Dermascopic hair disocclusion using inpainting

Paul Wighton[a,b,c] and Tim K. Lee[a,b,c] and M. Stella Atkins[a]

[a]School of Computing Science, Simon Fraser University, Burnaby BC, Canada;

[b]BC Cancer Research Center, Vancouver BC, Canada;

[c]Photomedicine Institute, Department of Dermatology and Skin Science, University of British Columbia and Vancouver Coastal Health Research Institute, Vancouver BC, Canada

ABSTRACT

Inpainting, a technique originally used to restore film and photographs, is used to disocclude hair from dermascopic images of skin lesions. The technique is compared to the conventional software DullRazor, which uses linear interpolation to perform disocclusion. Comparison was performed by simulating occluding hair on a dermascopic image, applying DullRazor and inpainting and calculating the error induced. Inpainting is found to perform approximately 33% better than DullRazor's linear interpolation, and is more stable under heavy occlusion. The results are also compared to published results from two other alternatives: auto-regressive (AR) model signal extrapolation and band-limited (BL) signal interpolation.

Keywords: Image restoration and enhancement, Disocclusion, Inpainting, DullRazor

1. INTRODUCTION

Malignant melanoma poses a significant risk. While the overall age-standardized mortality rate of cancer is decreasing in Canada, the mortality rate of melanoma continues to rise with a 2-3% annual increase in incidence rates over the last thirty years.[1] Prognosis for advanced melanoma remains poor, with a five-year survival rate of around 50% for lesions thicker than 3.5mm. However if melanoma is detected early (while the lesions thickness is less than 1.5mm), the ten-year survival rate is over 90%.[2] Therefore, the early diagnosis of melanoma is critical so that it can be completely excised while it is still localized. As a result, there is a considerable amount of ongoing research into the automated diagnosis of melanoma from digital dermascopic images. Nearly all published methods include a pre-processing step at the beginning of the analysis. Pre-processing is usually responsible for transforming the image into a convenient and workable form, as well as segmentation: the process by which the boundary of the skin lesion is determined. The disocclusion of hair is a vital part of pre-processing; occluding hairs can mislead the segmentation algorithm as well as any further analysis. Disocclusion can be broken into two distinct steps: 1) The pixels of the object to be disoccluded are first be identified, then 2) the underlying color of the scene is estimated for these pixels. This paper applies a method to improve upon the second stage of this process. A technique called inpainting is applied to increase the accuracy in estimating the underlying color of occlusion pixels.

Inpainting is the art of modifying an image in a way that is undetectable to the casual observer. It was originally used to restore damaged artwork, however recent automated inpainting methods have found many modern applications in areas such as restoration, disocclusion, increasing resolution,[3] filling 'holes' in 3D models[4] and wireless transmission.[5] It has largely been a manual process until Bertalmio, Sapiro et al.'s seminal paper in 2000.[6] They consulted conservators at the Minneapolis Institute of Arts in order to analyze how experts inpaint. While the experts disclaimed that there in no single correct way to inpaint, and that inpainting is largely a subjective process based on both the artwork and artist, they described the overall process shown in figure 1. Before beginning, the entire picture is analyzed to give guidance as to how specific areas should be filled. The inpainting process then involves: (1) Continuing structural elements into the gap; (2) filling the gap with the colour of the boundary; and (3) Adding texture.

Further author information: (Send correspondence to Paul Wighton.)

Paul Wighton: pwighton@sfu.ca

Tim K. Lee: tlee@bccrc.ca
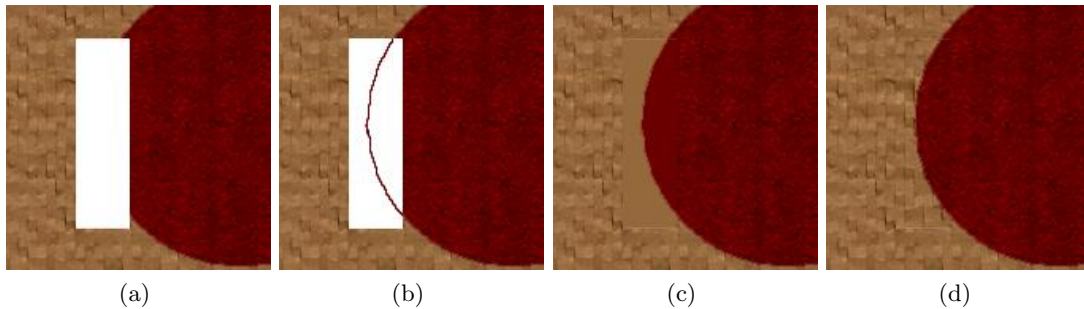
M. Stella Atkins: stella@cs.sfu.ca

Figure 1. How experts inpaint: a) The entire image is analyzed to determine how to fill in the gap (white). b) Structural elements are propagated. c) Colour is propagated. d) Texture is added.
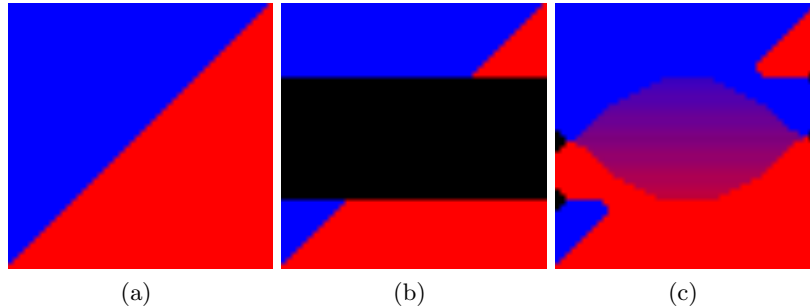

Figure 2. Phantom image demonstrating the limitations of DullRazor a) The original image b) The masked image intersecting the image gradient at a 45 degree angle. c) The result from DullRazor

Section 2 reviews previous work in dermascopic hair disocclusion as well as inpainting. Section 3 describes how the inpainting technique is compared to DullRazor;[7] the seminal work in dermascopic hair disocclusion. Section 4 reports qualitative and quantitative results of the comparison and section 5 concludes with discussion and comparison to other published methods.

2. PREVIOUS WORK

2.1 Dermacopic Hair Disocclusion

The seminal work in dermascopic hair disocclusion is Lee et al.'s freely available program DullRazor.[7] DullRazor first identifies dark occluding hair by means of a grayscale morphological closing operation applied separately to each of the colour channels.[8] Next, for each hair pixel, the two closest non-hair pixels in the direction perpendicular to the of the hair are determined and linear interpolation is used to estimate the underlying skin color across the width of the hair. While generally quite effective, there are some instances in which DullRazor performs poorly at estimating the replacement colour. Figure 2 illustrates one such instance: when a gradient occurs in the image that is at an angle to the direction of the hair. Additionally, there is no means to ensure smoothness among the pixels being replaced. This can cause artefacts in the resulting image, which may affect the segmentation process. Also, the linear interpolation technique is inherently one dimensional and therefore a reasonable direction in which to interpolate must be determined. This becomes more problematic as the amount of occluding hair increases, as illustrated in Figure 3.

She et al. propose two alternate methods to estimate the underlying color of hair pixels.[9] Their first method is an auto-regressive (AR) model signal extrapolation technique. This stochastic process uses the N previous outputs to recursively generate the current output. Their paper describes a 1-dimensional formulation of this model. They use either side of the hair as initialization points and extrapolate across the width of the hair in a similar fashion to DullRazor. As a result, this method suffers from the same problem as DullRazor: it is a one-dimensional technique being applied to a two-dimensional image and thus a reasonable extrapolation direction must be determined. Their second method is a band-limited (BL) signal interpolation technique. This method takes the Fourier transform of the image, sets the response outside of a defined region to zero, takes the
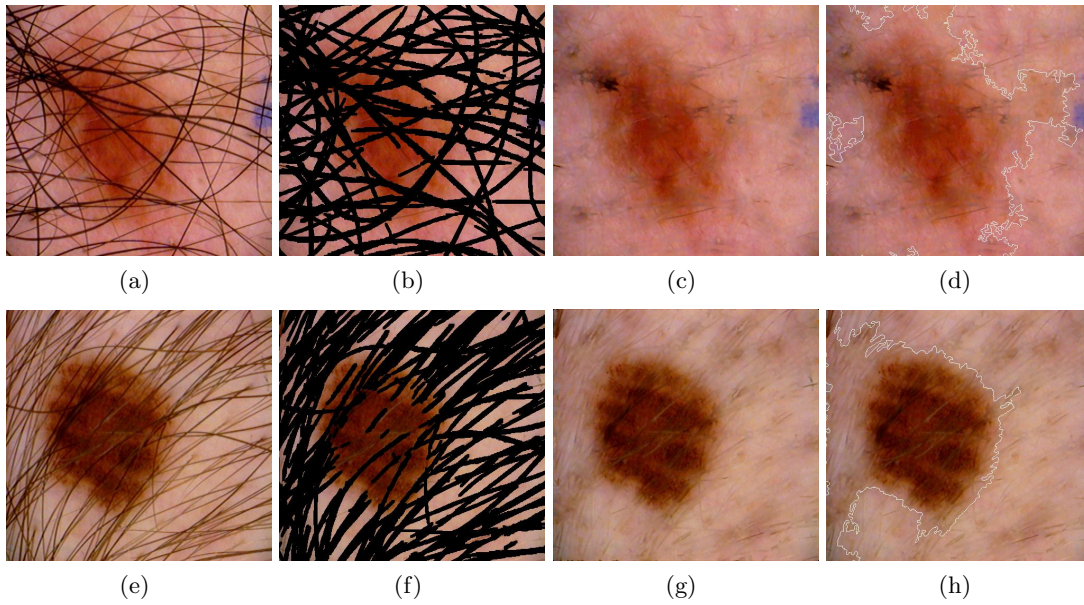
Figure 3. Difficult real-world examples of skin lesions with many occluding hairs. a), e) The original images b), f) The masked images after the 'hair pixels' have been identified. c), g) The results from DullRazor's linear interpolation. d), h) The resulting inaccurate segmentations.

inverse Fourier transform and updates the pixels within the masked region accordingly. This process iterates until convergence.

2.2 Inpainting

There now exists many inpainting algorithms for many different tasks. The original algorithm[6] inpainted structure only, resulting in filled regions lacking texture (such as in figure 1c) however more recent work includes algorithms to inpaint texture,[10] video[3] and to decompose an image into a structural component (which can be inpainted) and textural component (which can treated with texture synthesis techniques[11]). The choice of the most appropriate algorithm is a critical decision. While the texture inpainting algorithms tend to yield more perceptually pleasing results, this comes at the expense of making larger assumptions about the underlying data. Since accuracy in segmentation and analysis is our motivation and not perceptual consistency, such large assumptions will surely induce more bias into these downstream processes. Therefore the original inpainting algorithm[6] that inpaints structure only, was chosen.

The key observation common to all algorithmic inpainting techniques is that preservation of structural elements is crucial.[6] This is illustrated in figure 4. Let I represent the image in question, and Ω be a region of the image to be inpainted (no assumptions as to the topology of Ω is made). Let $\delta\Omega$ represent the boundary between the region Ω, and the rest of the image I. We now must propagate 'information' from the image I into the region Ω. The 'information' we choose to propagate is unimportant at this point; it could be colour, texture, smoothness, etc. What is important is to decide *how* to propagate this information. For a given pixel p on the boundary $\delta\Omega$, instead of propagating information in the direction of the boundary normal (n_p) as linear interpolation (and therefore DullRazor) does, the novelty provided by inpainting algorithms is that information is propagated in the direction of *structure*. The structural importance of an area is determined by the magnitude of the gradient. Since the gradient defines the direction of the greatest rate of change, the direction of propagation is chosen to be perpendicular to the gradient ($\nabla I_p^{\perp}$).

The original inpainting algorithm uses 'smoothness' as the information to propagate into the region.[6] The Laplacian is used as the measure of smoothness, which is defined for a given pixel p as:
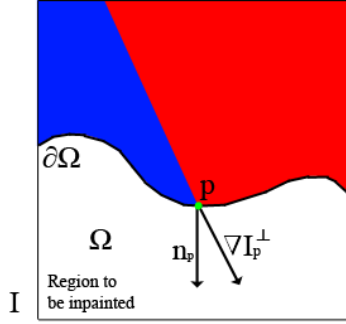
$$L(p) = I_{xx}(p) + I_{yy}(p) \qquad (1)$$

Figure 4. Algorithmic inpainting. For each pixel p on the boundary $\delta\Omega$, $\nabla I_p^\perp$ is chosen as the direction of information propagation (instead of n_p). The relative importance of information propagation is given by $\|\nabla I_p^\perp\|$

Where I_{xx} and I_{yy} denotes the second derivative in the x and y directions respectively. The problem is modeled as a differential equation. For each pixel along the boundary, smoothness is propagated in the direction of $\nabla I_p^\perp$ with a velocity of $\|\nabla I_p^\perp\|$. The formulation is solved using an iterative discreet scheme implementation. Additionally, after A iterations of inpainting, B iterations of anisotropic diffusion are performed.

3. METHODS

This section describes how the inpainting technique described in section 2.2 was compared to DullRazor. A dermascopic image of a benign nevus was selected. A certain number of simulated hairs were superimposed onto the lesion, and then removed using both DullRazor, and inpainting. Error was calculated by measuring the absolute difference between the resulting images and the original. This process was repeated several times, each time changing the number of simulated hairs to be generated. Section 3.1 describes the creation of the simulated hairs, section 3.2 describes how error is calculated and section 3.3 details the inpainting parameters used.

3.1 Simulating Hair

Simulated hairs were created by first generating a hair mask. A hair is generated by randomly selecting two points (x_1, y_1) and (x_2, y_2) in the image. Pixels that intersect this line segment are added to the hair mask. This is repeated n times. The resulting hair mask is then morphologically dilated with circle with a radius of 2. Once the hair mask has been created, pixels within the mask are set to black (RGB value of $(0, 0, 0)$). These images are then passed to the DullRazor and inpainting algorithms.

Additionally, tests were performed as She et al. describe.[9] They describe three tests each involving a single hair 3 pixels wide and 100 pixels long. Each tests uses a different hair orientation: horizontal, vertical or diagonal.

3.2 Calculating Error

Error is calculated by computing the sum of the absolute difference between the original and resulting images across all color channels. Let I_o and I_r represent an original and resulting image respectively. Let $I(x, y, c)$ represent the value of the image at location (x, y) in the c^{th} color channel. Total error (TE) is then:

$$TE = \sum_{\forall (x,y,c)} |I_o(x, y, c) - I_r(x, y, c)| \tag{2}$$

and per-pixel error (PPE) is the total error divided by the number of pixels in the hair mask.

Error is also calculated as in She et al.'s paper,[9] by computing the mean and standard deviation of the pixel error. Since they do not take the absolute value of the error, negative errors and therefore negative means are possible. A negative mean implies that the disocclusion algorithm tends to underestimate the underlying intensity of the skin. Similarly, a positive mean implies a tendency to overestimate. Additionally the standard deviation represents how consistent the algorithm is at performing estimations.

Hairs	% of pixels in mask	Inpainting			DullRazor		
		TE	PPE	Mean error	TE	PPE	Mean error
1	0.52	17679	13.55	-0.08±5.95	27069	20.74	1.48±15.26
4	1.99	58871	11.91	-0.15±5.37	78464	15.88	0.89±7.94
9	4.82	160114	13.35	-0.34±6.14	204996	17.09	0.45±8.10
16	6.69	213267	12.82	0.09±5.99	299334	17.99	0.32±8.65
25	9.27	287836	12.48	-0.10±5.76	413540	17.94	1.40±9.52
36	17.59	543969	12.43	0.04±5.70	717734	16.40	0.65±7.79
49	20.19	674724	13.43	-0.13±6.18	873449	17.39	0.56±8.42
64	25.81	830863	12.94	-0.13±6.01	1105938	17.22	0.67±9.05
81	32.67	1104086	13.58	0.21±6.34	1370704	16.86	0.86±8.56
100	37.88	1299338	13.79	0.02±6.37	1610241	17.08	0.77±8.74
121	44.18	1597838	14.53	0.13±6.73	1977745	17.99	1.23±9.97
144	46.55	1807199	15.60	-0.04±7.37	2275788	19.65	1.91±12.79
169	51.97	2017252	15.60	0.64±7.17	2482936	19.20	1.87±12.37
196	58.08	2475163	17.13	1.34±7.99	3722322	25.76	4.26±20.59

Table 1. Total error (TE), per pixel error (PPE) and mean error of DullRazor and inpainting with hair masks of varying sizes.

Direction	Inpainting	DullRazor	Linear Interpolation	AR extrapolation	BL interpolation
Horizontal	0.157 ± 3.92	-0.619 ± 3.58	-6.74 ± 35.68	-2.58 ± 15.24	-0.49 ± 9.42
Vertical	-0.583 ± 2.54	-0.485 ± 2.94	-5.50 ± 29.87	-2.42 ± 15.80	-0.09 ± 10.78
Diagonal	0.069 ± 3.38	1.06 ± 3.52	-3.76 ± 24.97	-2.31 ± 16.67	-1.14 ± 12.74

Table 2. Mean and standard deviation of pixel error for Inpainting, DullRazor as well a reproduction of She et al.'s published results[9] for Linear Interpolation, AR extrapolation and BL interpolation.

3.3 Inpainting Parameters

The inpainting algorithm was implemented in Matlab. Peter Kovesi's anisotropic diffusion implementation was used.[12] Certain parameters needed to be chosen for the inpainting algorithm. The number of iterations was set to 4000, A (the number of inpainting steps) was set to 100 and B (the number of steps of anisotropic diffusion) was set to 2. Parameters for anisotropic diffusion was set as follows: κ (conduction coefficient) was set to 50, and λ was set to 0.25. It is worth noting that very little time was spent tuning these inpainting parameters. It is very likely that the performance of inpainting would improve with further tuning.

4. RESULTS

Several trials were conducted, as described in section 3.1, each with varying amounts of simulated hair. Figure 5 illustrates the qualitative results of several trials. Error was calculated as described in section 3.2 and the results are summarized in table 1. Inpainting performs on average 32.7% better than the linear interpolation of DullRazor.

Single-hair tests described in section 3.1 designed to mimic she et al.'s testing procedure[9] were also performed. Our results, along with she et al.'s published results are summarized in table 2.

Additionally, real-world qualitative tests were performed by applying DullRazor and the inpainting algorithm to the heavily occluded images in figure 3, and the results are shown in figure 6.

5. CONCLUSIONS AND DISCUSSION

From table 1, it can be seen that inpainting outperforms DullRazor. On average, inpainting performs 32.7% than DullRazor. It is also more stable under heavy occlusion, which can been seen in figure 5 (Note the artefacts in figure 5s).
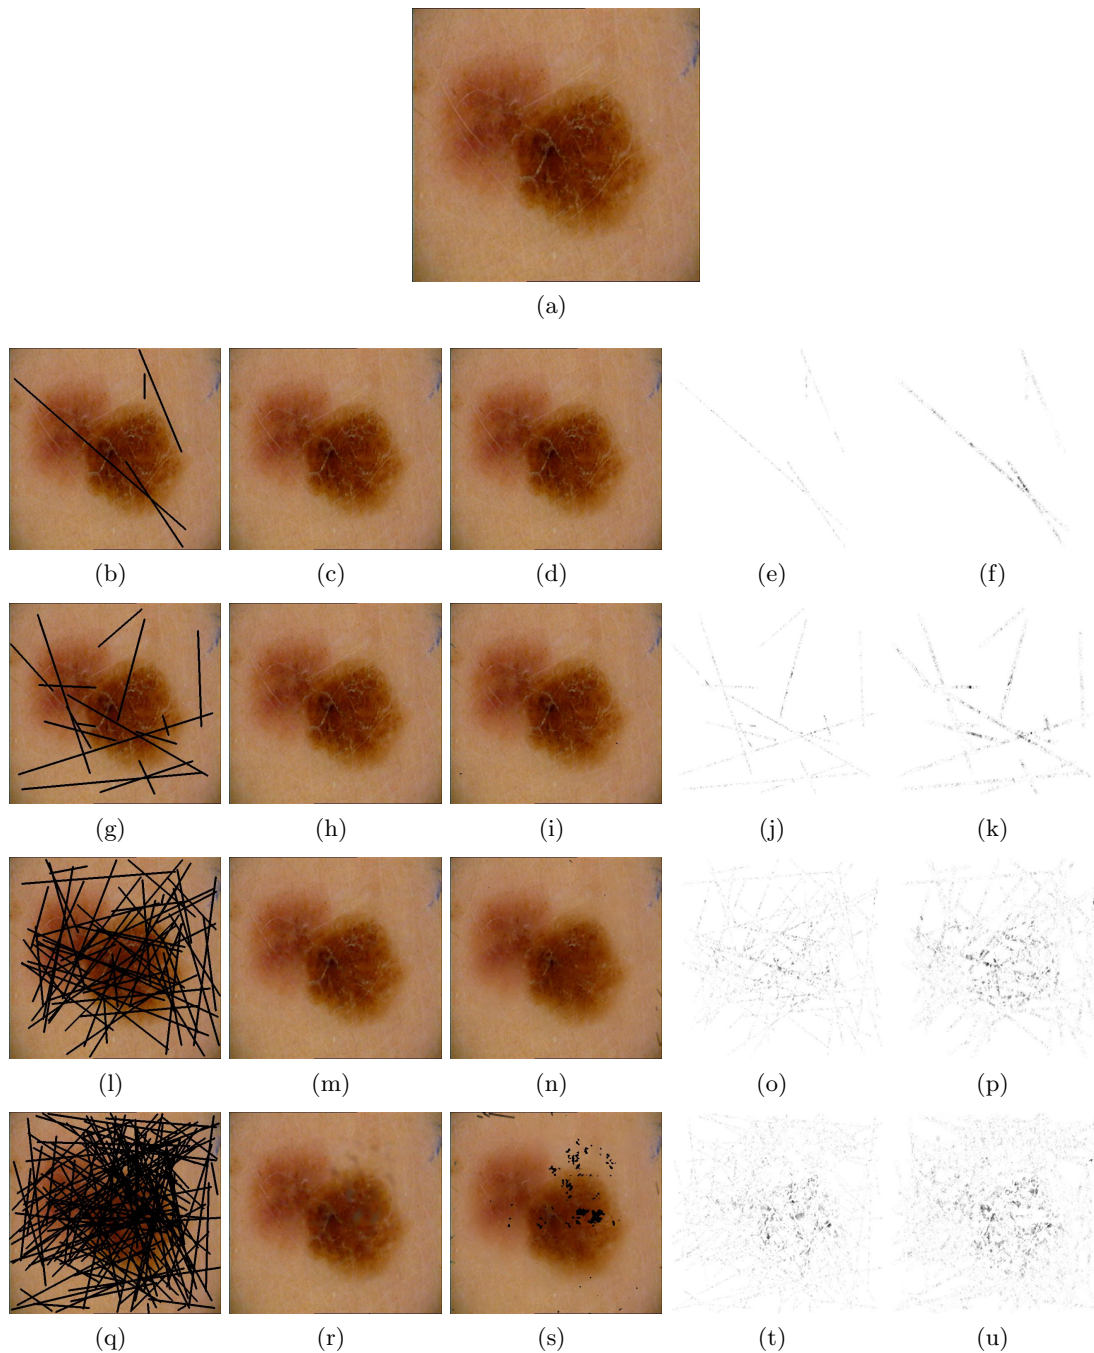
Figure 5. Qualitative results of applying inpainting and DullRazor to synthetic hair masks with 4, 16, 64, and 144 randomly generated hairs respectively. a) The original image (ground truth). 1st column: The image with generated hair masks; 2nd column: Results from inpainting; 3rd column: Results from DullRazor; 4th column: Inpainting difference map; 5th column: DullRazor difference map. (Difference maps are exaggerated for illustrative purposes)
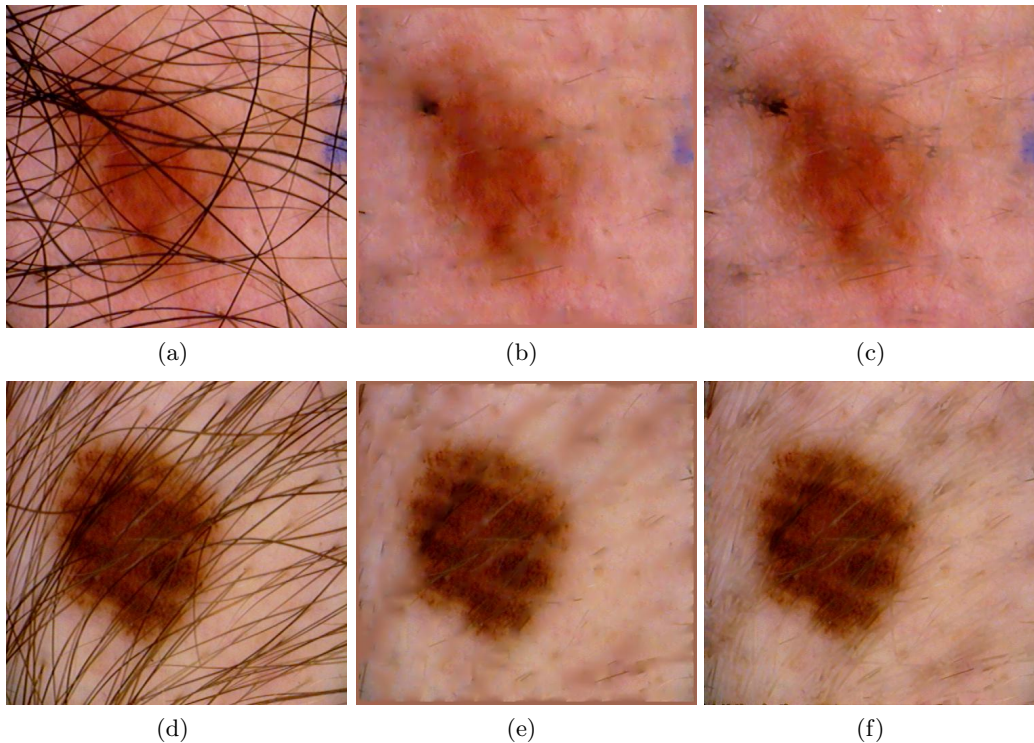
Figure 6. Real world qualitative results. a), d) Dermascopic images with many occluding hairs b), e) the results from inpainting c), f) the results from DullRazor

The results were also compared to She et al.'s results[9] in which a single simulated hair 3 pixels wide and 100 pixels long is disoccluded using linear interpolation as well as their proposed methods. Our results, as well as their reported results are reproduced in table 2. It is surprising that DullRazor seems to consistently outperform She et al.'s linear interpolation algorithm. The errors obtained for DullRazor are an order of magnitude smaller than those reported by She et al. It is unclear why this is the case. Moreover both DullRazor and inpainting seem to perform equally well under these test conditions. Additionally, DullRazor and inpainting outperform She et al.'s AR model technique. From this comparison, it is unclear whether DullRazor, Inpainting or the Band-Limited method is the superior algorithm, however the standard deviation of the error is significantly smaller for DullRazor and Inpainting. This implies that DullRazor and inpainting perform more consistent estimations. Also, since She et al. only performed tests with a single synthetic hair, it is unclear how their algorithms perform under heavy occlusion.

Much insight can also be gained by examining the qualitative results in figure 6. While it is unclear which results are more perceptually pleasing, the inpainting results don't contain the strong gradients found those of DullRazor. It is likely that this will aid in obtaining accurate segmentations. Also, it is interesting to compare the real world results in figure 6 to that of the synthetic results in figure 5. It appears that both DullRazor and Inpainting perform much better under synthetic conditions. Recall that hair disocclusion consists of two distinct steps: 1) Identifying 'hair pixels', and 2) estimating the underlying color of these 'hair pixels'. All the methods discussed in this paper (linear interpolation, inpainting, She et al.'s techniques) are all techniques to deal with the latter half of the hair disocclusion problem. While inpainting seems to be a viable approach to this part of the problem, it is clear from the staggering differences between figures 6 and 5 that accurately identifying 'hair pixels' (and possibly accounting for the penumbra effect) is currently the 'weak point' in the dermascopic hair disocclusion process.

ACKNOWLEDGMENTS

REFERENCES

1. Canadian Cancer Society/National Cancer Institute of Canada, "Canadian cancer statistics 2007." Toronto, Canada, 2007.
2. M. Lens and M. Dawes, "Global perspectives of contemporary epidemiological trends of cutaneous malignant melanoma," *British Journal of Dermatology* **150**(2), pp. 179–185, 2004.
3. M. Bertalmio, A. L. Bertozzi, and G. Sapiro, "Navier-stokes, fluid dynamics, and image and video inpainting," *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)* **01**, p. 355, 2001.
4. J. Verdera, V. Caselles, M. Bertalmio, and G. Sapiro, "Inpainting surface holes," *Image Processing, 2003. ICIP 2003. Proceedings. 2003 International Conference on* **2**, pp. II–903–6 vol.3, 14-17 Sept. 2003.
5. S. Rane, G. Sapiro, and M. Bertalmio, "Structure and texture filling-in of missing image blocks in wireless transmission and compression applications," *Image Processing, IEEE Transactions on* **12**(3), pp. 296–303, March 2003.
6. M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester, "Image inpainting," in *Siggraph 2000, Computer Graphics Proceedings*, K. Akeley, ed., pp. 417–424, ACM Press / ACM SIGGRAPH / Addison Wesley Longman, 2000.
7. T. Lee, V. Ng, R. Gallacher, A. Coldman, and D. McLean, "Dullrazor(r): A software approach to hair removal from images," *Computers in Biology and Medicine* **27**, pp. 533–543(11), November 1997.
8. T. K. Lee, "Measuring border irregularity and shape of cutaneous melanocytic lesions (phd thesis)."
9. Z. She, P. Fish, and A. W. Duller, "Improved approaches to hair removal from skin image," in *Proc. SPIE Vol. 4322, p. 492-499, Medical Imaging 2001: Image Processing.*, M. Sonka and K. M. Hanson, eds., **4322**, pp. 492–499, July 2001.
10. A. Criminisi, P. Perez, and K. Toyama, "Object removal by exemplar-based inpainting," *Computer Vision and Pattern Recognition, 2003. Proceedings. 2003 IEEE Computer Society Conference on* **2**, pp. II–721–II–728 vol.2, 18-20 June 2003.
11. A. A. Efros and T. K. Leung, "Texture synthesis by non-parametric sampling," *iccv* **02**, p. 1033, 1999.
12. P. D. Kovesi, "MATLAB and Octave functions for computer vision and image processing." School of Computer Science & Software Engineering, The University of Western Australia. Available from: <http://www.csse.uwa.edu.au/∼pk/research/matlabfns/>.