City Research Online

City St George's, University of London

Citation: Al-Arif, S. M., Asad, M., Knapp, K., Gundry, M. & Slabaugh, G. G. (2015). Hough Forest-based Corner Detection for Cervical Spine Radiographs. Paper presented at the Medical Image Understanding and Analysis Conference, MIUA 2015, 15-07-2015 - 17-07-2015, University of Lincoln, UK.

This is the published version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: https://openaccess.city.ac.uk/id/eprint/12276/

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to City Research Online policy.

City Research Online: http://openaccess.city.ac.uk/ publications@citystgeorges.ac.uk

Hough Forest-based Corner Detection for Cervical Spine Radiographs

S M Masudur Rahman Al-Arif[1]
S.Al-Arif@city.ac.uk
Mohammad Asad[1]
Karen Knapp[2], Micheal Gundry[2]
Greg Slabaugh[1]

[1] Department of Computer Science
City University
London, UK

[2] Department of Medical Imaging
University of Exeter,
Devon, UK

Abstract

The cervical spine (neck region) is highly sensitive to trauma related injuries, which must be analysed carefully by emergency physicians. In this work, we propose a Hough Forest-based corner detection method for cervical spine radiographs, as a first step towards a computer-aided diagnostic tool. We propose a novel patch-based model based on two-stage supervised learning (classification and regression) to estimate the corners of cervical vertebral bodies. Our method is evaluated using 106 cervical x-ray images consisting of 530 vertebrae and 2120 corners, which have been demarcated manually by an expert radiographer. The results show promising performance of the proposed algorithm, with a lowest median error of 1.98 mm.

1 Introduction

Evaluation of a cervical spine x-ray image is a pressing radiological challenge for an emergency physician, coupling subtle-to-detect pathologies with potentially severe clinical outcomes (neurological deficit, paralysis) [1]. Our overarching goal is to develop a computer aided detection (CAD) and analysis tool to assist human interpretation of lateral cervical spine radiographs. Towards this goal, we propose an algorithm to detect corners of cervical vertebral bodies. Applications of corner detection include analysing the alignment of the cervical spine [2], as well as initialising a segmentation technique, like an active shape model [3]. In this work, we propose a semi-automatic algorithm to find the corners of the cervical vertebra given limited input from a user. The knowledge of the possible location of the corners are learned from a training dataset and applied to find corners in test images. Our approach modifies the Hough Forest technique [4] based on a novel patch-based representation tailored to vertebral bodies imaged in a lateral view.

Although computerised analysis of radiographic images of the spine is difficult due to noise and low contrast, some related work appears in the literature. Tezmol et al. [5] proposed a Hough transform-based method to find the global position, orientation and size the cervical vertebrae in the image. Their method applies a brute force searching using a template for all possible combinations and votes in a four dimensional Hough space. Klinder et al. [6] proposed a 3D atlas-based method that can locate the whole vertebral column in CT scan and perform identification of different vertebrae of the spine with 70 to 85% success

© 2015. The copyright of this document resides with its authors.
It may be distributed unchanged freely in print or electronic forms.

in the cervical region. A probabilistic graphical model has been utilised in Dong et al. [7] to perform identification of the cervical vertebrae. Their work can also find a coarse size and orientation of the vertebrae based on the model. Glocker et al. [8] also worked on CT scans and applied a regression forest to localize vertebrae centres in arbitrary CT images. A hidden Markov model is also utilized to refine the results. Larhmam et al. [9] used template matching-based Hough transform to detect cervical vertebra centres, and reported a maximum detection of 97% for the C3 vertebra. In [10], Haar-like features are utilised with an Adaboost classifier to estimate the position, then perform a segmentation, of the cervical vertebrae in lateral X-ray images. Similar work has been done by Benjelloun et al. [3] based on active shape models (ASMs).

In our work, our goal is to find four corner points of each vertebral body. First, a coarse orientation is found based on user-provided clicks on the vertebrae centres. Then a machine learning approach is applied to locate the vertebral corners. Our work is based on the Hough Forest [4] algorithm, which is a variant of popular Random Forest [11] approach. Our contributions include a patch-based representation customised to vertebral body shapes. Each patch includes a class label, as well as vectors that point to the corner positions. A novel two stage prediction technique is introduced which utilizes both classification using patch labels and regression using corner vectors. Furthermore, we introduce an additional filtering stage inbetween the classification and regression prediction to aggregate votes robustly. Although the Hough Forest approach has been widely used in object detection and pattern recognition in outdoor images, in this paper we adapt the method for low contrast medical radiographs. To our knowledge, this is the first work that uses Hough Forest to determine the vertebrae corner locations.

2 Methodology

2.1 Data

Our data consists of 106 lateral cervical spine radiographs, acquired from Royal Devon & Exeter Hospital in association with University of Exeter. All imaging was performed in 2014, and the age of the patients varied from 17 to 96. Different (Philips, Agfa, Kodak, GE) radiograph systems were used for imaging. Resolution varied from 0.1 to 0.194 mm per pixel. The dataset is very challenging, as it contains normal, good contrast images to low contrast, abnormal images including patients with degenerative change. The data is anonymised and standard protocols have been followed for research purposes.

Each of these images was manually demarcated by an expert radiographer, who clicked on the corners of each vertebra. Examples are shown in Fig. 1.

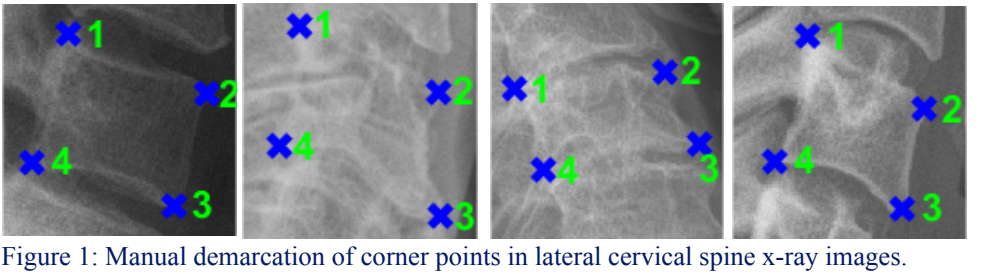

Figure 1: Manual demarcation of corner points in lateral cervical spine x-ray images.

2.2 Orientation and Normalization

Manually clicked vertebral centres are available as part of the CAD system being developed. For each vertebra, a vector is drawn from its centre to the centre of the vertebra above ($\mathbf{F}_a$) and below ($\mathbf{F}_b$). Then the orientation vector is the average of these vectors (see Fig. 2).

$$\mathbf{F} = (\mathbf{F}_a - \mathbf{F}_b)/2 \quad (1)$$

The vertebra size in pixels varies among images due to the difference in spatial resolution and patient size. In order to normalise these differences a normalisation constant, $N_s = |\mathbf{F}|$, will be used below.

2.3 Hough Forest

Hough Forest is a variant of Random Forest that performs classification and regression. It has been used to detect objects using small image patches. In this work, we exploit this idea to detect corners. Each vertebra is divided into N = 12 image patches as shown in Fig. 3a. This number balances the trade-off of having too few patches which may result in inadequate splitting in our decision trees, vs having too many patches where there are inadequate features due to a small patch size. First, a square region of interest (ROI) is defined around the centre of the vertebra such that the ROI ideally covers the whole vertebra (see Fig. 3b). The orientation and size of this ROI is determined by the vector **F** and normalisation constant Ns described above. Then the ROI is divided into 16 equal sized smaller regions, and the 12 boundary patches are taken. Non-boundary patches are not considered, as they usually contain a homogeneous intensity distribution. Each of these image patches is associated with two vectors and a class label. The label is the patch number (1 to 12) and abstractly encodes the patch position within the vertebra. Two vectors associated with each patch are: i) $\mathbf{d}_1$: the vector that points to the corner position and ii) $\mathbf{d}_2$: the vector that points to the centre of the vertebra. Both vectors originate from the patch centre. Figure 3 demonstrates the classes and vectors graphically. This approach differs from the Hough Forest technique, where the patches are created randomly from positive and negative example images, forming binary class labels. In contrast, our method creates patches in a more structured manner, where class labels abstractly denote the patch position.

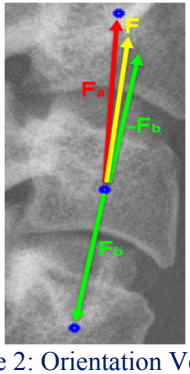

Figure 2: Orientation Vector

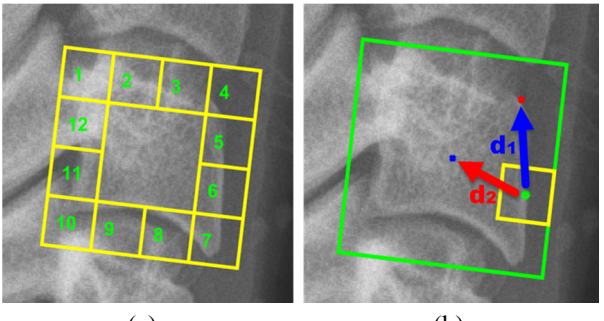

(a) (b)

Figure 3: Output Variables: (a) Patch Classes (b) Vectors

Separate forests are created for each corner. Once the patches are created for all the images, the patches are used to train the model. Each patch acts like a feature vector but instead of having a single output variable, we have a class label and a vector $\mathbf{d}_1$. The vector $\mathbf{d}_2$ is not used in training, however it is needed to back-project the predicted $\mathbf{d}_1$ on the image after testing. Standard expressions for entropy and information gains has been used [11].

2.4 Prediction

Once training is done, the testing is performed in a similar way. The user clicks at the centre the vertebrae of that image. Using this information the normalisation constant and orientation of each vertebrae is calculated, then each vertebra is divided into 12 image patches and each patch is then fed into the forest, where it ends up in a leaf node of a tree. The class is then calculated using Gaussian kernel density estimator for that tree. The maximal output is found, and the corresponding class is taken as the prediction of that tree. The process is shown in Fig. 4a. The estimator is applied to all the trees to find the final class prediction of the forest. The bandwidth of the Gaussians for the estimator is chosen empirically to be 1.0. This means Gaussians only from adjacent classes have high impact.

Based on the predicted class, a filtering process is initiated where we only consider the vectors $\mathbf{d}_1$ that are of the predicted class. Vectors belonging to other classes are discarded. For each of these vectors, a 2D Gaussian is fitted. All the distributions then summed up and normalized. The maxima of this distribution denotes the predicted $\mathbf{d}_1$ for the image patch. Figure 4b summarizes the process. Then this predicted $\mathbf{d}_1$ is added with the patch's $\mathbf{d}_2$ vector to point to the actual corner.

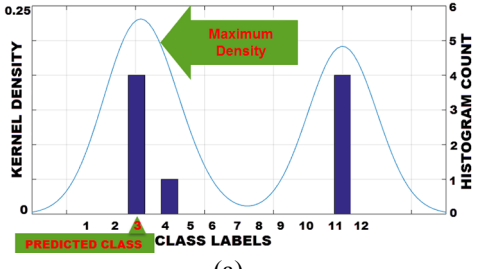

(a)

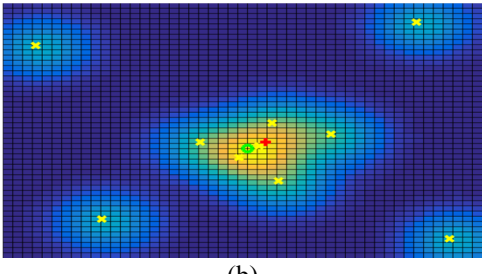

(b)

Figure 4: (a) Class prediction (b) Corner prediction: a heat map visualises confidence. Yellow crosses are filtered corners from leaf patches, the red plus is corner from manual demarcation and the green circle is the predicted corner.

3 Experiments

Experiments are conducted with a 10-fold cross validation scheme. Each time, 10% of all the images are chosen as test set and others are used for training the forest at each fold. This process is repeated until all the images are used as test image once. The Euclidean distances between the predicted corners and manually identified corners are computed. Then, the median, mean and standard deviation of these distances for all the cases are calculated. Different feature sizes and objective functions are considered to evaluate the performance of the algorithm.

A. **Feature Size:** Although the image sizes are normalised in mm, we experiment with different patch sizes in pixels, including 30 x 30, 10 x 10, 5 x 5 and 3 x 3.

Figure 5: Variation of patch sizes

B. **Objective function:** The standard Hough Forest chooses the objective function randomly between classification and regression. We performed: i) Randomly choosing between the Classification and Regression objective function, ii) Using only classification entropy and iii) Using only regression entropy.

4 Results and Discussion

Experiments have been performed in order to see which feature size yields the lowest possible error. The results are summarized in Table 1. The median, mean and standard deviation of the errors or distances are reported for each vertebra. The patch size is an important consideration. The lowest median error is found with the size 5 x 5. The reason for this can be understood, from the dimensionality of the feature vectors. A 30 x 30 image patch results in a feature vector of 900 values, which is too large for random feature selection to efficiently find discriminative features. Decreasing the patch size helps, however, as this gets too small, information is lost due to averaging (see Fig. 5). The average of the metrics show that a 5 x 5 patch size yields the lowest median error. Qualitatively, it can be said error less than 5 mm can be considered as a good prediction. The experiments were also repeated with gradient image patches, which had weaker performance with a lowest average median error of 3.66 mm for a 5 x 5 patch and results are withheld due to page constraints. Visual results are shown in Fig 6. Our corner detection approach is successful when the test vertebra is similar to the examples in the training data. But due to variation in anatomic data, errors occur in the corner prediction. This may be solved to some extent by increasing the training data, something we are currently working on. Test images from older patients are often very low contrast due to degenerative change. Disappearance of the intervertebral disk and presence of osteoarthritis make it difficult to analyse the scans even for expert radiographers. The algorithm may misclassify the image patch class and produce outlier predictions. Some examples of failed corner predictions are shown in last column of Fig. 6. With this best feature patch size, i.e. 5 x 5 pixels, new experiments were performed on the four corners of vertebra C3. But now either only classification or regression entropy is used. The results are reported in Table 2. It shows that the additional randomness introduced in the entropy selection process produces better results. The worst result is found when only regression entropy is used, because the vectors $\mathbf{d}_1$'s are too diverse and the algorithm cannot understand if the vectors are generated by the same class of image patches or not.

Table 1: Median, Mean and Standard Deviation (Std) of Errors in millimetres

Patch Size	30 x 30			10 x 10			5 x 5			3 x 3		
Vertebra	Median	Mean	Std	Median	Mean	Std	Median	Mean	Std	Median	Mean	Std
C3	3.01	4.38	4.36	2.81	4.22	4.34	2.92	4.27	4.15	3.04	4.48	4.39
C4	2.90	4.47	4.75	3.21	4.68	4.49	2.96	4.62	4.69	3.16	4.52	4.29
C5	4.02	5.74	5.05	3.96	5.16	4.42	3.46	5.31	4.74	4.19	5.43	4.41
C6	3.53	4.85	4.46	3.78	5.18	4.95	3.77	5.24	4.85	3.67	5.42	5.50
C7	1.98	2.80	2.70	2.09	2.83	2.53	2.16	2.81	2.35	2.04	2.82	2.51
Average	**3.09**	**4.45**	**4.27**	**3.17**	**4.41**	**4.14**	***3.05***	**4.45**	**4.16**	**3.22**	**4.54**	**4.22**

Table 2: Effect of Entropy Selection

Entropy		Random			Classification only			Regression only		
Vertebra	Corner	Median	Mean	Std	Median	Mean	Std	Median	Mean	Std
C3	1	3.10	4.70	4.54	3.01	4.15	3.87	2.37	4.85	5.50
	2	2.92	3.75	3.60	2.53	3.43	3.92	3.10	4.01	3.67
	3	2.55	4.08	4.15	3.25	4.51	4.31	3.51	5.07	5.12
	4	3.11	4.56	4.31	3.17	4.65	4.78	3.17	4.86	4.65
Average		**2.92**	**4.27**	**4.15**	**2.99**	**4.18**	**4.22**	**3.04**	**4.70**	**4.73**

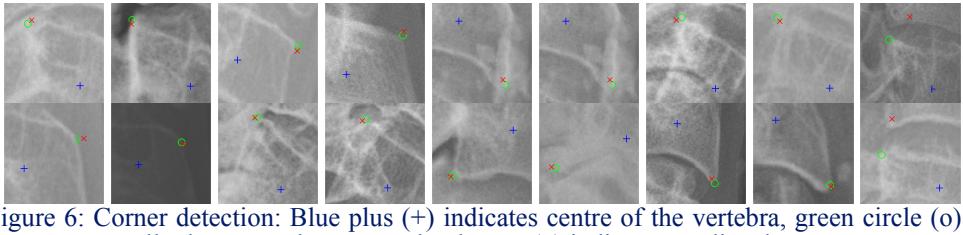

Figure 6: Corner detection: Blue plus (+) indicates centre of the vertebra, green circle (o) denotes manually demarcated corner and red cross (x) indicates predicted corners.

5 Conclusions

Automatic computer aided analysis of x-ray images is inherently challenging due to noise, low contrast, and anatomical variation between patients. This paper presents a machine learning approach to detect vertebral corners in cervical spine images, and is a first step towards automated alignment analysis and segmentation. In this work, we have explored the application Hough Forest in order to locate the corners of a vertebra. In the process, we have proposed a novel model, where the vertebra is divided into 12 images patches and each of these patch votes for possible corner location.

We have performed our experiment with very challenging, real life emergency room x-ray images. The images are diverse in size, shape, age and contrast. The results found with this approach are promising. In future, use of new features like [8] and [12] may be explored. We plan to include this part in an ASM framework to achieve full segmentation of the cervical vertebrae.

References

[1] Platzer, Patrick, Nicole Hauswirth, Manuela Jaindl, Sheila Chatwani, Vilmos Vecsei, and Christian Gaebler. "Delayed or Missed Diagnosis of Cervical Spine Injuries." *The Journal of Trauma: Injury, Infection, and Critical Care* 61.1 (2006): 150-55.

[2] Benjamin Narang ; Michael Phillips ; Karen Knapp ; Andy Appelboam ; Adam Reuben ; Greg Slabaugh; Semi-automatic delineation of the spino-laminar junction curve on lateral x-ray radiographs of the cervical spine . Proc. SPIE 9413, *Medical Imaging 2015: Image Processing*, 94132P (March 20, 2015); doi:10.1117/12.2082036.

[3] Benjelloun, Mohammed, Saïd Mahmoudi, and Fabian Lecron. "A framework of vertebra segmentation using the active shape model-based approach." *Journal of Biomedical Imaging* 2011 (2011): 9.

[4] Gall, Juergen, et al. "Hough forests for object detection, tracking, and action recognition." *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 33.11 (2011): 2188-2202.

[5] Tezmol, Abraham, et al. "Customized Hough transform for robust segmentation of cervical vertebrae from X-ray images." *Image Analysis and Interpretation, Fifth IEEE Southwest Symposium on*. IEEE, 2002.

[6] Klinder, Tobias, et al. "Automated model-based vertebra detection, identification, and segmentation in CT images." *Medical image analysis* 13.3 (2009): 471-482.

[7] Dong, Xiao, and Guoyan Zheng. "Automated vertebra identification from X-ray images." *Image Analysis and Recognition.* Springer Berlin Heidelberg, 2010. 1-9.

[8] Glocker, Ben, et al. "Automatic localization and identification of vertebrae in arbitrary field-of-view CT scans." *Medical Image Computing and Computer-Assisted Intervention–MICCAI* 2012. Springer Berlin Heidelberg, 2012. 590-598.

[9] Larhmam, Mohamed Amine, Said Mahmoudi, and Mohammed Benjelloun. "Semi-automatic detection of cervical vertebrae in X-ray images using generalized Hough transform." *Image Processing Theory, Tools and Applications (IPTA), 3rd International Conference on*. IEEE, 2012.

[10] Xu, Xi, et al. "Automatic segmentation of cervical vertebrae in X-ray images." *Neural Networks (IJCNN), The 2012 International Joint Conference on*. IEEE, 2012.

[11] Criminisi, Antonio, and Jamie Shotton. *Decision forests for computer vision and medical image analysis*. Springer Science & Business Media, 2013.

[12] Shotton, J.; Johnson, M.; Cipolla, R., "Semantic texton forests for image categorization and segmentation," *Computer Vision and Pattern Recognition*, 2008. CVPR 2008. *IEEE Conference on* , vol., no., pp.1,8, 23-28 June 2008