

Incremental Non-Rigid Structure-from-Motion with Unknown Focal Length

Thomas Probst¹, Danda Pani Paudel¹, Ajad Chhatkuli¹, and Luc Van Gool^{1,2}

¹ Computer Vision Lab, ETH Zürich, Switzerland

² VISICS, ESAT/PSI, KU Leuven, Belgium

{probstt, paudel, ajad.chhatkuli, vangool}@vision.ee.ethz.ch

Abstract. The perspective camera and the isometric surface prior have recently gathered increased attention for Non-Rigid Structure-from-Motion (NRSfM). Despite the recent progress, several challenges remain, particularly the computational complexity and the unknown camera focal length. In this paper we present a method for incremental Non-Rigid Structure-from-Motion (NRSfM) with the perspective camera model and the isometric surface prior with unknown focal length. In the template-based case, we provide a method to estimate four parameters of the camera intrinsics. For the template-less scenario of NRSfM, we propose a method to upgrade reconstructions obtained for one focal length to another based on local rigidity and the so-called Maximum Depth Heuristics (MDH). On its basis we propose a method to simultaneously recover the focal length and the non-rigid shapes. We further solve the problem of incorporating a large number of points and adding more views in MDH-based NRSfM and efficiently solve them with Second-Order Cone Programming (SOCP). This does not require any shape initialization and produces results orders of times faster than many methods. We provide evaluations on standard sequences with ground-truth and qualitative reconstructions on challenging YouTube videos. These evaluations show that our method performs better in both speed and accuracy than the state of the art.

1 Introduction

Given images of a rigid object from different views, Structure-from-Motion (SfM) [1–3] allows the computation of the object’s 3D structure. However, many such objects of interest are non-rigid and the rigidity constraints of SfM do not hold. The ever increasing number of monocular videos with deforming objects means provides a large incentive for being able to reconstruct such scenes. Such reconstruction problems can be solved with Non-Rigid Structure-from-Motion (NRSfM) which uses multiple images of a deforming object to reconstruct its 3D from a single camera. Another related approach computes the shape based on the object’s template shape and its deformed image, also termed as Shape-from-Template (SfT). While SfM is well-posed and has already seen several applications in commercial software [4, 5], non-rigid reconstruction has inherent theoretical problems. It is severely under-constrained without prior knowledge of the deformation or the shapes. In fact given a number of images, infinite possibilities of deformations exist that provide the same image projections. Therefore, one of the major challenges in NRSfM is to efficiently combine a realistic deformation constraint and the camera projection model to reduce the solution ambiguity.




Fig. 1: **Qualitative Results.** Comparison of our dense NRSfM method (bottom-right) to Ji et al. [6] (top-left) and Dai et al. [7] (top-right) on three different sequences.

A large majority of previous methods tackle NRSfM with an affine camera model and a low rank approximation of the deforming shapes [7–14]. However, such methods do not handle perspective effects and nonlinear deformations very well. In this paper we study the use of the uncalibrated perspective camera and the isometric deformation prior for non-rigid reconstruction. Isometry is a geometric prior which implies that the geodesic distances on the surface are preserved with the deformations. This is a good approximation for many real objects such as a human body, paper-like surfaces, or cloth. In SfT, the use of the isometric deformation prior with the perspective camera is considered to be the state-of-the-art [15–17] among the parameter-free approaches. In particular, [15, 18] also estimate the focal length while recovering the deformation. In NRSfM, some recent methods [6, 19] provide a convex formulation with the inextensible deformation for a calibrated perspective camera setup. The reconstruction is achieved by maximizing depth along the sightlines introduced in [20, 21] for template-based reconstruction. Although the methods use the perspective camera model and geometric priors for non-rigid reconstruction, their computational complexity does not allow reconstructing a large number of points. On the other hand, some recent dense methods using the perspective camera model have shown promising results, but they rely on piecewise rigidity constraints [22, 23] and shape initialization; this may be too constraining for several applications. Furthermore, methods using the perspective camera either rely on known intrinsics or cannot handle significant nonrigidity [24]. To the best of our knowledge, estimation of the unknown focal length has not been investigated in NRSfM for deforming surfaces.

In this paper we address the aforementioned issues with methods based on the convex relaxation of isometry. More precisely, we provide the following contributions: *a)* a method to ‘upgrade’ the non-rigid reconstruction obtained using incorrect camera intrinsics to the reconstruction of the correct one, *b)* a method to estimate intrinsics - all five entries in the case of SfT and the unknown focal length in the case template-less NRSfM *c)* an incremental method to add more points to the sparse 3D point-sets for consistent and semi-dense reconstruction *d)* online method of reconstruction by adding images. Besides being of immense practical concern and theoretical value, questions *a)* and *b)* have not been attempted for NRSfM for deforming objects. We provide a unified framework to solve the problems *a)* through *d)* using depth maximization and the relaxations of the isometry prior. We provide theoretical justification along with practical methods for intrinsics/focal length estimation as well as densification and online recon-

struction strategies. Despite being extremely challenging, we show the applicability of our method with compelling results. A few examples among them is shown in Fig 1.

1.1 Related Work

We discuss briefly the methods based on the isometry prior and the perspective camera model. This has been widely explored in the template-based methods [20, 21, 25]. In particular, [21] uses the inextensibility as a relaxation of the isometry prior in order to formulate non-rigid reconstruction as a convex problem by maximizing the depth point-wise. Several recent NRSfM methods [6, 19, 26–28] also use isometry or inextensibility with the perspective camera model. [27, 28] require the correspondence mapping function with its first and second-order derivatives limiting their application in practice. [19] improved upon [26] by providing a convex solution to NRSfM. They achieve this by maximizing pointwise depth in all views under the inextensibility cone constraints of [21] while also computing the template geodesics. Very recently a method [6] improving upon [19] suggested the use of maximization of sightlines rather than the pointwise depth. Both these methods have shown that moving the surface away from the camera under the inextensibility constraints can be formulated as a convex problem effectively reconstructing non-rigid as well as rigid objects. A different class of methods that use energy minimization approach on an initial solution also use the perspective camera model but with a piece-wise rigidity prior [22, 23]. However, all of these methods discussed here require the calibrated camera for reconstruction and do not provide any insights on how they can be extended to an uncalibrated camera. One notable exception is given by [24], however this approach is limited to dynamic scenes featuring a few independently moving objects [29, 30]. Yet another problem that has not been addressed in [6, 19] is the incremental reconstruction of a large number of points. Semi-dense or dense reconstruction as such is not possible here due to the high computational complexity of these methods.

2 Problem Modelling

We pose the NRSfM problem as that of finding point-wise depth in each view. We write the unknown depth as λ_i^l and the known homogeneous image coordinates as u_i^l , for the point i in the l -th image. A set of neighboring points of i is denoted by $\mathcal{N}(i)$. d_{ij} represents the template geodesic distance between point i and j , which is an unknown quantity for the NRSfM problem and a known quantity for the SfT problem. We define a nearest neighborhood graph as a set of fixed number of neighbors for each point i [19]. To represent the exact isometric NRSfM problem, we also introduce a geodesic distance function between two 3D points on the surface \mathcal{S} , $g_{\mathcal{S}}(x, y) : \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{R}$. Given the camera intrinsics K , the isometric NRSfM problem can be written as:

$$\begin{aligned} \text{Find } & K, \lambda_i^l \\ \text{s.t. } & g(K^{-1}\lambda_i^l u_i^l, K^{-1}\lambda_j^l u_j^l) = d_{ij}, \forall i, \forall j. \end{aligned} \quad (1)$$

(1) defines a non-convex problem and is also not tractable in its given form. It has been shown that with various relaxations [6, 19, 26], problem (1) can be solved for a known

K when different views and deformations are observed. In order to tackle the NRSfM problem with an unknown focal length we start with the observation that not all such solutions provide isometrically consistent shapes through all the views. We formulate our methods in the following sections.

3 Uncalibrated NRSfM

Given a known object template and a calibrated camera the NRSfM problem in (1) can be formulated as a convex problem by relaxing the isometry constraint with an inextensibility constraint [21] as below:

$$\begin{aligned} \max_{\lambda_i^l} \quad & \sum_l \sum_i \lambda_i^l, \\ \text{s.t.} \quad & \left\| K^{-1}(\lambda_i^l \mathbf{u}_i^l - \lambda_j^l \mathbf{u}_j^l) \right\| \leq d_{ij}, \quad \forall j \in \mathcal{N}(i). \end{aligned} \quad (2)$$

We are, however, interested on solving the same problem when both d_{ij} and K are unknown. Unfortunately, this problem is not only non-convex, but also unbounded. Therefore, we use two extra constraints on the variables K and d_{ij} such that the problem of (2), for unknown d_{ij} and K , becomes bounded.

$$\sum_i \sum_{j \in \mathcal{N}(i)} d_{ij} = 1, \quad K \leq \bar{K}. \quad (3)$$

Despite being bounded with the addition of (3), the reconstruction problem is still non-convex. More importantly, the maximization of the objective function favors the solution when K is as close as possible to \bar{K} . Therefore, we instead solve the reconstruction problem in (2) with a fixed initial guess \hat{K} and seek for the upgrade of both intrinsics and reconstruction later. Note that fixing the intrinsics makes the problem convex and identical to that in [19].

$$\begin{aligned} \max_{\lambda_i^l, d_{ij}} \quad & \sum_l \sum_i \lambda_i^l, \\ \text{s.t.} \quad & \left\| \hat{K}^{-1}(\lambda_i^l \mathbf{u}_i^l - \lambda_j^l \mathbf{u}_j^l) \right\| \leq d_{ij}, \quad j \in \mathcal{N}(i), \\ & \sum_i \sum_{j \in \mathcal{N}(i)} d_{ij} = 1. \end{aligned} \quad (4)$$

Now, we are interested to upgrade the solution of (4) such that the upgraded reconstruction correctly describes the deformed object in the 3D-space. In this work, the upgrade is carried out using a pointwise upgrade equation. In the following, we first derive this upgrade equation assuming that the correct focal length is known and then provide the theory and practical approaches to recover the unknown focal length.

3.1 Upgrade Equation

Let us consider, λ_i^l and $\hat{\lambda}_i^l$ are depths, of the point represented by \mathbf{u}_i^l , obtained from (2) and (4), respectively. The following proposition is the key ingredient of our work that relates $\hat{\lambda}_i^l$ to λ_i^l for the reconstruction upgrade.

Proposition 1. For $u_i^l \approx u_{\mathcal{N}(i)}^l$, $\hat{\lambda}_i^l$ can be upgraded to λ_i^l with the known \mathbf{K} using,

$$\lambda_i^l \approx \frac{\hat{\lambda}_i^l \|\hat{\mathbf{K}}^{-1} u_i^l\|}{\|\mathbf{K}^{-1} u_i^l\|}. \quad (5)$$

Proof. It is sufficient to show that every $j \in \mathcal{N}(i)$ satisfies $\|\hat{\mathbf{K}}^{-1}(\hat{\lambda}_i^l u_i^l - \hat{\lambda}_j^l u_j^l)\| \approx \|\mathbf{K}^{-1}(\lambda_i^l u_i^l - \lambda_j^l u_j^l)\|$. From (5), for any $u_i^l \approx u_{\mathcal{N}(i)}^l$, $\|\hat{\mathbf{K}}^{-1}(\hat{\lambda}_i^l u_i^l - \hat{\lambda}_j^l u_j^l)\|^2$ can be expressed as,

$$\begin{aligned} &\approx \|\mathbf{K}^{-1} u_i^l\|^2 \|\hat{\mathbf{K}}^{-1}(\lambda_i^l - \lambda_j^l) u_i^l\|^2 / \|\hat{\mathbf{K}}^{-1} u_i^l\|^2, \\ &= (\lambda_i^l - \lambda_j^l)^2 \|\mathbf{K}^{-1} u_i^l\|^2 \approx \|\mathbf{K}^{-1}(\lambda_i^l u_i^l - \lambda_j^l u_j^l)\|^2. \square \end{aligned} \quad (6)$$

Note that the condition $u_i^l \approx u_{\mathcal{N}(i)}^l$ is valid for any two sufficiently close neighbors. Such neighbors can be chosen using only the image measurements. More importantly, the assumption $u_i^l \approx u_{\mathcal{N}(i)}^l$ still allows depths λ_i^l and $\lambda_{\mathcal{N}(i)}^l$ to be different. This plays a vital role especially when the close neighboring points differ distinctly in depth, either due to camera perspective or high frequency structural changes. Although, (5) is only a close approximation for the reconstruction upgrade, its upgrade quality in practice was observed to be accurate. The following remark concerns Proposition 1.

Remark 1. As the guess on intrinsics $\hat{\mathbf{K}}$ tends to the real intrinsics \mathbf{K} , the upgrade equation (5) holds true for exact equality even when $u_i^l \not\approx u_{\mathcal{N}(i)}^l$. In other words,

$$\lim_{\hat{\mathbf{K}} \rightarrow \mathbf{K}} \lambda_i^l = \hat{\lambda}_i^l. \quad (7)$$

3.2 Upgrade Strategies

The upgrade equation presented in Proposition 1 assumes that the exact intrinsics \mathbf{K} is known. However, for uncalibrated NRSfM, \mathbf{K} is unknown. While the principal point can be assumed to be at the center of the image for most cameras [31], nothing can be said about the focal length. We henceforth, present strategies to estimate \mathbf{K} in two different scenarios of known and unknown shape template. We rely on the fact that isometric deformation, to a large extent, preserves local rigidity. This is reflected somewhat in the reconstruction obtained from (4). However, due to changes in the perspective and the extension of points along incorrect sightlines, the use of incorrect intrinsics produces reconstructions that are very less likely to remain isometric across different views. Similarly, an upgrade towards the correct intrinsics in that case produces reconstructions which satisfy the isometry better. This is also supported by the results in Section 6. There are various ways one can use isometry of the reconstructed surfaces to determine the correct intrinsics. A very simple method would be to use the fact that given reconstructed points that are dense enough, the correct intrinsics must preserve the local

euclidean distance. For $\hat{a}_i = \hat{\lambda}_i \|\hat{\mathbf{K}}^{-1} \mathbf{u}_i\|$, the euclidean distance between two upgraded neighboring 3D points, in any view as a function of intrinsics, can be expressed as,

$$\hat{d}_{ij}(\mathbf{K}) = \left\| \frac{\hat{a}_i \mathbf{K}^{-1} \mathbf{u}_i}{\|\mathbf{K}^{-1} \mathbf{u}_i\|} - \frac{\hat{a}_j \mathbf{K}^{-1} \mathbf{u}_j}{\|\mathbf{K}^{-1} \mathbf{u}_j\|} \right\|. \quad (8)$$

Now, we present techniques to estimate \mathbf{K} when the shape template is known (SfT), followed by a method to estimate the focal length for template-less case of NRSfM.

Template-based Calibration For the sake of simplicity, we present the calibration theory using only one image. This is also the sufficient condition for reconstruction when the shape template is known [21]. Recall that for SfT, d_{ij} in (4) are already known during the reconstruction process. For known template distance d_{ij} and the estimated euclidean distance after reconstruction upgrade $\hat{d}_{ij}(\mathbf{K})$, the intrinsics \mathbf{K} can be estimated by minimizing,

$$\Phi_{\mathcal{T}}(\mathbf{K}) = \sum_i \sum_{j \in \mathcal{N}(i)} \left(d_{ij} - \hat{d}_{ij}(\mathbf{K}) \right)^2. \quad (9)$$

Alternatively, one can also derive polynomial equations on the entries of the so-called Image of the Absolute Conic (IAC), defined as $\Omega = \mathbf{K}^{-\top} \mathbf{K}^{-1}$.

Proposition 2. *As long as the rigidity between any pair $\{\mathbf{u}_i, \mathbf{u}_j\}$ is maintained, either for any $\hat{\mathbf{K}}$ and $\mathbf{u}_i \approx \mathbf{u}_j$ or for any pair $\{\mathbf{u}_i, \mathbf{u}_j\}$ as $\hat{\mathbf{K}} \rightarrow \mathbf{K}$, the IAC can be approximated by solving,*

$$\mathbf{u}_i^{\top} \Omega \mathbf{u}_i \mathbf{u}_j^{\top} \Omega \mathbf{u}_j = \gamma_{ij} (\mathbf{u}_i^{\top} \Omega \mathbf{u}_j)^2, \quad (10)$$

for sufficiently many pairs, where,

$$\gamma_{ij} = \left(\frac{2\hat{a}_i \hat{a}_j}{\hat{a}_i^2 + \hat{a}_j^2 - d_{ij}^2} \right)^2. \quad (11)$$

We provide the proof in the supplementary material.

Note that (10) is a degree 2 polynomial on the entries of Ω . Since, Ω has 5 degrees of freedom, it can be estimated from 5 pairs of image points, using numerical methods.

The core idea of our template-based calibration consists of three steps: (i) a fixed number of hypothesis generation, (ii) hypothesis validation using the upgraded reconstruction quality, (iii) refinement of the best hypothesis.

Hypothesis generation: Given the template-based uncalibrated reconstruction from (4), we generate a set of hypotheses for camera intrinsics from randomly selected sets of minimal closest-point pairs. For every minimal set, we solve (10) for Ω to obtain these hypotheses. Then, the camera intrinsics \mathbf{K} is recovered by performing the Cholesky-decomposition on Ω .

Hypothesis validation: Each hypothesis is validated by computing its 3D reconstruction error. To do so, we first upgrade the initial reconstruction using the upgrade (5) for current hypothesis. Then, the reconstruction error is computed using (9). The hypothesis that results into minimum reconstruction error is chosen for further refinement.

Intrinsics refinement: Starting from the best hypothesis, we refine the intrinsics by minimizing the following objective function:

$$\mathcal{E}(\mathbf{K}) = \Phi_T(\mathbf{K}) + k_{(1,3)}^2 + k_{(2,3)}^2 + \left(1 - \frac{k_{(1,1)}}{k_{(2,2)}}\right)^2, \quad (12)$$

where, $k_{(i,j)}$ is the i^{th} -row and j^{th} -column entry of the normalized intrinsic matrix \mathbf{K} . Note that, we regularize the 3D reconstruction error $\Phi_T(\mathbf{K})$ by the expected structure of \mathbf{K} (i.e. principal point close to the center and unit aspect ratio). Our regularization term is often the main objective for existing autocalibration methods [31,32]. The minimization of objective $\mathcal{E}(\mathbf{K})$ can be carried out efficiently using locally optimal iterative refinement methods.

Now, we summarize our calibration method in Algo. 1.

Algorithm 1 $[\mathbf{K}] = \text{calibrateWithTemplate}(\hat{\mathbf{K}})$

1. Reconstruct 3D using (4) for known d_{ij} and the guess $\hat{\mathbf{K}}$.
 2. Select multiple sets of minimal closest-point pairs $\{u_i, u_j\}$.
 3. For each set,
 - (i) Generate hypothesis $\tilde{\mathbf{K}}$ by solving (10).
 - (ii) Upgrade the reconstruction for $\tilde{\mathbf{K}}$ using (5).
 - (iii) Compute the reconstruction error for $\tilde{\mathbf{K}}$ using (9).
 4. Among all sets, choose $\tilde{\mathbf{K}}$ with best reconstruction error.
 5. Refine the best hypothesis $\tilde{\mathbf{K}}$ using (12) to obtain \mathbf{K} .
-

Template-less Calibration As the self-calibration with the unknown template is extremely challenging, we relax it by considering that the principal point is at the center of the image and that the two focal lengths are equal. We assume that the intrinsics are constant across views. We then measure the consistency of the upgraded local euclidean distances, defined by (8), across different views. More precisely, we wish to estimate the focal length in \mathbf{K} by minimizing the following objective function,

$$\Phi(\mathbf{K}) = \sum_k \sum_{l \neq k} \sum_i \sum_{j \in \mathcal{N}(i)} \left(\hat{d}_{ij}^k(\mathbf{K}) - \hat{d}_{ij}^l(\mathbf{K}) \right)^2. \quad (13)$$

Ideally, it is also possible to derive polynomials on Ω , analogous to (10). This can be done by eliminating the unknown variable d_{ij} from two equations for two views of the same pair. Unfortunately, the equation derived in this manner does not turn out to be easily tractable. Alternatively, one can also attempt to solve the polynomials without eliminating variables d_{ij} – on both variables Ω and d_{ij} . However for practical reasons¹, we design a method assuming only one entry of Ω , corresponding to the focal length, is unknown. Under such assumption, we show in the supplementary materials that a polynomial of degree 4, one variable, equivalent to (10), can also be derived.

In this paper, we avoid making hypothesis on the focal length, since it is not really necessary. Unlike the case of template-based calibration, we address the problem of

¹For most of the cameras, it is safe to assume that their intrinsics have no skew, unit aspect ratio, and a principal point close to the image center.

template-less calibration iteratively in two steps: (i) focal length refinement, (ii) focal length validation. Henceforth for the template-less calibration, we make a slight abuse of notation by using K even for the intrinsics with only unknown focal length, unless mentioned otherwise.

Focal length refinement: Given an initial guess on focal length, its refinement is carried out by minimizing the objective function $\Phi(K)$ of (13) (optionally, on the full intrinsics). This refinement process finds a refined K which results a better isometric consistency of the reconstructions across views.

Focal length validation: The main problem of template-less calibration is to obtain the validity for the given pair of intrinsics and the reconstruction. In other words, if one is given all reconstructions from all possible focal lengths, it is not trivial to know the correct reconstruction. Especially when reconstructing using overestimated intrinsics with MDH, K allows the average depths to dominate the objective, while preserving the isometry. This usually leads to a flat and small scaled reconstruction [17]. Therefore an overestimated guess \hat{K} favors its own reconstruction over any upgraded one, while minimizing $\Phi(K)$. Relying on this observation, *we seek for the isometrically consistent reconstruction with the smallest focal length*, which works very well in practice. An algebraic analysis of our reasoning is provided in the supplementary material.

While searching for focal length, we use a sweeping procedure. On the one hand, if a reconstruction with the given focal length does not favor any upgrade, the sweeping is performed towards the lower focal length with a predefined step size, unless it starts favoring the upgrade. On the other hand, if the reconstruction favors the upgrade, we follow the suggested focal length update, until it suggests no more upgrade. The sought focal length is the one below which the upgrade is favorable, whereas above which it is not. Let $\delta(K_1, K_2)$ be gap in focal lengths of two intrinsics K_1 and K_2 , ΔK be a small step size which when added to an intrinsic matrix K increases its focal length by that step size. Our template-less calibration method is summarized in Algo. 2.

Algorithm 2 $[K] = \text{calibrateWithoutTemplate}(\hat{K})$

0. Set sweep direction $flag = 0$.
 1. Reconstruct 3D using (4) for the guess \hat{K} .
 2. Starting from \hat{K} , minimize $\Phi(K)$ in (13) to obtain K^* .
 3. IF $\delta(K^*, \hat{K}) \leq \epsilon$,
 - IF $flag == 0$, set $\hat{K} = K^* - \Delta K$ and goto step 1.
 - ELSE, return K^* .
 - ELSE, set and $flag = 1$, $\hat{K} = K^*$ and goto step 1.
-

We show in the experiment section, that the Algo. 2 converges in very few iterations. In every iteration, beside the reconstruction itself, the major computation is only required while minimizing $\Phi(K)$. Recall that, $\Phi(K)$ is minimized iteratively using a local method. During local search, the reconstruction for every update is required to compute $\Phi(K)$. Thanks to the upgrade equation, the cost $\Phi(K)$ can be computed instantly, without going through the computationally expensive reconstruction process.

3.3 Intrinsic Recovery in Practice

Although our reconstruction method makes inextensible shape assumption, the upgrade strategies use the piece-wise rigidity constraint. Despite the fact that the piece-wise rigid assumption is mostly true for inextensible shapes, it could be problematic in certain cases, for example, when the reconstructed points are too sparse. Therefore, some special care need to taken for a robust calibration.

Distance normalization and geodesics: Recall that the upgrade equation (5) is an approximation under the assumption that either the neighboring image points are sufficiently close to each other or a good guess \hat{K} is provided. When neither of these conditions are satisfied, the intrinsic obtained from energy minimization may not be sufficiently accurate. While a larger focal length may reduce the residual error, it also reduces individual distances creating disparities in the reconstruction scale of different views. Therefore, during each iteration of refinement, we fix the scale by enforcing,

$$\sum_i \sum_{j \in \mathcal{N}(i)} \hat{d}_{ij}^l(K) = 1, \forall l. \quad (14)$$

Another important practical aspect here is the use of geodesics $\hat{g}^l(i, j)$ instead of \hat{d}_{ij}^l in Eq. (13) or Eq. (9). When the scene points are sparse, using geodesics instead of the local euclidean distances may be necessary. We therefore choose to use geodesics computed from Dijkstra’s algorithm [33] instead of the local euclidean distances for stability.

Re-reconstruction and re-calibration: For a better calibration accuracy, especially when the initial guess \hat{K} is largely inaccurate, we iteratively perform re-reconstruction and re-calibration, starting from newly estimated intrinsic, until convergence. This has already been included in Algo. 2, which we also included on top of Algo. 1 in our implementation. In practice, only a few such iterations are sufficient to converge, even when the initial guess on intrinsic is very arbitrary.

4 Incremental Semi-dense NRSfM

The SOCP problem of (4) has the time complexity of $O(n^3)$. Therefore in practice, only a sparse set of points can be reconstructed in this manner. Here, we present a method to iteratively densify the initial sparse reconstruction, followed by online new view/camera addition strategy. Besides many obvious importance of incremental reconstruction, it is also necessary in our context: (a) to allow the selection of the closest image point pairs for camera calibration, (b) to compute 3D Geodesic distances for single view reconstruction.

4.1 Adding New Points

Let \mathcal{P} represents a set of sparse points reconstructed using (4). We would like to reconstruct a set of new points \mathcal{Q} with depths ζ_i^l , such that $\mathcal{Q} \cap \mathcal{P} = \emptyset$, consistent to

the existing reconstruction. This can be achieved by solving the following convex optimization problem,

$$\begin{aligned}
& \max_{\zeta_i^l, e_{ij}, \alpha} \quad \alpha \Lambda + \sum_l \sum_{i \in \mathcal{Q}} \zeta_i^l, \\
& \text{s.t.} \quad \left\| \hat{\mathbf{K}}^{-1}(\zeta_i^l \mathbf{u}_i^l - \alpha \lambda_j^l \mathbf{u}_j^l) \right\| \leq e_{ij}, \quad j \in \mathcal{N}_p(i), \\
& \quad \left\| \hat{\mathbf{K}}^{-1}(\zeta_i^l \mathbf{u}_i^l - \zeta_j^l \mathbf{u}_j^l) \right\| \leq e_{ij}, \quad j \in \mathcal{N}_q(i), \\
& \quad \sum_i \sum_{j \in \mathcal{N}_q(i)} e_{ij} = 1 - \alpha,
\end{aligned} \tag{15}$$

where, $\Lambda = \sum_l \sum_{i \in \mathcal{P}} \lambda_i^l$, $\mathcal{N}_p(i) = \mathcal{N}(i) \cap \mathcal{P}$, and $\mathcal{N}_q(i) = \mathcal{N}(i) \cap \mathcal{Q}$. The scalars α and $1 - \alpha$ represent the contributions of initial reconstruction \mathcal{P} and new reconstruction \mathcal{Q} , respectively. Note that the newly reconstructed points respect the inextensible criteria not only among themselves but also with respect to the initial reconstruction. This maintains the consistency between reconstructions \mathcal{P} and \mathcal{Q} . The incremental dense reconstruction process iteratively adds disjoint sets $\mathcal{Q}_1, \mathcal{Q}_2, \dots, \mathcal{Q}_r$ to the initial reconstruction \mathcal{P} , where \mathcal{P} encodes the overall shape and \mathcal{Q}_r represents the details.

4.2 Adding New Cameras

Adding a new camera to the NRSfM reconstruction is fundamentally a template-based reconstruction problem. If the camera is calibrated, one can obtain the reconstruction directly from (2). For the uncalibrated case, the camera can be calibrated first using (10), and the reconstruction upgraded from (4) using (5). It is important to note that the computation of accurate template geodesic distances d_{ij} , as required for template-based reconstruction, is possible only when the reconstruction is dense enough. This is not really a problem, thanks to the proposed incremental reconstruction method.

5 Discussion

Initial guess $\hat{\mathbf{K}}$: In all our experiments, we choose the initial guess $\hat{\mathbf{K}}$ by setting both focal lengths to the half of the mean image size and principal point to the image center.

Missing features: Feature points may be missing from some images due to occlusion or matching failure. This problem can be addressed during reconstruction by discarding all the variables corresponding to missing points together with all the inextensible constraints involving them as done in [19, 26].

Reconstruction Consistency: Alternative to (15), one can also think of reconstructing two overlapping sets \mathcal{P} and \mathcal{Q} such that $\mathcal{P} \cap \mathcal{Q} = \mathcal{R}$ independently. Then, the registration between them can be done with the help of \mathcal{R} from two sides. However, this is not only computationally inefficient due to the overlap, but also geometrically inconsistent.

6 Experimental Results

We conduct extensive experiments in order to validate the presented theory and to evaluate the performance, run time and practicality of the proposed methods.

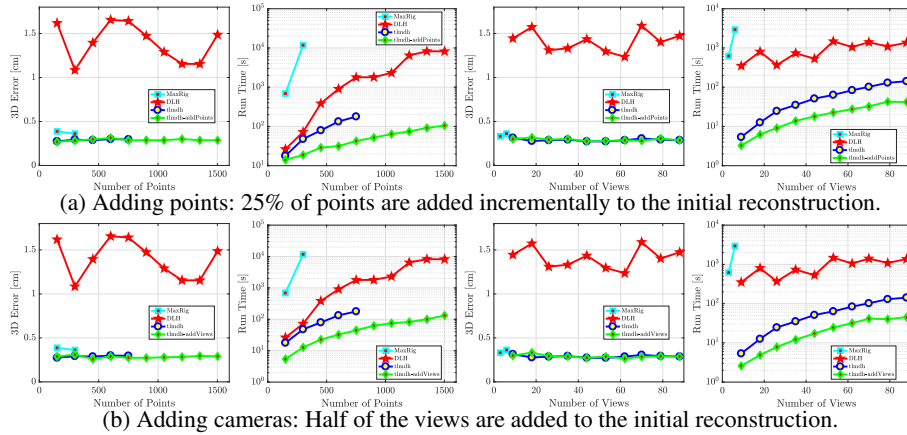


Fig. 2: **Incremental Semi-dense NRSfM.** Comparison of reconstruction error and run time on the *Hand* dataset. Left: varying number of points (number of views=88). Right: varying number of views (number of points=751). Run time shown in log scale.

Datasets. We first provide a brief descriptions of the datasets we use to analyze our algorithms. **KINECT Paper.** This VGA resolution image sequence shows a textured paper deforming smoothly [34]. The tracks contain about 1500 semi-dense but noisy points. **Hulk & T-Shirt.** The datasets contain a comic book cover in 21 different deformations, and a textured T-Shirt with 10 different deformations [35], in high resolution images. Although the number of points is low (122 and 85, resp.), the tracks have very little noise and therefore we obtain a very accurate auto-calibration. **Flag.** This semi-synthetic dataset is created from mocap recordings of deforming cloth [36]. We generate 250 points in 30 views using a virtual 640x480 perspective camera. **Newspaper.** This sequence² contains the deformation and tearing of a double-page newspaper, recorded with KINECT in HD resolution [19]. **Hand.** The Hand dataset [19] features medium resolution images. Dense tracking [37] of image points yield up to 1500 tracks in 88 views. The dataset consists of ground-truth 3D for the first and the last image of the sequence. **Minion & Sunflower.** These sequences are recorded with a static Kinect sensor [38]. Minion contains a stuffed animal undergoing folding and squeezing deformations. Sunflower however features only small translation w.r.t. the camera. We incrementally reconstruct more than 10,000 points for Minion, and 5,000 for Sunflower, as shown in Fig. 1. We are able to reconstruct the global deformation, and mid-level details such as the glasses of Minion. Unfortunately, due to the failure of optical flow tracking, we fail to reconstruct homogeneous areas and fine details. In Sunflower we can capture the deformation of the outside leaves, whereas finer details in the center of the blossom is not recovered due to insufficient change in viewpoint. **Camel³ & Kitten⁴.** We took two sequences from YouTube videos to show the incremental semi-dense NRSfM from

²The dataset was provided by the authors.

³<https://www.youtube.com/watch?v=PhpeadpZsa4>

⁴<https://www.youtube.com/watch?v=DIZM2OMNc7c>

uncalibrated cameras. The camel turns around its head towards the moving camera, providing enough motion to faithfully reconstruct the 3D motion of the animal. Fig. 1 shows the 3D structure of more than 3,000 points for one out of 61 views reconstructed. In the Kitten sequence (18,000 points for each 36 views), a cat performs both articulated and deforming motion with body and tail. Again, state-of-the-art optical flow methods struggle to maintain stable points tracks, especially on the head. Nevertheless, our method captures the general motion to a very good extent. In all of the above datasets, **DLH** fails to get the correct shape while **MaxRig** cannot reconstruct the shape faithfully as it cannot handle enough points. **Cap**. This dataset contains wide-baseline views of a cap in two different deformations [18]. The 3D template of the undeformed cap was obtained using SfM pipeline for the images from the first camera. Then, the second camera is calibrated using our template-based method.

6.1 Camera Calibration from a Non-rigid Scene

To measure the quality of our calibration results, we report the 3D root mean square error (RMSE), the relative focal length and principal point estimation error. Furthermore, we provide the number of iterations and the corresponding run times in Table 1.

Dataset	Number of		Run time [s]			Focal Estimation				Reconstruction Error E_{rec}			
	Points	Views	T_{orig}	N_{iter}	T_{total}	f_{init}	f_{GT}	f_{est}	Error %	f_{GT}	f_{est}		
Template-based focal length estimation													
KINECT Paper	301	23	2.3	-	16.8	-	528	590	11.74	3.00	0.54%	2.83	0.50%
Hulk	122	21	0.4	-	4.2	-	3784	4300	13.61	5.73	1.43%	5.53	1.37%
Flag	250	30	1.3	-	178.2	-	384	420	9.38	4.74	0.58%	4.54	0.56%
Cap	137	1	0.3	-	11.0	-	2039	2300	12.8	1.13	4.80%	1.13	4.80%
Template-less focal length estimation													
KPaper	301	23	5.8	3	110.1	280	528	540	2.27	4.44	0.80%	4.28	0.77%
Hulk	122	21	1.9	5	36.5	1641	3784	3800	0.40	2.76	0.67%	2.75	0.66%
T-Shirt	85	10	0.6	10	24.1	2000	3787	4000	5.63	3.52	1.10%	3.42	1.07%
Flag	250	30	2.6	6	185.4	280	384	400	4.17	5.24	0.64%	5.05	0.62%
Newspaper	441	19	24.5	5	523.6	750	1055	870	16.6	7.79	1.09%	9.27	1.30%

Table 1: **Focal Length Estimation from a Non-Rigid Scene.** We report the run-time, reconstruction error and relative focal length estimation error of our template-based and template-less NRSfM calibration methods. T_{orig} is the time needed to reconstruct with a given focal length, T_{total} the run time including calibration. For the template-less case, N_{iter} iterations were performed until convergence.

Template-based Camera Calibration. In the first part of Algo. 1 we generate hypotheses for K and choose the one with best isometric match with the template. We perform experiments on the KINECT Paper, Hulk and Flag dataset and report the results in Table 1. We observe a consistent improvement in reconstruction accuracy with the estimated intrinsics. The second part of Algo. 1 involves gradient-based refinement on the intrinsics by minimizing Eq. (12). To analyze this part, we conduct two experiments: First, we perform refinement on the initially estimated intrinsics f_{poly} . Here we can consistently improve reconstruction errors with the refined intrinsics. In the Hulk and Flag dataset, we also get a better estimate of the focal length. On KINECT Paper however, the focal length deteriorates, while reconstruction accuracy improves. This is most probably due to the noisy tracks in the sequence. Due to the effective regularization, the error in principal point is consistently low. In the second experiment, we gauge

the robustness of our refinement method. To this end, we simulate initial intrinsics by adding $\pm 20\%$ uniform noise independently on each of the entries of K_{GT} , and compare reconstruction error and the refined intrinsics shown in Table 2. We compare to Bartoli et al. [18] on the Cap dataset directly from the paper, since it is non-trivial to implement the method itself. We observed an error E_f of about 13% with our method, compared to 3.8%-7.3% reported by [18]. The slightly higher error in the Cap dataset can be partly attributed to the repeating texture that makes our image matches non-ideal. Overall we can observe a consistent improvement in almost all metrics, validating the robustness of the method and the assumptions it is based on.

Dataset	Template-based				Refined				Simulated initial K (10^3 samples avg.)					
	f_{GT}	f_{poly}	E_f	E_{rec}	f_{ref}	E_f	E_{PP}	E_{rec}	E_f	ΔE_f	E_{PP}	ΔE_{PP}	E_{rec}	ΔE_{rec}
KINECT Paper	528	590	11.74	2.83	604	14.45	0.04	2.73	8.87	-0.37	0.05	-10.96	3.82	-0.25
Hulk	3784	4300	13.61	5.53	4119	8.85	1.74	5.53	7.30	-2.36	1.77	-8.84	6.52	-0.01
Flag	384	420	9.38	4.54	414	7.98	0.05	4.34	8.61	-1.05	0.08	-10.45	6.05	+0.08
Cap	2039	2300	12.8	1.13	2360	13.1	2.33	1.13	9.18	-0.10	2.33	-8.42	1.48	-0.00

Table 2: **Calibration Refinement.** We compute the full calibration K by initializing with the template-based calibration f_{poly} , and test the robustness by adding synthetic noise on the K_{GT} . Reconstruction errors E_{rec} are in mm, others in %.

Template-less Camera Calibration. To visualize the dynamics of Algo. 2, we plot the error in isometry $\Phi(K)$ over focal length for each iteration on the Hulk dataset in Fig. 3 (a). Typically, less than 10 iterations are necessary for the method to converge. As we hypothesized above, Fig. 3 (b) empirically verifies that we can find the termination criterion for our sweeping strategy by thresholding the focal length change $\delta(K^*, \hat{K})$. Our method consistently recovers a correct estimate of the intrinsics as reported in Table 1. Moreover, the fact that we obtain better reconstruction accuracy in almost all datasets validates our approach of using the isometric consistency $\Phi(K)$.




Fig. 3: **Template-less Calibration (Algo. 2).** We iteratively search the smallest K that maximizes isometry.

6.2 Incremental Reconstruction

We first present experiments on the dense Hand dataset in Fig. 2. We compare to two state-of-the-art NRSfM approaches, **MaxRig** [6] and **DLH** [7], as well as the to batch version of our approach **tmdh** [19]. In the first row, we plot the performance of **tmdh-addPoints**: we start by reconstructing a random subset of $\max\{150, \frac{N}{4}\}$ points, and

Datasets	incr-tlmdh		tlmdh		p-isomet	p-isolh	DLH	o-kfac
KPaper	4.64	176.16s	5.41	605.06s	7.63	13.64	14.66	13.93
Hulk	2.99	0.80s	2.76	1.99s	10.76	14.54	22.98	-
T-Shirt	3.83	0.23s	3.53	0.47s	10.60	8.94	-	-
Cardboard	13.22	18.94s	14.56	34.35s	-	12.95	-	-
Rug	26.40	205.89s	26.60	542.39s	26.15	38.26	31.01	-
Table mat	15.99	5.54s	14.36	7.65s	14.21	20.71	17.51	16.24
Newspaper	10.79	89.27s	11.63	190.96s	18.40	37.21	24.94	30.74

Table 3: **Comparison of NRSfM methods.** Mean 3D errors in mm and run time comparison for batch and incremental reconstruction in real datasets.

incrementally add the remaining points in subsequent iterations according to Eq. (15). While achieving competitive reconstruction accuracy on par with **tlmdh**, we observe remarkable advantages in run time compared to all other methods. **MaxRig** shows good accuracy, but suffers from serious run time and memory problems. **DLH** on the other hand is slow and exhibits poor accuracy on this dataset, due to perspective and non-linear deformations. The second row of Fig. 2 shows the same experimental setup with **tlmdh-addViews**. Here, we reconstruct all points at once, but incrementally add the remaining 50% of views to the reconstruction of the first half. To this end, we compute the template from the first reconstruction and employ SfT. The graphs clearly show that **tlmdh-addViews** exhibits a favorable run time complexity without impairing the reconstruction accuracy. We provide more results in the supplementary material. Furthermore, we perform extensive experiments on a variety of additional datasets, and compare with the reconstructions of **p-isomet** [28], **p-isolh** [35], **DLH** [7], and **o-kfc** [39] in Table 3 obtained from [40]. Overall, we observe a significant advantage in accuracy and run time in particular compared to the best performing baseline **tlmdh**.

7 Conclusions

In this paper we formulated a method addressing the unknown focal-length in NRSfM and unknown intrinsics in SfT. Despite the computational complexity of convex NRSfM, we formulated an incremental framework to obtain semi-dense reconstruction and reconstruct new views. We developed our theory based on the surface isometry prior in the context of the perspective camera. We developed and verified our approach for intrinsics/focal-length recovery for both template-based and template-less non-rigid reconstruction. Essential to our method is a novel upgrade equation, that analytically relates reconstructions for different intrinsics. We performed extensive quantitative and qualitative analysis of our methods on different datasets which shows the proposed methods perform well despite addressing very challenging problems.

Acknowledgements

Research was funded by the EU’s Horizon 2020 programme under grant No. 645331–EurEyeCase and grant No. 687757–REPLICATE, and the Swiss Commission for Technology and Innovation (CTI, Grant No. 26253.1 PFES-ES, EXASOLVED).

References

1. Longuet-Higgins, H.: A computer algorithm for reconstructing a scene from two projections. *Nature* **293** (1981) 133–135 [1](#)
2. Nistér, D.: An efficient solution to the five-point relative pose problem. *IEEE Trans. Pattern Anal. Mach. Intell.* **26**(6) (2004) 756–777 [1](#)
3. Hartley, R.I., Zisserman, A.: *Multiple View Geometry in Computer Vision*. Second edn. Cambridge University Press, ISBN: 0521540518 (2004) [1](#)
4. PhotoScan, A.: *Agisoft PhotoScan User Manual Professional Edition, Version 1.2.* (2017) [1](#)
5. ReCap, A.: *ReCap 360 – Advanced Workflows.* (2015) [1](#)
6. Ji, P., Li, H., Dai, Y., Reid, I.: "Maximizing Rigidity" revisited: A convex programming approach for generic 3d shape reconstruction from multiple perspective views. In: *ICCV.* (2017) [2](#), [3](#), [13](#)
7. Dai, Y., Li, H., He, M.: A simple prior-free method for non-rigid structure-from-motion factorization. In: *CVPR.* (2012) [2](#), [13](#), [14](#)
8. Bregler, C., Hertzmann, A., Biermann, H.: Recovering non-rigid 3D shape from image streams. In: *CVPR.* (2000) [2](#)
9. Torresani, L., Hertzmann, A., Bregler, C.: Nonrigid structure-from-motion: Estimating shape and motion with hierarchical priors. *IEEE Trans. Pattern Anal. Mach. Intell.* **30**(5) (2008) 878–892 [2](#)
10. Del Bue, A.: A factorization approach to structure from motion with shape priors. In: *CVPR.* (2008) [2](#)
11. Garg, R., Roussos, A., Agapito, L.: Dense variational reconstruction of non-rigid surfaces from monocular video. In: *CVPR.* (2013) [2](#)
12. Fayad, J., Agapito, L., Del Bue, A.: Piecewise quadratic reconstruction of non-rigid surfaces from monocular sequences. In: *ECCV.* (2010) [2](#)
13. Agudo, A., Montiel, J., Agapito, L., Calvo, B.: Online dense non-rigid 3d shape and camera motion recovery. In: *BMVC.* (2014) [2](#)
14. Taylor, J., Jepson, A.D., Kutulakos, K.N.: Non-rigid structure from locally-rigid motion. In: *CVPR.* (2010) [2](#)
15. Bartoli, A., Pizarro, D., Collins, T.: A robust analytical solution to isometric shape-from-template with focal length calibration. In: *ICCV.* (2013) [2](#)
16. Ngo, T.D., Östlund, J.O., Fua, P.: Template-based monocular 3D shape recovery using laplacian meshes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **38**(1) (2016) 172–187 [2](#)
17. Chhatkuli, A., Pizarro, D., Bartoli, A., Collins, T.: A stable analytical framework for isometric shape-from-template by surface integration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **39**(5) (2017) 833–850 [2](#), [8](#)
18. Bartoli, A., Collins, T.: Template-based isometric deformable 3D reconstruction with sampling-based focal length self-calibration. In: *CVPR.* (2013) [2](#), [12](#), [13](#)
19. Chhatkuli, A., Pizarro, D., Collins, T., Bartoli, A.: Inextensible non-rigid shape-from-motion by second-order cone programming. In: *CVPR.* (2016) [2](#), [3](#), [4](#), [10](#), [11](#), [13](#)
20. Perriollat, M., Hartley, R., Bartoli, A.: Monocular template-based reconstruction of inextensible surfaces. *International Journal of Computer Vision* **95**(2) (2011) 124–137 [2](#), [3](#)
21. Salzmann, M., Fua, P.: Linear local models for monocular reconstruction of deformable surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **33**(5) (2011) 931–944 [2](#), [3](#), [4](#), [6](#)
22. Kumar, S., Dai, Y., Li, H.: Monocular dense 3d reconstruction of a complex dynamic scene from two perspective frames. In: *ICCV.* (2017) [2](#), [3](#)

23. Russell, C., Yu, R., Agapito, L.: Video pop-up: Monocular 3d reconstruction of dynamic scenes. In: ECCV. (2014) 2, 3
24. Xiao, J., Kanade, T.: Uncalibrated perspective reconstruction of deformable structures. In: Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1. Volume 2. (2005) 1075–1082 Vol. 2 2, 3
25. Bartoli, A., Gérard, Y., Chadebecq, F., Collins, T., Pizarro, D.: Shape-from-template. IEEE Trans. Pattern Anal. Mach. Intell. **37**(10) (2015) 2099–2118 3
26. Vicente, S., Agapito, L.: Soft inextensibility constraints for template-free non-rigid reconstruction. In: ECCV. (2012) 3, 10
27. Chhatkuli, A., Pizarro, D., Bartoli, A.: Stable template-based isometric 3D reconstruction in all imaging conditions by linear least-squares. In: CVPR. (2014) 3
28. Parashar, S., Pizarro, D., Bartoli, A.: Isometric non-rigid shape-from-motion in linear time. In: CVPR. (2016) 3, 14
29. Salzmann, M., Hartley, R., Fua, P.: Convex optimization for deformable surface 3-D tracking. In: ICCV. (2007) 3
30. Akhter, I., Sheikh, Y., Khan, S., Kanade, T.: Trajectory space: A dual representation for nonrigid structure from motion. IEEE TPAMI **33**(7) (2011) 1442–1456 3
31. Nistér, D.: Untwisting a projective reconstruction. International Journal of Computer Vision **60**(2) (2004) 165–183 5, 7
32. Chandraker, M., Agarwal, S., Kahl, F., Nistér, D., Kriegman, D.: Autocalibration via rank-constrained estimation of the absolute quadric. In: Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on, IEEE (2007) 1–8 7
33. Dijkstra, E.W.: A note on two problems in connexion with graphs. Numer. Math. **1**(1) (1959) 269–271 9
34. Varol, A., Salzmann, M., Fua, P., Urtasun, R.: A constrained latent variable model. In: CVPR. (2012) 11
35. Chhatkuli, A., Pizarro, D., Bartoli, A.: Non-rigid shape-from-motion for isometric surfaces using infinitesimal planarity. In: BMVC. (2014) 11, 14
36. White, R., Crane, K., Forsyth, D.: Capturing and animating occluded cloth. In: SIGGRAPH. (2007) 11
37. Sundaram, N., Brox, T., Keutzer, K.: Dense point trajectories by gpu-accelerated large displacement optical flow. In: ECCV. (2010) 11
38. Innmann, M., Zollhöfer, M., Nießner, M., Theobalt, C., Stamminger, M. In: VolumeDeform: Real-Time Volumetric Non-rigid Reconstruction. Springer International Publishing, Cham (2016) 362–379 11
39. Gotardo, P.F., Martinez, A.M.: Computing smooth time trajectories for camera and deformable shape in structure from motion with occlusion. IEEE Trans. on Pattern Analysis and Machine Intelligence **33**(10) (2011) 2051–2065 14
40. Chhatkuli, A., Pizarro, D., Collins, T., Bartoli, A.: Inextensible non-rigid structure-from-motion by second-order cone programming. IEEE Transactions on Pattern Analysis and Machine Intelligence **PP**(99) (2017) 1–1 14