Localisation of Vertebrae on DXA Images using Constrained Local Models with Random Forest Regression Voting

P.A. Bromiley, J.E. Adams and T.F. Cootes

Imaging Sciences Research Group, University of Manchester, UK.
`paul.bromiley@manchester.ac.uk`

Abstract. Fractures associated with osteoporosis are a significant public health risk, and one that is likely to increase with an ageing population. However, many osteoporotic vertebral fractures present on images do not come to clinical attention or lead to preventative treatment. Furthermore, vertebral fracture assessment (VFA) typically depends on subjective judgement by a radiologist. The potential utility of computer-aided VFA systems is therefore considerable. Previous work has shown that Active Appearance Models (AAMs) give accurate results when locating landmarks on vertebra in DXA images, but can give poor fits in a substantial subset of examples, particularly the more severe fractures. Here we evaluate Random Forest Regression Voting Constrained Local Models (RFRV-CLMs) for this task and show that, while they lead to slightly poorer median errors than AAMs, they are much more robust, reducing the proportion of larger failures by 66%. They are thus more suitable for use in computer-aided VFA systems.

1 Introduction

Osteoporosis is a common skeletal disorder defined by a reduction in bone mineral density (BMD) resulting in a T-score of <2.5 (i.e. more than 2.5 standard deviations below the mean in young adults), measured using dual energy X-ray absorptiometry (DXA) images [1]. It significantly increases the risk of fractures, most commonly occurring in the hips, wrists or vertebrae. Approximately 40% of postmenopausal caucasian women are affected, increasing their lifetime risk of fragility fractures to as much as 40% [1]. Osteoporosis therefore presents a significant public health problem for an ageing population.

Accurate identification of vertebral fractures is clinically important in the diagnosis of osteoporosis. Radiological assessment typically uses a semi-quantitative approach [2] requiring subjective judgement by a radiologist. Furthermore, only about one third of vertebral fractures present on images come to clinical attention; they are frequently not noted by radiologists, not entered into medical notes, and do not lead to preventative treatments [3]. Many of these cases involve images acquired for purposes other than VFA. However, a recent multicenter, multinational prospective study [4] has found a false negative rate of 34% in

VFA performed on lateral radiographs of the thoracolumbar spine. The potential utility of computer-aided VFA systems is therefore considerable.

Several authors have investigated the use of methods based on statistical shape models to segment vertebrae in both radiographs and DXA images (e.g. [5]) as a preliminary step for VFA. However, state-of-the-art results achieved using active appearance models (AAMs) [6] exhibit significant numbers of large errors due to model fitting failure, particularly on the more severely fractured vertebrae. This is the result of two effects. First, osteoporosis patients with vertebral fractures most commonly have only one or two fractures (e.g. [7]). Therefore, models encompassing multiple vertebrae must typically be trained on datasets containing more normal than fractured vertebrae, potentially introducing a bias against the most severe vertebral shape changes. Second, work on natural images of faces has shown that holistic methods such as AAMs, which rely on a single model of shape and intensity that covers all landmark points, tend to generalise poorly [8]. An alternative is to use a set of models, each covering an individual landmark point. The ambiguity inherent in the use of local image patches may be dealt with by imposing a global shape constraint (e.g. [9]). In particular, regression voting (RV) methods (e.g. [10]), especially those (e.g. [11] [12] [8]) based on Random Forests (RFs) [13] tend to be robust. The RFRV Constrained Local Model (RFRV-CLM) [11] [12], which uses a RF regressor for each point constrained by a global shape model, has been applied successfully to the annotation of landmark points both in facial (e.g. [11]) and clinical (e.g. [12]) images, and shows superior generalisation performance on facial images compared to the AAM [14].

The hypothesis investigated here is that the superior generalisation capability of the RFRV-CLM will lead to performance improvements, compared to AAMs, in terms of the number of large errors due to fit failures on DXA spinal images. RFRV-CLMs are applied to annotate vertebral landmark points in a dataset of 320 such images; this is the first time that RFRV-CLMs have been applied to this task. Extensive experiments were performed to investigate the effect of choice of parameters on performance. The results were compared to those achieved in [6] using AAMs on the same dataset, and show that RFRV-CLMs provide a considerable (66%) reduction in large errors across all vertebral classifications.

2 Method

The reader is referred to [11, 12] for full details of the RFRV-CLM algorithm; the key points are described below.

Constrained Local Models CLMs [15] build on previous work on Active Shape Models (ASMs) [16] and AAMs [17], providing a method for matching the points of a statistical shape model to an image. They combine global shape constraints with local models of the pattern of intensities. Given a set of training images with manual annotations $\mathbf{x}_l$ of a set of n landmark points $l = 1...n$ on each, a statistical shape model is trained by applying principal component

analysis (PCA) to the aligned shapes [17]. This yields a linear model of shape variation, which represents the position of each landmark point l using

$$\mathbf{x}_l = T_\theta(\bar{\mathbf{x}}_l + \mathbf{P}_l \mathbf{b} + \mathbf{r}_l) \qquad (1)$$

where $\bar{\mathbf{x}}_l$ is the mean position of the point in a suitable reference frame, $\mathbf{P}_l$ is a set of modes of variation, $\mathbf{b}$ are the shape model parameters, $\mathbf{r}_l$ allows small deviations from the model, and T_θ applies a global transformation (e.g. similarity) with parameters θ.

To match the model to a query image, $\mathbf{I}$, the overall quality of fit Q, of the model to the image is optimised over parameters $\mathbf{p} = \{\mathbf{b}, \theta, \mathbf{r}_l\}$

$$Q(\mathbf{p}) = \Sigma_{l=1}^n C_l(T_\theta(\bar{\mathbf{x}}_l + \mathbf{P}_l \mathbf{b} + \mathbf{r}_l)) \quad \text{s.t.} \quad \mathbf{b}^T \mathbf{S}_b^{-1} \mathbf{b} \leq M_t \quad \text{and} \quad |\mathbf{r}_l| < r_t \quad (2)$$

where C_l is a cost image for the fitting of landmark point l, $\mathbf{S}_b$ is the covariance matrix of shape model parameters $\mathbf{b}$, M_t is a threshold on the Mahalanobis distance, and r_l is a threshold on the residuals. M_t is chosen using the cumulative distribution function (CDF) of the χ^2 distribution so that 98% of samples from a multivariate Gaussian of the appropriate dimension would fall within it. This ensures a plausible shape by assuming a flat distribution for model parameters $\mathbf{b}$ constrained within hyper-ellipsoidal bounds [17]. In the original work [17], C_l was provided by normalised correlation with a globally constrained patch model.

RF Regression Voting in the CLM Framework In RFRV-CLM, C_l in Eq. 2 is provided by voting with a Random-Forest (RF) regressor. In order to train the RF for a single landmark point, the shape model is used to assess the global pose, θ, of the object in each image by minimising $|T_\theta(\bar{\mathbf{x}}) - \mathbf{x}|^2$. Each image is resampled into a standardised reference frame by applying $\mathbf{I}_r(i, j) = \mathbf{I}(T_\theta^{-1}(i, j))$. The model is scaled so that the width of the reference frame of the mean shape is a given value, w_{frame}. Sample patches of area w_{patch}^2 are generated and features $\mathbf{f}_j$ extracted from them at a set of random displacements $\mathbf{d}_j$ from the true position in the reference frame, $T_\theta^{-1}(\mathbf{x}_l)$, where $\mathbf{x}_l$ is the position of the point in the image. Haar-like features [18] are used, as they have proven effective for a range of applications and can be calculated efficiently from integral images. The displacements $\mathbf{d}_j$ are drawn from a flat distribution in the range $[-d_{max}, +d_{max}]$ in x and y. To allow for inaccurate initial estimates of the pose and to make the detector locally pose-invariant, the process is repeated with random perturbations in scale and orientation of the pose estimate. A RF [13] is then trained, using a standard, greedy approach, with the feature vectors $\mathbf{f}_j$ as inputs and the displacements $\mathbf{d}_j$ as regression targets. Each tree is trained on a bootstrap sample of N_s pairs $\{(\mathbf{f}_j, \mathbf{d}_j)\}$ from the training data. At each node, a random sub-set of n_{feat} features are chosen from this sample, and a feature f_i and threshold t that best split the data into two compact groups are selected by minimising an entropy measure [11]

$$G_T(t) = G(\{\mathbf{d}_i : f_i < t\}) + G(\{\mathbf{d}_i : f_i \geq t\}) \quad \text{where} \quad G(\{\mathbf{d}_i\}) = N log|\Sigma| \quad (3)$$

where N is the number of samples and Σ is their covariance matrix. Splitting terminates at either a maximum depth, D_{max}, or a minimum number of samples, N_{min}. The process is repeated to generate a forest of size n_{trees}.

RFRV-CLM Fitting Fitting to a query image is initialised via an estimate of the pose of the model e. g. from a small number of manual point annotations or a previous model, providing initial estimates $\mathbf{b}$ and θ (see Sec. 3). Eq. 2 is then optimised as follows. The image is resampled in the reference frame using the current pose $\mathbf{I}_r(i,j) = \mathbf{I}(T_\theta^{-1}(i,j))$. Cost images C_l are then computed by evaluating a grid of points in $\mathbf{I}_r$ over a region of interest around the current estimate of each point; the grid size is defined by a search range $[-d_{search}, +d_{search}]$, and the cost images are calculated for all landmark points independently. At each point $\mathbf{z}_l$ in the grid, the required feature values are extracted and the RF regressor R_l applied. R_l then casts a vote into an 2D accumulator array $C_l = V_l$ using $V_l(\mathbf{z}_l + \delta) \to V_l(\mathbf{z}_l + \delta) + v$. Each leaf node of the RF contains the mean $\bar{\mathbf{d}}$ and covariance $\mathbf{S}_d$ of the random displacements $\mathbf{d}_i$ from the true point position, in the reference frame, of its training samples. This supports several voting styles (v, δ); a single, unit vote at $\bar{\mathbf{d}}$, or probabilistic voting by weighting with $|\mathbf{S}_d|^{-0.5}$, or by casting a Gaussian spread of votes $N(\bar{\mathbf{d}}, \mathbf{S}_d)$.

The point positions are re-estimated by finding the lowest cost point within a disk of radius r of the current position in each cost image, applying the shape model and moving $\mathbf{b}$ to nearest valid point on the limiting ellipsoid if the shape constraint in Eq. 2 is violated, updating all point positions using $\mathbf{x}_l \to T_{\theta_r}(\bar{\mathbf{x}}_l + \mathbf{P}_l\mathbf{b} + \mathbf{r}_l)$, and iterating whilst reducing $r \to k_r r$. The initial disk radius r_{max} was set to the search range d_{search}, the search was terminated at $r_t = 1.5$ pixels (in the reference image), and k_r was set to 0.7. The optimisation is described in full in Algorithm 1.

3 Evaluation

A series of experiments was performed in order to optimise the various free parameters and options of the RFRV-CLM for application to the task of vertebral localisation in DXA images, and to compare the results to those achieved in [6] using AAMs. In order to facilitate this comparison, the same dataset and performance metrics were used. The dataset consisted of 320 DXA VFA images scanned on various Hologic (Bedford MA) scanners, obtained from: a) 44 patients from a previous study [19]; b) 80 female subjects in an epidemiological study of a UK cohort born in 1946; c) 196 females attending a local clinic for DXA BMD measurement, and for whom the referring physician had also requested VFA (as approved by the local ethics committee). Manual annotations of 405 landmark points were available for each image, covering the thoracic vertebrae from T7 to T12 and the lumbar vertebrae from L1 to L4. Each of these vertebrae in each image was also classified by an expert radiologist into one of five groups (normal, deformed but not fractured, and grade 1, 2 and 3 fractures according to the Genant definitions [2]; see Fig. 1).

SHAPE MODEL AND POSE PARAMETER OPTIMISATION

Input: r_{max}, r_t, k_r, $\mathbf{x}_l$ and C_l $\forall 1 \le l \le n$

1. Set $r \to r_{max}$, $\theta_r \to$ Identity, $\mathbf{x}_l \to \bar{\mathbf{x}}_l + \mathbf{P}_l\mathbf{b}$, $\mathbf{r}_l = 0$
2. While $r \ge r_t$
   (a) For every feature point l, find the best point $\hat{\mathbf{y}}_l$ in a disk of radius r around the current estimate
   $$\hat{\mathbf{y}}_l \to \arg\max_{\mathbf{y}_l : |\mathbf{y}_l - \mathbf{x}_l| < r} C_l(\mathbf{y}_l)$$
   (b) Fit the shape model to these best points to estimate shape and pose parameters $\{\mathbf{b}, \theta_r\}$ by solving
   $$\hat{\mathbf{y}}_l = T_{\theta_r}(\bar{\mathbf{x}}_l + \mathbf{P}_l\mathbf{b})$$
   (c) If $\mathbf{b}^T\mathbf{S}_b^{-1}\mathbf{b} > M_t$ then move $\mathbf{b}$ to nearest valid point on limiting ellipsoid
   (d) Update all feature point positions using $\mathbf{x}_l \to T_{\theta_r}(\bar{\mathbf{x}}_l + \mathbf{P}_l\mathbf{b} + \mathbf{r}_l)$
   (e) Set $r \to k_r r$ with $0 < k_r < 1$
3. Transform the resulting feature point positions into the image frame using T_θ with $\theta \to \theta \circ \theta_r$

Algorithm 1: Iterative model matching procedure to estimate the shape and pose parameters in the reference frame, given a set of feature point based cost images C_l.

Following [11] and [12], 2-stage, coarse-to-fine RFRV-CLMs were used and, in common with [6], individual models were trained for each vertebra; each covered the target vertebra and its two neighbours, or one neighbour for the top (T7) and bottom (L4) vertebrae. Fitting of the first-stage model was initialised using the approximate centre points of each of the vertebrae covered by the model. These were calculated from the full set of manual annotations as the centroid of the two central points on the upper and lower vertebral end-plates. This approach was adopted in order to avoid the significant reduction in random error, compared to individual manual annotations, that would occur if the centre points were calculated as the centroids of all manual annotations on each vertebra. Second-stage fitting was initialised using the results from the first stage. In order to compensate for the aperture problem present when annotating points on an extended edge, errors on automatic annotations were calculated as the mean, over each vertebra, of the minimum Euclidean distances between the automatic annotations and a Bezier spline through the manual annotations. This was applied to the points on the central vertebra in each model i.e. no use was made of the multiple fits for each point; see Sec. 4 for further comments.

RFRV-CLM Parameter Optimisation The free parameters of the RFRV-CLM, as described in Sec. 2, were divided into two sets; RF structure parameters (n_{trees}, n_{feat}, N_{min} and D_{max}), and image parameters (w_{frame}, w_{patch}, d_{max} and d_{search})[1]. These were optimised empirically and, in order to limit processor time requirements, a sequential approach was applied across both parameters and stages i.e. each first-stage parameter was optimised independently without

---

[1] Parameters w_{patch}, d_{max} and d_{search} were defined in the reference image.

Fig. 1: Example DXA spinal images. (a,b,c) 405-point manual annotations. (d,e,f) automatic annotation of the L2 vertebra (using the L1-L3 model), using the fully optimised, 2-stage RFRV-CLM. Example (a,d) shows grade 2 fractures on L2 and L3, (b,e) show a grade 3 fracture on L1, and (c,f) show a grade 3 fracture on L1 and a grade 1 fracture on L2.
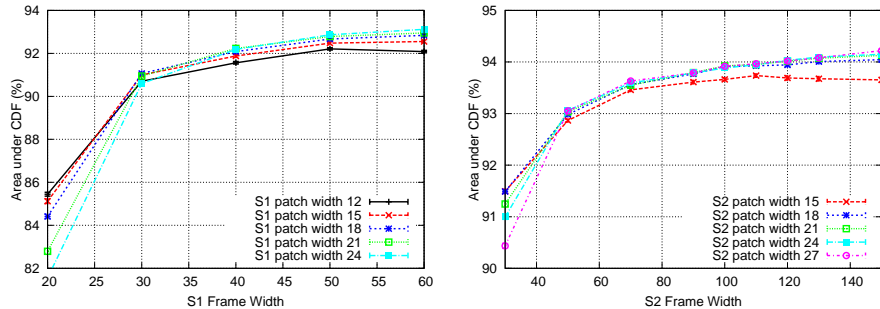
Fig. 2: Optimisation of the patch and frame width for first- (S1, left) and second-stage (S2, right) RFRV-CLMs on the L2 vertebra. Performance was measured as the %age area under the CDF of mean point-to-curve error across 160 test images. Error bars are too small to be visible at this scale.

applying the second-stage model; the optimal values were then fixed, and each second-stage parameter was optimised independently in a two-stage approach. Furthermore, the optimisation experiments were performed only on the L2 vertebra (i. e. the L1-L3 triplet model). L2 was chosen as it was the least obscured by confounding bony structures (ribs, scapulae, the iliac crest etc.) and imaging artefacts, minimising the contamination of the results with fitting failures. The optimisation results were then extended to the other vertebral levels by scaling the optimised w_{frame} using data on mean vertebral heights from [20], such that all image-based parameters were scaled. RF training includes a stochastic element, both in the random selection of data used to train each tree, and the random sub-selection from that data at each node. Therefore, each experiment was repeated five times in order to evaluate random errors.

For the sake of brevity, complete results are reported only for w_{frame} and w_{patch}, the parameters showing the greatest effect on performance; these are shown in Fig. 2. The graphs show the proportional area under the CDF of mean point-to-curve error on the L2 vertebra across the 160 test images. Performance generally increased with first-stage w_{frame}; however, d_{search} and w_{patch} are both defined in the reference frame, and so reducing w_{frame} increases the capture range. Therefore, w_{frame} was set using the point at which the performance increase ceased to be significant, giving 40 and 110 pixels for the first and second stages. Varying w_{patch} had a smaller effect on performance over most of the range tested, but showed more complex behaviour; values of 18 and 21 pixels were selected for the first and second stages, respectively, since these were either optimal or close to optimal over large portions of the tested ranges of w_{frame}. The remaining parameters were optimised in a similar way, giving (first stage, second stage): $n_{trees} = 2, 15$; $n_{feat} = 100, 200$; $N_{min} = 1, 1$; $D_{max} = 30, 30$; $w_{frame} = 40, 110$ pixels; $w_{patch} = 18, 21$ pixels; $d_{max} = 15, 15$ pixels; $d_{search} = 15, 10$ pixels (all pixel units, except for w_{frame} itself, were defined in the reference frame).
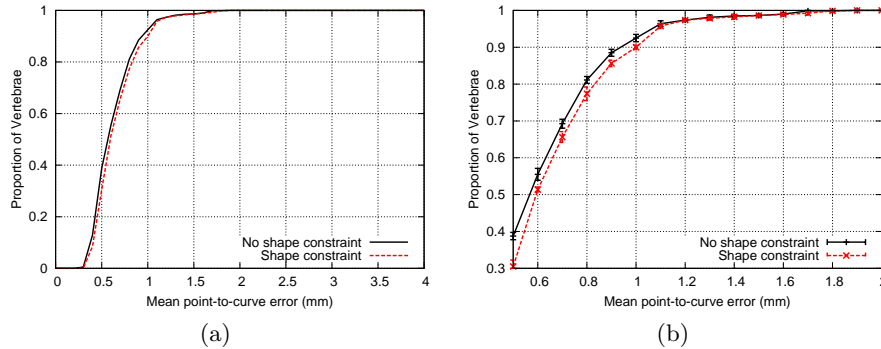
Fig. 3: Evaluation of the effect of the shape constraint in the final model stage using the L2 vertebra. Error bars are given as the standard deviation across five repeat experiments. Figure (b) shows an expanded view of Figure (a).

In general, the dependence of performance on parameter settings was weak over large ranges for most of the parameters.

Effects of the Shape Model Constraint A concern with the application of shape models is that the shape constraint may introduce a bias towards the mean of the training data; this is sub-optimal in clinical applications, where the pathological (i. e. outlying) cases are of most interest. In order to evaluate the effect, experiments were performed using the procedure and optimal parameters described above, both with and without the application of the shape model constraint during the fitting of the second stage of the RFRV-CLM, such that the shape constraint only aided in the approximate location of the global optimum; the final result was based on image information alone. The results are shown in Fig. 3 as the CDF of mean point-to-curve error for the L2 vertebra over the 160 test images. They show that the elimination of the shape model constraint resulted in a small and statistically insignificant change in performance, indicating that any shape model bias had an insignificant effect on the results given the error measure used here.

Effects of the Voting Style As described in Section 2, several methods for voting into the accumulator arrays were available. These alternatives were evaluated using the experimental procedure and optimised parameters described above, in order to determine whether probabilistic voting provided performance enhancements. The results are shown in Figure 4 as the CDF of mean point-to-curve error for the L2 vertebra over the 160 test images.

The results show that probabilistic voting provides no performance advantage. The performance of Gaussian voting was almost identical to that of single, unit voting. Single, weighted voting resulted in a small decrease in performance;

Fig. 4: Evaluation of various voting strategies for the L2 vertebra. Error bars are given as the standard deviation across five repeat experiments. Figure (b) shows an expanded view of Figure (a).

however, these differences were not statistically significant. Similar results have previously been found when applying RFRV-CLMs to facial images [11].

Performance across Multiple Vertebrae A set of leave-1/4-out experiments was performed to evaluate the RFRV-CLM on all vertebrae between T7 and L4 in all 320 images. The optimised parameters were derived from the L2 vertebra; they were adapted for the other vertebrae by scaling w_{frame} according to the ratio of mean vertebral heights in normal subjects from [20]. Shape model constraints were applied in all stages and single, unit voting was used.

The results are shown in Fig. 5. Annotation accuracy for T7 and L4 is lower than that for the other vertebrae, reflecting the fact that the corresponding models cover only two vertebrae, rather than three. The performance also decreases with increasing vertebral level; this may reflect the smaller size of the higher vertebrae and the presence of confounding bony structures (ribs, scapulae; see Fig. 1) in the thoracic region. However, mean errors of less than 2mm and 4mm were achieved for 95% and 99% respectively of the vertebrae at all levels. The results divided according to vertebral status show that, as expected, performance decreases with increasing severity of fracture i. e. increasing deformation relative to the mean shape. However, mean errors of <4mm were achieved for 95% of grade 3 fractures, and 100% of other classifications.

Table 1 provides numerical performance measures for the RFRV-CLMs, and compares them to the state-of-the-art results reported in [6], which applied AAMs to the same task and dataset. The AAM achieves better performance at the lower end of the CDF, as indicated by lower median errors, indicating smaller random errors on individual points. However, the RFRV-CLM achieves better mean errors for the more severely fractured vertebrae, and substantially lower numbers of vertebrae with mean errors >2mm regardless of classification,

Fig. 5: Evaluation of optimised, two-stage RFRV-CLMs on all vertebrae between T7 and L4. (Left) CDF of mean point-to-curve error for each vertebral level. (Right) CDF of mean point-to-curve error for each vertebral status.

Vertebra Status	AAM [6]				RFRV-CLM			
	%age of Sample	Mean (mm)	Median (mm)	%age Errors $> 2mm$	%age of Sample	Mean (mm)	Median (mm)	%age Errors $> 2mm$
Normal	84.9%	**0.55**	**0.40**	2.5%	83.25%	0.59	0.52	**0.68%**
Deformed	-	-	-	-	4.38%	0.72	0.57	3.57%
Grade 1	5.9%	**0.70**	**0.49**	4.8%	3.16%	0.73	0.60	**0.99%**
Grade 2	5.1%	0.92	**0.61**	10.2%	4.06%	**0.91**	0.80	**3.84%**
Grade 3	4.1%	1.19	**0.72**	16.5%	3.28%	**1.11**	0.90	**10.48%**

Table 1: Statistics of the mean point-to-curve errors on each vertebra for AAM and RFRV-CLM; bold numbers are the best results for each statistic/status.

indicating better performance at the higher end of the CDF i.e. a smaller number of fit failures (errors >2mm on 3.6% of all vertebrae for AAMs vs. 1.2% for RFRV-CLMs). Mean search time was 366 ms per model per image on a Dell Precision workstation with 2 Intel Xeon 5670 processors and 24 GB RAM, running OpenSuse 11.3 x64 (Linux kernel 2.6.34), using a single core.

4 Discussion and Conclusions

This paper has compared the performance of multi-stage, RFRV-CLMs to that of AAMs in the task of vertebral landmark point annotation in DXA spinal images. Several preliminary experiments were performed in order to optimise the various free parameters and options of the algorithm. In particular, no significant performance differences were observed, for the error metrics used here, either when implementing fully probabilistic regression voting or when eliminating the shape model constraint in the final stage of fitting, such that the result was

driven by image information alone. Application of fully optimised models to ten vertebral levels in 320 DXA spinal images showed that, whilst the AAM had smaller random errors, the RFRV-CLM produced lower mean errors for the more severe fracture classifications, and produced lower numbers of vertebrae with errors >2mm for all classifications. For comparison, the Genant method for vertebral fracture classification [2] defines grade 1, 2 and 3 fractures as vertebral height reductions of 20 to 25%, 25 to 40%, and >40%, respectively, and [20] measured mean vertebral heights varying from 22.97 ± 1.52mm for T7 to 35.62 ± 2.21mm for L4 in a sample of 108 normal women. Therefore, vertebral fractures are defined via height reductions of $\gtrsim 5mm$ regardless of grade or vertebral level. The results therefore indicate that, in the region of the CDF important for computer-aided VFA, RFRV-CLMs out-perform AAMs.

In this work, no use was made of the multiple fits to each vertebra provided by the overlaps of the models; only the points on the central vertebra in each were used. In future work, we intend to explore the combination of the multiple fits with goodness-of-fit measures both to improve the accuracy of the automatic annotation, in terms of random error, and to detect instances of fit failures i.e. systematic errors in individual fits.

References

1. Rachner, T.D., Khosla, S., Hofbauer, L.C.: Osteoporosis: now and the future. The Lancet **377**(9773) (2011) 1276–1287
2. Genant, H.K., Wu, C.Y., Kuijk, C.V., Nevitt, M.C.: Vertebral fracture assessment using a semi-quantitative technique. J Bone Miner Res **8**(9) (1993) 1137–1148
3. Cummings, S.R., Melton, J.: Epidemiology and outcomes of osteoporotic fractures. The Lancet **359**(9319) (2002) 1761–1767
4. Delmas, P.D., van de Langerijt, L., Watts, N.B., Eastell, R., Genant, H.K., Grauer, A., Cahall, D.L.: Underdiagnosis of vertebral fractures is a worldwide problem: the IMPACT study. J Bone Miner Res **20**(4) (2005) 557–563
5. Roberts, M.G., Cootes, T.F., Adams, J.E.: Vertebral morphometry: semi-automatic determination of detailed shape from DXA images using active appearance models. Investigative Radiology **41**(12) (2006) 849–859
6. Roberts, M.G., Cootes, T.F., Adams, J.E.: Automatic Location of Vertebrae on DXA Images Using Random Forest Regression. In: Proc. MICCAI 2012. Volume 7512 of LNCS., Springer-Verlag, Berlin (2012) 361–368
7. Leech, J.A., Dulberg, C., Kellie, S., Pattee, L., Gay, J.: Relationship of lung function to severity of osteoporosis in women. Am Rev Respir Dis **141**(1) (1990) 68–71
8. Dantone, M., Gall, J., Fanelli, G., Van Gool, L.: Real-time facial feature detection using conditional regression forests. In: Proc. CVPR. (2012) 2578–2585
9. Donner, R., Menze, B., Bischof, H., Langs, G.: Fast anatomical structure localization using top-down image patch regression. In: Medical Computer Vision. Recognition Techniques and Applications in Medical Imaging. Volume 7766 of Lecture Notes in Computer Science. (2013) 133–141
10. Valstar, M.F., Martinez, B., Binefa, X., Pantic, M.: Facial point detection using boosted regression and graph models. In: Proc. CVPR. (2010) 2729–2736

11. Cootes, T.F., Ionita, M.C., Lindner, C., Sauer, P.: Robust and Accurate Shape Model Fitting Using Random Forest Regression Voting. In: Proc. ECCV, Springer-Verlag, Berlin (2012) 278–291
12. Lindner, C., Thiagarajah, S., Wilkinson, J.M., arcOGEN Consortium, T., Wallis, G.A., Cootes, T.F.: Fully Automatic Segmentation of the Proximal Femur Using Random Forest Regression Voting. IEEE Trans. Med. Imag. **32**(8) (2013) 1462–1472
13. Breiman, L.: Random Forests. Machine Learning **45** (2001) 5–32
14. Sauer, P., Cootes, T., Taylor, C.: Accurate Regression Procedures for Active Appearance Models. In Hoey, J., McKenna, S., Trucco, E., eds.: Proc. BMVC. (2011) 30.1–30.11
15. Cristinacce, D., Cootes, T.: Automatic Feature Localisation with Constrained Local Models. Journal of Pattern Recognition **41**(10) (2008) 3054–3067
16. Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J.: Active Shape Models - Their training and application. Computer Vision and Image Understanding **61**(1) (1995) 38–59
17. Cootes, T.F., Edwards, G.J., Taylor, C.J.: Active appearance models. IEEE Transactions on Pattern Analysis and Machine Intelligence **23** (2001) 681–685
18. Viola, P., Jones, M.: Rapid object detection using a boosted cascade of simple features. In: Proc. CVPR, IEEE Computer Society (2001) 511–518
19. McCloskey, E., Selby, P., de Takats, D., Bernard, J., Davies, M., Robinson, J., Francis, R., Adams, J., Pande, K., Beneton, M., Jalava, T., Loyttyniemi, E., Kanis, J.A.: Effects of clodronate on vertebral fracture risk in osteoporosis: a 1-year interim analysis. Bone **28**(3) (2001) 310–315
20. Leidig-Bruckner, G., Minne, H.W.: The Spine Deformity Index (SDI); A New Approach to Quantifying Vertebral Crush Fractures in Patients with Osteoporosis. In: Vertebral Fracture in Osteoporosis. Osteoporosis Research Group, University of California (1995) 235–252