



Published in final edited form as:

Comput Med Imaging Graph. 2014 October ; 38(7): 606–612. doi:10.1016/j.compmedimag.2014.04.007.

Computer Aided Detection of Epidural Masses on Computed Tomography Scans

Jiamin Liu, Sanket Pattanaik, Jianhua Yao, Evrim Turkbey, Weidong Zhang, Xiao Zhang, and Ronald M. Summers

Imaging Biomarkers and Computer-Aided Diagnosis Laboratory, Radiology and Imaging Sciences Department, Clinical Center, National Institutes of Health, Bethesda, MD USA

Abstract

The widespread use of CT imaging and the critical importance of early detection of epidural masses of the spinal canal generate a scenario ideal for the implementation of a computer-aided detection (CAD) system. Epidural masses can lead to paralysis, incontinence and loss of neurological function if not promptly detected. We present, to our knowledge, the first CAD system to detect epidural masses on CT scans. In this paper, spatially constrained Gaussian Mixture Model (GMM) and supervoxel-based method are proposed for epidural mass detection. The detection is performed on the Gaussian level or the supervoxel level rather than the voxel level. Cross-validation on 40 patients with epidural masses on body CT showed that the supervoxel-based method yielded a significant improvement of performance (82% at 3 false positives per patient) over the spatially constrained GMM method (55% at 3 false-positives per patient).

1. Introduction

Masses in the epidural space of the spinal canal can cause discomfort, pain and even paralysis by compressing the spinal cord and nerve roots. Moreover, the presence of an epidural mass within the spinal canal is a strong predictor of metastatic disease. A retrospective study of 337 patients at the Mayo Clinic, for instance, revealed that 20% of all cases of spinal epidural metastases presented as the initial manifestations of malignancy [1]. Given the importance that early indicators of malignant cancers hold in the radiology community, the absence of a body of work on computer aided detection (CAD) of spinal canal lesions within the intradural and extra-dural space is quite surprising.

A CAD system designed to detect epidural masses within the constraints of the CT modality could prove invaluable. While confirmation of epidural tumors is almost always made using magnetic resonance imaging (MRI), due to its higher anatomic resolution and sensitivity to alterations of the central nervous system tissue, most patients will have received an examination using CT images. CT imaging remains the most prevalent radiologic modality

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

as it is rapid (current generation can take less than 1 minute), cost-effective, and can image over a large body area, effectively localizing many types of soft tissue tumors. Automated detection of an epidural mass is very challenging because of its low contrast to normal soft tissue in the spinal canal. Even a radiologist may fail to detect an epidural mass in CT scan, especially when the patient is being examined for unrelated symptoms. Figure 1 shows an example of a subtle epidural mass on the CT scan that is detectable as a slightly hyperattenuating region (Fig. 1(a)) when the image is viewed using an appropriate soft-tissue window width and level. The patient cohort examined in this study had an MRI confirming the presence of an epidural mass (Fig. 1(b)) within a month after the mass was detectable on a CT scan, a duration that could drastically affect patient outcome when concerning progressive disease.

A vast majority of tumors that encroach on the epidural spaces originate from the intravertebral foramina or the vertebral bodies surrounding the spinal canal [2]. In other words, the masses always extend contiguously from the radiopaque bony regions of the spine into the soft tissue and can be discriminated as a tissue exhibiting attenuation intermediate between bone and normal spinal canal soft tissues. In this work, two methods are implemented individually for epidural mass detection. (1) Spatial information is incorporated into conventional intensity based Gaussian mixture model (GMM) as a spatially-constrained GMM [3] for mass detection. A large number of Gaussians is used per class in the spinal canal to capture the local spatial feature. The intensity of a tissue is considered a global feature and is modeled by parameters linking all associated Gaussians. In this way, the mass *detection is performed on the Gaussian level rather than on the voxel level*. (2) Supervoxels or regions produced by bottom-up segmentation have better spatial support compared to voxels. Therefore a supervoxel-based mass detection is proposed and the *detection is performed on the supervoxel level rather than on the voxel level*. The performance of two methods are evaluated and compared. To our best knowledge, this work is the first computer aided detection of epidural mass on CT scans. A previous version of this work has been published [4].

2. System Description

This section presents the details of all the components of our mass detection system. The following is a summary of the steps taken by the system for each CT scan.

1. Region of interest (ROI): spinal canal is segmented and refined (Sect. 2.1) as the ROI of epidural mass detection.
2. Mass candidate generation: two methods are implemented individually, spatially constrained GMM (Sect. 2.2.1) and supervoxel (Sect. 2.2.2).
3. Rules for mass: three rules are used for early rejection of false positives (Sect. 2.3).
4. Classification: features are computed and selected for support vector machine (SVM) classification (Sect. 2.4).

2.1 Region of Interest

First, initial spine segmentation and localization is achieved by simple thresholding and region growing. We apply a threshold of 200 HU to mask out the bone pixels. A connected component analysis is conducted on the bone mask and the largest connected blob in the center of the image is retained as the initial spine segmentation.

The spinal canal links all vertebrae into a column. On a 2D slice, the spinal canal appears to be a low intensity oval region surrounded by vertebral structures. We apply a watershed algorithm to detect the potential spinal canal regions. The watershed regions surrounded by bone pixels are recorded as potential candidates for the spinal canal.

Then, a directed graph searching algorithm [5] is applied to find the longest path of potential spinal canal candidates, which is the spinal canal in our case. A four-part vertebra model was then used to locate the vertebral bodies, spinous processes, and left/right transverse processes, with rib structures used to separate vertebral segments.

Finally, curved planar reformations were then employed to segment the spinal canal, using the centerline as a backbone. More details about spinal canal segmentation can be found in our previous work [6].

The segmentation of the spinal canal, however, suffers from the ambiguity in discriminating between hyper-attenuating bony and spinal lesion regions within the canal. As a result, many cases under-segment the region of interest within which we expect to detect epidural masses. To contend with this complication, we use a snake [7] to refine the segmentation so that it covers the entire spinal canal. The spinal canal segmentation and refinement are demonstrated in Figure 2.

2.2 Mass Candidate Generation

2.2.1 Gaussian Level Detection—Intensity based K-Means clustering ($k=4$) is used for initial classification of tissues in the spinal canal. Selecting four different classes allowed us to delineate classes representative of normal intradural soft tissue, hypo-attenuating fatty tissue and vasculature, epidural masses, and the partial volume between the bone and soft tissue.

The spatially constrained GMM framework was modified from a method [3] employed to detect multiple sclerosis lesions from MRI images of the brain. To accommodate the spatial feature, we model an image as if its voxels were drawn independently from a mixture of many Gaussians:

$$f(x, I(x)) = \sum_{i=1}^n \alpha_i f_i(x, I(x) | \mu_i, \Sigma_i) \quad (1)$$

where x is the 3D position information included in the spatial vector (spatial parameters), $I(x)$ is the intensity vector (intensity parameters) associated with the voxel in position x , n is the number of Gaussians components in the mixture model, μ_i and Σ_i are the mean and the covariance of the i th Gaussian components f_i , and α_i is the i th mixture coefficient. The

spatial feature is incorporated into the probabilistic model. Each Gaussian component in the spatially constrained GMM represents a probabilistic model for a specific small area in the CT image, therefore $n \ll k$. Each Gaussian component is linked to a single tissue class and all the Gaussian components related to the same tissue class share the same intensity parameters. Assuming the intensity and spatial features are uncorrelated, we have

$$\mu_i = \begin{pmatrix} \mu_i^x \\ \mu_{\pi(i)}^I \end{pmatrix}, \Sigma_i = \begin{pmatrix} \Sigma_i^x & 0 \\ 0 & \Sigma_{\pi(i)}^I \end{pmatrix} \quad (2)$$

$\pi(i)$ is the tissue that is linked to the i th Gaussian component, μ_i^x and Σ_i^x are the spatial mean and covariance of the i th Gaussian component, and $\mu_{\pi(i)}^I$ and $\Sigma_{\pi(i)}^I$ are intensity mean and variance of class π ($\pi = 1, 2, 3, 4$) to which the i th Gaussian component belongs. Therefore, the Gaussian component f_i can be written as:

$$f_i(x, I(x)) = \sum_{i=1}^n \alpha_i(x; \mu_i^x, \Sigma_i^x) \times (I(x); \mu_{\pi(i)}^I, \Sigma_{\pi(i)}^I) \quad (3)$$

The main advantage of the spatially constrained GMM is to combine local spatial features with a global intensity feature, which makes it much more robust to noise than intensity based methods. The four tissue classes generated from the K-Means clustering served as the initialization for our GMM framework. Small clusters (<20 voxels) were defined as under-representative of the associated tissue class. From the remaining clusters, 1/20 voxels were selected as center of Gaussian components (Fig. 3(c)) in eq. (1). Each voxel within the cluster was then linked to its nearest Gaussian center. The component coefficient α_i was then initialized as the number of voxels in i th Gaussian component divided by the total number of voxels of all n Gaussian components.

After applying EM for estimation of the spatial and intensity parameters of each Gaussian, a Maximum-APosteriori (MAP) criterion was used to return the final label of the class. As a result, four classes are updated. Figure 3 shows the classification by K-Means (Fig. 3(b)) and GMM (Fig. 3(d)).

2.2.2 Supervoxel Level Detection—The advantages of employing supervoxels for mass detection are three fold: 1) supervoxels have better spatial support than voxels; 2) it significantly reduces the computation burden; and 3) supervoxels have more reliable statistics than voxels or Gaussian components of GMM.

To generate supervoxel representations, we apply the efficient graph-based segmentation technique [8] to oversegment the spinal canal into homogeneous regions. The graph-based supervoxel technique groups neighboring voxels based on their intensity differences, so that similar voxels are more likely to be grouped together. We specify the minimal region size in the supervoxel segmentation to be 40 voxels. There are about 1000 supervoxels within each

spinal canal (see Fig 4(b) for examples of produced supervoxels). Each supervoxel is represented by the mean and standard deviation of voxel intensity within the supervoxel. The supervoxels with high standard deviation (i.e. non-reliable homogenous regions such as partial volume region between two classes) are discarded from the future analysis.

K-Means clustering is then applied to the supervoxels with $k=3$ (normal soft tissue, hypo-attenuating fatty tissue and vasculature, and epidural masses). Figure 4(c) shows the K-Means clustering result of those qualified supervoxels within the spinal canal.

2.3 Rules for Early Rejection of False Positives

After the Gaussian level or supervoxel level mass candidates generation, a rule based system can be used for early rejection of false positives. A Gaussian or supervoxel for which all these conditions hold, is labeled as epidural mass. A rule set used in this work is the following:

1. Class: the detected Gaussian component or supervoxel should be in the epidural mass class.
2. Contrast: the CT intensity of detected mass usually has enough contrast compared to its nearby normal soft tissues, i.e. $\text{Mass}_{\text{CT}} - \text{SoftTissue}_{\text{CT}} > th$ HU. $th=40$ HU was used in our study.
3. Location: the detected mass should touch the boundary of spinal canal.

The first two rules reflect the general appearance of the epidural mass. The third rule incorporates contextual information by reflecting our expectation to find masses in the boundary of spinal canal. Note that at this point decisions are made at the Gaussian level or supervoxel level rather than the voxel level. The detected epidural masses are used for feature computation and classification.

2.4 Feature Computation and SVM

A comprehensive collection of texture features from the mass detections were computed in this work. Haralick Gray-Level Co-occurrence Matrix (GLCM) features [9] are widely used for analyzing image texture. The co-occurrence matrix stores the co-occurrence frequencies of the pairs of gray levels, which are configured by different distances and directions. We calculated the co-occurrence matrices for 4 offset distances and 13 directions on multiple planes, yielding 52 matrices for each mass detection. We then calculated 12 features from the matrix, including energy, entropy, correlation, contrast, variance, sum of mean, inertia, cluster shade, cluster tendency, homogeneity, maximal probability, and inverse variance. Thus, each detected mass has 624 Haralick GLCM features. We also extracted the volume of the masses, histograms of oriented gradients [10] and local binary patterns [11] features. Finally, an SVM committee [12] is employed for classification. The method involved bootstrap aggregation of features into seven SVM committees to improve on selection of features and avoid overfitting.

3. DataSet

The patient population consisted of patients who received a chest, abdomen, and pelvis CT scan within 1 month before or after receiving an MRI confirming the presence of an epidural mass. A cohort of 40 patients (27 scans with IV contrast and 13 without IV contrast; 23 patients with epidural masses and 17 control patients; 25 scans from General Electric, 9 from Siemens and 6 from Philips scanners), with a total of 42 epidural masses confirmed as visible in the CT images by radiologists, were examined in this study. The inter-slice spacing was 5–7 mm (90–150 slices), and the voxel spacing within an axial slice was in the range 0.7 – 0.9 mm (512×512 pixels). The detected epidural mass was marked as true-positive if the overlap between the detected mass and the mass demarcated as ground truth by a radiologist is more than 10% of the ground truth.

4. Results

The detection performance was evaluated using ten-fold cross-validation. The free response receiver operating characteristic (FROC) curves of the detection performance by using spatially constrained GMM and supervoxel are compared in Figure 5. We get a sensitivity of 55% at 3 false positives per patient on average for spatially constrained GMM method. We achieve a sensitivity of 82% at 3 false positives per patient on average for the supervoxel method (Fig. 5(b)). The difference in the methods was statistically significant ($p=0.02$) at the aforementioned operating points as determined using Fisher's exact test [13].

The FROC curves of the detection performance on different sizes of masses are shown in Figure 6. The mass size is represented by the area of spinal canal occupied by the mass in 2D (axial slice which has maximal mass). For the spatial constrained GMM (Fig. 6(a)), the system achieved a sensitivity of 72% at 2 false positives per patient for large masses and a sensitivity of 40% at 2 false positives per patient for small masses. For the supervoxel method (Fig. 6(b)), the system achieved a sensitivity of 87% at 2 false positives per patient for large masses (occupied area 30% of spinal canal area) and a sensitivity of 61% at 2 false positives per patient for small masses (occupied area between 5% and 30% of spinal canal area).

Figure 7 shows examples of true and false positive and false negative detections. Three masses are missed because there is not enough contrast with the nearby normal soft tissue. Beam hardening artifacts, disk bulging, and ligaments are the most common false positives which have similar appearance with epidural masses.

This CAD system is computationally efficient. The spinal canal segmentation takes about 2 minutes, mass candidate detection is about 1 minute, feature computation and classification take around 1 minute.

5. Discussion

In this work, spatially constrained GMM reported about two times as many false positives as the supervoxel method. For example, Figure 8(a) shows the false positives (disc bulging and ligament) generated from both GMM and supervoxel methods. The false positives (partial

volume between bone and normal soft tissue) generated from the GMM method only are shown in Figure 8(b) and (c). The main reason is because supervoxels have more reliable statistics than voxels or Gaussian components of GMM. These false positives were eliminated by subtracting supervoxels with high standard deviation of voxel intensity.

The false positive detections of our CAD system were mainly due to beam hardening artifacts in the upper thoracic level, intervertebral disk bulging or herniation to the epidural space, and hypertrophied ligamentum flavum (Fig. 7). The beam hardening artifacts in the upper thoracic level is well known and not avoidable. The bulging disk is an extension of disc out of its space as a part of normal aging process. The herniated disc is caused by protrusion of inner cartilage of the disc outside (i.e. epidural space) through a crack in the outer layer of the cartilage. These two conditions are space occupying lesions but they are not soft tissue masses. There are no significant density or texture differences between soft tissue masses and disc bulging/herniations. The differentiation between herniated disc and epidural mass is based on defining the anatomic level of the abnormality as being at the level of the intervertebral disc or not. The ligamentum flavum is a thin ligament within posterior portion of the spinal canal and connects the laminae of two adjacent vertebrae. Hypertrophy of this ligament is a degenerative condition and usually associated with facet joint hypertrophy and degenerative disc diseases.

The clinical significance of this work is that epidural masses are, in our experience, frequently missed as inadequate attention is paid to the spinal canal on routine body CT image interpretation. As epidural masses can be a first sign of malignancy or cause pain, paralysis or loss of function, it is important to detect them. For each patient, a radiologist needs to review all the axial CT images to detect epidural lesions. This step is required for each case even when the patient's diagnosis or symptoms are unrelated. Several there are a hundred images or more, this step is time consuming. The CAD can be run prior to the radiologist's review taking about 5 minutes of processing time and generates only a few images for the radiologist to review. Thus, our CAD system can potentially shorten the time spent for assessing epidural masses. In clinical practice, it will be useful if the CAD system detects 10 candidate lesions per patient with 90% sensitivity. The CAD system in the current study reaches this level of accuracy and may be applicable to clinical practice. This CAD system may reduce the incidence of missed epidural masses, although this would need to be demonstrated in a clinical trial.

6. Conclusions

We present the first CAD system to detect epidural masses in the CT imaging modality. Spatially constrained GMM and supervoxel-based methods are introduced and compared for the mass detection. The detection is performed on the Gaussian level or supervoxel level rather than the voxel level. The supervoxel-based method is more accurate and efficient than the spatially constrained GMM. We envision our CAD system could be used to direct the attention of a reader to a mass even when the reader is not primed for its detection.

Acknowledgments

This research was supported by the Intramural Research Program of the NIH Clinical Center.

References

1. Schiff D, O'Neill BP, Suman VJ. Spinal epidural metastasis as the initial manifestation of malignancy: Clinical features and diagnostic approach. *Neurology*. 1997; 49(2):452–456.2. [PubMed: 9270576]
2. Fine H, et al. Neoplasms of the Central Nervous System. *Cancer: principles & practice of oncology*. 2005:1834–1887.
3. Freifeld, O., Greenspan, H. *International Journal of Biomedical Imaging*. 2009. Multiple Sclerosis Lesion Detection Using Constrained GMM and Curve Evolution.
4. SP, Liu, J., Yao, J., Turkey, E., Zhang, W., Summers, R. Epidural Masses Detection on Computed Tomography Using Spatially-Constrained Gaussian Mixture Models. *Medical Image Computing and Computer-Assisted Intervention (MICCAI), Workshop on Computational Spine imaging*; 2013;
5. Cormen, TH., Leiserson, CE., Rivest, RL. *Introduction to Algorithms*. Mac Graw Hill; 1989.
6. Yao, J., O'Connor, SD., Summers, RM. Automated spinal column extraction and partitioning. 2006 3rd IEEE International Symposium on Biomedical Imaging: From Nano to Macro - Proceedings; 2006;
7. Xu CY, Prince JL. Generalized gradient vector flow external forces for active contours. *Signal Processing*. 1998; 71(2):131–139.
8. Felzenszwalb PF, Huttenlocher DP. Efficient graph-based image segmentation. *International Journal of Computer Vision*. 2004; 59(2):167–181.
9. Haralick RM, Shanmugam K, Dinstein I. Textural features for image classification. *IEEE Trans Syst, Man, Cybern*. 1973; 3(6):610–621.
10. Dalal N, Triggs B. Histograms of oriented gradients for human detection. *Proc CVPR*. 2005
11. Ojala T, Pietikainen M, Maenpaa T. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Trans Pattern Anal Mach Intell*. Jul; 2002 24(7): 971–987.
12. Yao J, Summers RM, Hara AK. Optimizing the committee of support vector machines (SVM) in a colonic polyp CAD system. *SPIE Medical Imaging*. 2005; 5746:384–392.
13. Carlson, J., Heckerman, D., Shani, G. Microsoft Research technical report MSR-TR-2009-53. 2009. False Discovery Rate for 2x2 Contingency Tables.

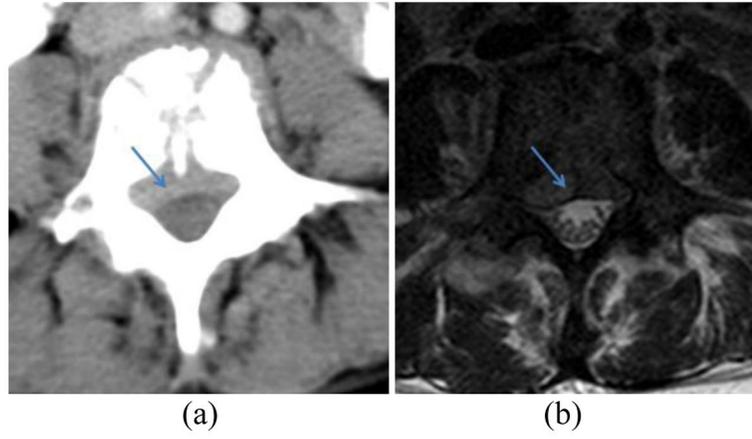


Fig. 1. A L4 level anterior epidural mass on CT (a) and MRI (b) results in spinal canal compromise.

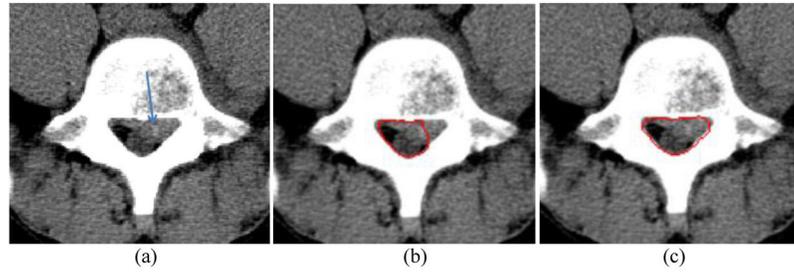


Fig. 2. Epidural mass (a) is not fully included in initial segmentation of the spinal canal (b), but is fully included in the refined canal segmentation (c).

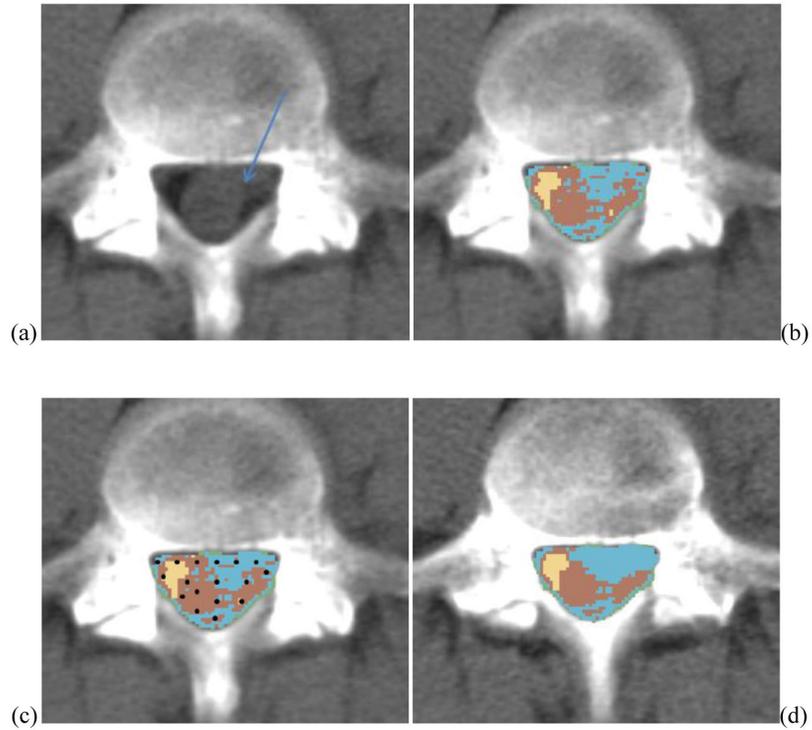


Fig. 3.

(a) Epidural mass (arrow) on CT confirmed by MRI. (b) K-Means clustering result (soft tissue in brown; fat in yellow; epidural mass in blue; partial volume between the bone and soft tissue in green). (c) Centers of Gaussian components (black dots) in the mixture model. (d) Clustering after spatially constrained GMM.

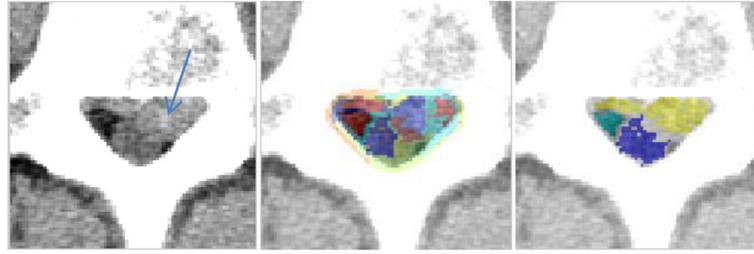


Fig. 4. Example of supervoxels and 3-class classification. (a) Epidural mass (arrow). (b) Supervoxels representation of spinal canal. (c) K-Means clustering result of supervoxels (3 classes: soft tissue in blue; fat in green; epidural mass in yellow).

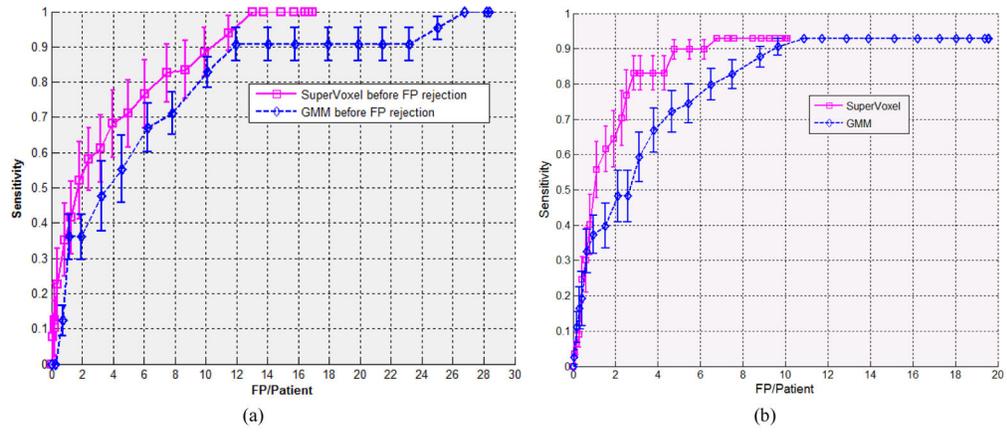


Fig. 5. 10-fold cross validation FROC of two methods. Before (a) and after (b) early rejection of FPs.

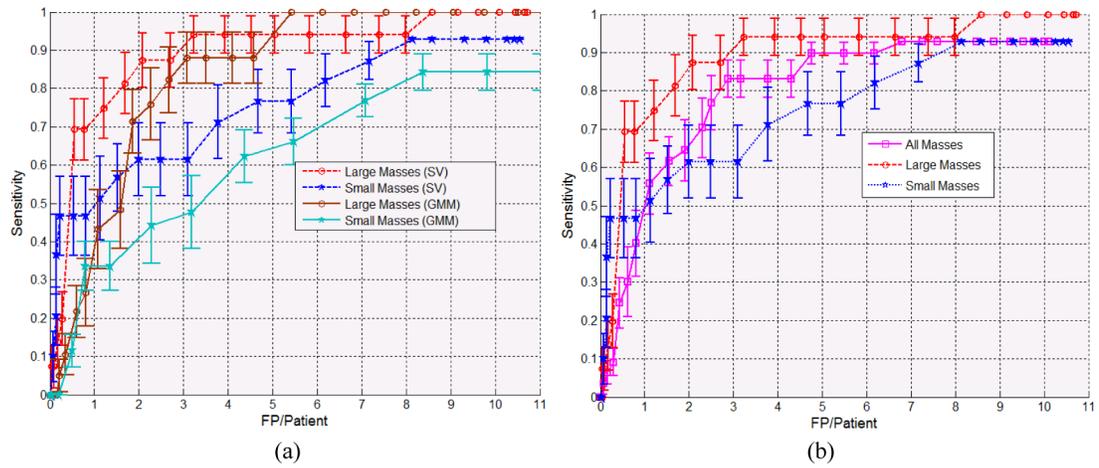


Fig. 6.

(a) FROC comparison on different sizes of masses (both supervoxel and spatially constrained GMM methods). (b) FROC comparison on different sizes of masses (supervoxel method). Large masses occupied 30% of spinal canal area. Small masses occupied between 5% and 30% of spinal canal area.

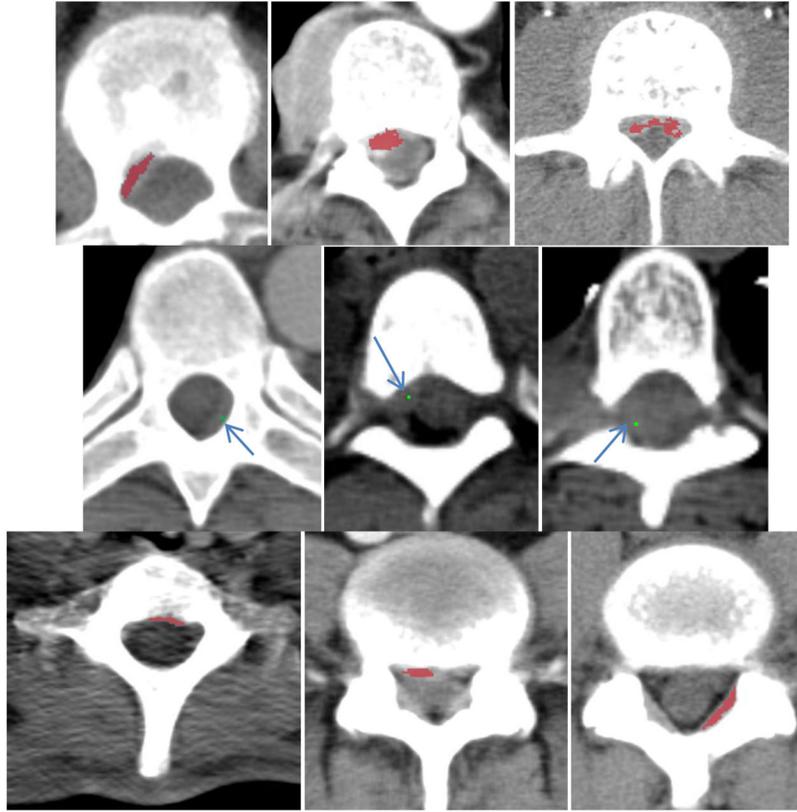


Fig. 7. Detection examples. 1st row: TP detections. 2nd row: missed epidural masses. 3rd row: FP detections (beam hardening artifacts, disk bulging, and ligament from left to right).

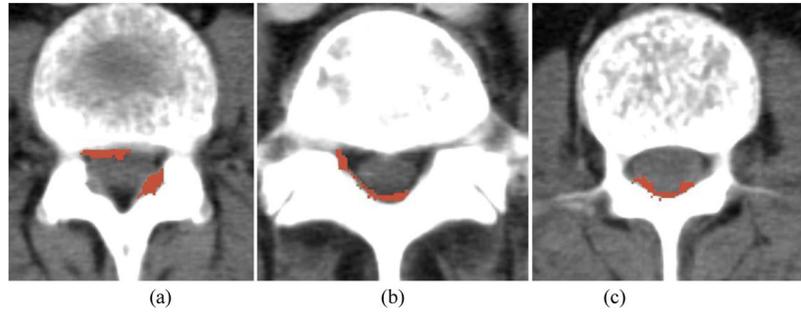


Fig. 8.

(a). False positives (disc bulging and ligament) from both GMM and supervoxel methods.
(b, c). False positives (partial volume between bone and normal soft tissue) from the GMM method only.