

In Defence of RANSAC for Outlier Rejection in Deformable Registration

Quoc-Huy Tran¹, Tat-Jun Chin¹, Gustavo Carneiro¹,
Michael S. Brown², and David Suter¹

¹ School of Computer Science, The University of Adelaide, Australia
{huy, tjchin, carneiro, dsuter}@cs.adelaide.edu.au

² School of Computing, National University of Singapore, Singapore
brown@comp.nus.edu.sg

Abstract. This paper concerns the robust estimation of non-rigid deformations from feature correspondences. We advance the surprising view that for many realistic physical deformations, the error of the mismatches (outliers) usually dwarfs the effects of the curvature of the manifold on which the correct matches (inliers) lie, to the extent that one can tightly enclose the manifold within the error bounds of a low-dimensional hyperplane for accurate outlier rejection. This justifies a simple RANSAC-driven deformable registration technique that is at least as accurate as other methods based on the optimisation of fully deformable models. We support our ideas with comprehensive experiments on synthetic and real data typical of the deformations examined in the literature.


1 Introduction

The goal of non-rigid registration is to align pixels in two or more images corresponding to an object which can move and deform smoothly, e.g., a beating heart, a waving t-shirt. The task is usually accomplished by estimating the transformation (e.g., a Radial Basic Function - RBF - warp) which maps pixels from one image to another. Representative applications include shape matching, segmentation in medical images, and retexturing of deformable surfaces.


A popular class of methods relies on detecting and matching salient features (keypoints) between the images, which are then used to learn the mapping parameters [1–4]. A critical issue in such *feature-based* methods is the identification and rejection of mismatches which unavoidably arise due to imperfect keypoint detection and matching. If no mismatches exist, estimating the transformation is trivial, e.g., by solving a linear system for a Thin Plate Spline (TPS) warp [5].

Common sense suggests that standard outlier rejection tools like RANSAC [6] are inapplicable, the fundamental obstacle being that the underlying transformation is of *unknown* and *varying* complexity [7, 8], i.e., the size of the minimal subset cannot be determined. It is also widely assumed that many realistic deformations (e.g., bending paper, rippling cloth) are too non-linear to be amenable to simple geometric modelling. Fig. 2(a) depicts such impressions of the data.


This paper advances the surprising view that, in practice, the scale of error of the mismatches are orders of magnitude larger than the effects of the curvature




(a) SIFT correspondences between a template and an input image. True matches are in green while incorrect matches are in red.




(b) Correspondences from (a) plotted using first-3 principal components.



(c) Another view of (b) with the fitted hyperplane shown in its “side view”.



(d) TPS estimated using matches lying within the hyperplane bounds in (c).



(e) TPS estimated using the true correspondences identified manually.

Fig. 1. Feature-based robust deformable registration using RANSAC

of the manifold containing the correct matches. Fig. 1 illustrates what we mean with images showing a sheet of paper bending — this kind of data is typically used in the literature, e.g., see [1, 8, 9]. SIFT [10] is first invoked to yield a set of correspondences $\mathcal{X} = \{\mathbf{x}_i\}_{i=1}^N$, where each $\mathbf{x}_i = [x_i \ y_i \ x'_i \ y'_i]^T \in \mathbb{R}^4$. Projecting the data onto the first-3 principal components reveals that the correct matches (inliers) are actually distributed compactly on a 2D affine hyperplane, *relative* to the gross error of the mismatches; see Figs 1(b) and 1(c). This means that we can robustly fit a hyperplane onto the data to dichotomise the inliers and outliers; Fig 1(d) shows the TPS warp estimated using the matches returned by RANSAC, which models the underlying warp very well. As we show later, this is characteristic of many of the physical deformations tested in the literature.

Our observation motivates the point that, for many types of deformations, a linear hyperplane is adequate to model the “correspondence manifold” for

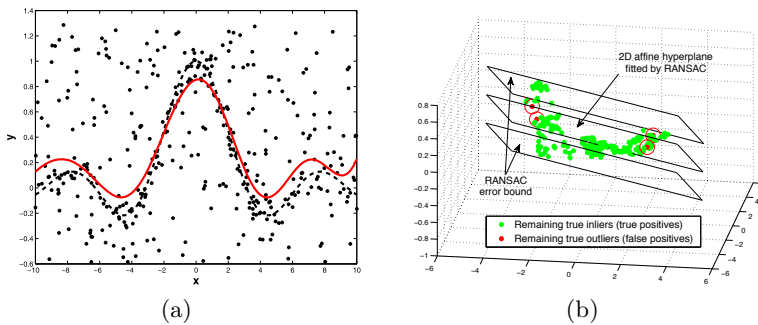


Fig. 2. (a) The characteristic and level of difficulty of the correspondence manifold targeted in [11]. This figure(a) is taken from [11]. (b) Data remaining after RANSAC, shown in the “local” scale of the manifold.

outlier removal. Any outliers remaining (i.e., false positives) are relatively benign rather than outright mismatches, and can usually be smoothed out by the regulariser of the warp estimator; see Fig. 2(b). Observe that the TPS in Fig. 1(d) is very similar to the “ground truth” TPS in Fig. 1(e) estimated using only the true inliers. It is worth noting that without further pixel-based refinement [9], warps estimated from keypoint matches alone cannot extrapolate well to correspondence-poor or occluded regions; see bottom right of Fig. 1(e).

In a sense our observation is not surprising, since PROSAC [12] - a variant of RANSAC - has been used as preprocessing to remove egregious mismatches or to provide affine initialisations for warp estimation [3, 4] (although it was not used in [1, 2], there are few obstacles to initialise with PROSAC/RANSAC there). However, it has always been assumed that due to the complexity of the inlier distribution, significant outliers will remain and it is vital to further optimise the warp robustly, e.g., by an annealing procedure which jointly identifies outliers and learns deformation parameters [3, 4]. Our aim is to show that such procedures overestimate the difficulty of the data, and basic RANSAC followed directly by (non-robust) warp estimation is sufficient.

Close to our work are recently proposed outlier rejection schemes for deformable registration [13, 11, 9]. In [11], SVM regression is used in conjunction with resampling to learn the correspondence manifold in the presence of outliers. In Section 3 of [9], local smoothness constraints are imposed (via Delaunay triangulation) to enable an iterative deformable outlier rejection scheme. These methods assume that substantial non-linearity of the data precludes the usage of RANSAC, which disagrees with our observation typified by Fig. 1(b). Using synthetic and real datasets, we convincingly show that basic RANSAC is at least as accurate as these approaches.

The rest of the paper is organised as follows: Sec. 1.1 surveys related work to put this paper in the context. Sec. 2 explains how RANSAC can be applied for outlier rejection, as well as presents detailed experiments on synthetic and real data. Sec. 3 investigates and compares the performance of our approach on retexturing of deformable surfaces, using publicly available sequences. We conclude and summarise our work in Sec. 4.

1.1 Related Work

Two major paradigms of image-based deformable registration can be distinguished: feature-based methods which rely on keypoint detection and matching [1–4], and pixel-based methods which operate on pixels directly [14, 15]. Feature-based methods are faster but less accurate, and cannot extrapolate well to correspondence-poor areas. However they are crucial for bootstrapping pixel-based methods which are more accurate but slower [9]. Since feature-based methods can only be relied upon to produce “rough” registration, it is desirable to keep this stage of the pipeline as simple and fast as possible. We argue that, on many datasets, bootstrapping based on RANSAC is sufficient.

More recently, methods capable of outlier rejection in feature-based deformable registration have been proposed ([13, 11], Section 3 of [9]). Li et al. [13, 11] proposed that outlier rejection amounts to robustly learning the “correspondence manifold” which, as depicted in Fig. 2(a), is assumed to be highly non-linear and mixed among uniformly distributed outliers. We show that such an assumption is overly pessimistic, since on many datasets the scale of the matching errors is extreme relative to the non-linearity of the manifold.

A parallel area is non-rigid structure from motion (NRSfM), where the aim is to recover the structure of objects that have deformed between views. A number of works assume the deformation to be piecewise rigid [16, 17], which is equivalent to recognising that the distribution of non-rigid data has low degrees of variation. Our work is different in that, towards the goal of outlier rejection for deformable registration, we propose that a *single and global* affine model (instead of a set of rigid or affine models) is sufficient for most correspondence data.

Note that our work is different to non-rigid point cloud or shape alignment, e.g., [18, 19], where the inputs are two sets of unmatched discrete points or landmarks, usually without accompanying image textures. This requires the joint estimation of the transformation and correspondence, whereas our work focusses on rejecting wrongly matched keypoints before non-rigid registration.

2 Outlier Rejection for Deformable Registration

In this section we describe how RANSAC can be applied to outlier rejection in deformable registration, and present experimental results to examine its efficacy.

2.1 The Correspondence Manifold

RBF warps have been applied extensively to model the deformation of various kinds of objects [5]. For deformations of 2D image features, it is common to use two separate RBF warps that share the same centres $\{\mathbf{c}_k\}_{k=1}^K$

$$\begin{bmatrix} x \\ y \end{bmatrix} \mapsto \begin{bmatrix} x' \\ y' \end{bmatrix} \quad \text{where} \quad \begin{aligned} x' &= [x \ y \ 1]^T \mathbf{a}_1 + \mathbf{w}_1^T \mathbf{l}(x, y) \\ y' &= [x \ y \ 1]^T \mathbf{a}_2 + \mathbf{w}_2^T \mathbf{l}(x, y), \end{aligned} \quad (1)$$

$\mathbf{l}(x, y)$ is a non-linear lifting function encapsulating the centres

$$\mathbf{l}(x, y) = [\phi(\|[x \ y]^T - \mathbf{c}_1\|_2) \dots \phi(\|[x \ y]^T - \mathbf{c}_K\|_2)]^T, \quad (2)$$

and $\phi(\cdot)$ is the RBF, e.g., Gaussian or TPS. Given a set of matching features $\mathcal{X} = \{(x_i, y_i), (x'_i, y'_i)\}_{i=1}^N$, the centres are taken as $\{(x_i, y_i)\}_{i=1}^N$. Learning an RBF involves estimating the affine parameters $\mathbf{a}_1, \mathbf{a}_2$ and the coefficients $\mathbf{w}_1, \mathbf{w}_2$ with regularisation to control the warp’s bending energy. For TPS warps this can be achieved by solving a linear system [5].

By regarding each correspondence as a point $\mathbf{x}_i = [x_i \ y_i \ x'_i \ y'_i]^T$ in the joint image space \mathbb{R}^4 , it can be shown that $\mathcal{X} = \{\mathbf{x}_i\}_{i=1}^N$ are samples from a smooth manifold [13]. It is clear that the manifold is two dimensional due to the two degrees of freedom of (x_i, y_i) . Assuming that the underlying warp is an RBF warp, we can express each point on the manifold as

$$\mathbf{x}_i = \underbrace{\begin{bmatrix} 1 & 0 \\ 0 & 1 \\ \mathbf{a}_{11} & \mathbf{a}_{21} \\ \mathbf{a}_{12} & \mathbf{a}_{22} \end{bmatrix} \begin{bmatrix} x_i \\ y_i \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ \mathbf{a}_{13} \\ \mathbf{a}_{23} \end{bmatrix}}_{\text{2D affine subspace}} + \underbrace{\begin{bmatrix} 0 \\ 0 \\ \mathbf{w}_1^T \mathbf{l}(x_i, y_i) \\ \mathbf{w}_2^T \mathbf{l}(x_i, y_i) \end{bmatrix}}_{\text{Non-linear deviation}}, \quad (3)$$

where \mathbf{a}_{pq} is the q -th component of the p -th affine parameter vector. In other words, the correspondence manifold “undulates” around a 2D affine subspace, and the deviation of each \mathbf{x}_i from the subspace is due to the data-dependent non-linear terms $\mathbf{w}_p^T \mathbf{l}(x_i, y_i)$; see Fig. 2(b).

Given a set of matched keypoints \mathcal{X} containing outliers, our premise is that the effects of the matching errors far outweigh the deviation of the true inliers from the affine component of the correspondence manifold. To illustrate this point, Fig. 3 plots the distribution of the orthogonal distances of the data in Fig. 1(a) to the RANSAC-fitted 2D affine hyperplane in Fig. 1(b). It is apparent that a clear separation exists between the inlier and outlier distribution.




Fig. 3. Distribution of distances to RANSAC-fitted 2D affine hyperplane

2.2 Outlier Rejection and Warp Estimation Using RANSAC

Our observations suggest that RANSAC is sufficient for outlier rejection in deformable registration. The goal is to robustly fit a 2D affine subspace onto \mathcal{X} . A minimal solution can be estimated from three data randomly sampled from \mathcal{X} (recall that each datum $\mathbf{x}_i = [x_i \ y_i \ x'_i \ y'_i]^T \in \mathcal{X}$ is a particular correspondence). Let $S = [\mathbf{x}_{s_1} \ \mathbf{x}_{s_2} \ \mathbf{x}_{s_3}] \in \mathbb{R}^{4 \times 3}$ be a random minimal subset with the data

concatenated horizontally. First, the mean of the sample $\boldsymbol{\mu}_S$ is subtracted from each column to yield \hat{S} , whose first-two left singular vectors $\mathbf{A}_{\hat{S}} \in \mathbb{R}^{4 \times 2}$ are then obtained. The pair $(\boldsymbol{\mu}_S, \mathbf{A}_{\hat{S}})$ is sufficient to characterise the affine subspace. The residual (orthogonal distance) of datum \mathbf{x}_i to the fitted subspace is

$$d(\mathbf{x}_i | \boldsymbol{\mu}_S, \mathbf{A}_{\hat{S}}) = \|\mathbf{x}_i - \mathbf{A}_{\hat{S}} \mathbf{A}_{\hat{S}}^T (\mathbf{x}_i - \boldsymbol{\mu}_S) - \boldsymbol{\mu}_S\|_2. \quad (4)$$

RANSAC iteratively generates a set of M 2D affine subspace hypotheses, each fitted on a randomly sampled minimal subset. The consensus of a hypothesis is the number of data with residual less than θ from the associated 2D affine subspace, and the hypothesis with the maximum consensus is returned. The inliers of the best hypothesis are then used to estimate the RBF warp.

A crucial parameter is the threshold θ . Firstly, to allow the usage of a constant θ for all datasets, we normalise the data such that the centroid of $\{(x_i, y_i)\}_{i=1}^N$ lies at the origin, and the mean distance of all points to the original is $\sqrt{2}$. The same normalisation is applied on the points $\{(x'_i, y'_i)\}_{i=1}^N$. The threshold parameter is then manually tuned and used for input images. *Note that an equivalent threshold on the error is required in the other methods* [1–4, 13, 11, 9] (e.g., r in [1, 2], σ in [3, 4], ξ in [11], and d_{TH} in [9]).

A second important parameter is the number of hypotheses M . To ensure with probability p that at least one all-inlier minimal subset is retrieved,

$$M = \frac{\log(1-p)}{\log(1-(1-\epsilon)^3)}, \quad (5)$$

where ϵ is the ratio of outliers among \mathcal{X} . For example, for $p = 0.99$ and $\epsilon = 0.5$, M is approximately 35. In practice the number of iterations used is several times larger than the predicted M . In our experiments we consistently set $M = 100$ for all datasets; as we show later this is still faster than other methods. Moreover, M can be further reduced by using guided sampling methods [12, 20] or the threshold θ can also be estimated automatically [21, 22], though we do not explore these options in our work.

2.3 Experiments on Synthetic Data

We first test the performance of RANSAC on synthetic data. A rectangular mesh is created with control points (RBF centres) distributed on a grid. Using the control points, a TPS warp is randomly generated following the method proposed in [23]. Inliers are produced by randomly sampling 100 positions on the template mesh and mapped using the synthesised TPS warp. The mapped points are then perturbed with Gaussian noise of std. dev. 5 pixels. We then randomly sample positions on the left and right “image” to form outliers. Fig. 4 shows data generated in this manner, with $\epsilon = 0.33$ (33% outliers). Parameter ν in the random warp generator controls the bending energy of the warp (see [23] for details). The effects of different values of ν are shown also in Fig. 4. Observe that for $\nu = 200$ and 500 the mesh is deformed seriously with self-occlusions.




Fig. 4. Top row: Template meshes. Bottom: Meshes warped using randomly generated TPS warps, with bending energy increasing from left to right. Green and red points indicate respectively inliers and outliers (correspondence lines not drawn for clarity).

We benchmark RANSAC against state-of-the-art outlier rejection methods for deformable registration: Iterative local smoothness test [9] (Section 3 of that article) and SVM regression with resampling [11]. We also compare against the class of annealed M-estimation methods [1–4]; since these methods are comparable in accuracy, it is sufficient to compare against [3] which offers the most efficient algorithm. Note that [1–4] can jointly optimise the warp identify outliers; here we concentrate on the aspect of outlier rejection/identification.

The ROC curve of each method is obtained by varying the threshold parameter and recording the resultant true positive rate (number of true inliers recovered over all true inliers) and false positive rate (number of true outliers misidentified as inliers over all true outliers). We set $\nu = 50, 100, 200$ and 500 , and for each ν , the outlier rate ϵ is set as 0.33 and 0.5 . For each combination of ν and ϵ , 100 random (and distinct) TPS warps are generated, and the ROC curves for each method are averaged over the 100 warps. Fig. 5 presents the results.

An apparent and expected trend is that as ν and ϵ increase, the accuracy of all methods decrease, with the method of [11] deteriorating the fastest, followed by [9]. The other two methods provide very comparable accuracies¹. The strength of our method, however, lies in its simplicity and efficiency. Table 1 presents the average running time of all methods for $\epsilon = 0.33$ and 0.5 , where RANSAC is clearly the fastest². The major factors affecting the speed of RANSAC are the outlier rate ϵ and the size of the minimal subset — since only three data are required for a minimal solution, RANSAC can tolerate large ϵ 's without significant sampling effort. On the other hand, the algorithms of [3, 11, 9] are more complicated and the run times scale with the data size.

¹ We were unable to secure the authors' own implementation of the competing algorithms. However the generally good performance of the competing methods implies that our implementation is correct. See code in supplementary material.

² Following a reviewer's comment, we have optimised our implementation of [3]. All methods were implemented and run in MATLAB, which makes the results in Table 1 an accurate picture for *relative* comparisons of run time.

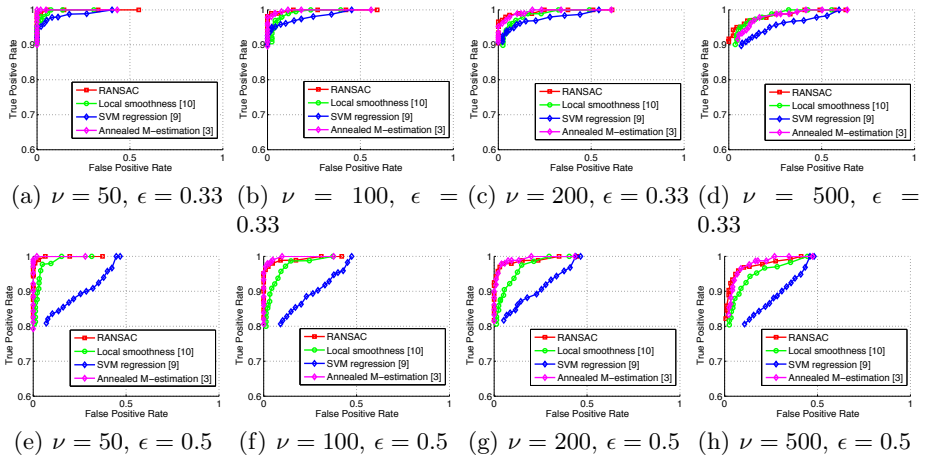


Fig. 5. ROC curves for outlier rejection on synthetic data

Table 1. Average run time (in seconds) for outlier rejection on synthetic data

	$\epsilon = 0.33$ (total 150 matches)	$\epsilon = 0.5$ (total 200 matches)
RANSAC	0.04	0.04
Local smoothness [9]	0.26	0.29
SVM regression [11]	0.06	0.09
Annealed M-estimation [3]	1.41	2.92

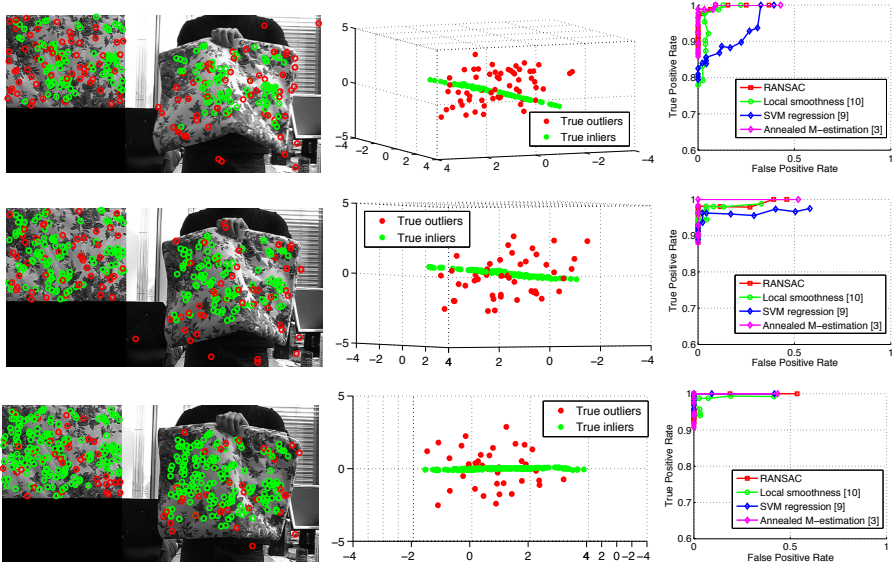
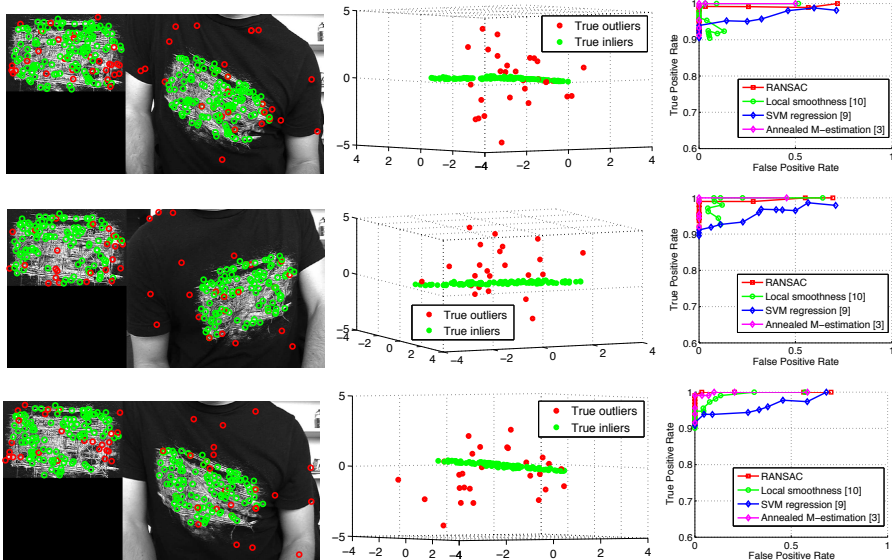


Fig. 6. Results on Frame 140 (145 matches, 41.38% outliers), Frame 160 (152 matches, 31.58% outliers) and Frame 178 (196 matches, 19.90% outliers) from the bedsheet sequence. Col 1: SIFT matches. Col 2: Data after PCA. Col 3: ROC curves.

Table 2. Average run time (in seconds) for outlier rejection on real data

Sequence name	bedsheet			tshirt			cushion		
Frame number	140	160	178	407	720	784	160	175	190
RANSAC	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04
Local smoothness [9]	0.26	0.29	0.28	0.21	0.17	0.19	0.28	0.26	0.22
SVM regression [11]	0.06	0.06	0.08	0.06	0.04	0.05	0.12	0.10	0.06
Annealed M-estimation [3]	1.34	1.52	3.36	1.66	1.40	1.33	3.04	2.30	1.70

**Fig. 7.** Results on Frame 407 (154 matches, 19.48% outliers), Frame 720 (127 matches, 18.90% outliers) and Frame 784 (136 matches, 19.85% outliers) from the tshirt sequence. Col 1: SIFT matches. Col 2: Data after PCA. Col 3: ROC curves.

2.4 Experiments on Real Data

We now test our method on real images. We used publicly available³ image sequences previously used for NRSfM (e.g., see [24]). In this experiment we chose 3 representative frames from the 3 hardest sequences (bedsheet, tshirt and cushion) as input images for outlier rejection. A subimage encapsulating a large portion of the surface was cropped from the first image of each sequence to form the template image. SIFT was invoked to produce keypoint matches, which we then manually categorised as true inliers and outliers. For RANSAC, 100 repetitions were performed on each input image and the average results (ROC curves) are reported. Figs. 6, 7 and 8 illustrate the results.

³ Obtained from <http://cvlab.epfl.ch/data/dsr/>

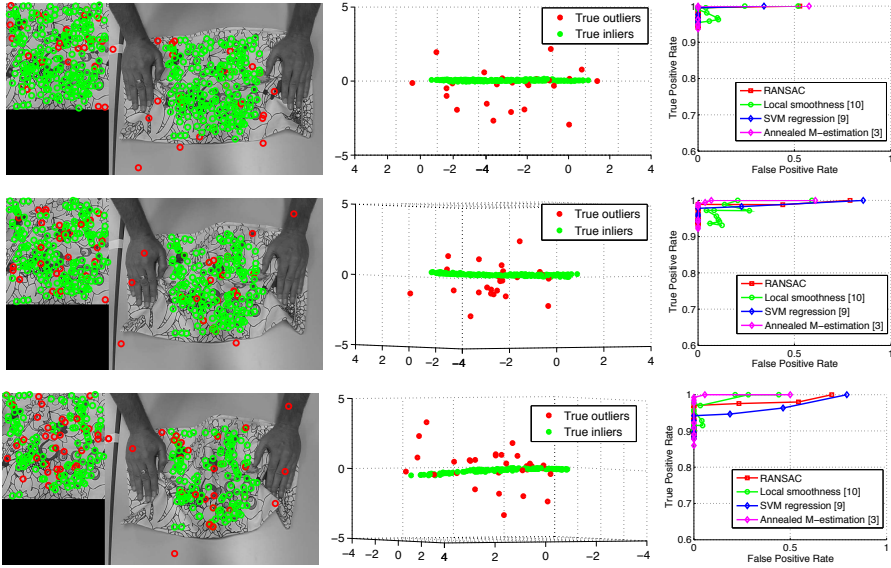


Fig. 8. Results on Frame 160 (234 matches, 10.68% outliers), Frame 175 (205 matches, 13.17% outliers) and Frame 190 (163 matches, 20.25% outliers) from the cushion sequence. Col 1: SIFT matches. Col 2: Data after PCA. Col 3: ROC curves.

Table 3. Number of vertices in warped mesh within 3 pixels away from corresponding vertices in the ground truth mesh


Sequence name	bedsheet			tshirt			cushion		
Frame number	140	160	178	407	720	784	160	175	190
RANSAC	603	728	810	667	660	653	667	666	649
Local smoothness [9]	400	518	675	475	339	217	667	552	645
SVM regression [11]	17	146	702	294	138	220	659	564	473
Annealed M-estimation [3]	648	810	810	667	667	663	667	666	667

The low-dimensional visualisations of all data show that again, relative to the outliers, the inliers are distributed compactly within a 2D affine hyperplane. Based on the ROC curves, a similar conclusion can be made on the accuracy of outlier rejection, i.e., annealed M-estimation [3] and RANSAC are the most accurate, followed by iterative local smoothness test [9] and SVM regression with resampling [11]. The run times of all methods are depicted in Table 2. Again, RANSAC is the fastest method, with constant run times across all images.


The data in which the gap in accuracy between annealed M-estimation [3] and RANSAC is the largest is Frame 190 of cushion (Fig. 8). In the next section we investigate the practical difference due to this disparity in accuracy. *Due to page limits, we provide outlier rejection and warp estimation results on all frames of the sequences (and on other sequences) as supplementary material.*

3 Retexturing Deformable Surfaces


Figs. 9, 10 and 11 provide qualitative comparisons of two best performing outlier rejection methods in Sec. 2.4. The warps for the meshes (for images used in Sec. 2.4) are obtained by first using RANSAC and annealed M-estimation [3] to reject outliers, and then using the remaining matches to estimate a TPS warp. The ground truth warp is obtained by estimating a TPS warp using only true inliers. The threshold value for RANSAC and annealed M-estimation is optimised using the ROC curves in Sec. 2.4. Note that annealed M-estimation can jointly identify outliers and estimate warps, however to yield comparable parameters (a different kind of warp and bending energy are used in [3]) we simply estimate a TPS warp using the inliers returned.



(a) Frame 140, grnd truth (b) Frame 140, RANSAC (c) Frame 140, method [3]




(d) Frame 160, grnd truth (e) Frame 160, RANSAC (f) Frame 160, method [3]




(g) Frame 178, grnd truth (h) Frame 178, RANSAC (i) Frame 178, method [3]

Fig. 9. Retexturing bedsheet images (best viewed on screen)


Both methods yield very close results to the ground truth, including Frame 190 of cushion in which the disparity in outlier rejection accuracy between RANSAC and annealed M-estimation is the largest (see Row 3 of Fig. 8). As mentioned in Sec. 1, false positives produced by RANSAC are normally benign outliers which can be smoothed out by the warp's regulariser.



(a) Frame 407, grnd truth (b) Frame 407, RANSAC (c) Frame 407, method [3]



(d) Frame 720, grnd truth (e) Frame 720, RANSAC (f) Frame 720, method [3]




(g) Frame 784, grnd truth (h) Frame 784, RANSAC (i) Frame 784, method [3]


Fig. 10. Retexturing tshirt images (best viewed on screen)

For quantitative benchmarking, we compute the goodness of each estimated warp as the number of vertices in the warped mesh which are within 3 pixels away from the corresponding vertices in the ground truth mesh. The results in Table 3 show that on several images annealed M-estimation is better than RANSAC in this measure — however, [3] imposes local smoothness constraints which help to “pin down” the position of each vertex relative to the others and this is beneficial for the goodness measure. This additional information is not provided to RANSAC. In any case, as shown in Figs. 9, 10 and 11, the practical differences between the two methods are minuscule.


A general problem for feature-based methods however is the lack of correspondences in certain areas of the surface. To deal with this issue, we track and propagate features in an image sequence. First, the template is divided into rectangular regions (e.g., 5×5 grid). If the number of matches in a region between the current frame and the template falls below a threshold, Mean Shift is initiated to track (pre-matched) features from the previous frame. All matches are then vetted by RANSAC before TPS warp estimation. Note that feature



(a) Frame 160, grnd truth (b) Frame 160, RANSAC (c) Frame 160, method [3]



(d) Frame 175, grnd truth (e) Frame 175, RANSAC (f) Frame 175, method [3]



(g) Frame 190, grnd truth (h) Frame 190, RANSAC (i) Frame 190, method [3]

Fig. 11. Retexturing cushion images (best viewed on screen)

tracking and propagation benefit all feature-based methods [1–4, 13, 11, 9] — *See supplementary material for the results.*

4 Concluding Remarks

We have provided in this paper (and supplementary material) extensive results supporting RANSAC as a viable and simple alternative for outlier rejection compared to more sophisticated approaches. Our premise and observation is that, relative to the extreme scale of gross mismatches, the distribution of inliers usually resembles a low-dimensional affine subspace. While we focus here on RANSAC, there are many approaches to robust fitting of linear manifolds. Some may have advantages over RANSAC and, in that regard, an important message of this paper is that the outlier detection issue with non-linear warping, can likely be done with a relatively cheap schemes.

References

1. Pilet, J., Lepetit, V., Fua, P.: Real-time non-rigid surface detection. In: CVPR (2005)
2. Pilet, J., Lepetit, V., Fua, P.: Fast non-rigid surface detection, registration and realistic augmentation. *IJCV* 76, 109–122 (2008)
3. Zhu, J., Lyu, M.R.: Progressive finite newton approach to real-time nonrigid surface detection. In: CVPR (2007)
4. Zhu, J., Hoi, C.H., Lyu, M.R.: Nonrigid shape recovery by gaussian process registration. In: CVPR (2009)
5. Bookstein, F.L.: Principal warps: thin-plate splines and the decomposition of deformations. *IEEE TPAMI* 11, 567–585 (1989)
6. Fischler, M.A., Bolles, R.C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM* 24 (1981)
7. Carneiro, G., Jepson, A.D.: Flexible spatial configuration of local image features. *IEEE TPAMI* 29, 2089–2104 (2007)
8. Bartoli, A.: Maximizing the predictivity of smooth deformable image warps through cross-validation. *J. Math. Imaging Vis.* 31, 233–244 (2008)
9. Pizarro, D., Bartoli, A.: Feature-based deformable surface detection with self-occlusion reasoning. *IJCV* (to appear)
10. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *IJCV* 60, 91–110 (2004)
11. Li, X., Hu, Z.: Rejecting mismatches by correspondence function. *IJCV* 89, 1–17 (2010)
12. Chum, O., Matas, J.: Matching with prosac - progressive sample consensus. In: CVPR (2005)
13. Li, X., Li, X., Li, H., Cao, M.: Rejecting outliers based on correspondence manifold. *Acta Automatica Sinica* 35, 17–22 (2009)
14. Bartoli, A., Zisserman, A.: Direct estimation of non-rigid registrations. In: *BMVC* (2004)
15. Gay-Bellile, V., Bartoli, A., Sayd, P.: Direct estimation of nonrigid registrations with image-based self-occlusion reasoning. *IEEE TPAMI* 32, 87–104 (2010)
16. Varol, A., Salzmann, M., Tola, E., Fua, P.: Template-free monocular reconstruction of deformable surfaces. In: *ICCV* (2009)
17. Taylor, J., Jepson, A., Kutulakos, K.: Non-rigid structure from locally rigid motion. In: CVPR (2010)
18. Belongie, S., Malik, J., Puzicha, J.: Shape matching and object recognition using shape contexts. *IEEE TPAMI* 24, 509–521 (2002)
19. Chui, H., Rangarajan, A.: A new point matching algorithm for non-rigid registration. *CVIU* 89, 114–141 (2003)
20. Choi, S., Kim, T., Yu, W.: Performance evaluation of RANSAC family. In: *BMVC* (2009)
21. Chen, H., Meer, P.: Robust regression with projection based m-estimators. In: *ICCV* (2003)
22. Rozenfeld, S., Shimshoni, I.: The modified pbm-estimator method and a runtime analysis technique for the ransac family. In: CVPR (2005)
23. Donato, G., Belongie, S.: Approximate Thin Plate Spline Mappings. In: Heyden, A., Sparr, G., Nielsen, M., Johansen, P. (eds.) *ECCV 2002, Part III. LNCS*, vol. 2352, pp. 21–31. Springer, Heidelberg (2002)
24. Salzmann, M., Fua, P.: Reconstructing sharply folding surfaces: a convex formulation. In: CVPR (2009)