

Vertebrae Detection and Labelling in Lumbar MR Images

Meelis Lootus¹, Timor Kadir², and Andrew Zisserman¹

Oxford University¹,
Mirada Medical²

`meelis@robots.ox.ac.uk`

`timork@mirada-medical.com`

`az@robots.ox.ac.uk`

<http://www.robots.ox.ac.uk/~vgg/>

Abstract. We describe a method to automatically detect and label the vertebrae in human lumbar spine MRI scans.

Our contribution is to show that marrying two strong algorithms (the Deformable Part Model (DPM) object detector of Felzenszwalb *et al.* [9], and inference using dynamic programming on chains) together with appropriate modelling, results in a simple, computationally cheap procedure, that achieves state-of-the-art performance. The training of the algorithm is principled, and heuristics are not required. The detections are performed in all slices of a sagittal scan.

The method is evaluated quantitatively on a dataset of 371 MRI scans, and it is shown that the method copes with pathologies such as scoliosis, joined vertebrae, deformed vertebrae and disks, and imaging artifacts. We also demonstrate that the same method is applicable (without re-training) to CT scans.

Keywords: Spine, HOG, MRI, Detection, Vertebrae, SVM

1 Introduction

The task dealt with in this paper is the following: given an MRI scan of the lumbar spine, localize and label all the vertebrae present in that image. The motivation for this work is that spine appearance, shape and geometry measurements are necessary for abnormality detection *locally* at each disk [2, 3, 7, 12, 17, 20] and vertebrae [11, 22] (such as herniation), as well as *globally* for the whole spine (such as spinal scoliosis).

In more detail, the input 3D image is a (sparsely spaced) stack of 2D sagittal images, and the output consists of labelled tight bounding boxes with labels around all the vertebrae in the image. Each bounding box is specified by its position, orientation, and scale. An example of the detection and labelling for a typical normal scan is shown in Figure 1.

This detection task is challenging for a number of reasons, including: (i) the repetitive nature of the vertebrae, (ii) varying image resolution and imaging protocols; artefacts, and (iii) large anatomical and pathological variation,




Fig. 1. The task. Given a 3D MR lumbar spine image comprising of a stack of sagittal 2D slices as input (the mid-slice is shown on the left), localize and label in that 3D image all the vertebrae that are present. The output (projected on the mid-slice on the right) consists of labelled tight bounding boxes around the vertebrae. Note that all the 2D slices in the 3D slice stack are searched for vertebrae candidates.

particularly in the lumbar spine. Various examples of challenging cases in our dataset are highlighted in Figure 2. The anatomy and pathology variation can affect both the local vertebrae / disks appearance (e.g. degraded disks – Figure 2 H), and the global layout of the spine (e.g. scoliosis – Figure 2 C).

Contributions. Our method brings together two strong algorithms – the Deformable Part Model of Felzenszwalb *et al.* [9] based on Histogram of Oriented Gradients (HOG) image descriptors [6] and efficient inference on graphical models [10, 8] – making the algorithm accurate, robust, and efficient on challenging spine datasets. The algorithm is also tolerant to varying MRI acquisition protocols, image resolutions, patient position, and varying slice spacing unlike related solutions in the literature. It localizes all the vertebrae present in a scan, and labels them correctly as long as the sacrum is present. Importantly, the method is applicable to standard MRI protocols.

The method has two distinct stages. First, vertebrae candidates are detected using a sliding window detector searching over position, scale, and angle (section 2.1). Second, a graphical model is fitted to the set of candidate detections to find the optimal spine layout and labelling based on the unary soft output score of the detector for each part, and a spatial cost between each pair of connected parts (section 2.2). The HOG descriptor captures the near rectangular shape of the vertebrae. We detect vertebrae rather than disks since the vertebrae shape is more consistent than the disk shape as the lumbar spine studies are more often




Fig. 2. Spine variation in our data. A collection of example images showing assorted image, anatomical and pathological modes of global variation of the spine shape, and local variation of the vertebrae, and the disks. Our algorithm is robust to all those variations. Abnormalities have been highlighted by the red arrows. (A) Normal spine with a zoom on a normal vertebra. (B) A low-resolution image. (C) A coronal view of a scoliotic spine, resulting in the spine not being cut by a single sagittal slice. (D) Top: a normal sacrum, with unambiguous L5, S1 labelling based on shape and S1 and L5 orientation. Bottom: a sacrum with ambiguous L5, S1 labelling based on their shape and orientation. (E) Joined vertebrae. (F-J) Pathologically deformed vertebrae and disks. (K) Magnetic susceptibility imaging artefacts.

aimed at targeting disk deformations, and more suitable to be modelled with HOG. Disk locations can easily be found after detecting vertebrae.

The closest previous work to ours is that of Oktay *et al.* [18]. They detect disks and vertebrae in the lumbar spine using a Pyramid HOG descriptor; however, they only detect six disks and vertebrae with their graphical model, require the existence of both T1 and T2 scans to first detect the spinal cord, and they have a separate HOG template for each vertebrae. In contrast, we demonstrate that just one generic vertebrae detector suffices for all vertebrae, and only require the T2 scan. Furthermore, they only use the mid-sagittal slices, making it only applicable to cases where all the spine parts are in the mid-sagittal slice, whereas we search for vertebrae in all the 2D slices in the 3D stack (not restricted to the mid-sagittal slice).

Ghosh *et al.* [13] also use HOG features [6], however they do not label the vertebrae and make strong use of heuristics and information from complementary axial scans. They detect disks rather than vertebrae. Zhan *et al.* [23] present a robust hierarchical algorithm to detect and label arbitrary numbers of vertebrae & disks in nearly arbitrary field of view scans, as long as one of four ‘anchor’ vertebrae (C2, T1, L1, S) are present. They first detect the ‘anchor’ vertebrae,

and then other ‘bundle’ vertebrae connected to it graphically. Although the method works very well within its domain, it requires isotropic 2.1mm resolution scans which limits its applicability severely. Our method is not limited to this domain and, in particular, does not require the high isotropic resolution.

A further extensive body of literature on spine localization and labelling exists. Many of these methods have only been demonstrated on CT, and most with relatively small test datasets. In almost all the papers, the algorithms work in two stages. First, some anatomical parts characteristic of the spine are detected (vertebrae [4, 5, 14] / disks [1, 13, 15, 19] / both [18, 23]). Second, a spine layout model is fitted to the candidates to determine the best hypothesis for the spine layout. The spatial configuration of the spine parts, and in some cases also their individual characteristics [14, 16, 23], are taken into account to both label the disks and/or vertebra, and localize the spine.

2 Method

We present a method to localize and label vertebrae in lumbar MR images using two HOG-based detectors and a graphical model. First, given a stack of sagittal MRI slices, vertebrae and sacrum candidates are detected using a Deformable Part Model (DPM) on each slice as described in Section 2.1. Next, after local non-maxima suppression, the vertebrae candidates comprising the spine are selected and labelled by fitting a graphical model, as explained in Section 2.2.

2.1 Spine Part Detection

The spine part (vertebrae) detection is implemented using two detectors based on the DPM framework of [9]. We learn one generic 2D detector for vertebrae bodies (VBs), and another more specific 2D detector for the sacrum part, comprising the VBs of the first two links of the sacrum. Both the models are visualized along with a set of training samples in Figure 3.

Training. Both the generic vertebrae body (VB) detector and the sacrum detector are trained using the DPM framework [9]. The positive training examples for the VB detector are tight bounding boxes around the vertebral bodies of T10...L5 vertebrae with the bounding box sides parallel to the vertebral facets as shown in Figure 3 A. The positive training examples for the sacrum detector are tight bounding boxes around the first two links of the sacrum, with one side parallel to the posterior side of the sacrum as shown in Figure 3 B. The bounding boxes for both the VB and the sacrum are defined by fitting a minimum bounding rectangle to landmarks on them – four for the VB and eight for the sacrum. Each training sample is extracted from the slice intersecting the middle of the respective vertebral body.

For the VB detector, four HOG templates are trained, each with a different aspect ratio. The HOG templates are each 6 cells high, and 6, 7, 8, and 9 cells wide, corresponding to aspect ratios between 1 and 1.5. The HOG cell size for

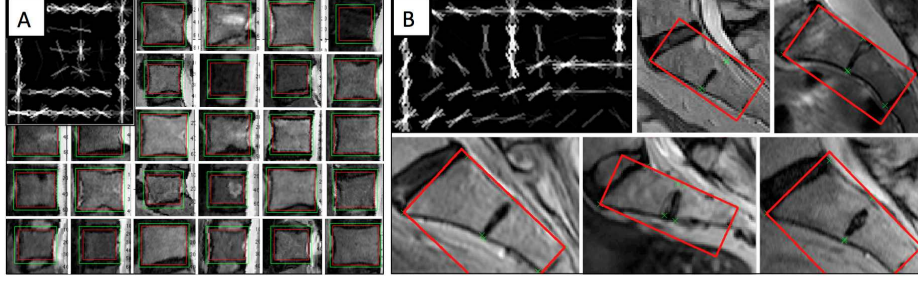


Fig. 3. The appearance model. Some training examples and a learned HOG template are shown for both the generic vertebrae body detector (A) and for the sacrum detector (B). The examples have been hand-annotated with tight ground truth bounding boxes as shown above and explained in Figure 5.

the VB model is 8×8 pixels. The HOG template for the sacrum detector is 9 cells high by 5 cells wide, with 8×8 pixel HOG cell size. The HOG feature vectors are 31-dimensional, with 18 contrast-sensitive, 9 contrast-insensitive direction bins; and 4 texture feature bins.

The HOG templates capture the rectangular shape of the vertebrae, with variations due to deformation, and the trapezoid shape of the first two links of the sacrum. The vertebrae exhibit wide size and resolution variation and are all scaled and warped to match one of the aspect ratios at training. The model is learned iteratively in several steps, with new positive samples mined by running the detector on the positive samples, collecting the strongest detections as new positives, and training a new detector using the new positives.

The negative samples for the vertebrae detector are first picked randomly from mid-slices with the vertebrae masked using a manually annotated polygon. Next, an iterative learning procedure is employed to pick hard negatives as false positive detections on the negative training images as detailed in [9].

Testing. During the candidate detection step at test time, a previously unseen sagittal scan is taken as input, and tight bounding boxes around vertebrae candidates are returned as output. The candidate search is performed in all slices of the scan. The VB and sacrum detector are run on each slice, searching over position, scale, and orientation. In the search over orientation, the scan rotated by -20° to 20° for vertebrae, and -60° to 0° for the sacrum, in 10 degree increments. A feature pyramid is calculated for each angle, with HOG cells placed densely next to each other. The feature pyramid has 10 levels per doubling of resolution (10 levels per octave), with the image resized and resampled to 2x the original size to 0.5x the original size from the finest to coarsest scale. All the detections at all positions, scales, orientations are collected and transformed onto the original test image coordinate system as shown in Figure 4.

A greedy non-maxima suppression algorithm is employed to remove most of the false positive detections in each slice as follows. First, the top-scoring




Fig. 4. Vertebrae Detection Pipeline. (A) Input image. (B) All detections at all rotation angles and scales. The green rectangles are generic vertebrae, and the red rectangles are sacrum candidates. (C) All detections, with top detections shown in thick blue line, and the “+” mark the ground truth vertebrae centre locations. (D) Output detection bounding boxes along with the ground truths and labels.

bounding box is retained, and all bounding boxes overlapping it more than 50% are discarded. Next, the second-highest scoring remaining bounding box is retained, and the discarding and retention process continues until all the remaining bounding boxes have at most 50% overlap.

Next, the remaining bounding boxes from all the slices are collected, and the non-maxima suppression process is repeated to retain only highest-scoring bounding boxes across all the slices that have at most 50% overlap between any two boxes. These bounding boxes are next passed as input to the Graphical Model as described in Section 2.2 in order to eliminate any remaining false positives, and to label the vertebrae.

2.2 Graphical Model for Spine Layout

We train a parts-based graphical model [8] connecting the vertebrae in a chain. The graphical model takes as input the detections after non-maxima suppression described in the previous section, and gives as output the placement and labels of all vertebrae in the image. The method deals with detections in multiple slices by ignoring the slice index in inference. However, detections in all slices are considered, and the slice index is returned in output. The spine layout is given as a configuration $L = (l_1, l_2, \dots, l_{n-1}, l_n)$ where l_i are the vertebra locations, with l_1 the C1 and $l_n = l_{25}$ the sacrum. The optimal configuration L^* of the graphical model is

$$L^* = \arg \min_L \left(\sum_{i=1}^n m_i(l_i) + \sum_{v_{i,j} \in G} d_{ij}(l_i, l_j) \right) \quad (1)$$

where l_i and l_j denote the vertebrae locations $l = (x_i, y_i, height_i, width_i, \theta_i)$ given by their location (x, y) , size $(height, width)$, and orientation θ_i . The best model fit minimizes the sum of the unary appearance mismatch terms m_i from the part detectors output and the spatial deformation cost d_{ij} for connected

pairs ij of parts, laid at l_i and l_j respectively. The last appearance term value m_{25} comes from the sacrum detector, and the rest of the appearance term values come from the universal vertebra detector. Since there might be fewer vertebrae visible in the scan than the model has, an additional “out-of-FOV” state is available for those vertebrae that indicates if they are not included in the scan. The appearance term m_i value for those vertebrae takes a constant penalty value learned on the training set as described in [21].

The spatial deformation cost is a sum of four box functions S , T , U , V on pairs of adjacent vertebrae in the chain:

$$d_{ij}(l_i, l_j) = S(A_i/A_j) + T(x_i - x_j) + U(y_i - y_j) + V(\theta_i - \theta_j) \quad (2)$$

where A_i , A_j are the areas, x_i , x_j & y_i , y_j the positions, and θ_i , θ_j the angles of the adjacent vertebrae i and j . The box functions take a low constant value if their argument values are within favourable distance of each other and a higher constant value if their arguments are outside that distance.

To speed up the fitting process, a Viterbi message passing scheme from [8] for fast inference in $O(nh^2)$ time is employed where n is the number of parts and h the number of candidates per part. Typically, there are around $h = 100$ candidate positions per part, plus an “out-of-FOV” state for each part.

Training. The edges for the box functions S , T , U , and V are found as the minimum and maximum argument values of those functions on the training set (e.g. the minimum and maximum x -distance between L1 and L2 for T , etc.).

Testing. At test time, the full model is fitted to the scan, with each part allowed to be visible, or out-of-FOV.

3 Experiments

3.1 Data, annotation and evaluation

The dataset consists of 371 MRI T2-weighted lumbar scans, acquired under various protocols. The scans contain normal and various abnormal cases as illustrated in Figure 2. The dataset is split into 80 training and 291 testing images. The scans have isotropic in-slice resolution varying from 0.34 to 1.64 mm with mean at 0.78, median at 0.84 mm; and varying slice spacing from 3mm to 5 mm, with 4mm in almost all scans. The scans range in fields of view, containing 7 to 23 vertebrae starting from the sacrum, with median at 10 per scan.

Annotation. The scans were hand-annotated with two types of ground truth as illustrated in Figure 5: (i) All the vertebrae centres in all the scans are marked with a point (“+” in Figure 5), and labelled with the vertebrae name; and (ii) all the training scans plus some test scans are annotated with a tight bounding box around each vertebra (Figure 5 A3, B3). The tight bounding boxes were defined by points (“x” in Figure 5) along the vertebrae boundaries as shown.




Fig. 5. The ground truth annotation process. A1-A3 show the generic vertebral, and B1-B3 the sacrum annotation process. There are two types of annotation: single point (the green “+” in the Figure – used for testing) and bounding box (the red rectangle – used for training). Given an input (A1, B1), the points (“+” and “x”) are hand-placed (A2,B2). The bounding box annotation is found as the minimal bounding rectangle to the “x” points around the vertebra / sacrum boundary. There are four boundary points for vertebrae (A) and eight for the sacrum (B).

Evaluation protocol. The detections are evaluated against vertebrae-centre and the sacrum-centre ground truth points. A positive detection for the sacrum is counted if a detected sacrum bounding box contains the sacrum ground truth point and does not contain any vertebrae centre ground truth points. A positive detection for the vertebrae is counted if a detected vertebra bounding box contains one and only one ground truth point for a vertebral body, including the sacrum. Note, this evaluation protocol ensures that the case where a large detection covers several vertebrae is not counted as positive.

3.2 Results

The algorithm is evaluated on a set of 291 lumbar spine test images with variable number of vertebrae visible. Example outputs are shown in Figure 6, and statistical results on localization error over the test set are plotted in Figure 7 and tabulated in Figure 8 by vertebrae type.

We achieve 84.1% correct identification rate overall, and 86.9% for the lumbar vertebrae. The mean detection error between the ground truth centre of the vertebrae and the centre of the detected bounding box is 3.3mm, with standard deviation 3.2mm. If the assigned labels are allowed to be shifted by one vertebra in either direction, the errors are 92.9% and 94.7% respectively. Typically, the full detection and labelling process from input to output takes less than a minute, with majority of time spent on candidate detection.

Independent sacrum detection (without graphical model) with local non-maxima suppression shows 98.1% recall at 48% precision. Independent general vertebrae detection (without graphical model) shows 97.1% recall at 9.1% precision.

Our method works well on very challenging examples with various anomalies illustrated earlier in Figure 2. The identification results compare favourably to other approaches in the literature, although direct comparison is not possible since the algorithms have been evaluated on different datasets. Glocker *et al.* [14] report median identification error of 81% with median localization error below

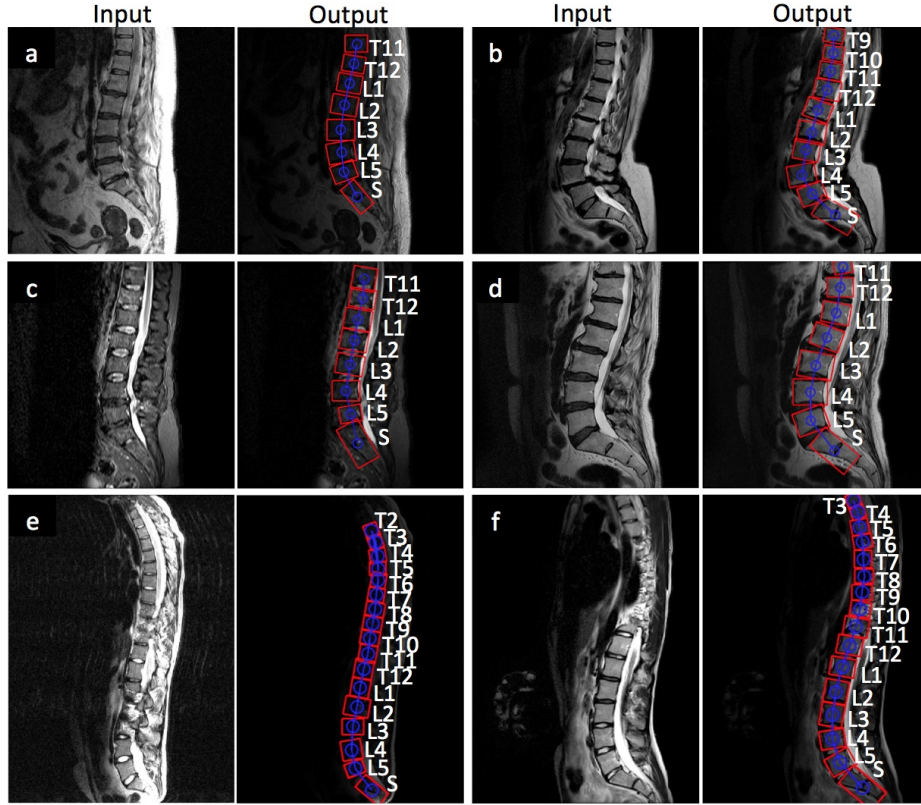


Fig. 6. Example results. Input and output are shown for six different scans a-f. The thick solid line rectangles show the detections for each vertebrae, along with their anatomical labels. Note how the algorithm is robust to varying Field of View, resolution, and anatomy. Note, for visualization purposes, only the mid-sagittal slice is shown for each scan, and all bounding boxes are projected onto it.

6mm on CT images. Zhan *et al.* [23] detect disks and vertebrae in isotropic MRI scans with 97.7% “perfect” labelling rate as assessed by a medic but do not report detection errors. Pekar *et al.* [19] report 83% correct labelling rate on 30 lumbar MRI scans. Our method correctly localizes the centres of vertebrae out of the mid-sagittal slice in scoliotic cases such as scan (f) in Figure 6.

Application to CT images. Although the method has been principally designed for MR images, it is directly applicable to CT images as shown in Figure 9. No retraining is required for detection on CT due to the high generalization of HOG detectors.




Fig. 7. Localization error by vertebrae type. Boxplots representing detection errors are shown. The error for a given vertebra type is calculated as the distance between the centre of the detected bounding box and the ground truth vertebra centre, divided by the mean width of that vertebra. The mean vertebrae widths are evaluated based on the bounding boxes on the training set. The horizontal line in the middle of each box is the median error, and the bottom and top of each box are the 25 and 75 percentile errors respectively. The bottom and top error bar end are the 5 and 95 percentile errors respectively, and the ‘+’ denote statistical outliers.

4 Conclusion

We have presented a HOG-based algorithm to localize vertebrae in lumbar MRI scans of the spine that is simple, accurate and efficient. We demonstrate robustness to severe deformations due to diseases, image artefacts, and a wide range of resolution, patient position, and acquisition protocols on a challenging clinical dataset. It is straightforward to extend the method to completely general FOVs if required, by taking other anatomical context into account [14].

Acknowledgements

We are grateful for discussions with Prof. Jeremy Fairbank and Dr. Jill Urban. Financial support was provided by ERC grant VisRec no. 228180 and by EPSRC. The data used in this research was obtained during the EC FP7 project HEALTH-F2-2008-201626.

References

1. R.S. Alomari, J.J. Corso, and V. Chaudhary. Labeling of lumbar discs using both pixel- and object-level features with a two-level probabilistic model. *IEEE TMI*, 2011.
2. R.S. Alomari, J.J. Corso, V. Chaudhary, and G. Dhillon. Desiccation diagnosis in lumbar discs from clinical mri with a probabilistic model. In *ISBI*, 2009.
3. R.S. Alomari, J.J. Corso, V. Chaudhary, and G. Dhillon. Computer-aided diagnosis of lumbar disc pathology from clinical lower spine mri. *International Journal of Computer Assisted Radiology and Surgery*, 5(3), 2010.

	C7	T1	T4	T7	T10	L1	L2	L3	L4	L5	S	Lumbar	Overall
Mean error (mm)	5.0	3.4	2.2	3.4	3.3	2.7	3.0	3.2	3.0	2.7	6.5	3.5	3.3
Std in error (mm)	0.0	4.4	1.4	2.8	2.8	2.1	2.4	2.8	2.4	2.1	5.9	3.0	3.2
Mean vert. width (mm)	19.4	19.6	23.9	27.9	32.7	36.0	37.6	38.4	38.8	37.8	31.4	36.6	30.0
Labelled count	1	2	9	25	151	241	249	255	254	254	264	1517	2293
Labelled count (+/- 1)	1	4	13	27	169	271	280	286	284	269	264	1654	2532
Total count	7	11	20	33	181	291	291	291	291	291	291	1746	2726
Labelling rate	14%	18%	45%	76%	83%	83%	86%	88%	87%	87%	91%	86.9%	84.1%
Labelling rate (+/- 1)	14%	36%	65%	82%	93%	93%	96%	98%	98%	92%	91%	94.7%	92.9%

Fig. 8. Localization errors. The mean and standard deviation (std) of localization errors are shown for the correctly detected and labelled vertebrae (identification rate 84% overall and 87% for lumbar). In addition the “count” – the number of vertebrae detected of each type – is provided, along with the mean width of each of the vertebrae in training set. By allowing the labelling to be correct to ± 1 vertebrae, the identification rates become 93% and 95% for all and lumbar vertebrae respectively. Note that some thoracic vertebrae results are omitted to save space.




Fig. 9. Detection on CT images with detectors trained on MRI. Detectors trained on MR images can also successfully localize vertebrae in CT scans, indicating the robustness of the method to varying image appearance.

4. M.S. Aslan, A. Ali, H. Rara, and A.A. Farag. An automated vertebra identification and segmentation in ct images. In *Proc. of 2010 IEEE 17th International Conference on Image Processing*, 2010.
5. M.P. Chwialkowski, P.E. Shile, D. Pfeifer, R.W. Parkey, and R.M. Peshock. Automated localization and identification of lower spinal anatomy in magnetic resonance images. *Computers and Biomedical Research*, 24(2), 1989.
6. N. Dalal and B Triggs. Histogram of Oriented Gradients for Human Detection. In *Proc. CVPR*, volume 2, pages 886–893, 2005.
7. D.F. Fardon and P.C. Milette. Nomenclature and classification of lumbar disc pathology. *Spine*, 26(5):E93–E113, 2001.
8. P. Felzenszwalb and D. Huttenlocher. Pictorial structures for object recognition. *IJCV*, 61(1), 2005.
9. P. F. Felzenszwalb, R. B. Grishick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part based models. *IEEE PAMI*, 2010.

10. M. Fischler and R. Elschlager. The representation and matching of pictorial structures. *IEEE Transactions on Computer*, c-22(1):67–92, Jan 1973.
11. S. Ghosh, R.S. Alomari, V. Chaudhary, and G. Dhillon. Automatic lumbar vertebra segmentation from clinical ct for wedge compression fracture diagnosis. In *SPIE*, 2011.
12. S. Ghosh, R.S. Alomari, V. Chaudhary, and G. Dhillon. Computer-aided diagnosis for lumbar mri using heterogeneous classifiers. In *ISBI*, 2011.
13. S. Ghosh, M.R. Malgireddy, V. Chaudhary, and G. Dhillon. A new approach to automatic disc localization in clinical lumbar mri: Combining machine learning with heuristics. In *ISBI*, 2012.
14. B. Glocker, J. Feulner, A. Criminisi, D.R. Haynor, and E. Konukoglu. Automatic localization and identification of vertebrae in arbitrary field-of-view ct scans. In *MICCAI*, 2012.
15. B.M. Kelm, M. Wels, K.S. Zhou, S. Seifert, M. Suehling, Y. Zheng, and D. Comaniciu. Spine detection in ct and mr using iterated marginal space learning. *Medical Image Analysis*, 2012.
16. 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.
17. S. Michopoulou, L. Costaridou, M. Vlychou, R. Speller, and A. Todd-Pokropek. Texture-based quantification of lumbar intervertebral disc degeneration from conventional t2-weighted mri. *Acta Radiologica*, 52(1):91–98, 2011.
18. A.B. Oktay and Y.S. Akgul. Simultaneous localization of lumbar vertebrae and intervertebral discs with svm based mrf. *IEEE TMI*, April 2013.
19. V. Pekar, D. Bystrov, H. S. Heese, S. P. M. Dries, S. Schmidt, R. Grewer, C.J.d. Harder, R.C. Bergmans, A.W. Simonetti, and A.M.v. Muiswinkel. Automated planning of scan geometries in spine mri scans. In *MICCAI*, 2007.
20. C.W.A. Pfirmann, A. Metzdorf, M. Zanetti, J. Hodler, and N. Boos. Magnetic resonance classification of lumbar intervertebral disc degeneration. *Spine*, 26(17):1873–1878, 2001.
21. V. Potesil, M. Lootus, A. El-Labban, and T. Kadir. Landmark localization in images with variable field of view. In *ISBI*, 2013.
22. M. Wels, B.M. Kelm, A. Tsymbal, M. Hammon, G. Soza, M. Sühling, A. Cavallaro, and D. Comaniciu. Multi-stage osteolytic spinal bone lesion detection from ct data with internal sensitivity control. In *SPIE*, 2012.
23. Y. Zhan, D. Maneesh, M. Harder, and X.S. Zhou. Robust mr spine detection using hierarchical learning and local articulated model. In *MICCAI*, 2012.