15, 1, 16]. and only a few data sets for quantitative evaluation of NRSfM methods are currently available.

With these limitations in mind, our motivation is to find a scalable NRSfM method, i.e., a method supporting point track matrices of different sizes, to make NRSfM robust to self-occlusions and easy to implement. Among existing approaches, it is hard to find one which fulfils all the properties stated above. Accordingly, our *contributions* in this paper can be summarised as follows:

- We design an NRSfM approach which is easy to implement, computationally efficient, can potentially run interactively and can be used for a variety of realistic data sets. We call the new method Scalable Monocular Surface Reconstruction, or SMSR (Sec. 3). In our method, optimal basis selection is performed automatically which reduces the overall number of unknowns. Moreover, a non-convex formulation for camera pose estimation is solved iteratively using singular value thresholding (SVT) and proximal gradient. Surface estimation is formulated as an optimisation problem where guaranteed convergence to global minima is achieved using alternating direction method of multipliers (ADMM) [25, 7].
- A novel pre-processing step for input matrix is proposed, i.e., the measurement matrix is reinitialised with smooth trajectory constraints. This provides consistently accurate shape recovery for a wide variety of data sets, both sparse and dense (Sec. 3.2).
- Due to lack of dense data sets with ground truth for quantitative evaluation, we propose a two-step scheme. First, 3D shapes are estimated using an accurate and computationally expensive Variational Approach (VA) [12] followed by applying a rotation pattern in the second step (Sec. 4.1). Eventually, we found realistic data sets which allow to emphasise strengths and advantages of the new method, whereas some competing methods perform less accurate on them.

- We perform an extensive evaluation of SMSR on a variety of sparse and dense data sets, both synthetic and real ones (Sec. 4). Across various evaluation scenarios, SMSR achieves *consistent state-of-the-art level accuracy* compared to the competing approaches.

2. Related Work

The first NRSfM approach based on the Low-Rank Shape Model (LRSF) was proposed by Bregler *et al.* [8] as an extension of the Tomasi-Kanade factorisation based rigid structure from motion [32]. In LRSF, every non-rigid shape is represented by a linear combination of unknown basis shapes. The basis ambiguity has plagued early NRSfM approaches. Despite the basis and the coefficients are ambiguous, it was proven that the shapes can be, nevertheless, reconstructed uniquely [2]. Dai *et al.* [10] proposed a method with a minimal number of prior assumptions but the performance of their method considerably varies with a dataset. Methods adopting trajectory space model (TSM) allowed to increase stability and accuracy of monocular surface recovery due to a trajectory basis defined a priori (*e.g.*, discrete cosine transform or DCT basis) [4, 19, 21]. A pre-defined basis is, at the same time a disadvantage, and the performance may likewise vary considerably depending on data. The first approach that used constraints in both metric and trajectory space was [19]. Russell *et al.* [29] demonstrated dense monocular surface recovery for the first time. The authors showed how existing piecewise NRSfM techniques could be improved for scalability with the number of points. Several methods allow determining an optimal basis automatically [12, 24, 15]. Garg *et al.* [12] proposed VA with a total variation (TV) spatial regulariser. VA achieves high accuracy for multiple data sets but strongly depends on initialisation and requires an additional connectivity data structure thereby increasing the number of inputs.

A class of methods aiming at the reconstruction of scenes with large deformations adopts the idea of piecewise processing. One of the recent methods by Lee *et al.* [23] achieves state-of-the-art accuracy on several sparse data sets. However, it scales poorly with the number of points and involves several computationally expensive steps. In [10], non-convex energy minimisation is performed directly by relaxing the matrix rank constraint to a nuclear norm using semidefinite programming solver.

Our method is an algorithmic mixture of several concepts encountered in several approaches. We believe that dense NRSfM methods should be well scalable with the number of points, *i.e.*, perform equally well and stable on a broad range of data sets, with small and large deformations. Thus, we adopt constraints in the metric and trajectory spaces and our formulation is most closely related to [23]. Compared to [19], we apply trajectory space constraints for smoothing of the measurement matrix and not

throughout the whole optimisation. As a distinctive characteristic, we use proximal gradient method for optimisation [5]. To find the final shapes, we coarsely follow Block Matrix Method (BMM) [10]. Instead of using a fixed point continuation algorithm to relax rank minimisation to a nuclear norm, we use ADMM, similar to the patch reconstruction step of Lee *et al.* [23]. In contrast to the latter method, our approach is not piecewise. One major advantage of using ADMM is a guaranteed global minimum for a convex optimisation problem. In our experiments, this property has been proven to be effective for the global dense setting.

3. Proposed Approach

SMSR consists of two main steps. First, the camera motion is estimated followed by 2D shape alignment. Next, a time-varying 3D surface is recovered. In the following, we describe the proposed algorithmic combination.

3.1. Problem Formulation

We formulate NRSfM as a factorisation approach assuming every non-rigid shape follows LRSF. The measurement matrix $\mathbf{W} \in \mathbb{R}^{2T \times N}$ combining coordinates of N tracked points throughout T frames is decomposed as a product of two matrices:

$$\underbrace{\begin{bmatrix} \mathbf{W}_1 \\ \mathbf{W}_2 \\ \vdots \\ \mathbf{W}_T \end{bmatrix}}_{\mathbf{W}} = \underbrace{\begin{bmatrix} \mathbf{R}_1 & & & \\ & \mathbf{R}_2 & & \\ & & \ddots & \\ & & & \mathbf{R}_T \end{bmatrix}}_{\mathbf{R}} \underbrace{\begin{bmatrix} \mathbf{S}_1 \\ \mathbf{S}_2 \\ \vdots \\ \mathbf{S}_T \end{bmatrix}}_{\mathbf{S}}, \quad (1)$$

where \mathbf{R} is a quasi-block-diagonal matrix of camera poses, \mathbf{S} is the block matrix with non-rigid shapes, $\mathbf{R}_i \in \mathbb{R}^{2 \times 3}$ and $\mathbf{S}_i \in \mathbb{R}^{3 \times N}$ are camera pose and shape for i^{th} frame respectively. Every frame in \mathbf{W} is registered to the centroid of the observed structure. Alternatively, the measurement matrix can be written as

$$\mathbf{W} = \mathbf{R} \underbrace{(\mathbf{C} \otimes \mathbf{I}_3)}_{\mathbf{S}} \mathbf{B} = \mathbf{M} \mathbf{B}, \quad (2)$$

where \otimes denotes Kronecker product, \mathbf{I}_3 is a 3×3 identity matrix, $\mathbf{B} \in \mathbb{R}^{3K \times N}$ denotes set of basis shapes of cardinality K and $\mathbf{C} \in \mathbb{R}^{T \times K}$ is the corresponding *coefficient matrix*; $\mathbf{M} \in \mathbb{R}^{2T \times 3K}$ is a combined block diagonal shape-coefficient matrix (an element of the motion manifold). We utilise singular value decomposition (svd) to find a $3K$ -rank approximation of \mathbf{W} as $\mathbf{W} \cong \mathbf{M}' \mathbf{B}'$. An "implicit" non-unique solution can be obtained up to an invertible matrix $\mathbf{Q} \in \mathbb{R}^{3K \times 3K}$ (corrective transformation):

$$\mathbf{M}' \mathbf{Q} \cong \mathbf{M}, \quad \mathbf{Q}^{-1} \mathbf{B}' \cong \mathbf{B}. \quad (3)$$

It can be noticed that $\mathbf{M}' \mathbf{Q}$ is a scaled orthogonal matrix. Let $\mathbf{M}'_{2i-1:2i} \in \mathbb{R}^{2 \times 3K}$ denote the $i^{th} \in \{1, \dots, T\}$ pair of

rows of \mathbf{M} and $\mathbf{Q}_k \in \mathbb{R}^{3K \times 3}$ denote the $k^{th} \in \{1, \dots, K\}$ column triplet of \mathbf{Q} . Then, for every $\mathbf{M}'_{2i-1:2i}$ submatrix, the following product holds:

$$\mathbf{M}'_{2i-1:2i} \mathbf{Q}_k = c_{ik} \mathbf{R}_i, \quad (4)$$

where c_{ik} is the element in i^{th} row and k^{th} column of \mathbf{C} . Let $\mathbf{F} = \mathbf{Q}\mathbf{Q}^T$ be a positive semi-definite matrix. Using orthogonality constraint, two systems of linear equations can be obtained for each i :

$$\begin{cases} \mathbf{M}'_{2i-1} \mathbf{F}_k \mathbf{M}'_{2i-1}{}^T = \mathbf{M}'_{2i} \mathbf{F}_k \mathbf{M}'_{2i}{}^T = c_{ik}^2 \mathbf{I}, \\ \mathbf{M}'_{2i-1} \mathbf{F}_k \mathbf{M}'_{2i}{}^T = 0. \end{cases} \quad (5)$$

The system of equations (5) can be rewritten as

$$\underbrace{\begin{bmatrix} \mathbf{M}'_{2i-1} \otimes \mathbf{M}'_{2i-1}{}^T - \mathbf{M}'_{2i} \otimes \mathbf{M}'_{2i}{}^T \\ \mathbf{M}'_{2i-1} \otimes \mathbf{M}'_{2i}{}^T \end{bmatrix}}_{\mathbf{A}_i} \text{vec}(\mathbf{F}_k) = 0, \quad (6)$$

where $\text{vec}(\cdot)$ is vectorisation operator defined as $\text{vec}(\mathbb{R}^{m \times n}) \Rightarrow \mathbb{R}^{mn \times 1}$. Here, we use the property:

$$\text{vec}(\mathbf{A}\mathbf{\Psi}\mathbf{\Upsilon}^T) = (\mathbf{\Upsilon} \otimes \mathbf{A}) \text{vec}(\mathbf{\Psi}), \quad (7)$$

which holds for real matrices \mathbf{A} , $\mathbf{\Upsilon}$ and $\mathbf{\Psi}$. We denote the matrix on the left side of Eq. (6) with \mathbf{A}_i . Assembling equations for all \mathbf{A}_i by stacking leads to a single equation

$$\mathbf{A} \text{vec}(\mathbf{F}_k) = 0, \quad (8)$$

where $\mathbf{A} = [\mathbf{A}_1^T, \mathbf{A}_2^T, \dots, \mathbf{A}_T^T]$. The optimisation problem in Eq. (8) — finding an optimal \mathbf{F}_k — can be solved with linear least-squares by minimising

$$\|\mathbf{A} \text{vec}(\mathbf{F}_k)\|^2 \quad (9)$$

[27]. The latter problem is non-convex due to the rank-3 constraint on \mathbf{F}_k , and can be efficiently solved using proximal gradient method [5] in an iterative manner. In each iteration, we solve the following subproblem:

$$\min_{\mathbf{F}_k} \|\mathbf{F}_k - \mathbf{F}_k^*\|^2, \text{ s.t. } \text{rank}(\mathbf{F}_k) = 3, \quad (10)$$

$$\text{vec}(\mathbf{F}_k^*) = \text{vec}(\mathbf{F}_k^0) - \frac{1}{L_{\mathbf{A}}} \mathbf{A}^T \mathbf{A} \text{vec}(\mathbf{F}_k^0). \quad (11)$$

In Eq. (10), \mathbf{F}_k^0 is an initial estimate of \mathbf{F}_k and $L_{\mathbf{A}}$ denotes Lipschitz constant, *i.e.*, the largest eigenvalue of $\mathbf{A}^T \mathbf{A}$. Once an optimal \mathbf{F} is found, the corrective transformation \mathbf{Q} is recovered by Cholesky decomposition. Further, \mathbf{R} is computed using Eq. (4) and (5) with a known \mathbf{Q} .

3.2. Smooth Shape Trajectory

For the enhanced accuracy, we add smoothness constraint on point trajectories. It enforces the non-rigid shapes

to change gradually over time, *i.e.*, a trajectory represented by a K -dimensional $t^{th} \in \{1, \dots, T\}$ vector $c_t|_{1:K} = c(t)$ must lie in a linear subspace which constraints a point $c_{t,k}$ to vary smoothly over time. Thus, we represent \mathbf{C} by K compact cosine series:

$$\mathbf{C} = \Omega_d [x_1, \dots, x_K] = \Omega_d \mathbf{X} \quad \text{with } x_k \in \mathbb{R}^d, \quad (12)$$

where $k \in \{1, \dots, K\}$ and d are numbers of low frequency DCT coefficients in the shape trajectory $\mathbf{X} \in \mathbb{R}^{d \times K}$ which represents \mathbf{C} compactly over a truncated DCT domain $\Omega_d \in \mathbb{R}^{T \times d}$. The advantage of the constraints in the trajectory space is a known DCT basis which reduces considerably the number of unknowns. This leads to a faster overall convergence. Using the representation in trajectory space, \mathbf{M} can be written as

$$\mathbf{M} = \mathbf{R}(\mathbf{C} \otimes \mathbf{I}_3) = \mathbf{R}(\Omega_d \mathbf{X} \otimes \mathbf{I}_3). \quad (13)$$

Using the pre-computed camera motion \mathbf{R} (see Sec. 3.1), \mathbf{M} is treated as a function of \mathbf{X} only. Further, \mathbf{B} can be computed as

$$\mathbf{B} = \mathbf{M}^\dagger \mathbf{W}, \quad (14)$$

whereby \dagger denotes Moore-Penrose pseudo-inverse operator [14]. The problem of estimating DCT shape trajectory \mathbf{X} is formulated as a problem of minimising the squared reprojection error in the Frobenius norm:

$$\min_{\mathbf{X}} \|\mathbf{W} - \mathbf{W}^*\|_{\mathcal{F}}^2, \quad \text{and } \mathbf{W}^* = \mathbf{M}\mathbf{B} = \mathbf{M}\mathbf{M}^\dagger \mathbf{W}. \quad (15)$$

\mathbf{X} is initialised to $\mathbf{X}_0 = [\mathbf{I}_K \ 0]$ which is $K \times K$ identity matrix with padding of additional zeros. The higher frequency DCT coefficients can be then estimated by iterative Gauss-Newton minimisation of Eq. (15). Shape basis matrix \mathbf{B} is not directly used in the next part of the pipeline. Instead, only 2D projection obtained from the newly computed matrix $\mathbf{S} = (\Omega_d \mathbf{X} \otimes \mathbf{I}_3)\mathbf{B}$ is used as the new measurement matrix \mathbf{W} . The described pre-processing step reinitialises the input matrix after imposing smooth trajectory constraint on the recovered shapes.

3.3. Non-rigid Shape Recovery

We use the new \mathbf{W} to recover the final shapes and follow the formulation proposed in [10]. A rearranged shape matrix

$$\mathbf{S}^\# = \begin{bmatrix} X_{11} \dots X_{1N} & Y_{11} \dots Y_{1N} & Z_{11} \dots Z_{1N} \\ \vdots & \vdots & \vdots \\ X_{T1} \dots X_{TN} & Y_{T1} \dots Y_{TN} & Z_{T1} \dots Z_{TN} \end{bmatrix} \quad (16)$$

with an additional constraint $\text{rank}(\mathbf{S}^\#) < K$ compactly represents the optimal non-rigid structure:

$$\min_{\mathbf{S}} \|\mathbf{S}^\#\|_*, \quad \text{s.t. } \mathbf{W} = \mathbf{R}\mathbf{S}, \quad (17)$$

where $\|\cdot\|_*$ denotes the nuclear norm. We assume the mean 3D component is dominant in $\mathbf{S}^\#$ and can be removed in temporal dimension. Eq. (17) can be modified as

$$\min_{\mathbf{S}} \|\mathbf{S}^\# \mathbf{P}\|_*, \quad \text{s.t. } \mathbf{W} = \mathbf{R}\mathbf{S}, \quad (18)$$

where $\mathbf{P} = (\mathbf{I} - \frac{1}{T}\mathbf{1}\mathbf{1}^T)$ is the orthogonal projection and $\mathbf{1}$ is a vector of ones being its only null space. Eq. (18) is an optimisation problem over a convex function. We use for optimisation ADMM which can be considered as the inner-loop version of augmented Lagrangian method [25]:

$$\begin{aligned} L(\mathbf{S}, \mathbf{Y}, \mu) &= \\ &= \|\mathbf{S}^\# \mathbf{P}\|_* + \mathbf{Y}^T(\mathbf{W} - \mathbf{R}\mathbf{S}) + \frac{\mu}{2} \|\mathbf{W} - \mathbf{R}\mathbf{S}\|_{\mathcal{F}}^2, \end{aligned} \quad (19)$$

where \mathbf{Y} denotes the vector of Lagrange multipliers and $\mu > 0$ is the positive step size. Eq. (18) is solved by the optimisation

$$\max_{\mathbf{Y}} \min_{\mathbf{S}} L(\mathbf{S}, \mathbf{Y}, \mu). \quad (20)$$

\mathbf{S} is updated iteratively by SVT [25] using coordinate descent, and Lagrange multipliers are updated using sub-gradient in the iteration step t :

$$\mathbf{S}^{t+1} = \arg \min_{\mathbf{S}} L(\mathbf{S}, \mathbf{Y}^t, \mu^t), \quad (21)$$

$$\mathbf{Y}^{t+1} = \mathbf{Y}^t + \mu^t(\mathbf{W} - \mathbf{R}\mathbf{S}^{t+1}). \quad (22)$$

μ grows geometrically in every iteration to impose the linearity constraint. One advantage of ADMM is a guaranteed global minimum for convex problems and we use the above modified formulation to estimate the final shapes.

4. Experimental Results

We perform experiments on both synthetic as well as real data and compare the proposed approach with several existing approaches, *i.e.*, Metric Projection (MP) [28] (for the dense data sets we also use accelerated MP [18]), Point Trajectory Approach (PTA) [3], smooth time trajectories approach by Gotardo *et al.* (CSF1) [19], complementary rank-3 spaces approach CSF2 [21], Dai *et al.* (BMM) [10], VA [12] and Lee *et al.* [23]. We run experiments on a system with 32 GB RAM and Intel Core-i5200U processor. For SMSR implementation we used *Matlab* programming tool [26]. For quantitative performance evaluation of the methods, we compute normalised mean 3D error e_{3D} given by

$$e_{3D} = \frac{1}{\sigma TN} \sum_{t=1}^T \sum_{j=1}^n e_{tj}, \quad \sigma = \frac{1}{3T} \sum_{t=1}^T (\sigma_{tx} + \sigma_{ty} + \sigma_{tz}), \quad (23)$$

where σ is the normalised variance and e_{tj} is the reconstruction error (*i.e.*, Euclidean distance for the j^{th} 3D point of t^{th} frame) between ground truth \mathbf{G} and recovered shape

\mathbf{X} given by $e_{tj} = \|\mathbf{G}_j^t - \mathbf{X}_j^t\|_{\mathcal{F}}^2$, where $\|\cdot\|_{\mathcal{F}}^2$ is the squared Frobenius norm. Before computing the error metrics, the Procrustes analysis is performed to align recovered shapes with the ground truth [22].

4.1. Dense Data Sets with Ground Truth

To generate more data sets for quantitative evaluation, we propose to employ the available dynamic 3D reconstructions. For a given input, we firstly reconstruct non-rigid shapes with VA [?]. Next, the 3D shapes are transformed to the fronto-parallel position relative to the image plane and projected onto it with a virtual orthographic camera. Thus, ground truth point correspondences are obtained. The 3D shapes and point correspondences can now be utilised as a data set for NRSfM evaluation with a known ground truth. We obtained three data sets in this way — *face* [12], *heart* [30] and *barn owl* [11]. While generating 2D measurement matrices, for all sequences we apply the recovered camera poses to the dynamic shapes.

Additionally, we use the actor data set of Beeler *et al.* with given dynamic 3D shapes [6]. The sequence was acquired using motion capture techniques and interpolation of key frame shapes; it contains $\approx 1.1 \times 10^6$ points and over 300 frames with subtle details of an actor's face with different facial expressions and emotions. We use every odd frame from the first 101 frames in the experiment and create two modified data sets, namely *Actor1* — with the introduced left-right head movement — and *Actor2* — with the introduced left-right up-down head movement. Similar to the *face*, *heart* and *barn owl* sequences, we project the 3D shapes by an orthographic camera and obtain 2D correspondences. In the similar manner — using the new sequences *Actor1* and *Actor2* — we created modified data sets *Actor1 Sparse* and *Actor2 Sparse*, as some implementations were only able to execute with less points ($\approx 3.8 \times 10^4$).

4.2. Evaluation on Sparse Data Sets

We evaluate SMSF on challenging synthetic data sets *Drink*, *Pickup*, *Stretch*, *Yoga* with added per frame angular displacement of 5 degrees [4]. They contain 41 tracked points and up to 1102 frames. Further, SMSR is evaluated on real sparse benchmark *Dinosaur* and *Dance* sequences introduced in [4].

In Table 1, we show that the performance of the proposed method attains very close to the respective best performing method for nearly all cases (for *Drink* and *Yoga* sequences, SMSR matches the best e_{3D} up to the second decimal digit). Moreover, CSF1 achieves e_{3D} of 0.023 and 0.602 for *Drink* sequence and *Dinosaur* sequence and our method achieves 0.028 and 0.467 for the respective sequences. Hence, SMSR shows robust behaviour and steady performance on all sparse benchmark data sets.

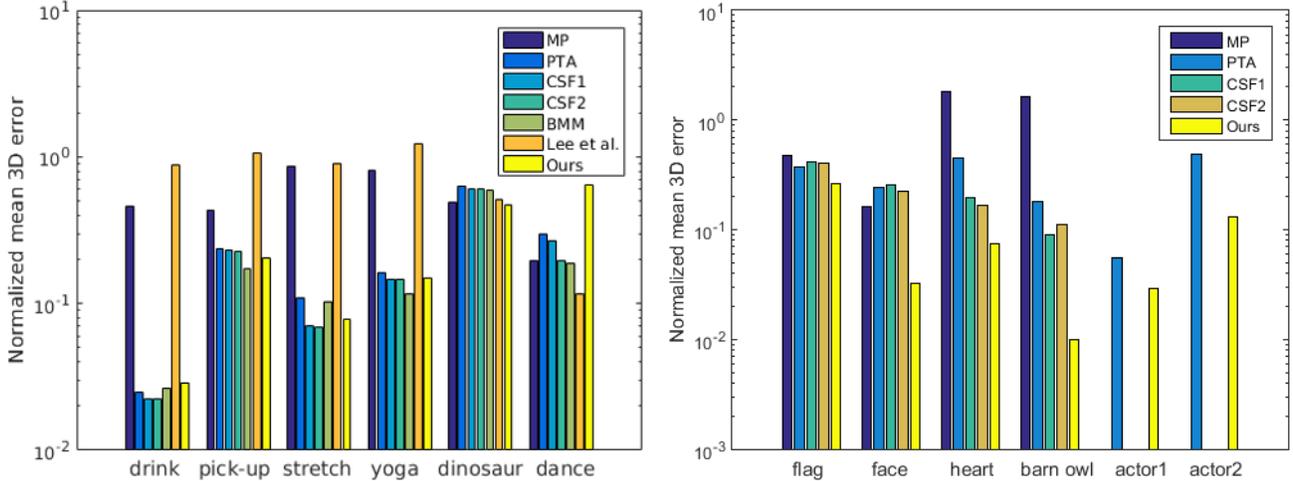


Figure 1: Comparison of the log scale normalised mean 3D error of SMSR (proposed) and several other methods over sparse and dense benchmark data sets. Missing bars in the figure for some approaches corresponds to N.A. in Table 2.

Method	Drink(K)	Pick-up(K)	Stretch(K)	Yoga(K)	Dinosaur(K)	Dance(K)
MP [28]	0.4604	0.4332	0.8549	0.8039	0.4894	0.1949
PTA [3]	0.0250(13)	0.2369(12)	0.1088(12)	0.1625(11)	0.6265(2)	0.2958(5)
CSF1 [19]	0.0223 (6)	0.2301(6)	0.0710 (8)	0.1467 (7)	0.6022(2)	0.2694(2)
CSF2 [21]	0.0223 (6)	0.2277(3)	0.0684 (8)	0.1465 (7)	0.6003(2)	0.1957(7)
BMM [10]	0.0266(12)	0.1731 (12)	0.1034(11)	0.1150 (10)	0.5874	0.1864(10)
Lee et al. [23]	0.8754	1.0689	0.9005	1.2276	0.5079	0.1156
SMSR (ours)	0.0287	0.2020	0.0783	0.1493	0.4678	0.6407

Table 1: Normalised mean 3D error e_{3D} for benchmark data sets [33, 3, 20]. K denotes number of bases which has led to the best result.

4.3. Evaluation on Dense Data Sets

We use in the experiments challenging dense synthetic mocap *flag* data set with ground truth surfaces [34, 13] with $\approx 10^4$ points and 60 frames. To the best of our knowledge, no real dense data set with the ground truth is available for quantitative evaluation of NRSfM yet. Aiming to fill this gap, we create further synthetic dense data sets with ground truth using publicly available sequences (see Sec. 4.1), *i.e.*, the *face* sequence with $\approx 28 \times 10^3$ points and 120 frames [12] (it contains sudden emotion changes between frames) and *heart* bypass surgery sequence [30] with $\approx 68 \times 10^3$ points and 30 frames. We used multi-frame optical flow [31] to compute dense correspondences. To demonstrate the scalable nature of SMSR, we have tested it on a dense real *barn owl* video acquired outdoors [11] with a higher number of points and frames ($\approx 2 \times 10^5$ points, 202 frames). Additionally, we use four synthetic face sequences introduced in [12] denoted by *Face Seq1* – *Face Seq4*. Each sequence contains 28887 points per frame. Note that the initial shape alignment by imposing smooth trajectories constraints (described in Sec. 3.2) takes longer time to converge where the total number of point tracks is high. As a result, for

these sequences, the initial shape alignment is disabled. Table 2 shows performance for mocap synthetic *flag* sequence [34] and real sequences [6, 12, 11]. The proposed method achieves the best performance and outperforms other approaches by a considerable margin. From Table 2 we can infer that other approaches exhibit high variation in performance across tested data sets and scale poorly for different deformations, number of points and frames. In contrast, SMSR performs robustly and is well scalable for sparse and dense data and different types of deformations.

In Table 3, we provide a quantitative comparison for several NRSfM methods on the newly created data sets *Actor1 Sparse*, *Actor2 Sparse* and synthetic face sequences *Face Seq1* – *Face Seq4*. The e_{3D} obtained by our approach remains lower than 0.04 for the modified actor data sets evidencing the consistent performance of our method. The qualitative performance evaluation for *Actor1 Sparse* sequence is shown in Fig. 2. SMSR accurately recovers actor’s face with distinct fine details while most of the other approaches fail. According to Table 3, SMSR achieves second best accuracy with e_{3D} being less than 0.18 while VA achieves best results for all synthetic face sequences, it

Method	<i>flag</i> (K)	<i>face</i> (K)	<i>heart</i> (K)	<i>barn owl</i> (K)	<i>Actor1</i> (K)	<i>Actor2</i> (K)
MP [28]	0.4756(8)	0.1604(2)	1.8231(4)	1.6277(4)	N.A.	N.A.
PTA [3]	0.3755(2)	0.2410(3)	0.4509(2)	0.1805(2)	0.0559(2)	0.4941(3)
CSF1 [19]	0.4202(4)	0.2581(2)	0.1941(2)	0.0896(2)	N.A.	N.A.
CSF2 [21]	0.4021(4)	0.2226(2)	0.1667(2)	0.1101(2)	N.A.	N.A.
SMSR (<i>ours</i>)	0.2631	0.0321	0.075	0.0099	0.0287	0.1318

Table 2: Normalised mean 3D error e_{3D} for benchmark data sets [34, 6, 12, 11]. K denotes number of bases which has led to the best result. N.A. is used where evaluation is not possible because of very high memory requirement or infeasible runtime. Likewise, BMM [10] and Lee *et al.* [23] could not be evaluated. The newly generated ground truth (see Sec. 4.1) does not contain connectivity data structure, thus running VA was not possible on these data sets.

Method	MP [28]	PTA [3]	CSF1 [19]	CSF2 [21]	VA [12]	SMSR (<i>ours</i>)
Actor1 Sparse(K)	0.5226(3)	0.0418 (2)	0.3711(2)	0.3708(2)	N.A.	0.0352
Actor2 Sparse(K)	0.2737(2)	0.0532 (2)	0.2275(3)	0.2279(3)	N.A.	0.0334
Face Seq1(K)	0.7251(2)	0.3933(2)	0.5325(3)	0.4677(3)	0.1058	0.1893
Face Seq2(K)	0.6633(2)	0.1871 (2)	0.9266(3)	0.7909(3)	0.1014	0.2133
Face Seq3(K)	0.5676(4)	0.1706(4)	0.5274(3)	0.5474(3)	0.0811	0.1345
Face Seq4(K)	0.5038(4)	0.2216(4)	0.5392(4)	0.5292(3)	0.0806	0.0984

Table 3: Normalised mean 3D error e_{3D} for the modified benchmark data set [6, 12]. K denotes number of bases which has led to the best result. BMM [10] and Lee *et al.* [23] could not be evaluated for *Actor1 Sparse* and *Actor2 Sparse* sequences because of infeasible memory requirement and very high runtime. The results of VA are included as presented in the publication [12]. N.A. is used for the newly generated data sets (see Sec. 4.1) which does not contain connectivity data structure, thus running VA was not possible on these data sets.

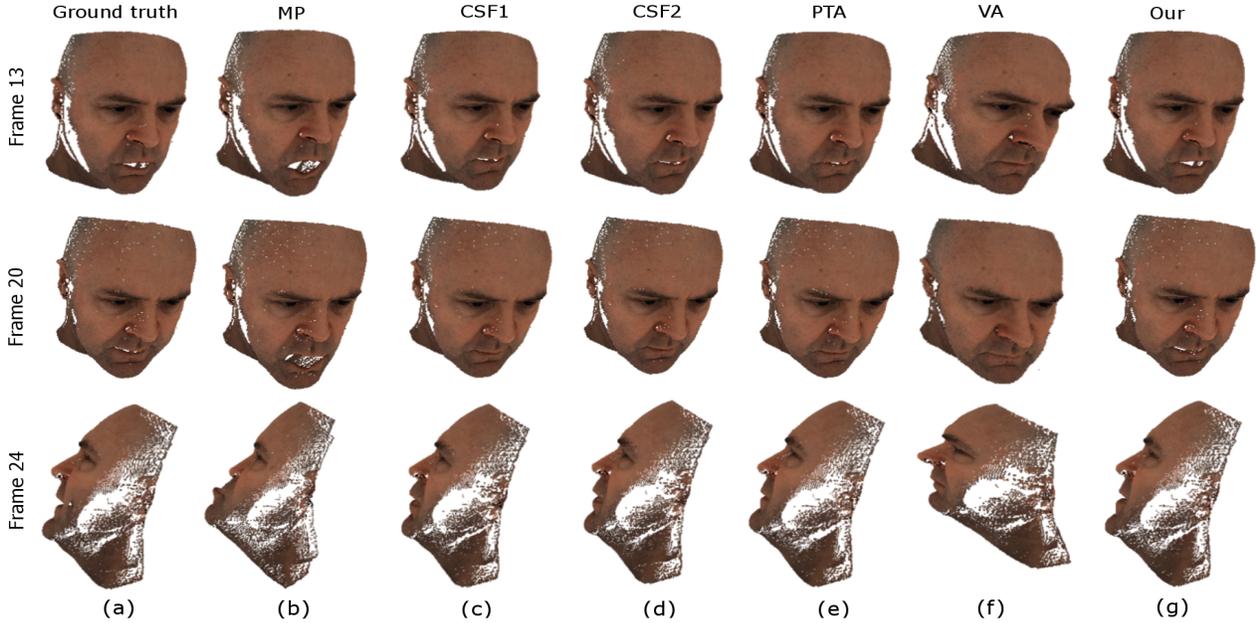


Figure 2: Qualitative performance evaluation between MP [28], CSF1 [19], CSF2 [21], PTA [3], VA [12] and our SMSR approach for the *Actor1 Sparse* sequence.

is found that these sequences contain overexaggerated facial deformations (*cf.* results of VA on sequences *Actor1 Sparse* and *Actor2 Sparse* with smaller and realistic deformations). We visualise 3D motion fields on the surface following the methodology proposed in [17]. Accordingly, in Fig. 3 we show two frames per data set in the top row and

the colour-coded magnitudes of the recovered 3D motion fields overlaid with the recovered shapes in the bottom row. Finally, Fig. 4 shows reconstructions obtained using various approaches — the third and fourth rows show that the SMSR reconstructions of *Face Seq1* – *Face Seq4* sequences remain very close to those of VA and ground truth.

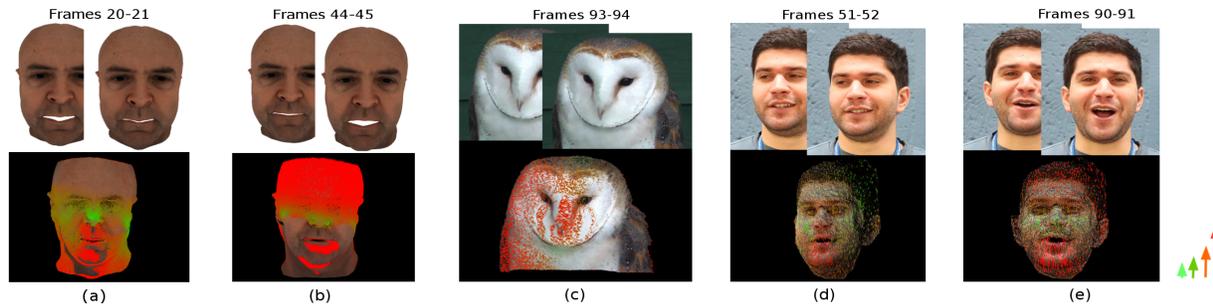


Figure 3: Visualisation of the 3D motion fields in the sense of the NRSfM-Flow framework [17] from the 3D shapes recovered by our SMSF method, for the dense benchmark data sets [6, 12, 11].

For all sparse benchmark data sets, SMSR takes less than 5 *sec* for the reconstructions, and between 40 and 300 *sec* for the dense data sets (*flag*, *heart*, *face*). Other sequences, *i.e.*, *barn owl* and *Actor* sequences can be reconstructed offline in feasible time (≤ 2500 *sec*) because of the high number of points. Note that our current implementation is not optimised. With hardware acceleration SMSR can be suitable for interactive applications.

4.4. NRSfM Challenge 2017

We tested SMSR on the NRSfM Challenge 2017 benchmark¹ [9]. This benchmark represents a set of five measurement matrices corresponding to five different monocular scene observations with small frame-to-frame changes — *Articulated Joints*, *Balloon Deflation*, *Paper Bending*, *Rubber Stretching* and *Paper Tearing* — and reports average root mean square (RMS) error² in *mm* for seventeen NRSfM methods with publicly available source code and methods submitted for evaluation. Measurement matrices for both perspective and orthographic cameras are provided. We reconstructed all sequences with the proposed SMSR and submitted to the portal (note that the ground truth 3D shapes are not publicly available). We use default parameters across all test sequences. SMSR achieves the average RMS error of 41.84 *mm* (ranging from 31.71 *mm* for *Balloon Deflation* to 58.12 *mm* for *Articulated Joints*) with the standard deviation of 9.2 *mm* — an intermediate overall accuracy. The low value of the standard deviation accentuates the scalable nature of the proposed SMSR. SMSR outperforms Lee *et al.* [23] and Dai *et al.* [10], though [28, 21, 19, 4] achieve lower errors on the new benchmark.

¹see <http://nrsfm2017.compute.dtu.dk/benchmark> for more details

² defined as $\frac{1}{T} \sum_{t=1}^T \frac{\|\mathbf{x}^t - \mathbf{G}^t\|_{\mathcal{F}}}{\mathbf{G}^t}$, with reconstructions \mathbf{X}^t and ground truth shapes \mathbf{G}^t

5. Conclusion

This paper presents a new method for scalable NRSfM which involves two main steps — camera motion estimation and 3D surface update. The primary advantage of the proposed SMSR approach is superior scalability and consistent performance across different data sets which are two core synergic properties which distinguish our SMSR method from the competing methods. SMSR imposes constraints in metric and trajectory spaces and involves various widely-used high-level operations in different steps including proximal gradient method, iterative Gauss-Newton optimisation, singular value thresholding, coordinate descent, Cholesky decomposition as well as multiple matrix-vector and matrix multiplications. We often observed a quick convergence of our method. The reconstruction error obtained in the extensive experiments by the proposed approach is low and consistent over a large variety of publicly available real and synthetic data sets including sparse and dense ones. To facilitate and enrich quantitative evaluation, we evaluated performance of several approaches on ground truth tracks obtained as projections of trail VA reconstructions. In the majority of the cases, SMSR achieves the lowest normalised mean 3D error among all tested methods or reaches close to the respective best performing method. One reason for this gain is the new pre-processing step, *i.e.*, smoothing of the shape trajectory. Another possible reason is the convergence property of the ADMM minimisation.

A current limitation of the proposed technique lies in handling large deformations such as those occurring during tissue bending and twisting. Extending SMSR for the handling of large deformations is an interesting direction for the future work. Another possible direction could be adopting the method for online applications and embedded devices.

Acknowledgement

This research was supported by the project DYNAMICS (01IW15003) of the German Federal Ministry of Education and Research (BMBF).

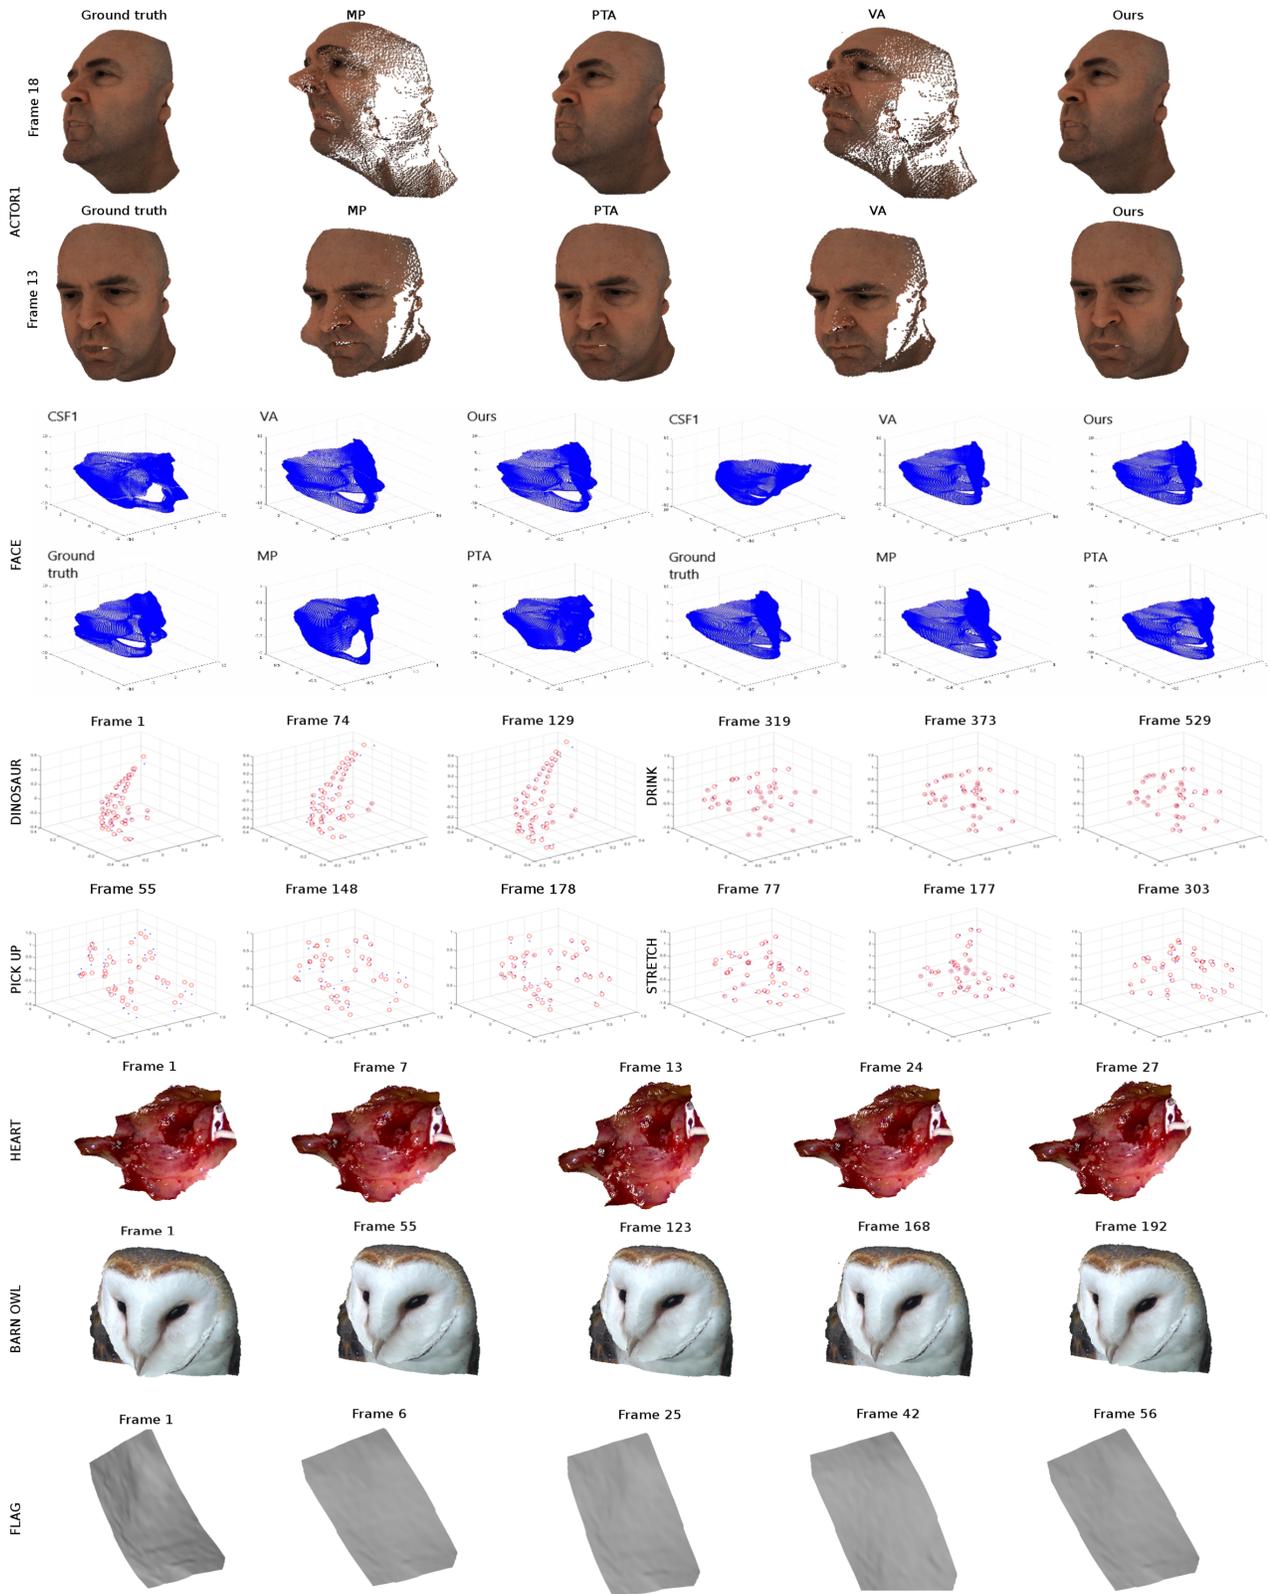


Figure 4: (Top 2 rows): qualitative evaluation on dense benchmark data set *Actor1* [6]; (rows 3–4): qualitative evaluation on dense benchmark data sets *Face1* (left) and *Face2* (right) [12]; (rows 5–6): qualitative evaluation on sparse benchmark data sets *Dinosaur*, *Pickup*, *Drink*, *Stretch* [33, 3]; (last three rows): quantitative evaluation of our approach on dense benchmark data sets *Heart* [30], *Barn Owl* [11] and *Flag* [34].

References

- [1] A. Agudo, J. Montiel, L. Agapito, and B. Calvo. Online Dense Non-Rigid 3D Shape and Camera Motion Recovery. In *British Machine Vision Conference (BMVC)*, 2014. **1**
- [2] I. Akhter, Y. Sheikh, and S. Khan. In defense of orthonormality constraints for nonrigid structure from motion. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1534–1541, 2009. **2**
- [3] I. Akhter, Y. Sheikh, S. Khan, and T. Kanade. Nonrigid structure from motion in trajectory space. In *Advances in Neural Information Processing Systems (NIPS)*, pages 41–48, 2009. **1, 4, 5, 6, 8**
- [4] I. Akhter, Y. Sheikh, S. Khan, and T. Kanade. Trajectory space: A dual representation for nonrigid structure from motion. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 33(7):1442–1456, 2011. **2, 4, 7**
- [5] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009. **2, 3**
- [6] T. Beeler, F. Hahn, D. Bradley, B. Bickel, P. Beardsley, C. Gotsman, R. W. Sumner, and M. Gross. High-quality passive facial performance capture using anchor frames. In *ACM Transactions on Graphics (TOG)*, 2011. **4, 5, 6, 7, 8**
- [7] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122, 2011. **1**
- [8] C. Bregler, A. Hertzmann, and H. Biermann. Recovering non-rigid 3d shape from image streams. In *Computer Vision and Pattern Recognition (CVPR)*, volume 2, pages 690–696, 2000. **2**
- [9] A. D. Bue, H. Aanaes, S. N. Jensen, and Y. Sheikh. Non-rigid structure from motion challenge 2017. <http://nrsfm2017.compute.dtu.dk/benchmark/>, 2017. **7**
- [10] Y. Dai, H. Li, and M. He. A simple prior-free method for non-rigid structure-from-motion factorization. *International Journal of Computer Vision (IJCV)*, 107(2):101–122, 2014. **2, 3, 4, 5, 6, 7**
- [11] P. Dinning. *Barn Owl at Screech Owl Sanctuary*. <https://www.youtube.com/watch?v=xmou8t-DHh0>, 2014. [online; accessed on 02.05.2017; usage rights obtained from the author]. **4, 5, 6, 7, 8**
- [12] R. Garg, A. Roussos, and L. Agapito. Dense variational reconstruction of non-rigid surfaces from monocular video. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1272–1279, 2013. **1, 2, 4, 5, 6, 7, 8**
- [13] R. Garg, A. Roussos, and L. Agapito. A variational approach to video registration with subspace constraints. *International Journal of Computer Vision (IJCV)*, 104(3):286–314, 2013. **5**
- [14] G. H. Golub and C. F. Van Loan. *Matrix computations*, volume 3. 2012. **3**
- [15] V. Golyanik, T. Fetzter, and D. Stricker. Accurate 3d reconstruction of dynamic scenes from monocular image sequences with severe occlusions. In *Winter Conference on Applications of Computer Vision (WACV)*, 2017. **1, 2**
- [16] V. Golyanik, T. Fetzter, and D. Stricker. Introduction to coherent depth fields for dense monocular surface recovery. In *British Machine Vision Conference (BMVC)*, 2017. **1**
- [17] V. Golyanik, A. S. Mathur, and D. Stricker. NRSfM-Flow: Recovering non-rigid scene flow from monocular image sequences. In *British Machine Vision Conference (BMVC)*, 2016. **6, 7**
- [18] V. Golyanik and D. Stricker. Dense batch non-rigid structure from motion in a second. In *Winter Conference on Applications of Computer Vision (WACV)*, 2017. **4**
- [19] P. F. Gotardo and A. M. Martinez. Computing smooth time trajectories for camera and deformable shape in structure from motion with occlusion. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 33(10):2051–2065, 2011. **1, 2, 4, 5, 6, 7**
- [20] P. F. Gotardo and A. M. Martinez. Kernel non-rigid structure from motion. In *International Conference on Computer Vision (ICCV)*, pages 802–809, 2011. **5**
- [21] P. F. Gotardo and A. M. Martinez. Non-rigid structure from motion with complementary rank-3 spaces. In *Computer Vision and Pattern Recognition (CVPR)*, pages 3065–3072, 2011. **2, 4, 5, 6, 7**
- [22] J. C. Gower. Generalized procrustes analysis. *Psychometrika*, 40(1):33–51, 1975. **4**
- [23] M. Lee, J. Cho, and S. Oh. Consensus of non-rigid reconstructions. In *Computer Vision and Pattern Recognition (CVPR)*, 2016. **1, 2, 4, 5, 6, 7**
- [24] M. Lee, C.-H. Choi, and S. Oh. A procrustean markov process for non-rigid structure recovery. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1550–1557, 2014. **2**
- [25] Z. Lin, M. Chen, and Y. Ma. The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices. *Technical Report UILU-ENG-09-2215*, University of Illinois at Urbana-Champaign, 2009. **1, 4**
- [26] MathWorks. *Matlab*. <https://www.mathworks.com/>, 2016. **4**
- [27] T. Morita and T. Kanade. A sequential factorization method for recovering shape and motion from image streams. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 19(8):858–867, 1997. **3**
- [28] M. Paladini, A. Del Bue, M. Stosic, M. Dodig, J. Xavier, and L. Agapito. Factorization for non-rigid and articulated structure using metric projections. In *Computer Vision and Pattern Recognition (CVPR)*, pages 2898–2905, 2009. **1, 4, 5, 6, 7**
- [29] C. Russell, J. Fayad, and L. Agapito. Dense non-rigid structure from motion. In *Second International Conference on 3D Imaging, Modeling, Processing, Visualization & Transmission*, pages 509–516, 2012. **1, 2**
- [30] D. Stoyanov. Stereoscopic scene flow for robotic assisted minimally invasive surgery. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 479–486, 2012. **4, 5, 8**
- [31] B. Taetz, G. Bleser, V. Golyanik, and D. Stricker. Occlusion-aware video registration for highly non-rigid objects. In *Winter Conference on Applications of Computer Vision (WACV)*, 2016. **5**

- [32] C. Tomasi and T. Kanade. Shape and motion from image streams under orthography: a factorization method. *International Journal of Computer Vision (IJCV)*, 9(2):137–154, 1992. [1](#), [2](#)
- [33] L. Torresani, A. Hertzmann, and C. Bregler. Nonrigid structure-from-motion: Estimating shape and motion with hierarchical priors. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 30(5):878–892, 2008. [1](#), [5](#), [8](#)
- [34] R. White, K. Crane, and D. A. Forsyth. Capturing and animating occluded cloth. In *ACM Transactions on Graphics (TOG)*, 2007. [5](#), [6](#), [8](#)
- [35] J. Xiao, J. Chai, and T. Kanade. A closed-form solution to non-rigid shape and motion recovery. *International Journal of Computer Vision (IJCV)*, 67(2):233–246, 2006. [1](#)