

Shape particle guided tissue classification

Marleen de Bruijne
IT University of Copenhagen
Denmark
marleen@itu.dk

Abstract

In many cases, the accuracy of statistical pixel classification can be improved by applying a spatially varying prior that can be derived from a shape model. We propose to represent the prior knowledge on the spatial distribution of tissue classes by a distribution of shape particles, each representing one possible distribution of tissue classes.

Classification and shape can then be optimized jointly by alternating a particle filtering step, in which the shape particle distribution is evolved under the influence of the current classification, with an update of the classification estimate using the shape distribution. Since a large number of shape hypotheses is used this method does not easily get trapped in local maxima. By applying shape models that are conditional on other, more easily discernible, objects in the image one can perform shape guided classification even if the shapes themselves are hardly visible.

The method is demonstrated on the task of detecting aortic calcifications in X-ray images, in which calcifications can only be present inside the aorta — mainly on the aortic wall — but the aorta itself is not visible.

1. Introduction

Statistical pixel classification is a popular tool for segmenting medical images [2, 3, 24–27]. It has been applied especially in tasks where the underlying shape of the object is not easily modeled, such as in white and gray matter segmentation in brain images, in blood vessels, and in pathologies of various kinds. Owing to image noise and inconsistencies in object appearance a segmentation based on pixel classification alone will usually result in a large number of misclassified pixels. Moreover, objects with similar appearance cannot be distinguished without using additional knowledge.

Prior knowledge about the fact that the object we try to segment is usually not a single pixel can be incorporated in various ways. Efforts to improve classification performance range from applying simple morphological post-

processing, adding position features to the appearance features, to performing contextual classification using Markov random fields and relaxation labeling methods (see for instance [22, 24, 30]). Especially in brain segmentation, registration to a digital atlas has been used extensively [2, 11, 25, 26]. These methods provide a rich description of possible spatial dependencies of tissue classes, but they rely on the (rigid or elastic) matching of an atlas to the image and therefore require the appearance of the entire image to be fairly consistent between subjects.

We propose to represent the prior knowledge on the spatial distribution of tissue classes by a distribution of shape ‘particles’ that evolves under the influence of image terms. The particles are sampled from a statistical shape model so as to constrain the allowed deformations. Each particle represents one possible distribution of tissue classes, and the complete collection of shape particles defines a spatially varying prior that is used in pixel classification. The particle cloud defining the spatially varying prior can be evolved similar to Monte Carlo methods known as ‘condensation’, ‘particle filtering’, or ‘factored sampling’ [9, 14, 17]. In each iteration, shape particles are weighted by their likelihood and a new shape set is constructed using weighted resampling and a small amount of random perturbation [8]. The joint optimum of classification and shape can be found by alternating the particle filtering step with an update of the classification estimate using the current shape distribution. The optimization by stochastic sampling makes the result relatively independent of the initialization and guarantees convergence provided that enough samples are used.

An important difference with the methods based on atlas registration is that the shape model can constrain the probability template even in images where the boundaries of the shape itself are not (or hardly) visible, as long as there exists a training set in which the shapes have been outlined. This can for instance be the case if images of another modality are available, or if a shape can only be determined by the existence of lesions.

In this work we focus on such an application — detecting calcified plaques in the lumbar aorta from standard ra-

diographs. A large number of structures in the image, e.g. bone and image artifacts, have a similar appearance as calcification and therefore prior information about the expected location of calcifications is needed for accurate detection. Obviously, aortic calcifications can only be present inside the aorta, but unfortunately the aorta itself is invisible in X-ray unless it contains calcifications. In this task registration-based methods would clearly fail. However, the shape and position of the aorta are strongly correlated to the shape and position of the spine, which is much easier detected in the image. We will exploit this correlation by applying a conditional model of aortic shape given the spine to guide appearance-based calcification detection.

Section 2 provides some background on aortic calcifications. The basic approach of shape particle guided classification is described in Section 3. The implementation for calcium detection, combining models of calcium distribution and a conditional model of the aorta given the spine, is described in Section 4. Section 5 presents experiments on 87 digitized X-ray images, and a conclusion is given in Section 6.

2. Aortic calcification

Calcifications in the abdominal aorta were shown to correlate with the presence — or future development — of calcifications at other sites such as the coronary arteries, and are an important predictor for future cardiovascular morbidity and mortality [10, 28, 29]. Accurate and reproducible measurement of the amount of calcified deposit in the aorta is therefore of great value in diagnosis, treatment planning, and the study of drug effects. Several automatic and semi-automatic calcium scoring methods have been proposed for CT [16, 18].

This paper aims at automatically detecting calcified plaques in the lumbar aorta from standard radiographs. Although CT is better suited for identifying and quantifying atherosclerosis, standard X-rays have the advantage that they are cheap and fast. Several approaches to manually quantifying the severity of aortic calcification in radiographs have been proposed, of which the antero-posterior severity score by Kauppila et al. [19] is the most popular. For this score, the lumbar part of the aorta is divided into four segments adjacent to the four vertebra L1-L4, and the severity of the anterior and posterior aortic calcification are graded individually for each segment on a 0-3 scale. The results are summed in a composite severity score ranging from 0 to 24. Such manual scoring systems have been successfully applied in epidemiological studies, but can not describe subtle changes in disease progression and are labor-intensive and prone to inter- and intra-observer variations.

An automated calcification detection scheme would allow for automatic scoring according to the current semi-quantitative standards as well as for continuous and likely

more precise quantification by for instance counting the number of calcified pixels or assessing the density of individual calcifications [4]. To our knowledge, no method by other authors currently exists for automatic detection of calcified plaques from standard radiographs.

We propose to combine pixel classification on the basis of local intensity features with a spatially varying prior that gives the probability of a pixel containing calcium, dependent on the position of a pixel within the aorta. In general, the shape and location of the aorta will not be known a priori, and since the aortic walls are only visible if calcium is present, automatic aorta segmentation can not be used as a first step to guide calcium detection. However, the shape and location of the aorta are strongly correlated with the shape and location of the spine [5]. We therefore use a set of training images in which the aorta and the vertebrae have been annotated to model the probability density function of aorta shape and location conditional on the spine shape. In a previous paper, we showed how samples from this shape model can be combined with a model of how the calcium is distributed within the aorta to define a spatially varying calcium prior dependent on the position and shape of the spine [7]. Application of this prior improved results of standard pixel classification. In this paper, we iteratively update the spatially varying calcium probability estimate and the classification to achieve their joint optimization.

The method requires the localization of the corner and midpoints of the first four lumbar vertebrae. Currently manual input is used here — obtained from a vertebral morphometry study on the same dataset — but these point positions could also be derived from an automatic spine segmentation, see e.g. [8, 23].

3. Shape particle guided classification

We optimize shape templates of class labels on an image using stochastic sampling optimization, similar as was done by de Bruijne and Nielsen in [8]. However, in this case the shape templates do not give a class label deterministically but define a probability for each class and each position in the template, and the estimated image labeling is updated during filtering.

Our aim is to find an image labeling C which is consistent with both the image I and a prior shape model $P(S)$. We will approach this by first obtaining the joint posterior probability distribution $P(C, S|I)$ for the labeling C and a set of shape samples S and then marginalizing over S .

To perform this joint optimization, one can use an iterative procedure:

Start with an initial labeling estimate C^0 and shape set $\mathcal{S}^0 = \{S_1, S_2, \dots, S_N\}$ sampled randomly from the shape prior

Iterate

1. Sample \mathcal{S}^t from $P(\mathcal{S}|C^{t-1}, I, \mathcal{S}^{t-1})$
2. Set estimate $C^t = \operatorname{argmax}_C P(C|I, \mathcal{S}^t)$

In the first step, a new shape set \mathcal{S}^t is sampled from the current shape set \mathcal{S}^{t-1} using likelihood weighting based on the current classification estimate. In the second step, the classification estimate is updated with the shape prior from the new shape set. In this way, the statistical classification is moderated with prior information from a shape model, whereas the variance in this shape prior decreases as image evidence for preferring some shapes above others is accumulated.

3.1. Initial classification estimate

We will use class probability density measurements as the image observations. The class probabilities are obtained using a pixel classifier trained to distinguish between pixels of different classes on the basis of local image descriptors. We have chosen a general scheme in which pixels are described by the outputs of a set of Gaussian derivative filters at multiple scales, and a k-NN classifier is used for probability estimation. The probability that a pixel with feature vector \mathbf{x} belongs to class ω is then given by

$$P(\omega|\mathbf{x}) = \frac{k_\omega}{k}, \quad (1)$$

where k_ω among the k nearest neighbors belong to class ω .

The initial labeling C^0 is the soft labeling based on pixel appearance alone

$$C^0(j, c) = P(\omega_c|\mathbf{x}_j) \quad (2)$$

which defines the probabilities for each of the classes c occurring in each pixel j .

3.2. Prior shape model

Any kind of shape model from which samples can be drawn can be inserted here. We will use the popular linear point distribution models (PDM) as proposed by Cootes and Taylor [6] to model the object shape variations observed in the training set.

In PDMs, shapes are defined by the coordinates of a set of landmark points which correspond between different shape instances. A collection of training shapes are aligned

using for instance Procrustes analysis [13] and a principal component analysis (PCA) of the aligned shapes yields the so-called *modes of shape variation* which describe a joint displacement of all landmarks. Each shape can then be approximated by a linear combination of the mean shape and these modes of variation. Usually only a small number of modes is needed to capture most of the variation in the training set.

3.3. Shape set resampling

Each shape S_i is associated with a class probability map $P(\omega_{cj}|S_i)$ which defines the probabilities for each of the classes occurring in each pixel. This class probability map can be based on expert knowledge of the problem at hand or it can be derived from a training set. Section 4.2 describes how the calcium probability we use in the experiments is obtained.

The shape set is updated using importance sampling and diffusion. Each shape S_i is assigned an importance weight w_i which reflects the degree of similarity between the current estimate of calcium probabilities and the calcium distribution that would be expected from the current shape. One possibility is to use the inner product between the two probability maps, as in:

$$w_i = \left[\frac{1}{Z} \sum_j \sum_c C^{t-1}(j, c) P(\omega_{cj}|S_i) \right]^\alpha, \quad (3)$$

where j are the pixels in the shape template (the image pixels that are covered by the shape template S_i), Z is a normalization factor which corrects for differences in template size, α is a tunable parameter controlling the speed of convergence of the algorithm, C^{t-1} is the current estimate of the probability that the pixel j belongs to the class ω_c , and $P(\omega_{cj}|S_i)$ is the probability that a pixel j is of class ω_c , given the aorta shape S_i .

A new shape set \mathcal{S}^t is generated from the current shape set \mathcal{S}^{t-1} using weighted sampling with replacement, and subsequently noise is added to the shapes that were selected multiple times in order to explore the solution space around these shapes. The noise represents a random diffusion.

In this iterative process, successful shapes multiply and unlikely shapes will vanish. Thus, the distribution of shape samples that make up the probability map ‘condenses’ in the presence of consistent image evidence. This leads to a more peaked prior probability density and a more precise labeling estimate.

3.4. Spatially varying classification

Each shape template defines a probability distribution of calcium given the shape, and the probability distribution

due to the shape collection is simply the average of all individual probability maps:

$$P(\omega_{cj}|\mathcal{S}) = \sum_i P(\omega_{cj}|S_i)P(S_i) \quad (4)$$

To simplify optimization, we assume that the two individual class probabilities, based on appearance and on position with respect to aorta shape, are independent. The classification estimate in each iteration is then given by the multiplication of the initial classification estimate with the current shape prior.

$$C^t(j, c) = P(\omega_{cj}|\mathbf{x}_j, \mathcal{S}^t) = P(\omega_c|\mathbf{x}_j)P(\omega_{cj}|\mathcal{S}^t). \quad (5)$$

Note that classification estimates as used during optimization are soft classifications, giving the vector of probabilities for each pixel. As the final classification estimate a hard classification is obtained from the probabilistic classification C by selecting an appropriate threshold.

4. Aorta shape particle guided calcification detection

This section describes our choice of shape prior $P(S)$ and class probability map $P(\omega_{cj}|S_i)$ that are specific to aortic calcification detection.

4.1. Conditional shape model

To construct a conditional shape model of the aorta given the spine, the spine and aorta landmarks are combined into one shape vector. The shapes are aligned using Procrustes alignment to minimize the sum of squared differences between the spine parts of the combined shape vectors. The distribution $P(S_1|S_2)$, i.e. the probability distribution of the expected aorta shape and pose S_1 for a given spine S_2 , can then be modeled as the Gaussian conditional density

$$P(S_1|S_2) = \mathcal{N}(\mu, K)$$

with

$$\mu = \mu_1 + \Sigma_{12}\Sigma_{22}^{-1}(S_2 - \mu_2)$$

$$K = \Sigma_{11} - \Sigma_{12}\Sigma_{22}^{-1}\Sigma_{21}$$

where μ_1 and μ_2 are the mean aorta and spine shapes respectively, and covariances Σ_{ij} are obtained from the covariance matrix of the combined model

$$\Sigma = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix}$$

as

$$\Sigma_{ij} = \frac{1}{n-1} \sum_n (S_{in} - \mu_i)(S_{jn} - \mu_j)^T.$$

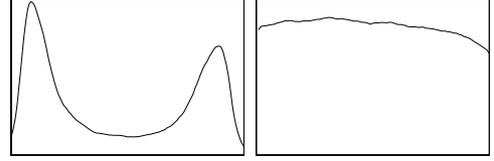


Figure 1. Cross-sectional profile (left) and longitudinal profile (right) of calcium distribution inside the aorta. Prevalence of calcium is higher at the aortic walls and near the aortic bifurcation.

$\Sigma_{12}\Sigma_{22}^{-1}$ is the matrix of regression coefficients of $(S_1 - \mu_1)$ on $(S_2 - \mu_2)$. Usually, Σ_{22} is not invertible owing to multi-collinearity in the landmark positions and unreliable due to chance covariance in a limited training set. Some regularization is therefore required. One option is to replace Σ_{22} by $\Sigma_{22} + \gamma I$, where γ is a positive and typically small constant. This approach is similar to regularization as was proposed for use in linear discriminant analysis [12] and equivalent to ridge regression [15]. As γ tends to infinity, the influence of the spine decreases; the remaining model is the original aorta model that describes the shape variation in the aorta independent of the spine.

An example of the modes of variation of such a conditional model is given in Figure 3. From this illustration, it is clear that the estimate of the posterior aortic wall, closer to the spine, is more strongly correlated with the spine landmarks and therefore exhibits smaller variation in the conditional model. The same holds for the upper part of the aorta which is more closely attached to the spine.

4.2. Class probability map

We use two classes in the probability templates, calcium and non-calcium. The non-calcium class is defined by a small boundary around the aorta. The calcium class has zero probability everywhere outside the aorta, and non-zero probability everywhere inside, but the magnitude varies with the position in the aorta.

It is well known that the distribution of calcification in the aorta is not uniform. The number of plaques increases towards the aortic bifurcation, and due to the projection imaging the majority of the plaques is visible along the anterior and posterior aortic walls and not in between.

If a large training set of example images with annotated aorta and calcifications was available, the probability of presence of calcium in each pixel could be estimated by labeling calcified pixels as 1 and non-calcified as 0, warping all aortas onto the mean aorta shape, and computing the average aorta label image.

If the training set is limited the above procedure will generalize poorly to unseen images; pixels inside the aorta may coincidentally have a very high or low probability of being

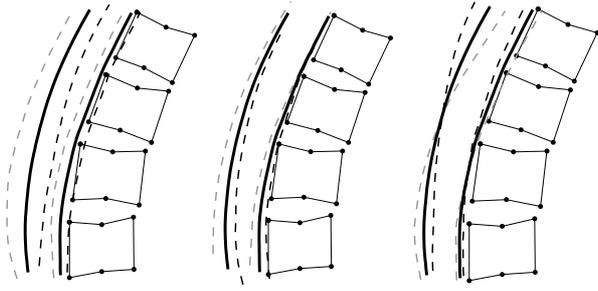


Figure 3. Modes of variation of the aorta given the known positions of vertebrae corner- and mid-points. The solid black line denotes the mean aorta shape, the mean shape ± 3 standard deviations are dashed in gray and black. From left to right the first three modes of variation are shown.

calcified. As a trade-off between generalizability and specificity we model the cross-sectional and longitudinal presence of calcium independently.

In a set of labeled training images, the part of the aorta adjacent to the first four lumbar vertebrae is selected and intensity profiles are sampled perpendicular to the vessel axis, reaching from the anterior to the posterior wall. All profiles are normalized to equal length and averaged to form a cross-sectional calcium prior distribution. For each image, one longitudinal profile is formed by summing the values in the individual profiles. An average longitudinal profile is computed by length normalizing and averaging the longitudinal profiles of all images.

For a given aorta shape, a calcium prior probability map can then be constructed by sweeping the cross-sectional prior profile along the axis, modulated with the longitudinal profile. The two profiles and an example of a calcium probability map are given in Figures 1 and 2.

5. Experiments

Leave-one-out experiments are performed on 87 lateral spine radiographs taken from a combined osteoporosis-atherosclerosis screening program. The dataset is diverse, ranging from uncalcified to severely calcified aortas. The original radiographs have been scanned at a resolution of 0.1 mm per pixel and were inverted for better visibility of calcific deposits. A medical expert outlined all calcifications adjacent to vertebrae L1 through L4 manually and also placed 6 points on each vertebra as is routinely done in quantitative vertebral morphology studies.

5.1. Parameter settings

Before further analysis the images were normalized to zero mean and unit variance. The appearance features used include the original image and the derivatives up to and including the third order computed at three different scales (1,

4.5 and 20 pixels). Training pixels were selected randomly from a region of interest including the aorta and its direct surroundings. The set of samples is normalized to unit variance for each feature, and k-NN classification is performed with an approximate k-NN classifier [1] with $k=25$. In all cases, results reported are accuracies of hard classification with the overall optimal threshold that is kept constant for all 87 images.

In the conditional shape model, 6 manually placed landmarks on each of the 4 vertebrae are used and 50 aorta landmarks are selected on each aortic wall by equidistant sampling along the manual outlines. The first 5 modes of shape variation are selected for the conditional shape model, $\gamma = 10^{-4}$ is used for regularization, and $N = 100$ aorta shapes are sampled randomly from the model to form the calcium prior probability map. The diffusion kernel applied during particle filtering is proportional to the shape model with 30% of the original variance, and the factor determining the speed of convergence was chosen as $\alpha = 2$. The optimization was run for a fixed number of 10 iterations.

5.2. Results

Performance of standard pixel classification and shape guided classification are measured in a region of interest surrounding the aorta. To make a fair comparison with pixel classification, we have chosen a region of interest (ROI) that can be selected using the location of the spine in the image and basic knowledge of distances between spine and aorta. The region is formed by connecting the anterior corner points on vertebrae L1-L4 with line segments and translating the thus obtained curve to the positions of minimum and maximum distance of the aorta with respect to this spine curve as were observed in the training set. Examples of ROIs obtained are given in Figure 4.

The pixel classification alone yields an average accuracy, defined as the percentage of correctly classified pixels within the ROI, of 96.7%. Combining this with the spatially varying prior results in an improvement to 96.9%. The difference may seem small, but because of the relatively small number of calcification pixels, differences of a few percent point in accuracy are important. The difference between standard classification and shape guided classification is highly significant with $p < 0.001$ in a paired t-test. The κ -statistic for observer agreement, which accounts for the imbalance of positives and negatives, increased from 0.09 for pixel classification to 0.23 for shape particle guided classification — an improvement from ‘slight’ to ‘fair’ agreement according to Landis and Koch [20]. The average overlap area of the mean aorta shape from the final particle distribution with the manually annotated aorta is 87% and the distance between the two contours is on average 2.7 mm.

Figure 4 shows examples of the classification results. At the threshold that is optimized for overall pixel classifica-

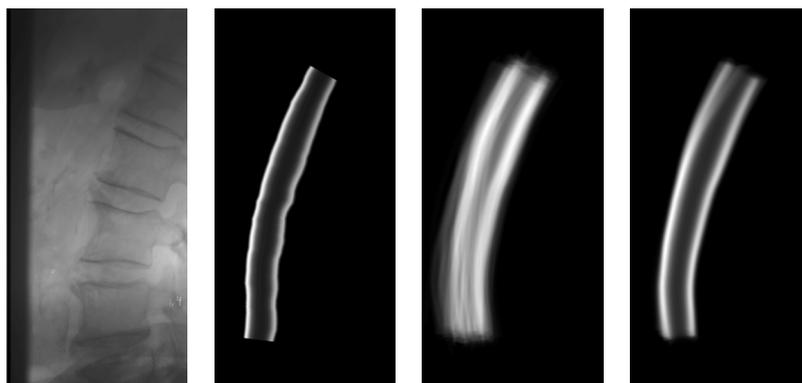


Figure 2. From left to right: Original (sub)image, inverted for better visibility of calcium; annotated aorta with constructed calcium probability map; calcium probability map from 100 random samples of the aorta conditional shape model; calcium probability map after 10 iterations of the joint optimization.

tion accuracy within the ROI, standard pixel classification misses many of the true positives and finds some false positives outside the aorta. Shape particle guided classification will normally produce positives only inside the aorta even at high threshold values, thus allowing for a higher sensitivity at the same specificity.

6. Conclusion

We propose an iterative, stochastic optimization algorithm that simultaneously and robustly optimizes a pixel labeling and an underlying shape distribution. This method could be used to incorporate high-level spatial information in pixel classification tasks.

The method is demonstrated, in connection with a model of aorta shape conditioned on the spine, on detecting calcifications in X-ray images of the aorta. Results of pixel classification without spatial prior were improved significantly.

7. Acknowledgements

We would like to thank L.B. Tankó and C. Christiansen of the Center for Clinical and Basic Research (CCBR A/S), Denmark, for providing the data sets and manual segmentations used in this study, L.A. Conrad-Hansen and M.T. Lund of the IT University of Copenhagen, Denmark, for providing the software to perform the manual segmentations, and CCBR A/S, Denmark, for funding.

References

- [1] S. Arya, D. Mount, N. Netanyahu, R. Silverman, and A. Wu. An optimal algorithm for approximate nearest neighbor searching. *Journal of the ACM*, (45):891–923, 1998.
- [2] J. Ashburner and K. Friston. Voxel-based morphometry – the methods. *Neuroimage*, 11(6):805–821, 2000.
- [3] A. Chung and J. Noble. Statistical 3D vessel segmentation using a Rician distribution. In *Medical Image Understanding and Analysis*, 1999.
- [4] L. Conrad-Hansen, M. de Bruijne, F. Lauze, L. Tankó, and M. Nielsen. Quantizing calcification in the lumbar aorta on 2-D lateral X-ray images. In Liu et al. [21], pages 409–418.
- [5] L. Conrad-Hansen, J. Raundahl, L. Tánko, and M. Nielsen. Prediction of the location of the posterior aortic wall using the first four lumbar vertebrae as a predictor. In M. Sonka and M. Fitzpatrick, editors, *Medical Imaging: Image Processing*, volume 5370 of *Proceedings of SPIE*. SPIE Press, 2004.
- [6] T. Cootes, C. Taylor, D. Cooper, and J. Graham. Active shape models – their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, 1995.
- [7] M. de Bruijne. A pattern classification approach to aorta calcium scoring in radiographs. In Liu et al. [21], pages 170–177.
- [8] M. de Bruijne and M. Nielsen. Shape particle filtering for image segmentation. In C. Barillot, D. Haynor, and P. Hellier, editors, *Medical Imaging Computing & Computer-Assisted Intervention*, volume 3216 of *Lecture Notes in Computer Science*, pages 186–175. Springer, 2004.
- [9] A. Doucet, N. de Freitas, and N. Gordon, editors. *Sequential Monte Carlo methods in practice*. Springer-Verlag, New York, 2001.
- [10] D. Eggen, J. Strong, and H. McGill. Calcification in the abdominal aorta: relationship to race, sex, and coronary atherosclerosis. *Archives of pathology*, 78:575–583, 1964.
- [11] B. Fischl, D. Salat, M. Albert, M. Dieterich, C. Haselgrove, A. van der Kouwe, R. Killiany, D. Kennedy, A. Montillo, N. Makri, B. Rosen, and A. Dale. Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron*, 33(3):341–355, 2002.

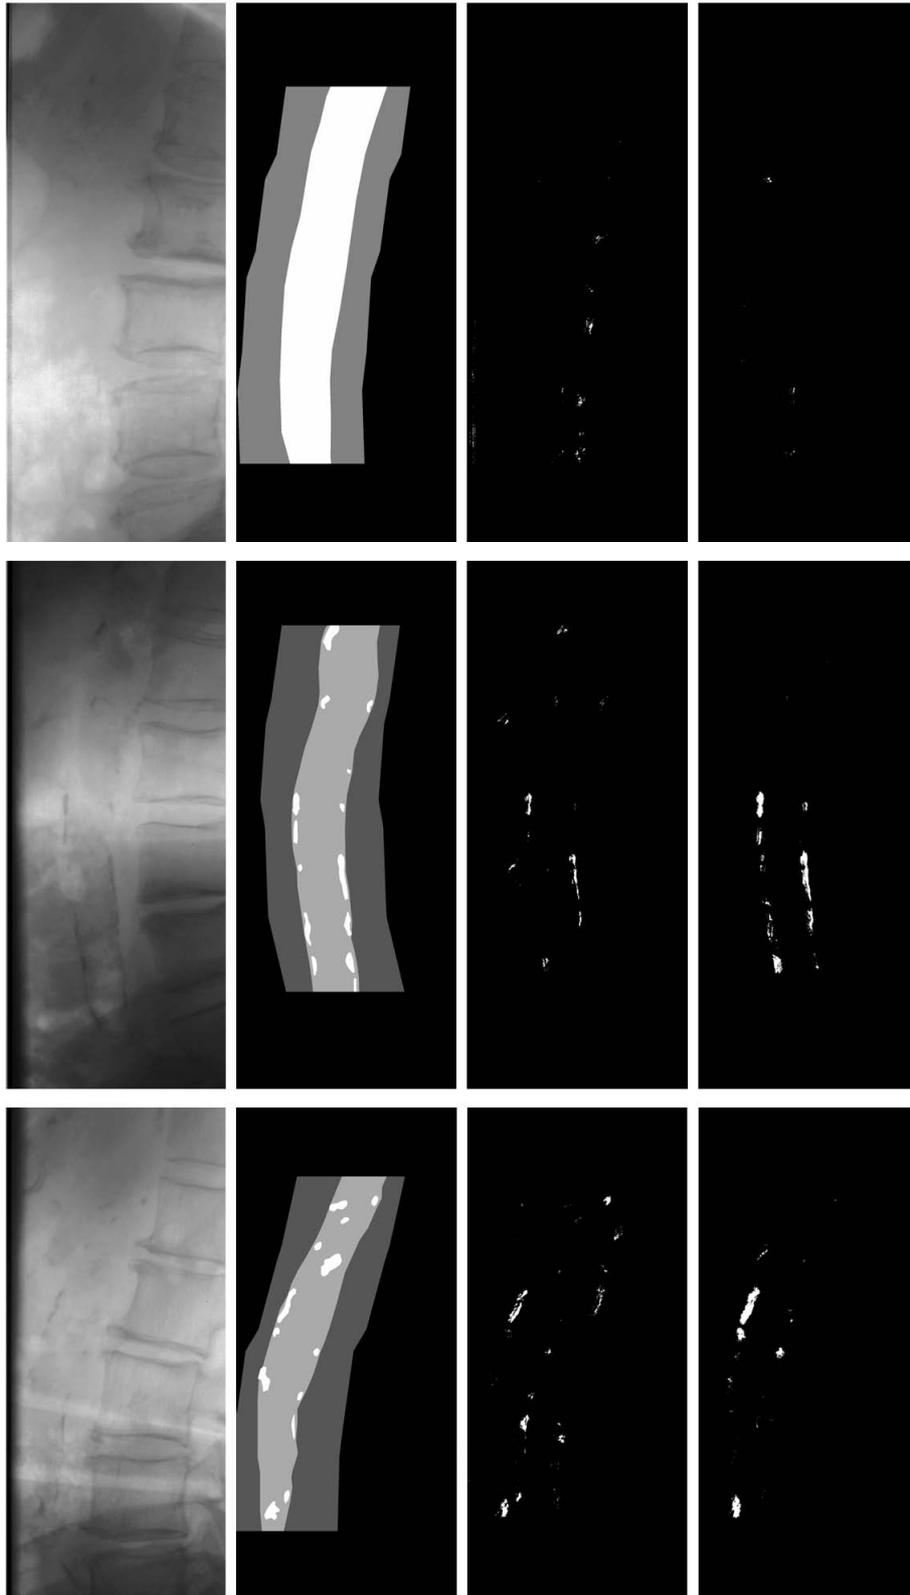


Figure 4. Examples of classifications obtained for images of varying degree of calcification. Each row gives the different results for one image, from left to right: Original (sub)image, inverted for improved visibility of calcium; Manual segmentation and ROI; Pixel classification alone; Shape particle guided classification.

- [12] J. Friedman. Regularized discriminant analysis. *Journal of the American Statistical Association*, 84(405):165–175, 1989.
- [13] C. Goodall. Procrustes methods in the statistical analysis of shape. *Journal of the Royal Statistical Society B*, 53(2):285–339, 1991.
- [14] U. Grenander, Y. Chow, and D. Keenan. *HANDS: A Pattern Theoretic Study of Biological Shapes*. Springer-Verlag, 1990.
- [15] A. Hoerl and R. Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, February 1970.
- [16] C. Hong, K. Bae, and T. Pilgram. Coronary artery calcium: Accuracy and reproducibility of measurements with multi-detector row CT — assessment of effects of different thresholds and quantification methods. *Radiology*, 227:795–801, 2003.
- [17] M. Isard and A. Blake. Visual tracking by stochastic propagation of conditional density. In B. Buxton and R. Cipolla, editors, *European Conference on Computer Vision*, volume 1064 of *Lecture Notes in Computer Science*, pages 343–356. Springer, 1996.
- [18] I. Isgum, B. van Ginneken, and M. Olree. Automatic detection of calcifications in the aorta from CT scans of the abdomen. *Academic Radiology*, 11(3):247–257, March 2004.
- [19] L. Kauppila, J. Polak, L. Cupples, M. Hannan, D. Kiel, and P. Wilson. New indices to classify location, severity and progression of calcific lesions in the abdominal aorta: a 25-year follow-up study. *Atherosclerosis*, 25(132(2)):245–250, 1997.
- [20] J. R. Landis and G. G. Koch. The measurement of observer agreement for categorical data. *Biometrics*, 33(1):159–174, Mar 1977.
- [21] Y. Liu, T. Jiang, and C. Zhang, editors. *Proceedings of the ICCV Workshop on Computer Vision for Biomedical Image Applications: Current Techniques and Future Trends (CVBIA)*, volume 3765 of *Lecture Notes in Computer Science*. Springer, 2005.
- [22] M. Loog and B. van Ginneken. Supervised segmentation by iterated contextual pixel classification. In *Image and Signal Processing*, Proceedings of the 16th International Conference on Pattern Recognition (ICPR 2002), pages II: 925–928, 2002.
- [23] P. Smyth, C. Taylor, and J. Adams. Vertebral shape: Automatic measurement with active shape models. *Radiology*, 211(2):571–578, 1999.
- [24] B. van Ginneken, M. Stegmann, and M. Loog. Segmentation of anatomical structures in chest radiographs using supervised methods: a comparative study on a public database. *Medical Image Analysis*, 10(1):19–40, Feb 2006.
- [25] K. van Leemput, F. Maes, D. Vandermeulen, and P. Suetens. Automatic segmentation of brain tissues and MR bias field correction using a digital brain atlas. In W. Wells, A. Colchester, and S. Delp, editors, *Medical Imaging Computing & Computer-Assisted Intervention*, volume 1496 of *Lecture Notes in Computer Science*, pages 1222–1229. Springer, 1998.
- [26] S. Warfield, M. Kaus, F. Jolesz, and R. Kikinis. Adaptive, template moderated, spatially varying statistical classification. *Medical Image Analysis*, 4:43–55, 2000.
- [27] W. Wells, III, W. Grimson, R. Kikinis, and F. Jolesz. Adaptive segmentation of MRI data. *IEEE Transactions on Medical Imaging*, 15(4):429–442, 1996.
- [28] P. Wilson, L. Kauppila, C. O'Donnell, D. Kiel, M. Hannan, J. Polak, and L. Cupples. Abdominal aortic calcific deposits are an important predictor of vascular morbidity and mortality. *Circulation*, 103:1529 – 1534, 2001.
- [29] J. Witteman, J. van Saase, and H. Valkenburg. Aortic calcification as a predictor of cardiovascular mortality. *Lancet*, 2:1120–1122, 1986.
- [30] Y. Zhang, M. Brady, and S. Smith. Hidden Markov random field model and segmentation of brain MR images. *IEEE Transactions on Medical Imaging*, 20(1):45–57, 2001.