Anantrasirichai, N., Achim, A. M., Kingsbury, N., & Bull, D. R. (2013). Atmospheric Turbulence Mitigation using Complex Wavelet-based Fusion. *IEEE Transactions on Image Processing*, 22(6), 2398-2408. <https://doi.org/10.1109/TIP.2013.2249078>

Peer reviewed version

Link to published version (if available):
[10.1109/TIP.2013.2249078](https://doi.org/10.1109/TIP.2013.2249078)

[Link to publication record in Explore Bristol Research](#)
PDF-document

(c) 2013 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other users, including reprinting/ republishing this material for advertising or promotional purposes, creating new collective works for resale or redistribution to servers or lists, or reuse of any copyrighted components of this work in other works.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Atmospheric Turbulence Mitigation using Complex Wavelet-based Fusion

N. Anantrasirichai, *Member, IEEE*, Alin Achim, *Senior Member, IEEE*, Nick Kingsbury, *Fellow, IEEE*, and David Bull, *Fellow, IEEE*

Abstract—Restoring a scene distorted by atmospheric turbulence is a challenging problem in video surveillance. The effect, caused by random, spatially varying, perturbations, makes a model-based solution difficult and in most cases, impractical. In this paper, we propose a novel method for mitigating the effects of atmospheric distortion on observed images, particularly airborne turbulence which can severely degrade a region of interest (ROI). In order to extract accurate detail about objects behind the distorting layer, a simple and efficient frame selection method is proposed to select informative ROIs only from good-quality frames. The ROIs in each frame are then registered to further reduce offsets and distortions. We solve the space-varying distortion problem using region-level fusion based on the Dual Tree Complex Wavelet Transform (DT-CWT). Finally, contrast enhancement is applied. We further propose a learning-based metric specifically for image quality assessment in the presence of atmospheric distortion. This is capable of estimating quality in both full- and no-reference scenarios. The proposed method is shown to significantly outperform existing methods, providing enhanced situational awareness in a range of surveillance scenarios.

Index Terms—image restoration, region-level fusion, DT-CWT, quality metrics.

I. INTRODUCTION

VARIOUS types of atmospheric distortion can influence the visual quality of video signals during acquisition. Typical distortions include fog or haze which reduce contrast, and atmospheric turbulence due to temperature variations or aerosols. In situations when the ground is hotter than the air above it, the air is heated and begins to form horizontal layers. When the temperature difference between the ground and the air increases, the thickness of each layer decreases and the air layers move upwards rapidly, leading to faster and greater micro-scale changes in the air's refractive index. This effect is observed as a change in the interference pattern of the light refraction. In strong turbulence, not only scintillation, which produces small-scale intensity fluctuations in the scene [1] and blurring effects are present in the video imagery, but also a shearing effect occurs and is perceived as different parts of

Fig. 1. Block diagram of image restoration for atmospheric turbulence

objects moving in different directions [2]. Examples of this effect are found at locations such as hot roads and deserts, as well as in the proximity of hot man-made objects such as aircraft jet exhausts. This is particularly a problem close to the ground in hot environments and can combine with other detrimental effects in long range surveillance applications, where images can be acquired over distances up to 20km [3].

Turbulence effects in the acquired imagery make it extremely difficult to interpret information behind the distorted layer. Hence, there has been significant research activity attempting to faithfully reconstruct this useful information using various methods. In practice, the perfect solution is however impossible, since the problem is ill-posed, despite being simply expressed with a matrix–vector multiplication as in Eq. 1.

$$\mathbf{I}_{obv} = \mathbf{D}\mathbf{I}_{idl} + \varepsilon \quad (1)$$

Here $\mathbf{I}_{obv}$ and $\mathbf{I}_{idl}$ are vectors containing the observed and ideal images, respectively. Matrix $\mathbf{D}$ represents geometric distortion and blur, while ε represents noise. Various approaches have attempted to solve this problem by modelling it as a point spread function (PSF), in which $\mathbf{D}$ is considered as a convolution matrix, and then employing deconvolution with an iterative process to estimate $\mathbf{I}_{idl}$. For the atmospheric distortion case, the PSF is generally unknown, so blind deconvolution is employed [4]–[6]. However, the results still exhibit artefacts since the PSF is usually assumed to be space-invariant.

It is obvious that removal of the visible spatio-temporal distortions is not possible with a single image. Hence all methods utilise a set of images to construct one enhanced image. Current multi-frame methods that address this problem are illustrated in Fig. 1, where most approaches employ all functions or a subset of them. The restoration process can be described by two main routes through the diagram. The first

Copyright (c) 2013 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

This work was supported by General Dynamics United Kingdom Limited. N. Anantrasirichai, Alin Achim and David Bull are with Visual Information Laboratory, University of Bristol, Bristol BS8 1UB, U.K. e-mail: (n.anantrasirichai@bris.ac.uk; alin.achim@bristol.ac.uk; dave.bull@bris.ac.uk).

N. Kingsbury is with the Department of Engineering, Signal Processing and Communications Laboratory, University of Cambridge, Cambridge CB2 1TN, U.K. (e-mail: ngk@eng.cam.ac.uk).

(green dashed line) employs an image registration technique with deformation estimation [7]–[12]. This process attempts to align objects temporally to solve for small movements of the camera and temporal variations due to atmospheric refraction. The image fusion block may optionally be employed (blue line) in order to combine several aligned images. Then, a deblurring process is applied to the combined image (this in itself is challenging since the blur is space-varying).

The other route (red solid line) employs image selection and fusion, often referred to as ‘*lucky region*’ techniques [13]–[17]. Those regions of the input frames that have the best quality in the temporal direction are employed. They are selected using an image quality metric, which is normally applied in the spatial frequency domain, to extract the best quality frames (minimally distorted and least blurred). These are then combined in an intelligent manner. Recently, this method has been improved by applying image alignment to those lucky regions [18]. Again, a deblurring process can be applied as a final step.

Effective mitigation of atmospheric turbulence is thus an important yet challenging problem. Model-based solutions are impractical and blind deconvolution methods suffer from spatial and temporal variation due to PSF. Furthermore, conventional registration methods are ineffective for large distortion and are also time-consuming. Finally conventional fusion methods require a large number of frames in order to select lucky regions. In this paper we introduce a new approach that overcomes these problems. Fusion is performed in the Dual Tree Complex Wavelet Transform (DT-CWT) domain since this provides near shift-invariance and good directional selectivity [19]. We also propose novel frame selection and ROI alignment methods for pre-processing region of interest (ROI) since this will frequently exhibit significant offsets and distortions between frames. Contrast enhancement is then used as the final step.

Our proposed algorithm is tested with real distorted sequences as well as with simulated sequences. The latter case includes heat distortion generated from gas burners and hence ground truth information is available. We also investigate a quality metric that is suitable for measuring restored image quality for atmospherically distorted content where generally the ground truth is not available. Existing no-reference assessment (NR) methods [20]–[23] are tested with our simulated sequences. The results however do not show high correlation with the objective results. Therefore we introduce a new NR measure based on machine learning.

The remaining part of this paper is organised as follows. The proposed scheme for mitigating atmospheric distortion is described in detail in Section II. A test methodology for objective assessment is introduced in Section III. The performance of the method is evaluated on a set of images and is compared with other techniques in Section IV. Finally, Section V presents the conclusions of the paper.

II. PROPOSED MITIGATION SCHEME

We propose a new fusion method for reducing the effects of atmospheric turbulence as depicted in Fig. 2. First, before

Fig. 2. Block diagram of the proposed method

Fig. 3. ROI alignment technique

applying fusion, a subset of selected images or ROIs must be aligned. Here we introduce a new alignment approach for distorted images. As randomly distorted images do not provide identical features, we cannot use conventional methods to find matching features. Instead, we apply a morphological image processing technique, namely erosion, to the ROI (or whole image) based only on the most informative frames. These are selected using a quality metric based on sharpness, intensity similarity and ROI size. Then, non-rigid image registration is applied.

We then employ a region-based scheme to perform fusion at the feature level. This has advantages over pixel-based processing since more intelligent semantic fusion rules can be considered based on actual features in the image, rather than on single or arbitrary groups of pixels. The fusion is performed in the Dual Tree Complex Wavelet Transform (DT-CWT) which employs two different real discrete wavelet transforms (DWT) to provide the real and imaginary parts of the CWT. Two fully decimated trees are produced, one for the odd samples and one for the even samples generated at the first level. This increases directional selectivity over the DWT and is able to distinguish between positive and negative orientations giving six distinct sub-bands at each level, corresponding to $\pm 15^\circ$, $\pm 45^\circ$, $\pm 75^\circ$. Additionally, the phase of a DT-CWT coefficient is robust to noise and temporal intensity variations thereby providing an efficient tool for removing distorting ripples. Finally, the DT-CWT is near-shift invariant - an important property for this application. After fusion, the effect of haze is reduced using locally-adaptive histogram equalisation.

For convenience, we refer to this algorithm as CLEAR (Complex waveLEt fusion for Atmospheric tuRbulence). Details of each step in our algorithm are described below.

A. ROI Alignment

Capturing video in the presence of atmospheric turbulence, especially when using high magnification lenses, may cause the ROI in each frame to become misaligned. The displacement between the distorted objects in the successive frames may be too large for conventional image registration,

Fig. 4. Average frame of *Number Plate* sequence by (a) using 200 original frames, (b) using 50 frames most similar (lowest MSE) to (a), and (c) using 200 frames after applying ROI alignment approach.

using non-rigid deformation, to cope with. Equally, matching using feature detection is not suitable since strong gradients within each frame are randomly distorted spatially. Hence, an approach using morphological image processing is proposed. The ROI (or ROIs) is manually marked in the first frame. Then the histogram, generated from the selected ROI and the surrounding area, is employed to find an Otsu threshold [24], which is used to convert the image to a binary map. An erosion process is then applied and the areas connected to the edge of the sub-image are removed. This step is performed iteratively until the area near the ROI is isolated. The same Otsu threshold with the same number of iterations is employed in other frames. The centre position of each mask is then computed. If there is more than one isolated area, the area closest in size and location to the ROI in the first frame is used. Finally, the centre of the mask in each frame is utilised to shift the ROI and align it across the set of frames (Fig. 3). Note that the frames with incorrectly detected ROIs will be removed in the frame selection process (section II-B). These frames are generally significantly different from others.

Fig. 4 demonstrates the improvement due to the proposed ROI alignment approach. The left image represents the average frame of the whole *Number Plate* sequence (see B4 in Table II and Fig. 9) and it reveals high variation due to camera movement which significantly impacts on image quality more than the turbulence. A simple alignment method might choose a subset of the original frames where the inter-frame distance is not too large. Fig. 4 (b) shows an example of this where, based on mean square error (MSE), 50 original frames with the smallest difference from the average image are selected. The alignment of these frames is consequently improved as shown by their average image in Fig. 4 (b). However, this truncation may lead to loss of some useful information. In contrast, the proposed ROI alignment approach keeps all frames so that an intelligent approach can be employed later to select the best subset. The average frame of the whole sequence of our method is shown in Fig. 4 (c), which is clearer and sharper than the other two images (Fig. 4 (a) and (b)). This means that the aligned frames are more stable thereby producing better image registration results.

B. Frame Selection

In CLEAR, not all frames in the sequence are used to restore the image since the low quality frames (e.g. the very blurred ones) would possibly degrade the fused result. A subset of

images are carefully selected using three factors: sharpness, intensity similarity and detected ROI size.

- **Sharpness** G_n is one of the most important image quality factors since it determines the amount of detail an image can convey. Here, the sharpness parameter G_n is computed from the summation of the highpass coefficient magnitudes. Intensity gradients can also be used as the result is insignificantly different from highpass coefficients.

- **Intensity similarity** S_n is employed to remove outliers. This operates under the assumption that most frames in the sequence contain fairly similar areas. Frames with significantly different content to others are likely to be greatly distorted. To compute S_n , the average frame of the whole sequence is used as a reference for calculating the mean square error (MSE) for frame n . Then MSE^{-1} represents the similarity of each frame. It should be noted that this approach is not robust to illumination changes.

- **Detected ROI size** A_n is the total number of pixels contained in the ROI. This is used because, from observation, larger ROIs are likely to contain more useful information.

The cost function C_n for frame n is computed using Eq. 2.

$$C_n = \frac{w_G G_n}{\lambda_G + |G_n|} + \frac{w_S S_n}{\lambda_S + |S_n|} + \frac{w_A A_n}{\lambda_A + |A_n|} \quad (2)$$

where w_k and λ_k are the weight and slope control of the factor $k \in \{G, S, A\}$, respectively. The sigmoid function is used here to prevent one factor dominating the others, e.g. a blocking artefact may cause significantly high values of sharpness, yet this frame should probably not be included in the selected data set. The λ_k is set to equal the mean of factor k so that at the mean value, its cost value is 0.5. The cost C_n is ranked from high to low. The Otsu method can then be applied to find how many frames should be included in the selected set.

C. Image Registration

Registration of non-rigid bodies using the phase-shift properties of the DT-CWT, as proposed in [25], is employed. This algorithm is based on phase-based multidimensional volume registration, which is robust to noise and temporal intensity variations. Motion estimation is performed iteratively, firstly by using coarser level complex coefficients to determine large motion components and then by employing finer level coefficients to refine the motion field. Fig. 5 shows an improvement in temporal direction (z) of the *Number Plate* sequence after applying the proposed ROI alignment and image registration.

D. Image Fusion

Due to its shift invariance, orientation selectivity and multi-scale properties, the DT-CWT is widely used in image fusion where useful information from a number of source images are selected and combined into a new image [26]–[28]. We employ a region-based scheme in the DT-CWT domain to implement image fusion at the feature level. In general, region-based fusion methods start the process by segmenting N images individually or jointly [26]–[29]. The segmentation map S_n of each image is then down sampled by 2 to give a

Fig. 5. The yz planes at column $x=160$ of each frame of *Number Plate*. Left: Distorted sequence corresponding to Fig. 4 (a). Top right: Aligned sequence corresponding to Fig. 4 (c). Bottom right: Registered sequence

Fig. 6. Region-based image fusion process in CLEAR.

decimated segmentation map $S_n^{\theta,l}, n \in N$ of level l and sub-band θ of the DT-CWT representation, where $\theta \in (1, \dots, 6)$. If $R_n = \{r_{n,1}, r_{n,2}, \dots, r_{n,T_n}\}$, is a list of all T_n regions of image n , a multi-resolution priority map P_n is generated as

$$P_n = \{p_{n,r_{n,1}}, p_{n,r_{n,2}}, \dots, p_{n,r_{n,T_n}}\} \quad (3)$$

for each region in each image n . Regions are then either chosen or discarded based on this priority. If $S_i \neq S_j$, a segmentation map in the fused image, S_F , is created such that $S_F = S_1 \cup S_2 \cup \dots \cup S_N$. Thus, where two regions $r_{i,p}$ and $r_{j,q}$ from image i and j overlap, both will be split into two regions, each with the same priority as the original.

The fusion rule, ϕ , can be defined to select the region with the maximum priority [26] or could employ a weighted average of these regions [28] to yield the wavelet coefficients of the fused image. Finally, the fusion image is obtained by performing the inverse transform on the fused wavelet coefficients.

In this paper, we have adapted the fusion technique explained above to specifically address the air-turbulence problem. The proposed region-based image fusion employed in CLEAR is illustrated in Fig. 6. For image segmentation, we employ the joint segmentation approach of O'Callaghan and Bull's joint morphological spectral unsupervised approach with a multiscale watershed segmentation from [30] to divide each image into similar regions, R_n . The lowpass DT-CWT coefficients of the fused image are simply constructed from the average of the lowpass values of all registered images, while the highpass (detail) coefficients are selected according to the priority P indicating the importance of each region. Here we employ the average magnitude of highpass coefficients in each

Fig. 7. Left: Region-based fusion result without the mask. Middle: Mask $B^{\theta,l}$. Right: Enhanced result with the mask

region, since wavelet coefficients having large absolute values contain information on the salient features of an image such as lines and texture. To produce sharper results compared to the results in [26]–[28], we operate on each sub-band separately. The priority P of region $r_n^{\theta} \in R_n$ in image n is computed with the highpass coefficients $d_n^{\theta,l}(x, y)$ of level l and sub-band θ as shown in Eq. 4, where $|r_n^{\theta}|$ is the size of such area used for normalisation. The fusion rule, ϕ , selects the region with maximum priority to construct the fused image.

$$P(r_n^{\theta}) = \frac{1}{|r_n^{\theta}|} \sum_{\forall l, (x, y) \in r_n^{\theta}} |d_n^{\theta,l}(x, y)| \quad (4)$$

The air-turbulence scenario differs from other image-fusion problems as the segmentation boundaries which separate inhomogeneous regions vary significantly from frame to frame (due to turbulence distortion). To provide the sharpest and most temporally consistent boundaries for each region, we use the maximum of DT-CWT coefficient magnitudes over all frames instead of selecting only one region based on $P(r_n^{\theta})$. To each boundary map $B^{\theta,l}$ (constructed from the multiscale watershed segmentation approach for each subband θ at level l), the dilation operation with a size of 1 pixel is applied. A 2D averaging filter is then applied to $B^{\theta,l}$ to prevent discontinuity after combining neighbouring areas. The DT-CWT coefficients, $d^{\theta,l}$, of the fused image can be written as in Eq. 5.

$$d^{\theta,l} = (1 - B^{\theta,l}) \sum_R \phi(d_1^{\theta,l}, d_2^{\theta,l}, \dots, d_N^{\theta,l}) + B^{\theta,l} \max(d_1^{\theta,l}, d_2^{\theta,l}, \dots, d_N^{\theta,l}) \quad (5)$$

The example of $B^{\theta,l}$ is illustrated in Fig. 7 (Middle). The improvement can be seen by comparing images in Fig. 7 (Left) and Fig. 7 (Right). The enhanced result shows sharper edges without boosting noise in homogeneous areas.

To reduce the distortion due to edge undulation, the phase of the complex wavelet coefficients plays an important role since it corresponds to the precise location of directional features in its support regions. Hence, the DT-CWT coefficients, $d^{\theta,l}(x, y)$, of the fused image are adjusted with a unit vector representing the average phase from all N frames used in the fusion process as shown Eq. 6.

$$\tilde{d}^{\theta,l}(x, y) = \frac{\sum_{n=1}^N d_n^{\theta,l}(x, y)}{\left| \sum_{n=1}^N d_n^{\theta,l}(x, y) \right|} |d^{\theta,l}(x, y)| \quad (6)$$

Although the turbulent motion also depends on wind velocity and direction, a simplifying assumption can be made, in that pixels deviate from their actual positions with approximately

Fig. 8. Fusion technique with phase adjustment shows improved result. Left: Fused image without phase adjustment. Right: Fused image with phase adjustment.

zero mean and with a quasi-periodic motion [9]. In this case, the average phase can be used. The improvement of the fusion result can be seen in Fig. 8. Clearly, the image with adjusted phases produces straighter lines which, in this case, are obviously closer to the truth.

To reduce noise, we apply a shrinkage function, A_s , derived as Maximum A Posteriori (MAP) estimators as in [31]. Then, if sharpening is required, a gain $A_g > 1$ can be applied to boost highpass-coefficient magnitudes. However, this simple technique may amplify any remaining noise. Therefore we enhance only those high-pass coefficients where they form regions of large magnitude. We create the binary map $Q^{\theta,l}$ for each subband, in which $Q^{\theta,l} = 1$ if $|d^{\theta,l}| > \tau$, where τ is a predefined threshold. Isolated pixels are subsequently removed from $Q^{\theta,l}$. The modified highpass coefficients are finally rewritten as Eq. 7.

$$d^{\theta,l} = (A_g^{\theta,l} Q^{\theta,l} + (1 - Q^{\theta,l})) A_s^{\theta,l} \tilde{d}^{\theta,l} \quad (7)$$

E. Post-Processing

1) *Contrast enhancement*: In many cases, atmospherically degraded images also suffer from poor contrast due to severe haze or fog. In such cases, pre- or post-processing is needed to improve image quality. Numerous techniques have been proposed for haze reduction using single images [32], [33]. Here we employ a simple and fast method using contrast limited adaptive histogram equalisation (CLAHE) [34]. The method enhances intensity locally, so it is suitable for applications which consider the ROI and its information content.

2) *Other possible enhancements*: Generally the embedded parameter A_g in our approach produces sharp results; however, in cases which are out-of-focus or which lack a ‘lucky region’, post-processing may be required to further sharpen the images. A number of sharpening methods exist, such as [35], [36]. However, if the constituent images are very poor, it is almost impossible to obtain a sharp result. Moreover, it may exhibit a halo effect due to over sharpening.

III. QUALITY ASSESSMENT

Image quality assessment is used to measure perceived image degradation, typically compared to an ideal or perfect image. This is important when assessing the performance of individual systems or for comparing different solutions. Image quality metrics can be classified according to the availability

of a reference (distortion-free) image, with which the distorted image is to be compared. Most existing approaches are classed as full-reference (FR), meaning that a complete reference image is available. Example FR methods include Peak Signal to Noise Ratio (PSNR), Multiscale Structural Similarity (MS-SSIM) [37], Visual Signal to Noise Ratio (VSNR) [38] and Perception-based Image Model (PIM) [39]. These metrics are employed for evaluating the performance of the proposed method in Section IV-A and IV-B4.

When a reference is not available, as is often the case for heat haze reduction, quality assessment becomes challenging, and is referred to as no-reference (NR) or blind quality assessment. This is described in the following section.

A. No-Reference Image Quality Assessment

The JPEG quality score was one of the first NR quality assessment methods introduced. It attempts to align image quality with HVS perception by characterising blockiness and blurring [20]. Subsequently, the JPEG2000 (JP2K) score was proposed for blind assessment of images compressed by wavelet based coding [21]. The Anisotropic Quality Index (AQI) is another NR metric based on measuring the variance of the expected entropy of a given image in a set of predefined directions [22]. Recently, the combination of five distortion types, namely JPEG, JPEG2000, white noise, Gaussian blur and fast fading, were used in the Blind Image Quality Index (BIQI) [23].

B. Proposed NR Method for Atmospheric Distortion

The methods described in section III-A do not work well with atmospheric distortion, since they are usually based on prior knowledge of the distortion characteristics and none are derived from spatially varying distortions. In this paper, we therefore introduce a new blind image quality assessment metric specifically for this scenario. We employ support vector regression (SVR) [40] to model and predict image quality scores using the features listed in Table I. There are three groups of features:

- **Individual scale** The magnitude of highpass coefficients relate to details and sharpness of the image, while the phase can be linked to edge information. We therefore employ the mean and variance of both values to compute the feature vectors at each scale level. We decompose the image into 3 levels using the DT-CWT.

- **Inter-scale** Weighted mean and variance at level l are computed using the magnitudes of the next coarse level to calculate a weight as shown in Eq. 8. Here $\hat{d}^{\theta,l} \in \hat{D}$ is the upsampled version of $d^{\theta,l}$ by 2 so that $\hat{d}^{\theta,l}$ has the same size as $d^{\theta,l-1}$. The weighted mean is computed as shown in Eq. 9, where $\kappa^{\theta,l}$ is the magnitude or the phase of complex coefficient $d^{\theta,l}$. The weighted variance is calculated from all $\kappa^{\theta,l}(x,y)$ that have $w^{\theta,l}(x,y) > 0.1$. This follows the assumption that the coefficients adjacent to strong edges (high $|d^{\theta,l}|$) are more important, since the distortions in those areas affect human perception more than others. Moreover, moderate atmospheric turbulence consistently manifests itself as clearly visible ripples along object edges.

TABLE I
FEATURES FOR MODELLING THE QUALITY OF THE IMAGES DISTORTED BY
ATMOSPHERIC TURBULENCE

Features	# dimensions
Individual scale	(3 scales)
Mean and variance of magnitude of CWT coefficients	6
Mean and variance of phase of CWT coefficients	6
Inter-scale	
Weighted mean and variance of magnitude	4
Weighted mean and variance of phase	4
Entropy of phase around high magnitude of coarser scale	2
Existing NR methods	
JPEG, JP2K, AQI, BIQI scores	4

$$w^{\theta,l}(x, y) = \frac{|\hat{d}^{\theta,l+1}(x, y)|}{\sum_{\forall \theta, (x, y)} |\hat{d}^{\theta,l+1}(x, y)|} \quad (8)$$

$$S_{\kappa}^l = \sum_{\forall \theta, (x, y)} w^{\theta,l}(x, y) \cdot \kappa^{\theta,l}(x, y) \quad (9)$$

We also utilise entropy to measure the randomness of the phase in the vicinity of strong edges, since a higher phase randomness can imply more edge-waviness. We employ the weighted average formula in Eq. 8 and 9 to compute entropy features. Here $\kappa^{\theta,l}(x, y)$ represents local entropy at position (x, y) calculated in a 5-by-5 neighbourhood.

- **Existing NR methods** Existing NR scores are also included in the regression to model image quality, since these have been developed for, and generally work well for, general natural images.

The SVR parameters are selected by cross validation. As the number of features is not large, principal component analysis is not required for dimensionality reduction. Instead, to find the best features for regression, different combinations of features were tested. We found that the combination including all features performs best. We refer to this new metric as QSVR.

IV. RESULTS AND DISCUSSIONS

First, we examine which NR methods (section III-A) are suitable for the turbulence case. Measurement values are compared with the chosen FR methods (PSNR, MS-SSIM, VSNR and PIM). Then the selected metrics are used to assess the results of our proposed atmospheric turbulence mitigation and to compare with existing methods.

A. Quality Metric Selection

We generated a number of image sequences containing objects distorted by turbulence using 8 gas hobs. All sequences were acquired with a Canon EOS-1D Mark IV camera with 105mm lens. The distance between the nearest gas hob and the camera was 1m, while the objects were 3.5m away from the camera. The space between each gas hob was 30cm. The flow of gas created temperature gradients leading to distortions in the scene. We captured 8 sequences containing different objects, including faces, common objects and text. The distortions in each video were varied by altering the

(a) Simulated datasets (A1-A8)

(b) Real datasets without ground truth (B1-B6)

Fig. 9. Distorted sequences. (a) Simulated datasets generated from gas burners. (b) Real datasets.

gas flow to produce three classes referred to as low, medium and high distortion with 100 frames each. Frames from each sequence are shown in Fig. 9 (a) and were used to investigate the performance of the NR methods.

The data for each sequence was divided randomly into training and validation sets. As subjective scores for these sequences do not exist, we attempt to find the NR metric that achieves the highest correlation with the FR scores. The values of the FR and the NR approaches are computed and normalised so that the measured qualities span 0 and 1 for better comparison. The performance of the NR methods is also assessed according to the correlation with the average of the FR scores. Fig. 10 shows the average correlation value for the 8 sequences using three statistical parameters: Linear Correlation Coefficient (LCC), Spearman Rank Order Correlations Coefficient (SROCC) and Root Mean Squared Error (RMSE). It can be concluded that only JP2K, AQI and our QSVR are suitable for the air turbulence problem.

The JPEG method is not well matched to our problem since it emphasises blocking artefacts. BIQI is also based on distortions present in compressed data rather than atmospheric distortions. It can be clearly seen from Fig. 10 that QSVR outperforms all other NR methods.

B. Results for Static Scenes

CLEAR has been tested using the three datasets summarised in Table II. The first set (A1-A8) contains the simulated data

Fig. 10. Correlation comparison between NR methods

TABLE II
DISTORTED SEQUENCES

type	#	name	colour type	resolution	# frames
Simulated datasets	A1	Books	RGB	1024×1024	100
	A2	Barcode	RGB	512×256	100
	A3	Back Car	RGB	256×256	100
	A4	Faces	RGB	512×512	100
	A5	Boxes	RGB	320×240	100
	A6	Plant	RGB	512×512	100
	A7	Front Car	RGB	512×512	100
	A8	Toys	RGB	1200×800	100
Real datasets without ground truth	B1	Hill House	grey	512×512	75
	B2	Hot Road	grey	320×240	175
	B3	Cold Car	RGB	720×576	75
	B4	Number Plate	RGB	320×240	200
	B5	Monument	RGB	512×512	100
	B6	Mirage	RGB	256×256	50
Real datasets with ground truth [10]	C1	Chimney	grey	237×237	100
	C2	Building	grey	237×237	100
	C3	Books	grey	109×109	100

Note that sequences B5 and B6 can be downloaded from www.bristol.ac.uk/vi-lab.

described in section IV-A (Fig. 9 (a)), while the other two datasets include real effects of turbulence in long range imaging. These real datasets have been captured without ground truth (B1-B6) and with ground truth (C1-C3) as shown in Fig. 9 (b) and Fig. 14 (first row), respectively. The distance between object and camera corresponding to B2-B6 is approximately 500m, 760m, 200m, 1.5km and 200m, respectively. The ROIs in the sequences B1 (house) and B5 (monument) were on the hill where there was a valley between the objects and camera, while ROIs of other sequences were about 50cm - 1.5m from the ground. Sequence B3 was captured in cold conditions, while the rest were shot in hot weather. Sequences B5 and B6 were collected by one of the authors using a Canon EOS-1D Mark IV camera with 400mm lens at a temperature of 46°C in dry desert conditions.

In the case of colour sequences, the images are converted into $YCbCr$ colour space and only the greyscale channel (Y) is processed. The output colour channels (C_b and C_r) are generated using the average values of all selected registered frames. This simple technique can be applied successfully because the distortions of the colours have less influence on human vision perception. At the end of the process, the greyscale fusion result is combined with the colour channels and converted back to the original colour space.

The parameters used in our experiments are as follows. The DT-CWT is applied with 3 or 4 decomposition levels

TABLE III
ASSESSMENT OF EACH STEP OF CLEAR

sequence	metrics	average original	ROI Align	Frame Select	Regis-tration	Fusion	CLAHE
Cold Car	JP2K	41.57	49.08	52.12	52.74	69.24	69.93
	QSVR	39.25	39.54	39.71	39.86	43.54	43.70
Number Plate	JP2K	42.35	51.35	51.60	55.29	78.02	78.49
	QSVR	32.17	39.96	40.21	40.23	42.87	43.15

when the image resolution is smaller or larger than 128×128 pixels, respectively. The weights for the frame selection w_G , w_S and w_A are set to 1. The value τ for cleaning the map $Q^{\theta,l}$ is individually set for each subband in each level to $\tau^{\theta,l} = |\bar{d}|^{\theta,l} + 0.5 \sigma^{\theta,l}$, where $|\bar{d}|^{\theta,l}$ and $\sigma^{\theta,l}$ are the mean and the standard deviation of the magnitude of the highpass subband. The gains $A_g^{\theta,l}$ are equal for each subband at the same level and are 2.5, 1.8, 1.2 and 1 for $l = 1, 2, 3$ and 4, respectively. For CLAHE, a window of 8×8 pixels is used to compute local histograms which are clipped at 1%.

1) *Relative Contribution of Each Stage in CLEAR*: To better understand why CLEAR produces excellent results, we have analysed the contribution of each stage in the algorithm. The results for 2 sequences, *Cold Car* and *Number Plate*, are shown in Table III. This presents the JP2K and QSVR metrics for each stage in the algorithm cumulatively. The results of ROI alignment, frame selection and registration stages are obtained from the average frame. As it is clearly seen, the fusion stage makes a major contribution toward our method's performance.

2) *Simulated datasets*: The sequences generated with gas burner turbulence (Fig. 9 (a)) are used to compare the performance of CLEAR with that of Shan's Blind Deconvolution (BD) [41] and with the space-variant overlap-add method (SVOLA) by Hirsch *et al.* [10]. In this case, an FR image quality assessment can be used. Fig. 11 shows the performance of the proposed scheme, BD and SVOLA applied to the *Toys* and *Plant* sequences. The blind quality assessment metrics (SVR and AQI) are also computed to demonstrate the correlation between FR and NR methods.

3) *Real datasets without ground truth*: Six sequences, Fig. 9 B1-B6, that exhibit significant turbulence distortions are employed in this trial. It should be noted that, since the *Number Plate*, *Mirage* and *Monument* sequences show a significant shift of the ROI between frames, we also artificially apply spatial shifts of between 1-20 pixels randomly to the other sequences. For comparisons with BD and SVOLA, it is important to note that these two methods have also benefitted from, the ROI alignment, frame selection and registration methods used in CLEAR. If BD and SVOLA are applied without this, their results are significantly inferior. Despite this benefit, the reconstructed images in Fig. 12 and Fig. 13 show that CLEAR can restore better detail and more easily readable text.

The objective results shown in Table IV support the subjective results. The proposed fusion approach achieves better JP2K, AQI and QSVR scores for all distorted sequences,

Fig. 11. Performance comparison between SVOLA and CLEAR using: Left: FR methods, Right: NR methods.

apart from *Number Plate* where the AQI value of SVOLA is slightly better (probably insignificant: $0.003 \approx 0.092\%$) than CLEAR. However, referring to Fig. 12 the subjective result of the *Number Plate* using the proposed approach clearly reveals more readable numbers. It should be noted that this sequence is highly distorted and the number on the plate is impossible to read in any single frame. It should be further noted that the quality values calculated from the whole frame are slightly different from those in Table IV, which relate solely to the ROI.

The subjective results clearly show that the CLEAR algorithm removes atmospheric distortion more efficiently than the other approaches. Shan's BD is inefficient for the sequences degraded by spatially-varying blur since the PSF is assumed to be similar for the entire image, while our method processes sets of homogeneous regions separately. Shan's method also takes four times longer to process than CLEAR mainly due to PSF estimation. SVOLA subdivides an image into overlapped regions and estimates the PSF separately; as a result, it provides better results compared to Shan's method. However, the computation time is even longer and the results are not as sharp as the proposed method. In addition, prior knowledge of PSF size is required for both previous methods.

4) *Real datasets with ground truth*: Three sequences, *Chimney*, *Books* and *Building* have been made available with their ground truth by Hirsch *et al.* [10]. Also, the results from their approach are available. The results are also compared to another atmospheric turbulence removal approach from Zhu [12]. The subjective and objective results for these sequences are shown in Fig. 14 and Table V, respectively. The PSNR, MS-SSIM, VSNR and PIM values reveal that CLEAR outperforms SVOLA and Zhu's approach, but the VIF value is highest for SVOLA. Interestingly, the subjective results shown in Fig. 14 reveal that our results are the most similar to the ground truth. The results of Zhu's fusion appear sharpened around structural features, but lose texture in homogenous areas thereby yielding the lowest objective scores.

Fig. 12. Reconstructed images from real sequences (B1-B5). Left: Shan's BD [41]. Middle: SVOLA [10]. Right: CLEAR. It should be noted that SVOLA and BD results have benefited from CLEAR's selection and registration processes.

C. Results for Sequences Containing Moving Objects

This section shows the potential of the proposed algorithm when applied to videos containing moving objects. Here, part of the moving object is manually selected as the ROI. A number of forward and backward frames are stored in a reference buffer. The forward frames are the future video frames, while the backward frames are the reconstructed frames using the proposed fusion process. The ROI of all frames in the buffer are aligned to the ROI of the current frame using the proposed ROI alignment method. Then, they are registered to their average. Subsequently, these registered frames are fused using the proposed method to remove distortions. Using the reconstructed frames as one of the references can lead to

Fig. 13. *Mirage* sequence (B6). Top-left: Magnified number plates of (from top) original frame, SVOLA, CLEAR and undistorted frames captured at closer distance). Top-right: Original image. Bottom-left: Result of SVOLA. Bottom-right: Result of CLEAR. Again it should be noted that BD and SVOLA results use CLEAR's preprocessing.

TABLE IV
OBJECTIVE RESULTS USING NR METHODS

sequence	method	JPEG	BIQI	JP2K	AQI	QSVR
Hill House	BD	10.07	2.22	52.61	2.13	54.32
	SVOLA	11.15	2.04	51.35	2.22	56.12
	CLEAR	11.05	2.24	61.71	2.99	56.88
Hot Road	BD	8.11	2.42	78.70	2.41	47.77
	SVOLA	9.25	4.09	78.43	2.56	47.81
	CLEAR	10.04	3.28	78.77	3.18	48.72
Cold Car	BD	9.02	1.87	58.40	0.16	42.26
	SVOLA	9.95	2.04	59.85	0.09	42.62
	CLEAR	10.44	2.57	69.93	1.90	43.70
Number Plate	BD	8.21	3.62	68.62	2.00	42.29
	SVOLA	8.99	3.21	58.79	3.26	42.75
	CLEAR	10.36	3.25	78.49	3.26	43.15
Mirage	BD	7.91	2.59	65.32	8.12	41.73
	SVOLA	10.64	3.99	66.83	9.39	42.49
	CLEAR	10.83	4.20	73.83	11.77	43.19
Monument	BD	8.65	2.47	72.84	0.88	45.23
	SVOLA	9.14	3.39	79.65	1.94	47.81
	CLEAR	9.93	3.32	79.98	3.44	48.13

error accumulation over long time periods. Therefore, periodic refresh of the reference buffer is required. Fig. 15 shows parts of the *Wine* sequence where the wine bottle is moved from the right to the left of the display over time. The online process employs two forward and two backward frames and the reference buffer is cleared every 20 frames. The results clearly show improvement in sharpness and motion smoothness. However, we acknowledge that more intelligent algorithms could be developed to remove motion jitter and to indicate the static background so that more reference

Fig. 14. Hirsch's Data (C1-C3). Top-Bottom: Original frame, Ground truth, SVOLA's results, ZHU's results and CLEAR's results

frames can be used to produce better results on such areas. Furthermore, alternative methods such as MOVIE [42] and AVM [39] may be more appropriate for quality assessment.

V. CONCLUSIONS

This paper has introduced a new method for mitigating atmospheric distortion in long-range surveillance imaging. Significant improvements in image quality are achieved using region-based fusion in the DT-CWT domain. This is combined with a new alignment method and cost function for frame selection to pre-process the distorted sequence. The process is completed with local contrast enhancement to reduce haze interference. CLEAR offers class-leading performance for off-line extraction of enhanced static imagery and has the potential to achieve high performance for on-line mitigation for full motion video – this is topic of ongoing research. Experiments with real data show superior performance compared with

TABLE V
OBJECTIVE RESULTS USING FR METHODS

sequence	method	PSNR	MS-SSIM	VSNR	VIF	PIM
Chimney	SVOLA	31.93	0.96	11.19	0.21	35.25
	ZHU	28.90	0.95	10.60	0.13	33.29
	CLEAR	32.02	0.96	12.98	0.19	35.58
Building	SVOLA	24.88	0.88	6.52	0.33	31.03
	ZHU	23.83	0.87	5.89	0.24	30.47
	CLEAR	25.18	0.91	6.95	0.26	31.34
Books	SVOLA	19.39	0.83	8.61	0.43	32.11
	ZHU	20.34	0.85	8.59	0.35	32.12
	CLEAR	25.36	0.94	10.58	0.39	32.57

Fig. 15. The Wine sequence. Top: from distorted video. Bottom: from reconstructed video. Left: the ROI of the 5th frame showing the barcode area which includes the row (y) of the middle pictures. The xz planes (z=temporal direction). Middle: at y=200. Right: at y=100.

existing methods. Using simulated data, full reference metrics clearly show the superiority of this method. We have also introduced a new metric, QSVR, based on support vector regression for blindly assessing image quality. This learning-based method shows higher correlation with the FR methods than existing NR methods.

ACKNOWLEDGMENT

The authors would like to thank Dr. Harish Bhaskar of Khalifa University, UAE, for his help in collecting sequences B5 and B6.

REFERENCES

- [1] L. C. Andrews, R. L. Phillips, C. Y. Hopen, and M. A. Al-Habash, "Theory of optical scintillation," *J. Opt. Soc. Am. A*, vol. 16, no. 6, pp. 1417–1429, June 1999.
- [2] H. S. Rana, "Towards generic military imaging adaptive optics," in *Proc. SPIE 7119, Optics and Photonics for Counterterrorism and Crime Fighting IV*, 2008.
- [3] Kylmar (2012), Advanced electro-optical systems. [Online]. Available: <http://www.kylmar.co.uk>
- [4] B. Davey, R. Lane, and R. Bates, "Blind deconvolution of noisy complex-valued image," *Optics Communications*, vol. 69, pp. 353–356, 1989.
- [5] E. Y. Lam and J. W. Goodman, "Iterative statistical approach to blind image deconvolution," *J. Opt. Soc. Am. A*, vol. 17, no. 7, pp. 1177–1184, July 2000.
- [6] S. Harmeling, M. Hirsch, S. Sra, and B. Scholkopf, "Online blind image deconvolution for astronomy," in *Proc of IEEE Conf Comp. Photogr.*, 2009.

- [7] J. Gilles, T. Dagobert, and C. Franchis, "Atmospheric turbulence restoration by diffeomorphic image registration and blind deconvolution," in *Proc of the 10th Inter Conf on Advanced Concepts for Intelligent Vision Systems*. Springer-Verlag, 2008, pp. 400–409.
- [8] C. S. Huebner, "Compensating image degradation due to atmospheric turbulence in anisoplanatic conditions," in *Proceedings of SPIE 7351, Mobile Multimedia/Image Processing, Security, and Applications*, 2009.
- [9] D. Li, "Suppressing atmospheric turbulent motion in video through trajectory smoothing," *Signal Processing*, vol. 89, no. 4, pp. 649 – 655, 2009.
- [10] M. Hirsch, S. Sra, B. Scholkopf, and S. Harmeling, "Efficient filter flow for space-variant multiframe blind deconvolution," in *Computer Vision and Pattern Recognition (CVPR)*, June 2010, pp. 607 –614.
- [11] J. Delpont, "Scintillation mitigation for long-range surveillance video," *Science real and relevant conference*, 2010.
- [12] X. Zhu and P. Milanfar, "Removing atmospheric turbulence via space-invariant deconvolution," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 35, no. 1, pp. 157 –170, Jan. 2013.
- [13] N. Joshi and M. Cohen, "Seeing Mt. Rainier: Lucky imaging for multi-image denoising, sharpening, and haze removal," in *Computational Photography, IEEE Inter Conf on*, Mar. 2010, pp. 1–8.
- [14] R. Gregory, "A technique of minimizing the effects of atmospheric disturbance on photographic telescopes," *Nature*, vol. 203, 1964.
- [15] Z. Wen, D. Fraser, and A. Lambert, "Bicoherence used to predict lucky regions in turbulence affected surveillance," in *Video and Signal Based Surveillance, IEEE International Conference on*, Nov. 2006, p. 108.
- [16] M. Aubailly, M. Vorontsov, G. Carhart, and M. Valley, "Automated video enhancement from a stream of atmospherically-distorted images: the lucky-region fusion approach," *SPIE*, vol. 7463, 2009.
- [17] P. J. Kent, S. B. Foulkes, J. G. Burnett, S. C. Woods, and A. J. Turner, "Progress towards a real-time active lucky imaging system," in *Technical Conference on Electro Magnetic Remote Sensing*, 2010.
- [18] C. S. Huebner and C. Scheffling, "Software-based mitigation of image degradation due to atmospheric turbulence," in *SPIE conference on Optics in Atmospheric Propagation and Adaptive Systems*, 2010.
- [19] I. Selesnick, R. Baraniuk, and N. Kingsbury, "The dual-tree complex wavelet transform," *Signal Processing Magazine, IEEE*, vol. 22, no. 6, pp. 123 – 151, Nov. 2005.
- [20] Z. Wang, H. Sheikh, and A. Bovik, "No-reference perceptual quality assessment of JPEG compressed images," in *Image Processing. 2002. Proceedings. 2002 International Conference on*, vol. 1, 2002, pp. 477–480.
- [21] H. Sheikh, A. Bovik, and L. Cormack, "No-reference quality assessment using natural scene statistics: JPEG2000," *Image Processing, IEEE Transactions on*, vol. 14, no. 11, pp. 1918 –1927, Nov. 2005.
- [22] S. Gabarda and G. Cristóbal, "Blind image quality assessment through anisotropy," *J. Opt. Soc. Am.*, vol. 24, no. 12, pp. B42–B51, 2007.
- [23] A. Moorthy and A. Bovik, "A two-step framework for constructing blind image quality indices," *Signal Processing Letters, IEEE*, vol. 17, no. 5, pp. 513 –516, May 2010.
- [24] N. Otsu, "A threshold selection method from gray-level histograms," *Systems, Man and Cybernetics, IEEE Transactions on*, vol. 9, no. 1, pp. 62 –66, Jan. 1979.
- [25] H. Chen and N. Kingsbury, "Efficient registration of nonrigid 3-d bodies," *Image Processing, IEEE Transactions on*, vol. 21, no. 1, pp. 262 –272, Jan. 2012.
- [26] J. J. Lewis, R. J. O'Callaghan, S. G. Nikolov, D. R. Bull, C. N. Canagarajah, and E. Basaeed, "Region-based image fusion using complex wavelets," in *Proc. 7th Inter. Conf. Info. Fusion*, 2004, pp. 555–562.
- [27] J. J. Lewis, R. J. O'Callaghan, S. G. Nikolov, and D. R. Bull, "Pixel- and region-based image fusion with complex wavelets," *Information Fusion*, vol. 8, no. 2, pp. 119–130, 2007.
- [28] T. Wan, N. Canagarajah, and A. Achim, "Segmentation-driven image fusion based on alpha-stable modeling of wavelet coefficients," *Multimedia, IEEE Transactions on*, vol. 11, no. 4, pp. 624 –633, June 2009.
- [29] G. Piella, "A general framework for multiresolution image fusion: from pixels to regions," *Information Fusion*, vol. 4, no. 4, pp. 259 – 280, 2003.
- [30] R. O'Callaghan and D. Bull, "Combined morphological-spectral unsupervised image segmentation," *Image Processing, IEEE Transactions on*, vol. 14, no. 1, pp. 49 –62, Jan. 2005.
- [31] A. Loza, D. Bull, N. Canagarajah, and A. Achim, "Non-gaussian model-based fusion of noisy images in the wavelet domain," *Comput. Vis. Image Underst.*, vol. 114, pp. 54–65, Jan. 2010.
- [32] K. He, J. Sun, and X. Tang, "Single image haze removal using dark channel prior," in *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, June 2009, pp. 1956 –1963.

- [33] P. Carr and R. Hartley, "Improved single image dehazing using geometry," in *Digital Image Computing: Techniques and Applications*, 2009. *DICTA '09*, dec. 2009, pp. 103–110.
- [34] K. Zuiderveld, *Contrast limited adaptive histogram equalization*. Academic Press Professional, Inc., 1994, pp. 474–485.
- [35] T. Gureyev, A. Stevenson, Y. Nesterets, and S. Wilkins, "Image deblurring by means of defocus," *Optics Communications*, vol. 240, no. 1-3, pp. 81–88, Oct. 2004.
- [36] W. Zhang and W.-K. Cham, "Single-image refocusing and defocusing," *Image Processing, IEEE Transactions on*, vol. 21, no. 2, pp. 873–882, Feb. 2012.
- [37] Z. Wang, E. Simoncelli, and A. Bovik, "Multiscale structural similarity for image quality assessment," in *Signals, Systems and Computers, 2003. Conference Record of the Thirty-Seventh Asilomar Conference on*, vol. 2, Nov. 2003, pp. 1398–1402 Vol.2.
- [38] D. Chandler and S. Hemami, "VSNR: A wavelet-based visual signal-to-noise ratio for natural images," *Image Processing, IEEE Transactions on*, vol. 16, no. 9, pp. 2284–2298, Sept. 2007.
- [39] F. Zhang and D. Bull, "A parametric framework for video compression using region-based texture models," *Selected Topics in Signal Processing, IEEE Journal of*, vol. 5, no. 7, pp. 1378–1392, Nov. 2011.
- [40] C. C. Chang and C. J. Lin, "LIBSVM: A library for support vector machines," *ACM Transactions on Intelligent Systems and Technology*, vol. 2, pp. 27:1–27:27, 2011.
- [41] Q. Shan, J. Jia, and A. Agarwala, "High-quality motion deblurring from a single image," *ACM Transactions on Graphics (SIGGRAPH)*, 2008.
- [42] K. Seshadrinathan and A. Bovik, "Motion tuned spatio-temporal quality assessment of natural videos," *Image Processing, IEEE Transactions on*, vol. 19, no. 2, pp. 335–350, Feb. 2010.

N. Anantrasirichai (S'04-M'07) received the B.E. degree in electrical engineering from Chulalongkorn University, Bangkok, Thailand, in 2000, the M.E. degree in telecommunications from the Asian Institute of Technology, Pathumthani, Thailand, in 2003, and the Ph.D. degree in electrical and electronic engineering from the University of Bristol, Bristol, U.K., in 2007. She is currently a research associate at the Visual Information Laboratory, University of Bristol. Her research interests include image and video coding, image analysis and enhancement, image fusion and medical imaging.

Alin Achim (S'99-M'04-SM'09) received the B.Sc. and M.Sc. degrees, both in electrical engineering, from University "Politehnica" of Bucharest, Romania, in 1995 and 1996, respectively and the Ph.D. in biomedical engineering from the University of Patras, Greece, in 2003. He then obtained an ERCIM (European Research Consortium for Informatics and Mathematics) postdoctoral fellowship which he spent with the Institute of Information Science and Technologies (ISTI-CNR), Pisa, Italy and with the French National Institute for Research in Computer Science and Control (INRIA), Sophia Antipolis, France. In October 2004 he joined the Department of Electrical & Electronic Engineering at the University of Bristol as a Lecturer. He became Senior Lecturer in 2010. His research interests are in statistical signal, image and video processing, multiresolution algorithms, wavelet analysis of digital images, image filtering and fusion, segmentation and classification. He has co-authored over 80 scientific publications, is an editorial board member of *Digital Signal Processing* (Elsevier), and has served on numerous conferences technical programme committees in signal and image processing. Alin has been organising committee member of the International Conference on Biosensing Technologies (Bristol 2009, Amsterdam 2011 and Barcelona 2013), of the Biomedical Imaging Workshop (Bristol 2010), and of the INGIMED X Conference (Bucharest 2009). He is member of the IET and a Senior Member of the IEEE.

Nick G. Kingsbury received the honours degree in 1970 and the Ph.D. degree in 1974, both in electrical engineering, from the University of Cambridge. He is a Fellow of the IEEE. From 1973 to 1983 he was a Design Engineer and subsequently a Group Leader with Marconi Space and Defence Systems, Portsmouth, England, specializing in digital signal processing and coding theory. Since 1983 he has been a Lecturer in Communications Systems and Image Processing at the University of Cambridge and a Fellow of Trinity College, Cambridge. He was appointed to Professor of Signal Processing in 2007 and is head of the Signal Processing and Communications Research Group. His current research interests include image analysis and enhancement techniques, object recognition, motion analysis and registration methods. He has developed the dual-tree complex wavelet transform and is especially interested in the application of wavelet frames to the analysis of images and 3-D datasets.

David R. Bull (M'94-SM'07-F'13) received the B.Sc. degree from the University of Exeter, Devon, U.K., the M.Sc. degree from the University of Manchester, Manchester, U.K., and the Ph.D. degree from the University of Wales, Cardiff, U.K., in 1980, 1983, and 1988, respectively. He holds the Chair in Signal Processing at the University of Bristol, Bristol, U.K. His previous roles include Lecturer with the University of Wales and Systems Engineer with Rolls Royce. He was the Head of the Electrical and Electronic Engineering Department at

the University of Bristol, from 2001 to 2006, and is currently the Director of Bristol Vision Institute. In 2001, he co-founded ProVision Communication Technologies, Ltd., Bristol, and was its Director and Chairman there until 2011. He has worked widely in the fields of 1-D and 2-D signal processing and has published some 400 academic papers and articles and has written two books. He has also given numerous invited/keynote lectures and tutorials. He holds numerous patents, several of which have been exploited commercially. His current research interests include problems of image and video communication and analysis for wireless, internet, military, broadcast and immersive applications. Prof. Bull is an IEEE Fellow and has been awarded two IEE Premium awards for his work.