



Temporally Coherent General Dynamic Scene Reconstruction

Armin Mustafa¹ · Marco Volino¹ · Hansung Kim¹ · Jean-Yves Guillemaut¹ · Adrian Hilton¹

Received: 11 July 2019 / Accepted: 4 August 2020 / Published online: 18 August 2020
© The Author(s) 2020

Abstract

Existing techniques for dynamic scene reconstruction from multiple wide-baseline cameras primarily focus on reconstruction in controlled environments, with fixed calibrated cameras and strong prior constraints. This paper introduces a general approach to obtain a 4D representation of complex dynamic scenes from multi-view wide-baseline static or moving cameras without prior knowledge of the scene structure, appearance, or illumination. Contributions of the work are: an automatic method for initial coarse reconstruction to initialize joint estimation; sparse-to-dense temporal correspondence integrated with joint multi-view segmentation and reconstruction to introduce temporal coherence; and a general robust approach for joint segmentation refinement and dense reconstruction of dynamic scenes by introducing shape constraint. Comparison with state-of-the-art approaches on a variety of complex indoor and outdoor scenes, demonstrates improved accuracy in both multi-view segmentation and dense reconstruction. This paper demonstrates unsupervised reconstruction of complete temporally coherent 4D scene models with improved non-rigid object segmentation and shape reconstruction and its application to various applications such as free-view rendering and virtual reality.

Keywords Dynamic 4D reconstruction · Segmentation · Reconstruction · 3D · Temporal coherence · Dynamic scenes

1 Introduction

Reconstruction of general dynamic scenes is of great importance in entertainment applications such as visual effects in film and broadcast production and for content production in virtual reality. The ultimate goal of modelling dynamic scenes from multiple cameras is automatic understanding

Communicated by Cristian Sminchisescu.

Electronic supplementary material The online version of this article (<https://doi.org/10.1007/s11263-020-01367-2>) contains supplementary material, which is available to authorized users.

✉ Armin Mustafa
a.mustafa@surrey.ac.uk

Marco Volino
m.volino@surrey.ac.uk

Hansung Kim
h.kim@surrey.ac.uk

Jean-Yves Guillemaut
j.guillemaut@surrey.ac.uk

Adrian Hilton
a.hilton@surrey.ac.uk

¹ Centre for Vision, Speech and Signal Processing (CVSSP),
University of Surrey, Guildford GU27XH, UK

of real-world scenes from distributed camera networks, for applications in robotics and other autonomous systems. Existing multi-view dynamic scene reconstruction methods either work in controlled environment with known background or chroma-key studio (Guillemaut and Hilton 2010; Goldluecke and Magnor 2004; Starck and Hilton 2007; Taneja et al. 2011) or require a large number of cameras (Furukawa and Ponce 2010). Extensions to more general outdoor scenes (Ballan et al. 2010; Kim et al. 2012; Taneja et al. 2011) use prior reconstruction of the static geometry from images of the empty environment. However these methods either require accurate segmentation of dynamic foreground objects, or prior knowledge of the scene structure and background, or are limited to static cameras and controlled environments. Scenes are reconstructed semi-automatically, requiring manual intervention for segmentation/rotoscoping, and result in temporally incoherent per-frame mesh geometries. Temporally coherent geometry with known surface correspondence across the sequence is essential for real-world applications and compact representation.

This paper addresses the limitations of existing approaches by introducing a methodology for unsupervised 4D temporally coherent dynamic scene reconstruction from multiple wide-baseline static or moving camera views without prior

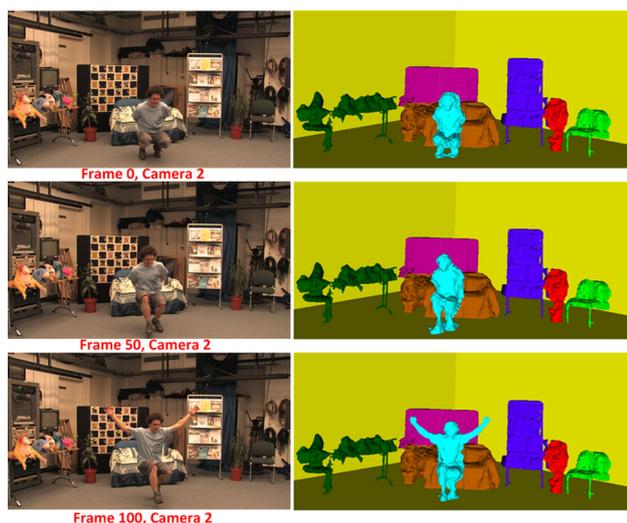


Fig. 1 Temporally consistent color-coded scene reconstruction and segmentation for Odzemek dataset (Color figure online)

knowledge of the scene structure or background appearance. The scene is automatically decomposed into a set of spatio-temporally coherent objects as shown in Fig. 1 where the resulting 4D scene reconstruction has temporally coherent labels and surface correspondence for each object. This temporally coherent dynamic scene reconstruction is demonstrated to work in applications for immersive content production such as free-viewpoint video (FVV) and virtual reality (VR). The contributions are summarized as follows:

- Unsupervised temporally coherent dense reconstruction and segmentation of general complex dynamic scenes from multiple wide-baseline views.
- Automatic initialization of dynamic object segmentation and reconstruction from sparse features.
- A framework for space–time sparse-to-dense segmentation, reconstruction and temporal correspondence.
- Robust spatio-temporal refinement of dense reconstruction and segmentation integrating error tolerant photo-consistency and edge information using geodesic star convexity.
- Robust and computationally efficient reconstruction of dynamic scenes by exploiting temporal coherence.
- Real-world applications of 4D reconstruction to free-viewpoint video rendering and virtual reality.

This paper presents a unified framework from two previously published papers, combining multiple view joint reconstruction and segmentation (Mustafa et al. 2015) with temporal coherence (Mustafa et al. 2016a) to improve per-frame reconstruction performance. In particular the approach estimates a 4D surface model with full correspondence over time. A comprehensive experimental evaluation with com-

parison to the state-of-the-art in segmentation, reconstruction and 4D modelling is also presented extending previous work. Application of the resulting 4D models to FVV rendering and content production for immersive VR experiences is also presented.

2 Related Work

Temporally coherent reconstruction is a challenging task for general dynamic scenes due to a number of factors such as motion blur, articulated, non-rigid and large motion of multiple people, resolution differences between camera views, occlusions, wide-baselines, errors in calibration and cluttered dynamic backgrounds. Segmentation of dynamic objects from such scenes is difficult because of foreground and background complexity and the likelihood of overlapping background and foreground color distributions. Reconstruction is also challenging due to limited visual cues and relatively large errors affecting both calibration and extraction of a globally consistent solution. This section reviews previous work on dynamic scene reconstruction and segmentation.

2.1 Dynamic Scene Reconstruction

Dense dynamic shape reconstruction is a fundamental problem and heavily studied area in the field of computer vision. Recovering accurate 3D models of a dynamically evolving, non-rigid scene observed by multiple synchronised cameras is a challenging task. Research on multiple view dense dynamic reconstruction has primarily focused on indoor scenes with controlled illumination and static backgrounds, extending methods for multiple view reconstruction of static scenes (Seitz et al. 2006) to sequences (Tung et al. 2009). Deep learning based approaches have been introduced to estimate shape of dynamic objects from minimal camera views in constrained environment (Huang et al. 2018; Wu et al. 2018) and for rigid objects (Stutz and Geiger 2018). In the last decade, focus has shifted to more challenging outdoor scenes captured with both static and moving cameras. Reconstruction of non-rigid dynamic objects in uncontrolled natural environments is challenging due to the scene complexity, illumination changes, shadows, occlusion and dynamic backgrounds with clutter such as trees or people. Methods have been proposed for multi-view reconstruction (Vo et al. 2016; Lei et al. 2009; Larsen et al. 2007) requiring a large number of closely spaced cameras for surface estimation of dynamic shape. Practical applications require relatively sparse moving cameras to acquire coverage over large areas such as outdoor. A number of approaches for multi-view reconstruction of outdoor scenes require initial silhouette segmentation (Wu 2013; Kim et al. 2012; Guan et al. 2010; Guillemaut

and Hilton 2010) to allow visual-hull reconstruction. Most of these approaches to general dynamic scene reconstruction fail in the case of complex (cluttered) scenes captured with moving cameras.

A recent work proposed reconstruction of dynamic fluids (Qian et al. 2017) for static cameras. Another work used RGB-D cameras to obtain reconstruction of non-rigid surfaces (Slavcheva et al. 2017). Pioneering research in general dynamic scene reconstruction from multiple handheld wide-baseline cameras (Ballan et al. 2010; Taneja et al. 2011) exploited prior reconstruction of the background scene to allow dynamic foreground segmentation and reconstruction. Recent work (Ngo et al. 2019) estimates shape of dynamic objects from handheld cameras exploiting GANs. However these approaches either work for static/indoor scenes or exploit strong prior assumptions such as silhouette information, known background or scene structure. Also all these approaches give per frame reconstruction leading to temporally incoherent geometries. Our aim is to perform temporally coherent dense reconstruction of unknown dynamic non-rigid scenes automatically without strong priors or limitations on scene structure.

2.2 Multi-view Video Segmentation

In the field of image segmentation, approaches have been proposed to provide temporally consistent monocular video segmentation (Grundmann et al. 2010; Papazoglou and Ferrari 2013; Narayana et al. 2013; Zhang et al. 2013). Hierarchical segmentation based on graphs was proposed in Grundmann et al. (2010), directed acyclic graph was used to propose an object segmentation (Zhang et al. 2013) and optical flow is used to consistently segment objects (Narayana et al. 2013; Papazoglou and Ferrari 2013). Recently a number of approaches have been proposed for multi-view foreground object segmentation by exploiting appearance similarity spatially across views (Djelouah et al. 2013; Kowdle et al. 2012; Lee et al. 2011; Zeng and Zeng 2004) and space–time similarity (Djelouah et al. 2015). However, multi-view approaches assume a static background and different colour distributions for the foreground and background which limits applicability for general scenes and non-rigid objects.

To address this issue we introduce a novel method for spatio-temporal multi-view segmentation of dynamic scenes using shape constraints. Single image segmentation techniques using shape constraints provide good results for complex scenes (Gulshan et al. 2010) (convex and concave shapes), but require manual interaction. The proposed approach performs automatic multi-view video segmentation by initializing the foreground object model using spatio-temporal information from wide-baseline feature correspondence followed by a multi-layer optimization framework. Geodesic star convexity previously used in single view seg-

mentation (Gulshan et al. 2010) is applied to constraint the segmentation in each view. Our multi-view formulation naturally enforces coherent segmentation between views and resolves ambiguities such as the similarity of background/foreground in isolated views.

2.3 Joint Segmentation and Reconstruction

Joint segmentation and reconstruction methods incorporate estimation of segmentation or matting with reconstruction to provide a combined solution. Joint refinement avoids the propagation of errors between the two stages thereby making the solution more robust. Also, cues from segmentation and reconstruction can be combined efficiently to achieve more accurate results. The first multi-view joint estimation system was proposed by Szeliski and Golland (1998) which used iterative gradient descent to perform an energy minimization. A number of approaches were introduced for joint formulation in static scenes and one recent work used training data to classify the segments (Zach et al. 2013). The focus shifted to joint segmentation and reconstruction for rigid objects in indoor and outdoor environments. These approaches used a variety of techniques such as patch-based refinement (Shin et al. 2013; Ozden et al. 2007) and fixating cameras on the object of interest (Campbell et al. 2010) for reconstructing rigid objects in the scene. However, these are either limited to static scenes (Zach et al. 2013; Hane et al. 2013) or process each frame independently thereby failing to enforce temporal consistency (Campbell et al. 2010; Guillemaut and Hilton 2010).

Joint reconstruction and segmentation on monocular video was proposed in Kundu et al. (2014), Atapour-Abarghouei and Breckon (2019), and Chen et al. (2019) achieving semantic segmentation of scene limited to rigid objects in street scenes. Practical application of joint estimation requires these approaches to work on non-rigid objects such as humans with clothing. A multi-layer joint segmentation and reconstruction approach was proposed for multiple view video of sports and indoor scenes (Guillemaut and Hilton 2010). The algorithm used known background images of the scene without the dynamic foreground objects to obtain an initial segmentation. Visual-hull based reconstruction was performed with known prior foreground/background using a background image plate with fixed and calibrated cameras. This visual hull was used as a prior and was optimized by a combination of photo-consistency, silhouette, color and sparse feature information in an energy minimization framework to improve the segmentation and reconstruction quality. Although structurally similar to our approach, it requires the scene to be captured by fixed calibrated cameras and a priori known fixed background plate as a prior to estimate the initial visual hull by background subtraction. The proposed approach overcomes these limitations allowing moving cam-

eras and unknown scene backgrounds. These methods are able to produce high quality results, but rely on good initializations and strong prior assumptions with known and controlled (static) scene backgrounds.

To overcome the limitations of existing methods, the proposed approach automatically initialises the foreground object segmentation from wide-baseline correspondence without prior knowledge of the scene. This is followed by a joint spatio-temporal reconstruction and segmentation of general scenes.

2.4 Temporally Coherent 4D Reconstruction

Temporally coherent 4D reconstruction refers to aligning the 3D surfaces of non-rigid objects in time for a dynamic sequence. This is achieved by estimating point-to-point correspondences for the 3D surfaces to obtain temporally coherent 4D reconstruction. 4D models allow to create efficient representation for practical applications in film, broadcast and immersive content production such as virtual, augmented and mixed reality. The existing approaches for multi-view reconstruction of dynamic scenes process each time frame independently. Independent per-frame reconstruction can result in errors due to the inherent visual ambiguity caused by occlusion and similar object appearance for general scenes.

3D scene flow (Menze and Geiger 2015) estimates frame-to-frame correspondence or exploits 2D optical flow (Wedel et al. 2011; Basha et al. 2010). These methods require an accurate estimate for most of the pixels which fails in the case of large motion. However, 3D scene flow methods align two frames independently and do not produce temporally coherent 4D models across the complete sequence to obtain a single surface model. Research investigating spatio-temporal reconstruction across multiple frames was proposed by Goldluecke and Magnor (2004); Larsen et al. (2007); Guillemaut and Hilton (2012); Mustafa et al. (2016b) exploiting the temporal information from the previous frames using optical flow. An approach for recovering space–time consistent depth maps from multiple video sequences captured by stationary, synchronized and calibrated cameras for depth based free viewpoint video rendering was proposed by (Lei et al. 2009). However these methods require accurate initialisation, fixed and calibrated cameras and are limited to simple scenes. Other approaches to temporally coherent reconstruction either require a large number of closely spaced cameras (Bailer et al. 2015) or bi-layer segmentation (Zhang et al. 2011; Jiang et al. 2012) as a constraint for reconstruction. Recent approaches for spatio-temporal reconstruction of multi-view data work on indoor data (Oswald et al. 2014).

The proposed framework addresses limitations of existing approaches and gives 4D temporally coherent reconstruction for general dynamic indoor or outdoor scenes with large

non-rigid motions, repetitive texture, uncontrolled illumination, and large capture volume. The proposed approach gives 4D models of complete scenes with both static and dynamic objects for real-world applications (FVV and VR) with no prior knowledge of scene structure.

2.5 Summary and Motivation

Image-based temporally coherent 4D dynamic scene reconstruction without constraints is a key problem in computer vision. Existing dense reconstruction algorithms need some strong initial prior and constraints for the solution to converge such as background, structure, and segmentation, which limits their application for automatic reconstruction of general scenes. Current approaches are also commonly limited to independent per-frame reconstruction and do not exploit temporal information or produce a 4D model with known correspondence. The approach proposed in this paper aims to overcome the limitations of existing approaches by initializing the joint reconstruction and segmentation algorithm automatically, introducing temporal coherence in the reconstruction and geodesic star convexity in segmentation to reduce ambiguity and ensure consistent non-rigid structure initialization at successive frames.

3 Methodology

An overview of the proposed framework for temporally coherent multi-view reconstruction is presented in Fig. 2 and consists of the following stages.

Multi-view video The scenes are captured using multiple video cameras (static/moving) separated by wide-baseline ($> 15^\circ$). Cameras can be synchronised either directly or in post-processing using the audio information. The cameras can be synchronized using time-code generator or later using the audio information. Camera intrinsics are known. Camera extrinsics (location, orientation) and scene structure are assumed to be unknown.

Sparse reconstruction Segmentation based feature detection (SFD) (Mustafa et al. 2015, 2019) is used to detect sparse features distributed throughout the scene for wide-baseline matching. SFD features are matched between views using a SIFT descriptor. The camera extrinsics and sparse 3D feature locations are then estimated for each time instant for the entire sequence (Hartley and Zisserman 2003).

Initial coarse reconstruction: Sect. 3.1 The sparse point cloud is clustered in 3D (Rusu 2009) with each cluster representing a unique foreground object. Automatic initialisation is performed without prior knowledge of the

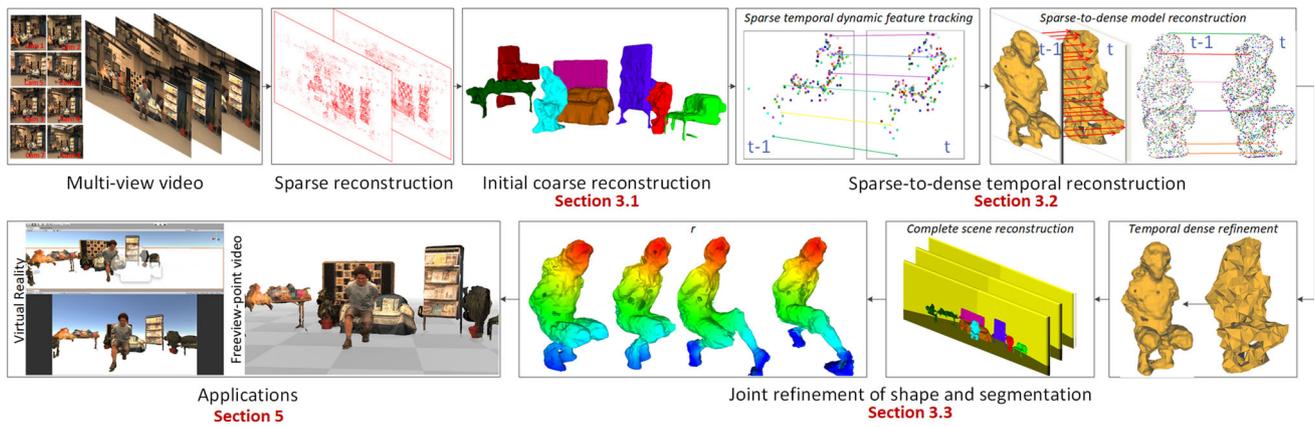


Fig. 2 Overview of temporally consistent scene reconstruction framework

scene structure or appearance to obtain an initial approximation for each object.

Sparse-to-dense temporal reconstruction: Sect. 3.2 Temporal coherence is introduced to initialize the coarse reconstruction and obtain frame-to-frame dense correspondences for dynamic objects. Sparse temporal correspondence helps in identifying dynamic objects and allows propagation of the dense reconstruction between time instants to obtain an initialization.

Joint refinement of shape and segmentation: Sect. 3.3 The initial estimate is refined for each object per-view through joint optimisation of shape and segmentation using a robust cost function combining matching, color, contrast and smoothness information with a geodesic star convexity constraint. A single 3D model for each dynamic object is obtained by fusion of the view-dependent depth maps using Poisson surface reconstruction (Kazhdan et al. 2006). Surface orientation is estimated based on neighbouring pixels.

3.1 Initial Coarse Reconstruction

For general dynamic scene reconstruction, we need to reconstruct and segment the objects in the scene. This requires an initial coarse approximation for initialisation of a subsequent refinement step to optimise the segmentation and reconstruction. Sparse point cloud clustering is used to segment objects, an overview is shown in Fig. 3. The dense reconstruction of the foreground objects and background are combined to obtain a full scene reconstruction at the first time instant. A coarse geometric proxy of the background is created. For consecutive time instants dynamic objects and newly appeared objects are identified and only these objects are reconstructed and segmented. The reconstruction of static objects is retained which reduces computational complexity. The optic flow and cluster information for each dynamic

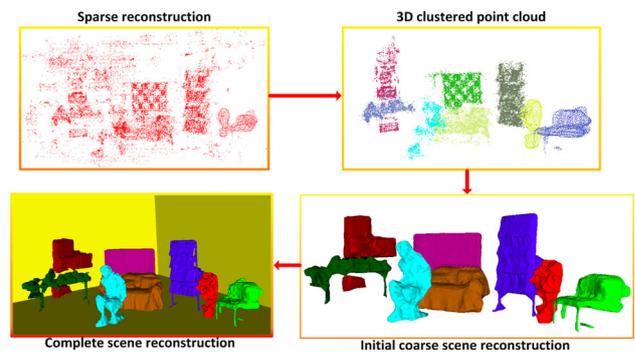


Fig. 3 Overview of stages for estimation of an initial dense scene reconstruction. For more details refer to Sect. 3.1

object ensures that we retain consistent labels for the entire sequence.

3.1.1 Background Reconstruction

Accurate reconstruction of the background is often challenging due to uniform appearance of large regions. A coarse geometric proxy of the background is created by computing the minimum oriented bounding box for the sparse 3D point-cloud using principal component analysis (Dimitrov et al. 2006). Different methods are used for background estimation for indoor and outdoor scenes. For outdoor scenes a plane is inserted at infinity perpendicular to the ground plane as there are no consistent constraints. For indoor scenes the Manhattan world assumption (Coughlan and Yuille 2000) is applied to estimate room structure. The process used for estimation of the background is described below:

- The centroid $\mathbf{A} = (a_0, a_1, a_2)$ and normalized covariance of the point-cloud are estimated to compute the eigenvectors $\vec{e} = (e_0, e_1, e_2)$ for the covariance matrix of the point-cloud. We define the reference system as $\mathbf{R} = (e_0, e_1, e_0 \times e_1)$ such that: $e_0 \times e_1 = \pm e_2$. Rotation

matrix \mathbf{R} and translation \mathbf{A} are used to map sparse points in the first frame and place a box in correct location.

- The minimum and maximum values of coordinates in the x , y and z directions for the transformed cloud are computed to determine the minimum oriented box width, height, and depth.
- Given a box centred at the origin with size defined above the rotation \mathbf{R} and translation $\mathbf{R} \times \mathbf{C} + \mathbf{A}$ is applied, where \mathbf{C} is the middle of the minimum and maximum points.

This background reconstruction is a rough geometric proxy estimate of the background of the scene but gives reasonable results for complete scene reconstruction.

3.1.2 Sparse Point-Cloud Clustering

The sparse representation of the scene is processed to remove outliers using the point neighbourhood statistics to filter outlier data (Rusu 2009). We segment the objects in the sparse scene reconstruction, this allows only moving objects to be reconstructed at each frame for efficiency and this also allows object shape similarity to be propagated across frames to increase robustness of reconstruction. Object segmentation increases efficiency and improve robustness of 4D models.

We use data clustering approach based on the 3D grid subdivision of the space using an octree data structure in Euclidean space to segment objects at each frame (Rusu 2009). In a more general sense, nearest neighbor information is used to cluster, which is essentially similar to a flood fill algorithm. We choose this data clustering because of its computational efficiency and robustness. The approach allows segmentation of objects in the scene and is demonstrated to work well for cluttered and general outdoor scenes as shown in Sect. 4.

Objects with insufficient detected features are reconstructed as part of the scene background. Appearing, disappearing and reappearing objects are handled by sparse dynamic feature tracking, explained in Sect. 3.2. Clustering results are shown in Fig. 3. This is followed by a sparse-to-dense coarse object based approach to segment and reconstruct general dynamic scenes.

3.1.3 Coarse Object Reconstruction

The process to obtain the coarse reconstruction for the first frame of the sequence is shown in Fig. 4. The sparse representation of each element is back-projected on the rectified image pair for each view. Delaunay triangulation (Fortune 1997) is performed on the set of back projected points for each cluster on one image and is propagated to the second image using the sparse matched features. Triangles with edge length greater than the median length of edges of all triangles

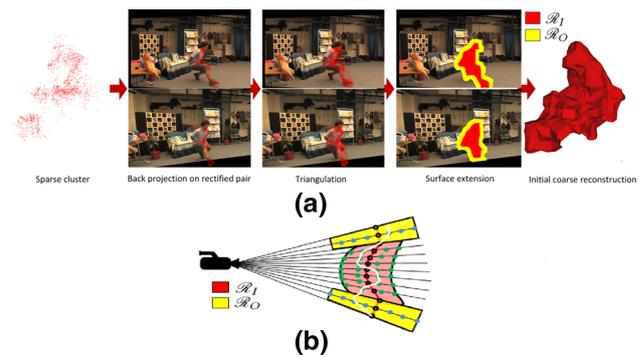


Fig. 4 Initial coarse reconstruction: **a** sparse-to-dense initial coarse reconstruction of the dynamic object in Odzemok dataset; and **b** white line represents the actual surface, Depth labels are represented as circles; blue circles depict depth labels in \mathcal{D}_O , green circles depict depth labels in \mathcal{D}_I and black circles depict the initial surface estimate (Color figure online)

are removed. For each remaining triangle pair direct linear transform is used to estimate the affine homography. Displacement at each pixel within the triangle pair is estimated by interpolation to get an initial dense disparity map for each cluster in the 2D image pair labelled as \mathcal{R}_I depicted in red in Fig. 4. The initial coarse reconstruction for the observed objects in the scene is used to define the depth hypotheses at each pixel for the optimization.

The region \mathcal{R}_I does not ensure complete coverage of the object, so we extrapolate this region to obtain a region \mathcal{R}_O (shown in yellow) in 2D by 5% of the average distance between the boundary points (\mathcal{R}_I) and the centroid of the object. To allow for errors in the initial approximate depth from sparse features we add volume in front and behind of the projected surface by an error tolerance, along the optical ray of the camera. This ensures that the object boundaries lie within the extrapolated initial coarse estimate. The tolerance for extrapolation may vary if a pixel belongs to \mathcal{R}_I or \mathcal{R}_O as the propagated pixels of the extrapolated regions (\mathcal{R}_O) may have a high level of error compared to error at the points from sparse representation (\mathcal{R}_I) requiring a comparatively higher tolerance. The calculation of threshold depends on the capture volume of the datasets and is set to 1% of the capture volume for \mathcal{R}_O and half the value for \mathcal{R}_I . This volume in 3D corresponds to our initial coarse reconstruction of each object and enables us to remove the dependency of previous approaches on static background plane and visual hull estimates. This process of cluster identification and initial coarse object reconstruction is performed for multiple objects in general environments. Initial object segmentation using point cloud clustering and coarse segmentation is insensitive to parameters. Throughout this work the same parameters are used for all datasets. The result of this process is a coarse initial object segmentation and reconstruction for each object.

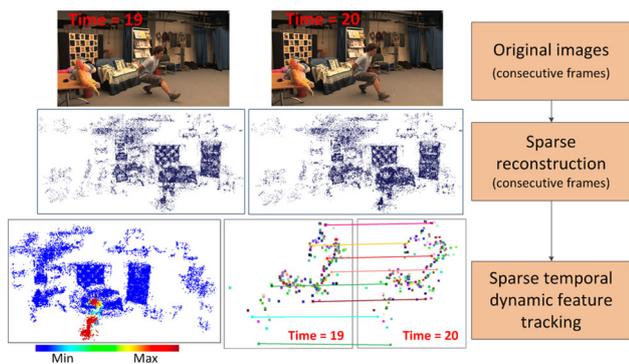


Fig. 5 Sparse temporal dynamic feature tracking algorithm: results on Odzemek dataset and Juggler dataset (Ballan et al. 2010) captured with only moving cameras; Min and Max is the minimum and maximum movement in the 3D points respectively

3.2 Sparse-to-Dense Temporal Reconstruction

Once the static scene reconstruction is obtained for the first frame, we perform temporally coherent reconstruction for dynamic objects at successive time instants instead of whole scene reconstruction for computational efficiency and to avoid redundancy. Dynamic objects are identified from the temporal correspondence of sparse feature points (Sect. 3.2.1), shown in Fig. 5. Sparse correspondence is used to propagate and obtain an initial model of the moving object for refinement (Sect. 3.2.2). The initial coarse reconstruction for each dynamic region is refined in the subsequent optimization step with respect to each camera view.

3.2.1 Sparse Temporal Dynamic Feature Tracking

Numerous approaches have been proposed to track moving objects in 2D using either features or optical flow. However these methods may fail in the case of occlusion, movement parallel to the view direction, large motions and moving cameras. To overcome these limitations we match the sparse 3D feature points obtained using SFD (Mustafa et al. 2015, 2019) from multiple wide-baseline views at each time instant. The use of sparse 3D features is robust to large non-rigid motion, occlusions and camera movement. SFD detects sparse features which are stable across wide-baseline views and consecutive time instants for a moving camera and dynamic scene. Sparse 3D feature matches between consecutive time instants are back-projected to each view. These features are matched temporally using SIFT descriptor to identify the corresponding moving points. Robust matching is achieved by enforcing multiple view consistency for the temporal feature correspondence in each view as illustrated in Fig. 6. Each match must satisfy the constraint:

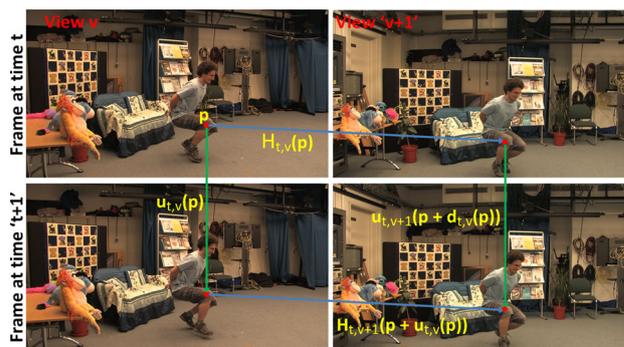


Fig. 6 Spatio-temporal consistency check for 3D tracking for Odzemek dataset

$$\|H_{t,v}(p) + u_{t,r}(p + H_{t,v}(p)) - u_{t,v}(p) - H_{t,r}(p + u_{t,v}(p))\| < \epsilon \tag{1}$$

where p is the feature image point in view v at frame t , $H_{t,v}(p)$ is the disparity at frame t from views v and r , $u_{t,v}(p)$ is the temporal correspondence from frames t to $t+1$ for view v . The multi-view consistency check ensures that correspondences between any two views remain temporally consistent for successive frames. Matches in the 2D domain are sensitive to camera movement and occlusion, hence we map the set of refined matches into 3D to make the system robust to camera motion. The Frobenius norm is applied on the 3D point gradients in all directions (Zhang et al. 2013) to obtain the ‘net’ motion at each sparse point. The ‘net’ motions between pairs of 3D points for consecutive time instants are ranked, and the top and bottom 5% values are removed followed by Median filtering to identify the dynamic features. New objects are identified per frame from the clustered sparse reconstruction and are labelled as dynamic objects.

3.2.2 Sparse-to-Dense Model Reconstruction

Dynamic 3D feature points are used to initialize the segmentation and reconstruction of the initial model. This avoids the assumption of static backgrounds and prior scene segmentation commonly used to initialise multiple view reconstruction with a coarse visual-hull approximation (Guillemaut and Hilton 2010). Temporal coherence also provides a more accurate initialisation to overcome visual ambiguities at individual frames. Figure 7 illustrates the use of temporal coherence for reconstruction initialisation and refinement. Dynamic feature correspondence is used to identify the mesh for each dynamic object. This mesh is back projected on each view to obtain the region of interest. Lucas Kanade Optical flow (Bouguet 2000) is performed on the projected mask for each view in the temporal domain using the dynamic feature correspondences over time as initialization. Dense multi-view wide-baseline correspondences from the previous

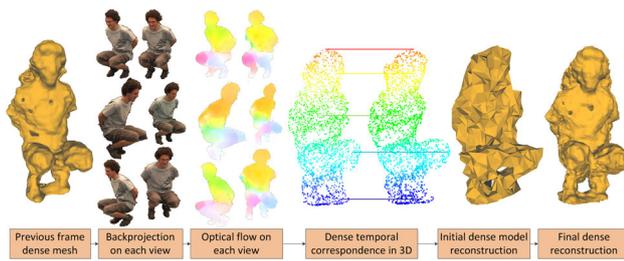


Fig. 7 Initial sparse-to-dense model reconstruction workflow



Fig. 8 Improvement in segmentation for the Odzemok dataset and reconstruction for the Juggler dataset with temporal coherence (highlighted in yellow) (Color figure online)

frame are propagated to the current frame using the information from the flow vectors to obtain dense points in the current frame. The matches are triangulated in 3D to obtain a refined 3D dense model of the dynamic object for the current frame.

For dynamic scenes, to allow the introduction of new objects and object parts we use information from the cluster of sparse points for each dynamic object. The cluster corresponding to the dynamic features is identified and static points are removed. This ensures that the set of new points not only contain the dynamic features but also the unprocessed points which represent new parts of the object. These points are added to the refined sparse model of the dynamic object. To handle the new objects we detect new clusters at each time instant and consider them as dynamic regions. Once we have a set of dense 3D points for each dynamic object, Poisson surface reconstruction (Kazhdan et al. 2006) is performed on the set of sparse points to obtain an initial coarse model of each dynamic region \mathcal{R} . The depth of region \mathcal{R} is refined per view for each dynamic object at a per pixel level.

The sparse-to-dense initial coarse reconstruction improves the quality of segmentation and reconstruction after the refinement. Examples of the improvement in segmentation and reconstruction for Odzemok [1] and Juggler (Ballan et al. 2010) datasets are shown in Fig. 8. As observed limbs of the people is correctly reconstructed by using information from the previous frames in both the cases.

3.3 Joint Refinement of Shape and Segmentation

The initial reconstruction \mathcal{R} and segmentation (\mathcal{R} projected in views) from dense temporal feature correspondence is refined using a joint optimization framework. View-dependent optimisation of depth is performed with respect to each camera which is robust to errors in camera calibration and initialisation. Calibration inaccuracies produce inconsistencies limiting the applicability of global reconstruction techniques which simultaneously consider all views; view-dependent techniques are more tolerant to such inaccuracies because they only use a subset of the views for reconstruction of depth from each camera view.

Our goal is to assign an accurate depth value from a set of depth values $\mathcal{D} = \{d_1, \dots, d_{|\mathcal{D}|-1}, \mathcal{U}\}$ and assign a layer label from a set of label values $\mathcal{L} = \{l_1, \dots, l_{|\mathcal{L}|}\}$ to each pixel p for the region \mathcal{R} of each dynamic object. Each d_i is obtained by sampling the optical ray from the camera and \mathcal{U} is an unknown depth value to handle occlusions. This is achieved by optimisation of a joint cost function (Guillemaut and Hilton 2010) for label (segmentation) and depth (reconstruction):

$$E(l, d) = \lambda_{data} E_{data}(d) + \lambda_{contrast} E_{contrast}(l) + \lambda_{smooth} E_{smooth}(l, d) + \lambda_{color} E_{color}(l) \quad (2)$$

where d is the depth at each pixel, l is the layer label for multiple objects and the cost function terms are defined in Sect. 3.3.2. The equation consists of four terms: the data term is for the photo-consistency scores, the smoothness term is to avoid sudden peaks in depth and maintain the consistency and the color and contrast terms are to identify the object boundaries. Data and smoothness terms are common to solve reconstruction problems (Bleyer et al. 2011) and the color and contrast terms are used for segmentation (Kolmogorov et al. 2006). This is solved subject to a geodesic star-convexity constraint on the labels l .

3.3.1 Shape Constraint for Joint Optimization

A novel shape constraint is introduced based on geodesic star convexity which has previously been shown to give improved performance in interactive image segmentation for structures with fine details (for example a person's fingers or hair) (Gulshan et al. 2010). Previous methods used shape constraints by enforcing star convexity prior to improve segmentation (Veksler 2008; Vicente et al. 2008). However star-convexity constraints fail for non-rigid objects (humans), as illustrated in Fig. 9. To handle complex objects the geodesic star convexity prior with multiple star centres was introduced in interactive segmentation for 2D objects (Gulshan et al. 2010). The notion of connectivity was extended from Euclidean to

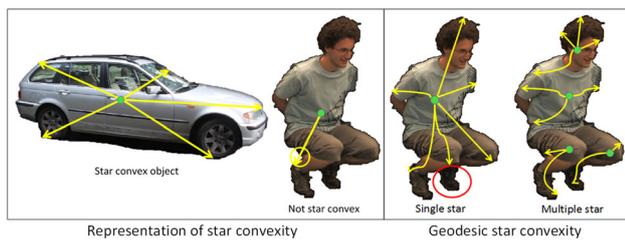


Fig. 9 **a** Representation of star convexity: the left object depicts example of star-convex objects, with a star center marked. The object on the right with a plausible star center shows deviations from star-convexity in the fine details, and **b** multiple star semantics for joint refinement: single star center based segmentation is depicted on the left and multiple star is shown on the right

geodesic space so that paths adapt to image data as opposed to straight Euclidean rays, thus extending visibility and reducing the number of star centers required. The union formed by multiple object seeds and geodesic paths form a geodesic forest (Gulshan et al. 2010).

In this work the geodesic star-convexity shape constraint is automatically initialised for each view from the initial segmentation and sparse features to constrain the energy minimisation for joint multi-view reconstruction and segmentation. The shape constraint is based on the geodesic distance as star centres to which the object shape is restricted. This allows complex shapes to be segmented. To automatically initialize the segmentation we use the sparse temporal feature correspondence as star centers (seeds) to build a geodesic forest.

The region outside the initial coarse reconstruction of all dynamic objects is initialized as the background seed for segmentation as shown in Fig. 11. The shape of the dynamic object is restricted by this geodesic distance constraint that depends on the image gradient. Comparison with existing methods for multi-view segmentation demonstrates improvements in recovery of fine detail structure as illustrated in Fig. 11.

In Eq. (2) a label l is star convex with respect to center c , if every point $p \in l$ is visible to a star center c via l in the image x which can be expressed as an energy cost:

$$E^*(l|x, c) = \sum_{p \in R} \sum_{q \in \Gamma_{c,p}} E_{p,q}^*(l_p, l_q) \tag{3}$$

$$\forall q \in \Gamma_{c,p}, \quad E_{p,q}^* = \begin{cases} \infty & \text{if } l_p \neq l_q \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

where $\forall p \in R : p \in l \Leftrightarrow l_p = 1$ and $\Gamma_{c,p}$ is the geodesic path joining p to the star center c given by:

$$\Gamma_{c,p} = \arg \min_{\Gamma \in P_{c,p}} L(\Gamma) \tag{5}$$

where $P_{c,p}$ denotes the set of all discrete paths between c and p and $L(\Gamma)$ is the length of discrete geodesic path as defined in (Gulshan et al. 2010). In the case of image segmentation the gradients in the underlying image provide information to compute the discrete paths between each pixel and star centers and $L(\Gamma)$ is defined below:

$$L(\Gamma) = \sum_{i=1}^{N_D-1} \sqrt{(1 - \delta_g)j(\Gamma^i, \Gamma^{i+1})^2 + \delta_g \|\nabla I(\Gamma^i)\|^2}$$

where Γ is an arbitrary parametrized discrete path with N_D pixels given by $\{\Gamma^1, \Gamma^2, \dots, \Gamma^{N_D}\}$, $j(\Gamma^i, \Gamma^{i+1})$ is the Euclidean distance between successive pixels, and the quantity $\|\nabla I(\Gamma^i)\|^2$ is a finite difference approximation of the image gradient between the points (Γ^i, Γ^{i+1}) . The parameter weights δ_g the Euclidean distance with the geodesic length. Using the above definition, one can define the geodesic distance as defined in Eq. (5).

An extension of single star-convexity is to use multiple stars to define a more general class of shapes. Introduction of multiple star centers reduces the path lengths and increases the visibility of small parts of objects like small limbs as shown in Fig. 9. Hence Eq. (3) is extended to multiple stars. A label l is star convex with respect to center c_i , if every point $p \in l$ is visible to a star center c_i in set $\mathcal{C} = \{c_1, \dots, c_{N_T}\}$ via l in the image x , where N_T is the number of star centers (Gulshan et al. 2010). This is expressed as an energy cost:

$$E^*(l|x, \mathcal{C}) = \sum_{p \in R} \sum_{q \in \Gamma_{c,p}} E_{p,q}^*(l_p, l_q) \tag{6}$$

In our case all the correct temporal sparse feature correspondences are used as star centers, hence the segmentation will include all the points which are visible to these sparse features via geodesic distances in the region \mathcal{R} , thereby employing the shape constraint. Since the star centers are selected automatically, the method is unsupervised. Comparison of segmentation constraint with geodesic multi-star convexity against no constraints and Euclidean multi-star convexity constraint is shown in Fig. 10. The figure demonstrates the usefulness of the proposed approach with an improvement in segmentation quality on non-rigid complex objects. The energy in the Eq. (2) is minimized as follows:

$$\min_{s.t.} \min_{l \in S^*(\mathcal{C})} E(l, d) \Leftrightarrow \min_{(l,d)} E(l, d) + E^*(l|x, \mathcal{C}) \tag{7}$$

where $S^*(\mathcal{C})$ is the set of all shapes which lie within the geodesic distances with respect to the centers in \mathcal{C} . Optimization of Eq. (7), subject to each pixel p in the region \mathcal{R} being at a geodesic distance $\Gamma_{c,p}$ from the star centers in the set \mathcal{C} , is performed using the α -expansion algorithm for a pixel p by iterating through the set of labels in $\mathcal{L} \times \mathcal{D}$

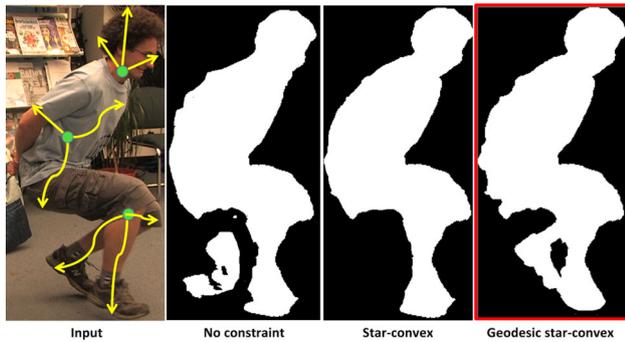


Fig. 10 Segmentation comparison results with no constraint, star convexity constraint and geodesic star convexity constraint for Odzemok dataset

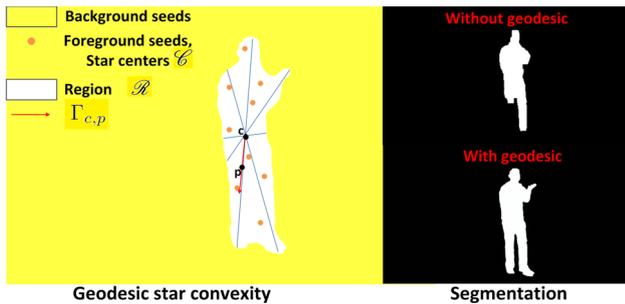


Fig. 11 Geodesic star convexity: a region \mathcal{R} with star centers \mathcal{C} connected with geodesic distance $\Gamma_{c,p}$. Segmentation results with and without geodesic star convexity based optimization are shown on the right for the Juggler dataset

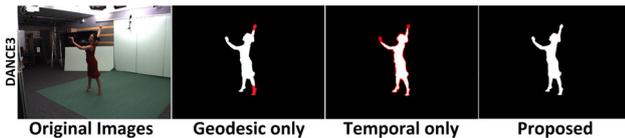


Fig. 12 Comparison of segmentation with introduction of temporal coherence, geodesic star convexity (GSC) and proposed method (GSC and temporal coherence) for Dance2 dataset

(Boykov et al. 2001). Graph-cut is used to obtain a local optimum (Boykov and Kolmogorov 2004).

The improvements in the results using geodesic star convexity in the framework is shown in Fig. 11 and by using temporal coherence is shown in Fig. 8. Figure 12 shows improvements using geodesic shape constraint, temporal coherence and combined proposed approach for Dance2 [2] dataset.

3.3.2 Energy Cost Function for Joint Optimization

For completeness in this section we define each of the terms in Eq. (2), these are based on previous terms used for joint optimisation over depth for each pixel introduced in Mustafa et al. (2015), with modification of the color matching term to improve robustness and extension to multiple labels.

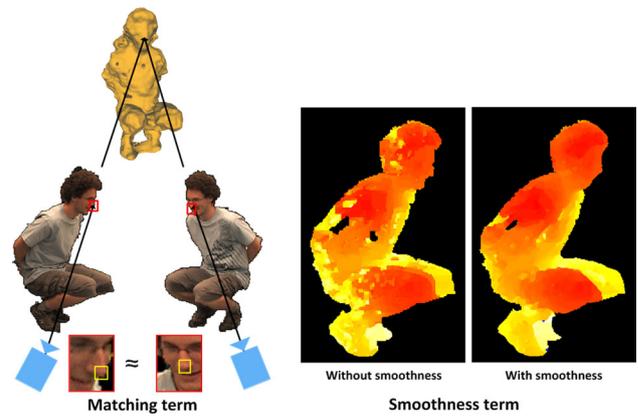


Fig. 13 Illustration of matching and smoothness term for the energy minimization

Matching term The data term for matching between views is specified as a measure of photo-consistency (Fig. 13) as follows:

$$E_{data}(d) = \sum_{p \in \mathcal{P}} e_{data}(p, d_p) = \begin{cases} M(p, q) = \sum_{i \in \mathcal{O}_k} m(p, q), & \text{if } d_p \neq \mathcal{U} \\ M_{\mathcal{U}}, & \text{if } d_p = \mathcal{U} \end{cases} \quad (8)$$

where \mathcal{P} is the 4-connected neighbourhood of pixel p , $M_{\mathcal{U}}$ is the fixed cost of labelling a pixel unknown and q denotes the projection of the hypothesised point P in an auxiliary camera where P is a 3D point along the optical ray passing through pixel p located at a distance d_p from the reference camera. \mathcal{O}_k is the set of k most photo-consistent pairs.

For textured scenes normalized cross correlation (NCC) over a squared window is a common choice (Seitz et al. 2006). The NCC values range from -1 to 1 which are then mapped to non-negative values by using the function $1 - NCC$. A maximum likelihood measure (Matthies 1992) is used in this function for confidence value calculation between the center pixel p and the other pixels q and is based on the survey on confidence measures for stereo (Hu and Mordohai 2012). The measure is defined as:

$$m(p, q) = \frac{\exp \frac{c_{min}}{2\sigma_i^2}}{\sum_{(p,q) \in \mathcal{N}} \exp \frac{-(1-NCC(p,q))}{2\sigma_i^2}} \quad (9)$$

where σ_i^2 is the noise variance for each auxiliary camera i ; this parameter was fixed to 0.3. \mathcal{N} denotes the set of interacting pixels in \mathcal{P} . c_{min} is the minimum cost for a pixel obtained by evaluating the function $(1 - NCC(\cdot, \cdot))$ on a 15×15 window.

Contrast term Segmentation boundaries in images tend to align with contours of high contrast and it is desirable to represent this as a constraint in stereo matching. A consistent

Table 1 Properties of all datasets where ‘Type’ represents whether the data is static or dynamic. In ‘No. of views’ S stands for static cameras and M for moving cameras

Datasets	Resolution	No. of views	Baseline	Type
Office [1]	1920 × 1080	8 (all S)	25°–35°	Dynamic
Juggler (Ballan et al. 2010)	960 × 544	6 (all M)	15°–45°	Dynamic
Dance1 [1]	1920 × 1080	8 (all S)	20°–30°	Dynamic
Odzemok [1]	1920 × 1080	8 (2 M)	15°–30°	Dynamic
Dance2 [2]	780 × 582	8 (all S)	35°–45°	Dynamic
Magician (Ballan et al. 2010)	960 × 544	6 (all M)	15°–45°	Dynamic
Couch (Kowdle et al. 2012)	640 × 480	9 (all S)	25°–30°	Static
Chair (Kowdle et al. 2012)	640 × 480	17 (all S)	5°–8°	Static
Car (Kowdle et al. 2012)	640 × 480	16 (all S)	5°–8°	Static

interpretation of segmentation-prior and contrast-likelihood is used from Kolmogorov et al. (2006). We used a modified version of this interpretation in our formulation to preserve the edges by using Bilateral filtering (Tomasi and Manduchi 1998) instead of Gaussian filtering. The contrast term is as follows:

$$E_{contrast}(l) = \sum_{p,q \in \mathcal{N}} e_{contrast}(p, q, l_p, l_q) \tag{10}$$

$$e_{contrast}(p, q, l_p, l_q) = \begin{cases} 0, & \text{if } (l_p = l_q) \\ \frac{1}{1+\epsilon} (\epsilon + \exp^{-C(p,q)}), & \text{otherwise} \end{cases}$$

$\|\cdot\|$ is the L_2 norm and $\epsilon = 1$. The simplest choice for $C(p, q)$ would be the squared Euclidean color distance between intensities at pixel p and q as used in Guillemaut and Hilton (2010). We propose a term for better segmentation as $C(p, q) = \frac{\|B(p)-B(q)\|^2}{2\sigma_{pq}^2 d_{pq}^2}$ where $B(\cdot)$ represents the bilateral filter, d_{pq} is the Euclidean distance between p and q , and $\sigma_{pq} = \left\langle \frac{\|B(p)-B(p)\|^2}{d_{pq}^2} \right\rangle$. This term enables to remove the regions with low photo-consistency scores and weak edges and thereby helps in estimating the object boundaries.

Smoothness term This term is inspired by (Guillemaut and Hilton 2010) and it ensures the depth labels vary smoothly within the object reducing noise and peaks in the reconstructed surface. This is useful when the photo-consistency score is low and insufficient to assign depth to a pixel (Fig. 13). It is defined as:

$$E_{smooth}(l, d) = \sum_{(p,q) \in \mathcal{N}} e_{smooth}(l_p, d_p, l_q, d_q) = \tag{11}$$

$$\begin{cases} \min(|d_p - d_q|, d_{max}), & \text{if } l_p = l_q \text{ and } d_p, d_q \neq \mathcal{U} \\ 0, & \text{if } l_p = l_q \text{ and } d_p, d_q = \mathcal{U} \\ d_{max}, & \text{otherwise} \end{cases}$$

Table 2 Parameters used for all datasets: λ_c represents $\lambda_{contrast}$

	λ_{data}	λ_c	λ_{smooth}	λ_{color}
Magician/Dance2	0.4	5	0.0005	0.6
Juggler	0.5	5	0.0005	0.4
Odzemok/Dance1/Office	0.4	3	0.001	0.6

d_{max} is set to 50 times the size of the depth sampling step for all datasets.

Color term This term is computed using the negative log likelihood (Boykov and Kolmogorov 2004) of the color models learned from the foreground and background markers. The star centers obtained from the sparse 3D features are foreground markers and for background markers we consider the region outside the projected initial coarse reconstruction for each view. The color models use GMMs with 5 components each for Foreground/Background mixed with uniform color models (Das et al. 2009) as the markers are sparse.

$$E_{color}(l) = \sum_{p \in \mathcal{P}} -\log P(I_p | l_p) \tag{12}$$

where $P(I_p | l_p = l_i)$ denotes the probability at pixel p in the reference image belonging to layer l_i .

4 Results and Performance Evaluation

The proposed system is tested on publicly available multi-view research datasets of indoor and outdoor scenes, details of datasets explained in Table 1. The parameters used for all the datasets are defined in Table 2. More information is available on the website.¹

¹ <http://cvssp.org/projects/4d/4DRecon/>.

Table 3 Static segmentation completeness comparison with existing methods on benchmark datasets (%)

Dataset	Kowdle	Djelouah	Guillemaut	Mustafa	Proposed
Couch	99.6 ± 0.1	99.0 ± 0.2	97.0 ± 0.3	98.5 ± 0.2	99.7 ± 0.3
Chair	99.2 ± 0.4	98.6 ± 0.3	97.9 ± 0.5	98.0 ± 0.5	99.1 ± 0.3
Car	98.0 ± 0.7	97.0 ± 0.8	95.0 ± 0.7	97.6 ± 0.3	98.6 ± 0.4

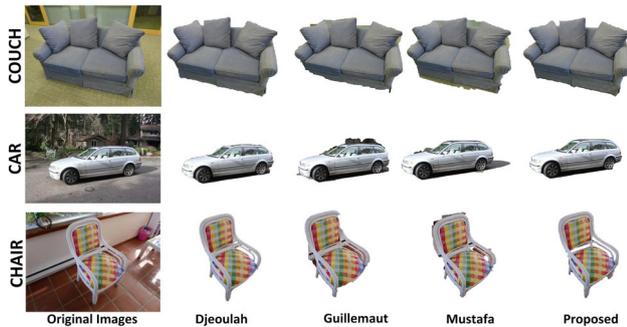


Fig. 14 Comparison of segmentation on benchmark static datasets using geodesic star-convexity

4.1 Multi-View Segmentation Evaluation

Segmentation is evaluated against the state-of-the-art methods for multi-view segmentation Kowdle et al. (2012) and Djelouah et al. (2013) for static scenes and joint segmentation reconstruction methods Mustafa et al. (2015) (per frame) and Guillemaut and Hilton (2012) (using temporal information) for both static and dynamic scenes.

For static multi-view data the segmentation is initialised as detailed in Sect. 3 followed by refinement using the constrained optimisation Sect. 3.3. For dynamic scenes the full pipeline with temporal coherence is used as detailed in Sect. 3. Ground-truth is obtained by manually labelling the foreground for Office, Dance1 and Odzemok dataset, and for other datasets ground-truth is available online. We initialize all approaches by the same proposed initial coarse reconstruction for fair comparison.

To evaluate the segmentation we measure completeness as the ratio of intersection to union with ground-truth (Kowdle et al. 2012). Comparisons are shown in Table 3 and Figs. 14, 15 for static benchmark datasets. Comparison for dynamic scene segmentations are shown in Table 4 and Figs. 16, 17. Results for multi-view segmentation of static scenes are more accurate than Djelouah, Mustafa, and Guillemaut, and comparable to Kowdle with improved segmentation of some detail such as the back of the chair. We also perform ablative analysis on Eq. (2) by removing E_{data} , E_{smooth} and $E_{contrast}$ terms. Results demonstrate that joint depth (E_{data} , E_{smooth}) and segmentation estimation improves the result and contrast information ($E_{contrast}$) helps in improving the quality of the segmentation.

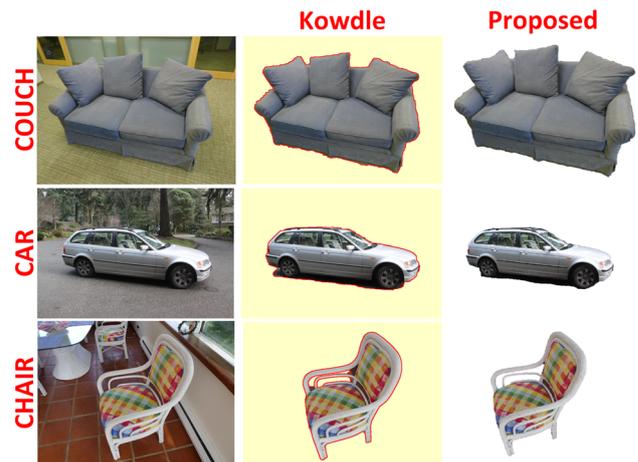


Fig. 15 Comparison of segmentation with Kowdle et al. (2012)

For dynamic scenes the geodesic star convexity based optimization together with temporal consistency gives improved segmentation of fine detail such as the legs of the table in the Office dataset and limbs of the person in the Juggler, Magician and Dance2 datasets in Figs. 16 and 17. This overcomes limitations of previous multi-view per-frame segmentation.

4.2 Reconstruction Evaluation

Reconstruction results obtained using the proposed method are compared against Mustafa et al. (2015), Guillemaut and Hilton (2012), and Furukawa and Ponce (2010) for dynamic sequences. Furukawa and Ponce (2010) is a per-frame multi-view wide-baseline stereo approach which ranks highly on the middlebury benchmark (Seitz et al. 2006) but does not refine the segmentation.

The depth maps obtained using the proposed approach are compared against Mustafa and Guillemaut in Fig. 18. The depth map obtained using the proposed approach are smoother with low reconstruction noise compared to the state-of-the-art methods. Figures 19 and 20 present qualitative and quantitative comparison of our method with the state-of-the-art approaches.

Comparison of reconstructions demonstrates that the proposed method gives consistently more complete and accurate models. The colour maps highlight the quantitative differences in reconstruction. As far as we are aware no ground-truth data exist for dynamic scene reconstruction

Table 4 Dynamic scene segmentation completeness (%): $P_{smooth} = Proposed - E_{smooth}$, $P_{contrast} = Proposed - E_{contrast}$, $P_{ds} = Proposed - E_{data} - E_{smooth}$

	Magician	Juggler	Odzemok	Dance1	Office	Dance2
Guillemaut	68.0	84.6	90.1	99.2	99.3	98.6
Mustafa	88.7	87.9	89.9	99.4	99.0	99.0
$P_{contrast}$	79.6	88.9	90.5	99.1	99.2	98.8
P_{smooth}	89.9	90.6	90.8	99.2	99.2	98.9
P_{ds}	78.3	86.2	90.0	99.3	99.1	98.7
Proposed	91.2	93.3	91.8	99.5	99.4	99.0

from real multi-view video. In Fig. 20 we present a comparison with the reference mesh available with the Dance2 dataset reconstructed using a visual-hull approach. This comparison demonstrates improved reconstruction of fine details with the proposed technique.

In contrast to all previous approaches the proposed method gives temporally coherent 4D model reconstructions with dense surface correspondence over time. The introduction of

temporal coherence constrains the reconstruction in regions which are ambiguous on a particular frame such as the right leg of the juggler in Fig. 19 resulting in more complete shape. Figure 21 shows three complete scene reconstructions with 4D models of multiple objects. The Juggler and Magician sequences are reconstructed from moving handheld cameras.

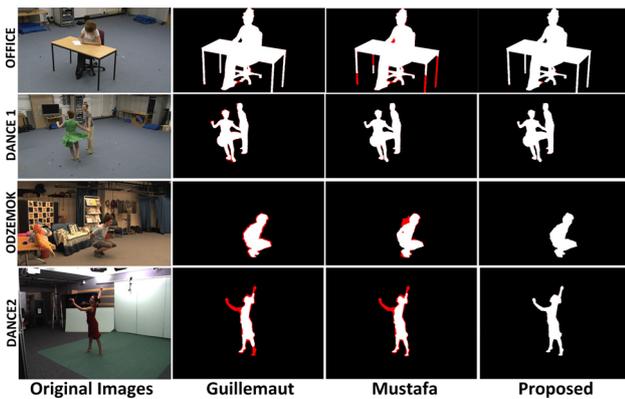


Fig. 16 Segmentation results for dynamic scenes (error against ground-truth is highlighted in red) (Color figure online)

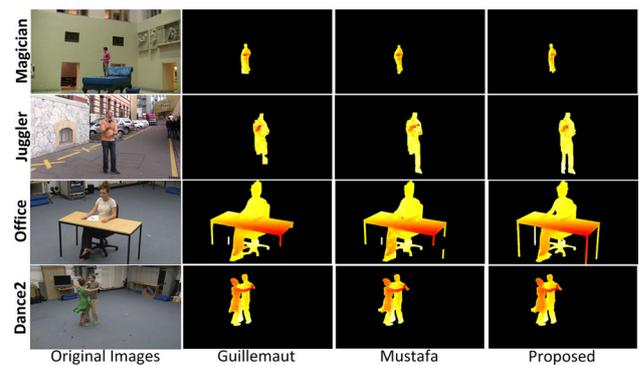


Fig. 18 Comparison of depth maps against existing methods for two indoor and two outdoor benchmark datasets

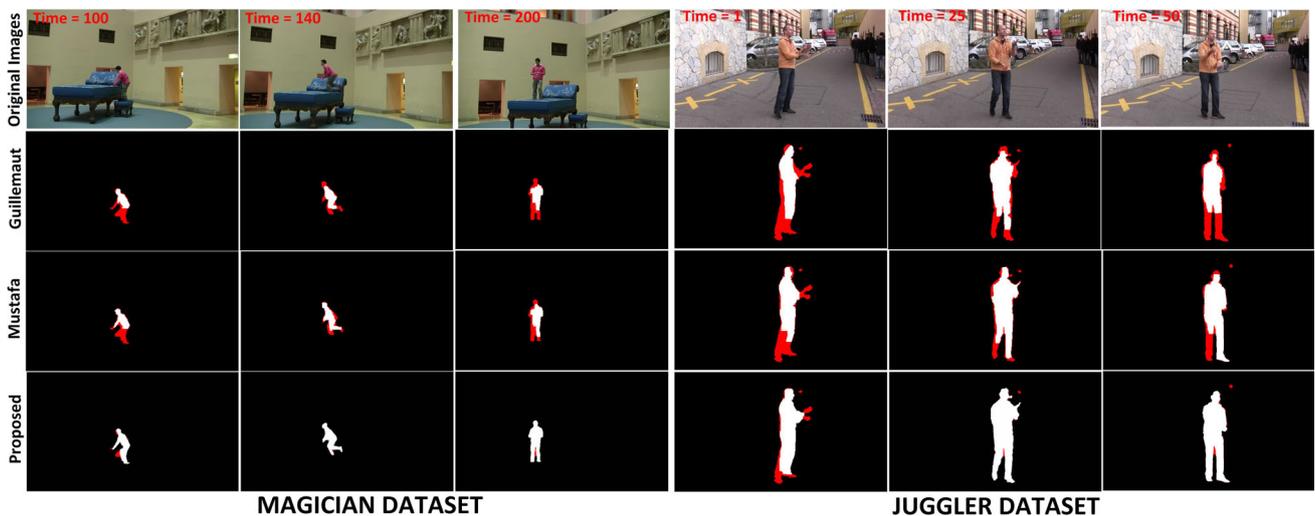


Fig. 17 Segmentation results for dynamic scenes on sequence of frames (error against ground-truth is highlighted in red) (Color figure online)

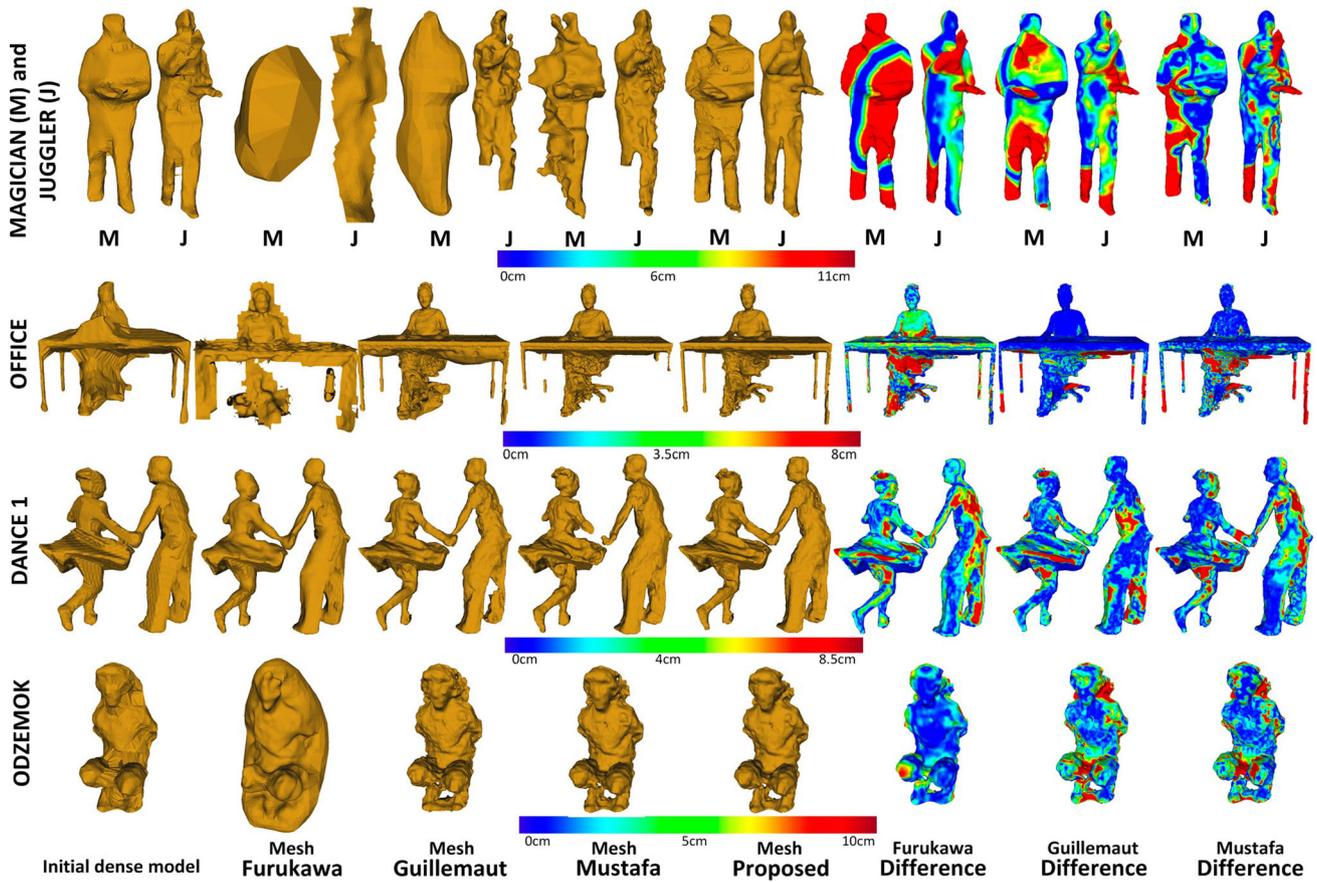


Fig. 19 Reconstruction result mesh comparison against state-of-the-art methods with errors shown in the last three columns

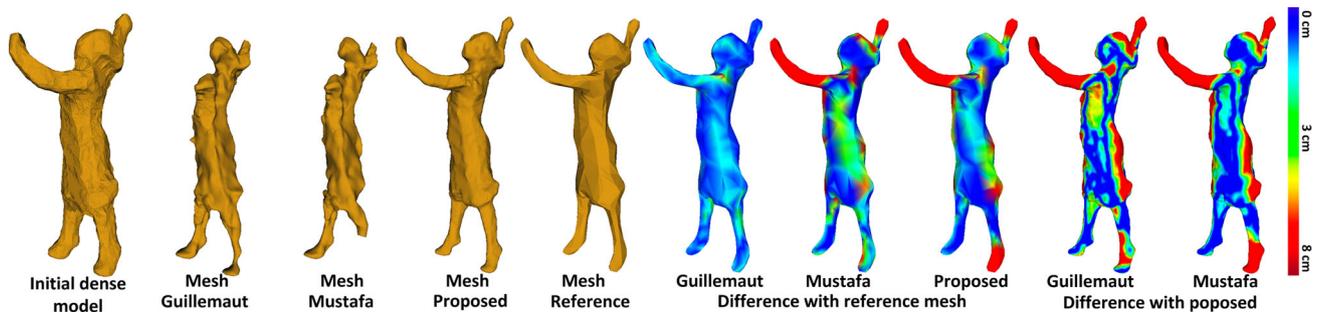


Fig. 20 Reconstruction result comparison with reference mesh and proposed for Dance2 benchmark dataset

Computational Complexity

Computation times for the proposed approach vs other methods are presented in Table 5. The proposed approach to reconstruct temporally coherent 4D models is comparable in computation time to per-frame multiple view reconstruction and gives a ~ 50% reduction in computation cost compared to previous joint segmentation and reconstruction approaches using a known background. This efficiency is achieved through improved per-frame initialisation based on temporal propagation and the introduction of the geodesic

star constraint in joint optimisation. Further results can be found in the supplementary material.

Temporal Coherence

A frame-to-frame alignment is obtained using the proposed approach as shown in Fig. 22 for Dance1 and Juggle dataset. The meshes of the dynamic object in Frame 1 and Frame 9 are color coded in both the datasets and the color is propagated to the next frame using the dense temporal coherence information. The color in different parts of the object is retained to the

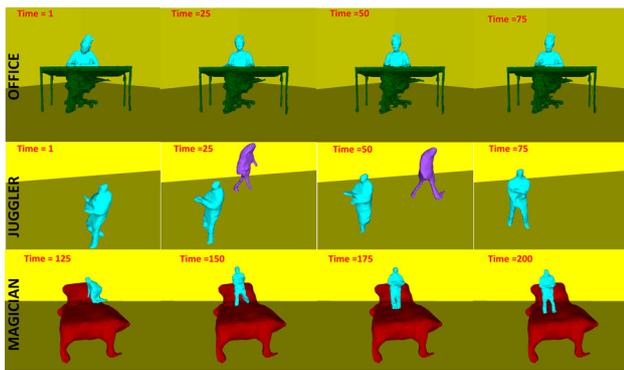


Fig. 21 Complete scene reconstruction with 4D mesh sequence

Table 5 Comparison of computational efficiency for dynamic datasets (time in seconds (s))

Dataset	Furukawa	Guillemaut	Mustafa	Ours
Dance1	326	493	295	254
Magician	311	608	377	325
Odzemok	381	598	394	363
Office	339	533	347	291
Juggler	394	634	411	378
Dance2	312	432	323	278

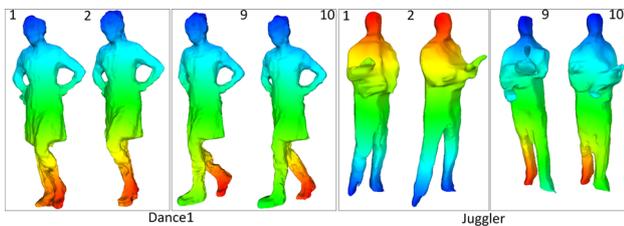


Fig. 22 Frame-to-frame temporal alignment for Dance1 and Juggler dataset

next frame as seen from the figure. The proposed approach obtains sequential temporal alignment which drifts with large movement in the object, hence successive frames are shown in the figure.

Limitations

As with previous dynamic scene reconstruction methods the proposed approach has a number of limitations: persistent ambiguities in appearance between objects will degrade the improvement achieved with temporal coherence; scenes with a large number of inter-occluding dynamic objects will degrade performance; the approach requires sufficient wide-baseline views to cover the scene. Background reconstruction is limited to coarse reconstruction of orthogonal planes based on the Manhattan world assumption.



Fig. 23 Application of proposed method for freeview-point video for Dance2 dataset

5 Applications

The 4D meshes generated from the proposed approach can be used for applications in immersive content production such as FVV rendering and VR. Unlike the previous methods proposed framework does not require strong prior assumptions and manual interactions to create 4D meshes for real-world applications. This section demonstrates the results of these applications.

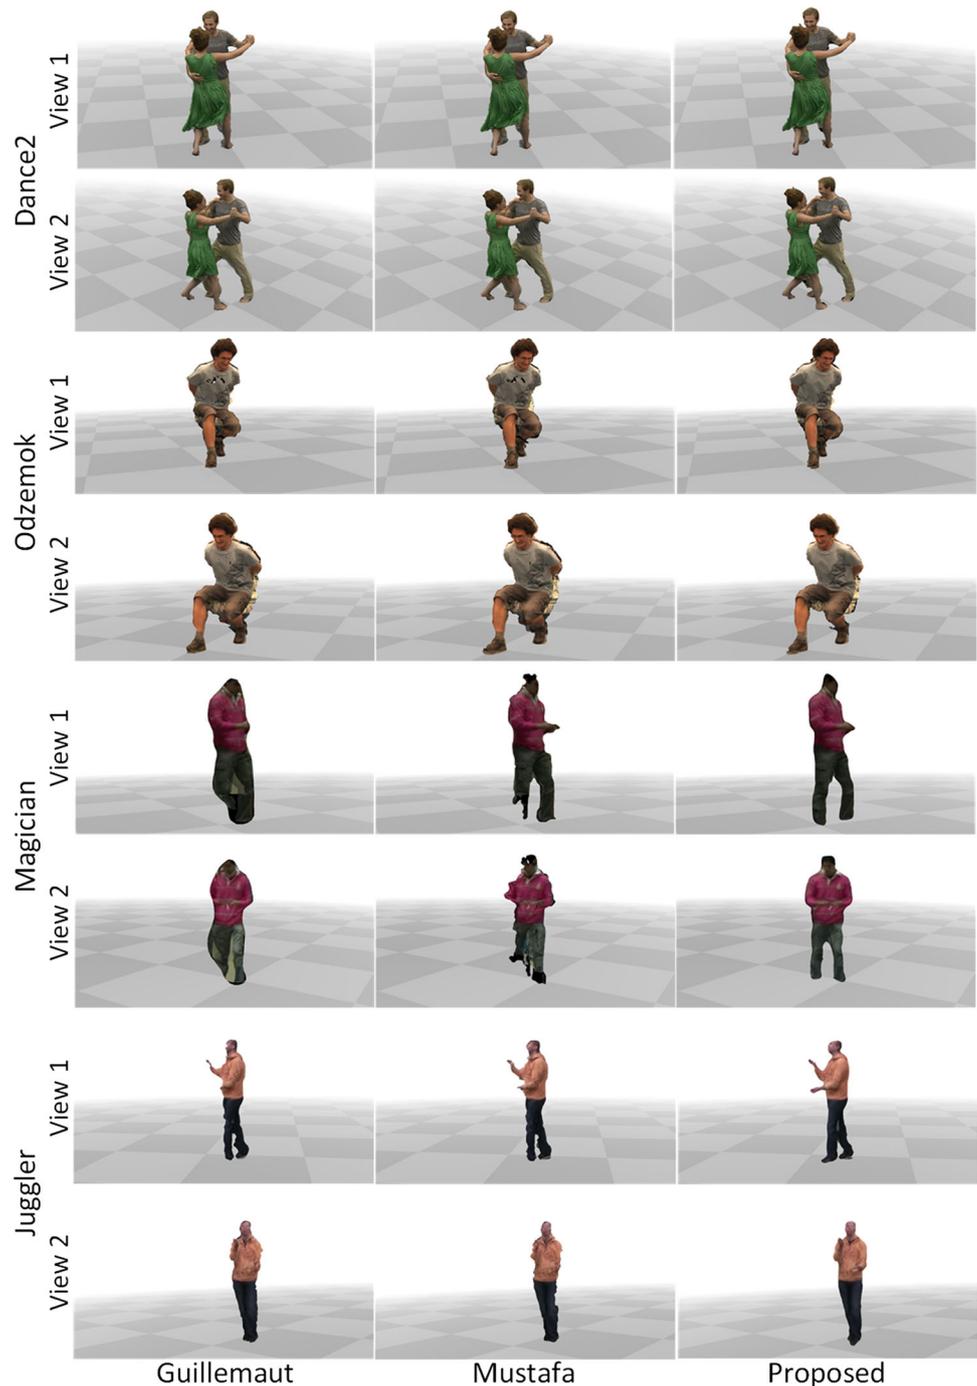
5.1 Free-Viewpoint Rendering

In FVV, the virtual viewpoint is controlled interactively by the user. The appearance of the reconstruction is sampled and interpolated directly from the captured camera images using cameras located close to the virtual viewpoint (Starck et al. 2009).

The proposed joint segmentation and reconstruction framework generates per-view silhouettes and a temporally coherent 4D reconstruction at each time instant of the input video sequence. This representation of the dynamic sequence is used for FVV rendering. To create FVV, a view-dependent surface texture is computed based on the user selected virtual view. This virtual view is obtained by combining the information from camera views in close proximity to the virtual viewpoint (Starck et al. 2009). FVV rendering gives user the freedom to interactively choose a novel viewpoint in space to observe the dynamic scene and reproduces fine scale temporal surface details, such as the movement of hair and clothing wrinkles, that may not be modelled geometrically. An example of a reconstructed scene and the camera configuration is shown in Fig. 23.

A qualitative evaluation of images synthesised using FVV is shown in Figs. 24 and 25. These demonstrate reconstruction results rendered from novel viewpoints from the proposed method against Mustafa et al. (2016a) and Guillemaut and Hilton (2010) on publicly available datasets. This is particularly important for wide-baseline camera configurations where this technique can be used to synthesize intermediate viewpoints where it may not be practical or economical to physically locate real cameras.

Fig. 24 Comparison of free-viewpoint rendering of proposed method against Mustafa and Guillemaut for four datasets



5.2 Virtual Reality Rendering

There is a growing demand for photo-realistic content in the creation of immersive VR experiences. The 4D temporally coherent reconstructions of the dynamic scenes obtained using the proposed approach enables the creation of photo-realistic digital assets that can be incorporated into VR environments using game engines such as Unity and Unreal

Engine, as shown in Fig. 26 for single frame of four datasets and for a series of frames for Dance1 dataset.

In order to efficiently render the reconstructions in a game engine for applications in VR, a UV texture atlas is extracted using the 4D meshes from the proposed approach as a geometric proxy. The UV texture atlas at each frame are applied to the models at render time in unity for viewing in a VR headset. A UV texture atlas is constructed by projectively texturing and blending multiple view frames

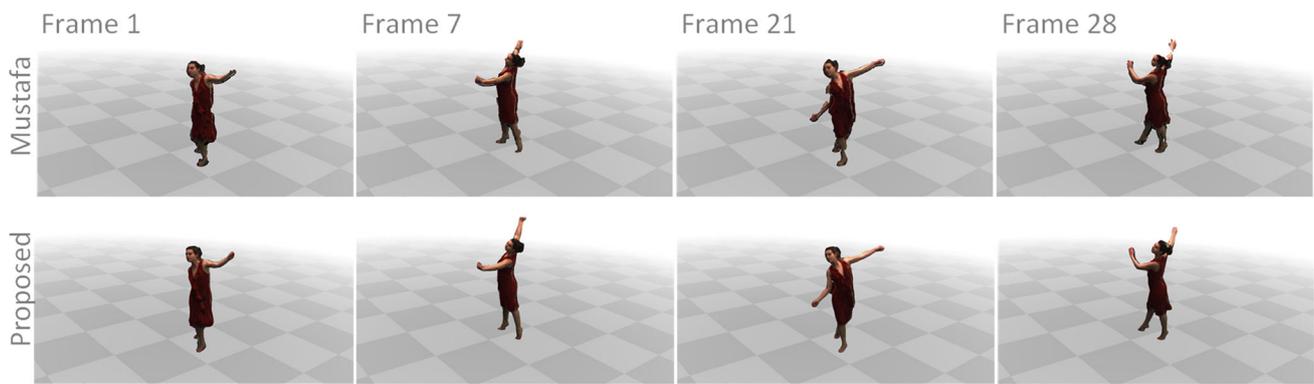


Fig. 25 Comparison of free-viewpoint rendering of proposed method against Mustafa for Dance1 sequence

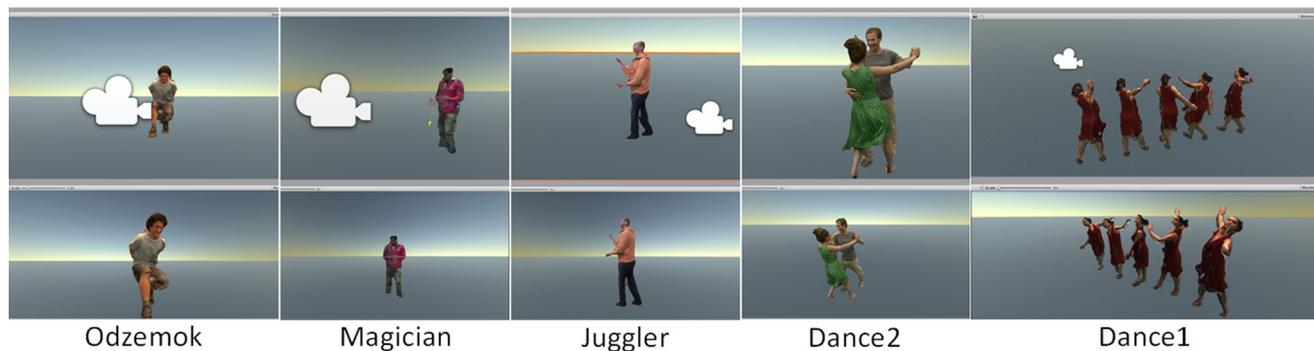


Fig. 26 Application of proposed method for virtual reality. Renderings in unity are shown for five datasets, including a sequence for Dance1



Fig. 27 UV texture atlases for different dynamic datasets to render in VR at a single time instance

onto a 2D unwrapped UV texture atlas, see Fig. 27. This is performed once for each static object and at each time instance for dynamic objects allowing efficient storage and real-time playback of static and dynamic textured reconstructions within a VR headset.

6 Conclusion

This paper introduced a novel technique to automatically segment and reconstruct dynamic scenes captured from multiple moving cameras in general dynamic uncontrolled environments without any prior on background appearance or structure. The proposed automatic initialization

was used to identify and initialize the segmentation and reconstruction of multiple objects. A framework for temporally coherent 4D model reconstruction of dynamic scenes from a set of wide-baseline moving cameras. The approach gives a complete model of all static and dynamic non-rigid objects in the scene. Temporal coherence for dynamic objects addresses limitations of previous per-frame reconstruction giving improved reconstruction and segmentation together with dense temporal surface correspondence for dynamic objects. A sparse-to-dense approach is introduced to establish temporal correspondence for non-rigid objects using robust sparse feature matching to initialise dense optical flow providing an initial segmentation and reconstruction. Joint refinement of object reconstruction and segmenta-

tion is then performed using a multiple view optimisation with a novel geodesic star convexity constraint that gives improved shape estimation and is computationally efficient. Comparison against state-of-the-art techniques for multiple view segmentation and reconstruction demonstrates significant improvement in performance for complex scenes. The approach enables reconstruction of 4D models for complex scenes which has not been demonstrated previously.

Acknowledgements This research was supported by the Royal Academy of Engineering Research Fellowship RF-201718-17177 and the EPSRC Platform Grant on Audio-Visual Media Research EP/P022529.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- (2009). 4d and multiview video repository. In *Centre for vision speech and signal processing*. UK: University of Surrey.
- (2014). 4d repository, <http://4drepository.inrialpes.fr/>. In *Institut national de recherche en informatique et en automatique (INRIA) Rhone Alpes*.
- Atapour-Abarghouei, A., & Breckon, T. P. (2019). Veritatem dies aperit-temporally consistent depth prediction enabled by a multi-task geometric and semantic scene understanding approach. In *CVPR*.
- Bailer, C., Taetz, B., & Stricker, D. (2015). Flow fields: Dense correspondence fields for highly accurate large displacement optical flow estimation. In *ICCV*.
- Ballan, L., Brostow, G. J., Puwein, J., & Pollefeys, M. (2010). Unstructured video-based rendering: Interactive exploration of casually captured videos. In *ACM transactions on graphics* (pp. 1–11).
- Basha, T., Moses, Y., & Kiryati, N. (2010). Multi-view scene flow estimation: A view centered variational approach. In *CVPR* (pp. 1506–1513).
- Bleyer, M., Rhemann, C., & Rother, C. (2011). Patchmatch stereo-stereo matching with slanted support windows. In *BMVC*.
- Bouguet, J. (2000). Pyramidal implementation of the Lucas–Kanade feature tracker. *Microprocessor Research Labs: Intel Corporation*.
- Boykov, Y., & Kolmogorov, V. (2004). An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *PAMI*, 26, 1124–1137.
- Boykov, Y., Veksler, O., & Zabih, R. (2001). Fast approximate energy minimization via graph cuts. *PAMI*, 23, 1222–1239.
- Campbell, N., Vogiatzis, G., Hernández, C., & Cipolla, R. (2010). Automatic 3D object segmentation in multiple views using volumetric graph-cuts. *Image and Vision Computing*, 28, 14–25.
- Chen, P. Y., Liu, A. H., Liu, Y. C., & Wang, Y. C. F. (2019). Towards scene understanding: Unsupervised monocular depth estimation with semantic-aware representation. In *CVPR*.
- Coughlan, J. M., & Yuille, A. L. (2000). The Manhattan world assumption: Regularities in scene statistics which enable Bayesian inference. In *NIPS* (pp. 845–851).
- Das, P., Veksler, O., Zavadsky, V., & Boykov, Y. (2009). Semiautomatic segmentation with compact shape prior. *Image and Vision Computing*, 27, 206–219.
- Dimitrov, D., Knauer, C., Kriegel, K., & Rote, G. (2006). On the bounding boxes obtained by principal component analysis. In *22nd European Workshop on Computational Geometry*.
- Djelouah, A., Franco, J. S., Boyer, E., Le Clerc, F., & Perez, P. (2013). Multi-view object segmentation in space and time. In *ICCV* (pp. 2640–2647).
- Djelouah, A., Franco, J. S., Boyer, E., Le Clerc, F., & Perez, P. (2015). Sparse multi-view consistency for object segmentation. In *PAMI* (p. 1).
- Fortune, S. (1997). Handbook of discrete and computational geometry. In *Chapter Voronoi diagrams and Delaunay triangulations* (pp. 377–388).
- Furukawa, Y., & Ponce, J. (2010). Accurate, dense, and robust multiview stereopsis. *PAMI*, 32, 1362–1376.
- Goldluecke, B., & Magnor, M. (2004). Space–time isosurface evolution for temporally coherent 3D reconstruction. In *CVPR* (pp. 350–355).
- Grundmann, M., Kwatra, V., Han, M., & Essa, I. (2010). Efficient hierarchical graph based video segmentation. In *CVPR*.
- Guan, L., Franco, J. S., & Pollefeys, M. (2010). Multi-view occlusion reasoning for probabilistic silhouette-based dynamic scene reconstruction. *IJCV*, 90, 283–303.
- Guillemaut, J. Y., & Hilton, A. (2010). Joint multi-layer segmentation and reconstruction for free-viewpoint video applications. *IJCV*, 93, 73–100.
- Guillemaut, J. Y., & Hilton, A. (2012). Space–time joint multi-layer segmentation and depth estimation. In *3DIMPVT* (pp. 440–447).
- Gulshan, V., Rother, C., Criminisi, A., Blake, A., & Zisserman, A. (2010). Geodesic star convexity for interactive image segmentation. In *CVPR* (pp. 3129–3136).
- Hane, C., Zach, C., Cohen, A., Angst, R., & Pollefeys, M. (2013). Joint 3D scene reconstruction and class segmentation. In *CVPR* (pp. 97–104).
- Hartley, R., & Zisserman, A. (2003). *Multiple view geometry in computer vision* (2nd ed.). Cambridge: Cambridge University Press.
- Hu, X., & Mordohai, P. (2012). A quantitative evaluation of confidence measures for stereo vision. *PAMI*, 34(11), 2121–2133.
- Huang, Z., Li, T., Chen, W., Zhao, Y., Xing, J., LeGendre, C., Luo, L., Ma, C., & Li, H. (2018). Deep volumetric video from very sparse multi-view performance capture. In *Proceedings of the European Conference on Computer Vision (ECCV)* (pp. 336–354).
- Jiang, H., Liu, H., Tan, P., Zhang, G., & Bao, H. (2012). 3D reconstruction of dynamic scenes with multiple handheld cameras. In *ECCV* (pp. 601–615).
- Kazhdan, M., Bolitho, M., & Hoppe, H. (2006). Poisson surface reconstruction. In *Eurographics symposium on geometry processing* (pp. 61–70).
- Kim, H., Guillemaut, J., Takai, T., Sarim, M., & Hilton, A. (2012). Outdoor dynamic 3-D scene reconstruction. *CSVT*, 22, 1611–1622.
- Kolmogorov, V., Criminisi, A., Blake, A., Cross, G., & Rother, C. (2006). Probabilistic fusion of stereo with color and contrast for bilayer segmentation. *PAMI*, 28, 1480–1492.
- Kowdle, A., Sinha, S., & Szeliski, R. (2012). Multiple view object cosegmentation using appearance and stereo cues. In *ECCV* (pp. 789–803).
- Kundu, A., Li, Y., Dellaert, F., Li, F., & Reh, J. M. (2014). Joint semantic segmentation and 3D reconstruction from monocular video. *ECCV*, 8694, 703–718.

- Larsen, E., Mordohai, P., Pollefeys, M., & Fuchs, H. (2007). Temporally consistent reconstruction from multiple video streams using enhanced belief propagation. In *ICCV* (pp. 1–8).
- Lee, W., Woo, W., & Boyer, E. (2011). Silhouette segmentation in multiple views. *PAMI*, 33(7), 1429–1441.
- Lei, C., Chen, X. D., & Yang, Y. H. (2009). A new multiview spacetime-consistent depth recovery framework for free viewpoint video rendering. In *ICCV* (pp. 1570–1577).
- Matthies, L. (1992). Stereo vision for planetary rovers: Stochastic modeling to near real-time implementation. *IJCV*, 8, 71–91.
- Menze, M., & Geiger, A. (2015). Object scene flow for autonomous vehicles. In *CVPR*.
- Mustafa, A., Kim, H., Guillemaut, J., & Hilton, A. (2015). General dynamic scene reconstruction from wide-baseline views. In *ICCV*.
- Mustafa, A., Kim, H., Guillemaut, J. Y., & Hilton, A. (2016a). Temporally coherent 4D reconstruction of complex dynamic scenes. In *CVPR, Oral*.
- Mustafa, A., Kim, H., & Hilton, A. (2016b). 4D match trees for non-rigid surface alignment. In *ECCV*.
- Mustafa, A., Kim, H., & Hilton, A. (2019). MSFD: Multi-scale segmentation-based feature detection for wide-baseline scene reconstruction. *IEEE Transactions on Image Processing*, 28, 1118–1132.
- Mustafa, A., Kim, H., Imre, E., & Hilton, A. (2015). Segmentation based features for wide-baseline multi-view reconstruction. In *3DV*.
- Narayana, M., Hanson, A., & Learned-Miller, E. (2013). Coherent motion segmentation in moving camera videos using optical flow orientations. In *ICCV* (pp. 1577–1584).
- Ngo, T., Nagahara, H., Nishino, K., Taniguchi, R., & Yagi, Y. (2019). Reflectance and shape estimation with a light field camera under natural illumination. *IJCV*, 127(11–12), 1707–1722.
- Oswald, M., Stöhrer, J., & Cremers, D. (2014). Generalized connectivity constraints for spatio-temporal 3D reconstruction. In *ECCV, 2014* (pp. 32–46).
- Ozden, K., Schindler, K., & Van Gool, L. (2007). Simultaneous segmentation and 3D reconstruction of monocular image sequences. In *ICCV* (pp. 1–8).
- Papazoglou, A., & Ferrari, V. (2013). Fast object segmentation in unconstrained video. In *ICCV* (pp. 1777–1784).
- Qian, Y., Gong, M., & Yang, Y. H. (2017). Stereo-based 3D reconstruction of dynamic fluid surfaces by global optimization. In *CVPR*.
- Rusu, R. B. (2009). Semantic 3D object maps for everyday manipulation in human living environments. Ph.D. thesis, Computer Science Department, Technische Universität München, Germany.
- Seitz, S., Curless, B., Diebel, J., Scharstein, D., & Szeliski, R. (2006). A comparison and evaluation of multi-view stereo reconstruction algorithms. In *CVPR* (pp. 519–528).
- Shin, Y. M., Cho, M., & Lee, K. M. (2013). Multi-object reconstruction from dynamic scenes: An object-centered approach. *CVIU*, 117, 1575–1588.
- Slavcheva, M., Baust, M., Cremers, D., & Ilic, S. (2017). Killingfusion: Non-rigid 3D reconstruction without correspondences. In *CVPR*.
- Starck, J., & Hilton, A. (2007). Surface capture for performance-based animation. *IEEE Computer Graphics and Applications*, 27, 21–31.
- Starck, J., Kilner, J., & Hilton, A. (2009). A free-viewpoint video renderer. *Journal of Graphics, GPU, and Game Tools*, 14(3), 57–72.
- Stutz, D., & Geiger, A. (2018). Learning 3D shape completion under weak supervision. *IJCV*, 128(5), 1162–1181.
- Szeliski, R., & Golland, P. (1998). Stereo matching with transparency and matting. In *ICCV* (pp. 517–524).
- Taneja, A., Ballan, L., & Pollefeys, M. (2011). Modeling dynamic scenes recorded with freely moving cameras. In *ACCV* (pp. 613–626).
- Tomasi, C., & Manduchi, R. (1998). Bilateral filtering for gray and color images. In *ICCV* (pp. 839–846).
- Tung, T., Nobuhara, S., & Matsuyama, T. (2009). Complete multi-view reconstruction of dynamic scenes from probabilistic fusion of narrow and wide baseline stereo. In *ICCV* (pp. 1709–1716).
- Veksler, O. (2008). Star shape prior for graph-cut image segmentation. In *ECCV* (pp. 454–467).
- Vicente, S., Kolmogorov, V., & Rother, C. (2008). Graph cut based image segmentation with connectivity priors. In *CVPR* (pp. 1–8).
- Vo, M., Narasimhan, S. G., & Sheikh, Y. (2016). Spatiotemporal bundle adjustment for dynamic 3D reconstruction. In *CVPR*.
- Wedel, A., Brox, T., Vaudrey, T., Rabe, C., Franke, U., & Cremers, D. (2011). Stereoscopic scene flow computation for 3D motion understanding. *IJCV*, 95, 29–51.
- Wu, C. (2013). Towards linear-time incremental structure from motion. In *3DV* (pp. 127–134).
- Wu, S., Huang, H., Portenier, T., Sela, M., Cohen-Or, D., Kimmel, R., & Zwicker, M. (2018). Specular-to-diffuse translation for multi-view reconstruction. In *ECCV*.
- Zach, C., Cohen, A., & Pollefeys, M. (2013). Joint 3D scene reconstruction and class segmentation. In *CVPR*.
- Zeng, G., & Quan, L. (2004). Silhouette extraction from multiple images of an unknown background. In *ACCV*.
- Zhang, D., Javed, O., & Shah, M. (2013). Video object segmentation through spatially accurate and temporally dense extraction of primary object regions. In *CVPR*.
- Zhang, G., Jia, J., Hua, W., & Bao, H. (2011). Robust bilayer segmentation and motion/depth estimation with a handheld camera. *PAMI*, 33(3), 603–617.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.