This is a reprint of material published in ISSN 0302-9743, ISBN 978-3-642-18420-8, Medical Computer Vision—Recognition Techniques and Applications in Medical Imaging (MICCAI 2011 Workshop), LNCS 6533, Beijing, China, Sept. 2010, pp. 96–105.

Detection of 3D Spinal Geometry Using Iterated Marginal Space Learning

B. Michael Kelm[1], S. Kevin Zhou[2], Michael Suehling[2], Yefeng Zheng[2], Michael Wels[1], and Dorin Comaniciu[2]

[1] Corporate Technology, Siemens AG, Erlangen, Germany,
michael.kelm@siemens.com,
[2] Siemens Corporate Research, Princeton, USA.

Abstract. Determining spinal geometry and in particular the position and orientation of the intervertebral disks is an integral part of nearly every spinal examination with Computed Tomography (CT) and Magnetic Resonance (MR) imaging. It is particularly important for the standardized alignment of the scan geometry with the spine. In this paper, we present a novel method that combines Marginal Space Learning (MSL), a recently introduced concept for efficient discriminative object detection, with a generative anatomical network that incorporates relative pose information for the detection of multiple objects. It is used to simultaneously detect and label the intervertebral disks in a given spinal image volume. While a novel iterative version of MSL is used to quickly generate candidate detections comprising position, orientation, and scale of the disks with high sensitivity, the anatomical network selects the most likely candidates using a learned prior on the individual nine dimensional transformation spaces. Since the proposed approach is learning-based it can be trained for MR or CT alike. Experimental results based on 42 MR volumes show that our system not only achieves superior accuracy but also is the fastest system of its kind in the literature – on average, the spinal disks of a whole spine are detected in 11.5s with 98.6% sensitivity and 0.073 false positive detections per volume. An average position error of 2.4mm and angular error of 3.9° is achieved.

1 Introduction

Examinations of the vertebral column with both Magnetic Resonance (MR) imaging and Computed Tomography (CT) require a standardized alignment of the scan geometry with the spine. While in MR the intervertebral disks can be used to align slice groups and to position saturation bands, in CT the reconstruction planes need to be aligned. In addition to the position and orientation of the disks, physicians are interested in labeling them (e.g. C2/C3, C5/T1, L1/L2, ...). Such a labeling allows to quickly determine the anatomical location without error-prone counting. As manual alignment is both time-consuming and operator-dependent, it is desirable to have a robust, fully automatic, and thus reproducible approach.

An automatic procedure for extracting the spinal geometry faces various challenges, however. Varying contrasts and image artifacts can compromise the detection of intervertebral disks based on local image features. Thus, a global spine model is required to robustly identify individual disks from their context. Such a model must also cope with missed detections and patients with an unusual number of vertebrae. Finally, the overall approach should run within seconds to allow clinical application.

In this paper we propose a novel approach that combines efficient local object detection based on Marginal Space Learning (MSL) [14] with a global probabilistic model that incorporates pose priors on the nine dimensional parameter spaces that encode position, orientation and scale of the individual disks. The whole approach follows the database-guided detection paradigm [4] and can thus be easily trained for spine detection in CT as well as MR acquired with different sequences.

1.1 Related Work

Recently, the detection and analysis of spinal geometry has regained interest. Boisvert et al. [1] present a model that describes the statistical variations of the spine in terms of sequential rigid transformations of the local vertebra coordinate systems. Using principal component analysis on the Riemannian manifold of rigid transformations they can extract clinically meaningful eigenmodes. Although relying on the same metrics we formulate a probabilistic spine model that is applied for detection rather than statistical analysis.

The detection of intervertebral disks in 3D MR scout scans has recently been addressed by Pekar et al. [8]. They propose a three-step approach using a special-purpose 2D image filter for disk candidate detection, followed by a customized spine tracking method and a final labeling step based on counting. Since their approach is designed to work on MR data only, it might not be easily adapted to CT image volumes.

Schmidt et al. [10] propose a trainable approach based on extremely randomized trees in combination with a complete graphical model. They employ an A*-search based inference algorithm for exact maximum a posteriori (MAP) estimation. The approach only considers the position of the intervertebral disks, while we also determine their orientations and scales. However, their parts-based 3D approach appears most related to ours and their results based on 3D T_1-weighted composed multi-station MR data can best be compared with ours.

Corso et al. [2] argue that a two-level probabilistic model is required to separate pixel-level properties from object-level geometric and contextual properties. They propose a generative graphical model with latent disk variables which they solve by generalized expectation maximization (EM). Although the approach only provides position estimates and has only been evaluated for lumbar disks in 2D T_2-weighted MR data, it could in principle be extended to full 3D estimation. But since EM only finds a local optimum of the expected log likelihood, which can render such an approach very sensitive to initialization, it is not clear

how the approach would scale to higher-dimensional estimation including 3D position, orientation, and scale.

2 Methods

Our approach can be subdivided into three major steps (cf. Fig. 1). To constrain the search range for the disks, the spine is roughly located within the given volume first. Second, disk candidates are generated with high sensitivity using a novel iterative extension of the MSL approach [14]. Finally, a global probabilistic spine model is used to select the most likely disk candidates based on their appearance and relative pose and to determine the appropriate label for each disk.

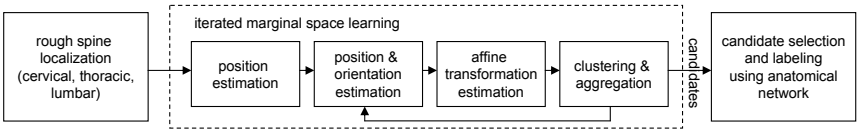

Fig. 1. Overall approach.

2.1 Global Probabilistic Spine Model

The typical spatial structure of the spine gives rise to a prior on the relative poses of the spinal disks. This has been modeled by the factor graph [7] depicted in Fig. 2. We have chosen a chain model with potentials considering position, orientation and scale of the spinal disks. Each of the (vector-valued) random variables $\boldsymbol{b}_1$ to $\boldsymbol{b}_N$ represents the pose of a certain spinal disk, thus $\boldsymbol{b}_s$ holds a 3D position $\boldsymbol{p}_s = [x_s, y_s, z_s]^T$, a unit quaternion $\boldsymbol{q}_s$ representing the orientation [6] and an anisotropic scale $\boldsymbol{s}_s = [s_s^x, s_s^y, s_s^z]^T$ for every disk $s \in \{1, \ldots, N\}$. Thus, a distribution over disk poses is defined by the log probability

$$\log \Pr(\boldsymbol{b}_1, \boldsymbol{b}_2, \ldots, \boldsymbol{b}_N \,|\, \boldsymbol{\Theta}, \boldsymbol{I}) = \sum_s V_s(\boldsymbol{b}_s \,|\, \boldsymbol{\theta}_s, \boldsymbol{I}) + \sum_{s \sim t} V_{st}(\boldsymbol{b}_s, \boldsymbol{b}_t \,|\, \boldsymbol{\theta}_{st}) - A \quad (1)$$

where A is the log partition function, $\boldsymbol{I}$ represents the image data and $\boldsymbol{\Theta} = \{\boldsymbol{\theta}_s, \boldsymbol{\theta}_{st}\}$ subsumes all model parameters which are detailed in the following.

The pair potential between two neighboring disk $\boldsymbol{b}_s$ and $\boldsymbol{b}_t$ combines relative position, relative orientation and relative scale terms:

$$V_{st}(\boldsymbol{b}_s, \boldsymbol{b}_t \,|\, \boldsymbol{\theta}_{st}) = V_{pos,st} + V_{rot,st} + V_{sca,st} \quad (2)$$

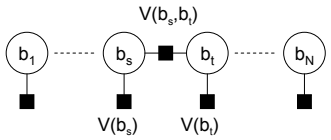

Fig. 2. Factor graph modeling the relation between the spinal disks.

Each of the terms is defined as a Gaussian pair potential, i.e.,

$$V_{pos,st}(\boldsymbol{b}_s, \boldsymbol{b}_t) = -\frac{1}{2}\, \boldsymbol{d}_{pos}^T(\boldsymbol{b}_s, \boldsymbol{b}_t)\, \boldsymbol{\Sigma}_{pos,st}^{-1}\, \boldsymbol{d}_{pos}(\boldsymbol{b}_s, \boldsymbol{b}_t) \tag{3}$$

$$V_{rot,st}(\boldsymbol{b}_s, \boldsymbol{b}_t) = -\frac{\alpha(\boldsymbol{q}_t \boldsymbol{q}_s^{-1} \boldsymbol{\mu}_{rot,st}^{-1})^2}{2\sigma_{rot,st}^2} \tag{4}$$

$$V_{sca,st}(\boldsymbol{b}_s, \boldsymbol{b}_t) = -\frac{1}{2}\, \boldsymbol{d}_{sca}^T(\boldsymbol{b}_s, \boldsymbol{b}_t)\, \boldsymbol{\Sigma}_{sca,st}^{-1}\, \boldsymbol{d}_{sca}(\boldsymbol{b}_s, \boldsymbol{b}_t) \tag{5}$$

with the rotation angle $\alpha(\boldsymbol{q}) = \alpha([q_0\, q_1\, q_2\, q_3]) = 2\arccos(q_0)$ and with $\boldsymbol{d}_{pos}(\boldsymbol{b}_s, \boldsymbol{b}_t) = \mathbf{R}_s^{-1}(\boldsymbol{p}_t - \boldsymbol{p}_s) - \boldsymbol{\mu}_{pos,st}$ and $\boldsymbol{d}_{sca}(\boldsymbol{b}_s, \boldsymbol{b}_t) = \boldsymbol{s}_t - \boldsymbol{s}_s - \boldsymbol{\mu}_{sca,st}$ where $\mathbf{R}_s$ is the rotation matrix associated with the quaternion $\boldsymbol{q}_s$. In summary, the pair potential parameters $\boldsymbol{\theta}_{st}$ are the mean parameters $\boldsymbol{\mu}_{pos,st}$, $\boldsymbol{\mu}_{rot,st}$, $\boldsymbol{\mu}_{sca,st}$ and the (co-)variance parameters $\boldsymbol{\Sigma}_{pos,st}$, $\sigma_{rot,st}$, $\boldsymbol{\Sigma}_{sca,st}$. To keep the number of estimated parameters small, both $\boldsymbol{\Sigma}_{pos,st}$ and $\boldsymbol{\Sigma}_{sca,st}$ are constrained to diagonal matrices.

Both, the position and the scale potentials are defined based on Euclidean distance. The required mean parameters $\boldsymbol{\mu}_{pos,st}$ and $\boldsymbol{\mu}_{sca,st}$ and the covariance matrices $\boldsymbol{\Sigma}_{pos,st}$ and $\boldsymbol{\Sigma}_{sca,st}$ are determined from the training data. The rotation potential in Eqn. (4) uses the intrinsic metric $\alpha(\boldsymbol{q})$ of the corresponding manifold $\mathcal{SO}^3$. Consequently, the mean rotation is determined as the Fréchet mean [9]. Collecting all instances of a certain disk pair $(\boldsymbol{b}_s, \boldsymbol{b}_t)$ into the training sample $\mathcal{P}_{st}$, the Fréchet mean for the corresponding rotation potential is determined as

$$\boldsymbol{\mu}_{rot,st} = \underset{|\boldsymbol{q}|=1}{\operatorname{argmin}} \sum_{(\boldsymbol{b}_s, \boldsymbol{b}_t) \in \mathcal{P}_{st}} \alpha(\boldsymbol{q}_t \boldsymbol{q}_s^{-1} \boldsymbol{q}^{-1})^2. \tag{6}$$

It can be efficiently computed using the eigen-decomposition proposed in reference [6]. The Gaussian variance is estimated with

$$\sigma_{rot,st}^2 = \frac{1}{|\mathcal{P}_{st}| - 1} \sum_{(\boldsymbol{b}_s, \boldsymbol{b}_t) \in \mathcal{P}_{st}} \alpha(\boldsymbol{q}_t \boldsymbol{q}_s^{-1} \boldsymbol{\mu}_{rot,st}^{-1})^2. \tag{7}$$

Finally, the single site potentials, which are determined by iterated marginal space learning as described in the following section, encode image-based likelihood, i.e.,

$$V_s(\boldsymbol{b}_s \,|\, \boldsymbol{\theta}_s, \boldsymbol{I}) = log(\Pr(\boldsymbol{b}_s \,|\, \boldsymbol{\theta}_s, \boldsymbol{I})). \tag{8}$$

Since the defined potentials are invariant under global rigid transformations (translation and rotation), the resulting distribution is insensitive towards different poses of the spine. Furthermore, models capturing only parts of the complete spine can be easily constructed by just omitting the superfluous disk variables. Since all potential parameters are determined independently (i.e., the likelihood decouples), no retraining is required and a probabilistic model appropriate for the current acquisition protocol, e.g., a lumbar spine protocol, can be assembled at runtime.

2.2 Iterated Marginal Space Learning

In principle, the defined potentials can be evaluated for every possible position, orientation and scale. However, performing an exhaustive search on the uniformly discretized nine dimensional parameter space (3 position, 3 orientation and 3 scale parameters) would require evaluating a huge number of single site as well as pair potentials. Such a direct approach would be computationally very expensive.

Hence, we adopt the MSL paradigm [14], a novel concept that has recently proven successful in numerous applications [3, 5, 13]. Instead of searching the whole nine dimensional parameter space, the MSL paradigm proposes a three-step approach. First candidate positions for the sought object are collected by using a probabilistic machine learning classifier to check every voxel location within a defined range. In the second step, a number of 3D orientation hypotheses that have been derived from the training set are evaluated by a second classifier using the the most likely object positions from the first step. Similarly, the last step estimates three scale parameters based on the candidates from the second step using a third classifier.

MSL has been designed to detect a single, specific object such as, for example, a particular organ or landmark. If multiple objects of the same type are to be detected, as in our case, the described MSL approach may end up with detections for the most salient disks only, i.e., many disks would be missed. Although the sensitivity could be improved by drastically increasing the number of considered candidates in each step, this is not practicable since MSL would then loose its computational efficiency.

We therefore propose a novel extension to MSL, iterative MSL (iMSL), to cope with multiple objects of the same type (cf. Fig. 3). It is designed to achieve a higher sensitivity than usual MSL at moderate computational costs. First, the position detector is evaluated in each voxel of the given image volume region. The N_0 most likely candidates are collected in the set of initial position candidates $\mathcal{P}_0$. Then, the best N_{pos} ($N_{pos} < N_0$) candidates from $\mathcal{P}_0$ are evaluated using the orientation detector whose top candidates are evaluated using the scale detector. The resulting set $\mathcal{D}_{sca}$ contains disk candidate detections with all estimated parameters. Using pairwise average-linkage clustering with Euclidean distance, clusters of candidate disks are obtained. The most likely N_A box candidates of each resulting cluster are averaged and added to the set of detected disk candidates $\mathcal{D}$. After removing all position candidates from $\mathcal{P}_0$ that are closer than a specified radius R to any of the detections in $\mathcal{D}$, orientation and scale

Input: R, N_0, N_{pos}, N_{ort}, N_{sca}
Output: Set $\mathcal{D}$ of detected disk candidates
$\mathcal{D} := \{\}$;
$\mathcal{P}_0 :=$ the N_0 most likely candidates according to the position detector;
repeat
  $\mathcal{P}_0 := \{p \in \mathcal{P}_0 : d(p,q) > R \; \forall q \in \mathcal{D}\}$;
  $\mathcal{D}_{pos} :=$ the N_{pos} most likely candidates from $\mathcal{P}_0$;
  $\mathcal{D}_{ort} :=$ the N_{ort} most likely candidates from $\mathcal{D}_{pos}$ according to the orientation detector;
  $\mathcal{D}_{sca} :=$ the N_{sca} most likely candidates from $\mathcal{D}_{ort}$ according to the scale detector;
  Perform hierarchical agglomerative clustering on $\mathcal{D}_{sca} \cup \mathcal{D}$;
  foreach *cluster* $\mathcal{C}$ **do**
    if $|\mathcal{C}| \geq N_A$ **then**
      Aggregate the top N_A candidates and add the resulting box to $\mathcal{D}$;
    end
  end
until $|\mathcal{P}_0| = 0$ *or* $|\mathcal{D}|$ *remains constant*;

Fig. 3. Pseudo-code for Iterated Marginal Space Learning (iMSL).

detection are repeated on the remaining position candidates until no position candidates are left or no new disk candidates are detected.

Like Zheng et al. [14] we employ the probabilistic boosting tree (PBT) classifier using Haar-like features for the position detector and steerable features for the orientation and scale detectors.

The probabilistic spine model described in the previous section is discretized using the disk candidates detected with iMSL. Each random variable $\boldsymbol{b}_s$ is transformed into a discrete random variable where each state represents one of the detected disk candidates. In order to allow for missed detections, an extra "missing" state is introduced. Note, that iMSL detects disk candidates with high sensitivity which usually results in more disk candidates than actual disks. The MAP estimate, i.e. the maximum of Eqn. (1)), provides the optimum assignment of a disk candidate to one of the disk variables according to the probabilistic spine model. Thus, only those disk candidates that form a valid spine are selected and are implicitly assigned a suitable label.

The MAP is efficiently computed by belief propagation where, due to the tree structure of the factor graph (cf. Fig. 2), a single forward-backward pass yields the exact solution [7]. An additional speed-up is obtained by constraining the search for disk candidates to the area of the spine. For this purpose, bounding boxes around the lumbar, thoracic and cervical regions of the spine are detected first using the usual MSL approach as described in reference [14].

3 Experimental Results

3.1 Data

Experiments have been conducted based on 3D T_1-weighted MR volumes (FL3D-VIBE sequence) from 42 volunteers. About one half of the volumes has been acquired on two 1.5T scanner models (MAGNETOM Avanto and MAGNETOM Espree, Siemens AG, Erlangen) with TR = 5/4ms, TE = 2ms and a flip angle of 10°. The other half has been obtained from two 3T scanner models (MAGNETOM Trio, MAGNETOM Verio, Siemens AG, Erlangen) with TR = 4/3ms, TE = 1ms and again a flip angle of 10°. Each of the volumes was recorded in a two station scan and subsequently combined to a volume covering the whole spine (approximately 860mm × 350mm × 190mm) with an isotropic resolution of 2.1mm. Susceptibility artifacts and intensity variations due to magnetic field inhomogeneities were present in the data. No bias field correction was performed.

3.2 Results

To obtain ground truth, each intervertebral disk has been annotated with four defined landmarks. From these, ground truth boxes have been derived for the intervertebral disks as well as the lumbar, thoracic and cervical spine regions. For disk detection, *i*MSL was employed with a cluster radius of $R = 6$mm, $N_0 = 3000$ initial position candidates and 500 detection candidates for the remaining detection estimation steps ($N_{pos} = 500$, $N_{ort} = 500$, $N_{sca} = 500$).

All evaluation results have been obtained using 10-fold cross validation, ensuring that training and testing data never stem from the same patient. Every ground truth annotation for which no disk within a distance of 10mm was detected, was counted as a missed detection. Overall, intervertebral disks have been detected with a sensitivity of 98.64% and only 0.0731 false positives per volume, yielding a positive predictive value of 99.68%. The overall processing time on a 2.2GHz dual core laptop computer was between 9.9s and 13.0s and 11.5s on average where most of the time was spent on disk candidate detection.

The accuracy of the detected intervertebral disks has been evaluated by the position distance and the angle between the disk plane normals of the detected intervertebral disks and the ground truth annotation (cf. Table 1). On average, a position error of 2.42mm (about 1 voxel) and an angular error of 3.85° was obtained.

Four examples from the MR data set are shown in Fig. 4. The right-most example shows a case where the volunteer has been instructed to lie down twisted in order to simulate a scoliotic spine. Still the proposed approach could locate and label all spinal disks reliably.

Some results on lumbar spine CT are shown in Fig. 5. While the complete probabilistic spine model as well as the *i*MSL detectors have been trained on CT data showing various regions of the spine, our approach allows to assemble an appropriate model for the lumbar spine without retraining.

	cervical	thoracic	lumbar	overall	cervical	thoracic	lumbar	overall
mean	2.09	2.41	2.86	2.42	4.86	3.38	3.80	3.85
median	1.84	2.18	2.68	2.19	3.89	2.90	3.37	3.17
lower quartile	1.40	1.56	1.88	1.58	2.52	1.85	2.08	1.97
upper quartile	2.63	3.00	3.63	3.05	6.68	4.48	5.03	5.02

Table 1. Disk detection results using 10-fold cross validation based on 42 T_1-weighted MR volumes. **Left**: position error [mm]. **Right**: angular error between normals [degree].

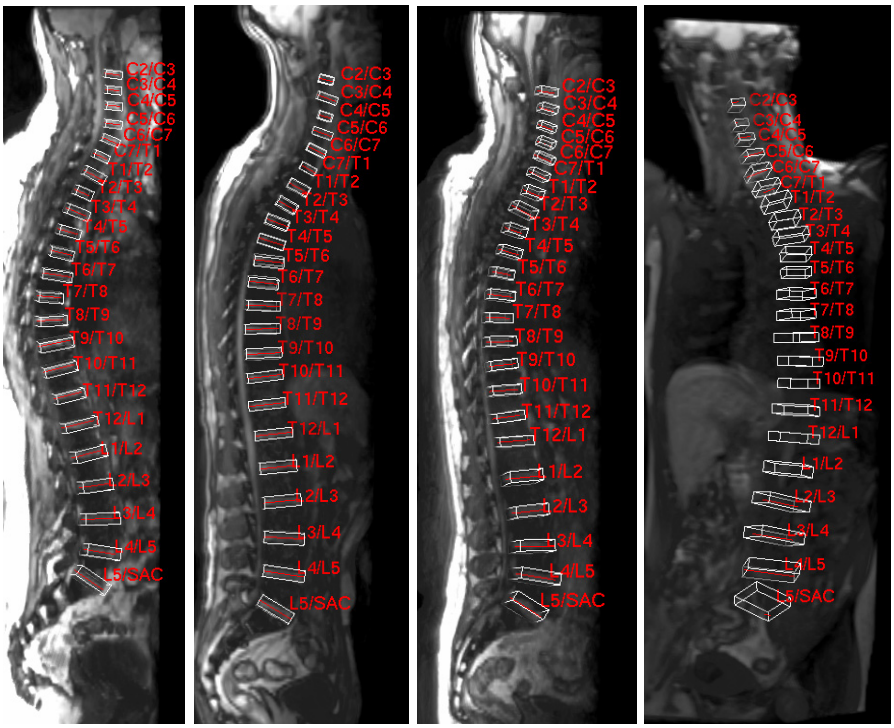

Fig. 4. Four examples from the MR data with detection results. Although the volunteer in the rightmost example lay down in an unusually twisted pose, all intervertebral disks were detected and labeled correctly.

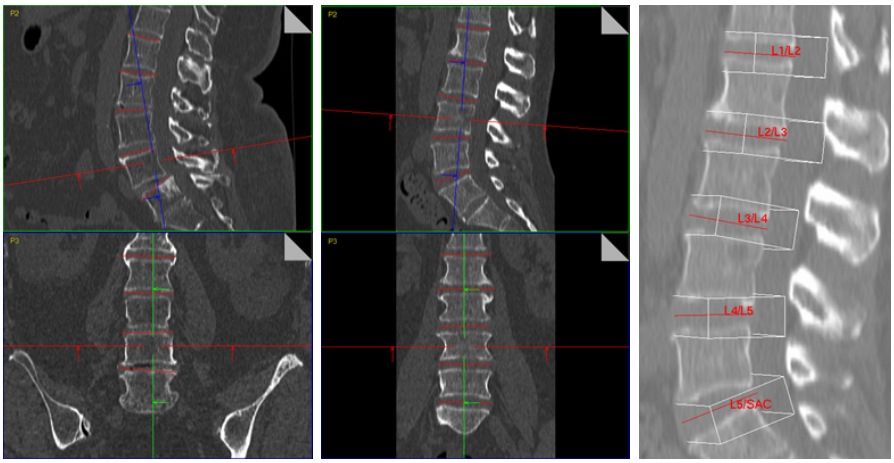

Fig. 5. Detection results for CT scans of the lumbar spine.

The results of our proposed method compare favorably with results presented in previous works. While with only 6s processing time the approach by Pekar et al. [8] runs faster than ours, it has lower sensitivity (95.6% before candidate selection) and does not provide orientation estimates.

Compared with the best cross validation results by Schmidt et al. [10], the results obtained with our approach are significantly better. While a competitive but still smaller sensitivity of 97% is reported, they only achieve a position error of 5.1mm. Furthermore, no orientation estimates are provided and the approach takes several minutes to run. Furthermore, in contrast to Schmidt et al. [10], we did not perform any posterior search at the positions of missing disks which could further increase our sensitivity.

While at the current state we did not perform systematic testing on data from patients with pathologies (e.g. scoliosis, stenosis, disk degeneration, herniation, desiccation), we are confident that our approach also works for disease cases. In this as well as other applications we have observed, that the MSL approach is very robust to imaging artifacts and unusual appearances of the sought object. Using *i*MSL, increases sensitivity and helps detect disks with very unusual appearance. Furthermore, since the global spine model is restricted to candidates provided by the disk detector, scoliotic abnormalities can be robustly handled. The volunteer with the twisted pose in Fig. 4 provides evidence towards this. Finally, simple retraining of our system with some abnormal cases added, enables the detectors as well as the prior model to handle them even more reliably.

4 Conclusion and Future Work

In this paper, we have presented a novel approach to the fully automatic detection of 3D spinal geometry and labeling of the intervertebral disks. The approach

uses an iterative extension of MSL for disk candidate detection along with an anatomical network that incorporates spatial context in form of a prior on the nine dimensional disk poses. Since the entire approach is learning-based, it can be trained for CT and MR alike.

Using 42 MR image volumes, superior sensitivity and accuracy was obtained than in previous works. With an overall processing time of only $11.5s$, the approach is also comparably fast and can be used as routine procedure for the automatic planning of scan geometries. Results on CT data show that the proposed approach can be adapted to different modalities. For this purpose, the graphical model can be adjusted to handle partial spine recordings that are commonly acquired with CT.

Apart from automatic scan alignment, the proposed system for detecting and labeling the intervertebral disks could be part of a computer-aided diagnosis system for analyzing pathologies of the intervertebral disks or the vertebrae. The detected bounding boxes could, for example, be used for initializing a detailed vertebra segmentation algorithm with subsequent analysis. Furthermore, the proposed system could support semantic body parsing and semantic annotation to automatically generate semantic location descriptions as frequently used by physicians for reporting [12, 11]. Both applications will be considered in future work.

References

1. Boisvert, J., Cheriet, F., Pennec, X., Labelle, H., Ayache, N.: Geometric variability of the scoliotic spine using statistics on articulated shape models. IEEE Trans. Med. Imag. 27(4), 557–568 (2008)
2. Corso, J.J., Alomari, R.S., Chaudhary, V.: Lumbar disc localization and labeling with a probabilistic model on both pixel and object features. In: Proc. MICCAI. pp. 202–210 (2008)
3. Feng, S., Zhou, S., Good, S., Comaniciu, D.: Automatic fetal face detection from ultrasound volumes via learning 3D and 2D information. In: Proc. CVPR. pp. 2488–2495 (2009)
4. Georgescu, B., Zhou, X.S., Comaniciu, D., Gupta, A.: Database-guided segmentation of anatomical structures with complex appearance. In: Proc. CVPR. pp. 429–436 (2005)
5. Ionasec, R.I., Voigt, I., Georgescu, B., Wang, Y., Houle, H., Hornegger, J., Navab, N., Comaniciu, D.: Personalized modeling and assessment of the aortic-mitral coupling from 4D TEE and CT. In: Proc. MICCAI. pp. 767–775 (2009)
6. Karney, C.F.: Quaternions in molecular modeling. J. Molec. Graph. Modelling 25, 595–604 (2007)
7. Kschischang, F.R., Frey, B.J., Loeliger, H.A.: Factor graphs and the sum-product algorithm. IEEE Trans. Inf. Theory 47(2), 498–519 (2001)
8. Pekar, V., Bystrov, D., Heese, H.S., Dries, S.P.M., Schmidt, S., Grewer, R., den Harder, C.J., Bergmans, R.C., Simonetti, A.W., van Muiswinkel, A.M.: Automated planning of scan geometries in spine MRI scans. In: Proc. MICCAI. pp. 601–608 (2007)
9. Pennec, X.: Intrinsic statistics on Riemannian manifolds: Basic tools for geometric measurements. J. Math. Imag. Vis. 25(1), 127–154 (2006)

10. Schmidt, S., Kappes, J., Bergtholdt, M., Pekar, V., Dries, S., Bystrov, D., Schnörr, C.: Spine detection and labeling using a parts-based graphical model. In: Proc. IPMI. pp. 122–133 (2007)
11. Seifert, S., Barbu, A., Zhou, S.K., Liu, D., Feulner, J., Huber, M., Sühling, M., Cavallaro, A., Comaniciu, D.: Hierarchical parsing and semantic navigation of full body ct data. In: Proc. SPIE Medical Imaging. pp. 725–732 (2009)
12. Seifert, S., Kelm, M., Möller, M., Mukherjee, S., Cavallaro, A., Huber, M., Comaniciu, D.: Semantic annotation of medical images. In: Proc. SPIE Medical Imaging. o.A. (2010)
13. Wels, M., Zheng, Y., Carneiro, G., Huber, M., Hornegger, J., Comaniciu, D.: Fast and robust 3-D MRI brain structure segmentation. In: Proc. MICCAI. pp. 575–583 (2009)
14. Zheng, Y., Barbu, A., Georgescu, B., Scheuering, M., Comaniciu, D.: Four-chamber heart modeling and automatic segmentation for 3-D cardiac CT volumes using marginal space learning and steerable features. IEEE Trans. Med. Imag. 27(11), 1668–1681 (2008)