Superpixel Soup: Monocular Dense 3D Reconstruction of a Complex Dynamic Scene

Journal Article**Author(s):**

Kumar, Suryansh; Dai, Yuchao; Li, Hongdong

Publication date:

2021-05-18

Permanent link:

<https://doi.org/10.3929/ethz-b-000428990>

Rights / license:

In Copyright - Non-Commercial Use Permitted

Originally published in:

IEEE Transactions on Pattern Analysis and Machine Intelligence 43(5), <https://doi.org/10.1109/tpami.2019.2955131>

Supervoxel Soup: Monocular Dense 3D Reconstruction of a Complex Dynamic Scene

Suryansh Kumar, *Member, IEEE*, Yuchao Dai, *Member, IEEE*, Hongdong Li, *Senior Member, IEEE*

Abstract—This work addresses the task of dense 3D reconstruction of a complex dynamic scene from images. The prevailing idea to solve this task is composed of a sequence of steps and is dependent on the success of several pipelines in its execution [1]. To overcome such limitations with the existing algorithm, we propose a unified approach to solve this problem. We assume that a dynamic scene can be approximated by numerous piecewise planar surfaces, where each planar surface enjoys its own rigid motion, and the global change in the scene between two frames is as-rigid-as-possible (ARAP). Consequently, our model of a dynamic scene reduces to a *soup* of planar structures and rigid motion of these local planar structures. Using planar over-segmentation of the scene, we reduce this task to solving a “3D jigsaw puzzle” problem. Hence, the task boils down to correctly assemble each rigid piece to construct a 3D shape that complies with the geometry of the scene under the ARAP assumption. Further, we show that our approach provides an effective solution to the inherent scale-ambiguity in structure-from-motion under perspective projection. We provide extensive experimental results and evaluation on several benchmark datasets. Quantitative comparison with competing approaches shows state-of-the-art performance.

Index Terms—Dense 3D reconstruction, perspective camera, as-rigid-as-possible, relative scale ambiguity, structure from motion.

1 INTRODUCTION

THE task of reconstructing 3D geometry of the scene from images —popularly known as structure-from-motion (SfM), is a fundamental problem in computer vision. An initial introduction and working solution to this problem can be found as early as 1970’s and 1980’s [2] [3] [4], which Blake *et al.* discussed comprehensively in their seminal work [5]. While this field of study in the past was largely dominated by sparse feature based reconstruction of a rigid scene [6] [7] [8] [9] [10] and a non-rigid object [11] [12] [13] [14] [15], in recent years, with the surge in computational resources, dense 3D reconstruction of the scene have been introduced and successfully demonstrated [16] [17] [1].

A dense solution to this inverse problem is essential due to its increasing demand in many real-world applications—from animation and entertainment industry to robotics industry (VSLAM). In particular, with the proliferation of *monocular* camera in almost all modern mobile devices has elevated the demand for sophisticated dense reconstruction algorithm. When the scene is static and the camera is moving, 3D reconstruction of such scenes from images can be achieved by using conventional rigid structure from motion techniques [8] [18] [19] [20]. In contrast, to model arbitrary dynamic scene can be very challenging. When the camera is moving and the scene is static under such settings, the elegant geometrical constraint can help explain the camera’s [7] [21], which are later used to realize the dense 3D reconstruction of the scene [19] [20] [17] [22]. However, such geometrical constraint may fail when multiple rigidly moving objects are observed by a moving camera. Although

Fig. 1: Dense 3D reconstruction of a complex dynamic scene, where both the camera and the objects are moving with respect to each other. The top left shows a sample reconstruction on messi sequence from Youtube Object dataset [23]. The top right shows the reconstruction on alley_1 sequence from the MPI Sintel dataset [24].

each of the individual rigid objects can be reconstructed up to an arbitrary scale (assuming motion segmentation is provided), the reconstruction of the whole dynamic scene is generally impossible, simply because the relative scales among all the moving shapes cannot be determined in a globally consistent way. Furthermore, since all the estimated motions are relative to each other, one cannot distinguish camera motion from the object motion. Therefore, prior information about the objects, or the scene, and their relation to the frame of reference are used to fix the placement of these objects relative to each other.

Hence, from the above discussion, it can be argued that the solution to 3D reconstruction of a general dynamic scene is non-trivial. Nevertheless, it is an important problem to solve as many real-world applications need a reliable solution to this problem. For example, understanding of a traffic scene, a typical outdoor traffic scene consists of both multiple rigid motions of vehicles, and non-rigid motion of the pedestrians. To model such scenarios, it is important to have an algorithm that can provide dense 3D information from images.

- Suryansh Kumar is with ETH Zürich and Australian National University. E-mail: sukumar@vision.ee.ethz.ch, suryansh.kumar@anu.edu.au.
- Yuchao Dai is with Northwestern Polytechnical University. E-mail: daiyuchao@gmail.com.
- Hongdong Li is with Australian National University and ARC Centre of Excellent for Robotic Vision. E-mail: hongdong.li@anu.edu.au.

Recently, Ranftl *et al.* [1] proposed a three-step approach to procure dense 3D reconstruction of a general dynamic scene using two consecutive perspective frames. Concretely, it performs object-level motion segmentation followed by per-object 3D reconstruction and finally solves for scale ambiguity. We know that in a general dynamic setting, the task of densely segmenting rigidly moving objects or part is not trivial. Consequently, inferring motion models for deforming shapes becomes very challenging. Furthermore, the success of object-level segmentation builds upon the assumption of multiple rigid motions, fails to describe more general scenarios such as “when the objects themselves are deforming”. Subsequently, 3D reconstruction algorithms dependent on motion segmentation of objects suffer.

Motivated by such limitations, we propose a unified approach that neither performs any object-level motion segmentation nor assumes any prior knowledge about the scene rigidity type and still able to recover scale consistent dense reconstruction of a complex dynamic scene. Our formulation instinctively encapsulates the solution to inherent scale ambiguity in perspective structure from motion which is a very challenging problem in general. We show that by using *two prior assumptions* —about the 3D scene and the deformation, we can effectively pin down the unknown relative scales, and obtain a globally consistent dense 3D reconstruction of a dynamic scene from its two perspective views. The two basic assumptions we used about the dynamic scene are:

- 1) The dynamic scene can be approximated by a collection of *piecewise planar surfaces* each having its own *rigid motion*.
 - 2) The deformation of the scene between two frames is *locally-rigid* but *globally as-rigid-as-possible*.
- *Piece-wise planar model*: Our method models a dynamic scene as a collection of *piece-wise planar* regions. Given two perspective images $\mathbf{I}$ (reference image), $\mathbf{I}'$ (next image) of a general dynamic scene, our method first over-segment the reference image into superpixels. This collection of superpixels are assumed approximation of the dynamic scene in the projective space. It can be argued that modeling dynamic scene per pixel can be more compelling, however, modeling of a scene using planar regions makes this problem computationally tractable for optimization or inference [25], [26].
 - *Locally-rigid and globally as-rigid-as-possible*: We implicitly assume that each local plane undergoes a *rigid motion*. Suppose every individual superpixel corresponds to a small planar patch moving rigidly in 3D space and dense optical flow between frame is given, we can estimate its location in 3D using rigid reconstruction pipeline [8], [27]. Since the *relative scale* of these patches is not determined correctly, they are floating in 3D space as a set of unorganized superpixel soup. Under the assumption that the change between the frame is not too arbitrary rather regular or smooth, the scene can be assumed to be changing as rigid as possible globally. Using this intuition, our method starts finding for each superpixel an appropriate scale, under which the entire set of superpixels can be assembled (glued) together coherently, forming a piecewise smooth surface, *as if* playing the game of “3D jig-

Fig. 2: Reconstructing a 3D surface from a soup of un-scaled superpixels via solving a 3D Superpixel Jigsaw puzzle problem.

saw puzzle”. Hence, we call our method the “SuperPixel Soup” algorithm (see Fig. 2 for a conceptual visualization).

In this paper, we show that our aforementioned assumptions can faithfully model most of the real-world dynamic scenarios. Furthermore, we encapsulate these assumptions in a simple optimization problem which are solved using a combination of continuous and discrete optimization algorithms [28], [29], [30]. We demonstrate the benefit of our approach on available benchmark dataset such as KITTI [31], MPI Sintel [24] and Virtual KITTI [32]. The statistical comparison shows that our algorithm outperforms many available state-of-the-art methods by a significant margin.

2 RELATED WORKS

The solution to SfM has undergone prodigious development since its inception [2]. Even after such a remarkable development in this field, the choice of algorithm depends on the complexity of the object motion and the environment. In this work, we utilize the idea of rigidity (locally) to solve dense reconstruction of a general dynamic scene. The concept of rigidity is not new in structure from motion problem [2] [33] and has been effectively applied as a global constraint to solve large scale reconstruction problem [18]. The idea of global rigidity to solve structure and motion has also been exploited to solve reconstruction over multiple frames via a factorization approach [10].

The literature on structure from motion and its treatment to different scenarios is very extensive. Consequently, for brevity, we only discuss the previous works that are of direct relevance to dynamic 3D reconstruction from *monocular* images. The linear low-rank model has been used for dense non-rigid reconstruction. Kumar *et al.* [34], [35] and Garg *et al.* [36] solved the task with an orthographic camera model assuming feature matches across multiple frames is given as input. Fayad *et al.* [37] recovered deformable surfaces with a quadratic approximation, again from multiple frames. Taylor *et al.* [38] proposed a piecewise rigid solution using locally-rigid SfM to reconstruct a soup of rigid triangles.

While Taylor *et al.* [38] method is conceptually similar to ours, there are major differences:

- 1) We achieve *two-view* dense reconstruction while [38] relies on multiple views ($N \geq 4$).
- 2) We use *perspective camera model* while they rely on an orthographic camera model.
- 3) We solve the scale-indeterminacy issue, which is an inherent ambiguity for 3D reconstruction under perspective projection, while Taylor *et al.* [38] method does not suffer from this, at the cost of being restricted to the orthographic camera model.

Recently, Russel *et al.* [39] and Ranftl *et al.* [1] used object-level segmentation for dense dynamic 3D reconstruction. In

contrast, our method is free from object segmentation, hence circumvent the difficulty associated with motion segmentation in a dynamic setting.

The template-based approach is yet another method for deformable surface reconstruction. Yu *et al.* [40] proposed a direct approach to capture dense, detailed 3D geometry of generic, complex non-rigid meshes using a single RGB camera. While it works for generic surfaces, the requirement of template prevents its wider application to more general scenes. Wang *et al.* [41] introduced a template-free approach to reconstruct a poorly-textured, deformable surface. Nevertheless, its success is restricted to a single deforming surface rather than the entire dynamic scene. Varol *et al.* [42] reconstructed deformable surfaces based on a piecewise reconstruction assuming overlapping patches to be consistent over the entire surface, but again limited to the reconstruction of a single deformable surface.

While the conceptual idea of our work appeared in ICCV 2017, this journal version provides (i) in-depth realization of our overall optimization (ii) Qualitative comparison with [1], Video-PopUp [39] as well as statistical comparison with deep-learning method [43]. (iii) Comprehensive ablation test showing the importance of each term in the overall optimization. (iv) Extensive performance analysis showing the performance with the variation in the number of superpixels, choice of k-nearest neighbor, choice of dense optical flow algorithm and change in the shape of the superpixel. (v) Detail discussion on the failure cases, choice of euclidean metric for nearest neighbor graph construction, and limitation of our work with possible direction for improvements.

3 MOTIVATION AND CONTRIBUTION

The formulation proposed in this work is motivated by the following endeavor in dense structure from motion of a dynamic scene.

3.1 Object level motion segmentation

To solve dense reconstruction of an entire dynamic scene from perspective images, the first step that is practiced usually is: Perform object-level motion segmentation to infer distinct motion models of multiple rigidly moving object in the scene. As alluded before, dense segmentation of moving object in a dynamic scene in itself is a challenging task. Also, non-rigidly moving object themselves may compose of a union of distinct motion models. Therefore, object-level segmentation build upon the assumption of per object rigid motion will fail to describe a general dynamic scene. This motivates us to develop an algorithm that can recover a dense-detailed 3D model of a complex dynamic scene from its two perspective images, *without object-level motion segmentation* as an essential intermediate step.

3.2 Separate treatment for rigid SfM and non-rigid SfM

Our investigation shows that the algorithms for deformable object 3D reconstruction often differs from a rigidly moving object. Not only solutions, but even the assumptions varies significantly e.g orthographic projection, low-rank shape [11] [12] [13] [15]. The reason for such inadequacy is perfectly valid due to the under-constraint nature of the

problem itself. This motivated us to develop an algorithm that can provide *i.e* "3D reconstruction of entire dynamic scene and the non-rigidly deforming object under similar assumptions and formulation."

Although to accomplish this goal for any arbitrary non-rigid deformation remains an open problem, experiments suggest that our framework under the aforementioned assumptions about the scene and the deformation, can reconstruct a general dynamic scene irrespective of the scene rigidity type. Thanks to the recent advancement in the dense optical flow algorithms [44] [45] which can reliably capture smooth non-rigid deformation over frames. These robust dense optical flow algorithms allow us to exploit local motion of deforming surfaces. Thus, our formulation is competent enough to bridge this gap between rigid and non-rigid SfM.

The main contributions of our work are as follows:

- 1) A framework which disentangles object-level motion segmentation for dense 3D reconstruction of a complex dynamic scene.
- 2) A common framework for dense two-frame 3D reconstruction of a complex dynamic scene (including deformable objects), which achieves state-of-the-art performance.
- 3) A new idea to resolve the inherent relative scale ambiguity problem in monocular 3D reconstruction by exploiting the as-rigid-as-possible (ARAP) constraint [46].

4 OUTLINE OF THE ALGORITHM

Before providing the details of our algorithm, we would like to introduce some common notations that are used throughout the paper.

4.1 Notation

We represent two consecutive images as $\mathbf{I}, \mathbf{I}' : \Omega \rightarrow \mathbb{R}^3$ $|\Omega \subset \mathbb{Z}^2$, also referred as reference image and next image respectively. Vectors are represented by bold lower case letter, such as ' $\mathbf{x}$ ' and matrices are represented by bold upper case letter such as ' $\mathbf{X}$ '. The subscript ' a ', ' b ' denotes anchor point and boundary point respectively, for e.g $\mathbf{x}_{ai}$, $\mathbf{x}_{bi}$ represents anchor point and boundary point corresponding to i^{th} superpixel in the image space. The 1-norm, 2-norm of a vector is denoted as $\|\cdot\|_1$ and $\|\cdot\|_2$ respectively. For matrices, Frobenius norm is denoted as $\|\cdot\|_F$.

4.2 Overview

We first over-segment the reference image into superpixels, then model the deformation of the scene by a union of piecewise rigid motions of these superpixels. Specifically, we divide the overall non-rigid reconstruction into a local rigid reconstruction of each superpixel, followed by an assembly process which glues all these individual local reconstructions in a globally coherent manner. While the concept of the above divide-and-conquer procedure looks simple, there is however a fundamental difficulty (of *scale indeterminacy*) in its implementation. Scale-Indeterminacy refers to the well-known fact that using a moving camera one can only recover the 3D structure up to an unknown scale. In our method,

the individual rigid reconstruction of each superpixel can only be determined up to an unknown scale, the assembly of the entire non-rigid scene is only possible if and only if these scales among the superpixels are solved—which is, however, a challenging open task itself.

In this paper, we show how this can be done using two very mild assumption §3.2. Under these assumptions, our method solves the unknown relative scales and obtains a globally-coherent dense 3D reconstruction of a complex dynamic scene from its two perspective views.

4.3 Problem Statement

To implement the above idea of piecewise rigid reconstruction, we first partition the reference image $\mathbf{I}$ into set of superpixels $\xi_I = \{s_1, s_2, \dots, s_i, \dots, s_N\}$, where each superpixel s_i is parametrized by its boundary pixels $\{x_{bi} = [u_{bi}, v_{bi}, 1]^T \mid b = 1, \dots, B_i\}$ and an *anchor point* x_{ai} corresponding to the centroid of the i^{th} superpixel in the image plane. Such a superpixel partition of the image plane naturally induces a piecewise-smooth over segmentation of the corresponding 3D scene surface. We denote this set of 3D scene surfaces as $\xi_W = \{\tilde{s}_1, \tilde{s}_2, \dots, \tilde{s}_i, \dots, \tilde{s}_N\}$. Although *surfel* is perhaps a better term, we nevertheless call it “3D superpixel” for the sake of easy exposition. We further assume each 3D superpixel ($\tilde{s}_i$) is a small 3D *planar patch* $\Pi_{\tilde{s}_i} = \{n_i, \tilde{x}_{ai}, \{\tilde{x}_{bi}\} : (n_i, \tilde{x}_{ai}) \in \mathbb{R}^3 \text{ and } \{\tilde{x}_{bi}\} \in \mathbb{R}^{3 \times B_i}\}$, which is parameterized by surface normal n_i , 3D anchor-point $\tilde{x}_{ai}$, and 3D boundary-points $\{\tilde{x}_{bi}\}$ (i.e these are the pre-images of x_{ai} and $\{x_{bi}\}$). Assume every 3D superpixel $\tilde{s}_i$ moves rigidly according $M_i = \begin{pmatrix} R_i & \lambda_i \hat{t}_i \\ 0 & 1 \end{pmatrix} \in SE(3)$, where R_i represents relative rotation, $\hat{t}_i$ is the translation direction, and λ_i the unknown scale.

After our notation and symbol introduction, we are in a position to put our idea in a more precise way: Given two intrinsically calibrated perspective images $\mathbf{I}$ and $\mathbf{I}'$ of a generally dynamic scene and the corresponding dense optical flow field, our task is to reconstruct a piecewise-planar approximation of the dynamic scene surface. The deformable scene surface in the reference frame (i.e, ξ_W) and the one in the second frame (i.e, ξ'_W) are parametrized by their respective 3D superpixels $\{\tilde{s}_i\}$ and $\{\tilde{s}'_i\}$, where each $\tilde{s}_i$ is described by its surface normal n_i and an anchor point $\tilde{x}_{ai}$. Any 3D plane can be determined by an anchor point $\tilde{x}_{ai}$ and a surface normal n_i . If one can estimate correct placement of all the 3D anchor points and all the surface normals corresponding to the reference frame, the problem is solved, since each element of ξ_W is related to ξ'_W via $SE(3)$ transformation (locally rigid).

The overall procedure of our method is presented in **Algorithm 1**.

4.4 Formulation

We begin by briefly reiterating some of our representation. We partition the reference image into a set ξ_I , whose corresponding set in the 3D world is ξ_W . Equivalently, ξ_I and ξ'_W are the respective sets for the next frame. The mapping of each element in the reference frame and next frame differs by a rigid transformation. Mathematically, $\xi_W \mapsto \xi'_W$ via $SE(3)$ transformation (also known as special euclidean

Algorithm 1 : SuperPixel Soup

Input: Two consecutive image frames of a dynamic scene and dense optical flow correspondences between them.

Output: 3D reconstruction for both images.

1. Divide the reference image into ‘N’ superpixels and construct a K-NN graph to represent the entire scene as a graph $G(V, E)$ defined over these superpixels §4.4.
2. Employ two-view epipolar geometry to recover the rigid motion and shape for each 3D superpixel §4.5.
3. Optimize the proposed energy function to assemble (or glue) and align all the reconstructed superpixels (“3D Superpixel Jigsaw Puzzle”) §4.5.2.

Note: The procedure of the above algorithm looks simple; there is, however, a fundamental difficulty of scale indeterminacy in its execution.

group), for instance $\tilde{x}'_{ai} = M_i \tilde{x}_{ai}$ where $\tilde{x}'_{ai} \in \xi'_W$ and $\tilde{x}_{ai} \in \xi_W$. In our formulation each 3D plane is described by $\Phi_{\tilde{s}_i} = \{(\Pi_{\tilde{s}_i}, M_i) \mid \forall i \in [1, N]\}$, where N is the total number of superpixels (see Fig. 3). Similarly, in the image space $\xi_I \mapsto \xi'_I$ through the plane-induced homography $s'_i = K \left(R_i - \frac{\lambda_i t_i n_i}{\lambda_i d_i} \right) K^{-1} s_i$ [8]¹. Here, K is the intrinsic camera matrix and d_i is the depth of the plane. Using these notations and definitions, we build a K-NN graph.

Build a K-NN graph: Using over-segmentation of the reference image ξ_I (which is the projection of a set of 3D planes $\Phi_{\tilde{s}_i}$) and Euclidean distance metric, we construct a K-NN graph $G(V, E)$ in the *image space* connecting each anchor point to its K-nearest anchor points. The graph vertices (V) are composed of anchor point that connects to other anchor points via graph edges (E). The distance between any two vertices ($E_i \subset E$) is taken as the Euclidean distance between them. Here, we assume Euclidean distance as a valid graph metric to describe the edge length between any two local vertices. Such an assumption is valid for local compactness (Euclidean spaces are locally compact). Interested reader may refer to [47] [48] [49] for comprehensive details. Here, ‘K’ is the number of nearest neighbor that is used to construct local graph structure. This K-NN graph relation helps to constrain the motion and continuity of the space (defined in terms of optical flow, depth). To impose a hard constraint, we build a K-NN graph using anchor point beyond its immediate neighbors (Fig. 4).

This K-NN graph is crucial in the establishment of local rigidity constraint which is the basis of our assumption. This graph structure allows us to enforce our assumption *i.e*, the shape to be as rigid as possible globally and rigid locally.

As-Rigid-As-Possible (ARAP) Energy Term: Our method is built on the idea that the correct scales of 3D superpixels can be estimated by enforcing prior assumptions that govern the deformation of the dynamic surface. Specifically, we require that, locally, the motion that each 3D-superpixel undergoes is rigid, and globally the entire dynamic scene surface must move as rigid as possible (ARAP). In other words, while the dynamic scene is globally non-rigid, its deformation must be *regular* in the sense that

1. scale λ_i is introduced both in the numerator and denominator for clarification that scale does not affect the homography transformation.

Fig. 3: Illustration shows the modeling of a continuous scene with a piece wise rigid and planar assumption. Each superpixel is composed of a set (Π_{s_i}, M_i) where Π_{s_i} contains geometric parameters such as normal, anchor point, boundary points of a plane in 3D and M_i contains the motion parameters i.e rotation and translation.

it deforms as rigidly as possible. To implement this idea, we define an ARAP-energy term as:

$$\begin{aligned} E_{\text{arap}} = & \sum_{i=1}^N \sum_{k \in \mathcal{N}_i} w_1(\mathbf{x}_{\text{ai}}, \mathbf{x}_{\text{ak}}) \left(\|\mathbf{R}_i - \mathbf{R}_k\|_F + \|\lambda_i \hat{\mathbf{t}}_i - \lambda_k \hat{\mathbf{t}}_k\|_2 \right) \\ & + w_2(\mathbf{x}_{\text{ai}}, \mathbf{x}_{\text{ak}}) \cdot \left| \|\tilde{\mathbf{x}}_{\text{ai}} - \tilde{\mathbf{x}}_{\text{ak}}\|_2 - \|\tilde{\mathbf{x}}'_{\text{ai}} - \tilde{\mathbf{x}}'_{\text{ak}}\|_2 \right|. \end{aligned} \quad (1)$$

Here, the first term favors smooth motion between the local neighbors, while the second term encourages inter-node distances between the anchor node and its K nearest neighbor nodes (denoted as $k \in \mathcal{N}_i$) to be preserved before and after motion (hence as-rigid-as-possible, see Fig. 4). We define the weighting parameters as:

$$w_1(\mathbf{x}_{\text{ai}}, \mathbf{x}_{\text{ak}}) = w_2(\mathbf{x}_{\text{ai}}, \mathbf{x}_{\text{ak}}) = \exp(-\beta \|\mathbf{x}_{\text{ai}} - \mathbf{x}_{\text{ak}}\|). \quad (2)$$

These weights are set to be inversely proportional to the distance between two superpixels. This is to reflect our intuition that, the further apart two superpixels are, the weaker the E_{arap} energy is. Although there may be redundant information in these two terms w.r.t scale estimation, we keep them for motion refinement §4.5.2. Note that, this term is only defined over anchor points, hence it enforces no depth smoothness along boundaries. The weighting term in E_{arap} advocates the local rigidity by penalizing over the distance between anchor points. This allows immediate neighbors to have smooth deformation over time. Also, note that E_{arap} is generally *non-convex*. This non-convexity is due to the second term in Eq. 1, where we have a minus sign between two l_2 norm terms. In Eq. 2 β is an empirical constant.

E_{arap} alone is good enough to provide reasonably correct scales, however, the piece-wise planar composition of a continuous 3D space creates discontinuity near the boundaries of each plane. For this reason, we incorporate additional constraint to fix this depth discontinuity and further refine motions and geometry for each superpixel via neighboring relations. We call these constraints as Planar Re-projection, 3D Continuity and Orientation Energy constraint.

Planar Re-projection Energy Term: With the assumption that each superpixel represents a plane in 3D, it must satisfy corresponding planar reprojection error in 2D image space.

- K-NN to i^{th} Anchor node in relation to reference frame
- K-NN to i^{th} Anchor node in relation to next frame
- i^{th} Anchor node ($\tilde{\mathbf{x}}_{\text{ai}}$)

Fig. 4: Demonstration of as rigid as possible constraint. Superpixel segmentation in the reference frame is used to decompose the entire scene as a set of anchor points. Schematic representation shows the construction of K-NN around a particular anchor point (shown in Red). We constrain the local 3D coordinate transformation both before and after motion (green shows K-NN the reference frame, yellow shows the relation in the next frame (after motion)). We want this transformation to be as rigid as possible.

This reprojection cost reflects the average dissimilarity in the optical flow correspondences across the entire superpixel due to motion. Therefore, it helps us to constrain the surface normal, rotation and translation direction such that they obey the observed planar homography in the image space. To infer any pixel inside a superpixel, we use the operator $\psi(\cdot)$, for e.g $\psi(s_i^j)$ will give the coordinates of j^{th} pixel inside s_i . Using it we define

$$\begin{aligned} E_{\text{proj}} = & \sum_{i=1}^N \frac{w_3}{|\psi(s_i)|} \sum_{j=1}^{|\psi(s_i)|} \|\psi(s_i^j)' - \mathbf{K}(\mathbf{R}_i - \frac{\mathbf{t}_i \mathbf{n}_i^T}{d_i}) \mathbf{K}^{-1} \psi(s_i^j)\|_2. \end{aligned} \quad (3)$$

Here, $\psi(s_i^j)$, $\psi(s_i^j)'$ is the optical flow correspondence of j^{th} pixel inside i^{th} superpixel in the reference frame and next frame respectively. The operator $|\cdot|$ represent the cardinal number of a set. w_3 is a trade-off scalar chosen empirically. A natural question that may arise is: *This term is independent of scale, then what's the purpose of this constraint? How does it help?* Kindly, refer to §4.5.2 for details.

3D Continuity Energy Term: In case of a dynamic scene, where both camera and the objects are in motion, its quite apparent that the scene will undergo some changes across frames. Hence, to assume unremitting global continuity with a piece-wise planar assumption, in a dynamic scene is unreasonable. Instead, local weak continuity constraint can be enforced —a constraint that can be broken occasionally [50] i.e., local planes are connected to few of its neighbors. Accordingly, we want to allow local neighbors to be piece-wise continuous. To favor this continuous or smooth surface reconstruction, we require neighboring superpixels to have a smooth depth transition at their boundaries. To do so, we

Fig. 5: 3D Continuity energy favors continuous surface for the planes that shares the common boundary points. a)-d) The lesser the E_{cont} is, smoother the surface becomes (color bar shows the energy).

define a 3D continuity energy term as:

$$E_{\text{cont}} = \sum_{i=1}^N \sum_{k \in \mathcal{N}_i} w_4(\mathbf{X}_{\text{bi}}, \mathbf{X}_{\text{bk}}) \left(\|\tilde{\mathbf{X}}_{\text{bi}} - \tilde{\mathbf{X}}_{\text{bk}}\|_F + \rho(\|\tilde{\mathbf{X}}'_{\text{bi}} - \tilde{\mathbf{X}}'_{\text{bk}}\|_F) \right) \quad (4)$$

where, $\mathbf{X}$, $\tilde{\mathbf{X}}$ represents the corresponding matrices in 2D image space and 3D Euclidean space ($\mathbf{X}_{\text{bi}} \in \mathbb{R}^{2 \times B_i}$, $\tilde{\mathbf{X}}_{\text{bi}} \in \mathbb{R}^{3 \times B_i}$, where B_i is the total number of boundary pixel for i^{th} superpixel). Since in our representation, geometry and motion are shared among all pixels within a superpixel, so regularization within the superpixel is not explicitly needed. Thus, we only concentrate on the shared boundary pixels to regularize our energy. Note that the neighboring relationship in E_{cont} is different from E_{arap} term. Here, the neighbors share common boundaries with each other.

To encourage the geometry to be approximately smooth locally if the object has similar appearance, we color weight the energy term along the boundary pixels. For each boundary pixel of a given superpixel, we consider its 4-connected neighboring pixels to weight. Using this idea for w_4 we obtain:

$$w_4(\mathbf{X}_{\text{bi}}, \mathbf{X}_{\text{bk}}) = \sum_{j=1}^4 \exp(-\beta \|\mathbf{I}(\mathbf{X}_{\text{bi}}) - \mathbf{I}(\zeta_j)\|_F) \quad (5)$$

which weigh the inter-plane transition by color difference. The symbol $\zeta_j \in \mathbb{R}^{2 \times B_i}$ is a set that contains the 4 connecting pixels to each i^{th} superpixel boundary pixel shared with k^{th} superpixel. The color based weighting term plays an important role to allow for “weak continuity constraint” i.e gradually allow for occasional discontinuity [50] [51].

To better understand the implication of E_{cont} constraint, consider two boundary points in the image space $a, b \in \mathbb{R}^2$. Generally, if these two points lie on a different plane, it will not coincide in the 3D space before and after motion. Hence, we compute the 3D distance between boundary pixels corresponding to both reference frame and next frame, which leads to our goal of penalizing distance along shared edges (see Fig. 5). Therefore, this term ensures the 3D coordinates across superpixel boundaries to be continuous in both frames. The challenge here is to reach a satisfactory solution for overall scene continuity, almost everywhere in

Fig. 6: a) Superpixelated reference image b) Individual superpixel depth with arbitrary scale (unorganised superpixel soup) c) recovered depth map using our approach (organised superpixel soup) d) ground-truth depth map.

both the frames [5]. In the Eq. 4 ρ is a truncation function defined as $\rho = \min(., \sigma)$ and similar to Eq. 2 β in Eq. 5 is a constant, chosen empirically.

Orientation Energy Term: To encourage the smoothness in the orientation of the neighboring planes, we added one more geometric constraint i.e, E_{orient} defined as follows.

$$E_{\text{orient}} = \sum_{i=1}^N \sum_{k \in \mathcal{N}_i} \rho_n(1 - \mathbf{n}_i^T \mathbf{n}_k) \quad (6)$$

Here neighbor index is same as 3D continuity term. ρ_n denotes the truncated l_1 penalty function which is defined as $\rho_n(x) = \min(|x|, n)$. Intuitively, it encourages the similarity between neighboring normal’s and truncate any value more than n .

Combined Energy Function: Equipped with all these constraints, we define our overall cost function or energy function to obtain a scale consistent 3D reconstruction of a complex dynamic scene. Our goal is to estimate depth (d_i), surface normal ($\mathbf{n}_i$) and scale λ_i for each 3D planar superpixel. The key is to estimate the unknown relative scale λ_i . We solve this by minimizing the following energy function:

$$\begin{aligned} \min_{\lambda_i, \mathbf{n}_i, d_i, \mathbf{R}_i, \mathbf{t}_i} E &= E_{\text{arap}} + \alpha_1 E_{\text{proj}} + \alpha_2 E_{\text{cont}} + \alpha_3 E_{\text{orient}} \\ \text{subject to } \sum_{i=1}^N \lambda_i &= 1, \lambda_i > 0. \\ \mathbf{R}_i &\in \mathbb{SO}(3), \|\mathbf{n}_i\|_2 = 1. \end{aligned} \quad (7)$$

The equality constraint on λ fixes the unknown freedom of a global scale. The constraint on $\mathbf{R}_i$ is imposed to restrict the rotation matrix to lie on $\mathbb{SO}(3)$ manifold. In our formulation, the rotation matrix represents the combined Euler 3D angles. Although there are other efficient representations for 3D rotation, we used $\mathbf{R}_i \in \mathbb{R}^{3 \times 3}$ matrix representation as it comes naturally via epipolar geometric constraint, hence, further post-conversion steps can be avoided. The constant $\alpha_1, \alpha_2, \alpha_3$ are included for numerical consistency.

4.5 Implementation

We partition the reference image into 1,000-2,000 superpixels [52]. Parameters such as $\alpha_1, \alpha_2, \alpha_3, \beta, \sigma$ were tuned differently for different datasets. To perform optimization of the proposed energy function (Eq. 7), we require initial set of proposals for motion and geometry.

4.5.1 Initial Proposal Generation

We exploit piece-wise rigid and planar assumption to estimate an initial proposal for geometry and motion. We start by estimating homography ($\mathbf{H}_{\pi_1}$) for each superpixel using dense feature correspondences. Piece-wise rigid assumption helps in approximate estimation of rotation and correct translation direction via triangulation and chirality check

[8] [6]. To obtain the correct normal direction and initial depth estimate, we solve the following set of equations for each superpixel:

$$\mathbf{H}_{\pi_i} = \mathbf{K}(\mathbf{R}_i - \frac{\mathbf{t}_i \mathbf{n}_i^T}{d_i}) \mathbf{K}^{-1} \quad (8)$$

The reason we choose this strategy to obtain normal is because a simple decomposition of homography matrix to the rotation, translation and normal can lead to sign ambiguity [42] [53]. Nevertheless, if one has correct rotation and direction of translation –which we infer from chirality check, then inferring normal becomes easy². Here, we assume the depth ‘ d_i ’ to be a positive constant and the initial arbitrary reconstruction is in the +Z direction. This strategy of gathering 9-dimensional variables (6-motion variable and 3-geometry variable) for each individual superpixel gives us a good enough estimate to get started with the minimization of our overall energy function³.

To initialize 3D vectors in our formulation we use the following well known relation:

$$\tilde{\mathbf{x}}_{ai} = \left[\left(\frac{u_{ai} - c_x}{f_x} \right), \left(\frac{v_{ai} - c_y}{f_y} \right), 1/\mathbf{n}_i^T \mathbf{K}^{-1} \begin{pmatrix} u_{ai} \\ v_{ai} \\ 1 \end{pmatrix} \right]^T (\lambda_i d_i) \quad (9)$$

where, (u_{ai}, v_{ai}) are image coordinates and (c_x, c_y, f_x, f_y) are camera intrinsic parameters which can be inferred from $\mathbf{K}$ matrix.

4.5.2 Optimization

With good enough initialization of the variables, we start to optimize our energy function Eq. 7. A global optimal solution is hard to achieve due to the non-convex nature of the proposed cost function (Eq. 7). However, it can be solved efficiently using interior-point methods [28] [29]. Although the solution found by the interior point method is at best local minimizer, empirically they appear to give good 3D reconstruction. In our experiments, we initialized all λ ’s with an initial value of $\frac{1.0}{N}$.

Next, we employ a particle based refinement algorithm to rectify our initial motion and geometry beliefs. Specifically, we used the Max-Product Particle Belief propagation (MP-PBP) procedure with the TRW-S algorithm [30] to optimize over the surface normals, rotations, translations and depths for all 3D superpixels using Eq. 10. We generated 50 particles as proposals for the unknown parameters around the already known beliefs to initiate refinement moves. Repeating this strategy for 5-10 iterations, we obtain a smooth and refined 3D structure of the dynamic scene.

$$\mathbf{E}_{\text{ref}} = \mathbf{E}_{\text{arap}} + \alpha_1 \mathbf{E}_{\text{proj}} + \alpha_2 \mathbf{E}_{\text{cont}} + \alpha_3 \mathbf{E}_{\text{orient}}. \quad (10)$$

5 EXPERIMENTS AND RESULTS

We evaluated our formulation both qualitatively and quantitatively on various standard benchmark datasets, namely MPI Sintel [24], KITTI [31], VKITTI [32] and You-Tube Object dataset [23]. All these dataset contains images of dynamic scene where both camera and objects are in motion w.r.t each

other. To test the reconstruction result on deformable objects we used Paper, T-shirt [42] [54] and Back Sequence [36]. For evaluating the result, we selected the most commonly used error metric *i.e.*, mean relative error metric.

Evaluation Metric: To keep the evaluation metric consistent with the previous work [1], we used mean relative error (MRE) metric for evaluation. MRE is defined as $\frac{1}{P} \sum_{i=1}^P |z_{\text{gt}}^i - z_{\text{est}}^i| / z_{\text{gt}}^i$. Here, $z_{\text{est}}^i, z_{\text{gt}}^i$ denotes the estimated and ground-truth depth respectively with P being the total number of points. The error is computed after re-scaling the recovered depth properly as the reconstruction is obtained up to an unknown global scale. Quantitative evaluation for the YouTube-Objects dataset and the Back dataset are missing due to the absence of ground-truth result.

To show that our same formulation works well for both rigid and non-rigid cases, we evaluated our method with different types of scene that contain *rigid, non-rigid, complex dynamic scene i.e., composition of both rigid and non-rigid*.

5.1 Experimental Setup and Results

Experimental setup and processing time: We partition the reference image using SLIC superpixels [52]. We used the current state-of-the-art optical flow algorithm to compute dense optical flow [44]. To initialize the motion and geometry variables, we used the the procedure discussed in §4.5.1. Interior point algorithm [28] [29] and TRW-S [30] were employed to solve the proposed optimization. We implemented our algorithm in MATLAB/C++. Our modified implementation (modified from our ICCV implementation [55]) takes an average of 15-20 minutes to provide the result for images of size 1024×436 . The processing time is estimated on a regular desktop with Intel core i7 processor (16 GB RAM) for 50 refinement particle per superpixel.

Results on MPI Sintel Dataset: We begin our analysis of experimental results with MPI Sintel dataset [24]. This dataset is derived from an animation movie featuring complex scenes. It contains highly dynamic sequences with large motions, significant illumination changes, and non-rigidly moving objects. This dataset has emerged as a standard benchmark to evaluate dense optical flow algorithm’s and recently, it has also been used in the evaluation of dense 3D reconstruction methods for a general dynamic scene [1].

The presence of non-rigid objects in the scene makes it a prominent choice for us to test our algorithm. It is a challenging dataset particularly for the piece-wise planar assumption due to the presence of many small and irregular shapes in the scene. Additionally, the presence of ground-truth depth map makes quantitative analysis much easier. We selected 120 pair of images to test our method that includes images from *alley_1, ambush_4, mountain_1, sleeping_1 and temple_2*. Fig. 7 shows some qualitative results on a few images taken from the sub-group of MPI Sintel dataset.

Results on VKITTI Dataset: The Virtual KITTI dataset [32] contains computer-rendered photo-realistic outdoor driving scenes which resemble KITTI dataset. The advantage of using this dataset is that it provides perfect ground-truths for many measurements. Furthermore, it helps to simulate algorithm related to dense 3D reconstruction with distortion-free and noise-free images, facilitating quick experimentation. We selected 120 pair of images from *0001_morning, 0002_morning, 0006_morning and*

2. The solution to the obtained normal must be normalized.

3. If the size of the superpixel is very small, use the neighboring superpixels optical flow to estimate motion parameters.

Fig. 7: Qualitative results using our algorithm in a complex dynamic scene. Example images are taken from MPI Sintel dataset [24]. **Top row:** Input reference image from *sleeping_1*, *sleeping_2*, *shaman_3*, *temple_2*, *alley_2* sequence (from left to right). **Middle row:** Ground-truth depth map for the respective frames. **Bottom row:** Recovered depth map using our method.

Fig. 8: Qualitative results using our algorithm for the outdoor scenes. Examples are taken from VKITTI dataset [32]. **Top row:** Input reference image. **Middle row:** Ground-truth depth map for the respective frames. **Bottom row:** Recovered depth map using our method.

0018_morning. Our qualitative results in comparison to the ground-truth depth map are shown in Fig. 8.

Results on KITTI Dataset: The KITTI dataset [31] features the real-world outdoor scene targeting autonomous driving application. The KITTI images are taken from the camera mounted on top of a car. It’s a challenging dataset as it contains scenes with large camera motion and realistic lighting condition. In contrast to the aforementioned datasets, it only contains sparse ground-truth 3D information which makes evaluation a bit strenuous. Nonetheless, it captures noisy real-world situation and therefore, it is well suited to test the 3D reconstruction algorithm for a general dynamic scene case. We selected 00-09 sub-category from odometry dataset to evaluate and compare our results. We calculated mean relative error only over the provided sparse 3D LiDAR points –after adjusting the global scale. Fig. 9 shows some qualitative results on few images.

Results on Non-Rigid Sequence We also tested our method on some commonly used dense non-rigid sequence namely *kinect_paper* [42], *kinect_tshirt* [42] and *back sequence* [36]⁴. Most of the benchmark approach to solve non-rigid structure from motion use multiple frames and orthographic camera model. Despite a two-frame method and perspective camera model, we are able to capture the deformation of non-rigid object and achieve its reliable reconstruction. Qualitative results for dense non-rigid object sequence are shown in Fig. 10. To compute the mean relative error, we align and scale our shape (fixing global ambiguity) w.r.t ground-truth shape.

5.2 Comparison

We compared the performance of our algorithm against several dynamic reconstruction methods, namely, Block Matrix Method (BMM) [12], Point Trajectory Approach (PTA) [56],

4. Note: The intrinsic matrix for back sequence is not available with the dataset, we estimated an approximate value of it using 2D-3D relation available from Garg *et. al.* [36].

Fig. 9: Qualitative results on KITTI Dataset [31]. The second row shows the obtained depth map for the respective frames. Note: Dense ground-truth depth data is not available with this dataset.

Fig. 10: Dense 3D reconstruction of the objects that are undergoing non-rigid deformation over frames. **Top row:** Input reference frame from *Back sequence* [36], *Paper sequence* [42] [54] and *t-shirt sequence* [42] [54]. **Bottom row:** Qualitative 3D reconstruction results for the respective deforming object.

Low-rank Reconstruction (GBLR) [57]), Depth Transfer (DT) [58], DMDE [1] and ULDEMv [43]. This comparison is made on the available benchmark datasets i.e., MPI Sintel (MPI-S), KITTI, VKITTI, kinect_tshirt (k_tshirt), kinect_paper (k_paper). Table 1 provides the statistical result of our method in comparison to the baseline approach on these datasets. Our method outperforms others in the outdoor sequence and provides a commendable performance for deformable sequence. Additionally, we performed a qualitative comparison on MPI Sintel [24], KITTI [31] and YouTube object dataset [23]. Fig. 11 and Fig. 12 provides the visual comparison result of our method to other competing methods. It can be observed that our method consistently delivers superior performance on all of these datasets. While compiling the results per frame comparison is also made over the entire sequence. Evaluation in the case of KITTI dataset is done only for the provided sparse 3D LiDAR points. Fig. 13(a), Fig. 13(b) and Fig. 14(c) shows per category statistical performance of our approach against other competing methods on the benchmark dataset.

5.3 Performance Analysis

Besides statistical comparison, we conducted other experiments to analyze the behavior of our algorithm. These experiments supply an in-depth understanding of the dependency of our algorithm on other input modules.

Performance with variation in number of superpixels: Our method uses SLIC based over-segmentation of the reference frame to discretize the 3D space. Therefore, the number of superpixels that represent the real-world plays a crucial

Method Dataset ↓	DT (SF)	GLRT (MF)	BMM (MF)	PTA (MF)	DMDE (TF)	Ours (TF)
MPI-S	0.4833	0.4101	0.3121	0.3177	0.297	0.1643
V-KITTI	0.2630	0.3237	0.2894	0.2742	-	0.0925
KITTI	0.2703	0.4112	0.3903	0.4090	0.148	0.1254
k_paper	0.2040	0.0920	0.0322	0.0520	-	0.0472
k_tshirt	0.2170	0.1030	0.0443	0.0420	-	0.0480

TABLE 1: Performance Comparison: This table lists the MRE errors. For DMDE [1] we used its previously reported results as its implementation is not publicly available. SF, MF, TF refers to single frame, multi-frame and two-frame based approach respectively. The reference to the method DT [58], GLRT [57], BMM [12], PTA [56], DMDE [1].

Fig. 11: Qualitative comparison of our method with DMDE [1] on MPI Sintel [24] and KITTI Dataset [24]. **Left to Right:** For each input reference image, we show its ground-truth depth map (GT Depth), depth map reported by DMDE [1] and depth map obtained using our approach. Note: Dense GT depth map for KITTI Dataset is taken from DMDE [1] work.

Fig. 12: Qualitative evaluation of our approach with the Video-PopUp [39]. Clearly, our method provides more dense and detailed reconstruction of the scene. In the second row t-shirt description is missing with Video-PopUp [39] approach. By contrast our method has no such holes. Note: The results presented here for Video-PopUp are taken from their webpage since the source code provided by the authors crashes frequently.

role in the accuracy of piece-wise continuous reconstruction. If the number of superpixels are very high the estimation of motion parameters becomes tricky and therefore, neighboring superpixels are used to estimate rigid motion which leads to computation challenges. In contrast, small number of superpixels are unable to capture the intrinsic details of a complex dynamic scene. So, a trade-off between the two is often a better choice. Fig. 14(a) shows the plot of depth error variation with the change in the number of superpixels.

Performance with regular grid as image superpixel: Under the piece-wise planar assumption, its not only the number of superpixels that affects the accuracy of reconstruction but also the type of superpixel pattern. To analyze this dependency, we took the worst possible case i.e to divide the reference image into approximately 1000 regular grid and compare its performance against 1000 SLIC superpixel. Our observation clearly shows a decline in the performance in comparison to SLIC superpixels. However, the difference in accuracy is not very significant (see Fig. 15).

Effects of K in K-NN Graph: In our method, the ARAP energy term is evaluated using the K nearest neighbor graph. Different K value leads to different 3D reconstruction result. An experiment on the flying dragon sequence is conducted to analyze the effect of varying K on the performance of our algorithm. The result of the flying dragon case is shown in Fig. 16. With the increase in K, the rigidity con-

straint is enforced in an increased neighborhood which directs the 3D reconstruction towards a globally rigid solution. On the other hand, a very small value of K fails to constrain the within object motion. In most of our experiments, we used a K in the range of 15 – 20 which achieved satisfactory 3D reconstruction. Also, increasing the value of K directly affects the overall algorithmic complexity.

Performance variation using different optical flow algorithm: As our method uses dense optical flow correspondences between frames as input, the performance of our method is directly affected by its accuracy. To analyze the sensitivity of our method to different optical flow methods, we conducted experiments by testing our method with the ground-truth optical flow and few state-of-the-art optical flow methods [44] [45]. In Fig. 14(b), we show the 3D reconstruction performance evaluated in RMSE⁵ with different optical flow as inputs. This experiment reveals the importance of dense optical flow in the accurate reconstruction of a dynamic scene. While ground truth optical flow naturally achieves the best performance, the difference in result using different state-of-the-art optical flow is not dramatic. Therefore, we conclude that our method can achieve reliable results with the available dense optical flow algorithm's.

6 LIMITATIONS AND DISCUSSION

The success of our method depends on the effectiveness of the piece-wise planar and as rigid as possible assumption. As a result, our method may fail if the piece-wise smooth model is no longer a valid approximation for the dynamic scene. For example, very fine or very small structures which are considerably far from the camera are difficult to recover under the piecewise planar assumption. Further, what about as rigid as possible assumption, *When as rigid as possible assumption may fail?* When the motion of the dynamic objects between consecutive frame is significantly large such that most of its neighboring relations in the reference frame get violated in the next frame. Additionally, if the non-rigid shape shrinks or expands over frames such as a *deflating or inflating balloon*, ARAP model fails. A couple of examples for such situations are discussed in Fig. 17. The other major limitation of our method is the overall processing time.

6.1 Discussion

1. *Direction to reduce the processing time of our algorithm:* Our algorithm is computationally expensive to execute on a regular desktop machine. This is due to the formulation for solving a higher-order graph optimization problem and particle-based refinement using TRW-S. To speed up the processing time, we are implementing some of the recent research work in the field of fast interior-point optimization and message-passing algorithm [59], [60] to our framework. We believe solving our optimization using these algorithms along with better computation capabilities can significantly reduce the processing time of our method.
2. *Suitability of euclidean distance metric between graph vertices:* Generally, the euclidean distance metric between graph

5. RMSE (Root Mean Square Error) = $\sqrt{\frac{1}{P} \sum_{i=1}^P (z_{gt}^i - z_e^i)^2}$, here z_e^i , z_{gt}^i denotes the estimated and ground-truth depth respectively and P is the total number of points.

Fig. 13: Quantitative comparison with our method with PTA [56], BMM [12], GLRT [57], DT [58] on benchmark datasets. The depth error is calculated by adjusting the numerical scale of the obtained depth map to the ground-truth value, to account for global scale ambiguity. (a)-(b) comparison on MPI Sintel [24], Virtual KITTI [32] and KITTI [31] dataset. These numerical values show the fidelity of reconstruction that can be retrieved on these benchmark datasets using our formulation.

Fig. 14: (a) Change in mean relative depth error with the change in number of superpixels. It can be observed that after 1000 superpixel the MRE more or less starts saturating with no significant effect on the overall accuracy. However, it was observed that the motion estimation becomes critical with the increase in number of superpixels. (b) Performance evaluation in RMSE (in meters) with the state-of-the-art optical flow methods in comparison to the ground-truth optical flow (MPI Sintel [24] dataset). (c) Mean Relative Depth Error comparison with a recently proposed unsupervised learning based approach (ULDEMV [43]) on KITTI dataset [31].

Fig. 15: Effects of superpixel pattern on the reconstruction of a dynamic scene. a) with SLIC as superpixels (MRE for the shown frame is 0.0912) b) with uniform grid as superpixels (MRE achieved for the given frame is 0.1442).

Fig. 16: Effect of parameter K in building the K -NN graph. Our algorithm results in good reconstruction if a suitable K is chosen, in accordance with the levels of complexity in a dynamic scene. (b) Ground-truth depth-map (scaled for illustration purpose). (c) when $K=4$, a reasonable reconstruction is obtained. (d) when $K=20$, regions tend to grow bigger. (Best viewed in color.)

Fig. 17: (a)-(b) are the reference frame and the next frame. It is a very challenging case for proper scale recovery with monocular images with dynamic motion. In both of these cases the motion of the girl between two consecutive frames is very large and therefore, the neighboring relations with the planes (say superpixels in image domain) in the consecutive frames gets violated. In such cases, our method may not be able to provide correct scales for each moving planes in 3D. In the first example, the complicated motion of the feet of the girl leads to wrong scale estimation. In the second example, the cart along with girl is moving w.r.t the camera. The hand of the girl has a substantial motion in the consecutive frames which leads to incorrect estimation of scale. (c)-(d) Ground-truth and obtained depth map respectively.

vertices works well under our piece-wise planar assumption of a dynamic scene. However, there are situations where it may not be an appropriate metric. For example: when the shape of the superpixels is affected by noise or modeling of curved spaces using a piece-wise planar graph structure. To handle such special cases its better to measure distance in an embedding space (isometric embedding) or use l_1 metric, etc. To be very precise, depending on the shape of the deforming structure over time, the choice of a suitable metric may vary. Interested readers are encouraged to study the field of intrinsic metric on graphs [61].

7 ABLATION ANALYSIS

To understand the contribution of different energy term in the overall optimization, we performed ablation analysis. Firstly, in the proposed optimization framework the *3D continuity* term is defined over boundaries between neighboring superpixels, which alone is not sufficient to constrain the motion beyond its immediate neighbors. Secondly, E_{proj} and E_{orient} has nothing to do with scale computation whatsoever. Hence, combining these three terms is not good enough to explain the correct scales for each of the object present in the scene. On the other hand, *as rigid as possible* term is defined for each superpixel's anchor point over the K -NN graph structure. However, it does not take into account the alignment of planes in 3D along the boundaries. As a result, the overall reconstruction suffers. Thus, this demonstrates that all the terms are essential for reliable dynamic 3D reconstruction. Fig. 18 illustrates the contribution of different terms toward the final reconstruction result. Table 2 provides numerical value showing the importance of different terms on the overall performance of our algorithm. It can be observed that the improvement in output due to normal orientation constraint is not very significant.

Fig. 18: Effect of using “as rigid as possible”, “Planar re-projection”, “3D continuity” and “Orientation” term. **Top row:** By enforcing the “as rigid as possible” term only, the recovered relative scales are correct but the reconstructed planes are misaligned with respect to their neighbors. **Middle row:** With the planar re-projection, 3D continuity and orientation term enforced, the resultant 3D reconstruction achieves continuous neighboring boundaries, however, the relative scales for every plane in 3D are not correct. **Bottom row:** By enforcing the “as rigid as possible” term along with all the other smoothness terms, we can handle both relative scales and 3D reconstruction for a complex dynamic scene.

Data	E_{arap}	$+E_{proj}$	$+E_{cont}$	$+E_{orient}$
alley_1	0.2248	0.2022	0.1697	0.1606
ambush_4	0.2381	0.2093	0.1701	0.1676
mountain_1	0.2127	0.1923	0.1492	0.1405
sleeping_1	0.2418	0.2026	0.1912	0.1823

TABLE 2: Contribution of each individual energy term to the overall optimization. Each column show the mean relative reconstruction error due the addition of the respective energy term. The + sign symbolizes the addition of the all the energy term (columns) left of it.

8 CONCLUSION

In this paper, we have explored, investigated and supplied a distinct perspective to solve one of the classical problems in geometric computer vision i.e., to reconstruct a dense 3D model of a *complex, dynamic, and generally non-rigid* scene from its two perspective images. This topic of research is often considered as a very challenging task in structure from motion. In spite of its reasonable challenges, we have demonstrated that dense detailed 3D reconstruction of dynamic scenes is, in fact possible, provided that certain prior assumptions about the scene geometry and the deformation in the scene are satisfied. Both the assumptions we used are mild, realistic and commonly satisfied by the real-world scenarios. Our comprehensive evaluation on the benchmark datasets shows that our new insight to solve dense monocular 3D reconstruction of a general dynamic scene provides better results than other competing methods. This said, we think more profound research on top of our idea may help in the development of sophisticated SfM algorithms.

We believe our algorithm provides a plausible new direction to perceive a complex dynamic scene with a single monocular camera. Lastly, we want to stress on the point that rigidity is a powerful concept in SfM, careful or acute extensions of the present approach may open up a new direction to advance dense 3D reconstruction from images.

Acknowledgements: This research is supported in part by the Australia Research Council ARC Centre of Excellence for Robotics Vision (CE140100016), ARC-Discovery (DP 190102261) and ARC-LIEF (190100080), The Natural Science Foundation of China grants (61871325, 61420106007, 61671387), the “New Generation of Artificial Intelligence” major project under Grant 2018AAA0102800, and ARC grant DE140100180 and in part by a research gift from Baidu RAL (ApolloScapes-Robotics and Autonomous Driv-

ing Lab). The authors gratefully acknowledge the Data Science GPU gift award by NVIDIA Corporation. We thank all the reviewers and AE for their constructive suggestions.

REFERENCES

- [1] R. Ranftl, V. Vineet, Q. Chen, and V. Koltun, “Dense monocular depth estimation in complex dynamic scenes,” in *CVPR*, 2016, pp. 4058–4066.
- [2] S. Ullman, “The interpretation of structure from motion,” *Proceedings of the Royal Society of London B: Biological Sciences*, vol. 203, no. 1153, pp. 405–426, 1979.
- [3] W. E. L. Grimson, *From images to surfaces: A computational study of the human early visual system*. MIT press, 1981.
- [4] H. C. Longuet-Higgins, “A computer algorithm for reconstructing a scene from two projections,” *Nature*, vol. 293, no. 5828, pp. 133–135, 1981.
- [5] A. Blake and A. Zisserman, *Visual reconstruction*. MIT press, 1987.
- [6] R. I. Hartley and P. Sturm, “Triangulation,” *Computer vision and image understanding*, vol. 68, no. 2, pp. 146–157, 1997.
- [7] R. I. Hartley, “In defense of the eight-point algorithm,” *T-PAMI*, vol. 19, no. 6, pp. 580–593, 1997.
- [8] R. Hartley and A. Zisserman, *Multiple view geometry in computer vision*. Cambridge university press, 2003.
- [9] C. Tomasi, “Pictures and trails: a new framework for the computation of shape and motion from perspective image sequences,” Cornell University, Tech. Rep., 1993.
- [10] C. Tomasi and T. Kanade, “Shape and motion from image streams under orthography: a factorization method,” *IJCV*, vol. 9, no. 2, pp. 137–154, 1992.
- [11] C. Bregler, A. Hertzmann, and H. Biermann, “Recovering non-rigid 3d shape from image streams,” in *CVPR*, vol. 2. IEEE, 2000, pp. 690–696.
- [12] Y. Dai, H. Li, and M. He, “A simple prior-free method for non-rigid structure-from-motion factorization,” *IJCV*, vol. 107, no. 2, pp. 101–122, 2014.
- [13] M. Lee, J. Cho, C.-H. Choi, and S. Oh, “Procrustean normal distribution for non-rigid structure from motion,” in *CVPR*, 2013, pp. 1280–1287.
- [14] S. Kumar, Y. Dai, and H. Li, “Multi-body non-rigid structure-from-motion,” in *Fourth International Conference on 3D Vision (3DV)*, 2016. IEEE, 2016, pp. 148–156.
- [15] S. Kumar, Y. Dai, and H. Li, “Spatio-temporal union of subspaces for multi-body non-rigid structure-from-motion,” *Pattern Recognition*, vol. 71, pp. 428–443, 2017.
- [16] R. A. Newcombe, D. Fox, and S. M. Seitz, “Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time,” in *CVPR*, 2015, pp. 343–352.
- [17] R. A. Newcombe, S. J. Lovegrove, and A. J. Davison, “Dtm: Dense tracking and mapping in real-time,” in *ICCV*. IEEE, 2011, pp. 2320–2327.
- [18] S. Agarwal, Y. Furukawa, N. Snavely, I. Simon, B. Curless, S. M. Seitz, and R. Szeliski, “Building rome in a day,” *Communications of the ACM*, vol. 54, no. 10, pp. 105–112, 2011.
- [19] J. L. Schönberger and J.-M. Frahm, “Structure-from-motion revisited,” in *CVPR*, 2016.
- [20] J. L. Schönberger, E. Zheng, M. Pollefeys, and J.-M. Frahm, “Pixel-wise view selection for unstructured multi-view stereo,” in *ECCV*, 2016.
- [21] V. M. Govindu, “Combining two-view constraints for motion estimation,” in *CVPR*, vol. 2. IEEE, 2001, pp. II–II.
- [22] B. Triggs, P. F. McLauchlan, R. I. Hartley, and A. W. Fitzgibbon, “Bundle adjustment: a modern synthesis,” in *International workshop on vision algorithms*. Springer, 1999, pp. 298–372.
- [23] A. Prest, C. Leistner, J. Civera, C. Schmid, and V. Ferrari, “Learning object class detectors from weakly annotated video,” in *CVPR*. IEEE, 2012, pp. 3282–3289.
- [24] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black, “A naturalistic open source movie for optical flow evaluation,” in *ECCV*. Springer, 2012, pp. 611–625.
- [25] M. Bleyer, C. Rother, P. Kohli, D. Scharstein, and S. Sinha, “Object stereojoint stereo matching and object segmentation,” in *CVPR* 2011. IEEE, 2011, pp. 3081–3088.
- [26] C. Vogel, K. Schindler, and S. Roth, “3d scene flow estimation with a piecewise rigid scene model,” *IJCV*, vol. 115, pp. 1–28, 2015.
- [27] —, “3d scene flow estimation with a rigid motion prior,” in *ICCV*. IEEE, 2011, pp. 1291–1298.

- [28] H. Y. Benson, R. J. Vanderbei, and D. F. Shanno, "Interior-point methods for nonconvex nonlinear programming: Filter methods and merit functions," *Computational Optimization and Applications*, vol. 23, no. 2, pp. 257–272, 2002.
- [29] H. Y. Benson and D. F. Shanno, "Interior-point methods for non-convex nonlinear programming: cubic regularization," *Computational optimization and applications*, vol. 58, no. 2, pp. 323–346, 2014.
- [30] V. Kolmogorov, "Convergent tree-reweighted message passing for energy minimization," *T-PAMI*, vol. 28, no. 10, pp. 1568–1583, 2006.
- [31] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The KITTI dataset," *Int. J. Rob. Res.*, vol. 32, no. 11, pp. 1231–1237, Sep. 2013.
- [32] A. Gaidon, Q. Wang, Y. Cabon, and E. Vig, "Virtual worlds as proxy for multi-object tracking analysis," *arXiv preprint arXiv:1605.06457*, 2016.
- [33] H. C. Longuet-Higgins, "A computer algorithm for reconstructing a scene from two projections," *Readings in Computer Vision: Issues, Problems, Principles, and Paradigms*, MA Fischler and O. Firschein, eds, pp. 61–62, 1987.
- [34] S. Kumar, A. Cherian, Y. Dai, and H. Li, "Scalable dense non-rigid structure-from-motion: A grassmannian perspective," in *IEEE, CVPR*, 2018, pp. 254–263.
- [35] S. Kumar, "Jumping manifolds: Geometry aware dense non-rigid structure from motion," in *IEEE, CVPR*, June 2019.
- [36] R. Garg, A. Roussos, and L. Agapito, "Dense variational reconstruction of non-rigid surfaces from monocular video," in *CVPR*, 2013, pp. 1272–1279.
- [37] J. Fayad, L. Agapito, and A. Del Bue, "Piecewise quadratic reconstruction of non-rigid surfaces from monocular sequences," *ECCV 2010*, pp. 297–310, 2010.
- [38] J. Taylor, A. D. Jepson, and K. N. Kutulakos, "Non-rigid structure from locally-rigid motion," in *CVPR. IEEE*, 2010, pp. 2761–2768.
- [39] C. Russell, R. Yu, and L. Agapito, "Video pop-up: Monocular 3d reconstruction of dynamic scenes," in *ECCV*, 2014, pp. 583–598.
- [40] R. Yu, C. Russell, N. D. Campbell, and L. Agapito, "Direct, dense, and deformable: Template-based non-rigid 3d reconstruction from rgb video," in *ICCV. IEEE*, 2015, pp. 918–926.
- [41] X. Wang, M. Salzmann, F. Wang, and J. Zhao, "Template-free 3d reconstruction of poorly-textured nonrigid surfaces," in *ECCV*. Springer, 2016, pp. 648–663.
- [42] A. Varol, M. Salzmann, E. Tola, and P. Fua, "Template-free monocular reconstruction of deformable surfaces," in *ICCV. IEEE*, 2009, pp. 1811–1818.
- [43] T. Zhou, M. Brown, N. Snavely, and D. G. Lowe, "Unsupervised learning of depth and ego-motion from video," *arXiv preprint arXiv:1704.07813*, 2017.
- [44] C. Bailer, B. Taetz, and D. Stricker, "Flow fields: Dense correspondence fields for highly accurate large displacement optical flow estimation," in *CVPR*, 2015, pp. 4015–4023.
- [45] Q. Chen and V. Koltun, "Full flow: Optical flow estimation by global optimization over regular grids," *arXiv preprint arXiv:1604.03513*, 2016.
- [46] O. Sorkine and M. Alexa, "As-rigid-as-possible surface modeling," in *Symposium on Geometry processing*, vol. 4, 2007.
- [47] D. Burago, Y. Burago, and S. Ivanov, *A course in metric geometry*. American Mathematical Society Providence, 2001, vol. 33.
- [48] R. Williamson and L. Janos, "Constructing metrics with the heine-borel property," *Proceedings of the American Mathematical Society*, vol. 100, no. 3, pp. 567–573, 1987.
- [49] W. Whiteley, "Rigidity and scene analysis," 2004.
- [50] G. E. Hinton, "Relaxation and its role in vision," 1977.
- [51] A. Blake, "The least-disturbance principle and weak constraints," *Pattern Recognition Letters*, vol. 1, no. 5–6, pp. 393–399, 1983.
- [52] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Süsstrunk, "SLIC superpixels compared to state-of-the-art superpixel methods," *T-PAMI*, vol. 34, no. 11, pp. 2274–2282, 2012.
- [53] E. Malis and M. Vargas, "Deeper understanding of the homography decomposition for vision-based control," Ph.D. dissertation, INRIA, 2007.
- [54] A. Varol, M. Salzmann, P. Fua, and R. Urtasun, "A constrained latent variable model," in *CVPR. IEEE*, 2012, pp. 2248–2255.
- [55] S. Kumar, Y. Dai, and H. Li, "Monocular dense 3d reconstruction of a complex dynamic scene from two perspective frames," in *ICCV*, Oct 2017, pp. 4649–4657.
- [56] I. Akhter, Y. Sheikh, S. Khan, and T. Kanade, "Trajectory space: A dual representation for nonrigid structure from motion," *T-PAMI*, vol. 33, no. 7, pp. 1442–1456, 2011.
- [57] K. Fragkiadaki, M. Salas, P. Arbeláez, and J. Malik, "Grouping-based low-rank trajectory completion and 3d reconstruction," in *Advances in Neural Information Processing Systems*, 2014, pp. 55–63.
- [58] K. Karsch, C. Liu, and S. B. Kang, "Depth transfer: Depth extraction from video using non-parametric sampling," *T-PAMI*, vol. 36, no. 11, pp. 2144–2158, 2014.
- [59] J. W. Pearson and J. Gondzio, "Fast interior point solution of quadratic programming problems arising from pde-constrained optimization," *Numerische Mathematik*, vol. 137, pp. 959–999, 2017.
- [60] S. Tourani, A. Shekhovtsov, C. Rother, and B. Savchynskyy, "Mplp++: Fast, parallel dual block-coordinate ascent for dense graphical models," in *The ECCV (ECCV)*, September 2018.
- [61] M. Keller, "Intrinsic metrics on graphs: a survey," in *Mathematical technology of networks*. Springer, 2015, pp. 81–119.

Suryansh Kumar is a Professor für Computer Vision at ETH Zürich. He received his Ph.D. in Engineering and Computer Science from the Australian National University, Canberra, Australia. He received his M.S in Computer Science and Engineering from International Institute of Information Technology, Hyderabad (IIIT-H) in 2013. Before joining Australian National University, he worked as a Visiting Scientist in the e-Motion Group at INRIA Rhône Alpes Grenoble. After that, he spend one year as a consultant engineer in a computer vision industry-Hyderabad, India. His research interests includes geometric computer vision, robotics, abstract algebra, machine learning and mathematical optimization. He won the best algorithm award in CVPR NRSFM Challenge 2017.

Yuchao Dai is a Professor in School of Electronics and Information at Northwestern Polytechnical University, Xi'an, China. He was an ARC DECRA Fellow with the Research School of Engineering at the Australian National University, Canberra, Australia. He received the B.E. degree, M.E degree and Ph.D. degree all in signal and information processing from Northwestern Polytechnical University in 2005, 2008 and 2012, respectively. His research interests include structure from motion, multiview geometry, deep learning, compressive sensing and optimization. He won the Best Paper Award in CVPR 2012, DSTO Best Fundamental Contribution to Image Processing Paper Prize in 2014 and Best Algorithm award in CVPR NRSFM Challenge 2017. He served as an Area Chair for WACV 2019/2020 and ACM MM 2019.

Hongdong Li is a Chief Investigator of the Australia Centre of Excellence for Robotic Vision, Australian National University. He is Associate Director (for Research) for ANU School of Engineering. He was a visiting professor with the Robotics Institute, CMU during sabbatical in 2017. His research interests include geometric computer vision, pattern recognition and machine learning, vision perception for autonomous driving, and combinatorial optimization. He is an Associate Editor for IEEE TPAMI, and served as Area Chair for recent years' CVPR, ICCV and ECCV. He was the winner of CVPR Best Paper Award 2012, Marr Prize (Honorable Mention) at ICCV 2017, IEEE ICPR and IEEE ICIP Best Student Paper Winner, DSTO Best Fundamental Contribution to Image Processing Paper Prize at DICTA 2014 and Best algorithm award in CVPR NRSFM Challenge 2017. He is a program co-chair for ACCV 2018 and ACCV 2022. His research is funded by Australia Research Council, CSIRO, General Motors, Ford Motors, Microsoft Research etc. He is Presentation Co-Chair for ICCV 2019 and AC for CVPR 2020 and ECCV 2020.