

Cervical Vertebral Corner Detection using Haar-like Features and Modified Hough Forest

S M Masudur Rahman Al Arif¹, Muhammad Asad¹, Karen Knapp², Michael Gundry² and Greg Slabaugh¹

¹ City University London

e-mail: {S.Al-Arif, Muhammad.Asad.2, Gregory.Slabaugh.1}@city.ac.uk

² University of Exeter

e-mail: {K.M.Knapp, mg361}@exeter.ac.uk

Abstract—The neck (cervical spine) is a flexible part of the human body and is particularly vulnerable to injury. Patients suspected of cervical spine injuries are often imaged using lateral view radiographs. Incorrect diagnosis based on these images may lead to serious long-term consequences. Our overarching goal is to develop a computer-aided detection system to help an emergency room physician correctly diagnose a patient's injury. In this paper, we present a method to localize the corners of cervical vertebrae in a set of 90 lateral cervical radiographs. Haar-like features are computed using intensity and gradient image patches, each of which votes for possible corner position using a modified Hough forest regression technique. Votes are aggregated using two dimensional kernel density estimation, to find the location of the corner. Our method demonstrates promising results, identifying corners with an average median error of 2.08 mm.

Keywords—Hough forest, random forest, classification, regression, cervical vertebrae, Haar-like features.

I. INTRODUCTION

The spine is one of the most important parts of the human body, as it not only provides support and movement for the body, it protects the spinal cord. Due to its degree of movement, the cervical spine, or neck region, is especially vulnerable to trauma. Cervical spine injuries (CSIs) can result from high energy impacts like automobile accidents, falls, dives into shallow water. Apart from these major accidents, minor injuries may also lead to CSI in elderly people and people with pre-existing bone abnormalities. However, evaluation of a cervical spine x-ray image is a major radiological challenge for an emergency physician, particularly those with less experience. Failure to establish a correct diagnosis may result in death or serious disability. Academic literature has reported up to 20% of CSI patients suffer tragic extension of their injuries due to delayed or missed diagnosis [1].

Early and accurate detection of a CSI is critical to plan appropriate treatment. Despite standardisation and advances in imaging, missed or delayed diagnosis of cervical spine injuries is still a common problem in emergency departments. In one study [1], the most common cause (accounting for 44%) of missed cervical spine injuries was misinterpretation of the images. Another study [2] resulted in a similar number (47%) of missed or delayed diagnosis due to misinterpretation. In most cases, junior staff responsible for initial radiologic examination failed to diagnose the injuries until experienced staff later performed a second evaluation of the radiographs. In [1], complications attributed to delayed or missed diagnosis ranged from motor and/or sensory neurologic deficits to complete quadriplegia. In other studies, 67% of patients with missed cervical fractures suffered neurological deterioration and nearly 30% of delayed CSI diagnosis developed permanent neurological deficit [3].

Therefore it is crucial to assist physicians in interpreting lateral radiographic images of the cervical spine through the state of the art advances in medical image processing. Our overarching goal is to develop a computer-aided detection system to assist the emergency

room physician with interpretation of cervical spine radiographs. As an important step towards this goal, this paper presents a semi-automatic method to detect vertebral corners.

Although computer analysis of radiographic images is difficult due to high amount of noise and low contrast, the literature includes some work to address similar issues in different imaging modalities. Tezmol et al. [4] proposed a Hough transform based method to find the global position, orientation and size of the cervical vertebral column in an x-ray image. Their method applies a brute force search using a template for all possible orientations, positions and sizes and votes in a four dimensional Hough space. The maximum vote indicates the best match to determine the orientation and position. Klinder et al. [5] proposed a 3D atlas based method that locates the whole vertebral column in CT scan. Their method then applies an active appearance model (AAM) to achieve the segmentation. It also performs identification of different vertebrae of the spine, with a reported 70% to 85% success rate for the cervical region. A probabilistic graphical model has been utilized in [6] to perform identification of the cervical vertebrae, along with their coarse size and orientation. Glocker et al. [7] also present results using CT scans, applying a regression forest to localize vertebrae centres in arbitrary CT images. A hidden Markov model is utilized to refine the results, achieving promising results. Larhamam et al. [8] used template matching based Hough transform to detect cervical vertebra centres, reporting a maximum detection of 97% for the C3 vertebra. In [9], this approach was improved utilizing k-means clustering techniques. Most of the work described above is based on 3D CT scans. In [10], Haar-like features are utilised in an AdaBoost classifier to estimate the position of the cervical vertebra in lateral X-ray images. Then a multi-resolution AAM is initialized on this position to segment the vertebrae, with results presented on the NHANES II dataset. In a similar vein, Benjelloun et al. [11, 12] apply an active shape model (ASM) to achieve the segmentation, however this method is sensitive to the initialization of the mean shape, based on a coarse orientation based on two user click points. However on more complex data with patients that have different neck articulations, as well as degenerative change poses a challenge. Vertebral edges are found using a Mahalanobis distance to achieve excellent segmentation results on the NHANES II dataset.

In this work, a two stage process is employed to find the corners of each cervical vertebra. First, the coarse orientation and size of each vertebra are calculated based on user click points. Image patches are then created, and for each, Haar-like features [13] are computed using the distribution of image intensity and gradient in these image patches. A modified Hough Forest [14], a variant of the Random Forest [15, 16] technique, is used to train a forest that performs both classification and regression to locate a possible corner location given a test image patch. The contributions of this work include our patch creation method based on a vertebral shape, application of Hough forest based on Haar-like features to the problem of vertebral corner detection in cervical radiographs, and a novel filtering method to localise the corner. The proposed algorithm is tested on a dataset of 90 radiographs, with an

average median error of 2.08 mm, demonstrating promising performance.

II. METHODOLOGY

A. Data

A total of 90 lateral cervical vertebrae radiographs form the dataset for this work. The images were provided by the Royal Devon & Exeter Hospital in collaboration with the University of Exeter. All scans were performed in 2014; the age of the patients varied from 17 to 96. Different (Philips, Agfa, Kodak, GE) radiographic systems were used to generate the images. Pixel spacing varied from 0.1 to 0.194 pixel per millimetre. Some images from the dataset are shown in Fig. 1. The dataset is very challenging, as it contains images with varying contrast as well as degenerative change and injury. The data is anonymised and standard protocols have been followed to use the images for research purposes. Each of the vertebrae in these images were manually annotated by an expert radiographer using a MATLAB GUI. A radiographer manually clicked on four corner points from which the centre of each vertebra was computed (as the centroid). Some manual annotations are shown in Fig. 2.

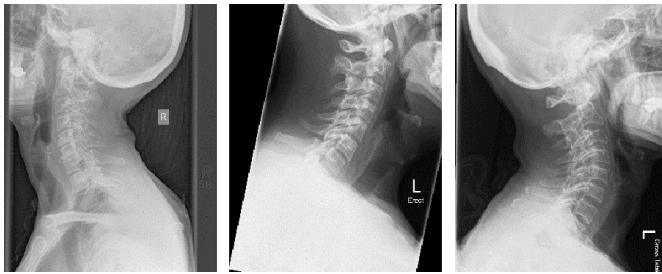


Fig. 1. Lateral cervical radiographs

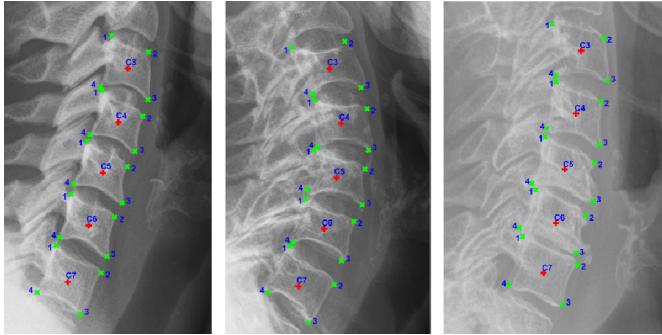


Fig. 2. Manually annotated vertebra corners and centers

B. Coarse Orientation and Size

The coarse orientation and size of the vertebrae are computed using the centre points. For each vertebra a vector is drawn from its centre to the centre of the vertebra above (\mathbf{F}_u) and below (\mathbf{F}_d). Then the orientation vector for that vertebra is the average of these vectors.

$$\mathbf{F} = (\mathbf{F}_u - \mathbf{F}_d)/2 \quad (1)$$

In case of the top vertebra $\mathbf{F} = -\mathbf{F}_d$ and for the bottom vertebra $\mathbf{F} = \mathbf{F}_u$.

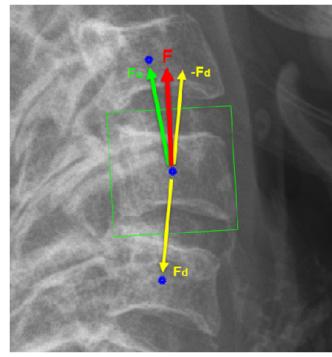


Fig. 3. Orientation and ROI

The vertebra sizes in pixels vary considerably among images due to the difference in spatial resolution of the images. The size also varies in millimetres from patient to patient because of natural variation amongst the human population. In order to normalize these differences the magnitude N of the vector \mathbf{F} is used as normalization constant, i.e. $N = |\mathbf{F}|$. N also represents the coarse size of the vertebra.

C. Image Patch, Class Label and Response Vector

Using the orientation vector \mathbf{F} and normalization constant N , a bounding box is generated to identify a region of interest (ROI) around the vertebra centre (green box in Fig. 3). This ROI is then divided into 16 image patches. Patches in the vertebra's centre are not considered as they contain homogeneous intensity distributions. Each of the boundary patches are assigned a class label and five vectors (see Fig. 4). The class labels (from 1 to 12 clockwise) abstractly encode the relative position of the patch with regard to the vertebra centre. Vector \mathbf{d}_2 represents the patch centre from the vertebra centre. Four response vectors (\mathbf{d}_{1P1} , \mathbf{d}_{1P2} , \mathbf{d}_{1P3} and \mathbf{d}_{1P4}) point to the four corners from the patch centre.

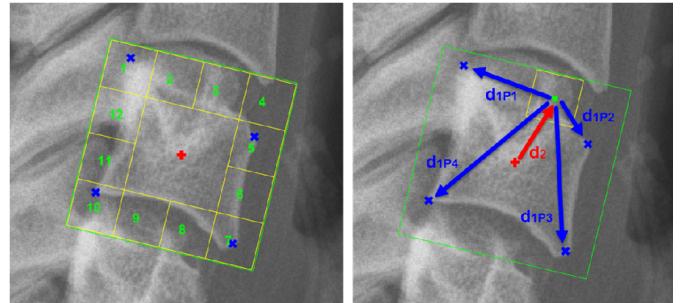


Fig. 4. Class labels and response vectors

For this work, these image patches are converted into feature vectors, \mathbf{H}_v , by using Haar-like features (see Fig. 5).

$$\mathbf{H}_v = [f_1, f_2, f_3, f_4, \dots, f_{10}] \quad (2)$$

The Haar-like features are chosen carefully to capture the variation seen in our vertebrae data. To calculate feature values from an image patch, the normalized intensity or gradient at each pixel is accumulated; the average signal intensity of the darker part is subtracted from the bright part.

$$f_i = \bar{I}_{\text{shaded}(i)} - \bar{I}_{\text{bright}(i)}, \text{ where } \bar{I}_x \text{ is the average intensity of area } x. \quad (3)$$

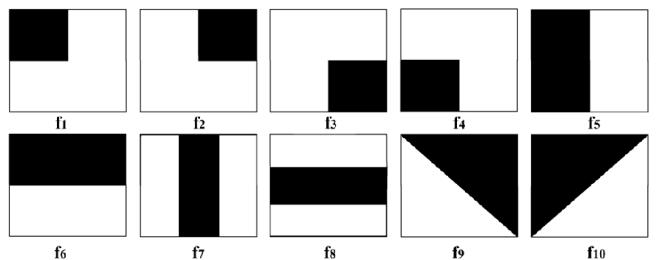


Fig. 5. Haar-like features.

D. Training Hough Forest

Hough forest performs both classification and regression. In our work, class labels from the image patches are used for classification and \mathbf{d}_1 vectors for different corners are used for regression. Each forest is trained only for one corner, and we use 100 trees in each forest. Each tree only accesses a random 25% of all the training image patches. At each split node one feature variable is randomly chosen and 20 randomly chosen thresholds are considered to split the data into its child nodes. The split occurs when it maximizes the information gain (IG) (Eqn. 4). Each branch of a tree terminates at a leaf node when maximum

depth, D_{max} , is reached or number of elements in a node becomes less than a threshold N_{min} .

The IG is calculated using the equation below [15]:

$$IG = H(S) - \sum_{i \in \{L,R\}} \frac{|S^i|}{|S|} H(S^i) \quad (4)$$

Where S is the set of data arriving at the split node, S^L is the data to be sent in the left child node, S^R is the data to be sent to the right child node and $H(S)$ is the entropy of the data S . $H(S)$ can be classification entropy or regression entropy. In Hough forest, the entropy is chosen randomly at each node between classification and regression.

1) Classification entropy: The classification entropy is calculated based on the output class labels at the considered node.

$$H(S) = -\sum_{c \in C} p(c) \log(p(c)) \quad (5)$$

Here C is the set of classes available at the node considered. In our case $C \subset \{1, 2, 3, 4, \dots, 12\}$.

2) Regression entropy: The regression entropy is computed based on the \mathbf{d}_1 vectors at that node.

$$H(S) \propto \log(|\Lambda(S)|) \quad (6)$$

Where S is the set of \mathbf{d}_1 vector arriving at the node and $\Lambda(S)$ is the covariance matrix S .

D. Prediction

After training, each leaf node contains a set of classes and corresponding \mathbf{d}_1 vectors. During testing, a vertebrae is divided into the 12 image patches (Fig. 4) following user clicks on the centre of each vertebrae. All 12 of these image patches are considered as test patch. The image patch is then converted into feature vector and fed into the forest. Each tree in a forest then returns a leaf node. Based on the training classes present on that leaf node, a class for the tree is predicted based on Eqn. 7.

$$\hat{C}_{tree} = \arg \max_c p(c|H_v) \quad (7)$$

These tree predictions are used to finally predict the final class from the forest.

$$\hat{C}_{forest} = \arg \max_{\hat{C}_{tree}} p(\hat{C}_{tree}|H_v) \quad (8)$$

A novel filtering process is then introduced where all the \mathbf{d}_1 vectors belonging to class *other* than \hat{C}_{forest} are discarded. Using the remaining vectors at each leaf node, a 2D kernel density estimation (KE2D) is initiated to determine $\hat{\mathbf{d}}_{1tree}$. The process is repeated to estimate $\hat{\mathbf{d}}_{1forest}$ from all the trees in the forest, which is essentially predicted vector, $\hat{\mathbf{d}}_{1patch}$, for the considered image patch.

$$\hat{\mathbf{d}}_{1tree} = KE2D(\hat{\mathbf{d}}_{1filtered}); \hat{\mathbf{d}}_{1filtered} = \{\forall \mathbf{d}_1 | c = \hat{C}_{forest}\} \quad (9)$$

$$\hat{\mathbf{d}}_{1patch} = KE2D(\{\hat{\mathbf{d}}_{1tree(1)}, \hat{\mathbf{d}}_{1tree(2)}, \hat{\mathbf{d}}_{1tree(3)}, \dots, \hat{\mathbf{d}}_{1tree(N)}\}) \quad (10)$$

where N is number of trees in the forest. Finally the same is done for all the 12 patches. And each patch votes for a possible corner position, $\hat{\mathbf{D}}_{patch}$, with respect to the vertebra centre.

$$\hat{\mathbf{D}}_{patch} = -\mathbf{d}_{2patch} + \hat{\mathbf{d}}_{1patch} \quad (11)$$

where vector \mathbf{d}_2 is known and used to generate the ROI and patches at the start of the prediction process. These 12 possible predicted corners votes are then fed into 2D kernel estimation, to compute the final prediction of the corner, $\hat{\mathbf{D}}$.

$$\hat{\mathbf{D}} = KE2D(\{\hat{\mathbf{D}}_{patch(1)}, \hat{\mathbf{D}}_{patch(2)}, \dots, \hat{\mathbf{D}}_{patch(12)}\}) \quad (12)$$

D. 2D Kernel Density Estimation

Two dimensional kernel density estimation is applied to find the single estimated vector location from a collection of vectors. All the vectors are first located in the 2D plane. Then a normalized 2D Gaussian distribution is fit to each vector location. The distributions are added together and maxima of the aggregated distribution is localized [13]. This location is returned as the predicted vector. An example of an aggregated distribution for the corner detection process is graphically shown in Fig. 6.

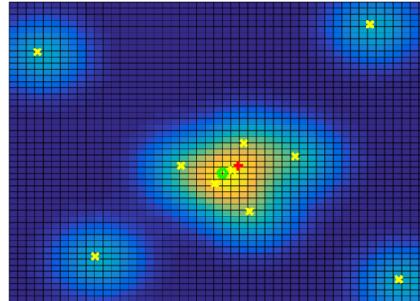


Fig. 6. The heat map visualizes confidence of the corner estimation. Yellow crosses are filtered corners from leaf patches, the red plus is the corner from manual demarcation and the green circle is the predicted corner.

III. EXPERIMENTS

In this work, we considered five cervical vertebrae, C3, C4, C5, C6 and C7. Each vertebral body has four corners, so in total 20 corners are studied for each patient. Each forest is trained only for one corner at a time. For each corner, we employed a 10 fold cross validation, where for each fold 9 images were withheld for independent testing and the remaining were used to train the algorithm. As described in Section II-C, image patches are converted into feature vectors by using ten Haar-like feature patches. As stated earlier, the image patch itself can be the normalized signal intensity and/or the gradient of image intensities. Based on this we experimented with three different feature vector types.

1. Intensity: Haar-like features are applied on normalized intensity of the image patch.
2. Gradient: Haar-like feature are calculated using the gradient of the normalized intensity of the image patch.
3. Mixed: Both the intensity and gradient are used to create a concatenated feature vector.

IV. RESULTS & DISCUSSION

The predicted corner locations are compared with manually annotated corner location. The error, measured as Euclidian distance in millimeters, is calculated for each corner. Finally the mean, median and standard deviation (std) of the error are computed from all 90 radiographs and reported in Table 1. Graphical results are shown in Fig. 7.

The results from Table 1 indicate that the mixed features perform the best, followed by intensity based features and the gradient based features. The reason might be the choice of the Haar-like feature patches. Most of our selected Haar-like features (Fig. 5) bear similarity to the intensity within an image patch. For example image patches of class 1, 4, 7 and 10 have an abstract resemblance with first four Haar-like patches (f_1, f_2, f_3 and f_4), and a similar conclusion can also be made for next two (f_5 and f_6) Haar-like features with image patch class of 2, 3, 5, 6, 8, 9, 11 and 12. The next two Haar-like patches (f_7 and f_8) are more related with gradient image patches as they look like vertical and horizontal edges. As we have most of the Haar-like patches similar to intensity image patches and fewer patches similar to gradient image patches, it is understandable why intensity vector performs better than gradient vectors. Similarly when mixed vectors are used, useful

information from both the intensity and gradient distribution are now available for selection.

Another interesting point is that, the errors for C7 are lower than other vertebra. A possible explanation is that the coarse orientation works better for C7, so the bounding boxes for C7 are more consistent. Whereas for other vertebra, because of different degree of bending in the cervical spine, the orientation of the bounding box varies among different patients. As a result image patches from C7 from all images have better similarity.

Representative examples of the corner prediction for all the vertebrae are shown in Fig. 7. Qualitatively it can be said that errors less than 3 mm are good. A few misdetections are also shown in the last two columns of Fig. 7. Similar to other image analysis techniques with radiographic images, the algorithm sometimes suffers from low contrast or absence of edges in the vertebra (Fig. 7. C6-E). However, since our proposed method uses voting from multiple patches,

sometimes it can predict with high accuracy even though some patches have low contrast. (Fig. 7. C7-E). Also some patients have different and comparatively rare bone structure. For example, C5-D, E and C6-E have comparatively longer vertebrae. Detections for these cases yield higher error and misdetection.

Comparison with other work is difficult as, similar vertebral corner detection process on radiographs are not found in the literature. However, Glockner et al. [7] worked on CT images, where center of the vertebra is located in 3D and compared with manual annotations. They reported a lowest average median error of 6.14 mm for cervical region. Our previous work [17] yielded a lowest average median of 3.05 mm using a normalized patch itself as a feature vector with different class prediction criteria. In comparison, in this work we have been able to reduce the error further to 2.08 mm.

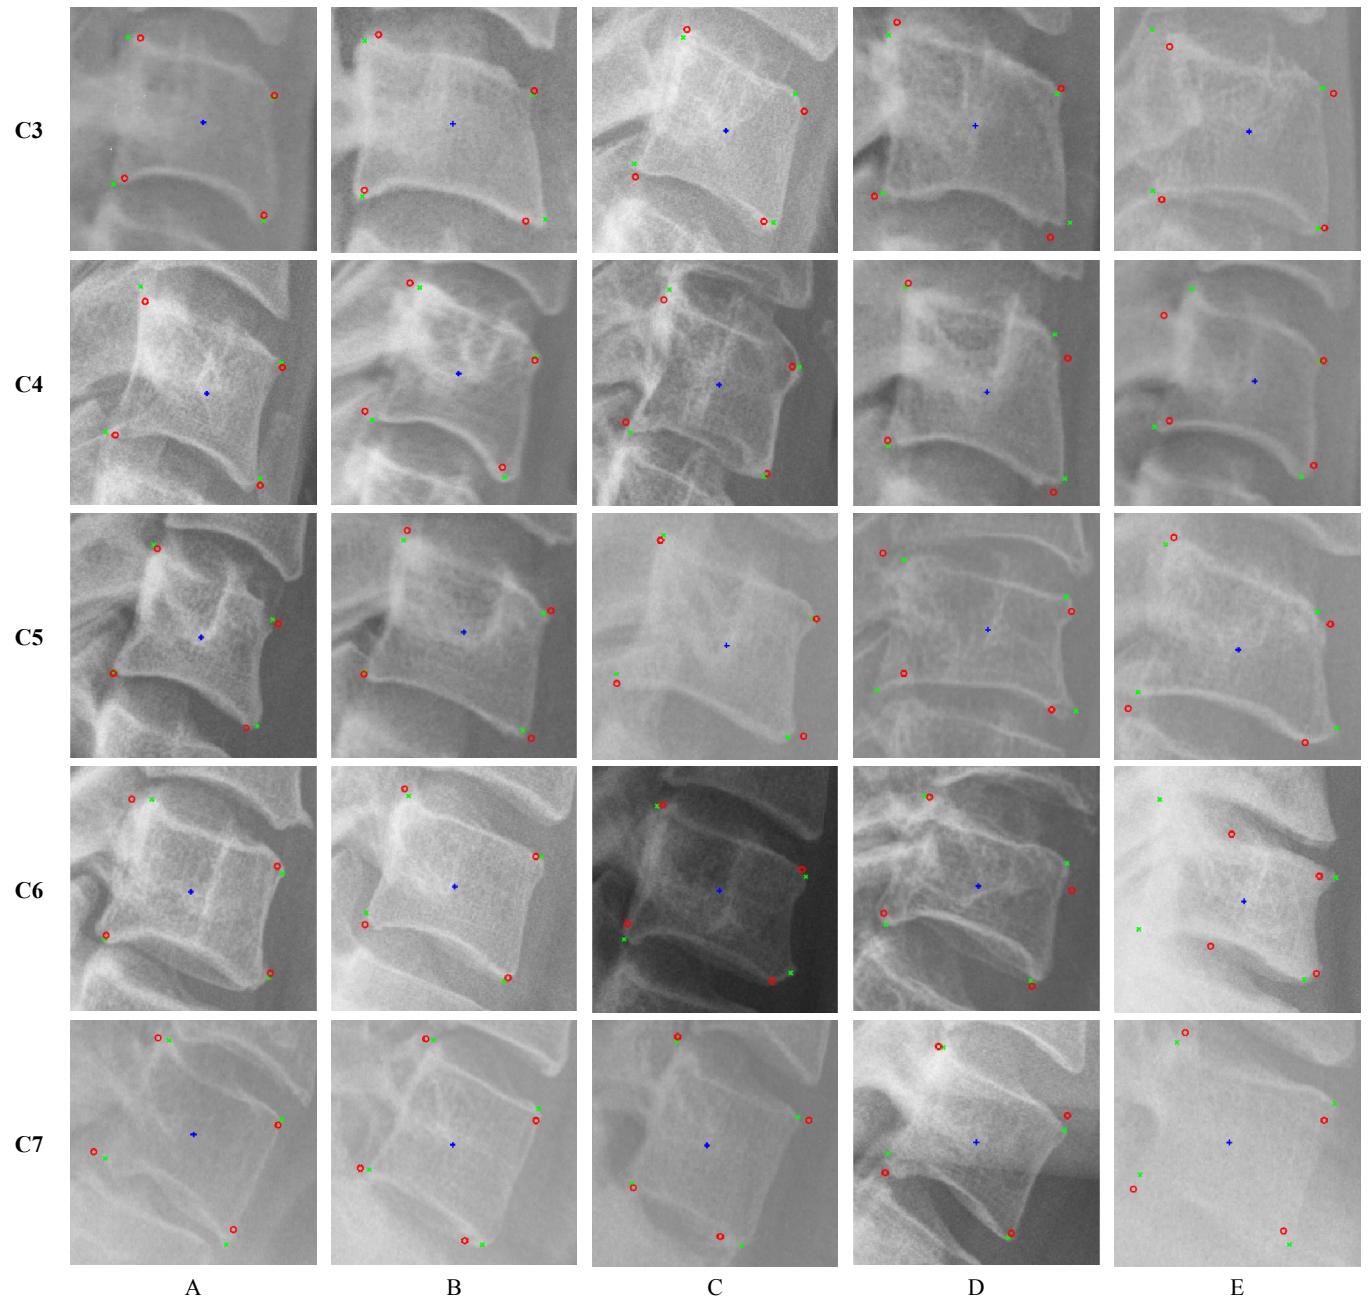


Fig. 7. Vertebral corner detection examples: the blue plus indicates vertebra centers, the red circle denotes predicted corner and green cross represents manually annotated corners

Table 1: Performance of different type of feature vectors

Vertebra	Corner	Type of Haar Features								
		Intensity			Gradient			Mixed: Intensity and Gradient		
		Median	Mean	Std	Median	Mean	Std	Median	Mean	Std
C3	1	2.49	3.30	3.06	2.78	3.96	3.73	2.07	2.55	2.14
	2	1.84	2.89	3.32	1.85	3.11	4.10	2.09	2.95	3.71
	3	2.49	3.37	3.77	2.92	4.72	4.82	2.51	3.66	3.71
	4	2.75	3.85	3.92	2.94	4.22	3.95	2.39	3.43	3.18
C4	1	2.57	3.64	3.55	3.36	4.63	4.47	2.14	3.38	3.44
	2	1.73	3.26	4.20	1.84	2.73	2.79	1.62	2.48	3.31
	3	2.01	3.28	3.64	2.79	4.18	4.28	2.21	3.17	2.99
	4	2.21	3.87	4.40	2.84	4.33	4.36	2.11	3.04	2.85
C5	1	2.10	3.28	2.83	2.78	4.96	5.01	2.21	3.52	3.46
	2	1.57	2.76	3.53	1.63	3.20	3.90	1.69	2.58	2.83
	3	2.53	3.82	3.87	2.91	4.88	4.68	2.45	4.15	4.89
	4	2.67	4.44	4.31	3.01	4.49	3.85	2.46	3.52	3.26
C6	1	2.66	3.71	3.48	2.75	4.01	3.72	2.42	3.04	2.63
	2	1.56	3.07	3.74	2.67	4.10	4.60	2.07	2.86	2.96
	3	2.76	4.07	3.62	3.67	5.44	5.09	2.87	4.09	4.12
	4	2.45	3.60	3.61	2.83	4.26	3.62	2.26	3.59	4.18
C7	1	1.23	2.00	2.23	2.66	3.37	2.88	1.13	1.77	2.00
	2	1.10	1.66	1.55	1.56	2.83	2.91	1.03	1.95	2.03
	3	1.61	2.26	2.04	2.33	3.32	2.74	2.08	2.63	2.25
	4	1.56	2.33	2.14	3.26	4.21	2.78	1.88	2.19	1.58
Average		2.10	3.22	3.34	2.67	4.05	3.91	2.08	3.03	3.08

IV. CONCLUSION

A semi-automatic cervical vertebral corner detection algorithm is proposed in this work. The proposed method is based on a modified Hough Forest architecture. In the process, a new model is introduced to generate image patches from vertebrae using manually clicked center points; whereas in original Hough Forest the patches are created randomly from positive examples. Haar-like feature are used in the framework. A novel two stage prediction method is introduced. First the class is predicted from a forest, then a filtering process is applied to exclude vectors belonging to other classes. Finally the corner position is predicted using a two dimensional kernel density estimation process.

The algorithm is tested on a dataset of 90 emergency room lateral cervical radiographs. The experiments are performed with a 10 fold cross validation scheme. Errors between the predicted corners and manually annotated corners are reported. A lowest average median error of approximately 2 mm is achieved.

In future, new features like displacement features [18] or bounding box features [19] can be used to compare their performance with Haar-like features. The next phase of the project is to segment the vertebra correctly. These predicted corners can be used to initialize mean vertebra models on the vertebra. Then a statistical shape model segmentation approach like Active Shape Model (ASM) [20] or Active Appearance Model (AAM) [21] can be utilized to achieve the final segmentation.

ACKNOWLEDGMENT

The authors would like to thank Andy Appelbaum and Adam Reuben from Royal Devon and Exeter Hospital, UK for their contributions to the project.

REFERENCES

- [1] P. Patrick, N. Hauswirth, M. Jaindl, S. Chatwani, V. Vecsei, and C. Gaebler. Delayed or Missed Diagnosis of Cervical Spine Injuries. *The Journal of Trauma: Injury, Infection, and Critical Care*, 61.1: 150-55, 2006.
- [2] J.W. Davis, L.P. David, B.H. David and R.C. Mackersie. The Etiology Of Missed Cervical Spine Injuries. *The Journal of Trauma: Injury, Infection, and Critical Care* 34.3: 342-46, 1993.
- [3] C.G.T. Morris and É. McCoy. Clearing the Cervical Spine in Unconscious Polytrauma Victims, Balancing Risks and Effective Screening. *Anaesthesia* 59.5: 464-82, 2004.
- [4] A. Tezmol, H. Sari-Sarraf, S. Mitra, R. Long and A. Gururajan. Customized Hough transform for robust segmentation of cervical vertebrae from X-ray images. *Fifth IEEE Southwest Symposium on Image Analysis and Interpretation*, pp. 224-228, New Mexico-USA, April 2002.
- [5] T. Klinder, J. Ostermann, M. Ehm, A. Franz, R. Kneser and C. Lorenz. Automated model-based vertebra detection, identification and segmentation in CT images. *Medical image analysis*, 13(3): 471-482, 2009.
- [6] X. Dong and G. Zheng. Automated vertebra identification from X-ray images. *Image Analysis and Recognition*. Springer Berlin Heidelberg, 1-9, 2010.
- [7] B. Glocker, J. Feulner, A. Criminisi, D.R. Haynor and Konukoglu, E. Automatic localization and identification of vertebrae in arbitrary field-of-view CT scans. *Medical Image Computing and Computer-Assisted Intervention-MICCAI*, Springer Berlin Heidelberg, pp. 590-598, 2012.
- [8] M.A. Larhamam, S. Mahmoudi, and M. Benjelloun. Semi-automatic detection of cervical vertebrae in X-ray images using generalized Hough transform. *3rd International Conference on Image Processing Theory, Tools and Applications (IPTA)*, pp. 396-401, Istanbul-Turkey, October 2012.
- [9] M.A. Larhamam, M. Benjelloun and S. Mahmoudi. Vertebra identification using template matching modelmp and K-means clustering. *International journal of computer assisted radiology and surgery* 9.2: 177-187, 2014.
- [10] X. Xu, H.W. Hao, X.C. Yin, N. Liu and S.H. Shafin. Automatic segmentation of cervical vertebrae in X-ray images. *The 2012 International Joint Conference on Neural Networks (IJCNN)*, pp. 1-8, Brisbane-Australia, June 2012.

[11] S. Mahmoudi and M. Benjelloun. A new approach for cervical vertebrae segmentation. *Progress in Pattern Recognition, Image Analysis and Applications*. Springer Berlin Heidelberg, 753-762, 2007.

[12] M. Benjelloun, S. Mahmoudi and F. Lecron. A framework of vertebra segmentation using the active shape model-based approach. *Journal of Biomedical Imaging*: 9, 2011.

[13] P. Viola and M. Jones. Rapid object detection using a boosted cascade of simple features. *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 1, pp. 511-518, Hawaii-USA, December 2001.

[14] J. Gall, A. Yao, N. Razavi and L.V. Gool. Hough forests for object detection, tracking, and action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 33.11: 2188-2202, 2011.

[15] L. Breiman. Random forests. *Machine learning* 45.1: 5-32, 2001.

[16] A. Criminisi and J. Shotton. Decision forests for computer vision and medical image analysis. *Springer Science & Business Media*, 2013.

[17] S.M.M.R. Al-Arif, M. Asad, K. Knapp, M. Gundry, and G. Slabaugh. Hough Forest-based Corner Detection for Cervical Spine Radiographs. *Medical Image Understanding and Analysis (MIUA '15), proceedings, Lincolnshire-UK, July 2015*.

[18] J. Shotton, T. Sharp, A. Kipman, A. Fitzgibbon, M. Finocchio, A. Blake, M. Cook and R. Moore. Real-time human pose recognition in parts from single depth images. *Communications of the ACM* 56.1: 116-124, 2013.

[19] A. Criminisi, D. Robertson, O. Pauly, B. Glocker, E. Konukoglu, J. Shotton, D. Mateus, A. Martinez Möller, S. G. Nekolla and N. Navab. Anatomy Detection and Localization in 3D Medical Images. *Decision Forests for Computer Vision and Medical Image Analysis*. Springer London, 193-209, 2013.

[20] T.F. Cootes, C.J. Taylor, D.H. Cooper and J. Graham. Active shape models-their training and application. *Computer Vision and Image Understanding* 61.1: 38-59, 1995.

[21] T.F. Cootes, G.J. Edwards and C.J. Taylor. Active appearance models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23.6: 681-685, 2001.