Selective Multi-Source Total Variation Image Restoration

Stephen Tierney
School of Computing and Mathematics
Charles Sturt University
Bathurst, NSW 2795, Australia
stierney@csu.edu.au

Yi Guo
Digital Productivity & Services Flagship
CSIRO
North Ryde, NSW 2113, Australia
yi.guo@csiro.au

Junbin Gao
School of Computing and Mathematics
Charles Sturt University
Bathurst, NSW 2795, Australia
jbgao@csu.edu.au

***Abstract*—This paper is concerned with automatically fusing multiple noisy and partially corrupted source images into a single denoised image. To create the fused image we minimise a convex objective function, which ensures spatial smoothness through total variation regularisation, and similarity to the source images via pixel-wise selective regularisation against each of the source images. We call this approach Selective Multi-Source Total Variation Image Restoration (SMTV). Applications of SMTV include noise removal in low-light conditions, enhancement of images from low quality or damaged imaging sensors and haze or cloud removal from satellite imagery. Experimental evaluation demonstrates that the fusion of multiple images results in a more accurate recovery than single image restoration.**

I. Introduction

Removal of noise and blurring effects in images has been a long standing problem in image processing. These unwanted artefacts can be introduced during both acquisition and transmission of the images. In fact the presence of noise is almost entirely unavoidable in virtually all image capture systems, particularly in low-light environments. These perturbations can have a detrimental effect on the ability on further image analysis applications such as segmentation, classification, object detection or tracking. For example recent work has shown [1] that imperceptible noise can seriously affect image classification results when using convolutional neural networks.

Due to the importance of image restoration a large variety of approaches have been proposed. These range from simple filtering approaches [2], [3], [4], to transformed domain filtering [5], [6], statistical methods [3] and more advanced methods which exploit of repeated image features [7], [8], [9].

Although the previous denoising methods have shown tremendous success in many cases, they are only applicable to single images. Therefore they suffer in the following two ways:

- recovery of all fine detail is not always possible from a single image
- recovery of large sections of a single image, which have been occluded or totally corrupted is impossible

To understand the first issue, consider that every imaging device will capture random noise along with the subject due to the image formation process. This effect is easily noticeable in low-light environments. This noise may completely destroy some parts of the image, in particular it overwhelms fine details such that recovery is impossible.

The second issue can be intuitively understood. If there is no information at all about the subject in a portion of the image, we cannot recover any information in that portion. For example weather and atmospheric effects such as cloud or dust in satellite imagery can completely obstruct the view of the ground. Depending upon the density of the cloud recovery of the surface may be impossible.

In this paper we present a method called Selective Multi-Source Total Variation Image Restoration (SMTV), which resolves both issues of single image restoration by restoring a single unified image from multiple noisy and partially corrupted images. We exploit the knowledge that random noise in images will differ from exposure to exposure and with enough exposures, the full image detail can be recovered. The same principle is applied to occlusions and corruptions, where many exposures with varied occlusions can reveal enough detail even when large portions of the source images are obstructed. We provide full source code and data used in our experiments[1].

II. Total Variation Based Image Restoration

Total Variation Restoration [3], [10], [11], [12] assumes the following image formation model

$$\mathbf{I} = P(\mathbf{A}) + \mathbf{N} \tag{1}$$

where $\mathbf{A}, \mathbf{N}, \mathbf{I} \in \mathbb{R}^{d \times N}$ are the original image, Gaussian noise and the observed image respectively, $N = m \times n$ is the number of pixels, m is the number of vertical pixels, n is the number of horizontal pixels, d is the number of colour channels and P is a linear operator representing a process such as blurring. Note that each channel of the images have been unrolled into a row vector.

To recover the original image $\mathbf{A}$ the following convex objective is used

$$\min_{\mathbf{A} \in B_{l,h}} \lambda \sum_{i}^{d} \mathrm{TV}(\mathbf{A}_{(i,:)}) + \frac{1}{2} \|P(\mathbf{A}) - \mathbf{I}\|_F^2 \tag{2}$$

where $B_{l,h}$ is the bounding constraint:

$$B_{l,h} = \{\mathbf{A} \in \mathbb{R}^{m \times n} | l \leq A_{ij} \leq h, \forall i, j\},$$

[1] http://github.com/sjtrny/smtv

978-1-4673-6795-0/15/$31.00 ©2015 IEEE

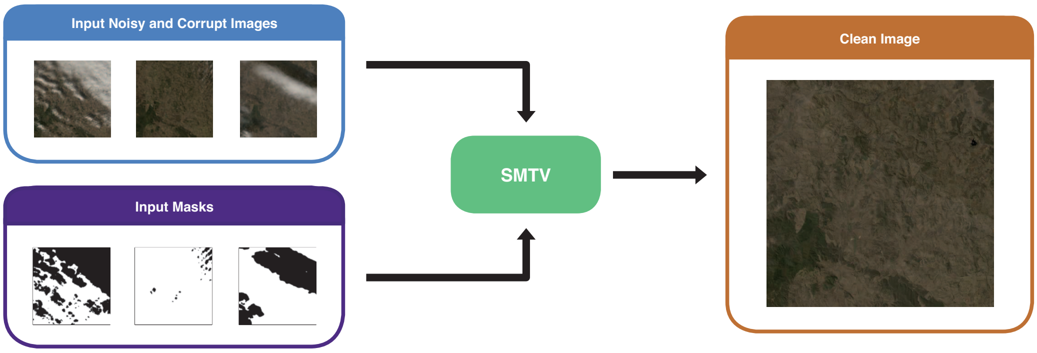

Fig. 1: Visual representation of the SMTV process.

$\mathrm{TV}(\mathbf{b})$ is the discrete total variation of $\mathbf{b}$ and $\mathbf{A}_{(i,:)}$ is the i'th row of $\mathbf{A}$ i.e. the i'th colour channel and l and h are the lower and upper bounds. Note that the total variation of each channel is considered seperarely. There are two versions of TV: ℓ_2 isotropic and ℓ_1 anisotropic. Isotropic TV is defined as

$$\mathrm{TV}_i(\mathbf{b}) = \sum_{x=1}^{m}\sum_{y=1}^{n}\sqrt{(b_{x+1,y}-b_{x,y})^2+(b_{x,y+1}-b_{x,y})^2}$$

and anisotropic TV is defined as

$$\mathrm{TV}_a(\mathbf{b}) = \sum_{x=1}^{m}\sum_{y=1}^{n}|b_{x+1,y}-b_{x,y}|+|b_{x,y+1}-b_{x,y}|$$

where $b_{x,y}$ is the pixel value of b at the original spatial pixel column and row coordinates x, y i.e. the row vector used in (2) has been reshaped to a matrix. Note that we have the following boundary condition for TV

$$\{b_{x+1,y}-b_{x,y}=0|x=n\}$$
$$\{b_{x,y+1}-b_{x,y}=0|y=m\}$$

since $m+1$ and $n+1$ exceed the image boundaries. In other words the total variation penalty encourages spatial smoothness through minimising the horizontal and vertical intensity differences between neighbouring pixels. Therefore (2) can be thought of as enforcing spatial smoothness while remaining as close to the original image as possible.

III. Selective Multiple-Source TV Restoration

Even though single image TV restoration methods have demonstrated relative success there are limitations to the amount of detail that can be recovered. For example consider taking a photograph with a digital camera sensor in a low-light environment. This photograph will contain a noticeable amount of random noise. This is particularly a problem for cameras in consumer technology such as point and shoot cameras, action cameras and mobile phones. As these devices shrink in size, so must the lens, camera sensor and the pixel pitch shrink, which leads to an increase in noise as there is less light available per pixel. In some circumstances this may mean that some fine image detail may be entirely lost and TV based methods will not be able to recover it.

There is another case where single image TV restoration struggles: large areas of corruption. Consider a satellite that has been tasked with capturing images of a specific patch of the Earth's surface. The view from the satellite may be partially or completely obstructed by aircraft, pollution, smoke and atmospheric effects such as cloud or fog. Recovery of the surface detail in these cases is often impossible, particularly when the atmospheric effects are totally opaque. Therefore the satellite may need to make multiple passes over the target area to produce an image as free as possible of such defects. However due to the satellites orbit and the nature of being a shared resources it is often the case that no clean image will be captured (in a short enough time span) of the target area.

These two problems are the central concern of this paper and we show how both can be simultaneously resolved using multiple images instead of a single image. Since in low-light photography the noise is random we exploit the fact that it will be different over multiple captures. When enough images are captured all detail within the scene can be resolved. A similar concept applies to the problem of occlusions, where if enough images are captured it is likely that each pixel is free of occlusion at least once among all captures. To perform this operation automatically we extend single image TV restoration to the case where multiple images are captured. Our objective function is the following:

$$\min_{\mathbf{A}\in B_{h,l}} \lambda\sum_{i}^{d}\mathrm{TV}(\mathbf{A}_{(i,:)})+\sum_{i}^{t}\frac{1}{2}\|(P(\mathbf{A})-\mathbf{I}_i)\mathbf{D}_i\|_F^2 \qquad (3)$$

where $\mathbf{I}_i$ is the i'th observed image of t total images and $\mathbf{D}_i$ is a diagonal matrix with known corrupt pixels set to 0 and all others set to w_i where w_i is a weight assigned to image i. We illustrate the SMTV process in Figure 1.

With this modification to the original TV restoration objective (2) we can now perform TV restoration from multiple source images. Furthermore through $\mathbf{D}_i$ we can easily ignore corrupt or undesirable pixels and simultaneously preference the influence of each image or even single pixels from each image. Where the type of corruption is known it is possible to use automatic methods to set $\mathbf{D}_i$. For example in satellite imagery there are many techniques such as FMASK [13] and ACCA [14], which can estimate cloud coverage at a pixel level. In other cases user interaction may be required to select regions of each image that are to be ignored.

IV. OPTIMISATION

To solve (3) we use LADMAP [15], which allows us to iteratively solve the objective. For the moment we consider only the anisotropic variant. See the following subsection for the changes needed when using the isotropic variant.

The anisotropic objective can be written as

$$\min_{\mathbf{A}\in B_{h,l}} \lambda\|\mathbf{AR}\|_1 + \sum_i^t \frac{1}{2}\|(P(\mathbf{A})-\mathbf{I}_i)\mathbf{D}_i\|_F^2 \tag{4}$$

where $\mathbf{R} \in \mathbb{R}^{N\times 2N}$ encodes the vertical and horizontal differences for each pixel i.e.

$$R_{ij} = \begin{cases} 1 & \text{if } i = \lceil \frac{j}{2} \rceil \\ -1 & \text{if } j \in \mathbb{O} \text{ and } i \in E_{\lceil \frac{j}{2} \rceil} \\ -1 & \text{if } j \notin \mathbb{O} \text{ and } i \in E_{\lceil \frac{j}{2} \rceil} \\ 0 & \text{otherwise.} \end{cases} \tag{5}$$

where $\mathbb{O}$ is the set of positive odd integers, E_i is a tuple containing the horizontal and vertical neighbour indices for pixel i and $\lceil x \rceil$ is the ceiling of x. Essentially $\mathbf{AR}$ is a matrix where the vertical and horizontal differences for each pixel alternate along the columns with the rows corresponding to each colour channel just as before.

First we introduce auxiliary variable $\mathbf{U}$

$$\min_{\mathbf{A}\in B_{h,l}} \lambda\|\mathbf{U}\|_1 + \sum_i^t \frac{1}{2}\|(P(\mathbf{A})-\mathbf{I}_i)\mathbf{D}_i\|_F^2 \tag{6}$$

$$\text{s.t.} \quad \mathbf{U} = \mathbf{AR} \tag{7}$$

then we can relax the constraints by constructing the Augmented Lagrangian form

$$\min_{\mathbf{A},\mathbf{U}} \mathcal{X}_{h,l}(\mathbf{A}) + \lambda\|\mathbf{U}\|_1 + \sum_i^t \frac{1}{2}\|(P(\mathbf{A})-\mathbf{I}_i)\mathbf{D}_i\|_F^2 + \langle \mathbf{Y}, \mathbf{U}-\mathbf{AR}\rangle + \frac{\mu}{2}\|\mathbf{U}-\mathbf{AR}\|_F^2 \tag{8}$$

where $\mathcal{X}_{h,l}(\mathbf{A})$ is the boundary indicator function and $\mathbf{Y}$ is a Lagrange multiplier. Then the process is to iterate over the following until convergence (see Algorithm 1 for full details):

1) Update $\mathbf{A}$

$$\min_{\mathbf{A}} \mathcal{X}_{h,l}(\mathbf{A}) + \sum_i^n \frac{1}{2}\|(P(\mathbf{A})-\mathbf{I}_i)\mathbf{D}_i\|_F^2 + \frac{\mu}{2}\|\mathbf{AR}-\mathbf{U}-\frac{1}{\mu}\mathbf{Y}\|_F^2$$

which we can rewrite as

$$\min_{\mathbf{A}} \mathcal{X}_{h,l}(\mathbf{A}) + \frac{\rho}{2}\|\mathbf{A}-(\mathbf{A}_k - \frac{1}{\rho}\partial G)\|_F^2$$

where $\partial G = \sum_i^n (P^*(P(\mathbf{A})-\mathbf{I}_i))\mathbf{D}_i + \mu(\mathbf{AR}-\mathbf{U}-\frac{1}{\mu}\mathbf{Y})\mathbf{R}^T$ and P^* is the adjoint of P. See [16], [17], [18] for detail. We have

$$A_{ij} = \begin{cases} l, & \overline{A}_{i,j} < l, \\ h, & \overline{A}_{i,j} > h, \\ \overline{A}_{i,j}, & \text{otherwise,} \end{cases} \tag{9}$$

where $\overline{\mathbf{A}} = \mathbf{A}_k - \frac{1}{\rho}\partial G$.

2) Update $\mathbf{U}$

$$\min_{\mathbf{U}} \lambda\|\mathbf{U}\|_1 + \langle \mathbf{Y}, \mathbf{U}-\mathbf{AR}\rangle + \frac{\mu}{2}\|\mathbf{U}-\mathbf{AR}\|_F^2$$

$$\min_{\mathbf{U}} \lambda\|\mathbf{U}\|_1 + \frac{\mu}{2}\|\mathbf{U}-\mathbf{AR}-\frac{1}{\mu}\mathbf{Y}\|_F^2$$

which is solved with the ℓ_1 thresholding operator [19], [20]

$$U_{ij} = \text{sign}(\overline{U}_{ij})\max(|\overline{U}_{ij}| - \frac{\lambda}{\mu}, 0) \tag{10}$$

where $\overline{\mathbf{U}} = \mathbf{AR} - \frac{1}{\mu}\mathbf{Y}$.

Algorithm 1 **Solving** (3) **by LADMAP**

Require: $\mathbf{P}^{N\times N}$, $\mathbf{I}_i^{d\times N}$, $\mathbf{D}_i^{N\times N}$, $\mathbf{R}$, λ, h, l, ϵ_1, ϵ_2,

1: Initialise: $\mathbf{A} = \mathbf{0}$, $\mathbf{U} = \mathbf{0}$, $\mathbf{Y} = \mathbf{0}$, $\mu = 0.1$, $\mu_{\max} = 10$, $\gamma^0 = 1.1$, $\rho = \|R\|_F$

2: **while** not converged **do**

3: Update $\mathbf{A}$ using (9)

4: Update $\mathbf{U}$ using (11) or (10)

5: Set q

$$q = \frac{\mu\sqrt{\rho}}{\|\mathbf{R}\|_F}\max(\|\mathbf{A}^{k+1}-\mathbf{A}\|_F, \|\mathbf{U}^{k+1}-\mathbf{U}\|_F)$$

6: Check stopping criteria

$$\frac{\|\mathbf{U}-\mathbf{AR}\|_F}{\|\mathbf{R}\|_F} < \epsilon_1, q < \epsilon_2$$

7: Update Lagrangian Multiplier

$$\mathbf{Y}^{k+1} = \mathbf{Y}^k + \mu^k(\mathbf{U}-\mathbf{AR})$$

8: Update γ

$$\gamma_1 = \begin{cases} \gamma^0 & \text{if } q < \epsilon_2 \\ 1 & \text{otherwise,} \end{cases}$$

9: Update μ

$$\mu^{k+1} = \min(\mu_{\max}, \gamma\mu^k)$$

10: **end while**

11: **return** $\mathbf{A}$

A. Isotropic Variant

The isotropic variant is solved in nearly the same way, with the only difference being the update rule for $\mathbf{U}$. The subproblem can be converted into the following proximal $\ell_{1,2}$ minimisation

$$\min_{\hat{\mathbf{U}}} \lambda\|\hat{\mathbf{U}}\|_{1,2} + \frac{\mu}{2}\|\hat{\mathbf{U}}-\mathbf{B}\|_F^2$$

where $\|\mathbf{U}\|_{1,2} = \sum_{i=1}^n \|\mathbf{U}_{(:,i)}\|_2$ and $\mathbf{B} = f(\mathbf{AR} - \frac{1}{\mu}\mathbf{Y})$ and $f(\mathbf{X}) : \mathbb{R}^{d\times 2N} \mapsto \mathbb{R}^{2\times dN}$ such that the two rows correspond to horizontal and vertical differences and columns to pixel indices repeated for each colour channel. The above problem has the following closed form solution [21]

$$\hat{\mathbf{U}}_{(:,i)} = \frac{\max(\|(\mathbf{B}_{(:,i)}\|_2 - \frac{\lambda}{\mu}, 0)}{\max(\epsilon, \|(\mathbf{B}_{(:,i)}\|_2)} \circ \mathbf{B}_{(:,i)} \tag{11}$$

where $\circ$ is the elementwise product.

Experiment	GAPG	SMTV Single Image	SMTV All Images
Noisy Images	0.01	0.05	0.1
Noisy and Corrupt Images	0.1	0.05	0.0001
Low-Light Photography	-	-	0.05
Haze and Cloud Removal	-	-	0.0001

TABLE I: Overview of λ value used for each experiment and algorithm.

V. Experimental Evaluation

In this section we present the results of SMTV in a number of image restoration tasks. First we consider a synthetic experiment for the recovery of noisy images. Second we consider the same task again with corrupt pixels. Third we demonstrate the application of SMTV to photography in low-light environments and last we show how SMTV can be used to restore satellite images which have cloud and haze occlusions.

In all experiments we use anisotropic TV. For evaluation of accuracy against ground truth images we use the Peak Signal to Noise Ratio (PSNR) metric, which is defined as

$$\text{PSNR} = 10 \log_{10} \left(\frac{s^2}{\frac{1}{mn} \sum_i^m \sum_j^n (X_{ij} - Y_{ij})^2} \right) \tag{12}$$

where $\mathbf{X}$ is the ground truth, $\mathbf{Y}$ is the estimation and s is the maximum possible value of elements in $\mathbf{X}$. Note that decreasing values of PSNR indicate increasing amounts of noise and this is a logarithmic scale.

MATLAB implementations of all tested algorithms were used on a machine with 3.4 Ghz i7 CPU and 32GB RAM. Values chosen for λ in each experiment can be found in Table I. In our experiments we compare SMTV against the state of the art algorithm for TV image restoration GAPG [12] and SMTV but on each image seperately. Full source code and all images can be found at http://github.com/sjtrny/smtv.

A. Noisy Image Restoration

In this subsection we present the results of restoring an image from multiple noisy observations of that image. We use four standard test images: barbara, cameraman, lena and peppers. For each image we make five copies which are corrupted with Gaussian noise i.e.

$$\mathbf{I}_i = \mathbf{A} + \mathbf{N}_i$$

where $\mathbf{A}$ is the ground truth image and $\mathbf{I}_i$, $\mathbf{N}_i$ are the ith observed image and Gaussian noise respectively. In this experiment the Gaussian noise has zero mean and standard deviation 10^{-1}.

Visual results for comparison can be found in Figure 2 and numeric results in Table II. Since we are generating multiple versions of this image we report the maximum PSNR and minimum running time for the single image denoising methods (GAPG and SMTV Single). For all images SMTV with all images clearly outperforms GAPG and Single Image SMTV in terms of accuracy. In comparison to the other methods SMTV is able to create very pleasing restorations. This is even more impressive considering when the severeness of the noise, which has been limited to a standard deviation of 10^{-3} in past work [12].

Image	Method	PSNR	Running Time (s)
Barbara	GAPG	24.88	**2.82**
	SMTV Single	25.57	3.12
	SMTV	**30.07**	4.52
Cameraman	GAPG	24.92	0.87
	SMTV Single	27.35	**0.77**
	SMTV	**31.89**	1.04
Lena	GAPG	28.41	**2.82**
	SMTV Single	29.01	3.03
	SMTV	**33.14**	4.72
Peppers	GAPG	25.95	0.85
	SMTV Single	28.06	**0.77**
	SMTV	**32.52**	1.11

TABLE II: Synthetic image denoising experiment: PSNR and running time values for all images.

Image	Method	PSNR	Running Time (s)
Barbara	GAPG	26.05	2.64
	SMTV Single	26.27	3.71
	SMTV	**46.44**	**2.35**
Cameraman	GAPG	26.07	0.85
	SMTV Single	29.17	0.92
	SMTV	**46.47**	**0.60**
Lena	GAPG	32.07	2.66
	SMTV Single	32.04	3.69
	SMTV	**46.46**	**2.56**
Peppers	GAPG	27.56	0.85
	SMTV Single	31.28	0.96
	SMTV	**46.47**	**0.60**

TABLE III: Synthetic denoising and corruption experiment: PSNR and running time values for all images.

B. Noisy and Corrupt Image Restoration

Similar to the previous subsection we subject our ground truth to images to noise however this time we set random pixels to 0 to simulate corruption or occlusion. Therefore the image formation model is of the form

$$\mathbf{I}_i = \mathcal{N}_\Omega(\mathbf{A} + \mathbf{N}_i)$$

where $\mathcal{N}_\Omega$ is defined as

$$\mathcal{N}_\Omega(A_{ij}) = \begin{cases} 0, & (i,j) \in \Omega, \\ A_{ij}, & \text{otherwise.} \end{cases} \tag{13}$$

and Ω is the set of corrupted indices. In this experiment the Gaussian noise has zero mean and standard deviation 10^{-2} and Ω is randomly generated and uniformly distributed.

Again we make five corrupted copies of each image and report the best recovery results for the single image algorithms. Visual results for comparison can be found in Figure 3 and numeric results in Table III. Overall SMTV dominates clearly in terms of reconstruction accuracy. Furthermore SMTV using all five observations is actually faster than GAPG and SMTV on single images. We believe this is due to SMTV with multiple observations converging faster since each observation acts in a way that averages or smoothes out the image further and reduces the total variation of the combined image.

C. Improved Low-Light Photography

In this section we present visual results of applying SMTV to low-light photography situations. Photographs taken in low-light environments typically contain a significant and noticeable amount of noise. This is due to the reduction in light

(a) Original Images

(b) One of the Five Noisy Images

(c) GAPG (Best of Five)

(d) SMTV

Fig. 2: Visual results for the synthetic image denoising experiment.

falling upon the imaging sensor. To circumvent this photographers often take longer exposures however this comes with it's own set of problems. For example moving scenery (such as a person or animal walking across the scene) will affect the entire image, whereas with SMTV we can discard pictures containing these defects. Or the camera sensor may suffer from problems such as hot pixels (pixels that give defective readings as they heat up over time) if the exposure setting is too long. For this experiment images of static scenes were taken indoors under standard room lighting conditions. The images were taken on an Apple iPhone 5s, which has an 8 megapixel camera and $f/22$ aperture. The iPhone's flash was disabled and the exposure time was set to $\frac{1}{15}$ of a second. Five images were taken of each scene using the iPhone's burst photo functionality. We then performed SMTV on a 500×500 pixel subsection for each scene. The images used for each scene and the corresponding recovered image can be found in Figure 4. Note that in the zoomed sections of the source images (first and second column) there is a lot of noise and edges are very blurry. SMTV removes all perceptible noise at this magnification level and enhances sharpness and clarity. For example the edges of individual blocks on the alien toy are sharpened and each block smoothed. In the mandarin image the dimples on the skin are enhanced and less blurry and the text on the label is clearer. Even the small boundary between the text shadow effect and the white bubble outline on the "2.P.H." logo is visible.

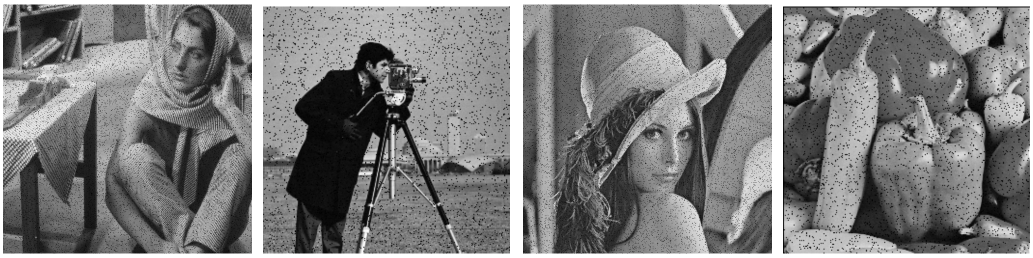

(a) One of the Five Noisy and Corrupt Images

(b) GAPG (Best of Five)

(c) SMTV

Fig. 3: Visual results for the synthetic image and corruption experiment.

D. Cloud Removal and Shadow Correction in Satellite Imagery

In this section we present visual results of applying SMTV to satellite imagery of Earth, which suffers from cloud occlusions of the Earth's surface and shadows from those clouds. Two seasonal datasets were generated from LANDSAT images of a section of Australian land. Each dataset consists of four images using the visible spectrum bands of the LANDSAT images at a total resolution of 4000×4000 pixels. The source images are freely available from EarthExplorer [22]. To obtain the cloud and shadow mask (corrupt pixels) for each image we used the FMASK algorithm [13]. Since the masks produced by FMASK are rough we then refined them using Guided Image Filtering [18]. Finally SMTV was applied to each dataset to restore images free of the aforementioned defects. See Figures 5 and 6 for results of both datasets. Since we cannot reproduce the full high resolution images here we select a 1000×1000 cropped section to show for each set. The average running time of SMTV for an image of this size was around 20 seconds.

VI. Conclusion

We proposed and evaluated a new method to automatically fuse multiple noisy and corrupted images into a single denoised image. SMTV has shown remarkable quantitive and qualitative improvement over traditional single image methods in denoising and image restoration tasks. Furthermore it has opened up new applications such as the improvement of photographs taken in low-light conditions and cloud affected satellite imagery.

Acknowledgment

This research was supported by Australian Research Council Discovery Projects funding scheme (project DP140102270) and the National Natural Science Foundation of China (Project 41371362).

References

[1] A. Nguyen, J. Yosinski, and J. Clune, "Deep neural networks are easily fooled: High confidence predictions for unrecognizable images," *CoRR*, vol. abs/1412.1897, 2014.

[2] C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," in *Sixth International Conference on Computer Vision*. IEEE, 1998, pp. 839–846.

[3] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D: Nonlinear Phenomena*, vol. 60, no. 1, pp. 259–268, 1992.

[4] L. Yin, R. Yang, M. Gabbouj, and Y. Neuvo, "Weighted median filters: a tutorial," *Circuits and Systems II: Analog and Digital Signal Processing, IEEE Transactions on*, vol. 43, no. 3, pp. 157–192, 1996.

[5] A. Chambolle, R. A. De Vore, N.-Y. Lee, and B. J. Lucier, "Nonlinear wavelet image processing: variational problems, compression, and noise removal through wavelet shrinkage," *Image Processing, IEEE Transactions on*, vol. 7, no. 3, pp. 319–335, 1998.

(a) Input Images (b) Zoomed Section of Input Image 1 (c) SMTV Restored Image

(d) Input Images (e) Zoomed Section of Input Image 1 (f) SMTV Restored Image

Fig. 4: Improvement of photographs taken under low-light conditions. This figure is best viewed in the electronic version. The left hand column shows the five source images, the middle column shows a zoomed crop of one of the source images and the right hand column shows the SMTV restored image. Note the removal of noise and increase in clarity of fine image details.

[6] J. Portilla, V. Strela, M. J. Wainwright, and E. P. Simoncelli, "Image denoising using scale mixtures of gaussians in the wavelet domain," *Image Processing, IEEE Transactions on*, vol. 12, no. 11, pp. 1338–1351, 2003.

[7] A. Buades, B. Coll, and J.-M. Morel, "A non-local algorithm for image denoising," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 2. IEEE, 2005, pp. 60–65.

[8] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-d transform-domain collaborative filtering," *Image Processing, IEEE Transactions on*, vol. 16, no. 8, pp. 2080–2095, 2007.

[9] S. Zimmer, S. Didas, and J. Weickert, "A rotationally invariant block matching strategy improving image denoising with non-local means," in *Proc. 2008 International Workshop on Local and Non-Local Approximation in Image Processing*, 2008, pp. 135–142.

[10] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM Journal on Imaging Sciences*, vol. 2, no. 1, pp. 183–202, 2009.

[11] J. M. Bioucas-Dias and M. A. Figueiredo, "A new twist: two-step iterative shrinkage/thresholding algorithms for image restoration," *Image Processing, IEEE Transactions on*, vol. 16, no. 12, pp. 2992–3004, 2007.

[12] W. Zuo and Z. Lin, "A generalized accelerated proximal gradient approach for total-variation-based image restoration¡? query valign="-12pt" copyedited file?¿," *Image Processing, IEEE Transactions on*,

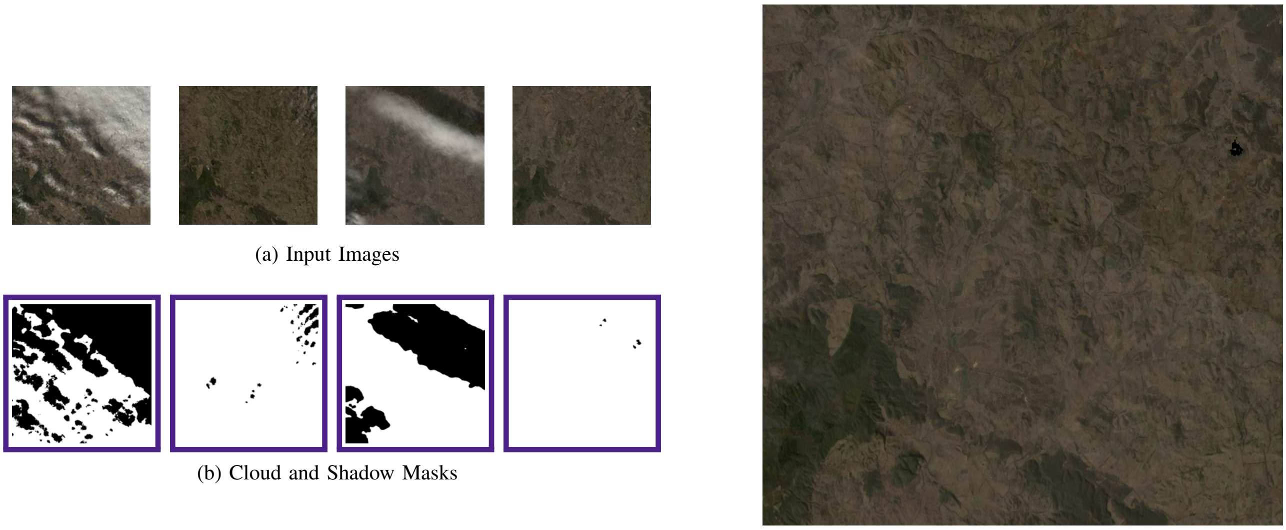

(a) Input Images

(b) Cloud and Shadow Masks

(c) Final Output

Fig. 5: Results of SMTV upon the LANDSAT 1 dataset. The cloud and shadow masks are combined with the input images in SMTV to produce a final image free of defects.

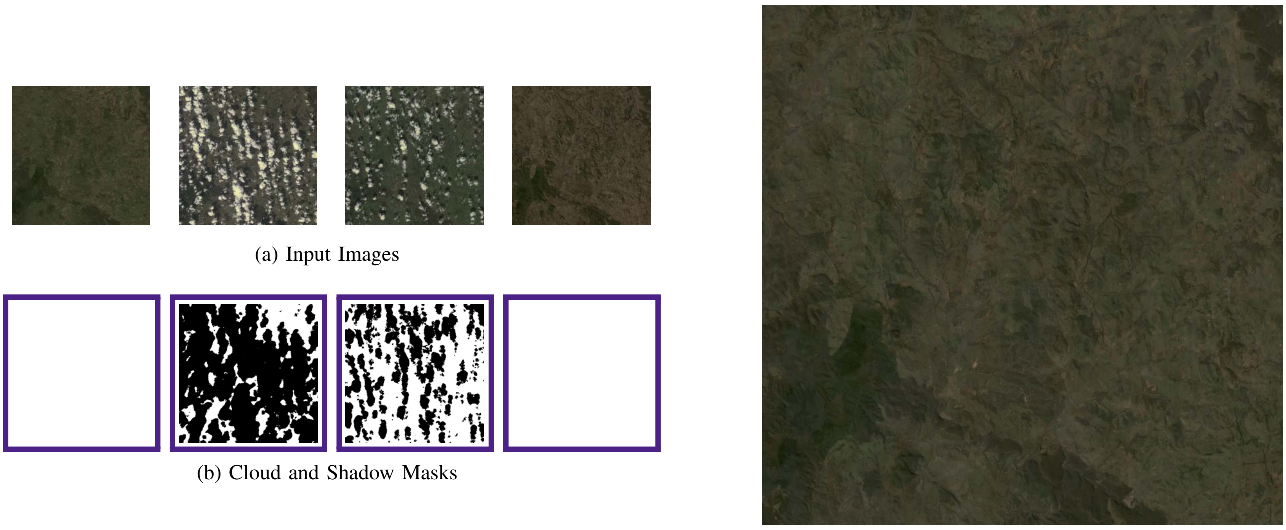

(a) Input Images

(b) Cloud and Shadow Masks

(c) Final Output

Fig. 6: Results of SMTV upon the LANDSAT 2 dataset. The cloud and shadow masks are combined with the input images in SMTV to produce a final image free of defects.

vol. 20, no. 10, pp. 2748–2759, 2011.

[13] Z. Zhu and C. E. Woodcock, "Object-based cloud and cloud shadow detection in landsat imagery," *Remote Sensing of Environment*, vol. 118, pp. 83–94, 2012.

[14] R. R. Irish, J. L. Barker, S. N. Goward, and T. Arvidson, "Characterization of the landsat-7 etm+ automated cloud-cover assessment (acca) algorithm," *Photogrammetric Engineering & Remote Sensing*, vol. 72, no. 10, pp. 1179–1188, 2006.

[15] Z. Lin, R. Liu, and Z. Su, "Linearized alternating direction method with adaptive penalty for low-rank representation," in *Advances in Neural Information Processing Systems*, 2011, pp. 612–620.

[16] S. Ji and J. Ye, "An accelerated gradient method for trace norm minimization," in *Proceedings of the 26th Annual International Conference on Machine Learning*. ACM, 2009, pp. 457–464.

[17] Y. Nesterov, "Smooth minimization of non-smooth functions," *Mathematical Programming*, vol. 103, no. 1, pp. 127–152, 2005.

[18] Y. Nesterov *et al.*, "Gradient methods for minimizing composite objective function," 2007.

[19] J. Liu and J. Ye, "Efficient l1/lq norm regularization," *arXiv preprint arXiv:1009.4766*, 2010.

[20] F. Bach, R. Jenatton, J. Mairal, and G. Obozinski, "Convex optimization with sparsity-inducing norms," *Optimization for Machine Learning*, pp. 19–53, 2011.

[21] G. Liu, Z. Lin, and Y. Yu, "Robust subspace segmentation by low-rank representation," in *International Conference on Machine Learning*, 2010, pp. 663–670.

[22] U. S. G. Survey. Earthexplorer. [Online]. Available: http://earthexplorer.usgs.gov/