Citation for published version:

Li, W, Cosker, D & Brown, M 2016, 'Drift robust non-rigid optical flow enhancement for long sequences', *Journal of Intelligent and Fuzzy Systems*, vol. 31, no. 5, pp. 2583-2595. <https://doi.org/10.3233/JIFS-169099>

DOI:

[10.3233/JIFS-169099](https://doi.org/10.3233/JIFS-169099)

Publication date:

2016

Document Version

Peer reviewed version

[Link to publication](https://doi.org/10.3233/JIFS-169099)

The final publication is available at IOS Press via <https://doi.org/10.3233/JIFS-169099>

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Drift Robust Non-rigid Optical Flow Enhancement for Long Sequences

Wenbin Li^a, Darren Cosker^b and Matthew Brown^c

^a *Department of Computer Science, University College London, WC1E 6BT, UK*

E-mail: w.li@cs.ucl.ac.uk

^b *Centre for the Analysis of Motion, Entertainment Research and Applications (CAMERA),*

University of Bath, BA2 7AZ, UK. E-mail: d.p.cosker@bath.ac.uk

^c *Google, Mountain View, CA 94043, US*

E-mail: m.brown@bath.ac.uk

Abstract. It is hard to densely track a nonrigid object in long term, which is a fundamental research issue in the computer vision community. This task often relies on estimating pairwise correspondences between images over time where the error is accumulated and leads to a *drift*. In this paper, we introduce a novel optimisation framework with an *Anchor Patch* constraint. It is supposed to significantly reduce overall errors given long sequences containing nonrigidly deformable objects. Our framework can be applied to any dense tracking algorithm, e.g. optical flow. We demonstrate the success of our approach by showing significant error reduction on 6 popular optical flow algorithms applied to a range of realworld nonrigid benchmarks. We also provide quantitative analysis of our approach given synthetic occlusions and image noise.

Keywords: Computer Vision, Dense Tracking, Anchor Patch, Optical Flow, Drift, Long Sequence

1. Introduction

Tracking a set of landmark points through multiple images is a fundamental research issue in computer vision. We define tracking in this work as the estimation of corresponding sets of vertices, pixels or landmark points between a reference frame and any other frame in the same image sequence. In the last two decades, optical flow has become a popular approach for tracking through image sequences [7,4,16,17,25] in fields [21,20,12,28]. Compared with feature matching methods e.g. [19], optical flow provides subpixel accuracy and dense correspondence between a pair of images. In this work, we focus in particular on improving tracking in image sequences using optical flow, and our contribution applies to this class of algorithm.

One of the main drawbacks of optical flow is *drift* [6, 15]. Errors accumulated between frames over time result in movement away from the correct tracking trajectory. Between single image pairs, this problem may not be noticeable. However, accumulation when tracking across long sequences can be particularly problematic. Several authors have previously attempted to re-

Fig. 1. A mesh is tracked through a long nonrigid sequence by using our method.

duce optical flow *drift* in tracking. DeCarlo *et al.* [7] introduce contour information on a human face to improve tracking stability, while Borshukov *et al.* [4] employ manual correction. More recently, Bradley *et al.* [5] proposed an optimisation method constrained by additional tracking information from multiview video sequences. Beeler *et al.* [2] then introduced the concept of anchor frames for human face tracking. In this approach, the sequence is decomposed into several clips based on anchor images which are visually similar to a reference frame. Their optimisation method shortens the tracking distance from reference frames to the target frame to help alleviate errors. However,

their approach is domain specific (faces), and assumes that the entire face will return to a neutral expression (the anchor) several times throughout the sequence. In general, it is difficult to label anchor frames on general object sequences with large displacement motion e.g. waving cloth, as there is usually significant deformation between the reference frame and the other frames. In addition, repeated patterns are typically not global as observed in a face (return to a neutral expression). Rather, they occur in smaller local regions at intermittent intervals.

In this work, we focus on tracking long video sequences using optical flow algorithms, and specifically concentrate on reducing *drift*. The general strategy of our approach is to shorten tracking distances for local regions throughout a long sequence. Our proposed framework combines long term feature matching with optical flow estimation. It may be applied to the tracking of general objects with large displacement motion, and results in a significant reduction in *drift*. We first detect *Anchor Frames* for a sequence (Sec. 4). This provides an initial set of start points for tracking the sequence. Our main contribution is extending this approach by proposing the concept of *Anchor Patches* (Sec. 5). These are corresponding points and patches throughout the sequence which are propagated directly from the reference frame. Our framework substantially reduces overall drift on a tracked image sequence, and may be applied to any optical flow algorithm in a straightforward manner. In our evaluation, we apply the proposed optimisation framework on 6 popular optical flow estimation algorithms to illustrate its applicability. We provide analysis of our method using 6 synthetic benchmark sequences (Sec. 7) generated using a method similar to [10], three of which are degraded by adding occlusion, gaussian noise and salt&pepper noise. In addition, we show its applicability on a popular publicly available real world facial sequence with manually annotated ground truth. We show that our proposed optimisation framework significantly improves tracking accuracy and reduces overall drift when compared against the baseline optical flow approaches alone.

This paper is organized as follows: In Sec. 2, an overview of our proposed optimisation framework is outlined. Sec. 3, 4, 5 and 6 give details of the four major steps in our framework. In Sec. 7, we evaluate our approach using 6 optical flow algorithms tested on 6 synthetic benchmark sequences and a real world facial sequence.

2. System Overview

Our proposed optimisation framework reduces overall optical flow drift given long image sequences, and provides additional robustness against other issues such as large displacements and occlusions. The major procedure is shown in Table 1. The aim of our *Anchor Patch optimisation Framework* (APO) is accurately tracking a mesh denoted by $M_R = (V_R, E_R, F_R)$ from a reference frame I_R to every other frame I_i in the sequence. $M_i = (V_i, E_i, F_i)$ denotes the corresponding mesh on frame I_i . In the following sections, the four major steps are discussed in detail.

3. Step One: Computing Optical Flow Fields

The first step is to compute an optical flow field between every frame and its successor over a long video sequence in both forward and backward directions (Fig. 2). In our evaluation, we consider application of our APO framework on a number of dense correspondence optical flow or tracking approaches, e.g. Brox *et al.* [6], Classic+NL [24] and ITV-L1 [26]. Let $\mathbf{w}_{i \rightarrow i+1}$ denote the optical flow field from frame I_i to frame I_{i+1} . Similarly we have $\mathbf{w}'_{i+1 \rightarrow i}$ denoting the optical flow field from frame I_{i+1} to frame I_i in the backward direction. The optical flow field between frame I_i and I_j where $i < j$ (Forward direction), is denoted by $\mathbf{w}_{i \rightarrow j}$ as $\mathbf{w}_{i \rightarrow j} = \sum_{i < j} \mathbf{w}_{i \rightarrow i+1}$. Similarly, the optical flow field between frame I_j and I_i where $i < j$ (Backward direction), is denoted by $\mathbf{w}'_{j \rightarrow i}$ as $\mathbf{w}'_{j \rightarrow i} = \sum_{j > i} \mathbf{w}'_{j \rightarrow j-1}$.

In order to evaluate the optical flow at a specific pixel $\mathbf{x} = (x, y)^T$, an *Error Score* $E(w)$ is proposed here, where $w = (u, v)^T$ is the optical flow vector at pixel $\mathbf{x}$. The pixel $\mathbf{x}$ in frame I_i is matched to pixel $\mathbf{x}' = (x', y')^T$ in frame I_{i+1} where $\mathbf{x}' = \mathbf{x} + w$. The *Error Score* $E(w)$ is calculated as the weighted *Root Mean Square* (RMS) error at a 3×3 pixel area centred on pixel $\mathbf{x}$ and $\mathbf{x}'$.

$$E(w) = \sqrt{\frac{\alpha_1 d(x, y) + \alpha_2 d_c(x, y) + \alpha_3 d_d(x, y)}{\alpha_1 + \alpha_2 + \alpha_3}}$$

$$d_d(x, y) = d(x-1, y-1) + d(x+1, y+1) + d(x-1, y+1) + d(x+1, y-1)$$

$$d_c(x, y) = d(x-1, y) + d(x+1, y) + d(x, y-1) + d(x, y+1)$$

$$d(x, y) = |I_i(x, y) - I_{i+1}(x+u, y+v)|^2 \quad (1)$$

Input: A reference frame, a triangle mesh and an image sequence
Step 1. Computing Optical flow fields (Sec. 3) <ol style="list-style-type: none"> 1.1 Compute optical flow fields in both forward ($\mathbf{w}_{i \rightarrow i+1}$) and backward ($\mathbf{w}'_{i+1 \rightarrow i}$) 1.2 Define the <i>Error Score</i> function
Step 2. Detect anchor frames and propagate the entire mesh to these frames (Sec. 4) <ol style="list-style-type: none"> 2.1 Match SIFT features from the reference to every other frame 2.2 Compute the general <i>Error Score</i> on matchings 2.3 Label the anchor frames from any frame with the low general <i>Error Score</i>
Step 3. Label anchor patches on non-anchor frames (Sec. 5) <ol style="list-style-type: none"> 3.1 Reuse the SIFT feature matching from 2.1 3.2 Propagate patches from the reference using <i>Barycentric Coordinate Mapping</i>
Step 4. Track remaining patches from anchor frames to non-anchor frames (Sec. 6) <ol style="list-style-type: none"> 4.1 Propagate patches from the reference to anchor frames (Sec. 6.1) <ol style="list-style-type: none"> 4.1.1 Compute concatenating optical flow field $\mathbf{w}_{R \rightarrow A}$ 4.1.2 Propagate patches from anchor frames to non-anchor frames using $\mathbf{w}_{R \rightarrow A}$ 4.1.3 Refine the patches using <i>Error Score</i> 4.2 Propagate patches from anchor frames to non-anchor frames (Sec. 6.2) <ol style="list-style-type: none"> 4.2.1 Track the patches from the reference to frame i using $\mathbf{w}_{A \rightarrow i}$ 4.2.2 Track the patches from the <i>Nearest Anchor Patches</i> to frame i 4.2.3 Eliminate the vertex position conflicts between 4.2.1 and 4.2.2
Output: A mesh tracked throughout the entire image sequence

Table 1

The major steps of the *Anchor Patch* optimisation framework.

Fig. 2. **Step One.** The optical flow fields are computed in both forward ($\mathbf{w}_{i \rightarrow i+1}$) and backward ($\mathbf{w}'_{i+1 \rightarrow i}$) directions between every adjacent images pair in the sequence where the first frame is labelled as a reference frame.

Where α_1 , α_2 and α_3 are weights for controlling the contribution of each pixel in the 3×3 area. Using this 3×3 kernel is supposed to give extra robustness against the subpixel accuracy and illumination changes [16,17,14,18]. In our experiments, all these weights are set as $\alpha_1 = 1$, $\alpha_2 = 0.25$ and $\alpha_3 = 0.125$ which refer to the distance from the centre pixel $\mathbf{x}$ of the area. This *Error Score* is intended to evaluate the optical flow at a specific pixel. We also use it to evaluate feature matching scores later in our framework.

4. Step Two: Labeling Anchor Frames

After obtaining our optical flow fields, anchor frames are then detected in a similar manner to Beeler *et al.* [2], with the difference that we employ

SIFT for feature matching as opposed to *Normalised Cross Correlation* (NCC), and additionally use our *Error Score* function (Sec. 3) to evaluate matches. The main procedure is as follows (Fig. 3):

- **Feature Capture.** A set of *SIFT* features S_R is detected in the reference frame I_R . Note that other features could be employed, but we select *SIFT* due to the general high accuracy and robustness. Here we apply the GPU version matching approach [9] to perform correspondence matching of *SIFT* feature sets S_R to feature set S_i of any other frames I_i .
- **Outlier Rejection.** The aim of this selection process is removing outliers from our feature matching on all the frames. Correspondence matches

Fig. 3. **Step Two.** The frames are detected as anchor frames (Red) because of the similar appearance to the reference (Blue). These anchor frames partition the entire sequence into several independent clips which allows tracking performing in parallel.

Fig. 4. The anchor frames are selected based on our general error score which is computed by comparing the reference frame to every other frame in our *Carton* benchmark sequence.

of the *SIFT* feature set S_R between the reference frame I_R and the target frame I_i are performed. We select the matches which meet $|\mathbf{x} - \mathbf{x}'| < \tau$ where $\mathbf{x}$ is feature position in I_R , $\{\mathbf{x} \in S_R, \mathbf{x} = (x, y)^T\}$; $\mathbf{x}'$ is the corresponding feature position in I_i ; τ is a threshold which is set as 30 pixels in our experiments. We find this simple outlier rejection strategy sufficient for most of cases in our experiments (Sec. 7). More sophisticated outlier rejection method such as [22] could also be employed.

- **General Error Score.** The general error score is computed for every image as the average of the overall *Error Score* $E(w)$ (Eq. (1)). Frames that contain the lowest general error score (below a specific threshold) are selected as anchor frames denoted I_A and the other frames are non-anchor frames. It is because that the general error score is supposed to quantise the general appearance deformation where low score presents the small appearance change. Fig. 4 shows this process on our *Carton* benchmark sequence.

After labeling anchor frames that are visually similar to reference frame, these are used as a basis to partition the entire image sequence into several independent *clips*. This also allows computation in the next steps to be performed in parallel. In addition, the mesh M_R is propagated from the reference frame I_R to each anchor frame I_A using *SIFT* matches and a direct optical flow field between them. More detail can be found in Sec. 6.1. The propagated mesh in an anchor frame is denoted $M_A = (V_A, E_A, F_A)$. Because of large dis-

placement motion between anchor frames, and the fact that many images in a deformable sequence may not return to a reference point, these alone are typically insufficient to provide reliable tracking. In the next section, the *Anchor Patch* concept will be introduced to overcome this issue.

5. Step Three: Labeling Anchor Patches

Fig. 6. Anchoring patches using *Barycentric Coordinate Mapping* and *SIFT* features.

The motivation of the original *Anchor Frame* method [2] is to provide multiple *Starting Points* for tracking. Since error accumulates, the technique is intended to reduce overall error accumulation across long image sequences. However, as mentioned in the previous section, large displacement motion and complex motion may yield a fact that most images in a video sequence have significant visual differences from the reference frame.

The main observation in long image tracking is that local spatial patterns throughout a sequence may be repeated - i.e. part of a cloth might return to the same

Fig. 5. **Step Three.** Anchor patches (blue patches) are label on non-anchor frames within every clip using *SIFT* feature matching and *Barycentric Coordinate Mapping* between reference frame and non-anchor frame.

position several times throughout a video. We take advantage of these repeating regions in order to track between shorter segments, and thus alleviate error accumulation. Apart from taking an entire image as anchor information, an *Anchor Patch* is defined as a set of individual vertices or a group of pixels in the non-reference frame (any other frame in the sequence), which are highly correspondent to a specific part of the reference. The benefit of using anchor patches is to provide additional information for correcting accumulated errors when tracking using optical flow. This technique can also reduce the impact of a low-quality anchor frame (i.e. the one is too dissimilar from the reference frame). Before anchoring patches on non-anchor frames, we first obtain a set of high-quality *SIFT* feature matches between the reference frame and non-anchor frames, i.e. those frames are not already labelled as the reference frame, or an existing anchor frame. This process proceeds as follows:

- **Feature Capture.** In order to save the computational time, we reuse the *SIFT* feature sets from *Step Two* (Sec. 4). Here the *SIFT* feature set is denoted as S_R in the reference frame I_R ; S_i presents a feature set of non-anchor frame I_i .
- **Matching Selection.** We also reuse the refined matchings from *Step Two* (Sec. 4). This process generates a matches set $\mathbf{m}_{R \rightarrow i}$ from S_R to S_i .

The set of matches $\mathbf{m}_{R \rightarrow i}$ is used as our initial basis for anchoring patches on non-anchor frames. In order to obtain final anchor patches, *Barycentric Coordinate Mapping* and *Error Refinement* are applied as follows:

5.0.1. Barycentric Coordinate Mapping

We suppose to determine the pixel position in a non-anchor frame which corresponds to the position of a vertex on the reference mesh M_R in I_R . These correspondences provide our baseline for stable tracking throughout the image sequence. Fig. 6 illustrates the process of anchoring patches where $v = (x, y)^T$ denotes a vertex in M_R ; $f_* = (x_*, y_*)^T$, and denotes *SIFT* features in the reference frame I_R . Sim-

ilarly, $f'_* = (x'_*, y'_*)^T$ denotes *SIFT* features in a non-anchor frame I_i . For the non-anchor frame I_i , we have $\{f_k \rightarrow f'_k \in \mathbf{m}_{R \rightarrow i}, k = 1, 2, 3, \dots\}$ which denotes previously obtained corresponding *SIFT* feature matches. We wish to calculate the new vertex position $v' = (x', y')^T$ in the non-anchor frame I_i . We do this by searching for the three nearest *SIFT* features f_* in a small 5×5 search window centred on the vertex of interest v . Next, v' is calculated by solving the *Barycentric Coordinate Mapping* equations as:

$$\begin{bmatrix} f_1 & f_2 & f_3 \\ f'_1 & f'_2 & f'_3 \end{bmatrix} \begin{bmatrix} \beta_1 \\ \beta_2 \\ \beta_3 \end{bmatrix} = \begin{bmatrix} v \\ v' \end{bmatrix} \quad (2)$$

Where β_* are intermediate variables that satisfy $\beta_1 + \beta_2 + \beta_3 = 1$. In practice we found this technique to provide an accurate transformation when applied to small region (5×5 pixel block). However, more sophisticated (although slower) interpolation methods could also be used. The process is performed on every vertex in M_R .

5.0.2. Error Refinement

After *Barycentric Coordinate Mapping*, candidate anchor patches denoted by v'_* are obtained in non-anchor frames I_i . We also have matches $v_* \rightarrow v'_*$, the strength of which can be evaluated using our error equation (1). Using this error, we select final anchor patches in a non-anchor frame I_i using $\{P(v'_*) | E(v_* \rightarrow v'_*) < \eta\}$ where η is a predefined threshold.

6. Step Four: Mesh Propagation

The objective of our optimisation framework is to track a mesh M_R from the reference frame to every other frame in an image sequence. Given tracking information from the previous sections, this process is separated into two steps: first, the mesh M_R is propagated from reference frame to anchor frames (Sec. 4

Fig. 7. **Step Four.** Tracking other patches from the anchor frame and nearest anchor patches within a clip where the blue patches are anchor patches, selected from *Nearest Anchor Patch*.

and 6.1). Second, the propagated mesh M_A is propagated from anchor frames to the non-anchor frames within the clip (Sec. 6.2).

6.1. Propagating from the reference frame to anchor frames

The mesh propagation process from the reference frame to the anchor frame is as follows:

- **Computing the optical flow field.** The optical flow field $\mathbf{w}_{R \rightarrow A}$ directly between the reference frame to the anchor frame is computed by sum up pairwise optical flow fields $\mathbf{w}_{i \rightarrow i+1}$ in between.
- **Matching selection.** We propagate the whole mesh M_R from the reference to the anchor frames. For every vertex in M_R , *high error matches* in anchor frames are eliminated (see Error Refinement).
- **Barycentric Coordinate Mapping.** The positions of those eliminated vertices are recomputed by applying *Barycentric Coordinate Mapping* to low error matches. The operation is shown in Fig. 6.

After this stage, information for every vertex in M_R is established from the reference frame to the anchor frame.

6.2. Propagating from anchor frames to non-anchor frames

The entire image sequence is partitioned into clips which are bound by different anchor frames. The propagation process can be individually performed within these clips in parallel. Within these clips, the anchor patches are supposed to improve overall tracking stability and accuracy. In order to use anchor patches in this process, we define *Nearest Anchor Patch* as follows. For vertex v in M_A , the *Nearest Anchor Patch* of v on frame I_i is the anchor patch $\{v'_{i+k} | v \rightarrow v'_{i+k}\}$ on non-anchor frame I_{i+k} which is nearest to I_i in the

image sequence. Fig. 8 shows an example where frame I_{i+k} is the frame which is nearest to frame I_i in image sequence and contains anchor patch v'_{i+k} matching to v in anchor frame I_A . The main tracking procedure proceeds (Fig. 7) as follows:

- **Mesh propagation.** In order to establish tracking information between anchor frames and non-anchor frame, the mesh M_A is first propagated from anchor frame I_A to non-anchor frames I_i using the previously calculated optical flow field $\mathbf{w}_{A \rightarrow i}$ from *Step One* (Sec. 3).
- **Anchor patches propagation.** The *Nearest Anchor Patch* of each vertex v in M_A is searched through the whole clip then propagated to non-anchor frame I_i using the optical flow field in the forward $\mathbf{w}_{* \rightarrow i}$ or backward $\mathbf{w}'_{i+k \rightarrow i}$ direction.
- **Conflict eliminating.** After propagating the mesh and nearest anchor patches to non-anchor frame I_i , there may be position conflict on some of the propagated vertices. As shown in Fig. 8, $\tilde{v}_i$ and $\tilde{v}'_i$ are not in the same desired position. In order to eliminate the conflict, the position of $\{v_i | v \rightarrow v_i\}$ matching to v can be calculated using the sum of all weighted candidate positions e.g. $\tilde{v}_i$ and $\tilde{v}'_i$ (Eq.3) based on the *Error Score*.

$$v_i = \frac{E(v \rightarrow \tilde{v}'_i)\tilde{v}_i + E(v \rightarrow \tilde{v}_i)\tilde{v}'_i}{E(v \rightarrow \tilde{v}'_i) + E(v \rightarrow \tilde{v}_i)} \quad (3)$$

Due to the fact that the anchor frames divide the overall sequence into smaller clips, this allows the mesh propagation in between to be calculated in parallel. In the next section we perform an evaluation of our framework.

7. Evaluation

We evaluate APO with a range of 6 popular optical flow estimation methods which are publicly avail-

Fig. 8. Vertex conflict can happen when mesh and anchor patches are propagated to target frame I_i . Here v'_{i+k} is an anchor patch that is strongly matched to v .

able from the *Middlebury Evaluation System* [1]. *Combined local-global Optical Flow* (CLG-TV) [8], *Large Displacement Optical Flow* (LDOF) [6] and *Classic+NL* [24] are state of the art while the *Horn and Schunck* (HS) [13], *Black and Anandan* (BA) [3,24], *Improved TV-L1* (ITV-L1) [26] are classic optical flow frameworks and also widely used. CLG-TV is a high speed approach that uses a combination of bilateral filtering and anisotropic regularization and also one of the top three algorithms in the normalized interpolation error test from Middlebury. LDOF is an integration of rich feature descriptors and variational optical flow and one of best current optical flow estimation algorithms for large displacement motion. Classic+NL provides high performance in the Middlebury evaluation by formalizing the median filtering heuristic and Lorentzian penalty as explicit objective functions in an *improved TV-L1* framework. The HS method is a pioneering technique optical flow. BA provides improvements to the HS framework by introducing robust quadratic error formulation. ITV-L1 is a recent and increasingly popular optical flow framework which uses a similar numerical optimisation scheme to Classic+NL. Our choice of a mixture of newer, state of the art methods, with older traditional approaches, is to highlight the fact that irrespective of the approach used, our APO framework provides significantly improved tracking in all cases.

For our evaluation, we compare the optical flow estimation methods previously mentioned – with and without our optimisation framework – on 7 long benchmark sequences with ground truth. Table 2 gives an overview of the benchmark sequences used in our evaluation. In previous work Garg *et al.* released to the community a set of ground truth data for evaluating optical flow algorithms over long sequences. This is as opposed to the Middlebury dataset, which just considers optical flow between pairs of images, and is there-

fore not applicable to our framework. The sequences of Garg *et al.* contains 60 frames and are generated using interpolated dense *Motion Capture* (MOCAP) data from real deformations of a waving flag [27]. As shown in Fig. 9, we use the same MOCAP data to generate a long video sequence and three other degraded sequences, each of which contains 237 frames of size 500×500 pixels. The three degraded sequences are generated in order to test the robustness of our APO framework under different image conditions. They are generated by individually adding synthetic occlusions, gaussian noise and salt & pepper noise with the same parameters described in [10]. In order to increase the diversity of the sequences, we include three other sequences. One is a *Talking Face Video* (Frank) sequence which contains 300 frames with 68 ground truth annotation points per frame. The other two are also synthetic benchmark sequences generated using MOCAP data of Salzmann *et al.* [23] from the carton and serviette deformations. One contains 266 frames of size 1024×768 while the other contains 307 frames of the same image size. In addition, we also consider the effect of the number of SIFT features detected in the frame, and how this affects overall tracking stability of the APO framework. All optical flow algorithms are applied with default parameter settings from their original papers.

Our baseline optical flow based tracking strategy – for each of the above algorithms – is performed as follows: First, the optical flow field is computed (in forward direction) for every pair of adjacent frames in the sequence. We then mark the initial tracking points in the first frame using the same ground truth points in the same frame of the sequence (Table 2). The correspondent points in the next frame are computed based on the optical flow field in between. This process is repeated until correspondent landmark points are obtained in every frame of the sequence. The aver-

	Information of the Benchmark Sequences						
	Original	Occlusion	Guass.N	S&P.N	Carton	Serviette	Frank
Image Size (pix.)	500 × 500	500 × 500	500 × 500	500 × 500	1024 × 768	1024 × 768	720 × 576
Sequence Length	237	237	237	237	266	307	300
Annotation Points	160	160	160	160	81	63	68
Avg. Feature Amount	364.80	358.32	566.13	1276.50	2498.01	3315.49	2071.11

Table 2

An overview of the benchmark sequences in our evaluation. That includes 4 attributes of image size (pixel), sequence length, number of ground truth annotation points per frame and average SIFT feature amount per frame.

Fig. 9. Our long nonrigid synthetic sequences with ground truth are rendered by using flag MOCAP data [27] and Garg et al. [11]. Images are adopted from [10]

age *Endpoint Error* (EE) [1] is then calculated against the ground truth annotation points. We then apply our APO framework using the same optical flow fields.). Note that the parameter values relevant to the APO framework are initially and experimentally selected, but then remain constant in all our evaluations.

Table 3(a) shows the measurement of average *Endpoint Error* (AEE) in pixels over all the frames of the sequences. We highlight the top three best AEE measures for each sequence using superscripts next to different values. Notice that APO significantly reduces the AEE compared to the baseline optical flow methods. Our optimisation framework yields the best AEE measure in all the cases. For instance, *ITV-L1* with APO performs the best in sequence *Original* while *LDOF* with APO yields the best result in sequence *Frank*. We also observe that although in the *Guass.Noise* and *S&P.Noise* sequences the improvement is less than in the unaltered sequences, the overall result is still an improvement with the addition of APO. We also observe that *LDOF* gives good results

even without APO. It is because that the *LDOF* framework takes into account both regular optical flow energy and the feature technique. The latter contributes additional accuracy.

Table 3(b) shows another experiment, in which we performs Garg et al. [10] and Pizarro et al. [22] on our benchmark sequences (results on *Carton*, *Serviette* and *Frank* are not available.) using a direct tracking strategy. Here we compute the optical flow fields directly from the reference to any other frames of the sequence. The annotation points are then directly tracked to the test frames using those flow fields. Note that the numbers in Table 3(b) may be slightly different from their original work [11]. It is because that, first, our sequences are extended to 237 frames which is around 3 times longer; second, we evaluate the tracking results of only 160 annotation points instead of all the pixels. We observe that the both state-of-the-art approaches (Garg et al. and Pizarro et al.) give higher accuracy than any other baseline methods in Table 3(a). The hidden conditions are (1) the tracking distance is minimum for Garg et al. and Pizarro et al. which very much reduces the accumulate errors; (2) both Garg et al. and Pizarro et al. shows high accuracy for nonrigid surface tracking in the record [10,22]. And all our sequences contain single nonrigid object. However, such direct tracking strategy cannot handle the situation where objects may be temporally out of the scene. In addition, the object appearance in the reference may be significantly different from the one in some other frames of the sequence. That brings extra difficulty to optical flow estimation. In this measure (Table 3(b)), our APO degrades to correct the tracking using the long term features. In the *Original* and *Occlusion* cases, the result shows that using APO still yields lower (or the same) errors over the baseline methods even those have given really minuscule measures (< 1.60 pixel). We also observe that APO slightly reduces the accuracy on the noisy cases because the extra noises affect the feature detection. Please note that the results are shown as "N/A" on *Carton*, *Serviette* and *Frank* as the baselines Garg et al. and Pizarro et al. are not available.

(a) Average *Endpoint Error* (AEE) comparison of different methods with our optimisation framework on the benchmark sequences.

Methods	Average <i>Endpoint Error</i> in pix (AEE)						
	Original	Occlusion	Guass.N	S&P.N	Carton	Serviette	Frank
BA [3]	6.14	8.03	11.02	7.79	10.56	5.18	17.57
BA + APO	1.72 ²	1.91 ²	7.89¹	5.04¹	2.77	1.56¹	6.60
CLG-TV [8]	8.59	10.93	20.28	33.93	28.94	32.17	19.29
CLG-TV + APO	2.25	2.97	12.31	18.99	6.95	9.43	7.05
HS [13]	29.16	30.44	29.74	29.43	27.69	37.90	31.27
HS + APO	11.68	12.88	17.79	17.21	10.25	10.03	14.19
LDOF [6]	6.21	6.39	16.24	24.14	6.33	5.51	14.73
LDOF + APO	1.75 ³	1.67¹	11.65	13.12	1.18¹	1.84 ²	3.12¹
Classic+NL [24]	7.07	10.61	12.65	9.50	5.72	6.62	17.32
Classic+NL + APO	2.15	3.18	8.31 ²	6.46 ²	1.34 ²	2.03 ³	3.44 ²
ITV-L1 [26]	5.73	8.25	17.29	14.49	5.34	7.11	17.91
ITV-L1 + APO	1.50¹	2.33 ³	9.53 ³	7.70 ³	1.70 ³	2.36	3.69 ³

(b) Average *Endpoint Error* (AEE) comparison of Garg et al. and Pizarro et al. on the benchmark sequences (directly tracking from the reference to any other frames).

Methods	Average <i>Endpoint Error</i> in pix (AEE)						
	Original	Occlusion	Guass.N	S&P.N	Carton	Serviette	Frank
Garg et al., PCA [10]	0.61	0.71	1.64	1.21	N/A	N/A	N/A
Garg et al., PCA [10] + APO	0.60	0.69¹	1.65	1.23	N/A	N/A	N/A
Garg et al., DCT [10]	0.59¹	0.74	1.86	1.54	N/A	N/A	N/A
Garg et al., DCT [10] + APO	0.59¹	0.73	1.85	1.53	N/A	N/A	N/A
Pizarro et al. [22]	0.71	0.79	0.99	0.98	N/A	N/A	N/A
Pizarro et al. [22] + APO	0.79	0.81	0.99¹	0.98¹	N/A	N/A	N/A

Table 3

Average *Endpoint Error* (AEE) comparison on our long benchmark sequences.

While we consider ourselves primarily with tracking over long sequences, the shorter sequences are considered as well. In Table 4, the AEE measures of various methods are compared on the first 30 frames of our benchmark sequences. We observe similar AEE measures as in the long sequence case (Table 2(a)). The APO framework significantly increases the tracking accuracy – outperforming the baseline tracking methods in all cases even given degradation (e.g. *Gauss.Noise* and *S&P.Noise*). Moreover, the *BA* with APO is also observed to overfit in the noisy sequences while *Classic+NL* with APO yields the best measures in both sequences of *Gauss.Noise* and *S&P.Noise*.

We also evaluate the effect on tracking accuracy by varying the number of selected features. Different numbers (50% and 0%) of features are randomly selected from the initial full detection feature set before performing *Anchor Patch* detection. Information on our total number of features can be found in Table 2, e.g. there are 364.80 features averagely on each frame of the sequence *Original*. Table 5 shows an AEE com-

parison given various numbers of features. We observe that AEE is improved given more features in all cases. Another interesting observation is that our optimisation framework provides lower error against the baseline tracking strategy even given sparse or no features (0% feature). Note that in this case, our APO framework defaults to using an optical flow method with just the *Anchor Frame* approach [2]. Also note – for example by comparing to Table 3 – that this indicates that the APO framework also provides significant tracking improvement over using anchor frames alone.

We also make the visual comparisons on two of our sequences, *Frank* and *Serviette*. The former is real world sequence with ground truth annotation points, while the latter is synthetic sequence overlaid with a ground truth mesh. In Fig. 10, we observe noticeable *drift* problems given the baseline optical flow tracking strategy. Also note that more details can be found in the corresponding video footage where we visually show that our framework significantly reduces the *drift*.

Methods	Average Endpoint Error (AEE) on the First 30 Frames						
	Original	Occlusion	Guass.N	S&P.N	Carton	Serviette	Frank
BA [3]	1.57	1.72	3.87	2.71	2.37	1.56 ³	8.76
BA + APO	1.41 ³	1.65	3.66 ³	2.13 ²	2.17	1.13 ¹	5.40
CLG-TV [8]	2.40	2.60	6.71	8.77	8.10	5.54	8.60
CLG-TV + APO	2.10	2.24	6.53	8.39	4.79	5.11	7.35
HS [13]	33.67	35.70	35.05	34.50	26.16	22.08	12.76
HS + APO	16.11	16.32	13.78	19.37	9.78	6.33	9.19
LDOF [6]	2.38	2.37	3.96	4.03	3.90	2.52	8.51
LDOF + APO	1.15 ²	0.97 ¹	3.75	2.66	0.89 ¹	1.44 ²	2.82 ¹
Classic+NL [24]	1.63	1.76	3.61 ²	2.51 ³	2.18	1.75	8.77
Classic+NL + APO	1.51	1.33 ²	3.54 ¹	1.99 ¹	1.24 ²	1.68	3.70 ³
ITV-L1 [26]	1.55	1.76	6.27	5.07	2.37	2.01	9.22
ITV-L1 + APO	0.99 ¹	1.31 ²	5.77	4.65	1.69 ³	1.71	3.48 ²

Table 4

Average Endpoint Error (AEE) comparison of different methods with our optimisation framework on the first 30 frames of the benchmark sequences.

Methods	Average Endpoint Error (AEE) on Different Feature Distributions						
	Original	Occlusion	Guass.N	S&P.N	Carton	Serviette	Frank
BA [3], No <i>APO</i>	6.14	8.03	11.02	7.79	10.56	5.18	17.57
<i>APO</i> , 100% Feature	1.72 ²	1.91 ²	7.89 ¹	5.04 ¹	2.77	1.56 ¹	6.60
<i>APO</i> , 50% Feature	3.64	4.71	8.06	6.12	5.89	2.98	10.63
<i>APO</i> , 0% Feature	5.12	6.44	9.23	7.21	8.69	4.35	12.69
CLG-TV [8], No <i>APO</i>	8.59	10.93	20.28	33.93	28.94	32.17	19.29
<i>APO</i> , 100% Feature	2.25	2.97	12.31	18.99	6.95	9.43	7.05
<i>APO</i> , 50% Feature	4.86	6.51	14.39	22.72	15.36	19.91	12.00
<i>APO</i> , 0% Feature	6.94	9.11	16.83	26.03	23.57	24.03	15.07
HS [13], No <i>APO</i>	29.16	30.44	29.74	29.43	27.69	37.90	31.27
<i>APO</i> , 100% Feature	11.68	12.88	17.79	17.21	10.25	10.03	14.19
<i>APO</i> , 50% Feature	18.13	20.28	20.66	19.91	17.39	25.99	23.45
<i>APO</i> , 0% Feature	24.73	27.11	23.97	23.40	24.09	33.11	29.17
LDOF [6], No <i>APO</i>	6.21	6.39	16.24	24.14	6.33	5.51	14.73
<i>APO</i> , 100% Feature	1.75 ³	1.67 ¹	11.65	13.12	1.18 ¹	1.84 ²	3.12 ¹
<i>APO</i> , 50% Feature	3.21	3.09	12.18	15.02	2.90	3.74	8.66
<i>APO</i> , 0% Feature	5.08	5.24	14.11	18.46	5.45	4.89	11.76
Classic+NL [24], No <i>APO</i>	7.07	10.61	12.65	9.50	5.72	6.62	17.32
<i>APO</i> , 100% Feature	2.15	3.18	8.31 ²	6.46 ²	1.34 ²	2.03 ³	3.44 ²
<i>APO</i> , 50% Feature	4.00	6.39	9.48	7.33	3.89	4.00	10.14
<i>APO</i> , 0% Feature	5.96	7.78	11.64	8.98	4.78	6.00	13.27
ITV-L1 [26], No <i>APO</i>	5.73	8.25	17.29	14.49	5.34	7.11	17.91
<i>APO</i> , 100% Feature	1.50 ¹	2.33 ³	9.53 ³	7.70 ³	1.70 ³	2.36	3.69 ³
<i>APO</i> , 50% Feature	3.59	5.17	10.93	8.47	3.41	5.00	10.11
<i>APO</i> , 0% Feature	4.77	6.92	12.50	10.31	4.43	5.95	14.29

Table 5

Average Endpoint Error (AEE) comparison on the benchmark sequences with varying feature distributions.

The computational consumption of our framework heavily relies on the supplementary optical flow method, because we need to calculate the optical flow fields

twice (forward and backward) for every pair of adjacent images. Apart from this, our framework can be implemented in a parallel computation fashion. An-

chor frames divide the sequence into clips which give multiple start points for tracking. In the implementation, a GPU version of SIFT approach [9] is applied for feature detection and matching (around 10 frames per second on our benchmarks). The whole framework is constructed under CUDA platform. Assuming all optical flow fields are obtained, our framework reach real-time efficiency (around 2 frames per second) on our benchmarks using on a 2.9Ghz Xeon 8-cores, NVIDIA Quadro FX 580, 16Gb memory computer.

8. Conclusion

In this paper, we have presented a novel optimisation framework using *Anchor Patches* constraint, which improves the tracking on mesh or sparse points through long image sequences. Our optimisation framework temporally anchors the image regions throughout the sequence in order to mitigate the effect of *Error Accumulation (Drift)*. In the evaluation, our approach combined with 6 popular optical flow algorithms and show significant improvement against baselines methods on 7 benchmark sequences. Such datasets include 6 synthetic benchmark sequences with realworld deformation and 1 realworld sequence.

9. Acknowledgements

We thank Ravi Garg and Lourdes Agapito for providing their GT datasets. We also thank Gabriel Brostow and the UCL Vision Group for their generous comments. The authors are supported by the EPSRC CDE EP/L016540/1 and CAMERA EP/M023281/1; and EPSRC projects EP/K023578/1 and EP/K02339X/1.

References

- [1] Simon Baker, Daniel Scharstein, J. Lewis, Stefan Roth, Michael Black, and Richard Szeliski. A database and evaluation methodology for optical flow. *International Journal of Computer Vision (IJCV'11)*, 92:1–31, 2011. ISSN 0920-5691.
- [2] T. Beeler, F. Hahn, D. Bradley, B. Bickel, P. A. Beardsley, C. Gotsman, R. W. Sumner, and M. H. Gross. High-quality passive facial performance capture using anchor frames. *ACM Transactions on Graphics (TOG'11)*, 30(4):75, 2011.
- [3] M.J. Black and P. Anandan. The robust estimation of multiple motions: Parametric and piecewise-smooth flow fields. *Computer vision and image understanding (CVIU'96)*, 63(1):75–104, 1996.
- [4] G. Borshukov, D. Pionni, O. Larsen, JP Lewis, and C. Tempelaar-Lietz. Universal capture: image-based facial animation for the matrix reloaded. In *ACM SIGGRAPH'05 Courses*, page 16. ACM, 2005.
- [5] D. Bradley, W. Heidrich, T. Popa, and A. Sheffer. High resolution passive facial performance capture. *ACM Transactions on Graphics (TOG'10)*, 29(4):41, 2010.
- [6] T. Brox and J. Malik. Large displacement optical flow: Descriptor matching in variational motion estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI'11)*, 33:500–513, 2011. ISSN 0162-8828.
- [7] D. DeCarlo and D. Metaxas. The integration of optical flow and deformable models with applications to human face shape and motion estimation. In *Computer Vision and Pattern Recognition (CVPR'96)*, pages 231–238, 1996.
- [8] M. Drulea and S. Nedevschi. Total variation regularization of local-global optical flow. In *Intelligent Transportation Systems (ITSC'11)*, pages 318–323. IEEE, 2011.
- [9] Brian Fulkerson and Stefano Soatto. Really quick shift: Image segmentation on a gpu. In *Trends and Topics in Computer Vision*, pages 350–358. Springer, 2012.
- [10] R. Garg, A. Roussos, and L. Agapito. A variational approach to video registration with subspace constraints. *International journal of computer vision (IJCV'13)*, 104(3):286–314, 2013.
- [11] Ravi Garg, Anastasios Roussos, and Lourdes Agapito. Dense variational reconstruction of non-rigid surfaces from monocular video. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR'13)*, pages 1272–1279, 2013.
- [12] Clement Godard, Peter Hedman, Wenbin Li, and Gabriel J Brostow. Multi-view reconstruction of highly specular surfaces in uncontrolled environments. In *International Conference on 3D Vision (3DV'15)*, pages 19–27. IEEE, 2015.
- [13] B.K.P. Horn and B.G. Schunck. Determining optical flow. *Artificial intelligence*, 17(1-3):185–203, 1981.
- [14] Wenbin Li. *Nonrigid Surface Tracking, Analysis and Evaluation*. PhD thesis, University of Bath, 2013.
- [15] Wenbin Li, Darren Cosker, and Matthew Brown. An anchor patch based optimisation framework for reducing optical flow drift in long image sequences. In *Asian Conference on Computer Vision (ACCV'12)*, pages 112–125, November 2012.
- [16] Wenbin Li, Darren Cosker, Matthew Brown, and Rui Tang. Optical flow estimation using laplacian mesh energy. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR'13)*, pages 2435–2442. IEEE, June 2013.
- [17] Wenbin Li, Yang Chen, JeeHang Lee, Gang Ren, and Darren Cosker. Robust optical flow estimation for continuous blurred scenes using rgb-motion imaging and directional filtering. In *IEEE Winter Conference on Application of Computer Vision (WACV'14)*, pages 792–799. IEEE, March 2014.
- [18] Wenbin Li, Yang Chen, JeeHang Lee, Gang Ren, and Darren Cosker. Blur robust optical flow using motion channel. *Neurocomputing*, 0(0):12, 2016.
- [19] D.G. Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision (IJCV'04)*, 60(2):91–110, 2004.
- [20] Zhihan Lv, Alex Tek, Franck Da Silva, Charly Empereur-Mot, Matthieu Chavent, and Marc Baaden. Game on, science-how video game technology may help biologists tackle visualization challenges. *PLoS one*, 8(3), 2013.
- [21] Zhihan Lv, Alaa Halawani, Shengzhong Feng, Haibo Li, and Shafiq Ur Rehman. Multimodal hand and foot gesture interaction for handheld devices. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 11(1):10, 2014.
- [22] Daniel Pizarro and Adrien Bartoli. Feature-based deformable surface detection with self-occlusion reasoning. *International Journal of Computer Vision (IJCV'12)*, 97:54–70, 2012. ISSN 0920-5691.

(a) Visual comparison of different methods on the frame 88 of the sequence *Frank*. (b) Visual comparison of different methods on the frame 192 of the sequence *Serviette*.Fig. 10. Visual comparison and AEE measures on sequences of *Frank* and *Serviette*.

- [23] M. Salzmann, R. Hartley, and P. Fua. Convex optimization for deformable surface 3-d tracking. In *International Conference on Computer Vision (ICCV'07)*, pages 1–8, 2007.
- [24] Deqing Sun, S. Roth, and M.J. Black. Secrets of optical flow estimation and their principles. In *Computer Vision and Pattern Recognition (CVPR'10)*, pages 2432–2439, 2010.
- [25] Rui Tang, Darren Cosker, and Wenbin Li. Global alignment for dynamic 3d morphable model construction. In *Workshop on Vision and Language (V&LW'12)*, 2012.
- [26] A. Wedel, T. Pock, C. Zach, H. Bischof, and D. Cremers. An improved algorithm for tv-l 1 optical flow. In *Statistical and Geometrical Approaches to Visual Motion Analysis*, pages 23–45. Springer, 2009.
- [27] R. White, K. Crane, and D.A. Forsyth. Capturing and animating occluded cloth. *ACM Transactions on Graphics (TOG'07)*, 26(3):34, 2007.
- [28] Jiachen Yang, Yancong Lin, Zhiqun Gao, Zhihan Lv, Wei Wei, and Houbing Song. Quality index for stereoscopic images by separately evaluating adding and subtracting. *PloS one*, 10(12),

2015.