Unsupervised Monocular Depth Reconstruction of Non-Rigid Scenes

Conference Paper**Author(s):**

Takmaz, Ayça ; Paudel, Danda Pani; Probst, Thomas ; Chhatkuli, Ajad ; Oswald, Martin R.; Van Gool, Luc

Publication date:

2021

Permanent link:

<https://doi.org/10.3929/ethz-b-000523758>

Rights / license:

In Copyright - Non-Commercial Use Permitted

Originally published in:

<https://doi.org/10.1109/3DV53792.2021.00091>

Funding acknowledgement:

820434 - ENergy aware BIM Cloud Platform in a COst-effective Building RENovation Context (EC)

Unsupervised Monocular Depth Reconstruction of Non-Rigid Scenes

Ayça Takmaz², Danda Pani Paudel¹, Thomas Probst¹, Ajad Chhatkuli¹, Martin R. Oswald^{2,3}, Luc Van Gool^{1,4}

¹Computer Vision Lab, ETH Zurich

²Department of Computer Science, ETH Zurich

³University of Amsterdam, Netherlands

⁴VISICS, ESAT/PSI, KU Leuven, Belgium

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

Abstract

Monocular depth reconstruction of complex and dynamic scenes is a highly challenging problem. While for rigid scenes learning-based methods have been offering promising results even in unsupervised cases, there exists little to no literature addressing the same for dynamic and deformable scenes. In this work, we present an unsupervised monocular framework for dense depth estimation of dynamic scenes, which jointly reconstructs rigid and non-rigid parts without explicitly modelling the camera motion. Using dense correspondences, we derive a training objective that aims to opportunistically preserve pairwise distances between reconstructed 3D points. In this process, the dense depth map is learned implicitly using the as-rigid-as-possible hypothesis. Our method provides promising results, demonstrating its capability of reconstructing 3D from challenging videos of non-rigid scenes. Furthermore, the proposed method also provides unsupervised motion segmentation results as an auxiliary output.

1. Introduction

Understanding the 3D structure of a scene can provide important cues for many tasks such as robot navigation [19], motion capture [42], scene understanding [33], and augmented reality [28]. While humans are exceptionally capable of inferring the non-rigid 3D structures from an image, geometric computer vision techniques either require a large amount of labeled data or are capable of learning only from rigid scenes [84]. However, in the real world, scenes often consist of non-rigid and dynamic elements. Thus, it is quite natural to seek for the ability of inferring the depth from an image of any given scene, regardless of whether it is highly dynamic or not. We refer to Fig. 1 for some examples.

The ability of humans to understand their environment geometrically and semantically is mainly acquired through childhood learning. In an attempt to emulate this ability, learning-based strategies have been applied to the problem of *depth from single-view*. In fact, many learning-based methods already offer very promising progress in this di-

Figure 1: **Unsupervised Depth.** Image-depth pairs; depth is estimated in an unsupervised manner using the proposed method.

rection. Among them, supervised methods [40, 12, 37] aim to reconstruct the depth of rigid and non-rigid parts during training, whereas many unsupervised methods [23, 84, 83, 25, 13, 45, 73, 43, 44, 26] mainly employ a training strategy which aims to reconstruct the rigid parts of a scene. Although a recent unsupervised method [39] also reconstructs objects translating on a ground plane, it has limitations in terms of modelling highly non-rigid scenes. Other unsupervised methods for non-rigid reconstruction exploit object-specific priors [59, 77, 80], or reconstruct sparse [49] or dense [64, 66] points of a single non-rigid object. Such methods, however, do not have the same applicability or motivation as that of single view scene depth estimation. Unlike the monocular setup, the calibrated stereo (or multi-camera) methods for learning single view depth [24, 48] do not suffer from ambiguities due to non-rigidity and can handle complete scene depth. However, they have significant limitations, particularly when acquiring training data via calibrated stereo pairs is not feasible.

In this work, we are interested in learning to predict dense depth from a single image using an unsupervised monocular pipeline. More importantly, we would like to reconstruct the depth of all parts of the scene during the training, irrespective of the rigidity/non-rigidity of these parts. Our unsupervised setup assumes that only calibrated monocular videos¹ with known intrinsics are available during training. Such an assumption is realistic as well as crucial for a wide variety of setups, ranging from consumer to surgical cameras, where depth or stereo acquisition for

¹Our method can potentially also be used for multi-view setups.

supervision is often impractical. In this context, learning depth from monocular videos of non-rigid scenes remains an unresolved problem. This is no surprise, given the challenges, *e.g.*, ill-posedness, ambiguities, inconsistent priors.

The success of the unsupervised depth learning methods for rigid scenes can be primarily attributed to the advancements in deep learning and the rigid reconstruction constraints used in such methods. This motivates us to explore the non-rigid 3D reconstruction literature employing various assumptions. Our goal is to keep in mind an overview of the literature and to use the gained insights to build our unsupervised monocular pipeline for depth reconstruction in non-rigid scenes. Our main **contributions** are threefold:

- We reformulate the Non-Rigid Structure-from-Motion (NRSfM) priors in a novel unified framework using the Euclidean distance matrix measures across views. This contribution is summarized in Table 1.
- The utility of the proposed framework is also demonstrated for unsupervised non-rigid monocular depth, by exploiting the as-rigid-as possible (ARAP) prior during CNN training. For the implementation of the ARAP prior, we define and employ a concept of *motion embeddings*. This contribution is illustrated in Fig. 3.
- Through experiments, we provide interesting new insights towards learning non-rigid scene depths in an unsupervised manner, detailed in our discussion section.

2. Non-Rigid Reconstruction Revisited

Since generic non-rigid 3D reconstruction from a monocular camera is an ill-posed problem, methods in the literature rely on some priors or assumptions about the scene. The most common scene priors can be broadly divided into the following four categories.

Low-Rank (LR). The low-rank prior assumes that non-rigid 3D structures can be expressed as a linear combination of finite basis shapes. The landmark work of *Costeira & Kanade* [16] developed for orthographic cameras (or slight variations) uses the LR prior for multi-body 3D scene reconstruction. Since, it been widely used in various cases for the 3D reconstruction of single [10, 9, 8, 50, 17, 21, 32, 70, 2, 49, 35] and multiple [79, 3] non-rigid objects.

Scene Motion (SM). Several notable works of *Shashua et al.* [62, 7, 76, 72] have shown that the known planar/linear motion prior can be exploited to reconstruct the 3D structure of dynamic scenes². Work of *Ozden et al.* [51] also tackled dynamic scene 3D reconstruction using the principle of non-accidentalness. An insightful work of *Hartley and Vidal* [27] reveals that the method of [76] can be indeed extended to the generic non-rigid case, using the low-rank structure and without an explicit motion prior, with only one

²The recent work [39] can be seen as an adaptation of [7].

severe ambiguity, which stems from the fact that the recovered camera motion is relative to the moving points. This in turn implies the unfavorable news that the low-rank prior alone is not sufficient to recover scale-consistent non-rigid 3D structures from monocular projective cameras, when the objects are independently moving in the multi-body setting.

Isometric Deformation (ID). To avoid the algebraic prior of LR, *Salzmann et al.* [61] and *Perriollat et al.* [57] introduced the geometric ‘ID’ prior of an object deforming isometrically, which makes non-rigid surface reconstruction possible when used with a known template. Since, the ID prior of the object has been used in many other works [63, 71, 14, 53, 15, 74, 58, 38]. *Taylor et al.* [68] introduced local-rigidity for non-rigid reconstruction, which is a form of the ID prior. Other related geometric priors have also been exploited in the literature [46, 60, 1]. When objects undergo severe deformations, relying on a geometric prior for the objects and avoiding any explicit camera motion estimation have demonstrated state-of-the-art results for non-rigid object reconstruction [55, 52, 58, 22, 47]. The benefit of not estimating the camera motion explicitly can be understood quite intuitively, as the estimated rigid camera motion is merely a motion with respect to the scene. Interested readers may recall the relevant work in [27]. In this regard, the works of *Li* [38], *Ji et al.* [31] and *Chhatkuli et al.* [15] stand out. While reconstructing inextensible structures, [15] formulates an explicit-motion-free problem, which turns out to be equivalent to the rigid case formulation of [38] in the absence of deformations. In fact, the formulation of [15] is shown to be effective for both rigid and non-rigid scenes. All three works [15, 31, 38] are based on the pairwise distance equality in 3D, while [38] uses the same pairwise sampling as [31].

As-Rigid-As-Possible (ARAP). Non-rigid shape modelling using the ARAP assumption is very prevalent in computer graphics [6, 29, 65]. The ARAP assumption maximizes rigidity while penalising stretching, shearing, and compression. ARAP concept has also been used in many works [54, 36] for monocular non-rigid 3D reconstruction. In particular, two works of *Parashar et al.* [54] and *Kumar et al.* [36] are noteworthy. Using the shape template of an object, [54] reconstructs the deformed volumetric 3D under the ARAP assumption. On the other hand, [36] demonstrates the practicality of ARAP in depth densification and refinement for non-rigid *scenes*, using multi-view setups. More interestingly to us, the ARAP prior is sufficient to resolve the scale ambiguity between freely moving parts, *i.e.* the issue previously discussed by *Hartley & Vidal* in [27].

In summary, the computation of camera motion becomes ambiguous in complex non-rigid scenes for a monocular camera. Nevertheless, the structure of both rigid scenes and non-rigid objects³ can still be recovered without any explicit

³Given that the assumed object’s prior is sufficient.

Figure 2: **Pairwise distance formulation.** The pair of 3D Euclidean distances in k -th and l -th views are given by $d_{ij}^k = \|\mathbf{x}_i^k - \mathbf{x}_j^k\|$ and $d_{ij}^l = \|\mathbf{x}_i^l - \mathbf{x}_j^l\|$, respectively. Every 3D point, say $\mathbf{x}_i^k = \lambda_i^k \mathbf{u}_i^k$, is expressed using the unknown depth and known homogeneous image coordinates, i.e. λ_i^k and $\mathbf{u}_i^k$.

camera motion estimation. On the other hand, the scale consistent reconstruction intrinsically requires some additional prior, such as the ARAP assumption used in this work.

3. Preliminaries

Notations. We denote matrices with uppercase and their elements with double-indexed lowercase letters: $\mathbf{A} = (a_{ij})$. Similarly, we write vectors and index them as: $\mathbf{a} = (a_i)$. The inequality $\mathbf{A} > 0$ refers to $(a_{ij} > 0)$, unless mentioned otherwise. We use special uppercase Latin or Greek letters for sets and graphs, such as $\mathcal{S}$ and $\mathcal{G}$. The lowercase Latin letters, as in a , are used for scalars. The set of neighbours of i within a radius r is given by $\mathcal{N}_r(i)$.

Non-rigid Structure-from-Motion (NRSfM). Our problem formulation is camera motion independent, similar⁴ to [38, 61, 15]. We pose the NRSfM problem as finding point-wise depth in each view. We use superscript k to denote the k -th image (among m) and subscript i to denote the i -th point (among n). As illustrated in Fig. 2, we represent the unknown depth as λ_i^k and the known homogeneous image coordinates as $\mathbf{u}_i^k$. The set of points in the k -th view is given as $\mathcal{X}^k = \{\mathbf{X}_i^k | \forall i \in 1 \dots n\}$. The Euclidean distance between point $\mathbf{X}_i^k$ and point $\mathbf{X}_j^k$ is denoted as d_{ij}^k . In our formulation, we use the Euclidean distance matrix:

Definition 3.1 (Euclidean distance matrix) *Euclidean Distance Matrix (EDM), say $\mathbf{E} \in \mathbb{R}^{n \times n}$, is a matrix representing the spacing of n points in 3-dimensional Euclidean space, say $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n] \in \mathbb{R}^{3 \times n}$, and the entries of $\mathbf{E}$ are given by, $e_{ij} = d_{ij}^2 = \|\mathbf{x}_i - \mathbf{x}_j\|^2$.*

Let $\Lambda^k = [\lambda_1^k, \lambda_2^k, \dots, \lambda_n^k] \in \mathbb{R}^n$, be the sought depth of the k -th view. We represent the 3D structure in the form, $\mathbf{X}^k(\Lambda^k) = [\lambda_1^k \mathbf{u}_1^k, \lambda_2^k \mathbf{u}_2^k, \dots, \lambda_n^k \mathbf{u}_n^k] \in \mathbb{R}^{3 \times n}$. Let us define the Gram matrix $\mathbf{G}(\Lambda^k) = (\mathbf{X}^k(\Lambda^k))^T \mathbf{X}^k(\Lambda^k) \in \mathbb{R}^{n \times n}$. The EDM for the k -th view is then given by:

$$\mathbf{E}(\Lambda^k) = \text{diag}(\mathbf{G}(\Lambda^k)) \mathbf{1}^T - 2\mathbf{G}(\Lambda^k) + \mathbf{1} \text{diag}(\mathbf{G}(\Lambda^k))^T. \quad (1)$$

⁴Similar also in mathematical terms.

NRSfM aims to estimate $\{\Lambda^k\}$ by imposing priors on $\mathbf{E}(\Lambda^k)$. For clarity, we present an example of the rigid case.

Example 3.2 (The rigid structure prior) *For noise- and outlier-free rigid scenes, the estimated $\{\Lambda^k, \Lambda^l\}$ must satisfy, $\mathbf{E}(\Lambda^k) - \mathbf{E}(\Lambda^l) = 0$, for all $k, l = 1, \dots, m$. Intuitively, the distance between any pair of points must be preserved across views. Method of [38] is a variant of this example.*

The prior of the above example is also referred as the *global rigidity* constraint [75, 20]. Weaker than the *global rigidity* is *local rigidity*. Interested readers can find a thorough study of local/global rigidity in [38]. Here, we are interested in the non-rigid isometric deformation prior [57, 68, 61, 15]. Let us introduce a weight matrix $\mathbf{W} \in \mathbb{R}^{n \times n}$, whose entries are $w_{ij} = 1$ if $j \in \mathcal{N}_r(i)$ for some radius r , and $w_{ij} = 0$ otherwise. Then the relaxed isometric prior of [61, 15, 58] aims to reconstruct the depth by solving,

$$\begin{aligned} \text{find } & \{\Lambda^k\}, \\ \text{s.t. } & \mathbf{W} \odot (\mathbf{E}(\Lambda^k) - \mathbf{E}(\Lambda^l)) = 0, \forall k, l, \\ & \Lambda^k > 0, \forall k. \end{aligned} \quad (2)$$

where $\odot$ represents the Hadamard-product. Intuitively, (2) aims to preserve the local Euclidean distances across views. The constraints $\Lambda^k > 0$ ensure positive depths in all views. It is important to note that for sufficiently large radii r the problem of (2) is equivalent to that of the Example 3.2.

4. This Work

One may use the formulation in Eq. (2) to learn to reconstruct both *rigid* scenes or non-rigid *objects*, for known $\mathbf{W}$ ⁵. We however, are interested in reconstructing *complex scenes* consisting of both, which requires additional priors (recall Section 2). Moreover, it is unclear how to obtain the weight matrix $\mathbf{W}$ in general cases. We use the ARAP assumption to address the scale ambiguity problem. In this process, the weight matrix is also learned along with the depth. If one aims to exploit other non-rigid priors, later in Section 4.2 we provide a thorough analysis in that direction.

4.1. Problem Formulation

We aim to estimate the depth Λ for each view, and the weight matrix $\mathbf{W}$ for given view-pairs. Our weights $w_{ij} \in [0, 1]$ can be interpreted as rigidity scores, between points $\mathbf{X}_i$ and $\mathbf{X}_j$. Since, the concept of rigidity is meaningful only for two (or more) views, our rigidity scores are computed accordingly. We consider two points to be rigidly connected, if their distance does not change across views. If one seeks for pair-wise rigidity on generic graphs, several interpretations can be derived based on the graph connectivity [75, 20, 38]. In the context of this paper, we formulate the ARAP assumption as follows:

⁵For example, [15] uses 2D neighbors to construct $\mathbf{W}$.

Definition 4.1 (As-rigid-as-possible) For a point set under deformation, the as-rigid-as-possible model assumes that every pair-wise distance of the fully connected graph between points, respects at least some degree of rigidity.

We can now formalize our problem statement as follows.

$$\begin{aligned} \min_{\{\Lambda^k\}, \{W^{kl}\}} \quad & \eta, \\ \text{s.t.} \quad & |W^{kl} \odot (E(\Lambda^k) - E(\Lambda^l))| \leq \eta \|W^{kl}\|_{1,1}, \quad (3) \\ & 1 \geq W^{kl} \geq \tau, \Lambda^k > 0, \forall k. \end{aligned}$$

Here, the positive scalars τ and η are the rigidity threshold and the rigidity adjusted maximum allowed distance error, respectively. Very often the point pairs from non-rigid objects respect rigidity. This is when many priors are best justified. In general, local rigidity does not imply global rigidity. For large fully connected graphs though, pair-wise rigidity means global rigidity. We relax this constraint by allowing different edges to have different rigidity scores. We like to draw the reader’s attention on two key (and somewhat related) aspects of our formulation: (1) image pair-wise rigidity scores; and (2) global connectivity.

Our motivation for using individual W^{kl} stems from the following observation: across all frame pairs, most of the edges respect rigidity at least once in the global sense (e.g. stopping objects or periodic motions). This global rigidity can be captured and propagated in the rest of the reconstruction for scale consistency, if the global connectivity is established using the fully connected graph. For local reconstruction, we can rely on local connectivity which often better respects rigidity. Note that the maximum allowed adjusted distance error, *i.e.* η , is measured in the normalized (by the weight matrix) form. This encourages rigid edges to bear higher weights, and vice versa. Such distribution of weights prioritizes the pairs holding rigidity to reconstruct rigidly. On the other hand, non-rigid pairs – if reconstructed correctly at their rigid instant (with highest weight) – can still satisfy the imposed constraints, as long as the extended distance error does not exceed η/τ .

4.2. Relation to Previous Works

Heretofore, we have discussed the relationship of our formulation in Eq. (3) to rigidity, ID, and ARAP priors. Now, we further explore the same with respect to LR and SM, for the completeness of our theoretical framework. To do so, let us denote $D^{kl} = E(\Lambda^k) - E(\Lambda^l)$ and its stacked version $\mathbf{D} = [(D^{12})^\top, \dots, (D^{kl})^\top, \dots, (D^{(m-1)m})^\top]^\top$. Our investigation leads to the following relationships between the formulation used in this paper and the prior of the low-rank purely on structures, without explicit camera motion.

Proposition 4.2 (EDM of Low-rank structures) For 3D structures that can be represented using b linear basis,

Method	Unsup.-Monoc.	Non-Rigid	Rigid prior	ARAP	Isom. deform.	Scene motion	Low rank	Constraints
Li [38]		✓						$w_{ij}^{kl} = 1$
Parashar et al. [54]		✓	✓					$1 \geq w_{ij}^{kl} \geq t \geq 0$
Kumar et al. [36]		✓	✓	✓				
Salzmann & Fua [61]		✓			✓			
Taylor et al. [68]		✓	✓		✓			$w_{ij}^{kl} = 1$, if $j \in \mathcal{N}_r(i)$,
Chhatkuli et al. [15]		✓			✓			$w_{ij}^{kl} = 0$, otherwise
Probst et al. [58]		✓			✓			
Shashua & Wolf [62]		✓	✓		✓			$w_{ij}^{kl} = s_e^{kl}$, if $e \in S^{kl}$,
Avidan & Shashua [7]		✓			✓			$w_{ij}^{kl} = 0$, otherwise
Vogel et al. [72]		✓			✓			
Li et al. [39]		✓	✓	✓	✓			
Costeira et al. [16]			✓		✓			
Wolf & Shashua [76]		*			✓			$W^{kl} = 1, \text{rank}(\mathbf{D}) \leq \frac{(b+1)(b+2)}{2}$,
Harley & Vidal [27]		*			✓			$\text{rank}(\mathbf{D}^{kl}) \leq \min(8, b+2)$
Agudo et al. [2]				*	✓			
Wu et al. [77]		✓	✓					
Ours		✓	✓	✓	✓	✓	✓	

Table 1: **Our formulation in relation to the existing works.** All five priors, including the rigid prior, can be used to train our method, given the suitability of priors on the structures of interest. Indices $\{i, j, k, l\}$ iterate through all, unless mentioned otherwise. The asterisk(*) indicates inexact or incomplete use of a prior.

the constraints $\text{rank}(\mathbf{D}^{kl}) \leq \min(8, b+2)$ for all k, l and $\text{rank}(\mathbf{D}) \leq \frac{(b+1)(b+2)}{2}$ are necessary to recover the low-rank structures, using Euclidean distance matrices.

Proof The proof is provided in the suppl. material. ■

SM prior can also be integrated in our formulation by enforcing rigidity scores. For example, rigid objects with linear constant velocity maintain highest rigidity with respect to parallel lines/planes⁶ along the velocity direction. In fact, the exact rigidity scores to the other points can be derived, if the velocity is also known. This is particularly interesting, if the motion prior and image semantics are known, *e.g.* a rigid car driving on the rigid road with stops, or a rigid surgical tool poking an isometric organ surface. We represent a set of edges whose rigidity scores can be measured (or bounded) in some or all image pairs, as $S^{kl} = \{s_e^{kl} | e \in \{e_{ij}^{kl}\}\}$. These scores then can also be used (if necessary with bounds) in our formulation. Needless to say, the SM prior can be used in conjunction with LR in practice [76, 27]. An overview of commonly used priors as discussed in Section 2 in relation to previous works is given in Table 1, and presented in a unified framework. Our method can be used for all priors unified in Table 1. However, different priors are useful depending on the scene.

Having provided this unified theoretical framework, we proceed with the experimental analysis of our proposed methodology. In this work, we are interested in reconstructing the depth of *all* parts of a scene, and in that direction we choose to focus on the ARAP prior which is suf-

⁶Which can be represented using a set of points on them.

Figure 3: **Unsupervised learning of depth reconstruction and motion embeddings.** Our model consists of a depth and a per-pixel motion embedding networks, which jointly learn depth map, and motion embeddings whose similarity expresses the rigidity between scene point pairs. A pre-trained flow network used for dense correspondences is included for completeness. The training objective is given in Eq. (6). After training, only a single-pass through the depth network is required for the inference of depth from an image.

ficient to resolve the scale ambiguity between freely moving parts, in contrast to priors such as the LR in multi-body settings. That being said, the other priors we unify in Table 1 can also be used within our framework to reconstruct scenes/objects operating under different priors, by computing rigidity scores with respect to the corresponding prior and using these scores in our formulation.

5. Learning Depth Using ARAP

We first use problem (3) to learn depth reconstruction from image data, given the dense correspondences between views k and l . We will then formalize our learning objective for a given image pair I^k and I^l of views k and l . Our optimization objective can be extended to multi-view images $\mathcal{I} = \{I^1, I^2, \dots, I^m\}$ of the same scene. An overview of our learning pipeline is shown in Fig. 3.

Objective and overview. Using a deep convolutional neural network $f_{\theta_d}(\mathcal{I}) : \mathbb{R}^{H \times B \times 3} \rightarrow [0, \mathbb{R}^+]^{H \times B}$ parameterized by θ_d , we wish to estimate the depth for a given RGB image of size $H \times B \times 3$, as the output of the network.

In order to train the depth network, we use the ARAP prior. For that purpose, we predict depths for the views k and l as: $\Lambda^k = f_{\theta_d}(I^k)$ and $\Lambda^l = f_{\theta_d}(I^l)$, respectively. In order to use the ARAP objective (3), we establish correspondences between views k and l with a pre-trained optical flow network $f_{\theta_f}(\mathcal{I} \times \mathcal{I})$, whose weights are frozen. We use the dense correspondences for the view pair (k, l) obtained from $f_{\theta_f}(I^k, I^l)$ to compute the difference of EDMs given by $E(\Lambda^k) - E(\Lambda^l)$. Next we describe how we obtain the rigidity scores W^{kl} for the view pair (k, l) .

Motion embeddings. Revisiting the initial optimization problem given in (3), we aim to estimate the weight matrix W^{kl} , whose entries $w_{ij}^{kl} \in [0, 1]$ represent the *rigidity* between points (X_i^k, X_j^l) , across the transformation between the views $(k \leftrightarrow l)$. To derive this weight matrix W^{kl} , we propose to learn *per-pixel motion embed-*

dings, whose similarity indicates the sought rigidity between points. In particular, we seek for a dense and complete map of motion embeddings, given a pair of views (k, l) . To this end, we make use of another neural network, $f_{\theta_m}(\mathcal{I} \times \mathcal{I}) : \mathbb{R}^{(H \times B \times 3) \times (H \times B \times 3)} \rightarrow [0, 1]^{H \times B \times v}$ parameterized by θ_m , to learn the motion embeddings:

Definition 5.1 (Motion embedding matrix) A motion embedding matrix, say $M^{kl} \in \mathbb{R}^{v \times n}$, is a matrix representing n points in v -dimensional Euclidean space, in which $M^{kl} = [m_1^{kl}, m_2^{kl}, \dots, m_n^{kl}] \in \mathbb{R}^{v \times n}$. The entry m_i^{kl} represents a **motion embedding vector** for each point X_i^k in the k -th view, representing the motion that the point undergoes across the given view pair (k, l) .

Having obtained the motion embedding matrix M^{kl} , we derive the entries of W^{kl} from the pairwise distances between the motion embedding vectors m_i^{kl} and m_j^{kl} . It should be noted that we expect a higher weight when the distance between the motion embedding vectors are smaller, and vice versa. With this consideration, we formulate the motion similarity scores for every pair (as illustrated in Fig. 4) as,

$$w_{ij}^{kl} = 1 - \tanh(\|m_i^{kl} - m_j^{kl}\|). \quad (4)$$

Edge Sampling. We aim to compute the *motion similarity scores*, w_{ij}^{kl} , ideally for all possible point pair combinations. However, performing this computation is not tractable considering the number of pixels. Therefore, from the graph for the k -th view, say $\mathcal{G}^k = (V = \mathcal{X}^k, E = \mathcal{X}^k \times \mathcal{X}^k)$, we uniformly sample a random subset of all edges (point pairs) for the loss computation. The sampled edges are denoted as $\mathcal{S}^k$. Using this subset $\mathcal{S}^k$, we compute a weight matrix $W_{\mathcal{S}}^{kl}$, whose entries are given by

$$w_{\mathcal{S}, ij}^{kl} = \begin{cases} w_{ij}^{kl} & (X_i^k, X_j^k) \in \mathcal{S}^k, \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

Figure 4: **Motion similarity scores.** For a given point pair (X_i^k, X_j^k) from the scene graph $\mathcal{G}^k$, we retrieve the associated motion embeddings m_i^{kl} and m_j^{kl} , from which we derive motion similarity scores. These scores express how much we expect the 3D distance to be preserved for each point pair.

Loss Formulation. Using (4) and (5), we reformulate the problem of (3) as a loss function to train our networks, parameterized by $\theta = \{\theta_d, \theta_m\}$, as follows:

$$\mathcal{L}_\theta(\Lambda^k, \Lambda^l, W_S^{kl}) = \frac{\|W_S^{kl} \odot (E(\Lambda^k) - E(\Lambda^l))\|_{1,1}}{\alpha \|W_S^{kl}\|_{1,1}}. \quad (6)$$

The normalization factor $\alpha = \|E(\Lambda^k) + E(\Lambda^l)\|_{1,1}$ is introduced for numerical stability and to avoid reconstructions with near-zero depth values. The constraints of (3) are imposed in the network output: Depths Λ^k are ensured to be positive by using a sigmoid on inverse depth output from the network. Similarly, the weights W^{kl} are bounded by using (6), followed by a τ -offset and clamping of the weight values to $[0, 1]$ in progression (to enforce the constraints of (3)). Computation of (6) can be efficiently performed by exploiting the sparsity of W_S^{kl} . During the calculation, we only iterate through the indices with non-zero entries of W_S^{kl} . We additionally incorporated a weight-norm regularization term in our training objective, i.e. $\beta \|W_S^{kl}\|_{1,1}$, to control the maximization of the weights. Further implementation details can be found in the supplementary. We summarize our loss computation process in Algorithm 1.

Network structure. Our depth network $f_{\theta_d}(\mathcal{I})$ consists of a ResNet-18 based encoder and a decoder. The input to our depth network is a single RGB image, and the output from the network is a single depth map of the same size. Our per-pixel motion embedding network $f_{\theta_m}(\mathcal{I} \times \mathcal{I})$ also consists of a multi-input ResNet-18 based encoder and a decoder. The input to the motion embedding network is a pair of images, whose channels are concatenated to create a single input tensor before being passed through the network. The output from the motion embedding network is a motion-embedding map, which, in our case, has 3 channels.

6. Experimental Results

We perform a variety of experiments to demonstrate the performance of our unsupervised monocular pipeline for non-rigid scenes. In addition to assessing the depth recon-

Algorithm 1 $[\mathcal{L}_\theta^{kl}] = \text{computeLossARAP}(l^k, l^l)$

- 1: Sample a set of edges $\mathcal{S}^k$ with their vertices $\mathcal{V}^k$.
 - 2: Estimate the motion embedding $M^{kl} = f_{\theta_m}(l^k, l^l)$.
 - 3: Compute W_S^{kl} using M^{kl} and (4) for $\mathcal{S}^k$.
 - 4: Establish (i, j) between (k, l) using $F^{kl} = f_{\theta_f}(l^k, l^l)$.
 - 5: Start loop $s = k, l$
 - 6: – Estimate the depth $\Lambda^s = f_{\theta_d}(l^s)$
 - 7: – Reconstruct 3D $X^s(\Lambda^s) = [\lambda_1^s u_1^s, \dots, \lambda_n^s u_n^s]$ for $\mathcal{V}^s$.
 - 8: – Compute the EDM $E(\Lambda^s)$ using $X^s(\Lambda^s)$.
 - 9: End loop
 - 10: Compute loss $\mathcal{L}_\theta(\Lambda^k, \Lambda^l, W_S^{kl})$ using (6).
 - 11: Return $\mathcal{L}_\theta(\Lambda^k, \Lambda^l, W_S^{kl})$.
-

Loss computation between two views for ARAP prior.

struction performance, we also evaluate the learned motion embeddings in terms of the motion segmentation task.

Datasets. In our experiments, we use *MPI Sintel* [11], *VolumeDeform* [30], and *Hamlyn Laparoscopic Video Dataset* [82]. From the *MPI Sintel* training subset, we select a set of 14 final-pass sequences with varying levels of motions. We also use the *VolumeDeform* dataset consisting of 8 sequences for further experiments on deforming scenes. Lastly, we evaluate our method on the *Hamlyn Centre Laparoscopic Video Dataset*, which consists of rectified stereo image pairs collected from a partial nephrectomy. Ground-truth depth maps for *VolumeDeform* are obtained from the provided depth recordings; and OpenSFM [5, 4] was used to obtain ground-truth depth for *Hamlyn* using the calibrated stereo pairs in the dataset. For *Sintel*, both ground-truth depth maps as well as optical flow maps are provided, which is not the case for the other two datasets. A pre-trained supervised optical flow network, RAFT [69], performs reasonably well on *VolumeDeform* for estimating dense correspondences. On *Hamlyn*, due to a large domain gap, we fine-tune an unsupervised model, DDFlow [41].

Training Details. We follow two different sets of experiments: test-time training or inference. As the small volumes of *MPI Sintel* and *VolumeDeform* datasets do not yet enable us to train a depth network that can generalize, we perform *test-time training* for each image sequence, i.e. we train separate models for each sequence, for our methods and the deep baselines. In this pipeline, for all of the methods, the training is rather used as an optimization procedure for solving for the depth in all frames, given an image sequence. For the *Hamlyn* dataset where the amount of data enables us to train models that could generalize, we present inference results. We use the original training and test split from the *Hamlyn* dataset, and further create train/val/test splits. We implemented our model in PyTorch [56]. In all experiments, the ResNet18-based depth encoder and motion-embedding encoder were initialized with weights from ImageNet [18] pre-training. Adam [34] opti-

Scene	Method	VolumeDeform [30] dataset							MPI Sintel [11] dataset						
		Abs Rel↓	Sq Rel↓	RMSE↓	RMSE _{log} ↓	$\delta < 1.25^\uparrow$	$\delta < 1.25^2^\uparrow$	$\delta < 1.25^3^\uparrow$	Abs Rel↓	Sq Rel↓	RMSE↓	RMSE _{log} ↓	$\delta < 1.25^\uparrow$	$\delta < 1.25^2^\uparrow$	$\delta < 1.25^3^\uparrow$
Mostly Rigid Seq.	PackNet [26]	0.691	2.879	10.544	1.739	0.425	0.591	0.667	0.355	9.278	10.923	0.544	<u>0.499</u>	0.789	<u>0.885</u>
	Li et al. [39]	0.871	5.267	4.284	0.757	0.214	0.496	0.680	0.412	3.798	3.256	0.663	0.476	0.676	0.817
	Ours w/ motion	<u>0.562</u>	<u>2.483</u>	<u>5.327</u>	<u>0.700</u>	0.294	0.523	<u>0.688</u>	0.213	1.400	6.724	<u>0.533</u>	0.566	0.807	0.886
	Ours w/o motion	0.549	2.441	<u>5.315</u>	0.694	<u>0.300</u>	<u>0.525</u>	0.704	<u>0.255</u>	<u>1.710</u>	<u>3.780</u>	0.528	0.493	<u>0.794</u>	0.883
Non-Rigid Seq.	PackNet [26]	0.774	3.595	4.196	2.680	0.242	0.386	0.535	1.480	14.737	7.465	0.882	0.327	0.508	0.640
	Li et al. [39]	1.217	9.254	4.391	<u>0.852</u>	0.211	0.419	0.594	1.312	15.242	6.278	0.895	0.320	0.498	0.618
	Ours w/ motion	<u>0.562</u>	<u>1.675</u>	2.576	0.533	0.492	0.673	0.823	<u>0.638</u>	<u>4.833</u>	<u>7.602</u>	0.833	0.442	0.615	0.713
	Ours w/o motion	0.519	1.363	2.767	0.553	<u>0.373</u>	<u>0.635</u>	0.817	0.617	4.686	7.681	<u>0.835</u>	0.406	0.594	0.708
All Seq.	PackNet [26]	0.673	3.623	6.585	2.227	0.299	0.465	0.576	1.134	13.057	8.529	<u>0.778</u>	0.380	0.594	0.716
	Li et al. [39]	0.992	6.855	4.268	0.777	0.211	0.465	0.652	1.012	11.427	5.271	0.818	0.372	0.557	0.684
	Ours w/ motion	<u>0.578</u>	<u>2.513</u>	5.210	<u>0.726</u>	0.400	0.597	<u>0.748</u>	0.507	3.777	7.332	0.741	0.480	0.674	0.766
	Ours w/o motion	0.540	2.301	<u>4.294</u>	0.721	<u>0.349</u>	<u>0.582</u>	0.753	<u>0.505</u>	<u>3.770</u>	<u>6.481</u>	0.741	<u>0.433</u>	<u>0.655</u>	<u>0.762</u>

Table 2: **Depth reconstruction evaluation for VolumeDeform [30] and MPI Sintel [11] datasets.** Sequences are classified into two categories: mostly rigid or non-rigid. For MPI Sintel, scenes with less than 10% dynamic pixels are labeled as mostly-rigid. VolumeDeform split is based on our qualitative analysis. Our method with motion embeddings results in superior performance in highly non-rigid scenes in both datasets. Best results are in bold, second best are underlined. Both of our methods generally perform better than our baselines, and for the non-rigid sequences, we can see the benefits of using our method with motion embeddings.

Figure 5: **Our qualitative results with depth and motion networks for MPI Sintel dataset.** Top to bottom: RGB, predicted depth, and predicted embeddings. Embedding colors encoding information about motion clusters are not related across sequences.

mizer with $\beta_1 = 0.9$, $\beta_2 = 0.999$ was used, in combination with a learning rate decay by 0.1 every 10 epochs. We follow a two-stage training. In the first stage, we jointly train the motion-embedding network and the depth network. In the second stage, we freeze the weights of the motion-embedding network, and re-initialize the depth network training. In the latter stage, we apply a τ -offset to the weights, to impose the constraint from (3). For the selection of the edges, we uniformly sample pairs of points with correspondences in the consecutive frame. We empirically find 100K pairs (edges) to provide a suitable trade-off

Scene	Method	Abs Rel↓	Sq Rel↓	RMSE↓	RMSE _{log} ↓	$\delta < 1.25^\uparrow$	$\delta < 1.25^2^\uparrow$	$\delta < 1.25^3^\uparrow$
		Rel↓	Rel↓	↓	↓	↑	↑	↑
Rigid	w/ motion	0.387	2.276	5.954	0.778	0.407	0.632	0.718
	w/o motion	0.393	2.335	5.529	0.787	0.431	0.636	0.725
Non-Rigid	w/ motion	0.702	1.828	1.618	0.503	0.348	0.608	0.795
	w/o motion	0.889	5.879	8.919	0.551	0.329	0.592	0.732

Table 3: **Metrics computed over the dynamic parts or static parts separately.** The separation between the dynamic and static parts was obtained using the ground-truth motion segmentation maps. Metrics are averaged over all images from all sequences.

between increased memory requirements and slow convergence. Further details are in the supplementary material.

Depth reconstruction results. Our qualitative results are demonstrated in Fig. 1, 5 and 6. We report our results by performing per-image median ground-truth scaling for each method, as introduced in [84]. For the *MPI Sintel* dataset, we evaluate the performance only where the GT depth is smaller than 50 m. For the other datasets, we evaluate the depth for every point where the GT is available. As shown in Table 2, 3 and 4, we report results from “Ours w/ motion”, and “Ours w/o motion”. “Ours w/ motion” refers to the setting we described in Section 5, where we learn the motion similarity scores serving the ARAP prior. “Ours w/o motion” refers to the same unsupervised pipeline, except that all rigidity scores are explicitly set to 1. With this modification, piecewise-rigidity is implied during training. In Table 3, we further compare the performance of our methods, by separately evaluating the performance in dynamic and static parts of the images. In Table 4, the results from the non-rigid reconstruction models *DLH* [17], *MDH* [58] and *MaxRig* [31] were obtained using the complete test sequence to perform reconstruction. The evaluations for *DLH* and *MaxRig* were performed only at the reconstructed points as these methods provide sparse reconstructions.

Motion segmentation results. We also evaluate the performance of our per-pixel motion embedding. As the *ground-truth* motion-embeddings are not available, we per-

Figure 6: **Depth reconstruction results for the Hamlyn dataset.** Images and depth predictions are presented. The images shown in grayscale here are used in RGB format during training.

Method	Abs	Sq	RMSE	RMSE _{log}	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
	Rel↓	Rel↓	↓	↓	↑	↑	↑
PackNet [26]	0.389	2.533	4.378	0.418	0.469	0.733	0.882
Li <i>et al.</i> [39]	0.344	2.078	4.169	0.385	0.484	0.772	0.907
DLH [17]	0.644	5.150	6.108	0.611	0.291	0.547	0.744
MDH [58]	1.222	14.210	4.530	0.412	0.572	0.829	0.945
MaxRig [31]	0.232	1.083	3.247	0.294	0.587	0.850	0.962
N-NRSfM [64]	0.392	2.548	5.295	0.730	0.361	0.626	0.779
Ours w/ motion	<u>0.217</u>	<u>0.941</u>	<u>3.120</u>	<u>0.279</u>	<u>0.608</u>	<u>0.886</u>	<u>0.977</u>
Ours w/o motion	0.213	0.921	3.065	0.272	0.618	0.891	0.980

Table 4: **Depth reconstruction performance evaluation for Hamlyn Dataset [82].** Both of our methods result in superior performance compared to previous methods.

Method	Mean	Mean
	ACC ↑	IoU ↑
Tanai <i>et al.</i> [67]	<u>0.872</u>	<u>0.728</u>
Yang <i>et al.</i> [81]	0.755	0.519
Ours w/ motion	0.912	0.731

Table 5: **Average overall pixel accuracy on MPI Sintel [11].** Overall pixel accuracy and intersection-over-union metrics were averaged over the sequences.

form the evaluation on *MPI Sintel* dataset using the available ground-truth motion segmentation maps [78]. With that purpose, we first calculate the static background embedding vector by taking channel-wise median of the image border embeddings, assuming the image border mostly consists of the static background. Then we separate the static part from the dynamic parts of the scene, by thresholding the distances between the static background embedding and each per-pixel motion embedding vectors. In Table 5, we compare the performance of our model with two recent works [67] and [81], and demonstrate our competitive results for the motion segmentation task in Fig. 7.

Discussion. Through numerous experiments, we demonstrated that our formulation with ARAP assumption performs well, even for highly dynamic scenes. However, for some scenes with small deformations, we have observed that the pairwise rigidity assumption performs better. Such behaviour is often due to the short video sequence-based supervision. Often, the scene rigidity may be maintained for a small duration. More importantly, large scene parts are often rigid, which offer better reconstruction when stronger priors are employed. However, rigid reconstruction can be performed rather easily, and despite the fact that most of the time dynamic and deformable objects are of higher interest, they are also more difficult to reconstruct without a suitable prior. In this direction, we investigated several of the com-

Figure 7: **Motion segmentation performance.** Our method gives competitive results despite motion segmentation being only an auxiliary output in our pipeline.

monly used priors, and demonstrated the utility of ARAP assumption for our pipeline. We believe that our framework can be very useful when the scene-specific priors are known for individual use-cases. Note that our formulation highly differs from existing unsupervised monocular depth methods. More specifically, the sole source of self-supervision in our case is the minimization of the geometric errors, and not the photometric errors. This is particularly beneficial for avoiding reconstruction errors that could stem from photometric errors, especially in the presence of deformations. This fact, along with the motion-independent formulation may explain why our pair-wise rigidity works better than a motion-explicit rigid unsupervised learning method [26] for non-rigid scenes. Even if we seek for pair-wise rigidity in non-rigid scenes, since the exact rigid reconstruction is not feasible, the optimal reconstruction is bound to be ARAP.

7. Conclusion

We demonstrated the possibility for unsupervised learning of depth using monocular videos of non-rigid scenes in order to infer depth from a single image. Unsupervised learning is enabled by established priors used in the NRSfM literature. Building upon existing works, we reformulated commonly used NRSfM priors within a unified framework suitable for neural network training. We further investigated the utility of the ARAP prior for dense depth reconstruction from a single view. Our experiments demonstrate improvements over the alternative methods on numerous datasets. The improvements can be attributed to our geometric error-based loss function (6) computed directly in the 3D reconstruction space. We assumed only the minimum necessary prior in this work, as we aimed to learn without any supervision (including additional scene priors). Depending on scene properties, different priors can be integrated into our pipeline in a similar way, boosting the performance further.

Acknowledgments. This research was partly supported by Innosuisse funding (Grant No. 34475.1 IP-ICT), by FIFA, and by the ETH Zurich project with Specta and EU H2020 ENCORE (Grant No. 820434).

References

- [1] Antonio Agudo, Lourdes Agapito, Begona Calvo, and Jose MM Montiel. Good vibrations: A modal analysis approach for sequential non-rigid structure from motion. In *Proceedings of the IEEE Conference on computer vision and pattern recognition*, pages 1558–1565, 2014. 2
- [2] Antonio Agudo, Melcior Pijoan, and Francesc Moreno-Noguer. Image collection pop-up: 3d reconstruction and clustering of rigid and non-rigid categories. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2607–2615, 2018. 2, 4
- [3] Ijaz Akhter, Yaser Sheikh, Sohaib Khan, and Takeo Kanade. Trajectory space: A dual representation for nonrigid structure from motion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(7):1442–1456, 2010. 2
- [4] Pablo Fernández Alcantarilla, Adrien Bartoli, and Andrew J. Davison. Kaze features. In Andrew Fitzgibbon, Svetlana Lazebnik, Pietro Perona, Yoichi Sato, and Cordelia Schmid, editors, *Computer Vision – ECCV 2012*, pages 214–227, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg. 6
- [5] Pablo Fernández Alcantarilla, Jesus Nuevo, and Adrien Bartoli. Fast explicit diffusion for accelerated features in non-linear scale spaces. In *BMVC*, 2013. 6
- [6] Marc Alexa, Daniel Cohen-Or, and David Levin. As-rigid-as-possible shape interpolation. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 157–164, 2000. 2
- [7] Shai Avidan and Amnon Shashua. Trajectory triangulation: 3d reconstruction of moving points from a monocular image sequence. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(4):348–357, 2000. 2, 4
- [8] Adrien Bartoli, Vincent Gay-Bellile, Umberto Castellani, Julien Peyras, Søren Olsen, and Patrick Sayd. Coarse-to-fine low-rank structure-from-motion. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008. 2
- [9] Matthew Brand and Rahul Bhotika. Flexible flow for 3d non-rigid tracking and shape recovery. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, volume 1, pages I–I. IEEE, 2001. 2
- [10] Christoph Bregler, Aaron Hertzmann, and Henning Biermann. Recovering non-rigid 3d shape from image streams. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition. CVPR 2000 (Cat. No. PR00662)*, volume 2, pages 690–696. IEEE, 2000. 2
- [11] Daniel Butler, Jonas Wulff, Garrett Stanley, and Michael Black. A naturalistic open source movie for optical flow evaluation. In A. Fitzgibbon et al. (Eds.), editor, *European Conf. on Computer Vision (ECCV)*, Part IV, LNCS 7577, pages 611–625. Springer-Verlag, Oct. 2012. 6, 7, 8
- [12] Weifeng Chen, Zhao Fu, Dawei Yang, and Jia Deng. Single-image depth perception in the wild. *CoRR*, abs/1604.03901, 2016. 1
- [13] Yuhua Chen, Cordelia Schmid, and Cristian Sminchisescu. Self-supervised learning with geometric constraints in monocular video: Connecting flow, depth, and camera. *CoRR*, abs/1907.05820, 2019. 1
- [14] Ajad Chhatkuli, Daniel Pizarro, and Adrien Bartoli. Non-rigid shape-from-motion for isometric surfaces using infinitesimal planarity. In *BMVC*, 2014. 2
- [15] Ajad Chhatkuli, Daniel Pizarro, Toby Collins, and Adrien Bartoli. Inextensible non-rigid shape-from-motion by second-order cone programming. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1719–1727, 2016. 2, 3, 4
- [16] João Paulo Costeira and Takeo Kanade. A multibody factorization method for independently moving objects. *International Journal of Computer Vision*, 29(3):159–179, 1998. 2, 4
- [17] Yuchao Dai, Hongdong Li, and Mingyi He. A simple prior-free method for non-rigid structure-from-motion factorization. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2018–2025. IEEE, 2012. 2, 7, 8
- [18] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Fei Fei Li. Imagenet: a large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 06 2009. 6
- [19] Sotirios Diamantas, Anastasios Oikonomidis, and Richard Crowder. Depth estimation for autonomous robot navigation: A comparative approach. In *2010 IEEE International Conference on Imaging Systems and Techniques*, pages 426–430, 2010. 1
- [20] Tolga Eren, Peter N Belhumeur, and A Stephen Morse. Closing ranks in vehicle formations based on rigidity. In *Proceedings of the 41st IEEE Conference on Decision and Control*, 2002., volume 3, pages 2959–2964. IEEE, 2002. 3
- [21] Katerina Fragkiadaki, Marta Salas, Pablo Arbelaez, and Jitendra Malik. Grouping-based low-rank trajectory completion and 3d reconstruction. In *Advances in Neural Information Processing Systems*, pages 55–63, 2014. 2
- [22] Mathias Gallardo, Daniel Pizarro, Toby Collins, and Adrien Bartoli. Shape-from-template with curves. *International Journal of Computer Vision*, 128(1):121–165, 2020. 2
- [23] Ravi Garg, Vijay Kumar B. G, and Ian D. Reid. Unsupervised CNN for single view depth estimation: Geometry to the rescue. In *ECCV*, 2016. 1
- [24] Clément Godard, Oisín Mac Aodha, and Gabriel J. Brostow. Unsupervised monocular depth estimation with left-right consistency. *CoRR*, abs/1609.03677, 2016. 1
- [25] Clément Godard, Oisín Mac Aodha, and Gabriel J. Brostow. Digging into self-supervised monocular depth estimation. *CoRR*, abs/1806.01260, 2018. 1
- [26] Vitor Guizilini, Rares Ambrus, Sudeep Pillai, and Adrien Gaidon. Packnet-sfm: 3d packing for self-supervised monocular depth estimation. *CoRR*, abs/1905.02693, 2019. 1, 7, 8
- [27] Richard Hartley and René Vidal. Perspective nonrigid shape and motion recovery. In *European Conference on Computer Vision*, pages 276–289. Springer, 2008. 2, 4
- [28] Florian Heinrich, Kai Bornemann, Kai Lawonn, and Christian Hansen. Depth perception in projective augmented reality: An evaluation of advanced visualization techniques. In

- 25th ACM Symposium on Virtual Reality Software and Technology, VRST '19, New York, NY, USA, 2019. Association for Computing Machinery. 1
- [29] Takeo Igarashi, Tomer Moscovich, and John F Hughes. As-rigid-as-possible shape manipulation. *ACM transactions on Graphics (TOG)*, 24(3):1134–1141, 2005. 2
- [30] Matthias Innmann, Michael Zollhöfer, Matthias Nießner, Christian Theobalt, and Marc Stamminger. Volumedeform: Real-time volumetric non-rigid reconstruction. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part VIII*, volume 9912 of *Lecture Notes in Computer Science*, pages 362–379. Springer, 2016. 6, 7
- [31] Pan Ji, Hongdong Li, Yuchao Dai, and Ian D. Reid. "maximizing rigidity" revisited: a convex programming approach for generic 3d shape reconstruction from multiple perspective views. *CoRR*, abs/1707.05009, 2017. 2, 7, 8
- [32] Imran Khan. Non-rigid structure-from-motion with uniqueness constraint and low rank matrix fitting factorization. *IEEE transactions on multimedia*, 16(5):1350–1357, 2014. 2
- [33] Byung-soo Kim, Pushmeet Kohli, and Silvio Savarese. 3d scene understanding by voxel-crf. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1425–1432, 2013. 1
- [34] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2014. cite arxiv:1412.6980 Comment: Published as a conference paper at the 3rd International Conference for Learning Representations, San Diego, 2015. 6
- [35] Chen Kong and Simon Lucey. Deep non-rigid structure from motion. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019. 2
- [36] Suryansh Kumar, Ram Srivatsav Ghorakavi, Yuchao Dai, and Hongdong Li. Dense depth estimation of a complex dynamic scene without explicit 3d motion estimation. *arXiv preprint arXiv:1902.03791*, 2019. 2, 4
- [37] Yevhen Kuznietsov, Jörg Stückler, and Bastian Leibe. Semi-supervised deep learning for monocular depth map prediction. *CoRR*, abs/1702.02706, 2017. 1
- [38] Hongdong Li. Multi-view structure computation without explicitly estimating motion. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 2777–2784. IEEE, 2010. 2, 3, 4
- [39] Hanhan Li, Ariel Gordon, Hang Zhao, Vincent Casser, and Anelia Angelova. Unsupervised monocular depth learning in dynamic scenes. *arXiv preprint arXiv:2010.16404*, 2020. 1, 2, 4, 7, 8
- [40] Chao Liu, Jinwei Gu, Kihwan Kim, Srinivasa G. Narasimhan, and Jan Kautz. Neural rgb->d sensing: Depth and uncertainty from a video camera. *CoRR*, abs/1901.02571, 2019. 1
- [41] Pengpeng Liu, Irwin King, Michael R. Lyu, and Jia Xu. DdfLOW: Learning optical flow with unlabeled data distillation. In *AAAI*, 2019. 6
- [42] Matthew Loper, Naureen Mahmood, and Michael J. Black. Mosh: Motion and shape capture from sparse markers. *ACM Trans. Graph.*, 33(6), Nov. 2014. 1
- [43] Chenxu Luo, Zhenheng Yang, Peng Wang, Yang Wang, Wei Xu, Ram Nevatia, and Alan L. Yuille. Every pixel counts ++: Joint learning of geometry and motion with 3d holistic understanding. *CoRR*, abs/1810.06125, 2018. 1
- [44] Xuan Luo, Jia-Bin Huang, Richard Szeliski, Kevin Matzen, and Johannes Kopf. Consistent video depth estimation, 2020. 1
- [45] Reza Mahjourian, Martin Wicke, and Anelia Angelova. Unsupervised learning of depth and ego-motion from monocular video using 3d geometric constraints. *CoRR*, abs/1802.05522, 2018. 1
- [46] Abed Malti, Richard Hartley, Adrien Bartoli, and Jae-Hak Kim. Monocular template-based 3d reconstruction of extensible surfaces with local linear elasticity. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1522–1529, 2013. 2
- [47] Abed Malti and Cédric Herzet. Elastic shape-from-template with spatially sparse deforming forces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3337–3345, 2017. 2
- [48] Ishit Mehta, Parikshit Sakurikar, and PJ Narayanan. Structured adversarial training for unsupervised monocular depth estimation. In *2018 International Conference on 3D Vision (3DV)*, pages 314–323. IEEE, 2018. 1
- [49] David Novotny, Nikhila Ravi, Benjamin Graham, Natalia Neverova, and Andrea Vedaldi. C3dpo: Canonical 3d pose networks for non-rigid structure from motion. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019. 1, 2
- [50] Carl Olsson and Magnus Oskarsson. A convex approach to low rank matrix approximation with missing data. In *Scandinavian Conference on Image Analysis*, pages 301–309. Springer, 2009. 2
- [51] K. Ozden, K. Cornelis, L. Van Eycken, and L. Van Gool. Reconstructing 3d independent motions using non-accidentalness. In *CVPR*, volume 1, pages I-819–I-825 Vol.1, 2004. 2
- [52] Shaifali Parashar, Adrien Bartoli, and Daniel Pizarro. Robust isometric non-rigid structure-from-motion. *arXiv preprint arXiv:2010.04690*, 2020. 2
- [53] Shaifali Parashar, Daniel Pizarro, and Adrien Bartoli. Isometric non-rigid shape-from-motion in linear time. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4679–4687, 2016. 2
- [54] Shaifali Parashar, Daniel Pizarro, Adrien Bartoli, and Toby Collins. As-rigid-as-possible volumetric shape-from-template. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 891–899, 2015. 2, 4
- [55] Shaifali Parashar, Mathieu Salzmann, and Pascal Fua. Local non-rigid structure-from-motion from diffeomorphic mappings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2059–2067, 2020. 2
- [56] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming

- Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems* 32, pages 8024–8035. Curran Associates, Inc., 2019. 6
- [57] Mathieu Perriollat, Richard Hartley, and Adrien Bartoli. Monocular template-based reconstruction of inextensible surfaces. *International journal of computer vision*, 95(2):124–137, 2011. 2, 3
- [58] Thomas Probst, Danda Pani Paudel, Ajad Chhatkuli, and Luc Van Gool. Incremental non-rigid structure-from-motion with unknown focal length. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 756–771, 2018. 2, 3, 4, 7, 8
- [59] Anurag Ranjan, Timo Bolkart, Soubhik Sanyal, and Michael J. Black. Generating 3d faces using convolutional mesh autoencoders. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 704–720, 2018. 1
- [60] Chris Russell, Rui Yu, and Lourdes Agapito. Video pop-up: Monocular 3d reconstruction of dynamic scenes. In *European conference on computer vision*, pages 583–598. Springer, 2014. 2
- [61] Mathieu Salzmann and Pascal Fua. Linear local models for monocular reconstruction of deformable surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(5):931–944, 2010. 2, 3, 4
- [62] Amnon Shashua and Lior Wolf. Homography tensors: On algebraic entities that represent three views of static or moving planar points. In *European Conference on Computer Vision*, pages 507–521. Springer, 2000. 2, 4
- [63] Shuhan Shen, Wenhuan Shi, and Yuncai Liu. Monocular 3-d tracking of inextensible deformable surfaces under l_2 -norm. *IEEE Transactions on Image Processing*, 19(2):512–521, 2009. 2
- [64] Vikramjit Sidhu, Edgar Tretschk, Vladislav Golyanik, Antonio Agudo, and Christian Theobalt. Neural dense non-rigid structure from motion with latent space constraints. In *European Conference on Computer Vision (ECCV)*, 2020. 1, 8
- [65] Olga Sorkine and Marc Alexa. As-rigid-as-possible surface modeling. In *Symposium on Geometry processing*, volume 4, pages 109–116, 2007. 2
- [66] Srinath Sridhar, Davis Rempe, Julien Valentin, Bouaziz Sofien, and Leonidas J Guibas. Multiview aggregation for learning category-specific shape reconstruction. In *Advances in Neural Information Processing Systems*, volume 32, 2019. 1
- [67] Tatsunori Tani, Sudipta N. Sinha, and Yoichi Sato. Fast multi-frame stereo scene flow with motion segmentation. *CoRR*, abs/1707.01307, 2017. 8
- [68] Jonathan Taylor, Allan Jepson, and Kiriakos Kutulakos. Non-rigid structure from locally-rigid motion. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2010. 2, 3, 4
- [69] Zachary Teed and Jia Deng. RAFT: recurrent all-pairs field transforms for optical flow. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part II*, volume 12347 of *Lecture Notes in Computer Science*, pages 402–419. Springer, 2020. 6
- [70] Jack Valmadre, Sridha Sridharan, Simon Denman, Clinton Fookes, and Simon Lucey. Closed-form solutions for low-rank non-rigid reconstruction. In *2015 International Conference on Digital Image Computing: Techniques and Applications (DICTA)*, pages 1–6. IEEE, 2015. 2
- [71] Sara Vicente and Lourdes Agapito. Soft inextensibility constraints for template-free non-rigid reconstruction. In *European conference on computer vision*, pages 426–440. Springer, 2012. 2
- [72] Christoph Vogel, Konrad Schindler, and Stefan Roth. 3d scene flow estimation with a piecewise rigid scene model. *International Journal of Computer Vision*, 115(1):1–28, 2015. 2, 4
- [73] Chaoyang Wang, José Miguel Buenaposada, Rui Zhu, and Simon Lucey. Learning depth from monocular videos using direct methods. *CoRR*, abs/1712.00175, 2017. 1
- [74] Xuan Wang, Mathieu Salzmann, Fei Wang, and Jizhong Zhao. Template-free 3d reconstruction of poorly-textured nonrigid surfaces. In *European Conference on Computer Vision*, pages 648–663. Springer, 2016. 2
- [75] Walter Whiteley. Rigidity and scene analysis, 2004. 3
- [76] Lior Wolf and Amnon Shashua. On projection matrices p_k to p_2 , $k=3, \dots, 6$, and their applications in computer vision. *International Journal of Computer Vision*, 48(1):53–67, 2002. 2, 4
- [77] Shangzhe Wu, Christian Rupprecht, and Andrea Vedaldi. Unsupervised learning of probably symmetric deformable 3d objects from images in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1–10, 2020. 1, 4
- [78] Jonas Wulff, Laura Sevilla-Lara, and Michael J. Black. Optical flow in mostly rigid scenes. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017*, pages 6911–6920, Piscataway, NJ, USA, July 2017. IEEE. 8
- [79] Jing Xiao and Takeo Kanade. Uncalibrated perspective reconstruction of deformable structures. In *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, volume 2, pages 1075–1082. IEEE, 2005. 2
- [80] Gengshan Yang, Deqing Sun, Varun Jampani, Daniel Vlasic, Forrester Cole, Huiwen Chang, Deva Ramanan, William T. Freeman, and Ce Liu. Lasr: Learning articulated shape reconstruction from a monocular video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15980–15989, June 2021. 1
- [81] Yanchao Yang, Antonio Loquercio, Davide Scaramuzza, and Stefano Soatto. Unsupervised moving object detection via contextual information separation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 8

- [82] Menglong Ye, Edward Johns, Ankur Handa, Lin Zhang, Philip Pratt, and Guang-Zhong Yang. Self-supervised siamese learning on stereo image pairs for depth estimation in robotic surgery. *CoRR*, abs/1705.08260, 2017. 6, 8
- [83] Zhichao Yin and Jianping Shi. Geonet: Unsupervised learning of dense depth, optical flow and camera pose. *CoRR*, abs/1803.02276, 2018. 1
- [84] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G. Lowe. Unsupervised learning of depth and ego-motion from video. In *CVPR*, 2017. 1, 7