Skirting Line Annotation via Deformation Modelling

Daniel Perez Banuelos, Raphael Falque, Timothy Patten and Alen Alempijevic
Robotics Institute, University of Technology Sydney, NSW, Australia
{daniel.a.perezbanuelos}@student.uts.edu.au,
{raphael.guenot-falque, timothy.patten, alen.alempijevic}@uts.edu.au

Abstract

Automating the process of wool handling has the potential to drastically improve the productivity of on-farm operations that would result in significant cost savings for wool growers. Towards this goal, we present a method to automatically extract the skirting line (*i.e.*, the separation between clean and contaminated wool) by comparing pre- and post-skirted RGB images of freshly shorn wool fleece. The intention is to provide annotation to support downstream learning methods. Our approach detects feature correspondences then performs non-rigid outlier rejection to overcome the challenge of deformation when the wool is handled. The final alignment, and hence identification of the skirting line, is achieved through the use of a non-rigid deformation method. A controlled experiment shows, quantitatively, that our approach outperforms a rigid registration baseline. We then demonstrate the applicability to the real use case by presenting qualitative results on images of skirted fleeces collected from a wool shed.

1 Introduction

Wool is highly sought after around the world as it is used in a range of diverse products such as clothing, insulation, carpeting, furniture and packaging. Australia contributes to a significant proportion of the world's wool market. In 2016-17, it was estimated that the value of Australian wool exports was $3.615 billion [DAWR, 2020].

Despite the commercial importance and significance of the wool industry in Australia, it remains largely manual with very little automation. In particular, the on-farm processes, such as shearing, handling and classing, are still currently performed by hand. Wool handling is a highly manual and repetitive task that is particularly amenable to automation. The task of identifying and removing contaminants in freshly shorn fleece, known as *skirting* [Hansford, 2012], is not only similar to other factory-like tasks that have now been automated across many industries including agriculture. Attempts have been made in the scientific community to automate wool inspection, e.g., [Zhang *et al.*, 2005a], however, they deal specifically with non-organic contaminants (such as foreign material).

Our previous work demonstrated that it is possible to detect *natural* contaminants in wool fleece directly from RGB images [Patten *et al.*, 2021]. However, learning to detect contamination requires a large volume of annotated data. Annotation is not only logistically difficult and time-consuming to obtain but is also highly unreliable due to the ambiguity of contamination in wool.

To address this issue, we present a method to automatically determine the delineation of clean and contaminated wool in freshly shorn fleece by analysing images of pre- and post-skirted fleeces. A major challenge is that the wool along the skirting line is significantly deformed during the skirting process. The handler uses one hand to hold down the fleece, acting as an anchor, while the other pulls the section that needs to be removed to achieve a tear along the desired skirting line. Thus, achieving fibre separation along the skirting line simultaneously stretches and compacts fibres on either side of the skirting line. To overcome the challenge, we propose an efficient pruning procedure of feature correspondences that allows for non-rigid deformation and then align the images through an optimisation process.

We quantitatively verify our method in a controlled experiment by attaching markers to pieces of wool and deforming them by hand. The markers serve as ground truth, and we show that our method achieves much better alignment to the deformed target in comparison to the rigid registration baseline. We then qualitatively show the applicability of our work using the Dubbo dataset from [Patten *et al.*, 2021], which consists of pairs of pre- and post-skirted images collected during real operations in a wool shed. Our results show a significantly more reasonable estimate of the transformation between the wool compared to rigid registration. Interestingly, the results further highlight the unreliability of human annotation, which erroneously marks portions of the wool that is to be removed but is still visibly present in the post-skirted images.

2 Related Work

Distinguishing contaminants in wool in RGB images has traditionally used adaptive thresholding [Zhang *et al.*, 2005b; Zhang *et al.*, 2005a; Su *et al.*, 2006]. This is limited as it relies on a significant colour difference between the wool and the contaminant. Unfortunately, this is not always the case and more robust and general methods need to be applied. Deep learning offers these capabilities, having achieved remarkable success in various computer vision tasks [Krizhevsky *et al.*, 2012]. Recently, deep learning has been used to detect contaminants in wool [Patten *et al.*, 2021] and in the similar domain to detect foreign fibres in cotton [Wei *et al.*, 2019; Wei *et al.*, 2020].

The current limitation for deploying deep learning is the requirement of a large volume of annotated data to supervise the learning, which in previous work was hand labelled [Patten *et al.*, 2021]. The motivation of our work is to develop a method to extract the annotation automatically. Given our setup and ability to collect data before and after contamination is removed, the problem translates into one of registering images to identify the missing components. Since wool is highly deformable, these images must be registered through non-rigid alignment.

In the field of medical imaging, similar theoretical problems can be found where scans from various modalities require non-rigid alignment. Classical approaches [Rueckert *et al.*, 1999] solve this problem by optimising a deformation field that maximise the similarity of the different images (in general quantified with mutual information) while enforcing a smoothness regularisation term over the deformation. Most of the literature builds upon the seminal work from Rueckert *et al.* by considering different approaches of the data preprocessing [Pluim *et al.*, 2003] or alternative optimisation methods [Klein *et al.*, 2007]. These methods tightly couple the deformation and the data association, implicitly defined through mutual-information, and are prone to fall into local minima.

More specifically, to non-rigid deformation, the field of computer graphics has extensively studied the problem of animating characters for movies and video games. Most of the work focus on the deformation of triangle meshes given a skeleton animation, which also provides some partial piece-wise rigidity [Kavan *et al.*, 2007]. Related to wool garments, [Sperl *et al.*, 2021] employs deformation optimisation to animate yarn-level cloth in real-time based on the deformation of the mesh. Other work optimises the position of triangles in a mesh purely based on handle-based inputs (*i.e.*, without skeletons). This can be obtained by minimising the distance of vertex handles' movement while simultaneously maximising the rigidity over the surface [Sorkine and Alexa, 2007].

3 Methodology

Given two different images, I_1 and I_2, with their respective masks, M_1 and M_2, our approach computes features for each image and stores them as two different sets ($\mathcal{F}_1$ and $\mathcal{F}_2$). From the putative matches between $\mathcal{F}_1$ and $\mathcal{F}_2$, we remove the outliers to find the relative transform that better aligns the features. This is done with a proposed as-rigid-as-possible (ARAP) filter that accounts for the non-rigid deformation present in the data. Given the filtered feature correspondences, we use a method based on embedded deformation [Sumner *et al.*, 2007] to deform the image (*i.e.*, to obtain an individual transformation which is applied to each pixel). Technical details of these steps are provided in the following sections.

3.1 Feature extraction

Our method uses the Scale-Invariant Feature Transform (SIFT) [Lowe, 2004] to compare the descriptors between the skirted and non-skirted wool. The SIFT method proposed by Lowe helps to solve for the rotation and change in viewpoint present in the data. The SIFT algorithm is composed of four steps. Firstly, differences of Gaussians identifies the locations and scales that are repeated across multiple views. Secondly, a keypoint localisation model is created to determine the location and scale of the data. Thirdly, an orientation is assigned to the keypoint regions based on the image gradients direction. Lastly, the gradient is measured at the selected scale in the area around each keypoint.

The Euclidean distance criterion is then used to match the feature vectors between the two images. The features are then stored as two sets $\mathcal{F}_1$ and $\mathcal{F}_2$ where $\mathcal{F}_i = \{\boldsymbol{f}_1^i, \ldots, \boldsymbol{f}_n^i\}$ with $\boldsymbol{f}_j^i \in \mathbb{R}^2$. $\mathcal{F}_1$ corresponds to the features from the original wool and $\mathcal{F}_2$ corresponds to the features from the skirted wool.

3.2 Non-rigid outliers rejection

As illustrated in Figure 1, the presence of outliers is a common problem when performing feature matching. Typically, this is solved using random sample consensus (RANSAC) [Fischler and Bolles, 1981] to obtain a subset of feature matches that provide a consensus with respect to a rigid transformation. While RANSAC can be reformulated for non-rigid cases [Tran *et al.*, 2012], it embeds

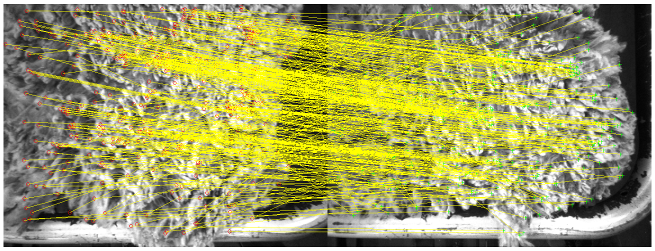

Figure 1: Feature matching based on SIFT descriptors. While most matches are correct, some of the matches are outliers and need to be removed.

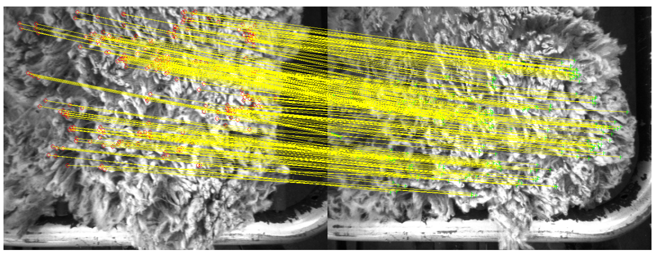

Figure 2: SIFT matches filtered by accounting for non-rigid deformation.

the computation of the non-rigid transformation in an iterative process, which is computationally ineffective. We propose a filter that prunes the feature matching iteratively with a small computational footprint. Our filter is initiated by building graphs on the features sets and performs the pruning based on the local rigidity difference between the two graphs.

The filter is initialised by building a directed graph that connects the feature points from $\mathcal{F}_1$ to its k neighbours. These connections are then concatenated into an edge set $\mathcal{E}_1$ and the first graph is then defined as $\{\mathcal{F}_1, \mathcal{E}_1\}$. The graph on the second set of features $\mathcal{F}_1$ is generated by copying the edges $\mathcal{E}_1$ providing the graph $\{\mathcal{F}_2, \mathcal{E}_1\}$.

For each feature node present in the graph, the local rigidity changes is computed based on the formulation of the rigidity as defined in ARAP [Sorkine and Alexa, 2007]. Given two nodes, $\boldsymbol{f}_i^1$ and $\boldsymbol{f}_i^2$, which correspond to two SIFT matches, the local rigidity for the i^{th} node is defined as

$$\mathbf{r}(i) = \sum_{j \in \mathcal{N}(i)} \left\| (\boldsymbol{f}_i^1 - \boldsymbol{f}_j^1) - \hat{\boldsymbol{R}}_i(\boldsymbol{f}_i^2 - \boldsymbol{f}_j^2) \right\|^2 \frac{1}{\left\| \boldsymbol{f}_i^1 - \boldsymbol{f}_j^1 \right\|^2} \tag{1}$$

where the left-hand part of the equation quantifies the deformation operated on the edge from $(\boldsymbol{f}_i^1 - \boldsymbol{f}_j^1)$ to $(\boldsymbol{f}_i^2 - \boldsymbol{f}_j^2)$ while accounting for a rigid rotation $\hat{\boldsymbol{R}}_i$ (computed using SVD similarly to [Sorkine and Alexa, 2007]). The right-hand part of the equation makes this rigidity term scale invariant.

To remove the outliers in the feature matching, the

Algorithm 1: Prunes the feature matches until the local rigidity falls under the threshold τ.

Input: $\mathcal{F}_1$, $\mathcal{F}_2$, τ
Output: $\hat{\mathcal{F}}_1$, $\hat{\mathcal{F}}_2$

```
function ARAP filter
    build the graph {F_1, E_1}
    copy the edges on {F_2, E_1}
    compute the rigidity terms r(i) with (1)
    while max(r) < τ do
        prune the nodes f_argmax(r) in F_1 and F_2
        update the graph {F_1, E_1}
        copy the edges on {F_2, E_1}
        compute the rigidity terms r(i) with (1)
```

---

feature node with the maximum ARAP residual is pruned from the graph. The graph and the rigidity terms are then updated and the pruning process is repeated iteratively until the ARAP residual falls below a defined threshold. An example of the result after feature matching pruning is shown in Figure 2 and the pseudo-code is given in Algorithm 1.

3.3 Non-rigid deformation

We propose the use of a non-rigid deformation method based on embedded deformation (ED) [Sumner *et al.*, 2007]. Given that this approach was originally designed for three-dimensional data, this provides a natural extension of our work on RGB-D data. Furthermore, ED is highly flexible with respect to the format of the data to deform.

We consider the input of the non-rigid deformation as a set of scattered points with coordinates defined in three dimensions (in cases where a standard 2D image is used as an alternative to an RGB-D input, the third dimension is stacked with a row of zeros). ED performs the non-rigid optimisation by solving the deformation on a graph, referred to as a deformation graph. The nodes of this graph, $\mathcal{G} = \{\boldsymbol{g}_1, \ldots, \boldsymbol{g}_\nu\}$, are obtained by heavily downsampling the input data and the edges are generated by connecting each node with its k neighbors.

Once the deformation graph is created, we optimise the local rotations $\boldsymbol{R}_i \in \mathbb{R}^{3\times3}$ and local translations $\boldsymbol{t}_i \in \mathbb{R}^3$ for each node in the deformation graph $\mathcal{G}$. This optimisation is performed by minimising an energy function that accounts for the pairwise distance between feature nodes E_{con}, the rotations E_{rot}, and the regularization E_{reg} such that

$$\underset{\boldsymbol{R}_1, \boldsymbol{t}_1, \ldots, \boldsymbol{R}_\nu, \boldsymbol{t}_\nu}{\operatorname{argmin}} \; w_{con}E_{con} + w_{rot}E_{rot} + w_{reg}E_{reg} + w_{rig}E_{rig} \tag{2}$$

where E_{con} is the Euclidean distance between the SIFT

(a) (b) (c) (d)

Figure 3: Data used for the controlled experiment. The sample (a) is annotated (b), stretched and compressed (c), and the annotation are then removed to enable evaluating SIFT feature correspondence (d).

feature matches defined as

$$E_{con} = \sum_{l=1}^{n} \|\mathcal{F}_{1,l} - \mathcal{F}_{2,l}\|_2^2 . \tag{3}$$

E_{rot} adds the errors of all the rotations matrices and is defined similarly to [Jiawen *et al.*, 2012] as

$$E_{rot} = \sum_{j=1}^{\nu} \left\| \boldsymbol{R}_j{}^T \boldsymbol{R}_j - I \right\|_F^2 . \tag{4}$$

The regularisation term E_{reg} enforce smoothness over the deformation and is defined as

$$E_{reg} = \sum_{i=1}^{\nu} \sum_{j=1}^{\mu} \|\boldsymbol{R}_j(\boldsymbol{g}_i - \boldsymbol{g}_j) + \boldsymbol{g}_j + \boldsymbol{t}_j - (\boldsymbol{g}_i + \boldsymbol{t}_i)\|_2^2 . \tag{5}$$

The energy function described in (2) is then minimised with Levenberg-Marquardt optimisation. Once the new position of the deformation graph nodes is known, the points of $\mathcal{P}_1$ are updated using

$$\boldsymbol{p}_i^* = \sum_{j=1}^{\mu} w_j(\boldsymbol{p}_i)[\boldsymbol{R}_j(\boldsymbol{p}_i - \boldsymbol{g}_j) + \boldsymbol{g}_j + \boldsymbol{t}_j], \tag{6}$$

with the neighbour's nodes $\boldsymbol{g}_j$ from $\mathcal{P}_i$ found using a search with a kD-tree. The weight for each vertex is defined as

$$w_j(\boldsymbol{p}_i) = (1 - ||\boldsymbol{p}_i - \boldsymbol{g}_j||/d_{max}), \tag{7}$$

where d_{max} is the maximum distance of the vertex to the $\mu + 1$ nearest node from $\mathcal{G}$.

4 Experiments

Our algorithm is tested on two different scenarios. The first is a controlled environment that is performed to quantitatively evaluate the performance of our approach. The second experiment is performed with data collected in the field and is used for a qualitative evaluation and discussion.

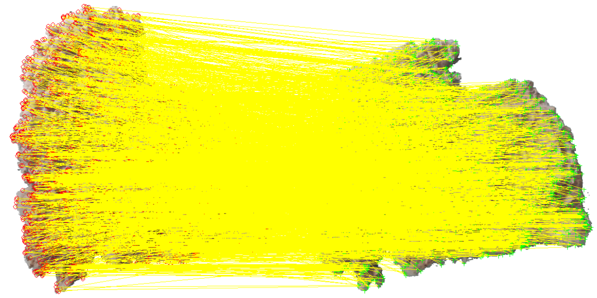

(a) Original SIFT matches.

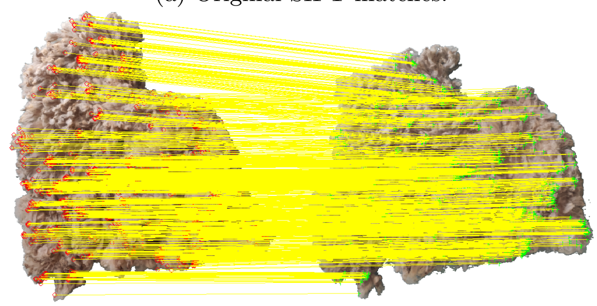

(b) Filtered SIFT matches.

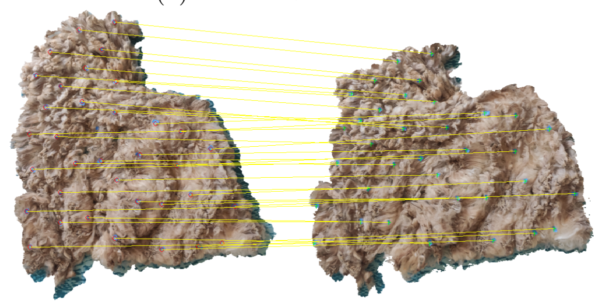

(c) Manual annotations.

Figure 4: SIFT features matching before (in (a)) and after (in (b)) the filtering process described in Section 3.2. In (c), we show the correspondences from the annotations.

Data collected for the control experiment are shown in Figure 3: a piece of wool shown in Figure 3(a), has been annotated (Figure 3(b)), stretched and compressed (Figure 3(c)), and the annotation has been removed to avoid corrupting the SIFT features extractions (Figure 3(d)). The annotation alongside with the boundary of the wool are used to quantify the output of the non-rigid deformation.

The matchings of the SIFT features are shown in Figure 4(a). As shown, many outliers are present in the feature matching and need to be removed. The filtered feature matching using the proposed ARAP filter are shown in Figure 4(b). The filtered matching visually corresponding with the manual annotations obtained from Figure 3(b) and 3(c) and are shown in Figure 4(c) for reference.

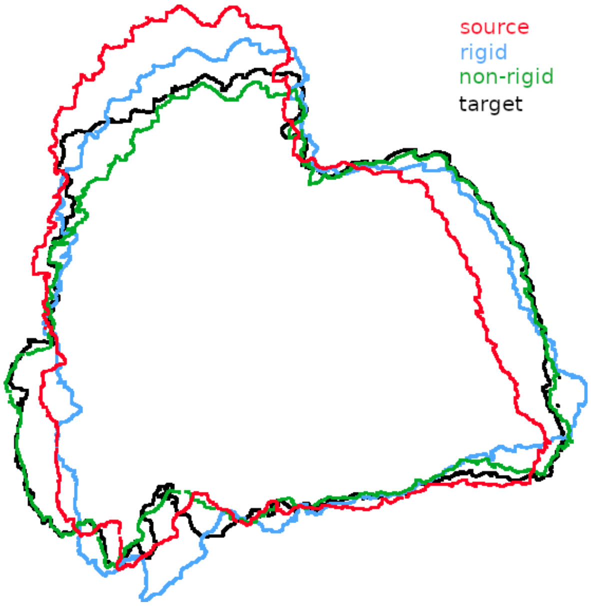

Figure 5: Contours using rigid and non-rigid deformation used for evaluation. The rigid deformation resulting from SIFT+RANSAC filter is shown in blue and the non-rigid deformation produced using SIFT+ARAP filter is shown in green. The original contour (red) and the target (black) are shown for reference.

Given the filtered feature matching, we then perform the non-rigid deformation using ED as discussed in Section 3.3. The updated contours of the deformed image are displayed in Figure 5 in green alongside the original contour (in red) and the targeted image (in black). We compare our approach to a standard method that includes SIFT matchings, RANSAC filtering and rigid deformation. The contour of this deformation is displayed in blue in Figure 5.

A quantitative evaluation is performed by analysing

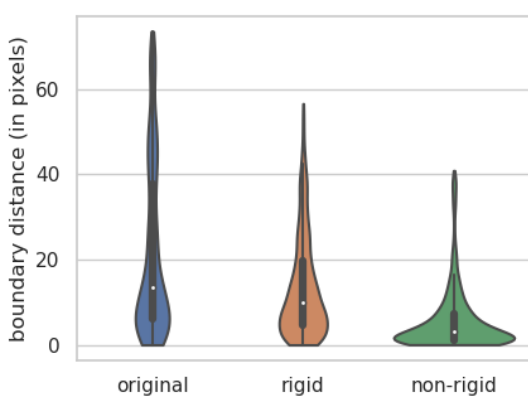

Figure 6: Distribution of the distance between the contours (shown in Figure 5). Our proposed method generates contours that are significantly closer to the target.

Table 1: Quantitative evaluation and ablation study: starting from the SIFT features matching, the outliers are filtered using RANSAC (rigid) or the proposed ARAP filter (non-rigid). The image is then deformed rigidly or non-rigidly. The measurement without deformation is provided for reference. Distances are in pixels. * these experiments were run 50 times given the non deterministic property of RANSAC.

		boundary		annotations	
filtering	deform	mean	std	mean	std
none	none	21.83	20.18	40.17	18.22
rigid*	rigid	21.16	17.98	36.63	19.39
rigid*	non-rigid	18.86	16.57	34.29	17.85
non-rigid	rigid	19.57	19.73	23.90	18.88
non-rigid	non-rigid	5.75	7.44	14.04	6.66

the distance between the deformed contours and the deformed manual annotations. The annotations (which are extracted manually using tags on the wool) provide information regarding the accuracy of the deformation across the whole surface. It provides us with an accurate quantification of the deformation error. The mean error and standard deviation of these annotations are reported in Table 1.

As a complementary quantitative evaluation, we measure the distance between the boundaries of the deformation, which closely relates to the aim of our work, *i.e.*, solving the detection of the skirting lines. The distribution of the distance between the boundaries is defined by finding the nearest point of the targeted boundary, in Euclidean space, for each point of the deformed boundary. These distributions are displayed in the violin plots of Figure 6 and the mean and standard deviation are given in Table 1.

The ablation study in Table 1, which considers rigid / non-rigid filtering and rigid / non-rigid deformation, shows that the problem needs to be tackled non-rigidly for both the features matching filtering and the deformation.

We also compare the proposed method with manual annotation made by expert wool handlers. These annotations are a reference for where the wool is to be skirted. The wool was then actually skirted, allowing us to compare the data before and after the process. The deformation of the wool after skirting has been manually assessed (with samples shown in Figure 8) and it shows convincing results regarding the quality of the deformation (*i.e.*, specific patterns in the wool are visible in the same location for both the pre- and post-skirted images). However, as shown in Figure 7, there is a significant difference between the manual annotation (which corresponds to a regular trimming of a few centimetres) and the actual skirting by the wool handler. The dis-

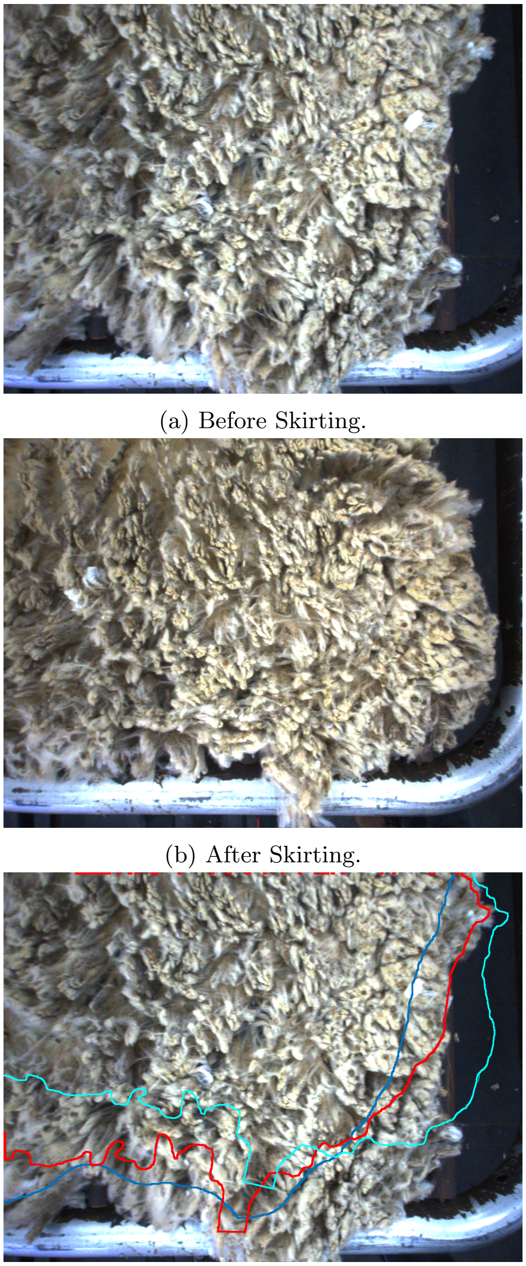

(a) Before Skirting.

(b) After Skirting.

(c) Contour comparison between methods.

Figure 7: Data used to compare the different methods. In (c) we demonstrate the accuracy of each method with the use of contours: Wool handlers annotation (blue), proposed method (red) and rigid deformation (cyan).

crepancy can be explained by the difficulty for a human to assess the quality of the wool in an image without manipulating it or having a close inspection. This issue, as highlighted here, strongly supports the requirement for our proposed method.

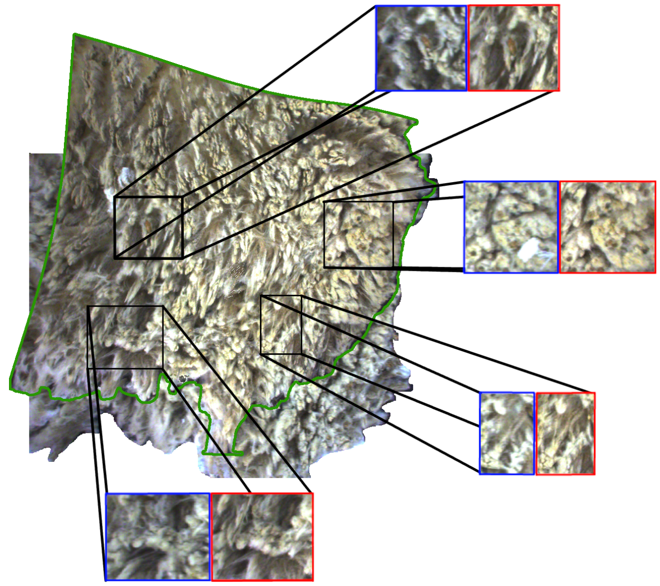

Figure 8: Qualitative evaluation of the morphing on data from the field. The deformed skirted image (which border is shown in green) is superposed to the original wool image. In blue, we show the wool image prior to skirting and in red the sample after skirting. The areas selected are strictly superposed.

5 Discussion and Future Work

We present a method to automatically determine the delineation of clean and contaminated wool in freshly shorn fleece by comparing pre- and post-skirted RGB images. The process of manually removing contaminants, known as skirting, has handlers pulling sections of the wool that need to be removed. Fibre separation along the skirting line simultaneously stretches and compacts the fleece, thus undergoes non-rigid deformation. To overcome the challenge, we apply an efficient pruning procedure of feature correspondences using the concept of as-rigid-as-possible (ARAP) and then align the images through an optimisation process for non-rigid deformation.

We quantitatively verify our method in a controlled experiment by attaching markers to pieces wool and deforming them by hand and show that our method achieves much better alignment to the deformed target in comparison to the rigid registration baseline. We qualitatively show the applicability of our work on images acquired during operations in a wool shed. Our results highlight the unreliability of human annotation, which erroneously marks portions of the wool that is to be removed but is still visibly present in the post-skirted images.

The results provide initial evidence of the validity of our approach, but there are a number of areas of future work that would need to be addressed in order to develop an end-to-end implemented system by exploiting a larger dataset beyond this preliminary study. Firstly, adequately thresholding for the ARAP filter, which depends on the density of the points and the size of the image. Secondly, evaluating the effect and constraints of the regularisation term that controls deformation.

Acknowledgments

This work is supported by funding from Australian Wool Innovation Limited, grant number ON-00713.

References

[DAWR, 2020] DAWR. Department of agriculture and water resources: Woolpool review. `https://www.agriculture.gov.au/ag-farm-food/meat-wool-dairy/wool`, 2020. Accessed: 2021-09-27.

[Fischler and Bolles, 1981] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981.

[Hansford, 2012] Kerry Hansford. Contamination. *University of New England, WOOL472:Wool biology and metrology (coursenotes)*, 2012.

[Jiawen *et al.*, 2012] Chen Jiawen, Izadi Shahram, and Fitzgibbon Andrew. KinEtre: Animating the world with the human body. *ACM symposium on User interface software and technology*, pages 435–444, 2012.

[Kavan *et al.*, 2007] Ladislav Kavan, Steven Collins, Jiří Žára, and Carol O'Sullivan. Skinning with dual quaternions. In *Proceedings of the Symposium on Interactive 3D Graphics and Games*, pages 39–46, 2007.

[Klein *et al.*, 2007] Stefan Klein, Marius Staring, and Josien PW Pluim. Evaluation of optimization methods for nonrigid medical image registration using mutual information and b-splines. *IEEE Transactions on Image Processing*, 16(12):2879–2890, 2007.

[Krizhevsky *et al.*, 2012] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, volume 25, 2012.

[Lowe, 2004] David G Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004.

[Patten *et al.*, 2021] Timothy Patten, Alen Alempijevic, and Robert Fitch. Learning image-based contaminant detection in wool fleece from noisy annotations. In *Proceedings of the International Conference on Computer Vision Systems*, pages 234–244, 2021.

[Pluim *et al.*, 2003] Josien PW Pluim, JB Antoine Maintz, and Max A Viergever. Mutual-information-based registration of medical images: a survey. *IEEE Transactions on Medical Imaging*, 22(8):986–1004, 2003.

[Rueckert *et al.*, 1999] Daniel Rueckert, Luke I Sonoda, Carmel Hayes, Derek LG Hill, Martin O Leach, and David J Hawkes. Nonrigid registration using free-form deformations: Application to breast MR images. *IEEE Transactions on Medical Imaging*, 18(8):712–721, 1999.

[Sorkine and Alexa, 2007] Olga Sorkine and Marc Alexa. As-rigid-as-possible surface modeling. In *Symposium on Geometry Processing*, volume 4, pages 109–116, 2007.

[Sperl *et al.*, 2021] Georg Sperl, Rahul Narain, and Chris Wojtan. Mechanics-aware deformation of yarn pattern geometry. *ACM Transactions on Graphics*, 40(4), 2021.

[Su *et al.*, 2006] Zhenwei Su, Gui Yun Tian, and Chunhua Gao. A machine vision system for on-line removal of contaminants in wool. *Mechatronics*, 16(5):243–247, 2006.

[Sumner *et al.*, 2007] Robert W. Sumner, Johannes Schmid, and Mark Pauly. Embedded deformation for shape manipulation. *ACM Transactions on Graphics*, 2007.

[Tran *et al.*, 2012] Quoc-Huy Tran, Tat-Jun Chin, Gustavo Carneiro, Michael S Brown, and David Suter. In defence of ransac for outlier rejection in deformable registration. In *European Conference on Computer Vision*, pages 274–287, 2012.

[Wei *et al.*, 2019] Wei Wei, Dexiang Deng, Lin Zeng, Chen Zhang, and Wenxuan Shi. Classification of foreign fibers using deep learning and its implementation on embedded system. *International Journal of Advanced Robotic Systems*, 16(4):1–20, 2019.

[Wei *et al.*, 2020] Wei Wei, Chen Zhang, and Dexiang Deng. Content estimation of foreign fibers cotton based on deep learning. *Electronics*, 9(11), 2020.

[Zhang *et al.*, 2005a] Liwei Zhang, A. Dehghani, Zhenwei Su, T. King, B. Greenwood, and M. Levesley. Development of a mechatronic sorting system for removing contaminants from wool. *IEEE/ASME Transactions on Mechatronics*, 10(3):297–304, 2005.

[Zhang *et al.*, 2005b] Liwei Zhang, Abbas Dehghani, Zhenwei Su, Tim King, Barry Greenwood, and Martin Levesley. Real-time automated visual inspection system for contaminant removal from wool. *Real-Time Imaging*, 11(4):257–269, 2005.