

Wavelet Enhanced Appearance Modelling

Mikkel B. Stegmann^a, Søren Forchhammer^b, Timothy F. Cootes^c

^aInformatics and Mathematical Modelling, Technical University of Denmark,
Richard Petersens Plads, Building 345v, DK-2800 Kgs. Lyngby, Denmark

^bResearch Centre COM, Technical University of Denmark,
Building 371, DK-2800 Kgs. Lyngby, Denmark

^cDivision of Imaging Science and Biomedical Engineering,
University of Manchester, Manchester, England

ABSTRACT

Generative segmentation methods such as the Active Appearance Models (AAM) establish dense correspondences by modelling variation of shape and pixel intensities. Alas, for 3D and high-resolution 2D images typical in medical imaging, this approach is rendered infeasible due to excessive storage and computational requirements. This paper extends the previous work of Wolstenholme and Taylor where Haar wavelet coefficient subsets were modelled rather than pixel intensities. In addition to a detailed review of the method and a discussion of the integration into an AAM-framework, we demonstrate that the more recent bi-orthogonal CDF 9-7 wavelet offers advantages over the traditional Haar wavelet in terms of synthesis quality and accuracy. Further, we demonstrate that the inherent frequency separation in wavelets allows for simple band-pass filtering, e.g. edge-emphasis. Experiments using Haar and CDF 9-7 wavelets on face images have shown that segmentation accuracy degrades gracefully with increasing compression ratio. Further, a proposed weighting scheme emphasizing edges was shown to be significantly more accurate at compression ratio 1:1, than a conventional AAM. At higher compression ratios the scheme offered both a decrease in complexity *and* an increase in segmentation accuracy.

Keywords: registration, compression, atlases, deformable models, active appearance models, wavelets, face images

1. INTRODUCTION

Generative models capable of synthesising photo-realistic images of objects have been shown to be powerful tools for image interpretation. Explicitly modelling the value of every pixel covering an object is feasible for low-resolution 2D images. However, for high-resolution 2D images, 2D and even 3D time-series, this approach is rendered infeasible due to excessive storage and computational requirements. This is especially unfortunate due to the many information-rich medical imaging modalities emerging and maturing these days. We therefore address this problem by applying wavelet compression to the widely used generative model; the Active Appearance Model (AAM).^{3,8} Refer to^{6,11} for reviews of various applications.

Since wavelets provide an approximate decorrelated representation of each training image, we aim at choosing consistent subsets of wavelet coefficients over the complete training set. The AAM is subsequently built using these subsets in order to lessen the computational and storage requirements. The major challenge in this context is to choose an appropriate wavelet basis and a subset that affects the final performance measure as little as possible. In this study we focus on the segmentation capabilities of AAMs.

The presented work is validated on an example of very complex biological variation in both shape and appearance. It consists of a cohort of face images. Examples of medical use of statistical face models include the work of Hutton et al.,⁹ which aims at a better understanding of syndromic facial dysmorphologies in a study of the Noonan syndrome.

Corresponding author is M. B. Stegmann, E-mail: mbs@imm.dtu.dk, Web: <http://www.imm.dtu.dk/~mbs/>.

2. RELATED WORK

Model size concerns have been present since the introduction of AAMs.^{3,8} The initial models contained $\sim 10,000$ pixels and were built on sub-sampled input images of approximately 200×300 pixels.⁶ Aiming at reducing computational costs Cootes et al.² used a sub-sampling scheme to reduce the texture model by a ratio of 1:4. The scheme selected a subset based on the ability of each pixel to predict corrections of the model parameters. When exploring different multi-band appearance representations Stegmann and Larsen¹² studied the segmentation accuracy of face AAMs at different scales in the range 850 – 92.118 pixels obtained by pixel averaging. Mitchell et al.¹⁰ introduced 3D AAMs and noticed that a current limitation was the requirement of heavy down-scaling of their input images (3D cardiac MRI and 2D+time ultra sound).

This paper extends directly on the original work of Wolstenholme and Taylor,^{13,14} which introduced the idea of modelling wavelet coefficient subsets in Active Appearance Models. They augmented an AAM with the Haar wavelet and carried out a set of experiments on brain MRI using a fixed model size compression ratio of 1:20 using both a multi-resolution and single resolution model. Other compression ratios were mentioned, however, no qualitative results were given for these.

3. ACTIVE APPEARANCE MODELS

Active Appearance Models are generative models capable of synthesising images of a given object class. By estimating a compact and specific basis from a training set, model parameters can be adjusted to fit unseen images and hence perform image interpretation. The modelled object properties are usually shape and pixel intensities (here denoted *texture*). Training objects are defined by marking up each example image with points of correspondence. Using prior knowledge of the optimisation space, AAMs can be rapidly fitted to unseen images, given a reasonable initialisation.

Variability is modelled by means of principal component analyses (PCA). Prior to PCA modelling shapes are Procrustes aligned and textures are warped into a shape-free reference frame and sampled. Let there be given P training examples for an object class, and let each example be represented by a set of N landmark points and M texture samples. Let \mathbf{s} and \mathbf{t} denote a synthesised shape and texture and let $\bar{\mathbf{s}}$ and $\bar{\mathbf{t}}$ denote the corresponding means. Synthetic examples are parameterised by \mathbf{c} and generated by

$$\mathbf{s} = \bar{\mathbf{s}} + \Phi_{s,c}\mathbf{c} \quad , \quad \mathbf{t} = \bar{\mathbf{t}} + \Phi_{t,c}\mathbf{c} \quad (1)$$

where $\Phi_{s,c}$ and $\Phi_{t,c}$ are composed of eigenvectors estimated from the training set dispersions. An object instance, $\{\mathbf{s}, \mathbf{t}\}$, is synthesised into an image by warping the pixel intensities of \mathbf{t} into the geometry of the shape \mathbf{s} .

An AAM is matched to an unseen image using a least-squares criterion by an iterative updating scheme. This is based on a fixed Jacobian estimate,⁴ or, originally, a principal component regression.³ For this work we have chosen the former approach in a setup treated in detail elsewhere.¹¹ For further details on AAMs refer to.^{3,4,6}

4. WAVELETS

Wavelets are a family of basis functions that decompose signals into both space and frequency. In the following we use the discrete wavelet transform (DWT), which can be viewed as a set of linear, rank-preserving matrix operations. In practice these are carried out in a convolution scheme known as the fast wavelet transform (FWT)¹ where an image is decomposed by a high-pass filter into a set of detail wavelet sub bands, and by a low-pass filter into a scaling sub band. These bands are then down-sampled and can be further decomposed. We use the dyadic (octave) decomposition scheme that successively decomposes the scaling sub band, yielding a discrete frequency decomposition. Alternative decomposition schemes include the wavelet packet basis where successive decompositions are carried out in the detail sub bands as well.

Figure 1 shows a two-level octave wavelet decomposition. The first, third and fourth quadrants are the detail sub bands and stem from the initial decomposition (level 1). The four sub-quadrants of the second quadrant stem from the second decomposition (level 2) with the scaling sub band at the top left corner.



Figure 1. The wavelet coefficients of a two-level octave decomposition using the Haar wavelet.

Wavelets are invertible which is typically achieved by orthonormality. Wavelet transforms can thus be considered a rotation in function space, which in addition – through successive decompositions – adds a notion of scale. This scale property lends itself nicely to progressive signal processing. Wavelets for image compression are designed to perform rotations that decorrelate image data by using vanishing moments. Wavelet coefficients close to zero can thus be removed with minimal impact on the reconstruction.

Bi-orthogonal wavelets will also be considered in the following. These are not strictly orthogonal and therefore have odd finite length and linear phase. They come in pairs of analysis and synthesis filters, together forming a unitary operation.

5. WAVELET ENHANCED APPEARANCE MODELLING

This section introduces a notation for wavelet compression and describes how it can be integrated into an AAM framework thereby obtaining a Wavelet enHanced Appearance Model (WHAM).

First, let an n -level wavelet transform be denoted by

$$\mathcal{W}(\mathbf{t}) = \mathbf{\Gamma}\mathbf{t} = \hat{\mathbf{w}} = [\hat{\mathbf{a}}^\top \hat{\mathbf{u}}_1^\top \cdots \hat{\mathbf{u}}_n^\top]^\top \quad (2)$$

where $\hat{\mathbf{a}}$ and $\hat{\mathbf{u}}$ denote scaling and detail wavelet coefficients, respectively. For 2D images each set of detail coefficients is an ensemble of horizontal, vertical and diagonal filter responses. Compression is now obtained by a truncation of the wavelet coefficients

$$\mathcal{C}(\hat{\mathbf{w}}) = \mathbf{C}\hat{\mathbf{w}} = \mathbf{w} = [\mathbf{a}^\top \mathbf{u}_1^\top \cdots \mathbf{u}_n^\top]^\top \quad (3)$$

where \mathbf{C} is a modified identity matrix, with rows corresponding to truncated coefficients removed.

As in¹⁴ a WHAM is built on the truncated wavelet basis, $\mathbf{w} = \mathcal{C}(\mathcal{W}(\mathbf{t}))$, rather than the raw image intensities in \mathbf{t} . This splits all texture-related matrices into scale-portion. For the texture PCA of wavelet coefficients we have

$$\mathbf{w} = \bar{\mathbf{w}} + \mathbf{\Phi}_w \mathbf{b}_w \Leftrightarrow \begin{bmatrix} \mathbf{a} \\ \mathbf{u}_1 \\ \vdots \\ \mathbf{u}_n \end{bmatrix} = \begin{bmatrix} \bar{\mathbf{a}} \\ \bar{\mathbf{u}}_1 \\ \vdots \\ \bar{\mathbf{u}}_n \end{bmatrix} + \begin{bmatrix} \mathbf{\Phi}_a \\ \mathbf{\Phi}_{u_1} \\ \vdots \\ \mathbf{\Phi}_{u_n} \end{bmatrix} \mathbf{b}_w \quad (4)$$




Figure 2. Two-level wavelet decomposition of a texture vector.

where Φ_w is the eigenvectors of the wavelet coefficient covariance matrix. Rearranging this into scaling and detail terms we get

$$\mathbf{a} = \bar{\mathbf{a}} + \Phi_a \mathbf{b}_w \quad (5)$$

and

$$\{\mathbf{u}_i = \bar{\mathbf{u}}_i + \Phi_{u_i} \mathbf{b}_w\}_{i=1}^n. \quad (6)$$

The texture model is thus inherently multi-scale and may be used for analysis/synthesis at any given scale. Motivations for doing so include robustness and computational efficiency. Compared to multi-scale AAMs this also gives a major decrease in storage requirements. However, since the PCA is calculated at full scale (including all detail bands) we can not expect this to be equal to the separate scale-PCAs of conventional multi-scale AAMs.

Using non-truncated orthogonal wavelets (i.e. $\mathbf{C} = \mathbf{I}$), $\mathcal{W}(\mathbf{t})$ is a true rotation in texture hyperspace. Hence the wavelet PCA is a rotation of the original intensity PCA, i.e. $\Phi_w = \Gamma \Phi_t$, iff \mathcal{W} is fixed over the training set. The PC scores are identical, $\mathbf{b}_w = \mathbf{b}_t$. If \mathbf{C} is chosen to truncate the wavelet basis along directions with near-zero magnitude, wavelet PC scores obviously resembles the original PC scores closely.

Direct usage of the sparse matrix Γ is excessively slow. Instead the fast wavelet transform (FWT) is applied. Figure 2 shows the stages of transformation. First, a normalised texture vector is rendered into its equivalent shape-free image. Secondly, the shape-free image is expanded into a dyadic image representation to avoid any constraints on n due to image size. This image is then transformed using FWT and rendered into the vector $\hat{\mathbf{w}}$ by masking out areas outside the octave representation of the reference shape. Eventually, $\hat{\mathbf{w}}$ is truncated into \mathbf{w} .

5.1. Free Parameters and Boundary Effects

To apply the above, the type of wavelet, \mathcal{W} , must be chosen and values for n and \mathbf{C} must be determined. Being a key issue, the estimation of \mathbf{C} is described in a later section.

The choice of wavelet type depends on two factors, i) the nature of the data, and ii) computational issues. Data containing sharp edges suggests sharp filters such as the Haar wavelet and smooth data requires smooth wavelets. However, the price comes in form of increased computational load induced by the extra filter-taps. Since the wavelets operate in a finite discrete domain, traditional boundary issues arise. Consequently, the longer, smoother wavelets suffer the most. To calculate filter responses for a full image, boundaries are typically extended by mirroring. The width of the boundary extension is half the width of the wavelet filter. For even-length wavelets such as Haar the symmetry point is placed between the edge pixel and the first boundary extension pixel. For odd-length wavelets the symmetry point is placed on the edge pixel. Normally this is carried out as a rectangular extension but in this application the extension should adapt to the shape of the texture image.

When using an out-of-the-box wavelet package for rectangular domains, a poor mans solution is to approximate this with a shape adaptive horizontal/vertical mirroring of the dyadic shape image. The width of the mirroring should be $\langle \text{filter width} \rangle \times 2^{n-1}$.

5.2. WHAM Building

Building a WHAM can be summarised in five major steps:

1. Sample all training textures into $\{\mathbf{t}_i\}_{i=1}^P$.
2. Transform $\{\mathbf{t}_i\}_{i=1}^P$ into $\{\hat{\mathbf{w}}_i\}_{i=1}^P$.
3. Estimate $\overline{\hat{\mathbf{w}}}$ and \mathbf{C} .
4. Truncate $\{\hat{\mathbf{w}}_i\}_{i=1}^P$ into $\{\mathbf{w}_i\}_{i=1}^P$.
5. Build an AAM on $\{\mathbf{w}_i\}_{i=1}^P$.

Further, all incoming textures in subsequent optimisation stages should be replaced with their truncated wavelet equivalents, i.e. $\mathcal{C}(\mathcal{W}(\mathbf{t}))$. The synthesis of a WHAM is the reverse of Figure 2, again with appropriate masking and mirroring. Truncated wavelet coefficients are reconstructed using the corresponding mean values:

$$\hat{\mathbf{w}}_{\text{synth}} = \mathcal{C}_{\text{synth}}(\mathbf{w}) = \mathbf{C}^T \mathbf{w} + \mathbf{C}_{\text{synth}} \overline{\hat{\mathbf{w}}} \quad (7)$$

where $\mathbf{C}_{\text{synth}}$ is a modified identity matrix, with rows corresponding to non-truncated coefficients zeroed.

5.3. A Note on Representation

Depending on the AAM implementation, the conversion from \mathbf{t} to the shape-free image might be avoided since this is often the canonical representation of \mathbf{t} after warping. The calculation of the truncation can also be done directly in image-space to avoid unnecessary transforms from vector to image-space, et cetera.

5.4. Wavelet Coefficient Selection

For segmentation purposes the optimal \mathbf{C} is given by the minimum average error between the optimised model shape and the ground truth shape over a test set of t examples

$$\arg \min_{\mathbf{C}} \left(\sum_{i=1}^t |\mathbf{s}_{i,\text{model}} - \mathbf{s}_{i,\text{ground truth}}|^2 \right), \quad (8)$$

subject to the constraint that \mathbf{C} has m rows. This gives rise to a compression of ratio $1 : M/m$.

Contrary to traditional image compression, quantization and entropy coding are not carried out. Consequently, this compression ratio does not commensurate to those in the compression literature.

Unfortunately, direct optimisation of Equation 8 is not feasible since each cost function evaluation involves building a complete model from the training set and a subsequent evaluation on a test set. Alternatively one can design \mathbf{C} based on prior beliefs.

The traditional approach when dealing with compression of training data ensembles, also taken by,¹⁴ is to let \mathbf{C} preserve per-pixel variance over the training set. This is accomplished by constructing \mathbf{C} to select the m largest coefficients from

$$\kappa = \sum_{i=1}^P (\hat{\mathbf{w}}_i - \overline{\hat{\mathbf{w}}}) \odot (\hat{\mathbf{w}}_i - \overline{\hat{\mathbf{w}}}), \quad (9)$$

where \odot denotes the Hadamard product (i.e. element-wise multiplication). We further assume spatial coherence and regularise this scheme prior to coefficient selection by a smoothing of κ (in the image domain) using a convolution with a suitable kernel.

In summary, this prior equates variance with interesting signal.

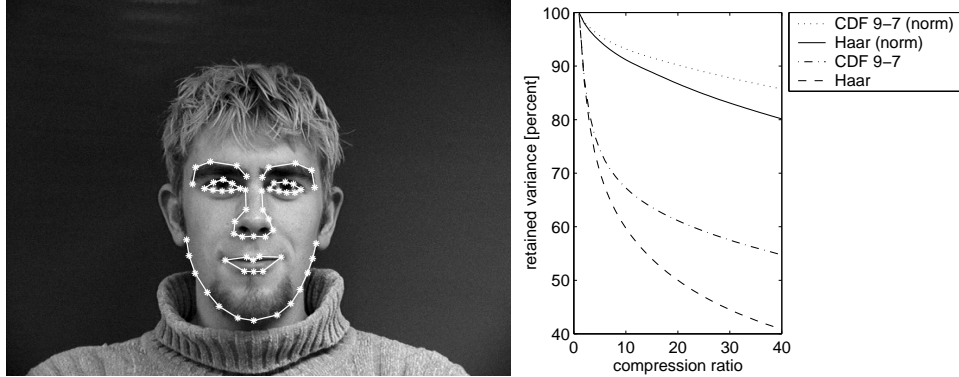


Figure 3. Left: Example annotation of a face using 58 landmarks. Right: Retained wavelet coefficient variance over the training set as a function of factor by which the number of wavelet coefficients have been reduced.

5.5. Signal Manipulation

An important added bonus from the wavelet representation is that certain signal manipulations become exceedingly simple. Due to the band separation, frequency response modifications are easily carried out. The simplest modification is to change the norm of the high- and low-pass filters used. To obtain orthonormality the filters must be normalised. With a minor abuse of notation this is denoted by $\|\mathcal{W}\|_2 = 1$. We call this the *normalised case*. To emphasise high-frequency content $\|\mathcal{W}\|_2$ must be less than one. In the following we propose to use the norm $\|\mathcal{W}\|_2 = 1/\sqrt{2}$ and call this the *weighted case*.

5.6. Extension to Multi-channel AAMs

Support for multi-channel images such as RGB,⁷ Edge AAMs⁵ or composite representations¹² can be implemented in different ways. A simple approach is to wavelet transform each texture band separately with non-coupled truncation schemes. However, this often leads to situations where one multi-band pixel is partially truncated and partially included into the model. If this is not desirable, a composite measure must be used when estimating \mathbf{C} , e.g. the sum of band variances at each pixel position. If pixels are correlated across bands as in e.g. RGB images, the wavelet decomposition can be carried out as a 3D wavelet transformation, resulting in higher compression (or better quality).

6. EXPERIMENTAL RESULTS

To assess the segmentation accuracy of a WHAM a set of experiments was carried out on 37 still images of people facing the camera acquired as 640×480 pixel, 8 bit, grey-scale images and annotated using 58 landmarks, see Figure 3 (left).

The average landmark distance from model to ground truth (pt.pt.) was used as performance measure. Model searches were initialised by displacing the mean configuration $\pm 10\%$ of its width and height in x and y from the optimal position. To reduce statistical fluctuation and reduce pessimistic bias leave-one-out evaluations were carried out. Two wavelet bases were evaluated; the orthogonal Haar wavelet (2 filter taps) and the widely used bi-orthogonal CDF 9-7 wavelet (9 and 7 filter taps). Both of these were tested in the normalised and the weighted case using compression ratios in the range 1:1 – 1:40 and compared to a standard AAM. This led to a total of $37 \times 2 \times 2 \times 9 + 37 \times 1 = 1369$ models. All experiments used three decomposition levels. Results are shown in Figure 4 and 5. The standard AAM contained 31224 pixels on average over all leave-one-out experiments.

In the normalised case we observe that the accuracy degrades gracefully with increasing compression ratio. Though consistently better on average at all ratios, the CDF 9-7 wavelet is not markedly better than the simple Haar wavelet in terms of accuracy. In the weighted case we observe a considerable improvement in accuracy for almost all compression ratios. The non-overlapping notches in the box plots in Figure 5 show that the median

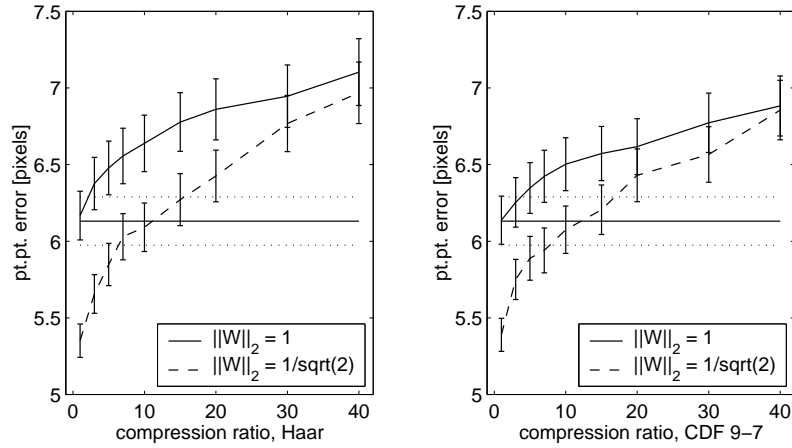


Figure 4. Average segmentation error vs. compression ratio. Error bars are one standard error. Horizontal lines show values for a standard AAM.

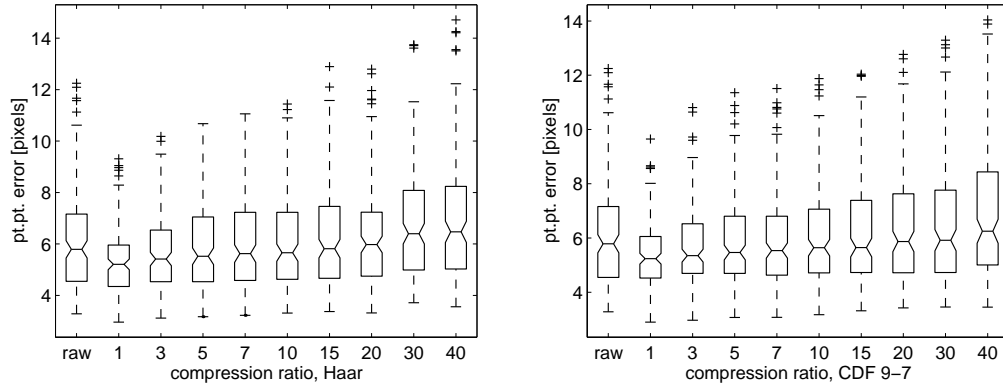


Figure 5. Boxplots of segmentation error in the weighted case vs. compression ratio including a standard AAM (raw). Whiskers are 1.5 IQR at maximum.



Figure 6. Selected wavelet coefficients for the face training set (CDF, ratio 1:10). The scaling coefficients are shown in the upper left. Left: $\|W\|_2 = 1$. Right: $\|W\|_2 = 1/\sqrt{2}$.




Figure 7. The first combined mode of texture and shape variation; $c_1 = \{-3\sigma_1, 0, 3\sigma_1\}$. Top: AAM (i.e. uncompressed). Middle: WHAM (Haar, ratio 1:10). Bottom: WHAM (CDF 9-7, ratio 1:10).

of the standard AAM is significantly worse than both the median of the uncompressed Haar and the CDF 9-7 wavelet at a 0.05 level.

To investigate this behaviour further, Figure 6 shows the selected wavelet coefficients using the CDF 9-7 wavelet at ratio 1:10 using both normalised and weighted filters. The normalised case tends to preserve mostly low-frequency content, while the weighed case distributes wavelets coefficients near image details more evenly across levels. Together with the nature of the weighting this leads to models with more emphasis on edges.

The amount of retained wavelet coefficient variance over the training set is plotted in Figure 3 (right). As expected, the CDF 9-7 performs consistently better than the Haar wavelet. Further, the normalised wavelets are seen to put relatively more variance into fewer components, when compared to the weighted.

Finally, Figure 7 shows the first combined mode of texture and shape deformation for an uncompressed model and two compressed models at compression ratio 1:10. Subtle blocking artefacts from the Haar wavelet are present, while the smooth 9-7 leaves a more visually pleasing synthesis result. Quantitatively, the 9-7 wavelet also offered a better reconstruction of the training textures in terms of mean squared error (MSE).

7. DISCUSSION

Although no significant difference was observed between the two wavelets in terms of segmentation accuracy; the CDF 9-7 wavelet performed consistently better. Further, the CDF 9-7 also provided the best synthesis quality. However, if low computational complexity is required the Haar wavelet should be chosen.

Our case study supports the expected behaviour that it is the high-frequency image content that provides the segmentation accuracy. Variance preserving filters are optimal for reconstruction but not necessarily for segmentation. Inherent scale-separation has been obtained in the presented unified AAM and DWT framework and enabled a very simple method for favouring high-frequency content. This is obtained while still retaining the robustness gained from a low frequency representation of the object.

This paper extends the previous work of Wolstenholme and Taylor where an AAM was augmented with the Haar wavelet and evaluated on brain MRI using a fixed compression ratio of 1:20. Our work validates their previous findings in a different case study using a more thorough evaluation methodology. In addition, the more recent CDF 9-7 wavelet is evaluated and compared to the Haar wavelet. Finally, we have demonstrated that the inherent frequency separation in wavelets allows for simple band-pass filtering, which enables compression schemes that both decrease complexity *and* increase segmentation accuracy.

8. FUTURE WORK

Future work should include investigation of more elaborate coefficient selection and sub band weighting schemes and the (straightforward) extension to n -D images. The latter could be applied to medical data such as 3D and 4D cardiac MRI, 3D brain MRI, ultra sound etc. Exploiting the inherent scale properties of a WHAM is also attractive in order to obtain computationally cheaper and more robust means of optimisation. Finally, earlier work on denoising in the wavelet domain could also be incorporated into this framework.

9. CONCLUSION

This work has described how an AAM framework can be augmented with wavelet compression to reduce model complexity to successfully cope with large medical data sets. Experimental validation has been carried out using two different wavelet bases at nine different compression ratios. It was found that segmentation accuracy degrades gracefully with increasing compression ratio. Our studies indicated that the CDF 9-7 wavelet should be considered as a worthy alternative to the Haar wavelet. Further, it was shown that a simple weighting scheme could enable wavelets to both decrease complexity *and* increase segmentation accuracy.

We anticipate that wavelet-based appearance modelling will become an important technique with many applications, in particular within medical image interpretation.

REFERENCES

1. M. Antonini, M. Barlaud, P. Mathieu, and I. Daubechies. Image coding using wavelet transform. *Image Processing, IEEE Transactions on*, 1(2):205–220, 1992.
2. T. F. Cootes, G. Edwards, and C. J. Taylor. A comparative evaluation of active appearance model algorithms. In *BMVC 98. Proc. of the Ninth British Machine Vision Conf.*, volume 2, pages 680–689. Univ. Southampton, 1998.
3. T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. In *Proc. European Conf. on Computer Vision*, volume 2, pages 484–498. Springer, 1998.
4. T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. *IEEE Trans. on Pattern Recognition and Machine Intelligence*, 23(6):681–685, 2001.
5. T. F. Cootes and C. J. Taylor. On representing edge structure for model matching. In *Proc. IEEE Computer Vision and Pattern Recognition – CVPR*, volume 1, pages 1114–1119. IEEE, 2001.
6. T. F. Cootes and C. J. Taylor. *Statistical Models of Appearance for Computer Vision*. Tech. report, University of Manchester, 2001.
7. G. J. Edwards, T. F. Cootes, and C. J. Taylor. Advances in active appearance models. In *Proc. Int. Conf. on Computer Vision*, pages 137–142, 1999.
8. G. J. Edwards, C. J. Taylor, and T. F. Cootes. Interpreting face images using active appearance models. In *Proc. 3rd IEEE Int. Conf. on Automatic Face and Gesture Recognition*, pages 300–5. IEEE Comput. Soc, 1998.
9. T. J. Hutton, B. R. Buxton, and P. Hammond. Dense surface point distribution models of the human face. *Proceedings IEEE Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA 2001)*, pages 153–60, 2001.
10. S. C. Mitchell, J. G. Bosch, B. P. F. Lelieveldt, R. J. van der Geest, J. H. C. Reiber, and M. Sonka. 3-D active appearance models: Segmentation of cardiac MR and ultrasound images. *Medical Imaging, IEEE Transactions on*, 21(9):1167–1178, 2002.
11. M. B. Stegmann, B. K. Ersbøll, and R. Larsen. FAME – a flexible appearance modelling environment. *IEEE Trans. on Medical Imaging*, 22(10):1319–1331, 2003.
12. M. B. Stegmann and R. Larsen. Multi-band modelling of appearance. *Image and Vision Computing*, 21(1):61–67, jan 2003.
13. C. B. H. Wolstenholme and C. J. Taylor. Using wavelets for compression and multiresolution search with active appearance models. *BMVC99. Proceedings of the 10th British Machine Vision Conference*, pages 473–82 vol.2, 1999.
14. C. B. H. Wolstenholme and C. J. Taylor. Wavelet compression of active appearance models. In *Medical Image Computing and Computer-Assisted Intervention, MICCAI*, pages 544–554, 1999.